

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/237005742>

Complex Systems '98 --- Complexity Between the Ecos: From Ecology to Economics

Book · January 1998

CITATIONS

3

READS

230

8 authors, including:



[B. I. Henry](#)

University of New South Wales

102 PUBLICATIONS 1,552 CITATIONS

[SEE PROFILE](#)



[Robert Marks](#)

University of New South Wales

91 PUBLICATIONS 705 CITATIONS

[SEE PROFILE](#)



[David G. Green](#)

Monash University (Australia)

176 PUBLICATIONS 2,378 CITATIONS

[SEE PROFILE](#)



[Steve Keen](#)

Kingston University

75 PUBLICATIONS 584 CITATIONS

[SEE PROFILE](#)

Complex Systems '98

Complexity Between the Ecos: From Ecology to Economics

Edited by

Russell Standish, Bruce Henry, Simon Watt, Robert Marks
The University of New South Wales, Sydney, Australia

Robert Stocker, David Green
Charles Sturt University, Albury, Australia

Steve Keen
The University of Western Sydney, Macarthur, Australia

and

Terry Bossomaier
Charles Sturt University, Bathurst, Australia

ISBN 0 7334 0537 1

©1998 Copyright of all material in this book remains with the relevant authors. Submission of articles to this proceedings grants the editors a right for the papers to be published in this book, and electronically in the journal *Complexity International*.

Once printed versions of this book have sold out, the book will be obtainable from the *Complexity Online Network* (<http://life.csu.edu.au/complex>; <http://parallel.hpc.unsw.edu.au/complex>).

Preface

Complex Systems '98 — Complexity Between the Ecos is the fourth in a series of conferences [1, 2, 3] devoted to the inter-disciplinary study of Complex Systems. This year, we bid particular welcome to a small, but enthusiastic and growing band of “dynamical economists”, who are working to apply methods and results from complexity theory in their profession. It is now nearly three years since Steve Keen and others organised a convivial conference entitled *Commerce, Complexity and Evolution* that drew together economics and commerce researchers with an interest in this area. This conference may also be considered a successor of that conference. In the intervening years, a selection of papers presented at that conference, enhanced by feedback between the researchers, has been collected and published as a volume of *International Symposia in Economic Theory and Econometrics* [4]. I hope that this book will prove a worthy successor to that tome.

The theme of this conference reflects the fact that there is considerable correspondence between matters biological (or ecological) and economics. Marshall wrote in his *Principles*:

The Mecca of the economist lies in economic biology rather than in economic dynamics.

That correspondence we now recognise as being *Complex Systems*. The organisation of this volume follows this correspondence, from applications of complexity to biology, through general complexity studies (divided into theoretical mathematical approaches and computer modelling) to arrive at applications of complexity to economics.

Complex Systems studies in its present form is perhaps 15 years old, with some of its constituent parts (eg chaos theory, or fractals) a decade or so older again. We have reached, I believe, a level of maturity in the field where applications dominate over theoretical concerns. We have stopped asking “What is *Complexity*?” (there are too many confusing answers here) and like the biologists (who have stopped asking “What is Life?”), we are now getting down to the business of working out what it is useful for. It is vitally important, though, that we continue to talk to each other, and keep up that tradition of inter-disciplinary research that has been the hallmark of Complexity studies to date. For this is one of the most exciting areas of scientific discovery in the closing years of the twentieth century.

References

- [1] D. G. Green and T. J. Bossomaier, editors, *Complex Systems — from Biology to Computation* IOS, Amsterdam, 1993.
- [2] R. J. Stonier and X. H. Yu, editors, *Complex Systems — Mechanism of Adaption* IOS, Amsterdam, 1994.
- [3] R. Stocker, H. Jelinek, B. Durnota and T. Bossomaier, editors, *Complex Systems — from Local Interactions to Global Phenomena* IOS, Amsterdam, 1996.
- [4] W. Barnett, C. Chiarella, S. Keen, R. Marks, and H. Schnabl, editors, *Commerce, Complexity and Evolution*, volume 11 of *International Symposia in Economic Theory and Econometrics*. Cambridge UP, Cambridge 1998.

Acknowledgements

It is impossible for any one person to perform an adequate editorial function over such a wide range of topics that complex systems theory covers. Therefore, I wish to thank the entire editorial team for their work in assessing the papers, and for performing numerous miscellaneous tasks required to run a successful conference. I also wish to thank my employer (UNSW) for allowing me to spend so much time on this project. Finally, I would like to thank my family for their patience during the final few hectic weeks prior to the conference.

Russell Standish
November 1998.

Contents

I	Invited Paper	1
1	Discrete Dynamical Networks and their Attractor Basins <i>Andrew Wuensche</i>	3
II	Biology	23
2	Emergent symmetry of local and global maps in the primary visual cortex: Self-organization of orientation preference <i>David M. Alexander, Phil Sheridan, Paul D. Bourke, Otto Konstandatos and James J. Wright</i>	25
3	Environmental informatics - a new paradigm for coping with complexity in nature <i>David G. Green and Nicholas I. Klomp</i>	32
4	Complex behavior in perceptual line length <i>Takuo Henmi and Michael L. Kalish</i>	40
5	Wide-band spectral power fluctuations characterize the response of simulated cortical networks to increasing stimulus intensity <i>John Klopp, Patrick Johnston, Valeriy Nenov and Eric Halgren</i>	55
6	Mechanism for Changing the Foraging Behavior in an Ant Colony Model <i>Mari Nakamura and Koichi Kurumatani</i>	66
7	The First Time-Derivative of the EEG: A Possible Proxy for the Order-Parameter for the Cerebral Cortex <i>J. W. Sleigh, D.A. Steyn-Ross and Moira Steyn-Ross</i>	74
8	Cellular Ecolab <i>Russell K. Standish</i>	80
III	Mathematical Tools	89
9	Effective Computation of 2D Coupled Map Lattices <i>Jacques Blanc-Talon</i>	91
10	Succinct representation of the Poincaré map for periodically driven differential equations <i>J.H.B. Deane and D.J. Jefferies</i>	101

11	Convergence and Aperiodicity in Fuzzy Cellular Automata: Revisiting Rule 90	110
	<i>P. Flocchini, F. Geurts, A. Mingarelli and N. Santoro</i>	
12	A Theoretical Framework for Abundance Distributions in Complex Systems	118
	<i>Stephan R.P. Halloy</i>	
13	The Variable Structure System: Intermittency, Chaos, and Trapping in Electronic Experiments and Simulations	130
	<i>D.J. Jefferies and J.H.B. Deane</i>	
14	Is There Meaning In Fractal Analyses?	144
	<i>Herbert F. Jelinek, Cameron. L. Jones and Matthew D. Warfel</i>	
15	A Categorical Representation of the State Transition Graph of Finite Cellular Automata	150
	<i>Jung-Hee, Park and Hyen-Yeal, Lee</i>	
16	Adaptive Stabilization of Lorenz Chaos	162
	<i>Yuping Tian and Xinghuo Yu</i>	
IV	Modelling	169
17	Effects of population size upon emergent group behavior	171
	<i>Cristobal Baray</i>	
18	Generous and Greedy Strategies	179
	<i>Bengt Carlsson and Stefan Johansson</i>	
19	A Complex Systems Approach to Simulating Human Behaviour Using Synthetic Landscapes	188
	<i>H. Randy Gimblett, Merton T. Richards and Robert M. Itami</i>	
20	Linguistic rule extraction from neural networks for high-dimensional classification problems	210
	<i>Hisao Ishibuchi, Manabu Nii and Kimiko Tanaka</i>	
21	Self-Organisation in a Simple Pursuit Game	219
	<i>Ashley Tews and Raymond Lister</i>	
V	Economics	227
22	A Model of Short- and Long Term Stock Market Behaviour	229
	<i>Trond Andresen</i>	
23	Learning about the Cobweb	244
	<i>Carl Chiarella and Xue-Zhong He</i>	
24	Troubles in Wonderland	258
	<i>Ric D. Herbert and Gareth D. Leeves</i>	
25	Competition between strategies for a market selection game	272

Hisao Ishibuchi, Chi-Hyon Oh and Tomoharu Nakashima

26	The Emergence and Collapse of the State: A Game Theoretic Analysis with Computer Simulations	282
	<i>Atsushi Iwasaki, Sobei H. Oda and Kanji Ueda</i>	
27	Complex Dynamics of Speculative Price	291
	<i>Taisei Kaizoji</i>	
28	The Nonlinear Dynamics of Debt Deflation	305
	<i>Steve Keen</i>	
29	The Emergence and Collapse of Market-dominant Products	320
	<i>Yusuke Koyama and Sobei H. Oda</i>	
30	An Explanation of Generic Behavior in an Evolving Financial Market	327
	<i>Shareen Joshi and Mark A. Bedau</i>	
31	The Complexity of Competitive Marketing Strategies	336
	<i>Robert E. Marks, David F. Midgley, Lee G. Cooper and G. M. Shiraz</i>	
32	The Application of Cellular Automata to the Theory of Consumers' Learning and Behavioral Interdependence	346
	<i>Sobei H. Oda, Kouhei Iyori, MIURA Ken and Kanji Ueda</i>	
33	Econolab	358
	<i>Russell K. Standish</i>	
34	Crystallisation of Two-Dimensional Cellular Automata	364
	<i>Tomoaki Suzudo</i>	
35	Agent based Iterated Multiple Lake Game with Local Governments	376
	<i>Tomohisa Yamashita, Keiji Suzuki and Azuma Ohuchi</i>	

Part I
Invited Paper

Discrete Dynamical Networks and their Attractor Basins

Andrew Wuensche
 Santa Fe Institute, 1399 Hyde Park Road,
 Santa Fe, New Mexico 87501 USA,
wuensch@santafe.edu,
<http://www.santafe.edu/~wuensch/>

Abstract

A key notion in the study of network dynamics is that state-space is connected into basins of attraction. Convergence in attractor basins correlates with order-complexity-chaos measures on space-time patterns. A network's "memory", its ability to categorize, is provided by the configuration of its separate basins, trees and sub-trees. Based on computer simulations using the software Discrete Dynamics Lab[19], this paper provides an overview of recent work describing some of the issues, methods, measures, results, applications and conjectures.

1 Introduction

Processes consisting of concurrent networks of interacting elements which affect each other's state over time are central to a wide range of natural and artificial systems drawn from many areas of science; from physics to biology to cognition; to social and economic organization; to computation and artificial life; to complex systems in general. The dynamics of these "decision making" networks depends on the connections and update logic for each element, resulting in complex feedback webs that are difficult to treat analytically. Understanding these systems depends on numerical simulations of idealized computer models known as discrete dynamical networks.

Cellular automata (CA) are a powerful yet simple class of network, characterized by a homogeneous rule and uniform nearest neighbour connections, providing models to study processes in physical systems such as reaction-diffusion[15], and self-organization by the emergence of coherent interacting structures[18]. By contrast, random Boolean networks (RBN) provide models for biological systems such as neural [3] and genetic[11] networks, where connections and rules must be less constrained. In addition, the idealized networks themselves hold intrinsic interest as mathematical/physical systems.

A key notion underlying network behavior is that state-space is organized into a number of basins of attraction, connecting states according to their transitions, and summing up the network's global dynamics, analogous to Poincaré's "phase portrait" which provided powerful insights in continuous dynamics.

The quality of dynamical behaviour of CA, from ordered to chaotic¹, is reflected by convergence in attractor basins, their characteristic in-degree, which influences the length of transients and attractor cycles. The in-degree of a state is its number of pre-images (predecessors). Bushy subtrees with high in-degree imply high convergence and order. Sparsely branching subtrees imply low convergence and chaos. In the case of RBN, attractor basins reveal how the network is able to hierarchically categorize state-space into separate basins, trees and sub-trees, the network's "memory". Changes to the network's wiring or rules change the memory categories, providing insights into learning[17, 20].

Traditionally, network dynamics has been investigated by running networks forward from many initial states to study space-time phenomenology[15], and for statistical measures on basins of attraction[8]. More recently, exact representations of basins of attraction and sub-trees have become accessible, where algorithms directly compute the pre-images of network states, allowing the network to be run "backwards" to disclose all possible historical paths[16, 17, 21]. Based on computer simulations using the software Discrete Dynamics Lab (DDLab)[19], this paper provides an overview of network architecture, the characteristics of space-time patterns, the methods and algorithms for reconstructing

¹The notion of "chaos" is used here by analogy only to its meaning in chaos theory, although there are many common properties, for example sensitivity to initial conditions.




Figure 1: A basin of attraction (one of 15) of a random Boolean network ($n=13$, $k=3$) shown in figure 15. The basin links 604 states, of which 523 are garden-of-Eden states. The attractor period = 7, and one of the attractor states is shown in detail as a bit pattern. The direction of time is inwards from garden-of-Eden states to the attractor, then clock-wise.

basins of attraction, and related parameters, measures, results, applications and conjectures, placing the dynamics along particular trajectories in the context of global dynamics.

2 Network Architecture

Discrete dynamical networks consist of a set of elements (cells) taking inputs from each other, and changing their cell-state according to some logical function on their inputs. The connectivity is usually sparse. The cell-state ranges over a discrete alphabet, in this paper just a binary alphabet (0 or 1) is considered. The updating is generally synchronous, though updating sequentially in a preset order or partial order is also of interest. A partial order is a sequence of sets of cells, where updating within each set is synchronous.


A CA is a very regular network, sometimes described as an artificial universe with its own physics. Cells take inputs from their nearest (and next nearest) neighbours (local “wiring”) according to a fixed neighbourhood template, so issues of network geometry and boundary conditions are crucial. The same logical rule is applied everywhere. Figure 2 shows neighbourhood templates for 1d, 2d and 3d as applied in DDLab. An RBN relaxes these constraints, allowing arbitrary “wiring” and rules, as in figure 3. The number of input wires available to each cell may also vary. However, an RBN architecture can be biased in countless of ways, described in section 4, to constrain wiring or rules. For example an RBN wiring with a constant rule, or local wiring with mixed rules can be a generalisation of a CA.

The wiring and rules can be tailored to very specific requirements, as in models of neural networks in the cortex[3]. The wiring can be constrained within a fixed distance from each cell, which confers meaning to network geometry and boundary conditions, whereas with completely arbitrary wiring the geometry just provides a convenient way of representing the network. A rule mix can be constrained to sample just a few rules or rules with a particular bias, as in genetic network models[5].


Hybrid networks can be constructed by putting an RBN within a CA or vice-versa. Networks of networks can be set up, with weak interactions so they perturb each other’s dynamics. The functionality for setting up networks in these ways is present in DDLab.

The network parameters can be listed as follows:

1d, $k=0-13$. The extra asymmetric cell in even k is on the right. The wiring is shown between two time-steps.



2d, $k=2-13$ ($k=0-1$ as in 1d). Note that $k=6$ and $k=7$ define an effectively triangular grid by changing between odd and even rows. The classical von Neumann and Moore neighbourhoods are indicated.



3d, $k=6-13$ ($k=0-5$ as in 2d), shown looking up into an axonometric cage.




Figure 2: 1d, 2d and 3d neighbourhood templates defined in DDLab. In 2d and 3d, to maximize symmetry, even k does not include the central target cell.

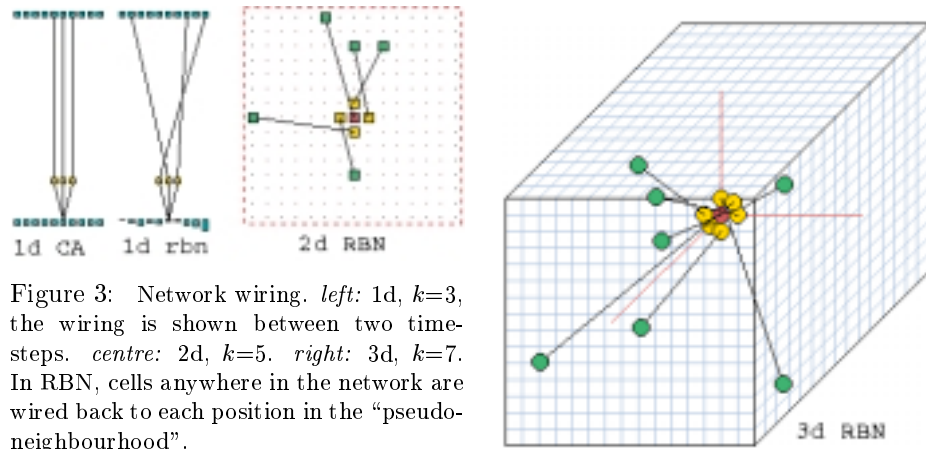


Figure 3: Network wiring. *left*: 1d, $k=3$, the wiring is shown between two time-steps. *centre*: 2d, $k=5$. *right*: 3d, $k=7$. In RBN, cells anywhere in the network are wired back to each position in the “pseudo-neighbourhood”.

size: The system size n , the number of cells in the network.

connectivity: The number of input wires per cell k , or the k -mix if the connectivity is not homogeneous. The connectivity is usually sparse, i.e. $k \ll n$.

neighbourhood: The neighbourhood template for CA, or the pseudo-neighbourhood for RBN, as in figure 2.

wiring: For RBN, how each cell is wired relative to its pseudo-neighbourhood.

rule: The rule for CA, or the rule scheme scheme for RBN. Rules are generally defined as look-up tables.

updating: The updating, usually synchronous. Alternatively sequential according to a defined order or partial-order.

geometry: The underlying geometry and boundary conditions, 1d, 2d, 3d, orthogonal or triangular, or some other geometry, for example a hypercube. This is essentially a function of the the neighbourhood template and wiring scheme. The geometry for graphically representing the network may not necessarily correspond to the underlying geometry.

A CA neighborhood, or RBN pseudo-neighborhood, of size k has 2^k permutations of values. The

most general expression of the Boolean function or rule is a lookup table (the rule-table) with 2^k entries, giving 2^{2^k} possible rules. Sub-categories of rules can also be expressed as simple algorithms, concise AND/OR/NOT logical statements (which could be implemented as combinatorial circuits), totalistic rules[14] or threshold functions.

By convention[14] the rule table is arranged in descending order of the values of neighborhoods, and the resulting bit string converts to a decimal or hexadecimal rule number. For example the $k=3$ rule-table for rule 30,

7	6	5	4	3	2	1	0	... neighbourhoods, decimal
111	110	101	100	011	010	001	000	... neighbourhoods, binary
0	0	0	1	1	1	1	0	... outputs, the rule table

The rule-table for other k values are set out in a corresponding way. $k \geq 4$ rules are referred to by their hexadecimal rule numbers. $k \leq 3$ rules are usually referred to by their more familiar decimal rule numbers.

For a given geometry, the behaviour space of CA depends on the size of rule-space, 2^{2^k} , though rule symmetries effectively reduce this number. For example, the $2^{2^3} = 256$ rules in $k = 3$ rule-space reduce to 88 equivalence classes[16]. The behaviour space of RBN is much greater, taking into account possible permutations of wiring and rule schemes, but there are also RBN equivalence classes relating to these permutations[10]. In general, the number of effectively different RBN of size n cannot exceed $(2^n)^{(2^n)}$ (see section 9).

3 Trajectories and space-time patterns

A state of a discrete dynamical network is the pattern of 0s and 1s at a given time-step. A trajectory is the sequence of states at successive time-steps, the systems *local* dynamics. Examples of 1d, 2d and 3d space-time patterns are shown in figures 4, 5 and 6. A time axis is only possible in representations of 1d or 2d systems. As well as showing cells as white(0) or black(1), an alternative presentation shows cells in colors (or shades) according to their look-up neighbourhood (figure 4). This allows the most frequently occurring colors to be progressively filtered to show up gliders and other space-time structures as in figure 5, which can be done interactively, on-the-fly, in DDLab for any CA. This is an alternative method to the “computational mechanics” approach[4].




Figure 4: Space-time patterns of a CA ($n=24$, $k = 3$, rule 90). 24 time-steps from an initial state with a single central 1. Two alternative presentations are shown. *Left*: cells by value, white=0 black=1. *Right*: cells colored (or shaded) according to their look-up neighbourhood. This allows filtering, and improves the clarity of space-time patterns in 2d and 3d.

3.1 Glider dynamics in CA

A large body of literature is devoted to the study space-time patterns in CA. “Glider” or particle dynamics, where coherent configurations emerge and interact, provide a striking instance of self-organization in a simple system. Such dynamics are classified as complex, in contrast to ordered or





Figure 5: Space-time patterns of the $k=3$ rule 54 ($n=150$) from the same initial state showing interacting gliders, which are embedded in a complicated background. *Left*: cells by value. *Right*: cells by neighbourhood lookup, with the background filtered.



(a) 2d 100x100 triangular grid

(b) 2d 56x56 +time

(c) 3d 20x20x20

Figure 6: Examples of 2d and 3d CA space patterns. (a) is an evolved time-step of a 2d CA on a $k=7$ triangular lattice with a reaction-diffusion rule. (b) is the 2d game-of-Life on a 56×56 grid, but with a time dimension added, similar to a 1d space-time pattern. The initial state is set with a number of gliders. (c) is a time-step of a 3d $k=7$ CA with a randomly selected rule and starting from a single central 1.

chaotic[14], a well know example being Conway’s 2d “game-of-Life”[1]. Because glider dynamics is relatively rare in CA rule spaces, much study has focused on the few known complex rules in 1d CA. However, an unlimited source of examples are now available, found by the methods described in sections 3.2 – 3.3.

Gliders are embedded within a uniform or periodic background or domain, and propagate at various velocities up the system’s speed of light set by the neighbourhood diameter. Gliders are interpreted as dislocation in the background or as the boundary reconciling two different backgrounds[4, 18]. Gliders may absorb or eject sub-gliders (glider-guns). Compound gliders may emerge made up of sub-gliders re-colliding periodically. Figure 7 shows some examples.

Glider dynamics has been interpreted as occurring at a phase transition in rule-space between order and chaos[9], relative to the rule parameters λ [9] and Z [16] (see section 6.2). Input-entropy provides a measure on space-time dynamics that allows the automatic classification of rule-space (see below).

3.2 Input entropy

Keeping track of the frequency of rule-table look-ups (the k -block frequency, or “look-up frequency”) in a window of time-steps, provides a measure, the variance of input-entropy over time, which is used to classify 1d CA automatically for a spectrum of ordered, complex and chaotic dynamics[22]. The method allows screening out rules that support glider dynamics and related complex rules, giving an unlimited source for further study. The method also shows the distribution of rule classes in the rule-spaces of varying neighbourhood sizes. The classification produced seems to correspond to our subjective view of space-time dynamics, and to global measures on the “bushiness” of typical sub-trees




Figure 7: Gliders, glider guns and compound gliders in $k=5$ 1d CA . (c) is a compound glider made up of two independent gliders locked into a cycle of repeating collisions. (d) is a glider with a period of 106 time-steps. (e) is a compound glider-gun.

in attractor basins, characterized by the distribution of in-degree sizes in their branching structure.

The look-up frequency can be represented by a histogram (figure 8) which distributes the total of $n \times w$ lookups among the 2^k neighbourhoods (shown as the fraction of total lookups), where n =system size, w =the window of time-steps defined and k =neighbourhood size. The Shannon entropy of this frequency distribution, the “input-entropy” S , at time-step t , for one time-step ($w=1$), is given by, $S^t = -\sum_{i=1}^{2^k} \left(\frac{Q_i^t}{n} \times \log \left(\frac{Q_i^t}{n} \right) \right)$, where Q_i^t is the look-up frequency of neighbourhood i at time t . In practice the measures are smoothed by being taken over a moving window of time-steps ($w=10$ in figure 8).

Figure 8 shows typical examples of ordered, complex and chaotic dynamics in 1d CA, with input-entropy plots and a snapshot of the lookup frequency histogram alongside. In ordered dynamics the entropy quickly settles at a low value with low or zero variance. In chaotic dynamics the entropy settles at a high value, but again with low variance. Both ordered and chaotic dynamics have low input-entropy variance. By contrast, in complex dynamics the entropy fluctuates erratically both up and down for an extended time, because glider collisions produce new gliders, often via a temporary zone of chaotic dynamics. Complex rules can be recognized by eye, subjectively. Input-entropy variance provides a non-subjective measure for recognizing complex rules automatically.

A related method of visualizing the entropy-variance is to plot input-entropy against the density of 1s relative to a moving window of time-steps. Each rule produces a characteristic cloud of points which lie within a parabolic envelope because high entropy is most probable at medium density, low entropy at either low or high density. Each complex rule produces a plot with its own distinctive signature, with high input-entropy variance. Chaotic rules, on the other hand, will give a flat, compact cloud at high entropy (at the top of the parabola). For ordered rules the entropy rapidly falls off with very few data points because the system moves rapidly to an attractor.

3.3 Automatically classifying rule-space

To distinguish ordered, complex and chaotic rules automatically, the mean input-entropy taken over a span of time-steps is plotted against the standard deviation of the input entropy. Figure 10 summarizes how random samples of $k=5, 6$ any 7 rules were classified by this method. For each rule, the data was gathered from 5 runs from random initial states, for 430 time-steps, discounting the first 30 to allow the system to settle, with $w=5$ as the size of the moving window of time-steps.

Chaotic rules are concentrated in the top left corner “tower”, ordered rules on the left with lower entropy. Complex rules have higher standard deviation, and are spread out towards the right. There is a fairly distinct boundary between ordered and chaotic rules, but a gradual transition from both towards the complex rules. As the standard deviation decreases glider interactions either become more frequent, transients longer, tending towards chaos, or less frequent, transients shorter, tending towards order. The plots for $k=6$ and $k=7$ rules indicate a greater frequency of chaotic rules at the expense of ordered and complex rules for greater k . The decrease in ordered rules is especially marked.

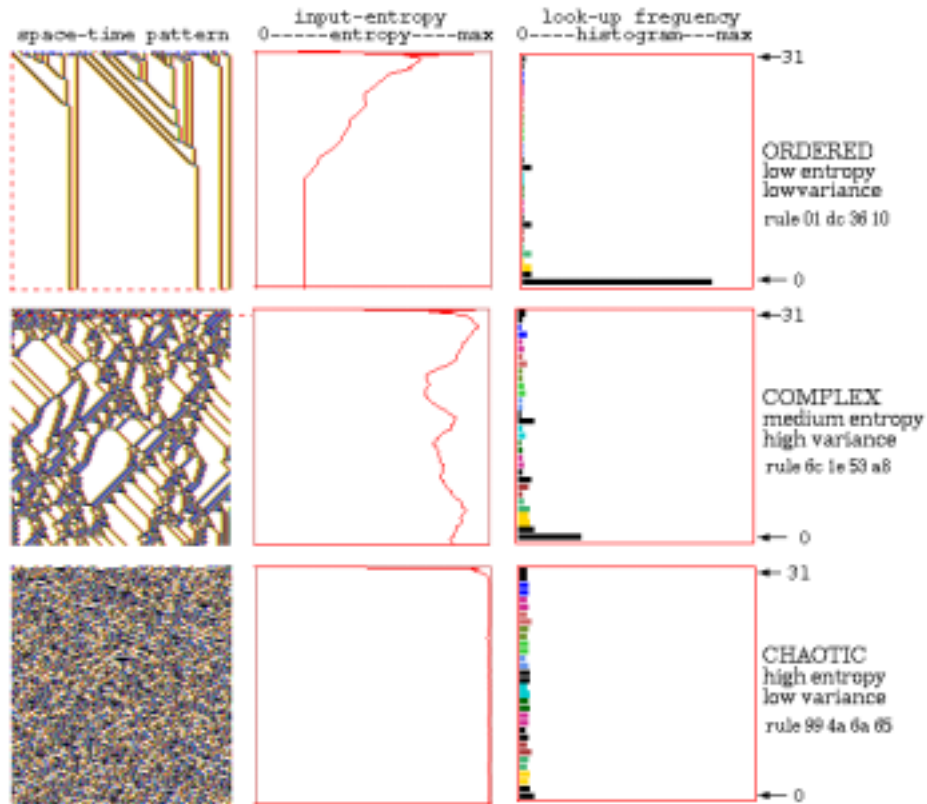


Figure 8: Typical 1d CA Space-time patterns showing ordered, complex and chaotic dynamics ($n=150$, $k=5$, rule numbers shown in hex). Alongside each space-time pattern is a plot of the input-entropy, where only complex dynamics (*centre*) exhibits high variance because glider collisions make new gliders.

To check whether the expected dynamics (recognized subjectively) corresponds to the measures as plotted, the dynamics of particular rules at different positions on the plots can be easily examined in DDLab, for example with a mouse click on the scatter plot. Preliminary scans confirm that the expected behaviour is indeed found, but further investigation is required to properly demarcate the space between ordered, complex and chaotic rules and to estimate the proportion of different rule classes for different k .

Input entropy is a local measure on the space-time patterns of typical trajectories. The distribution of the rule samples according to these local measures may be compared with global measures on convergence in attractor basins, G -density and the in-degree frequency, described in section 8. Preliminary results indicate a strong relationship between these global measures and the rule sample input-entropy plots.

4 RBN space-time patterns

In contrast to CA, glider dynamics in general cannot occur in RBN because of their irregular architecture. Figure 11 (left) shows glider dynamics degrading as local wiring is progressively scrambled. An alternative order-chaos notion in RBN is the balance between “frozen”, stabilized, regions and changing regions in the space-time pattern[8]. Stable regions are characteristic of RBN with low connectivity, $k \leq 3$, because rules which induce stability are relatively frequent in these rule-spaces. To induce stability for $k \geq 4$, where chaotic rules become overwhelmingly predominant, biases on rules must be imposed, low λ (see section 6.2) or a high proportion of “canalizing” inputs. In a rule’s lookup

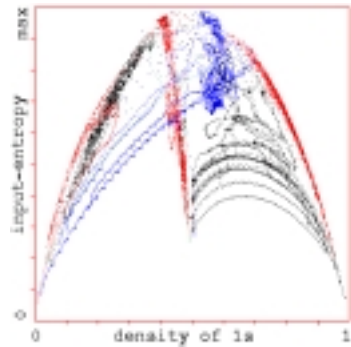


Figure 9: Entropy-density scatter plot. Input-entropy is plotted against the density of 1s relative to a moving window of time-steps $w=10$. $k=5$, $n=150$. Plots for a number of complex rules from the automatic sample (section 3.3) are shown superimposed, each of which has its own distinctive signature, with a marked vertical extent, i.e. high input-entropy variance. About 1000 time-steps are plotted from several random initial states for each rule.






Figure 10: *left*: Classifying a random sample of $k=5$ rules by plotting mean entropy against standard deviation of the entropy, with the frequency of rules within a 128×128 grid shown vertically. *below*: Equivalent plots for samples of $k=6$ and 7 rules.



table, an input wire is canalizing if a particular input (0 or 1) determines, by itself, the neighbourhood's output. A rule's degree of canalization can be from 0 to k , for the same output; for the network it is the percentage of all inputs that are canalizing, C . An RBN's order-chaos characteristics, for varying C , are captured by the measures illustrated in figure 12, and described below.

The "Derrida plot" [2], is analogous to the Lyapunov exponent in continuous dynamics, and measures the divergence of trajectories based on normalized Hamming distance H , the fraction of bits that differ between two patterns. Pairs of random states separated by H_t , are independently iterated forward by one (or more) time-steps. For a sample of random pairs, the average H_{t+1} is plotted against H_t , and the plot is repeated for increasing H_t (from 0 to 0.3 in figure 12). A curve above the main diagonal indicates divergent trajectories and chaos, below — convergence and order. A curve tangential to the main diagonal indicates a balance. A related measure is the distribution of "damage spread" resulting from a single bit change at a random position in a random state, for a sample of random states. The size of damage is measured once it has stabilized, i.e. not changed for say 5 time-steps. A histogram (figure 12) is plotted of damage size against the frequency of sizes. Its shape indicates order or chaos in the network, where a balance between order and chaos approximates to a power law distribution. Results by these measures for $k = 5$, indicate a balance at $C = 52\%$ (see figure 12). There are further measures on basins of attraction as in figure 14.

These methods are applied in the context of RBN models of genetic regulatory networks[8] dis-




Figure 11: Space-time patterns for intermediate 1d architecture, from CA to RBN. $n=150, k=5$, 150 time-steps from a random initial state. (a) Starting off as a complex CA (rule 6c1e53a8 as in figure 8), 4% (30/750) of available wires are randomized at 30 time-step intervals. The coherent pattern is progressively degraded. (b) A network with local wiring but mixed rules, vertical features are evident. (c) RBN, random wiring and mixed rules, with no bias, shows maximal chaotic dynamics.

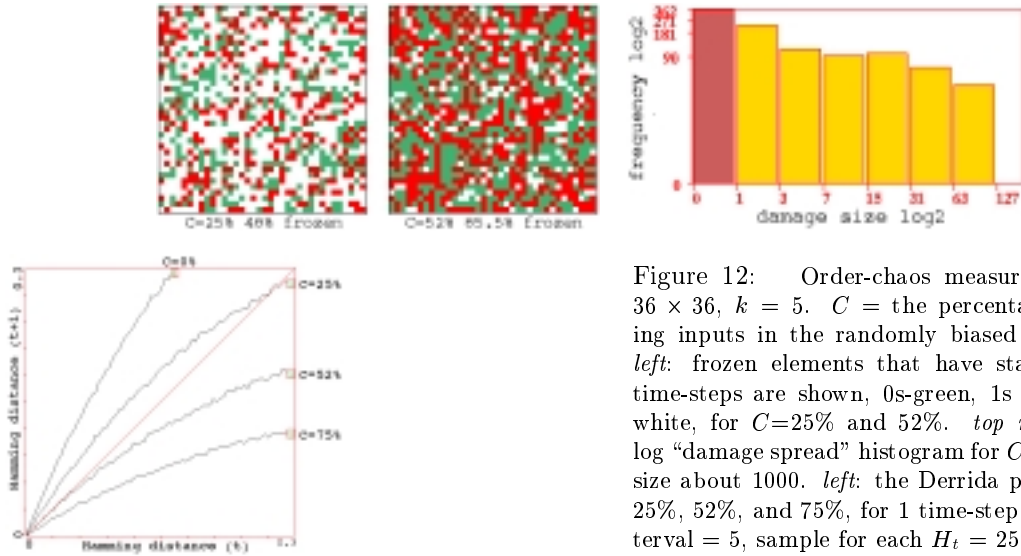


Figure 12: Order-chaos measures for a RBN $36 \times 36, k = 5$. C = the percentage of canalizing inputs in the randomly biased network. *top left*: frozen elements that have stabilized for 20 time-steps are shown, 0s-green, 1s red, otherwise white, for $C=25%$ and $52%$. *top right*: the log-log “damage spread” histogram for $C=52%$, sample size about 1000. *left*: the Derrida plot for $C=0%, 25%, 52%,$ and $75%$, for 1 time-step, $H_t=0-0.3$, interval = 5, sample for each $H_t = 25$.

cussed in section 10. The conjecture is that evolution maintains genetic regulatory networks marginally on the ordered side of the order-chaos boundary to achieve stability and adaptability in the pattern of gene expression which defines the cell type[5].

5 Basins of Attraction

The idea of basins of attraction in discrete dynamical networks is summarized in figure 13. Given invariant network architecture and the absence of noise, a discrete dynamical network is deterministic, and follows a unique (though in general, unpredictable) trajectory from any initial state. When a state that occurred previously is revisited, which must happen in a finite state-space, the dynamics becomes trapped in a perpetual cycle of repetitions defining the attractor (state cycle) and its period (minimum one, a stable point).

These systems are dissipative. A state may have multiple “pre-images” (predecessors), or none, but just one successor. The number of pre-images is the state’s “in-degree”. In-degrees greater than one require that transient states exist outside the attractor. Tracing connections backwards to successive pre-images of transient states will reveals a tree-like topology where the “leaves” are states without pre-images, known as garden-of-Eden states. Conversely, the flow in state-space is

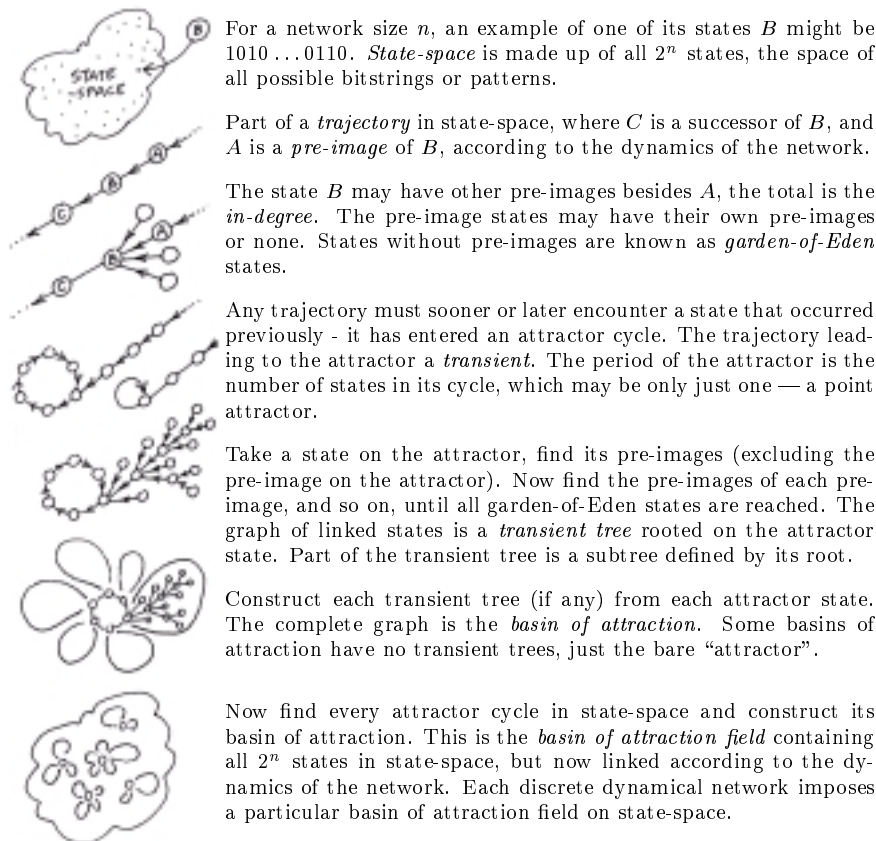


Figure 13: State space and basins of attraction.

convergent. Measures of convergence are G -density, the fraction of states that are garden-of-Eden, and the distribution of in-degrees, described in section 8. The set of transient trees rooted on the attractor is its basin of attraction (figure 1). The local dynamics connects state-space into a number of basins, the basin of attraction field, representing the systems global dynamics (figure 15).

6 Computing Pre-images

Attractor basins are constructed with algorithms that directly compute the pre-images of network states[16, 17, 21]. This allows the network’s dynamics, in effect, to be run backwards in time. Backward trajectories will, as a rule, diverge. Different reverse algorithms apply to networks with different sorts of connectivity. The most computationally efficient algorithm applies to 1d networks with local wiring, taking advantage of the regularity of connections. The wiring must be uniform, as for 1d CA, but the network may have a mix of rules. Analogous algorithms could be derived for 2d and 3d networks, but have not been implemented. An alternative algorithm is required for RBN with their non-local connections and possibly mixed k . This algorithm also applies to CA of any dimension or geometry, as CA are just a sub-class of RBN.

Provided $k \ll n$, these methods are in general orders of magnitude faster than the brute force method (section 9.1), constructing an exhaustive map resulting from network dynamics, a method which rapidly becomes intractable with increasing network size and so is limited to very small systems. However, the exhaustive method may be applied to all types of network, and also allows the attractor





Figure 14: Statistical data on attractor basins for a large network; a 2d RBN 20×20 , $k=5$, with fully random wiring and a fraction of canalizing inputs $C=50\%$. The histogram shows attractor types and the frequency of reaching each type from 12,360 random initial states, sorted by frequency. 46 different attractors types where found, their periods ranging from 4 to 102, with average transient length from 21 to 113 time-steps. The frequency of arriving at each attractor type indicates the relative size of the basin of attraction.

basins of random maps to be constructed, as described in section 9. The agreement of these three independent methods, and other checks, give considerable confidence in the accuracy of the pre-image computations.

Some basic information on attractor basin structure can be found by statistical methods, first applied by Walker[13], as shown in figure 14. These are also implemented in DDLab and are appropriate for large networks. Trajectories are run forward from many random initial states looking for a repeat in the network pattern to identify the range of attractor types reached. The frequency of reaching a given attractor type indicates the relative size of the basin of attraction, and other data are extracted such as the number of basins, and the length of transients and attractor cycles.

6.1 The CA reverse algorithm

Consider a 1d CA of size n (indexed $n - 1 \dots 0$) and neighbourhood k . To find the all pre-images of a state A , let P be a “partial pre-image” where at least $k - 1$ continuous bits (on the left) up to and including P_i , are known. Now find the next unknown bit to the right, P_{i-1} , consistent with the rule-table. (\bullet indicates known, \star unknown, bits),



If $k = 3$ (for example), the bitstring P_{i+1}, P_i, \star corresponds to two neighbourhood entries in the rule-table. When their outputs, T_1 and T_2 , are compared with each other and with A_i there are three possible consequences. The permutation is either deterministic, ambiguous or forbidden.

1. deterministic: if $T_1 \neq T_2$, then P_{i-1} is uniquely determined, as there is only one valid neighbourhood with the output A_i .
2. ambiguous: if $T_1 = T_2 = A_i$, then both 0 and 1 are valid solutions for P_{i-1} . The partial pre-image must be duplicated, with $P_{i-1} = 0$ in one version and $P_{i-1} = 1$ in the other.
3. forbidden: if $(T_1 = T_2) \neq A_i$, then P_{i-1} has no valid solution.

If forbidden (3) the partial pre-image P is rejected. If deterministic or ambiguous (1 or 2) the procedure is continued to find the next unknown bit to the right. However, in the ambiguous case (2), both alternative partial pre-images must be continued. In practice one is assigned to a stack of partial pre-images to be continued at a later stage. As the procedure is re-applied to determine each successive unknown bit towards the right, each incidence of ambiguous permutations will require another partial pre-image to be added to the stack. Various refinements can limit this growth.




Figure 15: The basin of attraction field of a random Boolean network ($n=13$, $k=3$). The $2^{13} = 8192$ states in state space are organized into 15 basins, with attractor periods ranging between 1 and 7. The number of states in each basin is: 68, 984, 784, 1300, 264, 76, 316, 120, 64, 120, 256, 2724, 604, 84, 428. Figure 1 shows the arrowed basin in more detail. Right: the network's architecture, its wiring/rule scheme.

cell	wiring	rule
12	10,1,7	86
11	6,2,9	4
10	10,10,12	196
9	2,10,4	52
8	5,6,8	234
7	12,5,12	100
6	1,9,0	6
5	5,7,5	100
4	4,11,7	6
3	8,12,12	94
2	11,6,12	74
1	6,5,9	214
0	12,9,6	188

The procedure is continued to the right to overlap the assumed start string, to check if periodic boundary conditions are satisfied; if so the the pre-image is valid. The procedure is re-applied to each partial pre-image taken from the partial pre-image stack, starting at the first unknown cell. Each time an ambiguous permutation (2) occurs a new partial pre-image must be added to the stack, but the stack will eventually be exhausted, at which point all the valid pre-images containing the assumed start string will have been found. The procedure is applied for 2^{k-1} start strings, assuming the different possible values of the first $k-1$ bits. The reverse algorithm is applied from left to right in DDLab, but is equally valid when applied from right to left. Examples are given in [16, 21].

6.2 The Z parameter

A by product of the CA reverse algorithm is the probability of the next unknown bit being *deterministic* (section 6.1(1)). Two versions of this probability are calculated from the rule-table. Z_{left} for the reverse algorithm applied from left to right, and Z_{right} for the converse. The Z parameter is the greater of these values. For $Z=1$ it can be shown[16] that for any system size n , the maximum in-degree, $I_{\text{max}} \leq 2^{k-1}$, because the next unknown bit is always uniquely determined, so the assumed start string of length $k-1$ may generate at most 2^{k-1} pre-images. If only one of Z_{left} or $Z_{\text{right}}=1$, $I_{\text{max}} < 2^{k-1}$, because at least one assumed start string must be forbidden (section 6.1(3)). At the other extreme, for $Z=0$, all state space converges on the state all-0s or all-1s in one step. For high Z , low in-degree (relative to system size n) is expected in attractor basins, growing at a slow rate with respect to n . Conversely, for low Z , high relative in-degree is expected growing quickly with respect to n . High Z

predicts low convergence and chaos, low Z predicts high convergence and order.

The 2^k neighborhoods of size k , each indexed $k-1 \dots 0$, each have an output T (0 or 1) which makes up the rule-table (section 2), and may be expressed as $a_{k-1}, a_{k-2}, \dots, a_1, a_0 \rightarrow T$. To calculate Z_{left} from the rule table, let n_k be the count of rule-table entries belonging to deterministic pairs, such that,

$$a_{k-1}, a_{k-2}, \dots, a_1, 0 \rightarrow T \text{ and } a_{k-1}, a_{k-2}, \dots, a_1, 1 \rightarrow \bar{T} \text{ (not } T)$$

The probability that the *next bit* is determined because of the above is given by, $R_k = n_k/2^k$. This is a first approximation of Z_{left} .

Let n_{k-1} be the count of rule-table entries belonging to deterministic 4-tuples (where “ \star ” may be 0 or 1), such that,

$$a_{k-1}, a_{k-2}, \dots, a_2, 0, \star \rightarrow T \text{ and } a_{k-1}, a_{k-2}, \dots, a_2, 1, \star \rightarrow \bar{T}$$

The probability that the *next bit* is determined because of the above is given by, $R_{k-1} = n_{k-1}/2^k$. This count is repeated if necessary for deterministic 8-tuples where $R_{k-2} = n_{k-2}/2^k$, 16-tuples, 32-tuples, \dots up to the special case of just one 2^k -tuple which occupies the whole rule-table. These are independent non-exclusive probabilities that the *next bit* is determined. The union of the probabilities $R_k \cup R_{k-1} \cup R_{k-2} \dots = Z_{\text{left}}$, is given by the following expression (the order of the probabilities makes no difference to the result),

$$\begin{aligned} Z_{\text{left}} &= R_k + R_{k-1}(1 - R_k) + R_{k-2}(1 - R_k + R_{k-1}(1 - R_k)) \\ &+ R_{k-3}(1 - (R_{k-2}(1 - R_k + R_{k-1}(1 - R_k)))) + \dots \end{aligned}$$

which simplifies to,

$$Z_{\text{left}} = R_k + R_{k-1}(1 - R_k) + R_{k-2}(1 - R_{k-1})(1 - R_k) + R_{k-3}(1 - R_{k-2})(1 - R_{k-1})(1 - R_k) + \dots$$

and may be expressed as² $Z_{\text{left}} = R_k + \sum_{i=1}^{k-1} R_{k-i} \left(\prod_{j=k-i+1}^k (1 - R_j) \right)$ where $R_i = n_i/2^k$, and n_i = the count of rule-table entries belonging to deterministic 2^{k-i} -tuples. A converse procedure gives Z_{right} , and the Z parameter = the greater of Z_{left} and Z_{right} . Examples are given in [16, 21].

By virtue of being a convergence parameter, Z is also an order-chaos parameter varying from 0(order) – 1(chaos). Z can be compared with Langton’s[9] well known λ parameter³. λ is an order-chaos parameter for CA which may have values greater than binary, and measures the density of “non-quiescent” outputs in a rule-table, so for just binary CA, $\lambda = c/2^k$ where c =the count of 1s a rule-table on k inputs. λ varies between 0 (order) – 0.5 (chaos) – 1 (order). To allow Z and λ to be compared, a normalized version of binary λ is defined[16], $\lambda_{\text{ratio}} = 2 \times c_{\text{min}}/2^k$ where c_{min} is the count of 0s or 1s in the rule-table, whichever is *less*. λ_{ratio} then varies from 0 (order) – 1 (chaos) just as Z .

Plots of G -density against both the λ_{ratio} and Z parameters, showing the discrepancies as well as similarities, are shown in figure 16, for the 256 $k = 7$ totalistic rules, which reduce to 136 non-equivalent rules in 72 clusters (having equal λ_{ratio} and Z). Points plotted in the top right corner of the λ_{ratio} graph represent λ_{ratio} values that do not correspond to behaviour as expected.

6.3 The RBN reverse algorithm

Consider an RBN of size n . Find all pre-images of a state A , $(A_{n-1}, A_{n-2}, \dots, A_0)$. Each network element A_i , has a pseudo-neighbourhood size k_i (assuming a mixed k network), indexed $k_i - 1, k_i - 2, \dots, 0$, a wiring scheme W_i , $(w_{k_i-1}, w_{k_i-2}, \dots, w_0)$, where w_j is a number between $n - 1$ and 0, the position of the wire connection from the j th branch of the pseudo-neighbourhood, and a rule-table R_i .

²Acknowledgment and thanks to Guillaume Barreau and Phil Husbands at COGS, Univ. of Sussex, for deriving this expression.

³Other versions of binary λ are “internal homogeneity” introduced earlier by Walker[13], and the P parameter, applied for RBN, which varies between 0.5 (chaos) – 1 (order). $P = c_{\text{max}}/2^k$ where c_{max} is the count of 0s or 1s in the rule-table, whichever is *more*, $P = 1 - \lambda_{\text{ratio}}/2$. A number of alternative order-chaos parameters have also recently been proposed, for example by Marty Zwick and Burton Voorhees.




Figure 16: G -density against both λ_{ratio} and Z for the set of $k = 7$ totalistic rules, $n=16$, for $Z \geq 0.25$. The complete basin of attraction field was generated for each rule and garden-of-Eden states counted.



Figure 17: Computing RBN pre-images. The changing size of a typical partial pre-image stack at successive elements. $n=24$, $k=3$.

To find the all pre-images of A , let P , $(P_{n-1}, P_{n-2}, \dots, P_0)$, be a candidate pre-image consisting of *empty* network elements, as yet unassigned to either 0 or 1. Starting with an element of A , A_i , assign bits from all valid pseudo-neighbourhoods in the rule-table R_i , i.e. that are consistent with A_i , to separate copies of P according to the wiring scheme W_i . As there will be a mix of 0s and 1s in R_i , only some of the 2^{k_i} possible pseudo-neighbourhoods will be valid. This will produce a stack of “partial pre-images” with some bits allocated and the remainder empty.

Now repeat the procedure for another element of A , say A_{i-1} , but this time independently for each partial pre-image previously created. If the allocation of a bit to a given partial pre-image conflicts with the bit already assigned, then the partial pre-image is rejected. Otherwise, the partial pre-image is added to the next generation of partial pre-images in a new stack. The allocation will be valid if it is made to an “empty” element, or to an allocated element with an equal bit. Valid allocation increases the size of the partial pre-image stack, conflicts reduce the size of the stack.

This procedure is repeated in turn for the remaining network elements of A . If the stack size is reduced to zero at any stage A has no pre-images. The algorithm works for any ordering of elements in A , though to minimize the growth of the partial pre-image stack, the order should correspond to the greatest overlap of wiring schemes. The changing size of the stack at successive elements can be displayed in DDLab, an example is shown in figure 17. When the procedure is complete, the final pre-image stack may still have empty network elements, which did not figure in any wiring scheme. These are duplicated so that all possible configurations at empty element positions are represented. The resulting pre-image stack is the complete set of pre-images of A without duplication.

The reverse algorithm for RBN works for networks with any degree of intermediate architecture between RBN and CA, including CA of any dimension. More detailed explanations of the algorithm are given in [17, 21].

7 Constructing and portraying attractor basins

To construct a basin of attraction containing a particular state, the network is iterated forward from the state until a repeat is found and the attractor identified. The transient tree (if it exists) rooted on each attractor state is constructed in turn. Using one of the reverse algorithms, the pre-images of the attractor state are computed, ignoring the pre-image lying on the attractor itself. Then the pre-images of pre-images are computed, until all “garden-of-Eden” states have been reached.

In a similar way, just a subtree may be constructed rooted on a state. Because a state chosen at random is very likely to be a garden-of-Eden state, it is usually necessary to run the network forward by at least one time-step, and use the state reached as the subtree root. Running forward by more steps will reach a state deeper in the subtree so allow a larger subtree to be constructed.

For CA, a considerable speedup in computation is achieved by taking advantage of “rotational symmetry”[16], a property of the regularity of CA and periodic boundary conditions, resulting in equivalent subtrees and basins.

Attractor basins are portrayed as state transition graphs, vertices (nodes) connected by directed edges. States are represented by nodes, by a bit pattern in 1d or 2d, or as the decimal or hex value of the state. In the graphic convention[16, 19], the length of edges decreases with distance away from the attractor, and the diameter of the attractor cycle approaches an upper limit with increasing period. The direction of edges (i.e. time) is inward from garden-of-Eden states to the attractor, and then clockwise around the attractor cycle, as shown in figure 1. Typically, the vast majority of states in a basin of attraction lie on transient trees outside the attractor, and the vast majority of these states are garden-of-Eden states.

8 Attractor basin measures




Figure 18: The G -density plotted against system size n , for the ordered, complex and chaotic rules shown in figures 8 and 19. The entire basin of attraction field was plotted for $n=7$ to 22, and garden-of-Eden states counted. The relative G -density and rate of increase with n provides a simple measure of convergence.

Measures on attractor basins include the number of attractors, attractor periods, size of basins, characteristic length of transients and the characteristic branching within trees. The last in particular gives a good measure of the convergence of the dynamical flow in state-space, where high convergence indicates ordered, and low convergence indicates chaotic dynamics.

The simplest measure that captures the degree of convergence is the density of garden-of-Eden states[18], G -density, counted in attractor basins or sub-trees, and the rate of increase of G -density with n as shown in figure 18. A more comprehensive measure is the in-degree frequency distribution, plotted as a histogram. The in-degree of a state is the number of its immediate pre-images. This can be taken on a basin of attraction field, a single basin, a subtrees, or on just part of a subtree for larger systems. Subtrees are portrayed as graphs showing trajectories merging onto the sub-tree root state.

Examples of in-degree histograms for typical sub-trees of ordered, complex, and chaotic rules are shown in figure 19. The horizontal axis represents in-degree size, from zero (garden-of-Eden states) upwards, the vertical axis represents the frequency of the different in-degrees. The system size $n=50$ for the complex and chaotic rules. For very ordered rules in-degrees become astronomical. The ordered rule shown is only moderately ordered, however the system size was reduced to $n=40$ to allow easier computation.

From the preliminary data gathered so far, the profile of the in-degree histogram for different classes of rule is as follows:

Ordered rules: Very high garden-of-Eden frequency and significant frequency of high in-degrees. High convergence.

Complex rules: Approximates a power law distribution. Medium convergence.

Chaotic rules: Lower garden-of-Eden frequency compared to complex rules, and a higher frequency of low in degrees. Low convergence.

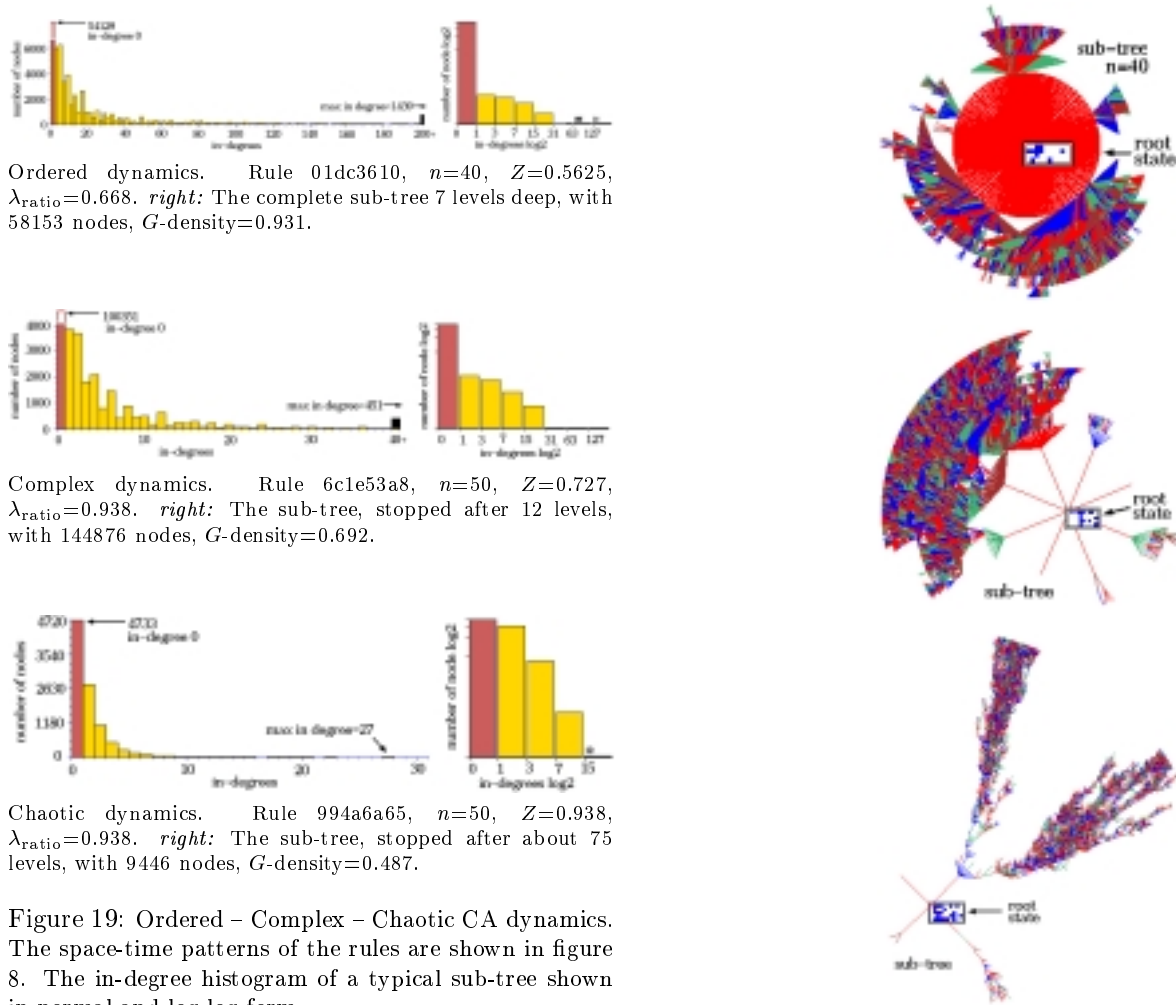


Figure 19: Ordered – Complex – Chaotic CA dynamics. The space-time patterns of the rules are shown in figure 8. The in-degree histogram of a typical sub-tree shown in normal and log-log form.

9 Random Maps

The attractor basins of discrete dynamical networks can be put into the wider context of random graph theory. CA belong to the set of RBN which in turn belong to the set of random directed graphs with out degree one, known as random maps. This is a mapping of the Boolean hypercube of sequences of length n (i.e. a set Q_2^n of size 2^n comprising all binary strings of length n), a mapping from $Q_2^n \rightarrow Q_2^n$. The structures found in random maps correspond to those in the attractor basins of discrete dynamical networks, where each separate component of the graph is made up of trees rooted on just one closed cycle. The structures can be computed in DDLab just as the attractor basins of CA or RBN.

A random map can be constructed by assigning a successor to each state in state-space, i.e. independently assign one successor (or image) V_* also belonging to Q_2^n , chosen at random (or with some bias) to each element V_i of the set Q_2^n . There are $(2^n)^{(2^n)}$ possible mappings. The mapping is represented below as 2^n pairs of strings (or states in state-space), where each image V_* represents a possibly different member of the set Q_2^n .

$$\begin{array}{ccccccc}
 V_{2^n-1} & V_{2^n-2} & \dots & V_i & \dots & V_2 & V_1 & V_0 \\
 \downarrow & \downarrow & & \downarrow & & \downarrow & \downarrow & \downarrow \\
 V_* & V_* & \dots & V_* & \dots & V_* & V_* & V_*
 \end{array}$$




Figure 20: The basin of attraction field of a typical unbiased random map, $n=12$. The $2^{12} = 4096$ states in state space are connected into 9 basins of attraction. The period (p) and size (s) of the biggest three (top row), including their percentage of state-space, are as follows: (1) $p=118$ $s=3204=78.2\%$. (2) $p=20$ $s=599=14.6\%$. (3) $p=32$ $s=194=4.74\%$. The field's G -density=0.37, this is a low value implying chaotic dynamics.

The list of images is likely to contain repeats, and if so some other members of Q_2^n must be missing from the list. Transitions to some arbitrary element V_x may thus be one-to-one or many-to-one, or may not exist. The latter is a garden-of-Eden state in the terminology of discrete dynamical networks. A representation of the particular mapping may be drawn as a basin of attraction field or fragment thereof just as for CA or RBN, and will have the same general topology of trees rooted on attractor cycles, as shown in figure 20.

Random maps provide the most general context for a discrete dynamical system and are equivalent to a fully connected RBN where $k = n$, (the neighbourhood = the network size). This follows because each cell in the RBN can be assigned an arbitrary output for any network state.

9.1 The Random Map Reverse Algorithm

The “brute force” reverse algorithm for finding the pre-images of states in random maps can also be applied to discrete dynamical networks, RBN and CA. The method depends on first constructing an exhaustive mapping $Q_2^n \rightarrow Q_2^n$. For discrete dynamical networks, the mapping is defined by iterating the network forward by one step from every state in state-space and filling in the image list accordingly. A list of 2^n pairs, each state and its image (successor), is held in a data structure. The pre-images of an arbitrary state S are found by scanning the image list; any occurrence of S in the list gives a pre-image, the state paired with S . If S does not occur in the list it has no pre-images, a garden-of-Eden state.

10 Biological networks

Genetic regulatory networks have been thought of as discrete dynamical networks, to explain how gene expression is able to settle into a number of distinct stable patterns or cell types, despite the fact that all eukaryotic cells in an organism carry an identical set of genes[5, 7, 11, 23]. The gene expression pattern of a cell needs to be stable but also adaptable. Section 4 described biases to RBN to achieve such a balance, and related measures.

Cell types have been interpreted as the separate attractors or basins of attraction into which network dynamics settles from various initial states. Trajectories leading to attractors are seen as

the pathways of differentiation. The attractor basins in RBN are idealized models for the stability of cell types against mutations, and also perturbations of the current state of gene activation. Figure 21 illustrates both effects. If a particular reference state (pattern of gene activation) undergoes a 1 bit perturbation, the dynamics may return to the same subtree, the same basin, or it may be flipped to another basin, a different cell type. In this case the basin of attraction field remains unchanged. Alternatively, the network itself may undergo a mutation (in the genotype), resulting in an altered basin of attraction field (the phenotype).

The examples in figure 21 are small so that the pattern at each node can be shown. Larger networks are affected in analogous ways. The consequences of a one bit mutation has a relatively smaller effect with increasing network size. However, a particular one bit mutation may cause drastic consequences whatever the size, such as breaking an attractor cycle. The consequences of moving a connection wire is usually greater than a one bit mutation in a rule.

10.1 Memory

Attractors classify state-space into broad categories, the network's "content addressable" memory in the sense of Hopfield[6]. Furthermore, state-space is categorized along transients, by the root of each subtree forming a hierarchy of sub-categories. This notion of memory far from the equilibrium condition of attractors greatly extends the classical concept of memory by attractors alone[17, 20].

It can be argued that in biological networks such as neural networks in the brain or networks of genes regulating the differentiation and adaptive behavior of cells, attractor basins and subtrees, *the network's memory*, must be just right for effective categorization. The dynamics need to be sufficiently versatile for adaptive behavior but short of chaotic to ensure reliable behavior, and this in turn implies a balance between order and chaos in the network.

A current research topic, known as the "inverse problem", is to find ways to deduce network architecture from usually incomplete data on transitions, such as a trajectory. This is significant in genetics, to infer the genetic regulatory network (modeled as RBN) from data on successive patterns of gene expression in the developing embryo[12]. In pattern recognition and similar applications in the area of artificial neural networks, solutions to the inverse problem would provide "learning" methods for RBN to make useful categories[17, 20].

11 Conclusion

Important insights may be gained by considering network dynamics in the context of attractor basins. Some methods of achieving this have been presented, including parameters and measures on particular trajectories that may be related to those on global dynamics. It is hoped that these methods may provide a basis for future research, both in theory and applications, in the many areas of complex systems where network dynamics plays a central role.

Acknowledgments: Thanks to Cosma Shalizi for suggestions, and to the many people who have contributed to this work over the years, notably Chris Langton and Stuart Kauffman.

References

- [1] Conway, J.H., (1982) "What is Life?" in *Winning ways for your mathematical plays*, Berlekamp, E., J.H. Conway and R. Guy, Vol.2, chap.25, Academic Press, New York.
- [2] Derrida, B., and D. Stauffer, (1986) "Phase transitions in Two-Dimensional Kauffman Random Network Automata", *Europhys. Lett.* 2, 739.
- [3] Douglas, R., and A. Wuensche, work in progress.
- [4] Hanson, J.E., and J.P. Crutchfield (1997) "Computational Mechanics of Cellular Automata, An example", *Physica D*, vol. 103, 169-189.
- [5] Harris, E.S., B.K. Sawhill, A. Wuensche, and S. Kauffman, (1997) "Biased Eukaryotic Gene Regulation Rules Suggest Genome Behaviour is Near Edge of Chaos", Santa Fe Institute Working Paper 97-05-039.

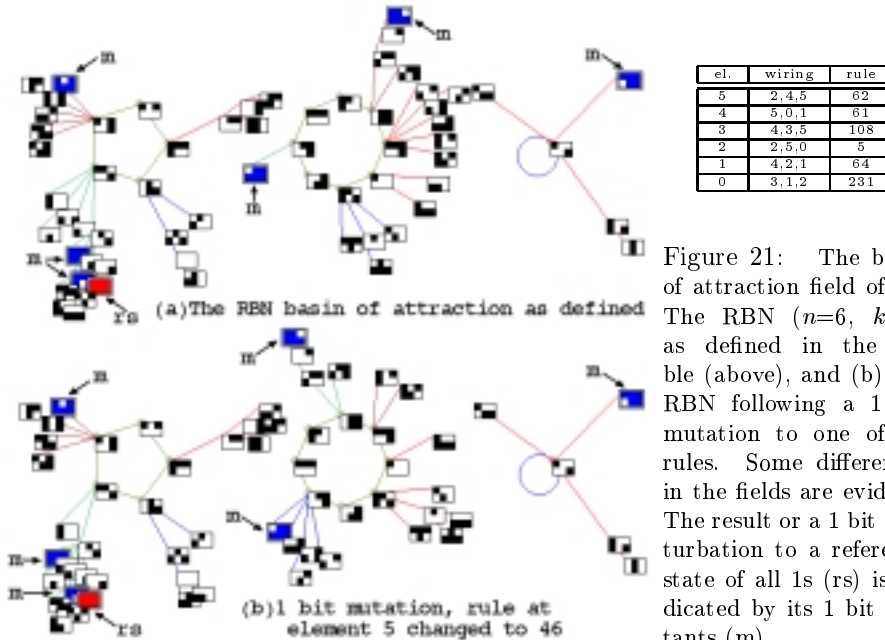


Figure 21: The basin of attraction field of (a) The RBN ($n=6$, $k=3$) as defined in the table (above), and (b) the RBN following a 1 bit mutation to one of its rules. Some differences in the fields are evident. The result of a 1 bit perturbation to a reference state of all 1s (rs) is indicated by its 1 bit mutants (m).

- [6] Hopfield, J.J. (1982) "Neural networks and physical systems with emergent collective computational abilities", *Proceedings of NAS* 79:2554-2558.
- [7] Kauffman, S.A., (1969) "Metabolic stability and epigenesis in randomly constructed genetic nets", *Journal of Theoretical Biology*, 22, 437-467.
- [8] Kauffman, S.A., (1993) "*The Origins of Order*", Oxford University Press.
- [9] Langton, C.G., (1990) "Computation at the Edge of Chaos, phase transitions and emergent computation", *Physica D* 42, 12-37.
- [10] Myers, J.E., (1997) "Random Boolean Networks — Three Recent Results", to be appear in *Complexity*.
- [11] Somogyi, R., and C. Sniegowski, (1996) "Modeling the Complexity of Genetic Networks: Understanding Multigenetic and Pleiotropic Regulation", *Complexity*, Vol.1/No.6, 45-63.
- [12] Somogyi, R., S. Fuhrman, M. Askenazi, A. Wuensche., (1997) "The Gene Expression Matrix", in *Proceedings of the World Congress of Non-Linear Analysis*, in press.
- [13] Walker, C.C., and W.R. Ashby, (1966) "On the temporal characteristics of behavior in certain complex systems", *Kybernetik* 3, 100-108.
- [14] Wolfram, S., (1984) "Universality and complexity in cellular automata", *Physica* 10D, 1-35.
- [15] Wolfram, S., ed. (1986) "*Theory and Application of Cellular Automata*", World Scientific.
- [16] Wuensche, A., and M.J. Lesser. (1992) "*The Global Dynamics of Cellular Automata*", Santa Fe Institute Studies in the Sciences of Complexity, Addison-Wesley.
- [17] Wuensche, A., (1994) "The Ghost in the Machine", in *Artificial Life III*, ed C.G. Langton, Santa Fe Institute Studies in the Sciences of Complexity, Addison-Wesley.
- [18] Wuensche, A., (1994) "Complexity in One-D Cellular Automata", Santa Fe Institute Working Paper 94-04-025.
- [19] Wuensche, A., (1996) "Discrete Dynamics Lab (DDLab)", <http://www.santafe.edu/~wuensche/ddlab.html>
- [20] Wuensche, A., (1996) "The Emergence of Memory", in *Towards a Science of Consciousness*, eds. S.R. Hameroff, A.W. Kaszniak, A.C. Scott, MIT Press.
- [21] Wuensche, A., (1997) "Attractor Basins of Discrete Networks", CSRP 461, Univ. of Sussex (D.Phil thesis).
- [22] Wuensche, A., (1998) "Classifying Cellular Automata Automatically", Santa Fe Institute Working Paper 98-02-018.
- [23] Wuensche, A., (1998) "Genomic regulation modeled as a network with basins of attraction", in *proceedings of Pacific Symposium on Biocomputing'98*, World Scientific.

Part II
Biology

Emergent symmetry of local and global maps in the primary visual cortex: Self-organization of orientation preference

David M. Alexander
Brain Dynamics Laboratory,
Mental Health Research Institute,
Melbourne, Vic, Australia
dalex@mhri.edu.au

Phil Sheridan
School of Information Technology,
Charles Sturt University,
Bathurst, NSW, Australia

Paul D. Bourke
Brain Dynamics Laboratory,
Mental Health Research Institute,
Melbourne, Vic, Australia

Otto Konstandatos
School of Mathematics,
University of Melbourne,
Melbourne, Vic, Australia

James J. Wright
Brain Dynamics Laboratory,
Mental Health Research Institute,
Melbourne, Vic, Australia

Abstract

We present a model of the cellular organization in the primary visual cortex which is based upon the idea that there are two visuotopic mappings, one global and the other local. The local visuotopic maps self-organize from a set of initially random inputs. The various response properties measured in V1, and the regular geometrical relationships between them, are explained in terms of the interaction of these two mappings. We describe computer modelling of orientation preference in V1 which relies on two assumptions: 1) the receptive fields develop due to a simple excitatory-centre / inhibitory-surround mechanism, and 2) any point in the global visuotopic map can reach any point in the non-granular layers via poly-synaptic routes. Hebbian learning is applied to these indirect inputs while the network is stimulated with a moving bar. The computer modelling shows that all the key geometrical features of orientation preference; singularities, linear zones and saddle-points, emerge consequent to the learning. This suggests that orientation preference is a byproduct of the double visuotopic mapping. More generally, the two mappings are hypothesized to allow those features of the visual field which tend to be spatially contiguous in the visual field (orientation, texture, colour, contrast) to be learned as response properties by neurons within V1.

1 Introduction

This paper presents a model of the organisation of the mammalian primary visual cortex. The core proposition is the existence of two, distinct (but topologically identical) mappings of the visual field into the primary visual cortex. The first mapping is the classical retinotopic, global projection of the visual field to V1, and the second is a local mapping of visual field inputs, multiple copies of which tile the non-granular layers of V1. This second mapping emerges due to Hebbian learning applied to diffuse, convergent projections. The only *a priori* requirement for self-organization of these input connections is the presence of a local connectivity in the non-granular layers which drives a simple excitatory-centre / inhibitory-surround mechanism.

Together these two mappings define the local-global symmetry (LGS) model of V1. The interaction between the two mappings supplies non-retinotopic response properties such as orientation preference. The LGS model of the primary visual cortex suggests there is a rather direct relationship between globally represented objects (e.g. an oriented line) and various other response properties (e.g. orientation preference) which have a local geometry.

The function of the LGS mapping is to allow each point in the retinotopic image access to a representation of the entire visual hemi-field. Each retinotopic point is then able to make associations with points in the visual image whose activity tends to coincide with that retinotopic point. Any oriented lines which pass through a particular retinotopic points can become associated with that retinotopic point. The process of self-organization in the model follows principles related to Phillips and Singer's[15] concept of *coherent infomax*, and may be regarded as a concrete demonstration of this concept in action in anatomical organization. Points in the visual image can discover those visual contexts which are predictively related to their own activity.

The key finding reported in this paper is that the local mappings of the visuotopic field spontaneously emerge under a robust range of conditions. The critical mechanism is a standard excitatory-centre / inhibitory-surround interaction between neurons, the so-called Mexican hat field. The only additional assumption is that any point in the global visuotopic map can influence the activity at any other point. Since these interactions are initially weak and can occur via poly-synaptic routes, this assumption is also reasonable.

Local mappings of the visuotopic field emerge robustly under a range of parameter variations. Assumptions of the model are:

1. cortical elements within a local neighbourhood interact via a standard excitatory-centre / inhibitory-surround (the so-called Mexican hat field) and this organization is *a priori* to the model
2. The retina projects to the cortex in the standard visuotopic manner, but additional to this direct topographic mapping, any retinal point may contribute some input to any point in the non-granular layers. These diffuse projections are initially weak and are considered to arise from converging and diverging polysynaptic pathways within V1.
3. Hebbian learning applies to synapses supplying cortical input to the non-granular layers via the indirect, diffuse pathway.

2 Response Property Geometry and Connectivity of V1

The response properties of V1 are well catalogued. These properties have been mapped through single-cell studies [11], metabolic transport studies [23, 21, 22] and optical imaging of the cortical surface [4, 8]. These studies have shown that V1 exhibits a distinctive tiling of various response properties [4]. Swindale[20] has described a set of canonical properties which a model of the geometry of the primary visual cortex should take into account. These properties include the spatial frequencies and organisation of: ocular dominance bands, cytochrome oxidase (CO) blobs, singularities and iso-orientation regions.

Layer 4C of the macaque primary visual cortex has a strict retinotopic organisation [5]. In the macaque, each hemi-retina becomes represented, through a complex logarithmic transformation, as an almond shaped map in layer 4C of V1 [23]. This transformation is of fundamental importance to the functionality of the visual recognition system, allowing computational simplification for operations such as rotation and scaling in two dimensions [17, 18].

Neurons in the non-granular layers of the primary visual cortex are organised into repeated units, roughly $800\mu\text{m}$ wide and $600\mu\text{m}$ high, called hypercolumns [11]. Each hypercolumn spans an entire range of orientation tunings and a left and right ocular dominance set. Located along the centres of ocular dominance bands are cytochrome-oxidase (CO) blobs. The blobs take up higher levels of CO staining because of their higher metabolic activity [10], and are responsive to colour and low contrast. Interblob regions are more selective for high contrast [21, 22].

The geometry of orientation preference reveals three predominant features: singularities, linear zones and saddle regions [4]. Orientation preference changes continuously around points, or singularities. These roughly circular "pinwheels" traverse 180° of all possible orientation preference. Between adjacent singularities, running parallel to the ocular dominance bands, are regions in which orientation




Figure 1: Map of orientation preference in the Tree Shrew. Orientation preference varies from 0 (white) to π (black). (Data provided by Bill Bosking).

Table 1: Analogies between the properties of the hemi-retinal image and the local organisation of receptive field properties in the non-granular layers. CO blobs are assumed to be the central visual fields representation in the local maps of the non-granular layers.

Global Mapping	Local Mapping
oriented lines	orientation preference
<u>fovea</u> \rightarrow <u>periphery</u>	<u>CO blob</u> \rightarrow <u>interblob</u>
decreasing ganglion cell density	decreasing metabolic activity
decreasing density of colour cones	decreasing colour selectivity
decreasing contrast sensitivity	low \rightarrow high contrast preference

preference changes slowly and continuously. These regions are called linear zones. Other regions between singularities show local minima in orientation preference in one direction and a local maxima in a perpendicular direction: so-called saddle points. Singularities surrounding saddle-points form mirror images of each other, through orthogonal reflection lines. The double reflections have the effect of allowing orientation selectivity to change continuously between adjacent pinwheels. An orientation preference map, in this case taken from the tree shrew, is shown in Figure 1.

Alexander et al [2] have noted strong analogies between the global properties of the hemi-retinal image and the response properties in the non-granular layers. These are given in Table 1. In particular, we assume these analogies with the hemi-retinal image apply to a geometrical unit in the non-granular layers corresponding to 1/4 of a hypercolumn. Such a unit includes one CO blob (or portion of an elongated CO blob) and one singularity and has the approximate dimensions in the macaque of $400\mu m \times 300\mu m$. Figure 2 shows a region of macaque visual cortex in which these geometrical units have been abstracted from imaging and staining data [1].

Anatomically, the mapping into layer 4C occurs via the Lateral Geniculate Nucleus (LGN) in the thalamus [6]. The non-granular layers receive direct input from 4C, as well as indirect inputs via various of the other non-granular layers [13, 14]. In addition, a large percentage of their inputs arrive via lamina 4A and 4B [7]. Of particular interest for the model presented in this paper are the lateral projections of lamina 4B spiny stellates within 4B itself. These efferent axons project laterally up to 4.5 mm in the macaque [7]. A similar pattern of projecting fibres is found in the squirrel monkey [16]. The lateral connections show periodic accumulations of denser terminal fibres every $375 - 400\mu m$ [7, 16]. These patches of connections form a radial pattern with a similar spatial periodicity to the CO blobs [16]. The fibres within 4B extend further than any other class of intrinsic fibres within the primary visual cortex. Only cells within layer 4B send out this class of long range lateral fibres and, at mid- to long-range distances from the cell body, all terminating boutons are within 4B itself. Most of the fibres within 4B are preferentially horizontal, rather than vertical [16].

The lateral connections within laminae 4A and 4B are proposed as the primary routes for the all to all mapping required for the development of local visuotopic maps. They guarantee that any point of the retinotopic image in 4C can project to any point in the non-granular layers within about 6 synaptic junctions. Convergence and divergence of the signal transmission at all way-stations of the



Figure 2: Abstracted map of the tiling of response properties across the surface of V1. Circles are local map representations of the foveal region (CO blobs or portions of elongated blobs), stars are singularities. Tile borders are defined largely by ocular dominance bands and reversals in the orientation preference gradient. (Original animal data from [3].)

visual pathway offers further potential diffusion of the retinal image to the non-granular layers.

In the LGS model, the global input-mapping supplies the retinotopic response properties of the non-granular layers. Other response properties of the local map representations are formed through an interaction of the local input-maps with the global retinotopic input-mapping into the non-granular layers. Here we distinguish between the *local input-map*, and the *local representational maps* (or local response property maps) the latter arising largely as a result of interactions between the two types of input-maps when driven by visual stimuli.


3 Modelling of the Local-Global Symmetry Mapping

The present modelling of V1 introduces a few simplifications into the visual system, and will not explicitly deal with ocular dominance, the cortical magnification factor, contrast and colour selectivity. The resulting simplified model focuses on the development of orientation selectivity, and aims to explain data from mammals such as the macaque, the tree shrew and the ferret. A more complete model is in development.

A moving bar was repeatedly swept across a simplified model retina. Activity in this retina then drove activity in the model non-granular layers, via the direct and indirect pathways. The effects local neighbourhood interactions, via a simplified excitatory-centre and inhibitory-surround mechanism, were then added. The present model is a static, time-averaged portrayal of what is never-the-less a *dynamical* wave-medium [24]. While more realistic dynamics will add to the explanatory scope of the LGS model, the present description captures the essential of the *structural* mechanisms involved in the formation of the local-input maps.

The activity of each non-granular unit (approximately a minicolumn, $30\mu m$ in diameter) was calculated in three steps. The symbols a , α , A and \mathbf{A} denote the activations at successive stages and include the effect of direct (retinotopic) input connections, local input-map connections, local excitatory influence and inhibitory surround, respectively and cumulatively. A single line was presented to the simulated retina, and the line swept across, each sweep taking place at randomly chosen angles. The activation, a , on the retina was fed directly to the model cortical minicolumns. No learning occurred in relation to these retinotopic projections, as they can be assumed to be largely innately provided. The local input-map connections, w_{jk} , were initially set to low values ($0 \leq w_{jk} \leq 0.01$ drawn from a uniform random distribution). Equation (1) describes the effects of the direct and local map connections on the activation of each minicolumn.

Figure 3: Grey-scale map of orientation preference. The map is produced in a manner analogous to the orientation preference maps in animals. Orientation preference varies from 0 (white) to π (black).



$$\alpha_k = a_k + \sum_{j=1}^v w_{jk} a_j + \epsilon, \quad a_j \in \{0, 1\}, \quad 0 \leq \sum_{j=1}^v w_{jk} \leq 0.5, \quad w_{jk} \leq 0 \quad (1)$$

where v is number of discrete retinotopic locations in the visual field. An amount of noise, ϵ , was added at this stage. The modelling results proved robust to a wide range of noise values ($0 \leq \epsilon \leq 0.1$). The excitatory-centre was implemented by simply averaging the activities of each cell with its immediate neighbours:

$$A_k = \frac{\alpha_k + \sum_{j=1}^e \alpha_j}{e + 1} \quad (2)$$

where e is the number of α_j 's within the excitatory radius. In the present modelling the excitatory radius was set to one minicolumn. The inhibitory-surround was implemented in the following fashion: each cell's activity, if lower than one of its neighbouring cells, was multiplied by the ratio of its own activity over the other cell's activity:

$$\mathbf{A}_k = \frac{A_k^i}{\prod_{j=1}^i A_k}, \quad (3)$$

where i is the number of A_j 's within the inhibitory radius that satisfy the condition $A_j > A_k$. The inhibitory radius was set to a distance of 7 minicolumns.

The learning rule was a simple variant of Hebbian learning. The local input-map connections were strengthened if the equivalent point in the retinotopic (input) layer was active and in proportion to the activation of the cell in the non-granular (output) layer.

$$\delta w_{jk} = d(\mathbf{A}_k - \mathbf{A}_{\min})a_j \quad (4)$$

where \mathbf{A}_{\min} was the minimum activation of the cells influenced by the direct retinotopic input and d was a constant affecting the learning rate. d could be varied over a broad range of values ($0.000003 \leq d \leq 0.01$) without any substantive effect on the simulation results.

3.1 Results

An example of the model cortex, activated by a bar, is shown in Figure 3. Results of the modelling, in the form of a grey-scale orientation preference map, are also shown. The simulation was run for 100,000 time-steps or 2,000 sweeps of the stimulus bar. The orientation preference map is produced in a manner analogous to animal studies. It shows singularities, saddle-points and linear zones.

Each local map is driven by retinotopic activity appropriate to its particular retinotopic location. The local input-map connections become organized according to the patterns of activity in the entire visual field which reliably coincide with activity at that retinotopic location. In the case of line orientation preference, each local map 'sees' only the subset of lines relevant to that local map, that is all lines which actually pass through that point. The double mapping has the effect of adding an orientation label to each 'pixel' in the global image of the line.

The model suggests very specifically the function of V1. It takes as inputs a set of active retinotopic locations. By creating a map within a map, points of high activity serve as feature labels. Information is supplied to later processing stages that, not only is this retinotopic point active, but it is active because of a particular line of orientation, θ . Each retinotopic location learns which other points in the visual field are relevant to its own activation.

4 Conclusions

This model described in this paper relies on four assumptions to explain orientation selectivity in the primary visual cortex. These are: a global, topology-preserving mapping from the retina to the non-granular layers via the laminae 4C (uncontroversial); an initial, diffuse, disorganized all to all mapping of the retinotopic inputs to the non-granular layers; that receptive field properties are organized through a centre-excitation / surround-inhibition mechanism, and; that the response properties, in particular orientation preference, result from the interaction of local and global mappings.

Orientation, texture, colour and contrast tend to be spatially contiguous elements of the visual field. The LGS mapping of visual inputs in V1 allows precisely these properties of the visual field to be learned as response properties, and their contiguity in visual space means the local maps will tend to be visuotopic. The local mapping to the non-granular layers makes the contents of the entire visual field available for association with each retinotopic location in the non-granular layer. It is a structural embodiment of Phillips and Singer's [15] coherent info-max.

The local mapping described in this paper has correspondences to the ice-cube model of Hubel and Wiesel [12]. Both assume the upper layers are tiled with a regular map of the feature space. The LGS model has the additional features that it 1) directly suggests the function of known anatomical connections, and 2) explains how the response properties are created.

The model assumes that the geometry of V1 serves to organise multiple response properties into two dimensions. It differs from other models of V1 geometry which, rather than mappings, use simulated annealing or other relaxation procedures to achieve this dimension reduction [9, 19]. The latter models have multiple free parameters [20]. By contrast, the LGS model of V1 contains a limited number of assumptions and the self organization is robust to over a fairly large range of learning and noise parameters.

Acknowledgements

The research was supported by the Australian Research Council.

References

- [1] D.M. Alexander and et al. Manuscript in preparation, 1998.
- [2] D.M. Alexander, P. Sheridan, P.D. Bourke, and O. Konstantatos. Global and local similarity of the primary visual cortex: mechanisms of orientation preference. In *Proceedings HELNET97*, 1997.
- [3] E. Bartfeld and A. Grinvald. Relationships between orientation-preference pinwheels, cytochrome oxidase blobs, and ocular-dominance columns in primate striate cortex. *Proceedings of the National Academy of Sciences USA*, 89:11905–11909, 1992.
- [4] G.G. Blasdel. Orientation selectivity, preference, and continuity in monkey striate cortex. *Journal of Neuroscience*, 12(8):3139–3161, 1992.
- [5] G.G. Blasdel and Fitzpatrick. Physiological organization of layer 4 in macaque striate cortex. *Journal of Neuroscience*, 4(3):880–895, 1984.

- [6] G.G. Blasdel and J.S. Lund. Termination of afferent axons in macaque striate cortex. *Journal of Neuroscience*, 3(7):1389–1413, 1983.
- [7] G.G. Blasdel, J.S. Lund, and D. Fitzpatrick. Intrinsic connections of macaque striate cortex: axonal projections of cells outside lamina 4c. *Journal of Neuroscience*, 5(12):3350–3369, 1985.
- [8] G.G. Blasdel and G. Salama. Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321:579–585, 1986.
- [9] R. Durbin and G. Mitchison. A dimension reduction framework for understanding cortical maps. *Nature*, 343:644–647, 1990.
- [10] J.C. Horton and D.H. Hubel. Regular patchy distribution of cytochrome oxidase staining in primary visual cortex of the macaque. *Nature*, 292:762–764, 1980.
- [11] D.H. Hubel and T.N. Wiesel. Receptive fields and functional architecture of the monkey striate cortex. *Journal of Physiology (London)*, 195:215–243, 1968.
- [12] D.H. Hubel and T.N. Wiesel. Functional architecture of macaque monkey visual cortex. *Proceedings of the Royal Society (B)*, 198:1–59, 1977.
- [13] E. A. Lachica, P. D. Beck, and V. A. Casagrande. Intrinsic connections of layer III of striate cortex in squirrel monkey and bush baby: Correlations with patterns of cytochrome oxidase. *J. Comparative Neurology*, 329:163–187, 1993.
- [14] R. Miller. Neural assemblies and laminar interactions in the cerebral cortex. *Biological Cybernetics*, 75:253, 1996.
- [15] W.A. Phillips and W. Singer. In search of common foundations for cortical computation. *Behavioral and Brain Sciences*, 20:657–722, 1998.
- [16] K.S. Rockland and J.S. Lund. Intrinsic laminar lattice connections in primate visual cortex. *Journal of Comparative Neurology*, 216:303–318, 1993.
- [17] E.L. Schwartz. Computational anatomy and functional architecture of striate cortex: a spatial mapping approach to perceptual coding. *Vision Research*, 20:645–669, 1980.
- [18] P. Sheridan and D.M. Alexander. Invariant transformations on a space-variant hexagonal grid. *Proceedings of Vision, Recognition, Action: Boston*, 1997.
- [19] N.V. Swindale. A model for the coordinated development of columnar systems in primate striate cortex. *Biological Cybernetics*, 66:217–230, 1992.
- [20] N.V. Swindale. The development of topography in the visual cortex: a review of models. *Network*, 7:161–247, 1996.
- [21] B.H. Tootell, M.S. Silverman, S.L. Hamilton, E. Switkes, and R. De Valois. Functional anatomy of the macaque striate cortex. v. spatial frequency. *Journal of Neuroscience*, 8(5):1610–1624, 1988.
- [22] B.H. Tootell, M.S. Silverman, S.L. Hamilton, R. De Valois, and E. Switkes. Functional anatomy of the macaque striate cortex III color. *Journal of Neuroscience*, 8(5):1569–1593, 1988.
- [23] B.H. Tootell, E. Switkes, M.S. Silverman, and S.L. Hamilton. Functional anatomy of the macaque striate cortex. II. retinotopic organization. *Journal of Neuroscience*, 8(5):1531–1568, 1988.
- [24] J.J. Wright. Simulation of EEG: Dynamic changes in synaptic efficacy, cerebral rhythms, and dissipative and generative activity in cortex. Submitted May 1998.

Environmental informatics - a new paradigm for coping with complexity in nature

David G. Green and Nicholas I. Klomp
School of Environmental and Information Science
Charles Sturt University

dgreen@life.csu.edu.au, <http://life.csu.edu.au/~dgreen/>

Abstract

The application of information technology to environmental issues is changing both theory and practice. The idea of “natural computation” provides new ways to understand environmental complexity across the entire range of scales, from individual phenotype to biogeography. Understanding the ways in which local interactions affect the global composition and dynamics of whole communities is crucial to the viability of strategies to manage ecosystems, especially in landscapes altered by human activity. Also environmental planning and management are increasingly dependent on accurate, up-to-date information that sets local decisions within a global context. The Internet makes it possible to combine environmental data from many different sources, raising the prospect of creating a global information warehouse that is distributed amongst many contributing sites.

1 Introduction

Humankind is in the midst of a crisis. For thousands of years people have exploited the environment as though it were an infinite resource — unchanging, predictable and inexhaustible. However the impacts of human activity are now felt everywhere. Conserving the world’s flora and fauna is one of the great challenges of our time. Loss of biodiversity, ecosystem degradation and pollution are just some of the environmental problems on planet Earth. With human population and industrialisation still increasing rapidly, it is becoming vital to place a check on these problems within the next few decades.

In the face of this rapidly changing situation, traditional ideas and approaches to environmental management are no longer enough. To manage (say) a national park adequately requires knowing much more than simply what is happening within the park. It demands that local issues be set in the context of the surrounding region, as well as national and international developments, global change, socioeconomic influences, and a host of other issues as well.

As planners and managers learn to cope with this new scenario, we are witnessing the development of a new paradigm that integrates traditional field ecology with modern technology. It is a paradigm that links scientific research to environmental planning and management. It links diverse and potentially massive sources of information, from field ecology to satellite imagery. Such a new approach can be invoked to address a host of practical problems, from land use planning to global warming.

In this brief account we try to achieve three goals. The first is to explain briefly the nature of complexity in the environment. Secondly, we argue that a new paradigm — environmental informatics — is emerging out of responses to the growing need to cope with this complexity. Finally we sketch out some of the “grand challenges”, both in research and in practice, that environmental informatics needs to address in the new millennium.

2 Complexity in the living world

2.1 Sources of environmental complexity

Even the simplest ecosystems are highly complex. Complexity in the environment is present for many reasons, but most many sources of complexity can be grouped into the categories described below.

- **The spatial scale**

Many influences of the world's environment come from sources outside the Earth's biosphere. The sun, the moon, meteors and geomagnetism all influence life on Earth. But even if considering only our biosphere, the sheer scale involved in global environmental management is immense. The planet's surface area totals over 509,000,000 square kilometres. Simply monitoring one factor (say, surface temperature) across such vast tracts is a huge task [27]. Thorough monitoring of all environmental factors or even rudimentary research of the entire surface of the Earth is currently impossible. Whilst modern technology can help (e.g. remote sensing), it generates huge volumes of data that must somehow be stored, collated and interpreted [19].

- **The temporal scale**

Many environmental processes occur over geological or evolutionary time. Even successional or micro-evolutional processes usually take place over time periods much longer than a human life (or the length of a typical ecological research project of 1-3 years!). This has led to many inaccurate, "time-blinkered" assumptions in ecology, such as stable community structures, climax states and other conclusions about balance and equilibria (see later).

- **The number of organisms**

Taxonomists have described about 1.5 million species [31]. The total number of species is not known, but is estimated to be somewhere between 10 million and 100 million. At the current pace it would take at least another 300 years of taxonomic research simply to document them all.

However it is not sheer numbers of species that make the living world complex, but rather the enormous variety of ways in which they combine and interact. For instance, suppose that 100 species inhabit a region; then there are 4,950 possible pairs of interacting species. However, when we look at possible combinations, the possibilities blow out to astronomical proportions. There are over 6.33×10^{19} ways in which we can select communities of 10 species at a time. For communities of 50 species at a time, this number rises to over 10^{93} combinations. This complexity increases further by orders of magnitude when the interactions of biotic and abiotic factors within an ecosystem are considered.

- **Criticality**

An important example of complexity, especially in landscapes, is the phase change between connected and fragmented population distributions [7]. For instance, if we remove small patches forest from a landscape then the forest as a whole retains its integrity. However if clearing continues (at random), then instead of small patches breaking off, the entire system remains connected until a critical point, whereupon it breaks down into many isolated fragments [8]. Such criticality or abrupt phase changes have now been documented in many natural systems [2], from pest and disease epidemics [10] to fire behaviour in forests [7].

- **Non-linear interactions and feedback loops**

One of the most important results to spring from ecological transect studies is that environmental factors alone do not fully explain the spatial distributions of organisms. For instance, competition between species often truncates distributions along an environmental gradient [25]. These results imply that ecosystems are not controlled in simple linear, fashion by external (*i.e.* abiotic) factors, but by interaction of biotic and abiotic factors within a system (eg. [2]).

Networks of interactions between species are a major source of complexity in ecosystems. Interactions between pairs of species can take many forms, such as predation and competition. Feedback loops are especially common in multi-species systems. In populations with seasonal reproduction, delays arising from feedback tend to produce cyclic behaviour. They can also lead to non-linear and chaotic dynamics [23].

Complexity in an ecosystem (as measured by species richness) does not necessarily imply stability [21, 22, 23]. One reason for this is that any random collection of interacting species is likely to contain at least one positive feedback loop, which destabilises the system and leads to local extinction of one or more species [28].

- **Human influence on natural systems**

Human influences on ecosystems tend to be disturbances that disrupt any semblance of equilibrium. Two of the most far-reaching of these disturbances have been land clearing and the introduction of exotic species. In many cases the effects are unintended side effects. Examples include wildfires, spread of diseases, pollution, salinisation, and desertification, to name just a few. In every case the disturbances force ecosystems away from equilibrium and can lead to local extinctions or other abrupt changes (e.g. [7, 8]).

In addition to making ecosystems more complicated (or more difficult to manage), humans influence the management of ecosystems by directing goals and agendas in ways that require environmental management decisions to be based on much more than ecological knowledge. This has led to the recent, rapid increase in the use of decision support models by land managers [19].

2.2 Some lessons of complexity

Although still in its infancy, complexity theory holds some important lessons for environmental science and management. Only some of these have been widely recognised so far. Taken together they highlight the need for new ways of doing research and management in ecology. Here I briefly summarise some of these lessons.

- **Local interactions can produce global effects**

Although reductionism has served science well, we must recognise that it fails badly in trying to make sense of environmental processes. Above we highlighted a few examples of the ways in which interactions between different populations can have unexpected effects. It has long been common practice in ecology to study individual populations separately (“autecology”), without reference to the ways they interact with other populations. Other reductionist practices include breaking down community level dynamics into studies of physiology and other responses at the level of individuals.

A good example is the way in which dispersal (interactions between sites in a landscape) can affect the dynamics of whole ecosystems [8]. Rare species tend to form clumped distributions which help them to persist in the face of superior competitors. Simulation studies suggest that this process provides one mechanism which maintains high diversity in tropical rainforests. Field studies confirm that rainforest contain just a few common, widespread species, and many rare species [14]. All of these rare species have clumped distributions. By not fully understanding such dynamics, conservation, management and research can be rendered ill-conceived or ineffective.

- **Systems can be inherently unpredictable**

Sensitivity to initial conditions is a well-known phenomenon in non-linear systems, and one of the hallmarks of chaos. It is especially common in ecology where so many interactions are non-linear [23]. As an example, consider what happens if the connectivity of a landscape is near the critical region mentioned earlier. Under such conditions the size and composition of connected patches becomes extremely variable so the outcome of processes that involve spread through a connected patch, such as fire, epidemics, and invasions, become inherently unpredictable [8, 10]. Likewise the addition of a single exotic species to an ecosystem alters the web of interactions between species, perhaps creating a potentially devastating positive feedback loop where none existed before [28].

The need to cope with unpredictability highlights the importance of tools such as simulation models. Simulation allows us to carry out virtual experiments. In environmental management

such experiments are often impossible to carry out in practice, either because they would take too long (e.g. forest succession) or because they would be too damaging (e.g. burning down an entire forest). Although exact prediction may be impossible simulation makes it possible to examine ways of dealing with many potential scenarios.

- **There is no balance of nature**

The idea that nature is in equilibrium – a cornerstone of much thinking within the environmental movement – arises from several sources. Perhaps the most important is the exceedingly long time scale of many processes in forest ecosystems, as mentioned earlier. Individual trees often live for many hundreds of years and simple succession – the replacement of one community by another – can take literally thousands of years to complete [6]. The fact that forests change so slowly gives the false impression that they are in equilibrium.

The equilibrium assumption underlies many ideas in theoretical ecology. For instance, MacArthur and Wilson’s theory of island biogeography suggested that for any island there is an equilibrium number of species that it can sustain [20]. However growing understanding of the large scale and long-term dynamics of ecosystems make equilibrium assumptions increasingly untenable. For example Clements’ theory assumed that succession leads to an equilibrium *climax* state [3]. This theory dominated plant ecology for most of the twentieth century, until evidence accumulated for other kinds of dynamics, such as chronically disturbed ecosystems [24] and long-term instabilities in vegetation history [5, 6].

Perhaps more importantly the urgent need to address environmental management, especially in disturbed ecosystems, is forcing ecologists to search for non-equilibrium models.

- **The Serendipity Effect**

Combining different datasets together often leads to unexpected discoveries. That is serendipity occurs. The probability of serendipity increases exponentially with the number of different datasets available. So large repositories of data are almost certain to be rich sources of new insights about environmental processes [9].

3 Natural computation

A new paradigm requires a new way of looking at the world. The increasing use of computers has stimulated a view in which the natural world is seen as a form of computation. The analogies are compelling. DNA is the code for life’s “program”. Organisms are akin to robots or agents, and animal communication is a form of information processing.

The links between biology (including ecology) and computing have been growing ever closer. Techniques such as genetic algorithms, cellular automata and neural networks clearly borrow on biological ideas. We have argued [10] that many algorithms can be improved by mimicking living systems more closely [18].

3.1 From genes to ecosystems

One of the major challenges for ecology is to bridge present gaps in our understanding in the spectrum of genotype, phenotype, population and community. Perhaps the least well understood is the link between genotype and phenotype, and thence to environmental processes. The obvious analogy for scholars of computing and complexity is that to understand how a computer program works it is not enough to understand what each line of code means. You also need to know how those lines of code are organised.

At present very little is known about the relationship between genetic composition and growth processes. Kauffmann [17] modelled genetic control over growth as a switching circuit in which genes

are ON-OFF switches that not only code for certain proteins but also affect other genes. However there has been very little other work of this kind.

L-system models [26] are now so sophisticated that they can faithfully reproduce the potential growth form of many plants. Virtual plants are now being used to carry out virtual experiments and could help to bridge the gap between laboratory experiments and field observations. A crucial step is to understand the link between growth form and taxonomic relationships. That is, how do genetic variations impact on the models?

3.2 Alife

One of the most relevant and important developments associated with natural computation is a new research field called artificial life ('Alife' for short). This is the study of life-like properties in computational systems.

One of the key ideas in Alife is that of an agent. An agent is a discrete entity that has certain computational capabilities, and can also interact both with its surroundings and with other agents. An important area of Alife research, and of advanced computing generally, is to study the properties and behaviour of multi-agent systems. This research is beginning to grow into a significant body of theory about systems of this kind.

For instance, in one early study, Hogeweg and Hesper (1983) showed that the observed social organisation of bumblebees arises as a natural consequence of the interaction between simple properties of bumblebee behaviour and their environment. For example, one rule they invoke is the TODO principle [12, 13]. Bumblebees have no intended plan of action, they simply do whatever there is to do at any given time and place. Similar interactions lead to order in many other animal communities, such as ant colonies and flock formation by birds.

4 Towards a new paradigm

For most of the Twentieth Century, conservation could be equated with national parks. However the rapidly growing scale of environmental alteration and increasing public awareness of environmental issues have highlighted the need for off-reserve conservation and environment management [4]. The range of off-reserve issues is now very broad. Some examples include: environmental impact assessment, state of the environment reporting, environmental monitoring, conservation of rare and endangered species, natural heritage planning, species relocation programs, land use planning, and environmental degradation.

Out of all the above activity has emerged an awareness that local decisions and priorities need to be set in a wider, and ultimately global context [30]. For instance to decide whether to log a patch of rainforest, you have to know how much other rainforest there is, what species will be put at risk, what the global costs and benefits are, etc. Conversely, every local area contributes data and experience that can be applied to other areas and can feed into setting global priorities and policies.

The new paradigm that is emerging treats environmental management as a host of activities all of which reinforce each other. Each area of activity is both enhanced and constrained by the global picture. The key to the success of the new approach is this two-way communication. Setting matters in context means having access to relevant and reliable information. During the 1990s governments have been very active in setting up regional, national, and international environmental information systems (e.g. [1, 9, 30]).

The growth of the Internet has played an integral part in this emerging paradigm. Up until recently most research was carried out as a series of isolated studies. However, by sharing data over the Internet, the results of previous studies can enrich subsequent research. The best examples are in genomic research, where the development of large, on-line databases means not only that new sequences can be interpreted by comparing them with whole families of existing data, but also that entirely new kinds of studies are possible in which researchers mine the databases for unsuspected

patterns and relationships. The challenge for ecology is to mobilise data from previous studies in similar fashion.

The essential advantage of the Internet (especially the World Wide Web) is its ability to combine information from many different sources in seamless fashion [9]. This has created an unprecedented opportunity for data sharing and cooperation on scales that were formerly deemed impossible. It also brings sharply into focus the need for coordination. The explosive growth of the Internet has led to massive confusion. Many organisations are duplicating facilities in inconsistent ways. There is an urgent need to develop for agreed protocols and standards regarding, data recording, quality assurance, custodianship, copyright, legal liability and indexing [9].

One of the most urgent needs is to develop a consistent framework for discussing environmental issues. One of the most basic problems is that we do not even have a comprehensive list of the world's species. Not only that, the taxonomic nomenclature has been confused and inconsistent. It is not surprising then that some of the first initiatives in on-line environmental information have focussed on putting together consistent reference lists of the world's species. For instance since 1993 the International Organization for Plant Information (IOPI) has been developing a checklist of the world's plant species [15]. This is now contributing to recent major initiatives in this area, including the Species 2000 Project [16] and the Global Biodiversity Information Facility (GBIF), which are international projects of the OECD's Megascience Forum [11]. The aim is to establish "... a common access system, Internet-based, for accessing the world's known species through some 180 global species databases ..."

A major challenge is to flesh out and complement the data that is now available with facilities that allow people to use it effectively. Along with data warehouses, we also need information systems to interpret and apply the information. For instance, foresters, faced with the need to demonstrate the environmental impact of logging operations, have developed simulation tools such as the visualisation program SmartForest. This program [29] integrates simulation models with geographic information to create views of future landscapes under selected scenarios.

5 Conclusion

Learning to conserve the world's living resources is one of the great challenges of our time. In a very real sense the future of humanity depends on finding a solution. It is not an easy problem to solve.

As we have seen here, achieving these goals will demand a much better understanding of environmental complexity than we have at present. Thus there is a need for greater dialogue between ecology and complexity studies. At present the extent of this dialogue is still small. With a few notable exceptions, most ecologists are largely unaware that the field of complexity even exists, and many researchers in (say) Alife are computer scientists who are unaware of the major issues and questions driving ecological research.

We can no longer pretend to manage nature in isolation from human activity. Human activity has expanded to affect virtually every ecosystem, everywhere. We have to learn to manage ecosystems that are not only out of equilibrium but also chronically disturbed and largely unpredictable. We can no longer confine conservation to "isolated", "natural" parks. Conservation needs to be incorporated into the ways we deal with all living systems in all environments.

Global conservation demands a much greater level of coordination than at present. This coordination includes two-way communication between the activities of different conservation agencies and groups. It also implies much greater planning because almost every socioeconomic activity potentially impinges on conservation. To achieve both of these ends, greater dialogue between ecologists and computer scientists is needed urgently.

References

- [1] Bolton, M.B. and Green, D.G. "Computers and Conservation — the Environmental Resources Information Network" 1991 *Trees and Natural Resources* **33**, 14–16.
- [2] Bull, C.M. and Possingham, H. "A model to explain ecological parapatry" 1995 *American Naturalist* **145(6)**, 935–947.
- [3] Clements, F.E. "Plant Succession" 1916 Carnegie Inst., Washington, Publ. No. 242.
- [4] Commonwealth of Australia "National Strategy for the Conservation of Australia's Biological Diversity" 1996 Australian Government Printing Service, Canberra.
- [5] Davis, M.B. "Pleistocene biogeography of temperate deciduous forests" 1976 *Geoscience and Man* **13**, 13–26.
- [6] Green, D. "Fire and stability in the postglacial forests of southwest Nova Scotia" 1982 *Journal of Biogeography* **9**, 29–40.
- [7] Green, D.G. "Simulated effects of fire, dispersal and spatial pattern on competition within vegetation mosaics" 1989 *Vegetatio*, **82**, 139–153.
- [8] Green, D.G. "Connectivity and complexity in ecological systems" 1994 *Pacific Conservation Biology*, **1(3)**, 194–200.
- [9] Green, D.G. "Databasing diversity — a distributed, public-domain approach", 1994 *Taxon*, **43**, 51–62.
- [10] Green, D. "Complexity in ecological systems" 1997 *Frontiers of Ecology — Making the Links* (ed. N. Klomp and I. Lunt) Elsevier, London, pp. 221–231.
- [11] Hardy, G. "OECD Megascience Forum Biodiversity Informatics Group" 1998 <http://www.oecd.org//ehs/icgb/BIODIV8.HTM>
- [12] Hogeweg, P. and Hesper, B. "Evolution as pattern processing: TODO as a substrate for evolution" 1991 In *From Animals to Animats* J. A. Meyer and S. W. Wilson (eds.) MIT Press, Boston, pp. 492–497.
- [13] Hogeweg, P. "As large as life and twice as natural: bioinformatics and the artificial life paradigm" 1993 *Complexity International*, Vol. 1. <http://www.csu.edu.au/ci/vol1/>
- [14] House, S.M. "Relationship between breeding and spatial pattern in some dioecious tropical rain-forest trees" 1985 Ph.D. Thesis, Australian National University.
- [15] International Organization for Plant Information (IOPI) "The Global Plant Checklist Project" 1993 <http://iopi.csu.edu.au/iopi/iopigpc1.html>
- [16] International Union of Biological Sciences (IUBS) "Species 2000" 1998 <http://www.sp2000.org/>
- [17] Kauffman, S. A. "The Origins of Order: Self-organization and Selection in Evolution" 1993 Oxford, Oxford University Press.
- [18] Kirley, M., Li, X. and Green, D.G. "Investigation of a cellular genetic algorithm that mimics landscape ecology" 1998 *SEAL '98*, to appear.
- [19] Klomp, N.I., Green, D.G. and Fry, G. "The role of technology in ecology" 1997 In *Frontiers of Ecology — Making the Links* N. Klomp and I. Lunt (eds.) Elsevier, London. pp. 299–309.

- [20] Macarthur, R.H. and Wilson, E.O. “The Theory of Island Biogeography” 1967 Princeton University Press, Princeton.
- [21] May, R.M. “Will a large complex system be stable?” 1972 *Nature*, **238**, 413–414.
- [22] May, R.M. “Stability and Complexity in Model Ecosystems” 1974 Princeton University Press, Princeton.
- [23] May, R.M. “Simple mathematical models with very complicated dynamics” 1976 *Nature*, **26**, 459–467.
- [24] Noble, I.R. and Slatyer, R. “Concepts and models of succession in vascular plant communities subject to recurrent fire” 1979 In *Fire and the Australian Biota* Gill, A.M., Groves, R.H and Noble, I.R. (eds.) Australian Academy of Science, Canberra. pp. 311–335.
- [25] Pielou, E.C. “Population and Community Ecology” 1974 Gordon and Breach, New York.
- [26] Prusinkiewicz, P., Hanan, J., Hammel, M. and Mech, R. “L-systems: from the theory to visual models of plants” 1997 In *Plants to Ecosystems - Advances in Computational Life Sciences* M.T.Michalewicz (ed.) CSIRO, Melbourne. pp. 1–27.
- [27] State of the Environment Advisory Council “Australia: State of the Environment” 1996 CSIRO, Melbourne.
- [28] Tregonning, K. and Roberts, A. “Complex systems which evolve towards homeostasis” 1979 *Nature*, **281**, 563–564.
- [29] UIUC Imaging Laboratory “SmartForest — an interactive forest visualizer” 1998 <http://www.imlab.uiuc.edu/smartforest/>
- [30] United Nations “Agenda 21: The United Nations programme of action from Rio” 1992 United Nations, New York.
- [31] Wilson, E.O. “The Diversity of Life” 1992 Penguin New York.

Complex behavior in perceptual line length

Takuo Henmi and Michael L. Kalish
 Department of Psychology
 University of Western Australia
 takuo@psy.uwa.edu.au
 kalish@boneyard.psy.uwa.edu.au

Abstract

There is ample evidence to show that nonlinear dynamical or chaotic properties underlie aspects of physiology, neurology, and even behavior. This paper presents a psychophysical “cascading” experiment in which the response is passed on to the next trial as the new stimulus. The time series of response is modeled by a nonlinear psychophysical model based on an existing recursive cubic polynomial function called the “ Γ recursion” originated by Robert Gregson. The responses in the cascading experiment are found to be classified into three categories, and some show the trace of chaos. However, the attempt to model the time series with the new model or the original Γ recursion resulted only in coarse approximations to the data. In spite of its inadequacy at simulating the time series itself, the new model managed to simulate the autocorrelation functions of the original data. These results suggest that the model we propose is in some sense within the same family of dynamical systems as the psychophysical dynamical system generating the observed data although it is necessary to develop more subtle nonlinear dynamical models.

1 Introduction

Both the biological mechanisms of life, and their behavioral manifestations, have been found in many of their aspects to show the properties of nonlinear dynamics, or deterministic chaos [12, 13, 20, 32, 37, 39]. Examples are found in analysis of biological systems with electroencephalogram [10, 23] and electrocardiogram [27], or also at neurological level [1, 2, 3, 4, 11, 31]. Psychophysics also provides examples of nonlinearity. The fundamental principle of psychophysics is to investigate the subjective intensity of a given stimulus, and to provide a mathematical description of the relationship between subjective intensity and stimulus magnitude in terms of a psychophysical law. For example, Stevens [40] proposed a “power law” which is defined as

$$\Psi = k\phi^\beta$$

where k and β are constants, ϕ denotes the stimulus magnitude, and Ψ denotes the subjective magnitude. That is, by assuming that changes in responses within the organism are directly proportional to changes in the level of external stimulation, the power law can plausibly produce the relationship of any stimulus magnitude and subjective intensity response. The power law is often referred to as a “linear” psychophysical law since the stimulus-response relationship can be represented with a straight line on log-log coordinates. The power law has been widely accepted as a good approximation of such relationships; however, it also has met with some skepticism. For example, Ross and Di Lollo [38] reported the constant failure of the power law in an experiment in which observers were asked to judge the magnitude of lifted weights. The failure of the power function lead some researchers to propose other expressions [29] while others (see chapter 6 of Uttal [43]) have proposed various “nonlinear” psychophysical functions. Indeed it is plausible to assume an essential nonlinearity in the stimulus-response relation if one assumes a nonlinearity in the biological substrate of perception.

Among the nonlinear psychophysical models that have been proposed, perhaps the leading example is the “ Γ recursion” function [14]. While other nonlinear functions have been applied to particular phenomena (e.g., Watson [46]), Gregson has shown that his Γ recursion, which is a complex valued cubic iterative function, can plausibly simulate a wide variety of psychophysical effects. Having

previously shown that the Γ recursion and its extended models are valid nonlinear psychophysical models at a practical simulation level [21], this research examines the ability of the one-dimensional model to simulate behavioral phenomena. In particular, we are interested in how a special property of the Γ , “cascading”, might be used to model dynamic aspects of line length perception, which, as reviewed below, is known to produce illusory effects under a number of conditions. The main characteristic of cascading is to transform the intensity of a response from one step of the simulation before applying it as a stimulus value at the next stage. This cascading process is vital for the model’s account of nonlinear psychophysical phenomena [14]. We suggest that the cascading process has a natural behavioral interpretation if we allow observers to produce their own stimuli. In particular, we asked observers to estimate the length of a briefly presented stimulus line over the course of a number of trials. In each trial but the first, in which a standard line was presented, the line presented to an observer was of the same length as they had indicated their subjective estimate of the line length to be on the preceding trial. If we assume nonlinear properties in human perception, then we should observe dynamics in this experiment like those predicted by a nonlinear cascaded model.

In this paper, we first describe the model to simulate the time series obtained by the experiment. The model is based on the cascaded Γ recursion [14], but since it is different in some respects, we refer to it as a cascaded cubic polynomial recursion (CCPR) model. We then briefly review the literature on perceived line length, and present the experiment. Apart from modeling with a nonlinear cascaded model, we also present a descriptive response time series analysis to see if the perceptual cascading process possesses nonlinear (in particular, chaotic) dynamics.

2 The Model

The Γ recursion is a well-established psychophysical function that Gregson has shown not only in a one-dimensional stimulus-response level [14], but also in a multi-dimensional vector form, $n\Gamma$ [15, 16, 17], and a lattice form, $(n \times n)\Gamma$ [18, 19]. A theoretical advantage of an iterative perceptual model such as Gregson’s is that it can reflect the iterative behavior of biological neural networks. As we mentioned above, we adopt a *cascaded cubic polynomial recursion* (CCPR) model which is based on Gregson’s one-dimensional Γ recursion. We therefore briefly review some properties of the Γ .

The Γ recursion is defined as

$$Y_{j+1} = -a(Y_j - 1)(Y_j + ie)(Y_j - ie) \quad i^2 = -1,$$

or equivalently

$$Y_{j+1} = -a(Y_j - 1)(Y_j^2 + e^2) \quad (1)$$

where a is real, ie is imaginary, and Y is complex. In order to apply this function as a psychophysical function, the input (stimulus) series U is scaled so that the parameter value a lies between 2 and 4. The imaginary component represents the internal activity level within the system. The parameter value e represents, roughly speaking, the sensitivity of the system to rates of change of inputs in time (not the well-known constant $e = 2.7182\dots$). It is constrained to the range between 0 and 0.5 ($0 < e < 0.5$) or $ae < 1.7$ in the region of $0.5 < e < 0.7$, to avoid an explosive condition. The initial value for this model, Y_0 , is fixed at the onset, namely, $Y_0 = (\zeta, \epsilon)$, $\zeta \leq 0.5$, $\epsilon < 10^{-8}$ (i.e. the initial condition $\approx 0.5 + 10^{-9}i$). After an arbitrary number of iterations (denoted as η), the real component of Y (denoted as $\text{Re}(Y)$ which must be in the region of $0 < \text{Re}(Y) < 1$ in order to prevent the explosion) represents the observable output or response magnitude. In short, the simulation is composed of;

$$U \Rightarrow a \Rightarrow \Gamma \Rightarrow \text{Re}(Y)_\eta.$$

The difference between Γ recursion and the cubic polynomial of the CCPR model is in their initial values. The initial value of the Γ recursion is $Y_0 = (\zeta, \epsilon)$ which gives the complex aspect to the model, but it is modified just to $Y_0 = (\zeta)$, i.e. $Y_0 = 0.5$) for the CCPR model. One can view the iterative

cubic polynomial adopted by the CCPR model as the special case of Γ recursion collapsed onto the real line.

In the “cascading” process, the output is transformed to the next stimulus, that is

$$a_k \Rightarrow (1) \Rightarrow Y_\eta \Rightarrow \text{Cascade} \Rightarrow a_{k+1} \Rightarrow (1) \Rightarrow \dots$$

Each point estimated with (1) corresponds to that of the response time series obtained by an actual experiment. More descriptively, a cascading model consists of two phases;

- estimate each element of the response time series (estimation phase), and
- transform the stimulus intensity and pass it back onto the estimation phase as the new stimulus intensity (cascading phase).

Mathematically, the CCPR model can be described as follows. The estimation phase is the same as (1), that is,

$$(Y_\eta)_k = -a_k((Y_{\eta-1})_k - 1)((Y_{\eta-1})_k^2 + e^2).$$

The cascading phase is composed of two parts. First, the next input is computed as

$$a_k = c(Y_\eta)_{k-1} + d,$$

where c and d are linear scaling constants, and η is the number of iterations. The second stage of the cascading phase concerns the initial value for the next estimation phase. For the first estimation phase, $(Y_0)_1 = 0.5$. The estimated value subsequently becomes the initial value of the following simulation, i.e., $(Y_0)_k = (Y_\eta)_{k-1}$ although $(Y_0)_k = \zeta + \text{Im}(Y_\eta)_{k-1} i$ in the original cascaded Γ recursion. If we make the time series obtained from an actual empirical (scaled) data set,

$$\mathbf{P} = \{p_1, p_2, p_3, \dots, p_n\},$$

then the CCPR model estimates each point and produces the corresponding time series,

$$\hat{p}_k = (Y_\eta)_k,$$

$$\hat{\mathbf{P}} = \{\hat{p}_1, \hat{p}_2, \hat{p}_3, \dots, \hat{p}_n\}.$$

The parameter values a_1 , e , c , and d are free, and can be set at the beginning to minimize the sum of the squared deviation,

$$\min(\sum (\mathbf{P} - \hat{\mathbf{P}})^2).$$

3 Background of Perceived Line Length

As was mentioned above, line length is often perceived distortedly, as in geometric illusions such as the Müller-Lyer illusion, the Ponzo illusion, the horizontal-vertical illusion, and the parallel-line illusion [5, 22, 24, 25, 33, 34, 35, 45], four of which are shown in Figure 1. Depending on context, then, the human visual process distorts the equal length of the horizontal parallel or the horizontal-vertical lines, so that they appear to be of unequal length. However the current research concerns distortion that happens in the course of successive comparisons of line length, rather than the perceptual distortion in any simultaneous comparison.

Woodworth [48] referred the error caused by successive comparison as “time error”, and noted that such error was usually negative—a negative time error happens when an observer underestimates the stimulus magnitude of an object (such as weight, frequency, etc). Woodworth’s review on successive comparison and time error referred to weight lifting, auditory, and some esthetic experiments, but not to line length estimation. We were unable to trace the first systematic experiment investigating this particular subject, but as early as 1957, Yokose, *et al.* [49] showed that the perceived vertical length




Figure 1: Examples of line length illusions. (a) Müller-Lyer illusion, (b) Ponzio illusion, (c) horizontal-vertical illusion, and (d) the parallel-lines illusion. The horizontal lines of (a) and (b), but the top lines appear longer in both cases. In (c), the vertical line seems longer than the horizontal line, but they are also the same length. The lines in (d) are different length. However, the presence of the other line has a normalizing effect. The top line appears shorter and the bottom line appears longer than it is.

of a line, in successive comparisons, depended on the exposure time of the test line. According to their results, the line was perceived to be shorter when the test line was presented for shorter durations (the minimum exposure duration for their experiment was 50 ms). Moreover the subjective length of the test line was shorter than its objective length. Erlebacher and Sekuler [9] conducted an analogous experiment, and their results replicated Yokose *et al.*'s finding; subjective length within the stimulus exposure duration was shorter than the objective length.

Tsal and Shalev [42] and Prinzmetal and Wilson [36] studied the effect of attention on subjective line length; they also mentioned the phenomenon just described above. That is, in a successive line length comparison experiment, the comparison lines were perceived to be shorter than the physical length of the standard line on average. Prinzmetal and Wilson [36] hypothesized that the cause of the underestimation could include a framing effect from the display monitor [28, 30]. Alternatively, the bias may have been due to one of their methods to control subjects' attention in their experiment. Although in a modified version of their experiment, Prinzmetal and Wilson were able to eliminate underestimation, they nonetheless noted that observers do show an overall tendency towards negative time error.

Previous research has also suggested the tendency to underestimate line lengths in successive comparisons (cited in Brigell and Uhlarik [6]), agreeing with Woodworth [48]. However, the fact of negative time error in estimating a line length raises some questions. For example, what will happen to the underestimation effect for stimuli near the lower limen of perception? If the stimulus is determined from the previous response, will stimuli converge to the lower limen, or will responses show some kind of more irregular behavior? We have hypothesized at first that the response time series in a cascading experiment would show a nonlinear behavior, however as we reviewed the previous literature on successive comparison, it should be adequate to modify the hypothesis. That is, if the human perceptual process possesses nonlinear dynamical properties, then:

- in accord with previous reports, the response would decrease (the comparison standard line would shrink as the trials proceed) up to the point when the line is near the margin of perceptible change and,
- the response time series after that point would show nonlinear or chaotic dynamics.

4 Experiment

4.1 Participants

Twenty participants (four males and 16 females) were recruited; nine psychology undergraduate students seeking to fulfil a course requirement, and eight psychology PhD students at the University of Western Australia together with three people of the experimenter's personal acquaintance. The participants' ages varied from 17 years to 38 years old. All had normal or correct to normal vision. Age and sex were not expected significantly to affect line length estimations (see Verrillo [44]), so the data from all observers was treated equally.

4.2 Apparatus

The apparatus used in this experiment is a PC with a 14-inch color display, and a two-button mouse device. The program is written in C++ in MS-DOS.

4.3 Procedure

Participants sat in front of a PC display monitor in a semi-darkened room. The chair was adjusted so that the observer's eyes were level with the center of the screen. The distance between the observer and the display was approximately 50 cm. The screen resolution was 640 pixels wide by 480 pixels high. The stimulus consisted of red vertical lines on a gray background. The lines were 3 pixels wide and could be adjusted from 0 to 440 pixels high. The stimulus line was presented 243 pixels from the left hand side of the screen for 750 ms. This was followed by a 350 ms inter-stimulus interval, during which time the screen was blank. Then the control line was presented 243 pixels from the right hand side of the screen (about 10.3 cm). The distance between the stimulus and control lines was 148 pixels (about 6.3 cm). After the observer made their judgement about the length of the line (described in detail below) there was a 700 ms delay before the next stimulus line was presented.

The origin of the both the standard and the control lines was fixed at the bottom of the display screen. Observers were asked to use a mouse to adjust the top of the control line to match the stimulus line in length, and then to click the left hand mouse button to record their response. The length of the standard line for the first trial was set so that the line extended from the bottom of the working screen to its center (220 pixels in length), while on subsequent trials, the standard line was set to the length of the control line from the previous trial, again anchored at the bottom of the screen.

The length of the control line, when it appeared on the screen, was set to be the length of the stimulus line (in pixels) plus or minus a random number of pixels between 0 and 40. Irrespective of the size of the control line determined from previous trials, line length was constrained between 0 and 440 pixels. Prior to the experiment, the experimenter advised the observers to be as precise as possible in their judgements. The experimenter also advised each observer to notify the experimenter if they accidentally clicked the mouse during the course of experiment in order to correct the mistakes immediately. Observers were naïve to the method used for determining the stimulus line length on each trial.

5 Results

The outcome of the experiment indicated that observers' responses tended to be of three distinct forms, described below, although, of course, no individual gave exactly the same responses as any other. The three categories of response series, of which an illustration over the course of 1000 trials is shown in Figure 2, were:

case1 The line length decreases, converging to the lower limen of a single pixel at the bottom of the screen (Figure 2(a) and 2(b)),




Figure 2: Changes in adjusted line length (in pixels) for 1000 trials in the cascading experiment. Observers’ data can be grouped into three distinct categories. In (a) and (b), the line converged to the lower limen. For (c) and (d), line length decreased almost uniformly, before oscillating within a small range around a given value that differed for each observer. Observers in the third category produced data that did not appear to converge to any particular value. Examples of the third category are shown in (e) and (f).

case2 The line length decreases up to certain point, but then oscillates irregularly within a narrow range (Figure 2(c) and 2(d)), and

case3 Neither of the above—the dynamics differ from **case1** and **case2** (Figures 2(e) and 2(f)).

The numbers of observers whose data could be classified into each case are 6, 10, and 4 respectively. The overall average of the positive and negative error (overestimation and underestimation of the stimulus magnitude) can easily be deduced by looking at the asymptotic response; that is, if the asymptotic response is below the starting point, then the overall average is negative and *mutatis mutandis* for responses above the starting point. Of the 20 observers we tested, only three showed a tendency towards positive errors, with the dynamics of their results obviously being of the form of **case3**. This result agrees with previous reports that a line indeed is perceived to be shorter than its physical length on average.

As for the dynamics of the time series, it is obvious that those in **case1** would be unlikely to show the trace of chaotic dynamics as they converged to the lower limen. Since the dynamics of those in **case2** consist of two phases we chose to pay particular attention to the second, oscillatory phase which we thought might be well modeled by a chaotic dynamical system. In regard to applying a chaos detecting algorithm to the time series, we have to be careful as Kantz and Schreiber [26] have drawn attention to the need to distinguish between stochastic noise and chaos. In order to determine whether the time series data that we collected was chaotic or stochastic noise, we adopted a method described by Khadra *et al.* [27]. Their algorithm is applicable to relatively short time series such as one in this experiment, as opposed to those methods that require a larger time series, such as those that need to compute the Lyapunov exponent. The algorithm for relatively short time series was introduced by Sugihara and May [41], but Ellner (1991, cited in Casdagli [7]) indicated their results were faulty. Khadra *et al.*’s algorithm [27], however, overcomes the problems that lie not only in the Sugihara and May’s [41] algorithm, but in other algorithms as well, making it more reliable. The basis of their algorithm is to state a null hypothesis that the given time series is not chaotic, and to derive the probability of obtaining some test statistic, calculated from the test data, under the null hypothesis. If this probability p is less than some predetermined α , then we reject the null hypothesis and assert that deterministic chaos is detected. Here, we will only summarize their algorithm, but the interested reader should review the technical and theoretical background of the algorithm in Khadra *et al.* [27]. The reader who is familiar with Efron’s bootstrap [8] will see that the algorithm developed

by Khadra *et al.* is a close relative. Khadra *et al.*'s algorithm is as follows:

1. Compute the median absolute forecasting error (MAE, denoted as Q_D) of the original time series with the forecasting algorithm with fixed embedding dimension d for fixed k nearest neighboring states described in Khadra *et al.* [27]. The delay factor (usually denoted as τ) is set to the minimal value (i.e. $\tau = 1$) in order to use all the data points. Since the method to choose optimal d and k was not specified in Khadra *et al.* [27], we adapted the parameters, over a small range ($2 \leq d \leq 7$, $3 \leq k \leq 7$), that minimize the MAE.
2. Create 128 surrogate data sets as described in Khadra *et al.* [27].
3. Compute the MAE for each surrogate data set.
4. Compute the mean (μ_s) and standard deviation (σ_s) of the 128 MAEs.
5. Compute χ , where $\chi = |Q_D - \mu_s|/\sigma_s$.
6. Compute the probability $p = \text{erfc}(\chi/\sqrt{2})$.

We applied the algorithm to those data in **case2** (last 800 points) and **case3** (1000 points), and disregarded **case1** data. Moreover, in order to show the validity of the algorithm, we applied it to three randomly generated time series that are given by,

- 1) Gaussian white noise (W-N),
- 2) a random walk where $Y_{k+1} = 0.99 Y_k + \text{Gaussian white noise}$ where $Y_0 = 0.5$ (R-W1), and
- 3) a random walk where $Y_{k+1} = Y_k + \text{random integer } l \text{ in } [-20, 20]$ where $Y_0 = 220$ (R-W2),

as well as a dummy time series generated by the CCPR model. The results of the analysis are summarised in Table 1. As can be seen from the table, evidence of chaotic dynamics was found in 8 instances out of the 14 data sets analyzed (two of such chaotic time series are shown in Figure 2(c) and (e)) whereas all the randomly generated data gave the converse results. The dummy series generated by the CCPR model also showed a sign of chaos that suggests its potential to model the data.

In order to test the CCPR model, we first fit it quantitatively to the data. We utilized the simulated annealing method to optimize the parameters a_1 , e , c , d of the CCPR model in terms of minimizing the sum of squared deviation between actual and predicted data values. Figure 3 shows the dynamics of a **case2** observer (Figure 2(c)) with the entire 800 points (a), first 400 points (b), first 200 points (c), and first 100 points (d) (solid line) along with the dynamics estimated by the CCPR model (dashed line). The response magnitude of each figures was scaled to $[0, 1]$ as the actual minimum value (in pixels) of the time series to be 0 and the maximum to be 1. As one can easily see, the simulated annealing method tended to determine the parameters so that the CCPR model would estimate close to the moving mean value throughout the series. The failure to fit the data with the CCPR model lead us to try simulating the data with the cascaded Γ recursion, that is, $Y_0 = 0.5 + 10^{-9} i$, $(Y_0)_k = 0.5 + \text{Im}(Y_\eta)_{k-1} i$. The results are shown in Figure 3 with dotted line. In this case too, the model could not simulate the data.

The next attempt is to minimize the sum of squared deviation of the autocorrelation function of the data and that of the time series generated by the CCPR model. The sample data taken here are two from **case2** with first 200 points dropped and two from **case3**—those four time series shown in Figure 2 (Figure 4–7(a)). The autocorrelation was computed at lags up to one quarter of the entire time series (i.e. 200 points for those in **case2**, and 250 points for **case3**) as suggested in Williams[47]. The autocorrelation function of the real data are shown in Figure 4–7(b). Those four figures roughly represent the range of autocorrelation functions of all the series. That is, for **case2**, the autocorrelation functions either decreased uniformly and rather unsteadily, or went down to negative then came back up to positive again, and for all **case3**, the functions decreased uniformly and rather smoothly. No distinction was detected in the autocorrelation function between chaotic and stochastic data.

		Cascading line length data ($d = 5, k = 7$)					
Observer	case	Q_D	μ_s	σ_s	χ	p value	Dynamics
1	2	0.643	0.401	0.104	2.327	0.020*	Chaotic
2	2	0.936	0.355	0.171	3.390	0.001*	Chaotic
3	2	0.191	0.437	0.171	1.441	0.150	Not chaotic
4	2	0.190	0.273	0.141	0.585	0.559	Not chaotic
5	2	4.326	0.325	0.099	40.573	6.87×10^{-360} *	Chaotic
6	2	0.110	0.402	0.135	2.165	0.030*	Chaotic
7	2	0.331	0.478	0.136	1.079	0.281	Not chaotic
8	2	0.501	0.323	0.073	2.423	0.015*	Chaotic
9	2	4.102	0.411	0.093	39.790	3.21×10^{-346} *	Chaotic
10	2	0.411	0.364	0.130	0.362	0.718	Not chaotic
11	3	4.336	0.371	0.152	26.152	9.46×10^{-151} *	Chaotic
12	3	0.744	0.390	0.187	1.990	0.047*	Chaotic
13	3	0.451	0.329	0.189	0.645	0.519	Not chaotic
14	3	0.419	0.511	0.321	0.285	0.776	Not chaotic
W-N		0.999	0.937	0.034	1.808	0.071	Not chaotic
R-W1		0.982	0.939	0.038	1.159	0.246	Not chaotic
R-W2		0.475	0.335	0.175	0.805	0.421	Not chaotic
CCPR		1.063	0.936	0.039	3.261	0.001*	Chaotic

* $p < 0.05$

Table 1: Chaos Detection Parameters for Line Length Experiment

In this case also, the parameters a_1 , e , c , d of the CCPR model were estimated by the simulated annealing method. The time series generated by the CCPR model are shown in Figure 4–7(c), and its autocorrelation function along with that of the real data are shown in Figure 4–7(d). Since the parameters were estimated to minimize the error in their autocorrelation function, the time series generated by the CCPR model do not necessarily bear resemblance to the original data. However, the CCPR model could produce close approximation of the observed spectrum at least in its overall shape.

6 Discussion

The primary interest of this study was to investigate how well we could model the dynamics of observers making iterated judgments of line length; first, with a descriptive time series analysis, and then with a cascaded nonlinear function based on Gregson’s Γ recursion. The descriptive analysis is essentially an attempt to detect chaos, if it exists, within observers’ responses. The use of CCPR on the other hand was an attempt to develop a model of observers’ responses that was capable of generating the same response train as the observers produced.

Our results from the descriptive analysis indicate that the observers, with the proportion of 8:6, show deterministic chaos in their successive judgements of line length. Furthermore, it is possible that the perceptions of some of those observers who made stabilized responses (**case1**) are actually chaotic, but that this is masked by the fact that a lower bound is imposed on the adjustment of the line by the physical limit of the screen. Nonetheless, the fact that it is some, and not all, of the observers who show evidence of chaotic behavior, suggests that there are considerable between-subject differences in the form of responses. That is, for this particular judgement task, individual behavior varies from chaotic to stochastic. Intuitively, the long-term behavior of other judgement tasks involving successive comparisons, such as weight lifting or judging the pitch of a tone, would show similar results.

Our attempts to detect chaos in observers’ responses were relatively successful. However, the presence or absence of chaos is not the most crucial question here—what one would really like to

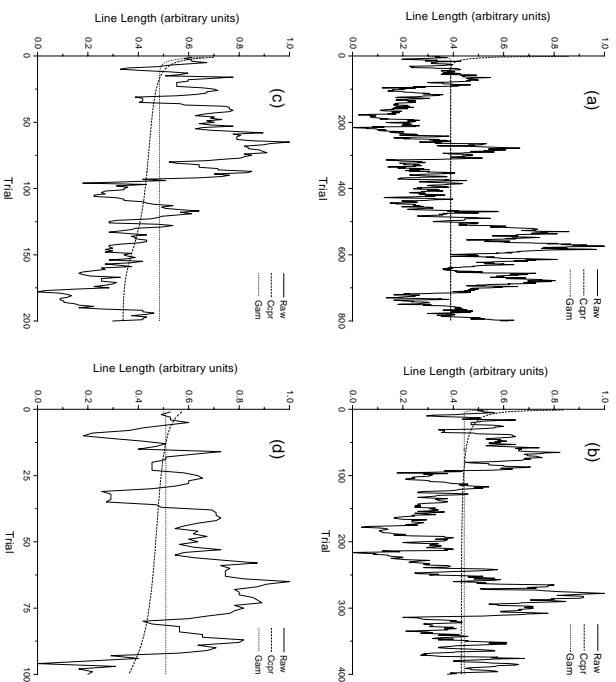


Figure 3: Example of an actual data set (solid line) with the corresponding optimised CCPR model parameters (dashed line) and the cascaded Γ recursion (dotted line). The graphs show the entire data set (a), the first 400 points (b), the first 200 points (c), and the first 100 points (d) for the observer shown in Figure 2(c). The line length is in arbitrary units.

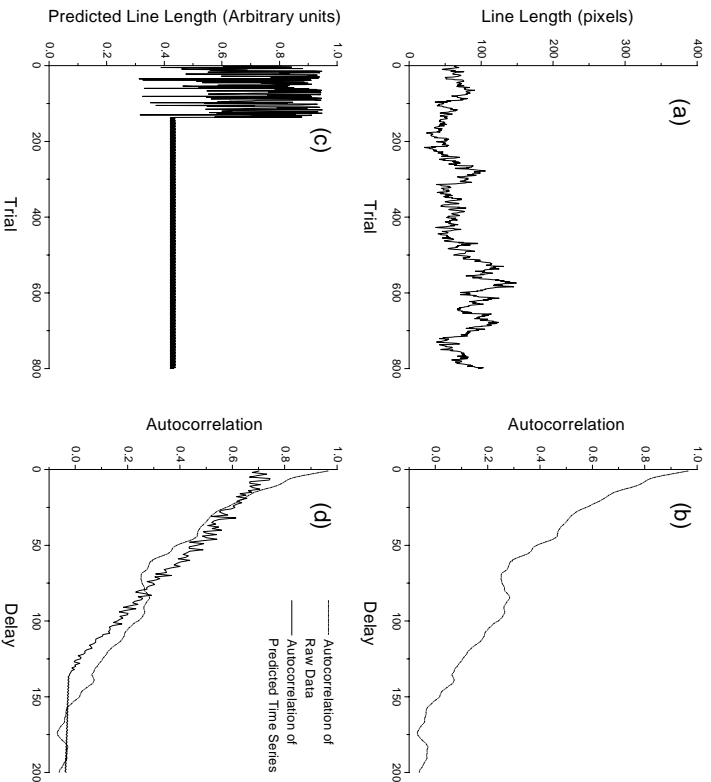


Figure 4: Comparison of autocorrelation functions for the data from the line length experiment and the CCPR model. (a) shows the actual time series of the observer shown in figure 2(c), and (b) shows its autocorrelation function. (c) shows the time series predicted by the CCPR model to minimize the error between the two autocorrelation functions. (d) shows the autocorrelation function of (c) along with the comparison (b).

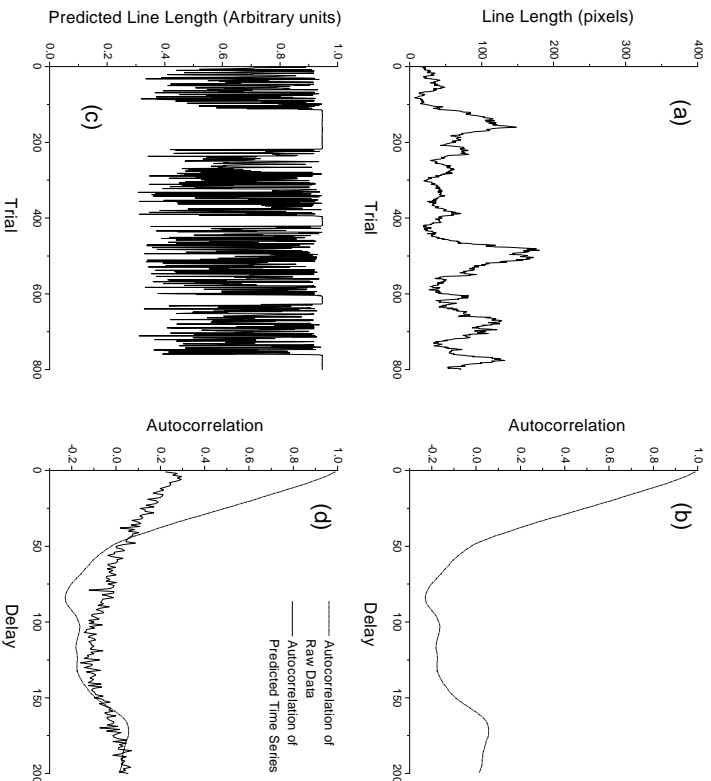


Figure 5: Comparison of autocorrelation functions for the data from the line length experiment and the CCPR model. (a) shows the actual time series of the observer shown in figure 2(d), and (b) shows its autocorrelation function. (c) shows the time series predicted by the CCPR model to minimize the error between the two autocorrelation functions. (d) shows the autocorrelation function of (c) along with the comparison (b).

know is the nature of the process underlying the phenomenon observed. It seems intuitively plausible that the process of iterated perception ensues from a simple underlying nonlinear dynamic, such as Gregson’s Γ recursion, which has proved useful in other instances of psychophysical modeling. But, demonstrating this rigorously, by matching the empirical data with data obtained from nonlinear dynamical equations, is a nontrivial matter.

We have made some initial efforts in this direction, using simulated annealing as a method of estimating optimal parameter values for the real Γ recursion and then for the CCPR. The real gamma failed utterly in this experiment, providing, as the optimal fit, a trajectory converging to a fixed point. The CCPR fared somewhat better, suggesting that CCPR is a more plausible model than the real gamma recursion; however the results here were not entirely satisfactory either. What the CCPR produced, with optimal parameters, was a smooth curve approximation, not unlike a moving average of the original data.

One could reasonably conclude from these results that the CCPR and cascaded Γ are unsuitable for modeling the kind of chaotic behavior that we found in some observers’ responses, irrespective of what parameter values one chooses. On the other hand, it turns out that the CCPR model did a fair job of simulating the autocorrelation function of the original data although the time series generated by the CCPR to fit the autocorrelation function do not show any particular resemblance to the original data. This result is not surprising since our attempt here is to fit the autocorrelation function as a whole, and not to show the overall trend of original time series. This could suggest that the CCPR is in some sense within the same family of dynamical systems as the psychophysical dynamical system generating the observed data. One could construct a somewhat plausible argument that the psychophysical phenomenon observed consists of the CCPR plus noise. However, we believe that this is not the case, and that the data reveals a deterministic chaos which is not revealed by the Γ and the CCPR models that have been the subjects of our computational experiments.

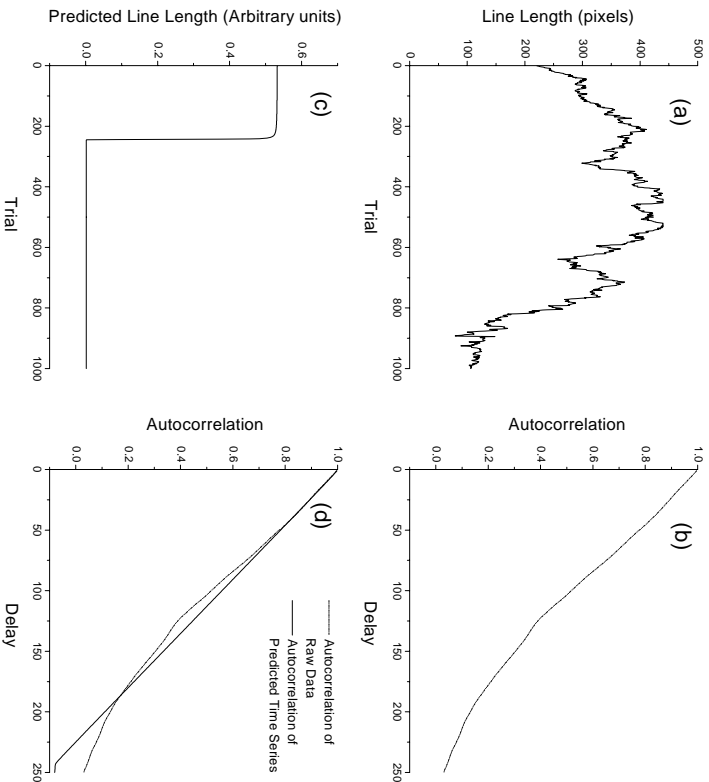


Figure 6: Comparison of autocorrelation functions for the data from the line length experiment and the CCPR model. (a) shows the actual time series of the observer shown in figure 2(e), and (b) shows its autocorrelation function. (c) shows the time series predicted by the CCPR model to minimize the error between the two autocorrelation functions. (d) shows the autocorrelation function of (c) along with the comparison (b).

The Γ model was introduced into psychophysics in order to provide a deterministic explanation for psychophysical phenomena which were previously written off as “just noise”. In this case, however, we have an apparently chaotic psychophysical phenomenon which neither the Γ nor the CCPR predicts in any simple way. The only clue we have is that the CCPR does better than Γ —it at least gives the trend of the data. Further study will be required to determine the form of the nonlinear dynamic that predicts not only the trend but the chaotic nature of the iterated perception phenomenon. Just as the classical psychophysical experiments led to the invention of Γ , this experiment must lead to the development of more subtle nonlinear dynamical models.

In addition to the primary results described above, our experiments also provided some secondary results, of particular note being the same tendency towards negative error that were found in earlier judgement tasks reviewed in Woodworth [48], and in the line length judgement experiments reviewed in the third section of this paper. Moreover, our experiments revealed three types of long term trends in the judgements of line length. Specifically, while some observers continued underestimating the length to the lower limit, the majority behaved irregularly when the line was very short; some even showed such irregularity throughout the experiment. In this experiment, the line was anchored at the bottom of the display with the line length being adjusted only from the top. If, however, one end of the line had been anchored at the center of the display with the line capable of being adjusted either upwards or downwards, then it is plausible to assume that the majority of **case1** observers would still have shown irregularity in their responses when the line was very short.

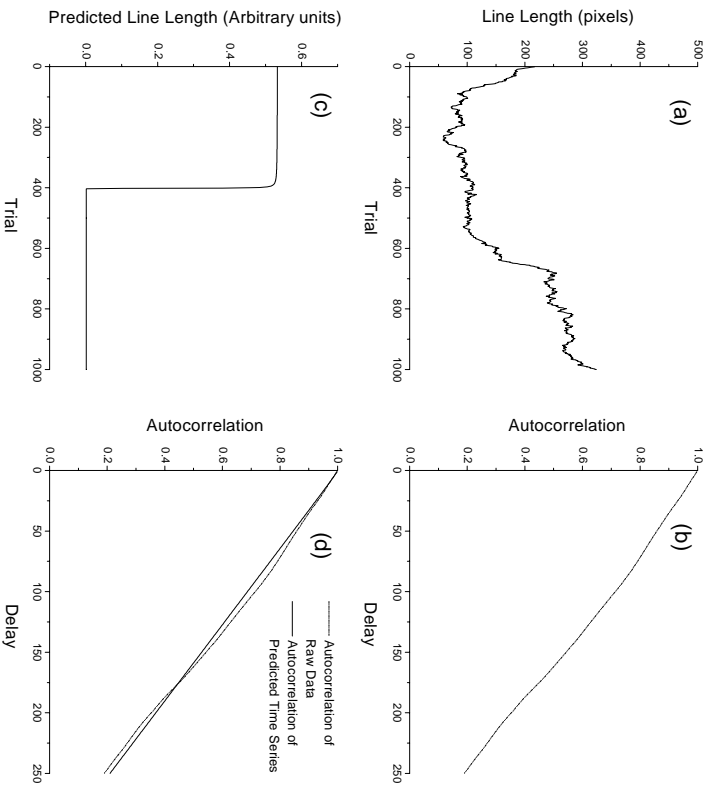


Figure 7: Comparison of autocorrelation functions for the data from the line length experiment and the CCPR model. (a) shows the actual time series of the observer shown in figure 2(f), and (b) shows its autocorrelation function. (c) shows the time series predicted by the CCPR model to minimize the error between the two autocorrelation functions. (d) shows the autocorrelation function of (c) along with the comparison (b).

7 Acknowledgements

We would like to thank Dr. Ben Goertzel at IntelliGenesis Corporation, Associate Professor Les Jennings at the Math Department, Mr. Mark Diamond and Mr. Jason Forte at the Psychology Department in the University of Western Australia for useful discussions.

References

- [1] Adachi, M. and Aihara, K. “Associative dynamics in a chaotic neural network”, 1996 *Neural Networks*, **10**, No. 1, 83–98.
- [2] Aihara, K. and Matsumoto, G. “Chaotic oscillations and bifurcations in squid giant axons”, 1986 *Chaos*, edited by Holden, 257–269.
- [3] Aihara, K. “Kaosu nyu-raru nettowa-ku” [Chaos neural network], 1990 *Kaosu: Kaosu Rinon no Kiso to Ouyou* [Chaos: Basis and Application of Chaos Theory] (Japanese), edited by Aihara, 289–317.
- [4] Aihara, K. “Chaos in neural response and dynamical neural network models: Toward a new generation of analog of computing”, 1994 *Towards the Harnessing of Chaos*, edited by Yamaguchi, 83–98.
- [5] Brigell, M., Uhlarik, J., and Goldhorn, P. “Contextual influences on judgements of linear extent”, 1977 *Journal of Experimental Psychology: Human Perception and Performance*, **3**, No. 1, 105–118.

- [6] Brigell, M. and Uhlarik, J. “The relational determination of length illusions and length aftereffects”, 1979 *Perception*, **8**, 187–197.
- [7] Casdagli, M. “Chaos and deterministic versus stochastic non-linear modelling”, 1991 *Journal of the Royal Statistical Society B*, **54**, No. 2, 303–328.
- [8] Efron, B. “Bootstrap methods: Another look at the jack-knife”, 1979 *Annals of Statistics*, **7**, 1–6.
- [9] Erlebacher, A. and Sekuler, R. “Perceived length depends on exposure duration: Straight lines and M-L stimuli”, 1974 *Journal of Experimental Psychology*, **103**, No. 4, 724–728.
- [10] Freeman, W. J. “The physiology of perception”, 1979 *Scientific American*, **Feb**, 34–41.
- [11] Freeman, W. J. “Neural networks and chaos”, 1994 *Journal of Theoretical Biology*, **171**, 13–18.
- [12] Glass, L. and Mackey, M. C., 1988 *From Clocks to Chaos: The Rhythms of Life*.
- [13] Glass, L. and Kaplan, D. “Time series analysis of complex dynamics in physiology and medicine”, 1993 *Medical Progress through Technology*, **19**, 115–128.
- [14] Gregson, R. A. M., 1988 *Nonlinear Psychophysics*.
- [15] Gregson, R. A. M. “Nonlinear psychophysics and Fechner’s Paradox”, 1989 *Mathematical and Theoretical Systems*, edited by Keats, 208–218.
- [16] Gregson, R. A. M. and Britton, L. A. “The size-weight illusion in 2-D nonlinear psychophysics”, 1990 *Perception and Psychophysics*, **48**, No. 4, 343–356.
- [17] Gregson, R. A. M. “Similarities derived from 3-d nonlinear psychophysics: Variance distributions”, 1994 *Psychometrika*, **59**, No. 1, 97–110.
- [18] Gregson, R. A. M., 1995 *Cascades and Fields in Perceptual Psychophysics*.
- [19] Gregson, R. A. M., submitted *Narrow parameter windows and analogues of contextual noise in nonlinear psychophysics*.
- [20] Halasz, M. F. “Nonlinear dynamics in behavioral systems”, 1994 *American Psychologist*, **50**, No. 2, 107–108.
- [21] Henmi, T. “Technical report submitted to The University of Western Australia, Perth Western Australia”, 1996.
- [22] Jaeger, T. “Contextual effects in the parallel lines illusion: Some implications for assimilation theory”, 1994 *Perceptual and Motor Skills*, **61**, 1263–1273.
- [23] Jeong, J., Joung, M. K., and Kim, S. Y. “Qualification of emotion by nonlinear analysis of the chaotic dynamics of electroencephalograms during perception of 1/f music”, 1994 *Biological Cybernetics*, **78**, No. 3, 217–225.
- [24] Jordan, K. and Schiano, D. J. “Serial processing and the parallel-lines illusion: Length contrast through relative spatial separation of contours”, 1986 *Perception and Psychophysics*, **40**, No. 6, 384–390.
- [25] Jordan, K. and English, P. W. “Simultaneous sampling and length contrast”, 1989 *Perception and Psychophysics*, **46**, No. 6, 546–554.
- [26] Kantz, H. and Schreiber, T., 1997 *Nonlinear Time Series Analysis*.

- [27] Khadra, L. M. Maayah, T. J., and Dickhaus, H. “Detecting chaos in HRV signals in human cardiac transplant recipients”, 1997 *Computers and Biomedical Research*, **30**, 188–199.
- [28] Künnapas, T. M. “Influence of frame size on apparent length of a line”, 1955 *Journal of Experimental Psychology*, **30**, No. 3, 168–170.
- [29] Krueger, L. E. “Reconciling Fechner and Stevens: Toward a unified psychophysical law”, 1989 *Behavioral and Brain Sciences*, **12**, 251–320.
- [30] Luo, C. R. and Wang, S. “Effects of figure context on the apparent length of a line”, 1997 *Perceptual and Motor Skills*, **85**, 551–558.
- [31] Mees, A., Aihara, K., Adachi, M., Judd, K., Ikeguchi, T., and Matsumoto, G. “Deterministic prediction and chaos in squid axon response”, 1991 Santa Fe, New Mexico: Santa Fe Institute, reference number 91-12-049.
- [32] Momose, K., Koyama K., and Uchiyama, A. “Nonlinear analysis of visual evoked potentials elicited by color stimulation”, 1997 *Methods of Information in Medicine*, **36**, 315–318.
- [33] Pietarinen, J. and Virsu, V. “Geometric illusion: II. Features of the method of magnitude estimation of length differences”, 1967 *Scandinavian Journal of Psychology*, **8**, 172–176.
- [34] Pollock, W. T. and Chapanis, A. “The apparent length of a line as a function of its inclination”, 1952 *Quarterly Journal of Experimental Psychology*, **4**, 170–178.
- [35] Prinzmetal, W. and Gettleman, L. “Vertical-horizontal illusion: One eye is better than two”, 1993 *Perception and Psychophysics*, **53**, No. 3, 81–88.
- [36] Prinzmetal, W. and Willson, A. “The effect of attention on phenomenal length”, 1997 *Perception*, **26**, 193–205.
- [37] Richards, W., Wilson, H. R., and Sommer, M. A. “Chaos in percepts?”, 1994 *Biological Cybernetics*, **70**, No. 4, 345–349.
- [38] Ross, J. and Di Lollo, V. “A constant failure of the power law for lifted weight”, 1970 *Perception and Psychophysics*, **8**, 289-290.
- [39] Smithson, M. “Judgement under chaos”, 1997 *Organizational Behavior and Human Decision Processes*, **69**, No. 1, 59–66.
- [40] Stevens, S. S., 1986 *Psychophysics: Introduction to Its Perceptual, Neural, and Social Prospects*, New introduction by Marks.
- [41] Sugihara, G. and May, R. “Nonlinear forecasting as a way of distinguishing chaos from measurement error in time series”, 1990 *Nature*, **344**, No. 19, 734–741.
- [42] Tsal, T. and Shalev, L. “Inattention magnifies perceived length: The attentional receptive field hypothesis”, 1996 *Journal of Experimental Psychology: Human Perception and Performance*, **22**, No. 1, 233–243.
- [43] Uttal, W. R., 1973 *The Psychobiology of Sensory Coding*.
- [44] Verrillo, R. T. “Stability of line-length estimates using the method of absolute magnitude estimation”, 1983 *Perception and Psychophysics*, **33**, No. 3, 261–265.
- [45] Virsu, V. “Geometric illusions: I. Effects of figure type, instruction, and pre- and internal training on magnitude and decrement of illusion”, 1967 *Scandinavian Journal of Psychology*, **8**, 161–171.

- [46] Watson, A. “A Riemann geometric explanation of the visual illusions and figural after-effects”, 1978 *Mathematical and Theoretical Systems*, edited by Keat, 139–169.
- [47] Williams, G. P., 1997 *Chaos Theory Tamed*.
- [48] Woodworth, R. S., 1950 *Experimental Psychology*, fifth edition.
- [49] Yokose, Z., Uchiyama, M., and Yokoyama, A. “Shiteki gensyou to shigeki jikan tonon kankei ni suite I—Mie no nagasa-ookisa no shincyou-kakudaikatei” [The relationship between visual phenomena and the stimulating time (I)—About the growing process of the perceived length and size], 1957 *Japanese Journal of Psychology* (Japanese), **28**, 10–17.

Wide-band spectral power fluctuations characterize the response of simulated cortical networks to increasing stimulus intensity

John Klopp
University of Utah, 729 Arapeen,
Center for Advanced Medical Technologies,
Salt Lake City, UT 84108.
jklopp@doug.med.utah.edu

Patrick Johnston

Valeriy Nenov

Eric Halgren

Abstract

The hippocampus is an anatomically distinct region of the medial temporal lobe that plays a critical role in the formation of declarative memories. Here we show that a computer simulation of simple compartmental cells organized with basic hippocampal connectivity is capable of producing stimulus intensity sensitive wide-band fluctuations of spectral power similar to that seen in real EEG. While previous computational models have been designed to assess the viability of the putative mechanisms of memory storage and retrieval, they have generally been too abstract to allow comparison with empirical data. Furthermore, while the anatomical connectivity and organization of the hippocampus is well defined, many questions regarding the mechanisms that mediate large-scale synaptic integration remain unanswered. For this reason we focus less on the specifics of changing synaptic weights and more on the population dynamics.

Spectral power in four distinct frequency bands were derived from simulated field potentials of the computational model and found to depend on the intensity of a random input. The majority of power occurred in the lowest frequency band (3–6 Hz) and was greatest to the lowest intensity stimulus condition (1% maximal stimulus). In contrast, higher frequency bands ranging from 7–45 Hz show an increase in power directly related with an increase in stimulus intensity. This trend continues up to a stimulus level of 15% to 20% of the maximal input, above which power falls dramatically. These results suggest that the relative power of intrinsic network oscillations are dependent upon the level of activation and that above threshold levels all frequencies are damped, perhaps due to over activation of inhibitory interneurons.

1 Introduction

The human brain is extensively divided into functionally discrete and overlapping regions. Communication within the brain in the form of electro-chemical interaction is never ending and changes dramatically with the state of the organism. The cellular membranes of individual neurons have electrical properties that fluctuate over time and populations of neurons that oscillate in synchrony produce electrical signals that can be detected at the scalp with EEG. Electroencephalography (EEG) is a convenient method for obtaining non-invasive real time measurements of neural processes. EEG is generated by correlated activity of groups of cells oriented perpendicular to the recording electrode [12]. In traditional scalp EEG the recording contacts are located outside of the brain and therefore only synchronous activity that involves a range from thousands to millions of cells and as many as 1010 synapses contribute to the EEG signal. As a general rule, rhythmic oscillations of the EEG are associated with idle or resting states. This has historical roots stemming from early EEG studies that associated increased synchronous activity with decreased mentation [4]. However scalp EEG presents a global view of brain waves and some EEG rhythms may play a critical role in coordinating cognitive processing between distant cortical areas [3]. Moreover, EEG studies that use intracranial recordings have revealed local event related and wide-band spectral power fluctuations that are specific both to the region of the brain and the timing of the task [9].

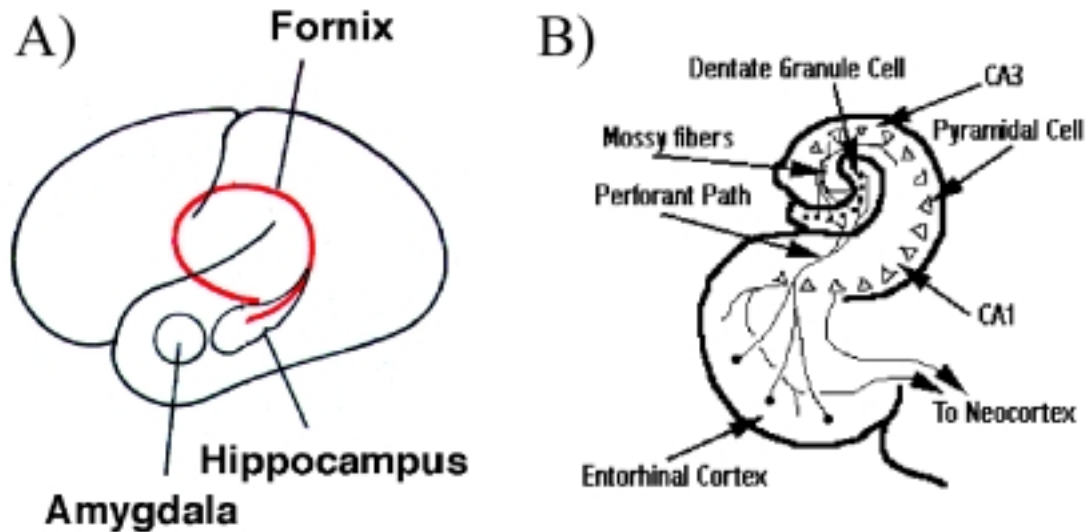


Figure 1: General anatomic location and connectivity of the hippocampal formation. A) Location of the hippocampal formation within the brain. The hippocampus is just posterior to the amygdala and runs along the inferior horn of the lateral ventricle in the medial temporal lobe. The fornix carries axons from CA3 pyramidal cells to the contralateral hippocampus. B) Structural subdivisions of the hippocampal formation and output to neocortex. The main direct input to the hippocampus originates from the perforant pathway that originates from the entorhinal cortex and granule cells in the dentate region.

We constructed a neural network model of the hippocampal region to test if increases in synaptic input evoked increases in spectral power as has been hypothesized from studies with intracranial EEG. The hippocampal region was chosen as a subject for modeling due to its relatively simple structure, its role in integrating information from many cortical areas and the large body of scientific literature that describes its form and function. The hippocampus is located along the inferior horn of the lateral ventricle in the medial temporal lobe and has direct and / or indirect reciprocal connections with all major neocortical multi-modal association areas as well as many sensory areas [2] (Figure 1A). Most of the neocortical input to the hippocampus is via the entorhinal cortex and dentate gyrus. The hippocampus is broadly divided into 2 major areas, CA3 and CA1, and we treat the dentate region as a simple input layer (Figure 1B). Taken together, the CA regions and the dentate region are commonly termed as the hippocampal formation. The synaptic connections between hippocampal cells are capable of rapid alteration of their post-synaptic response through activity dependent changes. This synaptic plasticity is Hebbian, in other words it relies upon depolarisation of the postsynaptic cell coincident with the activation of the presynaptic cell, and is specific for the activated synapse. These characteristics of activity dependent alteration of synaptic efficacy and wide spread reciprocal connectivity with neocortical associative areas are considered essential for the hippocampal formation's role in normal brain function as a temporary repository of declarative memory traces.

A number of recent neural models have identified the hippocampal formation as a temporary storage site for the declarative memory trace [8, 1, 13, 15, 10] These models have specifically identified the auto-associative network of pyramidal cells within the CA3 field as a prime candidate for the substrate of rapid short-term declarative memory formation and storage. Pyramidal cells of the hippocampal CA regions are organized in a regular laminar structure and are thereby capable of producing strong field potentials when groups of the cells are synchronously active. In order to simulate data acquisition, we included in our computer model simulated probes that sampled from

the neural network in a way comparable to intracranial EEG contacts placed within neural tissue. With this method we were able to collect data and apply analytical tools to the data in a manner similar to EEG collected from real neural tissue.

2 Computational Methodology

2.1 Modeling Platform

Computational models implemented at the Pittsburgh Supercomputing Center were written in PGENESIS [7], a parallel version of GENESIS (General Neural Simulation System) [5]. The standard GENESIS simulation package allows easy integration of simulation objects such as cellular compartments and cell-membrane channel conductances. PGENESIS allows a computer simulation to be distributed across multiple processors using Parallel Virtual Machine software and has been compiled and optimized for the CRAY T3E. The Cray T3E has 512 high-performance Digital Alpha 64-bit microprocessors (PEs), half running at 300 MHz and half running at 450 MHz. The 300 MHz PE's have a theoretical peak speed of 600 Mflops, and the 450 MHz PE's have a theoretical peak speed of 900 Mflops, bringing the machine's theoretical peak speed to 384 GFlops. The T3E's topology is that of a three-dimensional torus. Each processor runs a CHORUS-based microkernel. The memory is logically shared and physically distributed, with each PE having 128 MB.

2.2 The Single Cell Model

The model includes compartmental representations of excitatory pyramidal neurons, inhibitory interneurons and a large input layer of dentate granule cell spike generators. Pyramidal neurons consist of multiple compartments with a minimally branched dendritic morphology and incorporate fast Na^+ , $\text{K}^+(\text{dr})$, $\text{K}^+(\text{ahp})$ and Ca^{++} conductances (Figure 2A). Individual cells maintain compartmentalized membrane potentials that fluctuate with synaptic activation and neurons are capable of producing a variety of firing patterns and refractory periods including complex bursts and single spikes with a depolarisation envelope that is similar to real hippocampal pyramidal cells (Figure 3). Interneurons are modeled as a single compartment and include fast Na^+ and $\text{K}^+(\text{dr})$ currents giving a faster firing latency and a greater peak firing rate than the pyramidal cells (Figure 3B and Figure 4). Dentate cells are simple binary spike generators and are used for input into the network.

2.3 The Network Model

Of the CRAY T3E's 512 available PEs only 16 were required for the simulations presented here. The model consists of three regions (Figure 4). A population of 20736 spike-producing elements represent the dentate gyrus. Compartmental representations of 2592 pyramidal neurons and 144 interneurons represent the CA3 region. The CA1 region is twice the size of CA3 containing 5184 pyramidal and 288 inhibitory interneuron cells.

Synaptic interactions are modeled using a generalized alpha-function. Principle connections include sparse projections from the granule cells to CA3 pyramidal and interneurons, sparse, fast, recurrent excitatory connections within the CA3 region, diffuse fast feedback inhibitory and recurrent inhibitory connections with a slower time-course representing GABAB inhibition and Schaffer collaterals that project from CA3 pyramidal to CA1 pyramidal cells. Axonal delays were modeled with a realistic cable velocity of 0.67 meters per second. A narrow Gaussian variability of axonal delays was used to avoid network artifacts. Both types of CA3 cells receive feed forward excitation via sparse connections from the dentate gyrus spiking input elements. However, this 'mossy fiber' input does not yet account for non-associative LTP seen in these synapses. A standard Hebbian algorithm updates synaptic efficacy of the associative pyramidal connections. Contact probability diminishes with distance in a series of concentric ellipsoids (Figure 5A) and the processor workload is distributed in a design that biases output targets to remain on the originating processor and is also more accurate than




Figure 2: Schematic of simulated compartmental cells. A) Minimally branched dendritic compartments of the pyramidal cell contain various ionic membrane channels and active synaptic connections from other cells. The cell soma produces spike events when it reaches an electrical threshold. Spike events are passed along axonal lines to other cells where the spike event produces a depolarization of the target cell. B) The inhibitory interneurons are single compartment models with ionic membrane channels and active synaptic connections to and from other cells.

a completely distributed connectivity scheme (Figure 4 and Figure 5B). This design reduces inter-processor communication and optimizes the speed of simulation.

2.4 Simulation Conditions

Input to the neural network model was delivered through the layer of dentate spike generators. Random spatial and temporal activation of this region was altered from 0 to 100% where a value of 3% indicates that a random subset containing 3% of the total population of dentate spike generators were activated with zero auto-correlation in the dentate spike trains. Given that physiological levels of dentate activity are probably maximal at relatively low percentages, 1% increments were used from 0 to 10% and increasingly larger increments from 12 to 100% activation. In all, 19 simulations were performed with levels of dentate activation at: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 12, 15, 20, 30, 40, 60, 80 and 100%. Each simulation was allowed to run for 520 ms. This translated into approximately 2 hours of real time on the CRAY T3E.

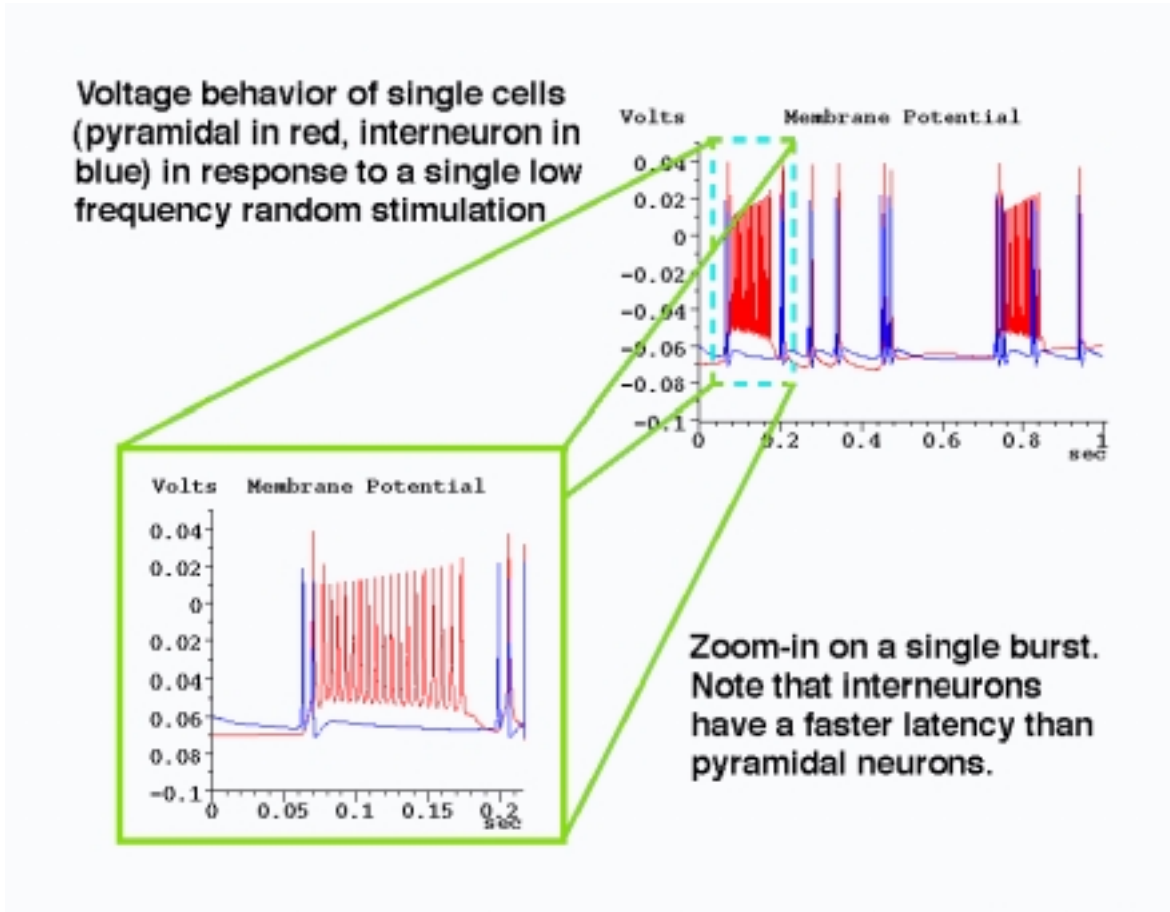


Figure 3: Voltage waveform behavior of single cells. Simulated membrane potentials of a pyramidal cell are shown in red and an inhibitory interneuron in blue. The inset shows a single pyramidal burst event. The time courses and refractory periods of individual cells are a critical factor contributing to the networks intrinsic resonance frequencies.

2.5 Data Collection and Analysis

Spike data was collected from all CA1 and CA3 pyramidal cells and interneurons. Simulated electrodes produced field potentials that sampled activity from pyramidal cells in the CA3 region using current sources and their distance from the electrode site to calculate the field. If one assumes that the medium surrounding neurons is of homogeneous resistivity and has no capacitance, then the field potential generated by a compartmental model can be calculated from the following equation:

$$F = \frac{1}{4\pi s} \sum_{i=1}^n \frac{I_{m_i}}{R_i} \text{ (see [11])}$$

where F is the field potential in volts, s is conductivity in $\Omega^{-1}m^{-1}$, I_{m_i} is the transmembrane current (Amperes) across the i^{th} neural compartment, and R_i is the distance from the i^{th} neural compartment to the recording electrode (i.e. the simulated electrode). Data analysis was performed on an SGI workstation using the S-Plus (Mathsoft). statistical language. Computation of spectral power was performed on 500 ms epochs of simulated field potential. Each epoch was analyzed using an un-normalized discrete Fourier transform. The 19 simulations yielded 76 power measurements for




Figure 4: Network diagram of the model. Input to CA3 is entirely mediated through spike generators in the dentate. Each region is simulated by multiple PEs on the CRAY T3E. In this example of the network connectivity the area simulated by two PEs are marked as '8' and '0'. Because area 8 of the dentate projects primarily to area 8 of the CA3 and area 8 of the CA3 projects primarily to area 8 of the CA1 inter-PE communication is reduced and the simulation speed is increased.

4 distinct frequency ranges. The frequency bands examined were low (3–6 Hz), alpha (7–12 Hz), beta (13–24 Hz) and gamma (25–45 Hz). Resulting power values were normalized to a percentage of the maximal obtained power measure.

3 Results

The percentage of cells that fired and the number of active cells producing single versus multiple spikes varied with the level of stimulus intensity. As a general trend, increasing the level of stimulation elicited an increase in the percentage of active CA3 pyramidal cells. Due to feed forward inhibition from the CA3 and dentate to CA1 interneurons increased stimulus levels resulted in fewer CA1 pyramidal cell spikes. Figure 6 illustrates spike activity over time for two simulation conditions (1% and 10% input). A large early activation or impulse response occurred in many CA3 cells to the higher stimulus conditions centered at around 120ms. Cells are displayed in an arbitrary numerical order and the spatial organization of the model is not preserved in Figure 6. In contrast, the longitudinal and transverse axes of the model are preserved in Figure 7. This figure shows the total number of spikes produced by any given cell and its appropriately arranged neighbors for the duration of the simulation in the 1% and 10% input conditions.

Power values were derived from simulated CA3 field potentials and were sensitive to the level of stimulus intensity. The majority of signal power occurred in the low frequency band and was maximal at an input level of 1%. Maximal alpha power occurred at an input level of 15% and was approximately 1/3 the amplitude of the maximal low frequency power. Maximal beta and gamma power occurred at input levels of 9% and 10% and were 0.1% and 0.025% of maximal low frequency power respectively. Power in all frequency bands was depressed relative to maximal values in response to stimulus conditions at and above 20% (Figure 8).

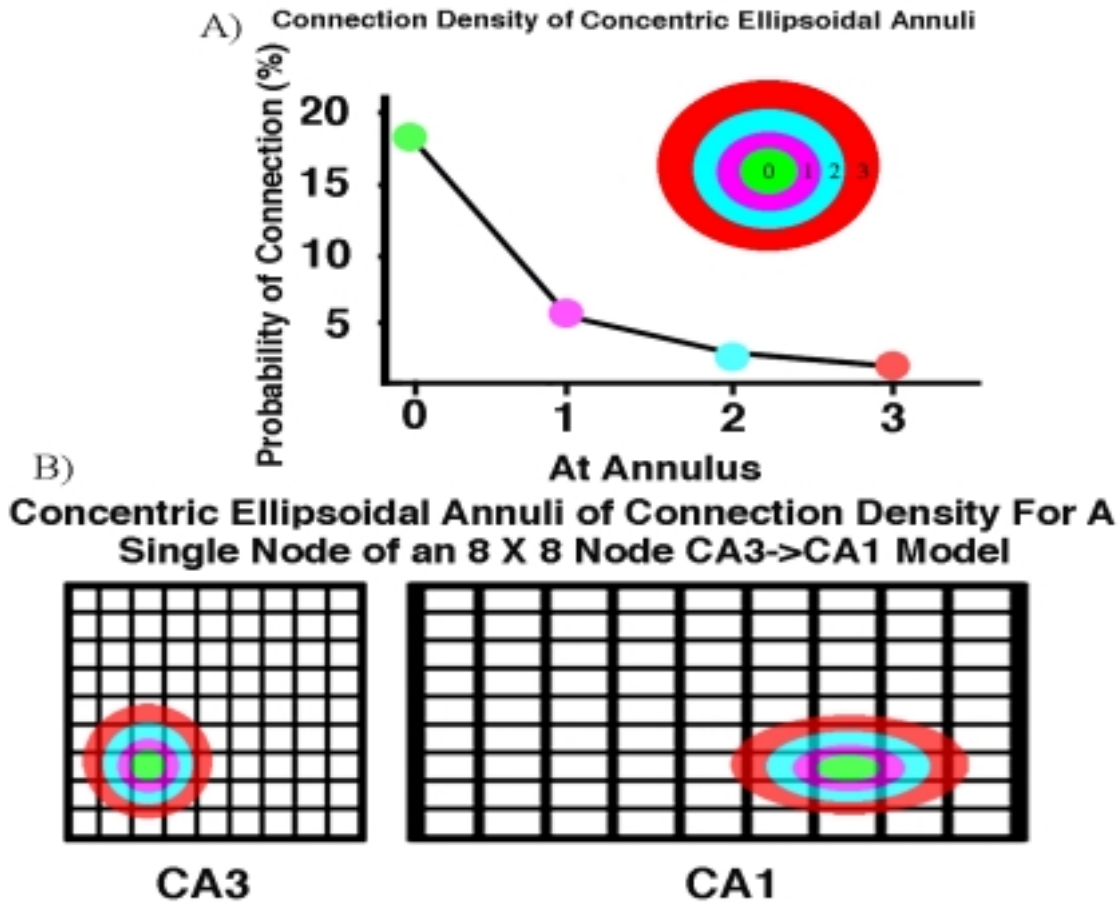


Figure 5: Diagrams of the pyramidal connectivity scheme. A) The probability of synaptic connection between two cells decreases with distance. B) The model is separated into nodes where a single PE is responsible for all simulation and message passing to and from a single node. Pyramidal connections within and between nodes are made within concentric ellipsoidal annuli. Inhibitory interneurons make simpler, more global connections with a higher probability of connectivity.

4 Conclusion

A common trait that is shared among a wide range of complex systems is the tendency to resonate. Many structures, due simply to the nature of their composition, are prone to oscillate at certain frequencies while other frequencies are damped out. In the field of engineering resonance can often have dangerous side effects such as collapsing bridges and exploding turbine generators. Much effort has gone into determining the oscillatory characteristics of neural systems. Do network oscillations occur as a property of individual neurons that oscillate intrinsically, or is it only through synaptic interactions between groups of cells (that individually would not oscillate) do oscillations emerge? While the answer to this question depends on the group of cells that one is studying, inter-neuronal feedback is considered the basis for cortical oscillations seen in the EEG rather than intrinsic autonomous oscillations of single cells [6, 14].

The input to the computational simulation was random with zero auto-correlation in the spike trains, as evidenced by the spatially pattern-less activity in Figure 7. However, vague evidence of rhythmic activity can be seen in Figure 6. This disposition to oscillate is born out in the spectral power measurements. Therefore, changes in the spectral power can be attributed to activation of

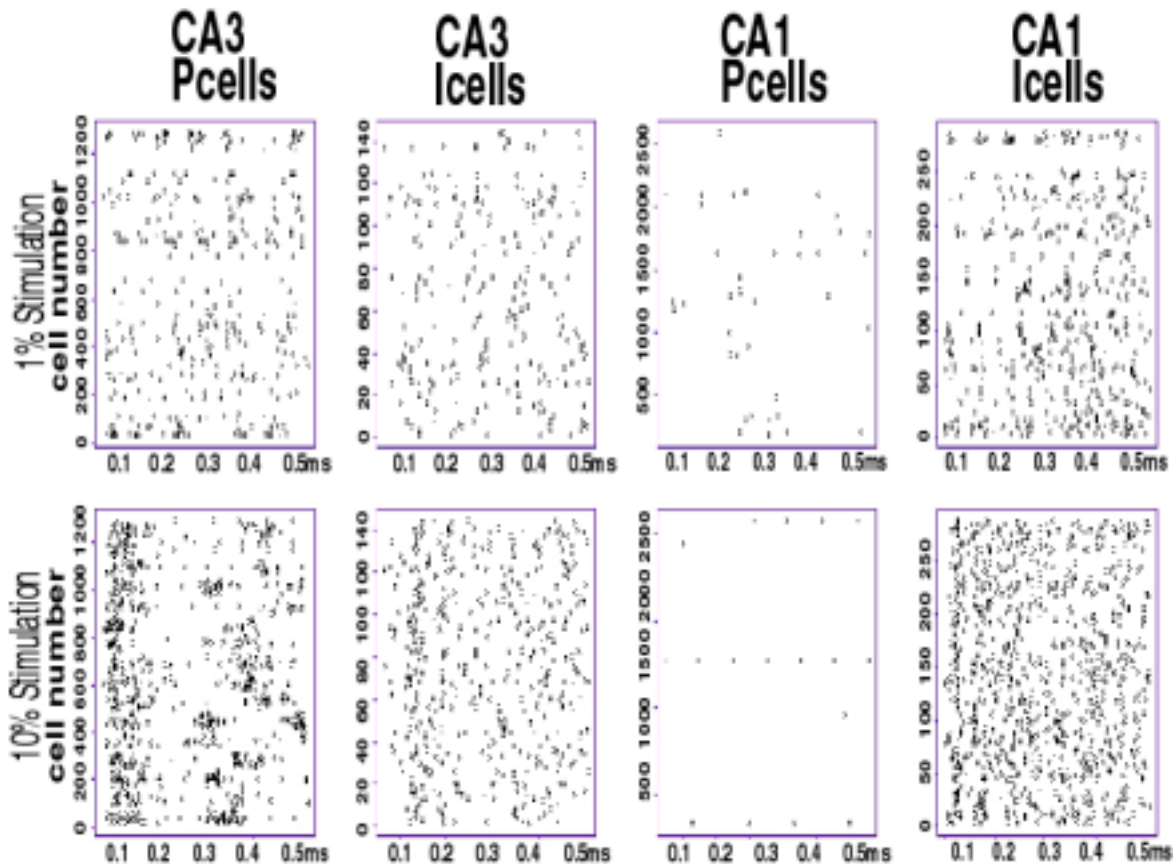


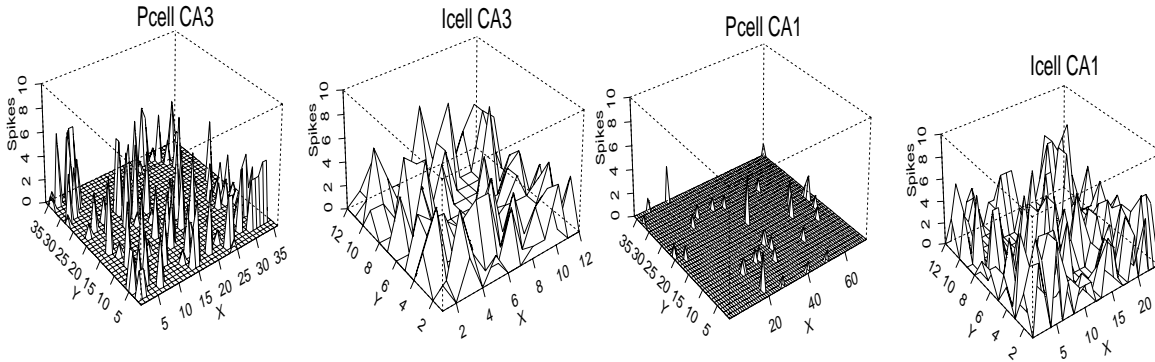
Figure 6: Spike activity over time for two simulation conditions. From top to bottom the percentage of maximal stimulation is 1% and 10%. From left to right spike events from each region are CA3 pyramidal, CA3 interneurons, CA1 pyramidal and CA1 interneurons. An early hypersynchronous impulse response in the CA3 pyramidal spike activity is apparent in the 10% stimulus conditions centered at about 120 ms.

intrinsic resonant frequencies. The most striking of these is in the low frequency and is highly active in response to the lowest stimulus condition of 1%. Above this level of stimulation low frequency power is greatly diminished and appears not to be influenced in any great measure by further increasing stimulus intensity.

Higher frequency components show strikingly different properties. From 7–45 Hz power increases with stimulus intensity in multiple frequency bands. This trend peaks at around 9% to 15% of the maximal stimulus intensity. Above this level of stimulation, power from 7–45 Hz is greatly diminished. This is, in a way, reassuring given that simulations that activate over 20% of the input layer in a period of 500 ms are not physiologically relevant, except perhaps under pathological conditions. These results suggest that a simple network of simulated cells is capable of producing changes in the spectral characteristics of its field potentials. Similar changes have been observed in the spectral response of real neural systems from intracranial EEG recordings acquired in the medial temporal lobe of human epileptic subjects [9].

Future simulations will be larger in order to attain a more realistic sparseness of connectivity. Results presented here were obtained from simulations using only 16 PEs on the CRAY T3E. Despite the potential to scale up the model to take advantage of all 512 PEs available on the T3E, increasing the size of the model causes an exponential increase in memory requirements and simulations beyond

1% Stimulus



10 % Stimulus

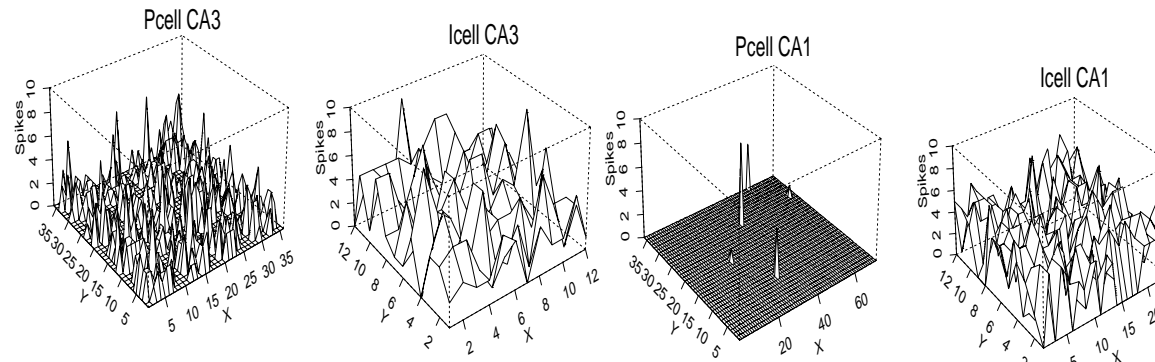


Figure 7: Spatial organization of cells and spikes from the same data presented in figure 7. The percentage of CA3 pyramidal cells that produce spikes increases from 8% to 31% respective to the increasing stimulus conditions (1% and 10% maximal stimulus). The number of CA1 pyramidal cells that produced spikes decreased with increasing stimulus conditions from 1, to 0.2%. Both CA3 and CA1 interneuron populations increased the number of cells that produced spikes with the increasing stimulus condition.

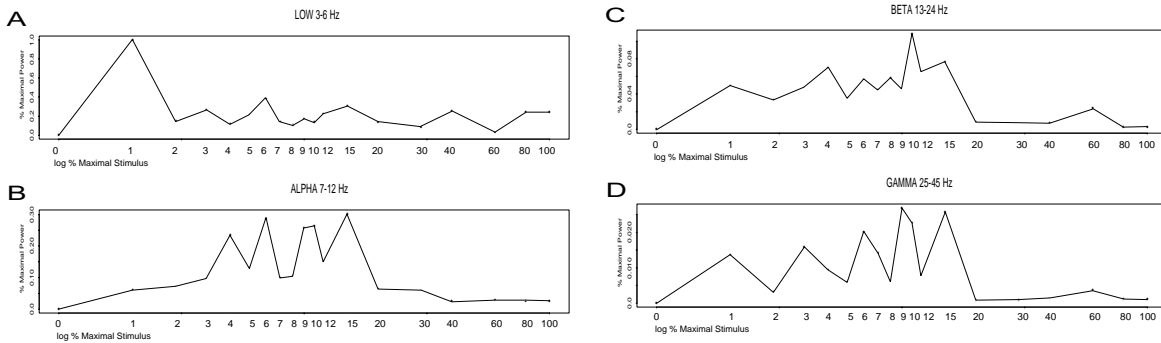


Figure 8: Event related spectral power results from 19 simulations where stimulus intensity was varied between 0 and 100%. Most power was contained in the low (3–6 Hz) frequency band. Higher frequencies (7–45 Hz) showed a wide-band sensitivity to stimulus intensity from 0 to 15%. Above this the power in all frequencies is largely diminished.

64 PEs quickly run out of available memory. In its current form the model has been run on as many as 64 PEs. However, these large simulations are time costly and have not been run enough to produce reliable results.

While some form of associative synaptic alteration lends to the realism of the model, in short simulations the changes of synaptic weights are probably negligible from the perspective of field potentials. Moreover, the computation of the Hebbian algorithm at every pyramidal synapse slows down the simulation. The next generation of this model will be tested with and without the Hebbian plasticity that was incorporated here. Longer simulations should be attempted to stabilize the network after the impulse response and also to see if exclusion of the Hebbian synapses has any impact on network dynamics.

The levels of stimulus were largely unrealistic and rather than cover a full range of 0 to 100%, future simulations will focus on a more realistic stimulus range of 0 to 6%. This will allow us to use smaller percentage steps between simulations and should yield a more realistic level of CA activity. Field potentials will be collected from the CA1 in addition to the CA3 in the next simulations to allow for comparison between the activity in these anatomically distinct regions. The spectral response of the CA3 region to the random input suggests that intrinsic resonance features of the CA3 are capable of reorganizing white noise activity into a more structured spectral landscape. It is possible, since CA1 is a major projection target for CA3, that field potentials have even greater spectral structure and sensitivity in the CA1 region.

Finally, although the hippocampus is a simplified cortical structure, observations here are generally applicable to neocortex. The human CA3 may contain as many as 3 million pyramidal cells with twice as many in the CA1 region [2]. Furthermore, the hippocampus contains a far more complex geometrical structure and synaptic connectivity than is simulated here. At best our model can be viewed either as a sparse sample of the hippocampal formation or a small percentage of the hippocampal formation. A properly scaled network model may reveal wide-band oscillatory patterns that would otherwise not appear in the current model.

Acknowledgements

This research was supported by the National Science Foundation under the Pittsburgh Supercomputing Center grant # IBN970004P and USPHS-NS18741, Human Frontiers Science Program, and ONR.

References

- [1] P. Alvarez and L. Squire. Memory consolidation and the medial temporal lobe: a simple network model. *Proc Natl Acad Sci U S A*, 91:7041–7045, 1994.
- [2] D. Amaral and R. Insausti. The hippocampal formation. In G. Paxinos, editor, *The Human Nervous System*, pages 711–755. Academic press : New York, 1990.
- [3] E. Basar and T. Bullock. *Induced Rhythms in the Brain*. Birkhauser : Boston, 1992.
- [4] H. Berger. Über das Elektrenkephalogramm des Menschen. *Arch. Psychiatr. Nervenkr*, 87:527–570, 1929.
- [5] J. Bower and D. Beeman. *The Book of GENESIS: exploring realistic neural models with the GEneral NEural Simulations System*. TELOS : New York, 1995.
- [6] W.J. Freeman. Predictions on neocortical dynamics derived from studies of paleocortex. In Basar and Bullock, editors, *Induced Rhythms of the Brain*, pages 183–201. Birhauser : Boston, 1992.

- [7] N. Goddard and G. Hood. Parallel genesis for large scale modeling. In J. Bower, editor, *Computational Neuroscience: Trends in Research*. Plenum Press, 1997.
- [8] E. Halgren. Human hippocampal and amygdala recordings and stimulation: Evidence for a neural model of recent memory. In L. Squire and N. Butters, editors, *The Neuropsychology of Memory*, pages 165–181. Guilford : New York, 1984.
- [9] J. Klopp, E. Halgren, K. Marinkovic, and V. Nenov. Focal fluctuations in event-related spectral power specific to faces in the human fusiform gyrus. To appear, 1998.
- [10] J. McClelland, B. McNaughton, and R. O’Reilly. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev*, 102:419–457, 1995.
- [11] P.L. Nunez. Electric fields of the brain. In *The Neurophysics of EEG*. Oxford University Press : Oxford, 1981.
- [12] P.L. Nunez. unknown reference. Reference missing, 1995.
- [13] W. Read, V. Nenov, and E. Halgren. Role of inhibition in memory retrieval by hippocampal area CA3. *Neurosci Biobehav*, 18:55–68, 1994.
- [14] F.H. Lopez Da Silva. Dynamics of electrical activity of the brain, local networks and modulating systems. In P.L. Nunez, editor, *NeoCortical Dynamics and Human EEG Rhythms*, pages 249–271. Oxford University Press : New York, 1995.
- [15] A. Treves and E. Rolls. Computational analysis of the role of the hippocampus in memory. *Hippocampus*, pages 374–391, 1994.

Mechanism for Changing the Foraging Behavior in an Ant Colony Model

Mari Nakamura
 Life Electronics Research Center
 Electrotechnical Laboratory (ETL)
 mari@etl.go.jp

Koichi Kurumatani
 Information Science Division
 Electrotechnical Laboratory (ETL)

Abstract

In this paper, we propose an improved ant colony model in which the foraging behavior of ants can be observed on a macro-scale as a result of micro-scale interacting behaviors among many individual ants. Simulation results of the new model have shown that the system changes its foraging behavior and selects the appropriate foraging strategy according to the food-supply rate. The mechanism for changing the foraging strategy is explained.

1 Introduction

In a previous paper [1], we proposed a model for observing the foraging behavior of an ant colony composed of many ants whose sensitivity to stimuli was limited due to the small area surrounding them. With this model, when the ants are collecting food from sites outside of their nest, they behave based on the following rules.

An ant looking for food (*i.e.*, in a *searching* task) walks randomly until it finds a food site or detects *recruitment pheromone* signal [2]. If it finds a food site, it changes its task to a *carrying* task. An ant involved in the carrying task retrieves a bit of food straight to its nest while laying recruitment pheromone on the ground. After carrying the food to the nest, it returns to the searching task. The laid recruitment pheromone gradually diffuses over a wide area. If other ants involved in the searching task detect this pheromone signal, they immediately change their task to the *recruited* task. Ants involved in the recruited task are attracted to the pheromone, and follow the pheromone's trail to the food site. After finding the food site, they change their task to the carrying task.

We performed a simulation of the above-described model. We observed a *recruitment competition* process among food sites; the process resulted in *over-concentrated* recruitment at a single site [1], [3]. During the competition, most of ants in the colony were assigned to the carrying or recruited tasks. This resulted in fewer ants involved in the searching task localized around the nest.

We suppose that ants were desensitized to the pheromone, in order to circumvent such an unwanted assignment.

Desensitization Supposition

When an ant perceives an over-concentration of recruitment around itself, it becomes desensitized to the recruitment pheromone for a certain period.

To incorporate this supposition into the model, we have established a rule stating that *when an ant perceives a strong pheromone signal, it becomes desensitized*. A simulation of the model was performed with this rule. We observed that when simulated with the appropriate parameter sets [1], the system organizes a stable, distributed recruitment at several food sites, and a concentrated recruitment at a single site. According to this rule, the system works to keep the strength of an organized pheromone pattern as constant as possible, and not to optimize its foraging. This is because the perception of over-concentration is determined by the presence of strong pheromone signal, not by food.

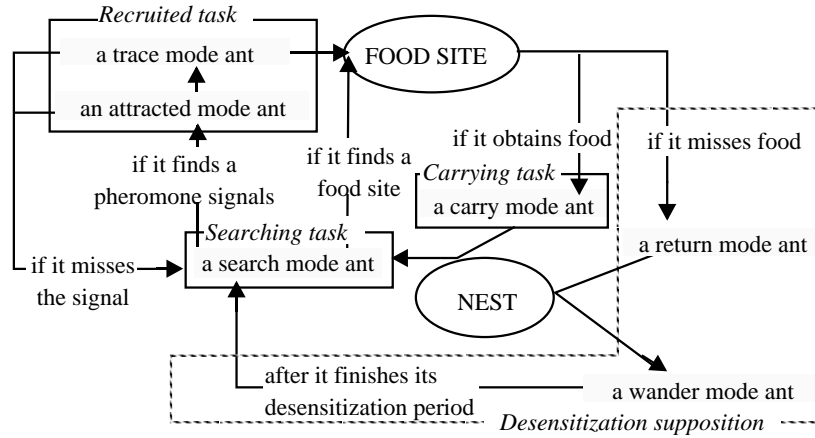


Figure 1: Mode transition rule of the desensitization model in this research.

In this paper, we changed the rule of desensitization to improve the foraging efficiency according to the food supply rate. The new rule states that *when an ant misses food at a food site, it becomes desensitized*. Under this rule, after the ants eat up a food site, they become able to search for new sites through pheromone signals. The construction of the model is given in chapter 2. Simulations of the model were made under different food-supply conditions. The results of the simulations showed that the system organized different recruitment patterns in response to the food-supply rate, that are corresponding to the change in foraging strategies. These results are stated in Chapter 3. In chapter 4, the mechanism responsible for organizing the patterns and changing the foraging strategies is explained.

2 Model

2.1 Behavior of an Ant

In this revised desensitization model, the behavior of ants is determined to increase the foraging efficiency (*i.e.*, amount of food collected by the colony). Ants change their behavior according to local situation around them, defined as below (illustrated in Fig. 1).

Searching task An ant in the *search* mode walks randomly. If it obtains or misses food at a food site, it changes its mode to the *carry* or the *return* mode, respectively. When it detects a pheromone signal (see 2.2), it changes its mode to the *attracted* mode.

Recruited task An ant in the *attracted* mode goes toward stronger pheromone. If it finds a pheromone trail, it changes its mode to the *trace* mode. If it misses a pheromone signal, it returns to the *search* mode.

An ant in the *trace* mode follows the pheromone trail toward the corresponding food site. If it obtains or misses food at the food site, it changes its mode to the *carry* or the *return* mode, respectively. If it misses the trail, it returns to the *search* mode.

Carrying task An ant in the *carry* mode goes toward its nest with food, while laying recruitment pheromone on the ground. After carrying the food to its nest, it returns to the *search* mode.

Desensitization An ant in the *return* mode goes toward its nest without laying pheromone. After it reaches its nest, it changes its mode to the *wander* mode.




Figure 2: Signal of recruitment pheromone laid by an ant.

An ant in *wander* mode walks randomly without perceiving any stimulus. After a desensitization period, it returns to the *search* mode.

2.2 Behavior of Recruitment Pheromone

Recruitment pheromone laid on the ground gradually diffuses over a wider area. This diffusion is formulated as follows;

$$(d/dt + \gamma_{\text{vap}})T(x, y) = 0 \quad \text{Evaporation:} \quad (1)$$

$$(d/dt - \gamma_{\text{dif}}\nabla)P(x, y, z) = \begin{cases} \gamma_{\text{vap}}T(x, y) & (z = 0) \\ 0 & (z > 0) \end{cases} \quad \text{Diffusion:} \quad (2)$$

In these differential equations, $P(x, y, z)$ denotes the *density of evaporated pheromone in the air*, and $T(x, y)$ denotes the *strength of the pheromone trail on the ground*. The region where $P(x, y, 0) \geq P_{\text{thr}}$ is defined as *an attracting zone*, and the region where $T(x, y) \geq T_{\text{thr}}$ is defined as *an active trail*, as shown in Fig. 2.

2.3 Parameter Settings

The simulated system in this paper has a nest at the center of the field and four food sites surrounding the nest. At every food site, a certain quantity of food is supplied per unit time. The other important parameters are listed below:

- The expanse of the simulated space covers $0 \leq x < 100\Delta x$, $0 \leq y < 100\Delta y$, and $0 \leq z < 3\Delta z$. For pheromone diffusion, the ground ($z = 0$) is a reflecting boundary and the other boundaries are absorbing boundaries.
- The evaporation and diffusion factors are given as $\gamma_{\text{vap}} = 0.15\Delta t$ and $\gamma_{\text{dif}} = 0.42(\Delta x^2/\Delta t)$. These factors determine the time constant of the trail evaporation and the effective range of the pheromone diffusion. P_{thr} and T_{thr} , the minimum perceivable pheromone and trail, are set so that the pheromone signal laid by an ant fades out after a few steps.
- The desensitization period is determined as 100 steps, that is long enough for wander mode ants to spread out over the field. The system is simulated over 2000 steps, to converge to a near equilibrium state. An ant moves within two grids in one step. The length of a step and that of a grid are $\Delta t = O(1 \sim 10\text{sec})$ and $\Delta x = \Delta y = \Delta z = O(0.01 \sim 0.1\text{m})$.

3 Simulation Results

Simulation results showed that the system organizes the following two types of recruitment patterns, in response to the food supply rate.




Figure 3: Concentrated recruitment to a single site. Fourteen food units are supplied per step at every site.

Case I: When the food supply is sufficient (food supply ≥ 7 units per step) Concentrated recruitment to a single food site is organized, as shown in Fig. 3. Most of the ants are assigned to the carrying and recruited tasks, and fewer ants are assigned to the searching task. The ratio of desensitized ants (*i.e.*, ants in the return and wander modes) is relatively small, as shown in Fig. 5. For a larger food supply, the ratio of desensitized ants is smaller. When food supply \geq thirteen units per step, no ant becomes desensitized because food supply at a single site exceeds maximum amount of foods collected by the colony per step.

Case II: When the food supply is insufficient (food supply ≤ 6 units per step) Stable distributed recruitment to several food sites is organized, as shown in Fig. 4. Ants in the colony are assigned to the searching task, as well as to the carrying and recruited tasks. The ratio of desensitized ants is relatively large. For a smaller food supply, both the ratio of desensitized ants and the number of food sites at which ants are recruited are larger, as shown in Fig. 5. In this case, the amount of food collected by the colony per step is larger than the food supply at a single site. When food supply rate is between three to six units per step, average collected foods, average rate of each mode ants are almost constant. With food supply $<$ two units per step, ants eat up all supplied foods, and rate of desensitized ants is large.

From the viewpoint of mathematical biology, the system autonomously changes its foraging strategy in response to the food supply, to increase the food-collecting efficiency. When the food supply is sufficient (case I), most ants go and return between a single food site and their nest, without wasting their time on random-walking. When the food supply is insufficient (case II), however a wide dispersion of many ants in the wander mode maintains recruitment at several food sites, from which the colony carries food.




Figure 4: Distributed recruitment to several sites. Four food units are supplied per step at every site.

4 Mechanism for Changing Foraging Strategies

4.1 General Explanation of the Mechanism

The mechanism responsible for controlling the changes in recruitment patterns and foraging strategies can be explained as follows. In a model without desensitization supposition, severe recruitment competition has been observed between food sites for search mode ants around the nest (details will be explained in 4.2), even with small food-supply rate. With the new desensitization rule, desensitization is induced to distribute search mode ants over the field against strong pheromone signals. As a result, wander mode ants spread over the field. After the desensitization period, these wander mode ants return to the search mode and are recruited to several food sites (details will be explained in 4.3). This interferes with the recruitment competition process, and causes stable distributed recruitment to several food sites.

The changes in organized patterns and foraging strategies are determined by the ratio of wander mode ants affecting the recruitment competition process. When the food supply is insufficient, the ratio is large enough to suppress the competition process, and when the food supply is sufficient, the ratio is too small to interfere with the process.

4.2 Explanation of Recruit Competition Process

At first, we consider the case of no desensitized ants, to explain the mechanism for the recruitment competition process (as illustrated in Figs. 6-1 and 6-2). In this section, the model is assumed to have only two food sites, to simplify the explanation.

In Fig. 6-1, S denotes the distribution of search mode ants who have just changed their mode from carry mode. S concentrate around the nest as illustrated in Fig. 6-1. The nest is located at $(0, 0)$, and the ends of two trails are located at $(\pm h, 0)$. The intersections between S and the two attracting zones at time i are defined as $S_1(i)$ and $S_2(i)$ where $S_1(i) + S_2(i) = 1$. $x(i)$ is defined as the drift of the boundary between $S_1(i)$ and $S_2(i)$.

$T_1(i + 1)$ and $T_2(i + 1)$, the strength of the trail to food site 1 and that to site 2 after a certain period of time, grow proportional to $N_1(i)$ and $N_2(i)$, *i.e.*, numbers of ants recruited to both sites.






Figure 5: Recruit pattern and foraging behavior versus food supply.

$N_1(i)$ and $N_2(i)$ are proportional to $S_1(i)$ and $S_2(i)$. Then, we obtain the following equation.

$$T_1(i+1)/T_2(i+1) = S_1(i)/S_2(i) \quad (3)$$

Let $u(i)$ denote $S_1(i)/S_2(i)$. $S_1(i)$ and $S_2(i)$ are determined as $1 - S(x(i))$ and $S(x(i))$. Then, $u(i)$ is determined by the following equation.

$$u(i) = \frac{S_1(i)}{S_2(i)} = \frac{1}{S(x(i))} - 1 \quad (4)$$

Let $P(x)$ denote the pheromone density at the boundary. $P(r)$ is proportional to $f(r)$ (r : distance from the pheromone trail, $f(r)$: a positive, monotonically decreasing function of r). Then, $P(x)$ is determined as follows;

$$\begin{aligned} P(x(i+1)) &\propto T_1(i+1)f(h-x(i+1)) = T_2(i+1)f(h+x(i+1)) \\ \therefore u(i) &= f(h+x(i+1))/f(h-x(i+1)) \end{aligned} \quad (5)$$

$u = (1/S(x)) - 1$ and $u = f(h+x)/f(h-x)$ are drawn graphically in Fig. 6-2. As $S(-h) = 0$ and $S(h) = 1$, the curves cross as shown. This figure indicates that x is finally driven to $\pm h$ and u is finally driven to 0 or ∞ . This means that the result is concentrated recruitment to either site.

4.3 Explanation of Distributed Recruit in Desensitized Model

Next, we consider the other case that all search mode ants return from desensitization (illustrated in Fig. 6-3), to explain how desensitization supposition enables stable distributed recruitment. In

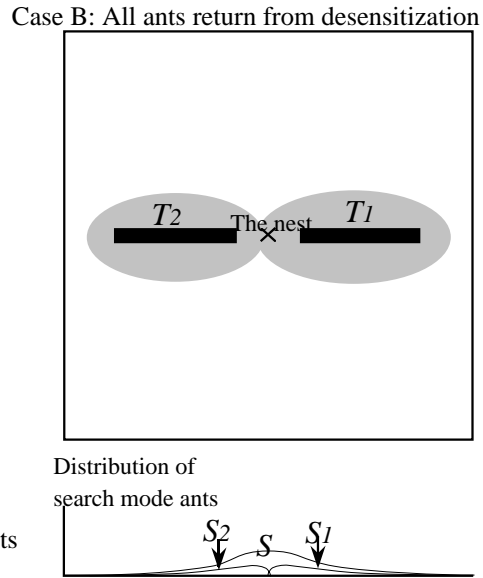

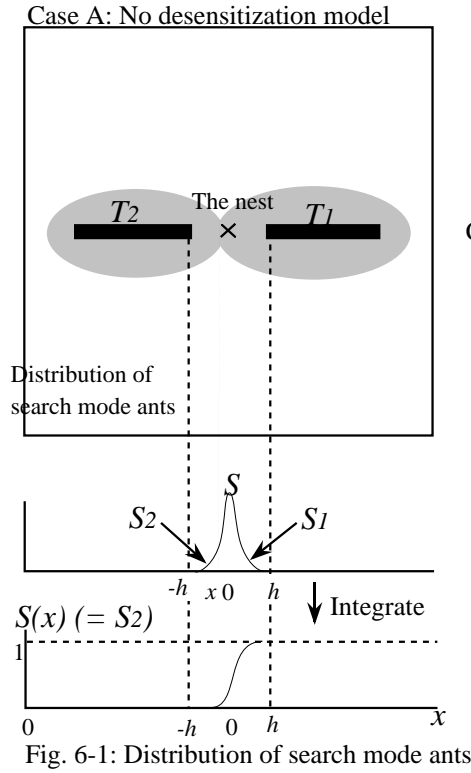


Figure 6: Mechanism for change of a recruitment pattern and foraging behavior.

Fig. 6-3, S denotes the distribution of search mode ants who have just changed their mode from the wander mode. S is non-zero over a wide area of the field. $S_1(i)$ and $S_2(i)$ denote the intersections between S and the two attracting zones at time i . In this case, there are two important differences from the previous case.

- $S_1(i)$ and $S_2(i)$ cannot cover the total S ($S_1(i) + S_2(i) < 1$). That is, many search mode ants are located outside of the attracting zones; gradually, they become trapped in the attracting zones after random-walking.
- The expanse of border between $S_1(i)$ and $S_2(i)$ is too small when compared with $S_1(i)$ or $S_2(i)$. Accordingly, the interference between the attracting zones is negligible when considering the growing process of the attracting zones.

In this case, each attracting zone grows independently until reaching the upper limit determined by the number of search mode ants returned from desensitization. This results in distributed recruitment to all sites equally.

5 Discussion

In this paper, the model is simulated under static food-supply conditions. This model will show better foraging behavior when simulated in dynamic food-supply condition, where search of new food sites is essential to foraging [4].

In this model, desensitization induces negative feedback from macro-scale foraging behavior of the system to micro-scale ants behavior. Desensitization induces *noise* into the system on demand, in order to control system's behavior properly. Similar mechanism for changing foraging strategies according to *noise* in ants' behaviors was proposed [5]. This system can escape from local minimum (overconcentrated state) by use of noise (desensitized ants). Some stochastic process may allow us to control of such a nonlinear distributed system.

6 Conclusion

In conclusion, an improved ant colony model that can increase the efficiency of foraging was proposed. In this model, the system shows change in recruitment patterns corresponding to change of foraging strategies, according to the food supply rate, and maintains a high food collecting efficiency.

7 Acknowledgments

One of authors, M. Nakamura, owes special thanks to Dr. Y. Kakui and Dr. T. Akiba, her colleagues in ETL, and to Mr. Y. Adachi in Kanazawa Institute of Technology.

References

- [1] Nakamura, M. and Kurumatani, K. "Formation Mechanism of pheromone pattern and Control of Foraging Behavior in an Ant Colony Model", 1997 *ALife V (The MIT Press)*, 67-74
- [2] Hoelldobler, B. and Wilson, E. "The Ants", 1990 (*the Belknap Press*)
- [3] Kurumatani, K. and Nakamura, M. "Generating Qualitative Equations about Macro-behaviors of Foraging in Ant Colony", 1996 *Proc. of the German Conf. on Bio-informatics, GCB'96*
- [4] Stephens, D. W. and Krebs, J. R. "Foraging Theory", 1986 (*Princeton University Press*)
- [5] Deneubourg, J. L. Pasteels, J. M. and Vergaenghe, J. C. "Probabilistic behavior in ants; A Strategy of errors?", 1983 *J. of Theor. Biol.* , **105**, 259-271

The First Time-Derivative of the EEG: A Possible Proxy for the Order-Parameter for the Cerebral Cortex

J.W. Sleigh
 Department of Anaesthetics
 Waikato Hospital
 Hamilton, New Zealand
 sleighj@hwl.co.nz

D.A. Steyn-Ross and Moira Steyn-Ross
 Department of Physics and Electronic Engineering
 University of Waikato
 Hamilton, New Zealand

Abstract

Many important aspects of the function of the cerebral cortex can be captured in a two dimensional lattice model. From this analogy, the change from the awake state to the unconscious state can be understood as a form of order/disorder phase transition. If this is so, there should exist an order-parameter that has zero value when the cortex is disordered (the anaesthetic state), and which rapidly climbs to an arbitrary positive value when the cortex becomes ordered (the awake state). Although the ‘spatially-measured soma potential’ \bar{v} of the cortex, relative to its unconscious state value \bar{v}_0 , can be considered to be the order-parameter, it is not possible to measure the mean soma potential directly. However, fluctuations in the soma potential give rise to the time-varying EEG signal $v(t)$ which is easily measured.

In this paper we hypothesise that the first time-derivative of these fluctuations dv/dt is sufficiently strongly linked to the mean soma potential that RMS or absolute values of dv/dt can serve as a proxy order-parameter for the cortex.

One-dimensional and two-dimensional cellular automaton (CA) simulations were run to test this hypothesis. The 1-D CAs showed an increase in $(dv/dt)_{\text{abs}}$ for more complex CAs (as defined by Wolfram’s classification). In the 2-D simulations there was a strong linear correlation between $(dv/dt)_{\text{rms}}$ and \bar{v} (having correlation coefficient $r = 0.92$). The EEG signal was recorded in 23 patients during induction of general anaesthesia. Overall the $(dv/dt)_{\text{rms}}$ value decreased by 70% when comparing the awake values with those from the anaesthetised state. We conclude that the (dv/dt) measure shows promise as an easily extracted proxy order-parameter, and that the transition from the awake state to a state of general anaesthesia is associated with a decrease in $(dv/dt)_{\text{rms}}$ consistent with the behaviour expected of an order-parameter.

1 Origin of the Electroencephalograph (EEG) Signal

The cerebral cortex has been modelled widely as a two-dimensional lattice. Its function depends on internal interactions between neurons (or neuronal assemblies), and input from the primitive brain. Generally, the neuronal interactions occur at synapses according to the following sequence:

- (i) Action potential (spike) in the pre-synaptic neuron, results in
- (ii) Neurotransmitter release, which causes
- (iii) Alteration of post-synaptic potential (PSP). Each perturbation is followed by an exponential return of the post-synaptic membrane to resting membrane potential.

The PSPs from an area of cortical pyramidal cells can be spatially summed to give the mean soma potential. The potentials from about 1 to 5 cm² of cortex can be recorded as the surface EEG signal. In practice only the alternating component of the mean soma potential can be captured in the EEG because of the problems associated with defining a suitable zero-voltage reference point which is reproducible, unambiguous, and insensitive to DC artifacts from muscle-voltages and other non-cortical sources.

The main question that this paper seeks to answer is: What features of the EEG signal best reflect the effectiveness of information flow in the underlying cerebral cortex? We have assumed that, in the

conscious state, there is accurate and unimpeded information flow between cortical neurons; whereas when a subject is unconscious (or anaesthetised), there is some impairment of this information flow.

It would seem plausible that the transition from consciousness to unconsciousness may be analogous to a phase transition from order to disorder. Thus, a graphic metaphor for the induction of general anaesthesia may be that of melting jelly. If the ordered/disordered phase transition is a realistic picture of the cortical alterations that occur with loss of consciousness, then there should exist an order parameter. There is theoretical justification to postulate that the mean soma potential is the appropriate order parameter for cerebral cortical function. There is also some indirect experimental evidence that is consistent with this hypothesis. It is known from recordings from individual neurons[2] and PET studies[1], that in the conscious state there is a high mean spike rate and a high neuronal metabolic rate. In contrast, in the anaesthetic state there is a low spike rate and a much reduced metabolic rate. If it were possible to measure the mean soma potential directly, we would be able to identify and track the conscious-to-unconscious transition. However, as stated above, it is not possible to extract the absolute value of the mean soma potential, but only its alternating component: in the form of the EEG signal. Therefore, it is necessary to develop an observable quantity that is highly correlated with the mean soma potential — to serve as an order-parameter by proxy.

For the time-varying EEG voltage signal $v(t)$, we postulate that the root mean-square (RMS) value of the first derivative of the EEG, $(dv/dt)_{\text{rms}}$, should be a good candidate, since this quantity, when squared, gives a measure of the energy contained in the rate-of-change of the EEG voltage. Intuitively, when many time-coincident PSP perturbations are summed, there will be a steep change in the EEG signal (i.e., large dv/dt) — unless there are equal numbers of inhibitory and excitatory perturbations. Because large numbers of activated neurons are required to produce large numbers of PSP perturbative events, we may posit some sort of positive correlation between $(dv/dt)_{\text{rms}}$ and the mean soma potential. The actual form of the function which relates these two quantities is unknown, so we have attempted to examine various data to see if we can uncover an empirical relationship.

In this paper we report on the use of the time-derivative statistic: firstly in cellular automaton (CA) simulations, and secondly, on EEG data from patients undergoing general anaesthesia.

2 Cellular Automaton Simulations

The CA model is a way of investigating, in an extremely simplified way, the behaviours of extended systems of interacting elements.

2.1 One-dimensional simulations

The initial investigations were designed to establish whether dv/dt , interpreted as ‘rate of state change’, has as any correlation with the complexity in the signal generated by a CA. On a small computer we set up various one-dimensional deterministic CAs ($k = 2$ states, range $r = 1$) consisting of 200 cells with periodic boundaries. The CAs were classified according to Wolfram[3]:

Class	Spatial Description	Pattern Evolution	Final State
1	spatially homogeneous	disappears	fixed
2	periodic structures	fixed finite size	regional
3	aperiodic	grows indefinitely	increasing size
4	propagating structures	grows and contracts	irregular

The essential distinction between classes is that in classes 1 and 2 information propagates a fixed distance, whereas in classes 3 and 4 information can propagate an infinite distance (at a fixed speed for class 3).

The simulations were coded in the C++ language. Each CA was allowed to evolve from a random initial pattern for a period of 200 time steps. We counted the number of cells which flipped state at




Figure 1: Box plots of the $(dv/dt)_{\text{abs}}$ for different classes of one-dimensional CA. The output from more complex CAs (classes 3 and 4) exhibit greatly increased $(dv/dt)_{\text{abs}}$ when compared with those from classes 1 and 2.

each step, and recorded the absolute value of cell ‘activity’ per time step:

$$\left(\frac{dv}{dt}\right)_{\text{abs}} = \frac{1}{N\Delta t} \sum_{i=1}^N |v_i - v_{i-1}|$$

where $N = 200$ steps, v_i is the number of active cells at time-step i , and $\Delta t = 1$ represents the cycle time.

As shown in fig. 1, the $(dv/dt)_{\text{abs}}$ values are much greater for classes 3 and 4 than for classes 1 and 2. Therefore, a broad conclusion would be that, for one-dimensional CAs, a high $(dv/dt)_{\text{abs}}$ value is associated with enhanced percolation of information across and through a system of interacting units. Conversely, a low $(dv/dt)_{\text{abs}}$ value would be associated with systems in which movement of information is constrained.

2.2 Two-dimensional simulation

This conclusion is further supported by the results from a 2-dimensional CA — which more closely simulates cortical function. Using a square lattice and a von Neumann neighbourhood, each cell can take on three states (quiescent, active, and refractory). At each time step an active cell can activate quiescent neighbours randomly according to a predefined ‘bond’ probability P_b . The active cell then moves to a refractory state. This model is very similar to the well-known ‘forest-fire’ model. The output ‘voltage’ v of the system is defined as the number of the active cells at each time step, thus v is analogous to the mean soma potential \bar{v} of an area of real cerebral cortex. It is generally accepted that the mean soma potential, relative to the unconscious state value, is the order parameter for the two-dimensional CA model.

The model was implemented in C++ using a 20×20 lattice. The array was continuously stimulated by activating a single cell at a randomly selected site at each time step. After allowing 200 time steps to eliminate any start-up transients, twenty runs of 200 time steps were performed and the average value for the $(dv/dt)_{\text{abs}}$ cell activity was recorded. This was repeated for a range of values for P_b (0.1 to 0.9).

Figure 2a shows an abrupt change (clear transition) in the mean activation \bar{v} of the system when P_b is reduced below about 0.5. Figure 2b demonstrates that this change in \bar{v} is closely mirrored by

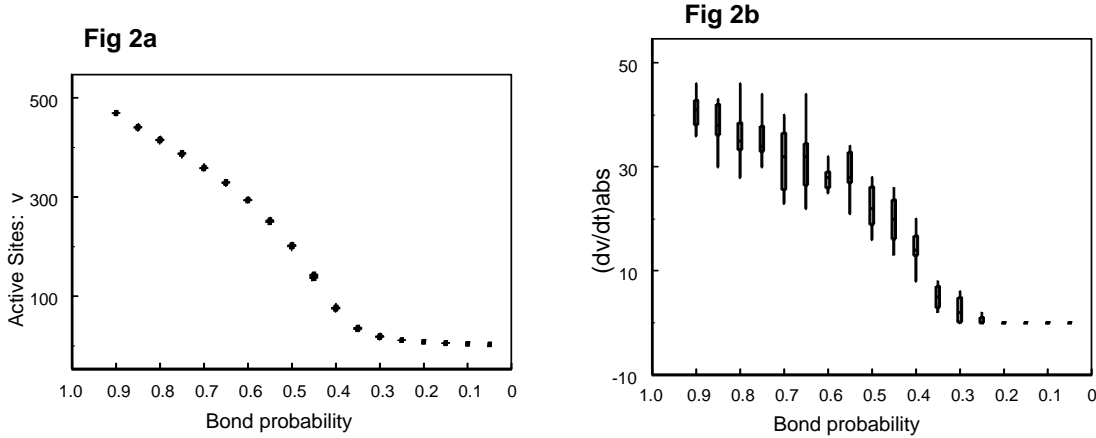


Figure 2: Results of the two-dimensional CA simulations at different values of bond probability (P_b). For each value of P_b , the box plot shows mean activation $v(t)$ (fig 2a), and rate of state change $(dv/dt)_{\text{abs}}$ (fig 2b). High values of P_b correspond to the awake state of the cortex, while low values correspond to the unconscious state. Both the $v(t)$ and $(dv/dt)_{\text{abs}}$ abruptly decrease as P_b decreases below a critical point.

a corresponding change in the $(dv/dt)_{\text{abs}}$ statistic. These data gave a good linear fit (correlation coefficient $r = 0.92$) between ‘rate of state change’ and ‘mean activation level’:

$$\left(\frac{dv}{dt}\right)_{\text{abs}} = 1.7 + 0.09\bar{v}$$

3 Changes in EEG during Induction of Anaesthesia

Obviously the details of the interactions of billions of neurons in a real cerebral cortex are vastly more complex than the 2-dimensional CA model, but to what extent has the essence been captured by the model? In order to try an answer this question we have analysed EEG data collected from 23 patients during induction of general anaesthesia.

The study was approved by the Waikato Ethics Committee and informed consent was obtained from adult patients prior to having minor surgery under general anaesthesia. The EEG data were collected using the ASPECT A-1000 monitor (Aspect Medical Systems, Natick, MA). While preparing for anaesthesia, five silver/silver chloride EEG pads (Zipprep, Aspect Medical Systems, USA) were attached to the patient’s forehead according to a standard montage: one as the ground (Fpz) and the other four as two separate bipolar channels (F7-T7, F8-T8) over the left and right prefrontal cortex. The low- and high-frequency filters were set to 0.45 Hz and 45 Hz respectively. Data from the EEG monitor were continuously transferred to an IBM computer for off-line analysis. The progression from the conscious state to a fully anaesthetised state was defined by noting the specific times of: (1) starting injection (start); (2) dropping a 50-ml syringe which was held between the thumb and fingertips (syringe drop); (3) insertion of laryngeal mask or endotracheal tube (tube); and (4) initial incision (incision).

As described above, the mean soma potential is not directly observable from the EEG record, so we attempt to extract a measure of the rate of cortical state change by computing the RMS value of the first time derivative of the EEG voltage. Since the EEG signal can be pictured as a sum of sinusoids, computing the derivative in the time domain is equivalent to frequency-scaling the amplitude of each frequency component in the Fourier domain. Thus $(dv/dt)_{\text{rms}}$ can either be computed directly from




Figure 3: An example of the abrupt changes in $(dv/dt)_{\text{rms}}$ during induction of anaesthesia.

the EEG time series,

$$\left(\frac{dv}{dt}\right)_{\text{rms}} = \left[\frac{1}{N-1} \sum_{i=1}^{N-1} \left(\frac{v_{i+1} - v_i}{\Delta t} \right)^2 \right]^{\frac{1}{2}}$$

or indirectly from the frequency-scaled power spectrum:

$$\left(\frac{dv}{dt}\right)_{\text{rms}} = \frac{1}{\sqrt{2}N} \left[\sum_{k=0}^{N/2} a_k^2 \omega_k^2 \right]^{\frac{1}{2}}$$

where $\omega_k = k(2\pi/N\Delta t)$ is the k th frequency harmonic, and a_k is the amplitude of that component. The RMS averages were computed for successive 1-sec epochs. The sampling rate was 256 samples/sec at a 14-bit resolution, giving $\Delta t = 1/256$ s, and $N = 256$. The two methods are equivalent, though the first (time-domain) method runs faster since no Fourier transformation is required. The second method is more convenient if analysis of specific frequency components is required.

Using these techniques our results showed a large decrease in $(dv/dt)_{\text{rms}}$ during the transition from the awake state to the anaesthetised state in almost all patients. An example is shown in figure 3. This transition reversed on awakening.

For the 23 patients, the mean value for $(dv/dt)_{\text{rms}}$ showed a significant decrease ($p = 0.007$, paired t -test) to 73% of the awake value at the point where the patient became asleep (syringe drop). This continued decreasing as the patient was intubated (tube) and up until the start of surgery (incision) by which time $(dv/dt)_{\text{rms}}$ had dropped to 30% of the awake value. Although these results show that there are statistically significant changes in the mean values for the whole group, there was marked variation in the patterns for some individuals. For example, it was common in the unconscious state for the $(dv/dt)_{\text{rms}}$ statistic to abruptly and momentarily shift from values close to zero to those values normally associated with being awake every second or so.

4 Discussion and Conclusions

We have shown that in simple two-dimensional CA models of cortical function, as the inter-cell bonding probability diminishes the average number of active cells \bar{v} decreases, as does the magnitude of their rate of state change, $(dv/dt)_{\text{abs}}$ (figs 2a,b). A similar decrease for the $(dv/dt)_{\text{rms}}$ statistic is found when applied to the EEG voltages for patients as they lose consciousness during the induction of anaesthesia. However, not surprisingly, the EEG data revealed more complicated patterns in some patients. These more complicated behaviours are difficult to interpret in terms of the simple ordered/disordered phase transition concept. In some cases there seems to be a form of intermittent ‘switching’ between the awake and unconscious phases. In a few cases there was very little change in the $(dv/dt)_{\text{rms}}$ during loss-of-consciousness. The significance of these aberrant examples is hard to interpret but the following possibilities should be considered:

- Was the signal corrupted by equipment noise or muscle-voltage (EMG) artifact?
- Is the $(dv/dt)_{\text{rms}}$ change purely an epiphenomenon?
- Were the patients never properly asleep?
- Was there very rapid switching between phases? Should RMS averaging be performed over shorter time scales?

These questions are the subject of ongoing research. We conclude that:

- RMS-averaged (dv/dt) shows some promise as an easily measured ‘proxy’ order indicator.
- The transition from the awake state to a state of general anaesthesia is associated with a decrease in $(dv/dt)_{\text{rms}}$. This decrease is qualitatively similar to that found in two-dimensional CA simulations. This change probably reflects the suppression of cortex ‘activation’ by anaesthetic agents and supports the notion that the transition from the awake state to that of general anaesthesia has features consistent with an ordered/disordered phase transition.

5 References

References

- [1] M.T. Alkire, R. J. Haier, S. J. Barker, N. K. Shah, J. C. Wu, and Y. J. Koa. Cerebral metabolism during propofol anesthesia in humans studied with positron emission tomography. *Anesthesiology*, 82:393–403, 1995.
- [2] J. Berg-Johnsen and I. A. Langmoen. Isoflurane hyperpolarises neurons in rat and human cerebral cortex. *Physiologica Scandinavica*, 130:679–685, 1987.
- [3] S. Wolfram. Universality and complexity in cellular automata. *Physica D*, 10:1–35, 1984.

Cellular Ecolab

Russell K. Standish
 High Performance Computing Support Unit
 University of New South Wales
 Sydney, 2052
 Australia
 R.Standish@unsw.edu.au
<http://parallel.hpc.unsw.edu.au/rks>

Abstract

In the original specification of the Ecolab model, spatial variation of the system's variables was included, with an additional *migration* operator introduced proportional to the divergence of the species density. Suitably discretized, the model then consists of a number of *cells*, each running the non-spatial Ecolab model, linked together by the migration operator. This paper reports on the implementation of such a cellular Ecolab.

Originally, it was hoped that computational parallelism might be exploited by providing parallel versions of the underlying array operations used. However, it was found that the cost of the irregular data communication involved in the sparse matrix multiplication destroyed any gain from parallelism. Thus a Cellular Ecolab offers a way of exploiting parallelism, that for certain situations as superlinear speedup is observed.

As an example of the type of problem addressable by a spatial Ecolab, the species-area rule is examined, although the results are not good enough at this stage to be definitive.

1 Introduction

In a previous paper[17], a model of an evolving Lotka-Volterra ecology called the *Ecolab model* was proposed, and has been studied in a number of subsequent papers[13, 12, 15, 14]. The defining equation (with slightly revised notation) is given by:

$$\dot{\mathbf{n}} = \mathbf{r} * \mathbf{n} + \mathbf{n} * \boldsymbol{\beta} \mathbf{n} + \boldsymbol{\gamma} * \nabla^2 \mathbf{n} + \boldsymbol{\mu}(\mathbf{r} * \mathbf{n}) \quad (1)$$

where \mathbf{n} is the species density, \mathbf{r} the effective reproduction rate (difference between the intrinsic birth and death rates in the absence of competition), $\boldsymbol{\beta}$ the matrix of interaction terms between species, $\boldsymbol{\gamma}$ the migration rate and $\boldsymbol{\mu}$ the mutation operator. All of these quantities (apart from $\boldsymbol{\beta}$, which is a matrix) are vectors of length n_{sp} , the number of species in the ecology. The operator $*$ denotes elementwise multiplication. The mutation operator returns a vector of dimensionality greater than n_{sp} , with the first n_{sp} elements set to zero — in effect expanding the dimensionality of the space, a key feature of this system. For a more detailed exposition of the various properties of the model, in particular, the precise form of the mutation operator, the reader is referred to the previous published papers, as well as the Ecolab Technical Report[16], which are all available from the Ecolab Web Site¹.

All previous studies of the Ecolab model have set $\boldsymbol{\gamma}=0$, and made the model independent of space. In this study, I implement Ecolab on a grid, to examine what effect migration has on evolutionary dynamics. Whilst it is conceptually possible for \mathbf{r} , $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$ and $\boldsymbol{\mu}$ to vary spatially, there is a problem when these values are generated through speciation. A new species comes into existence at one particular grid location. When an individual of that species migrates to a different grid location, what are its phenotypic parameters? Presumably completely random values or a well specified functional applicable to all species is incorrect. Here we might consider that different grid cells might have, for example, different topographic elevations. However, different species will respond in different ways to

¹<http://parallel.hpc.unsw.edu.au/rks/ecolab.html>

elevation. Until we have a solution to this problem, we have to study the simpler system where the phenotypic parameters are constant in space, but \mathbf{n} varies.

As an exercise of what sort of biological problem might be tackled by a spatial Ecolab model, the species-area relationship observed along an archipelago is an interesting one to investigate. The observed relationship[8] is one of a power law $n_{\text{sp}} \propto A^s$, yet there is no known theoretical reason why this should be. The importance of this rule is that it is often extrapolated to patchy land ecosystems, such as relating the size of a patch of rainforest to its diversity.

Already there is a complication in this model. May[9] proposed that random Lotka-Volterra webs would be unstable if

$$n_{\text{sp}} < \frac{1}{s^2 C} \quad (2)$$

where C is the connectivity, defined as the proportion of nonzero elements in β , and s is the interaction strength, defined as the standard deviation of the offdiagonal terms of β , divided by the average of the diagonal terms. Cohen and Newman[3] showed that May's criteria does not hold for Lotka-Volterra systems in general, only a smaller class related to the models May studied. However, the inverse relationship between species number and connectivity does appear to hold[11, 4, 2].

Although May considers stability of the ecosystem, whereas what is really important is the permanence of the system[5], the species number will depend on more factors than area. Keitt[7] examined the effect of "spatialising" random Lotka-Volterra systems on the stability criterion. He initialised an ensemble of Lotka-Volterra systems in an unstable configuration, and evolved the systems until no further extinctions were recorded (so in fact measuring a permanence criterion, rather than a stability criterion). He found that equation (2) fitted well the resulting distribution of end points for the non-spatial version. For a 100×100 grid, a generalised form of equation (2):

$$n_{\text{sp}} = \gamma s^\alpha C^\beta \quad (3)$$

with α , β and γ significantly different from the values of May's criterion (-2, -1 and 1 respectively) although the biggest difference was observed with the α values, which were substantially reduced.

So in this study, I again examine how the α , β and γ values from (3) vary with grid size. Without necessarily asserting the correctness of this, I also simulate systems at different grid sizes, starting from the same n_{sp} , s and C values to see what the total n_{sp} (the size of the union of species from all grid cells) will be as a function of area.

2 Implementation

The Ecolab modeling system represents the variables \mathbf{n} , \mathbf{r} etc. as dynamically allocated arrays `density`, `repro_rate` etc. As the system size varies, through the introduction of new species, or extinctions, new data storage is allocated for these variables, and the existing data is copied. To make the programming task simpler, these data structures are packaged into classes (called `array` and `iarray` (for integers)), that allow these objects to be manipulated in a *data parallel* fashion.

Each elementary operation of the model is coded as a TCL[10] command. Such commands include `generate`, which iterates

$$\dot{\mathbf{n}} = \mathbf{r} * \mathbf{n} + \mathbf{n} * \beta \mathbf{n},$$

`mutate` to add new species to the system, and the new `migrate` command, which handles migrations between cells. Additional commands are available to allow visualisation of what is going on in the model. To run a particular experiment, a TCL script is constructed, which set various parameters for the model, then executes an event loop, calling the above Ecolab commands.

Introducing a cellular structure to Ecolab required a major rewrite of the code. Each of the arrays `density`, `repro_rate` etc. now contain the entire multicellular ecosystem, packed in cell order. The `nsp` variable now becomes an array of values, containing the number of species in each cell, and can be used to determine what values belong to which cell. This allows the `generate` step, which iterates to be unchanged.

In order to keep track of the model's structure in a variety of situations, the model's variables (the \mathbf{n} , \mathbf{r} etc) are packaged up in a C++ class called a `global`. The main Ecolab model variables can be accessed by means of references which point into the global `global`'s (called `global_vars`) data structure. A `global` does not have many properties, but can be converted into a contiguous machine independent binary representation that may be saved to a file for checkpoint/restart, sent over the network to a client Ecolab program or sent to another processor in the parallel version of the code. One can also extract the variable corresponding to an Ecolab model variable: for example `g[density]` returns the iarray component of `g` that corresponds to the global `density` variable. So `global_vars[density]` is in fact synonymous with `density`. Coupled with functions that extract or insert globals corresponding to particular cell ranges (or even a single cell), this allows an easy way of accessing the model viewed as a single cell. The `mutate` and `migrate` algorithms consist of loops of the cells, with the code then working at the cell view level.

Parallelism also turns out to be relatively easy to implement using the MPI message passing interface². Each thread runs a separate copy of the Ecolab executable, which is a TCL interpreter. The main Ecolab script is executed in the master Ecolab process, the slave threads execute individual TCL commands as instructed by the master thread. Mostly, this is handled by the `PARALLEL` directive, which means run this TCL command on all processors simultaneously. Migration is handled by a template type operation — boundary cells on the processors are shifted circularly to the neighbouring processor, which are then used by the `get_neighbour` routine. `get_neighbour` returns a `global` containing the neighbour's cell, which can then be used for computing the Laplacian.

It should be pointed out that it is not necessary to have a parallel computer in order to run the code. All MPI calls are wrapped in `#ifdef MPI`, so one can selectively link to the MPI library. The algorithm treats a non-MPI code as though it is running on a single processor. With the MPI calls included, Ecolab typically runs about 10% slower on a single processor than when the calls are not included.

An important assertion that the code must satisfy is for the total number of individuals to be conserved by migration ($\sum_i n_i$ is constant). Since the migration operator returns a real valued vector, it needs to be probabilistically converted to the integer valued \mathbf{n} . One must ensure that the same dice roll is used for updating a cell and its neighbour across a cell boundary. To do this, an array of random numbers corresponding to cell boundaries is generated each time `migrate` is called, and distributed to all the processors.

Figure 1 shows the speedup over a single processor for a 30×40 grid for a typical Ecolab run. Note that superlinear speedup is obtained for the $n_p = 2$ and $n_p = 4$ cases. It is not actually clear why the $n_p = 3$ case performs so poorly.

It should also be noted that the Ecolab cellular structure is more general than geographical representation. In an economics model, the cells may correspond to separate sectors of a multisectoral model. Provided that the model has a slow and a fast time scale, with the fast time scale only operating within a cell, this cellular structure can be effectively used with parallel computers.

3 Simulations

The Ecolab model was run with different initial values of n_{sp} ($\{10, 20, \dots, 120\}$), different initial values of the connectivity ($\{2/n_{sp}, 4/n_{sp}, \dots, \lfloor 3n_{sp}/4 \rfloor / n_{sp}\}$) and interaction strength (the diagonal components of β was initialised by a uniform random variate between -1×10^{-3} and -5×10^{-4} , and the offdiagonal components to a uniform random variate between $-s \times 10^{-4}$ to $s \times 10^{-4}$, where s took on the values $\{1, 3, \dots, 9\}$ — *i.e.* interaction strength between 0.1 and 1.0). The other parameters of the system were kept fixed: \mathbf{n} was initialised to 100, \mathbf{r} initialised to a uniform random variate between -0.005 and 0.01 , μ (mutation rate) uniformly between 0 and 0.1, `sp_sep` 0.1 and γ (migration) set to 0.01. Each set of initial parameters was run as a separate thread for 10^5 timesteps with controls provided to limit the number of threads scheduled at any one time.

²<http://www.mcs.anl.gov/mpi/index.html>




Figure 1: Speedup curve for Ecolab running on a 30×40 grid with migration

This range of parameters reasonably filled the unit cube in (n_{sp}, C, s) space. The first 5000 timesteps were discarded to allow the system to settle down to its long time behaviour, then the values of (n_{sp}, C, s) were recorded each 1000 timesteps. Figures 2 to 4 show these data points, along with a minimal surface that is fitted to contain these data points. The hypertext version of this document contains links to VRML³ of these figures which allow the reader to interactively manipulate the figures.

The surfaces fitted are given by equation (3). The surface itself is fitted by minimising the sum of the volume underneath the curve and the sum of the distances of all points lying above the surface to the surface itself (the *data excess*). Clearly, there is a choice as to how much weight to give to minimising the volume and the minimising the data excess. By visualising the system using AVS⁴, it was found that weighting the data excess relative to the volume in the range 0.5 to 5 gave visually good fits, outside this range, the fit appeared quite poor. One could also compare the numerically fitted values against the values predicted by May's criterion 2 for the single cell case. In all the quoted values for α , β and γ in table 1, even weight was given to minimising the surface and the data.

The surface given by (3) is hyperbolic, and the volume bounded by it and the axes planes is infinite. In this work, I truncate the surface at some maximum value n_{max} . Figure 5 shows the connectivity-strength plane. In region A, the surface is below n_{max} , and in regions B and C, the surface is clamped to the value n_{max} . The locus D is given by $\gamma s^\alpha C^\beta = n_{max}$. The volume under the surface can then be found by integration, by adding the contributions from the 3 areas A, B and C:

$$\begin{aligned}
 \text{Vol.} &= \gamma \int_{\left(\frac{n_{max}}{\gamma}\right)^{\frac{1}{\beta}}}^1 C^\beta \int_{\left(\frac{n_{max}}{\gamma C^\beta}\right)^{\frac{1}{\alpha}}}^1 s^\alpha ds dC \\
 &+ n_{max} \int_{\left(\frac{n_{max}}{\gamma}\right)^{\frac{1}{\beta}}}^1 \left(\frac{n_{max}}{\gamma C^\beta}\right)^{\frac{1}{\alpha}} dC + n_{max} \left(\frac{n_{max}}{\gamma}\right)^{\frac{1}{\beta}}
 \end{aligned} \tag{4}$$

³<http://www.vrml.org>

⁴<http://www.avs.com>




Figure 2: Plot of sample data points in (n_{sp}, C, s) space for the 1×1 grid, with the fitted surface given by (2)




Figure 3: Plot of sample data points in (n_{sp}, C, s) space for the 8×8 grid, with the fitted surface given by (3), with the α , β and γ values shown in table 1

Grid Size	α	β	γ
1×1	-2.7	-0.98	1.0
2×2	-1.34	-1.99	0.5
4×4	-1.68	-1.53	0.5
8×8	-1.16	-1.2	0.88
16×16	-1.5	-0.85	1.01

Table 1: Table of α , β and γ values for different grid sizes

This formula was then integrated symbolically using GNUCalc⁵ and converted into C code. Care was taken of the special cases (when $\alpha = 1$ or $\beta = 1$) as in this case the symbolic integration involved logarithms, as opposed to powers.

The optimisation was performed using the Hooke algorithm[6, 1], and the results for the different grid sizes are listed in table 1.

The confidence in the values presented in table 1 was tested by visualising the data (Figures 2...4) and varying the α , β and γ to see how well the curve fitted the boundary of the observed data. It was found that β and γ were tightly constrained, with errors of no more than about 10%, whereas the α had a much larger error perhaps of order 40–50%.

Direct comparison with Keitt's work is difficult, as he worked with a 100×100 grid, and used a very different migration scheme. He reported a substantially lowered value of α (order 0.2), whereas β and γ had similar values to that of eq (2). It can be seen from table 1 that as grid sizes increase, the β and γ are tending back to May's values, whereas α is decreasing. Unfortunately, the confidence in the measured values is not that great, circumscribing these conclusions. At the present state of the art, the values are not well enough known in order to predict what the diversity (n_{sp}) does as a function of grid size, for constant C and s .

As an alternative, a simulation was performed by setting C and s to 1, and looking at the average

⁵<http://www.gnu.org/order/ftp.html>




Figure 4: Plot of sample data points in (n_{sp}, C, s) space for the 16×16 grid, with the fitted surface given by (3), with the α , β and γ values shown in table 1




Figure 5: Connectivity-Strength Plane. In region A, Keitt’s formula is less than n_{max} , in region B and C, the surface is clamped to n_{max} . The locus D is given by $\gamma s^\alpha C^\beta = n_{max}$.

diversity values for the entire grid over 50000 timesteps ignoring the first 5000, as a function of grid size. It is expected the the diversity should increase, as more niches open up with the greater number of cells. The result of this simulation is shown in figure 6. This was seen, but with a flat spot at intermediate grid sizes. Figure 7 shows the average connectivity and strengths values on the same run. Curiously enough, the curves are correlated with the diversity curve, not anti-correlated as might be expected from the hyperbolic nature of eq. (3).

4 Conclusion

This paper describes improvements made to the Ecolab code that allow for geographic variation in the evolutionary model. This improvement also allows for parallel processing to be exploited for the first time. An initial foray into examining the relationship between diversity and habitat area is described, but demonstrates that the issue is considerably more subtle than might initially be conceived. There are a number of theoretical issues that need to be resolved before a definitive answer is possible.

The migration algorithm used assumes that the phenotypic parameters of a species do not vary in space. This is an unfortunate restriction, and methods for dealing with propagating novel phenotypic information through space need to be developed.

5 Acknowledgments

The author wishes to thank the New South Wales Centre for Parallel Computing for the use of their Power Challenge Facility for running this code, and Sydney Vislab, in particular Ben Simons, for the use of AVS running on Silicon Graphics workstations, without which the visualisation of the surface fitting would have been impossible.




Figure 6: Diversity as a function of grid size




Figure 7: Connectivity and Strength as a function of grid size

References

- [1] Jr A.F. Kaupé. Algorithm 178: Direct search. *Comm. ACM*, 6:313, 1963.
- [2] J. E. Cohen, T. Luczac, C. M. Newman, and Z.-M. Zhou. Stochastic structure and nonlinear dynamics of food webs. *Proc. R. Soc. Lond. B*, 240:607–627, 1990.
- [3] J. E. Cohen and C. M. Newman. When will a large complex system be stable? *J. Theo. Bio.*, 113:153–156, 1985.
- [4] J. E. Cohen and C. M. Newman. Dynamic basis of food web organisation. *Ecology*, 1988.
- [5] J. Hofbauer, V. Hutson, and W. Jansen. Coexistence for systems governed by difference equations of Lotka-Volterra type. *J. Math. Biol.*, 25:553–570, 1987.
- [6] R. Hooke and T. A. Jeeves. Direct search solution of numerical and statistical problems. *J. ACM*, 8:212–229, 1961.
- [7] T. H. Keitt. Stability and complexity on a lattice: Coexistence of species in an individual-based food web model. Working Paper 97-02-016, Santa Fe Institute, 1997. <http://www.santafe.edu/~tkeitt/papers/foodweb/paper.html>.
- [8] R. H. MacArthur and E. O. Wilson. *The Theory of Island Biogeography*. Princeton UP, Princeton, 1967.
- [9] R. M. May. Will a large complex system be stable. *Nature*, 238:413–414, 1972.
- [10] J. K. Ousterhout. *Tcl and the Tk Toolkit*. Addison-Wesley, 1994.
- [11] S. L. Pimm. *Food Webs*. Chapman and Hall, London, 1982.
- [12] R. K. Standish. Ecolab: Where to now? In R. Stocker, H. Jelinek, B. Durnota, and T. Bossomeier, editors, *Complex Systems: From Local Interaction to Global Phenomena*, pages 263–271. IOS, Amsterdam, 1996. also *Complexity International*, vol. 3, <http://www.csu.edu.au/ci>.
- [13] R. K. Standish. The role of innovation within economics. In C. Chiarella, S. Keen, R. Marks, and H. Schnabl, editors, *Proceedings of the inaugural conference on Commerce, Complexity and Evolution*. UNSW, Feb 1996.

- [14] R. K. Standish. The role of innovation within economics. In W. Barnett, C. Chiarella, S. Keen, R. Marks, and H. Schnabl, editors, *Commerce, Complexity and Evolution*, volume 11 of *International Symposia in Economic Theory and Econometrics*. Cambridge UP, 1998.
- [15] R. K. Standish. Statistics of certain models of evolution. *Phys. Rev. E*, 1998.
- [16] Russell K. Standish. Ecolab documentation. Available at <http://parallel.acsu.unsw.edu.au/rks/ecolab.html>.
- [17] Russell K. Standish. Population models with random embryologies as a paradigm for evolution. In Russel J. Stonier and Xing Huo Yu, editors, *Complex Systems: Mechanism of Adaption*. IOS Press, Amsterdam, 1994. also *Complexity International*, vol. 2, <http://www.csu.edu.au/ci>.

Part III

Mathematical Tools

Effective Computation of 2D Coupled Map Lattices

Jacques Blanc-Talon
 CTA/GIP,
 16 bis, av. Prieur de la Côte d'Or, 94114 Arcueil, France
 blanc@etca.fr
<http://www.etca.fr/Users/Jacques.Blanc-Talon>

Abstract

Approximate identification of coupled map lattices is considered. The local dynamics is split into a local function, expanded in Hermite's polynomial series, and a coupling one which is the convolution product of the neighborhoods by a kernel. The local function fits the data set (an image) while the convolution kernel of the coupling function is adjusted to fit a set of some selected patterns, gathering the so-called "structural information". The structural information to be considered is shown to be the connected set of the zero-crossings of the Laplacian of the input image, which is computed by using Gauss kernels. A practical computation of a CML from stone patterns is shown.

1 Identification of CA and CMLs

Research on cellular automata (CA) is about a couple of decades old as it started, at least officially, approximately at the date of Wolfram's well-known papers [22]. Compared to the analytical study of global equations, cellular automata provide physicists with a mere bottom-up approach: global behaviors are generated from local rules and interactions over bounded domains. Generic CA have been proposed in various fields (see collective works [3, 8]) and the number of new related applications is still growing. The key point of this approach is that cellular automata can be efficiently implemented and tested: it is therefore an experimental approach in the best meaning of the word.

In brief, a cellular automaton is a discrete spatially-extended dynamical system: it is discrete in time, space and state space. However, this last feature could turn to be a stumbling block when matching with real experiments. Coupled map lattices (CMLs) were introduced by Kaneko [10] as a paradigm for the study of turbulence, convection and other similar problems occurring in physics. They may be considered as a generalization of CA since they are discrete in time and space but continuous in state space. Coupled map lattices are well-suited to the study of patterns growing under the action of repulsive and attractive forces [20], also called a reaction-diffusion process [16] and remain at the same time mathematically tractable. Despite the increase of their practical generality, only little attention has been paid to the inverse problem, namely the approximation of a data set; at the same time, the only works relevant to CA identification are probably [1, 8].

Actually, most of the work being available on CMLs is devoted to their analysis either by using statistical techniques [2] or by means of formal language theory [8, 9, 12]. One reason for this state of things is that both CMLs and CA are highly parameter sensitive which makes one wonder about the possibility of solving the inverse problem. Actually, any valid approximation scheme must answer a few fundamental questions, in the reverse order of importance:

1. How stable is the approximation with respect to the parameter tuning?
2. How stable is the approximation with respect to the noise of the data set?
3. Is there a one-to-one mapping between the expected approximation and the data set?
4. What information must the approximation converge to?

The first question is relevant to the well-known parameter sensitivity of coupled systems. One common technique for coping with it is to control the CML Jacobian, which is discussed in section 3: the system behaves in a chaotic manner as long as it does not approach a fixed point. This chaotic behavior is an advantage in applications where tuning by hand of the parameters must be forbidden, as in data encryption applications. Langton's conjecture, which assumes there is a critical parameter turning on the chaotic behavior [23] has been criticized in many papers [6].

Some hints for answering the second question can be found in Roy and Amritkar's paper [18]. The authors show that noise does not play a detrimental role by destroying small patterns but, on the contrary, allows formation of new structures in a mode they call stochastic resonance.

As we shall see below, the initial purpose of this study is certainly not to explain any physical process but on the converse to force a system to reach a point somewhere in its evolution. Though no straightforward relationship can be exhibited, the method still performs CML identification since the system is computed according to structural information; moreover, it is not constrained to accept the image as a stable point which leaves it full ability to evolve in different directions.

The information being actually caught by the CML has been called "structural information". It consists of the edges of the initial image, which lie in the connected set of the zero-crossings of the Laplacian and are computed by using Gauss kernels. This point is discussed in section 4.

The practical problem being addressed in this paper is to compute a general (deterministic) coupled map lattice which can yield a particular array of values, whose integer truncation can be displayed as an image. Possible applications fall within image compression, texture analysis and cryptography. As for CA, interaction domains to be considered are restricted to small neighborhoods around cells because physically meaningful interactions are limited in range. In the present application, this choice is still valid since the luminance of a given pixel depends almost only on its neighborhood. Since no additional information is assumed to be available about the initial array, a Gaussian noise over the array may be a good and plausible start according to the particular definition of the structural information.

The evolution equation is the sum of a local (site) function and a coupling function. About the site function, instead of taking the logistic map which has been extensively studied either in chains [5] or in arrays (e.g. [11] for the analysis of spatial structures in population dynamics) or, again, any other polynomial function of degree in \mathbb{N} or \mathbb{Q} [7], we shall try to remain as general as possible in keeping a power series within a convergence domain. Our approach which makes use of image processing techniques is rather innovative since the CML is computed straight from the image instead of being approximated by a CA like in [4]; moreover, there is only one CML for one or several images and its parameters do not evolve as in [14].

This paper is organized as follows. Section 2 introduces the theoretical work, a complete presentation of the method being given in section 3. In section 4, a new definition of CML structural information is given and discussed. Experiments from real images are shown in section 5, followed by the conclusion.

2 CML evolution equation revisited

In the following, a CML is 4-tuple

$$\mathcal{A} = (N, S, L, C, r)$$

where N is the dimension of the site array, S is the set of sites indexed by $\{1, 2, \dots, Q\}^N$ and L is the local transition function: $\mathbb{R}^2 \rightarrow \mathbb{R}^2$ "attached" to every site. r is the diameter of ball B (according to metric d in \mathbb{N}^N – see remark below) whose interior can interact with the behavior of the site at its center; C is the coupling function which governs that interaction, and thus the interaction of the whole ball on site x with respect to some measure μ is given by:

$$\int_{B(x,r)} C(\|x - y\|) dy.$$

On discrete lattices, different tessellations and related metrics can be used which lead to different topological properties. In the following, we consider only linear coupling functions leading to the evolution equation at site j :

$$y_j(k+1) = L(y_j(k)) + \beta \sum_{\substack{l \neq j \\ d(l,j) \leq r}} y_l(k).$$

Notwithstanding such an apparent naive definition, this system is able to capture very rich and complex behaviors. When d is the ‘‘city-block’’ (or ‘‘Manhattan’’) distance, an interesting remark is that the finite sum in the right member of previous equation can be rewritten as:

$$\sum_{l=l_0-L}^{l_0+L} \sum_{m=m_0-M}^{m_0+M} y_{l,m}(k) - y_{l_0,m_0} = \sum_{l=-L}^L \sum_{m=-M}^M \omega_{l,m} y_{l_0+l,m_0+m}(k)$$

with $\omega = 1$ for every $(l, m) \neq (0, 0)$ and $\omega = 0$ otherwise in the related box. This is the well-known formula of the convolution product $\omega \star y$ of a signal y by a kernel ω which leads us to consider the interaction of the neighborhood on the site as the one of a linear (or not!) system with a given transfer function. For sake of simplicity, we restrict now our attention to linear system with $M = L = 2r + 1$. What can we do about the local function?

Provided L is \mathcal{C}^0 and of finite energy with respect to the scalar product:

$$\langle f, g \rangle = \frac{1}{2^n n! \sqrt{\pi}} \int_{-\infty}^{\infty} f(x)g(x)e^{-x^2} dx$$

that is $\int_{-\infty}^{\infty} e^{-t^2} L^2(t)dt < \infty$, it can be expanded in Hermite’s polynomial series:

$$L(x) = \sum_{k=0}^{\infty} \alpha_k H_k(x) \text{ with } \alpha_k = \frac{1}{2^k k! \sqrt{\pi}} \int_{-\infty}^{\infty} H_k(x)L(x)dx. \tag{1}$$

The system is thus governed by the equation:

$$y(k+1) = \sum_{i=0}^{\infty} \alpha_i H_i(y(k)) + \omega \star y(k). \tag{2}$$

In the following, expansion is stopped at rank p . As we shall see, this particular choice of Hermite’s expansion of the local function is motivated by the definition and the practical computation of the structural information.

3 Computation of kernel and site functions

Let us consider an image \mathcal{I} of some process showing evidence of spatial structures we would like to model. We assume we don’t have any additional information about this image, that we know neither the image which initiated the process, nor whether the observed structures are stable or unstable, nor if the grid of the image is the support of the lattice. Reminding that the general inverse problem for CA is NP-hard, the CML inverse problem without any hints seems hopeless and it is not surprising that only a tiny effort has been devoted to it.

Instead of the rigorous mathematical problem of finding the right CML which generates image \mathcal{I} (as a stable point, if it has any), we can wonder:

Is there any CML which *may* generate a *close* image?

Of course, the proximity of the original image and the computed one must be expressed in terms of some structural distance, provided the favored structures gather the information we are interested in. Rewritten like this, the exact inverse problem becomes an approximation problem, which can be solved in several ways under the right assumptions. Our attempt, which hopes to remain as general as possible, is the following.

Square images are considered instead of rectangular ones for the sake of simplicity, without loss of generality. Let $\mathcal{I} = (g_{i,j})_{1 \leq i,j \leq P}$ be the original \mathbb{R} -valued image of size $P \times P$. The estimated computed image $\hat{\mathcal{I}}$ is given by the general equation 2 which can be rewritten as:

$$\hat{g}_{i,j} = \sum_{k=0}^p \alpha_k H_k(g_{i,j}) + \sum_{k=-r}^r \sum_{l=-r}^r [1 - \delta_{k,l}] \omega_{k,l} g_{i+k,j+l}$$

(δ is the Kronecker delta), or, with a little linear algebra:

$$\hat{\mathbf{g}} = L(\alpha, \mathbf{g}) + T(\omega) \mathbf{g} \quad (3)$$

or again, in explicit form:

$$\begin{pmatrix} \hat{g}_{i-r,j-r} \\ \dots \\ \hat{g}_{i,j} \\ \dots \\ \hat{g}_{i+r,j+r} \end{pmatrix} = \begin{pmatrix} \sum_{k=0}^p \alpha_k H_k(g_{i-r,j-r}) + \sum_{\substack{-r \leq i', j' \leq r \\ i' < 0 \text{ or } j' < 0}} g_{i+i',j+j'} \omega_{i',j'} \\ \dots \\ \sum_{k=0}^p \alpha_k H_k(g_{i,j}) \\ \dots \\ \sum_{k=0}^p \alpha_k H_k(g_{i+r,j+r}) + \sum_{\substack{-r \leq i', j' \leq r \\ i' < 0 \text{ or } j' < 0}} g_{i+i',j+j'} \omega_{i',j'} \end{pmatrix} + \begin{pmatrix} \omega_{0,0} & \dots & \omega_{r,r} & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ \omega_{-r,-r} & \dots & \omega_{0,0} & \dots & \omega_{r,r} \\ \dots & \dots & \dots & \dots & \dots \\ 0 & \dots & \omega_{-r,-r} & \dots & \omega_{0,0} \end{pmatrix} \begin{pmatrix} g_{i-r,j-r} \\ \dots \\ g_{i,j} \\ \dots \\ g_{i+r,j+r} \end{pmatrix}. \quad (4)$$

In equation 3, α and ω are vectors with entries $(\alpha_0, \dots, \alpha_p)$ and $(\omega_{-r,-r}, \dots, \omega_{0,0}, \dots, \omega_{r,r})$; please notice that $\omega_{0,0}$ may be different from 0. The vector $\mathbf{g} = (g_{i-r,j-r}, g_{i-r,j-r+1}, \dots, g_{i+r,j+r})$ is the input of the equation. \mathcal{N} is a Gaussian noise modelling the influence of the image outside of the rectangular box, its mean and variance are $\mu = E[1^T \mathbf{g}]$ and $\sigma = E[\mathbf{g}^T \mathbf{g}]$ which are easily found by direct computations over the set of neighborhoods. Thus, $T(\omega)$ is a Toeplitz matrix which shows interesting theoretical properties.

If the local function (*i.e.* the $\alpha_{i,j}$'s) were known, the coupling function would be given by the inversion of a large sparse matrix of size $Q^N \times Q^N$ or, equivalently, by the inverse of the mean Ω of all the matrices of size $(2r+1) \times (2r+1)$ over the image. The first method is called deconvolution and generally approached by the second one which yields a system like:

$$\omega = (\Omega(\alpha_0, \alpha_1, \dots, \alpha_p))^{-1} \mathbf{V}(\alpha_0, \alpha_1, \dots, \alpha_p) \quad (5)$$

or, in a less aesthetic form:

$$\begin{aligned} \omega &= \left(\sum_{(X(k), Y(k)) \in \mathcal{T}(\mathcal{I})} \sum_{i=r+1}^{P-r-1} \sum_{j=r+1}^{P-r-1} g_{X(k), Y(k)} g_{X(l), Y(l)} \right)^{(-1)} \\ &\quad \times \left(\sum_{(X(k), Y(k)) \in \mathcal{T}(\mathcal{I})} \sum_{i=r+1}^{P-r-1} \sum_{j=r+1}^{P-r-1} g_{X(k), Y(k)} \widehat{g}_{i,j} \right), \end{aligned} \quad (6)$$

with special indices $X(k) = i + r - \lfloor k/(2r+1) \rfloor$ and $Y(k) = j + r - (k \bmod r)$ ($\lfloor x \rfloor$ is the floor integer value of the real x). The sums are restricted to a set $\mathcal{T}(\mathcal{I})$ we shall discuss in the next section. Our problem is thus equivalent to the coupled subproblems:

$$\text{Minimize } \|\widehat{\mathcal{I}} - \mathcal{I}\| \iff \begin{cases} \omega = (\Omega(\alpha_0, \alpha_1, \dots, \alpha_p))^{-1} \mathbf{V}(\alpha_0, \alpha_1, \dots, \alpha_p), \\ \text{Minimize } p, \\ \frac{\partial}{\partial \alpha_k} \sum_{\mathcal{I}} (\widehat{g}_{i,j} - g_{i,j})^2 = 0 \quad 1 \leq k \leq p. \end{cases} \quad (7)$$

We haven't taken up the problem of finding the best p but the third subproblem is solved easily by minimizing the quadratic error:

$$\frac{\partial}{\partial \alpha_k} \sum_{\mathcal{I}} (\widehat{g}_{i,j} - g_{i,j})^2 = 0 \Rightarrow \sum_{\mathcal{I}} \left(\sum_0^p \alpha_k H_k(g) + (T(\omega) - 1)\mathbf{g} \right) H_p(g) = 0. \quad (8)$$

Inverting the summation order yields a linear equation $H\alpha = E(\omega)$ with H being the sum over the whole image \mathcal{I} of the Hermite's polynomials.

Proposition 9.1 *The solution of eq. (2) is approximated by the system:*

$$\alpha = H^{-1}E(\omega) \text{ and } \omega = \Omega^{-1}(\alpha)V(\alpha).$$

A standard Conjugate Gradient Method [21] has been used for minimizing this system; due to polynomials properties, the Hessian matrix used in the algorithm has again a special form. The starting point of the procedure is given as $\alpha_0 = \alpha_1 = 1, \alpha_k = 0, k \geq 2$ and ω as the Gaussian kernel from equation 9 below.

Since the behavior of the CML is assumed to admit the input image \mathcal{I} as a ‘‘rather stable’’ point, the Jacobian of the system over a small part of the image:

$$\left| \frac{\partial \widehat{g}(\alpha_1, \alpha_2, \dots, \alpha_n)}{\partial (\alpha_1, \alpha_2, \dots, \alpha_n)} \right|_{g_1, g_2, \dots, g_m}$$

has to be at most zero, at least *minimum in magnitude*. Its computation is a little tricky since it makes use of Hermite's polynomials recursion property and a discussion about it is out of the scope of this paper. On the other hand, the structural information of the image must remain invariant, *i.e.* under the action of the Gauss operator. This nuance on the meaning of a stable point should lead to a different analysis of stability.

4 A new definition of structural information

The finite sums in 7 are performed on a set $\mathcal{T}(\mathcal{I})$ we haven't explicitly defined yet: this is the set of connected structures within \mathcal{I} . According to our initial hypothesis, it collects the essential structural information needed to compute the CML. In [19], Roy and Amritkar define a structure as a region of space such that the difference in the values of close sites within this region is less than a given threshold. As such a simple definition could not be consistent with more advanced sophisticated structural models, we define a structure to be the curve on which the Laplacian of the image is zero.

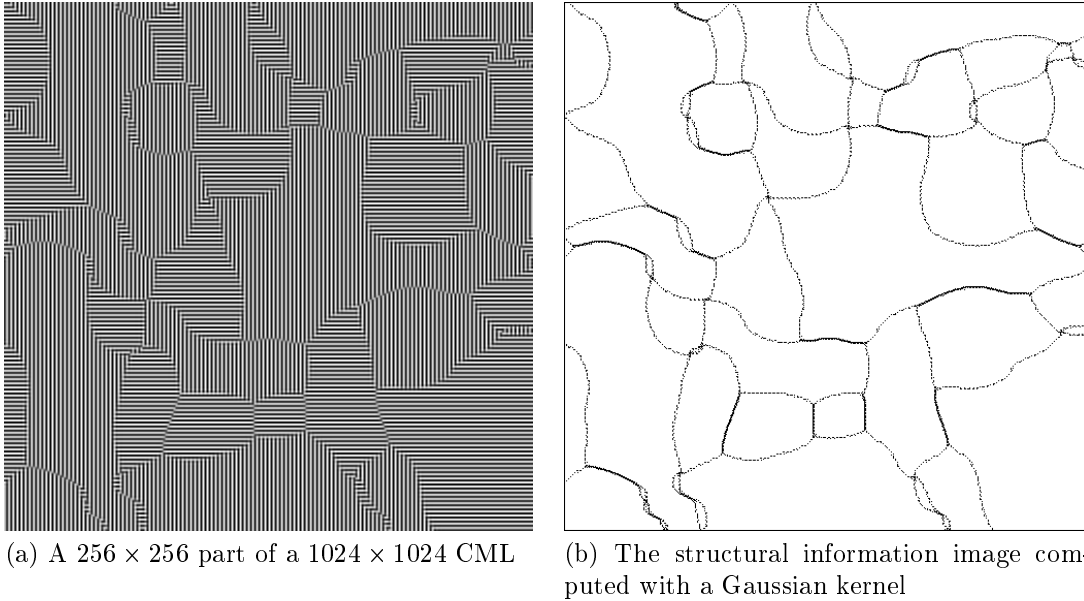


Figure 1: example of structures extraction.

First, the notion of structure has to be defined clearly; second, an explicit numerical method has to be found. A preliminary remark is that we consider the array of values of the CML as an image, in the common understanding of it. In that case, a structure can be defined according to the Mumford-Shah model which seeks simultaneously for a piecewise smoothed image $F(\mathcal{I})$ with a set γ of abrupt discontinuities called the edges. It can be shown [15] that the best edge detector is the system minimizing the following functional with respect to F :

$$\Gamma(F, \gamma) = \int_{\mathcal{I} \setminus \gamma} (|\nabla F|^2 + (F - \mathcal{I})^2) d\mu + \text{length}(\gamma).$$

The first part means that the smoothed image is actually smooth outside of the edge set, the second that it is a good approximation of the initial image and the third that we look for a minimal set, discarding trivial solutions. In many ways, minimization of this functional is linked to solving the heat equation a favored operator of which is the Gauss function. As a conclusion to a theoretical analysis of these arguments, we define the structural information in a CML as:

Definition 9.1 *The β -structural information set of a CML is the connected set of the zero-crossing of the Laplacian in the range $[-\beta, \beta]$.*

Non-connected edges are not taken into account since they vanish after a few iterations.

Numerous methods (may be a thousand!) exist in image processing for performing edge segmentation and edge connection. We have naturally been led to follow Marr and Hildreth's approach [13] which consists in smoothing \mathcal{I} by a Gaussian kernel, then in applying the Laplacian of the Gauss function, numerically computed by the difference of two Gaussians (DOG operator). For instance, figure 1 is a subimage of 20,000 iterations of $0.389H_1(g) - 0.016H_3(g) + Tg$ on an initial random image where T is the first approximation of a Gaussian kernel, i.e.

$$T = \begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}. \quad (9)$$

As in the definition of the evolution equation 2, the Gauss function plays a crucial role: more precisely, the structural information of the CML is mainly captured by the coupling function.

5 Experiments



Figure 2: natural rocks formations in Svalbard.

While honeymooning in Spitzberg in 1997, I arrived near New London (a smiling place which counted about 7 inhabitants at his best in the last century, when they were trying to extract marble from the stones under the ice) in a flat field covered with surprising formations, occurring in closed curves of medium size rocks put on a layer of smaller stones (see figure 2). Amazed, I took a picture of it and, back to the lab, tried to play with a simple CA simulation software so as to generate similar patterns. Of course I succeeded in doing something but was unable to know how valid my experiments were since I had no quality criterion.

Before applying the present identification process to the image, some preprocessing was made. First, the interesting part of the image showing the patterns was selected and mapped to a reference plane, after suppression of geometrical effects due to both perspective and camera optics (fig. 3.c); please note that this operation yields a triangular image. The structural information was then extracted from a rectangular part of the bottom of the image (the more accurate part); figure 3.d shows the structural information superimposed to the initial image.

For the structural information being computed, the computed solution is

$$\omega = \begin{bmatrix} 0.1, -0.1, 0.1 \\ -0.1, -0.2, -0.1 \\ 0.1, -0.1, 0.1 \end{bmatrix}, \quad \alpha = \begin{bmatrix} 9. \times 10^{-7} \\ 0.24 \\ 1.1 \times 10^{-6} \\ -0.17 \\ 10^{-8} \\ 9.22 \times 10^{-4} \end{bmatrix}.$$




Figure 3: preprocessing of the initial image.

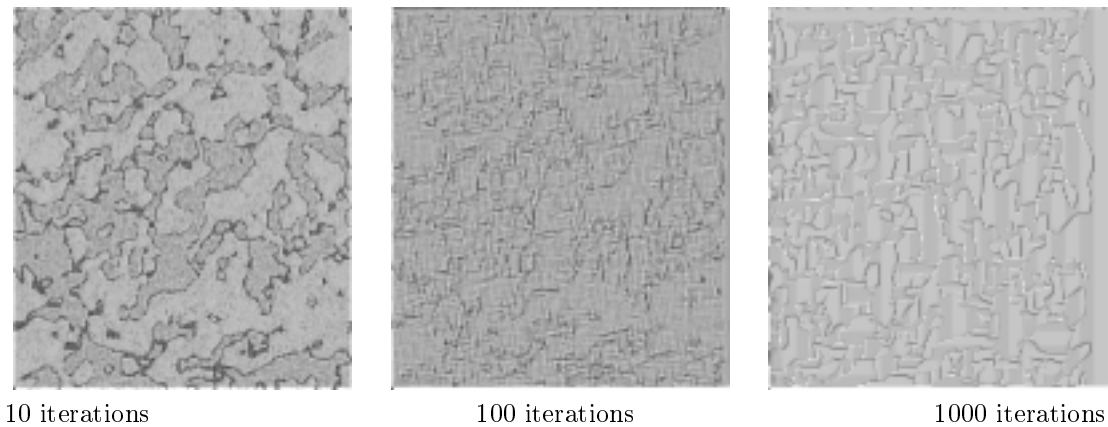


Figure 4: evolution of the computed CML.

Figure 4 shows respectively 10, 100 and 1,000 iterations of the previous function (the contrast has been enhanced so as to show tiny details). While the background changes dramatically from (e) to (f), one can check show that the patterns capturing the structural information remain quite stable.

6 Conclusion

Coupled map lattices were introduced about fifteen years ago as models of extended dynamical systems. One of their advantages is to show a very rich space-time dynamical behavior. However, most of the analytical studies in the literature have been restricted to low-order monic site functions since cubic or higher-degree polynomials are very hard to deal with analytically.

In the present case, we focused on the identification (or inverse) problem, considered in an approximation framework. This led us to define a notion of structural information coherent with information theory concepts and a new form of the evolution equation suitable for computational issues. A numerical identification algorithm has been proposed.

A lot of work remains to be undertaken in the direction of a reliable approximation tool, namely the study of the pseudo-stability of approximating CMLs and of their accuracy as well as the relationship

between patterns and the computed kernel, which was one of the primary goals. Also, global parameters should be estimated on practical experiments and compared to theoretical studies.

I would like to thank Ronan Thomas who performed the image registration.

References

- [1] Adamatzky, Andrew. “Identification of Cellular Automata”, Taylor and Francis, 1994.
- [2] Blanc-Talon, Jacques and Deniau, Laurent. “PCA and CA: a statistical approach for deterministic machines”, 1996, *Complexity International*, **2**.
- [3] Boccara, Nino and Goles, Eric and Martinez, Servet and Picco, Pierre. “Cellular Automata and Cooperative Systems”, 1993, **NATO ASI series 396**, Kluwer.
- [4] Chaté, Hugues and Manneville, Paul. “Coupled map lattices as cellular automata”, 1989, *J. of Statistical Physics*, **56**, 357–370.
- [5] Deane, J.H.B. and Jefferies, D.J. “The Behaviour of Coupled Map Chains”, 1996, *Complexity International*, **3**.
- [6] Dubacq, J.-C. and Durand, B. and Formenti, E. “Kolmogorov complexity and cellular automata classification”, 1997, Ecole Normale Supérieure de Lyon, Lyon, France.
- [7] Gottlieb, H.P.W.. “Properties of Some Generalised Logistic Maps with Fractional Exponents”, 1995 *Complexity International*, **2**.
- [8] Gutowitz, Howard. “Cellular Automata: theory and experiment”, 1991, MIT Press.
- [9] Hurd, Lyman Porter. “The application of formal language theory to the dynamical behavior of cellular automata”, 1988, Princeton University, PhD thesis.
- [10] Kaneko, K. “Period-doubling of kink-antikink patterns, quasiperiodicity in anti-ferro-like structures and spatial intermittency in coupled logistic lattice”, 1984, *Prog. Theoretical Physics*, **72:3**, 480–486.
- [11] Kendall, Bruce E. and Fox, Gordon A. “Spatial Structure, Environmental Heterogeneity, and Population Dynamics: Analysis of the Coupled Logistic Map”, 1997, *Theoretical Population Biology*.
- [12] Maass, Alejandro. “On sofic limit sets of cellular automata”, 1991, Université de Marseille, PhD thesis.
- [13] Marr, David and Hildreth, E. “Theory of Edge Detection”, 1980, *Proc. Roy. Soc.*, **B-207**, 187–214.
- [14] Mitchell, Melanie and Hraber, Peter T. and Crutchfield, James P. “Revisiting the edge of Chaos: Evolving Cellular Automata to Perform Computations”, 1993, Santa Fe technical report.
- [15] Morel, Jean-Michel. “Variational Methods in Image Segmentation”, 1995, *Prog. in Nonlinear Diff. Equ. and their Applications series*, **14**, Birkhäuser.
- [16] Rees, David and Bruteanu, Christian and Turpie, Duncan. “Reaction-Diffusion Textures using Coupled Map Lattices”, 1992, *ISSPA Int. Conf.*
- [17] F.C. Richards and T.P. Meyer and N.H. Packard. “Extracting cellular automata rules directly from experimental data”, 1991, in *Cellular Automata: Theory and Experiment*, H. Gutowitz editor, MIT North Holland.

- [18] Roy, Manojit and Amritkar, R.E. “Stochastic Coherence in Coupled Map Lattices”, 1997, *J. Phys.*, **48:271**.
- [19] Roy, Manojit and Amritkar, R.E. “Effect of noise on coupled chaotic systems”, 1997, *J. Phys.*, **48:271**.
- [20] Vaario, Jari and Shimohara, Katsunori. “On Formation of Structures”, 1995, in *ECAL'95: Advances in Artificial Life*, Morán, Federico and Moreno, Alvaro and Merelo, Juan J. and Chacón, Pablo editor, **Lect. Notes in Art. Int.**, **929**, Springer-Verlag, 421–435.
- [21] Press, William H. and Teukolsky, Saul A. and Vetterling, William T. and Flannery, Brian P. “Numerical Recipes in C, The art of Scientific Computing”, 1992, Cambridge Press.
- [22] Wolfram, Stephen. “Cellular Automata and Complexity”, 1994, Addison-Wesley.
- [23] Wootters, William K. and Langton, Chris G. “Is there a sharp phase transition for deterministic cellular automata?”, 1991, in *Cellular Automata: Theory and Experiment*, Gutowitz, Howard editor, MIT, North-Holland.

Succinct representation of the Poincaré map for periodically driven differential equations

J.H.B. Deane and D.J. Jefferies
 Department of Electronic and Electrical Engineering,
 University of Surrey,
 Guildford GU2 5XH,
 United Kingdom
 ees1jd/ees1dj@ee.surrey.ac.uk <http://www.ee.surrey.ac.uk>

Abstract

A differential equation, periodically driven with period T , defines the time evolution of the solution, a state vector $\underline{x}(t)$. The Poincaré, or time one, map is a function that relates $\underline{x}(t+T)$ to $\underline{x}(t)$. For most second and higher order nonlinear differential equations, the Poincaré map is not available in a closed form; it can generally only be inferred from numerical calculations.

In this paper, we derive an iterative representation of the Poincaré map for Duffing's equation. Our objectives are (a) to represent the mapping in as succinct a form as possible (compact enough to be published in this paper) and (b) to demonstrate that this map representation adequately reproduces the behaviour of Duffing's equation, for instance bifurcation diagrams, co-existing attractors and Poincaré sections. We succeed in these objectives, and our representation increases computation speed by a factor of 45 over traditional numerical calculations.

1 Introduction

In this paper we derive a representation of the Poincaré map for a periodically driven differential equation. We then compare the mapping in this form to results obtained by traditional numerical integration of the differential equation.

For the purposes of illustration, we confine ourselves to the well-known Duffing equation in the form:

$$\frac{d^2x}{dt^2} + \frac{3}{2} \frac{dx}{dt} + 40x(x^2 - 1) = A \sin 2\pi t \quad (1)$$

although our conclusions are valid for a wide class of differential equations. The period of the drive $T = 1$ and its amplitude is A . This differential equation has been studied intensively in the past because, despite its apparent simplicity, it can display coexisting solutions, and periodic and chaotic behaviour as A is varied.

Throughout what follows, it is convenient to re-write equation (1) as two coupled differential equations:

$$\begin{aligned} \dot{x} &= y \\ \dot{y} &= A \sin 2\pi t - 3y/2 - 40x(x^2 - 1), \end{aligned} \quad (2)$$

where \dot{x} is the time derivative of x . The state vector, which describes the state of the system at a time t , is then $\underline{x}(t) = [x(t), y(t)]$.

Of particular interest is the Poincaré map for Duffing's equation, defined as the function $\underline{\Phi}$ such that

$$\underline{x}(T) = \underline{\Phi}(\underline{x}(0)) \quad (3)$$

or, splitting $\underline{\Phi}$ into its components,

$$x(T) = \Phi_x[x(0), y(0)] \quad \text{and} \quad y(T) = \Phi_y[x(0), y(0)] \quad (4)$$




Figure 1: The Poincaré map, $\Phi_x[x(0), y(0)]$, obtained numerically. In this and the following two figures, $A = 13$, corresponding to a chaotic solution.

It can be shown [1] that the mapping so defined is unique, single valued and differentiable. The Poincaré map expresses all the essential information about solutions to Duffing's equation; for instance, for a value of A for which the differential equation has a period- n /chaotic solution, iterating the Poincaré map will also produce a sequence of points that repeats every n /is chaotic. Hence, if the Poincaré map were available in closed form, this could be iterated for an arbitrary initial condition \underline{x}_0 to unravel the behaviour of solutions of Duffing's equation (with the same initial condition) without resorting to numerical solutions.

What does $\underline{\Phi}(\underline{x})$ actually look like? Figures (1) and (2) show, respectively, $\Phi_x[x(0), y(0)]$ and $\Phi_y[x(0), y(0)]$. Figure (3) is a section through these maps at $y(0) = 2$. The figures suggest that $\underline{\Phi}(\underline{x})$ has a reasonably complex structure with large first- and higher-order derivatives. This in turn suggests the impracticality of representing the map by standard surface fitting methods, for example two-dimensional splines or two-variable polynomial curve fitting. In the case of splines, a large number would be required to reproduce the mapping faithfully, and in the polynomial case, a very high order polynomial would be needed. This is contrary to our stated objective of representing the Poincaré map in as succinct a form as possible.

2 Derivation of the representation

We now derive our representation of the Poincaré map. Figure (5) illustrates the process of analytical continuation [2] applied to this problem. The techniques is also known as cell-to-cell mapping [3].

2.1 Analytical continuation

It is necessary to look on the solution $x(t)$ as a function of *complex* time. This function has poles which are a short distance off the real t -axis, and their positions can be calculated [4] using, for instance, the ratio-like test of [5]. It is the existence of these poles that prevents us from writing the Poincaré map simply as a two variable power series in $x(0), y(0)$ — they are close enough to the real t -axis to prevent a single such series from converging over the whole range $t = 0$ to $t = T$.




Figure 2: The Poincaré map, $\Phi_y[x(0), y(0)]$.

We can estimate from real $x(t)$ data only, roughly how far the poles are from the real axis. Numerical evidence [4] and leading order analysis both suggest that all the poles in complex t for Duffing’s equation are of first order. They must also occur in conjugate pairs since $x(t)$ is real for real t . Hence, $x(t)$ for real t around the line joining such a pair, must be of the form

$$x(t) = \frac{r}{2i \operatorname{Im} z} \left[\frac{1}{t - z} - \frac{1}{t - z^*} \right]$$

where r is the residue of the poles and z, z^* are their positions. We require an estimate of $\operatorname{Im} z$. The turning point of $x(t)$ occurs at $t = \operatorname{Re} z$. Furthermore,

$$\ddot{x}(\operatorname{Re} z) = -\frac{2r}{(\operatorname{Im} z)^4}$$

and so

$$\operatorname{Im} z \approx \pm \sqrt{\frac{-2x(t)}{\ddot{x}(t)}} \Bigg|_{t=\operatorname{Re} z} \tag{5}$$

This formula requires only data for real t and no knowledge of the residue. Typical results are shown in figure 4. The calculation is only an estimate because in practice there are many pole pairs; the more isolated they are, the better the estimate.

Analytical continuation is a technique that allows us to work around the problem of convergence limited by the presence of poles. We assume that about a point $t = t_k$, there exists a series solution to a second-order differential equation with initial conditions $\underline{x}(t_k) = [x(t_k), y(t_k)]$, in the form:

$$x(t) = \sum_{n=0}^{\infty} a_n (t - t_k)^n, \tag{6}$$

with $a_0 = x(t_k)$ and $a_1 = y(t_k)$. The series (6) converges inside the circle C_k , centre t_k , of radius r_k ; r_k is the distance to the nearest pole in the complex time plane, as shown in figure (5). The calculation of the mapping then consists of the following steps:




Figure 3: Sections through the Poincaré map, $\Phi_x[x(0), -2]$ (left) and $\Phi_y[x(0), -2]$ (right).




Figure 4: Approximate nearest pole position, and $x(t)$ with its turning points marked, for $A = 13$. The pole positions are estimated from equation 5.




Figure 5: Analytical continuation is used to derive the mapping. The circles C_0, C_1, \dots with radii r_0, r_1, \dots , are the regions of convergence of the series solution to Duffing's equation about the points $t = t_0, t = t_1, \dots$. The crosses (X) represent the position of the poles in the complex plane, and occur in conjugate pairs.

1. Given an initial condition, $\underline{x}(0)$, at $t = t_0 = 0$, calculate the series for $\underline{x}(t)$ as described below.
2. Use this series to estimate $\underline{x}(t_1)$. A practical value of t_1 usually has to be found experimentally, since the radius of convergence of the series is rarely known in advance.
3. Repeat the above process expanding around the points $t = t_1, \dots, t = t_{N_c-2}$, where N_c , the number of circles used, is chosen so that an adequate representation of the map is obtained. Again, this value generally has to be obtained experimentally.

Hence, defining the function $\underline{\phi}_k$ as

$$\underline{x}(t_{k+1}) = \underline{\phi}_k[\underline{x}(t_k)],$$

the Poincaré map can be expressed as

$$\underline{\Phi}(\underline{x}) = \underline{\phi}_{N_c-2}(\underline{\phi}_{N_c-3}(\dots \underline{\phi}_1(\underline{\phi}_0(\underline{x})) \dots)) \tag{7}$$

We now explain how to derive the series for the $\underline{\phi}_k$, before demonstrating how the result in equation (7) works in practice.

2.2 The recursion formula

Substituting the series (6) into an m -th order differential equation results in a recursion formula for the coefficients a_m, a_{m+1}, \dots in terms of the m initial conditions. For a linear/nonlinear differential equation this recursion formula will be linear/nonlinear respectively.

In the case of the Duffing equation, the recursion formula is

$$a_{n+2} = (-1)^{\lfloor n/2 \rfloor} A \frac{[2\pi(t-t_k)]^n}{(n+2)!} \begin{cases} \sin 2\pi t_k & n \text{ even} \\ \cos 2\pi t_k & n \text{ odd} \end{cases} - \frac{3a_{n+1}}{2(n+2)} + \frac{40(a_n - c_n)}{(n+1)(n+2)} \quad (8)$$

where the c_n are defined by

$$\sum_{n=0}^{\infty} c_n (t-t_k)^n = \left[\sum_{n=0}^{\infty} a_n (t-t_k)^n \right]^3.$$

Nonlinearity enters the recursion formula (8) solely through the c_n . We can now calculate the series for $\underline{\phi}_k$, from which the Poincaré map for Duffing's equation can be built up recursively.

3 Results

A little experimentation shows that if we sum the series (6) up to and including $n = 5$, then $N_c = 10$ circles suffice for analytical continuation from $t = 0$ to $t = 1$, the period of the drive. The coefficients $a_0 \dots a_5$, for arbitrary A , are sufficiently simple algebraic expressions to be included in the Appendix. We now compare the Poincaré map constructed in the way described with numerical calculations. The latter were carried out using a variable-order, variable step routine from the NAG library, with a tolerance 10^{-8} .

3.1 Basin of attraction

The basin of attraction of Duffing's equation is infinite — that is, any initial condition will lead to a solution that eventually settles down to a small amplitude oscillation. On the other hand, our Poincaré map approximation in equation (7) will have only a finite basin of attraction. That is, there are initial conditions $x(0), y(0)$ that will lead to a divergent sequence with x, y tending to infinity. The basin of attraction of our Poincaré map approximation is defined as the region of the x, y -plane for which this does *not* happen, and is illustrated in figure 6.

3.2 Comparison of Poincaré sections

Another comparison we have carried out concerns a chaotic Poincaré section computed by both methods. Using the same initial condition and $A = 13$ in both cases, we calculated (x_0, y_0) to (x_{10000}, y_{10000}) . We have plotted only the first 100 values in figure 7 for clarity. Note that the actual point positions calculated by both methods diverge from each other (only to be expected in a chaotic system), but that the underlying structure of the attractor is insensitive to the method of calculation.

As an aside, it is interesting to note that our approximation runs about 45 times faster than the standard numerical method.

3.3 Comparison of bifurcation diagrams

The final comparison we report here shows that our map representation is valid over a range of bifurcation parameter A . The two bifurcation diagrams are shown in figure 8. The upper one is calculated by our algorithm and the lower one by a standard numerical integration, and the differences between them are seen to be small.

At $A \approx 17$ the two bifurcation diagrams are seen to differ; this is the result of the co-existence of two different solutions. Figure 9 shows that at $A = 17.1$, both the numerical solution and the mapping display (a) a period-1 solution and (b) a chaotic solution. Which one is observed depends on the initial conditions.




Figure 6: The basin of attraction — on and within the thick black curve — for our 5th-order, 10 circle approximation to the Poincaré map with $A = 13$. The Poincaré section is also shown and clearly lies well within the basin of attraction.




Figure 7: Comparison of Poincaré sections for $A = 13$.

4 Discussion and conclusions

We have derived and verified a compact, iterative representation of the Poincaré map for the Duffing equation. Our method can be regarded as a highly specialised method for solving that particular differential equation, and as a result of this, it is very efficient compared to standard, general numerical techniques. In practice, our method is around 45 times quicker. There is every reason to believe that the same technique can be applied to other periodically-driven differential equations.

We have verified that bifurcation diagrams and Poincaré sections can both be faithfully reproduced by our mapping, whose basin of attraction, although finite, is large enough for most applications we can envisage.

Finally, our work raises several interesting questions, among them:

Can any analytical results be derived from our representation?

Can a useful bifurcation analysis be carried out using our representation?

The expansion points used, t_k , were uniformly spaced; might there be a better arrangement? Is there an even more compact representation available? For instance, a Padé approximation to the series (6) might have some advantages.

A Appendix

The 5th-order approximation for the y -component of $\underline{\phi}_k(\underline{x})$ used in this paper is $a_0 + a_1 h + \dots + a_5 h^5$ where $h = t - t_k$. Writing $x = x(0)$, $y = y(0)$ for the initial conditions, we have $a_0 = y$, $a_1 = s + 40x(1 - x^2) - 3y/2$ and

$$a_2 = -\frac{3s}{4} + \pi c - 30x + 30x^3 - 60x^2y + \frac{169y}{8}$$

$$a_3 = -\frac{\pi c}{2} + \left(-\frac{2\pi^2}{3} + \frac{169}{24}\right)s + \frac{845x}{3} - \frac{329y}{16} - 20sx^2 - 40xy^2 +$$




Figure 8: Comparison of bifurcation diagrams. The upper diagram was produced with the algorithm described in this paper; the lower one was calculated numerically.




Figure 9: Co-existing solutions at $A = 17.1$. Continuous lines: numerically obtained period-1 and chaotic solutions; filled circle and crosses: Poincaré section of the same solutions.

$$\begin{aligned}
 & 60x^2y - \frac{3245x^3}{3} + 800x^5 \\
 a_4 = & \left(-\frac{\pi^3}{3} + \frac{169\pi}{48}\right)c + \left(-\frac{329}{64} + \frac{\pi^2}{4}\right)s + \frac{30001y}{384} - \frac{1645x}{8} - \\
 & 5(2\pi c - 3s)x^2 - 30sxy + \frac{6445x^3}{8} - 10y^3 - \frac{6535x^2y}{4} + 60xy^2 - \\
 & 600x^5 + 1800x^4y \\
 a_5 = & \left(\frac{30001}{1920} - \frac{169\pi^2}{120} + \frac{2\pi^4}{15}\right)s + \left(-\frac{329\pi}{160} + \frac{\pi^3}{10}\right)c + \left(\frac{30001}{48} - 6s^2\right)x - \\
 & \frac{82641y}{1280} + \left(6\pi c + \left[-\frac{2267}{4} + 4\pi^2\right]s\right)x^2 + (-16\pi c + 39s)xy - \\
 & 12sy^2 - \frac{657361x^3}{48} + 21y^3 + \frac{3867x^2y}{2} - \frac{1859xy^2}{2} + 600sx^4 - \\
 & 2100x^4y + 27470x^5 + 1680x^3y^2 - 14400x^7
 \end{aligned}$$

where $s = A \sin 2\pi t_k$ and $c = A \cos 2\pi t_k$.

The x -component of $\underline{\phi}_k(\underline{x})$ is then just $x + \int_0^h a_0 + a_1h + \dots + a_4h^4 dh$ to 5th-order in h .

References

- [1] MW Hirsch and S Smale, *Differential equations, dynamical systems, and linear algebra*, Academic Press, ISBN 0-12-349550-4 (1974)
- [2] TM MacRobert, *Functions of a complex variable*, Macmillan and co., London (1938)
- [3] C Hsu, *Cell-to-cell mapping*, Springer-Verlag, ISBN 0-387-96520-3 (1987)

- [4] K Konno and H Tatenno, *Duffing's equation in complex time and chaos*, **Progress in Theoretical Physics Letters**, vol. 72 no. 5, pp 1047–1049 (1984)
- [5] YF Chang and G Corliss, *Ratio-like and recurrence relation test for convergence of series*, **Journal of the Institute of Mathematics and its Applications**, vol. 25, pp 349–359 (1980)

Convergence and Aperiodicity in Fuzzy Cellular Automata: Revisiting Rule 90

P. Flocchini

Département d'Informatique,
Université de Québec à Hull,
Hull J8X 3X7, Canada,
flocchini@uqah.quebec.ca

F. Geurts

Service de Mathématiques de la Gestion,
Université Libre de Bruxelles, CP 210/01,
Boulevard du Triomphe, 1050 Bruxelles, Belgium,
fgeurts@smg.ulb.ac.be

A. Mingarelli

Department of Mathematics and Statistics,
Carleton University,
1125 Colonel By Drive, 538 Herzberg Building,
Ottawa K1S 5B6, Canada,
amingare@math.carleton.ca

N. Santoro

School of Computer Science,
Carleton University,
1125 Colonel By Drive, 538 Herzberg Building,
Ottawa K1S 5B6, Canada,
santoro@scs.carleton.ca

Abstract

Lifting Boolean to continuous CA by fuzzification of the disjunctive normal form local rule enables the analytical derivation of results for an interesting class of systems: Fuzzy CA, which include the Boolean evolution as a particular case. We concentrate on Fuzzy Rule 90, whose Boolean version has deserved some attention for the complex patterns it generates. We prove that aperiodicity is present in both versions, and a strong convergence to a fixed point is added in the fuzzy case, which extends and complements Jen's result on aperiodicity of Boolean CA. Moreover, we show that the complex pattern and the apparent chaotic behaviour observed in its Boolean evolution are not an inherent property of Rule 90 but rather an artifact of the fixed point being an extreme of one of the two value intervals.

1 Introduction

Boolean cellular automata (CA) have been introduced by von Neumann as models of self-organizing/reproducing behaviors [12]. Today, their applications range from ecology to theoretical computer science [5, 10, 13]. Recently, a *continuous* model of CA, namely *Fuzzy CA*, has been introduced to investigate complex behaviors [1, 2] and to study the impact of perturbations (e.g. noisy sources, computation errors, mutations, etc.) on the evolution of Boolean CA [4]. While Boolean CA have been extensively studied from many different angles, little is known about continuous CA, although related to Coupled Map Lattices which, on the contrary, have received more attention, both theoretically and practically [9].

One of the most studied problems in the CA community is the classification of long-term behaviors. For instance, in Wolfram's classification [13], the CA rule space is partitioned into four categories, based on the observed long-term attracting behavior of CA starting from random initial configurations. From [3], we know that any nontrivial classification is undecidable, but this restriction has not discouraged many authors to study subclasses of Wolfram's Classes 3 and 4, which respectively reveal complex dynamical and computational features. In particular, the chaotic Boolean Rules 90 and 18 have deserved a special attention, due to the fractal-like patterns they generate [7, 8, 11].

The goal of our work is the analytical study of dynamical properties of Fuzzy CA rules which generalize CA to continuous local states by using fuzzy operators instead of Boolean ones. The main advantage of lifting CA from Boolean to real numbers is that we can use classical tools of dynamical systems theory. In this paper, we focus on Fuzzy Rule 90, as its Boolean evolution, although linear, is considered to be complex (e.g., see [8, 11]). We show that the behavior of the fuzzy version, which has never been studied before, is actually very simple: independently of initial configurations,

the system converges to a fixed homogeneous configuration. In the particular case of finite support configurations, we generalize Jen's result [8, Prop. 2] on aperiodic sequences in Boolean Rule 90: we prove aperiodicity of temporal sequences in each site, but also in every diagonal and even every nontrivial temporal sequence of the spatio-temporal diagram of Fuzzy Rule 90.

The results of our analysis lead to an intriguing question. Since Fuzzy Rule 90 has *always* a simple behaviour (convergence to a fixed point), where do the complex pattern in its Boolean evolution come from? We solve the puzzle and show that the *apparent* chaotic behaviour of Boolean Rule 90 and its famous fractal-like patterns are only an artifact of having the fixed point $\frac{1}{2}$ as an extreme of the two discretized intervals (i.e., the result of visualizing the behaviour using only two values).

2 Basic definitions

A cellular automaton is a collection of cells arranged on a graph. All cells share the same local space (i.e., the set of values cells range in), the same neighborhood structure (i.e., the cells to which a cell is connected), and the same local function (i.e., the function defining the effect of neighbors on each cell, also called transition function or rule). The global evolution is defined by the synchronous update of all values according to the local function applied to the neighborhood of each cell. A configuration of the automaton is a description of all cell values.

Given a linear bi-infinite lattice of cells, the local Boolean space $\{0, 1\}$, the neighborhood structure $\langle \text{left neighbor, itself, right neighbor} \rangle$, and a local rule $g : \{0, 1\}^3 \mapsto \{0, 1\}$, the global dynamics of an *elementary CA* is defined by:

$$\begin{aligned} f & : \{0, 1\}^{\mathbb{Z}} \mapsto \{0, 1\}^{\mathbb{Z}} \\ \text{s.t.} & \quad \forall i \in \mathbb{Z}, f(x)_i = g(x_{i-1}, x_i, x_{i+1}). \end{aligned}$$

The *local rule* is defined by the 8 possible local configurations a cell can detect in its direct neighborhood:

$$(000, 001, 010, 011, 100, 101, 110, 111) \rightarrow (r_0, \dots, r_7),$$

where each triplet represents a local configuration of the left neighbour, the cell itself, and the right neighbour. In general, the value $\sum_{i=0:7} 2^i r_i$ is used as the name of the rule. The local rule of any Boolean CA is canonically expressed as a *disjunctive normal form*:

$$g(x_1, x_2, x_3) = \bigvee_{i|r_i=1} \bigwedge_{j=1:3} x_j^{d_{ij}}$$

where d_{ij} is the j -th digit, from left to right, of the binary expression of i , and x^0 (resp. x^1) stands for $\neg x$ (resp. x).

Definition 11.1 *A Fuzzy CA is obtained by fuzzification of the local function of a Boolean CA: in the disjunctive normal form, $(a \vee b)$ is replaced by $(a + b)$, $(a \wedge b)$ by (ab) , and $(\neg a)$ by $(1 - a)$. The resulting local rule is a real-valued function simulating the original function on $\{0, 1\}^3$, with $l(a, 0) = 1 - a$ and $l(a, 1) = a$:*

$$\begin{aligned} g & : [0, 1]^3 \mapsto [0, 1] \\ \text{s.t.} & \quad g(x_1, x_2, x_3) = \sum_{i=0:7} r_i \prod_{j=1:3} l(x_j, d_{i,j}). \end{aligned}$$

Example 11.1 *Consider rule 14 = 2 + 4 + 8:*

$$(000, 001, 010, 011, 100, 101, 110, 111) \rightarrow (0, 1, 1, 1, 0, 0, 0, 0).$$

The canonical expression of rule 14 is:

$$g_{14}(x_1, x_2, x_3) = (\neg x_1 \wedge \neg x_2 \wedge x_3) \vee (\neg x_1 \wedge x_2 \wedge \neg x_3) \vee (\neg x_1 \wedge x_2 \wedge x_3).$$

The fuzzification process after simplification yields:

$$g_{14}(x_1, x_2, x_3) = (1 - x_1) \cdot (x_2 + x_3 - x_2 \cdot x_3).$$

In the rest of this paper, unless specified otherwise, we will study Fuzzy Rule 90, defined by the local function

$$g : [0, 1]^3 \mapsto [0, 1]$$

$$\text{s.t. } g(x, y, z) = x + z - 2xz.$$

3 Homogeneous configurations

The first step of our analysis consists in a strong assumption on the values undertaken by the fuzzy cells of the automaton we consider: they are all initialized to the same value and, as the global dynamics is homogeneous, all subsequent configurations are homogeneous, too. This reduces the infinite-dimensional system to a one-dimensional one, the dynamics of which is analyzed in a straightforward way. The motivation is to get a clear idea of the local process.

The reduced function is

$$h : [0, 1] \mapsto [0, 1]$$

$$\text{s.t. } h(x) = 2x(1 - x).$$

It has two fixed points: 0 and $\frac{1}{2}$. The absolute value of the first derivative $h'(x) = 2(1 - 2x)$ evaluated in each of these fixed points gives $|h'(0)| = 2$ and $|h'(\frac{1}{2})| = 0$. Thus, 0 turns out to be repelling, and $\frac{1}{2}$ is an attractor. The second order Taylor expansion of h around x is

$$h(x + u) = h(x) + h'(x)u + \frac{1}{2}h''(x)u^2$$

$$= 2x(1 - x) + (2 - 4x)u - 2u^2$$

and, around the fixed points, we have $h(u) \approx 2u$ and $h(\frac{1}{2} + u) = \frac{1}{2} - 2u^2$. Initial conditions close to zero are essentially multiplied by two, and attracted by $\frac{1}{2}$: for any $u \in (-\frac{1}{2}, \frac{1}{2})$, $h(\frac{1}{2} + u)$ is strictly closer to $\frac{1}{2}$ than $\frac{1}{2} + u$.

Thus, this one-dimensional system behaves in a very simple way: 0 is a repelling fixed point; 1 is attracted to 0 in one iteration; $\frac{1}{2}$ is an attracting fixed point whose basin is the open interval $(0, 1)$.

4 Single values in zero backgrounds

In this second analysis, all cells but one are initially set to 0. This step is motivated by the classical analysis of finite support configurations of Boolean CA [8]. For example, starting from a single value $a = \frac{1}{4}$ in a zero background, the spatio-temporal evolution is represented in Table 1.

Time	Local states								
	...	-3	-2	-1	0	1	2	3	...
0	...	0	0	0	$\frac{1}{4}$	0	0	0	...
1	...	0	0	$\frac{1}{4}$	0	$\frac{1}{4}$	0	0	...
2	...	0	$\frac{1}{4}$	0	$\frac{3}{8}$	0	$\frac{1}{4}$	0	...
3	...	$\frac{1}{4}$	0	$\frac{7}{16}$	0	$\frac{7}{16}$	0	$\frac{1}{4}$...
4	...	0	$\frac{15}{32}$	0	$\frac{63}{128}$	0	$\frac{15}{32}$	0	...
⋮					⋮				

Table 1: Evolution from $\frac{1}{4}$ in a zero background.

More abstractly, assuming $a \neq 0$ and $g_a(x) = g(x, \bullet, a) = g(a, \bullet, x) = a + x(1 - 2a)$, we have the evolution of Table 2. What is the rule underlying these numbers, if any?

Time	Local states								
	...	-3	-2	-1	0	1	2	3	...
0	...	0	0	0	a	0	0	0	...
1	...	0	0	a	0	a	0	0	...
2	...	0	a	0	$h(a)$	0	a	0	...
3	...	a	0	$g_a(h(a))$	0	$g_a(h(a))$	0	a	...
4	...	0	$g_a^2(h(a))$	0	$h(g_a(h(a)))$	0	$g_a^2(h(a))$	0	...
⋮					⋮				

Table 2: Evolution from a in a zero background.

Definition 11.2 *The spatio-temporal diagram from an initial configuration x^0 is the double sequence $(x_i^t)_{i \in \mathbb{Z}, t \in \mathbb{N}}$ where t expresses time steps, and i denotes cell indices.*

Definition 11.3 *The j th diagonal is the sequence $(x_i^{i+2(j-1)})_{i \geq 0}$.*

Definition 11.4 *The light cone from a cell x_i^t is the set $\{x_j^{t+p} \mid p \geq 0 \wedge j \in \{i-p, \dots, i+p\}\}$.*

Proposition 11.1 *The second diagonal of the spatio-temporal diagram obtained by the evolution of Fuzzy Rule 90 from a single value $a \in (0, 1)$ in a zero background converges to $\frac{1}{2}$.*

PROOF. The first diagonal starting from the central non-zero value a is uniformly equal to a . The second diagonal, from $h(a)$, can be obtained by successive iterations of $g_a(x)$. It has exactly one fixed point $x = \frac{1}{2}$, independently of a . The absolute value of the slope of this linear function is smaller than one iff $a \in (0, 1)$. This means that the second diagonal converges to $\frac{1}{2}$ as time goes to infinity.

Let $f(t, i) = \binom{t}{\frac{t+i}{2}}$, where $\binom{a}{b} = \frac{a!}{b!(a-b)!}$.

Proposition 11.2 *The spatio-temporal diagram from a single value $a \in (0, 1)$ in a zero background is explicitly given by: $\forall t \in \mathbb{N}$,*

$$x_i^t = \begin{cases} \frac{1}{2}(1 - (1 - 2a)^{f(t,i)}) & \text{if } t+i \text{ is even and } i \in \{-t, \dots, t\} \\ 0 & \text{otherwise} \end{cases}$$

PROOF. By induction. The first case is straightforward: $x_0^0 = a$ and $\forall i \neq 0, x_i^0 = 0$. The inductive case is twofold.

- If $(t+1) + i$ is odd, $i > t+1$ or $i < -t-1$, then $x_{i-1}^t = x_{i+1}^t = 0$, whence $x_i^{t+1} = 0$.
- Otherwise, $(t+1) + i$ and $t + (i-1)$ are even, and $i \in \{-t-1, \dots, t+1\}$. Thus, if $-(t-1) \leq i \leq t-1$, then

$$\begin{aligned} x_i^{t+1} &= g(x_{i-1}^t, \bullet, x_{i+1}^t) \\ &= \frac{1}{2}(1 - (1 - 2a)^{f(t,i-1)+f(t,i+1)}) \\ &= \frac{1}{2}(1 - (1 - 2a)^{f(t+1,i)}). \end{aligned}$$

The two limit cases are $i-1 = -t-2$ and, symmetrically, $i+1 = t+2$. Let us examine the first one:

$$\begin{aligned} x_i^{t+1} &= x_{-t-1}^{t+1} = g(x_{-t-2}^t, \bullet, x_{-t}^t) = g(0, \bullet, x_{-t}^t) = x_{-t}^t \\ &= \frac{1}{2}(1 - (1 - 2a)^{f(t,-t)}) = a. \end{aligned}$$



Figure 1: Evolution from a random fuzzy configuration. Grey levels indicate different value ranges. Time evolves from top to bottom.

Thus, *any infinite sequence of non-zero terms* in the diagram, no matter how crazy it is, *converges to $\frac{1}{2}$* , provided that it is embedded in the light cone originating from the central a , and the sequence of non-zero time steps tends to infinity: the central column, any diagonal, any sequence containing horizontal segments or even backward loops, they all converge to $\frac{1}{2}$.

Corollary 11.1 *Let $x_0^0 = a \in (0, 1)$, $x_{i \neq 0}^0 = 0$, $i : \mathbb{N} \mapsto \mathbb{Z}$ and $\tau : \mathbb{N} \mapsto \mathbb{N}$ be two functions such that $i(0) = \tau(0) = 0$, $|i(j)| \leq \tau(j)$, $x_{i(j)}^{\tau(j)} \neq 0$, and $\lim_{j \rightarrow \infty} \tau(j) = \infty$. Then, the sequence $(x_{i(j)}^{\tau(j)})_{j \in \mathbb{N}}$ converges to $\frac{1}{2}$.*

Remark 11.1 *We have considered, both aperiodicity and convergence are present. Exactly as in nontrivial evolutions of Boolean Rule 90 [8], any nontrivial infinite sequence of states taken in the spatio-temporal evolution of Fuzzy Rule 90 from a single value in a zero background is aperiodic, due to Proposition 11.2. On the other hand, any such sequence converges to a specific value, here $\frac{1}{2}$, which was not the case of Boolean Rule 90.*

Remark 11.2 *If the initial configuration contains two consecutive values, a and b , in a zero background, the result is obtained by superposition of the individual diagrams obtained from a and b . Let (x_i^t) be such that $x_0^0 = a$ and $\forall i \neq 0, x_i^0 = 0$, (y_i^t) be such that $y_1^0 = b$ and $\forall i \neq 1, y_i^0 = 0$, then $(z_i^t = x_i^t + y_i^t)$ is the diagram starting from $z_0^0 = a, z_1^0 = b$ and $\forall i \in \mathbb{Z} \setminus \{0, 1\}, z_i^0 = 0$. Unfortunately, this property cannot be extended to larger finite support configurations.*

5 Infinite heterogeneous configurations

Here, we extend the analysis to evolutions starting from heterogeneous configurations, *i.e.*, initial values chosen arbitrarily in $(0, 1)$ (see Fig. 1). With such a weak assumption, we are, of course, not able to derive the closed-form expression of all x_i^t , but the last result mentioned in §4 still holds: we prove that any sequence of non-zero terms in the spatio-temporal diagram converges to $\frac{1}{2}$. We first establish useful lemmas, then we prove the main result.

Let us rewrite the local function g as follows:

$$\phi(x, y) = g\left(\frac{1}{2} + x, \bullet, \frac{1}{2} + y\right) = \frac{1}{2} - 2xy. \quad (1)$$

Lemma 11.1 *The function g is contracting around $\frac{1}{2}$:*

$$\forall x, y \in \left(0, \frac{1}{2}\right) \cup \left(\frac{1}{2}, 1\right), \left|g(x, \bullet, y) - \frac{1}{2}\right| < \min\left\{\left|x - \frac{1}{2}\right|, \left|y - \frac{1}{2}\right|\right\}.$$

PROOF. Let us change the variables: $x = \frac{1}{2} + v, y = \frac{1}{2} + w$, and rewrite the left-hand side expression: $\left|g(x, \bullet, y) - \frac{1}{2}\right| = \left|g\left(\frac{1}{2} + v, \bullet, \frac{1}{2} + w\right) - \frac{1}{2}\right| = \left|\phi(v, w) - \frac{1}{2}\right| = |2vw|$. Since $0 < |w| < \frac{1}{2}$, we have $|2vw| < |v| = \left|x - \frac{1}{2}\right|$. Symmetrically, $|2vw| < |w| = \left|y - \frac{1}{2}\right|$.

Lemma 11.2 *The function g is k -contracting around $\frac{1}{2}$ on $[\frac{1-k}{2}, \frac{1+k}{2}]$:*

$$\forall x, y \in [\frac{1-k}{2}, \frac{1+k}{2}], |g(x, \bullet, y) - \frac{1}{2}| \leq k \min\{|x - \frac{1}{2}|, |y - \frac{1}{2}|\}.$$

PROOF. Changing variables as in the previous lemma, we have $|g(x, \bullet, y) - \frac{1}{2}| = |2vw| \leq 2\frac{k}{2}|v|$, since $y = w + \frac{1}{2} \in [\frac{1-k}{2}, \frac{1+k}{2}]$.

The next example gives a convergence rate for purely temporal sequences of the diagram, *i.e.* columns, if one every other row is ignored, as g does not depend on its central argument. The proof easily follows from the previous lemma.

Example 11.2 *If $x^0 \in (0, 1)^{\mathbb{Z}}$ and $x_0^0 \in [\frac{1-k}{2}, \frac{1+k}{2}]$, then $|x_0^2 - \frac{1}{2}| \leq k^2|x_0^0 - \frac{1}{2}|$, and the sequence $(x_0^{2^t})_{t \in \mathbb{N}}$ converges to $\frac{1}{2}$.*

If connected paths are considered in the spatio-temporal diagram, the convergence rate is k .

Example 11.3 *If $x^0 \in (0, 1)^{\mathbb{Z}}$ and $x_0^0 \in [\frac{1-k}{2}, \frac{1+k}{2}]$, then $|x_{\pm 1}^1 - \frac{1}{2}| \leq k|x_0^0 - \frac{1}{2}|$, and any connected path $(x_{i(j)}^j)_{j \in \mathbb{N}}$ such that $i(0) = 0$ and $\forall j \in \mathbb{N}, i(j+1) = i(j) \pm 1$, converges to $\frac{1}{2}$.*

Finally, using Lemma 11.2 again, we generalize Corollary 11.1: any path in the spatio-temporal can be considered, and still lead to the same conclusion.

Theorem 11.1 *Let $x^0 \in [0, 1]^{\mathbb{Z}}$ be such that $\liminf x_i^0 > 0$ and $\limsup x_i^0 < 1$, $i : \mathbb{N} \mapsto \mathbb{Z}$ and $\tau : \mathbb{N} \mapsto \mathbb{N}$ be two functions, $\tau(0) = 0$, and $\lim_{j \rightarrow \infty} \tau(j) = \infty$. Then, the sequence $(x_{i(j)}^{\tau(j)})_{j \in \mathbb{N}}$ converges to $\frac{1}{2}$.*

PROOF. Let us define $k = 2 \max\{|\liminf x_i^0 - \frac{1}{2}|, |\limsup x_i^0 - \frac{1}{2}|\}$. We have of course $0 < k < 1$ and, $\forall i, x_i^0 \in [\frac{1-k}{2}, \frac{1+k}{2}]$. Thus, $|x_{i \pm 1}^1 - \frac{1}{2}| \leq k|x_i^0 - \frac{1}{2}| \leq k\frac{k}{2}$. The evolution does not depend on position i anymore; in fact, $\forall t, i, |x_i^t - \frac{1}{2}| \leq \frac{1}{2}k^{t+1}$, whence the result.

6 On the observation precision

In the previous sections we have shown that that Fuzzy Rule 90 has a very simple behaviour: it attracts everything to $\frac{1}{2}$. The result leads to the intriguing question: where do the complex pattern in its Boolean evolution come from?

In this section, we solve the puzzle and show that the *apparent* chaotic behaviour of Boolean Rule 90 and its famous fractal-like patterns are only an artifact of having the fixed point $\frac{1}{2}$ as an extreme of the two discretized intervals; *i.e.*, the result of visualizing the behaviour using only two colors) and not an inherent property of Rule 90.

We know that Fuzzy Rule 90 attracts everything to $\frac{1}{2}$. We also know that the values alternate around this point as they get closer to it: this is due to the minus sign appearing in equation (1). The corresponding “rule table” is detailed in Table 3.

The graphical representation will thus strongly depend on whether $\frac{1}{2}$ is in the middle or on a border of some discretization interval. Partition the interval $[0, 1]$ in a finite number of subintervals that are used as coarse-grained approximations of the real numbers.

If the fixed point $\frac{1}{2}$ is in the middle of the central subinterval (as in Fig. 1), the observation becomes homogeneous after a few steps, as all iterations get very close to the fixed point quickly.

On the contrary, if the fixed point is on the border of some subinterval, the values alternate around the two subintervals bordering on the fixed point; thus, the visualization will show the alternance between intervals (*i.e.*, using different grey levels, depending upon whether they are smaller or greater than $\frac{1}{2}$) creating the seemingly chaotic triangular patterns (as shown in Fig. 2 left). This is exactly what happens in the Boolean case, when the interval $[0, 1]$ is divided into two subintervals, one having

Behavior around $\frac{1}{2}$				Boolean Rule 90			
x	y	z	$g(x, y, z)$	x	y	z	$g(x, y, z)$
–	–	–	–	0	0	0	0
–	–	+	+	0	0	1	1
–	+	–	–	0	1	0	0
–	+	+	+	0	1	1	1
+	–	–	+	1	0	0	1
+	–	+	–	1	0	1	0
+	+	–	+	1	1	0	1
+	+	+	–	1	1	1	0

Table 3: Behavior around $\frac{1}{2}$ (left). The rule table is obtained from equation (1): – (resp. +) stands for “smaller than $\frac{1}{2}$ ” (resp. “greater than $\frac{1}{2}$ ”). Remark that this table corresponds to Boolean Rule 90 (right), if – (resp. +) is replaced by 0 (resp. 1).




Figure 2: Fuzzy (left) and Boolean (right) evolutions; the patterns in the fuzzy evolution are due to oscillations around $\frac{1}{2}$, observed via an appropriate discretization of $[0, 1]$.

$\frac{1}{2}$ as an extreme. In other words, the well-known observed dynamics of Boolean Rule 90 (see Fig. 2 right) describes only the behavior of Fuzzy Rule 90 around its fixed point!

In other words, if the fixed point is on the border of some subinterval, visualization will show an artifact of the discretization process; at the same time, it *will not show* the inherent convergence to $\frac{1}{2}$ which is the basic behaviour of Rule 90.

Summarizing, the same dynamics can be observed as very simple or totally complex, depending upon whether the attracting fixed point falls in the interior or on the border of some discretization interval. This very important phenomenon is not a specific feature of Fuzzy Rule 90; actually, when looking at the behavior of such systems by simulation, the descriptive precision used is of utmost importance on the observed complexity and it should deserve a deep investigation.

In the future, we will extend the tools and results presented here to a larger class of continuous CA rules. Instead of a case by case analysis, we could use abstraction techniques (building homomorphisms between known and new systems), functional transformations and composition operators (composing known rules to obtain new ones, and combining individual properties to get global ones homomorphically) preserving qualitative and/or quantitative dynamical properties.

Acknowledgements. P.F. is supported in part by *NSERC* and *FCAR* Canada; F.G. acknowledges the financial support of *FNRS* and *CFB* Belgium when visiting Carleton University; A.M. and N.S. are supported in part by *NSERC*.

References

- [1] G. Cattaneo, P. Flocchini, G. Mauri, and N. Santoro. Cellular automata in fuzzy backgrounds. *Physica D*, 105, 105-120, 1997.
- [2] G. Cattaneo, P. Flocchini, G. Mauri, and N. Santoro. Fuzzy cellular automata and their chaotic behavior. In *Proc. International Symposium on Nonlinear Theory and its Applications, Hawaii*, volume 4, pages 1285–1289. IEICE, 1993.
- [3] K. Culik II and S. Yu. Undecidability of CA classification schemes. *Complex Systems*, 2:177–190, 1988.
- [4] P. Flocchini and N. Santoro. The chaotic evolution of information in the interaction between knowledge and uncertainty. *Complexity International*, 2, 1995. Available from <http://www.csu.edu.au/ci/vol2/ci2.html>.
- [5] M. Garzon. *Models of Massive Parallelism. Analysis of Cellular Automata and Neural Networks*. Springer-Verlag, 1995.
- [6] F. Geurts. *Abstract Compositional Analysis of Iterated Relations — A Structural Approach to Complex State-Transition Systems*. Springer-Verlag, to appear.
- [7] J. E. Hanson and J. P. Crutchfield. The attractor-basin portrait of a cellular automaton. Technical Report 91-02-012, Santa Fe Institute, 1991.
- [8] E. Jen. Aperiodicity in one-dimensional cellular automata. *Physica D*, 45:3–18, 1990.
- [9] K. Kaneko. *Theory and Application of Coupled Map Lattices*. John Wiley & Sons Ltd, 1993.
- [10] C. G. Langton. Studying artificial life with cellular automata. In *Evolution, Games, and Learning*. North Holland, 1986.
- [11] H. Y. Lee and Y. Kawahara. On dynamical behaviors of cellular automata CA-60. *Bulletin of Informatics and Cybernetics*, 25(1–2):21–25, 1992.
- [12] J. Von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana, 1966.
- [13] S. Wolfram. *Theory and Applications of Cellular Automata*. World Scientific, 1986.

A Theoretical Framework for Abundance Distributions in Complex Systems

Stephan R.P. Halloy
 New Zealand Institute for Crop & Food Research
 PB 50034, Mosgiel, NEW ZEALAND
 halloys@crop.cri.nz

Abstract

A theoretical framework is proposed to explain how and where complex systems break up into agents or species. Splits lead to diversification and to abundance distributions which are similar to power functions on a rank-abundance representation, and to lognormal functions on a frequency-abundance representation. The combined manifestation of power and lognormal functions is a polo distribution, a situation toward which there seems to be a widespread tendency in complex systems (a polo pattern attractor).

Minimal complex system organisation requires three integrated hierarchical levels, the system, agents and particles. The tendency to polo emerges, or can be explained by, resource particle interaction, in which particles are attracted to each other according to their size and inversely to their distances. Simulation of this simple rule on a preliminary model leads toward polo abundance distributions. The level of abstraction allows the theoretical framework to be applicable to all fields where complex systems are found to have polo distributions. A clearer understanding of the rules and forces leading to diversification can have a range of applications in planning and management for conservation, agriculture, business, health and other areas dealing with complex systems.

1 Introduction

Evolution explains the mechanisms by which organisms change (mutation, recombination, selection, developmental constraints), but how taxa split, and the causes and timing of diversification are not clear. Both biological and economic systems are characterised by trends of increasing diversity. The abundance distributions of the elements (species, businesses, agents) of these systems tend toward characteristic abundance-rank and frequency-abundance patterns. The patterns approach a power function in the first case and a lognormal distribution in the second (polo for short).

Polo distributions are found in many complex systems, particularly with regards to aspects of size (volume, length, biomass) in what has been called the Dyar-Hutchinson rule by May[24], referring to living systems[29, 16, 15]. Similar distributions are also known in inanimate systems (e.g. nanoparticle sizes,[32]).

First I present a theoretical framework to attempt an explanation on how and where splits occur in a complex system, thus leading to both diversification and polo distributions. Then a simple resource attraction model is proposed to study the implications and effects of the theoretical concepts. The model is based on simple rules, allowing it to simulate reality and be tested.

2 Background

Many different models have been proposed to fit abundance distributions. These range from descriptive (e.g. a mathematical curve that fits empirical data) to attempting theoretical explanations of mechanisms. Preston[26, 27, 28], MacArthur[20, 21] and Sugihara's[34] lognormal models assume a flat system with no vertical hierarchy and no links (e.g. the canonical lognormal or the broken stick model). Kauffman's[17] NK models involve a number of elemental components linked in ways that lead to global system behaviour, without hierarchies. Bak et al.'s[3] self-organised criticality depends on the variations in size of agents by particle accretion, leading to a distribution of break-points

(avalanches or catastrophes) with a power or $1/f$ distribution in a cellular automaton or topological structure. Accretion is externally determined. Each cell can accumulate particles to a certain threshold after which it ‘collapses’, sending particles to neighbouring cells. If all cells are close to the threshold (critical state), this can lead to a large domino effect. The distribution of ‘avalanches’ tends to a power function. Avalanches represent the agents.

Many more attempts have been made to explain polo distributions on a case by case basis, focussing on a particular field of knowledge (e.g. [4, 37]). The latter models do not take into account the universality of polo distributions. None of the above models has proposed a conceptual framework to serve as a theoretical basis applicable to seek the mechanisms of diversification from an undifferentiated start.

3 Defining the processes and elements of a complex system — toward a theoretical framework

3.1 Processes

In an initial amorphous mass of resources (a ‘simple’ system), breaks may arise through unevenness of attraction. Such a process has been hypothesised more simply in the broken-stick type models mentioned above. Fragmentation at random as in the broken stick, or in crushing an object, tends to produce a lognormal distribution of fragment sizes. Such models and real systems imply external forces and no explicit relationality between elements. Here I postulate that the attraction between resource particles (interactions or links), in combination with some stochastic variation in their sizes, positions or both, leads to a rupture of the amorphous mass and clustering. This clustering changes the amorphous mass into a system with agents (the clusters) separated by boundaries where resources are rarer. The force of attraction is proportional to the mass of the resources, leading to a positive feedback. As agents grow, they attract ever more resources. But the attraction is also inversely proportional to distance. All agents in complex systems can be seen to respond in some ways to this process of attracting resources in proportion to their magnitude (in whatever units this may be measured) and inversely proportional to some measure of distance or difficulty to obtain that resource.

3.2 Elements

Complex system structure. Consider two basic types of systems: simple and complex. Simple systems can be thought of as a group of undifferentiated particles. There are two hierarchies: the particles and the system. Complex systems must have a structure requiring a minimum of three hierarchical levels: particles, agents and the system (figure 1). Clumping into agents is the result of the variable attractions between particles described above. Flexible boundaries delimit agents. I postulate that the size of such agents tends to a polo distribution in all complex systems where agents exchange resources in some way (relationality, competition or as used here, attraction) and agents are capable of indeterminate growth. The polo is thus a signature of complex systems.

There are two levels of complex systems according to whether significant amounts of resources flow through the system (e.g. dynamic biological and economic systems) or not (e.g. static physical systems such as the solar system). The static complex systems are at the lower end of the scale of complexity: the pattern is frozen. If resources flow through the system, it continues to evolve into a complex dynamical system. The boundary between simple and complex systems must be like a phase transition. Thus the distinction should be relatively clear in actual systems, constituting a test to the definition.

Systems where the growth of agents is constrained through some information (e.g. size of animals in a population, size of sand grains on the beach) tend to have a normal distribution of the elements and no hierarchy, whether the underlying agents are complex or simple. This




Figure 1: Minimum necessary structure of a complex system.

in turn depends on the character measured. The human population is normally distributed for size, as size is genetically constrained, but may in some societies be lognormally distributed for wealth when this is not culturally constrained.

Particles are the minimum units of resources. From an agent's viewpoint they are discrete packages of resources of variable size or 'mass'. Particles are analogous to individuals in a biological population, to quanta of light or space in a plant community, to particles of dust in the cosmos, or to economic elements. In an economic system the component elements are often called agents, but this term should not be confused with agents as defined here except when it refers to companies.

Resources. Any thing or process for which agents may compete. Resources are designated here as an abstract Mass (M), which can equally be seen to mean actual physical mass, biomass, space, time, energy, financial resources, etc.

Agents constitute the intermediate hierarchical level, which qualitatively distinguishes a complex system from a simple system. Agents arise when an initial undifferentiated mass of particles breaks up or coalesces (i.e. boundaries are formed) into a number of parts. Each agent contains or controls a number of particles. Agents are analogous to species or companies. The system boundary contains all the particles and all the agents.

Boundaries are formed where interactions are proportionately more important between the particles inside the agent than they are between them and particles outside. The same applies for boundaries between systems at a higher level. Boundaries fluctuate and have a certain degree of permeability. This is a fundamental aspect of complex systems which is often overlooked and leads to problems of sampling and definition (e.g. where are boundaries set from our perspective), and of successive nested hierarchical levels of complexity which vary in space and time (e.g. [11], or the SWARM simulation system,[?]).

3.3 Pattern attractor

Currently defined attractors are based on systems being represented by a single point moving through phase space[38]. It is necessary to define here a new type of attractor. Pattern attractors are defined as attractors where the system can only be represented by the relative position of more than one point in phase space. The phase space here is defined by two axes: frequency or rank, and abundance or magnitude. The series of points and their positions creates a pattern which defines the system's pattern phase state. The lognormal, or any other frequency distribution could be considered as possible pattern attractors. I postulate that the polo distribution is the main pattern attractor for the abundance distributions of complex systems.

By providing this definition, we circumvent the debate on the appropriate mathematical distributions to fit to natural systems, a question which is untestable as there are infinite numbers of natural systems, each modified by its own history. The question is not what distributions fit what empirical data in a given time slice, but toward what curves distributions may be tending. Comparative studies showing variations of goodness of fit to particular models will be more informative than a particular goodness of fit. Thus many natural distributions are rather bad approximations to lognormal models. However, if natural systems consistently approximate to lognormal models when left to their internal mechanisms, while distancing themselves from the lognormal when pressured by external forces, then we can suspect the presence of a lognormal pattern attractor. Considering a mathematical curve as an attractor becomes a testable hypothesis.

3.4 The Polo Pattern Attractor

The widespread occurrence of polo distributions, and the fact that in many cases where distributions shifted away from a polo pattern return toward that pattern, strongly suggest that the power function, the lognormal distribution, or their combination in the polo distribution act as pattern attractors for complex systems. The power functions and lognormal frequency distributions described in the literature cited above are mathematically distinct. However, both distributions are present in the same natural situations given certain restrictions on scale[34, 11, 36]. A data series with a lognormal distribution will also exhibit a power function rank distribution for the right part of its range, e.g. when its left side is veiled. Lognormal distributions found in nature are generally canonical and are often veiled[24, 28, 22, 6, 7]. Conversely, empirical data series fitted to a power function[40, 23, 22] also exhibit a pronounced drop or convexity at the lower right side, a distribution which resembles an exponential function but often indicates lognormality.

The distinction between exponential and lognormal can be seen in a frequency-abundance representation, where the exponential produces a straight horizontal line and the lognormal the typical bell-shaped curve.

The ubiquitous occurrence of polo distributions has often been ignored or downplayed. For example, Solé and Alonso[33] state ‘species abundance follows a power-law distribution and not a log-normal one, as it is usually assumed.’ This distinction between the two patterns is due to several causes:

- a large array of nomenclature for different expressions describing what is at least in part the same phenomenon, for example, power functions, fractals, $1/f$ noise, allometric species relations, Pareto principle, lognormal, etc.
- a multiplicity of different representations of abundance distributions[36].
- the amount of ‘noise’ in many natural systems, leading to most data only approximating power or lognormal distributions.
- sampling problems: where samples are too small the right part of the lognormal can be identical to a power function.

In their particular fields, polo patterns (in either of the forms stated above) have been variously used to determine the optimum business and marketing strategy for a company (e.g. Pareto principle) or to determine climates on the basis of vegetation structure and vice-versa (e.g. [29]). Polo patterns have been proposed as diagnostic indicators of ecosystem health ([?] cited in [11], [18]), predictors of vegetation changes and management tools for environmental risk assessment[31], determine sustainability of agricultural systems[12] and to calculate inputs needed to maintain a system away from its ‘harmonic’ polo pattern[13]. Such studies implicitly accept that the polo pattern acts as a system pattern attractor. Predictive powers could be refined if we could standardise representations and define the fundamental unifying principles leading toward the polo distribution. Then we will also be able to search whether there are some other pattern attractors distinguishable as power, lognormal or other. The challenge is not why complex systems have a polo distribution of abundance, which they

Coordinate	1	2	3	4	5	6	7	8	9
Particle size									
Time 0	0.9	1	0.9	1.1	0.9	1	1	0.9	1
Time 1	0.9	1	0	2.9	0	1	1	0.9	1
Time 2	1.9	0	0	2.9	0	1	1	0.9	1
Time 3	1	0	0	2.9	0	1	1	0.9	1.9
Time 4	1.9	0	0	2.9	0	1	1	0	1.9
Time 5	1.9	0	0	2.9	0	1	1	0	1.9

Table 1: Simplified linear resource attraction model. Mean particle size= 1, range of variation= 0.1 and two links (i.e. one on each side).

often do not, statistically speaking, but to define and quantify the forces that push complex systems toward a polo distribution and how they function.

4 A Resource Attraction Model to Explore the polo attractor.

A resource attraction model (RAM) was developed for the exploration of the theoretical framework described and to test whether these assumptions would lead to distributions of agent sizes approximating lognormal or power distributions or both. The model would also allow the exploration of sensitivity to initial conditions and variables and to compare its dynamics with empirical data.

4.1 Description

The model is based on a one-dimensional linear topology, but in principle can be expanded to more dimensions. Resource ‘masses’ are situated along this line with their position defined by a single coordinate. Each resource particle will then have a distance (d) from other neighbouring resources. The coordinate line can be made circular to avoid edge effects or a buffer can be set up. Space and distance here are fundamental abstractions applicable to any system. For example, in a biological community, links between individuals of all species are to some extent a function of spatial distances, but often more importantly of temporal, energy, matter and information ‘distances’.

Following the theoretical postulates, the model assumes that masses attract each other and that the force of this attraction is a function of their mass and an inverse function of their distance. The number of links (L) between particles or agents is a function of the individual force of attraction of each particle (determined by its mass), and its distance to other particles. In theory, L is infinite as all particles can affect each other even at great distances. As only those links which are above a certain strength are important to the agent dynamics, a subset of links to nearest neighbours represents effective links, or L_e .

Given an initial distribution of the two variables, particle size (M =mass) and particle position (coordinates), the subsequent positions are calculated by the attraction. The attraction or pull is calculated by the gravity analogy of M_1M_2/d^2 , where d is the distance calculated between the particles. At each iteration, masses move according to the magnitude and direction of this pull. The movement of particles to new cells results in clustering to varying degrees, depending on variables, thus forming agents (table 1).

The simulation in table 1 leads to a frozen distribution in just 4 steps, with agents oscillating in some cases between two positions. At time 1 one agent (on coordinate 4) has formed by accretion. Particles in coordinates 6 and 7 have swapped over, both remaining at a value of 1. Circularity is provided by the particles on the extreme left interacting with the particles on the extreme right. The rapid freezing is similar to some physical systems such as the solar system where no substantial quantities of new matter and energy enter the system, compared to the amount already there.

A more realistic situation for dynamic biological systems can be achieved by adding a flow of resources. This can be done by adding and subtracting resource particles at each step distributed uniformly or at random and with a given degree of size variation. This rain of particles simulates a continuous flow of energy or matter that leads to an increase in organisation or maintenance of a characteristic pattern in dynamic systems. With this addition the system remains dynamic. Agents evolve and become extinct, yet the abundance pattern still revolves close to a polo.

4.2 Basic rules, inputs and outputs

New Mass calculation: Actual mass \times relative mass + mass attracted from neighbours + rain – losses

Attraction calculation: $M_1 M_2 / d^2$

Relative mass: $M_1 / \text{Sum of neighbouring masses}$

Particle rain size: random or fixed around value decided

Proportion lost: $M_1 \times \text{loss coefficient}$

The input variables that the operator can modify are:

- total mass of resources
- total universe size (i.e. coordinate space)
- spatial distribution of particles
- number of effective links
- magnitude of the distance exponent
- resource rain, form of distribution and calculation
- particle mass, mean and range (i.e. variability of resource flow)
- resource loss (proportion)
- number of iterations
- minimum viable agent size

Additional constants and variables can be added and modified to explore and fine-tune the model to particular types of systems. Such variations provide ‘vibrations’ that can tip agents out of one cell, catastrophes that can disrupt or destroy them, or slight changes that may help them capture new particles.

The output parameters include:

- mean and deviation of the frequency distribution
- abundance-rank pattern and goodness of fit to power or other functions
- frequency-abundance pattern and goodness of fit to lognormal or other distributions
- total mass
- total agents
- agent dynamic fluctuations over time




Figure 2: Time sequence of agent development. Linear model starting from an initial uniform distribution of particles of randomly varying sizes, with particle rain at each step. Mean particle size 1, range 0.2, total mass 200–264, 10 links. Dotted: $t = 5$, striped: $t = 10$, full: $t = 20$.

5 Results of Simulations

The dynamics of the simple rules described above allow the exploration of the rule-space which leads to diversification. Results vary according to initial conditions, but remain similar within a wide range of situations. For example, the model can simulate the emergence of diversity from an initial amorphous mass of particles with uniform spatial distribution and some (e.g. 20%) variation in particle mass. Particles start to cluster into agents. Larger agents capture neighbouring particles or whole agents (take-over). Smaller agents grow slowly or become extinct. Occasionally, new agents arise (speciation) (figure 2).

The abundance distribution of agents tends to a power function with increasing slope toward the right in a log-log rank-abundance relation or a lognormal (figures 2–4). Also as in natural systems modifying the variables can lead to a veiled or truncated lognormal, which is then almost identical to a power function (see references under §3.4).

Despite the apparently chaotic dynamics of individual agents' growth and decline, the polo distribution attractor is robust for the RAM. Abundance distributions trend toward this attractor for a range of different variables, in many cases after very few steps (5–20). The similarity of two events arising from different initial sets of conditions in the rank-abundance graph (figure 3), but the clear distinction of the same two in the frequency-abundance graph (figure 4, \square = dotted line of figure 3, \square = bold line), suggests that the lognormal has greater diagnostic capability than the power function to discriminate the variables in action. For example, in systems which are large enough in relation to the particles, the dip in slope at the right end of the rank-abundance distribution may become an independent lognormal curve, showing that the system has split into two systems (e.g. as shown in the case of macro- and micro-economic plants in New Zealand agriculture; [12, 14]. Such splits are clearly apparent in frequency-abundance representation but not as obvious in abundance-rank representation.

Any distribution which starts as an orderly spatial array (e.g. an arithmetic progression: 1, 2, 3, 4, 5...), ends up concentrating all resources on one coordinate; there is no diversification. Thus, the way the particles are distributed with respect to their links (i.e. space-distance) is of importance to the development of diversification. A completely uniform initial distribution (i.e. all 1, or all 2) remains frozen in that same condition.




Figure 3: Rank-abundance distribution of agents. Agent size is the mass of resources captured by that agent. bold line: with particle rain, mean particle size 1, range 0.2, total mass 200–259, 20 steps, 10 links giving 45 agents of mass > 2.8 ; dotted line: without particle rain, mean particle size 1, range 2, total mass 210, 10 steps (frozen), 10 links giving 33 agents of mass > 2.8 .

6 Discussion

The theoretical model and the simple simulation described in this paper are clearly distinct. The theoretical framework attempts to define the fundamental components and concepts relating to complex systems. The computer model is an extension of a ‘pencil and paper’ exercise to visualise the consequences of the theory. The model is only one of a range of different approaches to test the theory.

6.1 Links

Kauffman[17] suggested that in a Boolean network the relation between number of links and number of elements was critical to the system dynamics. Only intermediate values lead to polo distributions. Too few links would make the system vary at random within a normal distribution (i.e. particles clump in random group sizes according to an initial random distribution). Too many links would freeze the system whichever way it started (historically determined). The RAM tends very often to polo, so how is the ratio of effective links to number of particles (L_e/N) close to the intermediate values needed? In Boolean networks, the number of links is a discrete number imposed from the outside. In the RAM the magnitude of L_e is a continuous function of the force determined internally by the rule of attraction. Boundaries to agents and to the system arise at the distance where the internal attraction reaches a certain ratio to external attraction. The boundary encloses N particles or resource mass. In this way N and therefore L_e/N are set within certain bounds by the internal rule.

6.2 Relationality: Attraction, repulsion, competition, cooperation

The theoretical framework may help clarify some debated aspects of the behaviour of complex systems.

The rule for attraction is analogous to gravity in a physical system, reflecting the observable fact that agents attract or pull in resources to grow whether they belong to physical, biological or economic systems. The agent’s attraction is a function of its mass and distance. The laws of increasing returns and of diminishing returns in economics are some of the manifestations of this phenomenon in the practical world. The law of increasing returns expresses the positive feedback between the size (e.g. market share) of a company and its increasing ability to capture more resources (see [1]). The law of diminishing returns observes that as an agent accumulates resources from a particular source, the initial accumulation is cheap and easy, but as the resource is depleted and/or competitors arise,




Figure 4: Frequency distribution of agent sizes. Frequency shows the number of agents with a given resource mass or size. Size shows the resource mass in each class. ∇ : without particle rain, mean particle size 1, range 0.2, total mass 200, 10 steps (frozen), 10 links giving 37 agents of mass > 2.8 ; \square : without particle rain, mean particle size 1, range 2, total mass 210, 10 steps (frozen), 10 links giving 33 agents of mass

> 2.8 ; \square : with particle rain, mean particle size 1, range 0.2, total mass 200-259, 20 steps, 10 links giving 45 agents of mass > 2.8 ; \square : with particle rain, mean particle size 1, range 0.2, total mass 200-264, 20 steps, 10 links giving 50 agents of mass > 2.8 ; \square : mean of two last series.

capturing the resources becomes increasingly costly. In the terminology of this model the resource is initially clustered, proximate and dense, leading to a high attraction, but at a later stage more distant and dispersed, with a higher entropy, and therefore more difficult to attract.

The attraction force can also be modelled as a repulsion force, as one is the inverse of the other (e.g. [10]). Other phenomena arise from such a force, without having been explicitly included. The notions of competition and cooperation arise as a result of attraction rather than as incompatible theoretical options. Agents compete to attract resources in between them, and cooperate to attract resources on either side of both.

6.3 Hierarchy and diversification

An aspect of behaviour that can be inferred from the framework is that once evolved to a large size, an agent can, in turn, become the basis for further differentiation on a new hierarchical level, and levels can cross over. For example, socio-economic systems arise as agents (e.g. nation-states) which break up into a series of new agents (provinces on a spatial scale, companies and ethnic groups on an economic and cultural scale) or join into super-systems such as nation-blocks or alliances.

6.4 Different one-dimensional self-gravitating models

Physicists have developed one-dimensional self-gravitating systems to explore multi-bodied dynamics of attraction [9, 30, 39]. Such models differ from the resource attraction model in that particles accelerate until they either bounce elastically or pass through each other. There is no accretion into larger agents. The dynamics of diversification and abundance variations which occur in the resource attraction model do not occur in these physical systems.

6.5 Convergence with other models

It is no coincidence that many models produce polo distributions starting from many different explanations (e.g. [40], John Conway's Game of Life, [5, 4, 25]), as the trend results from similar fundamentals such as neighbourhood interaction or competition for resources. Bak [2] suggested that such models evolve to criticality, as do natural systems (e.g. [19]). Whatever the underlying principles, there is a

need to describe these in a way valid for all systems. For example, the analogies with fundamental quantum theory are intriguing (‘quanta’ of energy or mass, and attraction) and are may not be coincidental. These analogies reflect deeper mathematical laws (e.g. [35, 8]). The theoretical framework proposed here is one attempt to identify such fundamental rules in a way which can be tested across natural systems. The resource attraction model behaves robustly because the rules define the system at each level independently of the user.

6.6 Issues to explore

The resource attraction model allows the exploration of a variety of important questions concerning the behaviour of complex systems. For example:

- explore the limit to diversification or polo: are there limits to diversification or to polo? If there is a limit, what is it? are most systems at the limit? if not, why?
- explore the pattern attractor phase space by modifying rules and variables (e.g. the variables leading to particular slopes or veil lines, conditions for splitting and diversification, or the collapse of diversity)
- under what conditions is the system most stable?

7 Conclusion

A theoretical foundation is proposed to explain the features and mechanisms of power and lognormal distributions so widely found in nature. The framework is based on fundamental elements which are observable in nature (e.g. attraction, distance), which lead to the breakup of resources into new agents (speciation, diversification) and emergent mechanisms to set boundaries. The level of abstraction allows the model to be applicable to any complex system fitting the proposed definitions.

The model proposes that relatively simple rules lead to the emergence of polo distributions in complex systems starting from a range of initial conditions. The prevalence of polo distributions suggests that they could be used as management tools in conservation, agriculture, economy and other disciplines dealing with complex systems (e.g. to determine ecosystem health, to predict agricultural inputs, to pre-empt weed invasions or to determine the cost of an egalitarian society). The fact that it is not used is to some degree due to the lack of an appropriate theoretical framework, to a confusion between fundamental rules and noise, and to non-standardised language leading to lack of communication.

Acknowledgements

I am grateful to B. Barratt, K. Dodds and J. Lockwood for allowing me to sound out these theories and discuss drafts of the manuscript; also to the editor, Russell Standish and an anonymous referee for helpful comments. I would also like to thank E. Jen for inviting me to visit the Santa Fe Institute where further insights were achieved.

References

- [1] W. B. Arthur. *Increasing Returns and Path Dependence in the Economy*. University of Michigan Press, Ann Arbor, 1994.
- [2] P. Bak. *How Nature Works: The Science of Self-Organized Criticality*. Oxford UP, Oxford, 1997.
- [3] P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality. *Phys. Rev. A*, 38(1):364–374, 1988.

- [4] N.D. Barlow. Size distributions of butterfly species and the effect of latitude on species sizes. *Oikos*, 71:326–332, 1994.
- [5] E. R. Berlekamp, J. H. Conway, and R. K. Guy. What is life? In *Winning Ways for Your Mathematical Plays*, volume 2. Academic Press, London, 1982.
- [6] J.H. Brown and B.A. Maurer. Macroecology: the division of food and space among species on continents. *Science*, 243:1145–1150, 1989.
- [7] J.H. Brown and P.F. Nicoletto. Spatial scaling of species composition: Body masses of North American land mammals. *Am. Nat.*, 138:1478–1512, 1991.
- [8] M. Chown. Anything goes. *New Scientist*, 158(2137):26–30, 1998.
- [9] S. Cuperman, A. Harten, and M. Lecar. A phase-space boundary integration of the Vlasov equation for collisionless one-dimensional stellar systems. *Astrophys. and Space Sci.*, 13:411–424, 1971.
- [10] S. Douady and Y. Couder. Phyllotaxis as a physical self-organized growth process. *Phys. Rev. Lett.*, 68:2098–2101, 1992.
- [11] S. Frontier. Diversity and structure in aquatic ecosystems. *Oceanography and Marine Biology Annual Review*, 23:253–312, 1985.
- [12] S. Halloy. Long term trends in the relative abundance of New Zealand agricultural plants. *Otago Conference Series*, 2:125–141, 1994.
- [13] S. Halloy. Indicadores de perturbación de ecosistemas: Especies claves y estructura de biodiversidad. In M. Liberman and C. Baied, editors, *Desarrollo Sostenible de Ecosistemas de Montaña: Manejo de Areas Frágiles en los Andes*, pages 73–81. UNU y PL-480, La Paz, Bolivia, 1997.
- [14] S. Halloy. The dynamic contribution of new crops to the agricultural economy — is it predictable? In *Proceedings of the Fourth National New Crops Symposium*, Phoenix, Arizona, 1998.
- [15] S. Halloy and A.F. Mark. Comparative leaf morphology spectra of plant communities in New Zealand, the Andes and the Alps. *J. Royal Soc. NZ*, 26(1):41–78, 1996.
- [16] G.E. Hutchinson and R. MacArthur. A theoretical ecological model of size distributions among species of animals. *Am. Nat.*, 93:117–126, 1959.
- [17] S.A. Kauffman. *The Origins of Order*. Oxford UP, New York, 1993.
- [18] P.G. Kevan, C. F. Greco, and S. Belaousoff. Log-normality of biodiversity and abundance in diagnosis and measuring of ecosystemic health: Pesticide stress on pollinators on blueberry heaths. *J. Applied Ecology*, 34:1122–1136, 1997.
- [19] D.R. Lockwood and J.A. Lockwood. Evidence of self-organized criticality in insect populations. *Complexity*, 2(4):49–58, 1997.
- [20] R. H. MacArthur. On the relative abundance of bird species. *Proc. Nat. Acad. Sci. USA*, 43:293–295, 1957.
- [21] R.H. MacArthur. On the relative abundance of species. *Am. Nat.*, 94:25–36, 1960.
- [22] A.E. Magurran. *Ecological Diversity and its Measurement*. Princeton UP, Princeton, New Jersey, 1988.
- [23] B.B. Mandelbrot. *The Fractal Geometry of Nature*. W.H. Freeman, N.Y., 1983.

- [24] R.M. May. The dynamics and diversity of insect faunas. In L.A. Mound and N. Waloff, editors, *Diversity of Insect Faunas*, pages 188–204. Blackwell, Oxford, 1978.
- [25] C. Pahl-Wostl. *The Dynamic Nature of Ecosystems. Chaos and Order Entwined*. John Wiley, Chichester, 1995.
- [26] F.W. Preston. The commonness, and rarity, of species. *Ecology*, 29:254–283, 1948.
- [27] F.W. Preston. The canonical distribution of commonness and rarity. *Ecology*, 43:185–215, 410–432, 1962.
- [28] F.W. Preston. Noncanonical distributions of commonness and rarity. *Ecology*, 61:88–97, 1980.
- [29] C. Raunkiaer. *The Life Form of Plants and Statistical Plant Geography*. Clarendon, Oxford, 1934.
- [30] G. B. Rybicki. Exact statistical mechanics of a one-dimensional self-gravitating system. *Astrophys. Space Sci.*, 14:56–72, 1971.
- [31] R.L. Schmoeyer, J.J. Beauchamp, C.C. Brandt, and F.O. Hoffman, Jr. Difficulties with the lognormal model in mean estimation and testing. *Environ. and Ecol. Stat.*, 3(1):81–97, 1998.
- [32] J. Soderlund, L.B. Kiss, G.A. Niklasson, and C.G. Granqvist. Lognormal size distributions in particle growth processes without coagulation. *Phys. Rev. Lett.*, 1998. in press.
- [33] R. V. Solé and D. Alonso. Random walks, fractals and the origins of rainforest diversity. Working Papers 98-07-060, Santa Fe Institute, 1998.
- [34] G. Sugihara. Minimal community structure: an explanation of species abundance patterns. *Am. Nat.*, 116:770–787, 1980.
- [35] Max Tegmark. Is the “Theory of Everything” merely the ultimate ensemble theory? <http://www.sns.ias.edu/~max/toe.html>, 1997.
- [36] R. Tokeshi. Species abundance patterns and community structure. *Adv. Ecol. Res.*, 24:111–195, 1993.
- [37] G.B. West, J.H. Brown, and B.J. Enquist. A general model for the origin of allometric scaling laws in biology. *Science*, 276:122–126, 1997.
- [38] A. Wuensche and M. Lesser. *The Global Dynamics of Cellular Automata — An Atlas of Basin of Attraction Fields of One-Dimensional Cellular Automata*. Addison-Wesley, Reading, Mass., 1992.
- [39] K. R. Yawn and B. N. Miller. Ergodic properties and equilibrium of one-dimensional self-gravitating systems. *Phys. Rev. E*, 56:2429–2436, 1997.
- [40] G.K. Zipf. *Human Behavior and the Principle of Least Effort — An Introduction to Human Ecology*. Addison-Wesley, Cambridge, Mass., 1949.

The Variable Structure System: Intermittency, Chaos, and Trapping in Electronic Experiments and Simulations

D.J. Jefferies and J.H.B. Deane
 School of Electronic Engineering,
 Information Technology, and Mathematics
 University of Surrey
 Guildford GU2 5XH
 UK

d.jefferies@ee.surrey.ac.uk

<http://www.ee.surrey.ac.uk/Personal/D.Jefferies/>

Abstract

The concept of a *Variable Structure System (VSS)*, in which the structure is determined by the dynamical response, is described and measurements and simulations on VSSs comprising electronic systems which display *bursty chaotic* or “*intermittent*” behaviour, and also on a network traffic transfer protocol, which is demonstrated to be a VSS in the same sense, are used to support the thesis that such bursty behaviour is common in such systems. An example of trapping in a two-centre system is given to show that the ideas can be extended to continuous-variable dynamical systems having piecewise-linear properties. The studies of these simple electronic systems provide insight for cases where similar behaviour of a time series is observed in other complex systems; some other VSSs are listed and their properties are considered.

1 Introduction

1.1 Background

This paper extends some previous work reported [1] at COMPLEX96 and developed [2] at the Circuits Conference ECCTD97 in Budapest, on simple electronic iterating circuits containing traps and snags. The present paper also relates that work to the behaviour of artificial network traffic models reported [3] at COMPLEX94 and in a paper [4] in the International Journal of Electronics where the mathematics is presented in full. These systems, together with a piecewise-linear dynamical system comprising a two-centre harmonic oscillator, are interpreted as examples of a *Variable Structure System (VSS)*. New observations are reported, of VSS-induced blocking behaviour on networks and of the *Crisis-induced intermittency* which has been observed *by direct experiment* in the electronic circuits described [1] at COMPLEX96, and is presented here in detail for the first time.

1.2 Purpose of the paper

We study systems in which the structure of the circuit or system depends on the dynamical response. In the present paper, observations of behaviours in complex systems where *abrupt changes* in the character of the motion (or of the time series) are observed, are presented. The possibility of sudden changes of behaviour is probably universal in VSSs, and we examine various other systems each of which which may be classified as a VSS.

Sudden changes are sometimes termed *intermittencies*, and there is a classic report of a type of intermittency, resulting in the interleaving of laminar flow with bursts of chaos, in the papers of Pomeau and Manneville [6][7]. Here we have borrowed the term “intermittency” to describe other kinds of sudden change in the time series; in the case of the electronic example, there are bursts of extended chaos interspersed with periods of less extended chaos; and in the traffic simulations, there are sudden changes in the mean lifetime of packets on the net.

The bursty trapping is interpreted as *evolution* in the system dynamics, in the sense that, as time progresses, the dynamics change discontinuously. Since in a VSS, the system dynamics control the structure of the system, it follows that in a system which traps, the system can be *engineered* to evolve. This paper makes the suggestion that artificially contrived VSSs may be used as test-beds for studies of the time-development of complex systems.

A common reaction on being presented with an abrupt change of behaviour is to seek out the *immediate cause*. Often this is done with a view to removing the immediate cause and thus to smooth out the behaviour, or to “control the chaos”, in order to make it more predictable and tractable. In many cases control of irregular behaviour is not possible; the attempts to control the chaos by contingent changes in parameters merely serve to turn the system into a VSS which can then display even more irregular behaviour.

The experiments we have made on the disparate systems, electronic experiments and network simulations, and the impacting oscillator study, all give a view of such behaviour; that the irregularities grow naturally out of the dynamics, which can occur in a closed autonomous system and not be consequences of externally time-varying parameters.

1.3 Definitions

For the systems considered in this paper we introduce the concept of a *controlled switch*. In its simplest form, this can be a gate or CMOS switch fed by the output of a simple analogue operational amplifier working as a comparator, which senses the size of a voltage or current, or other measurable quantity from a *transducer*. The actual circuits we had constructed to display intermittent behaviour contain no controlled switches explicitly embedded in their implementation, but we show that they behave isomorphically with systems which do contain such switches. That is, with some increase in complexity, they could be replaced by circuits containing only ideal amplifiers and controlled switches. We emphasise that the controlled switches are introduced *in addition* to the other linear electronic components; these systems are therefore not analogous to a network consisting solely of Boolean gates or elements. The controlled switches introduce *piecewise-linear* non-linear properties to the system.

The piecewise-linear transfer function system simulated previously [1] was implemented experimentally using the natural saturation properties of an operational amplifier, such that when the output reaches 12 volts, it saturates, and a further increase in input results in no further increase in output. In other words, the “differential gain” has fallen to zero. Such an amplifier may be replaced by an “ideal amplifier” with unrestricted input and output ranges, which is disconnected from the output line when the input reaches the appropriate threshold. The output line is instead connected via another controlled switch to a constant 12 volt source.

Alternatively, the various linear sections of the piecewise linear transfer function may be produced by amplifiers having gains A_1, A_2, A_3 appropriate to the slopes under consideration, switched in and out of circuit (with appropriate offsets O_1-O_5) by controlled switches driven from the input. To demonstrate the equivalence of the saturating amplifier transfer function (Figure 4) circuit to the VSS explicitly, Figure 1 shows a “controlled switch” version of the circuit; this version has not been implemented as it is an unnecessary complication from a constructional point of view.

A controlled switch may be used either to alter the system state space trajectory in a discontinuous manner, (for example, by adding an offset to the output) or alternatively to alter the configuration or structure of the system.

If controlled switches are used to alter the circuit configuration or structure, we arrive at the idea of a *variable structure system* [8] or VSS. In a VSS the structure of the circuit or system is not time-invariant, but fluctuates according to the instantaneous values of the state vector. Since the state vector in turn varies according to the structure of the dynamical system, a loop is established in which the chaos can develop unpredictably; an example is given below showing trapping, in which the chaos stops altogether, and in the network traffic example the global behaviour of the state vector over a complex system consisting of many interlinked simple VSSs is shown to evolve with time.

In digital electronics a device much used to reconfigure a system of gates is the FPGA, or Field

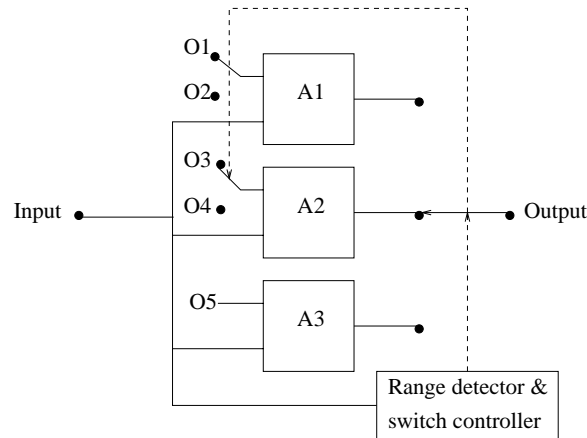


Figure 1: The transfer function circuit drawn explicitly as a VSS

Programmable Gate Array. However, we are not concerned only with restructuring the topology of *deterministic digital logic*, but also with altering the structure of *continuous analogue systems* which may contain noise which can push the signal which controls the switch across an analogue threshold. Thus there is the possibility of stochastic development as well as deterministic chaos in some of these VSSs.

In the circuits and systems described here, *feedback* is applied around a structure containing embedded controlled switches or their equivalents, as well as containing other logic or circuit components. These equivalents are shown explicitly by means of circuit block diagrams. For the simple electronic example, the feedback takes the form of a two-stage sample-and-hold analogue shift register, which transfers the output value to the input on a clock pulse and then acquires the resulting new output value for transfer on the next clock pulse. This is an *iteration circuit*. Such a system is a combination of the discrete and the continuous. Iterations are discrete but the variable being iterated is continuous with its intrinsic noise.

Thus, *adaptive behaviour* may be engineered in systems containing controlled switches, and in variable structure systems generally. Such behaviour need not be *cyclic*, repeating exactly, but can be *emergent* and result in progressive modifications to the system. For those people happier with biological terminology, we have the potential to apply *selection* to a fluctuating system, resulting in *evolution*. Here, our use of the term “evolution” implies no special adaptation to “fitness for purpose” or other measure of utility, but merely to the fact that the dynamics change as time progresses, and such changes can sometimes happen non-reversibly.

2 VSS examples — 1: Electronic iterating circuits used to demonstrate “crisis-induced” intermittency

A previous paper [2] introduced the behaviour of simple low-dimensional non-linear mappings under iteration. The mappings are designed to contain small *Features* which we called *traps* [1] or *snags* [2]. The trapping behaviour was discussed and reported [1] theoretically and by simulation in our paper at COMPLEX96. Here, we present the supporting experiments.

The Features give rise to the following behaviours: If the state variable enters a trap, the chaotic time series ceases, and the system remains at a fixed point. Thus a trap is an attracting fixed point of the system which is the end point of a chaotic transient; this can, in principle, last an arbitrarily long time. Normally, one expects some kind of progressive approach to a fixed point. In our systems, on the other hand, trapping is sudden, and happens without warning.

A snag consists of a chaotic attractor embedded in the larger chaotic environment. In the one-dimensional case considered in the earlier paper [2], there is a small probability of entering or leaving the snag from the larger domain. Thus the motion appears always chaotic, but “bursty” with clearly visible (Figure 5) differences between the bounds of the two competing attractors. Again, the transitions are sudden and happen without warning. In that paper, we attributed the bursty behaviour to the interaction of the added noise in the simulation with the features. Below, we show clear experimental evidence of true intermittency brought about by the overlap of the chaotic attractors. In the two-dimensional experiments presented here for the first time, it is not so easy to picture the structure of the 2-d features and so the experiments provide a valuable insight into the dynamical possibilities.

The term “crisis” refers to a point in the gradual change of a parameter which causes the snag to grow or move until its structure extends beyond the feature, and ejection (Figure 8) into the wider chaotic attractor becomes possible.

2.1 The iteration system

To recap, the basic one-dimensional building block for the two-dimensional system we have devised consists of a two stage analogue shift register providing feedback around a *transfer function* circuit which generates the mapping for the iteration. The transfer function circuit consists of an amplifier having different gains for different ranges of the input voltage. It is a *piecewise-linear* mapping of the input voltage on to the output voltage. The possible values of input and output voltage occupy the same overall span. Thus when the output is transferred to the input by the shift register, repeated transfers do not result in the voltages going out of the overall span. This is a classic chaotic system displaying *stretching* (gain size greater than unity) and *folding* (multiple values of input for each output). In the implementations we have made the maximum gain has size about 3.

The mapping lies in the interval $(0, 1); (0, 1)$ with four straight line sections. It is sketched in Figure 2. The most important point to note is that we have placed a central square box such that the iteration line (of unit slope) passes through diagonally opposite corners. Thus the wider square and the small square box may both be considered to be individual autonomous one-dimensional iterating mappings. If the motion passes from the large box to the small box, it will stay within the small box unless there is a method of ejecting it. This can be either by the addition of noise, or by extending the embedded *Feature* inside the small box so that it has sections which lie outside the small box boundary. As the Feature is enlarged, a crisis point occurs when the size is just sufficient that the mapping within the box touches the side. Inside the centrally placed square box of side ϵ (which we term *The region of the Feature*) is placed the Feature, which can be either a trap (fixed point or limit cycle) or a snag (chaotic attractor). The snag can be a *chaotic trap*. The size of the Feature and its region can be made arbitrarily small. The mapping is described mathematically by the following set of equations:-

$$f(x) = \begin{cases} 1 - 3x & 0 \leq x < \frac{1}{3} \\ \frac{3-3\epsilon}{1-3\epsilon}(x - \frac{1}{3}) & \frac{1}{3} \leq x < \frac{1}{2} - \epsilon \\ F(x) & \frac{1}{2} - \epsilon \leq x < \frac{1}{2} + \epsilon \\ 1 + \frac{3-3\epsilon}{1-3\epsilon}(x - \frac{2}{3}) & \frac{1}{2} + \epsilon \leq x < \frac{2}{3} \\ 3 - 3x & \frac{2}{3} \leq x < 1 \end{cases} \quad (1)$$

Here, $F(x)$ is the mapping of the Feature.

2.2 Features

A trap and a snag are shown in Figure 3.




Figure 2: The mapping containing the Feature.


Figure 3: A *trap* and a *snag*.

Figure 4: The mapping leading to crisis-induced intermittency.

If the discontinuities (turning points) in the snag overlap the boundaries of the square box, as in Figure 4, then the chaotic motion within the snag will occasionally escape into the larger region where it will stay until recapture.

2.3 Crisis-induced intermittency

In the case referred to in the last section, where the points of the snag just touch the sides defining the region of the Feature, a crisis occurs and arbitrarily small amounts of noise can eject the system from the Feature. If the points extend beyond the Feature, there will be a probability of escape even though there is no added noise. The motion will then consist of periods of chaos within the Feature, interspersed with intermittently occurring periods of motion outside the Feature. This behaviour we call *crisis-induced intermittency*. The intermittency sets in suddenly at the crisis point in the case of a hypothetical system with no added noise. However, in a real system with Gaussianly distributed added noise, the crisis point is less well defined as the snag size is increased. Other forms of snag are possible; here we have restricted ourselves to an easily implemented version. The division of the mapping into the Feature region and the wider attractor allows design of various kinds of snag, which when they extend beyond the Feature allow for adjustable probability of escape.

2.4 Simulation

Simulations of the trapping process were given in the COMPLEX96 paper [1]. To follow the evolution of the 1-dimensional system dynamics, in the presence of added noise, computer simulations of the iterating system were made with artificial additional Gaussian noise, g_n , having zero average, and standard deviation σ . The mathematics of the statistics of times between capture and release events have been reported [2].

2.5 Circuits

It was thought that a two-dimensional version of the simple iterating circuit should be investigated. Simulations are a little more difficult to think about so it was decided to conduct experiments. The circuits were designed and constructed from operational amplifiers, having differing adjustable gains,




Figure 5: Time series of 1-d crisis-induced intermittency




Figure 6: Time series of one channel of a 2-d crisis-induced intermittency

and with differing output ranges before saturation. In each of the two channels, one for each dimension, there were three saturating operational amplifiers whose outputs were combined in a summing amplifier. Control of the gains and the sizes of the Features was by ten-turn accurate potentiometers with precision dials to allow parameters to be recorded and reset. It was arranged to have offset controls in each channel for the position of the Features.

Two two-stage sample-and-hold circuits were used to transfer the outputs of the transfer function generator back to the input, on clock pulses provided by an external generator. A sum and difference *rotation matrix* circuit was used to mix the two channels thus generated, for the 2-dimensional experiments reported below. In order to keep the range of the input voltage and the range of the output voltage commensurate, the matrix circuits add and subtract 0.5 of the outputs of the individual channels. For a true “area preserving” matrix we would require a rotation such that 0.707 of the outputs were taken to add and subtract. Our matrix therefore rotates and contracts, by an area amount of a factor 2. However, since the individual gains of the transfer function circuits are close to magnitude 3, a small area expands by a factor 9 (3 in x times 3 in y) on passing through the transfer function part of the system so the loop area gain is $9/2$. Thus we can use this method to construct a chaotic system of arbitrarily large dimension. This is best implemented by using the transfer function circuit sequentially. In a 4-dimensional system one would store a 4-vector at input and output in individual 4-wide transfer and hold circuits.

2.6 Experiments

Five oscilloscope photographs of experiments which extend the 1-dimensional simulations reported [1][2] previously are presented below. It is important to confirm the results of chaotic simulations by real experiment, as it is not always clear what the effects of noise and imprecisions in the real electronic implementations will be. Therefore we have gone to some trouble to produce this experimental evidence to support our previous “thought experiments” and simulations on this system.

First, Figure 5 shows an experimental one-dimensional time series of crisis-induced intermittency from a snag that just extends beyond the region of the feature. This is a quite distinct phenomenon from the “noisy trapping” simulation presented [1] earlier, but nevertheless looks very similar. The release from the chaotic trap does not require added noise in this case. Since a single transfer function generator and shift register is used, it is noticeable that the snag recaptures the motion very quickly after its escape. This happens because the proportion of system space occupied by the snag is not particularly small.

Second, Figure 6 shows the time series for the 2-dimensional system, consisting of the two shift registers, the two transfer function generators, and the rotation matrix circuit. Here we observe that it takes significantly longer for the snag to recapture the motion. This is because the snag of size ϵ has area ϵ^2 which is of second order of smallness in the 2-d system. We also notice a phenomenon of *incomplete escape* from the snag which happens from time to time. Here “experimentalist’s license”




Figure 7: Motion inside a 2-d snag




Figure 8: leakage from a snag in crisis




Figure 9: A 2-d snag and the wider attractor; intermittent behaviour

has been used to adjust the overlap of the snag with the region of the feature in order to capture a telling picture. Size and offset controls have both been used. It is in this kind of example that the superiority of direct experiment over simulation is demonstrated.

Next, we show a photo of the average position in the 2 dimensions (Figure 7) for the case of system motion wholly contained within a 2-dimensional snag. Here, the snag has no part which extends beyond the region of the feature.

Next, a photo of the escape (Figure 8) from the 2-dimensional snag whose size and offset have been both adjusted to provoke a crisis. Incomplete escape can be seen, as can the nearby motion on part of the wider attractor.

Fifth and last, a photo (Figure 9) showing the time-average position of the 2-dimensional motion encompassing the entire space which contains both the intermittent snag and the wider attractor. If one can imagine the dynamics behind this picture one gets a good idea of what a crisis-induced intermittency looks like in more than one dimension.

3 VSS examples — 2: Two-centre and impacting systems

Experiments have been made on an electronic implementation of a two-centre harmonic oscillator. This oscillator implementation is shown in Figure 10 and consists of a simple two-integrator loop with damping, which models a simple damped linear second order system. There is a refinement in that the offset of the centre of the system is switched from $+X$ to $-X$ when the *displacement* variable x passes downwards through zero. The opposite happens when the displacement returns and passes upwards through zero.

This two-centre system is very similar to a damped impact oscillator, where the velocity changes discontinuously when the displacement reaches the wall position X . The impacting system models a mass-spring-damping system where the mass bounces elastically on a rigid wall. There is a complete isomorphism between the impacting system, and an adaptation of the two-centre system in which the driving function is inverted every time the displacement passes through zero. Thus we have two more

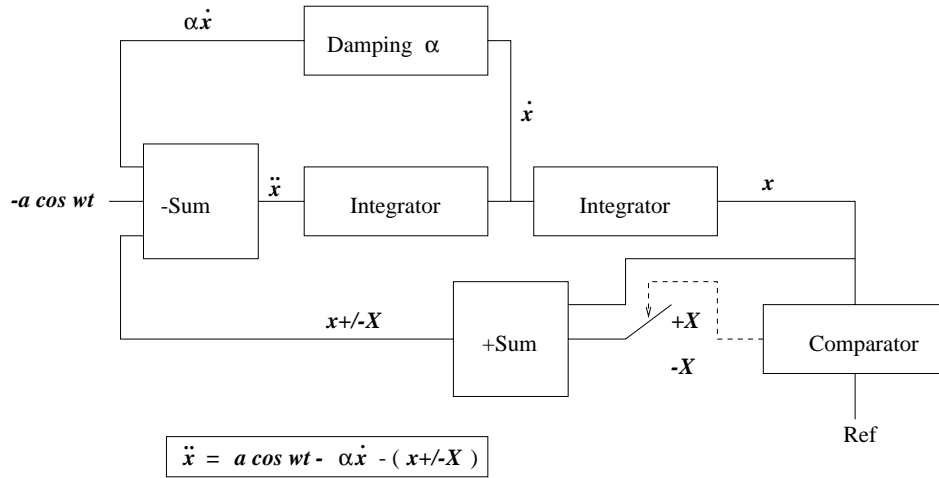


Figure 10: A two-centre VSS equivalent circuit diagram

VSS circuits to consider; intermittency has been observed in computer simulations and theoretical studies of the impact oscillator [12].

When, in the impacting system, the velocity discontinuously reverses, the time development of the linear part of the motion may be regarded as being suddenly advanced. This is not an exact modelling, but it explains the observation that when the drive frequency is adjusted to lie a little below the non-impacting resonant frequency of the system, the impacts can keep the motion more in phase with the drive and the phase plane portraits are on average larger than they would be if the wall were moved away to larger X . Thus the impacts are sustained until, happenchance, the chaotic fluctuations lead to an impact or impacts being missed, and the motion collapses onto the limit cycle of the linear system.

Since the two-centre system only differs from the impacting system by lacking an inversion of the drive, the phase relationship in the case of the two-centre system is reversed, and the orbits are, on average, larger for the drive frequency a little *above* resonance.

Thus we see in Figure 11 a time series development in which the chaotic fluctuations of the two-centre system eventually result in the subsequent zero-transitions not occurring and the system then falls onto a limit cycle (see the phase plane portrait for this system in Figure 12) on one side of the displacement origin at $x = 0$. In the analogous case of the impact oscillator, this is seen when the orbits no longer cross the wall position, and impacts then cease.

This behaviour is a classic case of trapping in a VSS. By adjusting the amplitude of the non-impacting limit cycle, it can be arranged for the average time to trap to become very large. The experimental photo shown in Figure 11 was the tenth exposure in a continuous sequence during which the system did not trap. The time span across the photo is 35 seconds; in this case the system took about 10 minutes to trap at an angular drive frequency of 10 radians per second. The size of the limit cycle is about 3/4 of the amplitude at which switching occurs. Thus we see that trapping can be arranged to be improbable; what is certain however is that for these conditions, this system will *always* trap if one waits long enough.

4 VSS examples — 3: Network dynamics - computer simulations

Here we present an example of bursty statistics from a different class of VSS in the area of irregular traffic flows on regular networks. First we briefly recap the system [3][4]. A square grid of N by N




Figure 11: Trapping in the system of figure 10; time series. The axes are time, horizontal; displacement, vertical.




Figure 12: The phase plane portrait of trapping in the two-centre system. The axes are displacement, horizontal; velocity, vertical.

cells is scanned from the NW corner to the SE corner in a raster (or other sequence) by a token, which alights on each cell in turn and enables development to occur. The grid is populated by packets numbering a fraction f of the N^2 cells; the packets have set out from transmitters on the W and S boundaries to make their way to receivers at the E and N boundaries respectively. The fraction f may be set by the program. When a receiver collects a packet it immediately returns it towards the transmitter whence it came without removing it from the grid. Thus f is a true constant of the motion and does not change. Packets contain a record of their intended destination. There can be only a single packet in a cell; thus a packet cannot move into an occupied cell.

If the token alights on a cell containing a packet, the packet attempts to move in the direction of the receiver, or as close to that as possible according to a simple protocol (see Figure 13). If it cannot move in any direction it stays put. The token then passes to the next square. This is the only contingent protocol in the system, in the sense that what happens next is dependent on the arrangement of packets in the neighbourhood of the token cell, and also on the direction to the intended receiver. The protocol may easily be implemented in hardware by a VSS, as can be seen in the figure (Figure 13). Here, the individual cell sites contain the hardware switches as shown, and provide the protocol for onward routing of the itinerant packets. It would also be possible, at some increase in complexity, to regard the itinerant packets as having hardware attached which determines their onward paths. The concepts developed schematically in the figure provide, in our opinion, a firm demonstration that our data network is a true variable structure system, in exactly the same sense as in the other examples.

Here, the controlled switch selects each of the neighbouring cells in turn, in the order selected by the global variable which has been set by the programmer. On finding the first empty cell, the switch moves the packet into the empty cell leaving the source cell vacant. We have not gone into the precise details of the electronics needed to do this; nevertheless it is a true variable structure hardware arrangement.

So the network study is a simulation of another kind of variable structure system. The state vector, of dimension fN^2 , consists of the positions and attributes of the pattern of populated squares on the grid. It evolves with time, and a single measure of the time development is taken to be the average lifetime of all the packets on the net since the last time they were reflected from a receiver. It is this variable which is plotted on the vertical axes in Figure 14, Figure 15, and Figure 18. The packets have to move in one of the directions set by the protocol, which is perfectly deterministic. However, the system is sufficiently large [3] that individual global (as opposed to local patterns in the region of the token) patterns will most probably never reappear in reasonable amounts of runtime. Thus the pattern and attributes of the occupied cells governing the development of the state vector between token passes almost certainly never repeats. It is truly variable and chaotic, and is a direct consequence of the large number of possible values of the state vector. The variable structure is also represented by the global patterns of packets on the net, affecting the development directly. The

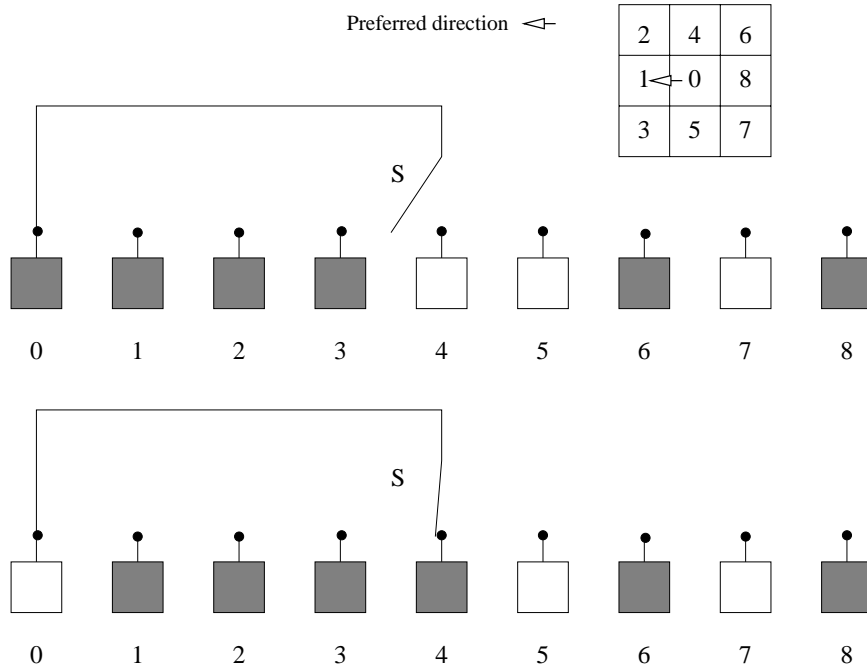


Figure 13: The local VSS for the network protocol

global development is the aggregate of N^2 decisions set by the individual VSS as described above.

Earlier papers [3][4] have investigated simple models of the network dynamics. Preliminary investigation of the phenomena reported there shows that traffic flows on the model network, whose structure is static and which supports constant traffic demand, displays bursty and intermittent behaviour for certain values of f ; a further example of the statistically self-similar time series produced is shown in Figure 14.

Another type of behaviour in this system is approximately periodic, showing cyclic oscillations between free flow and blocking behaviours. An example of the time series is shown below (Figure 15).

Clusters of packets form and dissolve with time, resulting in a *global* variable structure of clustered packets which impedes the traffic flow. This is a kind of collective variable structure which has arisen out of the aggregate of the behaviour of all the individual controlled switch decisions described above for the individual token cells.

This is illustrated in Figure 16 and in Figure 17 which show, respectively, clustering (in a region of impeded flow) and less-impeded flow, with the points on the time series diagram in Figure 14 labelled with arrows.

An even more striking example of clumping and free flow is provided in the time series (Figure 18), together with the blocked packet distribution (Figure 19) and the subsequent unblocked packet distribution (Figure 20), for a 30 by 30 grid with 39 percent of the sites occupied by packets. The qualitative behaviour seen here is universal and does not depend on details of the size of the grid or the position of the receive and transmit sites. We always see critical levels of loading at which bursty behaviour or periodic behaviour sets in.

It is interesting in this example that local VSS behaviour has resulted in a kind of *stochastic* variable structure across the network. We are observing collective behaviour of the dynamics resulting in global variable structure, that we did not set out with the intention of engineering when we wrote the protocol. Of course, there are many conditions for which the intermittent behaviour is not observed and the global VSS does not emerge; nevertheless the local implementation of the VSS always pertains.




Figure 14: Intermittent traffic blocking due to network dynamics




Figure 15: Cyclic traffic blocking on the network at a different load factor to that in figure 14




Figure 16: Packet bunching impeding the traffic flow at a peak in figure 15




Figure 17: Less packet bunching at a point of easy data flow at a minimum in figure 15




Figure 18: Packet blocking statistics, detail on a 30 by 30 grid with 39 percent loading




Figure 19: Packet bunching at a maximum latency (packet blocking) in figure 18.




Figure 20: The normal packet distribution for unblocked flow in figure 18.

5 Discussion

The variable structure scenario is of wide applicability. One might argue that many non-linear chaotic dynamical systems could be cast in the form of a VSS. While it is possible to construct a regularly behaved VSS, this VSS scenario does give us a way of approaching the engineering and construction of arrangements for producing irregular and bursty behaviour. An important and increasingly studied VSS is the Cellular Neural Network [13] described by Chua and Yang in 1988. This has important applications to signal processing and artificial vision, among others. Chua has suggested that CNNs are a test bench for the evolution of complexity, which he suggests emerges “on the edge of chaos” when the CNN is set up correctly. The non-linear elements are linked together by coupling and switching elements under software control. Various training algorithms are suggested for particular tasks. VLSI CNN chips are available and would be good candidates for studying the evolutionary behaviour of VSSs.

As we have seen, iterating electronic circuits, non-linear dynamical systems, and recursive network algorithms provide three different kinds of implementation of a VSS. In the case of the two-centre system with its integrators, the delay and memory through the integrator circuits provides the equivalent of the data propagation through the sample-and-hold circuits of the iteration system. In the case of the network example, the previous structure of the cells is stored in computer memory and dictates what happens on the subsequent token pass. Thus, to these authors at least, all these examples are transparently equivalent.

There is a good discussion of the stability of chaotic processes, and of the behaviour due to *interior crises* of the attractor in dynamical systems, in a paper by Kautz [14]. He says “when the unstable periodic orbit defines an interior basin, increasing α beyond α_c simply allows the attractor to expand into other parts of the basin of attraction.” At a critical value α_c of the bifurcation parameter α , the chaotic attractor collides with an unstable periodic orbit. Now it is a simple matter to arrange, electronically, for the parameter α to depend on a measure of the average size of the chaotic trajectory which is visiting a part of the attractor, and thus to engineer evolution.

There is also a sense in which simple electronic circuits, and their simulations, which rely on perfect diodes for their non-linear properties, may be modelled by VSS. A good example may be found in the two coupled nonlinear LC circuits of Wada [15] et al, in which they observe what they term *blowout bifurcation* and bursty behaviour similar to our observations, in their Figure 9.

Other examples may come to mind. For example, it is possible that the unpredictable and irregular behaviour of certain classes of computer software may be interpreted as being due to the variable structure in the processes running on deterministic hardware; such processes may interact and retain memory of their most recent invocation in a complex and unpredictable way, and be modellable in terms of a VSS in the software domain. The structure of the software (the bit-image in machine memory at startup) depends on the dynamical history such that even if the user tries to run the software on two successive occasions in *exactly the same way*, the bit image has changed, and on subsequent runs the behaviour is different.

6 Conclusions

Three different electronic systems have been studied either in simulation, or experimentally, or both. They all may be made to display sudden changes of the chaotic motion which we have interpreted in the light of the Variable Structure model. It is just possible that by finding universal properties such as this in disparate but structurally similar systems, ideas for the engineering of complex systems to either display or avoid such behaviour may emerge. It may even be possible to engineer such systems to display emergent or evolutionary properties.

References

- [1] Jefferies, D. J.; and Crawford, M. J.; "Traps in chaotic systems", in *Complex Systems 96 proceedings*, IOS Press, pp 253–260, 1996.
- [2] Jefferies, D. J.; and Deane, J. H. B.; "Noise, Traps, and Snags in iterating systems", in *Proceedings of ECCTD97, Budapest, September 1997, Technical University of Budapest*, pp 1014–1018, 1997.
- [3] Deane, J.H.B.; Jefferies, D.J.; and Smythe, C., "The dynamics of deterministic data networks", in *Complex Systems: Mechanisms of adaptation COMPLEX'94 Conference Proceedings, IOS Press 90 5199 186 X*, Queensland, Australia, pp 345-352, 1994.
- [4] Deane, J. H. B., Smythe, C., and Jefferies, D. J., "Self-similarity in a deterministic model of data transfer", *International Journal of Electronics*, vol. 80, no. 5, pp 677–691 (1996)
- [5] Jefferies, D.J.; "Crisis-induced intermittency in simple electronic circuits; experiments and simulations" to be published in *proceedings of NOLTA98, Crans-Montana, Switzerland, September 1998*.
- [6] Pomeau, Y., and Manneville, P.; "Intermittent transition to turbulence in dissipative dynamical systems." *Communications in Mathematical Physics*, vol. 74, pp 189-197, 1980.
- [7] Berge, P., Pomeau, Y., and Vidal, C.; Chapter 9 in *Order within Chaos*, John Wiley, pp 223-259, 1984.
- [8] Chen, Yon-Ping, and Cheng, Mu-Huo,; "Design of sinusoidal oscillators using the theory of variable structure systems", *International Journal of Electronics*, vol.72, pp 143-149, 1992.
- [9] Jefferies, D. J.; "The double potential well circuit; properties, simulation, and construction." in *Proceedings of the workshop NDES96 (Non-linear dynamics of electronic systems)*. Seville, Spain, pp 327-332, June 1996.
- [10] Hindmarsh, M. B.; and Jefferies, D. J.; "On the motions of the offset impact oscillator." *Journal of Physics A (Mathematical and General)*, vol.17, pp 1791-1803, 1984.
- [11] Budd, C., and Dux, F.; "Chattering and related behaviour in impact oscillators." *Phil. Trans. Roy. Soc. Lond. A*, vol. 347, pp 365-389, 1994.
- [12] Budd, C., and Dux, F.; "Intermittency in impact oscillators close to resonance." *Nonlinearity*, vol. 7, pp 1191-1224, 1994.
- [13] Chua, L.O., and Yang, L.; "Cellular Neural Networks: Theory and Applications." *IEEE Transactions on Circuits and Systems* vol. 35, pp 1215-1290 (1988). Also, Chua, L.O.; "CNN: A paradigm for complexity" book published by World Scientific, ISBN 981-02-3483-X, June 1998.

- [14] Kautz, R.L.;; “Global stability of the chaotic state near an interior crisis” *Structure, coherence and chaos in dynamical systems, edited by Christiansen, P. L., and Parmentier, R. D.*; publisher Manchester University Press, London, pp 207-225, (1989).
- [15] Wada, M, Nishio, Y, and Ushida, A.; “Blowout bifurcation and riddled basin in simple coupled chaotic circuits” in *ECCTD97 proceedings, Budapest, ,* volume 3 pp 1269-1273 (1997).

Is There Meaning In Fractal Analyses?

Herbert F. Jelinek
 Charles Sturt University,
 School of Community Health
 Albury, NSW 2640
 hjelinek@csu.edu.au

Cameron. L. Jones
 Swinburne University of Technology
 Centre for Mathematical Modeling,
 School of Mathematical Sciences
 P.O. Box 218 Hawthorn, Victoria 3122
 cjones@swin.edu.au

Matthew D. Warfel
 Mennonite Central Committee — Bolivia
 Casilla 59 Camiri, Bolivia
 caritascam@cidis.scbbs-bo.com

Abstract

We aim to clarify several fundamental terms used in fractal analysis and examine how the estimates of the fractal dimension can be made clearer to best serve as descriptive indices. The problem is essentially one of clarifying the semantics of the term ‘fractal’, since the syntax of calling something ‘fractal’ is often used with little regard to the principles underlying scaling theory. Towards this aim we discuss the use of language and the necessity to establish a linguistic base that serves as the template for communication across different disciplines.

1 Introduction

Our position deals with the use and understanding of language from a pragmatic standpoint. That is, we discuss the misconception of some vital definitions and terms in the field of fractal analysis. The paper is not about how we come to know but rather how we should communicate in order to gain knowledge and understanding. We offer practical solutions by identifying several terms associated with fractal analysis that need to be clarified. These terms, such as ‘fractal’ have been socialised and come to mean something in the literature that may be misleading and from an applied perspective not useful. Further, we clarify steps in the fractal analysis procedures, concentrating on the box-counting method as variations in sampling and preparing images for analysis and the analysis procedure itself can have non-trivial effects on the estimation and interpretation of D [3, 8, 12, 16, 19].

Applications of fractal analyses have been extensively used in diverse scientific, sociological and philosophical areas of research. Despite this large volume of literature, there is still a lack of clarity regarding the meaning of the terms used this field and therefore inferences drawn from the results are questionable[3, 18, 12, 19].

The concept of ‘fractal scaling’ is revolutionary but not new in that it offers a means to describe how things ‘are’ in terms of the object’s scaling characteristics[11, 21]. What we mean here is, that D quantifies ‘something’ in terms of shape, texture, size, number, colour, repetitiveness, similarity, randomness, regularity, heterogeneity, or any other adjectival descriptions that are commonly used to define the properties of some object or event. These descriptions can not be quantified using Euclidean geometry that idealises form. The properties mentioned above reflect the complexity of the object or event. Fractal analysis has provided a means of quantifying these properties as a measure of complexity or scale-dependency of the pattern. This is not to say that fractal analysis is the only means of quantifying complexity. Other analytical methods include Fourier analysis, fractal harmonics, polygonal harmonics and wavelets[13, 26]. In principal, fractal analysis should improve on the description of morphological features compared to conventional shape parameters. Common examples from the literature include heart rate irregularities, grazing effects on pastoral lands or stock market fluctuations[2, 16, 17]. It is our observation however, that conclusions drawn from fractal research remain at best tentative — with some research areas offering more conclusive results[1, 4].

The lack of conclusive results can, in many cases, be explained by the apparent lack of a linguistic base, that is a sound description of fractal theory and its relationship to the associated analysis procedures. Fractal research and discussion is characterised by the repetition of definitions and procedures that were intended to be vague[17].

Fractal analysis measures length as a process and therefore is defined as a limit which allows the image (if fractal) to be analysed in arbitrarily high resolution[22]. In practice any object is represented by a finite data set and the measurement is restricted to a finite magnification range. The image can be interpreted by D if it is assumed that the finite data set of the image reflects a fractal set and is self-similar[20]. If it is assumed that the image does not reflect an ideal fractal in a statistical sense (this is the case for biological images), than interpreting the image using D is meaningless. The fractal dimension may still be useful though by using it as a quantitative parameter that indicates complexity or the scale dependence of a pattern[14]. This fundamental concept is not made clear in the literature where D is taken to indicate fractality[19]. Communication becomes meaningful if all involved understand the terms. That is a transparent linguistic base exists.

2 Communication

A fractal set is a set in metric space for which the Hausdorff-Besicovitch dimension $D >$ topological dimension D_T . [17, p361]

The above quote is the most often quoted sentence found in journal articles to describe fractals, even though Mandelbrot states, on the next page, that this definition is rigorous, but also tentative. The definition is an example of communication that requires linguistic literacy as it requires an understanding of what is meant by metric space, Hausdorff dimension and topological dimension. It does not help in the understanding of fractal theory nor how this relates to fractal analysis. It would be more appropriate to point out that the definition applies to theoretical fractals and may be totally useless in practice if the image to be analysed does not reflect, albeit even statistically, a fractal set. In practice, the estimate of D quantifies scale invariance over a limited scaling range and does not indicate whether the image is fractal or not.

We suggest that communication occurs at different levels and that the literature can be divided into three main categories.

1. Theoretical mathematical research (those in the know who already possess a solid knowledge base in fractal and scaling theory). That is, those to whom the above definition by Mandelbrot means something quite concrete.
2. Applied science research such as the biological and social sciences (those with a solid knowledge base in other disciplines and using fractals).
3. Methodological analysts (those that develop tools for fractal analyses and aim at describing these clearly).

Once this division of the scientific/professional community is established it becomes obvious that a type of professional socialisation has taken place. Specific language used by each group establishes an identity within this group and marks it off as a specialist domain of knowledge and expertise[5]. To apply fractal analysis successfully it is important to obtain fractal literacy. The lack of clarity lies in the fact that many in the first category, being entrenched within a specific linguistic boundary and possessing fractal literacy, find it hard to communicate the necessary information to researchers in a different research field with a different subject literacy. Papers in the second category should contain research results obtained by applying fractal analysis correctly. As there are no guidelines for the 'correct' use of fractal analysis this is not possible and leads to errors in interpretation[3, 6, 19]. Fundamental concepts are often explained with little consideration of the diverse linguistic base of the audience and reflect the accepted meaning of terms within a specialised group. It is our assertion that




Figure 1: The Koch curve displaying the iteration process over several generations.

category 1 researchers have an obligation to disseminate their work to as wide an audience as possible. Significant studies in fractal theory have already occurred, yet such cumulative efforts must at some point become mainstream knowledge for all scientific disciplines, and eventually common knowledge.

For language to be a resource[10], with a potential to create meaning, it is important that novices are able to obtain the appropriate linguistic skills. Procedures used to determine the fractal dimension of images need to be made explicit. Some aspects of fractal analyses are like a black box in that one obtains a program, loads the image and obtains a fractal dimension. For the fractal dimension to be meaningful, the black box needs to be opened.

3 A Language Base for Novices

What then is a fractal? Fractals are characterised by a scaling law that relates two variables: the scale factor and the object being measured. This scaling relationship is described by a power law, which in turn describes the inherent physical attributes of the object being analysed[15, 27]. The exponent of the power law refers to the dimension. From our fractal linguistics perspective, fractal images are strictly limited visualisations of recursive equations that cannot be represented to the limit as an image on a computer. The Koch curve shown in Figure 1 is not a fractal, as the iteration level is not depicted at the level of infinity and has been referred to as prefractal[7].

A fractal object is the result of repeated transformations of a geometrical figure that leads to self-similar patterns. This self-similarity is a function of the scale invariance observed within the pattern between successive transformations. Viewing a fractal would therefore reveal identical patterns at different observation scales. In fact an important definition of fractals is that this self-similarity continues indefinitely.

Biological forms are not characterised by identical patterns at different scales and are statistically scale-invariant over a limited range[23]. Therefore the statement that a biological structure is fractal or not fractal is meaningless[19, 24]. The feature that makes fractal analysis interesting is that self-similarity can be described by a scaling function such as the power law. Representations of biological

objects, irrespective of whether they are represented as outlines, stick models or silhouettes can be analysed for their global scale-invariance and quantified by the fractal dimension. At this stage many researchers may not even have considered whether an image is fractal or whether, for their purposes, it is even important that the image is fractal[20]. If it is a fractal or a representation of a type of fractal, then what type of fractal is it [20]? If the fractal dimension is calculated, which of the ten plus alternatives of the general fractal dimension will it be[25, 27, 28, 29]? Referring back to the work of Richardson[21] shows that his aim was to demonstrate a relationship between the measuring scale and the attribute of the image being measured. D is then the parameter that quantifies this relationship. Thus whether the object is fractal or not may not have any bearing on the analysis of the object.

4 Conclusion

Language can be viewed as a resource for creating meaning[9, 10]. What a novice needs to do is construct a linguistic system that enables them to participate in the scientific debate. This is achieved, in our case, by helping novices to appropriate the fractal linguistic system. Specifically, knowledge needs to be encoded in a language that is understandable when transmitted from a specialist source. It should not be necessary for the learner to first decode language in order to learn anything from it. Our paper represents a start in this direction. Information must be transmitted in many different ways to incorporate different learners. Learners on the other hand must be able to freely communicate their thoughts as they attain the specific linguistic literacy.

We explored the meaning of the descriptor ‘fractal’ and one of its characteristics — ‘similarity’. Our aim was to demonstrate that an understanding of how language is used in this specialised field reflects on to the application and conclusions drawn from the research. In this respect, we discussed some of the considerations necessary when choosing an algorithm that determines the complexity of an image.

Acknowledgments

MDW would like to acknowledge the support of the National Science Foundation, Environmental and Ocean Systems Program (Grant No. BES-9312825) and of the principle investigator, Dr R.I. Dick. The work cited here was part of MDW’s Masters research.

References

- [1] D. Avnir, O. Biham, D. Lidar, and O. Malcai. Is the geometry of nature fractal? *Science*, 279:39, 1998.
- [2] J.B. Bassingthwaite, L.S. Liebovitch, and B.J. West. *Fractal Physiology*. Oxford UP, Oxford, 1994.
- [3] G.M. Bernston and P. Stoll. Correcting for finite spatial scales of self-similarity when calculating the fractal dimensions of real-world-structures. *Proc. R. Soc. Lond. B.*, 264:1531–1537, 1995.
- [4] A. Bunde and S. Havlin. *Fractals in Science*. Springer, Berlin, 1994.
- [5] B. Cambourne. Teaching literacy at the post-primary level: Is it part of the secondary teacher’s role? *Teaching & Learning On-line*, 1998. http://hsc.csu.edu.au/tlo/journal/discuss/brian_cambourne/.
- [6] F. Caserta, W.D. Eldred, E. Fernandez, R.E. Hausman, L.R. Stanford, S.V. Bulderev, S. Schwarzer, and H.E. Stanley. Determination of physiologically characterized neurons in two and three dimensions. *J. Neurosci. Meth.*, 56:133–144, 1995.

- [7] J. Feder. *Fractals*. Plenum, New York, 1988.
- [8] E. Fernandez, W.D. Eldred, J. Ammermüller, A. Block, W. von Bloh, and H. Kolb. Complexity and scaling properties of amacrine, ganglion, horizontal and bipolar cells in the turtle retina. *J. Comp. Neurol.*, 347:397–408, 1994.
- [9] J.P. Gee. What is literacy? In P. Shannon, editor, *Becoming Political: Readings and Writings in the Politics of Literacy Education*, pages 21–28. Heinemann, Portsmouth, NH, 1992.
- [10] M.A.K. Halliday. Language education: Interaction and development. the notion of context in language education. In *Proceedings of the International Conference held in Ho Chi Minh City, Vietnam*, 1991.
- [11] F. Hausdorff. Dimension und äußeres Maß. *Math. Ann.*, 79:157–179, 1919.
- [12] H.F. Jelinek and E. Fernandez. Neurons and fractals: how reliable and useful are calculations of fractal dimensions? *J. Neurosci. Meth.*, 81:9–18, 1998.
- [13] C.L. Jones. *Image analysis of fungal biostructure by fractal and wavelet techniques*. PhD thesis, Swinburne University of Technology, 1997.
- [14] N.C. Kenkel and D.J. Walker. Fractals in the biological sciences. *COENOSES*, 11(2):77–100, 1996.
- [15] H. Lauwerier. *Fractals: Images of Chaos*. Princeton University Press, New Jersey, 1991.
- [16] C. Loehle and B. Li. Statistical properties of ecological and geological fractals. *Eco. Mod.*, 85:271–284, 1996.
- [17] B.B. Mandelbrot. *The Fractal Geometry of Nature*. W.H. Freeman, N.Y., 1983.
- [18] J.D. Murray. Use and abuse of fractal theory in neuroscience. *J. Comp. Neurol.*, 361:369–371, 1995.
- [19] J. Panico and P. Sterling. Retinal neurons and vessels are not fractal but space filling. *J. Comp. Neurol.*, 361:479–490, 1995.
- [20] P. Pfeifer. Is nature fractal? *Science*, 279:784, 1998.
- [21] L.F. Richardson. The problem of contiguity: an appendix to statistics of deadly quarrels. *General Systems Yearbook*, 6:139–187, 1961.
- [22] K Sandau and H. Kurz. Measuring fractal dimension and complexity — an alternative approach with an application. *J. Microscopy*, 186(2):164–176, 1996.
- [23] M. Schroeder. *Fractals, Chaos and Power Laws*. W.H. Freeman, N.Y., 1991.
- [24] O.R. Shenker. Fractal geometry is not the geometry of nature. *Stud. Hist. Phil. Sci.*, 25(6):967–981, 1994.
- [25] H.E. Stanley and N. Ostrovsky. *On Growth and Form: Fractal and Non-Fractal Patterns in Physics*. Nijhoff, Dordrecht, 1986.
- [26] Z.R. Struzik. *From Coastline Length to Inverse Fractal Problem: the Concept of Fractal Metrology*. PhD thesis, University of Amsterdam, 1996.
- [27] H. Takayasu. *Fractals in the Physical Sciences*. Manchester University Press, Manchester and New York, 1990.

[28] M.D. Warfel. <http://www.cee.cornell.edu/~mdw/index.html>.

[29] M.D. Warfel. Characterization of particles from wastewater and sludge treatment facilities by size and morphology. Master's thesis, Cornell University, 1998.

A Categorical Representation of the State Transition Graph of Finite Cellular Automata

Jung-Hee, Park
 Department of Computer Science,
 Yangsan College,
 Kyungnam 626-800, South Korea
 PJH718@Chollian.net

Hyen-Yeal, Lee
 Department of Computer Science,
 Pusan National University
 Pusan 609-735, South Korea
 hylee@hyowon.cc.pusan.ac.kr

Abstract

One-dimensional cellular automata $A(m)$, $B(m)$, $C(m)$ and $D(m)$ with two states (0 and 1) and four different boundary conditions such as 0-0, 0-1, 1-0 and 1-1 respectively are studied on the self-reproduction of the state transition graph by the categorical representation. The simple recursive formulae of the fifteen rules such as rules 0, 3, 12, 15, 48, 51, 60, 192, 195, 204, 207, 240, 243, 252 and 255 among 256 rules were found, which are used to self-reproduce the state transition graph. It is also found that these fifteen rules satisfy that $A(m)$, $B(m)$, $C(m)$ and $D(m)$ are all equal or $A(m)$ is equal to $B(m)$ and $C(m)$ is equal to $D(m)$. Moreover, such finite cellular automata are classified according to the recursive formulae found.

1 Introduction

Cellular automata are discrete dynamical systems of simple construction but complex and varied behaviour. First introduced in 1948 by von Neumann[1] as potential models for biological self-reproduction, cellular automata have since been used as mathematical tools for studying a wide variety of problems. Attempts to find the recursive formulae which can self-reproduce the state transition diagram of the nearest-neighbor cellular automata with two states (0,1) and four different boundary conditions, which will be defined in section 2, in an algebraic method have been also made with particular tree and cycle expressions[2, 3, 4, 5]. However, many rules have still remained unexplored since it is not easy to find the recursive formulae of even one-dimensional cellular automata. In this paper it has been tried to find the recursive formulae for such automata by categorical representation with particular morphisms ϕ, ψ, τ and χ which will be defined in section 3[6, 7, 8]. The aim of this study is to self-reproduce the state transition diagrams of finite cellular automata from the cell size 0 (empty string) to given cell size m when the rule number R and the cell size m are given. The process which the state transition diagrams self-evolve when the cell size increases can be shown visually by a simple programming with the recursive formulae. It is also useful to identify the dynamical behavior of cellular automata. As a result it is found that the state transition graphs for all $A(m)$, $B(m)$, $C(m)$ and $D(m)$ of 15 rules such as rules 0, 3, 12, 15, 48, 51, 60, 192, 195, 204, 207, 240, 243, 252 and 255 among 256 rules can be self-reproduced by simple recursive formulae. In addition, the 256 rules are classified according to the patterns of the recursive formulae for generating the state transition graph. The paper is organized as follows. In section 2, we will denote a cellular automaton $(X_m, \delta_{a-b}^{m,R})$ and define the cellular automata $A(m)$, $B(m)$, $C(m)$ and $D(m)$ with four different boundary conditions such as 0-0, 0-1, 1-0 and 1-1, respectively. Section 3 defines the category $CA(X_m)$ of the evolution system of finite cellular automata and four particular morphisms ϕ, ψ, τ and χ . Section 4 proves the recursive formulae of $A(m)$, $B(m)$, $C(m)$ and $D(m)$ for rule 3 and section 5 classifies the 256 rules of one-dimensional cellular automata according to the patterns of the recursive formulae. Section 6 concludes this paper.




Figure 1: A part of the state transition diagram in the case of $m=5$ and $R=3$.

2 Finite Cellular Automata with Four Different Boundary Conditions

In this section, we will consider a cellular automaton $(X_m, \delta_{a-b}^{m,R})$ and define the finite cellular automata $A(m)$, $B(m)$, $C(m)$ and $D(m)$ with four different boundary conditions.

One-dimensional cellular automaton for any rule R , whose boundary condition is a-b (for a and b take 0 or 1) and the cell size (namely the size of the string at a site) is m , is a dynamical system $(X_m, \delta_{a-b}^{m,R})$. Here X_m is the set of states for the cellular automaton with cell size m , i.e. $X_m = \{x_1 x_2 \cdots x_{m-1} x_m \mid x_i \in \{0, 1\}\}$ and $X_0 = \epsilon$ (empty string) where m is zero or a positive integer and the state transition function $\delta_{a-b}^{m,R}$ is then defined in definition 2.1.

Definition 15.1 *The state transition function $\delta_{a-b}^{m,R}$ is defined as $\delta_{a-b}^{m,R} : X_m \rightarrow X_m$ such as*

$$\delta_{a-b}^{m,R}(x_1 x_2 \cdots x_{m-1} x_m) = f_R(a x_1 x_2) f_R(x_1 x_2 x_3) \cdots f_R(x_{m-1} x_m b) \quad (1)$$

where f_R is the triplet local transition function governed by the rule R for one-dimensional cellular automata.

For example, let us consider for rule 3. If the boundary condition is 0-0 and the cell size is 5, for the string $11001 \in X_5$

$$\delta_{0-0}^{5,3}(11001) = f_3(011) f_3(110) f_3(100) f_3(001) f_3(010) \quad (2)$$

By eq.(4) in section 4, the right hand side of eq.(2) becomes 00010. In other word, the string 11001 at a site is transmitted to the state string 00010 as shown in Figure 1.

Definition 15.2 *Define the cellular automata $A(m)$, $B(m)$, $C(m)$ and $D(m)$ with four different boundary conditions such as 0-0, 0-1, 1-0 and 1-1 respectively for each rule R as following:*

$$\begin{aligned} A(m) &= (X_m, \delta_{0-0}^{m,R}), \\ B(m) &= (X_m, \delta_{0-1}^{m,R}), \\ C(m) &= (X_m, \delta_{1-0}^{m,R}), \\ D(m) &= (X_m, \delta_{1-1}^{m,R}) \end{aligned}$$

Here $A(0)$, $B(0)$, $C(0)$ and $D(0)$ are empty graphs.

3 The Categorical Representation of Finite Cellular Automata

In this research, the recursive formulae which can self-reproduce the state transition graph of the cellular automata $(X_m, \delta_{a-b}^{m,R})$ from the previous system $(X_{m-1}, \delta_{a-b}^{m-1,R})$ for each rule R are founded. With this, the self-evolution system of finite cellular automata when the cell size m increases can be made.

Definition 15.3 The category $CA(X_m)$ of the evolution system of finite cellular automata has as objects maps $\delta_{a-b}^{m,R} : X_m \rightarrow X_m$ which we say dynamical system $(X_m, \delta_{a-b}^{m,R})$. Its morphisms are maps $h : (X_{m-1}, \delta_{a-b}^{m-1,R}) \rightarrow (X_m, \delta_{a-b}^{m,R})$ which satisfy the commutative diagram:

$$\begin{array}{ccc} X_{m-1} & \xrightarrow{\delta_{a-b}^{m-1,R}} & X_{m-1} \\ \downarrow & & \downarrow \\ X_m & \xrightarrow{\delta_{a-b}^{m,R}} & X_m \end{array}$$

Now, we will define some morphisms ϕ, ψ, τ and χ to express the recursive formulae which can self-reproduce the state transition graph of each cellular automata rule. Let X_m^0, X_m^1 be the set of states on cellular automata with cell size m , satisfied with

$$X_m = X_m^0 \cup X_m^1 \text{ and } X_m^0 \cap X_m^1 = \emptyset \quad (3)$$

where

$$X_m^0 = \{0x_1x_2 \dots x_{m-1} \mid x_1x_2 \dots x_{m-1} \in X_{m-1}\}$$

and

$$X_m^1 = \{1x_1x_2 \dots x_{m-1} \mid x_1x_2 \dots x_{m-1} \in X_{m-1}\}$$

Then the maps $\mathbf{0}$ and $\mathbf{1}$ are defined as

$$\begin{aligned} \mathbf{0} : \mathbf{X}_{m-1} &\rightarrow \mathbf{X}_m^0 : \mathbf{x}_1\mathbf{x}_2 \dots \mathbf{x}_{m-1} \rightarrow \mathbf{0}\mathbf{x}_1\mathbf{x}_2 \dots \mathbf{x}_{m-1}, \\ \mathbf{1} : \mathbf{X}_{m-1} &\rightarrow \mathbf{X}_m^1 : \mathbf{x}_1\mathbf{x}_2 \dots \mathbf{x}_{m-1} \rightarrow \mathbf{1}\mathbf{x}_1\mathbf{x}_2 \dots \mathbf{x}_{m-1} \end{aligned}$$

Definition 15.4 A morphism ϕ is defined as $\phi : (X_{m-1}, \delta_{a-b}^{m-1,R}) \rightarrow (X_m^0, \delta_{a-b}^{m,R})$ which satisfies the commutative diagram:

$$\begin{array}{ccc} X_{m-1} & \xrightarrow{\delta_{a-b}^{m-1,R}} & X_{m-1} \\ \mathbf{0} \downarrow & & \downarrow \mathbf{0} \\ X_m^0 & \xrightarrow{\delta_{a-b}^{m,R}} & X_m^0 \end{array}$$

This diagram means the following: The strings which prefix 0 to the strings in X_{m-1} are transmitted to the strings which prefix 0 to the images of X_{m-1} by $\delta_{a-b}^{m-1,R}$.

Definition 15.5 A morphism ψ is defined as $\psi : (X_{m-1}, \delta_{a-b}^{m-1,R}) \rightarrow (X_m^1, \delta_{a-b}^{m,R})$ which satisfies the commutative diagram:

$$\begin{array}{ccc} X_{m-1} & \xrightarrow{\delta_{a-b}^{m-1,R}} & X_{m-1} \\ \mathbf{1} \downarrow & & \downarrow \mathbf{0} \\ X_m^1 & \xrightarrow{\delta_{a-b}^{m,R}} & X_m^0 \end{array}$$

Definition 15.6 A morphism τ is defined as $\tau : (X_{m-1}, \delta_{a-b}^{m-1,R}) \rightarrow (X_m^0, \delta_{a-b}^{m,R})$ which satisfies the commutative diagram:

$$\begin{array}{ccc} X_{m-1} & \xrightarrow{\delta_{a-b}^{m-1,R}} & X_{m-1} \\ \mathbf{0} \downarrow & & \downarrow \mathbf{1} \\ X_m^0 & \xrightarrow{\delta_{a-b}^{m,R}} & X_m^1 \end{array}$$

Definition 15.7 A morphism χ is defined as $\chi : (X_{m-1}, \delta_{a-b}^{m-1,R}) \rightarrow (X_m^1, \delta_{a-b}^{m,R})$ which satisfies the commutative diagram:

$$\begin{array}{ccc} X_{m-1} & \xrightarrow{\delta_{a-b}^{m-1,R}} & X_{m-1} \\ 1 \downarrow & & \downarrow 1 \\ X_m^1 & \xrightarrow{\delta_{a-b}^{m,R}} & X_m^1 \end{array}$$

4 The Recursive Formulae for Generating the State Transition Graphs

The recursive formulae for self-generating the state transition graph of each cellular automata rule are shown in this section. In particular, the recursive formulae of rule 3 are proved. The triplet local transition function for rule 3 is as follows:

$$\begin{array}{cccccccc} 111 & 110 & 101 & 100 & 011 & 010 & 001 & 000 \\ \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{array} \quad (4)$$

In the following theorem, we want to show that $A(m)$ for rule 3 is partitioned by the digraphs derived by $A(m-1)$ and $C(m-1)$. Thus, for convenience, we consider τ and ψ as following:

$$\tau : (X_{m-1}, \delta_{0-0}^{m-1,R}) \rightarrow (X_m^0, \delta_{0-0}^{m,R}) \quad (5)$$

which satisfies the commutative diagram:

$$\begin{array}{ccc} X_{m-1} & \xrightarrow{\delta_{0-0}^{m-1,R}} & X_{m-1} \\ 0 \downarrow & & \downarrow 1 \\ X_m^0 & \xrightarrow{\delta_{0-0}^{m,R}} & X_m^1 \end{array}$$

and

$$\psi : (X_{m-1}, \delta_{1-0}^{m-1,R}) \rightarrow (X_m^1, \delta_{0-0}^{m,R}) \quad (6)$$

which satisfies the commutative diagram:

$$\begin{array}{ccc} X_{m-1} & \xrightarrow{\delta_{1-0}^{m-1,R}} & X_{m-1} \\ 1 \downarrow & & \downarrow 0 \\ X_m^1 & \xrightarrow{\delta_{0-0}^{m,R}} & X_m^0 \end{array}$$

By definition 2.2 and eq.(3), eq.(5) and eq.(6) can be rewritten as

$$\tau : A(m-1) \hookrightarrow A(m) \quad (7)$$

$$\psi : C(m-1) \hookrightarrow A(m) \quad (8)$$

since X_m^0 and X_m^1 are subsets of X_m .

Theorem 15.1 For rule 3, $A(m)$ is partitioned by $(A(m-1), \tau)$ and $(C(m-1), \psi)$. We denote it as $A(m) = \tau \cdot A(m-1) + \psi \cdot C(m-1)$ for convenience.

Proof 15.1 By eq.(5) and eq.(6), dynamical systems $(A(m-1), \tau)$ and $(C(m-1), \psi)$ are the same as $(X_m^0, \delta_{0-0}^{m,R})$ and $(X_m^1, \delta_{0-0}^{m,R})$, respectively. By definition 2.2, $A(m)$ is equal to $(X_m, \delta_{0-0}^{m,R})$. Thus theorem 15.1 means that $(X_m, \delta_{0-0}^{m,R})$ is partitioned by $(X_m^0, \delta_{0-0}^{m,R})$ and $(X_m^1, \delta_{0-0}^{m,R})$. By eq.(3), since $X_m = X_m^0 \cup X_m^1$ and $X_m^0 \cap X_m^1 = \emptyset$, our theorem is trivial.

Now let us consider τ and ψ as

$$\tau : (X_{m-1}, \delta_{0-1}^{m-1,R}) \rightarrow (X_m^0, \delta_{0-1}^{m,R})$$

which satisfies the commutative diagram:

$$\begin{array}{ccc} X_{m-1} & \xrightarrow{\delta_{0-1}^{m-1,R}} & X_{m-1} \\ 0 \downarrow & & \downarrow 1 \\ X_m^0 & \xrightarrow{\delta_{0-1}^{m,R}} & X_m^1 \end{array}$$

and

$$\psi : (X_{m-1}, \delta_{1-1}^{m-1,R}) \rightarrow (X_m^1, \delta_{0-1}^{m,R})$$

which satisfies the commutative diagram:

$$\begin{array}{ccc} X_{m-1} & \xrightarrow{\delta_{1-1}^{m-1,R}} & X_{m-1} \\ 1 \downarrow & & \downarrow 0 \\ X_m^1 & \xrightarrow{\delta_{0-1}^{m,R}} & X_m^0 \end{array}$$

Then these lead to theorem 15.2 by the similar way shown in theorem 15.1.

Theorem 15.2 For rule 3, $B(m)$ is partitioned by $(B(m-1), \tau)$ and $(D(m-1), \psi)$. We denote it as $B(m) = \tau \cdot B(m-1) + \psi \cdot D(m-1)$.

Again let us consider ϕ and ψ as

$$\phi : (X_{m-1}, \delta_{0-0}^{m-1,R}) \rightarrow (X_m^0, \delta_{1-0}^{m,R})$$

which satisfies the commutative diagram:

$$\begin{array}{ccc} X_{m-1} & \xrightarrow{\delta_{0-0}^{m-1,R}} & X_{m-1} \\ 0 \downarrow & & \downarrow 0 \\ X_m^0 & \xrightarrow{\delta_{1-0}^{m,R}} & X_m^0 \end{array}$$

and

$$\psi : (X_{m-1}, \delta_{1-0}^{m-1,R}) \rightarrow (X_m^1, \delta_{1-0}^{m,R})$$

which satisfies the commutative diagram:

$$\begin{array}{ccc} X_{m-1} & \xrightarrow{\delta_{1-0}^{m-1,R}} & X_{m-1} \\ 1 \downarrow & & \downarrow 0 \\ X_m^1 & \xrightarrow{\delta_{1-0}^{m,R}} & X_m^0 \end{array}$$

Then these lead to theorem 15.3 by the similar way shown in theorem 15.1.

Theorem 15.3 For rule 3, $C(m)$ is partitioned by $(A(m-1), \phi)$ and $(C(m-1), \psi)$. We denote it as $C(m) = \phi \cdot A(m-1) + \psi \cdot C(m-1)$.

Let us consider ϕ and ψ as

$$\phi : (X_{m-1}, \delta_{0-i}^{m-1,R}) \rightarrow (X_m^0, \delta_{1-i}^{m,R})$$

which satisfies the commutative diagram:

$$\begin{array}{ccc} X_{m-1} & \xrightarrow{\delta_{0-i}^{m-1,R}} & X_{m-1} \\ 0 \downarrow & & \downarrow 0 \\ X_m^0 & \xrightarrow{\delta_{1-i}^{m,R}} & X_m^0 \end{array}$$

and

$$\psi : (X_{m-1}, \delta_{1-i}^{m-1,R}) \rightarrow (X_m^1, \delta_{1-i}^{m,R})$$

which satisfies the commutative diagram:

$$\begin{array}{ccc} X_{m-1} & \xrightarrow{\delta_{1-i}^{m-1,R}} & X_{m-1} \\ 1 \downarrow & & \downarrow 0 \\ X_m^1 & \xrightarrow{\delta_{1-i}^{m,R}} & X_m^0 \end{array}$$

Then these lead to theorem 15.4 by the similar way shown in theorem 15.1.

Theorem 15.4 For rule 3, $D(m)$ is partitioned by $(B(m-1), \phi)$ and $(D(m-1), \psi)$. We denote it as $D(m) = \phi \cdot B(m-1) + \psi \cdot D(m-1)$.

The proofs of theorem 15.2, 15.3 and 15.4 are word for word the same as the proof of theorem 15.1, so are not reproduced.

Theorem 15.1, 15.2, 15.3 and 15.4 now yield:

Corollary 15.1 The state transition graphs of $A(m)$, $B(m)$, $C(m)$ and $D(m)$ for rule 3 can be self-reproduced by the interconnection of the four recursive formulae as following:

1. $A(m) = \tau \cdot A(m-1) + \psi \cdot C(m-1)$
2. $B(m) = \tau \cdot B(m-1) + \psi \cdot D(m-1)$
3. $C(m) = \phi \cdot A(m-1) + \psi \cdot C(m-1)$
4. $D(m) = \phi \cdot B(m-1) + \psi \cdot D(m-1)$

It is noted that these recursive formulae of cellular automata are irreversible. Thus cellular automata can be applied to cryptology with this property.

For example, let us consider the self-reproduction of the state transition graphs of $A(2)$ for rule 3. $A(2)$ is partitioned by the subgraphs $(A(1), \tau)$ and $(C(1), \psi)$ by theorem 15.1. In other word, $(A(1), \tau)$ is derived by transmitting the strings obtained by prefixing 0 to each string in $A(1)$ such as 00 and 01 to the strings obtained by prefixing 1 to $\delta_{0-0}^{1,3}(0) = 1$ and $\delta_{0-0}^{1,3}(1) = 0$ such as 11 and 10 respectively. $(C(1), \psi)$ is also constructed by transmitting the strings obtained by prefixing 1 to each string in $C(1)$ such as 10 and 11 to the strings obtained by prefixing 0 to $\delta_{1-0}^{1,3}(0) = 0$ and $\delta_{1-0}^{1,3}(1) = 0$ such as 00. These are shown in Figure 2. $A(2)$ in Figure 3 is composed of the subgraphs $(A(1), \tau)$ and $(C(1), \psi)$




Figure 2: The evolution process by the morphisms τ and ψ .



Figure 3: The state transition diagram of $A(2)$.

in Figure 2. These explain Theorem 4.1.

The state transition graphs of $A(m)$ and $B(m)$ for arbitrary cell size m can be reproduced by repeating this process. Similarly, $B(m)$ and $D(m)$ for rule 3 can be self-reproduced. As the same as the rule 3, the recursive formulae of all four cellular automata $A(m)$, $B(m)$, $C(m)$ and $D(m)$ for rules 0, 12, 15, 48, 51, 60, 192, 195, 204, 207, 240, 243, 252 and 255 were found in this research. These are shown in Appendix.

5 The Classification of Cellular Automata According to the Recursive Formulae

The recursive formulae for the state transition graphs of cellular automata $A(m)$, $B(m)$, $C(m)$ and $D(m)$ for each of 256 rules have investigated in this paper. 103 rules of 256 rules have at least more than two recursive formulae among four types of cellular automata $A(m)$, $B(m)$, $C(m)$ and $D(m)$.

These rules can be classified into the following three groups according to the patterns of the recursive formulae found as shown in table 1:

- Group 1. The rules that the recursive formulae of $A(m)$ and $B(m)$ only were found
- Group 2. The rules that the recursive formulae of $C(m)$ and $D(m)$ only were found
- Group 3. The rules that the recursive formulae of $A(m)$, $B(m)$, $C(m)$ and $D(m)$ all were found

The reason for the classification into the three groups is to show that the rules which have the recursive formulae belong to each of the three groups. In other words, there are no rules that the recursive formulae of $A(m)$ and $C(m)$ only or $A(m)$ and $D(m)$ only or $B(m)$ and $C(m)$ only or $B(m)$ and $D(m)$ only were found.

It is noted that even the rules such as rule 124, 147 and 195, which are included in the Wolfram's class 3 (namely have chaotic behaviors), have the recursive formulae.

Again the recursive formulae for $A(m)$ which are generated by $A(m-1)$ and $C(m-1)$ are classified into the following four types:

Group 1	Group 2	Group 3
19, 28, 31, 32, 35, 44, 64, 67, 76, 79, 80, 83, 92, 95, 96, 99, 108, 111, 112, 115, 124, 128, 131, 140, 144, 147, 156, 159, 160, 163, 175, 179, 188, 191, 208, 211, 220, 223, 188, 191, 208, 211, 220, 223, 224, 227, 236, 239	1, 2, 4, 5, 6, 7, 8, 9, 10, 11, 13, 14, 50, 52, 53, 55, 56, 57, 58, 59, 61, 62, 63, 194, 196, 197, 198, 199, 200, 201, 202, 203, 205, 206, 241, 242, 244, 245, 246, 247, 248, 249, 250, 251, 253, 254	0, 3, 12, 15, 48, 51, 60, 192, 195, 204, 207, 240, 243, 252, 255,

Table 1: The classification of the pattern of the recursive formulae found

Type 1	Type 2	Type 3	Type 4
19, 35, 67, 83, 99, 115, 131, 163, 179, 195, 211, 227, 243	28, 44, 12, 60, 76, 92, 108, 124, 140, 156, 188, 220, 236, 252	32, 48, 64, 80, 96, 112, 128, 144, 160, 192, 208, 224, 240	15, 31, 79, 95, 111, 147, 159, 175, 191, 207, 223, 239

Table 2: The classification of $A(m)$ and $B(m)$

- Type 1 : $A(m) = \tau \cdot A(m - 1) + \psi \cdot C(m - 1)$
- Type 2 : $A(m) = \phi \cdot A(m - 1) + \chi \cdot C(m - 1)$
- Type 3 : $A(m) = \phi \cdot A(m - 1) + \psi \cdot C(m - 1)$
- Type 4 : $A(m) = \tau \cdot A(m - 1) + \chi \cdot C(m - 1)$

The types of the recursive formulae for $B(m)$ are the same as those for $A(m)$, replacing $A(m - 1)$ by $B(m - 1)$ and $C(m - 1)$ by $D(m - 1)$. The rules which are included to each of these types are shown in table 2.

Similarly, $C(m)$ has the following four types of recursive formulae:

- Type 1 : $C(m) = \phi \cdot A(m - 1) + \psi \cdot C(m - 1)$
- Type 2 : $C(m) = \tau \cdot A(m - 1) + \psi \cdot C(m - 1)$
- Type 3 : $C(m) = \phi \cdot A(m - 1) + \chi \cdot C(m - 1)$
- Type 4 : $C(m) = \tau \cdot A(m - 1) + \chi \cdot C(m - 1)$

$D(m)$ has the same types as $C(m)$, replacing $A(m - 1)$ to $B(m - 1)$ and $C(m - 1)$ to $D(m - 1)$. The rules which are included to each of these types are shown in table 3.

Cellular automata can also be classified into the following three classes according to the relations among $A(m)$, $B(m)$, $C(m)$ and $D(m)$ as shown in table 4:

- Class 1. The rules satisfied with $A(m) = B(m)$ and $C(m) = D(m)$
- Class 2. The rules satisfied with $A(m) = B(m) = C(m) = D(m)$
- Class 3. The rules satisfied with $A(m) = C(m)$ and $B(m) = D(m)$

Type 1	Type 2	Type 3	Type 4
1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15	48, 50, 52, 53, 55, 56, 57, 58, 59, 60, 61, 63	192, 194, 195, 196, 197, 198, 199, 200, 201, 202, 203, 205, 206, 207	240, 241, 242, 243, 244, 245, 246, 247, 248, 249, 250, 251, 252, 253, 254

Table 3: The classification of $C(m)$ and $D(m)$

Class 1	Class 2	Class 3
3, 12, 15, 48, 60, 192, 195, 207, 240, 243, 252,	0, 51, 204, 255,	17, 34, 68, 85, 102, 136, 153, 170, 187, 221,

Table 4: The classification of the relations among $A(m)$, $B(m)$, $C(m)$ and $D(m)$

The 15 rules included in the class 1 and the class 2 are included in the group 3 of table 1. Moreover, in the class 1, $A(m)$ can be generated by the interconnection with $C(m)$ and $B(m)$ can be produced by the interconnection with $D(m)$ vice versa (namely $C(m)$ by $A(m)$, $D(m)$ by $B(m)$). However, each of $A(m)$, $B(m)$, $C(m)$ and $D(m)$ for the rules included in the class 2 has self-recursive formulae of itself (see Appendix).

6 Conclusion

In this paper, it has been tried to find the recursive formulae for the whole 256 rules of one-dimensional cellular automata. As a result, the recursive formulae of 103 rules among 256 rules were found. With these found recursive formulae, the self-evolution system of finite cellular automata can be made by a simple programming when the cell size increases. This system is also useful to identify the dynamical behavior of one-dimensional cellular automata. The morphisms ϕ , ψ , τ and χ considered the transition of only the first digit of state string on automaton configuration in this paper. The study on the transition of more than first digit remains as further work.

References

- [1] J. Von Neumann, *Theory of Self-Reproducing Automata*, A.W.Burks,ed, 1967
- [2] Yasuo Kawahara and Hyen Yeal Lee, "Period Lengths of Cellular Automata CAM-90 with Memory", *Journal of Mathematical Physics*, Vol.38, No.1, pp. 255-266 1997
- [3] Hyen Yeal Lee and Yasuo Kawahara, "On Dynamical Behaviors of Cellular Automata CA-60", *Bulletin of Informatics and Cybernetics*, Vol.25, No. 1-2, 1992
- [4] Hyen Yeal Lee, "*Studies on Dynamical Behaviors of Finite Cellular Automata*", Kyushu Univ. Ph.D Thesis, 1995
- [5] Hyen Yeal Lee and Yasuo Kawahara, "Transition Diagrams of Finite Cellular Automata", *Bulletin of Informatics and Cybernetics*, Vol.28, No. 1, pp. 47-69 1996
- [6] Heather, M. A. et. al. "Category theory:mathematics for the humanities?", *University of Newcastle upon Tyne: Computing Science. Technical Report Series 476* 1994
- [7] M. A. Arbib et.al.,*Arrows, Structures, and Functors*, Academic Press 1975
- [8] H. Herrlich et.al. *Category theory*, Herdermann Verlag Berlin 1979

A Appendix

For rule 0, the following recursive relations hold:

1. $A(m) = B(m) = C(m) = D(m)$
2. $A(m) = \phi \cdot A(m - 1) + \psi \cdot A(m - 1)$

$$3. B(m) = \phi \cdot B(m-1) + \psi \cdot B(m-1)$$

$$4. C(m) = \phi \cdot C(m-1) + \psi \cdot C(m-1)$$

$$5. D(m) = \phi \cdot D(m-1) + \psi \cdot D(m-1)$$

For rule 12, the following recursive relations hold:

$$1. A(m) = B(m), C(m) = D(m)$$

$$2. A(m) = \phi \cdot A(m-1) + \chi \cdot C(m-1)$$

$$3. B(m) = \phi \cdot B(m-1) + \chi \cdot D(m-1)$$

$$4. C(m) = \phi \cdot A(m-1) + \psi \cdot C(m-1)$$

$$5. D(m) = \phi \cdot B(m-1) + \psi \cdot D(m-1)$$

For rule 15, the following recursive relations hold:

$$1. A(m) = B(m), C(m) = D(m)$$

$$2. A(m) = \tau \cdot A(m-1) + \chi \cdot C(m-1)$$

$$3. B(m) = \tau \cdot B(m-1) + \chi \cdot D(m-1)$$

$$4. C(m) = \phi \cdot A(m-1) + \psi \cdot C(m-1)$$

$$5. D(m) = \phi \cdot B(m-1) + \psi \cdot D(m-1)$$

For rule 48, the following recursive relations hold:

$$1. A(m) = B(m), C(m) = D(m)$$

$$2. A(m) = \phi \cdot A(m-1) + \psi \cdot C(m-1)$$

$$3. B(m) = \phi \cdot B(m-1) + \psi \cdot D(m-1)$$

$$4. C(m) = \tau \cdot A(m-1) + \psi \cdot C(m-1)$$

$$5. D(m) = \tau \cdot B(m-1) + \psi \cdot D(m-1)$$

For rule 51, the following recursive relations hold:

$$1. A(m) = B(m) = C(m) = D(m)$$

$$2. A(m) = \tau \cdot A(m-1) + \psi \cdot A(m-1)$$

$$3. B(m) = \tau \cdot B(m-1) + \psi \cdot B(m-1)$$

$$4. C(m) = \tau \cdot C(m-1) + \psi \cdot C(m-1)$$

$$5. D(m) = \tau \cdot D(m-1) + \psi \cdot D(m-1)$$

For rule 60, the following recursive relations hold:

$$1. A(m) = B(m), C(m) = D(m)$$

$$2. A(m) = \phi \cdot A(m-1) + \chi \cdot C(m-1)$$

$$3. B(m) = \phi \cdot B(m-1) + \chi \cdot D(m-1)$$

$$4. C(m) = \tau \cdot A(m-1) + \psi \cdot C(m-1)$$

$$5. D(m) = \tau \cdot B(m-1) + \psi \cdot D(m-1)$$

For rule 192, the following recursive relations hold:

1. $A(m) = B(m), C(m) = D(m)$
2. $A(m) = \phi \cdot A(m-1) + \psi \cdot C(m-1)$
3. $B(m) = \phi \cdot B(m-1) + \psi \cdot D(m-1)$
4. $C(m) = \phi \cdot A(m-1) + \chi \cdot C(m-1)$
5. $D(m) = \phi \cdot B(m-1) + \chi \cdot D(m-1)$

For rule 195, the following recursive relations hold:

1. $A(m) = B(m), C(m) = D(m)$
2. $A(m) = \tau \cdot A(m-1) + \psi \cdot C(m-1)$
3. $B(m) = \tau \cdot B(m-1) + \psi \cdot D(m-1)$
4. $C(m) = \phi \cdot A(m-1) + \chi \cdot C(m-1)$
5. $D(m) = \phi \cdot B(m-1) + \chi \cdot D(m-1)$

For rule 204, the following recursive relations hold:

1. $A(m) = B(m) = C(m) = D(m)$
2. $A(m) = \phi \cdot A(m-1) + \chi \cdot A(m-1)$
3. $B(m) = \phi \cdot B(m-1) + \chi \cdot B(m-1)$
4. $C(m) = \phi \cdot C(m-1) + \chi \cdot C(m-1)$
5. $D(m) = \phi \cdot D(m-1) + \chi \cdot D(m-1)$

For rule 207, the following recursive relations hold:

1. $A(m) = B(m), C(m) = D(m)$
2. $A(m) = \tau \cdot A(m-1) + \chi \cdot C(m-1)$
3. $B(m) = \tau \cdot B(m-1) + \chi \cdot D(m-1)$
4. $C(m) = \phi \cdot A(m-1) + \chi \cdot C(m-1)$
5. $D(m) = \phi \cdot B(m-1) + \chi \cdot D(m-1)$

For rule 240, the following recursive relations hold:

1. $A(m) = B(m), C(m) = D(m)$
2. $A(m) = \phi \cdot A(m-1) + \psi \cdot C(m-1)$
3. $B(m) = \phi \cdot B(m-1) + \psi \cdot D(m-1)$
4. $C(m) = \tau \cdot A(m-1) + \chi \cdot C(m-1)$
5. $D(m) = \tau \cdot B(m-1) + \chi \cdot D(m-1)$

For rule 243, the following recursive relations hold:

1. $A(m) = B(m), C(m) = D(m)$
2. $A(m) = \tau \cdot A(m-1) + \psi \cdot C(m-1)$
3. $B(m) = \tau \cdot B(m-1) + \psi \cdot D(m-1)$
4. $C(m) = \tau \cdot A(m-1) + \chi \cdot C(m-1)$
5. $D(m) = \tau \cdot B(m-1) + \chi \cdot D(m-1)$

For rule 252, the following recursive relations hold:

1. $A(m) = B(m), C(m) = D(m)$
2. $A(m) = \phi \cdot A(m-1) + \chi \cdot C(m-1)$
3. $B(m) = \phi \cdot B(m-1) + \chi \cdot D(m-1)$
4. $C(m) = \tau \cdot A(m-1) + \chi \cdot C(m-1)$
5. $D(m) = \tau \cdot B(m-1) + \chi \cdot D(m-1)$

For rule 255, the following recursive relations hold:

1. $A(m) = B(m) = C(m) = D(m)$
2. $A(m) = \tau \cdot A(m-1) + \chi \cdot A(m-1)$
3. $B(m) = \tau \cdot B(m-1) + \chi \cdot B(m-1)$
4. $C(m) = \tau \cdot C(m-1) + \chi \cdot C(m-1)$
5. $D(m) = \tau \cdot D(m-1) + \chi \cdot D(m-1)$

Adaptive Stabilization of Lorenz Chaos

Yuping Tian

Faculty of Informatics and Communication,
Central Queensland University,
Rockhampton QLD 4702, Australia.

On leave from the Department of Automatic Control,
Southeast University, Nanjing, P.R. China

Xinghuo Yu

Faculty of Informatics and Communication,
Central Queensland University,
Rockhampton, QLD 4702, Australia

X.Yu@cqu.edu.au

Abstract

The problem of adaptive stabilization of the Lorenz system with unknown parameters is solved in this paper. The invariant manifolds theory is used in the development of the adaptive control strategy. Our control strategy is different from the conventional control strategies in the sense that it can deal with the case where the equilibrium of the system is not known. Conditions for stabilization of the Lorenz chaos are given and simulations presented.

1 Introduction

Controlling chaos has received more and more attention in the physics, mathematics and engineering communities. Various linear and nonlinear control techniques have been developed for chaotic systems, see, e.g., [1]–[5] and references therein. All the existing chaos control methods assume that the system parameters are known precisely. Although chaotic systems are deterministic dynamical systems, in practice, the constant system parameters may not be known either exactly or at all, making application of the existing control strategies problematic. In this paper, we propose an adaptive control strategy using the invariant manifold theory that is able to deal with the situation.

Adaptive control has been studied extensively and proved to be one of the effective design methods for linear and nonlinear systems with unknown parameters[6, 7, 8]. A basic assumption of these control strategies is the system equilibrium is fixed and known. The equilibrium of chaotic systems is usually a function of the system parameters. When these parameters are unknown, direct application of existing adaptive control strategies is not possible.

In this paper we investigate the adaptive chaos control issue by studying the adaptive control of the Lorenz chaos whose parameters are assumed to be unknown. We propose to use some stable manifolds (which will be adaptively adjusted) that exhibit a desired dynamical performance, and then design an adaptive control such that the manifolds are reached. An update law for parameter estimates is also proposed that is able to give exact estimation of the unknown parameters. A simulation is presented to show the effectiveness of the approach.

2 Adaptive Stabilization via Invariant Manifolds

The Lorenz system basically models the convection process in our atmosphere or the process such as heating a plot of water. The convection process was formally modeled and studied by Lorenz [10] as follows:

$$\dot{x}_1 = \sigma(x_2 - x_1), \quad (1)$$

$$\dot{x}_2 = \rho x_1 - x_1 x_3 - x_2, \quad (2)$$

$$\dot{x}_3 = -\beta x_3 + x_1 x_2. \quad (3)$$

The constant σ is called the Prandtl number and it is assumed that $\sigma > 1$. The constant ρ depends on the temperature difference and is assumed to be $\rho > 1$, meaning that the bottom of the region heated is always warmer than the top, and convection occurs. The constant β is related to the given space.

It is evident that the system has three equilibria:

$$a = (a_1, a_2, a_3)^T = \begin{bmatrix} (\beta(\rho - 1))^{1/2} \\ (\beta(\rho - 1))^{1/2} \\ \rho - 1 \end{bmatrix}, \quad (4)$$

$$b = (b_1, b_2, b_3)^T = \begin{bmatrix} -(\beta(\rho - 1))^{1/2} \\ -(\beta(\rho - 1))^{1/2} \\ \rho - 1 \end{bmatrix}, \quad (5)$$

$$c = (c_1, c_2, c_3)^T = (0, 0, 0)^T. \quad (6)$$

In this paper we assume that the constants σ , ρ and β are unknown. In this case the equilibrium a and b are also unknown. We only consider the problem of stabilization at the unknown equilibrium a . Stabilization at b is similar since the Lorenz system is symmetric with respect to the x_3 axis. Since the parameter σ does not influence the equilibrium, in the following we omit its estimation for efficiency sake.

To stabilize the Lorenz chaos we introduce two control inputs in the system as follows:

$$\dot{x}_1 = \sigma(x_2 - x_1), \quad (7)$$

$$\dot{x}_2 = \rho x_1 - x_2 - x_1 x_3 + u_1, \quad (8)$$

$$\dot{x}_3 = -\beta x_3 + x_1 x_2 + u_2. \quad (9)$$

The following two manifolds as invariant manifolds are selected:

$$M_1 = \{x \in \mathbb{R}^3 : x_2 - a_2 = 0\}, \quad (10)$$

$$M_2 = \{x \in \mathbb{R}^3 : x_3 - a_3 = 0\}. \quad (11)$$

It can be easily seen that M_1 is a stable manifold of the Lorenz system. Indeed, when $x_2 - a_2 = 0$, since $\sigma > 0$ and $\beta > 0$, from (7) and (9) it can be derived that $x_1 \rightarrow a_1$, $x_3 \rightarrow a_3$. Therefore the system is also stable on the intersection of the two manifolds:

$$M_c = M_1 \cap M_2 = \{x \in \mathbb{R}^3 : x_2 - a_2 = 0, x_3 - a_3 = 0\}. \quad (12)$$

Our goal is to design a control which drives the system trajectory along the two manifolds to the set M_c . Then, by the Lasalle's Invariant Set Principle [9] we can conclude that x will converge to the equilibrium a asymptotically where a is the only stable equilibrium of the set M_c .

To design such a stabilizing controller we construct a Lyapunov function as follows:

$$V(x) = \frac{1}{2} \psi^T(x) T \psi(x) \quad (13)$$

where $\psi^T(x) = [\psi_1(x) \quad \psi_2(x)] = [x_2 - a_2 \quad x_3 - a_3]$, $T \in \mathbb{R}^{2 \times 2}$ is a positively definite matrix. Now we want the following differential equations hold so that the controller can be designed:

$$T \dot{\psi}(x) + \psi(x) = 0. \quad (14)$$

Therefore the time derivative of $V(x)$ becomes

$$\dot{V}(x) = -\psi^T(x) \psi(x).$$

This implies that $\dot{V}(x) < 0$ for all $x \notin M_c$, meaning the system trajectory can be driven to M_c asymptotically. To find such a controller, substituting the system equations (7)–(9) into Eq.(14) gives

$$u = \begin{bmatrix} u_1 \\ u_2 \end{bmatrix} = - \begin{bmatrix} \rho x_1 - x_1 x_3 - x_2 \\ -\beta x_3 + x_1 x_2 \end{bmatrix} - T^{-1} \begin{bmatrix} x_2 - a_2 \\ x_3 - a_3 \end{bmatrix}. \quad (15)$$

If the parameters σ, ρ, β are known *a priori*, the controller (15) should be able to realise the stabilization. However, the controller (15) is not implementable because parameters σ, ρ and β are assumed to be unknown. This problem cannot be directly solved via the conventional adaptive control theory [7] because the basic assumption of the adaptive control is that the equilibrium is fixed and known. It is not the case in our situation. We shall develop a tailored adaptive control for the Lorenz chaos in the following. Because the parameters ρ and β are unknown, we use the modified two invariant manifolds M_1 and M_2 which are defined, respectively, as

$$\hat{\psi}_1(x) = x_2 - \hat{a}_2 \text{ and } \hat{\psi}_2(x) = x_3 - \hat{a}_3$$

where $\hat{a}_2 = \sqrt{\hat{\beta}(\hat{\rho} - 1)}$ and $\hat{a}_3 = \hat{\rho} - 1$. For the adaptive control design, we construct the following parameter-estimate-dependent Lyapunov function:

$$V(x, \hat{\rho}, \hat{\beta}) = \frac{1}{2} \hat{\psi}^T(x) T \hat{\psi}(x) + e^T \Gamma^{-1} e. \quad (16)$$

where $e = \theta - \hat{\theta}$ with $\theta = [\rho \ \beta]^T$ and $\hat{\theta} = [\hat{\rho} \ \hat{\beta}]^T$, and $\hat{\psi}^T(x) = [\hat{\psi}_1 \ \hat{\psi}_2]$ and $\Gamma \in \mathbb{R}^{2 \times 2}$ is a positive definite matrix. Denote

$$f(x) = \begin{bmatrix} \sigma(x_2 - x_1) \\ -x_2 - x_1 x_2 \\ x_1 x_2 \end{bmatrix}, \quad F(x) = \begin{bmatrix} 0 & 0 \\ x_1 & 0 \\ 0 & -x_3 \end{bmatrix}, \quad g(x) = \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

Then the time derivative of the Lyapunov function along the Lorenz dynamics is

$$\begin{aligned} \dot{V}(x, \hat{\rho}, \hat{\beta}) &= \hat{\psi}^T T \frac{\partial \hat{\psi}}{\partial x} (f(x) + F(x)\theta + g(x)u) + \hat{\psi}^T T \frac{\partial \hat{\psi}}{\partial \hat{\theta}} \dot{\hat{\theta}} - e^T \Gamma^{-1} \dot{e} \\ &= \hat{\psi}^T T \frac{\partial \hat{\psi}}{\partial x} (f(x) + F(x)\hat{\theta} + g(x)u) + \hat{\psi}^T T \frac{\partial \hat{\psi}}{\partial x} F(x)e + \hat{\psi}^T T \frac{\partial \hat{\psi}}{\partial \hat{\theta}} \dot{\hat{\theta}} - e^T \Gamma^{-1} \dot{e}. \end{aligned}$$

Let u consists of two component, i.e.

$$u = u_e + u_c.$$

Also let

$$\frac{\partial \hat{\psi}}{\partial x} (f(x) + F(x)\hat{\theta} + g(x)u_e) = -T^{-1} \hat{\psi}, \quad (17)$$

$$\hat{\psi}^T T \frac{\partial \hat{\psi}}{\partial x} g(x)u_c = -\hat{\psi}^T T \frac{\partial \hat{\psi}}{\partial \hat{\theta}} \dot{\hat{\theta}}, \quad (18)$$

$$\hat{\psi}^T T \frac{\partial \hat{\psi}}{\partial x} F(x)e = e^T \Gamma^{-1} \dot{e}. \quad (19)$$

Then

$$\dot{V} = -\hat{\psi}^T \hat{\psi} < 0$$

for all x such that $\psi \neq 0$. So the stability of the system state with respect to the invariant manifolds $\hat{\psi}_1 = 0$ and $\hat{\psi}_2 = 0$ is realized. From (17)–(19) we obtain

$$u_e = \begin{bmatrix} u_{e1} \\ u_{e2} \end{bmatrix} = - \begin{bmatrix} \hat{\rho} x_1 - x_1 x_3 - x_2 \\ -\hat{\beta} x_3 + x_1 x_2 \end{bmatrix} - T^{-1} \begin{bmatrix} x_2 - \hat{a}_2 \\ x_3 - \hat{a}_3 \end{bmatrix}, \quad (20)$$




Figure 1: 3D image of Controlled Lorenz Chaos

$$u_c = \begin{bmatrix} u_{c1} \\ u_{c2} \end{bmatrix} = - \begin{bmatrix} -\frac{\hat{\beta}}{2\sqrt{\hat{\beta}(\hat{\rho}-1)}} & -\frac{(\hat{\rho}-1)}{2\sqrt{\hat{\beta}(\hat{\rho}-1)}} \\ -1 & 0 \end{bmatrix} \Gamma \begin{bmatrix} x_1 & 0 \\ 0 & -x_3 \end{bmatrix} T \begin{bmatrix} x_2 - \hat{a}_2 \\ x_3 - \hat{a}_3 \end{bmatrix}, \quad (21)$$

and the update law for parameter estimation:

$$\dot{\hat{\theta}} = \begin{bmatrix} \dot{\hat{\rho}} \\ \dot{\hat{\beta}} \end{bmatrix} = \Gamma \left[\frac{\partial \hat{\psi}}{\partial x} F(x) \right]^T T \hat{\psi} = \Gamma \begin{bmatrix} x_1 & 0 \\ 0 & -x_3 \end{bmatrix} T \begin{bmatrix} x_2 - \hat{a}_2 \\ x_3 - \hat{a}_3 \end{bmatrix}. \quad (22)$$

Simulations were done to demonstrate the effectiveness of the proposed adaptive control strategy. The parameters for the Lorenz system were set to be $\sigma = 10$, $\rho = 28$ and $\beta = 8/3$. Also constant matrices were set to be

$$T = \begin{bmatrix} 1.43 & 0 \\ 0 & 0.50 \end{bmatrix}, \quad \Gamma = \begin{bmatrix} 1.05 & 0 \\ 0 & 0.25 \end{bmatrix},$$

and the initial condition was $x(0) = (-2.0 \quad -1.0 \quad 30.0)^T$, $(\hat{\rho}_0 \quad \hat{\beta}_0) = (33.0 \quad 4.0)$. Simulation results were shown in Fig.1 — Fig.3. Fig.1 presents the trajectory of the controlled Lorenz system. We can see that during the adaptation process of parameters, the trajectory spirals around and eventually converges to the desired equilibrium. Fig.2 shows the system state responses, and Fig.3 illustrates the convergence of the system parameter estimates. Apparently the parameter estimates converge to the true values of the Lorenz system.

3 Discussion and Conclusion

In this paper we have developed an adaptive control strategy for controlling Lorenz chaos. Our strategy is able to deal with the case where the equilibrium is unknown due to unknown system parameters which cannot be solved by the conventional adaptive control theory. The effectiveness of the approach has been demonstrated via computer simulations. A few remarks are offered here.




Fig. 2: System responses.

Note that the second manifold M_2 used in the proposed method is unstable. In fact, when $x_3 - a_3 = 0$, the dynamics of the Lorenz system is reduced to

$$\dot{x}_1 = \sigma(x_2 - x_1), \quad (23)$$

$$\dot{x}_2 = -x_2 + x_1, \quad (24)$$

$$\dot{x}_3 = x_1 x_2 - \beta(\rho - 1). \quad (25)$$

Differentiating (26) and (27) yields

$$(x_1 - x_2)' = -(\sigma + 1)(x_1 - x_2)$$

which means $x_1 \rightarrow x_2$. When x_1 and x_2 converge to a value such that $x_1 x_2 > \beta(\rho - 1)$, x_3 will blow up. The question is why do we need an unstable manifold for stabilization problem? To address this problem, let us consider the case when the system is controlled to the stable invariant manifold M_1 . In this case we can use a single control input as follows:

$$\dot{x}_1 = \sigma(x_2 - x_1), \quad (26)$$

$$\dot{x}_2 = \rho x_1 - x_2 - x_1 x_3 + u, \quad (27)$$

$$\dot{x}_3 = -\beta x_3 + x_1 x_2. \quad (28)$$

To make the manifold M_1 invariant, we can take a Lyapunov function as

$$V = \frac{1}{2}(x_2 - \hat{a}_2)^2 + \frac{1}{2}(\rho - \hat{\rho})^2 + \frac{1}{2}(\beta - \hat{\beta})^2. \quad (29)$$

Then, by using the design procedure presented in Section 2, we obtain the controller as

$$u_e = -(\hat{\rho} x_1 - x_1 x_3 - \hat{a}_2), \quad (30)$$

$$u_c = \frac{x_1 \hat{\beta}(x_2 - \hat{a}_2)}{2\sqrt{\hat{\beta}(\hat{\rho} - 1)}}, \quad (31)$$




Fig. 3: Parameter estimation.

and the update law for parameter estimation:

$$\begin{bmatrix} \dot{\hat{\rho}} \\ \dot{\hat{\beta}} \end{bmatrix} = \begin{bmatrix} (x_2 - \hat{a}_2)x_1 \\ 0 \end{bmatrix}. \quad (32)$$

We can see that $\hat{\beta}$ will not change. So when $t \rightarrow \infty$ the system state will converge to

$$x'_e = \begin{bmatrix} \sqrt{\hat{\beta}(\rho - 1)} \\ \sqrt{\hat{\beta}(\rho - 1)} \\ \frac{\hat{\beta}(\rho - 1)}{\beta} \end{bmatrix}.$$

Note that x'_e is neither the equilibrium of the actual Lorenz system nor the equilibrium of the Lorenz system under M_1 .

From the above discussion we can draw the following conclusions:

- When the equilibrium to be stabilized is a function of unknown parameters, the system stability is dependent on the convergence of parameter estimation;
- To ensure the convergence of parameter estimation we must introduce enough invariant manifolds to excite the system dynamics.

4 Acknowledgments

This work was supported by the National Natural Science Foundation and the National Key Project of China, and partly by the Australian Research Council.

References

- [1] G.Chen, "Chaos, bifurcation, and their control," in Wiley Encyclopedia of Electrical and Electronic Engineering, J.Webster (ed.), Wiley, New York, 1998.
- [2] G.Chen and X.Dong, From Chaos to Order: Perspectives, Methodologies, and Applications, World Scientific Pub., Singapore, 1998.
- [3] G.Chen X.Dong, "On feedback control of chaotic continuous time systems," IEEE Trans. on Circ.Sys. (I), Vol.40, pp.591-601, 1993.
- [4] E.Ott, C.Grebogi and J.A.Yorke, "Controlling Chaos," Physical Review Letters, Vol.64, pp.1196-1199, 1990.
- [5] X.Yu, "Controlling Lorenz chaos," International Journal of Systems Science, Vol.27, pp.355-359, 1996.
- [6] G.C.Goodwin and K.S.Sin, Adaptive Filtering Prediction and Control, Englewood Cliffs, NJ, Prentice-Hall, 1984.
- [7] K.S.Narendra and A.M.Annaswamy, Stable Adaptive Systems, Englewood Cliffs, NJ, Prentice-Hall, 1989.
- [8] M.Krstić, I.Kanellakopoulos and P.V.Kokotović, Nonlinear and Adaptive Control Design. Wiley, New York, 1995.
- [9] J.P.Lasalle, "Stability theory for ordinary differential equations," Journal of Differential Equations, Vol. 4, pp.57-65, 1968.
- [10] E.N.Lorenz, "Deterministic nonperiod, flow," Journal of Atmospheric Sciences, Vol.20, pp.130-141, 1963.
- [11] Y.Tian, "Controlling chaos via continuous nonlinear state feedback," Proc. of the 36th IEEE Conf. on Decision and Control, San Diego, Dec.10-12, pp.1502-1507,1997.

Part IV
Modelling

Effects of population size upon emergent group behavior

Cristobal Baray
Cognitive and Computer Science Departments
Indiana University
215 Lindley Hall
Bloomington, IN 47405
cbaray@cs.indiana.edu
<http://www.cs.indiana.edu/~cbaray/ecmas>

Abstract

Previous work defined a simple artificial world and evolved agents that utilized several effective communication schemes that aided the agents with a foraging task and predator avoidance. The agents were able to extend their average life span by coordinating their actions via undirected communication. The model did not force the agents to communicate — instead the model was designed to explore the types of communication schemes that could evolve and the situations that facilitated the evolution of communication. This work examines some of the assumptions within the previous work. Specifically, population size is altered to study the effectiveness of the communication scheme over varying conditions. This work shows that the population size can effect the group behavior and introduces a method for quantifying the emergent effects of individuals upon group behavior. The results show that the coordination techniques developed in the previous work are not always beneficial and that this cooperative model displays diminishing returns.

1 Introduction

Multi-agent systems offer many advantages over their single agent counterparts. The parallel nature of multi-agent systems facilitates scaling to handle larger, more diverse problems. The redundant property of multi-agent systems provides robustness with graceful degradation in the event of individual failures. Simple control architectures are often sufficient for the agents since multi-agent systems take advantage of emergent behaviors that arise from inter-agent actions and the interaction between the agents and their environment.

However, these advantages do not come for free. Coordinating the actions of the agents is not a trivial problem. Predicting and controlling group behavior is not a straightforward task because the relations between the system parameters and the group behavior are often complex. This research is focused on studying the parameters that effect coordination within multi-agent systems. Part of this process involves creating a method through which one can measure the effectiveness of a group's coordination. This paper describes one such method as it examines the relationship between group size and the group's performance.

2 Related Work

Distributed artificial intelligence has worked extensively on the problem of coordination ([1], [12]). However, the majority of the work in the field uses sophisticated individual agents, with belief systems and mechanisms that allow them to estimate the state of other agents and how to react to the information gathered. These agents are relatively complex and the behavior generated as they interact with each other is even more complicated. Design issues in these systems have remained hot research topics. Instead of adding to the complexity, Ferber [6] has worked with reactive agents. Reactive agents are stimulus-response machines, without any state or planning capabilities. Yet, in the work described by Ferber, the agents could only communicate by altering objects in the world. There were no facilities analogous to the auditory system. Mataric's [9] work on coordinating simplistic agents,

the agents are given a variety of primitive actions, but none of the actions involve communication. Coordination in these models is based solely upon visual input. Looking at animal societies, one notices that auditory communication can coordinate activity (elephants [13], vervet monkeys [15], and Belding's ground squirrels [16] are all examples of auditory animal communication systems). These societies can be used as inspiration for using undirected communication schemes to coordinate reactive agents.

Indeed, there has been work done on evolving communication systems with simplistic agents. MacLennan and Burghardt [7] describe their model within which they evolved a communication system. In their experiments one agent was in a position to receive environmental cues. This agent was then given the opportunity to communicate with other agents to inform them of the current state of the world. The others are rewarded for performing the proper action in response to the environmental cue provided to the single agent. With the aid of a genetic algorithm, agents evolved that were able to reliably respond appropriately to the environmental cue. Noble and Cliff [11] have reproduced and critiqued MacLennan and Burghardt's work. Their work can serve as a reminder that even in the simplest systems, small assumptions in design can have substantial side-effects in performance. Yanco and Stein [19] used two robots in a model similar to MacLennan and Burghardt's. The "leader" robot would receive an environmental cue then transmit a signal to the "follower". Their robots converged upon a communication scheme via reinforcement learning (the learning signal was supplied by humans observing the robots). Saunders and Pollack [14] describe a model that utilizes continuous channels of communication instead of discrete symbols. Their work evolved agents that would use one communication channel for recruitment in a search task. These systems all share the characteristic that communication occurs between only two agents, which could eliminate some of the advantages that communication systems have to offer.

Werner and Dyer [17] presented a model of a communication system evolved within a system of many agents interacting in a two dimensional environment. The females in their model were unable to move but were able to see and speak. The males in the model were able to move and hear but were blind and mute. After many generations of a genetic algorithm, the population of males and females converged upon representations that allowed them to efficiently find each other via the female's instructional signals. This model, along with some of the models mentioned earlier make distinctions between the sender and the receiver in the system. If an agent is designated as a sender, and others are made to receive the messages, then communication has to occur. In a sense, there is no other solution to the problems at hand besides communication.

Werner and Dyer [18] introduced a new model without these distinctions. They created 'BioLand' where the agents were modeled after Braitenberg's vehicles [4]. The model developed agents which displayed predator/prey dynamics. Though the agents were capable of communicating, the agents did not utilize a communication system. The lack of communication was attributed to the success of the agents' visual system. The agents were capable of visually sensing enough information for their survival and there was no need to communicate.

3 A New Model

Drawing from the previous simulations, as well as from animal societies, I designed a new artificial world to study undirected communication schemes and their potential coordination benefits. The world is updated in an asynchronous manner. An object is picked at random and then is activated. When the object is done performing its actions, another object is picked. Agents are allowed to roam freely about the world and since it is toroidal, they cannot fall off the edges. There is a coordinate system on the world, with each object only one spatial unit in size. Any number of objects can be at any spot in the world.

The agents have initial health values, which are reduced each time they are activated. When the agent's health value reaches 0 it is considered dead and removed from the world. Currently, only homogeneous populations of agents are introduced into the world. By creating a population of agents

from a single genetic representation, the agents can take advantage of the fact that they will only encounter agents that are identical to themselves. A biological analogy would be the similar innate behaviors found in a closely related group of animals.

In addition to the agents, the world contains areas spanning many spaces that either increase or decrease an agent's health. These areas are meant to be abstractions of food and danger. In a real world search task, these areas could represent various goal objects and novel obstacles. The areas appear periodically in random locations and would remain for varying durations. The areas are biased towards the dangerous type to create a less agent-friendly environment. Mobile predators are also part of this world and their sole purpose is to decrement the agent's health.

3.1 The Agent

The agents have 6 input channels. A single tactile channel responds to objects in contact with the agent. They have one visual channel, of a very limited range, that can sense immediately ahead, with a 40 degree field of view. They have four auditory channels, one for each direction. All the types of objects in the world have a different identifier, detectable on the visual and tactile channels. There are three distinct auditory signals recognized. The signals are discrete, in that they are heard or not heard. The clarity or strength of the signal does not degrade within an agent's sensory range.

The range for vision is five spatial units while the auditory system has a range of fifteen units. The values for the vision and auditory ranges attempt to model an agent with poor vision and decent hearing. Tall grass can limit visibility for smaller animals but their hearing is mostly unaffected and elephants use low frequency signals in order to communicate distances much further than they can see. There is no facility that the agents can use to maintain an internal sense of state. Severely limiting the abilities of the agents is an attempt to insure that the behaviors are emergent rather than an inherent ability. This work's long term goal is to study the development of intelligent systems, without necessarily developing intelligent agents.

The behavior of an agent is controlled by a simple production system. Each agent can have up to ten rules. The conditions of the production rules are combinations of possible input values logically OR'ed together. A logical NOT effecting the entire set of conditions is optional. If a production rule's condition is matched by environmental cues received by the agent, then one of seven actions is performed:

- Start moving forward (until the stop action is performed).
- Stop moving
- Orient away from stimulus (negative taxis)
- Orient towards stimulus (positive taxis)
- Emit signal 1, 2, or 3 (for one activation)

Putting it all together creates rules like these:

- if (see negative area) then start moving.
- if not (hear signal 1 behind OR see agent) then orient towards stimulus.
- if (hear signal 2 left OR hear signal 2 right) then stop moving.

If a rule's action is to orient towards or away from a stimulus, yet there is no stimulus in the condition (for instance, if not(see agent) then orient towards) then the agent will have a 10% chance of turning left or right. This interpretation of the action allows for the agents to perform a random walk.

Each time an agent is activated its health value is checked. If it is zero or less the agent is removed from the world. Otherwise, information is gathered from the local environment according to the limits of the agent's sensory system. This creates a list of inputs to use within the production system. The

inputs are stored in a limited sized buffer to simulate a finite mental capacity and an imperfect sensory system. Thus, there is no guarantee that all the inputs (or any specific input) from the environment would be available to the agent. The rule set then uses the inputs to trigger actions, with the actions taking effect immediately. Finally, the agent's health value is adjusted. One health unit is subtracted for the activation, and whenever an agent contacts a beneficial or dangerous area its health is also appropriately adjusted.

3.2 The Predator

The predators in the world are similar to the agents with only a few differences. Their sensory system has 3 visual channels (forward, left and right) with a greater range (15 units) and no auditory channels. They do not have the ability to emit sounds. If a predator is in contact with an agent, it will subtract a certain amount of health from the agent and add it to its own health. In order to make the predators a real threat, the predators' activation routine will loop probabilistically. This results in the predators being able to move several spaces per activation, essentially making them 'faster' than the agents. Also, the production rules that control the predators do not evolve, they were designed by hand and unchanged in order to remain a constant threat. Their rules instruct them to head towards any agent they see. If a predator does not detect any agents, it will move forward, occasionally turning to the left or right.

3.3 Previous Results

Early experiments [3] used a genetic algorithm to search through the space of rule sets to control the agents. Eventually a cooperative foraging behavior was evolved using a population of 15 identical agents. The foraging was accomplished by two rules — one that emitted a signal when food was found and another that oriented towards the signal when heard. Although this behavior does not directly help the sender, this communication does help extend the average life span of the group. The average life span is used to calculate the fitness of the homogeneous group's control set. The genetic algorithm then selects for these behaviors.

Later work [2] introduced the non-evolving mobile predators to the artificial world. Again, a genetic algorithm was used to evolve the rule sets further. A derivation of the foraging communication scheme was discovered. The general purpose recruitment scheme emits the attracting signal when food is found as well as when a predator is found. This behavior sounds counter-intuitive at first, but does extend the average life span of the agents more than the foraging call. This scheme takes advantage of the simple architecture of the predators. When the agents crowd the predators, the predators are not able to accurately track any particular agent. This distributes the damage done by the predator among all the nearby agents, as well as sometimes even disorienting the predators enough that they lose contact with the flock of agents entirely. Indeed, this phenomena is common in natural systems and is known as the "confusion effect" [5].

Another communication scheme evolved that used two different signals — one in response to food and one in response to predators. The agents orient towards the signal emitted for food and away from the signal generated for the predator. This scheme also had a successful derivation. Some rule sets found by the genetic algorithm would propagate the alarm call. That is, when the predator call was heard, not only would the agent orient away from the signal, but the agent would emit the alarm signal as well. This extended the effective range of the alarm call, providing the information of a threat to a wider audience.

4 Testing Performance over Population Sizes

The last two communication schemes mentioned above were examined in this work. They were chosen because it was not clear if they would continue to produce beneficial results as the number of the agents increased. Since the model uses undirected communication, population density effects the number of

recipients for each signalling. The more dense the population is, the more likely a larger number of agents will receive the signal. This behavior can lead to a noisy environment, where the abundance of signals becomes confusing and diminishes the value of the signals.

To test the effects of population size on group behavior, three agents were placed into the world with fifteen predators. The average life spans of the agents in the populations were recorded. To account for the variation in performance of the agents, this procedure was repeated nine more times and the values (average life span) were averaged over the 10 samples. This methodology was then applied with agent populations of size seven, ten, fifteen, twenty-two, thirty, forty-five, sixty, and seventy-five. The predators were kept constant at fifteen in order to keep the environment as constant as possible while the agent population size changed.

5 Results

The task of survival was chosen because it had the advantage that it can be accomplished as an individual. Foraging is a task that can be performed as an individual or as a group via communication. Allowing foraging to effect survival creates a relationship between cooperation and survival. The task doesn't demand cooperation or coordination, but individual performance can benefit from it, and it is a measurable benefit (increases in average life spans). When evaluating the performance of various population sizes, a similar metric is needed. This work proposes dividing the average life span of the group by the population size, creating the coordination advantage :

$$\text{Coordination Advantage} = \frac{\text{averageLifeSpan}}{\text{populationSize}}$$

This value determines the cases where the agents are actually coordinating their actions enough to benefit the group's behavior. This quantifies the effects of the emergent group phenomena as a measure for each agent's average contribution to the group's overall fitness. Figure 1 shows the Coordination Advantage values for the two models tested so far. The graph shows that the advantage is lost as the population size grows, reaches a minima, then starts to climb back up. However, the minima for the two schemes have different values and occur at different population sizes.

Population size plays an important role at the extremes of Figure 1. When the population is small, the agents would have to die very quickly in order for them to have a small Coordination Advantage (since the denominator is small). This is unlikely due to the imperfections of the predators — they just aren't that efficient. In smaller populations, the propagation of the alarm call proves to be a better strategy, as the signal range of the individuals is extended via the propagation.

When the population is very large, the Coordination Advantage begins to increase. This can be attributed to the confusion effect. The large number of agents makes it hard for the predators to track any particular agent. While other communication schemes explicitly attracted other agents to confuse the predators, the confusion effect can also be seen as a side effect of a large population of agents. The random distribution and movements of the agents can be enough to keep the predators from tracking any particular agent.

This leaves the propagation of the alarm call responsible for the differences in behavior for moderately sized groups. By propagating the alarm call, the agents are able to coordinate their activity better while the populations are in the low and middle ranges. In fact, all the population sizes of twenty-two or smaller are more productive per agent when the alarm calls are propagated. As the population size grows, less communication (no propagation of the call) is a more effective scheme.

Indeed, the extra communication in large populations that propagate calls degrades the performance of the group as a whole. The agents are propagating the call to too many other agents as the signal spreads too far. This creates a feedback effect, as the agents receive signals that were generated in response to their own, or even worse — that have wrapped around the world. The significance of the signal becomes questionable as an individual cannot tell how immediate the threat is. In some cases, the threat signal could drive the individual away from food sites, even when there is no nearby




Figure 1: The coordination advantages for the agent populations that don't propagate the alarm signal and that do propagate the alarm signal. The propagation of the signal is more efficient with smaller population sizes, but then loses its advantage as the population size grows.

threat. Therefore, propagating the alarm call isn't effective in large populations. But through the confusion effect, the agents are able to increase their Coordination Advantage in large populations, even though it is not explicit coordination.

In order to gain another measure of the effectiveness of the group, the amount of food foraged by the group was measured. This data doesn't rely on the abilities of the predator, but only on the ability of the agents to cooperate as a group. From one population size to the next, the amount of food gathered increases (due to more agents) so the change in harvest amount is measured instead, and that is averaged over a population size.

This data in Figure 2, shows that there are diminishing returns when it comes to foraging. The larger groups are again able to collect more as a whole, but the performance does not steadily increase. In this case there is not competition for the food, but instead, the agents are not communicating clearly. The excess communication deteriorates a large group's effectiveness, through inaccurate warning signals. As more agents are added to the system, the individuals become increasingly less productive. In smaller groups the propagating alarm call provides additional warning for the agents. This allows the individuals to live longer and have more chances to harvest. These agents end up being more productive individuals on average.

6 Discussion

The Coordination Advantage value proved useful in analyzing the differences between communication schemes. By looking at the graphs of performance over the population size, one can infer that the global trends (the U shaped curves) are due to environmental traits. The differences between the trends, the local changes, then can be attributed to the behavior differences. Via the Coordination Advantage metric this work has shown that there is an advantage to be gained at smaller population sizes with the propagation of the alarm call. However, the effort of propagating the call is wasted as the agent population increases.

Note that as each larger population dies off, the population size will enter a smaller range where communication does have benefits, which should increase their effectiveness. Yet in the end, some of




Figure 2: The amount harvested per individual changes with the population density. For instance, the average individual in the population of size 7 harvests more than the average individual in a population of size 3. The amount harvested increased by 6.14 items in the non-propagating case and 13 in the propagating case. As the population size increases, productivity of the individuals obeys the law of diminishing returns. The propagating model peaks earlier, showing that in small groups, the agents are able to forage more as they have a better warning system. In large groups, there is excess communication which degrades the signal propagating group's effectiveness.

the larger populations still have very bad Coordination Advantage values. This could be explained by the agents leaving food sources when they hear alarm calls. If alarm calls are propagated over relatively long distances, the agents would be leaving food sources for threats that are not nearby. This prevents the agents from extending their life spans — which decreases their Coordination Advantage value. The decrease in the amount of food foraged seems to support this hypothesis.

This model demonstrates that there are trade-offs between explicit coordination and population size. While the system harvests optimally at moderate population sizes, the average lifetime does not reflect that. This result reinforces the notion of coordination that arises as a side effect from the interactions between the agents and their environment. The environment does not have an intrinsic limit, creating the diminishing return. Instead the phenomena responsible for the diminishing returns arises from the interaction of the agents with each other. Their success is not due to their coordination, but instead the sheer number of agents in large populations is able to confuse the predators, making the predators less effective. This type of coordination is hard to plan for and can even be difficult to identify, but that does not mean it will not play a role in group behavior.

It would be of interest to further investigate the communicating animal societies, like Belding's ground squirrels and their alarm calls, to guide further research. In natural systems it isn't clear whether population size was limited by the effectiveness of communication or whether population size, limited by other forces, facilitates cooperative communication. This model provides a framework within which one can study the interplay between group size, communication, and group behavior, with the possibility of gaining insight into the behavior of natural systems. Regarding the design of multi-agent systems, this work also exposes some of the limitations of the systems. The systems cannot escape the law of diminishing returns, but we can study the agent-environment interactions to better predict the effects of system parameters on final returns.

7 Acknowledgements

This research is supported by the National Science Foundation under grants GER93-54898 and CDA93-03189.

References

- [1] N. M. Avouris and L. Gasser. *Distributed Artificial Intelligence : Theory and Praxis*. Kluwer Academic Press, Boston, MA, 1992.
- [2] C. Baray. Evolving cooperation via communication in homogeneous multi-agent systems. In H. Adeli, editor, *Intelligent Information Systems*, Grand Bahama Island, Bahamas, 1997.
- [3] C. Baray. On evolving communication in multiple agent systems. Technical Report 474, Indiana University, Department of Computer Science, 1997.
- [4] V. Braitenburg. *Vehicles*. MIT Press, Cambridge, MA, 1984.
- [5] L. C. Drickamer and S. H. Vessey. *Animal Behavior*. William C. Brown, Dubuque, IA, 1996.
- [6] J. Ferber. *Reactive Distributed Artificial Intelligence : Principles and Applications*. In O'Hare and Jennings [12], 1996.
- [7] B. J. MacLennan and G.M. Burghardt. synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior*, 2(2):161–188, 1994.
- [8] Maes, Mataric, Meyer, Pollack, and Wilson, editors. *From Animals to Animats 4*, Cambridge, MA, 1996. MIT Press.
- [9] M. J. Mataric. Designing emergent behaviors : from local interactions to collective intelligence. In Meyer et al. [10].
- [10] Meyer, Roitblat, and Wilson, editors. *From Animals to Animats 2*, Cambridge, MA, 1993. MIT Press.
- [11] J. Noble and D. Cliff. On simulating the evolution of communication. In Maes et al. [8].
- [12] G.M.P. O'Hare and N.R. Jennings, editors. *Foundations of Distributed Artificial Intelligence*. John Wiley and Sons, New York, NY, 1996.
- [13] K.B. Payne, W.R. Langbauer, and E.M. Thomas. Infrasonic calls of the asian elephant. *Behavioral Ecology and Sociobiology*, 18:297–301, 1986.
- [14] G. Saunders and J. Pollack. The evolution of communication schemes over continuous channels. In Maes et al. [8].
- [15] R.M. Seyfarth and D.L. Cheney. Monkey response to three different alarm calls : evidence of predator classification and semantic communication. *Science*, 210:801–803, 1980.
- [16] P.W. Sherman. Nepotism and the evolution of alarm calls. *Science*, 197:1246–1253, 1977.
- [17] G.M. Werner and M.G. Dyer. Evolution fo communication in artificial organisms. In Langton, Tayler, Farmer, and Rasmussen, editors, *Artificial Life II*, Reading, MA, 1992. Addison Wesley.
- [18] G.M. Werner and M.G. Dyer. Evolution of herding behavior in artificial animals. In Meyer et al. [10].
- [19] H. Yanco and L.A. Stein. An adaptive communication protocol for cooperating mobile robots. In Meyer et al. [10].

Generous and Greedy Strategies

Bengt Carlsson Stefan Johansson
 Department of Computer Science and Business Administration,
 University of Karlskrona/Ronneby,
 Soft Center, SE-372 25 Ronneby, Sweden
 bca/sja@ide.hk-r.se
 Magnus Boman
 Department of Computer and Systems Sciences,
 Stockholm University and the Royal Institute of Technology,
 Electrum 230, SE-164 40 Kista, Sweden
 mab@dsv.su.se

Abstract

We introduce generous, even-matched, and greedy strategies as concepts for analyzing games. A two person prisoner's dilemma game is described by the four outcomes (C,D), (C,C), (D,C), and (D,D). In a generous strategy the proportion of (C,D) is larger than that of (D,C), i.e. the probability of facing a defecting agent is larger than the probability of defecting. An even-matched strategy has the (C,D) proportion approximately equal to that of (D,C). A greedy strategy is an inverted generous strategy. The basis of the partition is that it is a zero-sum game given that the sum of the proportions of strategies (C,D) must equal that of (D,C). In a population simulation, we compare the prisoner's dilemma (PD) game with the chicken game (CG), given complete as well as partial knowledge of the rules for moves in the other strategies. In a traffic intersection example, we expected a co-operating generous strategy to be successful when the cost for collision was high in addition to the presence of uncertainty. The simulation indeed showed that a generous strategy was successful in the CG part, in which agents faced uncertainty about the outcome. If the resulting zero-sum game is changed from a PD game to a CG, or if the noise level is increased, it will favor generous strategies rather than an even-matched or greedy strategies.

1 Background

In the area of multi-agent systems (MAS), game theory [16] has proven useful, particularly as a tool for modeling the behavior of utility-based agents (see, e.g., [17]). In the quest for identifying and eventually inducing rational behavior in artificial agents, game theory has also been adopted as a normative theory for action. The main inspiration for this research has been the original axiomatic formulations of utility theory, starting with [19]. The difficulties involved in choosing a particular such axiomatisation as a blueprint for agent simulations led MAS researchers to simplify the assumptions of game theory. Confusion about the usefulness in practice of game-theoretic approaches in some MAS papers has led to criticism (cf. [15], [8], [11]). That said, simulation methods in MAS have been successfully connected to utility theory and economics, and generally to reasoning under uncertainty, and MAS simulation has matured into an important subtopic (see, e.g., [9]).

2 Methodology

In section 3, we introduce a generous-and-greedy model for strategies. There are at least four different questions that should be addressed when trying to implement this model:

1. Which kinds of strategies are involved?
2. Which kinds of games are played?

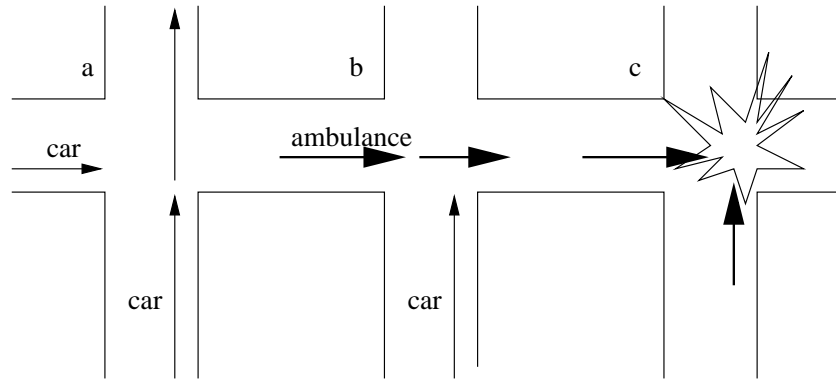


Figure 1: A traffic intersection situation.

3. What does a population of strategies look like?
4. What happens if the agents are uncertain about how to react against a strategy?

In sections 4-7 a traffic intersection example is described and simulated using both a population tournament and a noisy environment. We first look at questions 1 and 2, in section 5. Our main interest is to discuss dynamics, not to find the optimal solution for a certain kind of problem. We will look at 15 different strategies within two prisoner's dilemma-like games: the Iterated Prisoner's Dilemma (IPD) and the Iterated Chicken Game (ICG). In section 6, question 3 is treated as a population tournament. We start with the same amount of agents for each strategy and let the different agents compete within a population tournament. Finally, in section 7, we look at question 4. Introducing noise into the strategies simulates the "shaky hand principle". This means that the strategy changes to the opposite strategy for a given percentage of moves. We conclude with a short section on the implications of our results.

3 A Generous-and-Greedy Model for Strategies

The PD is a well-studied game, used in MAS [15] to create systems with a predicted cooperative behavior. When Axelrod and Hamilton ([4], [3]) analyzed the IPD, they found that a co-operating strategy, called *Tit-for-Tat* (*TfT*), did very well against strategies with more defect. This strategy has become an informal guiding principle for reciprocal altruism [18]. A *TfT* agent begins with cooperation and then imitates its opponent, in a game of unknown length. Axelrod describes this as being nice and forgiving against a defecting strategy that uses threats and punishments. Binmore [7] (p194-203) presents a critical review of *TfT*, and of Axelrod's simulation. He concludes that *TfT* is only one out of a very large number of equilibrium strategies and that it is not evolutionary stable. On the other hand, evolutionary pressures select equilibria for IPD in which the agents eventually tend to cooperate. Instead of highlighting niceness or some other similar property, we will analyze strategy quality strictly through proportions of (C,C), (C,D), (D,C), and (D,D). The notation (C,D) means that the first agent is playing cooperate against a second defecting agent, etc. We will next define informally a partition of the strategies, as an alternative to Axelrod's incomplete interpretation, in terms of nice, resistant to provocation, and evil strategies.

A generous strategy cooperates more often than its opponents do when they meet. This means that the proportion of (C,D) is larger than that of (D,C), i.e. the probability of facing a defecting agent is larger than the probability of defecting.

An even-matched strategy has the (C,D) proportion approximately equal to that of (D,C).

	Cooperate (C2)	Defect (D2)
Cooperate (C1)	$1.5T, 1.5T$	$2T, T$
Defect (D1)	$T, 2T$	$1.5T + qT, 1.5T + qT$

Table 1: Resource allocation as a time delay problem.

A **greedy strategy** defects more often than its opponents do when they meet, making it an inverted generous strategy.

The basis of the partition is that it is a zero-sum game on the meta-level in that the sum of proportions of the strategies (C,D) must equal the sum of the strategies (D,C). In other words, if there is a generous strategy, then there must also be a greedy strategy. The classification of a strategy can change depending on the surrounding strategies. Let us assume we have the following four strategies:

Always Cooperate (*AllC*) has 100% cooperate ((C,C) + (C,D)) when meeting another strategy. *AllC* will never act as a greedy strategy.

Always Defect (*AllD*) has 100% defect ((D,C) + (D,D)) when meeting another strategy. *AllD* will never act as a generous strategy.

Tit-for-Tat (*TfT*) always repeats the move of the other contestant, making it a repeating strategy. *TfT* naturally entails that (C,D) \approx (D,C).

Random plays cooperate and defect approximately half of the time each. The proportions of (C,D) and (D,C) will be determined by the surrounding strategies.

Random will be a greedy strategy in a surrounding of *AllC* and *Random*, and a generous strategy in a surrounding of *AllD* and *Random*. Both *TfT* and *Random* will behave as an even-matched strategy in the presence of only these two strategies as well as in a surrounding of all four strategies, with *AllC* and *AllD* participating in the same proportions. All strategies are even-matched when there is only a single strategy left. The described relation between strategies is independent of what kind of game is played, but the actual outcome of the game is a linear function of the payoff matrix.

4 A Traffic Intersection Example

Let us look at a traffic situation in an intersection using give right-of-way to traffic coming from the right (right-hand-rule). Drivers usually act in a cooperative mode and on average have to wait half of the time (Fig 1a). No supervisor or central control is needed to have a functional system. Rescue vehicles, like the fire brigade or an ambulance, can however use an emergency alarm to get access to the lane (Fig 1b). Let us suppose that if two ambulances both reach the intersection at the same time they will crash because they cannot hear the siren from the other vehicle (Fig 1c). If other cars begin to install sirens and behave as ambulances the whole traffic situation will collapse. The same thing happens if car drivers forget what is right and what is left. We treat such behavior as noise. Suppose it takes time T to cross the intersection. If an ambulance comes across a car, it will immediately get access to the lane. Two cars will on average need $1.5T$ to cross the intersection (we assume that there are no other time consuming delays). Two ambulances will get $1.5T + qT$, meaning that their disagreement will cause some extra costs.

Two similar games provide the foundations for this discussion of the applications of game theory in MAS: IPD and ICG. We could also have chosen, with a similar example, other PD like games like coordination game or compromise dilemma (see [13]). We will use this traffic intersection problem as an example of how to distribute (time) resources using a game theoretical model. Instead of ambulances we will talk about defecting agents that always want the resource immediately. The cars are cooperating agents that try to solve the resource allocation problem using the right-hand-rule.




Figure 2: Accumulative proportions of (C,C), (C,D), (D,C) and (D,D) for different strategies. Note how the four actions partition the space of possible actions.

5 Simulating the Traffic Intersection Example

For our simulation of the traffic intersection problem, we developed a simulation tool [12] in which 15 different strategies competed. Most of the strategies are described in ([1], [2]), see also Fig 2. All strategies handle the moves of the other agent and not the payoff value, since the latter does not affect the strategy. In a round-robin tournament, each strategy was paired with each different strategy plus its own twin, as well as with the *Random* strategy. Each game in the tournament was played on average 100 times (randomly stopped) and repeated 5000 times (see Fig 2). We interpret the proportions as a kind of fingerprint for the strategy in the given environment, independent of the actual value of the payoff matrix. For some of the strategies this is valid beyond doubt: As already noted, *AIC* and *AID* have 100% cooperate ((C,C) + (C,D)) and 100% defect ((D,C) + DD)), respectively, while *TFT* entails that (C,D) (D,C), for all payoff matrices. *AIC* definitely belongs to a group of generous strategies and so do *95% C*, *Cooperate (95% C)*, *Tit-for-two-Tats (T₂T)*, *Grofman*, *Fair*, and *Simpleton*, in this specific environment. The even-matched group of strategies includes *TFT*, *Random*, and *Anti-Tit-for-tat (ATT)*. Within the group of greedy strategies, *Feld*, *Davis*, and *Friedman* belong to a smaller family of strategies doing more cooperation moves than *Random*, i.e. having significantly more than 50% (C,C) or (C,D). An analogous family consists of *Joss*, *Tester*, and *AID*. These strategies cooperate less frequently than does *Random*. What will happen to a particular strategy depends both on the surrounding strategies and on the characteristics of the strategy. For example, *AIC* will always be generous while *95% C* will change to a greedy strategy when these two are the only strategies left. To see what these proportions mean to different payoff matrices, we recall our traffic intersection example and compare this to Axelrod's original matrix. Instead of using MaxAx, we will use a low score matrix, MinAx, shown in Fig 3 a. If q in Fig 3 b is between 0 and 0.5 it is a PD game, and if $q > 0.5$ it is a CG. The average payoff $E_{avg}(S)$ for a strategy S is a function of the payoff matrix and the distribution of the payoffs among the four outcomes (with the Greek letters referring to Fig 3 b):

$$E_{avg}(S) = p(C,C)\alpha + p(C,D)\beta + p(D,C)\gamma + p(D,D)\delta \quad (1)$$

We ran a simulation with the values for 1.5 + q equal to: 1.6; 1.9 (PDs have dashed lines); 2.1; 2.4; 3.0 (CGs have dotted lines), and compared this to the MinAx matrix (the solid line), see Fig 3. A lower

Strategy	First move	Description
<i>AllC</i>	C	Cooperates all the time
<i>95%C</i>	C	Cooperates 95% of the time
<i>Tf2T</i>	C	<i>Tit-for-two-Tat</i> , Cooperates until its opponent defects twice, and then defects until its opponent starts to cooperate again
<i>Grofman</i>	C	Cooperates if (C,C) or (D,D) was played, otherwise it cooperates with a probability of 2/7
<i>Fair</i>	C	A strategy with three possible states - “satisfied” (C), “apologising” (C), and “angry” (D). It starts in the satisfied state and cooperates until its opponent defects; then it switches to its angry state, and defects until its opponent cooperates, before returning to the satisfied state. If <i>Fair</i> accidentally defects, the apologising state is entered and it stays cooperating until its opponent forgives the mistake and starts to cooperate again [14]
<i>Simpleton</i>	C	Like <i>Grofman</i> , it cooperates whenever the previous moves were the same, but it always defects when the moves differed (e.g. (C,D)).
<i>TfT</i>	C	<i>Tit-for-Tat</i> . Repeats the moves of the opponent
<i>Feld</i>	C	Basically a <i>Tit-for-Tat</i> , but with a linearly increasing (from 0 with 0.25% per iteration up to iteration 200) probability of playing D instead of C
<i>Davis</i>	C	Cooperates on the first 10 moves, and then, if there is a defection, it defects until the end of the game
<i>Friedman</i>	C	Cooperates as long as its opponent does so. Once the opponent defects, <i>Friedman</i> defects for the rest of the game
<i>ATfT</i>	D	<i>Anti-Tit-for-Tat</i> . Plays the complementary move of the opponent
<i>Joss</i>	C	A <i>TfT</i> -variant that cooperates with a probability of 90%, when opponent cooperated and defects when opponent defected
<i>Tester</i>	D	Alters D and C until its opponent defects, then it plays a C and then <i>TfT</i> the rest of the iterations
<i>AllD</i>	D	Defects all the time

Table 2: Description of the different strategies.

a			b		
	C2	D2		C2	D2
C1	2 (a)	5 (b)	C1	1.5 (α)	2 (β)
D1	0 (c)	4 (d)	D1	1 (γ)	1.5+q (δ)

Table 3: A cost matrix for the Axelrod (a) and the resource allocation (b) matrices.




Figure 3: Outcome for the strategies in PD and CG. A lower score means a better result.




Figure 4: Population game without noise. Each bar shows the percentage of the total population for a strategy in a certain game.

score means a better result for the strategy. For the MinAx case a correlated value to the MaxAx is obtained by adding $E_{avg}(MaxAx) - E_{avg}(MinAx)$. The result is normalized to 1 for the sum of all the strategies in each game. In this example, with 15 different strategies, each strategy gets the value 0.0667 on average. None of the strategies in our simulation actually analyses its score and acts upon it. If we know the outcome of the competition between the strategies it is possible to calculate whatever payoff values are needed. This means that there is a linear correlation between the changes in scores between the games (see also [10]). Our choices of (D,D) are showing values near the borders 1.5 and 2.0 of the PD games and the border 2.0 of the CG. It is easy to extrapolate to another value, if desired. For all PD games (solid and dashed lines) there is a greedy strategy having a best score, but the result shows a large variation between different strategies. In the matrices of Axelrod and 1.9 PD, the strategies *Davis* and *Friedman* are doing best, while in 1.6 PD, *AID* is the winner. In CG, generous strategies are doing increasingly well with enhancements of the (D,D) value. This was expected, since there is an increase in the (D,D) value, and a linear payoff function was used.

6 A Population Tournament

Up until now nothing has been said about what happens if the number of agents within each strategy is allowed to vary. Maybe some vehicles after an unsuccessful trial want to change to a better strategy and ultimately find an optimal strategy for crossing the intersection. For our purposes it does not matter if we actually have ambulances and cars or if the vehicles behave like an ambulance in one intersection and as a car in another. A population tournament was held, letting each game continue until there was a single winning strategy left, or until the number of generations exceeded 10,000. For most of the games, one strategy won before reaching this limit (3150 generations were required on average). Each payoff matrix was used 100 times and the same (D,D) values were used as in the previous example. There were only four strategies not winning a single game (Fig 4). The most successful strategy was *Friedman*, which won the most games for three out of five different (D,D) values. Together with *Davis*, also a successful strategy, it belongs to the family of greedy strategies. For the PD part of the game *TfT* was successful. The generous strategies *Tf2T*, *Grofman*, *Fair*, and *Simpleton* form a rather successful family for the CG part. In Axelrod's matrix, the greedy strategies *Davis* and *Friedman*, together with *TfT*, are the winners. Notice that, because of the zero-sum nature of the game, all winners must become even-matched at the end. The initial observation of different kinds of strategies shows us how the strategies reached this even-matched state, and eventually why




Figure 5: The four most successful strategies in PD games with increasing noise.




Figure 6: The five most successful strategies in CG games with increasing noise.

they are successful.

7 Adding Noise

In the next simulation, we introduced noise on four levels: 0.01, 0.1, 1.0, and 10%. This means that the strategies changed to the opposite moves for this given percentage. The presence of uncertainty makes a huge difference as to which applications our results might have, and several writers (cf., e.g., [15]) have argued for the fact that noise as used here is an adequate representation of uncertainty. In Axelrod's simulation, *Tft* still won the tournament when 1% chance of misperception was added [3]. In other simulations of noisy environments, *Tft* has instead performed poorly [5]. The uncertainty represented by the noise reduces the payoff of *Tft* when it plays itself in the IPD. Instead of looking at all the different games we formed two different groups: PD, consisting of the Axelrod, 1.6D and 1.9D matrices, and CG consisting of 2.1D, 2.4D and 3.0D matrices. For each group we examined the five most successful strategies for different levels of noise. Fig 5 and 6 show these strategies for PD and CG when 0, 0.01, 0.1, 1.0, and 10% noise is introduced.

Among the four most successful strategies in PD there were three greedy and one even-matched strategy. In all, these strategies constituted between 85% (1% noise) and 60% (0.1%) of the population. *Tft* was doing well with 0.01% and 0.1% noise, *Davis* was most successful with 1% noise, and *AllD* with 10% noise. Three out of five of the most successful strategies in CG were generous. The total line in Fig 6 shows that five strategies constitute between 50% (no noise) and nearly 100% (0.1% and 1% noise) of the population. *Tft*, the only even-matched strategy, was the first strategy to decline as shown in the diagram. At a noise level of 0.1% or more, *Tft* never won a single population competition. *Grofman* increased its population until 0.1% noise, but then rapidly disappeared as noise increased. The same pattern was shown by *Simpleton* that declined after 1% noise level. Only *Fair* continued to increase when more noise was added, making it a dominating strategy at 10% noise together with the greedy strategy *AllD*.

8 Conclusions

Having illustrated the concepts of generous, even-matched, and greedy strategies we now return to the four questions posed in section 2. Which kinds of strategies are involved? Each strategy involved can be described using a "fingerprint" for each agent with a certain amount of (C,D) and (D,C) forming generous, even-matched, or greedy strategies. A new environment involves a new fingerprint for each

agent. Which kinds of games are played? The outcome of the game will depend on the payoff matrix involved. With the given interpretation of generous and greedy strategies it is natural to look at PD like games because they consist of co-operating and defecting behaviors. Different PD and CGs are the result of changes in the (D,D) value. For a certain set of strategies there is a linear correlation between the score of (D,D) and the score of each strategy. What does a population of strategies look like? A successful strategy has to do well against itself so, if the cost of the (D,D) value is high, we should expect generous or even-matched strategies to be successful. In CGs, cooperation proved to be increasingly fruitful, following an increase in the (D,D) value from 2.1 over 2.4 to 3.0. For strategies competing in a round-robin tournament, greedy and even-matched strategies did well in PD games, with *Friedman*, *Davis*, and *TfT* outscoring the other strategies in our traffic intersection example. What happens if the agents are uncertain about how to react against a strategy? We looked at an uncertain environment, free from the assumption of any existing perfect information between strategies, by introducing noise. Generous strategies were dominating the CG while greedy strategies were more successful in PD. In PD, *TfT* was successful with a low noise environment and *Davis* and *AllD* with a high noise environment. *Fair* was increasingly successful in CG when more noise was added. We conclude that the generous strategies are more stable in an uncertain environment in CG. Especially *Fair* and *Simpleton* were doing well, indicating these strategies are likely to be suitable for a particularly unreliable and dynamic environment. The same conclusion about generous strategies in PD, for another set of strategies, has been drawn by Bendor ([6],[5]). In our PD simulations we found *TfT* being a successful strategy when a small amount of noise was added while greedy strategies did increasingly better when the noise increased. This indicates that generous strategies are more stable in the CG part of the matrix both with and without noise. Given these results, and our chosen example, we recommend resource allocation agents to adapt a co-operating, generous strategy when the cost for a collision is high, or when different agents cannot be certain of the outcome of the game.

References

- [1] R. Axelrod. Effective choice in the prisoner's dilemma. *Journal of Conflict Resolution*, 24(1):379–403, 1980.
- [2] R. Axelrod. More effective choice in the prisoner's dilemma. *Journal of Conflict Resolution*, 24(3):3–25, 1980.
- [3] R. Axelrod. *The Evolution of Cooperation*. Basic Books Inc., 1984.
- [4] R. Axelrod and Hamilton W.D. The evolution of cooperation. *Science*, 211, 1981.
- [5] J. Bendor. Uncertainty and the evolution of cooperation. *Journal of Conflict resolution*, 37(4):709–734, 1993.
- [6] J. Bendor, R.M. Kramer, and Stout S. When in doubt...cooperation in a noisy prisoner's dilemma. *Journal of conflict resolution*, 35(4):691–719, 1991.
- [7] K. Binmore. *Playing Fair: game theory and the social contract*. MIT Press, 1994.
- [8] M. Boman and L. Ekenberg. Decision making agents with relatively unbounded rationality. In Nawarecki, editor, *Proceedings of DIMAS'95*, pages 28–35, 1995.
- [9] M. Bonatti, Y. M. Ermoliev, and A. A. Gaivoronski. Modeling of multi-agent systems in the presence of uncertainty: The case of information economy. Technical report, IIASA Working Paper WP-96-94, 1996.
- [10] B. Carlsson and S. Johansson. An iterated hawk-and-dove game. In Pagnucco, M. Wobcke, W. and Zhang, C., editors, *Agents and Multi-Agent Systems*, volume 1441 of *Lecture Notes in Artificial Intelligence*, pages 25–37. Springer Verlag, 1998.

- [11] Conte R. Castelfranchi, C. Limits of economic and strategic rationality for agents and multi-agent systems. *RAS Journal*, 1998. To appear.
- [12] Johansson, S. and Carlsson, B. Mini manual for the use of sits, simulation tool for strategies. Manuscript, available on request.
- [13] Carlsson, B. Johansson, S. and Boman, M. Modelling strategies as generous and greedy in prisoners dilemma like games. In *Proceedings of SEAL 98*, 1998.
- [14] K. Lindgren. Evolutionary phenomena in simple dynamics. In Farmer, J. D. Langton, C. G., Taylor, C. and Rasmussen, S., editors, *Artificial life II*. Addison Wesley, 1991.
- [15] B. Lomborg. Game theory vs. multiple agents: The iterated prisoner's dilemma. In Castelfranchi, C. and Werner, E., editors, *Artificial Social Systems*, volume 830 of *Lecture Notes in Artificial Intelligence*. Springer Verlag, 1994.
- [16] E. Rasmusen. *Games and Information*. Blackwell, second edition, 1994.
- [17] Rosenschein, J. S. and Zlotkin, G. *Rules of Encounter*. MIT Press, 1994.
- [18] R.L. Trivers. The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46:35–57, 1971.
- [19] Von Neumann, J. and Morgenstern, O. *Theory of Games and Economic Behavior*. John Wiley and Sons, New York, 1944.

A Complex Systems Approach to Simulating Human Behaviour Using Synthetic Landscapes

H. Randy Gimblett
 Associate Professor
 School of Renewable Natural Resources
 University of Arizona
 Tucson, Arizona, USA
gimblett@nexus.srn.arizona.edu
<http://nexus.srn.arizona.edu/~gimblett>

Merton T. Richards
 Associate Professor
 School of Forestry
 College of Ecosystem Science
 and Management
 Northern Arizona University
 P.O. 15018
 Flagstaff, Arizona, USA 86011-5018
richards@alpine.for.nau.edu

Robert M. Itami
 Environmental Planner
 Catchment Management and Sustainable Agriculture
 Department of Natural Resources and Environment
 P.O. Box 500
 East Melbourne 3002
Robert.Itami@nre.vic.gov.au

Abstract

This work introduces and explores the potential of using intelligent agent based modeling and simulation as a tool for examining the complex interactions between recreators and the environment, and interactions between recreators as a means to improving our understanding of the recreational use of wildland settings. In this research the concept of rule-driven autonomous agents as surrogates for human visitors is introduced. Agents are designed to represent the actions of the individual recreators (hiking, mountain bike riding, and pink jeep tour outfitters). Behavioural rules are derived from visitor surveys and interviews conducted in Broken Arrow Canyon, Arizona. The autonomous agents can be seen to dynamically move over a GIS based model of the Broken Arrow landscape. Line-of-sight calculations determine whether an individual agent is able to 'see' other agents and are used as method to record 'actual' and 'perceived' encounters with other agents. Using agent location maps combined with the underlying GIS data the agents can be observed moving across the landscape, pausing, changing pace, lingering at a view-point etc. A discussion focuses on analysing the resulting behaviours found in these simulations and additionally to explore the influence of alternative trail alignments on recreator movement, congestion and crowding. Some potential future directions for this research are discussed.

1 Introduction

Recreational uses of forest lands are among an extensive array of commodities and amenities that are increasingly demanded of forest managers. An in depth understanding of the relationships between recreational and other important uses is essential to effective ecosystem management[42]. Within the human dimension of ecosystem management, recreation and amenity uses of forest lands and the associated benefits of those uses, constitute an important component of management decisions. With the recent interest in the human dimension of ecosystem management, new opportunities are provided to improve upon recreation theory by developing new methods to collect, assess, model and simulate spatially relevant data of recreational use, needs, desires and behaviors in forest settings over time and incorporate these assessment methods into an ecosystem modeling framework.

The recreation assessment of forest lands in ecosystem management requires the interaction of four models: a model of desired recreation settings; a model that expresses the outcomes of recreation

behavior in those settings; a model of recreation behavior that predicts the number of users per unit of time in those settings and a model that minimizes conflicts within and between recreation groups from which personal, social, and economic value estimates can be made. The unique character of the recreational use of forest lands both requires this approach and makes it achievable[43].

While technology and its applications are growing by leaps and bounds, there still seems to be a considerable lack of knowledge and confusion in the area of recreation behavior in forest settings. The spatial orientation and temporal nature of encounters, conflicts, psychological states, experience opportunities and associated benefits between and within groups of recreators is still not well understood. Various authors[17, 33, 27, 50, 34, 45, 40, 32, 11, 16] have all focused on the the nature and extent of conflict between members of specific recreation groups, none to date have examined conflict from a spatial and temporal perspective. Work by [30, 29, 38, 47, 26, 52, 39, 49] have employed a variety of techniques including on-site surveys, recording devices and an experience sampling method[8, 28, 9, 10, 37] to quantify immediate psychological states and desires to get at this issue of dynamic, multimodal experience. Results of studies have varied but shed new light on the nature of the recreation experience. Recent work by [30, 29, 26] are among the few to have successfully used these methodologies to analyze a recreators dynamic experience patterns on-site and found that they varied predictably over the course of an outing and were strongly influenced by site characteristics and site management.

Recreation use and the values humans place on the amenity of public lands constitute an important component of management decisions and yet very little is known about the dynamic nature of the recreation experience, where recreators go in the environments, where conflicts occur and when crowding has a negative effect on the recreation experience and outcomes. The work presented in this paper is guided by the following ideas:

- That decision makers, such as natural resource managers would benefit from inexpensive, simulation techniques that could be utilized to explore dynamic recreation behavior, develop thresholds of use and test ideas or theories, before expensive management plans are implemented;
- If resource managers are to have confidence in the use and results of these simulations it is crucial that the design of the behavioral systems which are utilized in the simulations are grounded in observations and data captured of actual human behavior in the physical settings in which they naturally occur; and
- Integrated decision support models that explicitly relate the manageable characteristics of forest lands to recreational uses would greatly improve incorporation of the human dimension into ecosystem management.

2 Redrock Country — Expanding Pressures From Recreational Use: A Case Study

The focus of this paper is to present research undertaken to develop a new form of intelligent decision support and simulation system (IDSS) to assist natural resource managers in assessing and managing dynamic recreation behaviour, social interactions and resulting conflicts in wilderness settings using artificial intelligent agents in the Sedona Ranger District of the Coconino National Forest, Sedona, Arizona. The focus of this research is to utilize simulation techniques for exploring the complex interactions between recreators and the environment, and interactions among recreators as a means to improving the foundation of recreation theory.

Sedona, Arizona has been used in this work because it is typical of many special places which have become ever increasingly popular destinations for local, national and international tourists. The Sedona/Oak creek ecosystem covers nearly 200,000 acres from Sycamore Canyon on the West, to I-17 on the east; from beyond the Mogollan Rim on the north, to Beaverhead Flats and the savanna on the south. Sedona is well know by New Age enthusiasts for its “Spiritual Vortex”. This, combined

with the close proximity to the Grand Canyon, Monument Valley and the Navaho and Hopi Indian Reservations, make it an important tourist destination by visitors around the world.

Broken Arrow Canyon near Sedona, Arizona was used to capture visitor use data and demonstrate the prototype software to simulate conflicts between recreation groups over time. The Canyon is popular for day hikers, mountain bikers and people on commercial jeep tours because of the unique spectacular desert scenery of eroded red sandstone. The popularity of this canyon is a problem common to many popular wildland recreation destinations. People are “loving the place to death” by overuse. This overuse not only has negative impacts on the landscape but also in the quality of the experience people have when they visit. Crowding, conflicts between hikers, mountain bike enthusiasts and jeep tours can create negative experiences in what should be a spectacular and memorable landscape setting, but very little is known about where, why and how these impacts are occurring.

U.S. Forest Service who manages the resource, have been seeking guidance on what actions to take to protect the environment and provide the best possible recreation experience for a increasing diversity of visitors. While conventional survey techniques and public meetings have assisted in acquiring a better understanding of use, the spatial and temporal nature of the recreation experience still remain grossly misunderstood. To date, frameworks such as Recreation Opportunity Spectrum (ROS), Limits of Acceptable Change (LAC) and the growing interest in Benefits-Based Management (BBM) have provided managers with guidelines to assess the recreation opportunities, associated beneficial outcomes and identify where human-induced changes occur and to what degree they are acceptable. However, there no tools currently available for natural resource managers to study and quantify the complex spatial dynamic interactions and resulting impacts of recreational use over time.

The Recreation Behavior Simulator (RBSim), outlined later in this paper was developed to address these complex issues by using computer simulation technology. By simulating human behavior in the context of geographic space, it is possible to study the number and type of interactions that visitors will have within each group and between groups. Agent based modeling techniques are used to instill human-like behavior into artificial agents to explore recreation planning alternatives. If resource managers are to have confidence in the use and results of agent-based simulations it is crucial that the design of the behavioral systems of these agents is grounded in observations of actual human behavior in the physical settings in which they naturally occur. From the recreation behavior, management and conflict assessment perspective, a system of this nature could provide a better understanding of recreation conflict and provide a mechanism to test and assess new assumptions and theories of recreation behaviour (goal interference theory) and beneficial outcomes of experiences. In addition, the results of the simulations yield spatially-explicit, social setting data, could ultimately be used to strengthen and improve the overall predictability and mapability of ROS.

3 Rationale for Development of the Agent-Based Simulation Environment

There is a growing interest in the research community for using GIS for modeling spatially-explicit dynamic processes [2, 25, 46, 31, 22, 24, 4, 48]. The use of Individual-Based Models (IBM) is one of the popular approaches to modeling spatially-explicit ecological phenomena. IBMs according to [46] are “organisms-based models capable of modeling variation among individual and interactions between individuals.”

IBMs offer potential for studying complex behavior and human/landscape interactions within a spatial framework. Since spatial information about a phenomena is stored on a georeferenced coordinate system, space within a grid is implicit and relative to the origin of the grid [46]. IBMs offer some basic advantages over current cellular automata and other dynamic spatial modeling approaches for examining spatially explicit phenomena. Since space is continuous and location is explicit in IBMs, individuals can be simulated, independent of the environment. This provides the modeler with the ability to define an individual’s behavior, personality traits and interaction rules when encountering other individuals. Computer modeling of most ecological phenomena evolves in simulated time. Since

space is continuous and individuals are represented independently, temporal and spatial variability in IBMs can be handled asynchronously (individually updated) versus synchronously (global update) common to most raster-based GIS systems.

One form of individual-based modeling approaches that has recently gained popularity is artificial intelligent agents. Intelligent agents or what is referred to as ‘agent-oriented programming’ is being used to capture behavioral conditions and sets of intercommunication among and within agents that coexist in a environment. Several researchers [35, 41, 36, 51, 12, 5, 4, 24, 3, 1, 44, 19, 18] have taken advantage of the spatially-explicit IBMs, agents and GIS. Each of these researchers uses various techniques to build linkages between the simulation models and GIS.

Only recently have researchers seeking new ways to understand human/environment interactions been exploring simulation as a tool for developing models of human behavior. Recent studies by [15, 20, 7] clearly demonstrate the potential for agent-based modeling techniques to examine human/landscape interactions. These studies utilize a general model of multi-agent simulations based on computation agents that represent individual organisms (or groups of organisms) in a one to one correspondence. These studies seek to understand the process of evolution in the study of ecological and sociological systems. As Drogoul *et al.*[15] state “we are interested in the simulation of evolution of complex systems where interactions between several individuals at the micro level are responsible for measurable general situations observed at the macro level. When the situation is too complex to be studied analytically, it is important to be able to recreate an artificial universe in which experiments can be done in a reduced and simulated laboratory where all parameters can be controlled precisely.” This work and others examining emergent processes in societies is extremely exciting and is yielding interesting results that would have been hard to obtain without the use of such simulations.

Few studies to date that have explored emergent behavior in individuals or societies have utilized the power of GIS for representing the spatial worlds they reside in and interactions with those worlds. What is surprising is that none of these studies have taken advantage of the power of GIS. Since human behavior is inherently spatial, GIS can provide the worlds that individuals could respond to and function within. Currently, there is no GIS system with IBM capabilities. Dibble[13] states that “individual-based models do currently exist (Santa FE Institute, Swarm¹) and in many ways these systems may offer far deeper insights into human geographic phenomena than any current GIS.”

It is clear from an evaluation of current research using spatially-explicit IBMs, reactive agents and GIS to model animal/ landscape interactions that they offer a powerful alternative to previous modeling techniques for exploring emergent, complex, evolutionary processes. The ability to model the differences among groups, local interactions and variability in time and space, as well as the complex, decision making process of an individual, make IBMs an ideal technique for exploring human / landscape interactions.

4 Capturing and Defining Personality Traits and Interaction Rules for Simulated Recreators

To represent and simulate an individual’s behavior independent of the environment it requires an understanding of their personality traits (which include personal goals/intentions, expectations, length of visit, age etc.) and rules which define how they move and interact with other individuals they encounter and to the physical world or landscape where they are engaging in their favorite activity. This study consists of three phases; first to capture and analyze recreational use data for providing artificial agents with personalities and rules that closely reflect actual recreator behavior; second the development of an GIS-based agent simulator for mimicking recreator behavior overtime; finally testing of the agents in their simulated world under both typical use conditions and imploring alternative management strategies.

¹<http://www.santafe.edu/projects/swarm>

4.1 Defining Personality Profiles of Recreators from Visitor Use Data

An on-site visitor use survey was employed before and after recreational outing over a nine month period to capture data on recreational use, desired beneficial outcomes and conflicting recreational uses in the canyon. Trip motives, expectations, use density, reported contacts and place of encounters have been identified as key factors in a satisfactory recreational experience (Scenic Spectrums Pty Ltd. 1995). A binary measurement was used to solicit response on the type of benefits that were desired (trip motives and expectations) during their visit and to what degree they were able to obtain them. The focus was on recreation as essentially goal-directed behavior [14]. Expectations have been acknowledged as extremely important to goal-oriented approaches to recreation behavior. This measure coincided with Jacob and Schreyer's [34] goal interference definition of conflict. Visitors were asked if a range of benefits were desirable (*goals and intentions*) and whether they could obtain those benefits over time (*goal interference*). The benefit types used in this study are well documented in Bruns et al. [6] and Lee and Driver [38], based on research undertaken on other public lands. Desired benefits such as getting away from crowds, reduced stress and physical fitness are strong indicators of recreational satisfaction. Crowding has been shown to be one of the major predictors of user dissatisfaction. The survey was used to identify anything that either made the setting an ideal place for achieving, or interfered with acquiring the desired benefits. So negative detractors and the inability to obtain desired benefits together are used to measure goal interference and conflicts, and imply an inability to obtain desired recreational experiences leading to unsatisfactory outcomes [21, 23].

To derive meaningful recreator profiles of the visitors to Broken Arrow Canyon, cluster analysis first run on the recreation activity respondent data to isolate visitors by activity groups and then later used to aggregate visitors within each group based on desired benefits (goals and intentions). K-Means Cluster analysis allows one to specify the number of clusters desired or in the case of this research to explore the number of significant recreator types that could be found within each activity group. In addition, cross tabulation was used to calculate the frequency of which respondents within the classes derived previously identified the significance of each benefit type. Similarly they were asked to indicate their ability to obtain each of the benefits. This measure provides an indication of how often the respondents loaded on the benefit types by cluster and what particular benefits could not be obtained. This analysis was subsequently used to determine statistically relevant number of agent types within each activity class for subsequently programming artificial intelligent agents with these identical behavior traits. Since it is possible to derive hundreds of agent personality profiles, for purposes of demonstrating the method, this research aggregated agent classes into a reasonable number for final implementation.

Of the ($n=1041$) visitors sampled, three significant recreation use groups were identified; day-use hikers ($n=337$), mountain bikers ($n=393$) and Commercial Jeep outfitters (pink jeep tours) ($n=319$). While there was an extensive amount of visitor use data collected during the field work, only some is pertinent to this particular paper. For more detailed demographic data see Gimblett[21]. While there could be many combinations of personality traits derived from the visitor data collected, to demonstrate the utility of the agent modeling system the recreator patterns resulting from the cluster analysis were aggregated into two unique types for both the hikers and mountain bikers. These two types are referred to in this work as either a '*landscape*' or '*social*' recreator type. Each has significantly different desired benefits of their recreation experience.

Figure 1 illustrates the differences in the two recreator types. A *landscape recreator or agent type* is one that seeks out landscapes that are *physically challenging, avoid crowds subsequently leading to a reduction in stress*. This recreator type typically avoids others at all costs. This is evident by the extremely high desire to avoid crowds. In the exit interviews, recreators that were representative of this agent class indicated that they would only stop in locations where there are no other recreators and move as fast as possible along the trails. Physical exercise was a strong motivation in this recreation group and common to both hikers and mountain bikers as can be seen in Figure 2. These recreator types fall within the personal well being and health benefits class as identified in Bruns et al. [6].

A *social recreator or agent type* is more group oriented, one who seeks out those landscapes which




Figure 1: The Frequency of Responses of Aggregated Hikers to Desired Benefits




Figure 2: The Frequency of Responses of Aggregated Bikers to Desired Benefits

	Hiker Social Agents (Cluster 1,2,3)		Hiker Landscape Agent (Clusters 4,5)	
	of Agents	Ratio of Hikers	of Agents	Ratio of Hikers
Age Group 1–20	10	3%	13	3%
Age Group 21–40	101	26%	77	20%
Age Group 41–60	73	18%	64	17%
Age Group > 60	24	6%	29	7%
Proportion of Agents		53%		47%

Table 1: Aggregation of Hiker Agent types for Simulations based on Benefit Preferences and Age Group

are not necessarily physically challenging but tend to build *self-confidence*, provide more *opportunity to learn more about the natural and cultural history* of the area and interact with others who share these goals. This is evident in Figures 1 & 2 where their desires to obtain certain types of benefits are not as strong as those representative of landscape recreators. Social agents did not mind encountering other social recreators along the trail in the case of either hikers or mountain biking. During the exit interviews, recreators that represented this class indicated that liked social interaction while engaging in their favorite recreational activity and will spend longer periods of time wandering through the landscape, sitting in special locations, and contemplating life. This was again true of the mountain bikers as seen in Figure 2. Both of these recreator types have been identified over and over again in the applied recreation literature. Tables 1 & 2 outline the aggregation of agent clusters and their associated age groups into agent classes for both Hikers and Mountain Bikers. For more details on

	Biker Social Agents (Cluster 1,2,3)		Biker Landscape Agent (Clusters 4,5)	
	of Agents	Ratio of Bikers	of Agents	Ratio of Bikers
Age Group 1–20	41	4%	60	6%
Age Group 21–40	393	39%	426	42%
Age Group 41–60	38	4%	46	4%
Age Group > 60	4	.5%	1	.5%
Proportion of Agents		47.5%		52.5%

Table 2: Aggregation of Biker Agent types for Simulations based on Benefit Preferences and Age Group




Figure 3: The Frequency of Responses of Aggregated Hikers Ability to Obtain Benefits




Figure 4: The Frequency of Responses of Aggregated Bikers Ability to Obtain Benefits

the statistical analysis see Gimblett[21].

Figures 3 & 4 represent the same recreator agent classes as presented in Figures 1 & 2 but in response to the questions asking their ability to obtain the type of benefits that they desired. Figure 3 illustrates there was a reasonable agreement by both groups of recreators that they could satisfy their desired goals, except in the case of the landscape hikers who could not avoid crowds as opposed to the social agents who could or it simply did not matter to them. Figure 4 reports on the mountain biker's ability to obtain their desired benefits differed. The landscape bikers reported an inability to obtain their desire physical challenge benefits and strongly agreed that it was too crowded. Crowds could be the reason they were unable to achieve the physical benefit. The social bikers seemed to agree that most of their desired benefits could be obtained.

While these results are certainly not conclusive, they do provide a method for assessing the goals and intentions of the recreators visiting Broken Arrow Canyon and also provide a measure of how well they were able to meet those goals or satisfied with there recreational experience. While none on the visitors indicated they were totally unsatisfied with their experience, many seemed frustrated with the numbers of encounters they had with other recreators using the canyon.

5 Rules for Simulating Individual Behaviors and their Interactions with each other the Physical World

Rules for providing the simulated agents with social behaviors of human recreators were derived from what respondents told us about their experiences in the surveys, statistical analyses presented earlier and through interviews following their outing. The respondents were all asked to explain the types of behaviors that they exhibited along the trails when encountering other recreators. While the surveys clearly documented that visitors spent a minimum of two hours performing their activities, the desired benefit questions provided the goals and intentions for their visit. The maps provide a clear indication of where they rested, their final destination or where they stopped to view cultural and geologic features. Many of those recreators that subsequently fell into the social agent class indicated that they stopped at all the locations regardless of the numbers of other hikers or bikers that were present and stayed primarily on the appropriate trail for their activity. Some of mountain bikers and hikers who fell into the landscape agent classes indicated in both the survey and later in the exit interviews that they would stop at the cultural and geologic features only if there were a limited number of other recreators present. They also indicated that they would go out of their way to pass others along the trails quickly and avoid them if possible.

For the testing of the prototype agent simulator, Table 3 presents six rules were developed that conformed to what was learned about the intensity of use and interactions of both social and landscape

Behaviour Rules for Recreator Agents	
Rule 1	Hiker and mountain biker agents rest when down to 25% energy level
Rule 2	Landscape agents, pass other agents in front travelling 25% slower than you if you have at least 50% energy left;
Rule 3	Hiker and mountain biker landscape agents, slow down at landscape features if no other recreators;
Rule 4	Hiker, mountain biker and jeep agents stop at all landscape features;
Rule 5	Social hiker agents change their velocity to match other social hiker agents they encounter
Rule 6	Hiker and mountain biker agents will not stop at features if more than five other agents are present.

Table 3: Mobility Rules for Agents

	Landscape Agents	Social Agents
	Behaviour Rules	
Hiker Agents	Rules 1, 2, 3, 6	Rules 1, 4, 5
Mountain Biker Agents	Rules 1, 2, 3, 6	Rules 1, 4, 5
Jeep Agents ^a	Rule 4	

^aThe jeeps are assumed to move continuously throughout the landscape and stop at all features for interpretation

Table 4: Rules that Modify Agent Behaviour

types of recreators using Broken Arrow Canyon. To accommodate the solitude seeking and crowd avoidance desired benefit of the landscape recreator, a rule was developed that prohibits a landscape agent from stopping if there are more than five other agents present at the cultural or geologic features. Landscape agents are programmed to avoid crowds at all costs. They will speed up if they have 50% energy remaining, to pass other agents on the trails if they are within fifty meters and travelling slower. This rule conforms to what the some of the hikers and mountain bikers told us about their trail experiences and adds to the physical challenge that they sought.

To accommodate the needs of the social hiker and biker visitors, the corresponding agents are programmed to hike or ride to areas in the simulated landscape to learn more about natural features and to socialize with other agents. The agents in these classes generally spend at least two hours performing that activity. They have lower desires for extremely challenging physical fitness, but will seek out areas where they can spend time, such as at the cultural or geologic sites. If a social agent encounters a small group or perceives the ability to catch up to another social agent they will increase their speed to do so. They will remain with them throughout the duration of the simulation, unless they expended too much energy and will be forced to slow down and rest. These rules conform to what the hikers and mountain bikers reported about the type of behaviors they exhibit on the trails. Social agents will stop at all cultural or historic features no matter how many other recreators are present.

Since there are a four different recreator age groups being represented in the simulations, they all will move at different rates along the trail, some will run out of energy sooner (older ones) and will be forced to rest. The behavior rules for these social agents (rules 1, 4 & 5) are summarized in Tables 3 and 4.

The behavior of individuals involved in the jeep tours are not as robust as those hiking and mountain biking. Jeeps contain between four and six visitors on each trip. These visitors spend from two to three hours on the jeeps interacting with the driver and other tourists. Jeep agents are modelled as a group of passengers in one jeep which vary their jeep speed according to the topographic

Negative Encounters Between and Within Recreator Groups							
	Other Hikers/Hikers		Other Bikers		Other Jeeps		Freq. Resp ^a
	of Hikers	Freq.	of Bikers	Freq.	of Jeeps	Freq.	
Hikers	70	29%	75	30%	101	41%	246/338 (72%)
Bikers	80	58%	6	4%	53	38%	139/393 (35%)
Total	150		81		154		

^aThe number of times respondents identified negative detractors (conflicts) per total number of respondents in that group

Table 5: Recreation Conflicts Between and Within Recreator Groups

conditions. They speed up or slow down according to the degree of slope. Jeep agents conform to only one action rule that defines their behavior and that is Rule 4 (Table 3) to stop at all cultural or geologic features. The time they spend at these features is predetermined and conforms to what the jeep tour drivers typically spend at each location.

6 Simulating Conflicts

Solitude seeking is an important reported desire, goal or expectation in this study. The degree of interference with that goal is related to the number of encounters one has with other recreators. Table 5 illustrates that perceived negative encounters with certain types of recreators have some impact on the experience. The degree of that impact is not yet known. However, from the comments on the survey such as “seeing too many people”, “too many people”, “too many jeep tours on trails”, “seeing jeeps along the trails”, suggests that both visual and physical encounters are important measures of the degree to which a goal or desired benefit is interfered with. Based on the fact that both the hikers and mountain bikers reported that they were not able to obtain a crowd avoidance benefit, and the number of negative encounters that were reported, it was hypothesized that visual and physical encounters with other recreators interferes with their ability to acquire those desired benefits. For the purposes of this research, visual encounters within and between agent types are used as a measure of goal interference or the inability to achieve a desired or perceived benefit. The hypothesis being that the higher the degree of crowding induced encounters, the ability of the agent to obtain some of the other desired benefits may decline, lessening the perceived quality of the overall recreational experience.

From a management perspective what is needed is to identify the spatial locations along the trails where there are significant visual encounters. To accomplish this task, each agent keeps track of the number of encounters it has in each cell along the trail and also stores the type and number of visual encounters it has with other agents on other trails. These encounters are summarized, graphed and mapped to examine areas where there are levels of encounters that interfere with the recreator or agent’s goal to obtain a desired benefit.

7 RBSim — Recreation Behaviour Simulator

RBSim is a computer program that simulates the behavior of human recreators in high use natural environments. Specifically RBSim uses concepts from recreation research and artificial intelligence (AI) and combines them with geographic information systems (GIS) to produce an integrated system for exploring the interactions between different recreation user groups within geographic space. RBSim joins two computer technologies:

- Geographic Information Systems to represent the environment
- Autonomous human agents to simulate human behavior within geographic space.

RBSim is experimental at this stage, but demonstrates the potential of combining the two technologies to explore the complex interactions between humans and the environment[22, 24, 21, 23]. The implications of this technology should also be applicable to the study of wildlife populations and other systems where there are complex interactions in the environment.

RBSim uses autonomous agents to simulate recreator behavior. An autonomous agent is a computer simulation that is based on concepts from Artificial Life research. Agent simulations are built using object oriented programming technology. The agents are autonomous because once they are programmed they can move about the landscape like software robots. The agents can gather data from their environment, make decisions from this information and change their behavior according to the situation they find themselves in. Each individual agent has it's own physical mobility, sensory, and cognitive capabilities. This results in actions that echo the behavior of real animals (humans) in the environment.

The process of building an agent is iterative and combines knowledge derived from empirical data with the intuition of the programmer. By continuing to program knowledge and rules into the agent, watching the behavior resulting from these rules and comparing it to what is known about actual behavior, a rich and complex set of behaviors emerge. What is compelling about this type of simulation is that it is impossible to predict the behavior of any single agent in the simulation and by observing the interactions between agents it is possible to draw conclusions that are impossible using any other analytical process.

RBSim is important because until now, there have been no tools for recreation managers and researchers to systematically investigate different recreation management options. Much of the recreation research is based on interviews or surveys, but this information fails to inform the manager/researcher how different management options might affect the overall experience of the user. For example if a new trail is introduced, we might expect that conflicts might be reduced, but to what extent? If we go to a system of scheduling use, what is the impact on the number and frequency of users? More importantly when you have different, conflicting recreation uses, how do different management options increase or decrease the potential conflicts?

None of these questions can be answered using conventional tools. These questions all pivot around issues such as time and space as well as more complex issues such as inter-visibility between two locations. By combining human agent simulations with geographic information systems it is possible to study all these issues simultaneously and with relative simplicity.

8 RBSim Object Model

RBSim is developed using object oriented programming technology. Figure 5 shows a diagram of the principle components of the simulation program.

RBSim is comprised of five major components:

1. A Graphical User Interface for **Model Parameterization**. This is comprised of a set of forms for setting values for the remaining components described below.
2. **Output Classes** including:
 - (a) **the video display** showing a shaded relief map as a backdrop to the agent type and location displayed as graphic objects during the simulation run, and
 - (b) a **file object** for saving simulation statistics.
3. **Object Classes**, including
 - (a) the **recreator class** which represents a generic recreator class,
 - (b) a **trail class** which represents the trail as a list which contains the location, elevation and viewpoints at each trail location, and

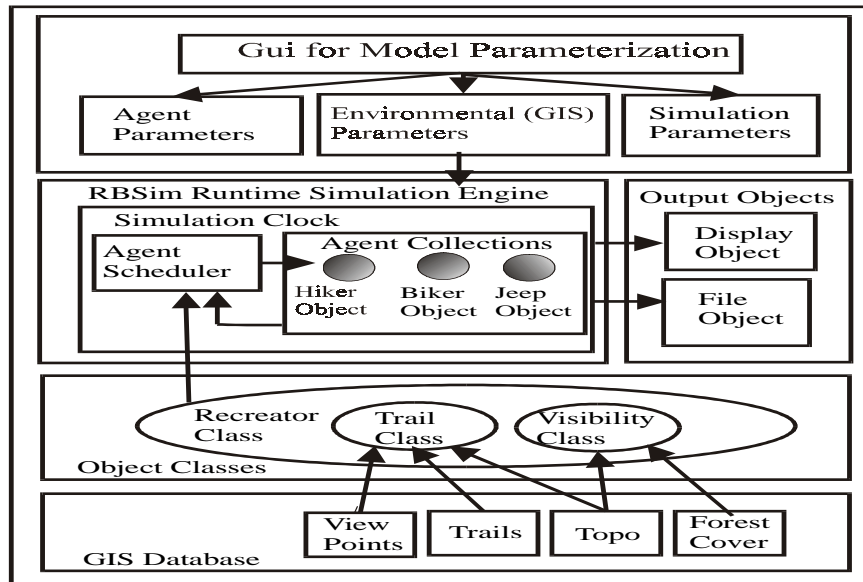


Figure 5: Recreation Behavior Simulation System (RBSim)

(c) the **visibility class** which provides the visual system for the recreator class

4. The GIS database which is used to parameterize the trail and visibility classes

8.1 Model Parameterization

RBSim allows the user to specify the following parameters for the recreation agents:

1. The total number of agents in each class (landscape hikers, landscape bikers, social hikers, social bikers, and jeeps)
2. The age distribution of recreators in the hiker and biker classes
3. The frequency within which each recreator begins a journey through the trail system.

8.2 The Recreator Class

The recreator class is the most complex object class. It is comprised of a set of properties for age and personality type (landscape or social agent). These properties determine the behavioral rules the agent will follow and the mobility and energy levels of the agent. Behavioral rules relate to how the agent responds to views, and the number and type of other recreation agents. These rules may result in the agent changing speed to overtake or catch up with other agents, slowing down and stopping to rest or spend time at important landscape features or viewpoints.

Hiker and Bike agents also have a system of energy levels programmed. Energy levels and speed of travel are related to the age of the agent. Very young, and older agents will move more slowly than agents in other age groups. In addition, as energy is expended during the simulation, these agents will also need to rest to rebuild energy levels. The length of resting time is determined by the estimated time it takes agents of different age groups to recover. Energy expended is calculated incrementally as the recreator moves along the trail. Uphill travel expends more energy than downhill travel. Resting times are randomized between preset time thresholds to represent variability between real human recreators.

8.3 The Visibility Class

Since much of the perception of crowding is based on visual contact as well as physical contact with other recreationists, a vision system is designed for the agents. The visibility class is a modification of standard GIS line of sight or intervisibility analysis. To reduce the computations required to check for visibility between two points, the visibility class checks for inter-visibility only between points occupied by other recreators referenced in the trail object. The line of sight is calculated taking into account intervening screening effects of topography and vegetation from the GIS databases for elevation and forest cover.

8.4 The Trail Class

Trails are specified for each agent, for each run. The trails are stored in the trail object which is constructed as a linear list of cells derived from a grid based GIS. For each trail cell the distance from the trail head, the elevation, and landscape features associated with the cell is stored. During the simulation run the trail object also stores the number recreators in each cell. This data structure is designed to minimize the computing time for agent navigation through the trail system. All agents of the same class (hiker, biker or jeep) share the same trail object. The trail object therefore acts as a “collective memory” for the agents in that respective class. Each agent can reference the trail object to determine the location of other agents on the trail and to determine the trail conditions. As the agent moves from one cell to the next, it de-references its location from the last cell (by subtracting one from the recreator count field for that cell) and references its location in the next cell. Since the hikers, bikers and jeeps follow different trails a unique trail object is created for each recreation type. To test management alternatives for new trails, the user may specify different trail files for each simulation run.

8.5 The Runtime Simulation Engine

The RBsim runtime simulation engine runs in discrete time steps. At each time step in the simulation, each recreator class (hikers, bikers and jeeps) is evaluated to determine if a new instance (agent object) of that class should be created. For each class of recreator a timer is set which begins incrementing from the start of the simulation run and is reset to zero each time a new recreator agent is generated. In the model parameterization the minimum and maximum times between agents is specified. A random time is generated between the minimum and maximum time each time a new agent is generated. A new agent of the respective class will be generated once the timer reaches the randomly generated time.

The new agent object is generated as an instance of the generic recreator class. When the agent is created, properties are set for age, personality, and agent type. These properties are set based on a randomly generated number (between 0 and 1) which sets the probability for each property. For instance, if 25 percent of the biker agents are of the landscape personality type and 75 percent are of the social personality type, then if the random number is between 0 and .25 the simulation engine will generate a landscape bike agent. If the number is greater than .25 and less than or equal to one, the simulation engine will generate a social bike agent. This same strategy applies to the age distribution as well.

Recreator Agents of the hiker, biker and jeep types are placed in collections for each type. The simulation engine then tracks each agent in each collection. Since the simulation engine is running on a synchronous clock, the order in which the agents are executed will affect consequences such as crowding and visibility. In order to avoid order effects from executing agent movement in a set sequence, the sequence is randomized within each collection for each iteration of the simulation. Each agent has a single method called “Move” which triggers the execution of the internal rules, energy levels and mobility for that agent. Once the agent has completed execution of all its behaviors for that time step, the runtime simulation engine then executes the move method for the next agent in

	Days of the Week Visitors Frequent Canyon							Time Period	
	Mon.	Tues.	Wed.	Thur.	Fri.	Sat.	Sun.	AM	PM
Hiker	4%	7%	5%	11%	15%	40%	18%	55%	45%
Mountain Bikers	6%	7%	4%	9%	16%	41%	17%	44%	56%
Jeeps	23%	24%	9%	4%	7%	14%	19%	40%	60%

Table 6: Frequency of Visits to the Canyon on a Weekly Basis and Time of Day

the randomized list for that iteration. This process continues in a loop until either all agents have completed their journey or the maximum time set for the simulation run is reached.

At the conclusion of the simulation each trail object writes its contents to the output file object. RBSim then returns control to the user.

9 Synthetic GIS World and Inherent Spatial Assessment Capabilities for Each Agent

The synthetic world that the simulated recreators utilize is a georeferenced, raster database consisting of 513 rows \times 522 columns, each cell 10 meters square. The database consists of topography, vegetation, adjacent primary and secondary roads, existing and proposed trails, trail head, jeep staging area, significant geologic features and scenic stops. These geographic themes were deemed important for this work, but many more could be incorporated as the sophistication of the modeling increases.

The approach taken in this research was to provide each agent with spatial analytical capabilities that is imperative to them processing information necessary for functioning in the simulated worlds. Each agent is provided with the ability to calculate distance or proximity to other agents and significant features in the landscape. Each calculates the percentage of slope from the topographic map and whether it is going up or down hill and in turn speeds up or slows down accordingly. They utilize neighborhood functions to identify trail cells or the location of significant geologic features and scenic stops. Most importantly each agent has visual capabilities for detecting other agents, how far away they are and whether they can or cannot be seen. This algorithm uses forest cover and topography as constraints to detecting other agents.

10 Applying RBSim for Simulating Typical Use Days and Management Alternatives in the Canyon

In order to test some of the ideas and concepts presented in this research and to determine the efficiency of the simulation system in identifying conflicting recreation behavior, a set of experiments were constructed. During the interview and survey phase of this research, visitors to Broken Arrow Canyon were asked in addition to the information already discussed, to record the month, day and time that they entered the canyon.

Table 6 presents a statistical summary of the visits of those recreators sampled over the duration of the study. As can be seen, peak times throughout the canyon are weekends. Over 40% of hikers and bikers frequent the canyon during these time periods, while 40% visit in the morning, with 60% in the afternoon. Week days according to our sampling were the highest for Jeep tours into the canyon. In order to test RBSim, many simulation runs were undertaken mimicking various peak and off-use times to examine the dynamic interactions of recreators and resulting visual and physical conflicts. This study reports on one of those experiments, a mid weekday (Wednesday) which typically is not a peak use period but contains a moderate number of hikers and bikers and relatively low jeep usage. The schedule of use during that typical day is reported in Table 7.




Figure 6: Runtime Simulation Interface with both the Original and Alternative Trail Layouts for Broken Arrow Study

In order to demonstrate a potential management action such as restricting biking use on a heavily used trail, two alternative bike trails were substituted for the original and the simulations rerun to evaluate the differences in recreational use and resulting perceived conflicts from all recreators perspectives (See Figure 6).

Wed	Biker Data (AM)				Hiker Data (AM)				Jeep Data (AM)			
10/11	1-20	21-40	41-60	> 60	1-20	21-40	41-60	> 60	1-20	21-40	41-60	> 60
9:00									1	1	0	1
9:00									2	0	0	1
9:00									0	2	0	1
9:00									2	0	0	1
9:00									0	2	0	1
10:30	1	2	0	0								
10:30	1	2	0	0								
11:30	0	1	0	0								
12:00	2	2	0	0								
12:00	1	2	0	0								
12:00	3	2	0	0								
1:30	0	2	0	0	0	0	6	7				
Percent	38%	62%	0%	0%	0%	0%	46%	54%				

Table 7: Typical mid-week entrance times by recreators into Broken Arrow Canyon




Figure 7: Graphed Results of Hiker Encounters with Other Agents from along Hiking Trail




Figure 8: Graphed Results of Biker Encounters with Other Agents from along Biking Trail

11 Initial Results of Simulations Runs

Table 7 is an example of mixed recreational use along the trails. Figures 7 illustrate the intensity of hiker encounters with other agents from the hiking trails. Figure 7 illustrates a significant number of encounters with both other hikers and bikers, hikers versus other hikers and bikers and jeeps. Encounters with bikers is high from the beginning of the simulation, peaks at Chicken Point and is chaotic until completing the journey. Hikers, on the other hand, peak at Chicken Point and then remain consistently high thereafter. What is of interest is that where the encounters with hikers peak, biker encounters drop off and visa versa.

As summarized in Table 5, over 40% of the negative encounters that occur to hikers are with jeeps and 30% with other hikers or bikers. It is interesting that even with the number of hikers, bikers and jeeps included in this simulation, that there are very few encounters with jeeps. On the other hand, the high amount of conflicts with bikers and other hikers may have a detrimental effect on the recreation experience. But because the place, time and duration of encounters that occur between biker and hiker agents are not consistent, this may tend reduce to accumulative impact of the encounters on those hiking.

Figures 8 illustrate biker encounters with hikers, jeeps and other bikers from along the biking trail. The patterns are similar to those found in Figure 7 except bike encounters steadily increase

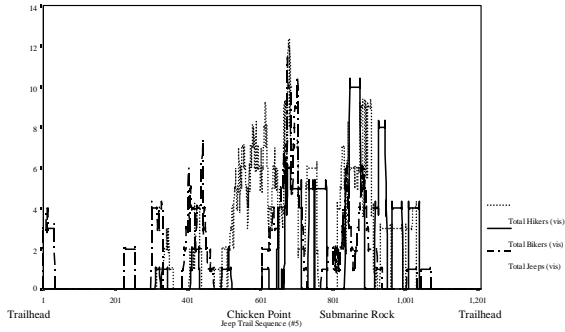


Figure 9: Graphed Results of Jeep Encounters with Other Agents from along Jeep Trail




Figure 10: Graphed Results of Hiker Encounters with Other Agents from along Hiking Trail




Figure 11: Graphed Results of Biker Encounters with Other Agents from along Biking Trail




Figure 12: Graphed Results of Jeep Encounters with Other Agents from along Jeep Trail

throughout the life of the simulation, but are not as high as in Figure 7. Biker’s encounters with hikers are more sporadic than is outlined in Figure 7 dropping off at the end of the simulation. Like above, there are virtually no encounters with jeeps.

Figure 9 illustrate a high number of jeep encounters with hikers and bikers and only a minimal number of encounters with other jeeps. The encounters occurring with hikers and bikers are concentrated around Chicken and Submarine Rock during the last half of the trip with virtually no encounters occurring for the first and last quarter of the simulation.

In summary it appears that with the increased number of hikers and bikers in the canyon that encounters with bikers are the most dominant impact. There are very few encounters with jeeps throughout all the simulations in experiment 4.

12 Simulations Using Alternative Bike Trail #1

As illustrated in Figure 10, selecting alternative bike trails can have a major impact on the number of encounters that occur along the trails. It can be seen in Figure 10 that when alternative bike route 1 is used in the simulations that the number of biker encounters that the hikers will have decreases significantly to the point that they are negligible after Chicken Point. When compared to Figure 7 and summarized in Table 8, by altering the trail layout the mean number of encounters has dropped by two thirds and the maximum number of encounters by half.

In Figure 11 the number of visual encounters with other recreators that bikers will have when using

	Exp. 4		Exp. 7 (4a)		Exp. 9 (4b)	
	Mean	Max	Mean	Max	Mean	Max
HVIS						
Hike	7.0	20	6.9	20	6.9	22
Bike	6.4	25	2.1	13	2.6	10
Jeep	2.9	15	2.7	17	2.7	18
BVIS						
Hike	5.1	18	1.4	17	1.9	18
Bike	5.5	24	4.5	12	2.0	12
Jeep	.51	7	.41	8	.67	7
JVIS						
Hike	2.1	12	2.1	7	2.3	15
Bike	1.2	10	.4	18	.33	5
Jeep	1.3	11	5.1	19	4.5	18

Table 8: Comparison Between Existing and Alternative Bike Routes for Experiments 4, 7 and 9

the alternative bike route reveals a dramatic decline in both hikers and jeeps, but a steady increase in number of bikers. In fact, an evaluation of Table 8 illustrates that visual encounters with hikers declines to one fifth of those that occurred in Figure 8, with the same number of hikers still using the trails. This strongly suggests that by using the alternative trail, the distribution of hikers and bikers within the canyon is more conducive to minimizing conflicts.

Figures 10 illustrate a significant number of encounters with jeeps from both bikers and hikers in the canyon. As in Figures 8 & 9, encounters with other agents declined. Of significance are the encounters with hikers and jeeps. But interestingly enough, increasing the number of bikers from eleven to twenty seven has little effect on the mean number of encounters that occur, but does effect the maximum. In other words while the number of encounters remains the same, the encounters are more evenly dispersed along the trail, rather than peaking at specific locations. From a management perspective, if the objective is to disperse the impacts of encounters over time and reduce high impact areas of conflict, then this alternative bike route would offer a solution to this problem.

13 Discussion

The results of this research illustrate however that the conflicts most often reported are from bikers having negative encounters with hikers. While jeeps are certainly considered to have a fairly high level of impact on both hikers and bikers, they are not as strong a determinant of a negative recreational experience as anticipated. Bikers and hikers continually clash in the canyons. What is of interest however is how often and where these encounters occur. An examination of the results of the agent simulation runs illustrates that bikers most frequently clash with other bikers. While bikers may have more encounters with other bikers, as reported in the survey, they do not see them as detracting from their experiences.

The agent simulations seem to be an excellent method for modeling recreator encounters and ultimately conflicts. While statistical results of the survey used in this study provide an indication of the average number of encounters (viewed as negative detractors), the agent simulations provide a dynamic view of these encounters and identify the spatially explicit locations where they occur. The effect of these encounters on the overall recreational experience is still unknown. However, this simulation environment provides a way to test and evaluate many scenarios of recreational use. While the survey provided a quantitative measure of the recreational experience, the simulation environment provides a dynamic, spatial representation of use and provides the added benefit of collecting and storing data on encounters over time. Both these data can be evaluated using conventional statistical

techniques and compared to explain commonalities and differences.

Of interest in this research, and showing the power of using simulation, is the impact of alternative routes on recreator encounters. An examination of the biker trail alternatives as suggested by the respondents to the survey, illustrate the importance of a well thought out trail design on recreational encounters. As can be seen in this research both alternative trail designs significantly reduce the number of encounters with other recreators. In fact, from before the turn around point to the completion of a biker's journey, they literally have no visual contact with any other recreator type. If hikers do have an accumulated negative effect on a bikers experience then it is clear that these alternative routes would alleviate this problem. This situation is identical when assessing encounters using the alternative jeep trail layout. Biker and hiker encounters with jeeps virtually disappear for over fifty percent of the journey. This is a substantial decline in encounters considering the significant number reported by the respondents. It is clear that the simulation environment can assist in evaluating existing and proposed trails in an attempt to minimize encounters and conflicts which ultimately lead to a decline in recreational experience.

What is imperative to emphasize in this work is that simulation of any dynamic behavior cannot be accomplished without such techniques as developed in this research. Simulation using personality traits and behavioral rules synthesized from human recreators provide a forum to evaluate and test a diversity of recreator use densities over time. These alternatives can be used to develop new facilities along the trails, and to redirect trail use to maximize user satisfaction while minimizing impact. Being capable of seeing the agents interacting under a variety of constraints can assist the manager in acquiring a better understanding of how human recreators use and interact on public lands.

This research has taken the first step forward to make linkages between GIS, Multiagent systems and Recreation Behavior Modeling. While this research has not directly dealt with goal interference theory, it does use it as a foundation for behavior modeling. It is assumed that perceived benefits, and obtaining or maintaining those benefits, directly correlates with the goal interference. Encounter and/or subsequent conflicts are the main cause of goal interference. The landscape recreator agent developed in this work was programmed specifically to avoid interference and when threatened passed other agents or avoided stopping at scenic lookouts. This technique allows one to reduce the amount of goal interference, while maximizing benefits. More than that, it allows one to experiment with artificial recreators to realistically determine thresholds of goal interference and devise management strategies to reduce it. This is one of the advantages of using simulation and the power of such multiagent environments.

14 Conclusions

This research advances our knowledge and understanding of natural resource assessment and intelligent simulation systems in the following ways: Extends the theoretical foundation of recreation and behavior by exploring the concept of benefits-based management for measuring desirable and obtainable benefits of leisure and assessing spatially-explicit visual and physical encounters among recreators in Broken Arrow Canyon; Extends the knowledge-base of the development, calibration and use of intentional, multi-autonomous agent systems in GIS represented worlds; Develops an entirely new form of intelligent decision support system (IDSS) to assist natural resource managers in assessing and managing human use of natural areas which could be easily extended into a number of other areas such as assessing impacts on wildlife habitats; Expands the existing capabilities of visual operators found in GIS for providing all mobile agents with visual capabilities. This vision system is used for controlling agent movement, goal-seeking, determining locations and distances of potentially conflicting agents and could be easily modified for identifying significant landscape features; Utilizes conventional social science survey techniques with automated field methodologies for calibrating agent movement; Develops a user friendly, parameterized interface for experimenting with alternative trail layouts and a diversity of agent configurations under a variety of conditions.

Much work can be undertaken to improve the predictability and reliability of the modeling frame-

work. To expand our understanding of the dynamic physiological and psychological experience patterns, sampling methods could be used. Dynamic experience patterns can be empirically measured including such factors as visual acuity, focus of attention, mood, psychological benefits, coping strategies, norms of behavior, and physiological changes at strategic locations within a stratified set of landscape settings found throughout the study site. A methodology employing these techniques that provided a way for the visitor to stop, record and photograph landscapes of importance would provide valuable information and lead to improved understanding of the dynamics of recreation experience. It is important however to ensure that wherever and whenever the visitor records such information, that their explicit location is captured as well so as to be able to link these changes to physiographic settings.

To improve the modeling of social interactions in a physical environment it is imperative that a more thorough understanding be acquired on how humans translate information from the environment into meaningful actions. Human-like agent simulations are no different. Once the spatial information is communicated to an artificial agent it must then be translated from its objective form into the symbolic and cognitive framework from which affective human responses are derived. This area of research needs considerable attention, but will provide meaningful outcomes.

15 Acknowledgments

We wish to thank the USDA Forest Service, Rocky Mountain Forest and Range Experiment Station and the Coconino National Forest for their assistance in facilitating this research effort. We also wish to thank Dr. B.L. Driver of the Rocky Mountain Station for his helpful review and oversight of this project. This research was supported in part by funds provided by the Rocky Mountain Forest and Range Experiment Station, Forest Service, U.S. Department of Agriculture.

Note: Instructions for obtaining a free copy of RBSim can be downloaded²

References

- [1] J. Anderson and M. Evans. Intelligent agent modeling for natural resource management. *Mathematical and Computer Modeling*, 20(8):109–119, 1994.
- [2] G. L. Ball. Ecosystem modeling in GIS. *Environmental Management*, 18(3):345–349, 1994.
- [3] F. Bousquet, C. Cambier, and P. Morand. Distributed artificial intelligence and object-oriented modeling of a fishery. *Mathematical and Computer Modeling*, 20(8):97–107, 1994.
- [4] D. Briggs, J. Westervelt, S. Levi, and S. Harper. A desert tortoise spatially explicit population model. In *Proceedings of the Third International Conference on Integrating GIS and Environmental Modeling*, Santa Fe, January 1996.
- [5] J. H. Brown. Modeling ecological patterns and processes using agent-based simulations and GIS. In *Third International Conference Integrating GIS and Environmental Modeling*, Santa Fe, January 1996.
- [6] D. Bruns, B. L. Driver, M. E. Lee, D. Anderson, and P. J. Brown. Implementing benefits-based management. In *The Fifth International Symposium on Society and Resource Management*, Fort Collins, Colorado, June 1994.
- [7] R. Conte and N. Gilbert. Computer simulation for social theory. In Gilbert and Conte, editors, *Artificial Societies: The Computer Simulation of Social Life*, pages 1–15. UCL Press, 1995.

²<http://nexus.srn.arizona.edu/~gimblett/rbsim.html>

- [8] M. Csikszentmihalyi and I. S. Csikszentmihalyi. *Optimal Experience: Psychological Studies of Flow in Consciousness*. Cambridge University Press, New York, 1988.
- [9] M. Csikszentmihalyi and R. Graef. The experience of freedom in daily life. *American Journal of Community Psychology*, 8:401–414, 1980.
- [10] M. Csikszentmihalyi and R. Larson. Validity and reliability of the experience sampling method. *The Journal of Nervous and Mental Disease*, 175:526–536, 1987.
- [11] S. E. Daniels and R. S. Krannich. The recreation opportunity spectrum as a conflict management tool. In Joanne Vining, editor, *Social Science and Natural Resource Recreation Management*, pages 164–179. Westview Press, Boulder, CO, 1990.
- [12] D. L. DeAngelis, D. M. Fleming, L. J. Gross, and W. F. Wolff. Individual-based modeling in ecology: An overview. In *Third International Conference Integrating GIS and Environmental Modeling*, Santa Fe, January 1996.
- [13] C Dibble. Representing individuals and societies in GIS. NCGIA Santa Barbara <http://www.geo.wvu.edu/www/i19/dibble>.
- [14] B. L. Driver and S. R. Tocher. Toward a behavioral interpretation of recreation, with implications for planning. In *Elements of Outdoor Recreation Planning*, pages 9–31. University Michigan Press, Ann Arbor, 1970.
- [15] A. Drogoul and J. Ferber. Multi-agent simulation as a tool for studying emergent processes in societies. In Gilbert and Doran, editors, *Simulating Societies: The Computer Simulation of Social Phenomena*, page 127. UCL Press, 1995.
- [16] A. Ewert, D. Chavez, and A. Magil. *Culture, Conflict, and Communication in Wildland-Urban Interface*. Westview Press, Boulder, CO, 1993.
- [17] A. S. Fege, C. McCarthy-Ryan, L. Munson, and R. Schreyer. Managing visitor conflicts. In *Proceedings of Managing America's Enduring Wilderness Resource: A Conference*, Minneapolis, MN, September 1989.
- [18] N Ferrand. Modeling and supporting multi-actor spatial planning using multi-agents systems. In *Third International Conference Integrating GIS and Environmental Modeling*, Santa Fe, January 1996.
- [19] N Ferrand. Multi-reactive-agents paradigm for spatial modeling. In *contribution to the European Science Foundation GISDATA program, Spatial Models and GIS*. Taylor and Francis, 1996. to be published.
- [20] N.V. Findler and R. M. Malyankar. Emergent behavior in societies of heterogeneous, interacting agents; alliances and norms. In Gilbert and Conte, editors, *Artificial Societies: The Computer Simulation of Social Life*, pages 212–236. UCL Press, 1995.
- [21] H.R. Gimblett. *Simulating Recreation Behavior in Complex Wilderness Landscapes Using Spatially-Explicit Autonomous Agents*. PhD thesis, University of Melbourne, Parkville, Vic. Australia, 1997.
- [22] H.R. Gimblett, B. Durnota, and R.M. Itami. Spatially-explicit autonomous agents for modelling recreation use in complex wilderness landscapes. *Complexity International*, 3, 1996. <http://www.csu.edu.au/ci>.
- [23] H.R. Gimblett and R. M. Itami. Modeling the spatial dynamics and social interaction of human recreators using GIS and intelligent agents. In *MODSIM 97 - International Congress on Modeling and Simulation*, December 1997.

- [24] H.R. Gimblett, R. M. Itami, and D. Durnota. Some practical issues in designing and calibrating artificial human agents in GIS-based simulated worlds. *Complexity International*, 3, 1996. <http://www.csu.edu.au/ci>.
- [25] D. G. Green. Spatial simulation of fire in plant communities. In P. Wise, editor, *Proceedings of National Symposium on Computer Modeling and Remote Sensing in Bushfire Prevention*, pages 36–41, Canberra, 1987. National Mapping.
- [26] W. E. Hammitt. Visual recognition capacity during outdoor recreation experiences. *Environment and Behavior*, 19(6), November 1987.
- [27] L. K. Harris, H. R. Gimblett, and W. W. Shaw. Multiple use management: Using a GIS model to understand conflicts between recreationists and sensitive wildlife. *Society and Natural Resources*, 8:559–572, 1995.
- [28] S. E. Hormuth. The sampling of experiences in situ. *J. Personality*, 54:262–293, 1986.
- [29] R. B. Hull and W. P. Stewart. The landscape encountered and experienced while hiking. *Environment and Behavior*, 27:404–426, 1995.
- [30] R. B. Hull, W. P. Stewart, and Y.K. Yi. Experience patterns: Capturing the dynamic nature of a recreation experience. *J. Leisure Research*, 24:240–252, 1992.
- [31] R. M. Itami. Simulating spatial dynamics: Cellular automata theory. *Landscape and Urban Planning*, 30:27–47, 1994.
- [32] M. I. Ivy, W. P. Stewart, and C. Lue. Explore the role of tolerance in recreation conflict. *J. Leisure Research*, 24:348–360, 1992.
- [33] G. R. Jacob. Conflict in outdoor recreation—the search for understanding. *Utah Tourism and Recreation Review*, 6, 1977.
- [34] G. R. Jacob and R. Schreyer. Conflict in outdoor recreation: A theoretical perspective. *J. Leisure Research*, 12:368–80, 1980.
- [35] K. M. Johnson. Using statistical regression analysis to build three prototype GIS wildlife models. In *Proceedings of GIS/LIS '92*, pages 374–386, 1992.
- [36] T. A. Kohler, C. R. van West, E.P. Carr, and C.G. Langton. Agent-based modeling of prehistoric settlement systems in the northern American southwest. In *Third International Conference Integrating GIS and Environmental Modeling*, Santa Fe, January 1996.
- [37] R. Larson and M. Csikszentmihalyi. The experience sampling method. In H. Reis, editor, *New directions for naturalistic methods in the behavioral sciences*, pages 41–56. Jossey-Bas, San Francisco, 1983.
- [38] M. E. Lee and B. L. Driver. Benefits-based management: A new paradigm for managing amenity resources. In *The Second Canada/US Workshop on Visitor Management in Parks, Forest, and Protected Areas*, Madison, WI, May 1992. U. Wisconsin-Madison.
- [39] M. Manfredo. The comparability of onsite and offsite measures of recreation needs. *J. Leisure Research*, 16:245–249, 1984.
- [40] P. L. Owens. Conflict as a social interaction process in environment and behavior research: The example of leisure and recreation research. *J. Environmental Psychology*, 5:243–259, 1985.
- [41] J. L. Rechel. Geographic information systems modeling of wildlife movements across fire and urban disturbed landscapes. *Bull Ecol. Soc. Am.*, 1992.

- [42] M. T. Richards and T. C. Daniel. Measurement of recreation and aesthetic resources in southwestern ponderosa pine forests. In A. Teclé and W. Covington, editors, *Multiresource Management of Southwestern Ponderosa Pine Forests: The Status of Knowledge*, chapter 7. USDA Forest Service, Southwestern Region, 1991.
- [43] M. T. Richards and H. R. Gimblett. Recreation assessment of forestlands in ecosystem management: A conceptual model. Technical report, U.S. Rocky Mountain Research Station, 1995.
- [44] H. Saarenmaa, J. Perttunen, J. Vakeva, and A. Nikula. Object-oriented modeling of the tasks and agent in integrated forest health management. *AI Applications in Natural Resource Management*, 8(1):43–59, 1994.
- [45] R. Schreyer. Conflict in outdoor recreation: The scope of the challenge to resource planning and management. In Joanne Vining, editor, *Social Science and Natural Resource Recreation Management*, pages 12–31. Westview Press, 1990.
- [46] R. L. Slothower, P. A. Schwarz, and K. M. Johnson. Some guidelines for implementing spatially explicit, individual-based ecological models within location-based raster GIS. In *Proceedings of the Third International Conference Integrating GIS and Environmental Modeling*, Santa Fe, January 1996.
- [47] W. P. Stewart. Influence of the onsite experience on recreation experience preference judgments. *J. Leisure Research*, 24:185–198, 1992.
- [48] D. R. B. Stockwell and D. G. Green. Parallel computing in ecological simulation. In A. Jakeman, editor, *Proceedings of the Simulation Society of Australia.*, pages 540–545, Canberra, 1989.
- [49] H. E. Tinsley and D. J. Tinsley. A theory of the attributes, benefits and causes of leisure experience. *Leisure Sciences*, 14(3):195–209, 1986.
- [50] A. E. Watson, M. J. Niccolucci, and D. R. Williams. The nature of conflict between hikers and recreational stock users in the John Muir wilderness. *J. Leisure Research*, 6:372–385, 1994.
- [51] J. Westervelt and L. D. Hopkins. Facilitating mobile objects within the context of simulated landscape processes. In *Third International Conference Integrating GIS and Environmental Modeling*, Santa Fe, January 1996.
- [52] D. Williams, G. Ellis, N. Nickerson, and C. Shafer. Contributions of time, format, and subject to variation in the recreation experience preference measurement. *J. Leisure Research*, 20:57–68, 1982.

Linguistic rule extraction from neural networks for high-dimensional classification problems

Hisao Ishibuchi Manabu Nii and Kimiko Tanaka
Department of Industrial Engineering
Osaka Prefecture University
hisaoi/manabu/kimiko@ie.osakafu-u.ac.jp

Abstract

In this paper, we show how linguistic rules can be extracted from trained neural networks for high-dimensional pattern classification problems. In our rule extraction method, antecedent linguistic values such as “*small*” and “*large*” are presented to a trained neural network for extracting linguistic rules. Since input values are handled as fuzzy numbers, outputs from the trained neural network are also calculated as fuzzy numbers by fuzzy arithmetic. When linguistic values in the antecedent part of a linguistic rule is presented to the trained neural network, the consequent class and the certainty grade are specified by the corresponding fuzzy output vector. In this manner, all combinations of antecedent linguistic values are examined as input vectors to the trained neural network for extracting linguistic rules. While our rule extraction method works very well for low-dimensional pattern classification problems, there exist several difficulties in the application to high-dimensional problems. One difficulty is the exponential increase of the number of possible combinations of antecedent linguistic values. Another difficulty is excess fuzziness of the calculated fuzzy output vector, which prevents our rule extraction method from appropriately specifying the consequent class and the certainty grade. From the viewpoint of the understandability of extracted knowledge, it is also another difficulty that a large number of linguistic rules are extracted from the trained neural network. In this paper, we show how these difficulties can be remedied in our rule extraction method. Simulation results on a real-world pattern classification problem with many continuous attributes show that classification knowledge can be extracted from trained neural networks in an understandable form.

1 Introduction

The main goal of the design of classification systems is to maximize the prediction ability for unseen patterns. Thus the performance of a classification system is usually measured by its classification rate on unseen patterns. When multi-layer feedforward neural networks are applied to classification problems, they are trained by numerical data for maximizing their prediction ability. In general, neural networks are handled as black box models in the application to classification problems. That is, we do not know why a trained neural network makes a particular decision with respect to the classification of an unseen pattern. In the context of knowledge discovery and data mining, there are two goals (i.e., prediction and description) where description tends to be more important than prediction [1]. While neural networks have high prediction ability, their description ability is not high. That is, connection weights of trained neural networks do not directly show any classification knowledge to users in a human-understandable form. Many approaches [2–7] have been proposed to the rule extraction from trained neural networks. Since the rule extraction is an extremely difficult task for arbitrary configured networks [3], usually special network architectures and/or learning algorithms were assumed in those approaches. High-dimensionality of pattern classification problems makes the rule extraction more difficult. In this paper, we tackle this extremely difficult task: extracting linguistic rules from arbitrary trained neural networks for high-dimensional pattern classification problems. We do not assume any special network architectures or learning algorithms. Linguistic rules are extracted from standard feedforward neural networks trained for high-dimensional pattern classification problems.

In our former study [8], we proposed an extraction method of linguistic rules from trained neural networks. While it worked very well for low-dimensional pattern classification problems, there exist several difficulties in the application to high-dimensional problems. The main aim of this paper is

to improve the applicability of our linguistic rule extraction method to high-dimensional problems. As in our former study [8], we use linguistic rules of the following form for an n -dimensional pattern classification problem with continuous attributes:

$$\text{Rule } R_p: \text{If } x_1 \text{ is } A_{p1} \text{ and } \dots \text{ and } x_n \text{ is } A_{pn} \text{ then Class } C_p \text{ with } CF_p, \quad (1)$$

where $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is an n -dimensional pattern vector, p indexes the number of rules, R_p is the label of the p -th rule, A_{pi} 's ($i = 1, 2, \dots, n$) are antecedent linguistic values such as “small” and “large”, C_p is a consequent class and CF_p is a certainty grade. We assume that the meaning of each linguistic value is specified by its membership function. That is, we handle linguistic values as fuzzy numbers. In Fig.1, we show membership functions of typical linguistic values. In computer simulations of this paper, we use the five linguistic values in Fig.1 and “don't care” as the antecedent linguistic values A_{pi} 's.




Figure 1: Membership functions of typical linguistic values. (S:small, MS:medium small, M:medium, ML:medium large, and L:large)

In our rule extraction method, we assume that a neural network has already been trained. While we use a standard back-propagation neural network [9] for describing our rule extraction method, it is applicable to any feedforward neural networks. In our method, an n -dimensional fuzzy vector of linguistic values in the antecedent part of a linguistic rule is presented to the trained neural network. The corresponding fuzzy output vector from the trained neural network is calculated by fuzzy arithmetic. The consequent class and the certainty grade of the linguistic rule are determined based on the fuzzy output vector. In this manner, all combinations of antecedent linguistic values are examined as fuzzy input vectors to the trained neural network for extracting linguistic rules. When we have six antecedent linguistic values (i.e., five linguistic values in Fig.1 and *don't care*), the total number of possible combinations of the antecedent linguistic values is 6^n where n is the number of input units of the trained neural network.

Our rule extraction method can not be directly applied to high-dimensional problems. One difficulty is the exponential increase of the number of possible combinations of the antecedent linguistic values. It is impossible to examine all the 6^n combinations when n is large. Another difficulty stems from the increase of excess fuzziness of fuzzy outputs. The fuzziness of fuzzy outputs is increased by the feedforward calculation in the trained neural network because fuzzy arithmetic is used for the calculation of the fuzzy input-output relation of each unit (see Nii & Ishibuchi [10]). The excess fuzziness tends to increase as the size of neural networks increases. Large excess fuzziness of the fuzzy output vector prevents our rule extraction method from appropriately specifying the consequent class and the certainty grade of each linguistic rule. From the viewpoint of the understandability of extracted knowledge, it is another difficulty in the application to high-dimensional problems that a large number of linguistic rules are extracted from trained neural networks. In the context of knowledge discovery and data mining, it is easy to generate a large number of rules, but most of the generated rules are not useful or interesting for users [11]. In this paper, we show how these difficulties can be remedied in our rule extraction method.

2 Outline of Our Rule Extraction Method

Let us assume that a multi-layer feedforward neural network has already been trained for an n -dimensional pattern classification problem with continuous attributes. We also assume that a set of linguistic values is given by domain experts for each attribute of the pattern classification problem. Our task is to extract linguistic rules in (1) from the trained neural network. For explaining our rule extraction method, we use a three-layer feedforward neural network with n input units, n_H hidden units and c output units where c is the number of classes involved in the pattern classification problem. For the simplicity of explanation, we use the five linguistic values in Fig.1 and ‘‘don’t care’’ as antecedent linguistic values for all the n attributes of the pattern classification problem.

In the basic form of our rule extraction method [8], all combinations of antecedent linguistic values are examined as fuzzy input vectors to the trained neural network. The consequent class and the certainty grade of each linguistic rule are determined based on the fuzzy output vector from the trained neural network corresponding to the fuzzy input vector of antecedent linguistic values. When a fuzzy input vector $\mathbf{A}_p = (A_{p1}, \dots, A_{pn})$ is presented to the neural network, the fuzzy input-output relation of each unit is written as follows [12]:

$$\text{Input units:} \quad O_{pi} = A_{pi}, \quad i = 1, 2, \dots, n, \quad (2)$$

$$\text{Hidden units:} \quad O_{pj} = f(\text{Net}_{pj}), \quad j = 1, 2, \dots, n_H, \quad (3)$$

$$\text{Net}_{pj} = \sum_{i=1}^n w_{ji} \cdot O_{pi} + \theta_j, \quad j = 1, 2, \dots, n_H, \quad (4)$$

$$\text{Output units:} \quad O_{pk} = f(\text{Net}_{pk}), \quad k = 1, 2, \dots, c, \quad (5)$$

$$\text{Net}_{pk} = \sum_{j=1}^{n_H} w_{kj} \cdot O_{pj} + \theta_k, \quad k = 1, 2, \dots, c, \quad (6)$$

where w_{ji} and w_{kj} are connection weights, θ_j and θ_k are biases, and $f(\cdot)$ is the sigmoidal activation function: $f(x) = 1/\{1 + \exp(-x)\}$. Our neural network architecture in (2)–(6) is the same as the standard feedforward neural network [9] except that the input and output of each unit are fuzzy numbers. As in various studies on fuzzified neural networks [12–14], the fuzzy input-output relation of each unit is defined by fuzzy arithmetic [15] and its numerical calculation is performed by interval arithmetic [16] on α -cuts of fuzzy numbers. In Fig.2, we illustrate the fuzzy nonlinear mapping by the sigmoidal activation function. The consequent class C_p and the certainty grade CF_p are determined by the fuzzy output vector $\mathbf{O}_p = (O_{p1}, \dots, O_{pc})$ calculated by (2)–(6).




Figure 2: Fuzzy nonlinear mapping by the sigmoidal activation function.

When a crisp input vector $\mathbf{x}_p = (x_{p1}, \dots, x_{pn})$ is presented to the trained neural network in (2)–

(6), we have a crisp output vector $\mathbf{o}_p = (o_{p1}, \dots, o_{pc})$. In this case, \mathbf{x}_p is usually classified by a single winner output unit as follows:

$$\text{If } o_{pk} < o_{ph} \text{ for all } k\text{'s } (k = 1, 2, \dots, c \text{ and } k \neq h) \text{ then } \mathbf{x}_p \text{ is Class } h, \quad (7)$$

where o_{pk} is the crisp output from the k -th output unit of the trained neural network.

By directly extending this decision rule to the case of the fuzzy input vector $\mathbf{A}_p = (A_{p1}, \dots, A_{pn})$, we have the following decision rule:

$$\text{If } O_{pk} < O_{ph} \text{ for all } k\text{'s } (k = 1, 2, \dots, c \text{ and } k \neq h) \text{ then } \mathbf{A}_p \text{ is Class } h, \quad (8)$$

where O_{pk} is the fuzzy output from the k -th output unit of the trained neural network. In this decision rule, the inequality relation $O_{pk} < O_{ph}$ between the fuzzy numbers O_{pk} and O_{ph} should be defined. First let us define the inequality relation $[O_{pk}]_\alpha < [O_{ph}]_\alpha$ between the α -cuts as

$$[O_{pk}]_\alpha < [O_{ph}]_\alpha \iff [O_{pk}]_\alpha^U < [O_{ph}]_\alpha^L, \quad (9)$$

where $[\cdot]_\alpha^L$ and $[\cdot]_\alpha^U$ denote the lower limit and the upper limit of the α -cut $[\cdot]_\alpha$ of a fuzzy number, respectively. From (9), we can see that the inequality relation $[O_{pk}]_\alpha < [O_{ph}]_\alpha$ holds when there is no overlap between $[O_{pk}]_\alpha$ and $[O_{ph}]_\alpha$. Using the inequality relation in (9) for α -cuts, the decision rule in (8) can be rewritten for classifying the α -cut $[\mathbf{A}_p]_\alpha$ of the fuzzy input vector \mathbf{A}_p as

$$\begin{aligned} \text{If } [O_{pk}]_\alpha < [O_{ph}]_\alpha \text{ for all } k\text{'s } (k = 1, 2, \dots, c \text{ and } k \neq h) \\ \text{then } [\mathbf{A}_p]_\alpha \text{ is Class } h. \end{aligned} \quad (10)$$

Using a pre-specified threshold value β , we define the inequality relation $O_{pk} < O_{ph}$ between the fuzzy numbers as

$$O_{pk} < O_{ph} \iff [O_{pk}]_\beta < [O_{ph}]_\beta. \quad (11)$$

Thus the decision rule for the fuzzy input vector \mathbf{A}_p in (8) is rewritten as

$$\begin{aligned} \text{If } [O_{pk}]_\beta < [O_{ph}]_\beta \text{ for all } k\text{'s } (k = 1, 2, \dots, c \text{ and } k \neq h) \\ \text{then } \mathbf{A}_p \text{ is Class } h. \end{aligned} \quad (12)$$

In Fig.3, we illustrate this decision rule. As shown in Fig.3(b), there are many cases where a single winner unit can not be specified by the decision rule in (12). In those cases, the classification of the fuzzy input vector \mathbf{A}_p is rejected. This means that the linguistic rule with the antecedent linguistic values $A_{p1}, A_{p2}, \dots, A_{pn}$ is not generated because its consequent part can not be uniquely specified.




Figure 3: Illustration of the decision rule for the fuzzy input vector. In (a), the fuzzy input vector is classified as Class 3. In (b), the classification of the fuzzy input vector is rejected.

When the fuzzy input vector \mathbf{A}_p is classified as a particular class by the decision rule in (12), a linguistic rule with the antecedent linguistic values $A_{p1}, A_{p2}, \dots, A_{pn}$ is generated. Its consequent

class C_p is determined as the classification result. The certainty grade CF_p is specified by examining the classification of the α -cut $[\mathbf{A}_p]_\alpha$ of the fuzzy input vector \mathbf{A}_p for various values of α . In our computer simulations, we used 100 values of α (i.e., $\alpha=0.01, 0.02, \dots, 1.00$). We specify the certainty grade CF_p as

$$CF_p = 1 - \min\{\alpha \mid [\mathbf{A}_p]_\alpha \text{ is classifiable by (10), } \alpha = 0.01, 0.02, \dots, 1.00\}. \quad (13)$$

In this manner, we can generate a large number of linguistic rules by examining all the possible combinations of antecedent linguistic values. In the case of Fig.3(a), the consequent class C_p is Class 3 and the certainty grade CF_p is 0.65.

3 Application to High-Dimensional Problems

While our linguistic rule extraction method worked very well for low-dimensional pattern classification problems (see Ishibuchi & Nii [8]), it can not be directly applied to high-dimensional problems. In this section, we illustrate several difficulties in the application to high-dimensional problems using wine data with 13 continuous attributes (available from UC Irvine database). The wine data consist of 178 samples from three classes. As a pre-processing procedure, we normalized all attribute values into real numbers in the unit interval $[0, 1]$ in order to use the same set of linguistic values in Fig.1 for all the 13 attributes. Thus the wine data were transformed into a three-class pattern classification problem in the 13-dimensional unit cube $[0, 1]^{13}$. We trained a three-layer feedforward neural network with 13 input units, 5 hidden units and 3 output units using all the 178 samples as training data. Our task in this section is to extract linguistic rules from the trained neural network for the wine data.

The main difficulty in the application of our linguistic rule extraction method to high-dimensional problems is the exponential increase of the number of possible combinations of antecedent linguistic values. For example, the total number of possible combinations of six linguistic values is over 13 billion (i.e., $6^{13} = 1.31 \times 10^{10}$) for the 13-dimensional wine data. Since it is impractical to examine such a large number of combinations as fuzzy input vectors to the trained neural network, we have to restrict the number of examined fuzzy input vectors within a tractable magnitude.

A simple but promising idea for restricting the number of possible combinations of antecedent linguistic values is to try to extract only general linguistic rules with a small number of antecedent conditions. Let us define the length of a linguistic rule by the number of its antecedent conditions excluding “*don't care*” attributes. For example, the length of the following linguistic rule is two:

$$\begin{aligned} &\text{If } x_1 \text{ is } \textit{small} \text{ and } x_2 \text{ is } \textit{don't care} \text{ and } x_3 \text{ is } \textit{don't care} \\ &\text{and } x_4 \text{ is } \textit{large} \text{ then Class 3 with } CF_p = 0.95. \end{aligned} \quad (14)$$

This linguistic rule can be rewritten as “If x_1 is *small* and x_4 is *large* then Class 3 with $CF_p = 0.95$ ” by omitting the “*don't care*” attributes. In Table 1, we show the number of possible combinations of antecedent linguistic values for each rule length in the case of the 13-dimensional wine data. From this table, we can see that the number of general rules (i.e., short rules) is much smaller than that of specific rules (i.e., long rules). Since general rules are more understandable for users than specific rules, the extraction of only general rules has another advantage: high understandability of extracted knowledge.

Rule length	1	2	3	4	5	10
Combinations	65	1,950	35,750	446,875	4,021,875	2,792,968,750

Table 1: The number of possible combinations of antecedent linguistic values for each rule length in the case of the 13-dimensional wine data.

We tried to extract linguistic rules of the length three or less from the trained neural network. That is, we examined 37765 fuzzy input vectors of antecedent linguistic values ($37765 = 65 + 1950 + 35750$;

see Table 1). In our linguistic rule extraction method, we specified the threshold value β as $\beta = 0.01$. While we tried to extract linguistic rules by presenting each of those 37765 fuzzy input vectors to the trained neural network, we could not extract any linguistic rule. This is because fuzzy output vectors from the trained neural network involved large fuzziness. For example, when we tried to extract a linguistic rule of the length two with the antecedent condition “If x_1 is *small* and x_{13} is *medium small*”, we presented the following fuzzy vector to the trained neural network:

$$\mathbf{A}_p = (\text{S}, \text{DC}, \text{DC}, \text{DC}, \text{DC}, \text{DC}, \text{DC}, \text{DC}, \text{DC}, \text{DC}, \text{DC}, \text{DC}, \text{MS}), \tag{15}$$

where “S”, “DC” and “MS” denote “*small*”, “*don’t care*” and “*medium small*”, respectively. Since all attribute values were normalized into real numbers in the unit interval $[0, 1]$, “*don’t care*” was handled as the unit interval $[0, 1]$. In Fig.4, we show the fuzzy output vector from the trained neural network corresponding to the fuzzy input vector \mathbf{A}_p in (15). As we can see from Fig.4, each fuzzy output has large fuzziness (i.e., the fuzzy outputs fully overlap one another). Such a large overlap prevents our rule extraction method from specifying the consequent class and the certainty grade.




Figure 4: Fuzzy outputs from the trained neural network calculated by fuzzy arithmetic.

When fuzzy arithmetic is used for calculating the fuzzy input-output relation of each unit, the fuzziness of fuzzy outputs increases by the feedforward calculation in neural networks. Usually the fuzzy-arithmetic-based calculation of fuzzy outputs involves large excess fuzziness (see Nii & Ishibuchi [10]). Such excess fuzziness degrades the performance of our linguistic rule generation method. Since the numerical calculation of fuzzy outputs is performed based on interval arithmetic on α -cuts of fuzzy inputs, we may decrease the excess fuzziness by improving the accuracy of interval arithmetic. In Nii & Ishibuchi [10], we used a straightforward subdivision method [16] for decreasing the excess fuzziness of fuzzy outputs from neural networks. This method is illustrated in Fig.5(a). In this method, each element of an n -dimensional interval vector (i.e., α -cut of an n -dimensional fuzzy vector) is subdivided into multiple subintervals (say, L intervals). This means that the interval vector is subdivided into L^n interval vectors. Here we encounter the curse of dimensionality, again. Since the straightforward subdivision method can not be applied to high-dimensional problems, we use a hierarchical subdivision method illustrated in Fig.5(b). In this method, only a single element of an interval vector is subdivided into two subintervals at a time.




Figure 5: Illustration of subdivision methods: (a) simple subdivision, and (b) hierarchical subdivision.

Using the hierarchical subdivision method, we calculated the fuzzy output vector from the trained neural network corresponding to the fuzzy input vector in (15). In Fig.6, we show the shape of each fuzzy output. In this case, we can generate the following linguistic rule:

$$\text{If } x_1 \text{ is } \textit{small} \text{ and } x_{13} \text{ is } \textit{medium small} \text{ then Class 2 with } CF_p = 0.81. \quad (16)$$

In this manner, we examined all the 37765 fuzzy input vectors for generating linguistic rules of the length three or less from the trained neural network. The number of extracted linguistic rules of each length is summarized in Table 2. We can see from Table 2 that a large number of linguistic rules were extracted.




Figure 6: Fuzzy outputs from the trained neural network calculated by a hierarchical subdivision method.

Rule length	1	2	3
Combinations	0	188	7193

Table 2: The number of extracted linguistic rules of each length.

It is very difficult for users to manually examine all the extracted 7381 rules ($7381 = 188 + 7193$; see Table 2). Moreover those extracted rules did not have high classification ability. When we classified all the 178 samples in the wine data by the extracted 7381 rules, we had a 80.9% classification rate. We can use genetic-algorithm-based rule selection methods [17-20] for selecting only a small number of significant rules from a large number of extracted rules. For example, a genetic algorithm selected 17 linguistic rules that have a 100% classification rate.

4 Conclusion

In this paper, we showed three difficulties in the application of our rule extraction method [8] to high-dimensional pattern classification problems. Those difficulties are (i) the exponential increase of possible combinations of antecedent linguistic values, (ii) the increase of excess fuzziness of fuzzy outputs by the feedforward calculation in neural networks, and (iii) a large number of extracted linguistic rules. We bypassed the first difficulty by extracting only general linguistic rules with a small number of antecedent conditions. That is, we did not try to extract specific linguistic rules with many antecedent conditions. The second difficulty was remedied by improving the accuracy of interval arithmetic on α -cuts of fuzzy input vectors. The third difficulty was resolved by selecting only a small number of significant linguistic rules by genetic algorithms.

As shown in this paper, our approach can find a small number of general linguistic rules from trained neural networks without assuming any specific network architectures or learning algorithms. High understandability of extracted knowledge and high applicability are two characteristic features of our approach.

Acknowledgement

This research was partially supported by Yazaki Memorial Foundation for Science and Technology.

References

- [1] Fayyad, U., Piatetsky-Shapiro, G., and Smyth, P. “Knowledge discovery and data mining: Towards a unifying framework”, 1996 *Proc. of 2nd International Conference on Knowledge Discovery & Data Mining*, 82–88.
- [2] Fu, L. “Rule generation from neural networks”, 1994 *IEEE Trans. on Systems, Man, and Cybernetics*, **24**, 1114–1124.
- [3] Sestito, S. and Dillon, T. “Knowledge acquisition of conjunctive rules using multilayered neural networks”, 1993 *International Journal of Intelligent Systems*, **8**, 799–805.
- [4] Towell G. and Shavlik, J. W. “Extracting refined rules from knowledge-based neural networks”, 1993 *Machine Learning*, **13**, 71–101.
- [5] Andrews, R., Diederich, J., and Tickele, A. B. “Survey and critique of techniques for extracting rules from trained artificial neural networks”, 1995 *Knowledge-Based Systems*, **8**, 373–389.
- [6] Hayashi, Y. “A neural expert system with automated extraction of fuzzy if-then rules and its application to medical diagnosis”, In Lippmann, R. P., Moody J. E., and Touretzky, D. S. (eds.), 1991 *Advances in Neural Information Processing Systems 3*. Morgan Kaufmann, San Mateo, USA, 578-584.
- [7] Matthews, C. and Jagielska, I. “Fuzzy rule extraction from a trained multilayered neural network”, 1995 *Proc. of IEEE International Conference on Neural Networks*, 744–748.
- [8] Ishibuchi, H. and Nii, M. “Generating fuzzy if-then rules from trained neural networks: Linguistic analysis of neural network”, 1996 *Proc. of IEEE International Conference on Neural Networks*, 1133–1138.
- [9] Rumelhart, D. E., McClelland, J. L., and the PDP Research Group. 1986 *Parallel Distributed Processing*, MIT Press, Cambridge.
- [10] Nii, M. and Ishibuchi, H. “Fuzzy arithmetic in neural networks for linguistic rule extraction”, 1998 *Proc. of 2nd International Conference on Knowledge-Based Intelligent Electronic Systems*, **2**, 387-394.
- [11] Liu, B., Hsu, W., and Chen, S. “Using general impressions to analyze discovered classification rules”, 1997 *Proc. of 3rd International Conference on Knowledge Discovery & Data Mining*, 31–36.
- [12] Ishibuchi, H., Fujioka, R., and Tanaka, H. “Neural networks that learn from fuzzy if-then rules”, 1993 *IEEE Transactions on Fuzzy Systems*, **1**, 85–97.
- [13] Buckley, J. J. and Hayashi, Y. “Fuzzy neural networks: A survey”, 1994 *Fuzzy Sets and Systems*, **66**, 1–13.
- [14] Ishibuchi, H., Morioka, K., and Turksen, I. B. “Learning by fuzzified neural networks”, 1995 *International Journal of Approximate Reasoning*, **13**, 327–358.
- [15] Kaufmann, A. and Gupta, M. M. 1985 *Introduction to Fuzzy Arithmetic*, Van Nostrand Reinhold, New York.

- [16] Moore, R. E. 1979 *Methods and Applications of Interval Analysis*, SIAM Studies in Applied Mathematics. Philadelphia.
- [17] Ishibuchi, H., Nii, M., and Murata, T. “Linguistic rule extraction from neural networks and genetic-algorithm-based rule selection”, 1997 *Proc. of IEEE International Conference on Neural Networks*, 2390–2395.
- [18] Ishibuchi, H., Nozaki, K., Yamamoto, N., and Tanaka, H. “Selecting fuzzy if-then rules for classification problems using genetic algorithms”, 1995 *IEEE Trans. on Fuzzy Systems*, **3**, 260–270.
- [19] Ishibuchi, H., Nozaki, K., Yamamoto, N., and Tanaka, H. “Construction of fuzzy classification systems with rectangular fuzzy rules using genetic algorithms”, 1994 *Fuzzy Sets and Systems*, **65**, 237–253.
- [20] Ishibuchi, H., Murata, T., and Turksen, I. B. “Single-objective and two-objective genetic algorithms for selecting linguistic rules for pattern classification problems”, 1997 *Fuzzy Sets and Systems*, **89**, 135–150.

Self-Organisation in a Simple Pursuit Game

Ashley Tews
School of Computing Science and
Electrical Engineering,
University of Queensland
Tel: +61-7-3365-3985
Fax: +61-7-3365-4999
tews@elec.uq.edu.au

Raymond Lister
School of Computing and
Information Technology
University of Western Sydney, Nepean,
PO Box 10 Kingswood NSW 2747 Australia
Tel: +61-47-360-610 Fax: +61-47-360-662
r.lister@uws.edu.au

Abstract

This paper concerns dynamic team formation in multi-agent systems, where each agent determines its own action by observing the other agents. The test bed is a simple pursuit game. Identical mobile agents learn to form teams to most effectively catch individual stationary targets. The decision-making component is not manually encoded, but instead it self-learns, without an external teacher, by the method of temporal differences. Agents do not negotiate, nor do they explicitly form or communicate commitments. Instead, agents determine their own actions by a behaviourist approach: they reason only from the relative positions of other agents, not from an interpretation of the intent of those other agents. Experimental results demonstrate that the agents do coordinate their activities, by forming teams. In this paper, self-organisation is manifested in three forms: the formation of teams by the pursuer agents, the temporal difference algorithm by which agents learn to cooperate, and in the representation of the learnt knowledge as the weights in a perceptron.

1 Introduction

Games have become a common type of micro-world for the study of cooperation in multi-agent systems. One popular type is the pursuit game, where agents work in teams to catch targets. Such a pursuit game was developed by Grinton, Sonenberg, and Sterling[1]. They described a system in which identical mobile pursuer agents cooperated to catch individual stationary targets. That test bed provided a discrete event simulation in a simple grid-based world of hexagonal cells that wrapped around at the edges. Each cell either contained one agent, or was empty. Pursuer agents caught targets by surrounding the target. Targets appeared at random in the world. They disappeared when they were either captured (or “killed”) or when they expired at a random age. Pursuer agents did not seek to maximise the number of kills to which they directly contributed. Instead, the agents worked together in an effort to maximise the score gained by the group of pursuers as a whole. To coordinate the actions of the agents, Grinton et al. used Jennings’[2] elaborate framework of commitments and conventions. Individual pursuers communicated explicit commitment proposals to all other pursuers. In all experiments performed by Grinton et al., the behaviour of the pursuers was handcrafted. There was not a learning component to the system.

Tan[5] demonstrated that agents in a related problem domain can learn to cooperate via a reinforcement learning algorithm. Sen, Sekaran, and Hale[3] demonstrated that agents can not only learn to cooperate via such a reinforcement learning algorithm, but they can do so without communicating intentions and commitments, and indeed without even being aware of each other. We refer to this as a behaviourist approach. This may at first seem an unlikely formula for emergent cooperation, but in some respects it is well suited to cooperation. Specifically, without commitments there can be no “lies” (ie. broken commitments). The behaviour of any agent is determined only by what is known with certainty.

However, both Tan and Sen *et al.* used the Q-learning algorithm [8] for reinforcement learning. This algorithm has inherent scaling problems, because it represents the system state as a vector of

integers. It follows that the algorithm is not always well suited to systems with many state variables, as the number of states rises exponentially with the number of state variables, and so it may be difficult for a learning system to generalise from searching a small portion of possible states. Tan restricted his work to problems with the order of a few hundred states. Sen *et al.* noted that their systems converged on sub-optimal solutions, and they attributed this to an incomplete exploration of the state space

In this paper, we present the results for a pursuit game where the agents are trained by the Temporal Difference reinforcement algorithm [4], and resource requirements increase polynomially with problem size, while maintaining good generalisation, because the decision making component of the system is a single-layer perceptron.

2 Details of the Simplified Test Bed

Our test bed was modelled on that of Grinton *et al.*, with the following minor differences, which were made as implementation expediences, that have no bearing on our primary interest, the formation of teams. The cells in our world are square, not hexagonal, and form a 10 by 10 grid without wrap around. Pursuer agents may occupy the same square, and they capture a target by occupying its square in sufficient numbers, rather than by surrounding the target cell. There are always a fixed number of targets on the board (a *target set*). The user determines the number of targets at system initialisation. In any given simulation run, all targets have an equal fixed lifetime, which is determined by the user. If, at the end of a target set's lifetime, three or more pursuers occupy one of the target squares, then that target is deemed to have been "killed" or "captured". Alternately, if a target square is occupied by less than three pursuers, then that target is deemed to have "expired". The instant that one target set disappears from the board, another target set appears immediately, with each target at a random location. As in the work of Sen *et al.*, we use a behaviourist approach. At each time interval, one pursuer re-evaluates the target towards which it will take a step (of one square), without reference to either its own history of moves, or the history of any other pursuer. The decision of that pursuer is based only upon the distances at that time interval of every pursuer to all targets.

3 The Method of Temporal Differences

3.1 Background

There is a fundamental problem that must be addressed by any game-playing system that self-learns: the only certain measure of performance is the final result of the game. Given that, how can a learning algorithm decide which moves were good and which were bad? In some circumstances, a certain move may have been good, but an earlier bad move had rendered the game unwinnable: how is the system to avoid penalising such "innocent" moves? Even if the system could correctly identify the bad moves, there is no external instruction as to the move it should have made. This *credit assignment problem* is not unique to game playing systems. In fact, most related work has been done with control problems in mind, where the archetypal problem is the balancing of a hinged pole on a horizontally translatable cart: how is the system to decide which movements of the cart were good and which movements caused the pole to fall?

One common approach is to build a learning system that contains two sub-systems: a controller and a predictor. At each temporal interval, the controller calculates its possible actions (eg. left/right cart movements of various strengths) and selects the action that the predictor estimates is most likely to lead to the desired outcome. Initially, the predictions are little more than random numbers, but the predictor learns to improve its performance by attempting to minimise the difference in successive predictions. If the prediction at time t is P_t , where $0 < P_t < 1$, then the error for that time interval is a function of the difference between successive predictions (ie. $P_{t+1} - P_t$). The actual final outcome has either $P_t = 0$ or $P_t = 1$, which may for example be interpreted as "player loses" or "player wins"

respectively. The training system seeks to minimise the sum, over all time intervals, of the difference between successive predictions. Sutton [4] formalised this *Method of Temporal Differences*. In our experiments reported below, we used Sutton's TD(0) form of the algorithm.

Tesauro [6] trained a Multilayer Perceptron to play Backgammon, using the method of temporal differences (TD). His most advanced systems are estimated to play at the level of a strong international master, after 1.5 million games of self-play [7].

In single agent environments, like Tesauro's Backgammon player, the program maintains a queue of scenarios, one scenario for each point in the game. A scenario consists of the game configuration at that time (eg. the position of pieces on a Backgammon board) and the P_t estimate given by the predictor for that configuration. When the game is completed, and the outcome is known, the difference in each pair of consecutive P_t values, and the related board configuration, are used to increment the training of the predictor.

3.2 TD in a Multi-Agent Domain

The application of TD to our simple pursuit game is complicated by the existence of multiple agents. In most work on TD learning, there is only a single agent. In our multi-agent game, an individual pursuer may elect to leave a particular target for other pursuers, and move toward another target instead. Unlike single agent domains, where the outcome of the game depends upon the actions of that agent alone, the determination of whether it was correct for a particular pursuer to leave a target depends upon the subsequent actions of other pursuers.

Our solution to the above problem is summarised in Figure 1. We maintain, for each pursuer, T queues of P_t values, where T equals the number of targets. The queue for each target contains successive predictions for the outcome if the pursuer was, at time t , to take a step toward that target, irrespective of which target the pursuer actually steps toward. The various P_t values are stored upon their respective target queues, as part of a scenario. Each scenario stored on these queues consists of 1) a P_t value and 2) the board configuration that gave rise to that P_t value, as defined by the Manhattan distance of each pursuer to that target.

When given the opportunity to move, a pursuer always takes a step toward one of the T targets. The choice of which target to move toward is made as follows. The pursuer makes T requests to the predictor, one for each possible move. The pursuer then steps toward the target with the highest P_t value.

Table 1 gives the final P_t value when there are less than six pursuers and two or more targets. It is the relative values between each P_t pair in each table row that is of primary concern, rather than the absolute value of any single P_t value, since a pursuer moves towards the target with the larger P_t value. Case 1 of Table 1 gives the values for all pursuers when a target was killed. Without loss of generality, let that be target X. With less than six pursuers, only one target can be killed, and at least three of the pursuers have correctly moved to the same target. Such behaviour is to be reinforced. Thus all pursuers have a final X queue outcome of one, and an Y queue outcome of zero. Case 2 describes what should happen when insufficient pursuers move to the same target. Since it is impossible to decide in general which pursuers have behaved incorrectly, no training is possible, hence the dashes in this row of the table.

Table 2 gives the final P_t value when there are six or more pursuers and exactly two targets. Under such circumstances, the pursuers must break up into two teams. As with the previous table, it is the relative P_t values in each row of this table that are of primary concern, since a pursuer moves towards the target with the larger predictive value. Cases 1 to 3 of Table 2 give the values for when both targets were killed. Such collective behaviour is to be reinforced. Thus each pursuer that was on a target (ie. cases 1 and 2) is further encouraged to repeat its behaviour in future. The dashes in case 3 indicate that training is not done with any pursuer that proved surplus to requirements. In case 4, insufficient pursuers moved to one of the targets, target Y. Consequently, all pursuers are encouraged to move toward target Y under such circumstances in the future; even those pursuers that participated in killing target X.

Initialise the predictor.

REPEAT

Place T targets on the board, each at random locations.

Place the requested number of pursuers on the board, in random locations.

Initialise as empty the T queues of each pursuer.

REPEAT

FOR each pursuer

DO

FOR each target

DO

Make a tentative move toward that target, get a P_t value from the predictor for that board configuration, add this scenario to the queue for this pursuer/target, and then retract the move.

END

Move the pursuer toward the target that gave the highest P_t

END

UNTIL the lifetime of this set of targets is over

FOR each pursuer

DO

Add to the T queues

1. the final P_t as given by Table 1 or 2, and
2. the final board configuration.

For the T queues, separately, update the predictor's weights, if that is appropriate, using Sutton's TD(0) algorithm.

END

UNTIL the requested number of iterations is reached.

Figure 1: The learning algorithm

Case No.	Final Outcome	P_t for X queue	P_t for other queues
1.	Target X killed, other targets expired	1.0	0.0
2.	All targets expired	—	—

Table 1: The final P_t value for target queues when there are less than six pursuers

Case No.	Final Outcome	P_t for X queue	P_t for Y queue
1.	Both targets killed, and pursuer on target X	1.0	0.0
2.	Both targets killed, and pursuer on target Y	0.0	1.0
3.	Both targets killed, and pursuer not on target	—	—
4.	Target X killed, target Y expired, all pursuers	0.0	1.0
5.	Both targets expired, all pursuers	—	—

Table 2: The final P_t value when there are six or more pursuers and two targets

4 The Neural Network Predictor

The predictor is a single-layer Perceptron (*i.e.* no hidden units). The single output unit is a $[0,1]$ sigmoid which is interpreted directly as a prediction P_t . Unlike most networks, our output unit does not have a bias (as preliminary analysis and experiments indicated that a bias is redundant in this particular problem). There are as many input units as there are pursuers. All input units represent the $[0,1]$ normalized Manhattan distance from a pursuer to the target in question. The first input unit is the distance of the pursuer about to move, the *focal* pursuer. The remaining input units are mapped at random to the remaining *non-focal* pursuers.

The random mapping to input units of the non-focal pursuers illustrates a subtlety of the method of temporal differences, that may account for unpublished anecdotes of the failure of TD implementations. When using TD, the implementor must always bear in mind that the predictor will merely learn some regularity in the environment, and not necessarily the regularity of interest to the implementor. In our early work, we mapped non-focal pursuers to input units in a deterministic fashion. Numbers designate all pursuers and all input units. The non-focal pursuers were assigned, in ascending order, to input units in ascending order. Thus, pursuer 1 was on the same input unit in all but the single case when it is the focal unit. Pursuer 2 was on the same input unit in all but two cases, when either it or pursuer 1 was the focal unit, and so on. Because of that deterministic mapping, when there were more than 3 pursuers, the system never the less contrived to form a team out of pursuers 1–3, irrespective of their placement on the board.

5 Results for a Greedy Benchmark

To best assess the performance of a cooperative algorithm, a non-cooperative benchmark is required, to establish a lower bound on the performance expected of the cooperative algorithm. An obvious candidate is a greedy algorithm, where each pursuer moves toward its closest target, irrespective of the positions of the other pursuers. Figure 2 summarises the performance of that algorithm, with differing numbers of pursuers, from 3 to 7, but exactly two targets. For each target life and each number of pursuers, 500 target pairs were sequentially placed at random on the board. The percentage of targets killed was then plotted. The lines join together the points for equal numbers of pursuers. The small box to the right of the Figure gives a legend that describes which line corresponds to a certain number of pursuers.

With five or less pursuers, optimal performance for any algorithm is 50%, since only one of the two targets can be killed. The greedy algorithm achieves optimality when there are five pursuers, as at least three of any five pursuers must have a common nearest target. For fewer than five pursuers, however, performance is significantly lower than optimal. The greedy algorithm performs under 40% with four agents, and under 20% with three agents.

With six or seven pursuers, optimal performance is 100%. The greedy algorithm only achieves a performance of around 60% with six agents, and around 70% with seven agents.

6 Results for the TD-Perceptron

Each training run of the neural network consisted of one thousand iterations of the outer loop in Figure 1.

6.1 Runs with three or more pursuers and two targets

We found that learning was quite fast: the system began making good decisions within the first ten iterations. Figure 3 describes the performance of the system over many training runs, with two targets. Each data point represents a single training run.




Figure 2: The performance of a greedy algorithm with two targets, and varying numbers of pursuers




Figure 3: The performance of the TD-Perceptron system, for two targets and varying numbers of pursuers

With five or fewer pursuers, and unlike the greedy benchmark, the neural network achieves optimal performance, thus demonstrating that the pursuers are behaving in a cooperative fashion.

With six or seven pursuers, the neural network approach does not perform optimally, but it does perform significantly better than the greedy algorithm. Whereas the greedy algorithm only achieves a performance of around 60% and 70% with six and seven agents respectively, the learning algorithm achieves 80% and 90% respectively, thus demonstrating once again that the pursuers are behaving in a cooperative fashion.

We performed experiments with multi-layer network topologies, with up to sixteen hidden units in a single layer. None of these alternative networks performed better than the single-layer networks.

6.2 Runs with three pursuers and more than two targets

We found that learning was slower than in the previous case, typically taking a large percentage of the one thousand training iterations. Figure 4 describes the performance of the system over many training runs, for systems with between two and ten targets.

Since performance is recorded as a percentage of targets killed, optimal performance is 50% for two targets, 33% for three targets, and so forth, down to 10% for 10 targets. Given sufficient target life, all training runs converged to exhibit such optimal performance.

We have attempted experiments with six or more pursuers and more than two targets. In such runs, pursuers must not only form teams but then also select a target for each team. Preliminary results suggest that this extended problem is not-linearly separable, and thus beyond the capacities of a single-layer perceptron.

7 Analysis of a Trained Network

The neural network predictor is used to assess the priority of moving a given pursuer toward a given target. The higher the output value, the more likely that the pursuer will be moved toward that target. All input units represent the $[0,1]$ normalized Manhattan distance of a pursuer to that target. Consequently, input values vary from zero, when the pursuer is on the target, to one. The first input unit is the distance of the pursuer in question (*i.e.* the pursuer about to move). The remaining input units are mapped at random to the remaining pursuers.

We examined the weights from a network trained in a six-pursuer two-target simulation run. Since the network is a single-layer Perceptron, all connections are from an input unit to the output unit.




Figure 4: The performance of TD-Perceptron for 3 agents and 2 to 10 targets. All runs exhibit optimal performance, given sufficient target life.

When all these weights were rounded to two significant figures, an obvious pattern is revealed. The weight from the first input unit to the output unit was -4.5 , while all the other weights were 1.0 .

Consider the weight from the first input unit. It is a large negative weight, so the further the nominated pursuer is from the nominated target, the larger the input value, hence the less likely it is that the pursuer will move toward that target.

The remaining input weights are identical (to two significant figures), so they have the same effect. They are positive, so the further the remaining pursuers are from the nominated target, the larger the input values, hence the more likely it is that the pursuer in question will move toward that target.

8 Conclusion

We have described a learning algorithm for a multi-agent system, tested it on a pursuit game, and demonstrated that the agents learn to cooperate. The cooperation manifests itself as dynamic team formation. It is perhaps surprising that such a simple system could learn to manifest dynamic team formation. There is no communication of commitments between agents, indeed there is no concept of commitment in the system, and the learning component of the system is merely a single-layer Perceptron. Cooperation is not programmed explicitly into the system, but emerges from the learning process.

We suspect there is scope for extending our basic multi-agent temporal difference learning algorithm to incorporate non-behaviourist concepts. For example, system performance may improve if the inputs to the neural network were doubled, to include an indication of the likelihood of each pursuer electing to move, at its next opportunity, toward the target in question. Such an indication would be a function of the most recent P_t value on that target queue for that pursuer. Such an indication might be interpreted as a measure of commitment.

References

- [1] C Grinton, L Sonenberg, and L Sterling. Exploring agent cooperation: Studies with a simple pursuit game. In *Advanced Topics in Artificial Intelligence, (10th Australian Conference on Artificial Intelligence, AI'97)*, pages 96–105. Springer, 1997.
- [2] N Jennings. Commitments and conventions: The foundation of coordination in multi-agent systems. *The Knowledge Engineering Review*, 8:223–250, 1993.
- [3] S Sen, M Sekaran, and J Hale. Learning to coordinate without sharing information. In *Proceedings of the National Conference on Artificial Intelligence*, pages 426–431, 1994.

- [4] R Sutton. Learning to predict by the method of temporal differences. *Machine Learning*, 3:9–44, 1988.
- [5] Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the 10th International Conference on Machine Learning*, pages 426–431, 1993.
- [6] G Tesauro. Practical issues in temporal difference learning. *Machine Learning*, 8:257–277, 1992.
- [7] G Tesauro. Td-gammon, a self-teaching backgammon program, achieves master-level play. *Neural Computation*, 1994.
- [8] C Watkins. *Learning With Delayed Rewards*. PhD thesis, Cambridge University Psychology Department, 1989.

Part V
Economics

Trond Andresen
 Department of Engineering Cybernetics,
 The Norwegian University of Science and Technology,
 N-7034 Trondheim, NORWAY.
 Trond.Andresen@itk.ntnu.no

A Model of Short- and Long Term Stock Market Behaviour

Abstract

Real-world stock markets are volatile and expresses such traits as overvaluation, psychological moods, cycles and crashes. This paper develops and explores a fairly simple model which expresses these traits. The model is continuous and non-linear. It is developed in stages. In the initial stages it applies to the price dynamics of one type of stock only. Later on it is applied to a weighted price index of different stocks, to try to capture the dynamics of a stock exchange as a whole. The model is not based on micro individual agents, but on the market as a whole displaying composite behavior that is argued to be the aggregate result of individual agent behaviour. The idea is to make some behavioural assumptions, and then adjust parameters and explore whether realistic qualitative traits of stock market dynamics show up. From this follows that the model can make no claims whatsoever to predict when and how much a given stock market will rise or fall. Its purpose is instead to gain qualitative insights into the mechanisms of stock market behaviour.

1 Introduction

This paper is a systems engineer's attempt to understand and construct a model of stock market dynamics. The model is generic, i.e. it is not an attempt to model or predict the behaviour of a specific stock exchange. The main assumptions behind the paper are as follows:

It is considered a fact of life that a significant share of stock market behaviour consists of following the herd, "noise trading", "trend chasing", "technical trading" etc. Thus we do not engage in the discussion of "how can such stock market agents prevail, won't they be weeded out since they are non-rational?", as for instance in DeLong et al. [1]. Instead we hold that non-rationality is an observable trait of real-world stock markets. This is in the tradition of J. M. Keynes [2], who states: "...all sorts of considerations enter into the market valuation which are in no way relevant to the prospective yield" (p.152), "... It might have been supposed that competition between expert professionals... would correct the vagaries of the ignorant individual... However, ... these persons are, in fact, largely concerned, not with making superior long-term forecasts of the probable yield of an investment over its whole life, but with foreseeing changes in the conventional basis of valuation a short time ahead of the general public... For it is not sensible to pay 25 for an investment of which you believe the prospective yield to justify a value of 30, if you also believe that the market will value it at 20 three months hence." (pp.154–55), "... The social object of skilled investment should be to defeat the dark forces of time and ignorance which envelop our future. The actual, private object of most skilled investment today is to beat the gun, as the Americans so well express it, to outwit the crowd, and to pass the bad, or depreciating, half-crown to the other fellow." (p. 155).

We also assume that there exist "fundamentals", or an "anchor" stock price corresponding to a sustainable yield—a yield that is also "reasonable", compared to alternative financial instruments like bonds. Furthermore, it is assumed that the aggregate of agents have some sort of feeling for what these fundamentals are (at least when stock prices are far away from them), and that this is crucial for long-range stock market dynamics. This is at odds with Davidson [3], who holds that any stock price is just as likely to prevail as another, based on the view that the future is completely unpredictable at any instant in time.

The model is of the top-down category, thus not based on a population of "micro" individual agents in the artificial life tradition. It is of a continuous non-linear differential equation type, of

the market as a whole, displaying composite behaviour that is argued to be the aggregate result of individual agent behaviour.

This paper is quite different from that of for instance W. Brian Arthur et al. [4], both in its top-down paradigm, and in important behavioural assumptions. Arthur et al. conclude that small bubbles and crashes will occur (how often is not clear since they in their paper do not put any scale on the time axis) because agents occasionally, because of willingness to *experiment*, will lock onto each other in bubble-like excursions away from a “correct” price that would otherwise follow from a pure rational expectations scenario. Related to Arthur et. al. is a paper of Brock and Hommes [5], who argue that price volatility is due to agents choosing “cheap” but destabilizing trading strategies near the correct price, but more expensive, better and stabilizing (thus moving the price towards the correct level) strategies when prices are obviously off the mark. Common for both papers is the view that dynamics and volatility can be explained solely through endogenous mechanisms and a basically “rational” type of agent behaviour. This paper, as stated above, assumes that agents have significant irrational behaviour traits, and furthermore that dynamics are also a consequence of exogenous inputs in the form of events and mood changes.

The large differences in approaches to stock market modeling are understandable. Agent behaviour in a stock market is extremely complicated and heterogeneous. Any such model must because of necessary and strong simplification be fairly speculative. Therefore several angles of attack on this problem should be explored, and this paper is one such attempt.

The dynamics of the market are considered to be driven by three main demand components: One due to the market’s valuation of the firm’s real-economic prospects (the “anchor” value of the stock, see above), another due to short-term herd mentality (over minutes and hours), and a third due to long-term public mood shifts (over several years).

The model’s dynamics stem mainly from these internal feedback loops, but the market is also assumed to be driven by a sequence of exogenous stochastic “event pulses”. This pulse process accounts for both changes in macroeconomic conditions (for instance a change in interest rates), fresh information about different firms (news about quarterly earnings and similar), and individual agents taking action for some exogenous reason (as opposed to “bandwagon”—or other endogenously based decisions).

The model will be developed in four stages:

- A model of a rational market (with perfect information) — for one category of stock.
- A market with imperfect information and bandwagon effects — one category of stock.
- A market with imperfect information, bandwagon effects and long-term optimistic or pessimistic moods — for a stock exchange index.
- A market with imperfect information, bandwagon effects, long-term optimistic or pessimistic moods, and panics — for a stock exchange index.

The presentation is based more on block diagrams than (equivalent) (differential) equations. This choice follows from a conviction that insights into feedback effects and dynamics come much easier this way.

2 The “rational” market

Consider the block diagram in figure 1.

The symbols in the diagram are defined as follows (denomination is shown in brackets, empty brackets means that the corresponding entity is dimensionless):

p_r = Real or sustainable value of the stock [], expressed by the price/earnings ratio it can yield in the long run. At p_r the stock is neither over- or undervalued. For convenience we will use the term “price” or “value” in the following, even if we are talking about the p/e ratio.




Figure 1: A “rational” market




Figure 2: exponential response—“rational market”

p = current market price (*i.e.* p/e ratio) of stock [].
 \dot{p}/p = per cent change in price per day [% / day]. The dot implies differentiation.
 s = differentiation operator [day^{-1}].
 n = net current demand for stock [number of units]. This demand may be negative, *i.e.* when there is a net surplus of stocks offered.
 c_1 = constant factor [% / (number of units·day)] transforming net demand into price increase rate. The total number of stocks issued is incorporated in this factor. There is, as indicated in the figure, saturation in price decrease rate, since the surplus offered cannot exceed the total number of stocks issued. In the model, this is translated into a maximum rate of price decrease.
 c_2 = constant factor [number of units / %] transforming relative price deviation into corresponding net demand. The total number of stocks issued is incorporated also in this factor.
 The model in figure 1 is a simple first order non-linear differential equation,

$$\frac{\dot{p}}{p} = c_1 n = c_1 c_2 \left(\frac{p_r - p}{p} \right) \tag{1}$$

which holds above the negative saturation limit for \dot{p}/p . If we consider only the dynamics from p_r to p , the system is linear within saturation limits. But we stick to the non-linear formulation since we need to access surplus demand n at later stages.

A rational market implies that all agents have the same perfect information about p_r . Any change in p_r is responded to by each agent in the same manner. However, action is here assumed to be dispersed in time. At this stage we assume that agents receive perfect information, but they do not receive it simultaneously, or act instantaneously after receiving it. It follows from (1) that the time span needed for price adjustment is inversely proportional to the factors c_1 and c_2 . If we now consider the case of a stepwise increase in p_r , for instance because a technological breakthrough has occurred in the firm, then the resulting price adjustment path is a first order stable exponential step response as shown in figure 2. We have no overshoot, no oscillations, no unpredictable excursions, just a smooth and asymptotically perfect price adjustment. This is of course a completely unrealistic representation of what occurs in the real world. At the same time, it should be noted that this is the way a stock market *ought to* work, reflecting real-economic changes impacting the listed firms, and nothing else.

3 A market with imperfect information and bandwagon effects

At this stage we introduce two new phenomena. The first is that individual agents act upon imperfect and differing information. Depending upon whether he or she overvalues or undervalues the stock, the agent will demand too much or too little, compared to what demand would have been, based on a correct valuation. We assume that the distribution of erroneous information over all agents — also when accounting for agents’ different influence in the market — is such that the mean of aggregate erroneous demand is zero, and that the demand error follows a normal distribution around this zero mean. We also assume that individual agents’ errors in demand change with time. The choice is then

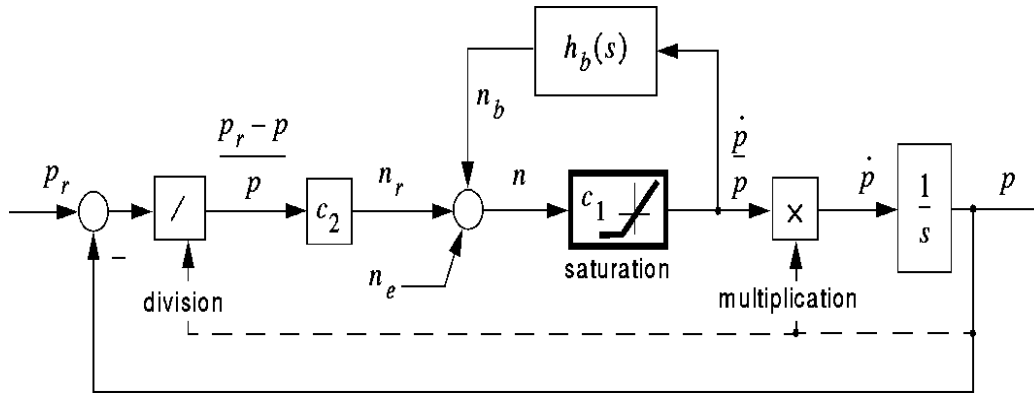


Figure 3: A market with bandwagon effects

to model aggregate demand error as a zero-mean normally distributed stochastic process (more on what sort of process later on).

The second phenomenon is the “bandwagon” effect. We explain this by referring to a block diagram of a stage 2 model, shown in figure 3:

Surplus aggregate demand is now assumed to consist of three components,

$$n = n_r + n_b + n_e \quad (2)$$

We have

n_r = Demand component due to agents being informed about the sustainable value of the stock.

n_b = Component due to agents watching price increase/decrease rate and doing “technical trading” based on this. The sustainable value of the stock has no direct influence on this component. Subscript b signifies “bandwagon”.

n_e = Component due to agents having different and erroneous information about the sustainable value of the stock. This is the zero mean stochastic process discussed above. Subscript e signifies “error”.

Furthermore, we have introduced a transfer function $h_b(s)$ in figure 3. This function decides the speculative component of market behaviour. There is a positive feedback through $h_b(s)$ from price increase rate to the surplus demand component n_b . If for instance \dot{p}/p is large and positive at a certain moment, many agents will jump on the bandwagon and buy now with the hope that prices will continue to rise. Of course some technical trading strategies are more elaborate than this, for instance action in counter-phase, *i.e.* buying when prices are falling in the expectation that they will rise later on. It is assumed however, that herd mentality is the dominant type of speculative behaviour. Involving Occam’s razor, the simplest transfer function that accounts for this, is

$$h_b(s) = \frac{K_b}{1 + T_b s} \quad (3)$$

Here T_b expresses the small time lag from acquiring price information to buying (or selling), that speculative action cannot get around. This lag is due to delays in acquiring information, considerations, and then getting the trading done. The gain K_b expresses how strong speculative action is, based on the available price change rate information.

Note that we say “action”, not “agents”. Individual agents may of course operate in a purely speculative or herd mode, others may again be pure “real investors”. But most have composite motives (real-economic more or less off the mark, and speculative). When the market as a whole is considered, however, this discussion becomes uninteresting, since the market as a whole must necessarily have a “composite motive”.

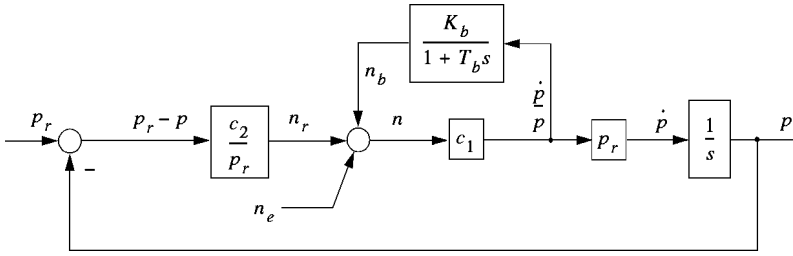


Figure 4: A linear model




Figure 5: Eigenvalues of linear model

As already stated, the surplus demand component n_e accounts for the aggregate effect of agents making erroneous and different assumptions about the stock value, but in such a way that the mean error in surplus demand is assumed to be zero. We now also incorporate the effects of differences in individual speculative behaviour into this noise process, since in reality each speculative agent will act according to a unique transfer function. It will not be linear and it will be complex, and its parameters and even structure will change with time. We have averaged out all this individual behaviour into the transfer function (3). We then posit that what is lost through this simplification may be assumed to be to a sufficient degree represented in the error process n_e . Thus this process is assumed to have two origins: Erroneous estimates of the stock's sustainable value, and modeling errors due to aggregation of the speculative (bandwagon) feedback path.

If we now consider a situation where the price of the stock is at its sustainable value, the market will have no real-economic incentive to trade. But trading will take place all the same. Individual more or less rational, more or less well-informed agents have their own assessment, and they trade also in this situation. In the language of our model, we may say that the error process is an exogenous input or disturbance that excites the system, so that the market is never in equilibrium, but fluctuates around it.

The model in figure 3 is non-linear. But if we consider small fluctuations around a constant sustainable value, it may be approximated by a linear model, see figure 4.

Since fluctuations are assumed to be small, we may ignore the saturation. And the blocks with division and multiplication in figure 4 may now be swapped with respectively constants $1/p_r$ and p_r . We have a linear system which is excited by the error process n_e . The transfer function from n_e to p is

$$h_{p,n_e}(s) = \left(\frac{p_r c_1}{T_b} \right) \frac{1 + T_b s}{s^2 + 2\zeta\omega_0 s + \omega_0^2} \tag{4}$$

where the undamped resonance frequency is

$$\omega_0 = \sqrt{\frac{c_1 c_2}{T_b}} \tag{5}$$

and the relative damping factor is

$$\zeta = \left(\frac{p_r c_1}{T_b} \right) \frac{1 + c_1 c_2 T_b - c_1 K_b}{2\sqrt{c_1 c_2 T_b}} \tag{6}$$

The two eigenvalues of the system are indicated in figure 5.

Consider a case where K_b is increased while T_b is held fixed, *i.e.* speculative action is stronger while the information/decision time lag remains the same. From (5) we see that ω_0 is independent of K_b , while ζ decreases with increasing K_b . In terms of figure 5, this means that the eigenvalues of the system move along the circle towards the imaginary axis. The system approaches the border of




Figure 6: Impulse response in price due to bandwagon feedback

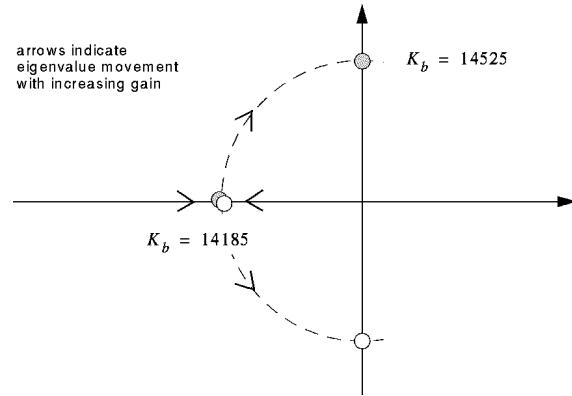


Figure 7: Eigenvalues move with increasing K_b

instability, which in the language of the market translates as increased volatility: For a given variance in the error process n_e , the variance in price will increase with K_b .

Figure 6 shows responses of the system to a small and short error process pulse, i.e. where some agents suddenly demand a certain amount of stock, even if the initial price is equal to the anchor value p_r . The amplitude of the pulse is 500 units of stock demanded out of 10.000. The pulse lasts 5 minutes, corresponding to $1/78$ of a trading day of 6.5 hours. (The 6.5 hour trading day is equal to that of the New York Stock Exchange).

The figure shows three price responses to this pulse, with parameters $\omega_0 = \text{constant}$, and $\zeta = 0, 0.4, 1$. Following (6), these values for ζ correspond to decreasing gain K_b .

In figure 6 — as opposed to figure 2 — the responses are from simulations where numerical values for system parameters have been chosen. This has been done by the following procedure: First, we choose a total number of stocks issued = 10.000, and a sustainable stock value $p_r = 10$. We may choose these values freely; the choices do not make any difference for the analysis to follow.

We assume that when all 10.000 units are demanded, respectively offered, on the market, this corresponds to a price change rate of $\pm 70\%$ per day. This decides the coefficient $c_1 = 7 \cdot 10^{-5}$. The lower saturation occurs for $n = -10.000$.

We then set $K_b = 0$, and are back to the stage one model. Only rational trading decisions are made, and we may posit some adjustment lag (see figure 2). We choose $T_r = 3$ [days], on the basis that such decisions are more carefully considered than technical trading decisions (see below), which are taken during fractions of a single day.

Since we have now chosen both T_r and c_1 , and $T_r = 1/(c_1 c_2)$ (see figure 2), we get $c_2 = 1/(c_1 T_r)$.

It now remains to decide the parameters K_b and T_b . Aiming for realism, we want the price dynamics due to the bandwagon loop to be very fast, in the order of a fraction of a day. And there should be at least one distinct overshoot (*i.e.* some volatility) before the response settles down.

We start by choosing ω_0 such that one day corresponds to fifteen full undamped ($\zeta = 0$) oscillations, see figure 6. When damping is increased to $\zeta = 0.4$, we get the impulse response delineated in bold in figure 6. We have some overshoot, indicating a certain amount of volatility. This is our choice for the dynamics of the bandwagon loop. The initial price response pulse due to technical traders jumping on the bandwagon lasts approximately 0.03 trading days = 12 minutes, and the reaction dies out after around 30 minutes. From the choice of ω_0, ζ the corresponding pair K_b, T_b is calculated from (5) and (6). (See end of paper for complete list of parameter values.)

At this stage, we emphasize that the above, and later, choices of parameter values, must necessarily be somewhat arbitrary. Our defense is that simulation experiments have demonstrated that similar

qualitative system behaviour shows up under a fairly wide spectrum of values.

There is an interesting insight that emerges from the model at this stage. Consider figure 7, which shows the the positions of eigenvalues (drawn as small circles) as a function of two values of gain

We observe that the system is marginally stable for $K_b = 14525$. On the other hand it is overdamped (non-volatile) for $K_b = 14185$! Since real stock markets are volatile, i.e. underdamped ($0 < \zeta < 1$), this suggests that there is some adaptive mechanism at work in the market to keep K_b in the surprisingly narrow band $14185 < K_b < 14525$, i.e. just under the instability border. We suggest that the mechanism is the following: On one hand, agents jump in to trade on volatile movements. But by their entry, they increase K_b and thus volatility. On the other hand they abstain when the market is too volatile (nervous market), thus decreasing K_b and volatility.

4 An index market with imperfect information, bandwagon effects and long-term optimistic or pessimistic moods

The model is to be further elaborated, but first we will argue that from now on we may consider it to represent not a specific stock, but a “composite stock” composed of stocks from all firms listed on a stock exchange, such that the composite stock price is proportional to the stock exchange index. We make some assumptions in connection with this:

- All categories of stock have similar dynamics.
- The p/e ratio for the composite stock is defined as being the total value of all stocks traded on the exchange, divided by the total sum of earnings.

These assumptions mean that the composite stock p/e ratio (from now on called an “index” or the “price” of the composite stock) will also fluctuate around p_r , with dynamics that are similar to those for one category of stock. The difference is that a price shock for one category of stock only, does not impact very strongly on the index.

The aggregation step from one category of stock to an index is comparable to the earlier step of aggregating all agents into one composite agent. We uphold all variable and parameter names and numerical values introduced for the stock model, with the note that they now pertain to the index. By now we are ready to consider an augmented model as shown in figure 8.

Again we hold that model imperfections and approximations, and fluctuations in demand for the composite stock implied by the index, is accounted for by a zero mean noise process as introduced earlier. But we have for the time being rescinded this stochastic excitation n_e , since the dynamics to be examined in this section will be shown not to depend upon being driven by an external input. We will reintroduce the noise input later on.

The additions to the model are indicated by the shaded area. Consider the block in the upper right-hand corner. The instantaneous price increase rate is an input to the block, a low pass filter with time lag T_f of a one year magnitude. The rationale for this filter is that the increase rate of “optimism” (or confidence, bullishness, positive “animal spirits” in the terminology of Keynes) is assumed proportional to the long-term trend in index increase. Hourly, daily, even weekly and (to some degree) monthly fluctuations are disregarded; there is a sluggishness in market mood. Price rise has to be persistent over a long time (years) before the market really picks up. On the other hand, when the price culminates and starts falling, the market will need a corresponding amount of time for such a change of affairs to sink in. The increase rate in optimism is set equal to the markets’ perception of the long-term index increase rate. Optimism is given a numerical value, and a range which is both positive and negative. (Thus pessimism corresponds to negative optimism.)

By now it should be clear why it has been necessary to transit from an individual category of stock to an index: Market mood is a function of the behaviour of the aggregate of all stocks, not one category only.

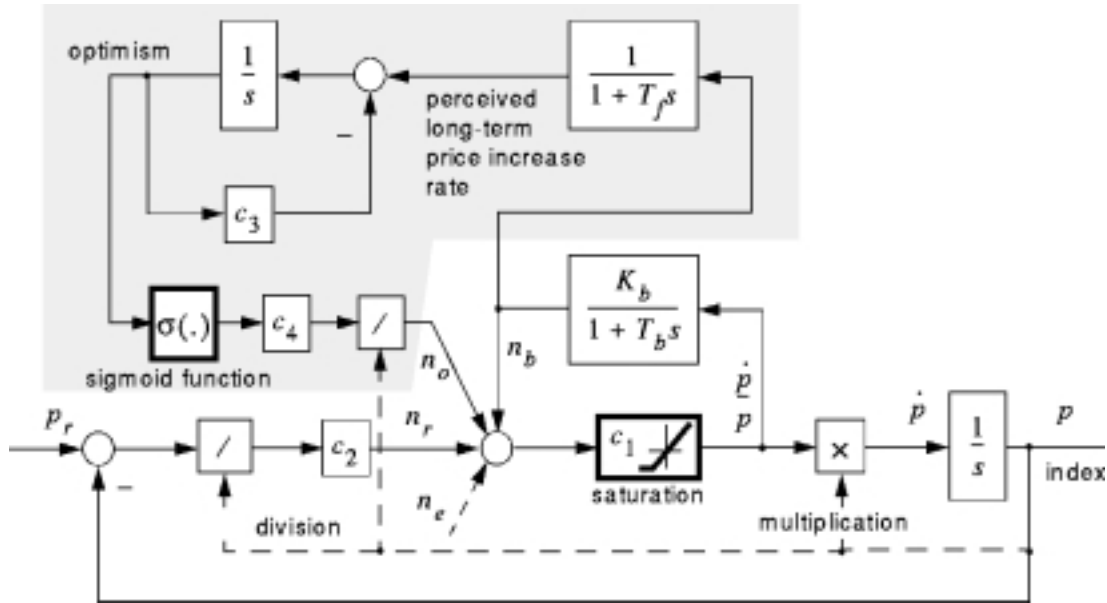


Figure 8: Model incorporating long-term mood dynamics




Figure 9: Optimism response to a pulse in price increase rate




Figure 10: Sigmoid function

In the absence of any perception of a long-term tendency for price to change (i.e. a flat price level over a long period), the current level of optimism will slowly erode to zero, through the factor c_3 . The argument for this is that the market will gradually forget its initial mood and tend towards a neutral attitude (zero optimism or pessimism) if the current mood is not maintained by a sustained increase or decrease in stock price.

We may now try an experiment. We isolate the two blocks in the upper part of the figure from the model, input a rectangular price increase pulse, and observe the response in optimism given by this model. We assume a one-year (defined as 250 trading days) constant price increase rate pulse. This pulse, and the corresponding response in optimism, is illustrated in figure 9.

The input pulse is not shown to scale. Parameter values for the simulation are chosen through a procedure described below. Note how optimism culminates after around 500 trading days (2 years). If we compare the time scales of figures 6 and 9, we observe a very large difference between the fast dynamics of the bandwagon loop, as opposed to the long-term mood loop. Thus we are dealing with a stiff system of differential equations.

To complete the explanation of the long-term mood loop, we now turn to the nonlinear sigmoid function $\sigma(\cdot)$. (The corresponding block in figure 8 is outlined in bold, to signify a nonlinear function.) The rationale for using the sigmoid is to account for upper and lower saturation in the system. The sigmoid function is shown in figure 10.




Figure 11: Cycles due to long-term mood dynamics

The function, including a gain c_4 introduced for convenience (see figure 8), outputs an additional demand component (measured in cash terms) which stems from the long-term mood of the market. It may be positive or negative. This money term translates into current additional composite stock demand n_o (“ o ” for optimism) through division by the current index. Upper soft saturation is assumed to make itself felt when the market is euphoric, i.e. optimism is at a maximum. The will to spend money on additional stock acquisition is there, but a very large amount of money has already been spent in the stock market, and fresh cash and credit is getting scarce. At the other extreme, when pessimism is at its maximum, agents holding stock are reluctant to sell it at the going bottom price, since there is very little to gain. This explains the lower saturation. The sigmoid function is expressed in the form

$$\sigma(x) = \left(\frac{p+n}{1 + \left(\frac{p}{n}\right) e^{-K_s[(p+n)/pn]x}} \right) - n \quad (7)$$

The parameters are:

p, n = The maximum saturation (asymptotic) value of $\sigma(x)$ is p . The minimum saturation is $-n$.

K_s = The slope of $\sigma(x)$ for $x = 0$. Thus K_s expresses the gain of $\sigma(x)$ for small excursions from a neutral mood, into optimism or pessimism.

Note that (7) has the necessary property $\sigma(0) = 0$, i.e. a neutral market mood results in zero additional demand. See also figure 10.

We have in this section introduced two first-order linear blocks and one non-linear relation, which together form the long-term mood loop. Figures 9 and 10 imply that numerical values for the parameters associated with this loop have been chosen. This is done in the following way: It is required that the system shall cycle regularly between euphoria and recession, but for the time not allowing panics to occur. The maximum price (i.e. p/e ratio) is set to approx. 25, the minimum to approx. 5. Furthermore, a full cycle (which is at this stage not shortened by panics on the downswing, see next section), shall last 10 years (one year is defined as 250 trading days). It is demanded that upswings shall have an essentially exponential growth shape, and that downswings shall be faster than upswings. Based on these conditions, all parameters have been decided together, through educated guessing and a comprehensive simulation-based trial-and-error process.

With the resulting choice of parameters, the cycles look like in figure 11. This is a stable limit cycle. In the next subsection we will see how this pattern is distorted by panics occurring near peak price levels, but we will also observe how this cycle remains an attractor that decides the fundamental long-run dynamics of the system. We emphasize that from this it follows that our model is *not* dependent upon a crash-and-subsequent-recovery mechanism for cycles to occur. The upswing is due to the spreading and self-reinforcing belief that “if I get in now, I can cash in my investment with a profit at a later stage”. But the upswing sooner or later has to culminate at some level, when the perception has spread sufficiently that current prices have grown too high in relation to the sustainable

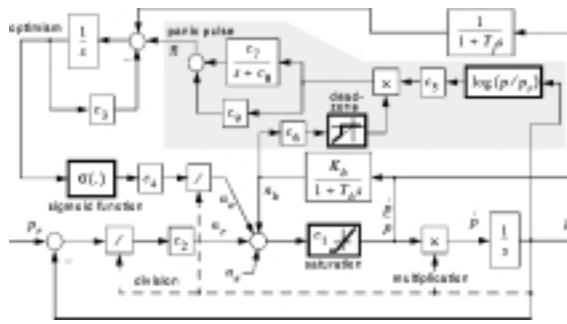


Figure 12: Model including panic mechanism




Figure 13: Two peaks with several crashes

value of the stock, and also due to increasing scarcity of additional money to invest in stocks. This reduces the growth rate of surplus demand and through this, the stock price growth rate. In the next round this reduces the growth rate in optimism, further reducing the demand growth rate, and so on. Optimism reaches a peak, and a feeling sooner or later overwhelms the optimistic mood of the upswing, and more stock is offered than demanded on the market. This further erodes optimism, and we are gradually into a downswing, which also needs time to build up momentum. This momentum causes the price to fall below sustainable value. But things will later pick up again when the earlier pessimistic mood is mostly forgotten, combined with a spreading recognition that the stock market is now generally undervalued. The upswing starts, and one cycle is completed. The period of a cycle is in the main decided by the inertia of market perception: It needs time to absorb a persistent tendency in price change, and it needs time to forget.

Note that our explanation for stock market cycles is a very psychological one. It has little connection with such other macroeconomic cycle explanations as Goodwin’s business cycle model [6] where cycles are due to worker-capitalist struggle over output, or build-up of indebtedness and related financial fragility (Minsky’s ‘financial instability hypothesis’ [7], as modeled and simulated by Keen [8], or Andresen [9]).

Our purpose is not, however, to contest the validity of these other models—it is to focus on one specific approach that may be contain some truth together with other approaches. But we suggest that the mechanisms presented here have become relatively more important as the financial sector has grown in size and influence, as unions have lost power, as the focus of the media and thus public opinion, has turned from finance as an instrument for enhancing production, to the financial market as a place for playing games for profit; i.e. Keynes’ “Casino”.

The reader may at this stage object that the downswing predicted by our model is very slow and well-behaved. Where are the panics, which may erase a substantial part of an index in a single trading day? In the next section we will extend the model to account for this.

5 A market with imperfect information, bandwagon effects, long-term optimistic or pessimistic moods, and panics

A modified model augmented with a panic mechanism is shown in figure 12. The additions to the model in the previous subsection (figure 8) are indicated by the shaded area in figure 12.

We designate this “the panic subsystem”. We note that this subsystem has two inputs: The ratio p/p_r , and relative price change rate, \dot{p}/p . The output is a negative pulse, designated π in the figure. Such a pulse has the effect of abruptly reducing optimism and thus the demand component n_o . We

will explain the panic subsystem by starting to the right in the shaded part of the figure, with the expression $\log(p/p_r)$. We have

$$\log(p/p_r) = \log(p) - \log(p_r) = \int_{p_r}^p \frac{dx}{x} = \int_{t_0}^{t_1} \frac{\dot{p}}{p}(t) dt \quad (8)$$

Here the current price and time is $p(t_1)$, and the system starts with $p(t_0) = p_r$, i.e. price level equal to sustainable, “fundamental” value. The logarithmic expression is equivalent to integrating relative price change rate \dot{p}/p , with initial value $p(t_0) = p_r$. Thus we have an expression that will be larger the further away from sustainable value the price is. This expression is assumed to give a measure of “wariness” in the market. By this we mean that agents will on the average be more sensitive to large downward blips in stock prices when the current prices are much higher than sustainable value. The wariness factor is multiplied with the output of a dead-zone block that ensures that only downward price blips above a certain magnitude are passed on (i.e. noticed by the market), as indicated by the graphics in that block. If prices are not too far away from sustainable value, however, even large downward movements are not considered to be danger signals. The filtered blips are passed on through a constant gain and transfer function in parallel, as shown in the left part of the shaded area, and then input to the optimism subsystem. The rationale for the transfer function, which has low-pass character with a time lag $1/c_8$ (of magnitude a couple of weeks), is to account for medium-term memory in the market of recent strong downward blips in an overvalued situation.

At this stage it should be noted that it is not self-evident that this panic sub-system will actually lead to panics and crashes. We have just made some reasonable behavioural assumptions about wariness etc. for agents, and implemented them in the model. But we will see that panics will occur.

Now to the task of deciding values for the parameters, this time for the panic sub-system. The procedure has been the same as described in the preceding sections: back and forth between simulations, educated guesses, parameter adjustments, new simulations. The resulting parameter value set is given at the end of this paper. The variance and character of the stochastic noise process n_e has also been decided as part of this selection process. Initially white noise was tried, which is uncorrelated with itself between sampling intervals (13 samples per trading day, i.e. a sample every half hour). This, however, gave price movements that displayed distinct swings only from hour to hour, but not over a couple of days. The white noise process was therefore substituted with white noise filtered through a first order low-pass filter with time lag = 3 trading days. This corresponds to a train of overlapping exponentially-tailed pulses exciting the system, and gave autocorrelated price dynamics over the week that (by visual inspection) resembled real-world index movements well. We emphasize that panics and crashes occur whether one employs white noise, correlated noise or a periodic pulse process to excite the system. Figure 13 shows a simulation with the chosen parameter values.

The simulations done were quite time-consuming: We wanted to simulate over a couple of peaks, where the time span is somewhat less than ten years = 2500 trading days, but where we also had to account for price movements which may change significantly from half hour to half hour. We thus had to use a stiff differential equation solver. To reduce simulation time, simulations were started from a system state somewhat into the peak (euphoric) phase of the first cycle. We observe from figure 13 that both the peak phases displays panics. Each panic is indicated by spikes, four during the first euphoric phase, two during the second. These spikes are the value of the variable π defined in figure 12 (note: the scale on the ordinate axis pertains to price only, the other variables are not shown to scale. For convenience, $-\pi$ is shown, so the spikes turn up positive in figure 13). π is zero when the negative price blip is within the limit in the dead-zone function, the “background noise level” that has to be surpassed before the market is alarmed. If it is bigger than this limit, however, its impact is decided by the value of the “wariness factor” (8), through multiplication with this factor. The product, a negative pulse, is then transmitted to the state that represents optimism, and this state is abruptly reduced. The trajectory for the optimism state is also given in figure 13. The abrupt decrease in optimism sets off a similarly abrupt decrease in demand, which again transmits an amplified negative price blip to the input of the dead-zone function. The loop is closed. We have positive feedback and a mechanism to explain panics. More on this further below.




Figure 14: Two peak phases over 300 days




Figure 15: Zooming in on the 10 days around a panic

Figure 14 shows magnified portions of the two peak parts of the price curve. Note the time scale, a range of 300 trading days.

We observe that the two big panics during each of the two peak phases result in respectively a 12% and 6% price fall in a very short time. We have by this established that the panic mechanism works as predicted.

To get a clearer picture, we stretch the time axis even further, to get a glimpse of dynamics over a few days. The result is shown in figure 15. One day is divided into 13 sampling intervals. An initial downwards blip that is large enough to pass the dead-zone function, lasts only one sampling interval. The further downside that may be observed for both peak phases, must therefore be a result of the aforementioned positive feedback loop.

If we ignore the panic events, and also trends, on the graphs in figures 14 and 15, and inspect the price excursions under normal conditions, we note that prices fluctuate between approx. ± 0.2 to ± 0.3 , which related to a level of 20 to 21, corresponds to ± 1.0 to ± 1.5 %. This is considered a realistic magnitude of day-to-day volatility.

If we compare the two euphoric phases in figures 13 to 14, we observe that they are quite different in character. This can only stem from differences in the random noise process that excites the system, since all other conditions at the start of a euphoric phase are similar. We will now discuss this proposition more closely by presenting a series of simulation runs where the only parameter that is changed, is the “seed” integer initiating the random generator. This means that we have exactly the same dynamic system, starting from exactly the same initial state, but excited by different realizations of the same stochastic process. The results for eight such runs are given in figure 16.

The integer listed on each graph is the corresponding seed integer. The upper left run is identical to the case already presented in figures 13 to 15. If we consider the eight graphs as a whole, the following observations may be made:

An early crash before price has gone to high, is a good thing in the sense that the crash will usually not be to big, as opposed to the case with seed = 2106, where a panic-free growth far above 20 results in a big 30% crash.

We also note that (small) crashes at an early stage contribute to reducing the danger of big crashes later. This, by the way, to some degree justifies the sometimes euphemistic term “correction” for crashes, so commonly used to calm market nerves.

Furthermore, we note that early crashes will not stop prices from continuing their rise later on, before culminating. This is due to the inertia of a general and still growing optimistic mood, a mood that needs more time and adversity to turn sour. On the other hand, if a big crash occurs when the euphoric phase is in a later and more mature stage, this crash will contribute to an earlier start of

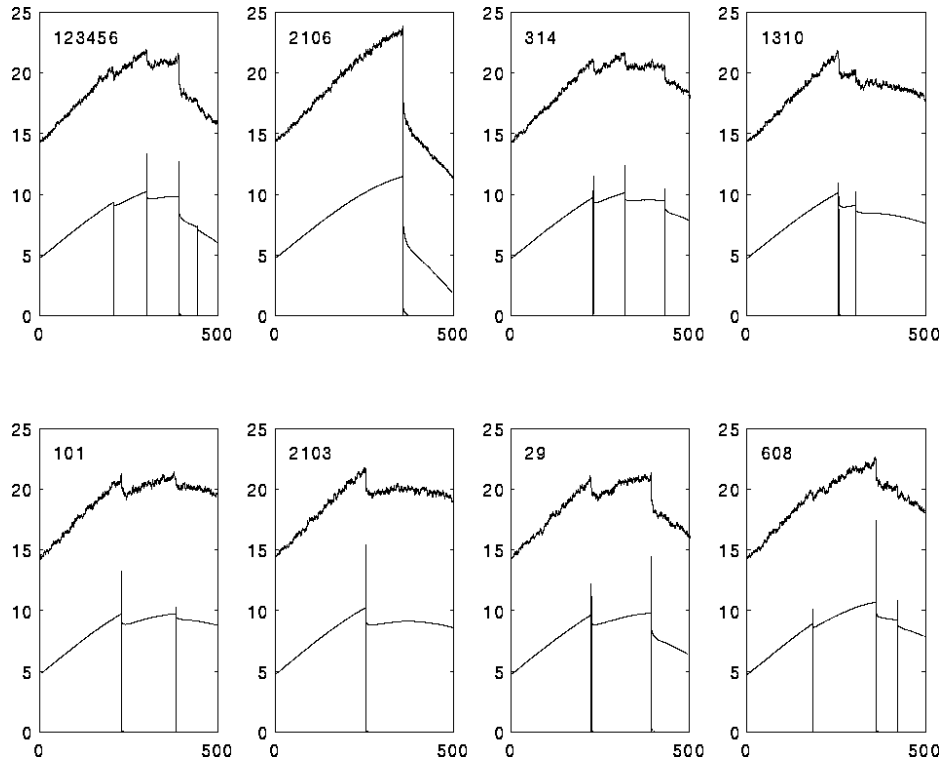


Figure 16: Euphoric phases for different seed integers

the inevitable general downside. Generally, the presence of panics and ensuing crashes will lead to a shorter period for the long-term cycle than what is displayed by the panic-free limit-cycle model in figure 11.

To gain insight into the character of the euphoric phase disturbed by panics, consider the graphs for the optimism state (also shown in the plots in figure 16). Note the “circular-saw-tooth” appearance of these graphs. After a crash optimism growth is slower, if the crash happens on the upswing. If the crash happens on the downswing, the downwards slide is steeper afterwards.

As mentioned, the case with seed = 2106 displays the biggest crash of the eight cases. At the bottom of the corresponding panic spike one may observe a small decaying exponential tail. This accounts for the medium term memory in the market of the crash, and stems from the low pass filter block with the coefficients c_7, c_8, c_9 in figure 12. Removing this block (and that memory effect), however, does not significantly change system dynamics.

6 Conclusion and suggestions

There are a couple of important mechanisms that are not incorporated in the current model. One phenomenon is “rallying”, which takes place over a couple of days, or even several weeks. This phenomenon may occur as a consequence of a feeling after a recent panic or strong fall, that now is the time to buy cheap because stocks have fallen too much and will rise—which of course is a self-fulfilling prophecy if enough agents think this way. Or it may take place due to exogenous impulses, for instance an announcement of lower Central Bank interest rates, or an announcement of a what is believed to be a credible IMF rescue operation towards an important country or group of countries.

In terms of our model, this may be taken care of by introducing additional parallel feedback

loops containing first- or second-order low pass transfer functions with time lags of a days/weeks magnitude, thus filling out the “gap” between the very short- and very long-term dynamics expressed by the corresponding loops in the current model. Such additional loops should enable the system to express swings over days and weeks, for instance rallying. Furthermore, one could look more closely into the modeling of exogenous impulses. In the current model, their impact all decay at the same rate, with a time constant of three days, and they arrive regularly (every half hour). A probably more realistic assumption is to let this noise be a Poisson process, thus generating impulses at irregular intervals, and more important—let these pulses decay at different rates, so that the model accounts for the fact that some news have more lasting influence on the market than other.

These suggested modifications will not, however, invalidate some insights and suggestions for further research that emerge from working with current model:

- A positive “bandwagon” feedback structure, and observed overshooting during fractions of a trading day, imply that the system is all the time very close to instability. Such a knife-edge balance can only be ensured by the market as a whole adaptively tuning the feedback gain in this loop.
- A pure psychological mood propagation mechanism, combined with a roof and a floor for demand, will by itself generate long-range cycles. The time lags of mood propagation and forgetting are decisive for the cycle period.
- Credible-looking panics and crashes can be generated by the mechanism described in this paper, which basically says that panics are triggered when two conditions are fulfilled: Gross overvaluation, and a random downward price blip that is so large that it “stands out”.
- By introducing market mood (“optimism”) as a system state, we have a possible means of modeling and simulating couplings between different stock markets. Obviously, what happens on the NYSE influences the other exchanges, and vice versa. Interactions can be realized by connections from price rate change in one model to the mood state of another separate stock exchange model.

A final note: I wish to express my thanks to Dr. Steve Keen, the University of Western Sydney, for stimulating discussions, help and suggestions.

7 List of parameters

- Sampling period (i.e. simulation step length) $T = 1/13 = 7.692308 \times 10^{-2}$

$T_r = 3$	$c_1 = 7.0 \times 10^{-5}$	$c_2 = 4.761905 \times 10^3$
$Kb = 1.424547 \times 10^4$	$T_b = 3.752636 \times 10^{-5}$	
- calculated from $\zeta = 0.4$, and $\omega_0 = 94.24778$

$T_f = 200$	$c3 = 1.4 \times 10^{-3}$	$c4 = 10^4$
$p = 15$	$n = 2.5$	$c5 = 0.1818182$
$c6 = 14.86250$	$c7 = 2.26 \times 10^{-2}$	$c8 = 0.15$
$c9 = 0.339$	deadzone = 0.18	
- variance of discrete white noise = 10^5
- filter time lag to make the noise correlated = 3,
- Initial values for the system’s three main states (faster system modes have initial values = 0):
 - optimism = 2.366744×10^{-1}
 - long-term mood filter state = 1.525×10^{-3}

– price = 1.434

- seed = 123456 in the first round, but is then varied as shown in figure 16.

References

- [1] J. Bradford De Long, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann (1991) “The Survival of Noise Traders in Financial Markets”, *Journal of Business* 64:1, pp. 1–20, January.
- [2] John Maynard Keynes (1973) “The General Theory of Employment, Interest and Money”, Macmillan, Cambridge University Press.
- [3] Paul Davidson (1998) “Volatile Financial Markets And The Speculator”, paper presented as the Economic Issues Lecture to the Royal Economic Society Annual Conference, Warwick, England, April, also in *The Journal of Economic Issues*, September.
- [4] W. Brian Arthur, John H. Holland, Blake LeBaron, Richard Palmer, and Paul Tayler (1997), “Asset Pricing Under Endogenous Expectations in an Artificial Stock Market”, in *The Economy as an Evolving Complex System II*, edited by W. B. Arthur, S. Durlauf, and D. Lane, Addison-Wesley.
- [5] William A. Brock and Cars H. Hommes (1997) “A Rational Route to Randomness”, *Econometrica*, vol. 65, pp. 1059-1095.
- [6] R. M. Goodwin (1967) “A Growth Cycle”, in *Capitalism and Economic Growth*, Cambridge University Press.
- [7] Hyman Minsky (1982) “The Financial Instability Hypothesis: An Interpretation of Keynes and an Alternative to ‘Standard Theory’”, in *Inflation, Recession and Economic Policy*, Wheatsheaf, Sussex.
- [8] Steve Keen (1995) “Finance and Economic Breakdown: Modelling Minsky’s ‘Financial Instability Hypothesis’”, *Journal of Post Keynesian Economics*, Vol. 17, No. 4.
- [9] Trond Andresen (1999) “The Dynamics of Long-range Financial Accumulation and Crisis”, forthcoming in *Nonlinear Dynamics, Psychology, and Life Sciences*, Vol. 3 No. 2, April.

Learning about the Cobweb

Carl Chiarella and Xue-Zhong He
 School of Finance and Economics
 University of Technology, Sydney
 PO Box 123 Broadway
 NSW 2007, Australia

carl.chiarella/tony.he1@uts.edu.au <http://www.bus.uts.edu.au>

Abstract

In this paper we consider how suppliers in a cobweb model may learn about their economic environment. Instead of assuming the one step backward-looking expectation scheme of the traditional linear cobweb model, we consider the subjective estimates of the statistical distribution of the market prices based on L -step backward time series of market clearing prices. With constant risk aversion, the cobweb model becomes nonlinear. Sufficient conditions on the local stability of the unique positive equilibrium of the nonlinear model are derived and, consequently, we show that the local stability region (of the parameters of the equation) is proportional to the lag length L . When the equilibrium loses its local stability, we show that, for $L = 2$, the model has strong 1 : 3 resonance bifurcation and a family of fixed points of order 3 becomes unstable on both sides of criticality. The numerical simulations suggest that the model has a simple global structure, it has no complicated dynamics as claimed recently by Boussard. However, complicated dynamics do appear when the model is modified with constant elasticity supply and demand.

1 Introduction

Consider the well-known cobweb model:

$$\begin{cases} p_{t-1,t}^* = aq_t + b & \text{(supply),} \\ p_t = \alpha q_t + \beta & \text{(demand),} \end{cases} \quad (1)$$

Here, q_t and p_t are quantities and prices, respectively, at period t , $p_{t-1,t}^*$ is the price expected at time t based on the information at $t - 1$, and $a, b, \beta (> 0)$ and $\alpha < 0$ are constants.

Instead of assuming the backward-looking expectation scheme $p_t^* = p_{t-1}$ as in Boussard [2], we rather assume that p_t^* is a random variable drawn from a normal distribution. Let \bar{p}_t^* and v_t^* be the mean and variance of p_t^* , respectively. With constant absolute risk aversion A , the marginal revenue certainty equivalent is ¹

$$\tilde{p}_t = \bar{p}_t^* - 2Av_t^*q_t. \quad (2)$$

Suppose a linear marginal cost, as in (1), so that the supply equation, under marginal revenue certainty equivalent becomes

$$\tilde{p}_t = aq_t + b. \quad (3)$$

Combining (2) and (3) and equating supply and demand gives the market clearing price in period t as a function of the subjective mean \bar{p}_t^* and variance v_t^*

$$p_t = \beta + \alpha \frac{\bar{p}_t^* - b}{a + 2Av_t^*}. \quad (4)$$

¹With constant absolute risk aversion A , we assume the certainty equivalent of the receipt $r = pq$ is $R(q_t) = \bar{p}_t^*q_t - Av_t^*q_t^2$. Then maximisation of this function with respect to q_t leads to the marginal revenue certainty equivalent $\tilde{p} = \frac{\partial R}{\partial q_t} = \bar{p}_t^* - 2Av_t^*q_t$

We assume agents form their subjective estimates of the mean and variance by considering past market clearing prices over a window of length L , that is

$$\bar{p}_t^* = \frac{1}{L} \sum_{i=1}^L p_{t-i}, \quad (5)$$

and

$$v_{t,L}^* = \frac{1}{L} \sum_{i=1}^L [\bar{p}_t^* - p_{t-i}]^2. \quad (6)$$

Let

$$x_{i,t} = p_{t-(i-1)}, \quad i = 1, 2, \dots, L.$$

Then, it follows from (5)–(6),

$$\begin{aligned} \bar{p}_{t+1}^* &= \frac{1}{L} [x_{1,t} + x_{2,t} + \dots + x_{L,t}], \\ v_{t+1}^* &= \frac{1}{L} \sum_{i=1}^L [\bar{p}_{t+1}^* - x_{i,t}]^2. \end{aligned} \quad (7)$$

Hence

$$x_{1,t+1} = p_{t+1} = \beta + \alpha \frac{\bar{p}_{t+1}^* - b}{a + 2Av_{t+1}^*}. \quad (8)$$

Because of the dependence of the subjective mean \bar{p}_t^* and variance v_t^* as price lagged L periods equation (8) is a difference equation of order L . It is more convenient to reduce it to a system of L first order difference equations. Let

$$\mathbf{x} = (x_1, \dots, x_L), \quad \mathbf{x}_t = (x_{1,t}, \dots, x_{L,t}),$$

and

$$f(\mathbf{x}) = \beta + \alpha \frac{(1/L)[x_1 + \dots + x_L] - b}{a + (2A/L) \sum_{i=1}^L ([x_1 + \dots + x_L]/L - x_i)^2}.$$

Then equation (8) can be written as the following difference system

$$\begin{cases} x_{1,t+1} = f(\mathbf{x}_t) \\ x_{2,t+1} = x_{1,t} \\ x_{3,t+1} = x_{2,t} \\ \vdots \\ x_{L,t+1} = x_{L-1,t} \end{cases} \quad (9)$$

2 Local Stability

One can see that the system (9) has a unique positive fixed point (p_o, p_o, \dots, p_o) satisfying

$$p_o = \beta + \alpha \frac{p_o - b}{a}$$

which implies $p_o = \frac{a\beta - \alpha b}{a - \alpha}$.

It can be verified that

$$\frac{\partial f}{\partial x_i}(p_o, \dots, p_o) = \frac{\alpha}{aL} \quad \text{for} \quad i = 1, \dots, L.$$

Therefore the Jacobian matrix of the system (9) at the steady state is given by the $L \times L$ matrix

$$J = \begin{pmatrix} \frac{\alpha}{aL} & \frac{\alpha}{aL} & \cdots & \frac{\alpha}{aL} & \frac{\alpha}{aL} \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{pmatrix}. \quad (10)$$

Denote $\gamma = -\frac{\alpha}{aL}$ and $D_i(L) = \det(\lambda I - J)_{i \times i}$ with $i = 1, 2, \dots, L$. We then have

$$D_1(\lambda) = \lambda + \gamma, \quad D_2(\lambda) = \lambda D_1 + \gamma = \lambda^2 + \gamma\lambda + \gamma,$$

and more generally,

$$D_i(\lambda) = \lambda D_{i-1} + \gamma, \quad i = 1, 2, \dots, L.$$

Thus

$$D_L(\lambda) = \lambda^L + \gamma\lambda^{L-1} + \cdots + \gamma\lambda + \gamma. \quad (11)$$

Using Jury's Test (see appendix A), we can derive the following local stability result. The proof of theorem 23.1 can be found in the appendix A.

Theorem 23.1 *The unique fixed point of the system (9) is locally stable if and only if $-L < \frac{\alpha}{a} < L$.*

It is interesting to notice that both the equilibrium (p_o, \dots, p_o) and the local stability condition of the system (9) are independent from the risk aversion A . As pointed out by Boussard [2], this is a peculiarity of the particular expectation hypothesis chosen here. Yet, under our assumption, $|\alpha/a|$ plays a key role on the local stability of the positive equilibrium of the system (9). Furthermore, the local stability condition $|\alpha/a| < L$ implies that the region of the parameter α/a on the local stability of the positive equilibrium of the system (9) is proportional to the lag length L . Theorem 23.1 tells us that larger time lags lead to larger region of stability (in terms of the parameter α/a).

3 Bifurcation Analysis

Let us consider the simplest case first, that is the case when $L = 2$. Then we have a system

$$\begin{cases} x_{1,t+1} = f(\mathbf{x}_t) \\ x_{2,t+1} = x_{1,t}, \end{cases} \quad (12)$$

where

$$f(\mathbf{x}_t) = \beta + \alpha \frac{(1/2)[x_1 + x_2] - b}{a + (A/2)[x_1 - x_2]^2}, \quad \mathbf{x} = (x_1, x_2) \in \mathbb{R}^2.$$

Let $y_i = x_i - p_o$ ($i = 1, 2$) so that the fixed point is at the origin. We then have

$$\begin{cases} y_{1,t+1} = g(\mathbf{y}_t) \\ y_{2,t+1} = y_{1,t}, \end{cases} \quad (13)$$

where

$$g(\mathbf{y}_t) = (\beta - p_o) + \alpha \frac{(1/2)[y_1 + y_2] + (p_o - b)}{a + (A/2)[y_1 - y_2]^2}, \quad \mathbf{y} \in \mathbb{R}^2.$$

The Jacobian of the system (13) at the origin is then given by

$$J_2 = \begin{pmatrix} \frac{\alpha}{2a} & \frac{\alpha}{2a} \\ 1 & 0 \end{pmatrix}.$$

From $L = 2$, we have $\gamma = -\frac{\alpha}{2a}$. It follows from Theorem 23.1 that the equilibrium (p_o, p_o) of the system (12) will lose its stability when γ passes through 1. Also, when γ is near 1, the Jacobian matrix J_2 has a pair of complex eigenvalues, say λ and $\bar{\lambda}$ with

$$\lambda = \lambda(\gamma) = \frac{1}{2}[-\gamma + i\sqrt{\gamma(4-\gamma)}] = \sqrt{\gamma}e^{i\theta},$$

where θ satisfies

$$\sin \theta = -\frac{\sqrt{\gamma}}{2}, \quad \cos \theta = \frac{\sqrt{4-\gamma}}{2}.$$

Let λ_o be the value of λ when $\gamma = 1$, that is,

$$\lambda_o = \frac{1}{2}[-1 + i\sqrt{3}].$$

Obviously, $\lambda_o^3=1$. For a map in \mathbb{R}^2 , according to Kuznetsov [6] (page 350), there is no "strong resonances" if there is an eigenvalue, say $\tilde{\lambda}$, satisfying $\tilde{\lambda}^q \neq 1$ for $q = 1, 2, 3, 4$. Otherwise, we say the map has a $1 : q$ **resonance** ($q = 1, 2, 3, 4$). Hence our map has a $1 : 3$ resonance. As pointed out by Hale and Kocak [4] (page 481), *the dynamics of such maps — strong resonances — can be exceedingly complicated and the answer is not yet completely known*. The complexity of such maps is illustrated in one of the Example 15.34 in Hale and Kocak [4] (pages 481–482). For more detailed discussion on the bifurcations of fixed points in discrete-time maps on \mathbb{R}^2 with both weak and strong resonances, we refer the reader to Iooss [5] when the maps involve one parameter and to Kuznetsov [6] when the maps involve two parameters.

The rest of this section is devoted to the study of generic bifurcations of the fixed points (p_o, p_o) of the map defined by (12). To keep the discussion simple, we will treat γ as the only parameter of the map.

To perform a standard normal form calculation (see Arrowsmith et. al. [1]) for the system (13), we write the function g in the following form

$$g(y_1, y_2) = \sum_{j,k;j+k \geq 1} g_{jk} y_1^j y_2^k,$$

with

$$\begin{cases} g_{10} = g_{01} = \frac{\alpha}{2a} \\ g_{11} = \frac{\alpha(p_o-b)A}{a^2} \\ g_{20} = g_{02} = -\frac{1}{2} \frac{\alpha(p_o-b)A}{a^2} \\ g_{30} = g_{03} = -\frac{1}{4} \frac{\alpha A}{a^2} \\ g_{21} = g_{12} = \frac{1}{4} \frac{\alpha A}{a^2}. \end{cases}$$

We now introduce complex coordinates

$$z = -iy_1 + iy_2\bar{\lambda}.$$

From which, we have

$$y_1 = \frac{(-i)[\lambda z + \bar{\lambda}\bar{z}]}{\sqrt{\gamma(4-\gamma)}}, \quad y_2 = \frac{z + \bar{z}}{\sqrt{\gamma(4-\gamma)}}.$$

Then the linear part of the map becomes

$$z_{t+1} = -i[-\gamma y_{1,t} - \gamma y_{2,t}] + i\bar{\lambda}y_{1,t} = -i[(\lambda + \bar{\lambda})y_{1,t} - \lambda\bar{\lambda}y_{2,t}] + i\bar{\lambda}y_{1,t} = \lambda z_t.$$

Hence the mapping can be written in a complex form

$$z_{t+1} = \lambda z_t + \sum_{j+k \geq 2} \frac{1}{j!k!} a_{jk} z_t^j \bar{z}_t^k,$$

where

$$\begin{cases} a_{10} &= g_{10} \frac{-i\lambda}{\sqrt{\gamma(4-\gamma)}} \\ a_{01} &= g_{01} \frac{-i\bar{\lambda}}{sqr\gamma(4-\gamma)} \\ a_{11} &= \frac{1}{\gamma(4-\gamma)} [-i(\lambda + \bar{\lambda})g_{11} - 2\lambda\bar{\lambda}g_{20} + 2g_{02}] \\ a_{20} &= \frac{2}{\gamma(4-\gamma)} [-i\lambda g_{11} - \lambda^2 g_{20} + g_{02}] \\ a_{02} &= \frac{2}{\gamma(4-\gamma)} [-i\bar{\lambda}g_{11} - \bar{\lambda}^2 g_{20} + g_{02}]. \end{cases}$$

The following Lemma on the normal form of the map (13) (with 1 : 3 resonance) can be found in Kuznetsov [6] (p.382).

Lemma 23.1 (Normal form map for 1:3 resonance) *The map (13) can be transformed by an invertible smooth and smoothly parameter-dependent change of variable, for all $\gamma = 1 + \epsilon$ with sufficiently small $|\epsilon|$, into the form*

$$\xi \mapsto \Gamma_\gamma(\xi) = \lambda(\gamma)\xi + B(\gamma)\bar{\xi}^2 + C(\gamma)\xi|\xi|^2 + O(|\xi|^4), \tag{14}$$

where

$$B(\gamma) = \frac{a_{02}(\gamma)}{2},$$

and

$$C(\gamma) = \frac{a_{20}(\gamma)a_{11}(\gamma)[2\lambda(\gamma) + \bar{\lambda}(\gamma) - 3]}{2[\bar{\lambda}(\gamma) - 1][\lambda^2(\gamma) - \lambda(\gamma)]} + \frac{|a_{11}(\gamma)|^2}{1 - \bar{\lambda}(\gamma)} + \frac{a_{21}(\gamma)}{2}.$$

It can be verified that $B(\gamma) = \frac{1}{2}a_{02}(\gamma) = (p_o - b)\frac{A}{a}\frac{1}{4-\gamma}(i\bar{\lambda} + 1)^2$. Hence, $B(1) = \frac{1}{18}\frac{A}{a}[\beta - b][2 + \sqrt{3} - i]^2$. Using a result from Iooss [5] (p.110, Theorem 1), we have the following bifurcation result.

Theorem 23.2 *For the system (12), if $\text{Re}(B(1)) \neq 0$ (which is true in our case), then there exists a single one-parameter family of fixed points of order 3 bifurcating from the positive equilibrium (p_o, p_o) . The positive equilibrium, which is stable for $\gamma < 1$, becomes unstable for $\gamma > 1$, where the family of fixed points of order 3 which bifurcates on both sides of criticality is unstable on both sides.*

Theorem 23.2 indicates the dynamic structure near the positive equilibrium $P_o(p_o, p_o)$ and the hyperbolic periodic points bifurcating from P_o near the critical value $\gamma = 1$. In the following, two most common numerical simulation techniques, phase diagrams and bifurcation diagrams, are used in the study of the global dynamics of the nonlinear model.

Numerical simulations

The case $L = 2$. In this part, we consider the case of $L = 2$, that is the system (12). The parameters a, b, A and β are selected to be fixed and α are varied to characterize the changing of γ . In the following discussion, we choose $a = 1, b = 0, A = 0.005$ and $\beta = 11$.




Figure 1: Pseudo-phase plot of (12) with $\alpha = -1.95$




Figure 2: Basin plot of (12) with $\alpha = -1.95$




Figure 3: Pseudo-phase plot of (12) with $\alpha = -1.95$ and the structure near the saddle point set S

Firstly, let $\alpha = -1.95$ so that $0 < \gamma < 1$. In this case, apart from the fixed positive equilibrium $P_o(p_o, p_o)$ with $p_o = 3.8596$, the system (12) has two sets of period three fixed points, denoted $S = \{S_1, S_2, S_3\}$ and $P = \{P_1, P_2, P_3\}$, where $S_1 = (4.2295, 2.7525)$, $S_2 = (4.2295, 4.2295)$, $S_3 = (2.7525, 4.2295)$ and $P_1 = (10.4993, -9.4737)$, $P_2 = (10.4993, 10.4993)$, $P_3 = (-9.4737, 10.4993)$. It follows from Theorem 23.2 that, when $\alpha = -1.95$ (so that $\gamma < 1$), P_o is locally stable and the period three point set S corresponds to the order 3 bifurcating from the positive equilibrium P_o .

Fig.1 shows the phase plot of $(x_{1,t}, x_{2,t}) = (p_t, p_{t-1})$, which is often called the pseudo-phase plot of the system. We select four initial values: $I_1 = (2.68, 4.42)$, $I_2 = (2.88, 3.8)$, $I_3 = (2.89, 3.8)$ and $I_4 = (4.6, 4.0)$. Numerical simulations in Fig.1 indicate that, solutions with the initial points I_3 and I_4 converge to the fixed point P_o , while the solutions with I_1 and I_2 converge to the period three point set P . One can choose other initial values to do the simulations, but it turns out that all the solutions with different initial values will converge to either P_o or P , as indicated in Fig.1. A more detailed numerical simulation on the basins of the attractors P_o and P are plotted in Fig.2, in which, all the solutions with initial values from the shaded area converge to P_o and the rest of the solutions converge to P .

In Fig.3 we enlarge the central part of Fig.1, we can then see clearly the structure of the bifurcating point set S . We select four initial values $(4.6, 4.0)$, $(4.7, 4.0)$, $(6.0, 4.0)$ and $(2.0, 2.8)$. As suggested by Iooss [5] (pp. 127–128), S is a set of saddle points.

Theorem 23.2 asserts the bifurcating behavior when γ is near the critical value 1. Now the question is whether the single one-parameter family of fixed points S of order 3 exists when γ moves away from




Figure 4: Solutions p_t of (12) with $\alpha = -1.82206$




Figure 5: Pseudo-phase plot of (12) with $\alpha = -2.0$

1. In fact one can check that, when α increases from -1.95 to -1.82207 , apart from the fixed equilibrium P_o , the two set of one-parameter (α) family of fixed points S and P of order 3 continue to exist and the distance between P and S , which is defined by $\max\{|P_i - S_i|, i = 1, 2, 3\}$, decreases. When $\alpha = -1.82206$, the system has only the positive fixed equilibrium P_o and the solutions with initial values $(6, 8.5)$, $(6, 8.6)$, $(6, 8.7)$ and $(6, 10)$ are plotted in Fig.4. This implies that there exists $\alpha_b \in (-1.82207, -1.82206)$, or equivalently there exists a $\gamma_b < 1$ such that, for $\gamma_b < \gamma < 1$, the structure of the solutions is given by Fig.1 and Fig.3; while for $\gamma < \gamma_b$ (and near γ_b), the structure is indicated by Fig.4. Noting that the solutions remain near an order 3 periodic solution before they converge to P_o with $p_o = 3.897861846$. An interesting finding is that, when we fixed the first initial value, say $x_o = 6$, and increase the second initial values, say y_o , from 10 up to near 30, the numerical steps needed for the convergence increases, after 30, the numbers of steps decreases.

As α decreases from α_b to -2 (but greater than -2), that is γ increases from γ_b to 1, the distance between two sets of one-parameter (α) families of fixed points S and P of order 3 increases and, correspondingly, the distance between S and P_o decreases (to zero). When $\alpha = -2.0$, that is $\gamma = 1$, the system has a fixed equilibrium P_o with $p_o = 11/3$ and an order 3 periodic set P . The phase structure in this case is indicated in Fig.5, in which four initial values $(2.2, 6)$, $(2.62, 6)$, $(2.3, 6)$ and $(3.67, 3.67)$ are selected. One can see that P is attracting and P_o is unstable and it has also the properties of the saddle point S with both stable and unstable manifolds.

Now we choose $\alpha = -2.2$ (so that $\gamma > 1$), then the system (12) has a fixed equilibrium P_o with $p_o = 3.4375$ and two sets of order 3 bifurcating points P and S . In Fig.6, we have the phase plot of the solutions with initial values $(3.43, 3.43)$, $(1.85, 1.5)$, $(1.8, 1.5)$, $(2.1, 2.25)$ and $(2.15, 2.25)$. It shows that P is the only attractor. 7shows the convergence of the order 3 periodic orbit P , in which the initial value $(1, 4)$ is selected.

General case. In general, near the critical value $\gamma = 1$, the system (9) has a periodic L orbit (fixed points of order L) bifurcating from the positive fixed equilibrium. The bifurcating periodic L orbit may have a similar behaviour as the set S as in the case of $L = 2$. 8and 9show the convergence of the unique fixed equilibrium of the system with $L = 10$ and $\alpha = -8.0(\gamma < 1)$, where initial value $(1.2, 1.3, 1.2, 1.3, 1.2, 1.3, 1.2, 1.3, 1.2, 1.3)$ is selected. When $L = 10$ and $\alpha = -12, \gamma > 1$, 10 shows the bifurcation of the positive equilibrium and the attractivity of a family of periodic 10 orbits. Numerical simulations show that the system (9) with $L > 2$ has similar dynamics to the one with $L = 2$.

Under the assumptions $\bar{p}_t^* = \hat{p} = constant$ and $\bar{p}_t^* = p_{t-2}$, respectively, Boussard [2] shows that these assumptions may result in the market generating chaotic price and quantity series. He suggested that it would be more rational to treat both prices and quantities as symmetrical and this is indeed




Figure 6: Pseudo-phase plot of (12) with $\alpha = -2.2$




Figure 7: Solution P_t of (3.1) with $\alpha = -2.2$




Figure 8: Pseudo-phase plot with $\alpha = -8$




Figure 9: Solution P_t with $\alpha = -8$

the basic assumption in this paper. Corresponding to our case when $L = 2$, he claimed that the main conclusions remain approximately the same. However, our results suggest that, under these more general symmetrical assumptions, the market generates simpler dynamic behaviour. In order to generate more complicated dynamics and chaotic motion, we need to replace p and q in equations (1) by their logarithms, which is also a natural solution to avoid negative prices and quantities that can arise under the linear supply and demand curves. This will be treated in the next section.

4 Constant Elasticity Supply and Demand Curves

The problem of making use of linear supply and demand curves is the occurrence of negative values for prices and quantities. One solution to this problem is to replace p and q by their logarithms. That is, we replace the demand equation in (1) by

$$p_t = \beta q_t^\alpha. \tag{15}$$

and the supply equation by

$$p_t^* = b q_t^a \tag{16}$$

where a, b and β are positive and α are negative constants.




Figure 10: Solution P_t with $\alpha = -12$

One can rescale the equations by letting ² $Q_t = q_t^a$. We then have

$$p_t = \beta Q_t^{\alpha/a} \tag{17}$$

and

$$p_t^* = bQ_t. \tag{18}$$

Under constant absolute risk aversion A , the certainty equivalent of the revenue $r = pQ$ is $\bar{p}_t^*Q_t - Av_t^*Q_t^2$. Thus the marginal revenue certainty equivalent is

$$\tilde{p}_t = \bar{p}_t^* - 2Av_t^*Q_t. \tag{19}$$

Suppose a “linear” (in terms of Q , not q) marginal cost so that the supply equation is

$$\tilde{p}_t = bQ_t. \tag{20}$$

These results lead to the supply equation

$$\bar{p}_t^* - 2Av_t^*Q_t = bQ_t; \tag{21}$$

that is

$$[b + 2Av_t^*]Q_t = \bar{p}_t^*. \tag{22}$$

Assume \bar{p}_t^* and v_t^* are formed as (5) and (6) in section 1, then from (22) and (15) the equality of supply and demand implies the market clearing quantity

$$\left[b + \frac{2A\beta^2}{L} \sum_{i=1}^L \left(\frac{1}{L} \sum_{k=1}^L Q_{t-k}^{\alpha/a} - Q_{t-i}^{\alpha/a} \right)^2 \right] Q_t = \frac{\beta}{L} \sum_{i=1}^L Q_{t-i}^{\alpha/a}. \tag{23}$$

Using equation(17), we can rewrite the equation (23) in terms of the price

$$p_t = \beta \left\{ \frac{\frac{1}{L} \sum_{i=1}^L p_{t-i}}{b + \frac{2A}{L} \sum_{i=1}^L \left(\frac{1}{L} \sum_{k=1}^L p_{t-k} - p_{t-i} \right)^2} \right\}^{\alpha/a}. \tag{24}$$

Let

$$u_{i,t} = p_{t-(i-1)}, \quad i = 1, 2, \dots, L. \tag{25}$$

²This rescaling is equivalent to a change in numeraire




Figure 11: Pseudo-phase plot with $\alpha = -1.9$




Figure 12: Pseudo-phase plot with $\alpha = -2.0$




Figure 13: Pseudo-phase plot with $\alpha = -2.1$




Figure 14: Pseudo-phase plot with $\alpha = -2.2$

Then the equation (24) can be written as the following L dimensional system of first order difference equations

$$\begin{cases} u_{1,t+1} &= \beta \left[\frac{(1/L) \sum_{i=1}^L u_{i,t}}{b + (2A/L) \sum_{i=1}^L ((1/L) \sum_{k=1}^L u_{k,t} - u_{i,t})^2} \right]^{\alpha/a} \\ u_{2,t+1} &= u_{1,t} \\ &\vdots \\ u_{L,t+1} &= u_{L-1,t}. \end{cases} \tag{26}$$

The system (26) has a unique positive equilibrium $u_1 = u_2 = \dots = u_L = b(\beta/b)^{a/(a-\alpha)}$. One can verify that, at the equilibrium point, the system (26) has the Jacobian matrix J as defined in section 2. Therefore, Theorem 23.1 holds for system (26) too.

Fig.11 is the phase plot of system (26) when $L = 2$ and $\alpha = -1.9$ (and hence $\gamma < 1$). We select three initial values $I_1(2.36, 2.4)$, $I_2(2.4, 2.4)$ and $I_3(1, 2.4)$. The solution with I_1 converges to the fixed equilibrium P_o and the solutions with I_2 and I_3 seem to converge to a bounded attractor, rather than P_o . Fig.12 shows the case when $\alpha = -2.0$ and the fixed equilibrium P_o is unstable. The corresponding attractor seems more complicated.

Fig.13 and Fig.14 show the case when $\gamma > 1$. It seems that the system has a strange attractor when $\gamma > 1$. It may have different shape for different α .

The above numerical simulations suggest that, for the system (26), the market generates more complicated dynamics. In particular, when $\gamma > 1$, the model may have chaotic behaviour. The




Figure 15: Bifurcation diagram

bifurcation plot of Q as a function of α is shown in Fig.15, which indicates the complicated dynamics of the system. Those simulations imply that the general behaviour of models built along this line is very different from what we have seen in the previous sections (certainly quite a different picture to the one suggested by Boussard [2]).

A Proof of Theorem 23.1

To give the proof of Theorem 23.1, we need introduce concepts of the *inners* of a matrix and the *positive innerwise* matrix, which can be found from the book by Elaydi [3] (pages 180–181).

Let $B = (b_{ij})_{n \times n}$ be a matrix. The *inners* of the matrix B are the matrix itself and all the matrices obtained by omitting successively the first and last rows and the first and last columns. A matrix B is said to be *positive innerwise* if the determinants of all its inners are positive.

We now consider the k th order scalar equation

$$x_{n+k} + p_1 x_{n+k-1} + p_2 x_{n+k-2} + \cdots + p_k x_n = 0, \quad (27)$$

where the p_i 's are real numbers. Obviously, the characteristic equation of the equation (27) is given by

$$p(\lambda) = \lambda^k + p_1 \lambda^{k-1} + \cdots + p_k. \quad (28)$$

The Schue-Cohn criterion defines the conditions for the characteristic roots of equation (28) to fall inside the unit circle. More precisely, the following Jury's test will be used in our proof to Theorem 23.1.

Theorem 23.3 *The zeros of the characteristic polynomial (28) lie inside the unit circle if and only if the following hold:*

- $p(1) > 0$
- $(-1)^k p(-1) > 0$,
- the $(k-1) \times (k-1)$ matrices

$$B_{k-1}^{\pm} = \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 \\ p_1 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & 0 \\ p_{k-3} & p_{k-4} & \cdots & 1 & 0 \\ p_{k-2} & p_{k-3} & \cdots & p_1 & 1 \end{pmatrix} \pm \begin{pmatrix} 0 & 0 & \cdots & 0 & p_k \\ 0 & 0 & \cdots & p_k & p_{k-1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & p_k & \cdots & p_4 & p_3 \\ p_k & p_{k-1} & \cdots & p_3 & p_2 \end{pmatrix}$$

are positive innerwise.

Now let us prove Theorem 23.1. What we need to show is that all the zeros of the characteristic polynomial $D_L(\lambda)$ defined by (11) lie inside of the unit circle if and only if $\gamma < 1$, that is, $D_L(\lambda)$ satisfies the three conditions in Theorem 23.3 if and only if $\gamma < 1$.

From $\gamma > 0$, it is easy to see that $D_L(1) = 1 + (L - 1)\gamma > 0$ and $(-1)^L D_L(-1) = 1 - \gamma$ if L is odd and $(-1)^L D_L(-1) = 1$ if L is even. Hence the first two conditions of Theorem 23.3 hold if and only if $\gamma < 1$. To show the third condition is satisfied, it is enough to show that, for $k = 1, 2, \dots, L - 1$, the matrix B_k^\pm with $p_1 = p_2 = \dots = p_L = \gamma$ are positive if and only if $\gamma < 1$.

Let $k = 2m$ be even. Then we have

$$B_k^+ = \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 & 0 & 0 & \cdots & 0 & \gamma \\ \gamma & 1 & \cdots & 0 & 0 & 0 & 0 & \cdots & \gamma & \gamma \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \gamma & \gamma & \cdots & 1 & 0 & 0 & \gamma & \cdots & \gamma & \gamma \\ \gamma & \gamma & \cdots & \gamma & 1 & \gamma & \gamma & \cdots & \gamma & \gamma \\ \gamma & \gamma & \cdots & \gamma & 2\gamma & 1 + \gamma & \gamma & \cdots & \gamma & \gamma \\ \gamma & \gamma & \cdots & 2\gamma & 2\gamma & 2\gamma & 1 + \gamma & \cdots & \gamma & \gamma \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \gamma & 2\gamma & \cdots & 2\gamma & 2\gamma & 2\gamma & 2\gamma & \cdots & 1 + \gamma & \gamma \\ 2\gamma & 2\gamma & \cdots & 2\gamma & 2\gamma & 2\gamma & 2\gamma & \cdots & 2\gamma & 1 + \gamma \end{pmatrix} \tag{29}$$

To evaluate the determinate of B_k^+ , we use (-1) to multiply the i -th columns and add to the $2m - (i - 1)$ -th columns, respectively, for $i = 1, \dots, m$. We then have

$$\begin{aligned} |B_k^+| &= \begin{vmatrix} 1 & 0 & \cdots & 0 & 0 & 0 & 0 & \cdots & 0 & \gamma - 1 \\ \gamma & 1 & \cdots & 0 & 0 & 0 & 0 & \cdots & \gamma - 1 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \gamma & \gamma & \cdots & 1 & 0 & 0 & \gamma - 1 & \cdots & 0 & 0 \\ \gamma & \gamma & \cdots & \gamma & 1 & \gamma - 1 & 0 & \cdots & 0 & 0 \\ \gamma & \gamma & \cdots & \gamma & 2\gamma & 1 - \gamma & 0 & \cdots & 0 & 0 \\ \gamma & \gamma & \cdots & 2\gamma & 2\gamma & 0 & 1 - \gamma & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \gamma & 2\gamma & \cdots & 2\gamma & 2\gamma & 0 & 0 & \cdots & 1 - \gamma & 0 \\ 2\gamma & 2\gamma & \cdots & 2\gamma & 2\gamma & 0 & 0 & \cdots & 0 & 1 - \gamma \end{vmatrix} \\ &= (1 - \gamma)^m \begin{vmatrix} 1 & 0 & \cdots & 0 & 0 & 0 & 0 & \cdots & 0 & -1 \\ \gamma & 1 & \cdots & 0 & 0 & 0 & 0 & \cdots & -1 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \gamma & \gamma & \cdots & 1 & 0 & 0 & -1 & \cdots & 0 & 0 \\ \gamma & \gamma & \cdots & \gamma & 1 & -1 & 0 & \cdots & 0 & 0 \\ \gamma & \gamma & \cdots & \gamma & 2\gamma & 1 & 0 & \cdots & 0 & 0 \\ \gamma & \gamma & \cdots & 2\gamma & 2\gamma & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \gamma & 2\gamma & \cdots & 2\gamma & 2\gamma & 0 & 0 & \cdots & 1 & 0 \\ 2\gamma & 2\gamma & \cdots & 2\gamma & 2\gamma & 0 & 0 & \cdots & 0 & 1 \end{vmatrix} \end{aligned} \tag{30}$$

Now for $i = 1, 2, \dots, m$, we first add the $2m - (i - 1)$ -th columns to the i -th columns, respectively. Then, multiply γ to the $2m - (i - 1)$ -th column and add to the all the first $m - 1$ columns. as a result, the upper left block matrix become a zero matrix and the down left block matrix has 2γ as non-diagonal elements and $2\gamma + 1$ as diagonal elements. Correspondingly,

$$|B_k^+| = (-1)^m (1 - \gamma)^m \begin{vmatrix} 2\gamma & \cdots & 2\gamma + 1 \\ \vdots & \ddots & \vdots \\ 2\gamma + 1 & \cdots & 2\gamma \end{vmatrix} \tag{31}$$

We use -1 to time the first column and add to all the rest columns. Then, use -1 to multiply the columns 2 to k and add them to the first column. As a result, we have a low triangle matrix with $(1, 1, \dots, 1, 2m\gamma + 1)$. Therefore,

$$\det(B_k^+) = (1 - \gamma)^m (L\gamma + 1). \quad (32)$$

Similarly,

$$B_k^- = \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 & 0 & 0 & \cdots & 0 & -\gamma \\ \gamma & 1 & \cdots & 0 & 0 & 0 & 0 & \cdots & -\gamma & -\gamma \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \gamma & \gamma & \cdots & 1 & 0 & 0 & -\gamma & \cdots & -\gamma & -\gamma \\ \gamma & \gamma & \cdots & \gamma & 1 & -\gamma & -\gamma & \cdots & -\gamma & -\gamma \\ \gamma & \gamma & \cdots & \gamma & 0 & 1 - \gamma & -\gamma & \cdots & -\gamma & -\gamma \\ \gamma & \gamma & \cdots & 0 & 0 & 0 & 1 - \gamma & \cdots & -\gamma & -\gamma \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \gamma & 0 & \cdots & 0 & 0 & 0 & 0 & \cdots & 1 - \gamma & -\gamma \\ 0 & 0 & \cdots & 0 & 0 & 0 & 0 & \cdots & 0 & 1 - \gamma \end{pmatrix} \quad (33)$$

To find the $\det(B_k^-)$, we expand it first by the last row and then by the first row and these lead to $\det(B_k^-) = (1 - \gamma) \det(B_{k-2}^-)$. Since $k = 2m$, it follows from the formula $\det(B_{2m}^-) = (1 - \gamma) \det(B_{2(m-1)}^-)$ that

$$\det(B_k^-) = (1 - \gamma)^m. \quad (34)$$

In conclusion, we have for $k = 2m$,

$$\det(B_k^+) = (1 - \gamma)^m (n\gamma + 1), \quad \det(B_k^-) = (1 - \gamma)^m. \quad (35)$$

Next we assume that $k = 2m + 1$. Then

$$B_k^- = \begin{pmatrix} 1 & \cdots & 0 & 0 & 0 & \cdots & -\gamma \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \gamma & \cdots & 1 & 0 & -\gamma & \cdots & -\gamma \\ \gamma & \cdots & \gamma & 1 - \gamma & -\gamma & \cdots & -\gamma \\ \gamma & \cdots & 0 & 0 & 1 - \gamma & \cdots & -\gamma \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & 0 & 0 & \cdots & 1 - \gamma \end{pmatrix} \quad (36)$$

It is easy to see that $\det(B_k^-) = (1 - \gamma) \det(B_{2m}^-)$. Using (35), we have

$$\det(B_k^-) = (1 - \gamma)^{m+1}. \quad (37)$$

On the other hand,

$$B_k^+ = \begin{pmatrix} 1 & \cdots & 0 & 0 & 0 & \cdots & \gamma \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \gamma & \cdots & 1 & 0 & \gamma & \cdots & \gamma \\ \gamma & \cdots & \gamma & 1 + \gamma & \gamma & \cdots & \gamma \\ \gamma & \cdots & 2\gamma & 2\gamma & 1 + \gamma & \cdots & \gamma \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 2\gamma & \cdots & 2\gamma & 2\gamma & 2\gamma & \cdots & 1 + \gamma \end{pmatrix} \quad (38)$$

To find the $\det(B_k^-)$, we multiply the i -th column by -1 and add to the $2m - (i - 1)$ -th column, respectively, for $i = 1, \dots, m$.

$$\det(B_k^+) = (1 - \gamma)^m \begin{vmatrix} 1 & \cdots & 0 & 0 & 0 & \cdots & -1 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \gamma & \cdots & 1 & 0 & -1 & \cdots & 0 \\ \gamma & \cdots & \gamma & 1 + \gamma & 0 & \cdots & 0 \\ \gamma & \cdots & 2\gamma & 2\gamma & 1 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 2\gamma & \cdots & 2\gamma & 2\gamma & 0 & \cdots & 1 \end{vmatrix}. \tag{39}$$

Similarly, one can use row operations to reduce the upper left $m \times m$ matrix to a zero matrix and correspondingly,

$$\det(B_k^+) = (-1)^{m+1}(1 - \gamma)^m \begin{vmatrix} \gamma & \cdots & \gamma & 1 + \gamma \\ 2\gamma & \cdots & 2\gamma + 1 & 2\gamma \\ \vdots & \ddots & \vdots & \vdots \\ 2\gamma + 1 & \cdots & 2\gamma & 2\gamma \end{vmatrix}. \tag{40}$$

Multiply the first column by -1 and add all the rest of the columns of $\det(B_k^-)$ and then, multiply the last column by $-\gamma$ and add to the first column, multiply -2γ to the columns $2, 3, \dots, m$ and add to the first column. We then add up with

$$\det(B_k^+) = (-1)^{m+1}(1 - \gamma)^m \begin{vmatrix} 0 & 0 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ \gamma(2m + 1) + 1 & -1 & \cdots & -1 \end{vmatrix}. \tag{41}$$

Therefore

$$\det(B_k^+) = (\gamma k + 1)(1 - \gamma)^m. \tag{42}$$

Then from (37) and (42), for $k = 2m + 1$,

$$\det(B_k^+) = (\gamma k + 1)(1 - \gamma)^m, \quad \det(B_k^-) = (1 - \gamma)^{m+1}. \tag{43}$$

Finally, it follows from (35) and (43) that B_k^\pm are positive if and only if $\gamma < 1$ and this completes the proof.

References

- [1] D. Arrowsmith, J. Cartwright, A. Lansbury, and C. Place, *The Bogdanov map: Bifurcations, model locking, and chaos in a dissipative system*, Int. J. Bifurcation and Chaos **3** (1993), 803–842.
- [2] J.-M. Boussard, *When risk generates chaos*, J. Econ. Behav. Organ. **29** (1996), 433–446.
- [3] S.N. Elaydi, *An introduction to difference equations*, Springer, New York, 1996.
- [4] J. Hale and H. Kocak, *Dynamics and bifurcations*, Texts in Applied Mathematics, vol. 3, Springer-Verlag, New York, 1991.
- [5] G. Iooss, *Bifurcations of maps and applications*, Mathematics studies, vol. 36, North-Holland, Amsterdam, 1979.
- [6] Y.A. Kuznetsov, *Elements and applied bifurcation theory*, Applied mathematical sciences, vol. 112, SV, New York, 1995.

Troubles in Wonderland

Ric D. Herbert and Gareth D. Leeves
 Modelling and Simulation Group
 University of Western Sydney
 Hawkesbury
 Richmond, NSW, 2753
 Australia
 r.herbert/g.leeves@uws.edu.au

Abstract

This paper examines interactions between the economy and the environment using the Wonderland model. Specifically we examine the issue of sustainable and non-sustainable development. We address this by adapting the original model to examine a wider range of possible scenarios. The paper also extends the model to include environmental taxes and exogenous shocks. We present the model using visual simulation tools.

1 Introduction

In this paper we look at the Wonderland model by [5] which aims to portray interactions between the economic, demographic and environmental systems. The model is a low-order model with particular emphasis on economic growth and the stock of natural capital. Our aim is to extend the model by considering how development paths of the economy and the environment are affected by environmental taxes. Importantly we also expand the model to include stochastic disturbances.

The model as developed by [5] is a discrete time model. The fast-slow dynamics of the model have been investigated in a continuous time formulation by [6] and by [2].

In further research, [7] use the sophisticated visualisation software AVS (by Advanced Visual Systems) to examine the dynamical behaviour of Wonderland. In this paper, we simulate and visualise these dynamics using the more accessible software MATLAB (by Mathworks).

2 The Model

2.1 Basic Model

Wonderland is a model of possible interactions between the economy, demographic change and the environment. The model enables one to observe whether particular development paths are sustainable or not. A development path is considered to be unsustainable, following the World Commission on Environment and Development [1], if higher standards of living in the current generation are obtained at the expense of future generations or lower current generation death rates means higher death rates in future generations. The equations of the Wonderland model can be grouped into four sections; economy, population, environment and environmental policy. We shall deal with each of these sections in turn. All variables are described in Table 1 and the parameters are set out in Table 2.

Economy

$$Y_{t+1} = Y_t [1 + \gamma - (\gamma + \eta)(1 - K_t)^\lambda] \quad (1)$$

$$I_t = Y_t - C_t \quad (2)$$

Equation 1 defines the economy's per capita output, Y , to grow exponentially, depending on its own previous level and the stock of natural capital, K . The value of natural capital can vary between

Variable	Description
Y	per capita output
I	net per capita output
B	the crude birth rate
D	the crude death rate
N	the population
F	the flow of pollutants
K	the natural capital stock
C	the pollution control expenditure
P	the quantity of pollution

Table 1: Model Variables

the values 0 and 1. If the stock value is 1 then all natural resources are undiminished by pollution. If the environment is totally polluted then it takes the value 0. The lower the stock of natural capital the lower the rate of per capita growth. The second equation in the economic section, Equation 2, states that net per capita output, I , is the difference between per capita output and per capita expenditures on pollution control, C .

Population

$$B_t = \beta_0 \left[\beta_1 - \left(\frac{e^{\beta I_t}}{1 + e^{\beta I_t}} \right) \right] \tag{3}$$

$$D_t = \alpha_0 \left[\alpha_1 - \left(\frac{e^{\alpha I_t}}{1 + e^{\alpha I_t}} \right) \right] [1 + \alpha_2(1 - K_t)^\theta] \tag{4}$$

$$N_{t+1} = N_t \left[1 + \left(\frac{B_t - D_t}{1000} \right) \right] \tag{5}$$

Population growth is represented by Equations 3 to 5. Growth in population, N , is measured as the difference between the crude birth rate, B , and death rate, D . Increases in net per capita output lead to decreases in the birth and death rates. The death rate is also influenced by the stock of natural capital, whereby decreases in the stock cause the death rate to rise.

One of the characteristics of this model of the population is that if the stock of natural capital is complete ($K = 1$), then for a constant level of economic output, the population will grow exponentially. Further, in this situation, in the extreme case when the economy completely collapses ($Y = 0$), then the population will still grow given the parameters in Table 2.

Environment

$$F_t = N_t Y_t P_t - \kappa \left(\frac{e^{\epsilon C_t N_t}}{1 + e^{\epsilon C_t N_t}} \right) \tag{6}$$

$$K_{t+1} = \frac{e^{\ln\left(\frac{K_t}{1-K_t}\right) + \delta K_t^\rho - \omega F_t}}{1 + e^{\ln\left(\frac{K_t}{1-K_t}\right) + \delta K_t^\rho - \omega F_t}} \tag{7}$$

$$C_t = \phi(1 - K_t)^\mu Y_t \tag{8}$$

$$P_{t+1} = \chi P_t \tag{9}$$

The environment is modelled by four equations. The first, Equation 6, describes the annual flow of pollutants, F . These are determined by the population, per capita output, per unit pollution and the

Parameter	Value	Section
γ	0.04	Economy
η	0.04	Economy
λ	2	Economy
α	0.09	Population
α_0	10	Population
α_1	2.5	Population
α_2	2	Population
β	0.08	Population
β_0	40	Population
β_1	1.375	Population
θ	15	Population
κ	1	Environment
δ	1	Environment
ϵ	0.02	Environment
ρ	0.2	Environment
ω	0.1	Environment
ϕ	0.5	Environment
μ	2	Environment
χ	0.96	Environment

Table 2: Model Parameters (for Dream Scenario)

amount spent on pollution control measures. The flow also depends on the effectiveness of pollution control measures denoted by the parameter κ .

Equation 7 specifies the interaction between the flow of pollution and the stock of natural capital. Obviously, natural capital is adversely affected by a higher pollution flow. However, the equation allows for natural capital to regenerate itself and offset the pollution flow of earlier periods.

The technologies of pollution production and control, P , are modelled by Equation 9. It is assumed that the technologies of pollution reduction is constantly improving over time at some prescribed rate χ , which reduces the pollution per unit. The equation shows exponential decay, when $\chi < 1$, in pollution per unit of production.

Pollution control expenditure, C , is related to two variables. Firstly, deterioration in the state of the environment, reflected in the stock of natural capital and not the current level of pollution flow, will cause increased expenditure. Secondly, an increase the level of per capita output leads to greater expenditure on control measures. The parameters χ and κ are important factors in determining the development path of the environment.

2.2 Model with Taxation

Suppose environmental taxes are introduced in Wonderland with the aim of reducing pollution. It is assumed that these taxes have the effect of reducing pollution per unit output, so that Equation 9, becomes

$$P_{t+1} = (1 - \tau)\chi P_t \quad (10)$$

where τ is the rate of environmental taxes.

The environmental taxes would also be expected to have the effect of reducing economic output, so that Equation 1 becomes

$$Y_{t+1} = Y_t \left[1 + \gamma - (\gamma + \eta)(1 - K_t)^\lambda - \frac{\gamma_0 \tau}{1 - \tau} \right] \quad (11)$$

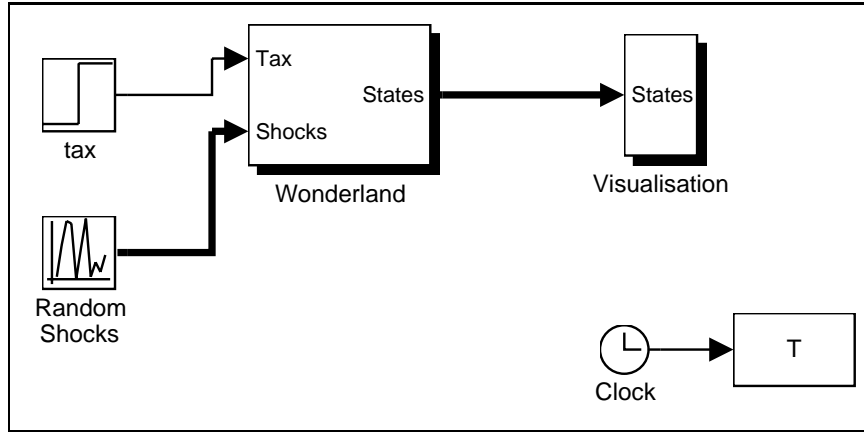


Figure 1: Overall Simulation Environment for Wonderland

It has been assumed that the effect of taxes on output is a scale factor ($\gamma_0 = 0.05$) less than on the rate of pollution output as investment by firms in pollution technology has a greater positive effect on pollution output than its negative effect on economic output. In addition, some of the ‘green’ technologies may be economically beneficial.

Modelling environmental taxes in Wonderland in this manner has a double effect on the flow of pollutants. This is because the taxes have the effect of reducing Y which reduces F . The taxes will also reduce F through the reduction in P . Thus the flow of pollutants is reduced through the decrease in economic output and the decrease in pollution per unit output. Reducing the flow of pollutants allows the stock of natural capital to regenerate. Thus such taxes have a powerful effect in Wonderland.

2.3 Economic and Environmental Disturbances

So far the model has been considered deterministic, it would be more realistic to consider the implications of stochastic disturbances to the economy and to pollution control technology. This is modelled by the introduction of independent white Gaussian noise to Equations 10 and 11.

2.4 State Form

The entire model can be represented in state form as

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t, \mathbf{p}) + H\varepsilon_t \tag{12}$$

where \mathbf{x} is the vector of states ($[Y, N, K, P]^T$), and \mathbf{u} is the vector of controls (which in this case is the parameter τ). Further, ε is the vector of white Gaussian noise, and H is the 4×4 matrix with the value 1 in the first and last elements of the leading diagonal and zeros elsewhere.

3 Visual Simulation of the Model

For the solution of the model the SIMULINK software was used [4]. This software has the advantage that simulation of the model can be undertaken in a convenient interactive environment. It also has the advantage that the flow of data and objects that make up the simulation can be visualised.

Figure 1 shows the overall simulation framework for the model and captures the state space form of Equation 12. The objects in the simulation show a step in taxation and random shocks as inputs;

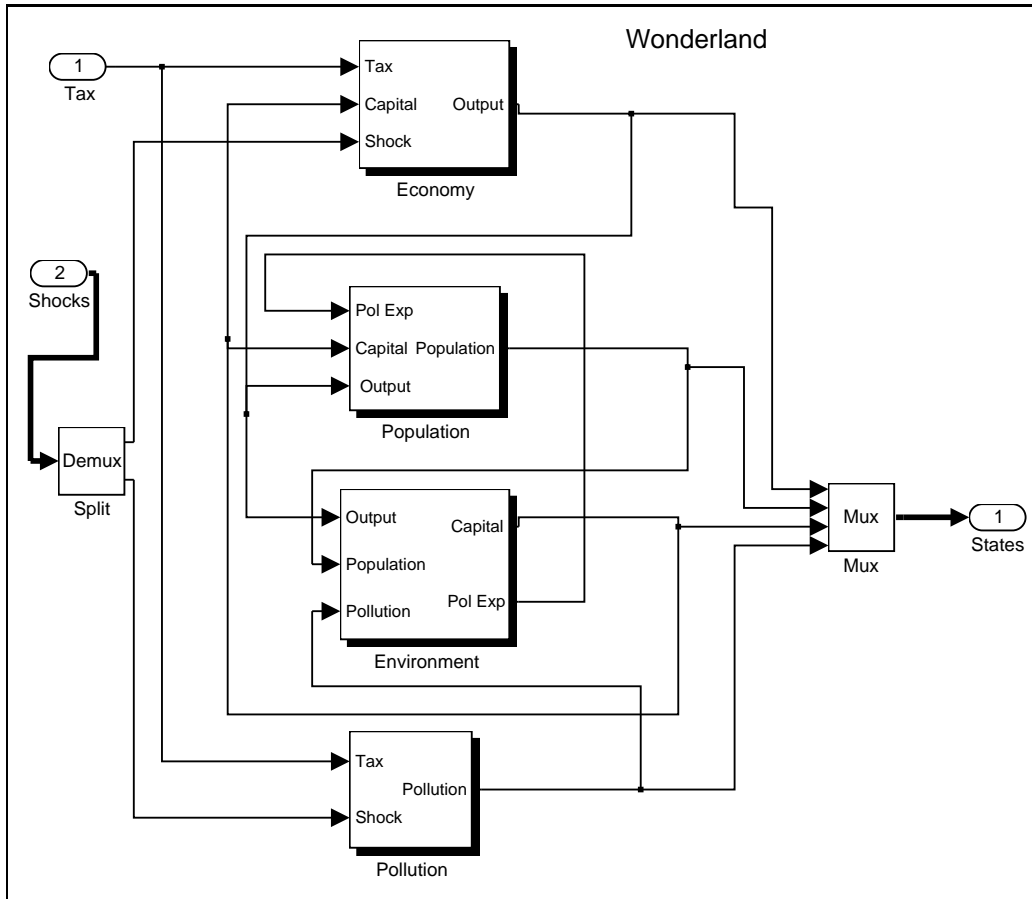


Figure 2: The Sectors of Wonderland

and the transfer of the output to a visualisation object. The simulation accepts easy modification of the tax rate, and variance of the shocks¹.

Figure 2 shows the sections of the model. The flow of data between the equations of the model can be traced in the Figure. Figures 3 through to 6 illustrate the objects and data flows in the sections of the model.

4 The Basic Model Scenarios

Following [5], for all scenarios, initially the economy, Y , and the population, N , and the pollution production technology are all at unity. The stock of natural capital, K , is at 0.98 giving a near complete stock. From this initial state position ($\mathbf{x}_0 = [1, 1, 0.98, 1]^T$), the states trajectories evolve over time. A long time horizon is used to show the effects upon future generations.

Depending upon certain parameter values the states of the model evolve differently generating a variety of scenarios. In this paper we consider the four scenarios shown in Figures 7 to 11, where each Figure presents a different state variable. The different scenarios are generated by differing values of χ , which determines the decrease in pollution per unit output over time; κ , which determines the effectiveness of pollution control; and ω , which determines the impact of pollution flow on the regeneration of natural capital.

¹More details of the capabilities of this software environment in economic models can be found in [3]

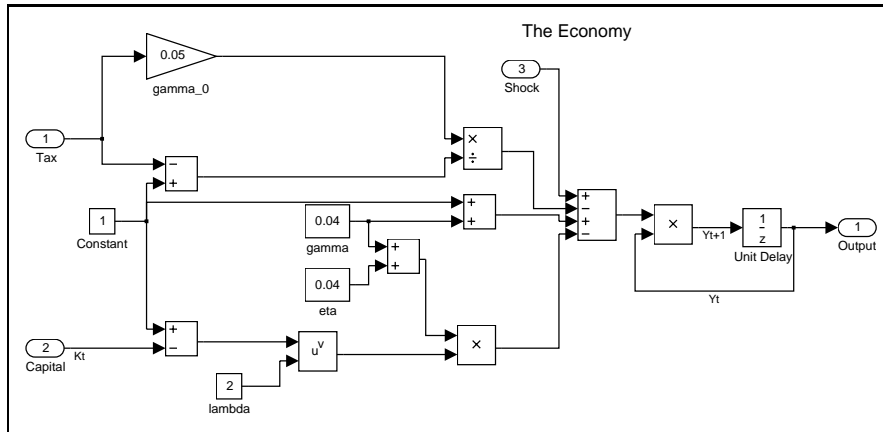


Figure 3: The Economy

4.1 Dream

The Dream scenario is the best situation for the inhabitants of Wonderland. The economy grows continually, the environment returns to and remains in a perfect condition, and the population grows and then stabilises. Under this scenario, there is a future of environmentally sustainable prosperity, all appears well in Wonderland.

This scenario uses the parameters in Table 2. Importantly, the rate of pollution per unit of output (χ) is set to fall by 4% per year. Under this sustainable scenario, both per capita output and net per capita output grow exponentially and natural capital remains undiminished.

The rapid recuperation of the stock of natural capital in this scenario is shown in Figures 9 and 11. The dynamics of natural capital are much faster than other states. The division of the fast-slow dynamical behaviour of the model is considered in detail in [7].

4.2 Horror

Under the Horror scenario, sustainable prosperity gives way to a catastrophic collapse of the economy, population and environment.

The only difference between this and the previous scenario is that χ is set such that rate of pollution per unit output falls by 1% per year. The results for Wonderland are tumultuous. Output and net output crash after a period of growth and this coincides with a complete collapse of the level of natural capital. This brings about a sharp rise in the death rate and the population begins to shrink. Clearly, this path is unsustainable. Earlier generations were living beyond their means both in terms of standard of living and their care for the environment. What is interesting is the abrupt change in the time path of events. For a long time the environment is able to assimilate the growth in the pollution flow, given the technology of pollution production and the amounts spent on pollution control. However, the continued growth of the economy and pollution flows overwhelm the environment's ability to cope. Moreover, continued economic and population growth have brought about a collapse rather than gradual decline. Thus, even with a constant rate of decrease in pollution per unit of output, a horror scenario occurs.

4.3 Escape 1

This scenario varies from the Horror scenario through the effectiveness of the economy's pollution control measures. A hundredfold increase in the parameter κ ($\kappa = 100$) reverses the fall in output through the expenditure control measures. However, this is not enough to halt a brief collapse in the

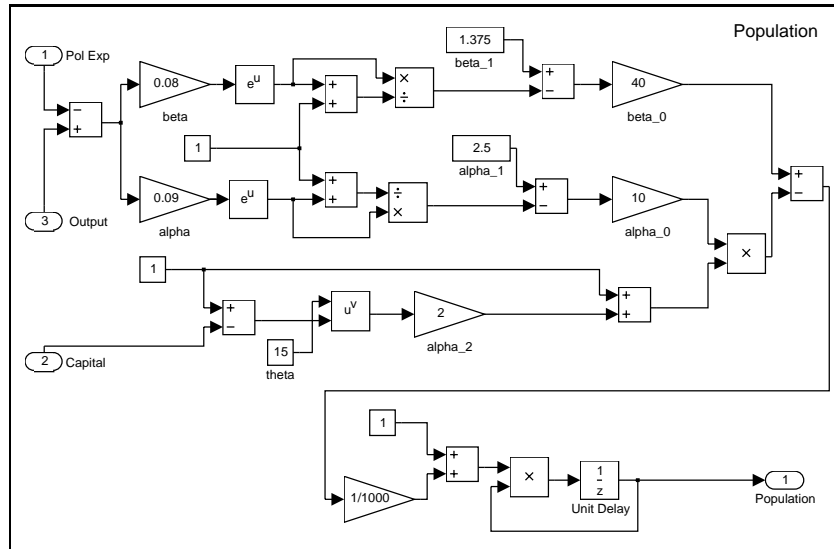


Figure 4: The Population

stock of environmental capital, but it does rebound back quickly. This scenario then repeats itself. These periodic shocks to the environment produce a staggered decline in population growth and cycles in per capita output growth.

4.4 Escape 2

This is the same as Horror scenario except that ω , the impact of the rate of flow of pollution on natural capital has increased. Interestingly, this scenario shows the resilience of natural capital regeneration even with longer periods of environmental collapse.

4.5 Wonderland Phase Portrait

The view of the dynamics of Wonderland can be seen by considering the phase portrait in Figure 11. The Figure considers the three states of economy, population and natural capital. The fourth state, pollution per unit output, is simply an exponential decay.

The first quadrant of the Figure shows the Dream scenario, with the dynamics of natural capital, and population quickly attaining their steady state, and the economy continually growing. In the Horror scenario (second quadrant), all appears well until the crash as all the states descend to the origin. In the first Escape scenario, the phase portrait shows loops as the environment depletes and regenerates, the population stabilises between falls, and the economy grows and crashes. In the second Escape scenario, Wonderland shifts between horror and dream scenarios.

An important implication for Wonderland is that there is only minor parameter value changes between the scenarios. The difference between the Dream and Horror is simply the rate at which pollution per unit of output decays over time. If technological improvements are such that the rate at which pollution per unit output of is 4% per year, then sustainable prosperity occurs in Wonderland. On the other hand, if this rate is 1% per year, then there are troubles in Wonderland.

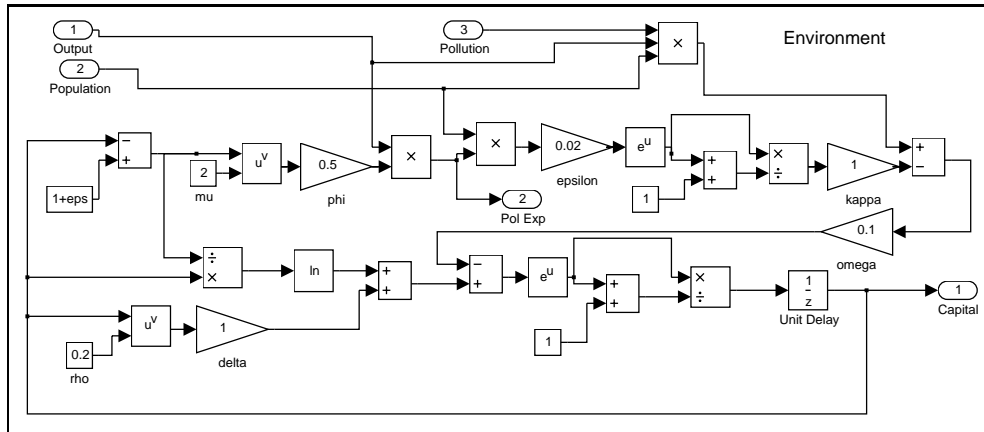


Figure 5: The Environment

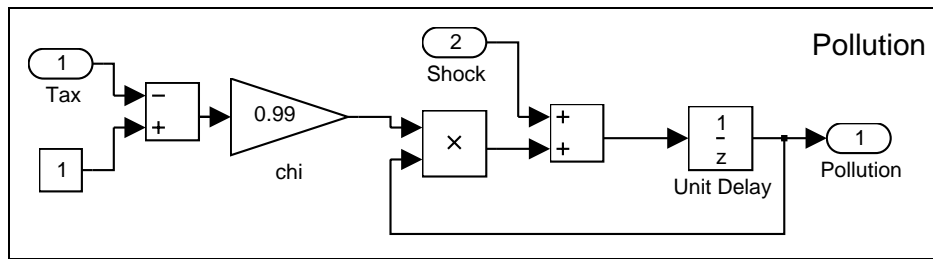


Figure 6: The Pollution Technology

5 Environmental Taxes

If Wonderland is in trouble (Horror scenario) then an interesting question that arises is: ‘what use can be made of environmental taxes to avoid the Horror scenario?’ Given the structure of the model, one approach is to introduce taxes that penalise firms for their pollution per unit of production. This should result in pollution per unit output falling, but may also have effects on economic output as firms transfer their resources into these technologies. The form in which these taxes have been introduced into the model was presented in Section 2.2 above.

Figure 12 presents the phase portrait for Wonderland under differing tax rates. In all cases, the Horror scenario is the starting point.

The first quadrant shows a low tax rate. Under this tax regime the Horror scenario remains, with the economy, environment and population building up to the sudden catastrophic crash. But with higher tax rates, as shown in quadrants two and three, the Horror scenario is transformed into a dream scenario.

The fourth quadrant illustrates what happens in Wonderland if the tax rate is further increased. In this case, the stock of natural capital is, as in the other cases ($\tau = 0.25$ and $\tau = 0.33$), complete, but the economy collapses. However the population increases. Through visual simulation it was established that a tax rate greater than 0.444 produced an economic collapse. The lower bound tax threshold which turns the Horror scenario to the Dream scenario is 0.02.

The tax structure used favours low taxes as the dislocation effect on output rapidly increases with rising taxes. Low taxes are enough to improve pollution flow and enable the stock of natural capital to regenerate, resulting in the Dream scenario.

Another question that may be considered is: ‘how does this tax structure respond when the model

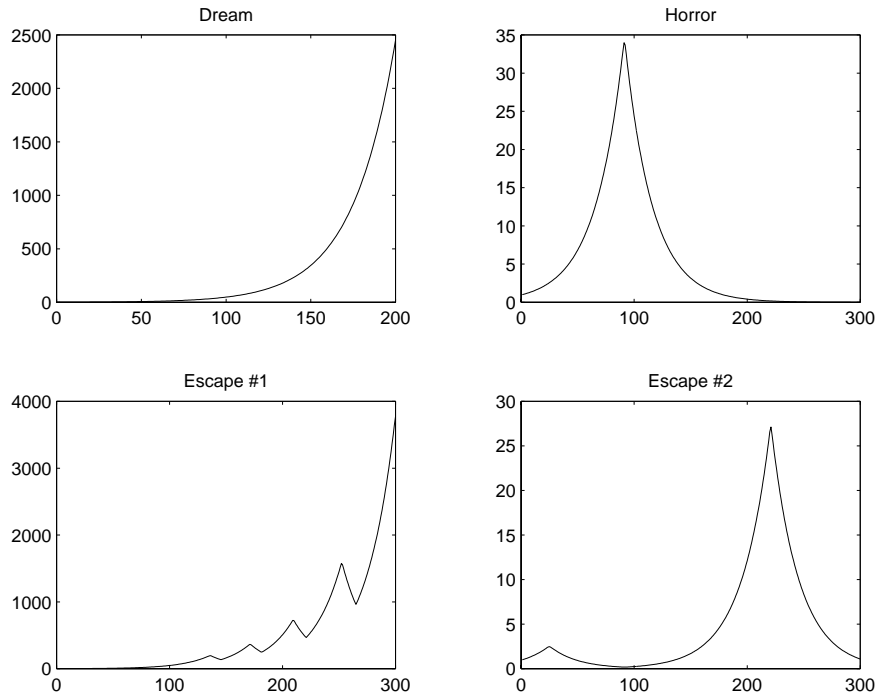


Figure 7: The Economy – Per capita Output (Y) over time under the differing scenarios.

is subjected to external disturbances?’.

6 External Disturbances

The introduction of external shocks into the model is described above in Section 2.3. The issue of concern is whether these shocks on the model mean that it is no longer possible for the tax structure to generate the dream scenario.

Starting with the parameters of the Horror scenario, Figure 13 illustrates the effects of these disturbances on per capita economic output (Y). The Figure plots the evolution of the economy over time to illustrate the continual nature of the shocks. These results are produced using zero means and a fixed variance of 0.001 for the pollution flow shock, and variances of 0.01 (the upper graph) and 0.1 (the lower graph) for the output shock. For these simulations a tax rate of 0.25 was used. It can be seen from the graphs that, within the time horizon, the characteristic dream pattern for output is maintained with the lower variance but not with the higher variance. In simulations with a zero tax rate, the Horror scenario occurs with the disturbance pattern under either variance.

7 Conclusion

This paper examined the issue of sustainable development using the Wonderland model. A characteristic of this model is that seemingly sustainable development becomes unsustainable. Consequently, we considered whether the introduction of an environmental tax could avert this situation. It was shown that this could be achieved. For a variety of tax rates the change to unsustainable development could be avoided. The structure of the tax is being considered in further work.




Figure 8: Population – N over time under the differing scenarios.

References

- [1] Gro Harlem Brundtland. *Our Common Future*. Oxford University Press, Oxford, UK, 1987. The World Commission on Environment and Development.
- [2] Eduard Gröller, Rainer Wegenkittl, Alexandra Milik, Alexia Prskawetz, Gustav Feichtinger, and Warren C. Sanderson. The geometry of Wonderland. *Chaos, Solitons and Fractals*, 7(12):1989–2006, 1996.
- [3] Ric D. Herbert and Rodney D. Bell. Visualisation in the simulation and control of economic models. *Computational Economics*, 10(2):107–118, May 1997.
- [4] Mathworks. *Using SIMULINK Version 2*. The Mathworks Inc, Natick, Massachusetts, USA, 1997.
- [5] Warren C. Sanderson. Simulation models of demographic, economic and environmental interactions. In Wolfgang Lutz, editor, *Population — Development — Environment*, chapter 3. Springer-Verlag, Berlin, 1994.
- [6] Rainer Wegenkittl, Eduard Gröller, and Werner Purgathofer. A guided tour to wonderland: Visualizing the slow-fast dynamics of an analytical dynamical system. Technical Report TR-186-2-95-17, Technical University of Vienna, Austria, 1995.
- [7] Rainer Wegenkittl, Eduard Gröller, and Werner Purgathofer. Visualizing the dynamical behaviour of Wonderland. *IEEE Computer Graphics and Applications*, 17(6):71–79, November/December 1997.




Figure 9: The Environment – Natural Capital (K) over time under the differing scenarios.

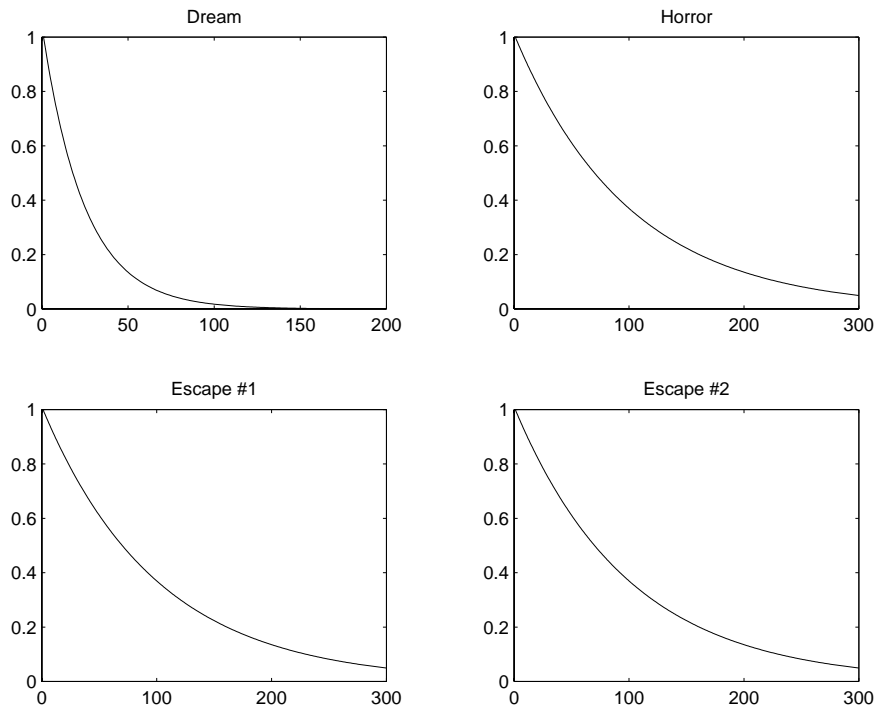


Figure 10: Pollution – Pollution per unit Output (Z) over time under the differing scenarios.

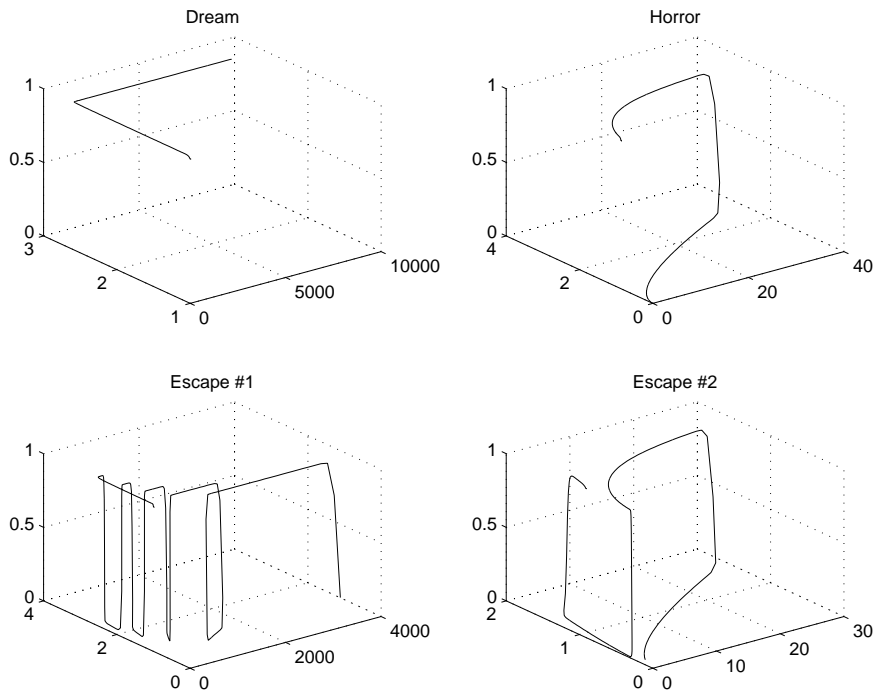


Figure 11: Phase Space – Output (Y), Population (N) and the Environment (K) under the differing scenarios. Output is in the x , Population in the y and the Environment in the z direction on the graphs.

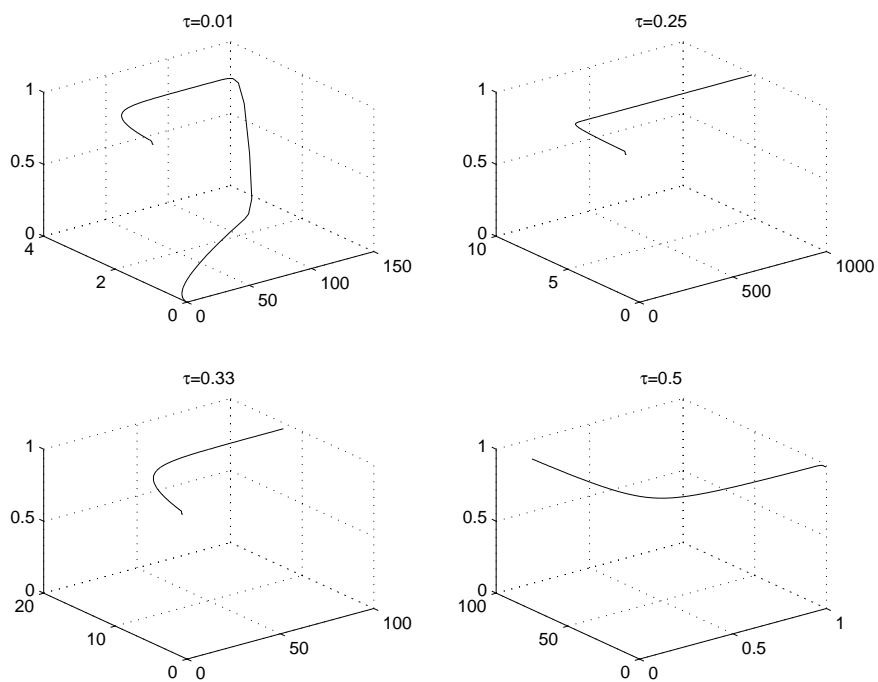


Figure 12: Wonderland Phase Space with fixed taxes – Output (Y), Population (N) and the Environment (K) under differing environmental tax (τ) rates. Output is in the x , Population in the y and the Environment in the z direction on the graphs.




Figure 13: The Economy (Y) under varying Disturbances. The upper diagram has a variance of 0.01, and the lower diagram has a variance of 0.1.

Competition between strategies for a market selection game

Hisao Ishibuchi Chi-Hyon Oh and Tomoharu Nakashima
Department of Industrial Engineering, Osaka Prefecture University
Gakuen-cho 1-1, Sakai, Osaka 599-8531, Japan
hisaoi/oh/nakashi@ie.osakafu-u.ac.jp

Abstract

Our game in this paper is a non-cooperative repeated game where many players compete with one another at several markets. At each iteration of our game, each player is supposed to choose a single market for maximizing his own profit obtained by selling his product at the selected market. It is assumed in our market selection game that the market price of the product is determined by the demand-supply relation at each market. For example, if many players bring their products to a particular market, the market price at that market becomes low. On the contrary, the market price is high if the total amount of products brought to the market is small. In this manner, the market price at each market is determined by the actions of all players. This means that the profit of each player depends on the actions of the other players. The main aim of this paper is to numerically analyze the competition between several strategies for our market selection game. We examine six strategies: a random selection strategy, a minimum transportation cost strategy, an optimal strategy for the previous actions, a mimic strategy of the nearest neighbor player, a Q -learning-based strategy, and a fuzzy Q -learning-based strategy. The performance of each strategy is examined by computer simulations with 100 players and five markets. First we examine the performance of each strategy by assuming that all the players use the same strategy. Next we examine the competition between two strategies, each of which are used by half of the players. Then we examine the competition between two strategies further by changing the number of players adopting each strategy. Finally we perform computer simulations of the competition between various strategies where the number of players adopting each strategy is increased or decreased according to the performance of the strategy.

1 Introduction

Strategies for repeated games have been mainly studied for the IPD (Iterated Prisoner's Dilemma) game. Axelrod [1] discusses many strategies based on the results of two computer tournaments for the IPD game. The evolution of strategies for the IPD game was studied in Axelrod [2], Lindgren [3], and Fogel [4]. In those studies, each player played the IPD game against all players in a current population. Nowak et al. [5] and Lloyd [6] examined a spatial version of the IPD game where a number of players were spatially fixed in a grid-world and they played the IPD game against only their neighboring players. While those studies on the IPD game involved many players with various strategies, the Prisoner's Dilemma game itself is very simple. That is, only two players repeatedly play the game against each other based on a simple 2×2 payoff matrix.

Our market selection game [7] involves much more players (e.g., 100 players) and a more complicated payoff mechanism than the IPD game. At each iteration of our game, each player is supposed to choose a single market from several ones (e.g., five markets) to sell his product at the market price of the selected market. The aim of the market selection is to maximize his own profit obtained by selling his product at the selected market. It is assumed in our game that the market price is determined by the demand-supply relation at each market. This means that the profit of a particular player depends on the actions of the other players. In this paper, we examine six strategies for our market selection game: a random selection strategy, a minimum transportation cost strategy, an optimal strategy for the previous actions, a mimic strategy of the nearest neighbor player, a Q -learning-based strategy, and a fuzzy Q -learning-based strategy. For examining the performance of each strategy by computer simulations, we set up an instance of our market selection game as shown in Fig. 1 where 100 players and five markets are randomly located. Our computer simulations in this paper are performed for the market selection game in Fig. 1.




Figure 1: Market selection game in our computer simulations.




Figure 2: Illustration of the action of each player.

2 Formulation of a Market Selection Game

We denote the number of players by n ($n = 100$ in Fig. 1). Each player is indexed by i where $i = 1, 2, \dots, n$. The number of markets is denoted by m ($m = 5$ in Fig. 1). Each market is indexed by j where $j = 1, 2, \dots, m$. Our market selection game is iterated as the IPD game. Let us denote the total number of iterations by T . Each iteration is indexed by t (i.e., $t = 1, 2, \dots, T$). We assume that each player has a single product to be sold at each iteration. The action of each player at each iteration is to select a single market where his product is sold. Let us denote the action of the i -th player at the t -th iteration of our game by x_{ij}^t where $j = 1, 2, \dots, m$; $t = 1, 2, \dots, T$; and

$$x_{ij}^t = \begin{cases} 1, & \text{if the } i\text{-th player chooses the } j\text{-th market,} \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Because each player is supposed to choose a single market from the given m markets for selling his product, the following relation holds:

$$\sum_{j=1}^m x_{ij}^t = 1, \text{ for } i = 1, 2, \dots, n; t = 1, 2, \dots, T. \quad (2)$$

We assume that all the players simultaneously perform the market selection at each iteration of our game. Thus no player knows the current actions of the other players when he chooses a market. This means that no player knows the optimal market selection for the current iteration of the game. In Fig. 2, we illustrate an example of the market selection by 100 players, each of which selects a single market from the given five markets.

We assume that the market price of the product is determined by the demand-supply relation at each market. For example, if many players bring their products to a particular market, the market price at that market becomes low. On the contrary, the market price is high if the total amount of products brought to the market is small. In this manner, the market price at each market is determined by the actions of all players. The total amount of products that are sold in the j -th market at the t -th iteration is calculated from (1) as follows:

$$X_j^t = \sum_{i=1}^n x_{ij}^t, \text{ for } j = 1, 2, \dots, m; t = 1, 2, \dots, T. \quad (3)$$

We assume that the market price of the j -th market at the t -th iteration is determined by the following linear demand-supply relation:

$$p_j^t = a_j - b_j \cdot X_j^t, \text{ for } j = 1, 2, \dots, m; t = 1, 2, \dots, T, \quad (4)$$

where a_j and b_j are positive constants that specify the demand-supply relation of the j -th market. In computer simulations of this paper, we use the same demand-supply relation for all the five markets:

$$p_j^t = 100 - 3 \cdot X_j^t, \text{ for } j = 1, 2, \dots, 5; t = 1, 2, \dots, T. \quad (5)$$

It is assumed that the cost c_{ij} for the transportation of the product from the i -th player to the j -th market depends on the distance between the player and the market. Let us denote the distance between the i -th player and the j -th market by d_{ij} . We assume that the transportation cost c_{ij} is given as follows:

$$c_{ij} = c \cdot d_{ij}, \text{ for } i = 1, 2, \dots, n; j = 1, 2, \dots, m, \quad (6)$$

where c is the transportation cost for the unit distance. In computer simulations, we specified the value of c as $c = 1$. Thus the transportation cost c_{ij} was specified as $c_{ij} = d_{ij}$. The introduction of the transportation cost makes the players different from each other because they are randomly placed in the $[0, 100] \times [0, 100]$ space. That is, some players located close to markets can easily enjoy high profits with small transportation costs while other players with no markets in their neighborhood suffer from high transportation costs.

Let us denote the profit (i.e., reward) of the i -th player at the t -th iteration by r_i^t . We define the profit r_i^t as follows when the i -th player chooses the j -th market for selling his product (i.e., when $x_{ij}^t = 1$):

$$r_i^t = p_j^t - c_{ij}, \text{ for } i = 1, 2, \dots, n; t = 1, 2, \dots, T. \quad (7)$$

It should be noted that the profit r_i^t of the i -th player depends on the actions of the other players through the market price p_j^t (see (3) and (4)). The aim of each player in our game is to maximize the total profit r_i over T iterations:

$$r_i = \sum_{t=1}^T r_i^t, \text{ for } i = 1, 2, \dots, n. \quad (8)$$

3 Various Game Strategies

In this paper, we examine the performance of six strategies for our market selection game: a random selection strategy, a minimum transportation cost strategy, an optimal strategy for the previous actions, a mimic strategy of the nearest neighbor player, a Q -learning-based strategy, and a fuzzy Q -learning-based strategy. Each strategy is explained in the following subsections.

3.1 Random Strategy

The simplest strategy for our market selection game is a random strategy. The random strategy can be implemented by specifying the market selection probability $\Pr(x_{ij}^t = 1)$ as

$$\Pr(x_{ij}^t = 1) = 1/m, \text{ for } i = 1, 2, \dots, n; j = 1, 2, \dots, m; t = 1, 2, \dots, T. \quad (9)$$

That is, each of the m markets is randomly selected with the same probability.

3.2 Minimum Transportation Cost Strategy

Another simple strategy is to minimize the transportation cost. That is, each player chooses the nearest market with the minimum transportation cost. This strategy, which is referred to as a minimum transportation cost strategy, is very effective when the transportation cost c for the unit distance is large.

3.3 Optimal Strategy for Previous Actions

Since each player simultaneously performs the market selection at each iteration of our game, no player knows the current actions of the other players when he chooses a market. This means that no player can perform the optimal market selection for the current iteration of the game. Each player, however, can calculate the optimal market for the previous actions of the other players. That is, a player can select a market by assuming that the other players choose exactly the same markets as in the previous iteration. This strategy is referred to as an optimal strategy for the previous actions. Of course, the optimal strategy for the previous actions is not always optimal for the current actions because the underlying assumption is not always valid. At the first iteration of our game, each player adopting this strategy chooses the nearest market with the minimum transportation cost. The next action of each player is chosen based on the previous actions of the other players.

3.4 Mimic Strategy of Nearest Neighbor Player

The optimal strategy for the previous actions requires the information about the demand-supply relations of all the m markets (i.e., a_j and b_j for all the m markets) and the actions of the other players at the previous iteration of the game. When such information is not available, we need simpler strategies that do not require a lot of information. One of such simple strategies is a mimic strategy of the nearest neighbor player where a player simply mimics the previous action of the nearest neighbor player. A player adopting this strategy randomly chooses a market at the first iteration of the game. The next action of the player is the same as the previous action of his nearest neighbor player.

3.5 Q -Learning-based Strategy

The four strategies described in Subsections 3.1 ~ 3.4 do not use any information about the profit obtained from each market during the previous iterations. In order to choose a market based on the profit obtained from each market during the previous iterations, we can use Q -learning [8], which is a well-known reinforcement learning scheme. In a Q -learning-based strategy, each player stores and updates a Q -value for each market during the execution of our repeated game (for detail, see [7]). Let Q_{ij}^t be the Q -value of the i -th player for the j -th market at the t -th iteration of our game. The Q -value for the selected market is modified after the t -th iteration as

$$Q_{ij}^{t+1} = \begin{cases} (1 - \alpha) \cdot Q_{ij}^t + \alpha \cdot r_i^t, & \text{if } x_{ij}^t = 1, \\ Q_{ij}^t, & \text{otherwise,} \end{cases} \quad (10)$$

where α is a positive learning rate. In our computer simulations, α was specified as $\alpha = 0.9$. The initial value of each Q -value was specified as $Q_{ij}^1 = 100$.

The market selection is performed based on the Q -value for each market. The selection probability $\Pr(x_{ij}^t = 1)$ of each market is defined by the roulette wheel selection with the linear scaling as follows:

$$\Pr(x_{ij}^t = 1) = \frac{Q_{ij}^t - \min\{Q_i^t\}}{\sum_{k=1}^m \{Q_{ik}^t - \min\{Q_i^t\}\}}, \quad (11)$$

where $\min\{Q_i^t\} = \min\{Q_{ij}^t \mid j = 1, 2, \dots, m\}$. In our computer simulations, this roulette wheel selection was used for the first 100 iterations of our game. After the 100 iterations, the market with the maximum Q -value (i.e., $\max\{Q_{ij}^t \mid j = 1, 2, \dots, m\}$) was always selected at each iteration of our game.




Figure 3: Membership functions of “low” and “high”.




Figure 4: Illustration of price changes.

3.6 Fuzzy Q-Learning-based Strategy

For handling continuous states and/or actions, Q -learning was extended to fuzzy Q -learning [9, 10]. In our market selection problem, the market price at each market in the previous iteration can be used as a continuous state (see [7]). That is, the estimation of the expected payoff from each market (i.e., the Q -value for each market) is conditioned by the market prices of all markets in the previous iteration. In our computer simulations, the market price of each market was partitioned into two linguistic values “low” and “high” in Fig. 3. These two linguistic values were used as antecedent fuzzy sets of fuzzy if-then rules of the following type:

$$\text{Rule } R_s : \text{ If } p_1^{t-1} \text{ is } A_{s1} \text{ and } \dots \text{ and } p_m^{t-1} \text{ is } A_{sm} \\ \text{ then } Q_{i1}^t = q_{si1}^t \text{ and } \dots \text{ and } Q_{im}^t = q_{sim}^t, \quad (12) \\ s = 1, 2, \dots, N,$$

where R_s is the label of the s -th fuzzy if-then rule, s is a rule index, p_j^{t-1} is the market price of the j -th market at the previous iteration, A_{sj} is an antecedent fuzzy set, Q_{ij}^t is a Q -value, q_{sij}^t is a consequent real number, and N is the number of fuzzy if-then rules. When we have two linguistic values (“low” and “high”) as antecedent fuzzy sets for each market price of the m markets, the number of fuzzy if-then rules for each player is $N = 2^m$. In our computer simulations with five markets (i.e., $m = 5$), the number of fuzzy if-then rules for each player is $N = 2^5 = 32$. The Q -value of each player for each market is calculated by a fuzzy reasoning method from the fuzzy if-then rules in (12). Let us define the compatibility grade of the previous market prices $\mathbf{p}^{t-1} = (p_1^{t-1}, p_2^{t-1}, \dots, p_m^{t-1})$ with the fuzzy if-then rule R_s by the product operator as

$$\mu_s(\mathbf{p}^{t-1}) = A_{s1}(p_1^{t-1}) \times \dots \times A_{sm}(p_m^{t-1}), \quad (13)$$

where $A_{sj}(\cdot)$ is the membership function of the antecedent fuzzy set A_{sj} . The Q -value of the i -th player for the j -th market at the t -th iteration is calculated as follows:

$$Q_{ij}^t = \frac{\sum_{s=1}^N \mu_s(\mathbf{p}^{t-1}) \cdot q_{sij}^t}{\sum_{s=1}^N \mu_s(\mathbf{p}^{t-1})}, \text{ for } i = 1, 2, \dots, n; j = 1, 2, \dots, m; t = 1, 2, \dots, T. \quad (14)$$

The consequent q_{sij}^t of each fuzzy if-then rule is adjusted in the same manner as in the Q -learning:

$$q_{sij}^{t+1} = \begin{cases} \{1 - \alpha \cdot \mu_s^*(\mathbf{p}^{t-1})\} \cdot q_{sij}^t + \alpha \cdot \mu_s^*(\mathbf{p}^{t-1}) \cdot r_i^t, & \text{if } x_{ij}^t = 1, \\ q_{sij}^t, & \text{otherwise,} \end{cases} \quad (15)$$

Table 1: Average profit by each strategy when all the players use the same strategy.

Strategy	Average Profit
Random selection strategy	-16.7
Minimum transportation cost strategy	3.7
Optimal strategy for the previous actions	-57.5
Mimic strategy of the nearest neighbor player	-20.5
Q -learning-based strategy	11.8
Fuzzy Q -learning-based strategy	9.8

where

$$\mu_s^*(\mathbf{p}^{t-1}) = \frac{\mu_s(\mathbf{p}^{t-1})}{\sum_{s=1}^N \mu_s(\mathbf{p}^{t-1})}. \quad (16)$$

The market selection based on the calculated Q -values in (14) is performed in the same manner as in the Q -learning.

4 Computer Simulations

In this section, we examine the performance of each strategy and analyze the competition between various strategies by computer simulations on the market selection game in Fig. 1.

4.1 Performance of Each Strategy

In this subsection, we assume that all the players use the same strategy. In computer simulations, we applied each strategy to the market selection game in Fig. 1. Our game was repeated 1000 times (i.e., $t = 1, 2, \dots, 1000$). For calculating the average profit per each iteration and each player, such a computer simulation of our repeated game was performed 100 times for each strategy because some strategies have stochastic nature based on randomization procedures. Average profit obtained by each strategy is summarized in Table 1. From this table, we can see that better results were obtained by the Q -learning-based strategy and the fuzzy Q -learning-based strategy. We can also see that the optimal strategy for the previous actions was the worst among the six strategies. This can be explained as follows. In general, this strategy chooses a market where the previous price was high. Thus many players choose the same market. As a result, the market price becomes low. The market price will be high again in the next round of the game because only a few players will choose that market after the low price. In computer simulations of this paper, we often observed such a cycle of price changes, which is illustrated in Fig. 4. If the price of a market was high (e.g., the point A in Fig. 4) in the previous round of the game, the total amount of products would be increased and the market price would be decreased (e.g., to the point B). The low price will decrease the total amount of products and increase the market price in the next round of the game (e.g., to the point C). Due to such cyclic concentration, each player could not obtain high profit on the average when all the players adopted the optimal strategy for the previous actions. As we will show in Subsections 4.2 and 4.3, this is not the case when the number of players adopting this strategy is small.




Figure 5: Illustration of price changes.


Figure 6: Simulation results of the competition between the fuzzy Q -learning-based strategy and the optimal strategy for the previous actions.

Table 2: Simulation results of the competition between two strategies.

Strategy	Strategy of the other 50 players					
	Random	Cost	Optimal	Mimic	Q	Fuzzy Q
Random	-16.7	-16.9	-16.9	-16.7	-16.7	-16.7
Cost	9.8	3.7	14.8	4.6	12.7	2.8
Optimal	-25.1	-21.0	-57.5	-2.1	-39.0	2.7
Mimic	-17.6	-3.1	4.9	-20.5	2.1	0.2
Q	13.4	13.2	8.9	13.0	11.8	12.8
Fuzzy Q	9.8	10.4	12.6	9.9	10.1	9.8

4.2 Competition between Two Strategies

In this subsection, we examine the competition between two strategies where each strategy is adopted by half of the players (i.e., 50 players). All combinations of two strategies were examined in the same manner as in Subsection 4.1. That is, the competition between each pair of strategies was examined by 100 independent trials of the repeated game with 1000 iterations. In each trial, 100 players were randomly divided into two groups of 50 players. Simulation results are summarized in Table 2. The bold-faced entries in this table are the same as Table 1 because those entries correspond to the situations where two strategies are the same. By focusing our attention on each row of Table 2, we can see that the average profit by each strategy (except for the random selection) was increased by competing with different strategies. For example, a positive average profit was obtained by the optimal strategy for the previous action when it competed with the fuzzy Q -learning-based strategy. On the other hand, by focusing our attention on each column, we can see that the fuzzy Q -learning-based strategy could bring positive profits to the other strategies except for the random selection.

As shown in Table 2, the performance of each strategy can be improved by competing with different strategies. By demonstrating this improvement more clearly, we performed similar computer simulations by changing the number of players adopting each strategy. We examined the competition between the optimal strategy for the previous actions and the fuzzy Q -learning-based strategy. The number of players adopting the fuzzy Q -learning-based strategy was specified as 1, 10, 20, \dots , 99. For each specification, average profit of each strategy was calculated in the same manner as in the above computer simulations. Simulation results are summarized in Fig. 5. From this figure, we can see that




Figure 7: Simulation results of the competition between the fuzzy Q -learning-based strategy and the mimic strategy of the nearest neighbor player.




Figure 8: Simulation results of the natural selection for strategies.

Table 3: Average profit by each strategy in the competition of the six strategies.

Strategy	Average Profit
Random selection strategy	-16.5
Minimum transportation cost strategy	13.7
Optimal strategy for the previous actions	11.2
Mimic strategy of the nearest neighbor player	5.8
Q -learning-based strategy	13.8
Fuzzy Q -learning-based strategy	9.6

the performance of each strategy was high when the number of that strategy was small. This means that the performance of each strategy was improved by the existence of different strategies. We also examined the competition between the mimic strategy of the nearest neighbor player and the fuzzy Q -learning-based strategy. Simulation results are summarized in Fig. 6.

4.3 Competition between Six Strategies

In this subsection, we examine the competition between the six strategies. In computer simulations, we specified the number of players adopting each strategy as follows:

- Random selection strategy: 17 players,
- Minimum transportation cost strategy: 17 players,
- Optimal strategy for the previous actions: 17 players,
- Mimic strategy of the nearest neighbor player: 17 players,
- Q -learning-based strategy: 16 players,
- Fuzzy Q -learning-based strategy: 16 players.

In the same manner as in the previous computer simulations, the performance of each strategy was calculated by 100 independent trials. In each independent trial, players adopting each strategy were randomly selected. Simulation results are summarized in Table 3. From the comparison between Table 1 and Table 3, we can see that the performance of some strategies was significantly improved by competing with many strategies.

We also applied the concept of natural selection to our computer simulation where one player

Table 4: Average profit by each strategy when all the players use the same strategy (different parameter specifications).

Strategy	Average Profit
Random selection strategy	-83.7
Minimum transportation cost strategy	10.6
Optimal strategy for the previous actions	12.0
Mimic strategy of the nearest neighbor player	-85.5
Q -learning-based strategy	8.2
Fuzzy Q -learning-based strategy	7.8

adopting the worst strategy changed his strategy into the best strategy during the generation update. The generation update was performed after the performance of each strategy was evaluated. That is, the number of players adopting the best strategy was increased by the generation update while the number of players with the worst strategy was decreased. For example, the number of Q -learning players was increased from 16 to 17 by the first generation update in the case of Table 3 while the number of random selection strategy players was decreased from 17 to 16. In this manner, we evolved the population of 100 players in our market selection game until 200-th generation. Simulation results are summarized in Fig. 7. Furthermore, we continued the generation update until the 1000-th generation. The minimum transportation cost strategy, the Q -learning-based strategy and the optimal strategy for previous actions coexisted with each other after such a large number of generation updates.

5 Conclusion

In this paper, we examined the performance of six strategies for a market selection game. Our game is a non-cooperative repeated game with many players and a complicated payoff mechanism. First we examined the performance of each strategy by computer simulations where all the players employed the same strategy. Next we examined the competition between two strategies where each of the two strategies was adopted by half of the players. For examining the competition between two strategies further, we also performed computer simulations where the number of players adopting each strategy was specified variously. Finally we examined the competition between six strategies. We also applied the concept of natural selection to our computer simulation for evolving the population of 100 players in our market selection game. By the computer simulations in this paper, we showed that the performance of each strategy was better in the competition with difference strategies than the situation where all the players employed the same strategy. Since the performance of a particular strategy adopted by a player strongly depends on the choice of strategies by the other players, it is very difficult to generally show which strategy works well. For example, while the optimal strategy for the previous actions was good in the competition between the six strategies in Subsection 4.3, its performance was the worst in Subsection 4.1 where all the players employed the same strategy. The performance of each strategy also strongly depends on the choice of parameter values in our market selection game. In the computer simulations of this paper, we specified the demand-supply relation and the unit transportation cost as $p_j^t = 100 - 3 \cdot X_j^t$ and $c = 1.0$, respectively. We also examined the performance of each strategy in the same manner as in Table 1 using different parameter specifications: $p_j^t = 100 - X_j^t$ and $c = 3.0$. In this case, we had totally different results as shown in Table 4. In this table, the best average result was obtained by the optimal strategy for the previous action, which was the worst strategy in Table 1. As we can easily see from the comparison between Table 1 and Table 4, simulation results on our market selection game strongly depends on parameter specifications. Examination of such dependency is one of future research topics on our market selection game.

Acknowledgments

This work was partially supported by Foundation for Fusion of Science & Technology.

References

- [1] Axelrod, R.: *The Evolution of Cooperation*. Basic Book. New York. (1984).
- [2] Axelrod, R.: "The Evolution of Strategies in the Iterated Prisoner's Dilemma," In Davis, L. (ed.). *Genetic Algorithms and Simulated Annealing*. Morgan Kaufmann. Los Altos. (1987) 32–41.
- [3] Lindgren, K.: "Evolution Phenomena in Simple Dynamics," In Langton, C. G., Taylor, C., Farmer, J. D., Rasmussen, S. (eds.). *Artificial Life II*. Addison-Wesley. (1991) 295–312.
- [4] Fogel, D. B.: "Evolving Behaviors in the Iterated Prisoner's Dilemma," *Evolutionary Computation*. **1** (1993) 77–97.
- [5] Nowak, M. A., May, R. M., and Sigmund, K.: "The Arithmetic of Mutual Help," *Scientific America*. **June** (1997) 50–53.
- [6] Lloyd, A. L.: "Computing Bouts of the Prisoner's Dilemma," *Scientific America*. **June** (1997) 80–83.
- [7] Ishibuchi, H., Nakashima, T., Miyamoto, H., Oh, C. -H.: "Fuzzy Q -Learning for a Multi-Player Non-Cooperative Repeated Game," *Proc. of 6th International Conference on Fuzzy Systems*. (1997) 1573–1579.
- [8] Watkins, C. J. C. H., Dayan, P.: " Q -Learning," *Machine Learning*. **8** (1992) 279–292.
- [9] Glorennec, P. Y.: "Fuzzy Q -Learning and Dynamical Fuzzy Q -Learning," *Proc. of 3rd International Conference on Fuzzy Systems*. (1994) 474–479.
- [10] Jouffe, L., Glorennec, P. -Y.: "Comparison between Connectionist and Fuzzy Q -Learning," *Proc. of 4th International Conference on Soft Computing*. (1996) 557–560.

The Emergence and Collapse of the State: A Game Theoretic Analysis with Computer Simulations

Atsushi Iwasaki
Graduate School of Science
and Technology,
Kobe University, 1-1 Rokkodai-cho,
Nada-ku, Kobe 657, Japan
iwasaki@mi-2.mech.kobe-u.ac.jp

Sobei H. Oda
Faculty of Economics,
Kyoto Sangyo University,
Motoyama, Kamigamo,
Kita-ku, Kyoto 603, Japan
oda@cc.kyoto-su.ac.jp

Kanji Ueda
Faculty of Engineering, Kobe University,
1-1 Rokkodai-cho, Nada-ku, Kobe 657, Japan
ueda@mech.kobe-u.ac.jp

Abstract

This paper describes how a state emerges and collapses that makes it possible for citizens to do something which they will not voluntarily. The model is the generalisation of Okada and Sakakibara (1991): people may voluntarily make a state, or a compulsory rule that force themselves to construct and maintain the public capital stock. The novelty of this paper lies in the assumption that productivity and benefit from the stock may differ from person to person. This paper presents both game theoretic analysis and results of simulations, which suggest that selfish but rational people may agree to make a state, which grows as the public capital stock accumulates but collapses when the stock reaches a certain level.

1 Introduction

The tragedy of commons is a well-known example of how non-altruistic people fail to cooperate for maintaining the public capital. We should like examine the possibility that people voluntarily make a state or a compulsory rule that force themselves to construct and maintain the public capital stock in the circumstances.

Our analysis is the generalisation of the analysis of Okada and Sakakibara (1991), which explains the emergence and collapse of a state in terms of game theory on the supposition that the productivity and the benefit from the public capital stock are common to all players. We shall examine the same problem in terms of game theory and computer simulations on the assumption that productivity and benefit from the public capital stock may differ from person to person. We shall see how non-altruistic people may agree to make a state, which grows as the public capital stock accumulates but collapses when the stock reaches a certain level.

This paper is organised in the following way. In Section 2 we shall explain the basic model, which is divided into four subgames: first each inhabitant announces whether he or she becomes a citizen or an outsider; then the citizen who proposes the smallest salary is chosen to be the enforcer who watches for tax evasion without making private business; then all citizens propose tax rates, of which the smallest is adopted as the tax rate; last tax payers pay taxes honestly or become tax evaders, whose income from private business will be all confiscated by the enforcer if tax evasion is found by him or her. In Section 3 we shall present a simple case and analyse all the subgames mathematically to find the subgame-perfect Nash equilibrium of the game by backward induction. In Section 4 we shall show some results of simulations for the case analysed in the previous case. Last in Section 5 we shall refer to the generalisation of our analysis.

2 The model

The outline of our model is as follows. There live n inhabitants in a valley irrigated by a canal. Inhabitants make shovels in the winter. They may exchange them for rice with foreigners living outside the valley and/or use them up for dredging the canal to increase their harvest of rice in the fall. The non-agricultural income of inhabitants is defined as the rice value of the shovels made by them, which may differ from person to person according to their skill, while their non-agricultural income is their harvest of rice, which depends on the location of their private rice field and the condition of the canal as the public capital.

In addition, for the sake of simplicity, let us assume the following. The condition of the canal is represented by its depth K . Defining a unit of shovels adequately, we assume that each additional input of non-agricultural input deepens the canal by one inch. The non-agricultural and the agricultural income of the i th inhabitant are γ_i and $\beta_i K$.

In the circumstances inhabitants may be faced with a difficulty in dredging the irrigation system. Suppose that the depth of the canal is \bar{K} without being dredged. The total income of the i th inhabitant is $\gamma_i + \beta_i \bar{K}$ if no inhabitants use their shovels to dredge the canal, while it equals $(1 - t_i)\gamma_i + \beta_i(\bar{K} + \sum_{j=1}^n t_j \gamma_j)$ if the j th inhabitant contributes 100 t_j percent of his or her non-agricultural income to the dredging of the canal. The latter can possibly be greater, but even then it may not automatically be realized. Certainly those inhabitants with $1 < \beta_j$ may voluntarily contribute all their non-agricultural income to the dredging of the canal, but those with $\beta_j < 1$ will contribute nothing; in terms of the game theory $t_j = 1$ and $t_j = 0$ are the dominant strategies for the respective groups.

Some organization or the system of monitoring and punishment may be required for making people contribute to the accumulation and maintenance of the public capital stock. In this paper we examine the, following scenario or the four-stage game.

Subgame 1

Each inhabitant announces whether he or she becomes a citizen or an outsider. Accordingly the n inhabitants in the valley: {Inhabitant i | $i \in N = \{1, 2, \dots, n\}$ } are divided into m ($0 \leq m \leq n$) citizens: {Inhabitant i | $i \in M \subset N$, } and $n - m$ outsiders: {Inhabitant i | $i \in L = N - M$ }, where $M \cup L = N$ and $M \cap L = \emptyset$. The citizens advance towards the following stages to determine their role or duty as well as the penalty which may be imposed on those who do not perform it, while the outsiders can enjoy all benefit from the public capital without making any contribution to its accumulation or maintenance.

Subgame 2

Every citizen announces the acceptable tax rate on non-agricultural income τ_i . The minimum τ_i is adopted as the tax rate of the state: $\tau^* = \min_{i \in M} \tau_i$. (Everyone can virtually dissolve the state by proposing $\tau_i = 0$.)

Subgame 3

Every citizen offers him/herself as the candidate for the enforcer who makes neither shovels nor rice to concentrate on monitoring the other citizens, by declaring the ratio of the enforcer's salary to the tax revenue of the state θ_i . The person who has proposed the minimum θ_i is elected as the enforcer and has salaries paid accordingly: if $\min_{i \in M} \theta_i = \theta_e = \theta^*$, Inhabitant e is the enforcer, whose salary is $\tau^* \sum_{j \in T} \gamma_j$ where Inhabitants i ($i \in T$) pay taxes honestly while Inhabitants j ($j \in U$) pay no taxes ($M = \{e\} \cup T \cup U$ and $\{e\} \cap T = T \cap U = U \cap \{e\} = \emptyset$). In addition to the salary, the enforcer can confiscate all non-agricultural income of the tax evaders he or she has found out, whose expected value is $\epsilon_{m-1}^e \sum_{j \in U} \gamma_j$ as his/her income. Here it is assumed that Inhabitant e can find out each tax evader at the probability of ϵ_l^e if he or she monitors l citizens.

Subgame 4

The $m - 1$ tax payers make shovels and rice, and pay or do not pay taxes. As the result, Inhabitant i expects the following income:

$$E_i = \begin{cases} E_i^e &= \theta^* \tau^* \sum_{j \in T} \gamma_j + \epsilon_{m-1}^e \sum_{j \in U} \gamma_j & \text{if } i = e \\ E_i^T &= (1 - \tau^*) \gamma_i + \beta_i \{ (1 - \theta^*) \tau^* \sum_{j \in T} \gamma_j + K \} & \text{if } i \in T \\ E_i^U &= (1 - \epsilon_{m-1}^e) \gamma_i + \beta_i \{ (1 - \theta^*) \tau^* \sum_{j \in T} \gamma_j + K \} & \text{if } i \in U \\ E_i^L &= \gamma_i + \beta_i \{ (1 - \theta^*) \tau^* \sum_{j \in T} \gamma_j + K \} & \text{if } i \in L \end{cases} \quad (1)$$

3 Game theoretic analysis

Let us assume the following generally. First, every inhabitant is rational and knows all the exogenous parameters: $\beta_i, \gamma_i, \epsilon_k^i, n$ and K where $1 \leq i \leq n$ and $1 \leq k \leq n - 1$, as common knowledge. Secondly, as to the range of the exogenous values we assume the following: $0 < \beta_i < 1$, $0 < \gamma_i$, $0 < \epsilon_k^i \leq 1$, $0 \leq K$ and $3 \leq n$.

A few remarks may be called for the last condition (the first condition implies — as stated in the last section — that no one voluntarily contributes to the accumulation of the public capital). The condition is a necessary condition for the emergence of a state. If $n = 2$ and a state is made, the only tax payer's income is — whether he or she honestly pays tax or not — is smaller than it would be if he or she were an outsider. Hence, even if both inhabitants agree to make a state at the first stage (Subgame 1), they both will say $\tau_i = 0$ at the next stage (Subgame 2).

Not being involved in the general solution of the model, we should only like to refer to a simple example in this paper. To put it concretely, we assume the following: $\epsilon_{m-k}^i = \epsilon$ for all $1 \leq i \leq n$ and $1 \leq k \leq n - 1$; $\gamma_i < \gamma_j$ for all $1 \leq i < j \leq n$; $\beta_i = \beta$. The first condition, which mainly relates to the dynamics, have little to do with the analysis of each game (each individual, who lives only for a period, do not take account for the effects of their behaviour on the next period). The second condition is not very restrictive either: since $1 \leq i \leq j \leq n$ can be assumed generally, it only implies that every inhabitant has different productivity. The third condition is the essential one that makes analysis simple.

Let us solve the model on the above-mentioned conditions by the backward induction, which is the usual procedure of the analysis of a dynamic game.¹

Subgame 4

At this stage, e, θ^*, τ^* and M are given. Inhabitant i ($i \in M$) will pay taxes honestly if and only if

$$(1 - \tau^*) \gamma_i + \beta \{ (1 - \theta^*) \tau^* \sum_{j \in T_i} \gamma_j + K \} \leq (1 - \epsilon) \gamma_i + \beta \{ (1 - \theta^*) \tau^* \sum_{j \in T_i} \gamma_j + K \}. \quad (2)$$

Here T_i , which represents the set of taxpayers who Inhabitant i expects will pay taxes honestly, can readily be removed from both side:

$$\theta^* \leq \hat{\theta} = 1 - \frac{1}{\beta} \left(1 - \frac{\epsilon}{\tau^*} \right). \quad (3)$$

¹Our model differs from that of Okada and Sakakibara (1990) in the following points. First, γ_i are all fixed to be unity in Okada and Sakakibara (1990) while they can differ from one another in our model. Secondly, the enforcer confiscate tax evaders' income in our model. This makes no difference in the example of the text, but it makes it possible that a rational inhabitant becomes a citizen to be a tax evader if β_i differ from one another (this is not allowed in Okada and Sakakibara (1990) either); see Concluding remarks.

Hence this condition is independent of i , either all taxpayers pay taxes honestly ($M = T$) or no taxpayers pay taxes ($M = U$). In either case E_i is uniquely determined for all i ($i \in N$):

$$E_i \begin{cases} = \theta^* \tau^* \sum_{j \in T} \gamma_j & \text{for } i = e \\ = (1 - \tau^*) \gamma_i + \beta \{ (1 - \theta^*) \tau^* \sum_{j \in T} \gamma_j + K \} & \text{for } i \in M \text{ if (3) is satisfied} \\ = \gamma_i + \beta \{ (1 - \theta^*) \tau^* \sum_{j \in M} \gamma_j + K \} & \text{for } i \in L \end{cases} \quad (4)$$

$$E_i \begin{cases} = \epsilon \sum_{j \in M} \gamma_j & \text{for } i = e \\ = (1 - \epsilon) \gamma_i + \beta K & \text{for } i \in M \text{ if (3) is not satisfied.} \\ = \gamma_i + \beta K & \text{for } i \in L \end{cases} \quad (5)$$

Subgame 3

At this stage τ^* and M are given. In other words e and θ^* must be determined as functions with respect to τ^* and M . We can prove the following:²

$$e = \min_{j \in M} j \quad (6)$$

$$\theta^* = \begin{cases} \hat{\theta} & \text{if } \underline{\theta} < \hat{\theta} < 1, \hat{\theta} < \bar{\theta} \\ \bar{\theta} & \text{if } \max[\underline{\theta}, \bar{\theta}] < \hat{\theta}, \bar{\theta} < 1 \\ 1 & \text{if } \underline{\theta} < \hat{\theta}, 1 < \min[\hat{\theta}, \bar{\theta}] \\ 0 & \text{if } \hat{\theta} < \underline{\theta}, \max_{j \in M} \gamma_j < \epsilon Z \\ 1 & \text{otherwise} \end{cases} \quad (7)$$

where

$$\underline{\theta} = \frac{(1 - \tau) \min_{j \in M} \gamma_j + \beta \tau (\sum_{j \in M} \gamma_j - \min_{j \in S} \gamma_j) + \beta K}{\tau \{ (1 + \beta) \sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j - \beta \min_{j \in S} \gamma_j \}} \quad (8)$$

$$\bar{\theta} = \frac{(1 - \tau) \min_{j \in S} \gamma_j + \beta \tau (\sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j) + \beta K}{\tau \{ (1 + \beta) \sum_{j \in M} \gamma_j - \min_{j \in S} \gamma_j - \beta \min_{j \in M} \gamma_j \}}. \quad (9)$$

Then

$$E_i \begin{cases} = \{ 1 - \frac{1}{\beta} (1 - \frac{\epsilon}{\tau^*}) \} \tau^* \sum_{j \in T} \gamma_j & \text{for } i = e \\ = (1 - \tau^*) \gamma_i + (\tau^* - \epsilon) \sum_{j \in T} \gamma_j + \beta K & \text{for } i \in M \text{ if } \theta^* = \hat{\theta} \\ = \gamma_i + (\tau^* - \epsilon) \sum_{j \in M} \gamma_j + \beta K \} & \text{for } i \in L \end{cases} \quad (10)$$

²Proof is all omitted for the following subgames. It is not difficult but rather involved.

$$E_i \begin{cases} = \bar{\theta} \tau^* \sum_{j \in T} \gamma_j & \text{for } i = e \\ = (1 - \tau^*) \gamma_i + \beta \left\{ (1 - \bar{\theta}) \tau^* \sum_{j \in T} \gamma_j + K \right\} & \text{for } i \in M \quad \text{if } \theta^* = \bar{\theta} \\ = \gamma_i + \beta \left\{ (1 - \bar{\theta}) \tau^* \sum_{j \in M} \gamma_j + K \right\} & \text{for } i \in L \end{cases} \quad (11)$$

$$E_i \begin{cases} = \tau^* \sum_{j \in T} \gamma_j & \text{for } i = e \\ = (1 - \tau^*) \gamma_i + \beta K & \text{for } i \in M \quad \text{if } \theta^* = 1 \\ = \gamma_i + \beta K & \text{for } i \in L \end{cases} \quad (12)$$

$$E_i \begin{cases} = 0 & \text{for } i = e \\ = (1 - \tau^*) \gamma_i + \beta \left\{ \tau^* \sum_{j \in T} \gamma_j + K \right\} & \text{for } i \in M \\ = \gamma_i + \beta \left\{ \tau^* \sum_{j \in M} \gamma_j + K \right\} & \text{for } i \in L \end{cases} \quad \text{if } \theta^* = 0. \quad (13)$$

3.1 Subgame 2

At this stage M is given. In other words τ^* must be determined as functions with respect to M .

Every citizen can offer his or her acceptable tax rate τ_i on his or her non-agricultural income and the minimum τ_i is adopted as the tax rate of the state (everyone can virtually dissolve the state by proposing $\tau_i = 0$). On this assumption we can prove the following:

$$\tau^* = \begin{cases} \min[\max[\tau^1, \tau^2, \tau^3], 1] & \text{if } \theta^* = \hat{\theta} \\ 0 & \text{if } \theta^* = \bar{\theta}, 0 < a \text{ and } 0 < \min_{j \in M} b_j \\ \min[\max[\tau^2, \tau^3, \tau^6, \tau^7, \tau^8], 1] & \text{if } \theta^* = \bar{\theta}, 0 < a \text{ and } \max_{j \in M} b_j < 0 \\ \min[\max[\tau^2, \tau^3, \tau^6], 1] & \text{if } \theta^* = \bar{\theta}, a < 0 \text{ and } 0 < \min_{j \in M} b_j \\ \min[\max[\tau^2, \tau^3, \tau^6, \tau^8], 1] & \text{if } \theta^* = \bar{\theta}, a < 0 \text{ and } \max_{j \in M} b_j < 0 \\ 0 & \text{if } \theta^* = \bar{\theta} \text{ and } \min_{j \in M} b_j < 0 < \max_{j \in M} b_j \\ 0 & \text{if } \theta^* = 0 \text{ or } \theta^* = 1 \end{cases} \quad (14)$$

where

$$\tau^1 \equiv \epsilon. \quad (15)$$

$$\tau^2 \equiv \frac{\beta \min_{j \in M} \gamma_j - \left\{ (1 + \beta) \sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j - \beta \min_{j \in S} \gamma_j \right\} \epsilon + \beta^2 K}{\sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j - \beta \min_{j \in S} \gamma_j}. \quad (16)$$

$$\tau^3 \equiv \frac{\beta \min_{j \in S} \gamma_j - \left\{ (1 + \beta) \sum_{j \in M} \gamma_j - \beta \min_{j \in M} \gamma_j - \min_{j \in S} \gamma_j \right\} \epsilon + \beta^2 K}{\sum_{j \in M} \gamma_j - \beta \min_{j \in M} \gamma_j - \min_{j \in S} \gamma_j} \quad (17)$$

$$\tau^4 \equiv \frac{(\beta + \epsilon) \min_{j \in M} \gamma_j - \epsilon \sum_{j \in M} \gamma_j + \beta^2 K}{(\beta - 1) \left(\sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j \right)} \quad (18)$$

$$\tau^5 \equiv \frac{\epsilon(\sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j) - \beta K}{\sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j - \max_{j \in M} \gamma_j} \quad (19)$$

$$\tau^6 \equiv \frac{\min_{j \in S} \gamma_j + \beta K}{\sum_{j \in M} \gamma_j} \quad (20)$$

$$\tau^7 \equiv \frac{\frac{(\min_{j \in M} \gamma_j + \beta K)\{(1 + \beta) \sum_{j \in M} \gamma_j - \beta \min_{j \in M} \gamma_j - \min_{j \in S} \gamma_j\}}{\sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j} - \min_{j \in S} \gamma_j - \beta K}{\beta(\sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j) - \min_{j \in S} \gamma_j} \quad (21)$$

$$\tau^8 \equiv \frac{-\beta \min_{j \in S} \gamma_j (\sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j) (\min_{j \in S} \gamma_j + \beta K)}{\max_{j \in M} \gamma_j \{ \sum_{j \in M} \gamma_j - \min_{j \in S} \gamma_j + \beta (\sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j) \} + \beta \sum_{j \in M} \gamma_j (\sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j)} \quad (22)$$

$$a \equiv \beta (\sum_{j \in M} \gamma_j - \min_{j \in S} \gamma_j) - \min_{j \in S} \gamma_j \quad (23)$$

$$b_i \equiv \gamma_i \{ \sum_{j \in M} \gamma_j - \min_{j \in S} \gamma_j + \beta (\sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j) \} + \beta \sum_{j \in M} \gamma_j (\sum_{j \in M} \gamma_j - \min_{j \in M} \gamma_j) \quad (24)$$

Then

$$E_i \begin{cases} = \hat{\theta} \tau^* \sum_{j \in T} \gamma_j & \text{for } i = e \\ = (1 - \tau^*) \gamma_i + \beta \{ (1 - \hat{\theta}) \tau^* \sum_{j \in T} \gamma_j + K \} & \text{for } i \in M \\ = \gamma_i + \beta \{ (1 - \hat{\theta}) \tau^* \sum_{j \in M} \gamma_j + K \} & \text{for } i \in L \end{cases} \quad \text{if } \tau^* = \max[\tau^1, \tau^2, \tau^3] \quad (25)$$

$$E_i \begin{cases} = \bar{\theta} \tau^* \sum_{j \in T} \gamma_j & \text{for } i = e \\ = (1 - \tau^*) \gamma_i + \beta \{ (1 - \bar{\theta}) \tau^* \sum_{j \in T} \gamma_j + K \} & \text{for } i \in M \\ = \gamma_i + \beta \{ (1 - \bar{\theta}) \tau^* \sum_{j \in M} \gamma_j + K \} & \text{for } i \in L \end{cases} \quad \text{if (29) is satisfied} \quad (26)$$

$$E_i \begin{cases} = \theta^* \sum_{j \in T} \gamma_j & \text{for } i = e \\ = \beta \{ (1 - \theta^*) \sum_{j \in T} \gamma_j + K \} & \text{for } i \in M \\ = \gamma_i + \beta \{ (1 - \theta^*) \sum_{j \in M} \gamma_j + K \} & \text{for } i \in L \end{cases} \quad \text{if } \tau^* = 1 \text{ and } (\theta^* = \hat{\theta} \text{ or } \theta^* = \bar{\theta}) \quad (27)$$

$$E_i \begin{cases} = 0 & \text{for } i = e \\ = \gamma_i + \beta K & \text{for } i \in M \\ = \gamma_i + \beta K & \text{for } i \in L \end{cases} \quad \text{if } \tau^* = 0 \text{ and } (\theta^* = \hat{\theta} \text{ or } \theta^* = \bar{\theta}) \quad (28)$$

where

$$\tau^* = \max[\tau^2, \tau^3, \tau^6, \tau^7, \tau^8] \text{ or } \tau^* = \max[\tau^2, \tau^3, \tau^6] \text{ or } \tau^* = \max[\tau^2, \tau^3, \tau^6, \tau^8]. \quad (29)$$

Subgame 1

Now we have only to determine M . All the 2^n subsets of N can be M . For each M , (25), (26), (27) and (25) determine all E_i . It is assured that there exists a Nash equilibrium If every individual simultaneously announces whether he or she be a citizen or not, it is assured that there exists a certain Nash Equilibrium: $\{x_1, x_2, \dots, x_n\}$ where X_i represents the probability at which Inhabitant i chooses to be a citizen. Nevertheless it is rather difficult to calculate it even numerically. We shall thus assume that Subgame 1 is not a simultaneous game but a dynamic game, where inhabitants announce whether their decision in a certain order. It is straightforward to calculate a subgame-perfect Nash equilibrium for the expanded form of Subgame 1, which is unique except for fluke, if the value of parameters and the initial capital stock are given.

4 Results of simulations

Let us continue to examine the example of the last section. Having obtained the subgame-perfect Nash equilibrium for the game of each year, we can see how it alters as time passes (or the public capital accumulates because it is assumed that only K is variable while all parameters remain constant through time.

Since it is rather complicated to express the dynamics of $E_i(t)$, we have made simulations, which show some general features. In this section let us examine the following case:

$$n = 7, \beta = 0.9, \gamma_i = 0.6 + 0.4 \frac{i}{n+1}, \epsilon = 0.2, K(0) = 0 \text{ and} \quad (30)$$

$$K(t) = K(t - 1) + (1 - \theta^*(t))\tau^*(t) \sum_{j \in S(t)} \gamma_j$$

As to the order of the announcement in Subgame 1, let us first assume that inhabitants announce whether they become a citizen or an outsider in order of their productivity. That is to say, Inhabitant 7 who has the largest γ_i declares first while Inhabitant 1 who has the smallest γ_i says last.

The left side of Figure 1 shows how E_i changes as time passes.

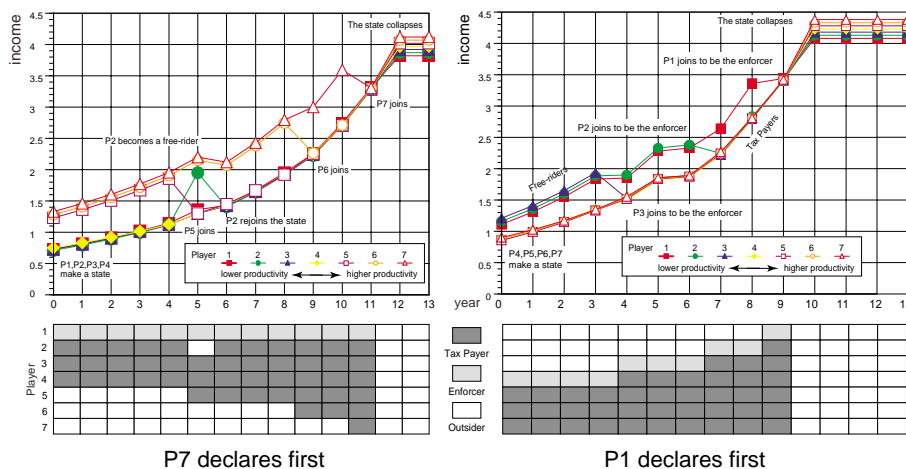


Figure 1: The dynamics of the state

In Year 0 Inhabitants 1, 2, 3 and 4 make a state and Inhabitant 4, who has the smallest productivity among them is elected to be the enforcer. This is because the value of parameters and the initial condition are such that only a state with four individuals can be formed. Knowing this, Inhabitants 5, 6 and 7 announce that they become outsiders, expecting that the remaining four will make a state.

In Year 5 Inhabitant 5 joins the state. This is because the total income of Inhabitants 1, 2, 3 and 4 can earn if no state is made has increased as the result of the accumulation of the public capital. They will not make a state, which does not make their income larger than it would be if no state is formed. Realises that a state is made if and only if he or she declares that he or she becomes a citizen (on the supposition that Inhabitants 1 and 2 become outsiders), Inhabitant 3 compares the income he or she obtains if no state is formed and what he or she gets if he or she joins the state of Inhabitants 1, 2, 3 and 4; he or she finds the latter is greater even though it is smaller than his or her income in the previous year.

Inhabitant 6 joins the state in Year 9 and Inhabitant 7 also become its citizen in Year 11. The reason they join the state is the same as Inhabitant 5 becomes a citizen in Year 5. It would also be obvious why they become a citizen even though their income temporarily decreases.

An exceptional phenomenon is observed when Inhabitant 5 joins the state: Inhabitant 2 leaves the state. This is because (he or she knows) Inhabitants 5, 4, 3 and 1 make a state; though Inhabitants 4, 3 and 1 do not agree to make a state with Inhabitant 1, they agree to make a state with Inhabitant 5 who has higher productivity. Nevertheless soon (actually in the next year) making a four-person state cannot attract any group of four individuals so that Inhabitant 2 becomes a citizen again to make a five-inhabitant state.

Let us also see the case where the order of announcement is reversed (the value of parameters and the initial condition remain unchanged). The right side of Figure 1 shows the dynamics of E_t , which is basically the same.

The dynamics of τ and θ are as follows: Both τ and θ periodically increase. Every time either




Figure 2: The dynamics of tax rate and salary ratio

value approaches to unity, the state increases population to decrease it; the state does not collapse till this trick does not work.

The dynamics of capital stock is shown in figure 3:




Figure 3: The dynamics of capital stock

5 Concluding remarks

Although no citizen evades taxes in the example analysed in this paper, it is not always the case. Certainly if a state is formed whether an inhabitant joins it or not, there is no point for him/her to join the state to evade taxes; the evader's income is less than or equal to the outsider's according as his/her tax evasion is found out or not. Nevertheless a state may be formed only if an inhabitant who will be a tax evader joins it; because it may increase the expected income of the enforcer who can confiscate tax evaders' income. Actually we have checked that this possibility is realised for certain combinations of the value of parameters. Such examples and the results of simulations for them will be discussed in Part Two of this paper.

References

- [1] Okada, A. and Sakakibara, K. 1990. The emergence of the state: A game theoretic approach to the theory of social contract, *Economic Studies Quarterly*, vol.42, no.4

Complex Dynamics of Speculative Price

Taisei Kaizoji
 Division of Social Sciences,
 International Christian University
 Osawa, Mitaka, Tokyo, 181-8585, Japan.
 kaizoji@icu.ac.jp

Abstract

In this paper we think of the security market as consisting of two types of investors: fundamentalists and bandwagon traders, and propose a heterogeneous agent model that represents speculative dynamics by using the Synergetic approach [17]. We show the characteristic patterns of speculative price (speculative bubbles and speculative chaos) which are generated by trading between the fundamentalists and bandwagon traders.

1 Introduction

A recent important development in the nonlinear dynamic theory is the discovery of *deterministic chaos* ([25], [29]). During the last few decades, the concept of deterministic chaos has received a great deal of attention from very diverse scientific fields.

The stock market crash of October 19, 1987 triggered a revolution of chaos in finance. After the 1987 market crash, a number of economists start thinking seriously about the possibility to apply theories on nonlinear dynamics and chaos to finance. Since chaotic dynamics is able to generate large movements which may look like stochastic processes at a first glance, with greater frequency than linear models, the idea that violent fluctuations of speculative prices are generated by some chaotic process, seems to be intuitively a right solution to bubbles and the crashes in the financial markets¹.

On the other hand, recently a number of structural asset pricing models have been introduced, emphasizing the role of heterogeneous beliefs in financial markets, with different groups of traders having different expectations about future prices. Most of these heterogeneous agent models are composed of two typical agent types. The first type is the *fundamentalists* or *arbitrageurs*, who believe that the security price is determined by the market fundamental values. The second type is the *noise traders*, sometimes called *chartists* or *technical analysts*, who may predict the future price using simple technical trading rules, extrapolation of trends and other patterns observed in past prices.

Another distinctive characteristics of the recent heterogeneous agent models has emphasized that heterogeneity in beliefs may lead to market instability and complicated dynamics, including periodic cycles and chaos in financial markets (*e.g.* [5], [6], [7], [9], [10], [11], [12], [13], [15], [16], [19], [26], [27], [28], [33], [38]). Among them, the heterogeneous agent models of Lux ([26], [27], [28]) constitute important examples of recently developed branch of literature on the heterogeneous market hypothesis. He formalized herd behavior or mutual mimetic contagion in speculative markets. A basic feature of the framework adopted by Lux's models is that heterogeneous agents are treated as a statistical ensemble. His mass-statistical formalisation of agents' attitudes and behavior follows a tradition in the so-called 'Synergetics' literature. The concept of synergetics originally developed by Haken [17], and applied to various problems from social sciences by [36], [37], [31]².

Lux [26] used a Synergetic approach to formalize the theory of non-rational bubbles and crashes advanced by [22] who highlights the importance of psychological factors and irrational factors in explaining historical financial crises. Lux [27] explained the so-called *leptokurtosis* of distribution of returns which is a basic stylized fact of both exchange rate and share price time series using the

¹It has been greatly debated, whether price fluctuations in the financial markets is random walk or chaos ([32], [20], [18]).

²Similar statistical approaches have developed to study various problems of social interactions among heterogeneous agents by [1], [2], [3], [4], [23] and [30].

Synergetic model with heterogeneous traders. Lux [28] extended his models in the earlier paper and found chaotic attractors within a broad range of parameter values. He also showed that the distributions of returns derived from chaotic trajectories of the model exhibit high peaks around the mean as well as leptokurtosis and become less leptokurtotic under time aggregation.

This paper presents a heterogeneous agent model of speculative dynamics which is based upon the Synergetic approach to finance. The basic structure of our model is similar to those of Lux models ([26],[27],and [28]). The main difference between Lux models and our model is as follows: in the Lux models the speculative dynamics is represented by the stochastic *differential* equations. In our model the speculative dynamics, by contrast, is represented by the stochastic *difference* equations.

We think of the security market as consisting of two types of traders: *fundamentalists* whose demand (or supply) is based on prices relative to the fundamental value, the so-called fundamentals prices, and *bandwagon traders* whose demand (or supply) is based on *positive feedback trading* strategies. Namely, the bandwagon traders involve buying the security when the price rises, and selling the security when the price falls³. Therefore, the dynamical properties of the speculative price depend both on the nature of positive feedback trading by bandwagon traders and *arbitrage* by fundamentalists. Following [21], we introduce the stochastic transition between the seller and the buyer. Under the circumstance that one cannot get information on the expectations formations and decision-making of all the traders, a probabilistic setting may be one of best means to formalize the behavior of a large number of heterogeneous traders. We will show that the characteristic patterns of speculative price (*speculative bubbles* and *speculative chaos*) might be generated from the Synergetic model with fundamentalists and bandwagon traders. We investigate three cases: (i) the case of the fundamentalists that the only fundamentalists exist in the security market, (ii) the case of the bandwagon traders that the only bandwagon traders exist in the security market, and (iii) the case of the coexistence that both the typical trader types participate in exchanges. To sum up the major results, we show that (i) arbitrage by the fundamentalists tends to stabilize the price and tends to converge its price into the fundamental price, and that (ii) the positive-feedback trading by the bandwagon traders tends to cause the instability of speculative dynamics, and particularly positive feedback trading reinforced by the bandwagon effect creates bubble-like price patterns and chaos, and that (iii) the combination of the arbitrage by the fundamentalists and the positive-feedback trading by the bandwagon traders generate various patterns of speculative price including speculative chaos, and large slowly decaying swings away from fundamentals price.

Section 2 describes the trading strategies of the fundamentalists and the bandwagon traders, and the transition probabilities of the investment attitude. Section 3 analyzes the characteristics of the speculative dynamics. Section 4 gives a brief summary and a remark on limitations of our model.

2 The Model

We think of a security market where many traders participate in trading. Traders are indexed $j = 1, 2, \dots, n$. x_{jt} denotes the investment attitude of trader j . The investment attitude x_{jt} is defined as follows: if trader j is the buyer of the security at period t , then $x_{jt} = +1$. If trader j , in contrast, is the seller of the security at period t , then $x_{jt} = -1$. There is a market-maker, such as the specialists in the New York Stock Exchange, and he/she compares the buying and selling orders by traders, and executes trading. If the aggregate demand for the security at period t exceeds the aggregate supply of the security at period t , then the market-maker raises the price of the security at period t , and vice versa. Hence, an adjustment process of the security price can be described as follows,

$$P_t - P_{t-1} = \lambda f(x_{1t}, x_{2t}, \dots, x_{nt}), \quad (1)$$

where $f(\cdot)$ denotes the excess demand function for the security and depends on all the traders' investment attitudes, and λ denotes the price adjustment speed determined by the market maker.

³For implications of positive feedback trading see [34].

The price change at period t , $(P_t - P_{t-1})$ becomes plus (minus) if the security market is in excess demand (excess supply) at period t .

Here, we assume that volume of trading per trader is the same quantity. Then the excess demand $f(\cdot)$ depends on the mean value of the investment attitudes x_{jt}

$$X_t = \sum_{j=1}^n \frac{x_{jt}}{n}. \quad (2)$$

If X_t is 0, then there exists the same number of the buyers or the sellers. The situation with $X_t > (<)0$ exhibit that more than half the number of the traders are the buyers (the sellers). In the extreme cases, $X_t = +1$ or $X_t = -1$ all the traders are the buyers or the sellers. For convenience of analysis, we assume that the excess demand function (2) is specified by the following linear function with respect to X_t :

$$P_t = P_{t-1} + \lambda\delta X_t \quad (3)$$

where δ denotes the trading volume per trader.

2.1 Fundamentalists and bandwagon traders

As mentioned in the introduction, the security market is composed of two groups of traders having the different trading strategies; that is, fundamentalists and bandwagon traders. The trading strategies of fundamentalists are described as follows: if the security price is below the fundamentals price, then they try to buy the security until the price is equal to the fundamentals price because they think that the security is undervalued. In contrast, if the security price is above the fundamentals price, then they try to sell the security until the price is equal to the fundamentals price because they think that the security is overvalued. To sum up, the fundamentalists' strategies are described as follows:

1. When $P^* - P_t > 0$, the fundamentalists become the buyer of the security at period $t + 1$.
2. When $P^* - P_t < 0$, the fundamentalists become the seller of the security at period $t + 1$.
3. When $P^* - P_t = 0$, the fundamentalists do not trade the security at period $t + 1$.

where P^* denotes the fundamentals price.

On the other hand, the trading strategies of the bandwagon traders are described as follows: the bandwagon traders try to buy the security after the price rises, and sell the security after the price falls, that is, they follow positive feedback strategies. Thus, the bandwagon traders' strategies are summarized as follows:

1. When $P_t - P_{t-1} > 0$, the bandwagon traders become the buyer of the security at period $t + 1$.
2. When $P_t - P_{t-1} < 0$, the bandwagon traders become the seller of the security at period $t + 1$.
3. When $P_t - P_{t-1} = 0$, the bandwagon traders do not trade the security at period $t + 1$.

As mentioned above, if the majority of the traders are the seller ($X_t > 0$), then the price rises ($P_t - P_{t-1} > 0$), and if the majority of the traders are the buyers ($X_t < 0$), then the market price falls ($P_t - P_{t-1} < 0$). Thus, it is possible that the bandwagon traders' strategies are rewritten as follows:

1. When $X_t > 0$, the bandwagon traders become the buyer of the security at period $t + 1$.
2. When $X_t < 0$, the bandwagon traders become the seller of the security at period $t + 1$.
3. When $X_t = 0$, the bandwagon traders do not trade the security at period $t + 1$.

2.2 The transition probability of the investment attitudes

In this subsection we formalize the transition probabilities of the investment attitudes. It is assumed that the transition probabilities of the investment attitude depends on $(P_t - P_{t-1})$ and $(P^* - P_t)$, because the fundamentalists' investment attitude depends upon the deviation of the price from the fundamentals price, $(P^* - P_t)$, and the bandwagon traders' investment attitude depends upon the change of the market price, $(P_t - P_{t-1})$.

We define the transition probability that a trader changes from the seller to the buyer as $g_{\uparrow}[(P_t - P_{t-1}), (P^* - P_t)]$, and, the transition probability that a trader changes from the buyer to the seller as $g_{\downarrow}[(P_t - P_{t-1}), (P^* - P_t)]$.

Furthermore, the transition probabilities are specified by the equations (4) and (5):

$$g_{\uparrow}[(P_t - P_{t-1}), (P^* - P_t)] = \nu(\mu + \exp[\alpha(P_t - P_{t-1}) + \beta(P^* - P_t)]), \quad (4)$$

$$g_{\downarrow}[(P_t - P_{t-1}), (P^* - P_t)] = \nu(\mu + \exp[-(\alpha(P_t - P_{t-1}) + \beta(P^* - P_t))]), \quad (5)$$

where the parameters, α and β are positive. α denotes the strength of the bandwagon traders' reaction upon the price changes, the so-called *bandwagon effect*, and β denotes the strength of the fundamentalists' reaction upon differences between the actual market price and the fundamental price. We call α the *bandwagon coefficient*, and β the *arbitrage coefficient*. These transition probabilities imply the following:

1. If $P_t - P_{t-1} > (<)0$, then $g_{\uparrow}[(P_t - P_{t-1}), (P^* - P_t)]$ increases (decreases), and simultaneously $g_{\downarrow}[(P_t - P_{t-1}), (P^* - P_t)]$ decreases (increases).
2. If $(P^* - P_t) > (<)0$, then $g_{\uparrow}[(P_t - P_{t-1}), (P^* - P_t)]$ decreases (increases), and simultaneously $g_{\downarrow}[(P_t - P_{t-1}), (P^* - P_t)]$ increases (decreases).

3 Speculative dynamics

With the transition probabilities (4) and (5), the time development of the mean values of the investment attitude X_t and the price $P(t)$ becomes the equations (6) and (7):

$$\langle P_t \rangle = \langle P_{t-1} \rangle + \lambda \delta \langle X_t \rangle, \quad (6)$$

$$\begin{aligned} \langle X_{t+1} \rangle &= \langle X_t \rangle + \xi[(1 - \langle X_t \rangle)g_{\uparrow}[(P_t - P_{t-1}), (P^* - P_t)] \\ &\quad - (1 + \langle X_t \rangle)g_{\downarrow}[(P_t - P_{t-1}), (P^* - P_t)]] \end{aligned} \quad (7)$$

where $\langle P_t \rangle$ and $\langle X_t \rangle$ denote the ensemble mean values of X_t and P_t . The equation (7) can be derived from the original stochastic system using the Master equation. On details of this derivation see Weidlich and Haag[36].

Now we have a dynamical system that is formed by the adjustment process of the price (6) and the dynamics of the investment attitudes (7). We investigate the dynamical properties of the system (6) and (7) below. We consider three typical cases: the case of the fundamentalists that the only fundamentalists exist in the security market, the case of the bandwagon traders that the only bandwagon traders exist in the market and the case that the two typical trader types coexist in the security market.

3.1 The case of the fundamentalists : $\alpha = 0, \beta > 0$

First, we consider the case that all the traders are the fundamentalist. In this case the dynamical system (6) and (7) rewritten by the following

$$\langle P_t \rangle = \langle P_{t-1} \rangle + \lambda \delta \langle X_t \rangle, \quad (8)$$

$$\begin{aligned}\langle X_{t+1} \rangle &= \langle X_t \rangle + \xi[(1 - \langle X_t \rangle)\nu(\mu + \exp[\beta(P^* - P_t)]) \\ &- (1 + \langle X_t \rangle)\nu(\mu + \exp[-\beta(P^* - P_t)])]\end{aligned}\quad (9)$$

For simplicity of analysis we specify the parameters as follows: $\xi\nu = 0.2$, $P^* = 100$, and $\mu = 0.5$.

The dynamical system has a unique fixed point, $(\langle P^* \rangle, \langle X^* \rangle) = (100, 0)$. We call the fixed point the *fundamental equilibrium*. We can demonstrate the following stability condition:

If the arbitrage coefficient $\beta < 3/\delta\lambda$, then the desired market equilibrium is locally stable.

The proof is trivial and omitted.

Therefore in the case of the fundamentalists the price is stabilized by the fundamentalists' arbitrage and converge to the fundamental equilibrium, provided that the arbitrage coefficient β is selected suitably.

3.2 The case of the bandwagon traders : $\alpha > 0$ and $\beta = 0$

Next, we consider the case that all the traders are the bandwagon traders. Substituting the equation (6) to the equation (7), the dynamical system is rewritten by the following equations

$$\langle P_t \rangle = \langle P_{t-1} \rangle + \lambda\delta\langle X_t \rangle, \quad (10)$$

$$\begin{aligned}\langle X_{t+1} \rangle &= \langle X_t \rangle + \xi\nu[(1 - \langle X_t \rangle)(\mu + \exp[\alpha\lambda\delta\langle X_t \rangle]) \\ &- (1 + \langle X_t \rangle)(\mu + \exp[-\alpha\lambda\delta\langle X_t \rangle])]\end{aligned}\quad (11)$$

Thus the map of $\langle X_t \rangle$, (11) is one-dimensional. We can demonstrate the following stability condition:

If the bandwagon coefficient $\alpha < 3.5/\delta\lambda$, then the origin is the unique equilibrium and locally stable.



The proof is trivial and omitted.

The parameters are specified as follows: $\delta\lambda = 0.5$, $\xi\nu = 0.2$, $P^* = 100$, and $\mu = 0.5$.

Figure 1, Figure 2 and Figure 3 illustrate the map (11) with the different values of the bandwagon coefficient α . These figures show that (i) for $\alpha < 7$ the origin is a unique equilibrium and is stable (Figure 1), and that (ii) for $\alpha = 7$ a *pitchfork bifurcation* at the origin occurs (Figure 2), and that (iii) for $\alpha > 7$ the origin becomes unstable, and simultaneously the two new equilibria, the *bull market equilibrium* and the *bear market equilibrium* are created, one above and one below the origin (Figure 3). Figure 4 is the bifurcation diagram for the map (11) where the bandwagon coefficient α varies smoothly from 7 to 8.2. Note that Figure 4 is created for some positive initial values of the investment attitude index, ($0 < \langle X_0 \rangle < 1$). This figure suggests the following bifurcation scenario with respect to the bandwagon coefficient α . If the bandwagon coefficient α is small, then the bull market equilibrium is stable for any positive initial values. If α is increased, then the bull market equilibrium becomes unstable and *period doubling bifurcations* occur. After infinitely many period doubling bifurcations the dynamics becomes chaotic. When α is further increased, the *symmetry-breaking bifurcation* [8] occurs, and the sudden increase in symmetry of the chaotic attractor.

Figure 5 and Figure 6 illustrate the time paths of $\langle X_t \rangle$ and $\langle P_t \rangle$ with $\alpha = 7.8$. The mean value of the investment attitude $\langle X_t \rangle$ fluctuates within the range of the bull market, [$0 < \langle X_t \rangle < 1$], so that the mean value of the price $\langle P_t \rangle$ keeps rising over time. Since the rise in the mean value of the price is caused by the bandwagon effect that are not justified by fundamentals, it seems reasonable to suppose that *speculative bubbles* occur in the security market. Figure 7 and Figure 8 illustrate the time paths of $\langle X_t \rangle$ and $\langle P_t \rangle$ with $\alpha = 8.06$ after the symmetry-breaking bifurcation. These figures show that the mean value of the investment attitude $\langle X_t \rangle$ fluctuates chaotically in the broad range of [$-1 < \langle X_t \rangle < 1$], and various rise and fall patterns of the mean value of the price $\langle P_t \rangle$ are created by the positive feedback trading of the bandwagon traders.

It follows from the numerical analysis that the existence of the bandwagon traders tends to destabilize the price, and positive feedback trading reinforced by the bandwagon effect give cause to bubble-like price patterns. When the bandwagon effect is further strong, the instability of the speculative dynamics is amplified, so that chaos of speculative price is caused.

Figure 1: The map (11) with $\alpha = 6$.Figure 2: The map (11) with $\alpha = 7$.

3.3 The case of coexistence of fundamentalists and bandwagon traders : $\alpha, \beta > 0$

Finally, we consider the case that both typical trader types exist in the security market. In this case the dynamical system is described by the original dynamical system (6) and (7). We specify the parameters as follows: $\delta\lambda = 0.5, \xi\nu = 0.2, P^* = 100$, and $\mu = 0.5$.

We observed that in the case of the bandwagon traders, the route to chaos passes through a cascade of period doublings. In the case of coexistence of both trader types we will see the existence of quasi-periodic transitions to chaos. First of all we consider the stability conditions for the fundamentals equilibrium $(\langle P^* \rangle, \langle X^* \rangle) = (100, 0)$. We can demonstrate the following stability conditions:



If $\alpha < 7$ and $\alpha - 0.5\beta < 2$, then the fundamental equilibrium is locally stable.

The proof is trivial and omitted. We see from the above stability conditions that if the bandwagon coefficient α is small, then the speculative dynamics is stable and the mean value of the price is converged into the fundamentals equilibrium by the arbitrage of fundamentalists. In other words, if the bandwagon effect is weak in the market, then the speculative dynamics is stabilized by the arbitrage of the fundamentalists.

To understand the global characteristics of the speculative dynamics in detail, we consider two cases: the cases with a large arbitrage coefficient and a small arbitrage coefficient.

First, we investigate the case with the large arbitrage coefficient ($\beta = 1$). In this case if the bandwagon coefficient α is above 7 in this case, then a Hopf bifurcation occurs at the fundamental equilibrium, and then a quasi-periodic orbit starting from the market fundamental equilibrium appears. Figure 9 and Figure 10 show two attractors in the $(\langle P_{t+1} \rangle, \langle P_t \rangle)$ plane and $(\langle X_t \rangle, \langle P_t \rangle)$ plane with $(\alpha, \beta) = (7, 1)$. In both the figures, the orbits converges to attracting invariant 'circle' created in the Hopf bifurcation. When the bandwagon coefficient α is increased from 7 to 7.4, the invariant circles break up the into strange attractors (Figure 11 and Figure 12). It follows from these figures that the transition occurs from quasiperiodicity to chaos when α is increased from 7 to 7.4, under $\beta = 1$. Figure 13 and Figure 14 show the time paths of $\langle X_t \rangle$ and $\langle P_t \rangle$ with $(\alpha, \beta) = (7.4, 1)$. The time series of the mean value of investment attitude $\langle X_t \rangle$ fluctuate irregularly within the broad range $[-1 < \langle X_t \rangle < 1]$. The corresponding time series of $\langle P_t \rangle$ is also chaotic.

Second, we investigate the case with the small arbitrage coefficient ($\beta = 0.01$). If the bandwagon coefficient α is above 7 in the case with $\beta = 0.1$, then a Hopf bifurcation occurs at the fundamental

Figure 3: The map (11) with $\alpha = 8$.Figure 4: The bifurcation diagram for the map (11) with respect to α .

equilibrium, and the orbits converges to attracting invariant ‘square’ created in the Hopf bifurcation (Figure 15 and Figure 16). As the bandwagon coefficient α is further increased, the invariant squares break up the into strange attractors (Figure 17 and Figure 18). Figure 19 and Figure 20 show the time series of the time paths of $\langle X_t \rangle$ and $\langle P_t \rangle$ with $(\alpha, \beta) = (7.4, 0.01)$. The mean value of the investment attitude $\langle X_t \rangle$ fluctuates irregularly within the range of $[0 < \langle X_t \rangle < 1]$ and suddenly falls into the range of $[-1 < \langle X(t) < 0]$, and then fluctuates erratically within the range, and again suddenly jumps up the range of $[0 < \langle X(t) < 1]$, and the same process is repeated. The corresponding time series of the mean value of the price repeats large slowly decaying swings away from fundamental price.

4 Conclusion

This paper has represented a Synergetic model that stresses the role of irrational sentiment of heterogeneous traders. In our model, price fluctuations are caused by an endogenous mechanism relating the fraction of the fundamentalists and the bandwagon traders to the strength of the arbitrage by the fundamentalists, and that of positive-feedback trading by the bandwagon traders. The important points as regards the endogenous mechanism that generates speculative dynamics are that (i) a large fraction of the fundamentalists, or the increasing strength of the fundamentalists’ reaction upon difference between actual price and the fundamentals price, tends to stabilize the speculative price, and in contrast, (ii) a large fraction of the bandwagon traders, or the increasing strength of the bandwagon traders’ reaction upon difference between the price at the present date and the price at the previous date, tends to destabilize the security price, and (iii) trading between the fundamentalists and the bandwagon traders generates various patterns of speculative dynamics including speculative bubble, speculative chaos, and large slowly decaying swings away from fundamental price.

These results seem to provide a useful analytic foundation to experiments on stock market behavior⁴.

⁴In the last decade a growing literature on experimental asset markets has emerged. A nice survey of this literature is given in [35].





Figure 5: **The time path of $\langle X_t \rangle$ with $\alpha = 7.8$.** Figure 6: **The time path of $\langle P_t \rangle$ with $\alpha = 7.8$: Speculative Bubble.**

However, the characterization of the traders is oversimplified and stylized in our model. We employ implicitly the extreme assumptions on the behavior of the traders: (i) the fundamentalists can know the exact value of the fundamentals price, and (ii) the bandwagon traders decide their investment attitude on the basis of the price change from the preceding period to the current period. The latter assumption implies that the bandwagon traders use only the data of the price at the present period and the preceding period in order to forecast the future price. As the result our model of speculative dynamics are formalized by two dimensional difference equations.

On the other hand standard finance models based on the efficient market hypothesis ([14]) assume that fundamentalists use all information available to them at present in order to perceive the fundamentals price. Similarly, in most of the heterogeneous agent models typical noise traders such as *chartists* or *technical analysts* are assumed to find price trends and other patterns observed in past prices from the long-term data of the prices, and then predict the future price using their technical trading rules. Therefore it is reasonable to suppose that the trading strategies of the traders will, at least, depend upon the long-run data of the prices. Whereas we recognize the importance of the problem on the time horizon in our model, it seems to us that the essential nature of speculative dynamics remains unchanged in more general frameworks. We leave it for future work to see whether this conjecture is true.

5 Acknowledgements

I would like to thank Yoko Yamaguchi, Nobuharu Miyatake, Thomas Lux for helpful comments on an earlier draft of this paper. Financial support by Japan productivity center for socio-economic development is gratefully acknowledged.

References

- [1] Aoki, M., 1994, New Macroeconomic Modeling Approaches: Hierarchical Dynamics and Mean




Figure 7: **The time path of $\langle X_t \rangle$ with $\alpha = 8.06$.**




Figure 8: **The time path of $\langle P_t \rangle$ with $\alpha = 8.06$: Speculative Chaos.**

Field Approximations, *Journal of Economic Dynamics and Control* 18, 865-877.

- [2] Aoki, M. 1996, *New Approaches to Macroeconomic Modeling: Evolutionary Stochastic Dynamics, Multiple Equilibria, and Externalities as Field Effect*, Cambridge University Press, New York.
- [3] Aoki, M. 1998, A Stochastic Model of Prices and Volumes in a Share Market with Two Types of Participants, mimeo.
- [4] Brock, W.A. 1993, Pathways to randomness in the economy: Emergent Nonlinearity and chaos in economics and finance, *Estudios Economicos* 8.
- [5] Brock, W. A., and Hommes, C.H., 1997, Models of complexity in economics and finance, In Heij, C., Schumacher, J.M., Hanzon, B., and Praagman, C. (Eds.) *System Dynamics in Economic and Financial Models*. Wiley, New York, Chapter 1, pp. 3-44.
- [6] Brock, W. A., and Hommes, C.H., 1997, Heterogeneous beliefs and routes to chaos in a simple asset pricing model, *Journal of Economic Dynamics and Control*, 22, 1235-1274.
- [7] Chiarella, C., 1992, The dynamics of speculative behavior, *Annals of operations research*, 37, 101-123.
- [8] Chossat, P. and Golubitsky, M. 1988, Symmetry-increasing bifurcation of chaotic attractors, *Physica D* 32, 423-436.
- [9] Dacorogena, M.M., Muller, U.A., Jost, C., Pictet, O.V., Olsen, R.B. and Ward, J.R., 1995, Heterogeneous real-time trading strategies in the foreign exchange market, *European Journal of Finance*, 1, 383-403.
- [10] Day, R.H. and Huang, W., 1990, Bulls, bears and market sheep, *Journal of Economic Behavior and Organization*, 14, 299-329.
- [11] DeGrauwe, P., DeWachter, H. and Embrechts, M., 1993, *Exchange rate theory: Chaotic models of foreign exchange markets*, Blackwell, Oxford.




Figure 9: **The quasi-periodic attractor with $(\alpha, \beta) = (7, 1)$.**




Figure 10: **The quasi-periodic attractor with $(\alpha, \beta) = (7, 1)$.**

- [12] DeLong, J.B., Shleifer, A., Summers, L.H. and Waldmann, R.J., 1990, Noise trader risk in financial market, *Journal of Political Economy*, 98, 703-738.
- [13] DeLong, J.B., Shleifer, A., Summers, L.H. and Waldmann, R.J., 1991, The survival of noise traders in financial markets, *Journal of Business* 64, 1-19.
- [14] Fama, E.F., 1970, Efficient capital market: A review of theory and empirical work, *Journal of Finance* 25, 383-417.
- [15] Frankel, J.A. and Froot, K.A., 1988, Chartists, fundamentalists and the demand for dollars, *Greek Economic Review* 10, 49-102.
- [16] Haltiwanger, J. and Waldmann, M., 1985, Rational expectations and the limits of rationality: an analysis of heterogeneity, *American Economic Review* 75, 326-340.
- [17] Haken, H., *Synergetics, An Introduction*, 1983, Springer-Verlag.
- [18] Hsieh, D., 1991, Chaos and nonlinear dynamics: Applications to financial markets, *Journal of Finance*, 46, 1839-1878.
- [19] Huang, W. and Day, R.H. 1993 Chaotically switching bear and bull markets: The derivation of stock price distributions from behavioral rules, in Day, R.H., and Chen, P. (eds.), *Nonlinear dynamics and evolutionary economics*, University Press, Oxford.
- [20] Jaditz, T., and Sayers, C., 1993, Is chaos generic in economic data ?, *International journal of bifurcation and chaos* 3-2, 745-755.
- [21] Kaizoji, T., 1998, A theory of speculative bubble, chaos and crash, *Simulation*, 17-2, 141-152.
- [22] Kindlberger, C.P., 1989, *Manias, Panics, and Crashes: A History of Financial Crises*, MacMillan, London.
- [23] Kirman, A., 1991, Epidemics of opinion and speculative bubbles in financial markets, in M. Taylor (ed.), *Money and financial market*, Macmillan, London.




Figure 11: The strange attractor with $(\alpha, \beta) = (7.5, 1)$. Figure 12: The strange attractor with $(\alpha, \beta) = (7.5, 1)$.

- [24] Kirman, A., 1993, Ants, Rationality, and Recruitment, *Quarterly Journal of Economics*, 108, 137-156.
- [25] Lorenz, E.N., 1963, Deterministic nonperiodic flow, *Journal of Atmospheric Sciences* 20, 130-141.
- [26] Lux, T., 1995, Herd Behavior, Bubbles, and Crashes, *The Economic Journal*, 105, 881-896.
- [27] Lux, T., 1997, Time Variation of Second Moments from a Noise/Infection Model, *Journal of Economic Dynamics and Control*, 22, 1-38.
- [28] Lux, T., The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions, 1998, *Journal of Economic Behavior and Organization*, 33, 143-165.
- [29] May, R.M., 1976, Simple mathematical models with very complicated dynamics, *Nature*, 261, 459-467.
- [30] Ramsey, J.B., 1996, On the existence of macro variables and macro relationships, *Journal of Economic Behavior and Organization* 30, 275-299.
- [31] Reiner, R., Munz, M., and Weidlich, W. 1988, Migratory dynamics of interacting subpopulations: regular and chaotic behavior, *System Dynamics Review*, 4, 1-2, 179-199.
- [32] Scheinkman, J. and LeBaron, B., 1989, Nonlinear dynamic and stock returns, *Journal of Business*, 62, 311-338.
- [33] Sethi, R., 1996 Endogenous regime switching in speculative markets, *Structural Change and Economic Dynamics*, 7, 99-118.
- [34] Shleifer, A., and Summers, L. H., 1990, The noise Trader approach to finance, *Journal of Economic Perspectives*, 4, 2, 19-33.




Figure 13: **The time path of $\langle X_t \rangle$ with $(\alpha, \beta) = (7.4, 1)$.**




Figure 14: **The time path of $\langle P_t \rangle$ with $(\alpha, \beta) = (7.4, 1)$.**

- [35] Sunder, S., 1995, Experimental asset markets: a survey, in Kagel, J. and A. Roth (eds.), *Handbook of experimental economics* chapter 6, Princeton University Press.
- [36] Weidlich, W. and G.Haag, *Concepts and Models of a Quantitative Sociology*, 1983, Springer-Verlag.
- [37] Weidlich, W., 1991, Physics and Science: The Approach of Synergetics, *Physics Reports* 204, 1-163.
- [38] Zeeman, E.C., 1974, The unstable behavior of stock exchange, *Journal of Mathematical Economics* 1, 39-49.




Figure 15: The quasi-periodic attractor with $(\alpha, \beta) = (7, 0.01)$.




Figure 16: The quasi-periodic attractor with $(\alpha, \beta) = (7, 0.01)$.




Figure 17: The strange attractor with $(\alpha, \beta) = (7.4, 0.01)$.




Figure 18: The strange attractor with $(\alpha, \beta) = (7.4, 0.01)$.




Figure 19: The time path of $\langle X_t \rangle$ with $(\alpha, \beta) = (7.4, 0.01)$.




Figure 20: The time path of $\langle P_t \rangle$ with $(\alpha, \beta) = (7.4, 0.01)$.

The Nonlinear Dynamics of Debt Deflation

Steve Keen

Department of Economics and Finance
 The University of Western Sydney Macarthur,
 Campbelltown, 2560
 S.Keen@uws.edu.au
<http://btwebsh.macarthur.uws.edu.au>

1 Introduction

After an era of moderate to high inflation lasting almost 30 years, the USA has recently begun to record falling consumer and intermediate goods prices. Prices in Japan have been static to falling (except for the immediate impact of the Asia crisis) since 1994, while several European countries have near zero rates of inflation.

Contemporaneously, several major East and Southeast Asian nations have experienced debt-induced economic collapses, and their largely foreign currency denominated private sector debts were subsequently dramatically amplified by market-driven currency devaluations. The “Asian crisis” will have international repercussions, both through a fall in their imports and the flow-on effects of their attempts to export their way out of their difficulties. The deflationary tendencies already evident in the West, and in particular the USA, could thus be amplified.

The prospect of a sustained period of falling prices thus appears likely in the West. While inflation is more likely than deflation in Asia due to the impact of the currency collapses, these devaluations will exacerbate the problems of excessive private sector debt. As has been most obvious in the case of Indonesia, the market value of a currency is largely determined by speculator expectations of the country’s ability to service its foreign currency debts.

These developments have led many observers to worry that a debt-deflationary process may have commenced in Asia, which may be partially transmitted to the West via aggressive exporting and consequent import price deflation. Significant public figures are also voicing the concern that policy experience gained during an era of inflation may be inappropriate during one of deflation. As Alan Greenspan commented, “deflation can be detrimental for reasons that go beyond those that are also associated with inflation”[4]. To have any inkling of what the economic future might have for us, we have to consider economic theories of deflation in the context of private debt.

2 Economic Theory

Compared to the wealth of economic argument about what can be done to control inflation, there is remarkably little economic theory devoted to deflation.¹ The main contributions were made by Fisher during the Great Depression[1], to some extent Keynes[9], and Minsky[11, 12, 13]. Minsky’s “Financial Instability Hypothesis” can be regarded as distilling the essence of these contributions.²

This is that a pure market economy is characterised by a fundamental asymmetry which can cause the debt to output ratio to rise over time, to levels which can be unsustainable. This asymmetry is easily put: firms incur debt to finance investment during booms, but have to repay that debt during slumps. Since the cyclical path of a capitalist economy is itself asymmetrical, this results in the level of debt “ratcheting up” during a sequence of trade cycles. Under fairly general conditions, this

¹Milton Friedman is the economist most identified with the drive to control inflation. What is less well known is that his preferred rate of inflation was actually minus 4 to 5 per cent a year[2]. The paper in which this result was derived completely ignored the existence of debt.

²See [13, especially Chs. 1] and [5, 6, 7, 8] for lengthier expositions.

process can reach a point at which the accumulated debt overwhelms the debt-financing capacity of the economy, thus leading to a Depression.

Conversely, according to Minsky's hypothesis, a mixed market-state economy avoids this tendency towards complete collapse because the counter-cyclical behaviour of government spending — rising during slumps and falling during booms — counters the tendency of the private sector to accumulate excessive debt. Minsky argues that increased taxation during a boom attenuates investment, thus limiting the tendency to over supply productive capacity and consequently borrow excessively during a boom, and that increased government spending during a slump provides additional cash flow to businesses which stops their debt exploding. This argument is in direct opposition to the attitude of conventional macroeconomic analysis to government spending, which has long since replaced the Keynesian advocacy of counter-cyclical government behaviour with an effective demonising of government deficits.

I have previously modelled these theories in the absence of price dynamics[5, 6, 8]. In this paper I extend the model to incorporate the impact of a variable price level, and inflation-dependent rates of interest.

3 The Basic Model

The foundation of this model is Goodwin's model of cyclical growth[3], which was itself based upon the Lotka-Volterra predator-prey model of species interaction on the one hand, and Marx's income-distribution/employment model of the trade cycle on the other. Over one century later, his arcane language notwithstanding, the best expression of this model is still that given by Marx:

a rise in the price of labor resulting from accumulation of capital implies . . . accumulation slackens in consequence of the rise in the price of labour, because the stimulus of gain is blunted. The rate of accumulation lessens; but with its lessening, the primary cause of that lessening vanishes, i.e. the disproportion between capital and exploitable labour power. The mechanism of the process of capitalist production removes the very obstacles that it temporarily creates. The price of labor falls again to a level corresponding with the needs of the self-expansion of capital, whether the level be below, the same as, or above the one which was normal before the rise of wages took place . . . To put it mathematically, the rate of accumulation is the independent, not the dependent variable; the rate of wages the dependent, not the independent variable. [10, pp580–581]

Goodwin showed that this could be modelled as a predator-prey system in which workers share of output played the role of predator, and the rate of employment the role of prey:³

$$\begin{aligned}\frac{d\omega}{dt} &= \omega \cdot (P(\lambda) - \alpha) \\ \frac{d\lambda}{dt} &= \lambda \cdot \left(\frac{1-\omega}{v} - \alpha - \beta \right),\end{aligned}\tag{1}$$

where ω is the wage to output ratio $\frac{W}{Y}$, $P(\lambda)$ is a nonlinear relationship between the rate of change of wages w and the rate of employment (known as the "Phillips curve"),⁴ λ the rate of employment or employment to population ratio $\frac{L}{N}$, α the rate of growth of labour productivity, β the rate of population growth, and v the capital to output ratio $\frac{K}{Y}$.⁵ As is well known, this model generates a stable limit cycle. The model also has an easy verbal explanation. The first equation says that workers'

³In that the linear term in the wages share of output equation is negative and that in the employment equation positive.

⁴See the Glossary for the form of this and other nonlinear functions used in this paper.

⁵I have extended the model to include variable capacity utilisation, so that the capital to output ratio is not a constant, and variable technical change, so that the rate of growth of labour productivity is not a constant[8]. Both these modifications makes the basic 2 sector model unstable, and increase the volatility of the final debt-deflation model.

share of output will grow if their wage demands (which are based on the level of employment) exceed the rate of growth of labour productivity; the second that the level of employment will grow if the rate of economic growth exceeds the sum of population and productivity growth.

4 First extension: Investment and Debt

The first step in extending this model is to replace the linear assumption that capitalists invest all their profits ($1 - \omega$ in the previous model is the profit to output ratio $\pi = \frac{\Pi}{Y}$) with the more realistic assumption that investment is a nonlinear function $k()$ of the rate of profit ($\frac{\Pi_n}{K} = \frac{\pi_n}{v}$), where π_n is profit net of interest payments.⁶ This does not disturb the underlying nature of the model, which still results in a stable limit cycle, but it sets the scene for the introduction of a finance sector.

Finance is introduced into the model by assuming the existence of a banking sector which exists solely to finance capitalist investment. The rate of change of debt in this system is thus simply interest on outstanding debt, plus new investment, minus gross profits:

$$\frac{dD}{dt} = r \times D + I_g - \Pi, \quad (2)$$

where $I_g = k(\pi) \times Y$ represents gross investment (in what follows, depreciation is introduced at the constant rate of γ p.a.). This produces the following three-dimensional system:

$$\begin{aligned} \frac{d\omega}{dt} &= \omega \cdot (P(\lambda) - \alpha) \\ \frac{d\lambda}{dt} &= \lambda \cdot \left(\left(\frac{k(\pi_n)}{v} - \gamma \right) - \alpha - \beta \right) \\ \frac{dd}{dt} &= d \times \left(r - \left(\frac{k(\pi_n)}{v} - \gamma \right) \right) + k(\pi_n) - \pi, \end{aligned} \quad (3)$$

where d is the debt to output ratio $\frac{D}{Y}$ and π_n is the net profit share of output:

$$\pi = 1 - \omega - r \times d \quad (4)$$

As is well known, a three dimensional system introduces the possibility of chaotic behaviour, and this particular model follows the inverse tangent route to chaos first identified by Pomeau and Manneville[14].

5 Perturbation Analysis

As is easily shown, with the functional form chosen for the ‘‘Phillips curve’’, the equilibrium value of employment is:⁷

$$\lambda_e = \frac{\ln(\alpha - C) - A}{B} = 97.12\% \quad (5)$$

There is an equilibrium value for profit share:

$$\pi_e = \frac{\ln(v \times (\gamma + \alpha + \beta) - G) - E}{F} = 16.18\%, \quad (6)$$

⁶The term Π will be retained for gross profit or output minus wages throughout. The term π_n will signify gross profit minus all other outgoings, which at this stage means interest on outstanding debt. In the next section, Π_n will signify gross profit minus interest payments, and taxation minus subsidies.

⁷With the parameter values used in the following simulations, which were derived by a nonlinear regression of Phillips’s original data against the rate of unemployment. The same regression is used later in the price level section.




Figure 1: Wages Share and Employment near Equilibrium




Figure 2: Debt to Output Ratio near Equilibrium

which corresponds to a rate of profit of approximately 5.4% and, given the investment function, an investment share of output of 16.5%. The equilibrium value for the debt to output ratio is:

$$d_e = \frac{k\left(\frac{\pi_n}{v}\right) - \pi_n}{k\left(\frac{\pi_n}{v}\right) - \gamma} = 7.02\% \quad (7)$$

Since the profit share is a linear combination of ω and d (equation (4)), this gives the curious result that at the equilibrium, workers' share of output and "bankers' share" are in direct opposition to each other, whereas "capitalists' share" is constant. The actual expression is:

$$\omega_e(r) = 1 - \pi_n - r \times d_e \quad (8)$$

This is, unremarkably, significantly different to standard economic models of income distribution, which argue that remuneration reflects relative factor productivity and which are not equipped to deal with a return to accumulated debt.

This equilibrium vector is locally stable but globally unstable, a significant echo of Fisher's intuition in 1933 that the market system has an equilibrium which "though stable, is so delicately poised that, after departure from it beyond certain limits, instability ensues" [1, p339].

Even at this basic level, the model contains some important insights into the role of debt in a market economy, and the impact of the rate of interest in a model in which, in contrast to the standard IS-LM model, debt is explicitly accounted for.

Conventional IS-LM analysis argues that an increase in the interest rate will reduce investment (which is portrayed as a monotonically decreasing function of the interest rate, in contrast to this model's argument that the rate of profit determines the level of investment) and thus growth; however any impact on the accumulation of debt is ignored. The final equation in (3) indicates that, when debt is explicitly accounted for, it is possible for debt to overwhelm the system, even though the equilibrium rate of profit significantly exceeds the rate of interest.

The approximate 5 year period of the cycles should also be noted: this is similar to those of the basic 2 dimensional Goodwin model.

Figure 2 shows the time path of the debt to output ratio, which rises in a cyclical fashion initially, but then also tapers towards its equilibrium value.

The phase diagram in Figure 3 and the period interactions shown in Figure 4 give a clear picture of the dynamics in this 3 dimensional system.

The initial conditions of slightly higher than equilibrium debt, workers share of output and employment leads to a downturn, as investment stagnates due to the resulting low rate of profit. The excess




Figure 3: Wages Share, Employment and Debt near Equilibrium




Figure 4: Period Interactions of Wages Share, Employment and Debt near Equilibrium




Figure 5: Wages share and employment far from Equilibrium




Figure 6: Debt to Output far from Equilibrium

of profit over investment leads to debt being reduced, but the downturn eventually leads to falling wage demands, and this leads to a boost in investment well before debt is fully repaid. Debt then rises with rising employment as investment boosts output, only to lead eventually to rising wage demands which cut into profits and once again cut off investment. The cycle then continues, with the system tapering towards a stable equilibrium debt to equity ratio, wages share and rate of employment.

Conversely, as Figures 5 to 7 indicate, at rates of interest which exceed the equilibrium rate of growth, the equilibrium vector is unstable. From a conventional macroeconomic point of view, this system would appear to be stable right up until the final crisis, since conventional macroeconomics dismisses the issue of income distribution as a topic for microeconomic analysis, and ignores the role of debt. This simulation began with all variables .05 below their equilibrium values:

However, the equilibrium analysis of this model would indicate cause for concern, since the secular trend towards decreasing workers' share of output would indicate that debt must still be rising, as is evidenced in Figure 6. Eventually, the level of accumulated debt becomes so high that repayments on outstanding debt eliminate all profit, leading to a collapse in output and hence a Depression:

The phase diagram of this simulation makes the bifurcation in system behaviour as the interest rate rises graphically apparent. What was previously a stable "volcano" shaped phase diagram becomes




Figure 7: Wages share, employment and debt far from Equilibrium




Figure 8: Period Interactions of Wages Share, Employment and Debt far from Equilibrium

an unstable “vortex” in which debt overwhelms the other system variables:

There is a superficially unremarkable explanation for this phenomenon: with a rate of interest higher than the rate of growth, it is not surprising that debt eventually smothers economic growth. However this is a nonlinear system, and with initial conditions at a further remove from the equilibrium, it is possible for the model to undergo a debt blowout even when the rate of interest is significantly lower than the equilibrium rate of growth. Figure 8 shows the behaviour of the model with a 3 per cent rate of interest when wages share is initially .1 below, employment .05 below, and debt .05 above the equilibrium vector of:

$$\begin{bmatrix} .83184 \\ .97123 \\ 21057 \end{bmatrix} .$$

The behaviour of this model thus clearly supports the Fisher-Keynes-Minsky contention that a pure market economy is fundamentally unstable, in that it is prone to fall into a debt-induced Depression from which there is no escape, bar “resetting the debt clock” via wholesale bankruptcy and debt moratoria. The next extension similarly supports Minsky’s claim that the government sector’s behaviour provides a homeostatic balance which controls and possibly eliminates this tendency to Depression.

6 Second extension: A Government Sector

Minsky’s contention that countercyclical behaviour by government stabilises the market by constraining its tendency to debt accumulation is explored by introducing government spending and taxation as functions respectively of the rate of employment and the profit share of output.⁸ This extension requires new definitions for profit share and net profit share:

$$\begin{aligned} \text{The gross profit share of output:} & \quad \pi = 1 - \omega \\ \text{The net profit share of output:} & \quad \pi_n = 1 - \omega - t + g - r \times d \\ \text{The government spending function:} & \quad \frac{dG}{dt} = g(\lambda) \times Y \\ \text{The government taxation function:} & \quad \frac{dT}{dt} = \tau(\pi_n) \times Y \end{aligned}$$

⁸The rate of profit is $\frac{\pi}{v}$, where v is a constant in this model. To simplify exposition, I used the profit share as the argument to the investment function rather than the rate of profit




Figure 9: Bifurcation in the equilibrium government debt




Figure 10: Far from equilibrium dynamics at low interest

where $g(\lambda)$ and $\tau(\pi)$ are as defined in the glossary. This extension results in the following five dimensional model of a mixed market-state economy:⁹

$$\begin{aligned} \frac{d\omega}{dt} &= \omega \cdot (w(\lambda) - \alpha) \\ \frac{d\lambda}{dt} &= \lambda \cdot \left(\left(\frac{k(\pi_n)}{v} - \gamma \right) - \alpha - \beta \right) \\ \frac{dd}{dt} &= d \times \left(r - \left(\frac{k(\pi_n)}{v} - \gamma \right) \right) + k(\pi_n) - \pi + t - g \end{aligned} \tag{9}$$

$$\begin{aligned} \frac{dg}{dt} &= g(\lambda) - g \times \left(\frac{k(\pi_n)}{v} - \gamma \right) \\ \frac{d}{dt}t &= \tau(\pi_n) - t \times \left(\frac{k(\pi_n)}{v} - \gamma \right) \end{aligned} \tag{10}$$

The behaviour of this model is consistent with Minsky’s hypothesis. The most intriguing aspect, from a complex systems point of view, is that the addition of a government sector transforms a system which was locally stable (about the equilibrium) but globally unstable, into a system which is locally unstable but globally stable. At least half the eigenvalues of the linearised version have positive real part for all values of r , yet rather than leading to breakdown, the model is constrained by a chaotic limit cycle, as the following simulations indicate.

The second intriguing feature of this model is the relationship between government debt and the interest rate. As with the previous model, the equilibrium wages share of output is a negative linear function of the interest rate, but in addition the level of government debt is a rectangular hyperbolic function of the interest rate (see Figure 9):

$$d_g = \frac{t - g}{r - \left(\frac{k(\pi_n/v)}{v} - \gamma \right)} \tag{11}$$

⁹The paper does not consider government debt, because at this level there is no consideration of the redeployment of net government spending into aggregate demand. The level of government debt can easily be incorporated into the model via the relation , and its level (as a proportion of output) is visible in these simulations in the gap between γ and t .




Figure 11: Far from equilibrium dynamics at low interest




Figure 12: Far from equilibrium dynamics at low interest

Thus if the prevailing (real) rate of interest is below the rate of growth of output, then with the equilibrium values for t and γ given by the parameter values used in these simulations, the equilibrium value of government debt is negative. Equally, if the rate of interest exceeds the rate of growth, the equilibrium value is positive. While the actual values differ substantially from equilibrium values because of the system's far from equilibrium dynamics, this negative/positive bifurcation remains in any simulation. Figures 10–12 show the behaviour of the model with an interest rate of 3% and a .01 deviation of all system variables from the equilibrium vector.¹⁰

The phase diagram in Figure 13 makes it clear that the dynamics are now governed by a chaotic limit cycle.

The model behaviour on the other side of the bifurcation point differs in one highly significant way: whereas government debt stabilised at a low rate of interest, at a high rate of interest government debt continues to grow cyclically but exponentially. rising government deficits have been a feature of post-WWII economies, especially since the adoption of a “fight inflation first” strategy in the mid-70s in an attempt to control the rate of inflation. The cornerstone of this policy was tight monetary policy — which meant high real interest rates. Figures 14 to 17 demonstrate the behaviour of the model with an interest rate of 5% and a .01 deviation of all values from the equilibrium vector.

The apparent paradox in Figure 15 — the coincidence of a positive overall government burden on the economy and yet a growing accumulated government deficit — is explained by the impact of the high rate of interest on the current level of outstanding debt, and the already high level of debt implied by starting from the equilibrium position. However a different initial condition with low or negative initial government debt could easily result in a surplus being accumulated by the government (see [5]), as opposed to the deficit shown here.

7 Third Extension: Commodity Prices

Fisher argued that debt accumulation on its own would not be sufficient to cause a depression, but instead would give rise to cycles. However the model above indicates the accumulation of debt alone can lead to a depression — as the end product of a series of business cycles — as the fundamental asymmetry that firms incur debt during booms but have to repay it during slumps asserts itself. Deflation is thus not essential to the occurrence of a depression, but it would accelerate the process, and exacerbate its depth by its impact upon the rate of bankruptcy. Similarly, Minsky's argument

¹⁰ $(w, l, dk, g, t) = (0.300604985584, 0.971225057244, 0.070191124862, -0.145020153379, 0.390427727909, -35.696525419245)$




Figure 13: Far from equilibrium dynamics at high interest




Figure 14: Far from equilibrium dynamics at high interest

that capital goods prices are expectations-driven[13, p64, 80] implies that pro-cyclical movements in capital goods will exacerbate the accumulation of debt, thus hastening the onset of a depression in a market economy.

These issues can be explored by revising the basic system of equations to include consumer prices (P_c) and capital goods prices (P_k). We start with an income shares equation in nominal (money) terms:

$$Y = W + r \times D + \Pi, \quad (12)$$

where wages can be decomposed into a real wage, a consumer price index, and the level of employment (L):

$$W = w \times P_c \times L. \quad (13)$$

The wage change relation is now in money terms:

$$\frac{dW}{dt} = \frac{d}{dt} (w \times P_c) = W \times w(\lambda). \quad (14)$$

On the other hand, the relations between labor and output, and output and capital, must now be expressed in real terms:

$$Y_r = \frac{Y}{P_c}; \quad L = \frac{Y_r}{a} = \frac{Y}{P_c \times a}; \quad K = v \times Y_r = v \times \frac{Y}{P_c} \quad (15)$$

The introduction of a capital goods price index affects the amount paid by firms for investment goods, but the change in physical productivity continues to depend on the real increment to capital. A distinction is thus required between nominal investment (I_n) which affects bank balances, and real gross investment (I_r) which affects the capital stock:

$$I_n = P_k \times k(\pi) \times Y; \quad I_r = k(\pi) \times Y \quad (16)$$

This results in the following system of equations:

$$\begin{aligned} \frac{d\omega}{dt} &= \omega \times \left(P(\lambda) - \frac{1}{P_c} \times \frac{dP_c}{dt} - \alpha \right) \\ \frac{d\lambda}{dt} &= \lambda \times \left(\frac{k(\pi)}{v} - \gamma - \alpha - \beta \right) \\ \frac{dd}{dt} &= d \times \left(r - \frac{1}{P_c} \times \frac{dP_c}{dt} - \left(\frac{k(\pi)}{v} - \gamma \right) \right) + k(\pi) \times \frac{P_k}{P_c} - \pi \end{aligned} \quad (17)$$




Figure 15: Far from equilibrium dynamics at high interest

Leaving aside the issue of functional forms for the rate of change of the price indexes, this set of equations confirms Fisher's and Minsky's insights concerning the impact of commodity price deflation and capital goods prices. As can be seen from the debt relation, a high rate of commodity price inflation reduces the real debt burden, as Minsky emphasises, while conversely price deflation will lead, as Fisher asserts, to an amplification of the real debt burden. The rate of debt accumulation also depends on the ratio of the capital goods price index to the consumer price index, and since the $\frac{P_k}{P_c}$ ratio will rise during a boom, this will accelerate the process of debt accumulation. The price system thus apparently increases the instability of the market economy.

To proceed, functional forms for the rate of change of the price indexes must be provided. This introduces one of the most vexing issues in economics, since despite the confidence of economists that they resolved the issue of price determination in the "marginal revolution", their theory of price setting has been under attack since its inception[15] and is clearly invalid in a dynamic setting.

The theory argues that price is set by the interaction of supply and demand, where the demand price falls as price rises and the supply price rises under the pressure of diminishing marginal returns. This generates a function for profit Π as the gap between total revenue TR and total costs TC , whose maximum with respect to quantity occurs where "marginal revenue" equals "marginal cost":

$$\begin{aligned}\Pi &= TR - TC \\ \frac{d}{dQ}\Pi &= \frac{d}{dQ}(TR - TC) = MR - MC \\ \frac{d}{dQ}\Pi &= 0 \text{ where } MR = MC\end{aligned}\tag{18}$$

Sraffa's 1926 critique was directed at one of the foundations of this latter argument, that any resource could be regarded as fixed in any realistic analysis of production in a modern economy. His critique is amplified when one introduces a realistic notion of time, as opposed to Marshall's deceit that time could discretely be divided according to the variability of inputs. The neoclassical price-setting schema is clearly static: it tells how to maximise profit with respect to quantity (and thus determine price), but not how to maximise profit with respect to time. Clearly the latter objective is primary in a dynamic setting, and it can easily be determined via the chain rule:

$$\frac{d}{dt}\Pi = \frac{d\Pi}{dQ} \times \frac{dQ}{dt} = (MR - MC) \times \frac{dQ}{dt}\tag{19}$$

This equation indicates that, regardless of the nature of $\frac{dQ}{dt}$, the neoclassical “profit-maximisation” price will result in a zero rate of growth of profits over time — which is hardly the objective of any existing corporation. Far from being the obvious, simple but insightful rule which most economists believe it to be, their theory of profit maximisation (and hence price determination) is intellectually equivalent to the advice that the cheapest way to drive from point A to point β is to travel at zero kilometres per hour. An alternative pricing model must therefore be used. One candidate which is tractable at both the micro and macro level is the Kaleckian proposition that prices are set by a markup on prime costs, where these in turn are the wages bill and depreciation:¹¹

$$P_c = \frac{Y}{Y_r} = \frac{(1 + \theta) \times (W \times L + \gamma \times K)}{Y_r} \quad (20)$$

In a simple model with a fixed capital to output ratio, the depreciation component has no impact on the rate of change of prices, so that the rate of change of prices is entirely a function of the rate of change of wages and the wages share of output:

$$\frac{d}{dt}P_c = P_c \times (1 + \theta) \times (P(\lambda) - \alpha) \times \omega \quad (21)$$

When this is substituted into the model, it results in the following 4 dimensional system:

$$\begin{aligned} \frac{d\omega}{dt} &= (P(\lambda) - \alpha) \times (\omega - (1 + \theta) \times \omega^2) \\ \frac{d\lambda}{dt} &= \lambda \times \left(\frac{k(\pi)}{v} - \gamma - \alpha - \beta \right) \\ \frac{dd}{dt} &= d \times \left(r - (1 + \theta) \times (P(\lambda) - \alpha) \times \omega - \left(\frac{k(\pi)}{v} - \gamma \right) \right) + k(\pi) \times \frac{P_k}{P_c} - \pi \\ \frac{d}{dt}P_c &= P_c \times (1 + \theta) \times (P(\lambda) - \alpha) \times \omega \end{aligned} \quad (22)$$

One final modification is necessary before proceeding to simulations: given a price level, the rate of interest is no longer a real rate but a nominal one, and must therefore be allowed to vary with respect to the rate of inflation. This extension is not straightforward, since the nominal rate of interest has the crucial peculiarity that it must be positive, with a minimum rate set exogenously. The interest rate also responds to inflation in a lagged fashion.¹² These aspects of the interest rate are captured in two functions,¹³ one specifying the lagged reaction of interest rates to the inflation rate, the other ensuring that the interest rate cannot be negative even when the rate of inflation is. While the real world relationships are bound to be more complex than these, they enable a first-pass at modelling the complex relationship between prices and interest rates. Equation (23) specifies the lagged relationship between the inflation component of the rate of interest and the rate of inflation:

$$\frac{d}{dt}r_i = -\frac{1}{T} \times r_i + \frac{1}{T} \times \left(\frac{1}{P} \times \frac{dP}{dt} \right), \quad (23)$$

where r_i is the inflation-determined component of the rate of interest, and T is the time lag between changes to the inflation rate and changes to the rate of interest. Equation (24) specifies the nonlinear summation of the base and inflation-determined components of the rate of interest:

$$r = r_b + \frac{1}{2} \times \left(r_i + \sqrt{r_i^2 + \chi} \right) \quad (24)$$

where χ is a curvature factor which also puts the actual rate above the base rate at zero inflation. A base rate of 3% and a value for χ of .000009 results in the following interest rate and inflation relationship:

¹¹Kaleckians also argue that price is set, at least to some extent, to finance planned investment, though there clearly are competitive limits to this process.

¹²Lagged responses are also relevant in many other parts of this model, and will be introduced in further research.

¹³I am grateful to Trond Andresen for assistance on these issues.




Figure 16: Inflation Rate and Interest Rate Function




Figure 17: A Debt-Deflation

The full model is now:

$$\begin{aligned}
 \frac{d\omega}{dt} &= (P(\lambda) - \alpha) \times (\omega - (1 + \theta) \times \omega^2) \\
 \frac{d\lambda}{dt} &= \lambda \times \left(\frac{k(\pi)}{v} - \gamma - \alpha - \beta \right) \\
 \frac{dd}{dt} &= d \times \left(r - (1 + \theta) \times (P(\lambda) - \alpha) \times \omega - \left(\frac{k(\pi)}{v} - \gamma \right) \right) + k(\pi) \times \frac{P_k}{P_c} - \pi \quad (25) \\
 \frac{d}{dt}P_c &= P_c \times (1 + \theta) \times (P(\lambda) - \alpha) \times \omega \\
 \frac{d}{dt}r_i &= \frac{-1}{T} \times r_i + \frac{1}{T} \times (1 + \theta) \times (P(\lambda) - \alpha) \times \omega
 \end{aligned}$$

This model is capable of demonstrating “Fisher’s Paradox”, that deflation can mean that “the more debtors pay, the more they owe” [1, p344]. The process of deflation turns a low nominal rate of interest into a high real rate, and the depressing effect of debt repayment commitments on investment causes output to plummet, thus accelerating the blowout in the debt to output ratio. A true Debt-Deflation results:

There are several notable aspects to this extended model. Firstly, price dynamics almost completely subsume the income distribution dynamics of the previous models.¹⁴ This result improves upon the realism of the basic Goodwin predator-prey cycle, since one well-known stylised fact is the relative stability of income shares, which display a secular trend but little cyclical behaviour. Secondly the range of behaviours that the model can demonstrate are dramatically extended over the basic two demonstrated by the non-price model. Other initial conditions can result in: bouts of cyclical employment and inflation behaviour before either a stable outcome of a debt-induced breakdown; sustained inflation with relatively constant income shares and restrained debt to output ratios; sustained deflation with a secular collapse in workers’ share of output; and undoubtedly many more cases which will be uncovered by more systematic simulation explorations.

Thirdly, breakdown now begins at quite realistic values of the debt to output ratio, and the actual collapse can precede the blowout in debts to some extent because of the depressing effect on investment of high real rates of interest, caused by the process of deflation.

¹⁴Indeed in this particular simulation price effects do completely subsume income distribution, with worker’s share of output remaining constant throughout. This is however an artefact of the particular initial conditions chosen.

8 Conclusion

From a complex systems point of view, the addition of price dynamics to the basic Keen-Minsky model is a double-edged sword. On the one hand, it provides an additional source of potential long-term stability, with inflation countering the tendency towards the accumulation of debt. On the other, it can accelerate the process of collapse — and possibly dramatically reduce the stable region around the system’s equilibrium. A full answer to this question will have to await future research.

From an economic point of view, this model demonstrates many of the facets of Fisher’s Debt-Deflationary “creed”. Given that a debt-deflationary process is well under way in East and Southeast Asia — exacerbated in some instances by severe exchange rate movements which this model is not as yet equipped to consider — the model contains several important insights for economic management.

Firstly, contrary to conventional economic wisdom, a debt-deflation is a possibility: the events in Asia are not necessarily just the result of peculiar institutional arrangements of those countries.

Secondly, either inflation or government deficits may be necessary to overcome a debt-induced collapse — though such relatively harmless means of escape from the abyss may be rendered ineffective in a world in which finance is international and exchange rates are market-driven.

Thirdly, as is now becoming obvious even to our most conservative politicians — if not conservative economists — finance can play a destabilising role in a capitalist economy. Deregulated finance is a recipe for crisis, not efficiency.

References

- [1] Fisher, I., 1933. “The debt deflation theory of great depressions”, *Econometrica*, 1: 337–355.
- [2] Friedman, M., 1969. “The Optimum Quantity of Money”, *The Optimum Quantity of Money and other essays*, Aldine Pub. Co, Chicago.
- [3] Goodwin, R.M., 1967. “A Growth Cycle”, in Feinstein, C.H. (ed.), *Socialism, Capitalism and Economic Growth*, Cambridge University Press, Cambridge, 54–58. Reprinted in Goodwin, R.M., 1982, *Essays in Dynamic Economics*, MacMillan, London.
- [4] Greenspan, A. 1998. Director, Federal Reserve Bank of the United States of America, Speech to . . . , <http://www.bog.frb.fed.us/boarddocs/speeches/19980103.htm>
- [5] Keen, S., 1995. “Finance and Economic Breakdown: Modelling Minsky’s Financial Instability Hypothesis”, *Journal of Post Keynesian Economics*, Vol. 17, No. 4, 607–635.
- [6] Keen, S., 1996. “The Chaos of Finance”, *Economies et Sociétés*, 30, special issue Monnaie et Production (10): 55-82.
- [7] Keen, S., 1998. “Minsky’s Financial Instability Hypothesis”, *Encyclopaedic Dictionary of Political Economy*, Macmillan, London.
- [8] Keen, S., 1998b. “The nonlinear economics of debt deflation”, in Barnett, W., Chiarella, C., Keen, S., Marks, R., Schnabl, H., (eds.) *Commerce, Complexity and Evolution*, Cambridge University Press, Cambridge (forthcoming).
- [9] Keynes, 1936. *The General Theory of Interest, Money and Employment*, Macmillan, London.
- [10] Marx, K. 1867, *Capital, A critical analysis of capitalist production*, Volume I, Progress Publishers, Moscow, 1954.
- [11] Minsky, H., 1975. *John Maynard Keynes*, Columbia University Press, New York.

- [12] Minsky, H., 1977. "The Financial Instability Hypothesis: an interpretation of Keynes and an alternative to 'standard' theory", *Nebraska Journal of Economics and Business*, reprinted in Minsky 1982, 59–70.
- [13] Minsky, H., 1982. *Can "It" Happen Again?*, ME Sharpe, New York.
- [14] Pomeau, Y. and Manneville, P., 1980. "Intermittent transition to turbulence in dissipative dynamical systems", *Communications in Mathematical Physics*, 74: 189–197.
- [15] Sraffa, P., 1926. "The Law of returns under competitive conditions", *Economic Journal*, 36: 538–550.
- [16] Sraffa, P., 1960. *Production of Commodities by Means of Commodities: Prelude to a Critique of Economic Theory*, Cambridge University Press, Cambridge.

Glossary

Term	Definition	Formula
Y	Output	
π	The gross profit to output ratio	$\pi = \frac{\Pi}{Y} = \frac{Y-W}{Y}$
a	The level of labour productivity	
α	The rate of growth of labour productivity	
a_0	The initial level of labour productivity	
N	The level of population	
β	The rate of growth of population	
N_0	The initial level of population	
w	The real wage rate	
W	The nominal wage rate	
$P(\lambda)$	The Phillips curve function	
A, B, C	Constants in the Phillips curve function	$P(\lambda) = e^{A+B \times \lambda} + C$
λ	The rate of employment	
π_n	The net profit to output ratio	$\pi_n = \frac{Y-W-r \times D}{Y}; \pi_n = \frac{Y-W-r \times D-T+G}{Y}$
ω	The wages to output ratio	
γ	The rate of depreciation	
$k(\pi_n)$	The investment function	$k(\pi_n) = e^{D+E \times \pi_n} + F$
D, E, F	Constants in the investment function	
D	The level of debt (not used directly in simulations)	
d	The Debt to output ratio	
g	The government subsidies to output ratio	
$g(\lambda)$	The government subsidies function	$g(\lambda) = e^{G+H \times \lambda} + I$
G, H, I	Constants in the government subsidies function	
G	The level of government subsidies (not used directly in simulations)	
$\tau(\pi_n)$	The taxation function	$\tau(\pi_n) = e^{J+K \times \pi_n} + L$
J, K, L	Constants in the taxation function	
T	The level of taxes (not used directly in simulations)	
T	The time lag in the interest rate function	
D_k	The level of capitalist debt	
D_g	The level of government debt	
d_k	The capitalist debt to output ratio	
d_g	The government debt to output ratio	
r	The rate of interest	
r_i	The inflation component of the rate of interest	
r_b	The base rate of interest	
χ	The curvature factor in the interest rate function	

The Emergence and Collapse of Market-dominant Products

Yusuke Koyama
 Graduate School of Economics,
 Kyoto University
 yus@ma1.seikyoin.ac.jp

Sobei H. Oda
 Faculty of Economics,
 Kyoto Sangyo University,
 Motoyama, Kamigamo,
 Kita-ku, Kyoto 603, Japan
 oda@cc.kyoto-su.ac.jp

Abstract

This paper describes how marketable commodities change as new products appear one after another. In these circumstances sellers and buyers must follow some heuristic rules. This paper presents a small model and some results of simulations, which suggest that although the dynamics of sellers' and buyers' surpluses is mostly explained by the value of their behavioural parameters, which product is marketable in the long run can be affected crucially by the order in which new products appear and certain stochastic variables.

1 Introduction

Recently such phenomena as *increasing returns*[2] and *winner-take-all markets*[5] has attracted much attention in the economic literature. Among the reason for those results or features of modern market competition often mentioned, are decreasing average production cost and positive network externality among producers and consumers. Actually they play an important — sometimes crucial — role in a number of modern markets, but *increasing returns* and *path dependency* can possibly come from other factors. In this paper we shall present a small model and simulating it to show that such phenomena can be generated if market participants have bounded rationalities.

Our chief concern is competition among products for daily consumption such as packages of corn flakes and tissue paper, which can be found at any convenience store and supermarket. Consumers' behaviour and the shops' strategy in the markets for these commodities are by no means simple, but we could summarise them in the following way.

Consumers make decision only at the point of purchase[6]. They do not spend much time for purchasing such daily consumption goods as coffee, toothpaste, margarine and cereals (Dickson and Sawyer [4, p.47] report that the average time time between arriving at and departing from the product category display was less than twelve seconds).

Hence shops have only to keep a few products for each commodity from which consumers can choose a satisfactory one.¹ What shopkeepers must do is to keep not the best product for each customer but a limited number of products which are purchased by most customers.

In Section 2 we shall present a small model of such markets where a shopkeeper and consumers follow the heuristic rules mentioned above (we consider only cases where there is only one shop in the neighbourhood). Although the model presupposes very simple behaviour rules of market participants, their interactions are too complicated to be solved mathematically. We shall thus show some results of simulations in Section 3. The main results could be summarised in the following way: that although the dynamics of consumers' surplus and the shopkeeper's profit is mostly explained by the value of their behavioural parameters, which product is marketable in the short and long run can be affected crucially by the order in which new products appear and consumers' stochastic behaviour.

¹We refer to Coke, Diet Coke, Pepsi, Diet Pepsi, etc. as products, whilst calling Cola a commodity. We avoid using *brand*, which may convey some information (say image) which is not taken into account in this paper. Incidentally we suppose in the text that a manufacturer may supply more than one product categorised as the same commodity.

2 The Model

Let us describe our model. There is a village where a grocer and M consumers live. The grocer sells an assortment of products of a commodity (say cereal food), from which consumers may buy a product (say a package of corn flakes) every week. Each product has its property, or the combination of two objectively measurable properties (say calories and sweetness), while every consumer has her favourite property, or the combination of the properties that she likes most. The surplus which consumer i obtains from the consumption of product j is determined by $U - d(i, j) - P$, where U , $d(i, j)$ and P stand for the maximum utility a consumer can obtain from the consumption of the commodity, the Euclidean distance between consumer i 's favourite property and product j 's property, and the price of the commodity respectively. Here U and P are assumed to be common to all consumers and products.

The grocer has two shelves and introduces N new products every T weeks. (Hereinafter we refer to a period of three weeks at the beginning of which new products appear as a *month*.) He uses Shelf One exclusively for those commodities which he has introduced in the present month while keeping Shelf Two for those commodities which he introduced in previous months.

Every week every consumer goes to the grocer's shop. She immediately goes to one of his two shelves. Then she picks up a product from the shelf and checks its property. If she finds it satisfactory, or if she can get positive surplus from its consumption, she buys it; otherwise she takes another product from the same shelf and checks its property.² She repeats this until she finds a purchasable product or she has checked all or *sufficiently many* products on the shelf in vain. (Here *sufficiently many* is referred to *at most* S , where S is a positive constant common to all consumers.) In the former case she buys the product and leaves the store (if she finds a purchasable product before checking all or S products on the shelf, she never goes on searching for a preferable product), while in the latter case she leaves the shop without purchasing anything (no consumer is supposed to search both shelves.)

The above-mentioned consumers' behaviour is not deterministic. We assume that every consumer chooses Shelf One or Shelf Two with equal probability and that she picks up Product X from the shelf according to its occupying space $\alpha(X)$.³ If there are 50 packages of Product A, 30 packages of Product B and 20 packages of Product C on the shelf she has chosen, $\alpha(A) = 0.5$, $\alpha(B) = 0.3$ and $\alpha(C) = 0.2$ so that she first picks up A, B and C at 50, 30 or 20 percent probability respectively; if she takes up A and not be satisfied with it, she may try another product, which will be B at 60 percent probability and C at 40 percent probability.

The grocer shelves products at the beginning of every week. When new products do not appear, he deals with his two shelves independently but according to the same principle. Let us designate $\beta(X)$ as the ratio of the sales of Commodity X to the total sales from the shelf where it is (since the price is common to all products, it makes no difference whether $\beta(X)$ is measured in monetary terms or in quantitative terms). The grocer removes those products where β was smaller than a certain value $E \in (0, 1)$ last week, restocking the remaining products proportional to their last week's β . As an example let us suppose that there were only three products A, B and C on a shelf last week and that their sales were 540, 360 and 100 respectively. If $E = 0.15$, the grocer takes away C to put A and B with $\alpha(A) = 0.6$ and $\alpha(B) = 0.4$ for this week. (The grocer keeps the same value of α for all products for a week: as soon as a package of a product is sold, he puts another package of the commodity on the shelf.)

The grocer combines his two shelves only at the beginning a month. Then he calculates — as he does at the beginning of every week — β for each shelf and — unlike in the usual shelving — select those products whose β is less than $2E$ and puts them on Shelf Two according to their last week's β , putting all new products equally on Shelf One: $\alpha = \frac{1}{N}$ for each new product.

The following example illustrates the above-mentioned monthly shelving. Suppose that the sales

²Obviously unrealistic is our assumption that consumers know perfectly the property of those products which they have not bought before. Thus we have also made simulations where consumers can measure $d(i, j)$ only with certain errors, say ten percent. However, we have not found any significant differences between such cases and cases mentioned in the text.

³This assumption is based on the empirical studies of Borrin and Faris [3], and Agrawal and Schorling [1].




Figure 1: Distribution of Consumers' Favourite Property Points




Figure 2: Distribution of Products' Property Points

of the last week of the last month are as follows: from Shelf One 15 packages of Product A, 12 packages of Product B and 3 packages of Product C; from Shelf Two 15 packages of Product D, 15 packages of Product E and 10 packages of Product F. Then $\beta(A) = 0.5$, $\beta(B) = 0.4$, $\beta(C) = 0.1$, $\beta(D) = \beta(E) = 0.375$ and $\beta(F) = 0.25$. If $E = 0.15$, at the beginning of this month, the grocer keeps A, B, D and E with $\alpha(A) = 0.3 = \frac{0.5}{0.5+0.4+0.375+0.375}$, $\alpha(B) = 0.24$, $\alpha(D) = 0.23$ and $\alpha(E) = 0.23$.

3 Simulations

3.1 The Value of Parameters and the Initial Condition

Having done a number of simulations with different conditions, in this paper we should like to mention only the cases where there live 200 consumers ($M = 200$); five new products appear every three weeks ($N = 5$ and $T = 3$); the period of simulation is 60 weeks or 20 months; the price of a product is 30 ($P = 35$); the maximum surplus from consuming a product is 60 ($U = 60$); consumers' favorite property points cluster around (20, 80), (80, 60) and (20, 20) as Figure 1 shows; products' property points are widely scattered as Figure 2 shows.

All simulations in this section presuppose the distribution of property points mentioned in Figures 1 and 2. In addition, though not explicitly shown in Figure 2, the 100 products appear in a fixed order with respect to the product properties in every simulation. Maintaining all the above-mentioned conditions and values, we shall examine the following four cases: $E = 0.05$ and $S = 2$ (the solid line with circles); $E = 0.05$ and $S = 3$ (the solid line with squares); $E = 0.15$ and $S = 2$ (the bold line with circles); $E = 0.15$ and $S = 3$ (the bold line with squares). Each case is designated as the line mentioned in the parentheses in all the figures of this section. For each of the four cases we have done one hundred simulations with different random seeds which determine the consumers' probabilistic choice of shelves and products. All the lines in the figures of this section represent the average weekly values calculated from the one hundred simulations.

We shall show the results of simulations for the basic model in the next subsection, mentioning the results for the generalised model in the conclusion.

3.2 The Results of Simulations

Let us start our analysis by examining Shelf One. Figures 3 and 4 represent, for each week of a month, the average number of products on the shelf and sales (the number of packages sold) from the shelf respectively.

We can see from the figures that shelving works: all four lines representing the number of products are decreasing while all four lines designating sales are increasing. In fact the grocer's shelving and




Figure 3: The number of Products on Shelf One (Weekly)




Figure 4: Sales from Shelf One (Weekly)




Figure 5: The Number of Products on Shelf Two (Monthly)




Figure 6: Sales from Shelf Two (Monthly)

consumers' behaviour make marketable products more marketable in turn: If $\beta(X)$ increased last week, the grocer accordingly increases $\alpha(X)$ this week, which will increase the chance where X will be picked up by consumers so that $\beta(X)$ will increase.

Let us now examine Shelf Two. Unlike products on Shelf One, which appear with random property and are replaced every month, products on Shelf Two may stay for more than one months if they are marketable. Actually the grocer may probably expect that sales from Shelf Two increases in the long run as the result of shelving. To examine whether shelving works as the grocer expects, we may examine the average number of products on Shelf Two (Figure 5) and the weekly sales from the (Figure 6).

Let us examine cases where $E = 0.15$. In Figure 5 both bold lines are often below the horizontal line of 3. This suggests that there often remain less than three products on Shelf Two (although it is not apparent in the figure, there usually remain only products for the second and third weeks of a month particularly in the second half of the simulation period). Seeing that consumers' favorite property points are distributed around three distant points, we can guess that about one third of those who have chosen the shelf (about $\frac{M}{2} = 100$ persons) may not be able to find a purchasable product. In fact both bold lines fall below the level of $66.7 = \frac{2}{3} \times 100$ in the long run.

Let us examine how over-elimination of products or awkward shelving, the shelving which discards all products whose properties are near the centre of the distribution of consumers' favorite properties,




Figure 7: Total Consumers' Surplus (Monthly)




Figure 8: Average Purchasers' Surplus (Monthly)

is made. We shall see it is crucial that 0.15 is only marginally smaller than $\frac{1}{6}$.

Over-elimination of products can be the direct result of common weekly shelving. Let us suppose that there are four products A, B, C and D on a shelf. If the properties of A and B are very close to two of the three centres of the consumers' favorite properties respectively while C and D are equally near but not very close to the other centre, both $\beta(C)$ and $\beta(D)$ will be a little smaller than $\frac{1}{6}$. If they are less than 0.15, both products are removed from the shelf.

Monthly shelving can be awkward too. Suppose that at the end of the last month both Shelves One and Two contain three products which are purchasable by the most of the three groups of consumers respectively. If $E = 0.05$, at the beginning of this month all six products will be on Shelf Two and hopefully a less favourable product may be removed for each group of consumers through the preceding weekly shelving. Nevertheless if $E = 0.15$, the less preferable product or both products to a group of consumers may be discarded by the shelving at the beginning of this month. Moreover if one or two products remain for each group of consumers, they have such properties as mentioned in the example of the previous section.⁴

Let us now examine cases where $E = 0.05$. We can see from the bold lines in Figure 5 and 6 that for both $S = 2$ and $S = 3$ there are about four products on Shelf Two while sales from the shelf are rather different between the two cases. The difference in sales is understandable if sales are calculated on the supposition that there are four products and that if consumers check all the four products, two thirds of them find a commodity is purchasable while the other one third find two purchasable products. If one hundred consumers search for a purchasable product only twice ($S = 2$), $(\frac{2}{3} \times \frac{1}{2} + \frac{1}{3} \times \frac{5}{6}) \times 100 = 61$ persons will find a purchasable products. Similarly for $S = 3$ the total sale will be $(\frac{2}{3} \times \frac{3}{4} + \frac{1}{3} \times 1) \times 100 = 83$. Both are quite near to the long-run average sales in Figure 6.

Now let us examine the consumers' surplus. Figure 7 shows the total consumers' surplus while Figure 8 represents the per-buyer surplus, which is obtained by dividing the total surplus by the number of purchasers.

We can see from Figure 7 that the total consumers' surplus increases if S increases from 2 to 3, and from Figure 8 that it is not because the surplus which a consumer obtains from buying a product increases but because the number of purchasers increases. This is quite understandable. If each consumer can check more products, more consumers may probably find a satisfactory product,

⁴A few remarks: First, in our simulation more products are eliminated by monthly shelving than by common weekly shelving. This is because the criteria for elimination is doubled: $2E = 0.3$; otherwise many products could survive monthly shelving. Yet even then α cannot but be reduced to half them so that most of them will be removed by the succeeding common shelving. (Actually those products which would not be removed by it if $2E = 0.3$ for the preceding monthly shelving may be eliminated, because those products with $\alpha < 0.15$ remain on the shelf and accordingly α for the other products are smaller.) What is special for monthly shelving is not the doubling of E but reducing α by half. Secondly, if $E > 0.17$, such cases where no commodities remain on either shelf are observed in our simulations. In this meaning $\frac{1}{6}$ sets the maximum value for E .




Figure 9: Surviving Rates of Products




Figure 10: Elimination Tournament among Products

but in our simulations it is assumed that those who have found a purchasable product cease searching; in this case few consumers can buy a more desirable product.

Lastly let us examine the dynamics of each product. Let us define the index for survivability for each product as $V(i) = \frac{\sum_{k=0}^{100} v(i, k)}{101 - \lfloor k/5 \rfloor}$ where $v(i, k)$ represents the number of simulations where Commodity i exists on a shelf for Week k ($0 \leq v(k) \leq 100$) while $\lfloor x \rfloor$ stands for the minimum integral number that is smaller than x . Figure 9 shows the relation between $V(i)$ (the vertical axis) and $D(i)$ the distance between the product's property point and the nearest distribution centre of consumers' favorite points (the horizontal axis) for products 1 to 50, which appear in the first ten months.

We can see a negative correlations between $V(i)$ and $D(i)$. We could thus safely conclude that those products which are purchasable by more consumers are more likely to survive. Nevertheless this is such a general tendency that it cannot predict which product dominates in the long run precisely. As an example let us look at the seven leftmost products: $A_1, A_2, A_3, A_4, A_5, A_6$ and A_7 . The order of survivability is $A_2, A_1, A_3, A_6, A_7, A_5$ and A_4 .

The main reason for the discrepancy between $v(i)$ and $d(i)$ is that all commodities do not appear at the same time. Figure 10 shows $v(i, k)$ for the five products that are the nearest to $(80, 60)$: A, B, C, E and E ($E = 0.15$ and $S = 2$). Products A and B appear at the same time with marginally different D ($D(A) < D(B)$); the marginal difference is fatal for Product B, which is soon removed by shelving. Nevertheless Product A is not the product that has the greatest chance to survive in the long run: with a larger $D(C)$, Product C has a greater chance. If Products A and C were introduced at the same time with the same α , Product B would soon be eliminated, just as Commodity B is. Hence β of products on Shelves One and Two would be such that at the beginning of a month α for Commodity B would be zero or very small; as stated above such awkward shelving is observed for $E = 0.15$ in particular (yet, though less frequently, awkward shelving is often observed for a smaller E).

Whether and how long each product is marketable is crucial to producers of products, but scarcely does it affect the total consumers' surplus and the grocer's profit. Our model produces a phenomenon which Frank and Cook [5, p.3] refer to in the best-seller book: When only barely perceptible quality margins spell the difference between success and failure, the buying public may have little at stake in the battles that decide which products win. But to the manufacturers the stakes are often enormous — the difference between liquidation and the continuation of multibillion-dollar annual revenues.

4 Concluding Remarks

Needless to say, our model is so simple that there remains much room for generalisation. We should here briefly mention the generalisation of the model.

We have made and examined a model where consumers have short-term memory. To put it concretely the model assumes the following consumers' behaviour. If a consumer purchased nothing last week, her memory is empty and does what she is supposed to do in the basic model. Yet if she bought a product last week, she remembers it and when she enters the grocer's shop, she asks him if she can buy it this week too. If he answers in the negative, she does what she would do if she did not buy a product last week. If he answers in the affirmative, she immediately buys it with a certain probability F ($0 < F < 1$) or goes to either shelf with the remaining probability $1 - F$. In the latter case she looks not for a purchasable product but for a preferable one: Consumer i buys Product j not when $0 < U - d(i, j) - P$ but when $d(i, j) < d(i, l)$ where Product l is the product she bought last week. If she has checked all or sufficiently many products on the shelf in vain, she buys the one she bought last week rather than leaving the shop without buying anything. There is no other difference between this model and the basic model defined in Section 2.

Simulations show the following. First, consumers' obtain more surplus when they have short-term memory. This is simply because a consumer never fail to gain at least the same surplus as she got last week if she can purchase the same product as she bought last week. Secondly, those commodities which are introduced in earlier weeks have greater chance to survive in the long run. This is also quite understandable: those products which are stored in consumers' memory has an advantage.

References

- [1] D. Agrawal and C. Schorling. Market share forecasting: An empirical comparison of artificial neural networks and multinomial logit model. *J. Retailing*, 72:383–407, 1996.
- [2] W.B. Arthur. *Increasing Returns And Path Dependence in The Economy*. The University of Michigan Press, 1994.
- [3] N. Borin and P. Farris. A sensitivity analysis of retailer shelf management models. *J. Retailing*, 71:153–171, 1995.
- [4] P.R. Dickson and A.G. Sawyer. The price knowledge and search of supermarket shoppers. *J. Marketing*, 54:42–53, 1990.
- [5] R.H. Frank and P.J. Cook. *The Winner Take All Society*. Free Press, 1995.
- [6] Nielsen Marketing Research. *Category Management. Positioning your Organization to Win*. NTC Business Books, Chicago, 1992.

An Explanation of Generic Behavior in an Evolving Financial Market

Shareen Joshi

Current Address: Santa Fe Institute,
1399 Hyde Park Rd.,
Santa Fe NM 87501

Departments of Mathematics and Economics
Reed College
sjoshi@reed.edu

Mark A. Bedau

Department of Philosophy
Reed College
mab@reed.edu

Abstract

The Santa Fe Artificial Stock Market [13, 4] is an agent-based artificial model in which agents continually explore and develop expectational models, buy and sell assets based on the predictions of those models that perform best, and confirm or discard these models based on their performance over time. The purpose of this paper is to classify the different types of behavior that emerge in the market as a function of evolutionary learning rate, and to explain these emergent behaviors. We observe four different types of behavior, which are distinguished by their effects on the volatility of prices, the complexity of strategies, and the wealth earned by agents over time. We also show that the differences between these behaviors may be attributed to variations in the rate at which agents revise their trading rules and the subsequent types of rules—technical or fundamental—that emerge in the market.

1 Introduction

Financial markets are complex. Their booms and crashes [15, 16, 17], distinct moods [1], and nonlinearities [14, 8, 9] all blunt the analytical tools of traditional economic theory. Reexamination of financial market behavior with the new techniques of agent-based economic modeling is now suggesting that this type of complexity may be an intrinsic property of such systems [13, 4, 10, 7].

The Santa Fe Artificial Stock Market, developed by Brian Arthur, John Holland, Blake LeBaron, Richard Palmer, and Paul Taylor at the Santa Fe Institute, provides a compelling example of how simple endogenous forces can cause complex market behavior. Arthur et al. [13, 4] showed that varying the rate at which individual agents learn new investment strategies reveals two different kinds of overall market behavior. If investment strategies evolve slowly, the market showed behavior generally consistent with the prediction of traditional economic theory. But if the strategies were allowed to evolve more quickly, the market showed the kind of instabilities and statistical properties typically observed in real-world markets. Their work suggests that the cause of the complex behavior of financial markets may involve the rate at which investment strategies evolve.

This paper follows up on the work of Arthur et al. by taking a closer look at the kinds of behavior exhibited by the Santa Fe Stock Market model. We systematically study how the market's behavior depends on the rate of evolutionary learning, classify the various behaviors that emerge, and attempt to explain these behaviors. The main novelty of the present study is the light shed on market behavior by the historical patterns in the activation of investment strategies.

2 The Santa Fe Artificial Stock Market

The artificial stock market we study here was developed by Brian Arthur, John Holland, Blake LeBaron, Richard Palmer, and Paul Taylor [13, 4]. The market consists of a population of heterogeneous agents that buy, sell, and hold stocks and bonds. An agent's buy, sell, and hold decisions are made on the basis of that agent's beliefs about whether the stock's dividend is likely to go up or down, and those beliefs are determined by a set of market forecasting rules that are continually being

assessed as to accuracy. Over time an agent's set of market forecasting rules evolve under the action of a genetic algorithm.

The following sections provide a brief introduction to the Santa Fe Artificial Stock Market model. More detailed descriptions are available elsewhere [13, 4]. When mentioning some of the model parameters below, we indicate the specific parameter values we used in the work reported here with typewriter font inside brackets [like this].

2.1 The Market

The market contains a fixed number N [25] of agents that are each initially endowed with a certain sum of money (in arbitrary units) [1000]. Time is discrete. Each time period each agent must decide whether to invest her money in a risky stock or in a risk-free asset analogous to a real world Treasury Bill. The risk-free asset is in infinite supply and pays a constant interest rate r [10%]. The risky stock, issued in N shares, pays a stochastic dividend that varies over time. The stock's dividend stream is an exogenous stochastic process whose present value is unknown to the agents.

Agents apply their market forecasting rules to their knowledge of the stock's price and dividend history to perform a risk aversion calculation and decide how to invest their money at each time period. The price of the stock rises if the demand for it exceeds the supply, and falls if the supply exceeds the demand. Each agent in the market can submit either a bid to buy shares, or an offer to sell shares—both at the current price p_t —or neither. Bids and offers need not be integers; the stock is perfectly divisible. The aggregate demand for the stock cannot exceed the number of shares in the market. The agents submit their decisions and offers to the market specialist—an extra agent in the market who controls the price so that his inventory stays within certain bounds. The specialist announces an initial trial price, collects bids and offers from agents at that price, from these data announces a new trial price, and repeats this process until demand and supply are equated. The market clearing price serves as the next period's market price.

2.2 Agents and Market Forecasting Rules

Agents possess a constant absolute risk-aversion utility function of the form $U(c) = -\exp(-\lambda c)$, where λ [0.5] measures the extent of risk aversion and $0 < \lambda \leq 1000$. At each time period each agent determines the number of shares and risk-free bonds that maximizes her utility of consumption. The outcome of this decision depends on the agent's estimate of the profitability of the stock and bond.

The agents make their investment decisions by using a set of hypotheses or rules about how to forecast the market's behavior. At each time period, each agent considers a fixed number [100] of forecasting rules. The rules determine the values of the variables a and b which are used to make a linear forecast of next period's price:

$$E(p_{t+1} + d_{t+1}) = a(p_t + d_t) + b$$

where p_t is the trial price and a and b are the forecasting parameters. The forecasting rules have the following form:

$$\text{IF (the market meets condition } D_i) \text{ THEN } (a = k_j, b = k_l)$$

where D_i is a description of the state of the market and k_j and k_l are constants.

Market descriptors (D_i) match certain states of the market by an analysis of the price and dividend history. The descriptors have the form of a boolean function of some number [12] of market conditions. The set of market conditions in each rule is represented as an array of bits in which 1 signals the presence of a certain condition, 0 indicates its absence, and # indicates "don't care". The breadth and generality of the market states that a rule will recognize is proportional to the number of # symbols in its market descriptor; rules with descriptors with more 0s and 1s recognize more narrow and specific market states. As these strings are modified by the GA, the number of 0s and 1s might go up, allowing

them to respond to more specific market conditions. An appropriate reflection of the complexity of the population of forecasting rules possessed by all the agents is the number of specific market states that the rules can distinguish, and this is measured by the number of bits that are set in the rules' market descriptors.

There are two different kinds of market conditions: those pertaining to trends in the stock price, which are recognized by *technical* trading bits, and those pertaining to the relationship between the stock's price and its fundamental value, which are recognized by *fundamental* trading bits. So, there are two (overlapping) kinds of rules, depending on whether their descriptors have technical or fundamental bits set. Technical trading rules are activated when the current state of the market meets some condition pertaining to a price trend (e.g., the condition that the current stock price exceeds the average price over the past fifty time periods). Fundamental trading rules are activated when the current state of the market meets a condition pertaining to the relation between the stock's price and fundamental value (e.g., the condition that the the current stock price times the interest rate divided by the most recent stock dividend exceeds 0.75). This method of modeling expectation formation makes it is possible to track exactly which descriptor bits (technical or fundamental) are being used by agents in the model, and this allows us to study the conditions under which technical trading emerges in the market.

An example may help clarify the structure of market forecasting rules. Suppose that there is a twelve bit market descriptor, the first bit of which corresponds to the market condition in which the price has gone up over the last fifty periods, and the second bit of which corresponds to the market condition in which the price was 75% higher than its fundamental value. Then the descriptor 10##### matches any market state in which the stock price has gone up for the past fifty periods and the stock price is not 75% higher than its fundamental value. The full decision rule

IF 10##### THEN ($a = 0.96, b = 0$)

can be interpreted as "If the stock's price has risen for the past fifty periods and is now not 75% higher than its fundamental value, then the (price + dividend) forecast for the next period is 96% of the current period's price."

If the market state in a given period matches the descriptor of a forecasting rule, the rule is said to be *activated*. A number of an agent's forecasting rules may be activated at a given time, thus giving the agent many possible forecasts to choose from. The agent decides which of the active forecasts to use by measuring each rule's accuracy and then choosing at random from among the active forecasts with a probability proportional to accuracy. Once the agent has chosen a specific rule to use, the rule's a and b values determine the agent's investment decision.

2.3 The Genetic Algorithm

A genetic algorithm (GA) provides for the evolution of the population of forecasting rules over time. Whenever the GA is invoked, it substitutes new forecasting rules for a certain fraction [5%] of the least fit forecasting rules in each agent's pool of rules. A rule's fitness is determined by both how well it has performed and by how complex it is (the GA has a bias against complex rules). Applying the genetic operators of mutation, crossover, and inversion to the most successful rules in the agent's rule pool creates the new rules, with more accurate rules producing more offspring. New rules are assigned an initial accuracy by averaging the accuracy of their parent rules.

The only market parameter that we varied in the results described below is the waiting time between invocations of the GA. We term this waiting time between GA invocations the *GA interval*. So, if the GA is invoked every time period, GA interval is 0; if the GA is invoked every 1000 time periods, GA interval is 1000; if the GA is never invoked, GA interval is 300000 (this was the total length of the simulation).

The model contains another mechanism for changing an agent's rules. If some agent's rule is not activated (thus not considered for use) by an agent for a significant number of time periods [1000],




Figure 1: Left: Variance of the stock price time series as a function of GA interval. A line showing the mean variance values at each GA interval overlays a scatter plot of variance values from all the simulations at various GA intervals. The far left of the GA interval scale represents interval zero. Note that variance is very low at the two boundary conditions (very small and very large intervals), and that between those boundary conditions variance is proportional to GA intervals. Right: Average final wealth of investors in the market as a function of GA interval. A line showing the average final wealth values at each GA interval overlays a scatter plot of wealth values from all the simulations at various GA intervals. Comparison with variance of the price stream (above) shows that investor accumulated wealth is inversely proportional to variance of the stock price stream between the two boundary conditions.

then one of the bits in the rule that is set is changed to a # so that it matches a broader set of market states. This makes it more likely to be activated and used by agents in the market.

3 Experimental Methods

We systematically studied how the behavior of the market depends on a key model parameter identified in earlier work—the interval between successive invocations of the genetic algorithm (GA), which we will term the “GA interval.” Previous experiments with GA interval [13, 4, 12] simultaneously varied the probability of crossover and the accuracy updating parameter.¹ Here, we fix the crossover probability at 0.3 and the accuracy updating parameter at .01. All simulations were run for 300,000 time periods in order to make the results independent of the initial random assignment of forecasting rules and to allow the asymptotic properties to emerge. We collected statistics on stock prices, stock trading volumes, accumulated wealth of agents, and number of bits set (technical and fundamental) in forecasting rules.

In order to explain the behavior we observed, we also collected data on the activation histories of various rules during a simulation. The activation history at time period t is the number of times a particular rule has been activated until time period t (summed over time). If a hypothesis is activated but not used, in one way or another it will eventually be removed by the genetic algorithm. So a rule’s activation history us a rough indication of the number of times it has actually been used by an agent in the market.

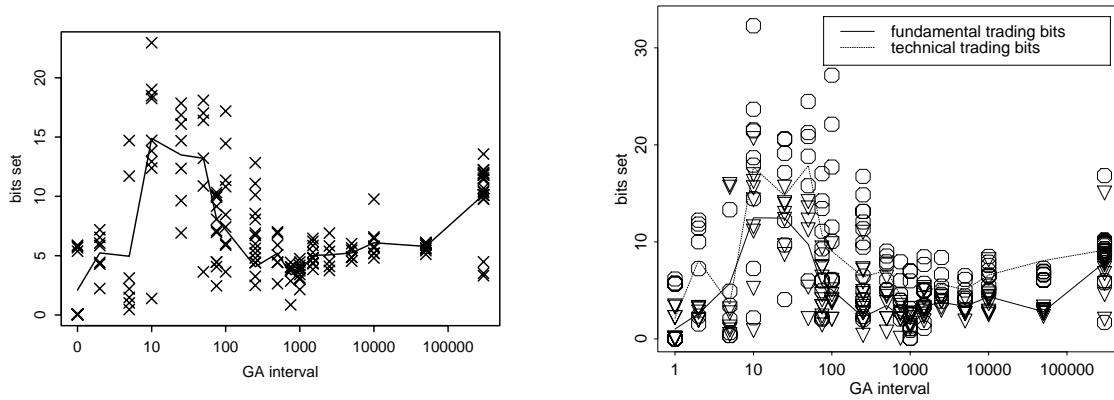


Figure 2: Number of bits in each agent’s pool of trading strategies that are set to non-null values (a measure of strategy complexity) as a function of the GA interval. A line showing the average number of bits set at each GA interval overlays a scatter plot of data from all the simulations. Left: all bits are graphed together. Right: technical trading bits (open triangles) and fundamental trading bits (open dots) are graphed separately. The number of bits set is normalized (*i.e.*, divided) by the total number of bits available. The number of bits set at very large GA intervals simply reflects the number of bits set in the initial population of strategies; the GA cannot change the strategy bits if it virtually never runs. When the GA interval does significantly change the complexity the strategies, large interval GA *lowers* it, small interval GA *raises* it, and very small GA interval *lowers* it.

4 Results

We observed four distinct types of behavior in the model, corresponding to four kinds of evolutionary learning. Two have been previously noted [13, 4]; the other two are boundary conditions. The differences between the four kinds of behavior can be seen in the volatility of prices, the wealth earned by agents (Figure 1), the total number of bits that are set in the forecasting rules, the relative number of technical and fundamental bits set (Figure 2), and the activation histories of the rules used by agents (Figures 3 and 4). Other differences (not shown here) can be seen in the mean prices, the trading volumes, and the deviations of the stock price from its fundamental value. The four classes of behavior can be summarized as follows, starting with the two boundary conditions:

¹Unpublished results involving variation in GA interval alone have been mentioned in a footnote in [4].




Figure 3: Activation of the investors’ individual trading strategies as a function of time, at three GA intervals. Left: GA interval is 300,000; the GA never runs. Middle: GA interval is 10,000; the GA runs 30 times in 300,000 time periods. Right: GA interval is 1,500; the GA runs 200 times in 300,000 time periods.




Figure 4: Activation of the investors' individual trading strategies as a function of time, at three GA densities. Left: GA interval is 250; the GA runs 1250 times in 300,000 time periods. Middle: GA interval is 5; the GA runs 60,000 times in 300,000 time periods. Right: GA interval is 0; the GA runs every time period.

Class I: No evolution so no rule switching. When the GA is never invoked (GA interval is the length of the simulation, i.e. 300,000 time periods), the agents have no choice but to stick with the pool of hypotheses with which they were initially endowed. The main characteristics of this regime are low volatility of prices, low accumulated wealth, and similar levels of fundamental and technical trading.

Class II: Too fast evolution prevents rule switching. When the GA is invoked at every time period (GA interval is 0), the prices are very stable, the complexity of strategies is very low, there is no significant difference between technical and fundamental trading, and wealth earned is high.

Class III: Slow evolution enables only slow rule switching. When the GA interval is moderately low ($1000 \leq \text{interval} \leq 10000$), price volatility is moderately low, the complexity of forecasting rules is low, wealth earned is high, and technical trading is low. In previous work the model authors noted that this class of behavior is consistent with the predictions of the theory of Rational Expectations and the efficient markets hypothesis in finance, so they called this the Rational Expectations (RE) regime [13, 4].

Class IV: Fast evolution encourages frequent rule switching. When the GA interval is moderately high ($100 < \text{interval} \leq 1000$), prices are volatile, the complexity of strategies is very high, wealth earned is low, and there is significant technical trading. The model authors observed that prices in this class of behavior deviate significantly from their fundamental values, bubbles and crashes occur frequently and the market shows statistical properties similar to real world stock markets [13, 4]. They called class IV the Complex Regime.

Classes I and II are very similar but we classify them separately because their behavior has significantly different causes. In Class II the GA is invoked at each time step and so the pool of decision rules is constantly changing, whereas in Class I the GA is never invoked and the pool of rules undergoes no changes at all. The behavior seen in Class II arises from a market that appears to be somewhat chaotic, even though it resembles a regime that is the exact opposite.

It is important to note that the classes described above are separated by periods of transition. At GA interval of 5 for example, the market shows characteristics of Class II and Class III behavior. The time series data of stock prices, wealth, technical and fundamental trading and the complexity of strategies appear to belong to class III, and the underlying behavior resembles both Class II and Class III (Figure 4).² An interesting topic for future research is to investigate the exact nature of the transition between these classes.

²The activation history graph (Figure 4) shows that the set of strategies used by agents is quite stable over time. This makes it similar to Class II. But unlike class I, some other strategies are also used (though not as frequently as the set of stable strategies). This makes it resemble Class III.

5 Discussion

The four different classes of behavior described above may be attributed to the effects of GA invocation rates on agent's evolutionary learning. Evolutionary learning affects the rate at which the agents switch between trading strategies. At the boundary conditions (GA interval 0 and GA interval 300,000) evolutionary learning is virtually nonexistent and so there is no significant evolution of trading strategies. Since the agents' trading strategies are relatively stable, so is the price series in the market. By contrast, when the GA interval is moderately low or moderately high, evolutionary learning is significant and this leads the agents' trading strategies to evolve, and this in turn makes the market less stable.

The speed at which agents switch strategies also affects the *type* of rules that they use: technical trading is significantly higher when the GA interval is moderately small. One explanation of this effect, developed below, depends on the connection between the "breathing time" a new rule enjoys before being scrutinized by the GA. Arthur et al. provide an additional explanation of this effect [4]: When GA interval is small, the agents switch rules often enough that it becomes likely for similar technical trading rules to be used by other agents in the population. Technical trading rules, when used by enough agents, can become self-fulfilling prophecies—if enough people believe the stock price is due to increase and buy the stock as a result, their demand for the stock will drive up the price—thus leading to market bubbles and crashes. Market volatility is roughly proportional to the presence of technical trading, so the regimes with less technical trading are significantly more stable.

In class I with GA interval at or near 300,000, the same pool of market forecasts available to the agents virtually never changes. The number of technical and fundamental bits set in the population of forecasting rules (Figure 2) reflects the complexity of the rules randomly assigned at the start of the simulation. In addition, as Figure 3 (top) shows, the rate at which different forecasting rules are activated by the market states is quite constant over time, and presumably the rules the agents actually use is similarly constant. In fact, fully a quarter of all of the available rules are activated virtually every time period, and thus contribute to the slope 1 line in the Figure. The agents' behavior becomes quite stable and predictable, which makes the market stable and predictable in turn, as Figure 1 (top) shows. (We are unsure why average final wealth in this regime varies as observed in the bottom of Figure 1.)

In class II with GA interval at or near 1, the GA's continual operation causes continual flux in the population of rules available to the agents. Yet, as Figure 3 (bottom) shows, virtually always the same subset of forecasting rules is activated. Furthermore, close to 95% of the available rules contribute to the slope 1 line representing these continually activated rules. The rules the agents actually use are chosen from these continually reactivated rules, of course, so Figure 3 (top) shows that the agents' trading strategies are stable over time. Thus, although there is a continual flux in the population of rules, the subset of rules actually used virtually never changes. The same subset (5%) of rules is continually replaced by the GA. Thus in class II the genetic algorithm only generates useless hypotheses so the rules being used never changes. As in class I, this stability of forecasting strategies makes the market relatively stable and predictable, as Figure 1 (top) shows. Figure 2 shows that class II evolution produces simpler strategies. This is probably due to the built-in cost of set bits, i.e., the evolutionary bias toward simpler strategies. If evolution cannot build useful strategies, as class II evolution evidently cannot, then simpler strategies should prevail. (We are unsure how to explain the variation in average final wealth seen in the bottom of Figure 1.)

Class III behavior appears when the GA interval is moderately large, roughly $1000 \leq \text{interval} \leq 10000$. The GA is invoked frequently enough for evolutionary learning to significantly improve the agents' strategies, unlike in the boundary conditions which cannot support evolutionary learning. The accumulated wealth in Figure 1 (bottom) shows the value of the strategies that evolutionary learning can produce. Only 4% of the rules are continuously activated—they are the rules that contribute to the slope 1 line in Figure 3 (bottom)—so the rules the agents actually use continue to evolve over the course of the simulation. The agents switch their investment strategies, but only relatively slowly. At the same time, the waiting time between GA invocations is long enough that newly generated

rules have a relatively long time to prove their worth before they face selection pressure from the GA. This means that evolutionary learning has an opportunity to discover those forecasting rules that are successful only in the long run (technical trading rules that identify very long-term trends or fundamental trading rules that do well only over the long haul). To the extent that agents are using rules that are successful only over the long haul, their rule use will tend to be fairly stable over time. This explanation would predict the kind of rough correlation between GA interval and price stream variance visible in the class III portion of Figure 1 (top), and the agents' risk aversion explains class III's inverse correlation between price stream variance and average final wealth (Figure 1). Evidently, these rules that focus on the long-term are not especially complex, so the GA bias toward simpler rules probably explains the relatively low complexity of class III rules (Figure 2).

Class IV behavior happens when the GA interval is moderately small, roughly $\times 1000 \leq \text{interval} \leq 100$. Figure 1 (bottom) shows that agents are able to accumulate some significant wealth, so the GA interval is not so low that it disables evolutionary learning. Yet the waiting time between GA invocations is short enough that rules must prove their worth relatively quickly to avoid succumbing to the GA. This sort of evolutionary learning favors rules that perform well in the short run. As with class III, only 4% of the rules are continuously activated; Figure 4 (top) shows that the subset of rules that the agents actually is continually evolving. Agents are switching their investment strategies relatively quickly. This instability in investment strategies used causes instability in the stock price (Figure 1 top), and the market becomes less predictable than in any other regime. Given the agents' risk aversion, this market instability drives the price down (Figure 1 bottom), Figure 2 shows not only that the rules produced in class IV are relatively complex and use more trading bits than those in any other class; the complexity of the quickly evolving trading strategies provides enough value to outweigh the GA's built-in bias toward simple rules. In class IV, and only in class IV, evolutionary learning supports the emergence of significantly complex strategies, and complex technical trading strategies in particular.

6 Summary and Conclusion

Varying the interval of the GA in the Santa Fe Stock market results in the appearance of four distinct kinds of market behavior. These correspond to four different rates of evolutionary learning. Evolutionary learning controls the rate at which agents switch between different rules in the population of rules. It also affects the types of different strategies (technical or fundamental) that evolve over time. Differences between rates of switching between rules and the types of rules that evolve in these classes lead to differences in the volatility of prices, wealth earned by agents, the complexity of strategies, the types of strategies that evolve in the market over time and the activation history of rules.

At low GA intervals, the frequent switching between strategies as well as the significant usage of technical trading rules results in high price volatility, increases in the complexity of strategies and lower overall wealth. At longer GA intervals, the infrequent switching between rules as well as the lower usage of technical trading rules results in lower price volatility, the usage of strategies of lower complexity and higher overall wealth. At the boundary conditions the usage of the same pool of rules over time leads to very low volatility and almost equal usage of technical and fundamental rules.

In conclusion, this paper has classified the various types of behavior in the Santa Fe Stock market and provided an explanation for the differences between observed behaviors. Given the resemblance of Class IV behavior to real world financial markets [4, 12], we hope that our results are also a step toward explaining the complexity of real world financial markets. Current and future work in this area includes quantifying evolutionary activity in this model using neutral models and evolutionary activity statistics [5, 6], and also studying the emergence of technical trading in financial markets[11].

Acknowledgments. Thanks to the authors of the Santa Fe Artificial Stock Market, especially Richard Palmer and Blake LeBaron, for making their source code available to us and helping us to use it productively. For helpful discussion, thanks to Doynne Farmer, Blake LeBaron, Norman Packard, Richard Palmer, and Jeffrey Parker.

References

- [1] Arthur, W. B. 1989. Positive feedbacks in the economy, *Scientific American* (February), 92–99.
- [2] Arthur, W. B. 1992. On learning and adaptation in the economy, Santa Fe Institute Working Paper 92-07-038.
- [3] Arthur, W. B. 1994. Inductive reasoning and bounded rationality, *American Economic Review, Papers and Proceedings* **84**(2)b, 406-11.
- [4] Arthur, W. B., J. H. Holland, B. LeBaron, R. Palmer, P. Tayler. 1997. Asset pricing under endogenous expectations in an artificial stock market, in W. B. Arthur, D. Lane, and S. N. Durlauf, (eds.), *The Economy as an Evolving, Complex System II*, Menlo Park: Addison-Wesley. Also published as Santa Fe Institute Paper 96-12-093.
- [5] Bedau, M.A. N.H. Packard. 1992. Measurement of evolutionary activity, teleology and life, in Langton, C.; C. Taylor, D. Farmer, and S. Rasmussen (eds.), *Artificial Life II*, Addison-Wesley.
- [6] Bedau, M.A., E. Snyder N. H. Packard. 1998. A classification of long-term evolutionary dynamics, in C. Adami, R. K. Belew, H. Kitano, and C. E. Taylor (eds.) *Artificial Life VI*, MIT Press, MA: Cambridge.
- [7] Beltratti, A., S. Margarita. 1992. Simulating an artificial adaptive stock market. Mimeo, Turin University.
- [8] Brock, W.A. 1988. Non-linearity and complex dynamics in economics and finance, in P.W. Anderson, K. J. Arrow and D. Pine (eds.), *The Economy as an Evolving Complex System: The Proceedings of the Evolutionary Paths of the Global Economy Workshop*, Santa Fe Institute Studies in the Science of Complexity Proceedings, vol. V, Redwood City, CA: Addison-Wesley.
- [9] Brock, W.A., P. de Lima. 1996. Nonlinear time series, complexity theory, and finance, forthcoming in G. Maddala, and C. Rao (eds.), *Handbook of Statistics Volume 14: Statistical Methods in Finance*, New York: North Holland.
- [10] Caldarelli, G., M. Marsili, Y.C. Zhang. 1997. A prototype model of a stock exchange, Manchester University Working Paper 9709118.
- [11] Joshi, S., J.P. Parker, M.A. Bedau. Technical trading creates a prisoner's dilemma: results from an agent-based model, preprint.
- [12] LeBaron, L. 1995. Experiments in evolutionary finance, Working Paper, Department of Economics, November 1995, University of Wisconsin-Madison.
- [13] Palmer, R. G., W. B. Arthur, J.H. Holland, B. LeBaron, P. Tayler, 1994. Artificial economic life: a simple model of a stock market. *Physica D* **75**, 264–274.
- [14] Ruelle, D. 1988. Can Non-linear dynamics help economists? in P. W. Anderson, K. J. Arrow and D. Pine (eds.), *The Economy as an Evolving Complex System: The Proceedings of the Evolutionary Paths of the Global Economy Workshop*, Santa Fe Institute Studies in the Science of Complexity Proceedings, vol. V, Redwood City, CA: Addison-Wesley.
- [15] Shiller, R. 1981. Do stock prices move too much to be justified by subsequent dividends?, *American Economic Review*, **71**, 421-36.
- [16] Shiller, R. 1984. Stock prices and social dynamics, *Brookings Papers on Economic Activity*, **2**, 457-510.
- [17] Shiller, R. 1989. *Market Volatility*, Cambridge, MA: MIT Press.

The Complexity of Competitive Marketing Strategies

Robert E. Marks and David F. Midgley
Australian Graduate School of Management
University of New South Wales
Sydney NSW 2052
bobm/davidm@agsm.unsw.edu.au

Lee G. Cooper
Anderson Graduate School of Management
University of California Los Angeles
Los Angeles CA 90095
USA
lcooper@agsm.ucla.edu

G. M. Shiraz
Dept. of Artificial Intelligence
University of New South Wales
Sydney NSW 2052
hossein@cse.unsw.edu.au

Abstract

Genetic algorithms (GAs) have been used extensively in engineering and computer science to optimize specific functions, especially those which exhibit non-convexities and so are not amenable to calculus-based methods of optimization. A parallel use of GAs has been to solve algorithmic problems. A third domain in which GAs have been used is that of searching for mappings which optimize a repeated procedure, which also reveals their complexity. An offshoot of this has been their use in what has been called co-evolution of mappings. This paper reports results from a project in which GAs have been used to, first, to derive mappings which may explain the behavior of brand managers in an oligopolistic retail market for coffee, second, to attempt to improve on the historical profits of these brand managers, pitted in weekly competition with each other, vying for sales and profits with their different brands of ground, sealed coffee on the supermarket shelves, and, third, to reveal how the artificial agents' performance is positively related to their complexity. As well as advancing the practice of GAs, with separate populations competing, the work also advances our understanding of modeling players in repeated oligopolistic interactions, or games.

1 Introduction

The theory of oligopolistic behavior (that is, the behavior of sellers in a market with a small number of sellers, but many buyers, so that one seller's actions will affect the profits of other sellers, and vice versa) has mainly been approached from the point of view of searching for Nash equilibria in players' actions, that is, a combination of actions, where each player's actions are the best he can do for himself, given that the other players' actions are the best they can do for themselves. Such a combination is self-reinforcing, since no single player has an incentive to alter his actions.

The project reported here, however, has been concerned with trying to explain and to improve upon the historical behavior and profits of a group of sellers, as recorded in supermarket scanner data, and using a market model to predict one-shot (weekly) profits of each player, given the marketing actions of all players. The data have been described in a recent article[9]. Briefly, each player has a choice of weekly actions: price per pound, coupons, in-store promotional displays, and featured local advertising. The CASPER market model[2], estimated from historical data, is used to identify each of the several firms' weekly profits, given all brand managers' actions.

We modeled the brand managers, the players, as stimulus-response automata[6], where the response is the player's marketing actions for the next week, and the stimulus is the state of the market this week, which we took to be a function of all players' actions this week and last week and several weeks past. The reason we believe that managers remember past actions is that this means they can respond to movements (aggressive or conciliatory) in other players' pricing.

For instance, it turns out that historically most prices and most sales have been made when prices are low. So if one brand, say Folgers, were pricing aggressively low last week, and raises its price this week, this could be a signal that it is becoming less aggressive, and might like reciprocation from its rival brands. If the brand managers are able to remember more than two weeks of marketing actions, then they may respond not just to rising or falling prices of their rivals, but also how quickly these prices are rising or falling. These issues are explored at greater length in [7].

2 Modeling the Managers

We model each manager as a finite automaton that responds to the state of the market with a set of marketing actions. To do this we need a set of rules, which are here represented by a binary string, following the Axelrod/Forrest representation [1]. Each string becomes an individual in a population of strings as artificial brand managers, and each string's average profit after a series of repeated interactions with the other artificial brand managers can be used as its "fitness" for the GA [10].

To be specific, say there are p players, each with a possible actions per week, and m weeks of memory, then the total number of possible states is given by

$$\text{number of possible states} = a^{mp}. \quad (1)$$

This number increases rapidly: with three players, four actions, and one week of memory there are 64 possible states, but increasing memory to two weeks increases the number of possible states to 4,096.

Moreover, the length of the bit-string is only equal to the number of possible states in the unlikely event that a player can choose only from two possible actions, which can then be coded as zero or one. If, however, the player can choose from four actions, then the bit length doubles, and from eight actions it trebles, so that each possible state corresponds to three bits, which code for eight possible actions.

We are modeling the brand managers as boundedly rational: bounded in terms of their perceptions of reality, which is really saying that it is costly to perceive reality finely [7]; bounded in terms of their memory (which is another way of saying that their perception is limited because costly); and bounded in terms of the possible actions they can make. None the less, we found that our simple finite-automaton artificial brand managers could outperform their historical flesh-and-blood forbears [9]. In showing this, we were able to develop strings (using the GA to search through the space of possible mappings from history to actions) that represented real strategies in asymmetric markets (asymmetric because the brands historically faced different costs, evoked different responses from customers, and chose from different sets of possible actions).

This line of research does not merely pit each bit-string against a complex and sometimes noisy environment, as had been done by others, in looking at artificial players in repeated games [1]. We co-evolved the players, so that each string was being tested for its fitness against the consequences of other strings, which in turn were being tested for their fitness [5]. This may be a good example of Szpiro's "surfing in a seascape" [12].

2.1 The Agents' Choices

Given the problem of the curse of dimensionality, with rapid growth in the length of the bit strings modeling the agents, the question at first was how could we model the market interactions with the smallest sacrifice of realism? We focussed on the three most active brands in the market: Folgers, Maxwell House, and Chock Full O' Nuts, although later we have increased the number of strategic players.

We assumed that the decision to use coupons was equivalent to a reduction in price. Moreover, we chose at first to use only four possible prices, instead of the range available to the historical managers (from \$1.50 per pound to about \$3.00 per pound). For each of the three players we examined the historical pricing decisions to arrive at the brand-specific sets of four possible prices per player. At the

same time, realising that other marketing actions (advertising feature and aisle display) were highly correlated with price, we factored those into the four pricing actions. (Only when the price is low did the historical players use feature or display, presumably to move more stock at an attractive price; see [9] for more discussion.)

To begin with, we modeled the players as remembering the actions of all three players of only one week ago, although this was relaxed later. With three players, each with four possible actions per week, and one week's memory, equation (1) tells us there are 64 possible states. With four possible actions, each state must map to two bits on the player's string. When, following [1], we use six bits for the phantom memory used in the first round (effectively endogenising the initial conditions of the simulation), each player is modeled with a 134-bit string. Not only are 134-bit strings easy to simulate, but the 75 weeks of historical data provided sufficient to evolve effective strings of this length.

Although it would have been possible to link the CASPER market model (which derives each brand's weekly profit, given the other brands' actions) to the GA, we found that computing the market response functions for each iteration of the game took an excessive time, and we had problems in marrying the compiled CASPER model with the compiled evaluation function of the GA. Moreover, with only 64 possible states, it seemed more elegant to derive three $4 \times 4 \times 4$ payoff matrices off-line (one per asymmetric brand), and to compile them into the GA as look-up routines. This was done, although later we would have to increase the dimensions of this array quite considerably.

2.2 The Genetic Algorithm

There is no need to describe the workings of GAs in 1998. There are many books[10, 3] and articles doing this. Suffice it to say that in our earlier work[9] we adapted GAucsd, the U.C. San Diego version of John Grefenstette's GENESIS[11]. We describe below the extensions that we have made to it in order to examine the phenomena under review.

3 Experiments

The results of the experiments described below are reported in more detail in [9] and [8]. Our purpose here is to discuss the extensions made to the GAucsd to accommodate our models and the performance of the artificial agents.

3.1 Unconstrained Agents

Despite some expectations that collusion would occur at a high price (price is the most powerful of the several marketing actions available to the sellers, and we concentrate on it here), we found convergence, with all brands pricing at their lowest historical prices. This result was consistent with the historical observation that most sales and most profits occur at low prices with promotions, because of such behavior as stockpiling and brand-switching. Ground coffee in vacuum sealed cans has a storage life of up to seven weeks. Moreover, the historical market was mature, with no external shocks on either the supply side or the demand side.

3.2 Institutional Constraints

Unfortunately, these results were unrealistic since historically only one brand a week has priced at the low promotional level to which all brands had converged. The supermarket chain whose scanner data we were using had managed to maximize its profits while not exhausting demand. Its policy was to constrain the brands: only one brand promoting with low prices in any week, and no brand promoting with low prices in two successive weeks.

We mimicked this. Ties in which two or more brands respond to the state of the market via their mapping strings by both promoting at low prices were broken by random choice, the loser pricing arbitrarily high. In order to speed up the simulations, we determined that we could examine the

genotype (the structure of each artificial brand's bit-string) to see whether that string's low promotion price this week would be followed by a similar price next week, rather than waiting for the simulation to reveal the particular realization of the player's phenotype (its response behavior). This "filtering" of strings greatly speeded up the simulation, since strings whose structures revealed illegal successive promotions were given arbitrarily low fitness, and their characteristics were excluded from future generations of strings by the GA. After 20 generations (with a population size of 25), most illegal strings had vanished, and the last had usually disappeared by generation 44.

Although the brands' behavior was closer to that seen historically [9], we found that, because the market model CASPER had been written and estimated for a single week's interaction, the overall levels of low, promotional prices were leading, with brand switching, to demand saturation.

3.3 Demand Saturation

While the retail coffee market is very volatile in the short run, it is very stable in the long run [9]. We pro-rated the weekly total by the degree of over-saturation of the past seven weeks, chosen to approximate the average interpurchase interval for this product. We first calculated the total sales volume per week, a function of the actions of the three strategic brands and the remaining non-strategic brands (whose behavior was assumed to be static). We then calculated the average total sales volume over the previous seven weeks and with a figure for the historical average total sales volume in this market calculated the percentage degree of saturation. If this was above 100%, then the total sales volume for the latest week was reduced by the degree of saturation. (In steady state, this procedure means that total sales volume must equal the historical average.) Then the profits of the three strategic brands were reduced from the limits now placed on each brand's sales volume.

The results of this experiment are to be seen in Figure 2 of [9]. The experiment results in a greater degree of competition than observed historically, owing to the immediacy of the simulation laboratory, in which brands immediately respond to others' actions last week. The artificial brand managers thus generated average weekly profits from 3.5 to 9.7 times higher than did the historical brand managers.

3.4 Tests Against History

How well had our best artificial agents learned (or evolved) to play the game which models the oligopolistic market for coffee we are examining? In order to answer this question, we took the most profitable agents from the previous series of experiments (after 100 generations of the GA) and tested each in turn against the historical actions of their two strategic rivals. The historical actions of the five non-strategic brands were also used, but our artificial agents as modeled were blind to these actions.

This was achieved by taking a string, designating it as a particular brand, say Maxwell House, and allowing it to respond to the historical actions of the two rivals brands over a 52-week period of history. Since the historical brand managers had had a much larger range of prices and other actions to choose from (although the artificial player's range spanned the historical range), we used a rough partitioning of the historical actions into four intervals, to which the artificial agent responded [7]. Its performance was measured by its average profits over this period, calculated weekly by CASPER, with the historical actions of the other strategic and non-strategic players as input. Since the GA's population size was 25, there were 25 possible strings: only later did we separate the players into distinct populations to be evolved in parallel by the amended GA.

The results are detailed in [9]. For two brands (Folgers and Chock Full O' Nuts) most of the strings performed better than their historical counterparts did; for Maxwell House only two of the 25 strings did (although they were 20% more profitable, none the less). Maxwell House historically was the most profitable of the three brands, so perhaps the artificial agents faced a higher performance hurdle.

A criticism of this experiment is that it is an "open-loop" regime: although the artificial agent responds to the historical actions, week by week, as it had been bred to do by the GA, the historical actions are fixed, with no possibility of responding to the artificial agent's action last week.

Another criticism, which we address in Section 3.5 below, is that we were using a single population of strings in the GA. When the problem is static, a single population of strings provides many possible solutions (Holland’s “implicit parallelism” [4]), but when we engage in coevolution with asymmetrical players, as here, there is no reason to believe that “one size fits all”, especially since the same state may best trigger quite different responses in different brands.

Because of these concerns, we concluded that what was impressive about these results was not that our artificial agents could outperform their historical counterparts, but that very simple agents (with only four possible actions and one week’s memory) could generate reasonable performance in the noisy coevolutionary environment.

3.5 Multiple Population Simulations

As mentioned, despite the fact that we were coevolving asymmetric agents, we — in common with all other users of the GA — had been using a single population. As well as making it much harder for the GA to search for fitter mapping strings (consider: a single string might perform well as one brand but badly as another), a single population means that, through the genetic recombination of the GA, strings may be communicating genotypically, as well as phenotypically via their fitness (profitability) in the repeated interaction. Tony Curzon Price has called this “incest”, in a personal communication.

We have extended GAucsd to include multiple populations of bit strings, so that the fitness of any string is dependent upon all strings in the other strategic players’ populations. As well as making things less noisy for the GA, having distinct populations means that the strings are interacting only via their phenotypic behavior, and not at the genotypic structural level, since the populations are entirely separate, as far as the GA knows.

Amending the GAucsd software was not a trivial exercise, since three or four players may be interacting many times in determining each string’s fitness (its average weekly profits). One of us (Shiraz) took the opportunity to streamline the logic of the fitness evaluation functions, by recording the other strings’ performances during the round-robin interactions, so that the new code with three populations is almost as fast as the old code with a single population.

Because of the stochastic nature of the simulations, we have performed Monte Carlo simulations (50 runs each) to compare the convergence and profits of the common-population GA (25 strings, 50 simulations each) with those of the distinct-population GA (three populations of 25 strings each, 50 simulations each).

Comparing Table 1 with Table 2, we see that the distinct-population GA generates more profitable strings and converges faster than does the common-population GA.

In aggregate, the improvements to average weekly profit are only about 4%, but this summary statistic masks interesting brand-specific outcomes: with distinct string populations, Folgers’ profits increase by 3% and Maxwell House’s by 24%, while Chock Full O’ Nuts’ profits fall by 16%. Distinct populations allow the Maxwell House strings to better capitalize on that brand’s strengths.

The distinct-population GA allows the brands to differentiate themselves more in terms of the patterns of weekly response, as [8] reports. Moreover, when testing strings from the distinct-population GA against history (see Section 3.4 above), we found that strings coevolved using the distinct-population GA did better against history than did strings coevolved using the common-population GA.

Indeed, we conclude that moving to distinct populations has generally resulted in higher-performing strings, both when coevolving and when competing against the historical actions of brand managers, and that distinct populations also result in greater heterogeneity in the performance of each brand’s artificial agents.

3.6 Four Strategic Players

With the rewritten, multi-population GA code, it was relatively easy to extend the simulations to a fourth strategic player, at some cost in terms of the complexity of the bit strings, which grew in length

	Actions				Action Profit
	Low Price	High Price			
<i>Pattern 1</i>	1	2	3	4	
21 runs ^a					
Folgers	1* ^{bc}	98	0	1	\$1,022
Maxwell House	32*	7	14	47	\$631
Chock Full O' Nuts	0*	100	0	0	\$633
<i>Pattern 2</i>	1	2	3	4	
11 runs					
Folgers	0*	97	2	1	\$1,011
Maxwell House	33*	4	10	53	\$625
Chock Full O' Nuts	0*	98	0	2	\$630
<i>Pattern 3</i>	1	2	3	4	
1 runs ^d					
Folgers	46*	52	0	2	\$1,082
Maxwell House	30*	0	34	36	\$623
Chock Full O' Nuts	0*	50	0	50	\$707

^apatterns of competition are computed during the hundredth generation from all combinations of 25 agents playing 52-week games.

^brow percentages

^casterisks identify the actions constrained by store policy.

^dbest performing of remaining patterns.

Table 1: Patterns of competition among evolved agents — common population and 4 actions

	Actions				Action Profit
	Low Price	High Price			
<i>Pattern 1</i>	1	2	3	4	
25 runs ^a					
Folgers	1* ^{bc}	92	3	4	\$1,093
Maxwell House	47*	0	3	50	\$804
Chock Full O' Nuts	2*	91	3	4	\$527
<i>Pattern 2</i>	1	2	3	4	
16 runs					
Folgers	1*	94	2	4	\$1,092
Maxwell House	47*	1	3	48	\$804
Chock Full O' Nuts	1*	91	3	4	\$527
<i>Pattern 3</i>	1	2	3	4	
1 run ^d					
Folgers	2*	92	0	6	\$1,045
Maxwell House	46*	0	4	50	\$830
Chock Full O' Nuts	48*	44	4	4	\$580

^apatterns of competition are computed during the hundredth generation from all combinations of 25 agents playing 52-week games.

^brow percentages

^casterisks identify the actions constrained by store policy.

^dbest performing of remaining patterns.

Table 2: Patterns of competition among evolved agents — 3 distinct populations and 4 actions

Action	Folgers			Maxwell House			Chock Full O' Nuts		
	Price	Feature	Display	Price	Feature	Display	Price	Feature	Display
1	*\$1.87 ^a	*95	*69	*\$1.96	*95	*69	*\$1.89	*100	*77
2	\$2.07	83	0	\$2.33	83	0	\$2.02	100	65
3	\$2.38	0	0	\$2.46	0	0	\$2.29	0	0
4	\$2.59	0	0	\$2.53	0	0	\$2.45	0	0
1	*\$1.62	*67	*67	*\$1.60	*97	*97	\$1.64	0	0
2	*\$1.83	*97	*96	*\$1.87	*94	*91	*\$1.89	*97	*97
3	\$1.96	0	0	*\$2.06	*88	*76	*\$1.89	*98	*29
4	*\$2.03	*79	*77	\$2.33	79	0	\$2.01	0	0
5	*\$2.04	*85	*0	\$2.38	54	0	*\$2.02	*97	*62
6	\$2.22	96	33	\$2.52	0	0	\$2.31	0	49
7	\$2.57	0	0	\$2.53	0	53	\$2.33	0	0
8	\$2.78	0	0	\$2.59	0	13	\$2.49	0	0

^aAsterisked actions are subject to store policy.

Table 3: Sets of four and eight possible actions.

from 134 bits (three players, four actions, one-week memory) to 520 bits (including the initial week's phantom memory).

Although Hills Bros., the fourth player, was a niche player, with smaller profits than the other brands, its inclusion results in significant and complex changes in the behavior and profitability of the three major brands. The details can be read in [8]. The impacts were greater than we had anticipated, but our approach allows us to analyze the changes using a methodology based on a detailed, realistic, and empirically grounded model of consumer response.

3.7 Eight Actions per Player

We had chosen the number of four possible actions per player for convenience in our initial work, but were pleased with the results we obtained with our constrained strings none the less. But rather than exogenously imposing our decisions on the artificial managers, we would prefer them to learn which actions were most profitable, given the actions of their rivals. By increasing the number of possible actions to eight, we hoped to give the artificial managers the opportunity of demonstrating that the four actions used previously were robust, and that our assumption of a mature oligopoly was correct. Table 3 shows the four and eight possible actions by specific player.

Doubling the number of possible actions implies further complexity: from 520 bits per string to 12,312 bits per string. Of each brand's eight actions, we chose six from an historical analysis, to which we added the brand's highest observed price and lowest promotional price, thus providing each artificial manager with a much richer set of possible actions than previously.

Although in early generations of the GA simulation each of the eight actions is used with a similar frequency, by the hundredth generation (25 individuals per population) the artificial managers fall into one of two patterns of competitive interaction, as revealed by 50 Monte Carlo runs, both of which employ many fewer than eight actions. See Tables 4 and 5.

The managers have learnt the two or three actions that are most profitable for them, given the behavior of their rivals. Against the historical actions of actual brand managers, the artificial managers do at least as well as their historical counterparts. See [8] for details.

3.8 Co-evolution: Sophisticates against Primitives

Unlike the use of GAs to solve static problems, where the fitness scores of the simulation improve as generations pass, when the strings model artificial managers competing against other evolving artificial managers — co-evolution — fitness scores may not improve from generation to generation. Rather

<i>Pattern^a</i>	Actions							
	Low price				High Price			
	1	2	3	4	5	6	7	8
Folgers	*8 ^{bc}	*7	11	*8	*6	13	11	36
Maxwell House	*6	*7	*6	15	12	13	12	29
Chock Full O' Nuts	11	*7	*6	13	*7	13	12	31

^apatterns of competition computed over the first four generations of one simulation.
^brow percentages total to 100%
^casterisks identify the actions constrained by store policy

Table 4: Frequency of actions over the first four generations

<i>Pattern 1</i>	Actions							
	Low price				High Price			
	1	2	3	4	5	6	7	8
27 runs ^a								
Folgers	*20 ^{bc}	*3	11	*20	*1	25	1	20
Maxwell House	*3	*1	*3	61	2	12	2	15
Chock Full O' Nuts	3	*34	*0	10	*0	8	4	
<i>Pattern 2</i>	1	2	3	4	5	6	7	8
14 runs								
Folgers	*27	*7	11	*5	*0	31	1	18
Maxwell House	*1	*1	*3	66	9	8	5	8
Chock Full O' Nuts	1	*30	*0	*7	1	11	4	

^apatterns of competition are computed during the 100th generation from all combinations of 3 by 25 agents playing 52-week games.
^brow percentages total to 100%
^casterisks identify the actions constrained by chain policy.

Table 5: Frequency of actions during the hundredth generation

than engaging an evolved string in the open-loop competition against the frozen patterns of behavior of its historical rivals, as reported in Section 3.4 above, we take a string (the “sophisticate” from the hundredth generation and play it against rival strings (the “primitives” from the eighth generation. Table 6 presents the results.

Since the sophisticates have had many more generations to learn and adapt than have the primitives, we should expect them to score better against primitive than against sophisticated rivals. But, using the original three brands and 50-run Monte Carlo simulations, we found that for two of the three brands the sophisticates do not compete effectively with the primitives, a phenomenon that Bernhard Borges has dubbed the Holyfield-Tyson effect.

Is this due to genetic drift, where the gene pool of a small population may change randomly, when specific genes (positions on our strings) are not useful in scoring well? To test this conjecture, we increased the size of each population from 25 to 250, which means that each string now has to

Best Sophisticate	Change in Folgers	Change in Maxwell House	Change in Chock Full O'Nuts
Folgers	-15.01	41.42	42.03
Maxwell House	2.03	-20.04	37.77
Chock Full O'Nuts	13.93	-28.99	82.34

Table 6: Mean changes in average weekly profits with best sophisticate

Best Sophisticate	Change in Folgers	Change in Maxwell House	Change in Chock Full O'Nuts
Folgers	-87.11	75.13	-55.66
Maxwell House	-101.87	-512.51	155.45
Chock Full O'Nuts	-63.19	-42.08	-23.77

Table 7: Mean changes in average weekly profits with best sophisticate after 160 generations, population of 250

compete against 250^2 combinations, instead of 25^2 , and there are ten times as many strings to test, a thousand-fold increase in the number of three-way interactions per generation. Convergence is also likely to be much slower. We did not attempt Monte Carlos: a single simulation run took weeks rather than hours to complete. Table 7 presents the results.

The results of our large-population simulations [8] appear to eliminate genetic drift as an explanation, but, given the length of the cycles of convergence, we cannot rule out the emergence of higher-performing sophisticates after the hundredth generation. Moreover, we were able to in the time available to examine a model with three players and four possible actions. Would an eight-action model, allowing the artificial agents greater degrees of freedom as discussion in Section 3.7 above, demonstrate genetic drift? Our prior is no.

4 Conclusions

Although we believe that our papers provide much insight into the historical patterns of oligopolistic rivalry in a mature market, as well as revealing how historical brand managers might learn to improve their profitability and competitiveness by consideration of the patterns and strategies learnt by the artificial brand managers via the GA simulation of coevolution, we have focussed here on our contributions to the use of GAs in competition analysis.

We have shown that it is possible and appropriate to use multi-population GAs when co-evolving asymmetric artificial agents. We have shown that the GA can effectively used for bit-string agents of very high complexity. We have shown the potential of GAs to be used in exploring the patterns and strategies of asymmetrical rivals in a mature oligopoly.

References

- [1] R. Axelrod. The evolution of strategies in the iterated prisoner's dilemma. In L. Davis, editor, *Genetic Algorithms & Simulated Annealing*. Pittman, London, 1987.
- [2] L.G. Cooper and M. Nakanishi. *Market Share Analysis: Evaluating Competitive Marketing Effectiveness*. Kluwer, Boston, 1988.
- [3] D.B. Fogel. *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence*. IEEE Press, New York, 1995.
- [4] J.H. Holland. *Adaptation in Natural and Artificial Systems*. MIT Press, Cambridge, second edition, 1992.
- [5] R.E. Marks. Breeding optimal strategies: optimal behavior for oligopolists. In J. David Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms*, pages 198–207, San Mateo, Calif., June 1989. George Mason University, Morgan Kaufmann Publishers.
- [6] R.E. Marks. Repeated games and finite automata. In J. Creedy, J. Eichberger, and J. Borland, editors, *Recent Developments in Game Theory*, pages 43–64. Edward Elgar, London, 1992.

- [7] R.E. Marks. Evolved perception and behavior in oligopolies. *J. Economic Dynamics and Control*, 22(8–9):1209–1233, July 1998.
- [8] R.E. Marks, D.F. Midgley, and L.G. Cooper. Refining the breeding of hybrid strategies. Australian Graduate School of Management Working Paper, Sydney, 1998.
- [9] D.F. Midgley, R.E. Marks, and L.G. Cooper. Breeding competitive strategies. *Management Science*, 43(3):257–275, 1997.
- [10] M. Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, 1996.
- [11] N.N. Schraudolph and J.J. Grefenstette. A user’s guide to GAUCSD 1.4. Technical Report CS92-249, UCSD CSE Department, La Jolla, CA, 1992.
- [12] G. Szpiro. The emergence of risk aversion. *Complexity*, 2:31–39, 1997.

The Application of Cellular Automata to the Theory of Consumers' Learning and Behavioral Interdependence

Sobei H. Oda
 Faculty of Economics,
 Kyoto Sangyo University,
 Motoyama, Kamigamo,
 Kita-ku, Kyoto 603, Japan
 oda@cc.kyoto-su.ac.jp

Kouhei Iyori
 Graduate School of Science
 and Technology,
 Kobe University,
 1-1 Rokkodai-cho,
 Nada-ku, Kobe 657-8501, Japan
 iyori@mi-2.mech.kobe-u.ac.jp

MIURA Ken
 Graduate School of Science and
 Technology,
 Kobe University,
 1-1 Rokkodai-cho,
 Nada-ku, Kobe 657-8501, Japan
 miuraken@mbx.kyoto-inet.or.jp

Kanji Ueda
 Faculty of Engineering,
 Kobe University,
 1-1 Rokkodai-cho,
 Nada-ku, Kobe 657-8501, Japan
 ueda@mech.kobe-u.ac.jp

Abstract

This paper describes competition between such products as operating systems or applications for which demand is crucially affected by network externalities and consumers' learning by doing. In fact consumers' reservation prices for such products seem strongly depend on how long they have used (consumers' learning by doing) and how many of others will use them (network externality). By giving memory to each cell and introducing a rivalling product to the Cellular automata model of Oda[3], this paper presents a CA+Agent model and results of its simulations. Among them: that only one product may survive competition for most initial conditions if consumers' reservation prices for products do not increase very much by learning; that very small difference in the initial condition may change the final winner.

1 Introduction

When you buy an application, you may probably take account of how long you have used it and how many of others will use it. Even if an application with higher performance is available, you may hesitate to change it from the one you are familiar with, suspecting understandably that mastering a new application may require considerable time and effort. You may however abandon the use of your favorite application if increasingly many your friends and colleagues use another one, fearing naturally that adhering to it may make it difficult to exchange data and programs with others. We should like to examine such markets where consumers consider these things: consumers' learning by doing and network externalities, which are not taken into account in the standard consumer's theory but can be crucial in the so-called information-oriented society.

Our analytical tool is computer simulation as in our previous work (Oda et al [3]), which examines how the consumption of a commodity diffuse if there are network externalities among consumers' utility. The model of this paper is a generalisation of the previous cellular automata model. In fact, apart from probabilistic determination of each consumer's neighbourhood (Markus and Hess [2]), we have only introduced two factors to the CA model: that there exist two rivalling products or services and that consumers' reservation prices for them increase as they use them longer. The latter makes our model comparable to CA+Agent models (Epstein and Axtell [1]) rather than the life game (Poundstone [4]), where each cell simply reacts its neighbouring cell's action in the previous period.

Although they do not move around, cells of our model are agents which have memory to represent their skill for using each product acquired from their experience,

We shall explain our model in Section 2 and mention a few results of its simulation in Section 3. The chief concern of this paper is to detect so-called butterfly effects, or such cases where small differences in the initial conditions result in the large divergence of the final state. To put it concretely, we shall examine some cases whose initial conditions are slightly different from one another on the assumption that every consumer gains the same utility from either product if his or her experience and the number of his or her neighbours who use it are the same.

We shall first examine cases where each consumer's reservation price for either product is affected only by the number of his or her neighbours who used it in the previous period. This is the same setting as in our previous paper except that there are two rivaling products. We shall show some results of simulations to demonstrate that at least for certain combinations of parameters only one product can survive in the long run: if the sales of two products increase almost at the same rate till every consumer buys either product, after that the sales of a product gradually decrease to zero while the market share of the other product increases to one hundred percent. Although it might be determined theoretically which product is the final winner at the initial point in time, it is virtually impossible to predict which is in the diffusion process, not to mention at the initial time.

We shall then introduce consumers' learning by doing (using). Each consumer now has incentive to continue to a product he or she has used if he or she is surrounded by users of the other product. It increases, as is readily imagined, the probability that both products coexist in the long run. Actually this may give an explanation for why the Beta system has driven out of the market for home video tape recording exhaustively by the VHS system, while Mac OS still survives under the dominance of Microsoft Windows. However, even in the circumstances, small changes in the initial condition and the value of parameters can drastically affect the long-run shares of both products. We shall show some results to suggest it.

In the last section, we shall briefly mention the generalisation of our model and analysis. Since the model of this paper contains much more parameters than that of the previous paper, we have checked simulations only for limited combinations of parameters for rather restrictive purposes. We shall refer to some generalisation and simulations which we are making or planning.

2 The model

Let us suppose that there are M^2 consumers in a closed society. Every consumer has a personal computer for which two operating systems are available. To use an OS, each consumer must make a new or renewal contract at the beginning of a week (time t) with the firm which supplies it through networks for the week (Week t). We designate $X(m, n, t) = 1$ if Consumer m ($m = (m_1, m_2)$, $1 \leq m_1 \leq M$ and $1 \leq m_2 \leq M$) contracts with the supplier of OS n ($n = 1$ or 2) for Week t ($t = 0, 1, 2, \dots$) and $X(m, n, t) = 0$ if he or she does not.

Let us define the utility which Consumer m obtains from using his or her computer for Week t as:

$$U(m, t) = \max_{n \in \{1, 2\}} (X(m, n, t)U(m, n, t) + \alpha X(m, 3 - n, t)U(m, 3 - n, t)) \quad (1)$$

where α is a constant ($0 \leq \alpha < 1$) while $U(m, n, t)$ represents Consumer m 's utility from using only with OS n . In other words: $U(m, t) = 0$ if $X(m, 1, t) = X(m, 2, t) = 0$; $U(m, t) = U(m, 1, t)$ if $X(m, 1, t) = 1$ and $X(m, 2, t) = 0$; $U(m, t) = U(m, 2, t)$ if $X(m, 1, t) = 0$ and $X(m, 2, t) = 1$; $U(m, t) = U(m, 1, t) + \alpha U(m, 2, t)$ if $U(m, 2, t) \leq U(m, 1, t)$ and $X(m, 1, t) = X(m, 2, t) = 1$; $U(m, t) = U(m, 2, t) + \alpha U(m, 1, t)$ if $U(m, 1, t) < U(m, 2, t)$ and $X(m, 1, t) = X(m, 2, t) = 1$. Here the first case presumes that only with an OS each consumer can use his or her computer for himself or herself and exchange information with others through networks. The last two cases presuppose that some applications run under either OS so that using two operating systems are not twice as useful as using one.

We assume that consumers' utility from using a firm's applications consists of three terms:

$$U(m, n, t) = U_{\min} + \theta(U_{\max} - U_{\min})L(m, n, t) + (1 - \theta)(U_{\max} - U_{\min})N(m, n, t) \quad (2)$$

where U_{\min} , U_{\max} and θ are given constants ($0 \leq U_{\min} \leq U_{\max}$ and $0 \leq \theta \leq 1$).

The first term U_{\min} stands for the basic utility that a beginner can readily obtain from standing alone computer usage.

The second term $\theta(U_{\max} - U_{\min})L(m, n, t)$ represents the effect of consumers' learning by doing: one can obtain more utility from the same OS as he or she uses it longer. Here $L(m, n, t)$ stands for the skill for using OS n that Consumer m has acquired till time t , which is defined as:

$$L(m, n, t) \begin{cases} = L(m, n, 0) & \text{for } t = 0 \\ = \lambda \sum_{k=1}^t (1 - \lambda)^k \max(X(m, n, t - k), \beta X(m, 3 - n, t - k)) & \text{for } t \geq 1 \end{cases} \quad (3)$$

where λ and β are given constants ($0 \leq \lambda \leq 1$ while $L(m, n, 0)$ are all given as initial conditions ($0 \leq L(m, n, 0) \leq 1$). That is to say, $L(m, n, t)$ increases by $\lambda(1 - \lambda)^k$ if Consumer m used OS n k weeks before, while the experience of using the other OS is counted as 100β percent of that of using OS n .¹ Here λ represents the speed of skill depreciation: the greater it is, the sooner acquired skill become obsolete or is forgotten. In addition, since as is readily checked $0 \leq L(m, n, t) \leq 1$ $\lim_{t \rightarrow \infty} L(m, n, t) = 1$ if $X(m, n, 0) = X(m, n, 1) = \dots = 1$ and $\lim_{t \rightarrow \infty} L(m, n, t) = 0$ if $X(m, n, 0) = X(m, n, 1) = \dots = 0$, we can regard the second terms as the product of the degree of skill accumulation $L(m, n, t)$ and its absolute weight on the total consumer's utility $\theta(U_{\max} - U_{\min})$.

The third term $(1 - \theta)(U_{\max} - U_{\min})N(m, n, t)$ stands for the effect of network externality, which is determined by

$$N(m, n, t) = \frac{\sum_{i \in \Omega(m)} \max(X(i, n, t), \gamma X(i, 3 - n, t))}{|\Omega(m)|} \quad (4)$$

Here γ is a given constant ($0 \leq \gamma \leq 1$); $\Omega(m)$ represents the set of Consumer m 's virtual neighbours:

$$\Omega(m) = \{\text{Consumer } i \mid \text{dis}(i, m) < R\} \quad (5)$$

where R is a given constants ($1 < R$); $|\Omega|$ stands for the number of Consumer m 's neighbours. The economic reasoning behind these expressions is as follows. Even if a consumer can freely exchange data and programs and communicate with anyone else, he or she will do so only with those who share similar interests. We have thus assume that Consumer i may exchange information with his or her neighbours in the virtual space of networks.²

The economic reasoning behind these expressions is as follows. Even if a consumer can freely exchange data and programs and communicate with anyone else, he or she will do so only with those who share similar interests. We have thus boldly assume that Consumer i may exchange information with his or her neighbours in the virtual space of networks.³ Needless to say virtual neighbours may not be actual neighbours and *vice versa*; an actress may regard another company's director as her virtual neighbour whine not considering her neighbour economist so.

¹A few remarks. First in addition to skill accumulation those who have a stock of data and programs they have made on an OS have a reason for continuing to use it. The increase of $L(m, n, t)$ by the past use of OS n may be considered to reflect it. Secondly $0 \leq \beta \leq 1$ if the experience of an OS does not hinder learning and using another OS. The authors who are long Macintosh users know it is not always the case.

²A few remarks: Strictly speaking, the definition of $\Omega(m)$ is more complicated: adopting the method of Markus and Hess (1990), we assume that those cells which are partially covered by the disk whose center is m and whose radius is R stochastically belong to $\Omega(m)$. The introduction of probability to the definition of the margin of neighbourhood may have changed some properties of the CA model of Oda et al (1997); see the next section. Secondly it is assumed that a consumer regards him/herself as his/her neighbour. Although we have not made simulations on the assumption that he or she does not for the present model, we have checked the previous model. The results of simulations suggest that it would not make any difference in qualitative terms whether consumers count themselves as their neighbours or not, though it can make the results of individual simulations quite different.

³It is assumed in the text every consumer regards him/herself as his/her neighbour. Though having made simulations under the condition that he or she does not, we have not yet found any significant difference in qualitative terms.

We can find some similarities in the second and the third term. First, since $0 \leq N(m, n, t) \leq 1$, we can regard the third term as the product of the degree of network externality $N(m, n, t)$ and its absolute weight on the total consumer's utility $(1 - \theta)(U_{\max} - U_{\min})$. (As is readily understood, θ and $1 - \theta$ represent the relative importance of the effect of skill accumulation and that of the effects of network externality.) Secondly, β and γ play a similar role: β is smaller if consumers can use both operating systems in a more similar way, while γ is smaller if users of different operating systems can more easily and completely transmit information between them. Thirdly, $\Omega(m)$ corresponds to λ : the former sets the contemporary boundary to network externality while the latter limits the benefit from past experience.

In theory at time 0 Consumer m might be able to determine $X(m, n, k)$ for all n and k ($k = 1, 2, \dots$) to maximise the present value of his or her surplus: $\sum_{k=0}^{\infty} (1 + r)^k (U(m, k) - C(m, k))$ where r stands for the interest rate while $C(m, k)$ stands for the cost which Consumer m must pay at time k according to the value of $X(m, n, k)$. In practice, however, at the beginning of Week t Consumer m must find it difficult to see his or her utility for the week $U(m, t)$, because it contains $N(m, n, t)$, which depends on his or her neighbours' $X(i, n, t)$ ($i \in \Omega(m)$), which they cannot determine optimally without knowing the behaviour of their neighbours (including Consumer m 's): $X(m, n, t)$ ($i \in \Omega(\Omega(m))$), \dots

Let us assume not the hypothetical auctioneer or consumers' unbounded rationality but a simple adaptive behaviour: consumers calculate $N(m, n, t)$ on the supposition that $X(i, n, t) = X(i, n, t - 1)$ for all $i \in \Omega(m)$. In other words we assume that at time t Consumer m expects the following utility for Week t :

$$\hat{N}(m, n, t) \begin{cases} = \hat{N}(m, n, 0) & \text{fort } = 0 \\ = \frac{\sum_{i \in \Omega(m)} \max(X(i, n, t-1), \gamma X(i, 3-n, t-1))}{|\Omega(m)|} & \text{fort } \geq 0 \end{cases} \quad (6)$$

Here $\hat{N}(m, n, 0)$ are given as initial conditions ($0 \leq \hat{N}(m, n, 0) \leq 1$). We also define $\hat{U}(m, n, t)$ by replacing $N(m, n, t)$ with $\hat{N}(m, n, t)$ in 2 and $\hat{U}(m, t)$ by replacing $U(m, n, t)$ with $\hat{U}(m, n, t)$ in 1.

Having seen how consumers expect their weekly utility at the beginning of each week, we may now consider their costs. Although there are some formulations of $C(m, t)$ which can be made only in a CA+Agent model (see the last section and Appendix Two, where some results of simulations with such formulations are mentioned), we simply assume the following in the text:

$$C(m, t) = X(m, 1, t)P_1 + X(m, 2, t)P_2 \quad (7)$$

where P_n is a positive constant that represents the weekly rent for OS n . In other words we assume that rents which are assumed to be constant through time are the only cost for using computers; all applications are free or bundled with operating systems.

We can now explain consumers' weekly decision-making. At time t Consumer m calculates:

$$\hat{V}(m, n, t) = \hat{U}(m, t) - C(m, t) \quad (8)$$

for all the four possible combination of the value of $X(m, 1, t)$ and $X(m, 2, t)$: (0, 0), (0, 1), (1, 0) and (1, 1), and chooses the combination that maximises $\hat{V}(m, n, t)$ as $(X(m, 1, t), (m, 2, t))$.

3 Simulations

We shall show some results of simulations. Having made a number of simulations with different value of parameters and initial conditions and checked general tendencies, we shall show some results of simulations only for the following value of parameters and initial conditions in this section: $M = 50$, $P_1 = P_2 = 0.25$, $R = 2$, $U_{\min} = 0.2$, $U_{\max} = 0.4$, $\alpha = \beta = \gamma = 0$, $\lambda = 0.5$ and $L(m, n, 0) = 0.5$ for all m and n . In other words, we show some results of simulations for different value of θ and the sets of initial uses of OS1 and OS2: $I(1) = \{m | X(m, 1, 0) = 1\}$ and $I(2) = \{m | X(m, 2, 0) = 1\}$. Here we assume that $I(1) \cap I(2) = \emptyset$.

For convenience we define new symbols: $X(n, t)$ stands for the number of OS n users for Week t ($X(n, t) = \sum_{\text{all } m} X(m, n, t)$) while $X(n)^*$ represents the number of OS n users in the long-run equilibrium $X(n)^* = \lim_{t \rightarrow \infty} X(n, t)$ if it exists.

3.1 Examples 1 and 2: cases where $\theta = 0$ and $I_1 \neq \emptyset = I_2$

These are cases where there exists only OS1 (because as is readily checked $X_2(t) = 0$ for all $t \geq 0$) and consumers do not accumulate skill by using computers (because $\theta = 0$). That is to say, this is a case examined by Oda et al [3].

Examples 1 and 2 describe how $X(1, t)$ changes from slightly different $I(1)$: $I(1)_1$ and $I(1)_2$ respectively.⁴ Here all initial users of OS1 in Example 1 are its initial users in Example 2: $I(1)_1 \subset I(1)_2$.

From the figures of examples we can see that small differences in initial conditions may result in large difference in the long-run consequence: $X(1)^*$ changes from zero to the total population² if $I(1)$ changes from $I(1)_1$ to $I(1)_2$.

This is so-called butterfly effect, which is also observed by Oda et al [3]. A few remarks may be called for. Although $0 < X(1)^* < M^2$ for a certain range of $X_1(0)$ in the examples of Oda et al [3], there does not exist a range of $X_1(0)$ resulting in $0 < X(1)^* < M^2$ in the examples above. It is partially because of the difference in the value of parameters; even in the present model the range exist for other combinations of the value of parameters. Another reason is that the margin of neighbourhood statistically changes in the present model so that clusters of users cannot remain separated in the long run from nearby clusters of users.

3.2 Examples 3, 4 and 5: cases where $\theta = 0$, $I_1 \neq \emptyset$ and $I_2 \neq \emptyset$

These are cases where there exist two rivalling operating systems: OS1 and OS2. Since all the parameters are common to both operating systems, their success or fail in the market depends on the initial distribution of users $I(1)$ thoroughly.

Examples 3, 4 and 5 describe how $X(1, t)$ and $X(2, t)$ changes from slightly different $I(2)$: $I(2)_3$, $I(2)_4$ and $I(2)_5$ respectively. Here $I(1) = I(2)_3$ in all the three examples and $I(2)_3 \subset I(2)_4 \subset I(2)_5$. In addition, although it is not shown graphically if $I(1) = \emptyset$. That is to say, we examine three cases where either OS is diffused to be dominate the whole market if the other OS does not exist.

The figures of the examples unmistakably show that the long-run consequence can change drastically if the initial distribution of users slightly alter: OS1 takes all the market if $I(2) = I(2)_3$; OS1 and OS2 share the market $I(2) = I(2)_4$; OS2 monopolises the market if $I(2) = I(2)_5$. A similar note to Subsection 1 may be mentioned here. The final patterns of OS1 users and OS2 users of Example 4 is rather simpler than the final user-non-user patterns mentioned by Oda et al [3]. It is partially because of the difference in the value of parameters; even in the present model the final patterns can be more complex for other combinations of the value of parameters and initial conditions. In addition the vibrating margins of neighbourhood makes the frontier of the users of different operating systems smoother.

In particular Example 3 may be noteworthy. OS1 users and OS2 users rapidly increase almost at the same rate till every consumer uses either product, but then the former gradually decrease and disappear in the end. Nevertheless neither products' properties nor consumers' behaviour has changed when the market is saturated. Both the rapid diffusion of OS1 and its fade-out are explained by the same value of parameters and the same utility functions, for which OS1 is destined at the initial point in time.

⁴ $I(n)_l$ means that the initial distribution of OS n users in Example l .

3.3 Examples 6, 7, 8, 9 and 10: cases where $\theta = 0.02$

Let us consider cases where consumer can make better use of an operating system as they use it longer. Examples 6, 7, 8, 9 and 10 correspond to Examples 1, 2, 3, 4 and 5 respectively.

From these examples we can see that the introduction of consumers' learning by using decrease both probability and effect of butterfly effect. As the result, the probability that $0 < X(1)^* < M^2$ for $I(2) = \emptyset$ as well as $0 < X(1)^* < M^2$ and $0 < X(2)^* < M^2$ for $I(2) \neq \emptyset$ increases.

If $I(2) = \emptyset$ and $\theta = 0.02$, it is scarcely observed that $X(1)^* = 0$. This is because two consumers who are neighbours each other can continue to use the same OS if they are surrounded by non computer users; in the circumstances their reservation price for the OS can be positive for ever (see Example 6).

Certainly they may abandon it if they are encompassed by uses of the other OS; even if they can receive positive surplus from continuing to use the same OS, they can gain more surplus from using the other one (see Examples 8 and 10).

However if an OS is inferior in the number of users, some of the users may make a small cluster. Such an island of users is encroached by the dominant OS users and disappears in the long run if $\theta = 0$ (Example 5). Nevertheless it can survive if $\theta = 0.02$ (Example 9). With small network externality consumers continue to use the same OS.

4 Concluding Remarks

We have not exhaustively examined our CA+Agent model yet. We are now making simulations for different value of parameters and initial conditions.

As to the generalisation of the model, we are considering to introduce changes in weekly rents for operating systems P_n . There certainly exists increasing return (decreasing average cost) in the production of an operating system: in comparison with the cost for making a new operating system and its updating, additional expenditure for having another consumer use it is considerably small. Hence a smaller rent can be imposed on an OS which will be used in a longer period by more people to collect its production cost.

In addition our CA+Agent model can describe local increasing return, which cannot be done by aggregate models. To use a computer not only an OS but applications are necessary. However, one need not buy or rent all applications; people have only to use those applications which are necessary for what they want to do. This implies that the computer user's cost $C(m, t)$, or the cost for using operating systems and applications, may most probably differ from user to user. Yet it may not be very different among virtual neighbours who are supposed to share similar interests and exchange data and programs frequently. We have already developed a CA+Agent model to describe such situation and making simulations.

References

- [1] Epstein, J.M. and Axtell, R. 1996. *Growing Artificial Society; Social Science From the Bottom Up*, Brooking Institution Press
- [2] Markus, M. and Hess, B. 1990. Isotropic cellular automaton for modeling excitable media, in *Nature*, vol.347, no.6288, 56-58
- [3] S.H. Oda, K. Miura and K. Ueda. 1997. The application of cellular automata to network externalities in consumer's theory: a generalisation of life game, in Langton, C.G. and Shimohara, K. (ed) *Artificial Life Five* in Langton, C.G. and Shimohara, K. (ed), MIT Press, 473-480
- [4] Poundstone, W. 1985: *The Recursive Universe; Cosmic complexity and the limits of scientific knowledge*

A Appendix

Let us assume that OS2 is an all-purpose operating system while OS1 is designed for a special purpose. In the circumstances, although the determining equation of $U(m, 2, t)$ can be common to all consumers, $U(m, 1, t)$ may probably differ from person to person: those who (do not) use a computer for the special purpose will (not) prefer OS1. In this appendix, to describe such cases, let us assume that $U(m, 2, t)$ is determined by (2) while U_{\max} is not common to all $m = (m_1, m_2)$ for OS1. To put it concretely, we assume that $U_{\max}(m^1) < U_{\max}(m^2)$ if $|m_1^2 - \frac{M}{2}| < |m_1^1 - \frac{M}{2}|$ (all the other parameters including U_{\min} are assumed to be constants common to OS1 and OS2).

Examples 11 and 12 show the basic effect of this modification. Example 11 show the case where both $U(m, 1, t)$ and $U(m, 2, t)$ are determined by (2). The initial condition is chosen that both OS equally share the market in the long run. Example 12 represents the case where $U(m, 1, t)$ differs from person to person. We can see rough gray-black-gray stripes in the final map of the market; which is the natural result of the above-mentioned modification of U_{\max} in the determination equation of $U(m, 1, t)$. A general tendency which might be less apparent *a priori*, is that the final share of OS1 is larger. Although $\max_m U_{\max}^1(m) - U_{\max}^2 = U_{\max}^2 - \min_m U_{\max}^1(m)$, the modification is usually favourable to OS1.



Examples 13 and 14 explain the advantage of a purpose-built product too. They both show the dynamics from the same initial condition. Example 13 is the case where OS1 is an all-purpose operating system like OS2; the number of the initial OS1 users is so small that OS1 users disappear in the long run. Example 14 is the case where OS1 is a purpose-designed operating system. Unless the number of the initial OS1 users is excessively small (Example 15), core OS1 users survive in the long run.

On the other hand Examples 16, 17 and 18 suggest the disadvantage of a purpose-built product. Examples 16 and 17 show the dynamics from the same initial condition. Example 16 is the case where OS1 is an all-purpose operating system like OS2; the number of the initial OS1 users is so large that OS1 monopolies the market in the long run. Example 17 is the case where OS1 is a purpose-designed operating system. The existence of users who can only find marginal benefit from using OS1 prevents OS1 from monopolising the market. To monopolise the market, the number of the initial OS1 users must be much larger; see Example 18.



Example 1.





Example 2.





Example 3.




Example 4.





Example 5.





Example 6.





Example 7.





Example 8.



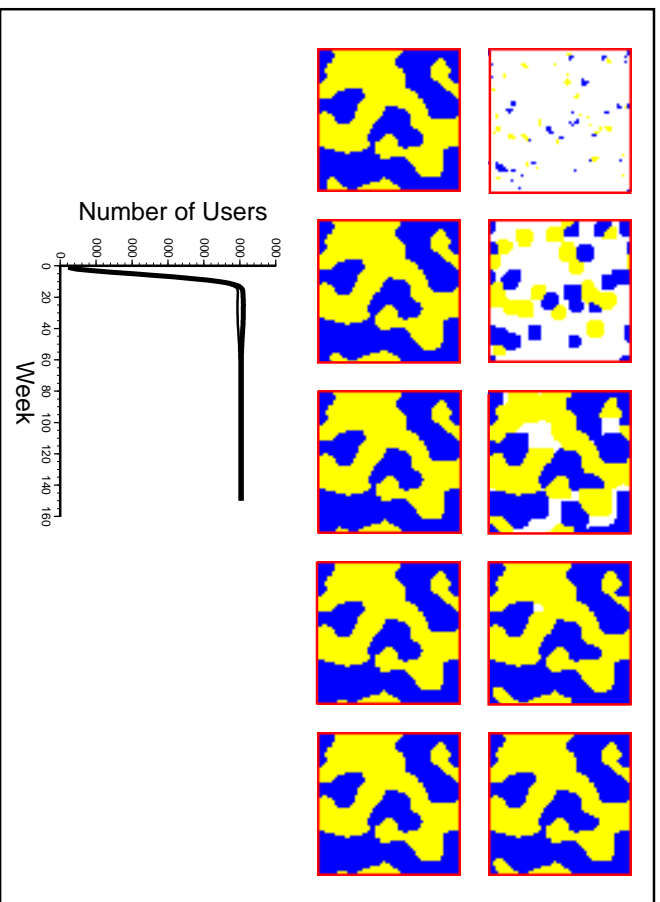
Example 9.



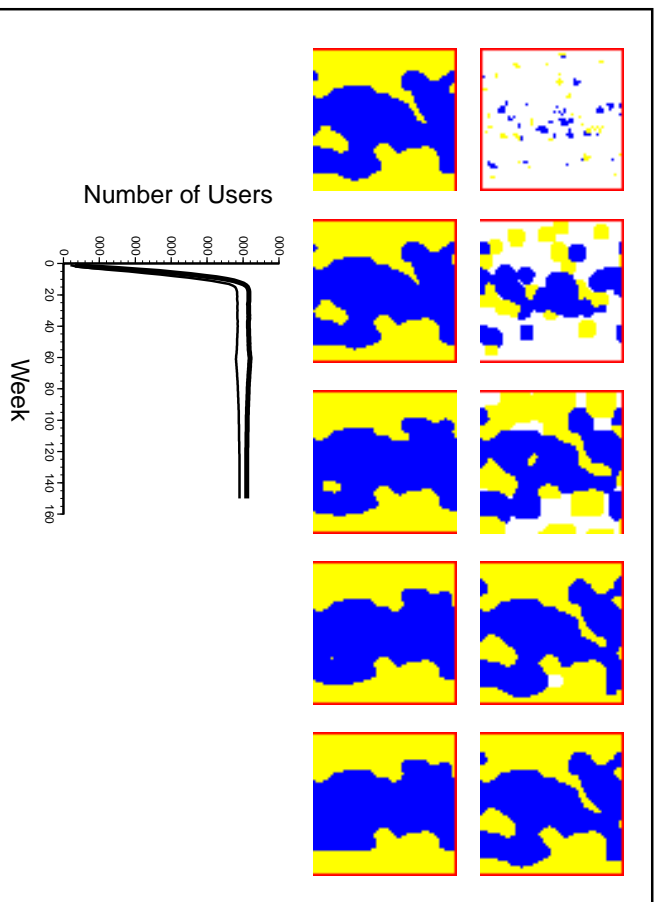
Example 10.



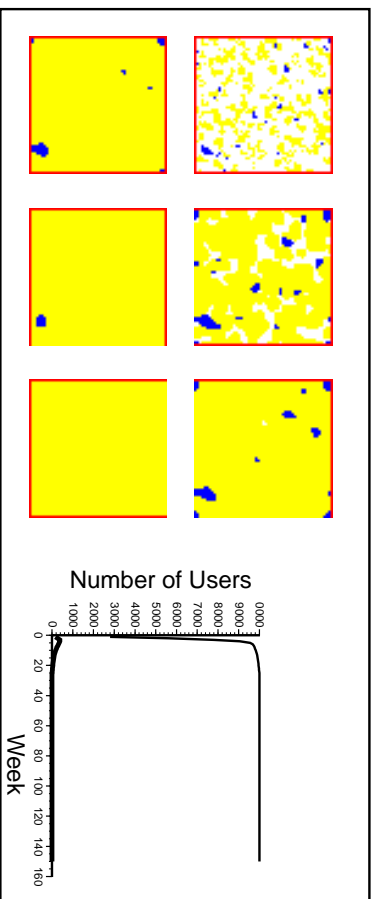
Example11



Example12



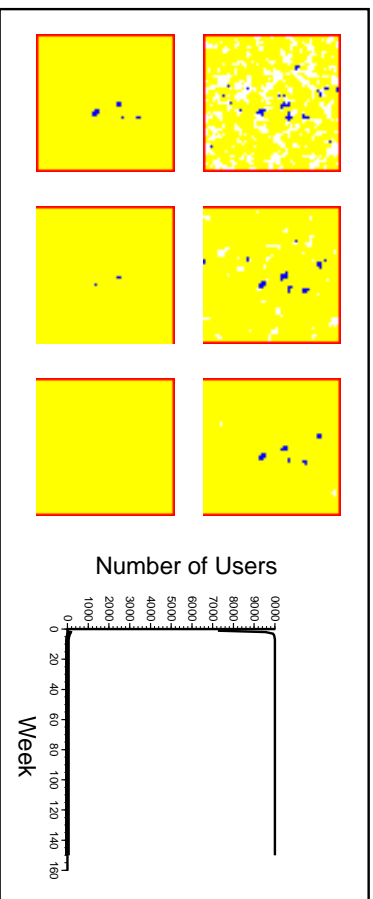
Example13



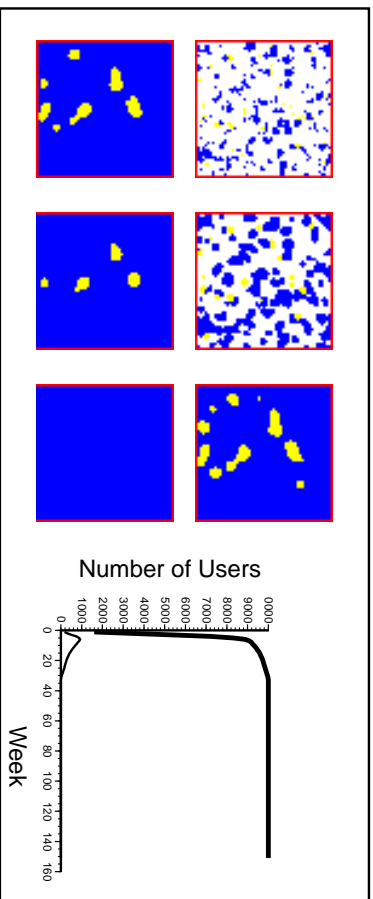
Example14



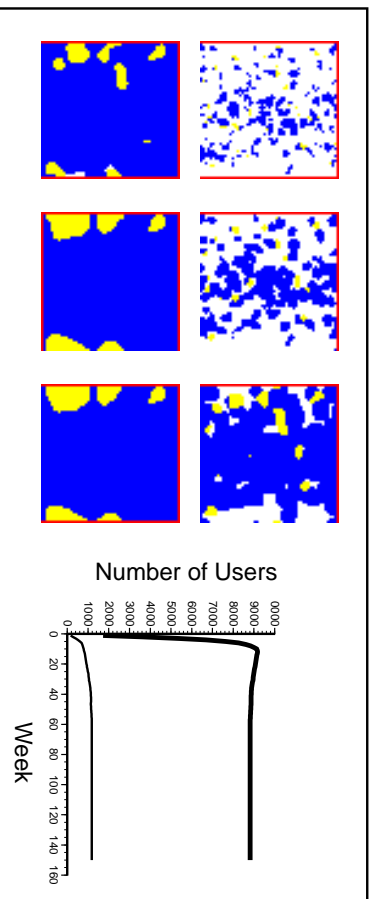
Example15



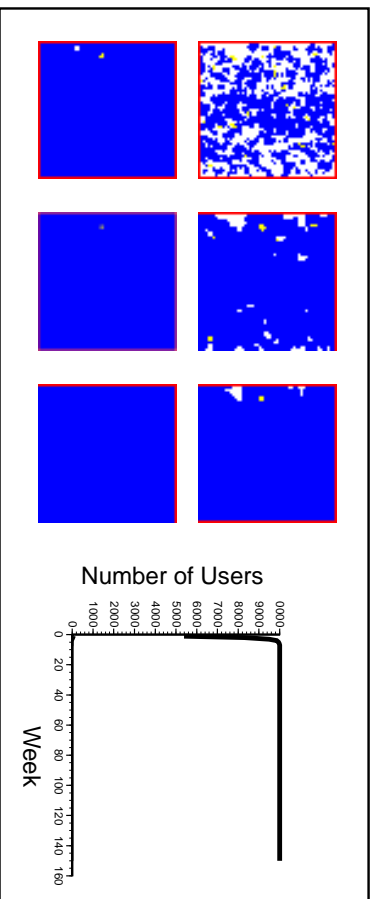
Example16



Example17



Example18



Econolab

Russell K. Standish
 High Performance Computing Support Unit
 University of New South Wales
 Sydney, 2052
 Australia
 R.Standish@unsw.edu.au
<http://parallel.hpc.unsw.edu.au/rks>

Abstract

The paper presents an evolutionary dynamic economic model based on von Neumann's work. Von Neumann's balanced growth solution is shown to be the long time average behaviour of the dynamic system, and that growth is found to be constrained by the productivities of key processes, called the *generators* of the economy. Some discussion is given about how to implement this model using the Ecolab modeling tool.

1 Introduction

In [8], I argue for an evolutionary economics model similar to the Ecolab model described in [10, 7]. Whilst current state of the art economic models are static equilibrium models, in recent times dynamic models incorporating nonequilibrium economics and chaos have become feasible. However, these models have a large number of fixed (and possibly unknowable) coefficients that make them problematic. By allowing for the evolution of these coefficients, one should be able to qualitatively model aspects of real economies. The contributions that evolutionary thinking can make to economic theory have been expanded in a sequel to [8], due to appear shortly[9]. This paper introduces the model defined in that paper, and then links it to the Ecolab software.

2 Building an Economic Dynamics

Many inferential similarities can be drawn between the biological evolutionary model of Ecolab and the processes of a capitalist economy. The obvious analogy for a biological species is a product, and for Darwinian evolution the process of technological change. I consider a model economics (*Econolab*) based on the insights of von Neumann, one of the founders of complexity theory, who introduced *von Neumann Technology* in the late 1930s[11, 12]. In this model economy, there is a set of *commodities* labeled $i \in \mathcal{N} = \{1 \dots N\}$, and a set of *technologies* or *processes* labeled $m \in \mathcal{M} = \{1 \dots M\}$. Each process has an activity z_m , input coefficients a_{mi} and output coefficients b_{mi} , such that in one time step, $a_{mi}z_m$ of commodity i (amongst others) is consumed to produce $b_{mj}z_j$ of commodity j (amongst others). The coefficients a_{mi} and b_{mi} may be zero for some values of m and i , corresponding respectively to processes that do not require a particular input, or do not produce a particular output. This differs from von Neumann's original approach, and is more in line with that of Kemeny, Morgenstern and Thompson[6]. Blatt[2] gives a good introduction to this model, discussing its flexibility in dealing with a range of economic processes. In the words of Blatt (p67):

The von Neumann work is a great achievement of mathematical model building in dynamic economics. It is the best available theory of capital and of rate of return.

To relate this work back to the Ecolab ecological model, the input/output coefficients a_{mi}/b_{mi} are fixed like the r_i, β_{ij} in the Ecolab model, and z_m is a free variable like n_i . In von Neumann's work, the dynamics is imposed in the form of an exponential growth condition:

$$z_m(t+1) = \alpha z_m(t) \quad \forall m \in \mathcal{M} \quad (1)$$

Rather than assuming a particular form for the dynamics, we should be looking for a first order differential equation (or its difference equation equivalent) that describes the dynamics. Consider the capital K_m associated with process m . The rate of change of this capital may be written:

$$\dot{K}_m = z_m \left(\sum_{i=1}^N b_{mi} p_i - \sum_{i=1}^N a_{mi} p_i \right) \tag{2}$$

assuming the market clears perfectly, where p_i is the price of commodity i . This has introduced two new sets of free variables K_m and p_i , for which we need to find closure relations. Clearly, activity is limited by the availability of capital (we do not allow the possibility of credit here):

$$\sum_{i=1}^N a_{mi} p_i z_m \leq K_m \tag{3}$$

For simplicity, let us assume that each process invests a fixed proportion of its capital into production, i.e.

$$\sum_{i=1}^N a_{mi} p_i z_m = \kappa_m K_m, \quad \exists \kappa_m : 0 < \kappa_m \leq 1 \tag{4}$$

Substituting (4) into (2) gives

$$\dot{z}_m = \kappa_m \left(\frac{\sum_{i=1}^N b_{mi} p_i}{\sum_{i=1}^N a_{mi} p_i} - 1 \right) z_m. \tag{5}$$

If price is a fixed quantity (as assumed in von Neumann theory) then (5) is equivalent to the *ansatz* (1).

Now the price acts like a regulator in the economy. Here I draw inspiration from the thermostat of Nosé and Hoover[3] that regulates the temperature of a non-equilibrium steady state system in a heat bath:

$$\dot{p}_i = \pi_i \left(\frac{\text{demand}}{\text{supply}} - 1 \right) p_i = \pi_i \left(\frac{\sum_{m=1}^M a_{mi} z_m}{\sum_{m=1}^M b_{mi} z_m} - 1 \right) p_i. \tag{6}$$

von Neumann assumes that demand never exceeds supply, and if supply exceeds demand (i.e. a surplus), then the commodity is free ($p_i = 0$). This would imply $\dot{p}_i = 0$, freezing prices. In effect this makes the system very stiff — equation (6) softens the dynamics with π_i controlling the stiffness.

This does beg the question of how demand can exceed supply, yet still be satisfied. It is analogous to a species population density exceeding the ecosystem's carrying capacity in Ecolab. The above model really presupposes the existence of reservoirs of commodities (rather like the EU grain and beef mountains), that can also take care of surpluses, in an analogous way that the heat bath acts as a source and sink of heat in a thermodynamic system. This may not be true of all markets, but is certainly valid in some. The parameters π_i scale between the situation of large reservoirs which would have soft dynamics and the situation of small to non-existent reservoirs which would have stiff dynamics.

Kemeny, Morgenstern and Thompson assume that $\forall i, \exists m : b_{mi} > 0$, i.e. every commodity has a process that produces it. I want to remove that restriction in the following way. Let there be a set of commodities $\mathcal{R} \subset \mathcal{N} : \forall i \in \mathcal{R}, \forall m \in \mathcal{M}, b_{mi} = 0$ (called resources) which have a finite reserve r_i , and a finite renewal rate s_i . For example iron ore has effectively zero renewal rate, whereas rainwater does not. Then clearly

$$r_i(T) = r_i(0) + \int_0^T s_i - \sum_m a_{mi} z_m dt \tag{7}$$

To include the effect of bounded resources, alter equation (6) to read

$$\dot{p}_i = \pi_i \left(\frac{\text{demand}}{\text{supply}} - 1 \right) p_i = \pi_i \left(\frac{\sum_{m=1}^M a_{mi} z_m}{r_i} - 1 \right) p_i, \quad \forall i \in \mathcal{R}. \quad (8)$$

When r_i is large, the effective resource price drops to zero, as it does in current economies. Each resource will need to be processed (eg iron ore needs to be extracted from the ground, and shipped to the smelters), this then corresponds to a process z_m that has $a_{mi} > 0$. As extraction of the resource starts to exceed availability, the price rises dramatically and the only thing that will prevent the escalation of prices to infinity is that another resource (with a more expensive process — eg mining asteroids in the case of iron ore, or towing icebergs from Antarctica in the case of freshwater) will become economically viable, setting a new price “equilibrium”.

3 The Balanced Growth Solution

Consider initially a fully connected economy. This means that ultimately, every commodity depends on every resource $i \in \mathcal{R}$. To make this more precise, let $A(\mathcal{S} \subset \mathcal{N}) = \{j : \sum_m a_{mj} b_{mi} > 0, \forall i \in \mathcal{S}\}$ be the set of inputs required to make a set \mathcal{S} of commodities. Then a fully connected economy is one where $\exists n \in \mathbb{N} : \forall i, A^n(\{i\}) = \mathcal{R}$. We can then replace this economy with its n th iterate, i.e. one where each process $m \in \mathcal{M}$ now has input coefficients

$$a_{mi}^{(n)} = \sum_{\substack{m_1, \dots, m_n \in \mathcal{M} \\ i_1, \dots, i_n \in \mathcal{N}}} a_{mi_1} b_{m_1 i_1} a_{m_1 i_2} b_{m_2 i_2} \cdots a_{m_n i}.$$

Then $\forall i \in \mathcal{N} \setminus \mathcal{R}, \forall j \in \mathcal{R}, \sum_{m \in \mathcal{M}} b_{mi} a_{mj}^{(n)} > 0$. Since ($j \in \mathcal{R}$) each have finite production rates, so does every commodity i , given by the minimum over $\sum_m b_{mi} a_{mj}^{(n)} \dot{r}_j$.

The only way to obtain a perpetual growth economy (with fixed numbers of commodities and processes) is to assume à la Kemeny, Morgenstern and Thompson that $\mathcal{R} = \emptyset$. Then there must be *generators* of the economy, i.e. sets $\mathcal{G} \subset \mathcal{N}$ so that $\forall i, j \in \mathcal{G}, \exists n \in \mathbb{N} : i \in A^n(\{j\})$. In words, this is saying that \mathcal{G} is a completely self-sufficient subeconomy, it produces all its own inputs. The completely connected case corresponds to asserting that there is a unique generator \mathcal{G} .

For convenience, write equations (5) and (6) in the following form:

$$\dot{z}_m = z_m \beta_m \quad (9)$$

$$\dot{p}_i = p_i \rho_i \quad (10)$$

with

$$\beta_m = \kappa_m \left(\frac{\sum_{i=1}^N b_{mi} p_i}{\sum_{i=1}^N a_{mi} p_i} - 1 \right) \quad (11)$$

$$\rho_i = \pi_i \left(\frac{\sum_{m=1}^M a_{mi} z_m}{\sum_{m=1}^M b_{mi} z_m} - 1 \right) \quad (12)$$

The long time behaviour of the system can be found by a time average

$$\lim_{T \rightarrow \infty} \frac{\ln z_m(T)}{T} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \frac{\dot{z}_m}{z_m} dt = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \beta_m dt = \beta'_m \quad (13)$$

$$\lim_{T \rightarrow \infty} \frac{\ln p_i(T)}{T} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \frac{\dot{p}_i}{p_i} dt = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \rho_i dt = \rho'_i \quad (14)$$

Let $\mathcal{Z} = \{m : \beta'_m \geq \beta'_l \ \forall l \in \mathcal{M}\}$ be the set of maximal growth rates, and $\mathcal{P} = \{i : \rho'_i \geq \rho'_j \ \forall j \in \mathcal{N}\}$ be the maximal inflation rate. Now if $b_{li} > 0 \ \exists l \in \mathcal{Z}$, then

$$\rho'_i = \pi_i \left(\frac{\sum_{m \in \mathcal{Z}} a_{mi} \hat{z}_m}{\sum_{m \in \mathcal{Z}} b_{mi} \hat{z}_m} - 1 \right) < \infty, \quad (15)$$

where $z_m(t) \rightarrow \hat{z}_m e^{\beta'_l t}$. However, should $b_{mi} = 0 \ \forall m \in \mathcal{Z}$, and $a_{li} > 0 \ \exists l \in \mathcal{Z}$, then

$$\rho_i = \pi_i \left(\frac{\sum_{m=1}^M a_{mi} z_m}{\sum_{m=1}^M b_{mi} z_m} - 1 \right) \rightarrow \pi_i \left(\frac{z_l \sum_{m \in \mathcal{Z}} a_m}{\sum_{m \notin \mathcal{Z}} b_{mi} z_m} - 1 \right) \rightarrow \infty. \quad (16)$$

So p_i grows “superexponentially” in this case. Substituting this back into equation (11), we find

$$\beta'_l = \kappa_l \left(\frac{\sum_{i \in \mathcal{P}} b_{li} p_i}{\sum_{i \in \mathcal{P}} a_{li} p_i} - 1 \right) = -\kappa_l, \quad (17)$$

i.e. the economy collapses under the weight of hyperinflation of its inputs. This is an uninteresting solution as this situation can only happen transitorily. The complement of the above cases is simply that $a_{mi} = b_{mi} = 0 \ \forall m \in \mathcal{Z}$. In this situation, the processes of greatest growth are entirely disconnected from the commodities of interest in the economy, so one should start again using $\mathcal{M} \setminus \mathcal{Z}$ as the set of processes.

Exactly equivalent relations exist for ρ'_i , i.e. if $a_{mi} = 0 \ \forall i \in \mathcal{P}$, and $b_{mj} > 0 \ \exists j \in \mathcal{P}$, then process m diverges superexponentially, causing a price crash for the whole economy. However, if $a_{mj} > 0 \ \exists j \in \mathcal{P}$ then

$$\beta'_m = \kappa_m \left(\frac{\sum_{i \in \mathcal{P}} b_{mi}}{\sum_{i \in \mathcal{P}} a_{mi}} - 1 \right) < \infty, \quad (18)$$

This result, that the inputs for processes in the maximal growth set must be contained within the outputs can be expressed as $A(\mathcal{Z}_o) \subset \mathcal{Z}_o$, where $\mathcal{Z}_o = \{i \in \mathcal{N} : b_{mi} > 0, \exists m \in \mathcal{Z}\}$. Now $\forall i \in \mathcal{Z}_o$, the sequence $\{i\}, A(\{i\}), \dots, A^n(\{i\})$ are all subsets of \mathcal{Z}_o , where $n > N$. However, $\exists n_1, n_2 : A^{n_1}(\{i\}) \cap A^{n_2}(\{i\}) \neq \emptyset$ as there are only N elements to choose from, so $\mathcal{Z}_o \cap \mathcal{G} \neq \emptyset$. Furthermore, since $\forall i, j \in \mathcal{G}, \exists n : i \in A^n(\{j\})$, then $\mathcal{G} \subset \mathcal{Z}_o$, i.e. all goods in the economy generator are produced from processes at the maximum growth rate. Since each commodity is connected to the generator, the whole economy has a unique growth rate, i.e. the time average of equations (5) and (6) is the von Neumann solution.

This result can be readily generalised to incompletely connected economies, by positing a number of generators $\mathcal{G}_1, \dots, \mathcal{G}_n$, each with its own growth rate. Then each commodity grows with the rate given by the minimum growth rate of all the generators it is connected to. This is the KMT solution.

Because of the averaging property (13), Jansen’s condition for permanence[4] can be applied to each generator \mathcal{G}_i , with the centre of motion transformed away by $z_m = \tilde{z}_m e^{\beta t}$ and $p_i = \tilde{p}_i e^{\rho t}$ where β and ρ are the unique growth rates for the generator. Permanence for the economy as a whole is the logical and of permanence for each generator, and having a positive return for each process m ($\beta_m > 0 \ \forall m$).

4 Adding Evolution

Now that we have an economic dynamics established, we need to consider how to graft on an evolutionary process. By direct analogy with Ecolab, it is clear that when a process exhausts its capital ($K_m = 0$), it forever remains that way, so this is equivalent to extinction in ecosystems. Adding new processes and commodities is conceptually easy. Blatt p57–58[2]:

What about technological progress? This can be included by assuming that the list of activities $m = 1, 2, \dots, M$ is not final, but new activities may be invented and hence become

available for use, as time goes on. This makes the total number of processes a function of time: $M = M(t)$. There is no need to remove obsolete processes from this list, since such processes may run at zero activity level. Although this way of handling technological progress exists in principle, we are not aware of any actual theoretical work making use of this idea. Von Neumann himself developed his theory on the basis of an unchanged technology (all input coefficients, output coefficients and the number of processes M are constant in time), and his successors have done the same. The inclusion of technological progress appears to us to be a highly interesting avenue for further exploration.

The difficulty is deciding how to choose new coefficients a_{mi}, b_{mi}, κ_i and π_i when a new process is added. There is no genotype of a process — the closest thing to it is Dawkins’s meme, and there is no genetic algorithm theory of the meme. Clearly new processes arise evolutionarily, with the new processes modeled on the old. The new coefficients will be varied randomly about the old values according to some kind of central distribution.

Recent results from Ecolab indicate that the emergent dynamics of the system is rather insensitive to the specific type of mutation algorithm chosen. Work is currently under way to classify exactly what effects different assumptions make.

In 1962, Arrow[1] pointed out that the cost per unit for production of an artifact falls as an inverse power of the number of units produced:

$$\text{cost/unit} \propto N^{-a}$$

This power law is most likely a consequence of the statistical properties of the underlying “fitness” landscape, as it can be seen in Kauffman NK model[5]. Presumably an evolutionary algorithm that searches process (and commodity) space according to the same power law would be optimally matched to generating change. However, it can also be pointed out that large changes of process are likely to cost proportionally more than smaller changes. As any research budget is finite, the distribution of process improvements must be finitely integrable (have a finite area underneath the curve), which the power law distribution is not, but the normal (Gaussian) distribution is.

5 Implementation in Ecolab

The dynamical equations are relative easy to implement. There are now two sparse matrices \mathbf{a} and \mathbf{b} , that are no longer square, but their product $\mathbf{a}*\mathbf{b}$ is square. Introducing a transpose operator \mathbf{tr} (which trivially swaps `row` and `col`, the dynamical equations can be implemented in the following lines of code:

```
dz = kappa * ( (a*p)/(tr(b)*p) - 1 ) * z;
dp = pi * ( (b*z)/(tr(a)*z) - 1 ) * p;
z += dz; p += dp;
```

For the mutation operator, and initial proposal is for processes to spawn off new processes (this could be imagined to be companies spawning new companies, or companies developing new product lines), proportional to their activity (which happens to be proportional to their capitalisation). The new values of κ and π will be generated by taking a value r from some positive distribution, such as an exponential, then evolving the parameters according to

$$\kappa' = \exp(\log(\kappa) + r) \tag{19}$$

Most of the other tools in Ecolab will port easily — the connect plot widget will display the products $\mathbf{a}*\mathbf{b}$ and $\mathbf{tr}(\mathbf{b})*\mathbf{tr}(\mathbf{a})$.

6 Conclusion

Econolab is a well defined economics model that will provide a first step toward understanding how economies evolve. Watch this space!

References

- [1] K. Arrow. The economic implications of learning by doing. *Review of Economic Studies*, 29:166, 1962.
- [2] J. M. Blatt. *Dynamic Economic Systems: A Post-Keynesian Approach*. M. E. Sharpe, New York, 1983.
- [3] W. G. Hoover. *Phys. Rev. A*, 31:1695, 1985.
- [4] W. Jansen. A permanence theorem for replicator and lotka-volterra systems. *J. Math. Biol.*, 25:411–422, 1987.
- [5] Stuart A. Kauffman. *The Origins of Order: Self Organization and Selection in Evolution*. Oxford UP, Oxford, 1993.
- [6] J. G. Kemeny, O. Morgenstern, and G. L. Thompson. A generalization of the von Neumann model of an expanding economy. *Econometrica*, 24:115–135, 1956.
- [7] R. K. Standish. Ecolab: Where to now? In R. Stocker, H. Jelinek, B. Durnota, and T. Bossomeier, editors, *Complex Systems: From Local Interaction to Global Phenomena*, pages 263–271. IOS, Amsterdam, 1996. also *Complexity International*, vol. 3, <http://www.csu.edu.au/ci>.
- [8] R. K. Standish. The role of innovation within economics. In C. Chiarella, S. Keen, R. Marks, and H. Schnabl, editors, *Proceedings of the inaugural conference on Commerce, Complexity and Evolution*. UNSW, Feb 1996.
- [9] R. K. Standish. The role of innovation within economics. In W. Barnett, C. Chiarella, S. Keen, R. Marks, and H. Schnabl, editors, *Commerce, Complexity and Evolution*, volume 11 of *International Symposia in Economic Theory and Econometrics*. Cambridge UP, 1998.
- [10] Russell K. Standish. Population models with random embryologies as a paradigm for evolution. In Russel J. Stonier and Xing Huo Yu, editors, *Complex Systems: Mechanism of Adaption*. IOS Press, Amsterdam, 1994. also *Complexity International*, vol. 2, <http://www.csu.edu.au/ci>.
- [11] J. von Neumann. A model of general economic equilibrium and a generalization of Brouwer’s fixed point theorem. In K. Menger, editor, *Ergebnisse Eines Mathematischen Kolloquiums*. Vienna, 1937.
- [12] J. von Neumann. A model of general economic equilibrium. *Review of Economic Studies*, 13:1–9, 1945.

Crystallisation of Two-Dimensional Cellular Automata

Tomoaki Suzudo

Control and Artificial Intelligence Laboratory
 Japan Atomic Energy Research Institute
 suzudo@clsu3a0.tokai.jaeri.go.jp
<http://www.jaeri.go.jp>

Abstract

The cellular automata (CAs) which asymptotically lead to uniform regular patterns, namely crystalline CAs are discussed. The two-dimensional and two-states-per-cell CAs with five neighbour structure are particularly focused. First, various CA evolution patterns are analysed by the entropy function, and it is suggested that there are two types of CAs with static or chaotic evolution. Analyses using a newly-introduced rule parameter, called μ parameter, ensure that the crystalline CAs appear at the phase transition point between static and chaotic CAs. The observed phenomena are probably considered as a generalised model of liquid/solid phase transition.

1 Introduction

Many systems composed of countless elements in nature have their own temporal and/or spatial patterns, e.g. galaxies, the weather, biological and ecological systems and societies. These patterns appear without explicit pressure or constraint from outside the systems, and such phenomena are called self-organisation. The objective of studies of self-organisation is to understand the underlying properties of such phenomena in general, and has been explored by many scientists, e.g. [1, 2, 3, 4].

Crystallisation from a liquid state, which is a common example of self-organisation, governs physical characteristics of solid-state materials and is of great importance in science and engineering. The traditional explanation of this phenomenon starts from the microscopic laws applicable to their components, the quantum mechanics in this case. However, as seen in Figure 1, it is possible to embody similar dynamics by exploiting dynamical systems artificially defined in the discrete space with a simple rule of evolution. This fact indicates that the essence of crystallisation in general is not quantum force but something else: This paper attempts to describe what this is.

Cellular Automaton (CA) is a common tool for studying such self-organising phenomena, e.g. [5, 6, 7, 8, 9], and was adopted also in this paper. Due to the simplicity of method, CA is expected to yield the results which could be applicable to crystallisation in general. Another advantage of using CAs is easy implementation of the computer algorithm.

This paper discusses two-dimensional CAs causing crystalline patterns from theoretical and empirical perspectives, and focuses on their statistical features to extract some general properties of such phenomena.

2 Notation and the scope of analysis

Consider an integer value $a_{i,j}(t) \in \{0, 1\}$ assigned at a two-dimensional discrete site (i, j) , $i \in \{0, 1, \dots, N\}$ and $j \in \{0, 1, \dots, N\}$, where $t \in \{0, 1, 2, \dots\}$ is a time step. The site value $a_{i,j}$ at the next time step is deterministically given by the mapping

$$a_{i,j}(t) = f[a_{i-1,j}(t-1), a_{i+1,j}(t-1), a_{i,j-1}(t-1), a_{i,j+1}(t-1), a_{i,j}(t-1)], \quad (1)$$

where f is an arbitrary function which specifies the CA rule, hence the value of a given site depends on the last values of five ‘‘neighbour’’ sites. This neighbourhood pattern is called five neighbour structure, and is sometimes referred to as the von Neumann neighbourhood. The total number of neighbour




Figure 1: Crystalline structure created from a random configuration by a particular CA rule.

state patterns is $2^5 = 32$, therefore a CA rule is composed of 32 mapping functions. In other words, there are 32 rule “entries” per rule. In this paper the value-0 state (or the “empty” state) is considered as a quiescent state, which means

$$f(0, 0, 0, 0, 0) = 0, \quad (2)$$

is satisfied, i.e. the value-1 state (or “occupied” state) is never generated only by “empty” neighbours. This requirement is useful when the CA is used to model interacting particle systems. In addition, the CA space is assumed to be isotropic, i.e. the rule is always rotationally and reflectionally symmetric. If a CA rule is defined such that each site value at the next time step is calculated by the “home” site value and the sum of the last value of the other neighbour sites such as

$$a_{i,j} = g[a_{i,j}(t-1), a_{i-1,j}(t-1) + a_{i+1,j}(t-1) + a_{i,j-1}(t-1) + a_{i,j+1}(t-1)], \quad (3)$$

the rule is called an *outer-totalistic rule*, and provides a subset of symmetric CA rules. The Game of Life [10, 11] is an example of the outer-totalistic rule with the nine neighbour structure.




Figure 2: An example of expanding CAs given by the rule outputting 0 except for $g(0, 1) = 1$.

It is surmised that the occupied cells in the background of empty cells always multiply toward the four directions when

$$g(0, 1) = 1 \quad (4)$$

is satisfied, as in Fig. 2 for example (see **Appendix A** for the reason). The analyses in this paper focus on CAs satisfying eq. (4), otherwise the occupied cells may disappear, or may become localised, as in Fig. 3. Such CAs must be eliminated from our scope in advance.

Let us tentatively swap the empty and occupied states in Fig. 2, and consider empty cells expanding into an occupied-state background. The expandable condition for an empty state can be derived by reversing all the input and the output states of eq. (4) such as

$$g(1, 3) = 0. \quad (5)$$

By taking account of this condition it is possible to remove the CAs whose evolution leads to localized empty cells such as in Fig. 4. Together with eq. (4), the condition of eq. (5) ensures that both occupied and empty cells do not become localised. In this way, the two types of cells are mixed with each other so that the CA pattern may become crystal-like, as in Fig. 1.

If eq. (2) is not satisfied, i.e. if

$$g(0, 0) = 1 \quad (6)$$

is satisfied, then the “empty” background as seen in Fig. 2 will change to an occupied-state background one time-step later. As mentioned above, rules that include such a mapping are not considered in this paper, but the reverse of eq. (6)

$$g(1, 4) = 1 \quad (7)$$

does not violate the quiescent condition, and prohibits a large space filled with occupied sites such as in Fig. 4. Therefore, as with eq. (5), eq. (7) can also avoid the localisation of empty cells in the background of occupied cells.

From the discussion above, CA rules satisfying

$$g(0, 1) = 1 \text{ and } \{g(1, 3) = 0 \text{ or } g(1, 4) = 0\} \quad (8)$$

are expected to lead random configurations to become neither homogeneous nor localised patterns of both empty and occupied cells, which was confirmed by empirical studies. Consequently the CAs in this condition are those which can potentially cause a crystalline structure, and are examined in the next section.

3 Analyses

3.1 Statistical measures of the CA

To understand the crystalline structure of CAs, it is useful to look at the entropy associated with their spatial pattern. The information entropy H is defined as

$$H \equiv - \sum_k P^k \log(P^k), \quad (9)$$

where P^k is a probability of the event k [12], and this can be used to evaluate the randomness (or regularity) of statistical variables. Consider four adjacent sites such as (i, j) , $(i + 1, j)$, $(i, j + 1)$ and $(i + 1, j + 1)$; there are $2^4 = 16$ possible patterns for this local patch. The entropy of the spatial pattern of CA configuration H_s at an arbitrary time step τ can be defined by

$$H_s(\tau) \equiv - \sum_k P_s^k(\tau) \log_{16}(P_s^k(\tau)), \quad (10)$$

where $P_s^k(\tau)$ is the probability for a particular pattern of the local patch at the time step τ . Note that a base of 16 is used for the logarithmic function as the entropy assumes a position between nil and unity. Typical CA configurations varying with H_s are shown in Fig. 5. CAs with high H_s have random spatial patterns, whereas those with low H_s have regular patterns, i.e. crystalline structure. Figure 6 shows the distribution H_s for all 768 cases satisfying eq. (8). The figure indicates that low- H_s (i.e. crystalline) CAs are in the minority and most CAs have random spatial patterns.

Similarly to H_s , it is possible to define the entropy associated with the temporal pattern. Consider an arbitrary site (i, j) . There are four mapping patterns to the next time step per site, i.e. $0 \rightarrow 0$, $0 \rightarrow 1$, $1 \rightarrow 0$ and $1 \rightarrow 1$, and assume we sample the mapping pattern of the site during a certain period. From the probability of each mapping pattern, the entropy of the temporal CA development at the site can be defined by

$$H_t^{(i,j)} \equiv - \sum_k P_t^{k,(i,j)}(\tau_0, \tau_1) \log_4(P_t^{k,(i,j)}(\tau_0, \tau_1)), \quad (11)$$

where τ_0 and τ_1 are the respective initial and final time steps of the data collection. The base of 4 is used for the logarithmic function as the entropy assumes a position between nil and unity. By collecting the entropy of all sites the average entropy of the temporal CA development ($H_t(\tau_0, \tau_1)$) can be given. Figure 7 shows the distribution of the CAs on the $H_s - H_t$ space, where the data H_s are

collected at the time step 2000 and those for H_t are collected during 1900 – 2000 time steps. The H_t clearly classifies the CAs into two groups, random and regular time-developing groups. Some points with moderate H_t correspond to CAs which are still transient even after 2000 time steps, and are expected to drop to the low- H_t region soon or later according to their past tracks. The high- H_t CAs always have a high- H_s value, and these are generally recognized as *chaotic*. The H_s of *static* (low- H_t) CAs, on the other hand, have various values and it is difficult to clearly separate crystalline CAs from non-crystalline ones. The graph also indicates that the crystalline structures are stable once they are produced, because all the crystalline (low- H_s) CAs have regular time-development (low- H_t) after the transient period.

3.2 The μ parameter

If a rule is configured such that all cell states never change whatever state the neighbour sites have, that is,

$$g(0, i) = 0 \text{ and } g(1, i) = 1 (i = 0, \dots, 4), \quad (12)$$

then it is expected that the evolution leads to a constant development. On the other hand, if a rule is configured such that all cell states perpetually change whatever state the neighbour sites have, that is,

$$g(0, i) = 1 \text{ and } g(1, i) = 0 (i = 0, \dots, 4), \quad (13)$$

then the CA evolution becomes periodic with the period two. These are the two extreme cases, and in the case of general rules, the mapping may or may not change cell states depending on the neighbour state. As a descriptor of how the cell state is changeable, we consider a parameter uniquely determined for each rule

$$\mu = \frac{2^5 - m}{2^5}, \quad (14)$$

where m is a total number of the rule entries which update the cell state unchanged. The value of eq. (14) is hereafter called the μ parameter. For instance, the μ parameters for the rules of eqs. (12) and (13) are nil and unity respectively. Note that the μ parameter is different from the λ parameter [13].

Figure 8 shows the frequency distribution of the μ parameter for static and chaotic CAs. CAs in the static group have a μ -value near the extrema, and the value for those in the chaotic group is in the central range of μ . Thus the graphs imply that the transition from static to chaotic takes place as the μ parameter changes from μ_0 (or μ_1) to the moderate value, where μ_0 and μ_1 are the smallest and largest values of μ parameter satisfying eq. (8) ($\mu_0 \approx 0.156$, $\mu_1 \approx 0.968$).

Consider CAs with $\mu_0 \leq \mu \leq 0.5$. If a rule has μ close to μ_0 , we can presume that from the definition of μ that there are many kinds of local patterns which remain unchanged through the updating by the rule. Therefore, CA evolutions starting from a random configuration remain random, although these time-developments are regular; these CAs are thought to be located at the high- H_s and low- H_t range in Fig. 7. As the μ parameter increases the number of such constant local patterns decreases, and accordingly the CA's spatial pattern after the transient period becomes less random. At the same time, the position of the CA on the $H_s - H_t$ space (Fig. 7) shifts gradually to the low- H_s and low- H_t range. If μ becomes so large that only a few kinds of local patterns can remain unchanged, these particular patterns must be spread all over the CA space through the continuation of this process, because each local pattern keeps changing unless it fits one of the particular patterns: This is considered to be the mechanism of crystallisation. For this reason, relatively longer transient times can be observed for the crystalline CAs compared to the others. There is a certain point in the μ parameter, say μ_c , at which all constant local patterns disappear. The rules with $\mu \geq \mu_c$ therefore lead to random time development and random spatial configuration, and correspond to the high- H_s and high- H_t group (namely the chaotic group) in Fig. 7. The value of μ_c is, however, not uniquely determined, because the largest μ which leads to the crystalline structure is dependent on the type of patterns.

A similar discussion is possible for the CA rules with $0.5 \leq \mu \leq \mu_1$. The crystalline structure in this case is determined not by the constant local pattern but by the period-2 local pattern.

Figure 9 shows the $\mu - H_s$ distribution of CAs with $H_t < 0.1$, namely static CAs. Crystalline CAs have moderate values of μ , while the non-crystalline ones have values near the extrema. These results also support the above scenario.

To sum up, the μ parameter statistically parameterises the CA rules, and is a key value of the crystallisation and the phase transition between static and chaotic performances. As mentioned in [13], the λ parameter is related to the four Wolfram Classes [6]:

1. Evolution leads to a homogeneous configuration.
2. Evolution leads to a set of separated simple stable or periodic structures.
3. Evolution leads to a chaotic or aperiodic pattern.
4. Evolution leads to a complex localized structure, sometimes long-lived.

Because the two parameters are independent, these combinations can give more detailed classifications. As already mentioned in section 2, the classes 1,2 and 4 are eliminated from the scope of this study, and the above analysis suggests that the remaining CAs are classified into three groups: Two more classes (crystalline and noncrystalline-and-static classes) can be added. These six classes are roughly located in the landscape of the rule space drawn on the $\lambda - \mu$ plane as seen in Fig. 10.

It is worthwhile to compare such CA dynamics to crystallisation in nature. The static and chaotic classes correspond to solid and liquid phases, respectively; the μ parameter corresponds to temperature. The initial condition is “liquid” as shown in Fig. 1. The “liquid” is cooled rapidly when μ is low, and a non-crystalline structure (or amorphas) is produced. If the “liquid” is slowly cooled at around melting point, i.e. at μ_c , crystal is produced. The initial “liquid” continues to be liquid if μ is high. The major difference is only that the μ_c is not uniquely determined like the melting point. Consequently, the macroscopic properties of the crystallisation in CAs is fairly similar to those in nature. The microscopic mechanism of crystallisation is commonly understood to be that existing crystalline structures allow a local pattern (a molecular) join if it is matching in structure, whereas those unsuitable (impurities) are likely to remain in solution. This is clearly similar to the mechanism of CA crystallisation as explained above.

4 Concluding remarks

Two-dimensional CAs causing the crystalline structure were studied. The spatial and temporal pattern entropy functions successfully characterise them, and classified them into two groups, static and chaotic. The background for this classification was investigated using the μ parameter, the probability for a cell-state to remain unchanged at the next time step; the μ parameter is uniquely determined by the rule definition. Theoretical and experimental studies confirmed that the μ parameter statistically parameterises the CA rules, and that it is a key value to the phase transition between static and chaotic performances. In addition, it was surmised that crystallisation appears when the rule has a μ parameter near the critical point between static and chaotic CA evolutions. The observed phenomena are probably considered as a generalised form of crystallisation, including ones in nature and computers.

A The expandability of CAs

One-dimensional CAs with 3-neighbour structure are defined by

$$a_i(t) = f_1[(a_{i-1}(t-1), a_i(t-1), a_{i+1}(t-1)], (a_i \in \{0, 1\}, i \in Z) \quad (15)$$

and rules satisfying the quiescent condition considered as

$$f_1(0, 0, 0) = 0, \quad (16)$$

and

$$f_1(1, 0, 0) = f_1(0, 0, 1) = 1. \quad (17)$$

Let us assume an initial condition of a single occupied site at the position of $i = 0$ with the background of empty sites. When $t = 1$, two cells at the position of $i = 1$ and -1 are occupied. Similarly, two cells at $i = 2$ and -2 become occupied when $t = 2$. Generally, when $t = \tau$, the sites at $i = \tau$ and $-\tau$ have an occupied state, and a “triangle” is formed in the spatio-tempo CA space, as seen in Fig. 11. The sites inside the triangle have either an occupied or an empty state, but it is never totally empty because an isolated occupied cell at $i = \tau$ again becomes another “seed” of the triangle. In this way, occupied cells spread all over the CA space, and this property is not influenced by the other entries such as $f_1(1, 1, 0)$. The evolution from the random configuration also leads to expansive dynamics as seen in Fig. 11a. Therefore the condition of eq. (17) ensures the expandability of the 1-dimensional CA. Equation (4) is the extension of this condition to two dimensions and ensures the expanding dynamics of 2-dimensional CAs.

References

- [1] Haken, H. “Synergetics -An Introduction”, 1977 *Springer-Verlag, Berlin*.
- [2] Nicolis, G. and Prigogine, I. “Self-organization in Nonequilibrium systems”, 1977 *Wiley, New York*.
- [3] Nicolis, G. and Prigogine, I. “Exploring complexity”, 1989 *R. Piper GmbH and Co. KG Verlag, Munchen*.
- [4] edited by Nijhout, H. F., Nadel, L. and Stein, D. L. “Pattern formation in the physical and biological sciences”, 1997 *Addison-Wesley Publishing Company*.
- [5] Wolfram, S. “Universality and complexity in cellular automata”, 1983 *Rev. Mod. Phys.*, **55**, **3**, 601–644.
- [6] Wolfram, S. “Statistical mechanics of cellular automata”, 1984 *Physica D*, **10**, 1–35.
- [7] Packard, N. and Wolfram, S. “Two-dimensional cellular automata”, 1985 *J. Statistical Physics*, **38**, **5/6**, 901–946.
- [8] Dewdney, A. K. “Computer recreations”, 1989 *Sci. Am.*, **261 August**, 102–105.
- [9] Fisch, R., Gravner, J. and Griffeath, D. “Cyclic cellular automata in two dimensions”, 1990 *in: T.E. Harris Festschrift (Birkhauser, Basel)*.
- [10] Gardner, M. “Mathematical games”, 1970 *Sci. Am.*, **223 October**, 120–123.
- [11] Gardner, M. “Mathematical games”, 1971 *Sci. Am.*, **224 February**, 112–117.
- [12] Korn, G.A. and Korn, T. M. “Mathematical handbook for scientists and engineers”, 1968 *McGraw-Hill Book Company*.
- [13] Langton, C. G. “Computation at the edge of chaos: phase transitions and emergent computation”, 1990 *Physica D*, **42**, 12–37.




Figure 3: The evolution leads to localized occupied cells: The rule outputs 0 except for $g(0,3) = g(1,1) = g(1,2) = 1$.





Figure 4: The evolution leads to localized empty cells: The rule outputs 0 except for $g(0,1) = g(0,2) = g(1,2) = g(1,3) = g(1,4) = 1$.

Figure 5: Typical CA patterns with various H_s .




Figure 6: Distribution of the entropy of spatial pattern of CAs after the long term evolution




Figure 7: Distribution of the CAs on the $H_s - H_t$ space.




Figure 8: The μ parameters of static ($H_t < 0.3$) and chaotic ($H_t > 0.7$) CAs.




Figure 9: H_s of static CAs varying with the μ parameter




Figure 10: The landscape of the rule space characterized by the μ and λ parameters: The graph in the range of $0.5 < \mu < 1$ ($0.5 < \lambda < 1$) is produced by the symmetric transformation of that for $0 < \mu < 0.5$ ($0 < \lambda < 0.5$) at $\mu = 0.5$ ($\lambda = 0.5$).




Figure 11: One-dimensional expandable CAs given by the rule outputting 0 except for $f_1(1, 0, 0) = f_1(0, 0, 1) = 1$.

Agent based Iterated Multiple Lake Game with Local Governments

Tomohisa Yamashita , Keiji Suzuki and Azuma Ohuchi
Lab. of Harmonious Systems Eng.,
Research Group of Complex Systems Eng.,
Graduate School of Eng.,
Hokkaido University, North-13, West-8,
Sapporo 060-8628, Japan
tomohisa@complex.eng.hokudai.ac.jp
<http://ses3.complex.eng.hokudai.ac.jp>

Abstract

In this paper, we extend the environment management game, “The Lake”, to the agent based iterated game. In the original game, the factories as players can make coalitions to save the cost for water purification. Although the whole coalition formation is a desirable solution in the original game, the rational factories fail to make whole coalitions and some of them become free loaders. To discuss more realistic situations and dynamic coalition formative processes, we introduce a new game model based on multiple lakes with local governments as new players. The local governments will affect the factories by levying taxes and imposing penalties. Although the objective of the local governments to maximize revenues, the competition between local governments cause dilemmas to restrain their selfish behaviors. Therefore, the governments have to adjust the tax and penalty adaptively according to the number of factories at their lakes. Under the adjustments of tax and penalties, the factories will be expected to make coalitions. In the experimental results, the evolutionary acquired strategies of these players are shown to produce the complex coalition formation processes in iterated games. Especially, we compared the single lake model and the multiple lake model to discuss the effects of the government dilemmas.

1 Introduction

In this paper, we propose the agent based iterated game of coalition formation for discussing what kinds of local interactions of agents produce complex formation processes according to their situations. In game theory, a coalition is a subset of players that has a binding agreement concerning their strategies. While in agent based game simulations, cooperative or competitive relations between agents develop, as in the prisoners’ dilemma[1] or the market game [4], coalition based game simulations have not received attention until now. The relationship of coalition between agents will be expected to be a more complex phenomenon and require more complex decision processing in the agents.

However, the concept of coalition in game theory is mainly focused with its structure and the sharing of payoff according to rational decisions in static environments. Thus, there is no game theory model based on interaction of agents forming coalitions.

To develop the coalition formative game based on agents, we extend “The Lake” into an iterated multiple lake game with local governments. “The Lake”[5] is known as one of the environment management games. In the original game, the factories as players wish to make coalitions to decrease the cost of water purification. This cost will change according to the number of factories who treat the used water before discharging them into the lake. However, if some factories make the coalition to treat the used water, other factories can decrease the cost without treating the water. Therefore, it is hard to make the whole coalition without preventing some factories becoming free-loaders.

Our interest is in what kind of interactions will prevent factories becoming free-loaders and how to adjust such interactions autonomously. Therefore, we introduce local governments as new players and multiple lakes for factories to select for their operations. The local government in each lake levies a tax on all factories around the lake and imposes a penalty on those in the anti-treating coalition.

Although the objective of local governments is to increase total revenue, the factories also can select the lake to decrease the total cost including the tax and penalty. Therefore, the local governments

can't raise the tax and penalty indefinitely, because they obtain no revenue if there is no factory sitting at their lake. As a result, the local governments are in a dilemma about whether to raise or reduce taxes and penalties, while factories will participate in the coalition that gives them the best payoff in the lake with the lowest pollution rate. Therefore, the interactions among the local governments and factories will be expected to show complex formative processes.

Throughout this proposed game model, the formative process is simulated with the agents as the factories and local governments. To find adaptive strategies based on expected payoff, evolutionary operations are applied to the strategies during each iteration of this game.

In the next section, the original game model, i.e., "The Lake", is introduced. An outline of the extended game model is proposed in section 3. The implementation as an evolutionary iterated game model is described in section 4. Finally, the simulation results are shown in section 5. Section 6 concludes this study.

2 Background

The Lake[5] is one of the environment managing games[3, 6]. In this game, there are m factories around a lake. In order to use the water for their work, they have to purify the water before they use it. The problem for the factories is that they should treat used water before discharging it back into the lake or not, according to the total cost for the purification. Here, it costs an amount B for a factory to treat its wastes before discharging it into the lake. It costs an amount $(m - s)D$ for a factory to purify its own water supply, if $(m - s)$ is the number of factories that do not treat their waste. If we assume that $D < B < mD$, some factories will wish to make a coalition because they can have a chance to decrease their costs by treating the waste water before discharge. Namely, the payoff (cost) for each player is altered by the number of the treating coalition, s , as follows;

$$v(S) = \begin{cases} -smD & \text{if } s < B/D \\ -sB - s(m - s)D & \text{if } s \geq B/D \end{cases} \quad (1)$$

According to this function, the strategy of the anti-treating coalition of $m - s$ factories, named *anti-tc*, is not to treat the waste. While, the treating coalition of s factories, named *tc*, treats its wastes only in the case that they can reduce the cost of purifying their own water supply. The condition whether the treating coalition treats its wastes or not is defined as follows;

$$\begin{cases} \text{treating} & \text{if } s < B/D \\ \text{not treating} & \text{if } s \geq B/D \end{cases} \quad (2)$$

We assume that the players of this game play for purpose of maximizing their own payoff according to their individual rationality. In this case, we can easily imagine that the desired result is for all of the factories to participate in *tc*. However, such desirable result will not appear if a subset of the factories has made *tc*. Namely, if other players are participating in *tc*, a player should participate in *anti-tc* because it will reduce the cost of purifying its own water supply by not treating its waste. So the optimal strategy is to participate in *anti-tc*, in other words free-load the treating coalition. In another case, if all the other players participate in *anti-tc*, a player should also participate in *anti-tc* because it is useless for it to spend for the others. So the optimal strategy is also to participate in *anti-tc*. Therefore the optimal strategy on this assumption is to participate in *anti-tc*. There is no player who participates in *tc* when all players act on the optimal reaction principle. Thus, we can't stop the players becoming the free-loaders and finally all players abandon to treating their waste under this game definition. In order to avoid this undesirable situation, what kind of interactions should be added in this game? Furthermore, our interest is that the adaptive agents as the players can be aware of the strategies for achieving the desirable situation or not, based on their local interactions.

In next section, we propose the extended game for evaluating the coalition formation process based on the agents.




Figure 1: Multiple lake with local governments

3 Multiple Lake Game with Local Governments

In the original game model, it is hard to make the whole coalition and prevent factories from becoming free-loaders. Now, our interest is what kind of interactions will prevent the factories becoming the free-loaders and how to adjust such interactions autonomously to make the whole coalition. Therefore, we introduce local governments as new players and multiple lakes for the factories to site their operations.

3.1 Introduction of local government

In order to prevent factories from becoming free-loaders, “local governments” as new players are introduced into the game. The local government wishes to levy a tax on all factories around the lake and imposes a penalty on the anti-treating coalition.

As a result, the characteristic function of factories is determined by the tax, the penalty and the size of coalition. The balance of these factors will change the decision making of the factories even if they obey individual rationality. The revenue of the local government is the sum of the tax and the penalty levied on the factories. That is, the purpose of the local government is to maximize its own revenue by adjusting the rates of the tax and the penalty.

3.2 Multiple Lakes

If only one local government exists the game, the government will raise the tax and the penalty indefinitely. To avoid this monopoly situation, some kind of relation is required to restrain the selfish activity of the local government. Thus, we introduce multiple lakes in order to introduce competition between the local governments assigned to the lakes. Now, the strategy of the factories is also extended. They can select not only the one of two coalitions, but also the lake to site their operations.

Why does the existence of some local governments cause a competitive situation? If some local governments levy high tax and impose high penalty to obtain high revenue, the factories under those governments will move to other lakes. The local governments hence can't raise the tax and the penalty indefinitely, because they will obtain no revenue if there are no factories selecting their lake. As a result, the local governments are in the dilemma of whether to raise or reduce the tax and the penalty.

To describe the strategy of the factories, we define the whole coalition set as $\{tc_1, anti-tc_1, \dots, tc_n, anti-tc_n\}$.

3.3 Introduction of the pollution rate and the cost of movement

Agent societies have been used as the basis for studies of social and economic behavior. Therefore, it seems unrealistic that the agent society does not include a limit to the natural resources and restraints of behavior by cost.




Figure 2: Iterating process of Multiple Lake Game

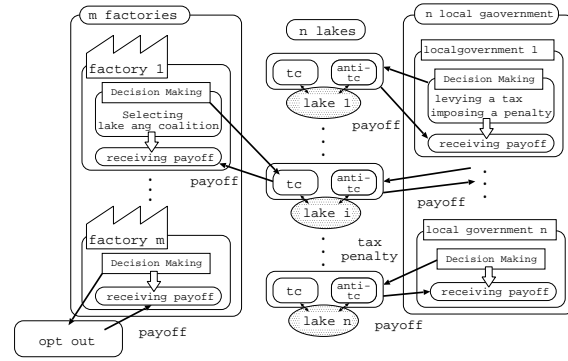


Figure 3: Outline of Agent based Multiple Lake Game

We assume that the natural resources of the lakes are limited to form a more realistic game model. There are two cases in which the pollution rate of one lake goes up. One case is that the size of *anti-tc* is over one constant number. The other case is that the sum of the size of *tc* and that of *anti-tc* is over another constant number. Above mentioned D as the cost of treating its own water in the lake raises according to the pollution rate going up. The pollution rate will fall down according to the size of *tc* and that of *anti-tc* going down.

To reflect reality in this game model, we assume that if factories move to other lake, there will be a constant cost associated with the movements. With this restriction, we expect the inhibition of frequent movements of the factories without the adaptive coalition formative process.

3.4 Introduction of the strategy to *opt-out*

By introducing multiple lakes, local governments are put in a dilemma. Although, if all local governments raise both their taxes and penalties at the same time, the factories can't evade through their selection of the operating place. To protect the factories from such situation, "*opt-out*" as additional strategy for the factories is introduced. Choosing this strategy means not operating on any lake. Namely, the factories in the *opt-out* have to refuse to interact with others, having neither the benefits nor the risks. If a factory chooses this strategy, the player has to pay a constant cost no matter how other players behave.

4 Agent based Implementation

4.1 Outline of agent based Iterated Game

Figure 2 shows the iterating process of the proposed game. Figure 3 shows the outline of agent based multiple lake game.

Two kinds of agents as players exist in this game. One has the role of the factories. A factory agent selects a lake as a site of operations initially. Then, it has to decide which coalition it should participate in. The role of the local government agent is to decide how much tax to levy and impose a penalty against the coalitions on its lake. According to their decisions, each of the agents receives a payoff. This decision making process is iterate until time T . After this process, the evolving phase is applied to find adaptive strategies for each agents. That is, the strategy of each player will be modified by evolutionary operators based on the received payoff. Here, one generation in this game is defined as the series of the decision making processes and the following evolving phase. In each generation, K , we will expect the agents to produce adaptive coalition formations based on their local interactions.

In this game model, the number of the lakes is set to n . In describing the following strategies of the factories, we define the coalition set as follows.

$$\{\{S_i, \overline{S}_i\}, opt-out\} \quad i = 1, \dots, n \quad (3)$$

S_i denotes the treating coalition of lake i . Similarly, \overline{S}_i express the anti-treating coalition of lake i . We define the size of treating coalition as $|S_i|$, the size of anti-treating coalition as $|\overline{S}_i|$.

Pollution rate of lake

To represent the limitation of resources, we introduce the pollution rate of a lake. The concept of the rate may reflect the natural purifying ability of a lake.

In this model, we formulate the pollution rate with the following four elements. One is the average size of a treating coalition in the k th generation, $\overline{|S_i(k)|}$. Another is the average size of an anti-treating coalition in the k th generation which is represented by $\overline{|\overline{S}_i(k)|}$.

The others are the thresholds concerning a limited size of an anti-treating coalition, L_1 , and the threshold concerning a limited sum of factories, L_2 .

The pollution rate is fixed during each generation. Then, the rate is renewed based on these parameters before starting the next generation.

That is, the pollution rate in the k th generation is defined as follows;

$$Po_i(k) = \begin{cases} (100 \times Po_i(k) + 3)/100 & \text{if } \overline{|S_i(k)|} + \overline{|\overline{S}_i(k)|} \geq L_1 \quad \text{and} \quad \overline{|\overline{S}_i(k)|} \geq L_2 \\ (100 \times Po_i(k) + 1)/100 & \text{if } \overline{|S_i(k)|} + \overline{|\overline{S}_i(k)|} \geq L_1 \quad \text{or} \quad \overline{|\overline{S}_i(k)|} \geq L_2 \\ (100 \times Po_i(k) - 1)/100 & \text{if } \overline{|S_i(k)|} + \overline{|\overline{S}_i(k)|} < L_1 \quad \text{and} \quad \overline{|\overline{S}_i(k)|} < L_2 \quad \text{and} \quad Po_i(k) > 1 \\ 1 & \text{if } \overline{|S_i(k)|} + \overline{|\overline{S}_i(k)|} < L_1 \quad \text{and} \quad \overline{|\overline{S}_i(k)|} < L_2 \quad \text{and} \quad Po_i(k) \leq 1 \end{cases} \quad (4)$$

Not only the factories but also the local governments have to adapt to this natural change in the environment, because the pollution rate will affect the cost of purifying the water in the iterated process.

4.2 The strategy and the characteristic function

The strategy and the payoff function of local government

The purposes of the local governments are to maximize their revenues. They have to adjust the tax and the penalty rates because they are caught in a dilemma. To adjust the taxes and the penalties, they raise or reduce from current values. Here $T_i(k, t)$ is the tax value and $Pen_i(k, t)$ is the penalty value levied by the government of the lake i . Namely, the strategies of the local governments change the values with -1 (decreasing), ± 0 and $+1$ (increasing) from the current values. At the t th game in the k th generation, the current tax $\Omega_{T_i}(k, t)$ is set as $\Omega_{T_i}(k, t) = \{T_i(k, t-1) - 1, T_i(k, t-1), T_i(k, t-1) + 1\}$. The current penalty Ω_{Pen_i} is $\Omega_{Pen_i}(k, t) = \{Pen_i(k, t-1) - 1, Pen_i(k, t-1), Pen_i(k, t-1) + 1\}$.

The payoff for the local government assigned the lake i depends on the size of the coalitions in the lake. Here, the size of treating coalition of lake i is represented as $|S_i|$. Similarly, the size of anti-treating coalition is denoted by $|\overline{S}_i|$. The payoff function for the local government is defined as follows;

$$g(T_i(k, t), Pen_i(k, t)) = (|S_i| + |\overline{S}_i|)T_i(k, t) + |\overline{S}_i|Pen_i(k, t) \quad (5)$$

The strategy and a characteristic function of factory

To operate the factories, the factories have to pay the cost for treating and purifying water. The factories can select which lake to operate from, and which coalition to participate in. The strategy set of factory j is represented as

$$\Pi_j = \{\{S_i, \overline{S}_i\}, opt-out\} \quad i = 1, \dots, n \quad (6)$$

Therefore, at the t th game in the k th generation, denote the strategy of the factories j as $\pi_j(k, t)$. Here, we abbreviate $\pi_j(k, t)$ to π_j and describe the complementary vector of strategies:

$$(\pi_1, \dots, \pi_{j-1}, \pi_{j+1}, \dots, \pi_m),$$

as π_{-j} .

According to the participating coalition, the costs of the coalitions are varied. That is, the coalition value depends on not only the size of the coalitions but also the strategies of the local governments. Furthermore, the pollution rate is affected to the value. Therefore, the characteristic function of S_i can be constructed as follows;

$$v(S_i) = \begin{cases} \alpha - |S_i|(|\overline{S}_i| + \beta)Po_i(k)D - T_i(k, t) - Pen_i(k, t) & \text{if } |S_i|/|\overline{S}_i| < B/D \\ \alpha - |\overline{S}_i|(B - |S_i|)D Po_i(k) - T_i(k, t) & \text{otherwise} \end{cases} \quad (7)$$

Similarly, the characteristic function of \overline{S}_i can be defined as follows;

$$v(\overline{S}_i) = \begin{cases} \alpha - |S_i|Po_i(k)D - T_i(k, t) - Pen_i(k, t) & \text{if } |S_i|/|\overline{S}_i| < B/D \\ \alpha - (|S_{2i}|)^2Po_i(k)D - T_i(k, t) - Pen_i(k, t) & \text{otherwise} \end{cases} \quad (8)$$

As sharing method of the coalition value to the members, the payoff of the factories is defined by the coalition value divided with the number of the members. The reason is that all factories are equal partners in the coalitions. Here, an element of a coalition in the set of π is represented as S . Therefore, the payoff function of the factory participating in S is defined as follows;

$$f_j(\pi_j, \pi_{-j}) = \begin{cases} v(S)/|S| & \text{if coalition selected} \\ \text{const.} & \text{if the } opt-out \text{ strategy is selected} \end{cases} \quad (9)$$

4.3 Evaluation of strategy

In this game, the players make decisions on the basis of the continuously updated values, “expected payoffs”, that relate to the reactive decisions in the iterated game and the evolution of the strategies.

The expected payoff is updated as follows; A player has received a payoff from the coalition of which it participates in or levies a tax and imposes a penalty. It has recorded this payoff, and it has updated its payoff count to reflect the receipt of the new payoff.

Suppose, for example, that a player v receives a payoff P . Player v immediately updates its current payoff count N_v , via the assignment statement,

$$N_v \leftarrow N_v + 1 \quad (10)$$

According to the updating of the payoff count, the weight value, w_v , is updated via the assignment statement,

$$w_v \leftarrow \frac{N_v}{N_v + 1} \quad (11)$$

Finally, the expected payoff U_v is updated via the assignment statement,

$$U_v \leftarrow w_v U_v + (1 - w_v)P \quad (12)$$

The expected payoff is used for determining the reactive strategy in the agent. The detail of the determination of the reactive strategy is describe in the following.

4.4 Reactive Strategy with FSM

In this game, all of the players wish to maximize their sum of payoff in the repeated situations. Therefore, the agents have to make the next decision based on the previous results. The previous results are stored as the expected payoff in each agent. In order to determine the next decision based on the expected payoff, we investigate the finite state machines (FSMs) like as the iterated prisoner’s dilemma [2, 4]. Namely, the structure of each FSM represents the strategy of each agent.

The FSMs of player v is represented as $M_v = (I_v, O_v, Q_v, \delta_v, \lambda_v)$.

The input value, I_v , is assigned by current expected payoff, like as Figure 4.




Figure 4: Structure FSM for reactive strategy.

The output value, O_v , is determined by the transition label that consist of a pair as I_v and O_v . The output value denotes the next decision.

The number of states, Q_i , is a constant. The transition function, δ_v , determines the next state from the input value and the current state. The output function, λ_v , determines the output value from the input value and the current state. These two functions are characterized in each player throughout the evolutionary operations based on the expected payoff.

Each factory has two FSMs. One FSM determines a lake that it participates in (or *opt-out*). Another FSM determines which coalition the factory should participate in. Each local government also has two FSMs. One FSM determines whether to raise, remain or reduce current tax. Similarly, the other determines the current penalty.

4.5 Evolution of Strategy with GA

In this paper, the agents try to acquire an adaptive strategy throughout the local interactions. In order to acquire an adaptive strategy, a genetic algorithm (GA) is applied to the FSMs in each agent. The chromosomes in the GA encode the transition function δ_v and the output function λ_v of the FSMs in agents with bit strings. In applied GA, the chromosomes are arranged in torus plane. Namely, an ecological type of GA is used for evolving the strategies. The fitness of the chromosome is just the sum of the agent’s payoffs. As evolutionary operators, we utilize the conventional two-point crossover and the mutation with a certain probability.

5 Simulation

In this paper, we extend the original single lake game to the multiple lake game with the local governments. Namely, we expect that the competition between the governments assigned to the multiple lakes produce a dilemma and prevent unlimited raising of taxes even if the objective of the governments is to maximize their revenues. To confirm this expectation, we observe the transitions of coalition formations, taxes and penalties in the case of one lake and three lakes.

5.1 The case of one lake

The parameters for this simulation are set as follows;




Figure 5: The size of *tc*, *anti-tc* and *opt-out*




Figure 6: The average payoff of all factories




Figure 7: The penalty levied by local government




Figure 8: The size of *tc* and *anti-tc* in lake 1

number of factories	$F = 100$
number of lakes	$L = 1$
number of local governments	$G = 1$
number of generations	Gen= 8000
number of iterations per generation	$I = 15$
cost for treating its wastes	$B = 10$
cost for purifying its own water supply	$D = 2$
coefficient of characteristic function	$\alpha = 600$
coefficient of characteristic function	$\beta = 5.0$
initial penalty	Pen = 0
initial tax	$T = 0$
initial expected payoff	$U_v = 0$
memory weight	$w = 0.7$
initial pollution rate	Po(1) = 1.0
limit of the size of anti-treating coalition	$L_1 = 5$
limit of the sum of factory	$L_2 = 45$
cost of moving	$C_{move} = 20$
payoff of choosing to <i>opt-out</i>	Payoff _{opt} = 40
crossover rate	$R_{crossover} = 0.5$
selection rate	$R_{selection} = 0.5$
mutation rate	$R_{mutation} = 0.05$

The game was simulated with this set of parameter values. The coalition formative processes are shown in Figure 5 to Figure 7. Figure 5 shows the number of *tc* and *anti-tc* and the number of the factories to *opt-out*. Figure 6 shows the average payoff of all factories. Figure 7 shows the penalty levied by government.

The detailed story in this simulation is described as follows;

0 — 100 generations

In the initial phase of the game, the size of *anti-tc* increased because the payoff of *anti-tc* is more than that of *tc*. This situation caused the increase of the pollution rate immediately. In addition, the local government changed the strategy to raise the penalty (Fig.7). Thus, the payoff to the participant in *anti-tc* rapidly decreases in the initial generations (Fig.6).




Figure 9: The size of tc and $anti-tc$ in lake 2Figure 10: The size of tc and $anti-tc$ in lake 3 and the number of $opt-out$ 

Figure 11: The average payoff of all factories




Figure 12: The penalty by levied by local governments

100 — 7000 generations

The remarkable changes from the situation in the initial generations rarely occurred. The exception was only that the local government discharged the high penalty at some times. In this situation, the payoff of $anti-tc$ was close to the payoff of $opt-out$. Therefore, almost all the factories selected $opt-out$ and the size of $anti-tc$ was less than twenty.

5.2 The case of three lakes

The parameters for this simulation are the same as the parameters of one lake model, except for the number of lakes and local governments. These two parameters for this simulation are set as follows;

$$\begin{array}{ll} \text{number of lakes} & L = 3 \\ \text{number of local governments} & G = 3 \end{array}$$

The coalition formation processes in case of three lakes are shown in Figure 8 to Figure 12. Figure 8 to Figure 10 show the progress of the size of tc and $anti-tc$ and the number of $opt-out$. Figure 11 shows the average payoff of all factories. The penalties levied by local governments are shown in Figure 12.

The detailed story in this simulation is described as follows;

0 — 600 generations

At first, the game started with the size of $anti-tc$ increasing, because the payoff of $anti-tc$ was more than that of tc . Following this situation, all local governments immediately raised the penalties. Thus the factories couldn't get the enough payoff even if they participate in any $anti-tc$.

600 — 800 generations

In these generations, we could observe first the size of tc exceeding that of $anti-tc$ in lake 3. According to the successful construction of tc , the average payoff of all factories started to increase. Following this situation, the size of $anti-tc$ in lake 3, namely free-loaders, increased soon. As a result, the size of tc in lake 3 became to zero. We can find that this short story of lake 3 quite resembles the story of the original game model.

800 — 2100 generations

After the successful coalition formation and its collapsing, small *anti-tcs* appeared on each lake and the other factories went to *opt-out*. Only government 3 tried to change the situation by varying penalty (Fig.12). However, this change only caused the participation rate in *anti-tc* to temporarily increased at the expense of *opt-out* (Fig.10). The averaged payoff of all players stayed low during these generations, since the payoff from *anti-tc* and the payoff from *opt-out* was nearly equal (Fig.11).

2100 — 3500 generations

Almost all factories withdrew from lake 1 and lake 2 (Fig.8, Fig.9). They moved to *opt-out* and lake 3. From this period, the factories on lake 3 came to form *tc* gradually. At generation 2700, a whole coalition that all factories participate in is successively formed (Fig.10). Even when the free-loaders appeared around in 3000 generation, the whole coalition recovered and remained during these generations. Therefore, the average payoff of the factories reached its maximum value (Fig.11).

3500 — 4300 generations

In previous generations, while the successful coalition was being formed, the pollution rate of lake 3 grew. Therefore, *tc* collapsed at generation 3500 (Fig.10). From this generation to the end of this period, all factories and governments seem to fall into confusion. For example, *tc* was emerged in lake 3 at one point, however the coalition immediately lost. Then *anti-tc* in each lake and *opt-out* were increased in turn.

4300 — 5900 generation

At end of the previous period, over the half of factories chose to *opt-out*. From this situation, all factories tried to make *tc* in each lake. However, the pollution rate in lake 3 was lower than other lakes, *tc* in lake 3 had grown faster than other *tcs*. Therefore, the *tc* in lake 3 absorbed all factories and a whole coalition formed again. Here, whole coalition formation also means the pollution rate rising rapidly similar to the previous whole coalition formation. Thus, the whole coalition in lake 3 broke down at generation 4700 and half of factories moved to *anti-tc* in other lakes. Although the coalition in lake 3 had halved in size, factories remaining in the coalition kept successfully until generation 5900.

5900 - 7000 generation

At generation 5900, the pollution rate became too high to keep the coalition in lake 3. Therefore, not only *tc* but also *anti-tc* disappeared from lake 3 even if the government 3 was decreasing the penalty. Some of the factories operating on lake 3 moved to other lakes. However, these factories moved again to *opt-out* since the both governments of lake 1 and lake 2 were keeping high penalties to the end of this period. Thus, the number of *opt-out* reached to maximum number in this simulation.

5.3 Discussion

By increasing the number of lakes from one to three, a complex coalition formative process, which includes whole coalition formation, can be observed. This phenomenon does not occur in case of one lake. Therefore, we can confirm that the agent based coalition formative process works well under the competition between the local governments.

In this game model, it is obvious that the average payoff of all factories will increase when *tcs* are formed. Concerning local government, we set that the objective of it is to maximize its revenue by levying a tax and imposing a penalty. Therefore, the average payoff of the local government doesn't depend on whether factories can form *tc* or not. However, not only the average payoff of all factories but also that of local governments increased in this simulation like as follows.

case of one lake	the average payoff	6,067
case of three lakes	the average payoff of government 1	8,977
	the average payoff of government 2	9,740
	the average payoff of government 3	5,398
	the sum of average payoff of three local governments	24,114
	the average payoff of governments	8,038

Why does the payoff of the factories and the governments increase? The increase of the payoff of the factories is obviously caused by tc formed. For the government, the increase of their payoff is mainly caused by the decrease of $opt-out$. In the case of one lake, the government can behave selfishly because it is just a monopoly. In the case of three lakes, the local governments have to compete with each other for getting the factories to operate on their lakes. This competitive situation seems to prevent the governments taking more payoff because they can't get the high revenue by simply raising the taxes and the penalties. If the governments raise the taxes and the penalties indefinitely, the factories on their lakes will move to other lakes and they may not select the lake again. On the other hand, even if the local government reduce the tax and the penalty to collect the factories, the behavior of this government isn't communicated to the factories in other places in this game model. Therefore, the remaining possibility for the government to increase the payoff under competition may be to exploit the dynamics of the coalition formative process around the multiple lakes. The exploitation of the dynamics of the coalition formative process are suggested in the simulation as follows.

In these simulations, we have observed that imposing of high penalty value has no influence on making tc in any situation. The results of two simulations showed that the factories prefer to select $opt-out$, in spite of its payoff being always low, rather than participate in any coalitions. For example, less than twenty factories select $anti-tc$ in the single lake model and others select $opt-out$. This is reason why the payoff of $opt-out$ is same as the payoff of $anti-tc$ when a maximum penalty is imposed by the local government. Therefore the local government in the single lake model will acquire the average payoff 6000 from the small $anti-tc$. To grow the size of coalitions, the payoff for the factories needs to exceed the payoff of small $anti-tc$. In this game model, only in the case of that more than fifty factories participate in tc will the payoff exceed that of $anti-tc$. Therefore the local government in the single lake model has no ability to produce the such a tc by blocking the existence of the certain size of $anti-tc$.

In the case of the multiple lake simulation, similarly, when three local governments keep maximum penalties, less than twenty factories select $opt-out$ and the others participate in $anti-tc$. In this simulation, tc emerged sometimes on the lake 3 when the lake has no blocking coalitions to form tc . Once tc is formed, the size of tc increases rapidly because the payoff of tc is higher than that of the $anti-tc$ which has a high penalty imposed. Therefore, in summary, the dynamics of coalition formative process with sufficient number of lakes can decrease the size of $anti-tc$ as blocking coalitions in each lake and help the emergence of tc .

6 Conclusion

We proposed the game model extending "The Lake" as an agent based iterated coalition formative game. To extend this game model, we introduced "local government" as new player, multiple lakes, the pollution rate, the move cost and the strategy to $opt-out$ into this game.

Throughout the simulations comparing the single lake model and the multiple lake model, we can confirm that the agent based coalition formative process in the multiple lake model worked well. Especially, the dilemma situation of the local government makes them exploit and help the dynamics of coalition formative process to form the desired coalitions.

Further development will be required to understand the complex of coalition formative process based on such type of simulations.

References

- [1] R. Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.
- [2] J. Batali and P. Kitcher. Evolutionary dynamics of altruistic behavior in optional and compulsory versions of the iterated prisoner's dilemma. *Artificial Life*, 5:343–348, 1995.
- [3] G. Hardin. The tragedy of the commons. *Science*, 162:1243–1248, 1969.
- [4] D. McFadzean and L. Tesfatsion. A C++ platform for the evolution of trade networks. Working Paper, Dept. of Economics, Iowa State University, Ames, 1995.
- [5] L.S Shapley and M. Shubik. On the core of economic systems with externalities. *The American Economic Review*, 59, 1969.
- [6] S. Thiagarajan. Garbage: A card game that simulates the trade-off between competition and concern. *Simulation and Games*, 22:112–115, 1991.

Author Index

David M. Alexander, 25
Phil Sheridan, 25
Trond Andresen, 229

Cristobal Baray, 171
Mark A. Bedau, 327
Jacques Blanc-Talon, 91
Paul D. Bourke, 25

Bengt Carlsson, 179
Carl Chiarella, 244
Lee G. Cooper, 336

J.H.B. Deane, 101, 130

P. Flocchini, 110

F. Geurts, 110
H. Randy Gimblett, 188
David G. Green, 32

Eric Halgren, 55
Stephan R.P. Halloy, 118
Xue-Zhong He, 244
Takuo Henmi, 40
Ric D. Herbert, 258

Hisao Ishibuchi, 210, 272
Robert M. Itami, 188
Atsushi Iwasaki, 282
Kouhei Iyori, 346

D.J. Jefferies, 101, 130
Herbert F. Jelinek, 144
Stefan Johansson, 179
Patrick Johnston, 55
Cameron. L. Jones, 144
Shareen Joshi, 327

Taisei Kaizoji, 291
Michael L. Kalish, 20
Steve Keen, 305
MIURA Ken, 346
Nicholas I. Klomp, 32
John Klopp, 55

Otto Konstandatos, 25
Yusuke Koyama, 320
Koichi Kurumatani, 66

Hyen-Yeal, Lee, 150
Gareth D. Leeves, 258
Raymond Lister, 219

Robert E. Marks, 336
David F. Midgley, 336
A. Mingarelli, 110

Mari Nakamura, 66
Tomoharu Nakashima, 272
Valeriy Nenov, 55
Manabu Nii, 210

Sobei H. Oda, 282, 320, 346
Chi-Hyon Oh, 272
Azuma Ohuchi, 376

Jung-Hee, Park, 150

Merton T. Richards, 188

N. Santoro, 110
G. M. Shiraz, 336
J.W. Sleight, 74
Russell K. Standish, 80, 358
D.A. Steyn-Ross, 74
Moira Steyn-Ross, 74
Tomoaki Suzudo, 364
Keiji Suzuki, 376

Kimiko Tanaka, 210
Ashley Tews, 219
Yuping Tian, 162

Kanji Ueda, 282, 346

Matthew D. Warfel, 144
James J. Wright, 25
Andrew Wuensche, 3

Tomohisa Yamashita, 376
Xinghuo Yu, 162

Concept Index

- 1987 crash, 291
- adaptive behaviour, 349
- adaptive control, 162
- agent, 35, 119, 336
- agent based modelling, 171, 188, 327, 346, 376
- allometric species relation, 121
- anaesthesia, 77
- ant colony, 66
- arbitrageurs, 291
- artificial intelligence, 171, 196
- artificial life, 36, 197
- Asian crisis, 305
- attractor, 3, 101, 112, 133, 237, 249, 292
- autecology, 34
- avalanche, 119
- AVS, 83, 258
- Axelrod/Forrest representation, 337

- balanced growth, 358
- bandwagon traders, 230, 292
- bankruptcy, 312
- basin of attraction, 3, 106, 141
- behaviour, 40
- behaviourism, 219
- belief, 327
- belief systems, 171
- benefits-based management, 190
- bifurcation, 101, 106, 244
 - Hopf, 297
- biodiversity, 32
- BioLand, 172
- biosphere, 33
- block diagram, 230
- boolean network, 125, 131
- bounded rationality, 337
- box-counting, 144
- brain, 55
- Braitenburg's vehicles, 172
- Broken Arrow Canyon, 188
- broken-stick model, 119
- Brund report, 258
- butterfly effect, 347, 350

- C++, 44, 75, 81
- canalizing, 9
- cascaded cubic polynomial recursion model, 41
- CASPER market model, 336
- categorical representation, 150
- cellular automata, 3, 6, 35, 74, 91, 110, 150,
 - 190, 346, 364
 - continuous, 110
 - fuzzy, 110
- chaos, 3, 7, 34, 40, 45, 92, 106, 130, 131, 162,
 - 202, 253, 291, 307, 364
- chaotic dynamics, 33, 124
- chartists, 291
- chicken game, 179
 - iterated, 180
- chromosome, 382
- climax, 35
- co-evolution, 336
- coalition, 376
- cobweb model, 244
- communication, 171
- community, 35
- complex dynamics, 6
- complex systems, 3, 61, 118, 130, 191, 311
- complexity, 3, 32, 34, 144, 191
- complexity theory, 358
- computer software, 141
- confusion effect, 175
- connectivity, 4, 25, 34, 55, 81
- conservation, 32
- consumer's utility, 349
- consumers, 320
- consumers' learning, 346
- content addressable memory, 20
- controlled switch, 131
- coordination, 171
- copyright, 37
- cortex, 4, 55, 74
 - visual, 25
- coupled map lattices, 91, 110
- crisis, 133
- critical point, 77

- criticality, 33, 244
- crossover, 382
- cryptology, 155
- crystallisation, 364
- cytochrome-oxidase blobs, 26

- data mining, 211
- debt, 305
- decision support system, 34, 189, 205
- deflation, 305, 312
- depression, 312
- Derrida plot, 10
- differential equation, 68, 101, 163, 229, 292, 359
- differentiation, 20
- diffusion, 68
- diminishing returns, 171
- diminishing returns, law of, 125
- Discrete Dynamics Lab, 3
- diversification, 118
- diversity, 118
- DNA, 35
- Duffing's equation, 101
- Dyar-Hutchinson rule, 118
- dynamical system, 40, 91, 110, 119, 130, 150, 294, 364

- Ecolab, 80, 358
- ecological, 382
- ecology, 35, 80
- Econolab, 358
- economy generator, 361
- ecosystem, 32, 188
- edge detector, 96
- edge of chaos, 141
- electrocardiogram, 40
- electroencephalogram, 40, 55, 74
- electronic, 130
- embryology, 20
- emergence, 142, 171, 191, 327
- entropy, 7, 126, 364, 367
- environment, 32, 188, 260
- environmental gradient, 33
- environmental informatics, 32
- environmental management, 32
- equilibrium, 20, 34, 68, 162, 180, 233, 244, 295, 307, 309, 350, 358
- eukaryotic cell, 19
- evolution, 80, 118, 132, 142, 171, 191, 327, 333, 358, 376
- exotic species, 34

- $1/f$ distribution, 119
- features, 132

- feedback loop, 33
- Field Programmable Gate Array, 132
- finite state machines, 382
- Fisher's Paradox, 316
- fitness, 337, 382
- foraging, 66, 171
- Fourier analysis, 144
- Fourier transform, 78
- FPGA, *see* Field Programmable Gate Array
- fractal, 110, 115, 144
 - dimension, 144
- fundamentalists, 291
- fuzzy, 210, 272

- game, 272
- Game of Life, 7, 126, 346, 365
- game theory, 282, 376
- Γ recursion, 40
- garden of Eden states, 11
- GAUcsd, 338
- Gaussian, 261
- GENESIS, 57
- genetic, 173
- genetic algorithm, 35, 172, 174, 216, 330, 336, 362, 382
- genetic networks, 3
- genotype, 35, 339
- Geographical Information System, 37, 188
- GIS, *see* Geographical Information System
- glider, 6–8
- Global Biodiversity Information Facility, 37
- global warming, 32
- goal interference theory, 205
- Goodwin model, 238, 306
- government, 376
- grand challenges, 32
- Great Depression, 305

- Hamming distance, 10
- Hausdorff-Besicovitch dimension, 145
- Hebbian learning, 25, 56
- hierarchy, 3, 118
- hippocampus, 55
- Hopfield, 20
- horizontal-vertical illusion, 42
- hypercolumns, 26

- image analysis, 91, 144
- increasing returns, 320
- increasing returns, law of, 125
- individual model, 190
- inflation, 305
- information technology, 32

- information theory, 98
- inners of a matrix, 254
- intermittent, 130
- International Organization for Plant Information, 37
- Internet, 32, 36
- invariant manifolds, 162
- IS-LM model, 308
- island biogeography, 35, 81

- Jury's test, 254

- Kaleckian, 315
- Kauffman's *NK* model, 36, 118, 362
- knowledge discovery, 211

- L-system, 36
- Lake, The, 376
- λ parameter, 15, 368
- land use, 32
- Langton's conjecture, 92
- language, 144
- learning, 219, 327, 333
- legal liability, 37
- leptokurtosis, 291
- Life, Game of, *see* Game of Life
- limit cycle, 237
- limits of acceptable change, 190
- linguistic rules, 210
- local-global symmetry model, 25
- lognormal, 118
- Lorenz system, 162
- Lotka-Volterra, 80, 306
- Lux model, 291
- Lyapunov exponent, 10, 45
- Lyapunov function, 163

- Müller-Lyer illusion, 42
- MacOS, 347
- market, 229, 272, 305, 320
- market clearing, 244
- Marshall, 314
- Marx, 306
- Matlab, 258
- meme, 362
- Mexican hat field, 26
- Microsoft Windows, 347
- migration, 80
- Minsky's "Financial Instability Hypothesis", 238, 305
- models
 - BioLand, 172
 - broken-stick, 119
 - CASPER, 336
 - CCPR, 41
 - DDLab, 3
 - Ecolab, 80, 358
 - Econolab, 358
 - GAucsd, 338
 - GENESIS, 338
 - Goodwin, 238, 306
 - IS-LM, 308
 - Kauffman's *NK*, 36, 118, 362
 - RBSim, 190
 - Santa Fe Artificial Stock Market, 327
 - SmartForest, 37
 - Swarm, 120, 191
 - The Lake, 376
 - von Neumann Technology, 358
 - Wonderland, 258
- monopoly, 378
- Monte Carlo, 340
- morphogenesis, 11
- MPI, 82
- multi-agent systems, 36, 171, 179, 191, 219
- multisectoral, 82
- Mumford-Shah model, 96
- mutation, 362, 382

- Nash equilibria, 282, 336
- neighbourhood
 - Moore, 5
 - von Neumann, 5
- network, 3
- network externality, 320, 346
- network traffic, 130, 137
- neural network, 4, 35, 56, 210
 - cellular, 141
- neural networks, 3
- neurology, 40
- neurons, 74
- niche, 342
- noise traders, 291
- nonlinear dynamics, 40
- Nosé-Hoover thermostat, 359

- object oriented programming, 197
- OECD, 37
- oligopoly, 336
- operating system, 347
- order, 3, 7

- parallel-line illusion, 42
- parallelism, 57, 80
- Pareto principle, 121
- path dependency, 320

- pattern attractor, 120
- pattern classification, 210
- payoff, 276, 379
- payoff matrix, 182
- perceptron, 219
- permanence, 81, 361
- phase transition, 74, 119, 364
- phenotype, 35, 80, 339
- pheromone, 66
- Phillips curve, 306
- physiology, 40
- Poincaré, 3, 101
- Poisson, 242
- pollution, 32, 260, 377
- polo distribution, 118
- Ponzo illusion, 42
- population, 35
- power law, 10, 40, 118, 146, 362
- Prandtl number, 163
- predator
 - avoidance, 171
- predator-prey, 172, 306
- prefractal, 146
- Prisoner's Dilemma
 - Iterated, 272
- prisoner's dilemma, 179, 376
 - iterated, 180
- producers, 320
- protocols, 37
- psychophysical, 40
- PVM, 57
- pyramidal cells, 56

- Q learning, 272
- quality assurance, 37
- quantum mechanics, 127, 364

- random Boolean networks, 3
- random graph, 18
- random walk, 173, 291
- rank-abundance, 118
- rationality, 282, 376
- RBSim, *see* Recreation Behaviour Simulator
- reaction-diffusion, 3, 7, 91
- reactive agents, 171, 191
- Recreation Behavior Simulator, 190
- Recreation Behaviour Simulator, 196
- recreation opportunity spectrum, 190
- recreators, 188
- reductionism, 34
- robot, 35, 172

- S-Plus, 59

- Santa Fe Artificial Stock Market, 327
- Schue-Cohn criterion, 254
- self-organisation, 3, 6, 25, 110, 364
- self-organised criticality, 118
- self-similarity, 146
- shaky hand principle, 180
- sigmoidal, 212
- simulated annealing, 47
- Simulink, 261
- SmartForest, 37
- snag, 132
- software robots, 197
- sparse connectivity, 4
- sparse matrix, 80
- speciation, 80
- species-area rule, 80
- Sraffa, 314
- stability, 34, 81, 164, 167, 244, 308
- standards, 37
- strange attractor, 253, 296
- strategy
 - generous, 179
 - greedy, 179
- structural information, 95
- superlinear speedup, 80, 82
- supply and demand, 244, 280
- sustainable development, 258
- Swarm, 120, 191
- Synergetics, 291

- TCL, 81
- technical analysts, 291
- technological progress, 362
- Tit for Tat, 180
- traffic intersection, 180, 181
- traffic simulation, 130, 137
- tragedy of commons, 282
- trap, 132

- unbounded rationality, 349
- utility, 321

- V1, 25
- variable structured system, 130
- VHS, 347
- von Neumann, 110, 150, 358
- von Neumann neighbourhood, 76, 364
- von Neumann Technology, 358
- VRML, 83

- wavelets, 144
- Wolfram classification, 75, 110, 156, 369
- Wonderland model, 258

World Wide Web, 37