Visualisation

Projects at iVEC@UWA

Paul Bourke

iVEC @UWA

Outline

- Introduction to iVEC.
- Introduction to science/data visualisation and supporting displays.
- Present various representative visualisation projects over the last 2 years.

- **Graphical representation of data**
  Project: Tornado simulation
- **Volume visualisation**
  Project: Pausiris mummy

2011-2012

- **Dark (matter)**
  Visualisation of large scale cosmological simulations

- **High resolution image capture and display**
  Projects: Beacon Island - Rock Art
- **Automatic 3D reconstruction from photographs**
  Projects: Rock art - Dragon Gardens - Ngintaka story

2013

- **360 degree video recording in cultural heritage**
  Projects: Ngintaka story - Mah Meri rituals

Introduction to iVEC

http://ivec.org

- A partnership between the 5 key research organisations in Western Australia.
  - Edith Cowan University
  - Curtin University
  - Murdoch University
  - Commonwealth Scientific and Industrial Research Organisation
  - The University of Western Australia

- Facilitates research at the partners by providing advanced computing: hardware, software and expertise.

- Five programs
  - **Education**
    Provides year-round training modules and runs an interns program each summer.
  - **eResearch**
    Supporting researchers maximise the benefits of digital technology within their discipline.
  - **Industry and government uptake**
    Facilitate relationships between iVEC and government and industry.
  - **Supercomputing technology and applications**
    Collaborates with and encourages the uptake of supercomputing by researchers.
  - **Visualisation**
    Supports visualisation through expertise and specialist infrastructure.

Visualisation

- Definition:
  Visualisation is the process of applying advanced computing techniques to data to facilitate insight into the underlying structures, relationships and processes.

- Definition for my mother to tell her friends:
  “Turning data into images and animations to aid interpretation/understanding”.

- Very interesting field: requires a wide range of skills
  - computer programming
  - algorithms in computer graphics
  - mathematics
  - realtime / interactive APIs and technologies
  - human / computer interfaces
  - knowledge of human perception theory
  - creativity and design

Outcomes

- Informing researchers.
- Conveying research outcomes to peers.
- Educational resources for University courses, the general public, and school children.
- Exhibitions in museums and art galleries.

- Research outcomes for data visualisation include
  - Uncovering something new.
  - Understanding some aspect of the data faster.
  - Finding errors.

Visualisation @ iVEC

- Three 1/2 FTE funded positions: UWA - Curtin - CSIRO
- Budget to support visualisation activities of researchers at any of the iVEC partners.
- Compute infrastructure dedicated or optimised for challenging visualisation projects.
- Displays to support visualisation.

Curtin & ECU

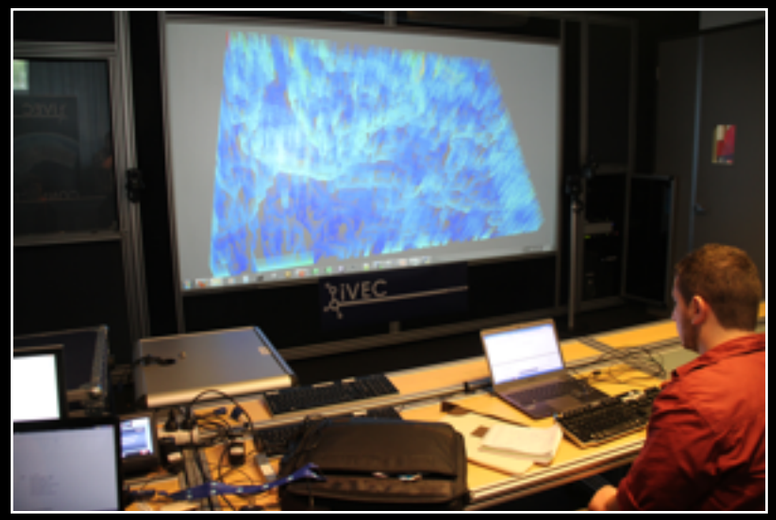

CSIRO

UWA

Murdoch

ECU

Visualisation @ iVEC

- Capture infrastructure:
  - Stereo3D video camera
  - High resolution video cameras
  - Specialist camera rigs
  - Structured light camera
  - 360 video camera
- Software tools and expertise.

Stereoscopic 3D cameras

High resolution video cameras

360 degree video capture

3D scanners

Displays

- As the name suggests, visualisation most often uses the sense of vision to convey information to the human brain.
- As such it makes sense to leverage the capabilities of the human visual system, three main areas:

- Stereopsis: the sense of depth we perceive due to having horizontally displaced eyes.

- Peripheral vision: the sense of "being there", of being immersed.

- Visual fidelity: ability to resolve detail at scale.

Tiled display (Fidelity)

iDome (Immersion)

Other senses ...

- Sonification: visualising using our sense of hearing.
- Classic examples:
  - Hospital pulse measurement: "The machine that goes PING".
  - Geiger counter.
- Two most common approaches are to map some variable to the waveforms (eg: amplitude or frequency modulation) or to map to instruments (eg: midi).
- Pure sonification can be difficult, often just sounds like noise or really bad music.
- More commonly used to accompany and reinforce the visuals.
- Good example is sonification of nuclear tests from 1945-1996 by Japanese artist Isao Hashimoto (http://www.youtube.com/watch?v=cjAqR1zlCA0)
- Sonification of pulsars

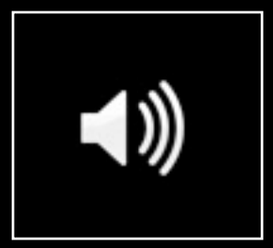

Midi instrument, equal tempered scale

Sense of touch

- Force feedback has been used for some time to allow data to be “felt”, haptics.
- Commodity example is joystick vibration in car driving games.
- Used extensively in remote surgery - eg: force feedback scalpel.
- More recently it has been possible to make physical models that can then be explored physically.
- Exploring data in the same way as we explore objects in real life.

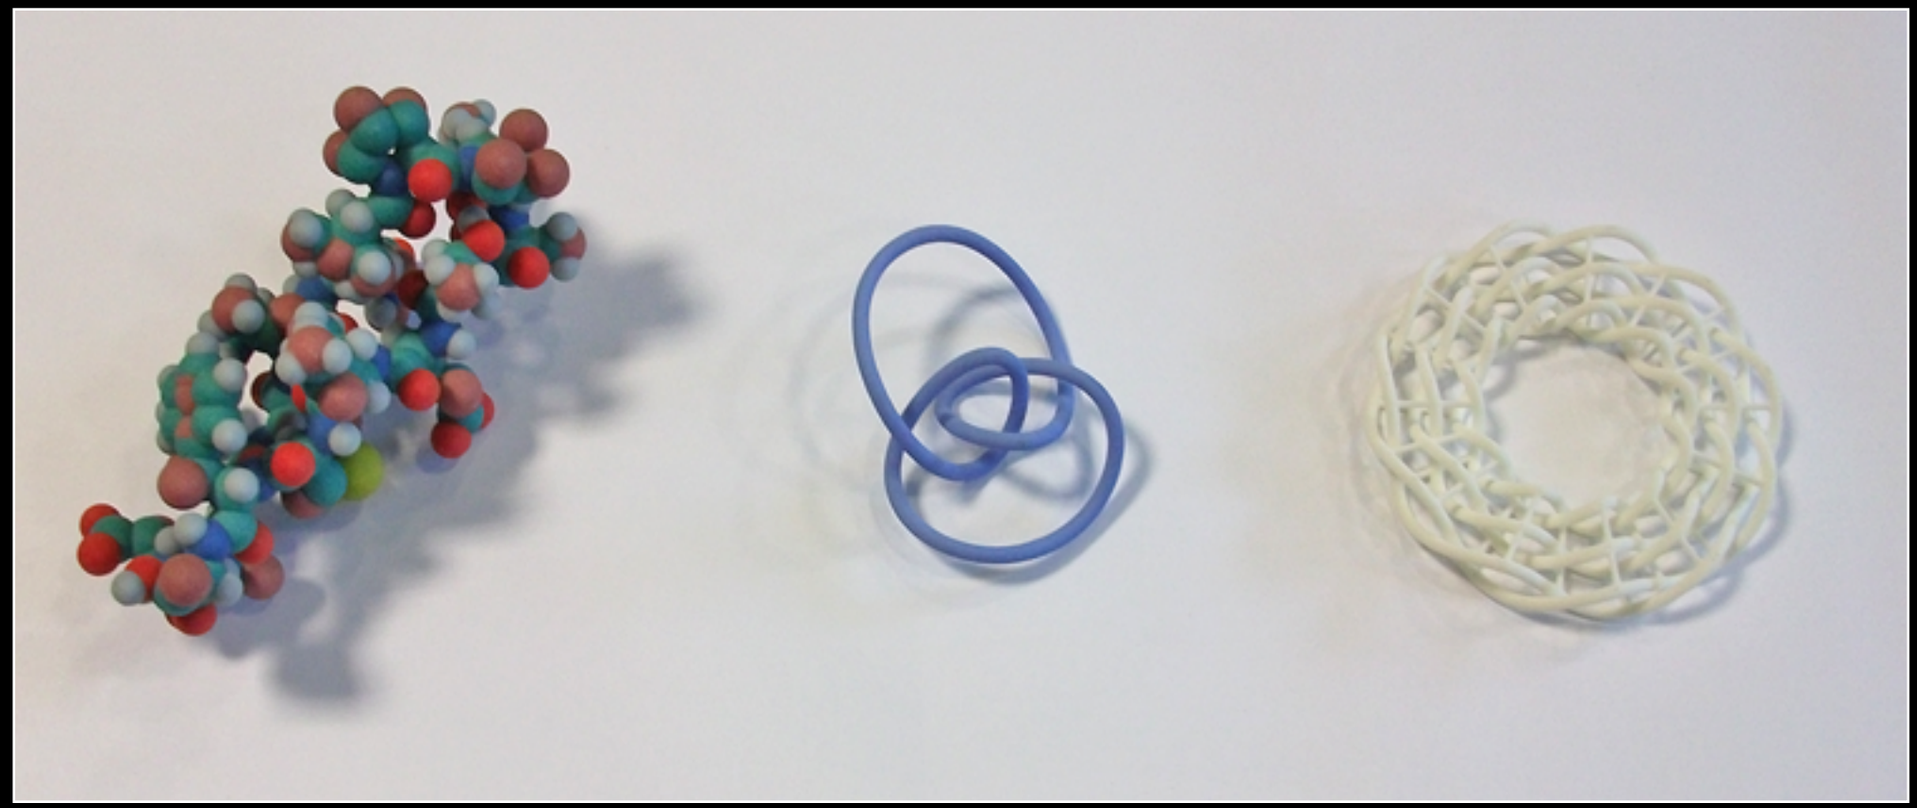

Peptide Knot theory Neuroscience

Mapping data to graphical elements

- A large part of what occurs in visualisation is mapping variables to graphics.
- Sometimes the mappings are obvious/intuitive, other times more freedom is possible.
- Colour often used to represent scalar quantities.
- “Glyphs” is the term given to graphical elements whose characteristics reflect a number of variables. Direction, volume, strength ...

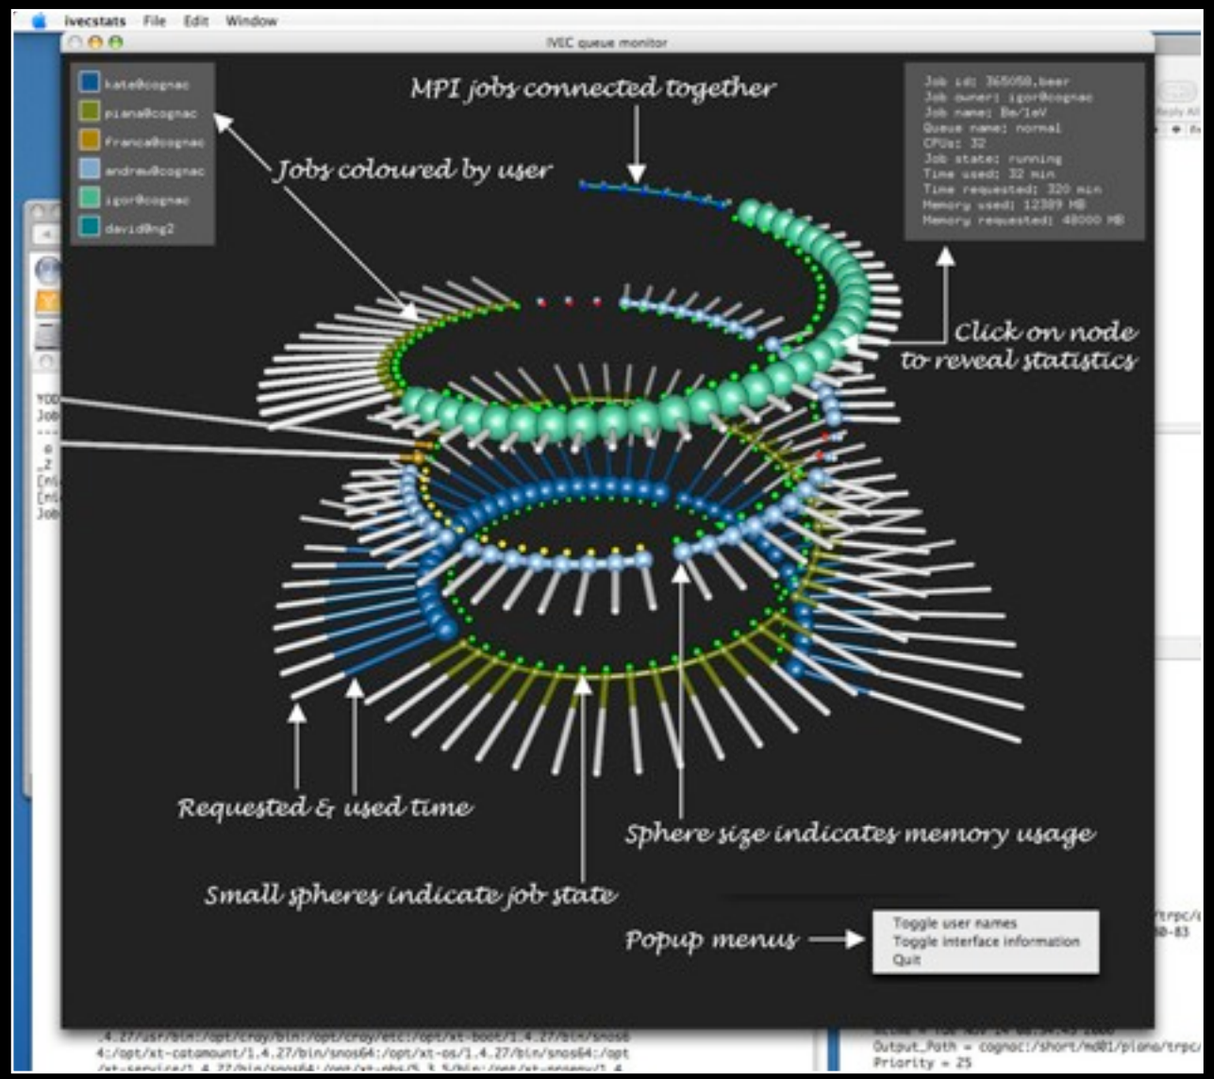

First, a distinction between data and illustrative visualisation

Movie

DNA

Courtesy Drew Berry, WEHI

Obvious literal mappings

Conventional (informed) mappings

Movie

Being inside the data

Movie

Visualising mathematics in vision research

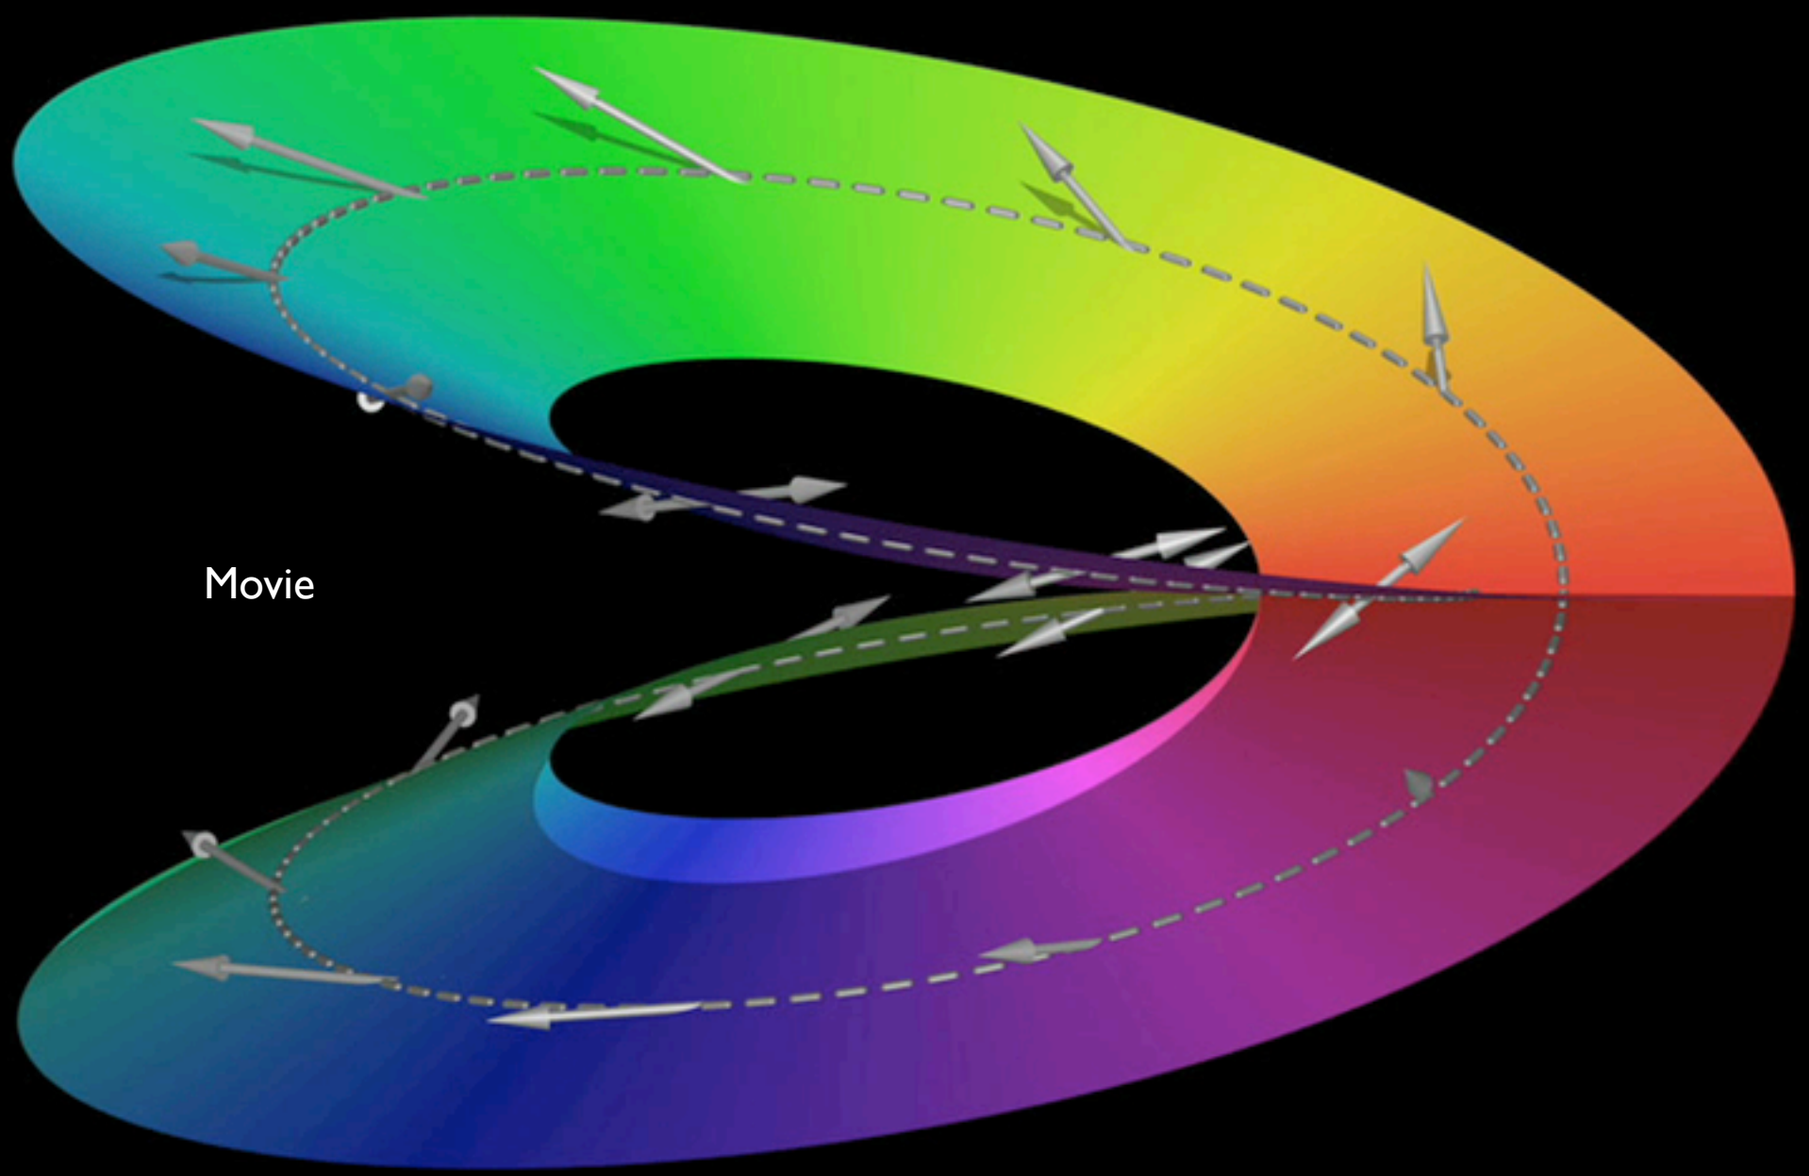

Movie

Movie

Tornado simulation

Volume visualisation

- A digital image contains some quantity sampled on a regular grid on a 2D plane.

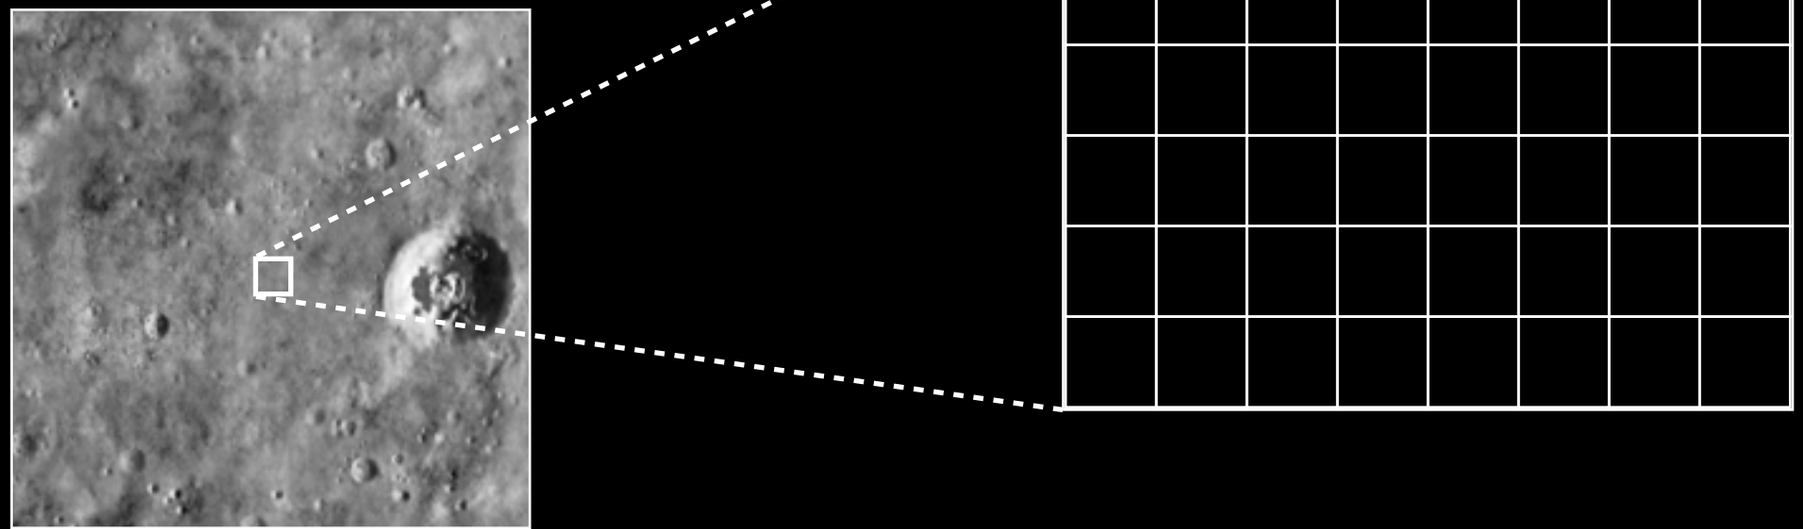

- In a volumetric dataset there is some quantity sampled on a regular 3D grid.

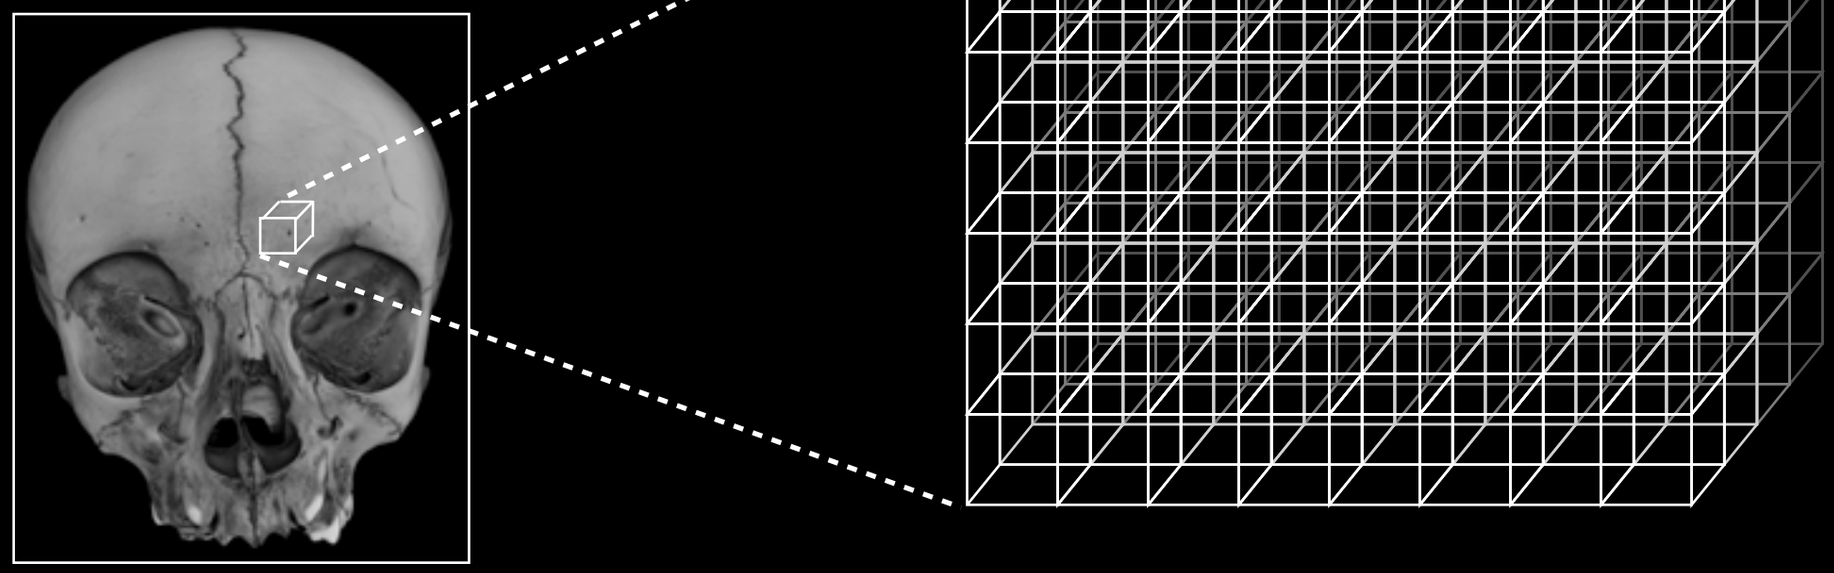

Volumetric data

- Volumetric datasets have been a common data type in many areas of science for some time.
- Traditionally one thinks about medical data, for example MRI.
- Other scanning and 3D imaging technologies include CT (MicroCT) and CAT scans. There are many others.
- Volumetric data also arises from numerical simulations.
  Quite common in astronomy and engineering (finite element calculations).
- In scanned volumetric datasets the quantity per voxel depends on the scanning technology. For example: MRI essentially gives water content, CT gives density.
- For volumetric datasets derived from simulation there can be multiple variables per voxel.

Medical research (MRI)

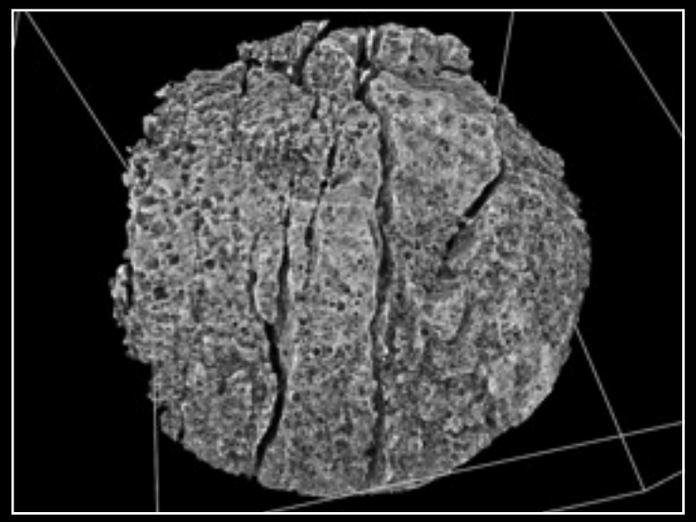

Geology (CT)

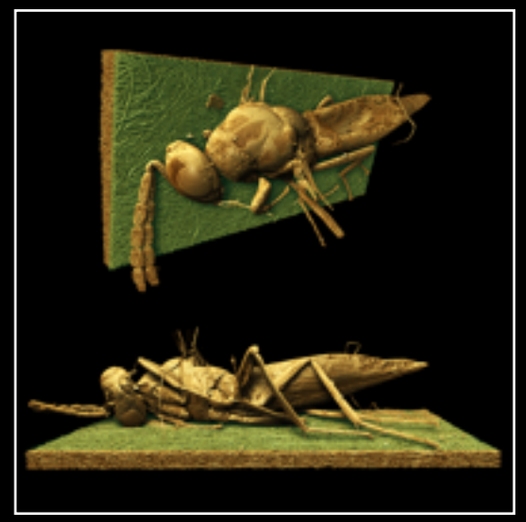

Entomology

Volume visualisation

- The process of exploring and revealing the structure/interior of a volumetric dataset.
- The general approach involves a mapping between voxel values and colour/opacity.
- Realtime volume visualisation generally requires hardware assistance, notable graphics cards.
- Has always been a demanding area in visualisation, the data volumes researchers wish to visualise has always been ahead of the technology.
- Still the case with huge volumes from MicroCT scanners and Synchrotrons.

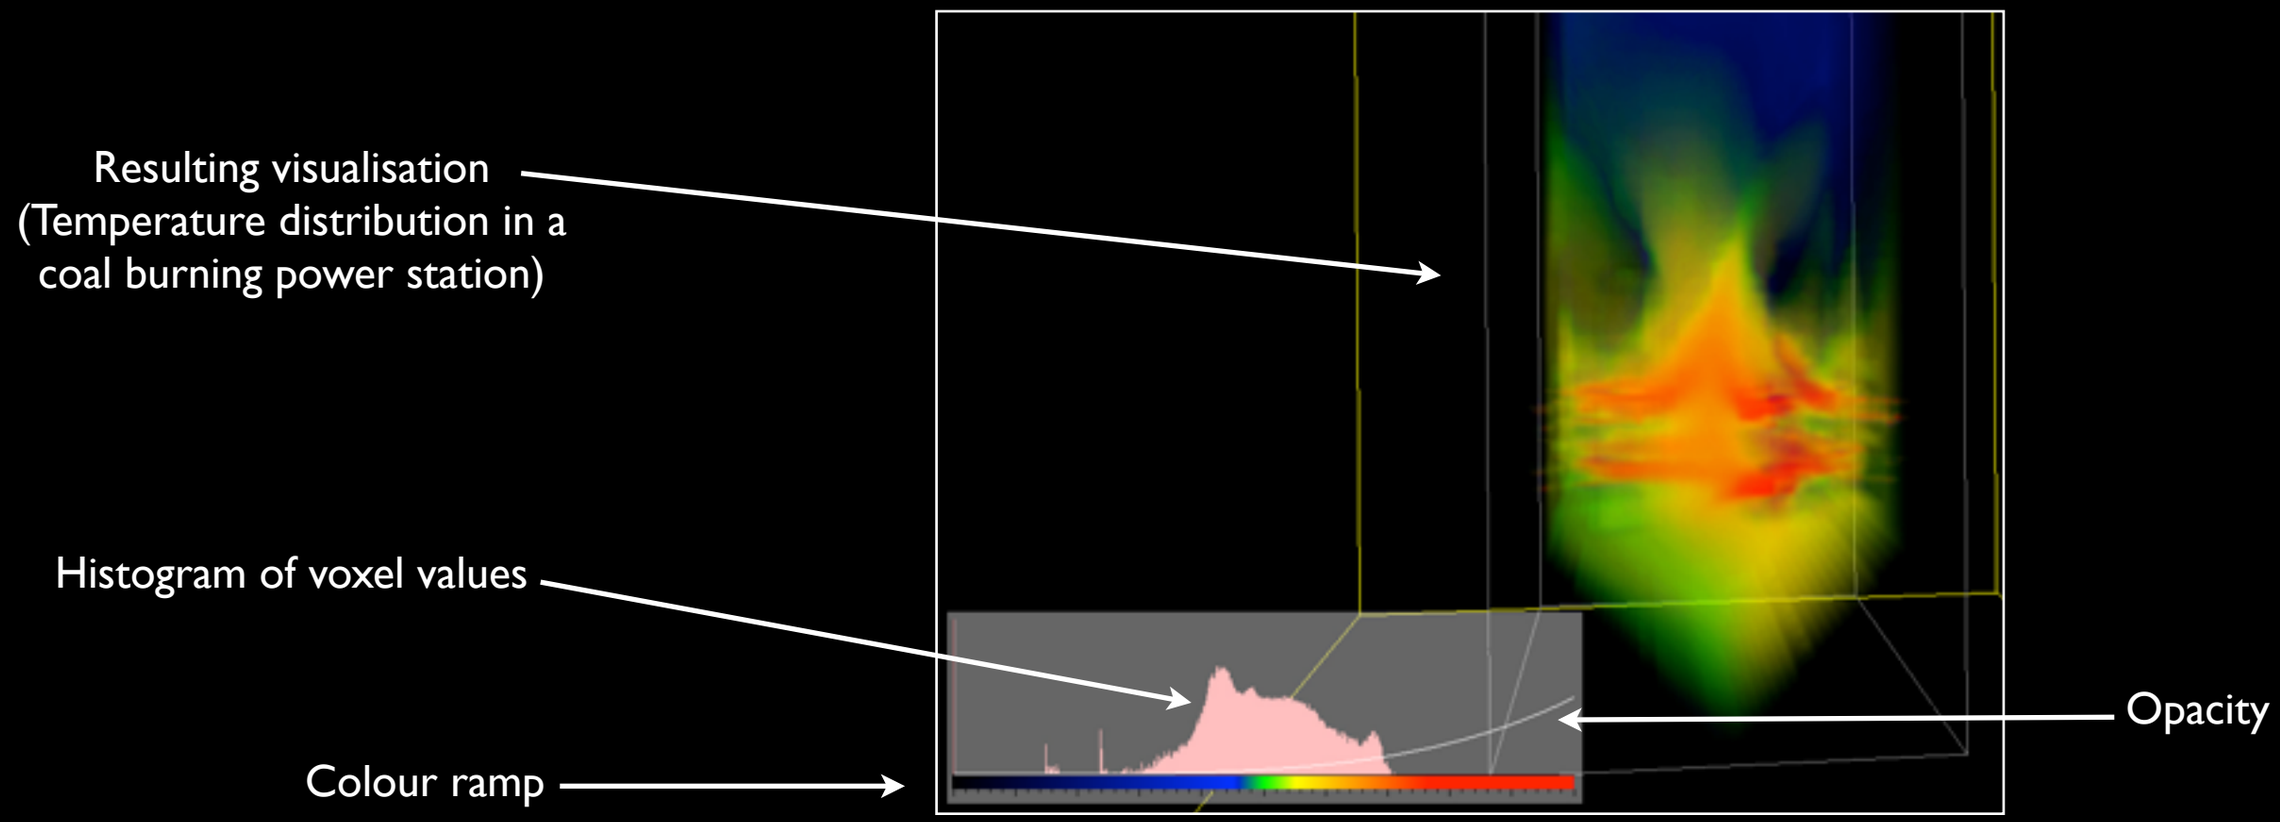

Example: Western Gorilla (Male)

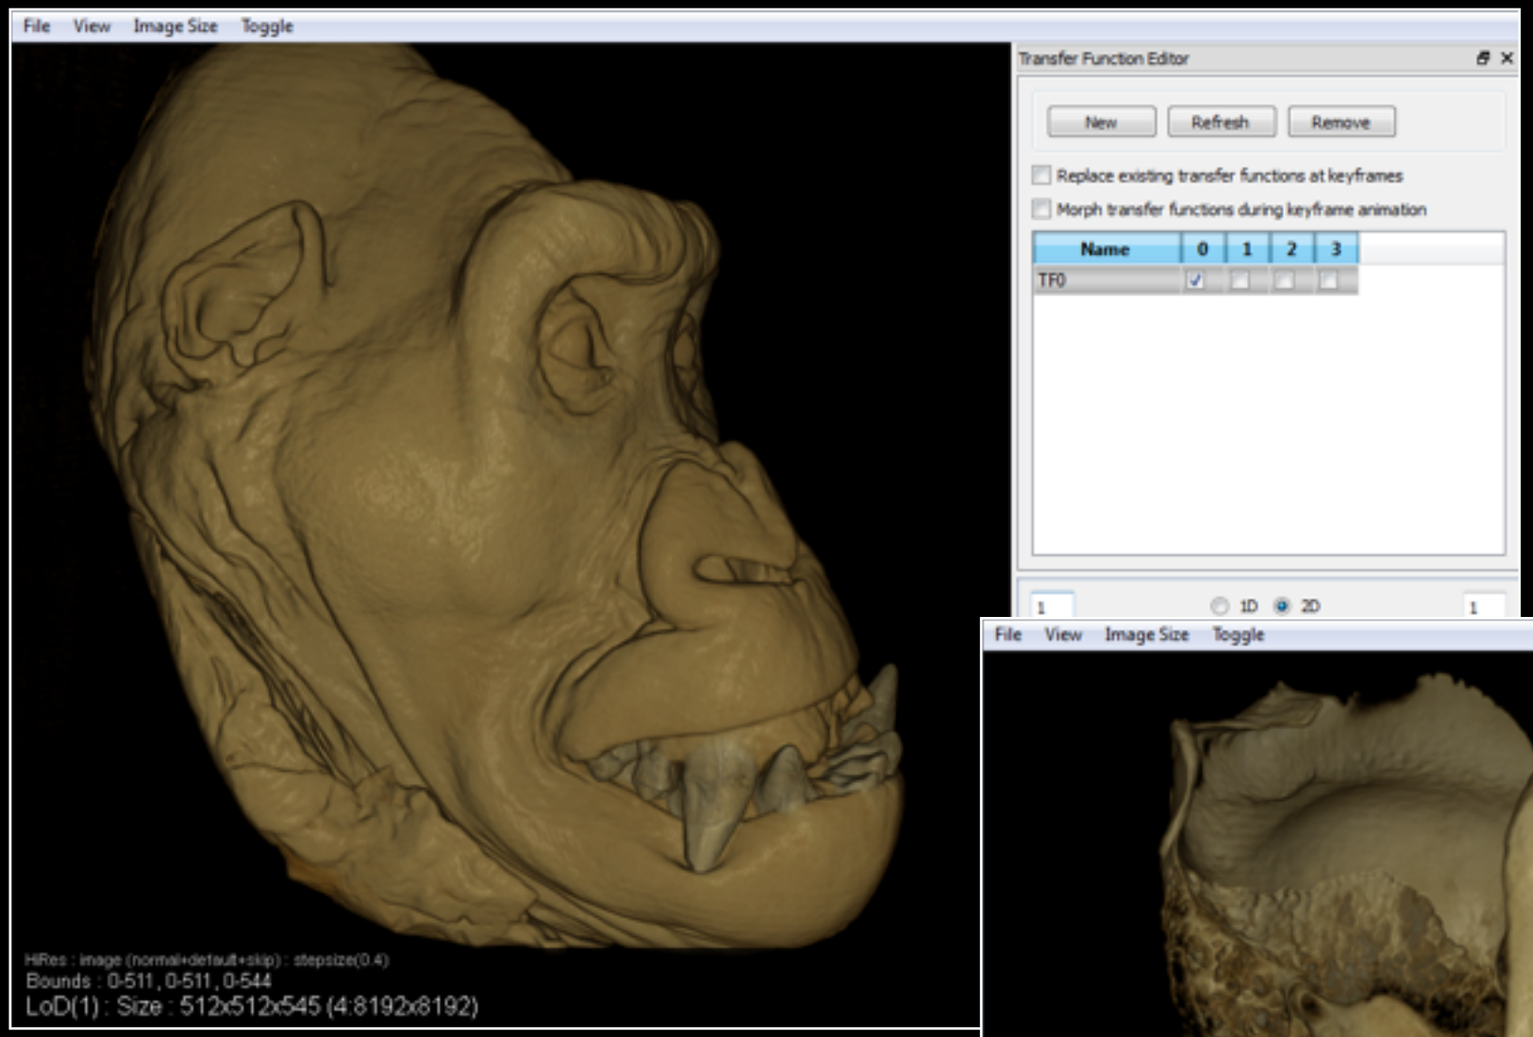

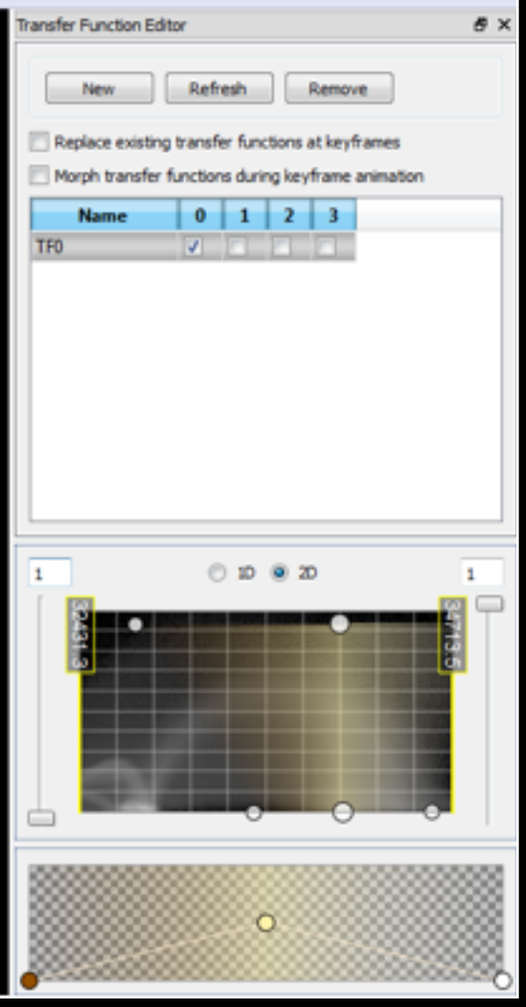

Drishti

Example: Rabbits liver

Movie

Courtesy Ajay Limaye, ANU

Pausiris mummy
Museum of New and Old Art, Hobart

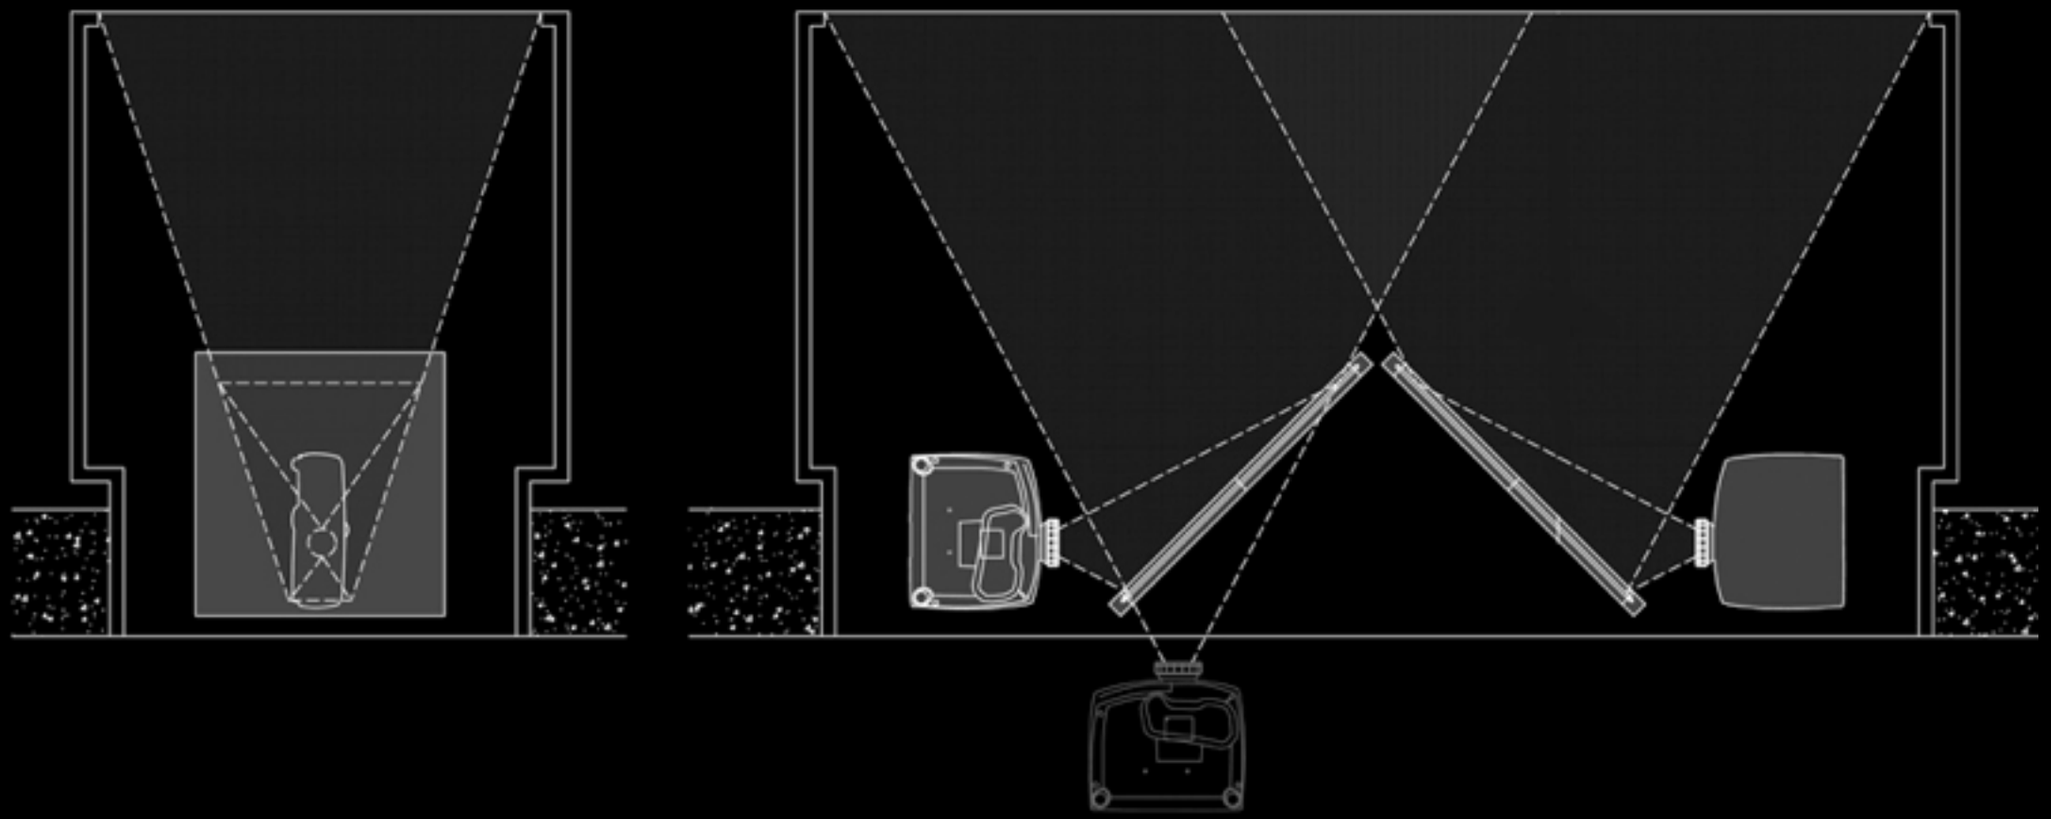

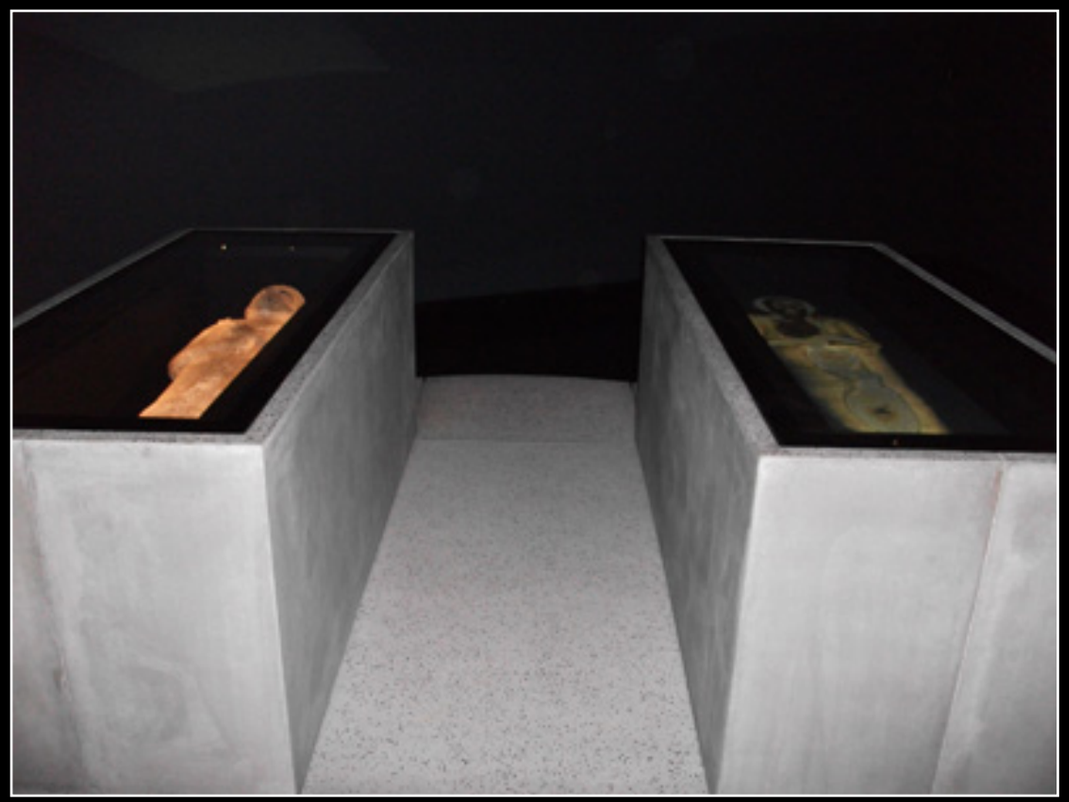

Visualisation of simulation data

- Project: Dark Matter cosmological simulations
- Example of visualisation starting out as being provided only for the researchers and being repurposed for an internationally distributed planetarium production.
- Characteristics:
  - Large numbers of points, minimum 200 million, maximum 1 billion (COSMOS).
  - Generally three types of particles: Dark Matter, Stars, Gas.
  - Relative numbers of each type of particle may vary over time.
  - Each point has a region of influence, smoothing kernel.
  - Typically have multiple parameters per particle.
    Interest here in position, velocity (for time interpolation), mass, smoothing radius.
- Outcomes
  - Explore pipelines appropriate for these types of data.
  - High resolution and quality animations for researchers post simulation.
  - High impact images and animations for public outreach.
  - Support for multiple projection types: orthographic, perspective, fisheye, spherical.
- Data volumes such that the visualisations were also performed on the supercomputers creating the simulations.

Not just points!

- If it were just “points” it would be much simpler.
- Region of influence are 3D functions of radius, similar to a “point spread function” in optics. Note this is used within the simulation software so not an arbitrary choice for the visualisations.
- For particles without a smoothing kernel (eg: stars) a Gaussian is used which allows the same pipeline to be employed. Use a single standard deviation, mass determines the amplitude.

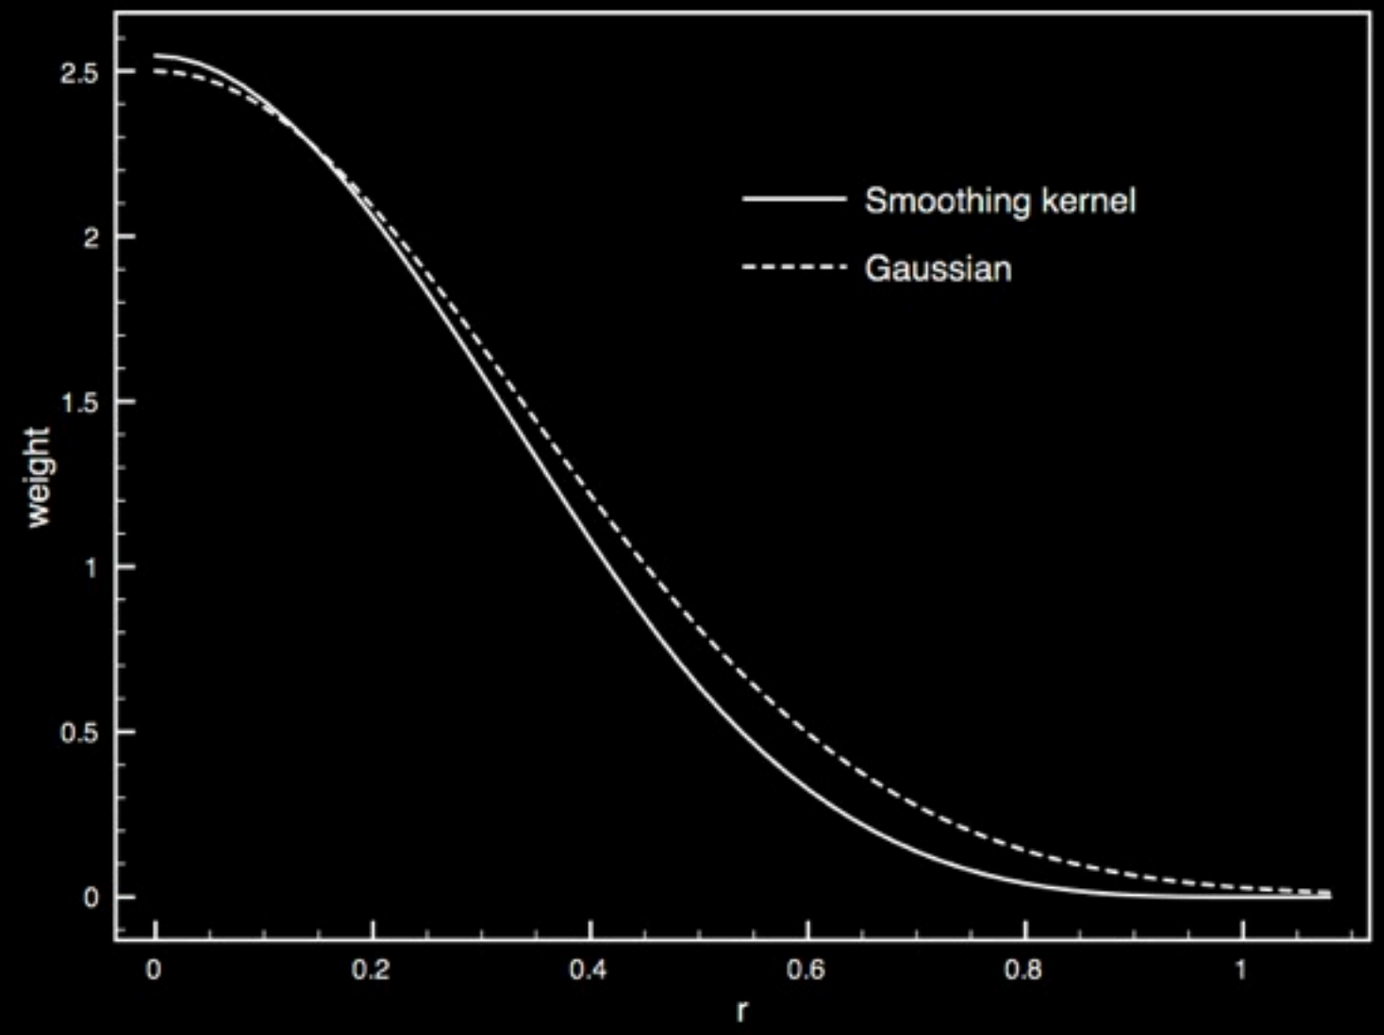

$$W(r) = \begin{cases} \dfrac{8\left(1 - 6\left(1 - \frac{r}{s}\right)^2 + 6\left(1 - \frac{r}{s}\right)^3\right)}{\pi s^3} & r < \frac{s}{2} \\ \dfrac{16\left(1 - \frac{r}{s}\right)^3}{\pi s^3} & \frac{s}{2} \le r < s \\ 0 & r \ge s \end{cases}$$

Control over the time/quality trade-off

- Common requirements of very large datasets is to be able to render simply and fast, better quality as time becomes available.
- Be able to manage interactive performance as well as high quality rendered imagery.
- Implemented smoothing kernel by sampling (regular or stochastic) in 3D. Points are then projected onto plane, cylinder, or spherical surface. The image is then a histogram the projected points contribute their kernel weighted mass to.
- Advantage of being able to form image frames with speed/quality trade-off.

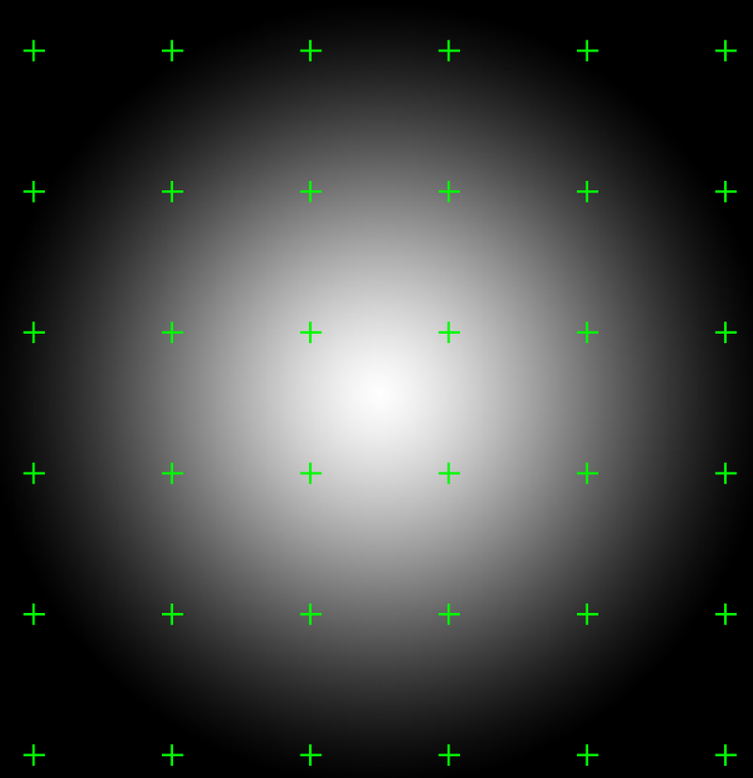

Uniform sampling

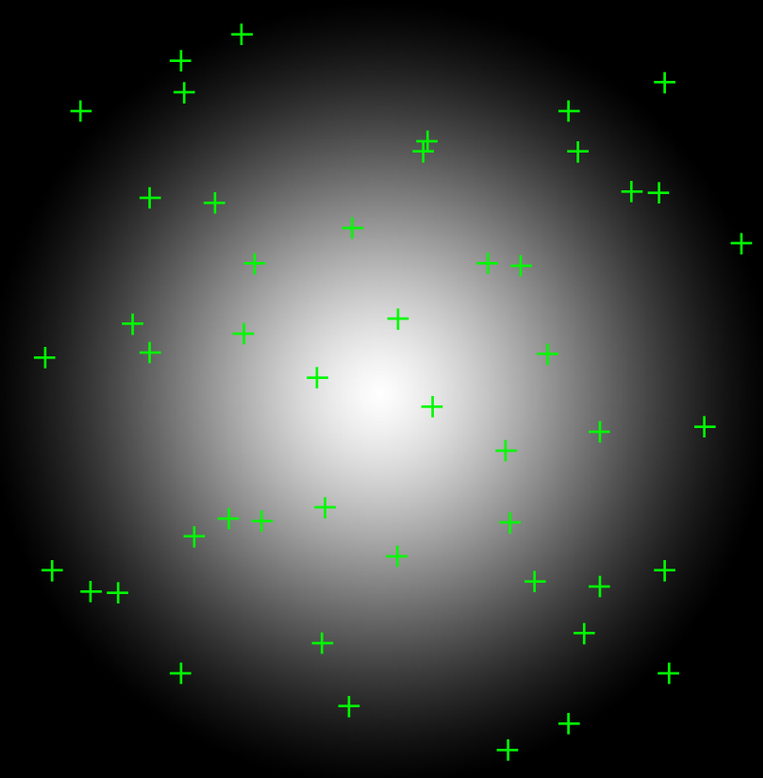

Random sampling

Parallel as with all HPC

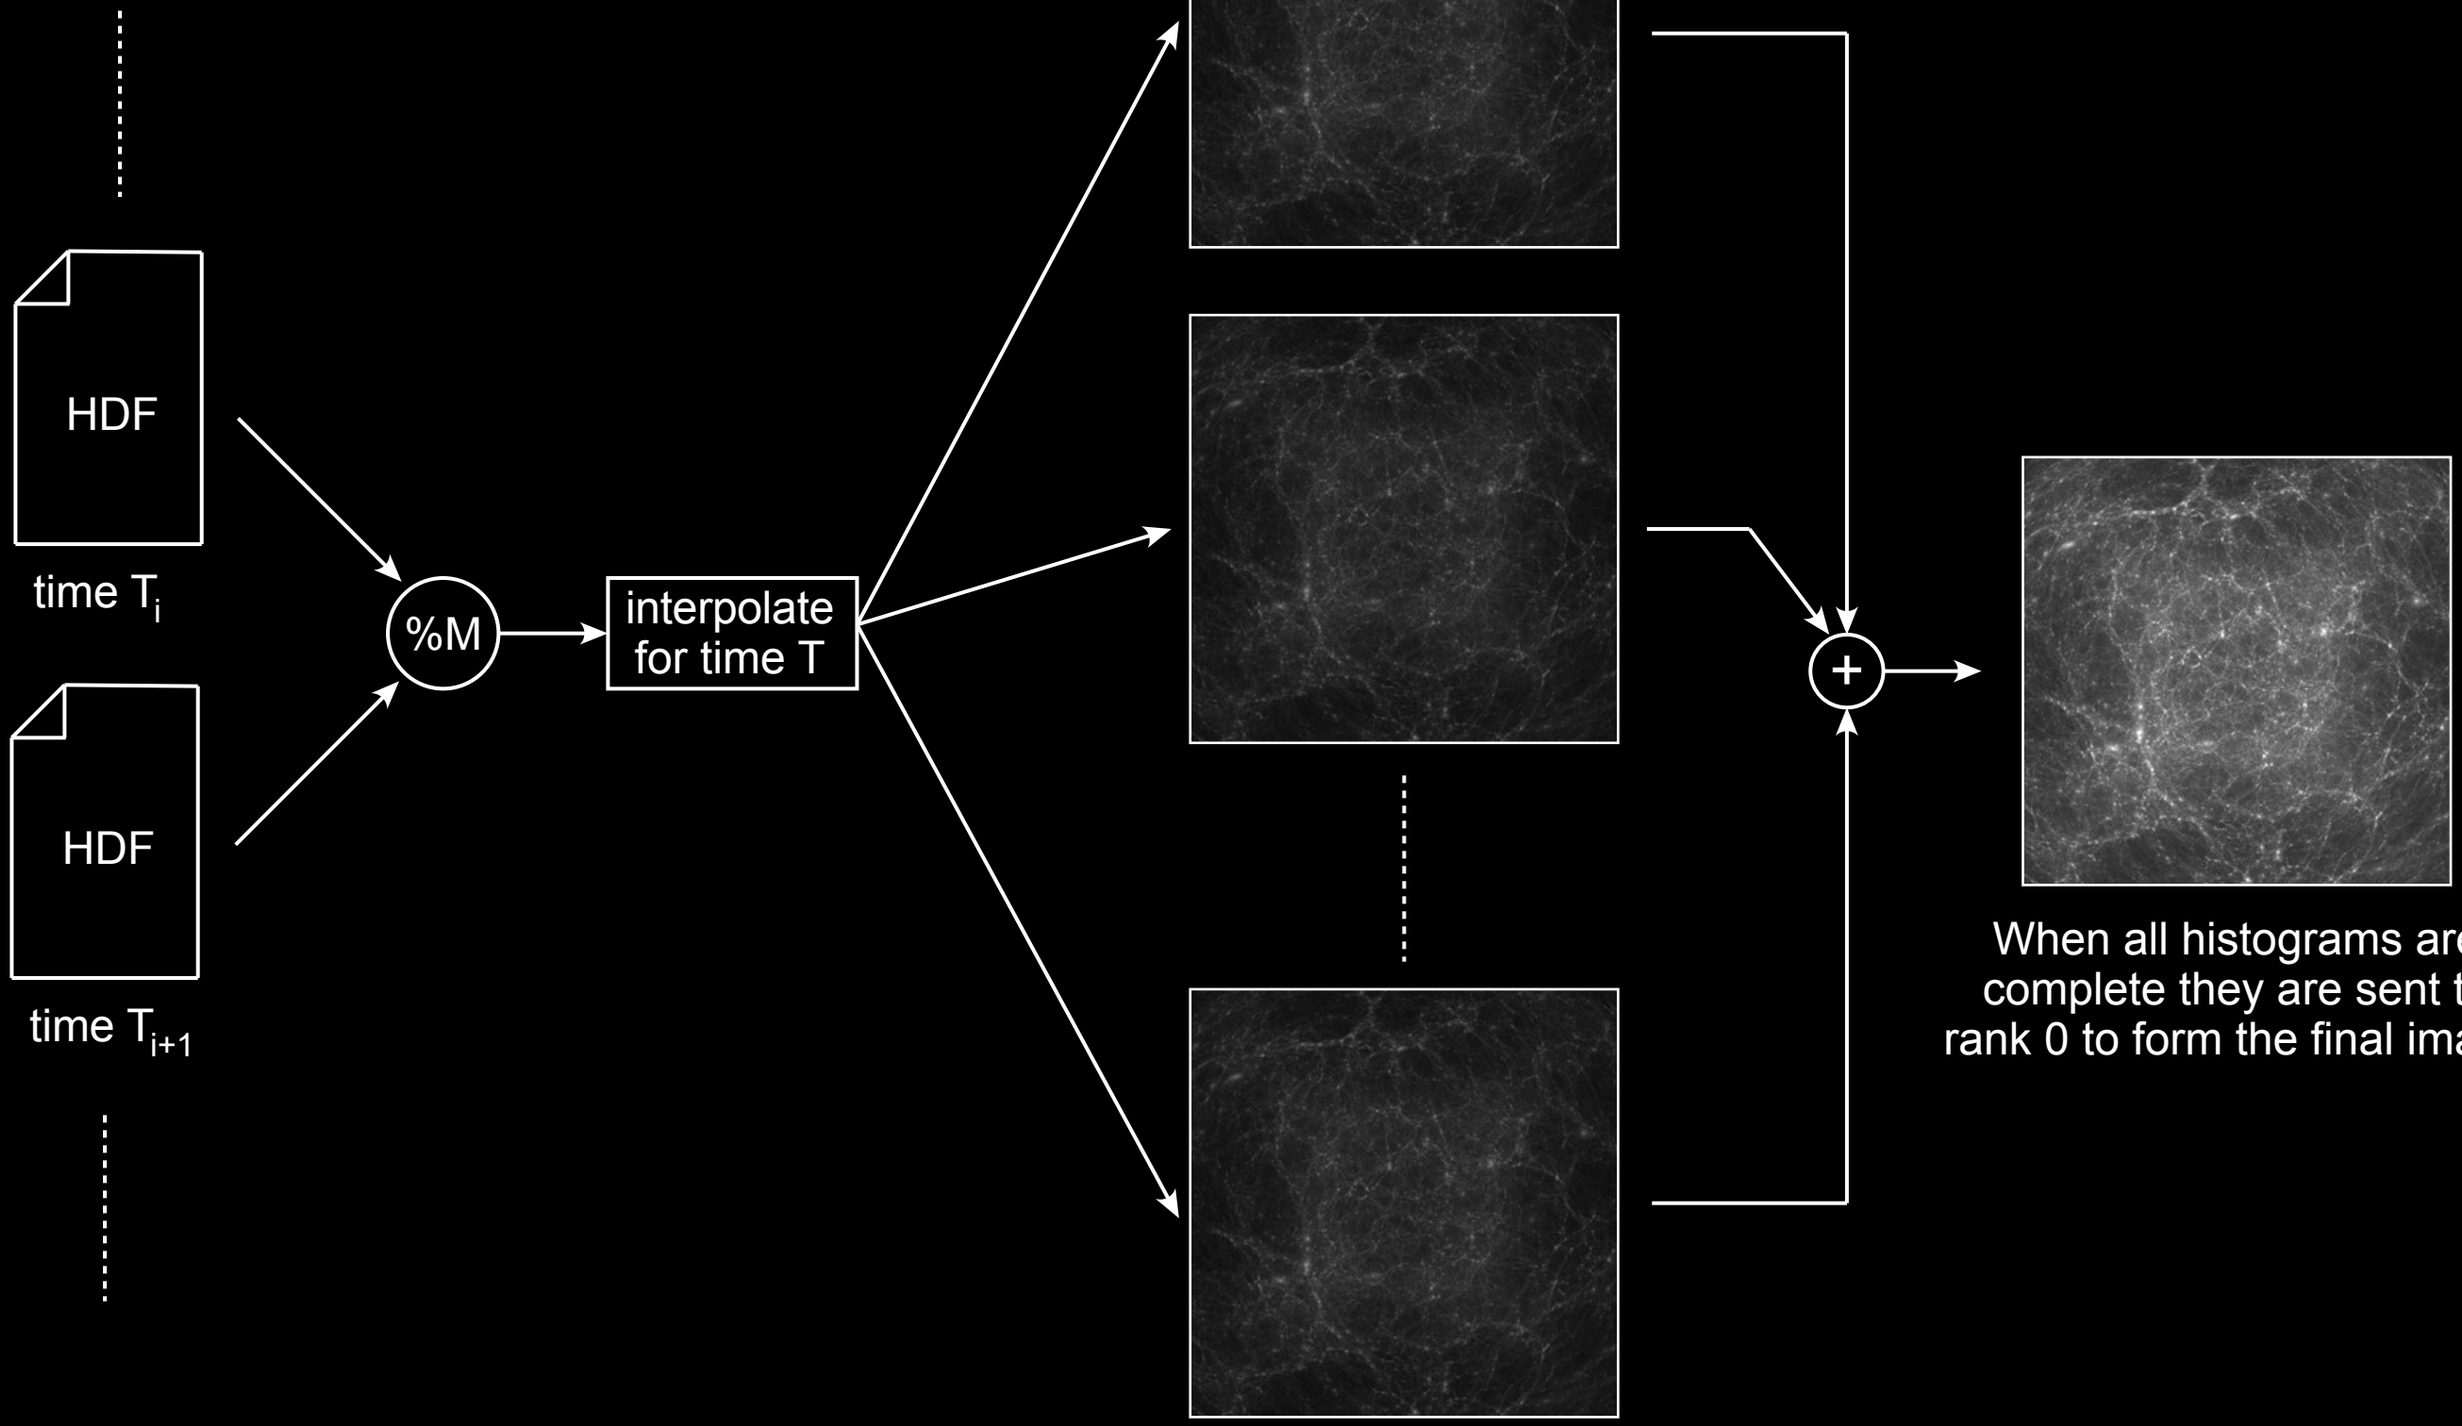

Cosmos

- Simulation within a cubic region (periodic bounds) of the Universe just after the Big Bang.
- 600 million light years on each side of the cube.
- Shows dark matter collapsing over 14 billion years of cosmic time, forming filaments and collapsing haloes of the Cosmic Web.
- Note there is no smoothing kernel here, the images look smooth and continuous due to the 1 billion+ particles per time step.
- Even at 3Kx3K, if the whole dataset is in shot then on average there are over 100 points per pixel (if they were distributed uniformly).
- The final image is essentially a histogram formed on the projection plane.
- Original simulation computed on vayu (NCI).
  Used 1024 cores, 2.8TB RAM, took 19 hours (~20,000 CPU hours)
  Rendering performed on epic (iVEC).

Movie

KINETIC

- Simulation of the formation of a Spiral Galaxy similar to our own Milky Way but about half the current age.

- The Gas follows the Dark Matter along the filaments.

- Computed on epic machine (iVEC).
  Used 1024 cores, 2.05TB RAM, took 470 hours (~500,000 CPU hours).
  Rendering performed on epic (iVEC).

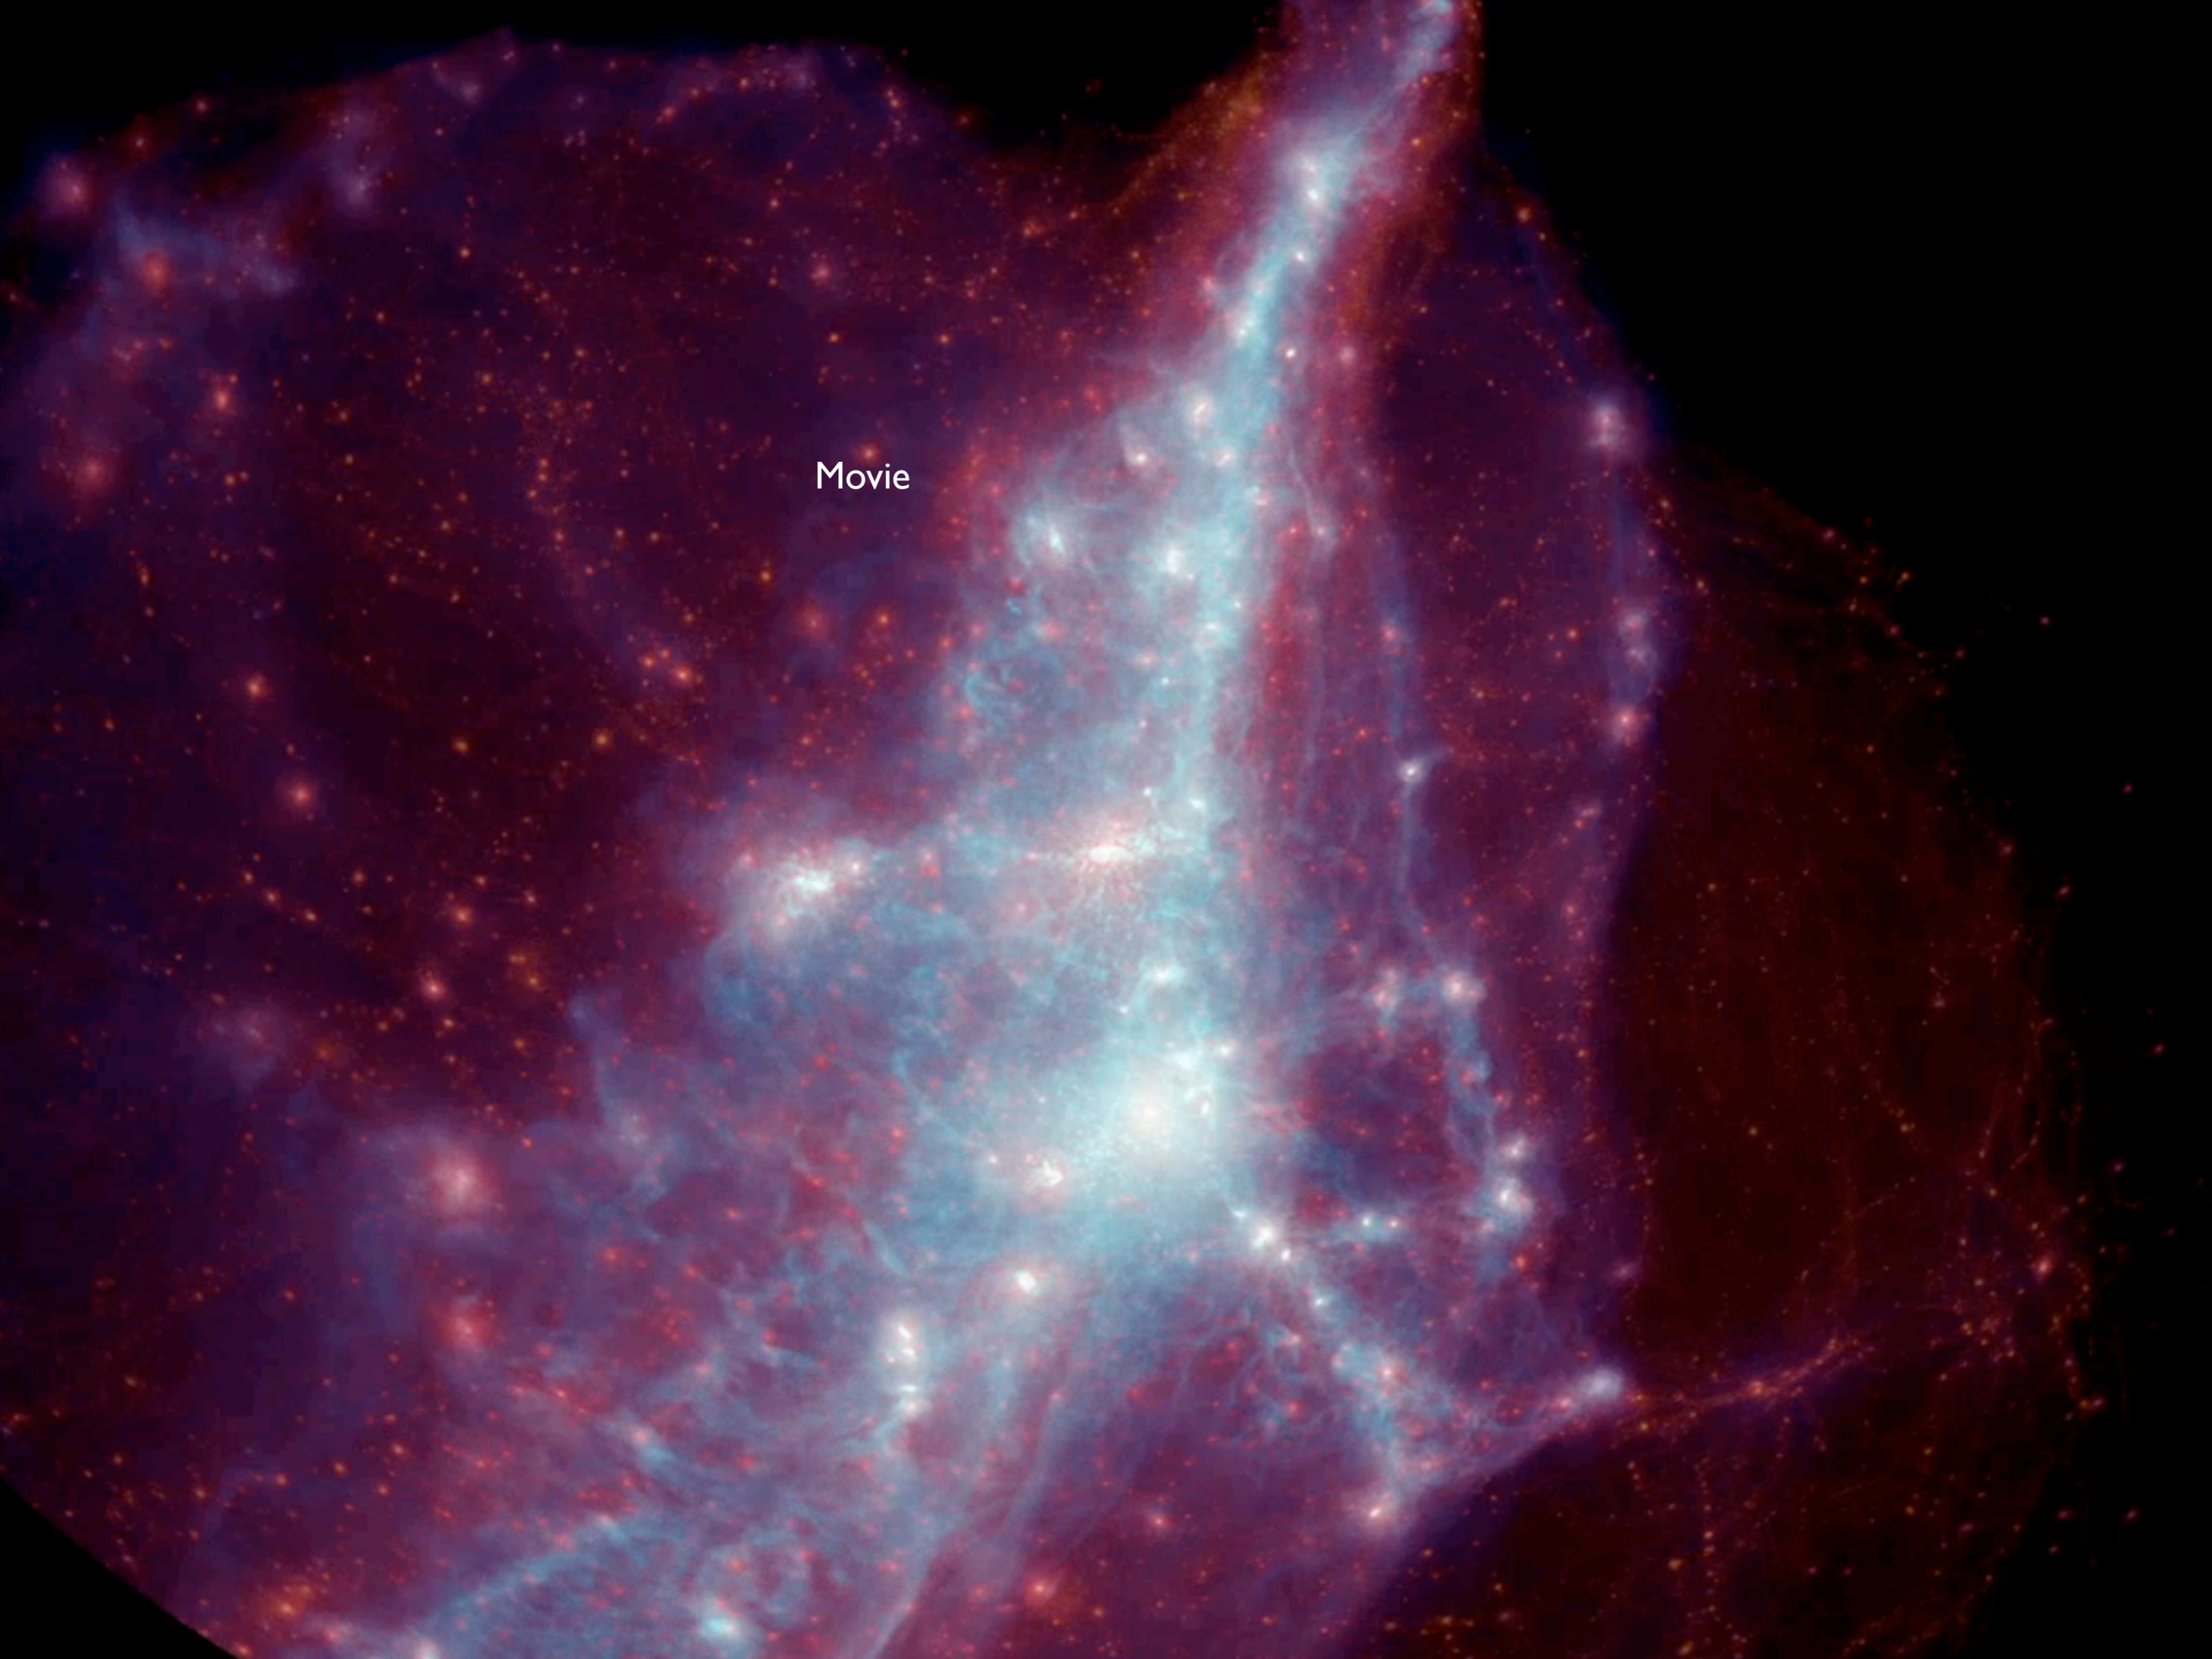
Movie

Dark

- Example of science visualisations being used in public outreach.
- http://darkthemovie.info
- *DARK is a fulldome movie that explains and explores the nature of Dark Matter, the missing 80% of the mass of the Universe.*
  The search for Dark Matter is the most pressing astrophysical problem of our time – the solution to which will help us understand why the Universe is as it is, where it came from, and how it has evolved over billions of years – the unimaginable depths of deep time, of which a human life is but a flickering instant.
- Currently showing in 16 countries, translated into 5 languages ... and counting.


Movie
rev
Zero emission
Electric Vehicle
WESTERN AUSTRALIA
GALAXY

High resolution image capture and display

- Imaging sensor resolution is only growing modestly.
  Current commodity SLR cameras are around the 20 to 30 MPixel range.

- Arbitrary high sensor density means the lens quality may be the final limiting factor to higher resolution.

- How does one capture imagery at higher than sensor resolution?

- Solution is to join a large number of photographs, each of a smaller area, together.
  A widely used technique from astronomy to microscopy.

- Motivations
  - Capture imagery from site where access is problematic.
  - Capture imagery of greater research value.
  - Acquire as image of the entire object as well as detail.

Hubble ultra deep field (HUDF)

- 10,000 galaxies
- 1mm x 1mm @ 1m
- 7000 x 7000 pixels

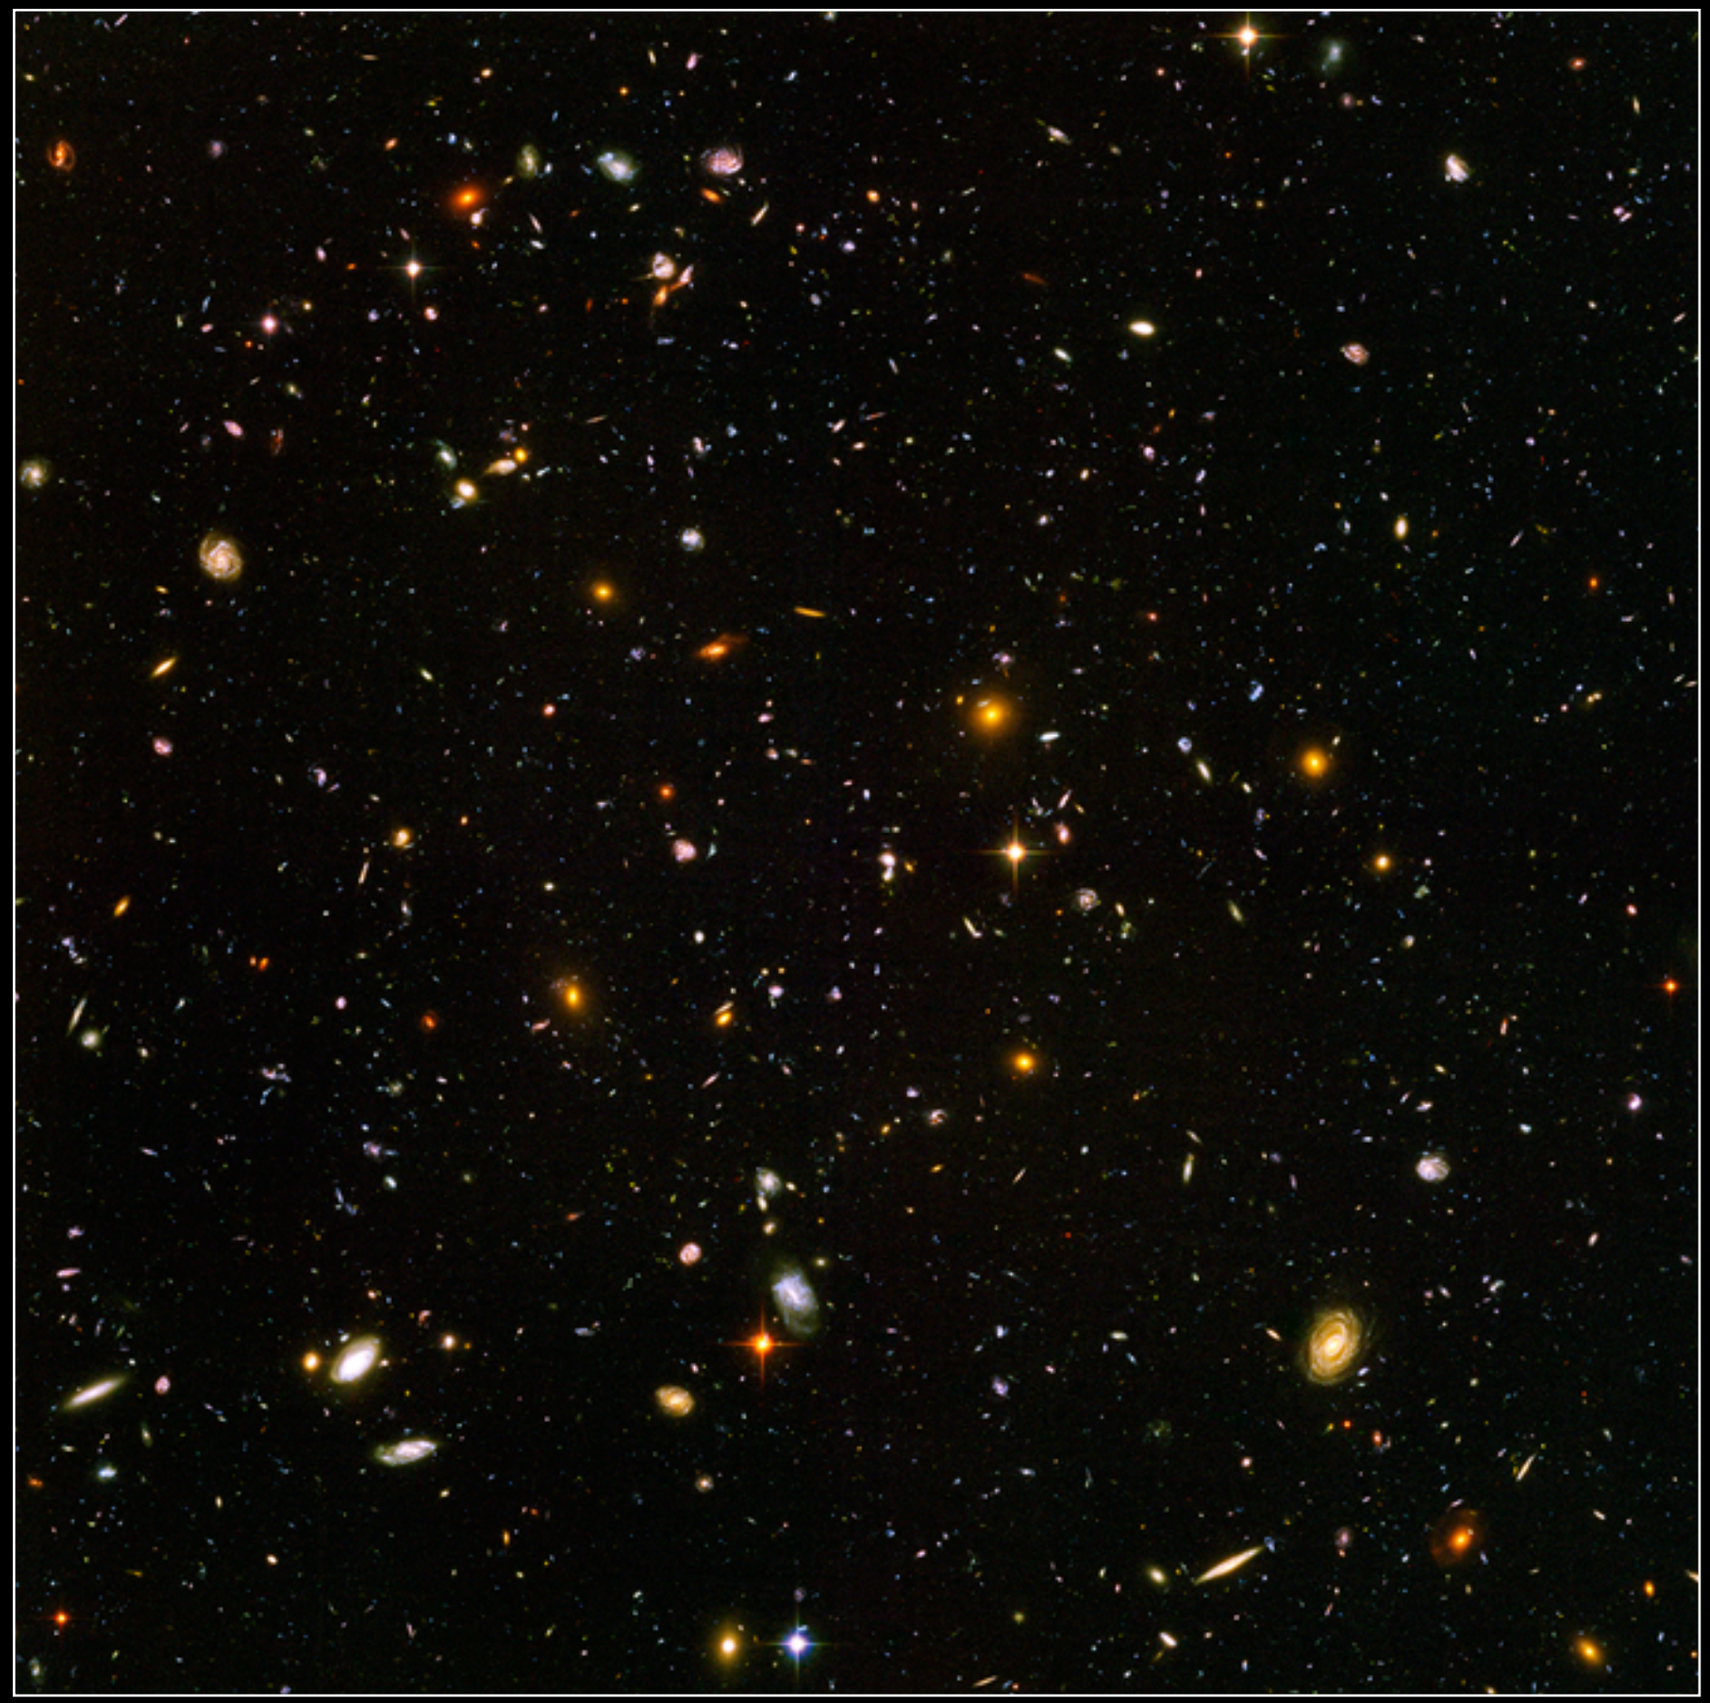

Optical microscopy

- Rat neuron
- 11,000 x 10,000 pixels
- 4x4 tile

CMCA, UWA

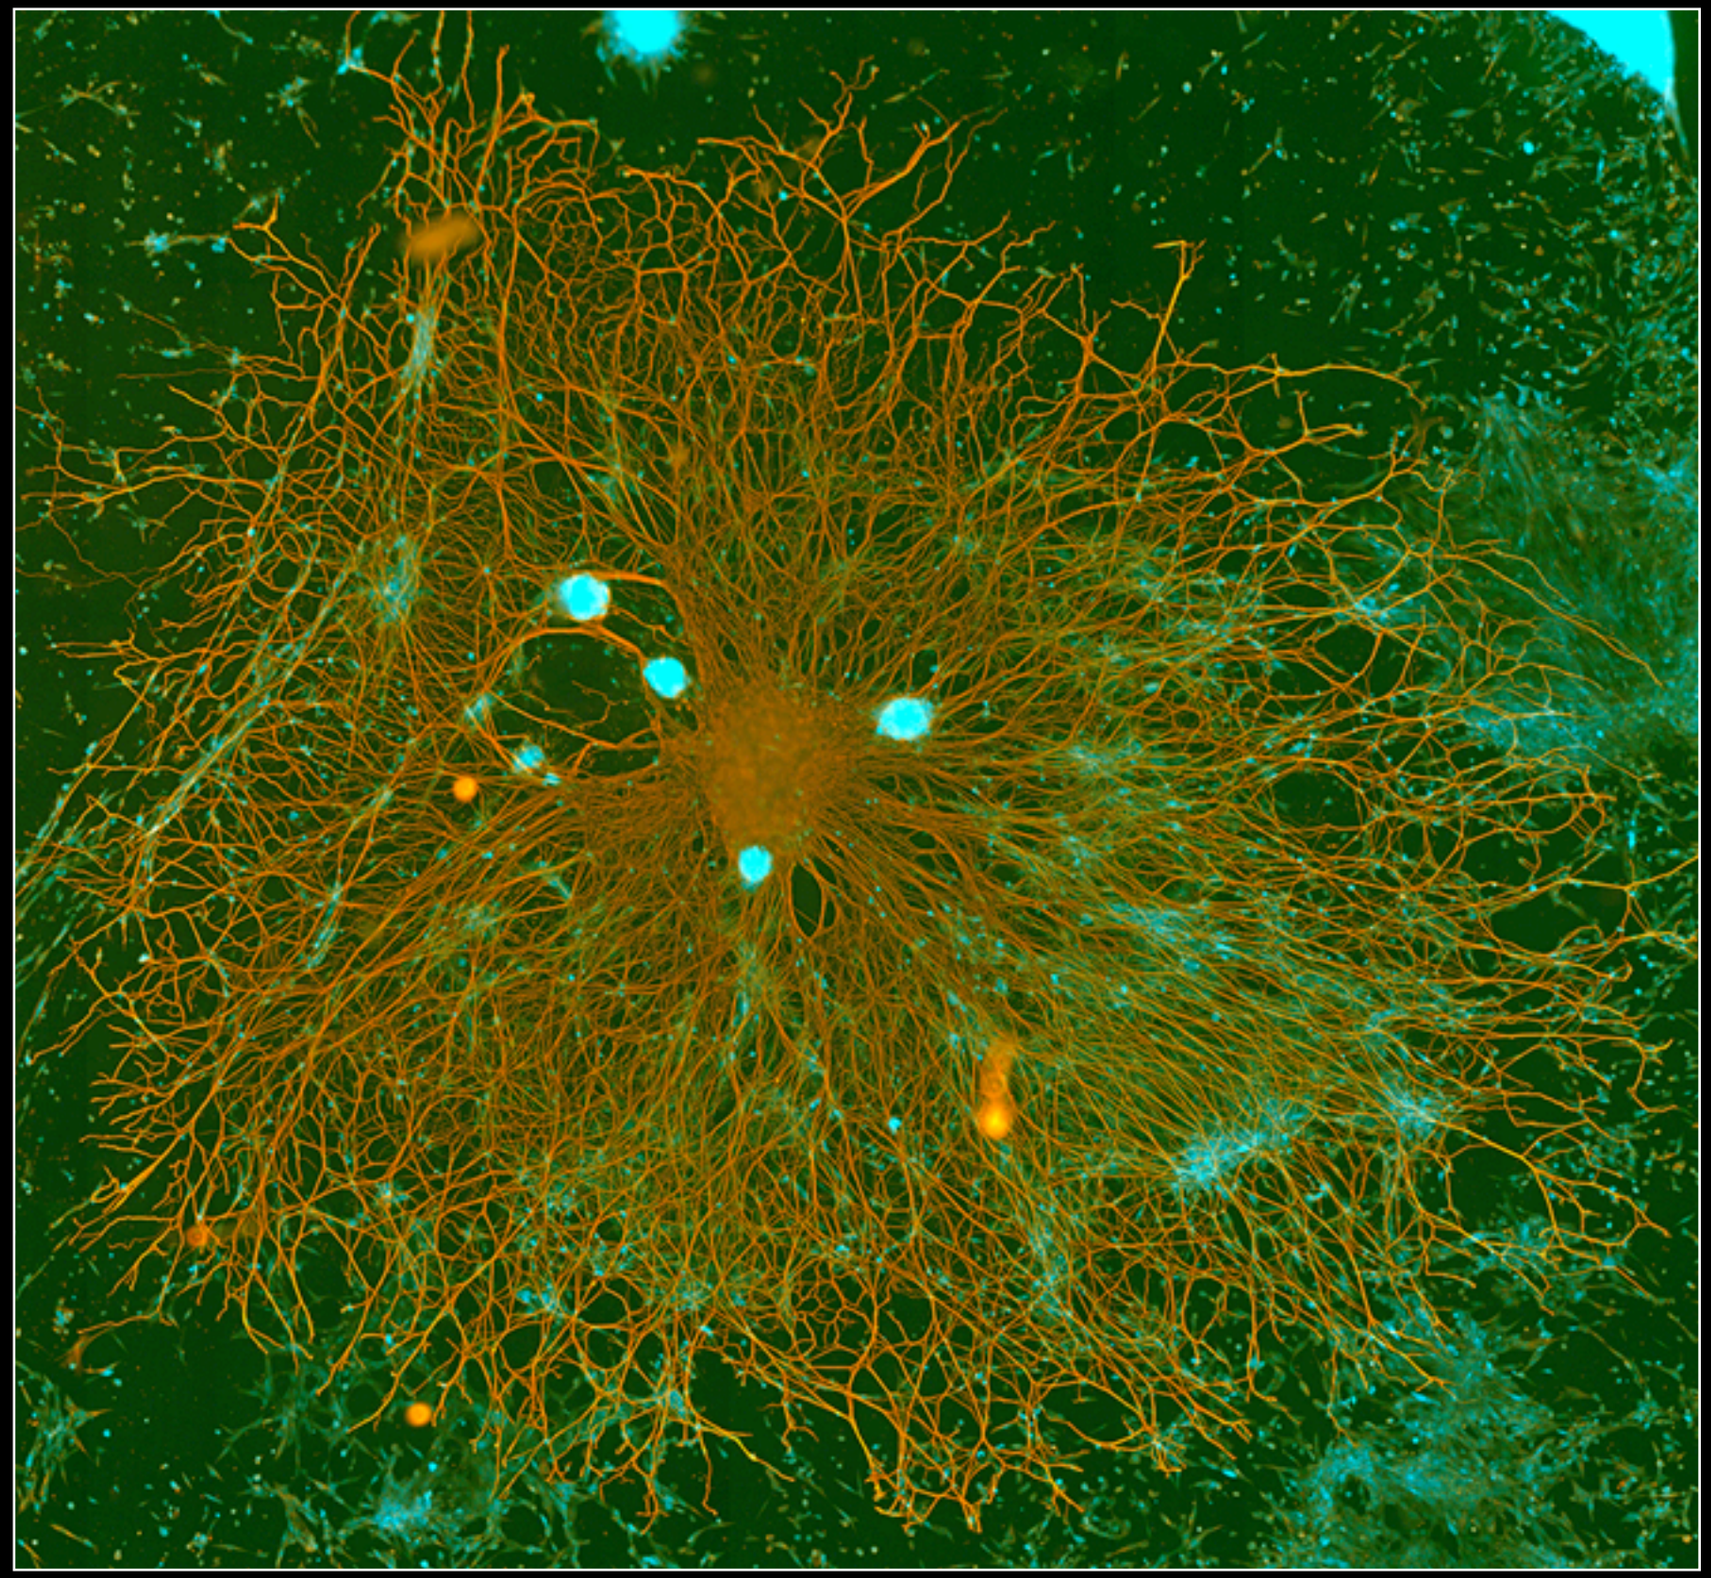

Geology

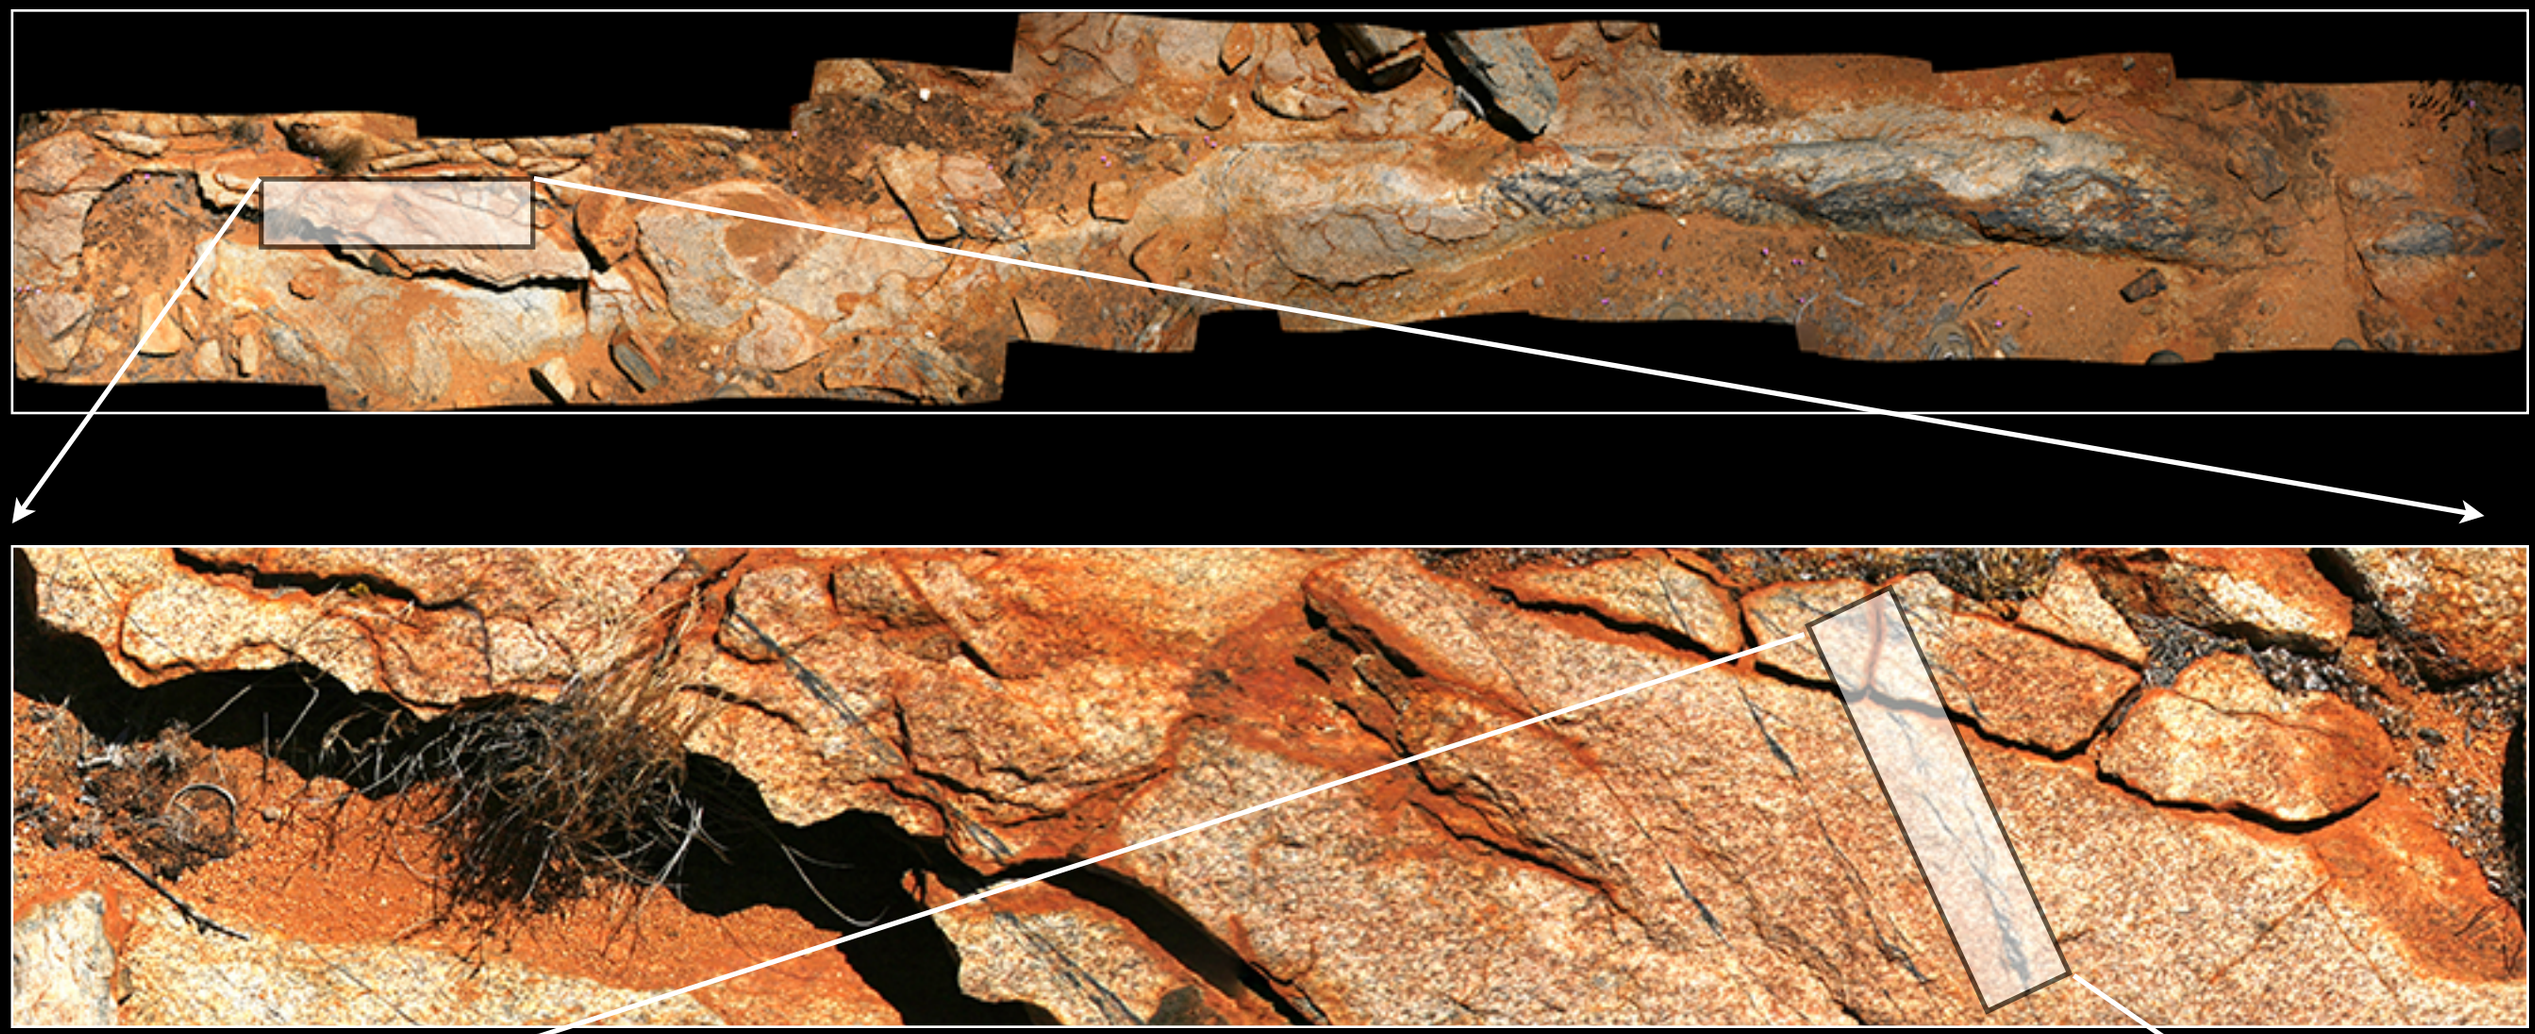

Basic technique

- Motorised camera rigs save the time of manual shoot and move camera.
- Final resolution only limited by zoom lens field of view and camera distance.

13 x 3 grid

40,000 x 10,000 pixels

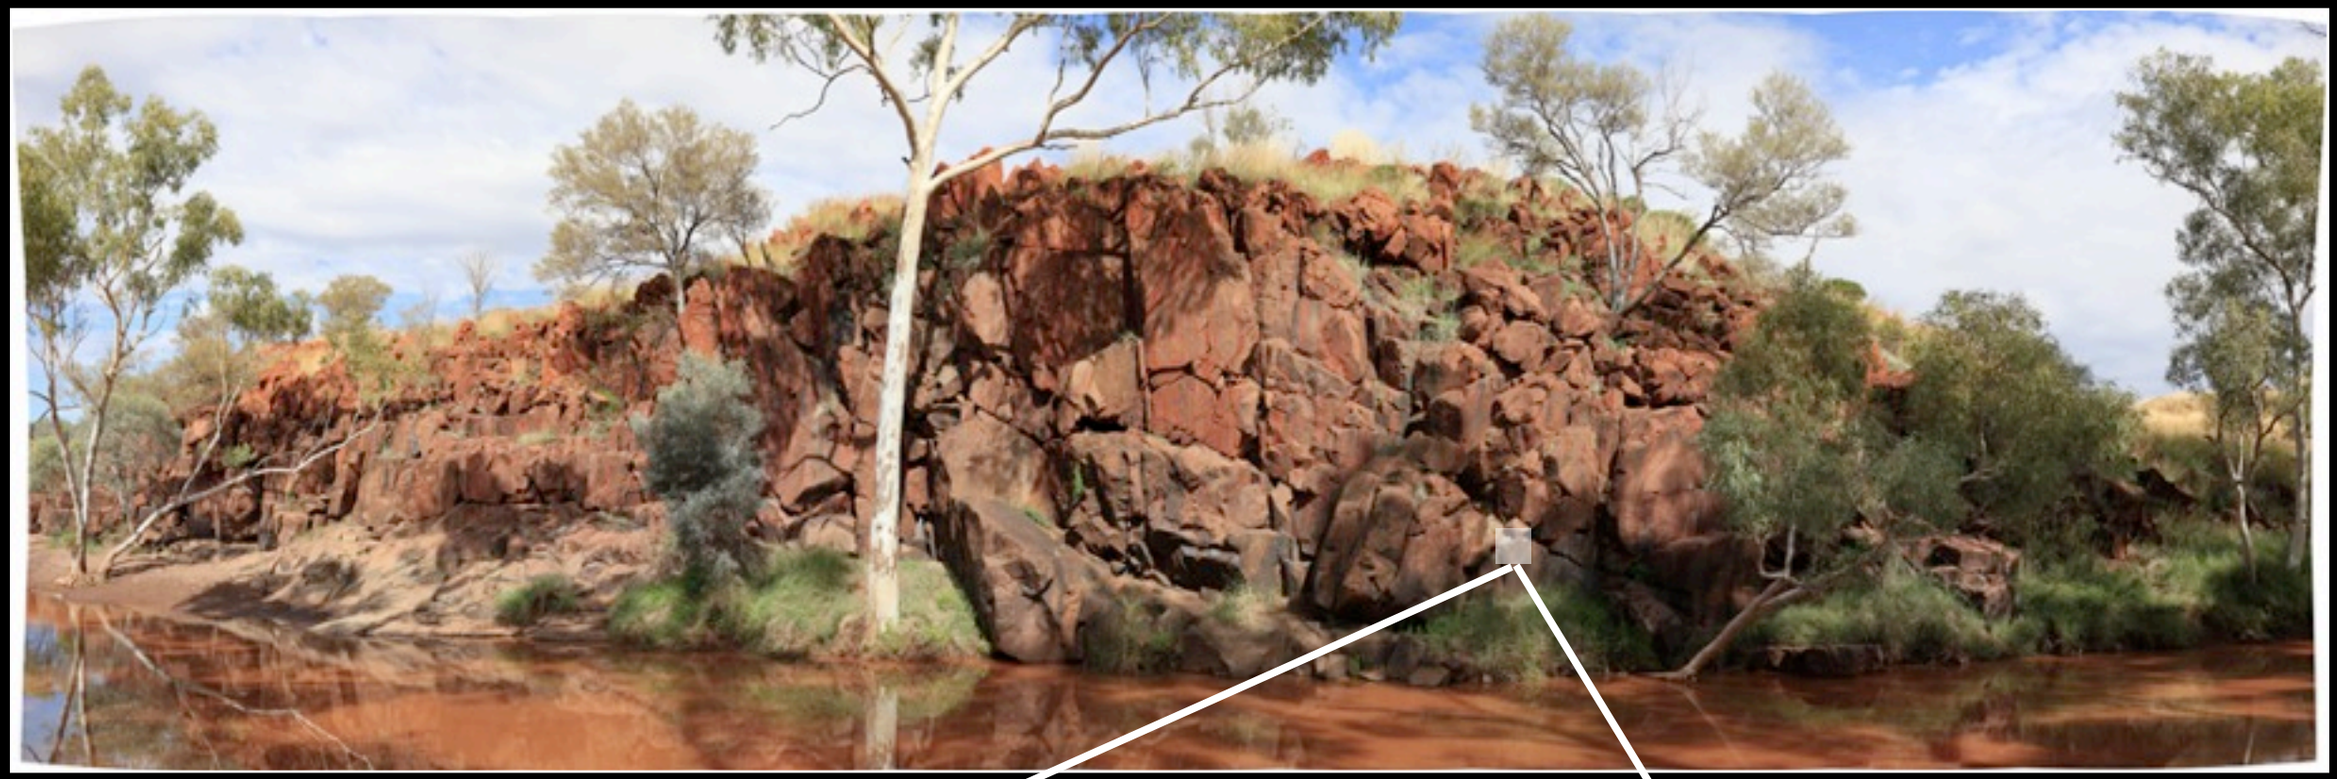

15 x 4 grid

Single 20MPixel image

Movie

Software challenges

- Are some issues with software to view/edit these images.
- Most image file formats for example don't support more than 32,000 pixels in width or height.
- Most viewers expect all the image to be in memory, may not be possible.
- There are "large image formats" and viewers capable of pyramidal multi-resolution formats.

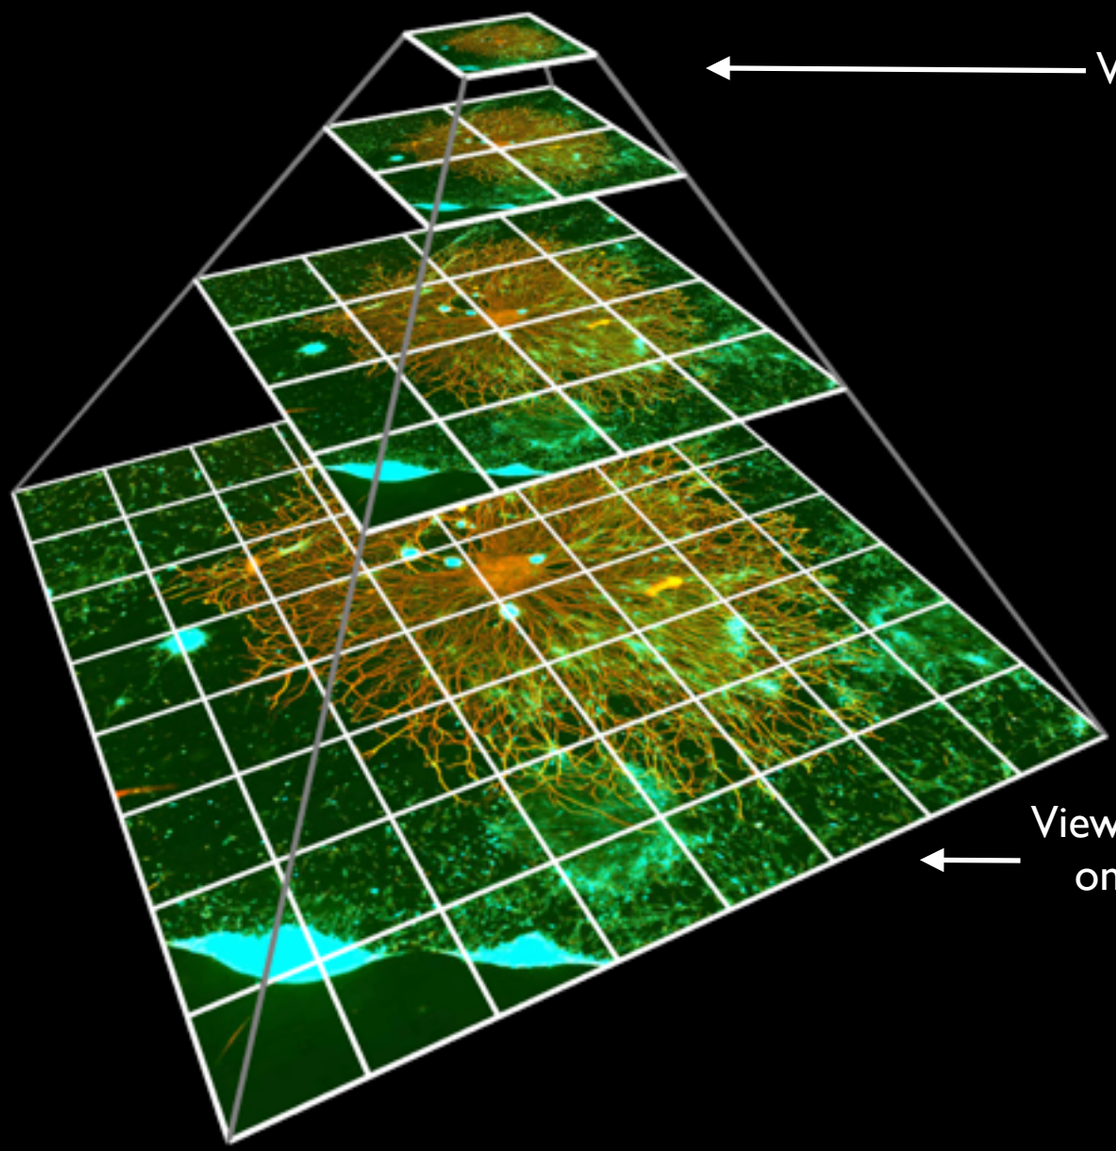

Viewing the entire dataset zoomed out

Viewing a portion of the dataset zoomed in,
only need a subset of the available tiles.

Beacon Island

- Site where the Batavia ship wreck survivors/victims came ashore.
- Project to record the site as it currently stands before fisherman huts are removed.
- Expect additional grave sites under the concrete slabs.

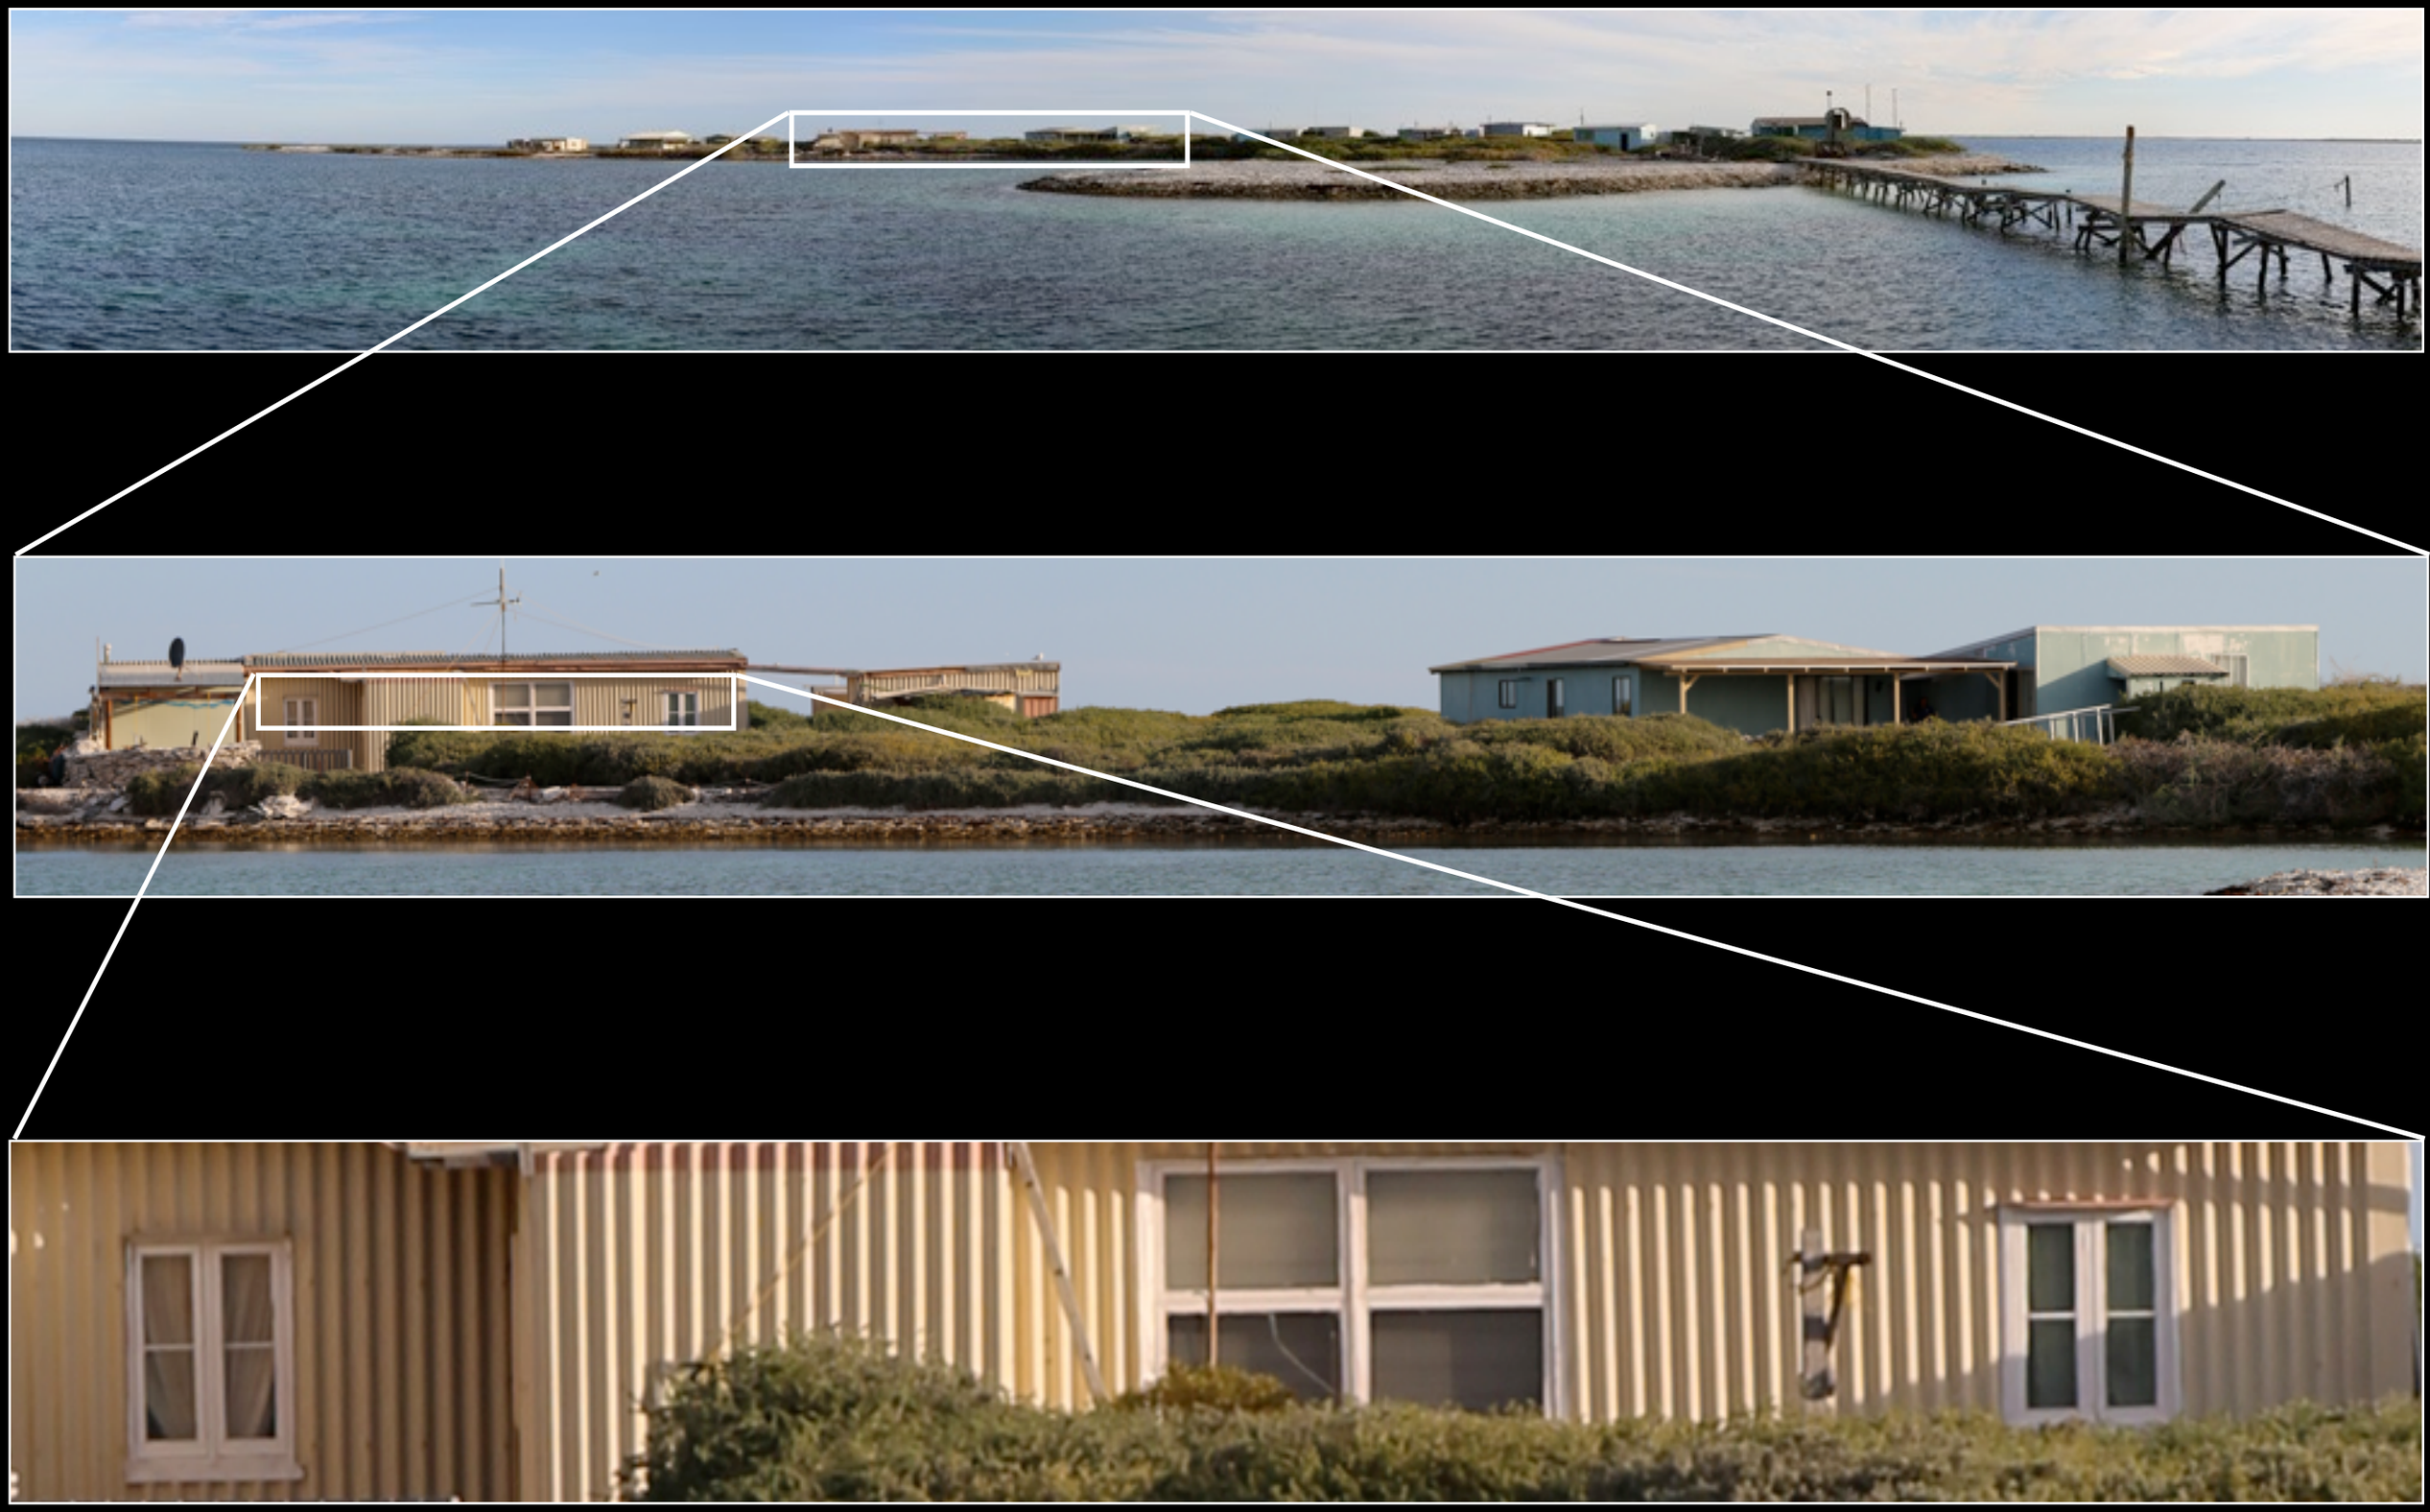

Automatic 3D reconstruction from photographs

- Photogrammetry: the general term give to deriving some 3D quality from a series of images.
- Traditionally used in landscape mapping, mostly 2.5 structure.
- Due to a number of recent algorithms it is being applied to the capture of full 3D objects.
- Fairly old history in Western Australia in mining and geology.
  For example, a cost effective way of determining volume of rock extracted in mining.
- Motivation and characteristics
  - Capturing 3D models of significant objects, richer data than just photographs.
  - Non-intrusive capture.
  - No specialist capture hardware.
  - Delivers texture and structure.
  - Fast acquisition of objects to populate virtual worlds.

Mine pit modelling

Movie

Geology

Movie

Rock Art: Wanmanna

Movie

Just 3 photographs!

Movie

Dragon gardens

- Heritage gardens in Hong Kong.
- Built by industrialist Wing Fat.
- Popularly known as the site for the 1974 James Bond movie “The man with the Golden Gun”.

Scene from the movie

Gardens today

Scene from the movie

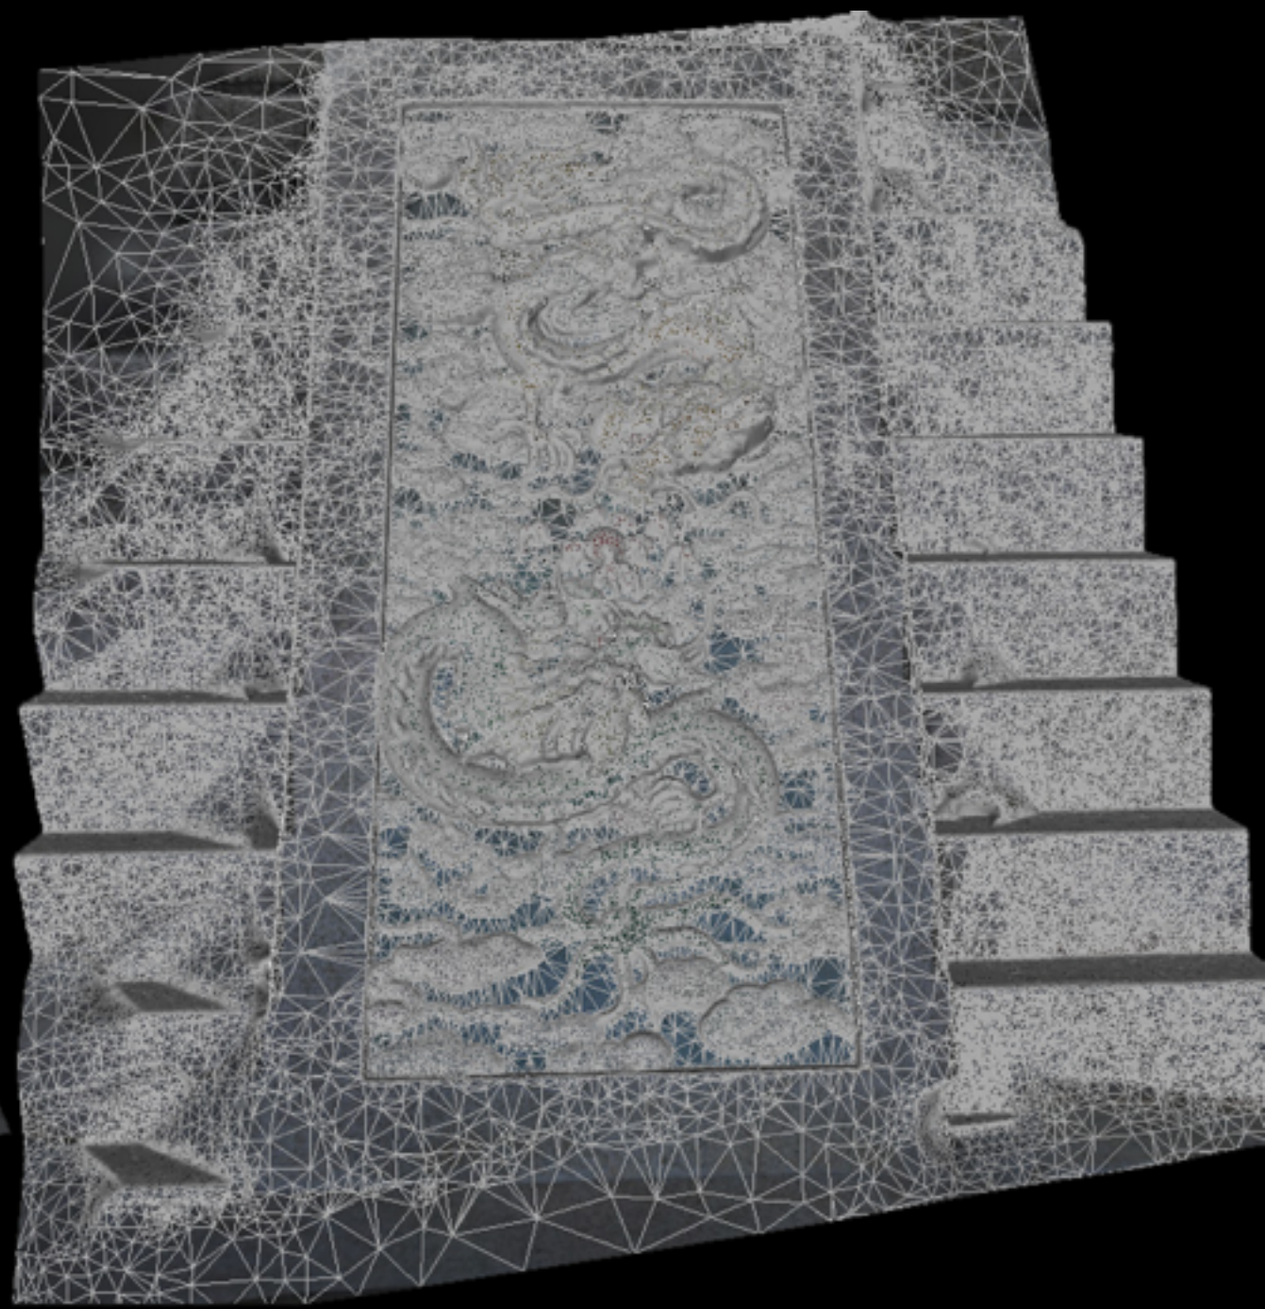




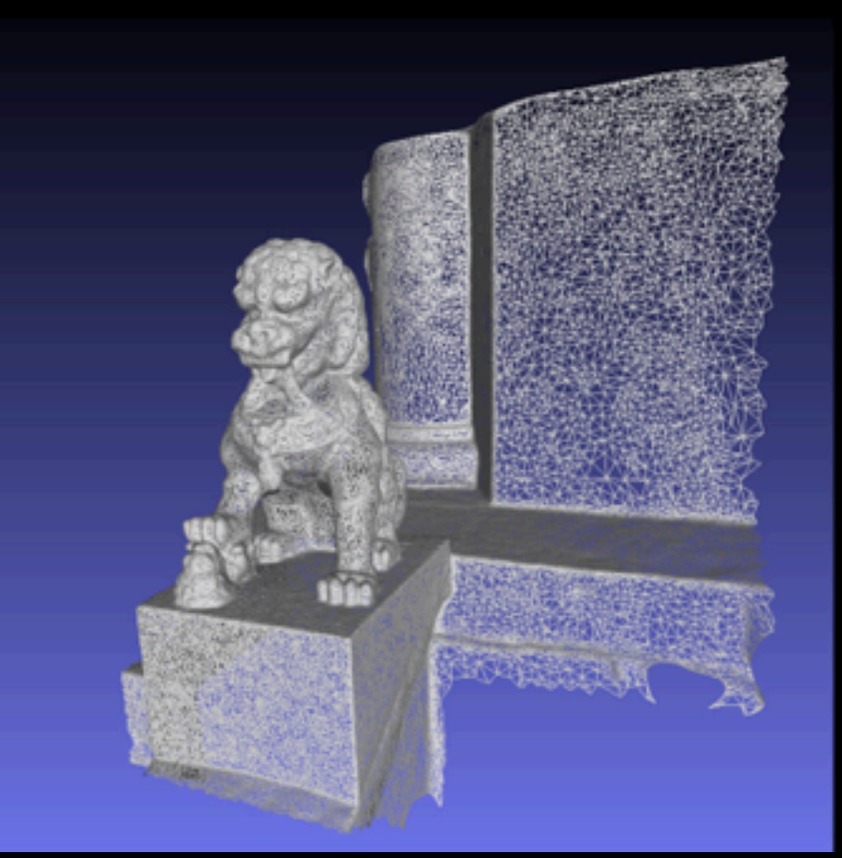
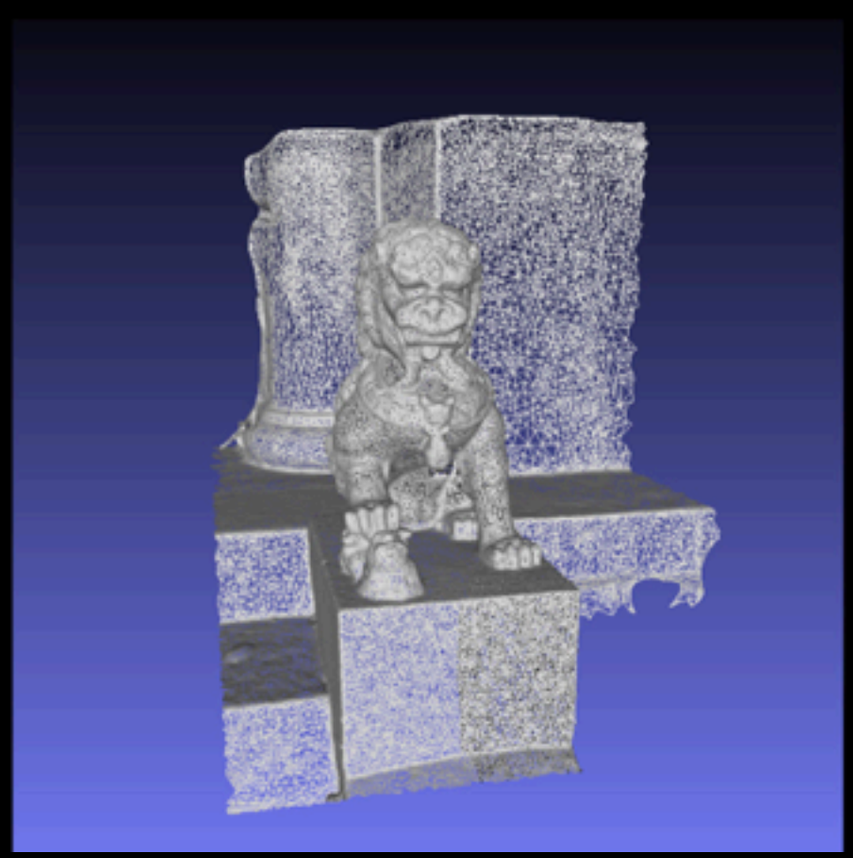
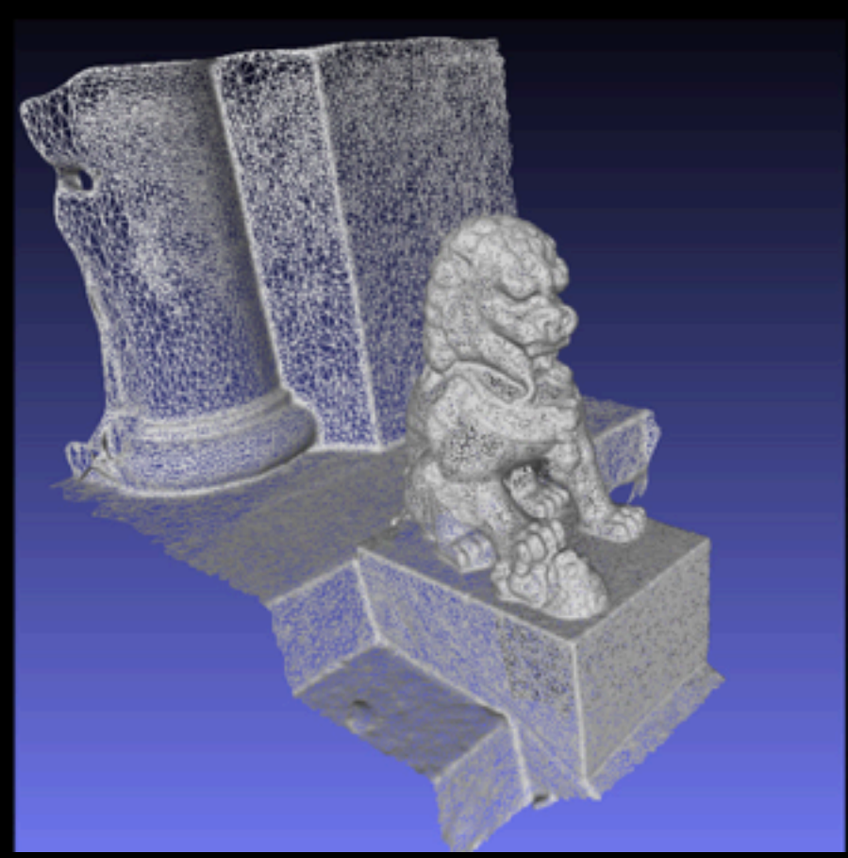

Ngintaka

- Story of grinding stone stolen by Ngintaka (A lizard).
- The story of Ngintaka is told at places across the landscape.

Grinding stone

Movie

Headdress

Movie

360 degree video recording and presentation

- Motivations:
  - Record everything occurring around a central point, nothing "off camera".
  - To capture data that can presented in a way as to create the sense of "being there".
- Applications of visualisation in cultural heritage: giving an insight into what it is like in a different culture, place or time.
- Two requirements
  - To capture the underlying video asset.
  - To present in an interactive environment, preferably one that fills the viewers FOV.
- One might imagine HMD (Head Mounted Displays) to be the vehicle but they still present very much a tunnel vision experience, at least for the affordable models.
- The approach often used for this work is a hemispherical or cylindrical display.
- A hemisphere or 1/2 hemisphere still requires navigation, a full cylinder doesn't, the video happens all around the viewer.
- Two examples
  - Ngintaka indiginous story/dance
  - Mah Meri dance/ritual

LadyBug camera

- PtGrey has produced 360 degree x 150 degree video cameras for some time.
  Distinction with most other 360 capture devices is the resolution.
- Target security and surveillance applications.
  Operator can see a full 360 field of view in a single camera shot.
- Remote operations.
- Performance recording and analysis.
- Sports science, presenting scenarios that more fully engage the human visual field.

Performance - Anatomy of a spherical projection

5400 x 2700 pixels

0 degrees longitude

360 degrees longitude

North pole, latitude: 90

Lower 40 degrees
not captured.

South pole, latitude -90

Cylindrical projection

Can derive cylindrical projection of any vertical field of view

Fisheye projections (Infinite number)

Remote operations

Movie

Sports science

Example: Ngintaka (Cultural heritage)

- One location for the Ngintaka story after the grind stone has been recovered is in a cave.
- The belly of Ngintaka.

360 x 140 degree video example

Movie

Navigable example in the iDome

Movie

Mah Meri

- Tribe in West Malaysia.
- Traditional healing involves a priestess determining the form of the sickness.
- A carver fashions a mask representing that entity.
- A ritual dance is performed around the patient who wears the mask.
- The mask is then released into the river, taking the illness with it.
- As such, not many masks in existence.
- A repertoire of up to 400 illness spirits.

Sample LadyBug footage

Movie

Fisheye view

Movie

Ritual area

Priestess building

Sure enough, came down with an as yet undiagnosed illness the next day, have only fully recovered in the last week or so.

Closing Message: Don’t mess with the local spirits!