3D reconstruction from photographs

(Photogrammetry)

Paul Bourke

http://paulbourke.net/EPFL2019/






?

Outcomes

- A overview of the technology
- When it is appropriate to use it
- Expectations of the outputs
- Workflow if you were to use the technique

Why?

- Creating richer objects (compared to photographs) for recordings in archaeology and heritage.
- Create geometric models suitable for analysis, eg: in geology or geoscience.
- Create models for virtual environments.
- Create models for augmented reality.
- Assets for educational and public outreach.

History

- Photogrammetry is the general term given to deriving geometric information from a series of photographs.
- Initially largely used for aerial surveys, deriving landscape models.
  Originally only used a stereoscopic pair, that is, just two photographs.
- More recently it has grown through the domain of machine vision, for example: deriving a 3D model of a robots environment.
- Big step forward was the development of SfM algorithms: structure from motion.
  This generally solves the camera parameters and generation of a 3D point cloud.
- Most common implementation is called Bundler: “bundle adjustment algorithm allows the reconstruction of the 3D geometry of the scene by optimizing the 3D location of key points, the location/orientation of the camera, and its intrinsic parameters”.

Order of business

- Application examples
- When
- Hardware
- Photographic technique
- Software workflow
- Other topics
- Post production
- Worked example
- Your turn

Application examples

- Heritage
- Archaeology
- Forensics
- Medical training
- Replication, 3D printing

Weld range rock shelter

Mowunjum rock art

Beacon Island, Batavia shipwreck site

Dragon gardens in Hong Kong

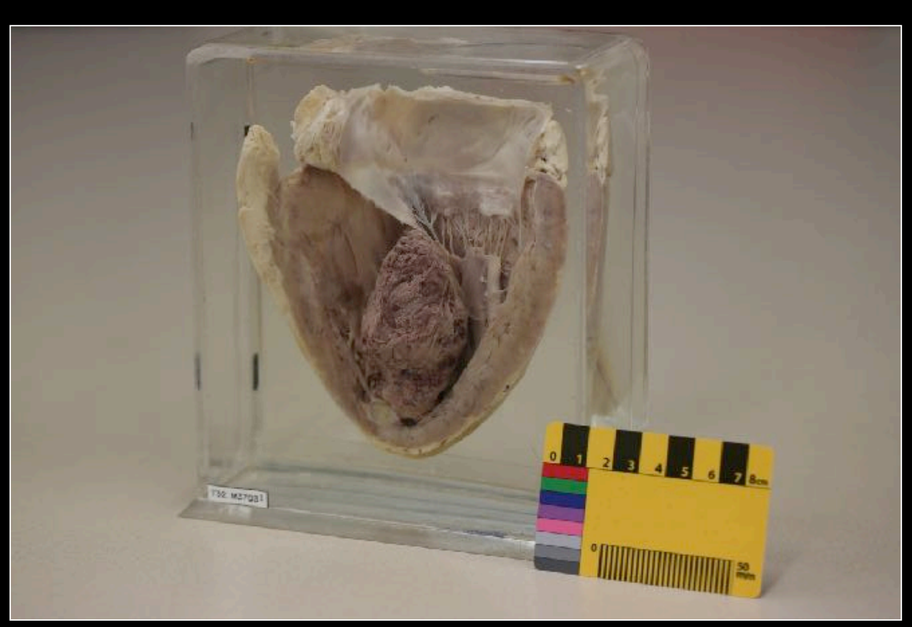
0 1 2 3 4 5 6 7 8cm
0 50 mm

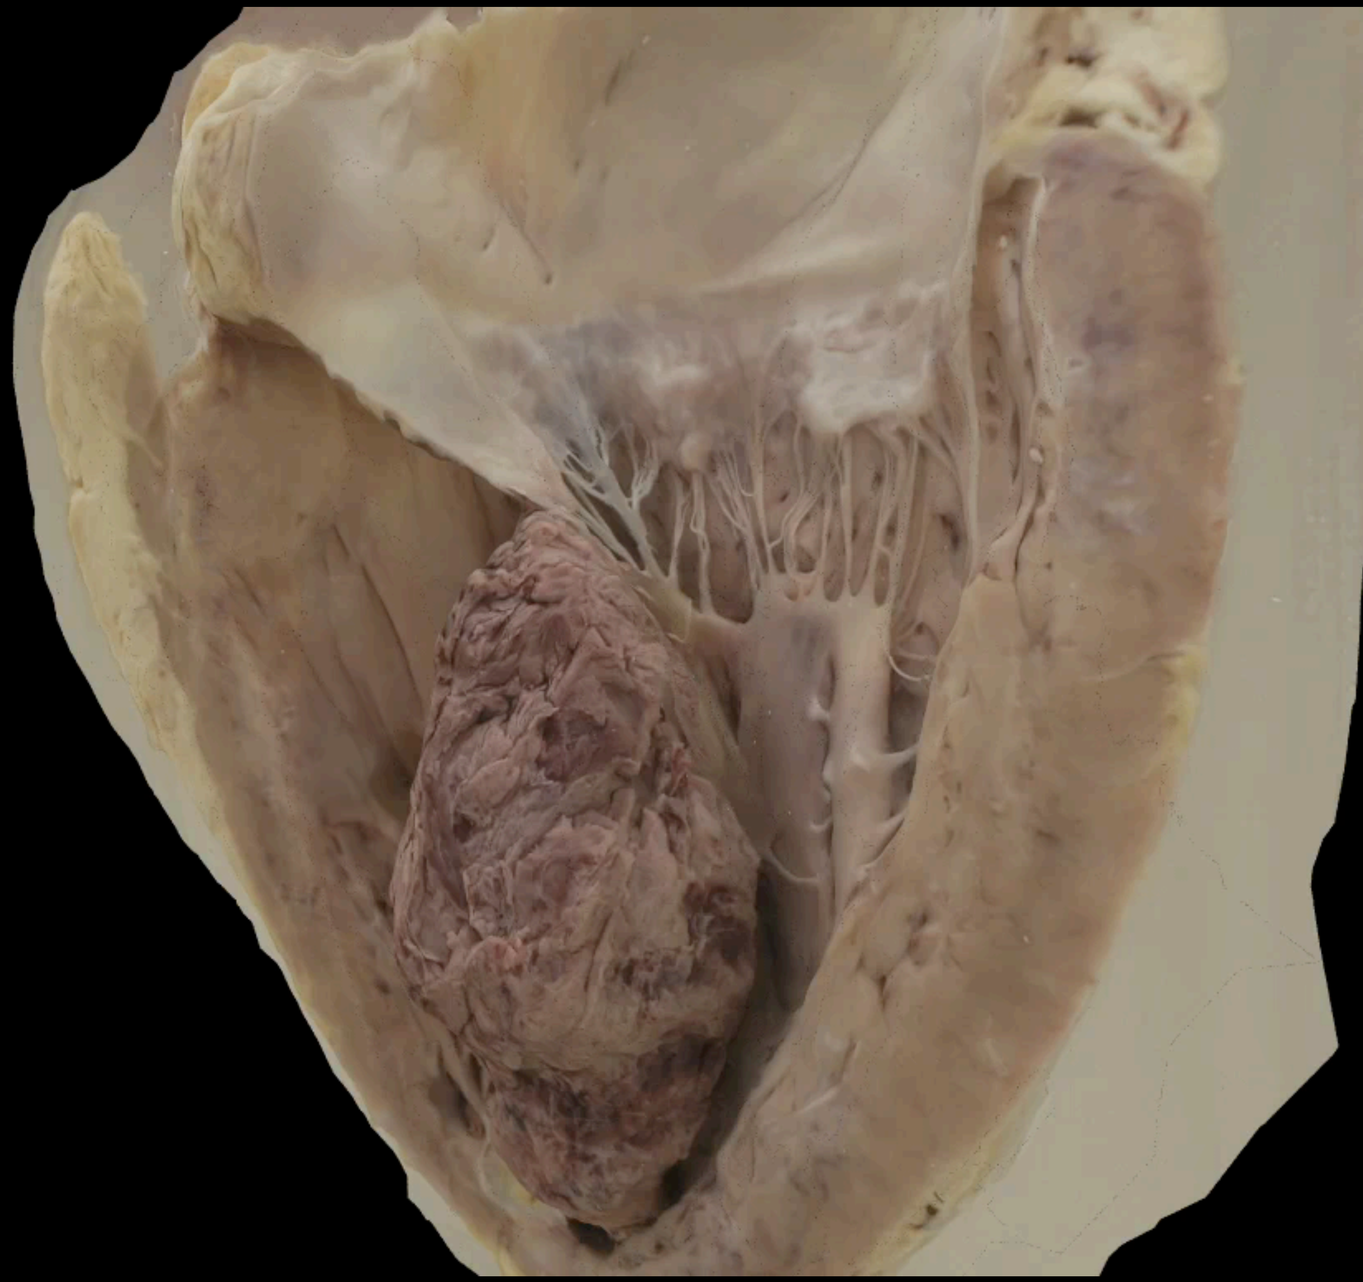

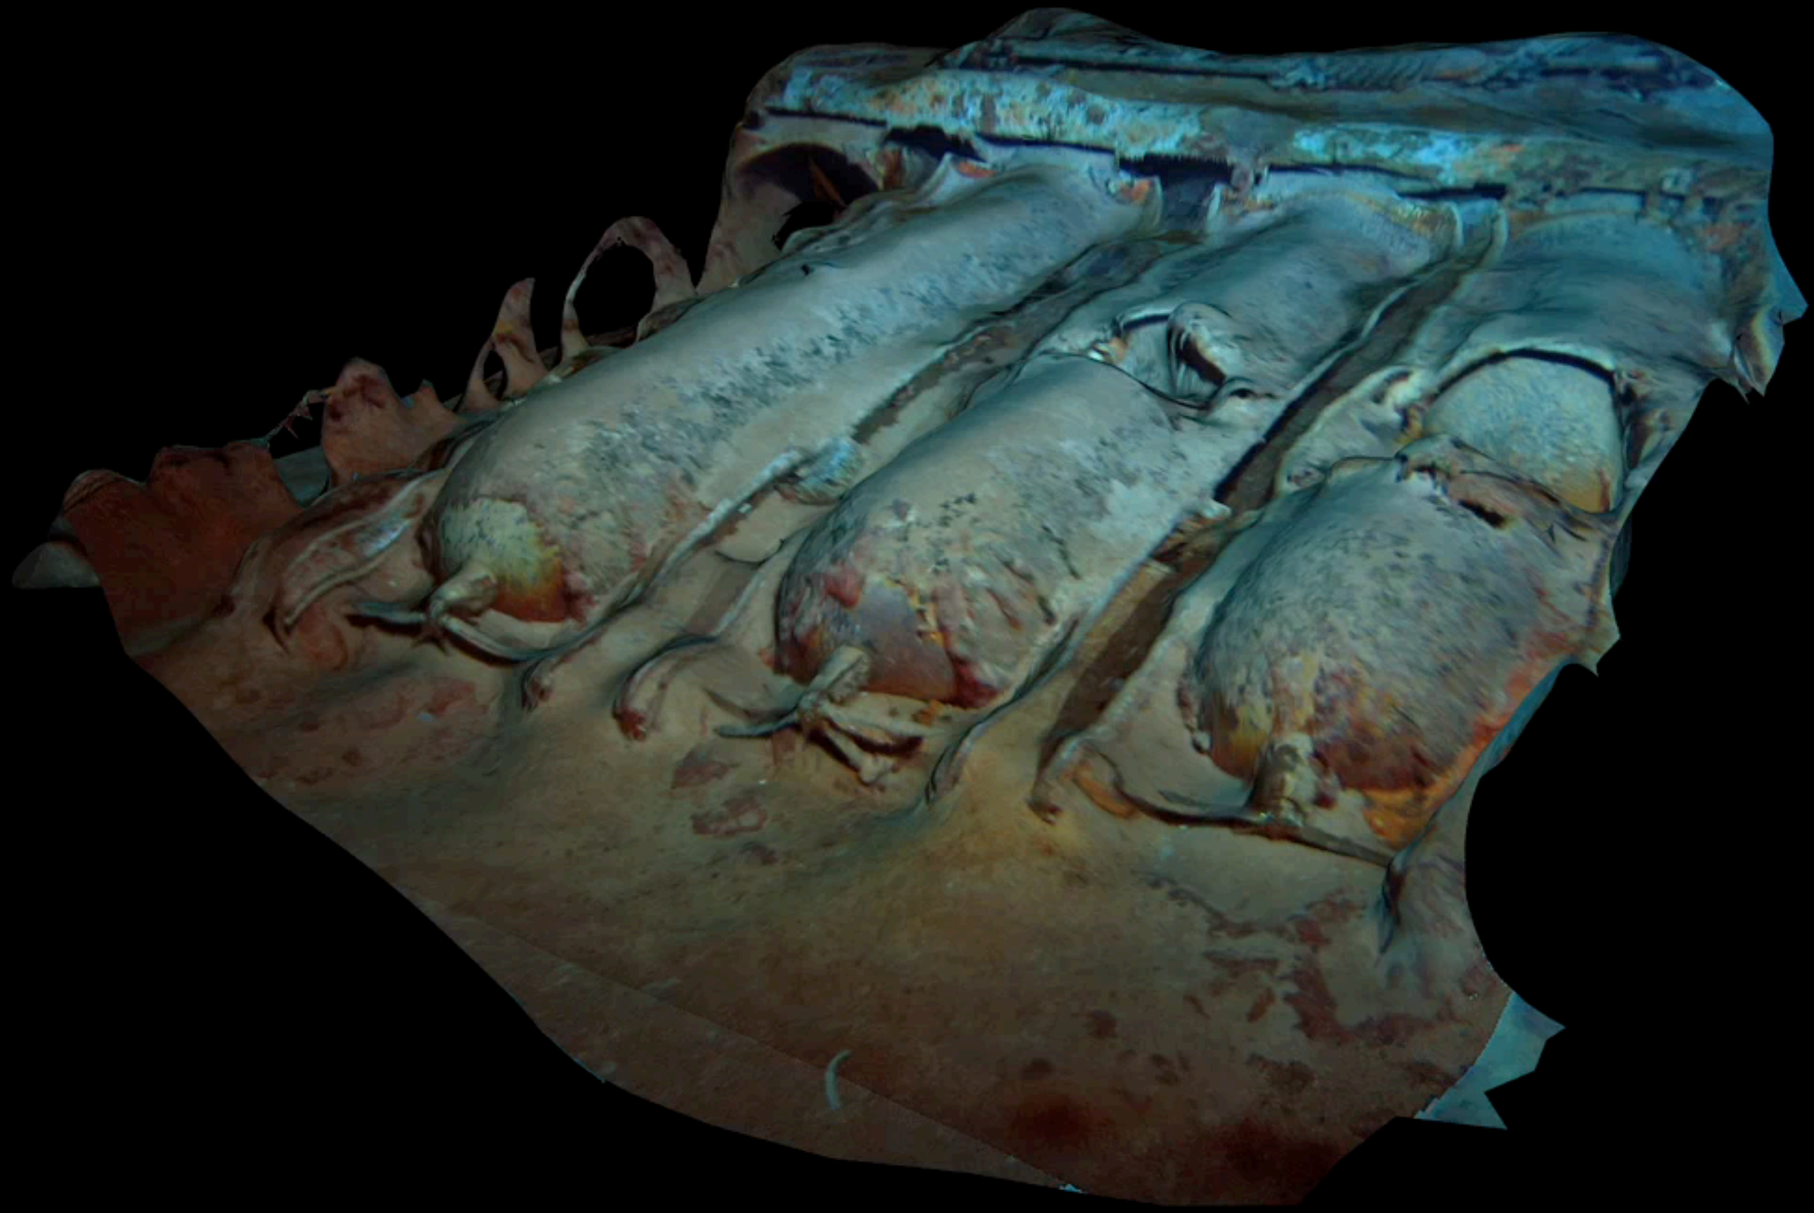

HMAS Sydney


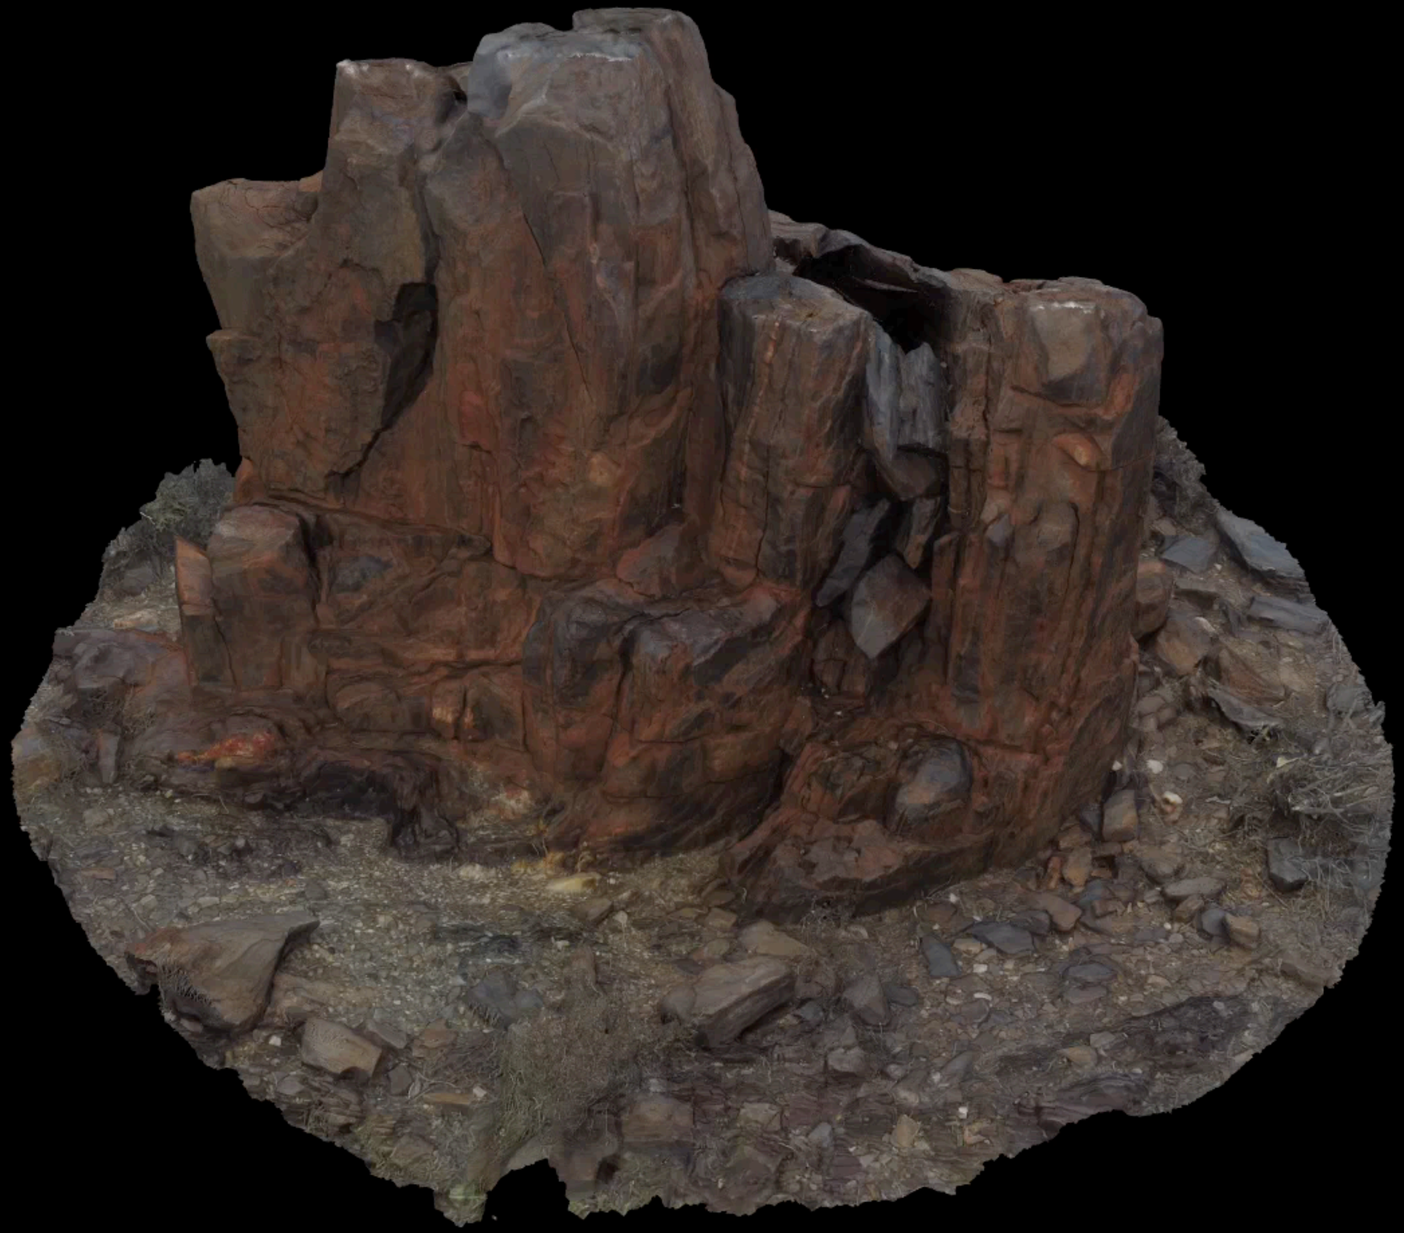

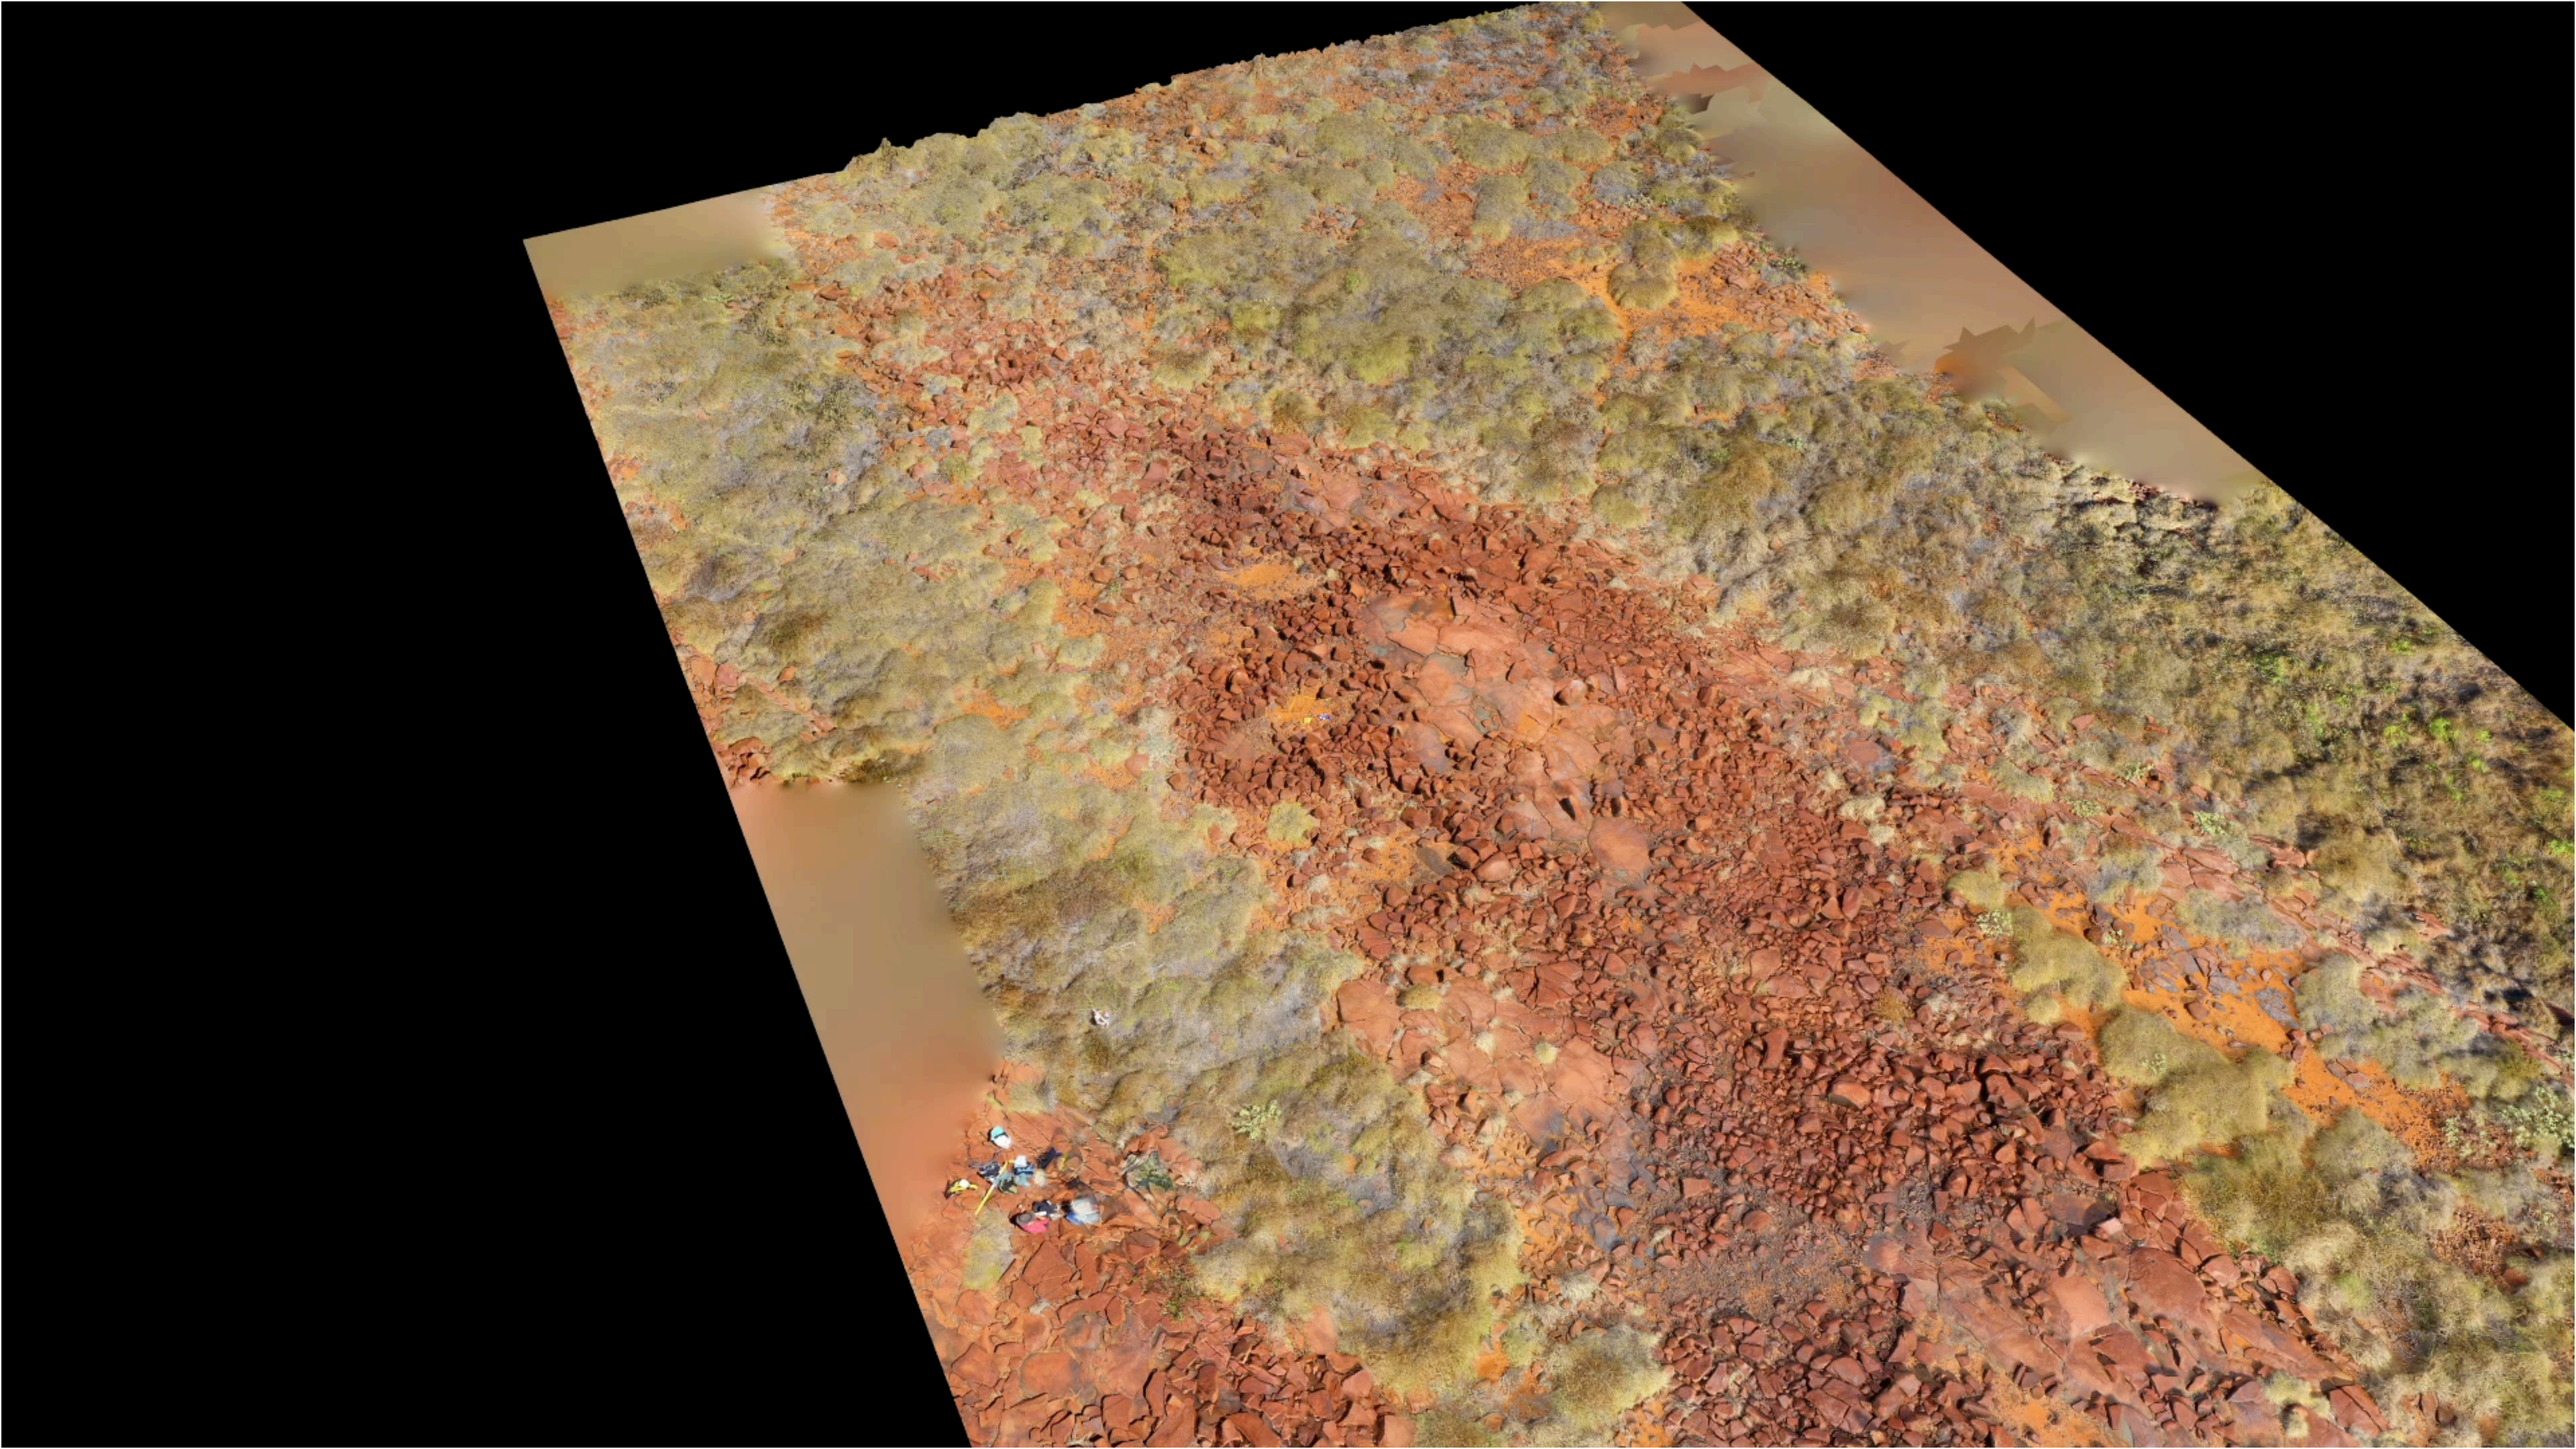

When

- Best suited to organic shapes that are otherwise hard to model.
- Provides an unbiased model compared to human modeller.
- Requires feature point extraction (see later), so well suited to texture rich and noisy surfaces.
- Some models the technique is less suited to
  - Plain flat surfaces
  - Highly reflective or mirror surfaces
- Not suited to extremely thin objects
- Not suited to objects that are moving, unless using a multiple camera rig.

Yes

No

Yes


THIS UNDERCROFT
IS DEDICATED TO
SOCRATES
WHO SOUGHT TRUTH
ALWAYS BY THE PATH
OF OPEN DISCUSSION
AND FREE ENQUIRY
MAY HIS SPIRIT
PRESIDE HERE
AT ALL TIMES

No

Yes

No

Beacon Island

The University of Western Australia

A small change in camera position results in a large difference in visible objects.

Hardware

- Camera.
- Prime lens, fixed focal point.

  Generally better results for better quality camera and optics.

- Tripod for long exposures in low light.
- Possibly a turntable for small objects.

Other methods

- Laser scanning, return of flight. (LIDAR - light detection and ranging)
- Depth cameras, also typically return of flight but different wavelengths.
  eg: Primesense (eg: Kinect)
- Structured light.
  eg: Artec Scanner
- Light field cameras.

There is no winning technology. Each has limitations and applications where they may be better than other techniques.

Photographic technique

- Sharp in focus photographs with good depth of focus.
- High resolution photographs.
- Use manual settings for consistent textures: white point, ISO, exposure time, aperture.
- Photograph all visible parts of the model (One cannot reconstruct what is not photographed!) + even spread of photographs.
- Every camera position different.
  Opposite to panoramic photography where one tries to photograph from the exact same position.
- Overall aim: any point on the model should be visible from at least 3 photographs, preferably more.
- Different strategies depending on the type of object.
- Camera orientation does not matter.


DSCF2714.JPG
DSCF2715.JPG
DSCF2716.JPG
DSCF2717.JPG
DSCF2718.JPG
DSCF2719.JPG

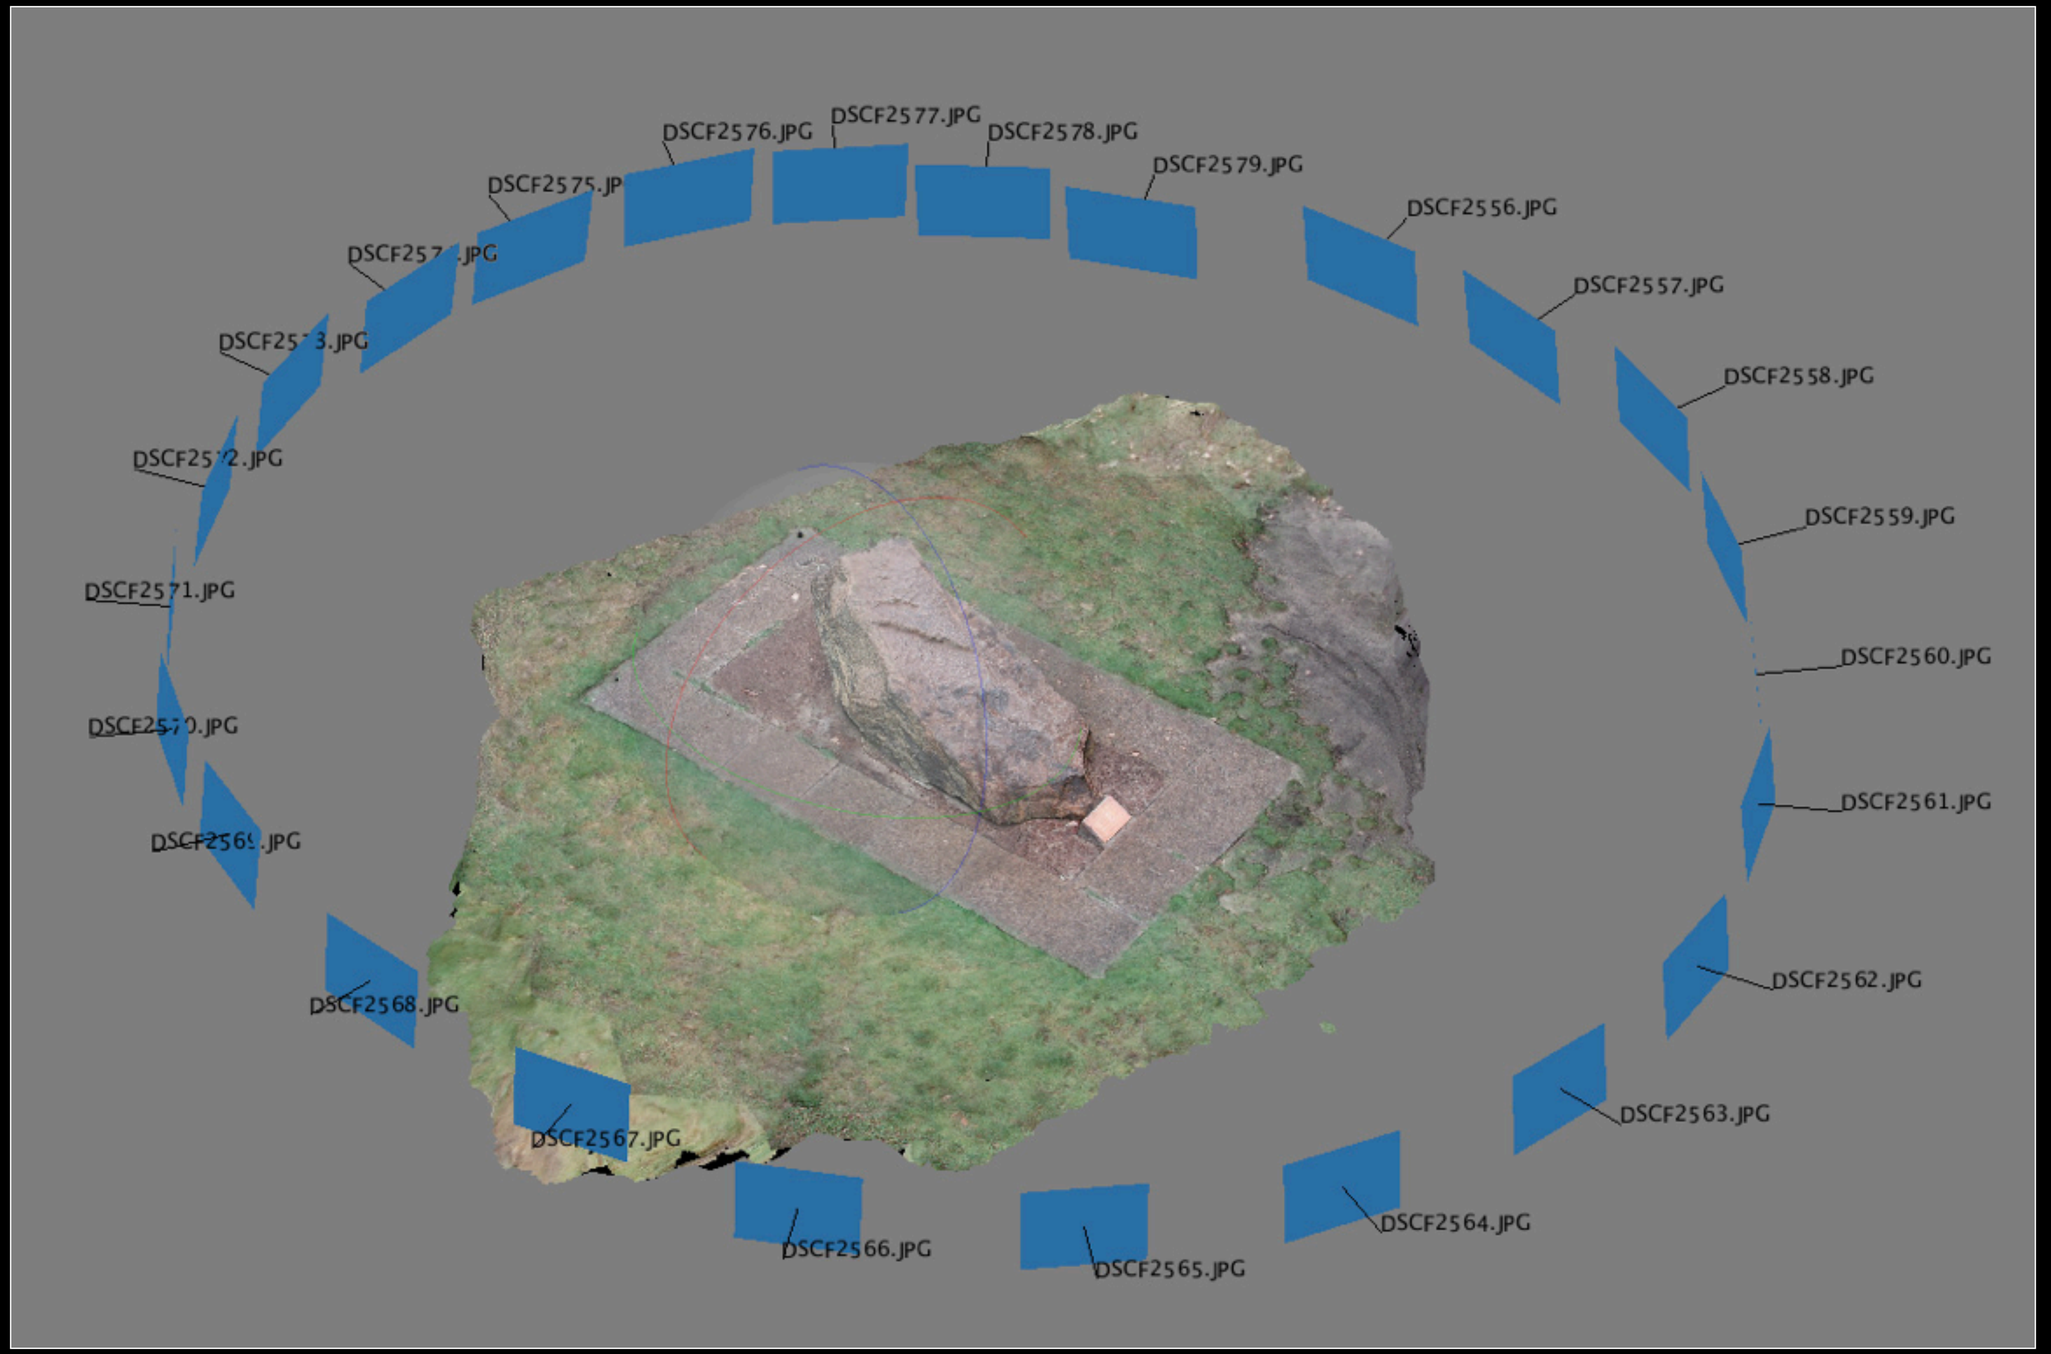
DSCF2576.JPG
DSCF2577.JPG
DSCF2578.JPG
DSCF2579.JPG
DSCF2556.JPG
DSCF2557.JPG
DSCF2558.JPG
DSCF2559.JPG
DSCF2560.JPG
DSCF2561.JPG
DSCF2562.JPG
DSCF2563.JPG
DSCF2564.JPG
DSCF2565.JPG
DSCF2566.JPG
DSCF2567.JPG
DSCF2568.JPG
DSCF2571.JPG

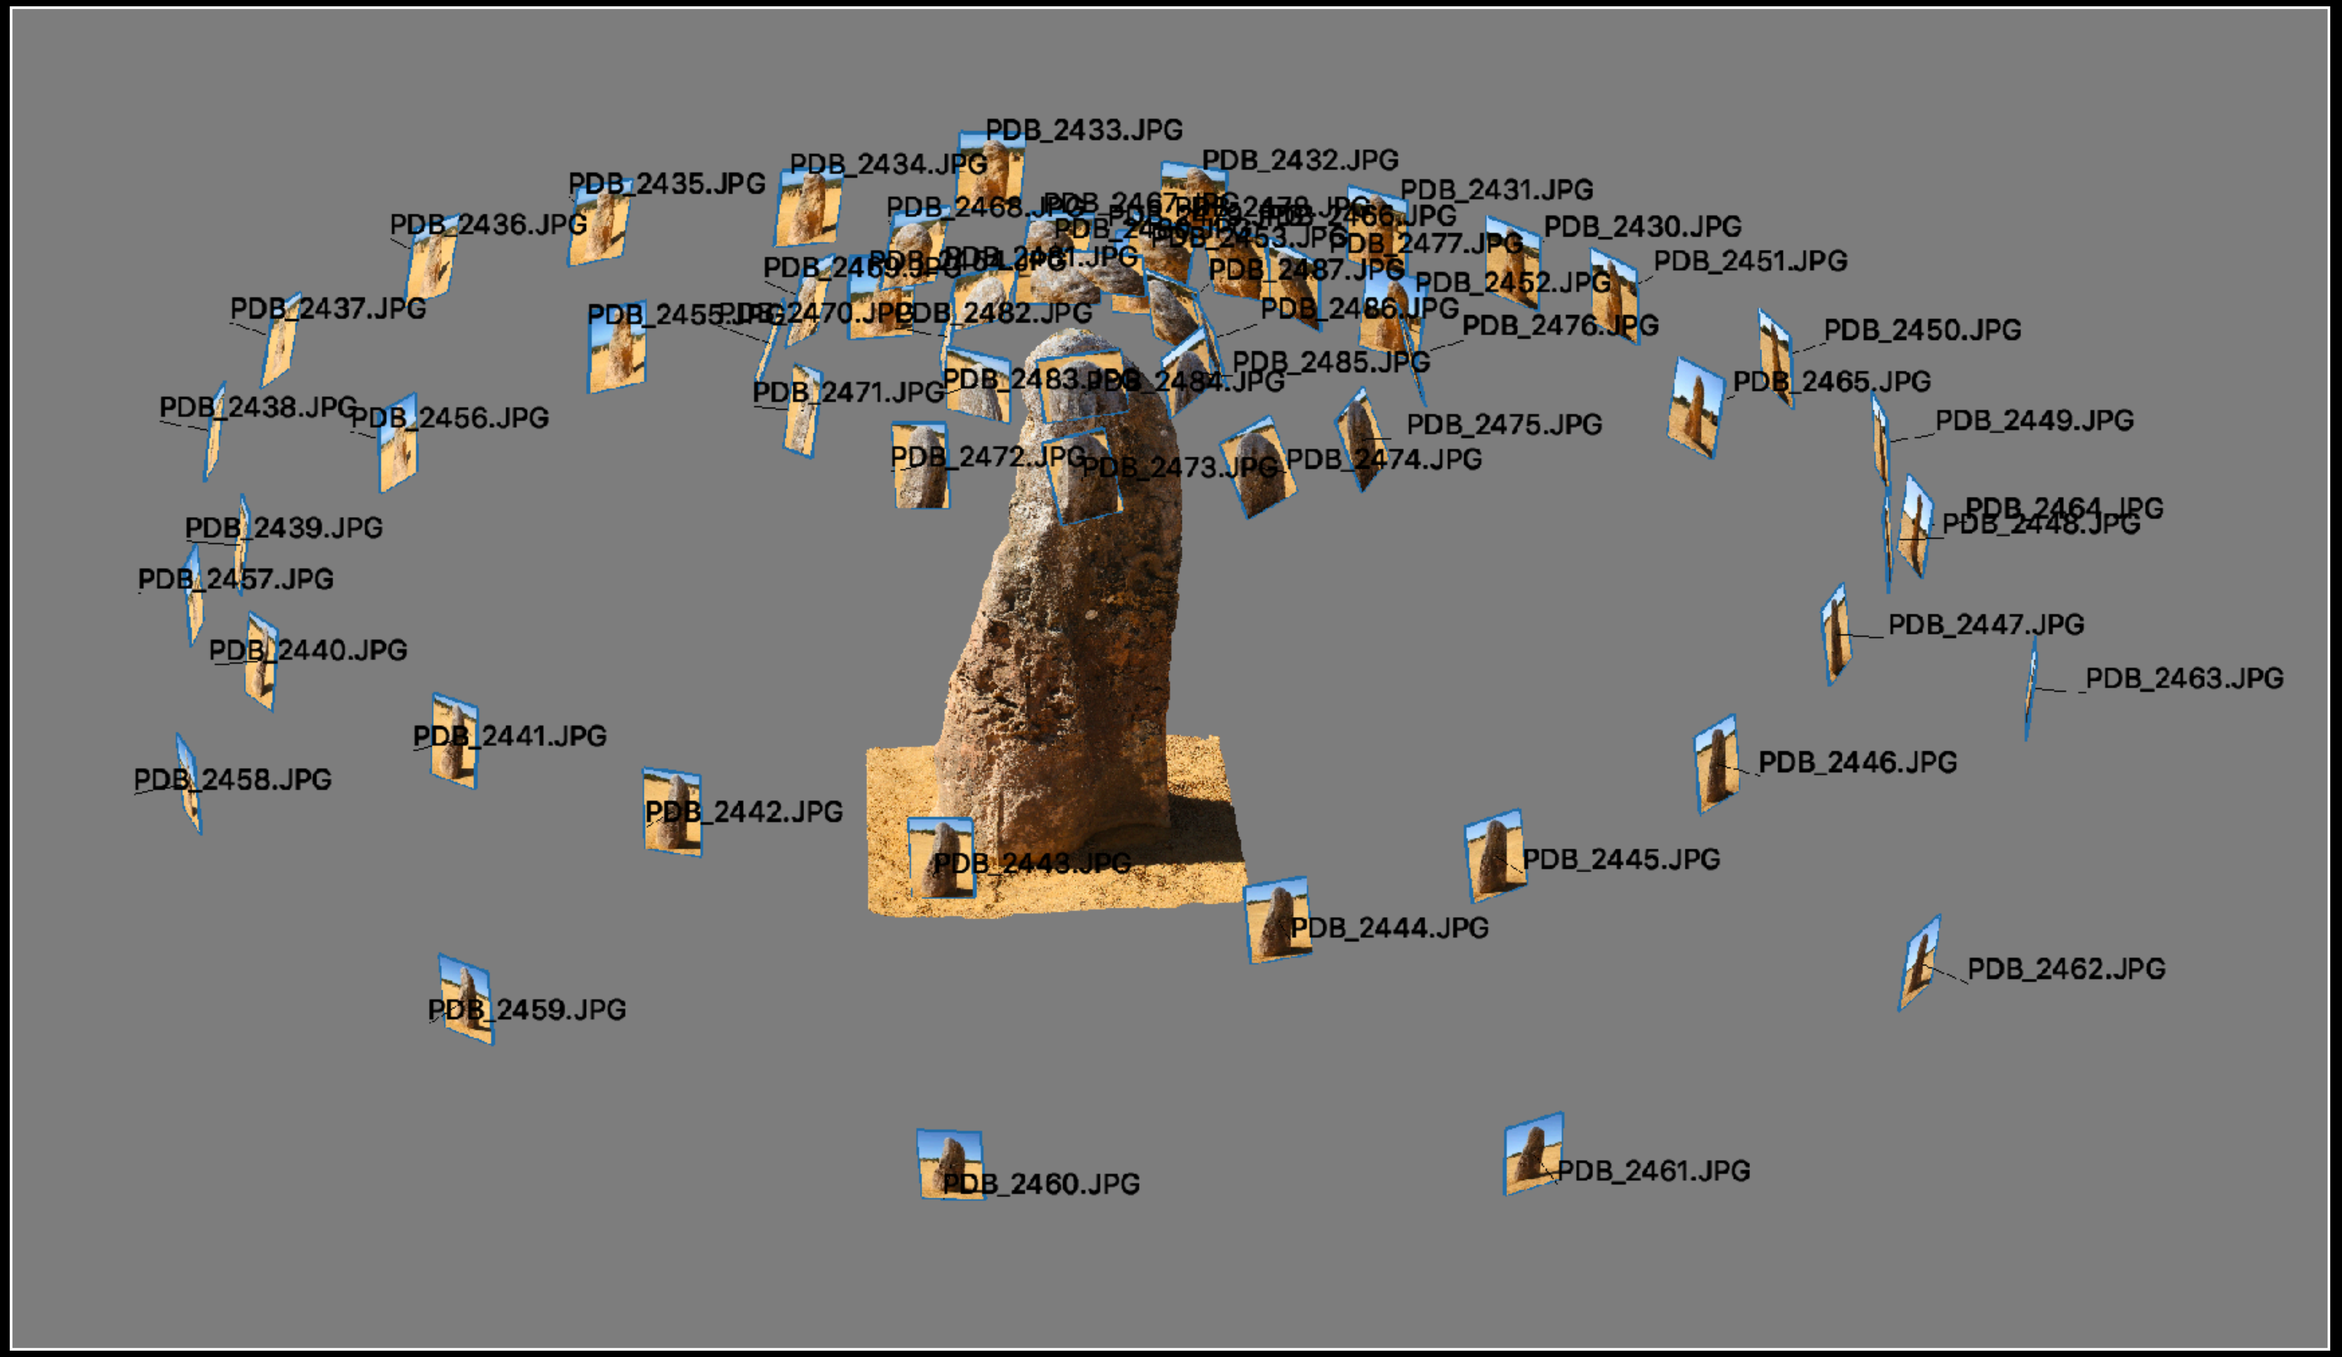
PDB_2433.JPG
PDB_2434.JPG
PDB_2435.JPG
PDB_2436.JPG
PDB_2437.JPG
PDB_2438.JPG
PDB_2439.JPG
PDB_2457.JPG
PDB_2440.JPG
PDB_2458.JPG
PDB_2441.JPG
PDB_2442.JPG
PDB_2459.JPG
PDB_2460.JPG
PDB_2443.JPG
PDB_2444.JPG
PDB_2445.JPG
PDB_2461.JPG
PDB_2446.JPG
PDB_2462.JPG
PDB_2463.JPG
PDB_2447.JPG
PDB_2464.JPG
PDB_2448.JPG
PDB_2449.JPG
PDB_2465.JPG
PDB_2450.JPG
PDB_2451.JPG
PDB_2430.JPG
PDB_2431.JPG
PDB_2432.JPG
PDB_2452.JPG
PDB_2476.JPG
PDB_2475.JPG
PDB_2474.JPG
PDB_2485.JPG
PDB_2486.JPG
PDB_2487.JPG
PDB_2477.JPG
PDB_2468.JPG
PDB_2469.JPG
PDB_2455.JPG
PDB_2470.JPG
PDB_2482.JPG
PDB_2471.JPG
PDB_2483.JPG
PDB_2484.JPG
PDB_2472.JPG
PDB_2473.JPG

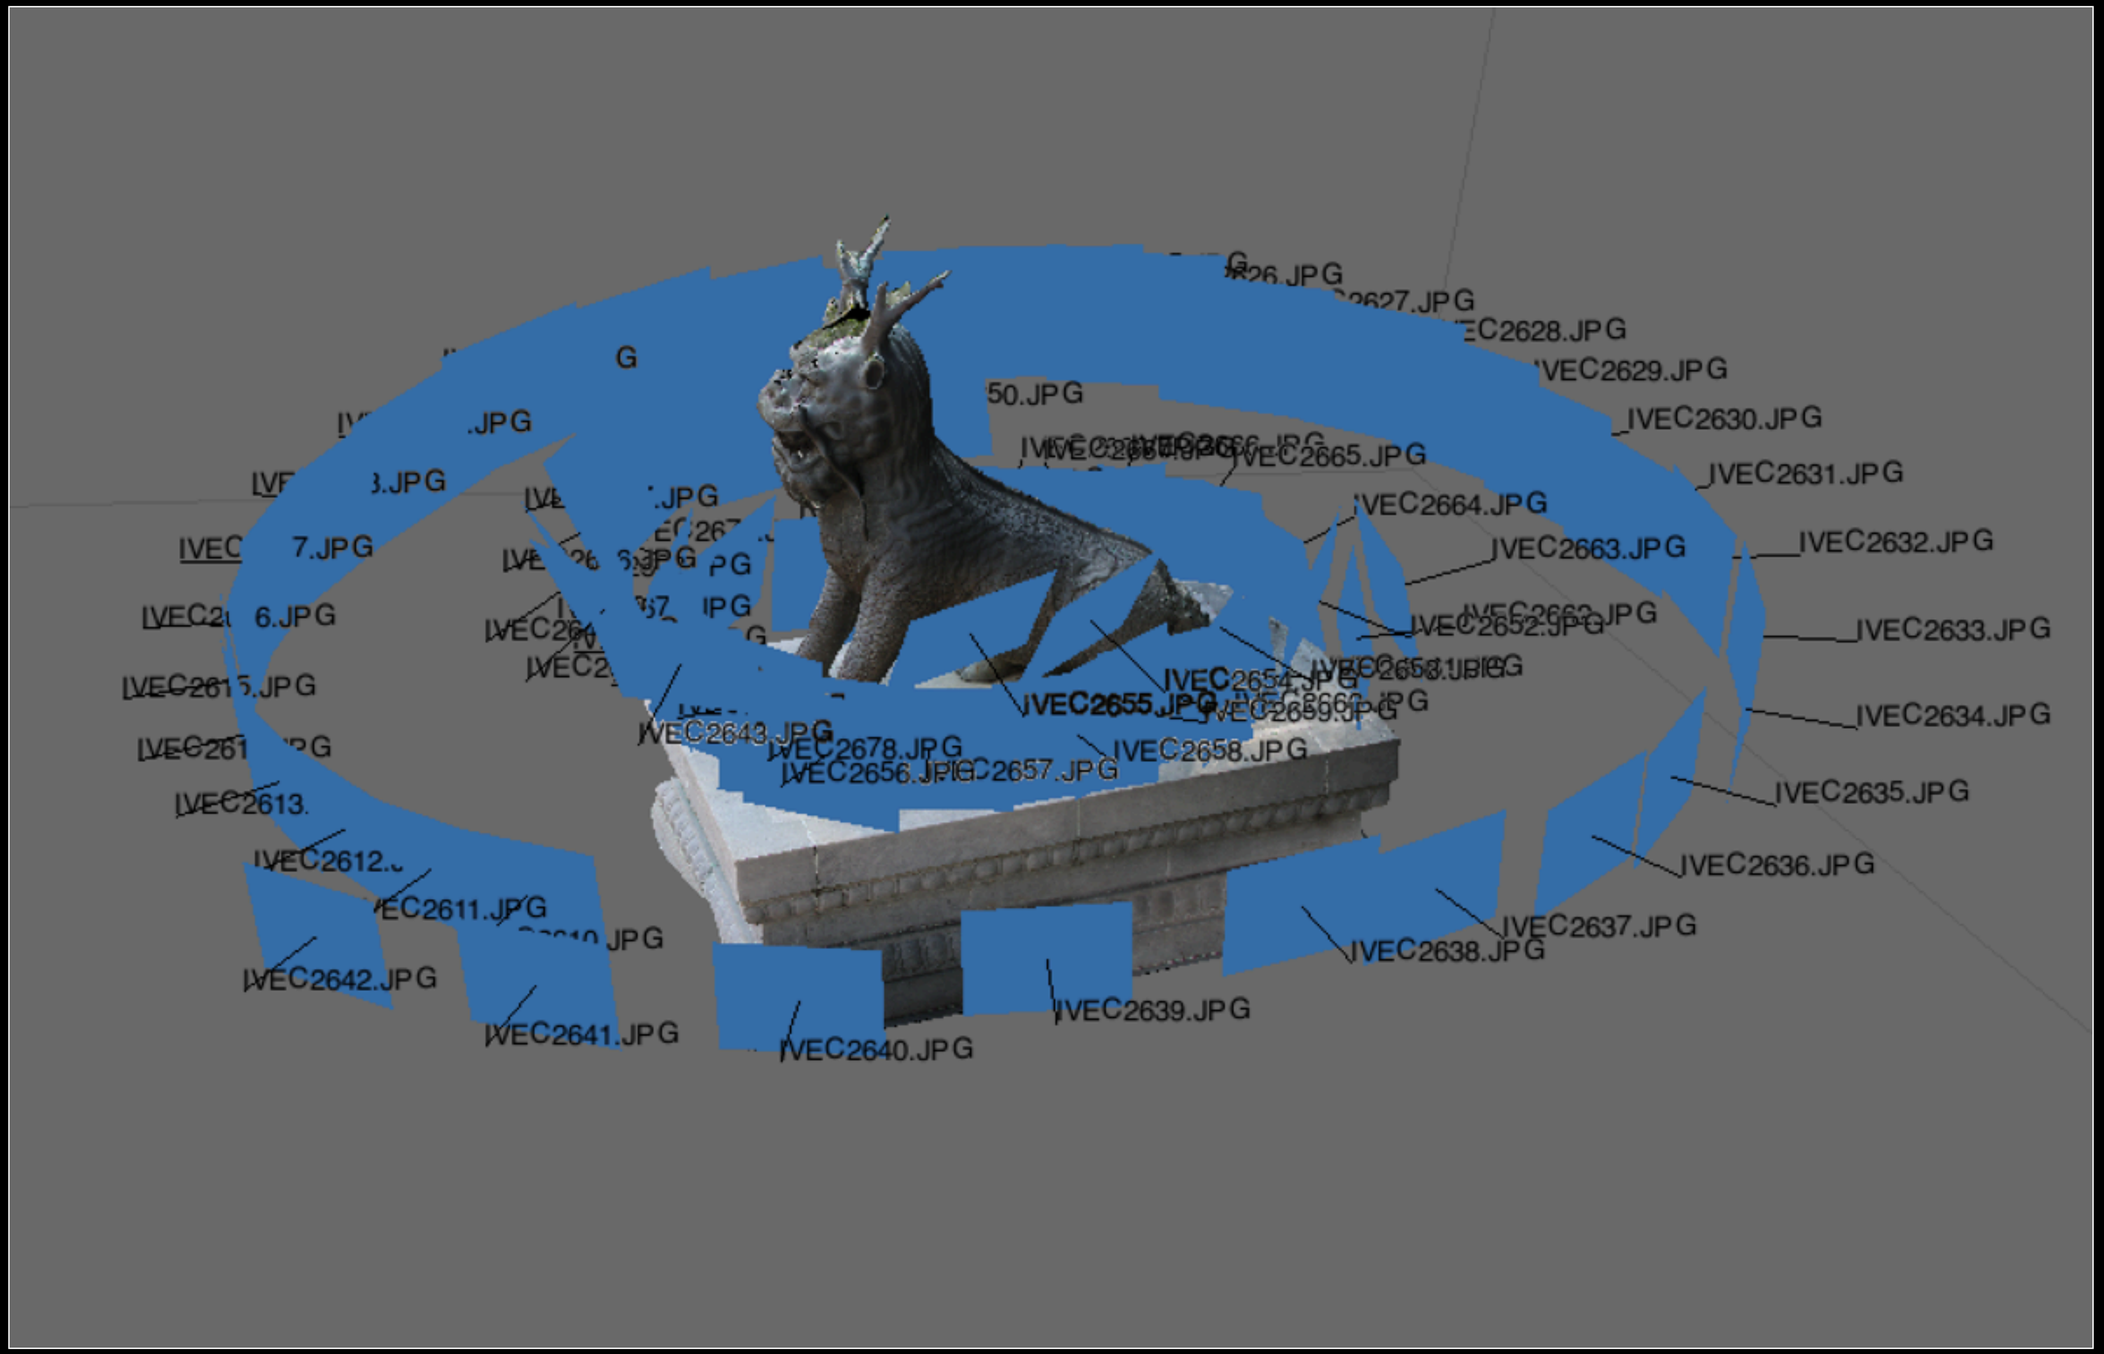
IVEC2628.JPG
IVEC2629.JPG
IVEC2630.JPG
IVEC2631.JPG
IVEC2632.JPG
IVEC2633.JPG
IVEC2634.JPG
IVEC2635.JPG
IVEC2636.JPG
IVEC2637.JPG
IVEC2638.JPG
IVEC2639.JPG
IVEC2640.JPG
IVEC2641.JPG
IVEC2642.JPG
IVEC2611.JPG
IVEC2643.JPG
IVEC2678.JPG
IVEC2656.JPG
IVEC2657.JPG
IVEC2658.JPG
IVEC2655.JPG
IVEC2654.JPG
IVEC2652.JPG
IVEC2663.JPG
IVEC2665.JPG

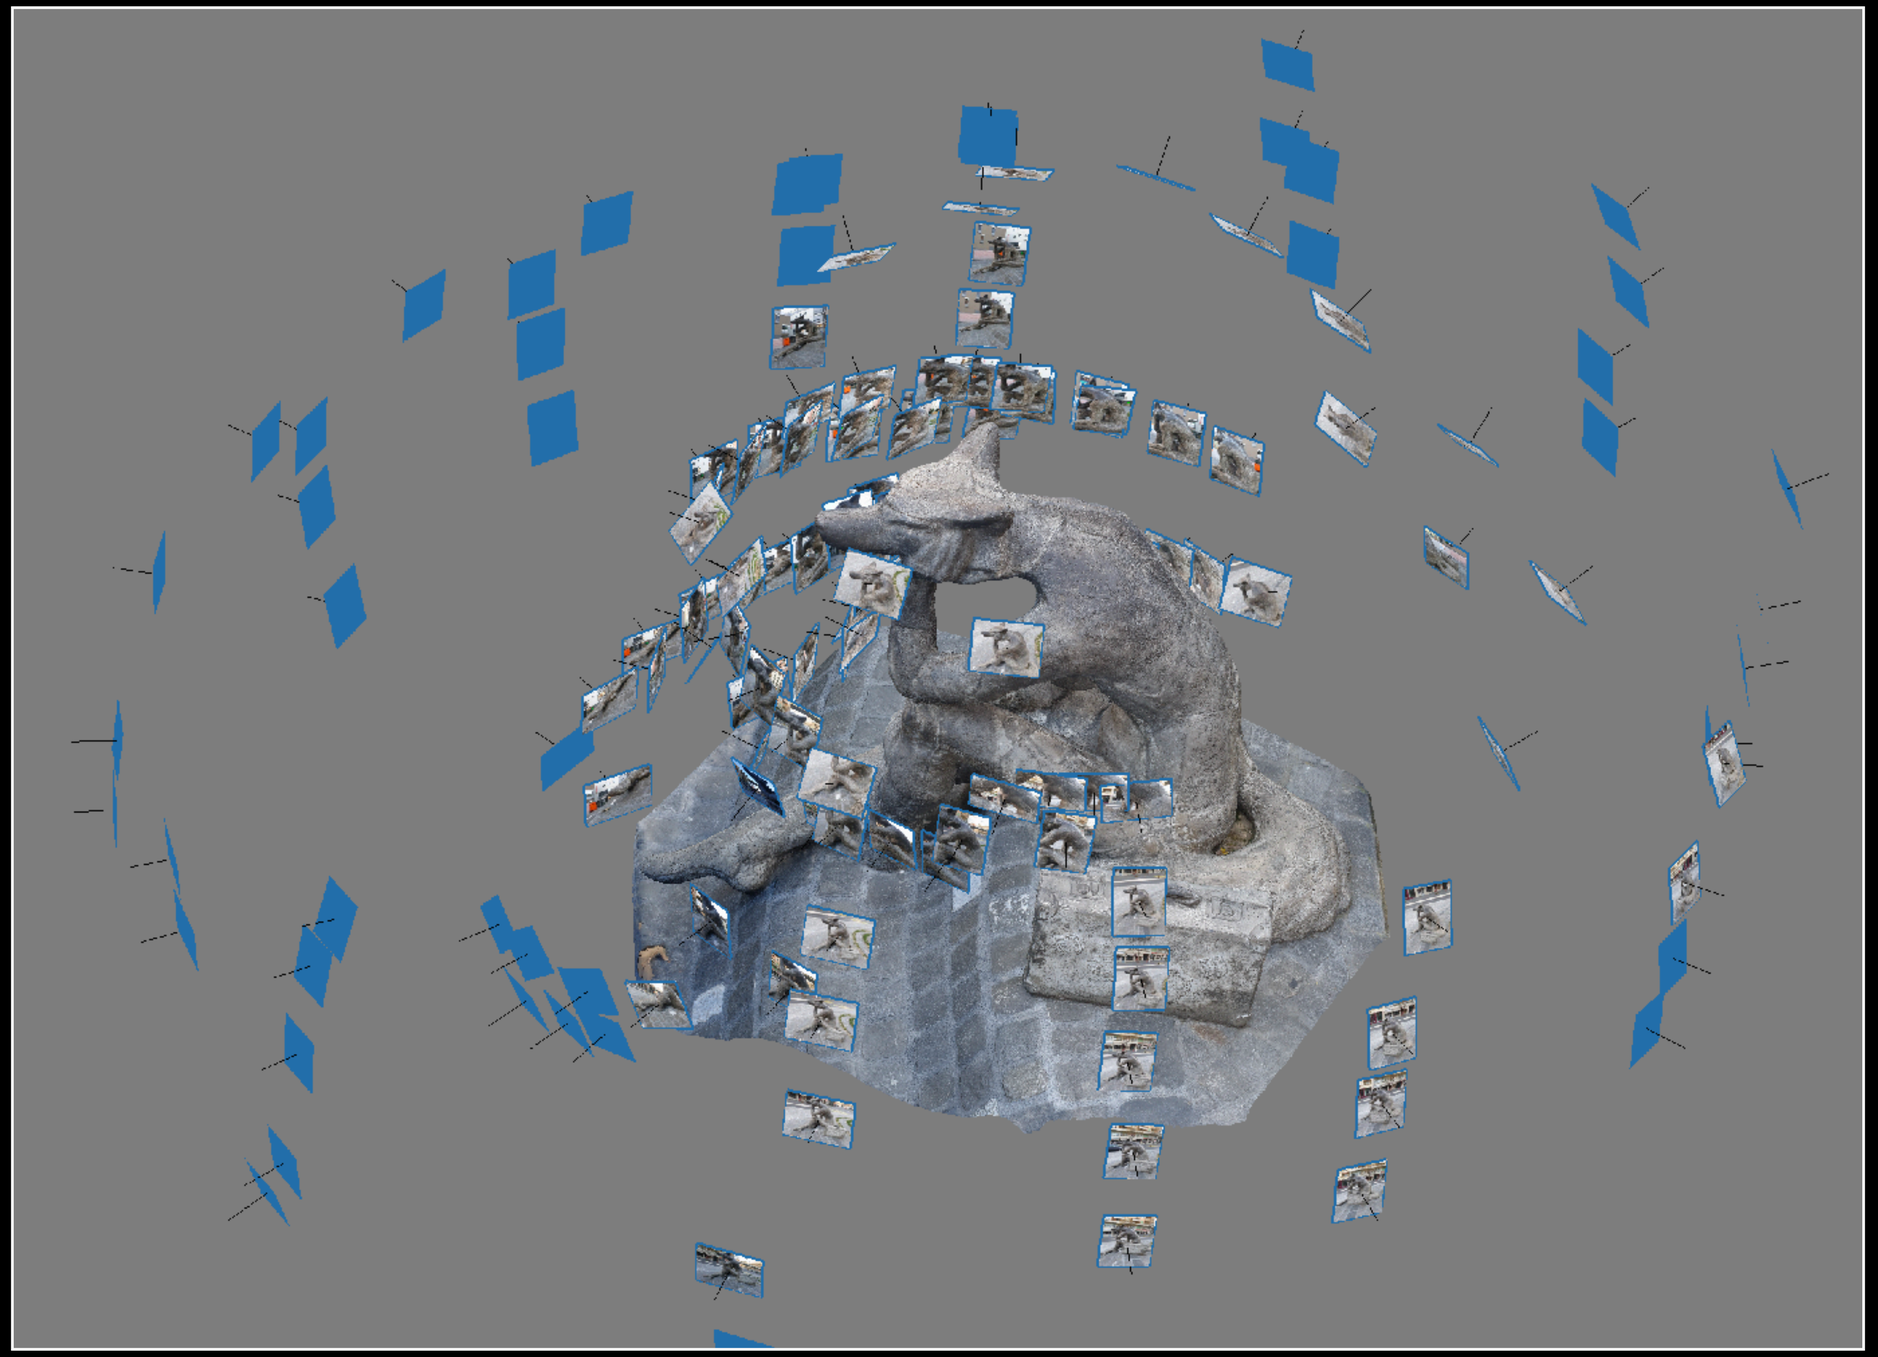

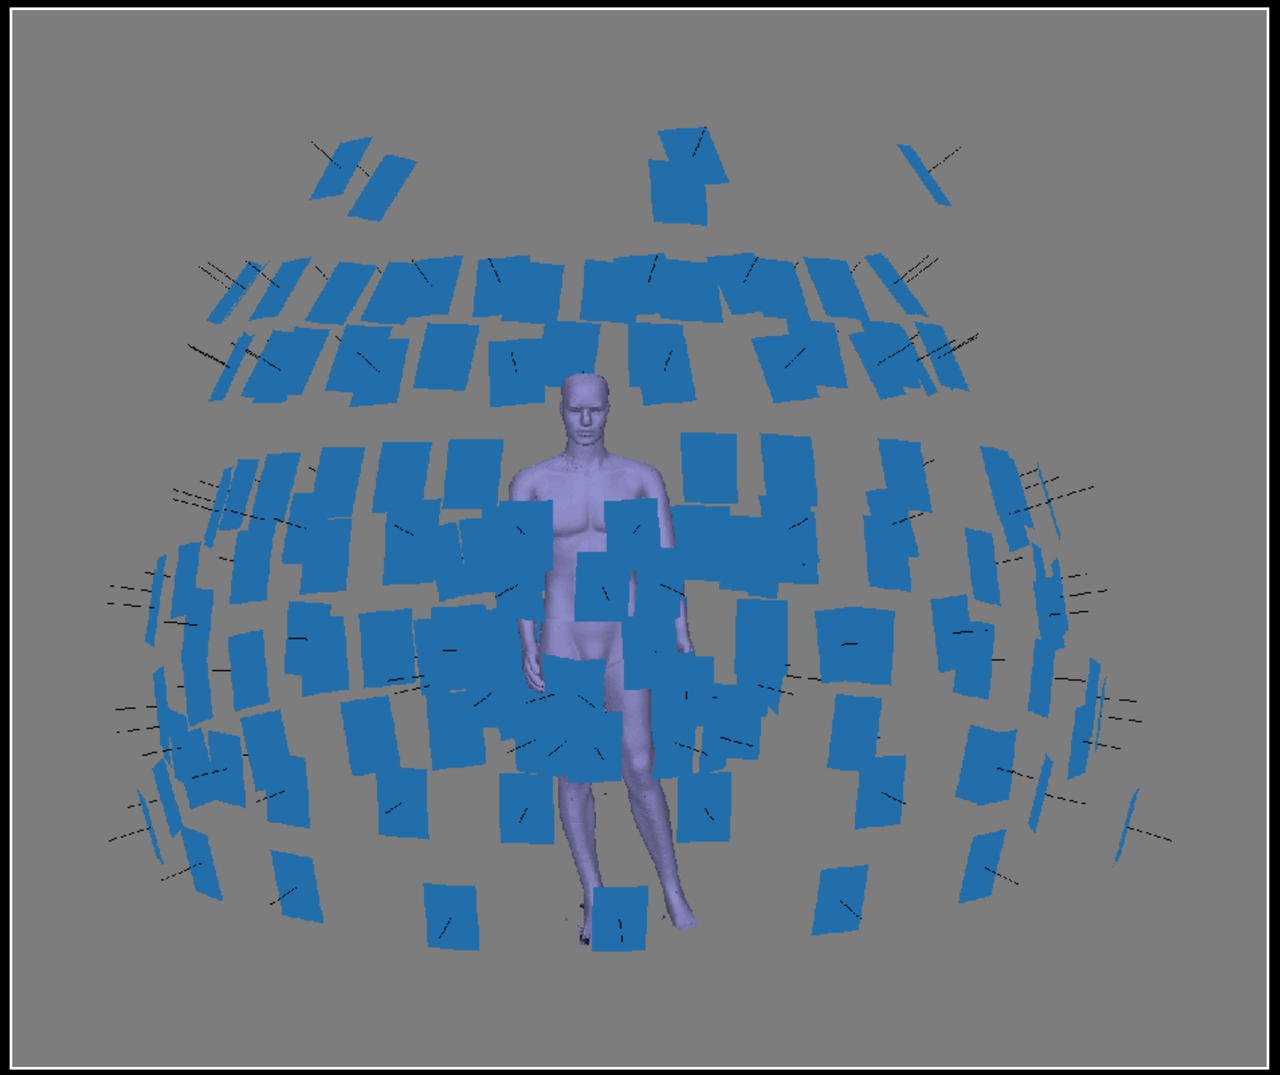

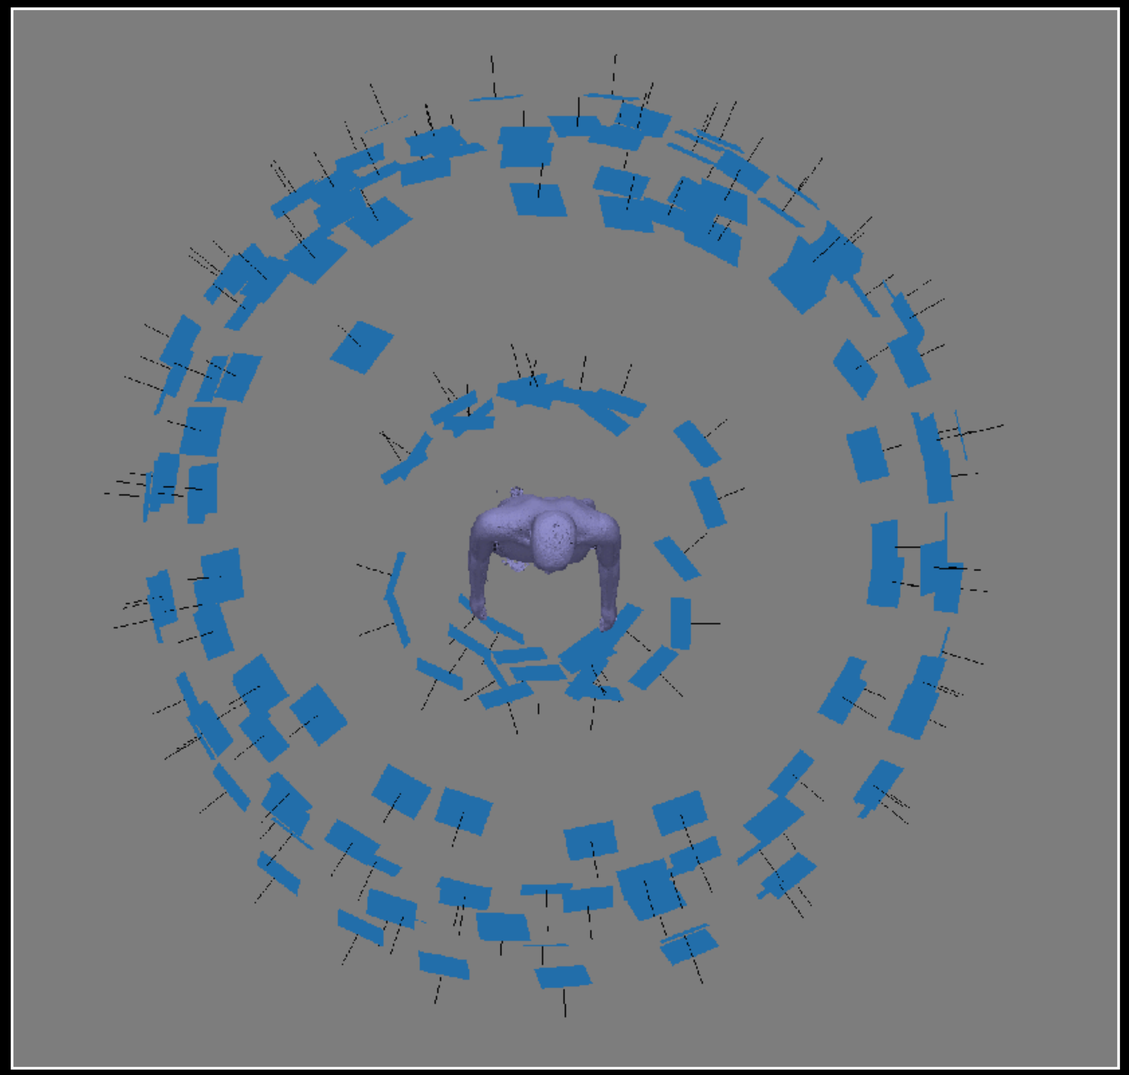

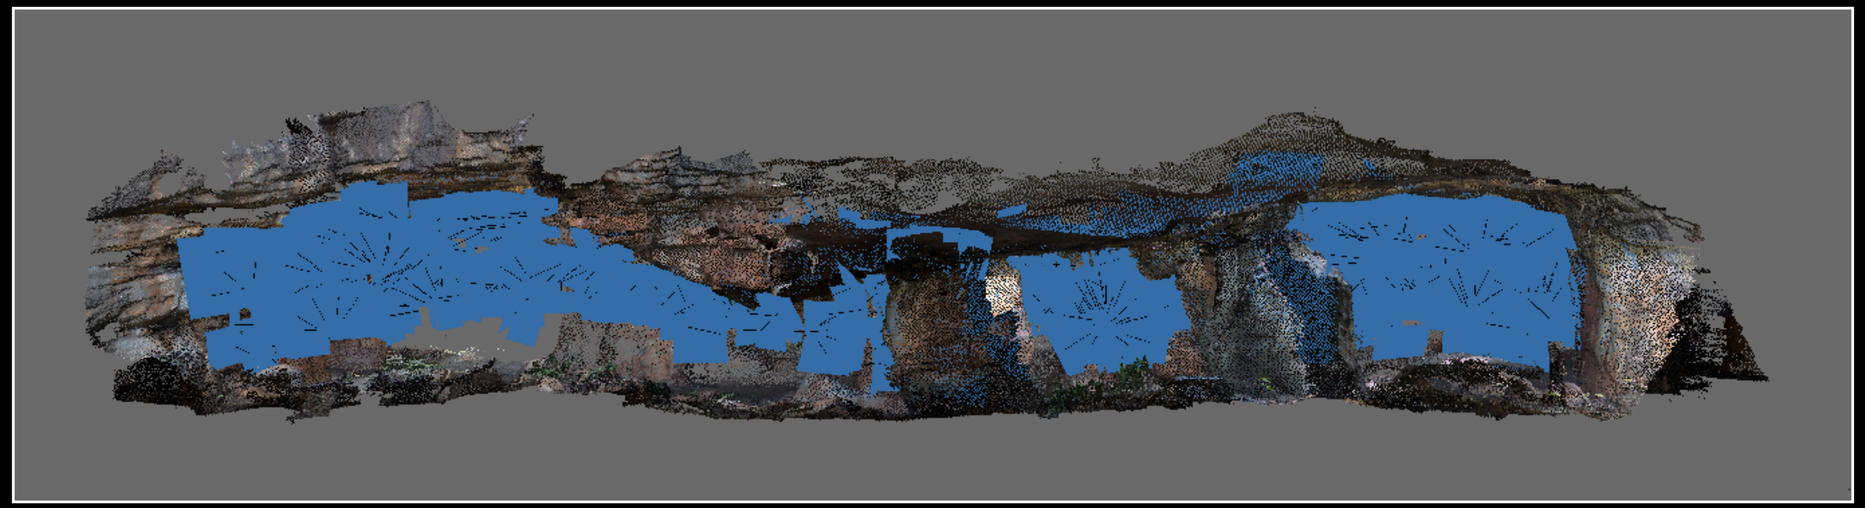

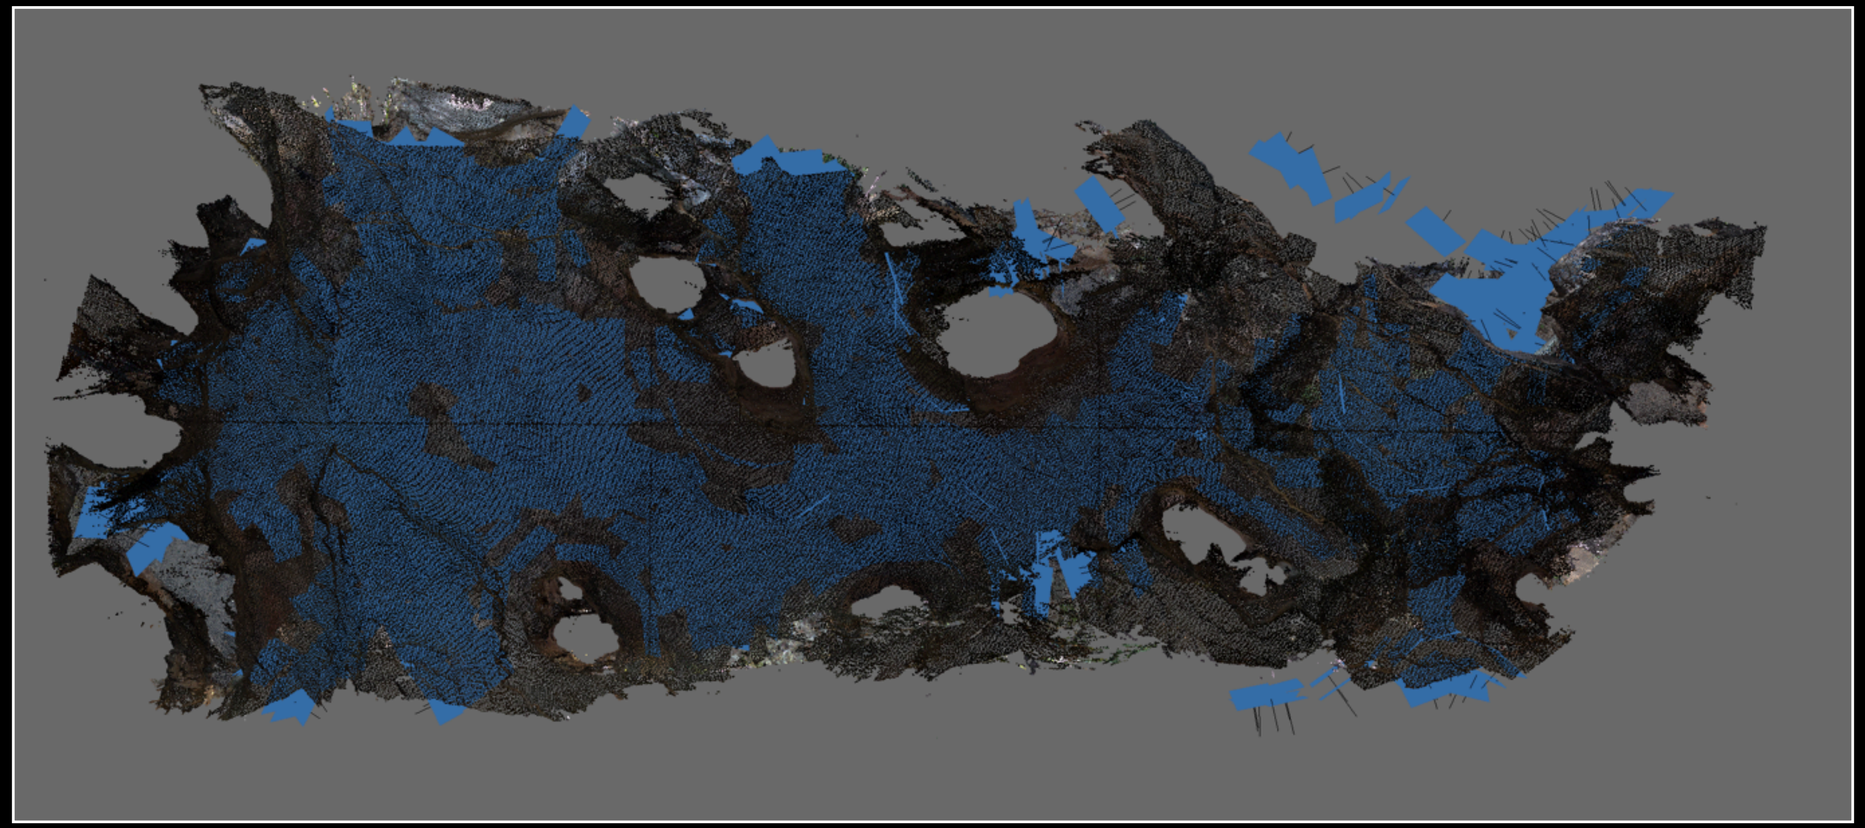

Software workflow

- A number of products in the market
  - PhotoModeller Scanner
  - Remake
  - Apero
  - Colmap (free)
  - Alice Meshroom (free)
  - Opensource pipelines
  - MetaShape, was Photoscan

  …….

- Distinguishing features
  - Degree of automation and human guidedness
  - Implement whole pipeline from photographs to textured mesh
  - Degree of algorithm control
  - Number of photographs that can be handled
  - Computer or cloud based

- We will focus here on perhaps the world leader, MetaShape (Was PhotoScan) from Agisoft.

Typical pipeline

- Feature point detection.
- Bundler for sparce point cloud and camera position and poses.

  Possible human intervention to filter sparce points
- Dense point cloud given the extra information available from camera positions.

  Possible human intervention to clean dense cloud
- Mesh the dense point cloud.

  Possible human intervention to edit mesh
- Preproject photographs from the camera positions onto the mesh.

  Possible human intervention to adjust texture image file

Rear and side walls

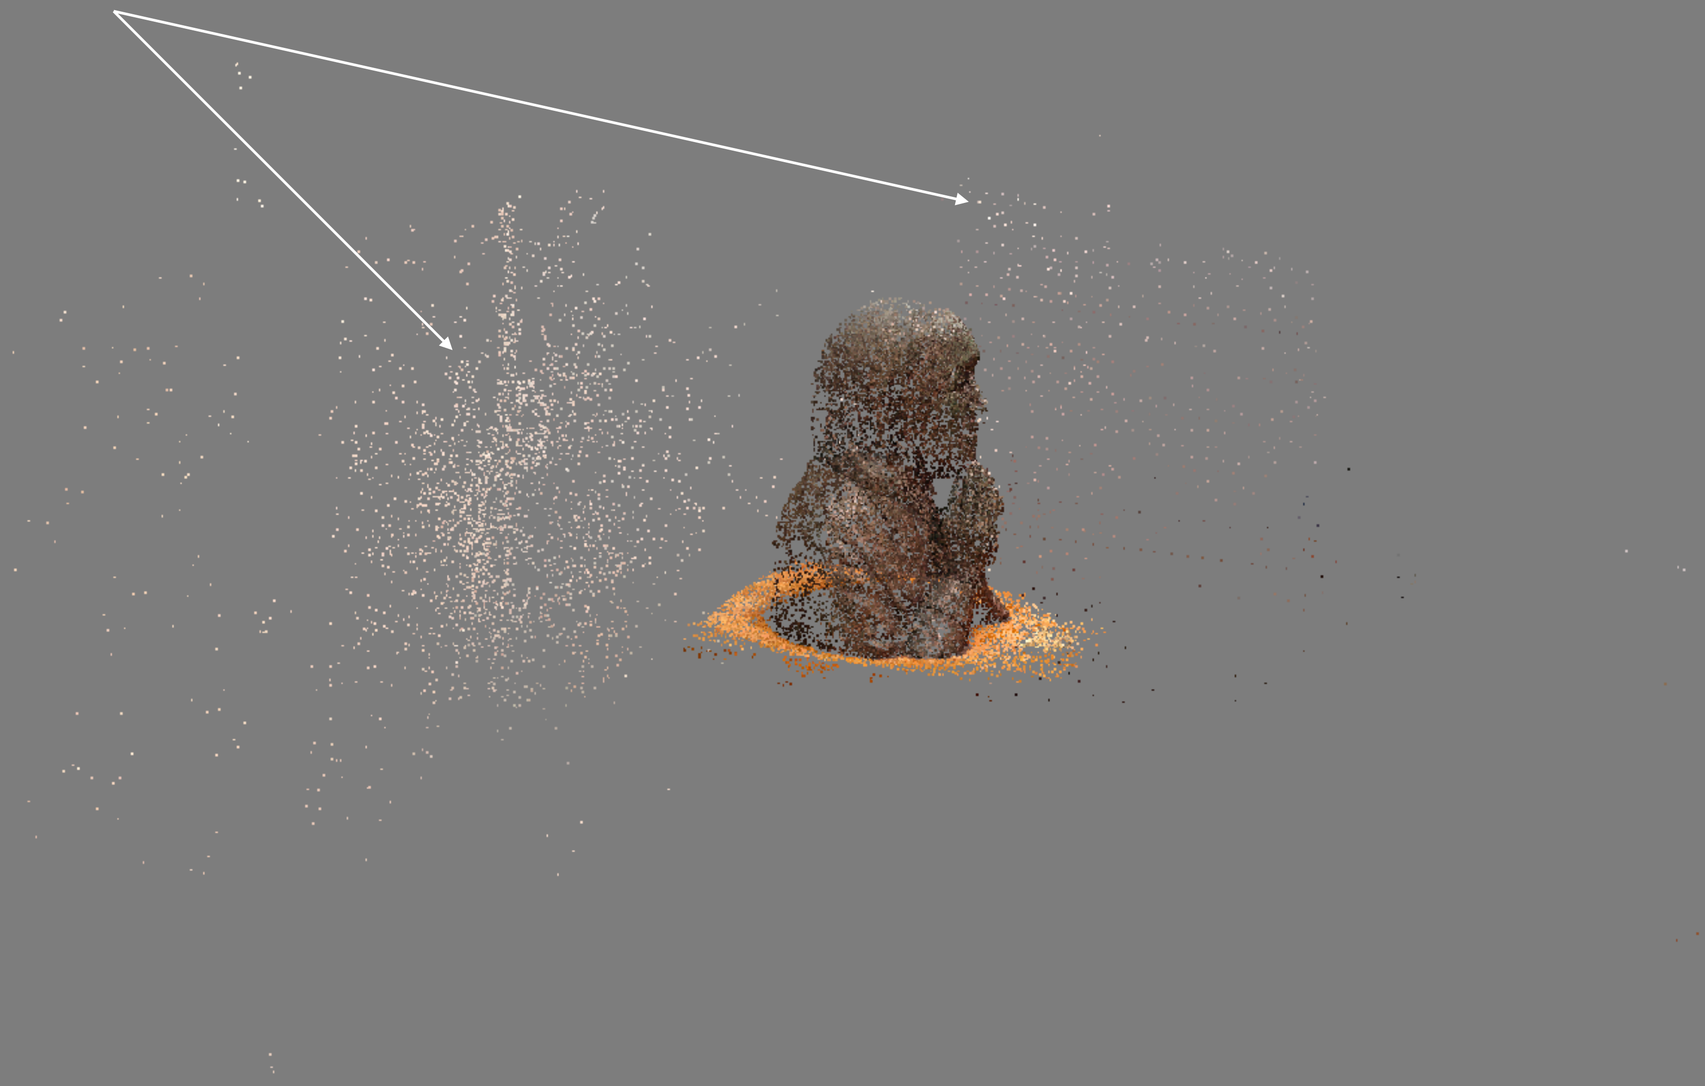

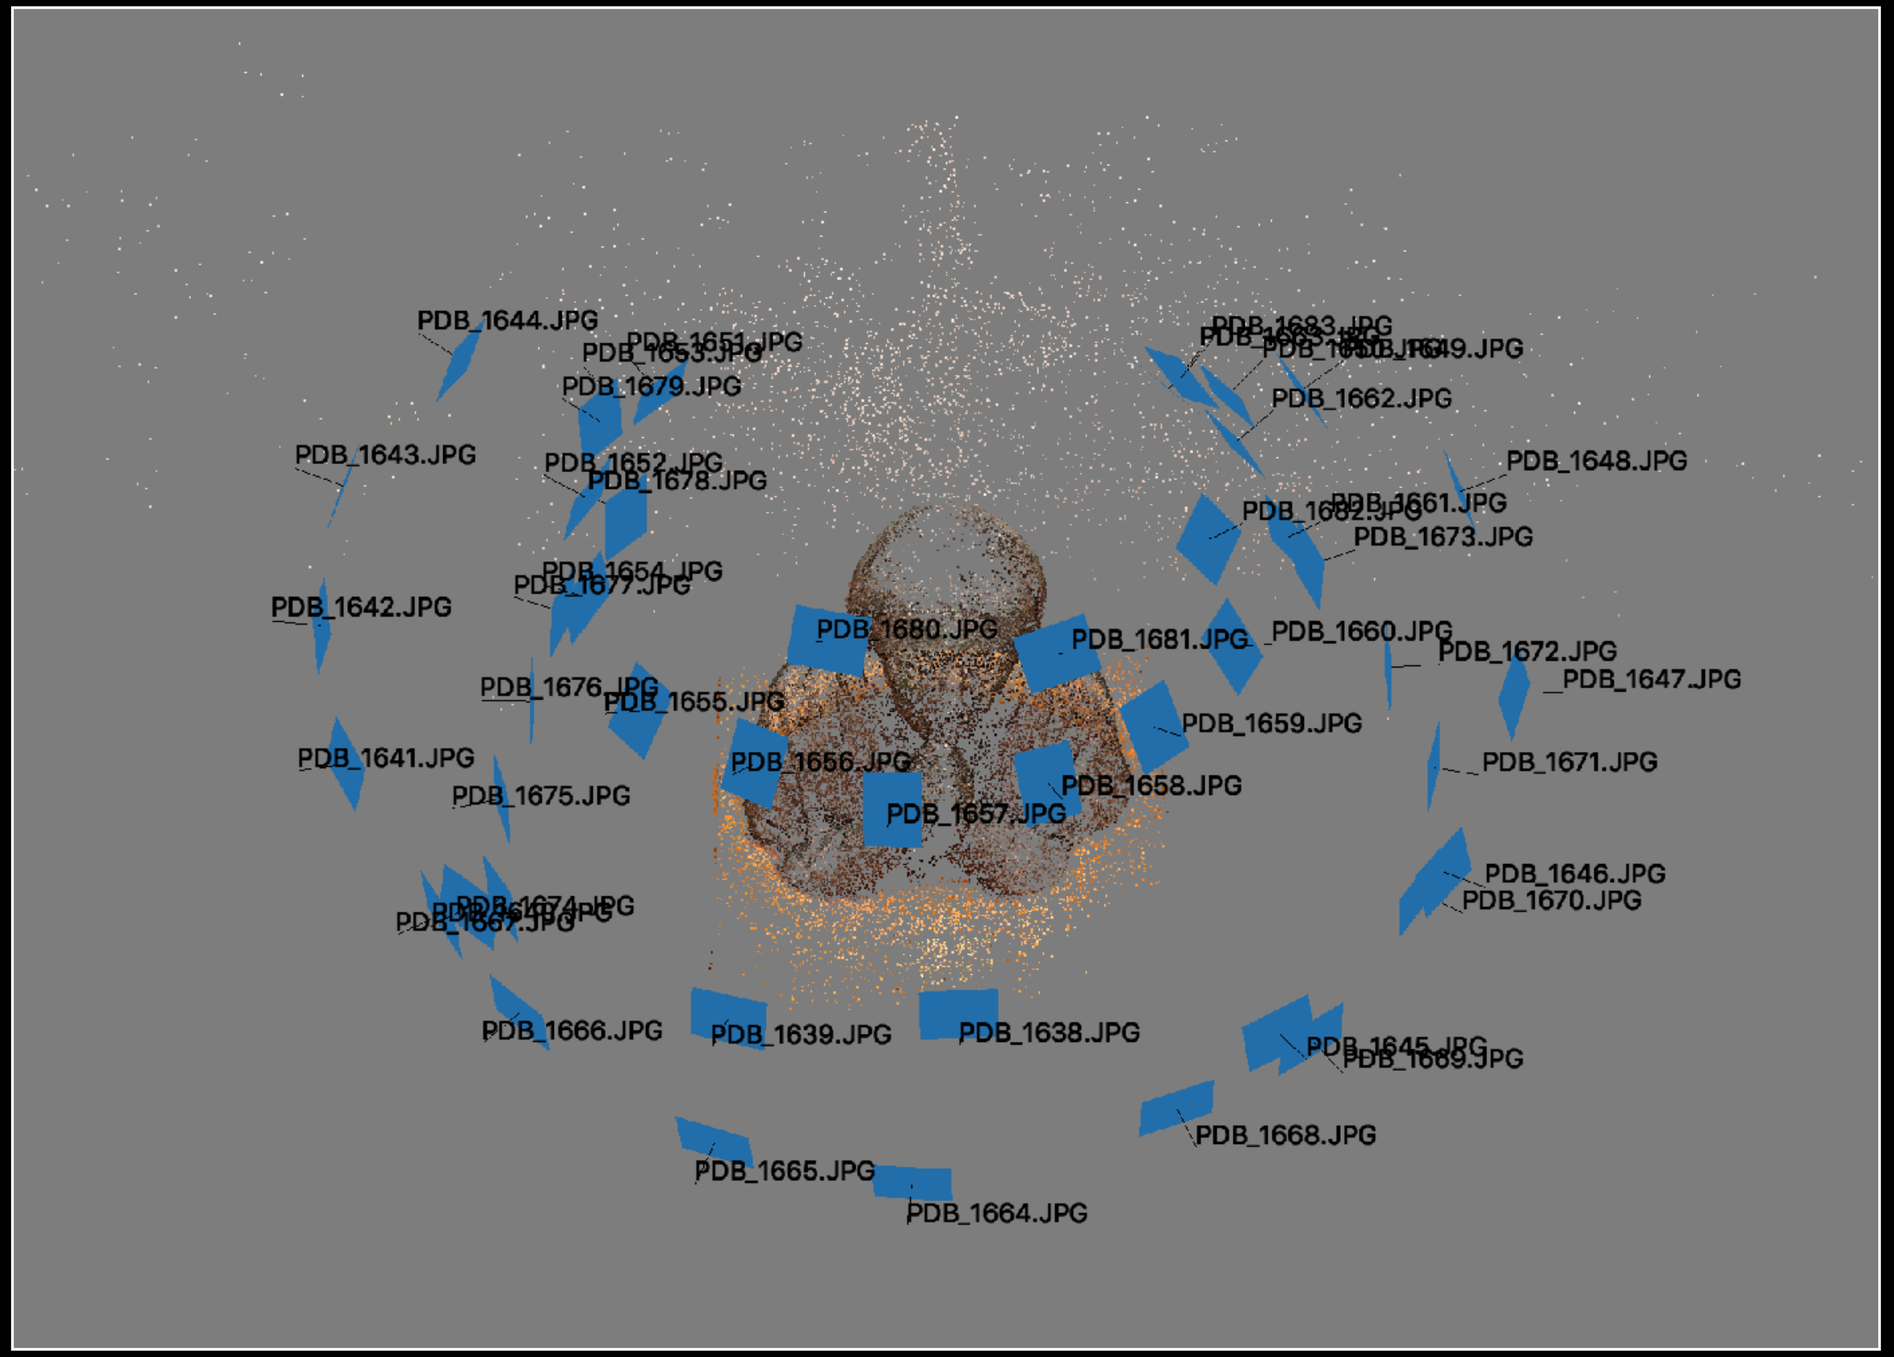
PDB_1644.JPG
PDB_1651.JPG
PDB_1653.JPG
PDB_1679.JPG
PDB_1643.JPG
PDB_1652.JPG
PDB_1678.JPG
PDB_1654.JPG
PDB_1677.JPG
PDB_1642.JPG
PDB_1676.JPG
PDB_1655.JPG
PDB_1641.JPG
PDB_1675.JPG
PDB_1674.JPG
PDB_1667.JPG
PDB_1666.JPG
PDB_1639.JPG
PDB_1638.JPG
PDB_1665.JPG
PDB_1664.JPG
PDB_1668.JPG
PDB_1645.JPG
PDB_1669.JPG
PDB_1646.JPG
PDB_1670.JPG
PDB_1671.JPG
PDB_1647.JPG
PDB_1672.JPG
PDB_1660.JPG
PDB_1659.JPG
PDB_1658.JPG
PDB_1657.JPG
PDB_1656.JPG
PDB_1680.JPG
PDB_1681.JPG
PDB_1662.JPG
PDB_1648.JPG
PDB_1661.JPG
PDB_1673.JPG

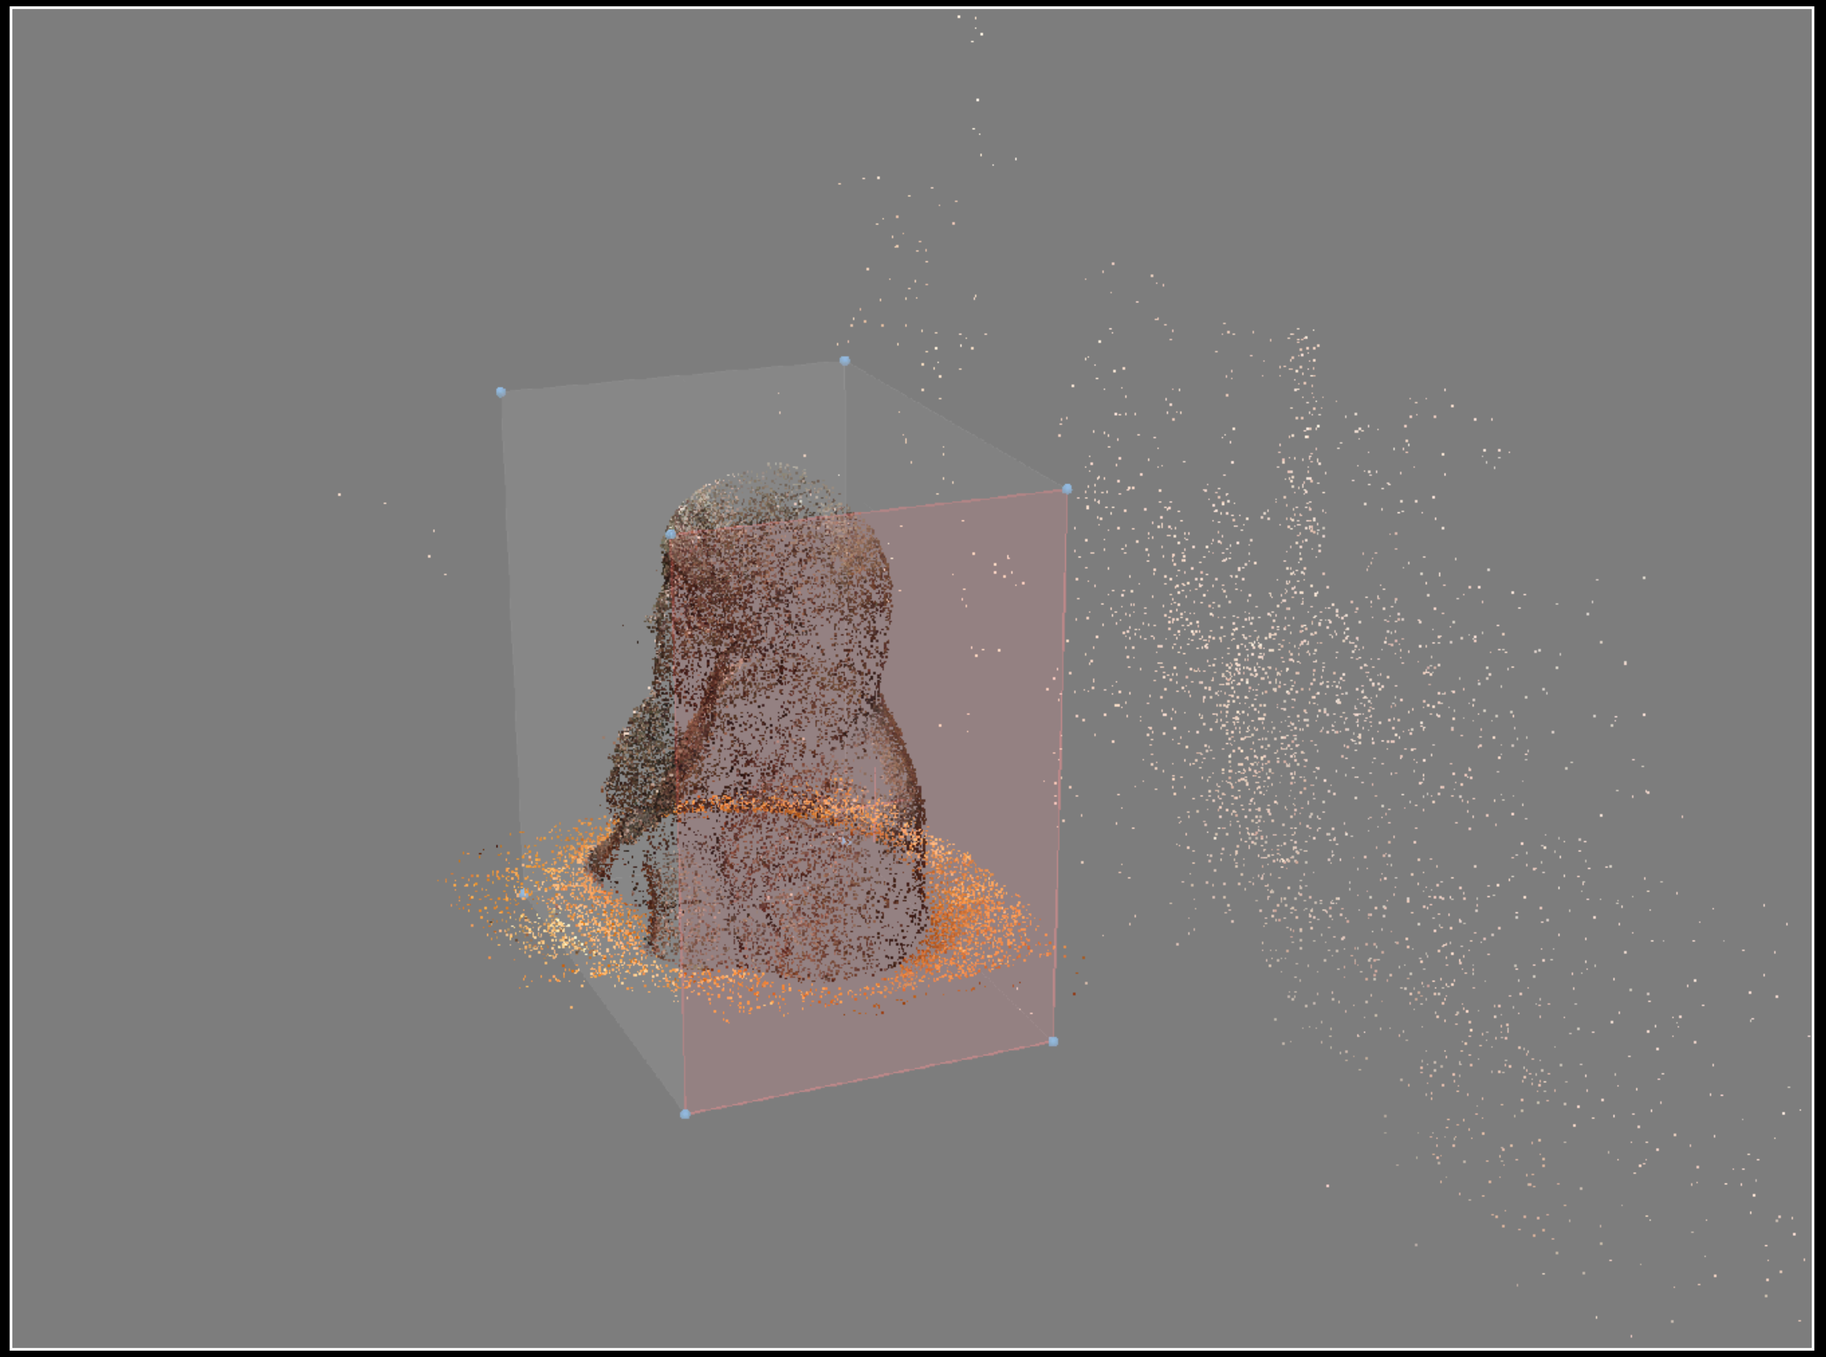

Stray points at the back and top due to lower photographic cover

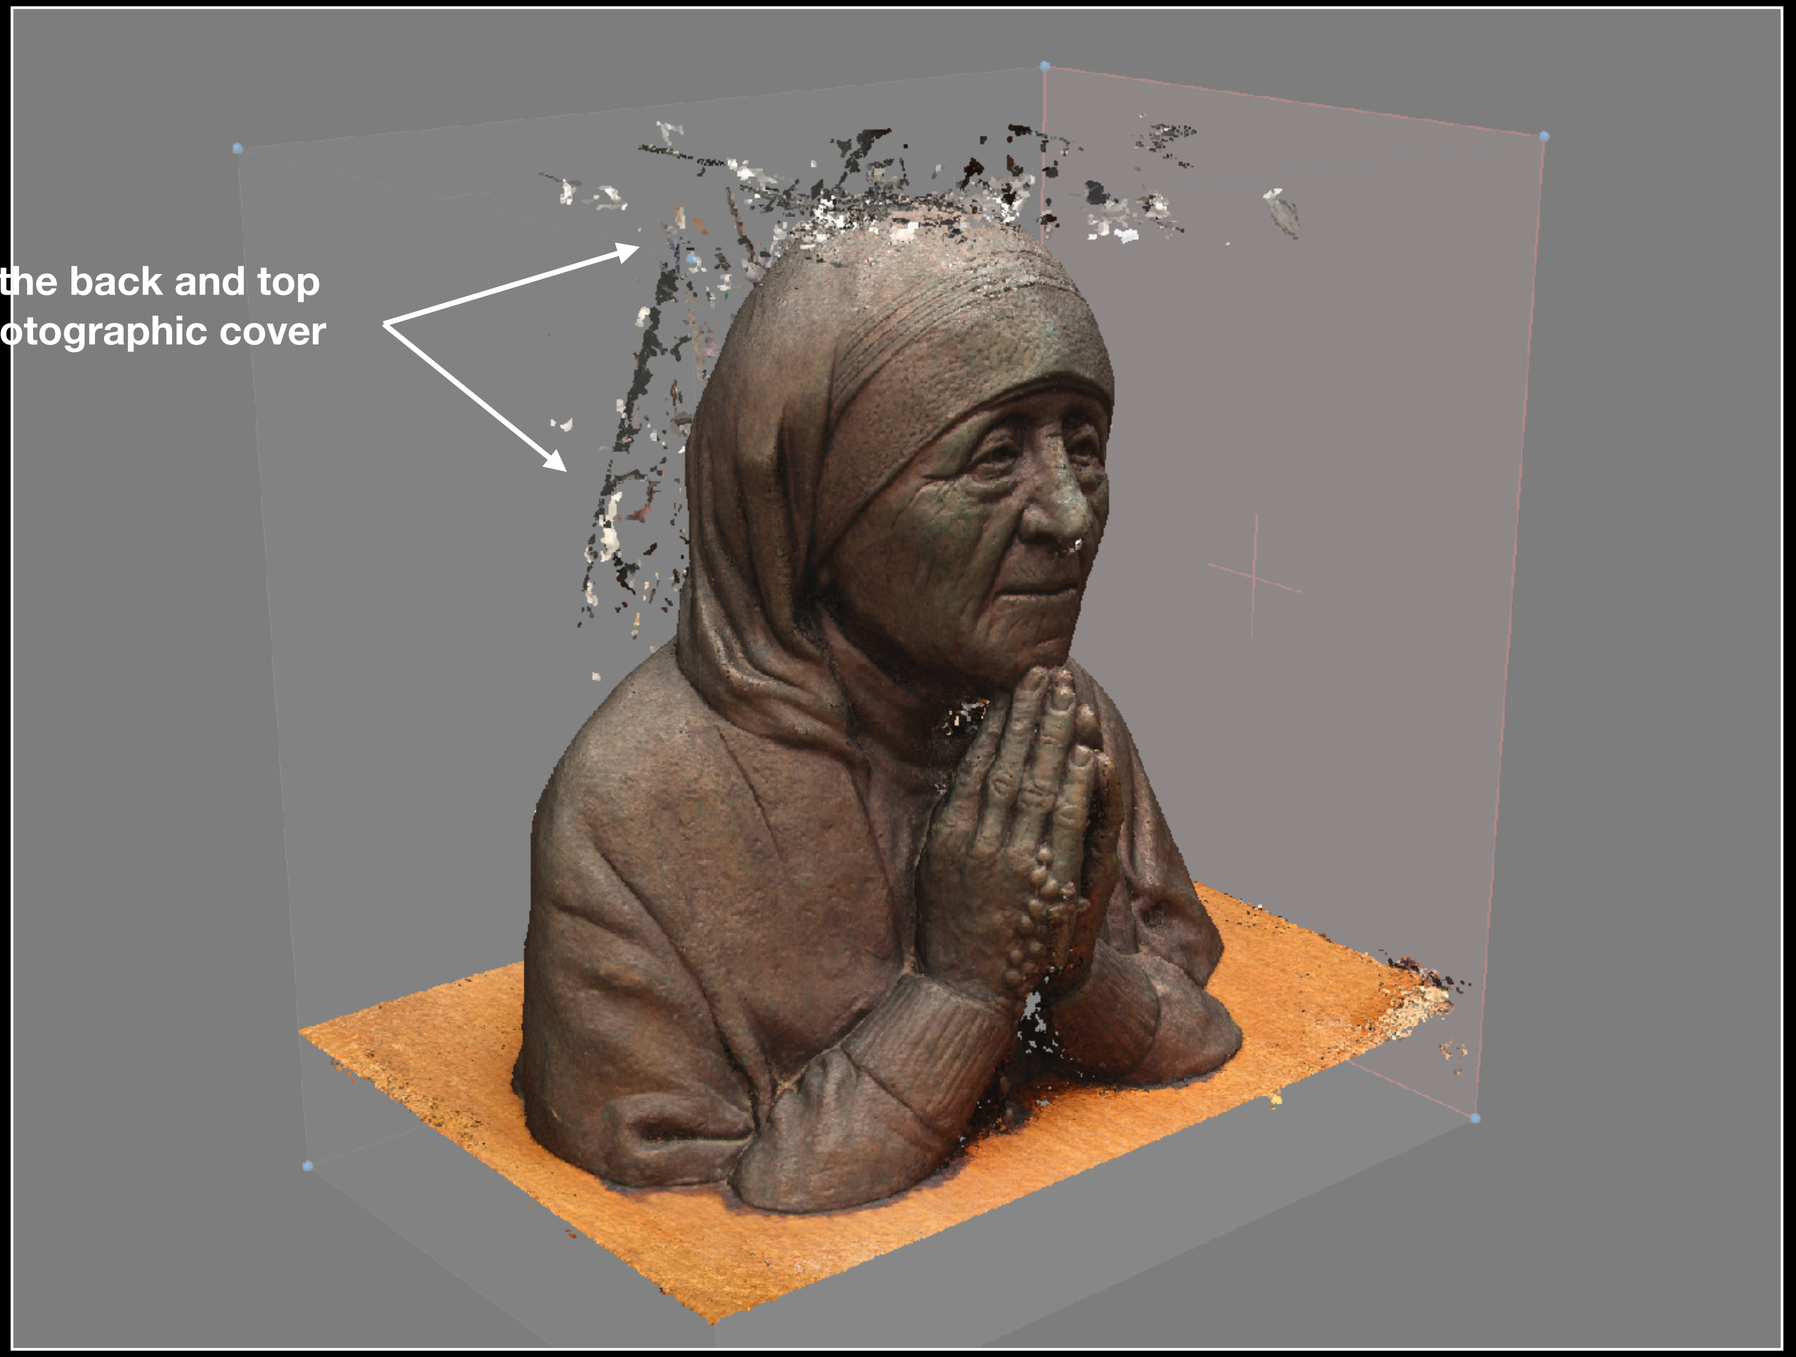

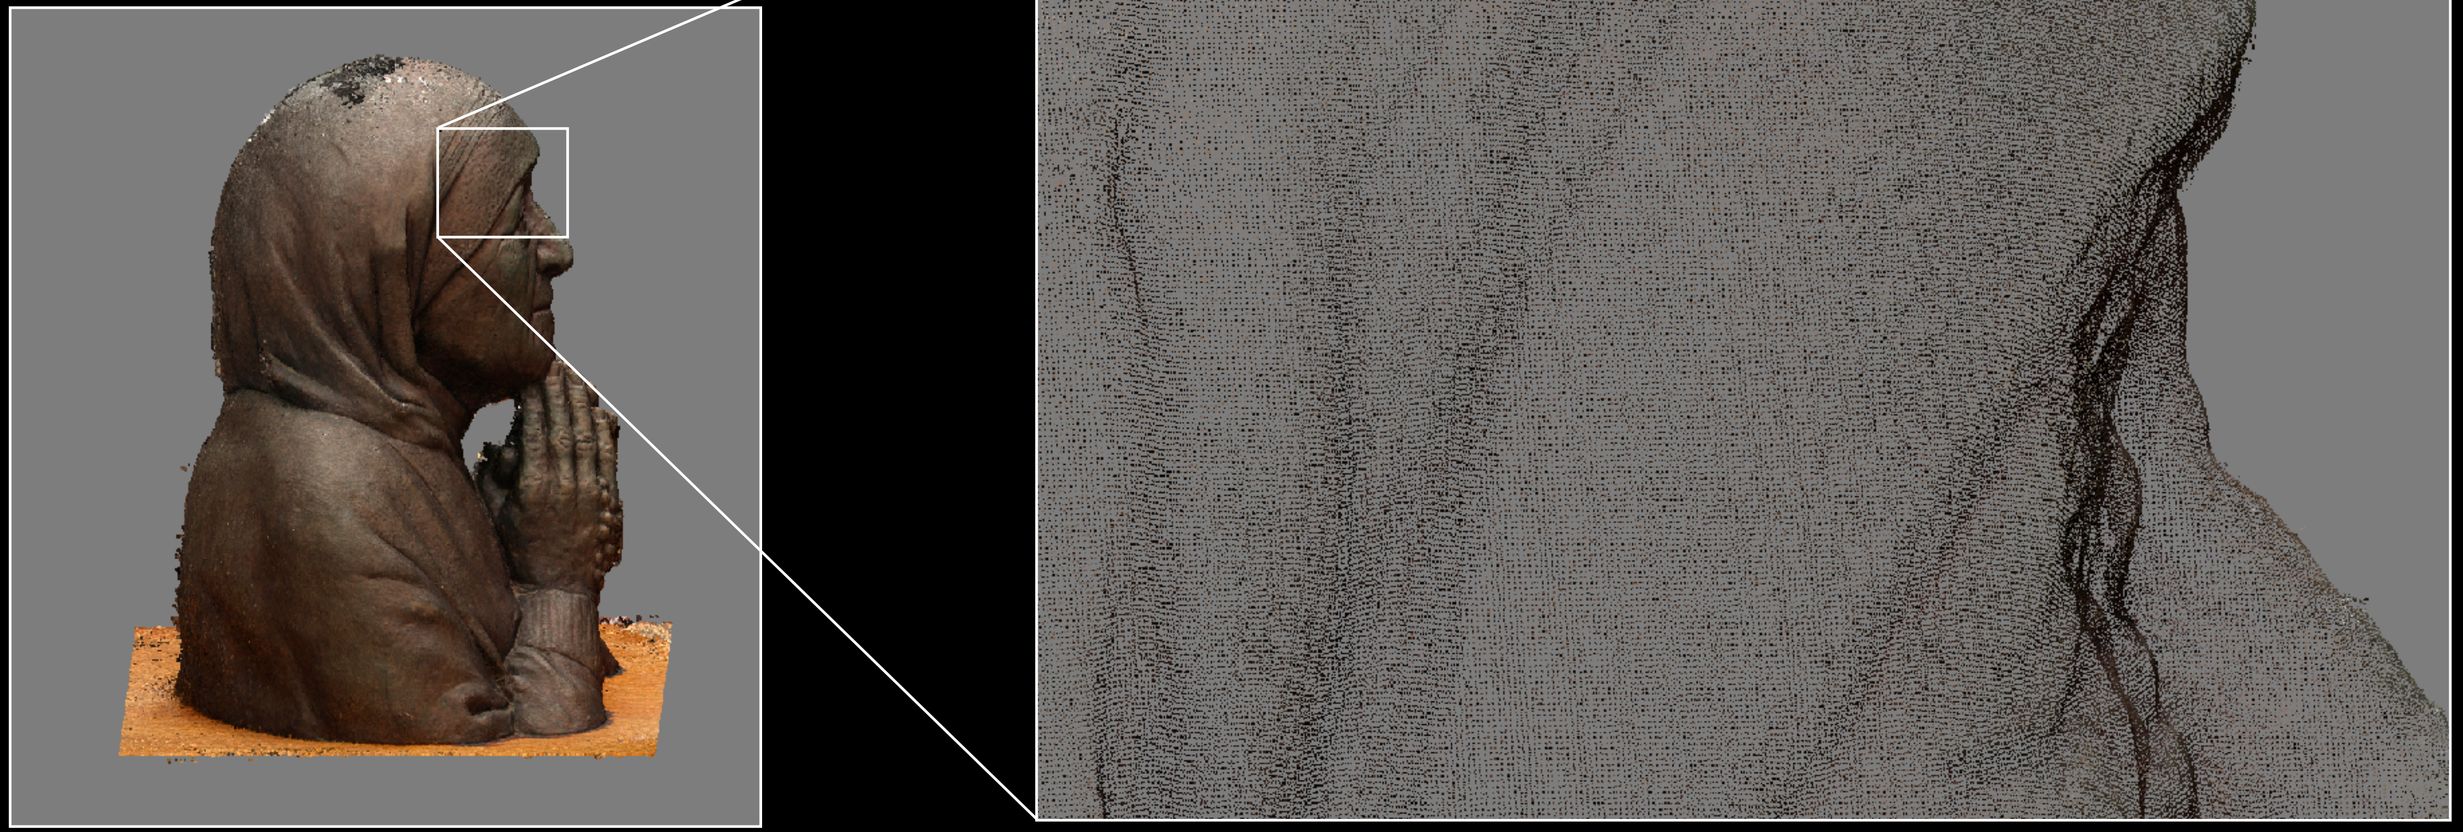

3,500,000 points

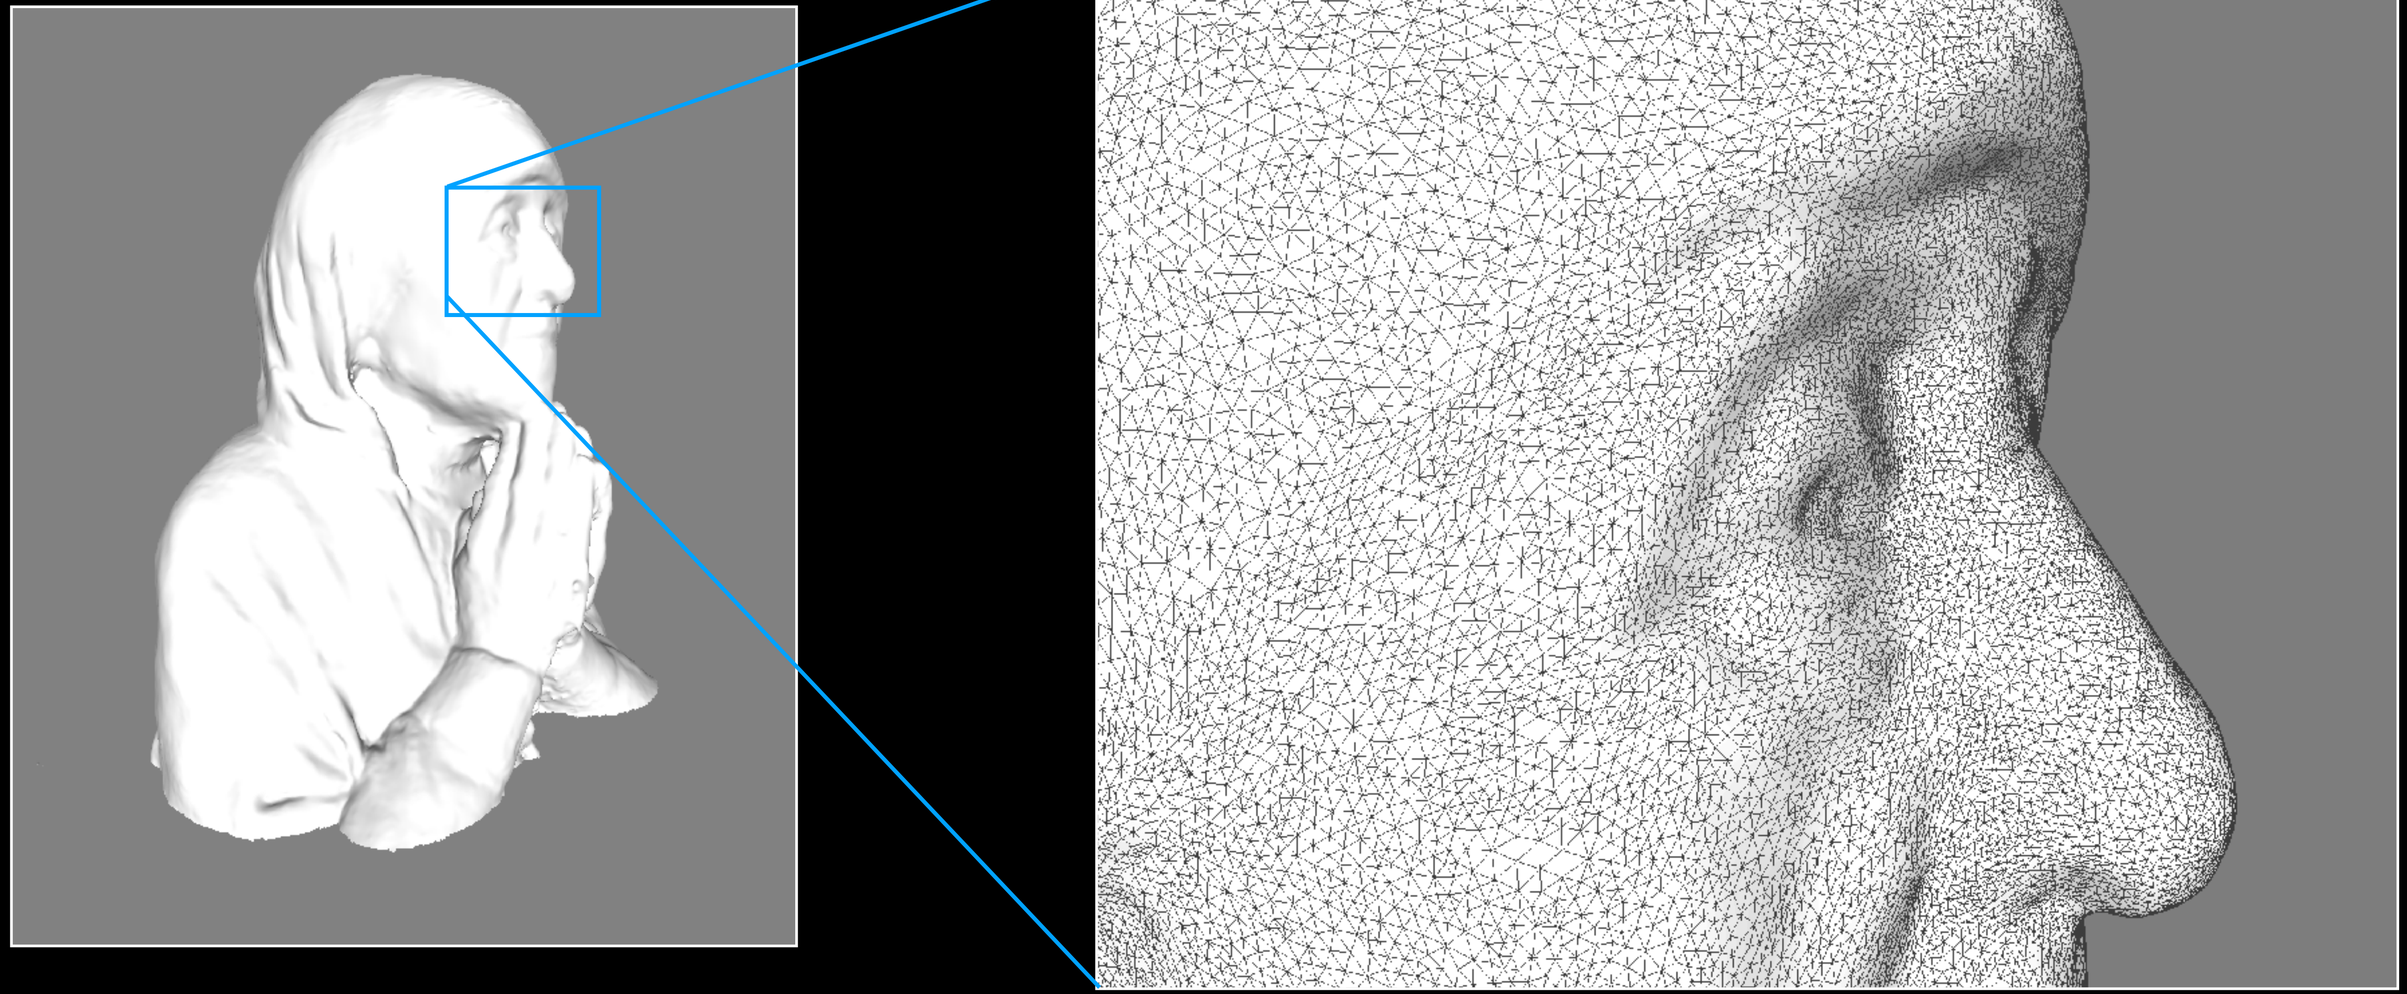


ZBrushCore 4R7 P4DynaMesh_Sphere_128 • Free Mem 15.926GB • Active Mem 457 • Scratch Disk 78 • ZTi QuickSave See-through 0 Menus DefaultZScript
Alpha Brush Color Document Draw Edit File Light Material Preferences Render Stroke Texture Tool Transform Zplugin
Home Page LightBox Draw Move Scale Rotate
Mrgb Rgb M Zadd Zsub
Rgb Intensity 100 Z Intensity 25
Draw Size 64 Focal Shift 0
ActivePoints: 410,456 TotalPoints: 410,456
SelectRect Rect Alpha Off Texture Off MatCap Gray Alternate FillObject
BPR Persp Floor Local L.Sym Xyz Frame Move Zoom3D Rotate Line Fill PolyF Transp Ghost Solo
Tool
Copy Tool Paste Tool Import Export Clone Make PolyMesh3D
smoother. 16 R
smoother smoother Cylinder3D
Subtool Geometry Preview Surface Deformation Masking Polygroups Display Properties
ClayBuildup Standard DamStandard Move TrimDynamic hPolish Pinch Inflate IMM Primitives 3 Paint Tran

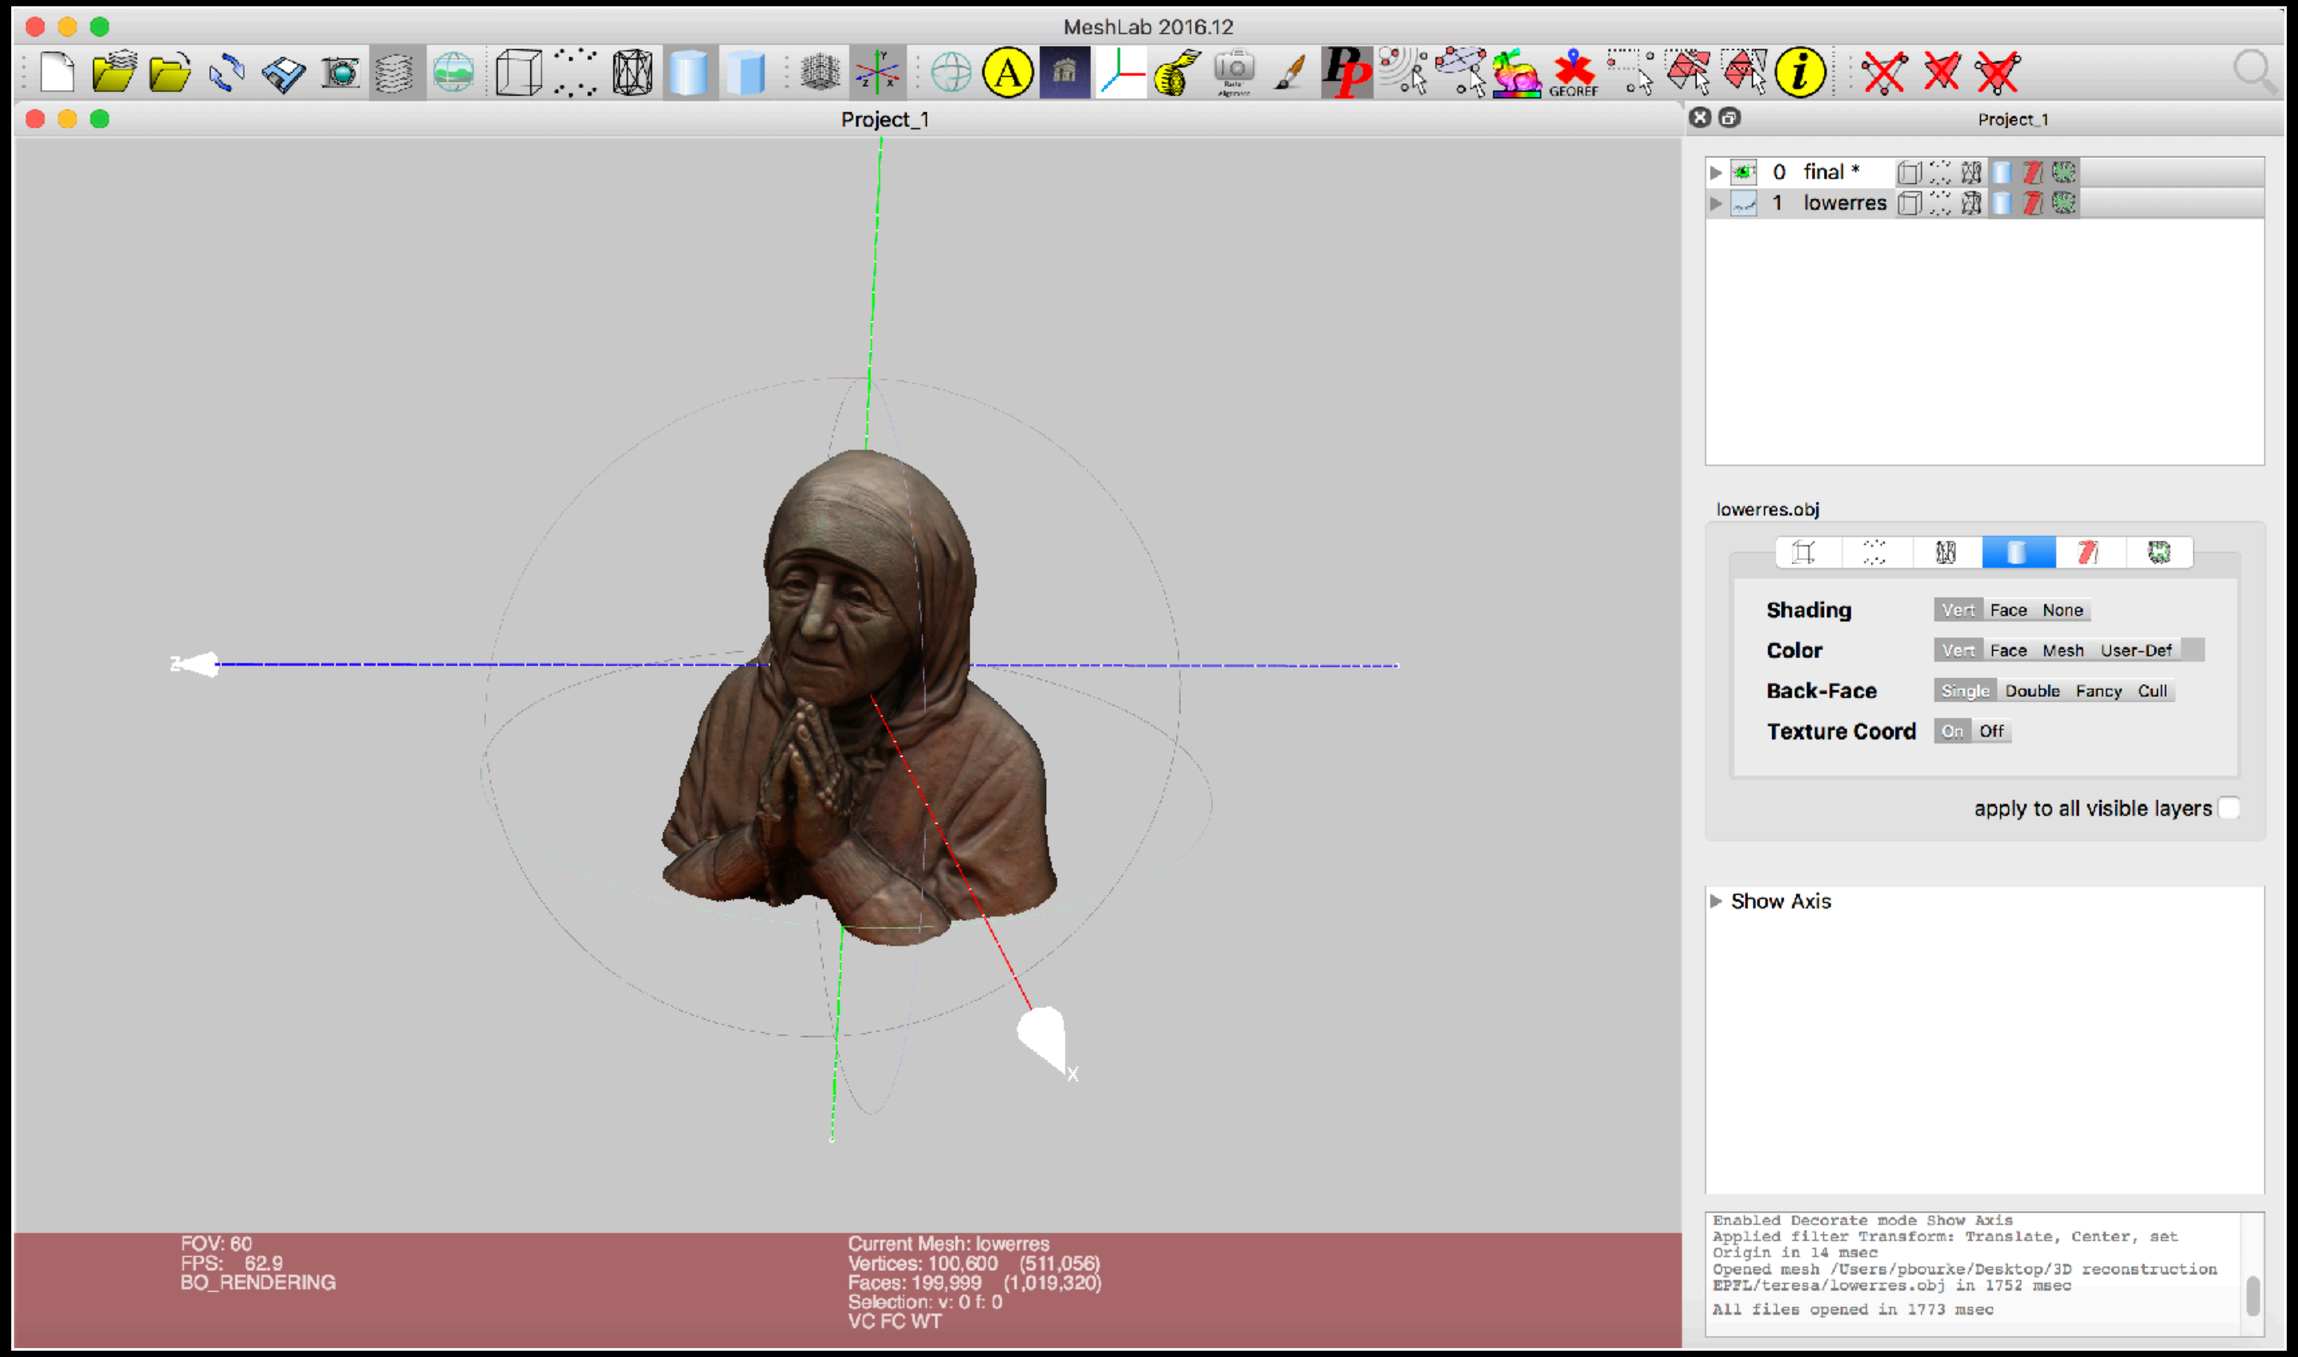
MeshLab 2016.12
Project_1
Project_1
0 final *
1 lowerres
lowerres.obj
Shading
Vert Face None
Color
Vert Face Mesh User-Def
Back-Face
Single Double Fancy Cull
Texture Coord
On Off
apply to all visible layers
Show Axis
Enabled Decorate mode Show Axis
Applied filter Transform: Translate, Center, set Origin in 14 msec
Opened mesh /Users/pbourke/Desktop/3D reconstruction EPFL/teresa/lowerres.obj in 1752 msec
All files opened in 1773 msec
FOV: 60
FPS: 62.9
BO_RENDERING
Current Mesh: lowerres
Vertices: 100,600 (511,056)
Faces: 199,999 (1,019,320)
Selection: v: 0 f: 0
VC FC WT

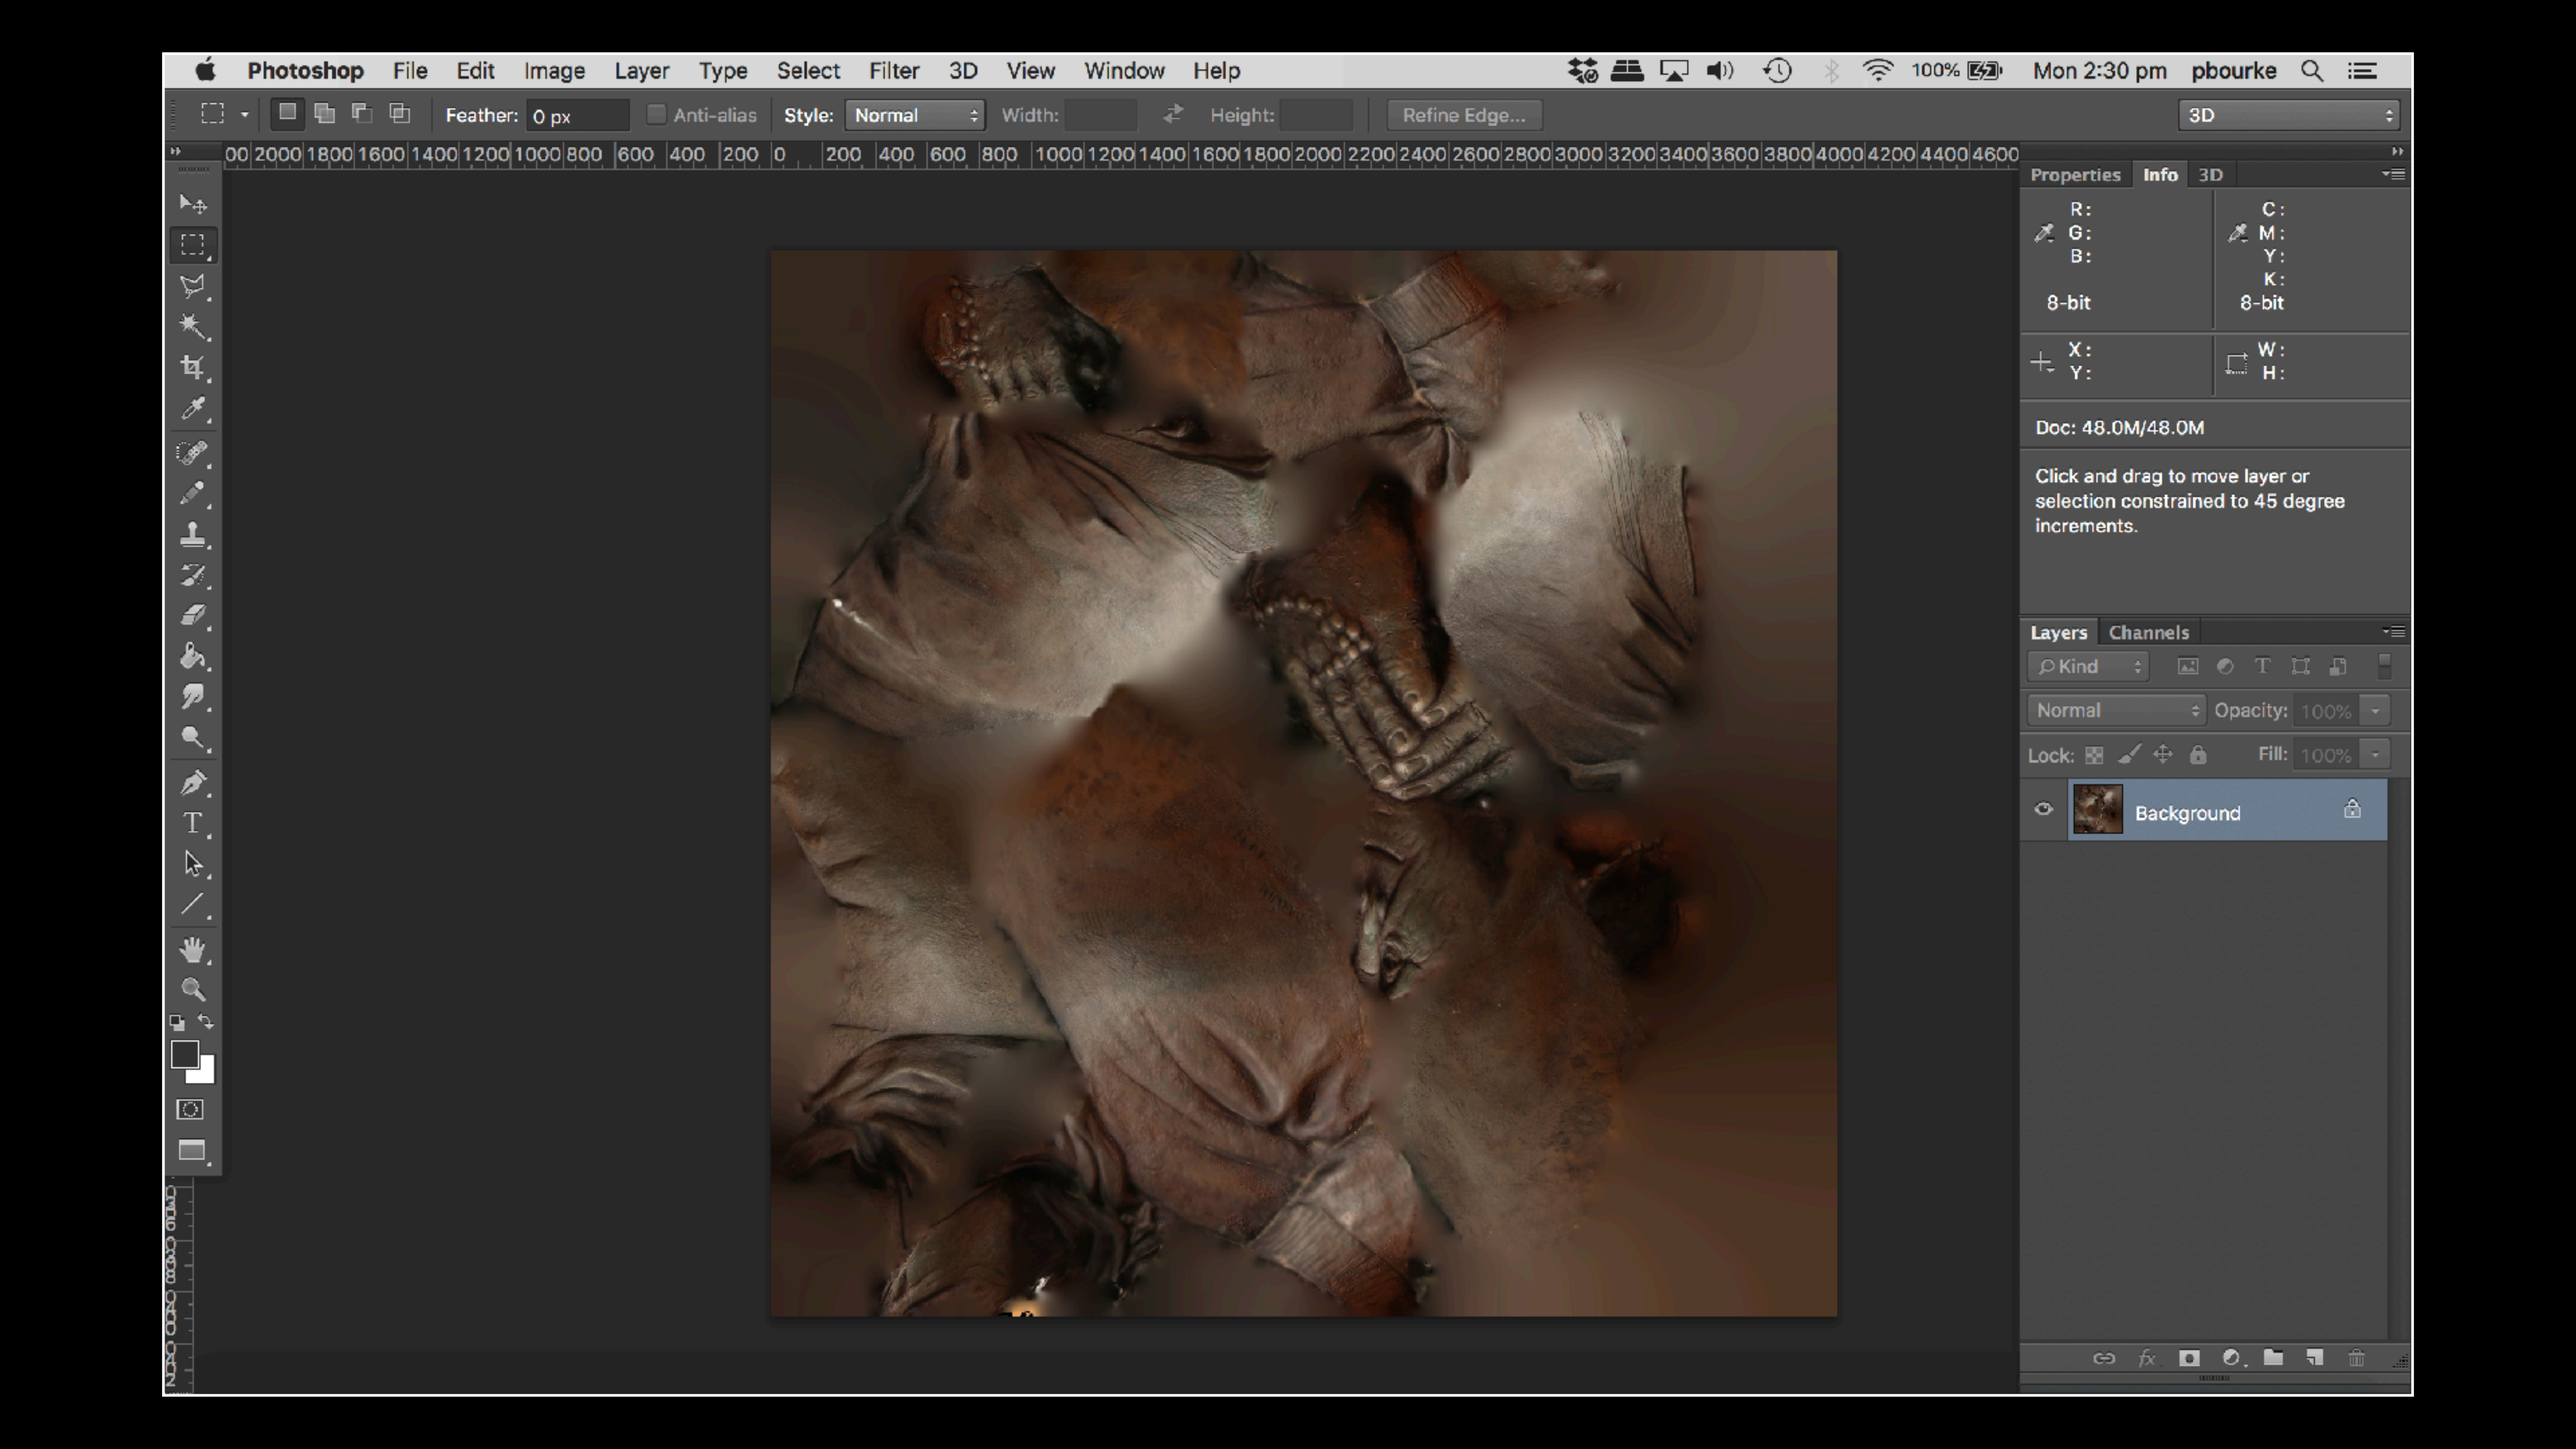
Photoshop File Edit Image Layer Type Select Filter 3D View Window Help
100% Mon 2:30 pm pbourke
Feather: 0 px
Anti-alias
Style: Normal
Width:
Height:
Refine Edge...
3D
Properties Info 3D
R:
G:
B:
8-bit
C:
M:
Y:
K:
8-bit
X:
Y:
W:
H:
Doc: 48.0M/48.0M
Click and drag to move layer or selection constrained to 45 degree increments.
Layers Channels
Kind
Normal
Opacity: 100%
Lock:
Fill: 100%
Background

Sketchfab EXPLORE BUY 3D MODELS FOR BUSINESS Search 3D models UPLOAD

MANAGE THIS MODEL

Status Published

Privacy Public edit

Download No Free Paid

Price

License Standard

Understand licenses

Additional file Upload a file

Textures, other 3D formats, etc.

Accepted formats: .zip, .rar, .7z

Max size: 500MB

PROPERTIES 3D SETTINGS MORE

Mother Teresa

3D Model

pbourke PRO

177 7

Secure payment

Support from sellers

Access to future versions

https://sketchfab.com/models/81d815efb13345f98b8f9f0798bb6ace/edit

Gmail | Paul Bourke | paypal | Things | Transient | Google Maps | Trello | Slack | FB | Other Bookmarks

Currently editing **Mother Teresa**

Help Center | EXIT | SAVE SETTINGS

SAVE VIEW

WORLD SCALE

Scale

2.8128

USE 1 UNIT = 1 METER | RESET SCALE

Tips for world scale

Character is 1.75m tall. A floor tile is 1m².

A smaller world makes your model look larger

A large world makes your model look smaller

FLOOR

Display floor for teleportation

Floor level

-0.06m

Tips for floor level

Position the floor right below large models

Position the floor further down for smaller objects

INITIAL VIEWING POSITION


MOTHER TERESA

Other topics

- Resolution
- Masking
- Scale bars, colour correction
- Lens calibration
- Merging multiple scans
- Turntables, moving camera vs moving object
- Texture editing
- Data formats
- Limitations

Resolution, Accuracy

- The first question many people ask is how accurate is it?

- Not easy to quantify
  1. Not all parts of a model are equally accurate.
  2. Not always possible to measure ground truth to compare to?
  3. Accuracy can depend on material characteristics of the model itself.
  4. How can you be sure best practice and best technology have been used?

- We have tested three methods to determine accuracy
  1. Reconstruct and compare key measures with known object.
  2. Perform ensemble reconstructions from large image sets, compare variation.
  3. Compare with other scanning techniques: laser scanning, CT, structured light.
  4. Visual comparison of zoomed in photographs of real and reconstructed.

- No absolute scale but use one length as reference.
- Model from 60 images.
- Subsequent measurements accurate to maximum 2mm, average 1mm.

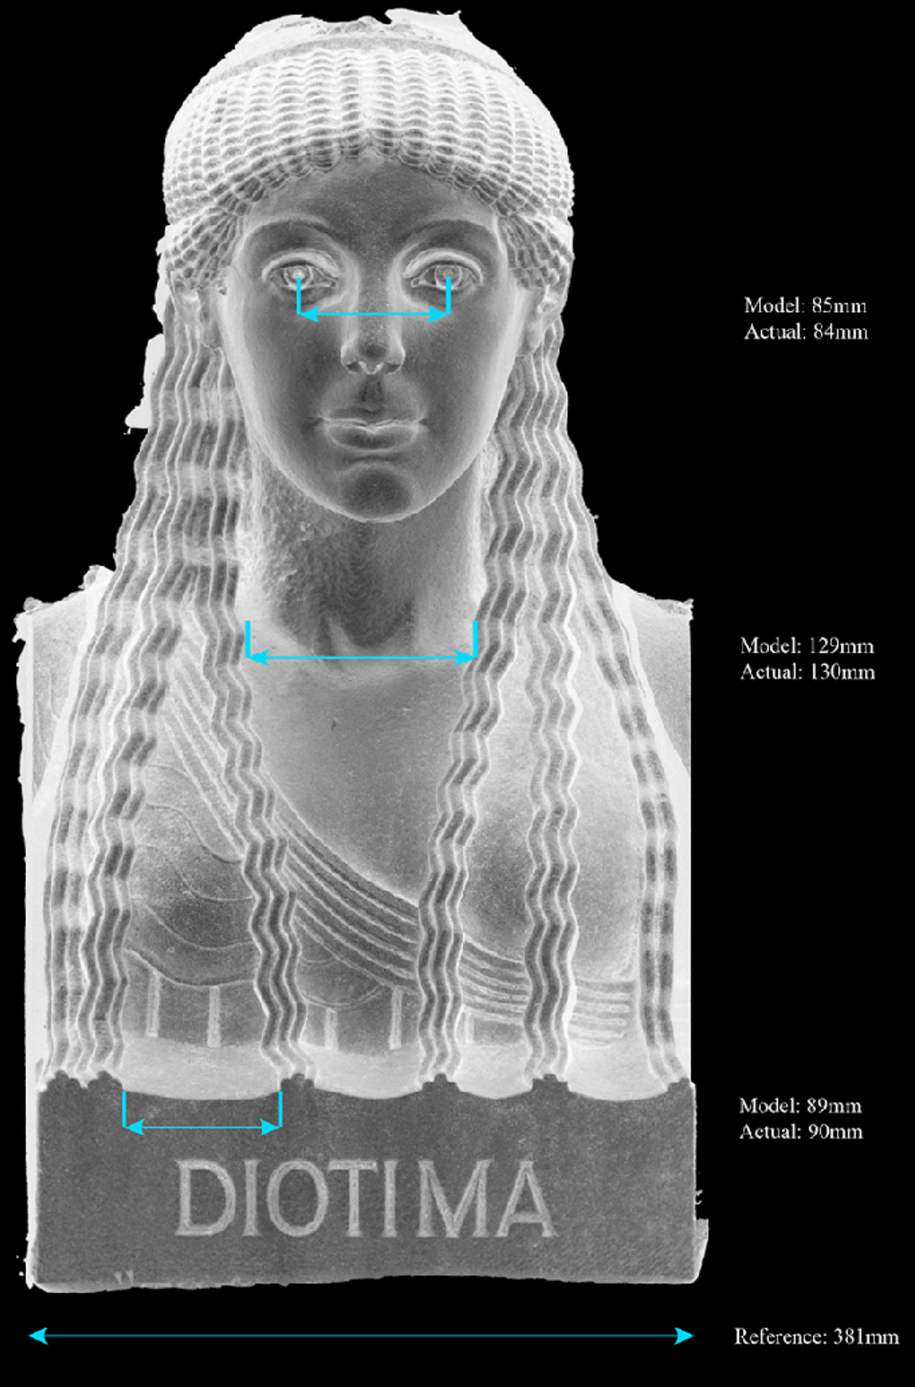

Original photograph

Reconstructed model

Shaded to
emphasise
surface variation

Original photograph

Reconstructed model

Shaded to
emphasise
surface variation

Original

Mesh and texture resolution

- What geometric (mesh) resolution is required to capture the (important) features.

- Texture resolution is similar to the more familiar image resolution, what texture image size is required to capture the visual detail at the intended zoom level.

- Requirements vary depending on the end application.
  - Realtime environments prefer low geometric complexity and high texture detail
  - Analysis generally requires high geometric detail
  - For a digital record one typically wants high geometric and texture detail

	Geometric resolution	Texture resolution
Gaming	Low	High
Analysis	High	Don’t care
Education	Medium	High
Archive/heritage	High	High
Online	Low/Average	Low/average

1,000,000 triangles

100,000 triangles

10,000 triangles

Masking

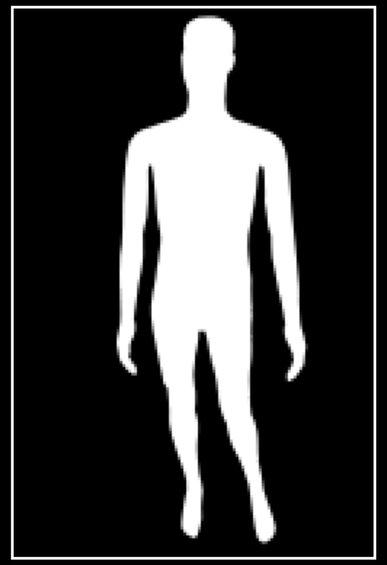

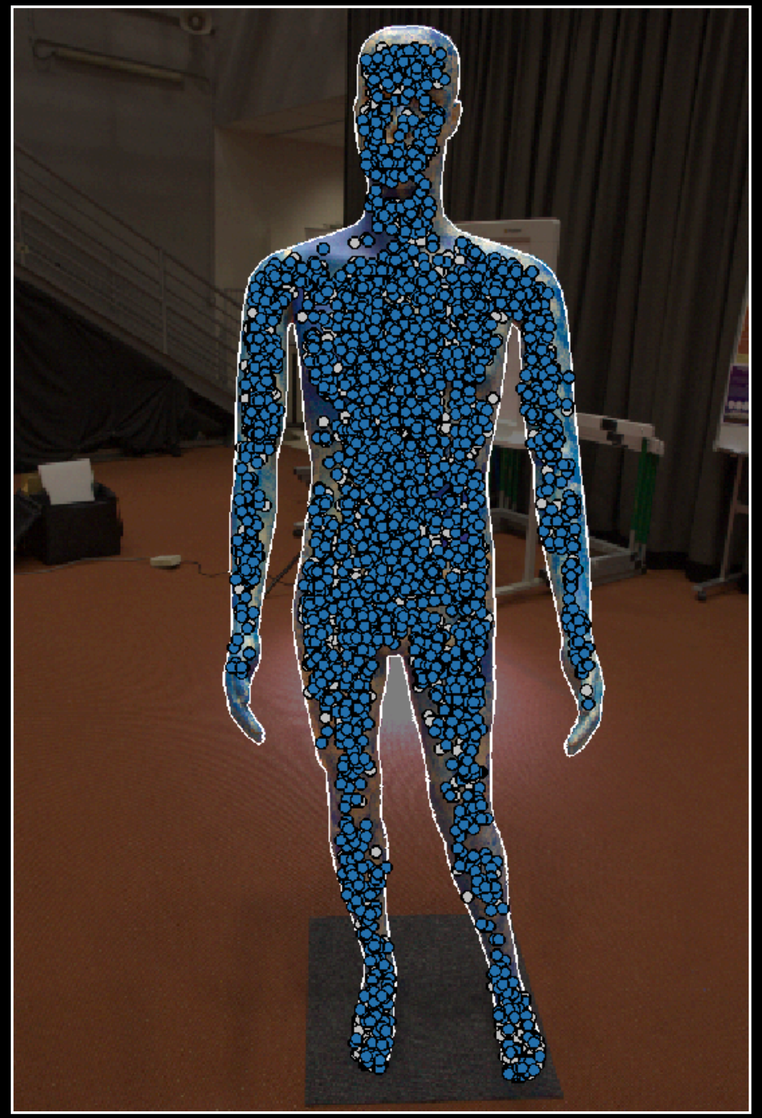

Scale bars

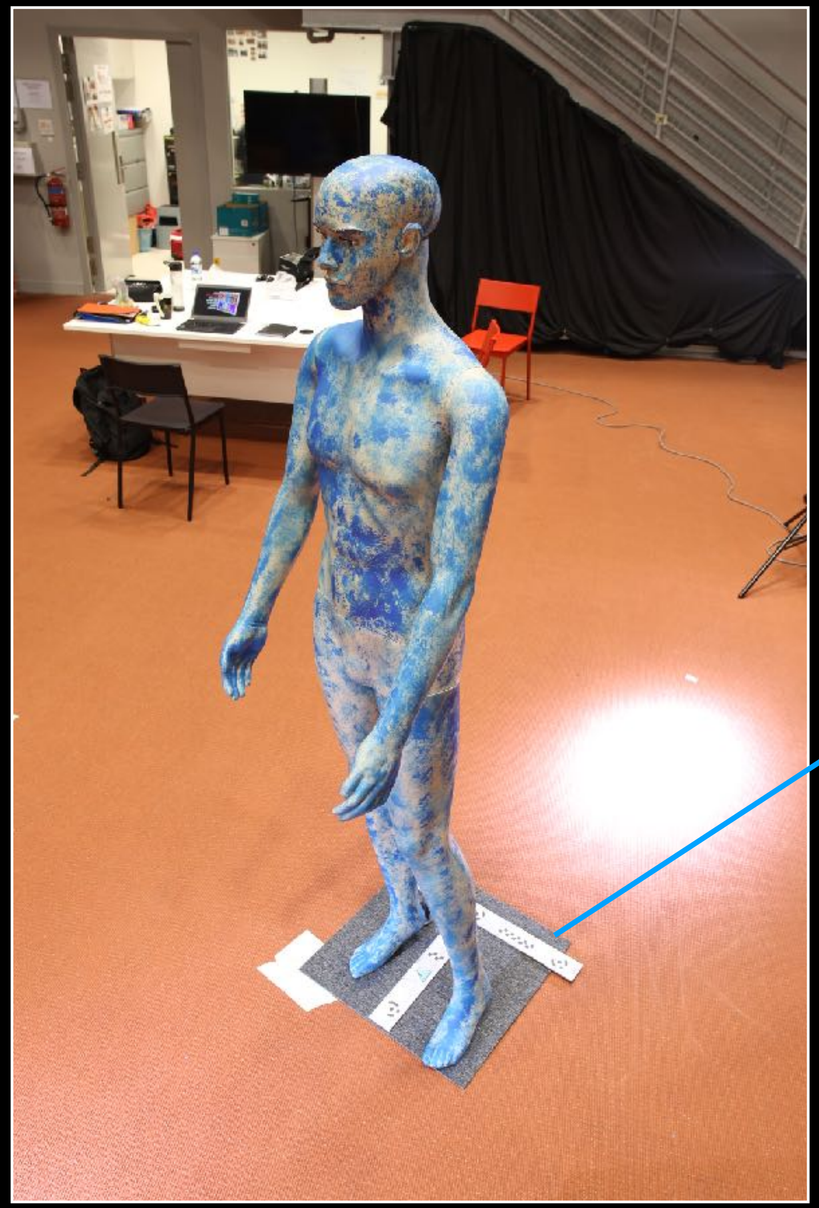

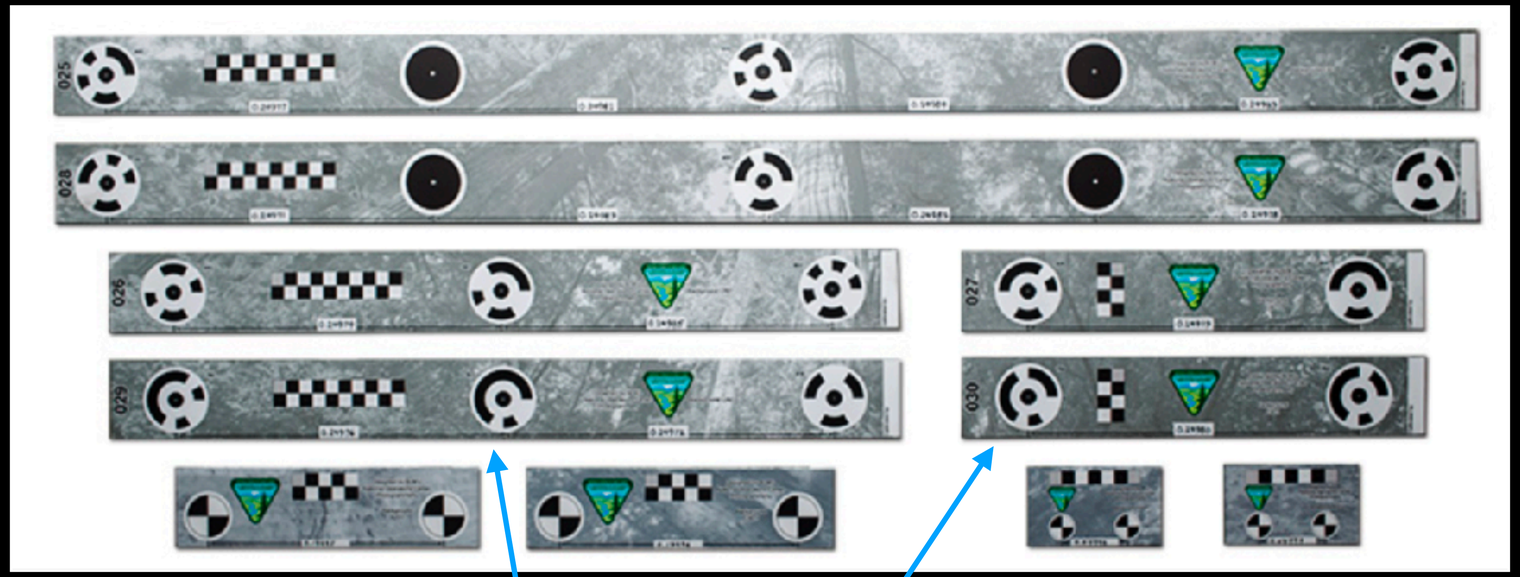

Automatically detected markers

Lens calibration

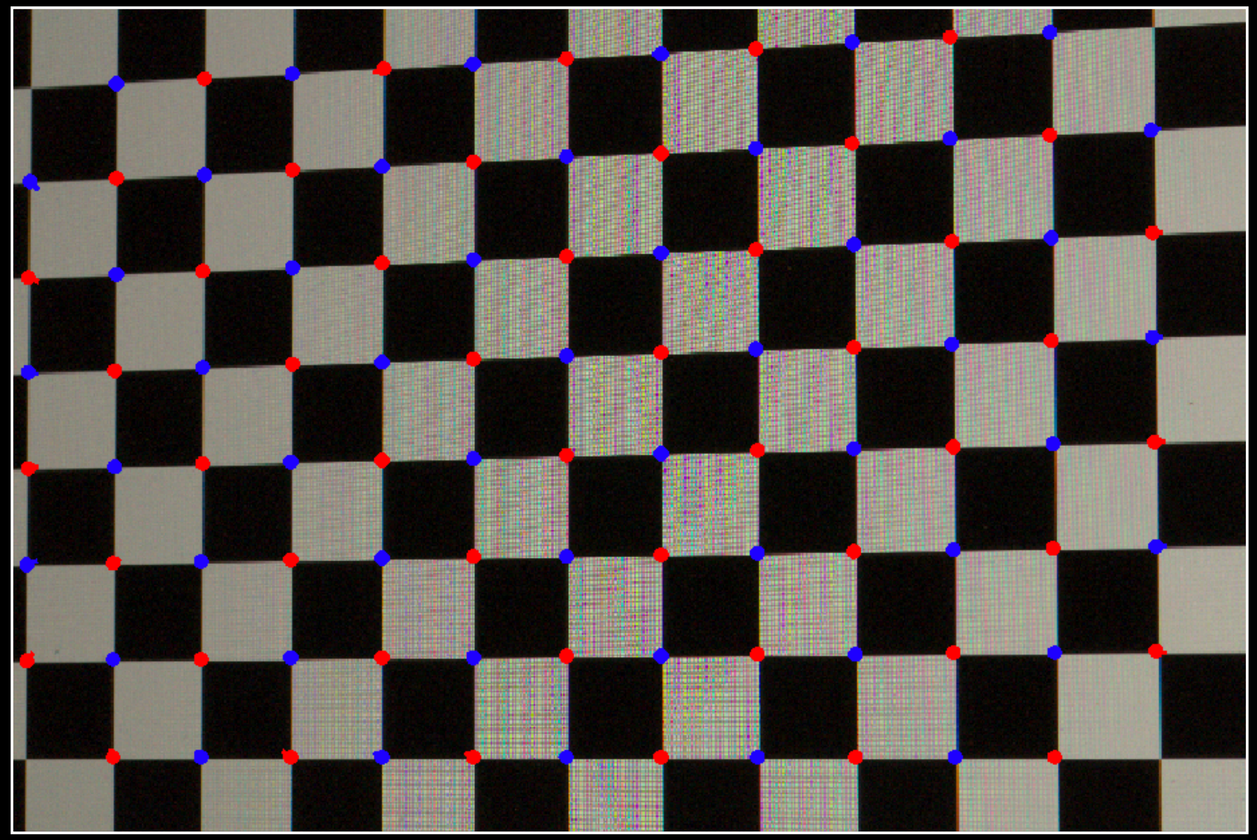

Merging multiple scans

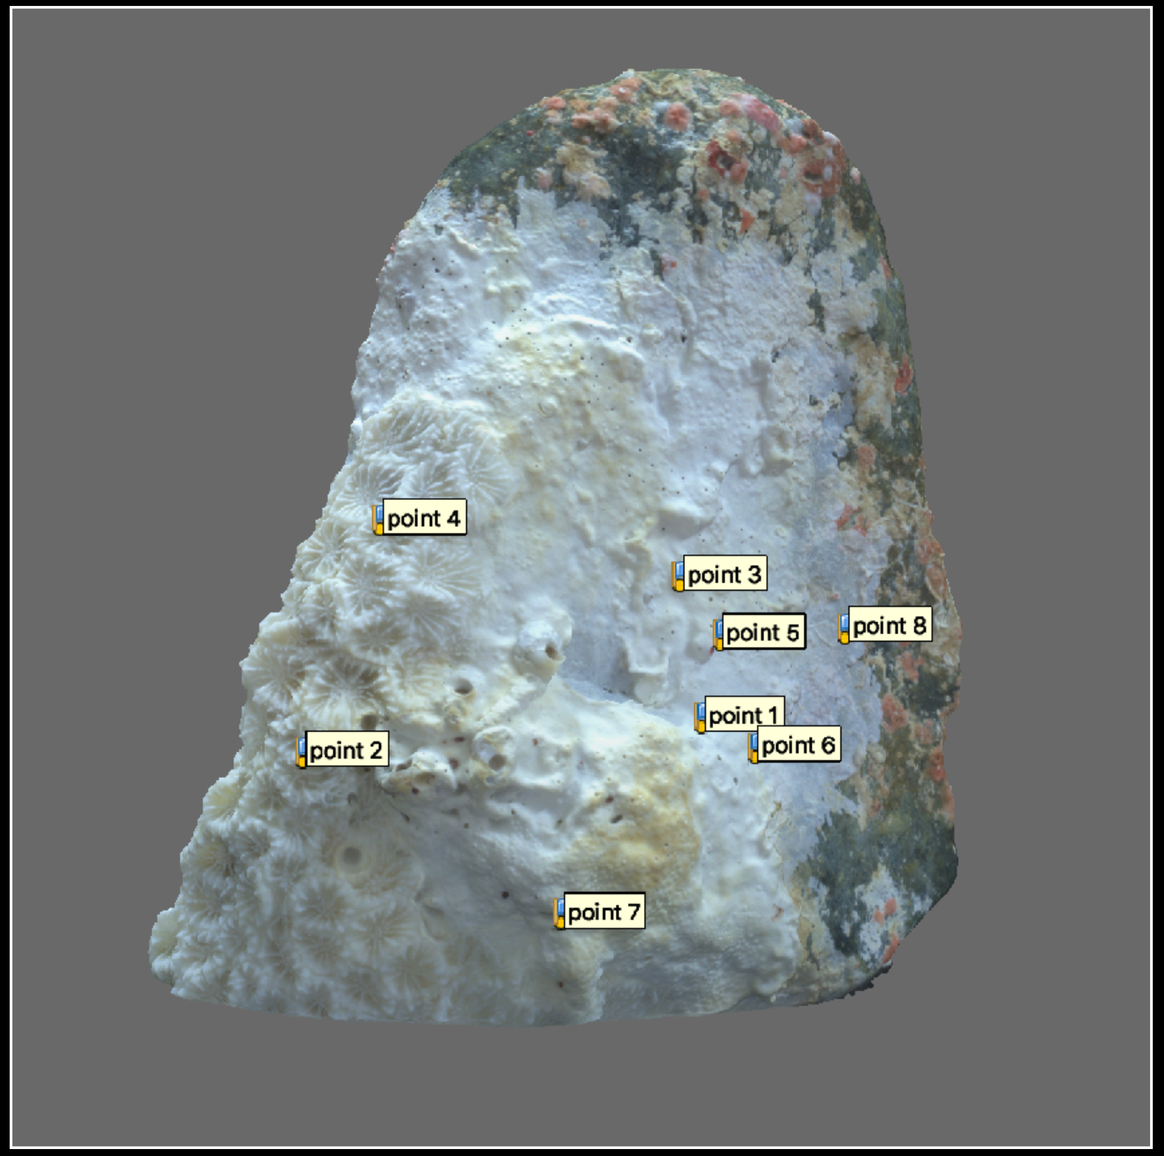
point 4
point 3
point 5
point 8
point 1
point 6
point 2
point 7

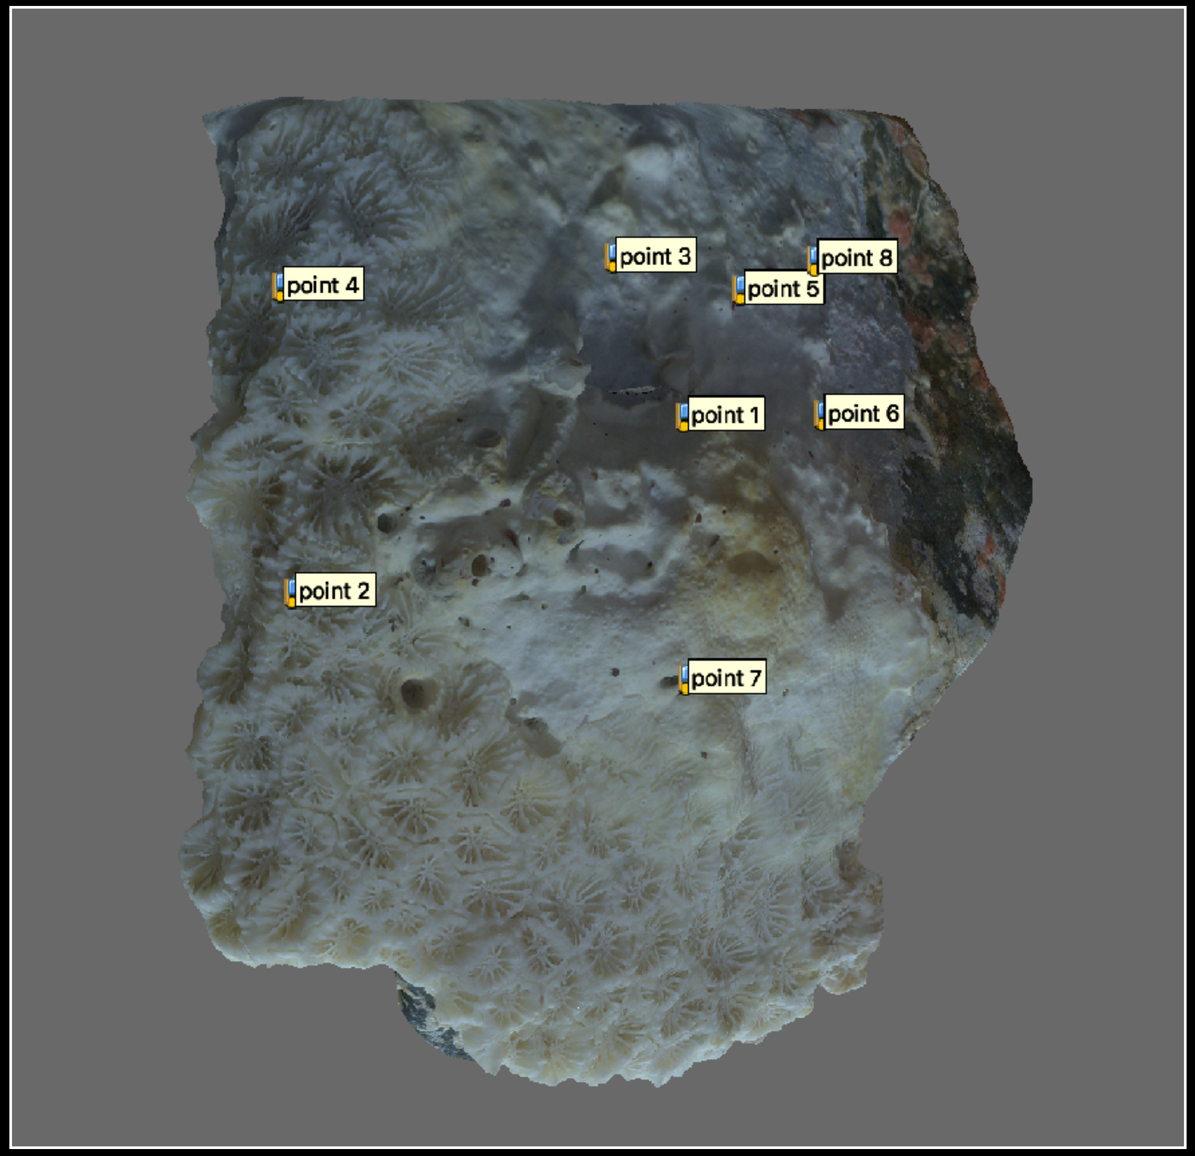
point 4
point 3
point 8
point 5
point 1
point 6
point 2
point 7

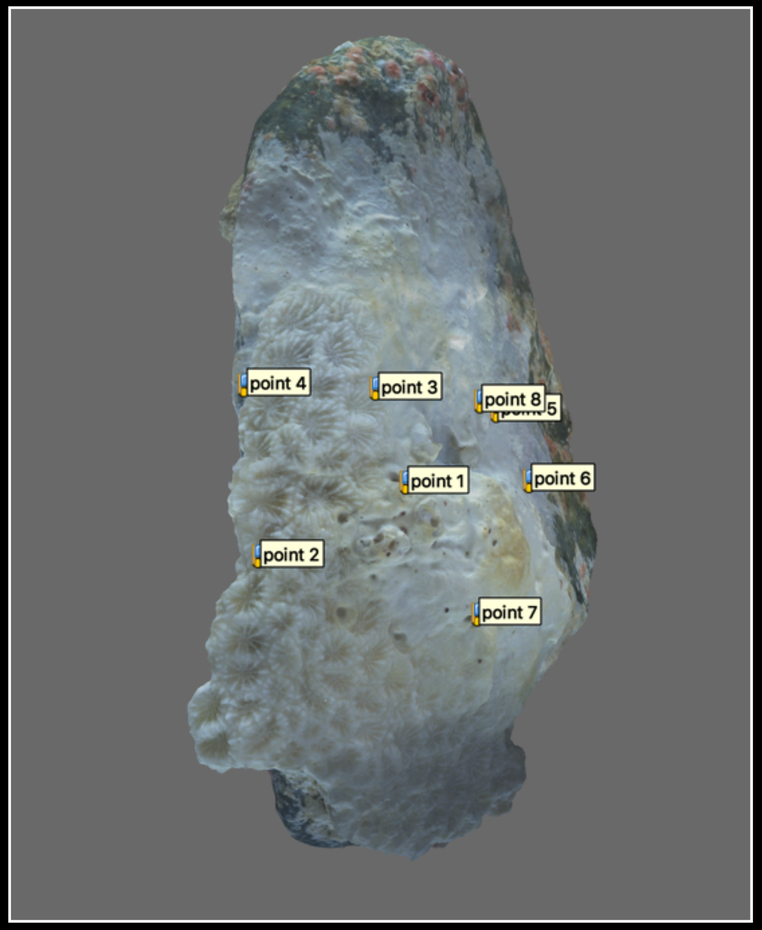
point 4
point 3
point 8
point 1
point 6
point 2
point 7

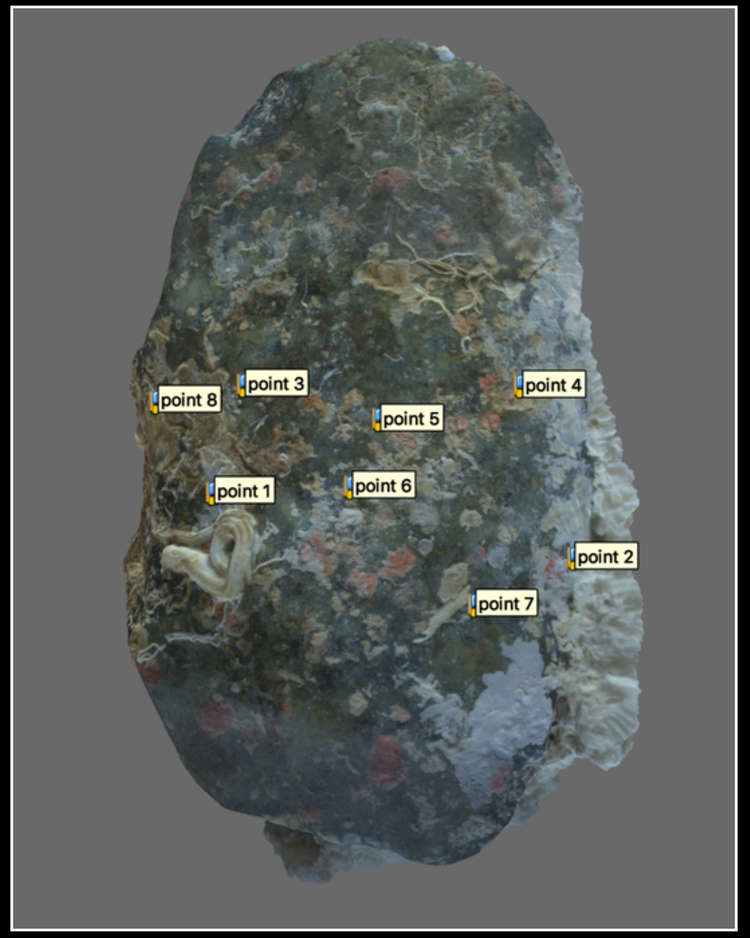
point 8
point 3
point 4
point 5
point 1
point 6
point 2
point 7

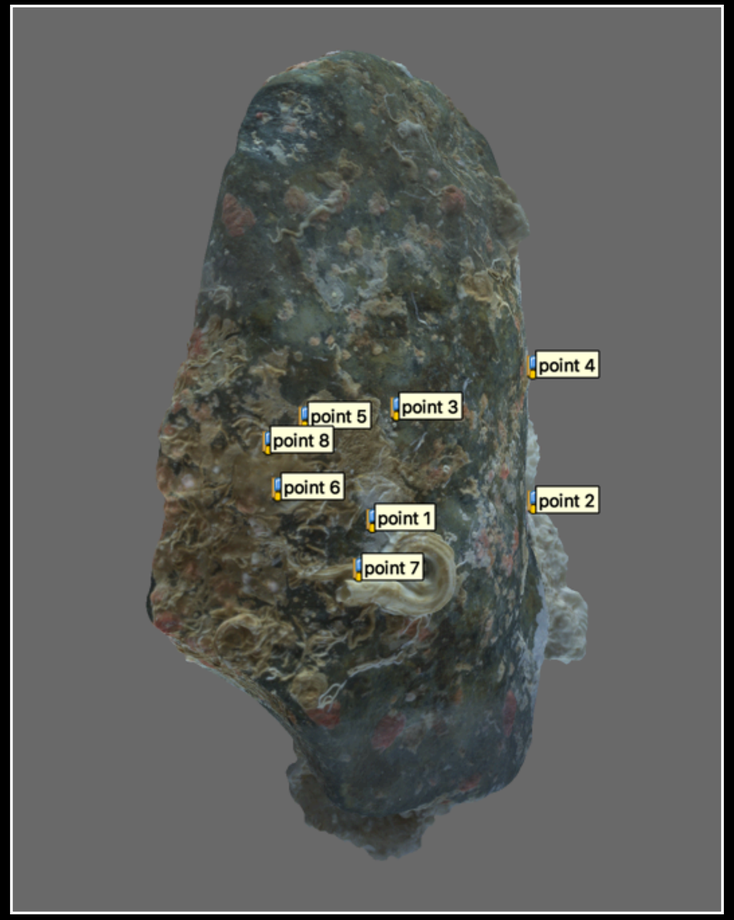
point 4
point 5
point 3
point 8
point 6
point 2
point 1
point 7

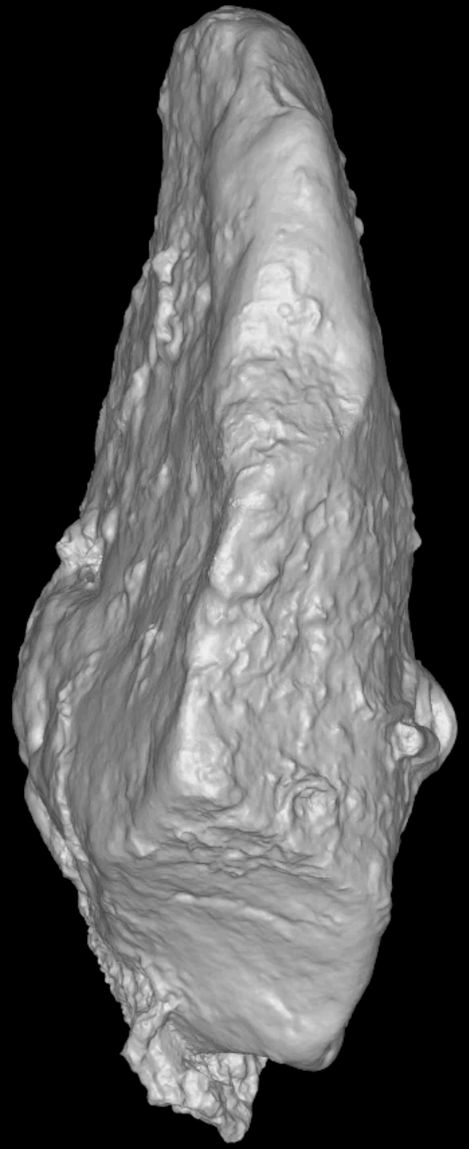

Turntables

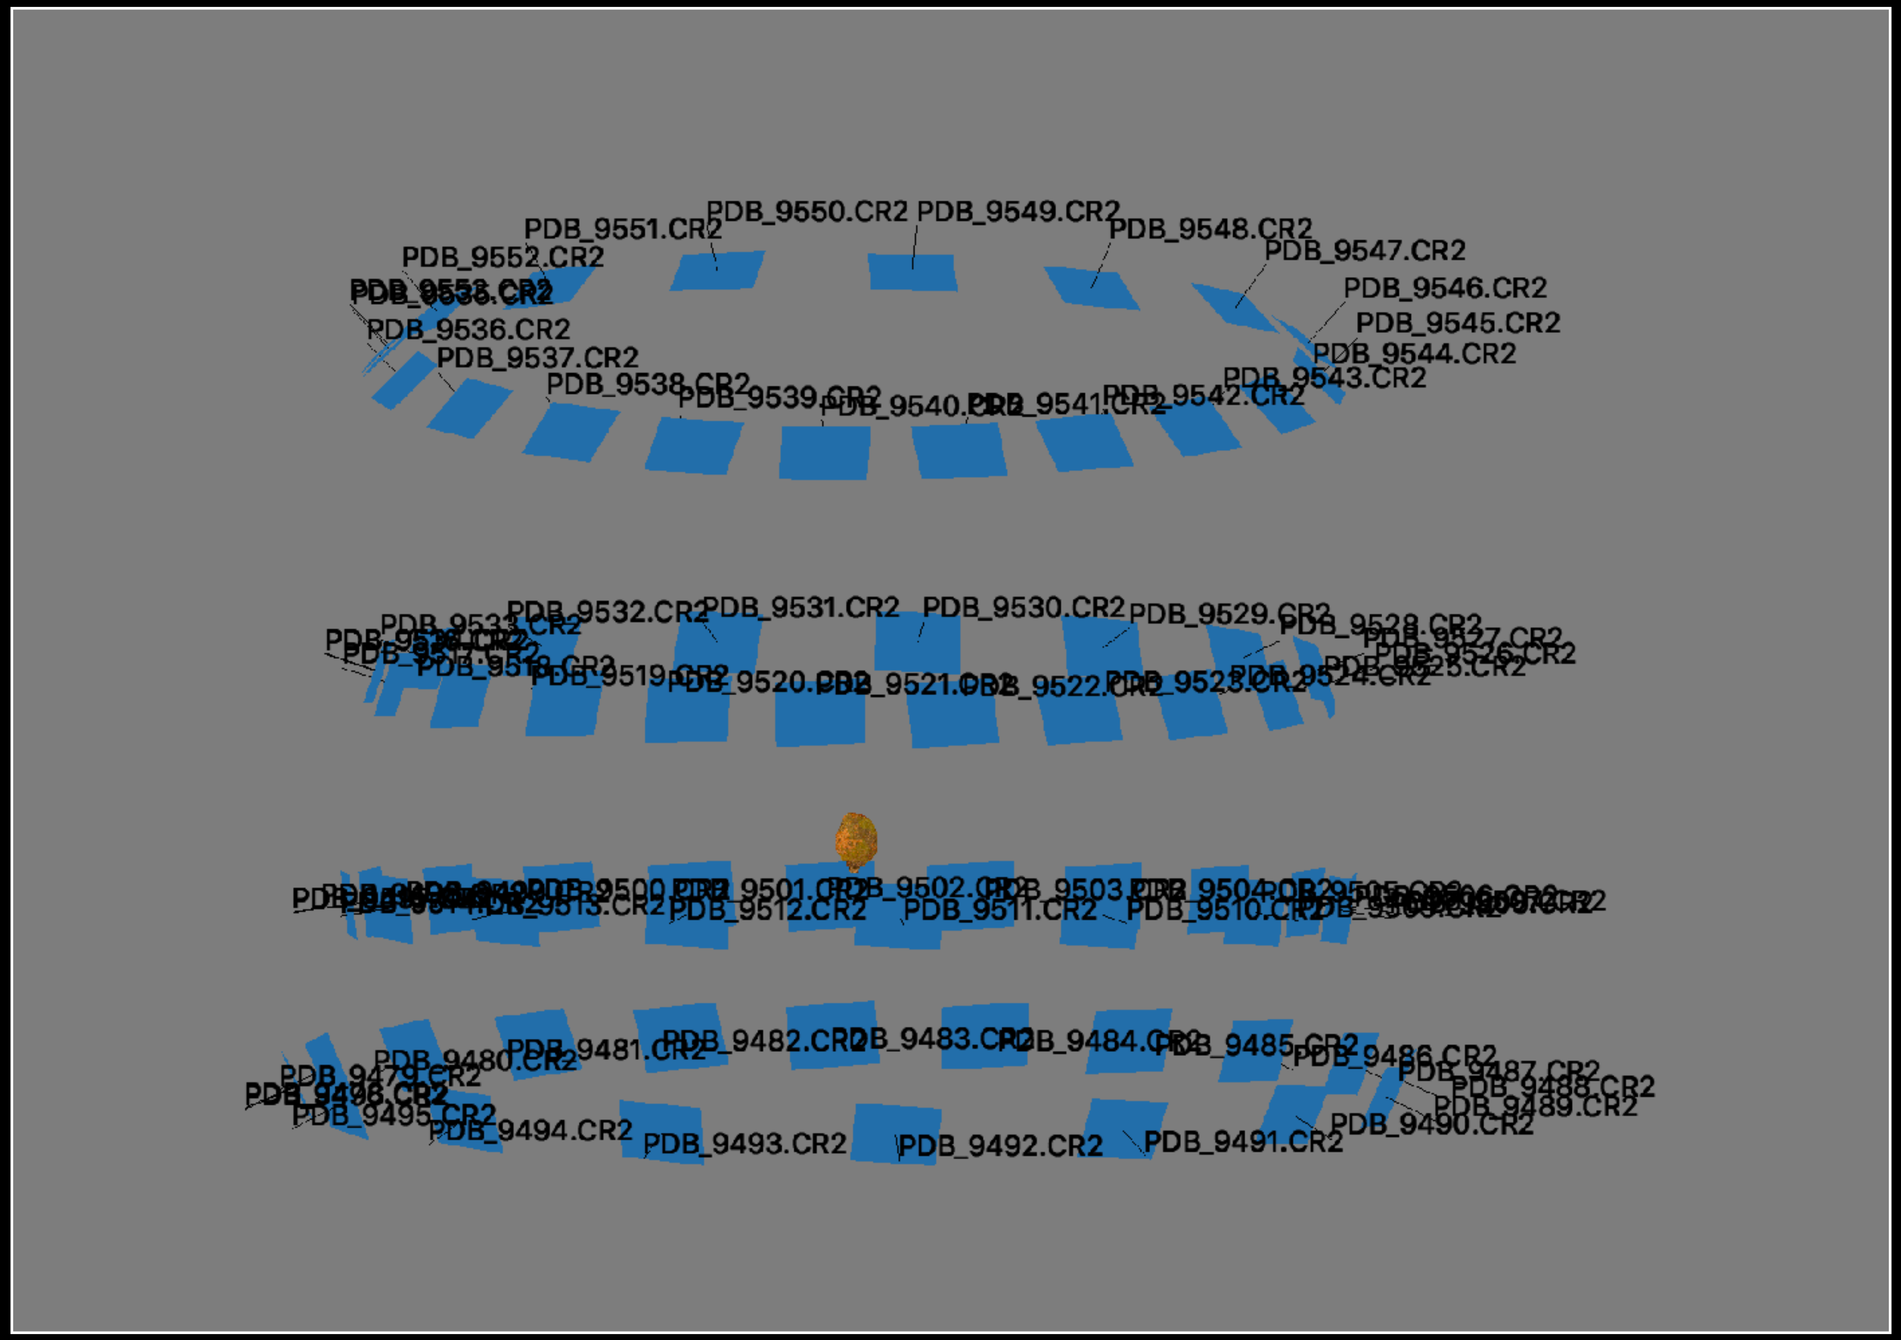
PDB_9550.CR2
PDB_9549.CR2
PDB_9551.CR2
PDB_9548.CR2
PDB_9552.CR2
PDB_9547.CR2
PDB_9546.CR2
PDB_9536.CR2
PDB_9545.CR2
PDB_9537.CR2
PDB_9544.CR2
PDB_9538.CR2
PDB_9543.CR2
PDB_9539.CR2
PDB_9542.CR2
PDB_9532.CR2
PDB_9531.CR2
PDB_9530.CR2
PDB_9529.CR2
PDB_9528.CR2
PDB_9527.CR2
PDB_9526.CR2
PDB_9520.CR2
PDB_9522.CR2
PDB_9523.CR2
PDB_9524.CR2
PDB_9512.CR2
PDB_9511.CR2
PDB_9510.CR2
PDB_9481.CR2
PDB_9482.CR2
PDB_9480.CR2
PDB_9486.CR2
PDB_9487.CR2
PDB_9488.CR2
PDB_9489.CR2
PDB_9495.CR2
PDB_9494.CR2
PDB_9493.CR2
PDB_9492.CR2
PDB_9491.CR2
PDB_9490.CR2


PDB_9478.CR2
PDB_9479.CR2
PDB_9480.CR2
PDB_9481.CR2
PDB_9482.CR2
PDB_9483.CR2
PDB_9484.CR2
PDB_9485.CR2
PDB_9486.CR2
PDB_9487.CR2
PDB_9488.CR2
PDB_9489.CR2


PDB_9478.CR2
PDB_9479.CR2
PDB_9480.CR2
PDB_9481.CR2
PDB_9482.CR2
PDB_9483.CR2
PDB_9484.CR2
PDB_9485.CR2
PDB_9486.CR2
PDB_9487.CR2
PDB_9488.CR2
PDB_9489.CR2


BOLING


JOBY

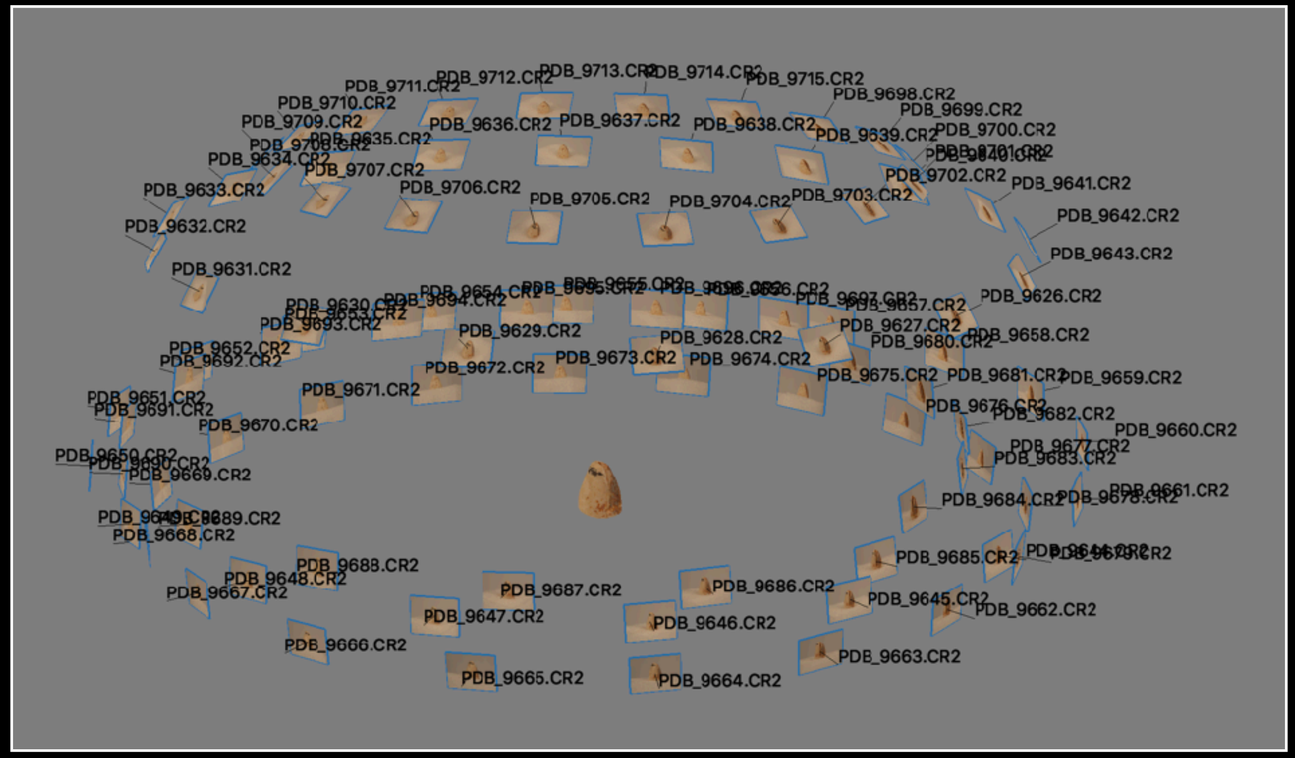

Top half + overlap

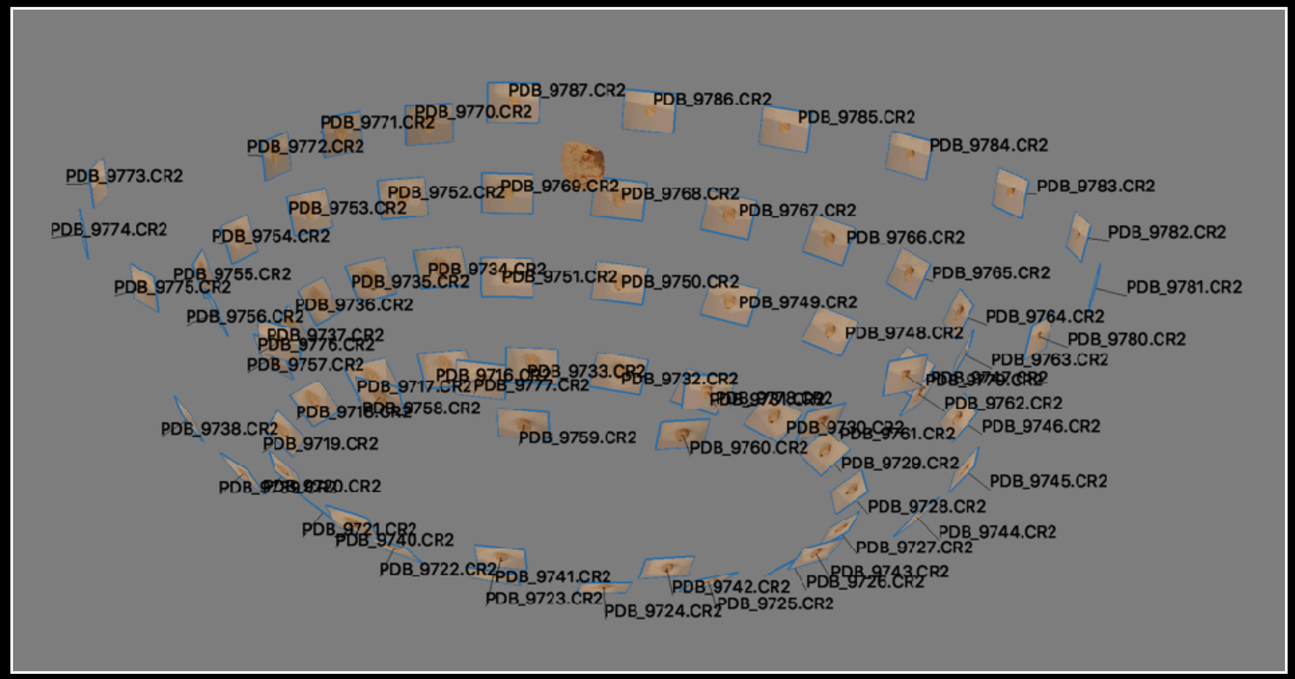

Bottom half + overlap

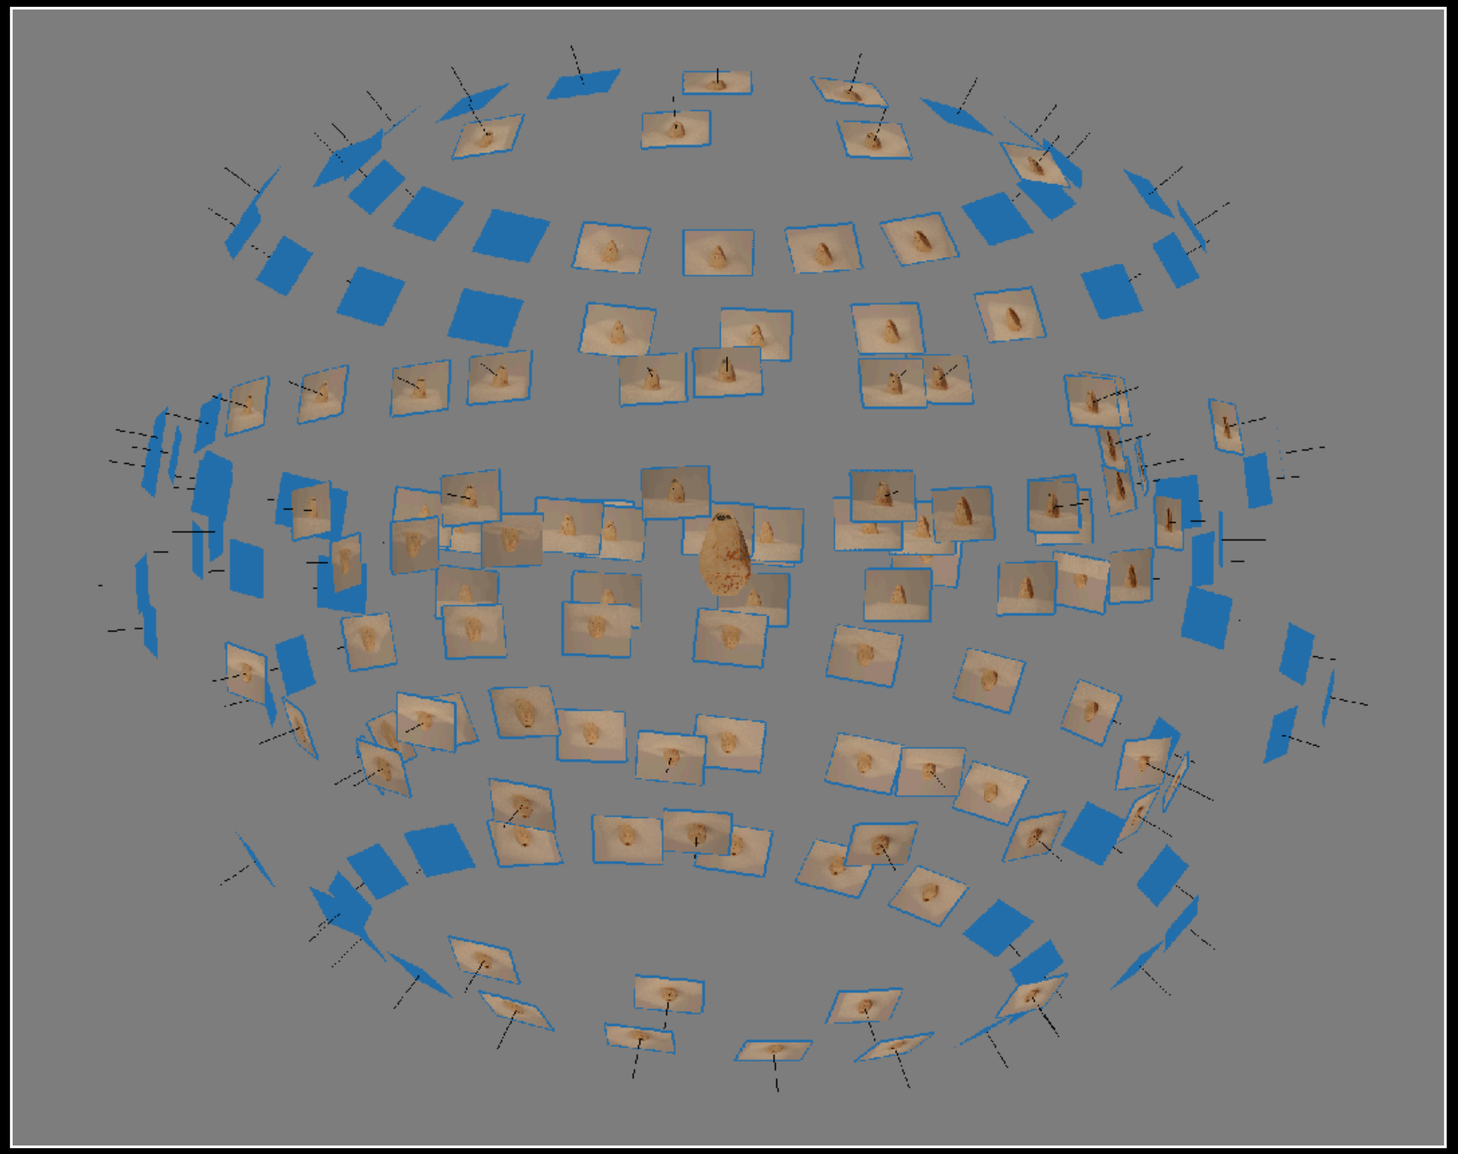

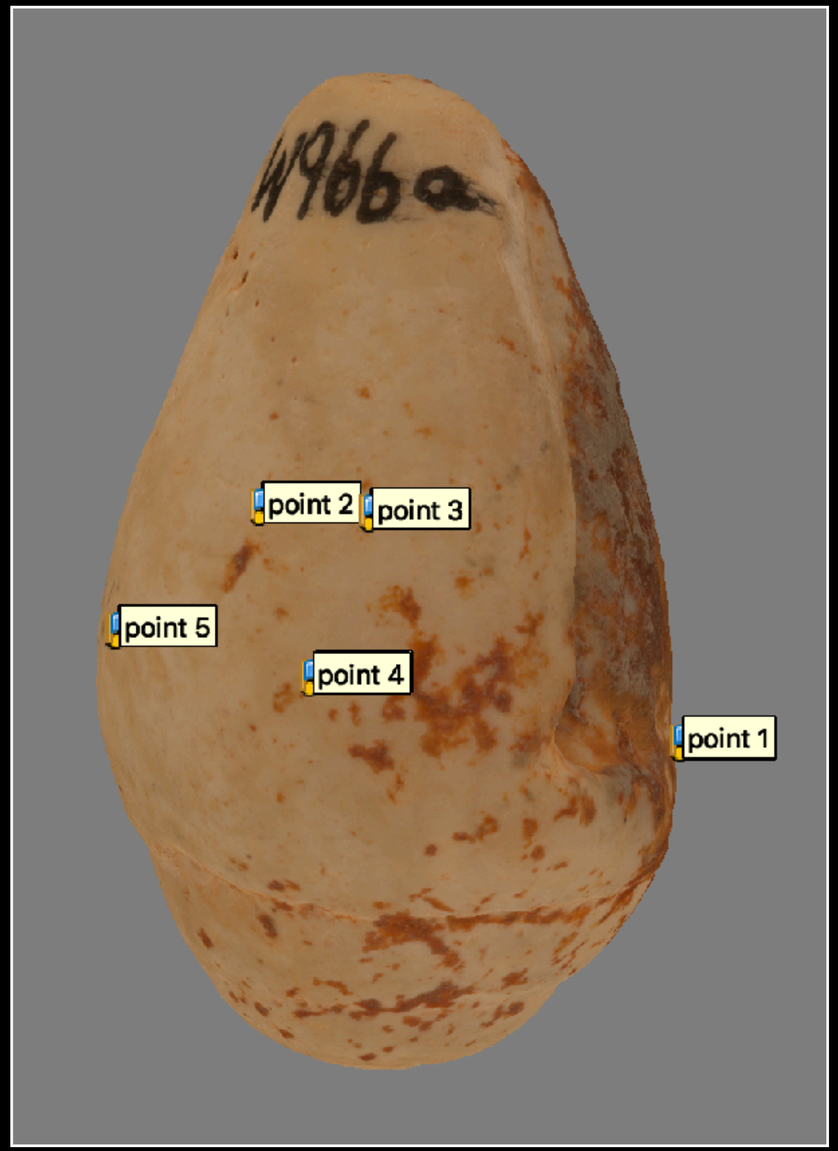
point 2
point 3
point 5
point 4
point 1

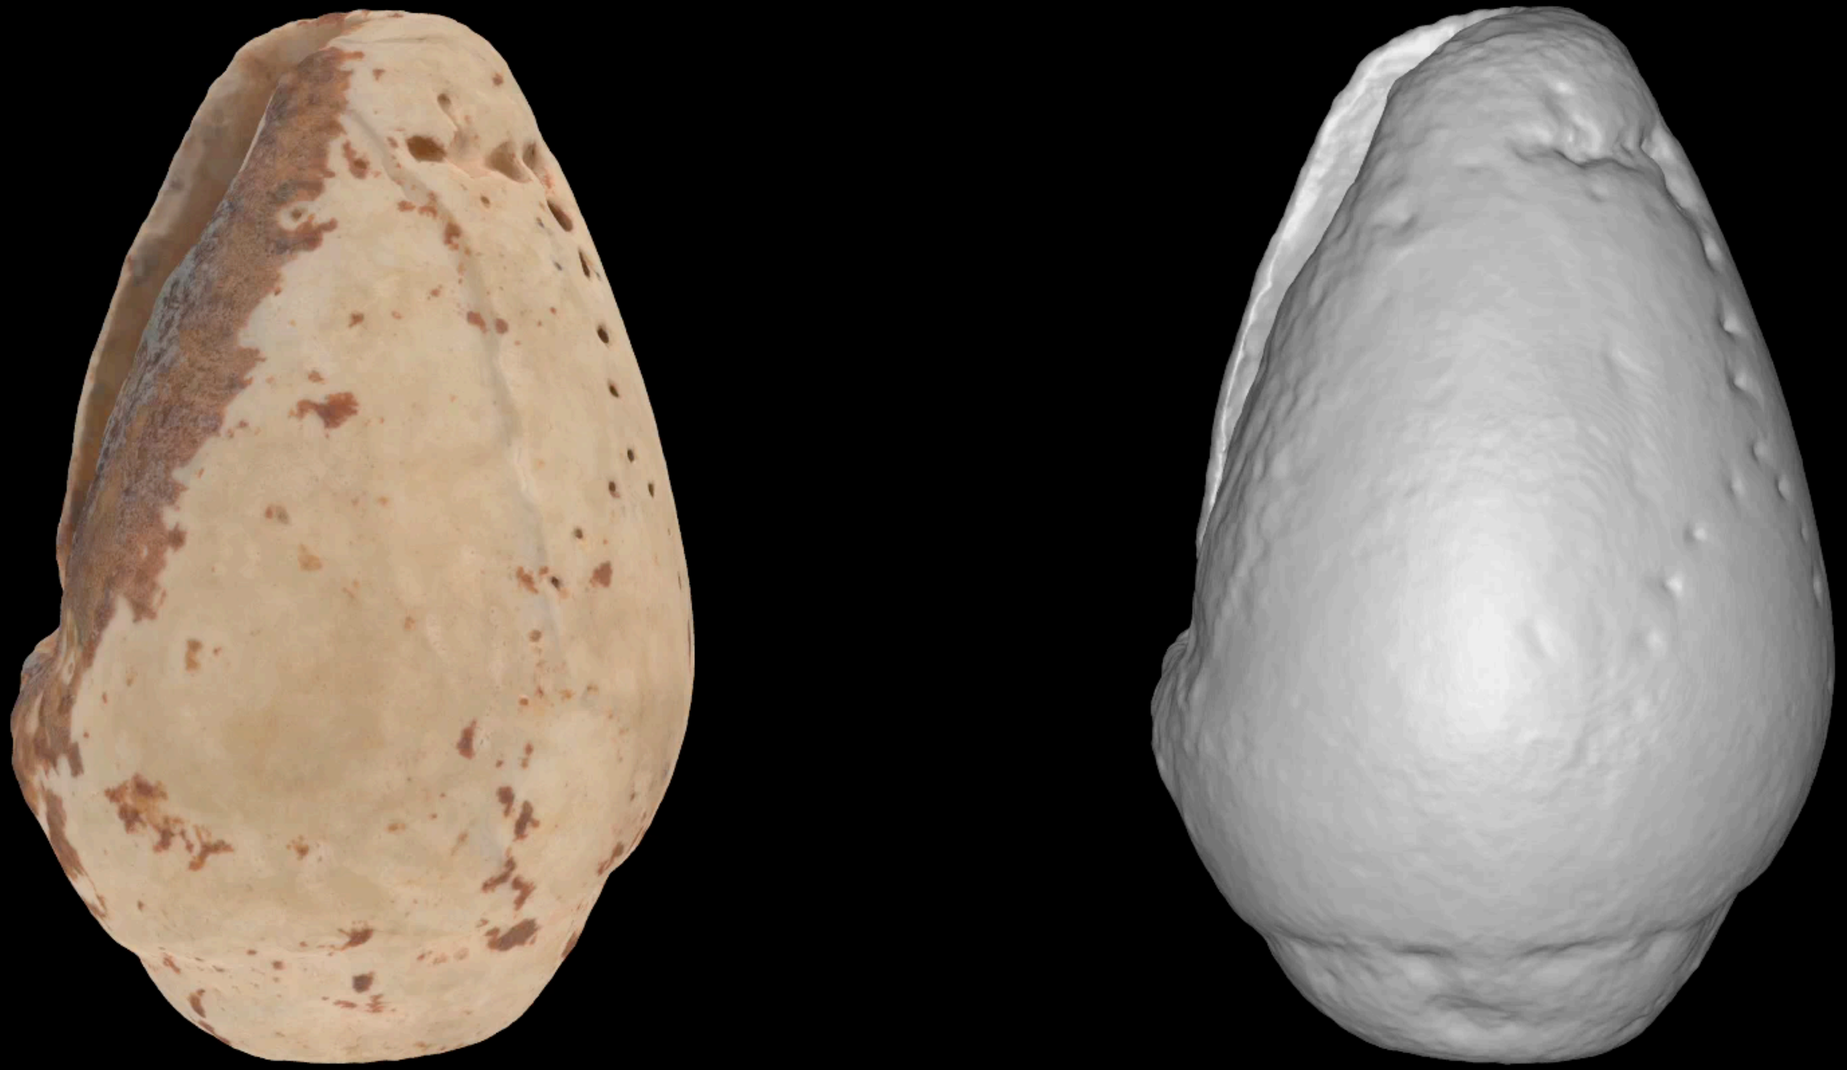

Texture editing

- Lots of applications to editing the texture file created.
  - Colour matching to a colour standard. Swatches in the reconstructed scene.
  - Applying filters for effect, eg: sharpening, increasing vibrance and so on.
  - Annotating, eg: labels, contour lines, segmentation.
  - Exposing shadow detail.

- Sometimes the baked in shadows can be removed, or differences in shading reduced.

- Holes in the model will generally not have any texture so may be filled in.

- Textures are not always clean and may need manual correction.

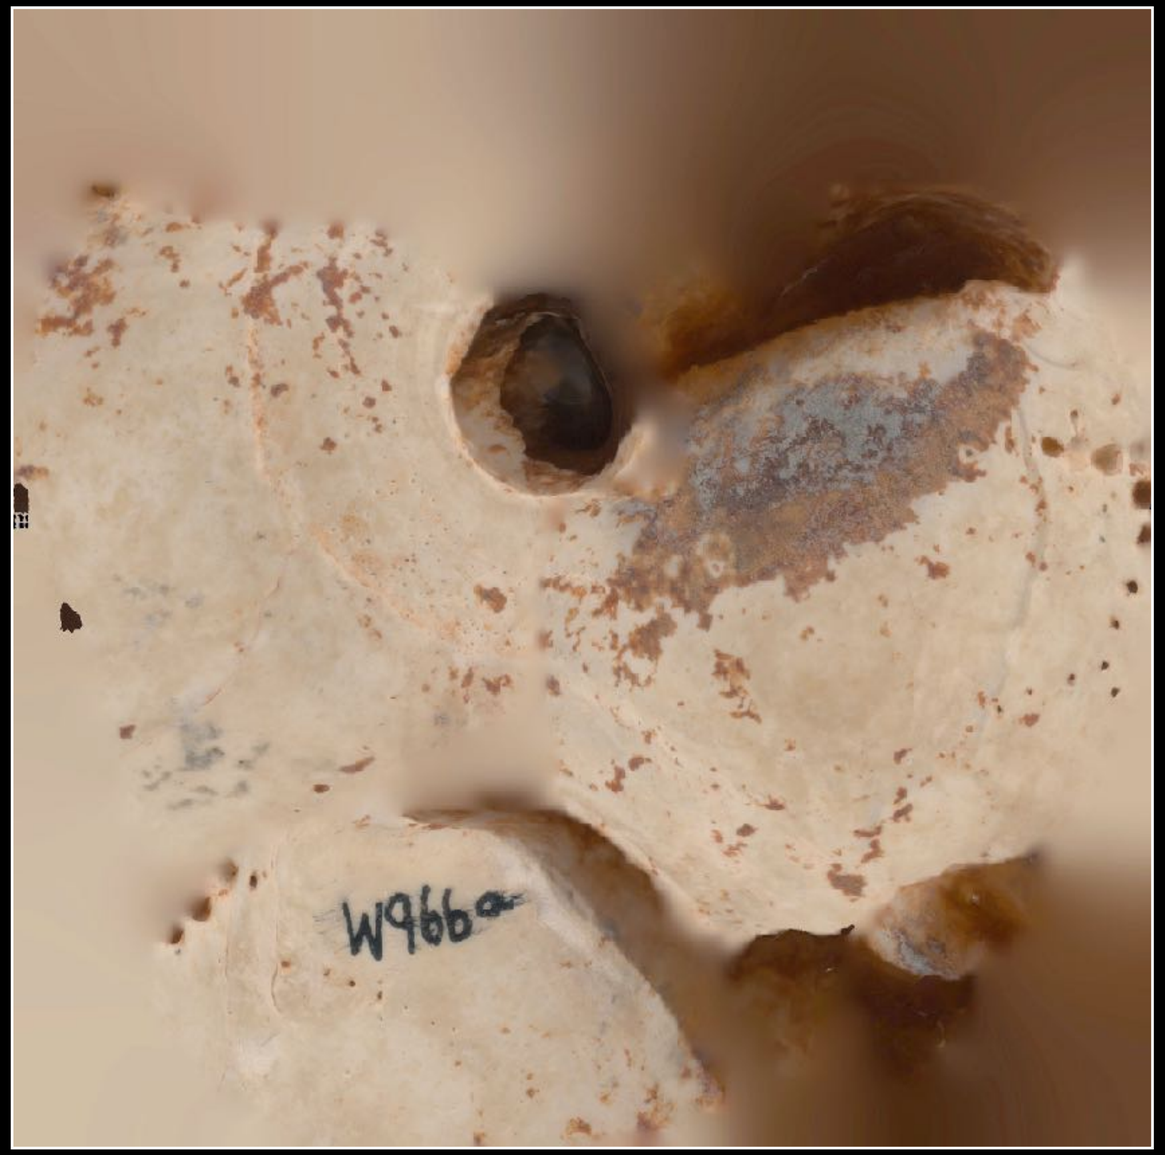
W966a


W966a

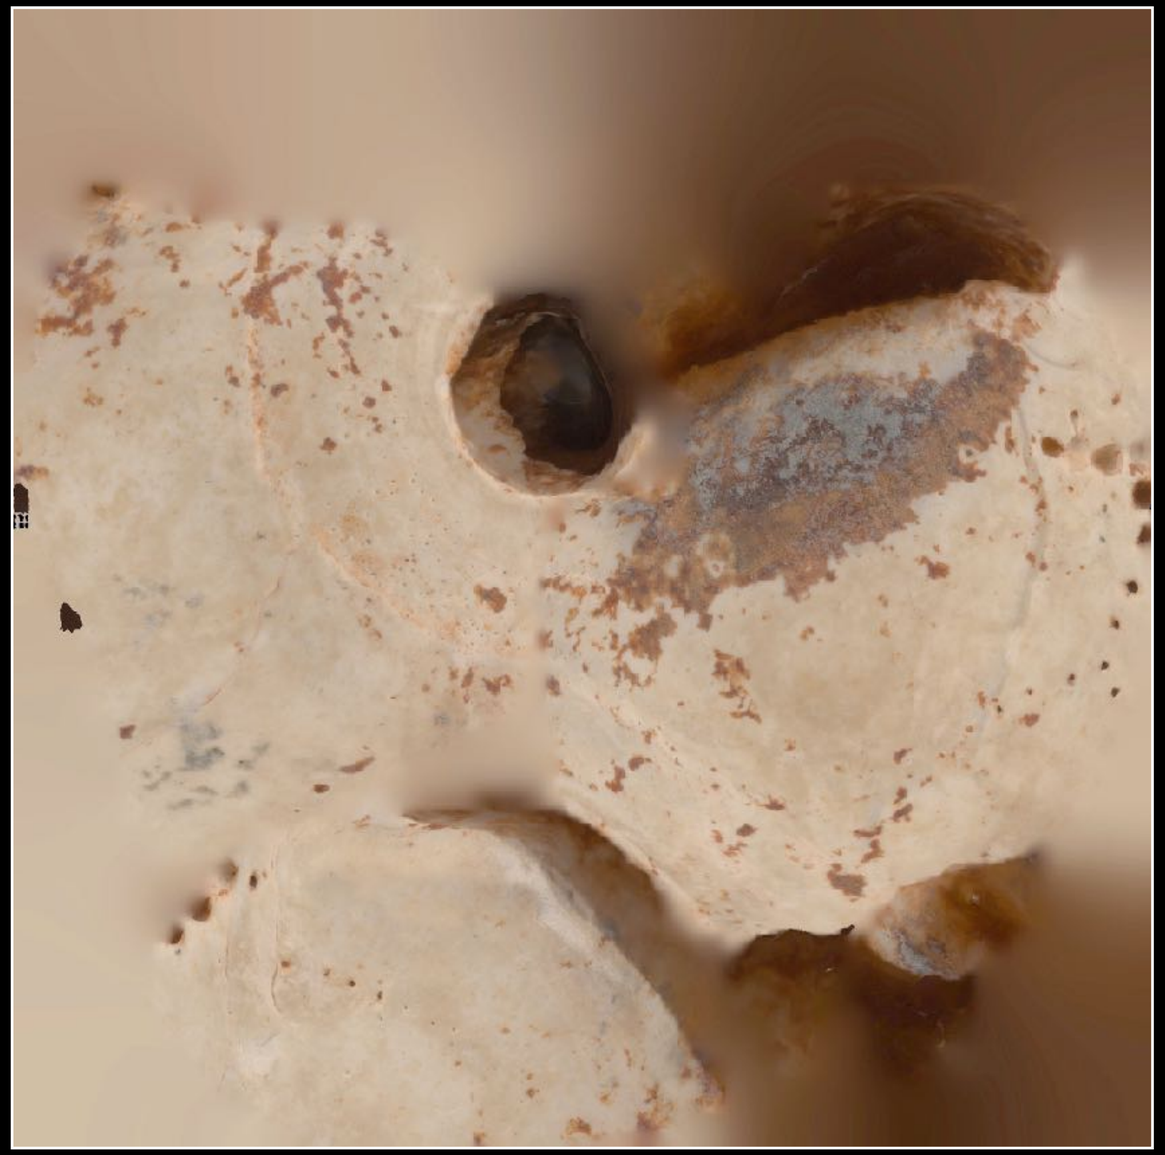

Data formats

- Requirements: unstructured triangular mesh.
  - mesh (vertices - edges - triangles - polygons)
  - texture coordinates
  - image based textures

- Common options
  - 3ds (3DStudioMax)
  - vrml, x3d
  - obj (Wavefront)
  - dae (collada)

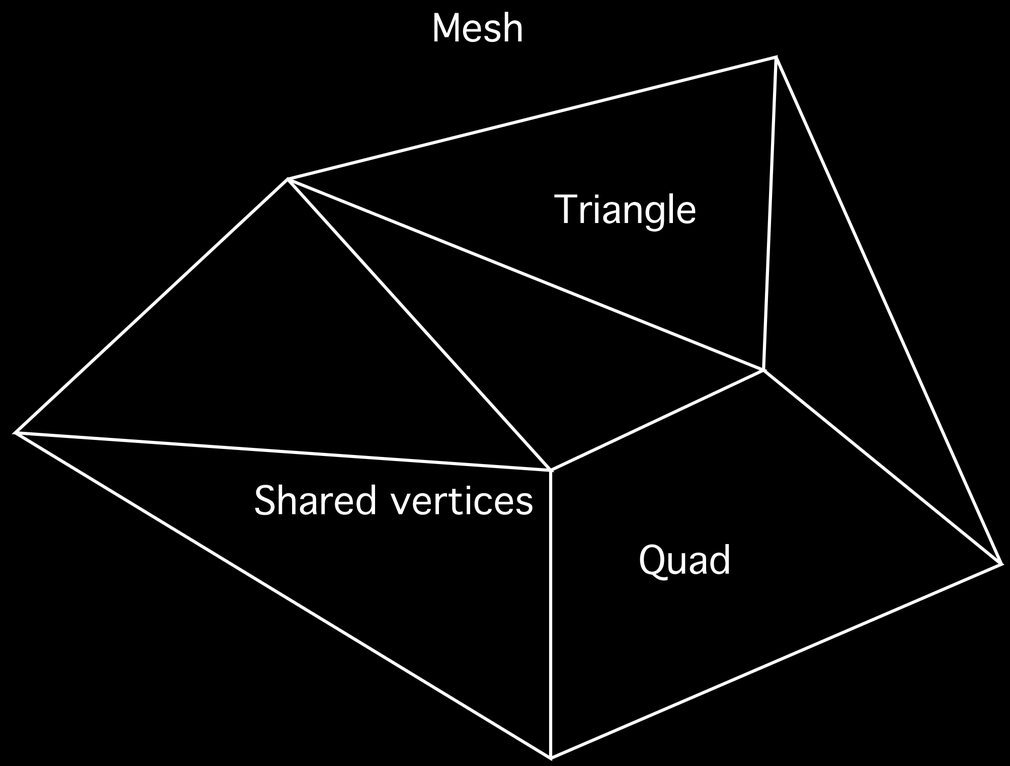

- Pretty much standardised on obj, desirable characteristics.
  - text only so human readable
  - relatively easy to parse by software for post processing or custom utilities
  - well supported by commercial 3D applications (import/export)
  - shared vertices so no chance of numerical holes
  - supports multiple texture materials and images

Anatomy of an OBJ description

Consists of 3 parts
- obj file: vertex, face, normals, texture coordinates
- materials file: colours and/or texture filename
- texture file(s): image

texture file

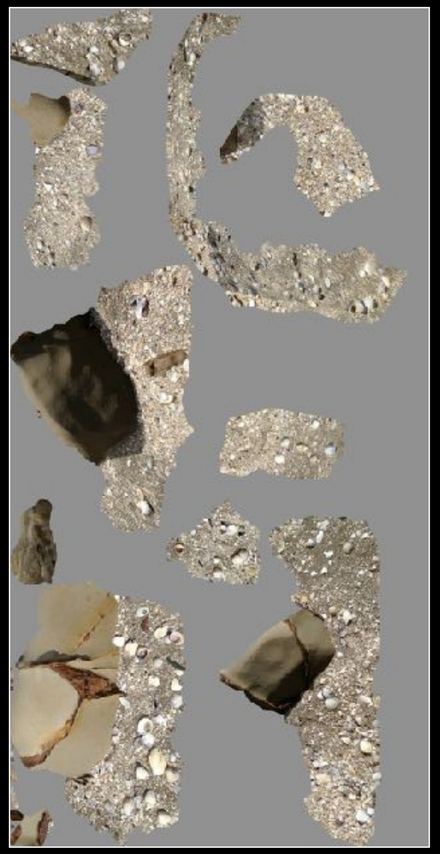

material file

```
newmtl rockmaterial
Ka 0.2 0.2 0.2
Kd 0.752941 0.752941 0.752941
Ks 1.000000 1.000000 1.000000
Tr 1.000000
illum 2
Ns 0.000000
map_Kd stone_tex_0.jpg
```

material name

obj file

```
mtllib stone.obj.mtl

v 7.980470 5.627900 3.764240
v 8.476580 2.132000 3.392570
v 8.514860 2.182000 3.396990
  :
  :
vn -0.502475 -1.595313 -2.429116
vn 1.770880 -2.076491 -5.336680
vn -0.718451 -4.758880 -3.222428
  :
  :
vt 0.214445 0.283779
vt 0.213670 0.287044
vt 0.211291 0.287318
  :
  :
usemtl rockmaterial
f 5439/4403/5439 5416/4380/5416 7144/6002/7144
f 5048/4013/5048 6581/5437/6581 5436/4400/5436
f 5435/4399/5435 5049/4014/5049 5436/4400/5436
  :
  :
```

material filename

vertices

normals

texture coordinates

triangles

vertex index

normal index

texture coordinate index

Limitations

- Require static objects otherwise a multiple camera rig required.

- Shadows and shading are baked into the textures.
  Shadows of the photographer.
  Diffuse lighting ideal = cloudy skies

- In the field, one cannot always photograph from desired locations.
  Poles - Drones - Cranes - Ladders

- Cannot capture structures that are below a few pixels wide in photographs.

- Reflective, specular surfaces problematic.
  Dusting - Painting - Polarisation technique

- Smooth featureless surfaces problematic.
  Dusting - Painting

Camera rigs for human scanning

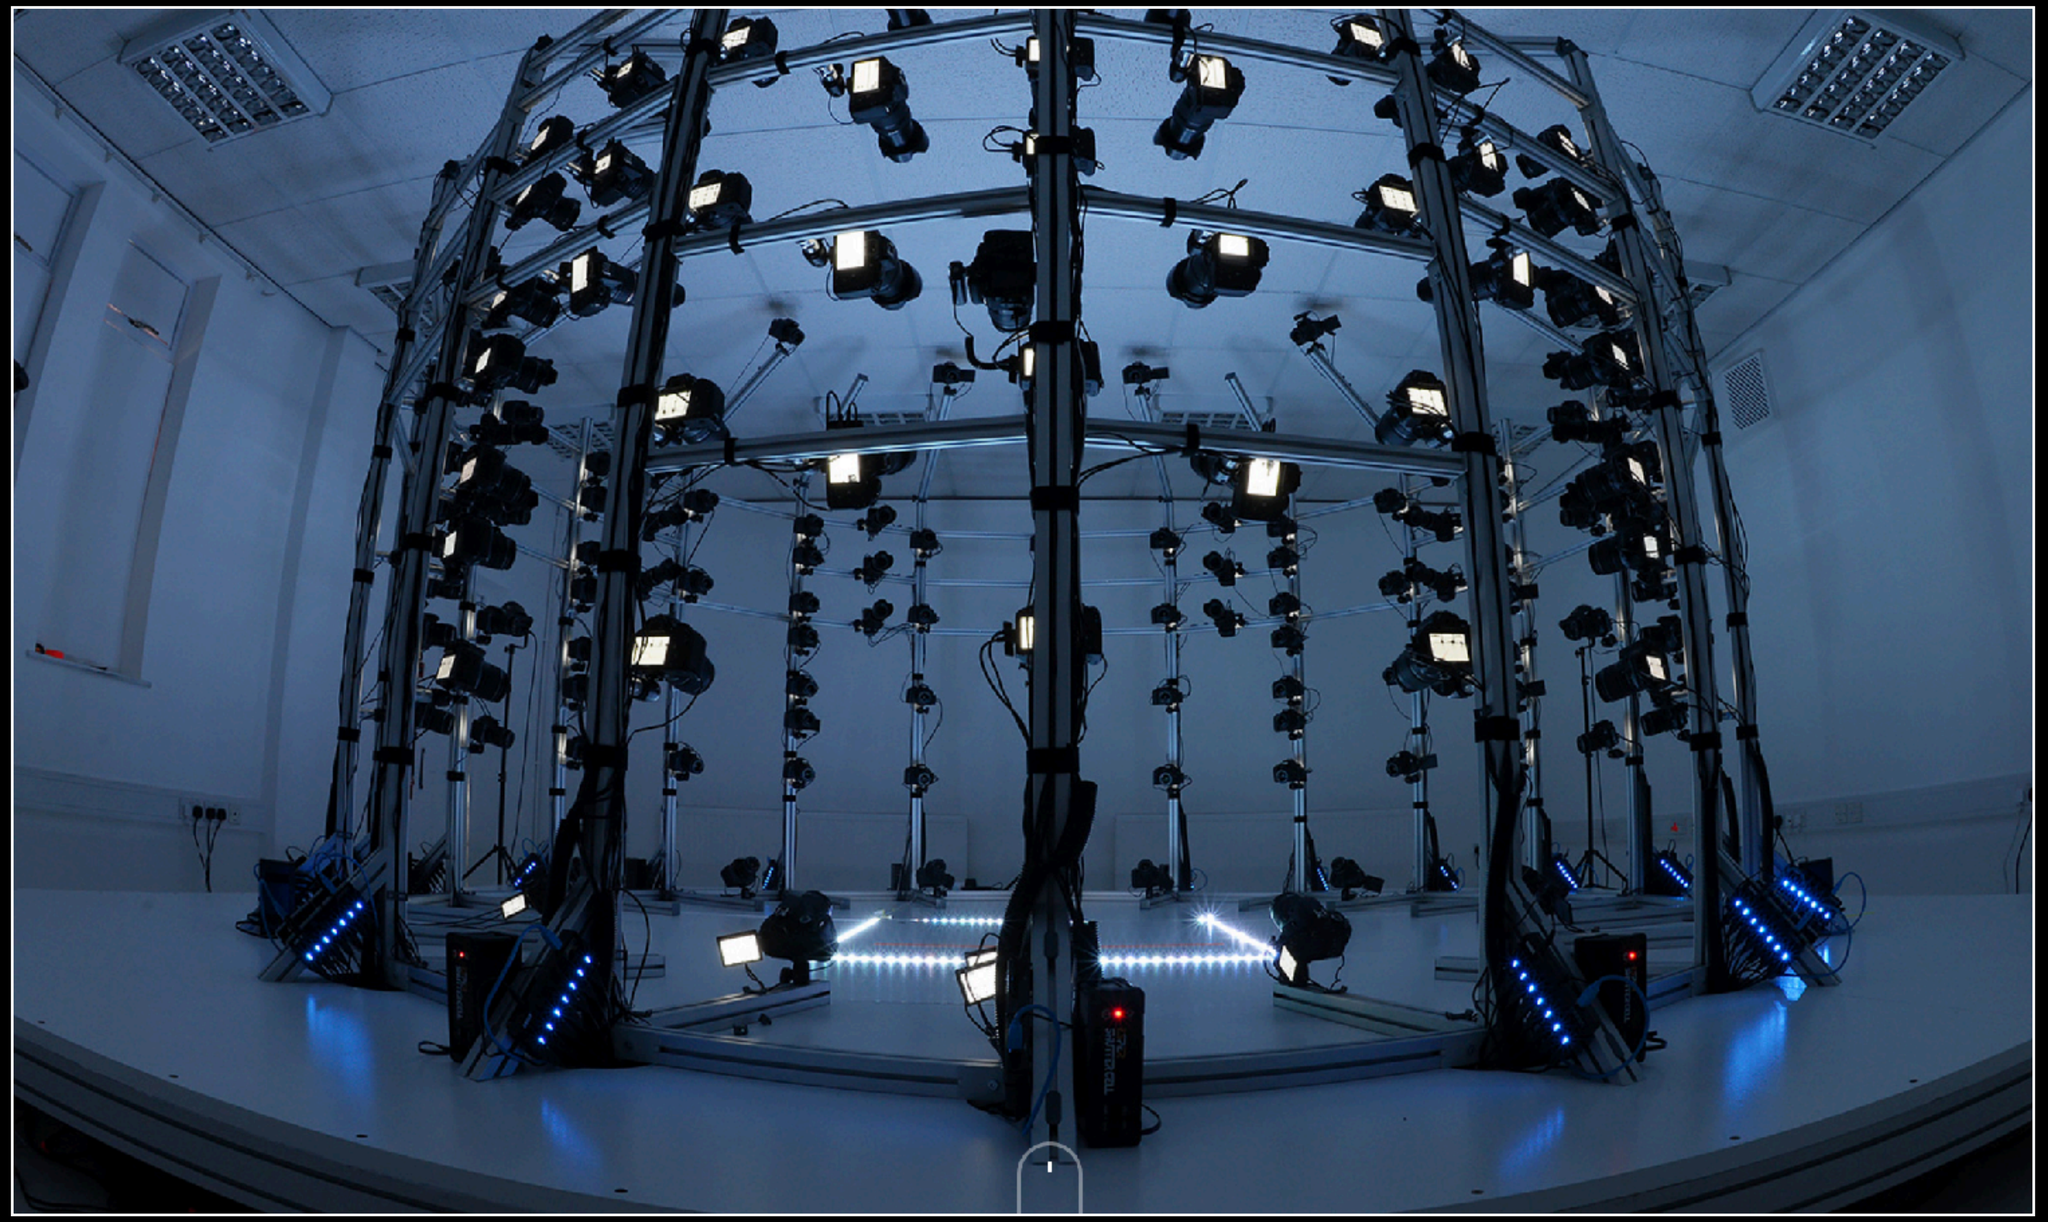


Artec Group

Shadows and Shading

Access to camera positions

Post Production

- Mesh simplification
- Thickening (3D printing)
- Relighting
- Hole closing (“watertight” for 3D printing)
- Removing shrapnel
- Per vertex editing
- Smoothing

Mesh simplification

- Meshes directly from the reconstruction (generated from the dense point cloud) are generally inefficient. Often need to reduce them for realtime applications and/or web based delivery.
- Also used to create multiple levels of details (LOD) for gaming and other realtime applications.
- The goal is easy to understand: remove mesh density where it will make minimal impact on the mesh appearance. For example, don’t need high mesh density in regions of low curvature.
- Most common class of algorithm is referred to as “edge collapse”, replace an edge with a vertex.
- A texture and geometry approximation ... need to estimate new texture coordinate at new vertices.
- Need to preserve the boundary.

Edge collapse

- Most edge collapse algorithms involve replacing an edge with a vertex
  - How to choose the edges to remove is the “trick”.
  - Where to locate the new vertex so as to minimise the effect on the surface.
  - How to estimate the new texture coordinate.

- Number of triangles reduces by 2 on each iteration.

- Can calculate the deviation of the surface for any particular edge collapse.
  Choose edges that result in the smallest deviation.
  For example: remove edges on flat regions, retain edge in regions of high curvature.

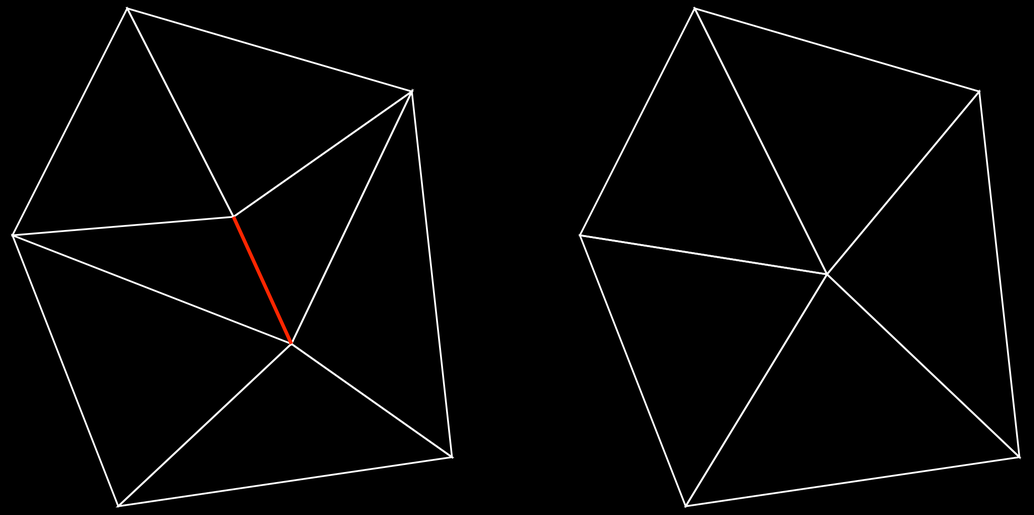

Red edge removed, results in two fewer triangles

1,000,000 triangles

100,00 triangles

1,000,000 triangles

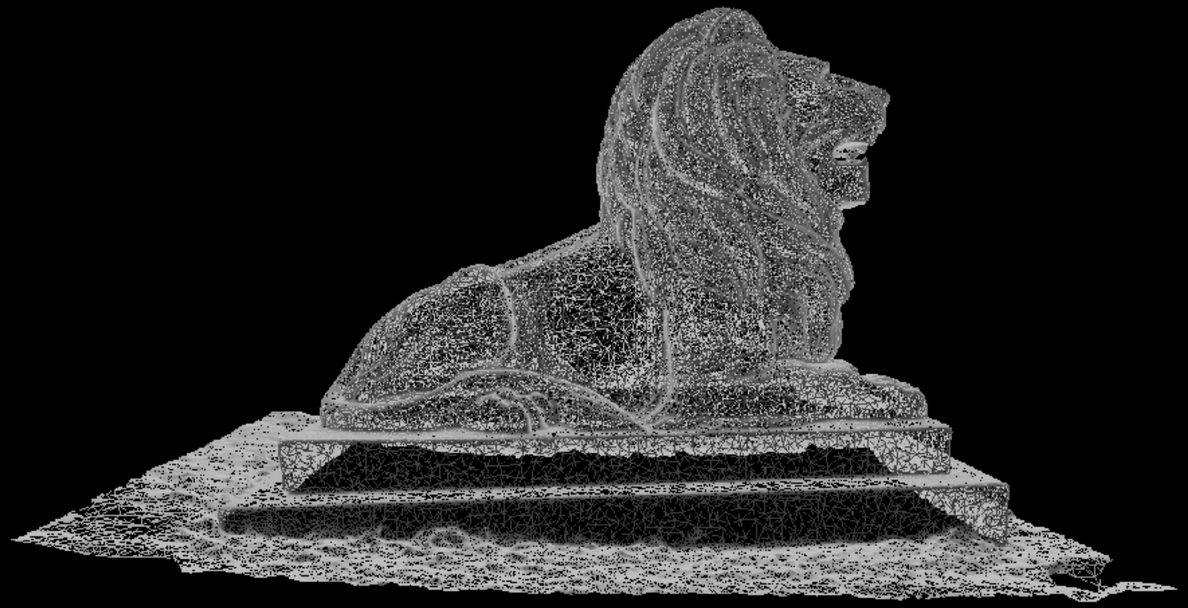

100,00 triangles

Mesh thickening

- Cases exist where one does not want idealised “infinitely thin” surfaces.
- Double sided rendering in realtime APIs is not quite the same visual effect as physical thickness.
- Required to create physical models, see rapid prototyping later.
- Perhaps the most common algorithm is known as “rolling ball”.

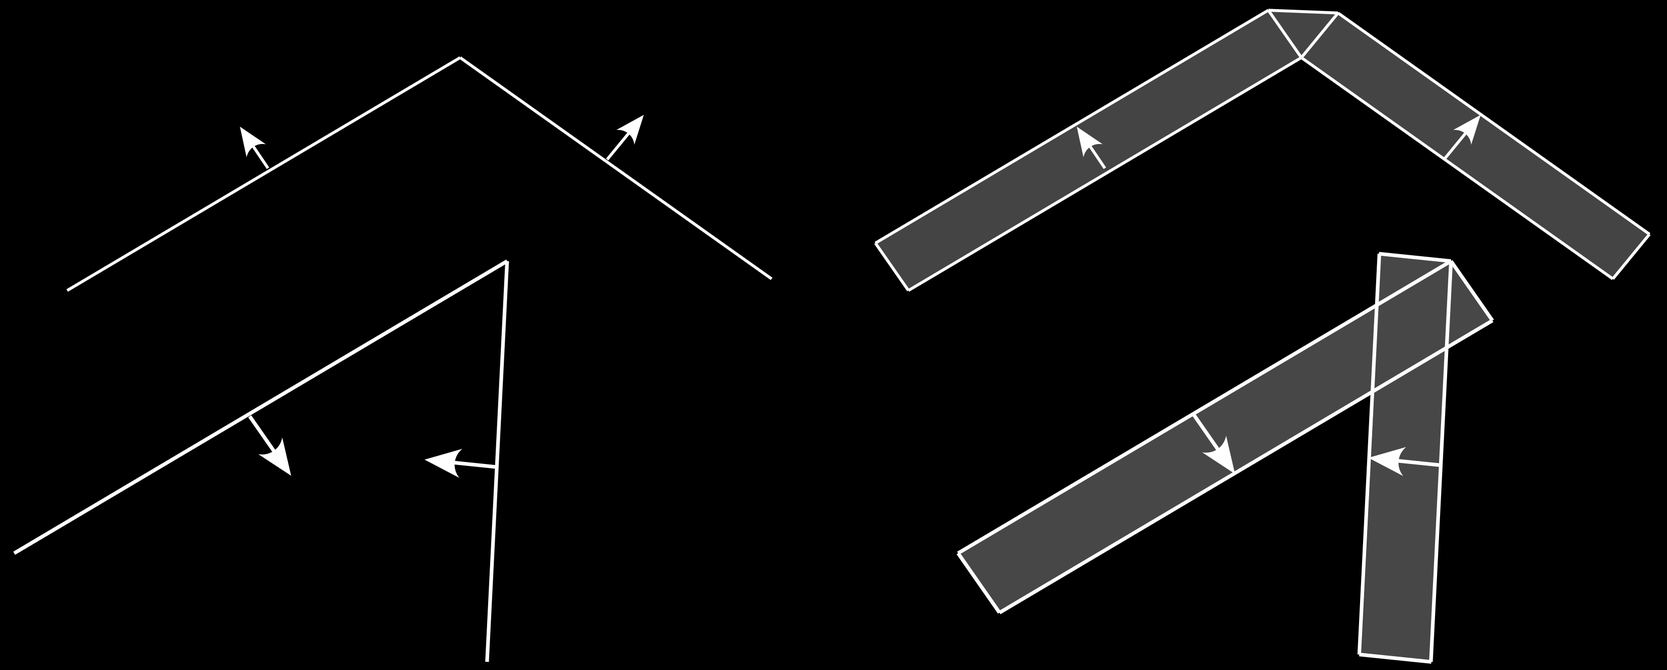

Thin joints arise at regions of high curvature

Get “poke-through” with sharp concave interiors

Ball rolling algorithm

- Solution is called “rolling ball” thickening.
- Imagine a ball rolling across the surface, form an external mesh along the ball path.
- Implemented in Blender as a modifier called “solidify”.

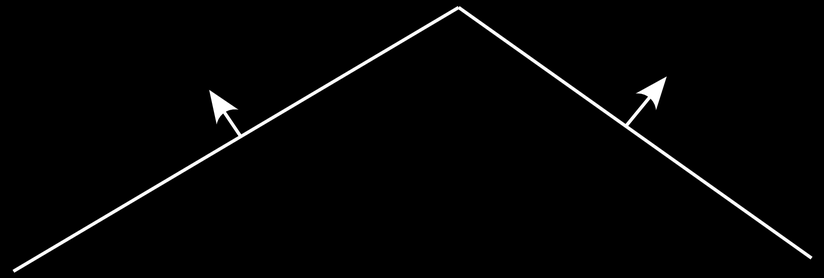

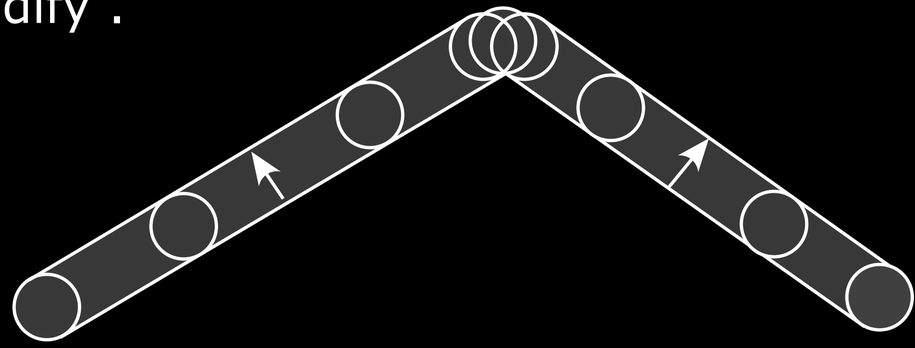

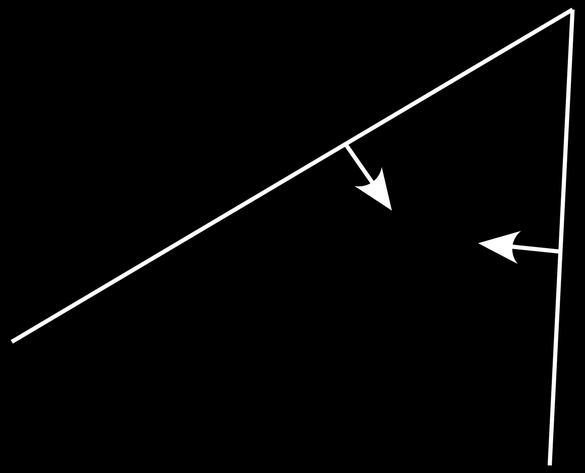

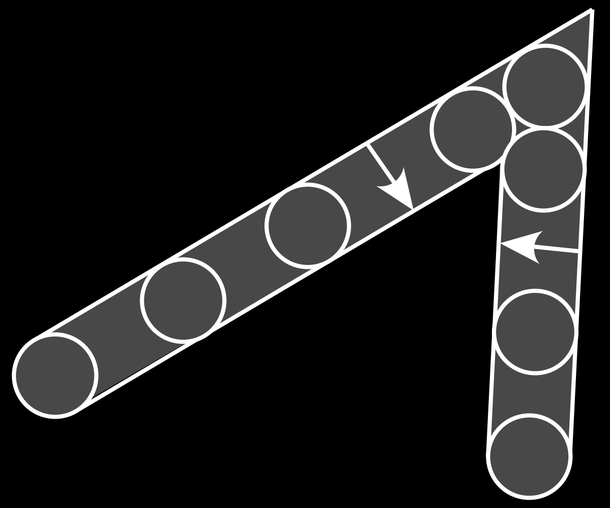

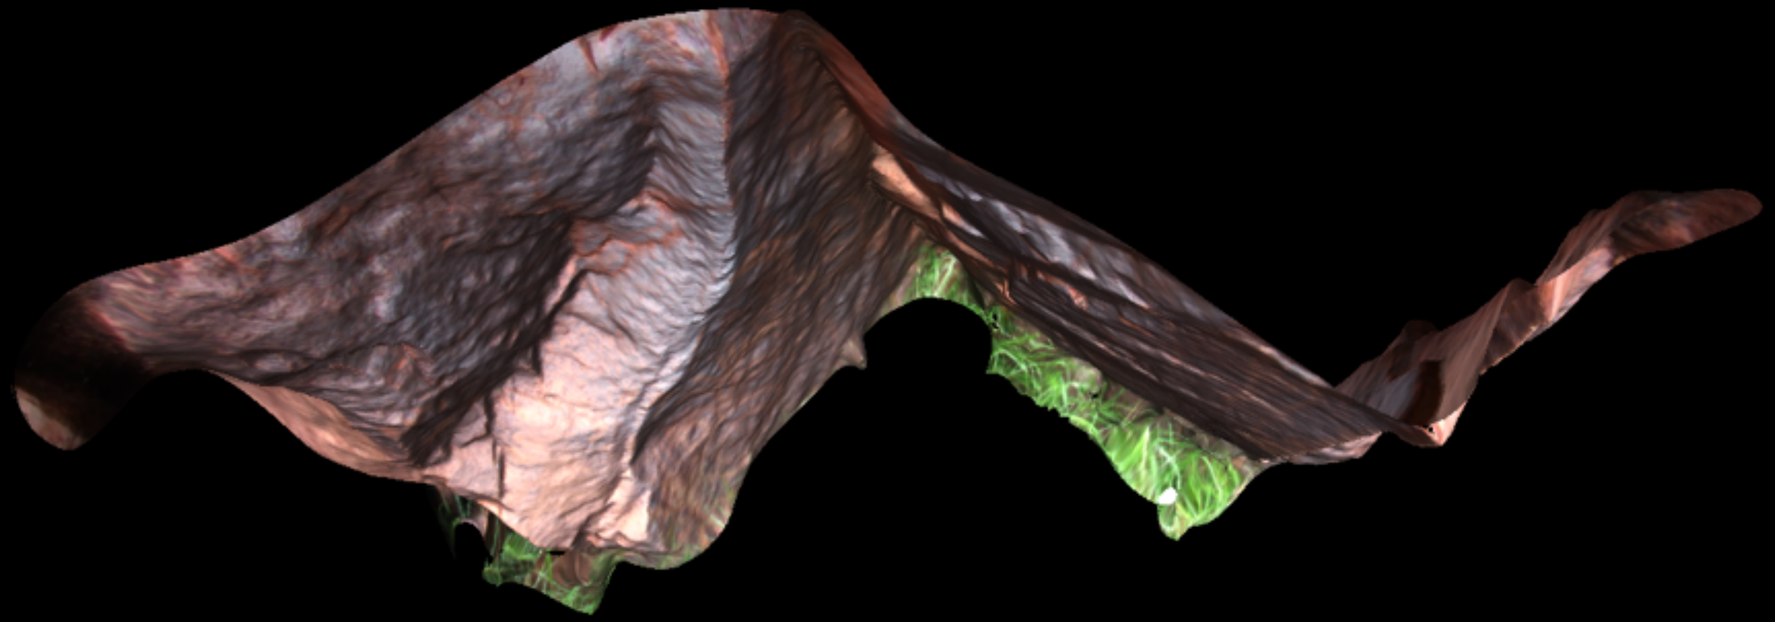

Infinitely thin surface
Unrealisable

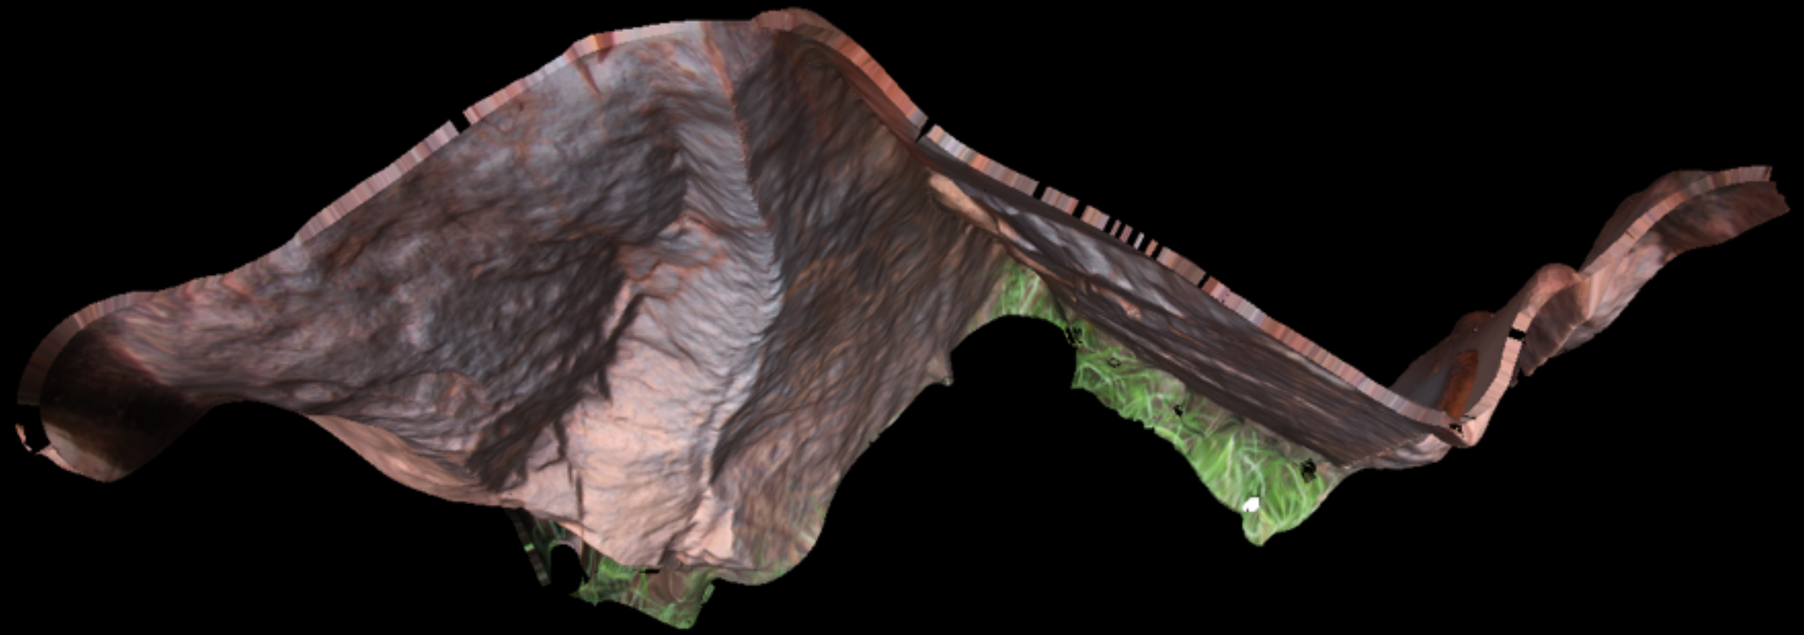

Thickened “realisable” surface

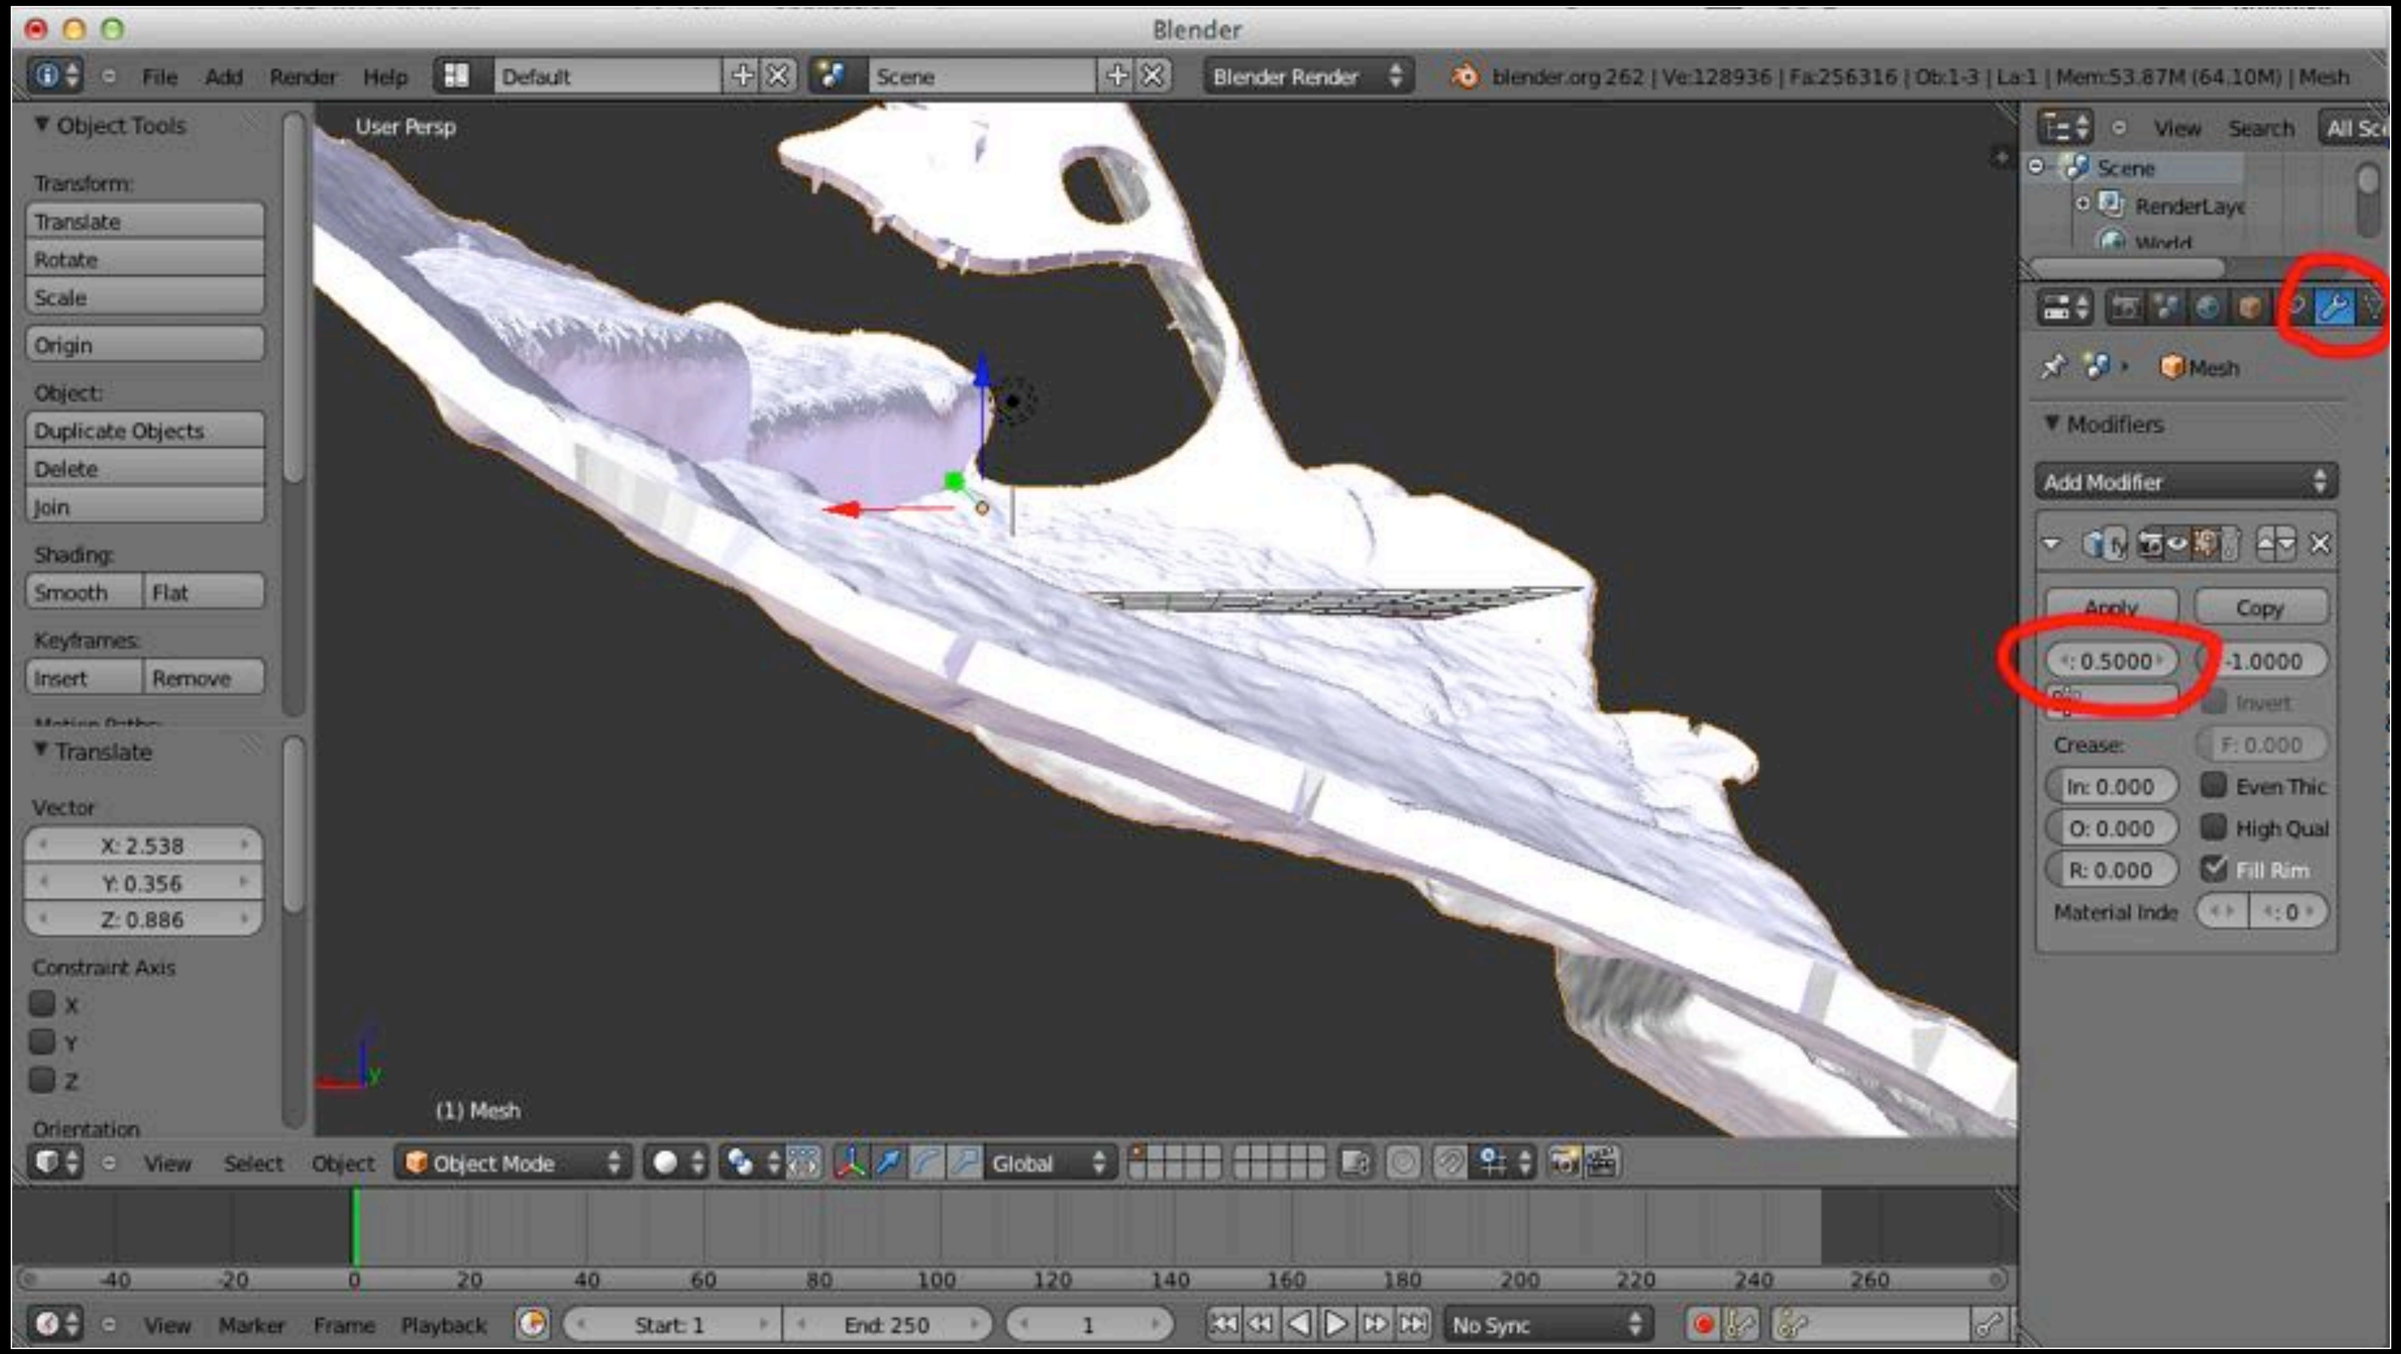
Blender
File Add Render Help
Default
Scene
Blender Render
blender.org 262 | Ve:128936 | Fa:256316 | Ob:1-3 | La:1 | Mem:53.87M (64.10M) | Mesh
Object Tools
Transform:
Translate
Rotate
Scale
Origin
Object:
Duplicate Objects
Delete
Join
Shading:
Smooth
Flat
Keyframes:
Insert
Remove
Translate
Vector
X: 2.538
Y: 0.356
Z: 0.886
Constraint Axis
X
Y
Z
Orientation
User Persp
(1) Mesh
View Search
Scene
Mesh
Modifiers
Add Modifier
Apply
Copy
: 0.5000
-1.0000
Crease:
In: 0.000
O: 0.000
R: 0.000
Even Thic
High Qual
Fill Rim
Material Inde
: 0
View Select Object
Object Mode
Global
-40 -20 0 20 40 60 80 100 120 140 160 180 200 220 240 260
View Marker Frame Playback
Start: 1
End: 250
1
No Sync


G07
Paul Bourke

Relighting

Post Production tools

- Smoothing.

  Z-brush

- Algorithmic. eg: mesh density reduction

  MeshLab

- Almost any 3D editing package you may be familiar with that can handle unstructured meshes.

  3DStudioMax - Maya - Cinema4D - Blender ….

Worked example

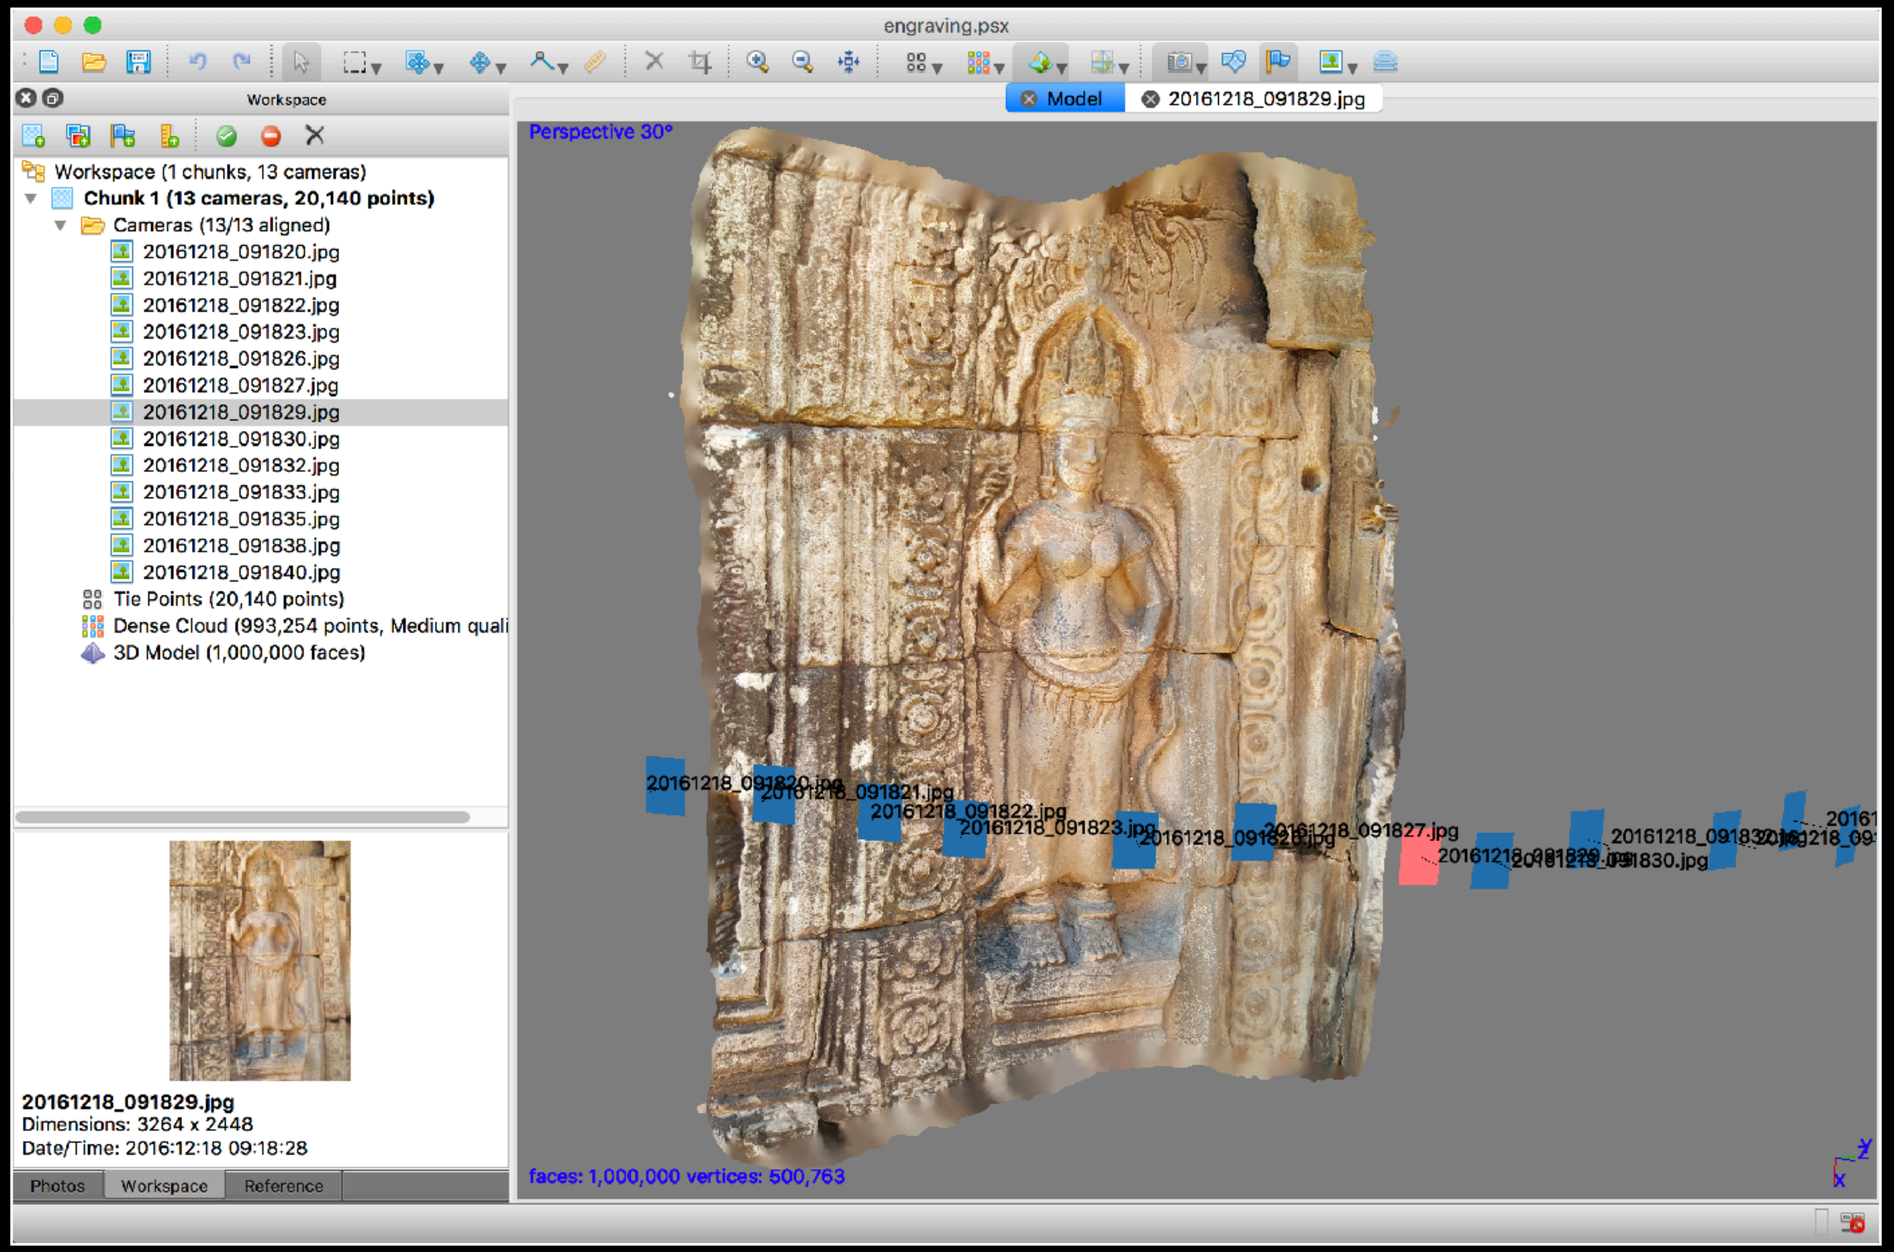

Your turn

- Free 30 day trial
  https://www.agisoft.com/downloads/installer/
- MeshLab (free)
  http://www.meshlab.net/
- This is about experimenting
  Start off with a modest number of photographs
  so compute time will be modest.