a pipeline for efficiently reconstructing heritage objects from range data

Kevin Cain, Todd Gill, Mark Eakle

end to end pipeline for processing archeological reconstructions

. evaluate the artifacts/sites under study

. plan your field work

. perform multi-resolution capture

. range alignment & merging

. surface simplification and normal extraction

. normal and color inpainting

. synthetic reconstruction

. rendering

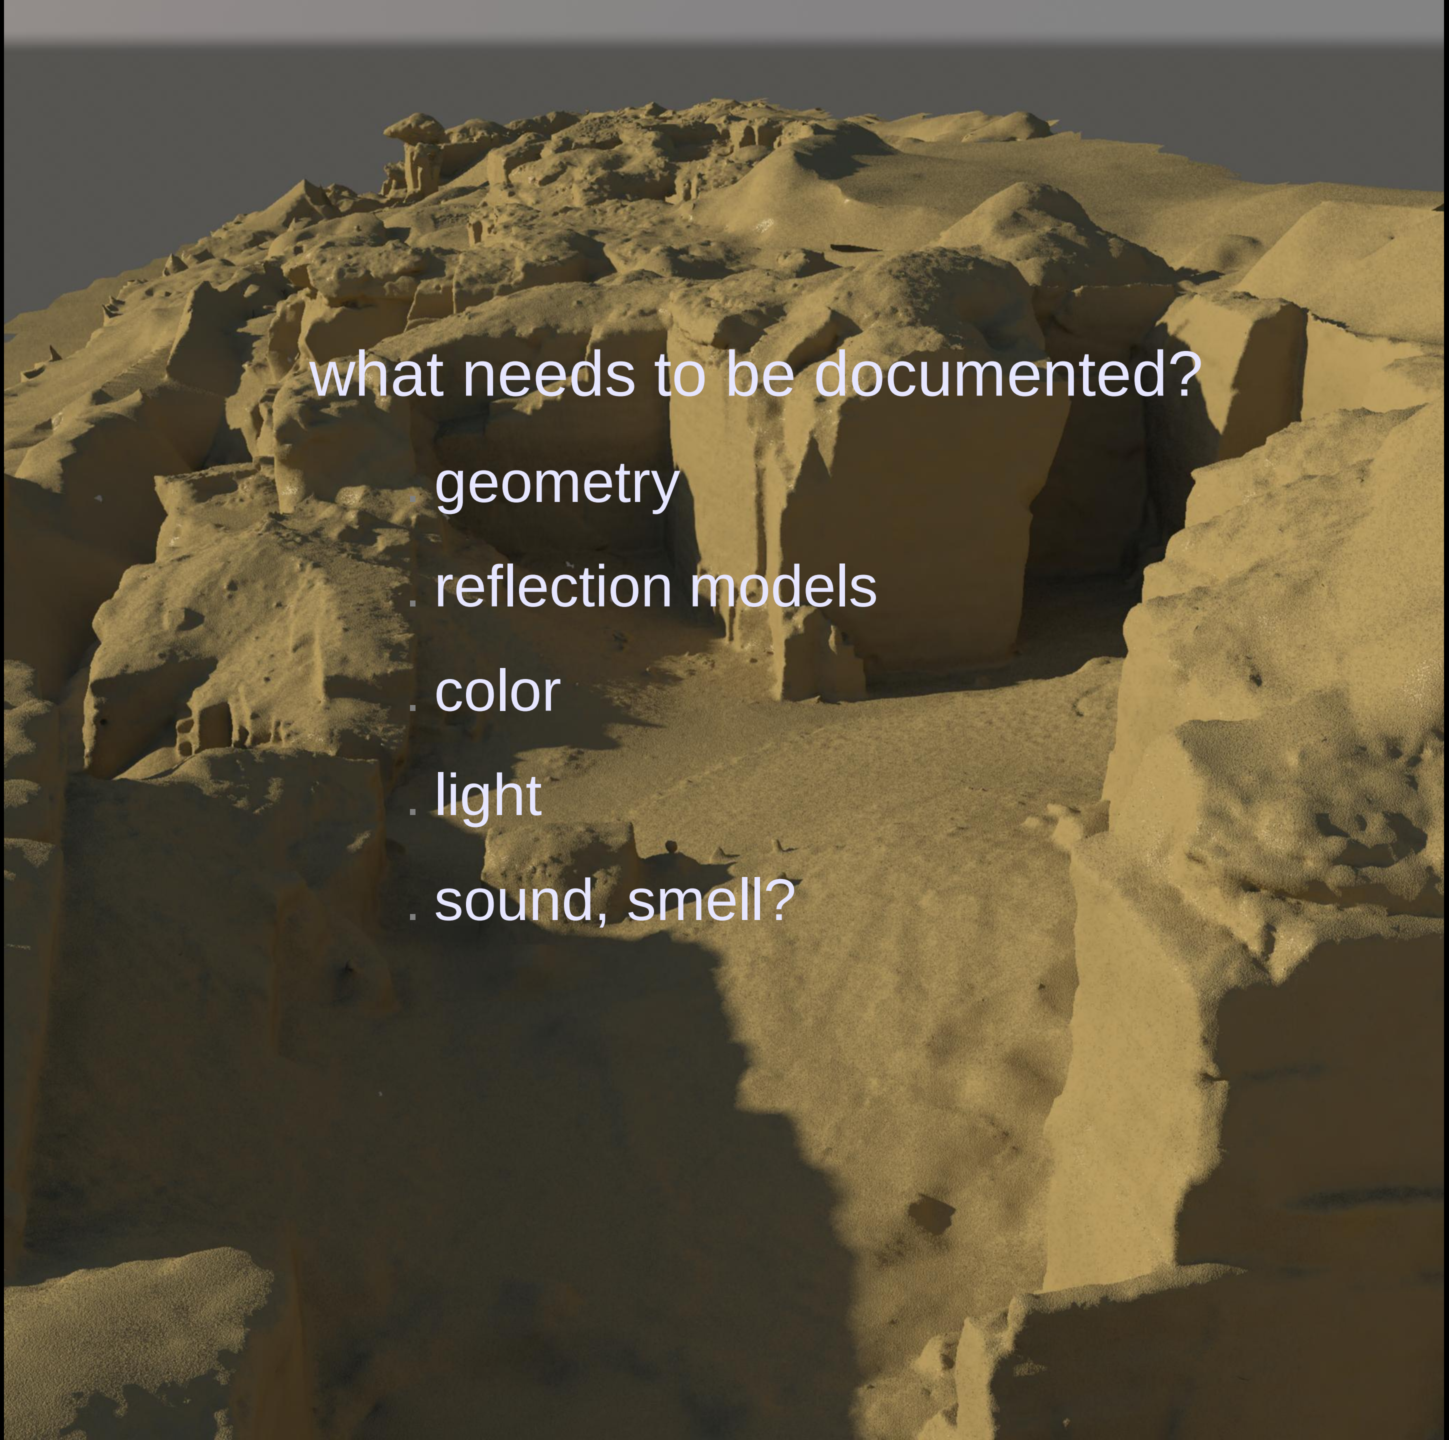

what needs to be documented?

- geometry
- reflection models
- color
- light
- sound, smell?

end to end pipeline for processing archeological reconstructions

. evaluate the artifacts/sites under study

. plan your field work

. perform multi-resolution capture

. range alignment & merging

. surface simplification and normal extraction

. normal and color inpainting

. synthetic reconstruction

. rendering

plan field acquisition

for heritage, inexpensive and flexible techniques often work best

dozens of range capture options exist:

. optical triangulation

. time of flight

. stereo imaging

. spatial carving

. photometric stereo

field work notes

. plan for extreme field conditions

. consider heat, cold, vibration, movement

. transportation and portability are critical

. manage data capture

. field backups are sometimes difficult, but not optional

end to end pipeline for processing archeological reconstructions

. evaluate the artifacts/sites under study

. plan your field work

. perform multi-resolution capture

. range alignment & merging

. surface simplification and normal extraction

. normal and color inpainting

. synthetic reconstruction

. rendering

why 'multi-resolution'?

heritage reasons: scenes often cross scale boundaries

+

technical reasons: high+low frequency data support one another

this may seem obvious, but often people think in terms of their equipment (and budget), rather than the needs of the scene

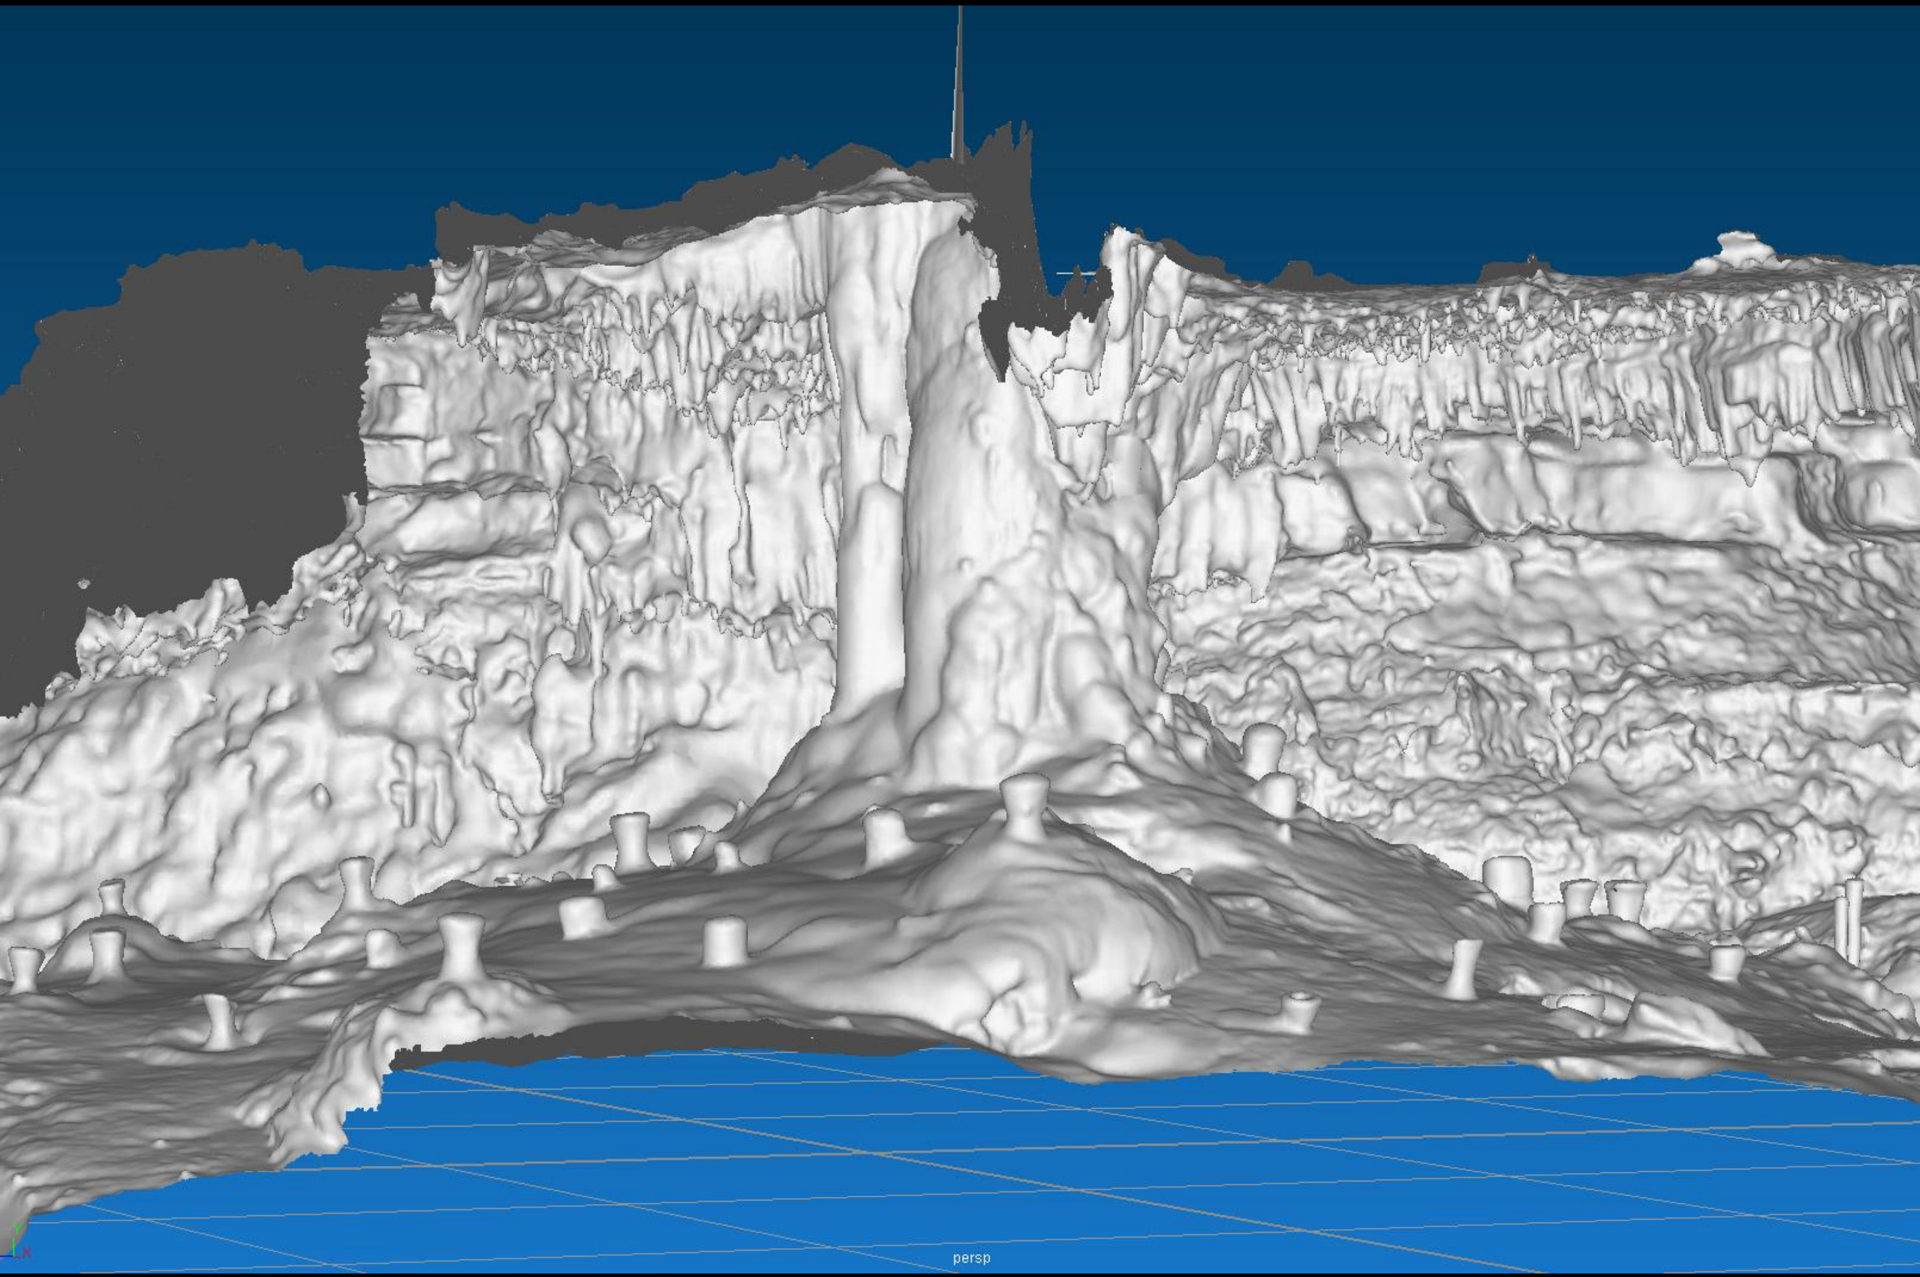
persp

photometric stereo,
one kind of high-frequency capture technique for CH

At pixel (i, j), want surface normal $\mathbf{n}_{ij}$
If Lambertian, observed image pixel is

$$I_{ij} = \alpha_{ij}\mathbf{n}_{ij} \cdot \mathbf{l}$$

Each new image $I^k_{..}$ has a new light position $\mathbf{l}^k$, providing many constraints on $\alpha_{ij}\mathbf{n}_{ij}$

[Note: for RGB, α is a 3-vector]

[Cain, Paterson 2007]

consider a pixel

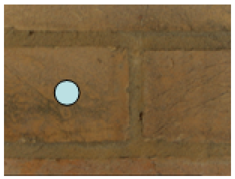

$$I^1_{ij} = \alpha_{ij}\mathbf{n}_{ij} \cdot \mathbf{l}^1$$

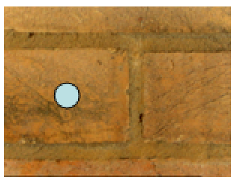

$$I^2_{ij} = \alpha_{ij}\mathbf{n}_{ij} \cdot \mathbf{l}^2$$

$\vdots$

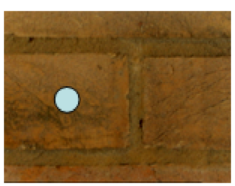

$$I^n_{ij} = \alpha_{ij}\mathbf{n}_{ij} \cdot \mathbf{l}^n$$

- One equation per pixel per image (greyscale)

$$\begin{pmatrix} I^1_{ij} \\ I^2_{ij} \\ \vdots \\ I^n_{ij} \end{pmatrix} = \begin{pmatrix} \mathbf{l}^{1\top} \\ \mathbf{l}^{2\top} \\ \vdots \\ \mathbf{l}^{n\top} \end{pmatrix} \alpha_{ij}\mathbf{n}_{ij}$$

$$\underset{n \times 1}{\mathbf{I}_{ij}} = \underset{n \times 3}{\mathbf{L}_{ij}} \; \underset{3 \times 1}{\alpha_{ij}\mathbf{n}_{\mathbf{ij}}}$$

So $\alpha_{ij}\mathbf{n}_{ij} = \mathbf{L}^{\dagger}_{ij}\mathbf{I}_{ij}$

[Cain, Paterson 2007]

[Cain, Paterson 2007]

[Cain, Paterson 2007]

[Cain, Paterson 2007]

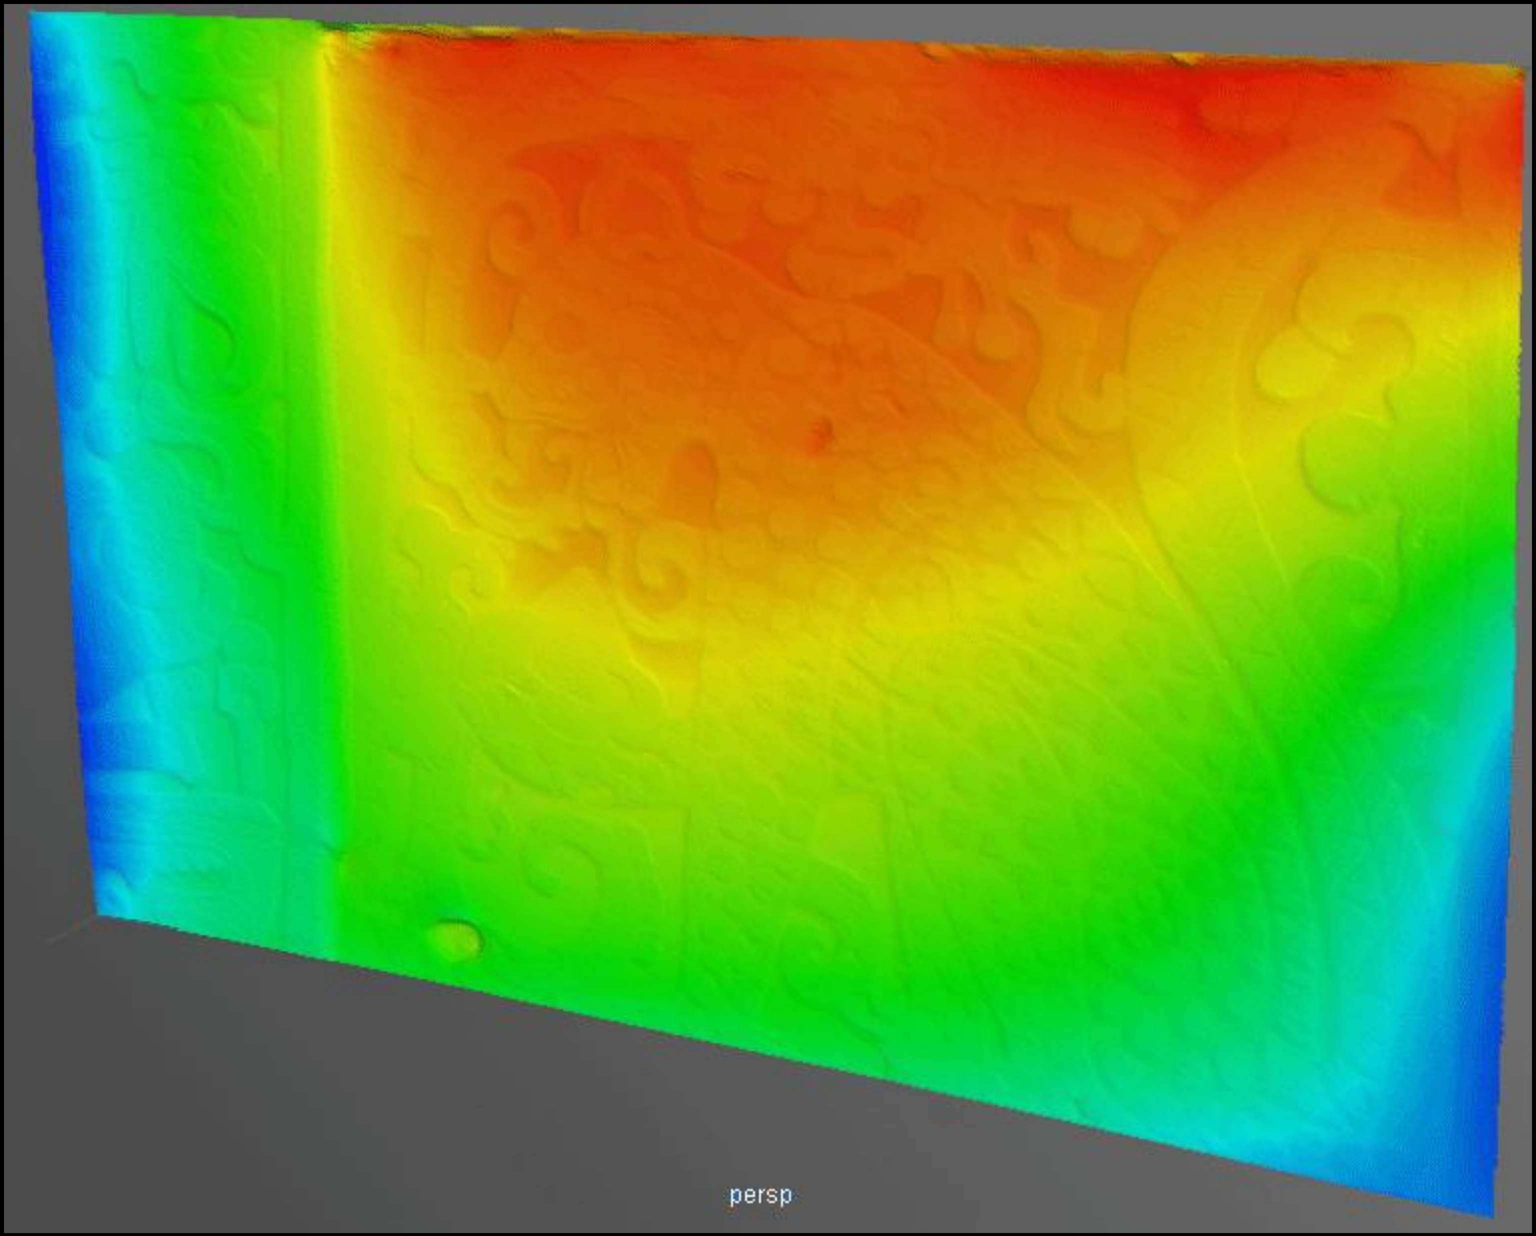

[Cain, Paterson 2007]

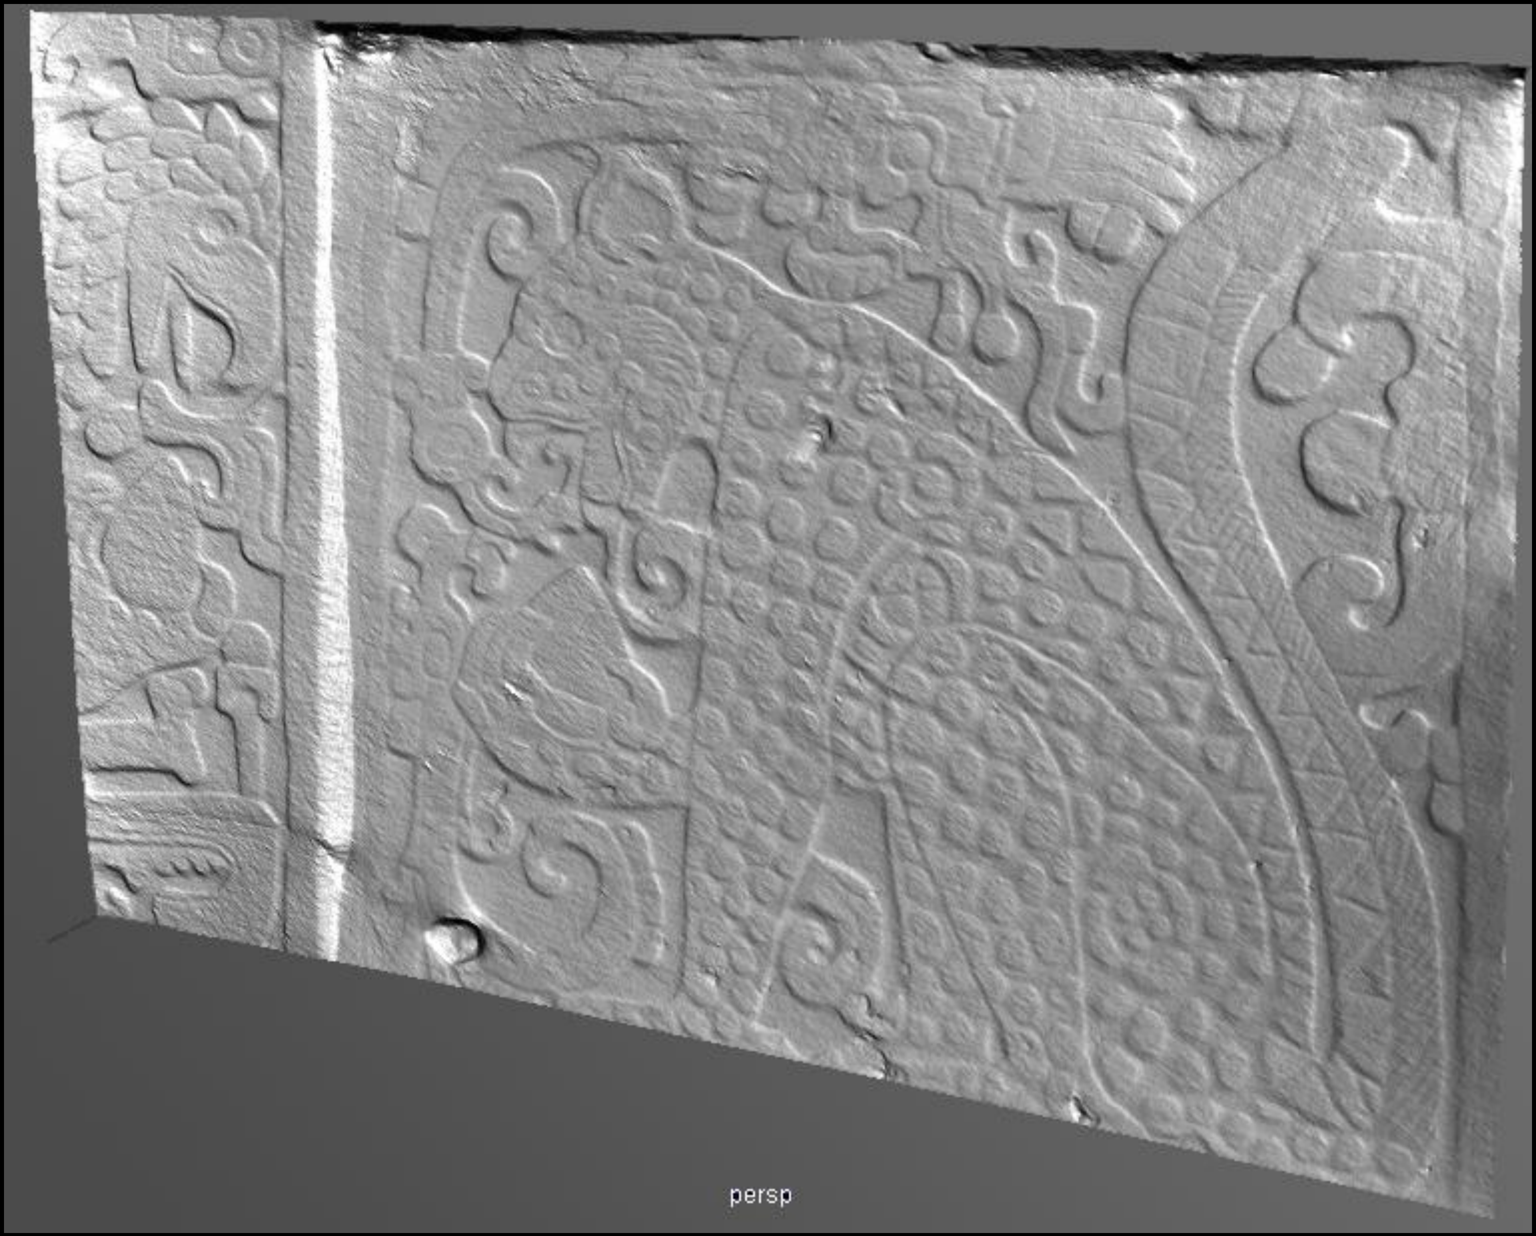

[Cain, Paterson 2007]

[Cain, Paterson 2007]

GLSL + Standard Textures

[Paterson 2009]

"ProcSet_03" (3082240 Vertices, Lamda = 0.000010) (Mesh BRDF view) (Diffuse Only)

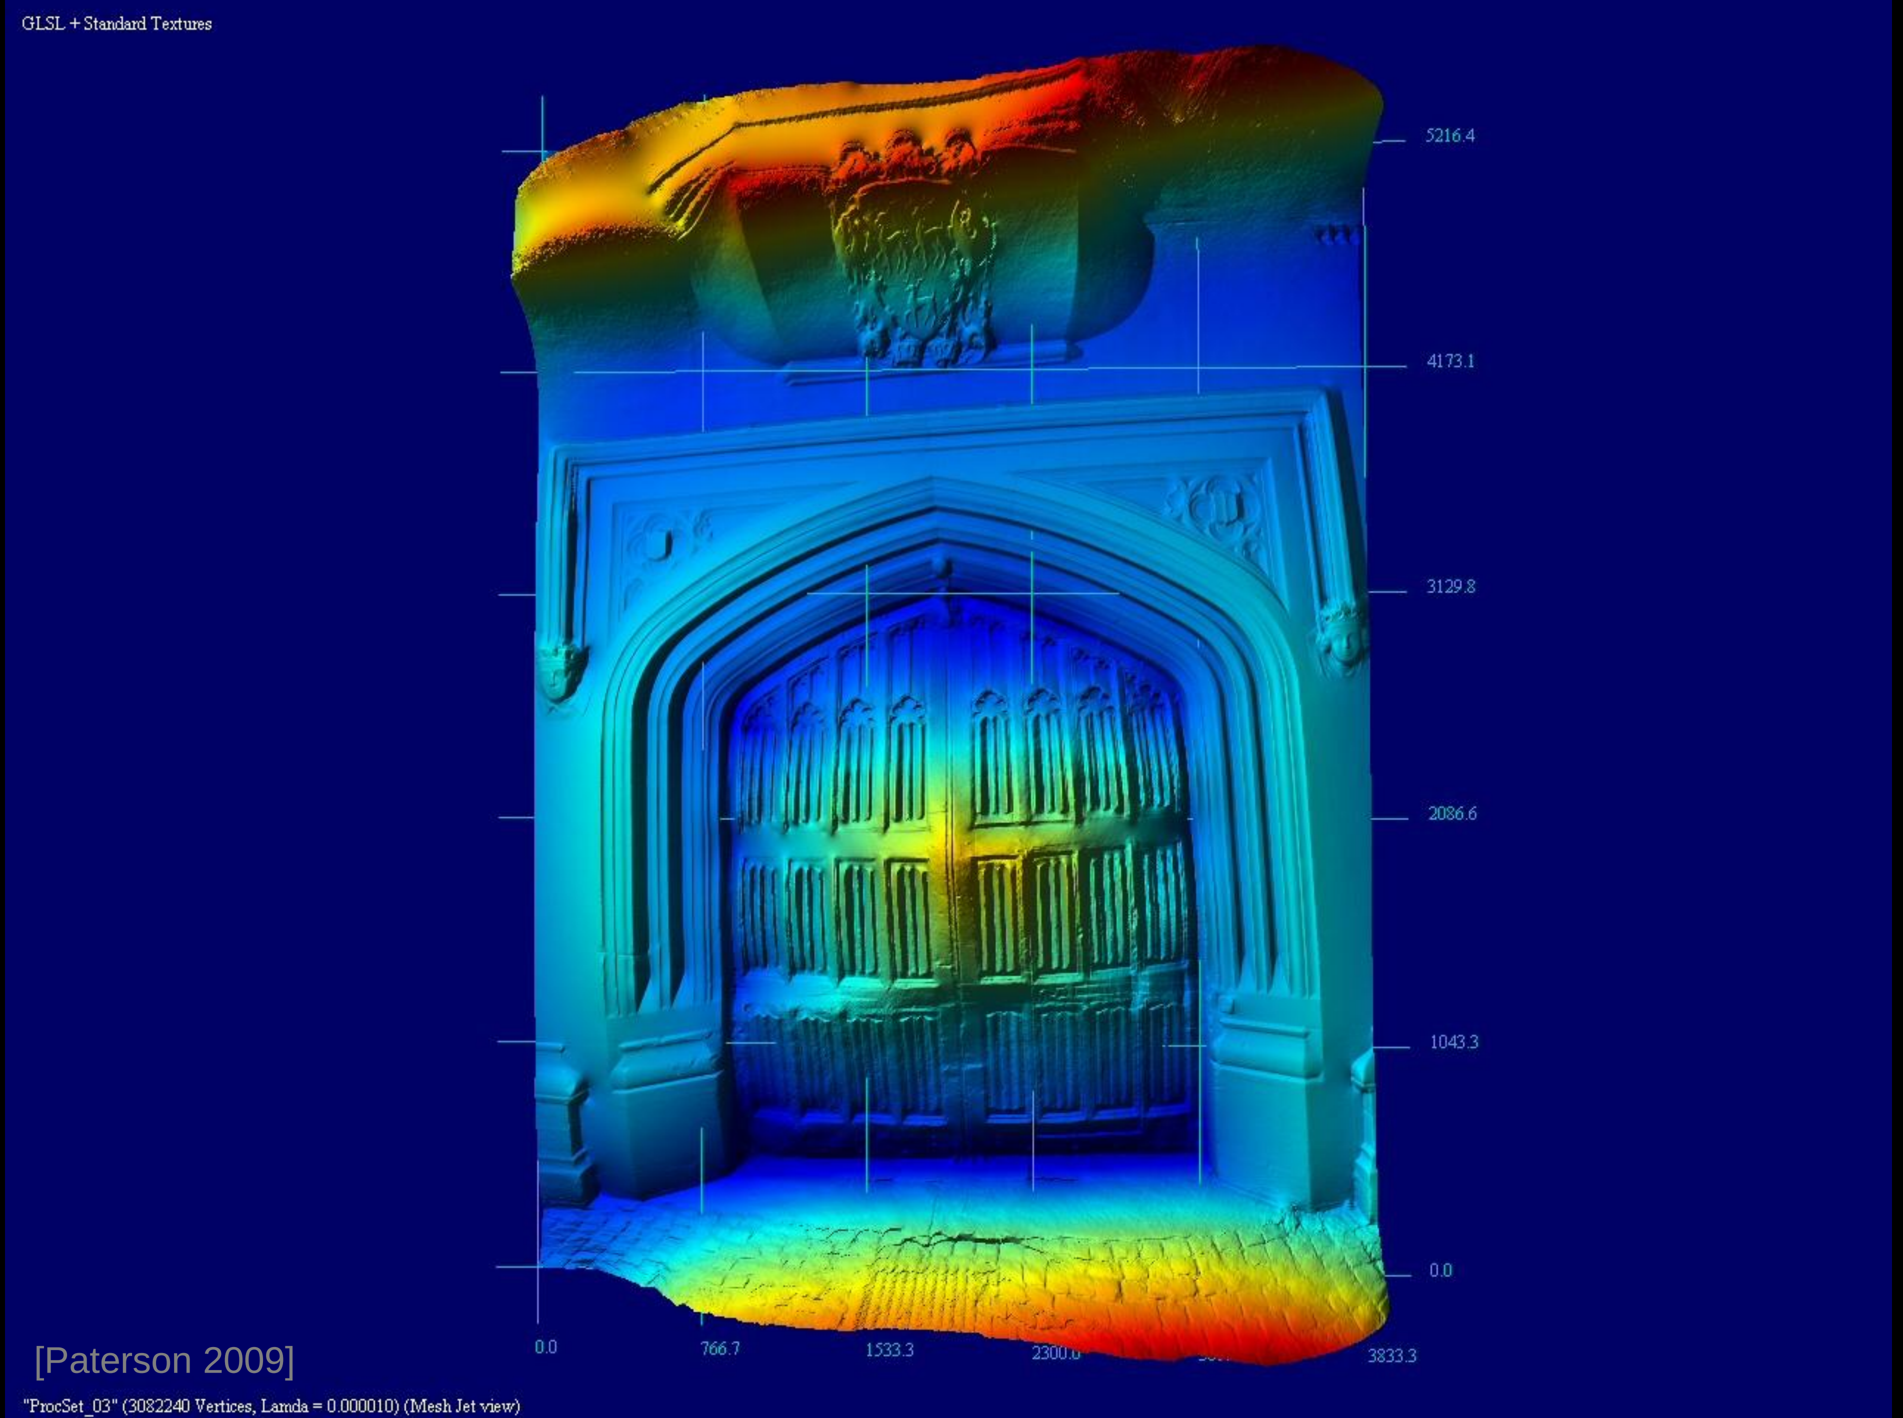

[Paterson 2009]

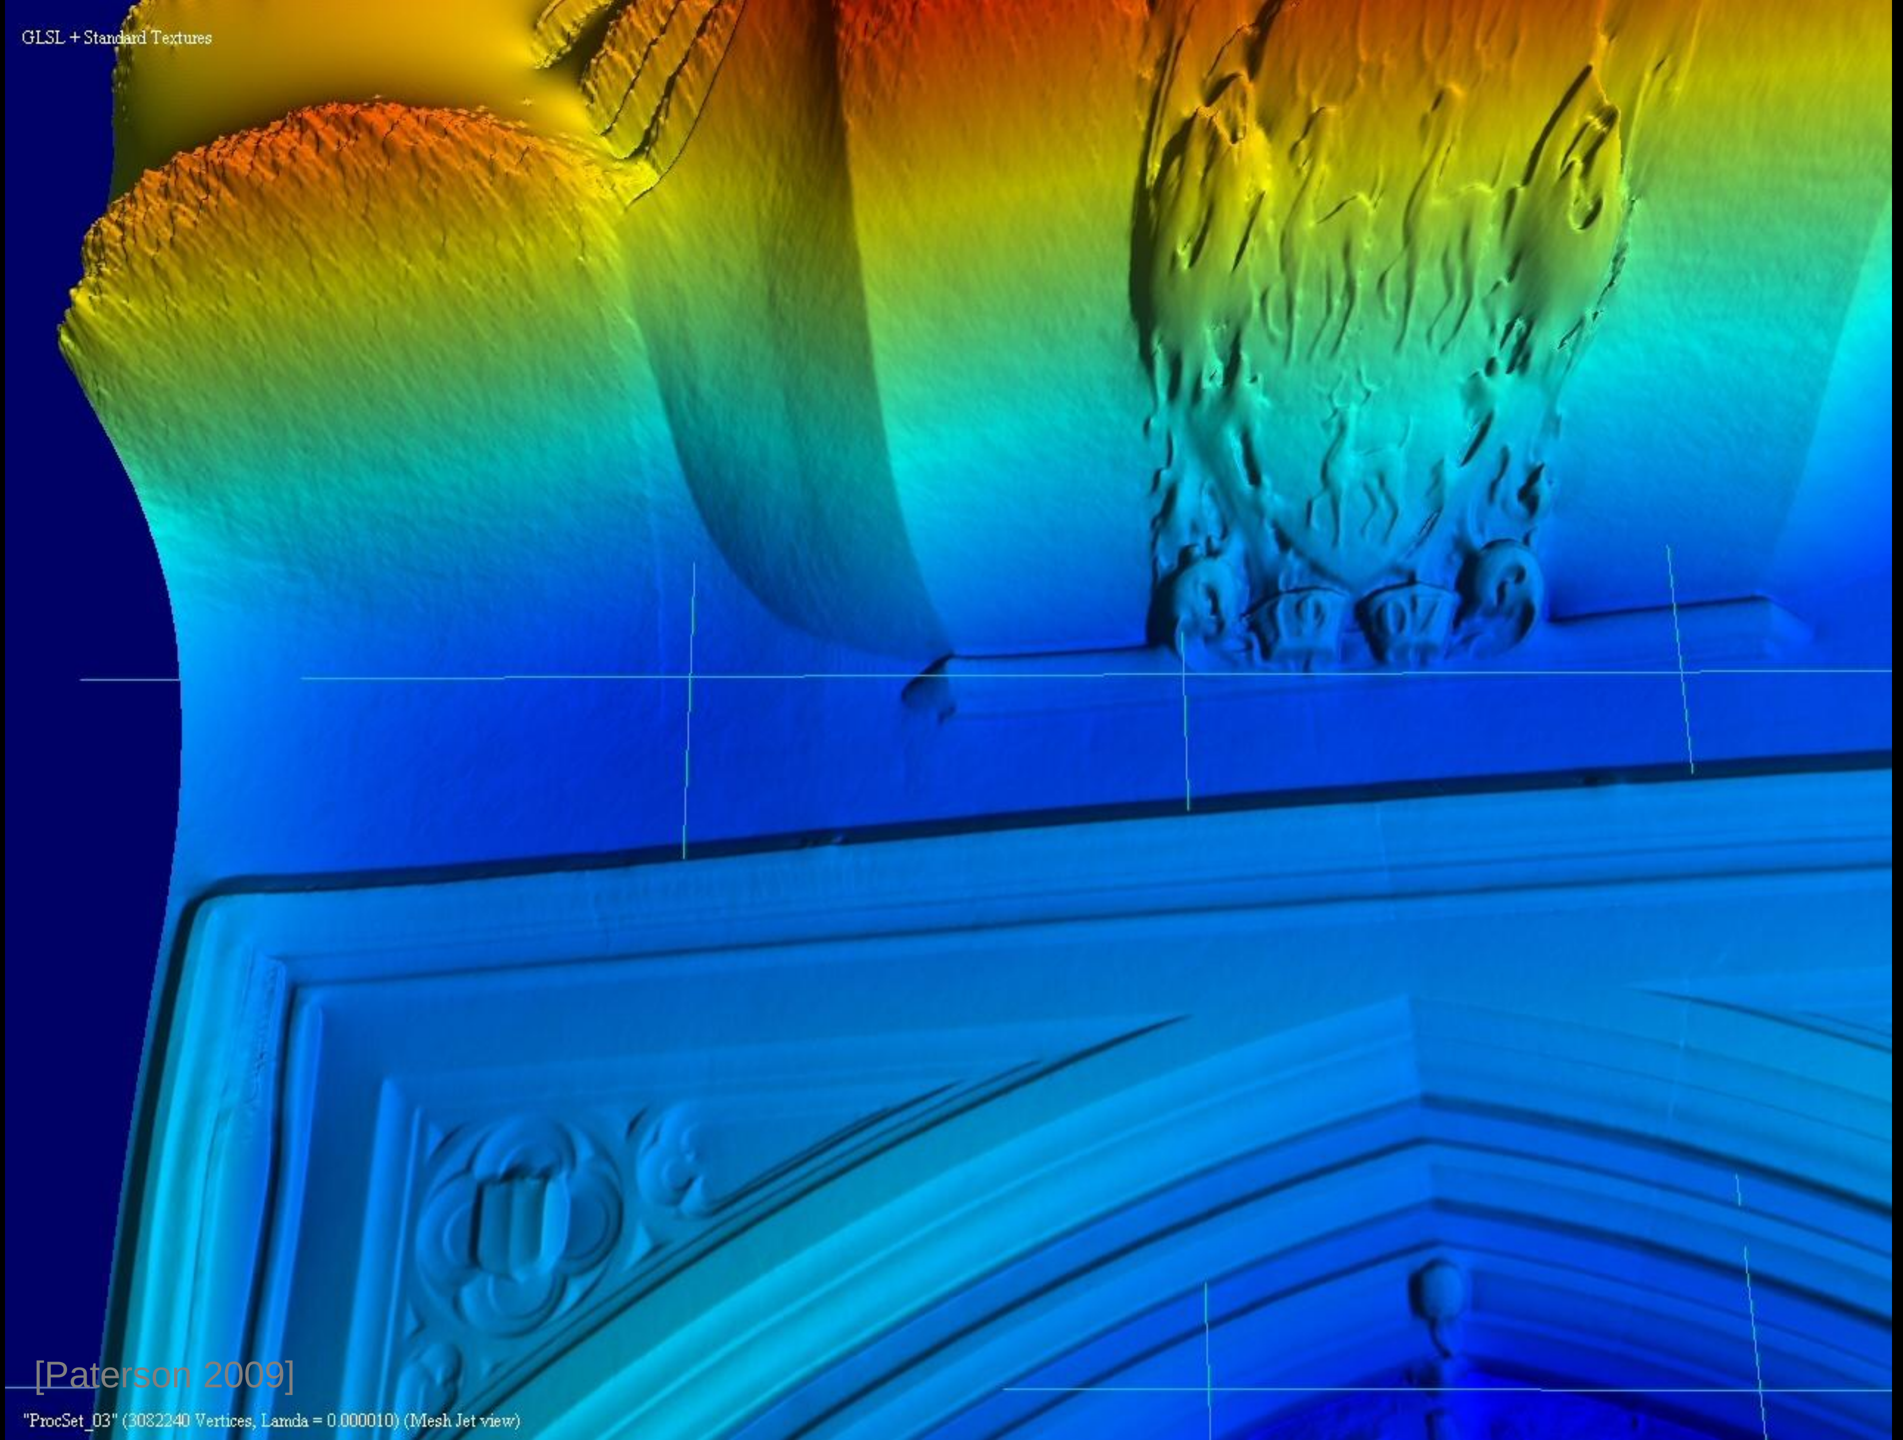

[Paterson 2009]

end to end pipeline for processing archeological reconstructions

. evaluate the artifacts/sites under study

. plan your field work

. perform multi-resolution capture

. range alignment & merging

. surface simplification and normal extraction

. normal and color inpainting

. synthetic reconstruction

. rendering

range data alignment, redux

- initial local alignment via ICP [Levoy, Curless 1998]
- global alignment [Brown 2006]

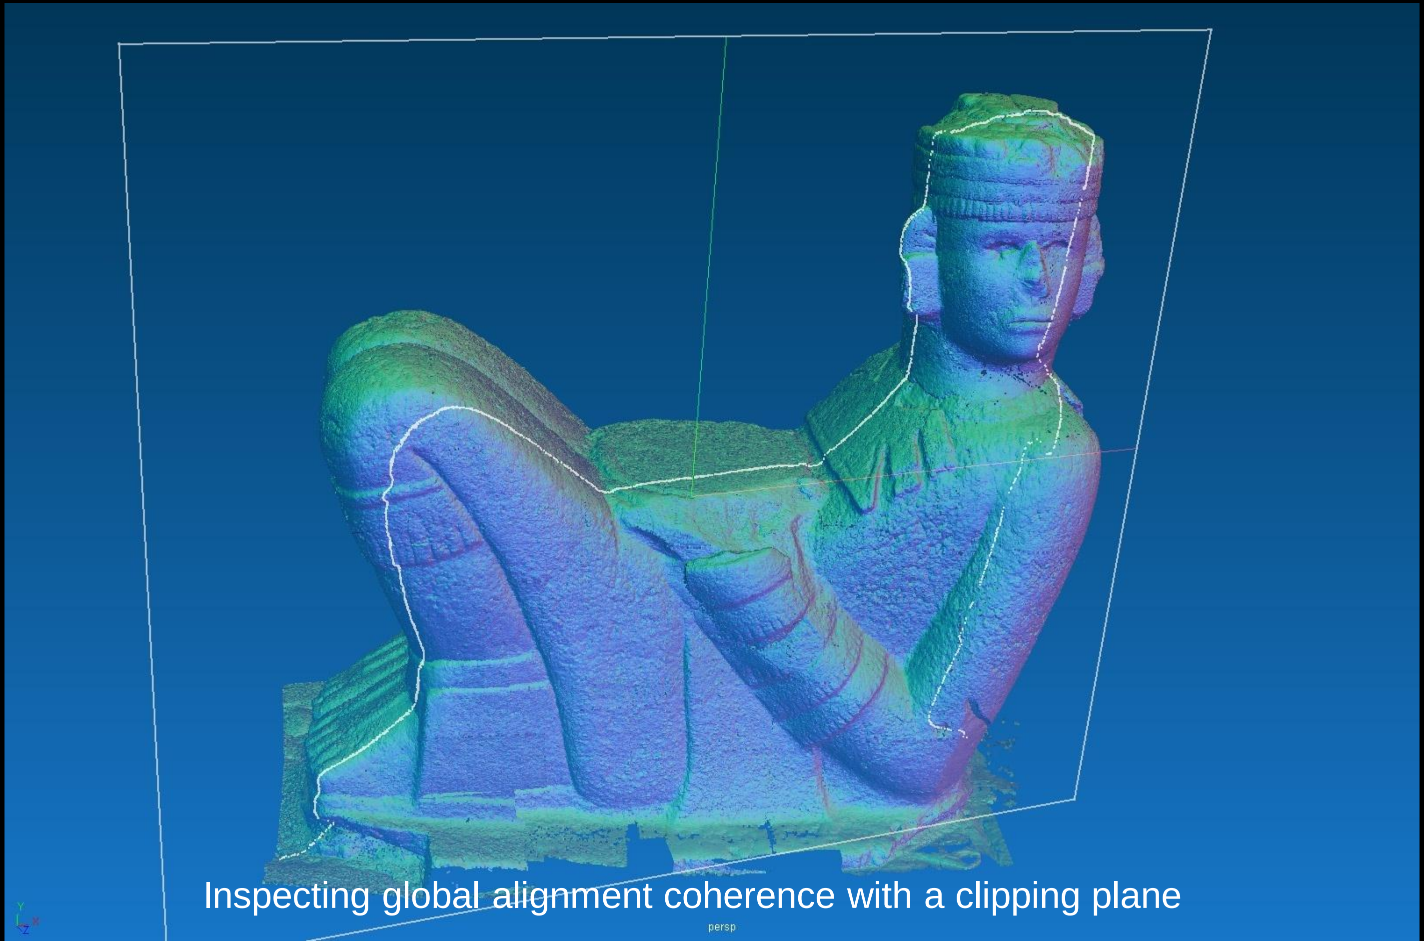

Inspecting global alignment coherence with a clipping plane

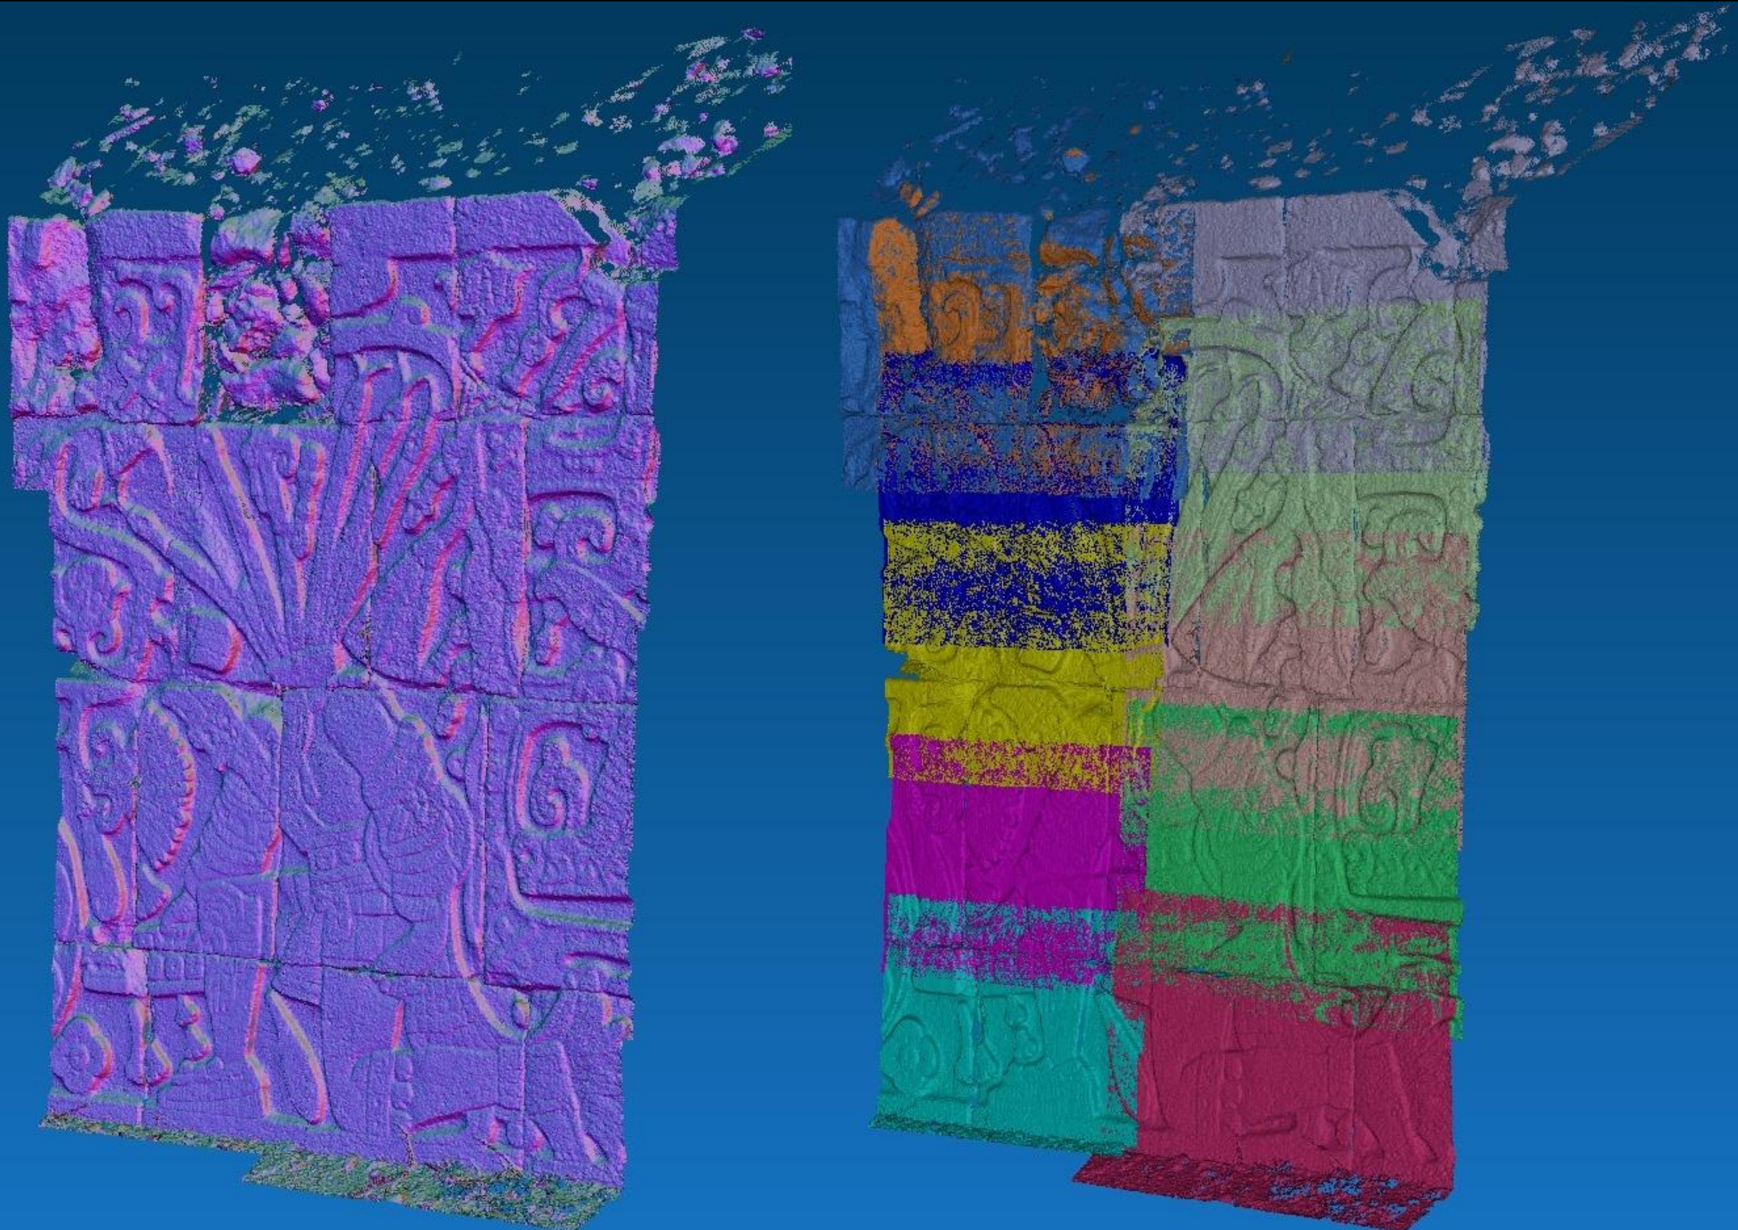

real world alignment caveats: avoiding local minima, open loops

Problem: minimizing pairwise correspondence yields a 'clamshell' opening where surfaces should be closed

- distribution of alignment error is large and local, not small and global
- akin to loop closure problem in SLAM

commonly, the solution is to globally relax error as Lanman, Taubin showed in their course Wednesday

real world alignment caveats: avoiding local minima, open loops

Problem: minimizing pairwise correspondence yields a 'clamshell' opening where surfaces should be closed

- . distribution of alignment error is large and local, not small and global
- . akin to loop closure problem in SLAM

commonly, the solution is to globally relax error as Lanman, Taubin showed in their course Wednesday

we propose: compute ICP, then *maximize* pair error

- . also, it is useful to iterate with top 20% of pairs, first, then expand incrementally up to 50% of pairs

point rendering – G. Drettakis

alignment – K. Cain

designing for multi-resolution models

- plan to insert detailed models within a simplified scaffold
  - pay penalties for large models only where needed
- where possible, extract normals from complex geometry

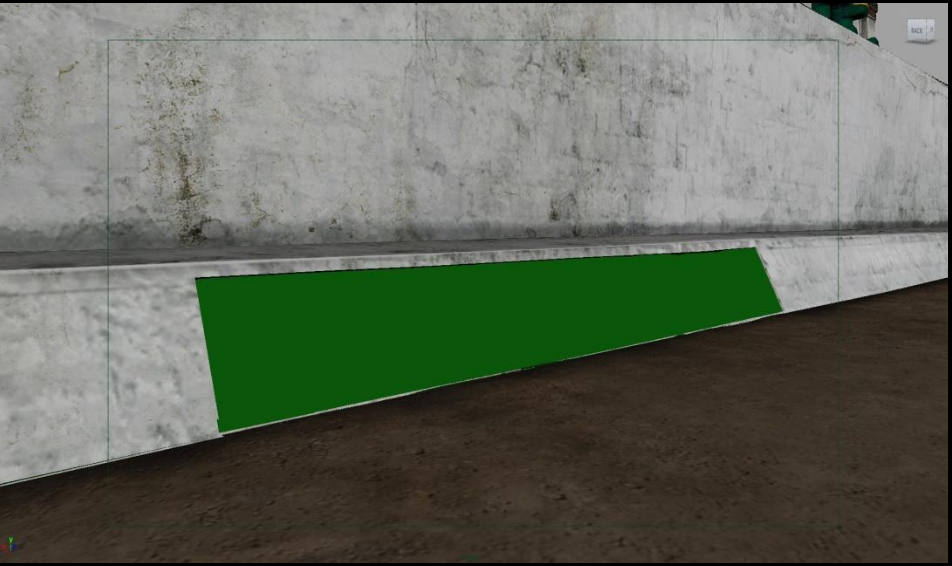

constraining high-frequency data to a lower-frequency scaffold

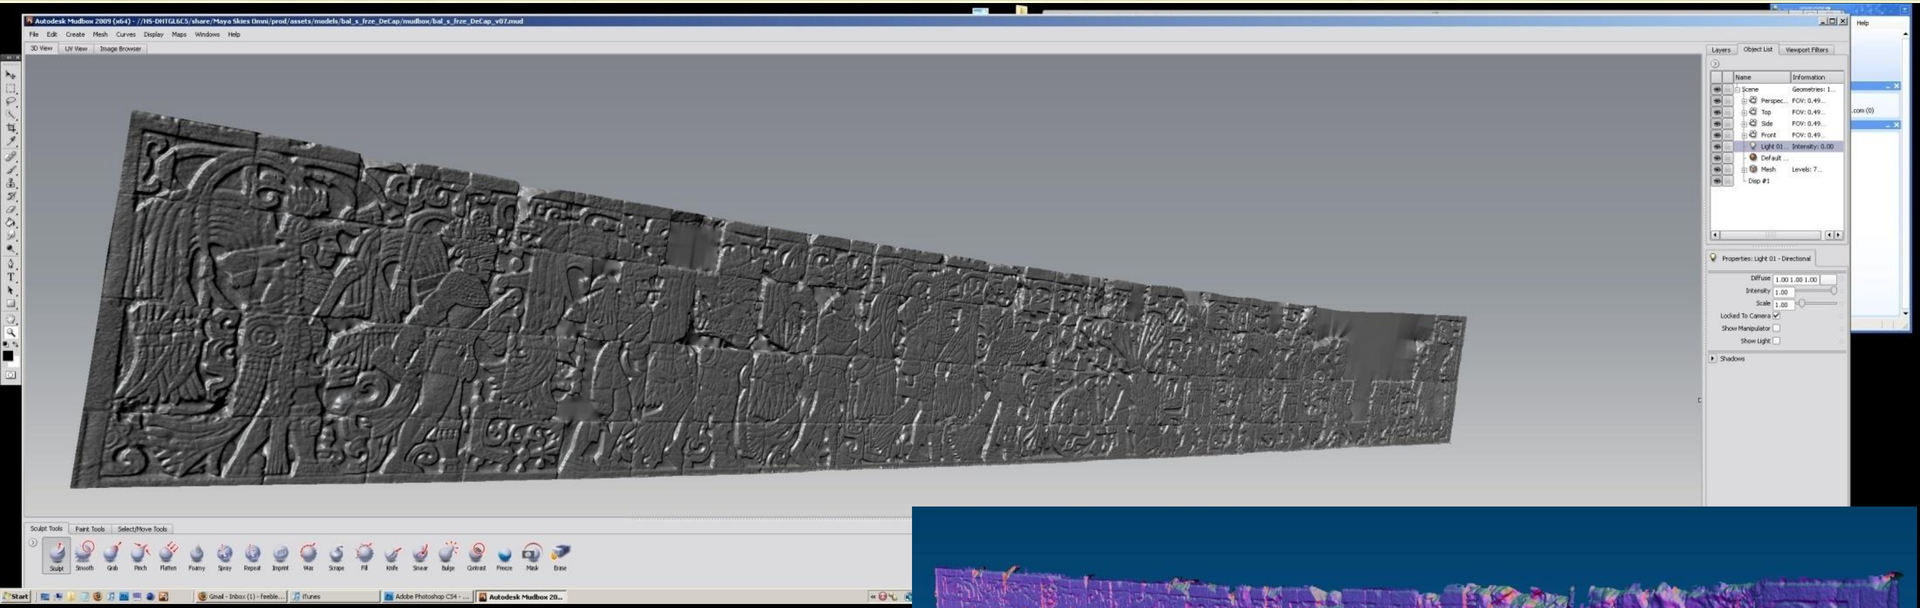

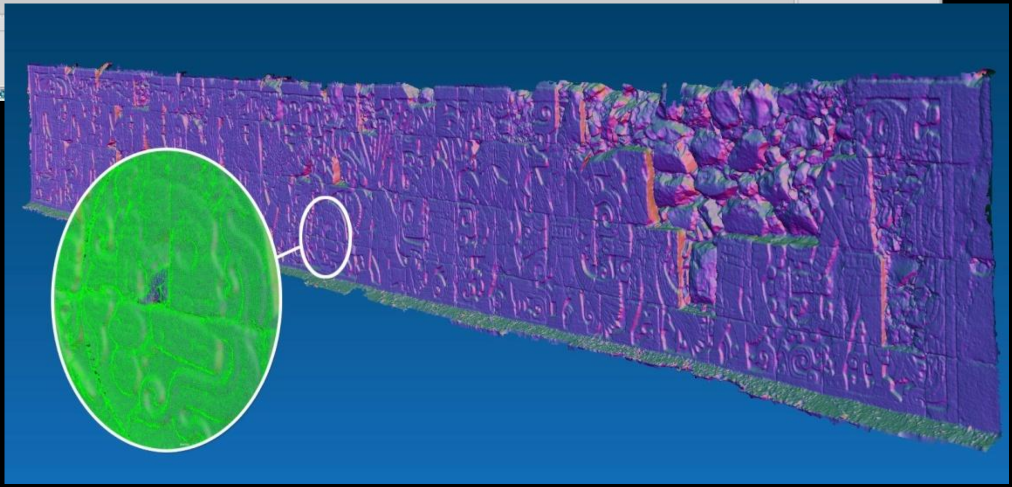

accumulated alignment error for high-frequency data benefits from a control grid that is accurate over comparatively large regions of the scene under study

multiscale data alignment

it's easy to get an order of magnitude differences in the sampling density from different data sources:

- . 7,000 x 5,000 normal map overlaid on a 640x480 range grid
- . ~.5mm optical triangulation mesh aligned to a ~5mm LIDAR mesh
- . .03mm RGB from photos aligned to a ~1mm triangulation mesh

unfortunately, most alignment approaches break when delta in input sampling distribution > an order of magnitude

ICP registration [Besl, McKay 1992]

problem:

arbitrary subdivision tends to discourage helpful movement during multiscale alignment

ICP registration [Besl, McKay 1992]

problem:

arbitrary subdivision tends to discourage helpful movement during multiscale alignment

some solutions:

. constrain alignment with sparse correspondences

. weight pair participation in the alignment via gauss curvature at coarse scale

or, allow user to paint feature regions which will participate in iterating global alignment


does this also work with image-based techniques of the kind we have seen today?

a qualified yes: accurate scale must be established by fiducials, or absolute reference

. photometric embroidery [Debevec 2007, 2009]

. sparse grid of triangulation control points as constraints [Horovitz, Kiryati 2004]

. GPS assist to image-based modeling [Seitz 2009]

end to end pipeline for processing archeological reconstructions

. evaluate the artifacts/sites under study

. plan your field work

. perform multi-resolution capture

. range alignment & merging

. surface simplification and normal extraction

. normal and color inpainting

. synthetic reconstruction

. rendering

after point/mesh alignment

- clean up
- isosurface extraction
- decimation
- control mesh creation based on scan mesh
- organized texture atlas assignment
- displacement map transfer to control mesh from scan mesh
- create sculpt mesh by baking displacement on hi-rez subdivided control mesh
- digital sculpting and reconstruction on sculpt mesh via input from archeological review and inference
- output modified base sculpt mesh (original control mesh)
- evaluate UV stretching from geometry modifications - re-UV base sculpt mesh making final base mesh
- displacement map transfer to base final mesh from original hi-rez sculpt mesh1
- create final mesh by baking displacement on hi-rez subdivided base final mesh
- normal map transfer to lower rez levels of final mesh (level0, level1, etc for different camera needs)
- texture maps can be created from painting directly on hi-rez final mesh in a 3D paint package
- texture maps can also be creatred from working with ambient occlusion, normal and UV snapshot in 2D photo editor, viewable in 3D package

minimal texture map output should be diffuse, specular, bump(fine tune control) and normal(gross control)

after point/mesh alignment

- clean up
- isosurface extraction
- decimation
- **control mesh creation based on scan mesh**
- organized texture atlas assignment
- displacement map transfer to control mesh from scan mesh
- create sculpt mesh by baking displacement on hi-rez subdivided control mesh
- **digital sculpting and reconstruction on sculpt mesh via input from archeological review and inference**
- output modified base sculpt mesh (original control mesh)
- evaluate UV stretching from geometry modifications - re-UV base sculpt mesh making final base mesh
- displacement map transfer to base final mesh from original hi-rez sculpt mesh1
- create final mesh by baking displacement on hi-rez subdivided base final mesh
- normal map transfer to lower rez levels of final mesh (level0, level1, etc for different camera needs)
- texture maps can be created from painting directly on hi-rez final mesh in a 3D paint package
- texture maps can also be creatred from working with ambient occlusion, normal and UV snapshot in 2D photo editor, viewable in 3D package

minimal texture map output should be diffuse, specular, bump(fine tune control) and normal(gross control)

end to end pipeline for processing archeological reconstructions

. evaluate the artifacts/sites under study

. plan your field work

. perform multi-resolution capture

. range alignment & merging

. surface simplification and normal extraction

. normal and color inpainting

. synthetic reconstruction

. rendering

isosurface extraction

an isosurface is a polygonal approximation of an implicit function associated with a voxel volume

'Poisson Reconstruction' [Kazhdan, Bolitho, Hoppe 2006]

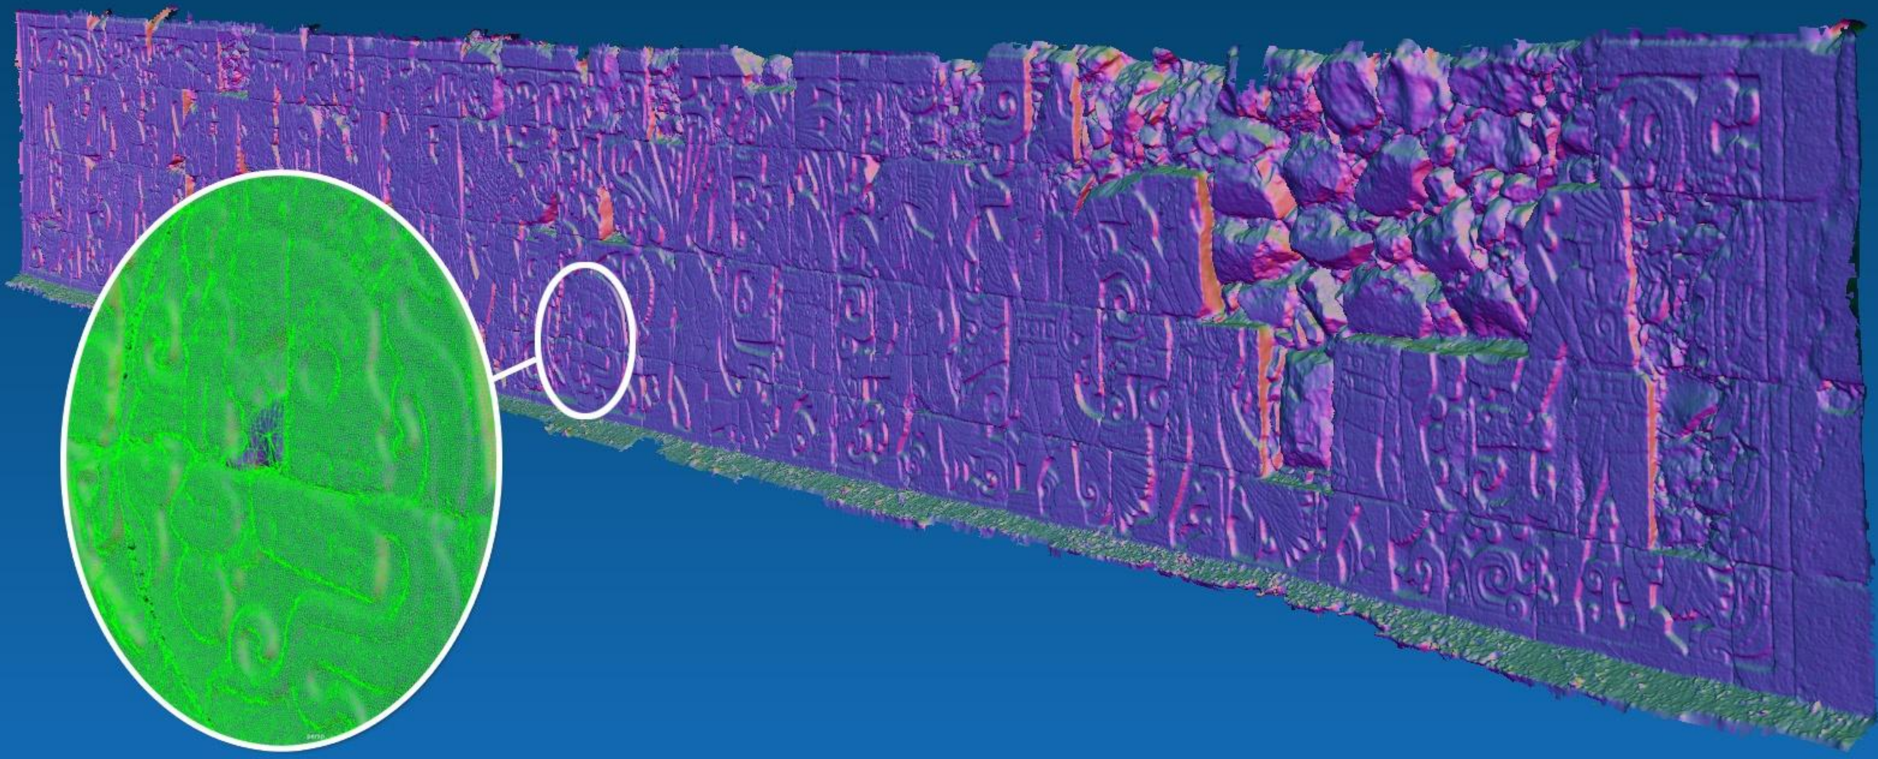

archaeological intervention

the Venice Charter asks us to limit interpolation / editing of data
i.e. Laplacian smoothing, hole filling

- compute fill geometry via tangency along hole boundaries
- volumetric diffusion (Davis, Marschner, Garr, Levoy 2002)

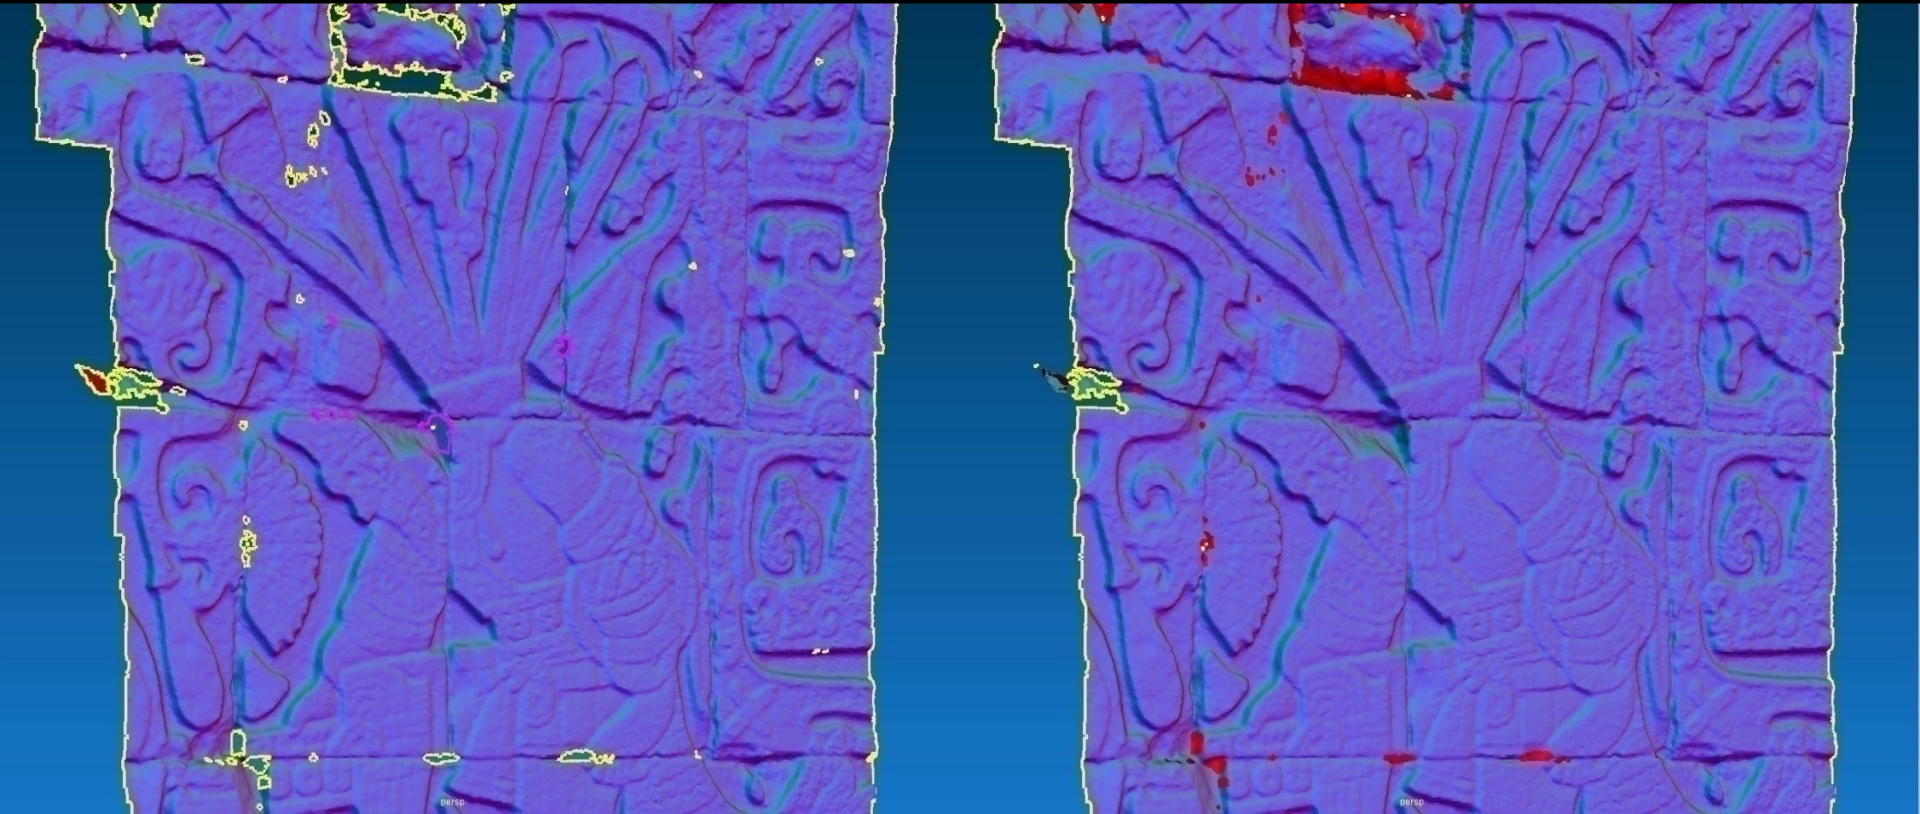

mesh simplification

adaptive decimation (below, left)

. preserve triangle count for areas with high local curvature

. aggressively decimate comparatively flat regions

mesh culling using camera normal criteria (below, right)

. data collected at oblique angles is generally less usable

. dot product with camera, edge length

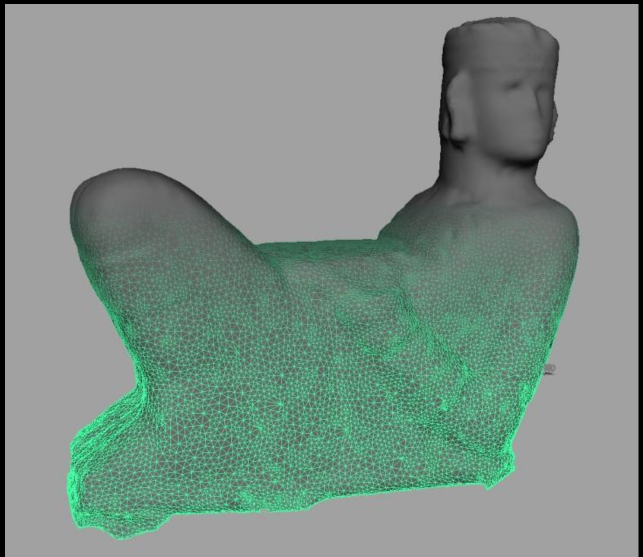

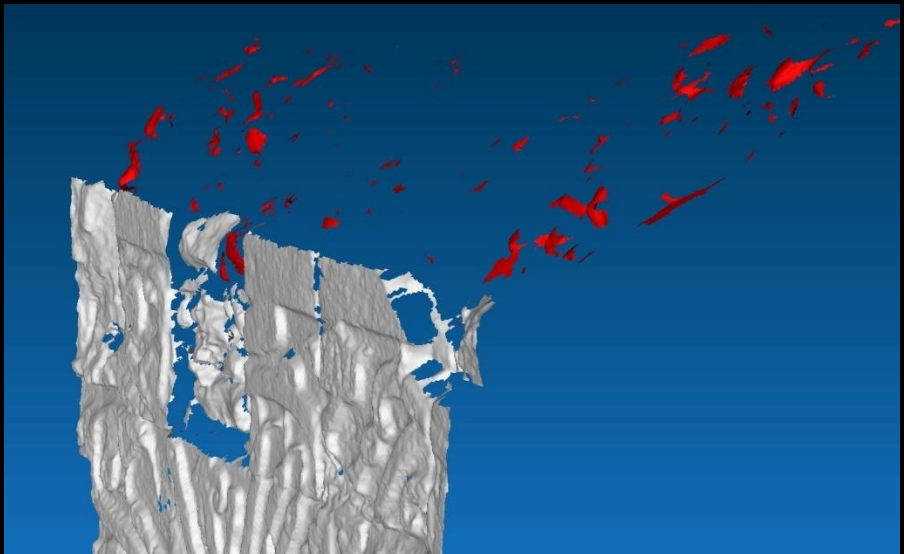

transfer surface normals to proxy cmesh

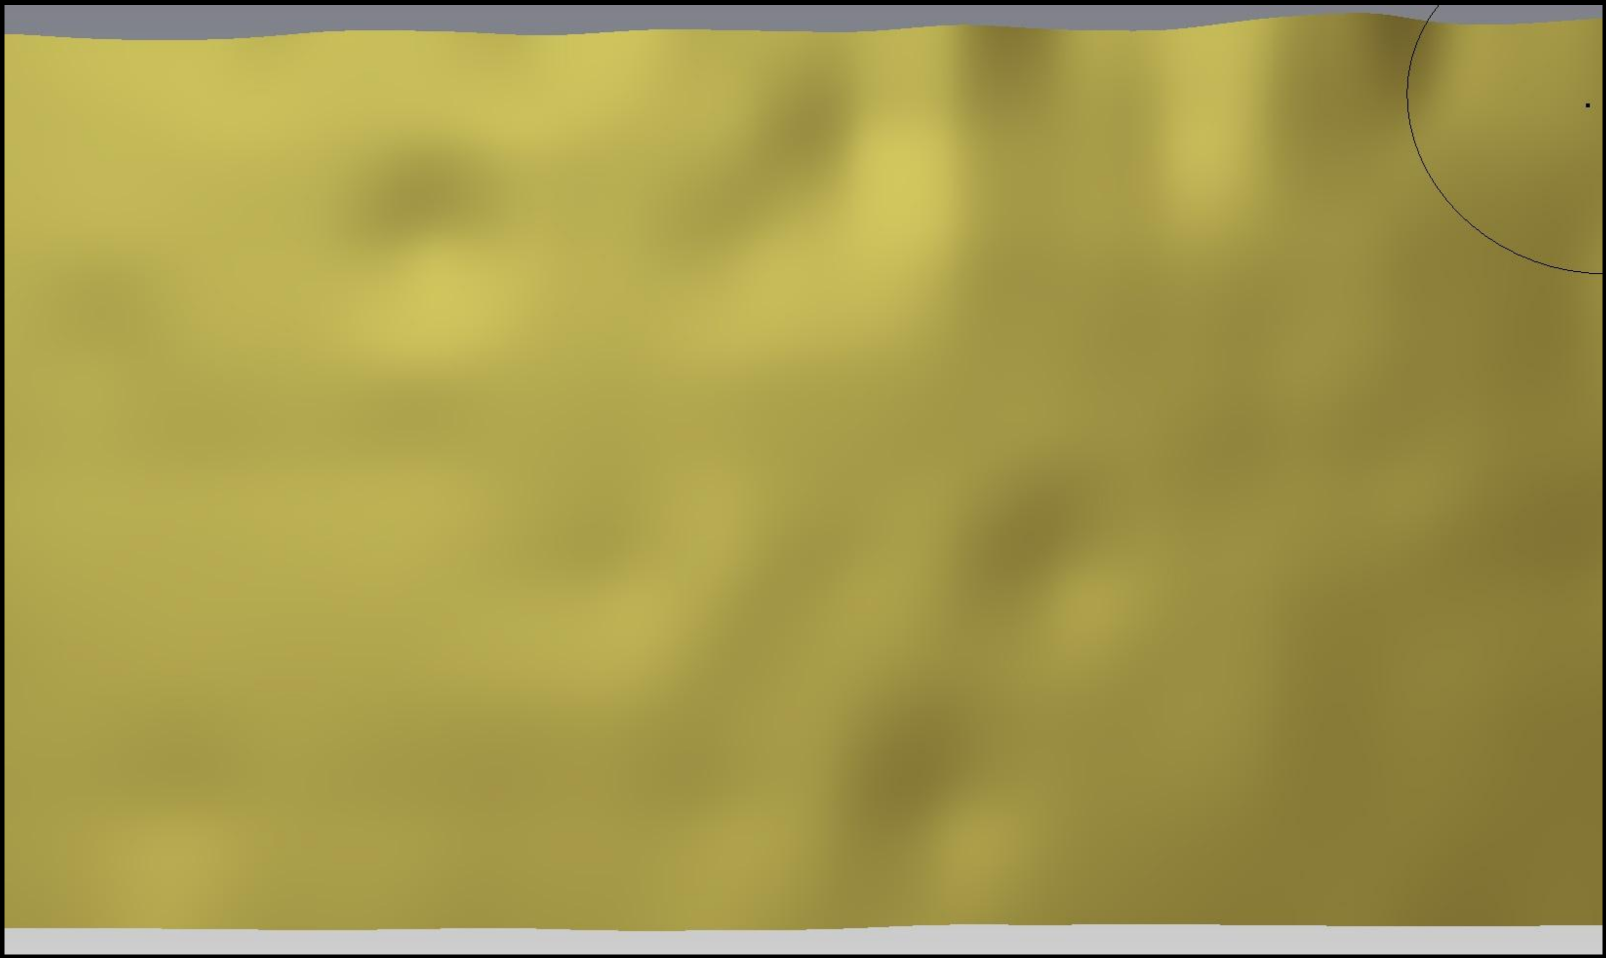

base mesh + extracted normals

base mesh + normal map inpainting

base mesh + normal map inpainting

end to end pipeline for processing archeological reconstructions

. evaluate the artifacts/sites under study

. plan your field work

. perform multi-resolution capture

. range alignment & merging

. surface simplification and normal extraction

. normal and color inpainting

. synthetic reconstruction

. rendering

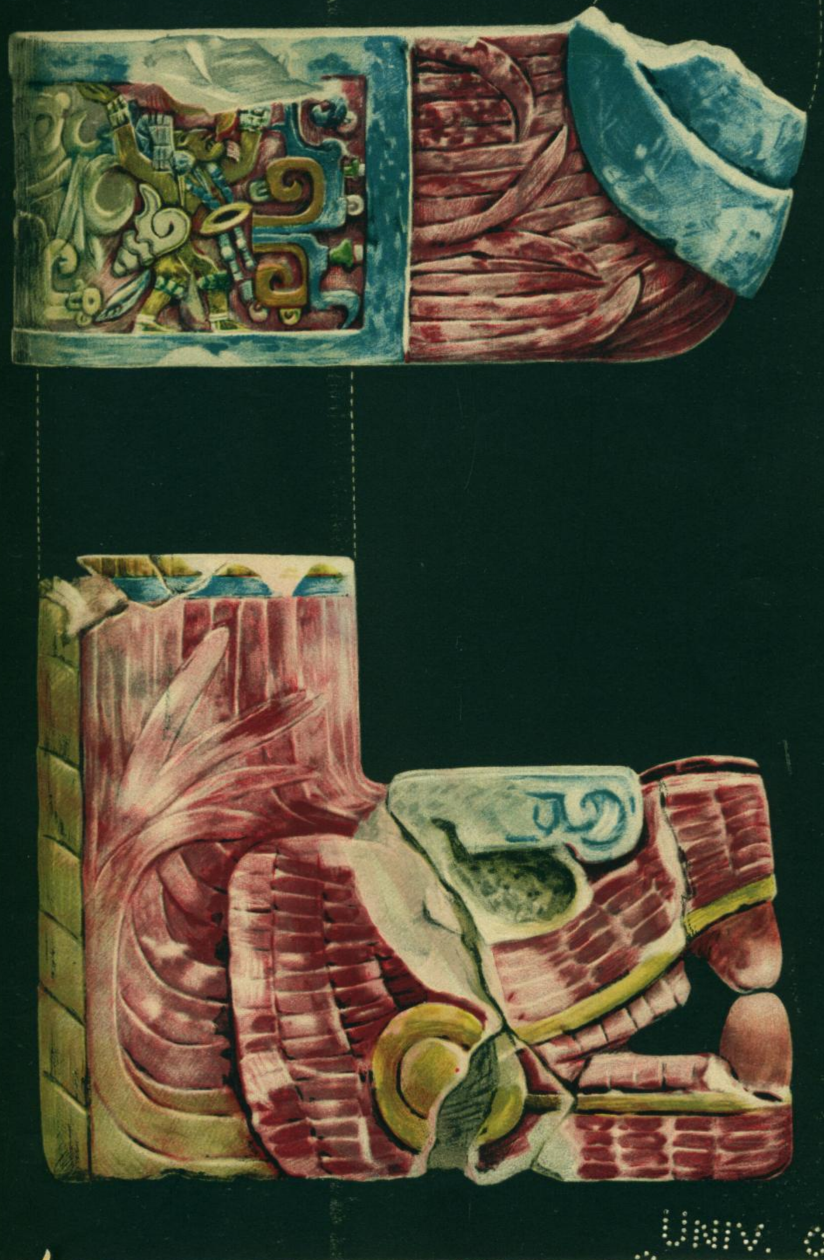

WATER COLOR BY JEAN CHARLOT

LITH. A. HOEN & CO., INC.

Head and Tail of Serpent Column, Temple of the Chac Mool

UNIV. OF CALIFORNIA

texture reconstruction

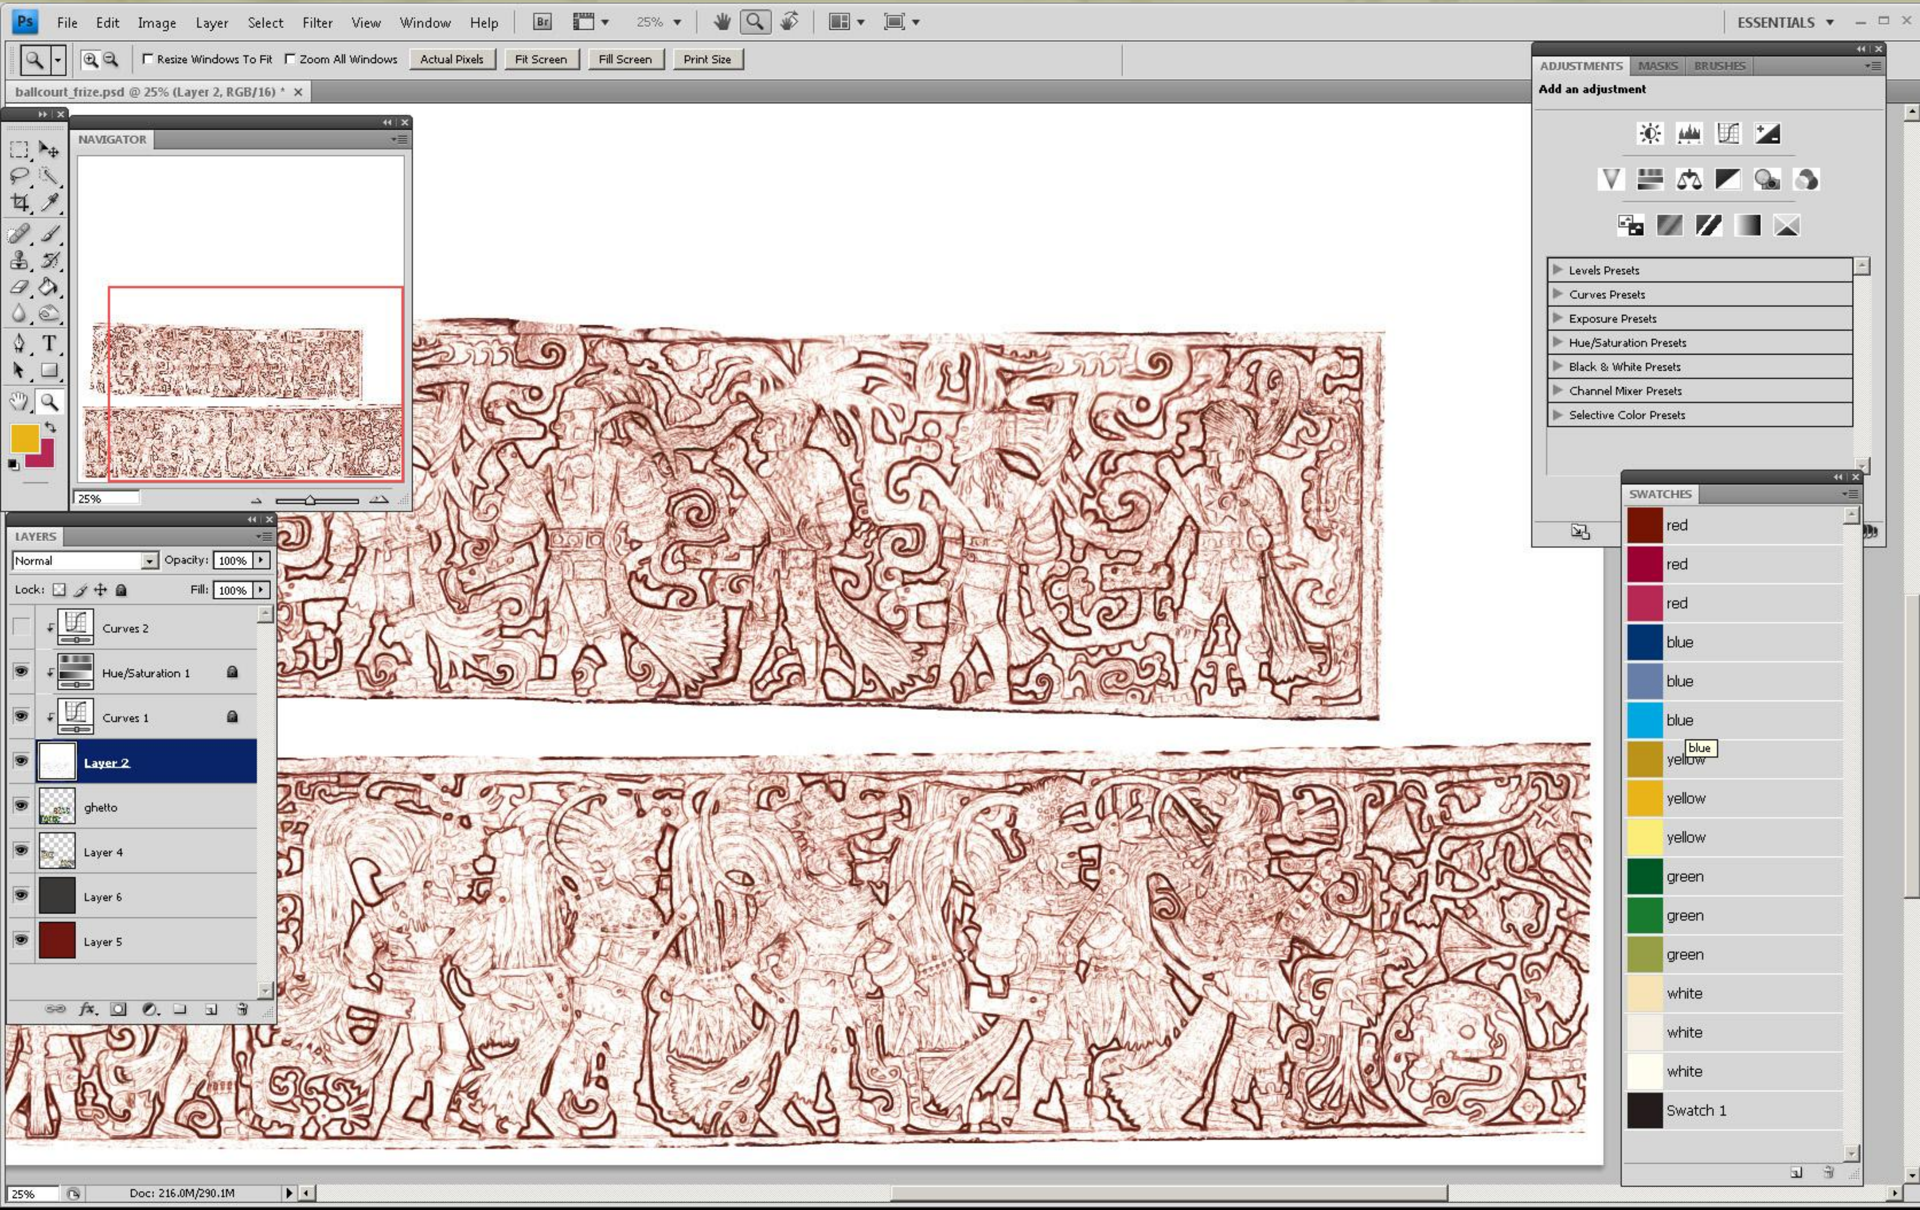

texture reconstruction

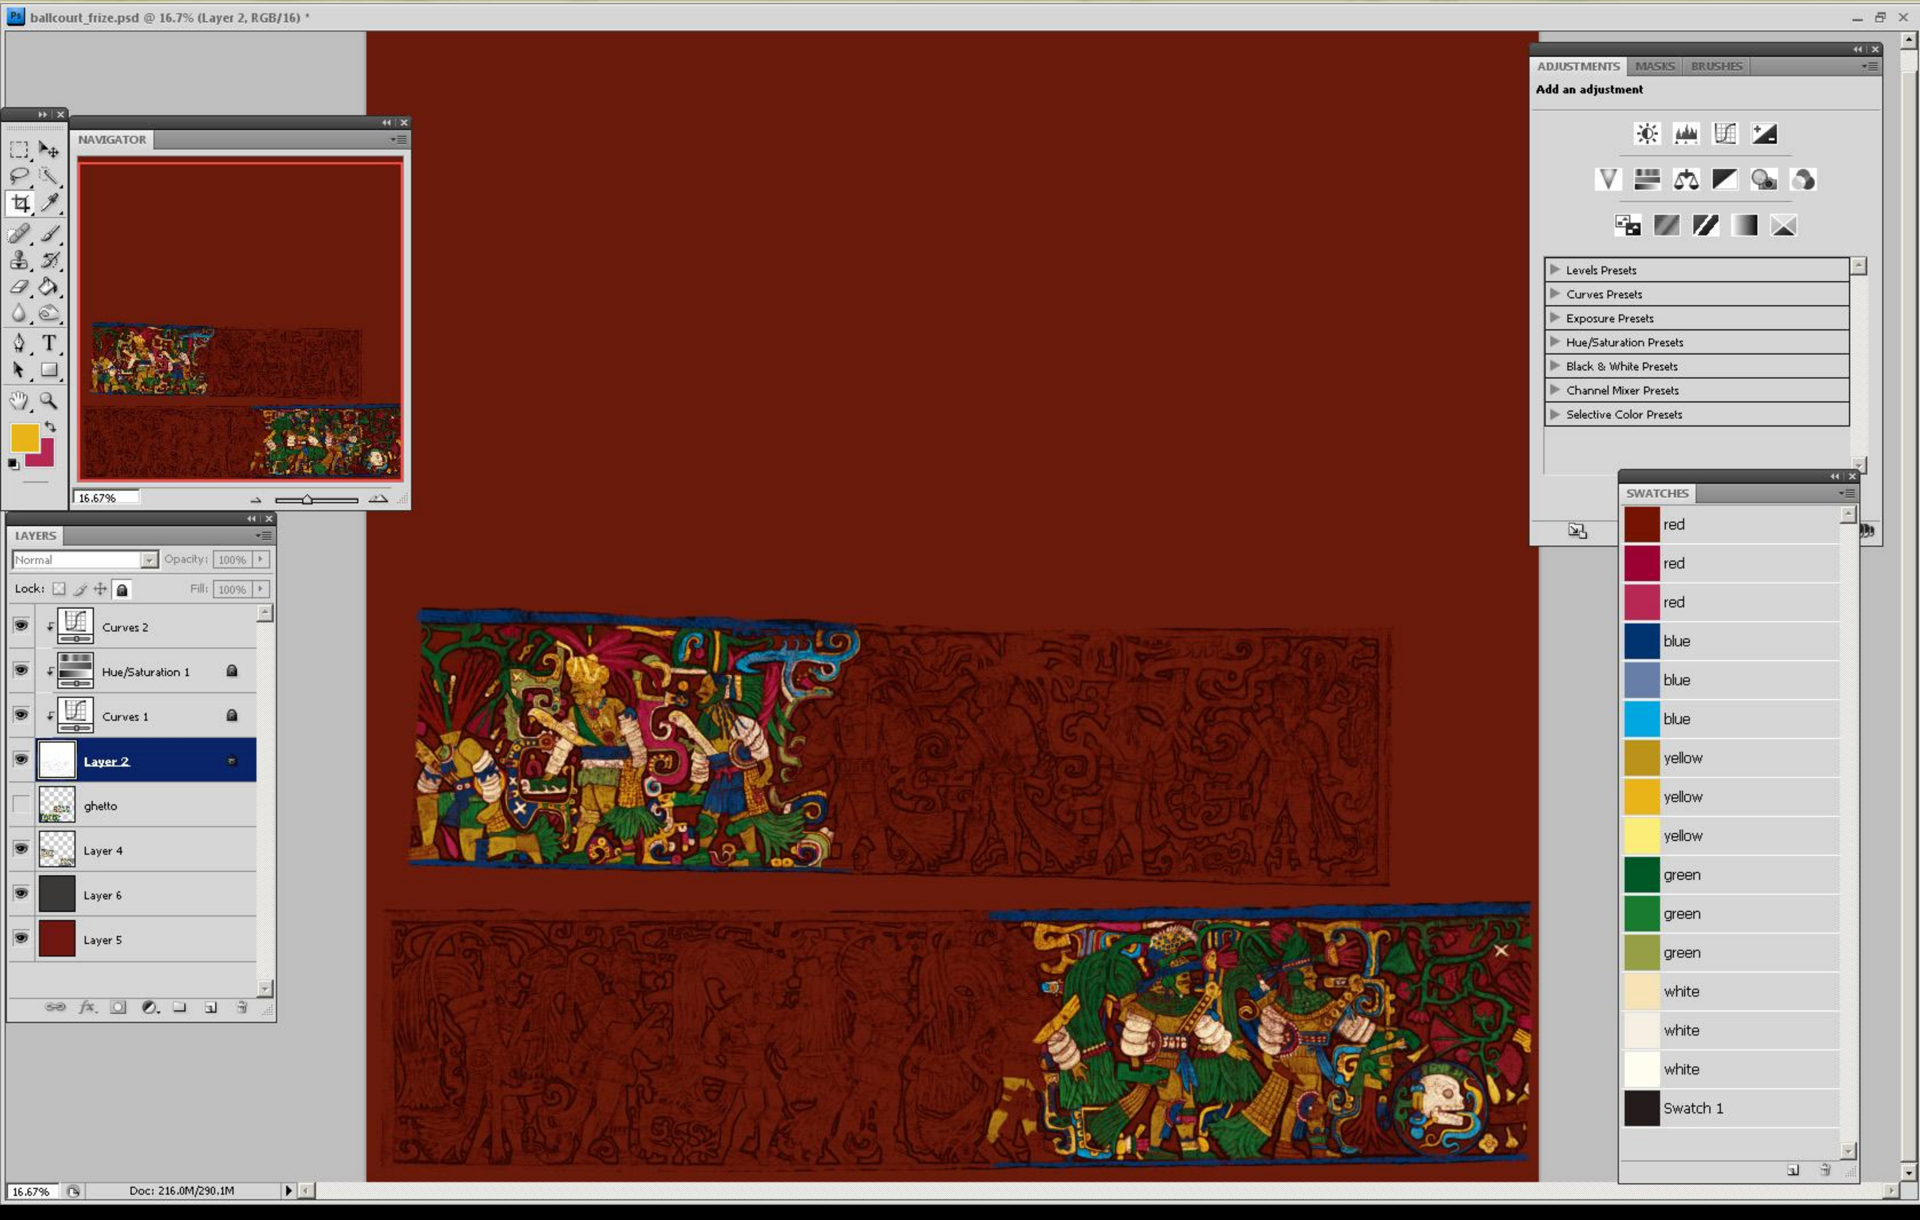

texture reconstruction

sculpture reconstruction via normal inpainting

end to end pipeline for processing archeological reconstructions

. evaluate the artifacts/sites under study

. plan your field work

. perform multi-resolution capture

. range alignment & merging

. surface simplification and normal extraction

. normal and color inpainting

. synthetic reconstruction

. rendering

texture reprojection

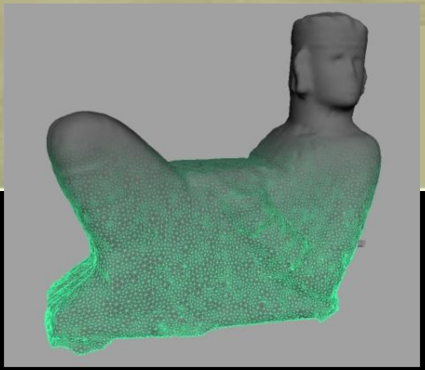

photo reprojection is useful when representing the current state of an artifact

. generate texture atlas for a simplified mesh

. recover camera position for an input set of images

. project color from the input set onto the mesh

























frequent pitfalls and deficiencies in the use of 3D acquisition techniques for CH

Paolo Cignoni