Surface Reconstruction from Point Clouds by Transforming the Medial Scaffold

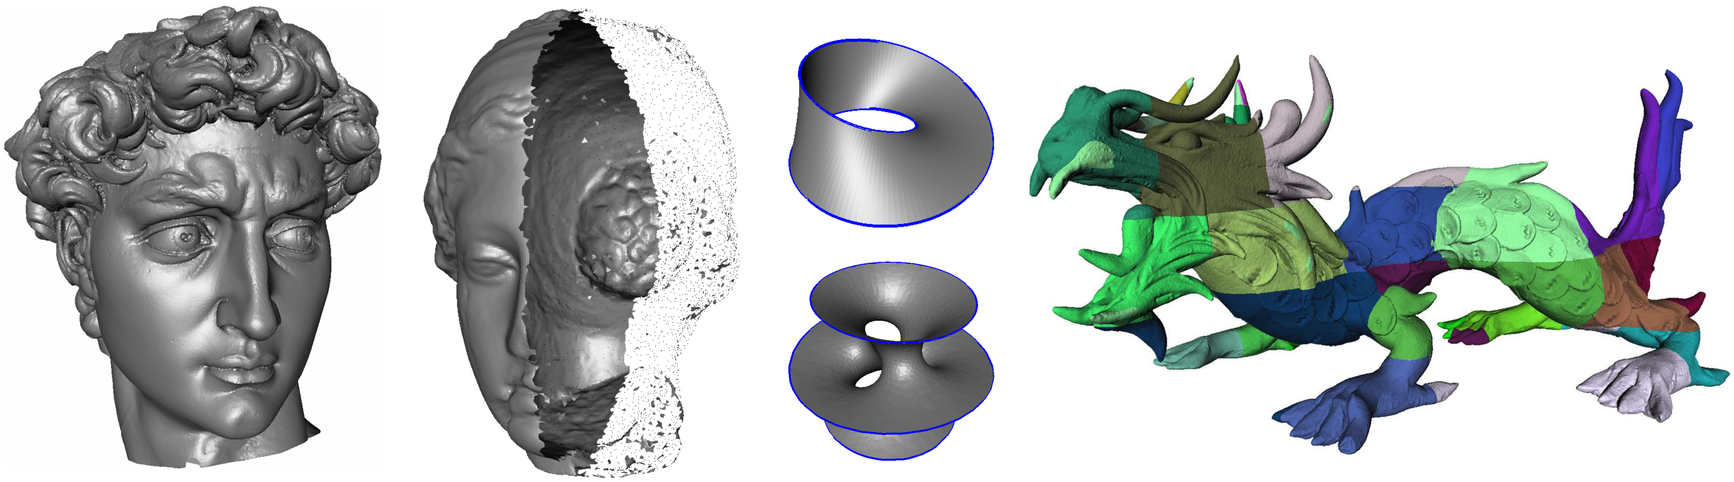

BROWN

Laboratory for Engineering Man/Machine Systems

Ming-Ching Chang

Frederic Fol Leymarie

Benjamin B. Kimia

Goldsmiths
UNIVERSITY OF LONDON

Problem: surface reconstruction with minimal assumptions

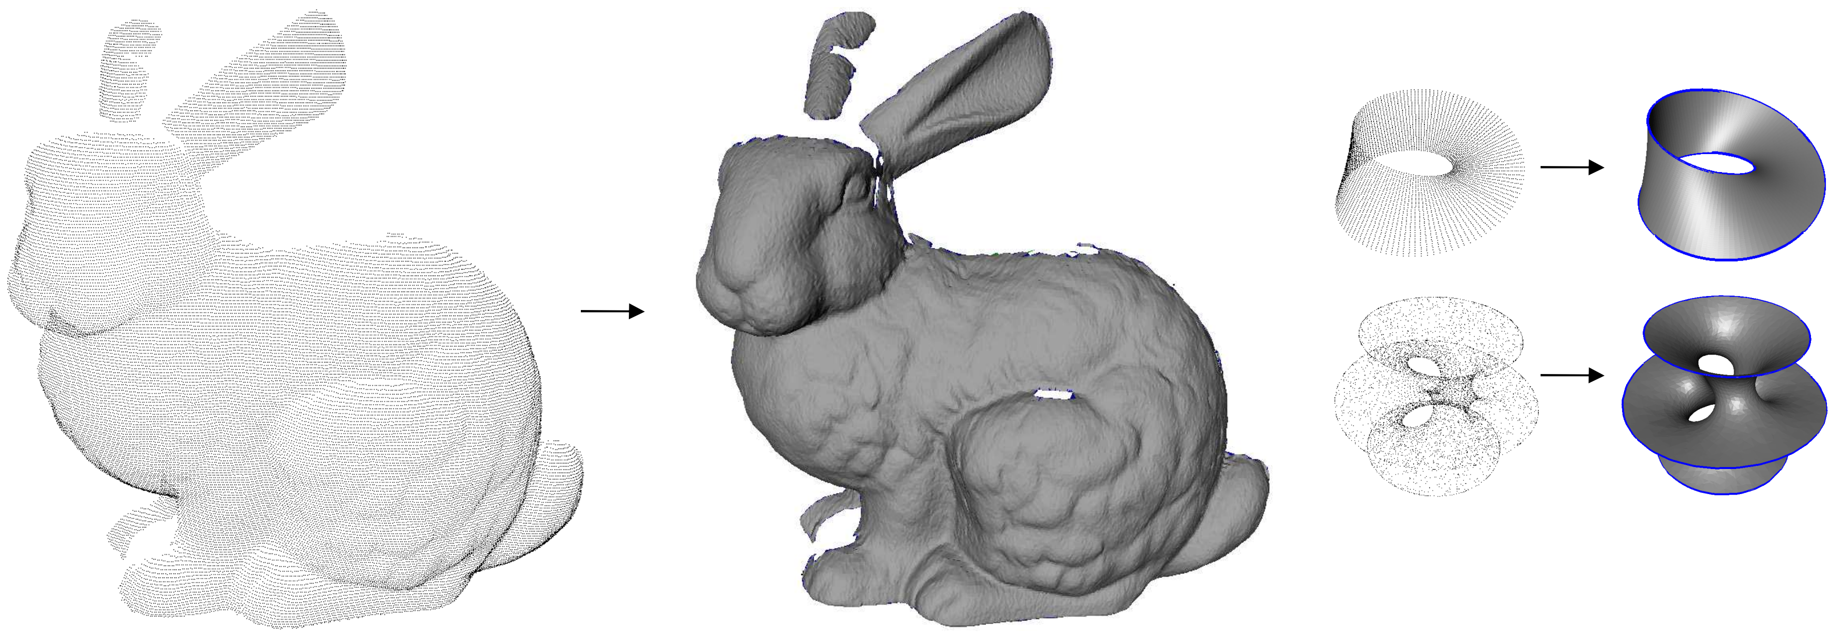

Context: reconstruct a surface mesh from *unorganized* points, with a "minimal" set of assumptions:

the samples are nearby a "possible" surface

(thick volumetric traces not considered here).

Benefit: reconstruction across many types of surfaces.

Goal: surface reconstruction with minimal assumptions

To find a *general* approach, applicable to various topologies,

without assuming strong *input constraints*, *e.g.:*

- No surface normal information.
- Unknown topology (with boundary, for a solid, with holes, non-orientable).
- No a priori surface smoothness assumptions.
- Practical sampling condition: non-uniformity, with varying degrees of noise.
- Practical large input size (> millions of points).

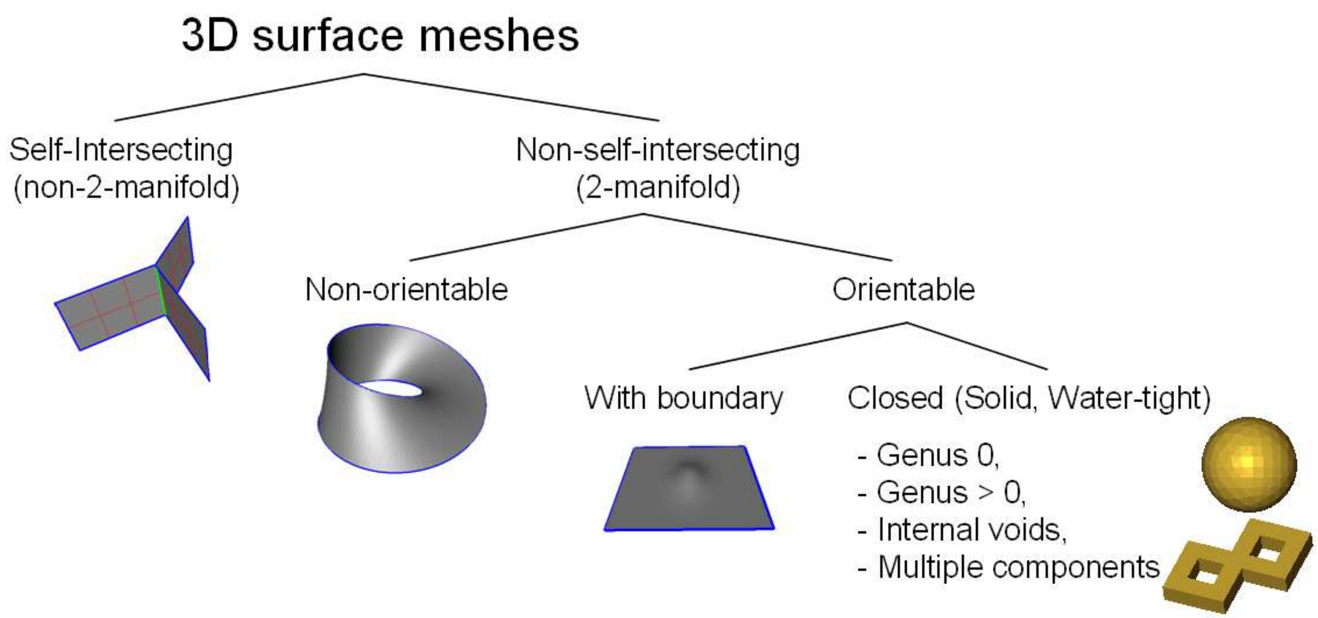

Goal: surface reconstruction with minimal assumptions

- Surface normal : not accurate, or problem locally solved
- Unknown topology : practical (e.g., holes, in CAD)
- No smoothness : practical (sharp features)
- Non-uniformity, noise : practical acquisition
- Large input size : scalable

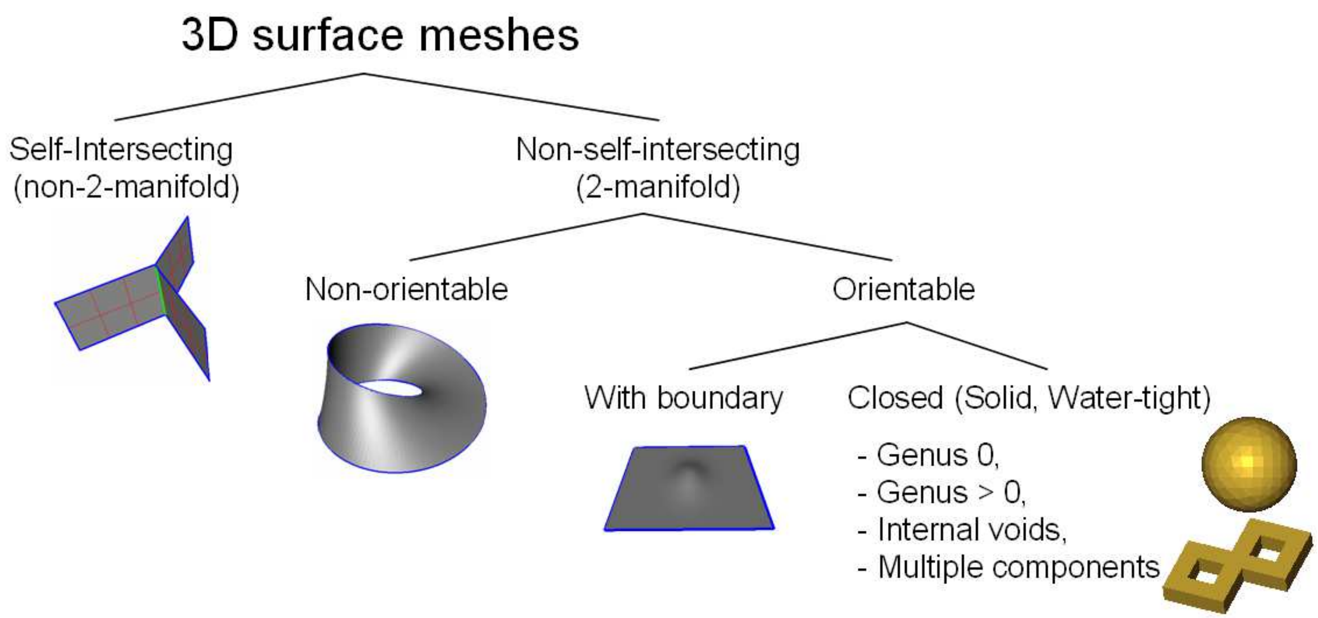

How: Literature Overview

- Implicit distance functions
  - Locally approximate the distance function by blending primitives.
  - Globally approximate the distance function by volumetric propagation.
- Propagation based (region growing) methods
- Voronoi / Delaunay geometric constructs
  - Incremental surface-oriented.
  - Volume-oriented.

Many methods have additional assumptions in addition to *unorganized* points:

- Surface normal: imply knowing the surface locally.
- Surface enclose a volume (distance field): a strong *global* information.

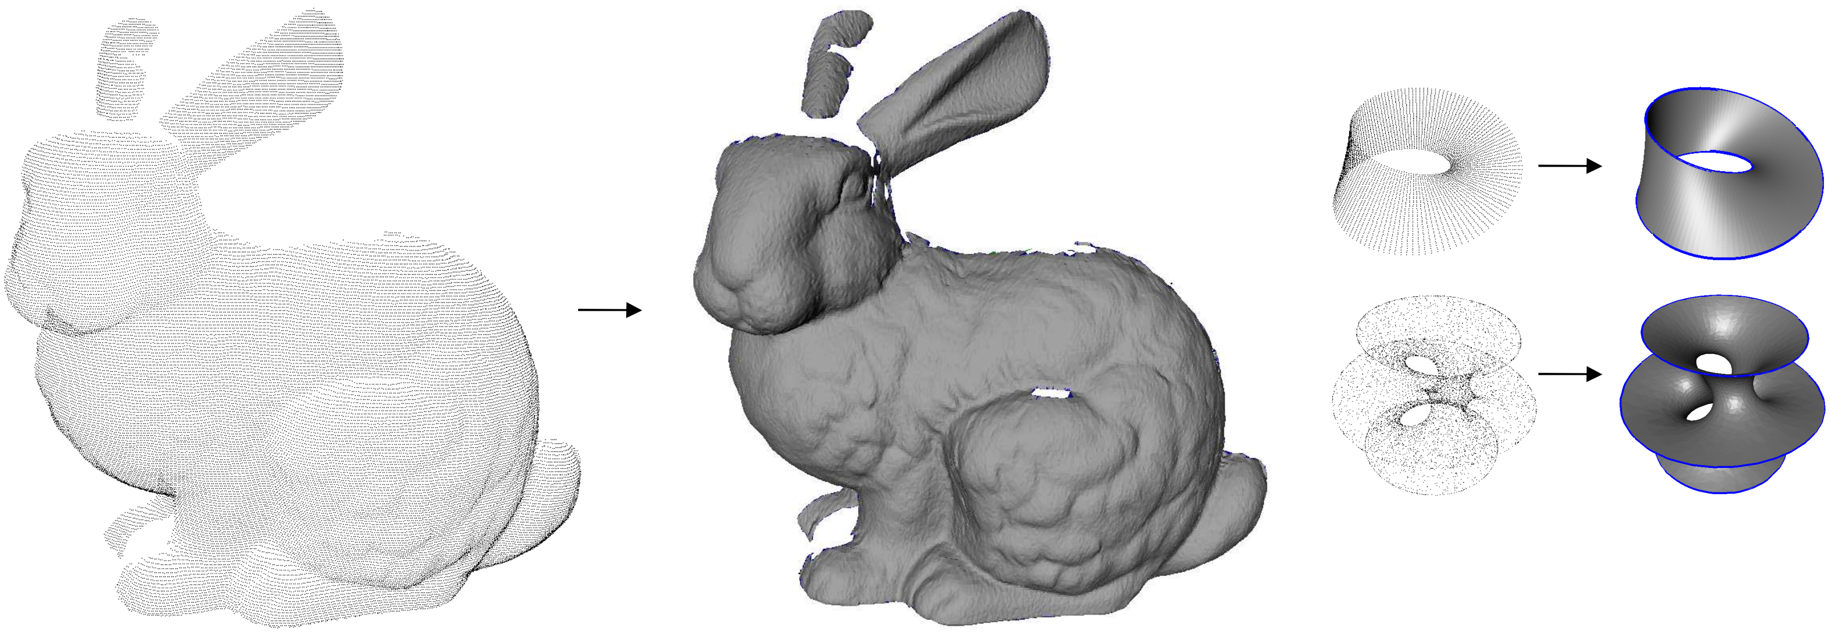

How we solve it: Find an Inverse of Sampling:

Relate the sampled surface with the underlying (unknown) surface and try to *invert* (recover) the sampling process…

How: Overview of Our Approach (2D)

Not many clues from the assumed loose input constraints.

- Work on the shape itself to recover the sampling process.

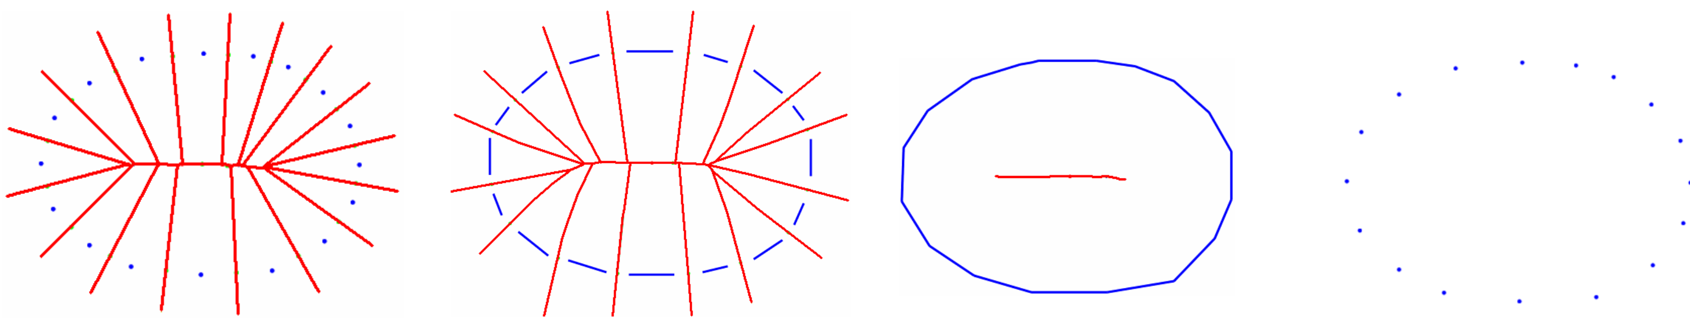

Key ideas:

- Relate the sampled shape with the underlying (unknown) surface by a sequence of shape deformations (growing from samples).
- Represent (2D) shapes by their medial "shock graphs". [Kimia *et al.*]
- Handle shock *transitions* across different shock topologies to recover gaps.

How: Sampling / Meshing as Deformations

Schematic view of sampling: infinitesimal holes grows, remaining are the samples.

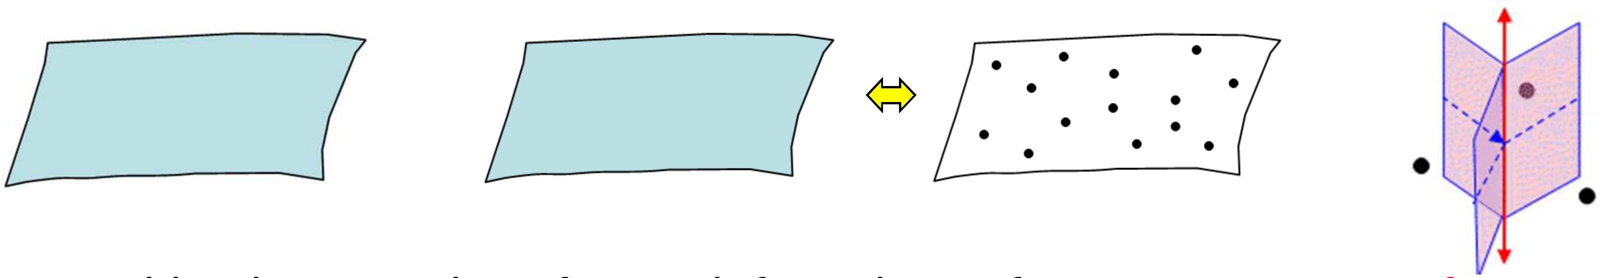

We consider the removing of a patch from the surface as a Gap Transform.

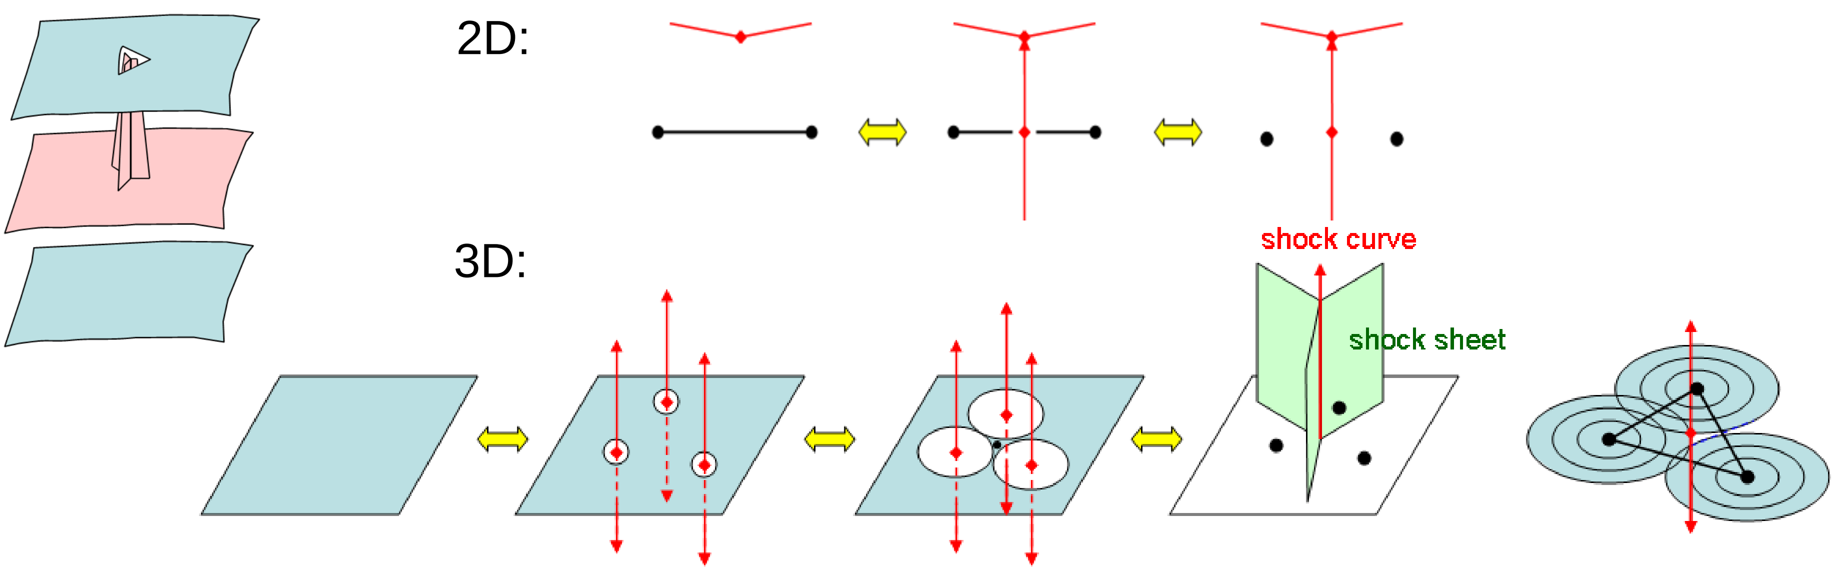

Surface

Infinitesimal holes

Holes merging, surface between holes degenerating into points

The shock curve with 3 shock sheets are used to recover gaps.

How: Medial Scaffolds for **3D** Shapes

A graph structure for the 3D Medial Axis

Classify shock points into 5 general types, and organized into a hyper-graph form [Giblin, Kimia PAMI'04]:

- Shock Sheet: A_1^2
- Shock Curves: A_1^3 (***Axial***), A_3 (***Rib***)
- Shock Vertices: A_1^4, A_1A_3

A_k^n: contact (max. ball) at *n* distinct points, each with *k+1* degree of contact.

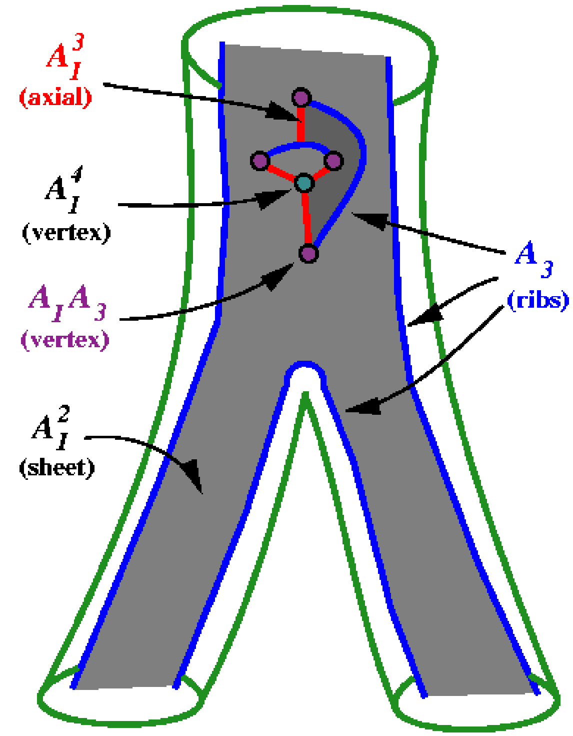

A special case of input of points

the Medial Scaffold consists of only:

A_1^2 Sheets, A_1^3 Curves, A_1^4 Vertices.

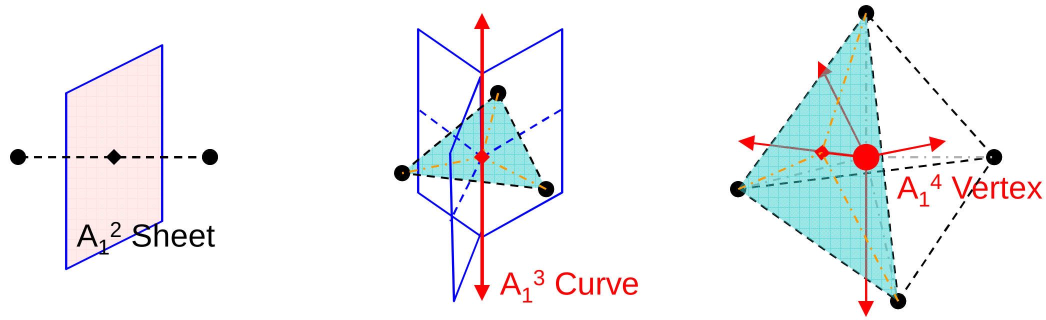

How: Medial Scaffolds for **3D** Shapes

A graph structure for the 3D Medial Axis

- *Augmented* Medial Scaffold (MS+): hyper-graph [Leymarie PAMI'07]:
- *Reduced* Medial Scaffold (MS): 1D graph structure
  - Shock sheets are seen as redundant (loops in the graph).

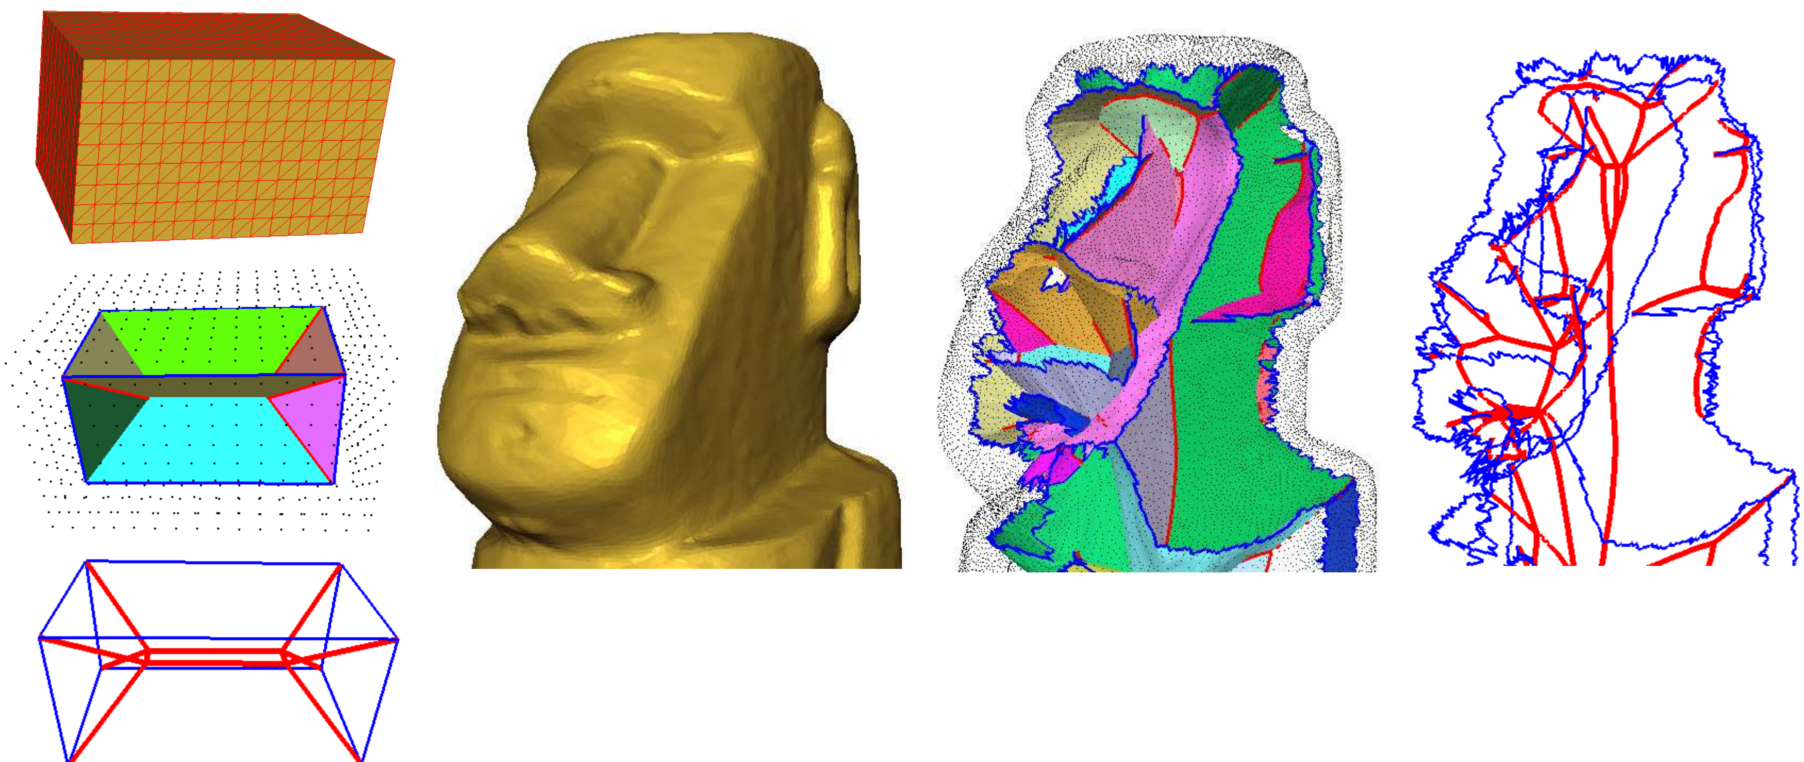

Easter island statue point data courtesy of Yoshizawa *et al.*

How: Organise/Order Deformations (**2D**)

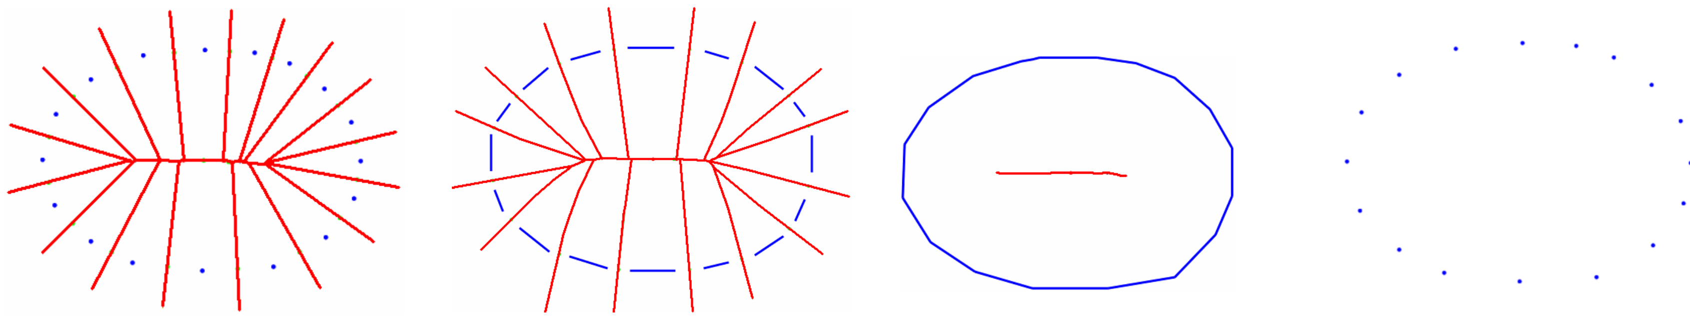

A B

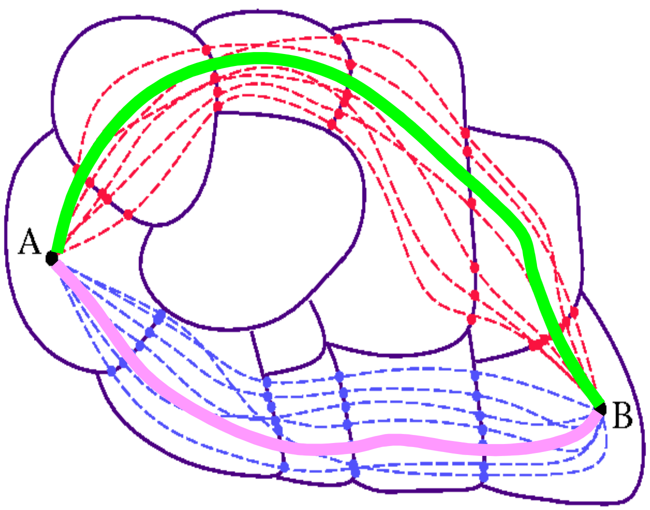

Deformation in shape space

NB: A & B share **object** symmetries. Symmetries due to the **sampling** need to be identified.

How: Organise/Order Deformations (**3D**)

- Recover a mesh (connectivity) structure by using Medial Axis transitions modelled via the Medial Scaffold (MS).
  - Meshing as *shape deformations* in the '*shape space*'.
- The Medial Scaffold of a point cloud includes both the symmetries due to sampling and the original object symmetries.
  - Rank order Medial Scaffold *edits* (gap transforms) to "segregate" and to simulate the recovery of sampling.

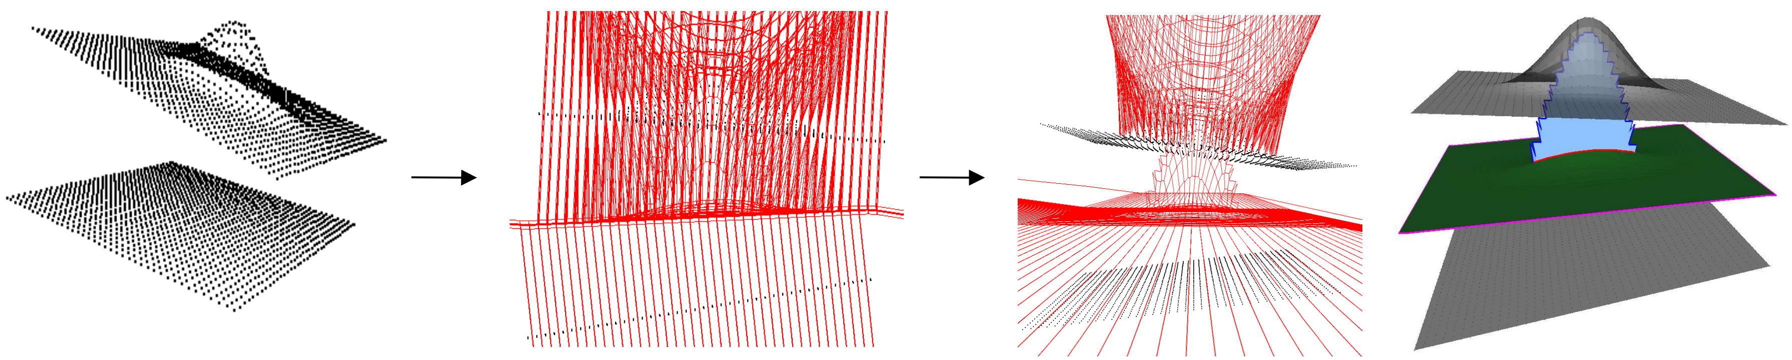

Sampling recovery

Object symmetry

Meshed Surface + Organized MA

Shock Segregation [Leymarie, PhD'02]

Algorithmic Method

Enough theory…

Here is how we order symmetries (and thus gap transforms) in practice.

Algorithmic Method

- Consider Gap Transforms on *all* A_1^3 shock curves in a ranked-order fashion:
  - best-first (greedy) with error recovery.
- Cost reflects:
  - Likelihood that a shock curve (triangle) represents a surface patch.
  - Consistency in the local context (neighboring triangles).
  - Allowable (local surface patch) topology.

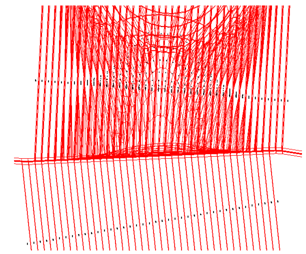

A_1^3 shock curve

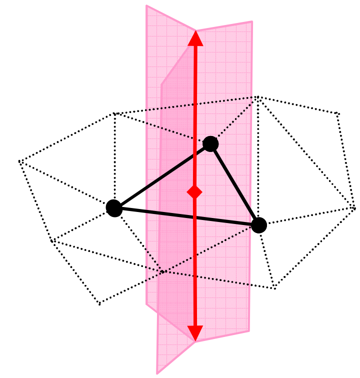

Three A_1^2 shock sheets

3 Types of A_1^3 shock curves (dual Delaunay triangles):
Represented in the MS by “singular shock points” (A_1^3-2)

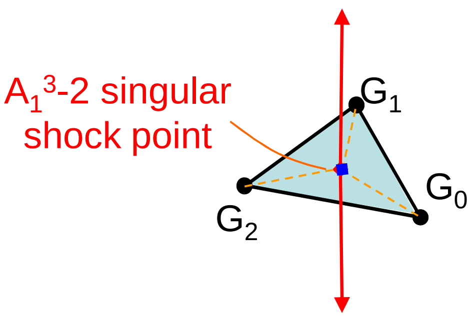

Type I

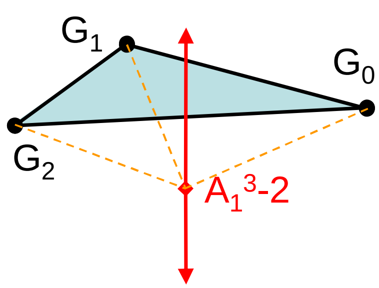

Type II

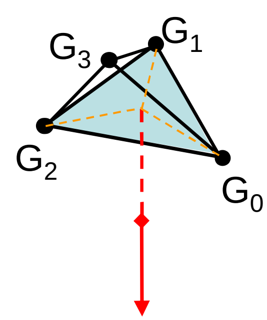

Type III

(unlikely to be correct candidate)

Algorithmic Method

How we **order gap transforms**:

- Favor small “compact” triangles.
- Favor recovery in “nice” (simple) areas, *e.g.*, away from ridges, corners, necks.
- Favor simple local continuity (similar orientation).
- Favor simple local topologies (2D manifold).
- BUT: allow for error recovery!

Ranking Isolated Shock Curves (Triangles)

Triangle geometry:

$$
\begin{aligned}
D &= \max(d_1, d_2, d_3) \\
P &= d_1 + d_2 + d_3 \\
m &= (d_1 + d_2 - d_3)(d_3 + d_1 - d_2)(d_2 + d_3 - d_1) \\
A &= \sqrt{(P \cdot m)/16} \qquad \text{(Heron's formula)} \\
C &= 4\sqrt{3} \cdot A/(d_1^2 + d_2^2 + d_3^2), \qquad \text{(Compactness, Gueziec's formula, } 0<C<1\text{)}
\end{aligned}
$$

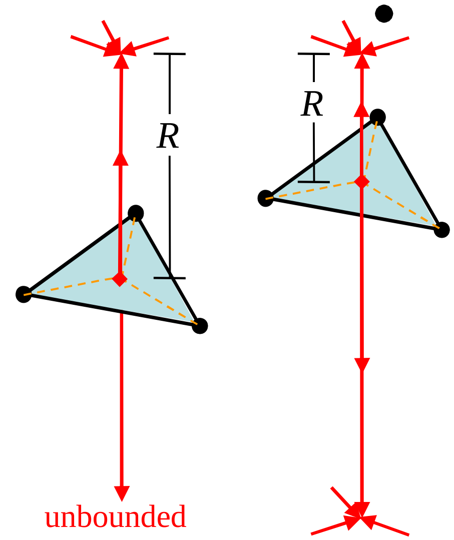

The side of smaller shock radius is more salient.

Cost: favors *small compact* triangles with large shock radius *R*.

$$
\rho_1 = \begin{cases} \frac{P}{R} \cdot \frac{1}{C^2}, & \text{if } D < d_{\max} \\ \infty, & \text{if } D \geq d_{\max} \end{cases}
$$

R: minimum shock radius

d_{max}: maximum expected triangle, estimated from d_{med}

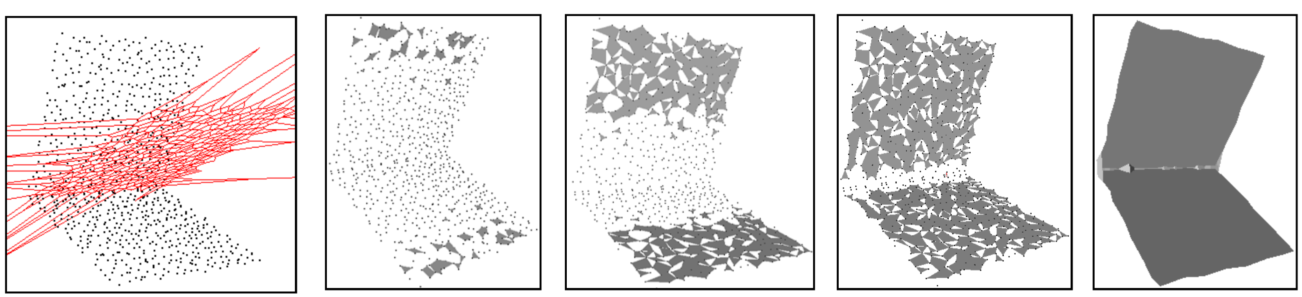

Surface meshed from confident regions toward the sharp ridge region.

Estimate the Sampling Scale

The maximum expected triangle size (d_{max}) can be estimated from shock radius distribution analysis.

Distribution of the A_1^3-2 radii of all shock curves corresponding to:

All triangles in the original Stanford bunny mesh.

All triangles of shock curves of type I and type II in the (full) Medial Scaffold of the point cloud.

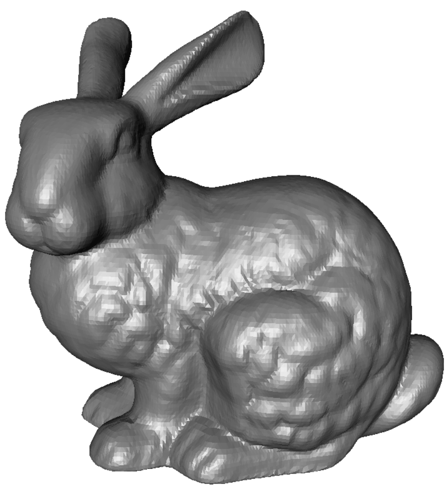

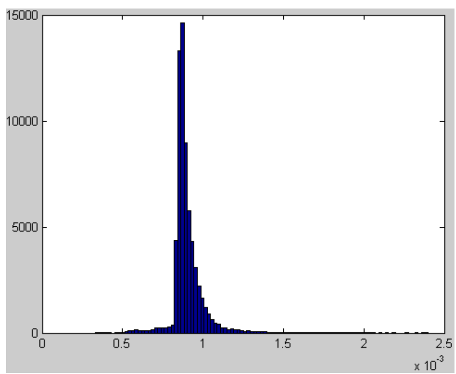

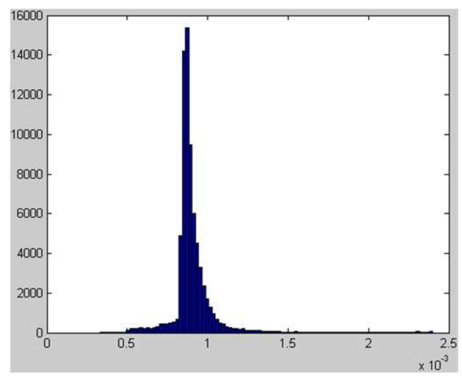

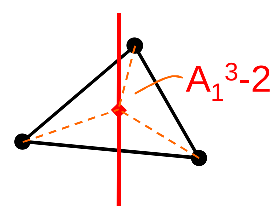

The median of the A_1^3-2 distribution (d_{med}) approximates its peak.

Cost Reflecting Local Context & Topology

Cost to reflect smooth continuity of edge-adjacent triangles:

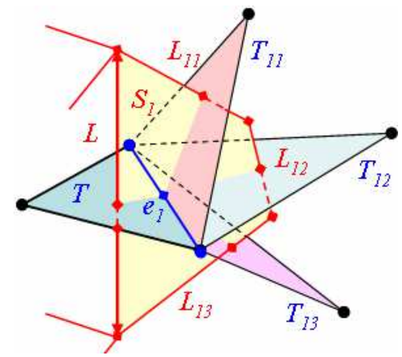

$$\rho_2 = \frac{d}{R} \cdot \frac{1}{C^2} \cdot f(\theta)\,,$$

$$f(\theta) = [\exp^{\theta} - 1]^2 - 1 \begin{cases} \theta = 0,\ f(\theta) = -1 \\ \theta = 40^\circ,\ f(\theta) \simeq 0 \\ \theta = 80^\circ,\ f(\theta) \simeq 8.24 \end{cases}$$

Typology of triangles sharing an edge:

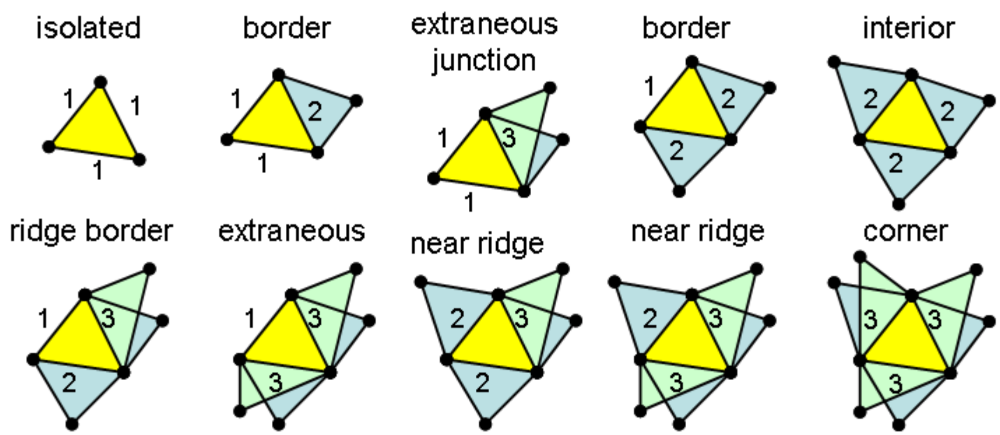

Typology of mesh vertex topology

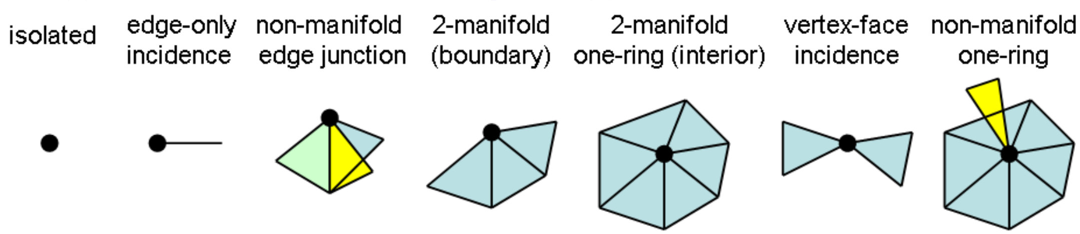

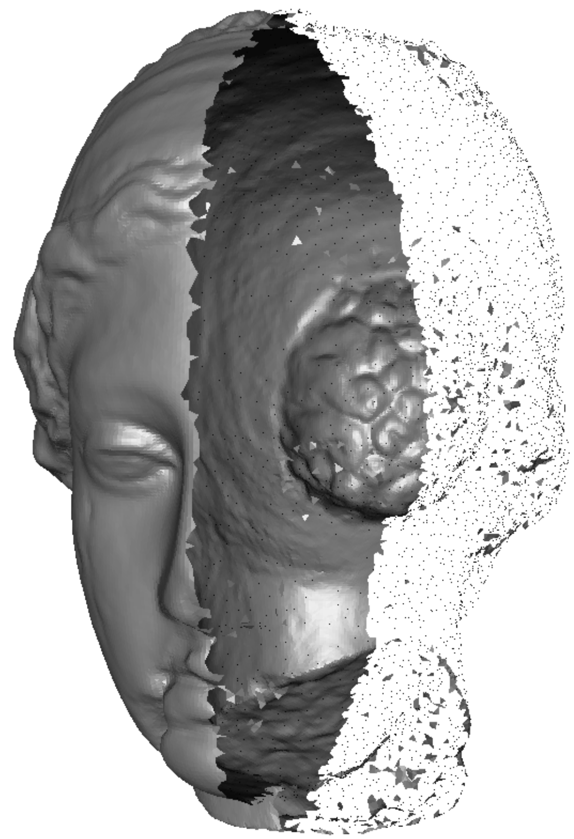

Point data courtesy of Ohtake *et al*.

Strategy in the Greedy Meshing Process

Problem: Local ambiguous decisions → errors.

Solutions:

- Multi-pass greedy iterations

  First construct confident surface triangles without ambiguities.

- Postpone ambiguous decisions
  - Delay related candidate Gap Transforms close in rank, until additional supportive triangles (built in vicinity) are available.
  - Delay potential topology violations.
- Error recovery
  - For each Gap Transform, re-evaluate cost of both related *neighboring (already built)* & *candidate* triangles.
  - If cost of any existing triangle exceeds top candidate, undo its Gap Transform.

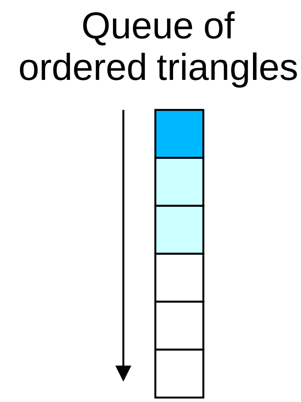

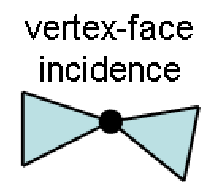

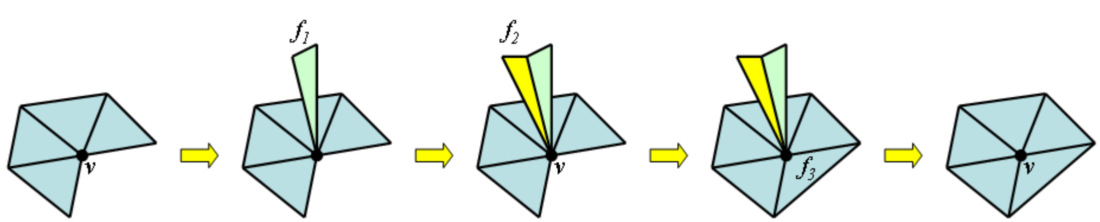

Summary of Our Approach

- We relate an object and its sampling by navigating the “shape space” (of deformations).
- We organize this navigation by gap transforms on the Medial Scaffold.
- We select a path by ordering these transforms and allowing for error recovery.

Show Time!

- Some results
- Other issues:
  - Validation,
  - Using a priori information,
  - Dealing with large inputs,
  - Sampling quality,
  - Running time.
- Conclusions

Results: Surface with Various Types

Water-tight surface

With sharp ridges
(discontinuous curvature)

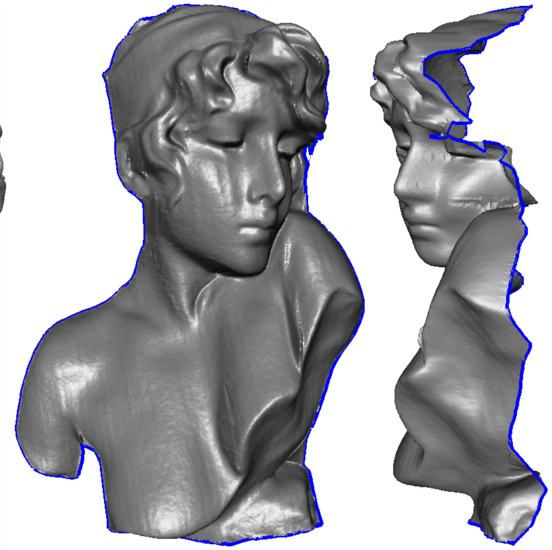

With boundary

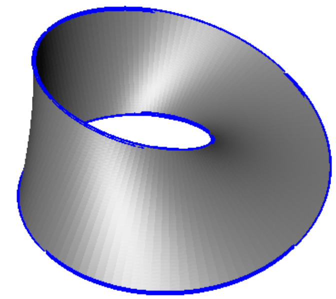

Non-orientable

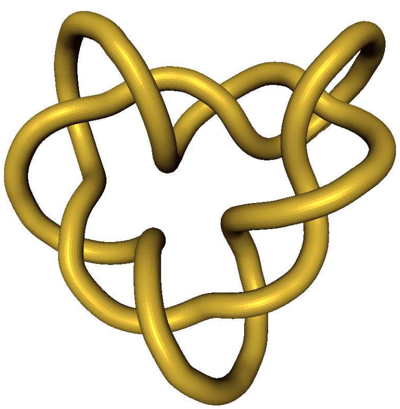

Closely knotted

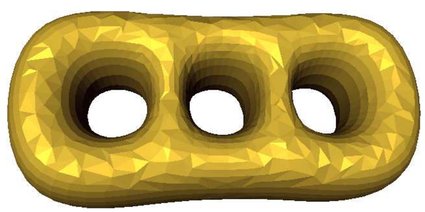

With multiple holes

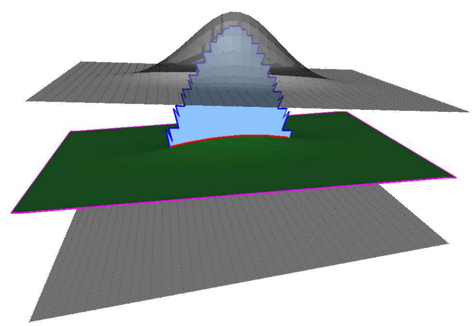

Multiple components

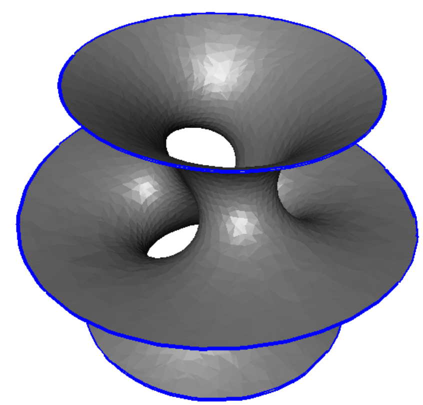

Multiply punctured

Gold: water-tight surface: Blue: mesh boundary.
Dataset are courtesy of Cyberware, Stanford data repository, Stony Brook archive, H. Hoppe.

Result: Videos on Meshing Algebraic Surfaces

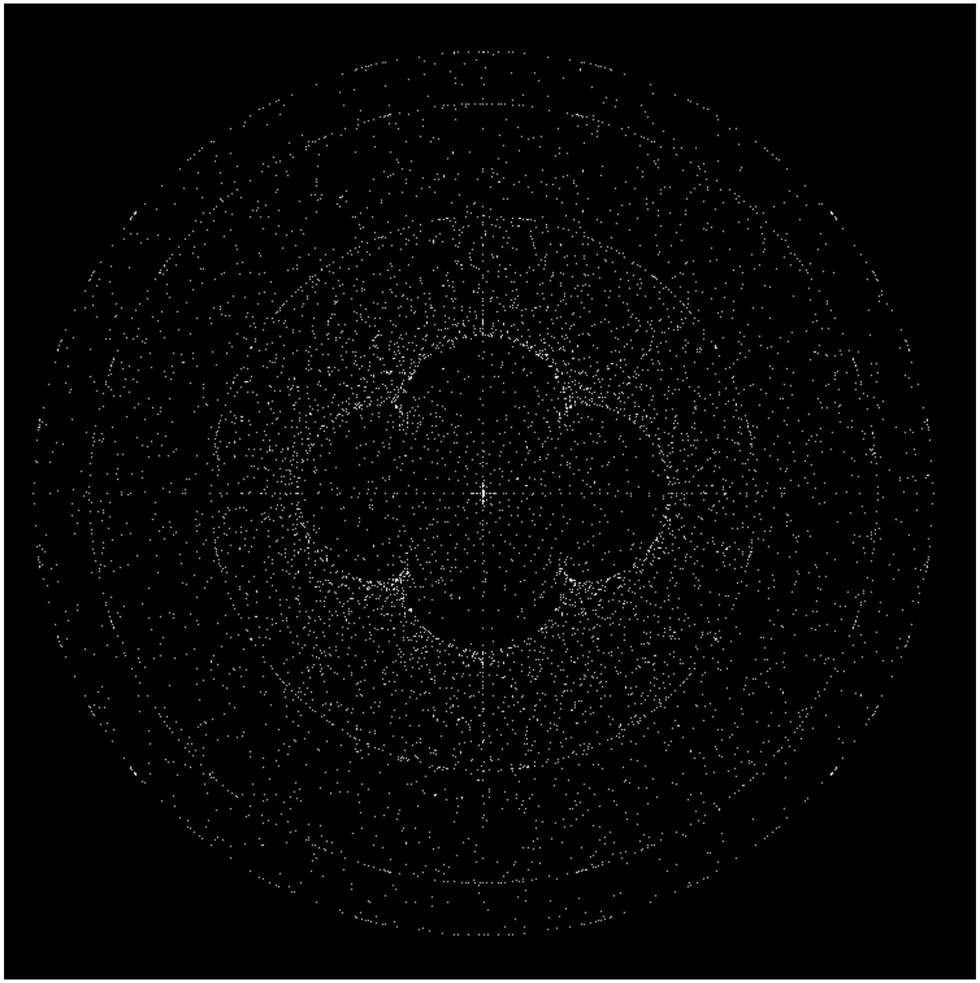

Costa minimum surface (courtesy of H. Hoppe)

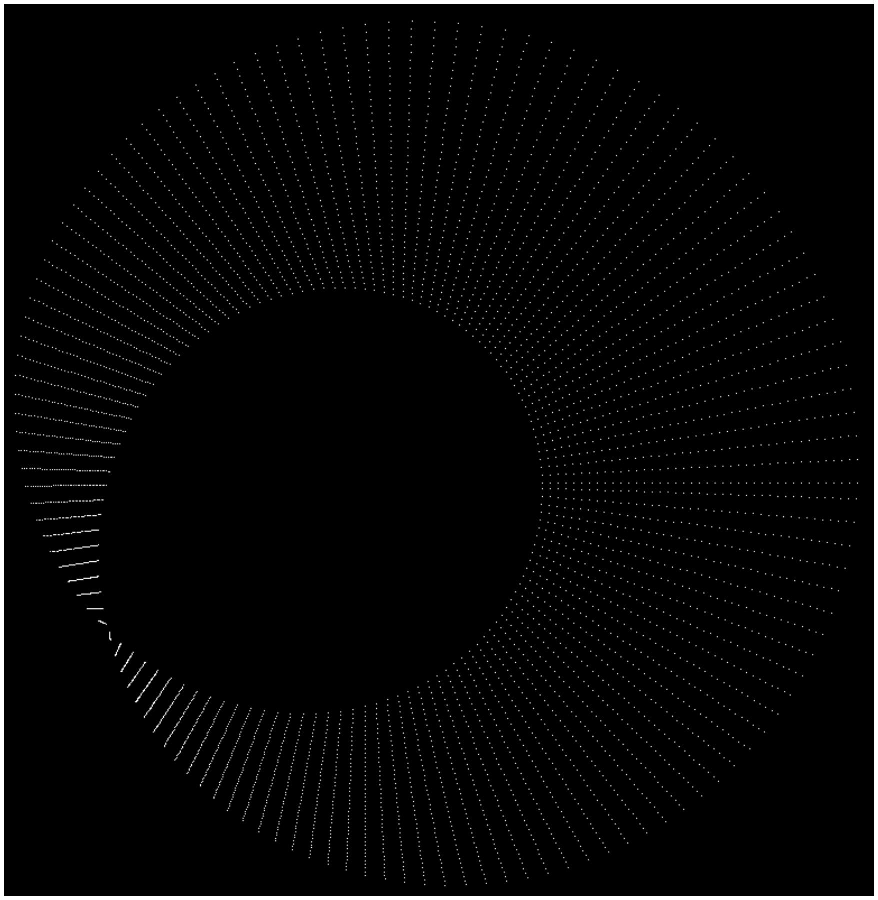

Mobius strip

Result: Video on Meshing the Rocker Arm

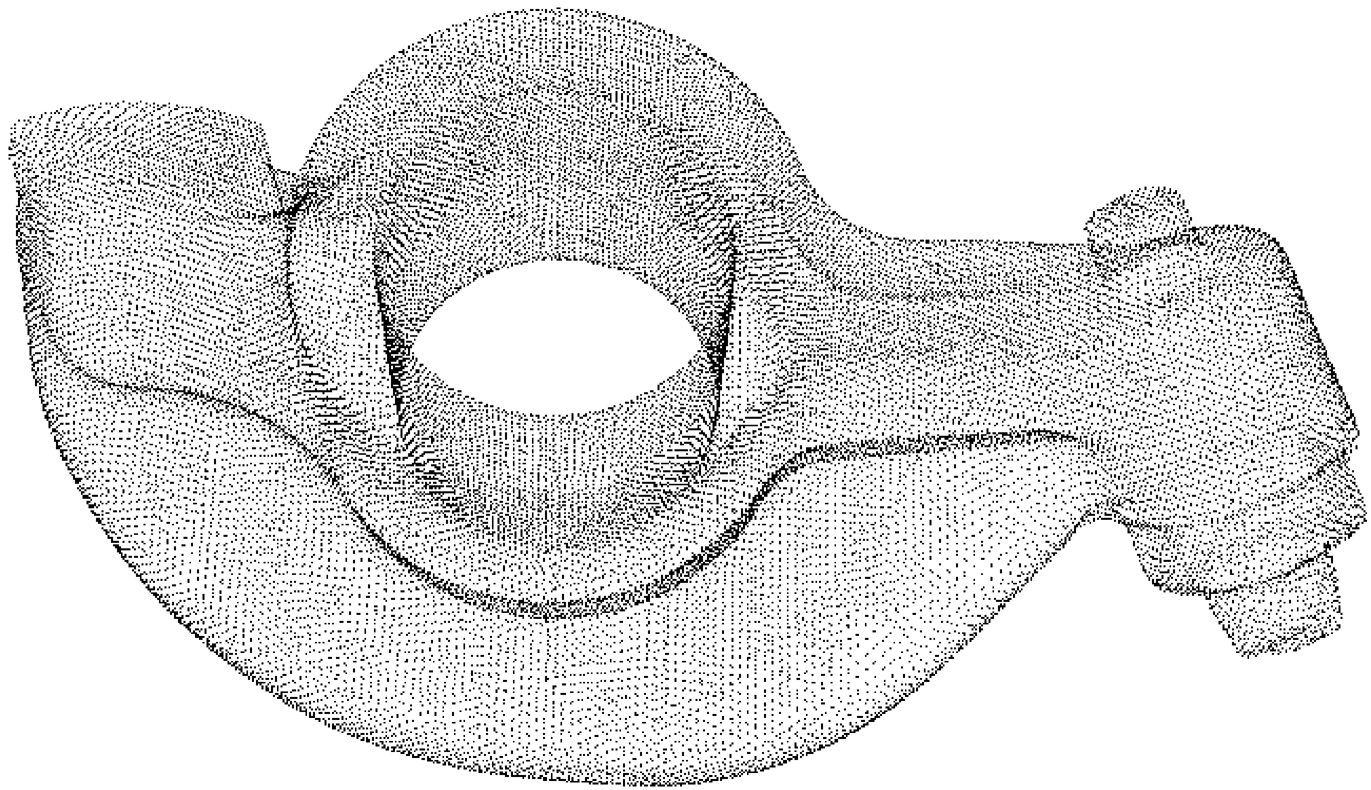

Flat smooth regions are meshed prior to the ridges/corners.

The rocker arm data courtesy of Cyberware.

Validation

- Superimpose our meshing result on the original mesh.

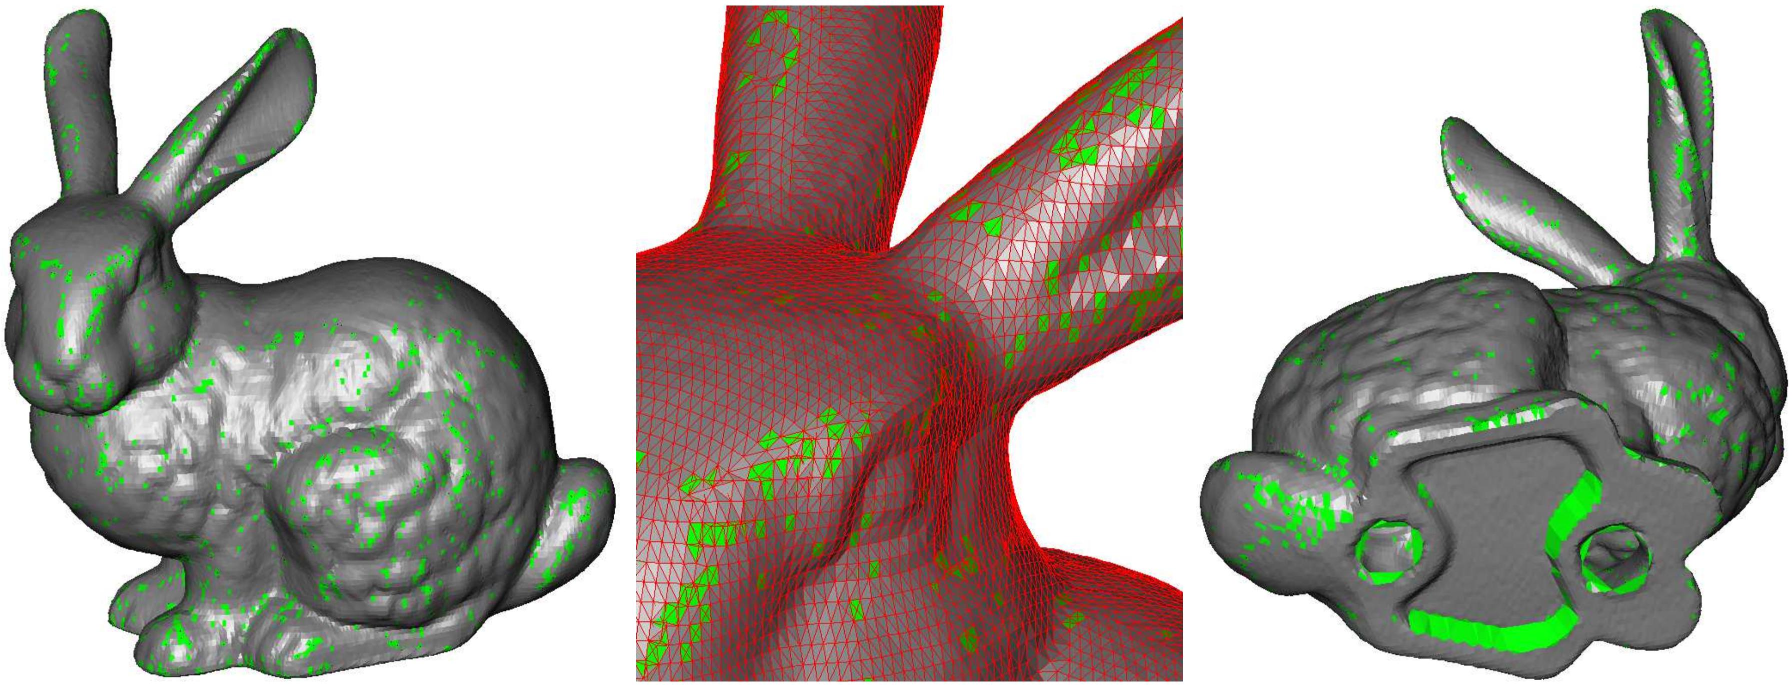

Color: Original mesh in gray.

Difference of reconstructed triangles in green.

Other Issues

- Validation,
- **Using a priori information,**
- Dealing with large inputs,
- Sampling quality,
- Running time.

Re-mesh / Repair a Partial Mesh

- In the case that existing triangles (in addition to the points) are know *a priori*:
  - Assign high priority to existing triangles.
  - Let candidates compete in the greedy algorithm.
  - Similar if surface normal is available.

4,102 points sampling a mechanical part (courtesy of H. Hoppe)

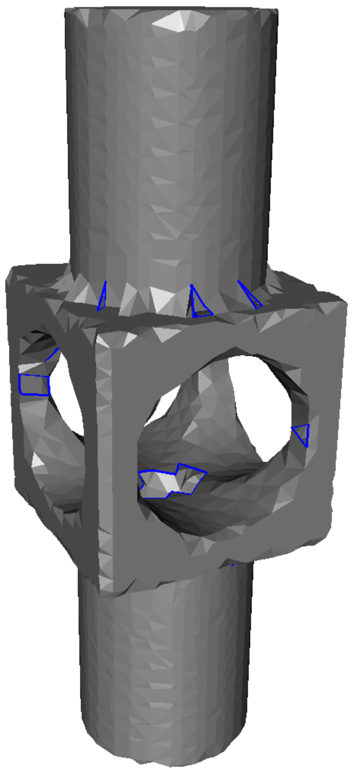

Meshing result of an implementation of ball pivoting algorithm (BPA) containing holes / topological errors.

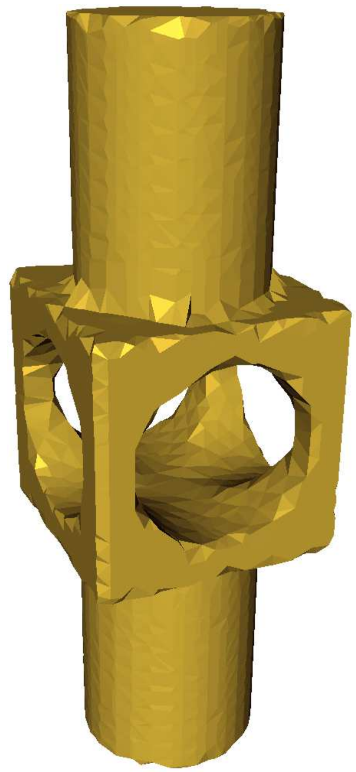

Re-mesh results of our algorithm (a solid)

Handle Large Datasets (Millions of Points)

- No strong constraints (topology, boundary, volume, *etc.*) on input.
- Divide input into buckets (or any full partition of space).
- Mesh surfaces in each bucket.
- Stitch surfaces by applying the same algorithm again.

Meshing Stanford Asian Dragon (3.6M points). Related to [Dey *et al.*'01]: Super Cocone.

Result of Stitching After Meshing in Buckets

Result: Bucketing + Stitching Video

120,965 points, divided into 20,000 points per bucket.

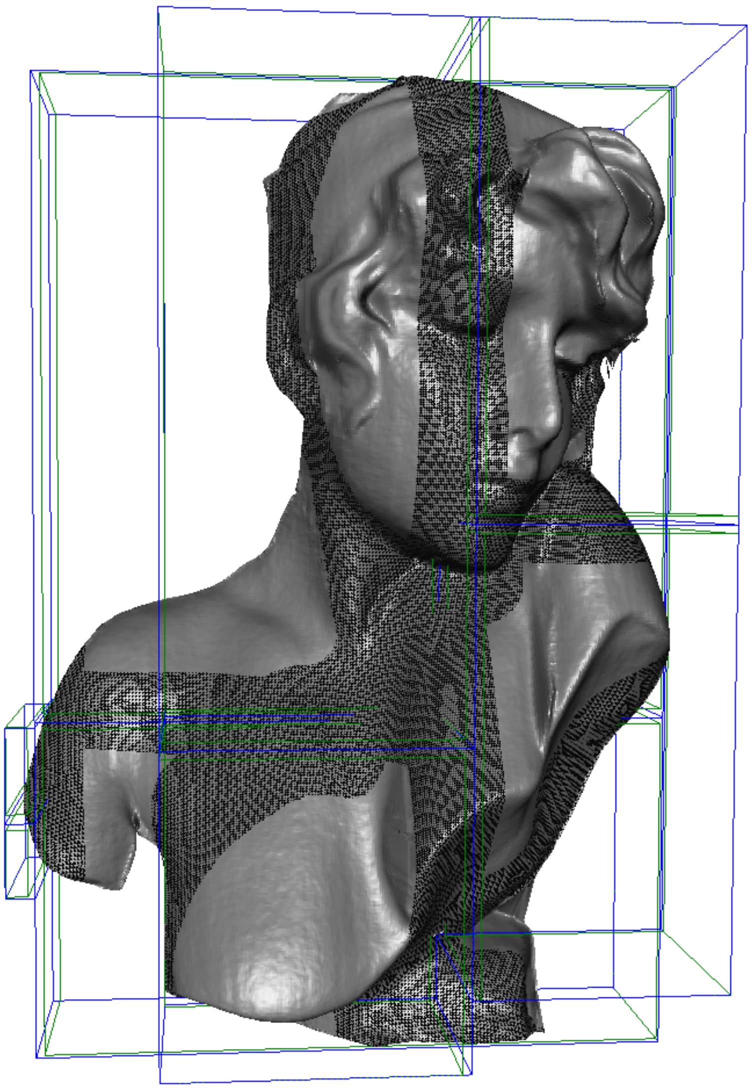

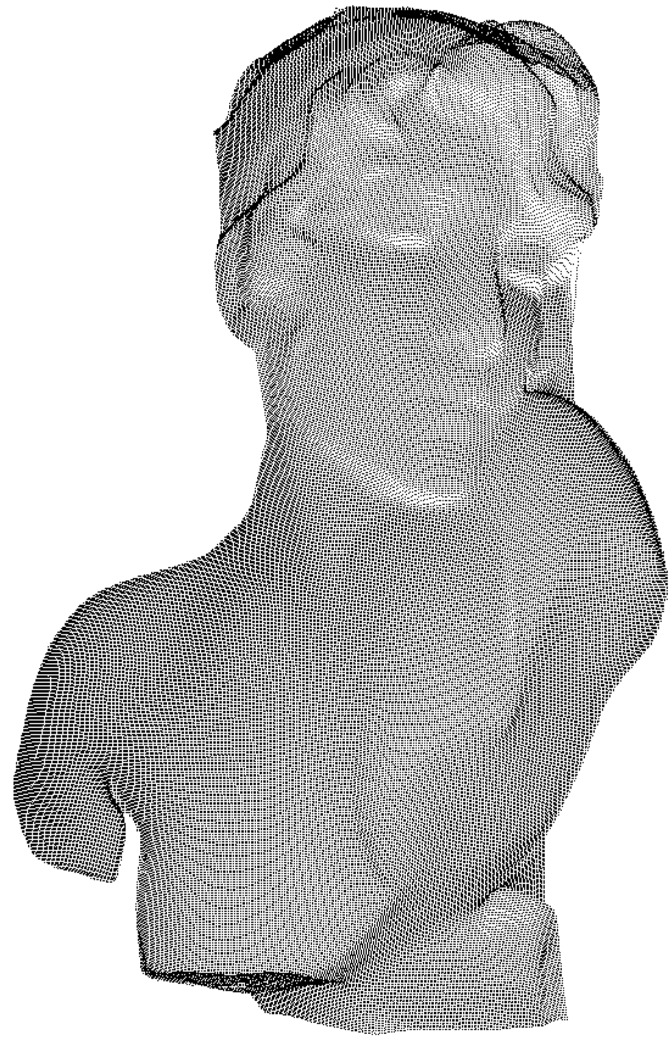

Sapho dataset courtesy of Stony Brook archive.

Meshing Stanford Thai Statue (5M points)

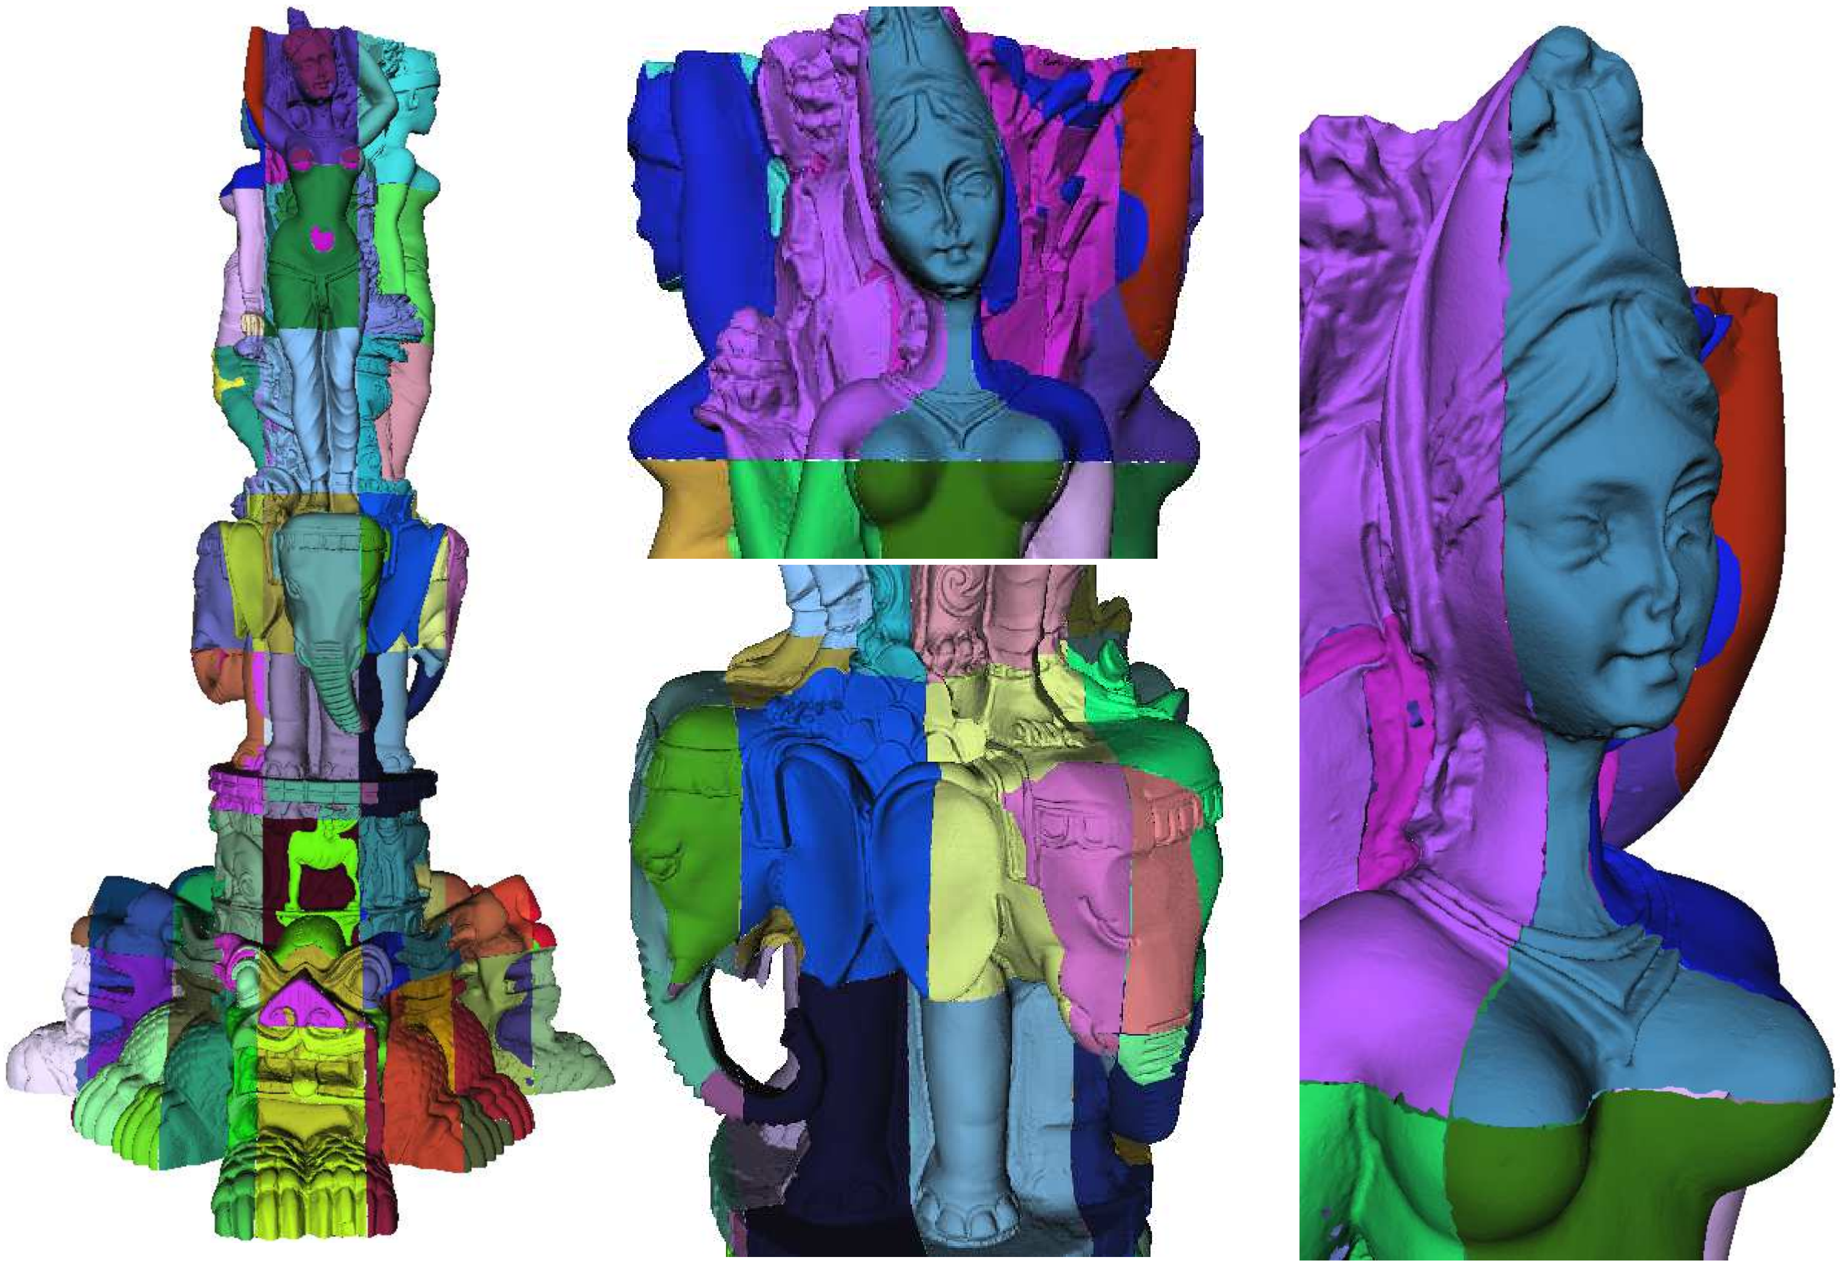

Dealing with sampling quality

Input of non-uniform and low-density sampling:

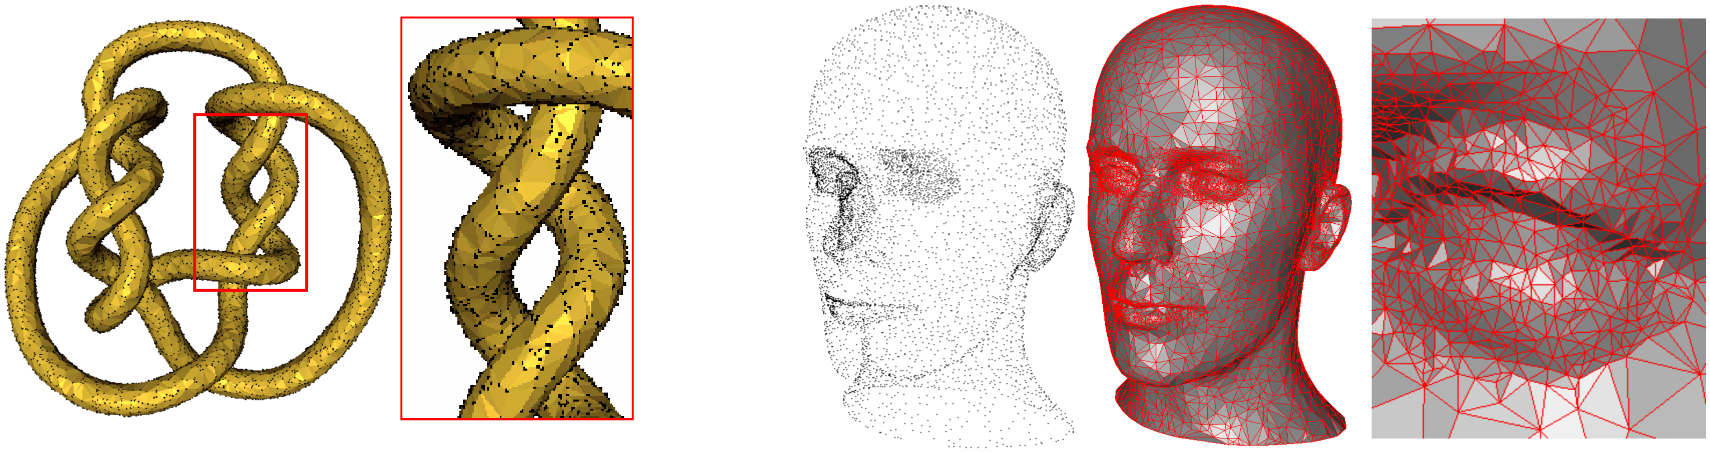

Response to additive noise:

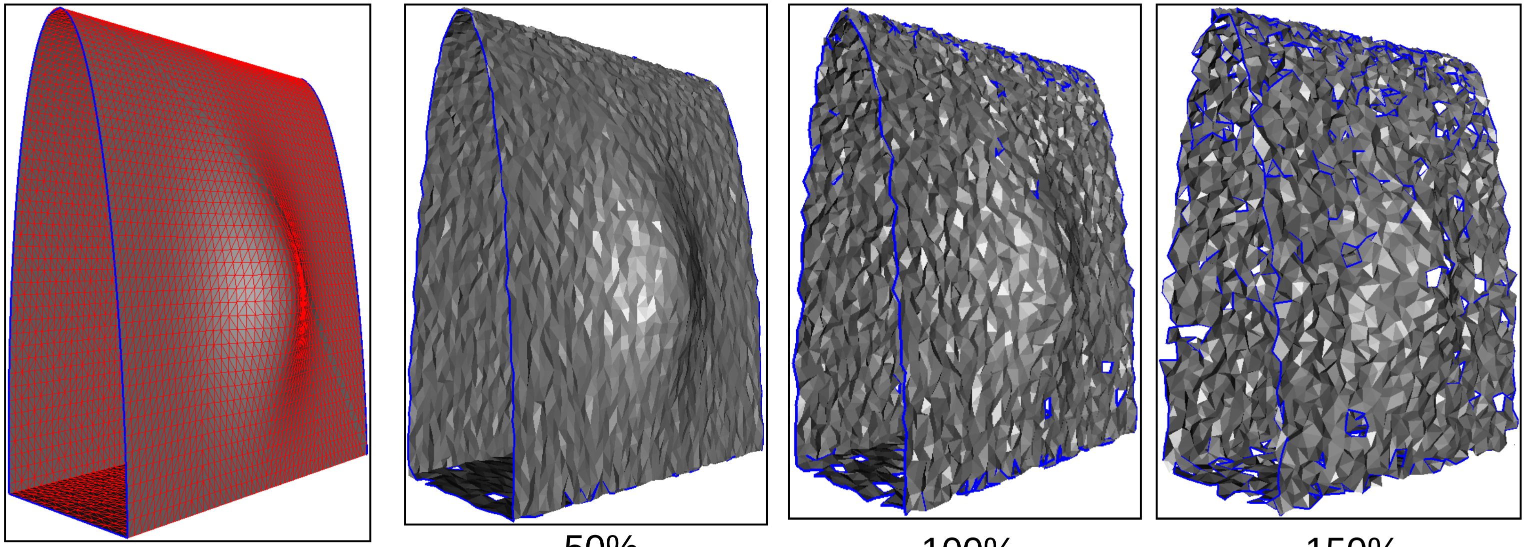

50% 100% 150%

Surface Meshing Running Time

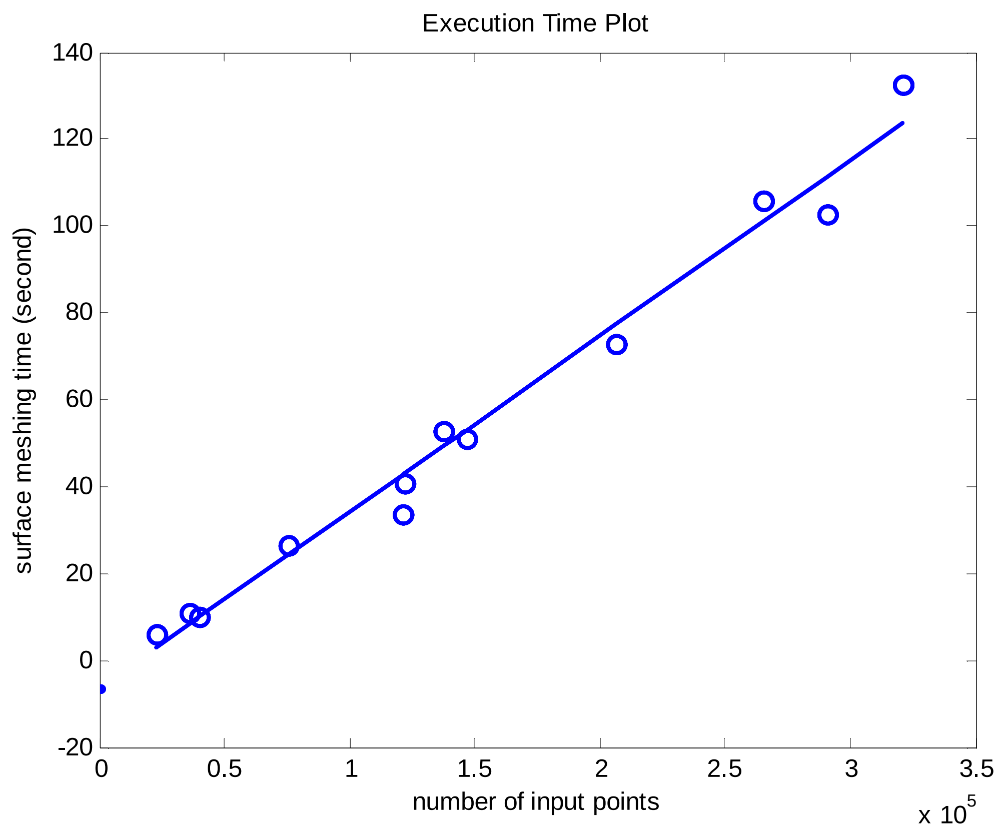

- Roughly linear to the number of samples.
- Performance similar to other recent Delaunay filtering methods.

Conclusions

- Surface reconstruction from point clouds.
  - Handle a great variety of surfaces of practical interest.
  - With little restrictions on input.
- Mesh surface by applying min. cost Gap Transforms in best-first manner, considering:
  - Geometrical suitability of candidate Delaunay triangles.
  - Shock type, shock curve radius.
  - Continuity from neighbors.
  - Mesh topology.
- Multiple-pass greedy algorithm with error recovery.
- Potential to handle arbitrarily large datasets.

Future Work & Discussions

- Additional Shock Transforms to handle *all* shock transitions.
  - Better greedy error recovery.
  - Medial Axis regularization: application to shape manipulation, segmentation, recognition.
- Surface meshing: theoretical guarantees.

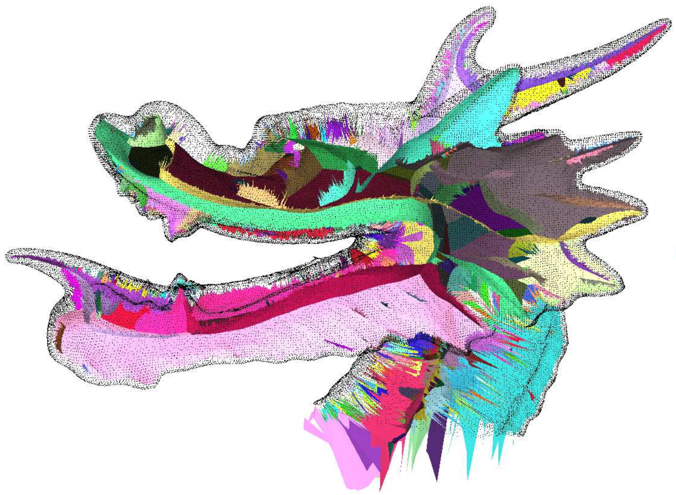

Medial Scaffold (MS+)

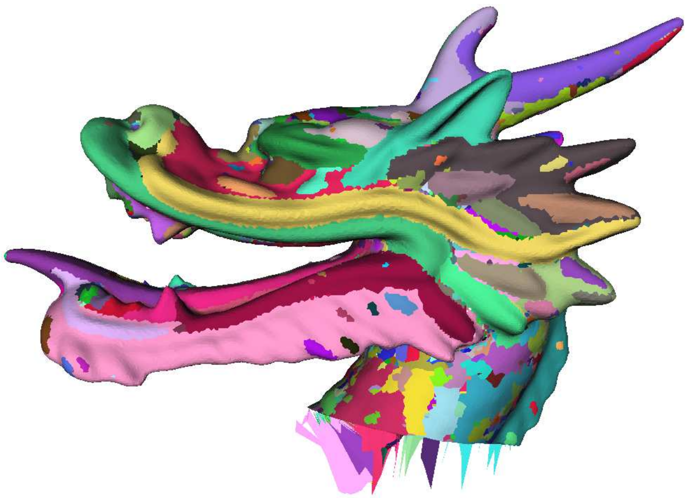

Corresponding surface patches

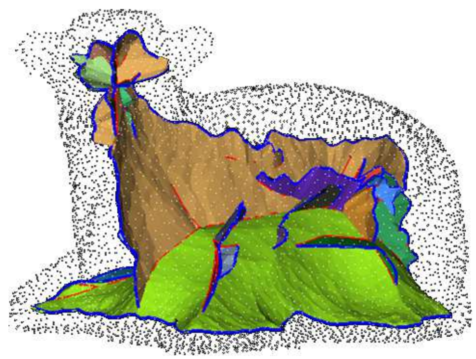

Regularized MS+

Acknowledgments:

Support from NSF. Coin3D (OpenInventor) for visualization/GUI.

Stanford, Cyberware, MPII, Stony Brook archive for 3D data.