Modeling Problems in 3D Photography

Werner Stuetzle
Professor and Chair, Statistics
Adjunct Professor, Computer Science and Engineering
University of Washington

Joint work with A. Certain, T. DeRose, T. Duchamp, M. Eck, H. Hoppe, H. Jin, M. Lounsbery, J.A. McDonald, J. Popovic, K. Pulli, D. Salesin, R. Szeliski

Project funded by NSF (Probability and Statistics, Computational Mathematics, Computer Systems Architecture, Computer Integrated Engineering), and industry contributions.

Outline of Talk

- Introduction to 3D photography
- Statistical aspects of 3D photography
- A specific problem: Fitting triangular meshes
- Other contributions
- Open problems / future work

Introduction to 3D photography

Emerging technology aimed at

- capturing
- viewing
- manipulating

digital representations of shape and visual appearance of 3D objects.

Will have large impact because 3D photographs can be

- stored and transmitted digitally
- viewed on computer screens
- used in computer simulations
- manipulated and edited in software
- used as templates for making electronic or physical copies

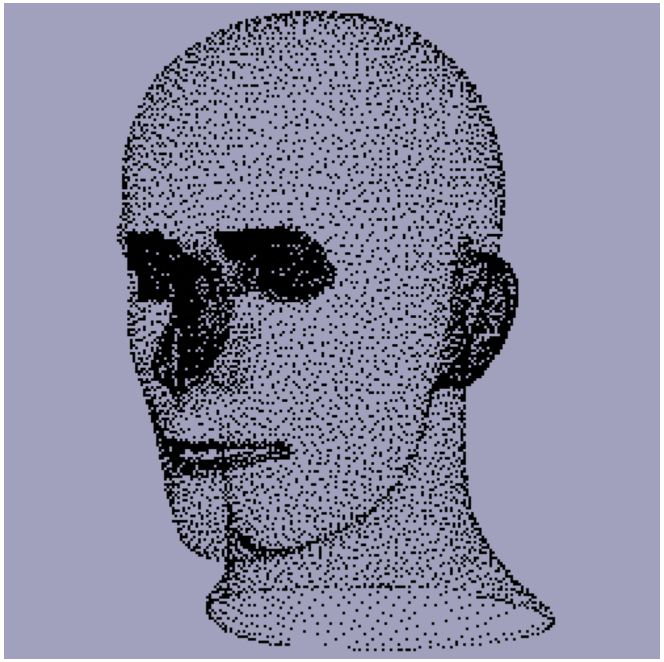

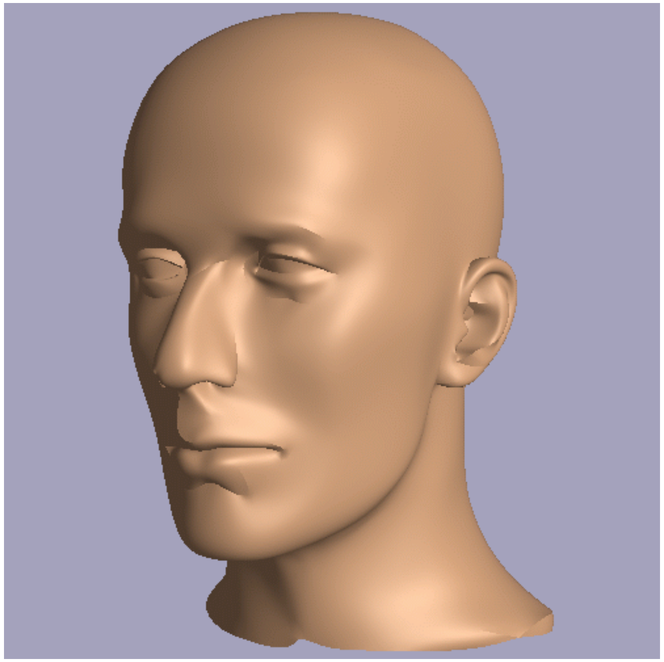

(Potential) Applications

- Reverse engineering
- Industrial design
- Quality control
- Marketing on the Internet
- Plastic surgery
- Mass customization of shoes, pants, shirts, etc.

Industrial design

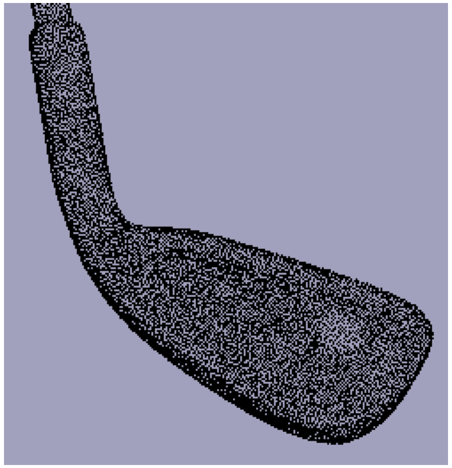

Scan of golf club
(Laser Design)

Computer model

Mass customization of garments

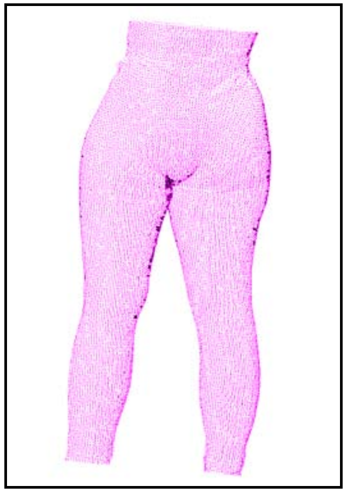

Scan of lower body
(Textile and Clothing Technology Corp.)

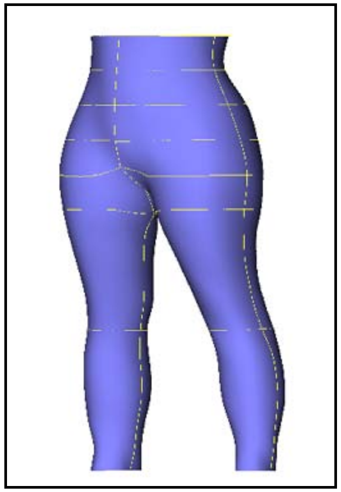

Fitted template
(Dimension curves drawn in yellow)

Components of a 3D Photography System

Sensors to collect data on the shape and "color" of the target object.

Modeling methods to convert raw data into a surface model that can be efficiently viewed and manipulated.

Note:

- Sensors typically produce sample of points scattered around target object. Some also supply RGB values.
- "Color" is a complex property described by *bidirectional reflectance distribution function*

Statistical aspects of 3D photography

Modeling shape $\Rightarrow$ estimating 2D manifold in 3D space from scattered data;

Modeling "color" $\Rightarrow$ Estimating functions on 2D manifolds of arbitrary topological type.

Talk will focus on modeling shape — *manifold estimation*

Modeling "color" is still in its infancy — lots of room for improvement.

Regression vs manifold estimation

The generic regression problem:

Given: Sample $(\underline{x}_1, y_1), \ldots, (\underline{x}_n, y_n)$, $\underline{x}_i \in R^p$, $y_i \in R$, assumed to be generated from the model $y_i = f(\underline{x}_i) + \epsilon_i$, with ϵ_i independent, and $\mathbf{E}(\epsilon_i) = 0$

Goal: Estimate $f(\underline{x}) = \mathbf{E}(Y \mid \underline{x})$.

Regression (or function estimation) is one of the fundamental problems in statistics.

In last 20 years, lots of interest in nonparametric methods that make only very general assumptions ("smoothness") about f: Kernel smoothing, spline smoothing, Cart, Projection Pursuit, ...

Little work on regression for more general domains (sphere).

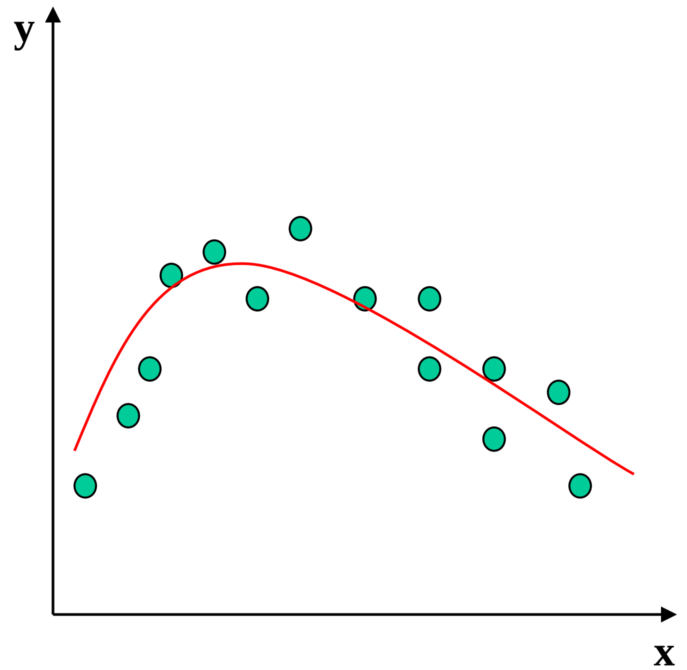

The manifold estimation problem:

Given: Sample $\underline{x}_1, \ldots, \underline{x}_n$ assumed to be generated according to $\underline{x}_i = \gamma_i + \epsilon_i$, with $\gamma_i \in \Gamma$, Γ (2D) manifold.

Goal: Estimate Γ.

More difficult than regression!

Have to choose representation:

- *Parametric*, as the embedding of a suitably chosen domain into R^3, vs
- *Implicit*, as the zero set of a function $R^3 \to R$.

We chose parametric representation.

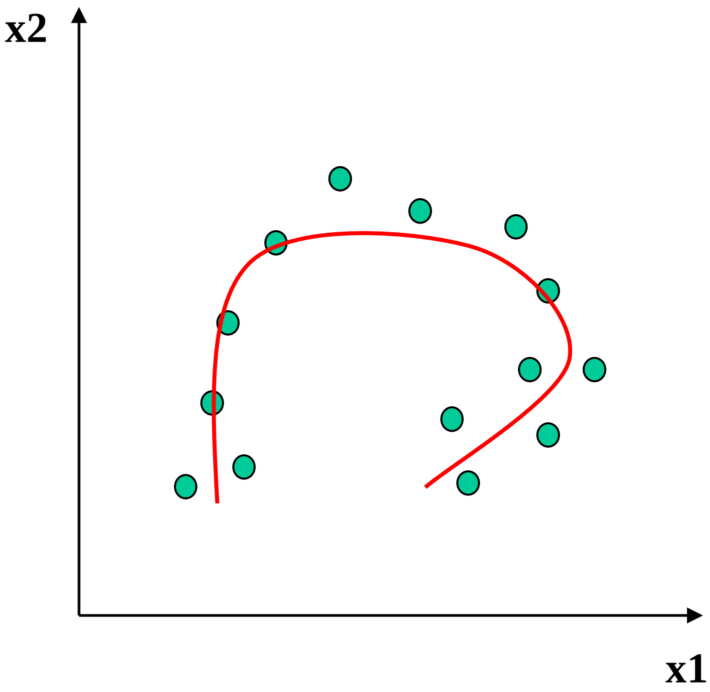

Estimation consists of two steps:

1. Estimate topological type and construct parametrization domain Λ of this type.
2. Estimate the coordinate functions $\underline{x}(\lambda)$ of the embedding.

Model: Observations $\underline{x}_1, \ldots, \underline{x}_n$ generated according to

$$\underline{x}_i = \underline{x}(\lambda_i) + \epsilon_i$$

If $\lambda_1, \ldots, \lambda_n$ were known, then estimating $\underline{x}(\lambda)$ would be a multi-response regression problem over a non-euclidean domain.

In our case, however, $\lambda_1, \ldots, \lambda_n$ are unknown and have to be estimated.

Strategy:

- Find initial, fairly crude surface estimate. Use its topological type as the estimate for the topological type of the manifold.
- Use initial estimate as starting guess for fitting method (analogous to stepwise regression) that generates an accurate and parsimonious surface model.

We have developed two methods two methods to generate initial estimate.

Both generate a triangular mesh (manifold consisting of planar triangles glued together along their edges)

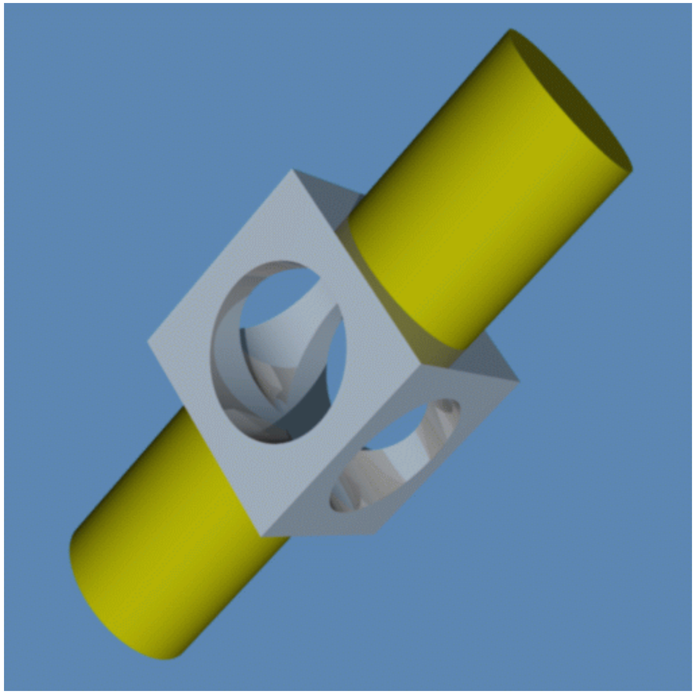

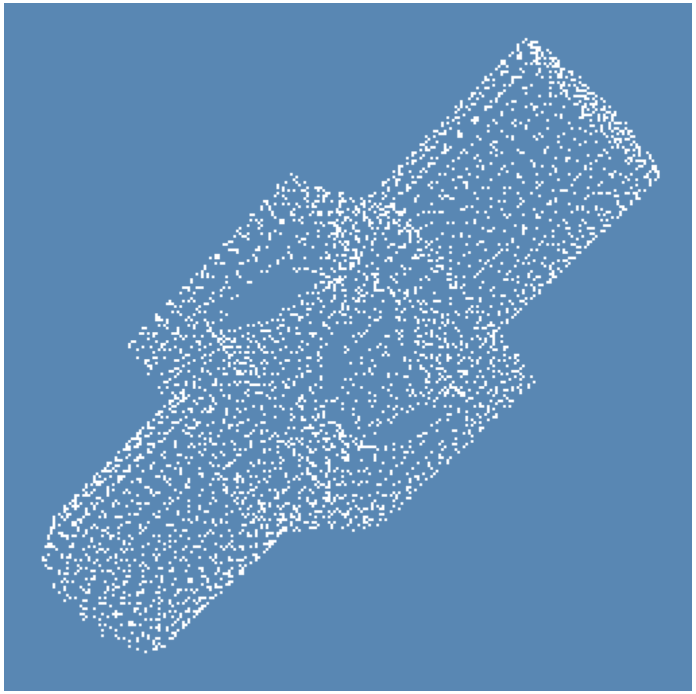

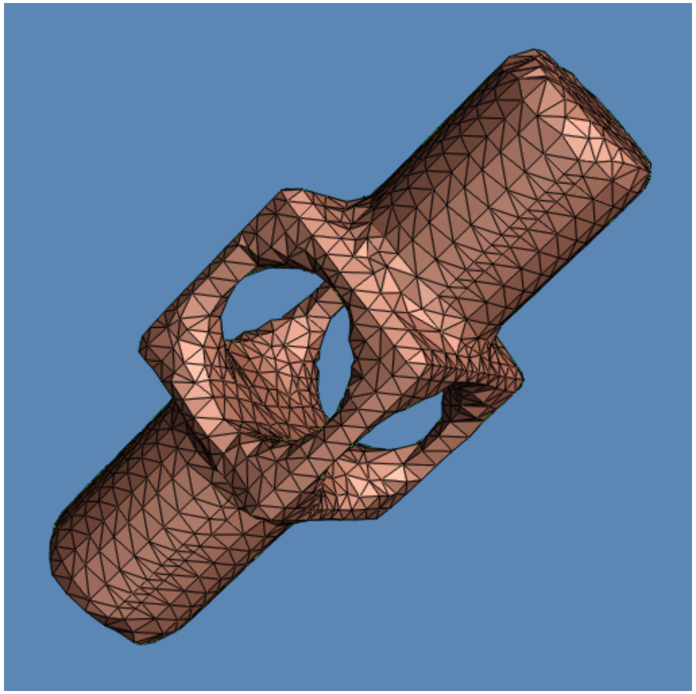

Top left: (simulated) mechanical part
Top right: point sample
Bottom: initial surface estimate

Fitting triangular meshes

Given:

- Points $\underline{x}_1, \ldots, \underline{x}_n$ scattered around unknown manifold Γ.
- Initial estimate for Γ in the form of a triangular mesh M_0.

Goal: Modify initial mesh M_0 to

- improve fit to data, and
- (possibly) reduce the number of faces

while leaving the topological type fixed.

Outline of mesh fitting

- Define *energy* of a mesh M, incorporating both lack of fit and complexity. Contribution of complexity determined by *vertex cost* c_{rep}.
- Minimize energy over number, connectivity, and position of the vertices.

Large value of c_{rep} results in optimal mesh with few vertices.

Representation of meshes

A mesh is a pair $M = (K, V)$:

- The simplicial complex K determines the connectivity of the vertices, edges, and faces of the surface.
- The vertex positions $V = (\underline{v}_1, \ldots, \underline{v}_m)$ determine the geometric shape of the surface in R^3.

Recall:

- K is a collection of subsets of $\{1, \ldots, m\}$
- $\{1\}, \ldots, \{m\} \in K$
- $\{i, j\} \in K$ iff $\underline{v}_i$ and $\underline{v}_j$ are connected by edge in S.
- $\{i, j, k\} \in K$ iff $\underline{v}_i, \underline{v}_j, \underline{v}_k$ are vertices of a face of S

K characterizes topological type of S.

Definition of energy

$$
\begin{aligned}
E(K,V) &= E_{dist}(K,V) + E_{rep}(K) + E_{spring}(K,V) \\
&= \sum_{i=1}^{n} d^2(\underline{x}_i, \Phi_V(|K|)) \\
&\quad + c_{rep}\, m \\
&\quad + \sum_{\{j,k\}\in K} \kappa \|\underline{v}_j - \underline{v}_k\|^2
\end{aligned}
$$

Motivation for spring energy E_{spring}:

- guarantees existence of minimum
- prevents parts of the surface from wandering away from the data
- tends to prevent self-intersection

Energy minimization

Goal: Minimize

$$E(K,V) = E_{dist}(K,V) + E_{rep}(K) + E_{spring}(K,V)$$

over the set $\mathcal{K}$ of simplicial complexes homeomorphic to K_0, and the vertex positions V.

Method: Decompose problem into two nested sub-problems:

- inner optimization over V fixed K
- outer optimization over $K \in \mathcal{K}$.

Optimization over vertex positions

Define *geometric realization* $|K| \subset R^m$:

- simplicial surface in R^m
- obtained by identifying 1-simplices $\{1\}, \ldots, \{m\}$ with $\underline{e}_1, \ldots, \underline{e}_m$.

Define linear map $\Phi_V : R^m \to R^3$ by

$$\Phi_V(\underline{e}_i) = \underline{v}_i, \;\; i = 1, \ldots, m.$$

For every point $\underline{y} \in S$ there is unique point $\underline{b} \in |K|$ with

$$\Phi_V(\underline{b}) = \underline{y}$$

$\underline{b}$: barycentric coordinates of $\underline{y}$

Note: $\underline{b}$ has at most 3 non-zero components.

Simplicial complex K

vertices: $\{1\}$, $\{2\}$, $\{3\}$

edges: $\{1, 2\}$, $\{2, 3\}$, $\{1, 3\}$

faces: $\{1, 2, 3\}$

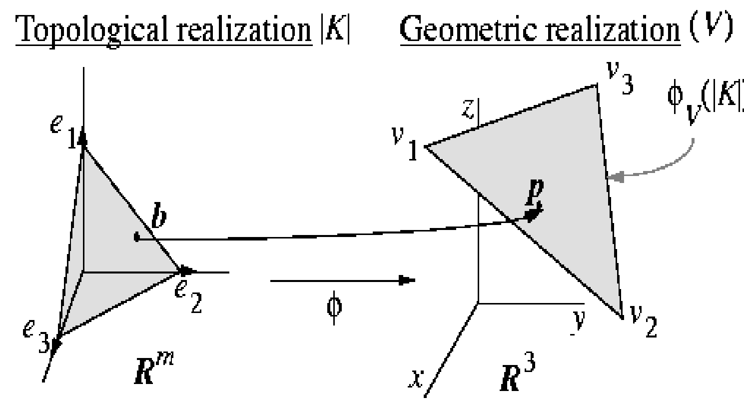

Restatement of fitting problem

Find vertices V to minimize

$$E(V) = \sum_{i=1}^{n} d^2(\underline{x}_i, \Phi_V(|K|)) + \sum_{\{i,j\}\in K} \kappa \|\underline{v}_i - \underline{v}_j\|^2$$

Note: $d^2(\underline{x}_i, \Phi_V(|K|))$ is itself solution of optimization problem:

$$\begin{aligned} d^2(\underline{x}_i, \Phi_V(|K|)) &= \min_{\underline{b}_i \in |K|} \|\underline{x}_i - \Phi_V(\underline{b}_i)\|^2 \\ &= \min_{\underline{b}_i \in |K|} \|\underline{x}_i - \sum_{j=1}^{m} b_{ij}\underline{v}_j\|^2 \end{aligned}$$

Restate fitting problem: Minimize new objective function

$$E(V,B) = \sum_{i=1}^{n} \|\underline{x}_i - \Phi_V(\underline{b}_i)\|^2 + \sum_{\{j,k\}\in K} \kappa \|\underline{v}_j - \underline{v}_k\|^2$$

over vertex positions $\underline{v}_1, \ldots, \underline{v}_m$ and barycentric coordinates $\underline{b}_1, \ldots, \underline{b}_n$.

Optimization method

Want to minimize

$$E(V,B) = \sum_{i=1}^{n} \|\underline{x}_i - \Phi_V(\underline{b}_i)\|^2 + \sum_{\{j,k\}\in K} \kappa \|\underline{v}_j - \underline{v}_k\|^2$$

over vertex positions $\underline{v}_1, \ldots, \underline{v}_m$ and barycentric coordinates $\underline{b}_1, \ldots, \underline{b}_n$.

Suggests alternating minimization scheme:

- For fixed $\underline{v}$'s, can find optimal $\underline{b}$'s by projection
- For fixed $\underline{b}$'s, can find optimal $\underline{v}$'s by solving 3 *linear* LS problems

Note:

- can use continuity between iterations in projection step
- LS problems are large but sparse — use conjugate gradients

Optimization over simplicial complexes

$E(K) = \min_V E(K, V)$: Energy of the optimal embedding of K.

Goal: Minimize $E(K)$ over the simplicial complexes $K \in \mathcal{K}$.

Method: Define three *elementary operations* on simplicial complexes:

1. Edge collapse
2. Edge split
3. Edge swap

Legal move: elementary operation leaving the topological type of K unchanged.

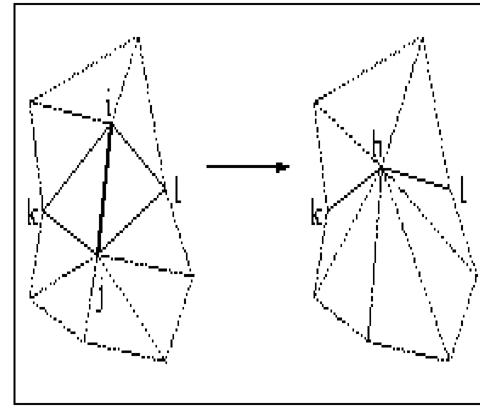

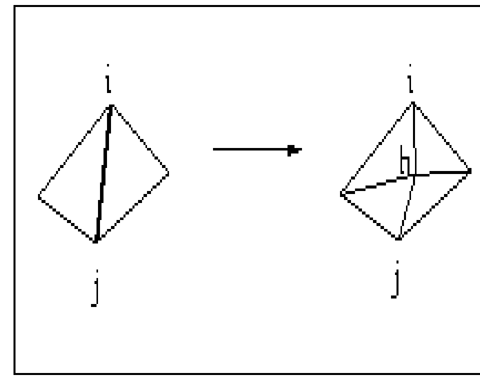

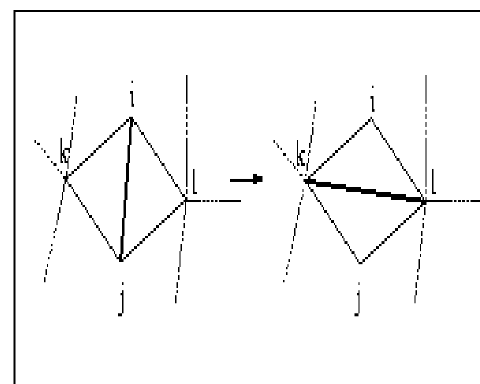

Prop: Every simplicial complex $K \in \mathcal{K}$ can be reached from K_0 through a sequence of legal moves.

Search for a sequence of legal moves leading from K_0 to a local minimum of $E(K)$

Random descent:

- Randomly select a legal move $K \Rightarrow K'$
- If $E(K') < E(K)$ accept the move, otherwise repeat the selection
- If a large number of trials fails to reduce $E(K)$, terminate the search

Result of mesh fitting

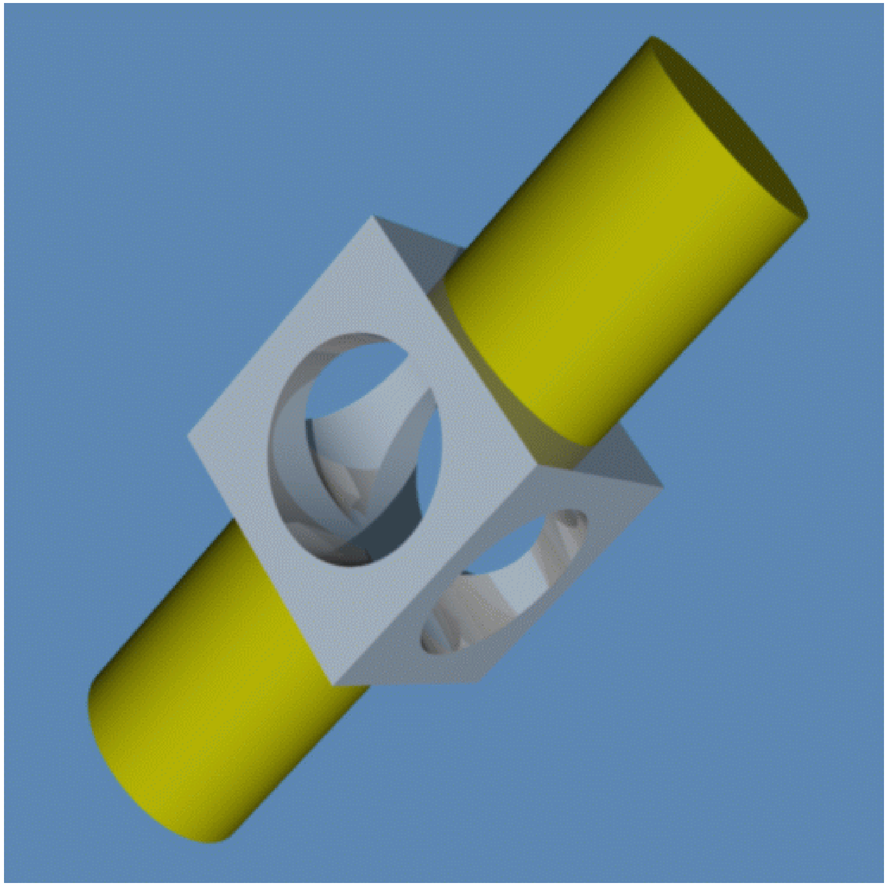

Original

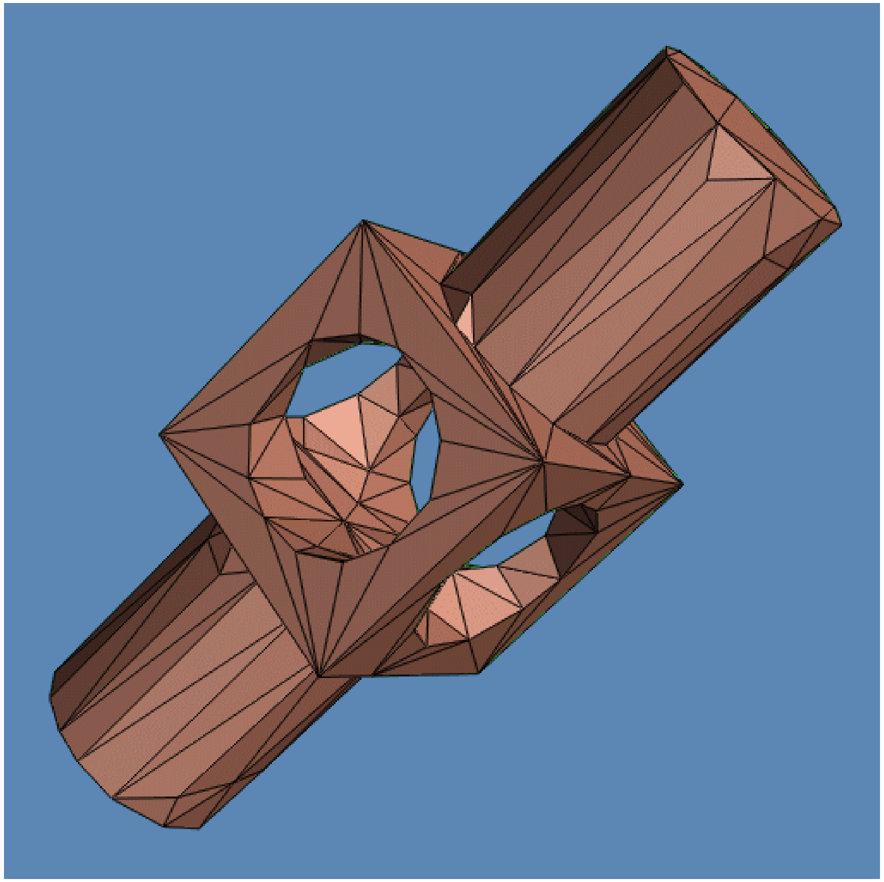

Fitted mesh

Result of mesh fitting

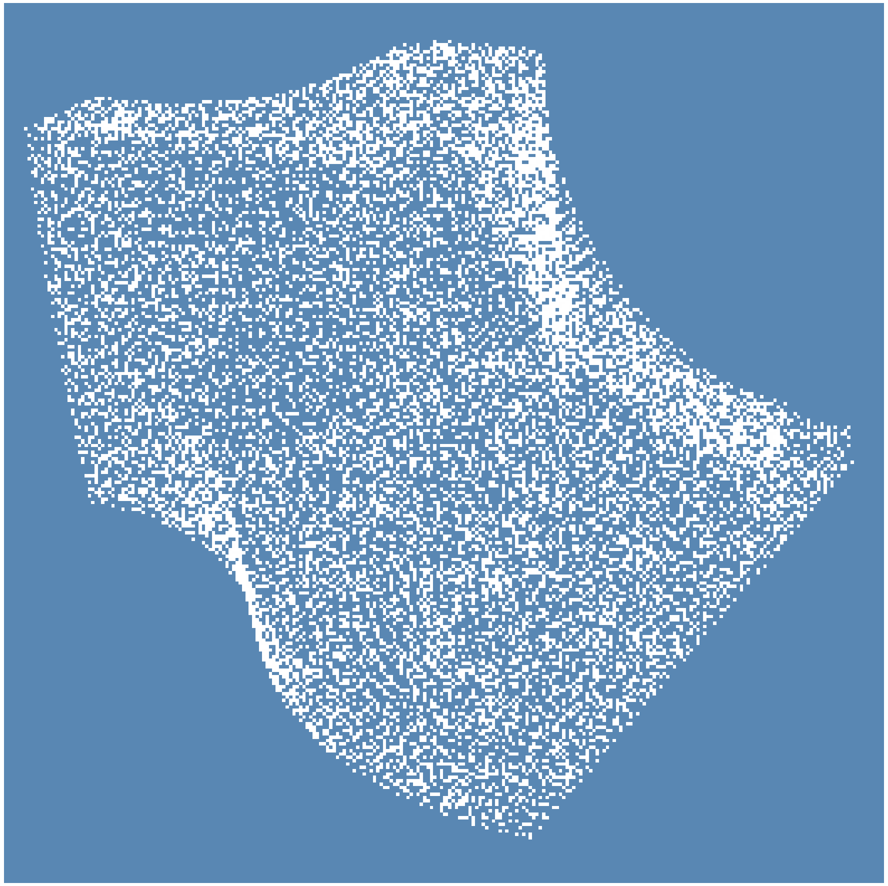

Scan of fan disc
(Laser Design)

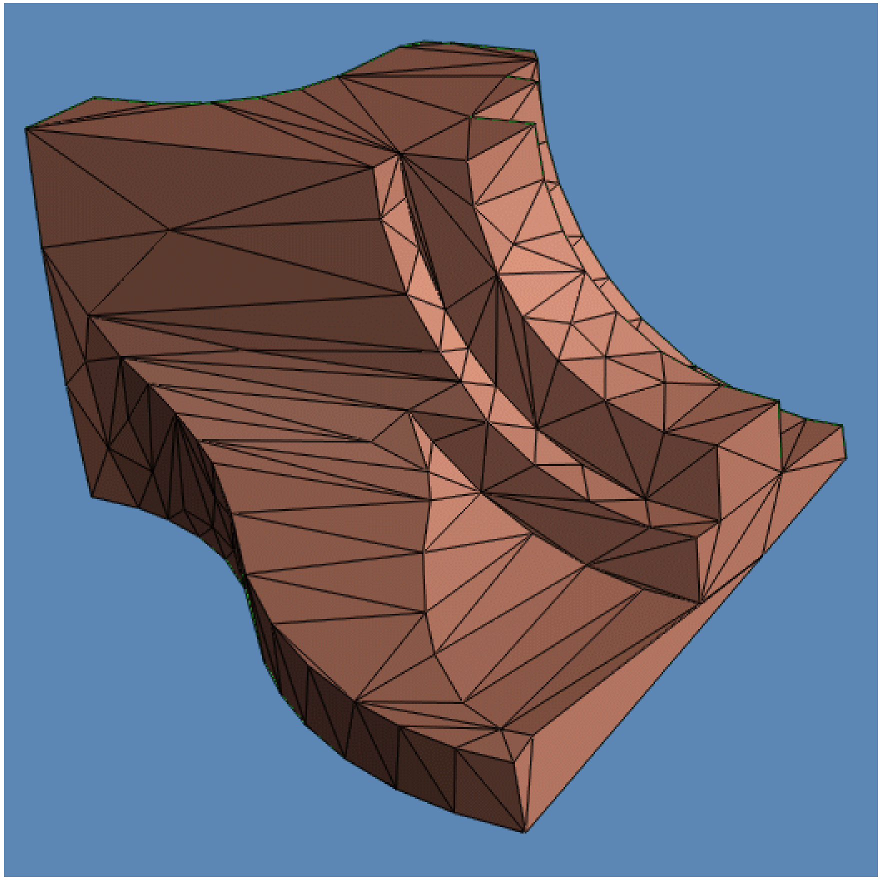

Fitted mesh

Result of mesh fitting

Scan of golf club head
(Laser Design)

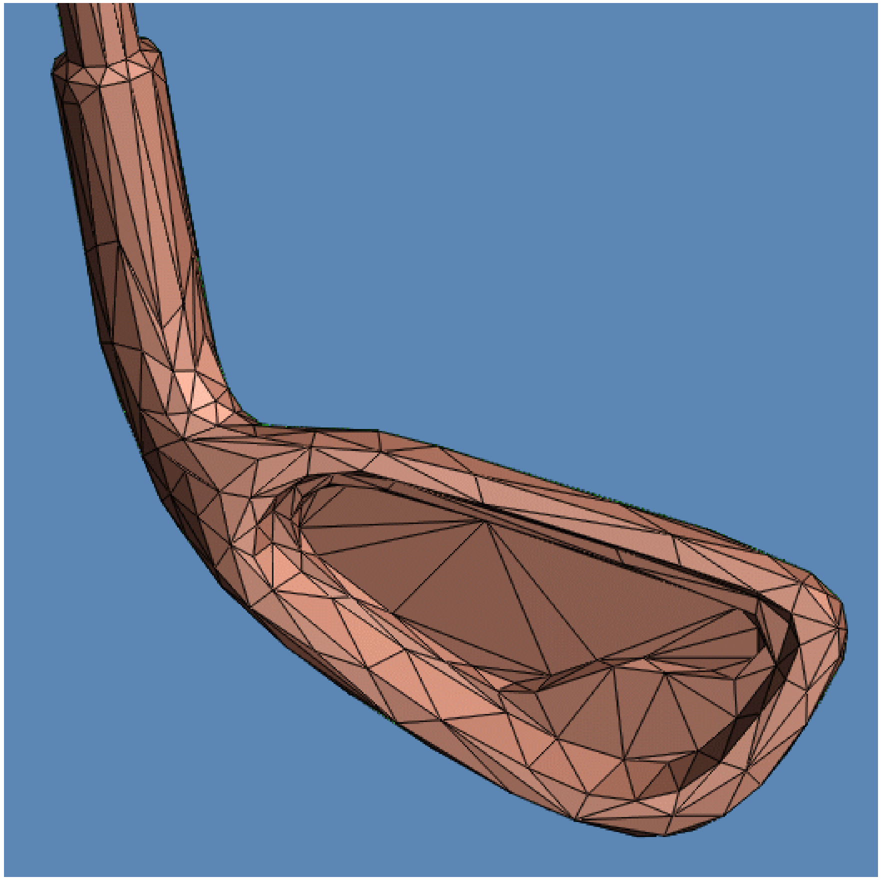

Fitted mesh

Other contributions

Modeling manifolds with piecewise smooth subdivsion surfaces

Motivation

- Many objects of interest are piecewise smooth
- Piecewise smooth models can lead to more accurate and parsimonious representation than meshes

Contributions:

- Proposed a new surface representation — piecewise smooth subdivision surfaces.
- Devised a fitting algorithm that estimates the number of vertices and connectivity of control mesh, the positions of the vertices and the presence and location of sharp features.

Fitting piecewise smooth subdivision surface

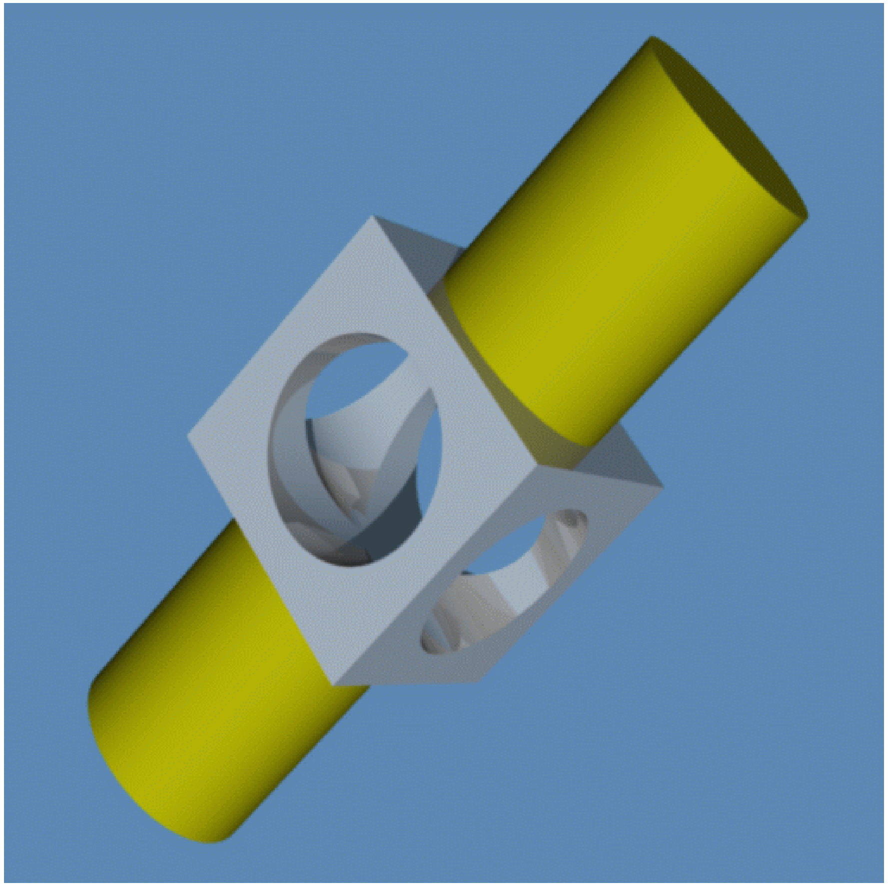

Original

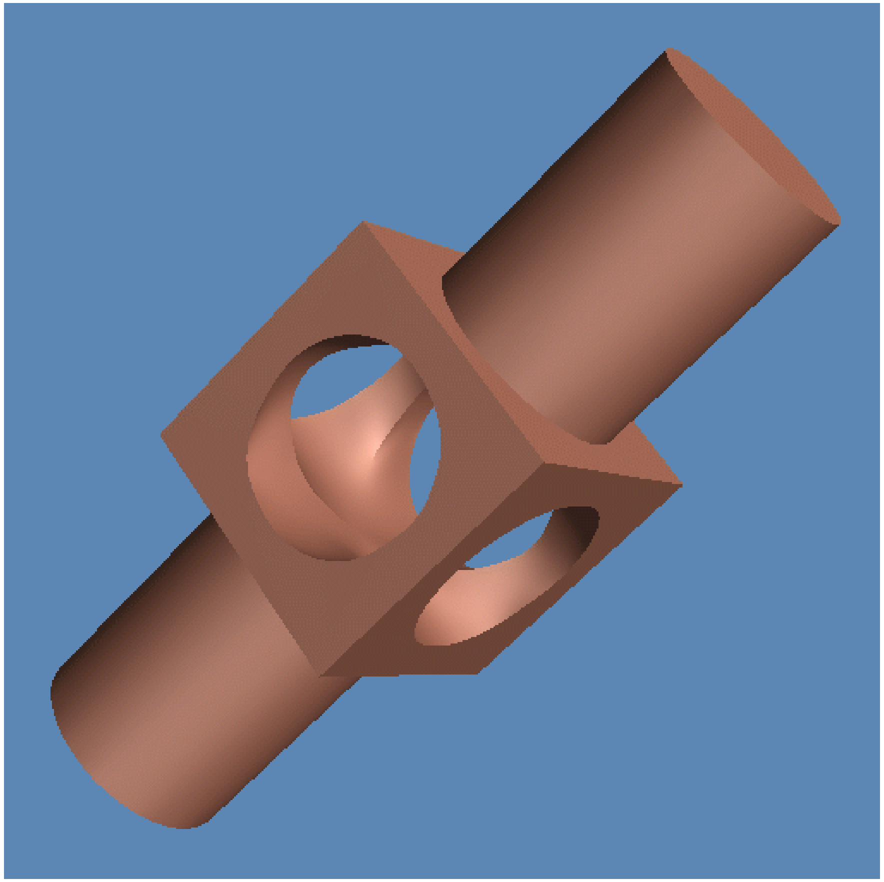

Subdivision surface

Fitting piecewise smooth subdivision surface

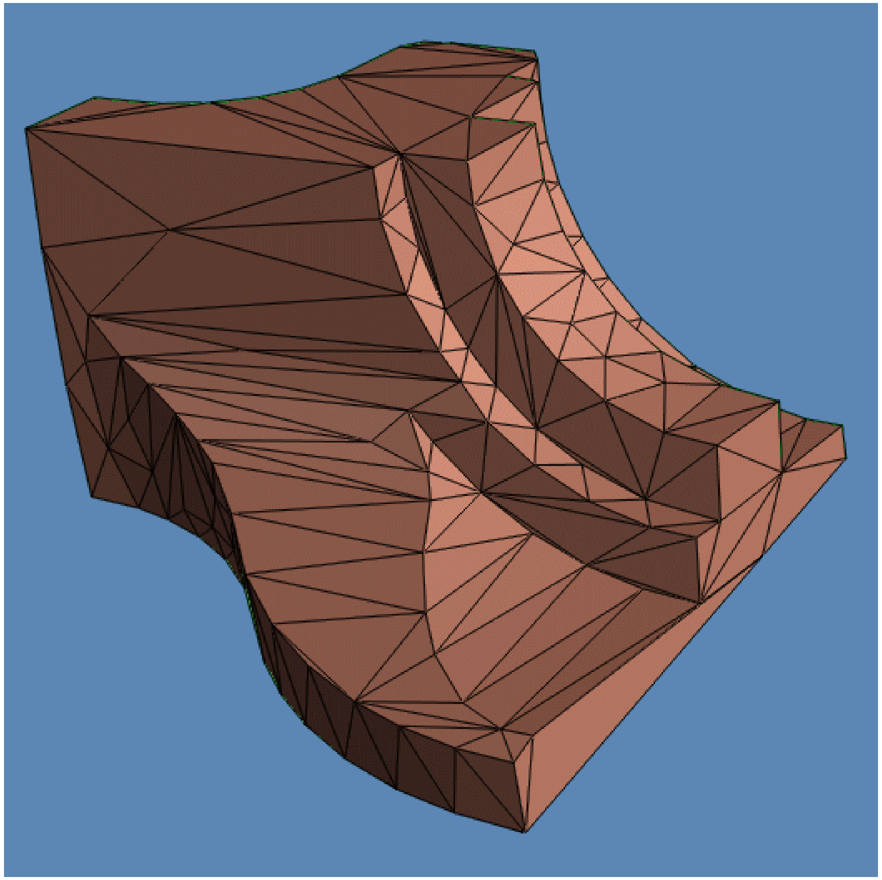

Mesh fitted to fan disc scan

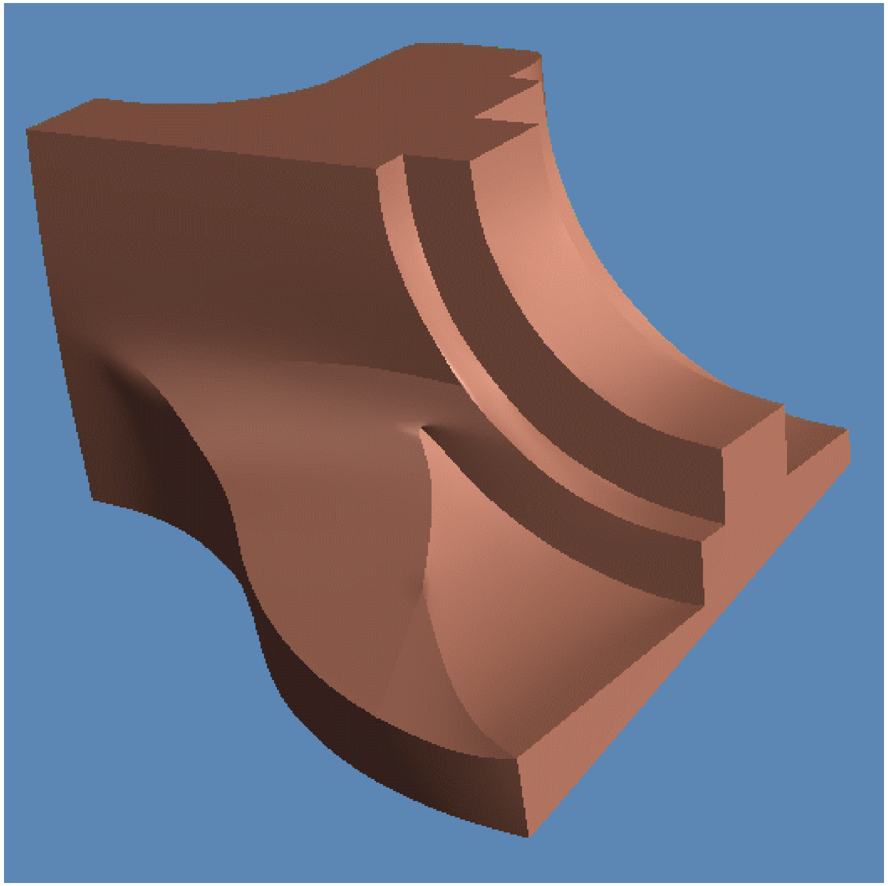

Subdivision surface

Fitting piecewise smooth subdivision surface

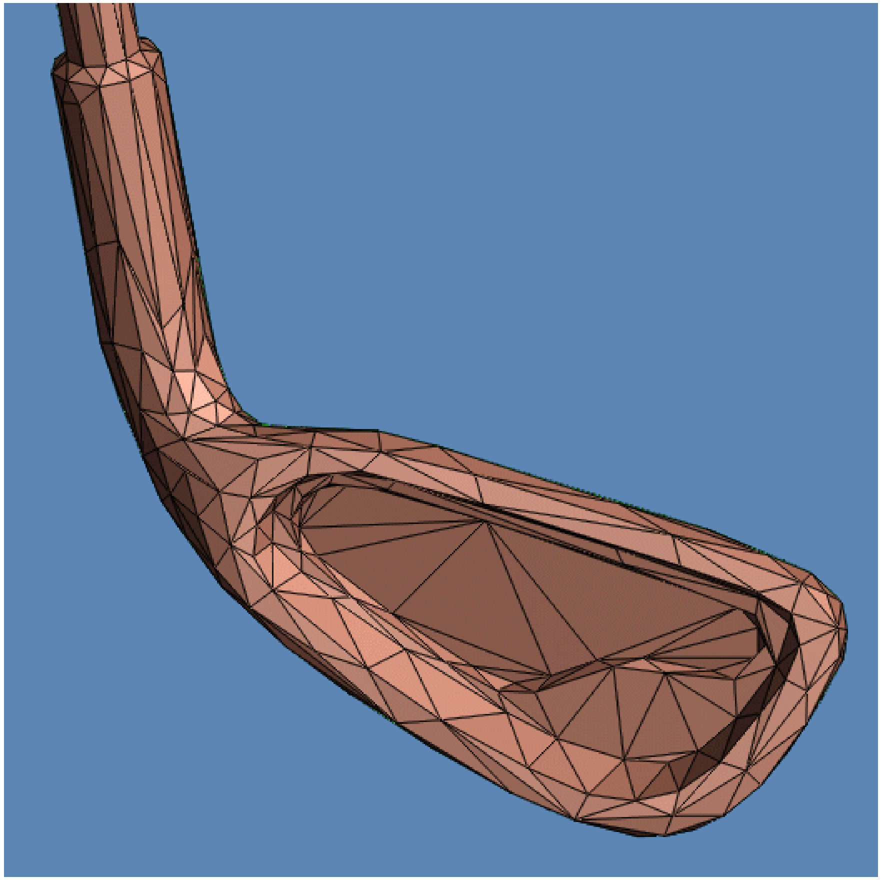

Fitted mesh

Subdivision surface

Multiresolution analysis of meshes

MRA is a widely used tool in signal processing, image analysis.

Supports denoising (Donoho and Johnston's *Waveshrink*), compression, progressive transmission.

Problem with applying MRA to meshes:

MRA is a tool applicable to functions, whereas meshes are manifolds.

Solution

- Parameterize the mesh – construct:
  - Parametrization domain K_0
  - Geometry function $f^{geom} : K_0 \mapsto R^3$
  - Color function $f^{color} : K_0 \mapsto R^3$
- Expand the color and geometry functions into wavelets.

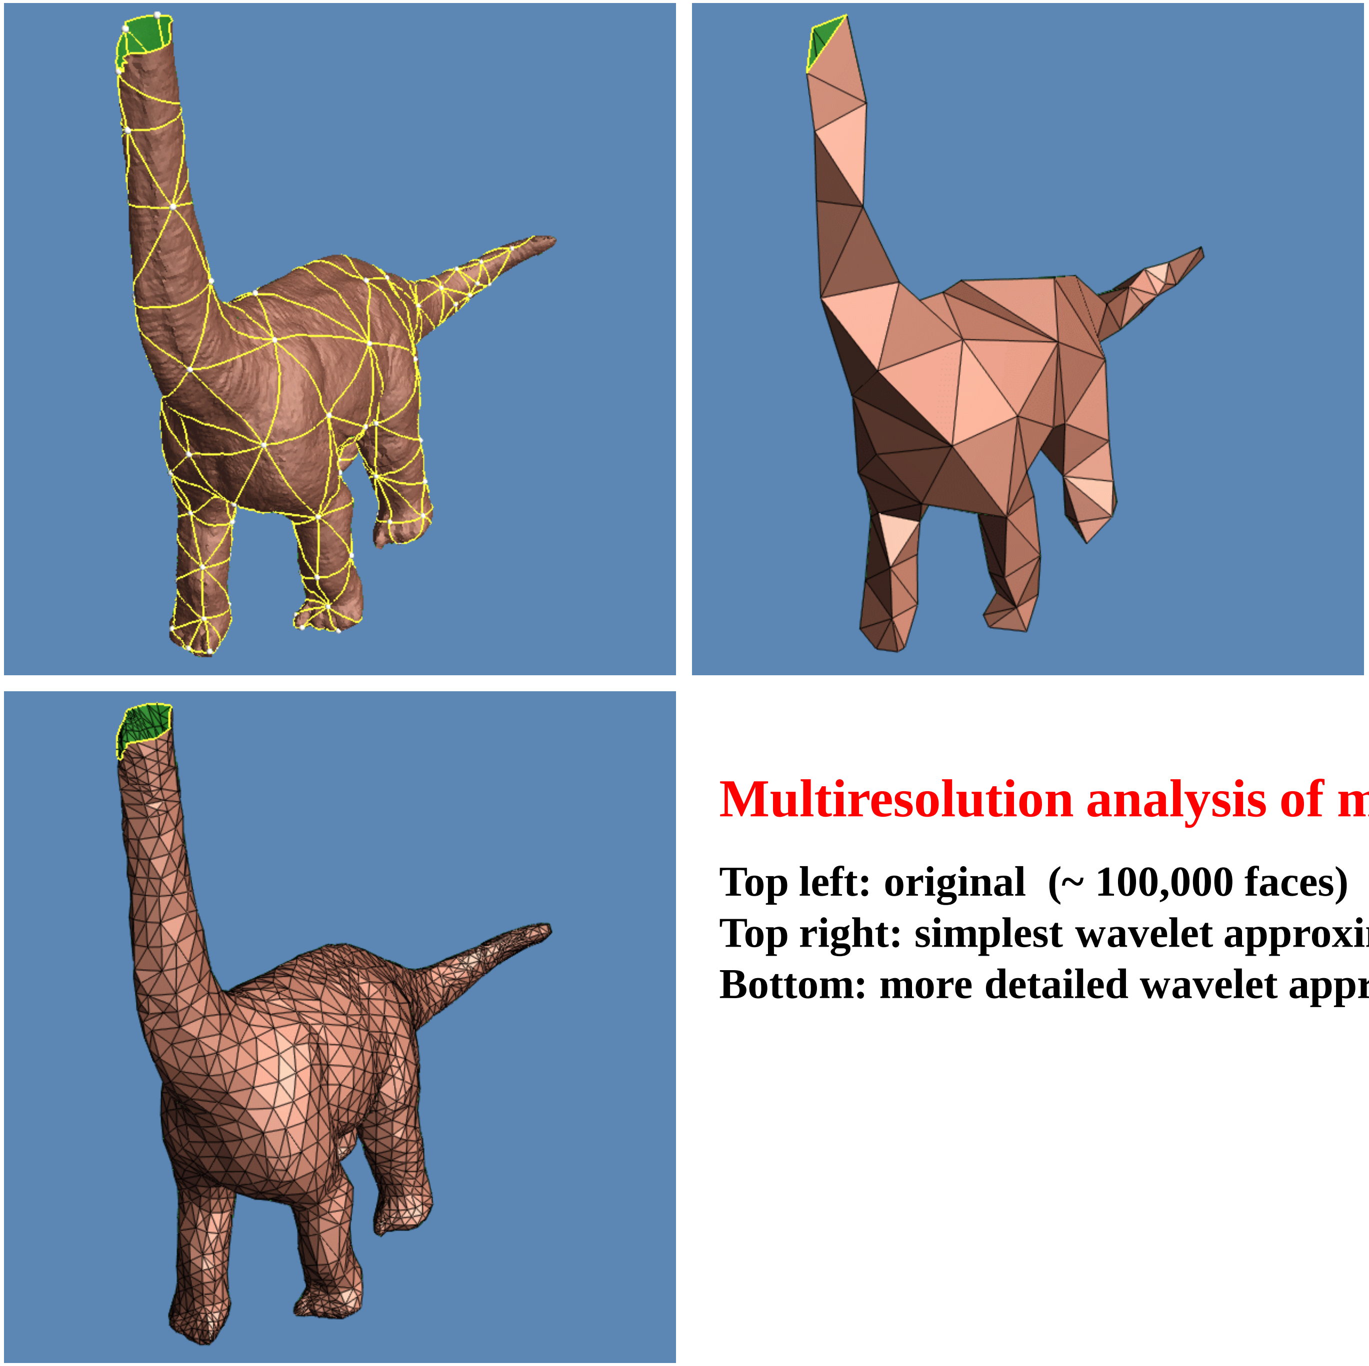

Multiresolution analysis of meshes

Top left: original (~ 100,000 faces)
Top right: simplest wavelet approximation
Bottom: more detailed wavelet approximation

Multiresolution analysis of meshes supports

- Compression
- Progressive transmission
- Level-of-detail control
- Performance tuned viewing
- Texture mapping

in *simple, unified* and *theoretically sound* way.

Future work

- Automatic choice of model complexity in manifold estimation (analogous to C_p, cross-validation, little Bootstrap)
- "Spline smoothing" on manifolds
- Using high-resolution color images to get high-fidelity models of target objects

Clearly, lots of room to

- Develop new statistical methodology and theory, and
- Apply it to a important and challenging technology area.

Thank you for your patience

Sensors (Laser scanners, shape cameras)

Simplest idea: Active light scanner using triangulation