Structure Discovery in Sampled Spaces

Leonidas Guibas
Computer Science

Gunnar Carlsson
Mathematics

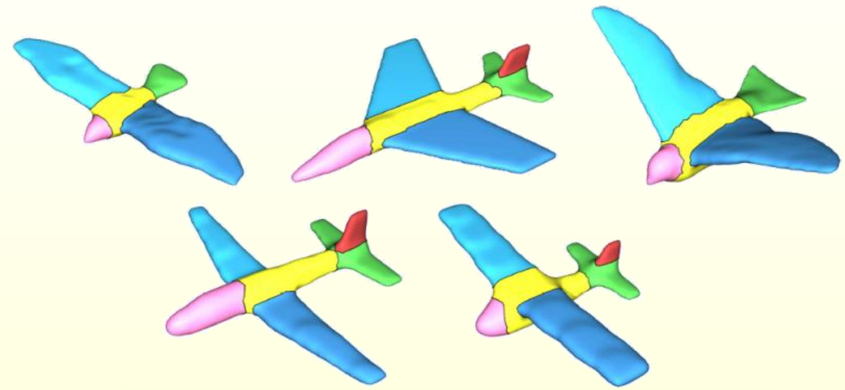

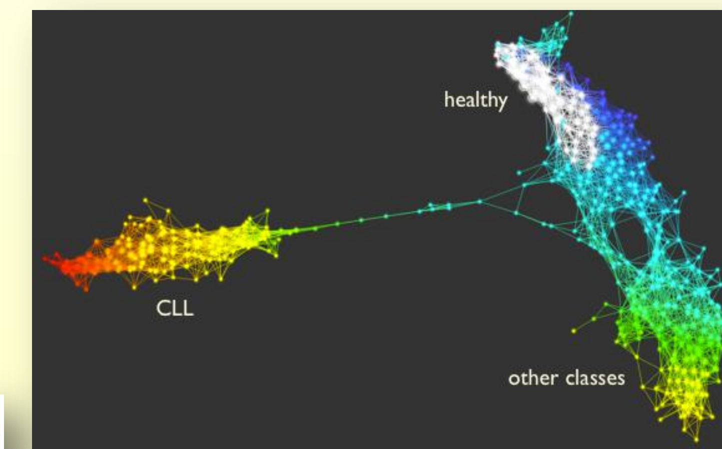

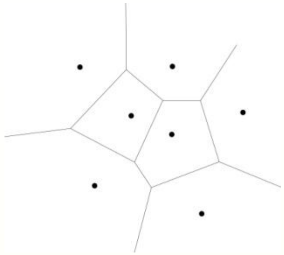

Project Goals

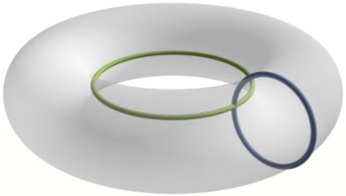

- Bring tools from Computational Geometry and Topology to the analysis and visualization of massive, distributed data sets
- Perform global structure discovery on such data
- Exploit this discovered structure in enabling visual exploration and human interaction with the data

- Recent focus: relationships between data sets

Finding Correspondences and Maps Between Data Sets

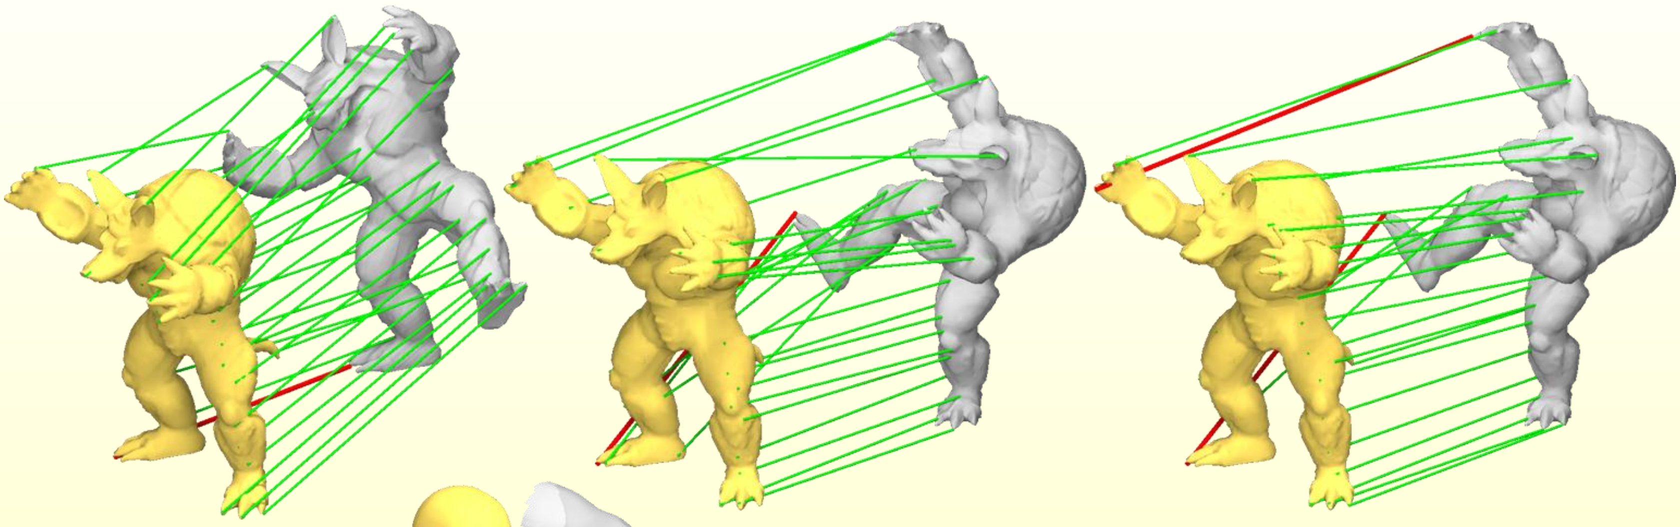

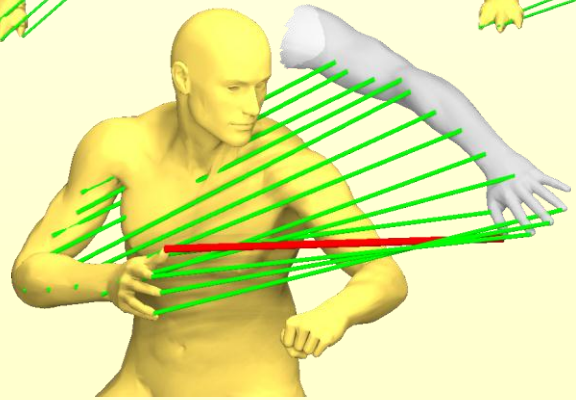

Typically, by solving an optimization problem

Understanding Data via Maps

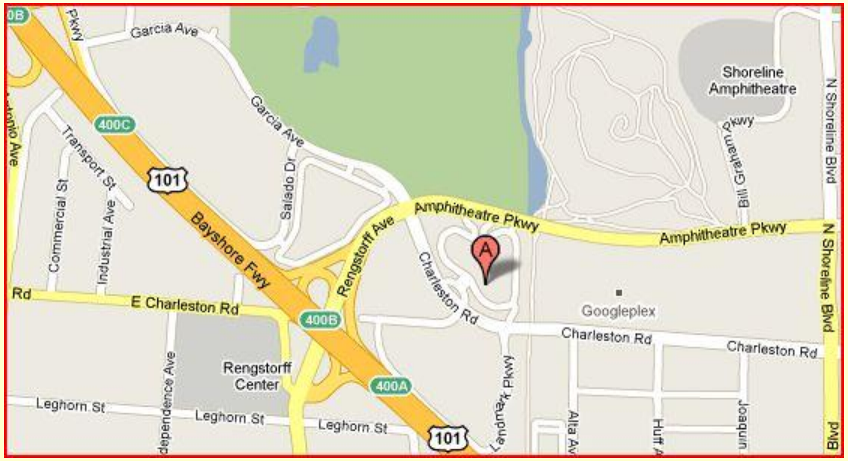

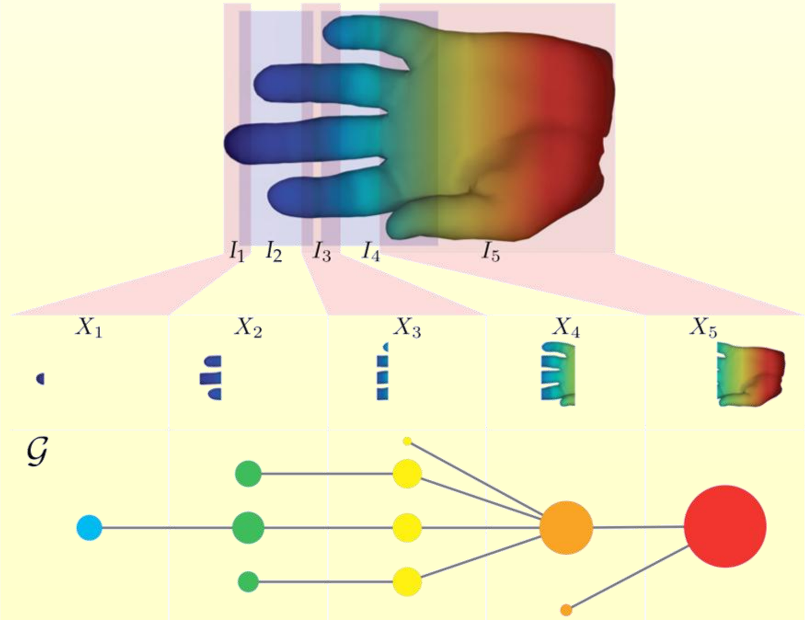

Maps, at What Scale?

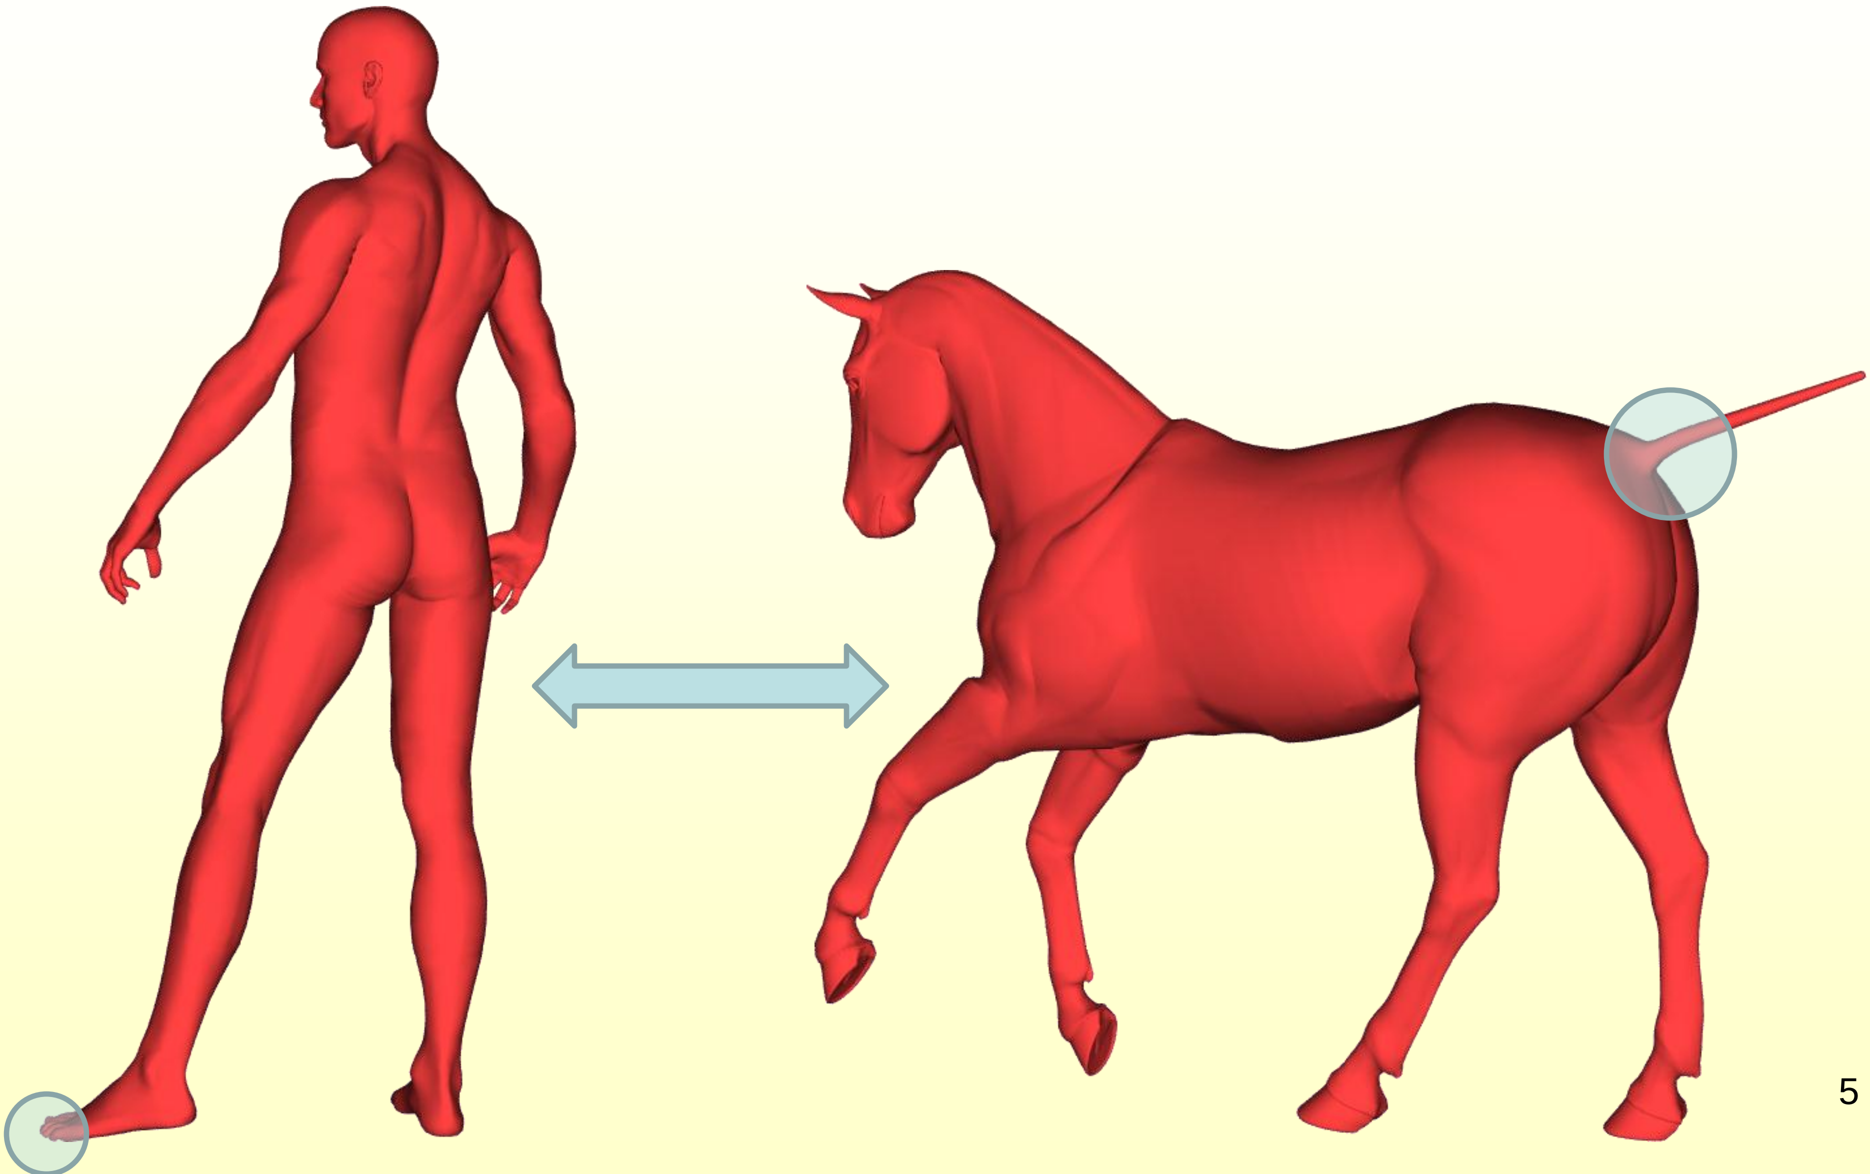

Joint Understanding Goals

- To understand the relationships between data sets, pairwise as well as in higher order combinations
- To extract the shared structure as well as the variability across the entire collection

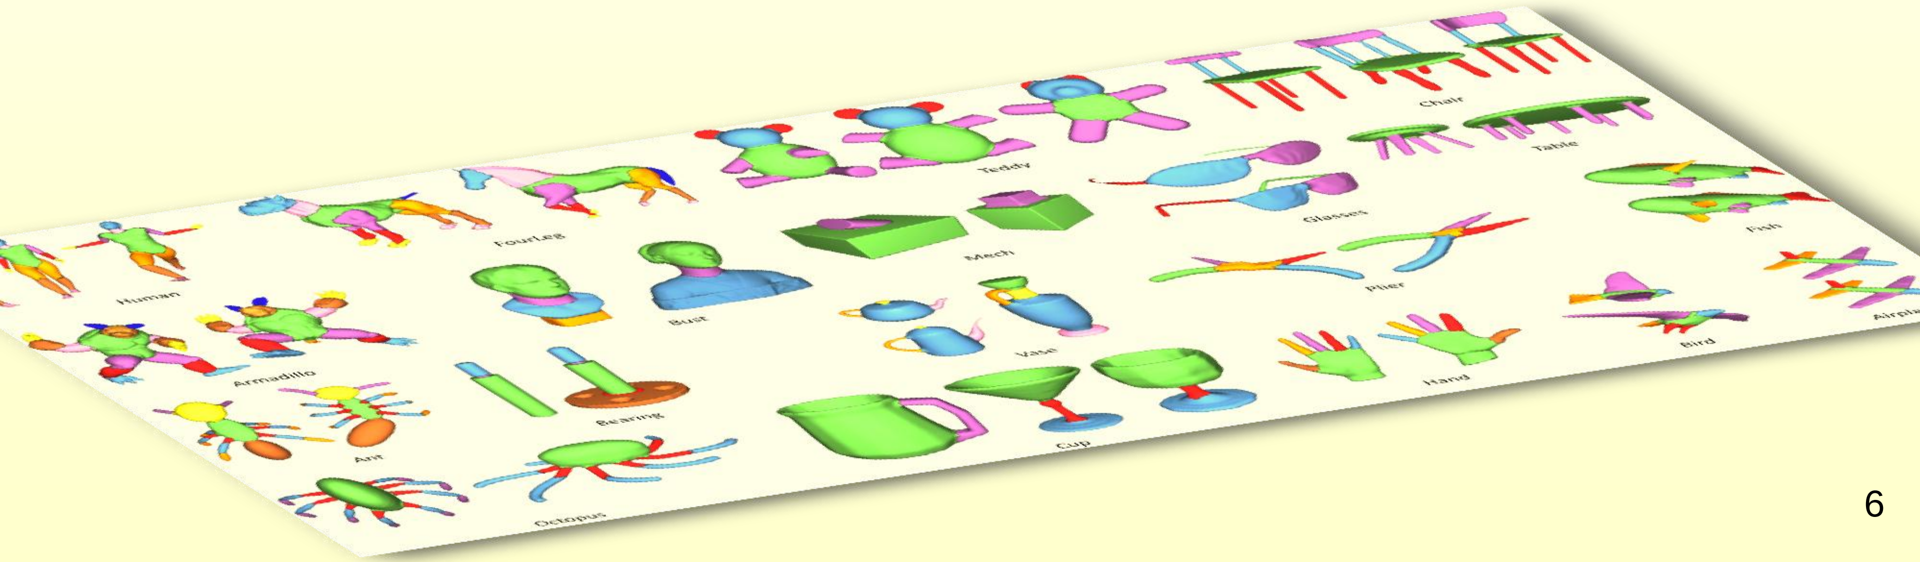

Talk Outline

- Multi-way data set relationships
  - Consistent segmentation — Poster
  - Map networks — Poster
  - Shape space navigation
- Other topics
  - Fuzzy maps — Poster
  - Metric reconstruction — Poster
  - Mapper cancer data analysis
- PNNL collaboration

Topology

Spectral Methods

Machine Learning

Geometry

Optimization

3D scans and meshes

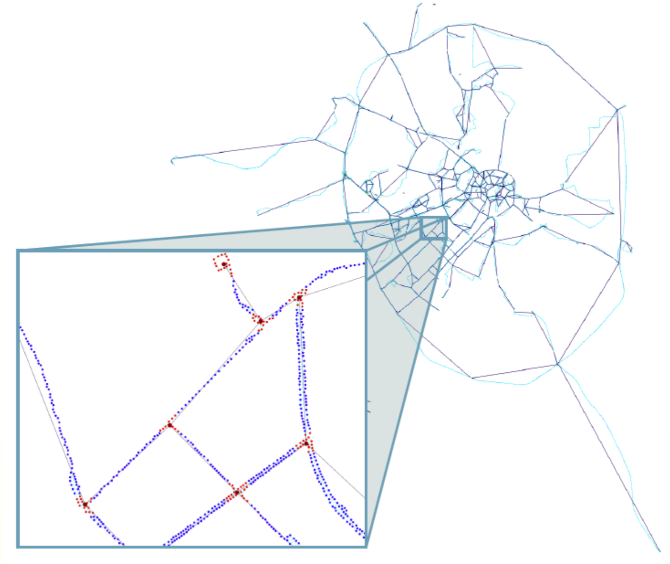

GPS vehicle traces

Diverse Data Sets

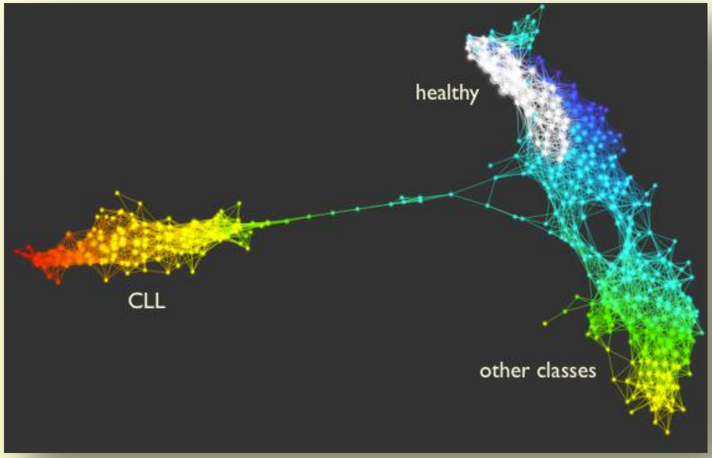

Microarray genomic data

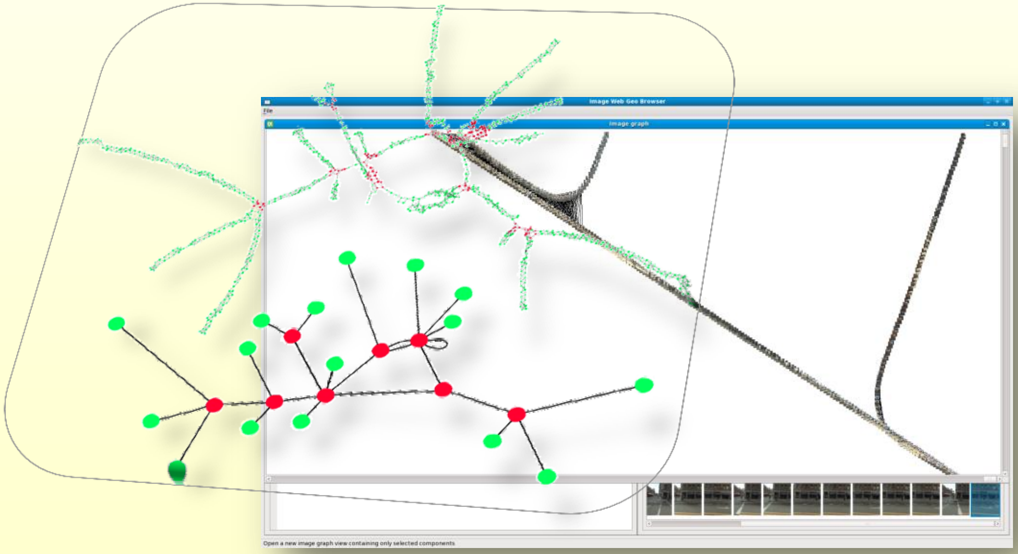

Image collections

Poster

Joint Shape Segmentation via Linear Programming

[with Huang, Koltun, SiggraphAsia '11]

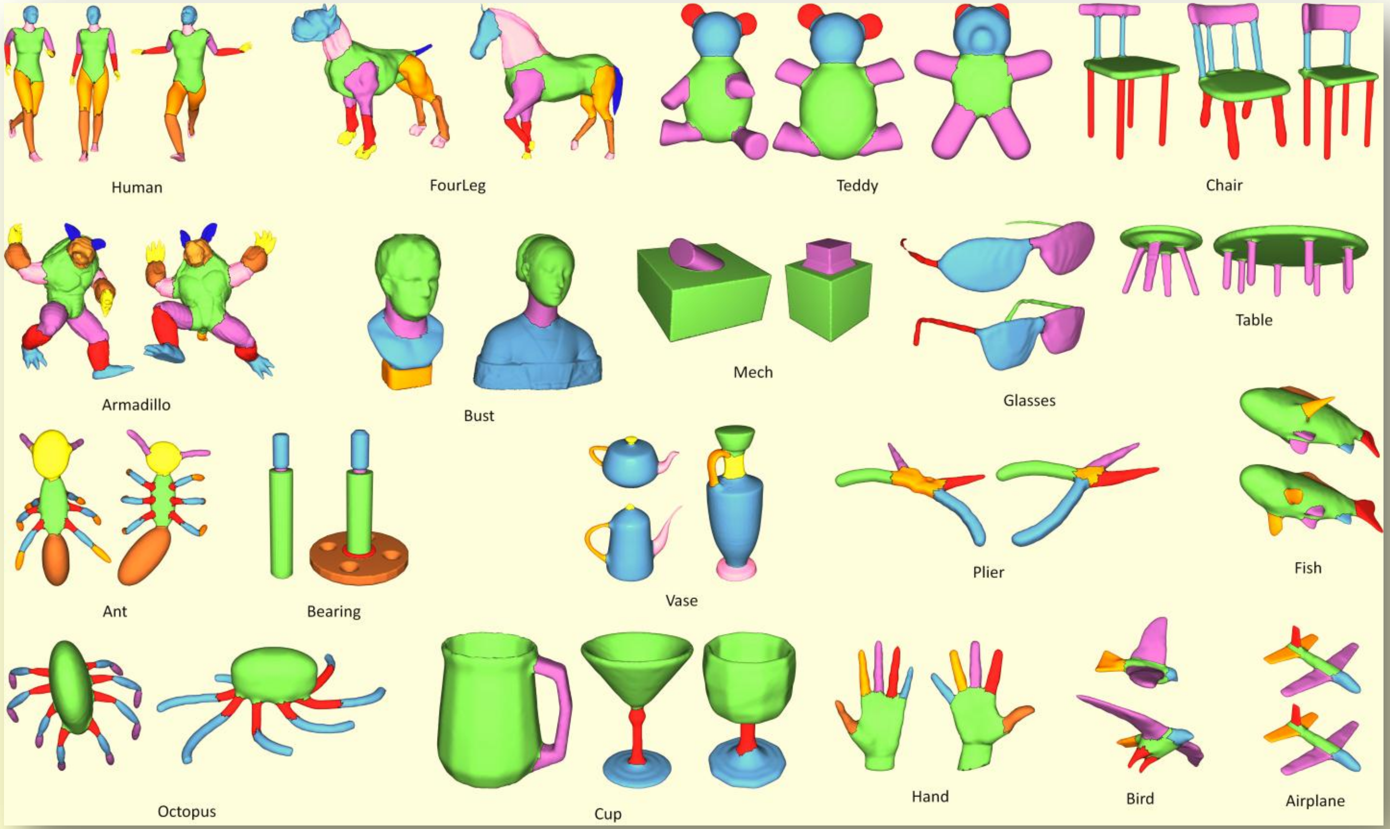

Shapes Have Semantics Beyond Surface Geometry

- Surface geometry alone may not capture all that is important about the shape
- Internal structure
- Function or use

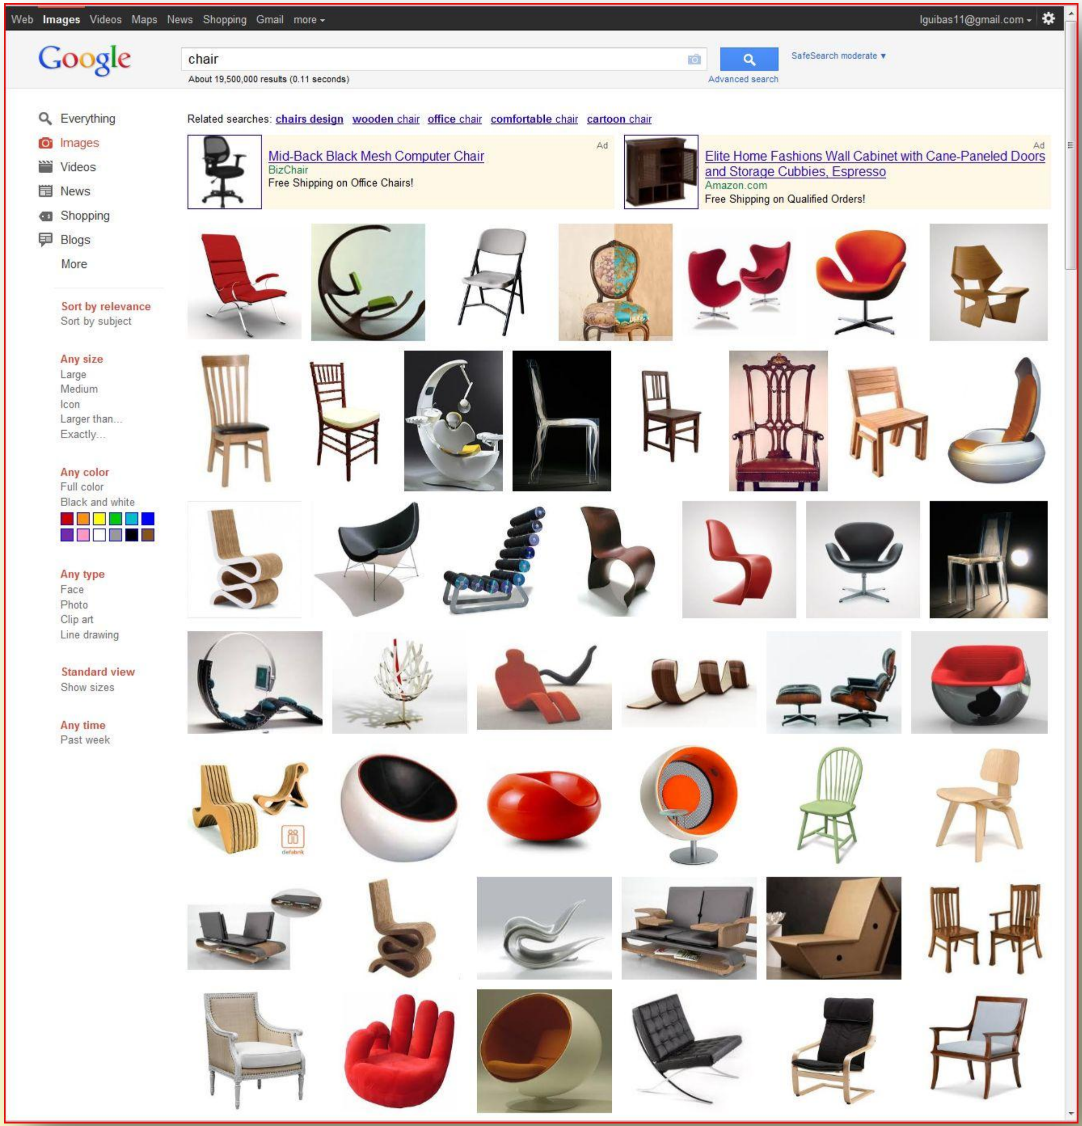

Why Joint Segmentation

Single shape segmentation

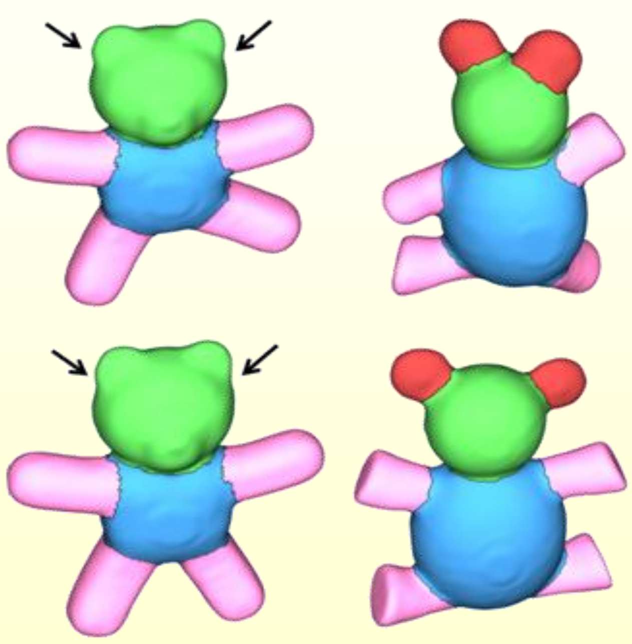

Joint shape segmentation

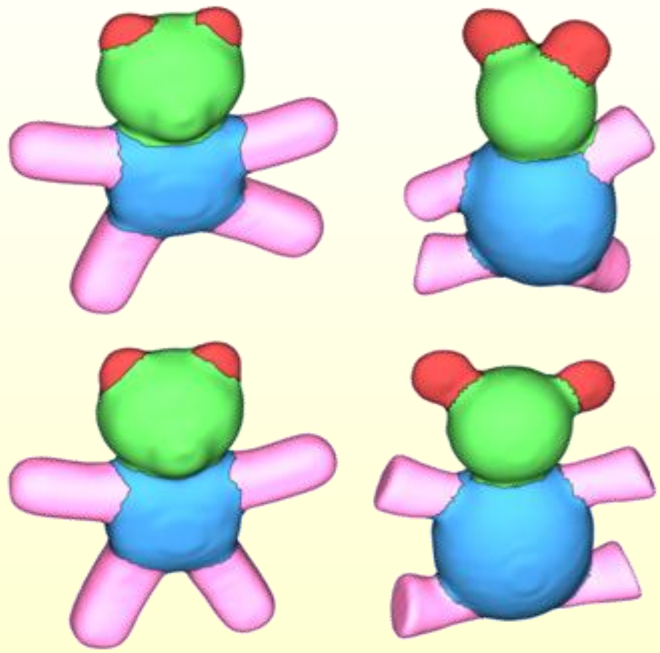

Low saliency

Why Joint Segmentation

Single shape segmentation

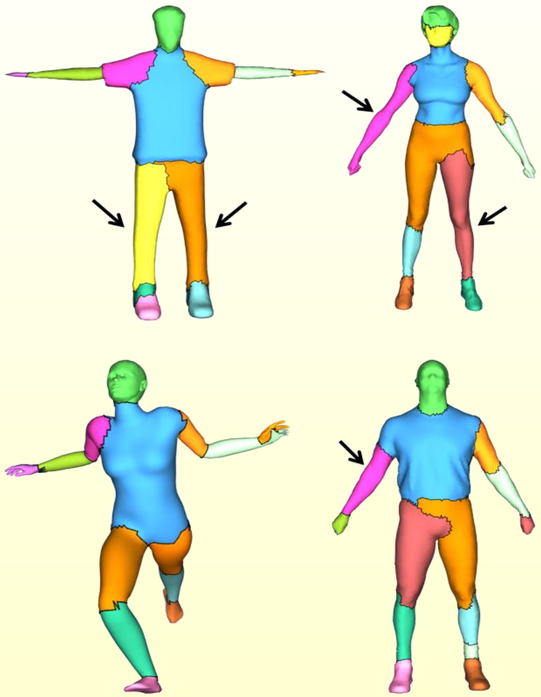

Joint shape segmentation

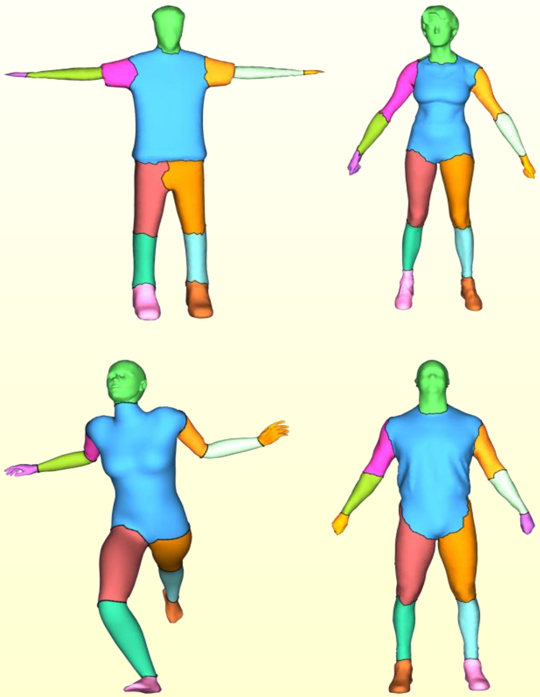

Consistency

Why Joint Segmentation

Single shape segmentation

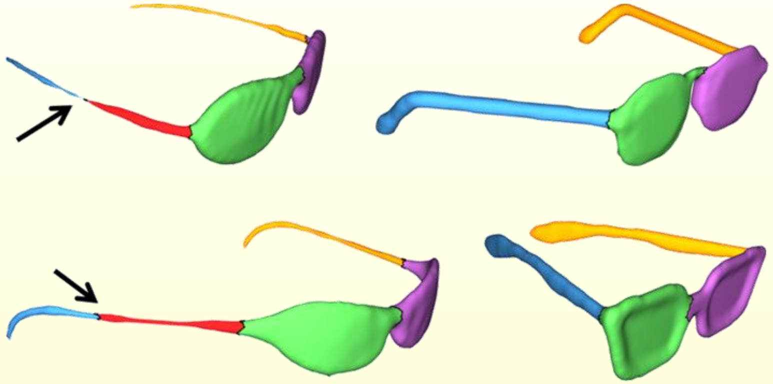

Joint shape segmentation

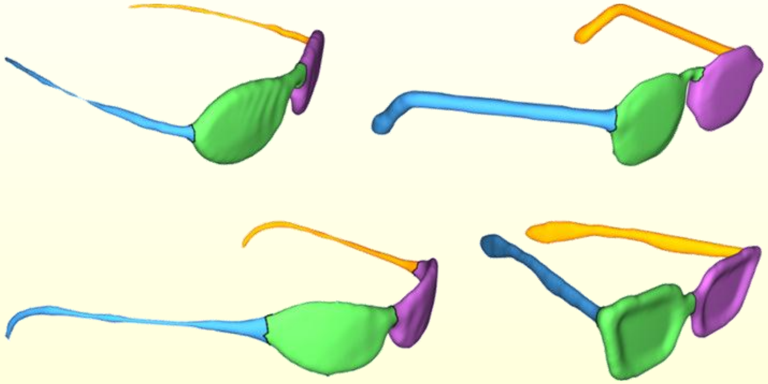

Extraneous geometric clues

Segmentation Evaluation: Princeton Segmentation Benchmark [Chen et al. 09]

- 380 shapes in 19 categories
- Manual segmentations for each shape (4300 in total)
- Evaluation metrics

Overview

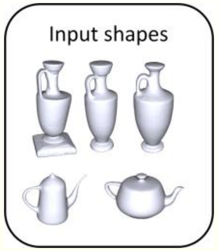

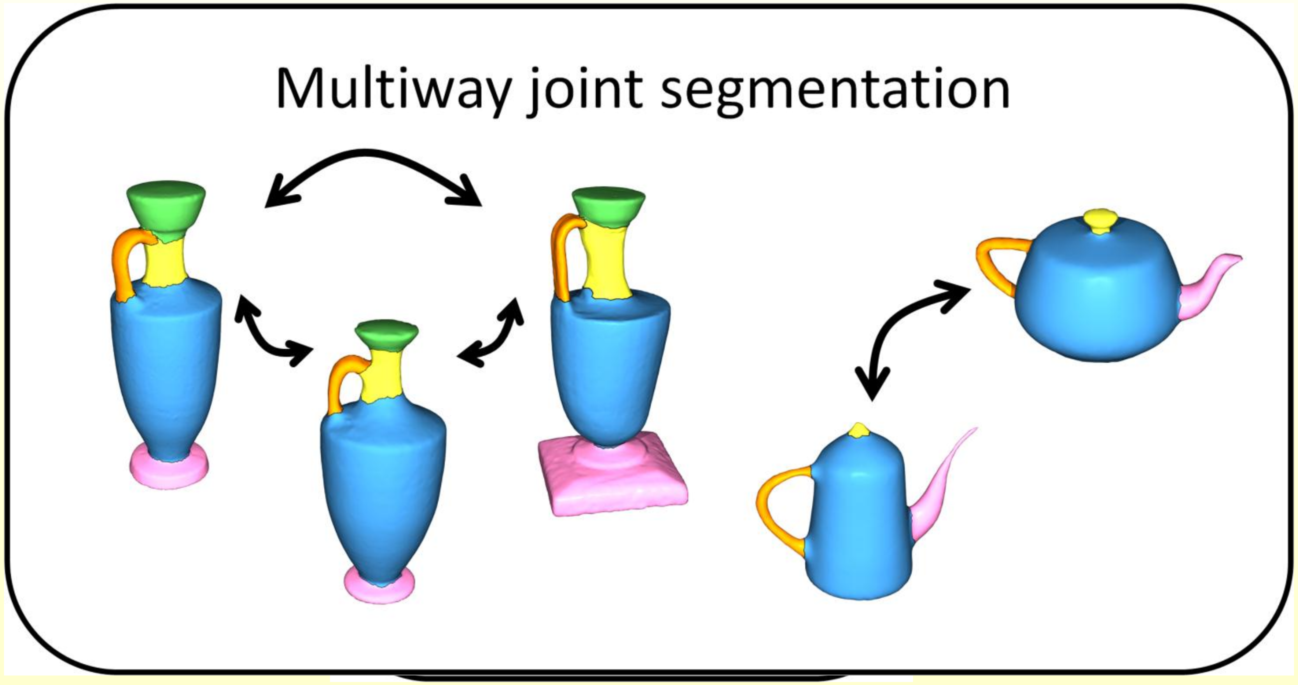

Initial Segments

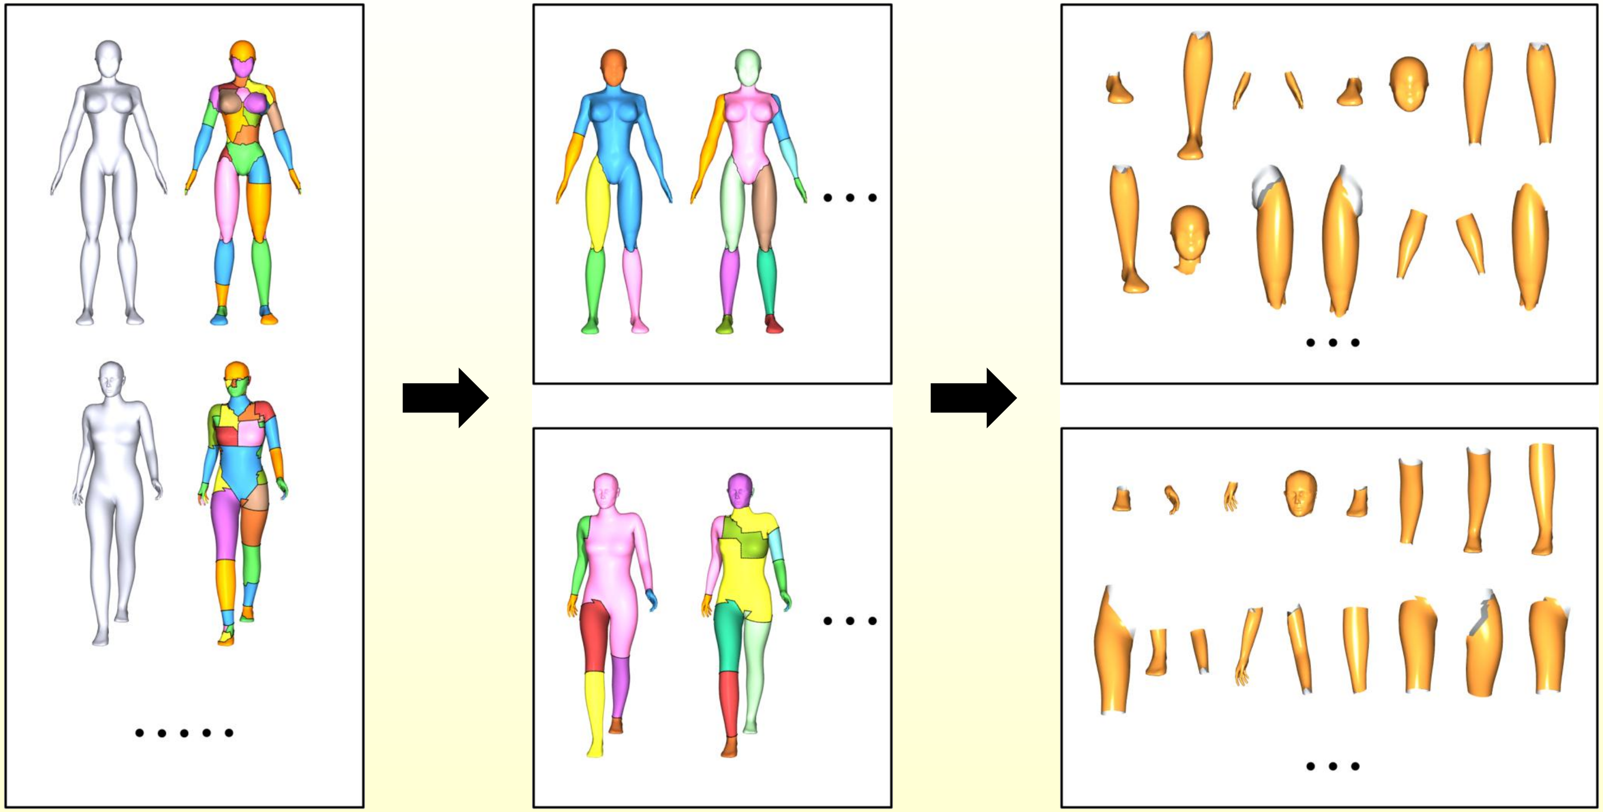

Patches, N-cuts [SM97]

Randomized Segmentations [GF08]

Initial Segments

Segment Weights

Frequency in randomized segmentations

Most similar segment on each other shape

$$w_s = \sum_{j=1}^{n} w_{(s,s_j^*)} r(s_j^*)$$

Geometry based similarity score

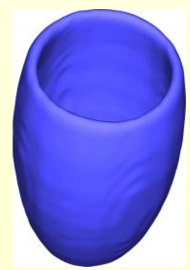

Importance diffusion

Pair-wise Co-Segmentation

- Optimize over segmentations and mappings between them
  - Each segmentation is given by a subset of initial segments
  - Directed maps

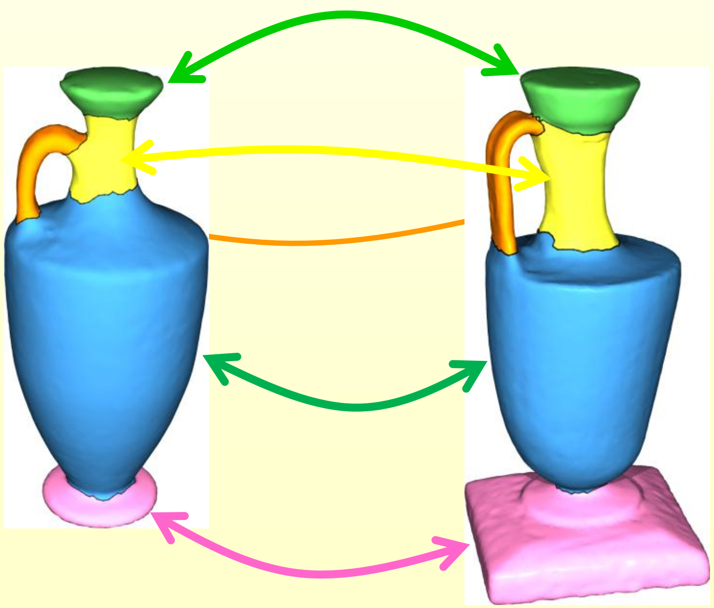

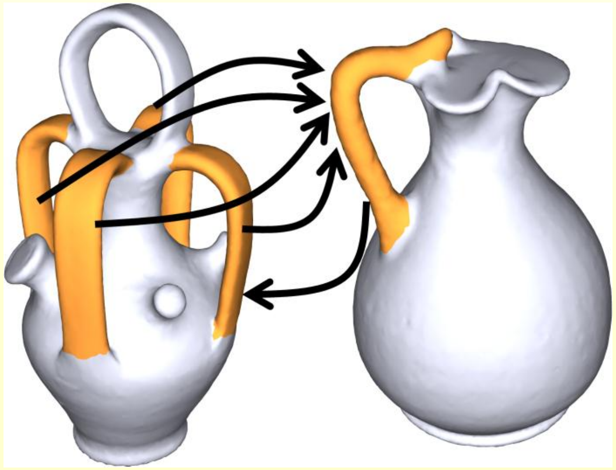

Objective Function

$$\max_{S_1 \subset \mathcal{I}_1, S_2 \subset \mathcal{I}_2} \quad \text{seg}(S_1) + \text{seg}(S_2)$$

$$+\left(\max_{\mathcal{M}_{12} \in S_1 \times S_2} \text{consistency}(\mathcal{M}_{12}) + \max_{\mathcal{M}_{21} \in S_2 \times S_1} \text{consistency}(\mathcal{M}_{21})\right)$$

$$\text{seg}(S_i) = \sum_{s \in S_i} \overline{w}_s = \sum_{s \in S_i} \overline{\text{area}}(s) w_s$$

$\overline{w}$ () + $\overline{w}$ () + $\overline{w}$ () + $\ldots$

$$\text{consistency}(\mathcal{M}_{12}) = \lambda \sum_{c \in \mathcal{M}_{12}} \overline{w}_c + \mu \sum_{(c,c') \in \mathcal{A}_{12}} \overline{w}_{(c,c')}$$

[Anguelov et al. ‘04]

$\overline{w}$ ($\bullet \longrightarrow \bullet$)

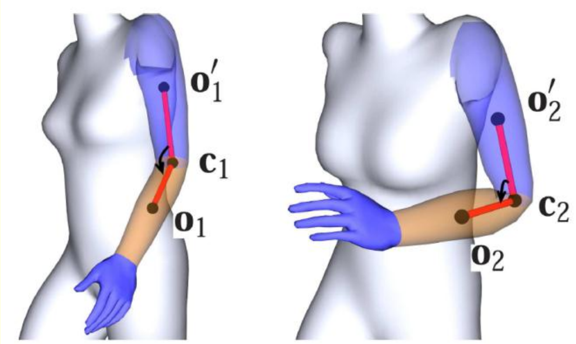

Constraints

- Segmentation constraints
  - Each patch is in exactly one selected segment

$$|\text{cover}(p)| = 1, \quad \forall p \in \mathcal{P}_i$$

the set of all segments that cover patch p

- Mapping constraints
  - An injective map from the segmentation of one shape to another

$$\mathcal{M}_{12} \subset \text{Injective}(S_1 \times S_2)$$
$$\mathcal{M}_{21} \subset \text{Injective}(S_2 \times S_1)$$

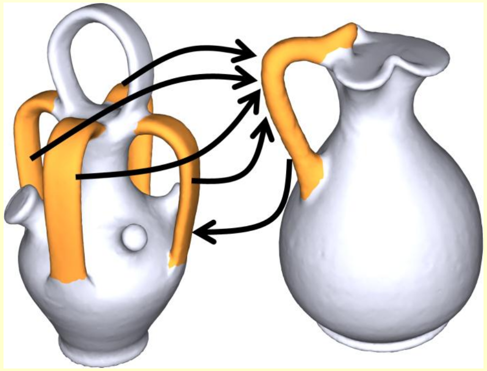

Integer Programming Formulation

- Assign segments and correspondences with binary indicator variables
  - Encode all possible configurations

$$x_s = \begin{cases} 1 & s \in S_1 \cup S_2 \\ 0 & \text{otherwise} \end{cases} \qquad y_c = \begin{cases} 1 & c \in \mathcal{M}_{12} \cup \mathcal{M}_{21} \\ 0 & \text{otherwise} \end{cases}$$

Map constraints:

$$\sum_{s' \in I_j} y_{(s,s')} \leq x_s, \quad \forall s \in I_i \qquad \qquad B_{ij}\mathbf{y}_{ij} \leq D_{ij}\mathbf{x}_i$$

$$\Longrightarrow$$

$$y_{(s,s')} \leq x_{s'}, \quad \forall (s,s') \in C_{ij} \qquad \qquad B'_{ij}\mathbf{y}_{ij} \leq D'_{ij}\mathbf{x}_j$$

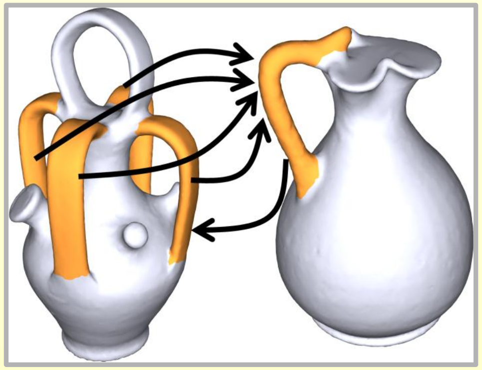

Binary Integer Programming Formulation

- Assign segments and correspondences with binary indicator variables

$$\max \sum_{i\in\{1,2\}} \mathbf{x}_i^{\mathsf{T}} \mathbf{w}_i^{\text{seg}} + \sum_{ij\in\{12,21\}} \left(\lambda \mathbf{y}_{ij}^{\mathsf{T}} \mathbf{w}_{ij}^{\text{corr}} + \mu \sum_{(c,c')\in\mathcal{A}_{ij}} y_c y_{c'} \overline{w}_{(c,c')} \right)$$

s.t. $\boxed{A_1\mathbf{x}_1 = 1 \quad \text{Seg. constraints} \quad A_2\mathbf{x}_2 = 1}$

$$\boxed{\begin{array}{lll} B_{12}\mathbf{y}_{12} \le D_{12}\mathbf{x}_1 & \text{Mapping} & B_{21}\mathbf{y}_{21} \le D_{21}\mathbf{x}_2 \\ B'_{12}\mathbf{y}_{12} \le D'_{12}\mathbf{x}_2 & \text{constraints} & B'_{21}\mathbf{y}_{21} \le D'_{21}\mathbf{x}_1 \end{array}}$$

and $x \in \{0, 1\}$ $\quad \forall x \in \mathbf{x}_1, \mathbf{x}_2, \mathbf{y}_{12}, \mathbf{y}_{21}$

Linear Programming Relaxation

- Linearize the objective function [Kumar et al. 09]

$$z_{(c,c')} = y_c y_{c'}$$

- Relax the variables

$$\max \sum_{i \in \{1,2\}} \mathbf{x}_i^{\mathsf{T}} \mathbf{w}_i^{\mathsf{seg}} + \sum_{ij \in \{12,21\}} (\lambda \mathbf{y}_{ij}^{\mathsf{T}} \mathbf{w}_{ij}^{\mathsf{corr}} + \mu \mathbf{z}_{ij}^{\mathsf{T}} \mathbf{w}_{ij}^{\mathsf{adj}})$$

$$\begin{array}{llll}
\text{s.t.} & A_1 \mathbf{x}_1 = 1 & & A_2 \mathbf{x}_2 = 1 \\
& B_{12} \mathbf{y}_{12} \le D_{12} \mathbf{x}_1 & & B_{21} \mathbf{y}_{21} \le D_{21} \mathbf{x}_2 \\
& B'_{12} \mathbf{y}_{12} \le D'_{12} \mathbf{x}_2 & & B'_{21} \mathbf{y}_{21} \le D'_{21} \mathbf{x}_1 \\
& E_{12} \mathbf{z}_{12} \le F_{12} \mathbf{y}_{12} & & E_{21} \mathbf{z}_{21} \le F_{21} \mathbf{y}_{21} \\
\text{and} & 0 \le x \le 1 & & \forall x \in \mathbf{x}_1, \mathbf{x}_2, \mathbf{y}_{12}, \mathbf{y}_{21}, \mathbf{z}_{12}, \mathbf{z}_{21}
\end{array}$$

Multi-way Joint Segmentation

- Combines objective functions of pairs of similar shapes
  - Threshold on values of objective functions

$$\max \sum_{i=1}^{n} \mathbf{x}_i^{\mathsf{T}} \mathbf{w}_i^{\mathsf{seg}} + \frac{n}{|\mathcal{E}|} \sum_{(i,j)\in\mathcal{E}} (\lambda \mathbf{y}_{ij}^{\mathsf{T}} \mathbf{w}_{ij}^{\mathsf{corr}} + \mu \mathbf{z}_{ij}^{\mathsf{T}} \mathbf{w}_{ij}^{\mathsf{adj}})$$

$$\text{s.t.} \quad A_i \mathbf{x}_i = \mathbf{1},\ \mathbf{0} \leq \mathbf{x}_i \leq \mathbf{1} \quad \text{for all } 1 \leq i \leq n$$

$$\text{and} \quad B_{ij}\mathbf{y}_{ij} \leq D_{ij}\mathbf{x}_i,\ B'_{ij}\mathbf{y}_{ij} \leq D'_{ij}\mathbf{x}_j,\ E_{ij}\mathbf{z}_{ij} \leq F_{ij}\mathbf{y}_{ij},$$

$$\mathbf{0} \leq \mathbf{y}_{ij} \leq \mathbf{1},\ \mathbf{0} \leq \mathbf{z}_{ij} \leq \mathbf{1} \quad \text{for all } (i,j) \in \mathcal{E}$$

Block coordinate descent for efficiency

Rand Index Scores on PSB

	SD	RC	Supervised	**Joint**	**JointAll**	Human
Average	17.2	15.3	10.7	**10.5**	**10.1**	10.3

- Significantly better than single shape segmentations
- Comparable or slightly better than supervised segmentation
- JointAll is slightly better than Joint

SD: shape diameter
RC: randomized cuts
Supervised

Joint: JS within a class
JointAll: JS over full DB

Rand Index Scores on PSB

Joint wins when per category shape variation is big

Armadillo

	SD	RC	Supervised	**Joint**	**JointAll**	Human
Armadillo	8.9	9.2	8.4	**7.4**	**7.4**	8.3

Rand Index Scores on PSB

Joint wins when per category shape variation is big

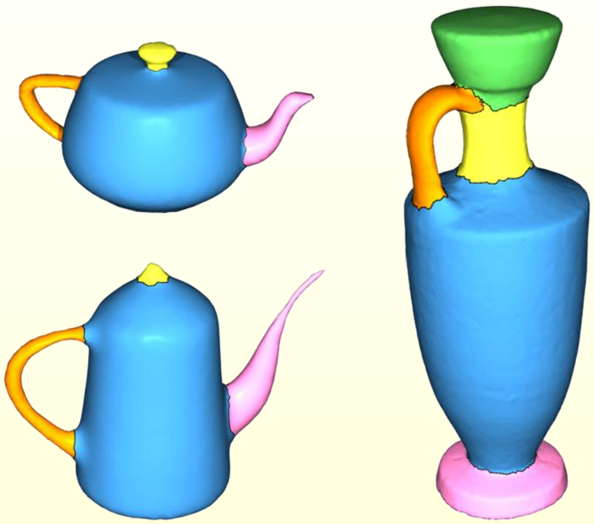

Vase

	SD	RC	Supervised	**Joint**	**JointAll**	Human
Vase	23.6	12.7	17.1	**13.5**	**13.2**	14.4

Rand Index Scores on PSB

Less benefit when per category shape variation is small

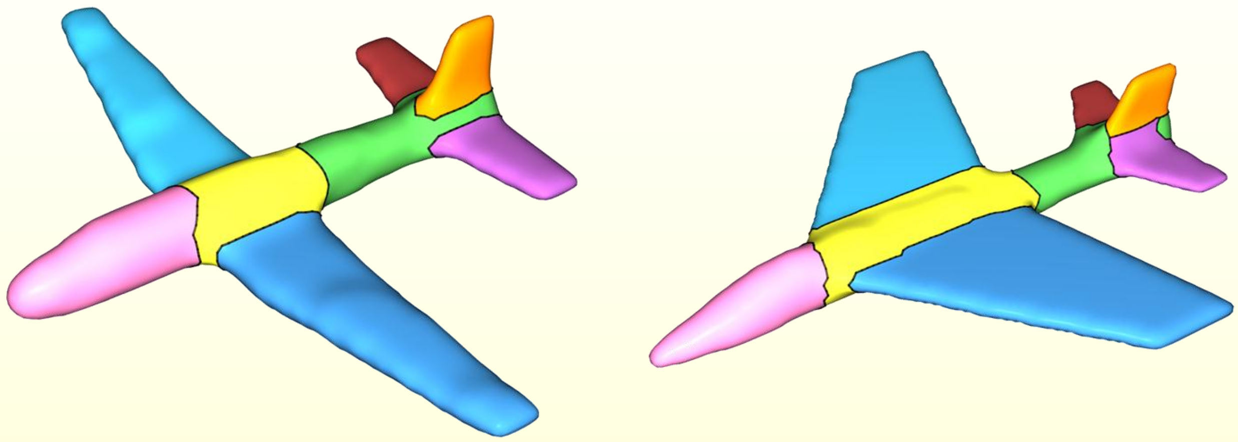

Airplane

	SD	RC	Supervised	**Joint**	**JointAll**	Human
Airplane	9.3	13.4	8.2	**12.9**	**10.2**	9.2

Joint versus JointAll

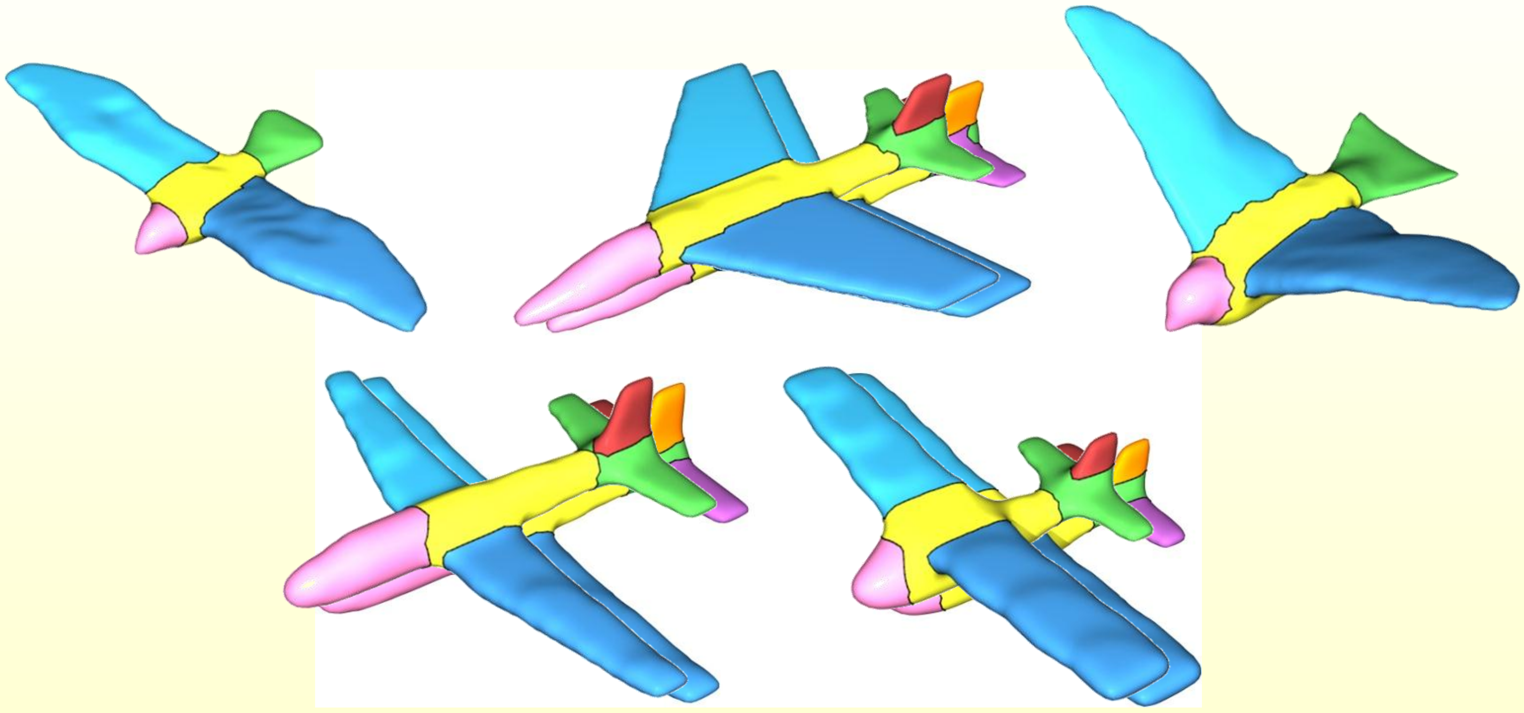

Joint versus JointAll

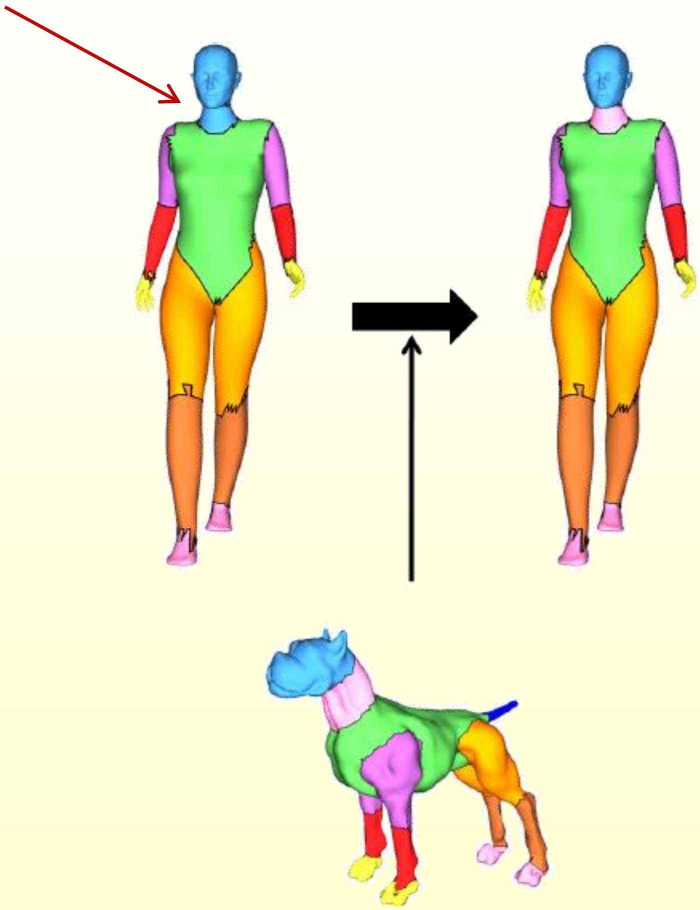

Versus Supervised Segmentation [KBS10]

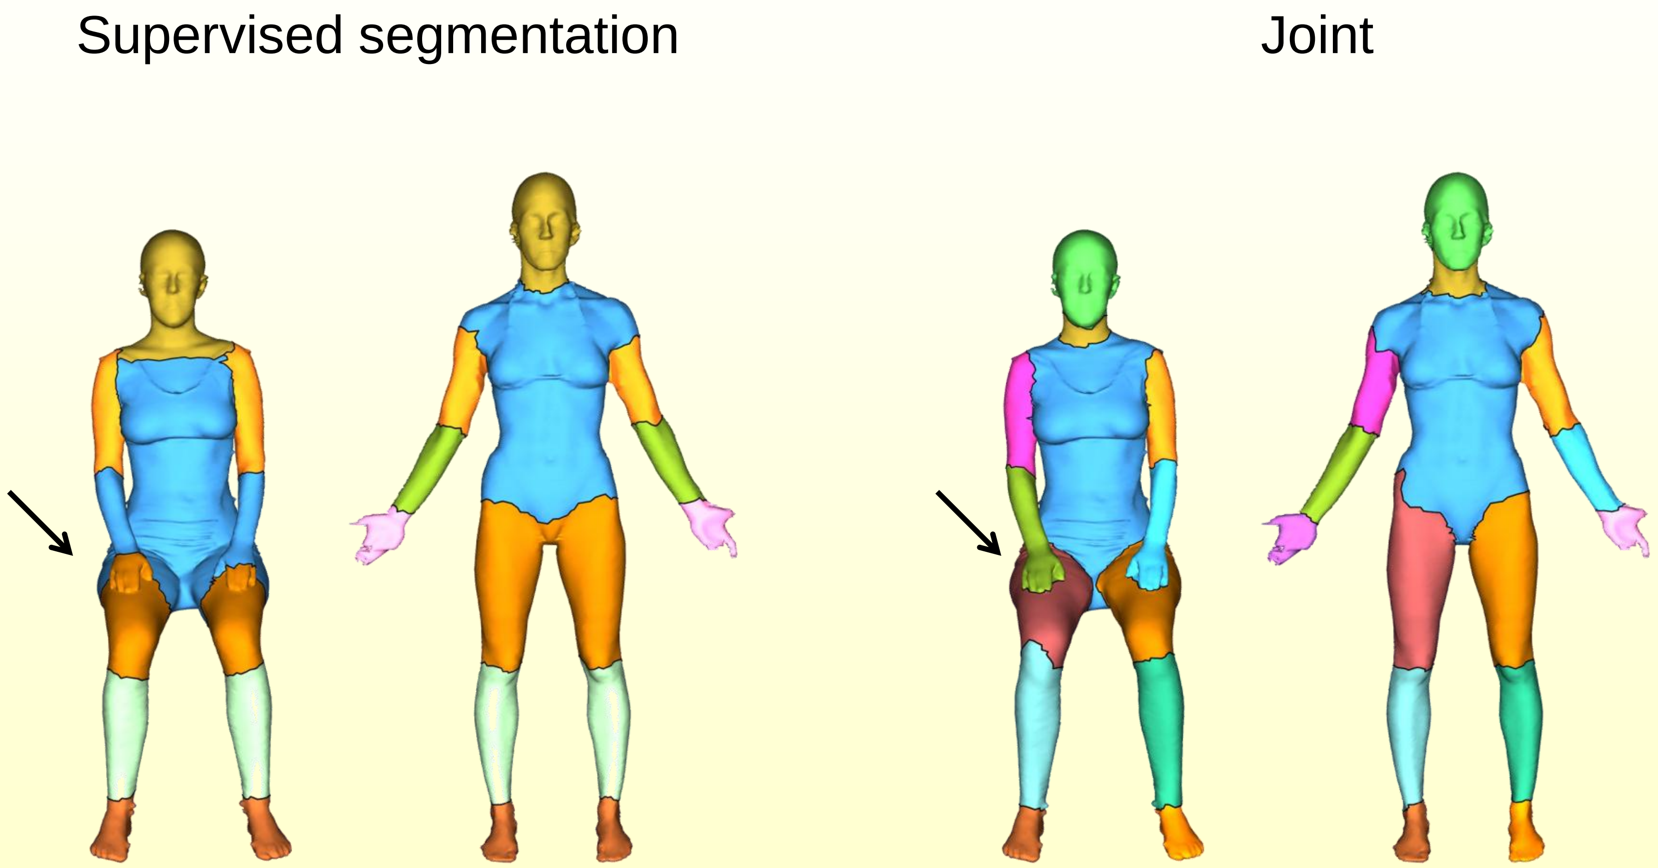

Rand Index Scores on PSB

Failure case

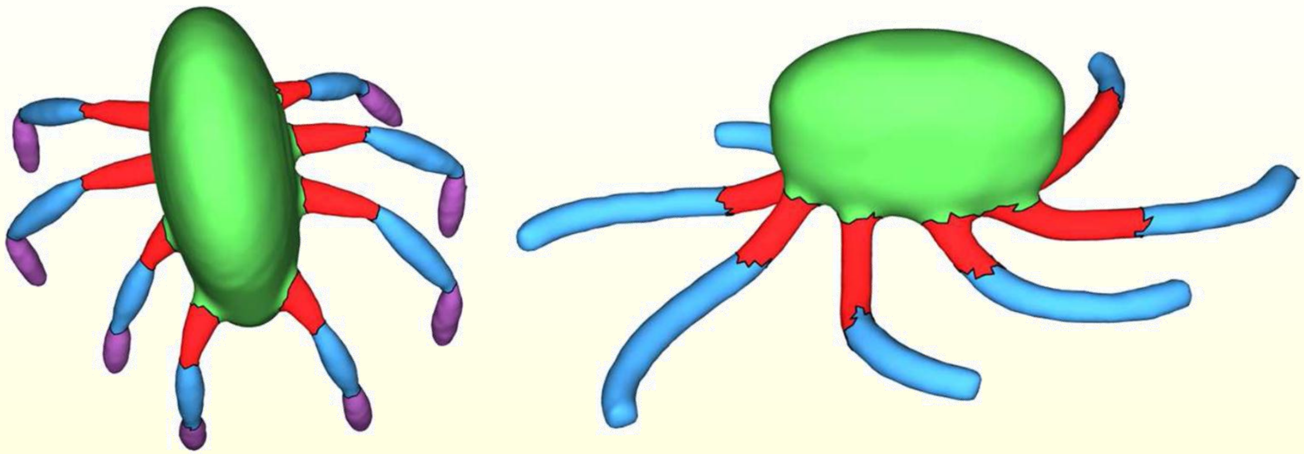

Octopus

	SD	RC	Supervised	**Joint**	**JointAll**	Human
Octopus	4.8	6.4	1.8	**6.7**	**7.2**	2.4

Different Levels of Similar Shapes

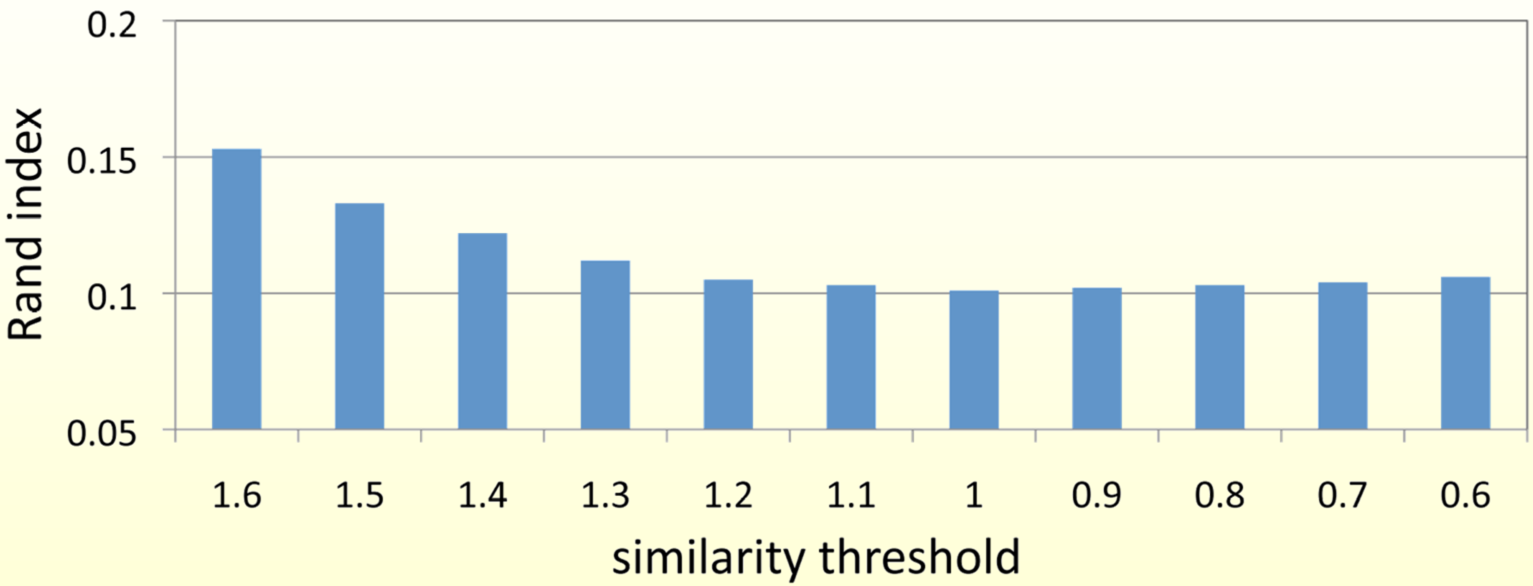

very similar shapes
less similar shapes

The Lessons

- By segmenting shapes jointly, we capture better semantic notions of shape parts
  - Less influenced by local geometry artifacts
  - The truth has less places to hide

Poster

Graphs of Map Systems

[with Nguyen, Ben-Chen, Welnicka, Ye, SGP ’11]

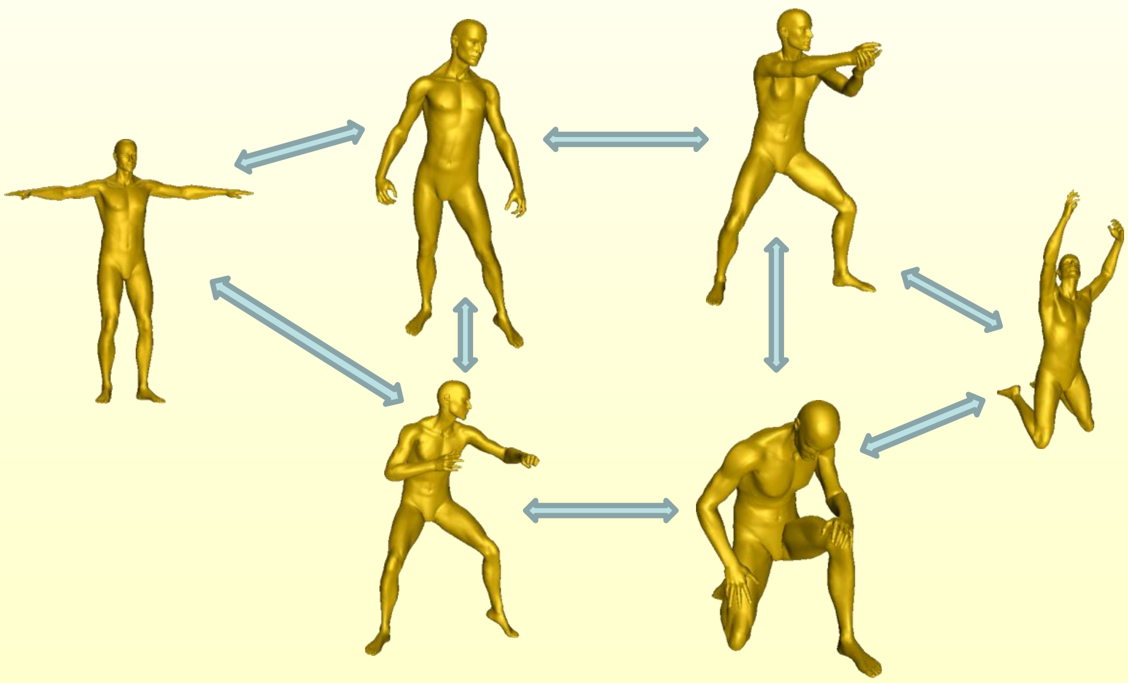

Optimal Maps Can Be Ambiguous or Unstable

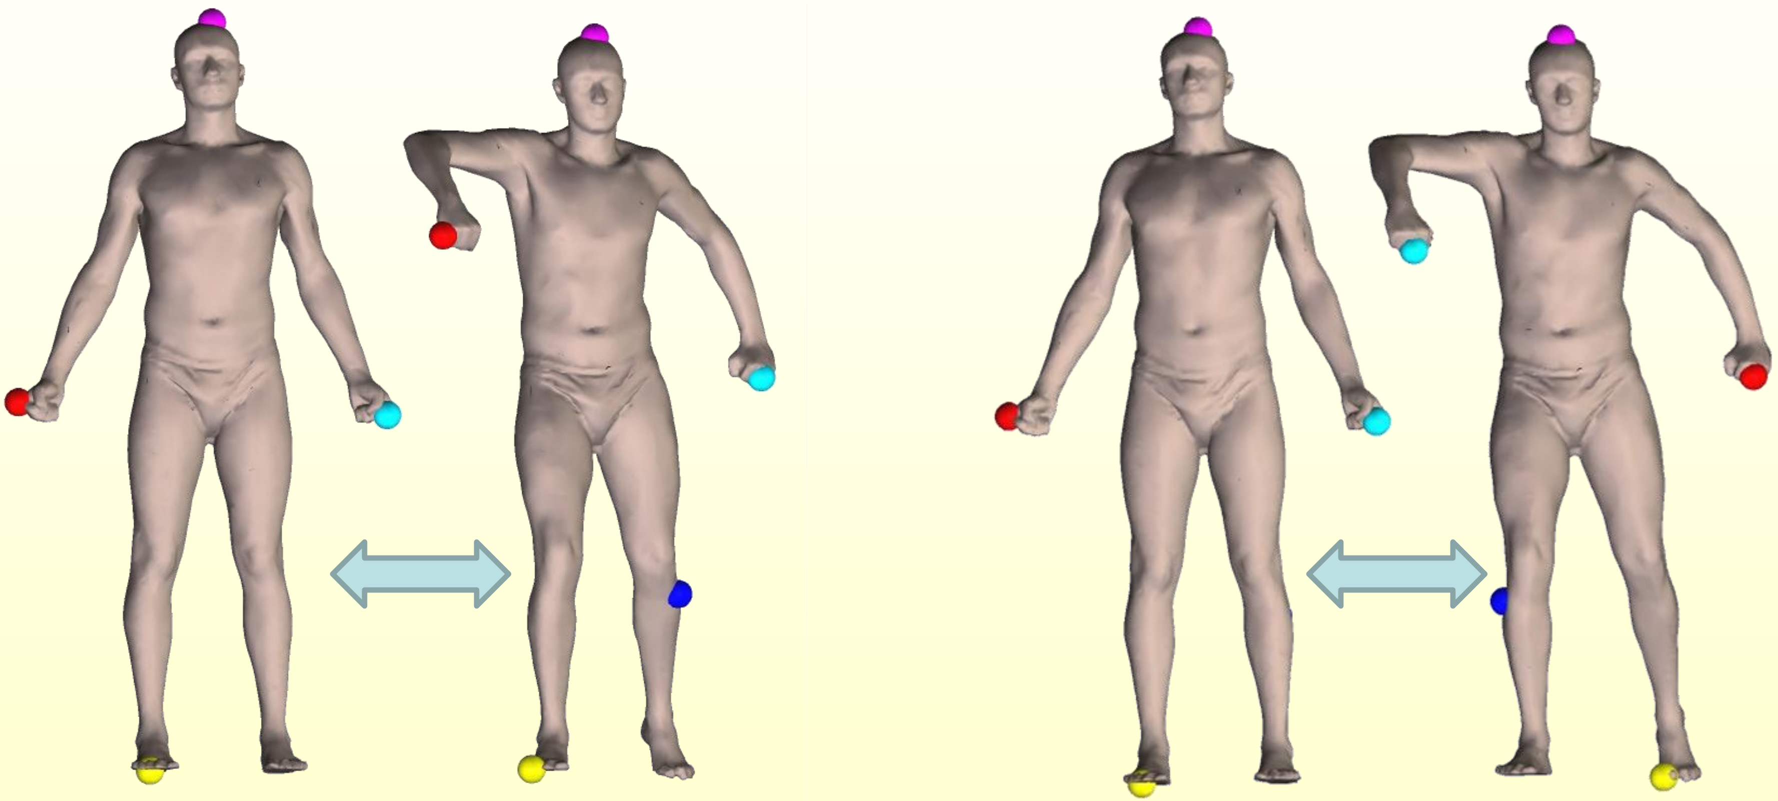

Equally good isometric maps

Problem Statement

- Input
  - A collection of related shapes
  - A collection of maps between all pairs of shapes
  - A distance measure on each shape
- Output
  - A collection of improved maps between all pairs of shapes.
  - Improved in the sense of being more
    - Accurate (close to ground truth)
    - Consistent (with each other)

Network Representation

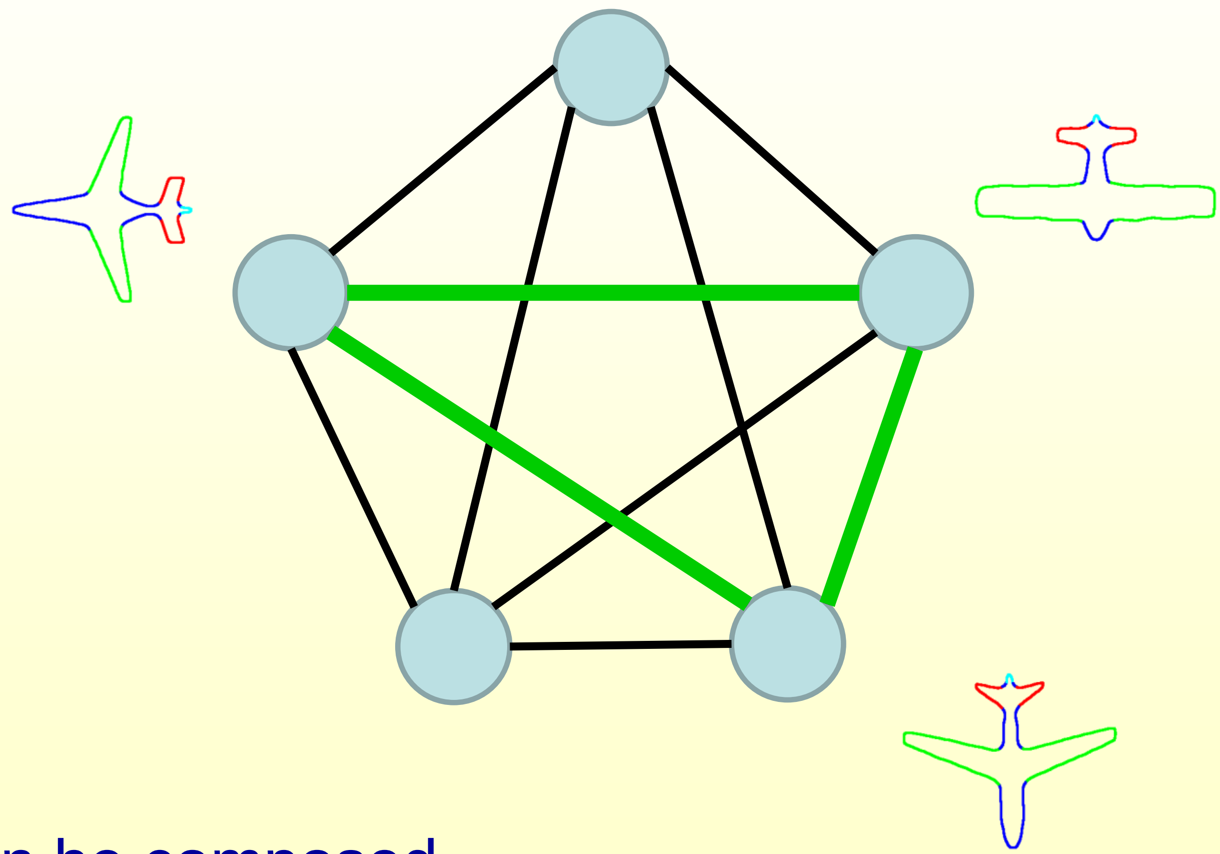

Maps can be composed.
For self-maps, easy to tell if they are good.

Approach: From Consistency to Accuracy

- Cycle consistency tells us something about accuracy
- Remove the inconsistencies we find
- Repeat using the improved collection

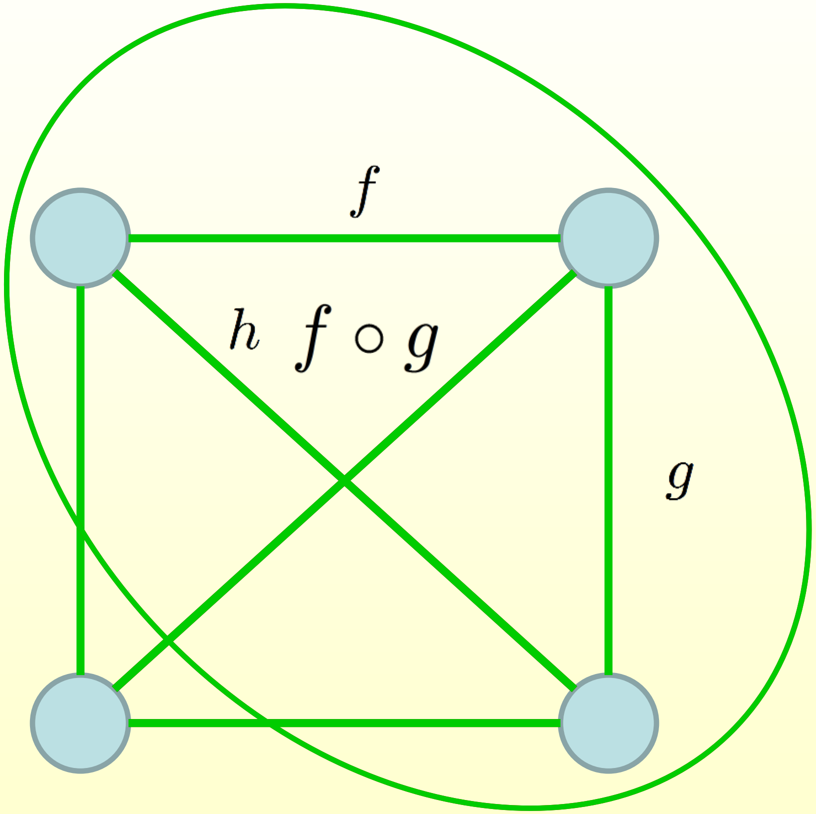

3 cycles

Proposal – Linear Program

- For each 3-cycle γ in the graph, compute the distortion C_γ
- Solve the following linear program to assign errors c_e to the edges:
  - Minimize $\displaystyle\sum_{e\in E} W_e c_e$
  - Subject to $\displaystyle\sum_{e\in\gamma} c_e \geq C_\gamma \ \forall\gamma$

    $c_e \geq 0 \ \forall e \in E$
  - Where $W_e = 1/(\displaystyle\sum_{\gamma:e\in\gamma} C_\gamma)$

L^1 concentrates the error on few edges

Proposal – Map Replacement

- LP gives us a weighted graph
- Remove bad maps: replace with shortest paths
- New collection of maps
- Run the LP again?

Convergence - Experimental

Map Type | LP Weights | Final accuracy

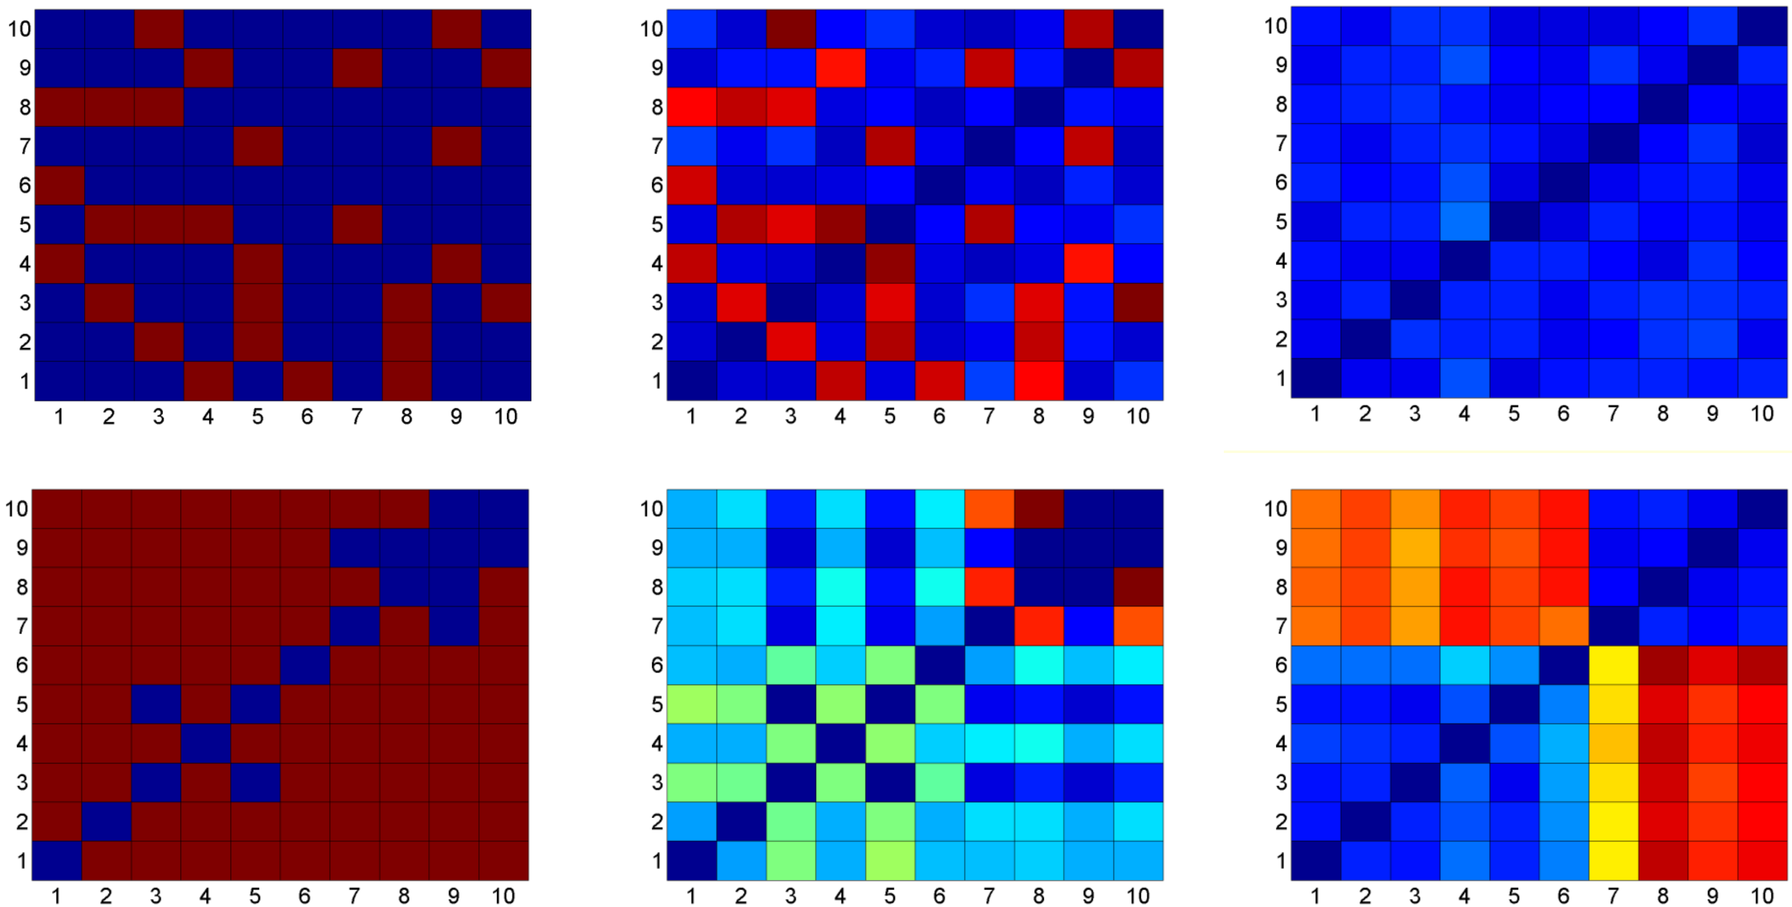

Results – 2D (DTW)

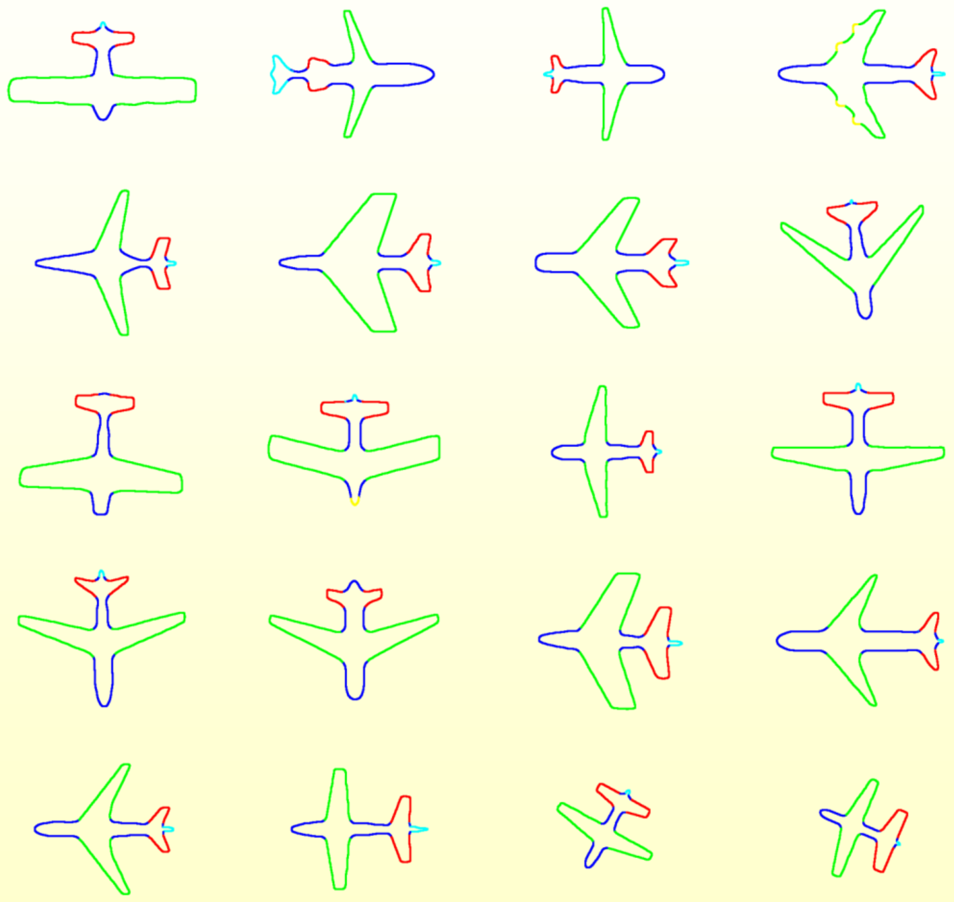

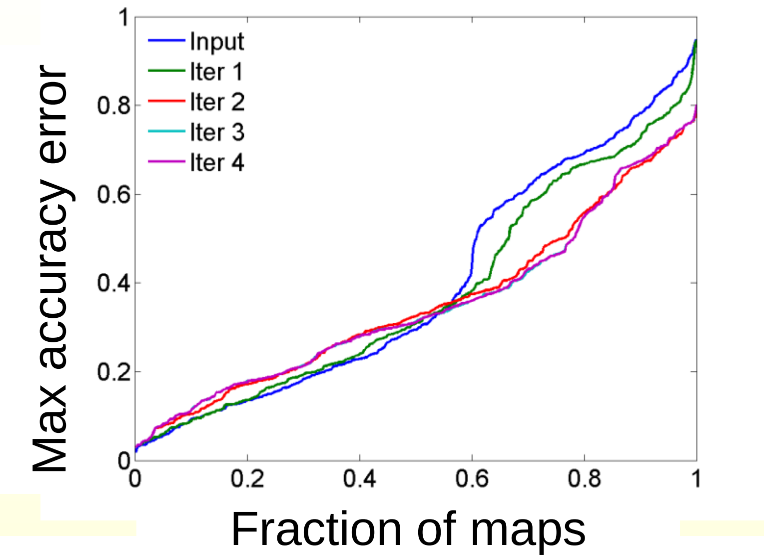

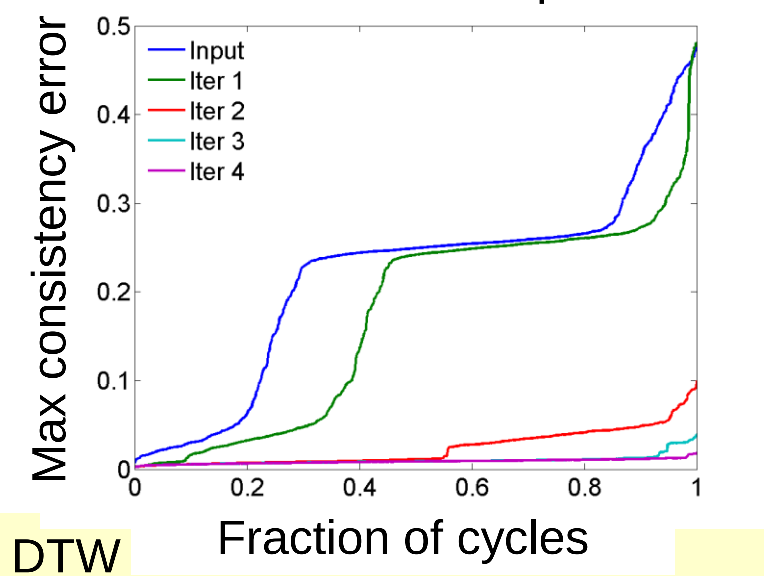

Princeton data, correspondences computed with DTW

Results – 2D (DTW)

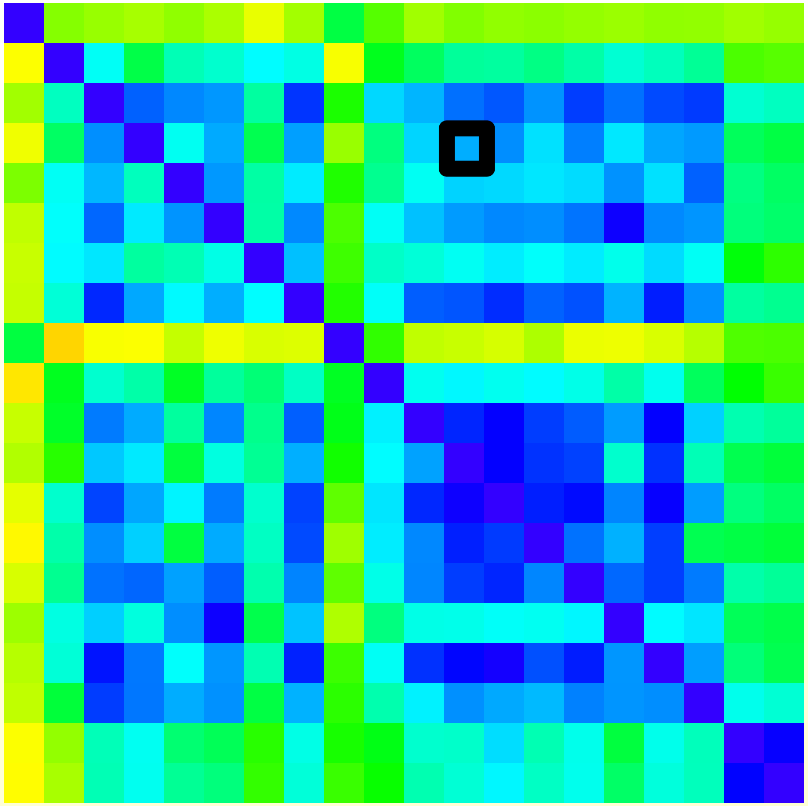

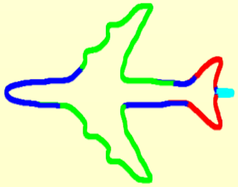

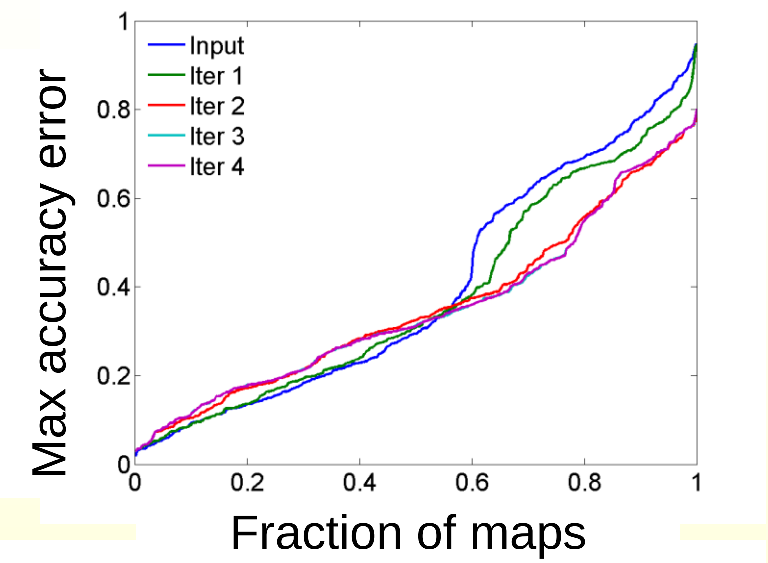

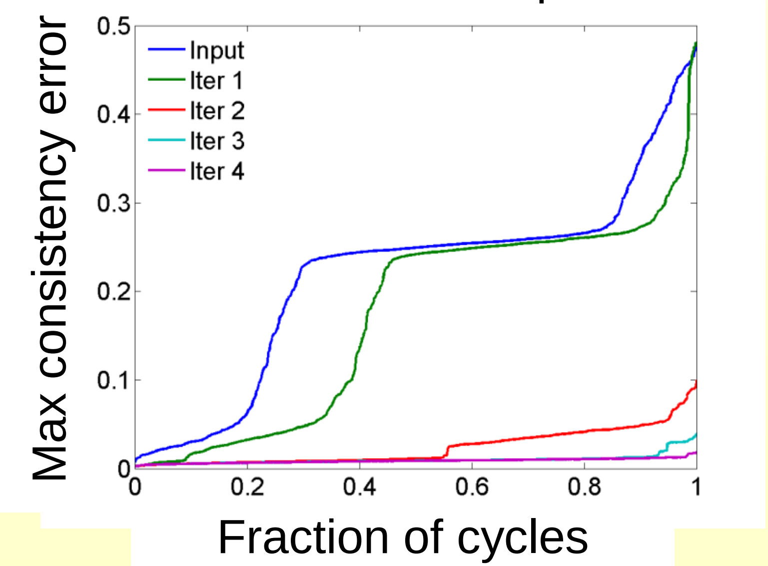

Results – 3D (Heat Kernel)

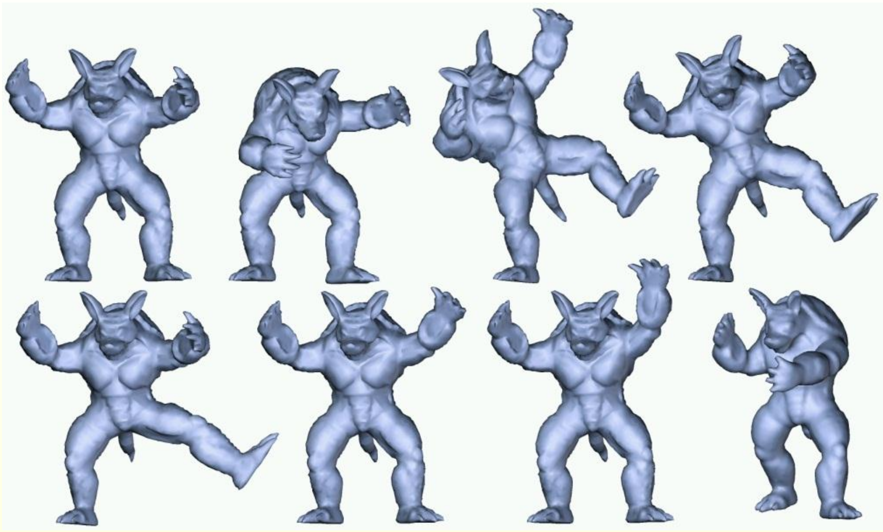

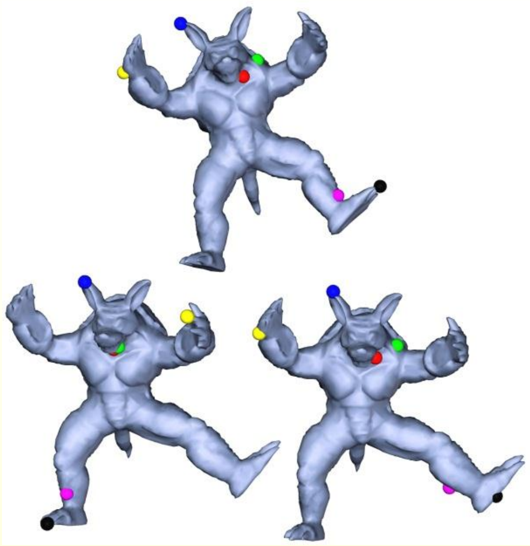

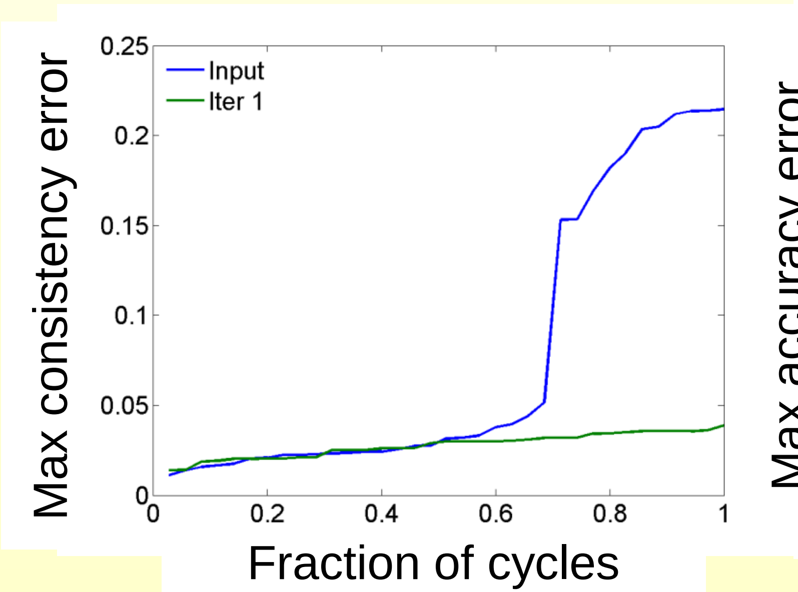

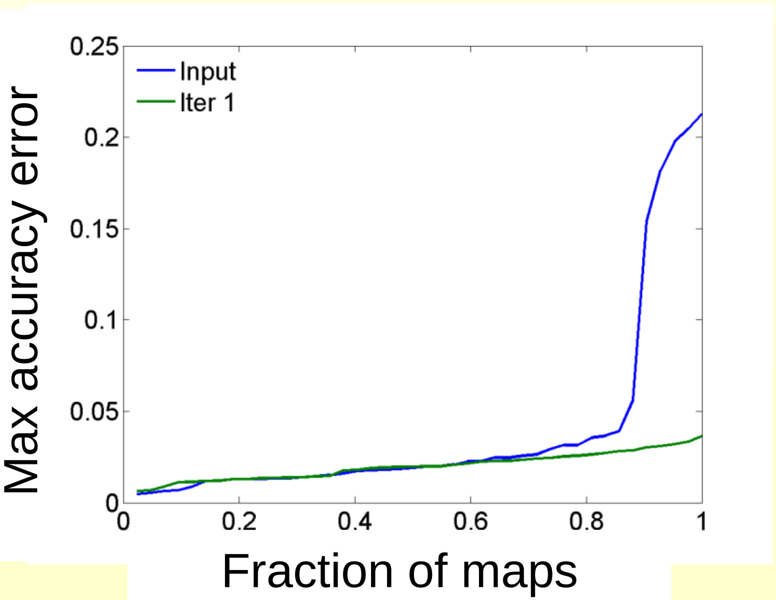

Results – 3D (Blended Maps)

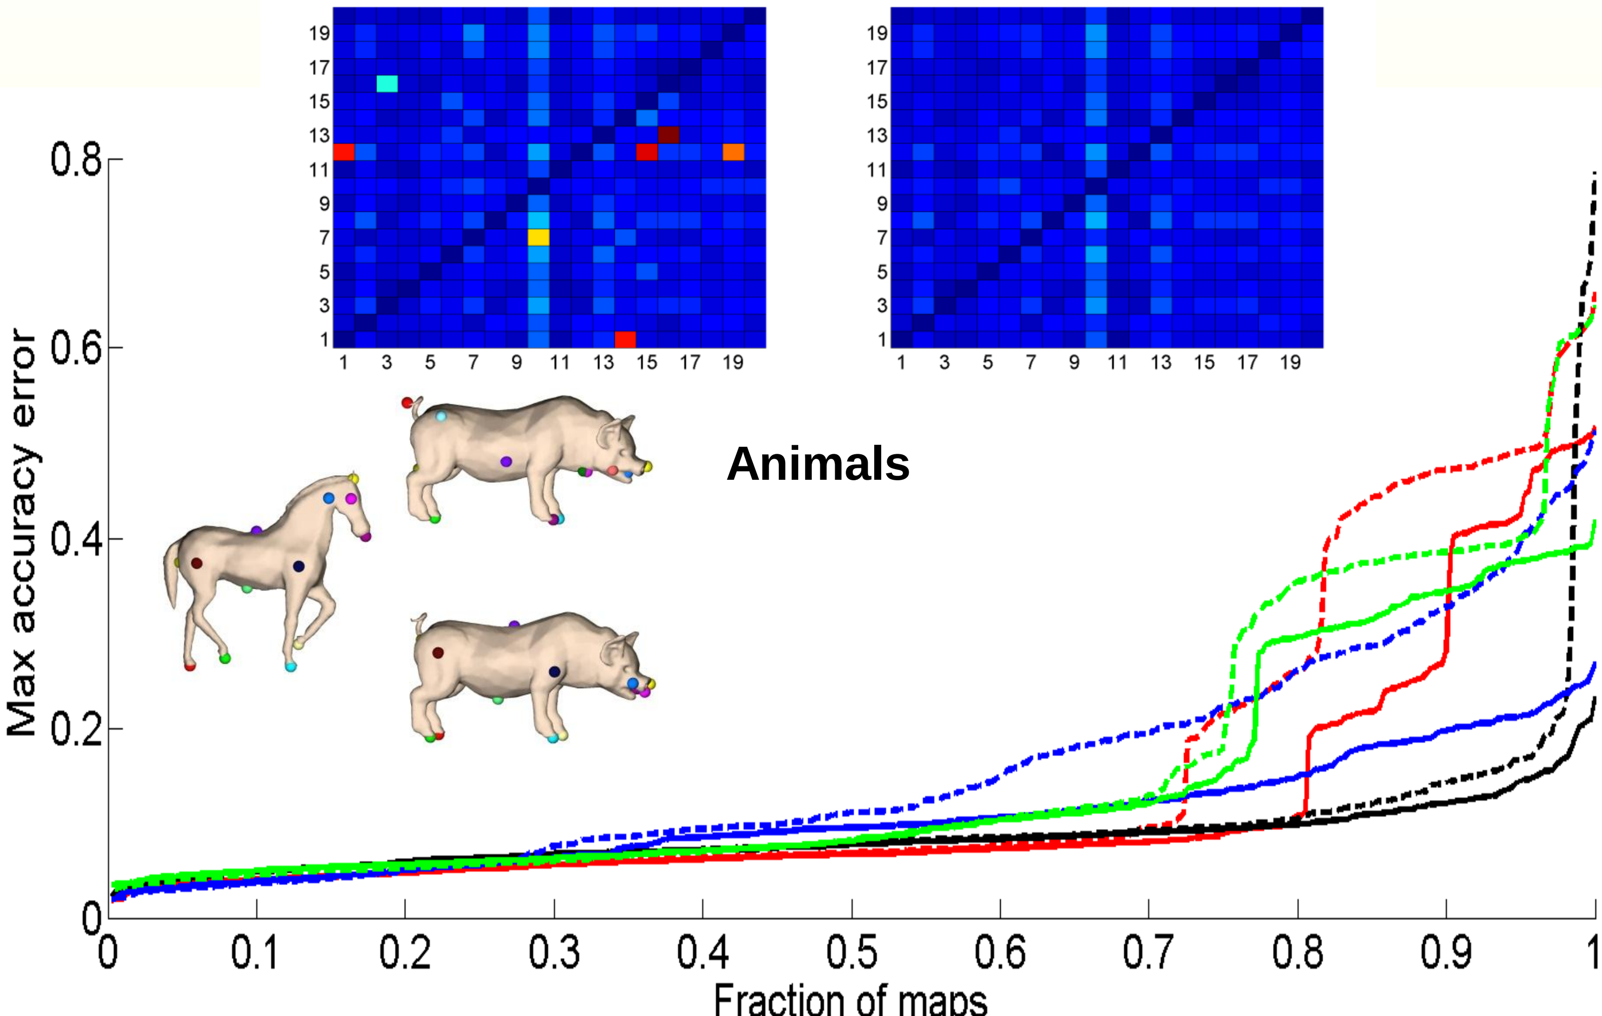

A Shape Morphing Result – 3D

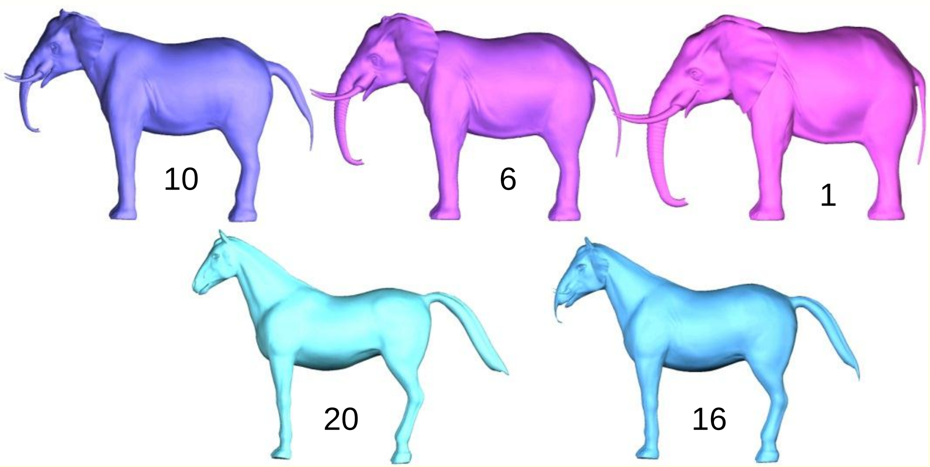

Correspondences computed with Mobius voting + GMDS

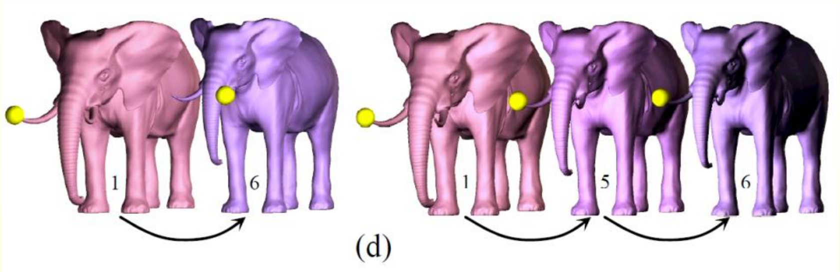

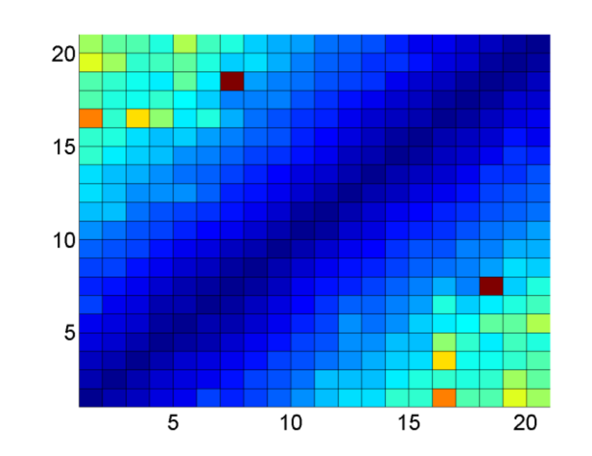

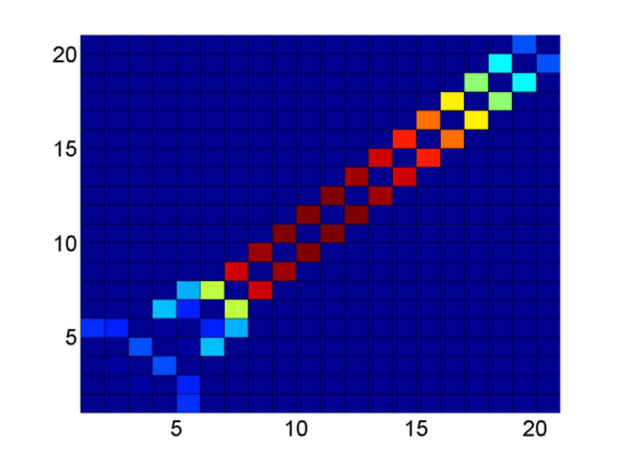

use frequency

The morph sequence is recovered ...

Phenotype genealogies

The Lessons

- Map networks are more powerful than graphs because maps can be composed

- They assist in the estimation of the consistency of shapes in a collection and thus can be used to used to understand the overall structure of the collection

Exploration of Continuous Variability in Shape Collections

[with Ovsjanikov, Li, Mitra, Siggraph '11]

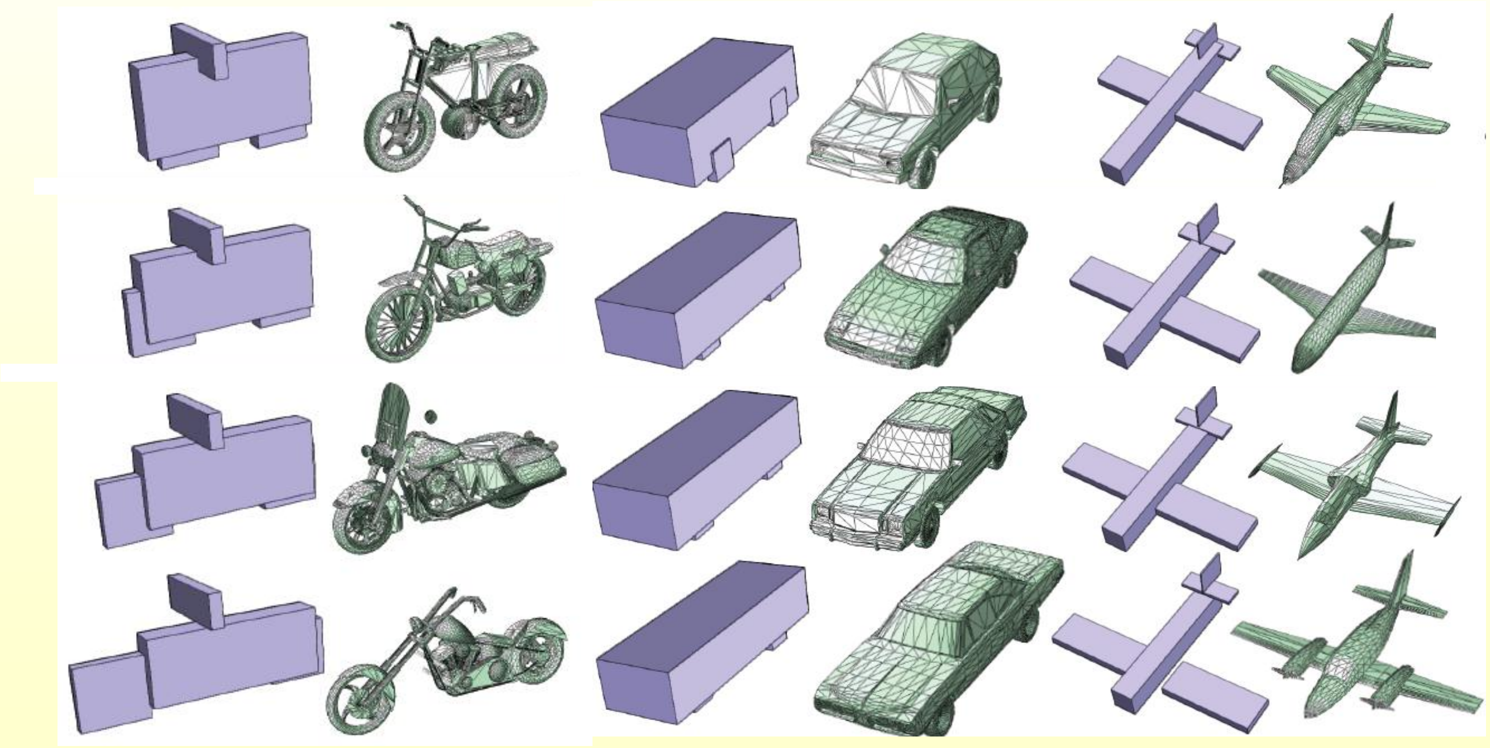

No correspondences or maps

Large Shape Repositories

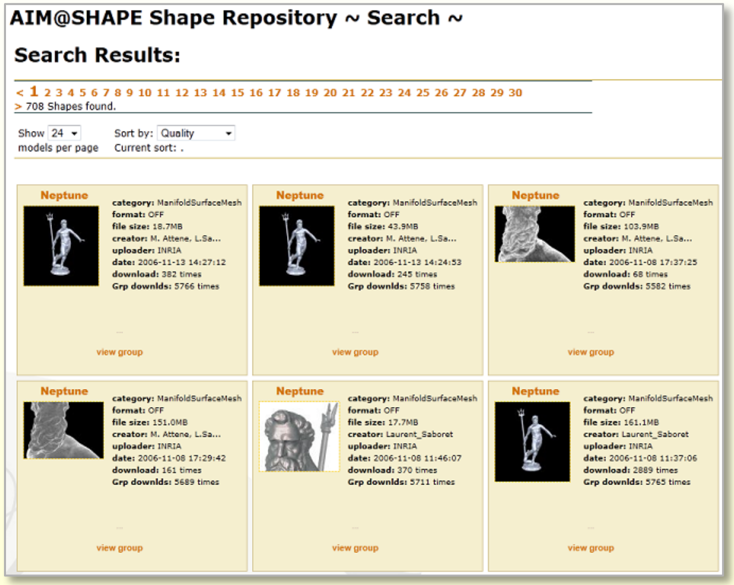

TURBO SQUID

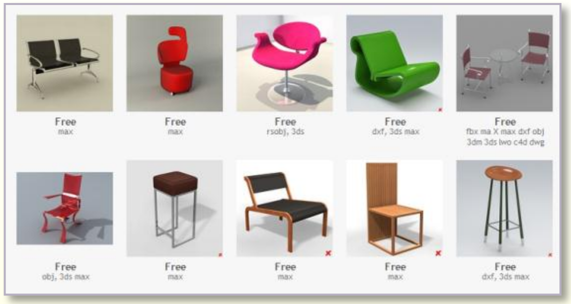

Google
3D warehouse

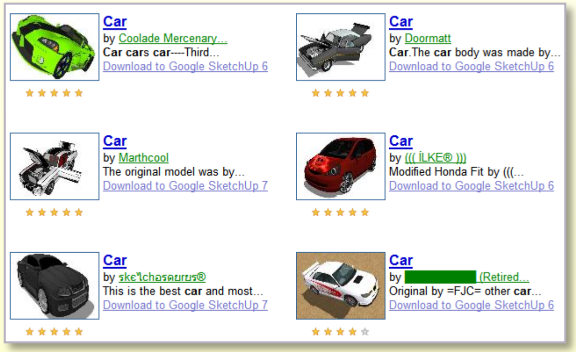

- Millions of models available
- Incorporating 3D models into workflows is challenging
  - difficult to know what is there

Text-Based Exploration

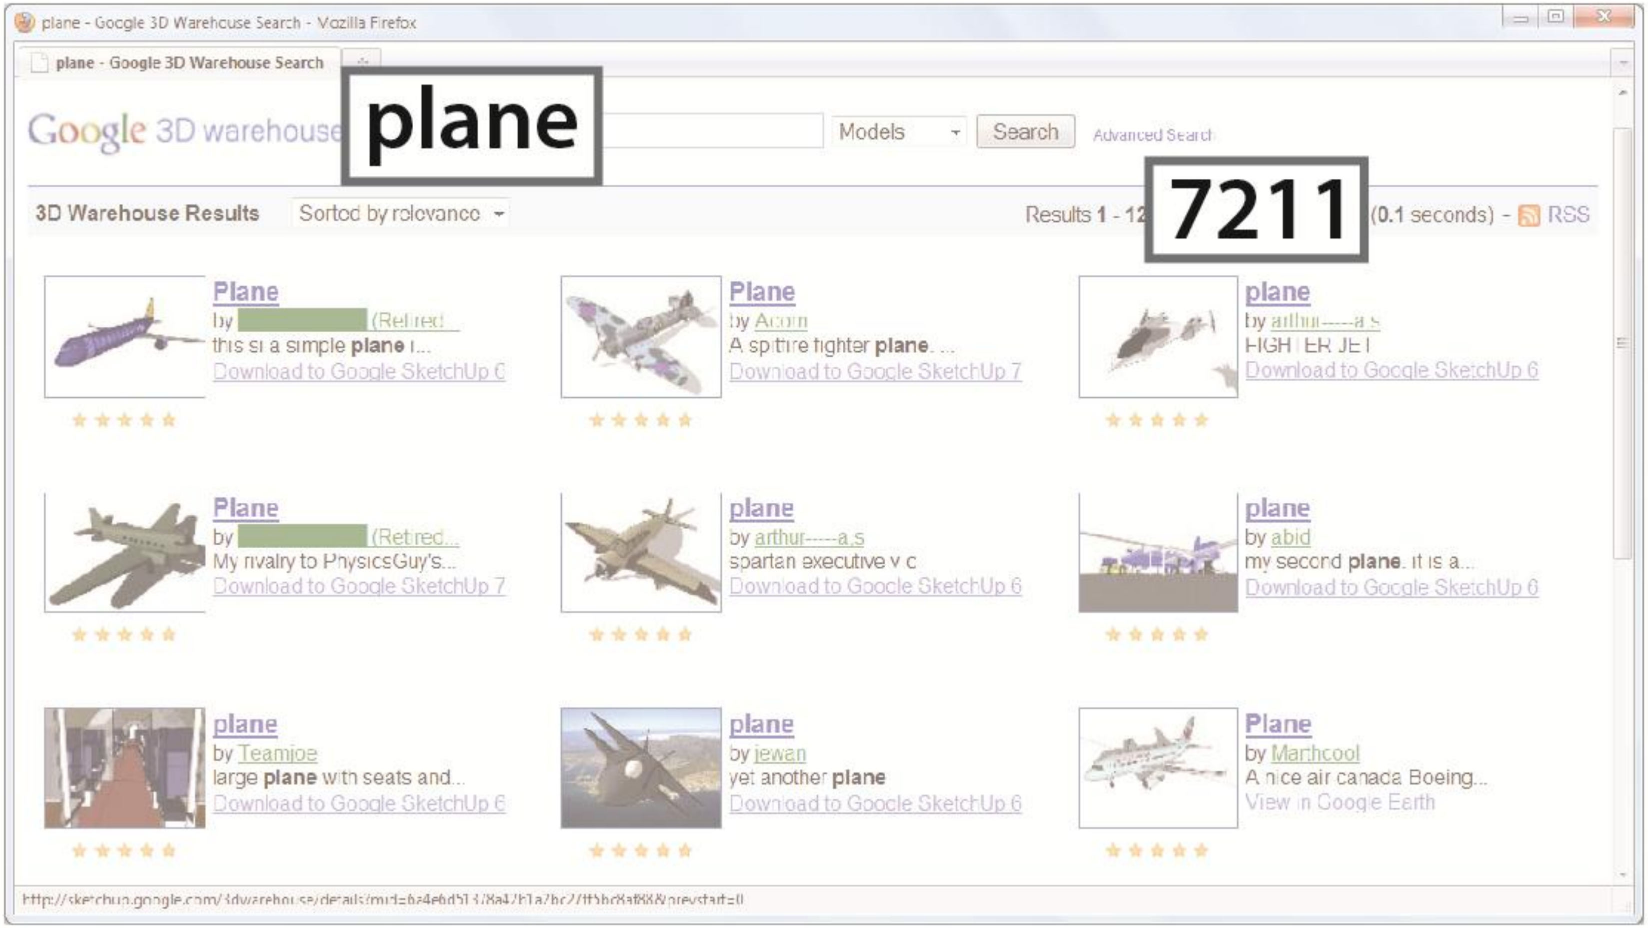

The Approach

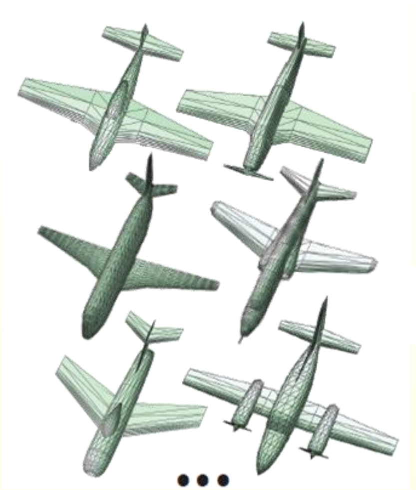

Input collection

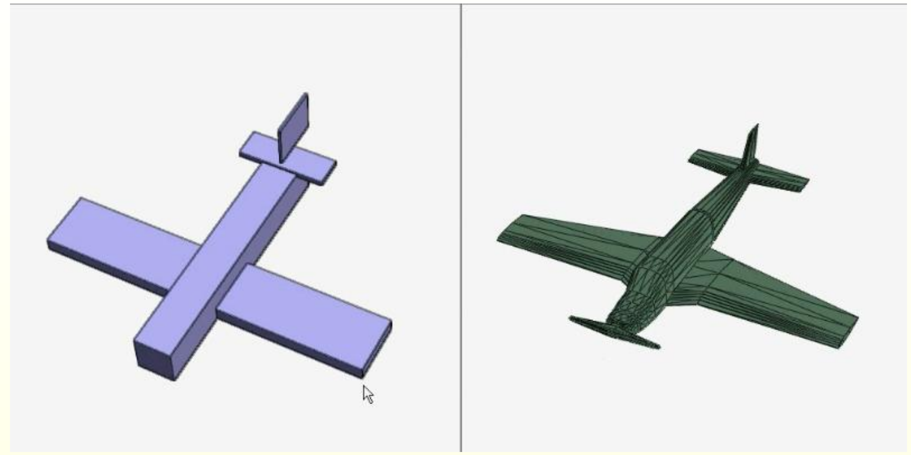

Exploration via deformation

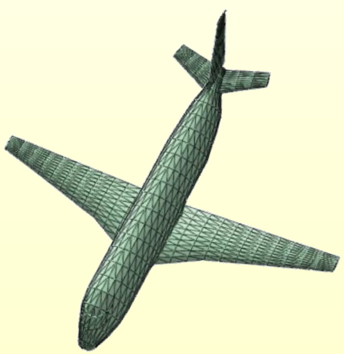

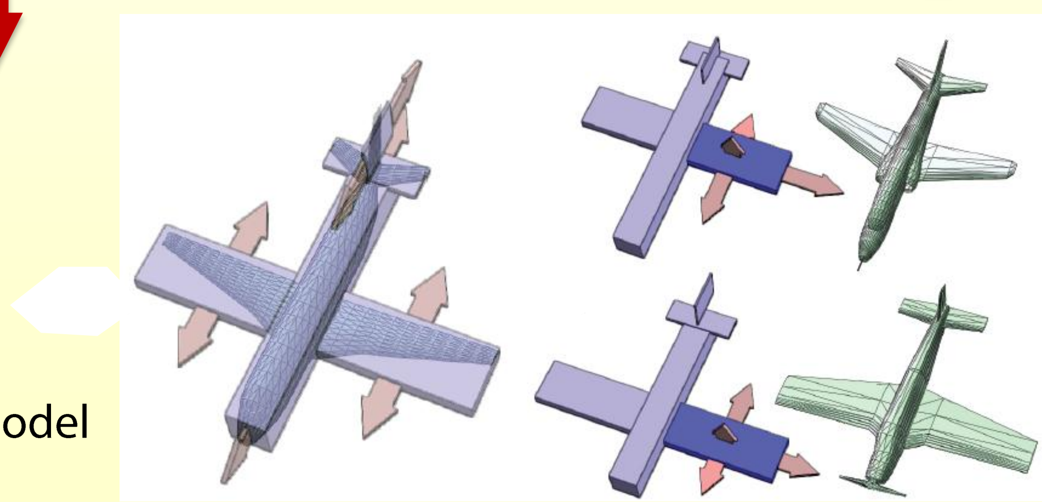

Template

with Deformation Model

$$\begin{tikzcd} S_1 \arrow[r] \arrow[d] & S_2 \arrow[d] \\ D_1 \arrow[r] & D_2 \end{tikzcd}$$

Analysis Stages

- Convert to descriptor space

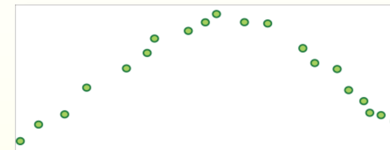

- Select template

- Deform to fit observed variability

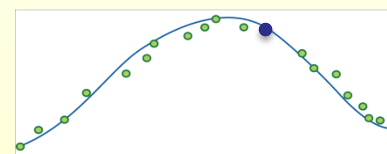

- Generate morphable model

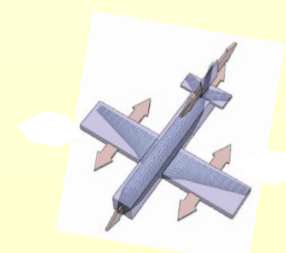

Deformation

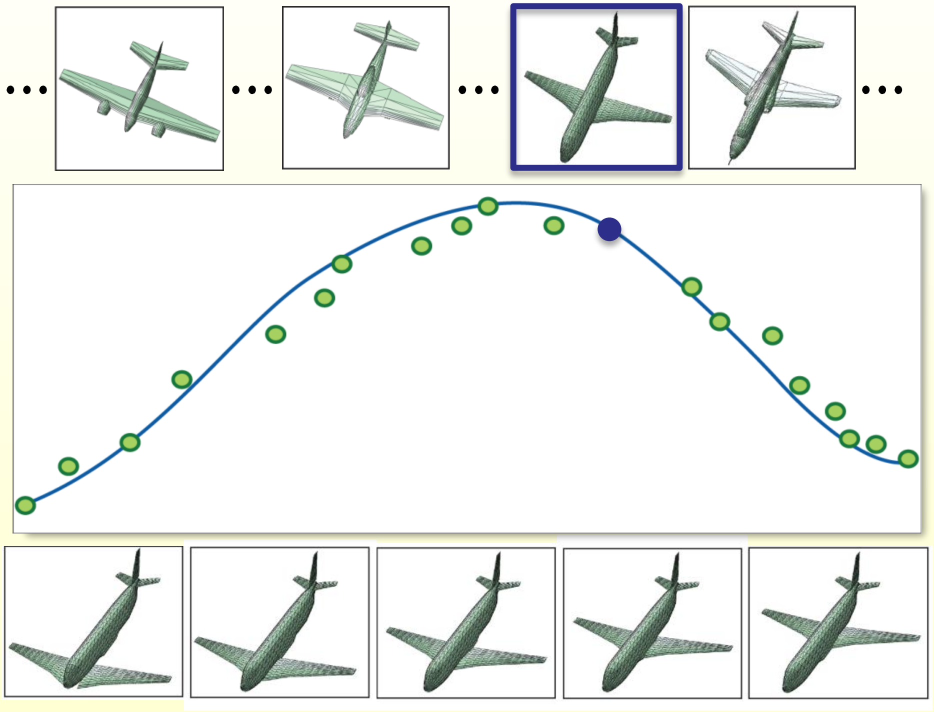

Descriptor space

But no orderings, no correspondences, no segmentations ...

Template Deformation Model

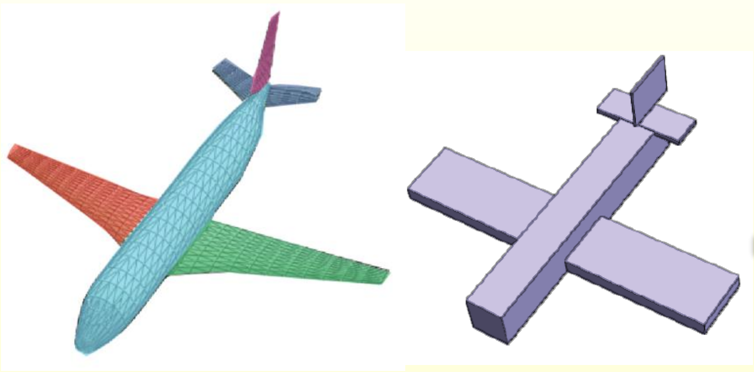

$$D \in \mathbb{R}^{6C}$$

D defines scale and translation for each of the C template components

$M_1(D)$

denotes shape M_1 deformed by D

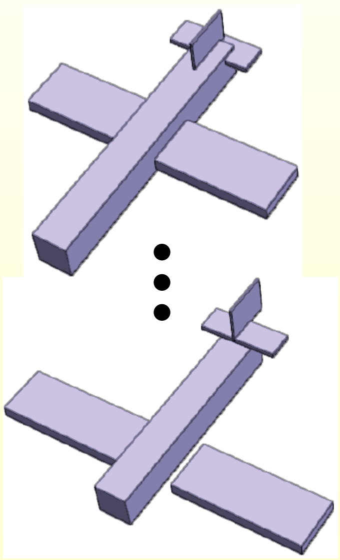

Choosing a Template Shape

- Remove outliers
- Compute mean descriptor
- Take closest shape (restrict number of mesh components)

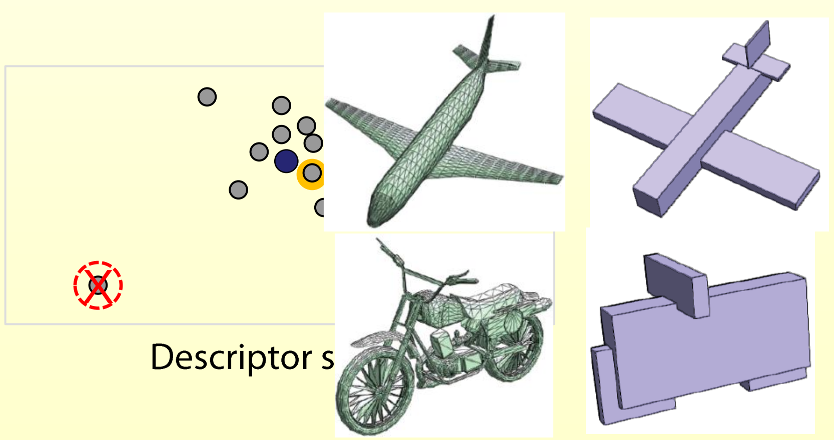

Deformable Model

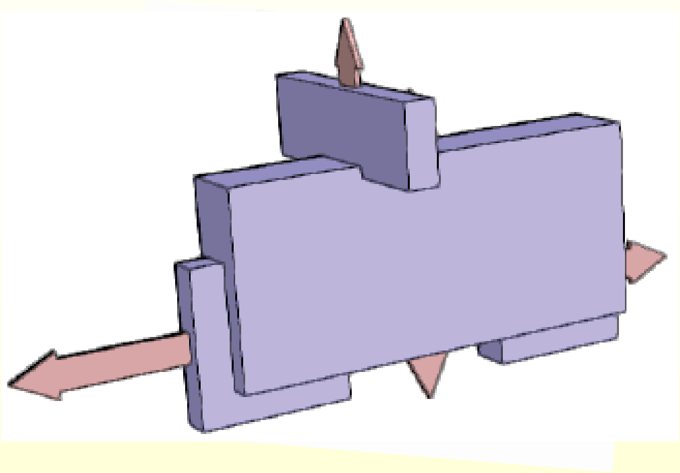

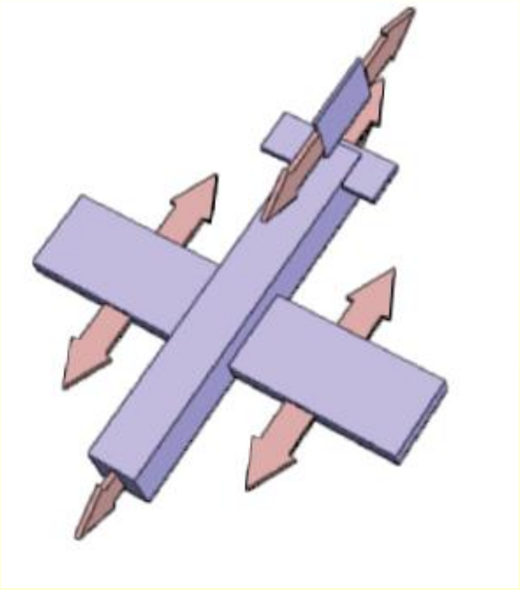

Exploration

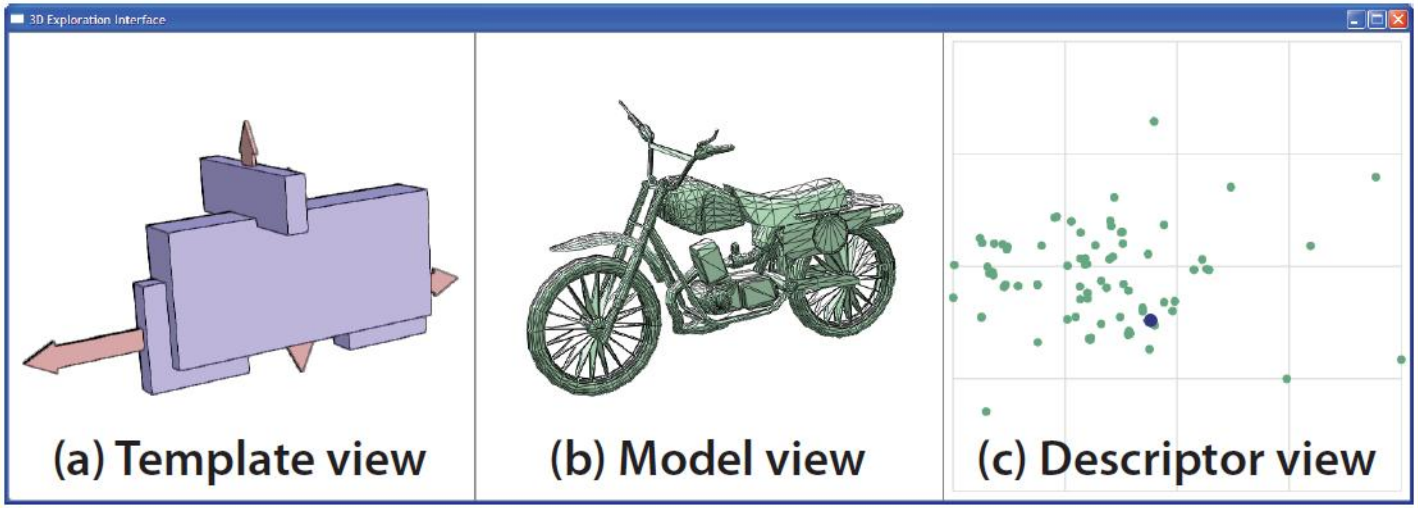

The Lessons

- Within a class, shape variability can be learned -- even without correspondences
- Shape collection navigation is just as important as shape search

Cancer Data Analysis via Mapper [Calrsson Group]

- Analysis of cancer genomic data to identify high survival groups using topological methods

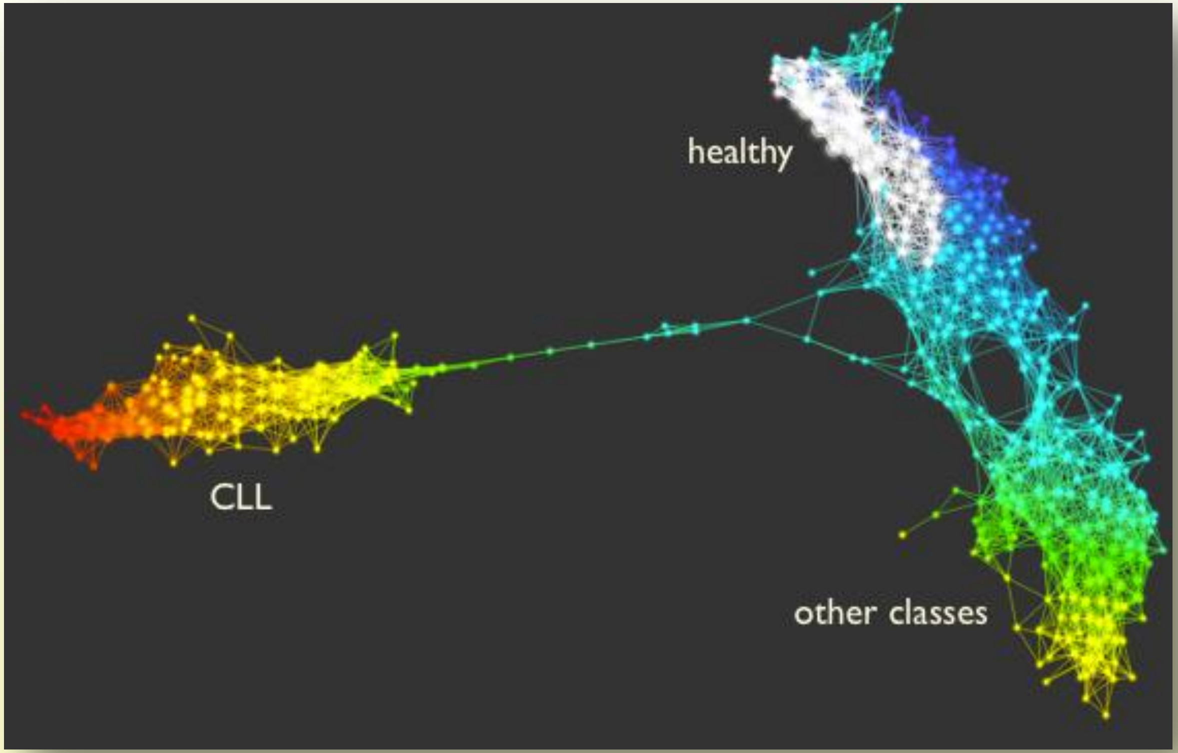

- Methods also applicable to social network analysis

PNNL Collaboration

- Topic: Morphological signatures for predicting nanoparticle biological interactions
- Shape of a nanoparticle affects:
  - Cellular internalization
  - Adhesion to surfaces
  - Transport in the body

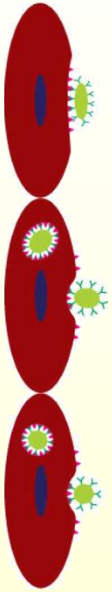

Collaboration with

Nathan Baker Lab

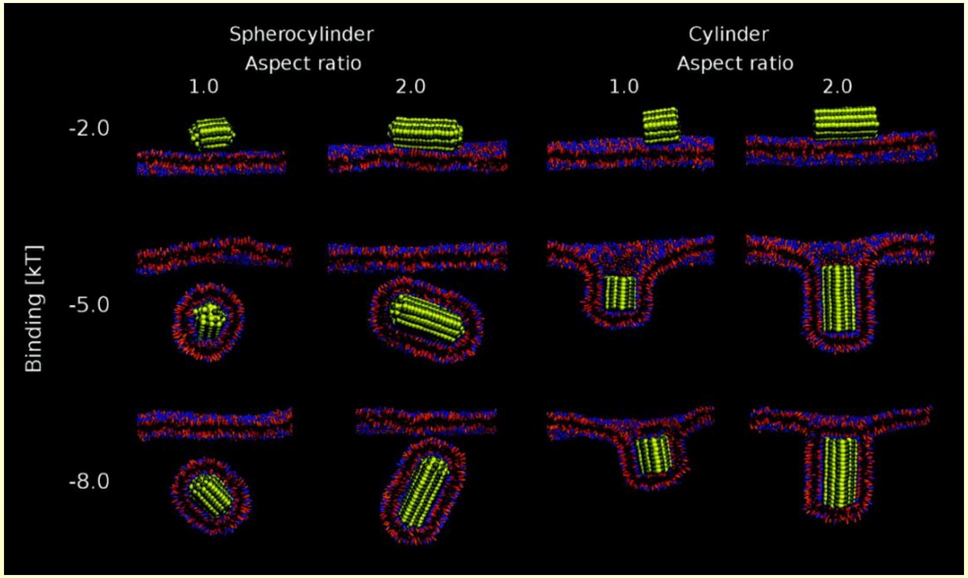

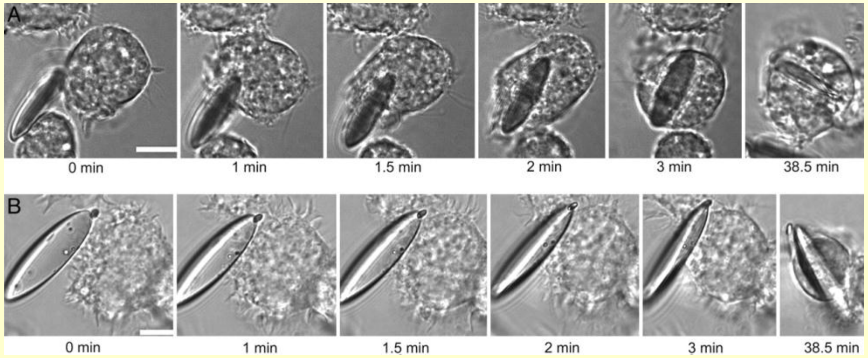

[From Vácha *et al*.: Endocytosis is suppressed for particles with sharp edges]

[From Champion *et al*.]

Acknowledgements

- Collaborators:
  - Current and past students: Daniel Chen, Peter Huang, Andy Nguyen, Maks Ovsjanikov
  - Current and past postdocs: Mirela Ben-chen, Adrian Butscher
  - Senior: Fred Chazal, Vladlen Koltun, Wilmot Li, Niloy Mitra

- Sponsors:

FODAVA

National Science Foundation
WHERE DISCOVERIES BEGIN

Homeland Security