Data-Driven RGBD Reconstruction

- Qixing Huang
- Nov. 12th 2018

Geometry Reconstruction

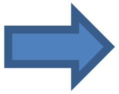

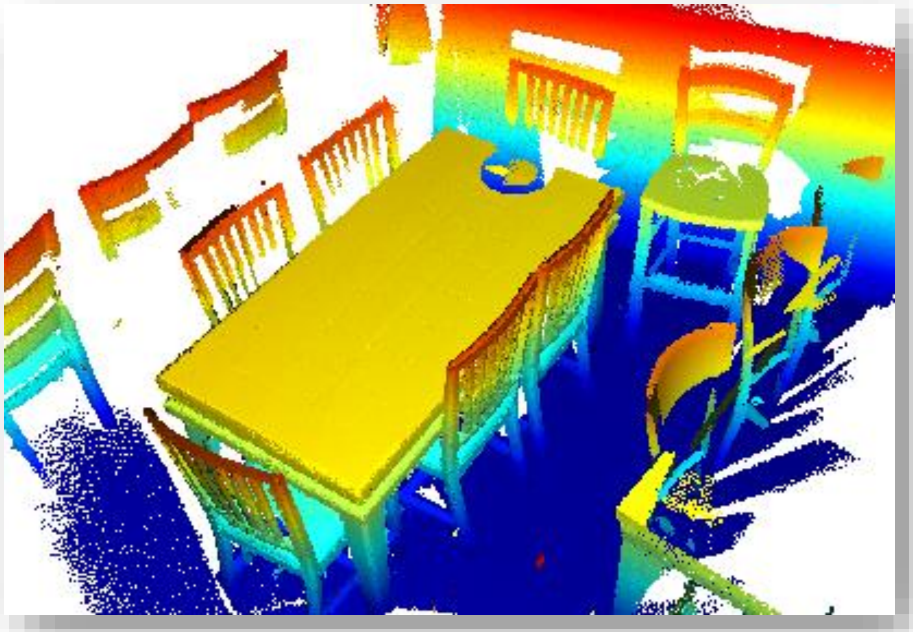

A Standard Approach

- Scanning
- Registration
- Reconstruction

A standard pipeline

A standard pipeline

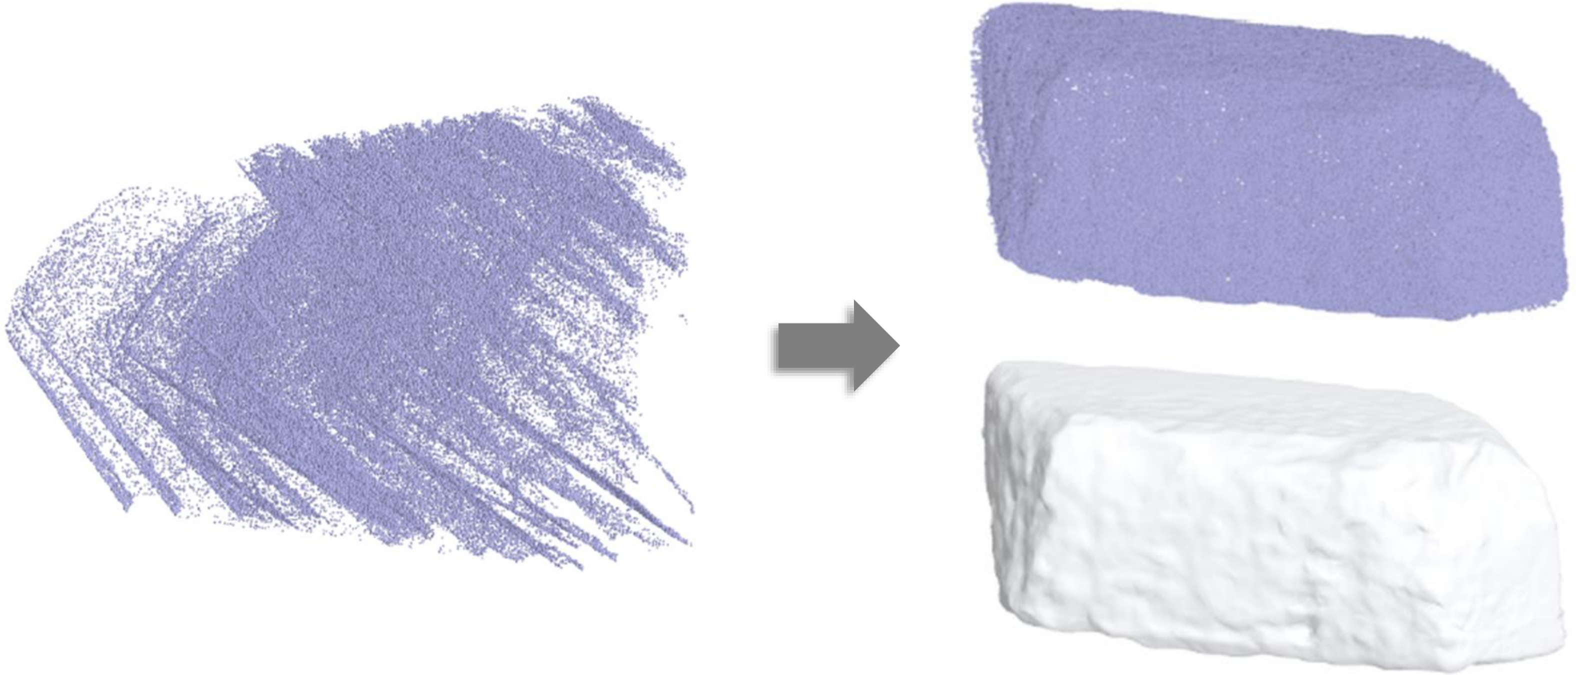

A standard pipeline

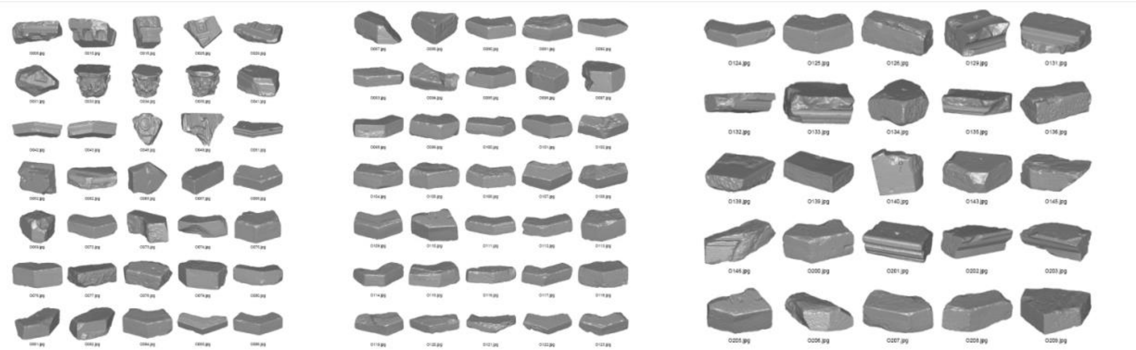

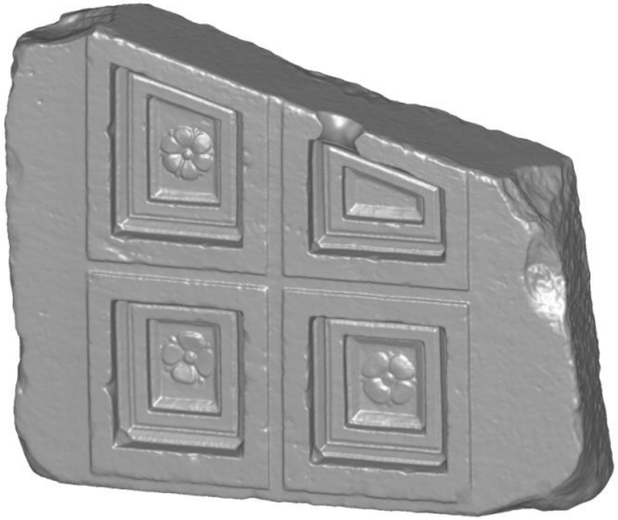

Limitation I – complete observation

Data-Driven
Geometry Reconstruction

The Big Bang in internet 3D models

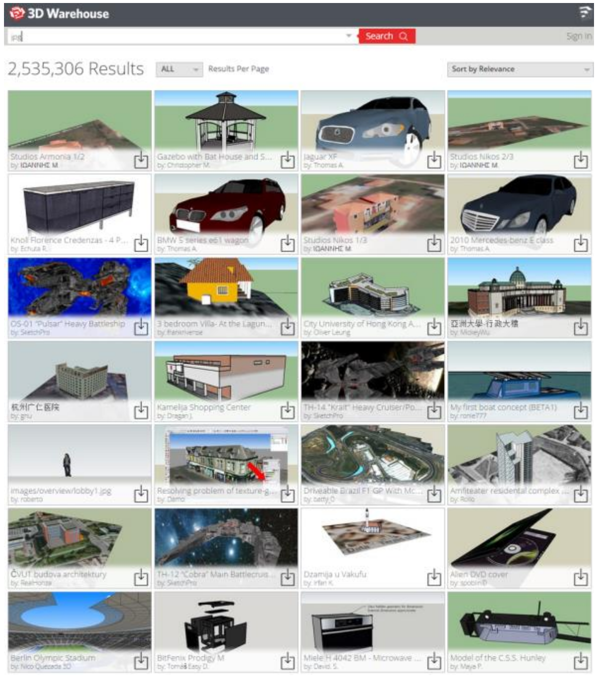

3D Warehouse

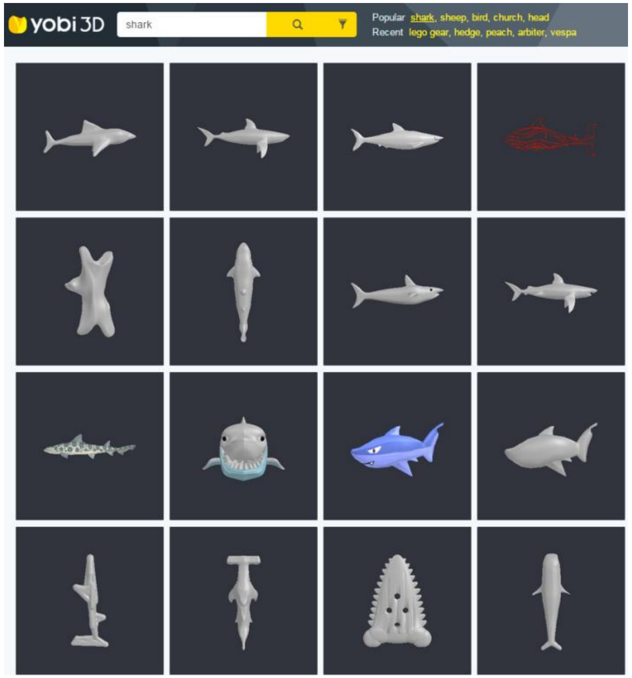

Yobi3D

3M models in more than 4K categories

Image-based shape retrieval

20 years ago

10 years ago

now

Single-view image based shape modeling

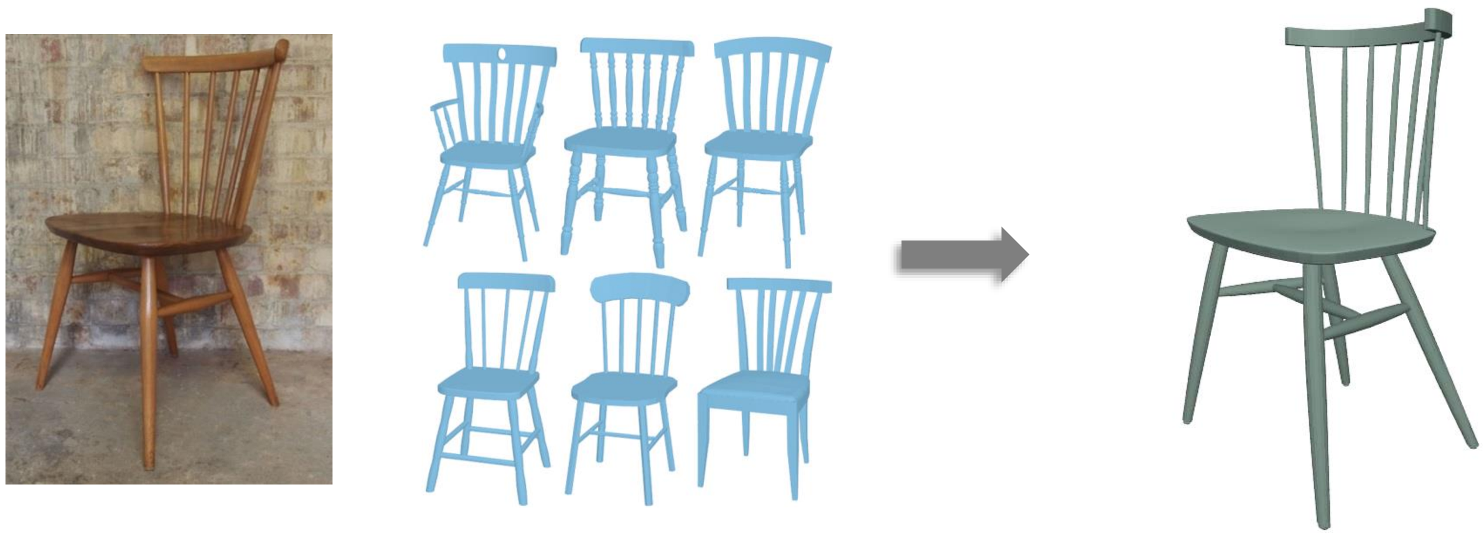

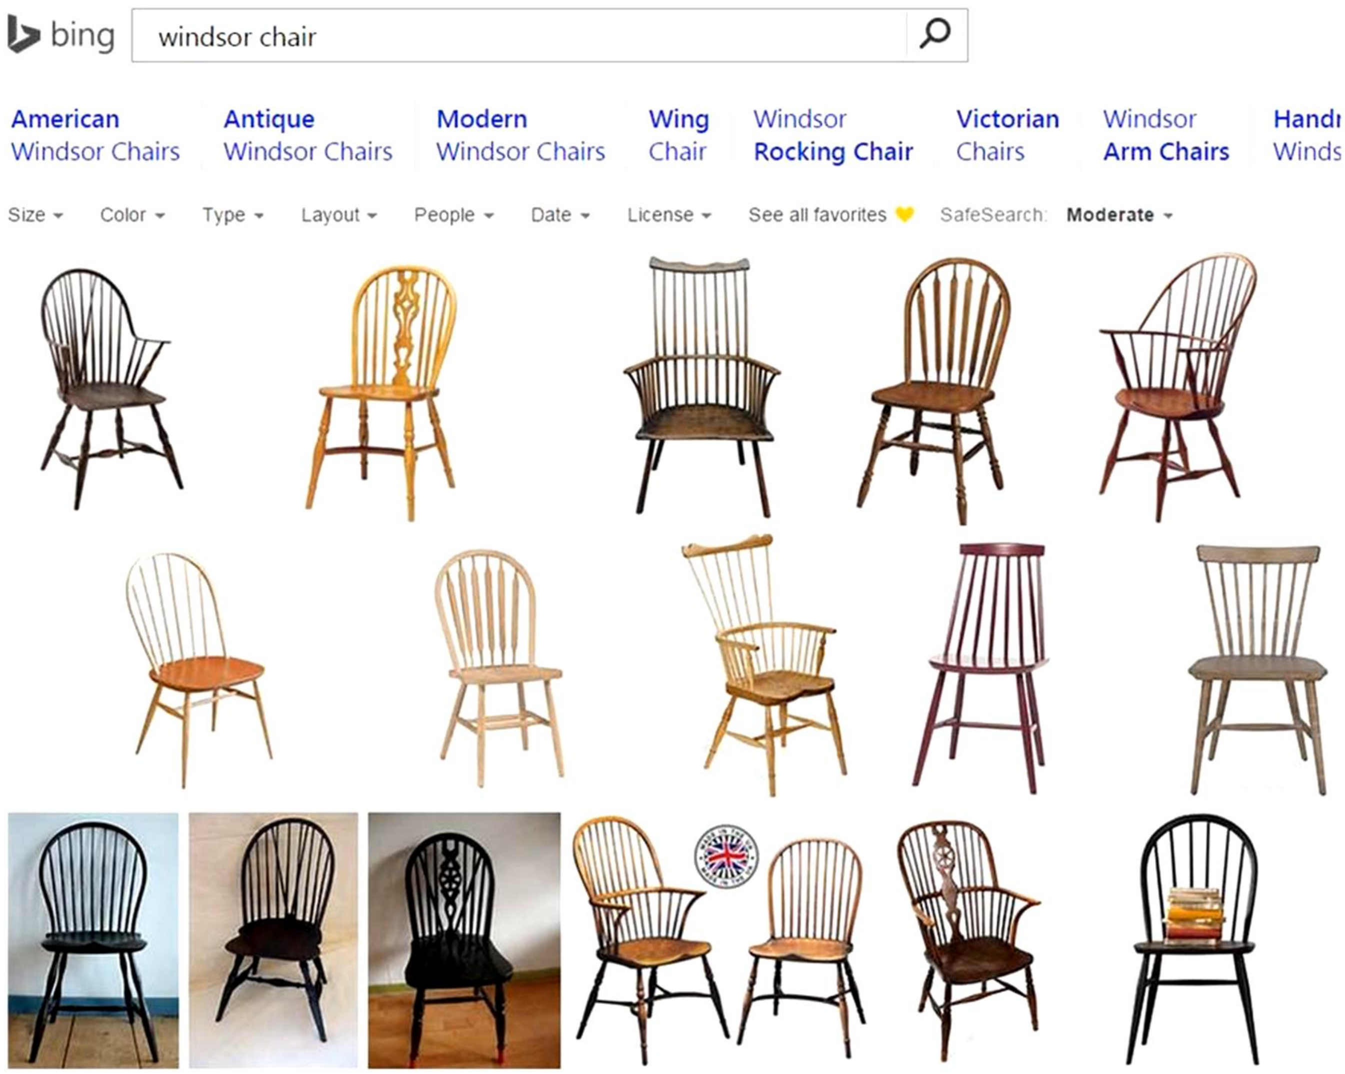

bing
windsor chair
American Windsor Chairs
Antique Windsor Chairs
Modern Windsor Chairs
Wing Chair
Windsor Rocking Chair
Victorian Chairs
Windsor Arm Chairs
Size
Color
Type
Layout
People
Date
License
See all favorites
SafeSearch: Moderate

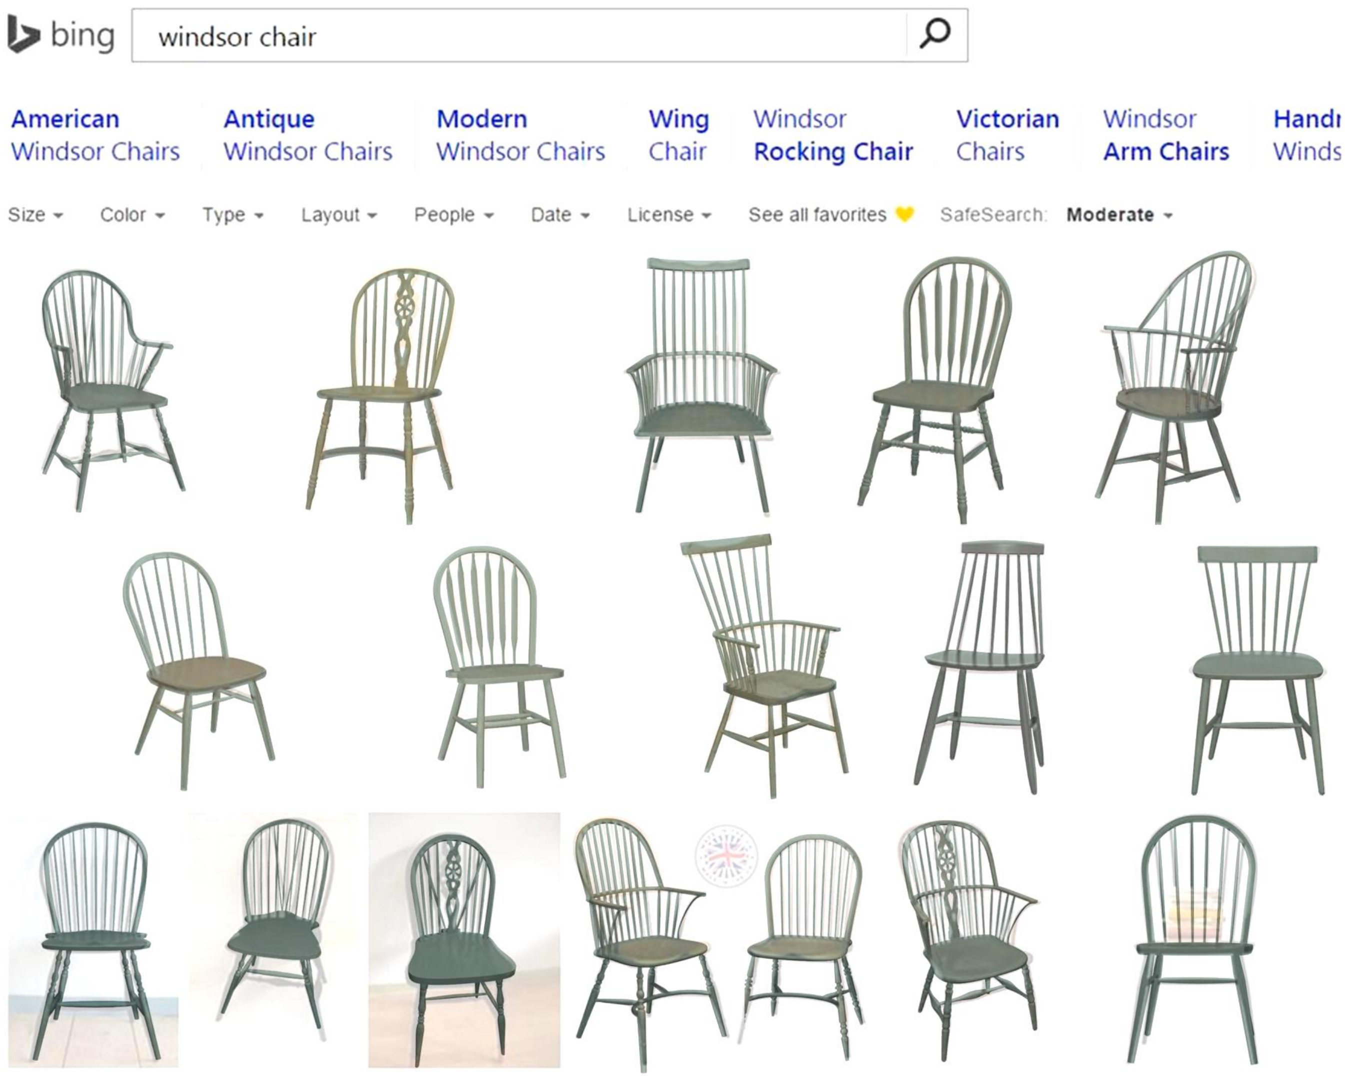
bing
windsor chair
American Windsor Chairs
Antique Windsor Chairs
Modern Windsor Chairs
Wing Chair
Windsor Rocking Chair
Victorian Chairs
Windsor Arm Chairs
Size
Color
Type
Layout
People
Date
License
See all favorites
SafeSearch: Moderate

The benefit of data-driven geometry processing

- Partial observation

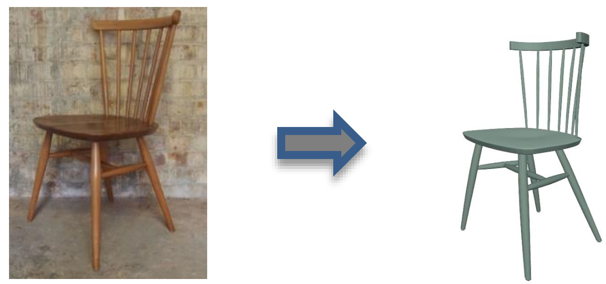

- Structural information

Classification

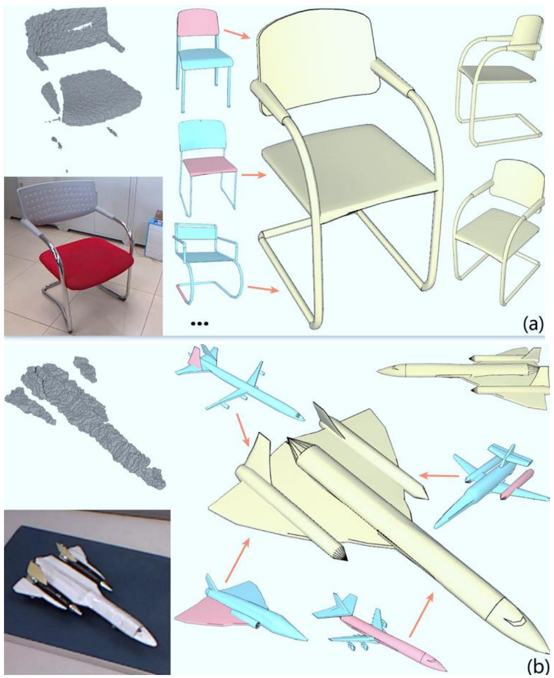

Shen et al. 12

Nearest Neighbor

Parametric Methods

Nearest Neighbor

Example-Based Scan Completion

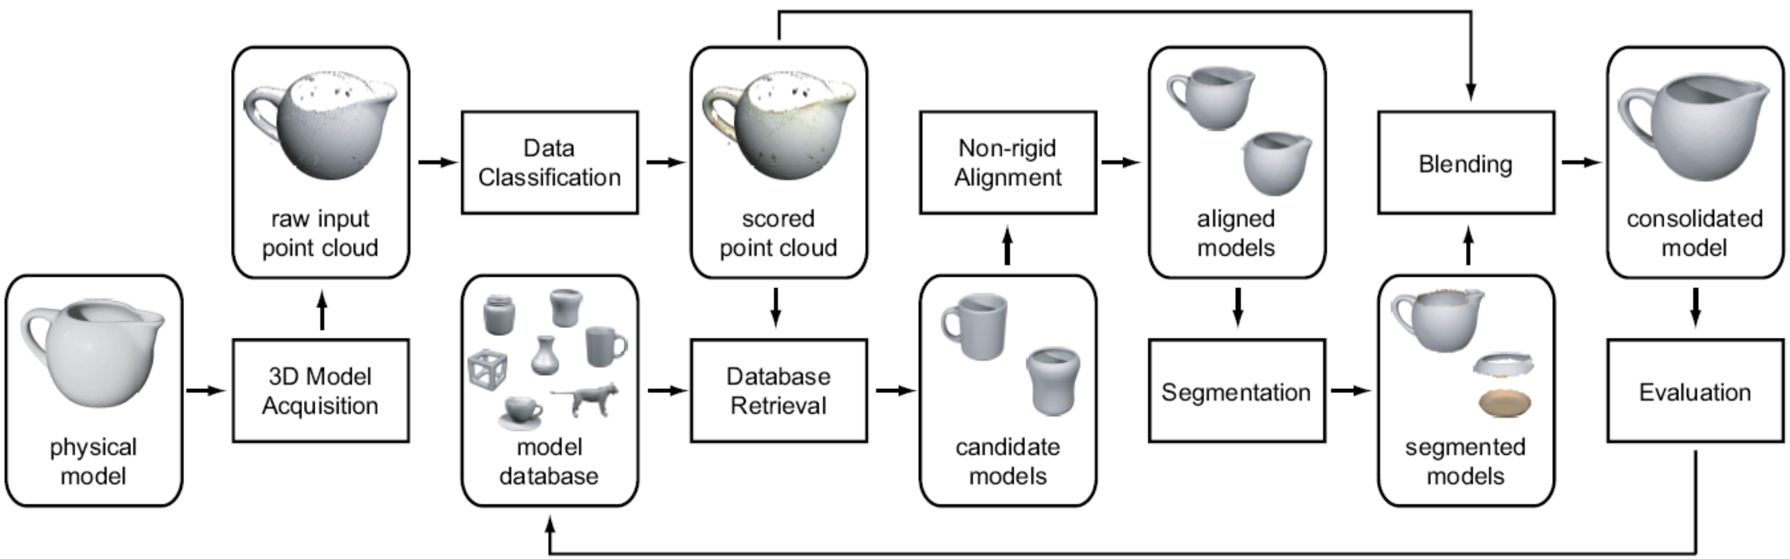

Example-Based 3D Scan Completion, SGP'05

Examples

Final Model

Context Models

Deformed Models

Part-based Shape Reconstruction [TOG'12]

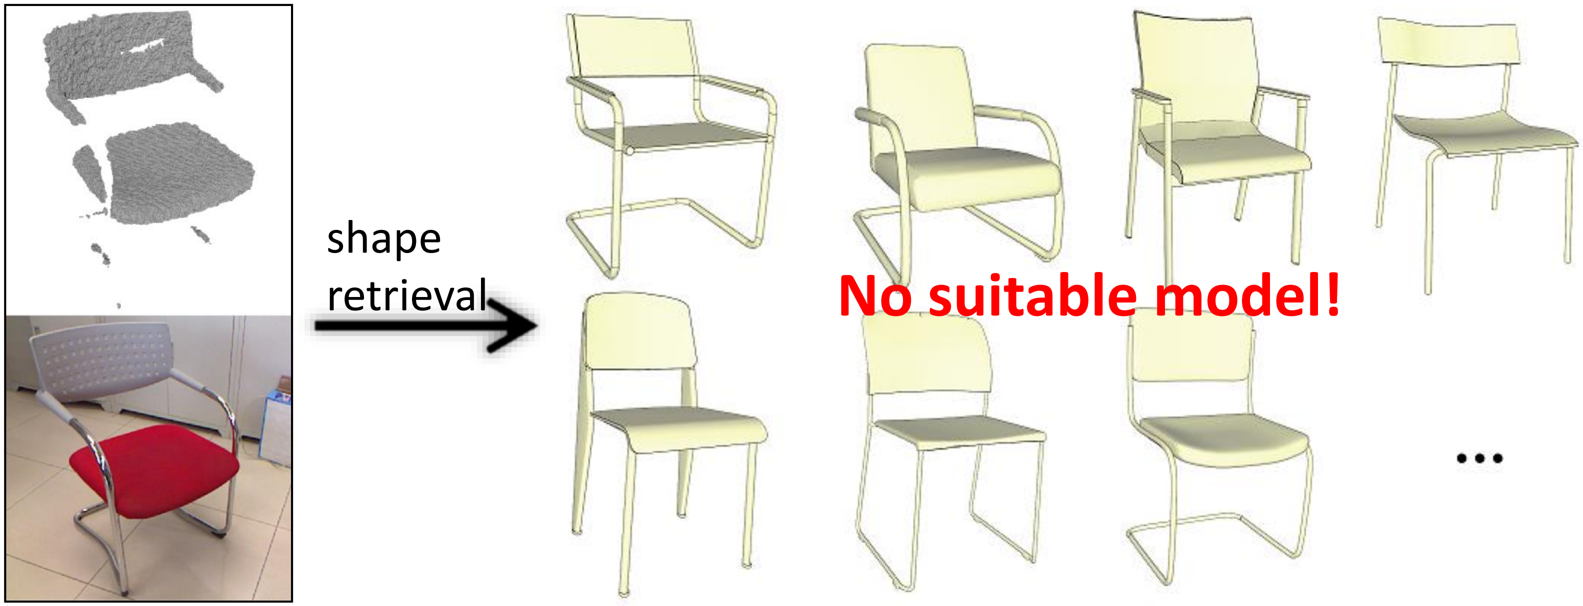

Structure Recovery by Part Assembly TOG'12

Recover the structure by part assembly

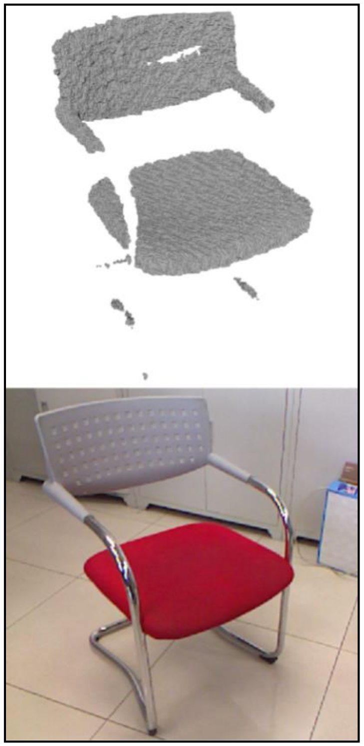

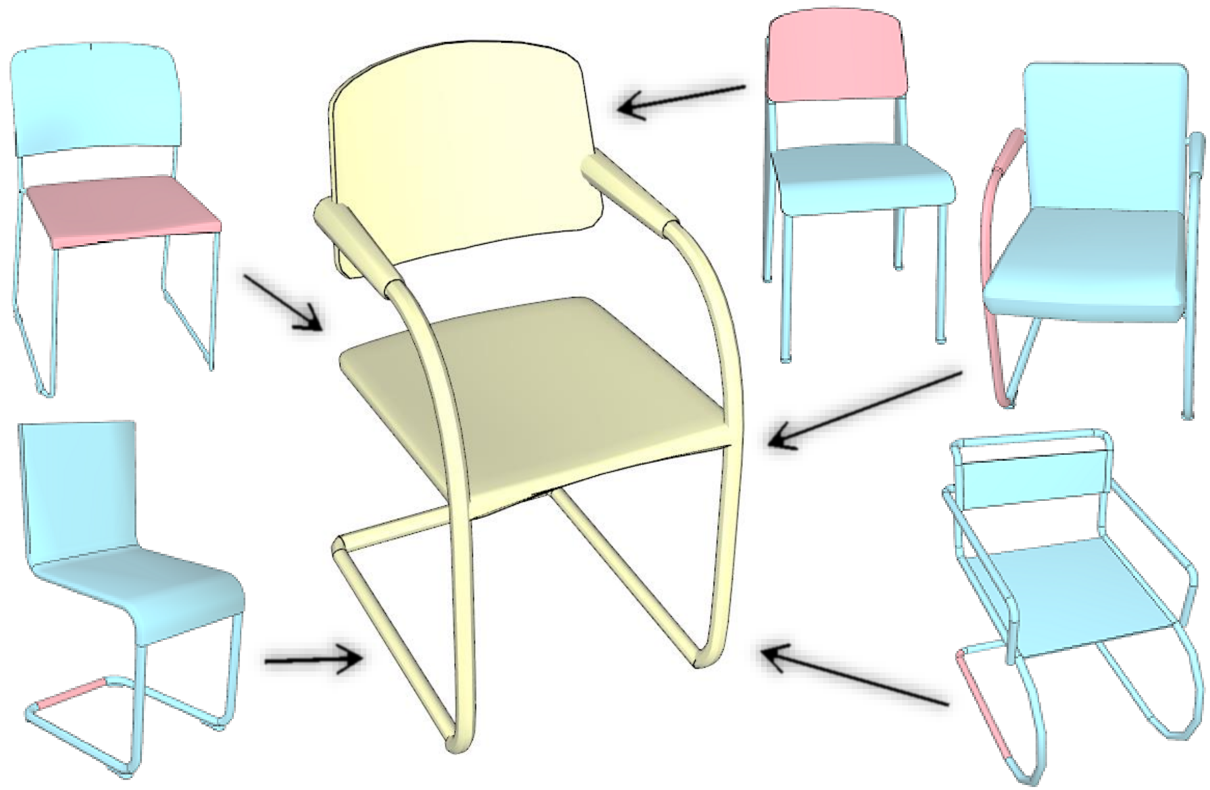

Algorithm Overview

Algorithm Overview

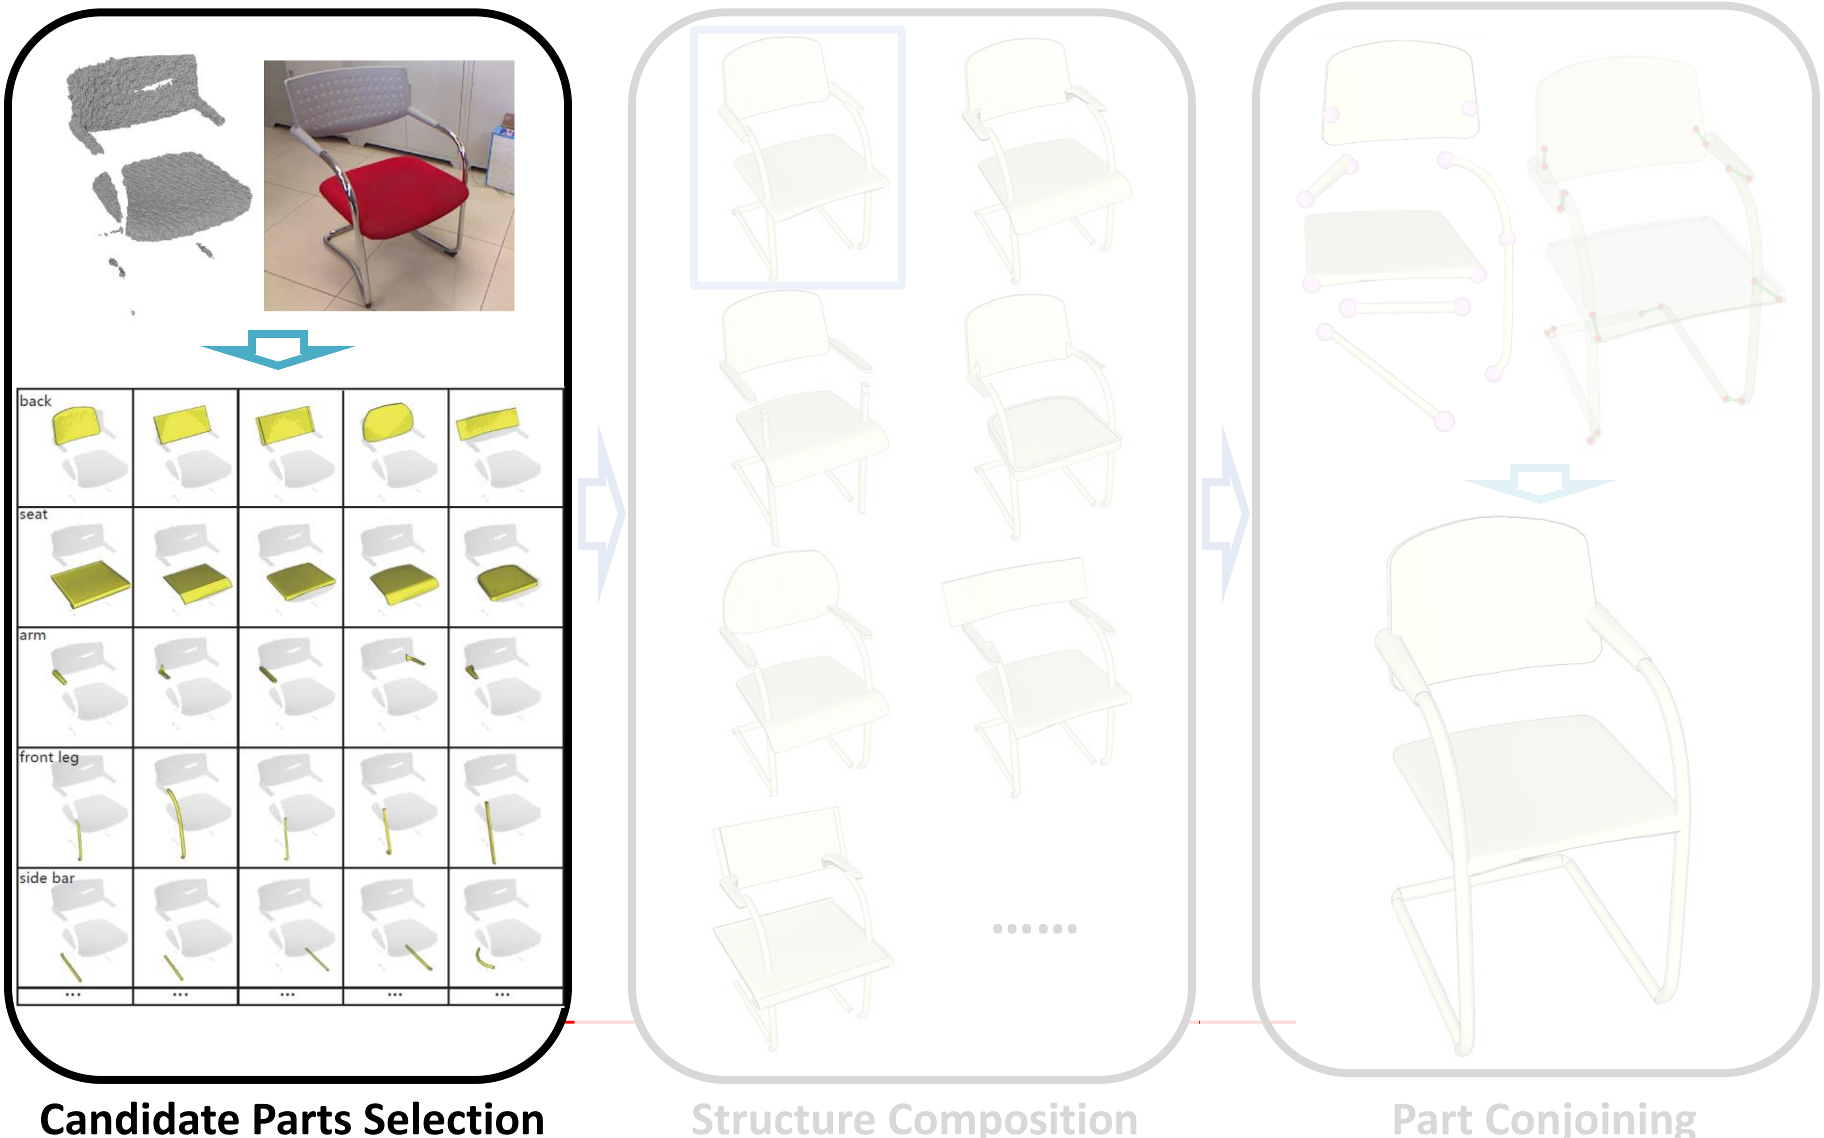

Algorithm Overview

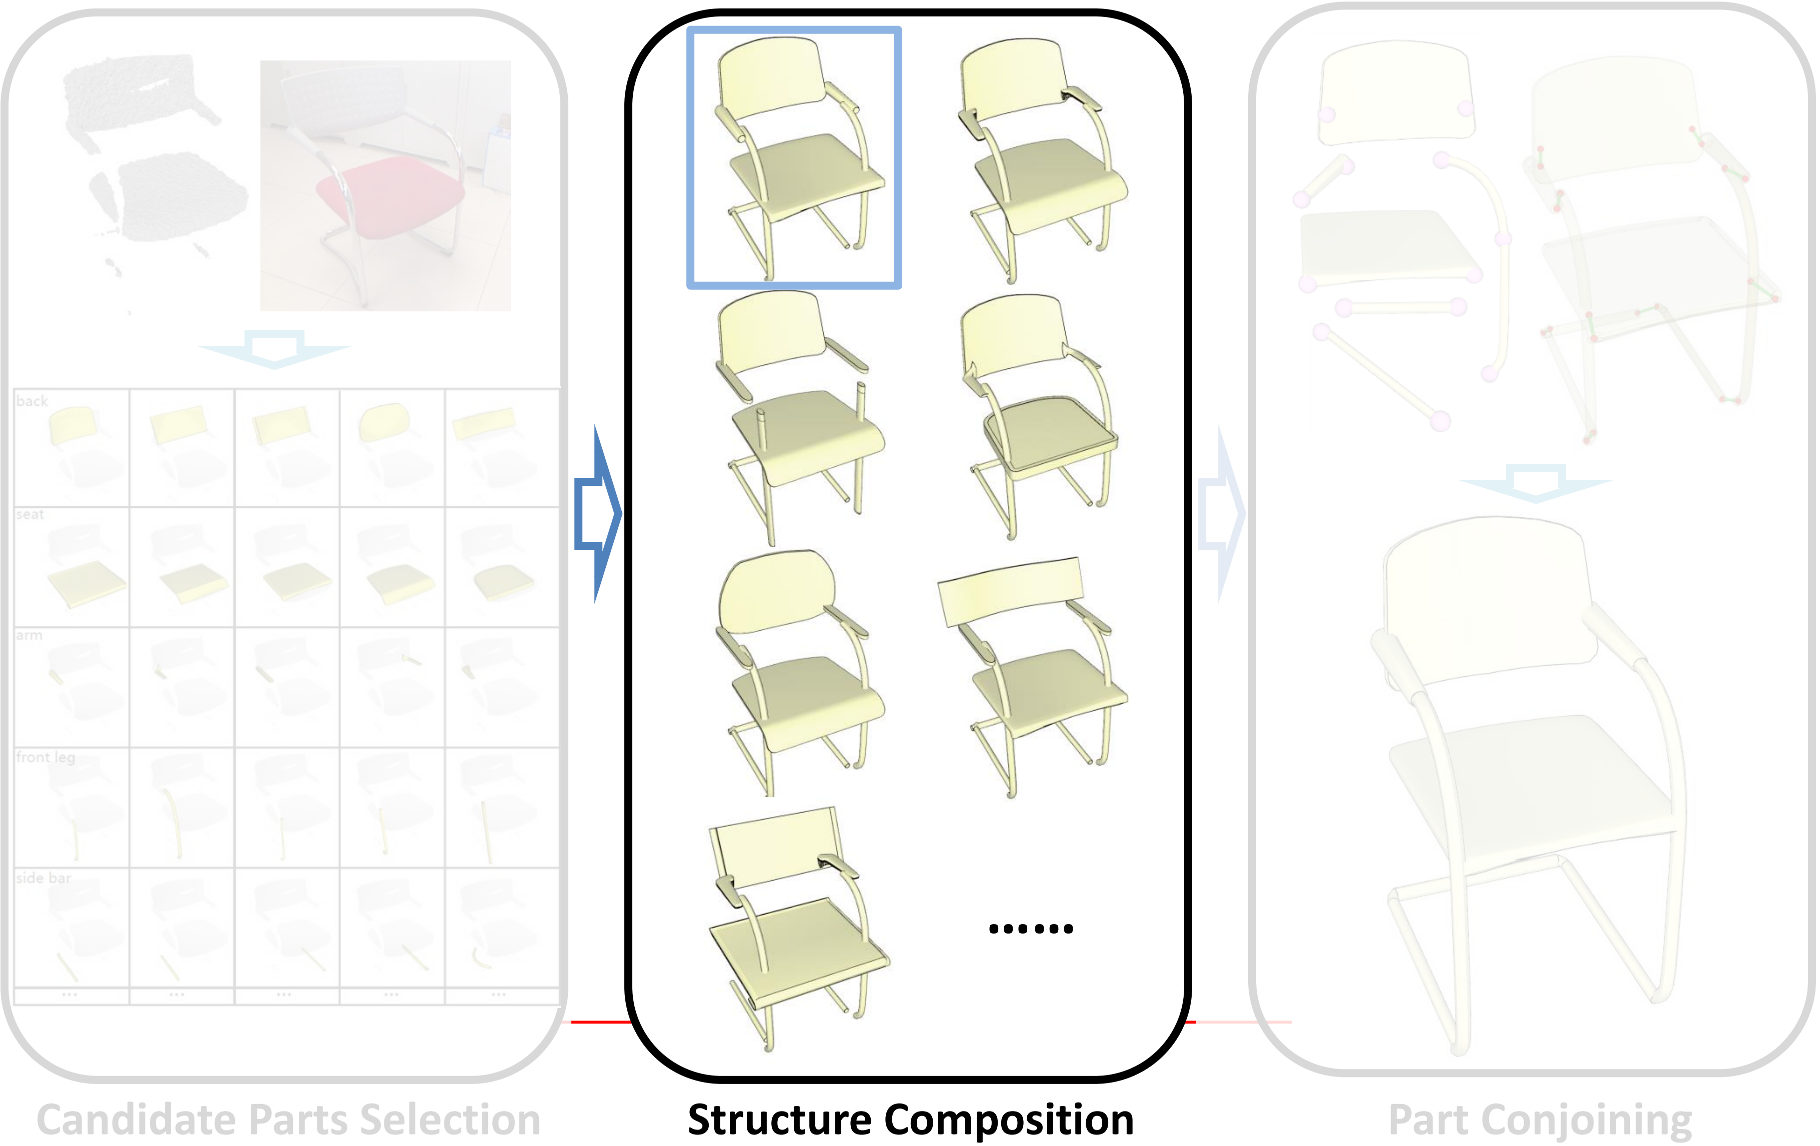

Algorithm Overview

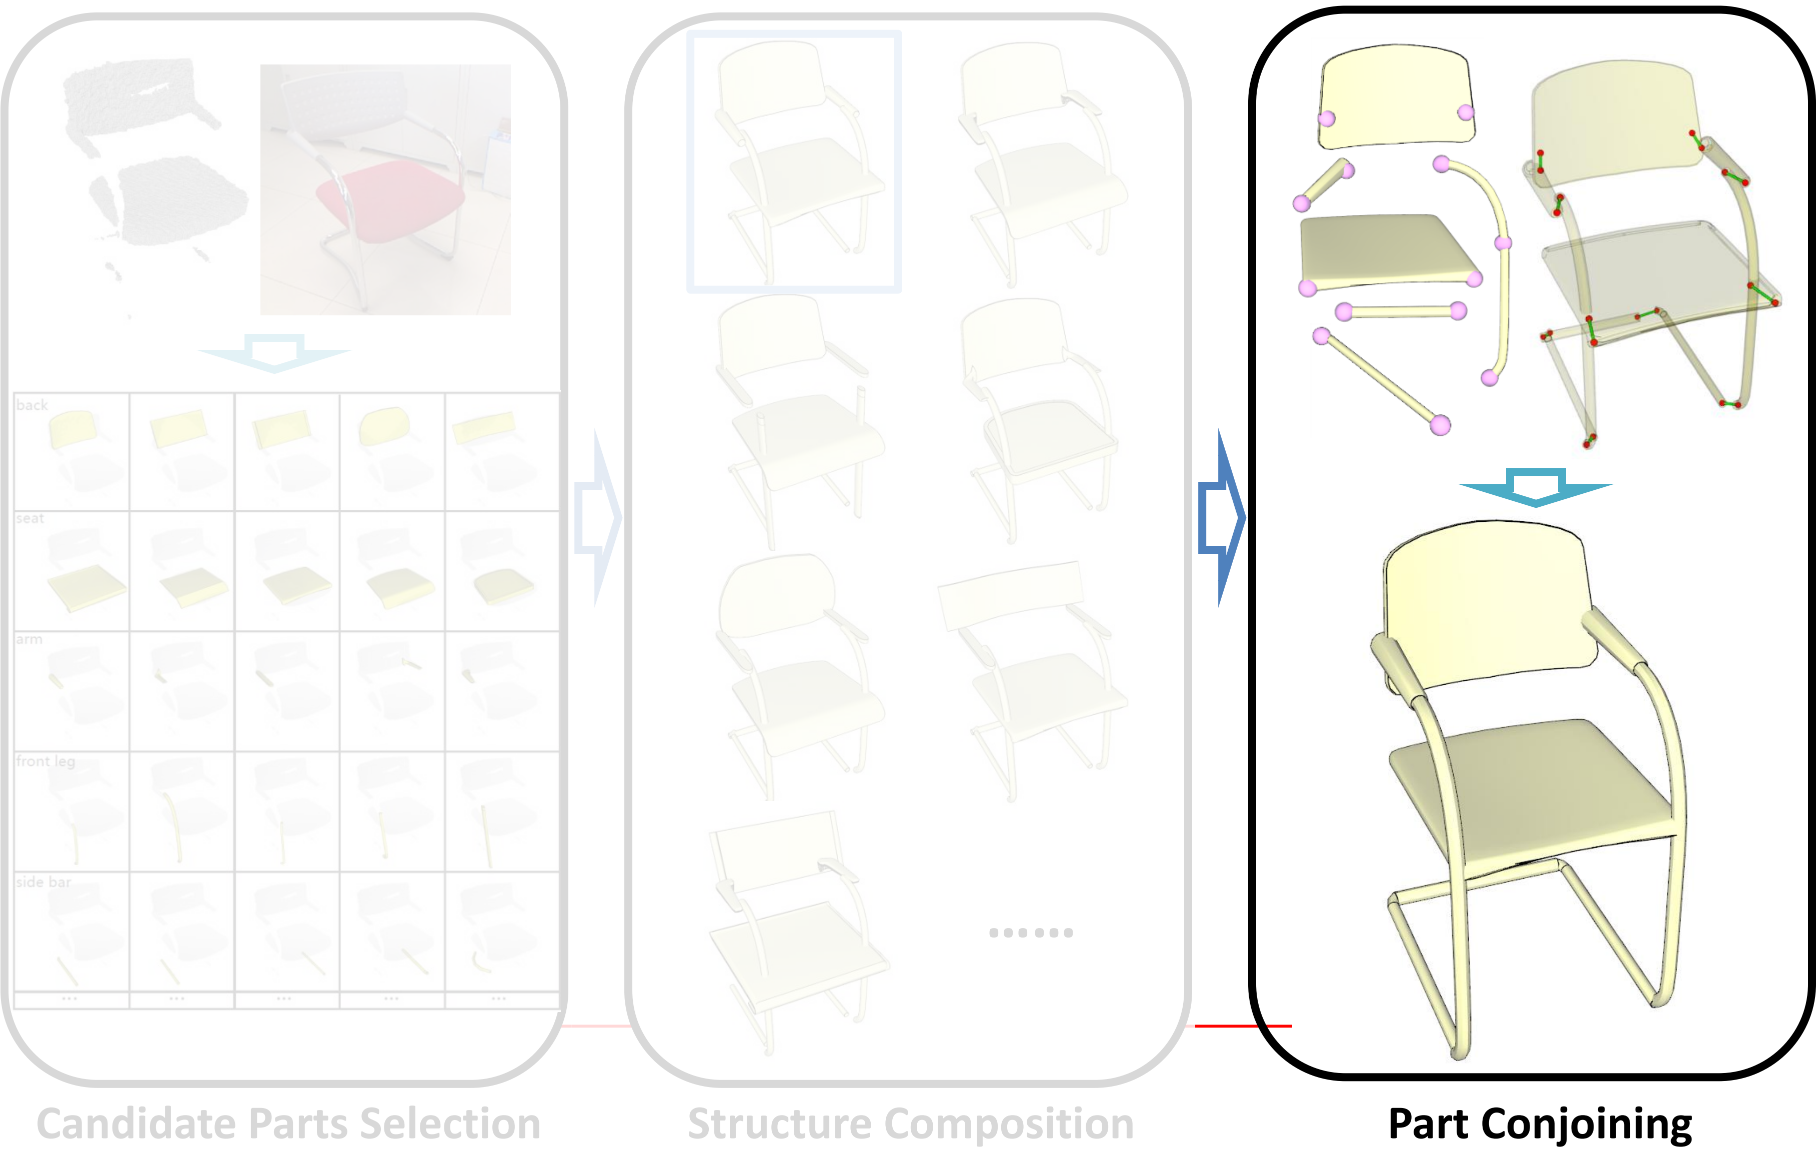

Results: Chairs

- 70 repository models, 11 part categories

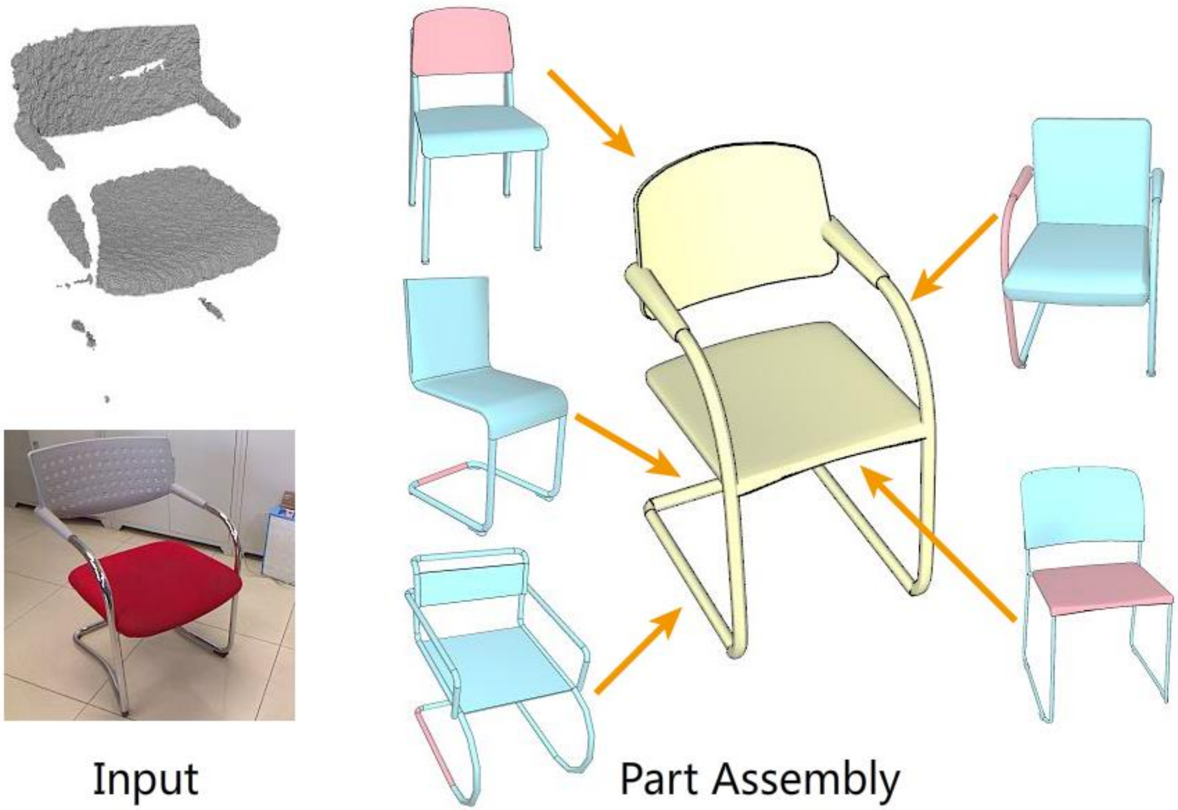

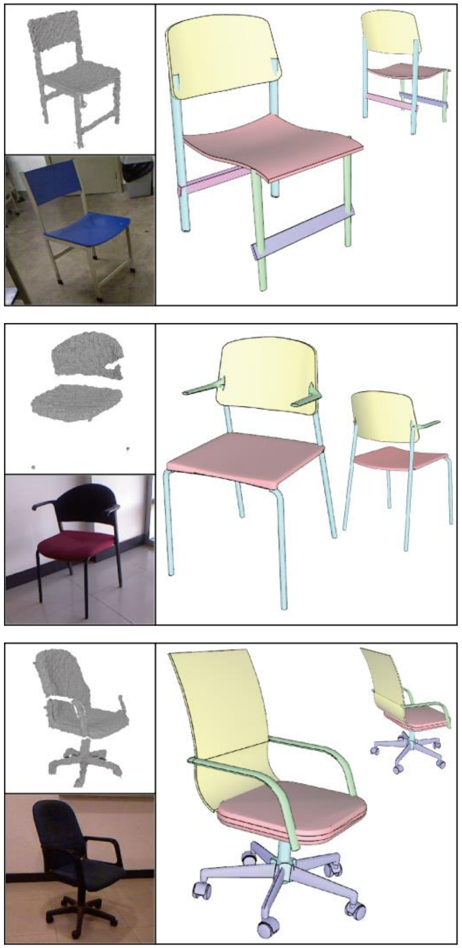

Results: Tables

- 61 repository models, 4 part categories

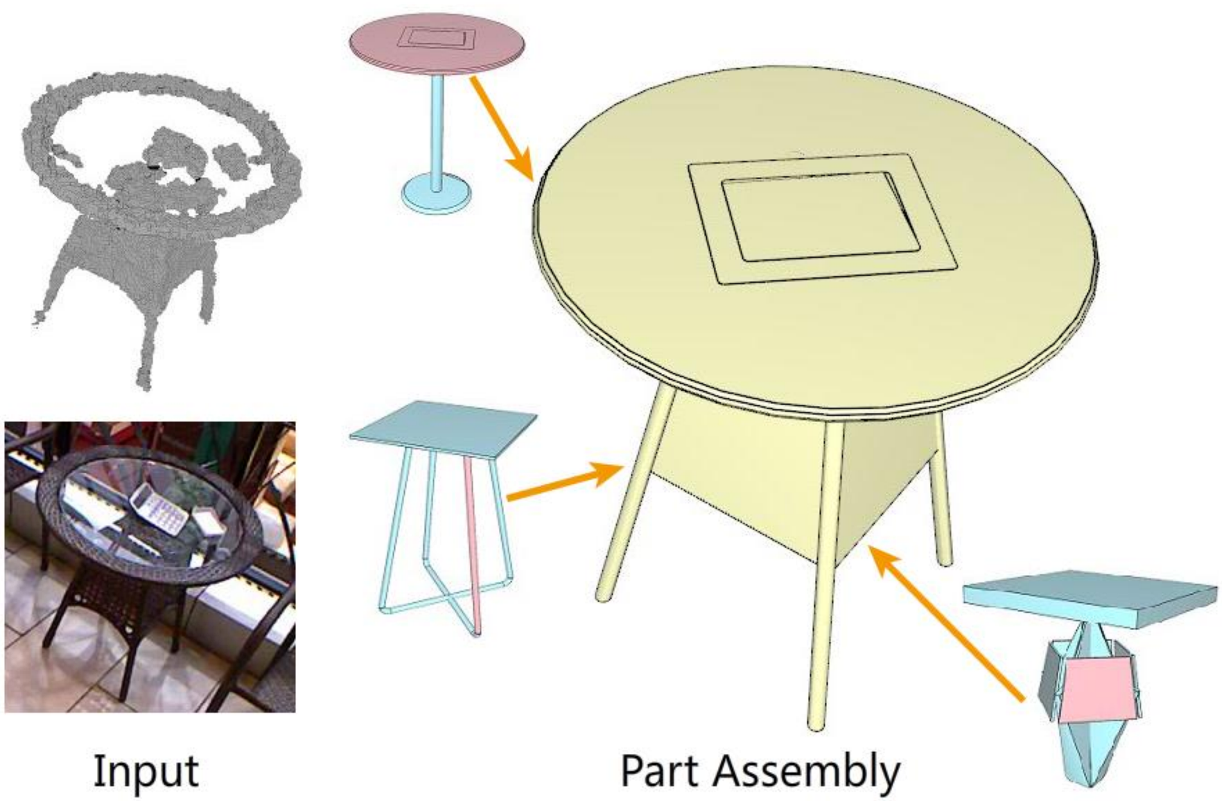

Results: Bicycles

- 38 repository models, 9 part categories

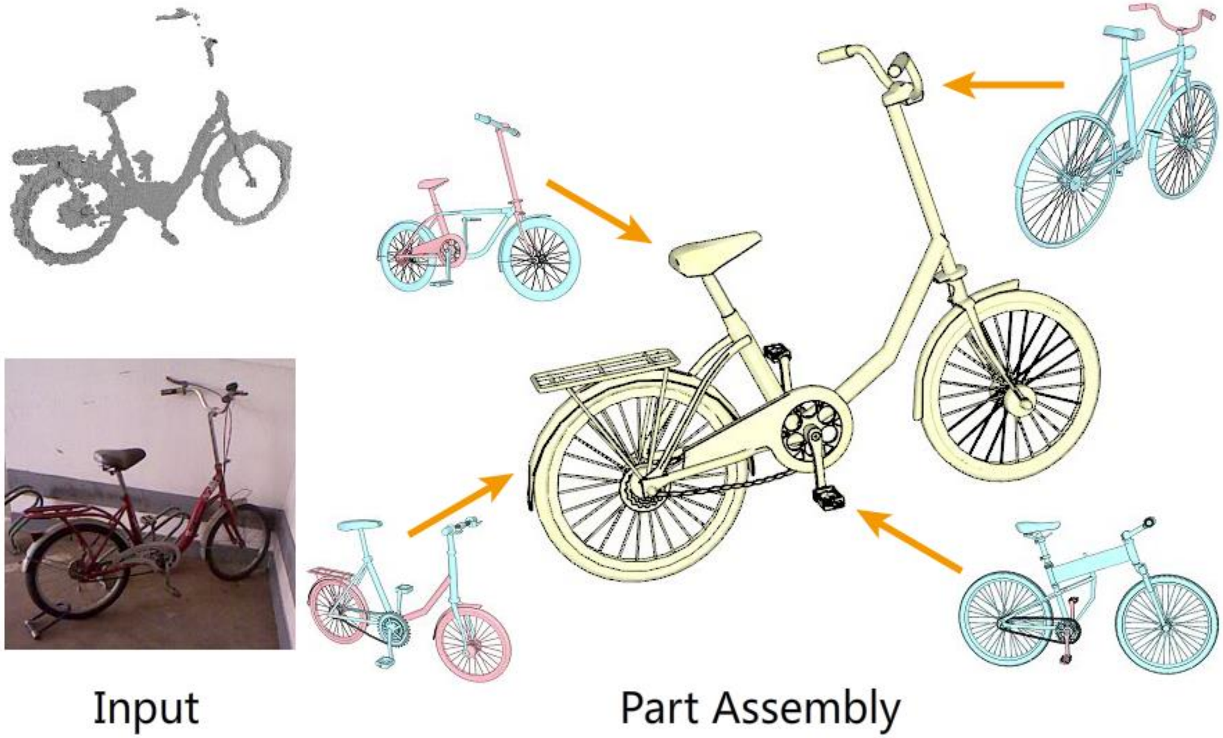

Results: Airplanes

- 70 repository models, 6 part categories

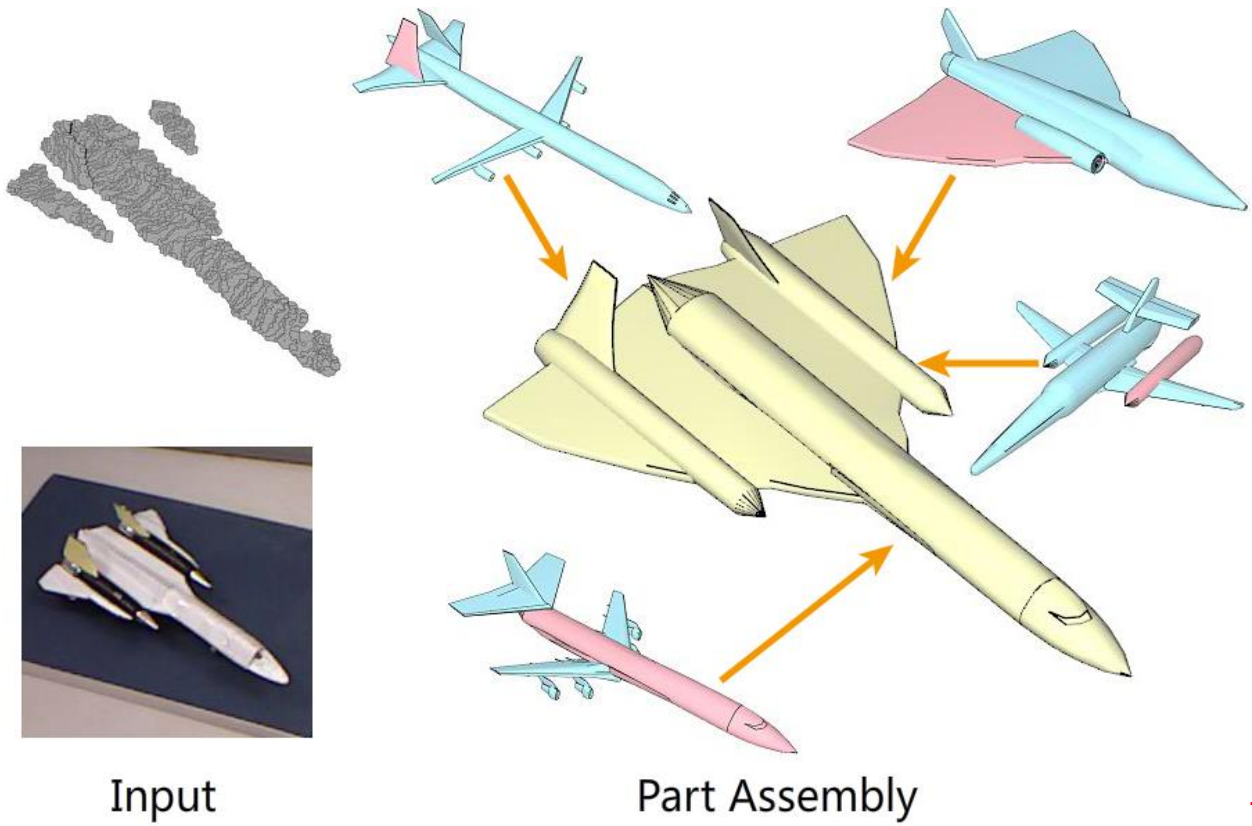

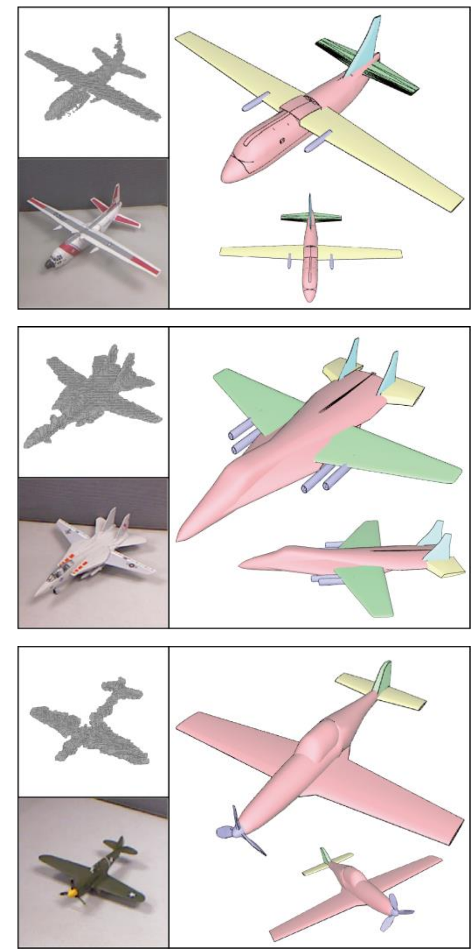

Discussion

- Hard to make it fully automatic --- many parameters to tune
- More data -> better algorithm
- Easy to add user interaction

Parametric Methods

A Morphable model for the synthesis of 3D faces

- Start with a catalogue of 200 3D Cyberware scans

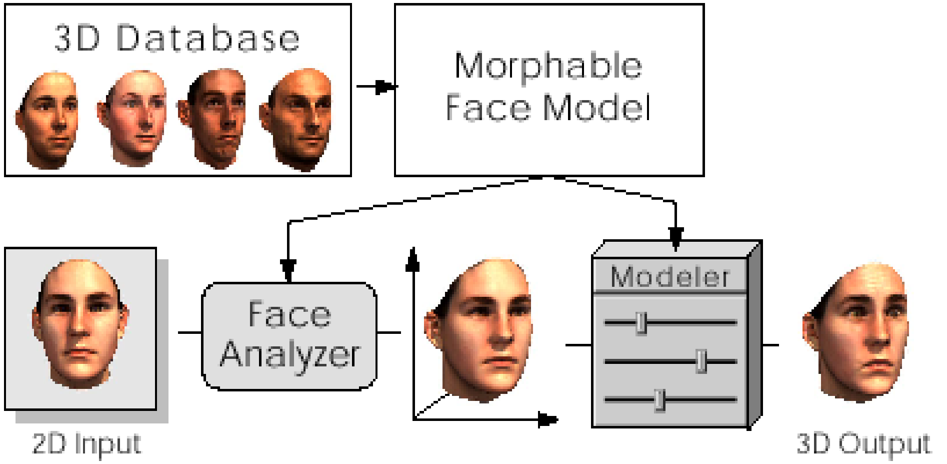

- Build a model of *average* shape and texture, and principal *variations*

Morphable 3D face model

$$\mathbf{S}_{mod} = \sum_{i=1}^{m} a_i \mathbf{S}_i \,, \quad \mathbf{T}_{mod} = \sum_{i=1}^{m} b_i \mathbf{T}_i \,, \quad \sum_{i=1}^{m} a_i = \sum_{i=1}^{m} b_i = 1.$$

$$\vec{a} = (a_1, a_2 \ldots a_m)^T \qquad \vec{b} = (b_1, b_2 \ldots b_m)^T$$

$$S_{model} = \overline{S} + \sum_{i=1}^{m-1} \alpha_i s_i \,, \quad T_{model} = \overline{T} + \sum_{i=1}^{m-1} \beta_i t_i \,, \qquad (1)$$

The probability for coefficients $\vec{\alpha}$ is given by

$$p(\vec{\alpha}) \sim exp[-\frac{1}{2} \sum_{i=1}^{m-1} (\alpha_i / \sigma_i)^2], \qquad (2)$$

Adding attributes

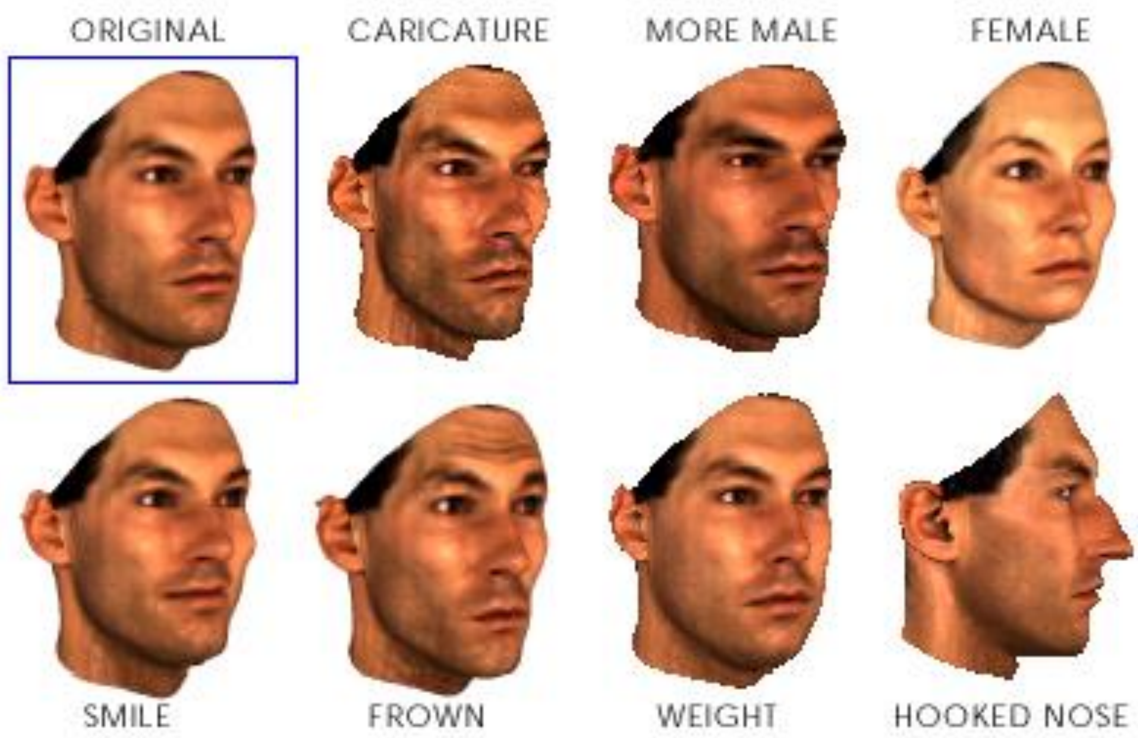

$$\Delta S = \sum_{i=1}^{m} \mu_i (S_i - \overline{S}) \qquad \Delta T = \sum_{i=1}^{m} \mu_i (T_i - \overline{T})$$

Reconstruction from single image

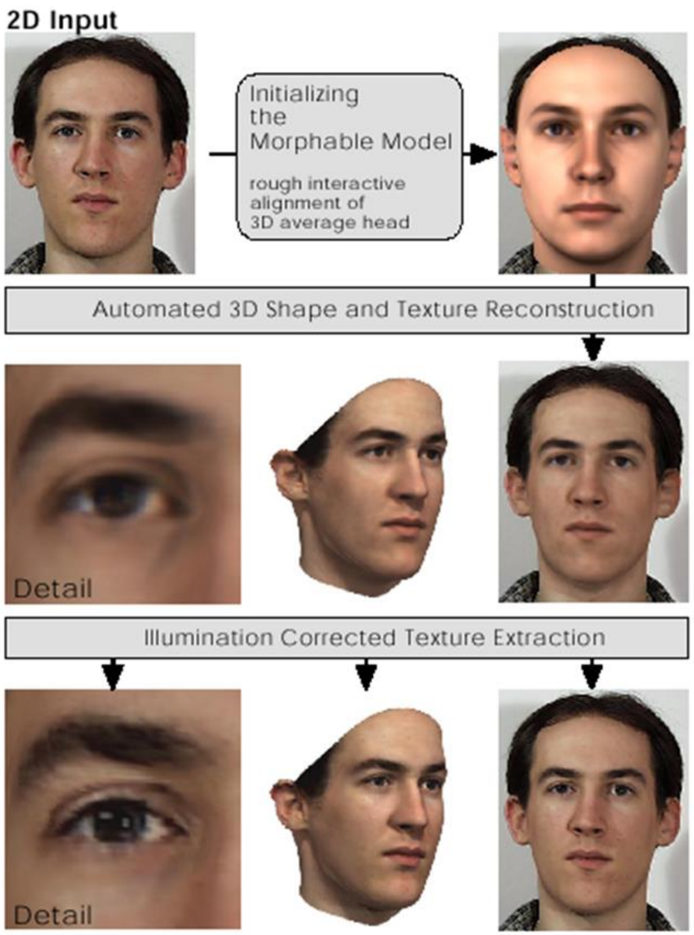

Phong illumination model

$$E_I = \sum_{x,y} \|\mathbf{I}_{input}(x,y) - \mathbf{I}_{model}(x,y)\|^2$$

$$E = \frac{1}{\sigma_N^2} E_I + \sum_{j=1}^{m-1} \frac{\alpha_j^2}{\sigma_{S,j}^2} + \sum_{j=1}^{m-1} \frac{\beta_j^2}{\sigma_{T,j}^2} + \sum_j \frac{(\rho_j - \bar{\rho}_j)^2}{\sigma_{\rho,j}^2}$$

SCAPE: Shape completion and animation of people --- joint pose and shape model

[Anguelov et al 05]

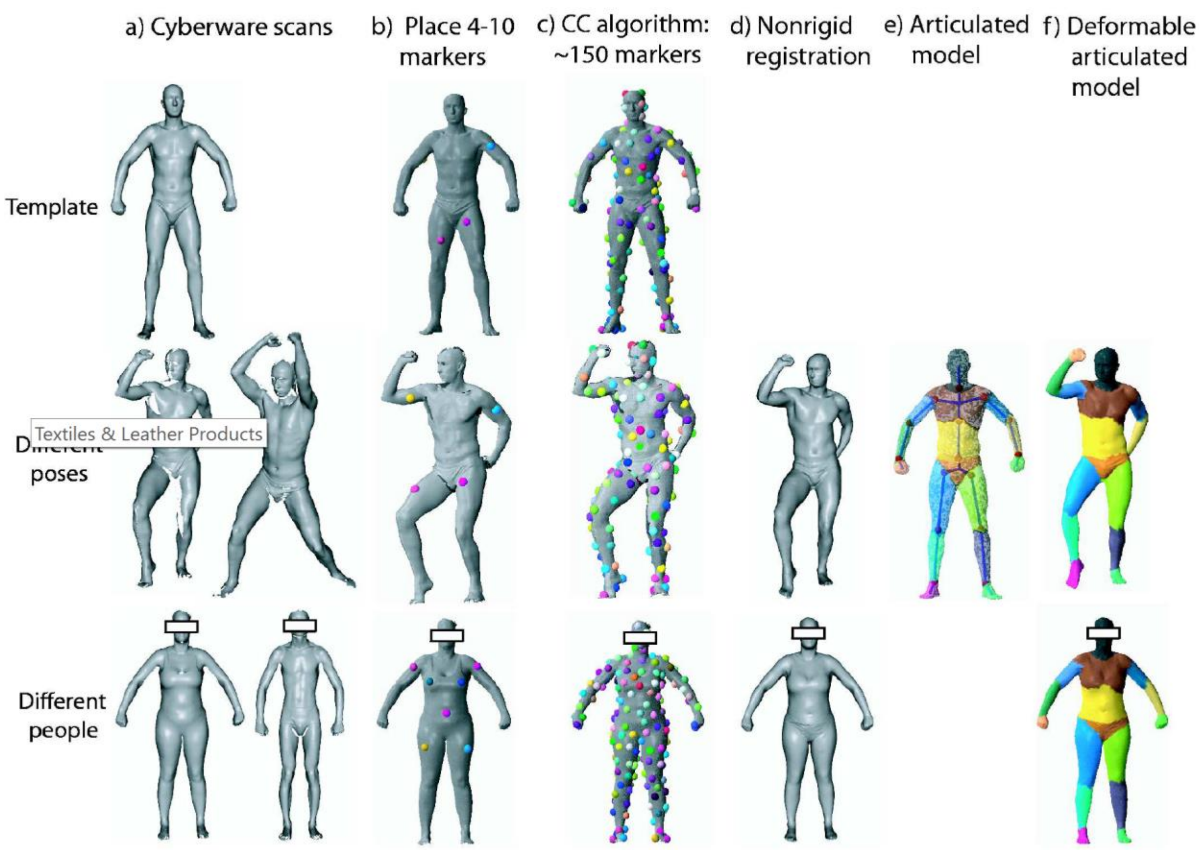

SCAPE: Shape completion and animation of people --- joint pose and shape model

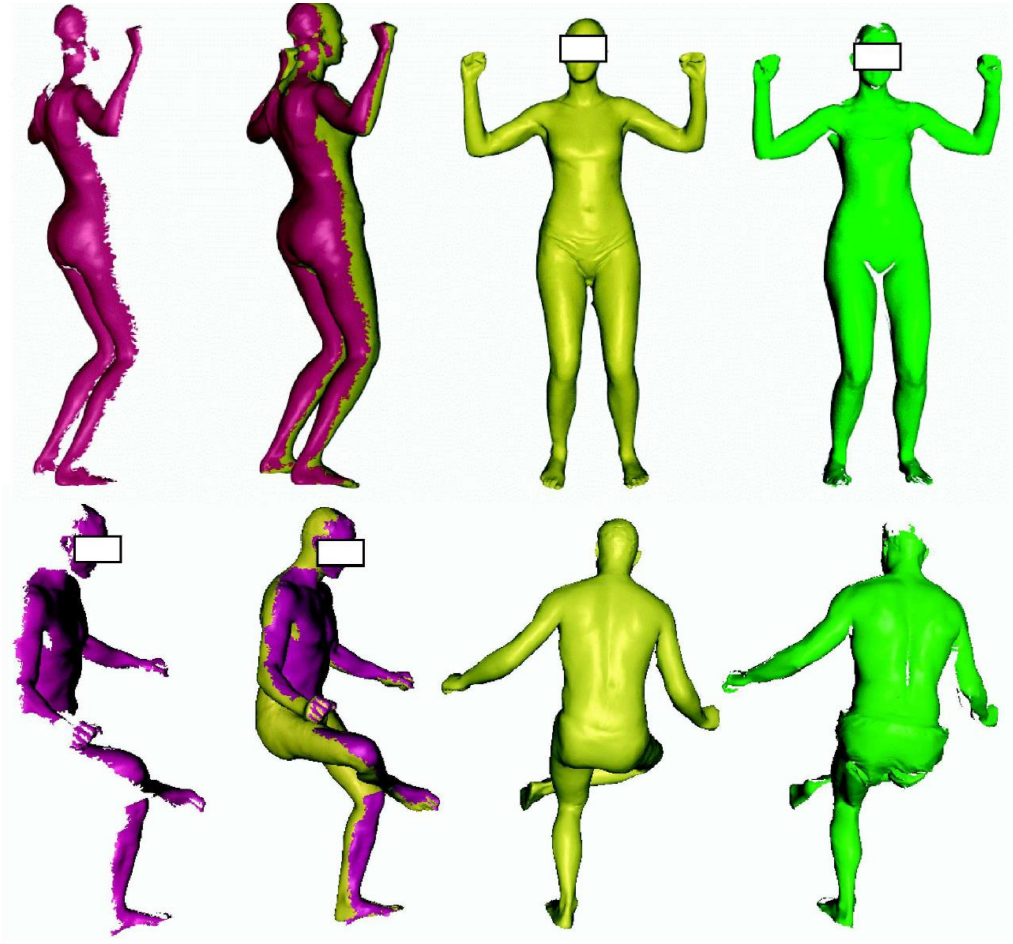

Data-Driven Shape Modeling

Modeling By Example [Funkhouser et al. 04]

Figure 6: Results of shape similarity queries where the query provided to the system is (top) the chair with the legs selected, and (bottom) the chair with the arms selected.

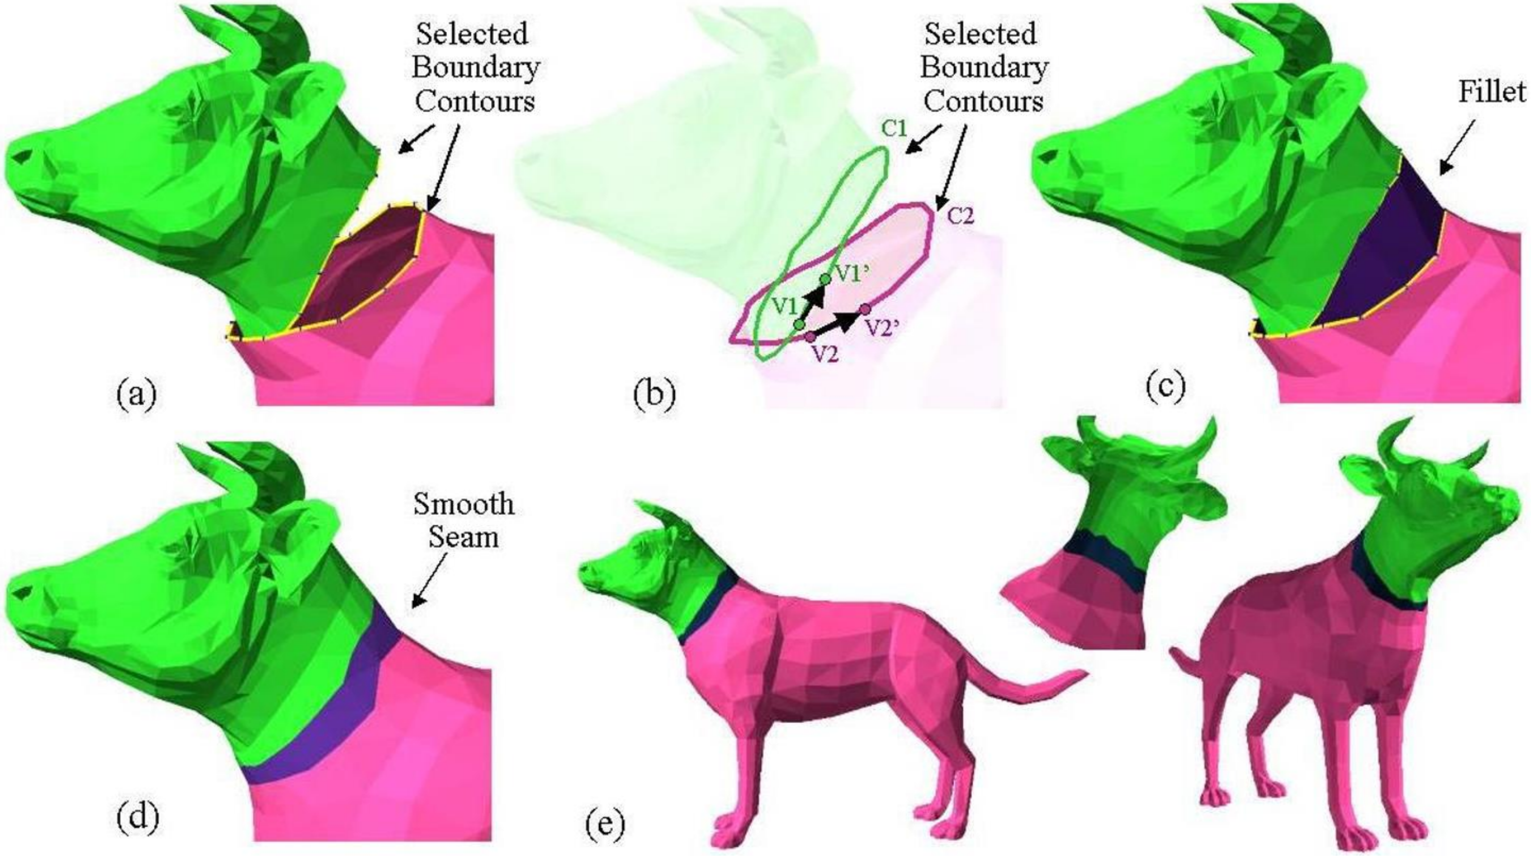

Figure 8: Attaching the head of a cow to the body of a dog: (a) a boundary contour is selected on each part ($C1$ and $C2$); (b) the pair of closest points ($V1$ and $V2$) is found and the local direction near those points is used to determine the relative orientation of the contours; (c) a fillet is constructed attaching the contours; (d) the mesh is smoothed in the region nearby the seams of the fillet. (e) the result is a smooth, watertight seam.

Data-Driven Suggestions for Creativity Support in 3D Modeling [Chaudhuri and Koltun' 11]

Basic idea

- Automatically suggest ways in which the user can extend a basic shape, to stimulate creative exploration

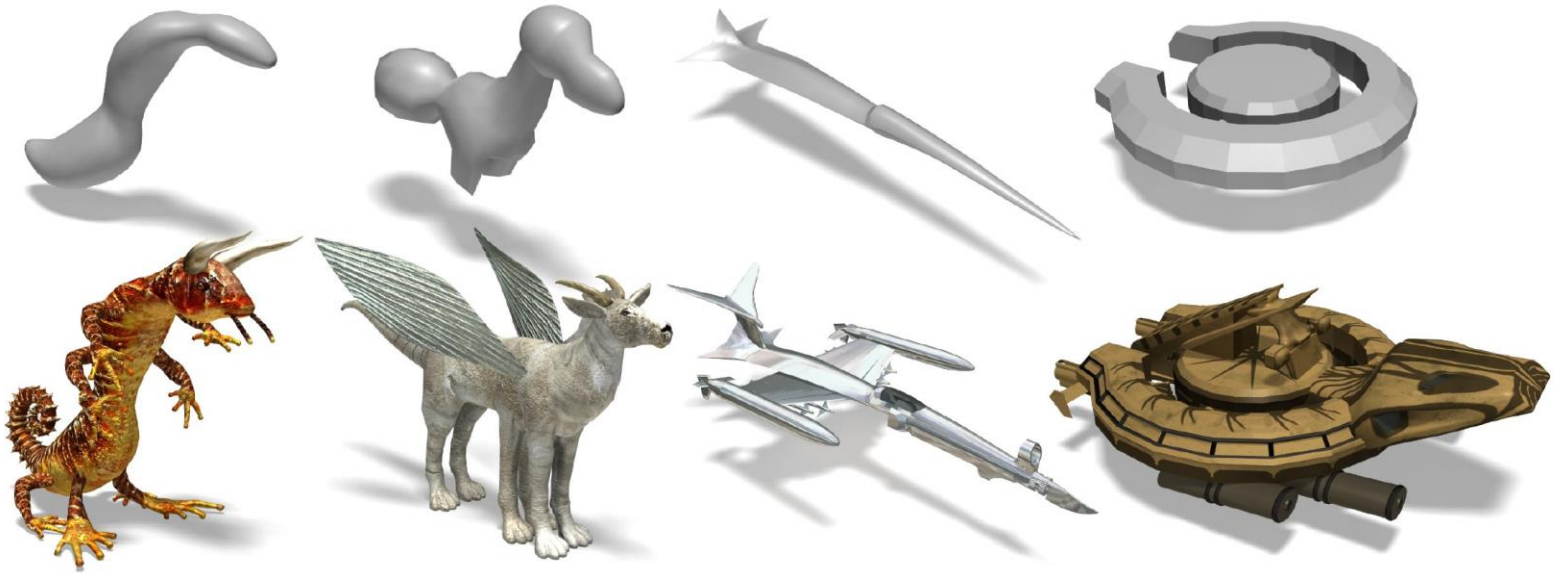

Overview

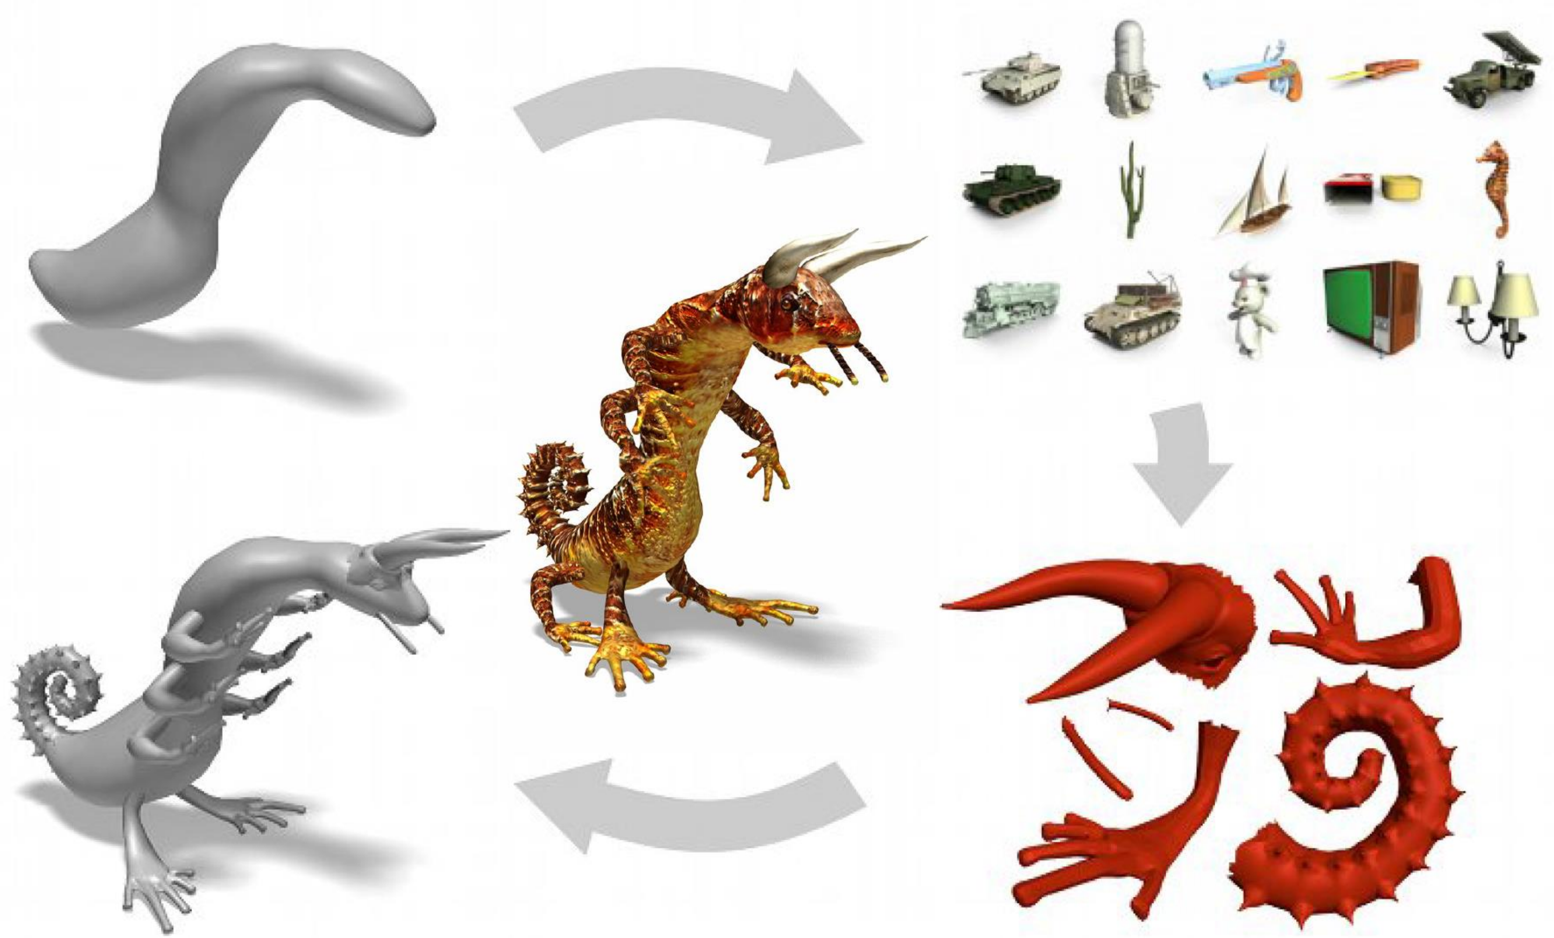

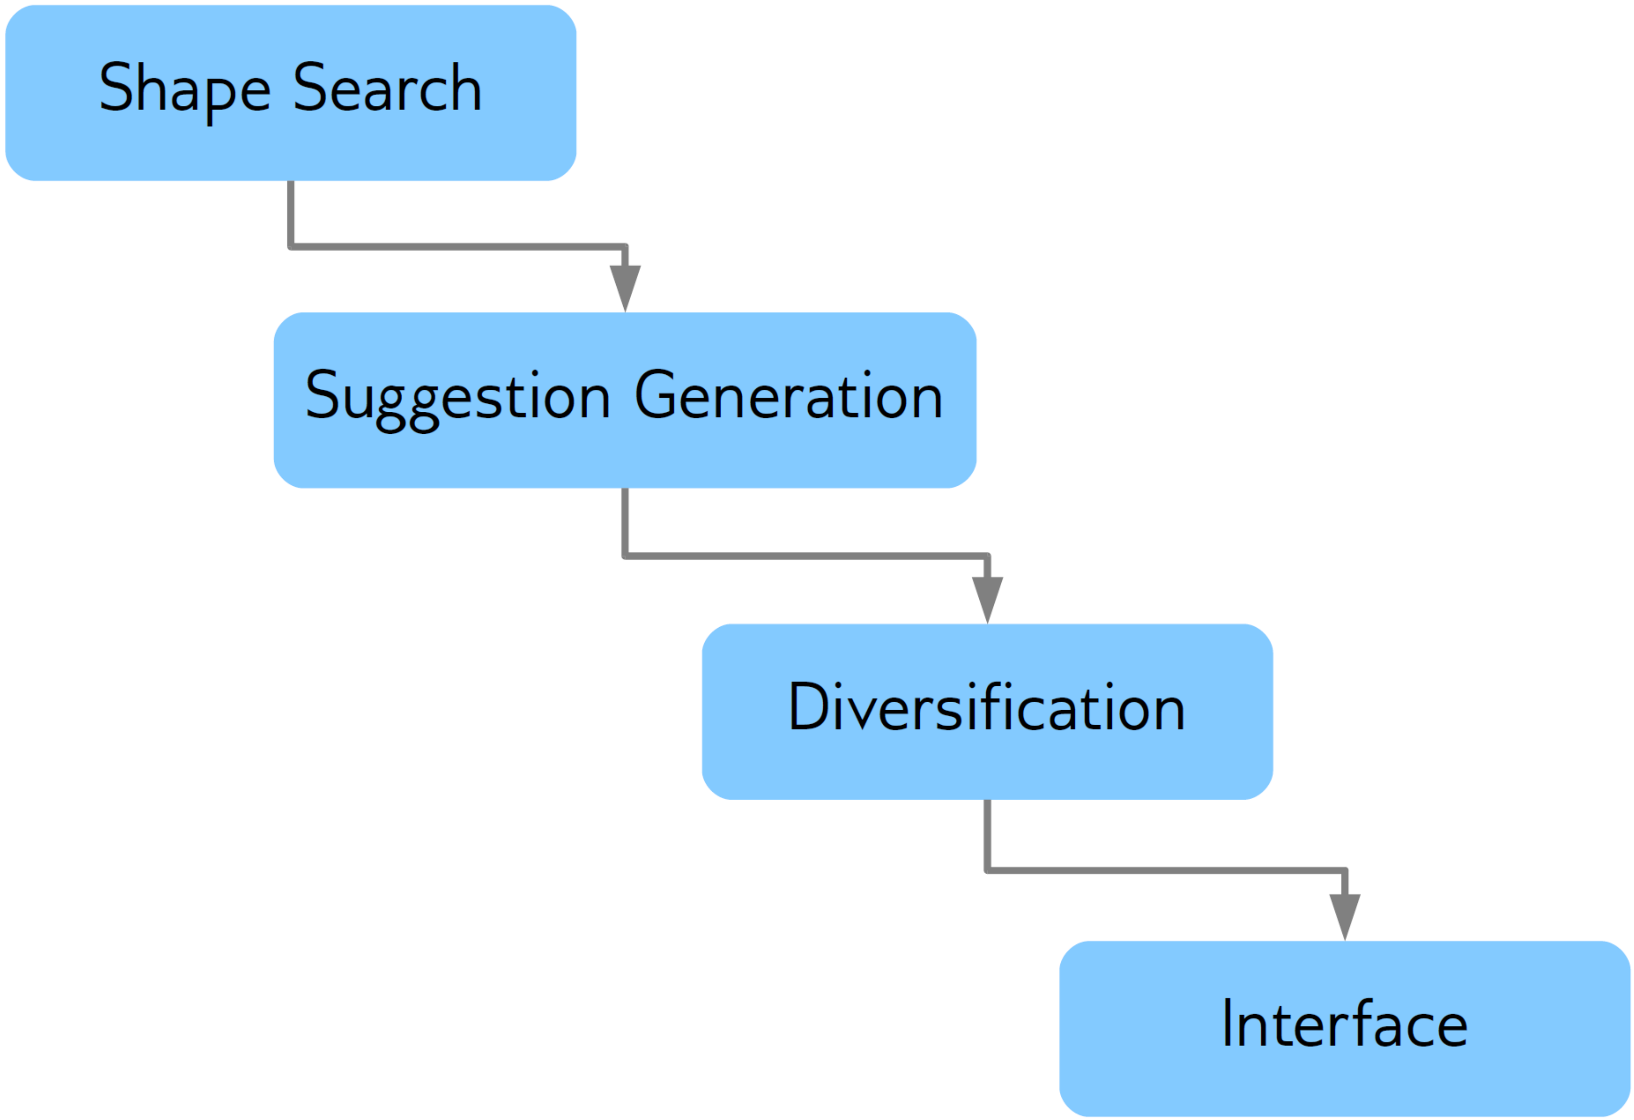
Shape Search
Suggestion Generation
Diversification
Interface

D^3 histogram

- Bin pairs of sample points on the shape
- Bins indexed by the distance between a pair of points, and the shape diameter (local thickness) of each point
- Comparison by histogram intersection and pyramid matching, for partial and approximate matches

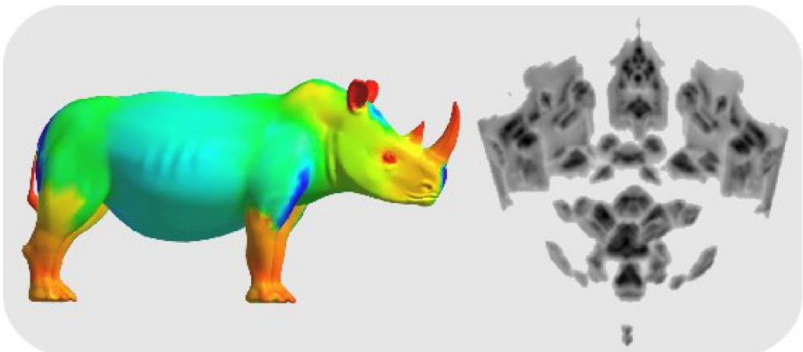

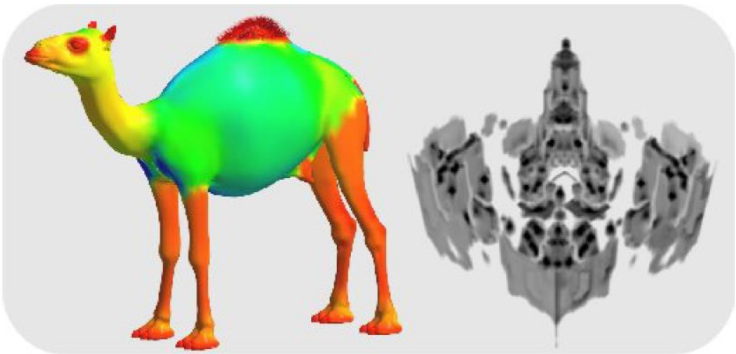

Histogram Intersection

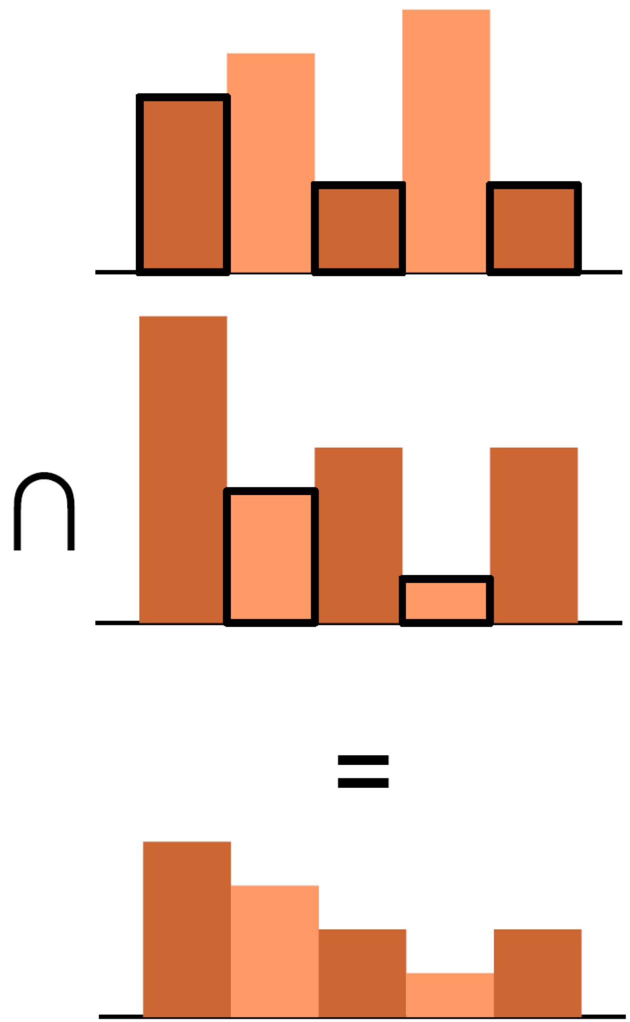

Suggestion Generation

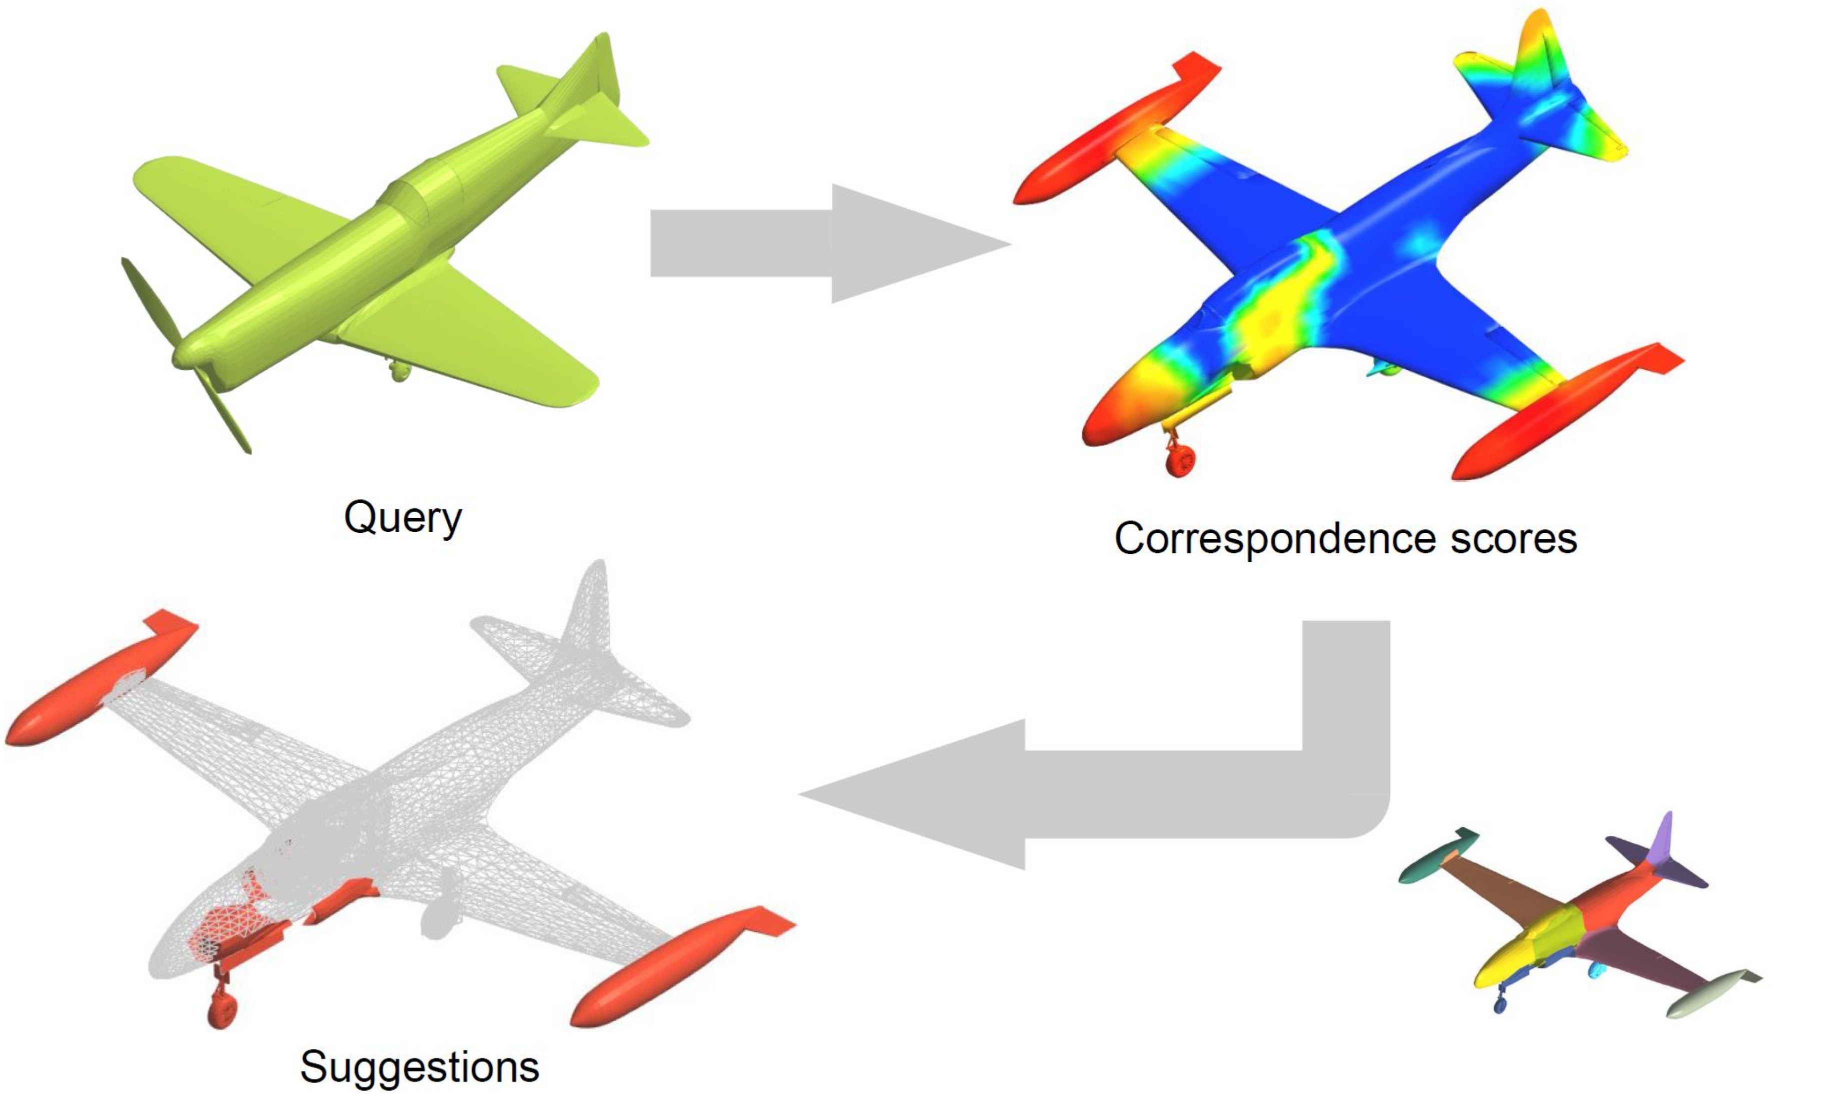

Segmentation

- Prior segmentation of database models based on shape diameter and approximate convexity
- No need for compatible segmentation of query

Diversification

- **Problem:** Large databases contain many near-identical shapes
  - If one is a good match, so are its twins
  - Most of the top-ranked options look the same
- Maximal Marginal Relevance (MMR) breaks up long runs of similar results in a ranked list *[Carbonell and Goldstein '98]*

Results

Results

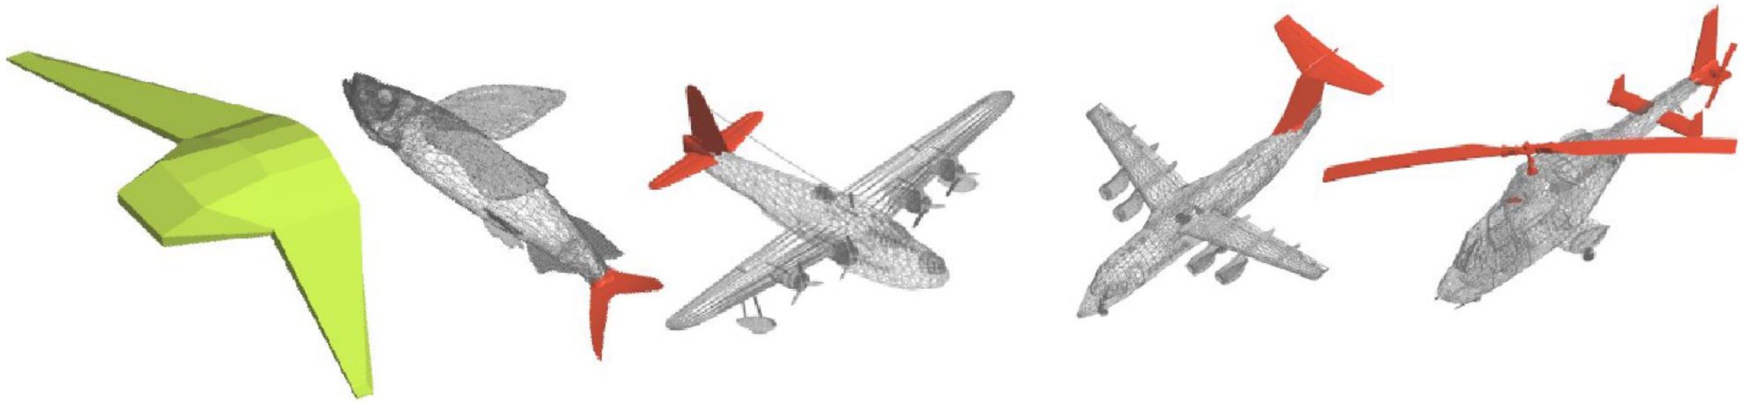

Results

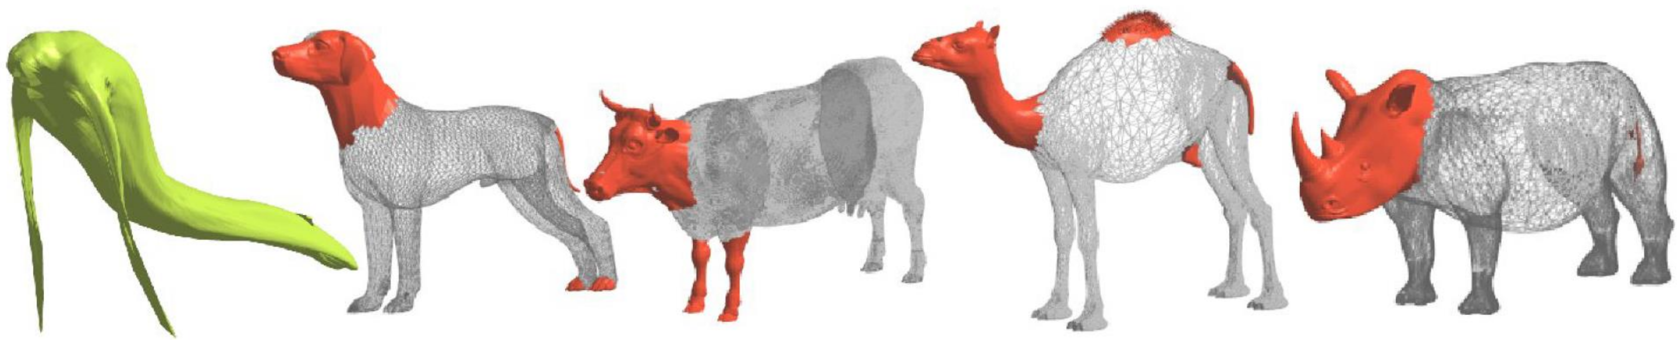

Results

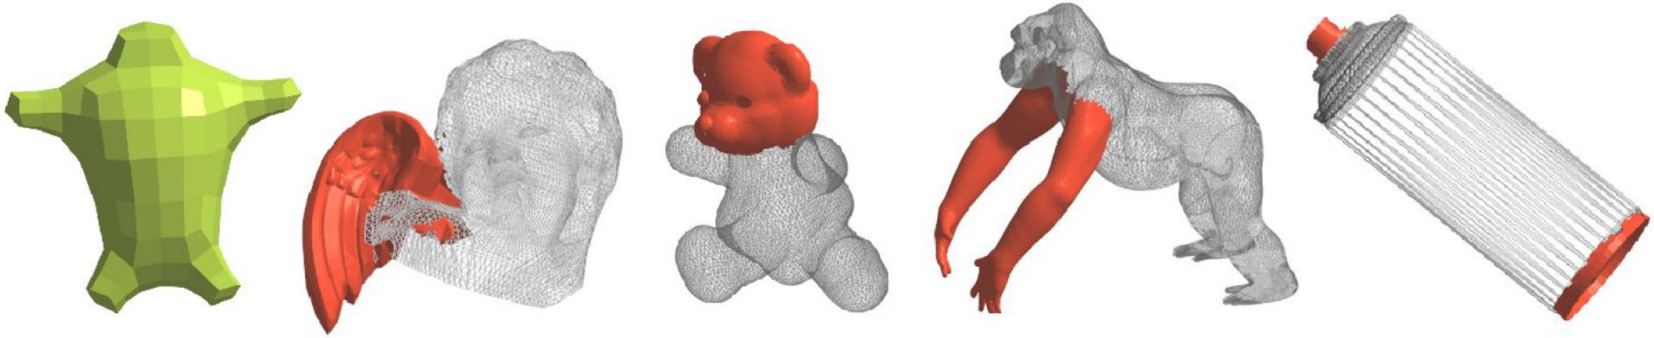

Results

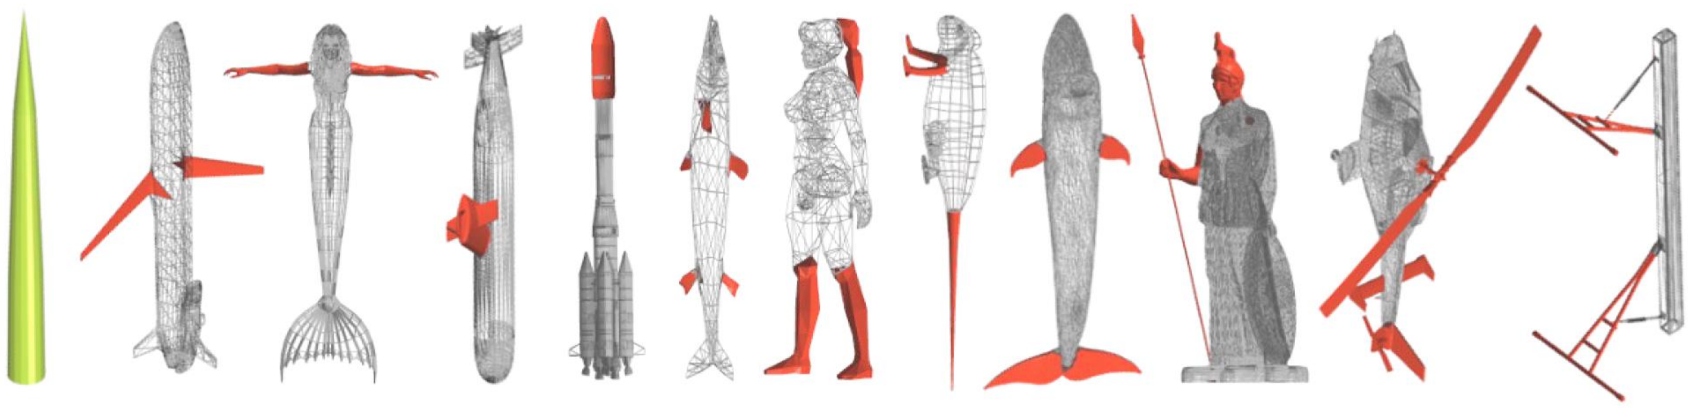

Results: Creatures

Results: Aircraft

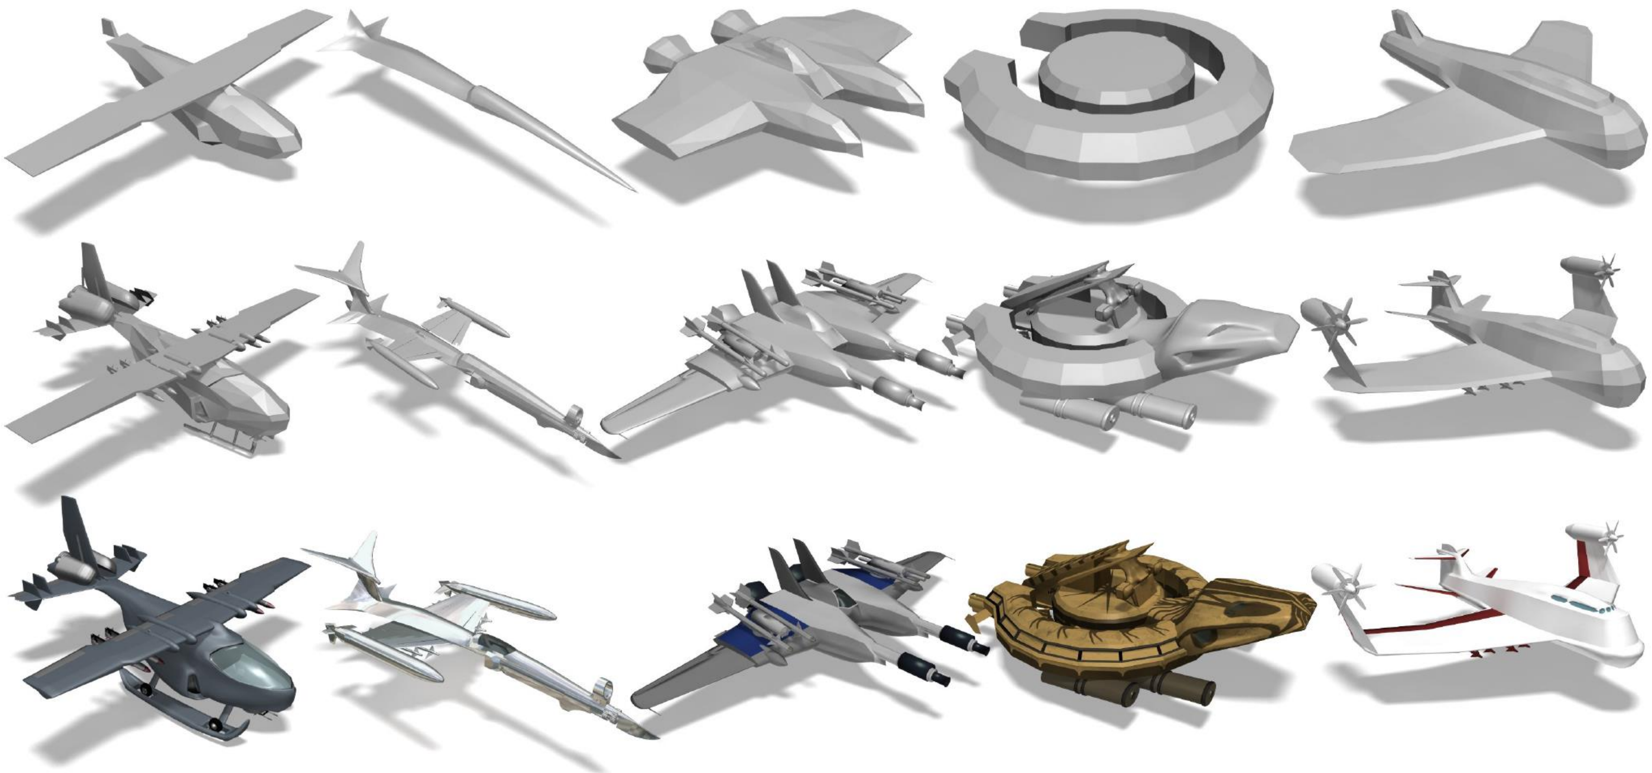

Exploratory Modeling with Collaborative Design Spaces
[Talton et al. 09]

91 dimension tree-space [Weber and Penn 95]

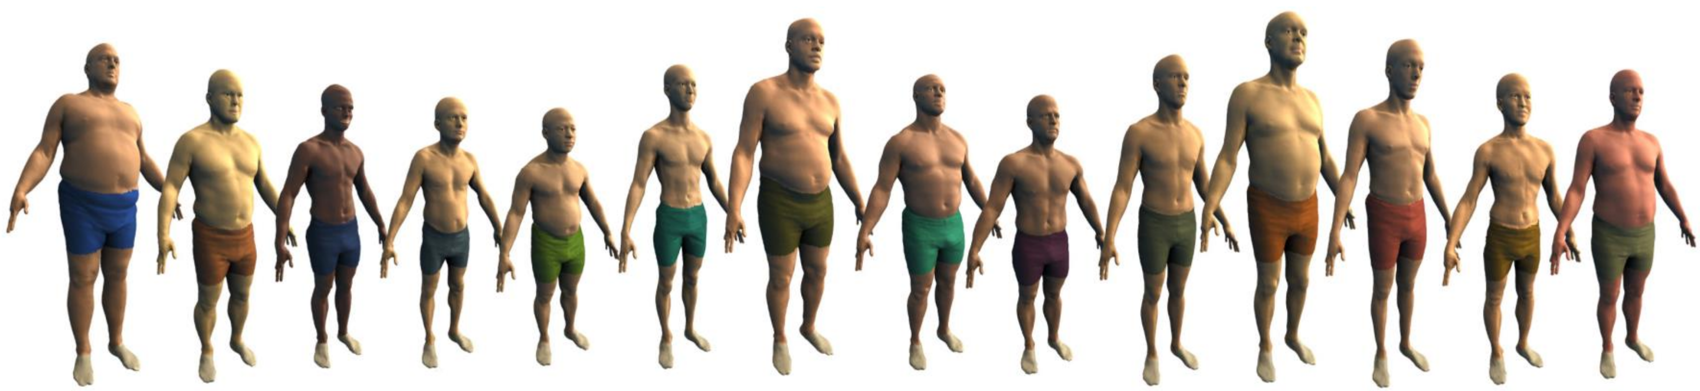

130 dimension human-space [Allen et al. 03]

Density Estimation from data

$$\hat{f}(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^{N} K_i(\mathbf{x})$$

$$K_i(\mathbf{x}) = \mathcal{G}(\mathbf{x}; \mathbf{x}_i, \mathbf{\Sigma}_i) =$$
$$\frac{1}{(2\pi)^{n/2} |\mathbf{\Sigma}_i|^{1/2}} \exp\left[-\frac{1}{2} (\mathbf{x} - \mathbf{x}_i)^T \mathbf{\Sigma}_i^{-1} (\mathbf{x} - \mathbf{x}_i)\right]$$

$$\mathbf{\Sigma}_{s,t} = \sum_{i=1}^{N} \omega_i \left[(\mathbf{x}_i)_s - (\mathbf{x})_s\right] \left[(\mathbf{x}_i)_t - (\mathbf{x})_t\right]$$

$$\omega_i = \frac{\mathcal{G}(\mathbf{x}_i; \mathbf{x}, \alpha\|\mathbf{x} - \mathbf{x}_{d(k)}\|^2 \mathbf{I})}{\sum_{j=1}^{N} \mathcal{G}(\mathbf{x}_j; \mathbf{x}, \alpha\|\mathbf{x} - \mathbf{x}_{d(k)}\|^2 \mathbf{I})}$$

Sampling

Local sampling

$$\frac{1}{\varphi}\,\mathcal{G}\big(\mathbf{x};\mathbf{x}_0,\boldsymbol{\Sigma}_0\big)\cdot\hat{f}(\mathbf{x})$$

Constrained sampling

$$\hat{f}(\mathbf{x}_1 \mid \mathbf{x}_2) = \frac{1}{N}\sum_{i=1}^{N} K_i(\mathbf{x}_1 \mid \mathbf{x}_2) = \frac{1}{N}\sum_{i=1}^{N} G(\mathbf{x}_1;\mathbf{x}_{i_{1|2}},\boldsymbol{\Sigma}_{i_{1|2}})$$

$$\begin{aligned}
\mathbf{x}_{i_{1|2}} &= \mathbf{x}_{i_1} + \boldsymbol{\Sigma}_{i_{12}}\boldsymbol{\Sigma}_{i_{22}}^{-1}(\mathbf{x}_2 - \mathbf{x}_{i_2}) \\
\boldsymbol{\Sigma}_{i_{1|2}} &= \boldsymbol{\Sigma}_{i_{11}} - \boldsymbol{\Sigma}_{i_{12}}\boldsymbol{\Sigma}_{i_{22}}^{-1}\boldsymbol{\Sigma}_{i_{21}}.
\end{aligned}$$

Figure 5: *(Left) Typical points sampled from the computed density functions of trees (top) and humans (bottom). (Right) Typical points chosen uniformly at random from these parametric spaces.*

Exploratory Modeling with Collaborative Design Spaces

Jerry O. Talton   Daniel Gibson   Lingfeng Yang
Pat Hanrahan   Vladlen Koltun

Stanford University

Probabilistic model for presenting relevant components

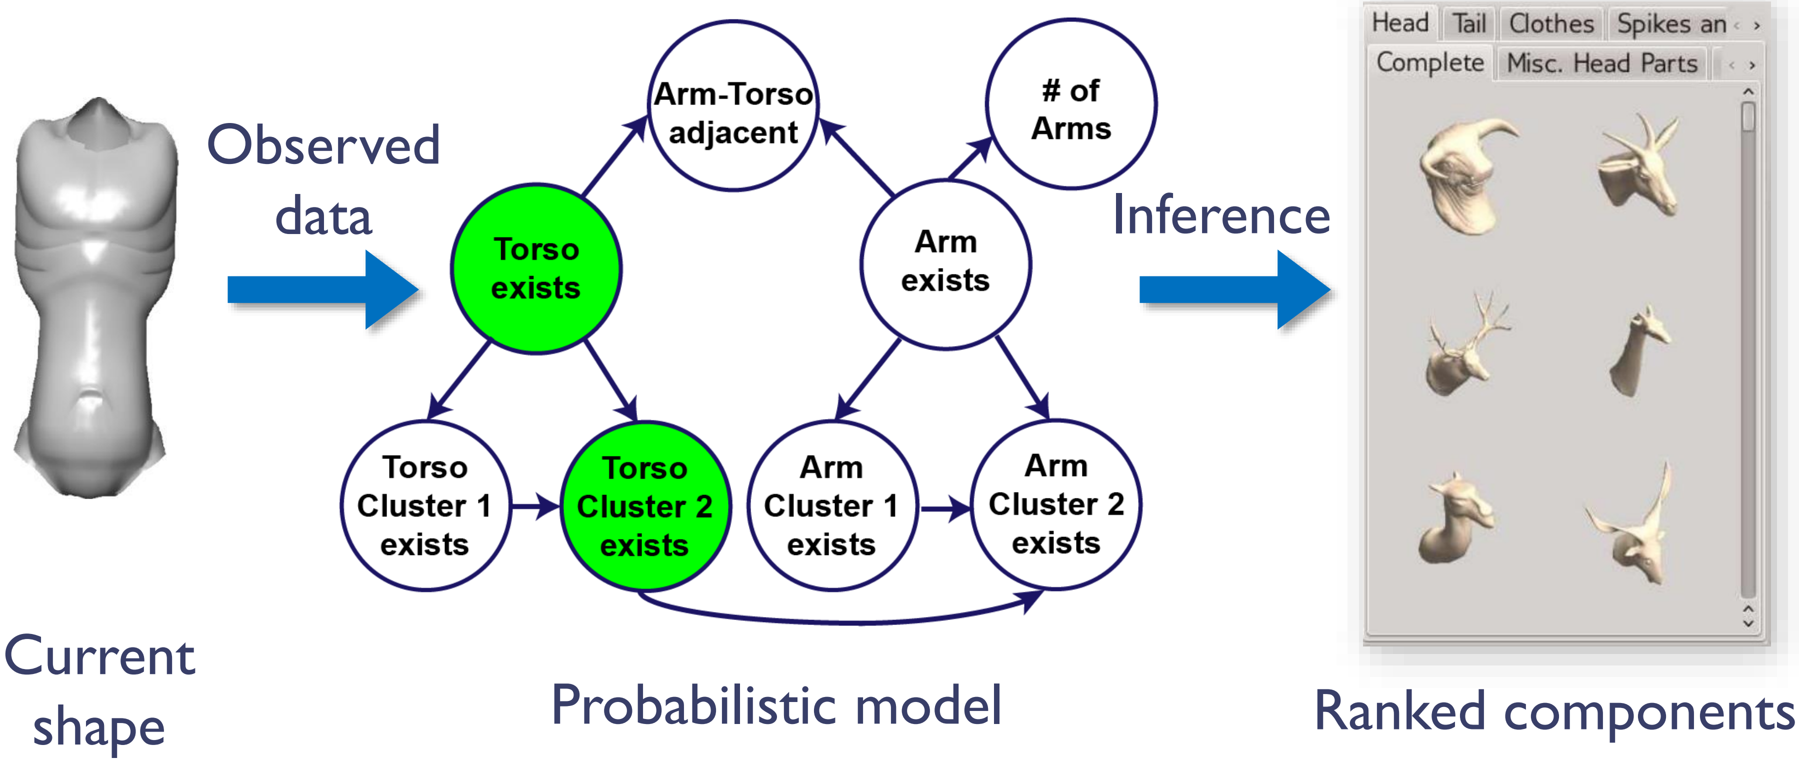

The model is learned from an input shape repository

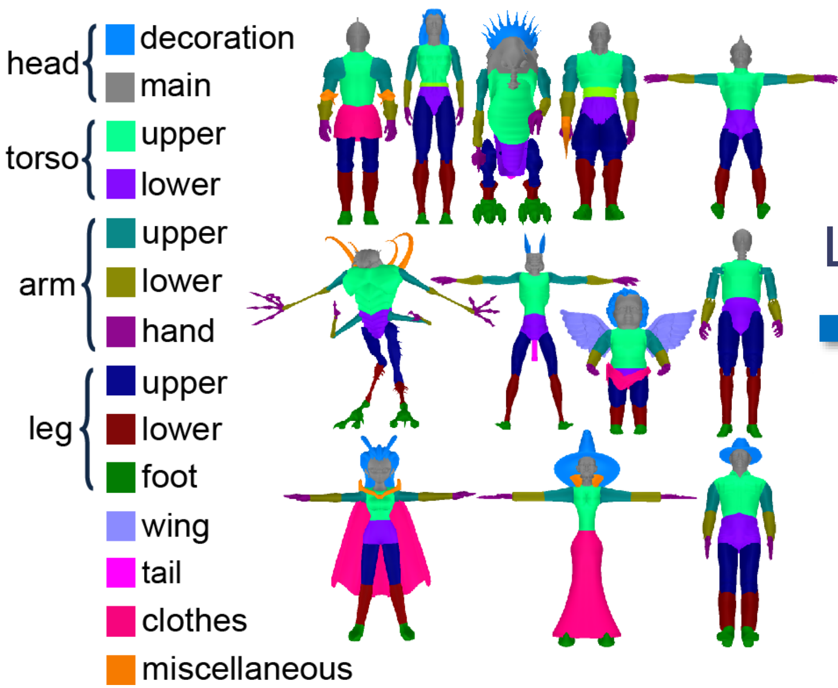

Input repository

Learning

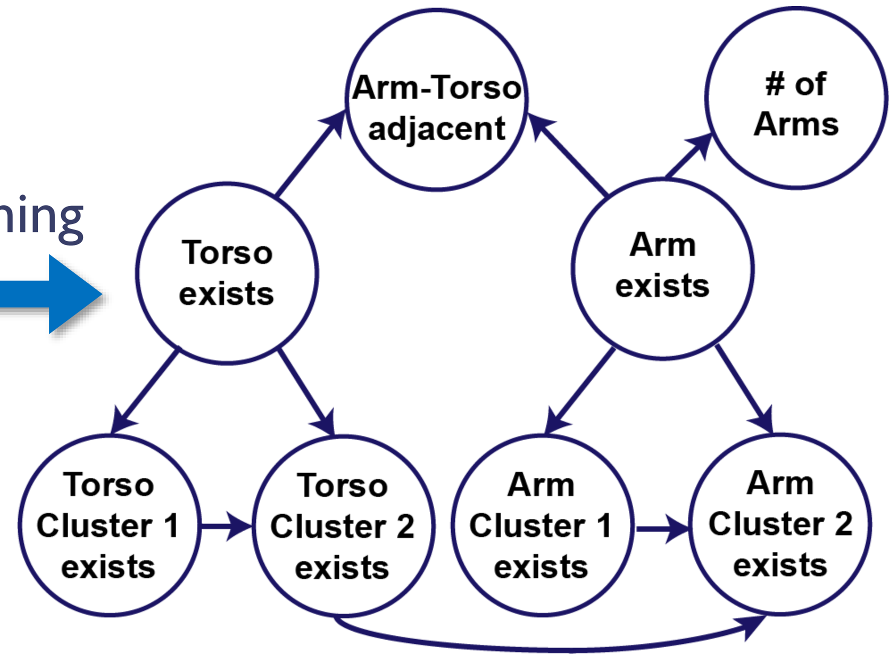

Probabilistic model

Formulation

The probabilistic model: a Bayesian Network

Shape attributes → Random variables $\boldsymbol{X} = \{x_i\}$

Dependencies between attributes → $P(\boldsymbol{X}) = \prod_i P\left(x_i \mid parents(x_i)\right)$

Represent with DAG

$$P(\boldsymbol{X}) = P(x_1)P(x_2 \mid x_1)P(x_3 \mid x_1)$$

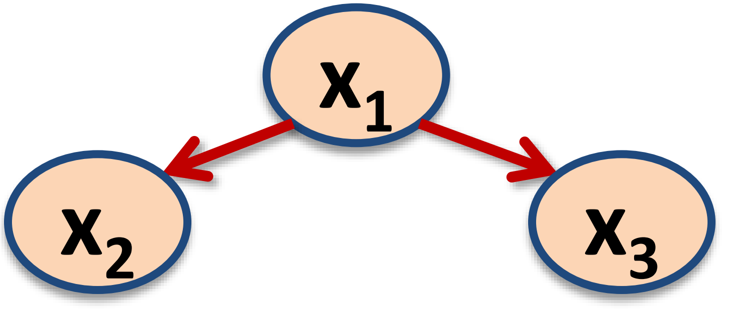

Random variables E_l

Existence of component from category l

Arm(s) exist

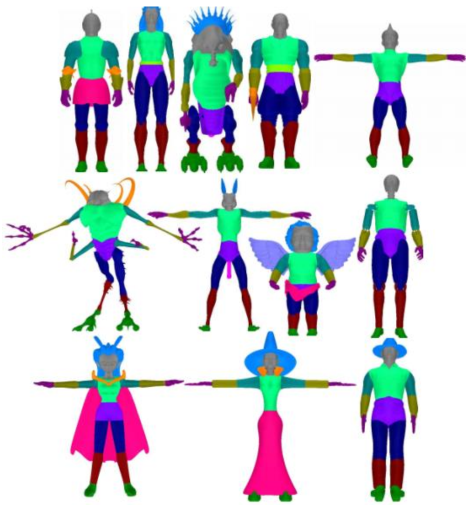

Torso(s) exist

Random variables N_l

Number of components from category l

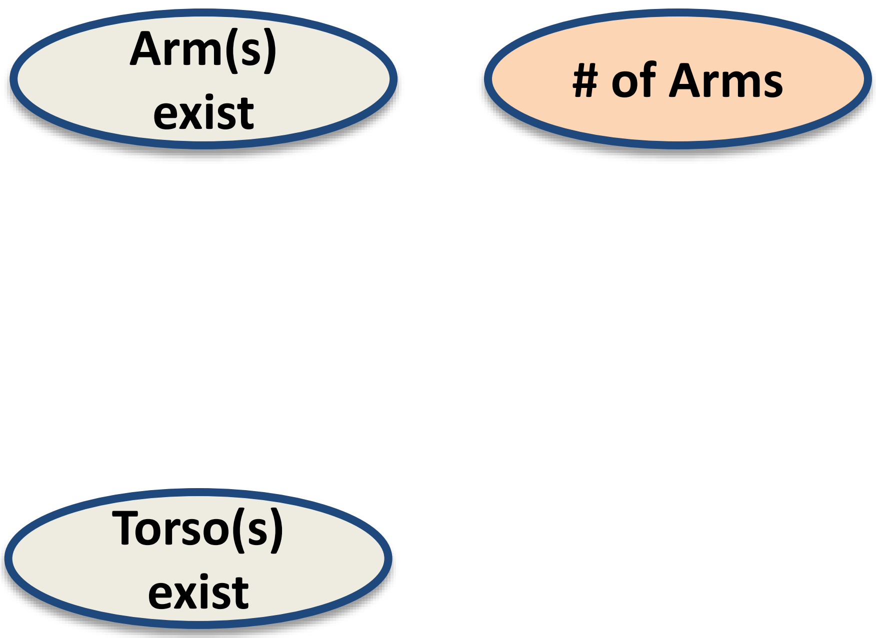

Random variables $A_{l,l'}$

Adjacency between components from categories l and l'

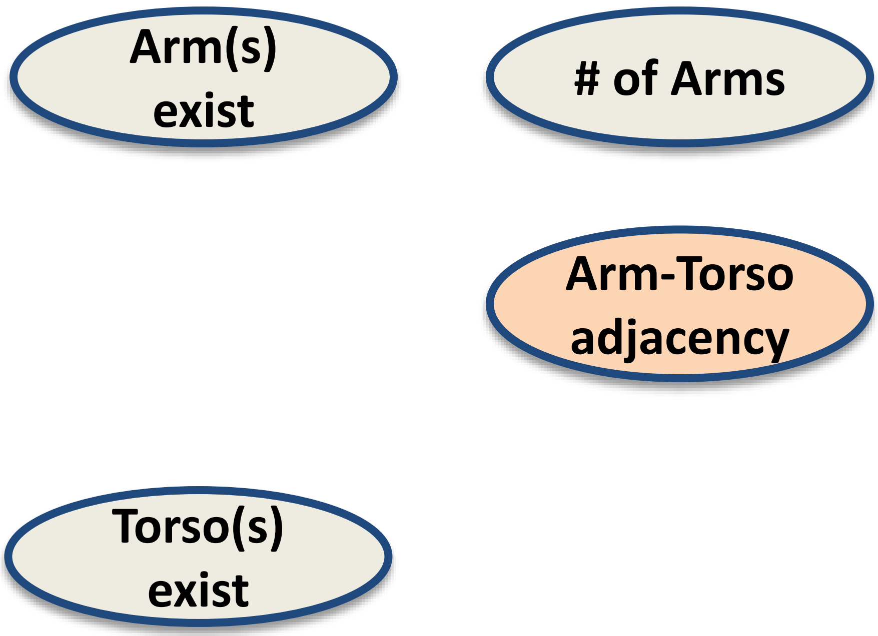

Random variables $R_{l,l'}$

Symmetry relation between components from categories l and l'

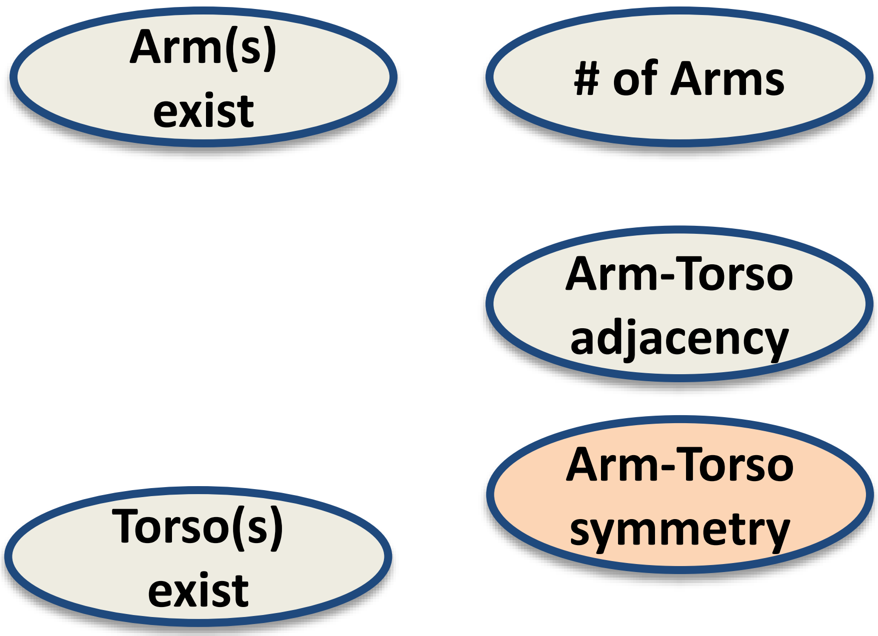

Random variables $S_{s,l}$

Existence of component from style cluster s of category l

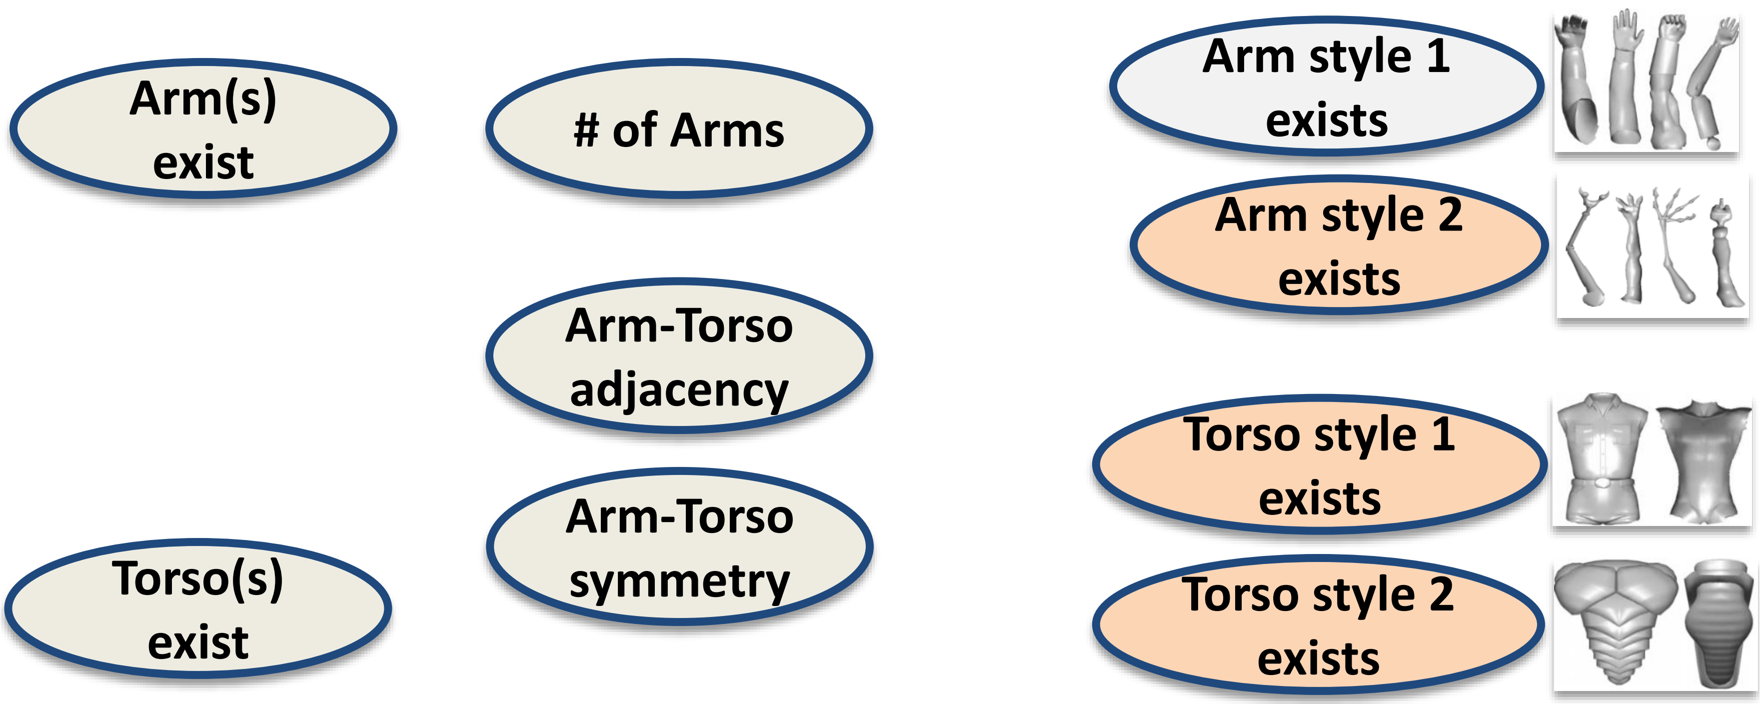

Dependencies between random variables

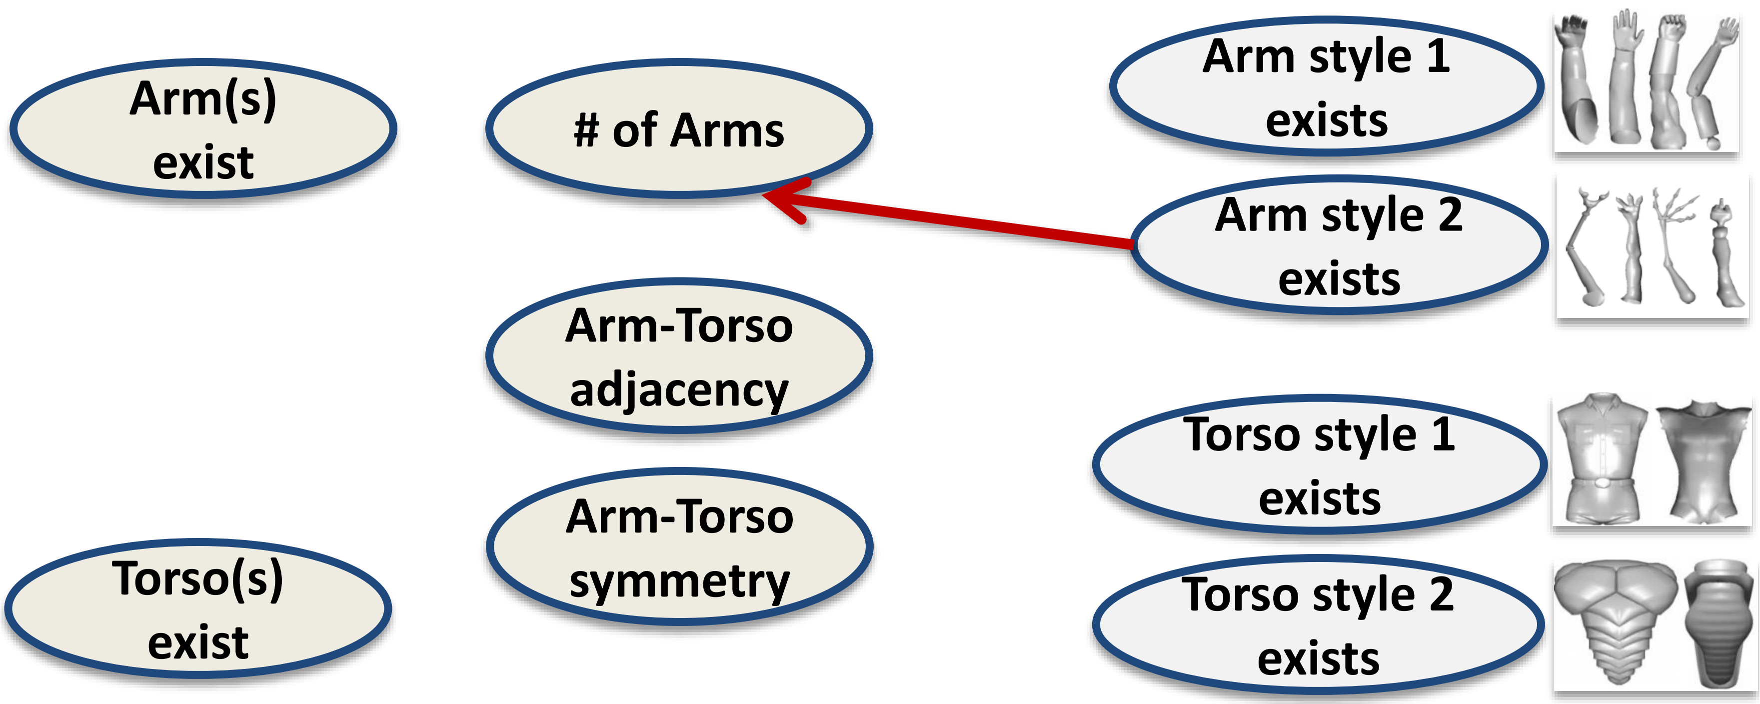

Conditional probability tables

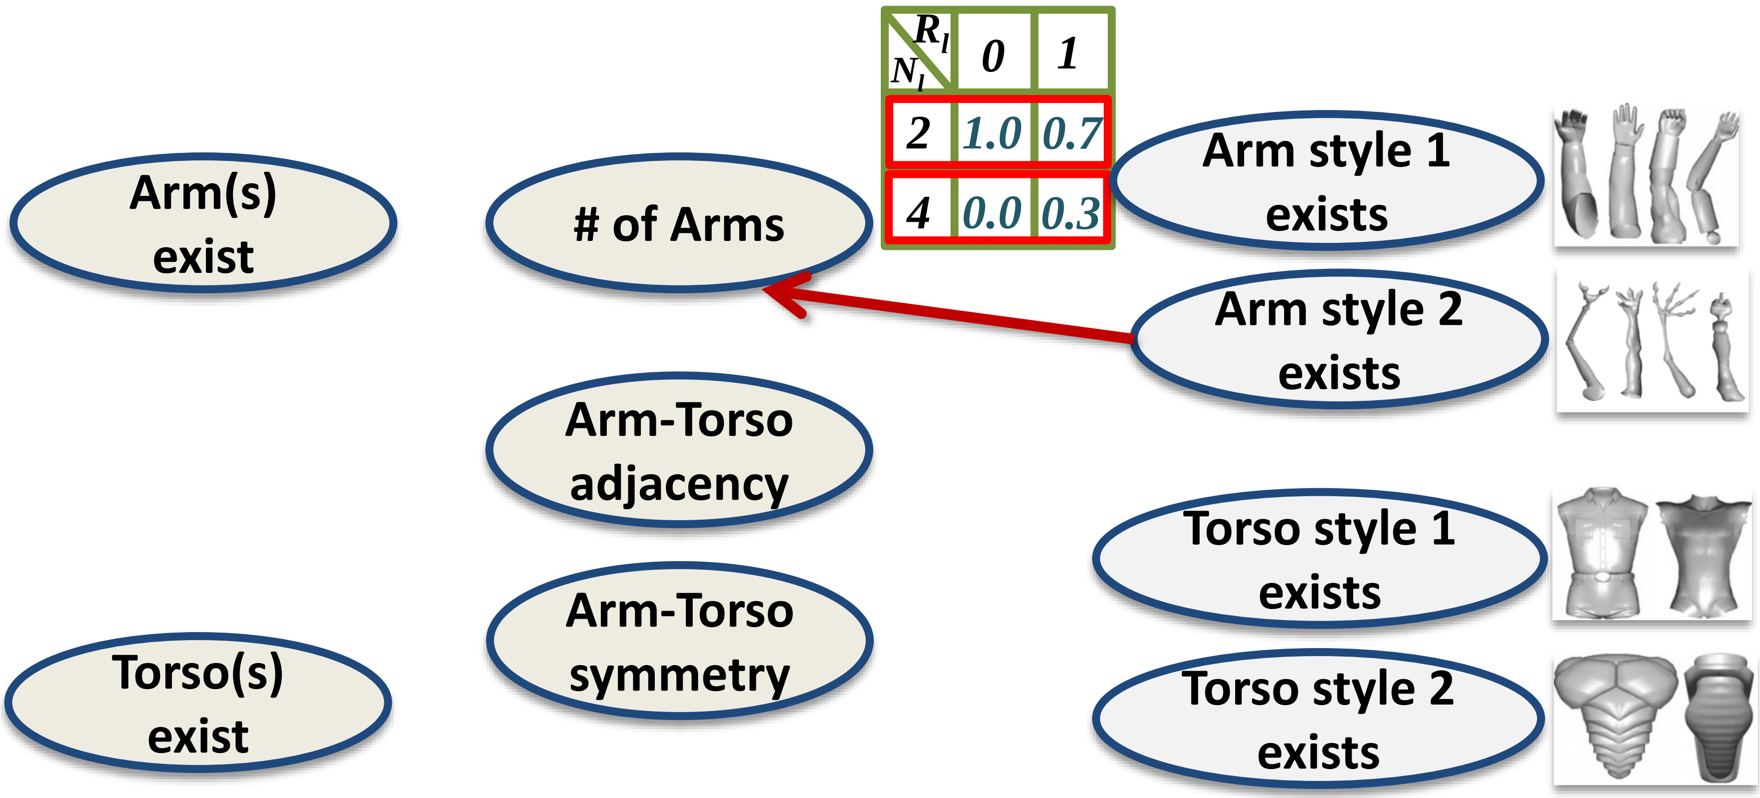

N_l \ R_l	0	1
2	1.0	0.7
4	0.0	0.3

Dependencies between random variables

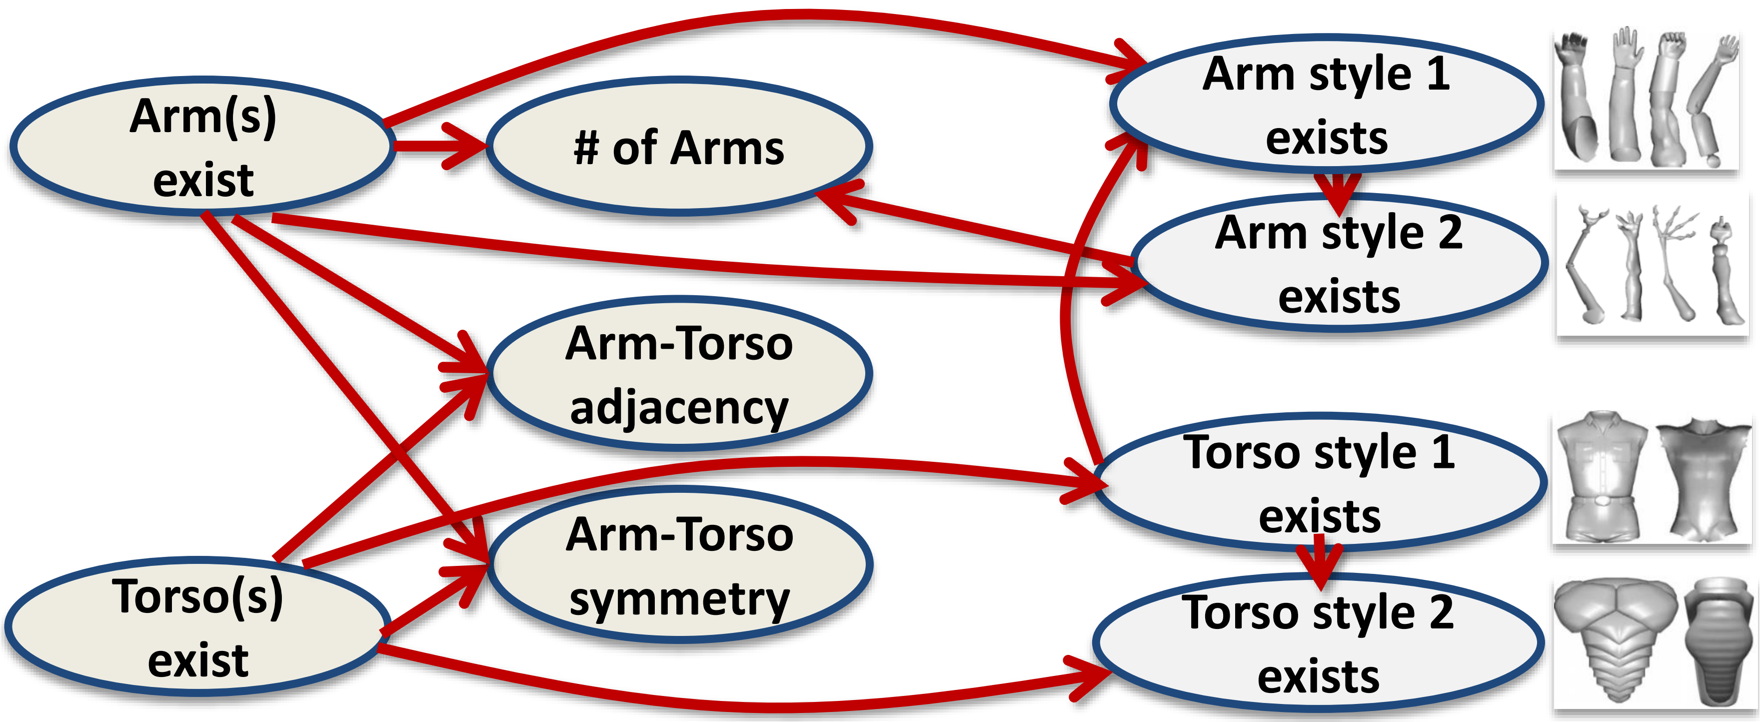

Learning

Learning the CPTs and the graph structure

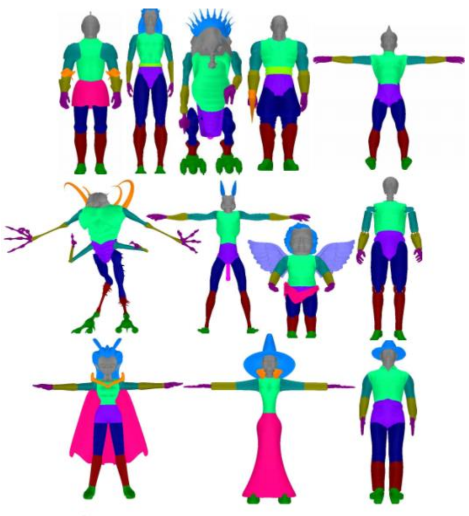

[Kalogerakis et al. 2010]
(modified)

Learning the CPTs and the graph structure

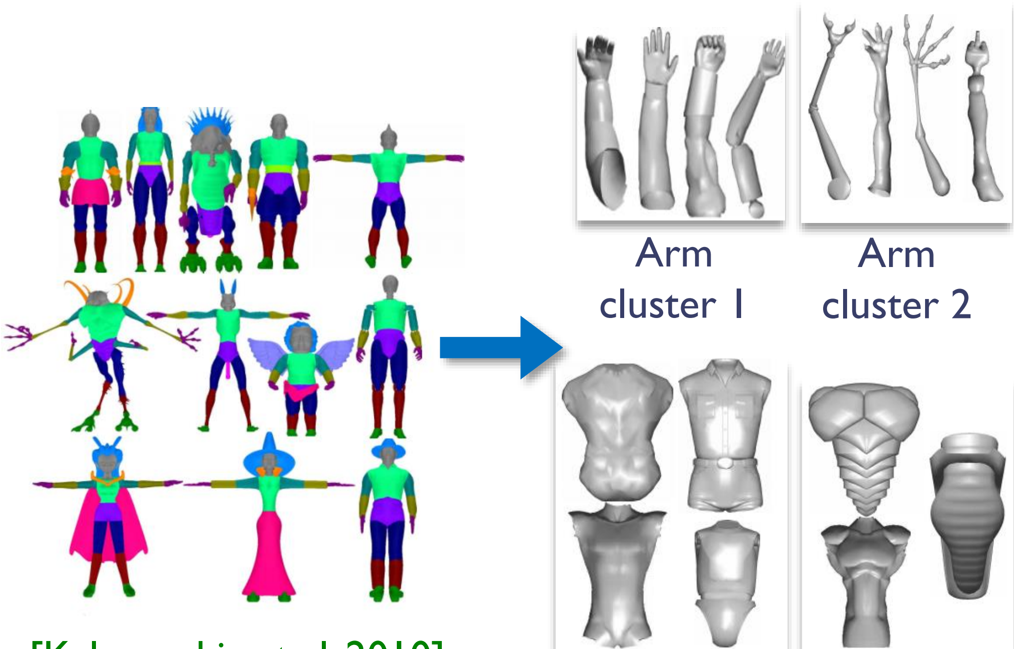

[Kalogerakis et al. 2010]
(modified)

Learning the CPTs and the graph structure

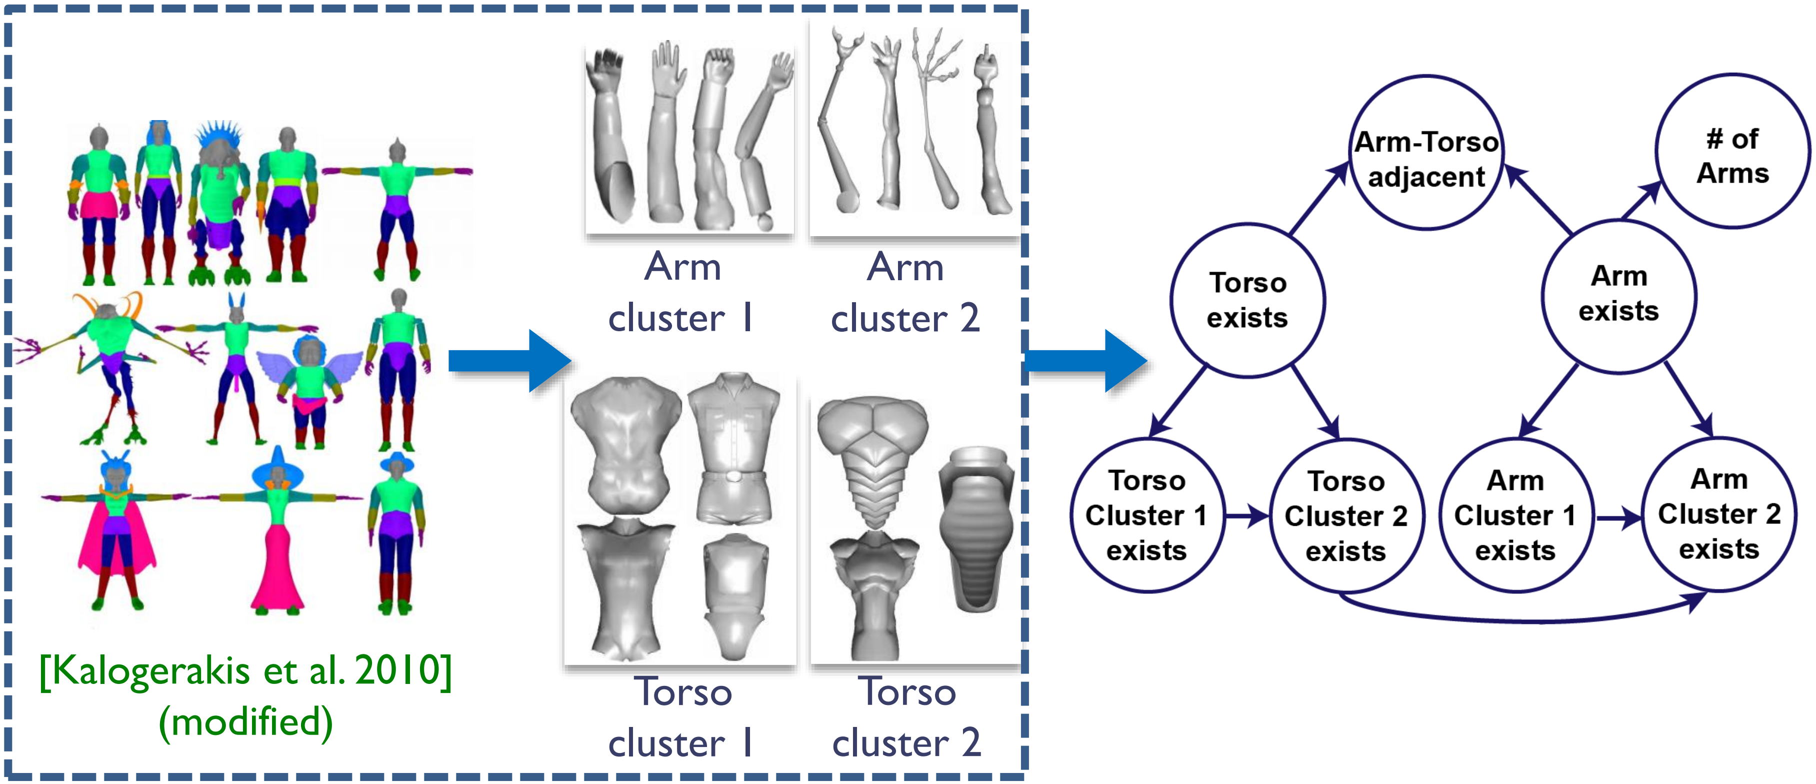

Structure and parameter learning

Maximize Bayesian Information Criterion

$$BIC = \log P(D \mid G, \boldsymbol{\theta}) - \frac{1}{2} v \log n$$

Structure and parameter learning

Maximize Bayesian Information Criterion

$$BIC = \boxed{\log P(D \mid G, \boldsymbol{\theta})} - \frac{1}{2} v \log n$$

Likelihood term

D: training data

G: graph structure

θ: CPT entries

Structure and parameter learning

Maximize Bayesian Information Criterion

$$BIC = \log P(D \mid G, \boldsymbol{\theta}) - \boxed{\frac{1}{2} v \log n}$$

Penalize model complexity

v: # of independent CPT entries

n: # of training shapes

Optimized using local search heuristics (adding, removing and flipping edges)

Inference

Inference

Inference

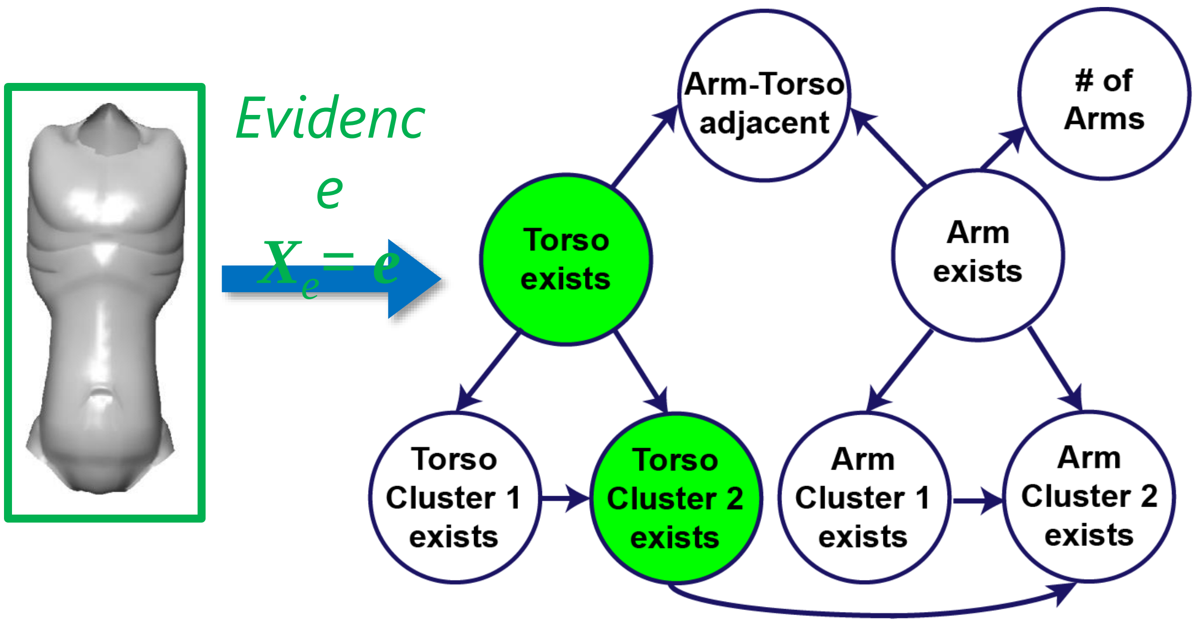
Evidence
$X_e = e$
Arm-Torso adjacent
# of Arms
Torso exists
Arm exists
Torso Cluster 1 exists
Torso Cluster 2 exists
Arm Cluster 1 exists
Arm Cluster 2 exists

Inference

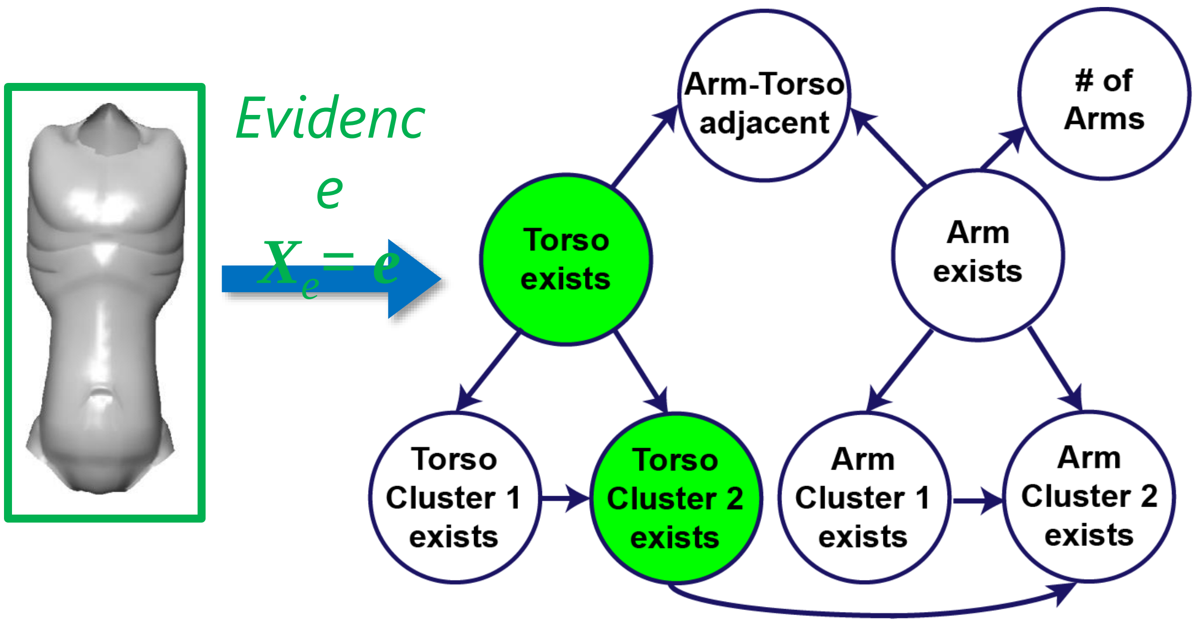

P( $| \, X_e = e$) ?

Inference

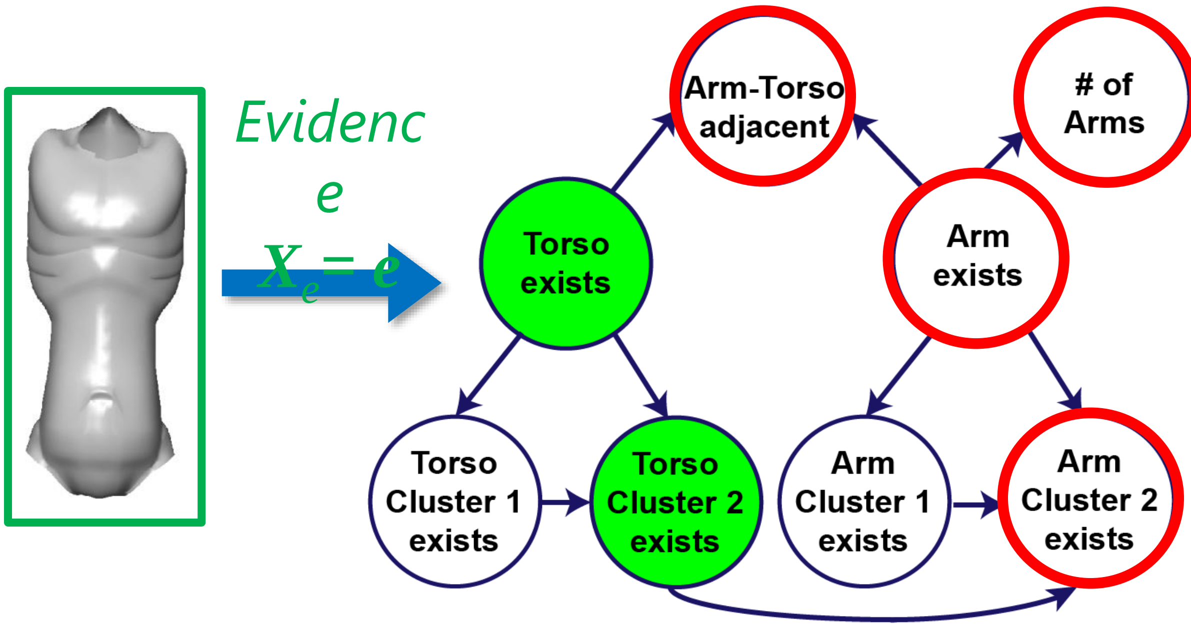

$$P(\ X_q = q \mid X_e = e\ )$$

Particle-based inference

Examples of shapes created by users