Co-Segmentation of a Set of Shapes

Ligang Liu
Graphics&Geometric Computing Lab
USTC
http://staff.ustc.edu.cn/~lgliu

Analyzing a **set** of shapes (Co-analysis)

Co-Segmentation

- More knowledge can be inferred from multiple shapes rather than an individual shape

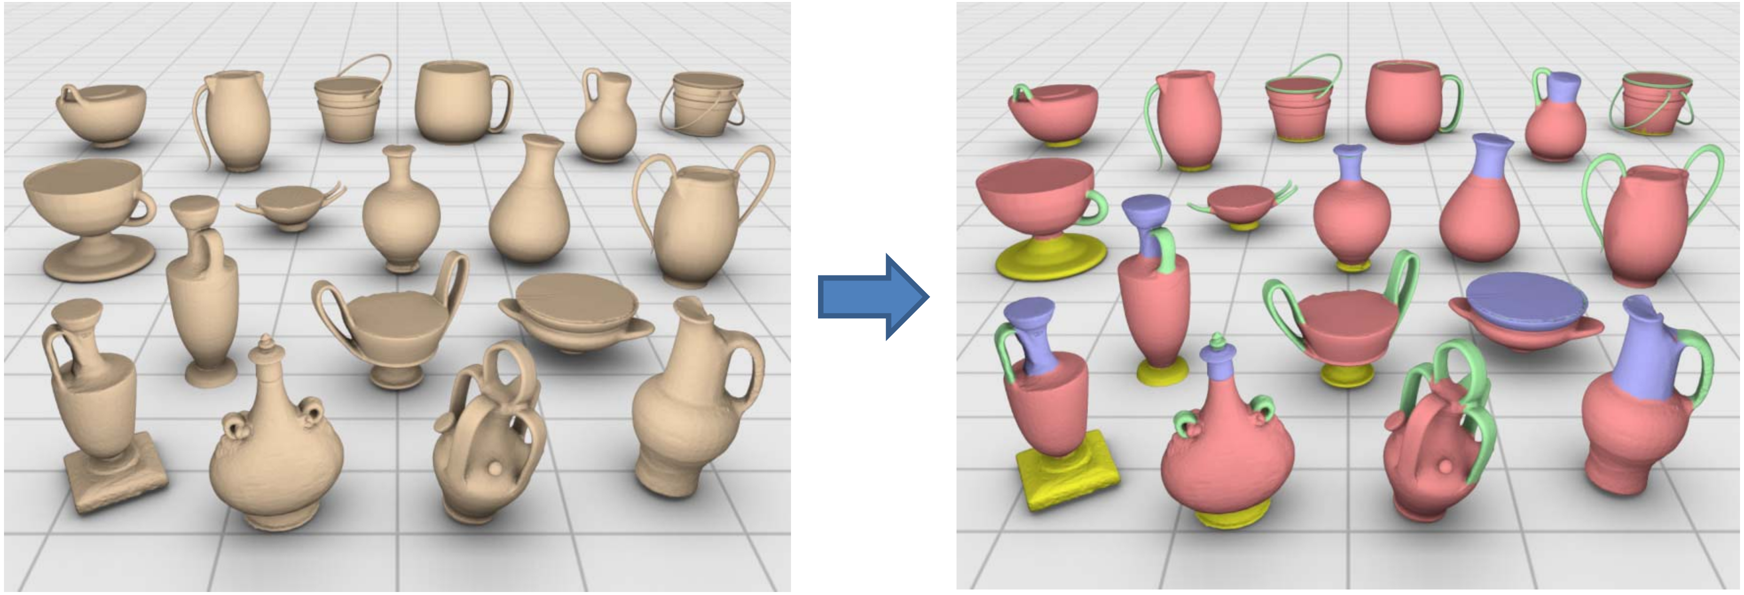

Papers

- Supervised approaches
  - Learning 3D Mesh Segmentation and Labeling [Siggraph 2010]
  - Prior Knowledge for Part Correspondence [Eurographics 11]
- Unsupervised approaches
  - Consistent Segmentation of 3D Models [SMI 09]
  - Style-Content Separation by Anisotropic Part Scales [Siggraph Asia 2010]
  - Joint Shape Segmentation with Linear Programming [Siggraph Asia 2011]
  - Unsupervised Co-Segmentation of a Set of Shapes via Descriptor-Space Spectral Clustering [Siggraph Asia 2011]

Consistent Segmentation of 3D Models

Aleksey Golovinskiy and Thomas Funkhouser
Princedon University
SMI 2009

Consistent Segmentation

- The first paper to segment a set of models consistently
- Simultaneously segments models and creates correspondences between segments

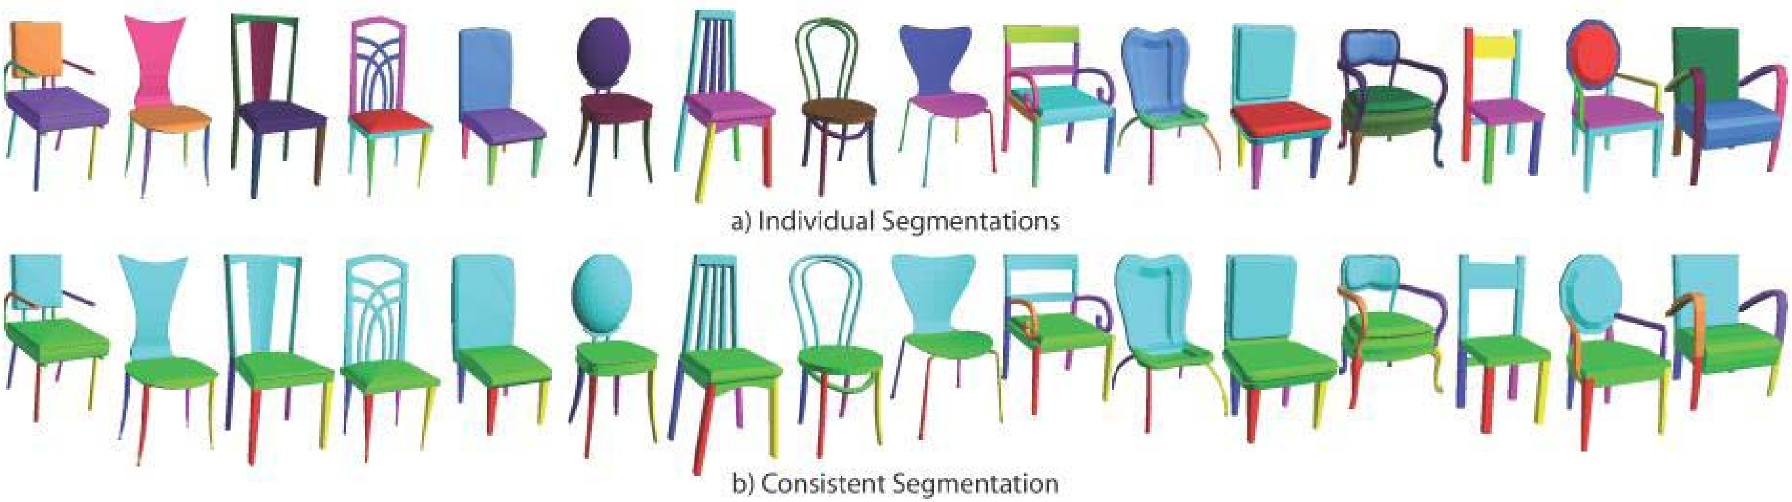

Algorithm

- Two main steps
  - create a graph that contains as nodes the faces of the models
    - Adjacency Edges
    - Correspondence Edges
  - segmentation by clustering its nodes into disjoint sets
    - weakly connected to the rest of the mesh
    - more inter-mesh correspondences to the other parts in the cluster than to parts outside of the cluster

Graph Construction

- Adjacency Edges
- Correspondence Edges
  - Alignment - PCA

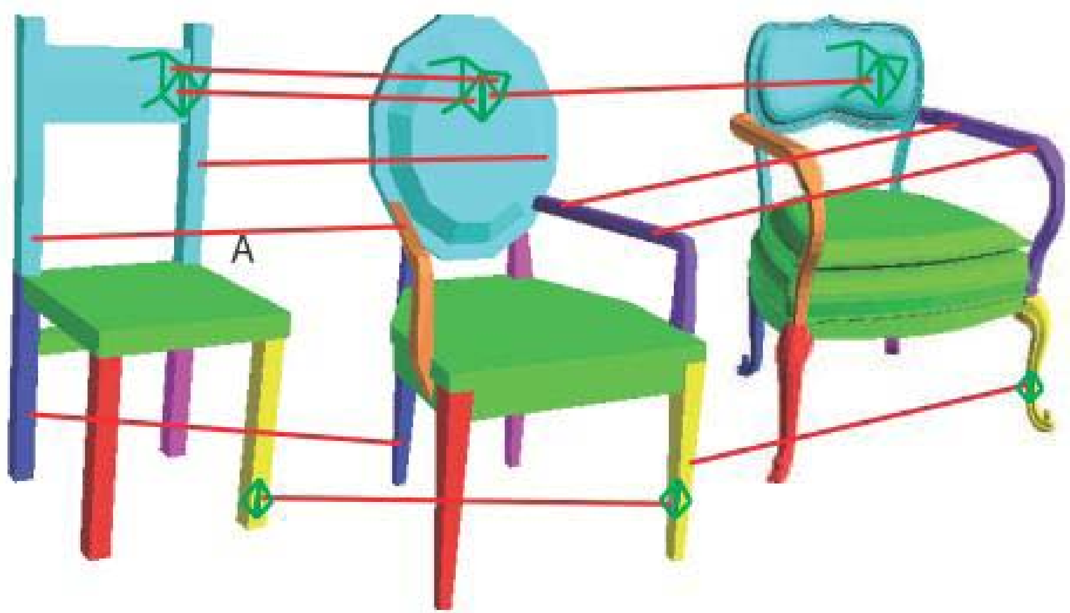

Clustering

- Over-segment each mesh independently
- Merge the segments by hierarchical clustering

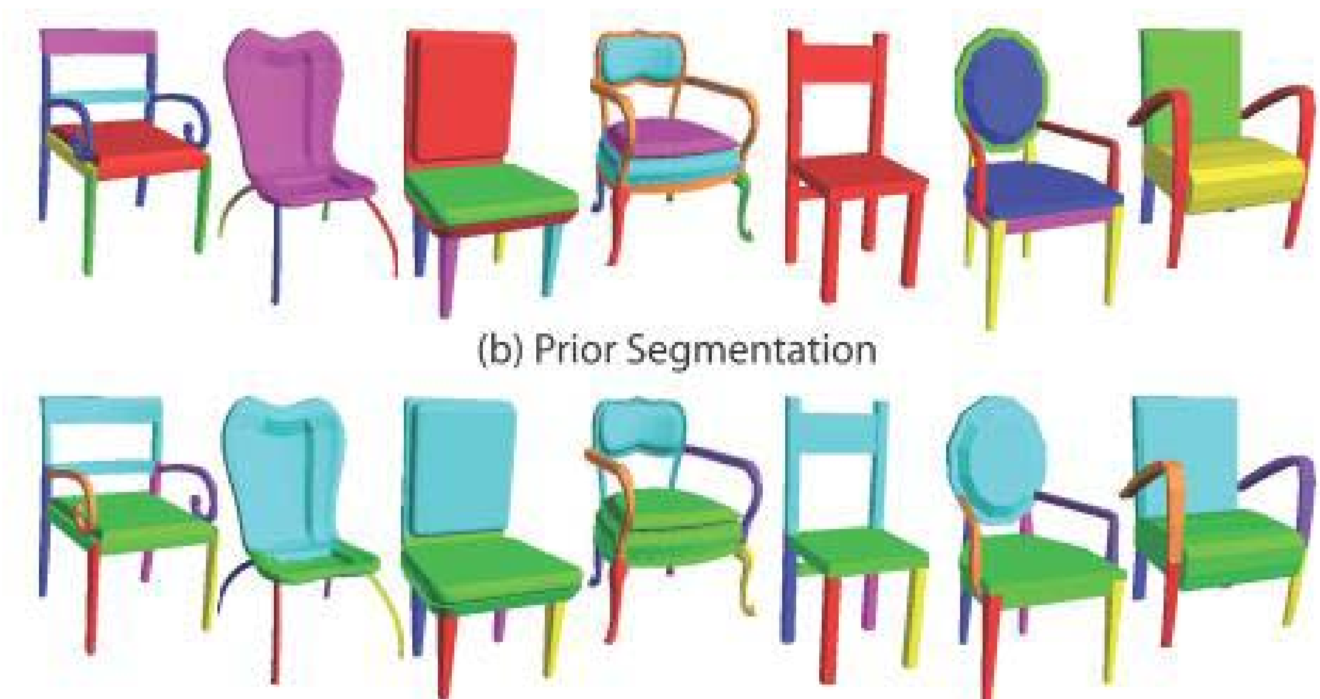

(b) Prior Segmentation

(c) Output Consistent Segmentation

Results

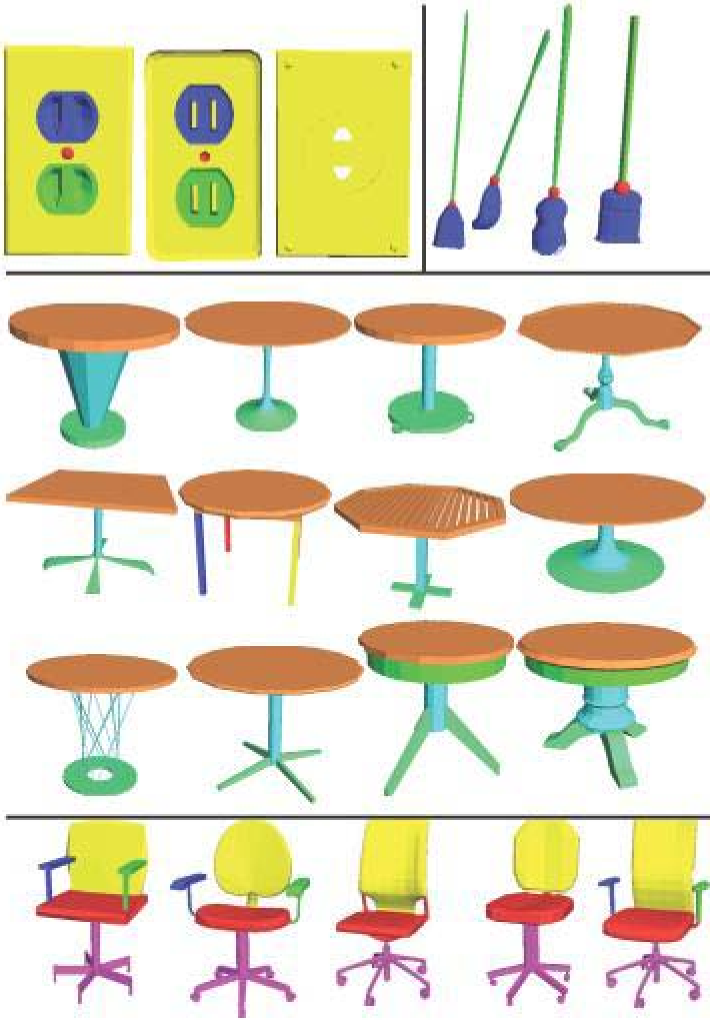

Consistent Segmentations

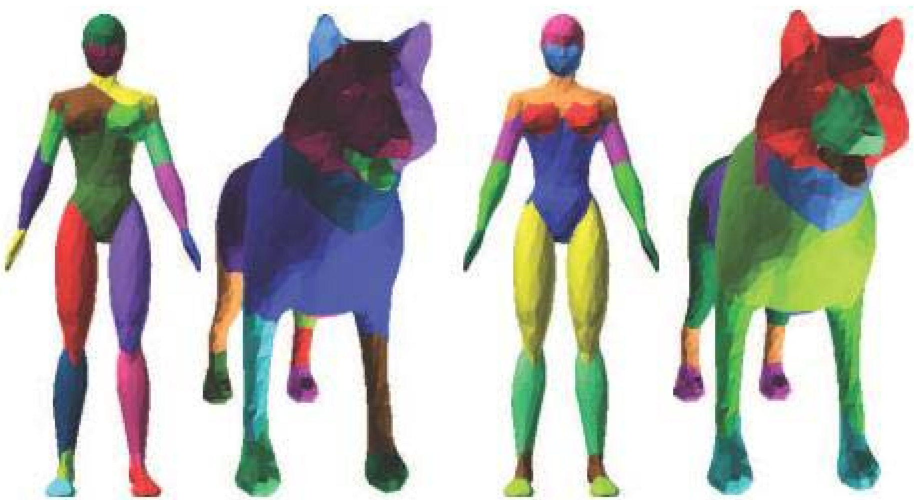

(a) Non-symmetric Segmentations (b) Symmetric Segmentations

Limitations

- created through a global similarity alignment
- consider low-level cues: adjacency and point correspondences

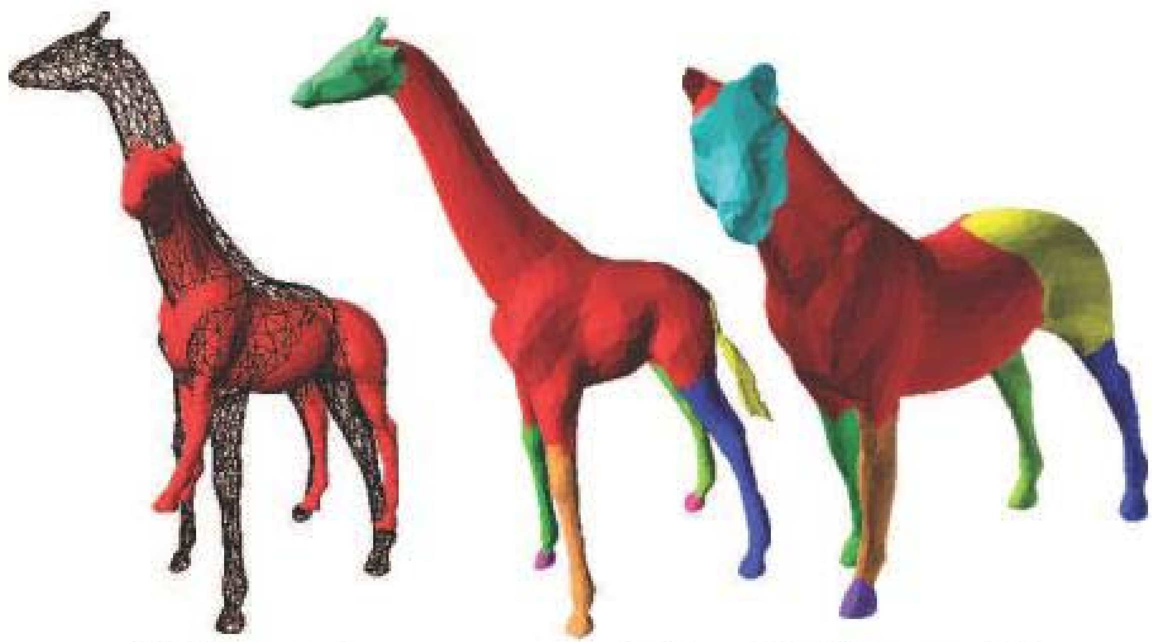

(a) Alignment (a) Consistent Segmentation

Learning 3D Mesh Segmentation and Labeling

Evangelos Kalogerakis, Aaron Hertzmann, Karan Singh

University of Toronto

Siggraph 2010

Segmentation and Labeling

- A data-driven approach to simultaneous segmentation and labeling of parts in 3D meshes

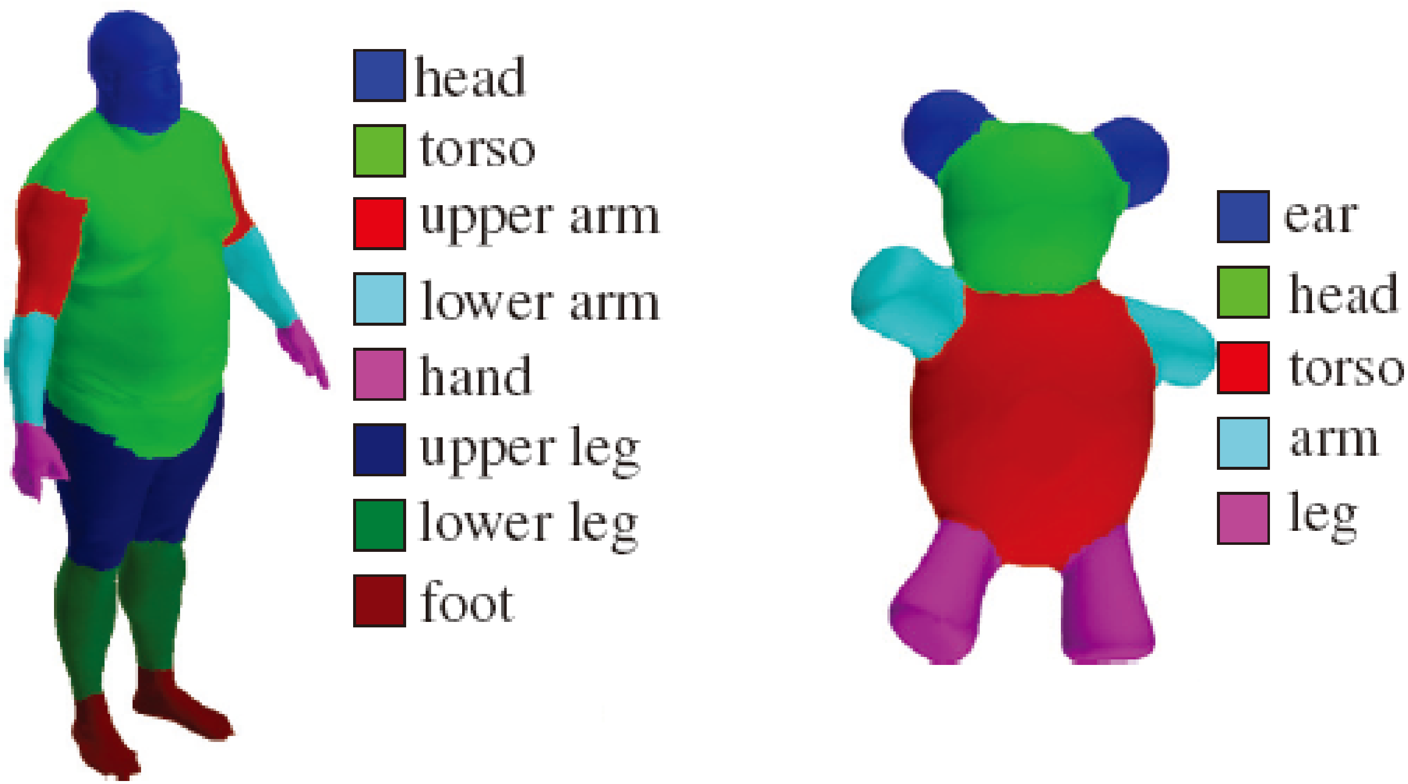

Training Set

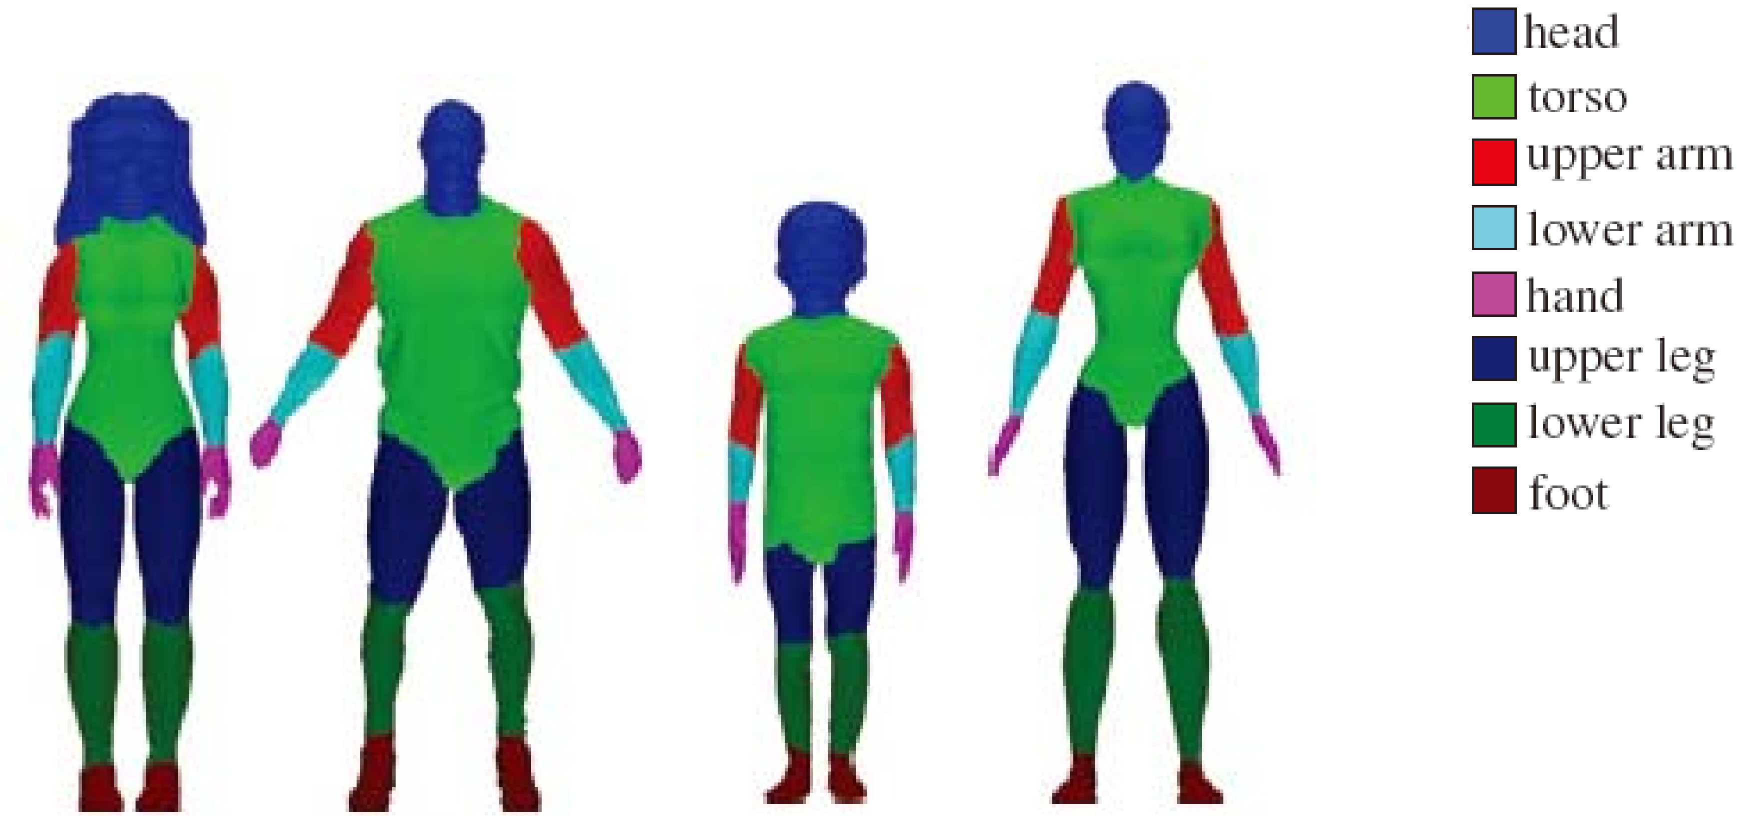

Algorithm

- **Goal**: to label each mesh face i with a label $l \in \mathcal{C}$
- **Define**:
  - each face i : *unary features* $\mathbf{x}_i$
  - each adjacent pair of faces : *pairwise features* $\mathbf{y}_{ij}$
  - objective function:

$$E(\mathbf{c};\theta) = \sum_i a_i E_1(c_i;\mathbf{x}_i,\theta_1) + \sum_{i,j} \ell_{ij} E_2(c_i,c_j;\mathbf{y}_{ij},\theta_2)$$

- **Conditional Random Field(CRF):**

$$P(\mathbf{c}|\mathbf{x},\mathbf{y},\theta) = \exp(-E(\mathbf{c};\theta))/Z(\mathbf{x},\mathbf{y},\theta)$$

- **JointBoost & hold-out validation**

Models

- Markov Random Field (MRF)
- Markov Chain
- Hidden Markov Model (HMM)
- Maximum Entropy Markov Model (MEMM)
- Conditional Random Field (CRF)

Energy Term

Unary Energy Term $E_1(c; \mathbf{x}, \theta_1) = -\log P(c|\mathbf{x}, \theta_1)$

Pairwise Energy Term $E_2(c, c'; \mathbf{y}, \theta_2) = L(c, c')\, G(\mathbf{y})$

- the label-compatibility term $L(c, c')$ measures the consistency between two adjacent labels.
- The geometry-dependent term $G(\mathbf{y})$ measures the likelihood of there being a difference in labels, as a function of the geometry alone.

$$
\begin{aligned}
G(\mathbf{y}) = & -\kappa \log P(c \neq c' | \mathbf{y}, \xi) \\
& -\lambda \log\left(1 - \min(\omega/\pi, 1) + \epsilon\right) + \mu
\end{aligned}
$$

Boosting

- **Committee**: combinations of models

$$y_{\mathrm{COM}}(\mathbf{x}) = \frac{1}{M}\sum_{m=1}^{M} y_m(\mathbf{x}).$$

- **Boosting**:
  – a powerful technique to produce a form of committee.
  – the sequential minimization of an exponential error function.

consider the expected error:

$$\mathbb{E}_{\mathbf{x},t}\left[\exp\{-ty(\mathbf{x})\}\right] = \sum_t \int \exp\{-ty(\mathbf{x})\}p(t|\mathbf{x})p(\mathbf{x})\,\mathrm{d}\mathbf{x}$$

$$\Rightarrow \quad y(\mathbf{x}) = \frac{1}{2}\ln\left\{\frac{p(t=1|\mathbf{x})}{p(t=-1|\mathbf{x})}\right\}$$

Results

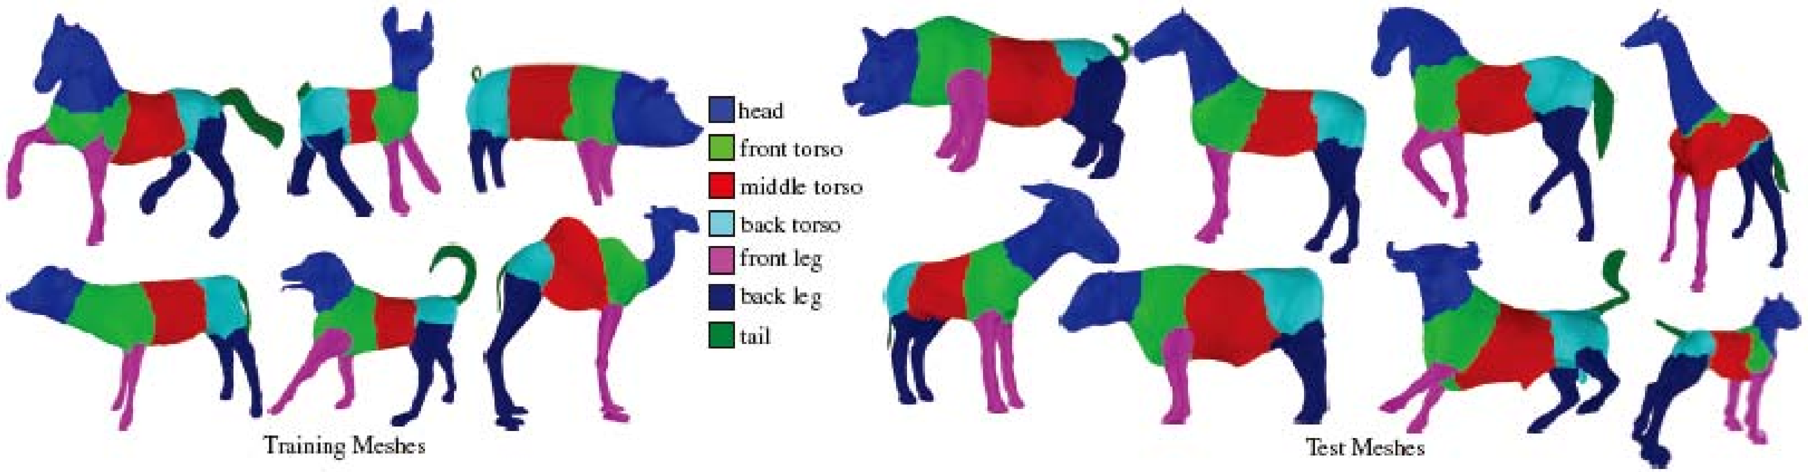

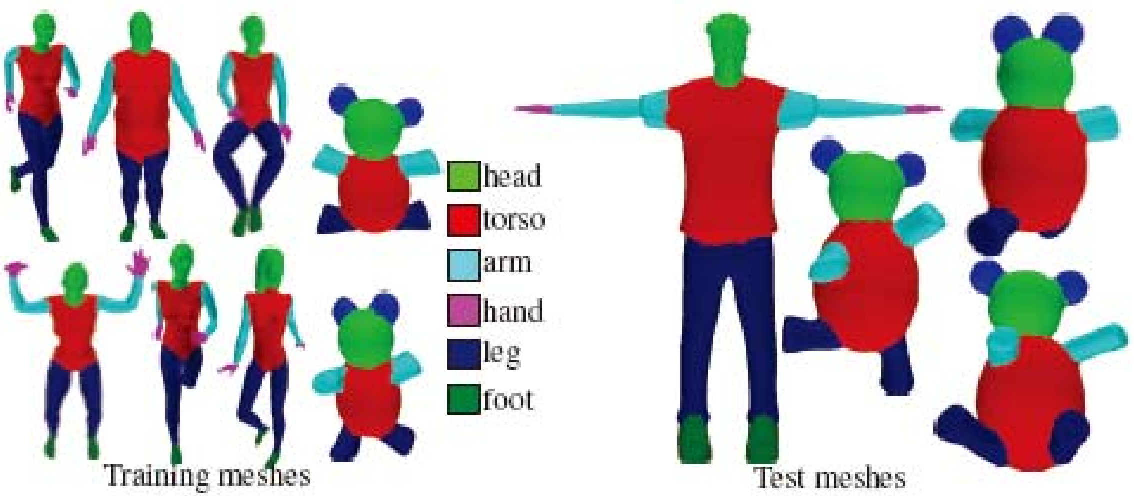

Merging categories

Results

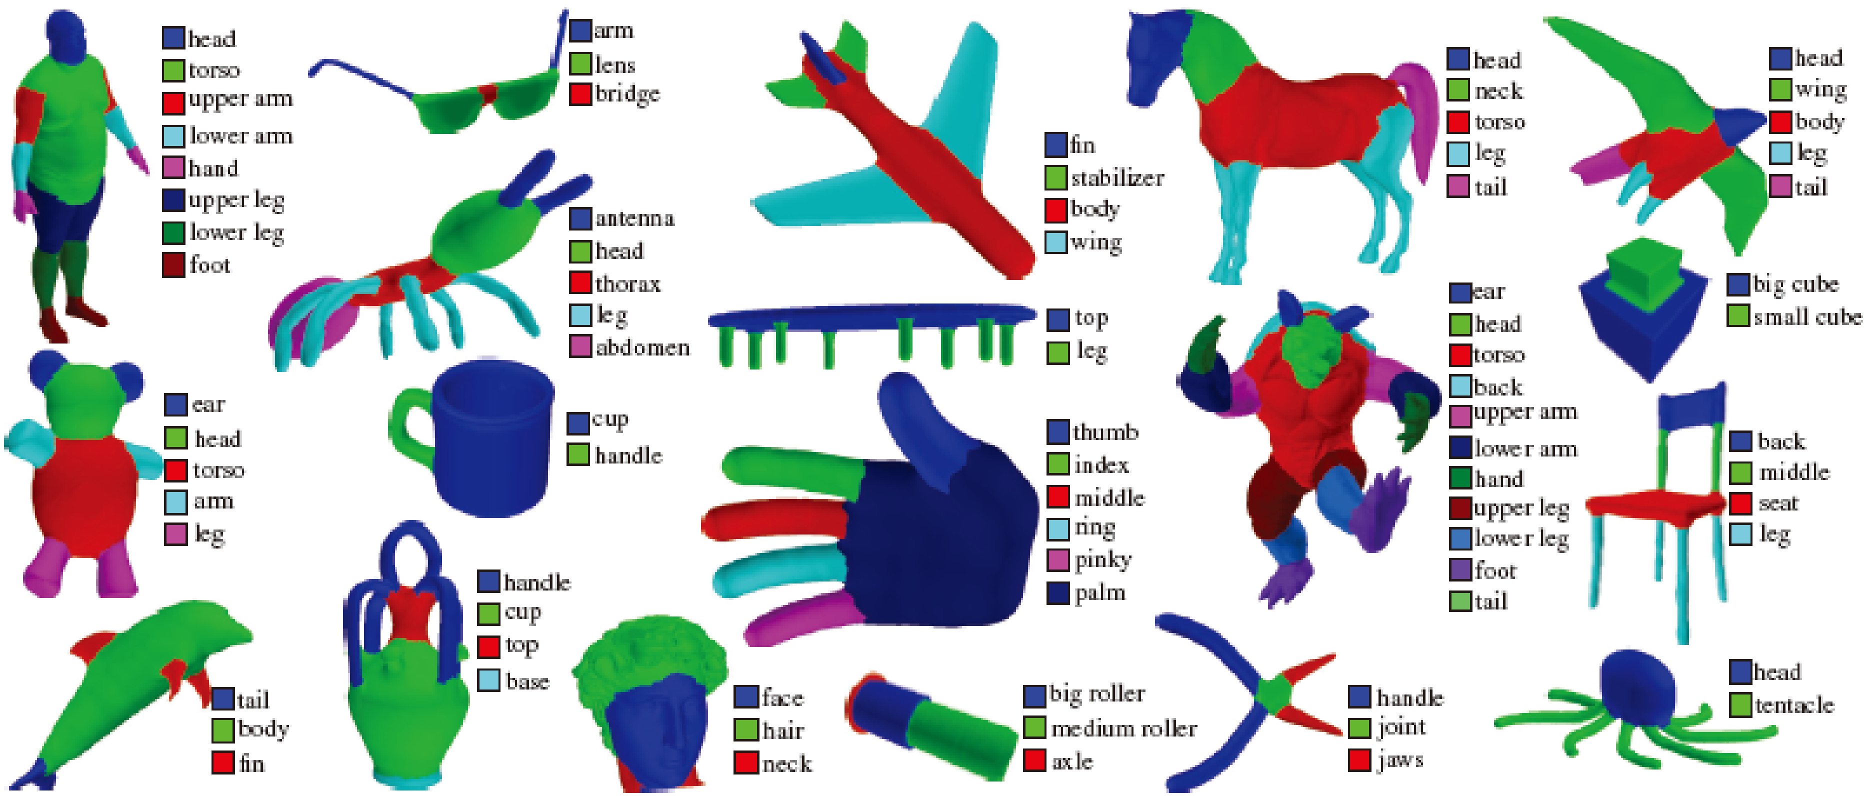
head
torso
upper arm
lower arm
hand
upper leg
lower leg
foot
arm
lens
bridge
antenna
head
thorax
leg
abdomen
fin
stabilizer
body
wing
head
neck
torso
leg
tail
head
wing
body
leg
tail
top
leg
ear
head
torso
arm
leg
cup
handle
thumb
index
middle
ring
pinky
palm
ear
head
torso
back
upper arm
lower arm
hand
upper leg
lower leg
foot
tail
big cube
small cube
back
middle
seat
leg
handle
cup
top
base
tail
body
fin
face
hair
neck
big roller
medium roller
axle
handle
joint
jaws
head
tentacle

Style-Content Separation by Anisotropic Part Scales

Kai Xu*† Honghua Li*† Hao Zhang* Daniel Cohen-Or‡ Yueshan Xiong† Zhi-Quan Cheng†
*Simon Fraser University †National University of Defense Technology ‡Tel-Aviv University

Siggraph Asia 2010

Replicating the “style” of an input

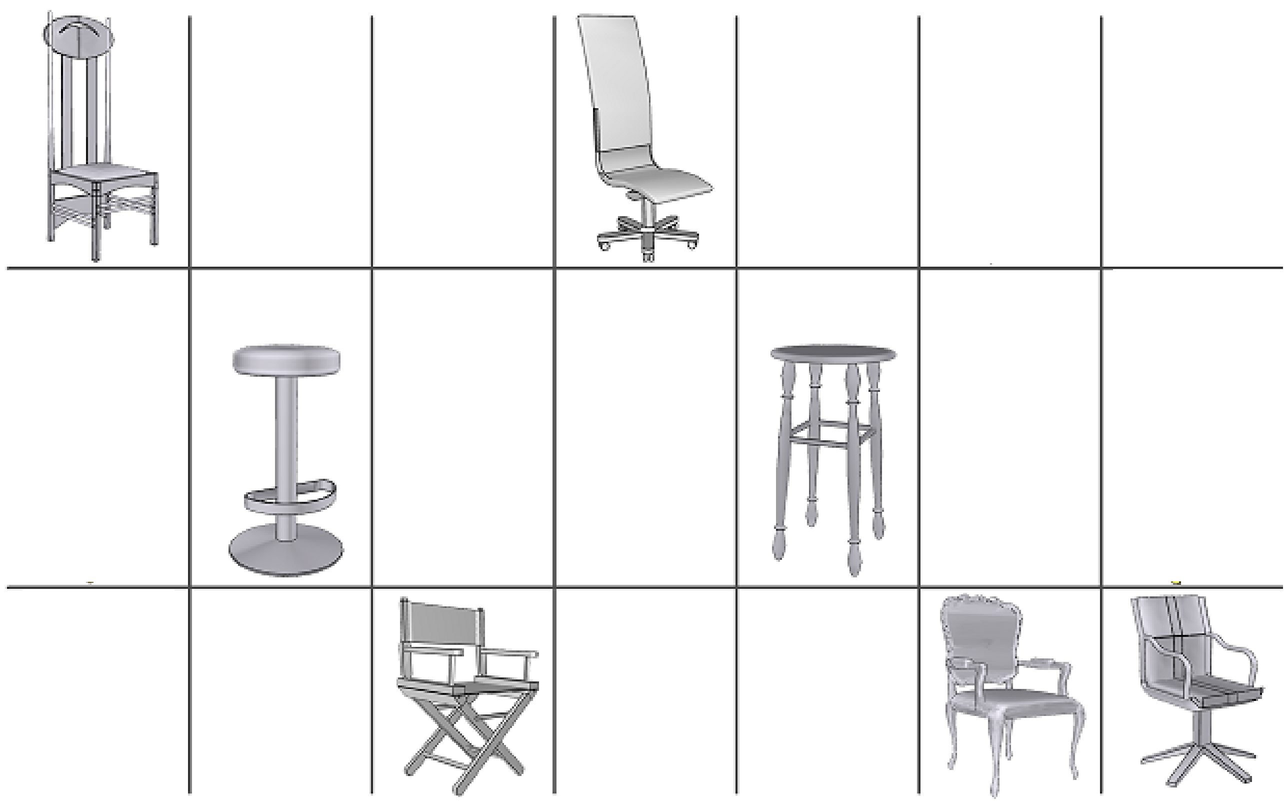

Replicating the “style” of an input

Style-Content Separation

- Co-analysis of a set of man-made 3D objects
- Creation instances derived from the set
  - more of the same

Co-Segmentation

- Over-segmentation
- Graph construction [SMI 2009]
- Inter-style correspondence
- Content classification

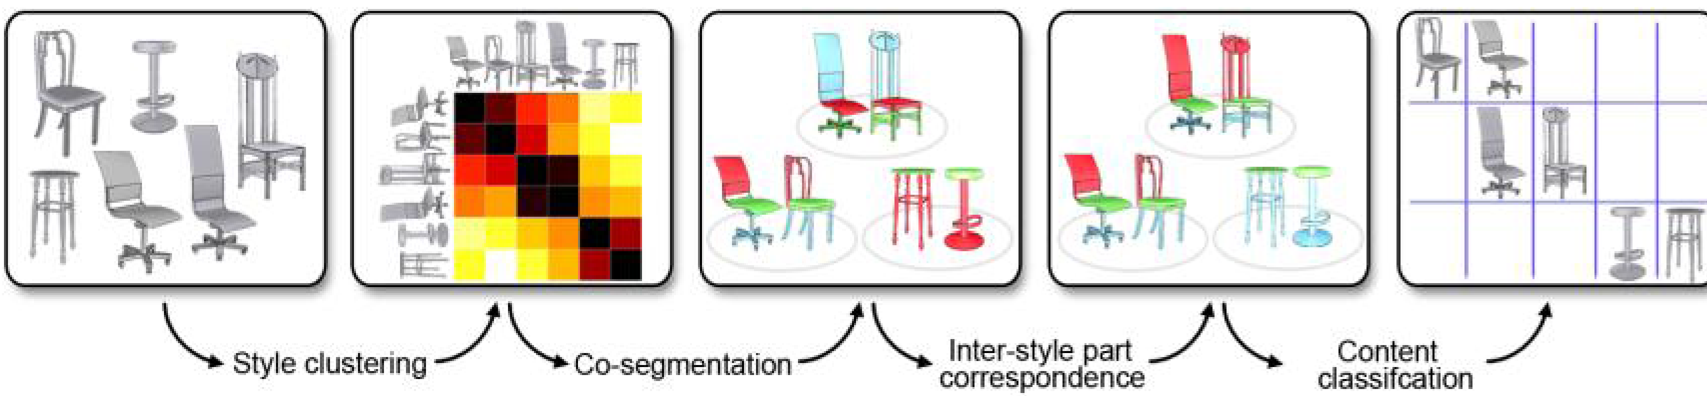

Figure 2: *An overview of the co-analysis pipeline.*

Results

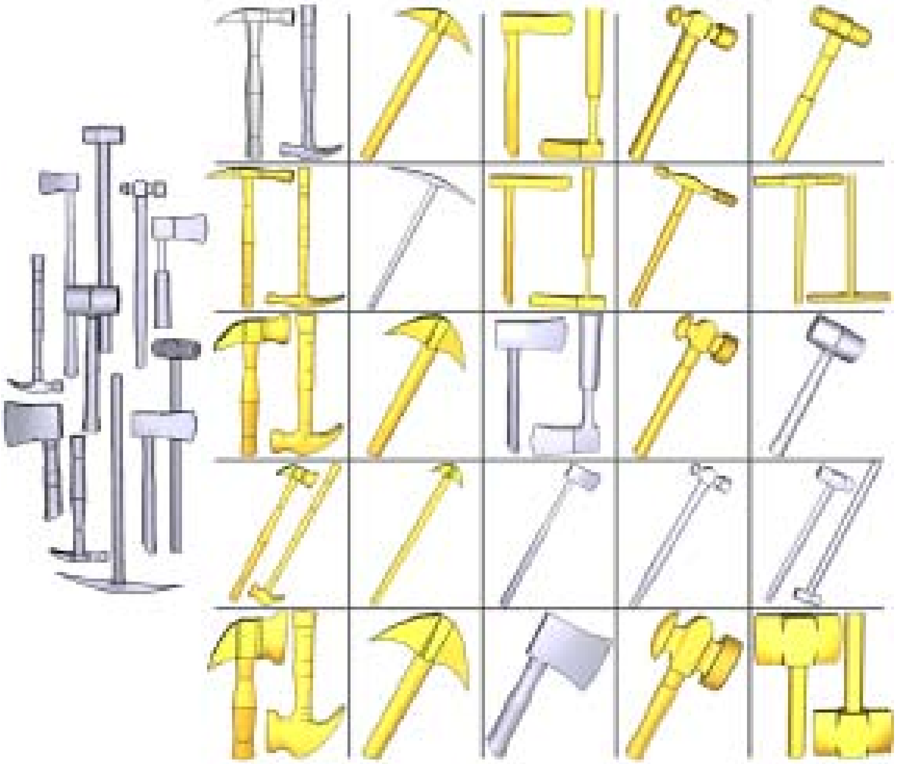

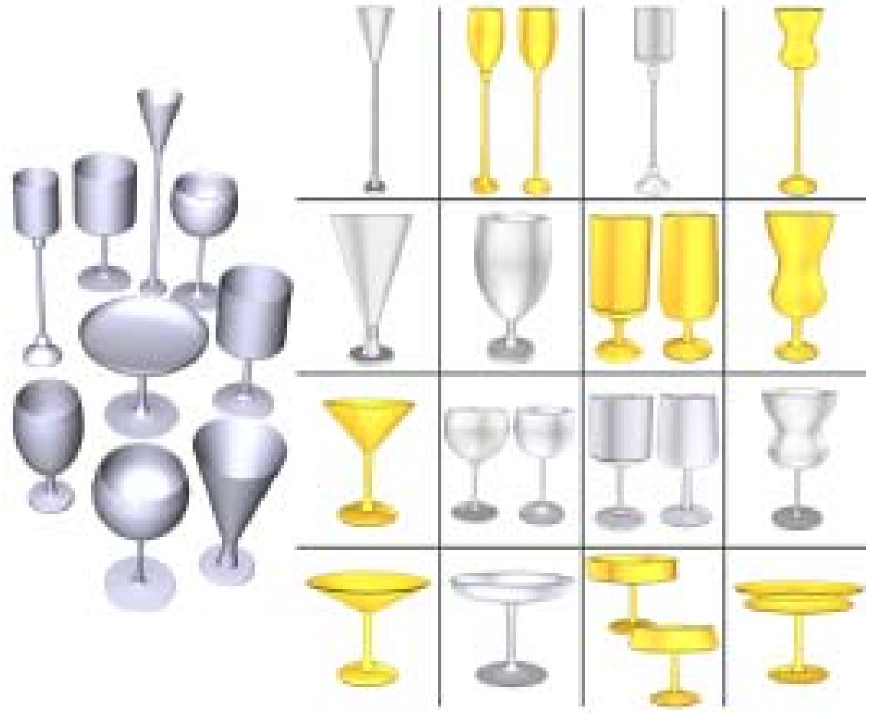

Results

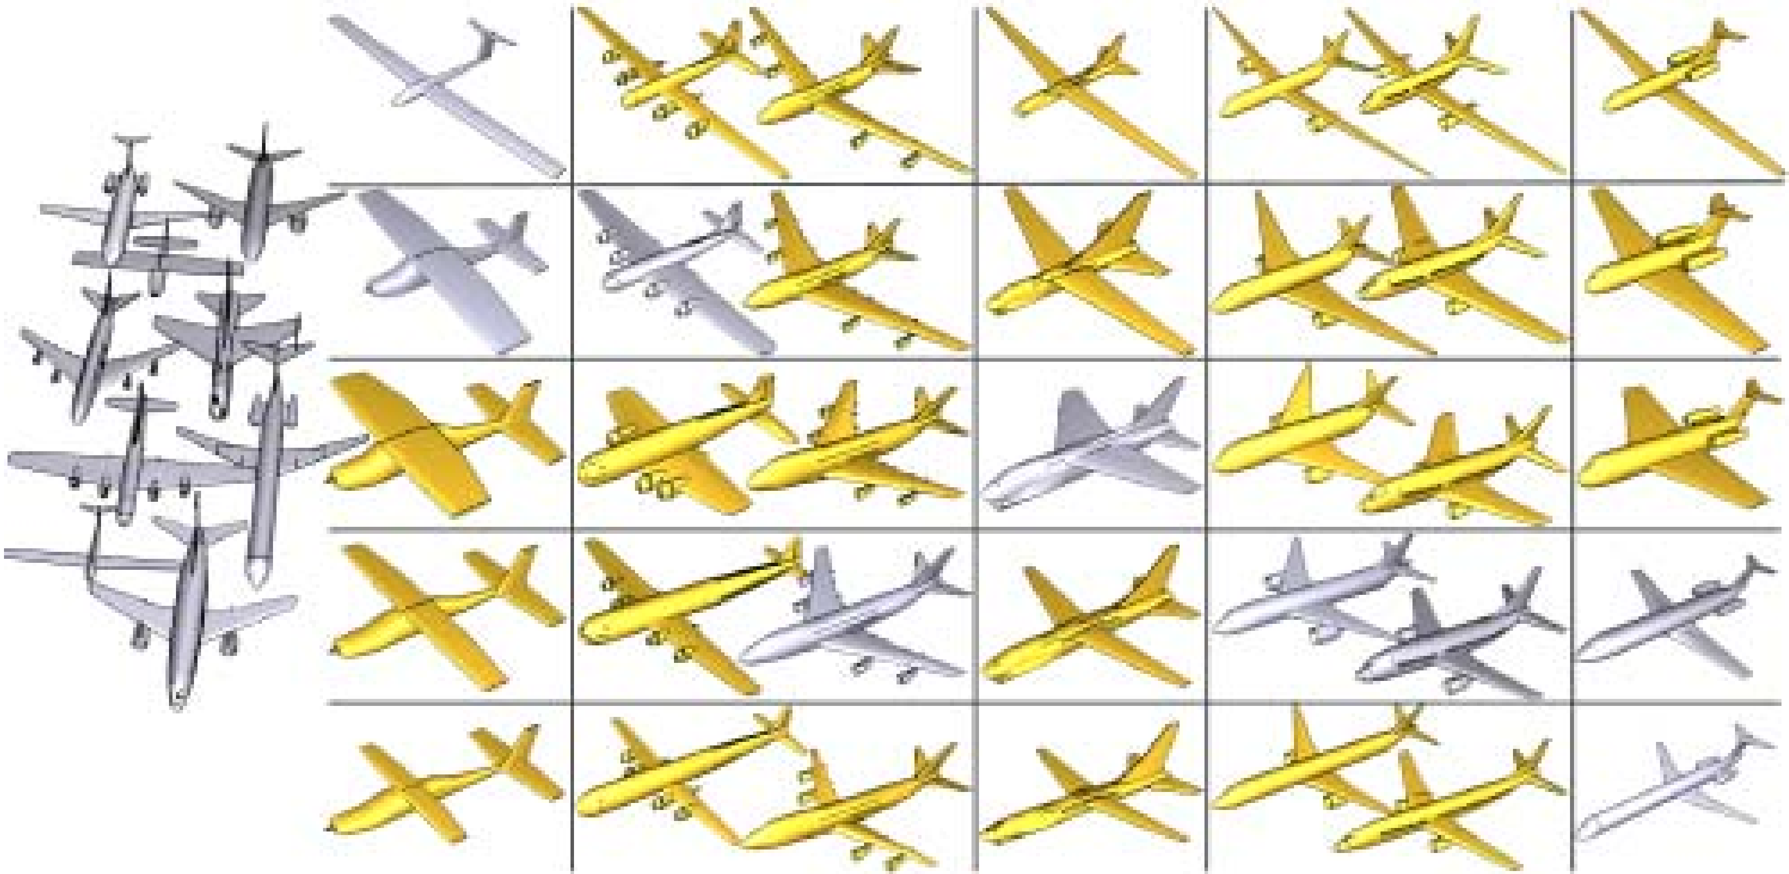

Prior Knowledge for Part Correspondence

Oliver van Kaick[1], Andrea Tagliasacchi[1], Oana Sidi[2], Hao Zhang[1], Daniel Cohen-Or[2], Lior Wolf[2], and Ghassan Hamarneh[1]

[1] School of Computing Science, Simon Fraser University, Canada
[2] School of Computer Science, Tel Aviv University, Israel

Eurographics 2011

Goal

- Models
  - large variations in the geometry or topology of the corresponding parts
- Content-driven analysis -> Knowledge-driven analysis
  - Prior knowledge on the parts would play a more dominant role than geometric similarity
- Joint labeling scheme

Overview

- Content-driven analysis
  - shape descriptors
  - $D_S = \{D_{s1}, D_{s2}, \ldots, D_{sn}\}$ and $D_T = \{D_{t1}, D_{t2}, \ldots, D_{tm}\}$,
  - $C = A(D_S, D_T)$
- Incorporation of semantic knowledge
  - Pre-segmented training set
- Probabilistic semantic labeling
- Joint labeling and part correspondence

Prior Knowledge

- Training set
  - Manually segmented shapes
  - semantic labeling of each part
- Shape descriptors
  - Kalogerakis et al. [KHS10]
  - Curvature, SDF, geodesic distances, and binning of face areas into geodesic shape contexts
- Classifier training
  - Train a classifier K_l with the descriptors D_l
  - "gentleboost" algorithm [KHS10]

Probabilistic Semantic Labeling

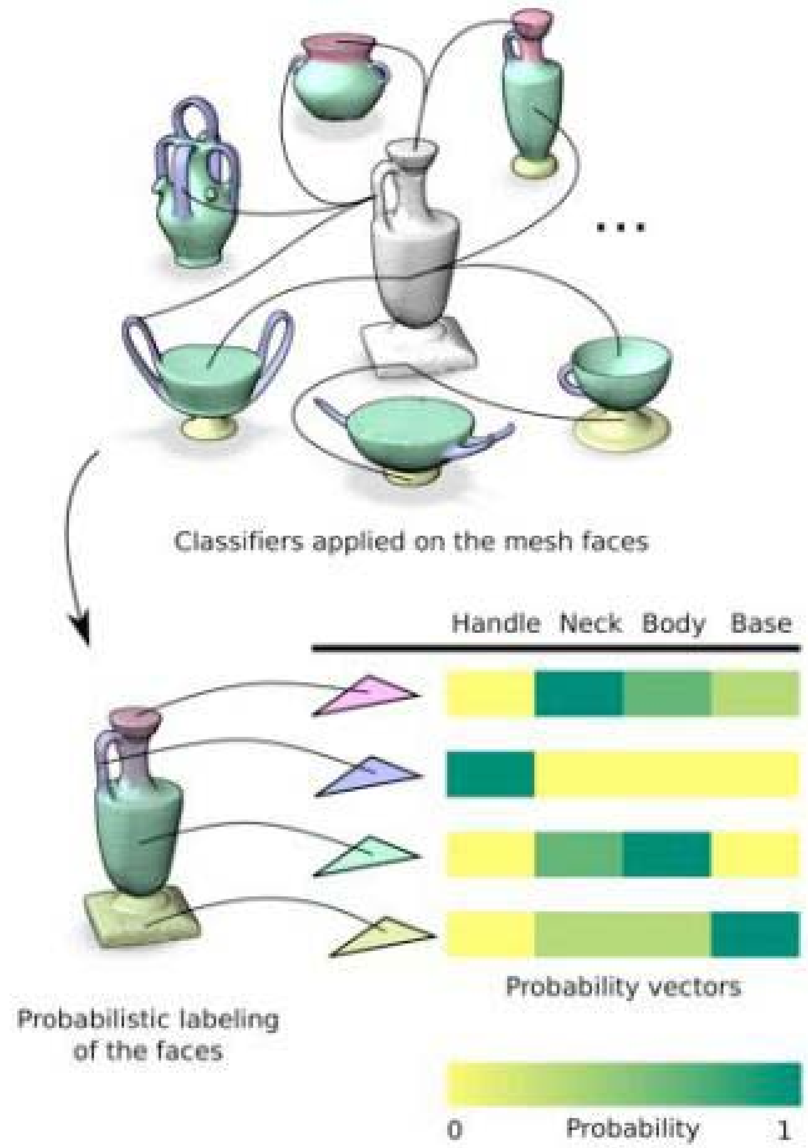

Joint Labeling

- Labeling energy

$$\mathcal{E}(\mathbf{l}) = \sum_{i \in V} \mathcal{U}(i, l_i) + \sum_{ij \in E_{\text{intra}}} \mathcal{B}_{\text{intra}}(i, j, l_i, l_j) + \sum_{ij \in E_{\text{inter}}} \mathcal{B}_{\text{inter}}(i, j, l_i, l_j),$$

$$\mathcal{U}(i, l_i) = -a_i \log P(l_i | \mathbf{x}_i)$$

$$\mathcal{B}_{\text{intra}}(i, j, l_i, l_j) = L(l_i, l_j)[\lambda \alpha_{ij} + \mu \ell_{ij}]$$

$$\mathcal{B}_{\text{inter}}(i, j, l_i, l_j) = L(l_i, l_j)[\nu \sigma_{ij}]$$

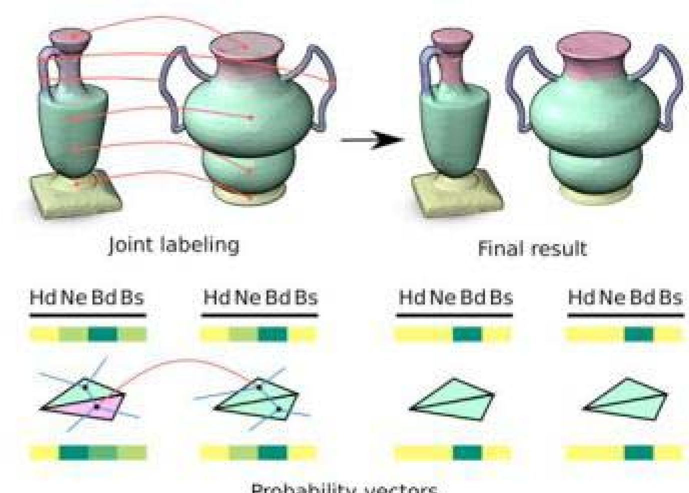

Results

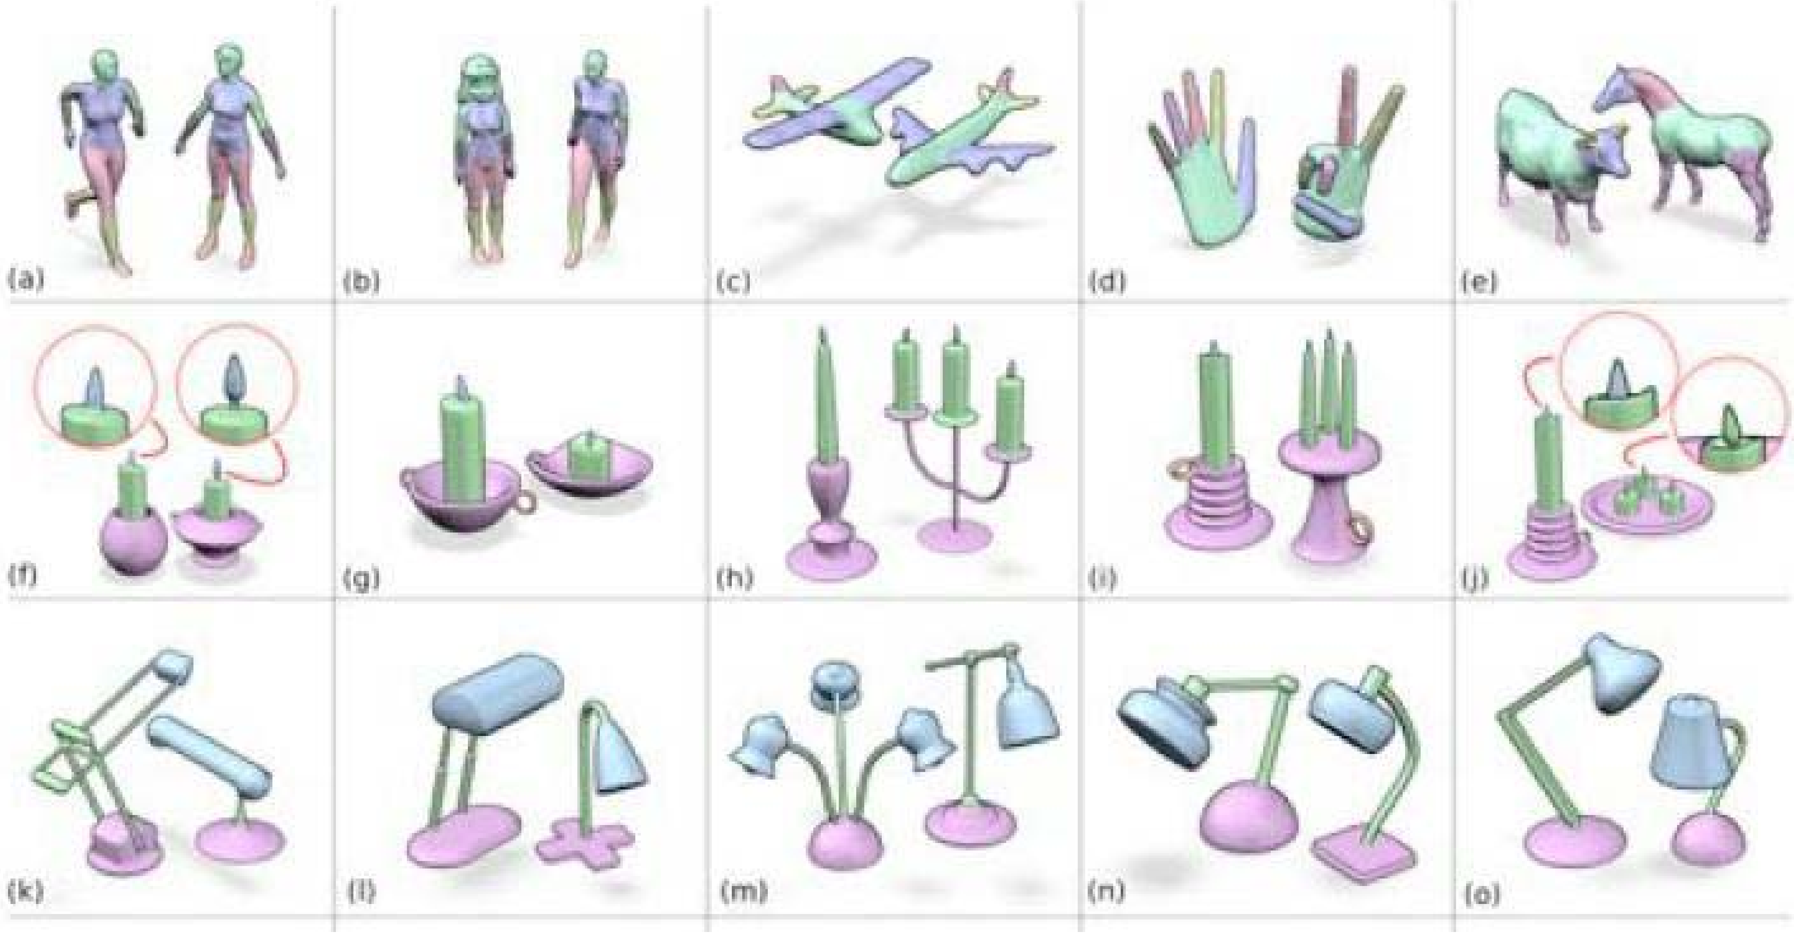

Conclusion

- Prior knowledge based
  - Reliance of low-level shape descriptors
- Computational cost

Joint Shape Segmentation with Linear Programming

Qixing Huang, Vladlen Koltun, Leonidas Guibas
Stanford University
Siggraph Asia 2011

Joint Segmentation

- Segments the shapes jointly, utilizing features from multiple shapes to improve the segmentation of each

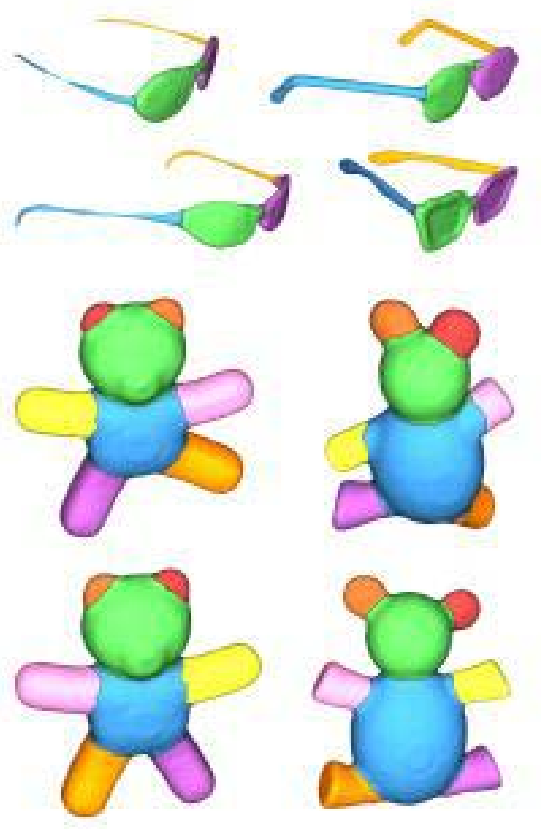

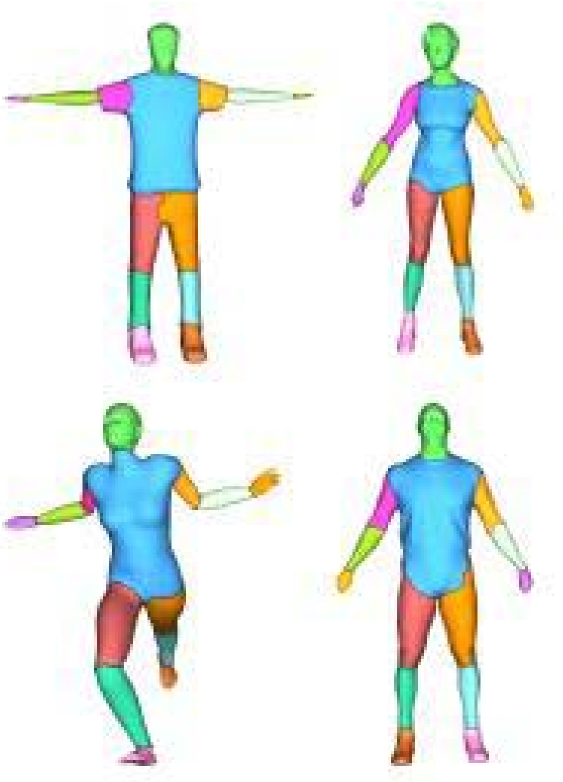

Overview

- Initial segments - randomized clustering
- Pairwise joint segmentation - identify pairs of similar shapes
- Multiway joint segmentation - increase segmentation quality across pairs of similar shapes

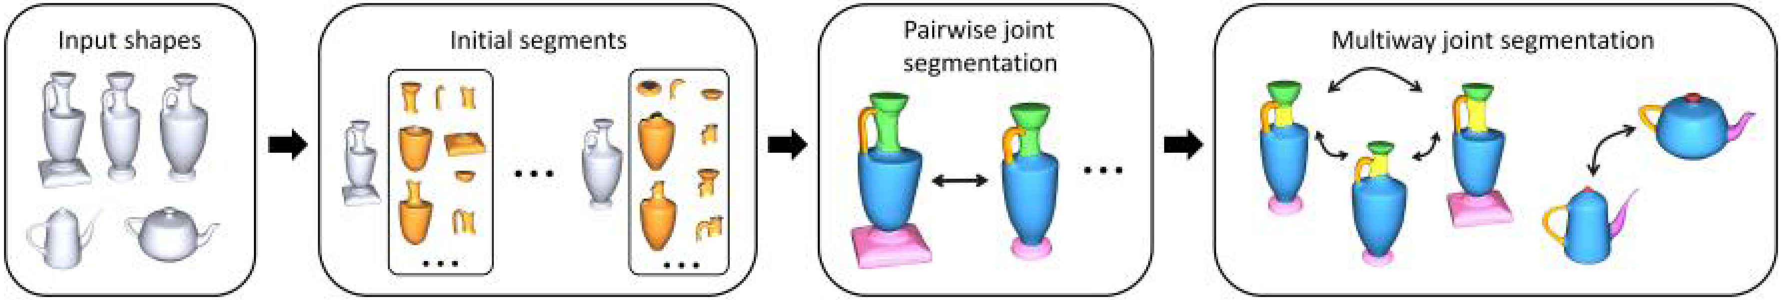

Figure 2: *Overview of our approach. In the first stage, we produce a set of initial segments for each shape. In the second stage, each pair of shapes is jointly segmented in order to identify similar shapes. In the third stage, a global optimization is performed over segmentations of all shapes together with correspondences between similar shapes.*

Results

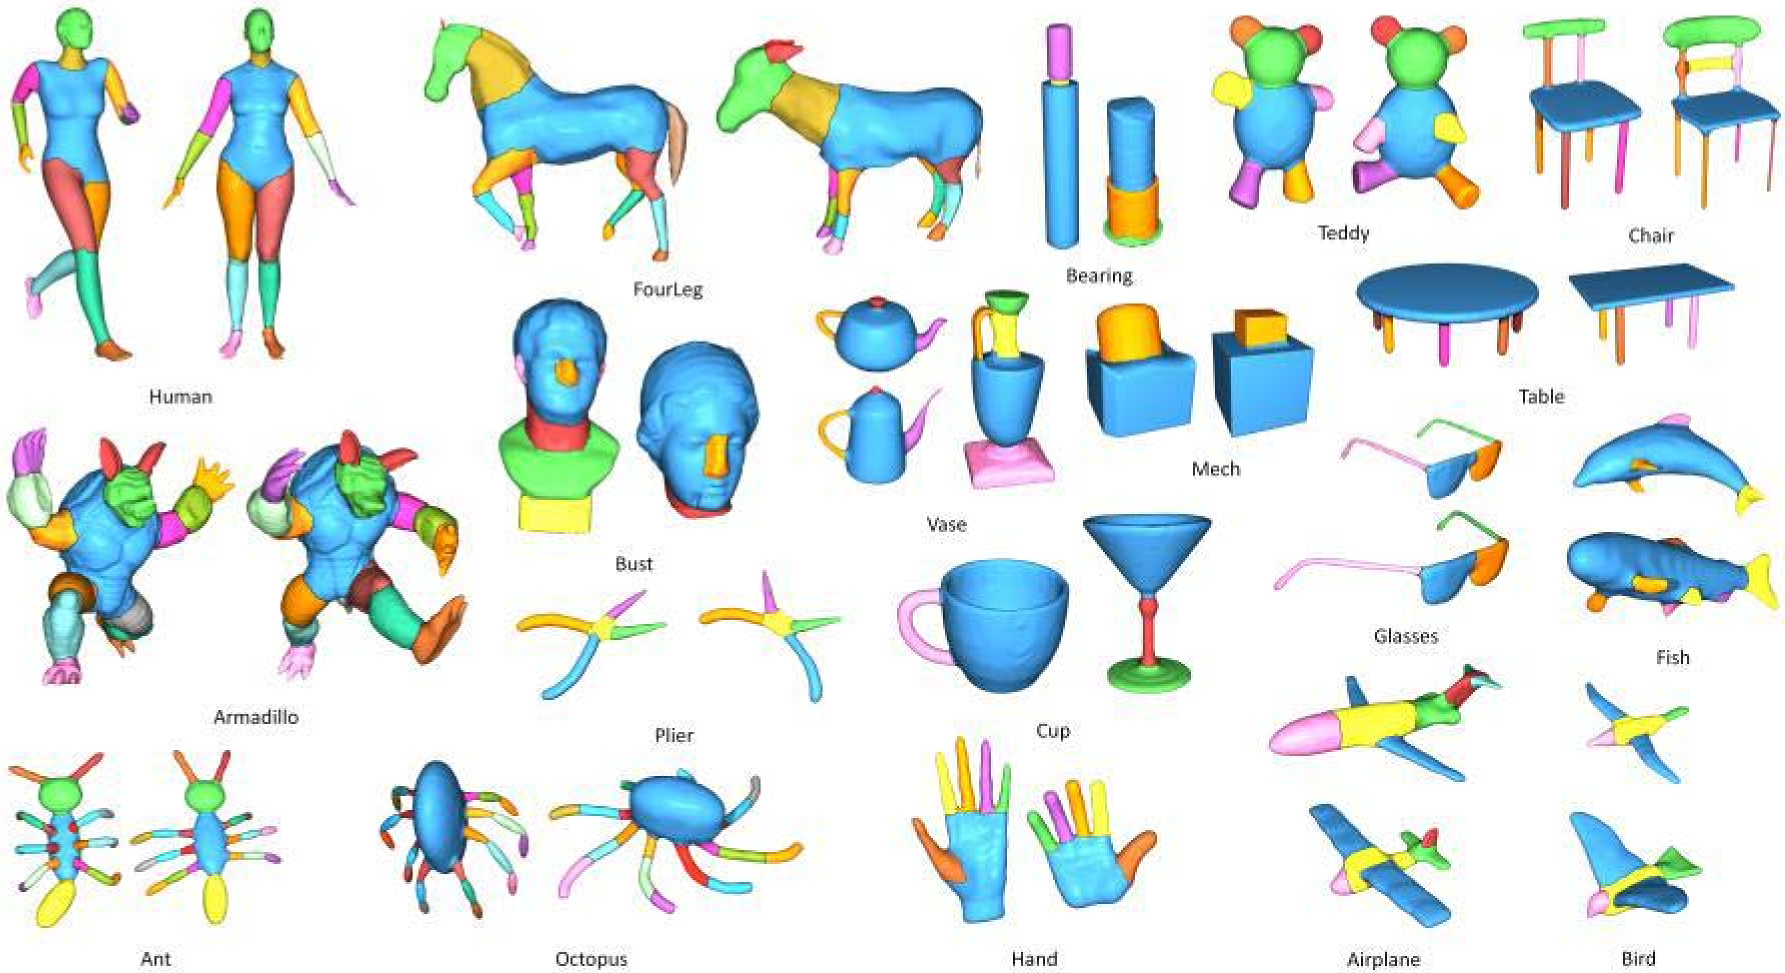
Human
FourLeg
Bearing
Teddy
Chair
Table
Bust
Vase
Mech
Armadillo
Glasses
Fish
Plier
Cup
Ant
Octopus
Hand
Airplane
Bird

Conclusion

- Segments in the final segmentation of each shape are generated from the initially computed patches
- Computational cost

Unsupervised Co-Segmentation of a Set of Shapes via Descriptor-Space Spectral Clustering

Oana Sidi* Oliver van Kaick† Yanir Kleiman* Hao Zhang† Daniel Cohen-Or*
*Tel-Aviv University †Simon Fraser University

Siggraph Asia 2011

Overview

- input: a set of meshes from a given family, $G = \{V, E\}$
- derive a statistical model for each class of parts:

$$p(f|c_i) = p(f, \mu_i, \Sigma_i) = C\, e^{-\frac{1}{2}(f-\mu_i)^T \Sigma_i^{-1}(f-\mu_i)} \qquad p(c_i|f) \propto p(f|c_i)p(c_i)$$

- define a collection of energies over each mesh:

$$\mathcal{E}(l) = \sum_{u \in V} \mathcal{E}_D(u, l_u) + \sum_{uv \in E} \mathcal{E}_S(u, v, l_u, l_v)$$

  - $\mathcal{E}_D(u, l_u) = -\omega \log(p(c_{l_u}|u))$
  - $\mathcal{E}_S(u, v, l_u, l_v) = \begin{cases} 0, & \text{if } l_u = l_v \\ -\log(\theta_{uv}/\pi)\, l_{uv}, & \text{otherwise} \end{cases}$

  where l_{uv} is the length of the edge between the faces corresponding to u and v, and θ_{uv} is the dihedral angle between the two faces.

graph-cuts

Statistical Model

- Extract shape descriptor at the face-level
  - the upright orientation vector [Fu et al. Sig08]
  - the angle between the normal of the face
  - the geodesic distance from the base of the shape to the face
- Per-object segmentation: *mean-shift algorithm*
  - Segment-level descriptor: for a segment s_i
    - the histogram of face-level descriptor: h_i^d
    - the segment area: a_i
    - the overall geometry of the segment: $g_i = [\mu_l \; \mu_p \; \mu_s]$

$$\mu_l = \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \lambda_3}, \quad \mu_p = \frac{2(\lambda_2 - \lambda_3)}{\lambda_1 + \lambda_2 + \lambda_3}, \quad \mu_s = \frac{3\lambda_3}{\lambda_1 + \lambda_2 + \lambda_3}, \quad \text{with} \quad \lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$$

λ_1, λ_2, and λ_3 are the three eigenvalues obtained when applying principal component analysis to all the vertices that are part of the segment.

Descriptor-space Spectral Clustering

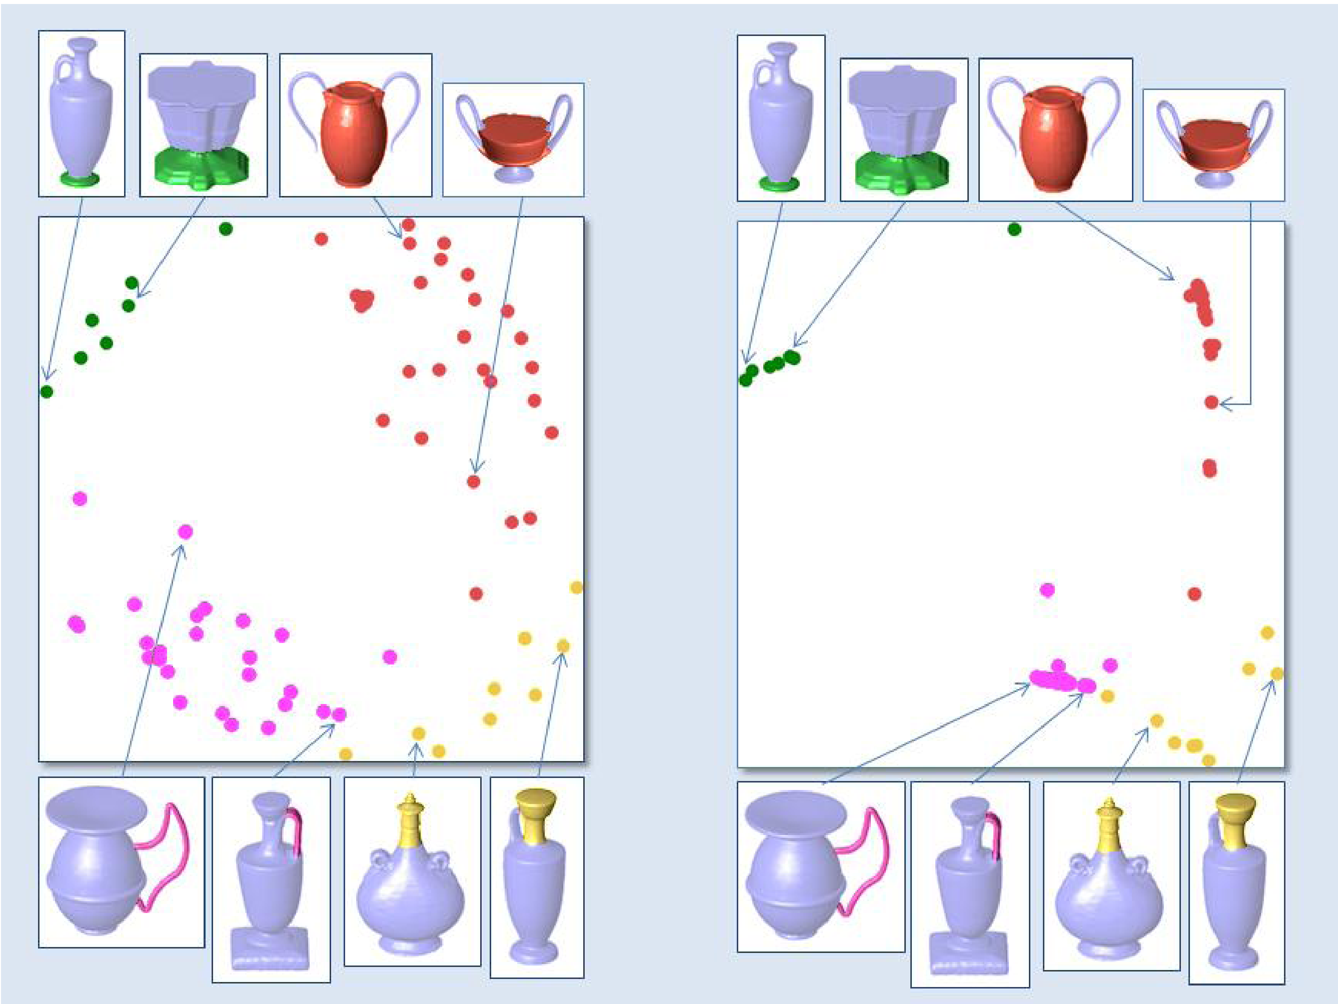

(b) Diffusion maps

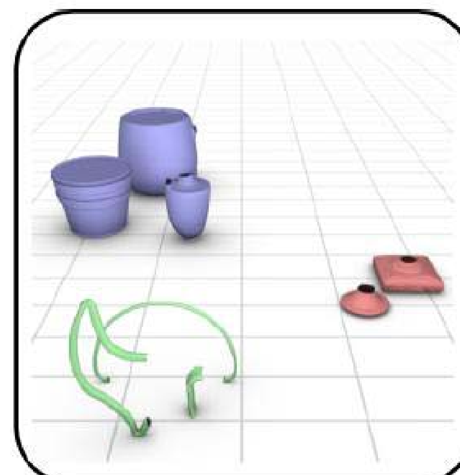

(c) Clustering

Results

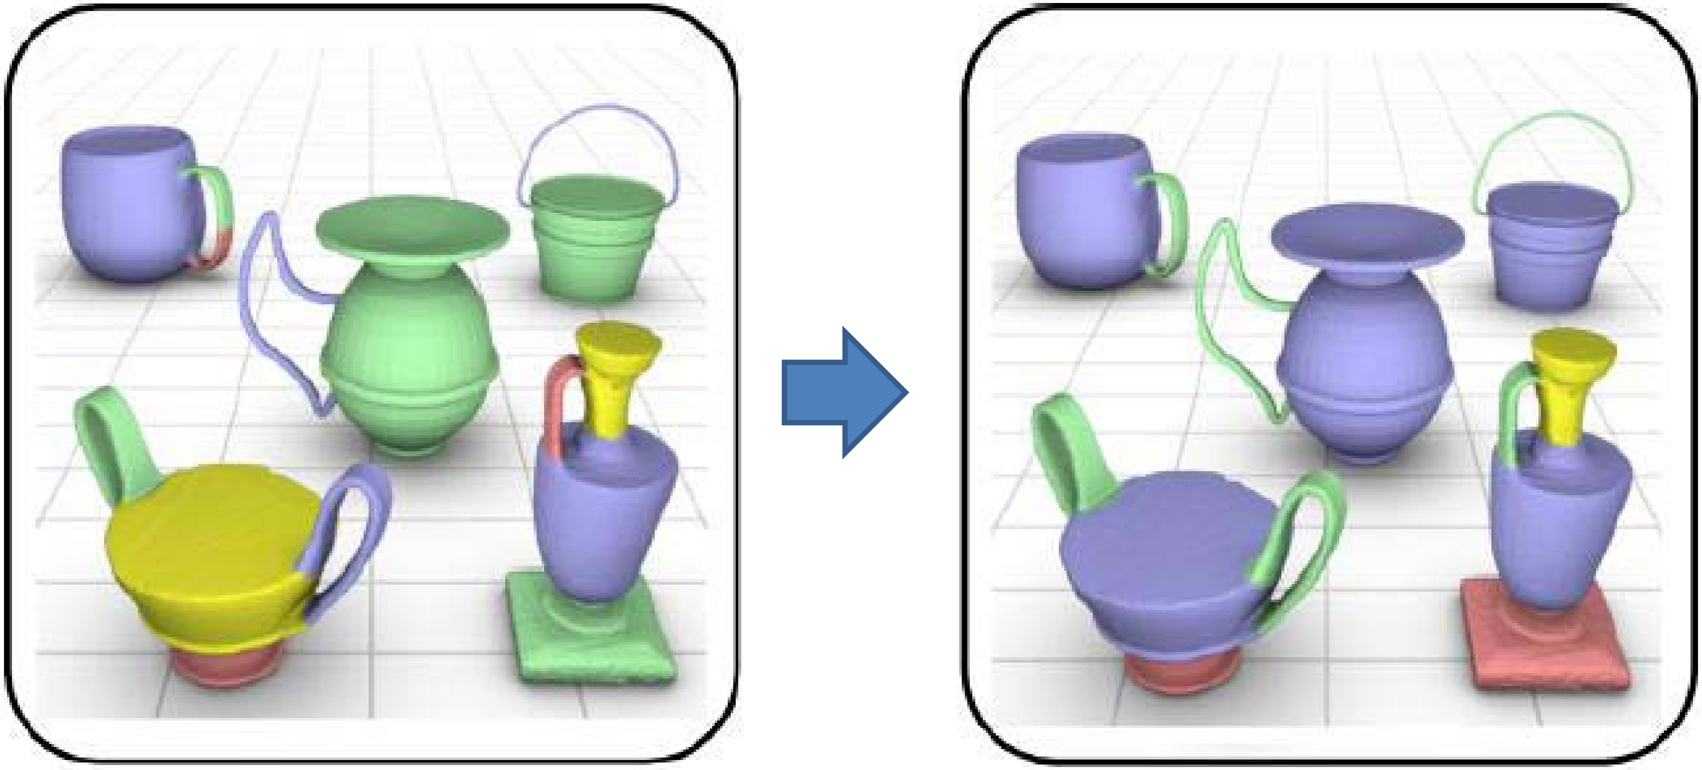

Results

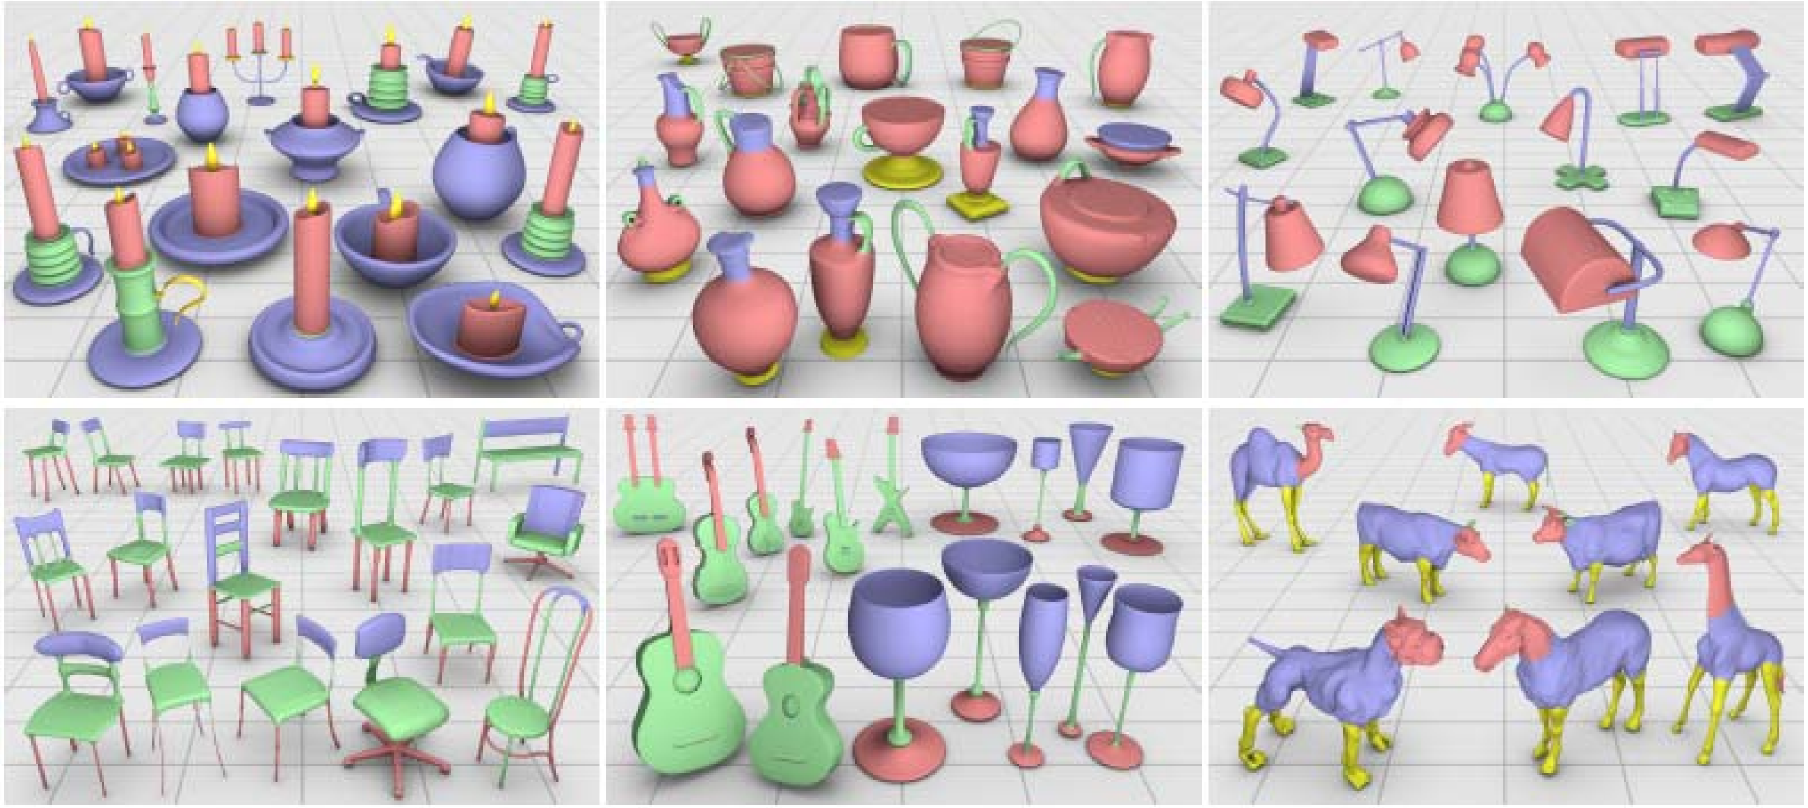

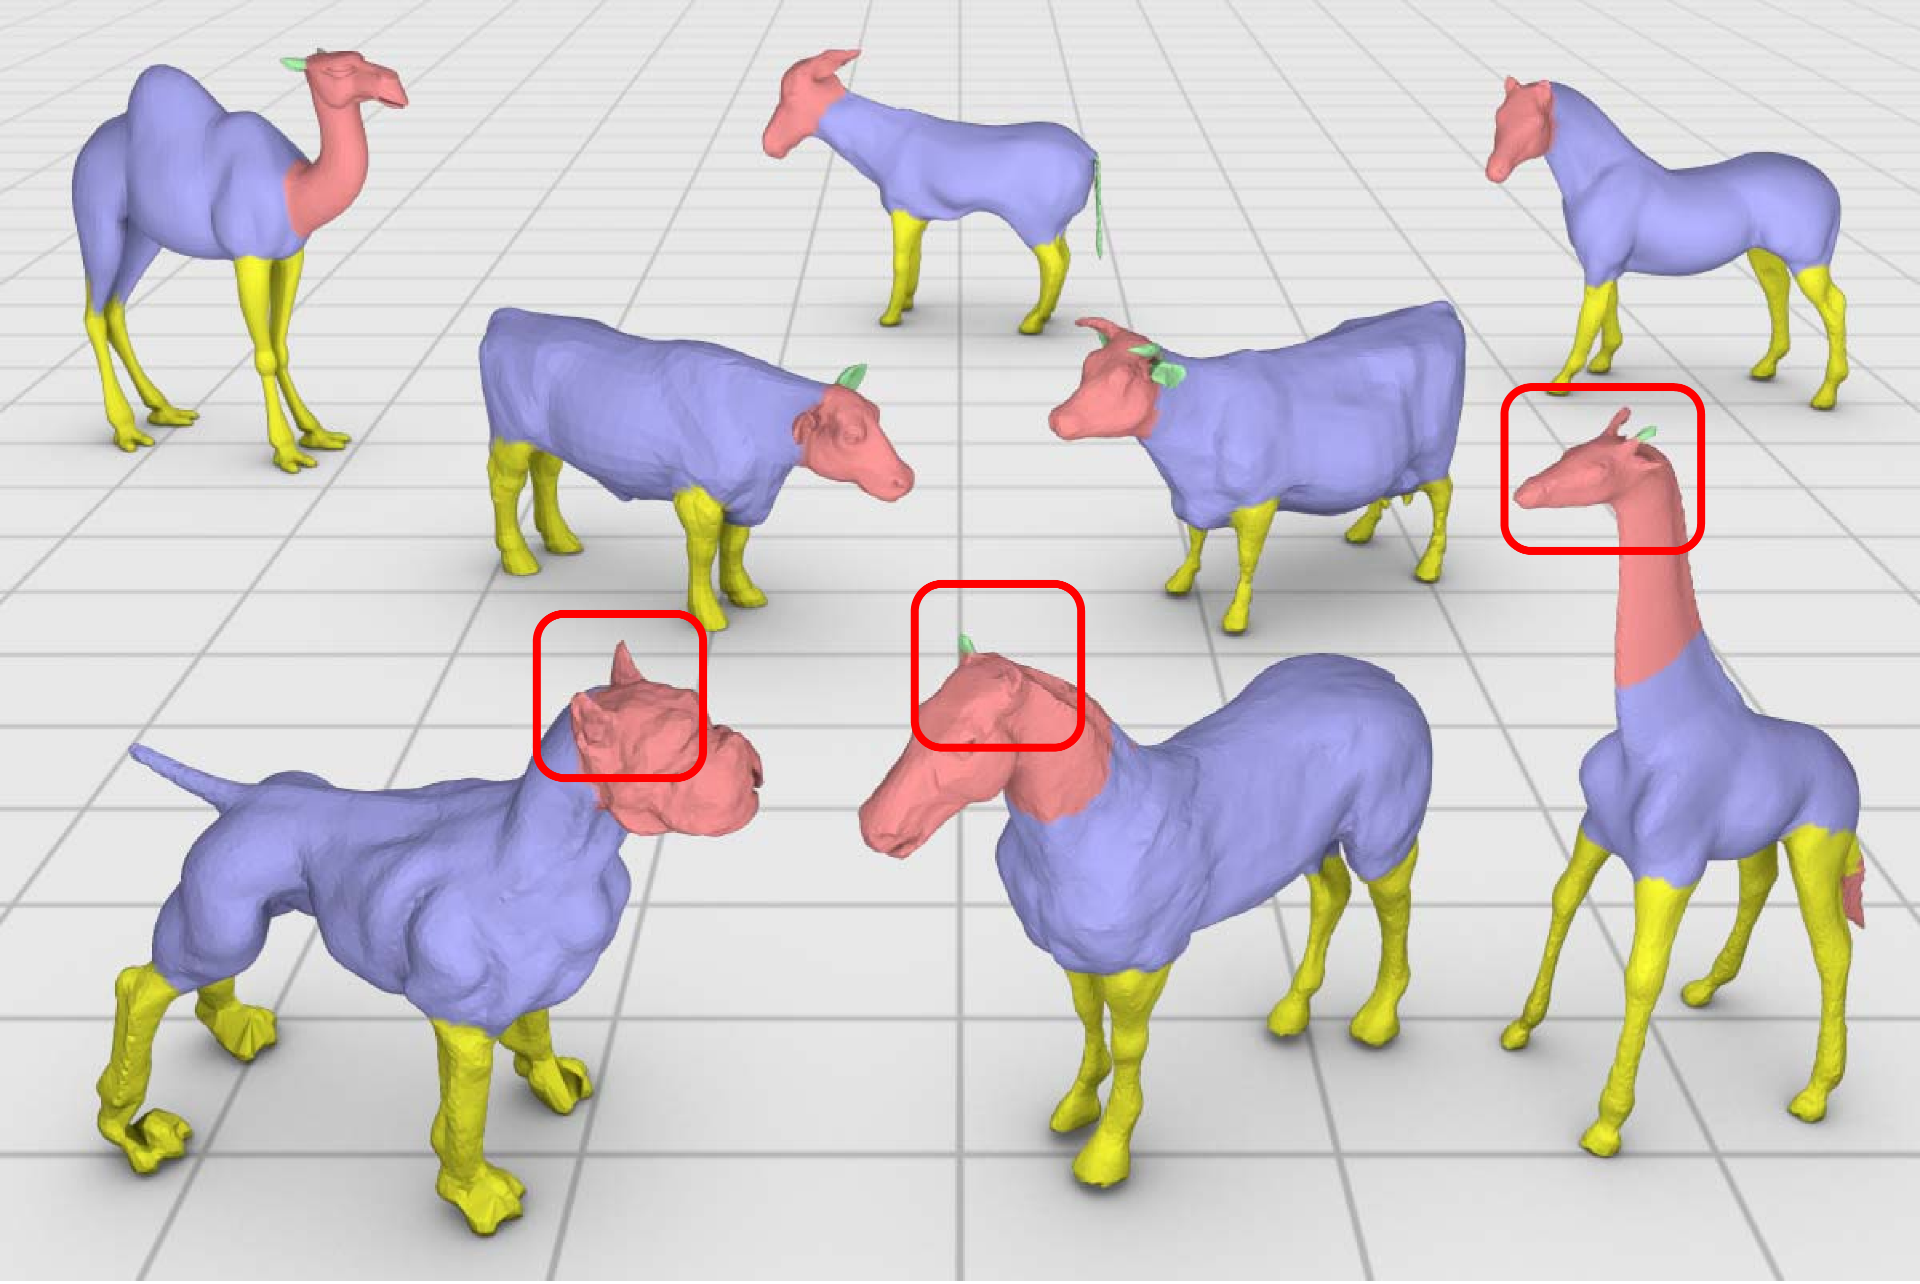

Limitations: lack of third parties

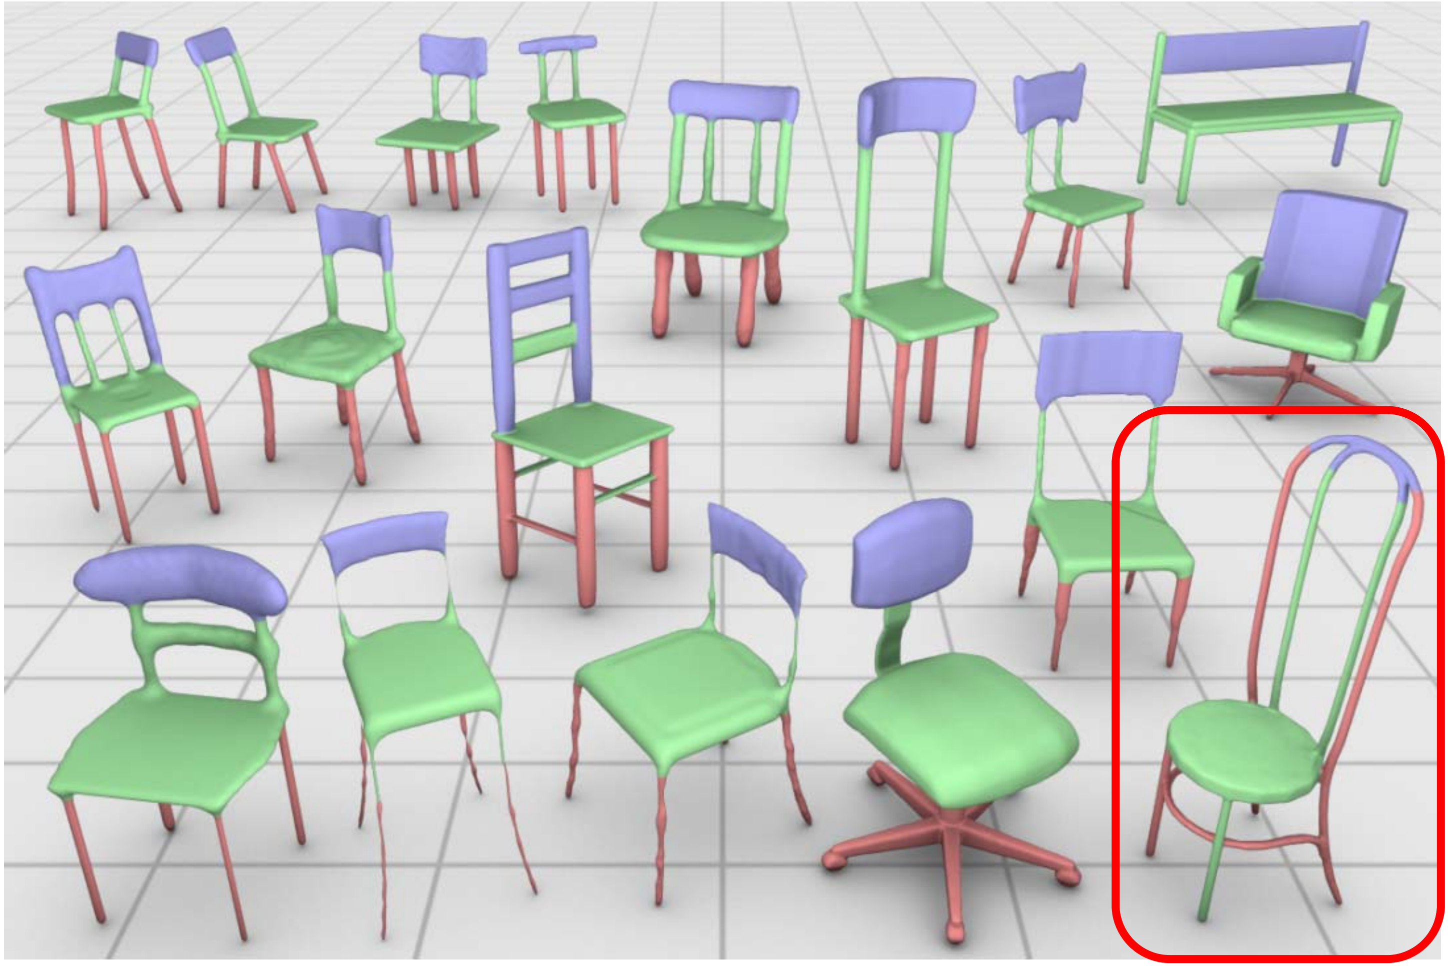

Limitations: incorrect connections

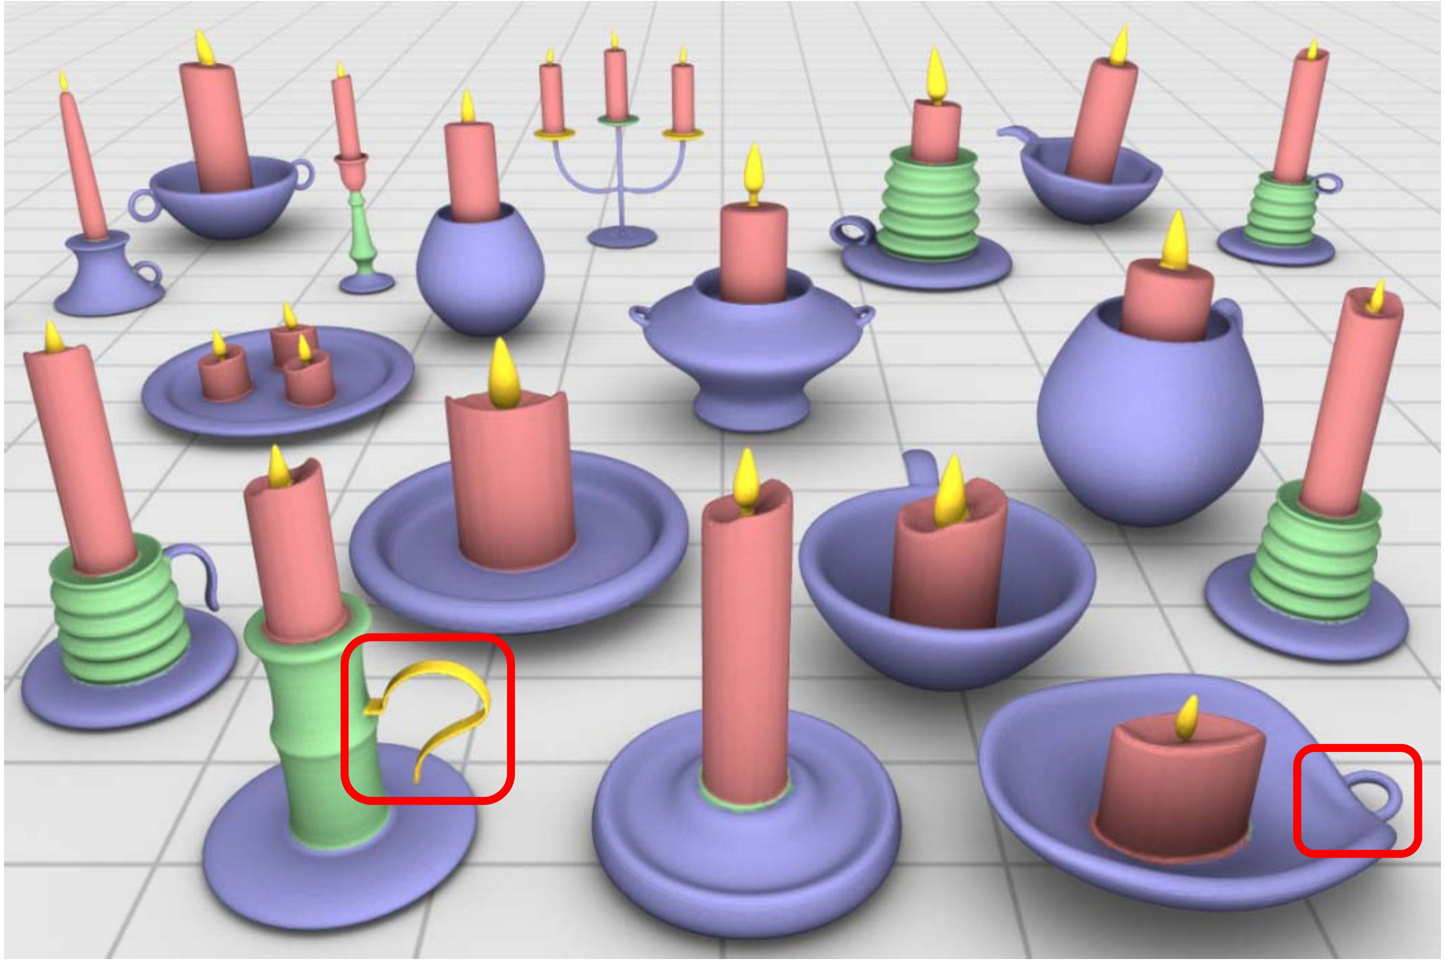

Conclusion

- Co-segmentation of a set
- Descriptor-space clustering
- Diffusion maps
- Third-party connections

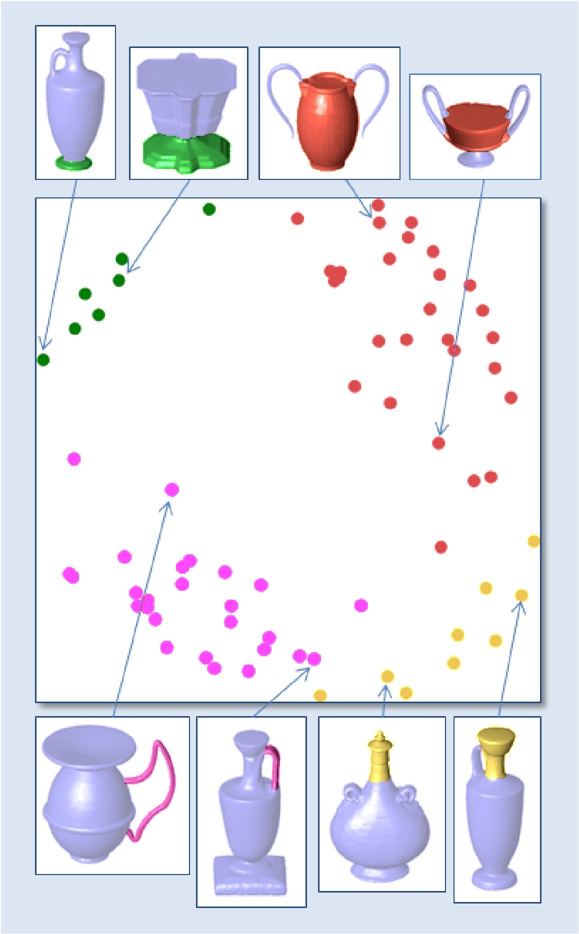

Summary

- Co-analysis and co-segmentation
  - Understanding a set of 3D shapes
- Approaches
  - Supervised
  - Unsupervised
  - Semi-supervised

- A hot research topic!

Bigger Question

- How much and what can we learn from a set?

Discussion