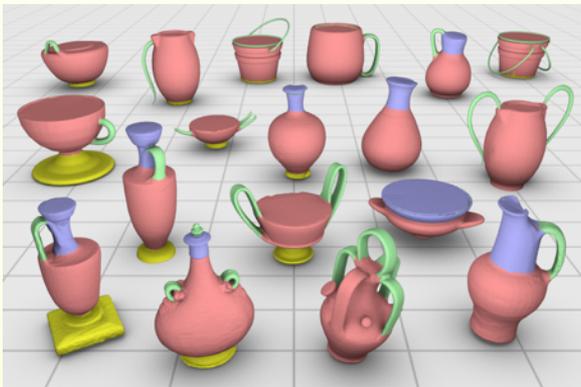
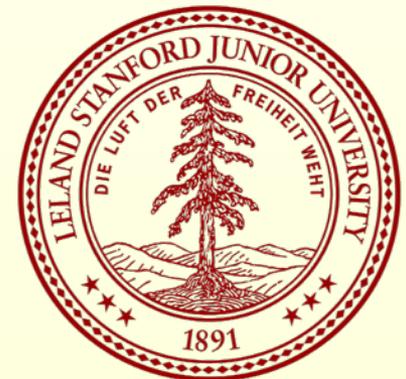
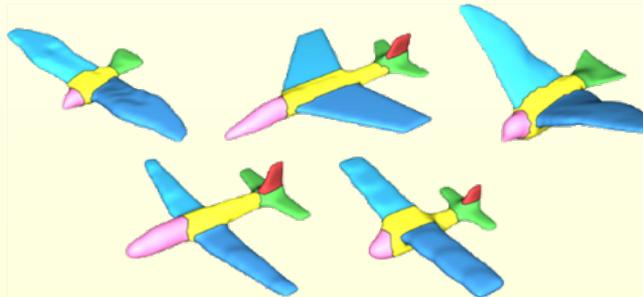


Shape segmentation and labelling



Anastasia Dubrovina
Computer Science Dept.
Stanford University



Motivation - abundance of 3D data



Shape Modeling



Mechanical CAD



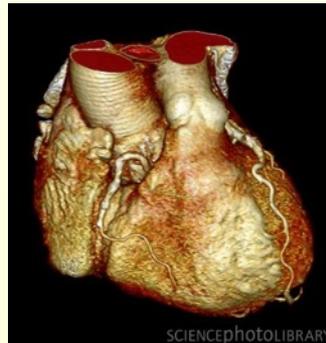
ShapeNet



Buildings



Cultural Heritage

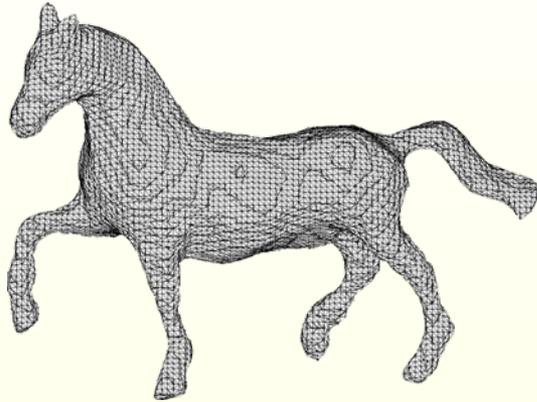


Medicine



Molecular Biology

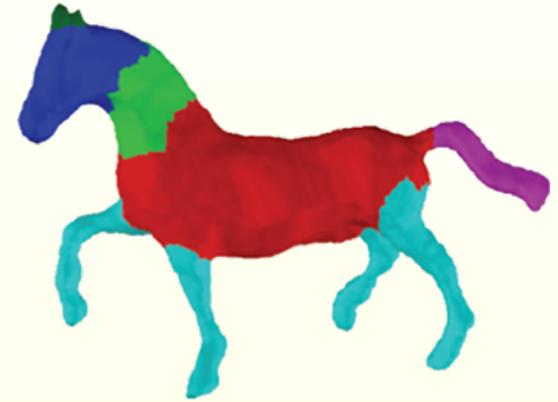
Shape segmentation / labeling



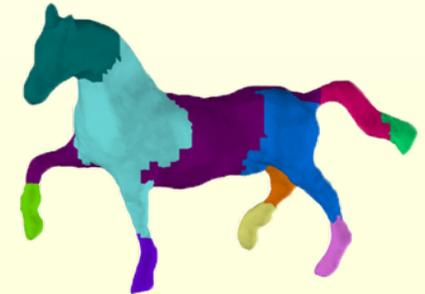
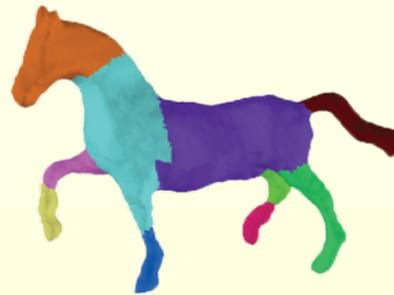
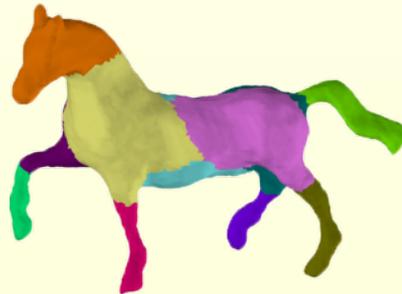
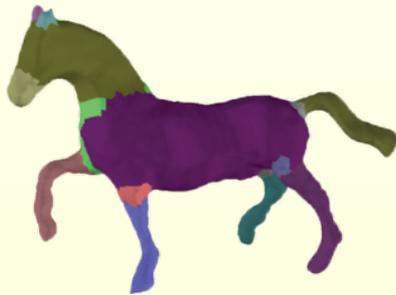
Input shape



**Segment shape into
(semantically)
meaningful part**

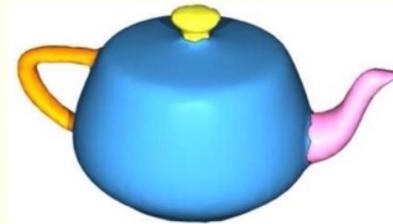


Segmented shape



Importance of shape segmentation

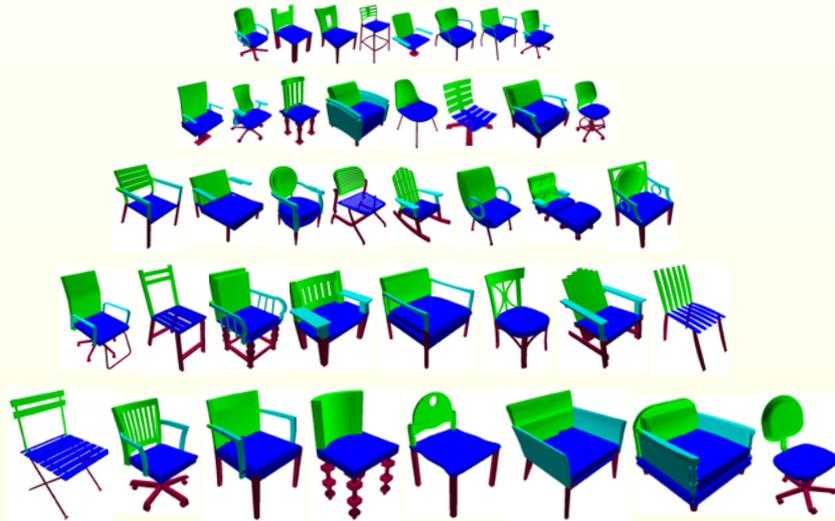
"How can we decompose a 3D model into parts?"



Psychological research indicates that recognition and shape understanding are based on structural decomposition of the shape into smaller parts [Hoffmann et al. 84,97]

Applications in other shape analysis tasks such as shape matching and shape recognition

Applications



Modeling by example
Funkhouser et al., 2004

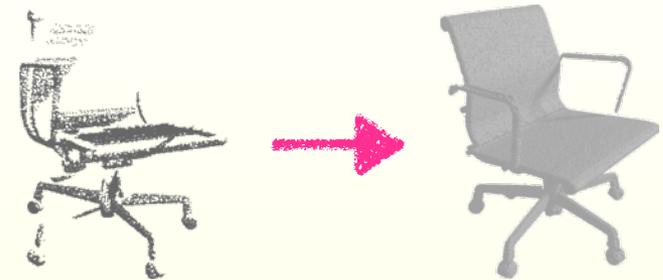
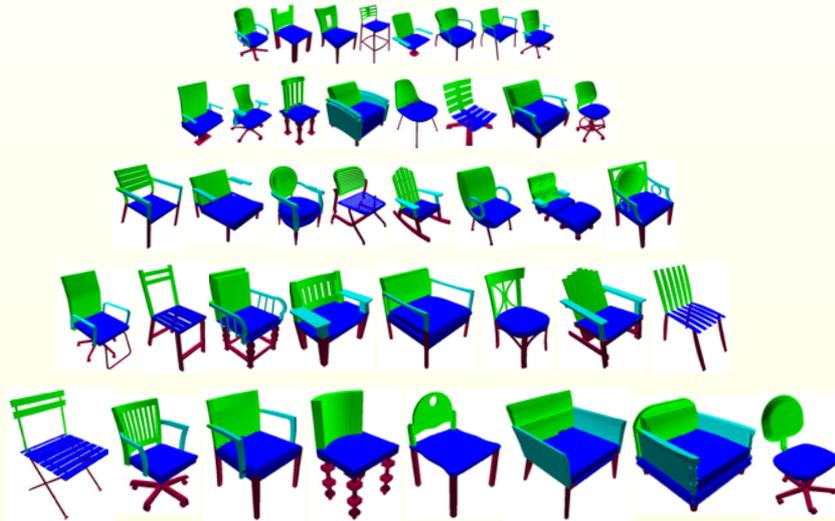


Prior knowledge for part correspondence
van Kaick et al., 2011



Part Based Shape Synthesis
Kalogerakis et al. 2012

Applications



Training Data

Data-driven shape completion
Sung et al., 2015

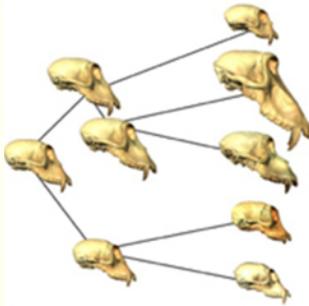


Learning to segment new synthetic
and scanned shapes
Kalogerakis et al. 2017



Image Based Modeling
Huang et al. 2015

More applications



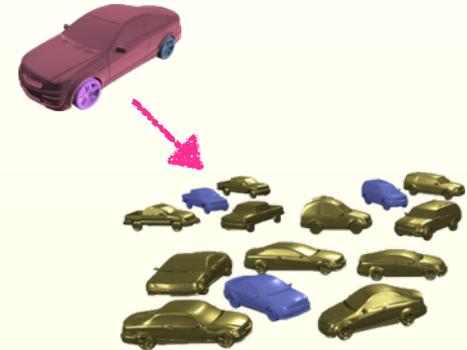
[Wiley et al.05]

Paleontology



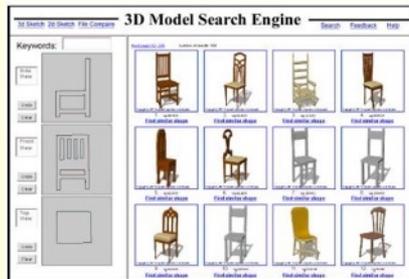
[Cooper et al.10]

Protein folding



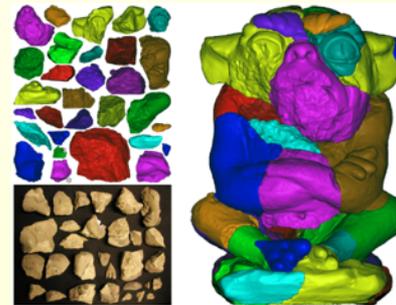
[Yumer and Kara 2014]

Learning custom deformation handles



[Funkhouser et al.05]

Product search



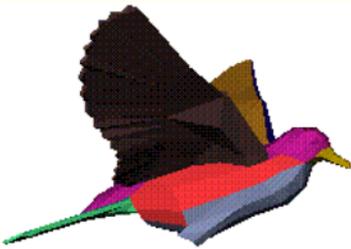
[Huang et al.06]

Solving puzzles

Lecture outline

- Shape segmentation - problem statement 
- Approaches
 - Overview
 - Supervised
 - Learning using co-segmented shape datasets
 - Unsupervised
 - Clustering in descriptor space
 - Semi-supervised
 - Active learning with human in the loop
- Datasets

Single-shape segmentation: geometry-driven



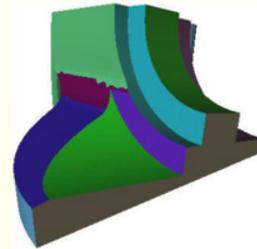
[Shalfman et al. 2002]

K-Means



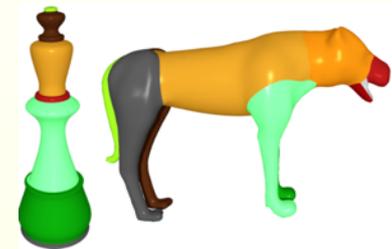
[Katz et al. 05]

Core Extraction



[Attene et. al 2006]

Fitting Primitives



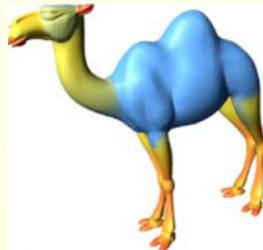
[Lai et al. 08]

Random Walks



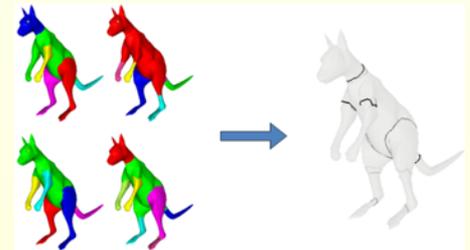
[Golovinskiy and Funkhouser 08]

Normalized Cuts



[Shapira et al. 08]

Shape Diameter Function

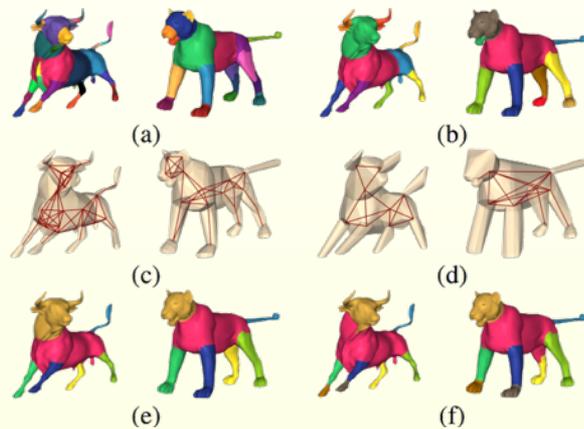


[Golovinskiy and Funkhouser 08]

Randomized Cuts

Consistent segmentation in shape collections

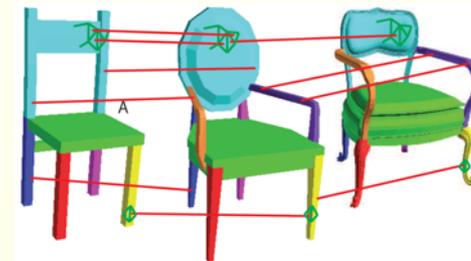
- Can more knowledge be inferred by analyzing a set of shapes?



Kreavoy et al. 2008

Consistent segmentations for
model composition

- Explore geometric cues
- Require shapes in correspondence



[Golovinskiy and Funkhouser 09]

Consistent segmentations



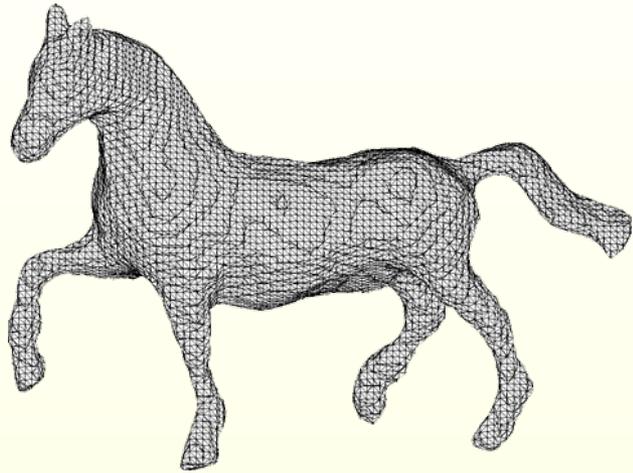
[Xu et al. 10]

Co-segmentation

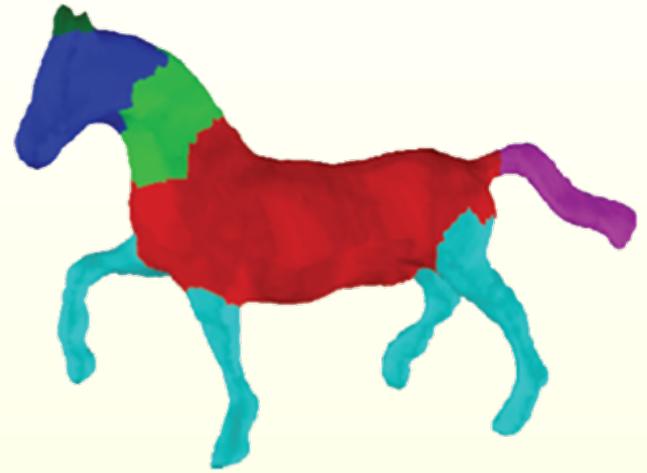
Supervised segmentation and labelling

SUPERVISED SHAPE SEGMENTATION

Supervised shape segmentation

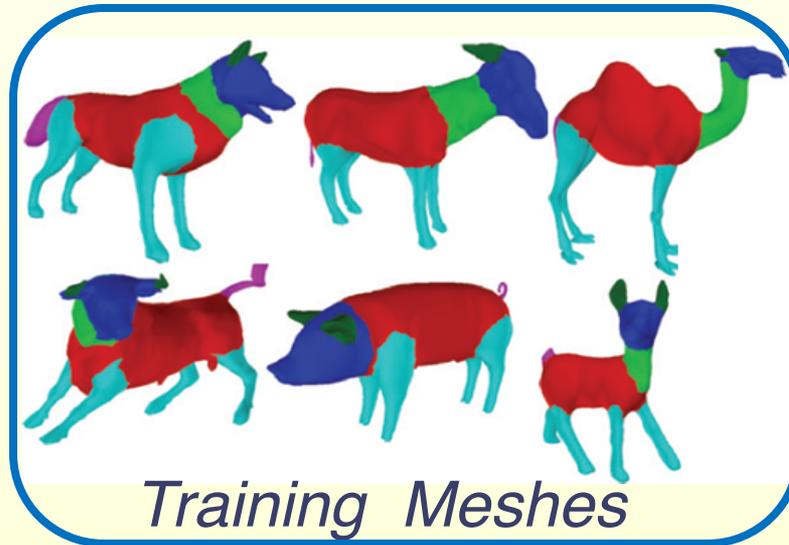


Input shape



Labeled shape

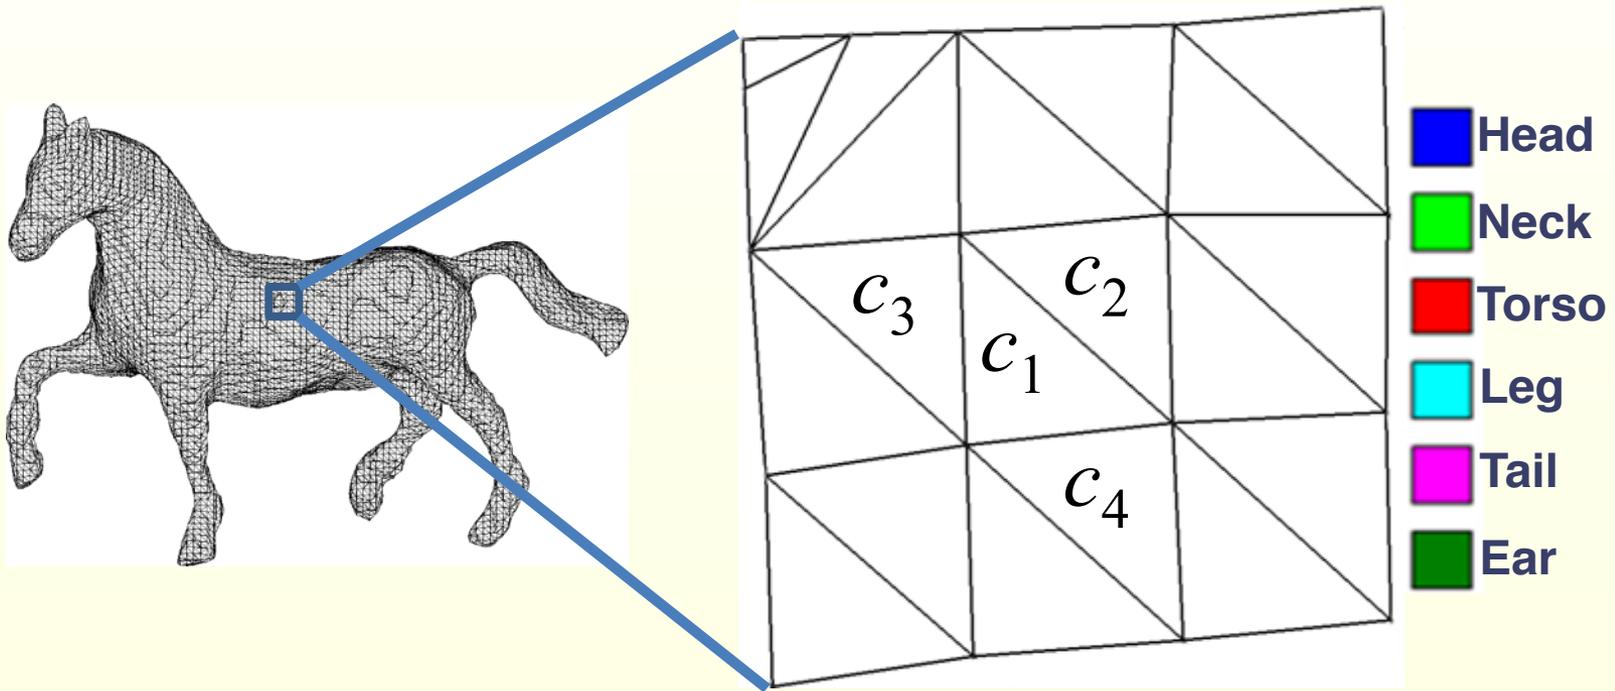
Learn from examples



Training Meshes

- Head
- Neck
- Torso
- Leg
- Tail
- Ear

Labeling problem statement

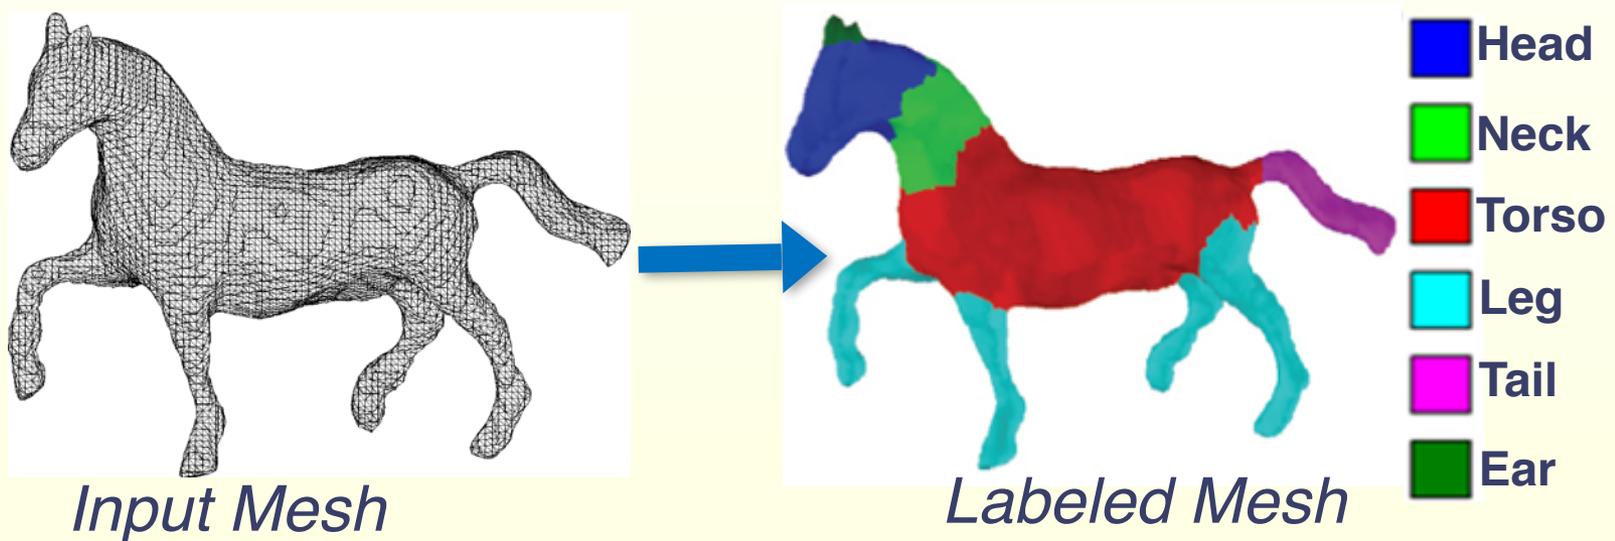


$$c_1, c_2, c_3 \in C$$

$$C = \{ \textit{head}, \textit{neck}, \textit{torso}, \textit{leg}, \textit{tail}, \textit{ear} \}$$

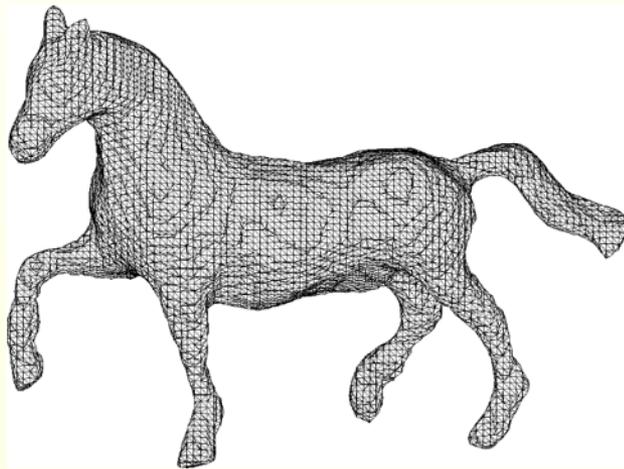
Conditional Random Field for Labeling

- “Learning 3D Mesh Segmentation and Labeling”,
[Kalogerakis et al. 2010]



$$c^* = \arg \min_c \left\{ \sum_t \alpha_i \underbrace{E_1(c_i; \mathbf{x}_i)}_{\text{Unary term}} + \sum_{t,j} l_{ij} E_2(c_i, c_j; \mathbf{y}_{ij}) \right\}$$

Conditional Random Field for Labeling



Input Mesh



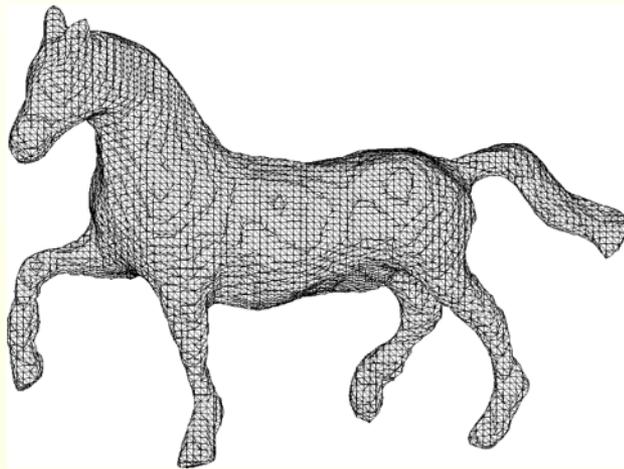
Labeled Mesh

-  Head
-  Neck
-  Torso
-  Leg
-  Tail
-  Ear

$$c^* = \arg \min_c \left\{ \sum_i \alpha_i E_1(c_i; \mathbf{x}_i) + \sum_{i,j} l_{ij} E_2(c_i, c_j; \mathbf{y}_{ij}) \right\}$$

*Face
features*

Conditional Random Field for Labeling



Input Mesh



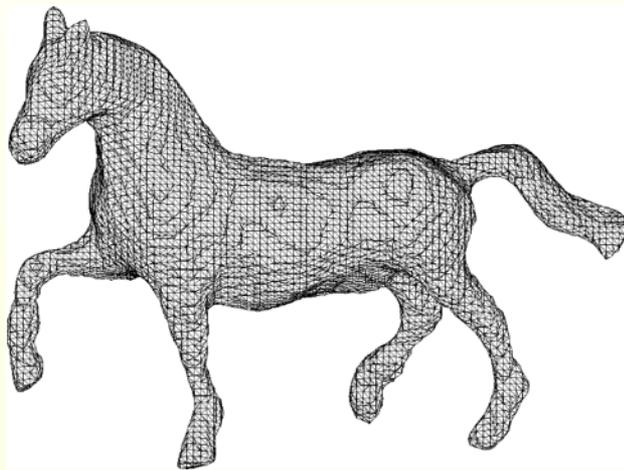
Labeled Mesh

-  Head
-  Neck
-  Torso
-  Leg
-  Tail
-  Ear

$$c^* = \arg \min_c \left\{ \sum_i \alpha_i E_1(c_i; \mathbf{x}_i) + \sum_{i,j} l_{ij} E_2(c_i, c_j; \mathbf{y}_{ij}) \right\}$$

Face Area

Conditional Random Field for Labeling



Input Mesh



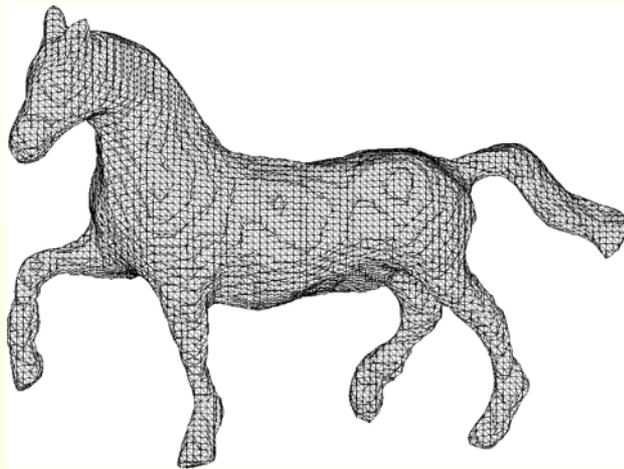
Labeled Mesh

- Head
- Neck
- Torso
- Leg
- Tail
- Ear

$$c^* = \arg \min_c \left\{ \sum_i \alpha_i E_1(c_i; \mathbf{x}_i) + \sum_{i,j} l_{ij} \boxed{E_2(c_i, c_j; \mathbf{y}_{ij})} \right\}$$

Pairwise Term

Conditional Random Field for Labeling



Input Mesh



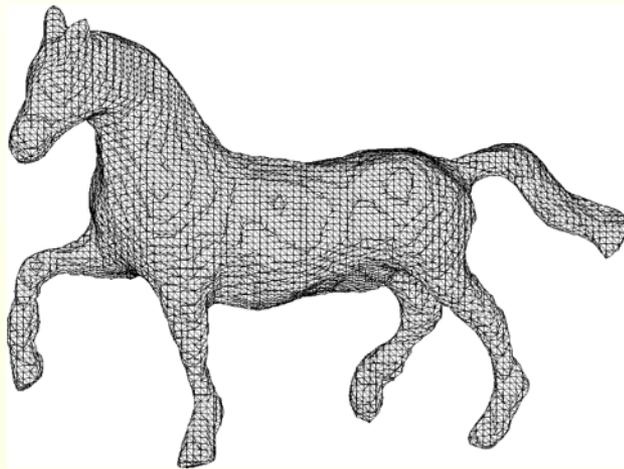
Labeled Mesh



$$c^* = \arg \min_c \left\{ \sum_i \alpha_i E_1(c_i; \mathbf{x}_i) + \sum_{i,j} l_{ij} E_2(c_i, c_j; \mathbf{y}_{ij}) \right\}$$

*Edge
Features*

Conditional Random Field for Labeling



Input Mesh



Labeled Mesh

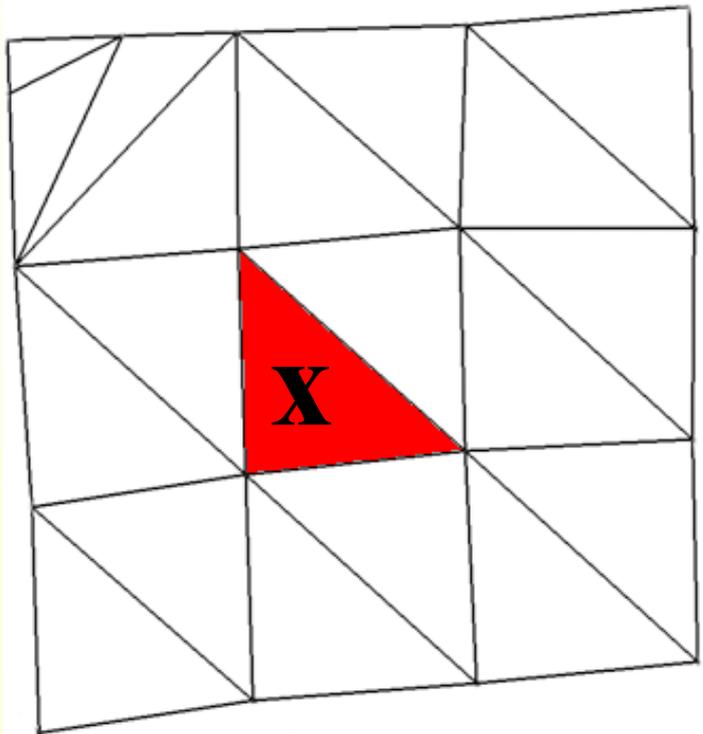
-  Head
-  Neck
-  Torso
-  Leg
-  Tail
-  Ear

$$c^* = \arg \min_c \left\{ \sum_i \alpha_i E_1(c_i; \mathbf{x}_i) + \sum_{i,j} \boxed{l_{ij}} E_2(c_i, c_j; \mathbf{y}_{ij}) \right\}$$

**Edge
Length**

Feature vector

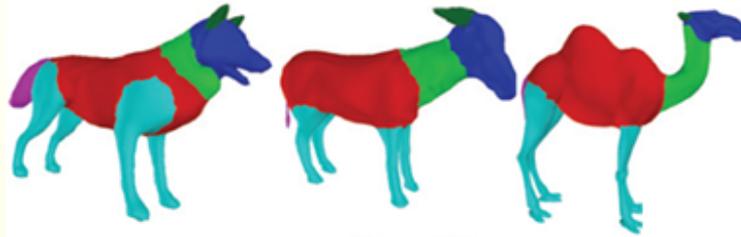
$$\mathbf{x} \in \mathfrak{R}^{375+35|C|} \rightarrow P(c | \mathbf{x})$$



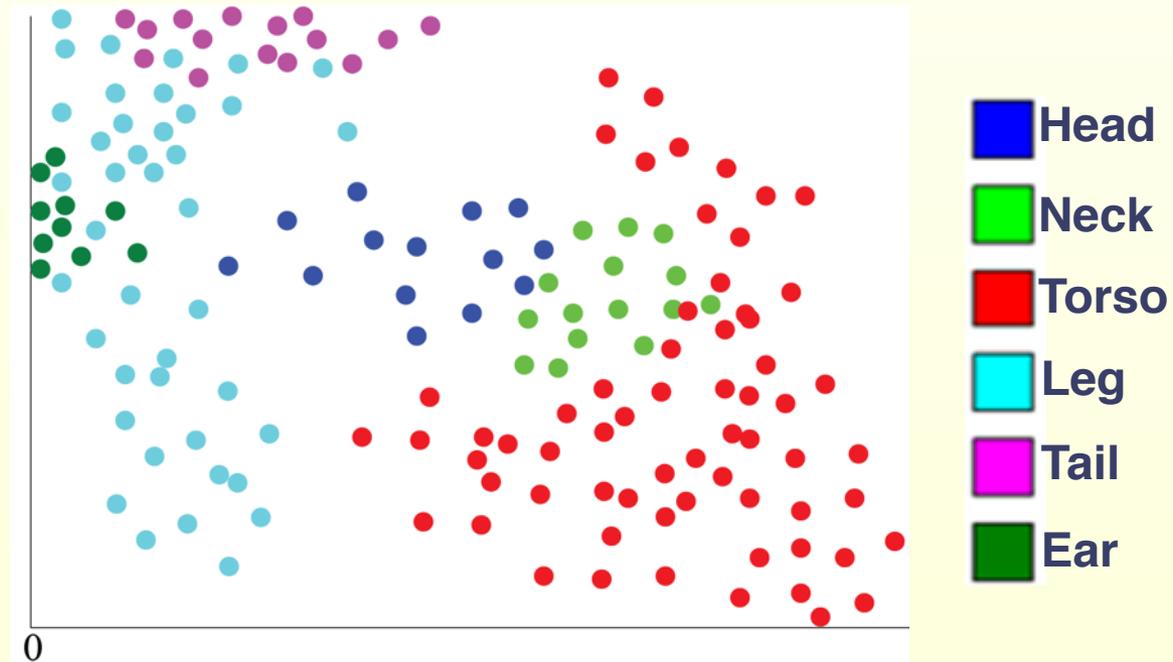
- surface curvature
- singular values from PCA
- shape diameter
- distances from medial surface
- average geodesic distances
- shape contexts
- spin images
- contextual label features

Using more features helps

Learning a classifier

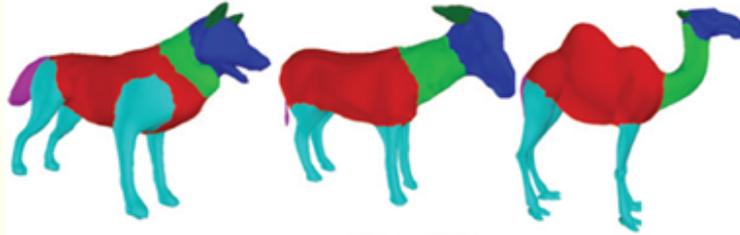


$\rightarrow \{(\mathbf{x}_i, c_i)\}$

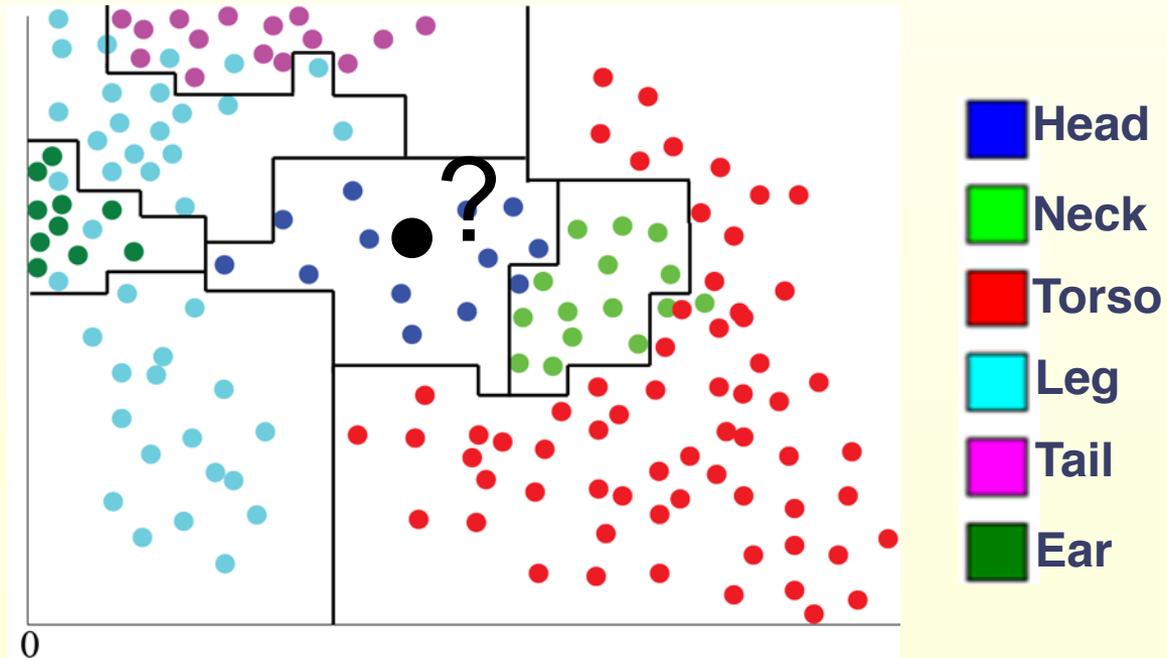


Learning a classifier

We use



Classifier [Torralba et al. 2007]
 $\rightarrow \{(\mathbf{x}_i, c_i)\}$



Joint-Boost classifier [Torralba et al. 2007]

- Properties
 - Automatic feature selection

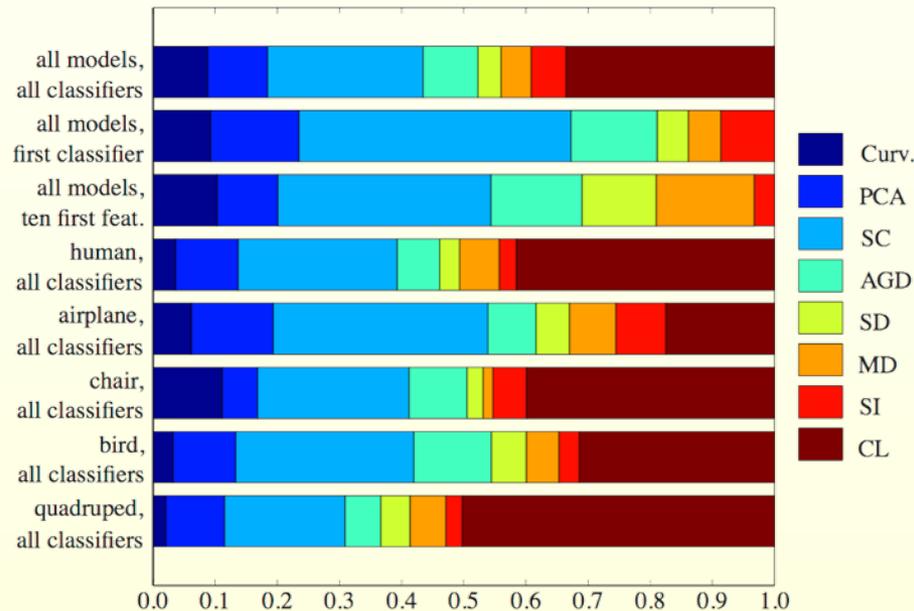
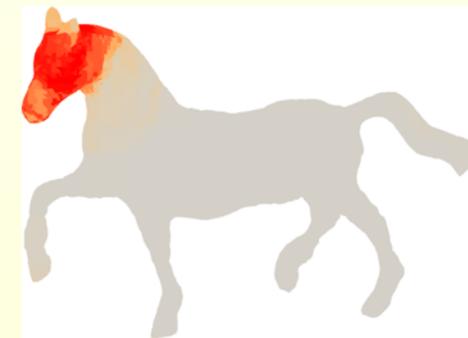
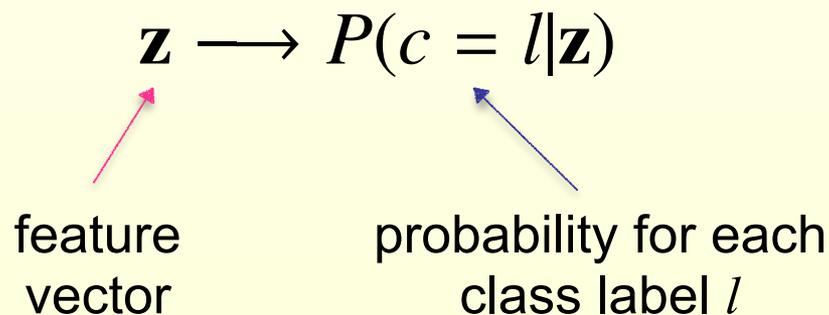


Figure 5: Percentages of features used by JointBoost for different cases. See text for details. Legend: Curv.=curvature, PCA=PCA singular values, SC=shape contexts, AGD=average geodesic distances, SD=shape diameter, MD=distance from medial surface, SI = Spin Images, CL = contextual label features.

Joint-Boost classifier [Torralba et al. 2007]

- Properties
 - Automatic feature selection
 - Can handle large number of input features
 - Fast sequential learning algorithm
 - Shares features among classes
 - Produces output probabilities



$P(\text{head} | \mathbf{x})$

Joint-Boost classifier [Torralba et al. 2007]

- Algorithm outline
 - Classifier composed of *decision stumps*

$$h(\mathbf{z}, l; \phi) = \begin{cases} a & z_f > \tau \text{ and } l \in \mathcal{C}_S \\ b & z_f \leq \tau \text{ and } l \in \mathcal{C}_S \\ k_l & l \notin \mathcal{C}_S \end{cases}$$

← Set of classes

Algorithm parameters ϕ : f, a, b, τ , the set \mathcal{C}_S , and k_l

Joint-Boost classifier [Torralba et al. 2007]

- Algorithm outline
 - Classifier composed of *decision stumps*

$$h(\mathbf{z}, l; \phi) = \begin{cases} a & z_f > \tau \text{ and } l \in \mathcal{C}_S \\ b & z_f \leq \tau \text{ and } l \in \mathcal{C}_S \\ k_l & l \notin \mathcal{C}_S \end{cases}$$

- Scores class labels by thresholding f -th entry of a feature vector \mathbf{z}
- Probability of a given class l

$$H(\mathbf{z}, l) = \sum_j h(\mathbf{z}, l; \phi_j)$$

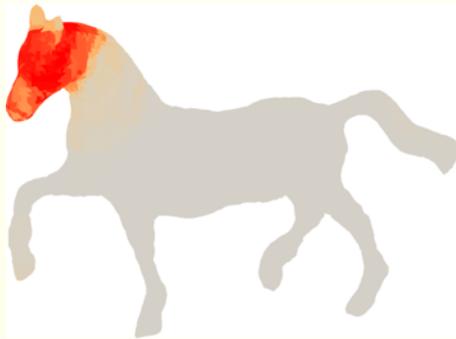
All parameters

$$P(c = l | \mathbf{z}, \xi) = \frac{\exp(H(\mathbf{z}, l))}{\sum_{\ell \in \mathcal{C}} \exp(H(\mathbf{z}, \ell))}$$

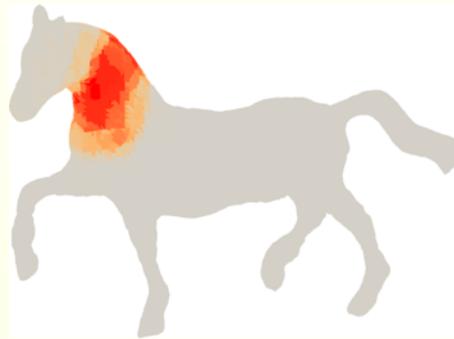
Soft-max

Labeling energy: unary term

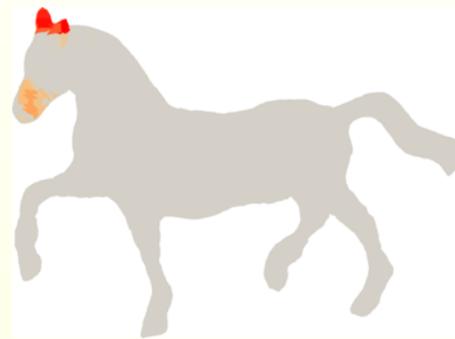
$$E_1(c; \mathbf{x}) = -\log P(c | \mathbf{x})$$



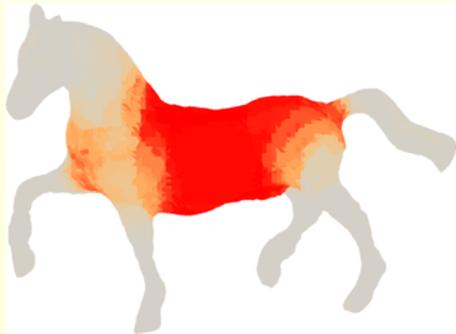
$P(\textit{head} | \mathbf{x})$



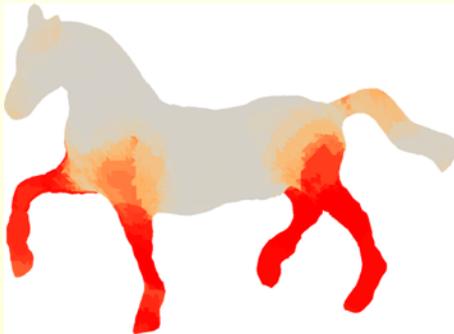
$P(\textit{neck} | \mathbf{x})$



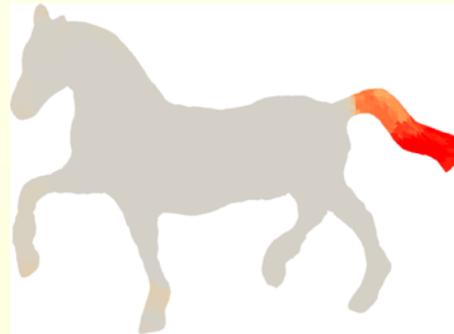
$P(\textit{ear} | \mathbf{x})$



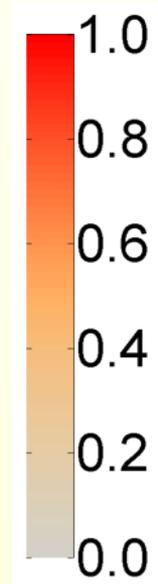
$P(\textit{torso} | \mathbf{x})$



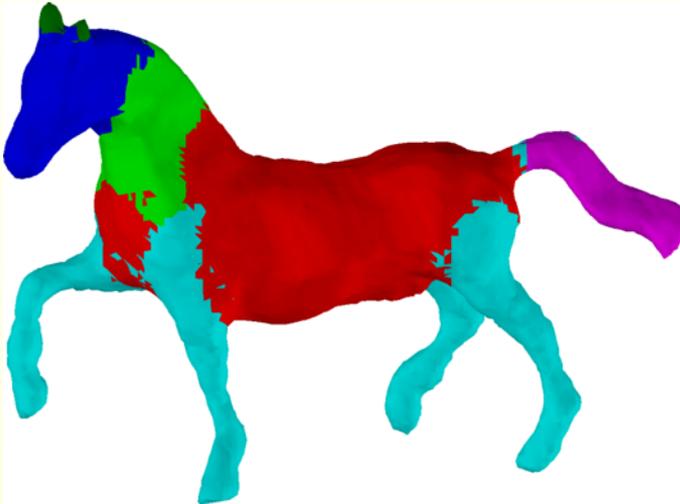
$P(\textit{leg} | \mathbf{x})$



$P(\textit{tail} | \mathbf{x})$



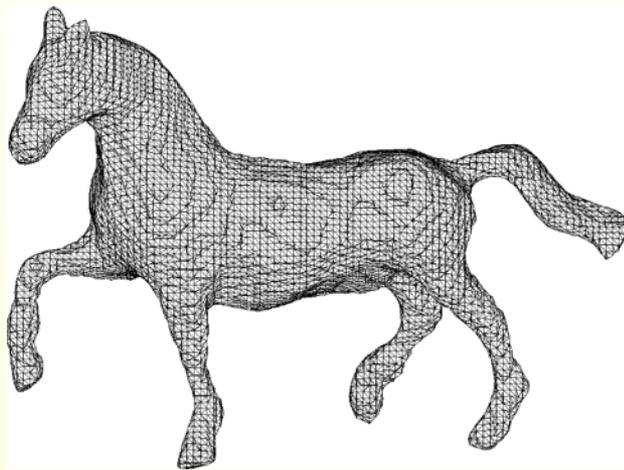
Labeling using unary term alone



$$\arg \max_c P(c | \mathbf{x})$$

Most-likely labels

Conditional random field labeling



Input Mesh



Labeled Mesh

-  Head
-  Neck
-  Torso
-  Leg
-  Tail
-  Ear

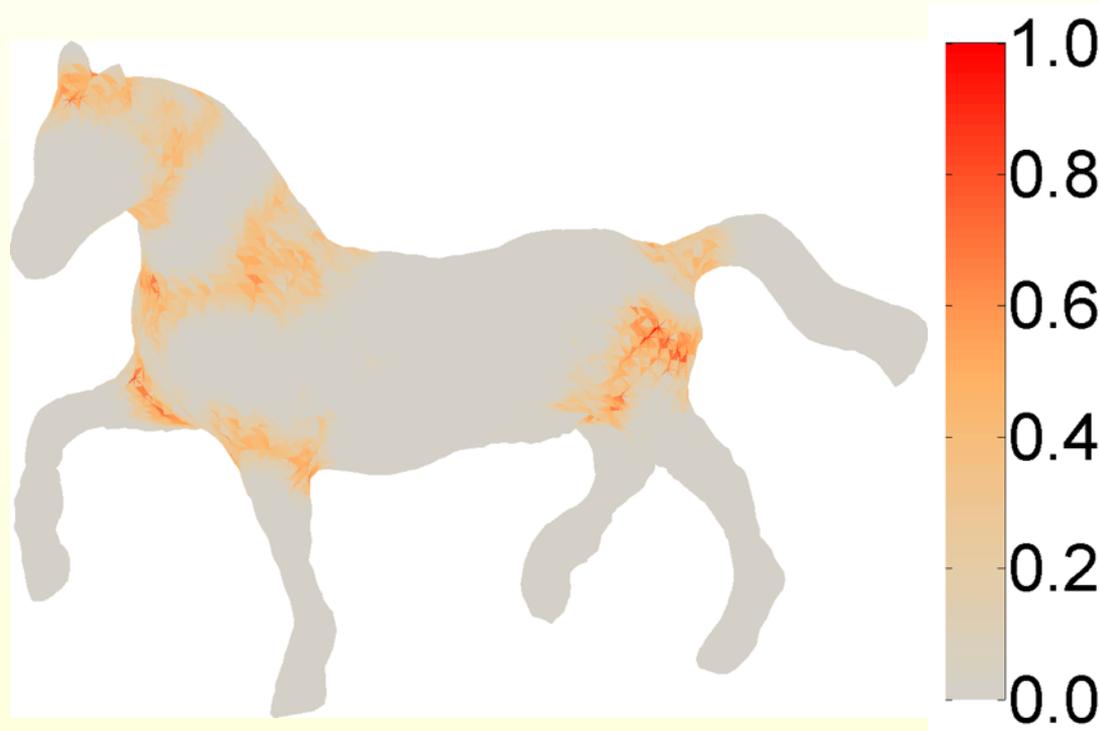
$$c^* = \arg \min_c \left\{ \sum_i \alpha_i E_1(c_i; \mathbf{x}_i) + \sum_{i,j} l_{ij} \boxed{E_2(c_i, c_j; \mathbf{y}_{ij})} \right\}$$

Pairwise Term

Pairwise Term

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

Geometry-dependent term: likelihood of label change based on geometry alone



Pairwise Term

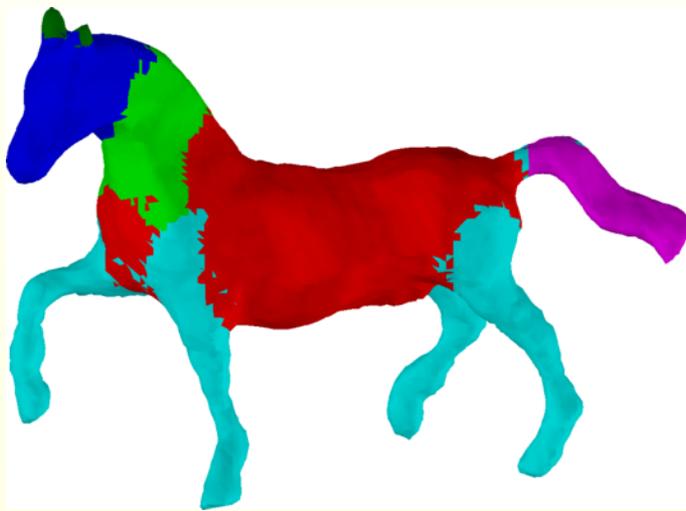
$$E_2(c, c'; \mathbf{y}, \theta_2) = G(\mathbf{y}) L(c, c')$$

Learned label compatibility term: learned with JointBoost from geometric 191-D features

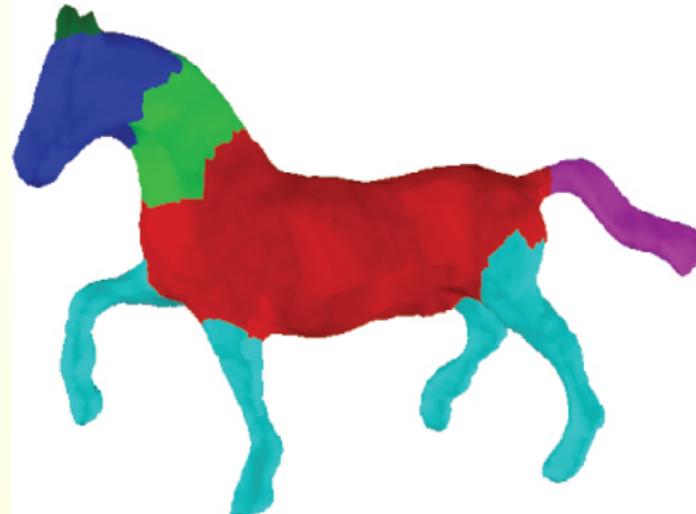
$L(c, c') =$

	Head	Neck	Ear	Torso	Leg	Tail	
	0	.45	.07	1	∞	∞	Head
	.45	0	∞	1	∞	∞	Neck
	.07	∞	0	∞	∞	∞	Ear
	1	1	∞	0	1	.56	Torso
	∞	∞	∞	1	0	∞	Leg
	∞	∞	∞	.56	∞	0	Tail

Full CRF result



Unary term classifier



Full CRF result

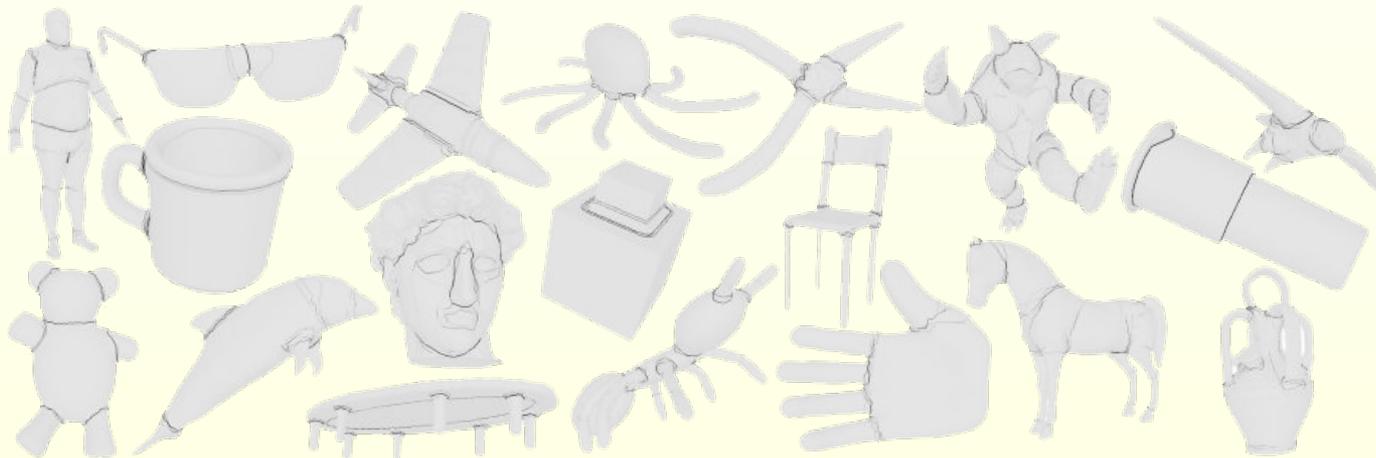
- Head
- Neck
- Torso
- Leg
- Tail
- Ear

Learning and inference

- Learn unary classifier and $G(y)$ with Joint Boosting [Torralba et al. 2007]
- Hold-out validation for the rest of parameters
- CRF inference - using graph cuts [Boykov et al. 2001]

Dataset used in experiments

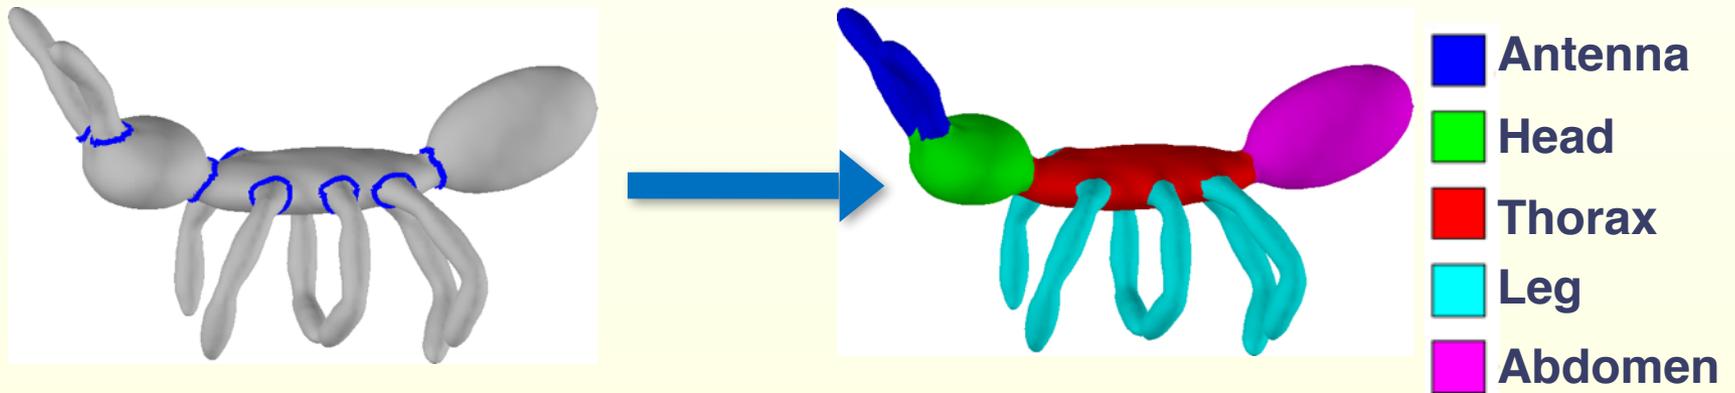
- Princeton Segmentation Benchmark (PSB) [Chen et al. 2009]
<http://segeval.cs.princeton.edu>



- 11 human generated segmentations (“probabilistic” ground truth)
 - 380 watertight meshes across 19 categories

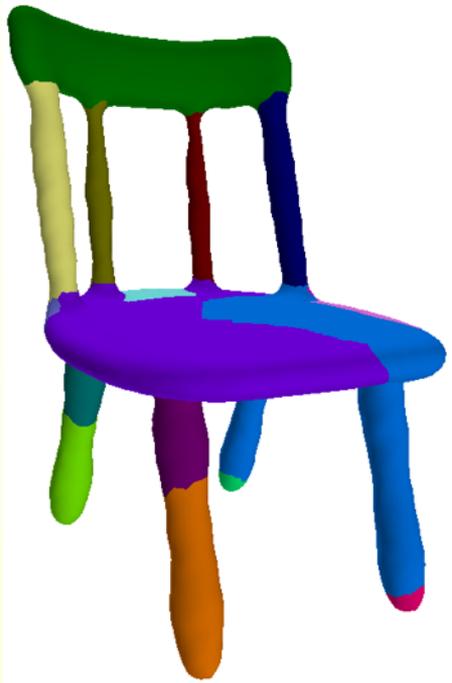
Dataset used in experiments

- We label 380 meshes from the PCB [Chen *et al.* 2009]

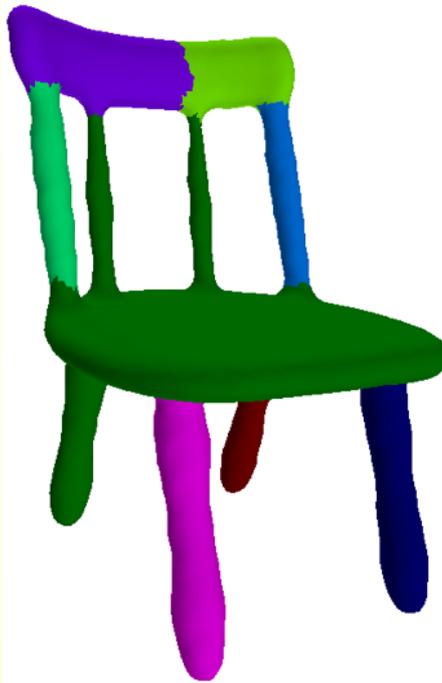


- Each of the 19 categories is treated separately

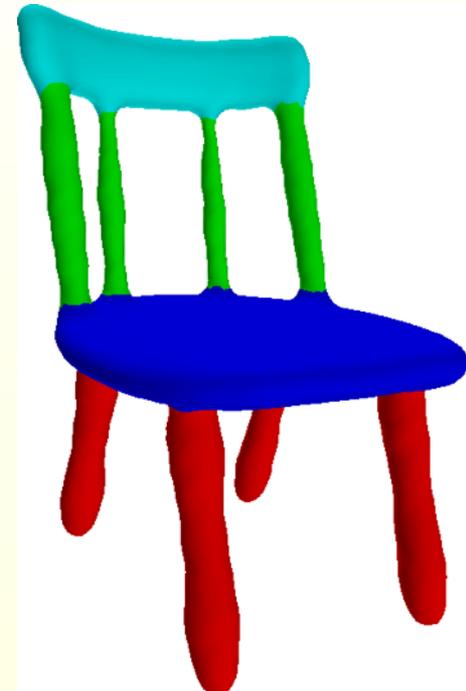
Segmentation Comparisons



Shape Diameter
[Shapira et al. 10]



Randomized Cuts
[Golovinskiy and Funkhouser 08]



Our approach

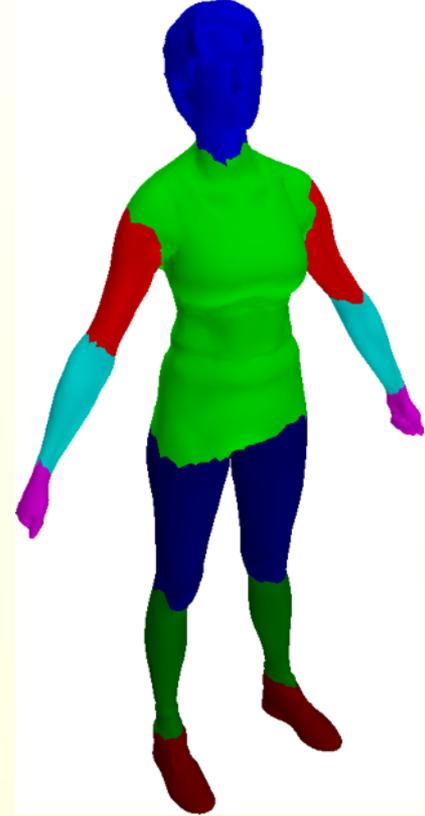
Segmentation Comparisons



Shape Diameter
[Shapira et al. 10]

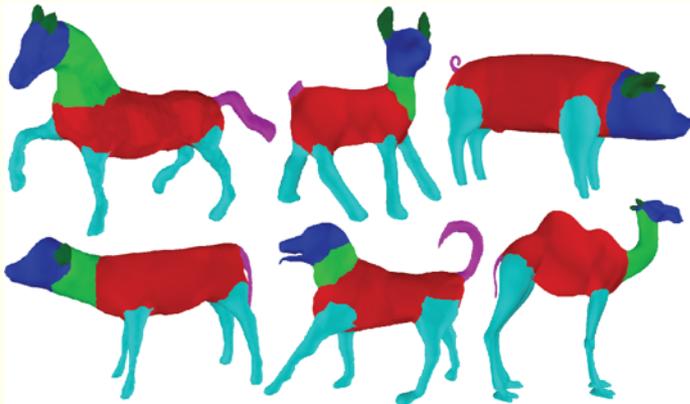


Randomized Cuts
[Golovinskiy and Funkhouser 08]



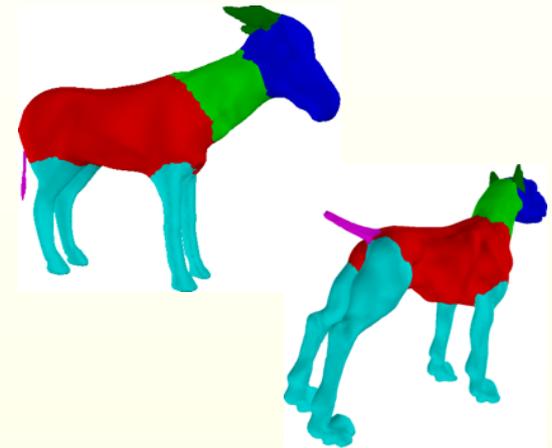
Our approach

Learning different segmentation styles

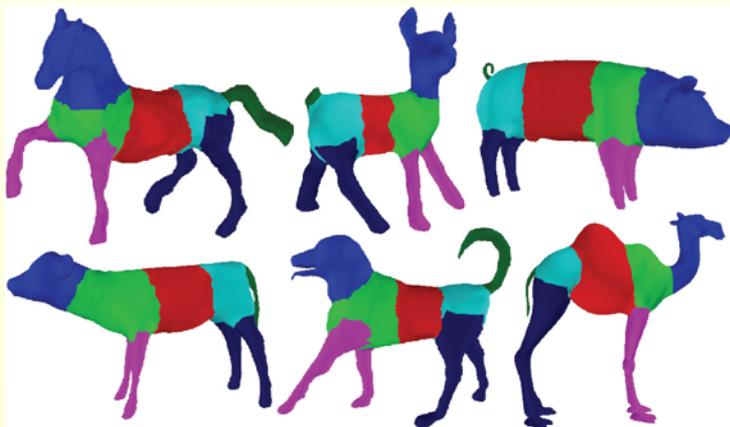


Training Meshes

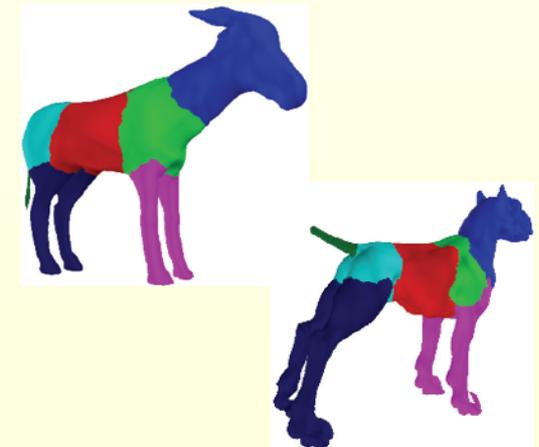
- Head
- Neck
- Torso
- Leg
- Tail
- Ear



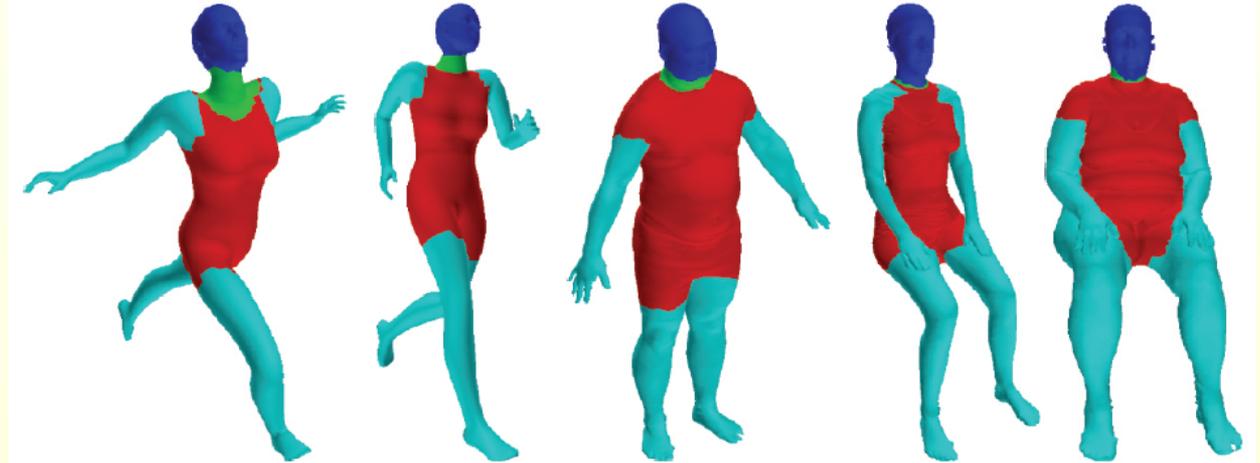
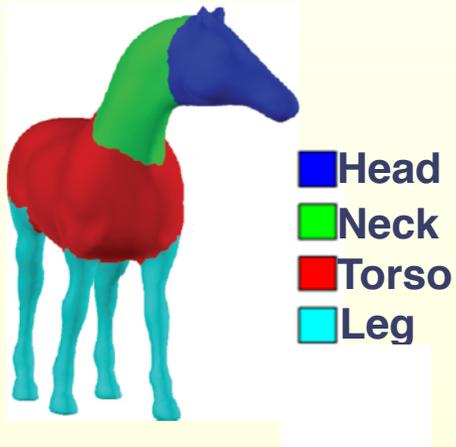
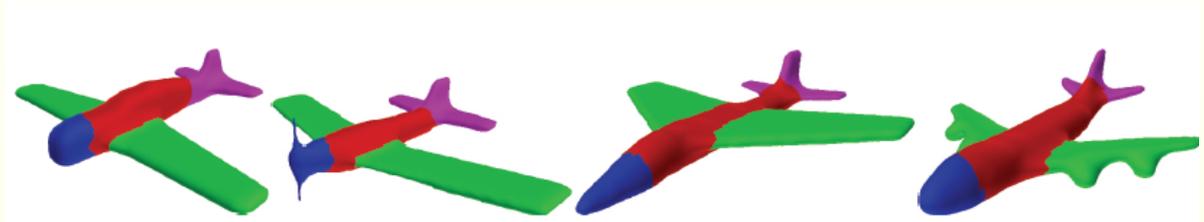
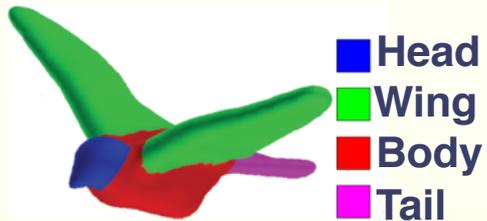
Test Meshes



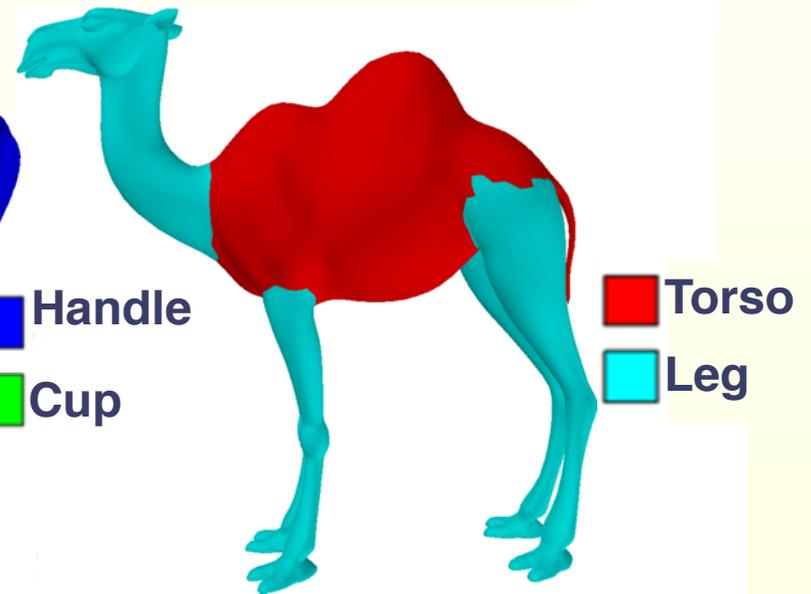
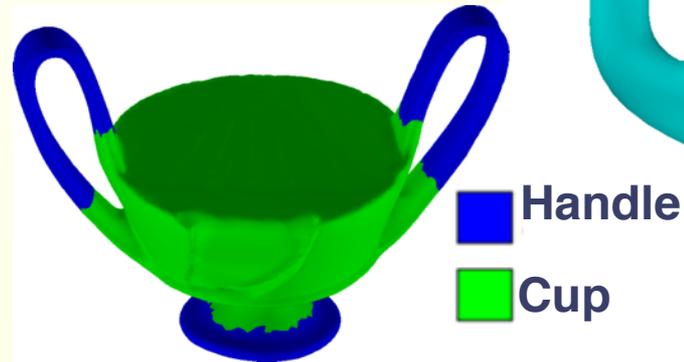
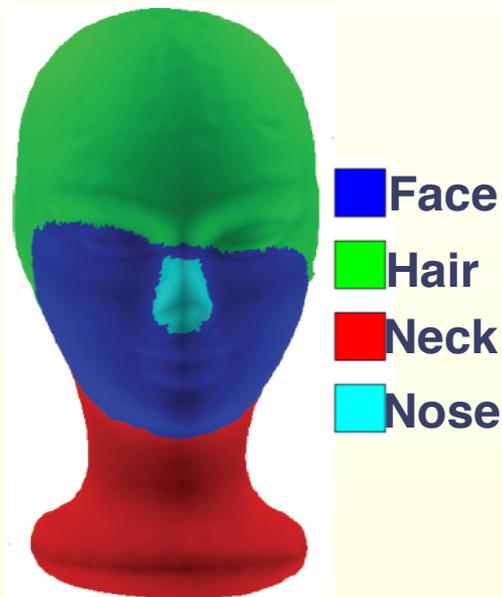
- Head
- Front Torso
- Middle Torso
- Back Torso
- Front Leg
- Back Leg
- Tail



Generalization to different categories



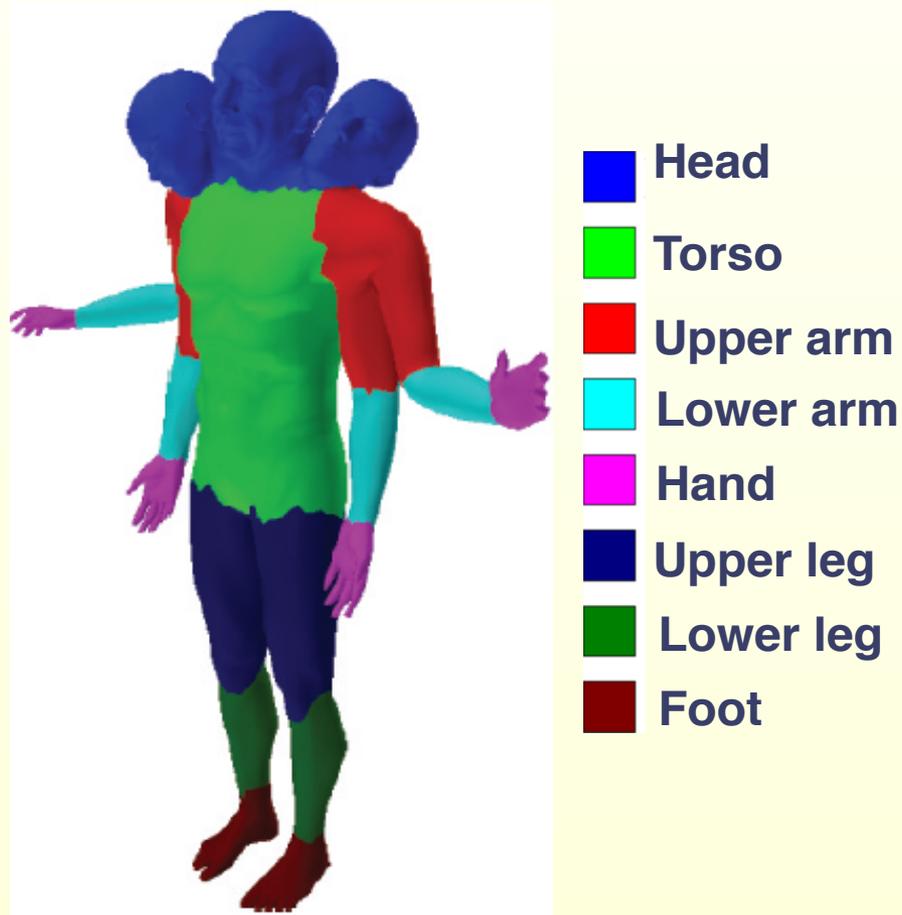
Failure cases



Limitations (common to most approaches)

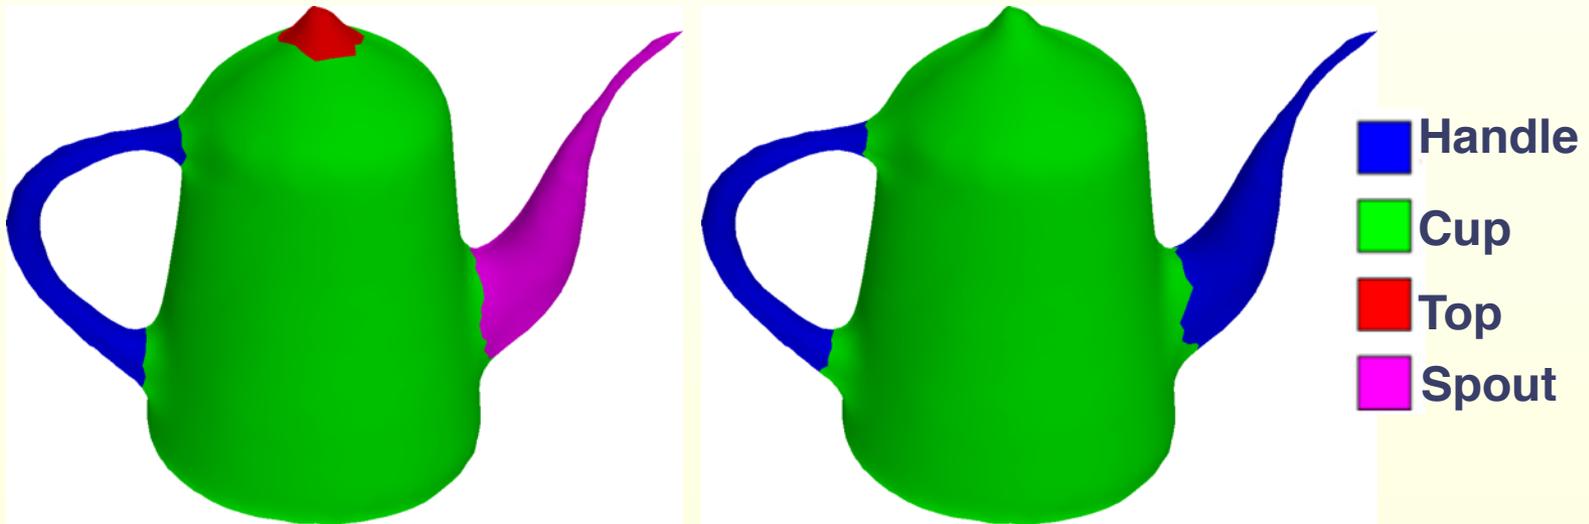
- Adjacent segments with the same label are merged

Can't do part
instance
segmentation



Limitations

- Results depend on having sufficient training data

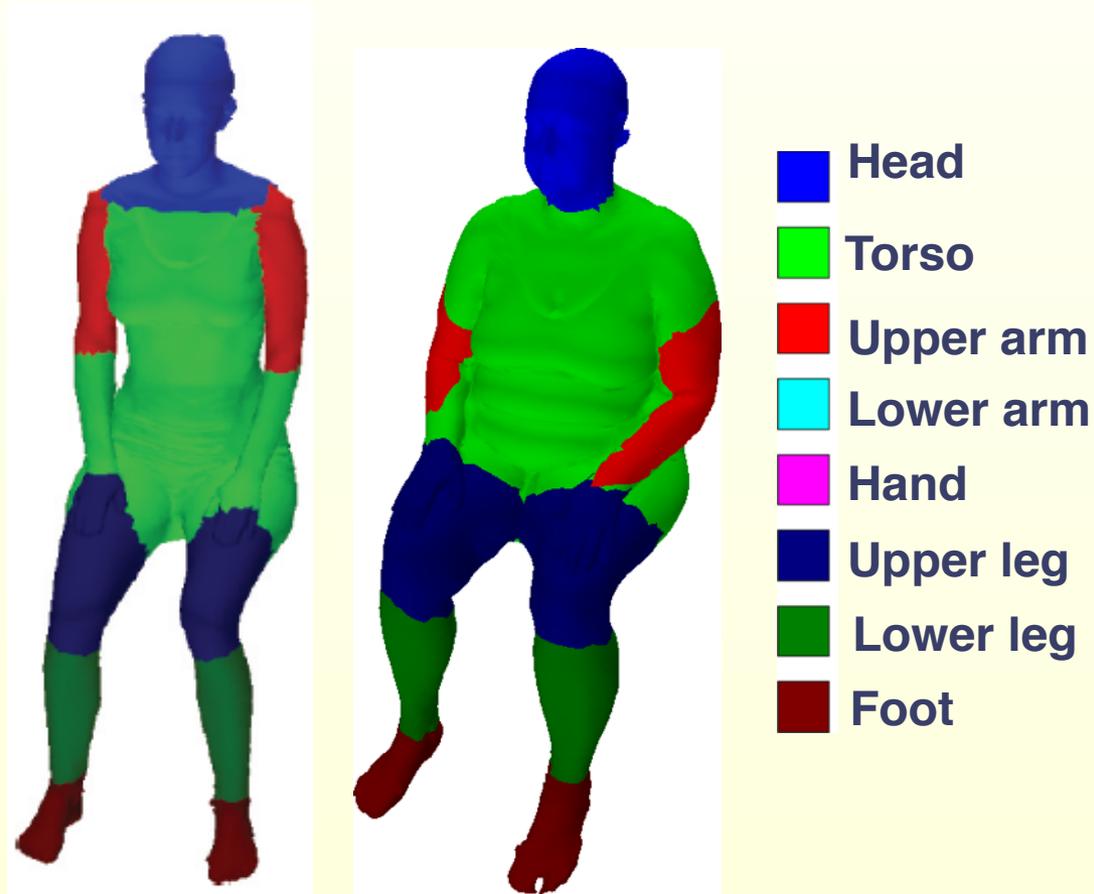


19 training meshes

3 training meshes

Limitations

- Many features are sensitive to topology

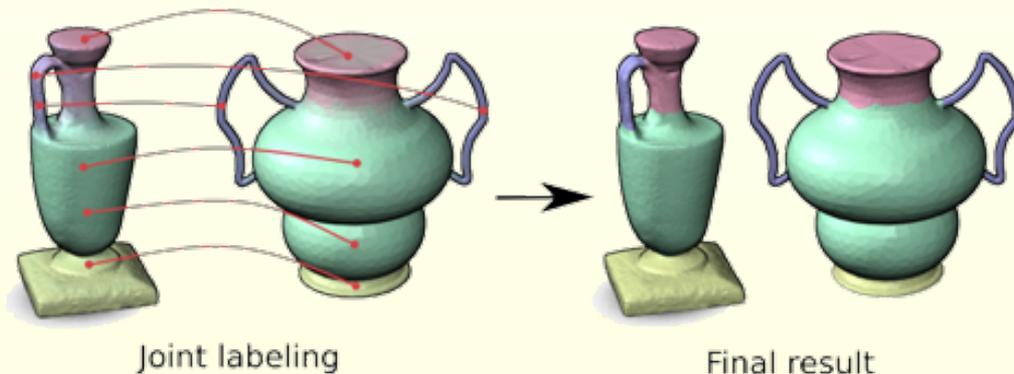


More on learning shape labeling

- A follow-up: “**Prior Knowledge for Part Correspondence**”
[van Kaick et al. 2011]

More on learning shape labeling

- A follow-up: “**Prior Knowledge for Part Correspondence**” [van Kaick et al. 2011]
 - Similar problem formulation, combined with per-face mesh alignment



Questions so far?

Unsupervised segmentation

UNSUPERVISED CO-SEGMENTATION OF SHAPE COLLECTIONS

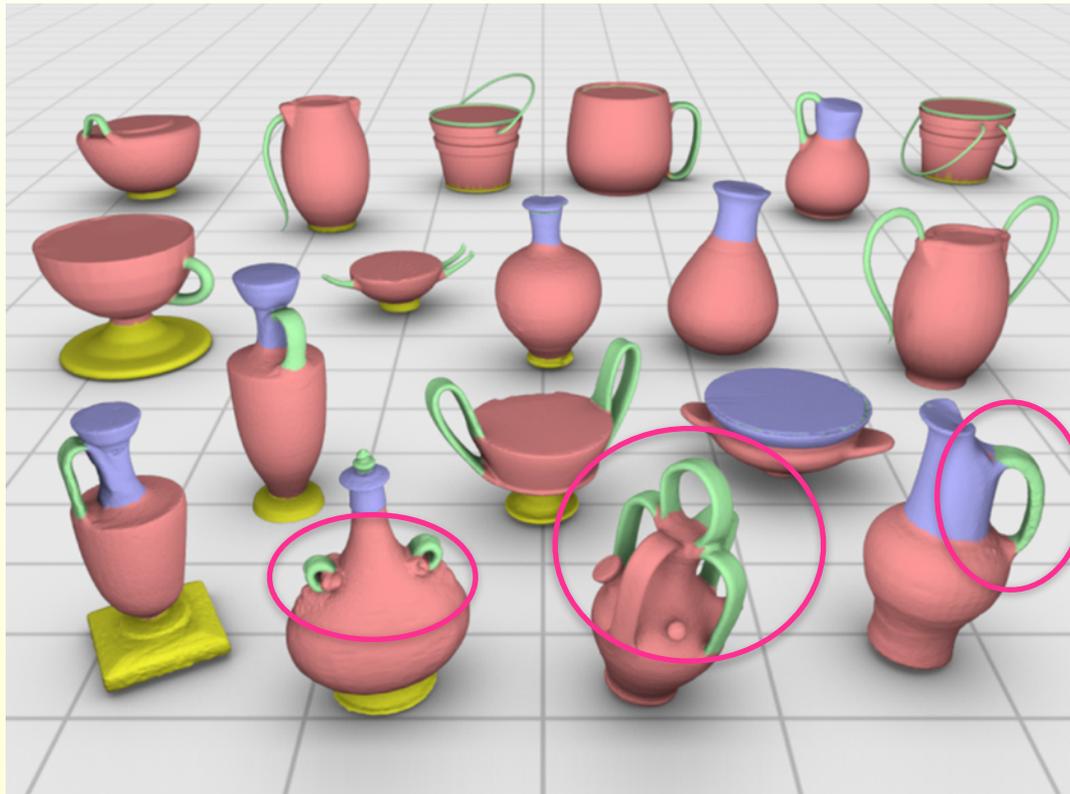
Unsupervised co-segmentation in shape collections

- A collection of shapes sharing a general structure and functionality, but not necessarily similar corresponding parts



Unsupervised co-segmentation in shape collections

- Our goal: derive consistent segmentation across the collection
- Note the geometric and topological differences of parts!



E.g. handles

Unsupervised co-segmentation in shape collections

- **“Unsupervised Co-Segmentation of a Set of Shapes via Descriptor-Space Spectral Clustering”** [Sidi et al. 2011]

Main ideas

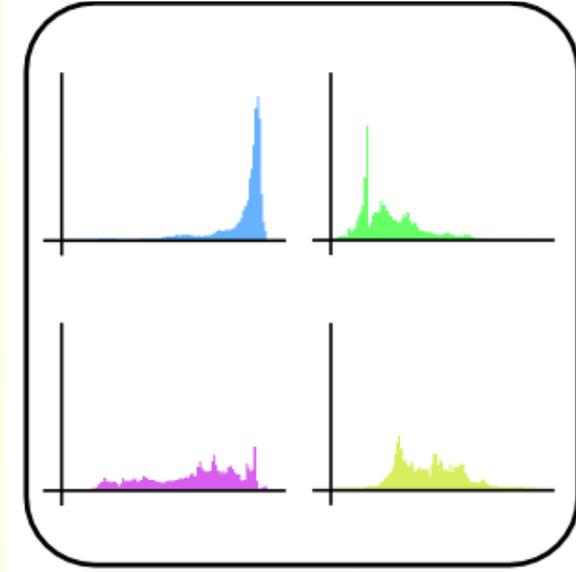
- Treat the segmentation problem as clustering in the space of shape descriptors rather than spatial coordinates
 - Pro: doesn't require spatially aligned shapes
- Spectral clustering using diffusion maps
- Create statistical models of shape part descriptors using obtained clusters
- Refine individual segmentations using graph cuts

Unsupervised co-segmentation in shape collections

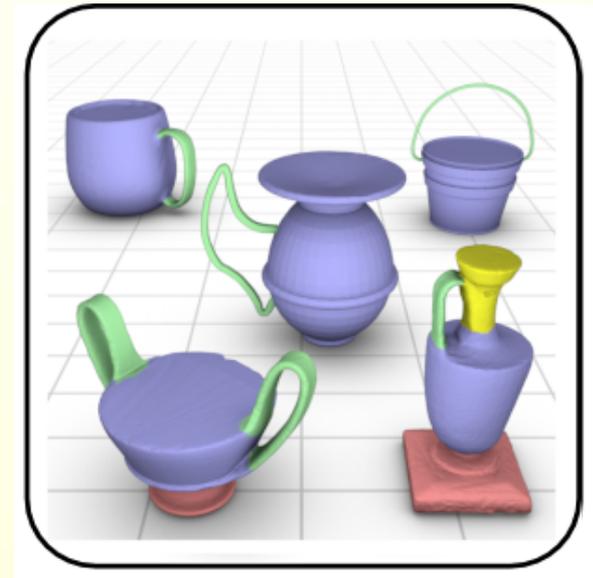
- Algorithm outline



1. Pre-segmentation



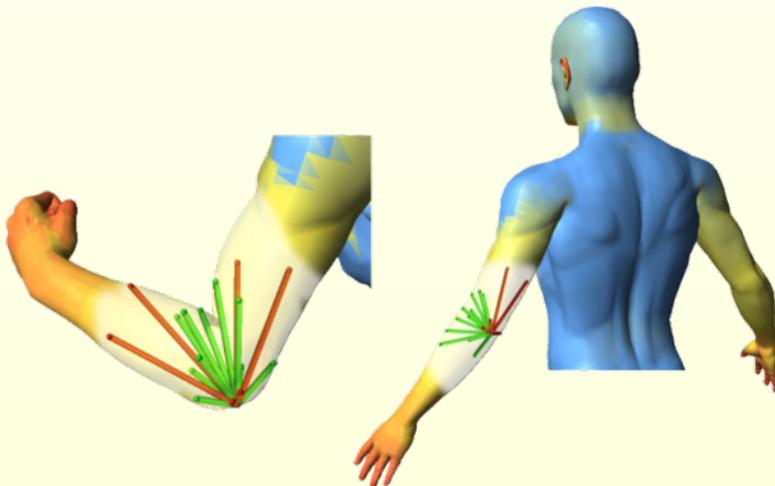
2. Segmentation by statistical descriptor embedding



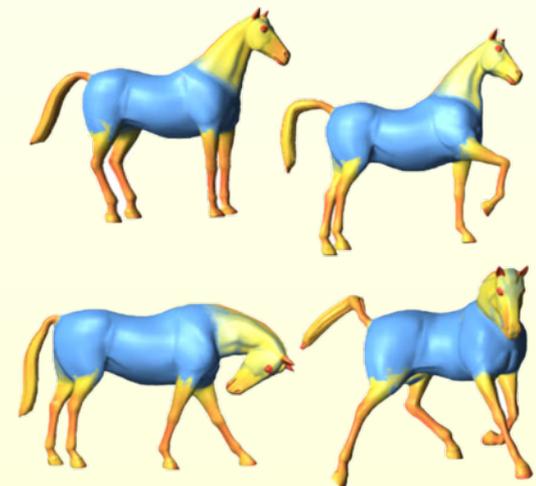
3. Clustering

Step 0: compute shape descriptors

- Input: scale-normalized shapes in upright orientation
- **Face-level descriptors**
 - Geodesic distance from the base of the shape
 - Angle between face normal and upright direction
 - Shape diameter function (SDF) [Shapira et al. 2009]



Examples of the cone of rays shot to the opposite side of the shape



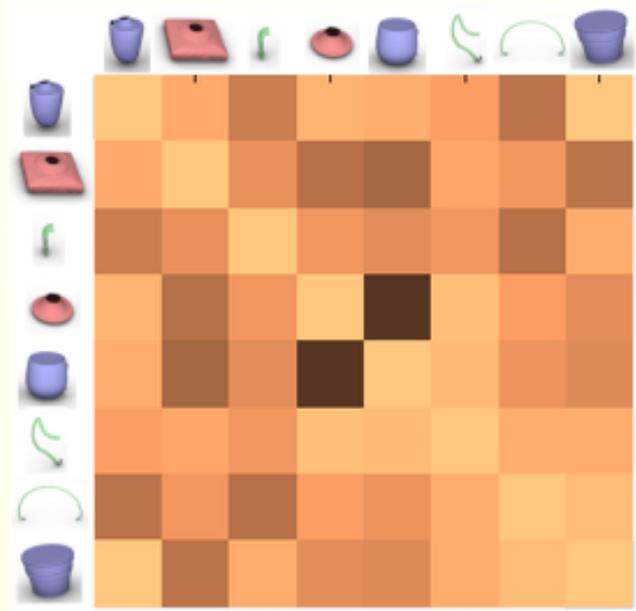
SDF - weighted average of all ray lengths (see the paper for details)

Step 1: pre-segmentation

- Initial (imperfect) segments obtained by *mean-shift clustering* [Comaniciu and Meyer 2002] applied to face descriptors
- **Segment-level descriptors**
 - Histograms of face-level descriptors per segment h_i^d - for descriptor d and patch i
 - Segment area normalized by total shape area a_i
 - Segment geometry descriptor, derived from eigenvalues of PCA applied to segment vertices g_i

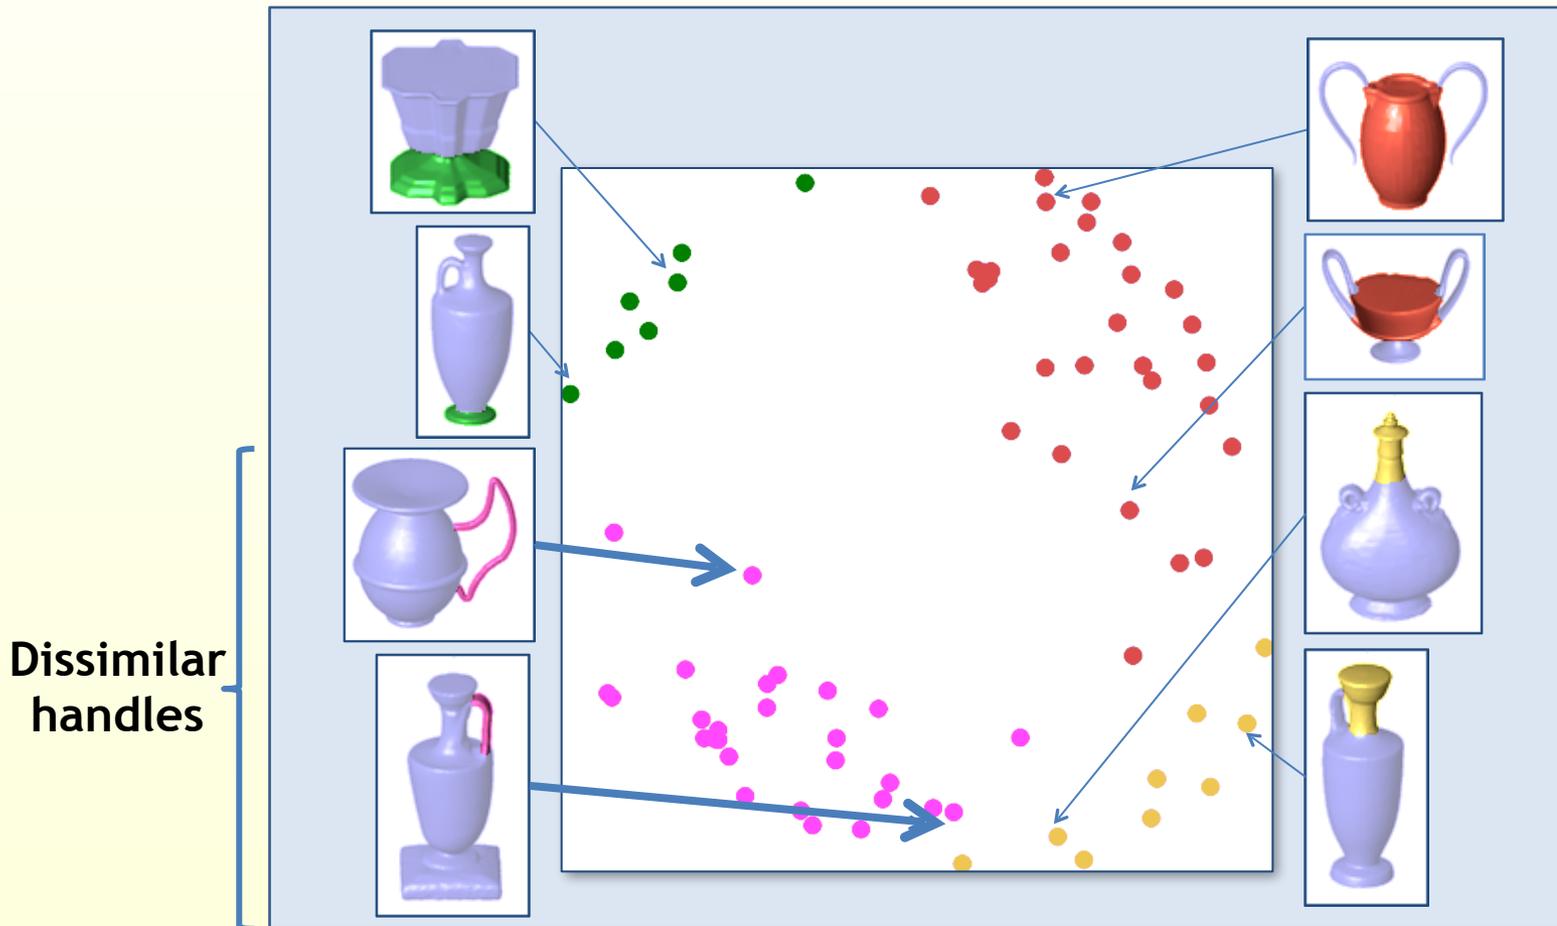


Step 2: diffusion descriptor embedding



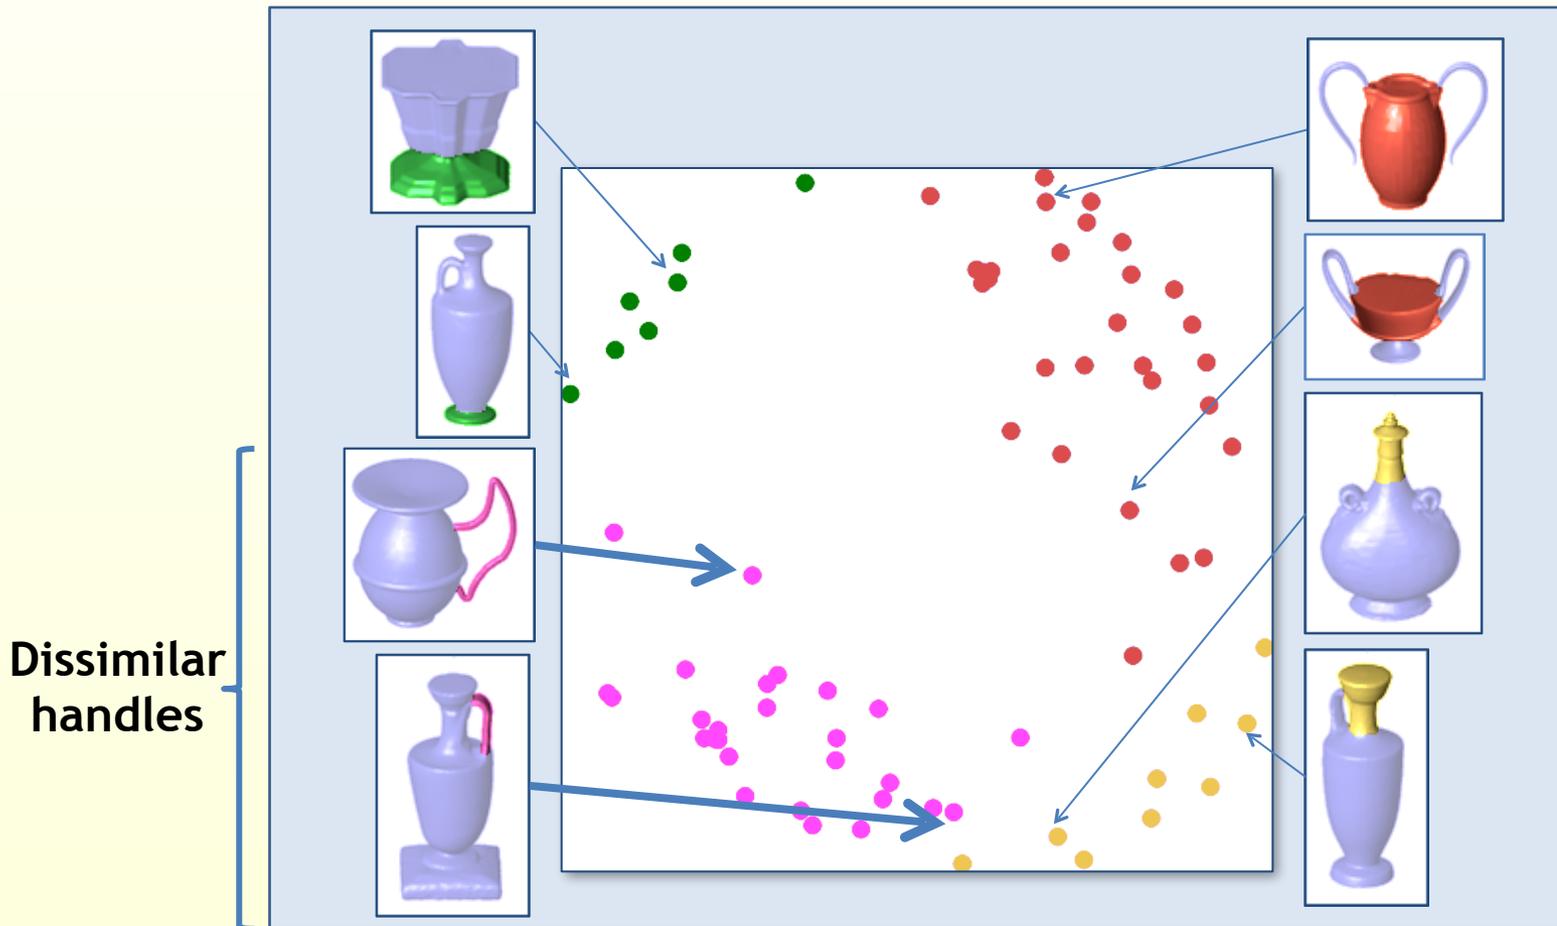
Step 2: descriptor embedding - motivation

- Applied to segment descriptors



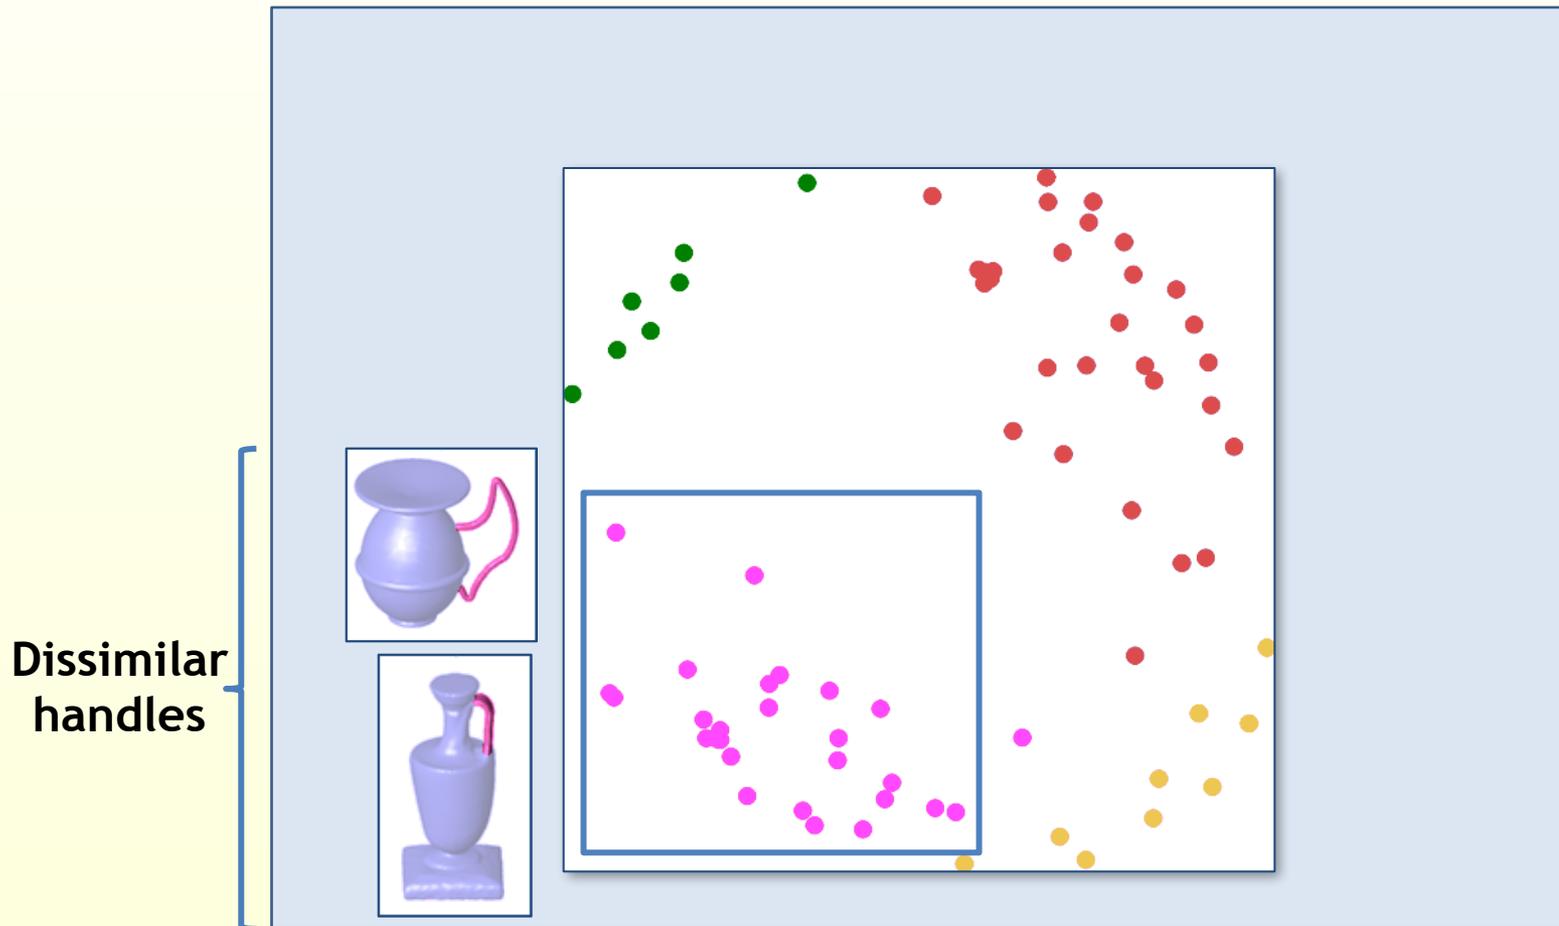
Step 2: descriptor embedding - motivation

- Applied to segment descriptors



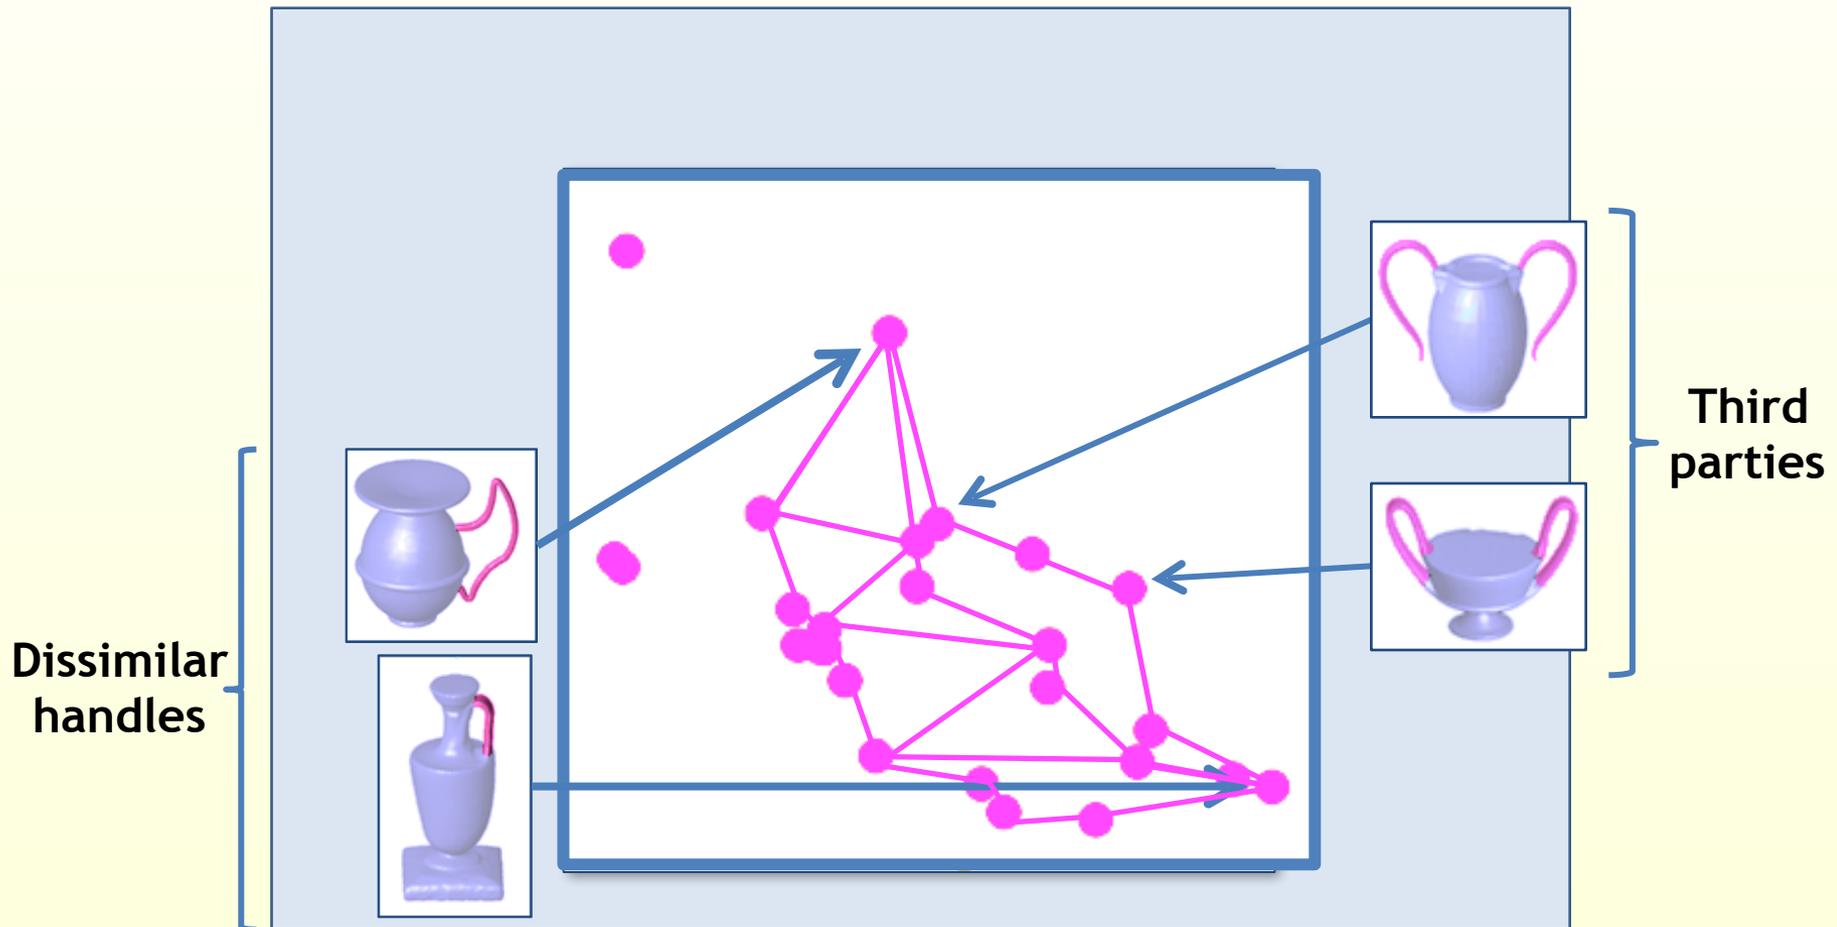
Step 2: descriptor embedding - motivation

- Applied to segment descriptors



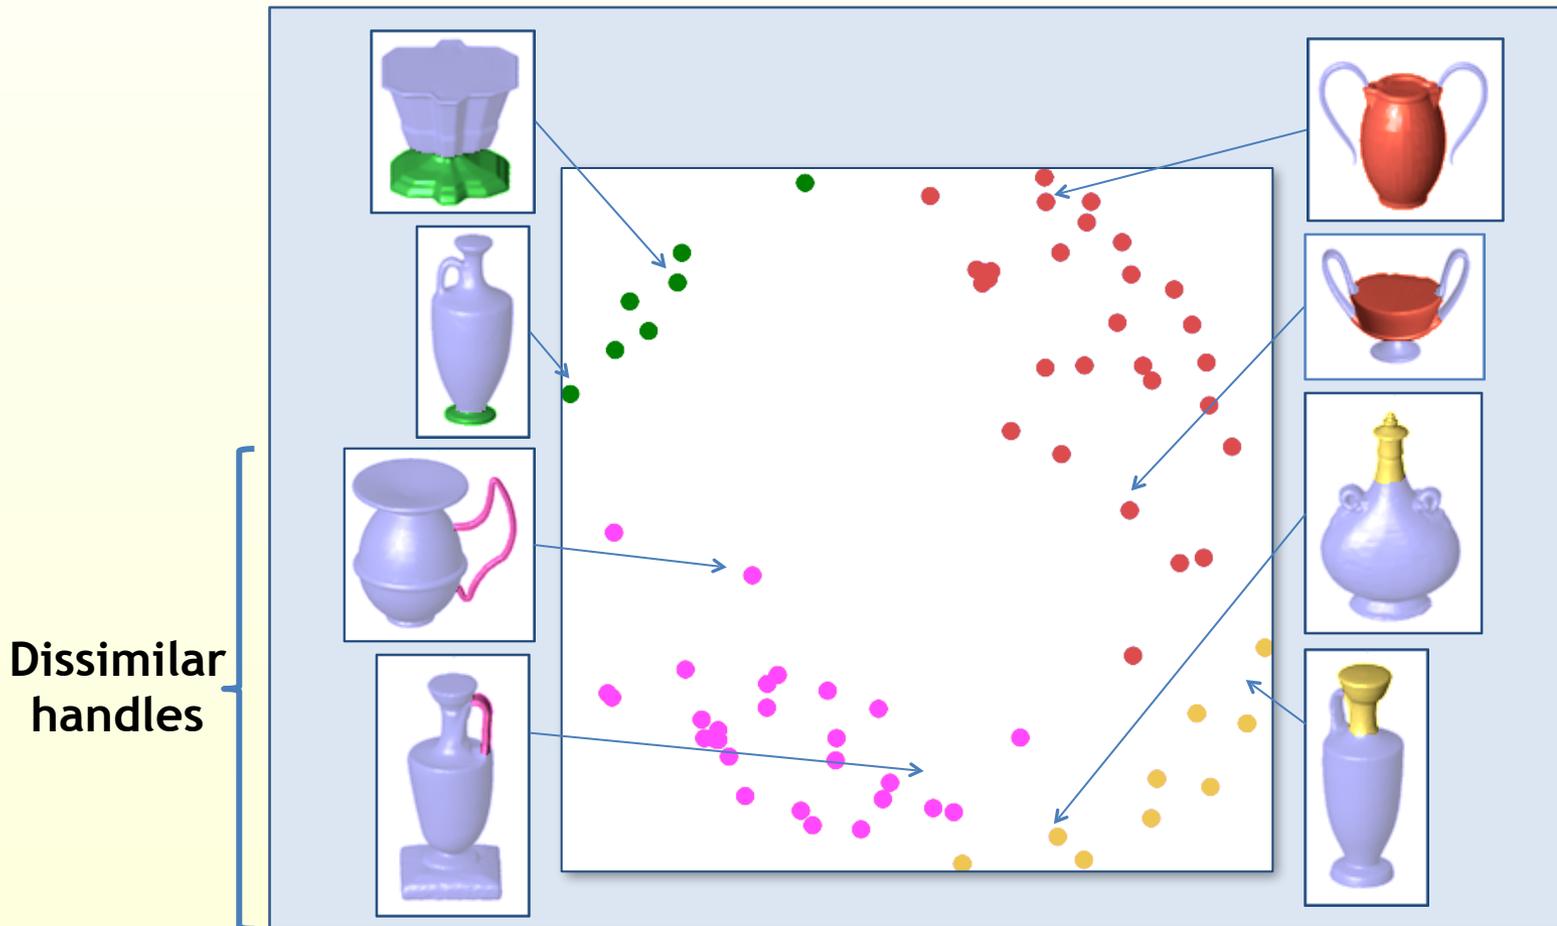
Step 2: descriptor embedding - motivation

- Third party connections



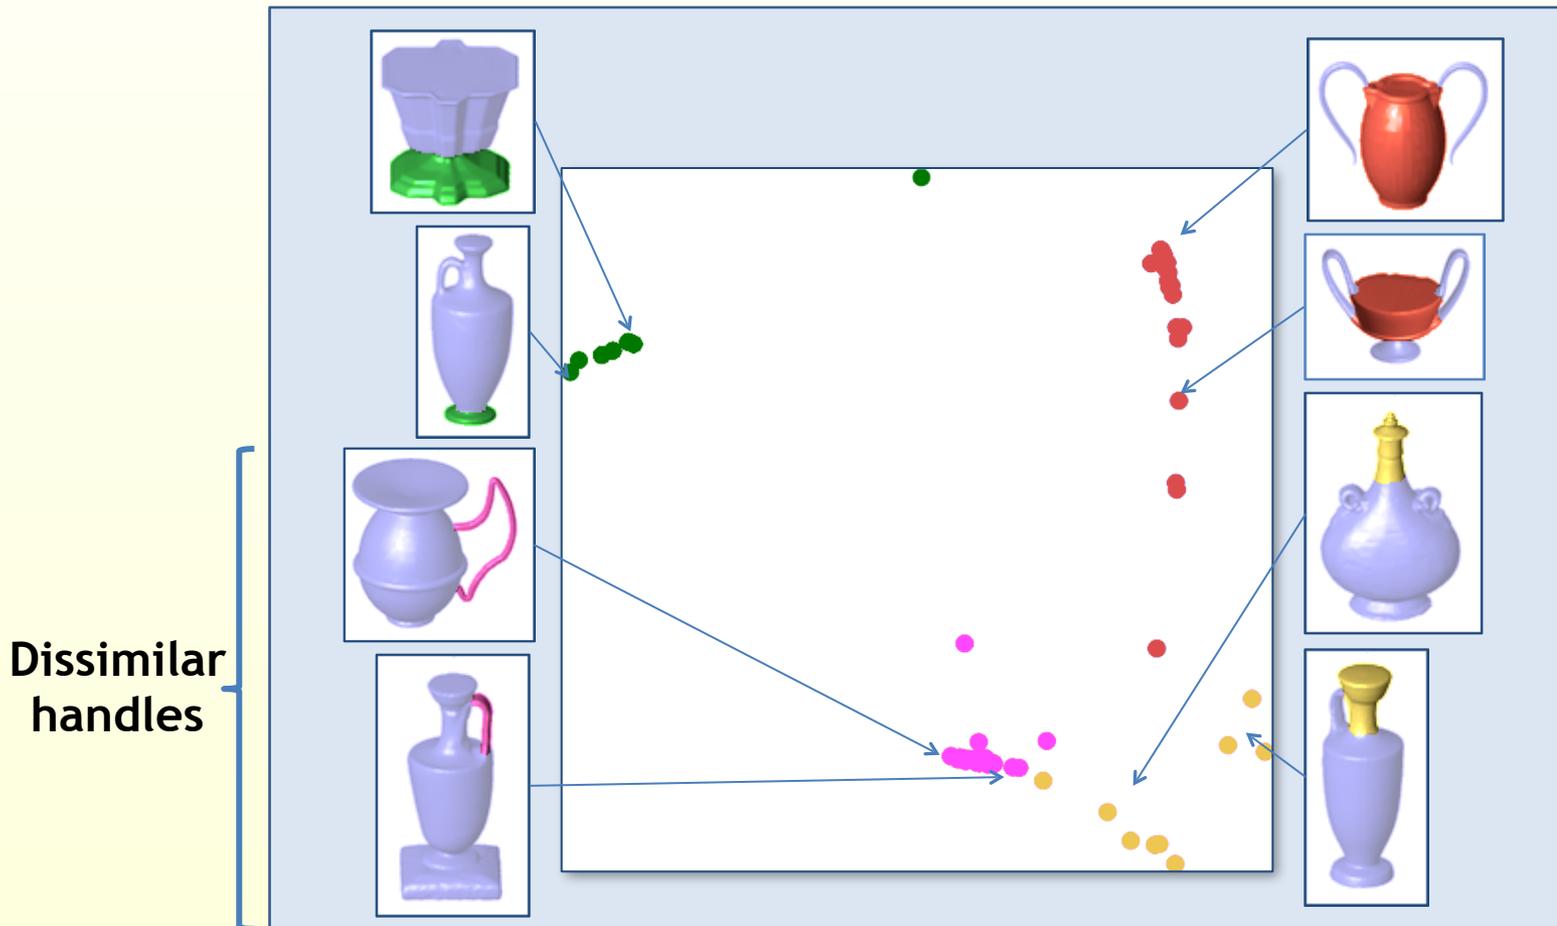
Step 2: descriptor embedding - motivation

- Diffusion map embedding via third parties

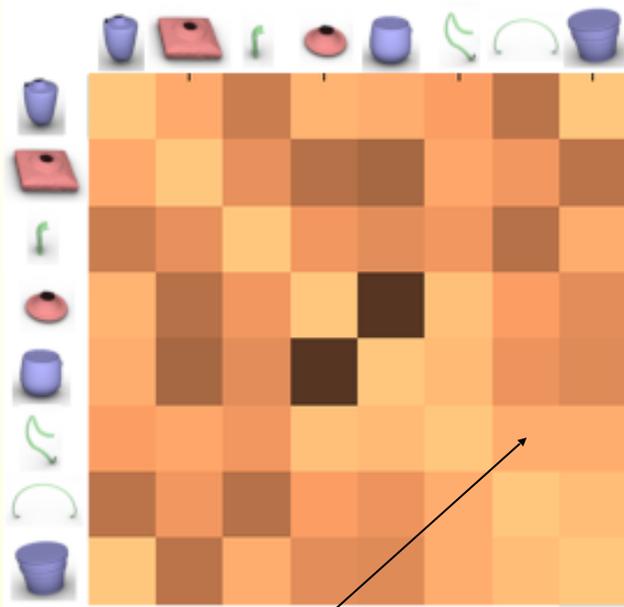


Step 2: descriptor embedding - motivation

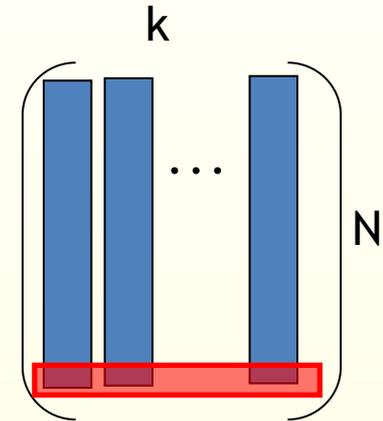
- Diffusion map embedding via third parties



Step 2: diffusion descriptor embedding



* more details
in the paper



k Leading
eigenvectors
normalized

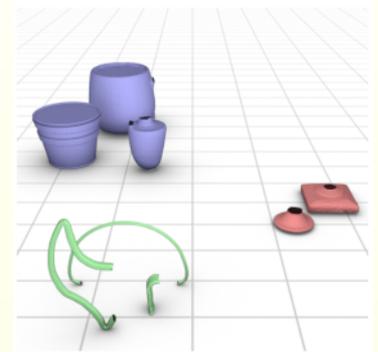
e.g. $A_{ij} = e^{-\frac{D(s_i, s_j)^2}{2\sigma^2}}$

3 eigenvectors were used in the paper

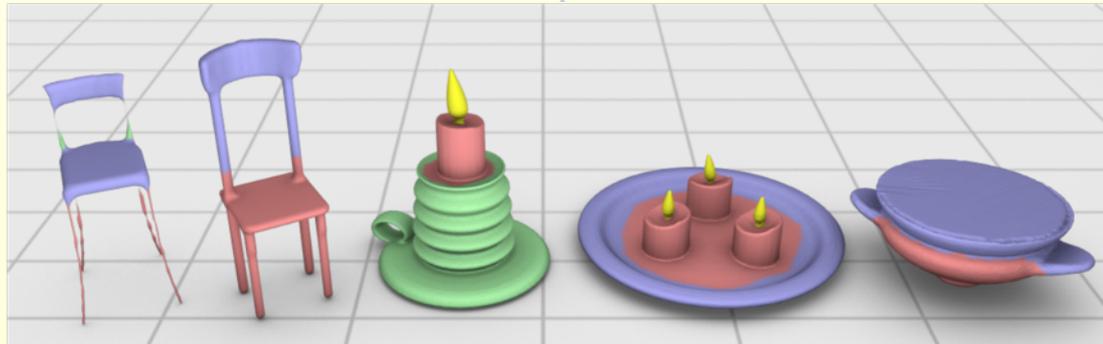
$$D(s_i, s_j) = \sqrt{\sum_{d=1}^{n_d} \text{EMD}^2(h_i^d, h_j^d) + |a_i - a_j|^2 + \|g_i - g_j\|_2^2}$$

Step 3: clustering in the embedding space

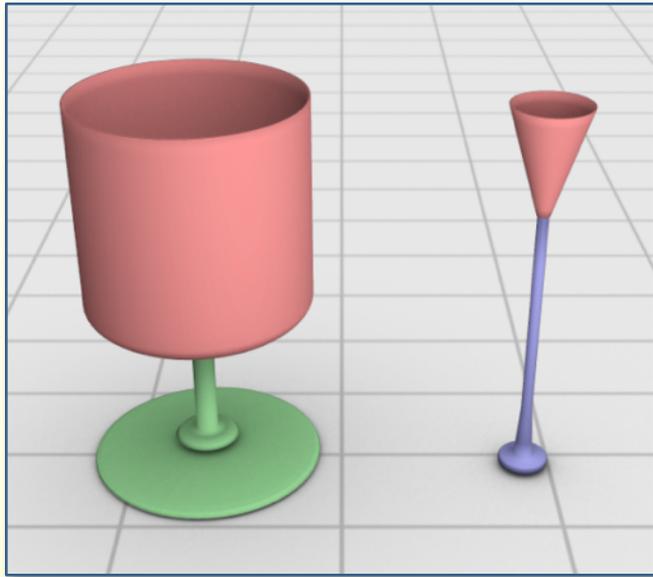
- Agglomerative hierarchical clustering [?] in the embedding space
- Number of classes given by the user



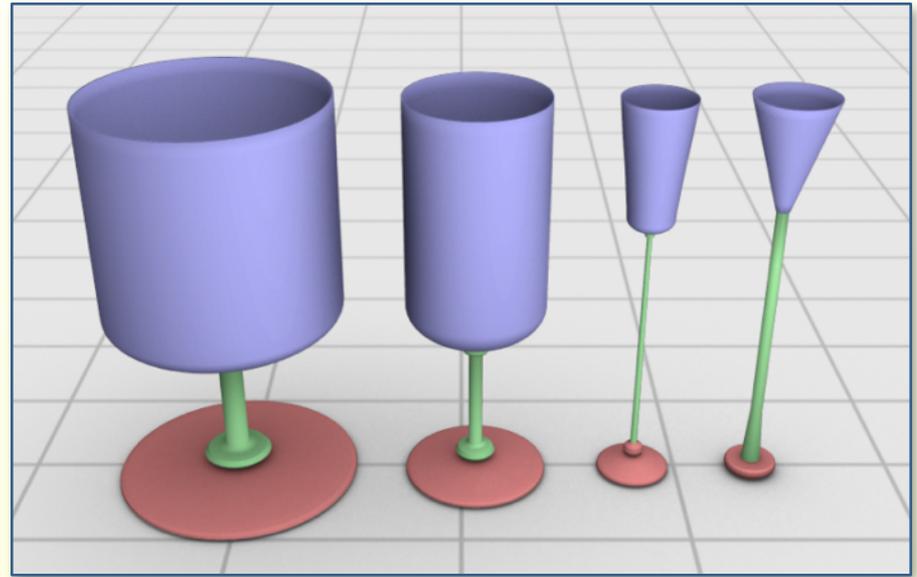
- Clustering gives an initial co-segmentation



Step 3: more results



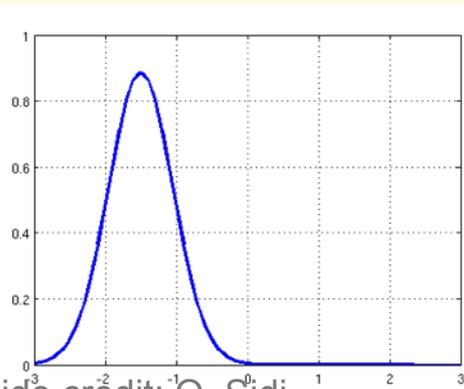
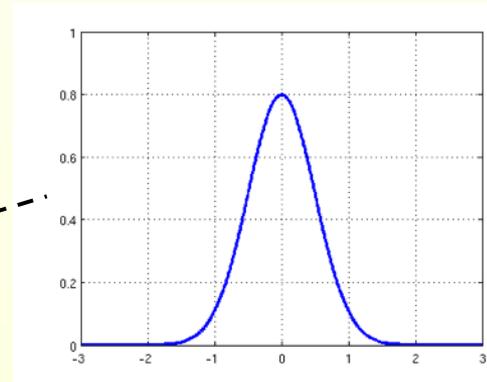
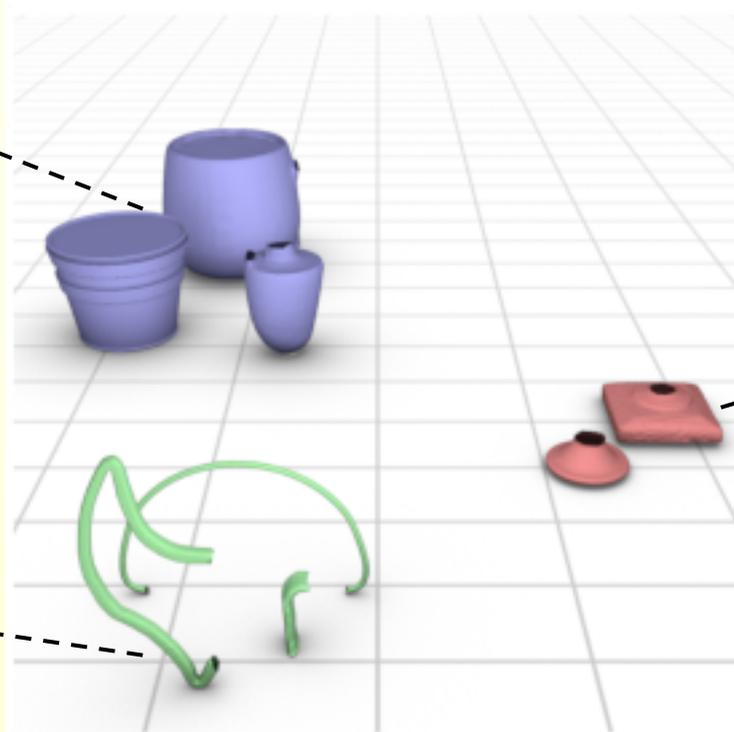
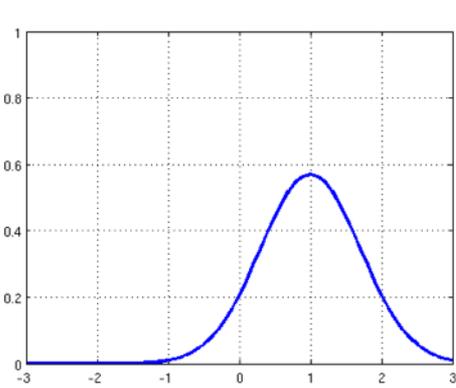
Initial shapes



Descriptor space clustering

Step 4: statistical model and labeling

- Goal: create statistical models of part descriptors to refine the segmentation

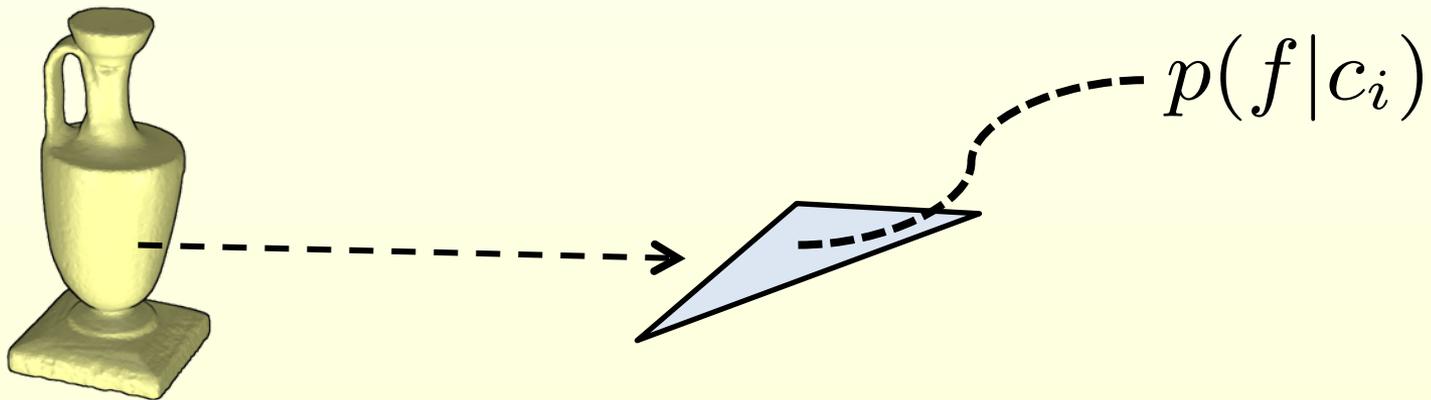


Step 4: statistical model and labeling

- Fit three-dimensional Gaussians to descriptors, using Expectation Maximization (EM)

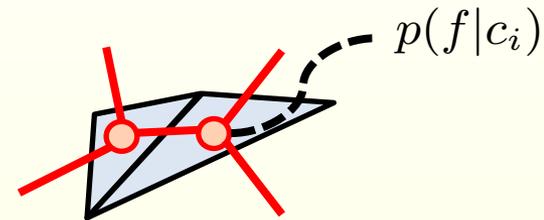
$$p(f|c_i) = p(f, \mu_i, \Sigma_i) = C e^{-\frac{1}{2}(f - \mu_i)^T \Sigma_i^{-1} (f - \mu_i)}$$

- Inference: probability of face belonging to class



Step 4: refined co-segmentation

- Graph cuts labeling based on learned statistical model
- Graph (G, E) given by triangle connectivity
 - Nodes in G correspond to triangles
 - Edges in E correspond to adjacent faces
- Energy functional



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

- where l_u and l_v are the labels assigned to nodes u and v ; \mathcal{E}_D and \mathcal{E}_S are the data and smoothness energy terms

Step 4: refined co-segmentation

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

How well label / matches face u ?

- Data term:

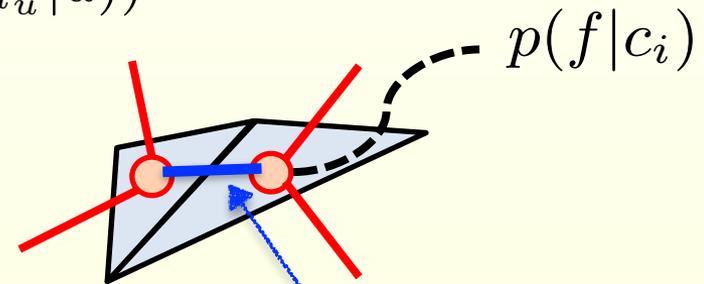
$$\mathcal{E}_D(u, l_u) = -\omega \log(p(c_{l_u} | u))$$

Is it a good place for a cut?

- Smoothness term:

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

(similar to [Shapira et al. 2009])

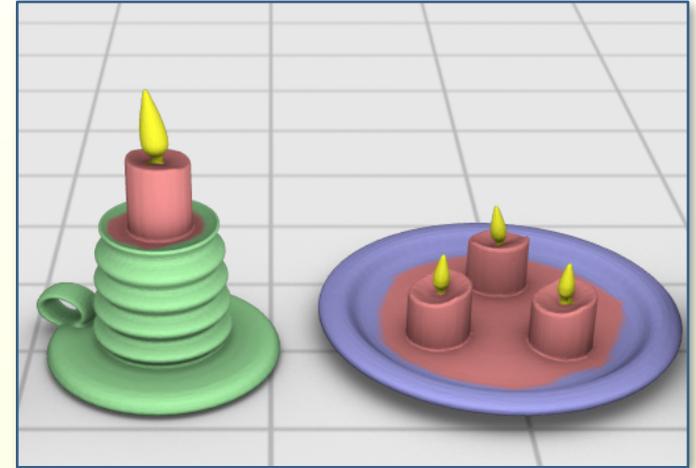
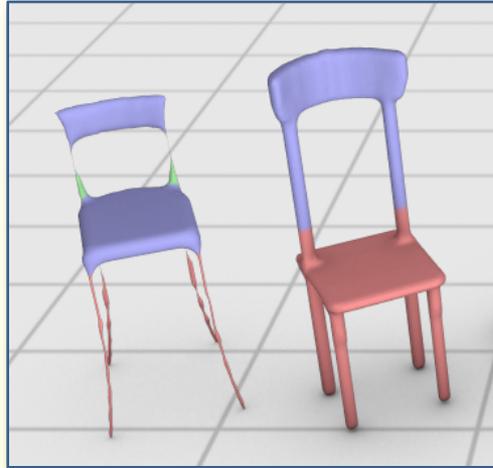


edge length

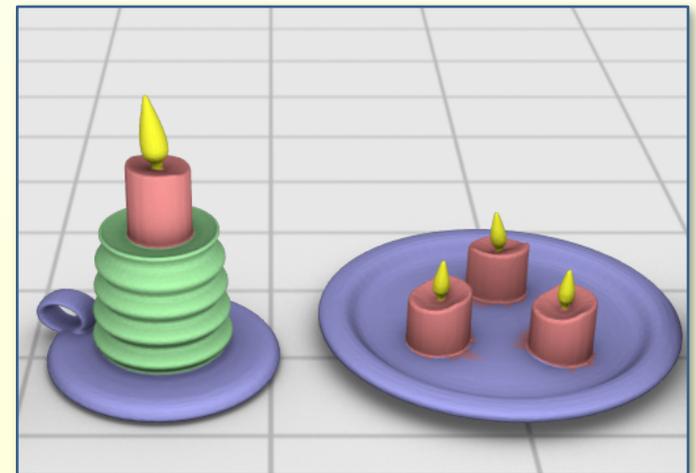
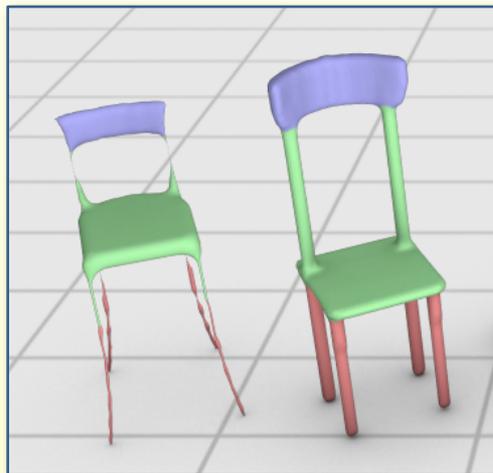
dihedral angle between faces

Step 4 result - final co-segmentation

Clustering
result



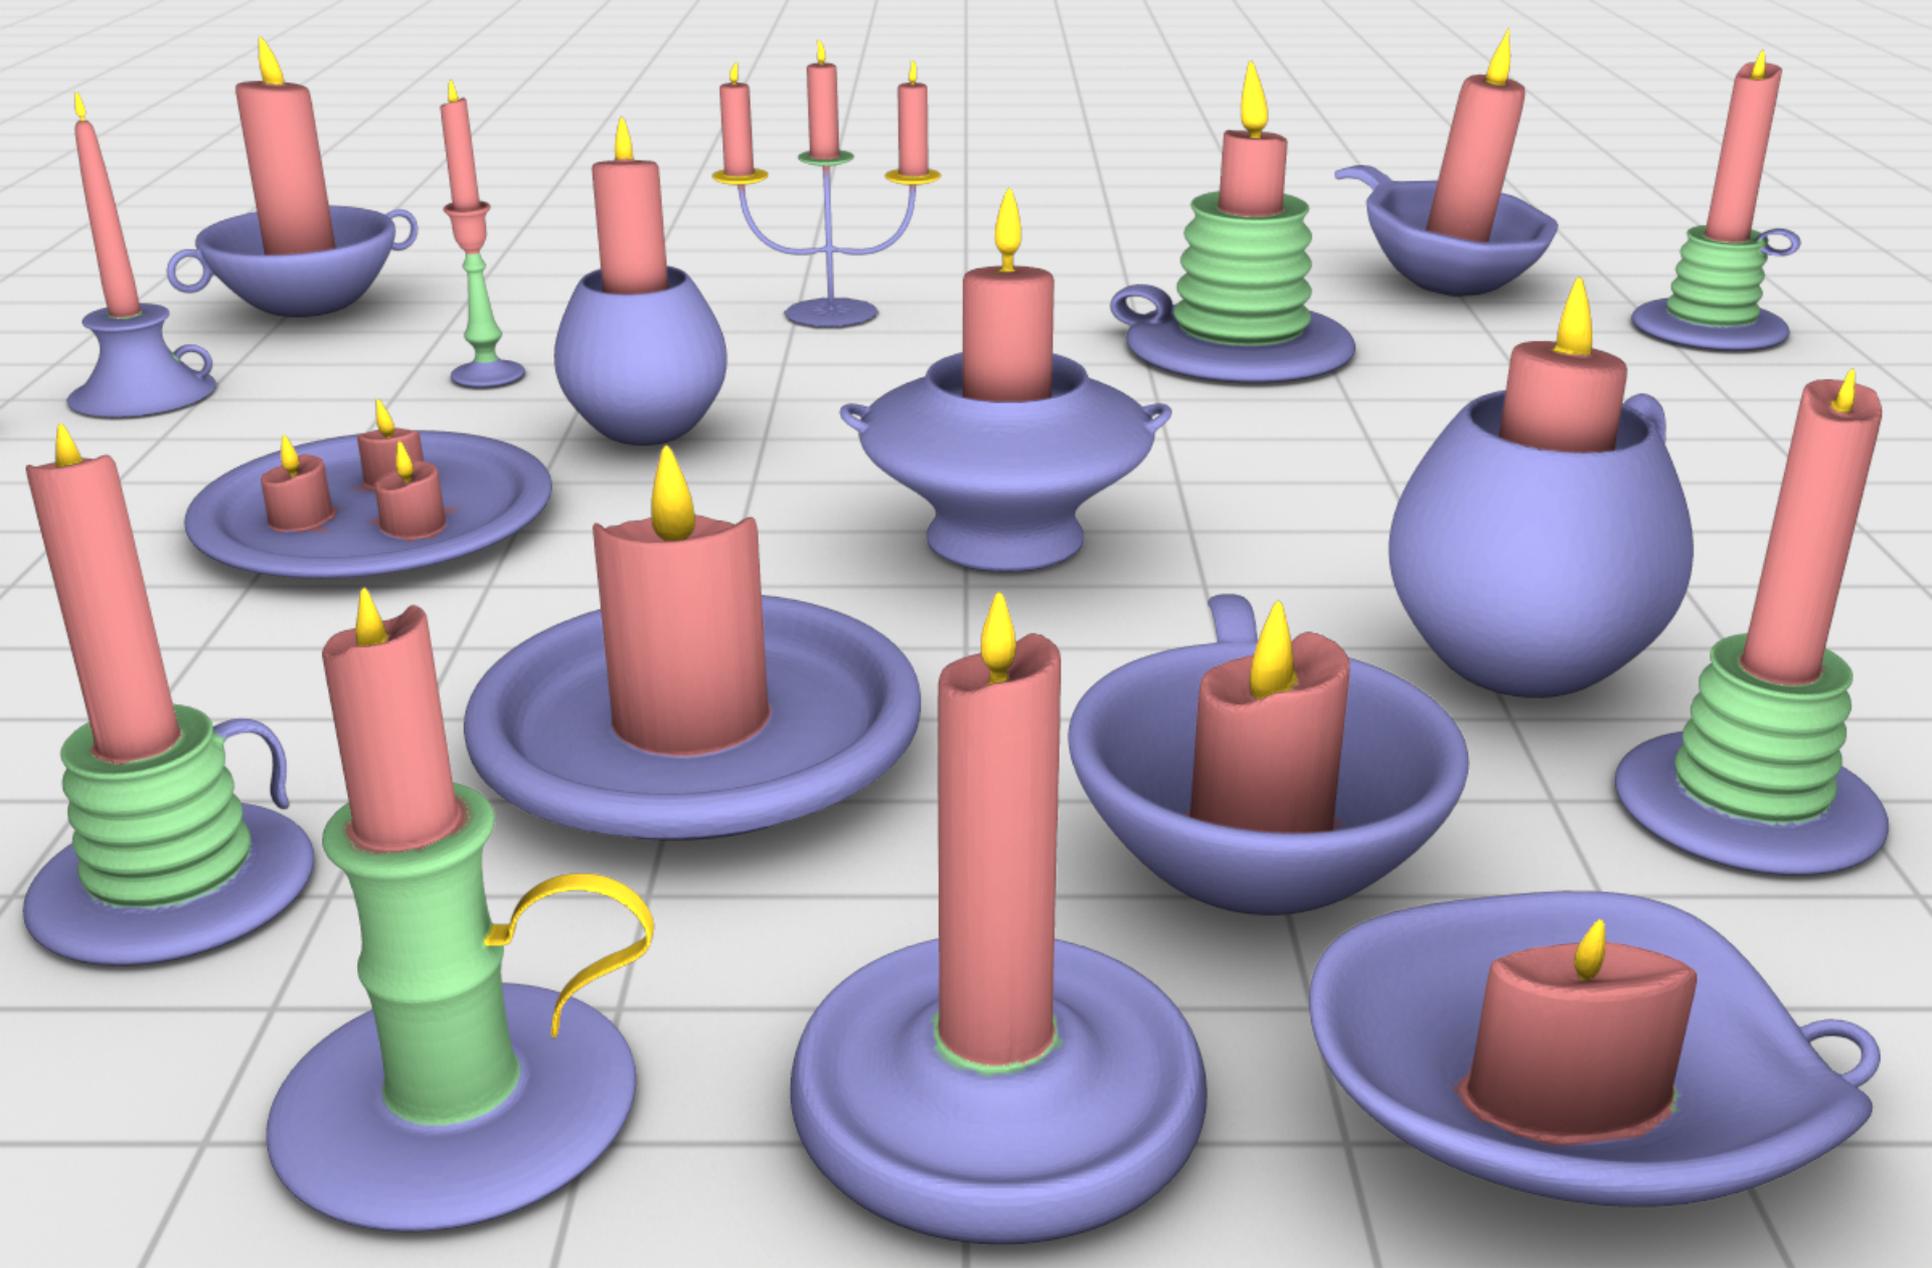
Graph cut
result



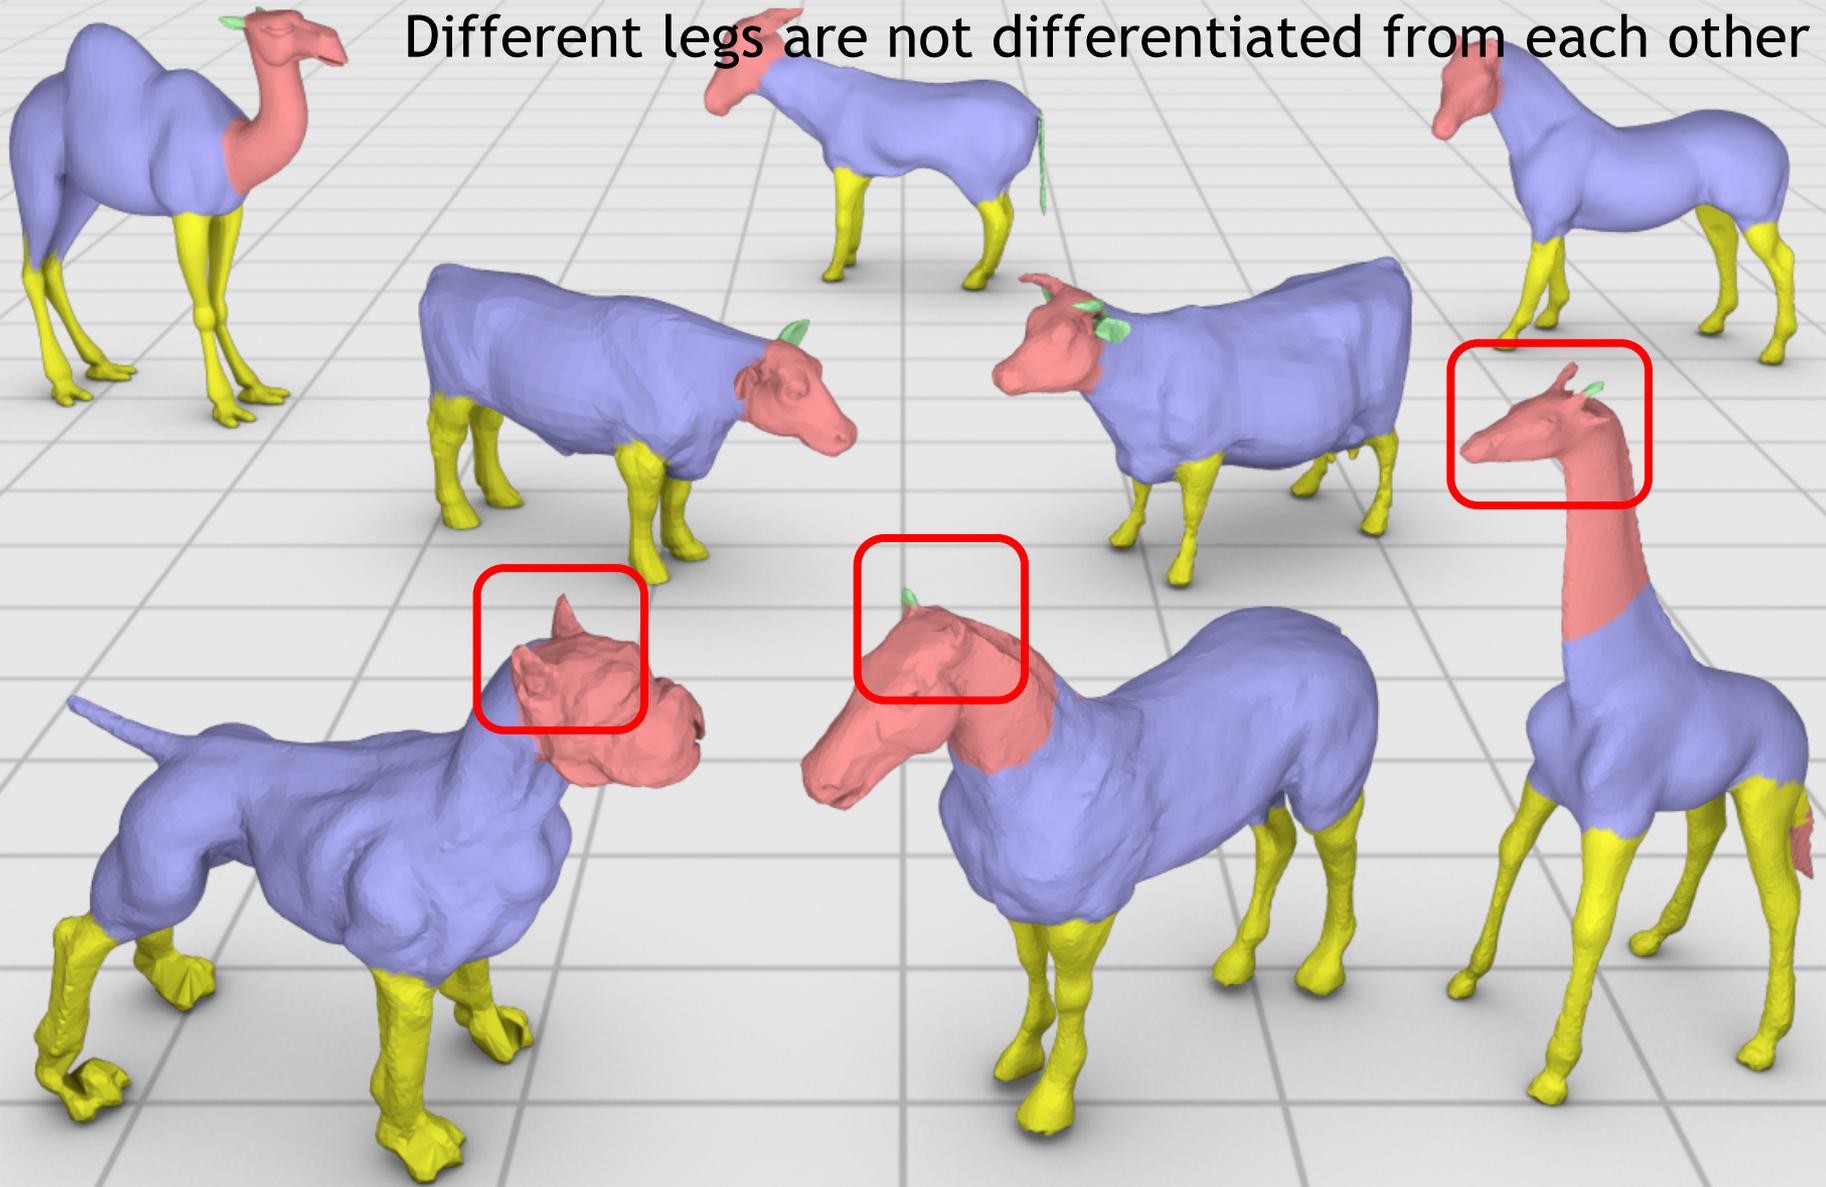
Results

- Dataset: seven sets of shapes
 - Man-made shapes
[van Kaick et al. 2011]
 - Organic shapes
[Kalogerakis et al. 2010, Chen et al. 2009]
- Variability in geometry and part composition



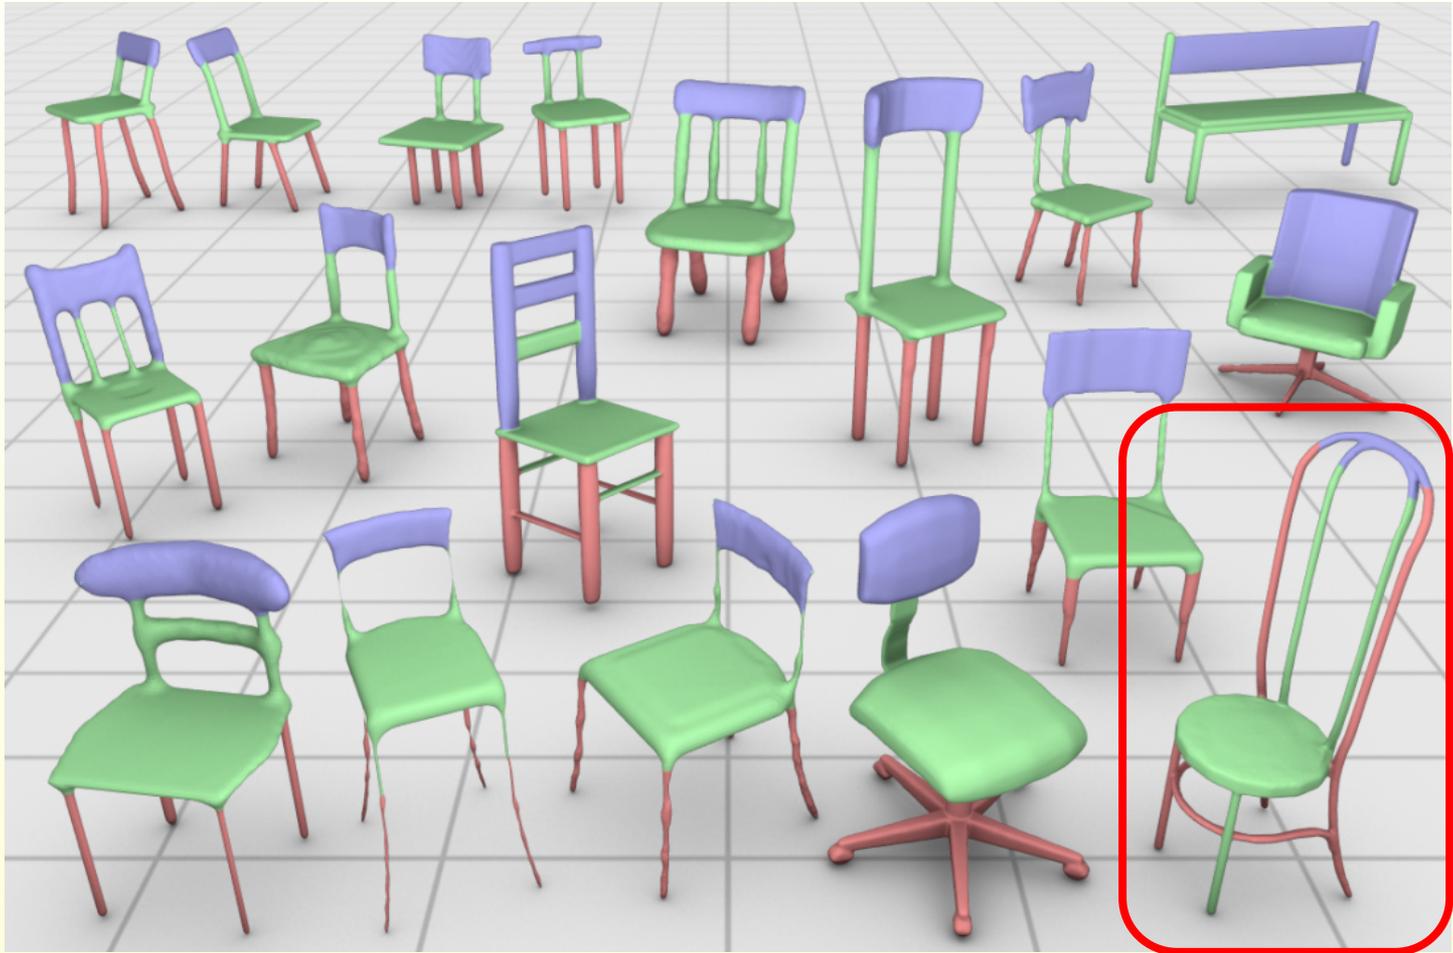


Different legs are not differentiated from each other

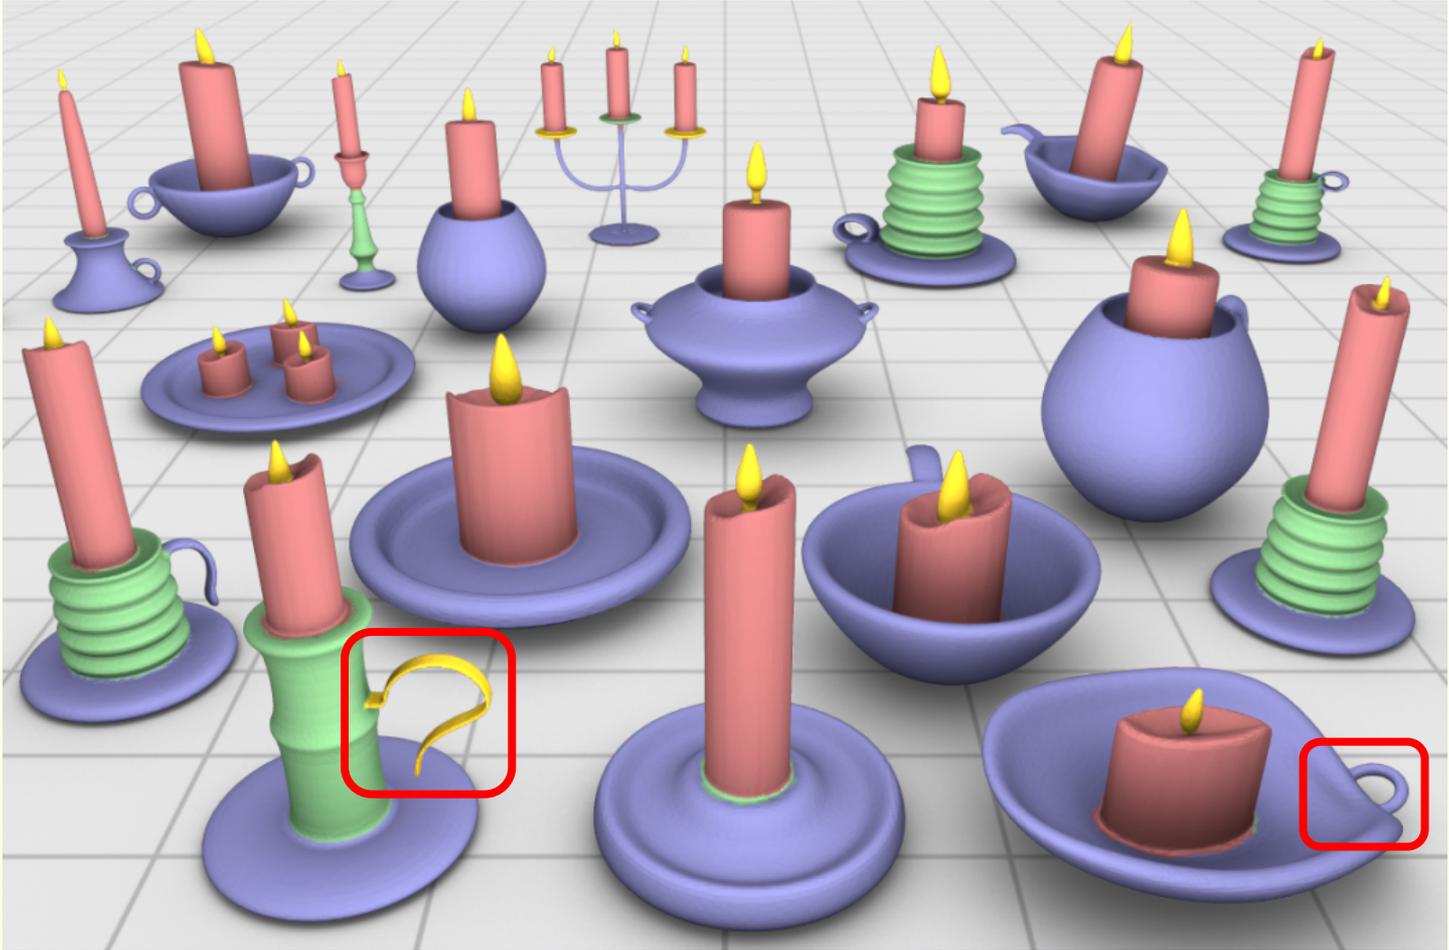


Limitation: small components like ears and horns are not consistently segmented

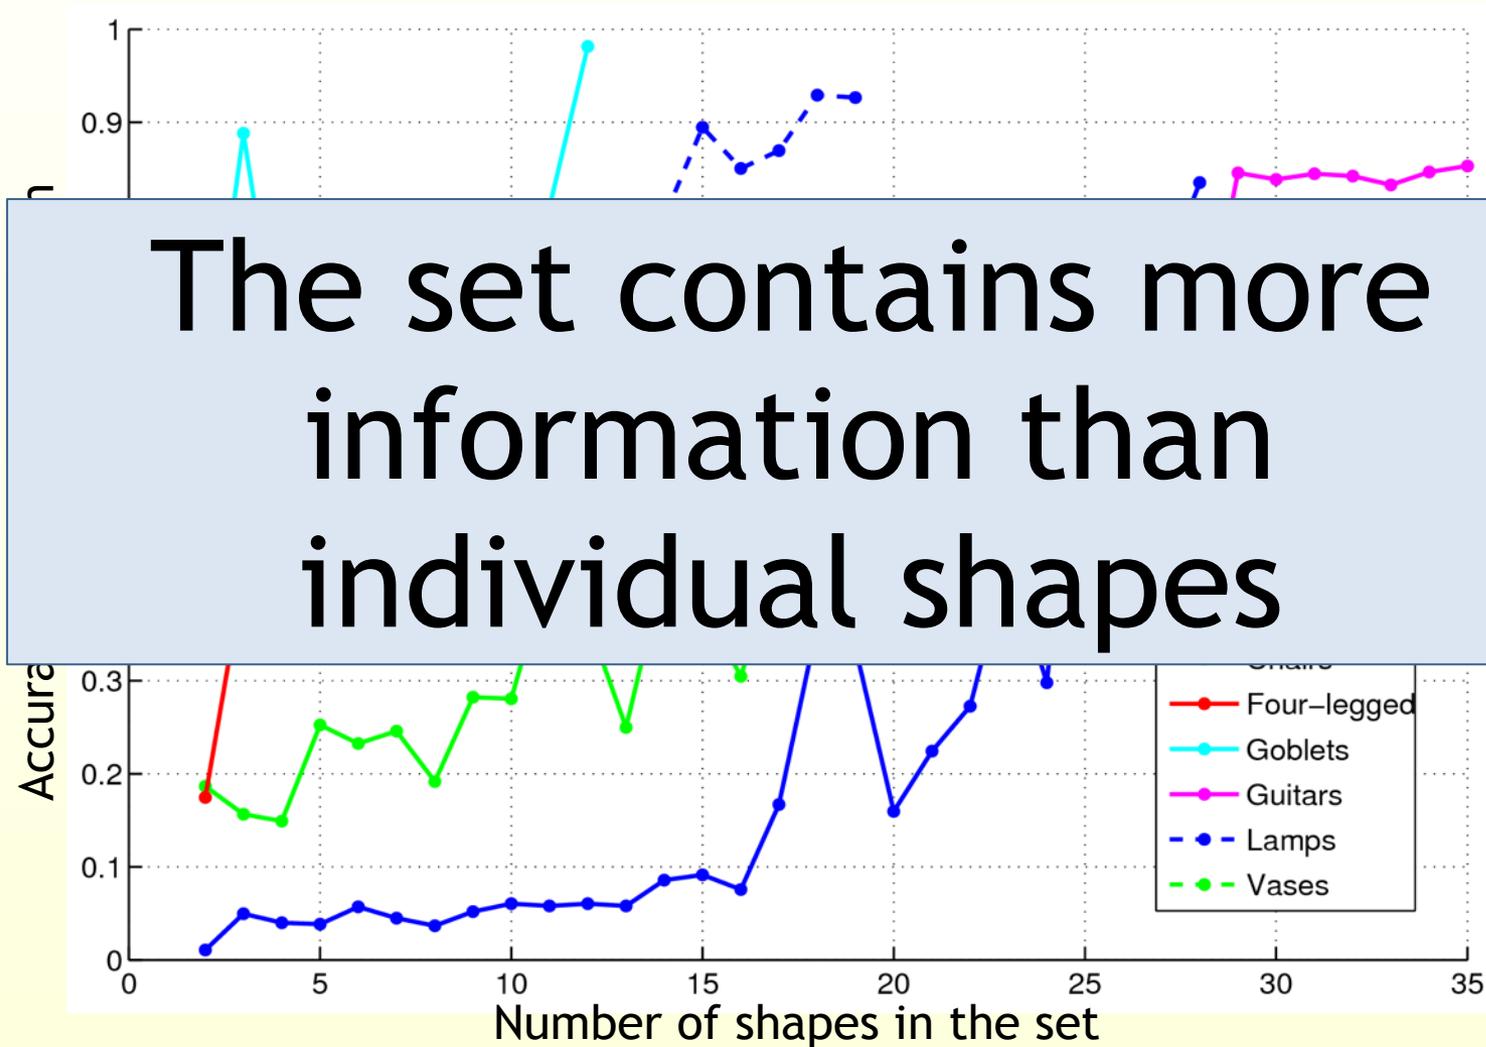
Limitations: lack of third parties



Limitations: incorrect connections

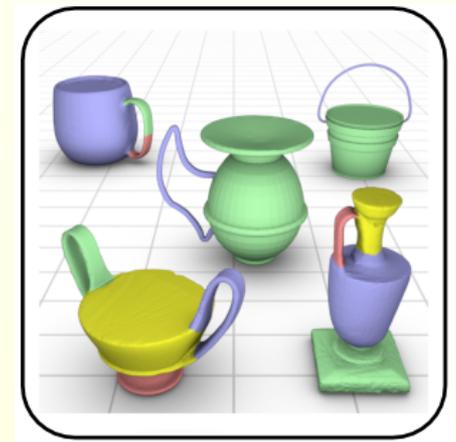


Effect of the set on the accuracy



Descriptor clustering - different flavor

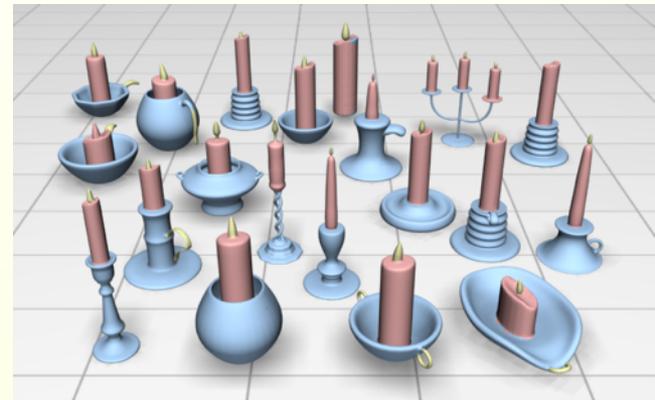
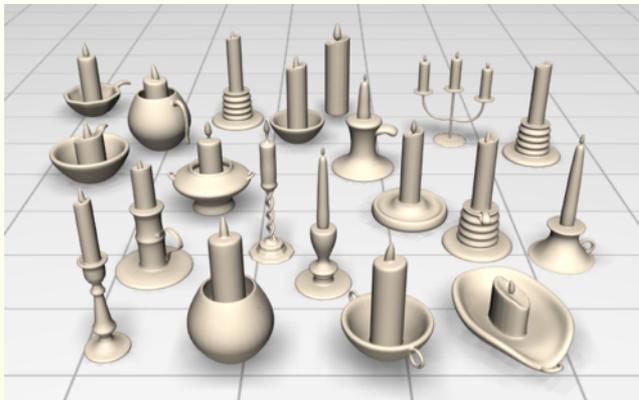
- A follow-up work: “**Co-Segmentation of 3D Shapes via Subspace Clustering**”, Hu et al. 2012
- Previous method [Sidi et al. 2011]
 - Relies on pre-segmentation
 - ▶ affects quality of final result
 - Uses SDF + simple descriptors
 - ▶ can we use more advanced descriptors?



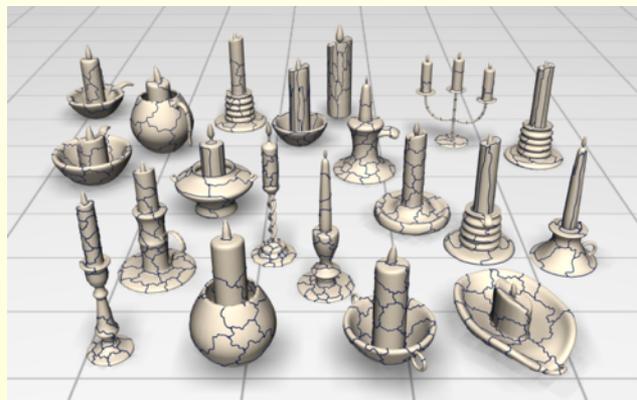
Pre-segmentation

Co-segmentation via subspace clustering

- Utilize over-segmentation to create initial segments



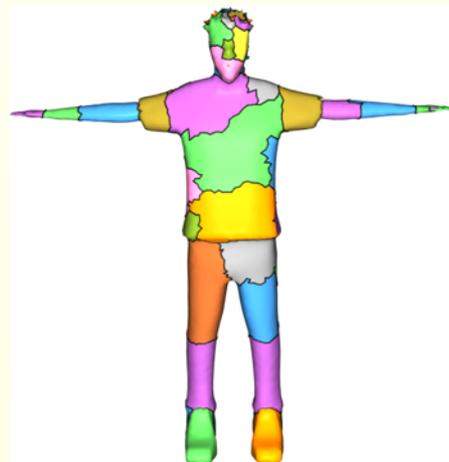
Over-segmentation
[Huang et al. 2011]



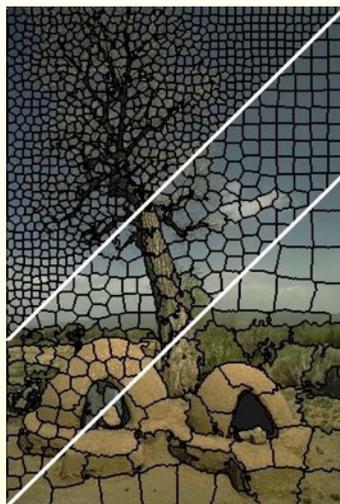
Unsupervised

Over-segmentation

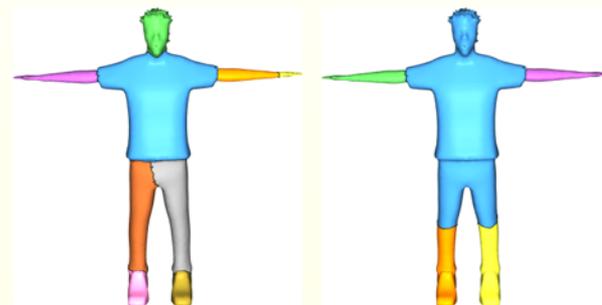
- Segmentation parameterization [Huang et al. 2011]



Patches
[Golovinskiy and Funkhouser 08]



Super-pixels
[Ren and Malik 03]



...
Randomized Cuts



...
Initial Segments

Co-segmentation via subspace clustering: key observation

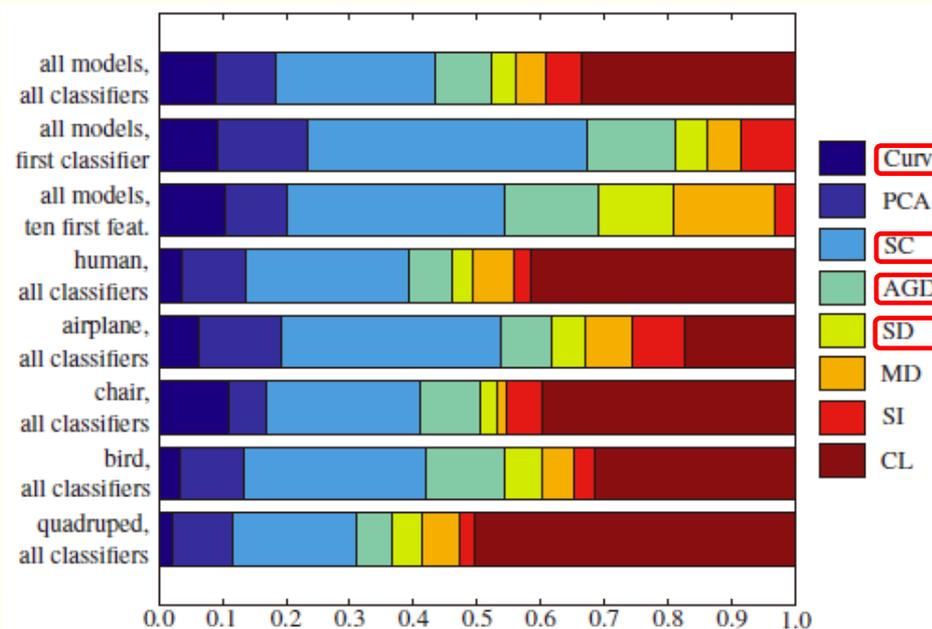
- “two parts of models perceived as corresponding may not necessarily be similar in all features and may even significantly differ on some”



Chair legs are quite similar in AGD features (upper right) while they differ a lot in SDF features (lower right)

Choice of features

- Different sets favor different features
 - Single feature is not enough

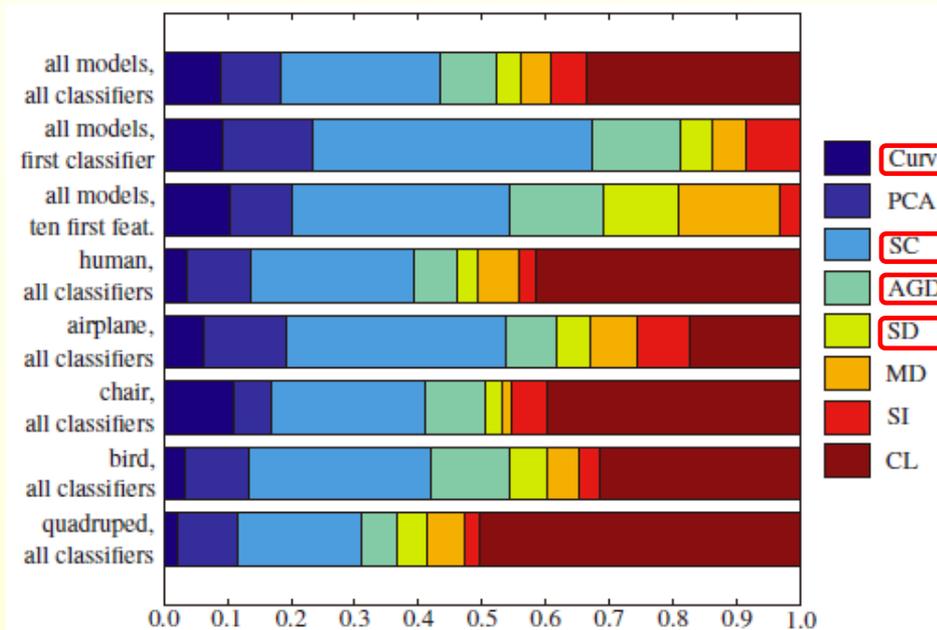


[Kalogerakis et al. 2010]

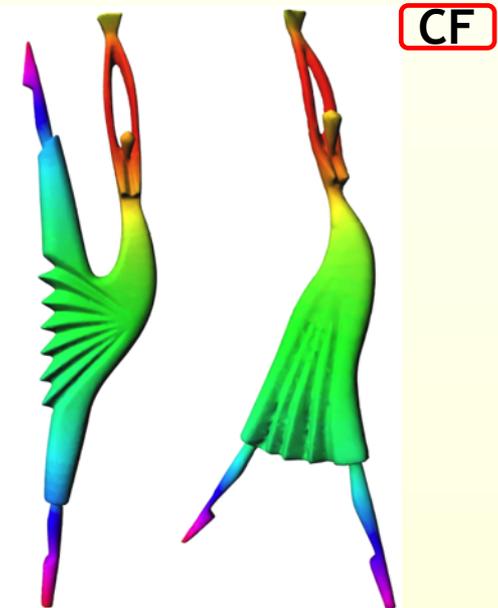
- Curvature
- Shape context
- Average geodesic distance
- Shape diameter (function)

Choice of features

- Different sets favor different features
 - Single feature is not enough



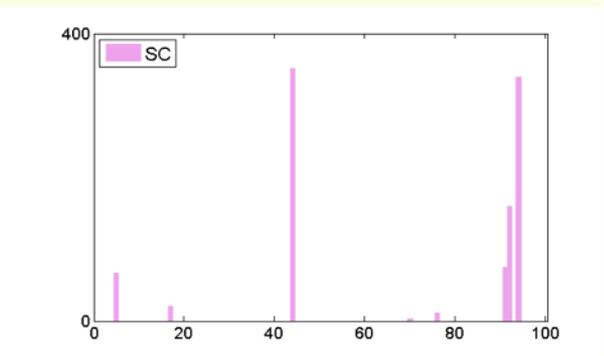
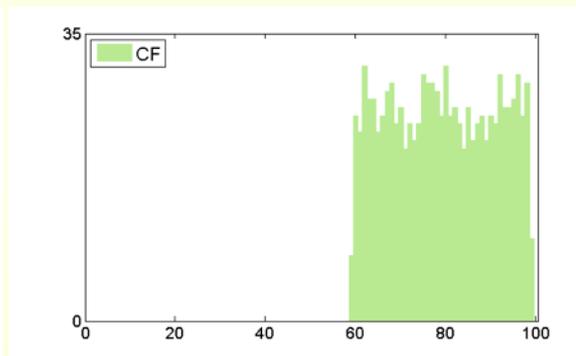
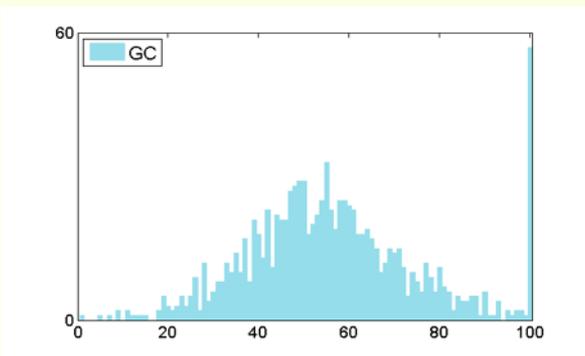
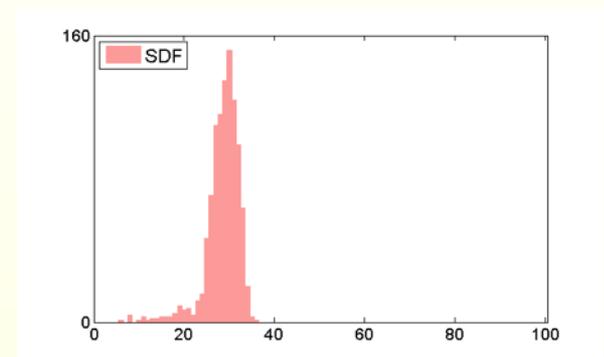
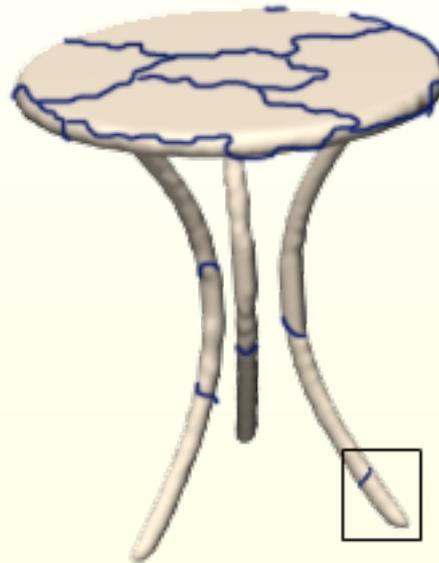
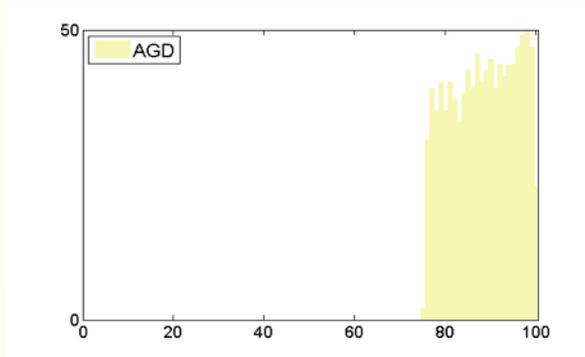
[Kalogerakis et al. 2010]



[Ben-Chen and Gostman 2008]

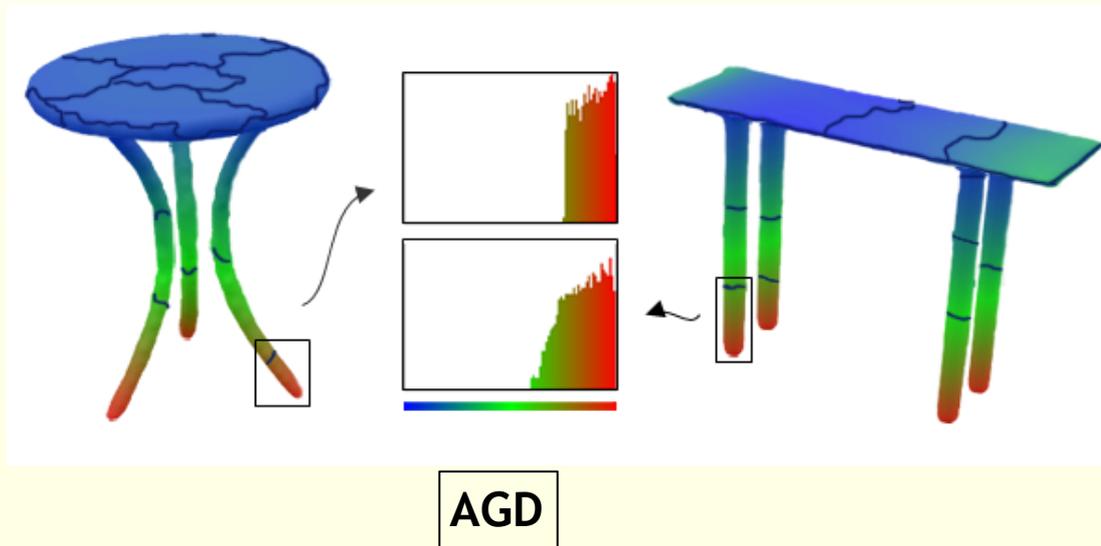
- Conformal factor

Patch features given by histograms



Co-segmentation via subspace clustering: key observation

- Idea 1: corresponding patches lie in common subspaces of some features (but not necessarily all features)



Co-segmentation

↓ *Key idea 1*

Subspace clustering

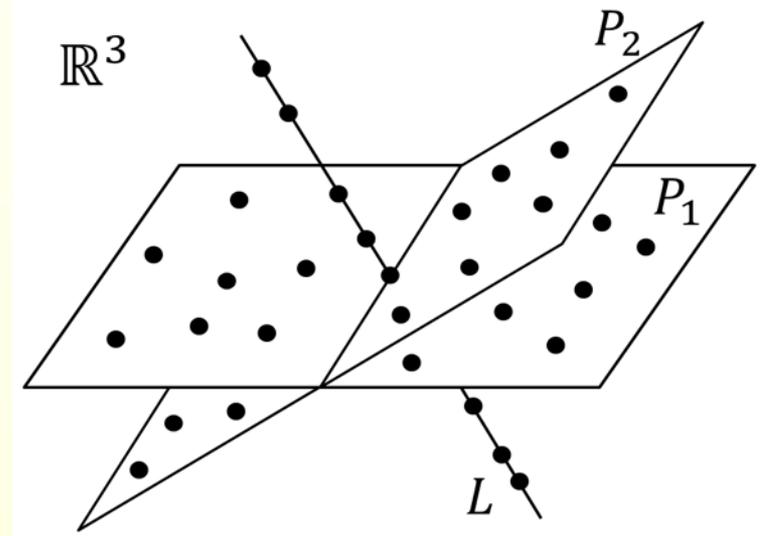
Key idea 2

- Idea 2: “two patches are considered as similar if they are similar in a subset of feature spaces and have large similarity measurement in each of these feature spaces”

Background: subspace clustering

[Vidal 2010]

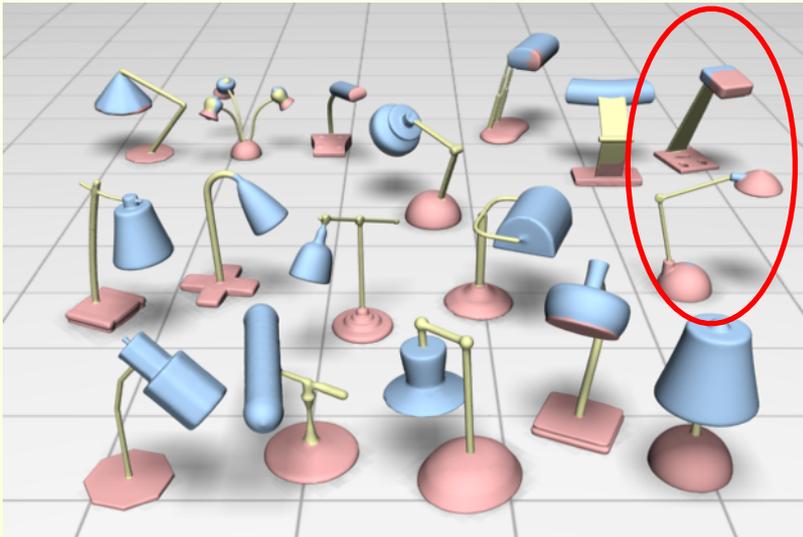
- Input:
 - high dimensional datasets having low intrinsic dimensions
 - $\{x_j\}_{j=1,\dots,N}, x_t \in \mathbb{R}^D$



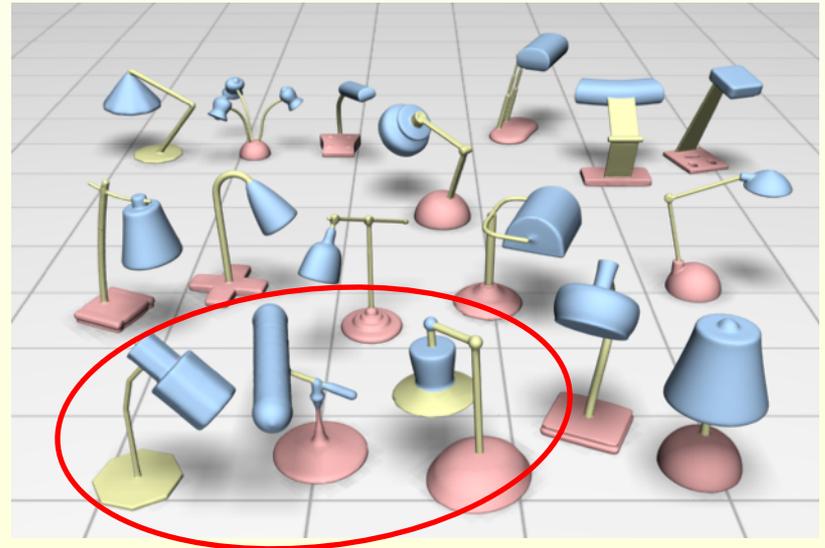
- Output:
 - multiple low-dimensional linear subspaces
 - L, P_1, P_2

Sub-space clustering: technical details

- Sub-space clustering performed via quadratic programming
- Smart penalty allows to find a subset of similar patch pairs, and a subset of features to measure their similarity
- Refine part boundaries using graph cuts



Our algorithm



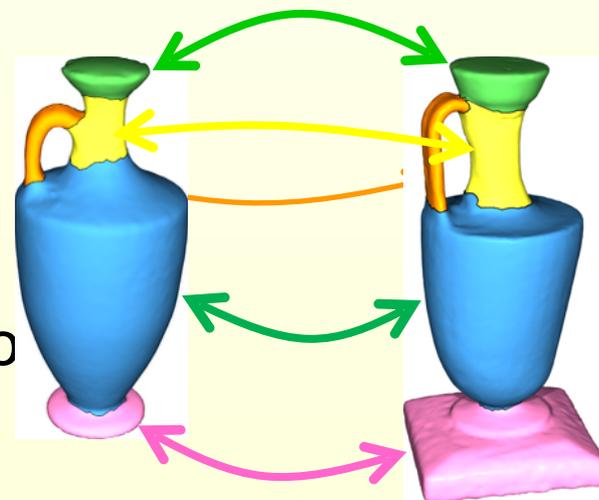
[Sidi et al. 2011]

Unsupervised co-segmentation as optimization problem

- “**Joint Shape Segmentation with Linear Programming**”
[Huang et al. 2011]
- Objective

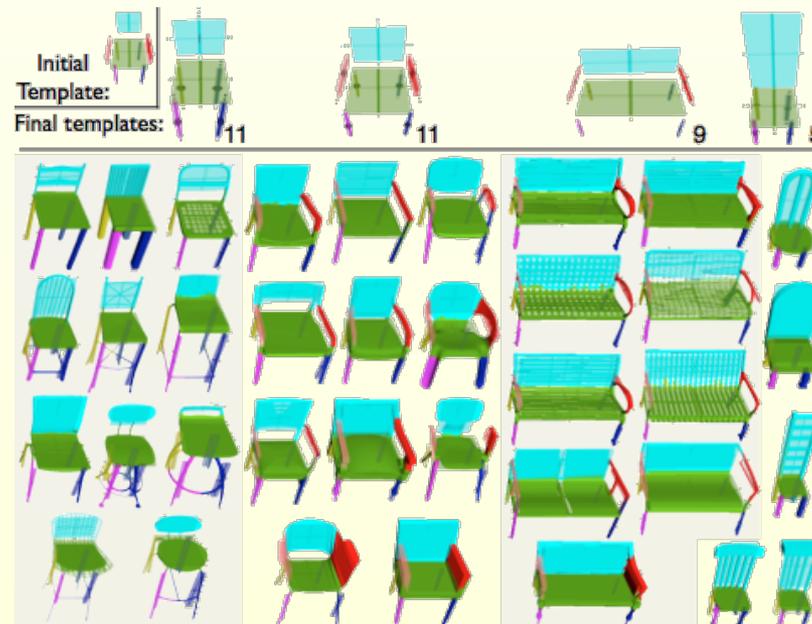
$$\max_{S_1, S_2} \text{score}(S_1) + \text{score}(S_2) + \text{consistency}(S_1, S_2)$$

- Outline
 - Segmentation parameterization
 - Segmentation score
 - Consistency score
 - 0-1 linear programming formulation



Learning part-based templates

- “Learning Part-based Templates from Large Collections of 3D Shapes” [Kim et al. 2013]
 - Jointly solve for deformations, part segmentation and inter-model correspondence

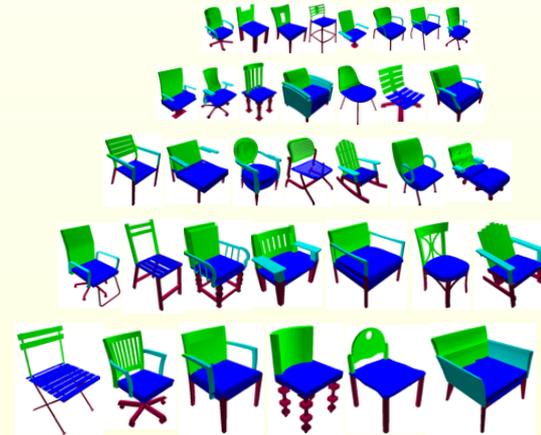


Semi-supervised segmentation

ACTIVE LEARNING WITH HUMAN IN THE LOOP

Active segmentation learning

- Data-driven method perform well when
 - They rely on big datasets
 - With detailed annotations (e.g., parts labels)
 - Annotations are curated
- Manual annotation
 - Time consuming
 - ▶ Not suitable for large collections
- Automatic annotation
 - Never perfectly reliable



Manual Annotation

Time consuming

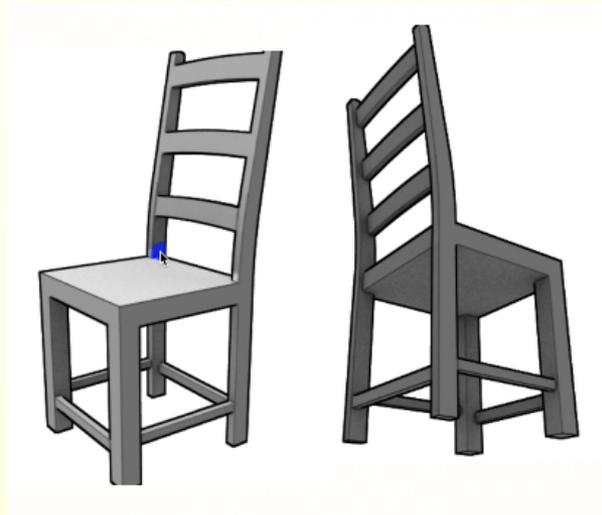
SHAPENET

3D Warehouse

yobi 3D

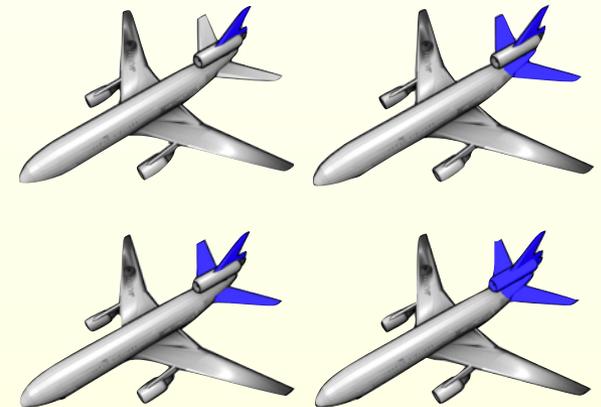
TURBOSQUID

Low efficiency



Not satisfactory accuracy

Airplane tail



- Millions of 3D models
- Constantly evolving

- 3D annotation takes lots of time

- Semantic ambiguity

Manual Annotation

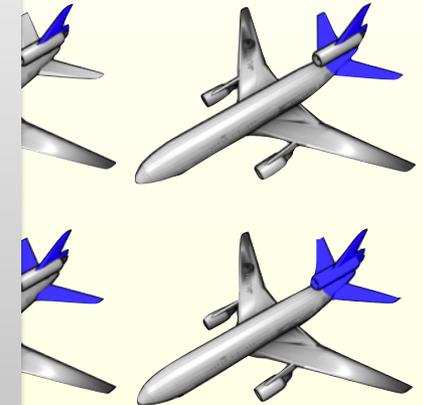
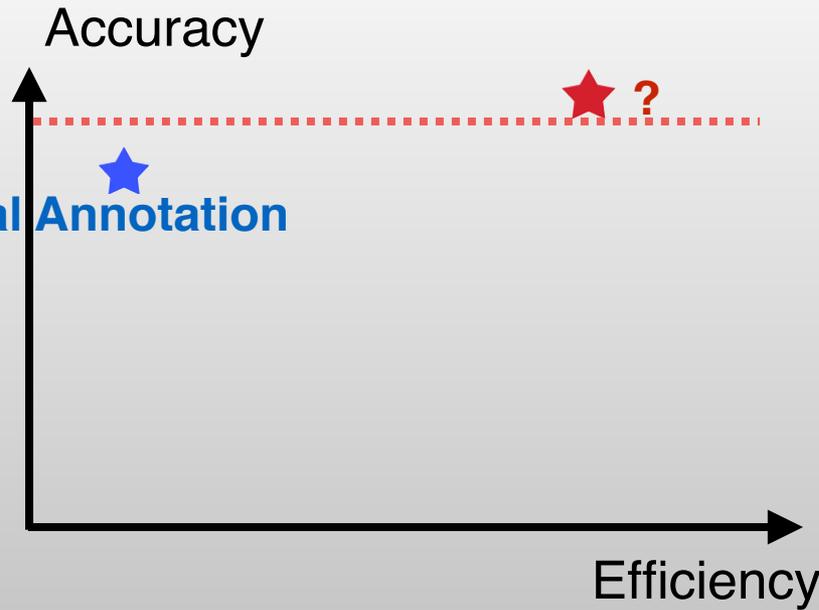
Time con

satisfactory
accuracy

airplane tail

SHAPE

Manual Annotation



3D Wa

yok

TURBO

- Millions of 3D models
- Constantly evolving

- 3D annotation takes lots of time

- Semantic ambiguity

Manual Annotation

Time con

satisfactory
accuracy

airplane tail

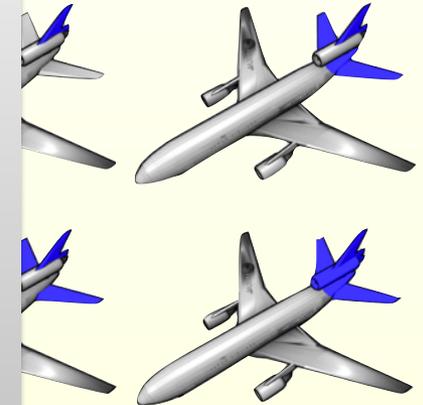
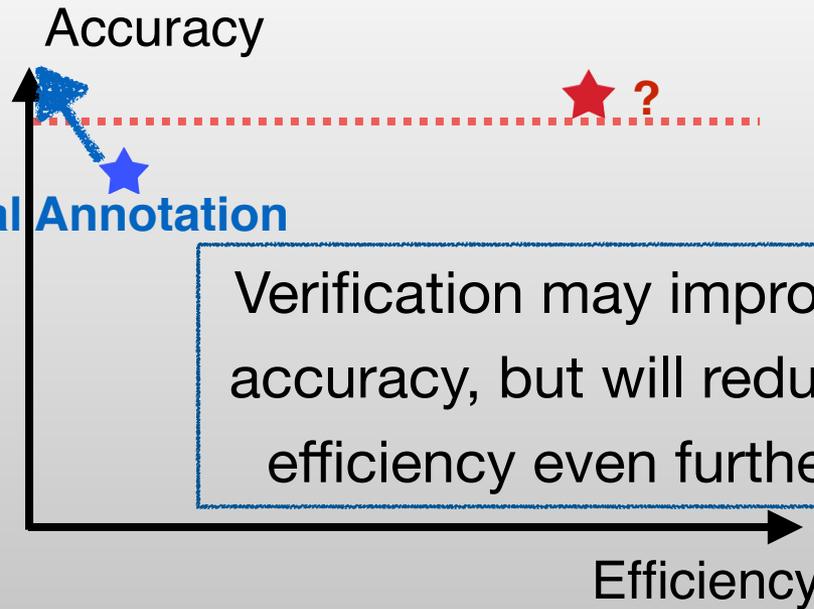
SHAPE

3D Wa

yok

TURBO

Manual Annotation



- Millions of 3D models
- Constantly evolving

- 3D annotation takes lots of time

- Semantic ambiguity

Automatic annotation

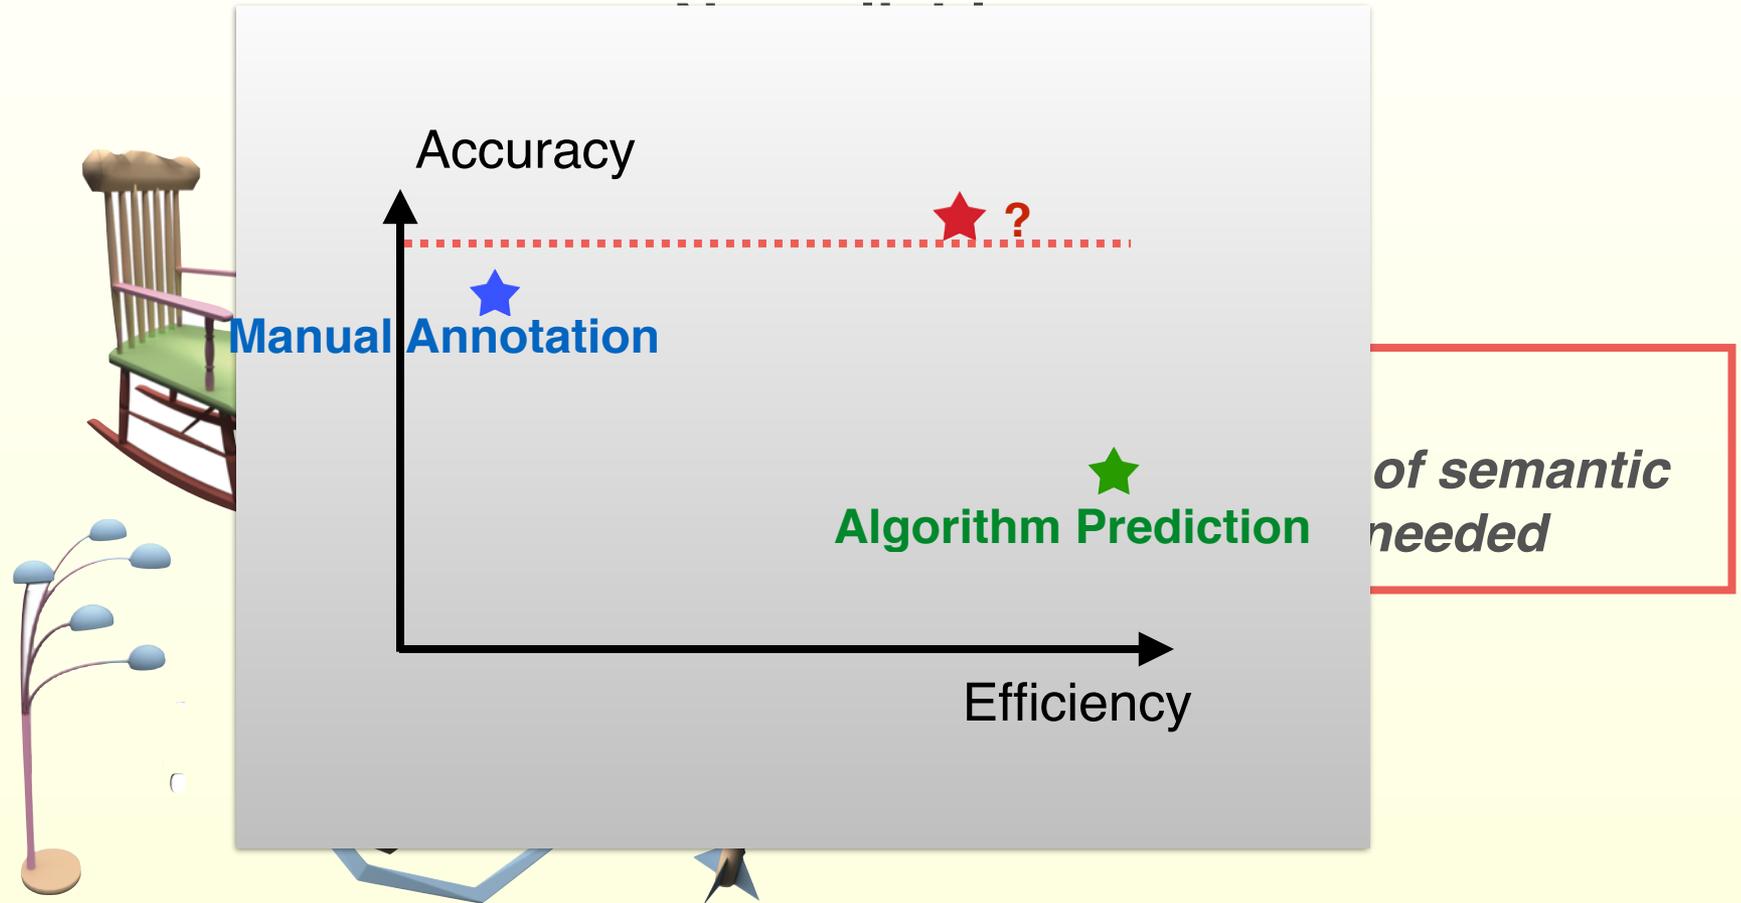
Automatic annotation

Not reliable

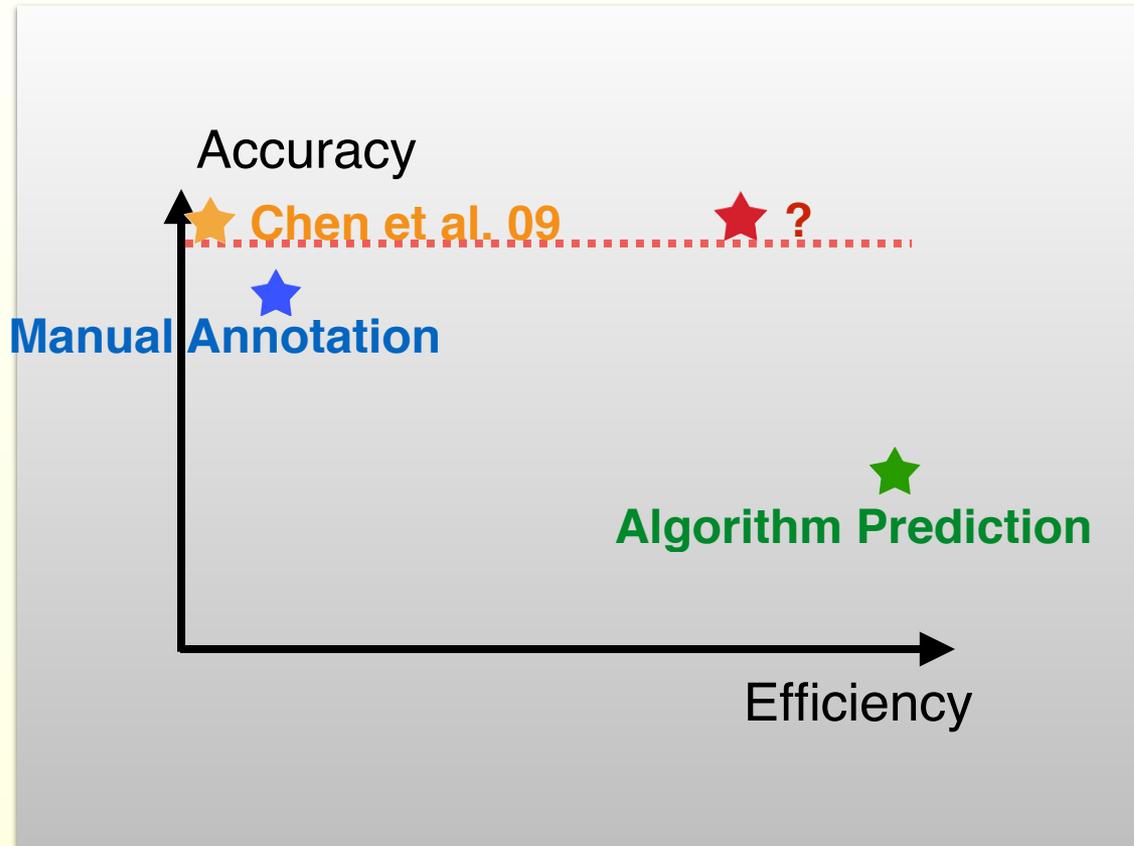


- *Large variation*
- *Different levels of semantic knowledge are needed*

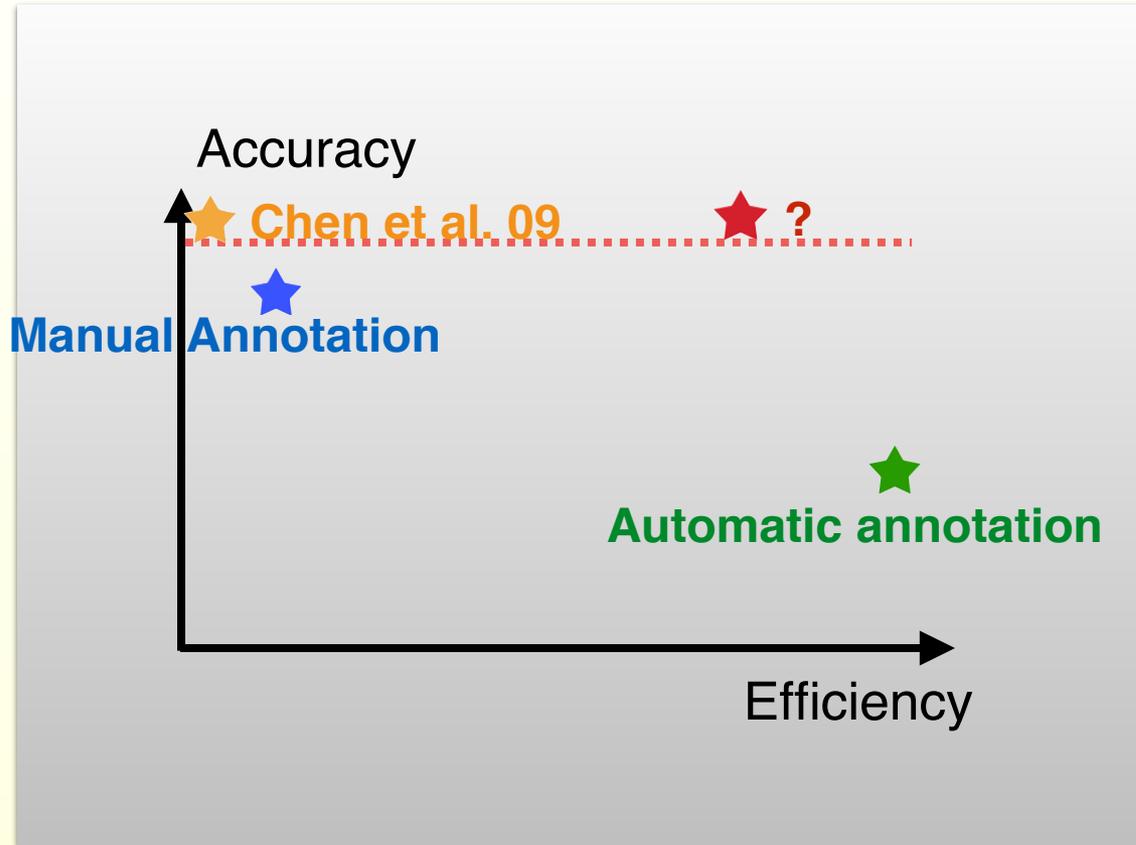
Automatic annotation



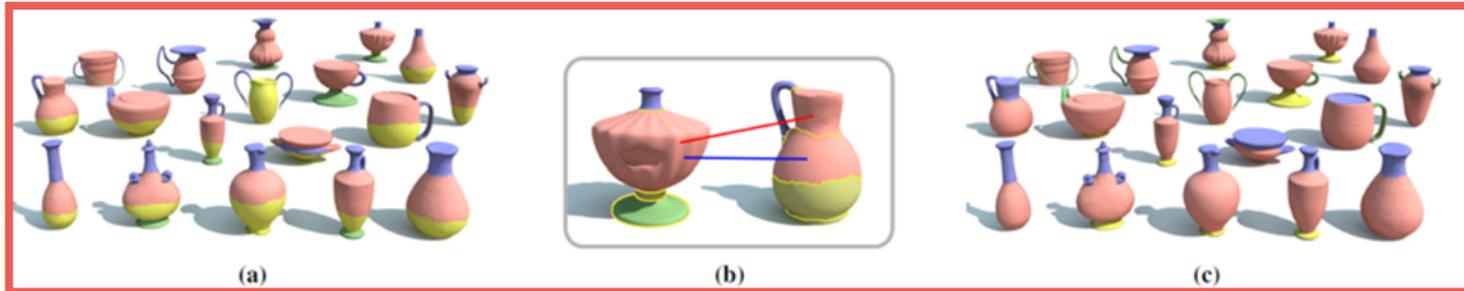
Automatic annotation



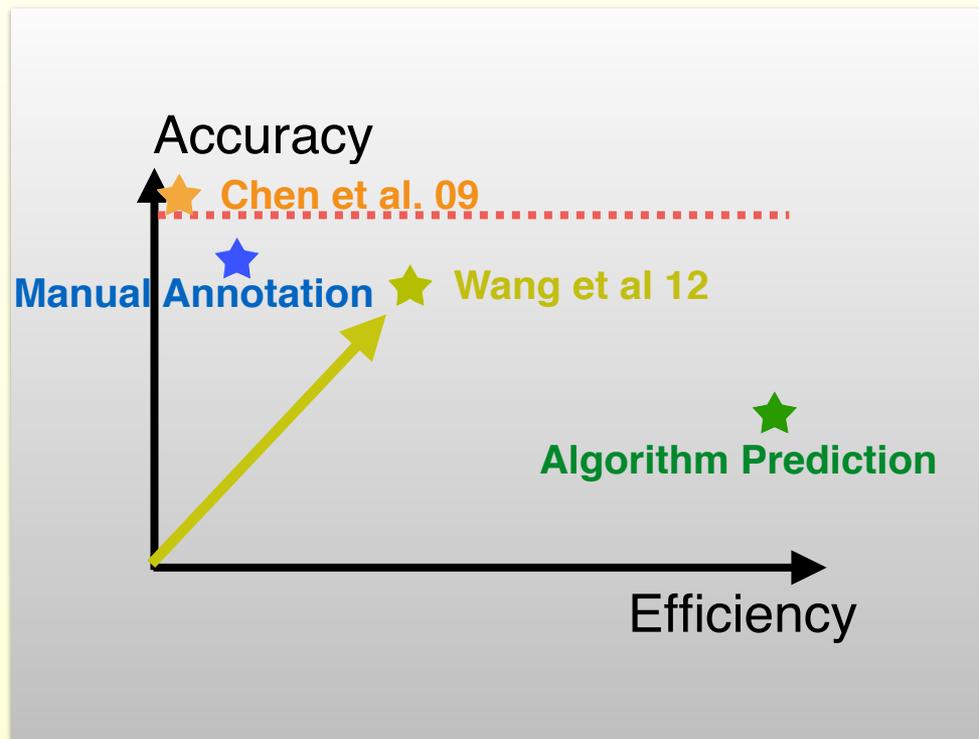
Automatic annotation



Problem Definition

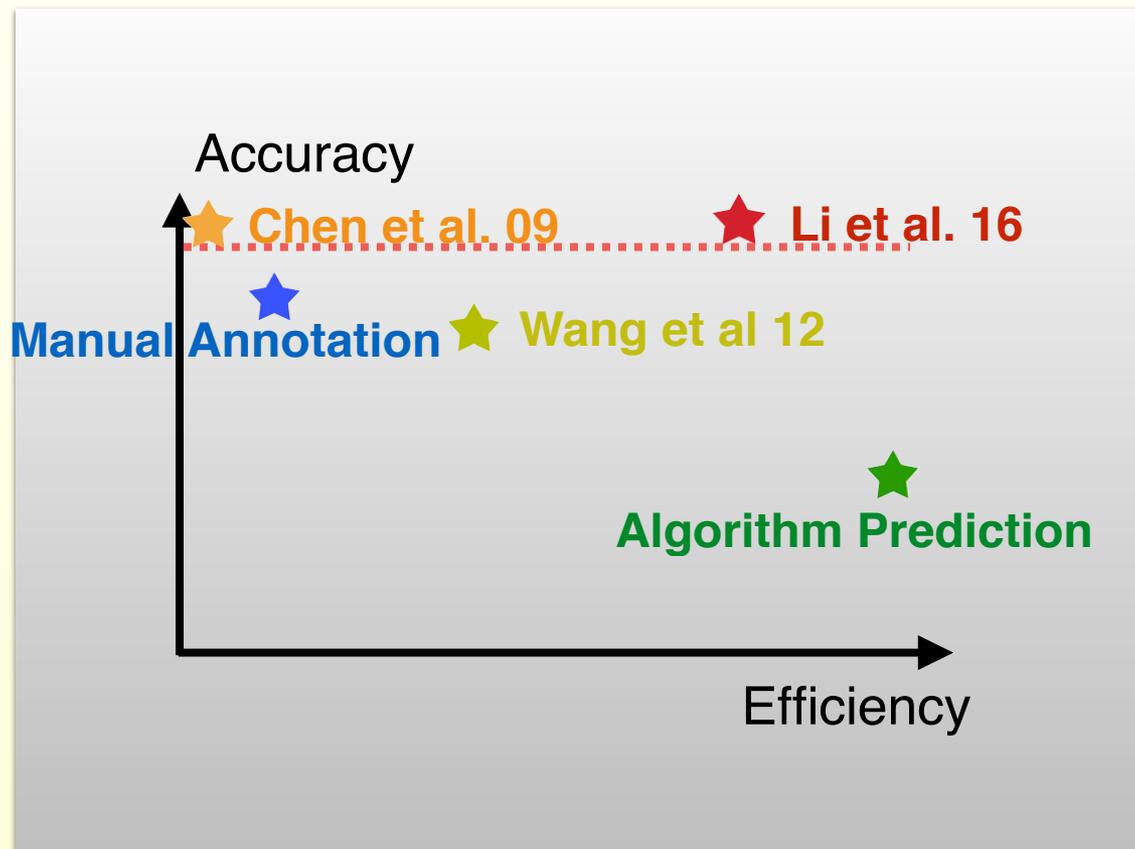


Active Co-Analysis of a Set of Shapes [Wang et al. 12]

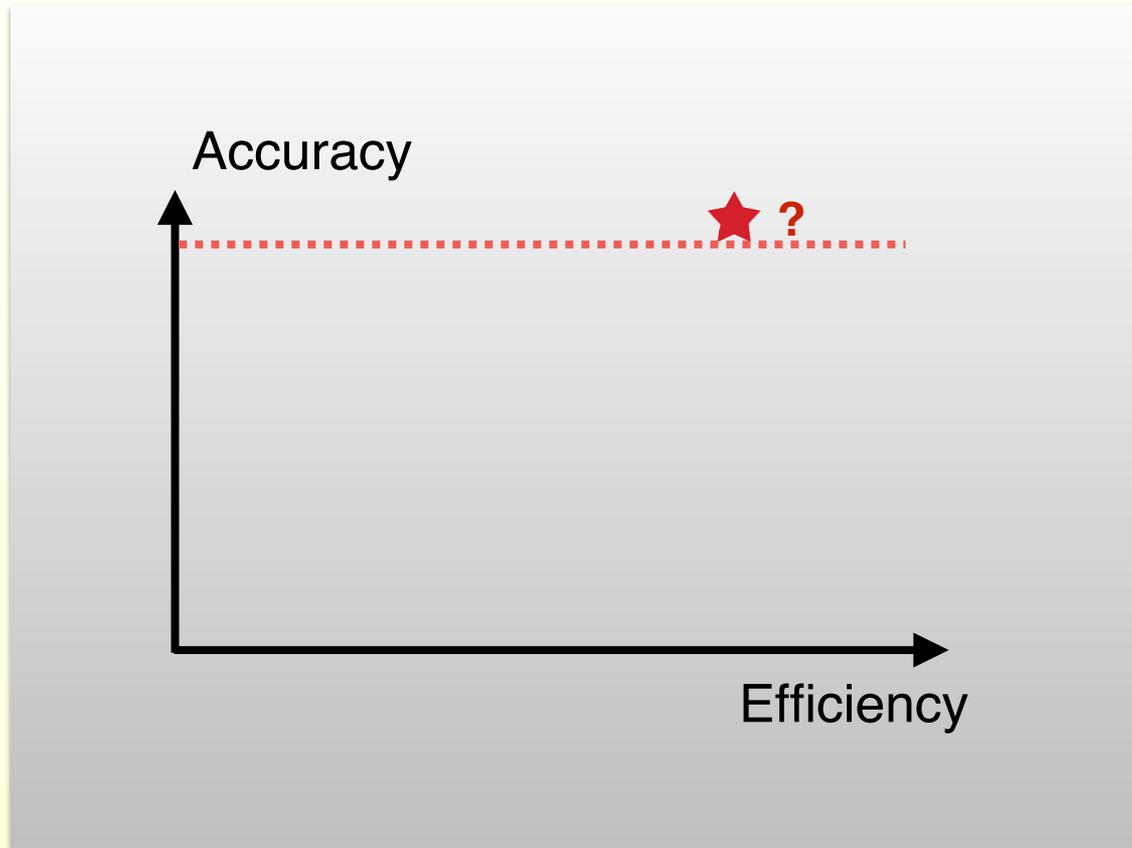


Problem Definition

- “A Scalable Active Framework for Region Annotation in 3D Shape Collections” [Li et al. 2016]

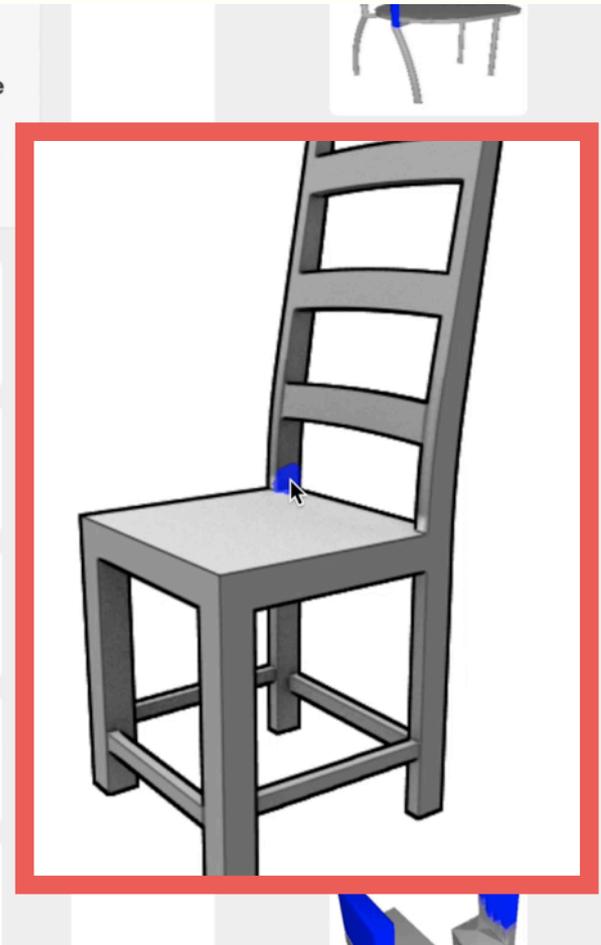


Problem Definition



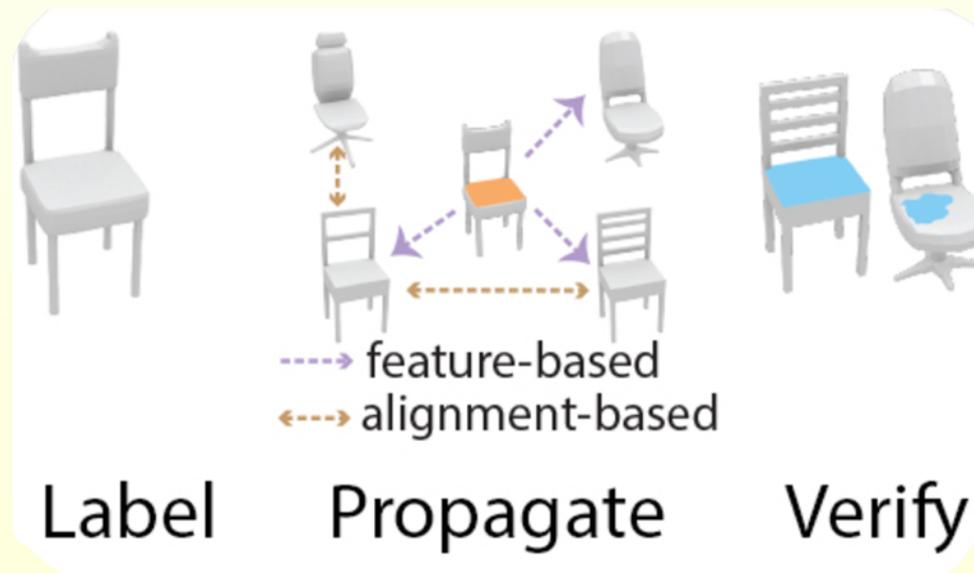
Main idea: verification is much more efficient than labeling

Instruction: Please pick up the images whose **back** is **NOT** highlighted correctly. Please use the example images as a reference. Remember to click on the bad images! Notice images **without back** and at the same time without any part highlighted should be treated as good images and you should **NOT** click on them. Images **with back** but at the same time without any part highlighted should be treated as bad images and you should click on them.

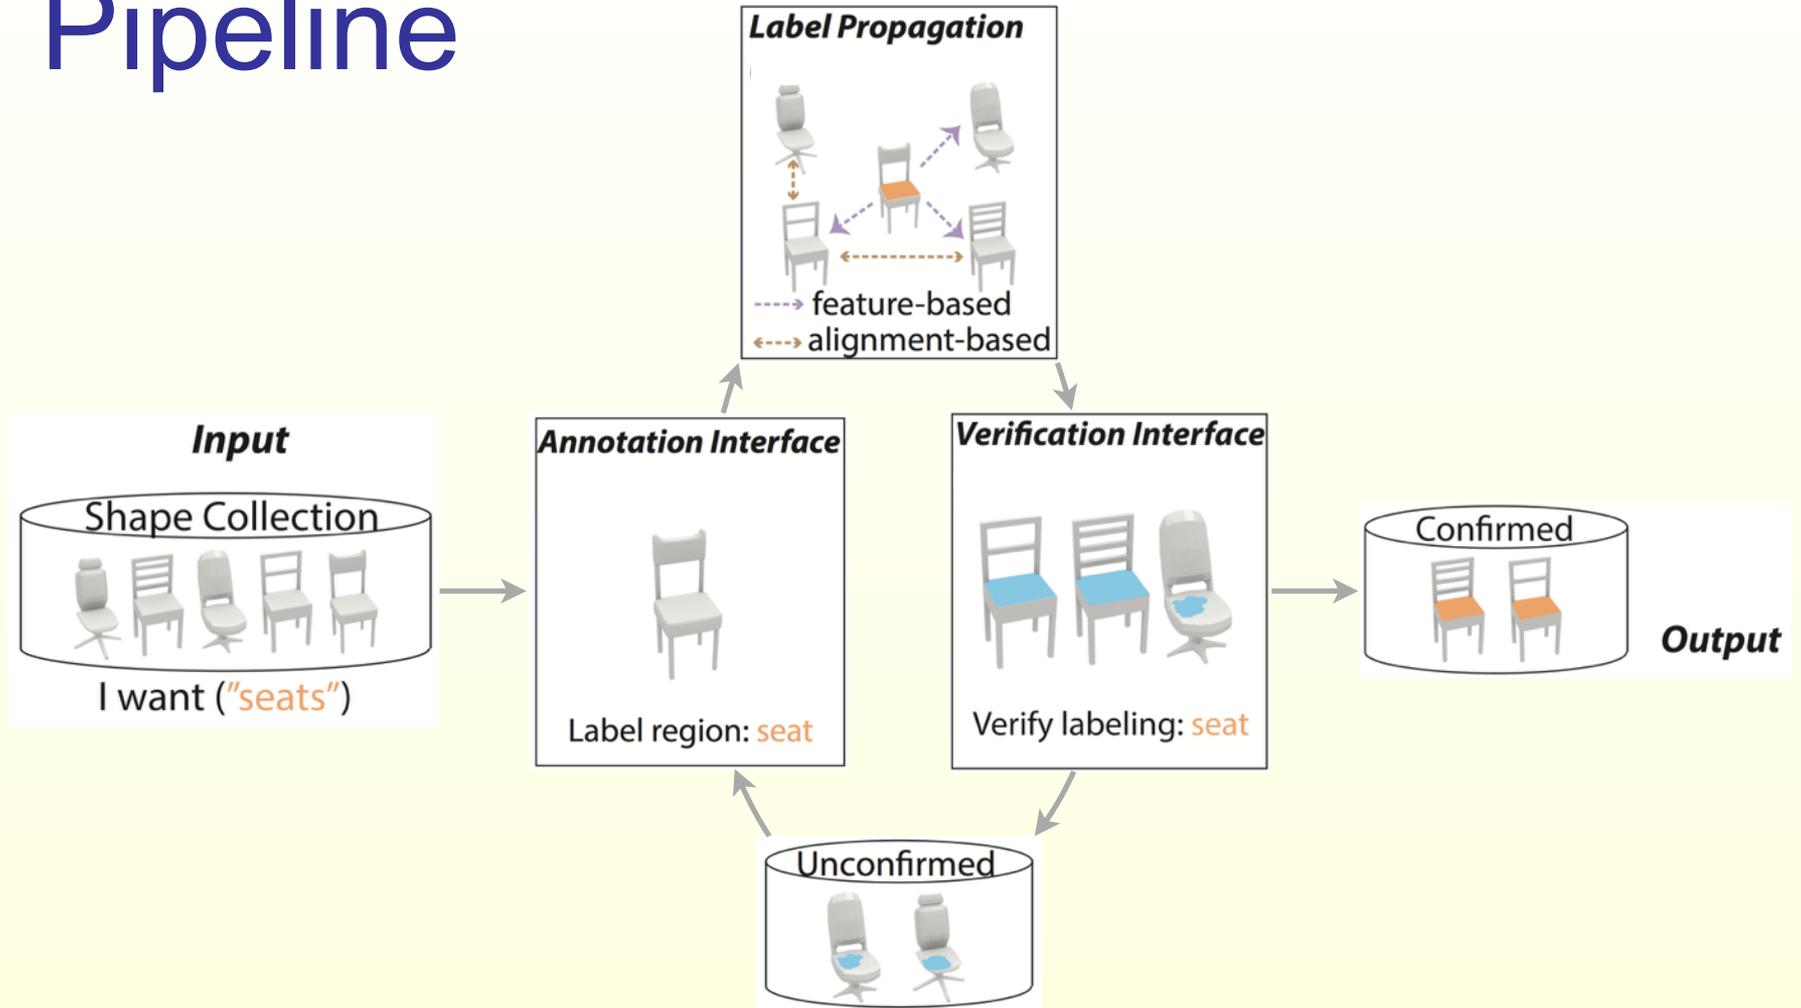


Approach overview

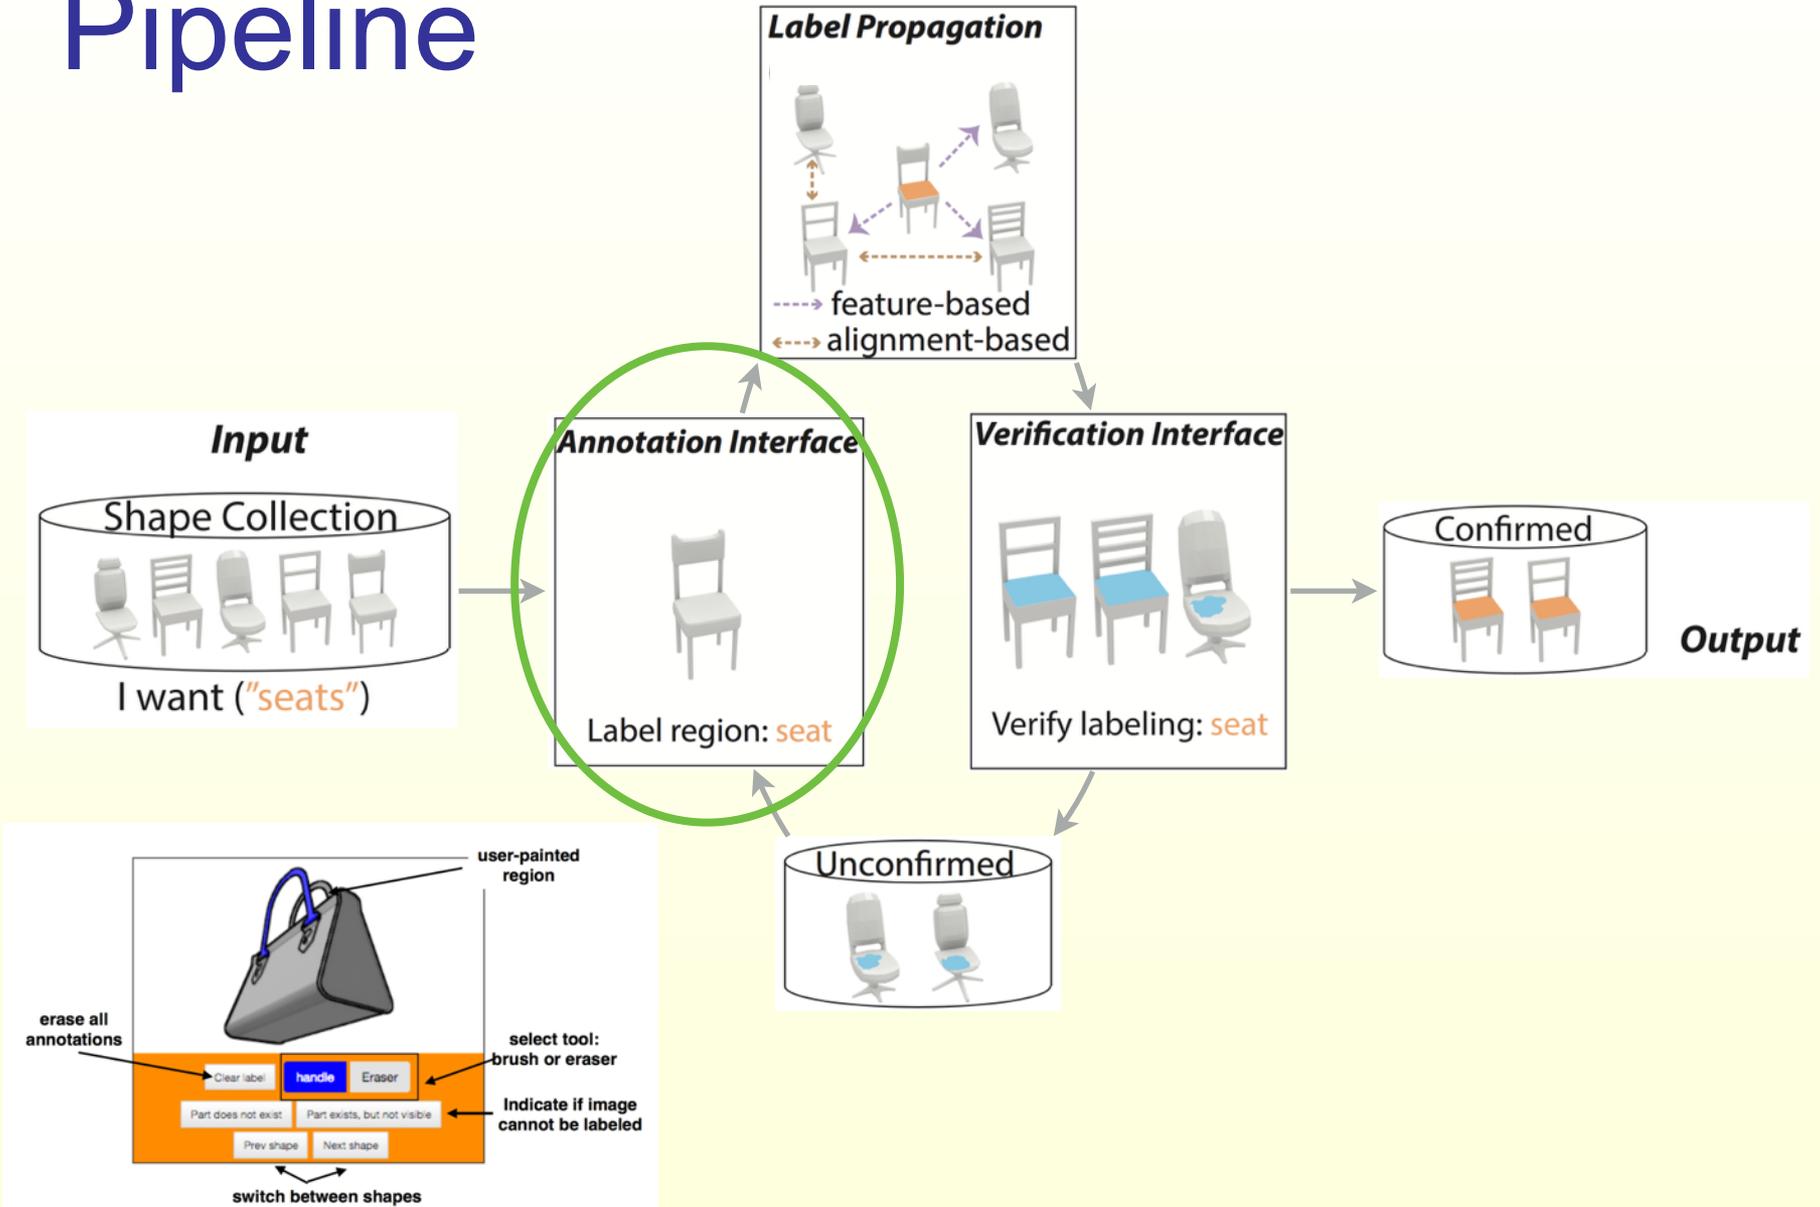
- Unified framework which combines
 - human annotation
 - automatic prediction by correspondence propagation
 - batch verification by expert



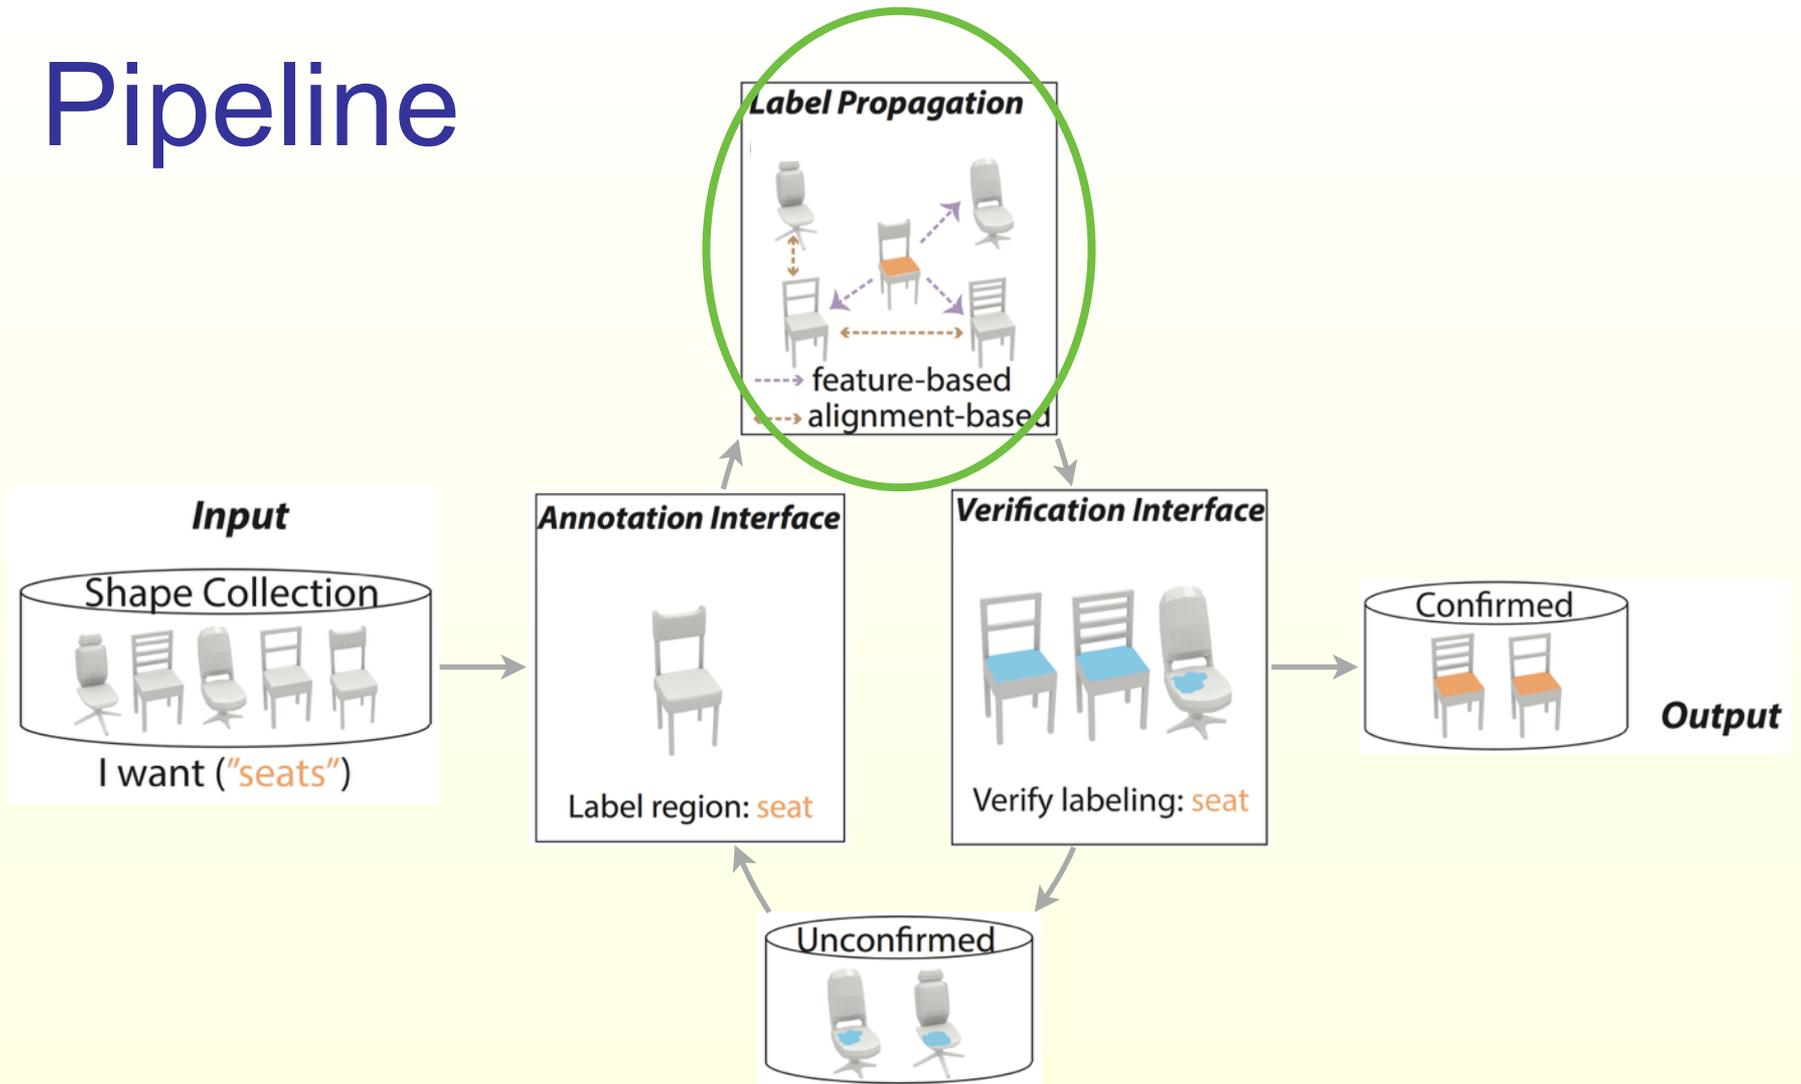
Pipeline



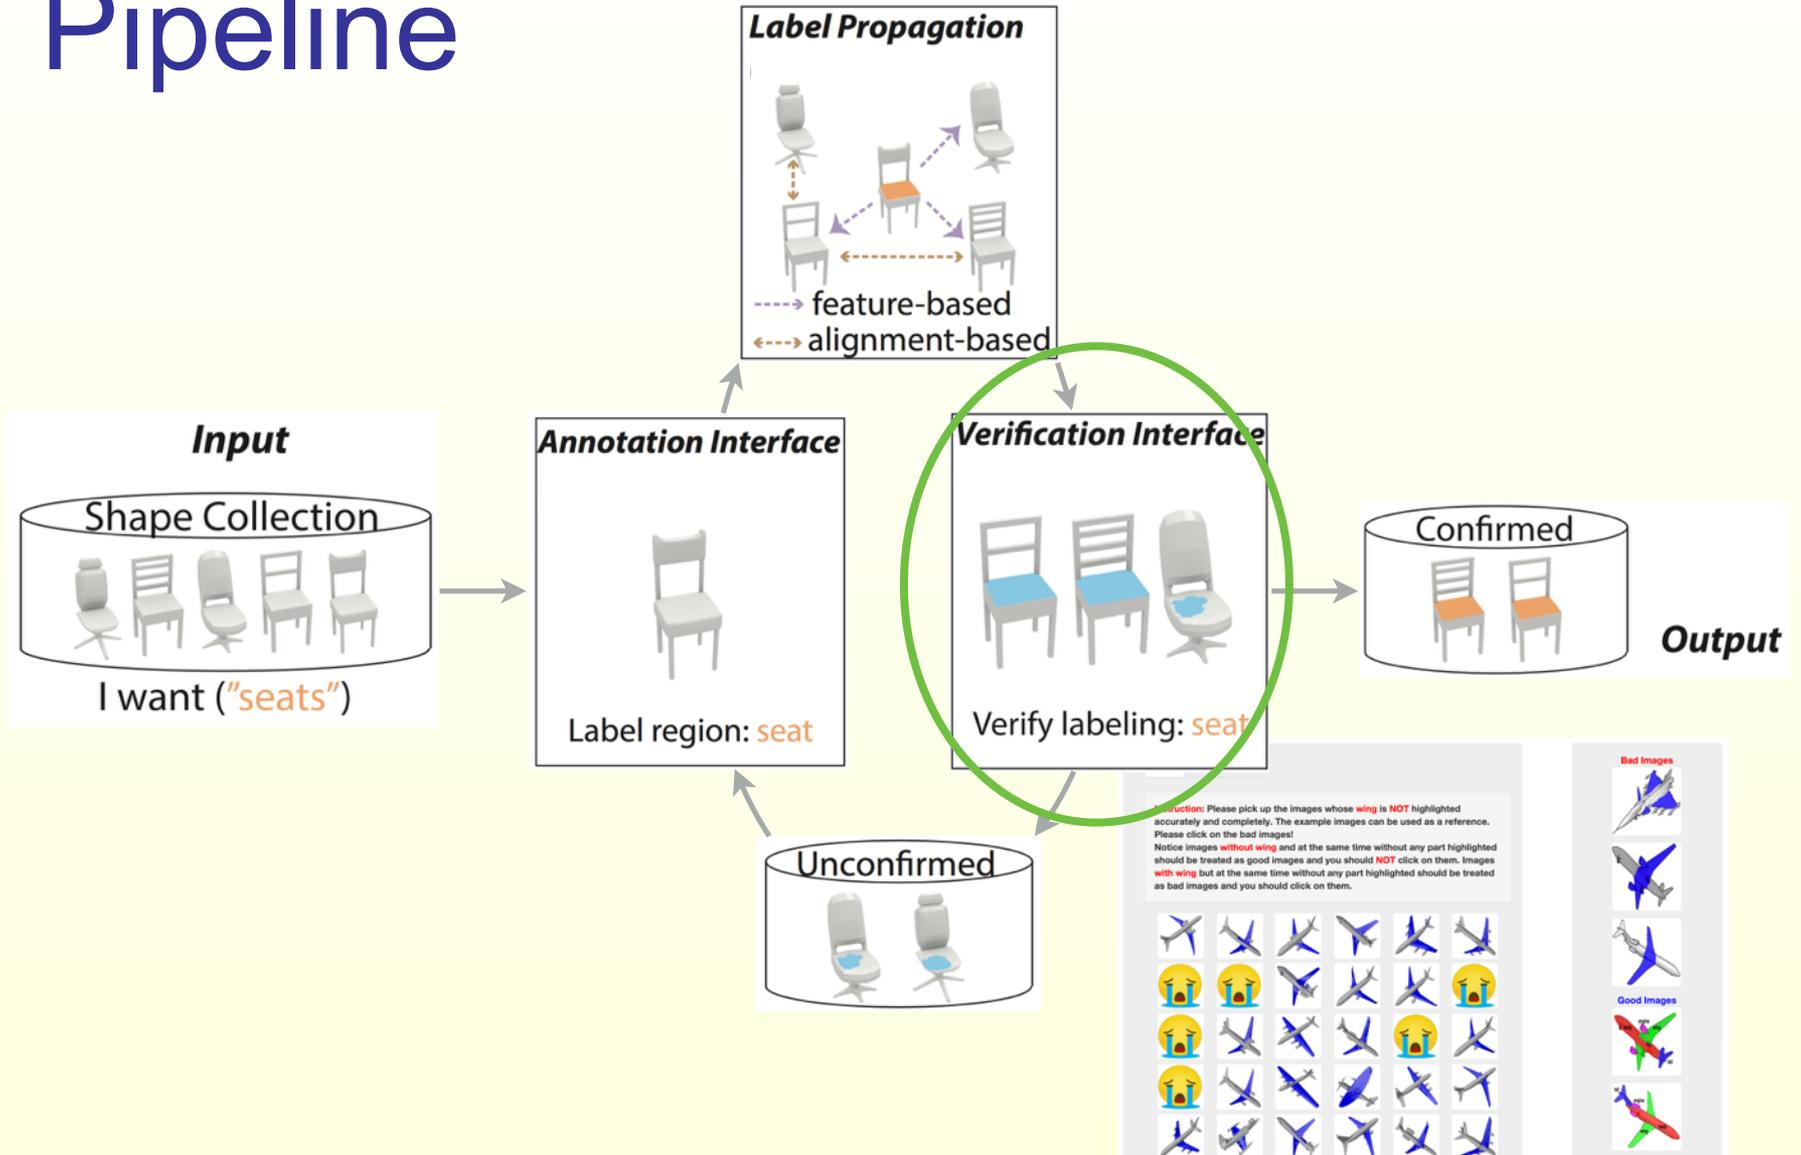
Pipeline



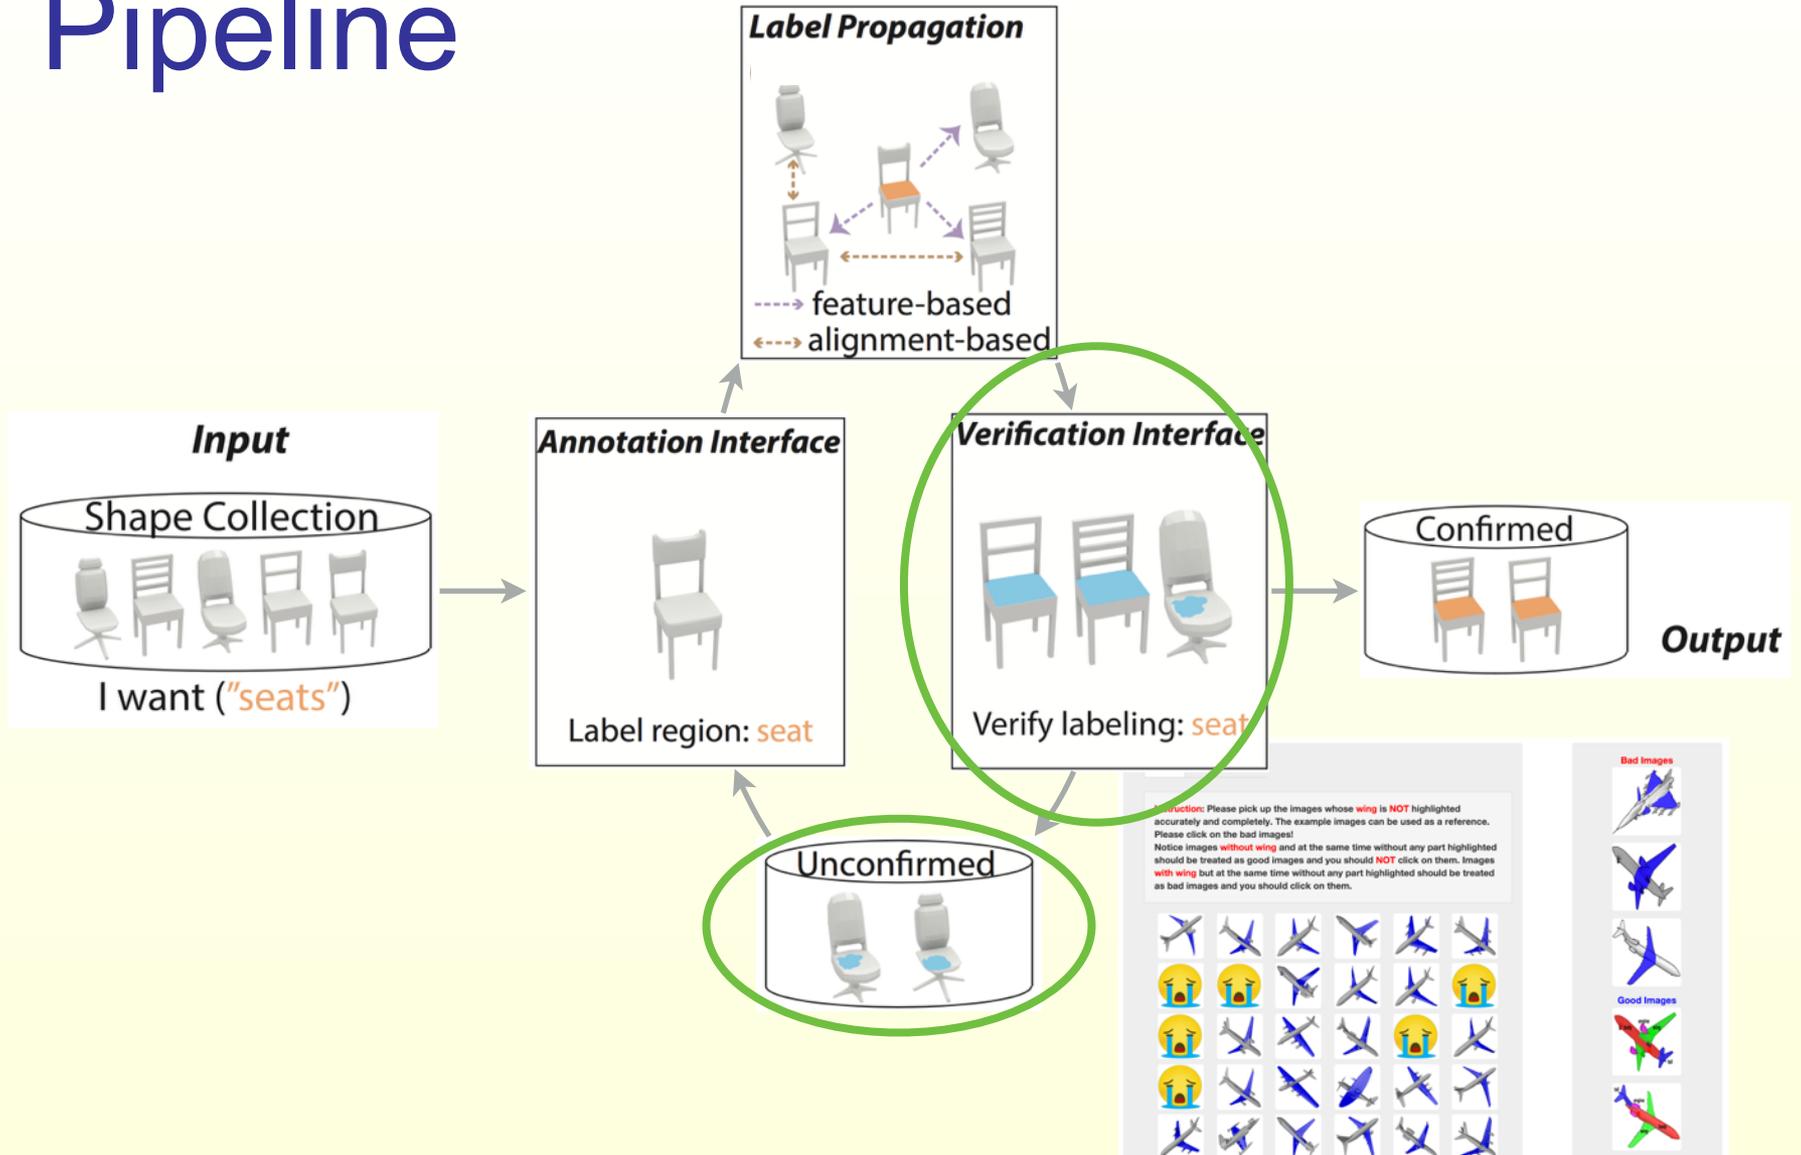
Pipeline



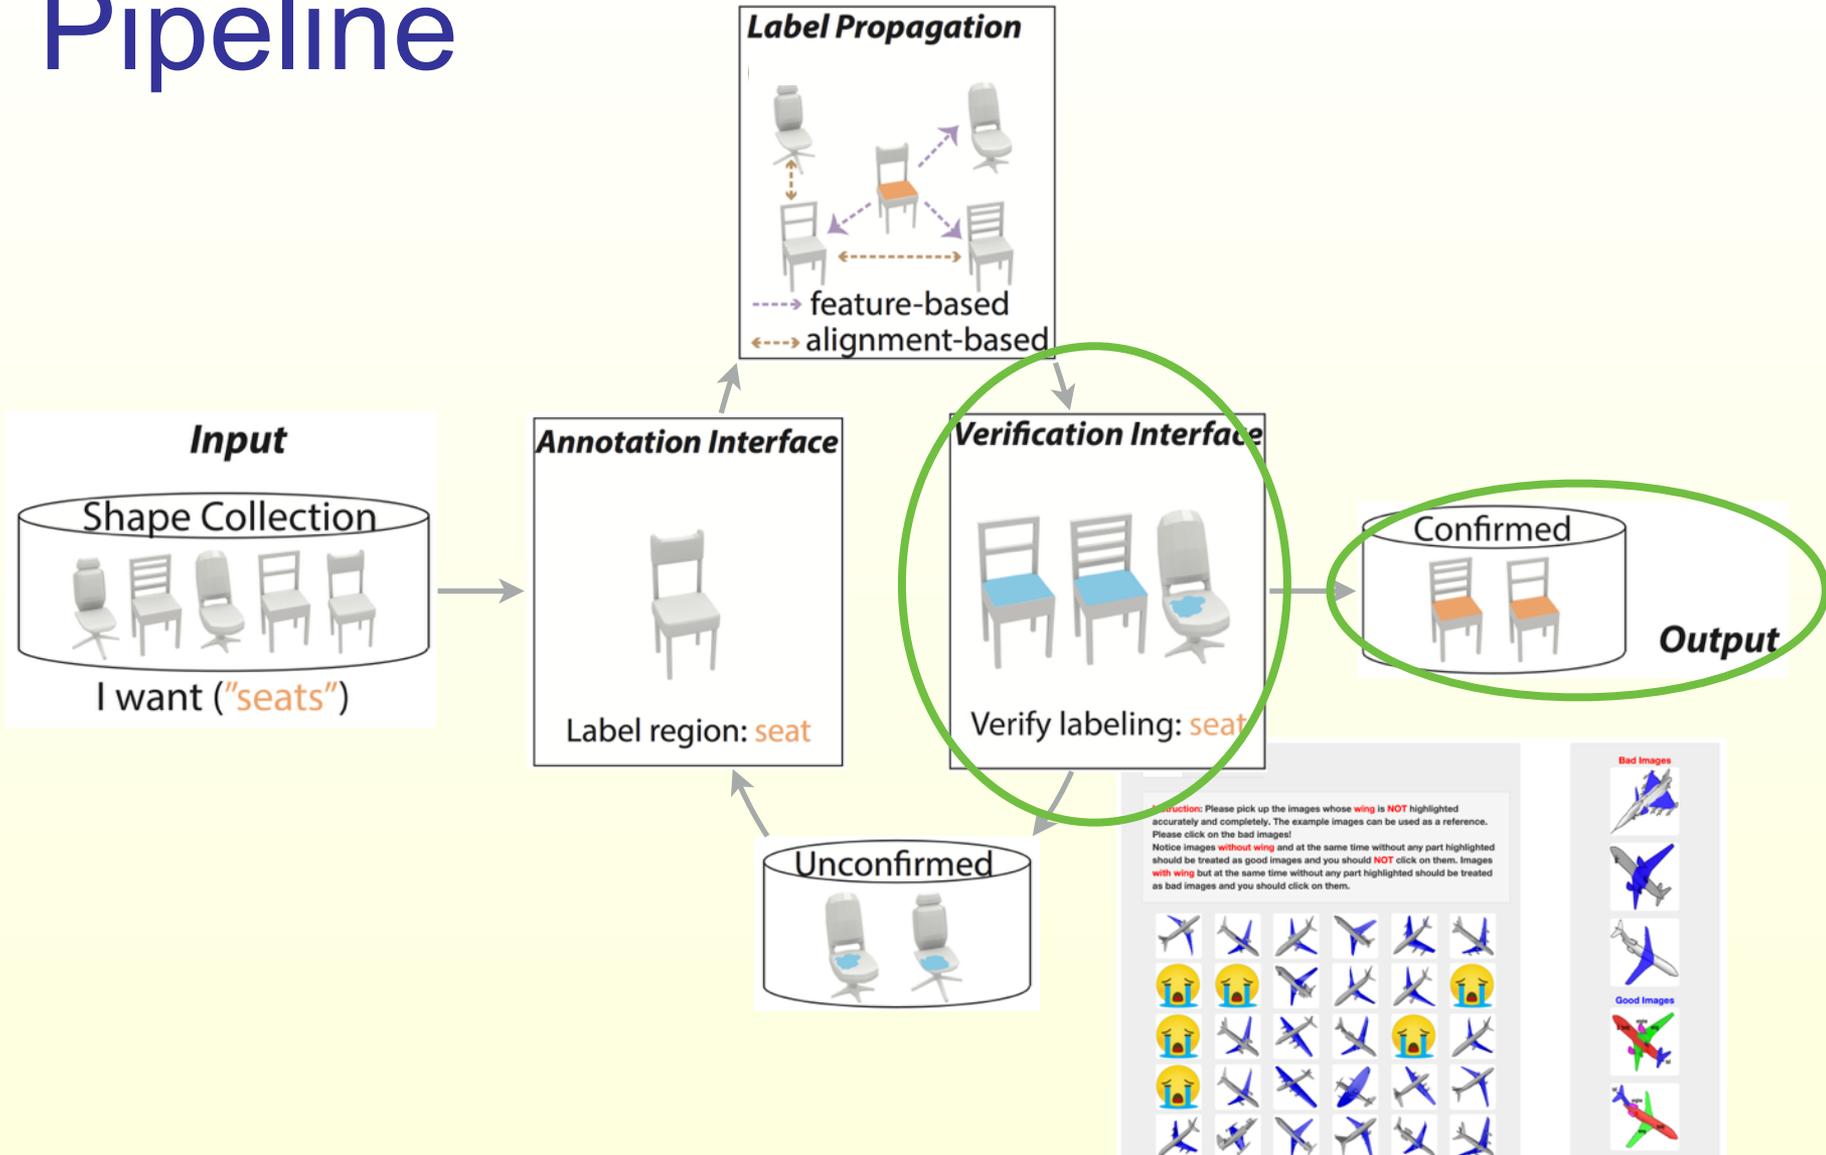
Pipeline



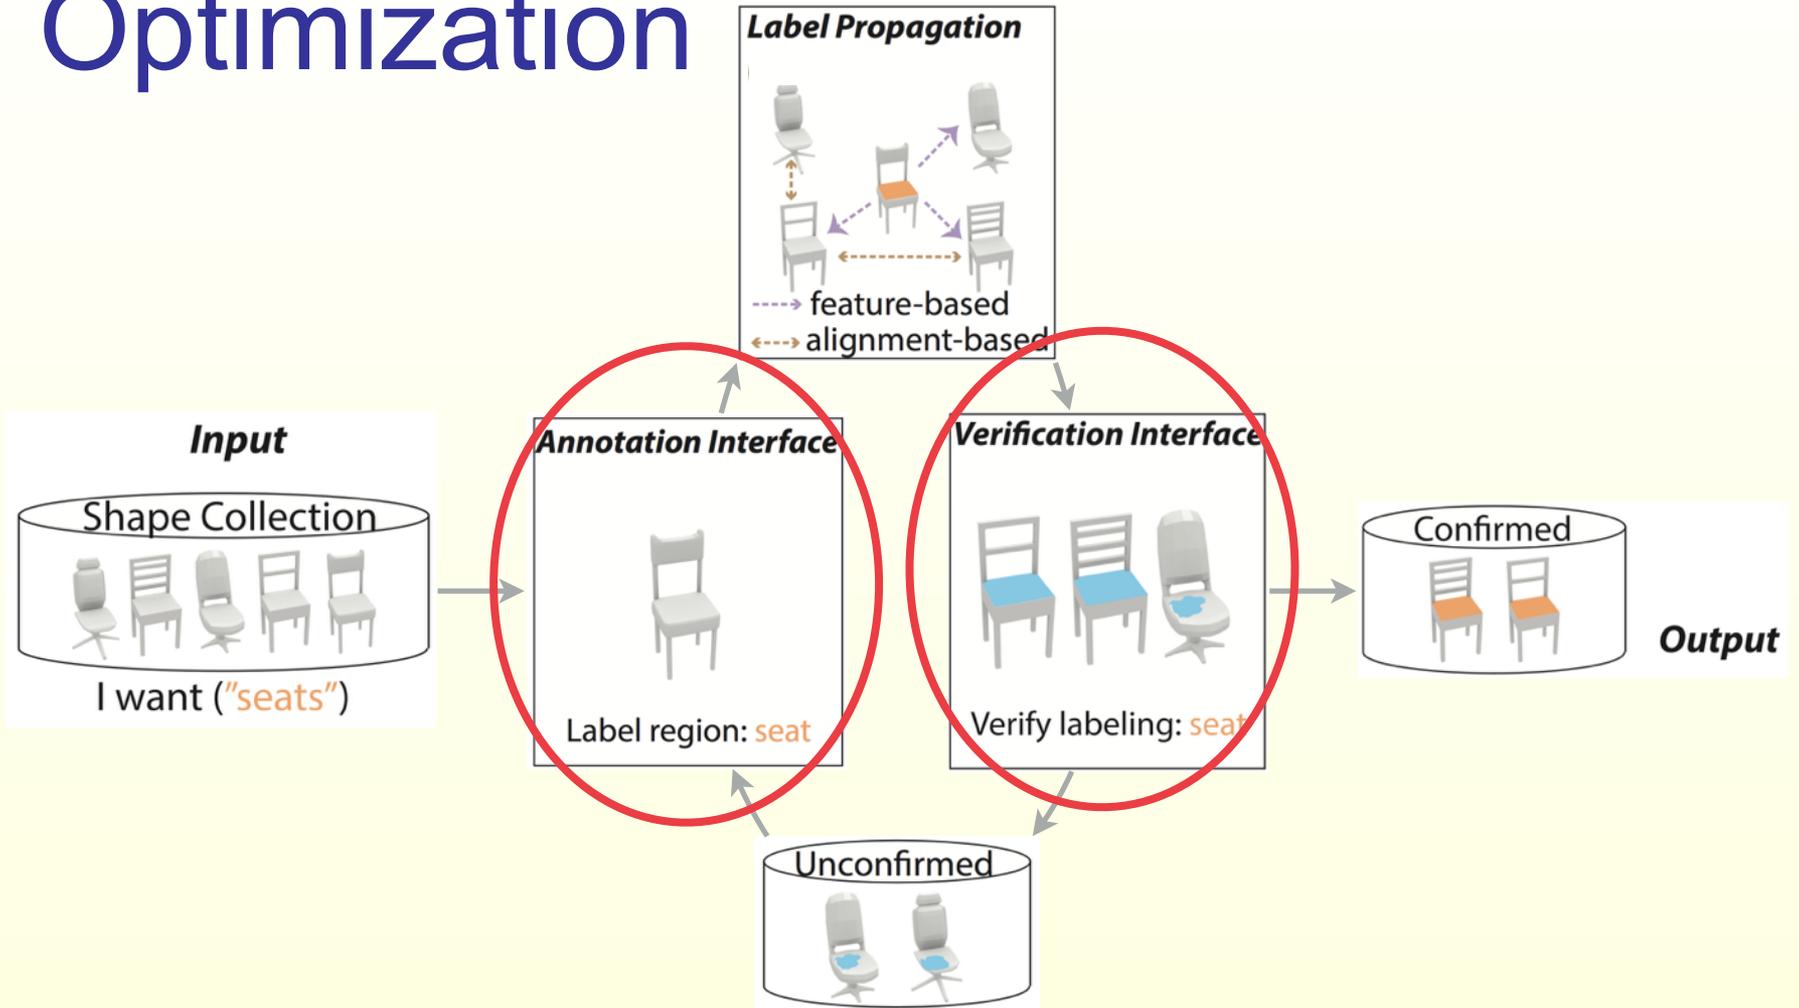
Pipeline



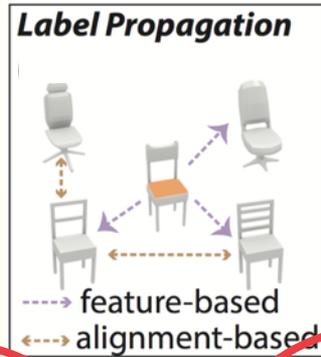
Pipeline



Optimization



Optimization



Input

Shape Collection

Annotation Interface

Verification Interface

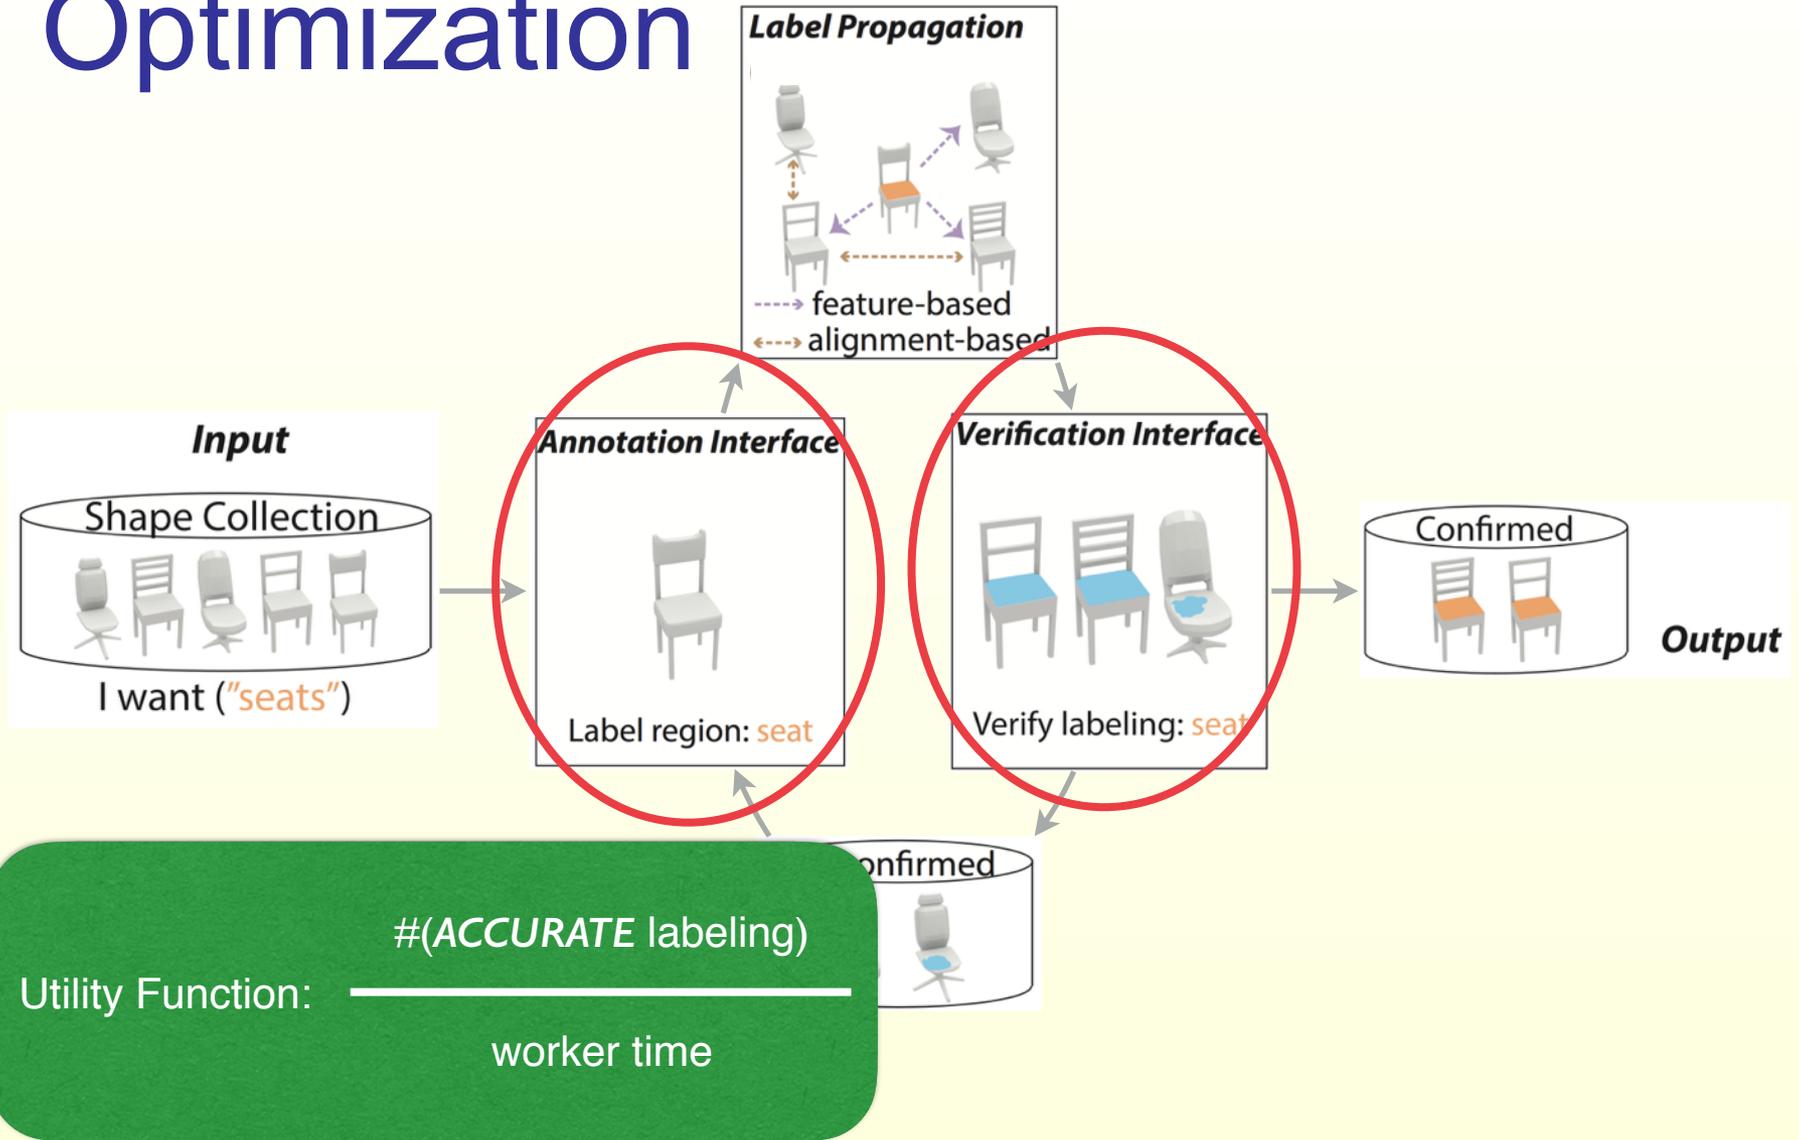
How to issue **Annotation** and **Verification** tasks

to optimize the **Utility Function**

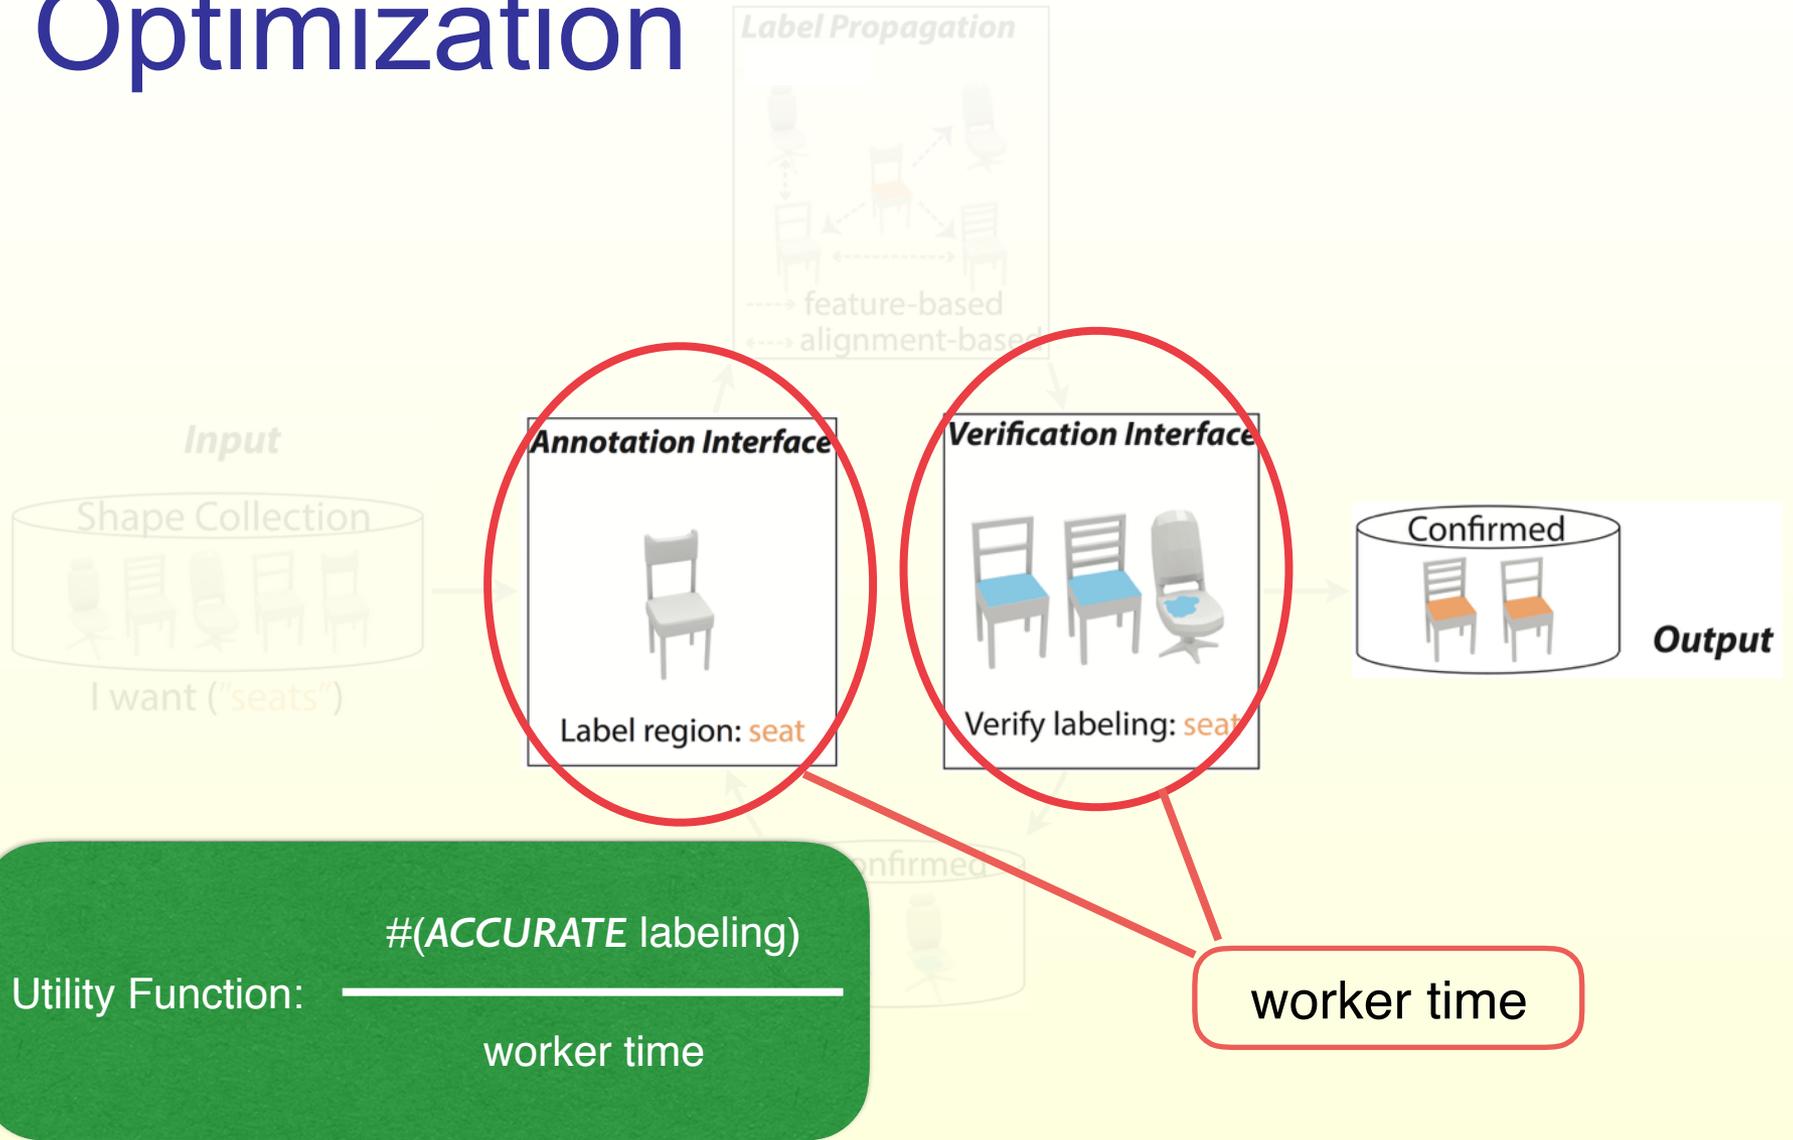
Unconfirmed

Utility Function: $\frac{\#(\text{ACCURATE labeling})}{\text{worker time}}$

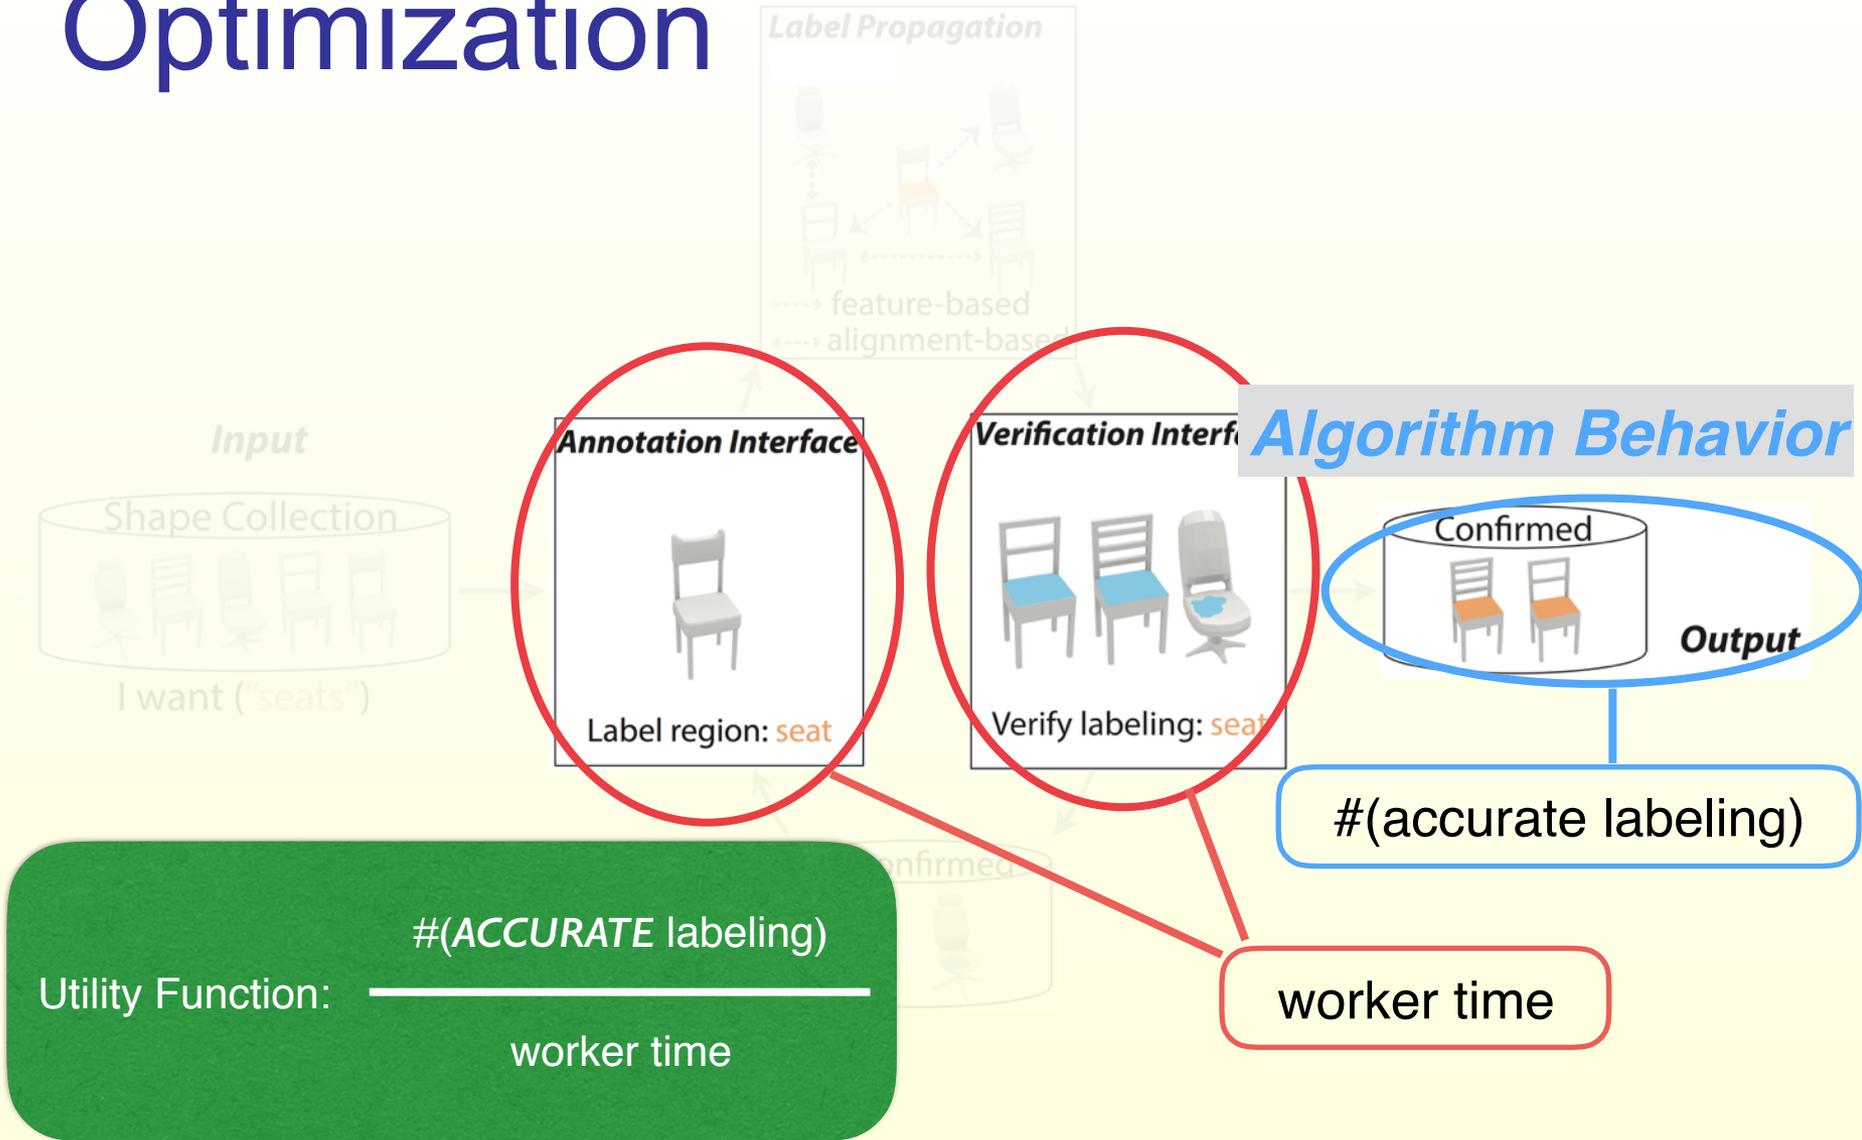
Optimization



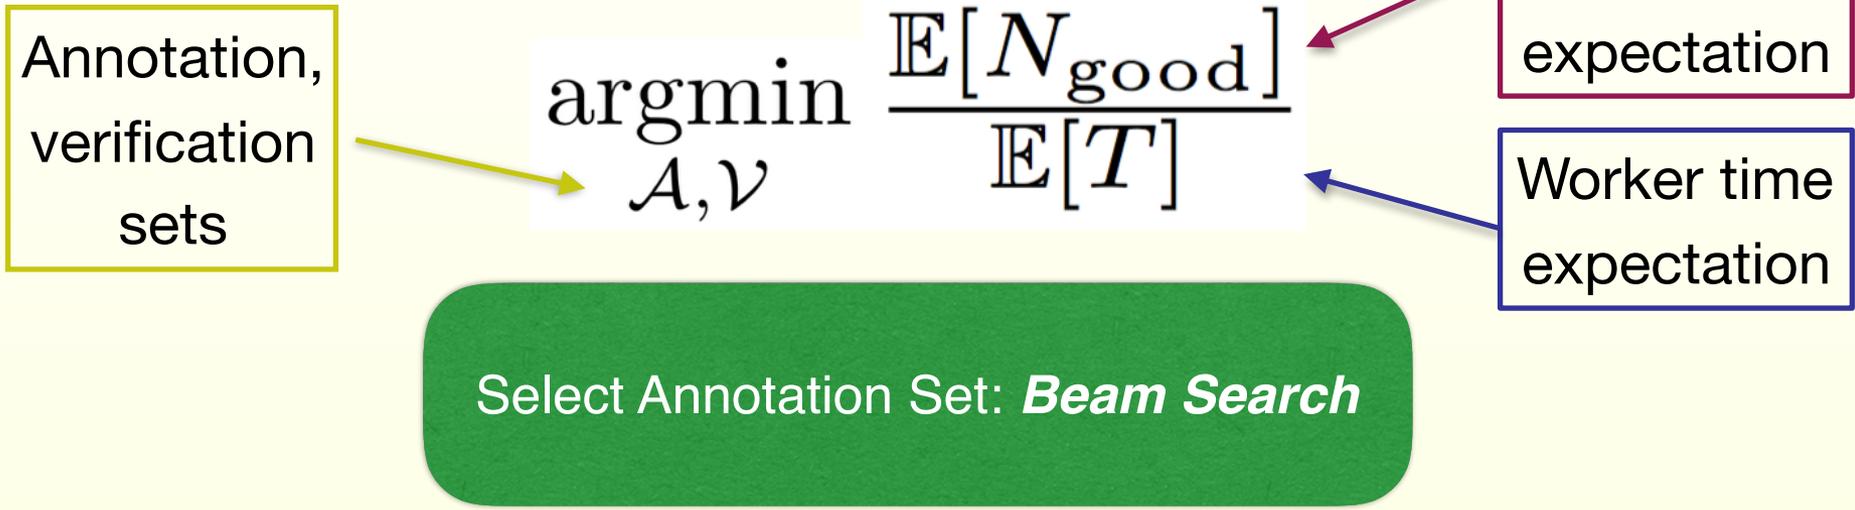
Optimization



Optimization



Utility Function



Utility Function

Annotation,
verification
sets

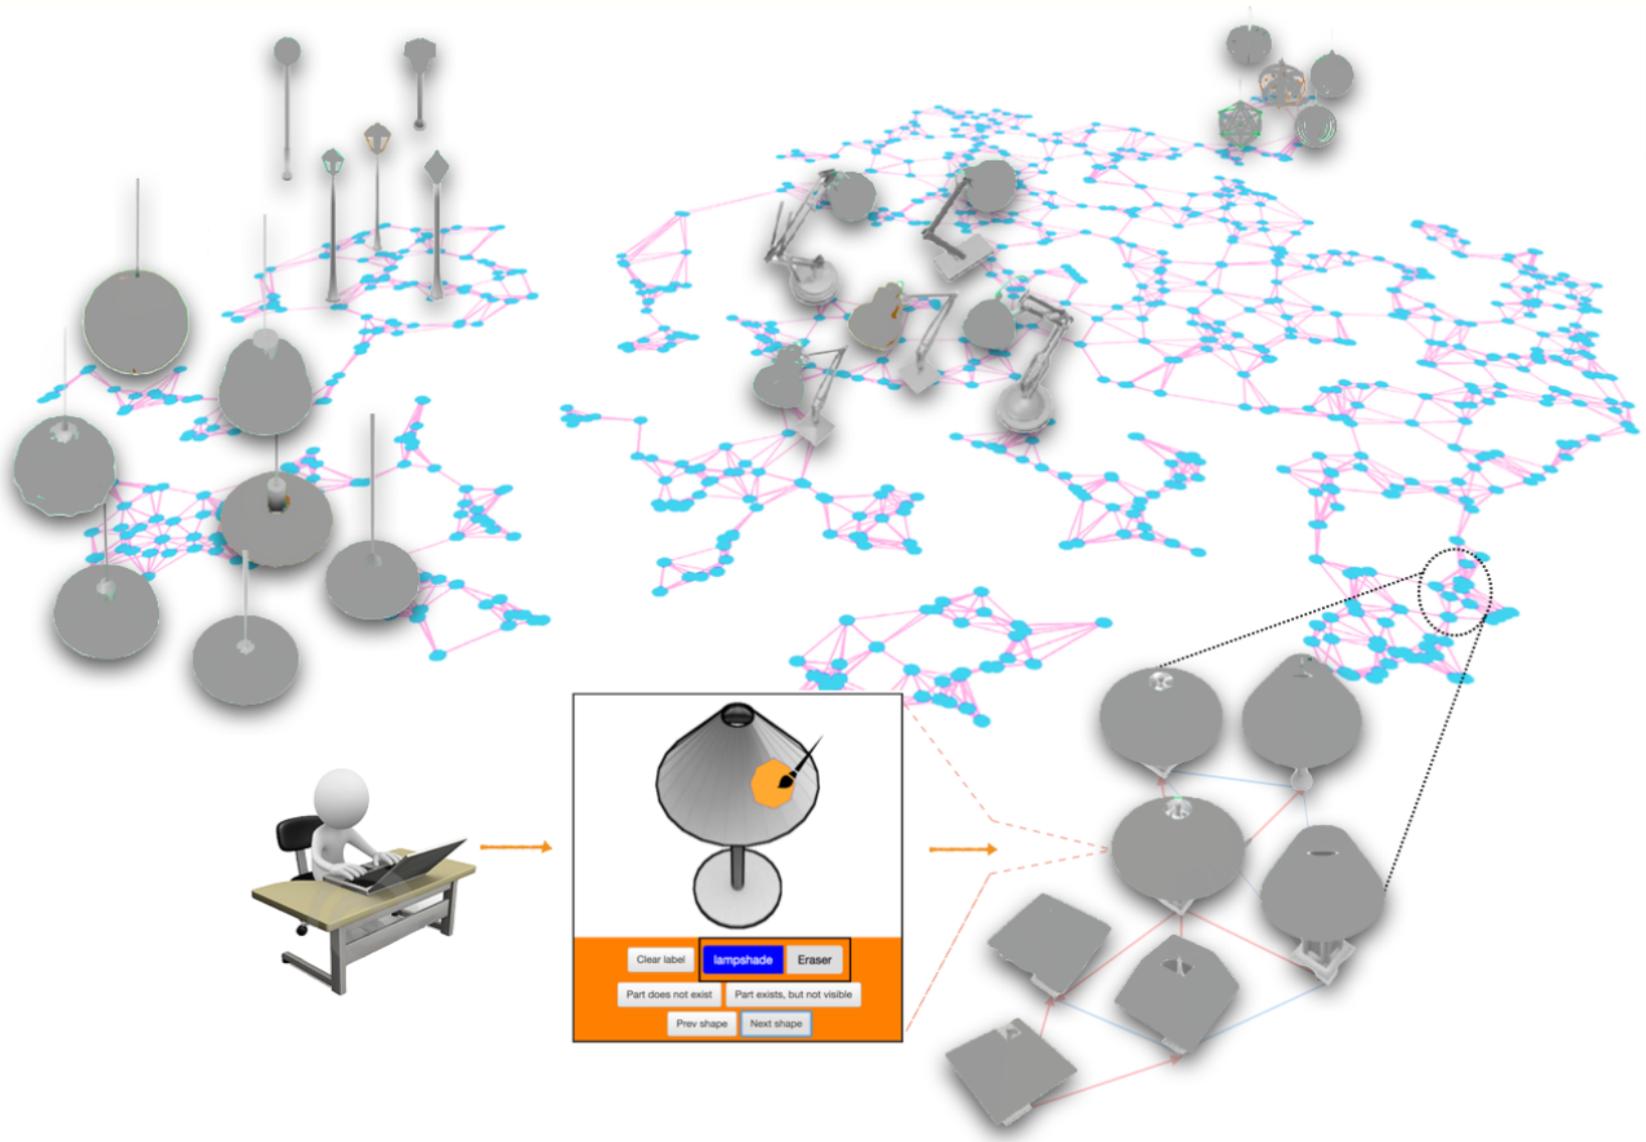
$$\operatorname{argmin}_{\mathcal{A}, \mathcal{V}} \frac{\mathbb{E}[N_{\text{good}}]}{\mathbb{E}[T]}$$

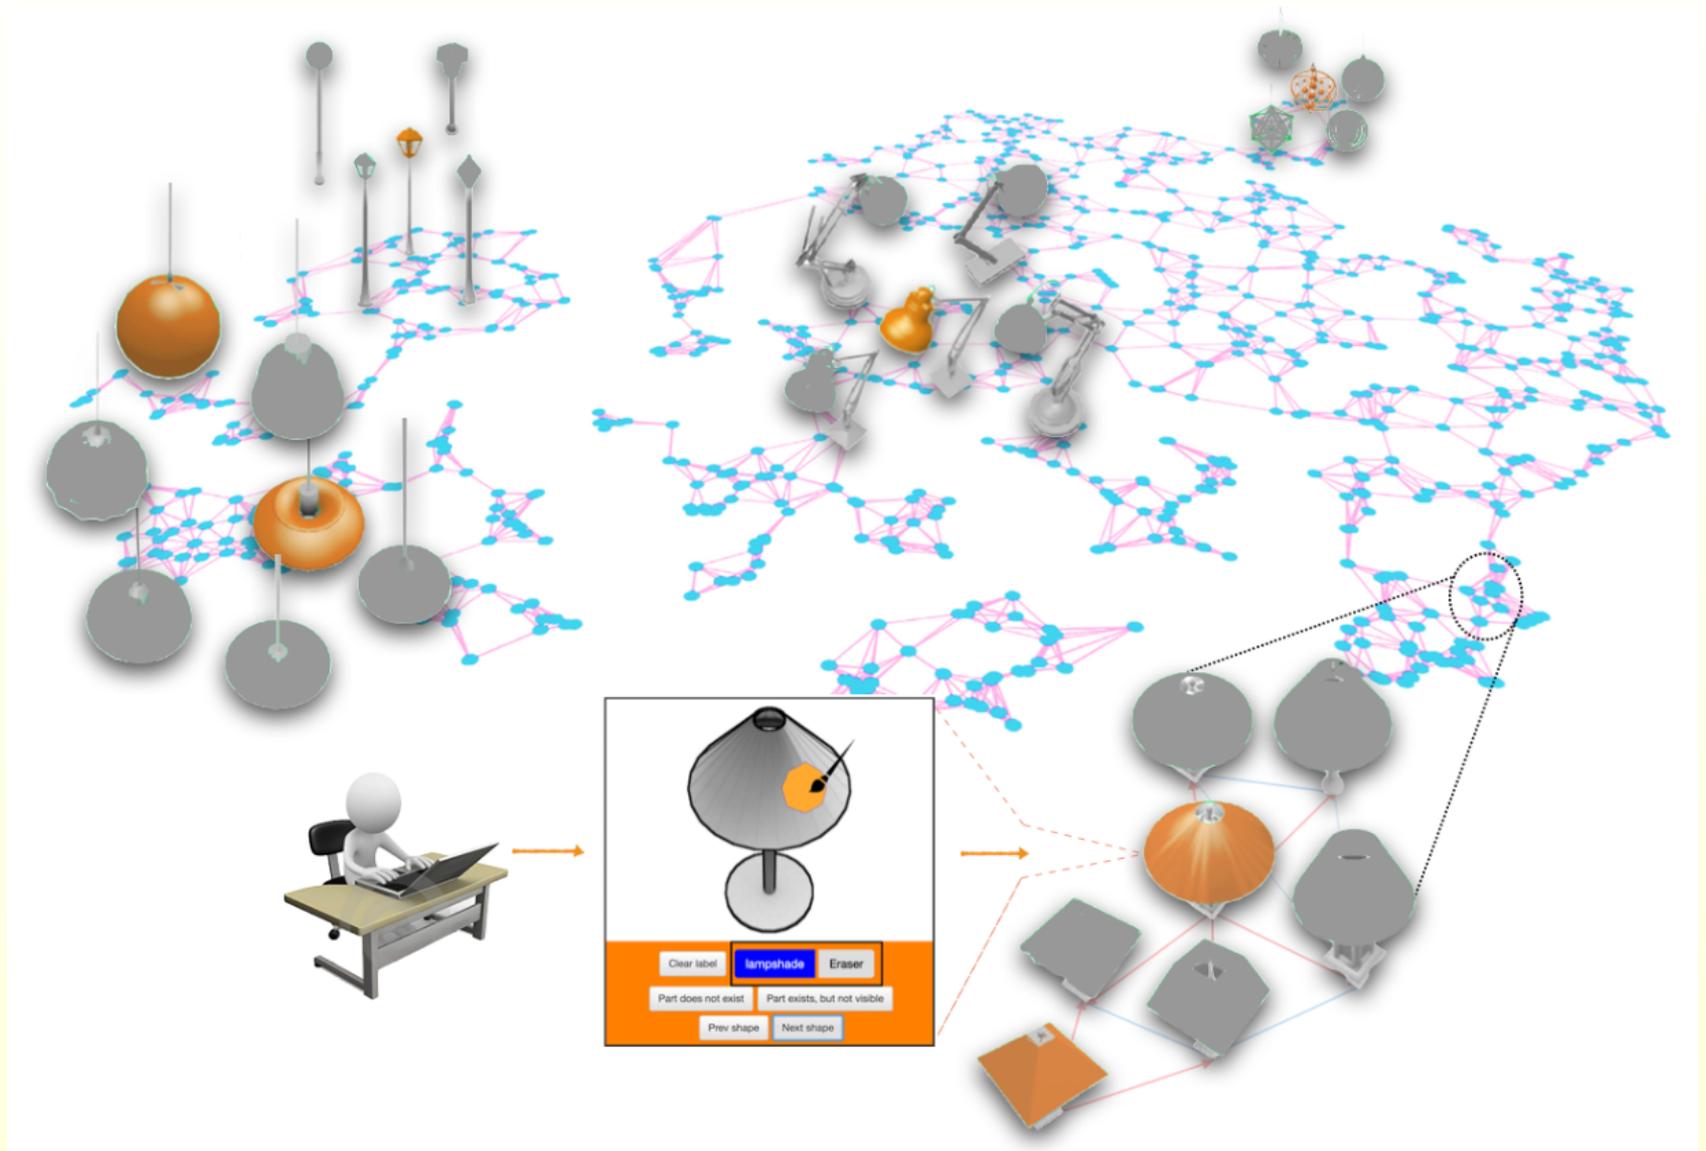
accurate
labels
expectation

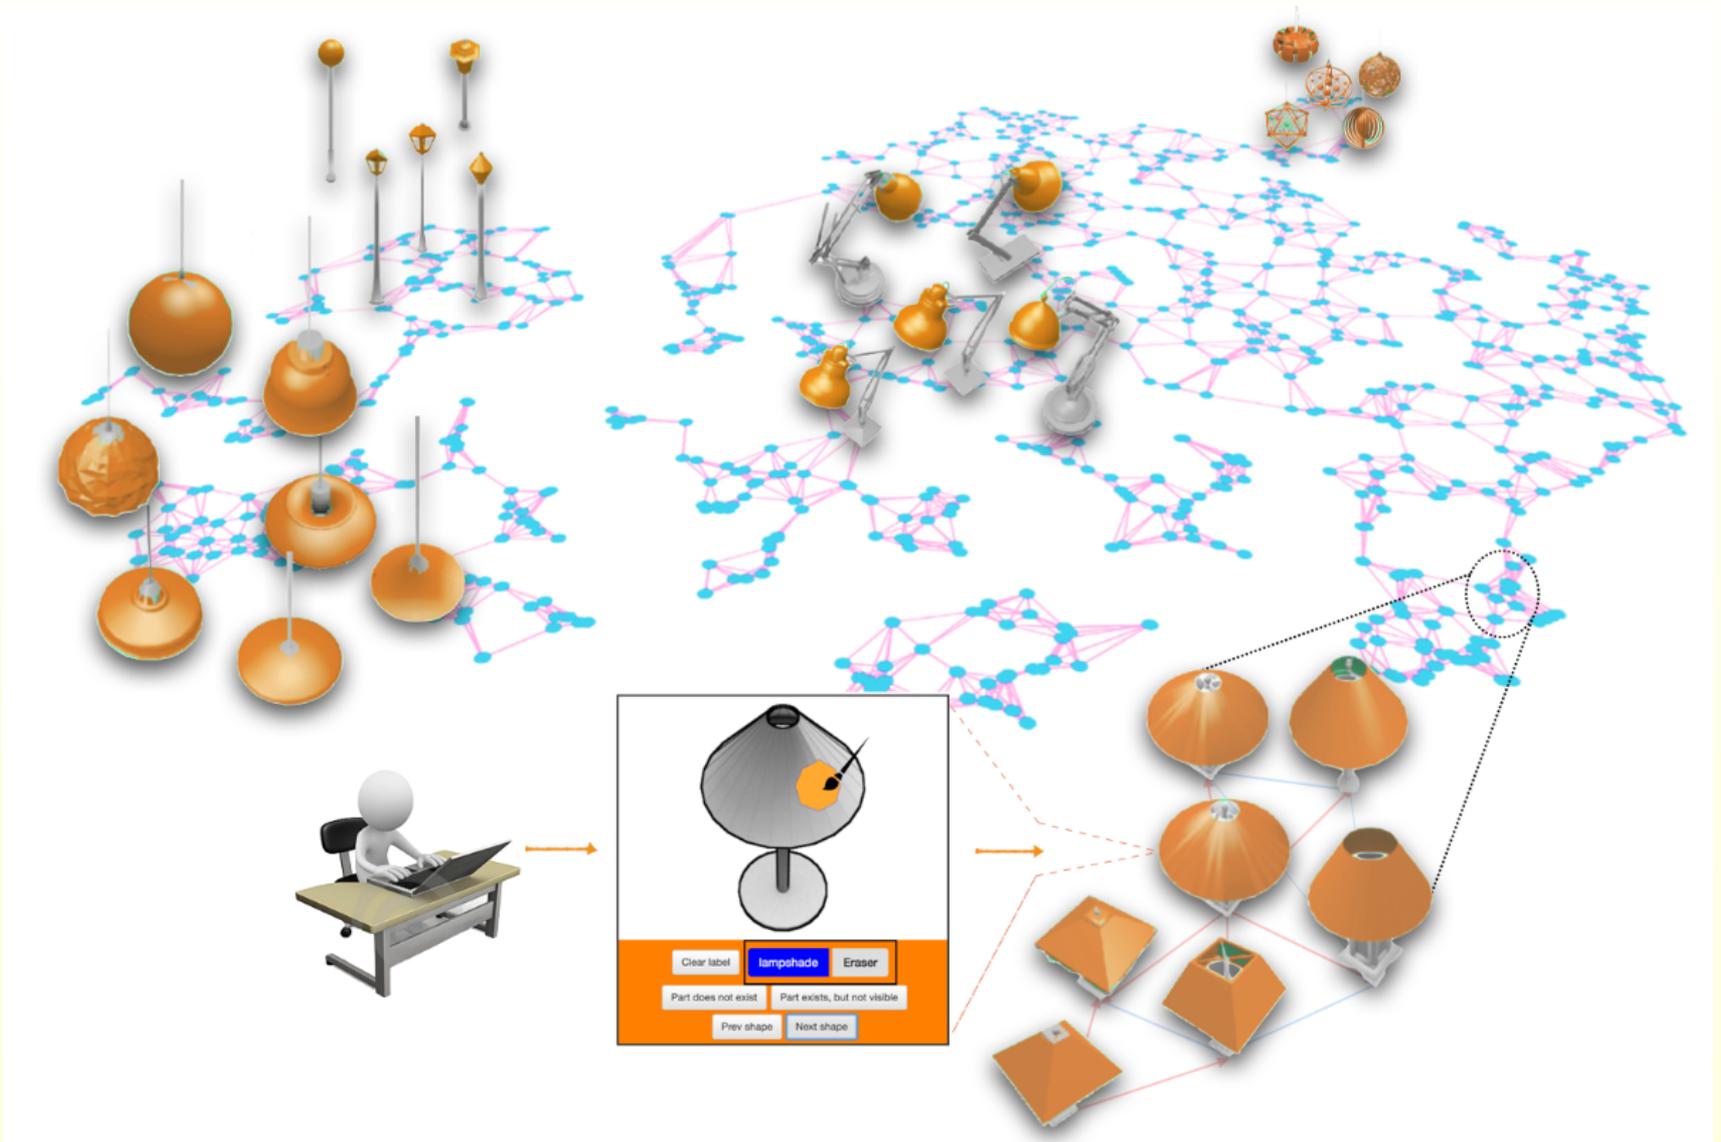
Worker time
expectation

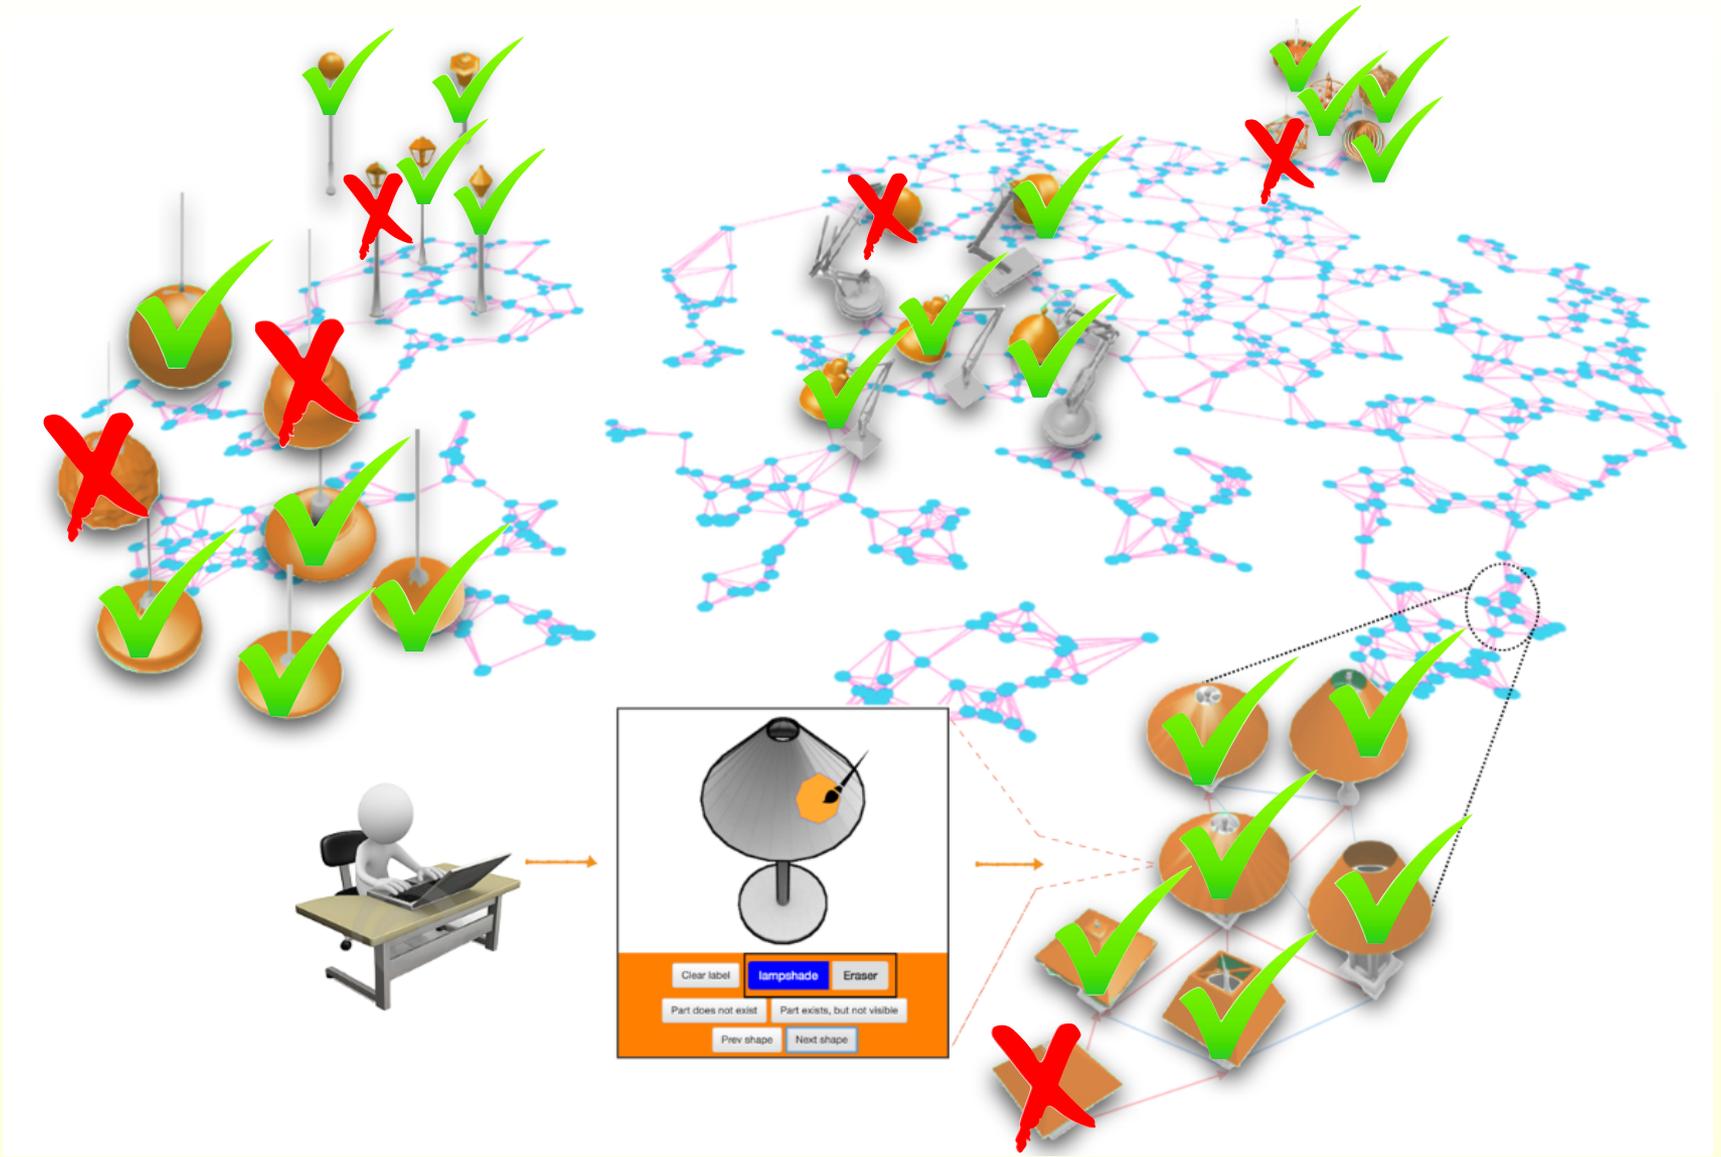
Select Annotation Set: *Beam Search*

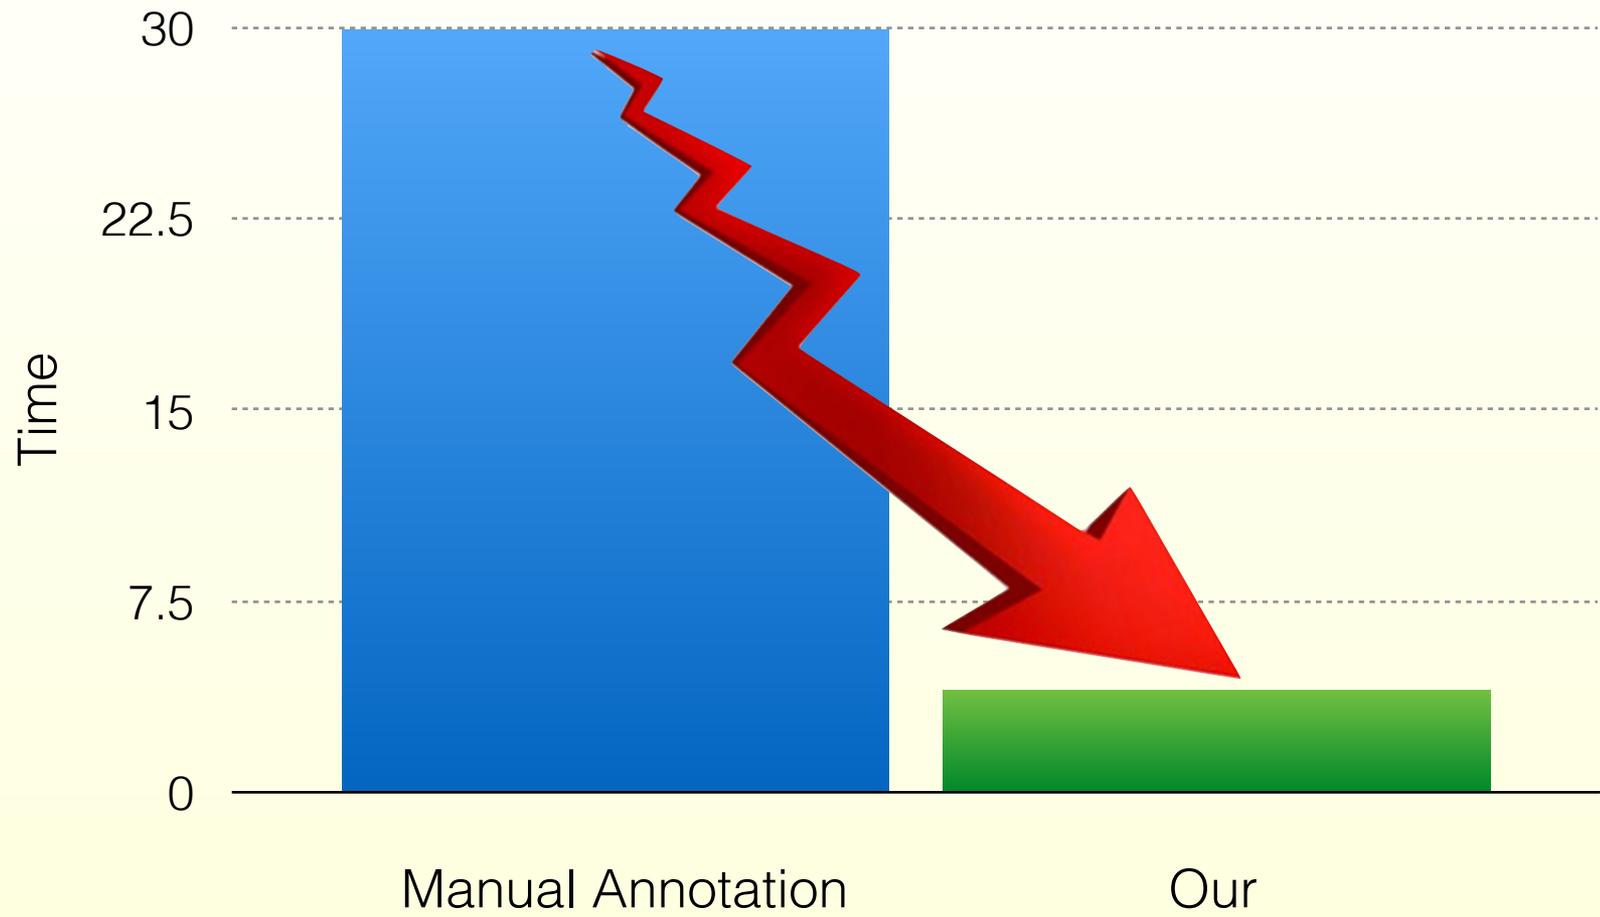
Select Verification Set: *Greedy Search*





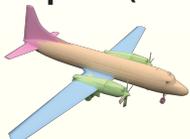






Results

airplane (4027)



- wings
- body
- tail
- engine

bag (83)



- handle
- body

guitar (793)



- body
- head
- neck

chair (6742)



- seat
- back
- arm
- leg

earphone (73)



- headb
- earph

mug (2)



- hand

knife



- ha
- bl

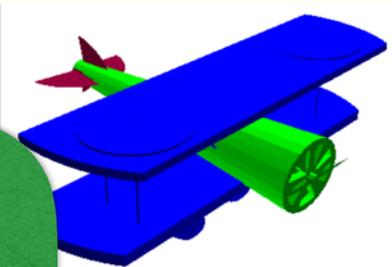
cap (56)



motorbike (336)



~30,000 shapes
~90,000 parts



car (7496)



- roof
- wheels
- hood

skateboard (152)



- deck
- wheel

Evaluation

Part labeling on COSEG dataset [Wang et al. 12]

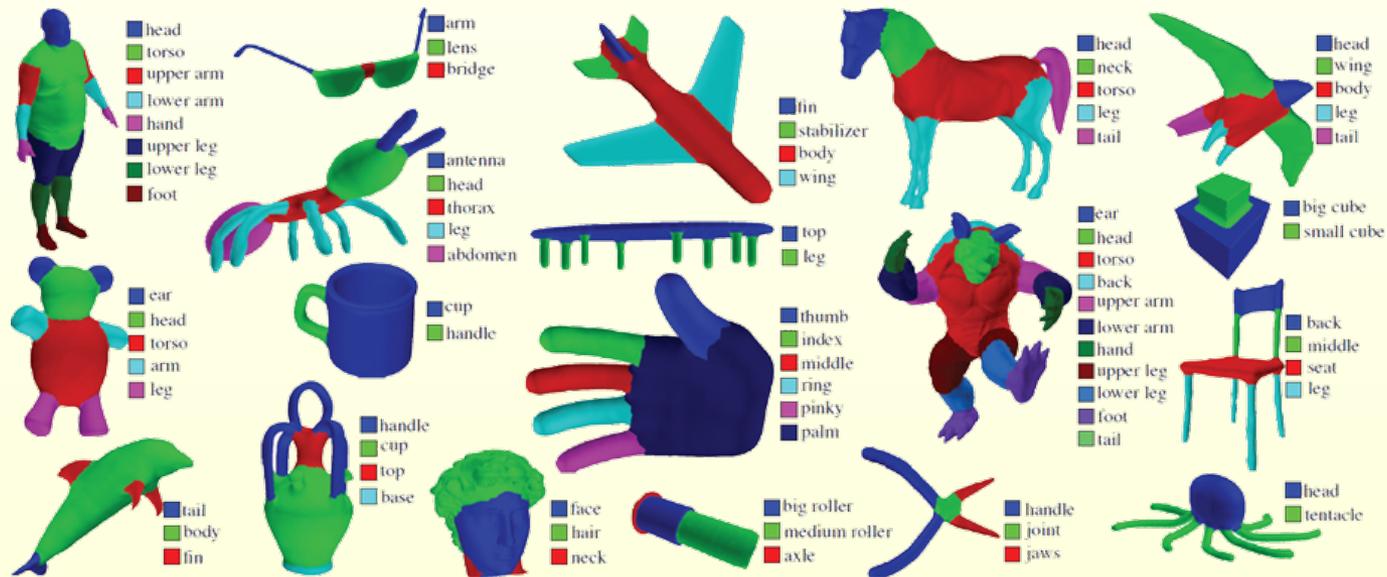
#Mesh	Time ACA (min)	Time Ours (min)
882	152.3	73.6

For your reference

SEGMENTATION/LABELING DATASETS

Segmentation/labeling datasets

- Segmentation/labeling datasets links:
 - Labeled PSB dataset (L-PSB) [Kalogerakis et al. 2010]
<http://people.cs.umass.edu/%7Ekalo/papers/LabelMeshes/>



Segmentation/labeling datasets

- Segmentation/labeling datasets links:
 - Shape COSEG dataset [Wang et al. 2012]
http://irc.cs.sdu.edu.cn/~yunhai/public_html/ssl/ssl.htm



- 1090 watertight meshes, obtained using active segmentation
- 2-5 parts / mesh

Segmentation/labeling datasets

- Segmentation/labeling datasets links:
 - Labeled ShapeNet subset [Yi et al. 2016]
http://web.stanford.edu/~eric/yi/project_page/part_annotation/index.html



~30,000 shapes, ~90,000 parts

Bibliography

- Oliver van Kaick, Andrea Tagliasacchi, Oana Sidi, Hao Zhang, Daniel Cohen-Or, Lior Wolf, Ghassan Hamarneh, "Prior Knowledge for Part Correspondence", Computer Graphics Forum (Proc. Eurographics), vol. 30, n. 2, pp. 553-562, 2011.
- Funkhouser, Thomas, et al. "Modeling by example." ACM Transactions on Graphics (TOG). Vol. 23. No. 3. ACM, 2004.
- Kalogerakis, Evangelos, Aaron Hertzmann, and Karan Singh. "Learning 3D mesh segmentation and labeling." ACM Transactions on Graphics (TOG) 29.4 (2010): 102.
- Sidi, Oana, et al. Unsupervised co-segmentation of a set of shapes via descriptor-space spectral clustering. Vol. 30. No. 6. ACM, 2011.
- Hu, Ruizhen, Lubin Fan, and Ligang Liu. "Co-Segmentation of 3D Shapes via Subspace Clustering." Computer Graphics Forum. Vol. 31. No. 5. Blackwell Publishing Ltd, 2012.
- Wang, Yunhai, et al. "Active co-analysis of a set of shapes." ACM Transactions on Graphics (TOG) 31.6 (2012): 165.
- Kraevoy V., Julius D., Sheffer A.: Model composition from interchangeable components. In Proc. Pacific Graphics (2007), pp. 129–138. 3

Bibliography

- Golovinskiy, Aleksey, and Thomas Funkhouser. "Consistent segmentation of 3D models." *Computers & Graphics* 33.3 (2009): 262-269.
- Shapira, Lior, Ariel Shamir, and Daniel Cohen-Or. "Consistent mesh partitioning and skeletonisation using the shape diameter function." *The Visual Computer* 24.4 (2008): 249.
- COMANICIU, D., AND MEER, P. 2002. Mean shift: a robust approach towards feature space analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* 24, 5, 603–619.
- Qixing Huang, Vladlen Koltun, and Leonidas Guibas. Joint-Shape Segmentation with Linear Programming. *ACM Transactions on Graphics* 30(5) (Proc. SIGGRAPH Asia 2011)
- Li Yi, Vladimir G. Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas. A Scalable Active Framework for Region Annotation in 3D Shape Collections . *ACM Transaction on Graphics* 35(6) (Proc. Siggraph Asia 2016).