Learning 3D mesh segmentation and labeling

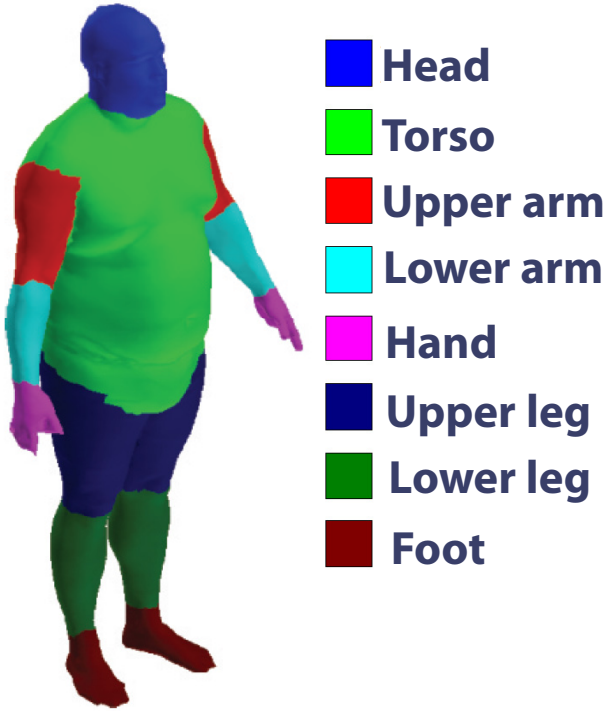

Evangelos Kalogerakis, Aaron Hertzmann, Karan Singh

University of Toronto

Goal: mesh segmentation and labeling

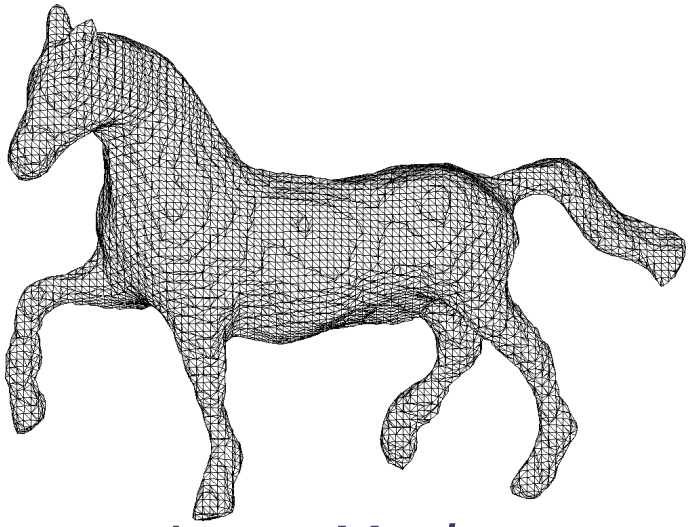

Input Mesh

Goal: mesh segmentation and labeling

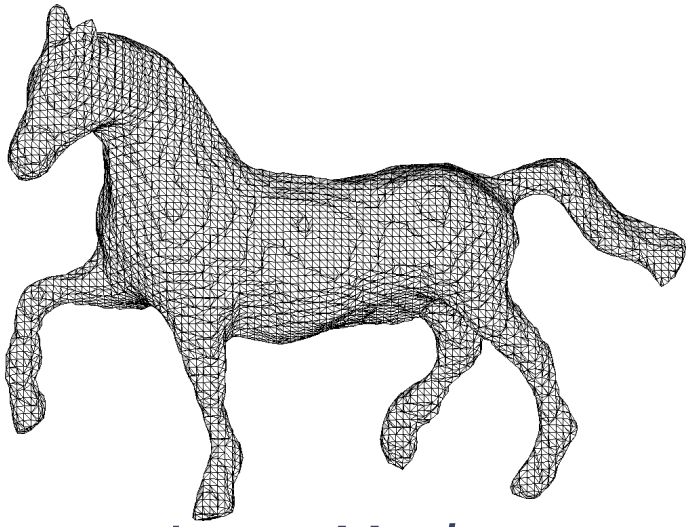

Input Mesh

Labeled Mesh

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

Goal: mesh segmentation and labeling

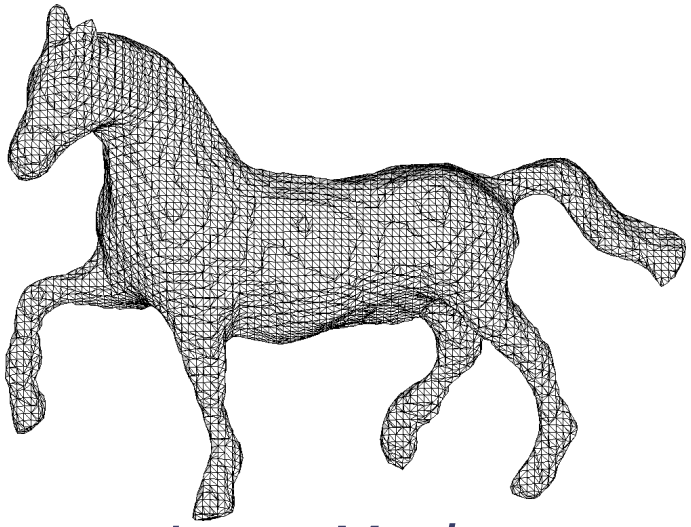

Input Mesh

Labeled Mesh

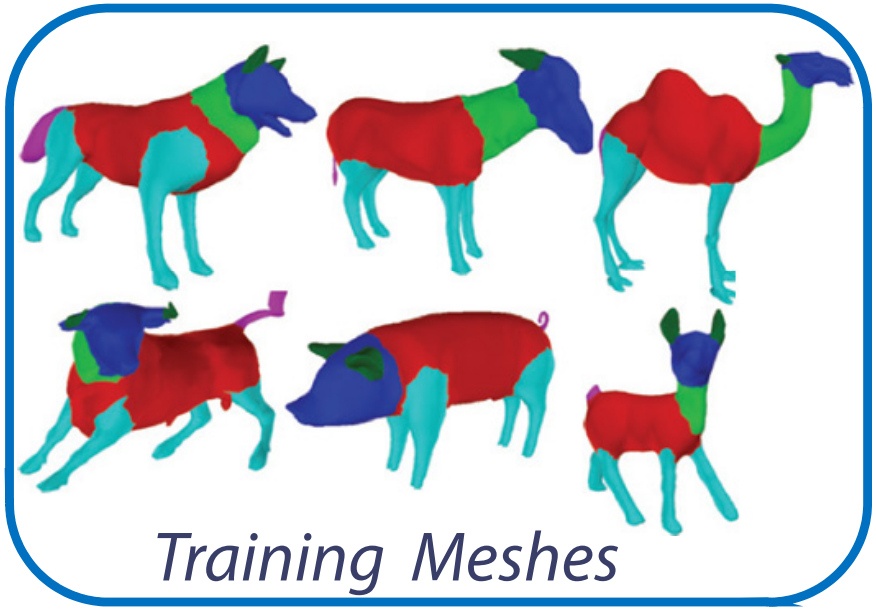

Training Meshes

Related work: mesh segmentation

[Mangan and Whitaker 1999, Shlafman *et al.* 2002, Katz and Tal 2003, Liu and Zhang 2004, Katz *et al.* 2005, Simari *et al.* 2006, Attene *et al.* 2006, Lin *et al.* 2007, Kraevoy *et al.* 2007, Pekelny and Gotsman 2008, Golovinskiy and Funkhouser 2008, Li *et al.* 2008, Lai *et al.* 2008, Lavoue and Wolf 2008, Huang *et al.* 2009, Shapira *et al.* 2010]

Surveys:

[Attene et al. 2006, Shamir 2008, Chen et al. 2009]

Related work: mesh segmentation

Shape Diameter
[Shapira et al. 10]

Randomized Cuts
[Golovinskiy and Funkhouser 08]

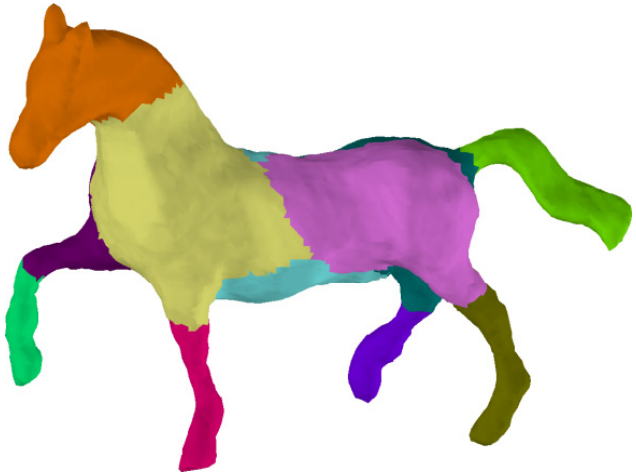

Random Walks
[Lai et al. 08]

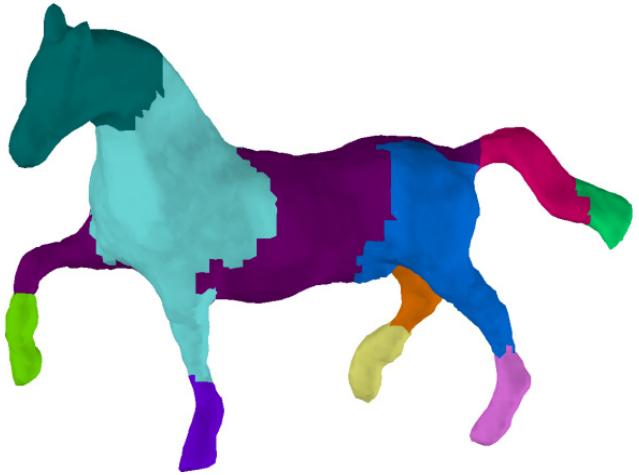

Normalized Cuts
[Golovinskiy and Funkhouser 08]

Is human-level segmentation even possible without higher-level cues?

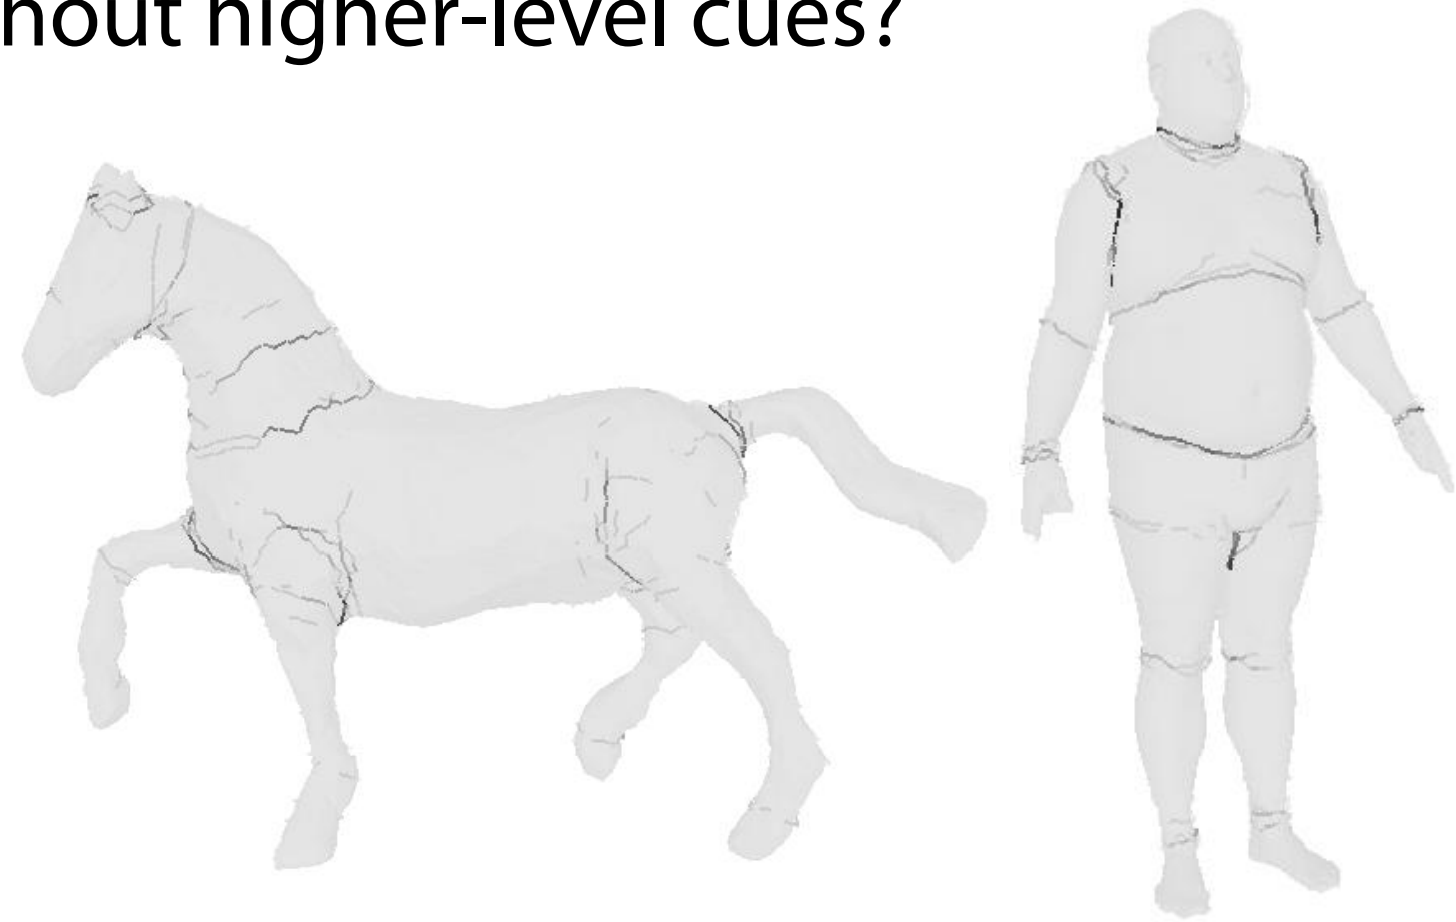

[X. Chen et al. SIGGRAPH 09]

Is human-level segmentation even possible without higher-level cues?

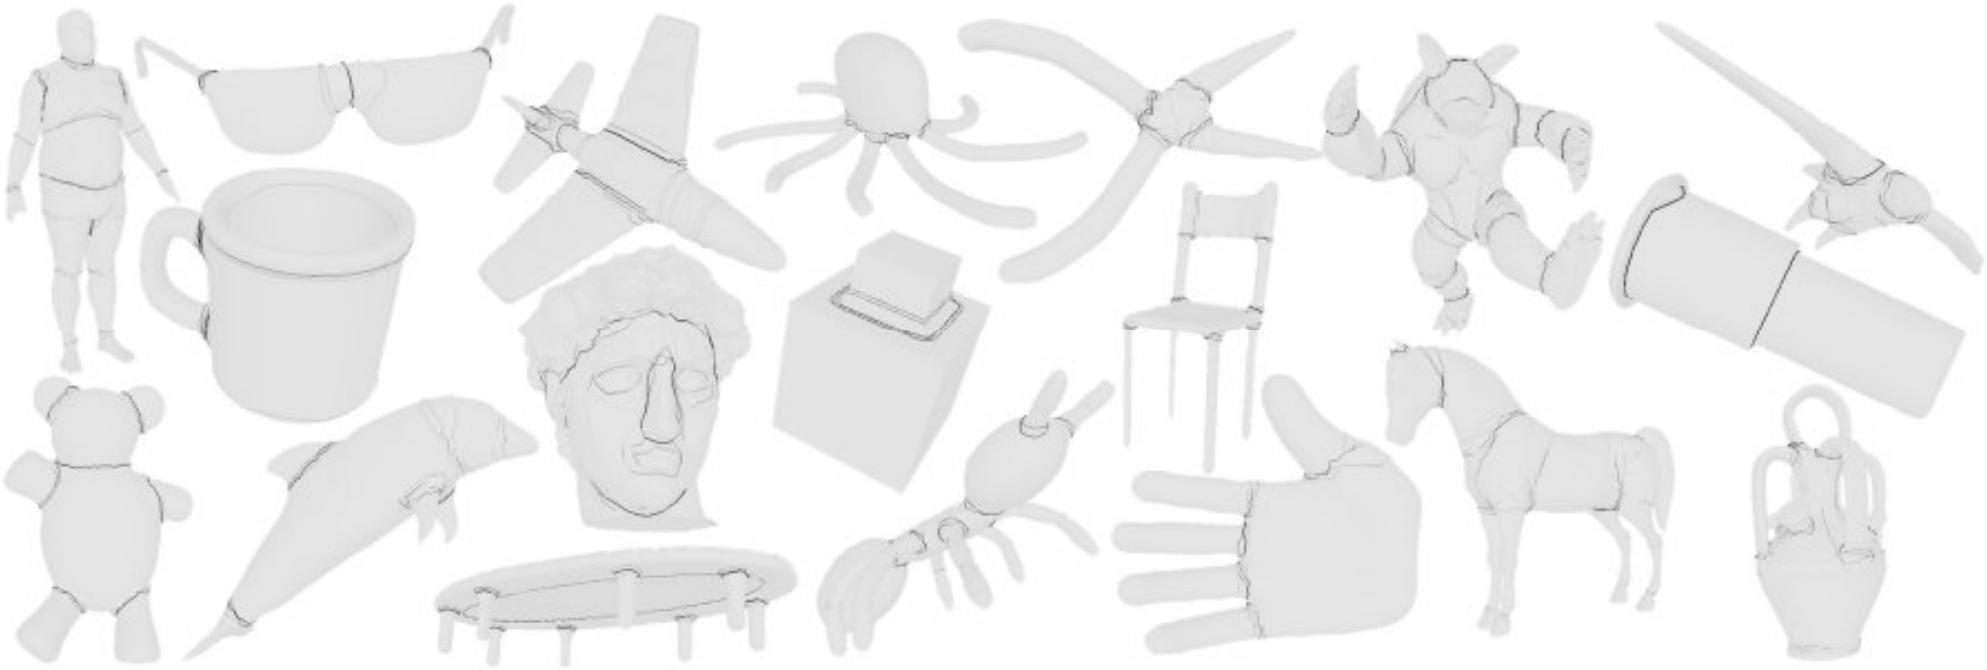

[X. Chen et al. SIGGRAPH 09]

Image segmentation and labeling

[Konishi and Yuille 00, Duygulu et al. 02, He et al. 04, Kumar and Hebert 03, Anguelov et al. 05, Tu et al.05, Schnitman et al. 06, Lim and Suter 07, Munoz et al. 08,…]

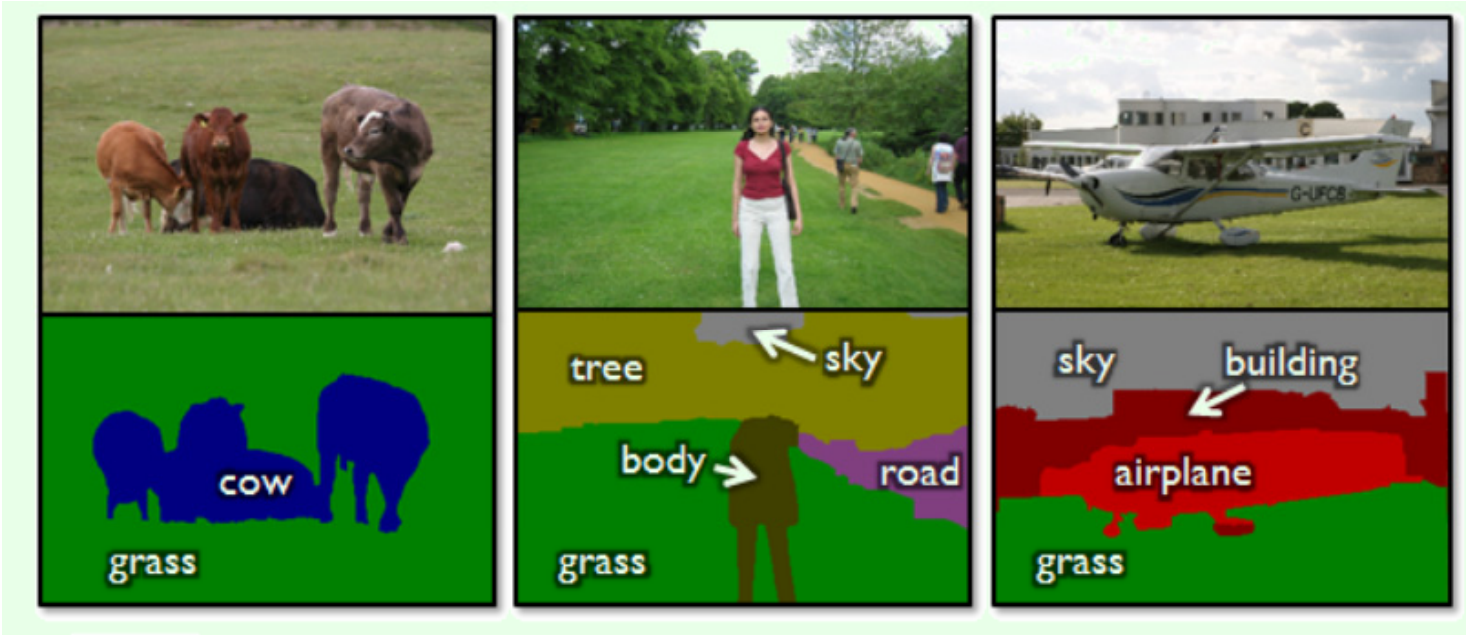

Textonboost
[Shotton et al. ECCV 06]

Related work: mesh segmentation & labeling

Consistent segmentation of 3D meshes
[Golovinskiy and Funkhouser 09]

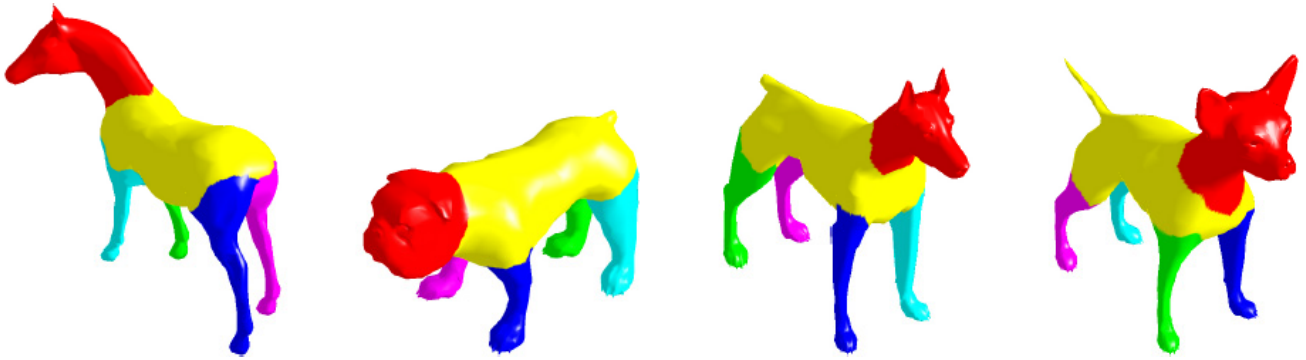

Multi-objective segmentation and labeling
[Simari et al. 09]

Learning mesh segmentation and labeling

Learn from examples

Significantly better results than state-of-the-art

No manual parameter tuning

Can learn different styles of segmentation

Several applications of part labeling

Labeling problem statement

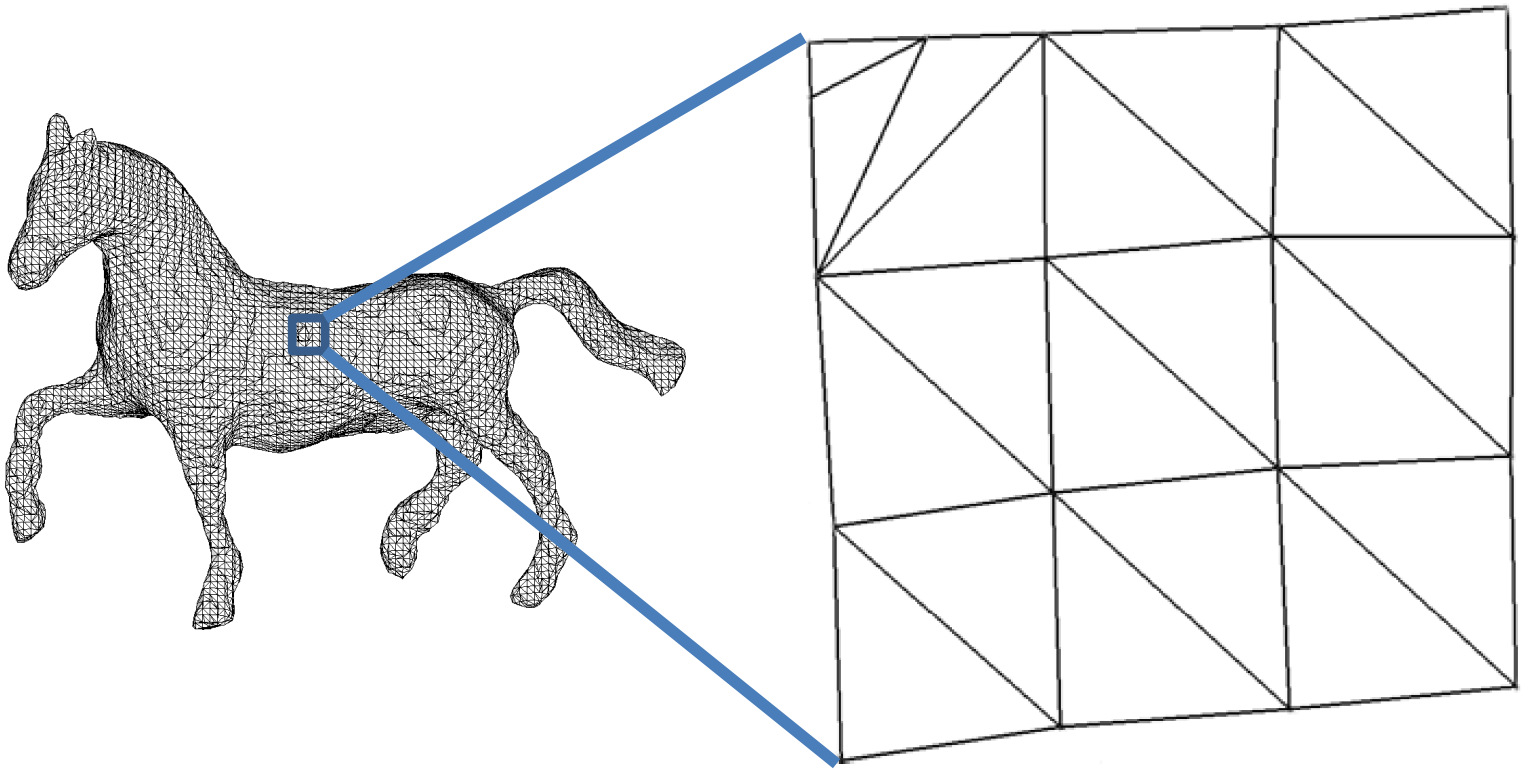

Labeling problem statement

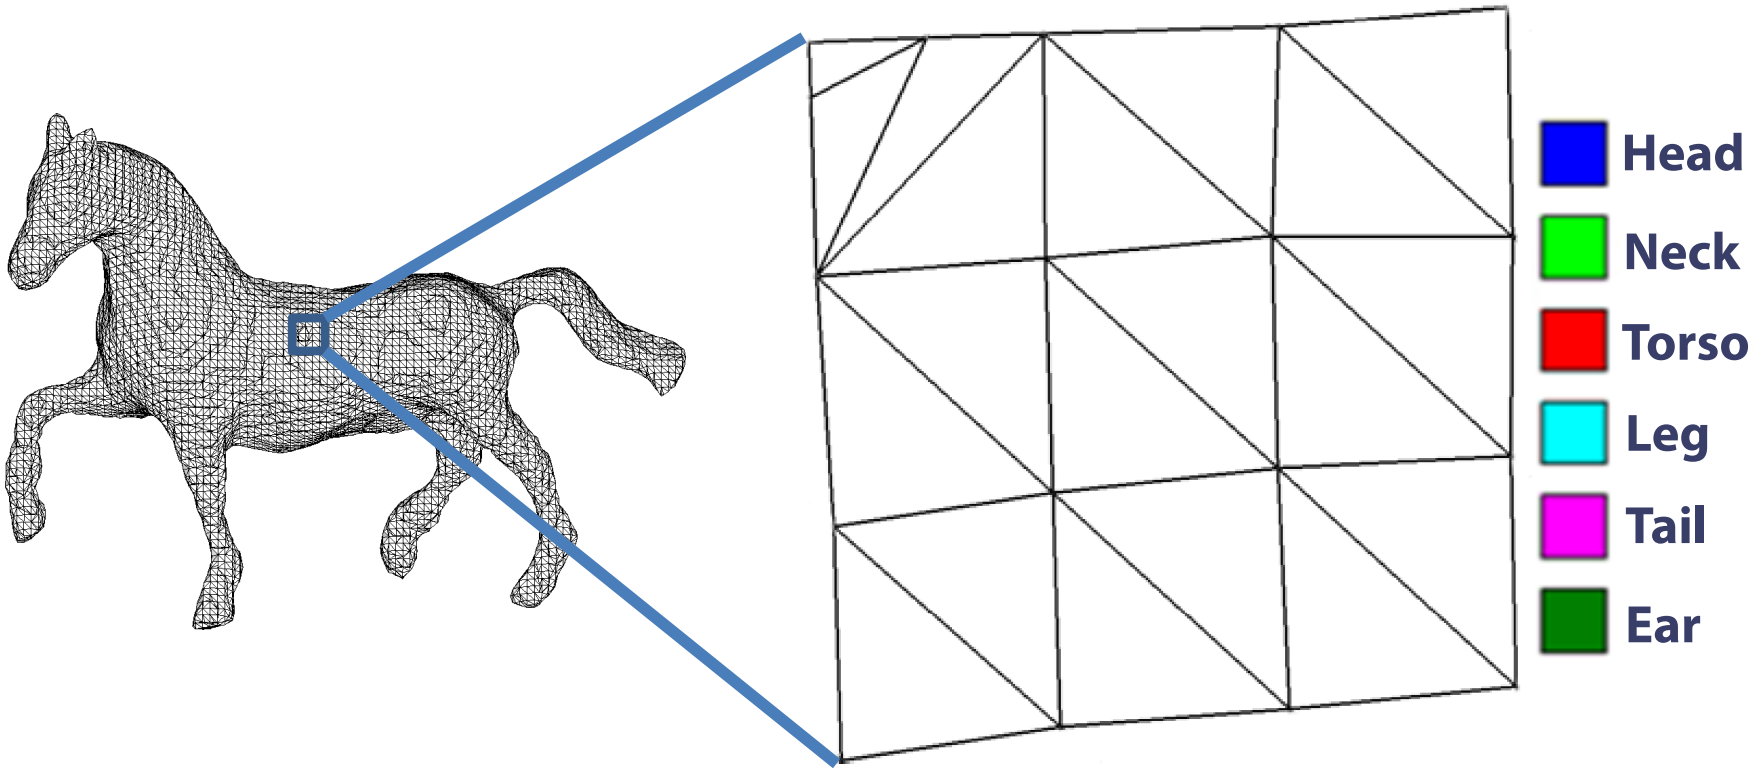

Labeling problem statement

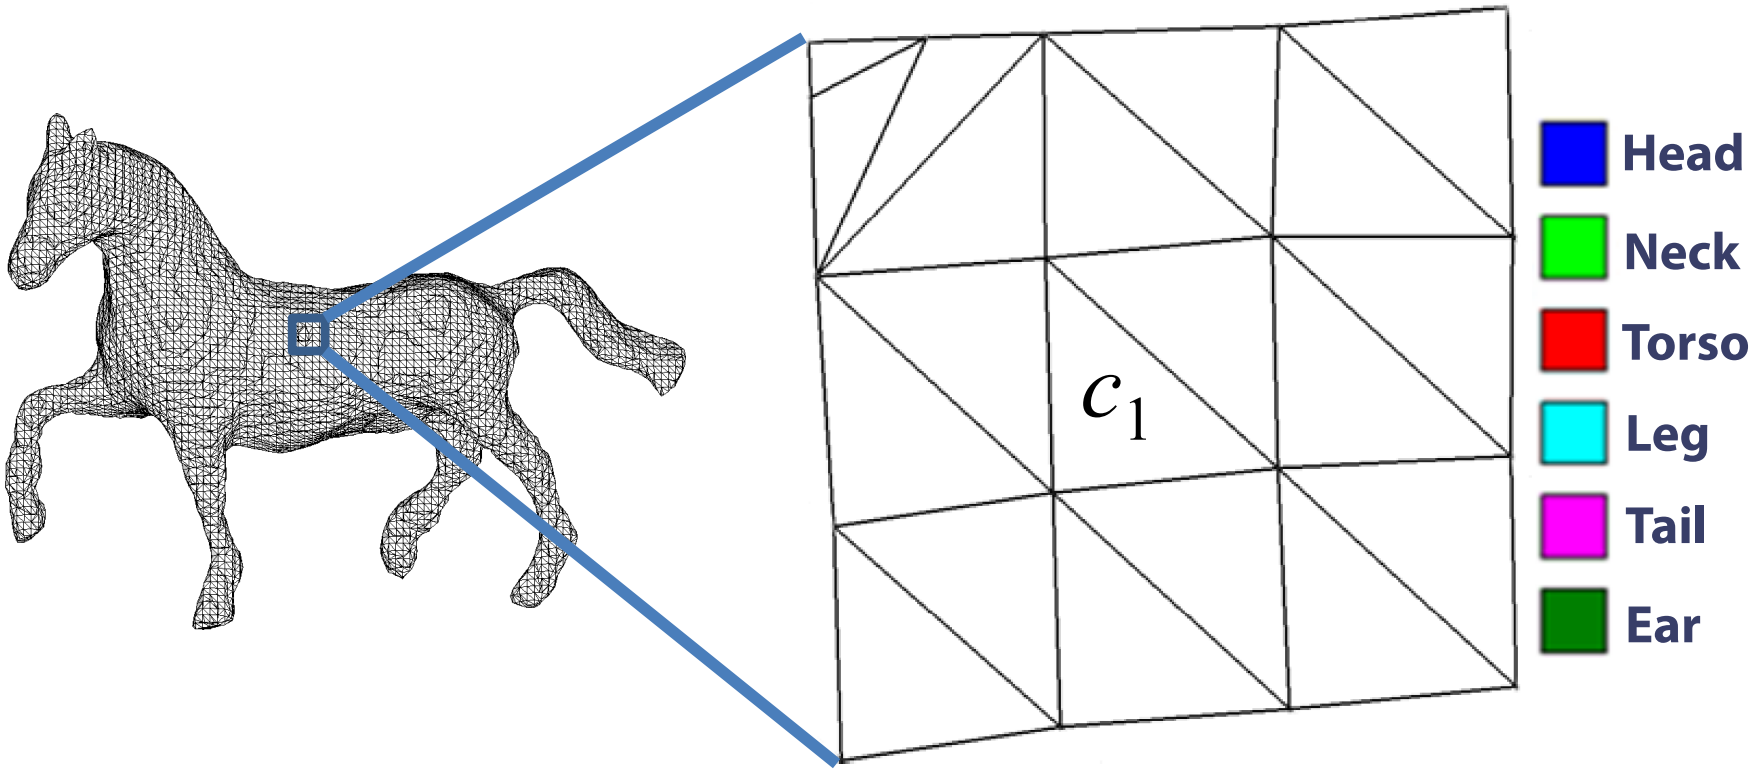

Labeling problem statement

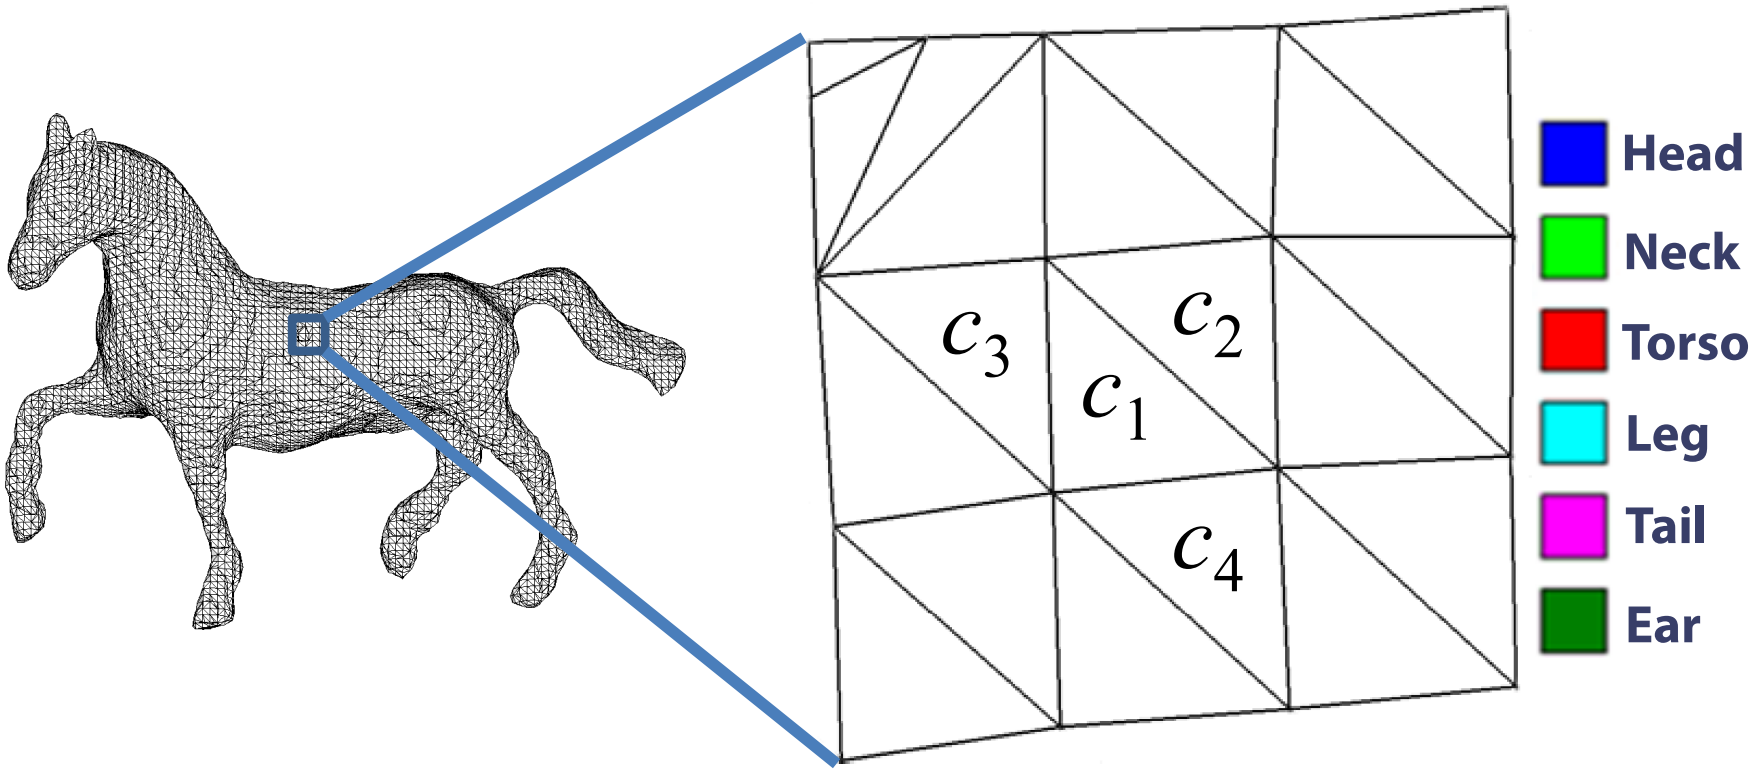

$$c_1, c_2, c_3 \in C$$

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

Conditional Random Field for Labeling

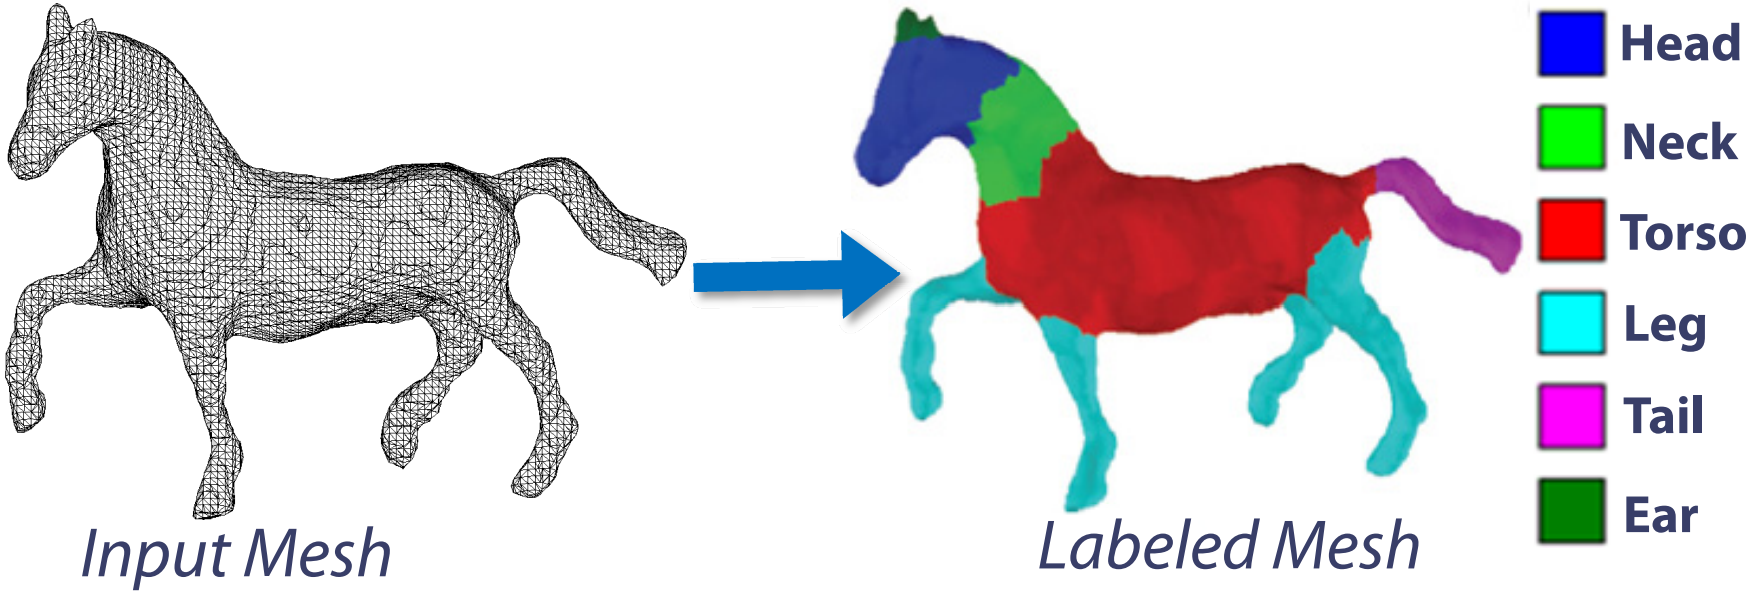

Conditional Random Field for Labeling

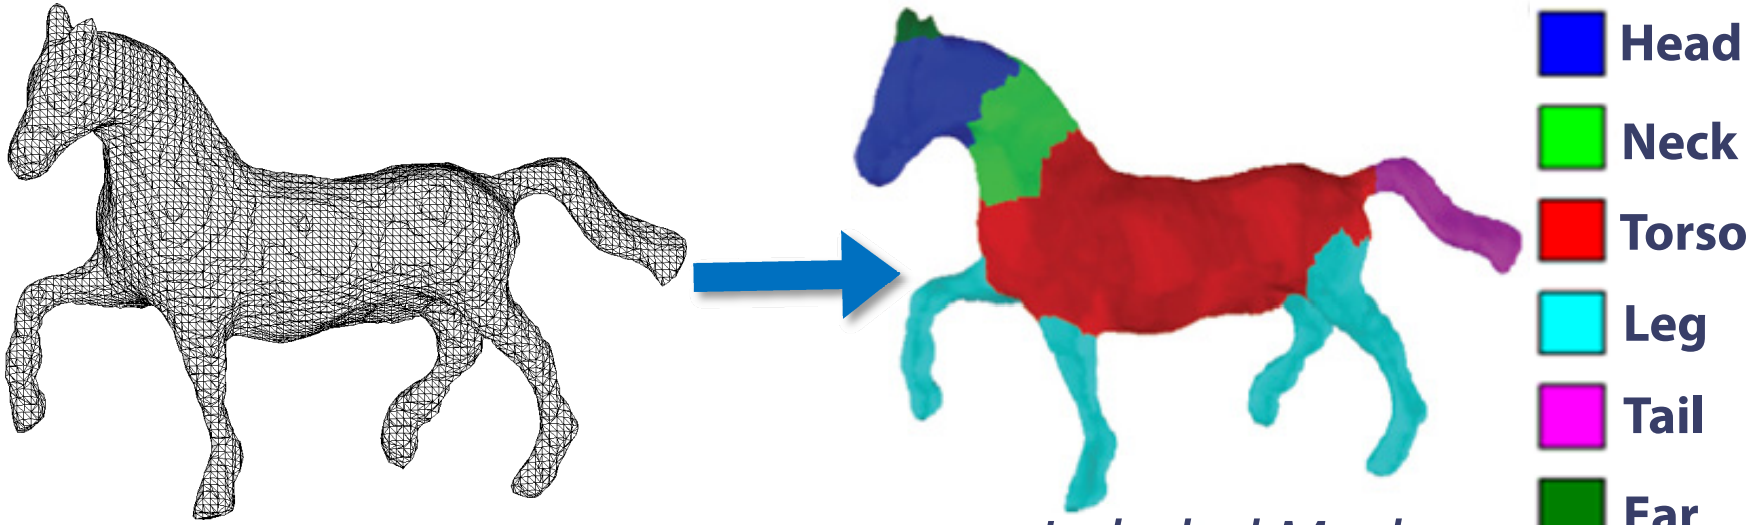

$$c^* = \arg\min_{\mathbf{c}} \left\{ \sum_{i} \alpha_i \mathrm{E}_1(c_i; \mathbf{x_i}) + \sum_{i,j} l_{ij} \mathrm{E}_2(c_i, c_j; \mathbf{y_{ij}}) \right\}$$

Conditional Random Field for Labeling

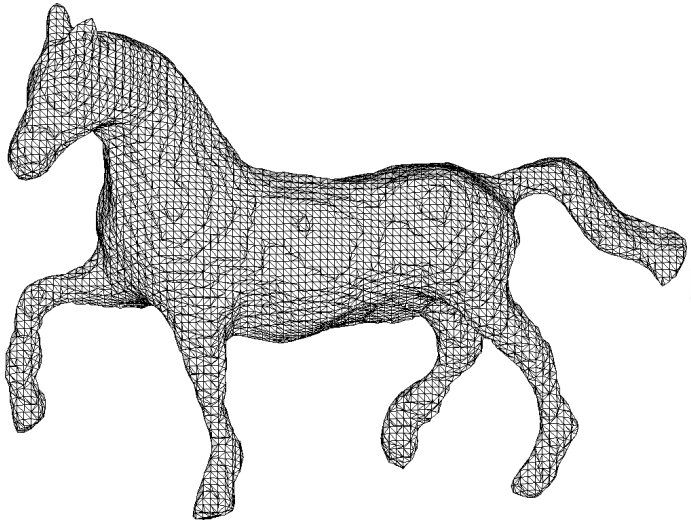

Input Mesh

Labeled Mesh

$$c^* = \arg\min_{\mathbf{c}} \left\{ \sum_i \alpha_i \mathrm{E}_1(c_i; \mathbf{x_i}) + \sum_{i,j} l_{ij} \mathrm{E}_2(c_i, c_j; \mathbf{y_{ij}}) \right\}$$

Unary term

Conditional Random Field for Labeling

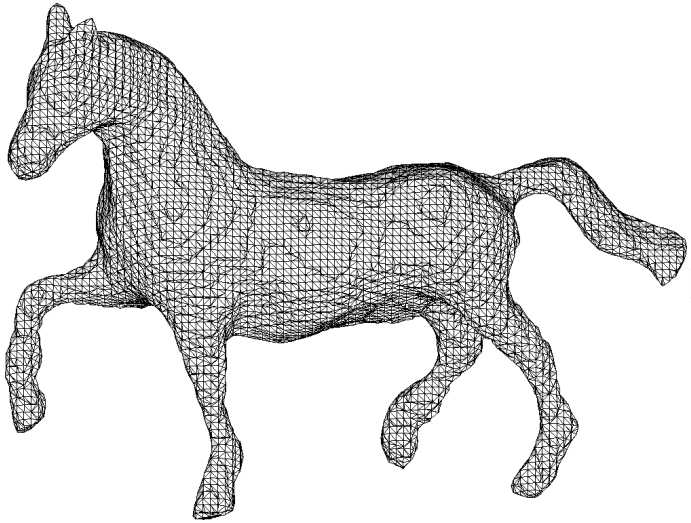

Input Mesh

Labeled Mesh

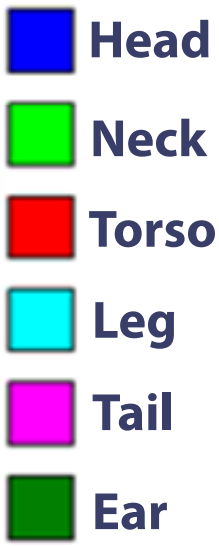

$$c^* = \arg\min_{\mathbf{c}} \left\{ \sum_{i} \alpha_i \mathrm{E}_1(c_i; \mathbf{x_i}) + \sum_{i,j} l_{ij} \mathrm{E}_2(c_i, c_j; \mathbf{y_{ij}}) \right\}$$

Face features

Conditional Random Field for Labeling

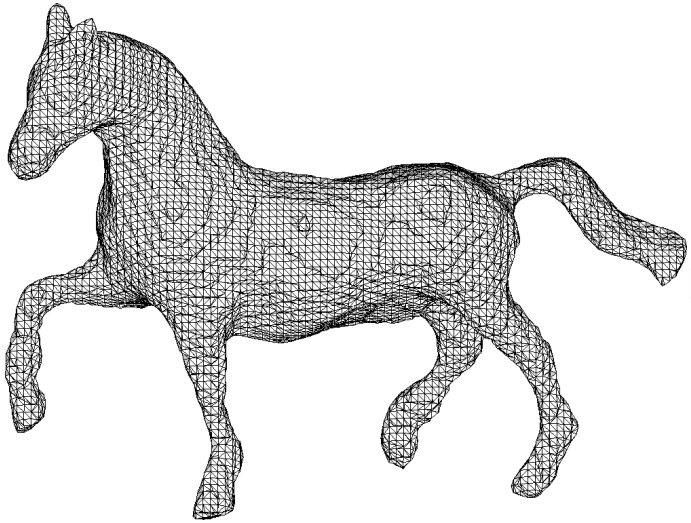

Input Mesh

Labeled Mesh

$$c^* = \arg\min_{\mathbf{c}} \left\{ \sum_{i} \alpha_i \mathrm{E}_1(c_i; \mathbf{x_i}) + \sum_{i,j} l_{ij} \mathrm{E}_2(c_i, c_j; \mathbf{y_{ij}}) \right\}$$

Face Area

Conditional Random Field for Labeling

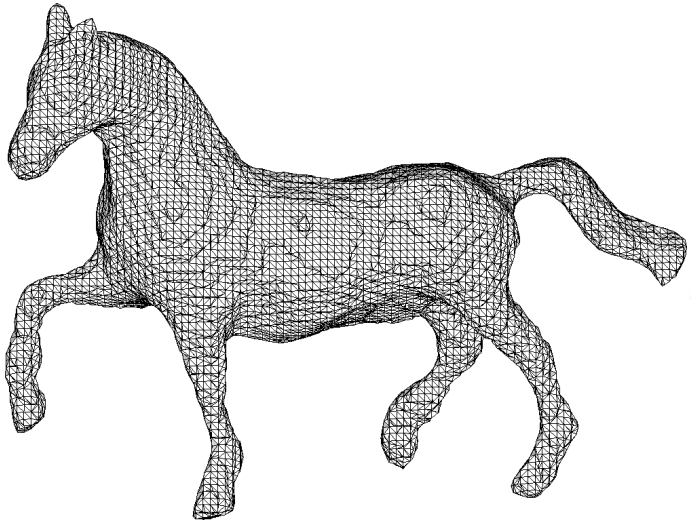

Input Mesh

Labeled Mesh

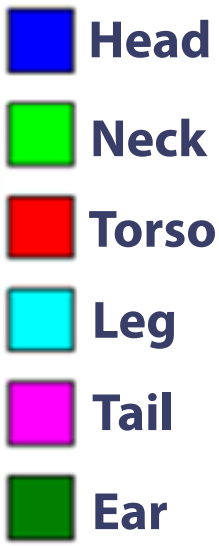

$$c^* = \arg\min_{\mathbf{c}} \left\{ \sum_{i} \alpha_i \mathrm{E}_1(c_i; \mathbf{x_i}) + \sum_{i,j} l_{ij} \mathrm{E}_2(c_i, c_j; \mathbf{y_{ij}}) \right\}$$

Pairwise Term

Conditional Random Field for Labeling

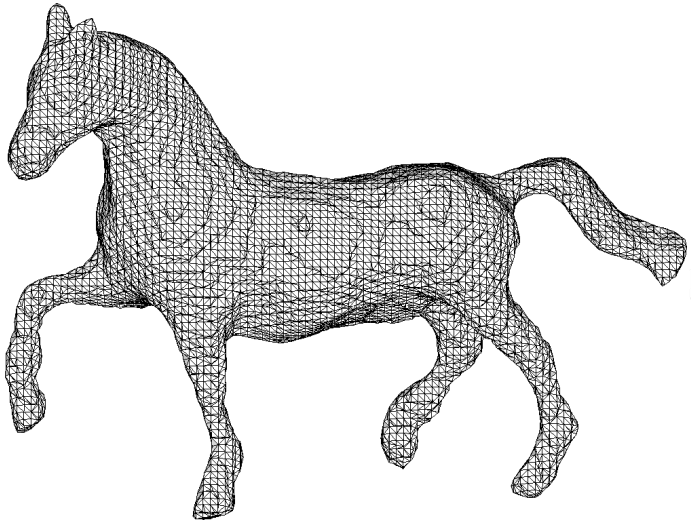

Input Mesh

Labeled Mesh

$$c^* = \arg\min_{\mathbf{c}} \left\{ \sum_{i} \alpha_i \mathrm{E}_1(c_i; \mathbf{x_i}) + \sum_{i,j} l_{ij} \mathrm{E}_2(c_i, c_j; \mathbf{y_{ij}}) \right\}$$

Edge Features

Conditional Random Field for Labeling

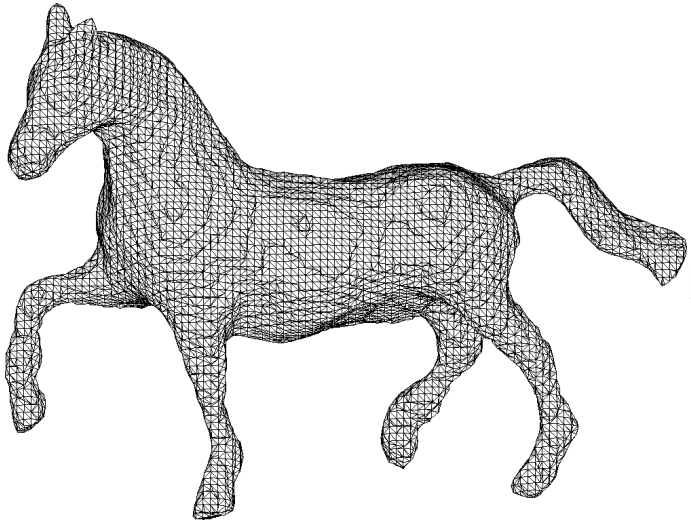

Input Mesh

Labeled Mesh

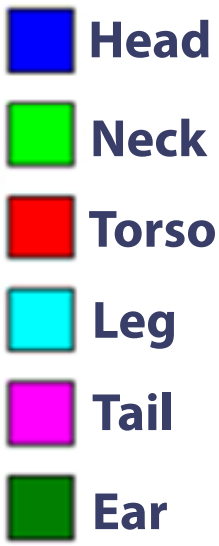

$$c^* = \arg\min_{\mathbf{c}} \left\{ \sum_{i} \alpha_i \mathrm{E}_1(c_i; \mathbf{x_i}) + \sum_{i,j} l_{ij} \mathrm{E}_2(c_i, c_j; \mathbf{y_{ij}}) \right\}$$

Edge Length

Conditional Random Field for Labeling

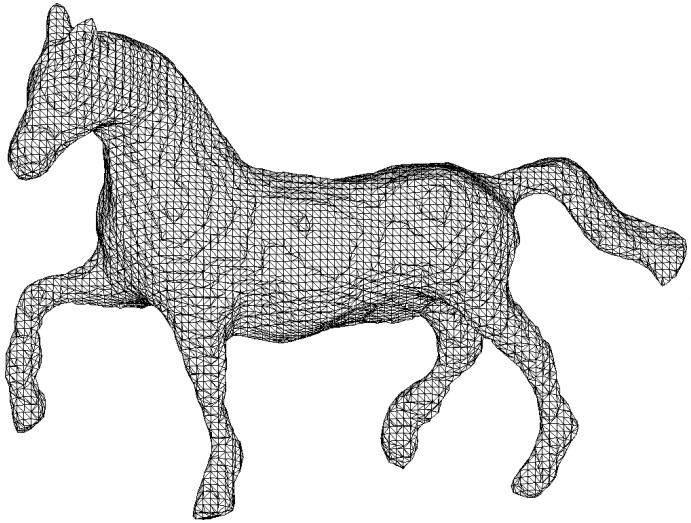

Input Mesh

Labeled Mesh

$$c^* = \arg\min_{\mathbf{c}} \left\{ \sum_i \alpha_i \mathrm{E}_1(c_i; \mathbf{x_i}) + \sum_{i,j} l_{ij} \mathrm{E}_2(c_i, c_j; \mathbf{y_{ij}}) \right\}$$

Unary term

Feature vector

$$\mathbf{x} \in \Re^{375+35|C|} \rightarrow \mathrm{P}(c \mid \mathbf{x})$$

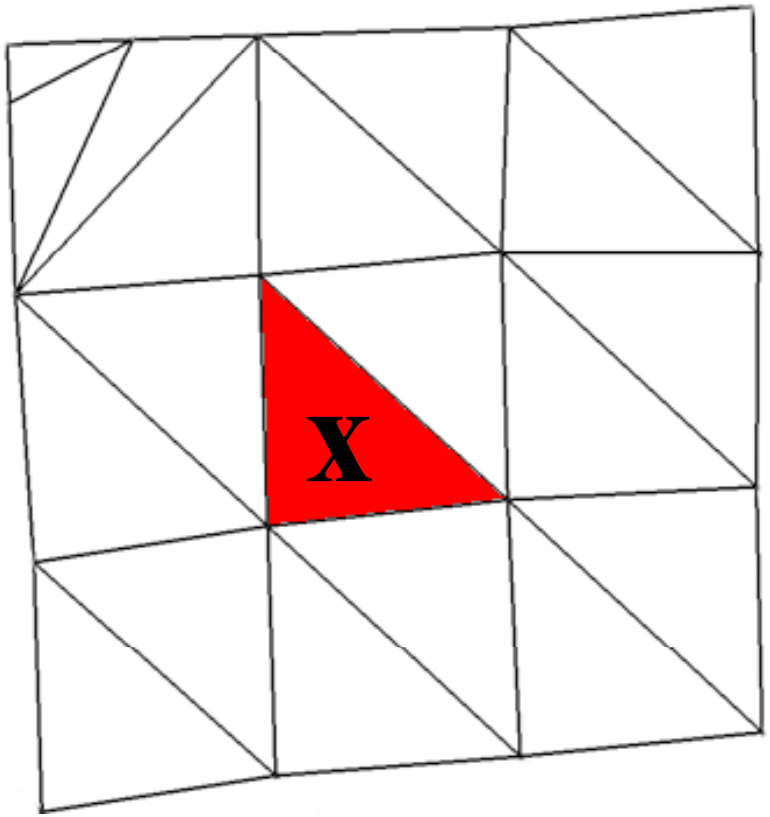

Feature vector

$$\mathbf{x} \in \Re^{375+35|C|} \rightarrow \mathrm{P}(c \mid \mathbf{x})$$

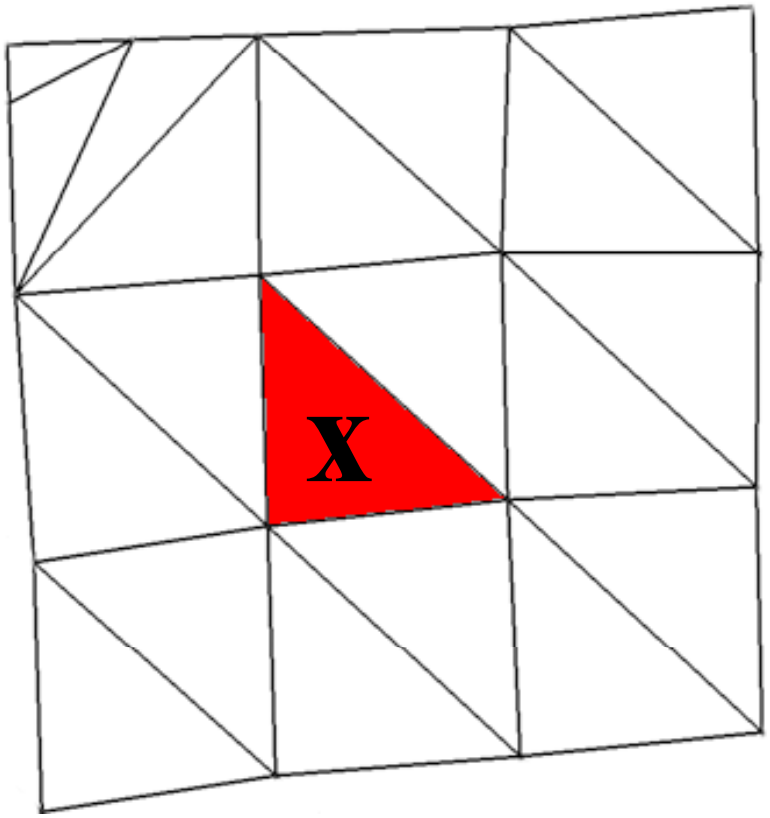

surface curvature

singular values from PCA

shape diameter

distances from medial surface

average geodesic distances

shape contexts

spin images

contextual label features

Learning a classifier

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

Learning a classifier

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

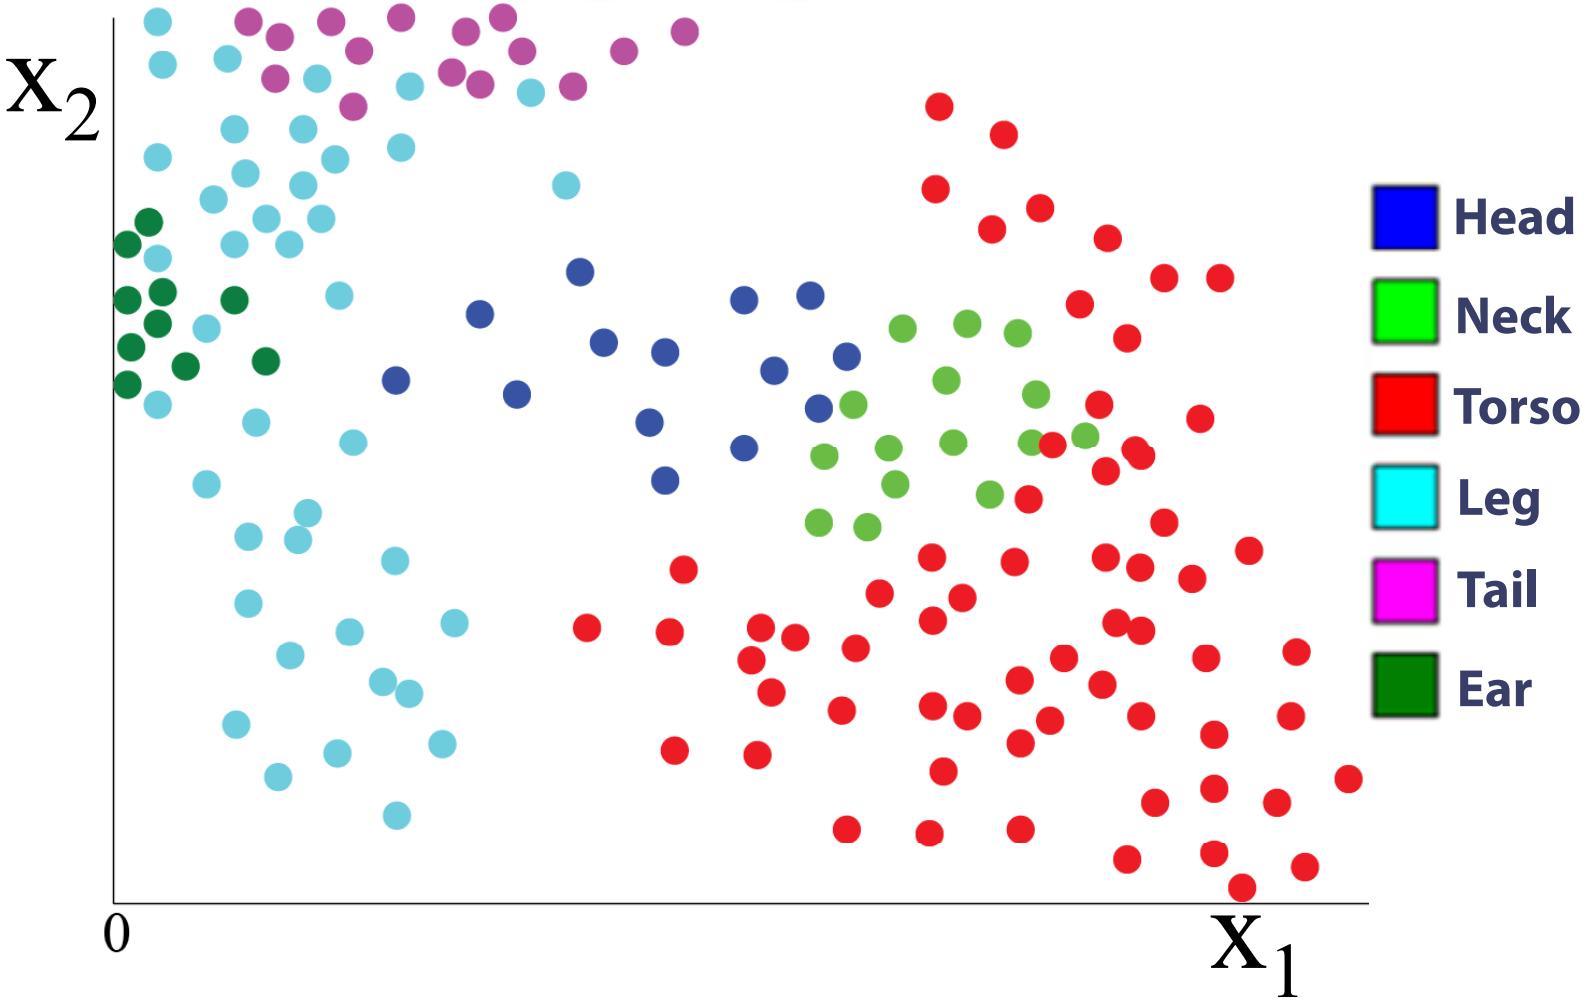

Learning a classifier

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

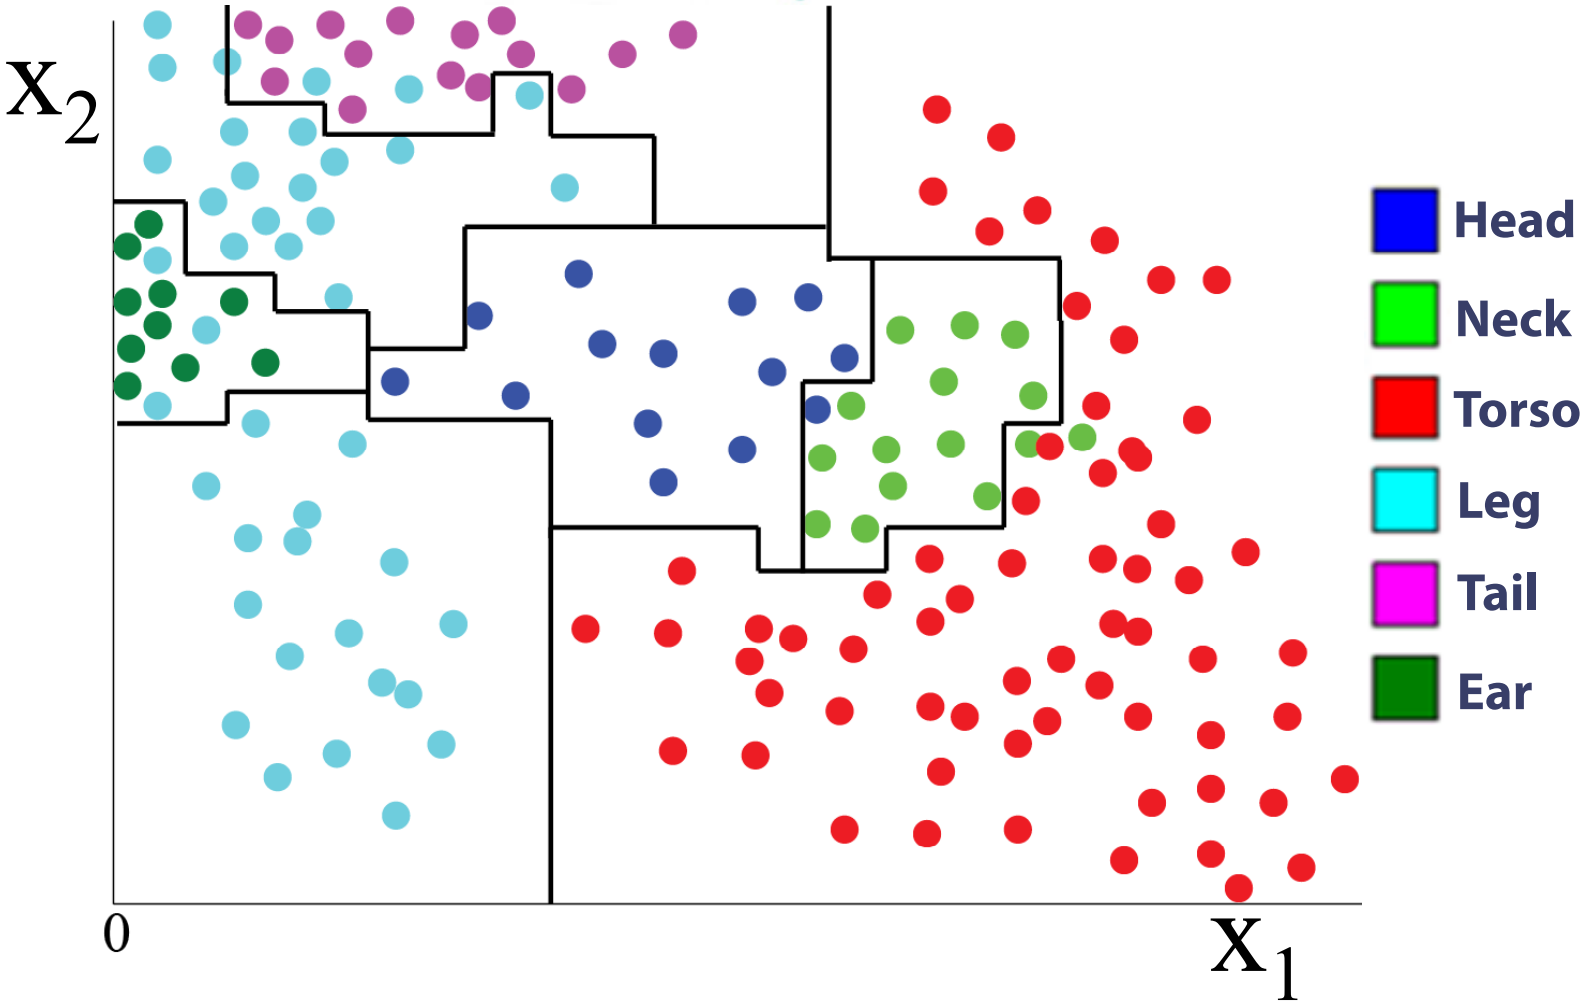

Learning a classifier

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

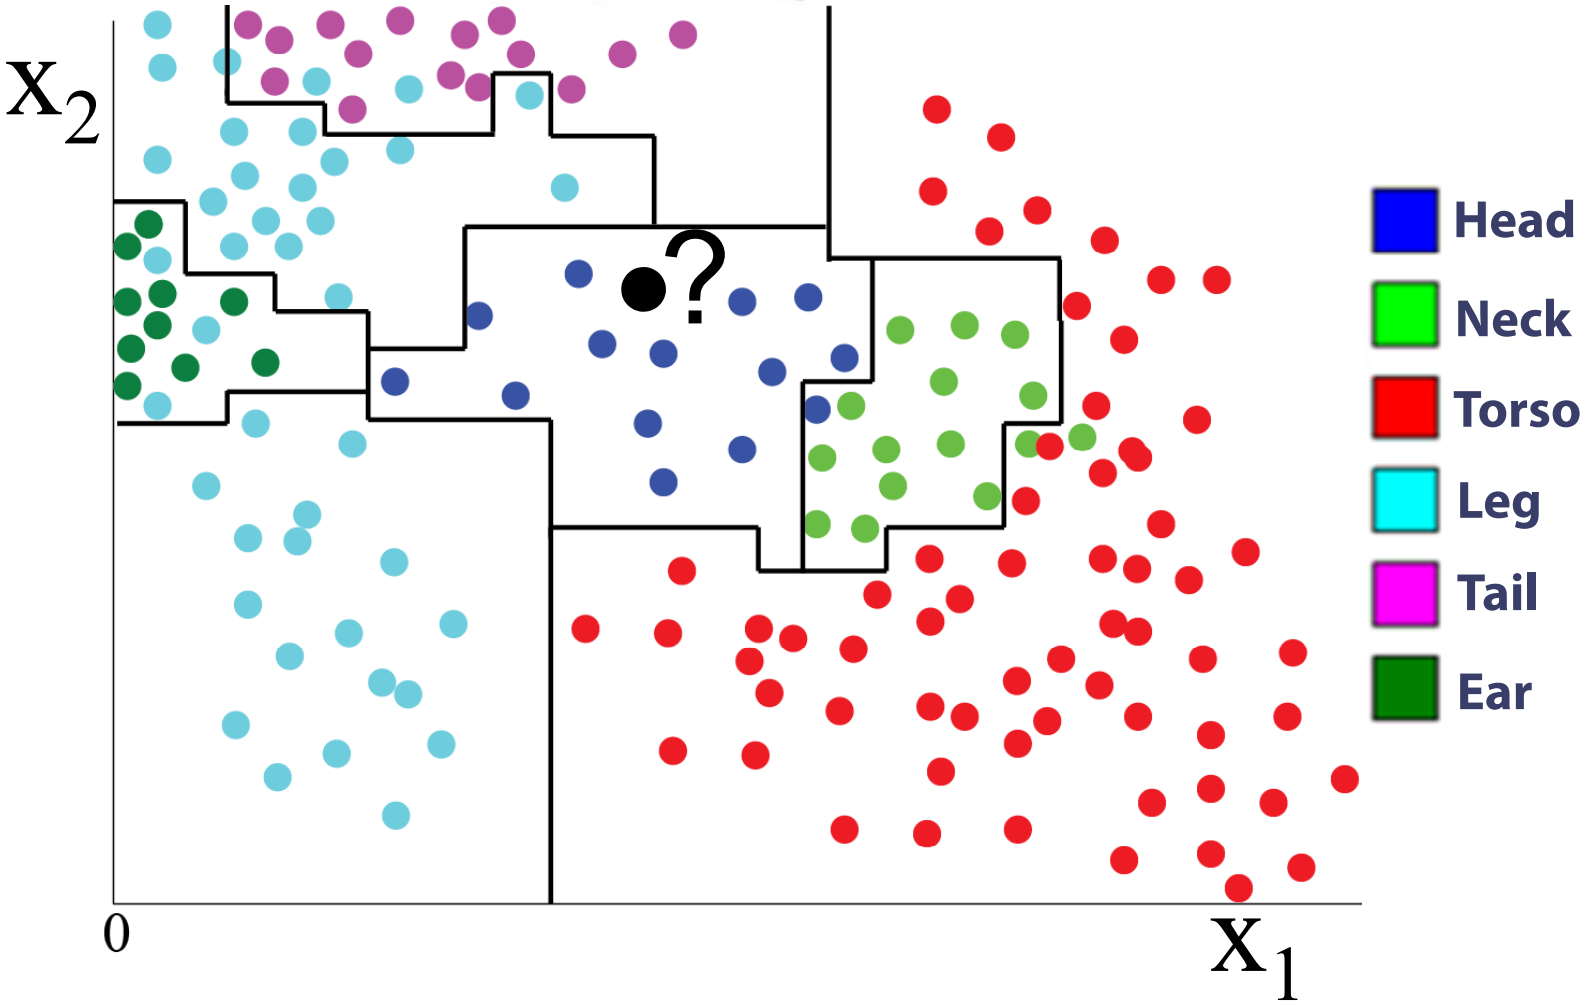

Learning a classifier

We use the Jointboost classifier [Torralba et al. 2007]

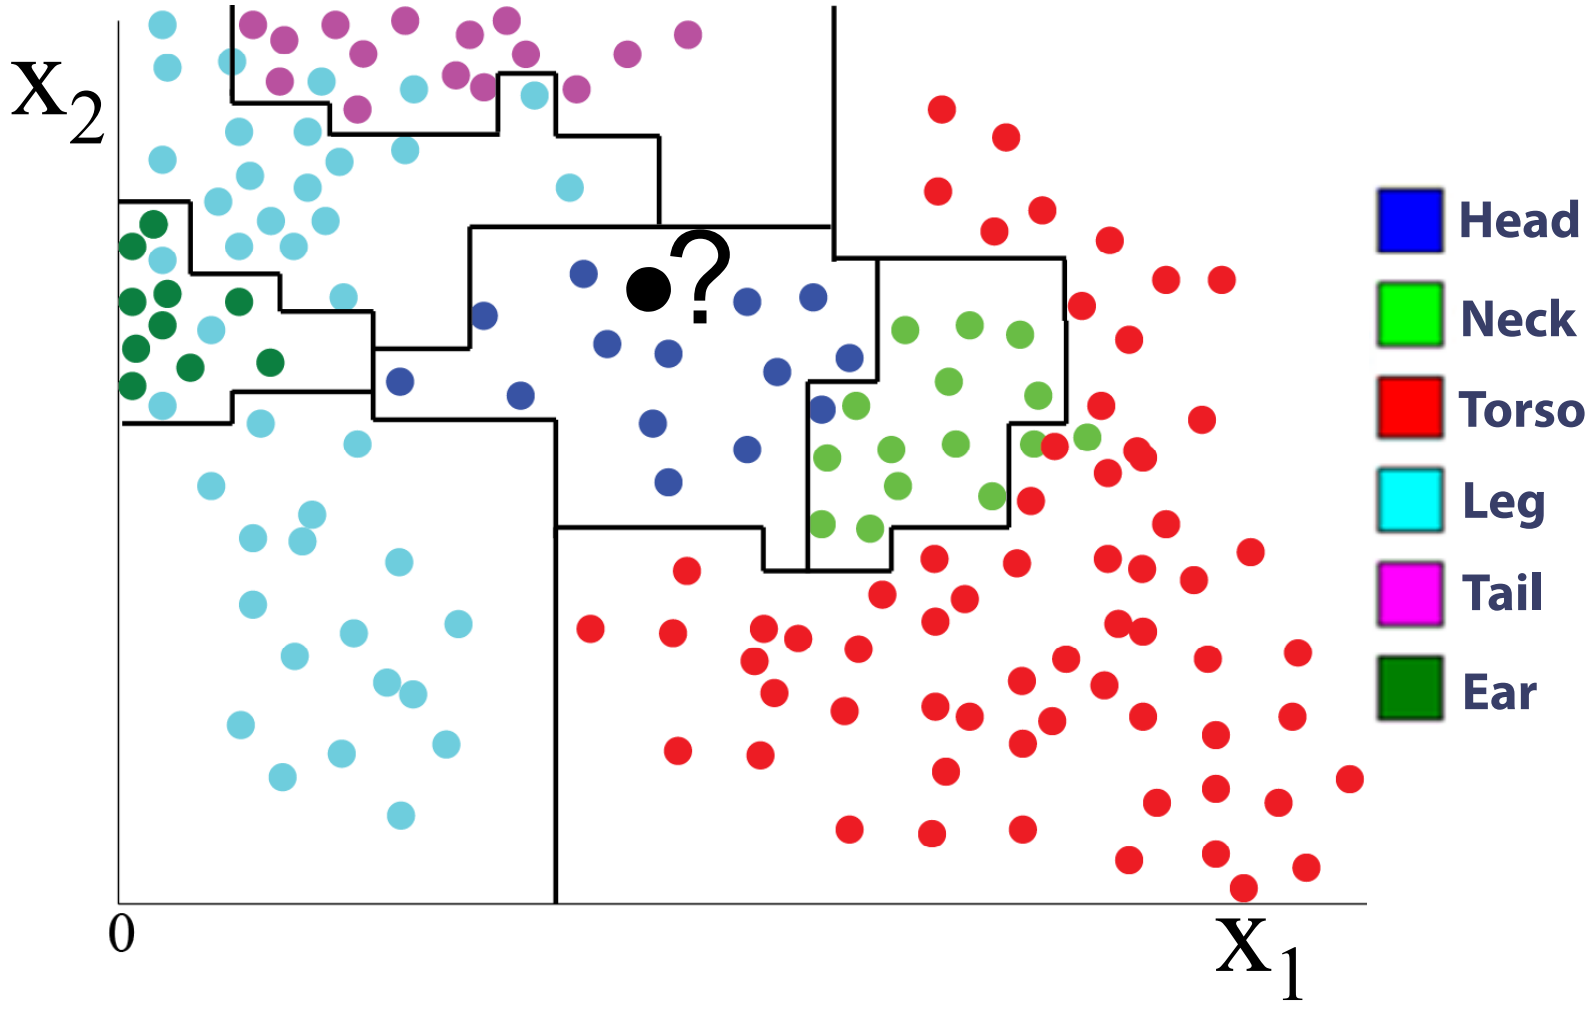

Unary term

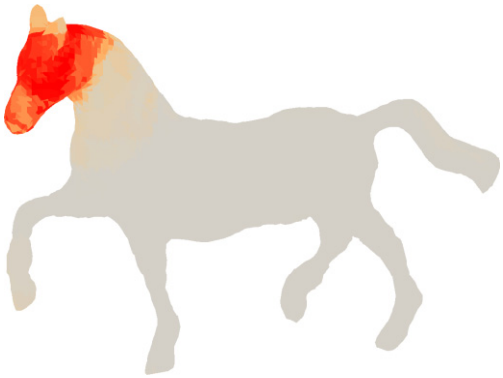

$\mathrm{P}(head \mid \mathbf{x})$

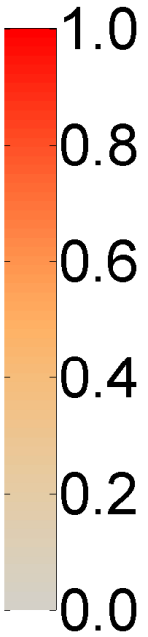

Unary term

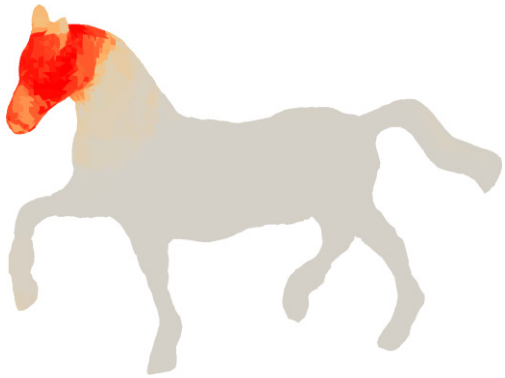

$\mathrm{P}(head \mid \mathbf{x})$

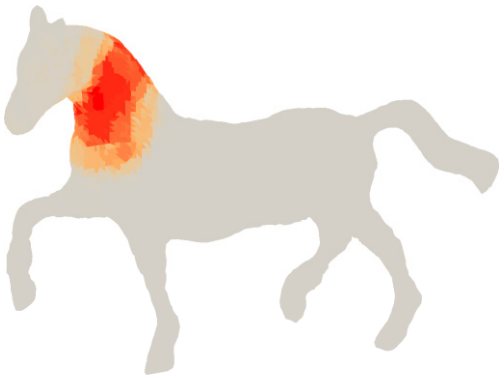

$\mathrm{P}(neck \mid \mathbf{x})$

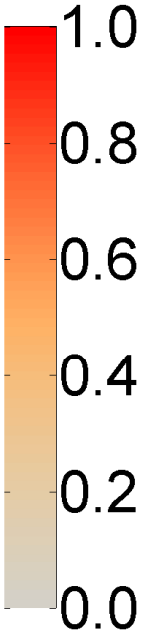

Unary term

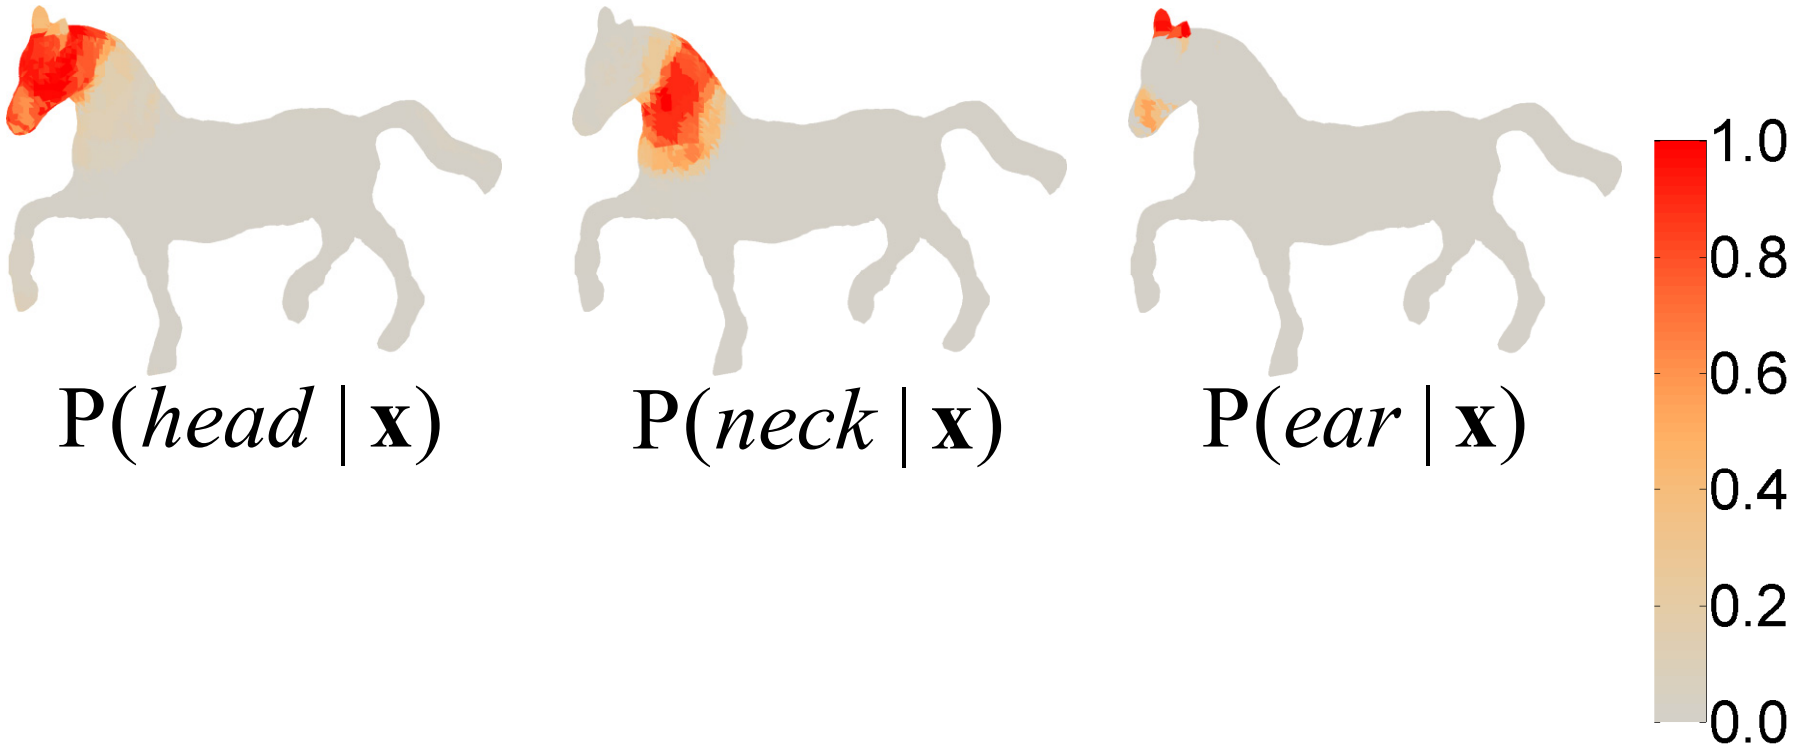

$\mathrm{P}(head \mid \mathbf{x})$ $\mathrm{P}(neck \mid \mathbf{x})$ $\mathrm{P}(ear \mid \mathbf{x})$

Unary term

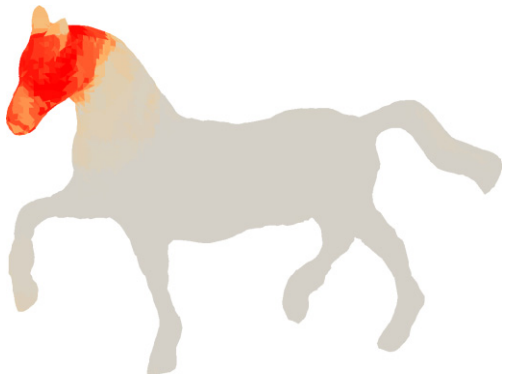

$\mathrm{P}(head \mid \mathbf{x})$

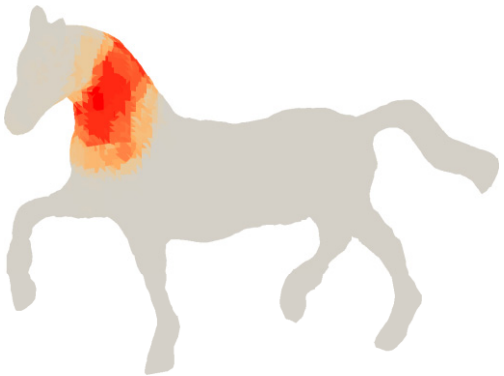

$\mathrm{P}(neck \mid \mathbf{x})$

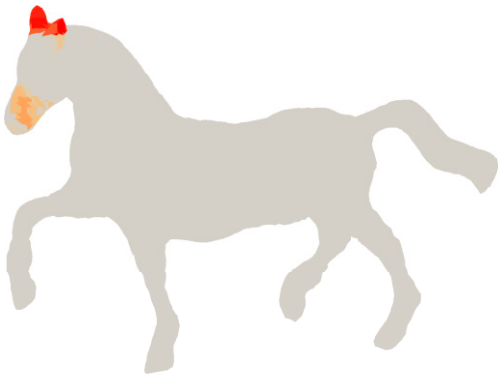

$\mathrm{P}(ear \mid \mathbf{x})$

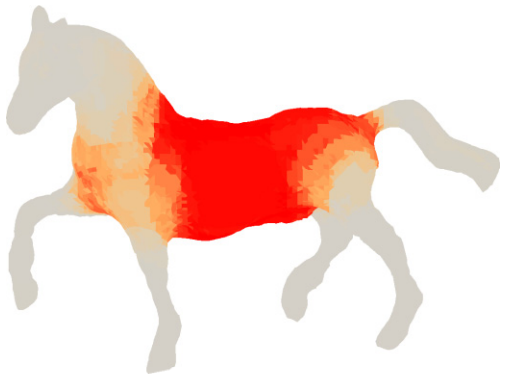

$\mathrm{P}(torso \mid \mathbf{x})$

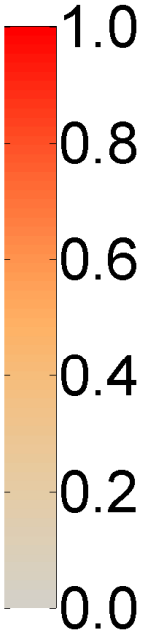

Unary term

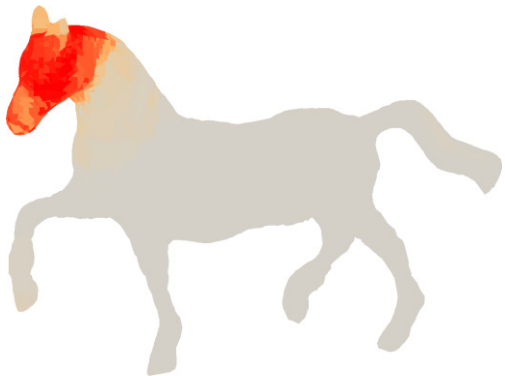

$\mathrm{P}(head \mid \mathbf{x})$

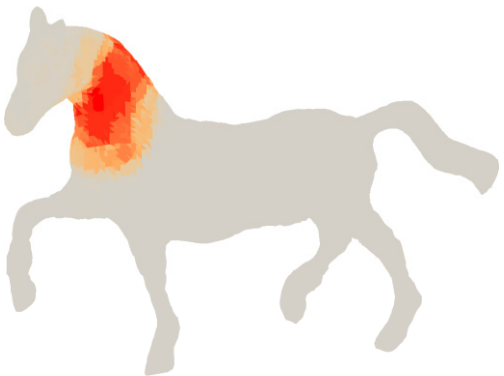

$\mathrm{P}(neck \mid \mathbf{x})$

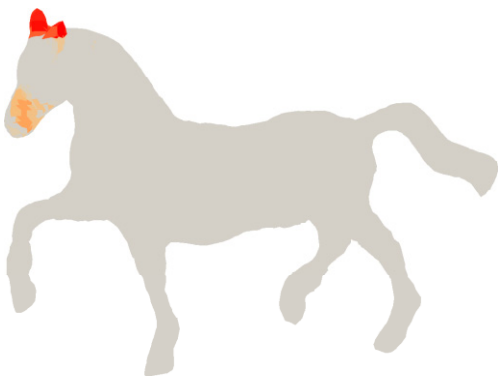

$\mathrm{P}(ear \mid \mathbf{x})$

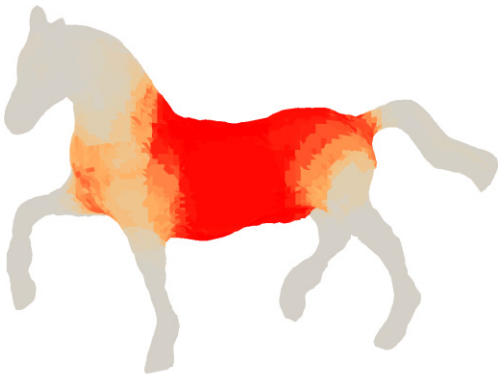

$\mathrm{P}(torso \mid \mathbf{x})$

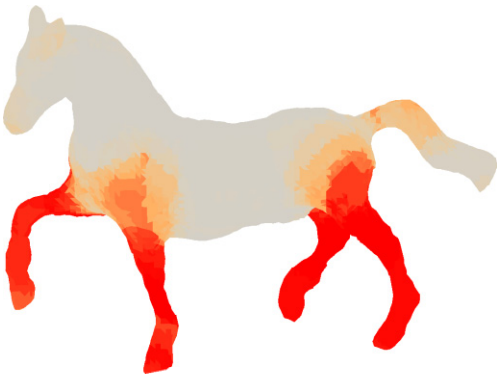

$\mathrm{P}(leg \mid \mathbf{x})$

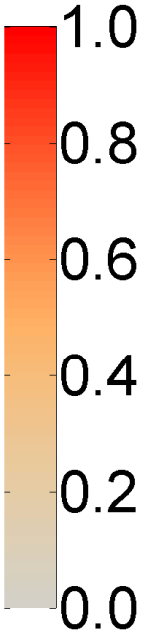

Unary term

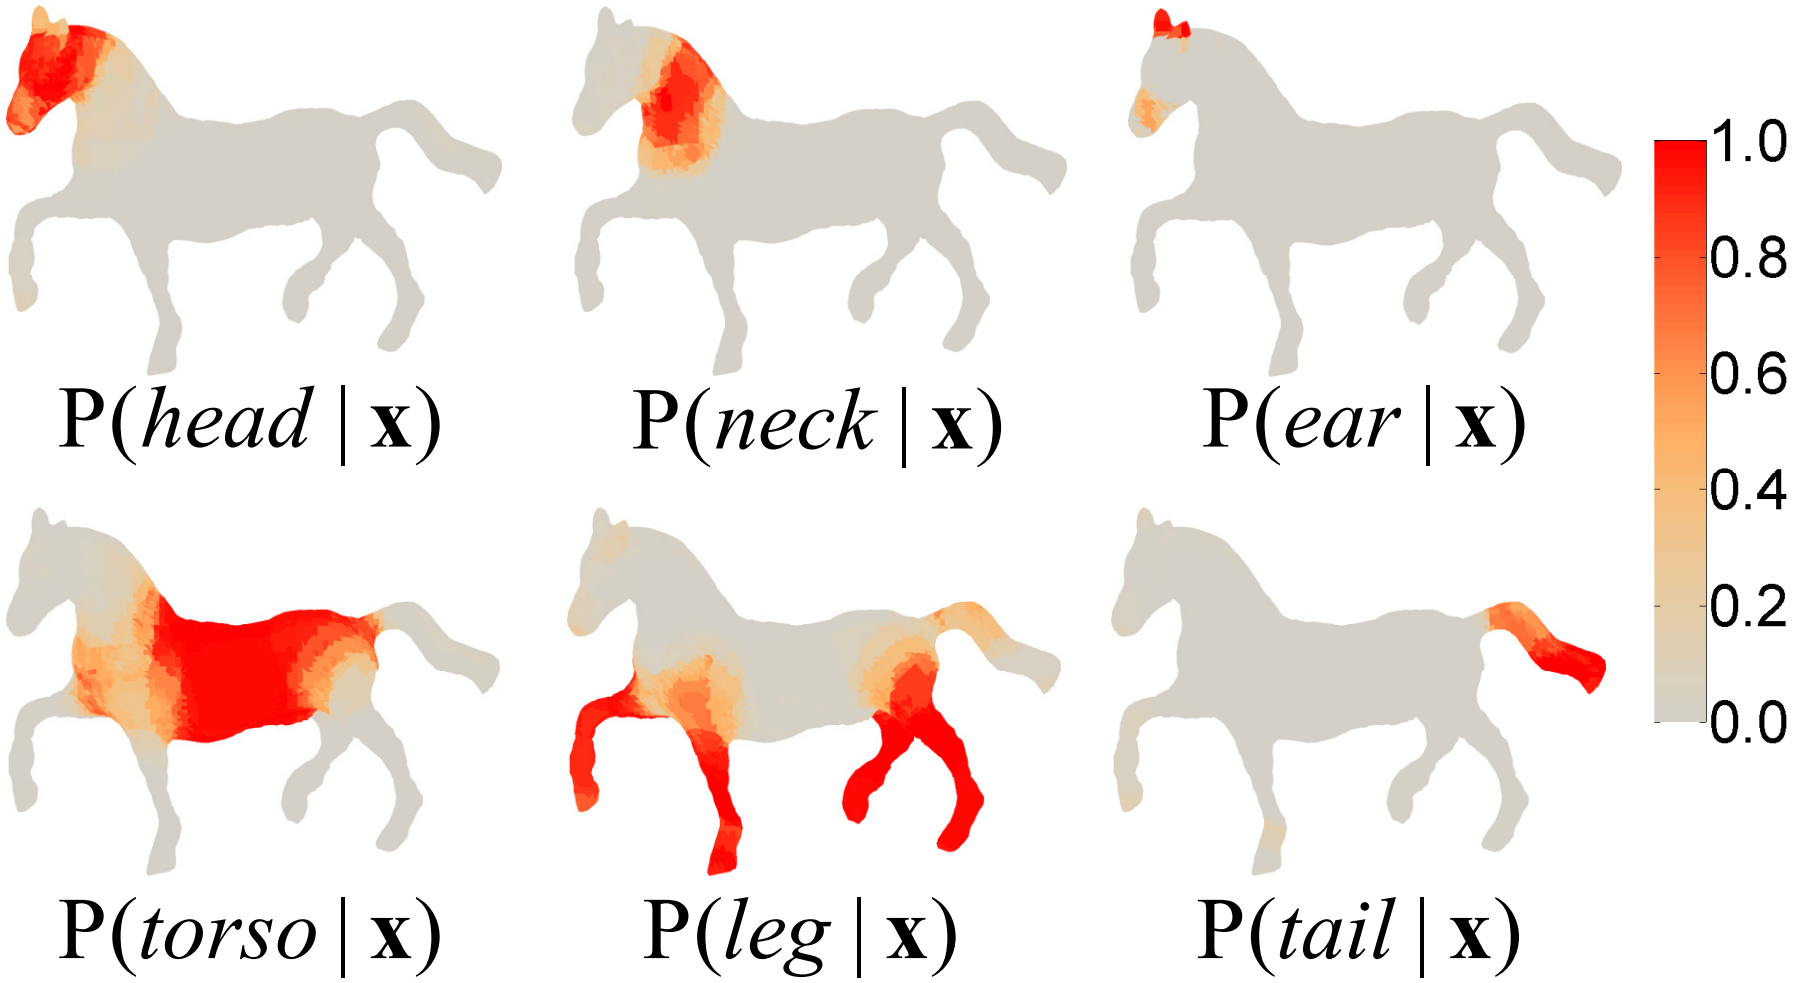

Unary term

$$\mathrm{E}_1(c;\mathbf{x}) = -\log \mathrm{P}(c \mid \mathbf{x})$$

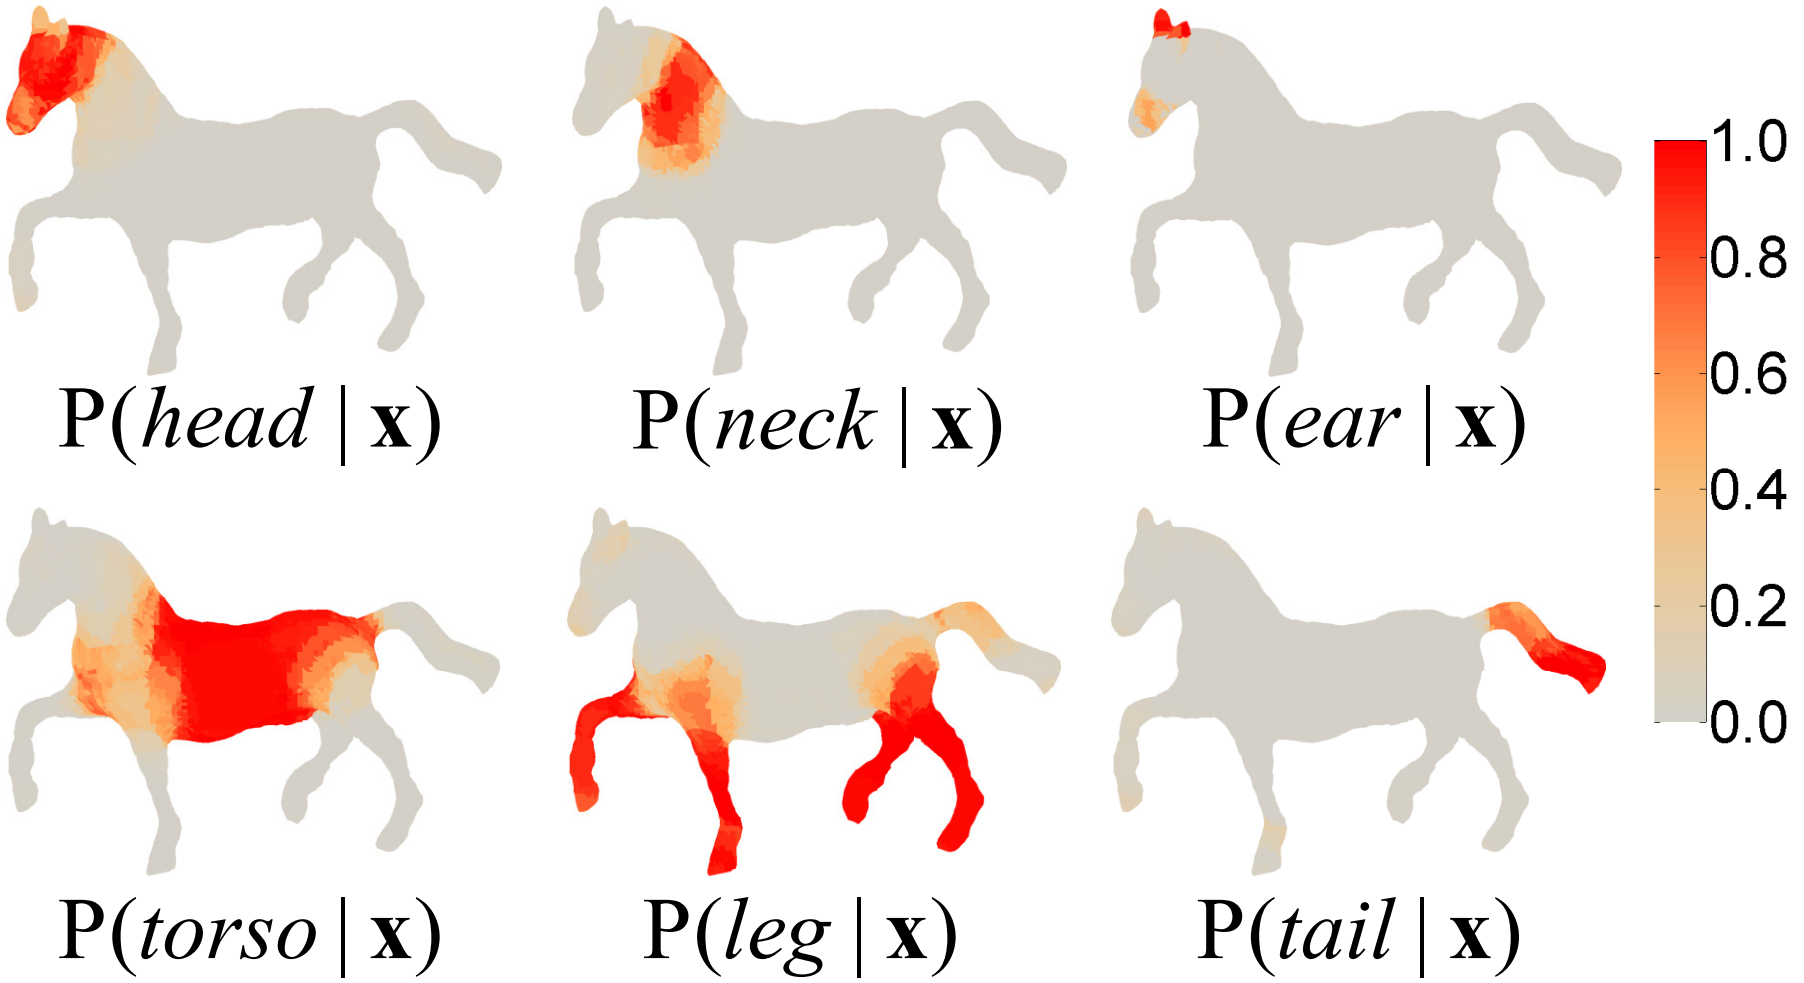

$\mathrm{P}(head \mid \mathbf{x})$ $\mathrm{P}(neck \mid \mathbf{x})$ $\mathrm{P}(ear \mid \mathbf{x})$

$\mathrm{P}(torso \mid \mathbf{x})$ $\mathrm{P}(leg \mid \mathbf{x})$ $\mathrm{P}(tail \mid \mathbf{x})$

Unary Term

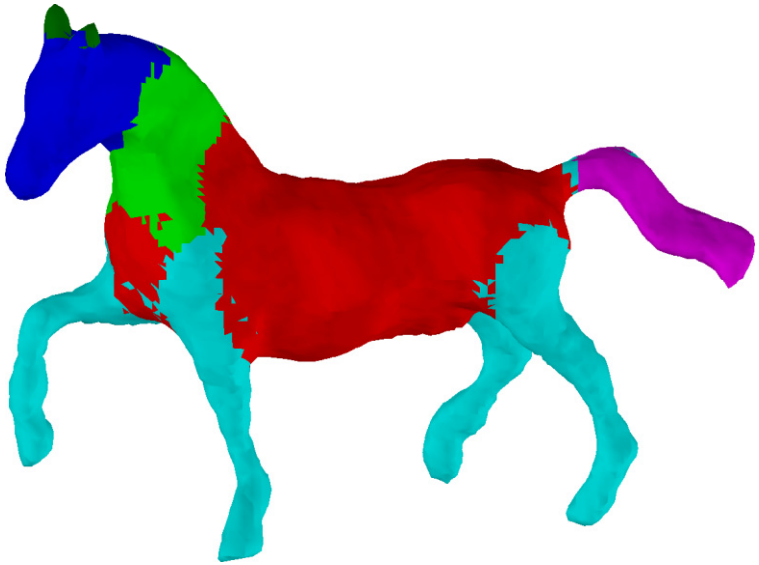

Most-likely labels

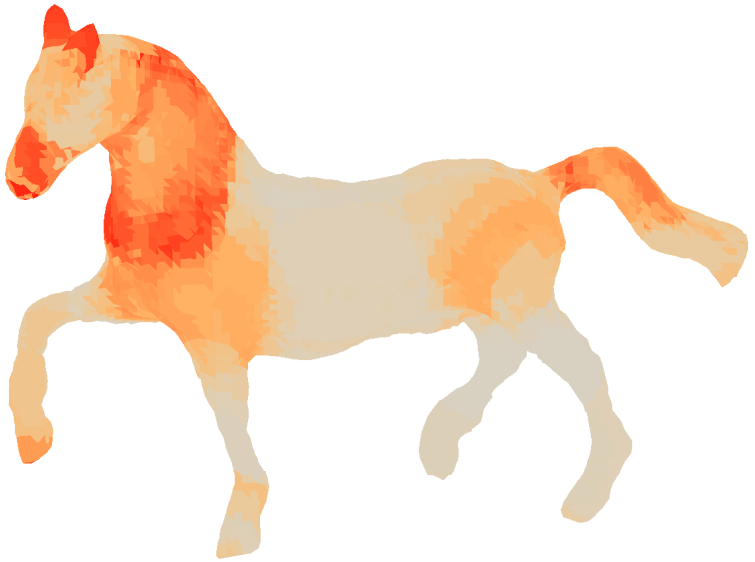

Classifier entropy

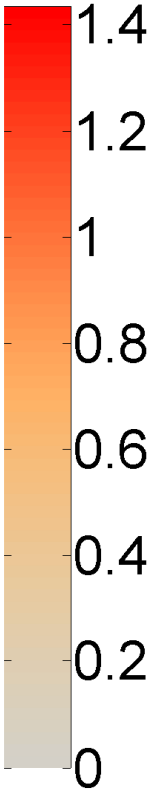

Our approach

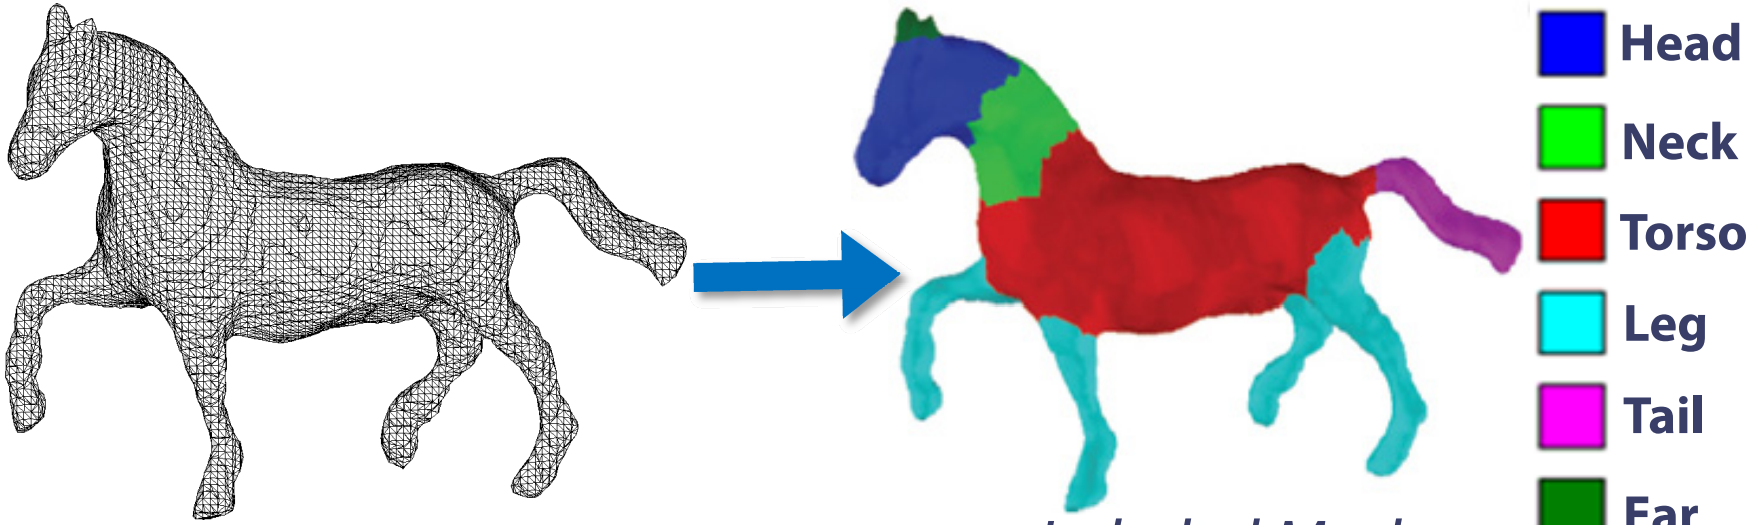

$$c^* = \arg\min_{\mathbf{c}} \left\{ \sum_i \alpha_i \mathrm{E}_1(c_i; \mathbf{x_i}) + \sum_{i,j} l_{ij} \mathrm{E}_2(c_i, c_j; \mathbf{y_{ij}}) \right\}$$

Pairwise Term

Pairwise Term

$$\mathrm{E}_2(c, c'; \mathbf{y}, \theta_2) = G(\mathbf{y}) L(c, c')$$

Geometry-dependent term

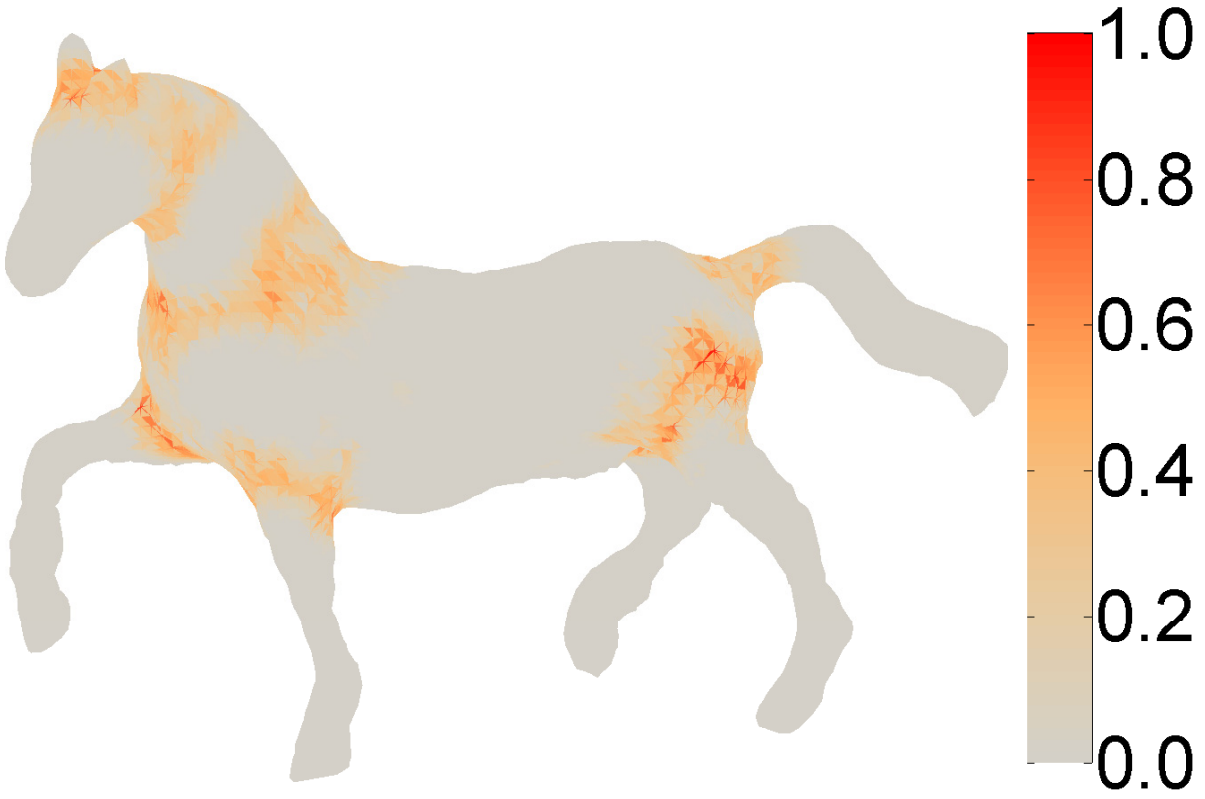

Pairwise Term

$$\mathrm{E}_2(c, c'; \mathbf{y}, \theta_2) = G(\mathbf{y}) L(c, c')$$

Label compatibility term

Pairwise Term

$$\mathrm{E}_2(c, c'; \mathbf{y}, \theta_2) = G(\mathbf{y}) L(c, c')$$

Label compatibility term

$L(c, c') =$

	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**

Pairwise Term

$$\mathrm{E}_2(c, c'; \mathbf{y}, \theta_2) = G(\mathbf{y}) L(c, c')$$

Label compatibility term

$$L(c, c') = \begin{array}{c} \\ \text{Head} \\ \text{Neck} \\ \text{Ear} \\ \text{Torso} \\ \text{Leg} \\ \text{Tail} \end{array} \begin{array}{c} \begin{array}{cccccc} \text{Head} & \text{Neck} & \text{Ear} & \text{Torso} & \text{Leg} & \text{Tail} \end{array} \\ \begin{bmatrix} 0 & .45 & .07 & 1 & \infty & \infty \\ .45 & 0 & \infty & 1 & \infty & \infty \\ .07 & \infty & 0 & \infty & \infty & \infty \\ 1 & 1 & \infty & 0 & 1 & .56 \\ \infty & \infty & \infty & 1 & 0 & \infty \\ \infty & \infty & \infty & .56 & \infty & 0 \end{bmatrix} \end{array}$$

Pairwise Term

$$\mathrm{E}_2(c,c';\mathbf{y},\theta_2) = G(\mathbf{y})L(c,c')$$

Label compatibility term

$$L(c,c') = \begin{array}{c} \\ \text{Head} \\ \text{Neck} \\ \text{Ear} \\ \text{Torso} \\ \text{Leg} \\ \text{Tail} \end{array} \begin{array}{c} \begin{array}{cccccc} \text{Head} & \text{Neck} & \text{Ear} & \text{Torso} & \text{Leg} & \text{Tail} \end{array} \\ \begin{bmatrix} 0 & .45 & .07 & 1 & \infty & \infty \\ .45 & 0 & \infty & 1 & \infty & \infty \\ .07 & \infty & 0 & \infty & \infty & \infty \\ 1 & 1 & \infty & 0 & 1 & .56 \\ \infty & \infty & \infty & 1 & 0 & \infty \\ \infty & \infty & \infty & .56 & \infty & 0 \end{bmatrix} \end{array}$$

Full CRF result

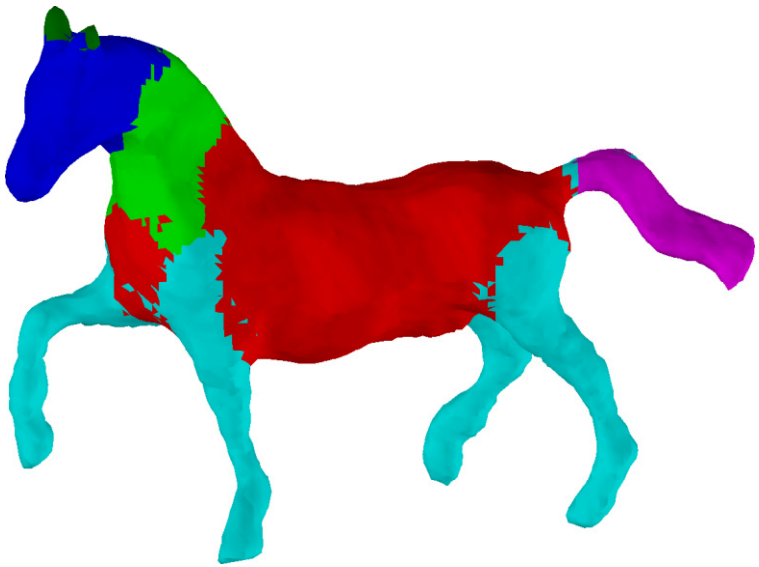

Unary term classifier

Full CRF result

Learning

Learning

Learn unary classifier and $G(\mathbf{y})$ with Joint Boosting
[Torralba et al. 2007]

Learning

Learn unary classifier and $G(\mathbf{y})$ with Joint Boosting
[Torralba et al. 2007]

Hold-out validation for the rest of parameters

Dataset used in experiments

We label 380 meshes from the Princeton Segmentation Benchmark **[Chen *et al.* 2009]**

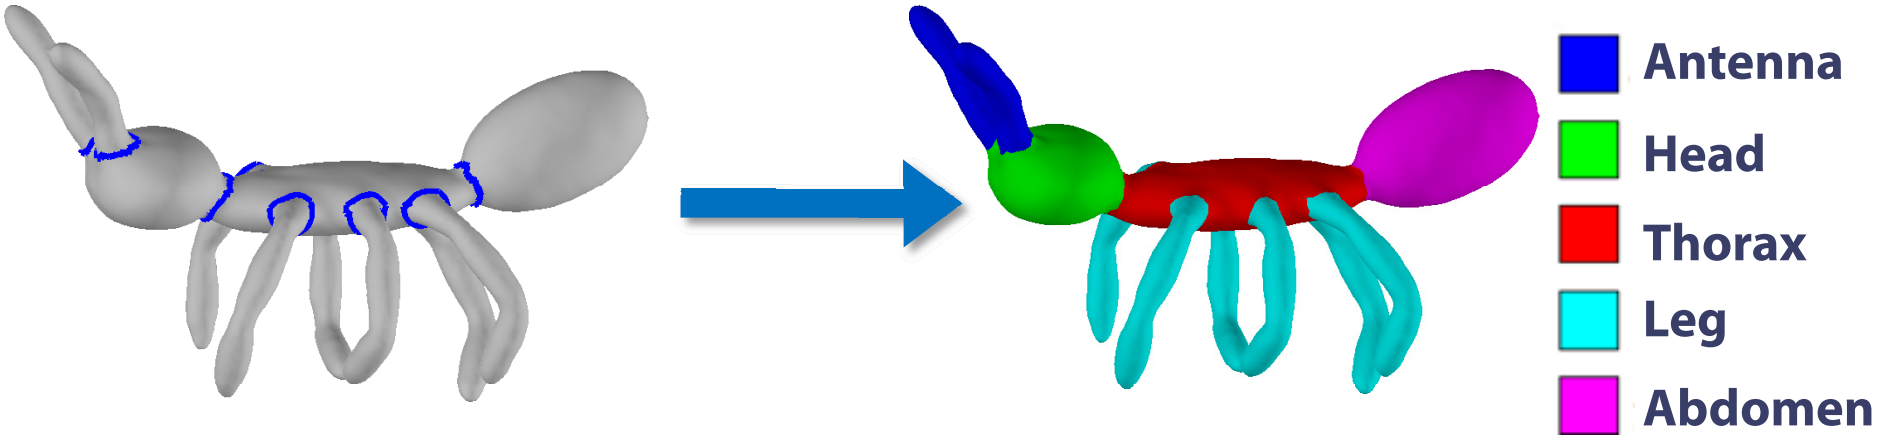

Dataset used in experiments

We label 380 meshes from the Princeton Segmentation Benchmark **[Chen *et al.* 2009]**

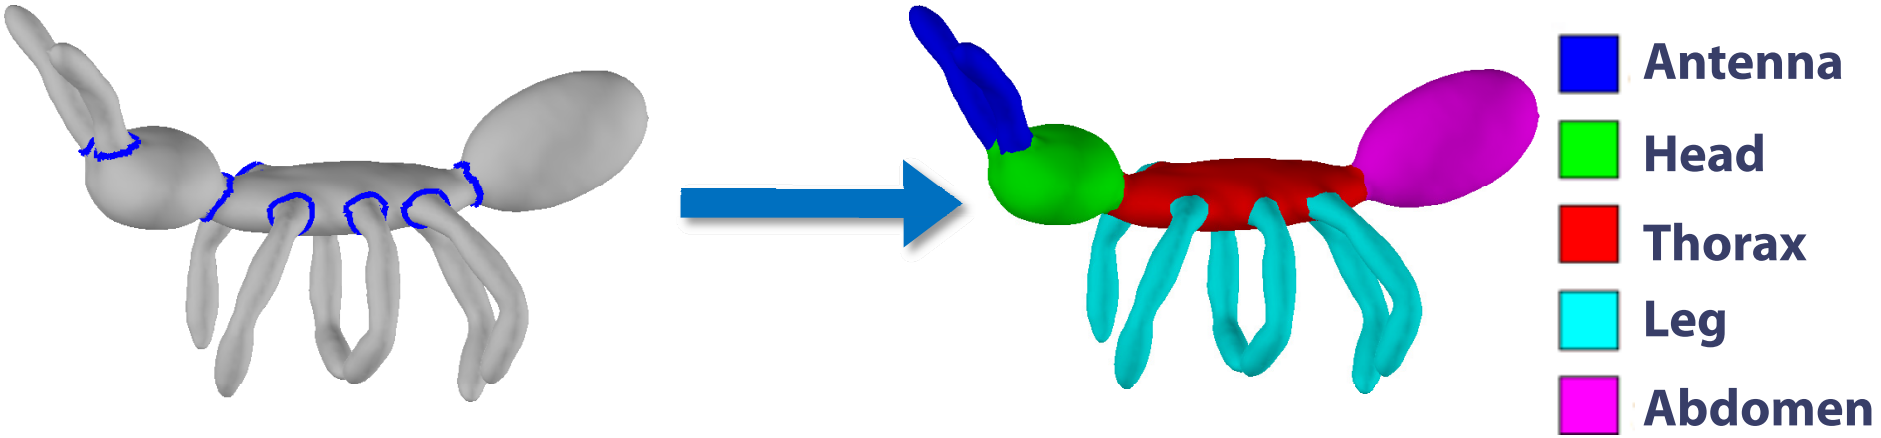

Each of the 19 categories is treated separately

Quantitative Evaluation

Labeling

- **6%** error by surface area

Quantitative Evaluation

Labeling

- **6%** error by surface area
- No previous automatic method

Quantitative Evaluation

Labeling

- **6%** error by surface area
- No previous automatic method

Segmentation

- Our result: **9.5%** Rand Index error

Quantitative Evaluation

Labeling

- **6%** error by surface area
- No previous automatic method

Segmentation

- Our result: **9.5%** Rand Index error
- State-of-the art: **16%** **[Golovinskiy and Funkhouser 08]**

Quantitative Evaluation

Labeling

- **6%** error by surface area
- No previous automatic method

Segmentation

- Our result: **9.5%** Rand Index error
- State-of-the art: **16%** **[Golovinskiy and Funkhouser 08]**
- With 6 training meshes: **12%**

Quantitative Evaluation

Labeling

- **6%** error by surface area
- No previous automatic method

Segmentation

- Our result: **9.5%** Rand Index error
- State-of-the art: **16%** **[Golovinskiy and Funkhouser 08]**
- With 6 training meshes: **12%**
- With 3 training meshes: **15%**

Labeling results

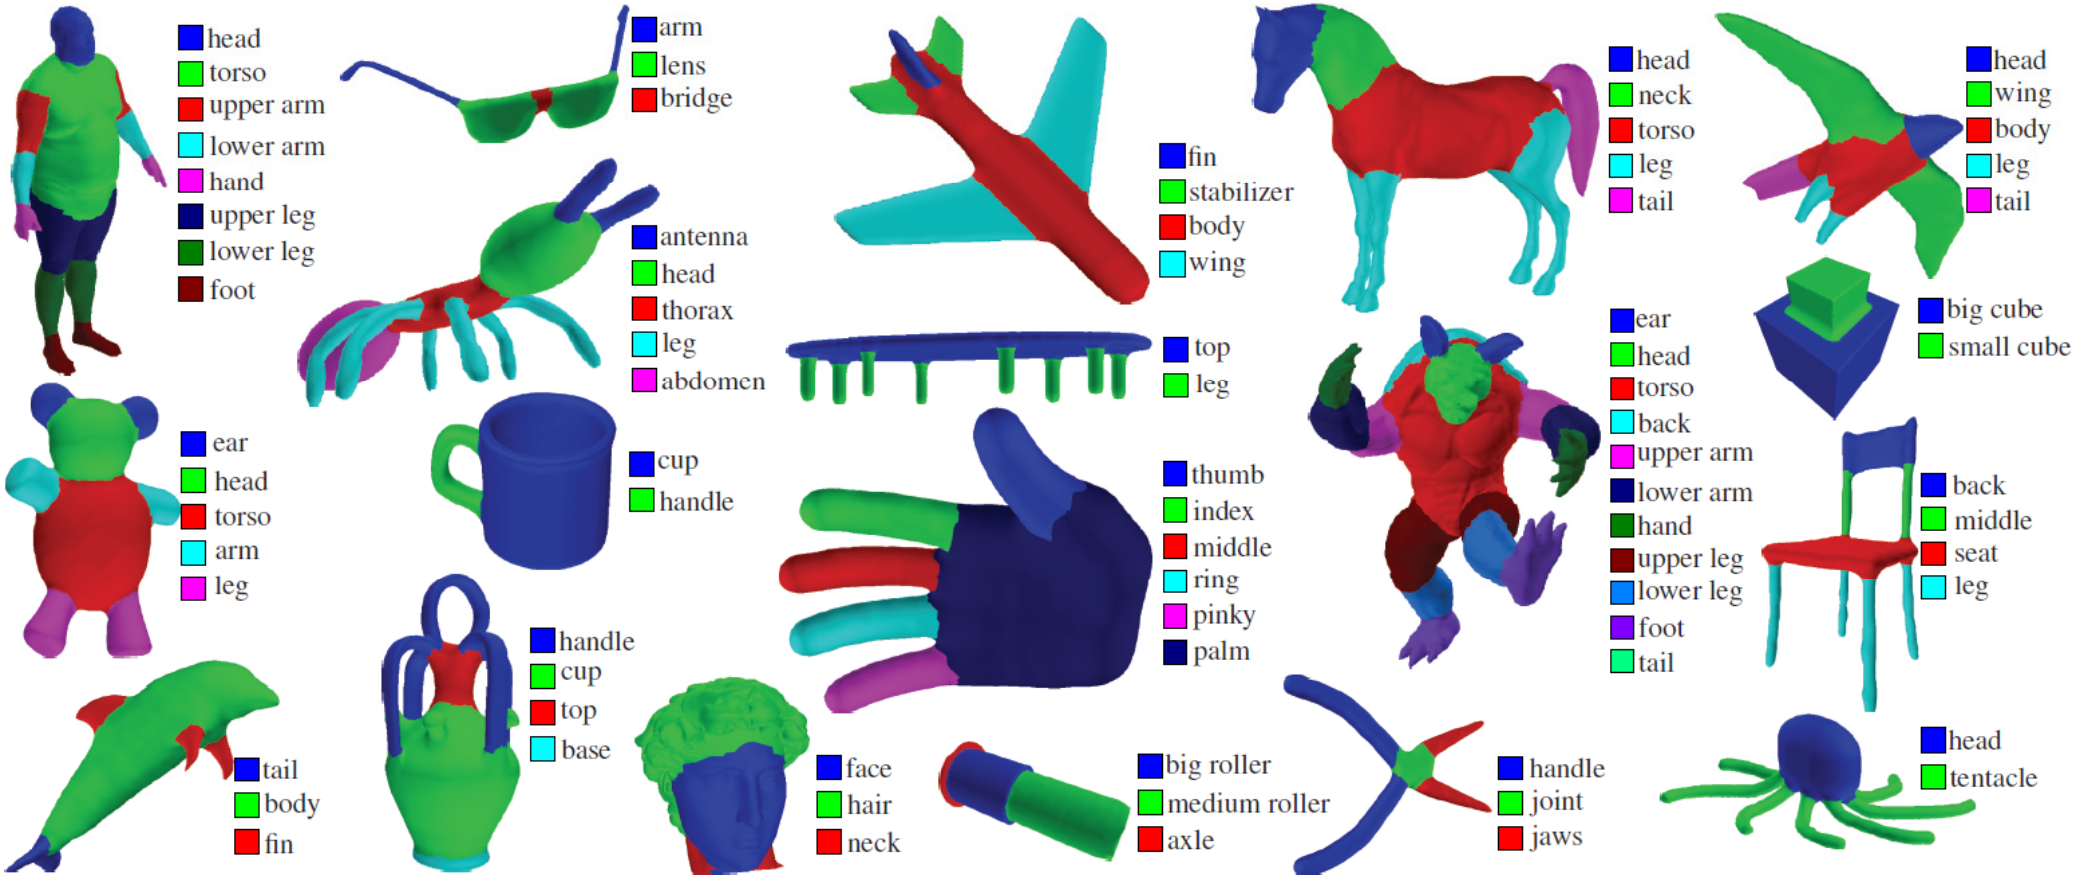

Segmentation Comparisons

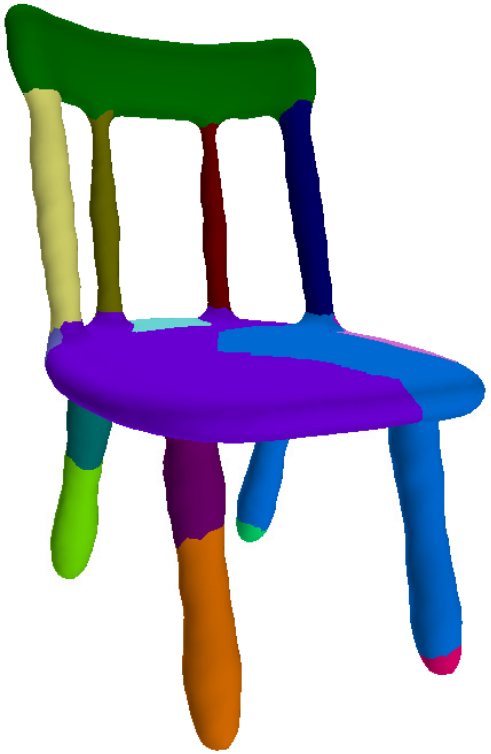

Shape Diameter [Shapira et al. 10]

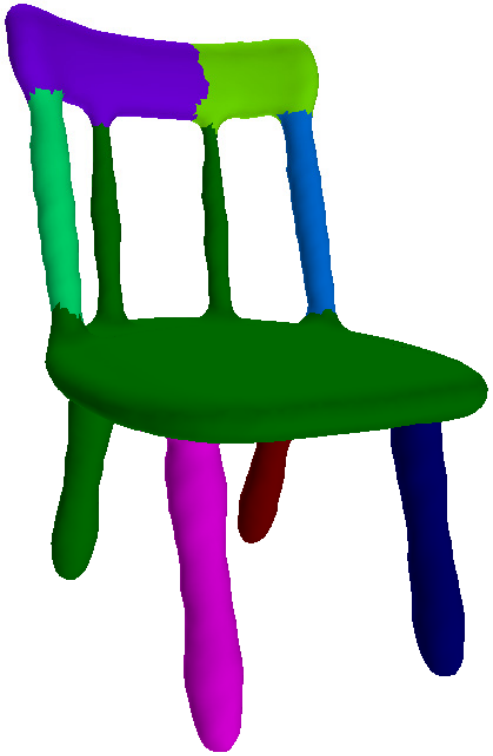

Randomized Cuts [Golovinskiy and Funkhouser 08]

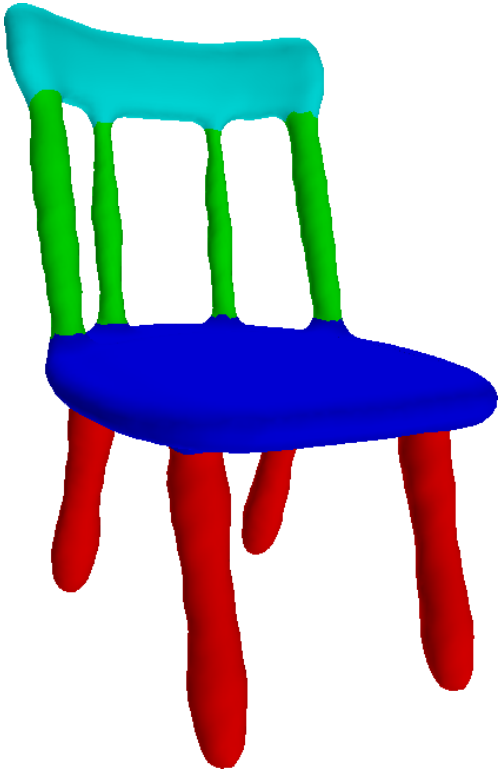

Our approach

Segmentation Comparisons

Shape Diameter [Shapira et al. 10]

Randomized Cuts [Golovinskiy and Funkhouser 08]

Our approach

Learning different segmentation styles

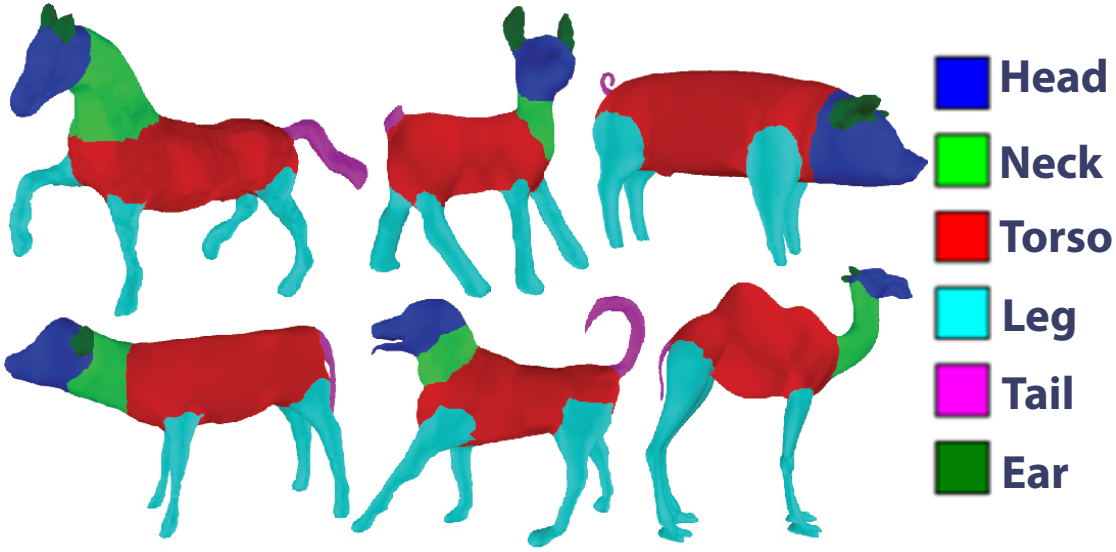

Training Meshes

Learning different segmentation styles

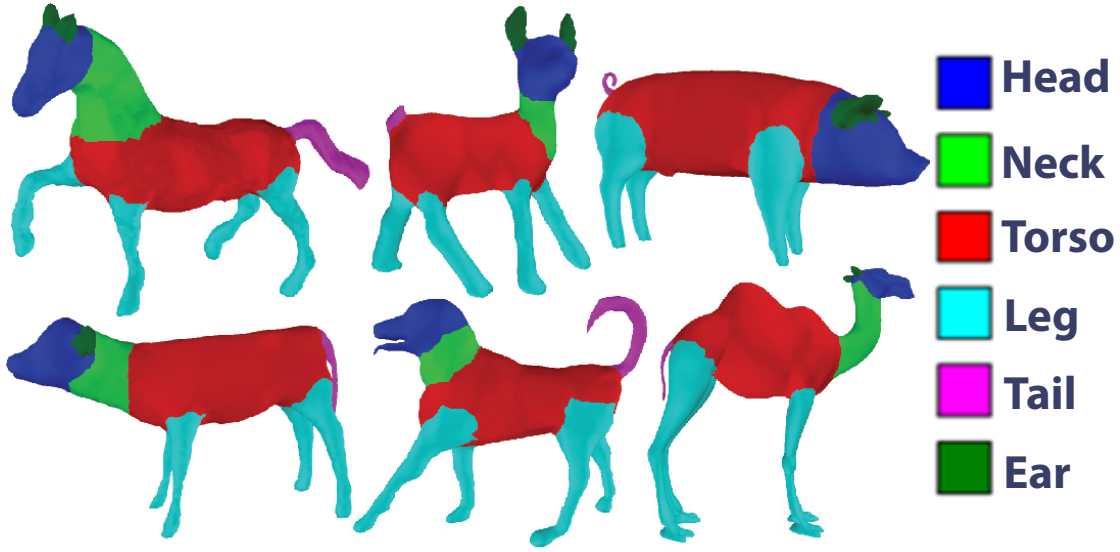

Training Meshes

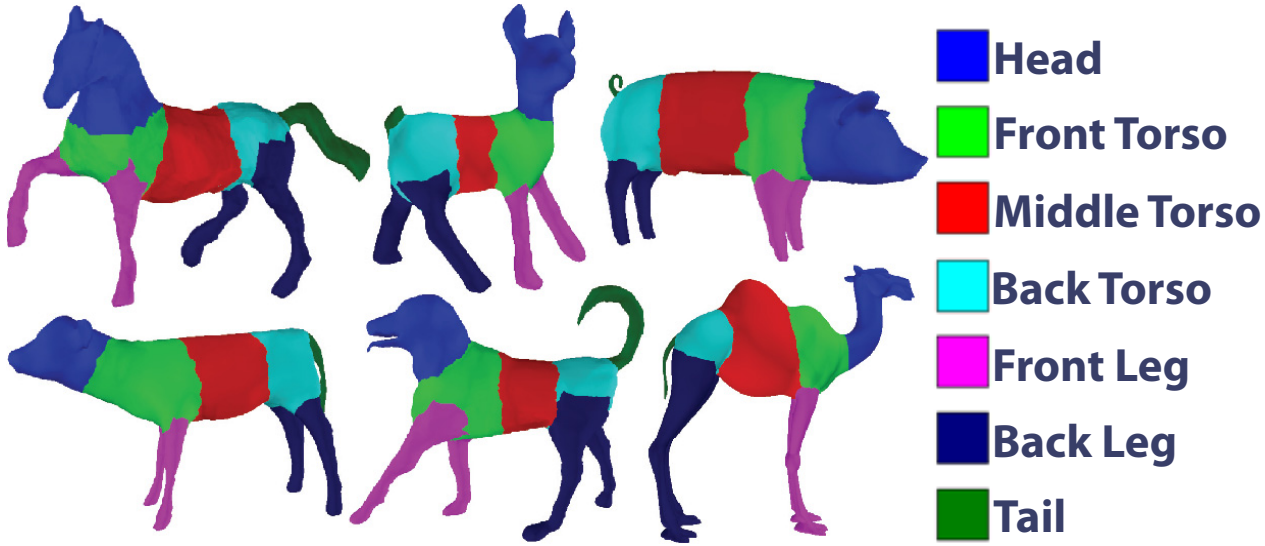

Learning different segmentation styles

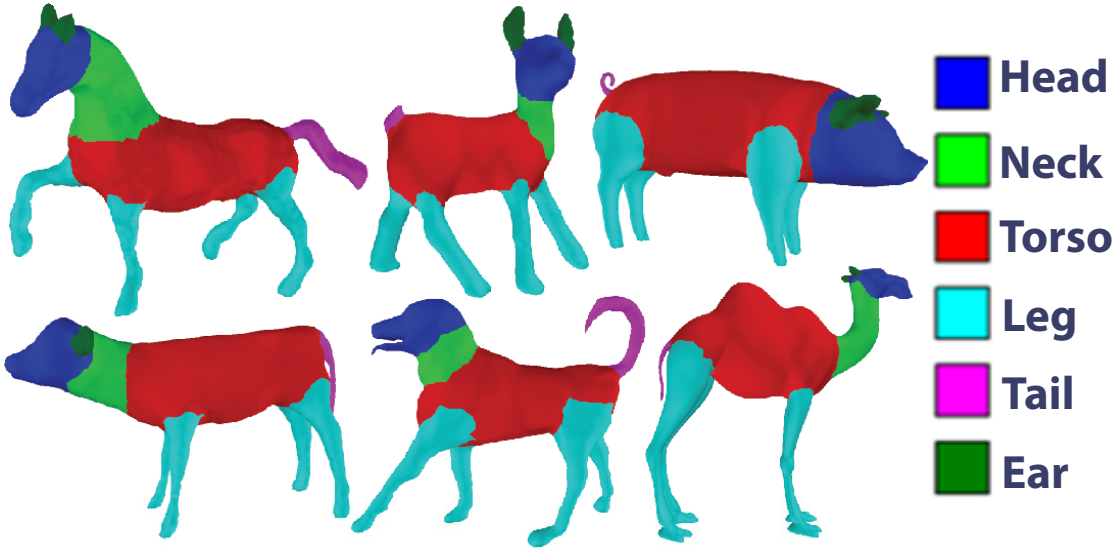

Head
Neck
Torso
Leg
Tail
Ear

Training Meshes

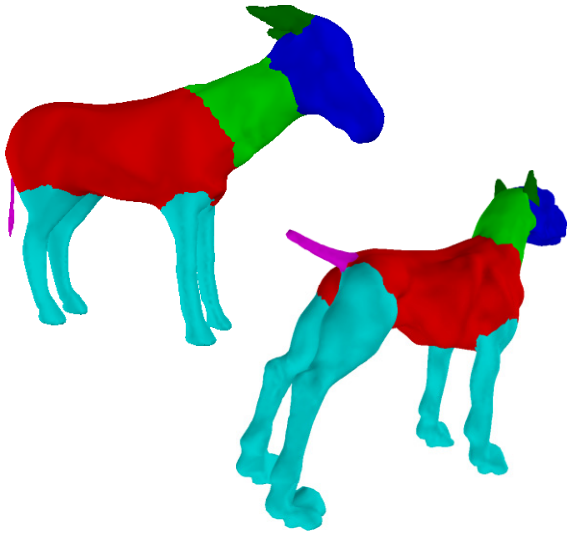

Test Meshes

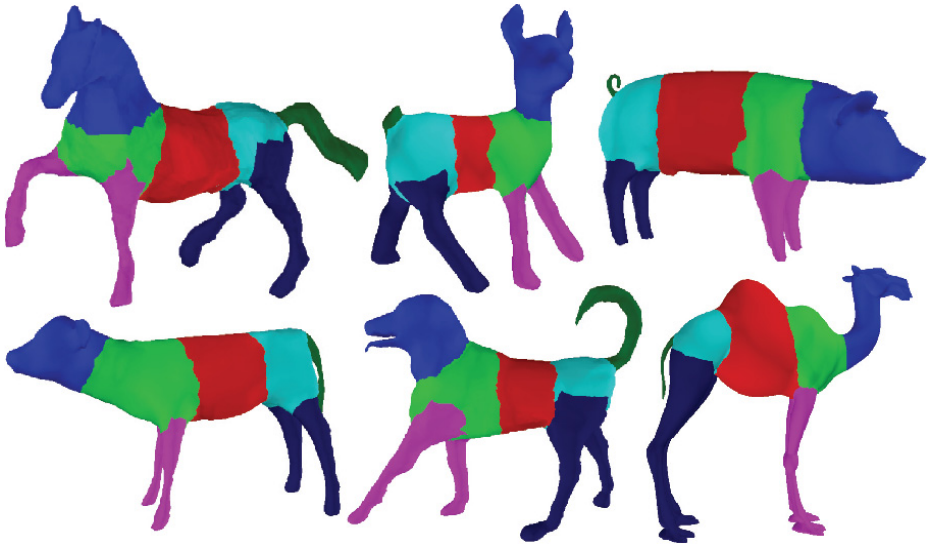

Head
Front Torso
Middle Torso
Back Torso
Front Leg
Back Leg
Tail

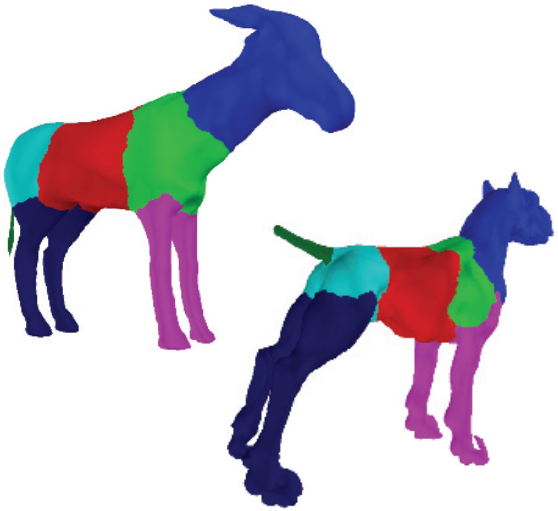

Generalization to different categories

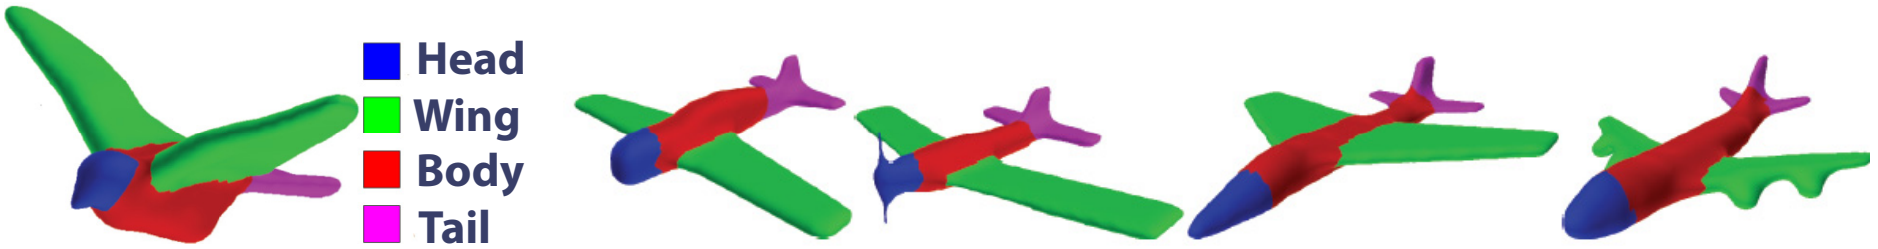

Generalization to different categories

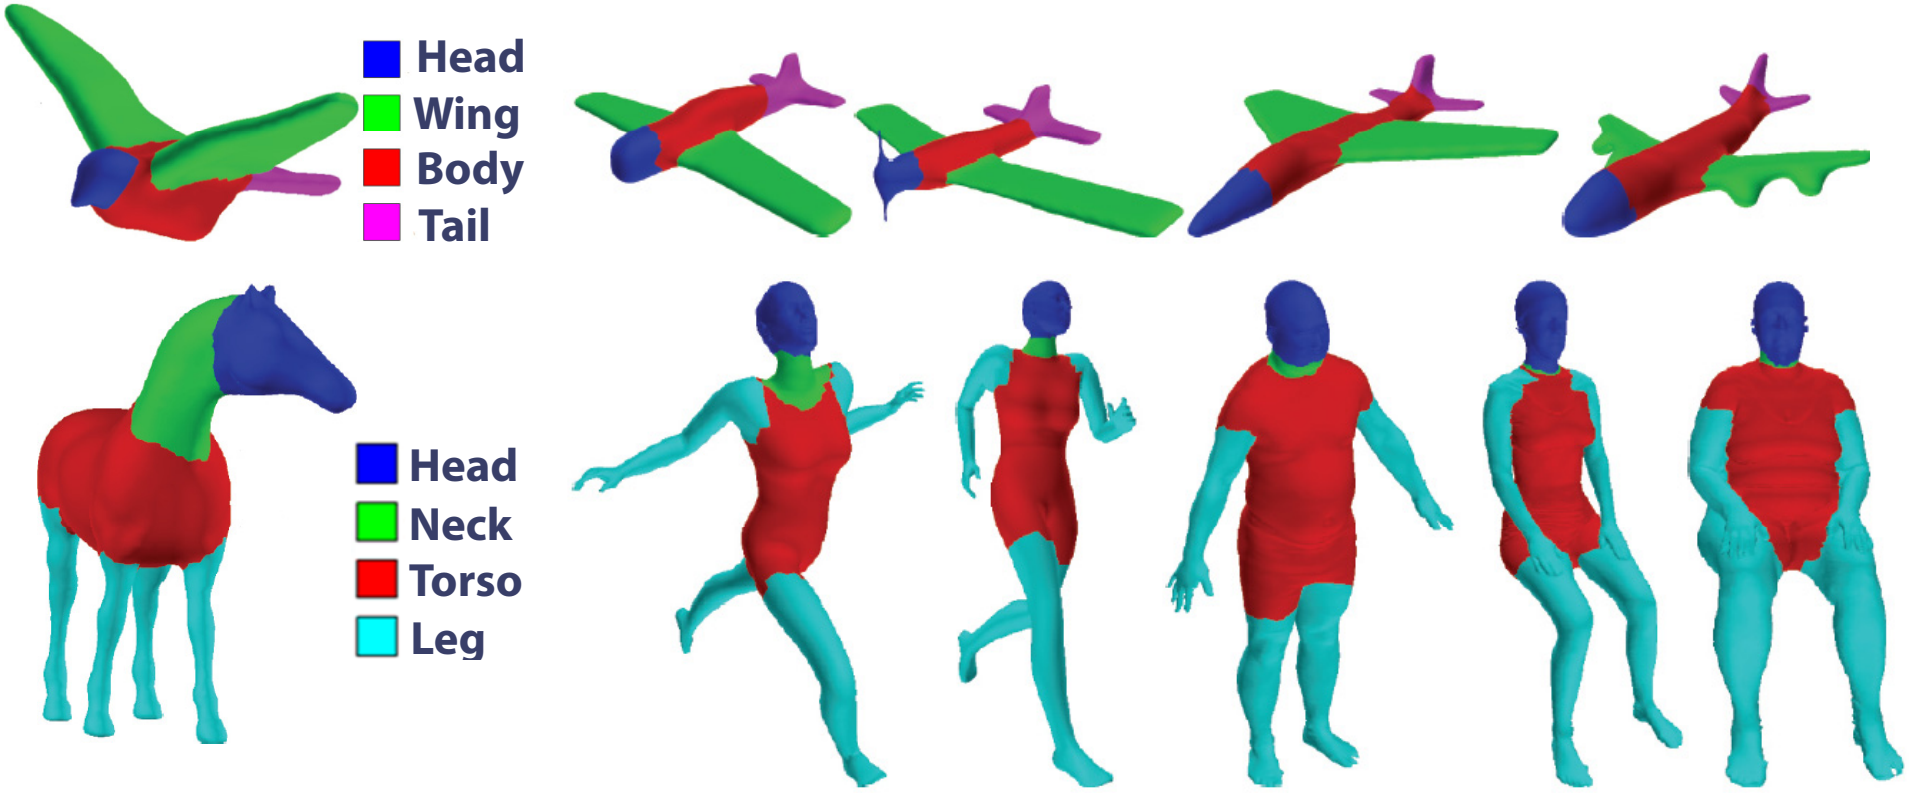

Failure cases

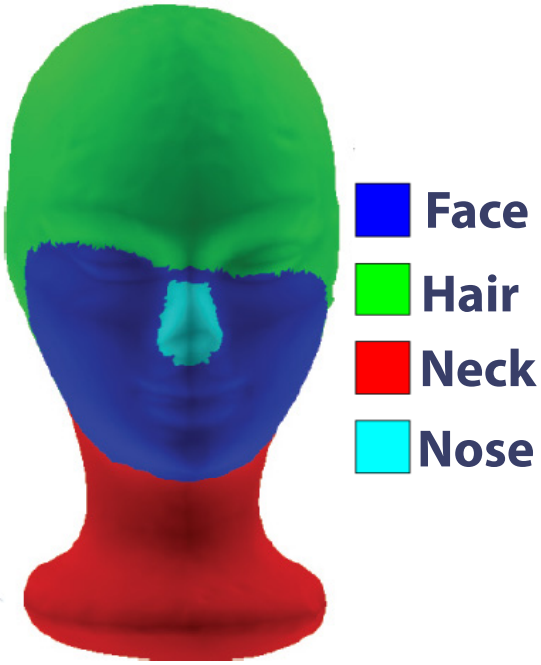

Failure cases

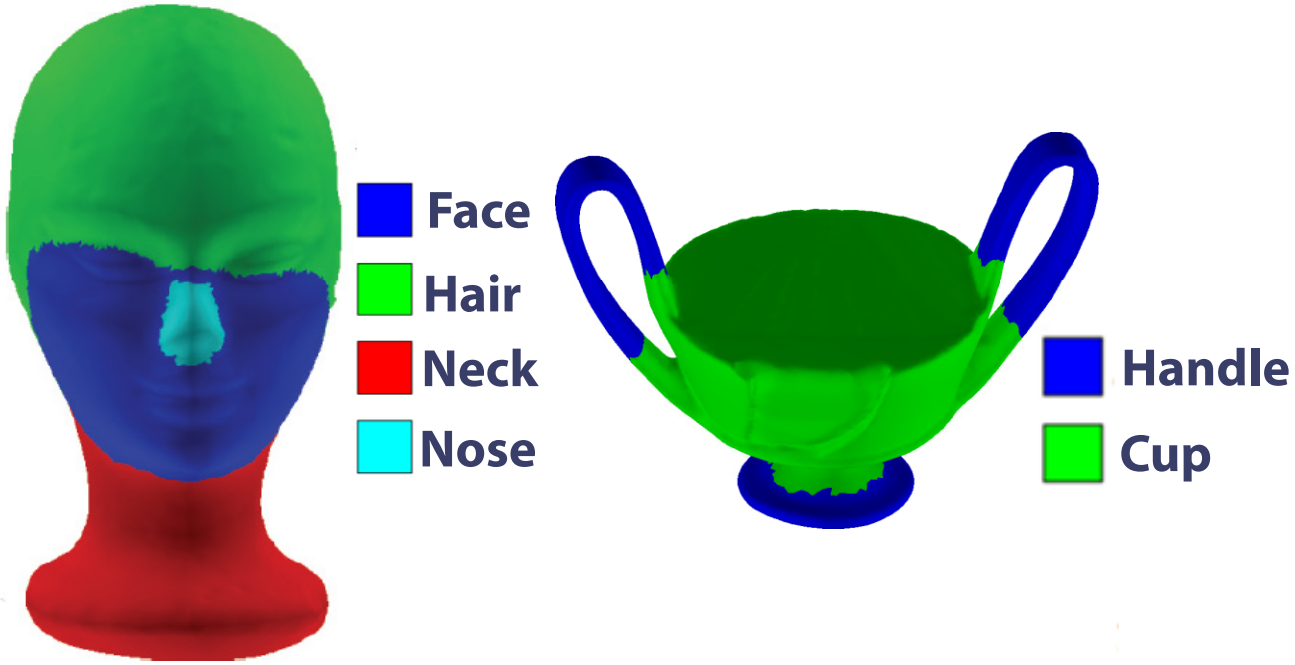

Failure cases

Limitations

Adjacent segments with the same label are merged

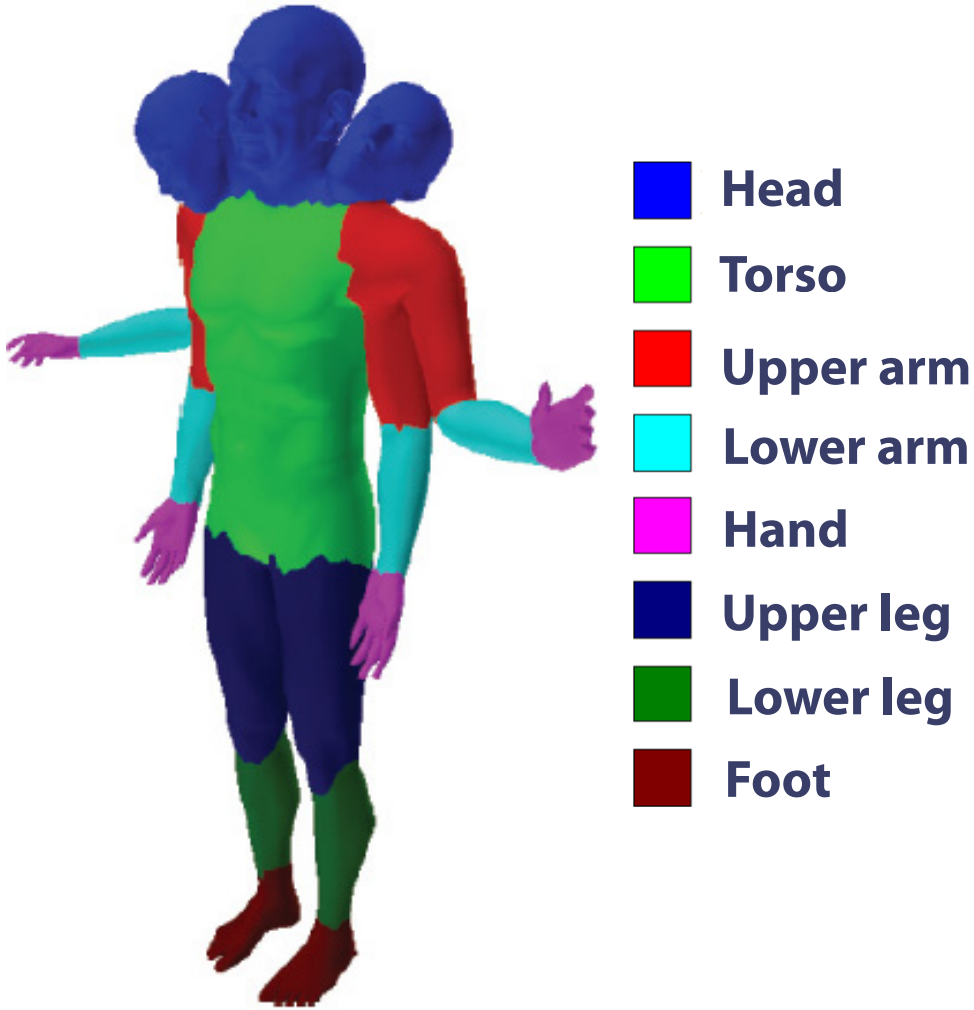

Limitations

Results depend on having sufficient training data

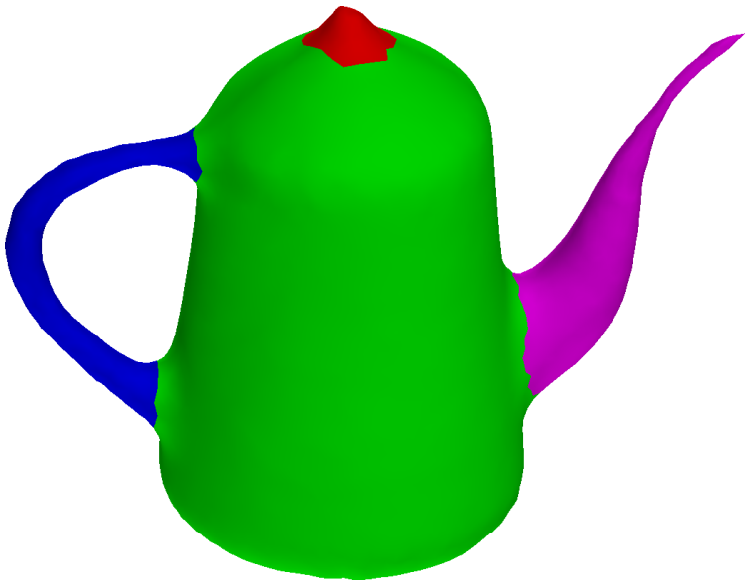

19 training meshes

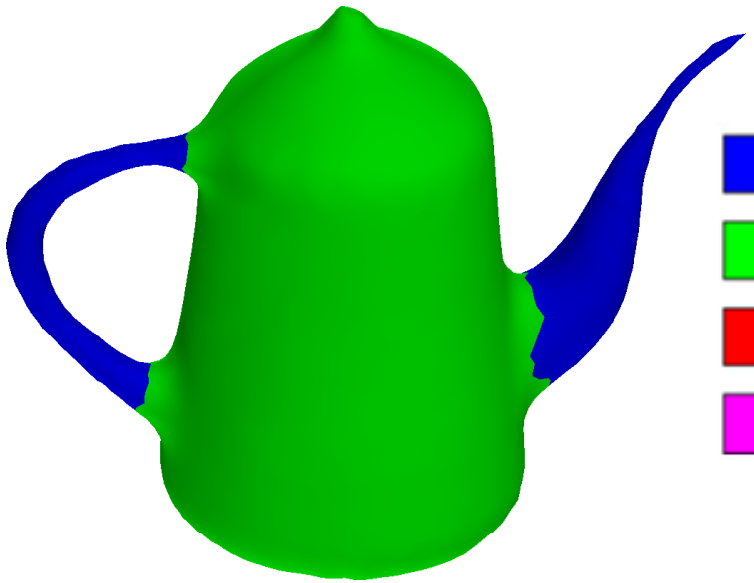

3 training meshes

Limitations

Many features are sensitive to topology

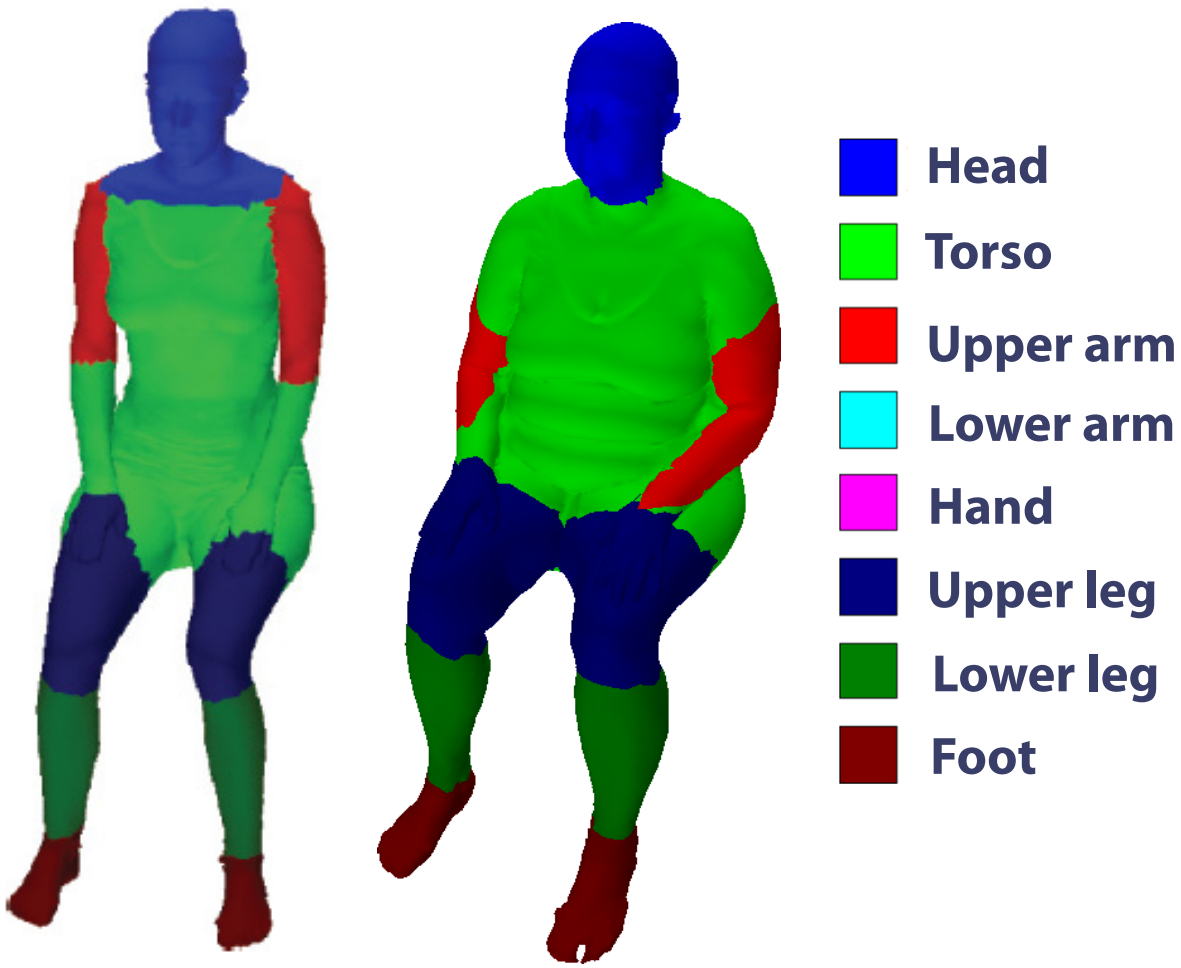

Applications: Character Texturing, Rigging

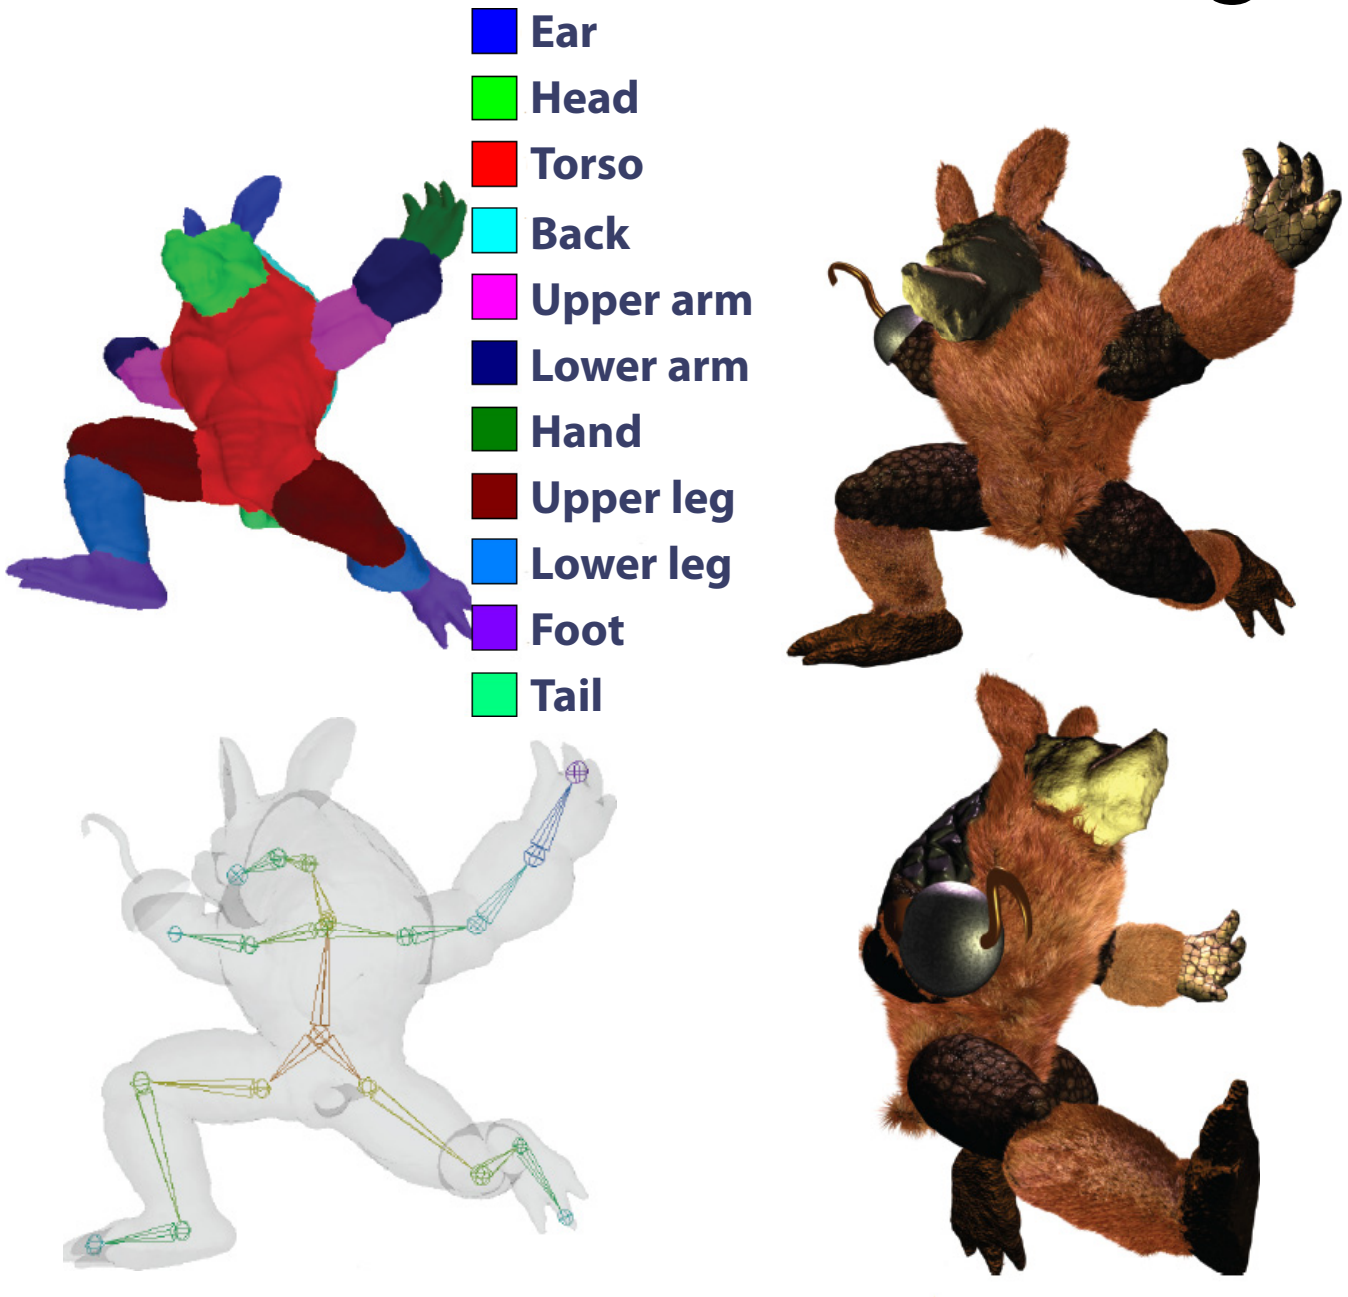

Summary

- Use prior knowledge for 3D mesh segmentation and labeling

Summary

- Use prior knowledge for 3D mesh segmentation and labeling
- Based on a Conditional Random Field model

Summary

- Use prior knowledge for 3D mesh segmentation and labeling
- Based on a Conditional Random Field model
- Parameters are learned from examples

Summary

- Use prior knowledge for 3D mesh segmentation and labeling
- Based on a Conditional Random Field model
- Parameters are learned from examples
- Applicable to a broad range of meshes

Summary

- Use prior knowledge for 3D mesh segmentation and labeling
- Based on a Conditional Random Field model
- Parameters are learned from examples
- Applicable to a broad range of meshes
- Significant improvements over the state-of-the-art

Thank you!

Acknowledgements: Xiaobai Chen, Aleksey Golovinskiy, Thomas Funkhouser, Szymon Rusinkiewicz , Olga Veksler, Daniela Giorgi, AIM@SHAPE, David Fleet, Olga Vesselova, John Hancock

Our project web page:

http://www.dgp.toronto.edu/~kalo/papers/LabelMeshes/

NSERC CRSNG

Canada Foundation for Innovation

Fondation canadienne pour l'innovation

Ontario

MINISTRY OF RESEARCH AND INNOVATION

MITACS