Neural Mesh Parameterization: DAWand and Beyond

Richard Liu[1], Vova Kim[2], Noam Aigerman[2], Rana Hanocka[1]

1 2 Adobe

About Me

PhD student @ UChicago 3DL Lab

Advisor: Rana Hanocka

Research interests: deep/classical geometry processing, 3D vision, 3D generative AI

I Like To Climb Things

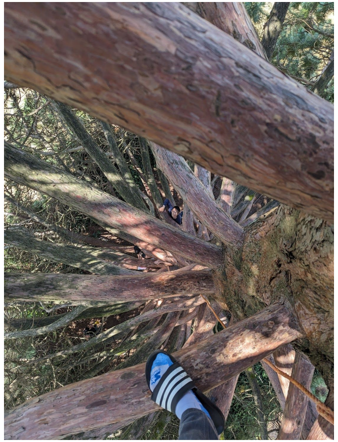

Mesh Parameterization Background

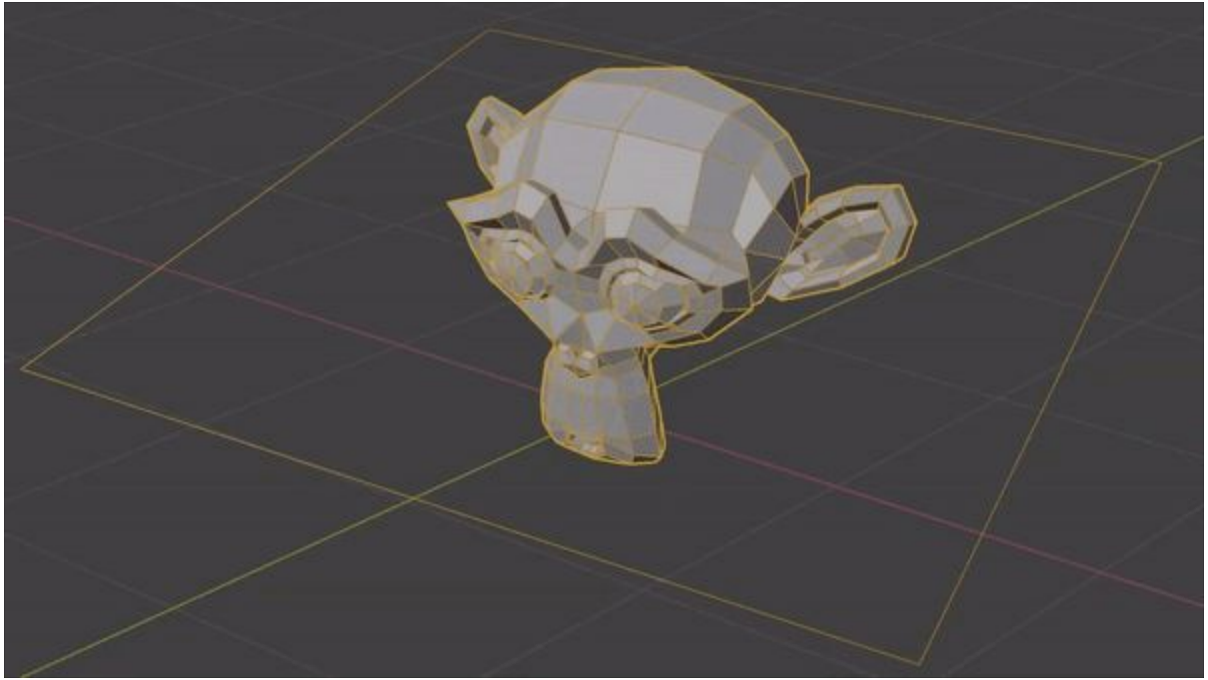

UVs are essential for 3D modeling

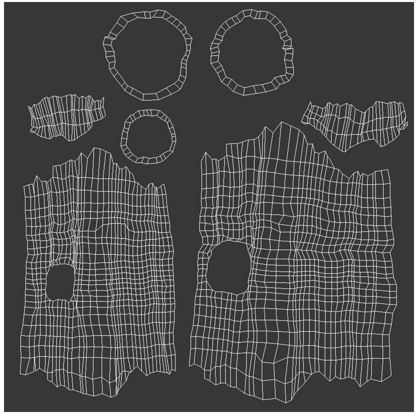

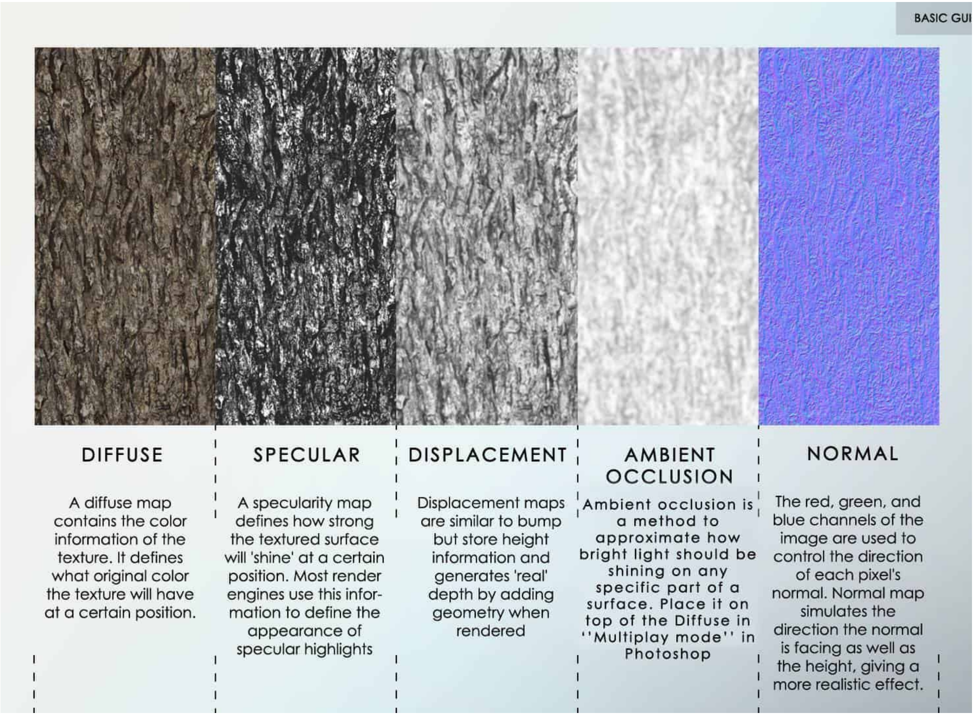

UVs are essential for 3D modeling

Remeshing

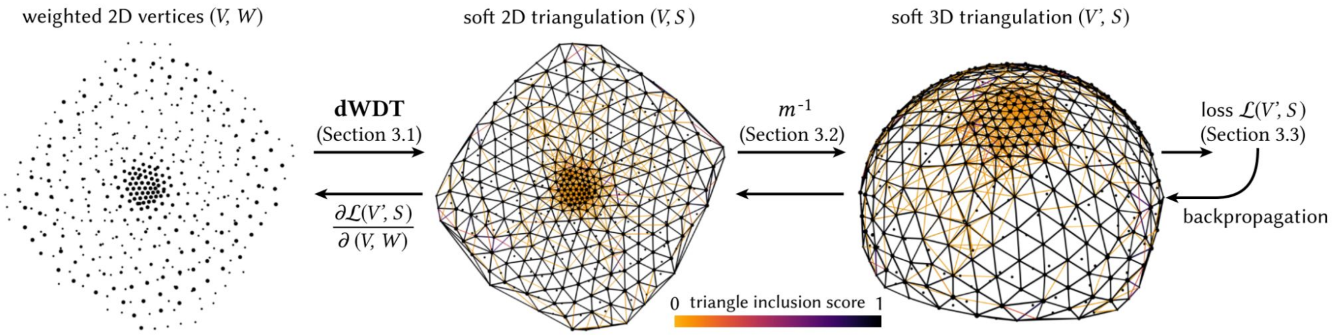

"Differentiable surface triangulation." Marie-Julie Rakotosaona, Noam Aigerman, Niloy J. Mitra, Maks Ovsjanikov, and Paul Guerrero. 2021.. ACM Trans. Graph. 40, 6, Article 267 (December 2021).

Simplification

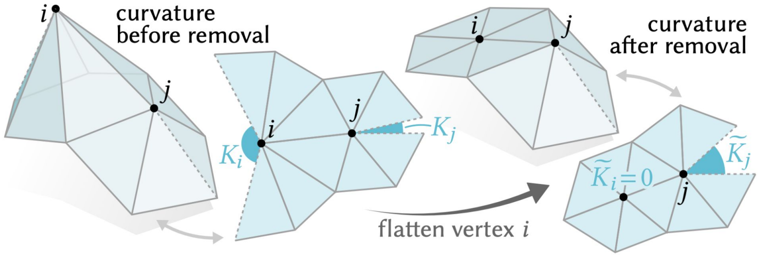

"Surface Simplification using Intrinsic Error Metrics." Hsueh-Ti Derek Liu, Mark Gillespie, Benjamin Chislett, Nicholas Sharp, Alec Jacobson, and Keenan Crane. 2023. ACM Trans. Graph. 42, 4, Article 118 (August 2023)

Global vs Local parameterization

Global Parameterization

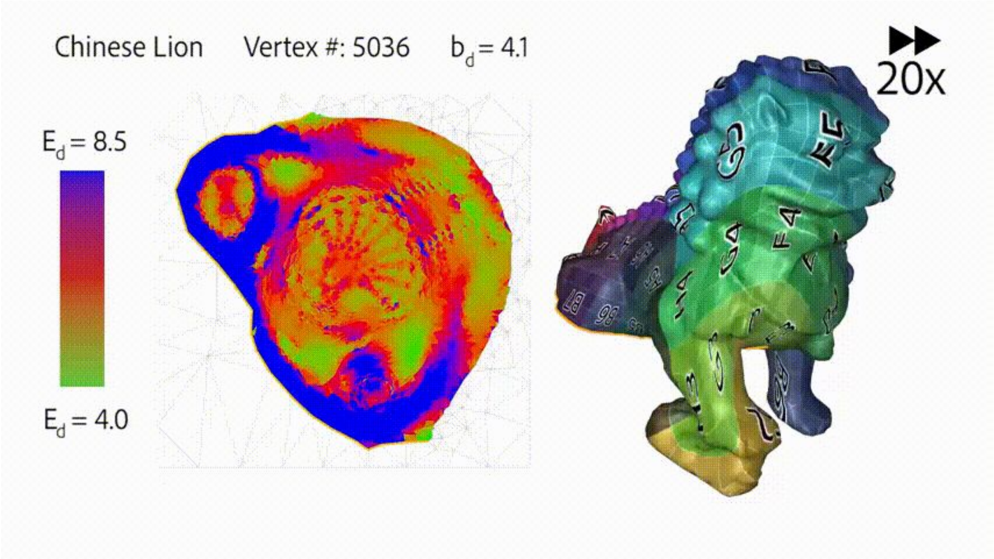

Li et al. 2018, “OptCuts: Joint Optimization of Surface Cuts and Parameterization” (SIGGRAPH 2018)

Local Parameterization

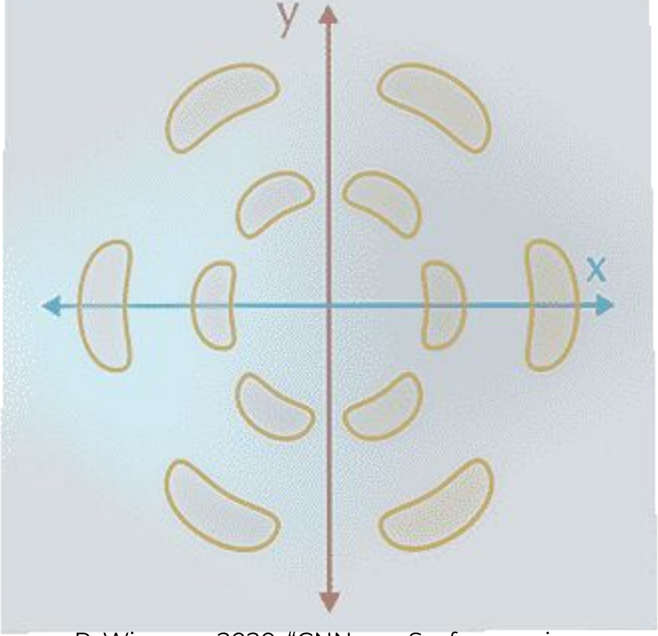

R. Wiersma 2020, “CNNs on Surfaces using Rotation-Equivariant Features” (SIGGRAPH 2018)

Global vs Local parameterization

Global Parameterization

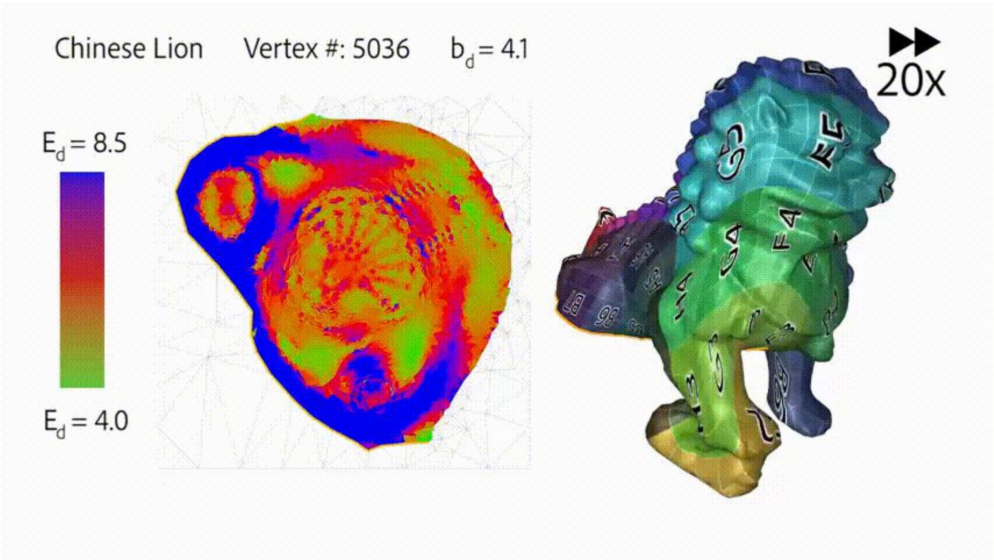

Li et al. 2018, “OptCuts: Joint Optimization of Surface Cuts and Parameterization” (SIGGRAPH 2018)

Local Parameterization

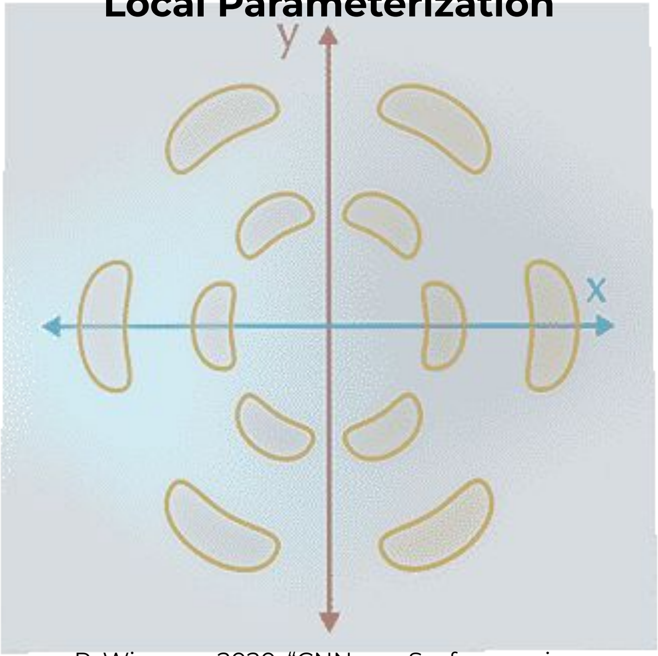

R. Wiersma 2020, “CNNs on Surfaces using Rotation-Equivariant Features” (SIGGRAPH 2018)

DA-Wand Motivation: Decaling

Segmentation for Parameterization Objectives

Maximize size of segmentation

Minimize parameterization distortion

TODO: REPLACE WITH ANIMATION OF GROWING DISTORTION WITH LOGMAP

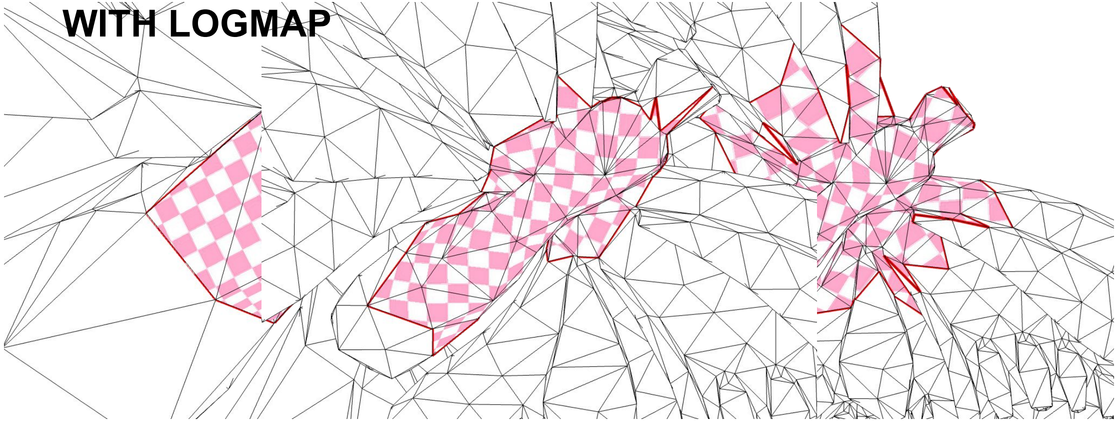

Parameterization distortion increases with segmentation size

Limitations of Existing Methods

LogMaps

Segmentation stops at high curvature boundaries

DCharts

Segmentation results in high distortion parameterization

DA-Wand

DA-Wand

Our method learns **distortion-aware** segmentations with a wide receptive field.

DA-Wand Overview

Technical challenges
- Lack of training data
- Energy crafting
- Initialization

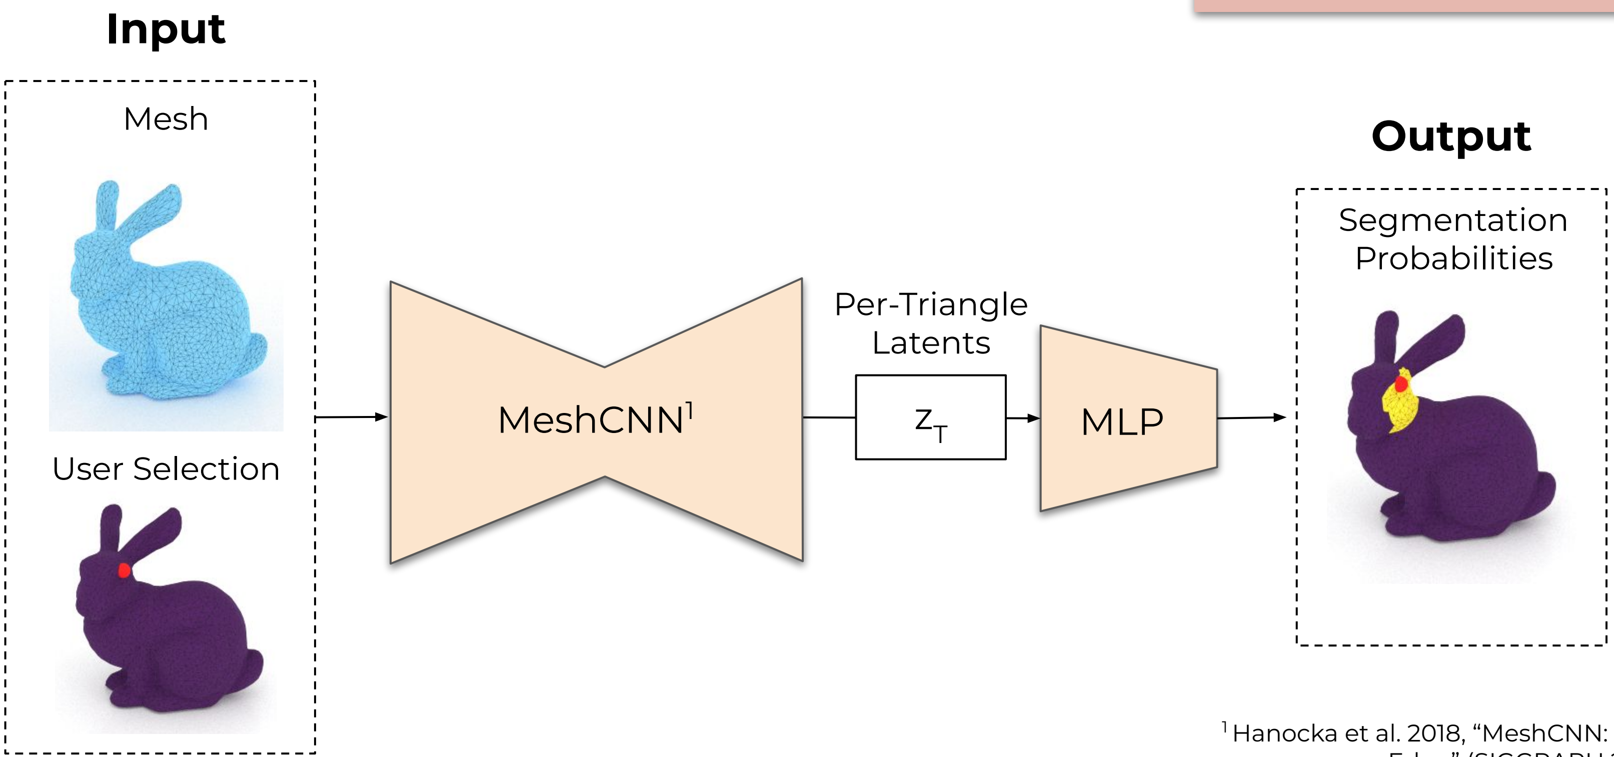

[1] Hanocka et al. 2018, “MeshCNN: A Network with an Edge” (SIGGRAPH 2019)

DA-Wand Overview

Technical challenges
- **Lack of training data**
- Energy crafting
- Initialization

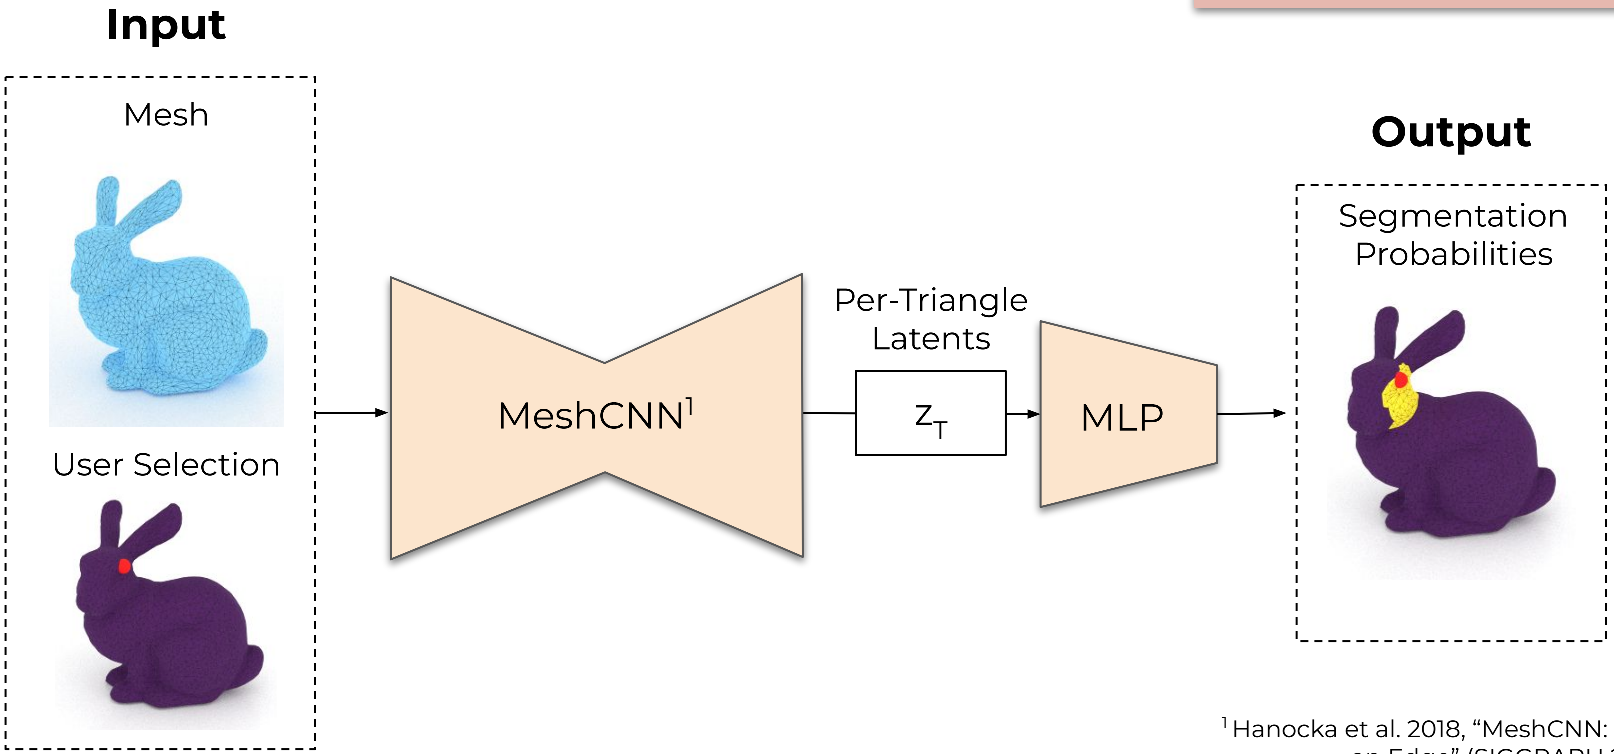

[1] Hanocka et al. 2018, "MeshCNN: A Network with an Edge" (SIGGRAPH 2019)

DA-Wand Overview

Technical challenges

- **Lack of training data**
- Energy crafting
- Initialization

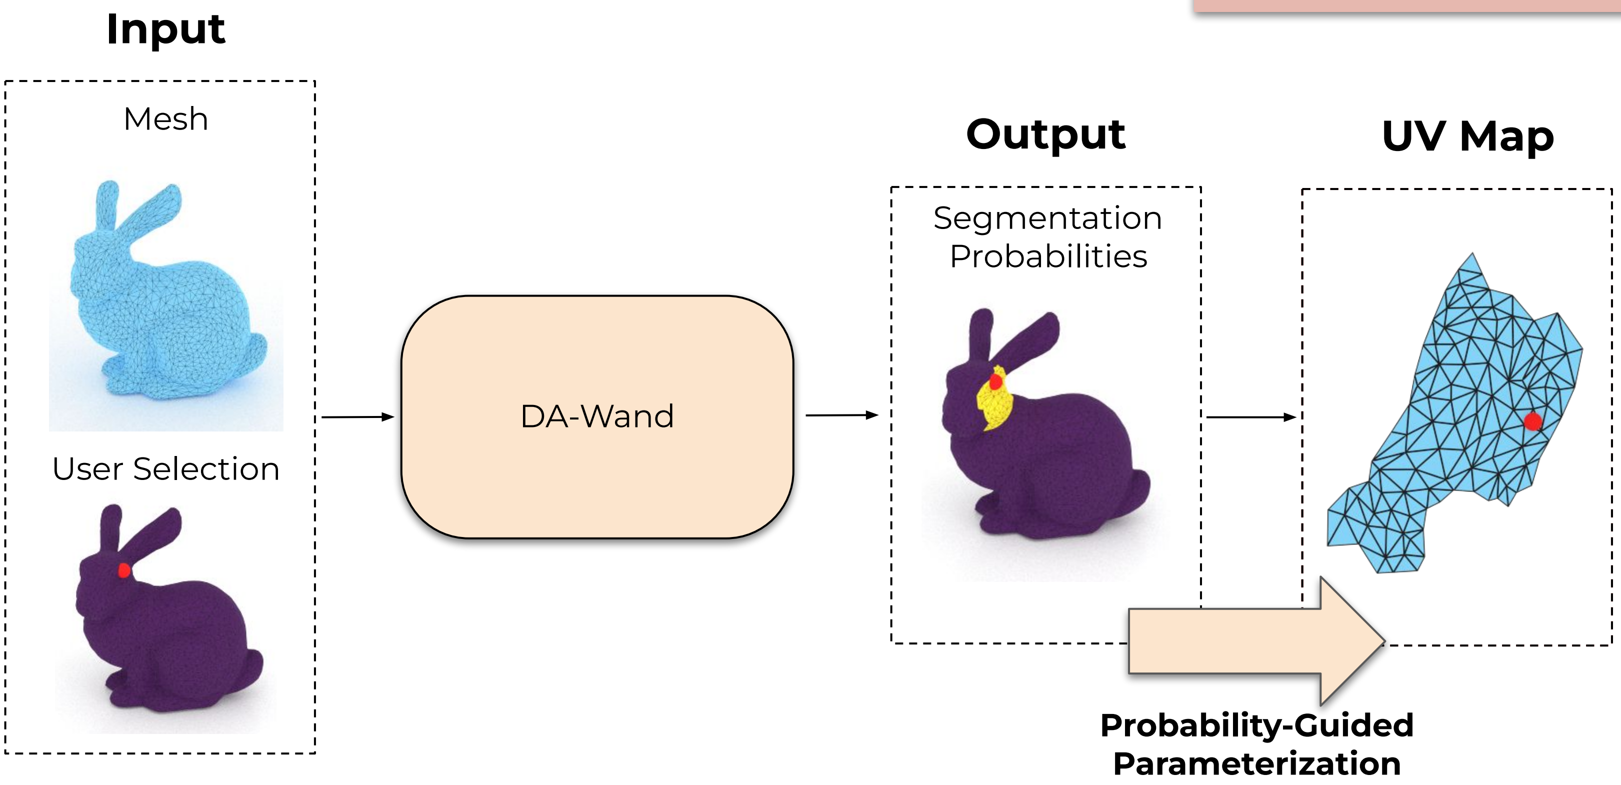

Probability Guided Parameterizatio

Technical challenges
- **Lack of training data**
- Energy crafting
- Initialization

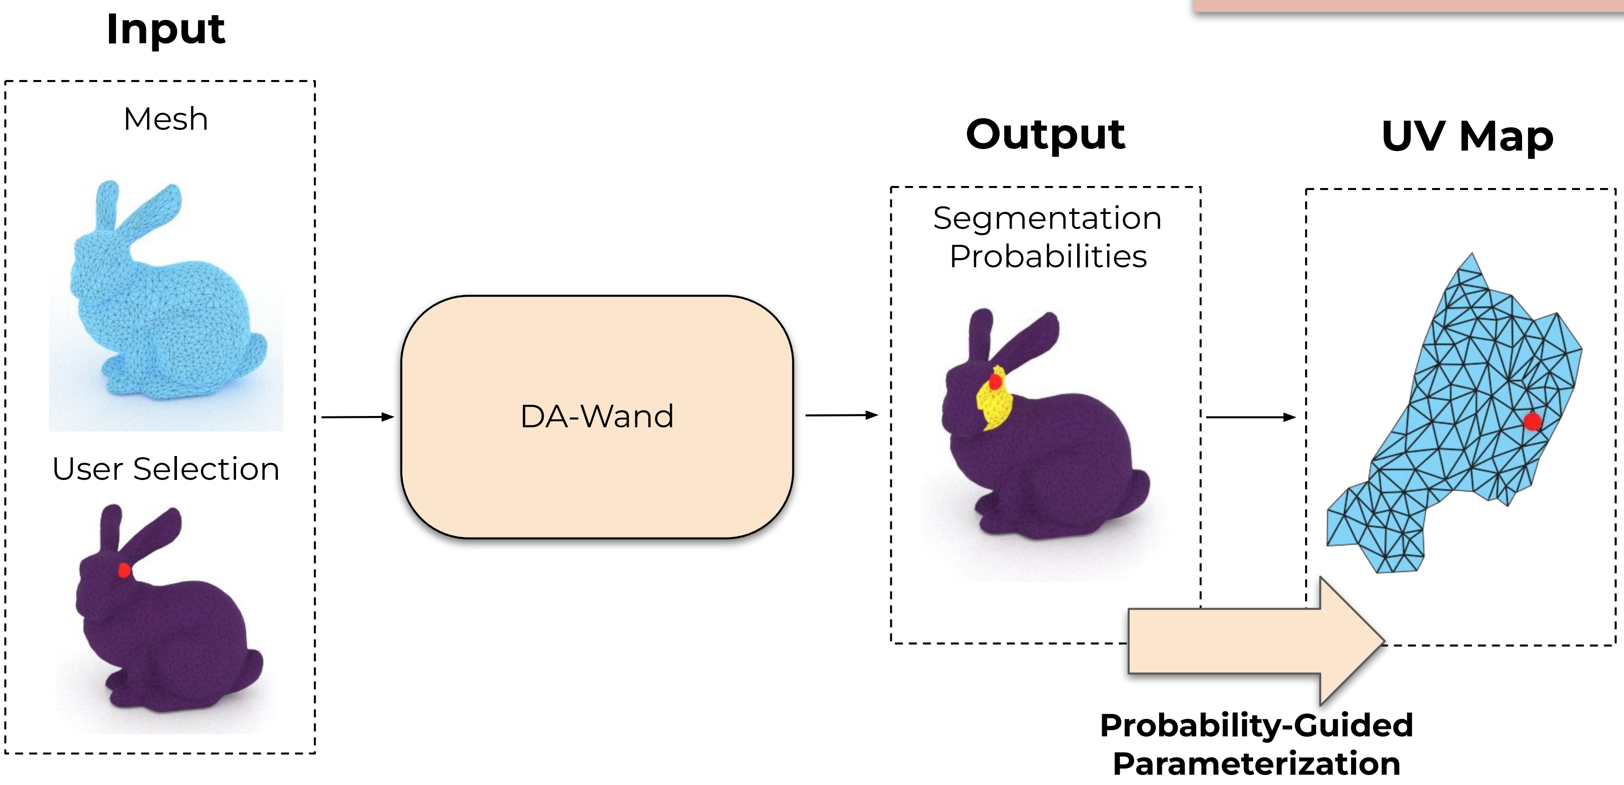

Probability-Guided Parameterization

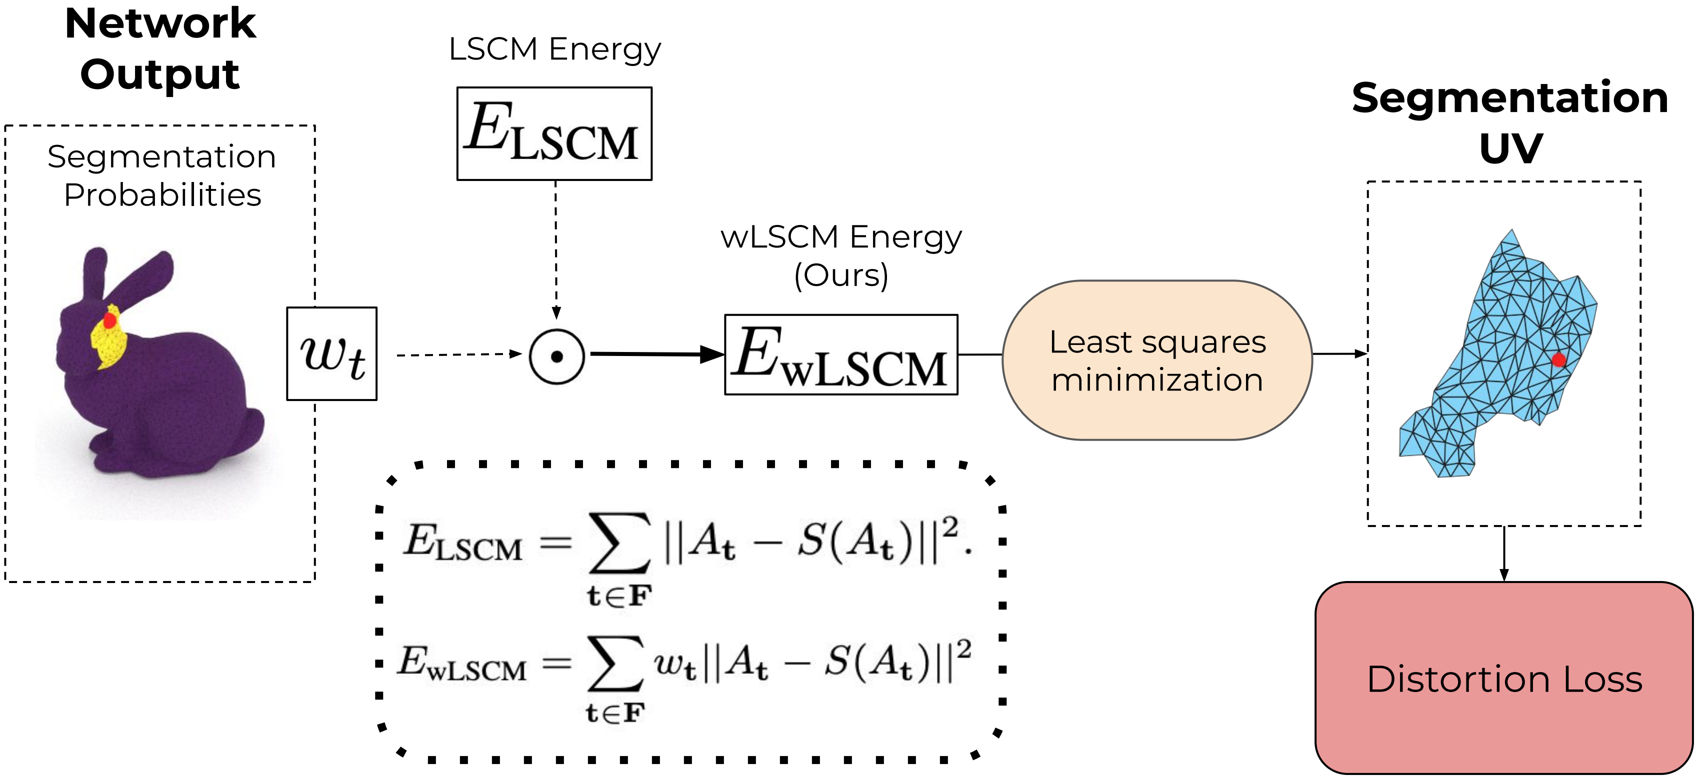

$$E_{\text{LSCM}} = \sum_{\mathbf{t} \in \mathbf{F}} ||A_{\mathbf{t}} - S(A_{\mathbf{t}})||^2.$$

$$E_{\text{wLSCM}} = \sum_{\mathbf{t} \in \mathbf{F}} w_{\mathbf{t}} ||A_{\mathbf{t}} - S(A_{\mathbf{t}})||^2$$

Probability-Guided Parameterization

$$E_{\mathrm{LSCM}} = \sum_{\mathbf{t}\in\mathbf{F}} ||A_{\mathbf{t}} - S(A_{\mathbf{t}})||^2.$$

$$E_{\mathrm{LSCM}} = \sum_{\mathbf{t}\in\mathbf{F}} ||A_{\mathbf{t}} - S(A_{\mathbf{t}})||^2.$$

$$E_{\mathrm{wLSCM}} = \sum_{\mathbf{t}\in\mathbf{F}} w_{\mathbf{t}}||A_{\mathbf{t}} - S(A_{\mathbf{t}})||^2$$

TODO: Explain all the parameters in the energy. Replace A_t with J_t. Copy formula for affine transformation from paper. Add R^2x2 dimensions to clarify size of matrices.

Animate parameter by parameter.

wLSCM Equivalence to LSCM

Theorem: As wLSCM weights converge to a binary segmentation, the wLSCM solution for triangles with nonzero weights converges to the LSCM solution for those triangles.

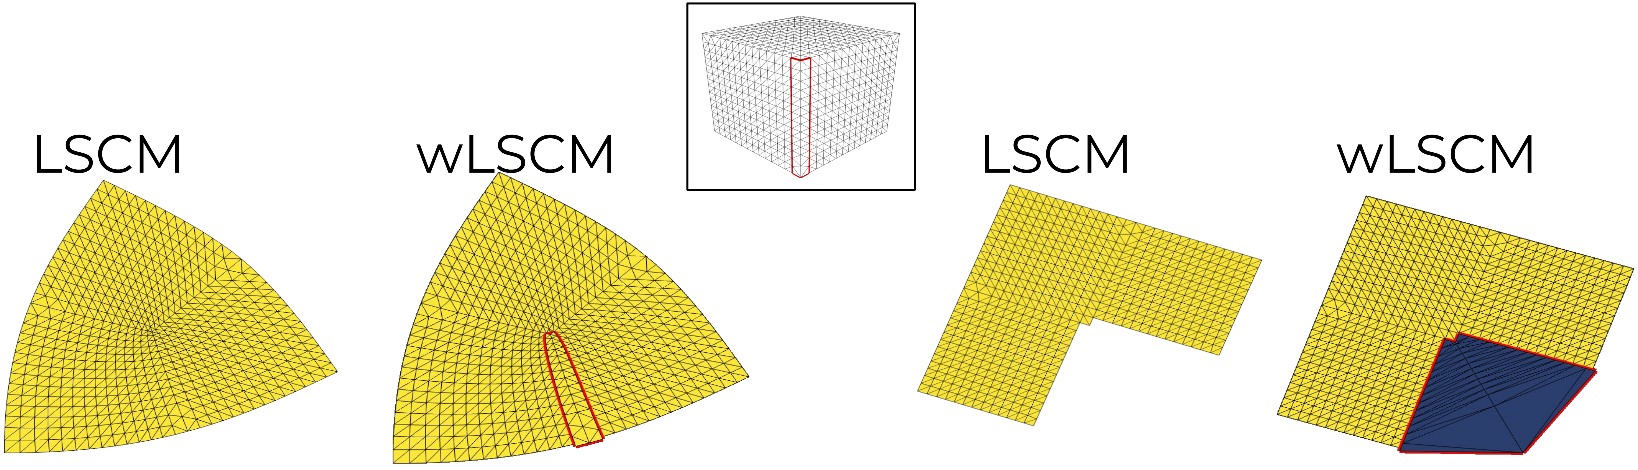

wLSCM Equivalence to LSCM

Theorem: As wLSCM weights converge to a binary segmentation, the wLSCM solution for triangles with nonzero weights converges to the LSCM solution for those triangles.

Weights: 1.0000

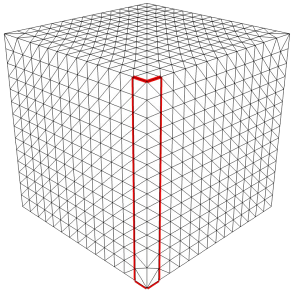

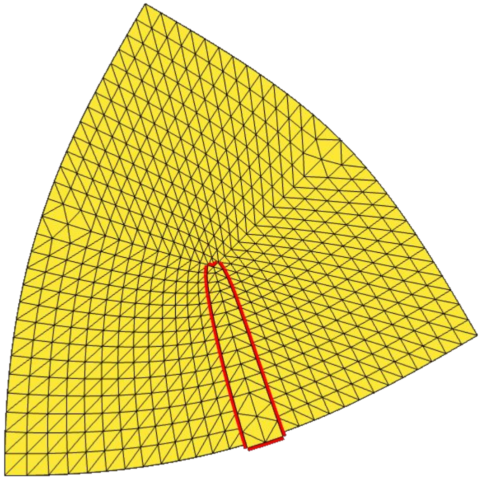

Recap: Differentiable Parameterization

TODO: visual that shows global parameterization with the segmented region mapped

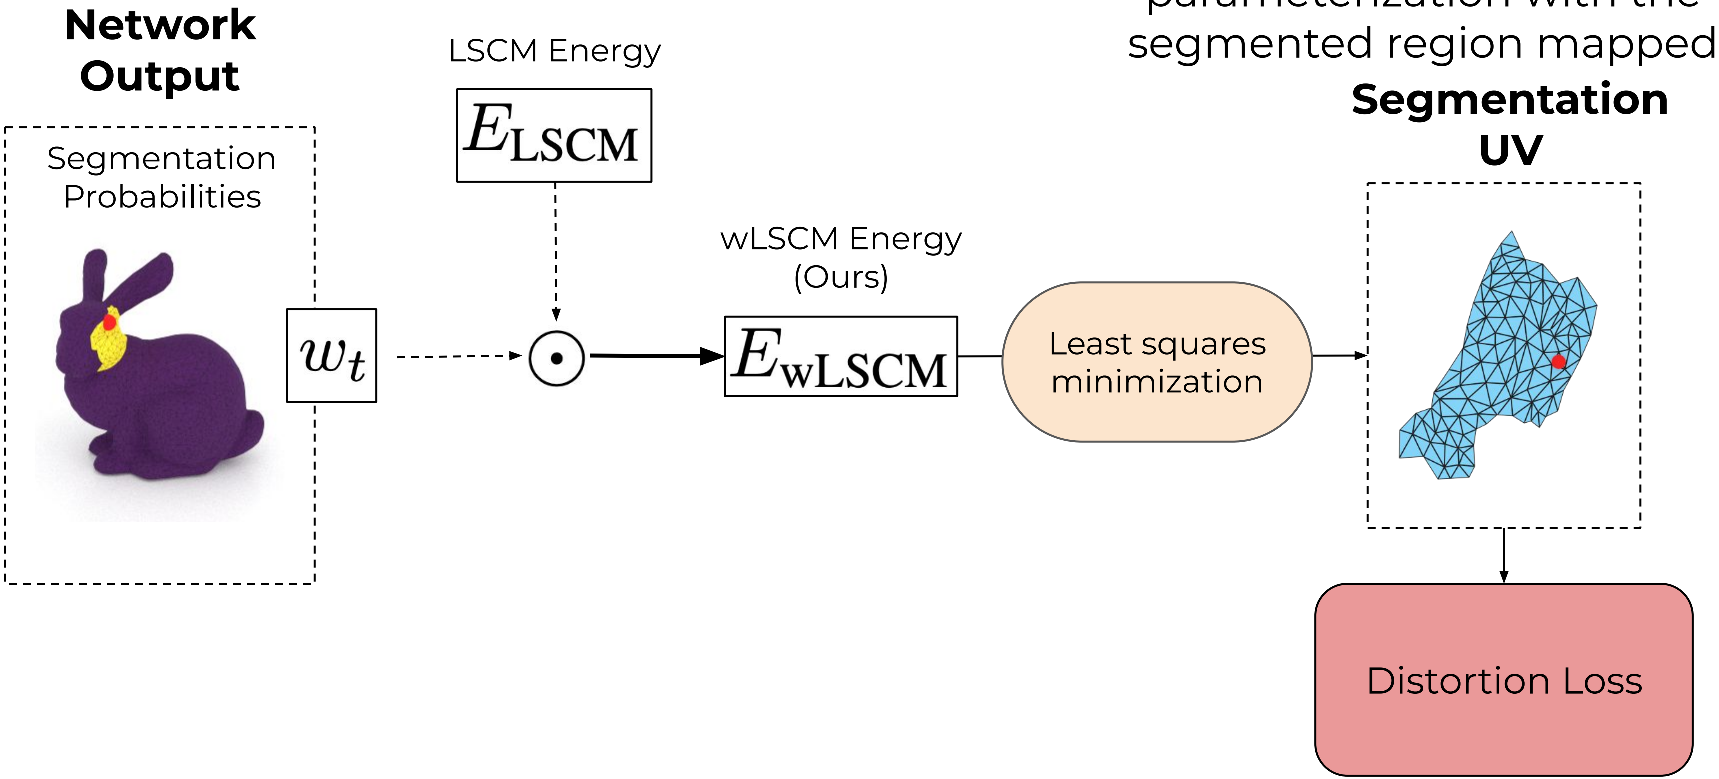

Local Minima Solutions to Segmentation Problem

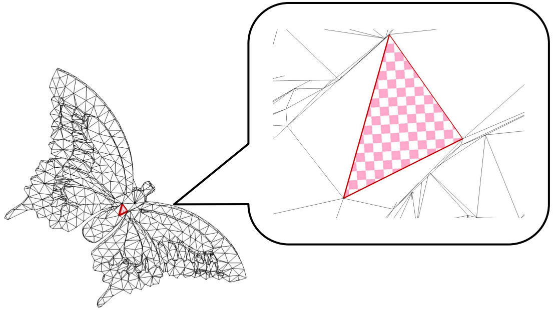

single triangle: no distortion

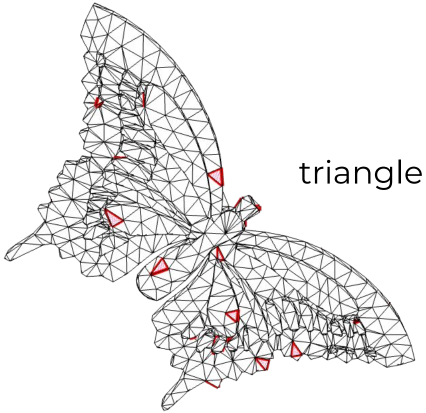

triangle soup: no distortion

DA-Wand Technical Roadmap

Technical Challenges

- Lack of training data
- **Energy formulation**
- Regularization
- Random initialization gives non-meaningful parameterizations => bad gradients

Size and Compactness Priors

Thresholded-distortion loss: maximize # triangles underneath a given distortion threshold (TODO: thresholded-distortion loss => show parameter that controls distortion-awareness => ARAP energy for isometric optimization)

$$L^{\text{threshold}}(D) = \sum_t \frac{A_t}{\sum A_t} \left(1 - e^{-(D_t/\gamma)^\alpha}\right)$$

Smoothness loss: coplanar triangles should have same segmentation weight (TODO: show graphcuts => focuses on sharp edges/geometry aware)

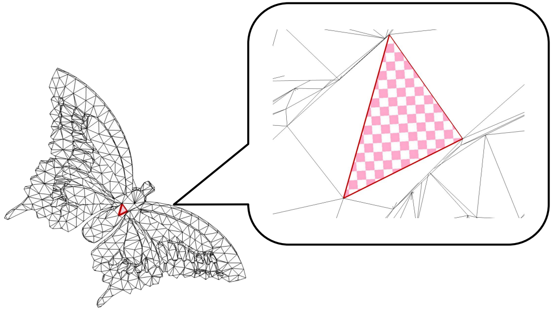

single triangle: no distortion

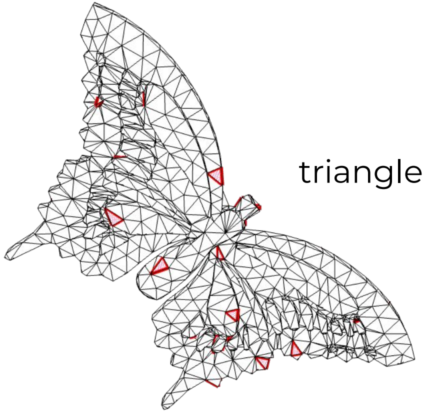

triangle soup: no distortion

Distortion Self-Supervision Pipeline

TODO: recap the full pipeline + how to generate the self-supervised training data (sample points on natural shapes)

Q. Zhou & A. Jacobson. 2016, “Thingi10K: A Dataset of 10,000 3D-Printing Models”

DA-Wand Technical Roadmap

Technical Challenges

- Lack of training data
- Energy formulation
- **Random initialization gives non-meaningful parameterizations => bad gradients**

Synthetic Supervision

Synthetic Dataset

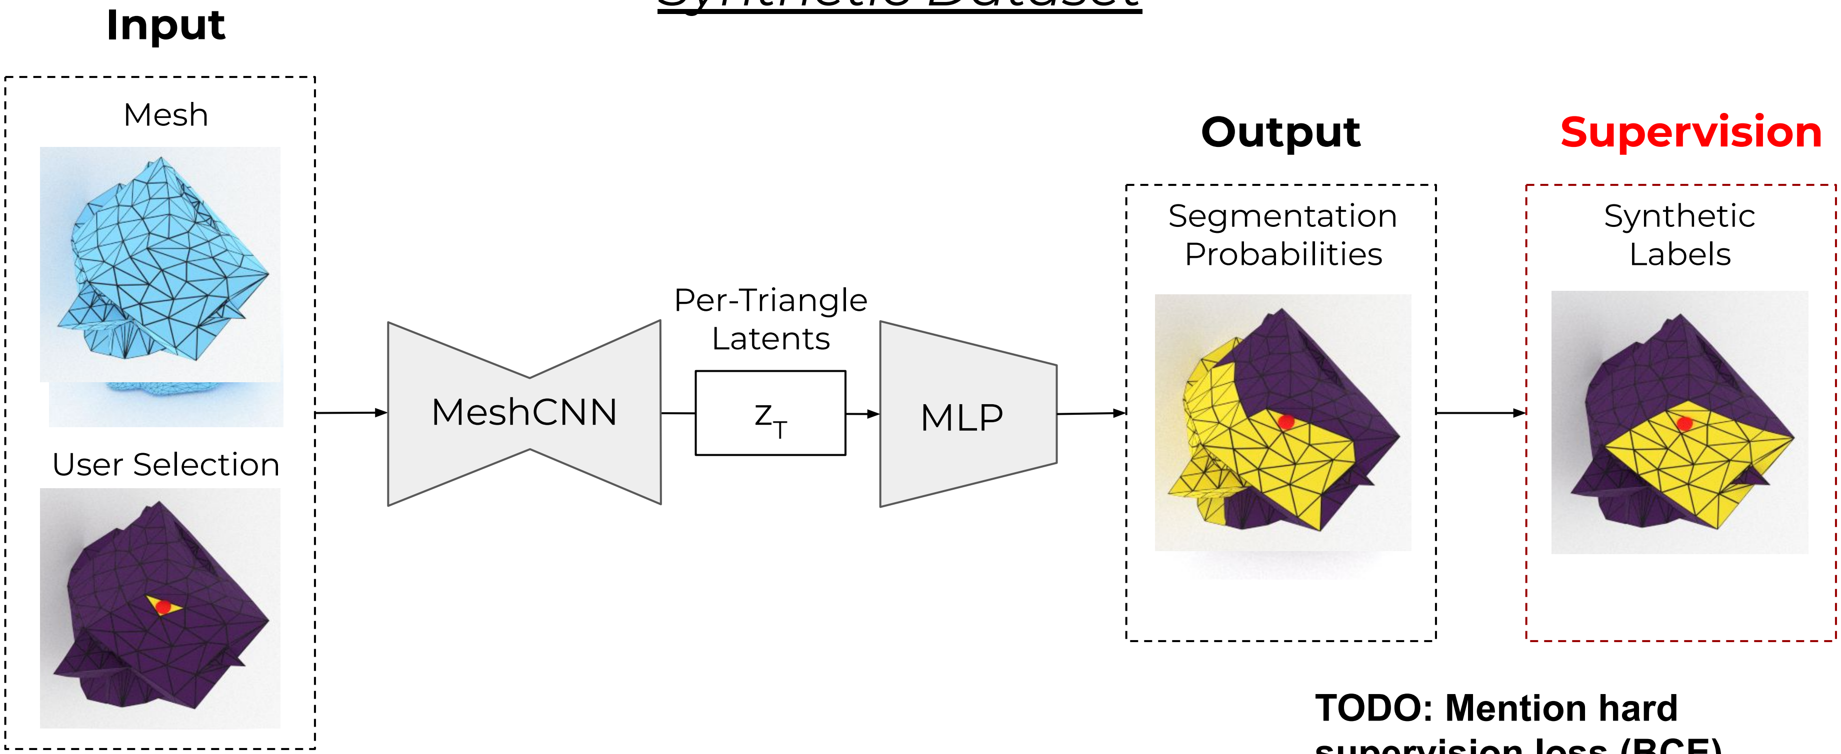

TODO: Mention hard supervision loss (BCE)

Synthetic Dataset Generation

Synthetic Segmentation Dataset

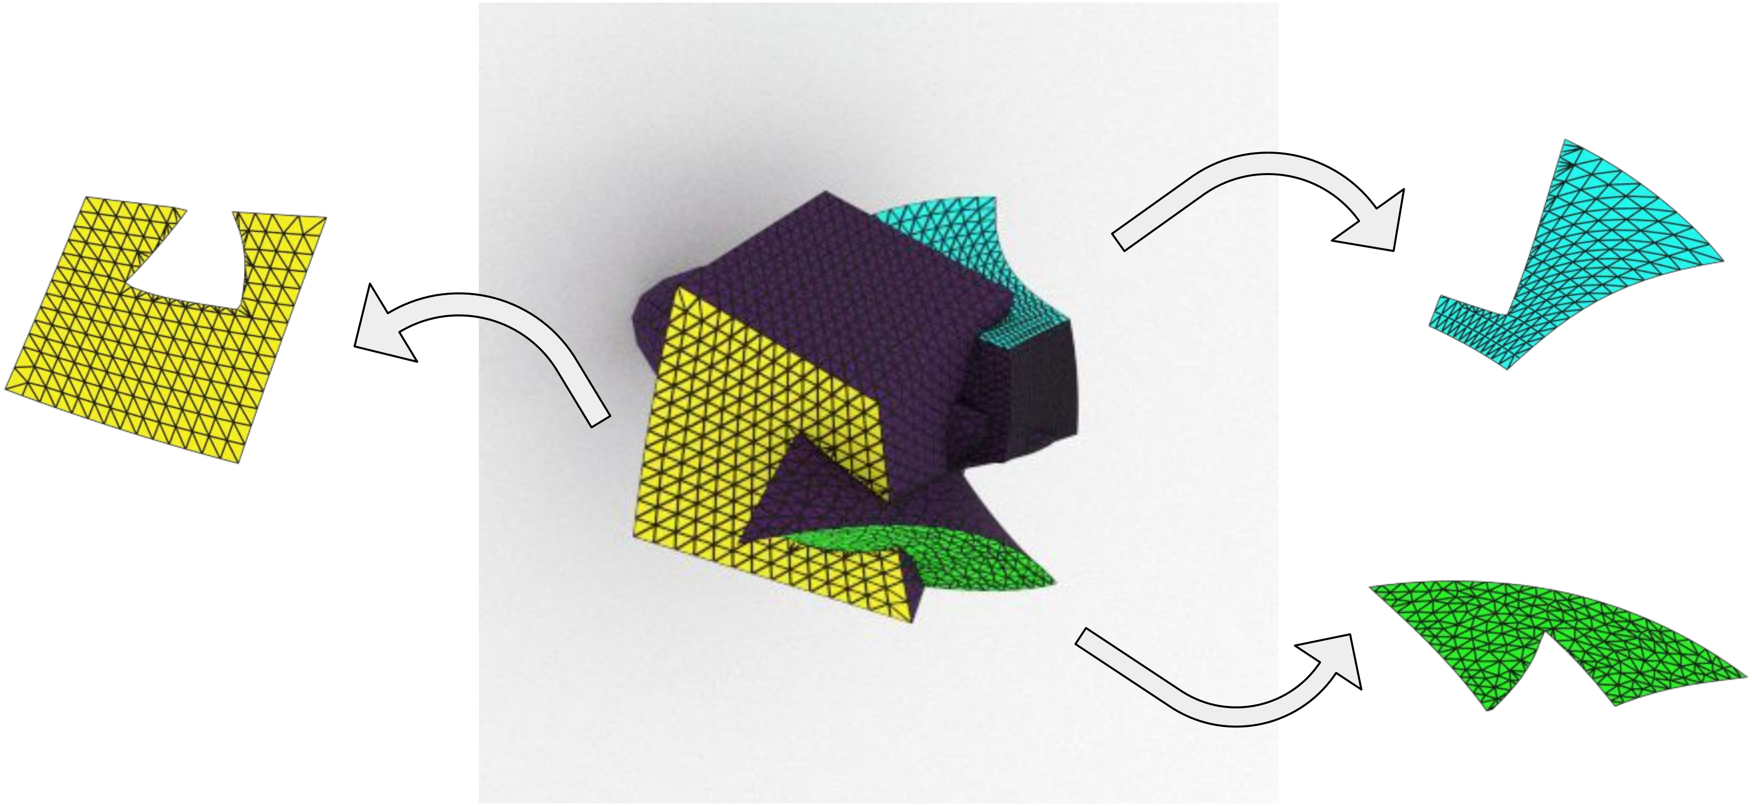

DA-Wand Pipeline Recap

Phase 1

TODO

Phase 2

TODO

Segmentation Results

Thingi10k

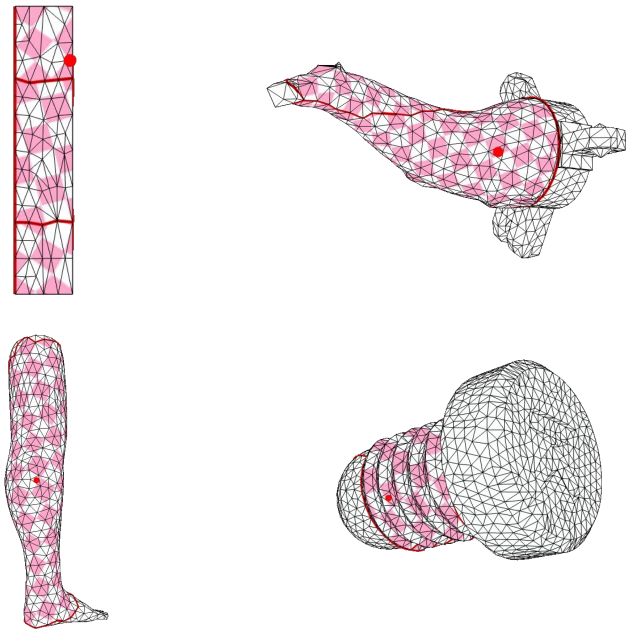

Parameterization Benchmark

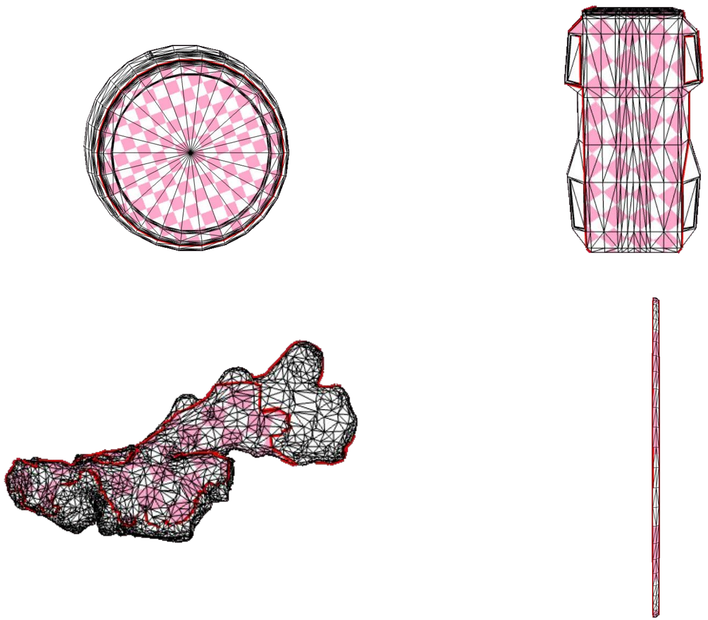

Interactive Application

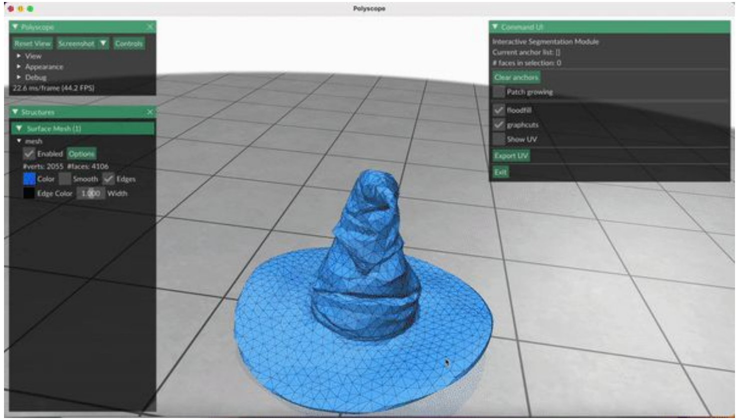

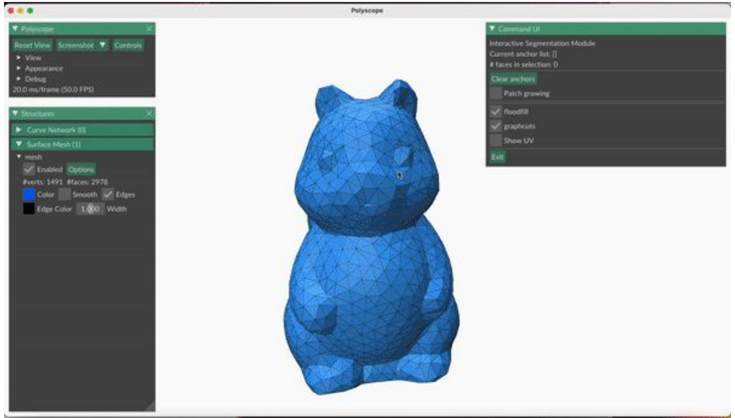

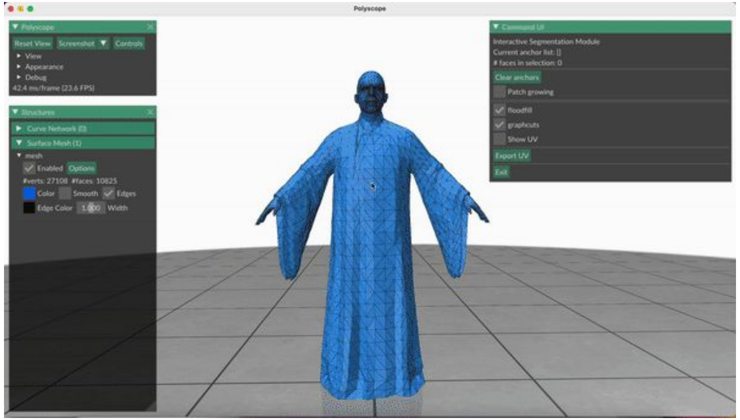

No GPU required

Summary

Motivation:

- Differentiable parameterization as a function of segmentation weights
- Unsupervised training
- Regularizing energies to enforce compactness + distortion bounds
- Data-driven to get distortion-aware segmentations
- Fully convolutional arch allows for interactive segmentation on medium-sized meshes

DA-Wand learns to solve (local) segmentation problem

Beyond DA-Wand: Moving to Global Parameterization

Mesh parameterization involves two entangled combinatorial problems.

- Segmentation
  - Minimize cuts i.e. discontinuities

Influence each other

- Flattening
  - Minimize distortion

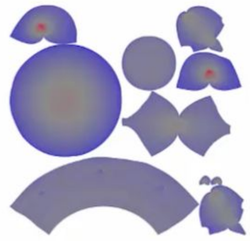

Data-driven solution for global mesh parameterization?

Limitations of Existing Methods

- Little to no global surface analysis
- Optimization-based: sensitive to initialization
- Tend to get stuck in local minima

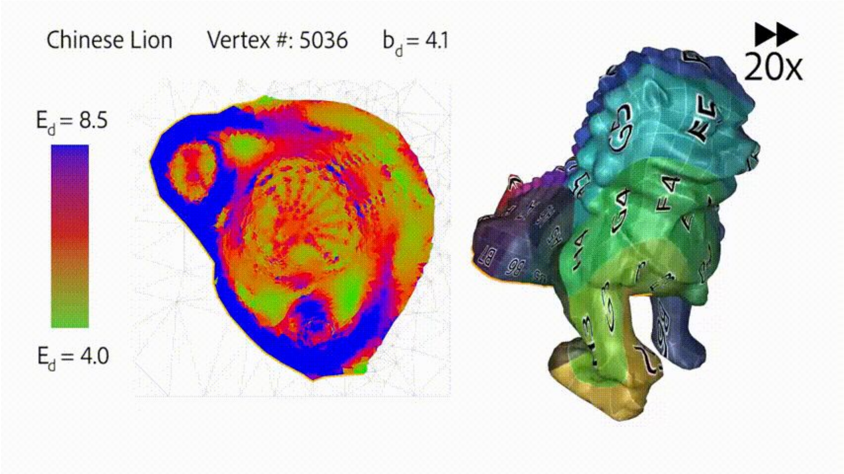

Li et al. 2018, "OptCuts: Joint Optimization of Surface Cuts and Parameterization" (SIGGRAPH 2018)

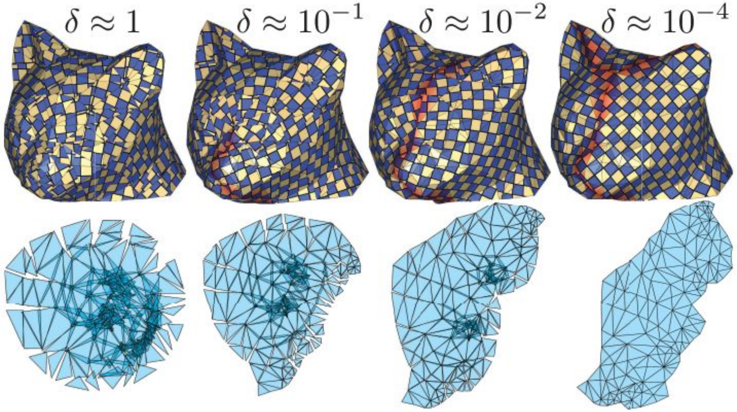

Poranne et al. 2017, "Autocuts: Simultaneous Distortion and Cut Optimization for UV Mapping" (SIGGRAPH 2017)

NeuralUV: Key Idea

DA-Wand: injects segmentation weights into the LSCM linear system to produce a “soft” LSCM system (UV map is smooth function of segmentation prediction)

$$E_{\mathrm{wLSCM}} = \sum_{\mathbf{t} \in \mathbf{F}} w_{\mathbf{t}} ||A_{\mathbf{t}} - S(A_{\mathbf{t}})||^2$$

NeuralUV: injects edge cutting weights into the Poisson solve to product “soft” Poisson system. (UV map is smooth function of **edge weights** and **predicted Jacobians**).

$$\Phi^* = (\mathcal{L} + \underbrace{\mathcal{W}}_{\textit{Edge weights}})^{-1} A \nabla^T \underbrace{J}_{\textit{Predicted Jacobians}}$$

Soft Poisson

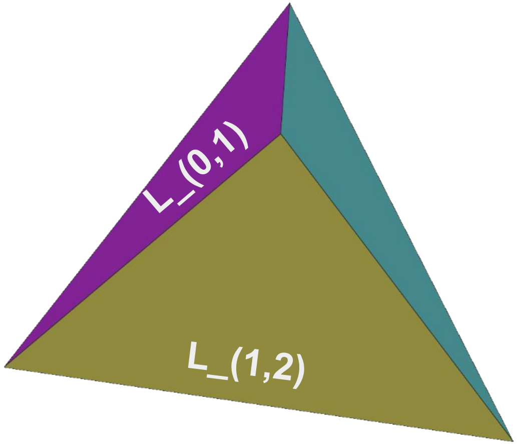

Soft Poisson

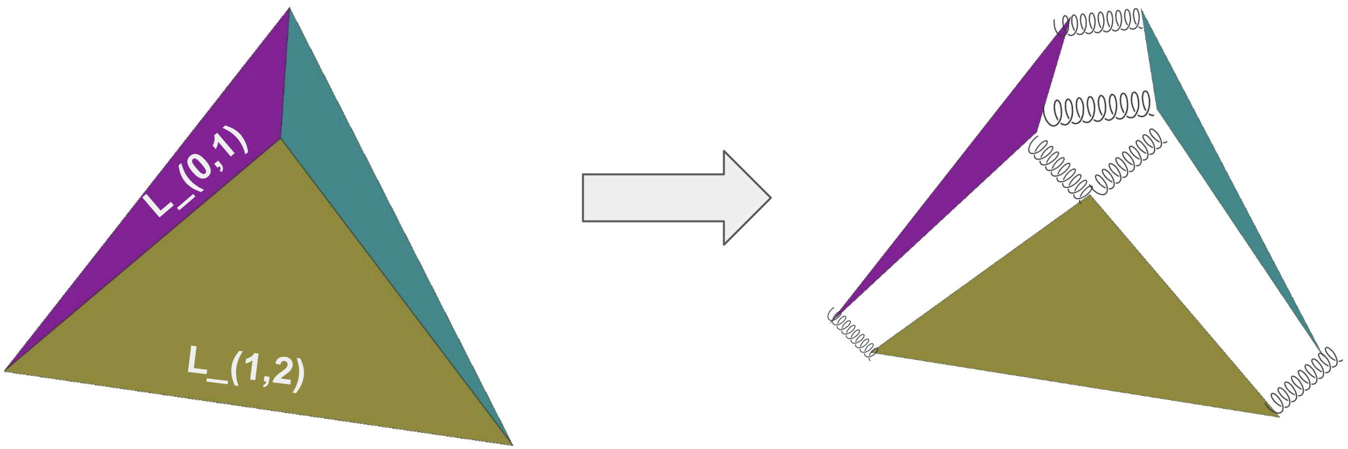

Soft Poisson

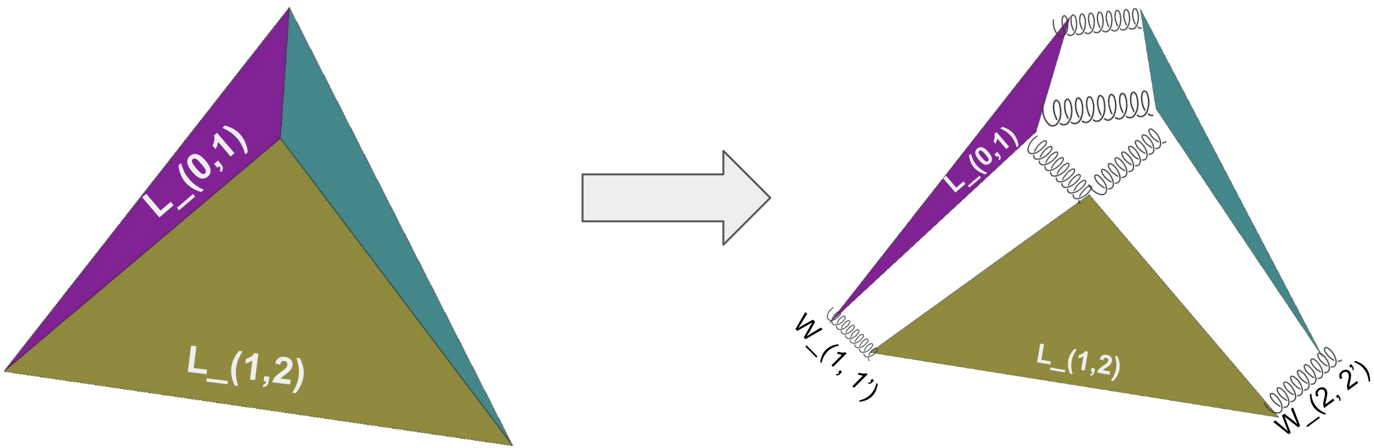

Thank you!

DA-Wand Objective Recap

Input

Mesh

User Selection

Output

Segmentation Probabilities

Output: binary segmentation probabilities

Desired Segmentation Properties

- Large
- Compact
- Distortion-aware
- User-input conditioned

Technical Challenges

- Lack of training data
- Energy crafting
- Regularization
- Random initialization gives non-meaningful parameterizations => bad gradients