Neural Mesh Editing

Rana Hanocka

Assistant Professor

University of Chicago

Neural Representations of Geometry Show Promise

NeRF [ECCV 2020]

Gaussian Splatting [SIGGRAPH 2023]

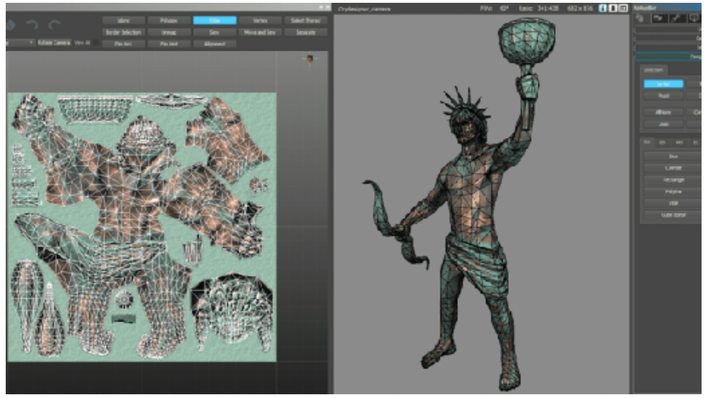

Meshes are the industry standard

Entire graphics pipeline is built around meshes!

3D content creation by editing meshes!

(as opposed to NeRFs) Because:

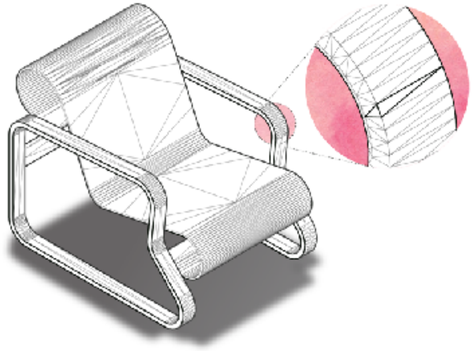

Meshes accurately represent sharp features topology

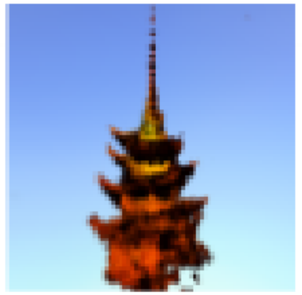

Mesh

Extraction

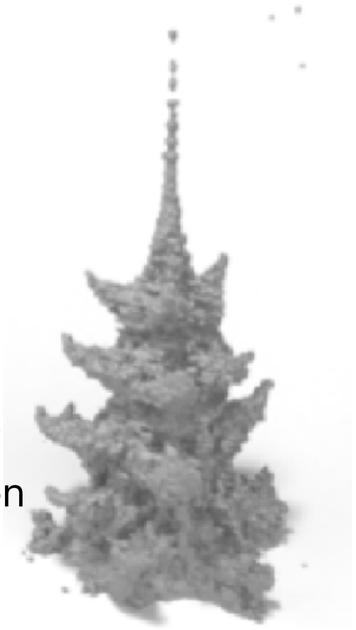

Need to convert to a mesh

NeRF → Mesh isn't easy

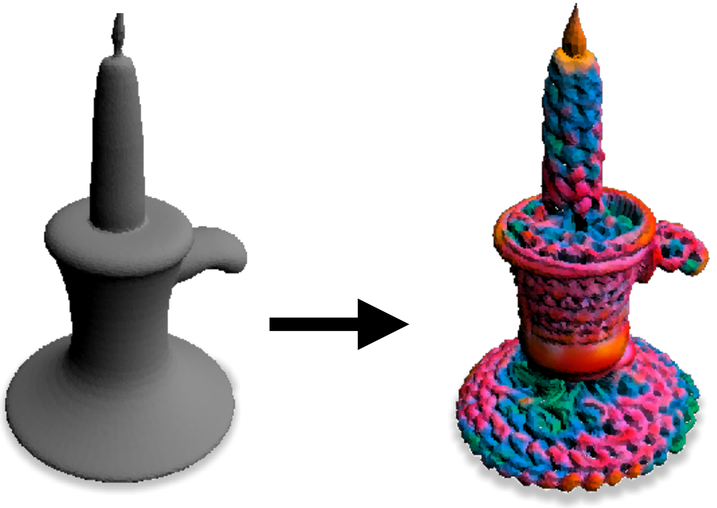

Text2Mesh [CVPR 2022]

Provides user/artist control

Where do we get the 3D datasets from?

Deep learning models are data-hungry!

Supervised or unsupervised datasets

Large unsupervised datasets still requires curation!

Difficult to obtain big 3D training datasets

Difficult to create 3D Data

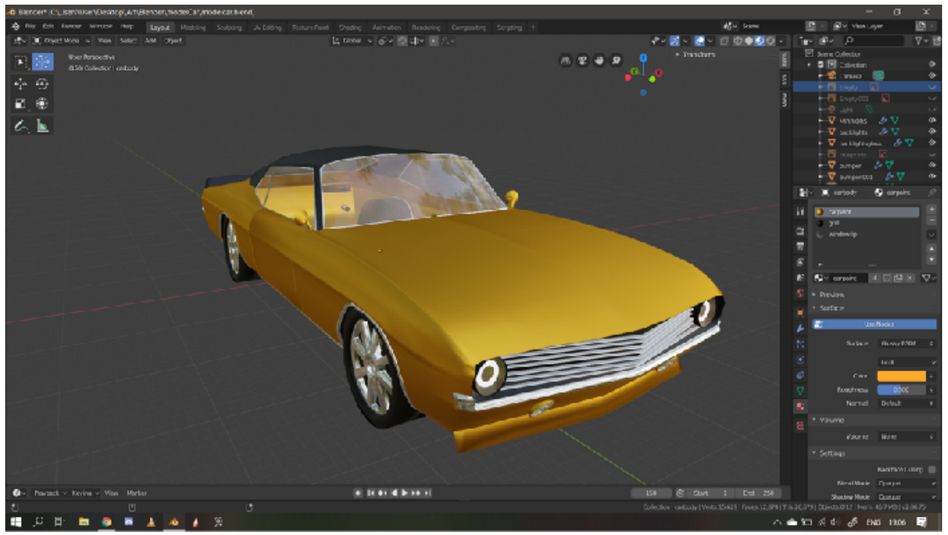

Blender

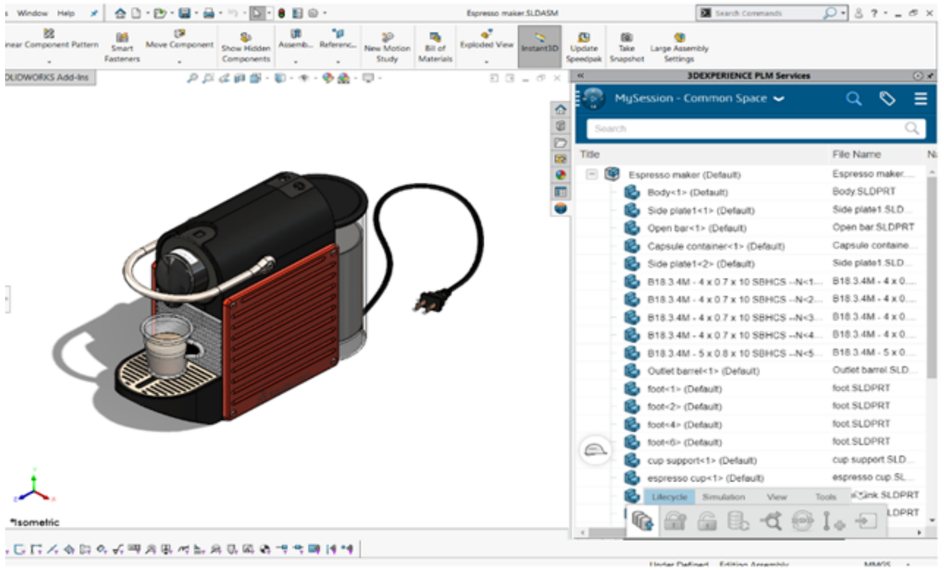

SolidWorks

High visual quality ≠ compatible for geometric computation

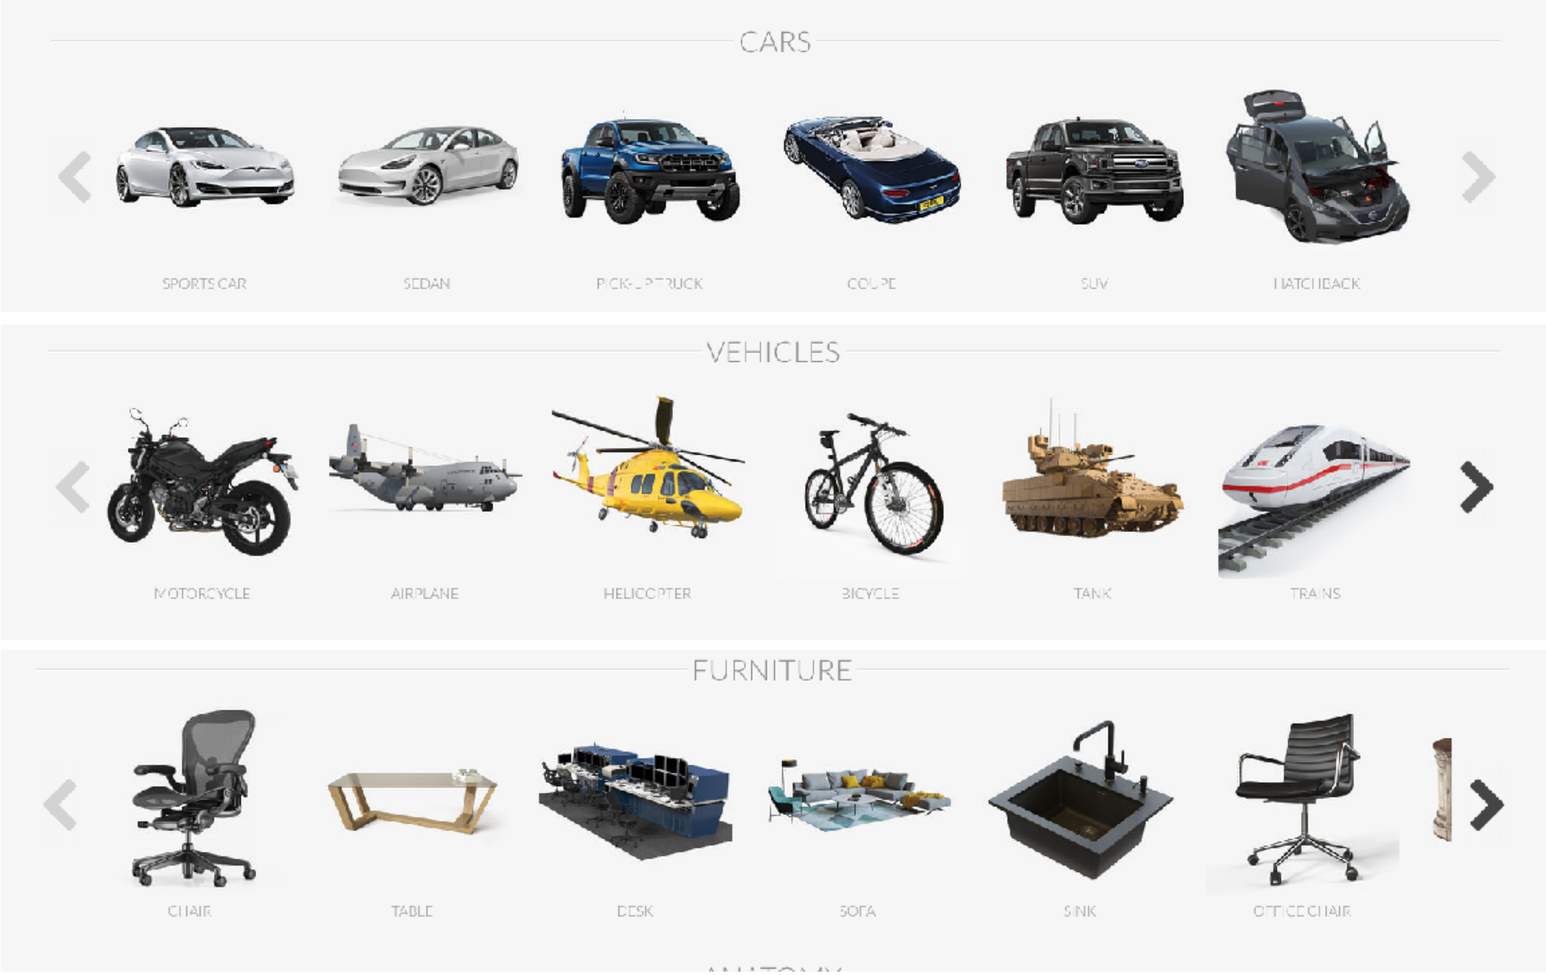

Big 3D data challenges

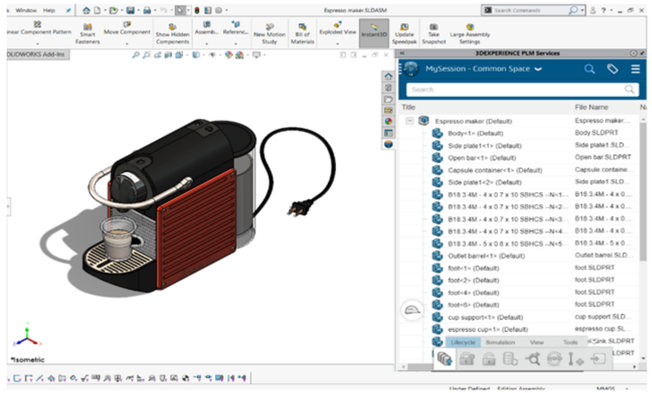

Time-consuming to create

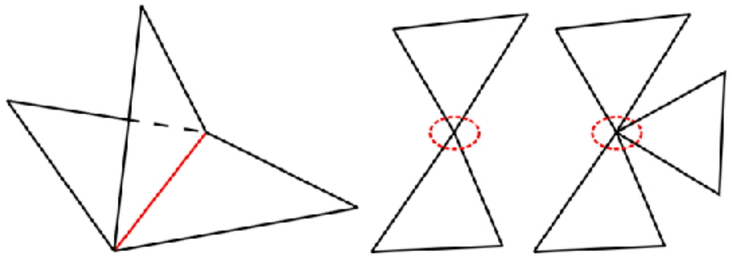

High-bar for geometric computation

Moving beyond relying on large 3D datasets

Pre-trained 2D models

Leveraging pre-trained 2D models to learn in 3D

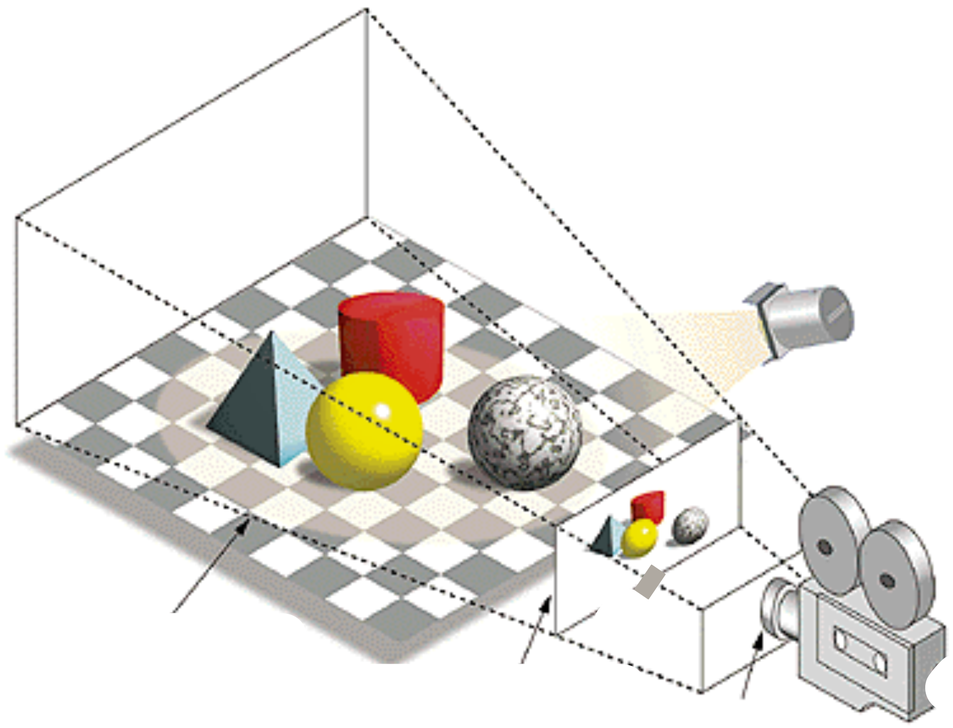

Connecting 2D to 3D via rendering

Take advantage of big 2D datasets

Images are how we perceive the 3D shape

Neural Mesh Editing <u>without</u> 3D data!

Stylization

Text2Mesh [CVPR 2022]

Localization

3D Highlighter [CVPR 2023]
3D Paintbrush [CVPR 2024]

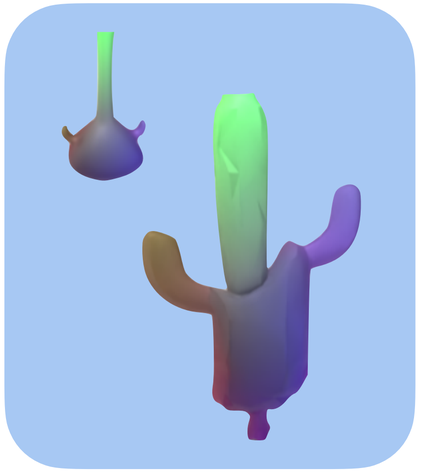

Deformation

TextDeformer [SIGGRAPH 2023]
MeshUp [3DV 2025]
Geometry in Style [CVPR 2025]

Segmentation

iSeg [SIGGRAPH Asia 2024]

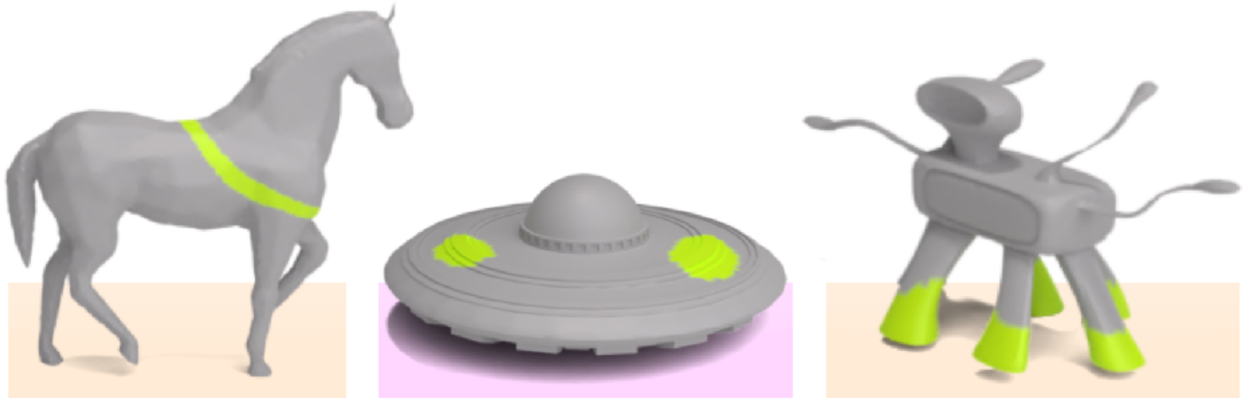

Pre-trained image models for localization

3D Highlighter: Localizing Regions on 3D Shapes via Text Descriptions [CVPR 2023]

Dale Decatur

Itai Lang

Rana Hanocka

Result from our prior work: Text2Mesh

Key question: can we extract the underlying analysis inherent in the synthesis process?

No explicit segmentation, but can we tease it out?

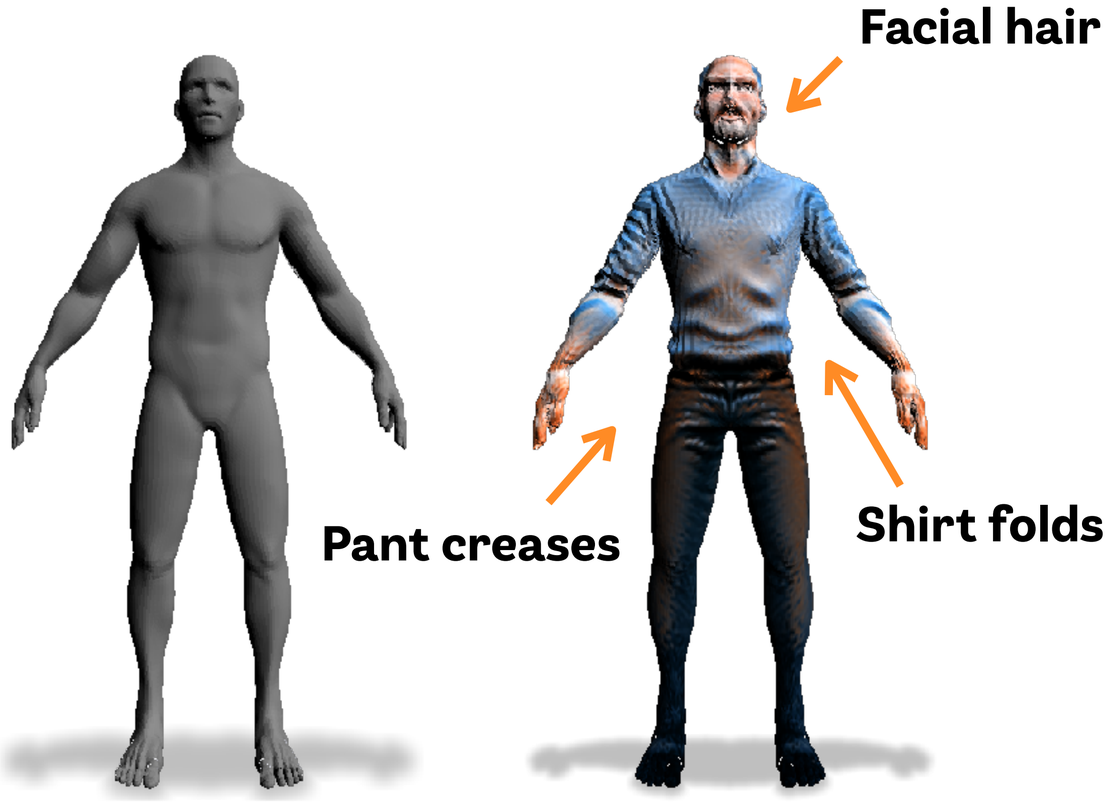

Text2Mesh: Text-Driven Neural Stylization for Meshes [CVPR 2022]

Our 3D Highlighter localizes semantic regions on a shape

using text as input

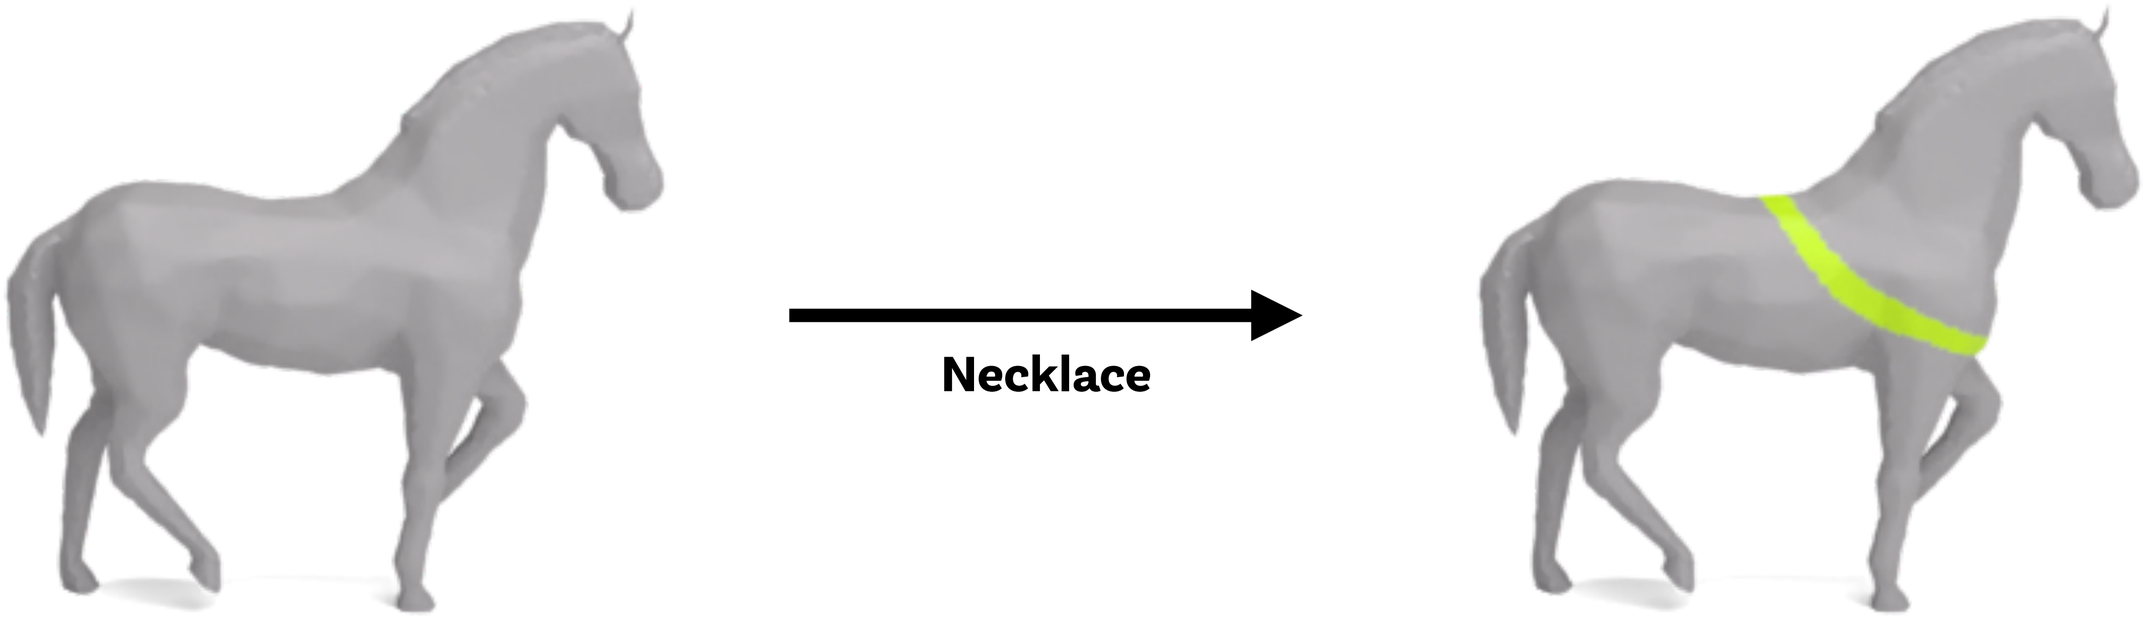

Our 3D Highlighter: provides the vertices that correspond to that region!

3D Highlighter reasons about where text-specified regions belong on a shape

Poncho

Necklace

Headphones

Explicitly perform analysis & use for synthesis

Previously: Text2Mesh

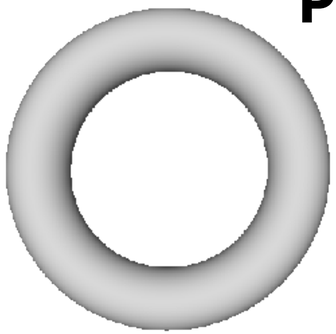

Mesh

→

Neural Style Field

→ *Style Attributes*

Stylized Mesh

Now: 3D Highlighter

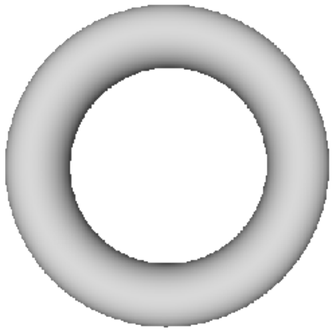

Mesh

→

Neural Localization Field

→ *Localization*

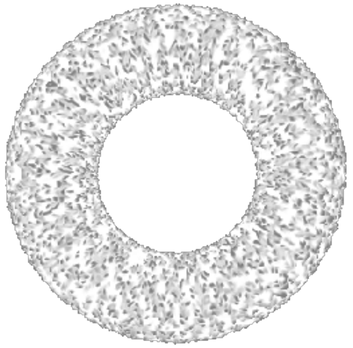

Localization

→

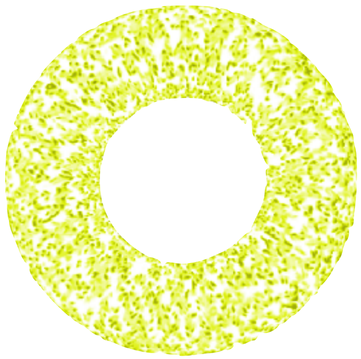

Localization-Explicit Appearance

The gist of our neural highlighter

How to supervise?
- *no 3D dataset!*
- *no pairs of text & 3D styles*

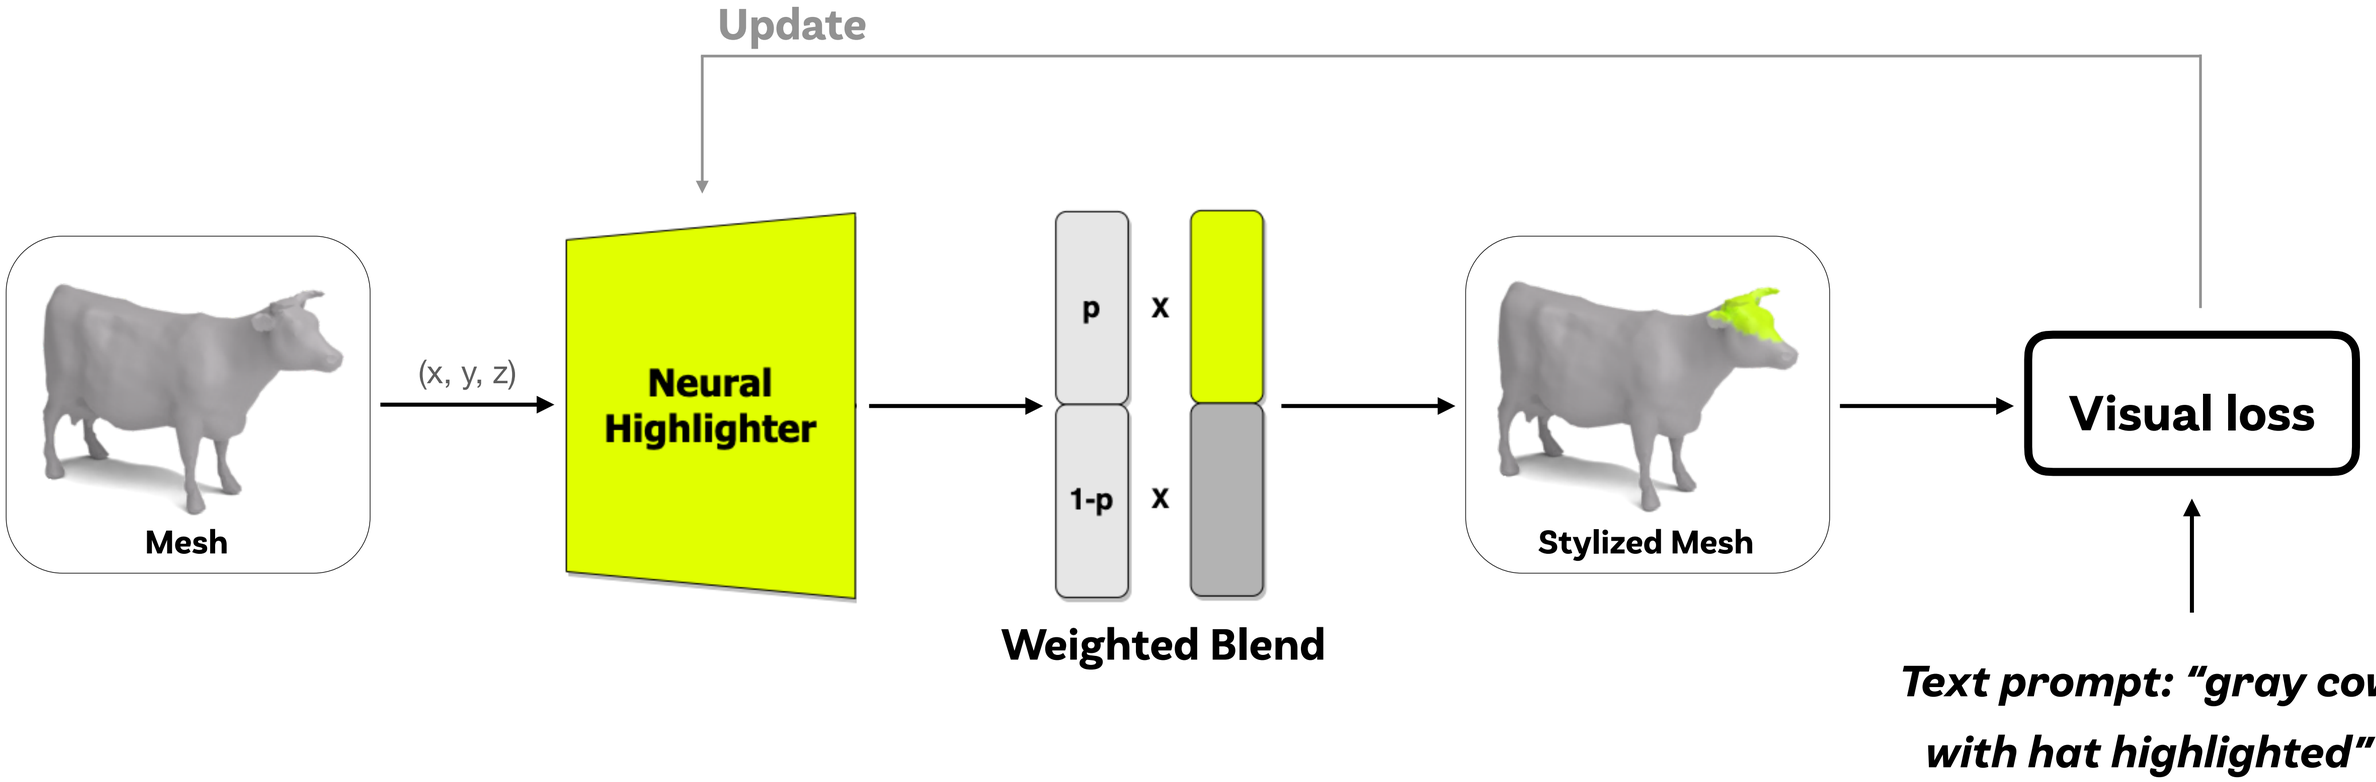

Supervise using 2D Renderings & CLIP

Spoiler from the future: diffusion (SDS) also works

Using 3D Highlighter for Shape Editing

Localized editing

Controllable Compositionality

Geometric edits

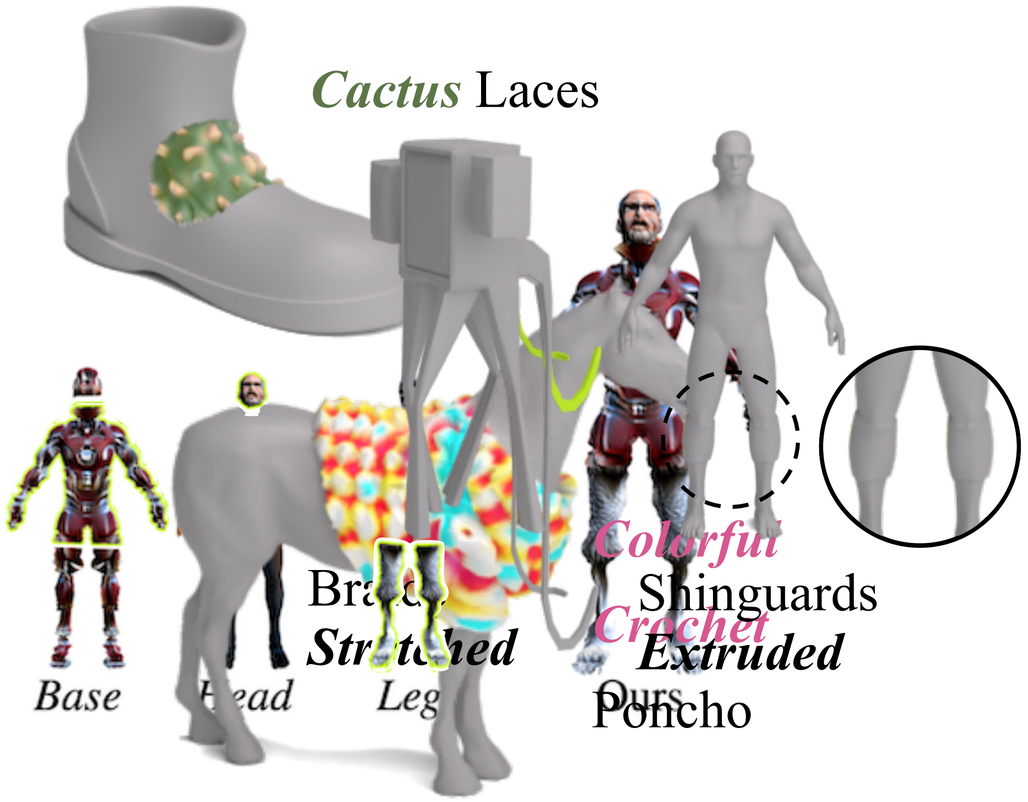

THREEDLE

THE UNIVERSITY OF CHICAGO

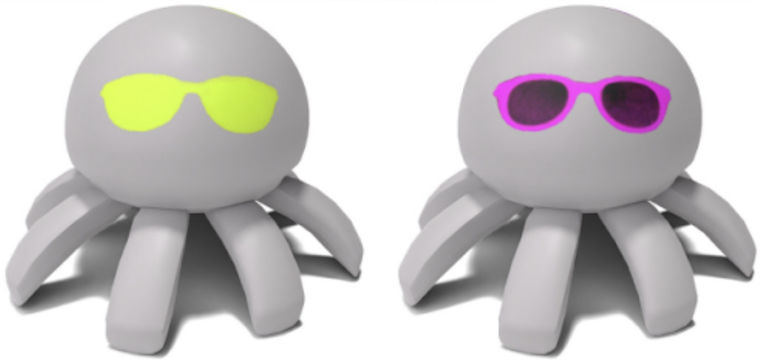

Pre-trained image models for local texture edits

3D Paintbrush: Local Stylization of 3D Shapes with Cascaded Score Distillation [CVPR 2024]

Dale Decatur

Itai Lang

Kfir Aberman

Rana Hanocka

Idea: we can synthesizing textures in tandem with localization

... and get better / more fine-grained localizations as a result!

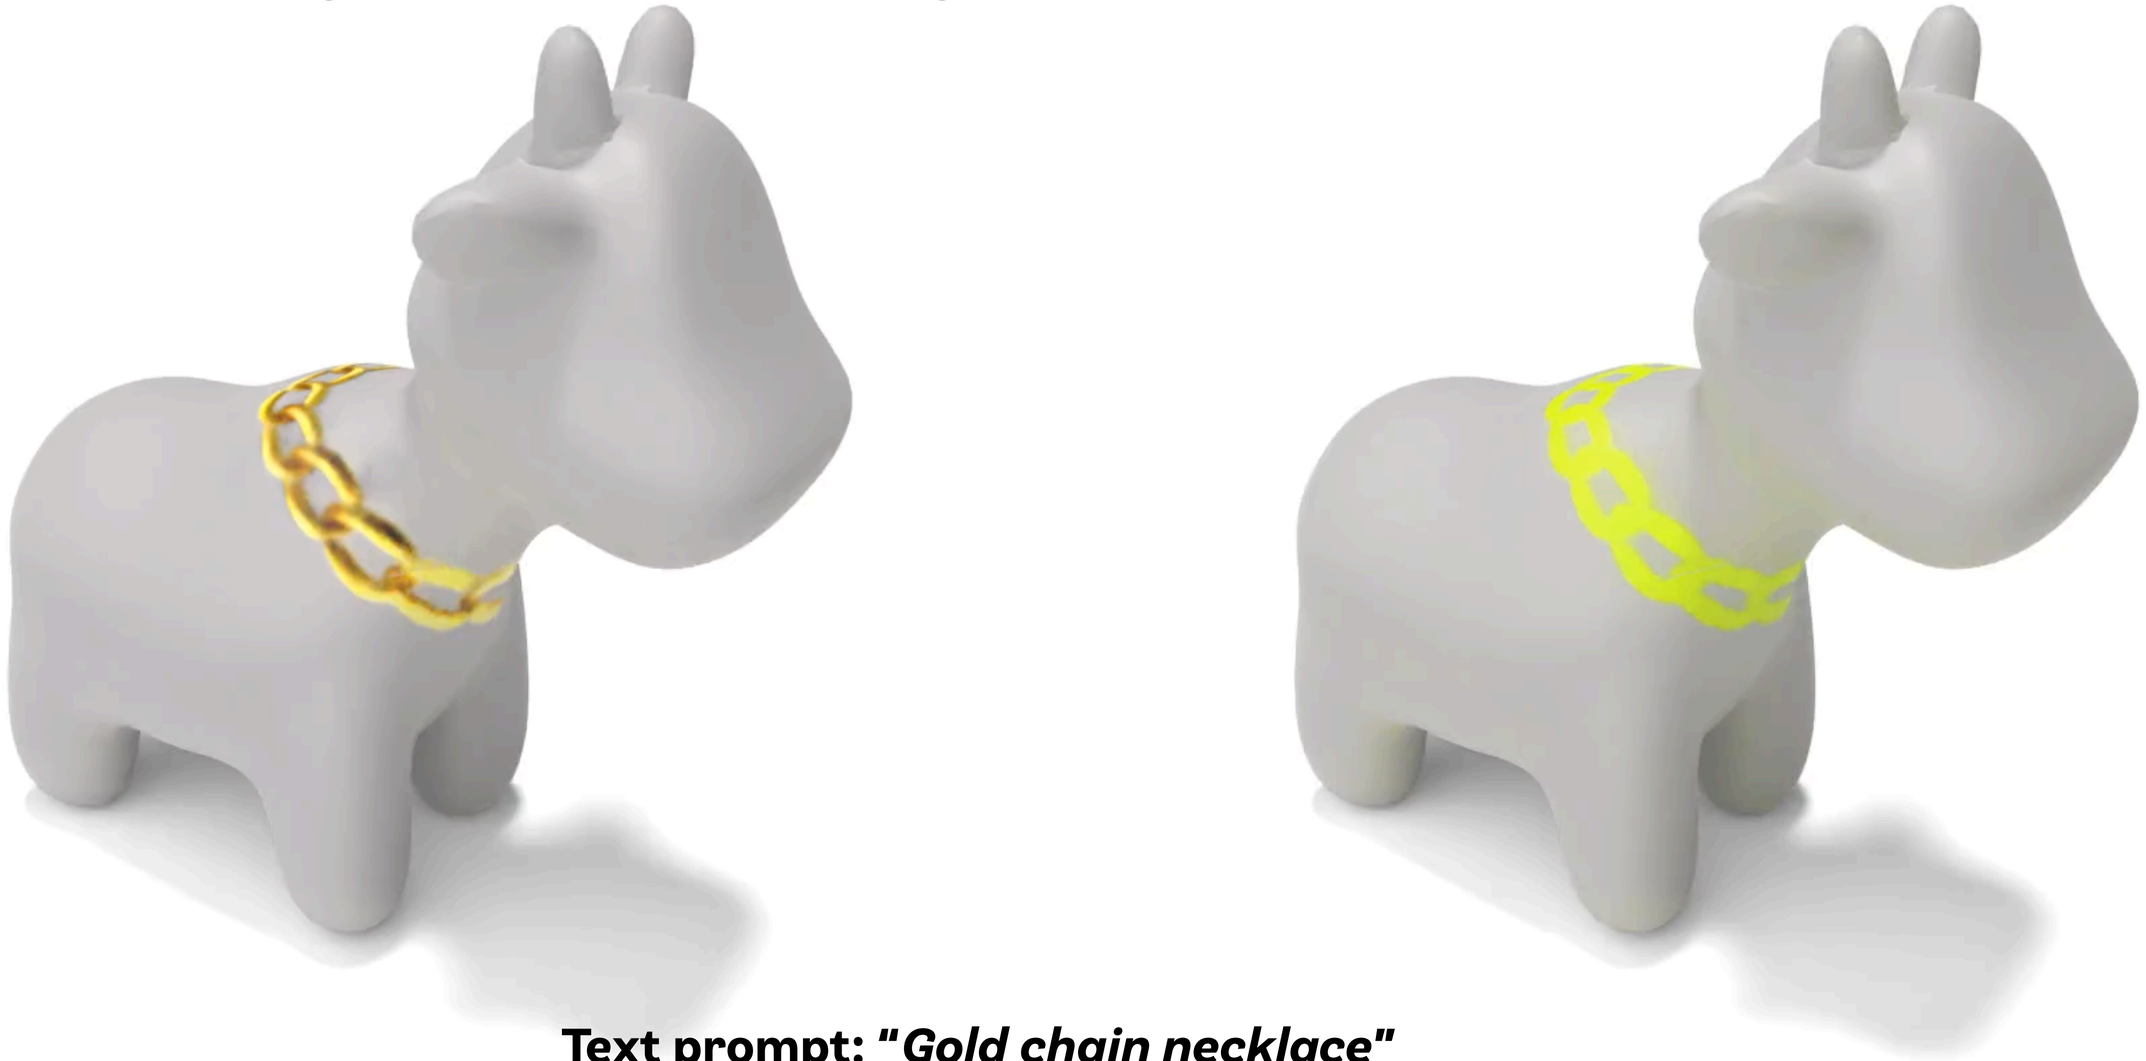

Text prompt: "*Gold chain necklace*"

Synthesizing local texture regions

Precise composition of multiple local textures

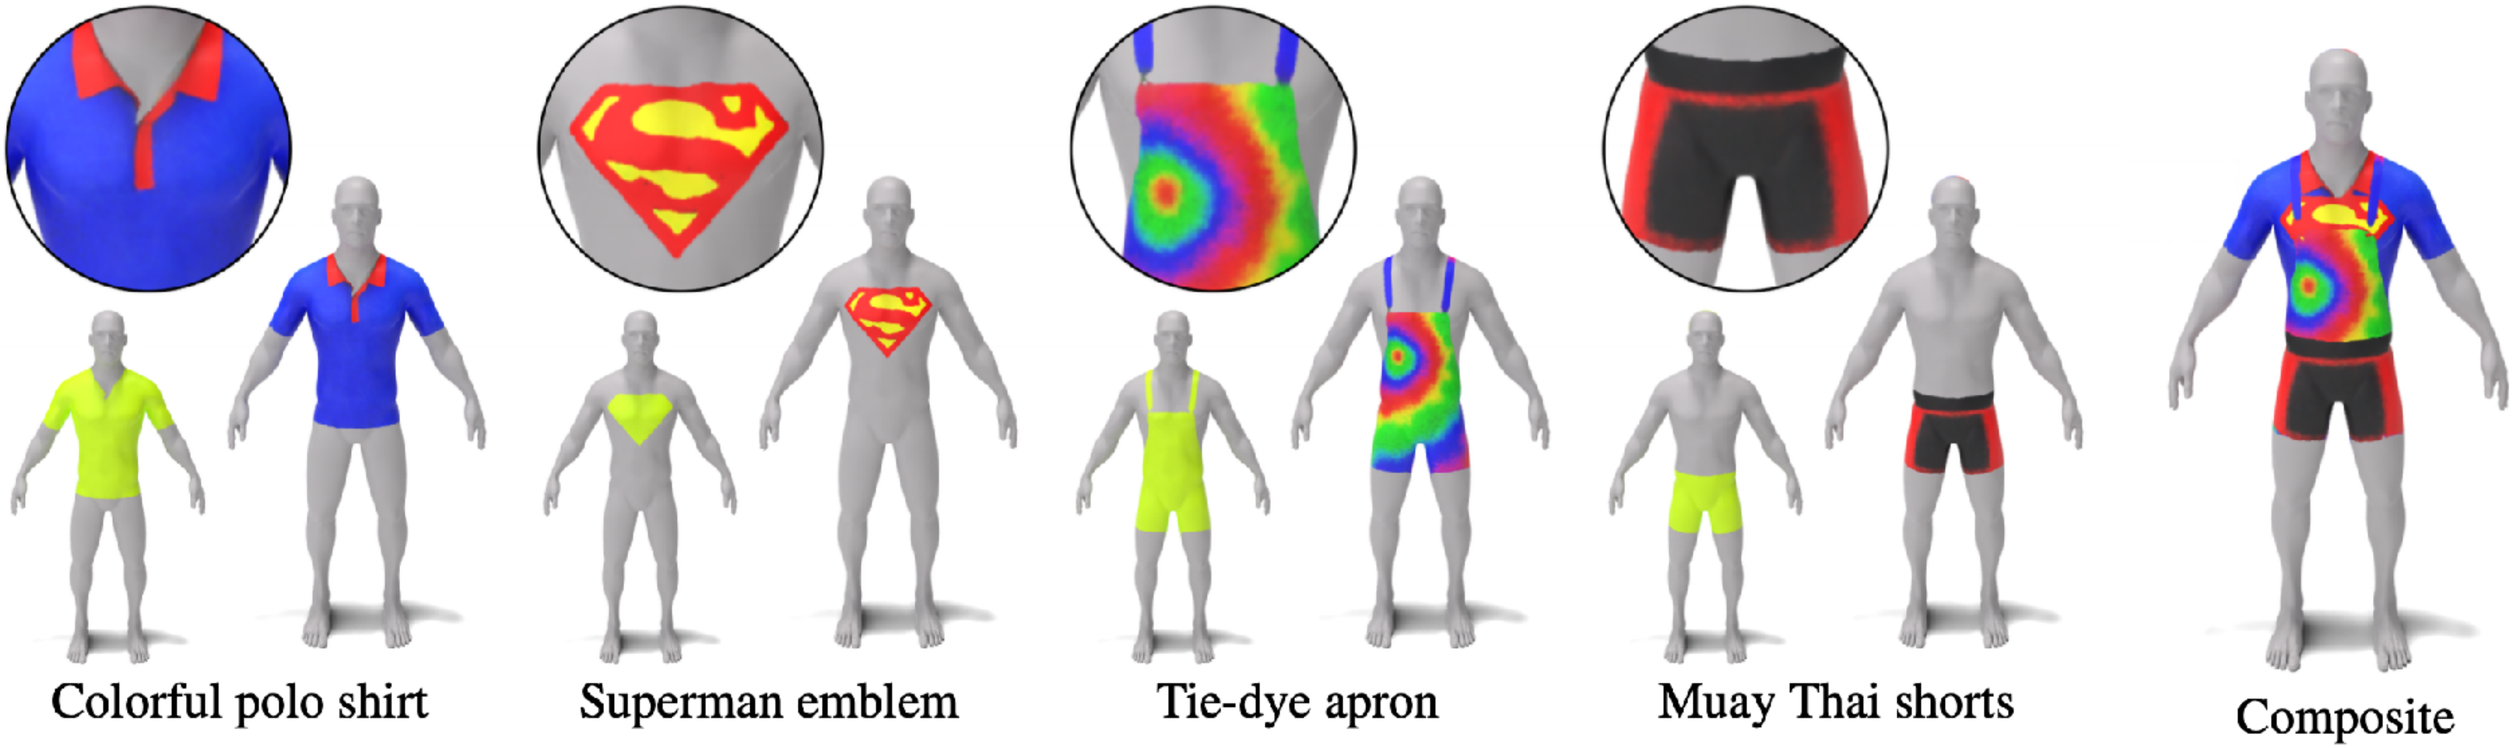

Explicit fine-grained segmentation masks offer additional level of control

Gist of 3D Paintbrush

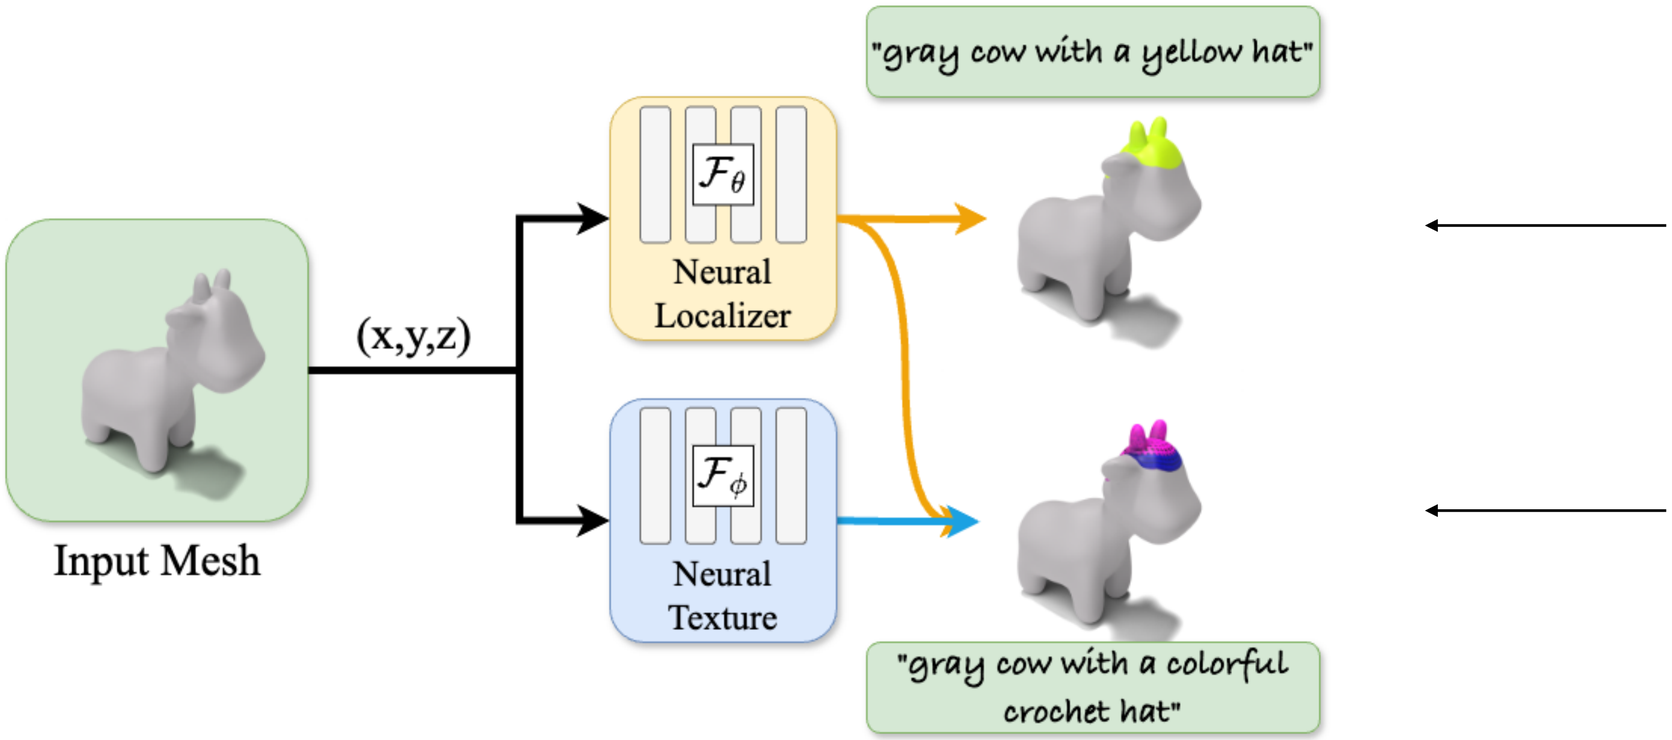

Cascaded Score Distillation (CSD)

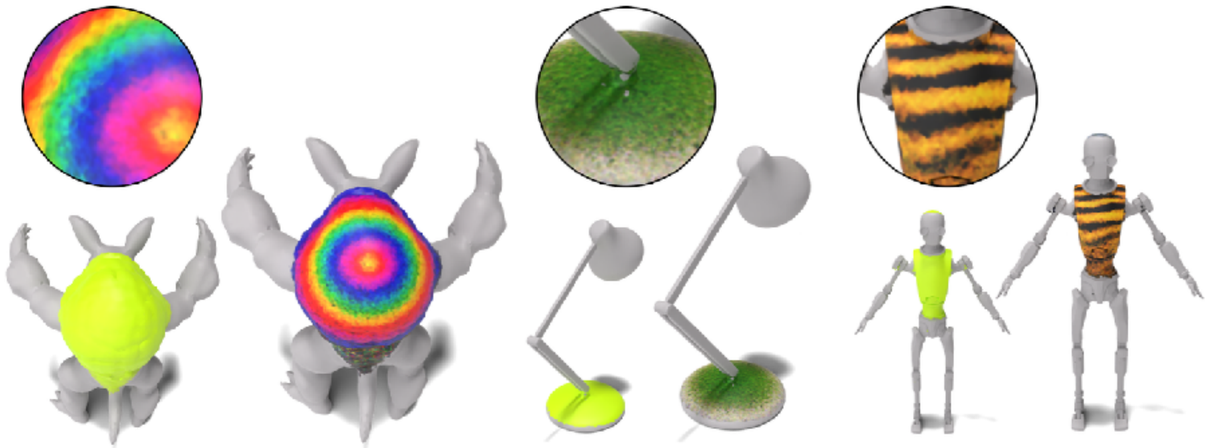

SDS

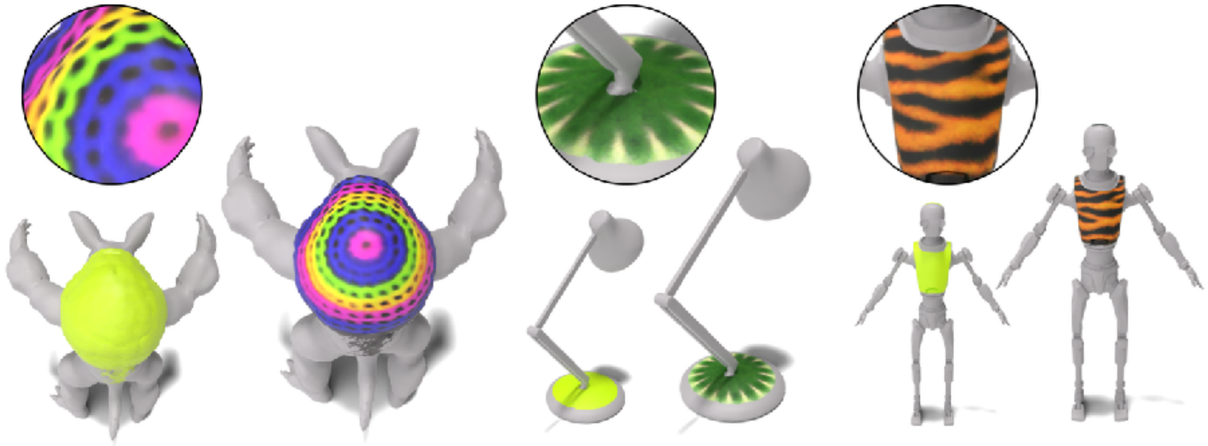

CSD

Existing methods use SDS – the first stage of cascaded model

CSD: Leverage multiple stages of a cascaded diffusion model

The effect of different diffusion stages

Text prompt: “Fancy gold watch”

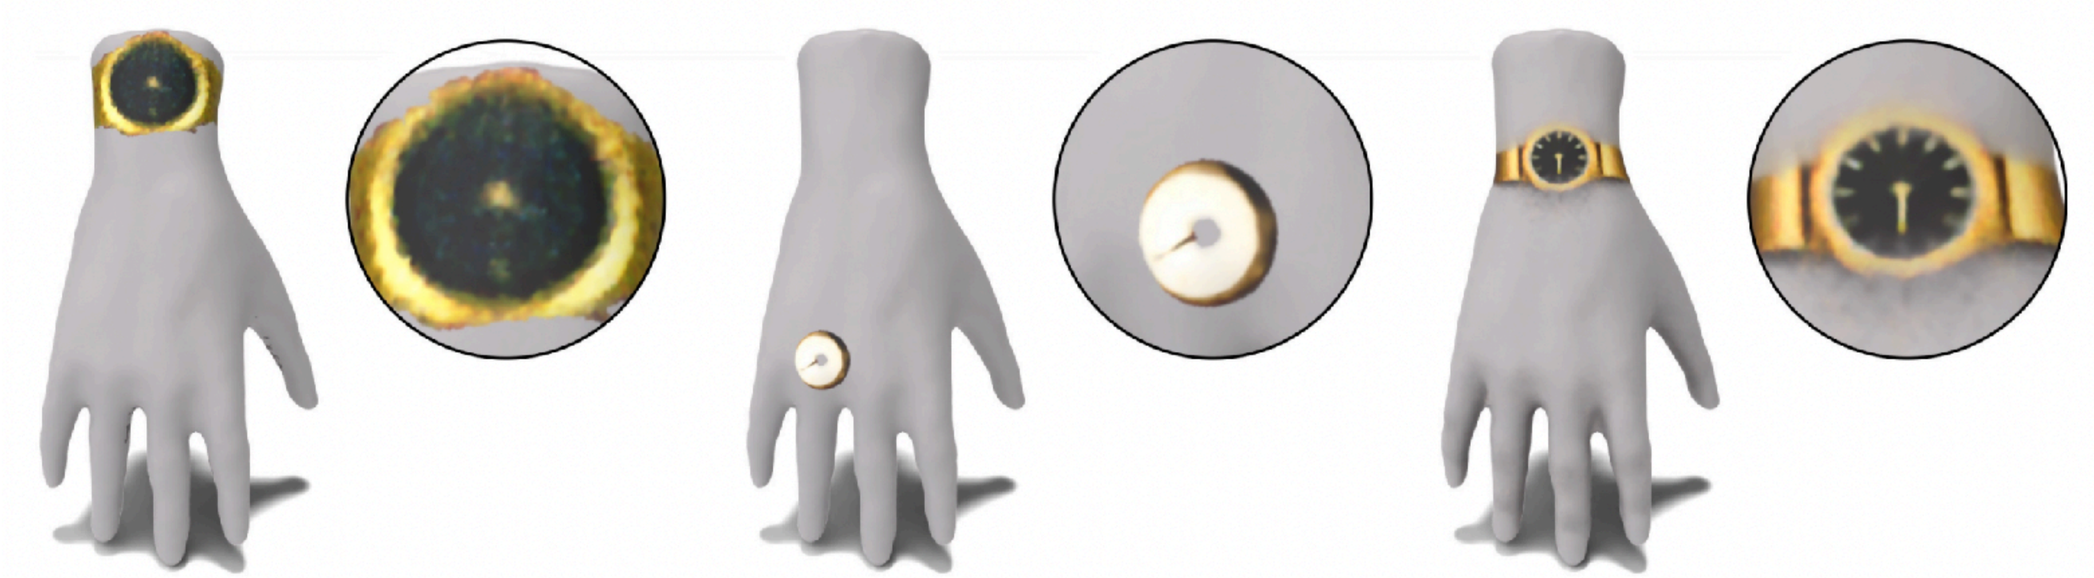

Only stage 1

Only stage 2

CSD

3D Paintbrush assorted results

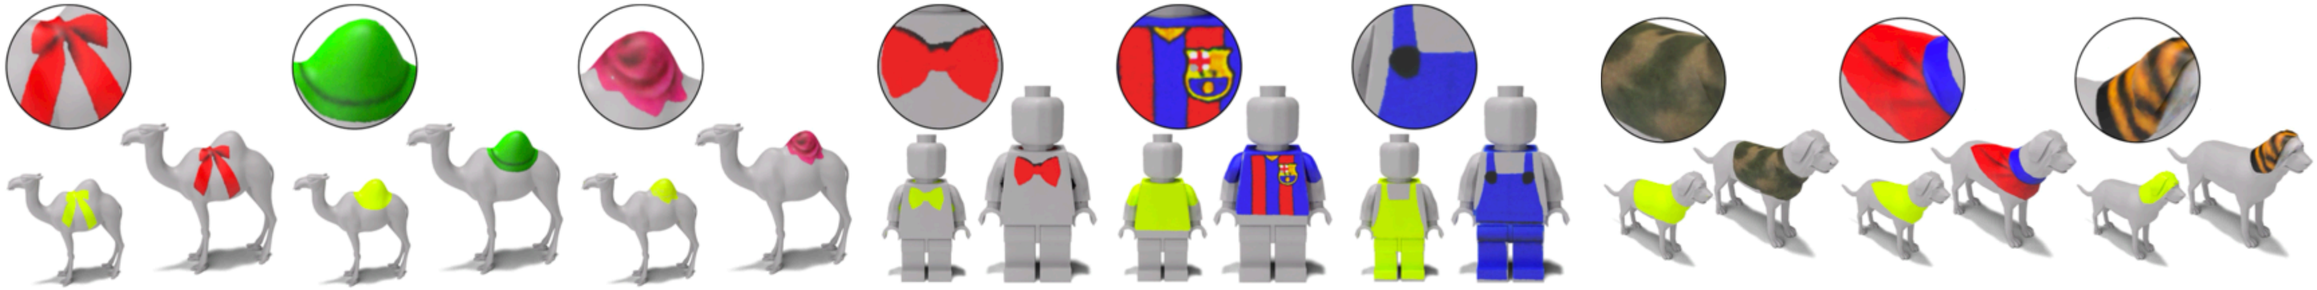

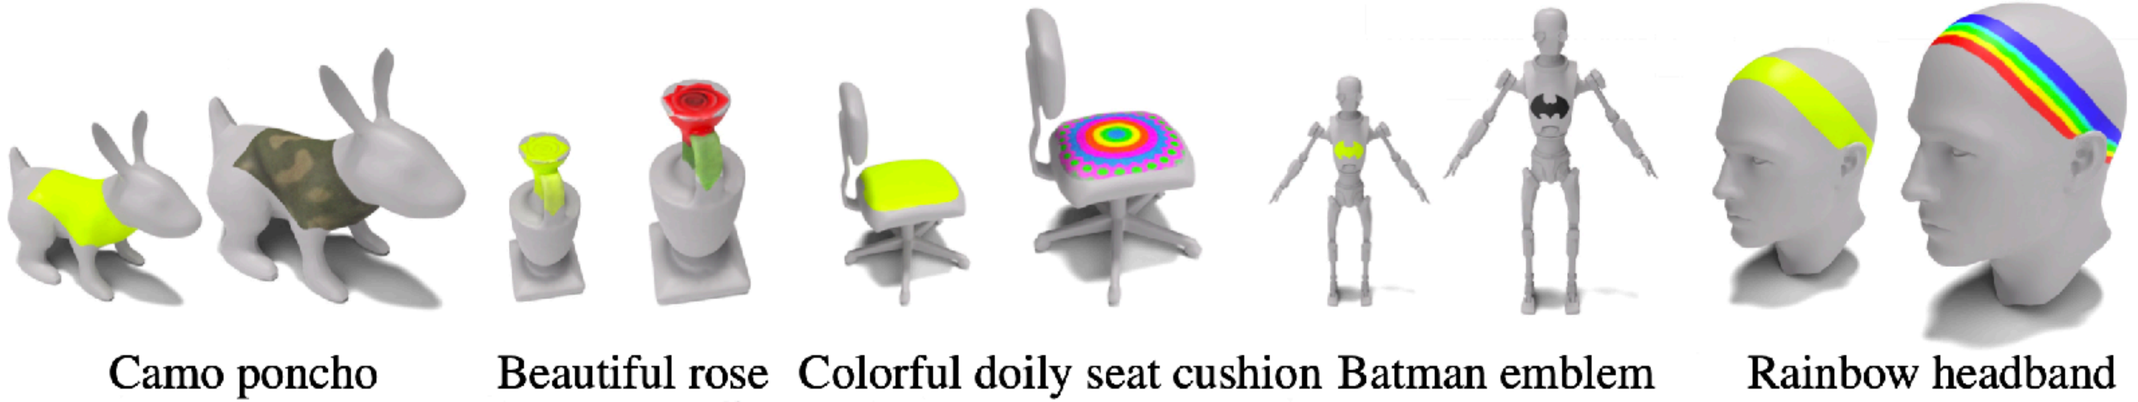

3D Paintbrush Results

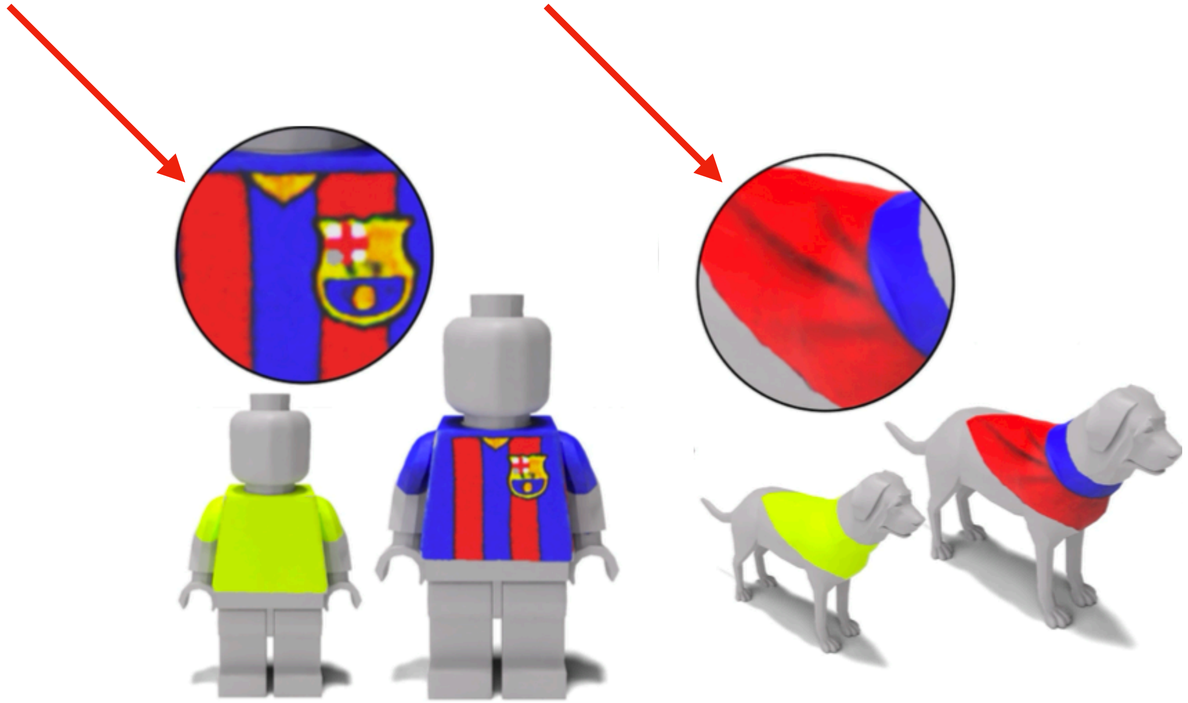

Barcelona jersey

Superhero cape

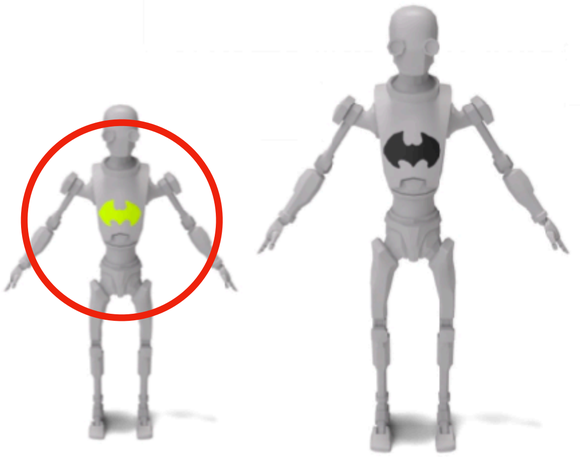

Batman emblem

Red flip flops

Text-driven localization summary

- Extract the underlying analysis inherent in the synthesis process
- Explicit and fine-grained segmentation masks give additional control
- Exploit pre-trained 2D foundation for segmentation in 3D

Neural Mesh Editing <u>without</u> 3D data!

Stylization

Text2Mesh [CVPR 2022]

Localization

3D Highlighter [CVPR 2023]
3D Paintbrush [CVPR 2024]

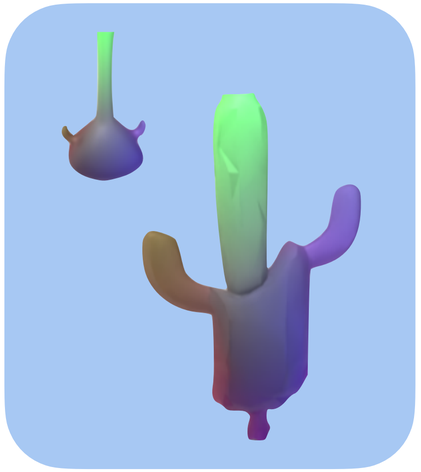

Deformation

TextDeformer [SIGGRAPH 2023]
MeshUp [3DV 2025]
Geometry in Style [CVPR 2025]

Segmentation

iSeg [SIGGRAPH Asia 2024]

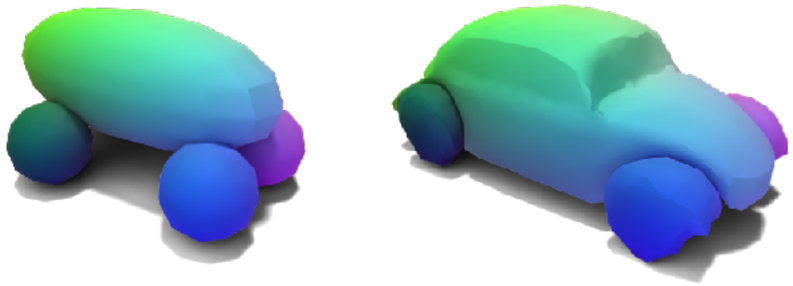

Pre-trained image models for deformation

MeshUp: Multi-Target Mesh Deformation via Blended Score Distillation [3DV 2025]


Hyunwoo Kim


Itai Lang


Noam Aigerman


Thibault Groueix


Vladimir G. Kim


Rana Hanocka

Expressive deformations of meshes using text prompts

General deformation framework, capable of highly expressive mesh manipulations!

Input 3D object

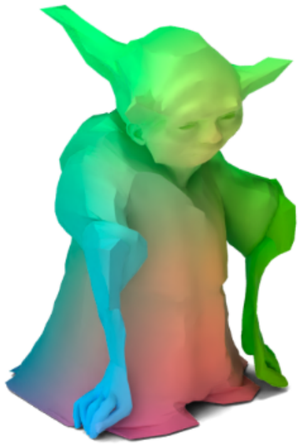

Yoda

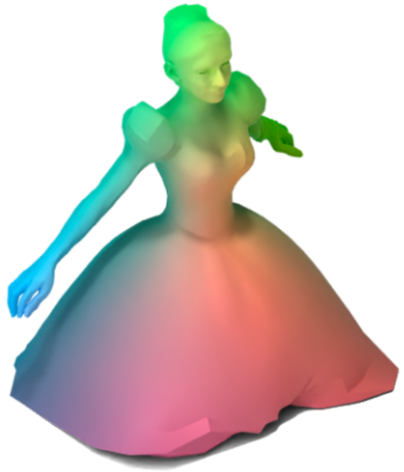

Cinderella

Compared to text-to-3D “from scratch” approaches

Renderings might look beautiful but extracting the underlying mesh in post-process often results in artifacts, floaters, missing parts, undesirable triangulation, etc

Magic3d

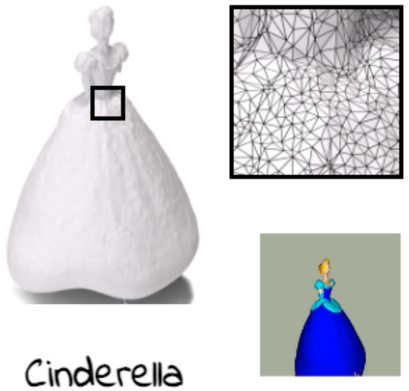

TextMesh

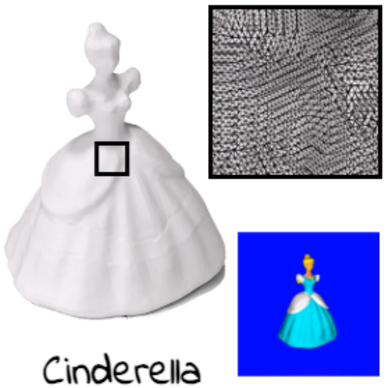

ProlificDreamer

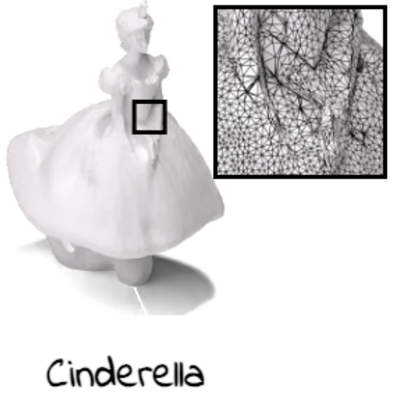

MVDream

Ours

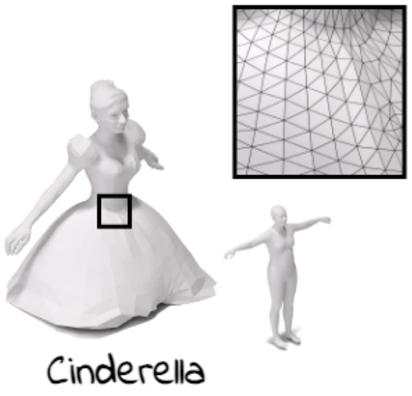

Gist of MeshUp (single-target)

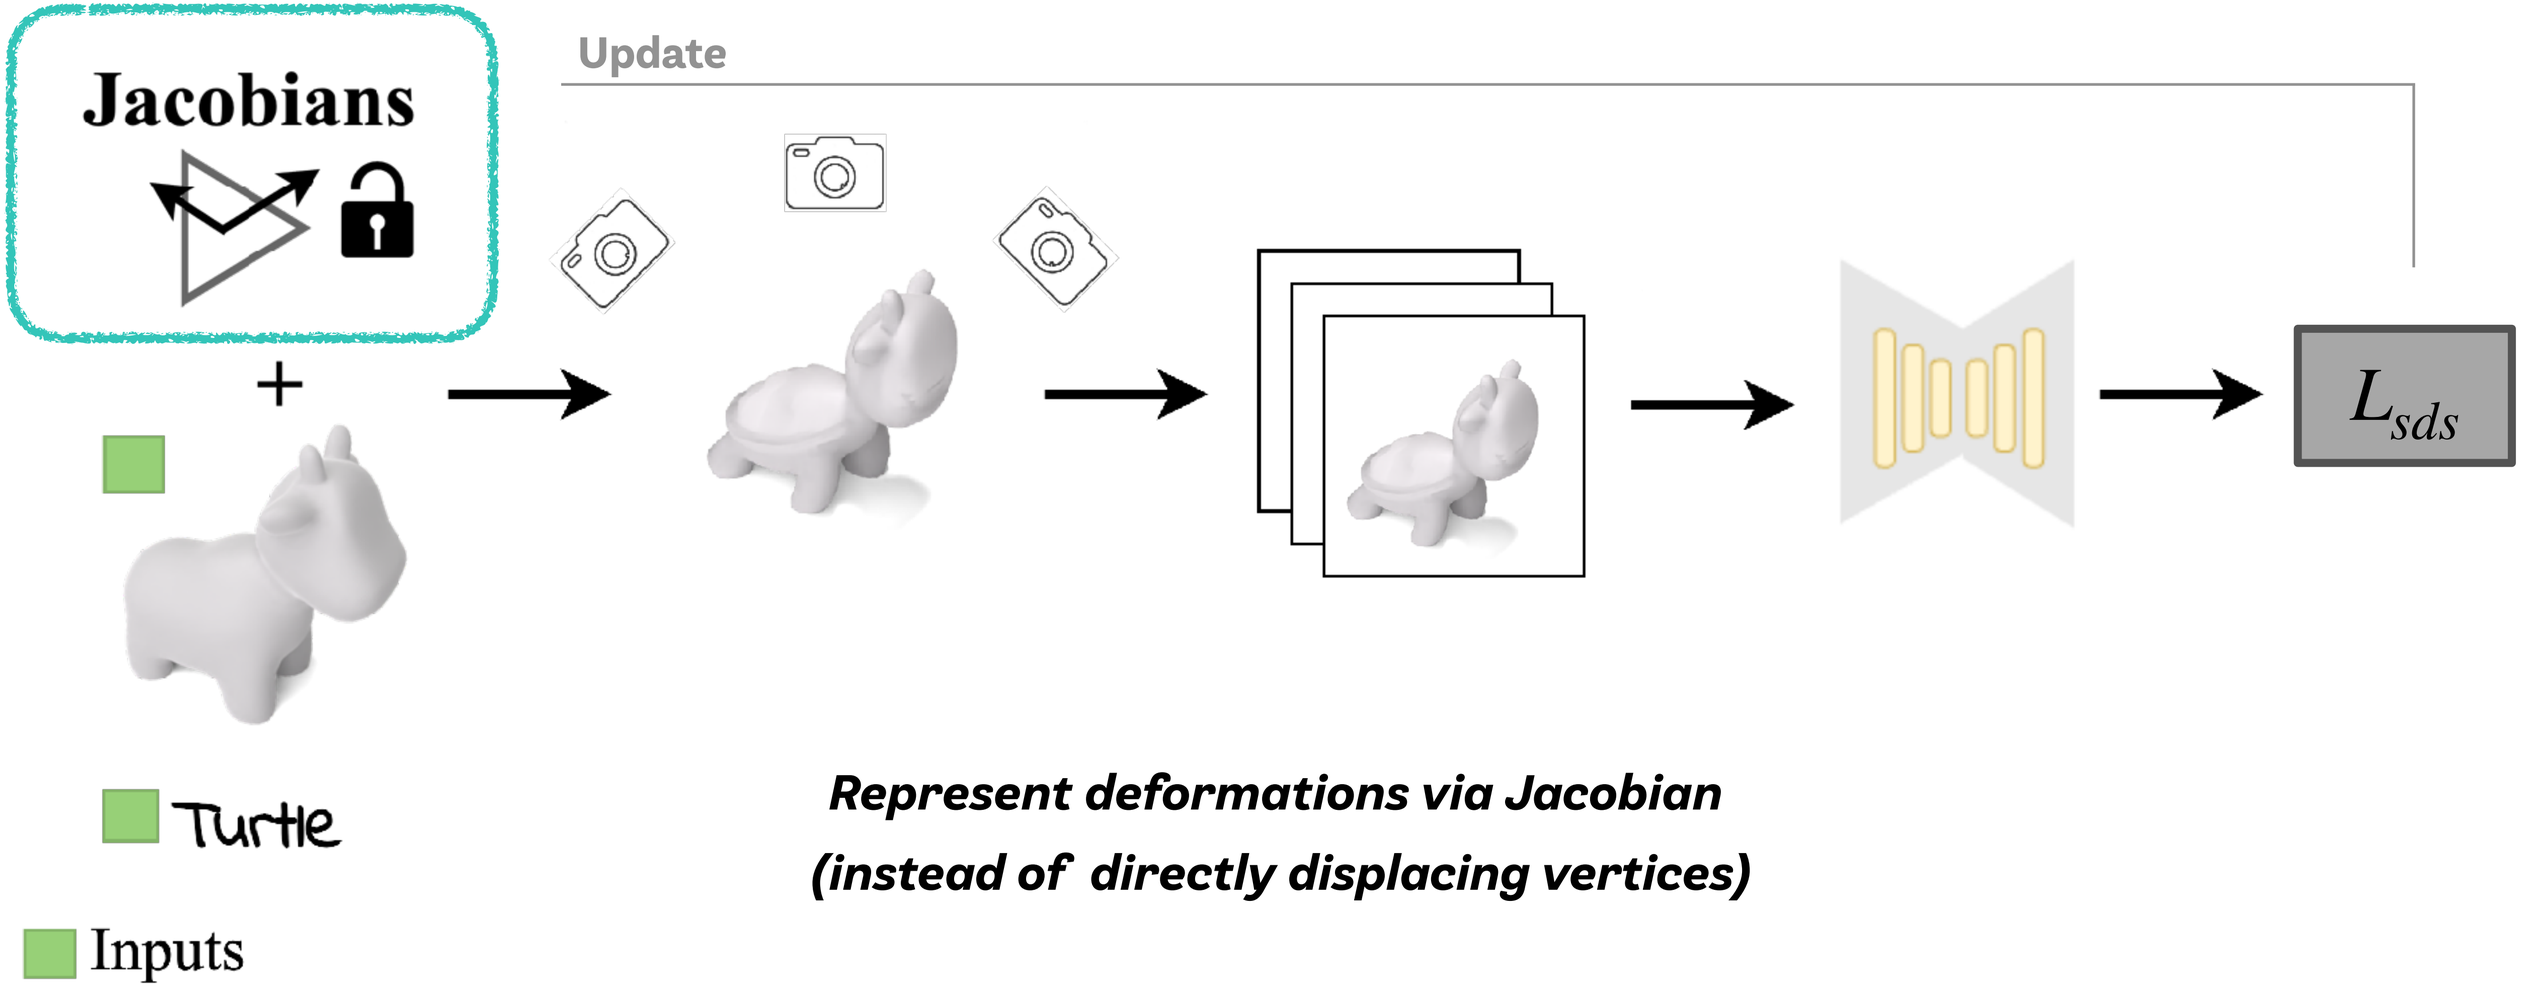

Represent deformations via Jacobian

(instead of directly displacing vertices)

Mix deformation concepts

Input 3D object

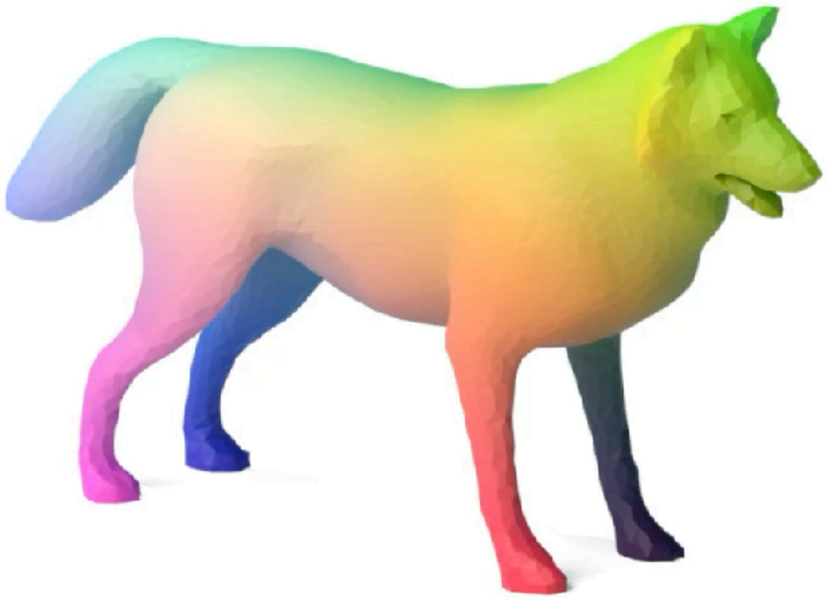

Multi-target deformation pipeline

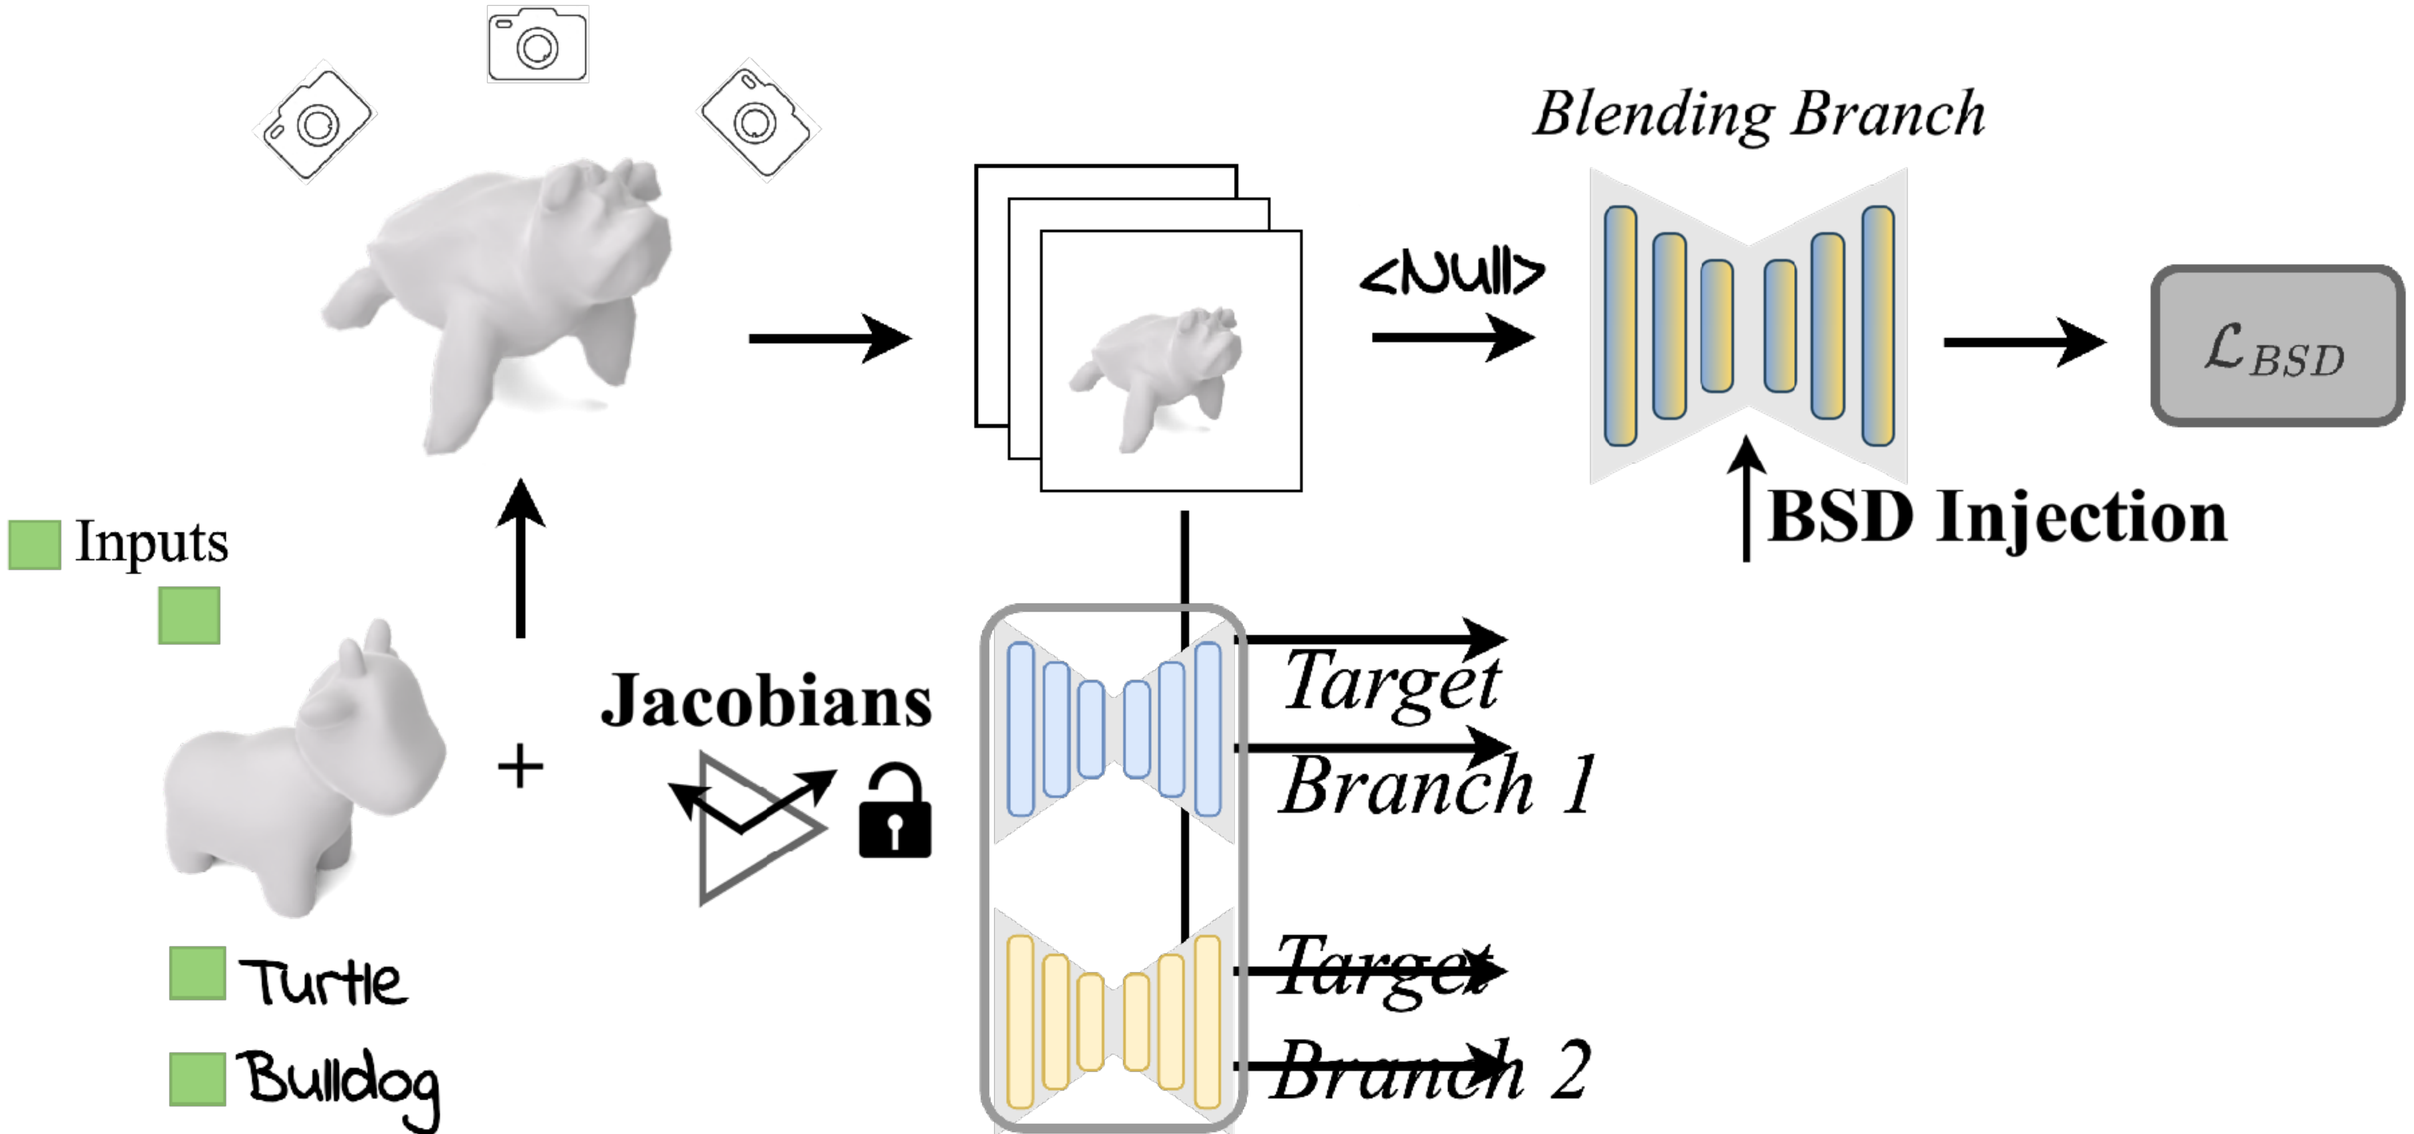

Local control over deformation

Input 3D object

Use image prompts as input for deformation

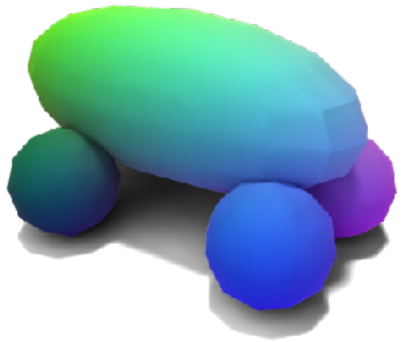

Input 3D object

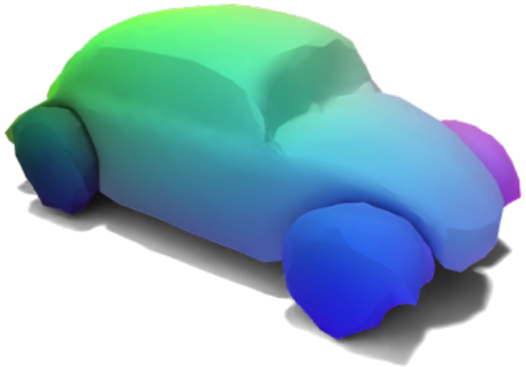

Deformed

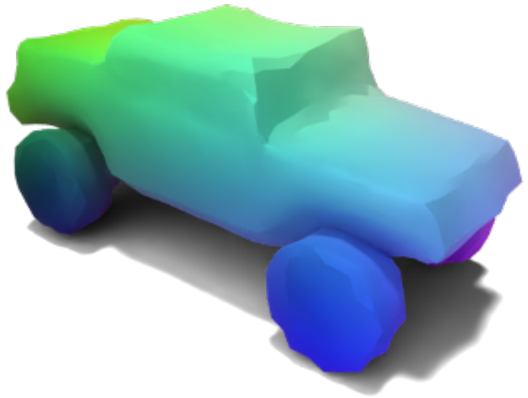

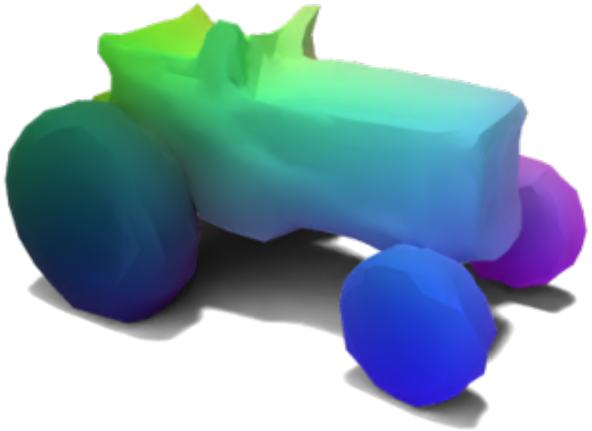

Target Image

THREEDLE

THE UNIVERSITY OF CHICAGO

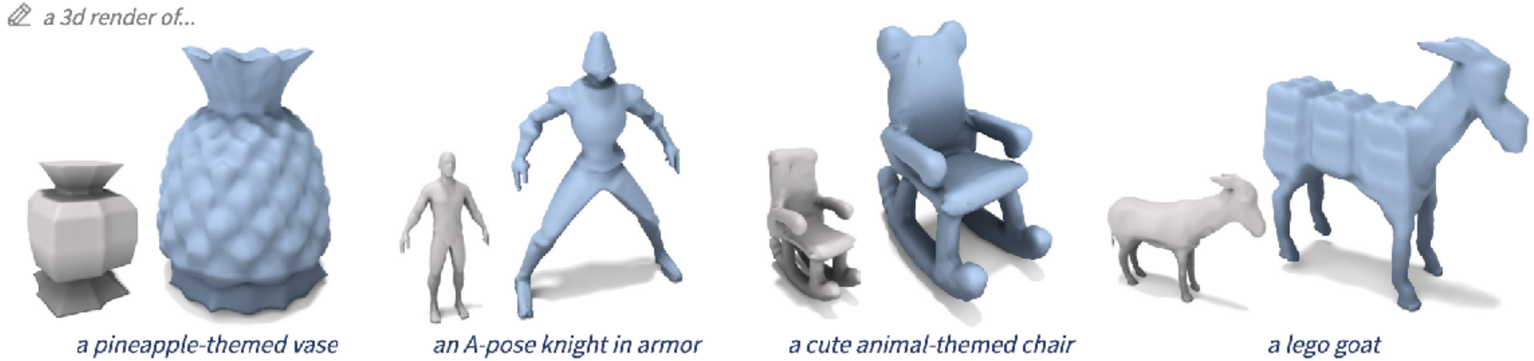

Geometry in Style: 3D Stylization via Surface Normal Deformation

CVPR 2025


Nam Anh Dinh


Itai Lang


Hyunwoo Kim


Oded Stein


Rana Hanocka

Identity-preserving stylization of mesh geometry

Input 3D object

Deformed

Identity-preserving stylization of mesh geometry

Gain more control over the generation process

How: reformulate and restrict the underlying deformation procedure

Input 3D object

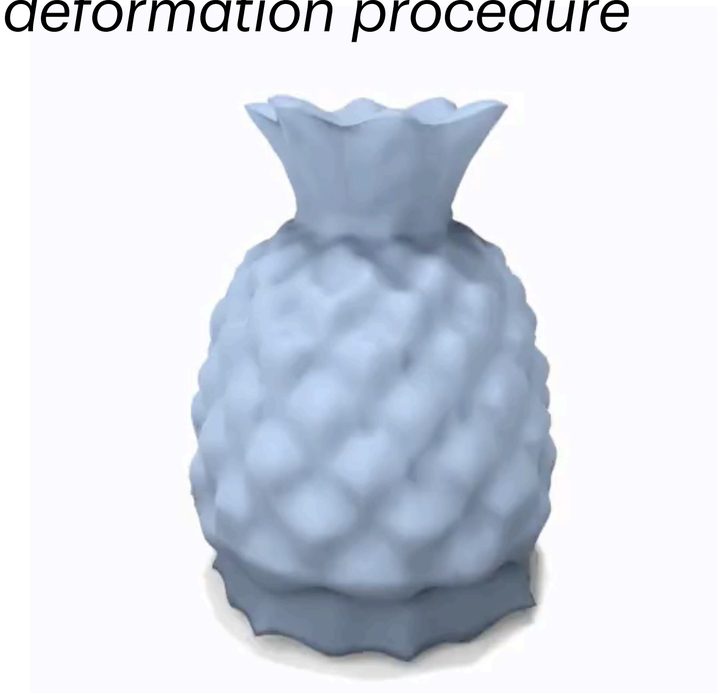

Deformed

One idea: use normals

Use normals to represent the space of deformations

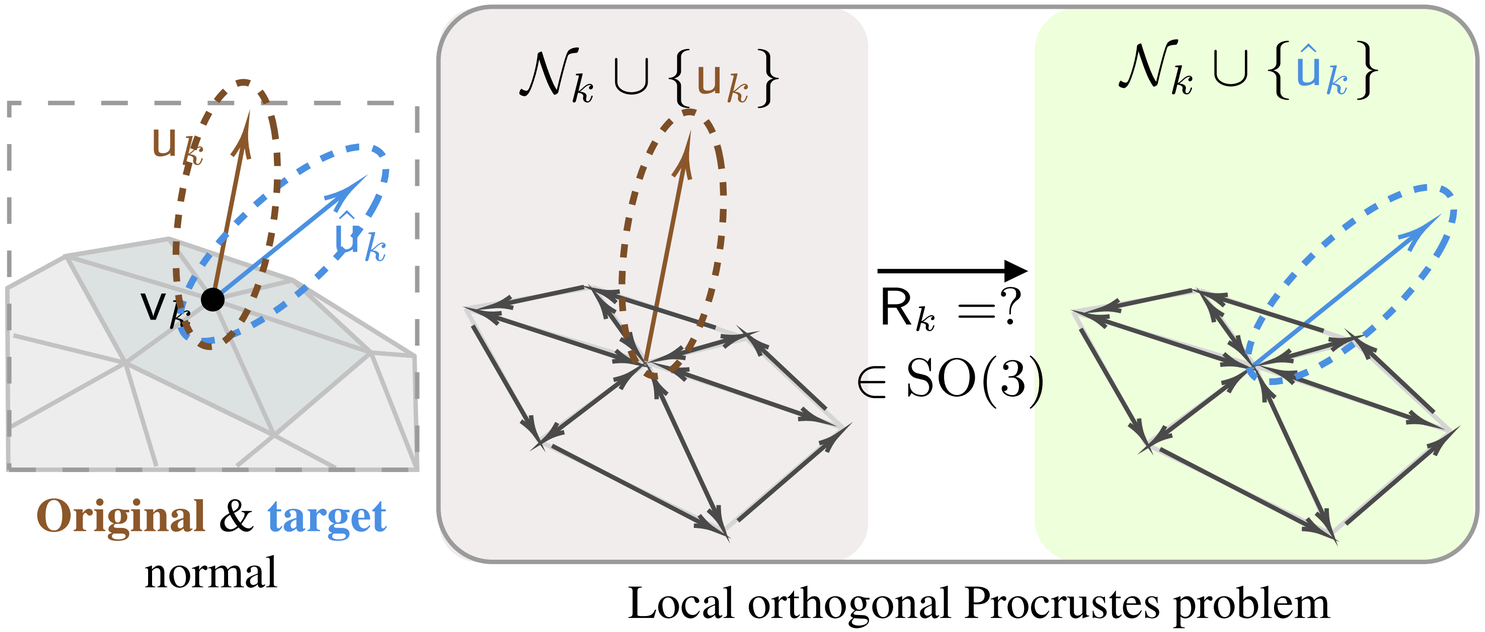

Original & **target** normal

Local orthogonal Procrustes problem

Recover a "best fit rotation" to the target normal via a local Procrustes solve

Lambda controls strength of deformation

More restrictive than Jacobians,
but preserves shape identity while still expressive

dARAP: Differentiable ARAP

a differentiable one-step adaptation of classical As-Rigid-As-Possible deformation

All runs in a single forward pass!

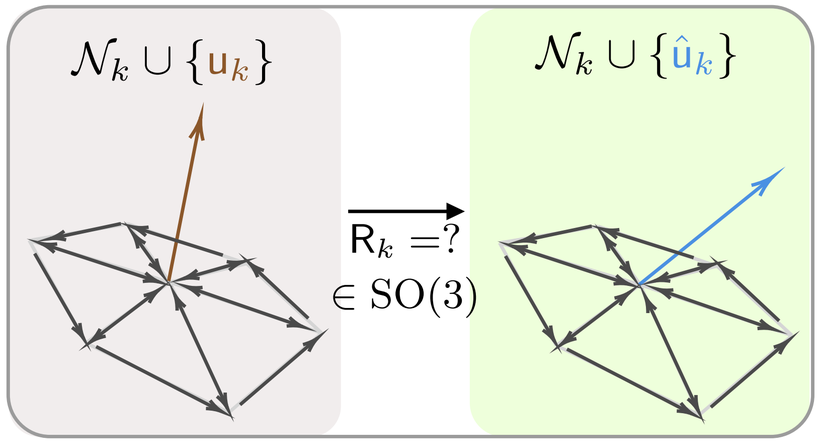

Local orthogonal Procrustes problem

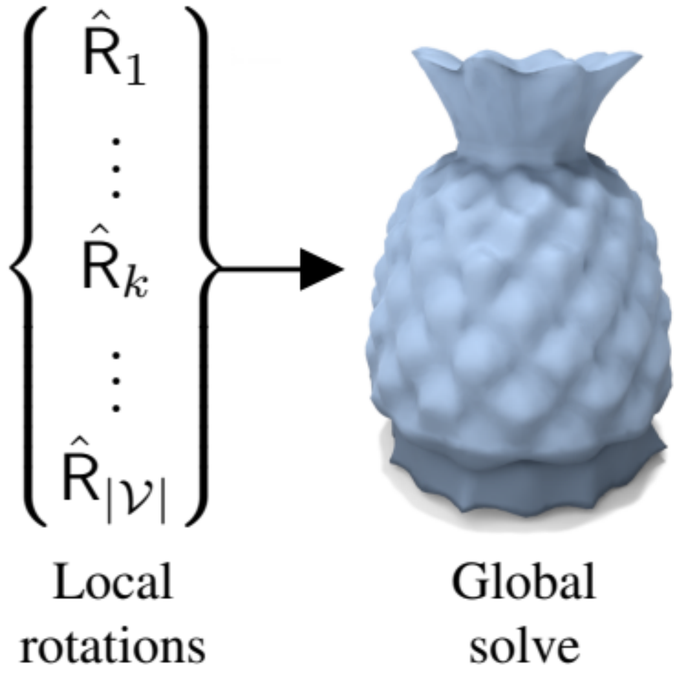

<u>Local</u>: solve for optimal rotation based on target normal

<u>Global</u>: solve for overall deformation from rotations

Gist of our approach

Optimize a normal vector per vertex

How to supervise?

- no 3D dataset!

- no pairs of text & 3D styles

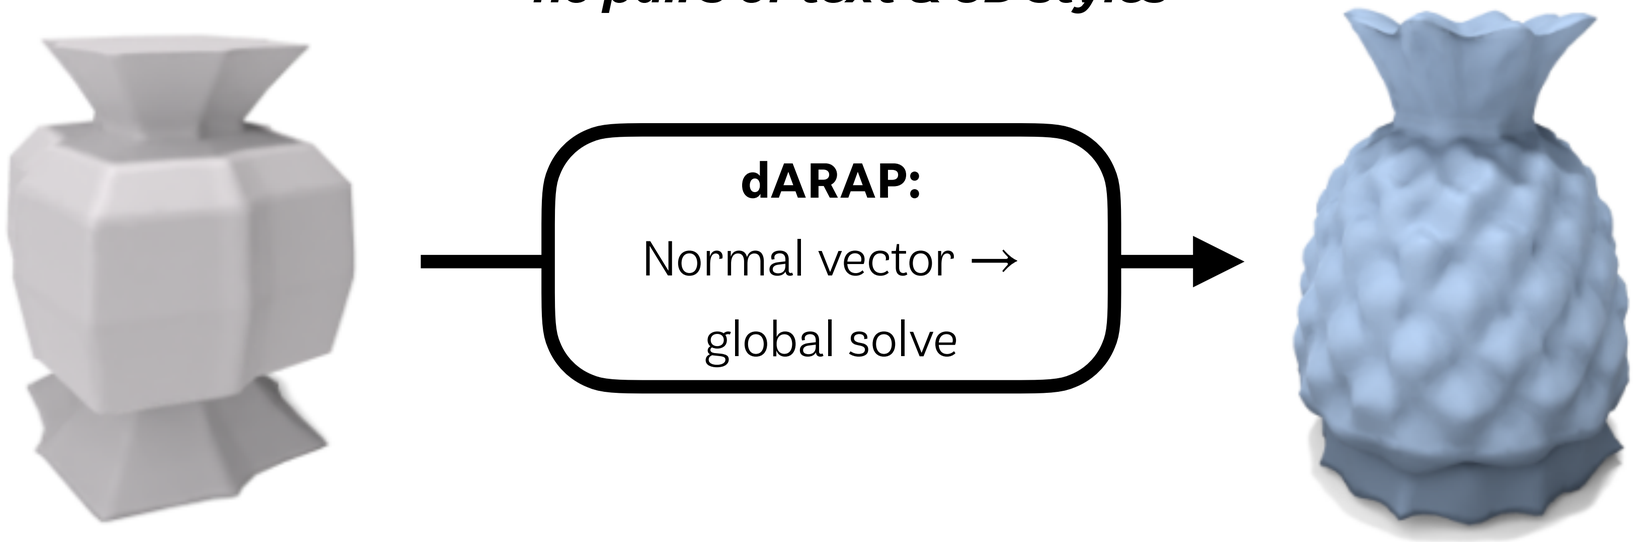

a pineapple-themed vase

Key idea: Differentiable render & supervise with a visual loss

Geometry in Style

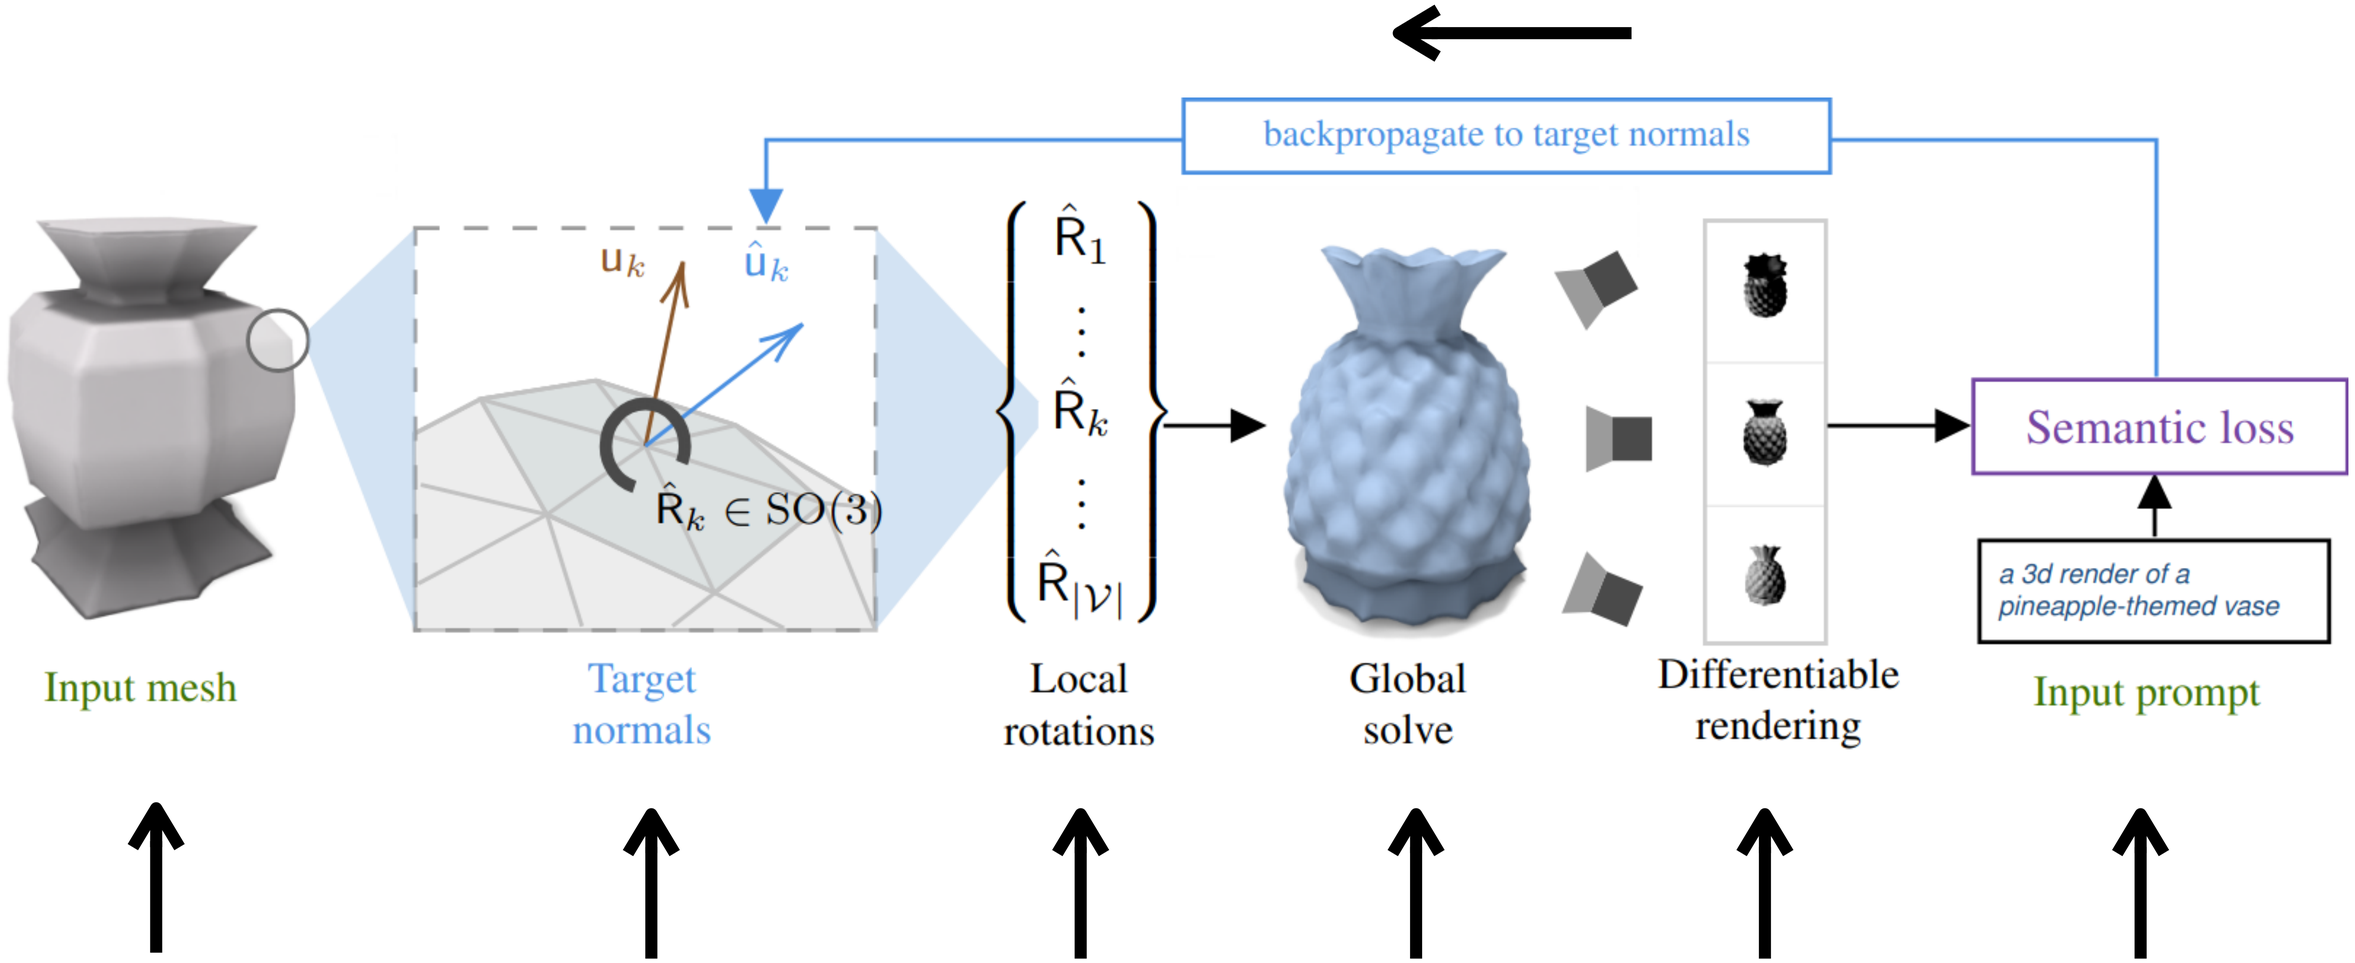

Control the strength of deformation – in post

a 3d render of a robot lion

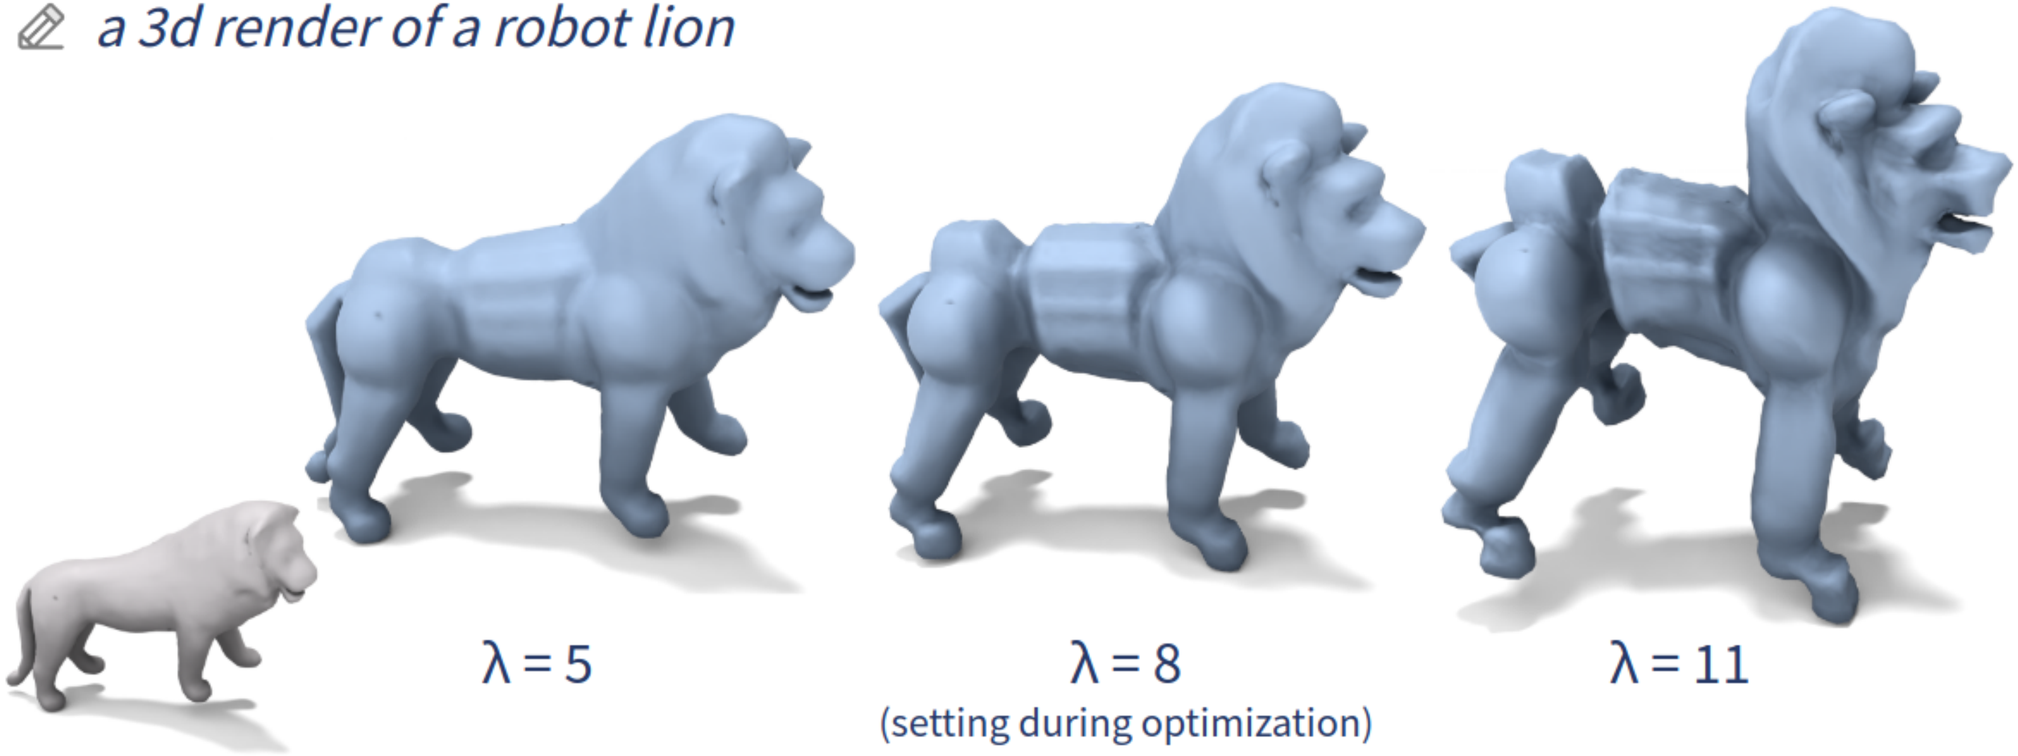

Controllable deformation region

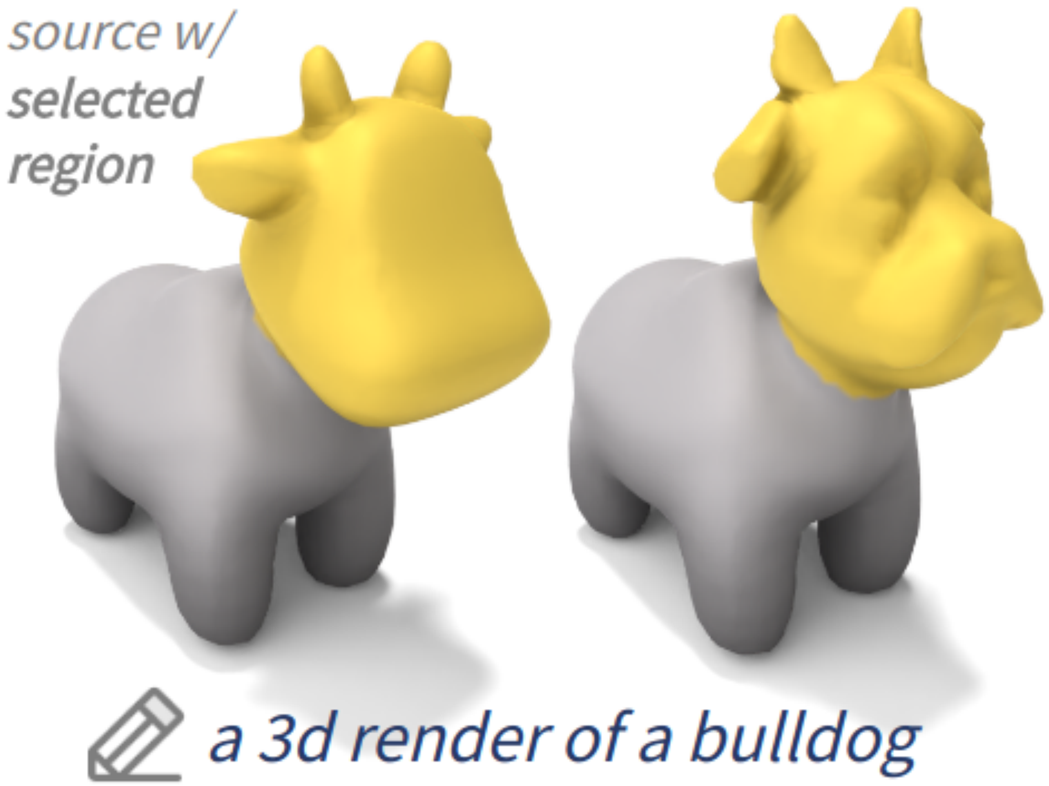

Identity-preserving stylization of mesh geometry

Input 3D object

Deformed

Identity-preserving stylization of mesh geometry

Input 3D object

Deformed

Different shapes, same style

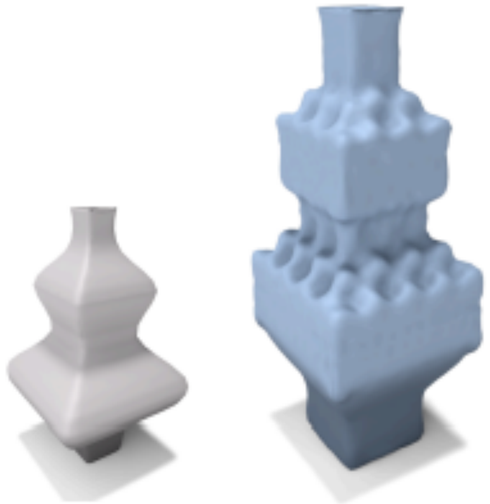

vase made of lego bricks

vase made of lego bricks

vase made of lego bricks

Same shape, different styles

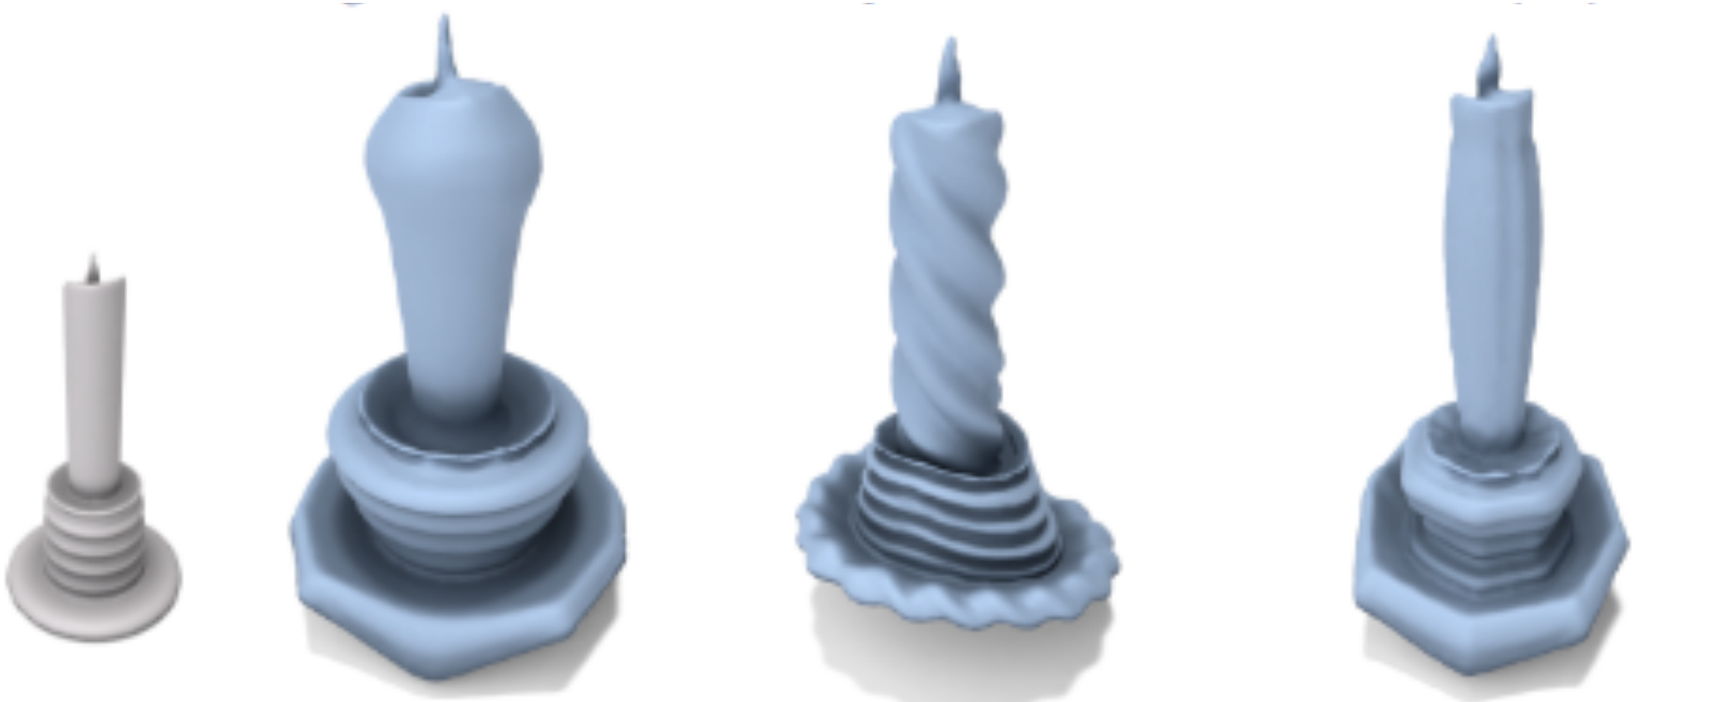

retro game joystick *braided pillar candle* *upright fountain pen*

Summary of Neural Deformations

- We can achieve highly expressive text-specified deformations
- Deformations via Jacobians are highly flexible (double edged sword)
- Deformations via surface normal is more restrictive, preserving the identity of the input

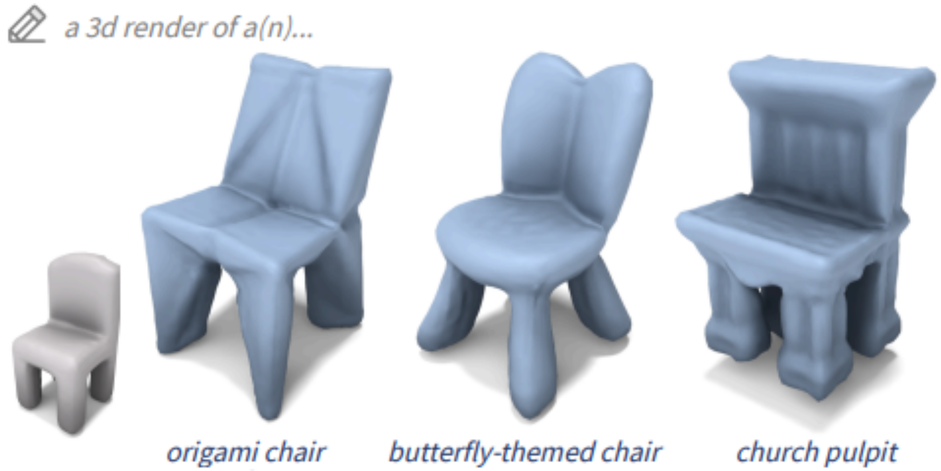

Neural Mesh Editing <u>without</u> 3D data!

Stylization

Text2Mesh [CVPR 2022]

Localization

3D Highlighter [CVPR 2023]
3D Paintbrush [CVPR 2024]

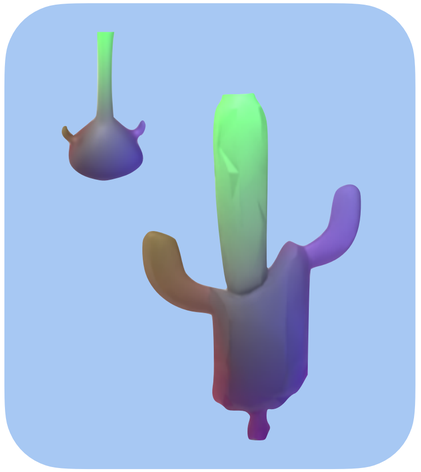

Deformation

TextDeformer [SIGGRAPH 2023]
MeshUp [3DV 2025]
Geometry in Style [CVPR 2025]

Segmentation

iSeg [SIGGRAPH Asia 2024]

The future of mesh editing without 3D datasets

More tasks in geometry processing

What other underlying properties can we extract during synthesis?

Use these methods to generate supervised data to bootstrap training feedforward networks

3DL @ UChicago

"Threedle"

Computer Science Building @ UChicago

UChicago Campus

City of Chicago

Thank you!

Rana Hanocka
PI

Dale Decatur
PhD Student

Richard Liu
PhD Student

Nam Anh Dinh
PhD Student

Itai Lang
Postdoctorate

Brian Kim
Undergraduate

a chair made of stained glass

THE UNIVERSITY OF CHICAGO

Links to code & papers available on our website!
https://3dl.cs.uchicago.edu/