Neural Shape Processing

Vladimir (Vova) Kim

Adobe Research, Seattle

Motivation

3D modeling of **high-quality** content that is

- Diverse and unique
- Detailed

Artist-generated Model
[Iron Throne by Tornado Studio]

Motivation

3D modeling of **high-quality** content that is

- Diverse and unique – interpolation of training data
- Detailed -- coarse

Artist-generated Model
[Iron Throne by Tornado Studio]

Challenges with Neural Generation:

Poursaeed et al., ECCV 2020

Neural Shape Processing

Modify existing shapes instead of generating from scratch

Retrieval

Deformation

Detailization

Shape Retrieval

Most prior techniques focus on finding geometrically-similar shape

Deformation-Aware Shape Retrieval

Retrieve shape that can be deformed to the query

Target

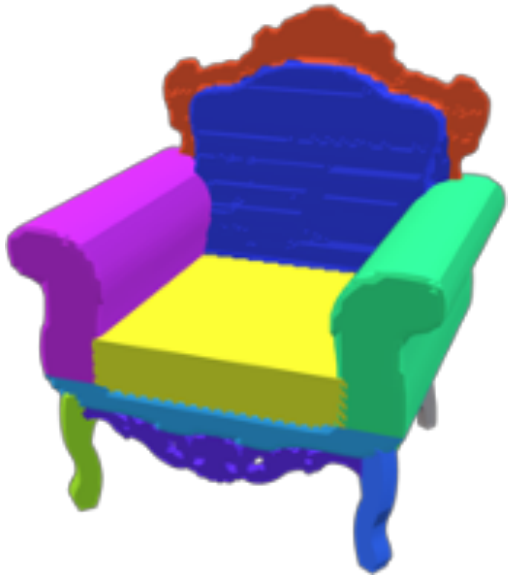

Most similar retrieval

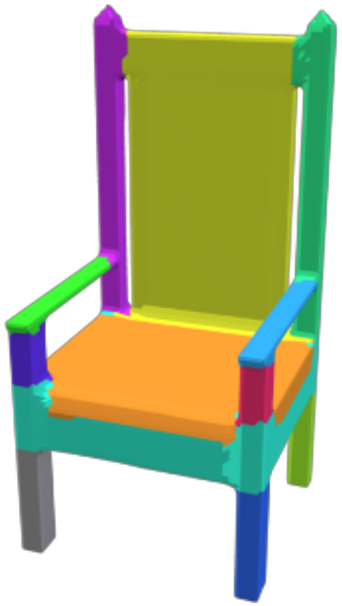

Deformation-aware retrieval

Deformation-Aware Shape Retrieval

Retrieve shape that can be deformed to the query

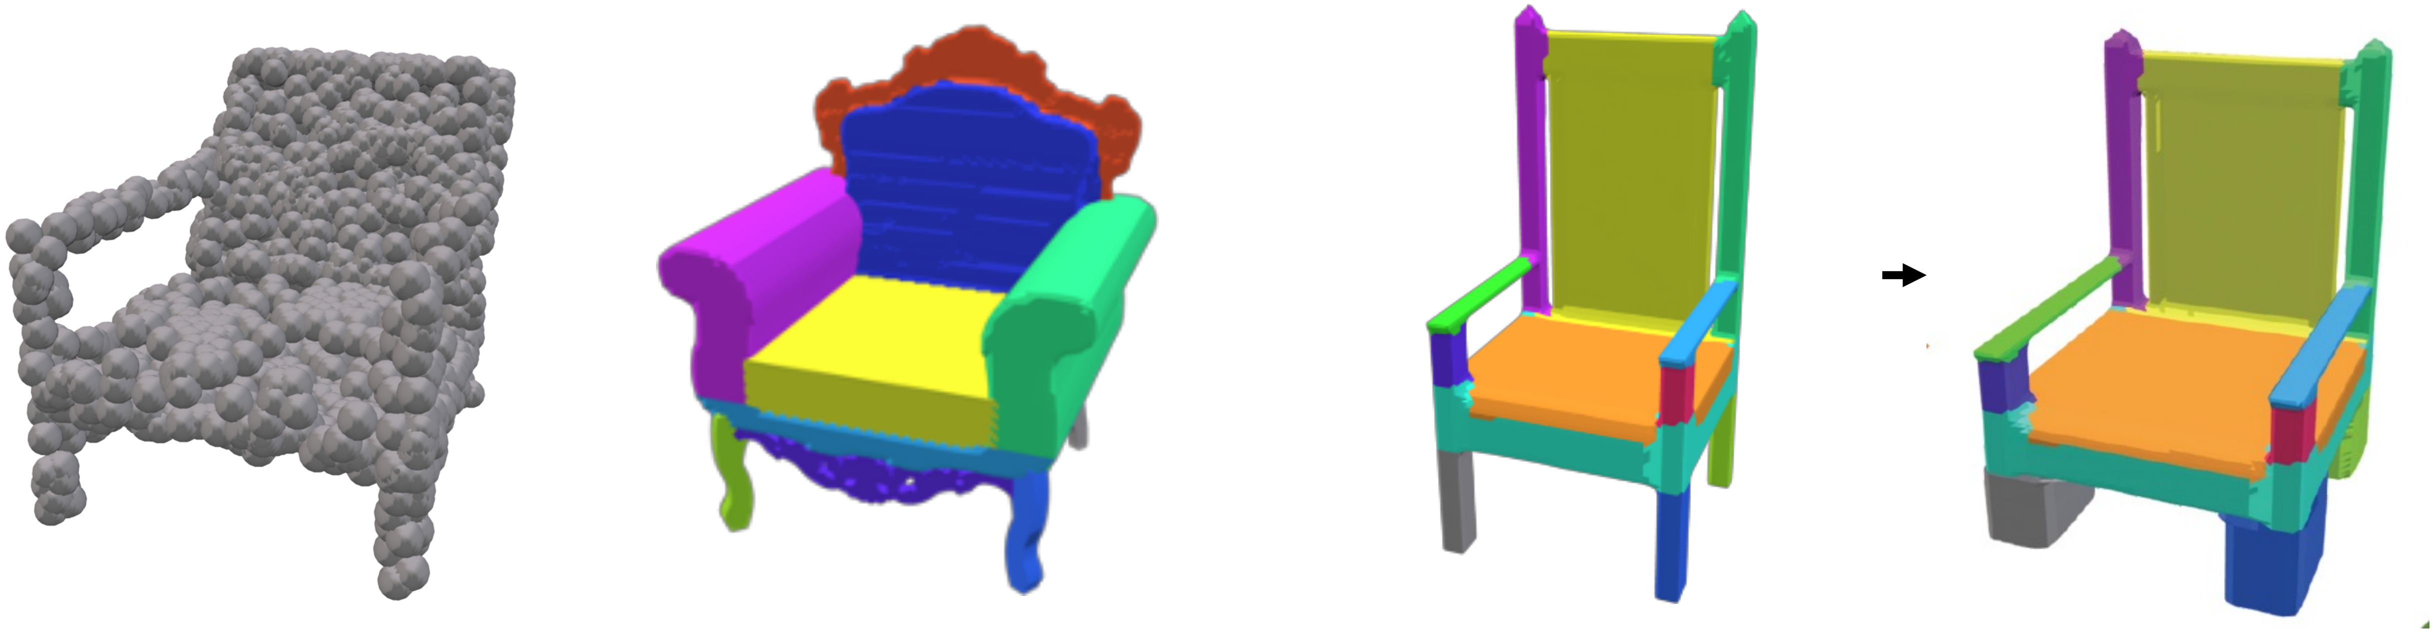

Target | Most similar retrieval | Deformation-aware retrieval | Deformed

Structure-aware Neural Deformation

Parameterize deformations based on part structure

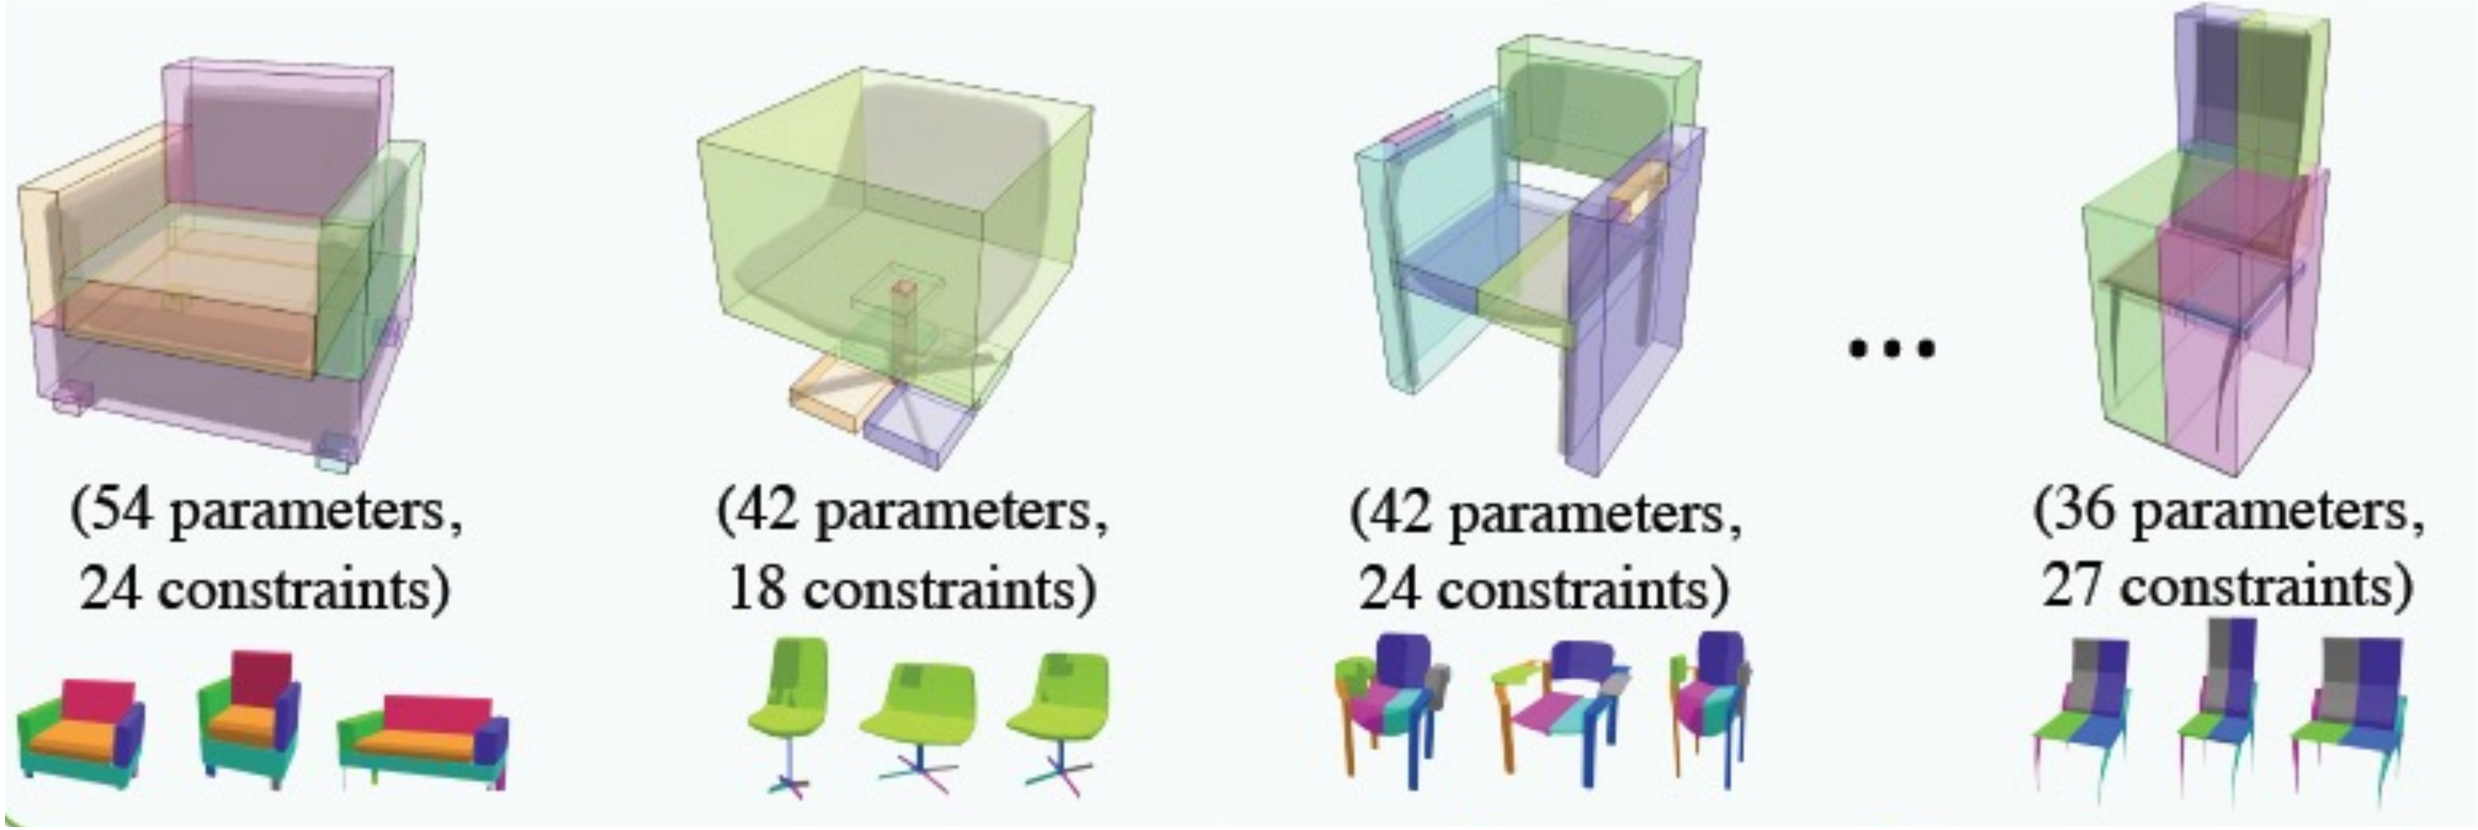

Joint Retrieval and Deformation Training

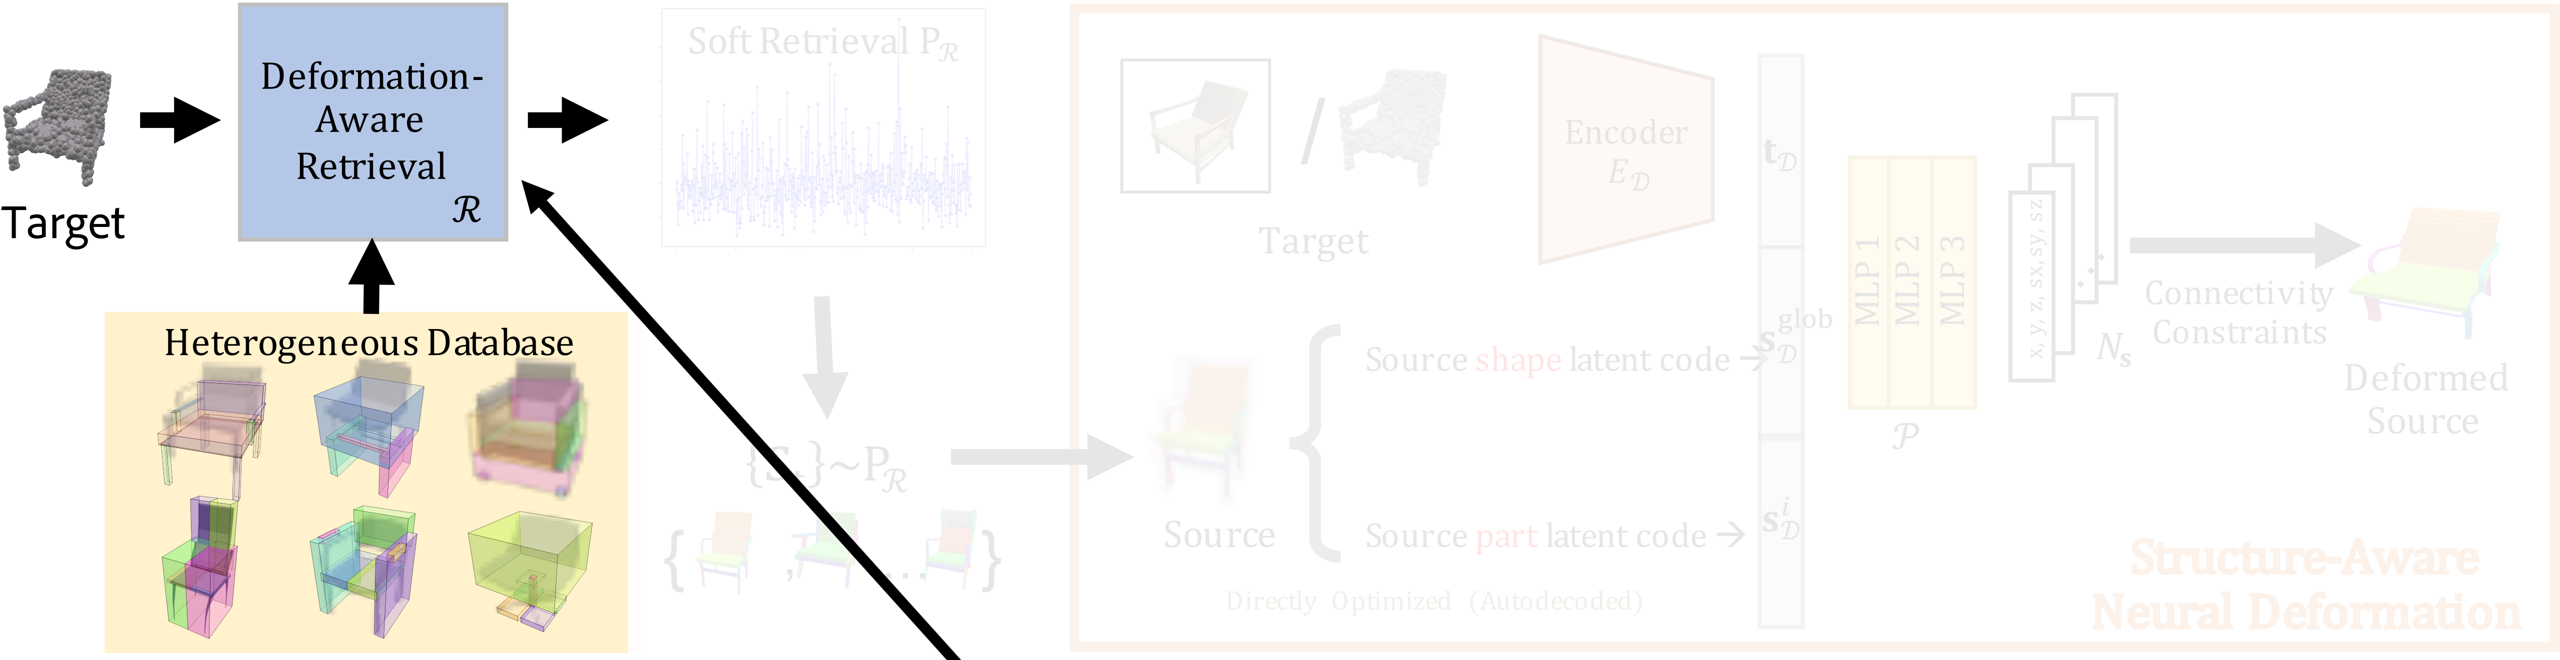

Neural Encoder

Fixed database
(codes are learned)

Joint Retrieval and Deformation Training

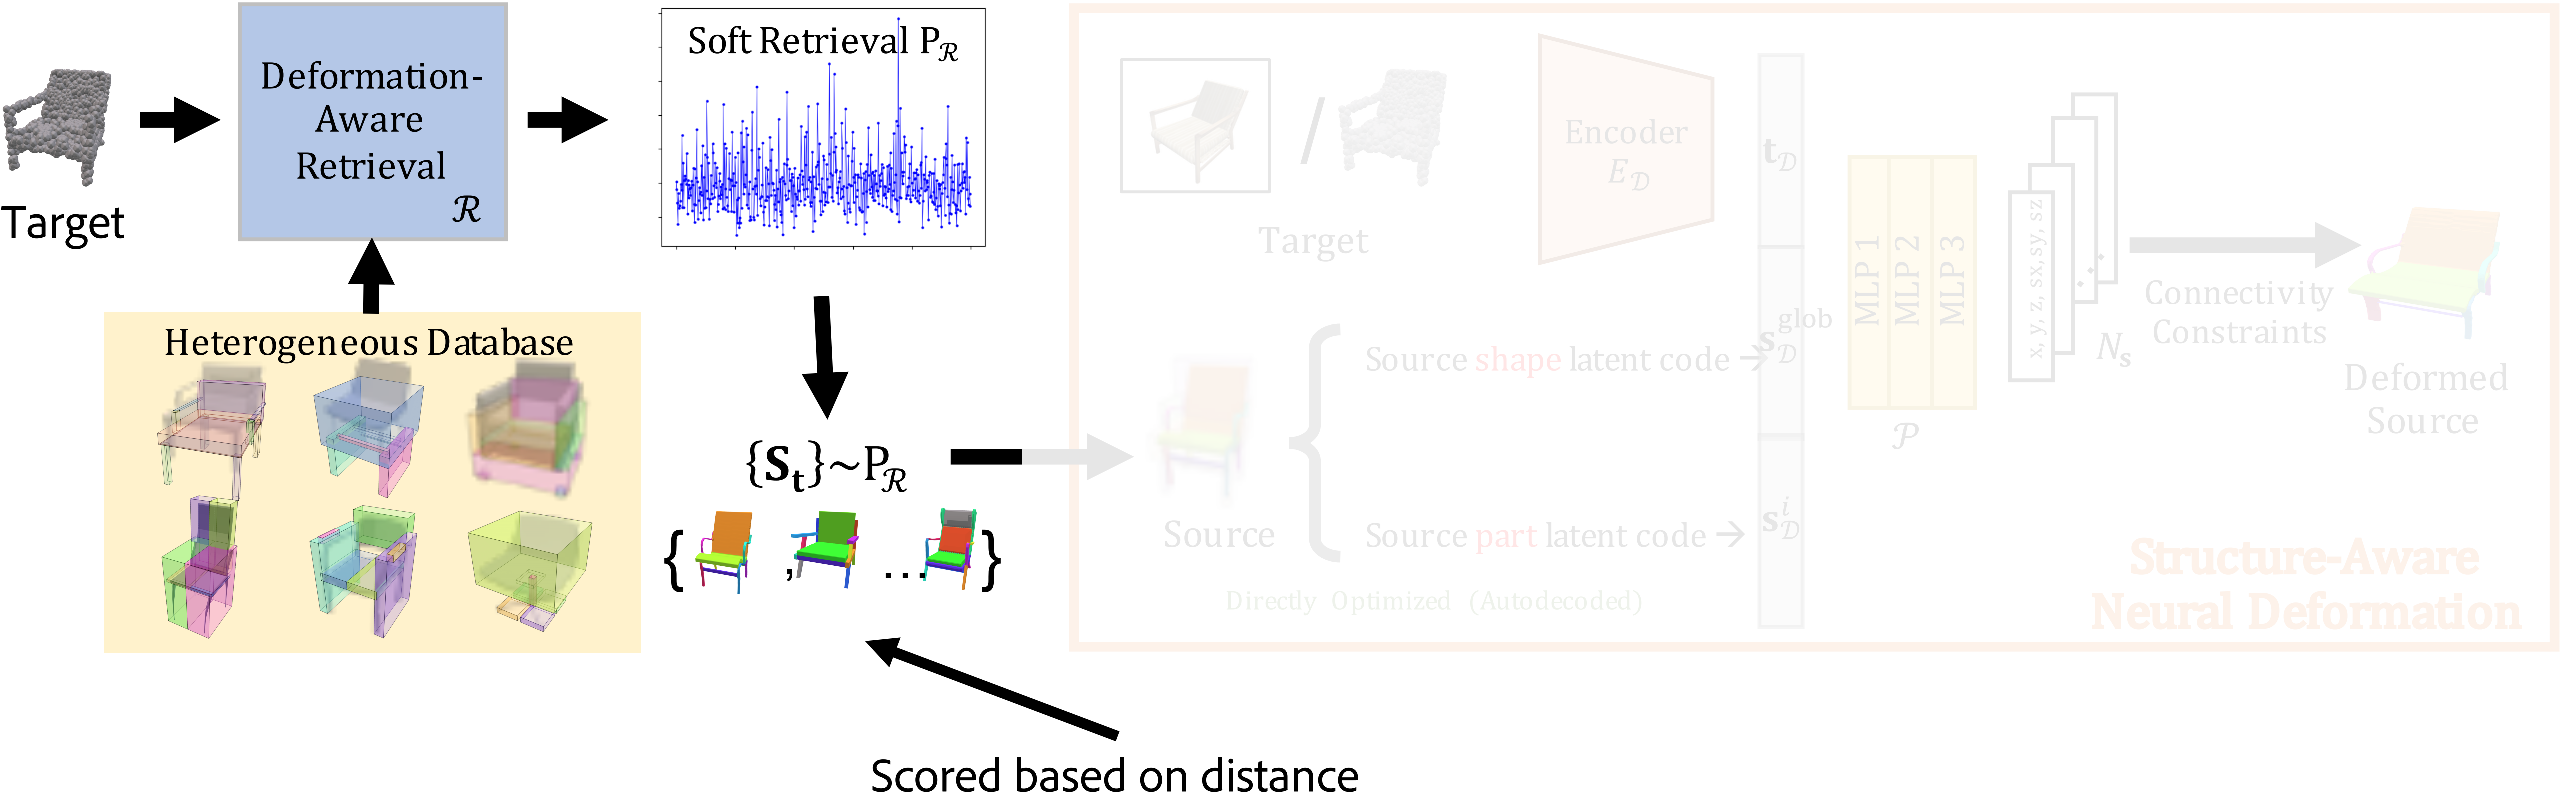

Joint Retrieval and Deformation Training

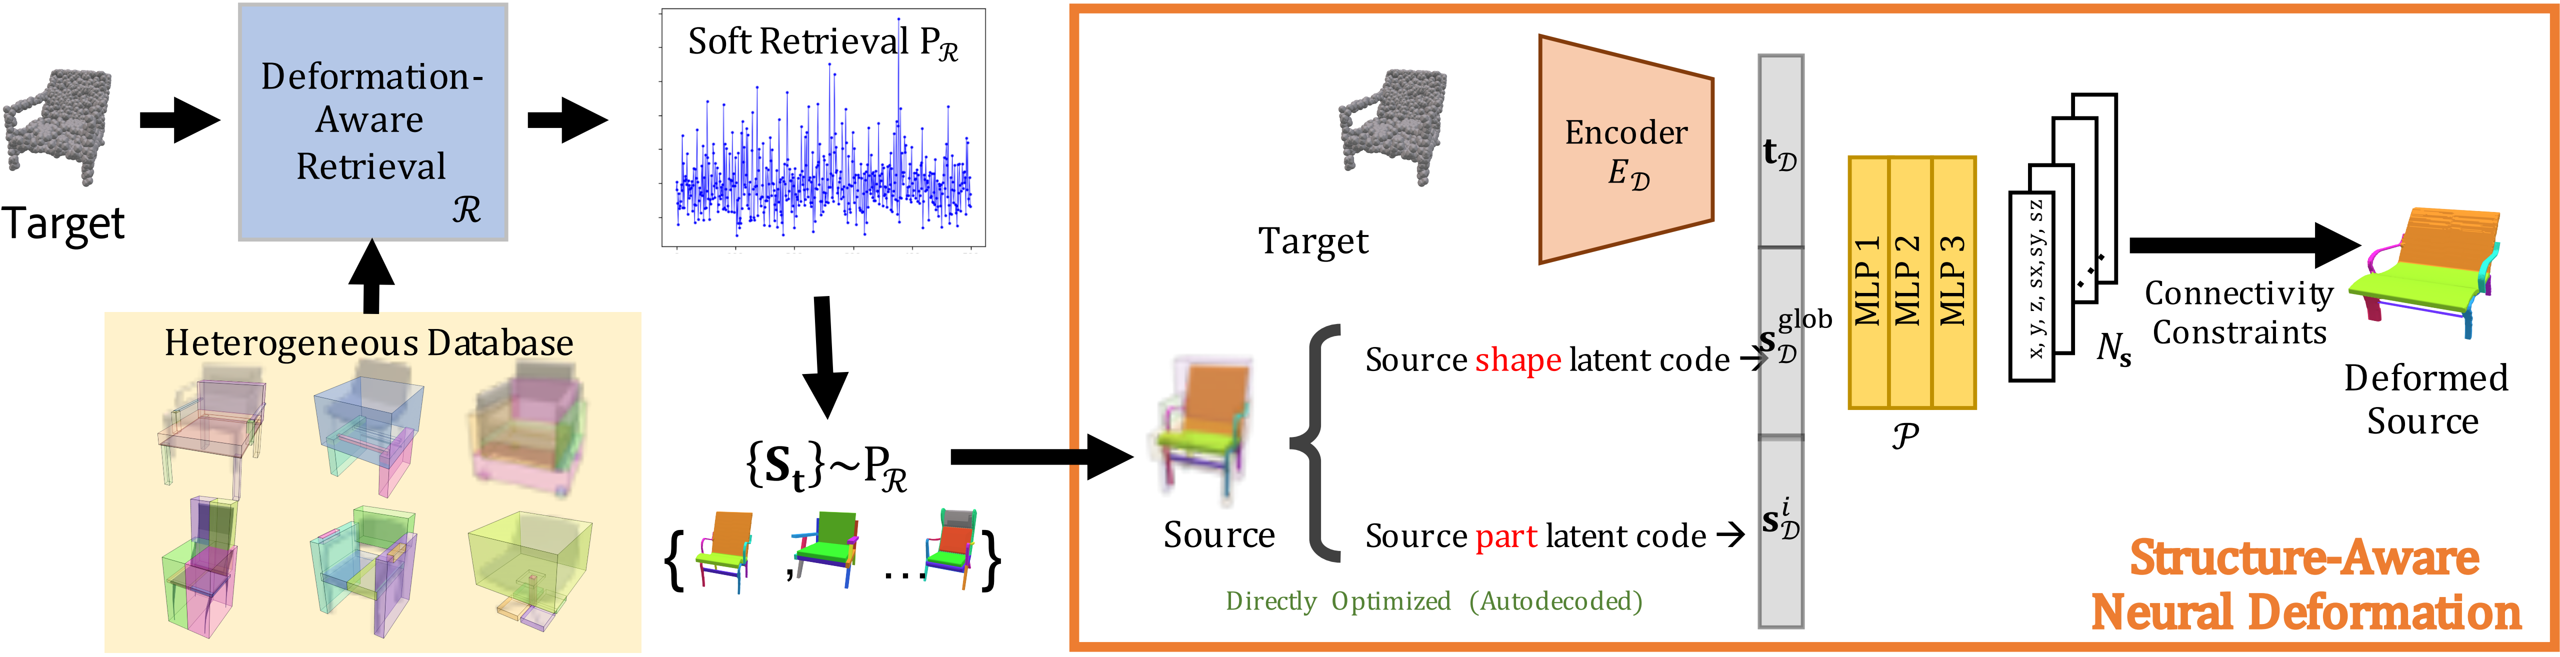

Deform top K sources to the target
each part has a learn-able code

Example Retrieval and Deformation from Scans

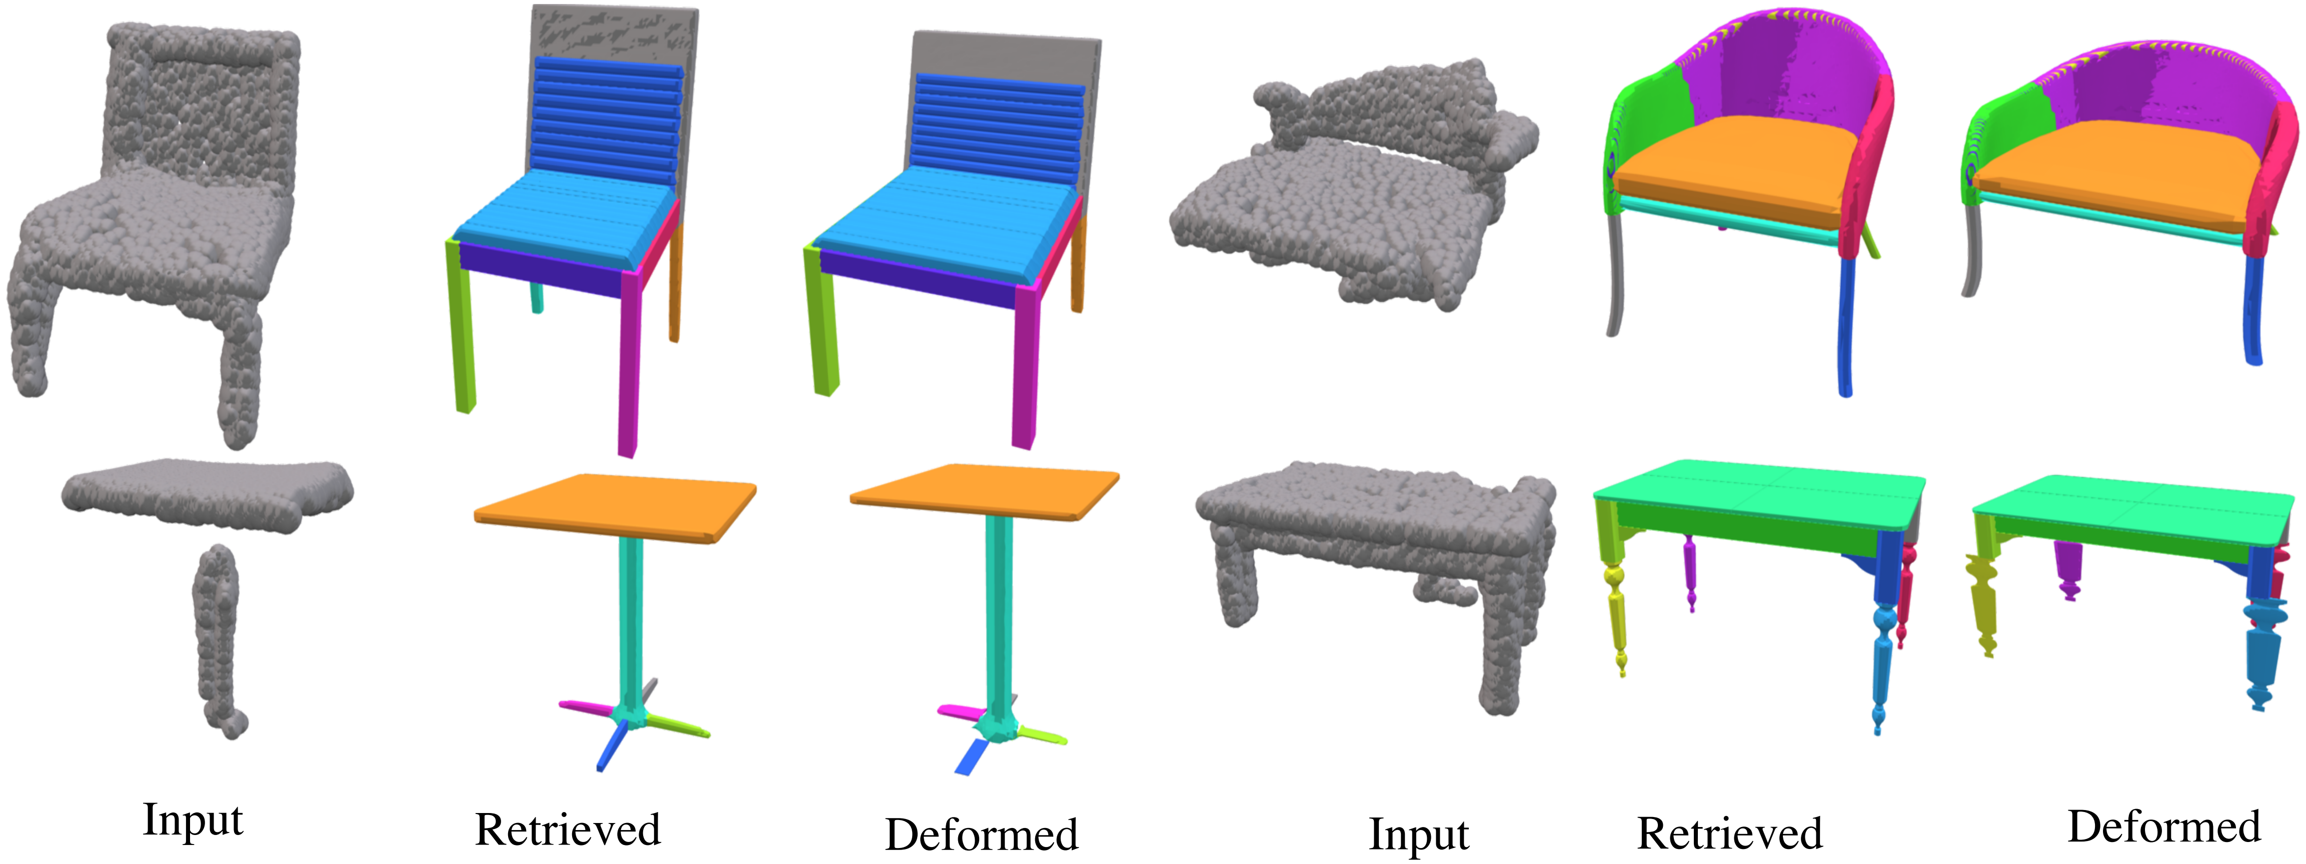

Example Retrieval and Deformation from Scans

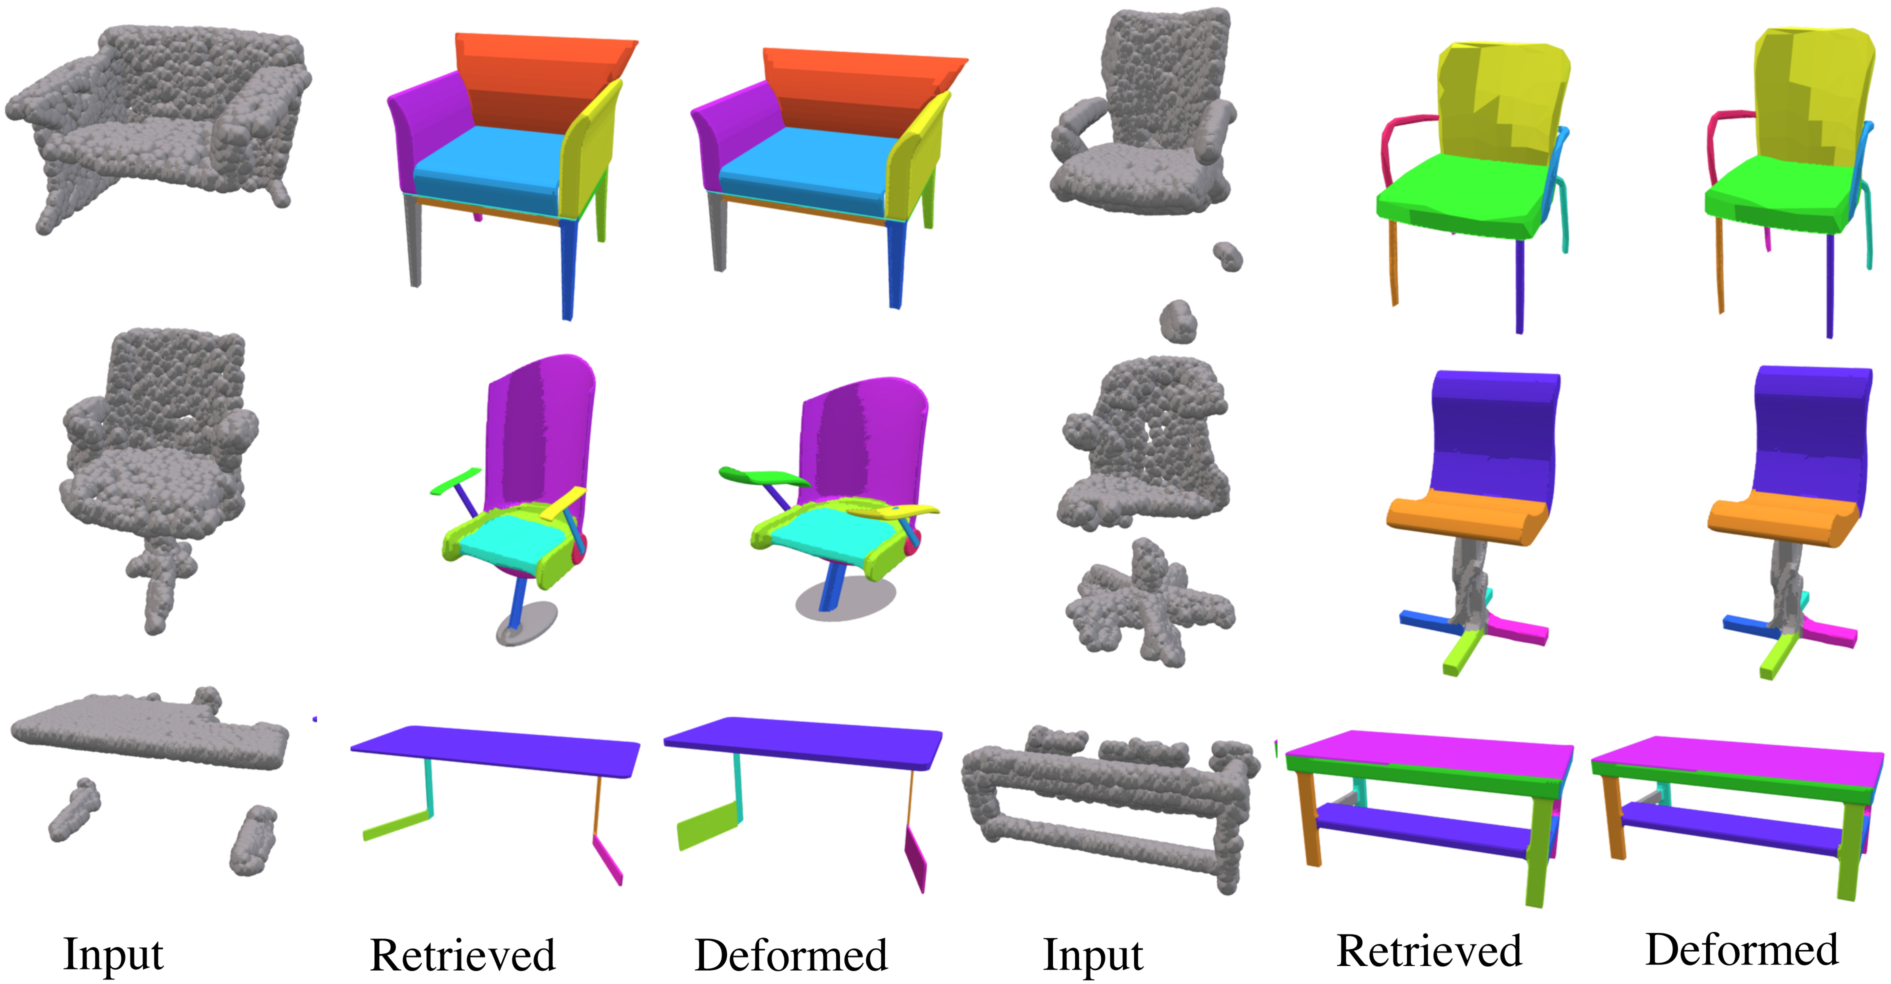

Example Retrieval and Deformation from an Image

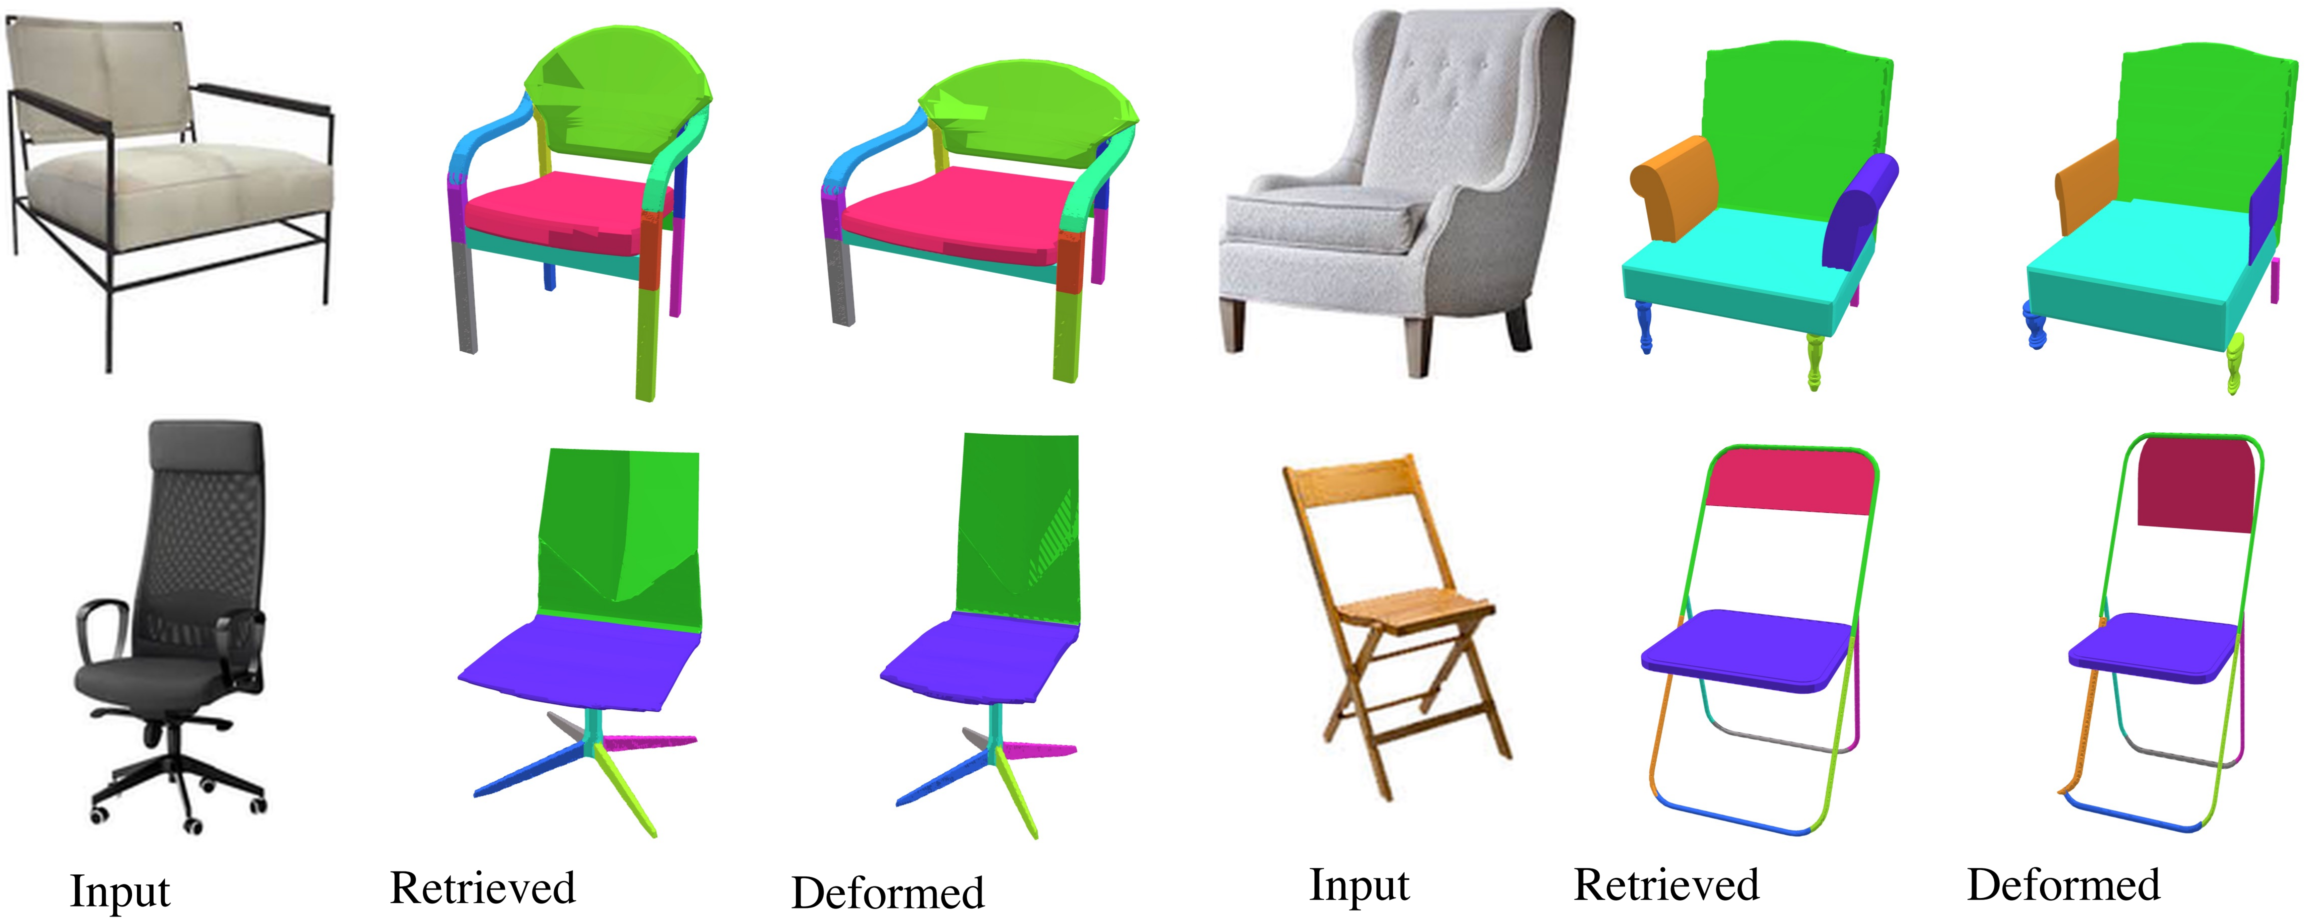

Example Retrieval and Deformation from an Image

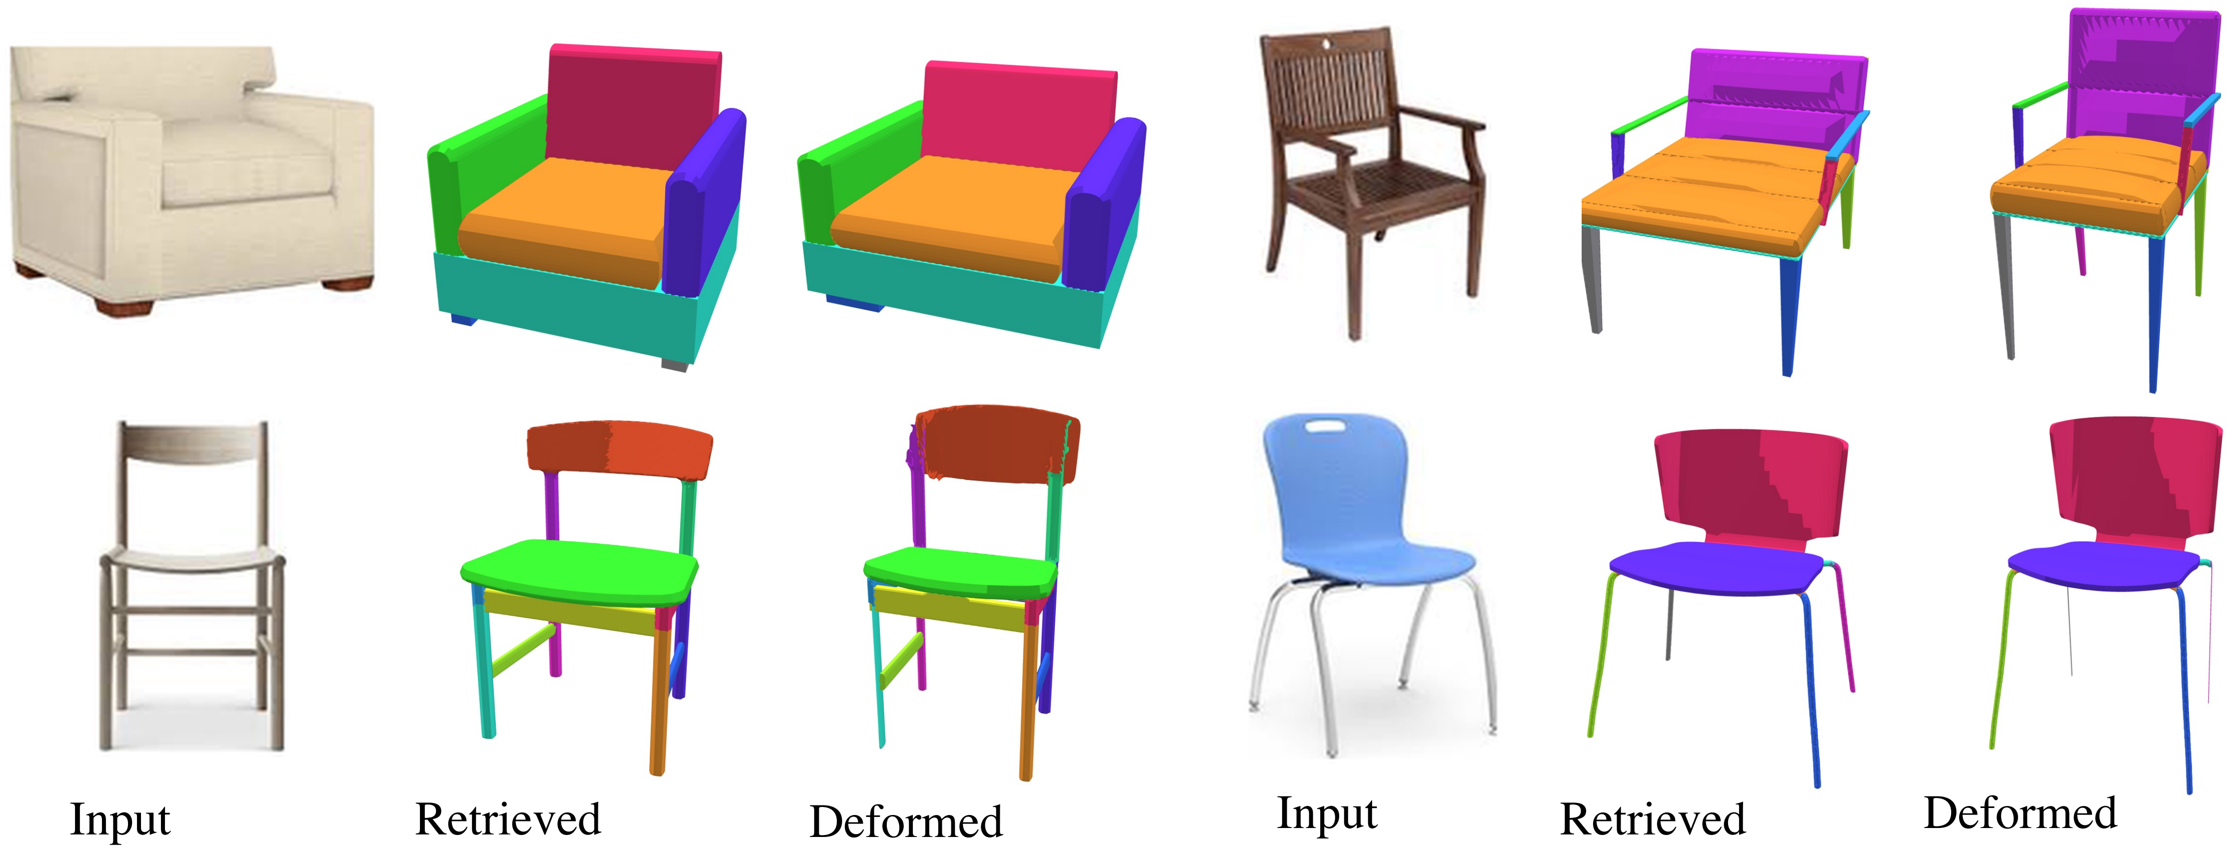

Comparisons

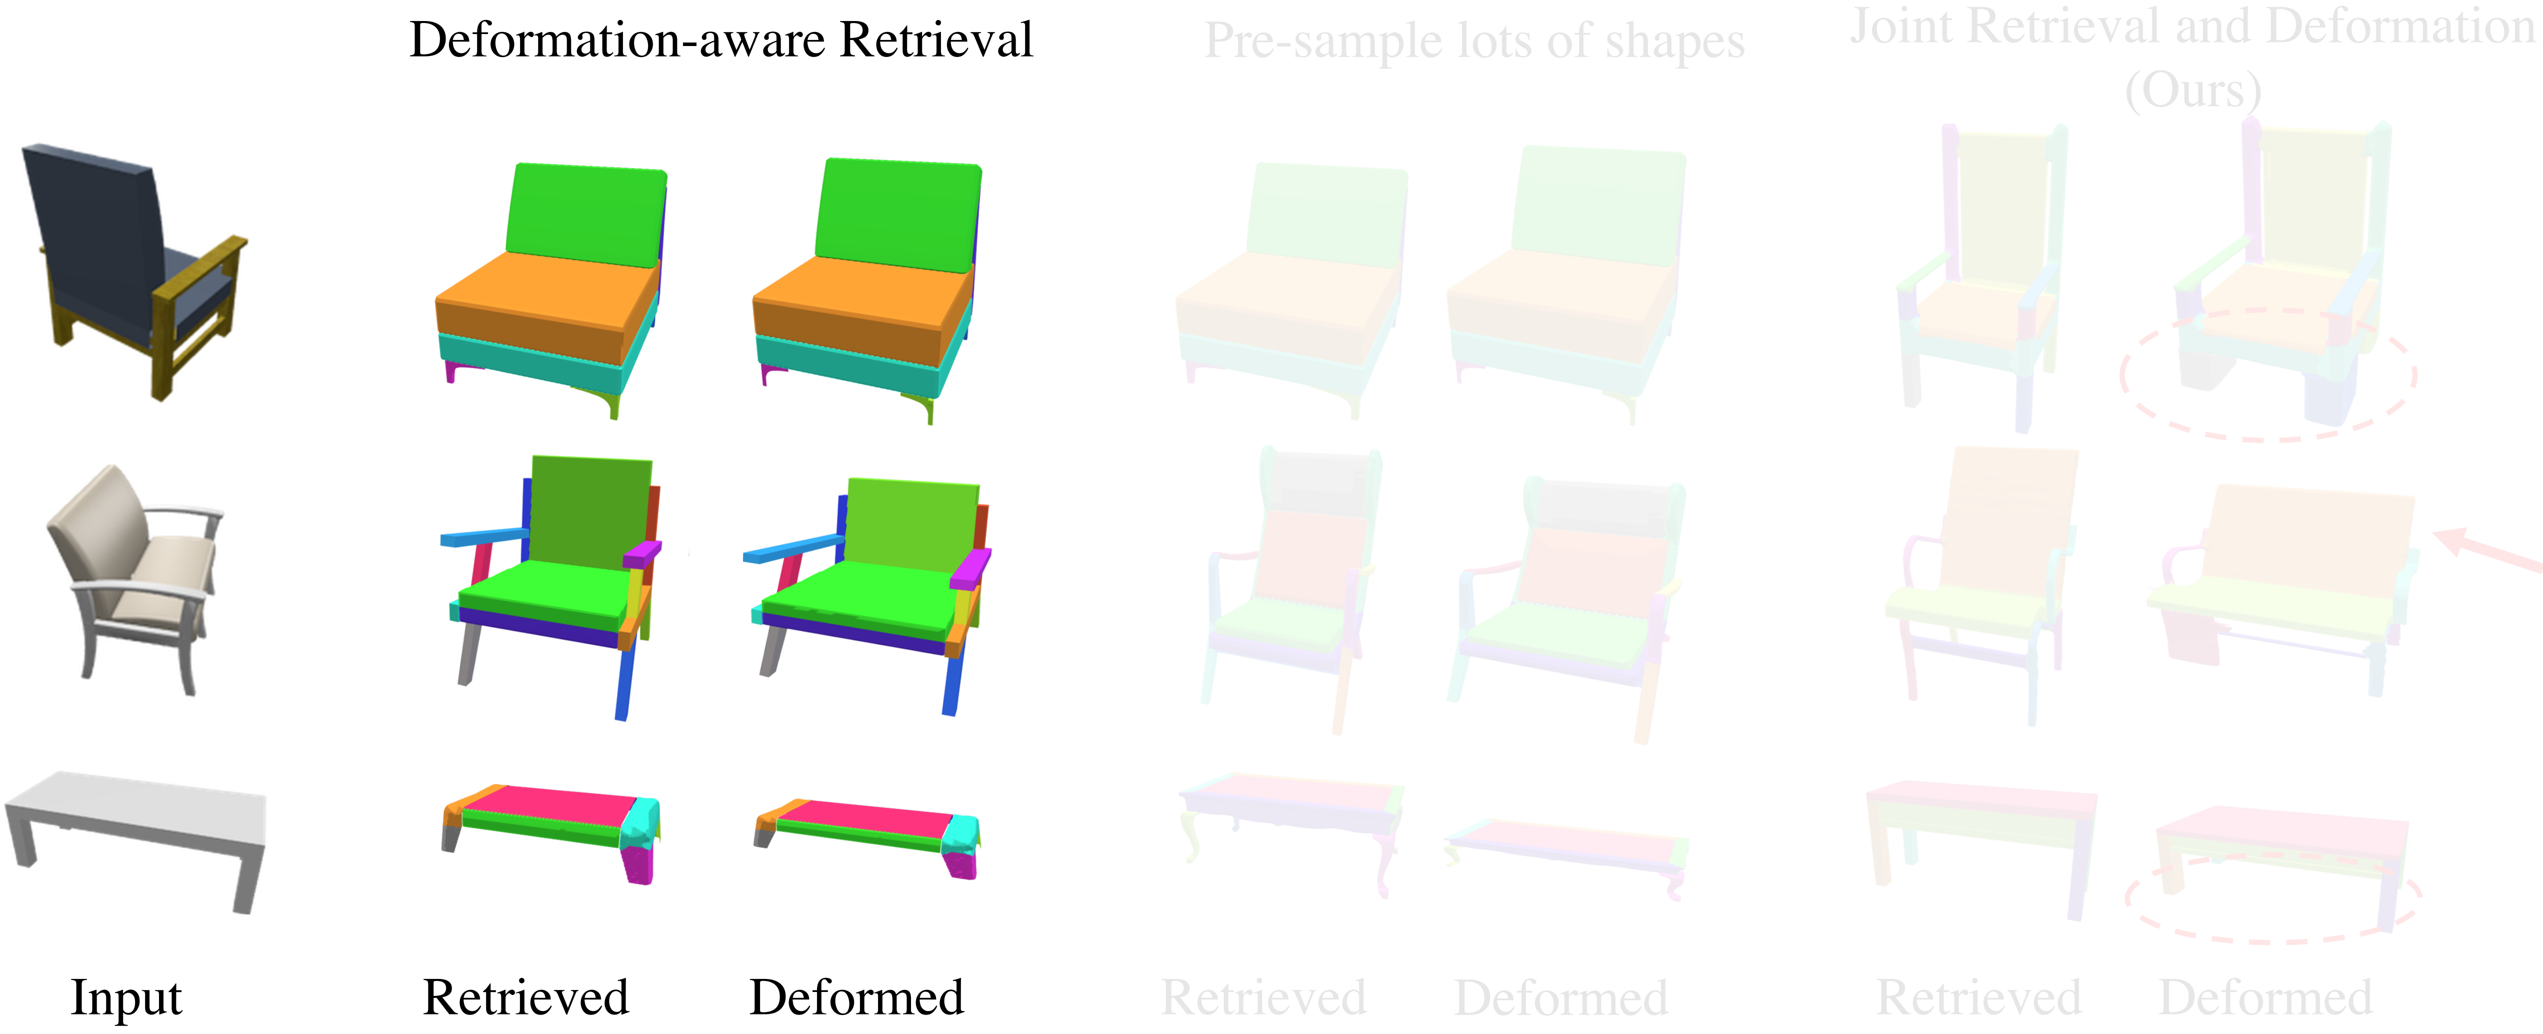

Comparisons

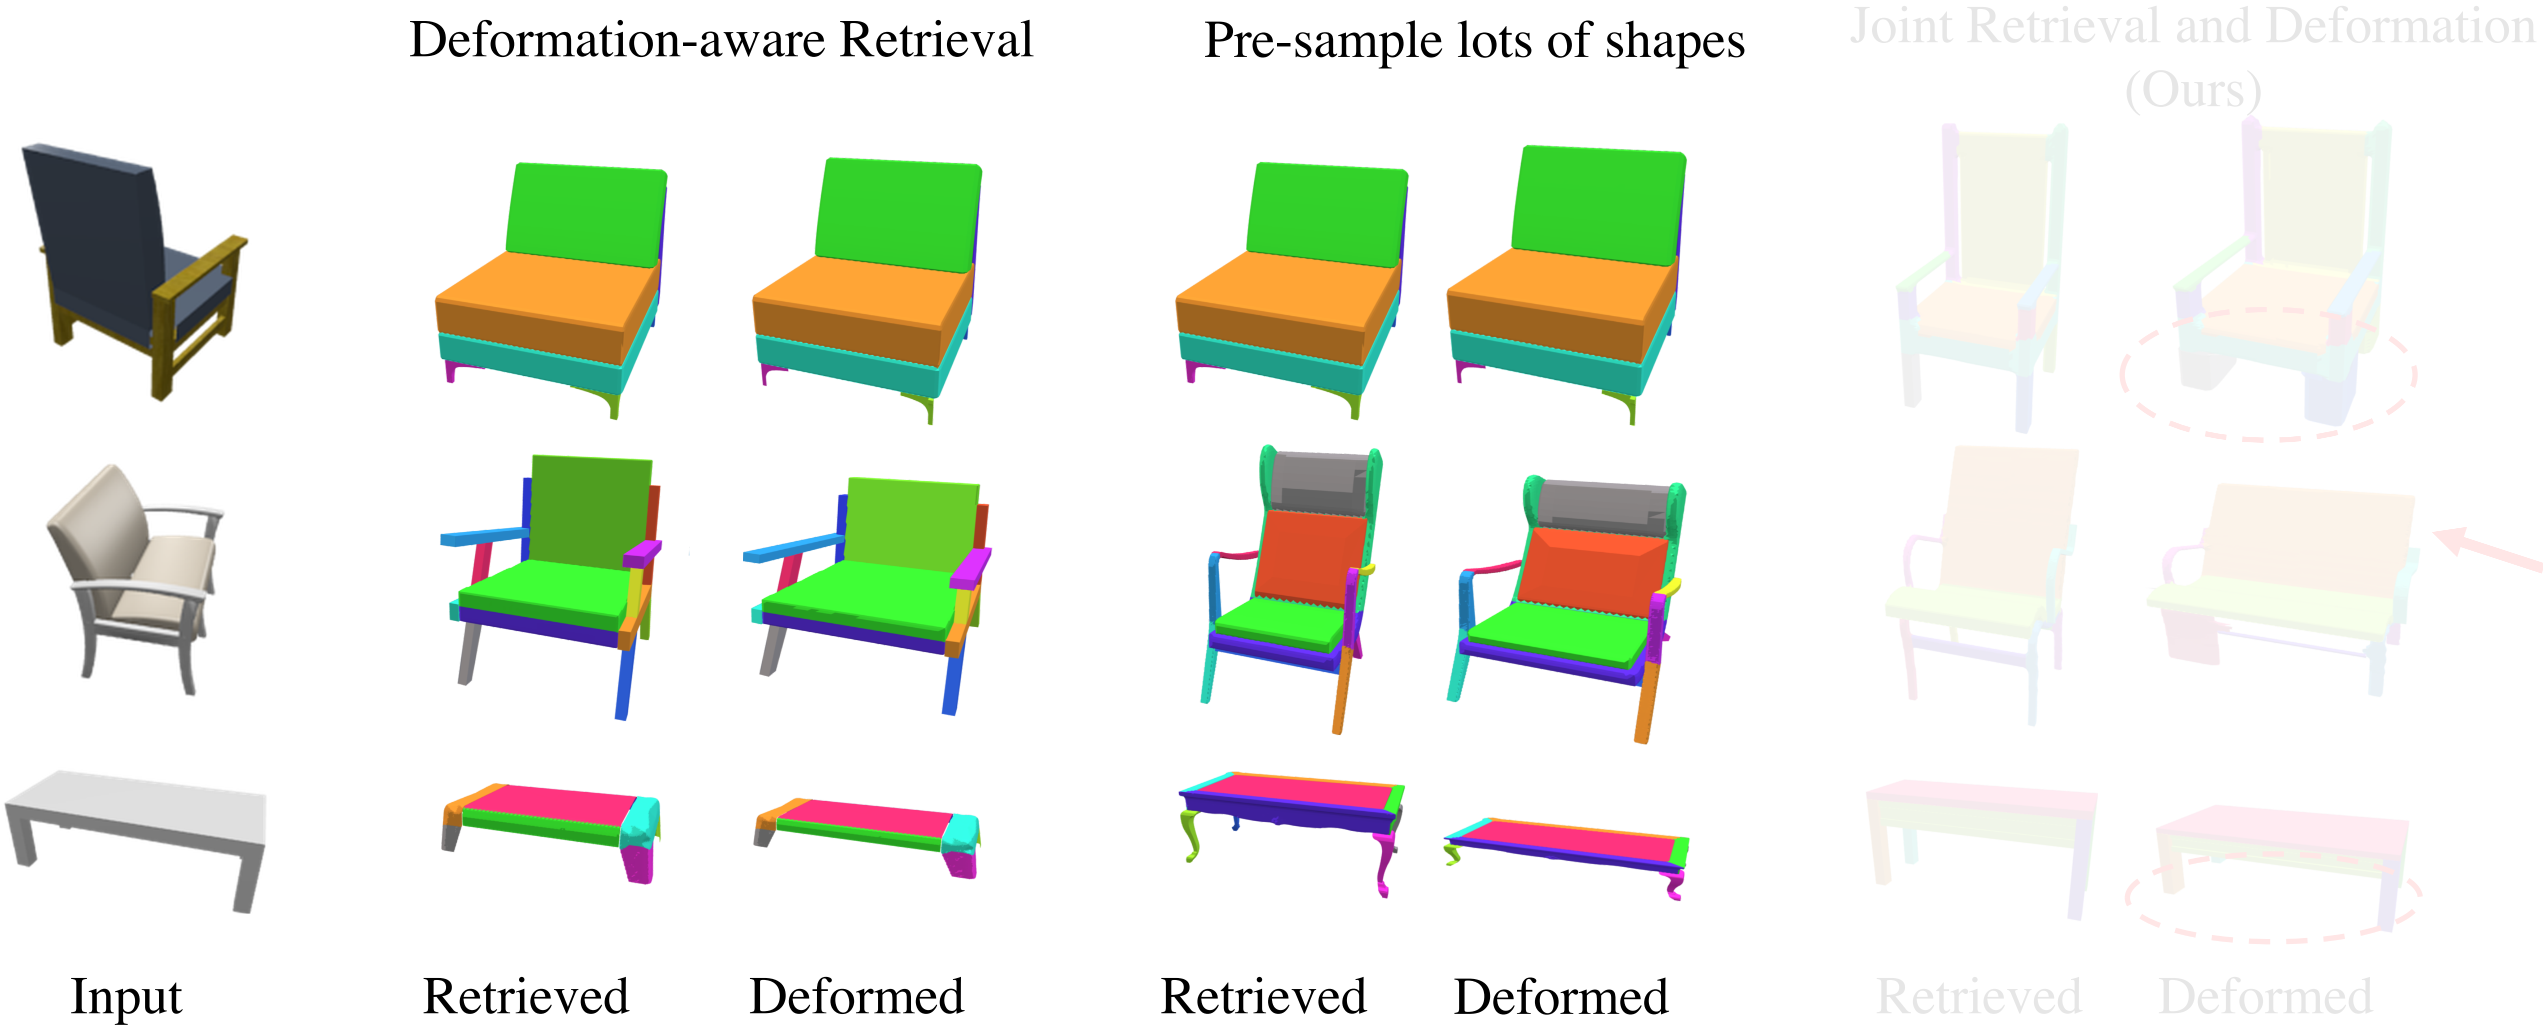

Comparisons

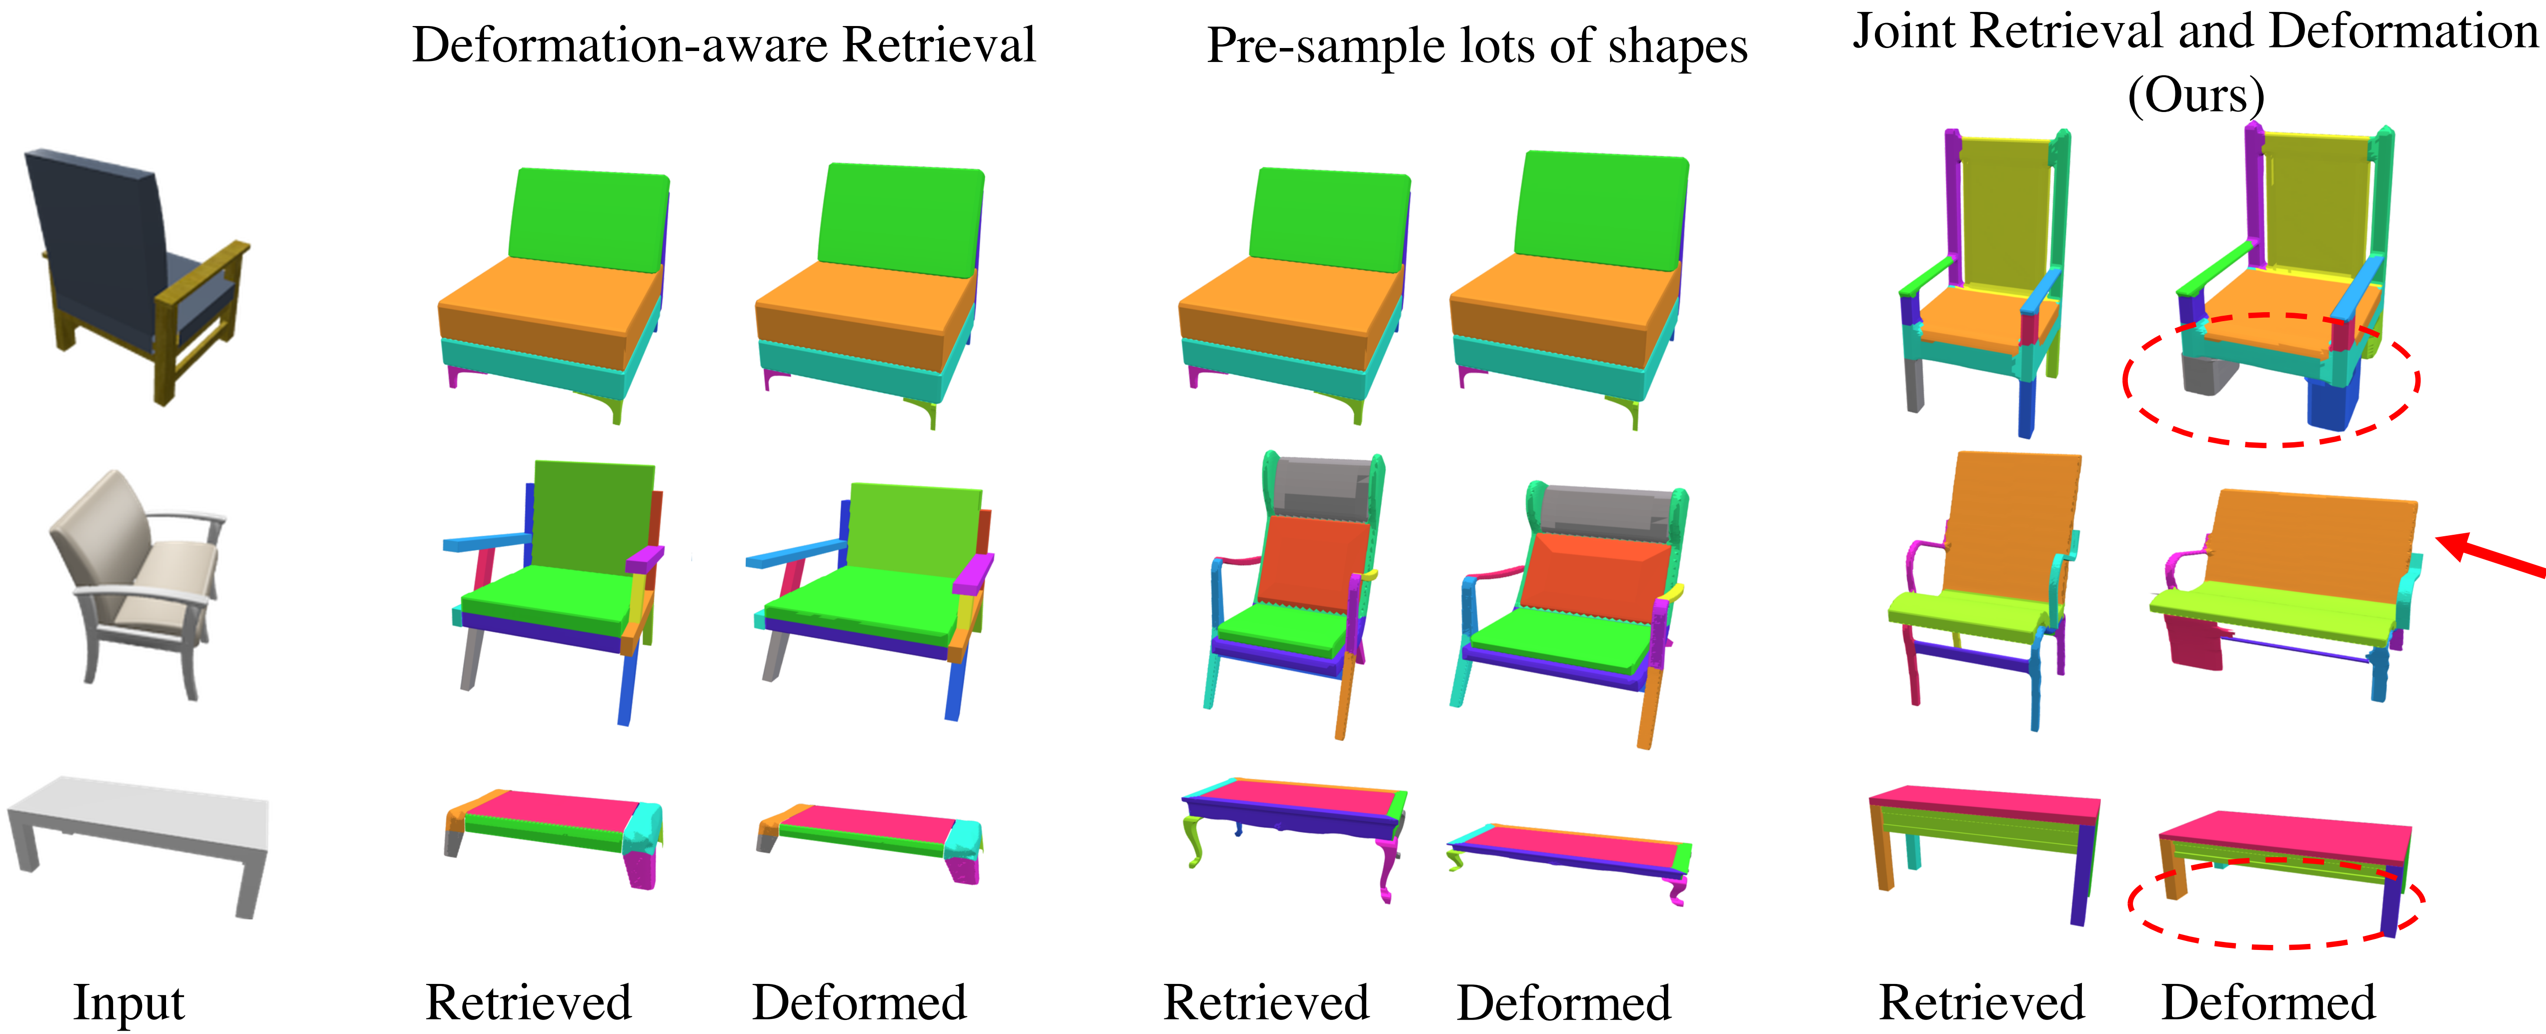

Comparisons

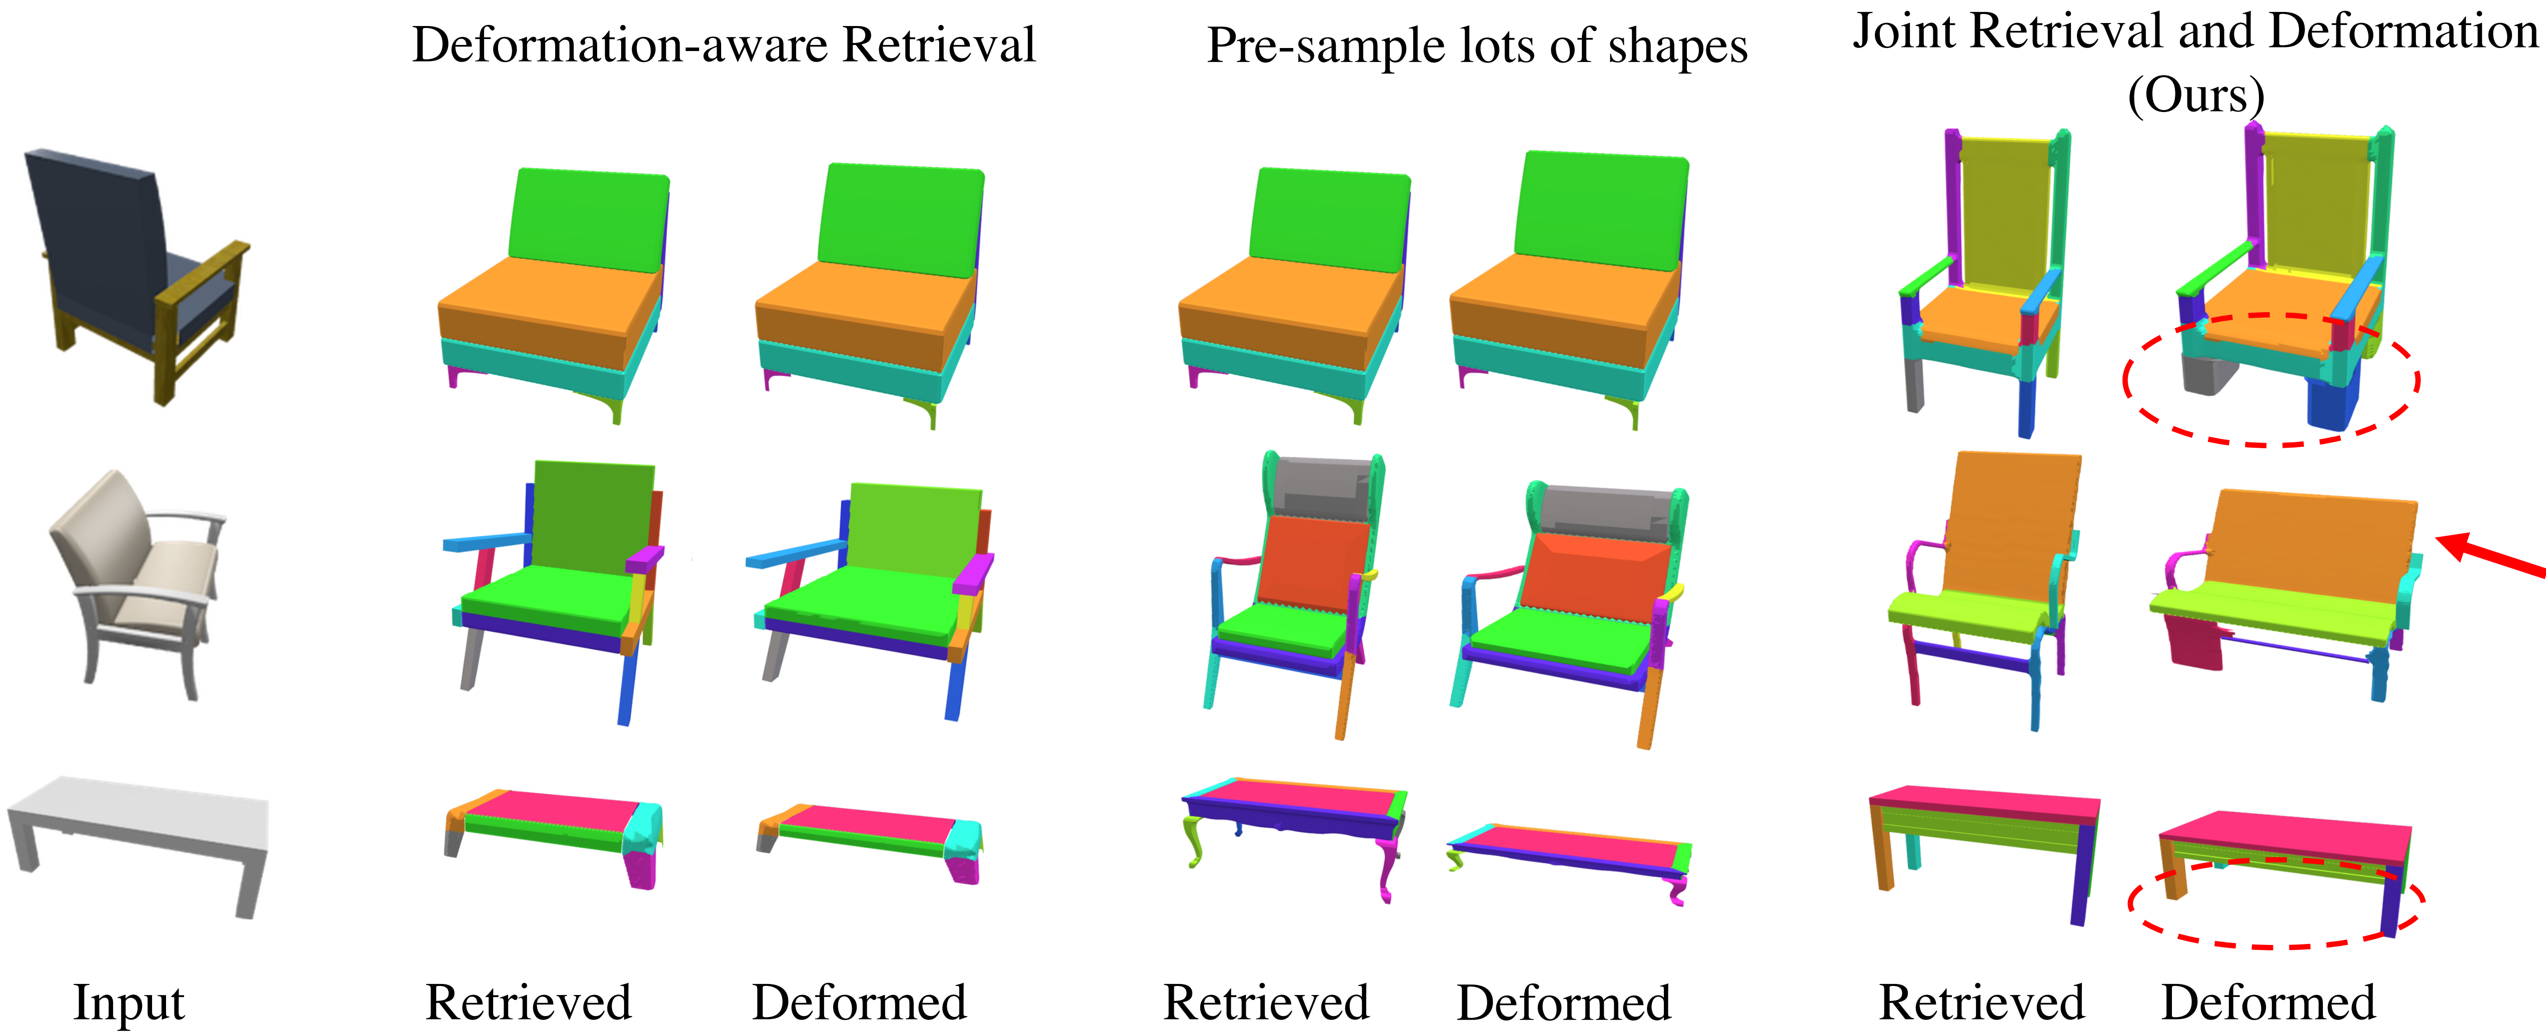

Connected components instead of true segments

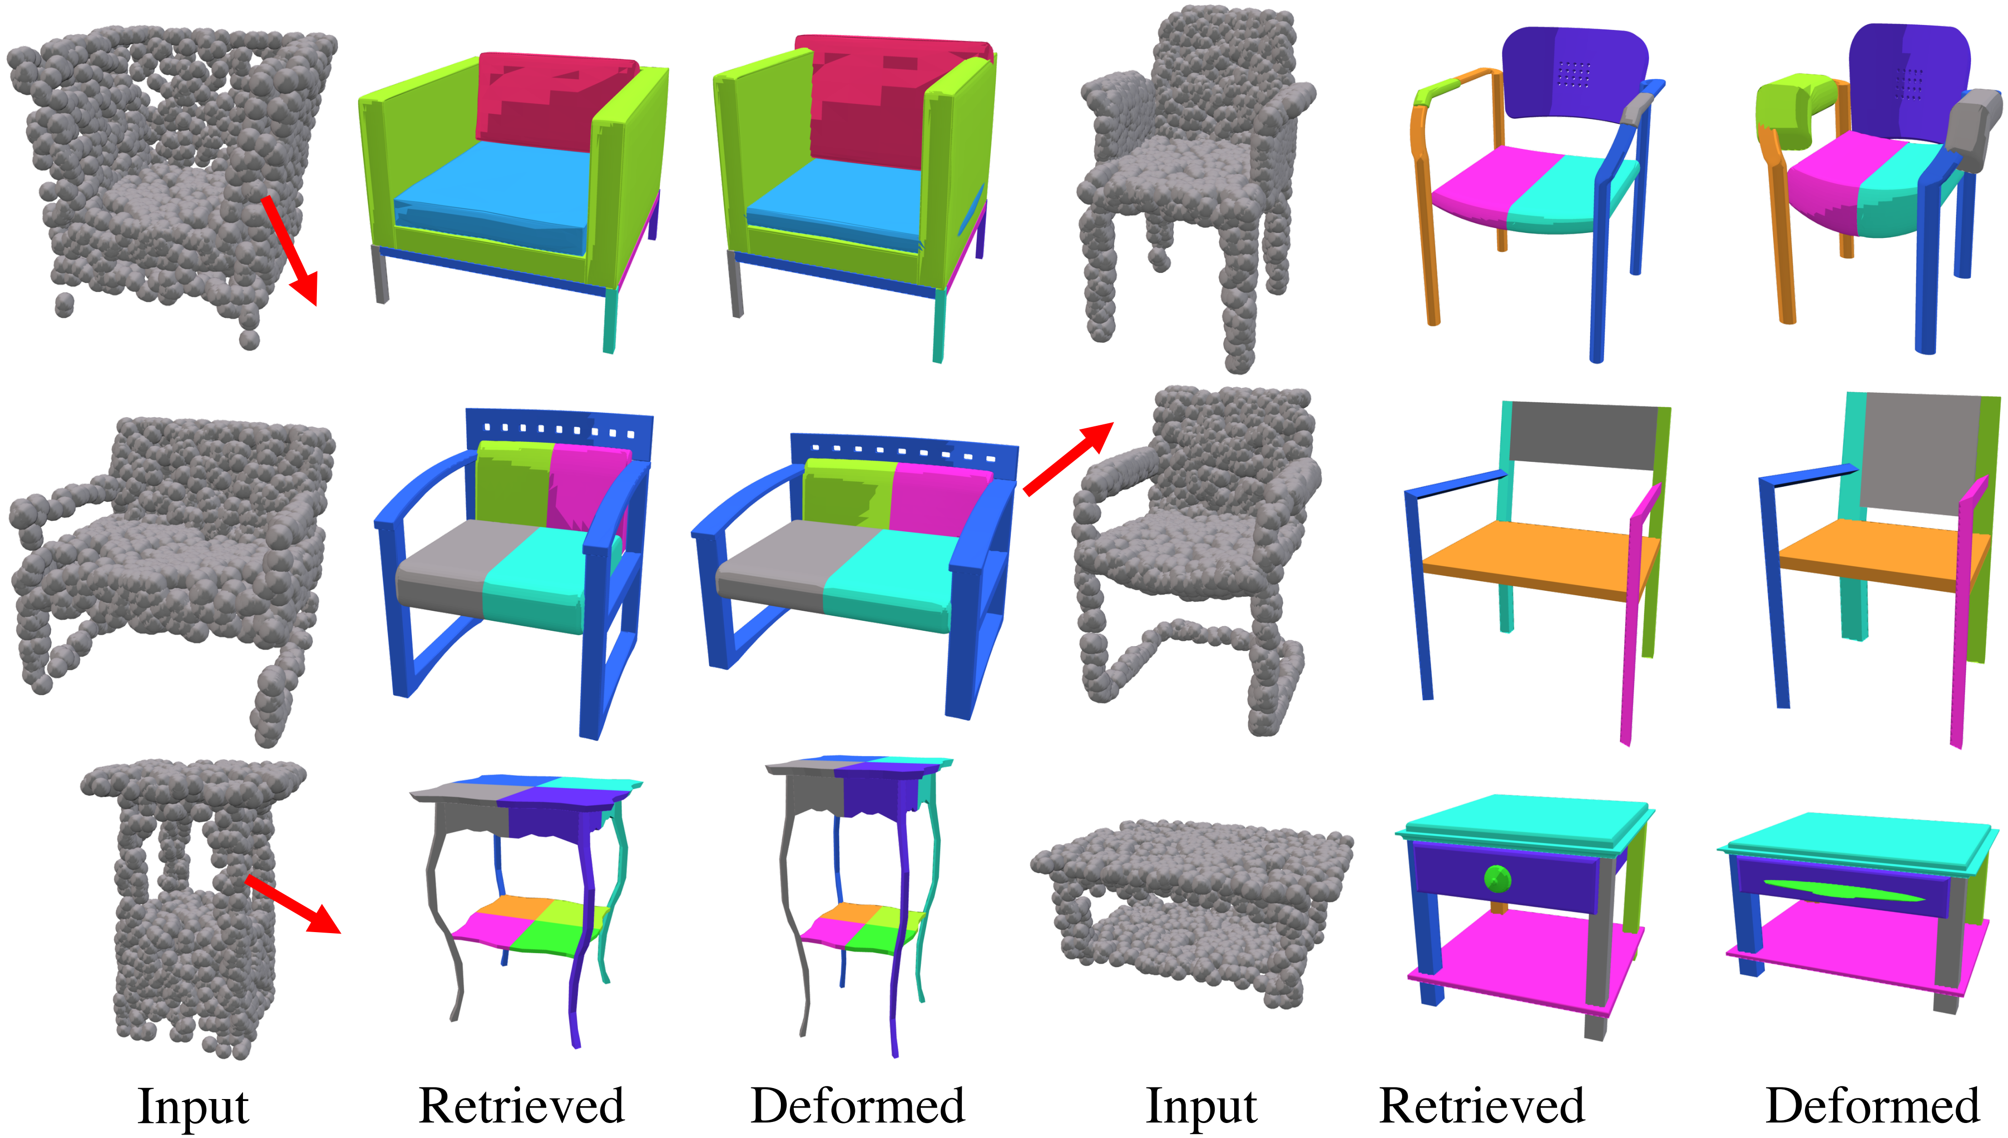

Neural Shape Processing

Modify existing shapes instead of generating from scratch

Retrieval

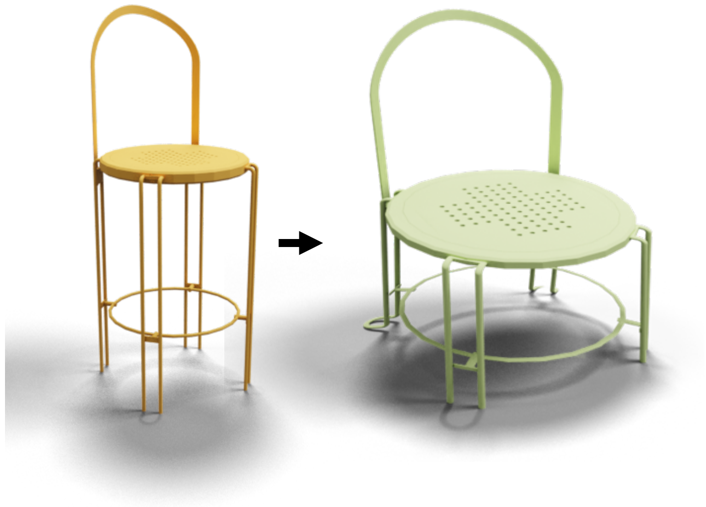

Deformation

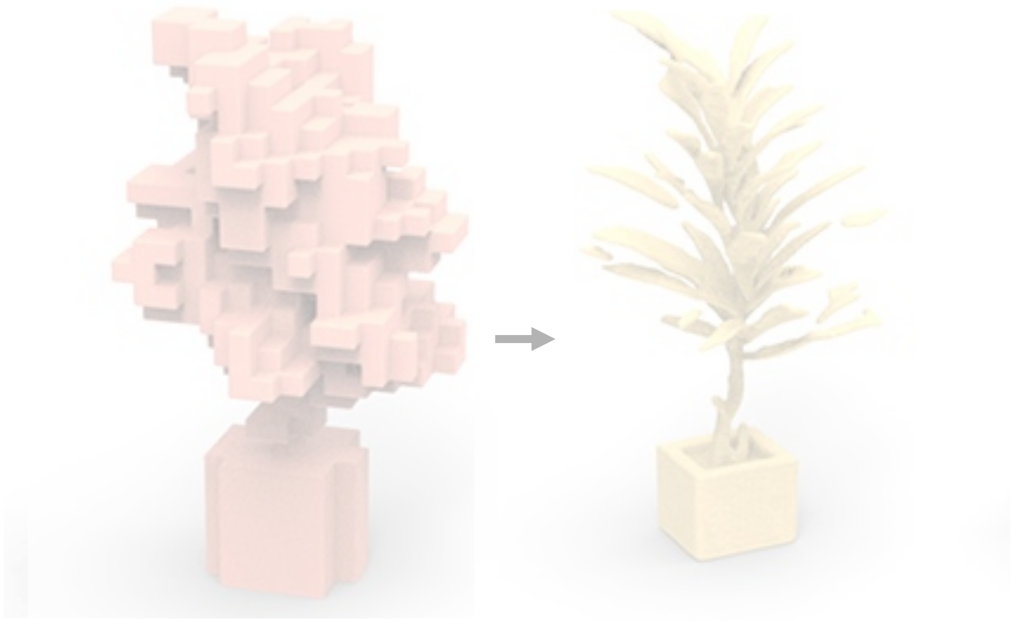

Detailization

Goal: Detail-Preserving Shape Deformation

Deform the source to match the target while preserving the details

Source Mesh

Target Shape

Deformed Source Mesh

Limitations of Direct Neural Deformation

$$f_\theta : z_{\text{source}} \times z_{\text{target}} \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$$

Source

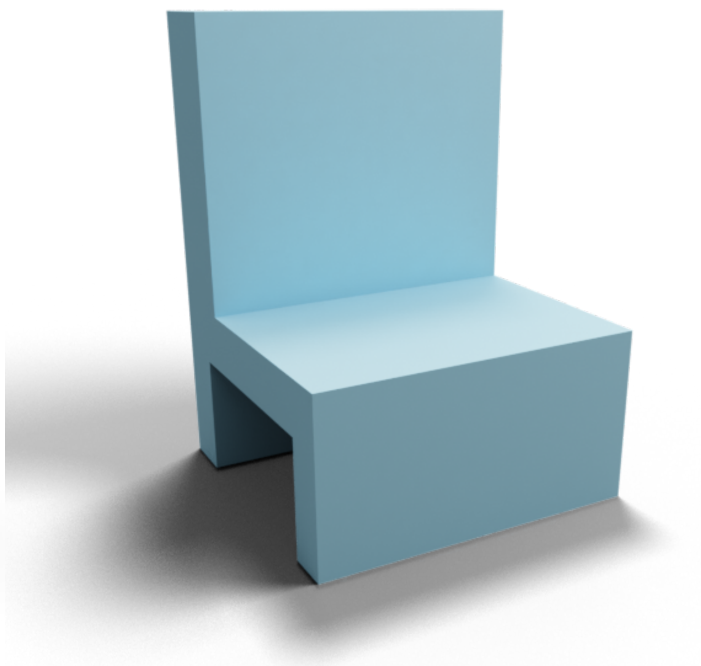

Target

Groueix et al. CGF 2019

Details are not preserved:
- too many degrees of freedom
- regularization energy fights reconstruction

Neural **Cage-based** Deformation

- Step 1: Learn to predict cage parameters

$$f_\theta : z_{\text{source}} \times z_{\text{target}} \rightarrow \boxed{\mathcal{C}_{\text{init}} \times \mathcal{C}_{\text{deformed}}}$$

Predict cage parameters with a neural network

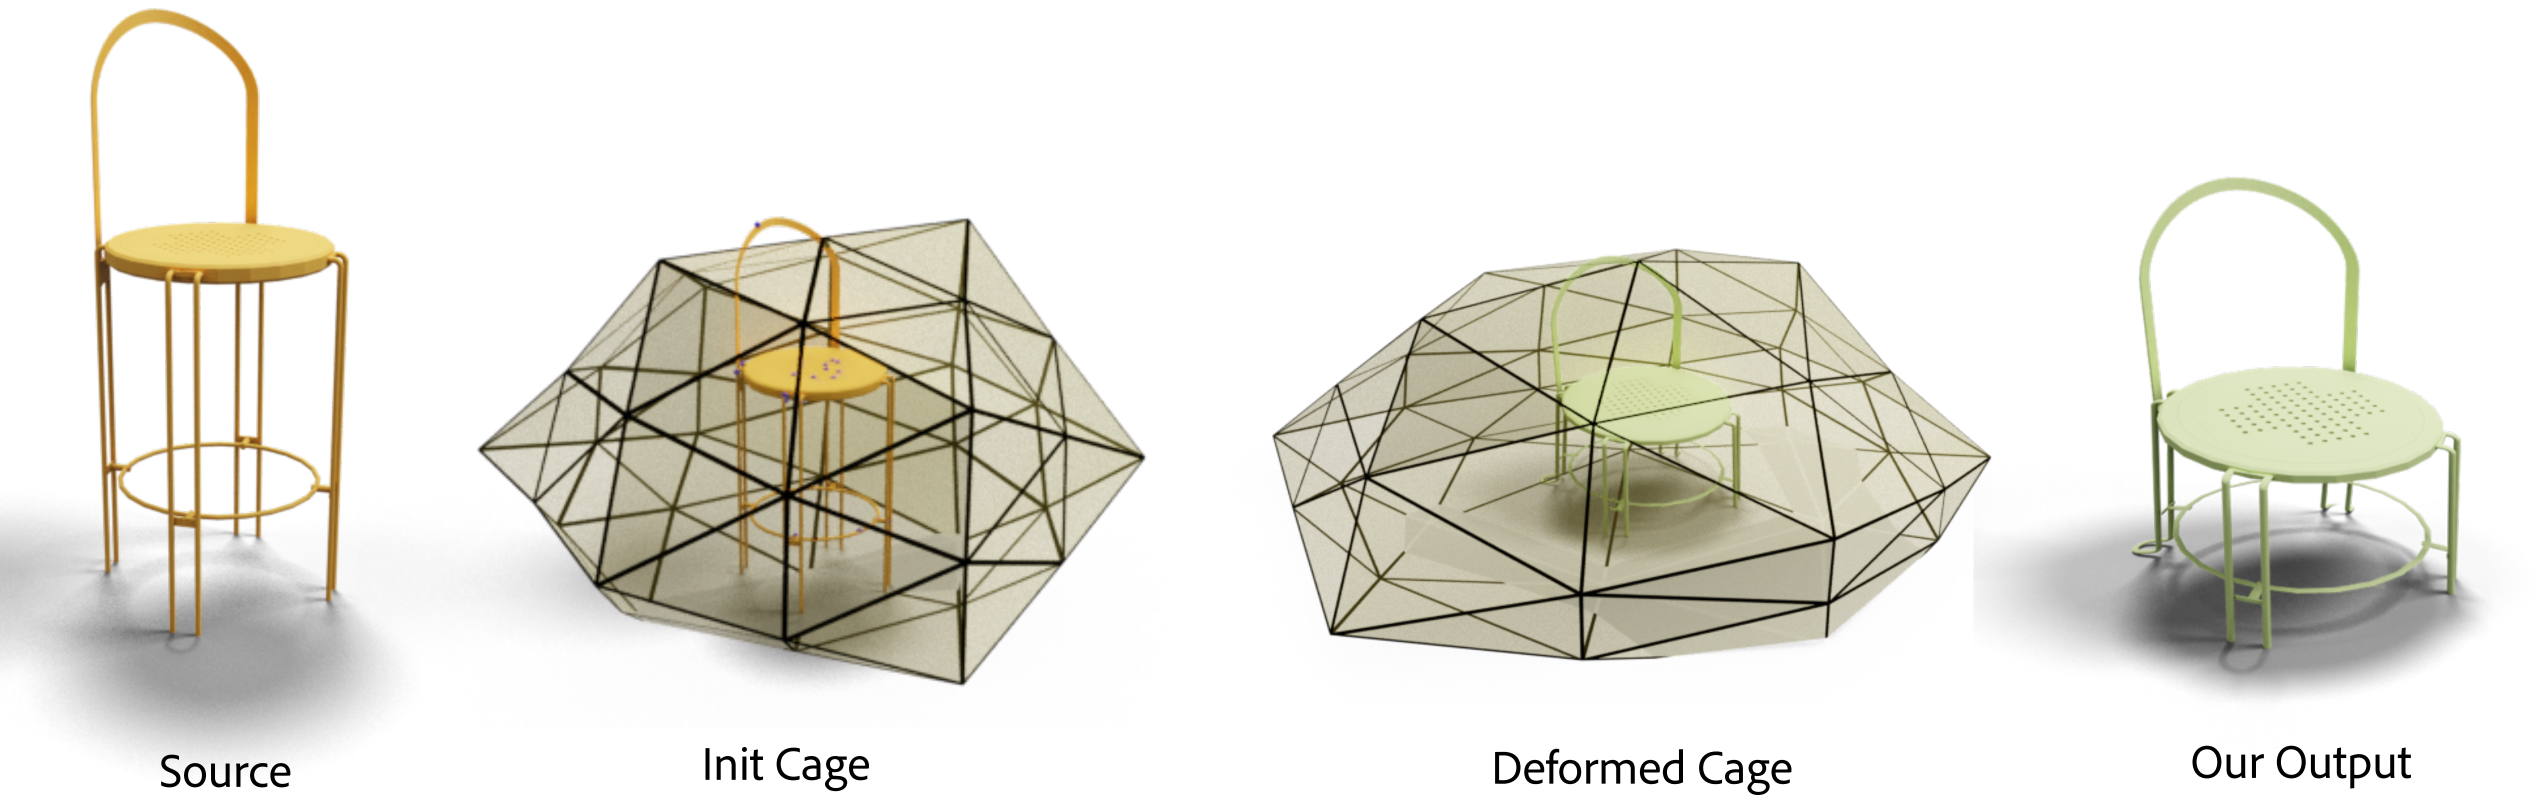

Neural **Cage-based** Deformation

- Step 1: Learn to predict cage parameters

$$f_\theta : z_{\text{source}} \times z_{\text{target}} \rightarrow \mathcal{C}_{\text{init}} \times \mathcal{C}_{\text{deformed}}$$

- Step 2: Use classical cage-based deformation technique

$$\boxed{CBD(\mathcal{C}_{\text{init}}, \mathcal{C}_{\text{deformed}}) : \mathbb{R}^3 \rightarrow \mathbb{R}^3}$$

Deform the source mesh via a differentiable cage-based deformation layer

Source

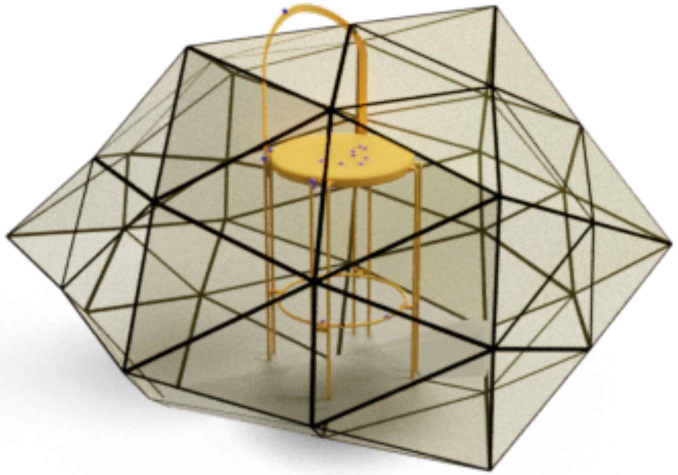

Init Cage

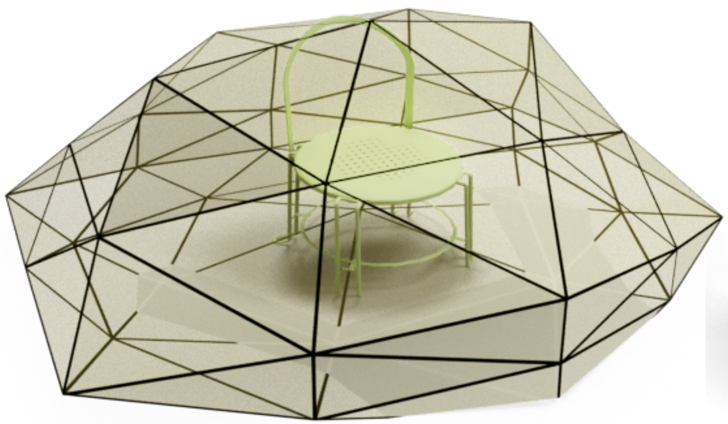

Deformed Cage

Our Output

Application: Stock Amplification

Create shape variations by picking random source/target pairs

Application: Deformation Transfer

Transfer a pose from a target to the source mesh

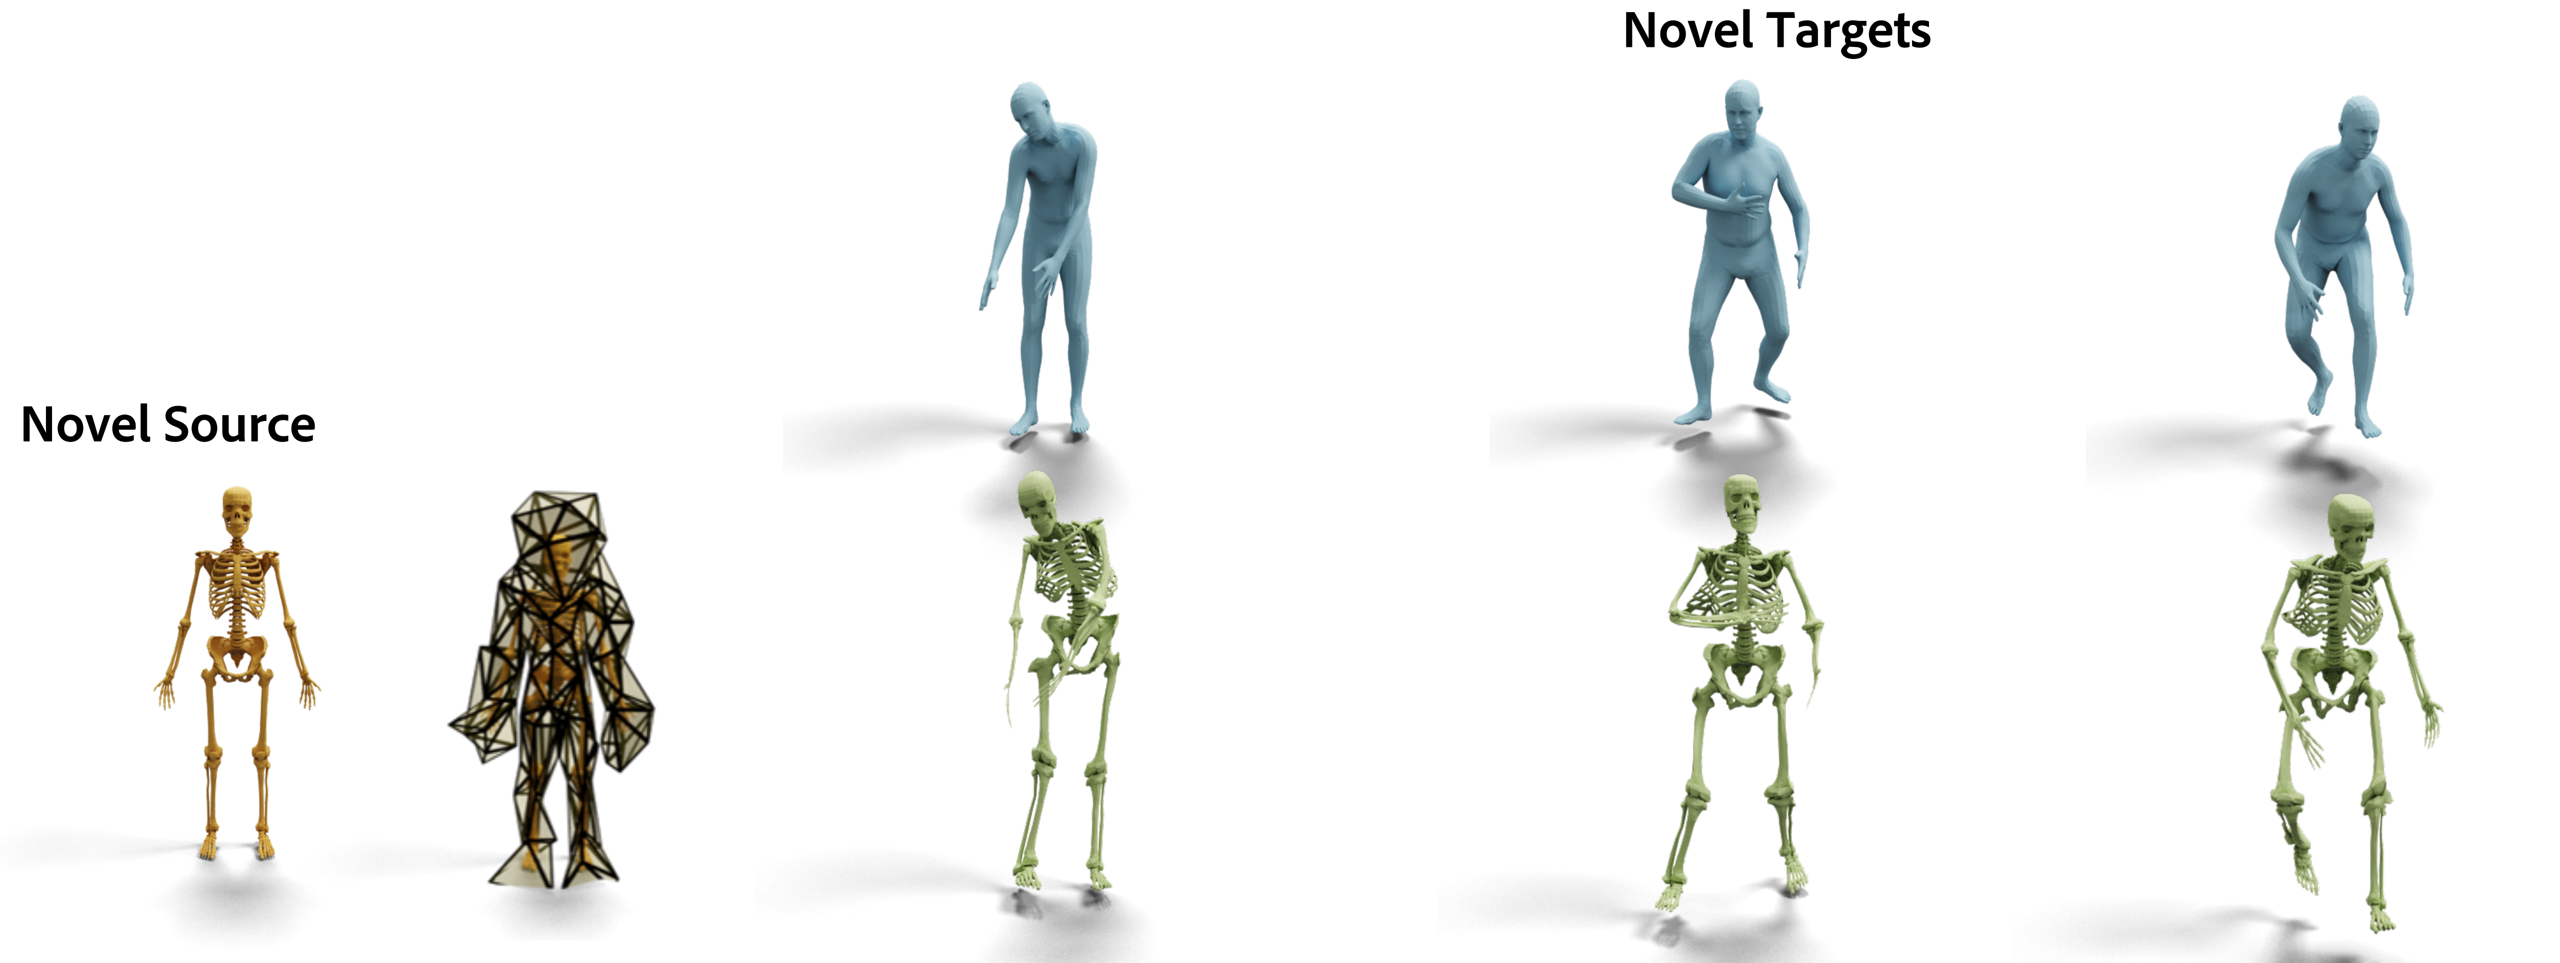

Application: Deformation Transfer

Transfer a pose from a target to the source mesh

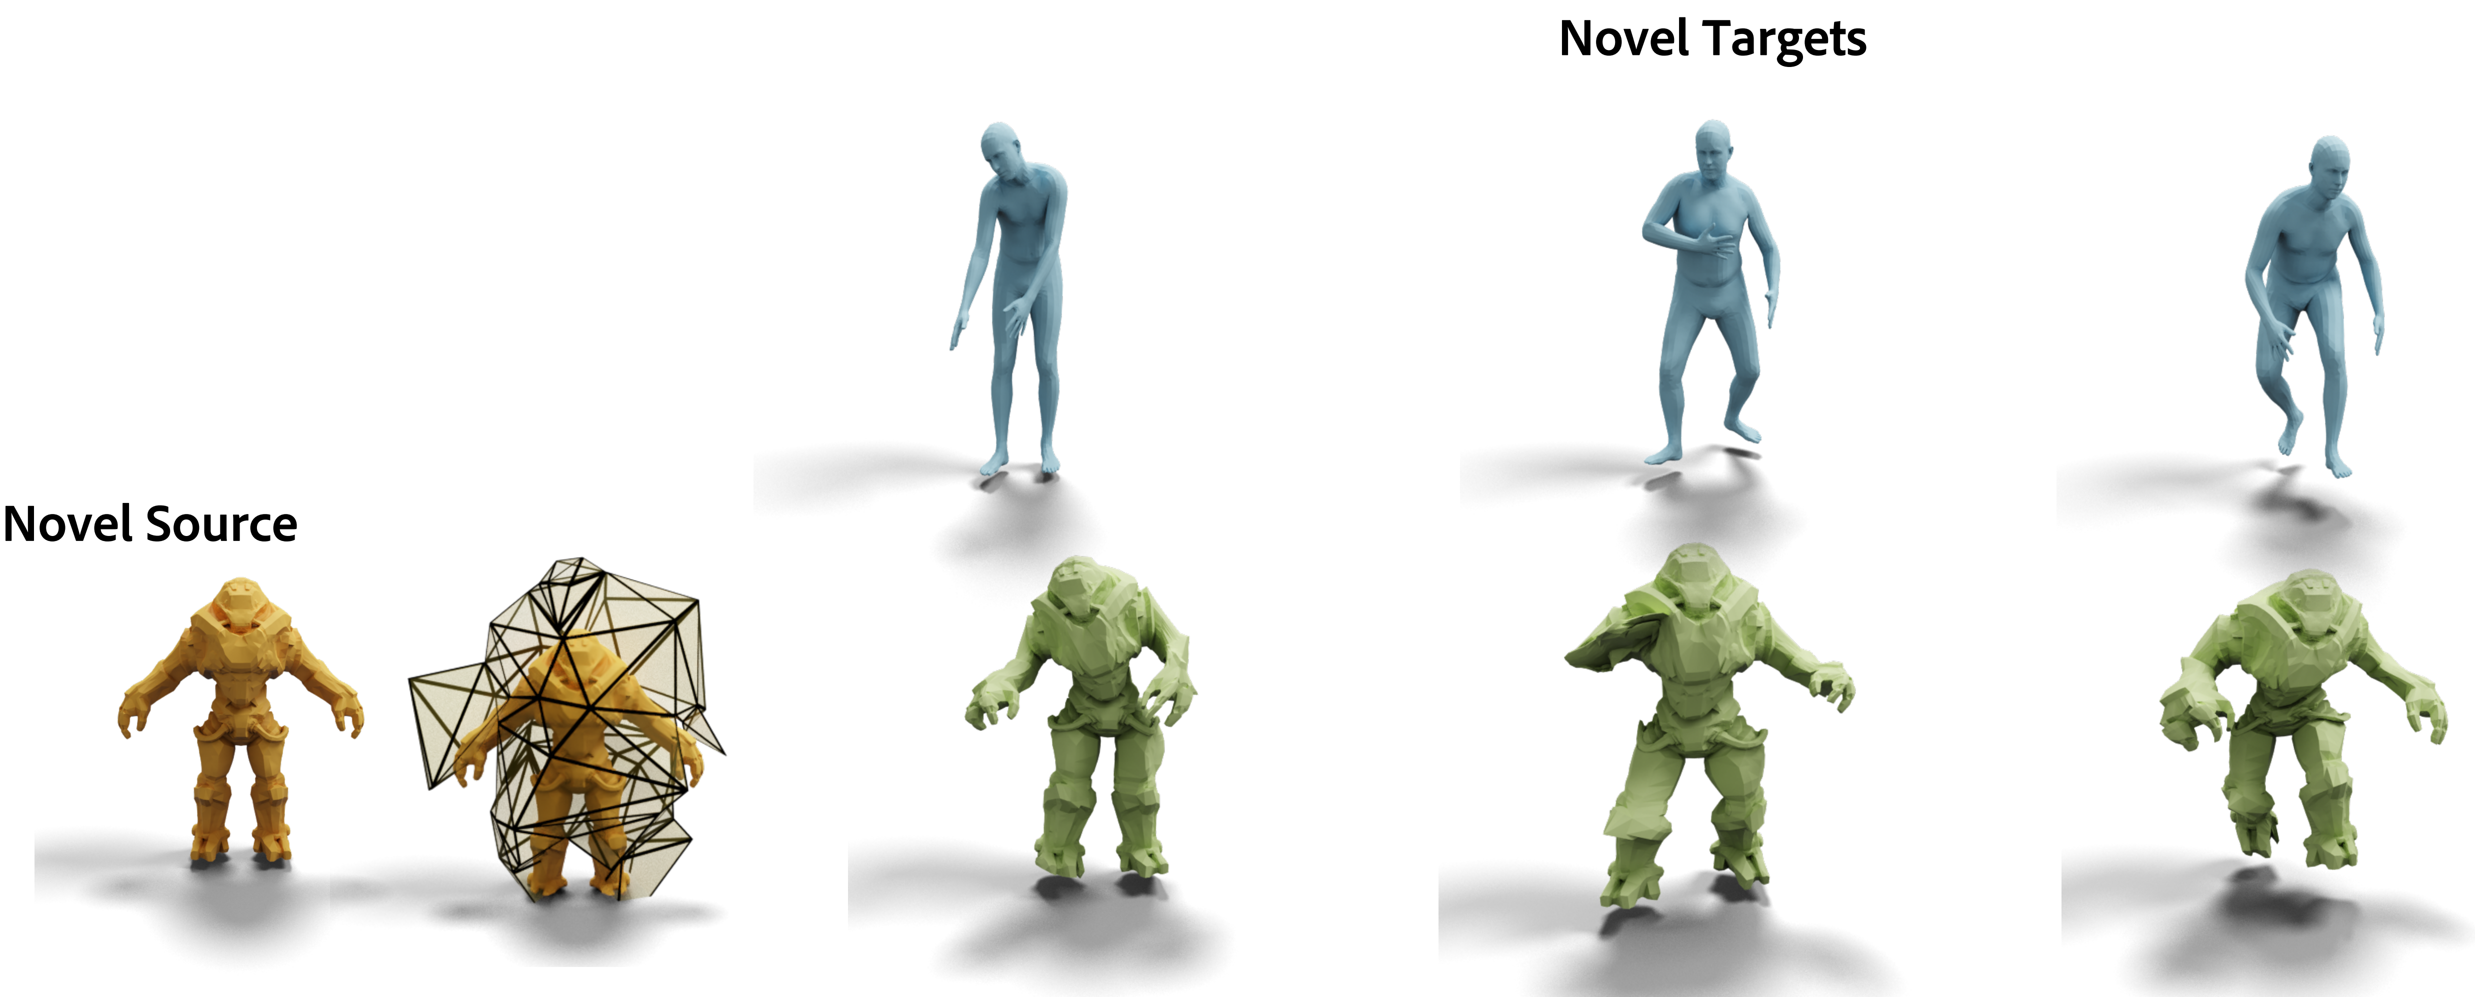

Cage-free Gradient Domain Deformation

- Hard to learn cages for complex shapes
- Cage-based deformation is not mesh-specific – it just maps the volume

Cage-free Gradient Domain Deformation

- Hard to learn cages for complex shapes
- Cage-based deformation is not mesh-specific – it just maps the volume
- Reminder: learning a map directly is prone to noise – hard to preserve details

$$f_\theta : z_{\mathbf{source}} \times z_{\mathbf{target}} \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$$

Cage-free Gradient Domain Deformation

- Hard to learn cages for complex shapes
- Cage-based deformation is not mesh-specific – it just maps the volume
- Reminder: learning a map directly is prone to noise – hard to preserve details

$$f_\theta : z_{\text{source}} \times z_{\text{target}} \times \mathbb{R}^3 \to \cancel{\mathbb{R}^3}\ \mathbb{R}^{3\times 3}$$

Predict a matrix

Cage-free Gradient Domain Deformation

- Hard to learn cages for complex shapes
- Cage-based deformation is not mesh-specific – it just maps the volume
- Reminder: learning a map directly is prone to noise – hard to preserve details

$$f_\theta : z_{\text{source}} \times z_{\text{target}} \times \mathbb{R}^3 \to \cancel{\mathbb{R}^3}\ \mathbb{R}^{3\times 3}$$

Predict a matrix

$$\big\downarrow \pi$$

$$\mathbb{R}^{3\times 2}$$

Project to Jacobian

Cage-free Gradient Domain Deformation

- Hard to learn cages for complex shapes
- Cage-based deformation is not mesh-specific – it just maps the volume
- Reminder: learning a map directly is prone to noise – hard to preserve details

$$f_\theta : z_{\text{source}} \times z_{\text{target}} \times \underset{\mathcal{S}}{\cancel{\mathbb{R}^3}} \rightarrow \cancel{\mathbb{R}^3}\ \mathbb{R}^{3\times 3}$$

Use points
on the surface
(triangle centroids,
Intrinsic features)

Predict a matrix

$$\mathbb{R}^{3\times 3} \Big\downarrow \pi \quad \mathbb{R}^{3\times 2}$$

Project to Jacobian

Neural Jacobian Fields Pipeline

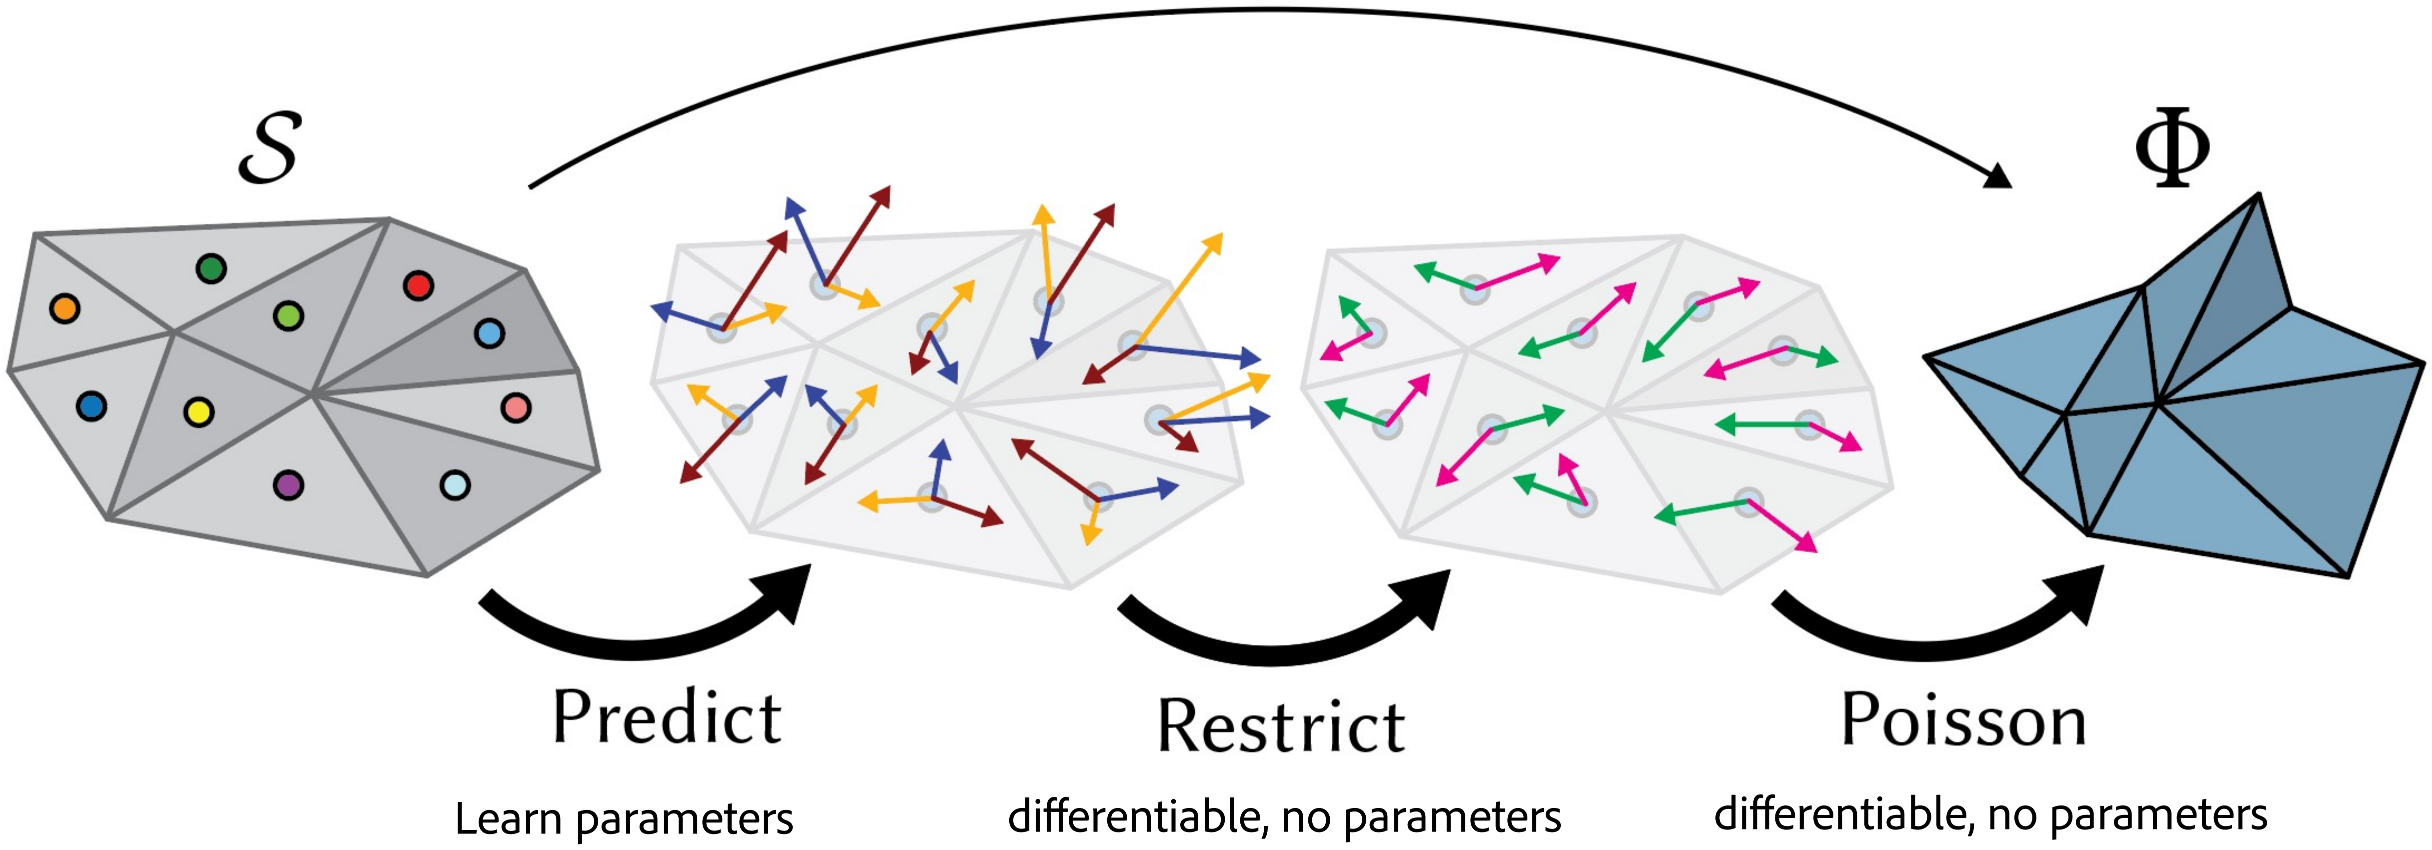

Neural Jacobian Fields Pipeline

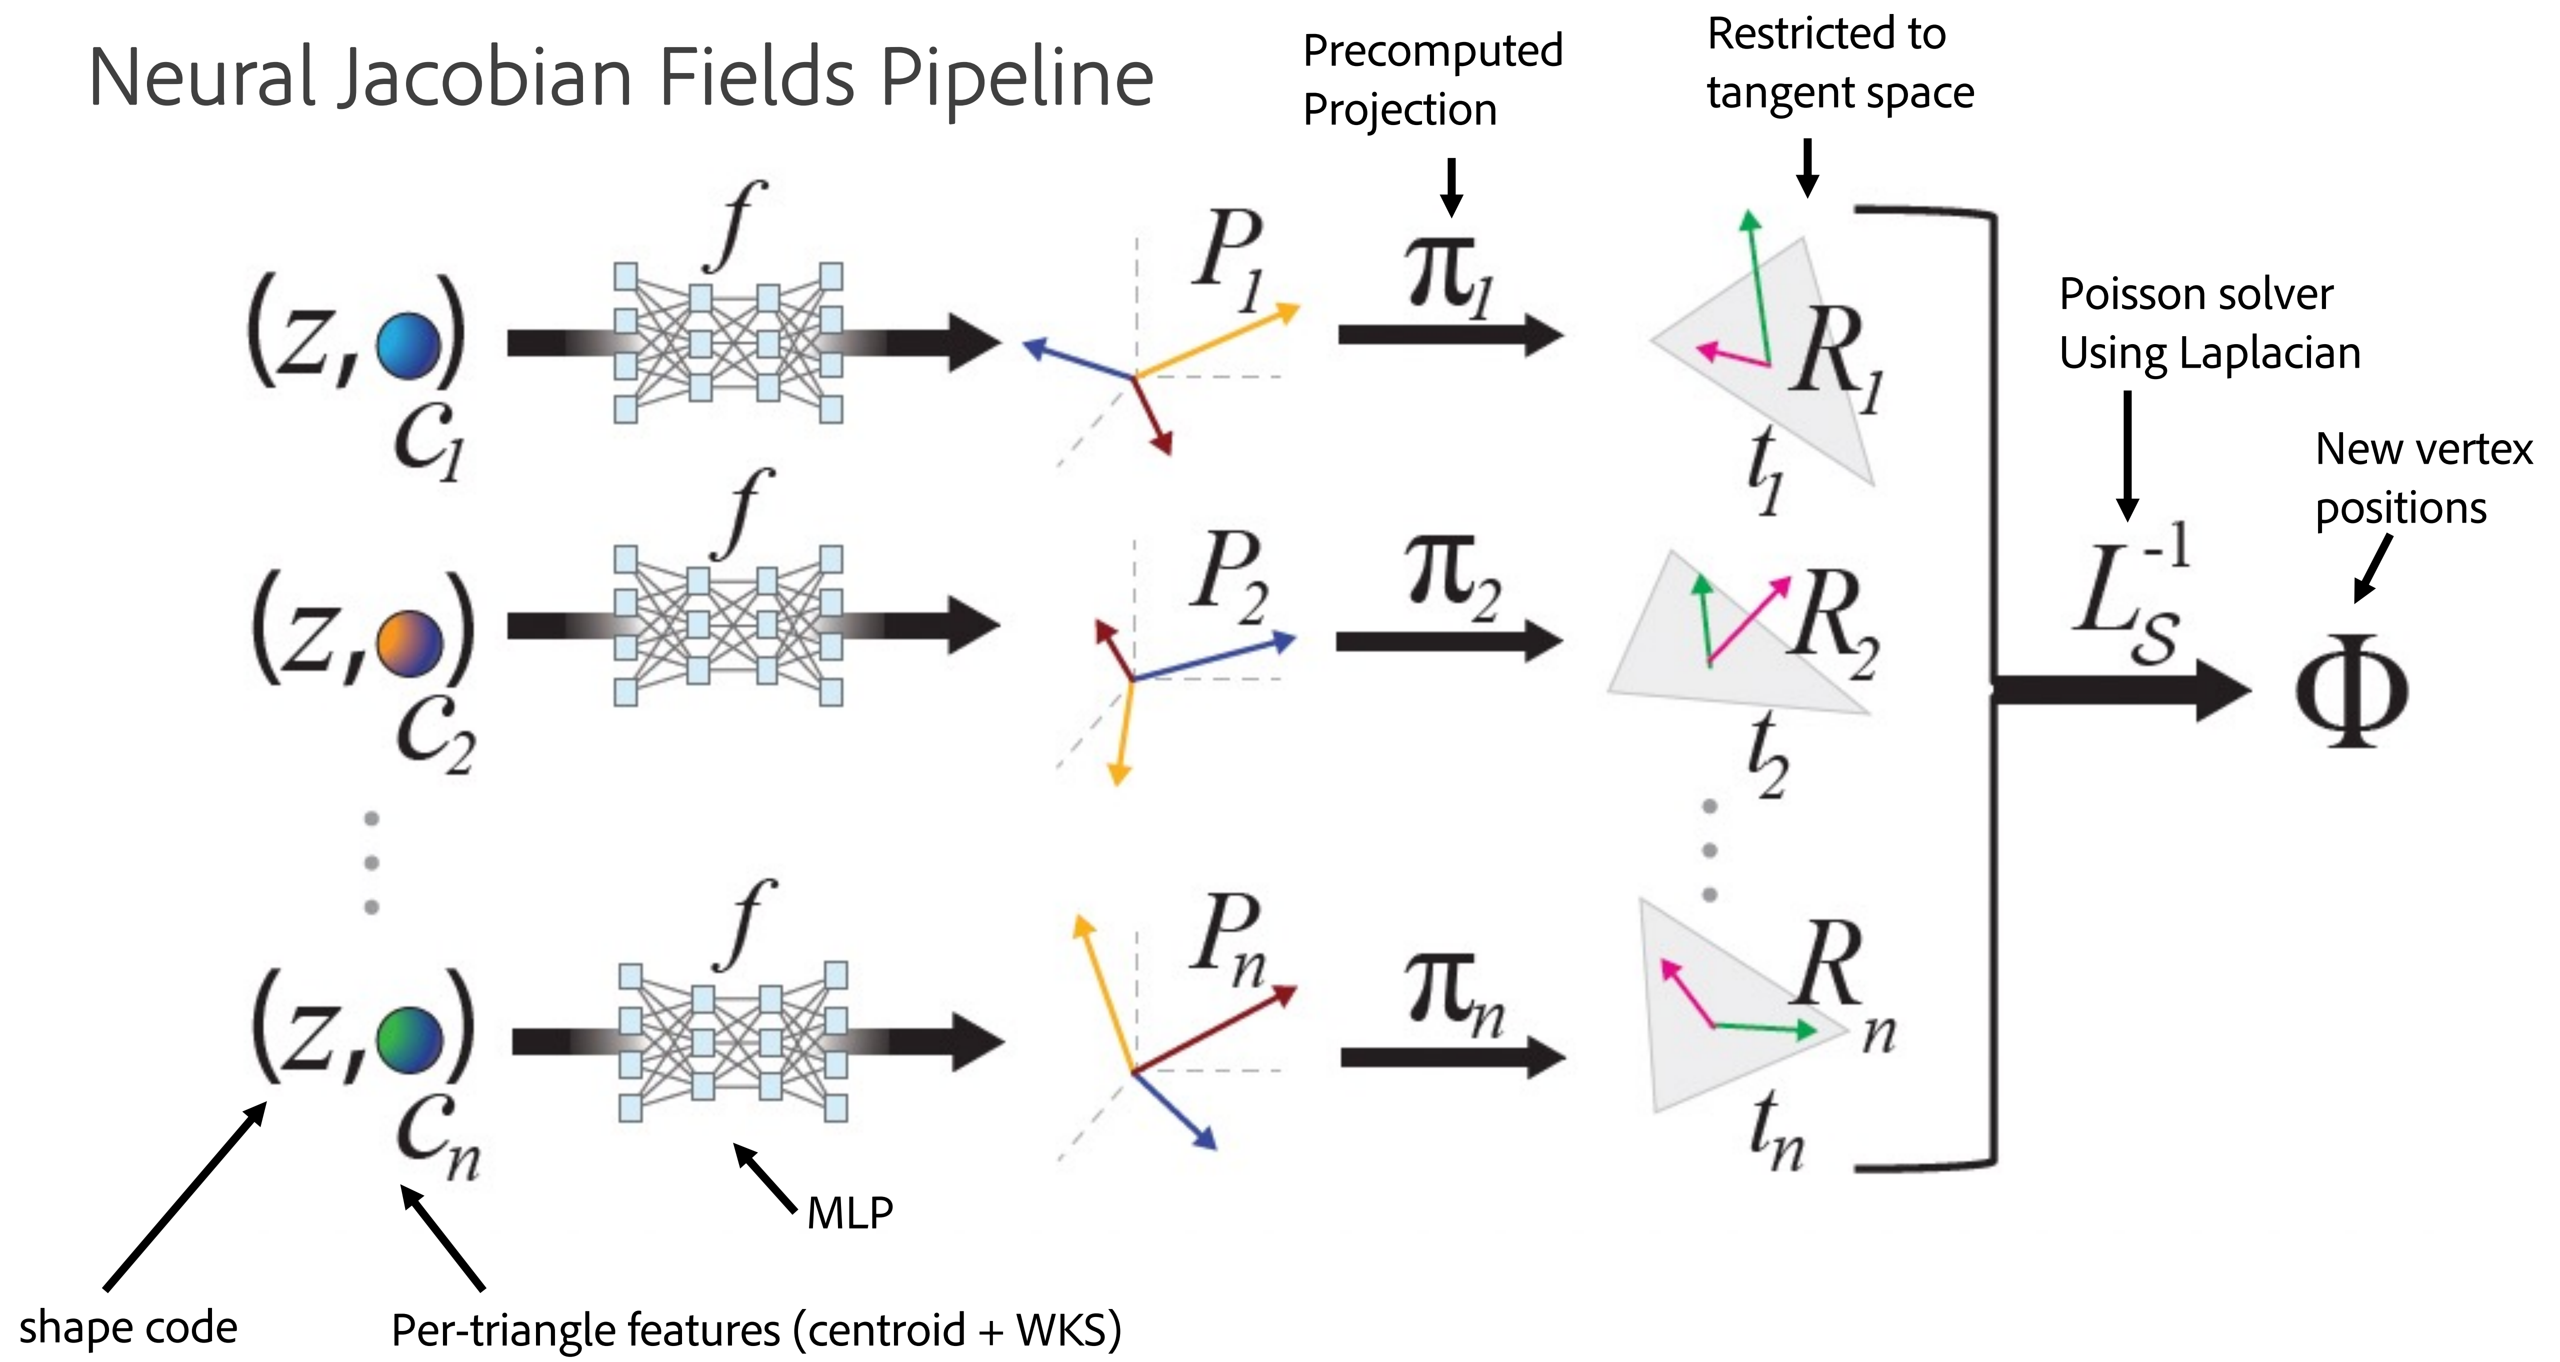

Training Neural Jacobian Fields

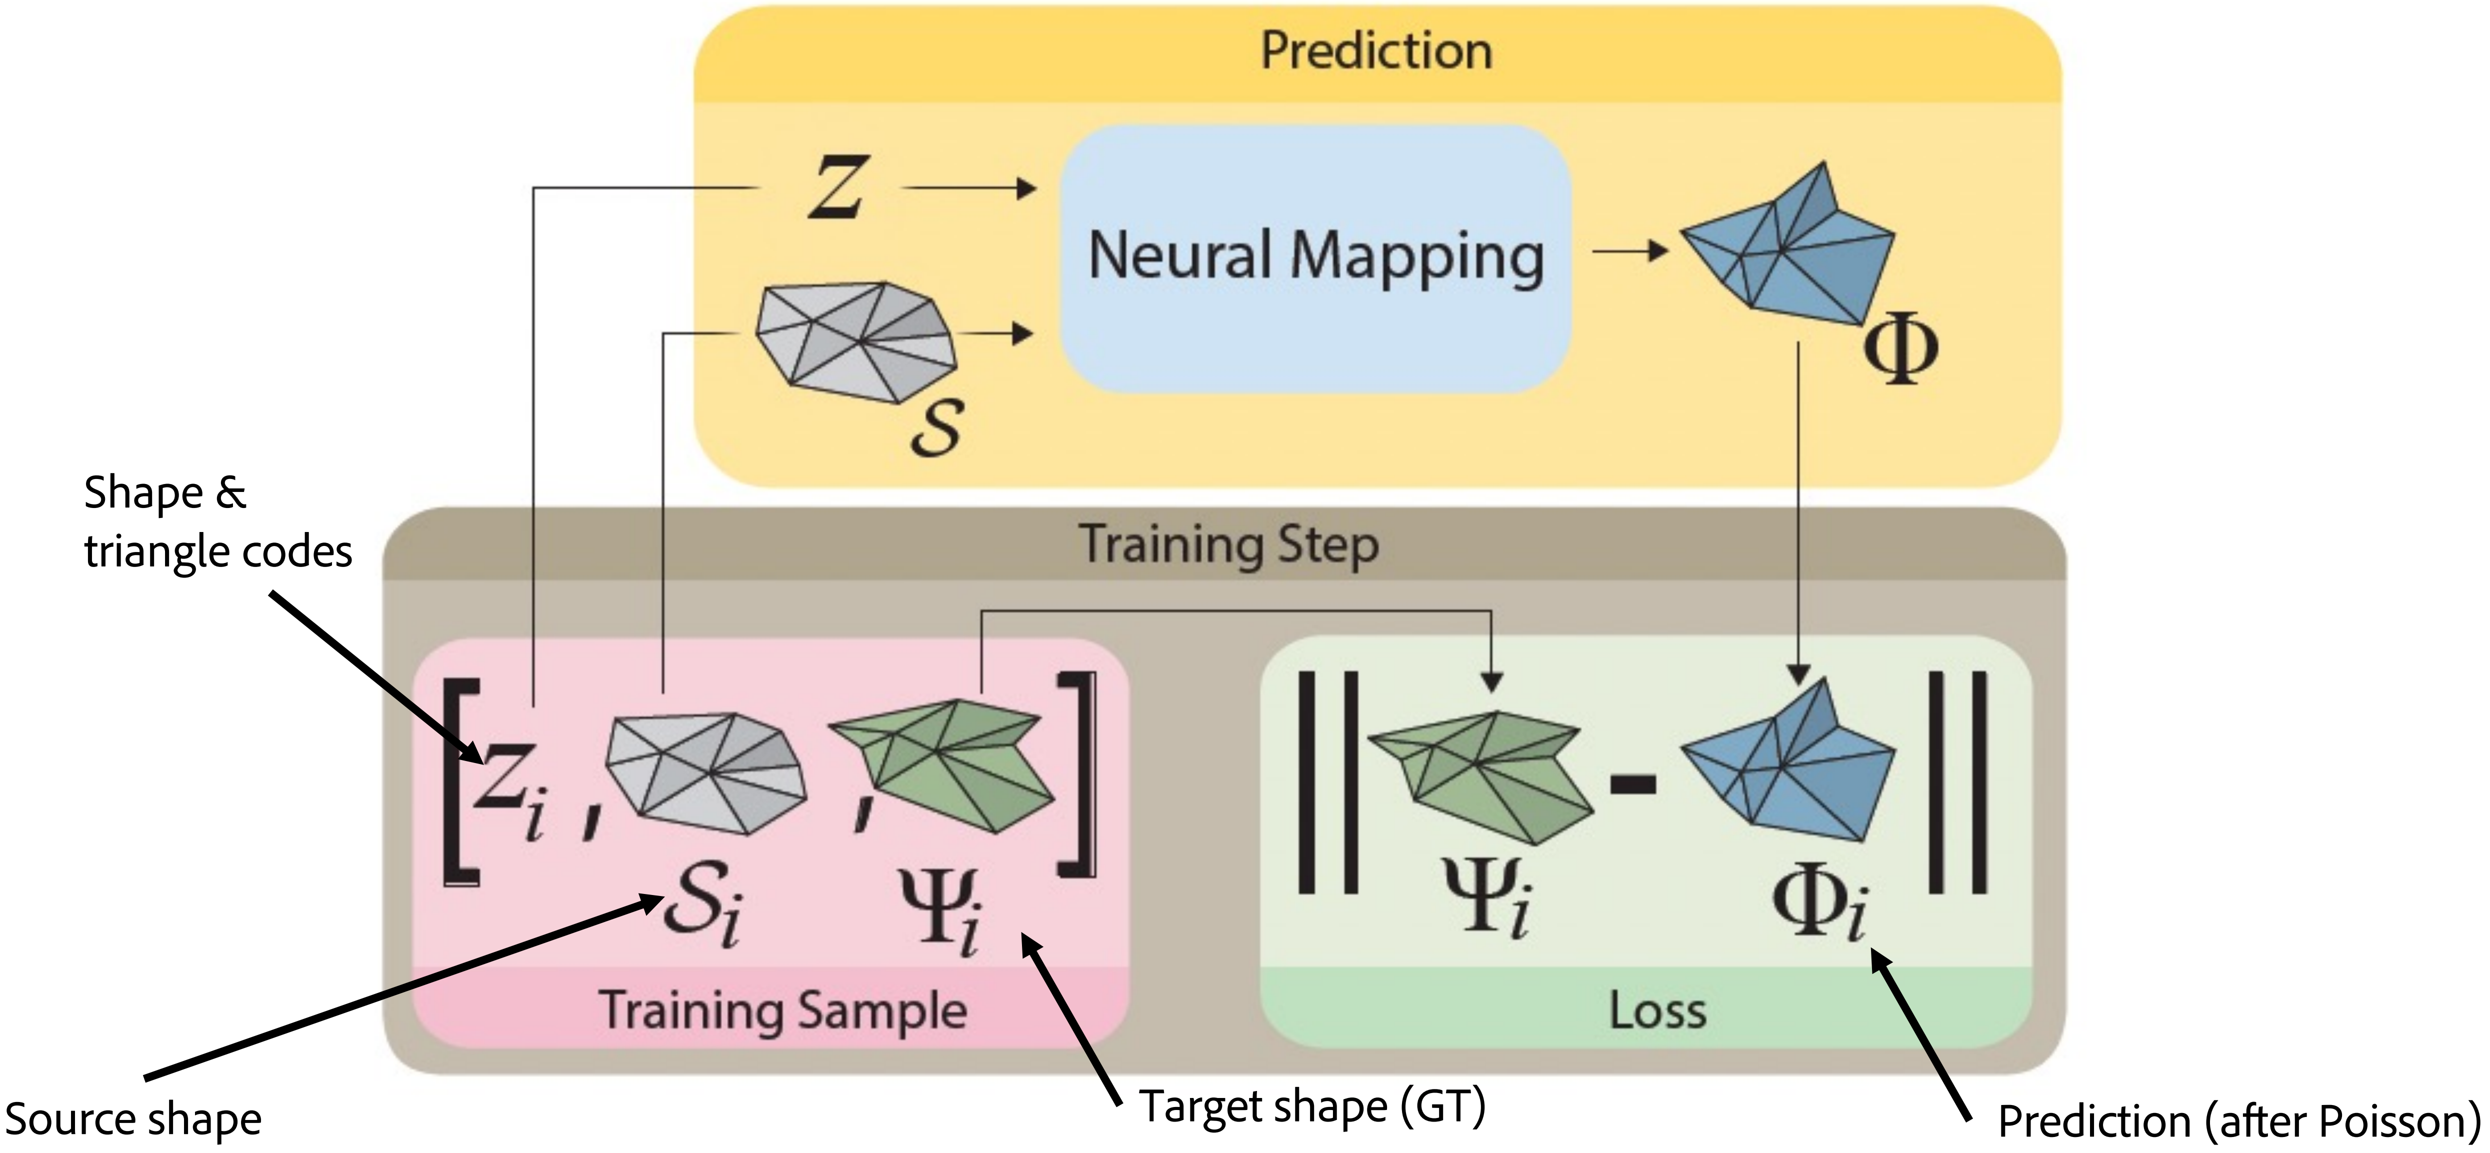

Application: Cageless Deformation Transfer

Only trained on humans, no extra input was needed for Big Buck Bunny

Partial Registration

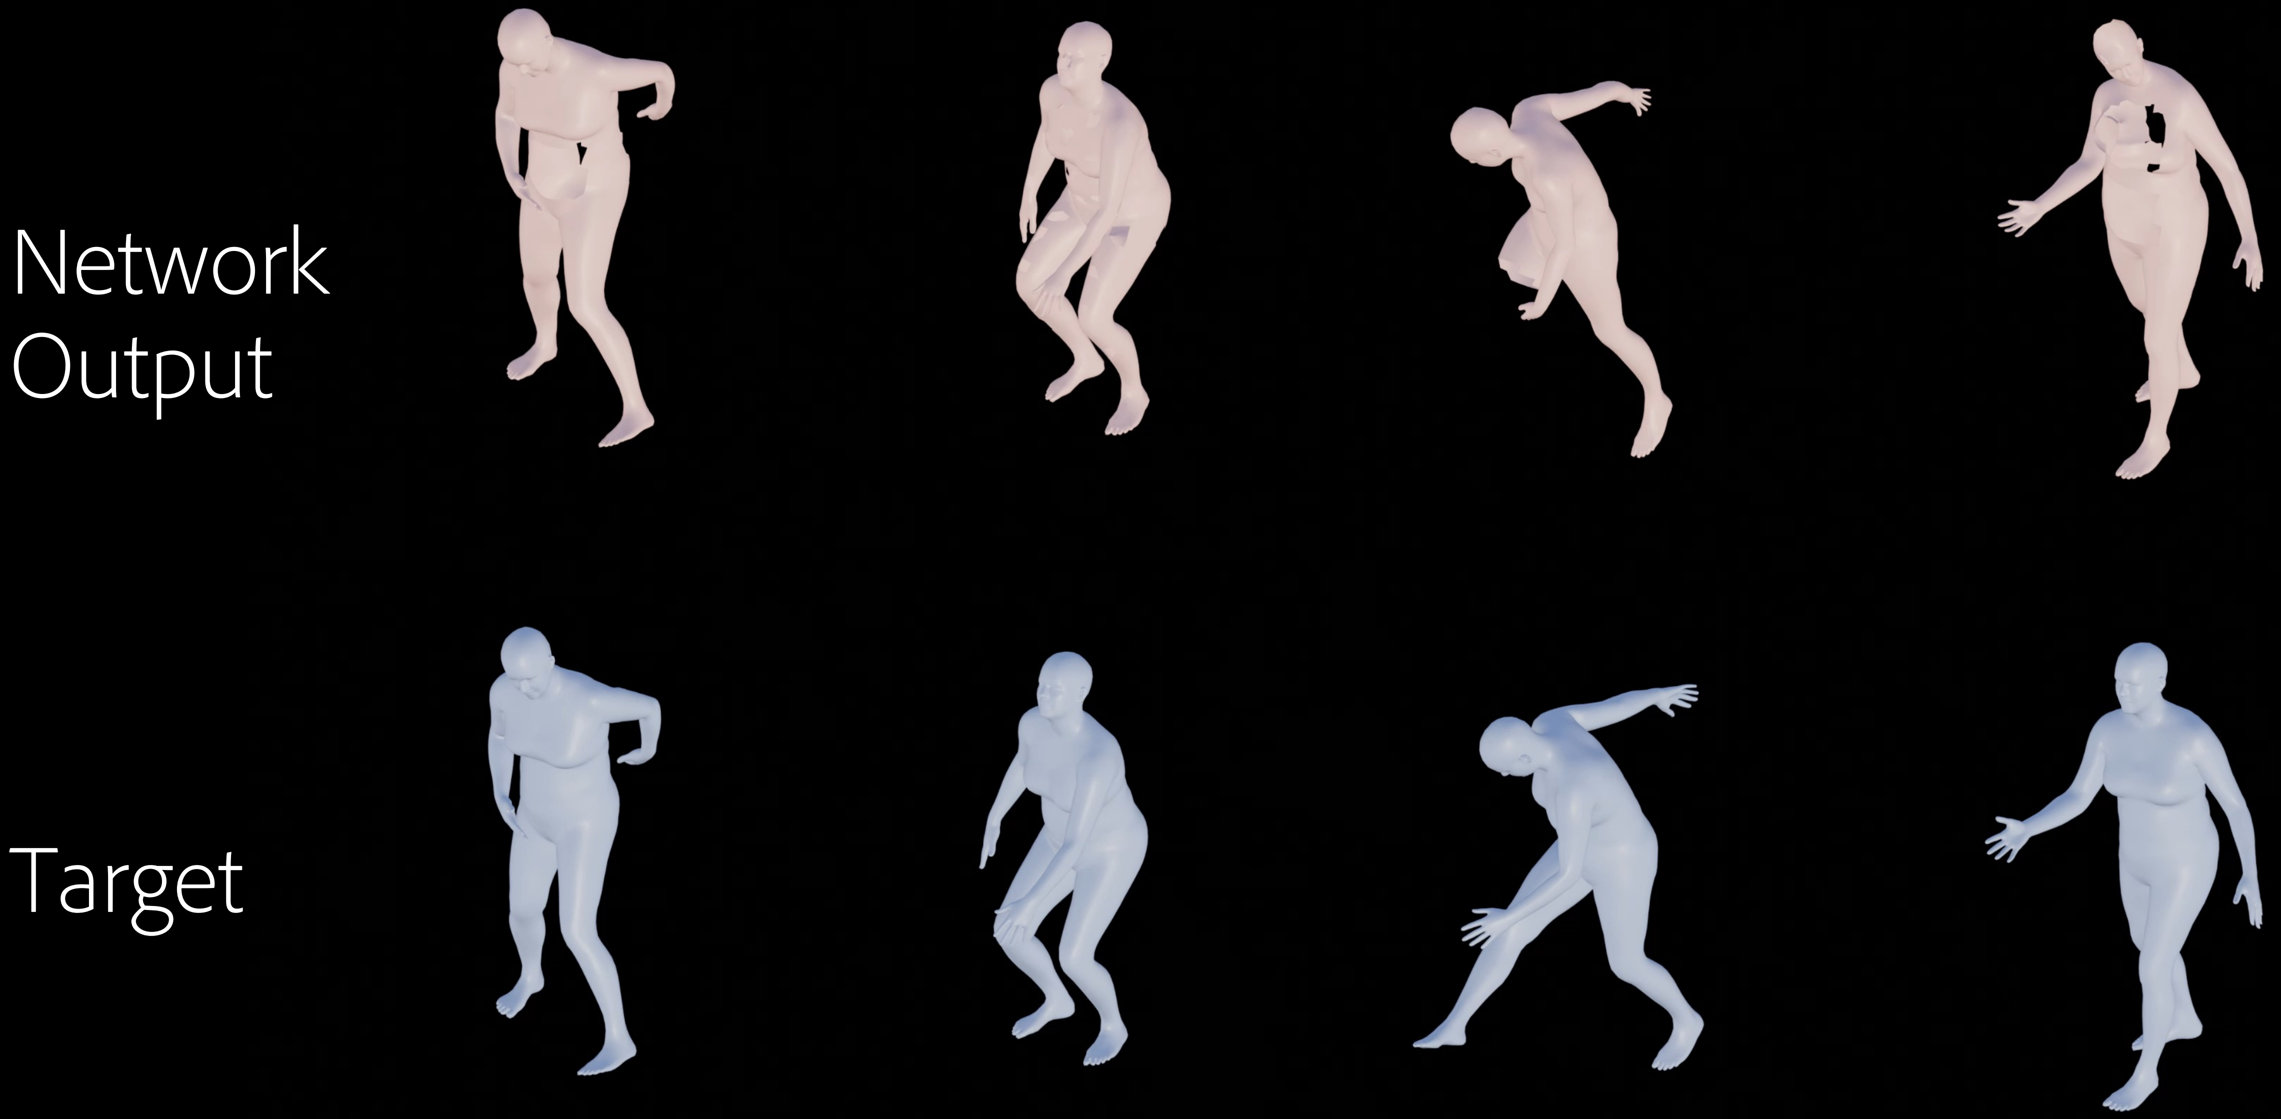

Morphing

Application: Learn to AutoUV

Supervised on SLIM parameterizations

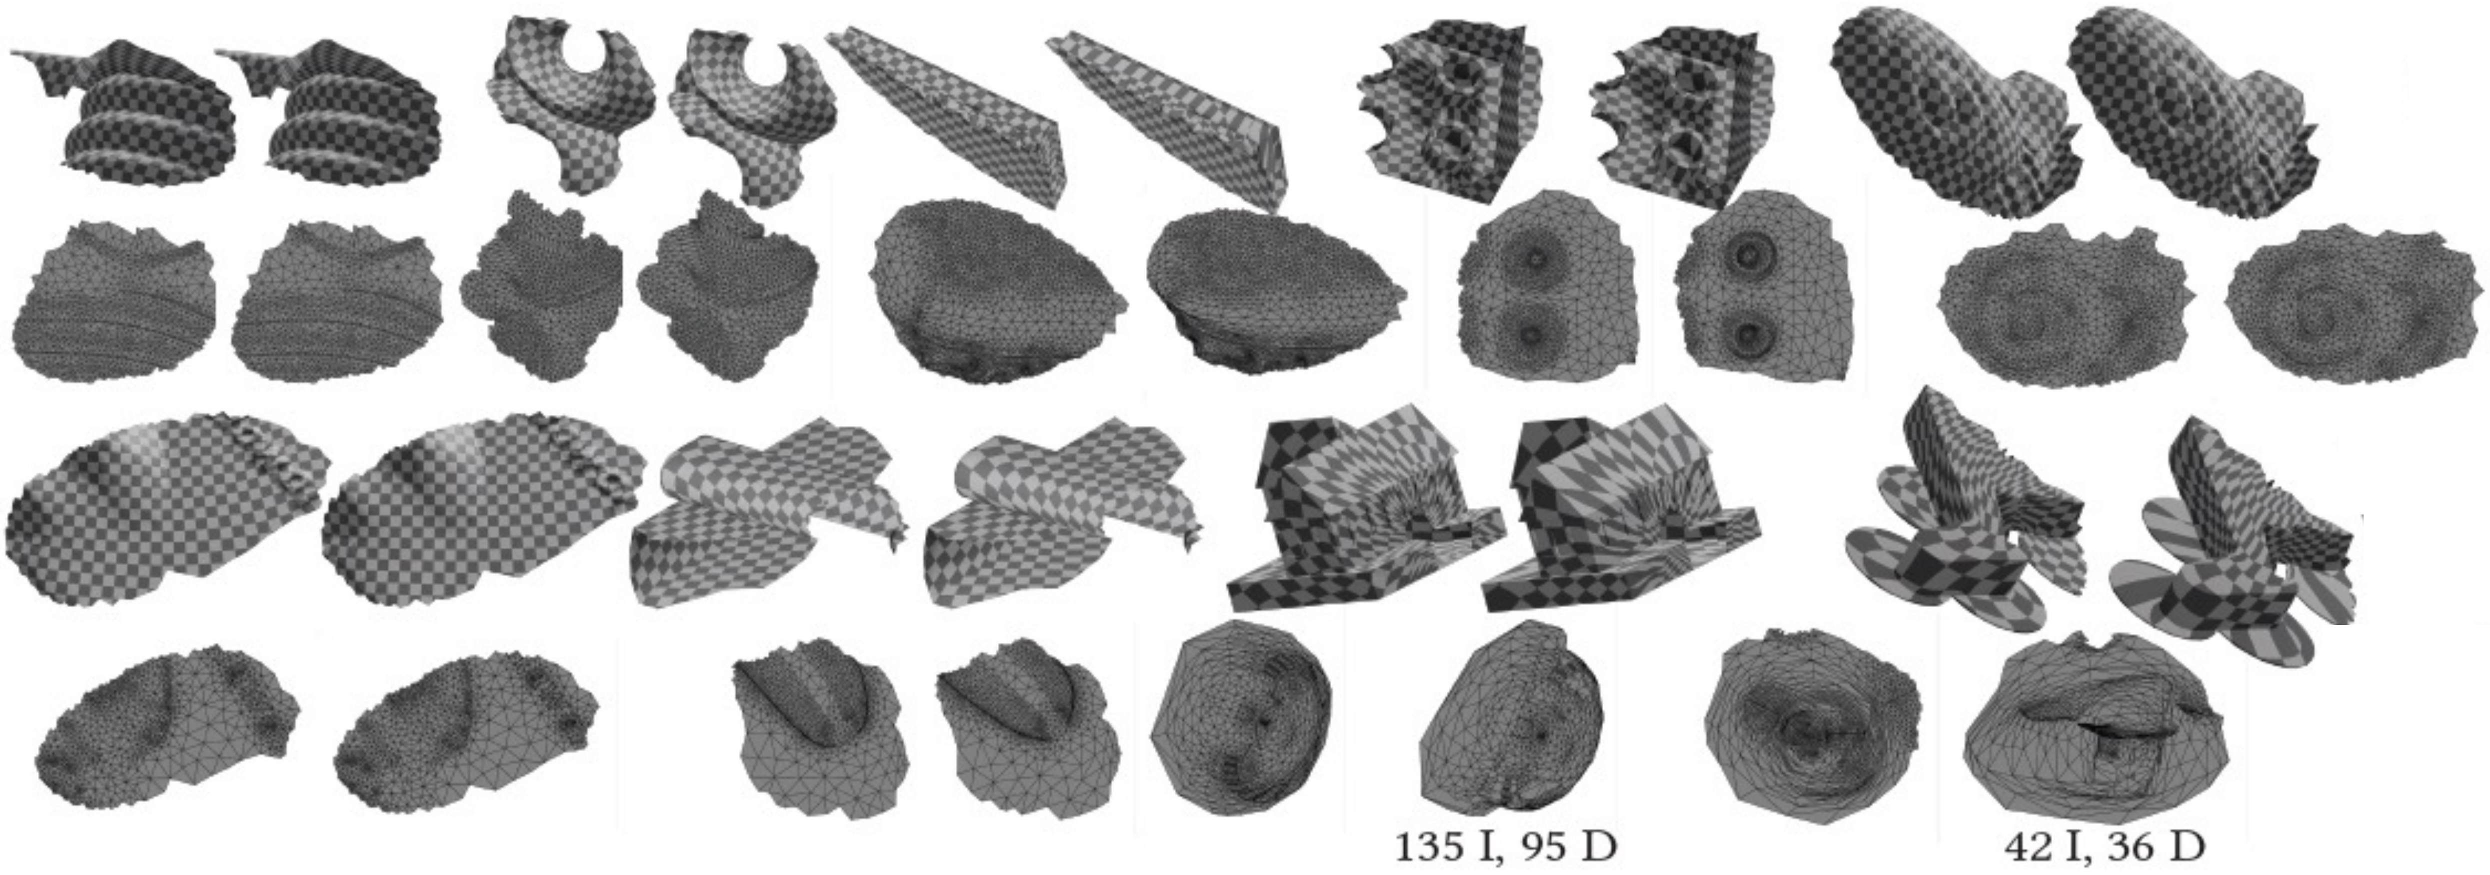

Application: Learn Collision-based Deformation

Using setup of Romero et al. 2021

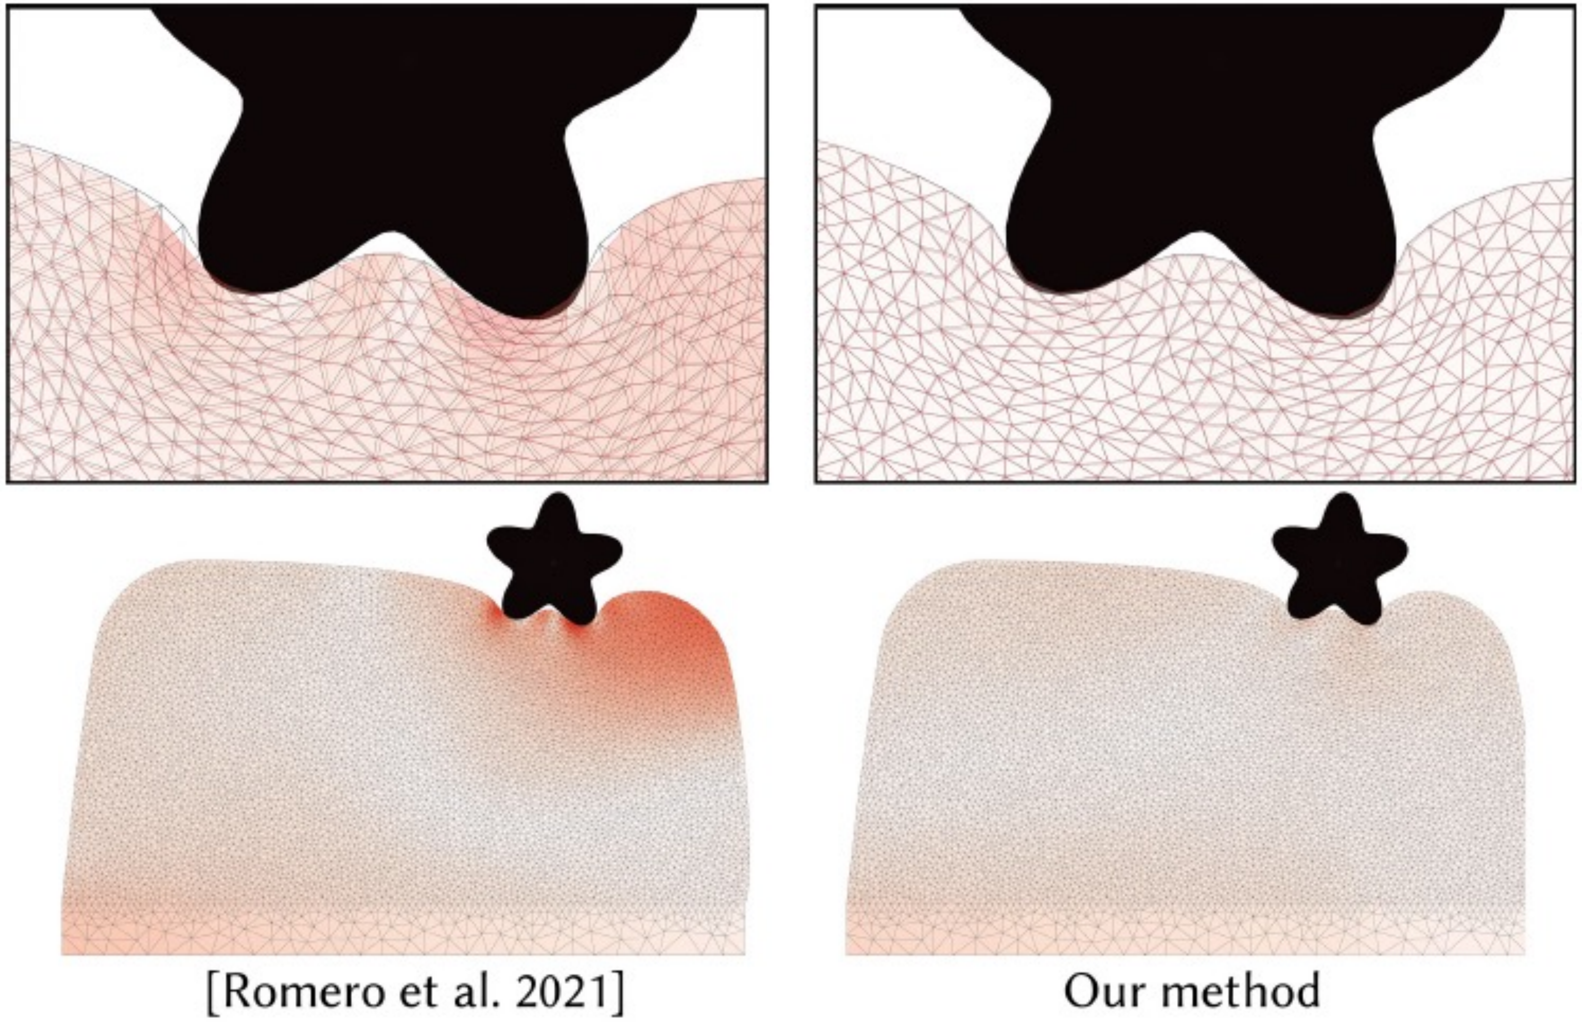

Neural Shape Processing

Modify existing shapes instead of generating from scratch

Goal: Detail Transfer

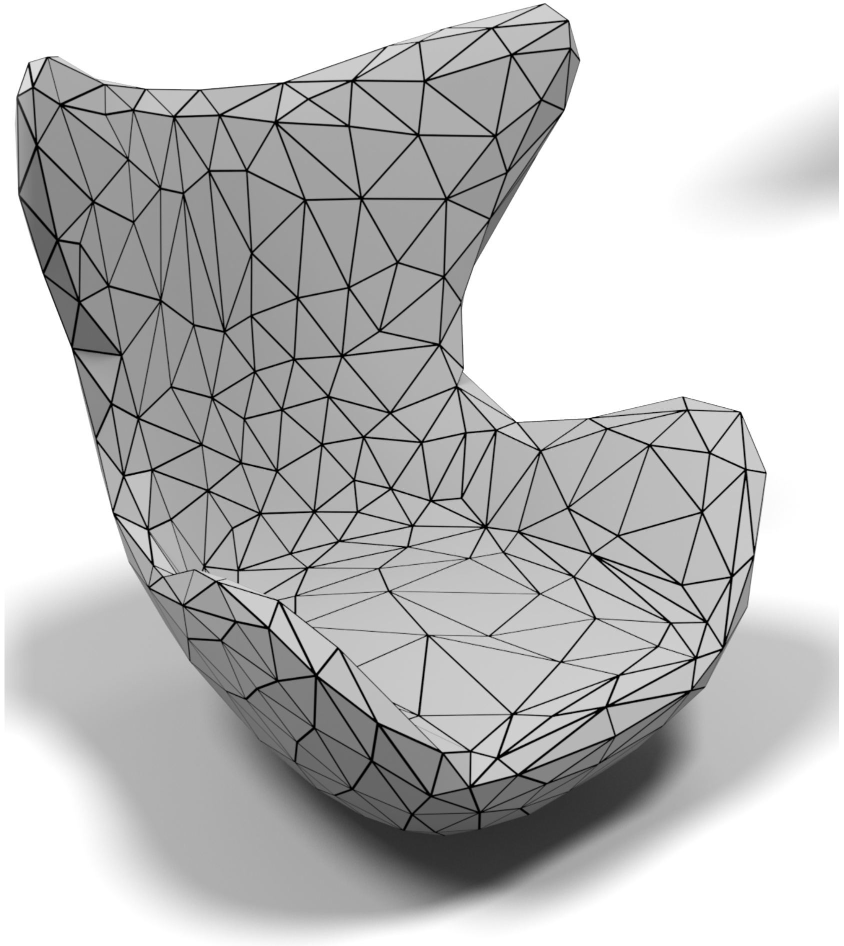

Input

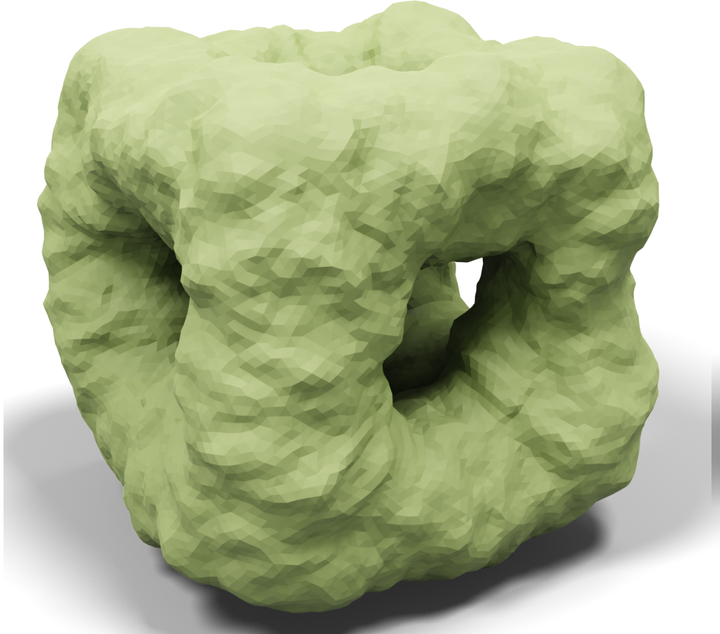

Target Details

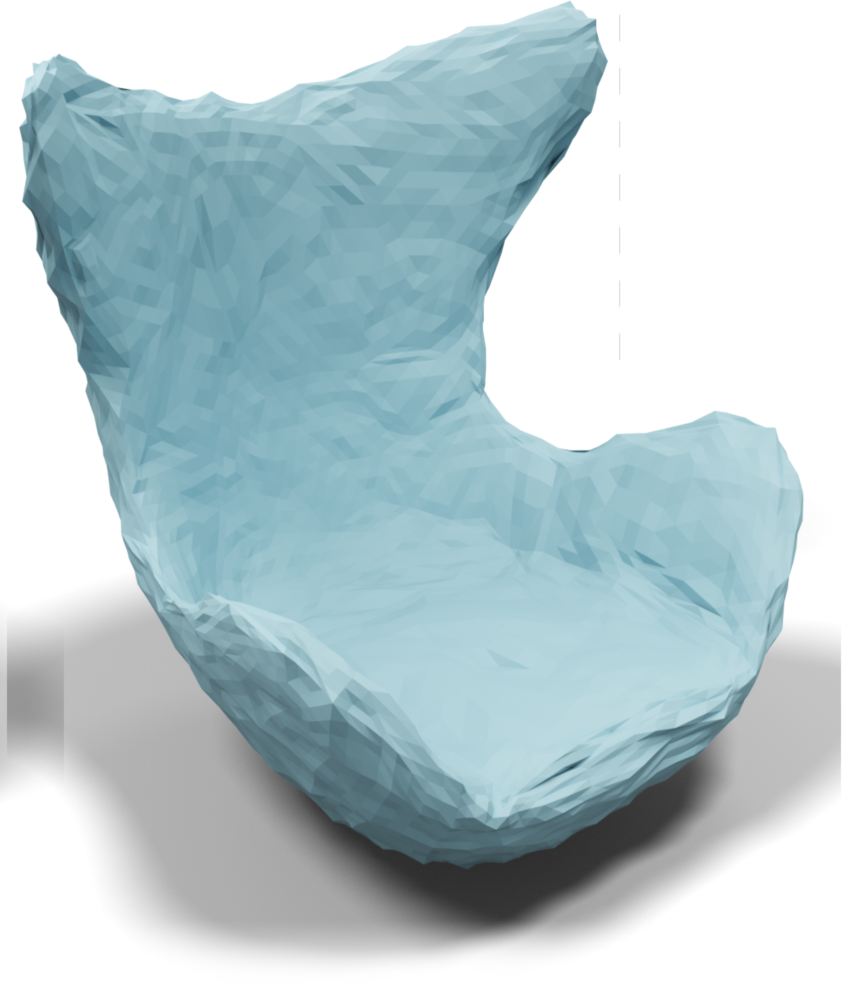

Output

Our Approach

- Decimate high-res mesh with target details to create training data
- Learn local up-sampling filters

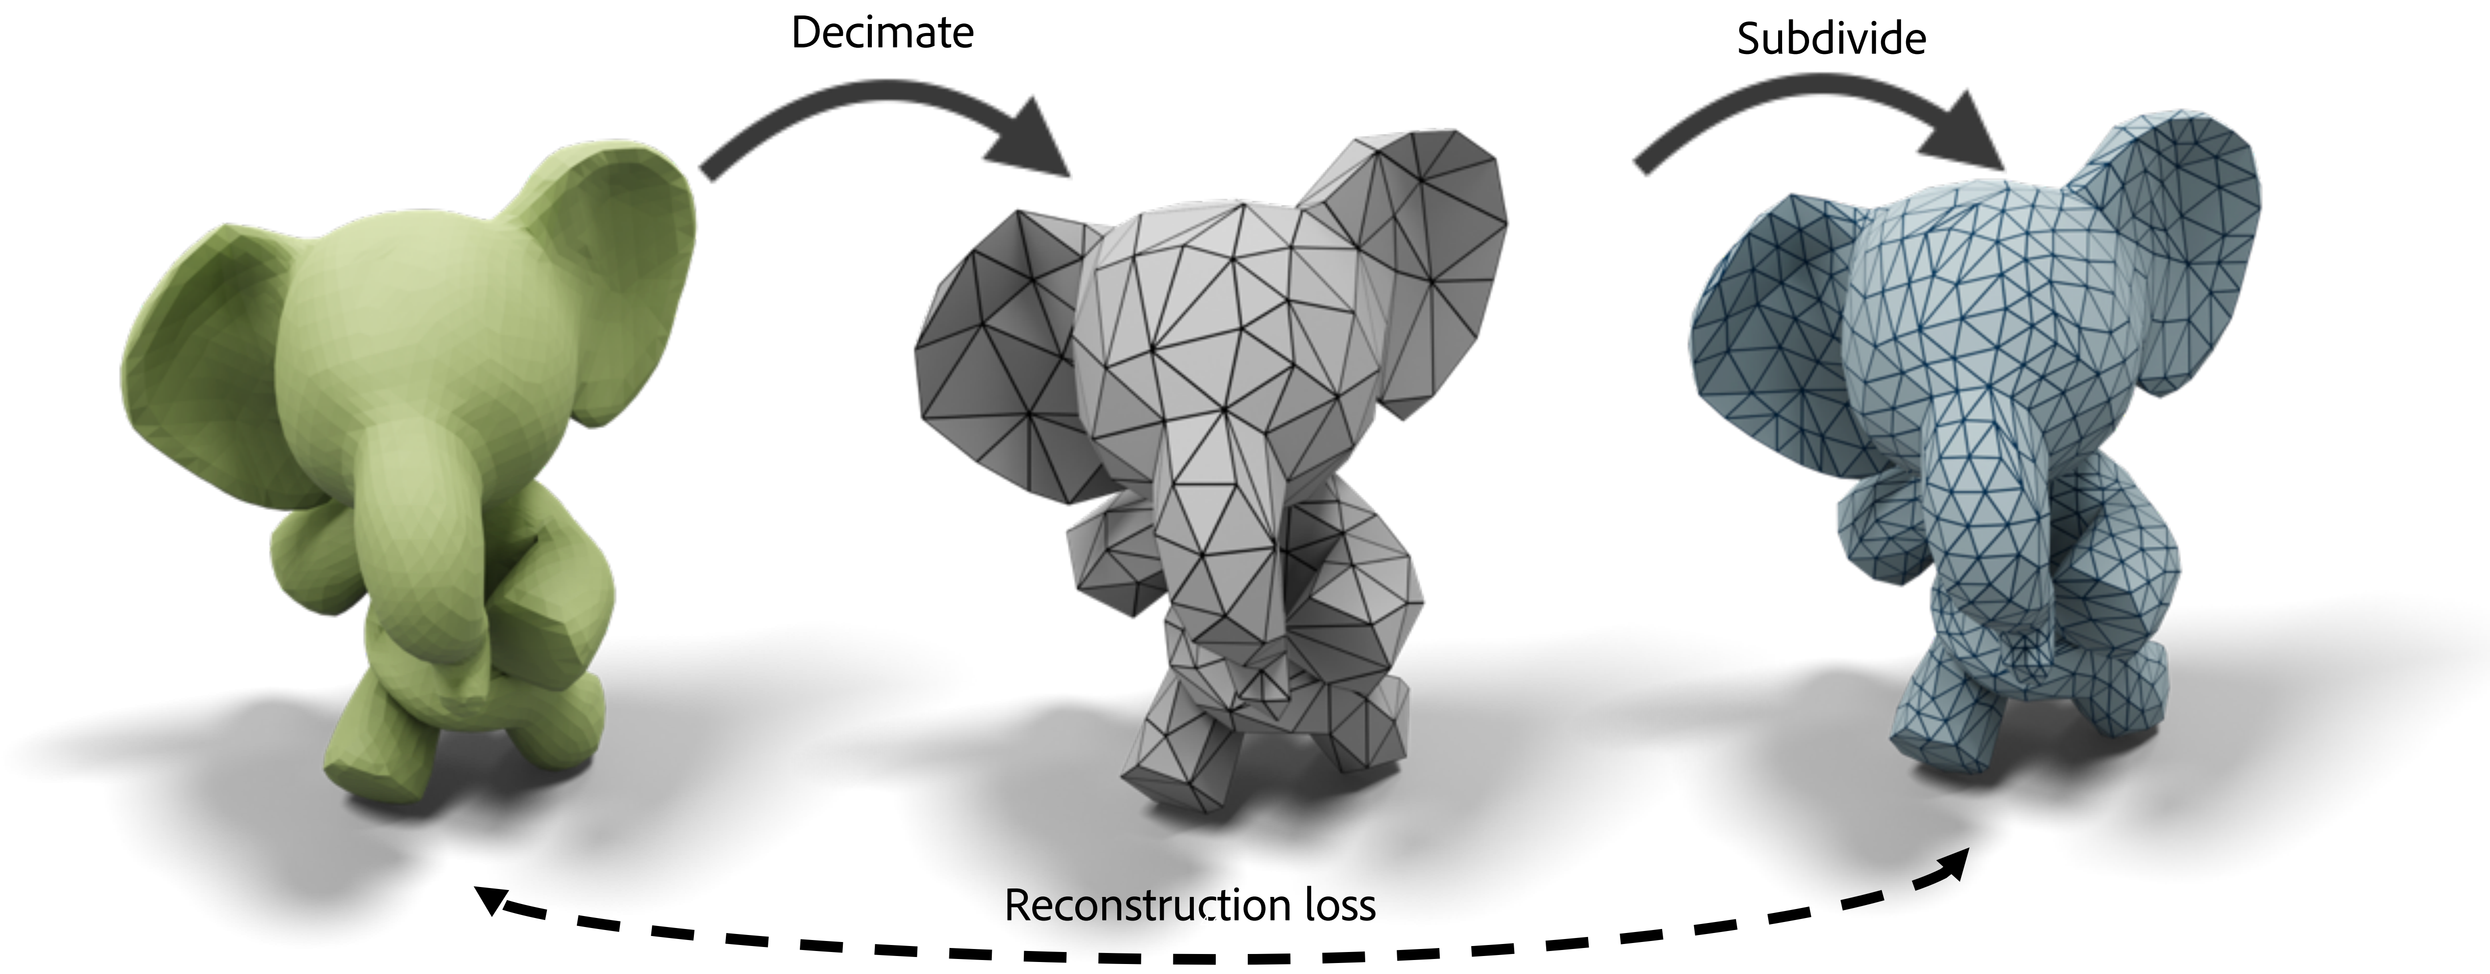

Maintaining Bijective Mapping

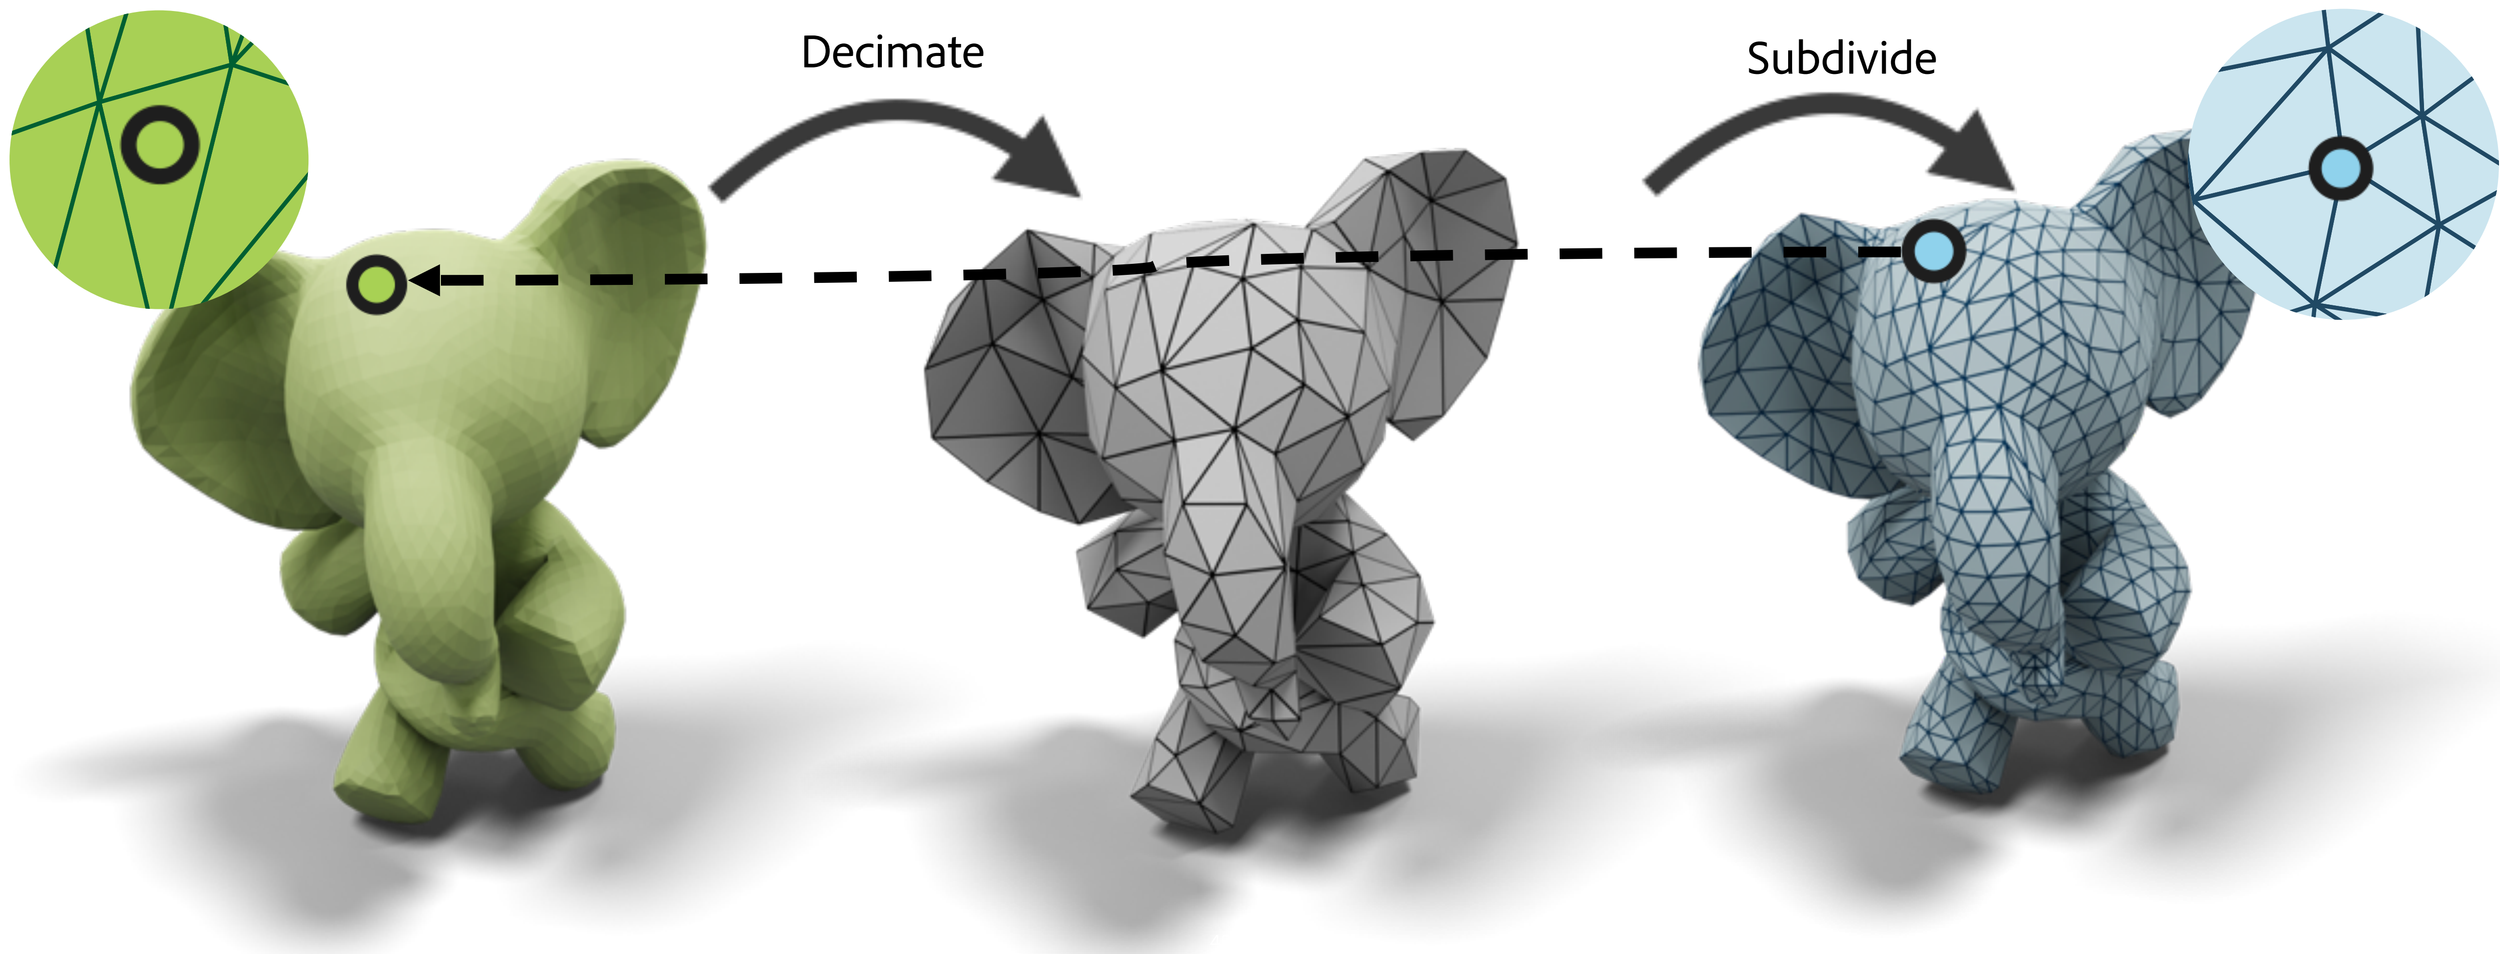

Maintaining Bijective Mapping

- Record barycentric coordinates during subdivision

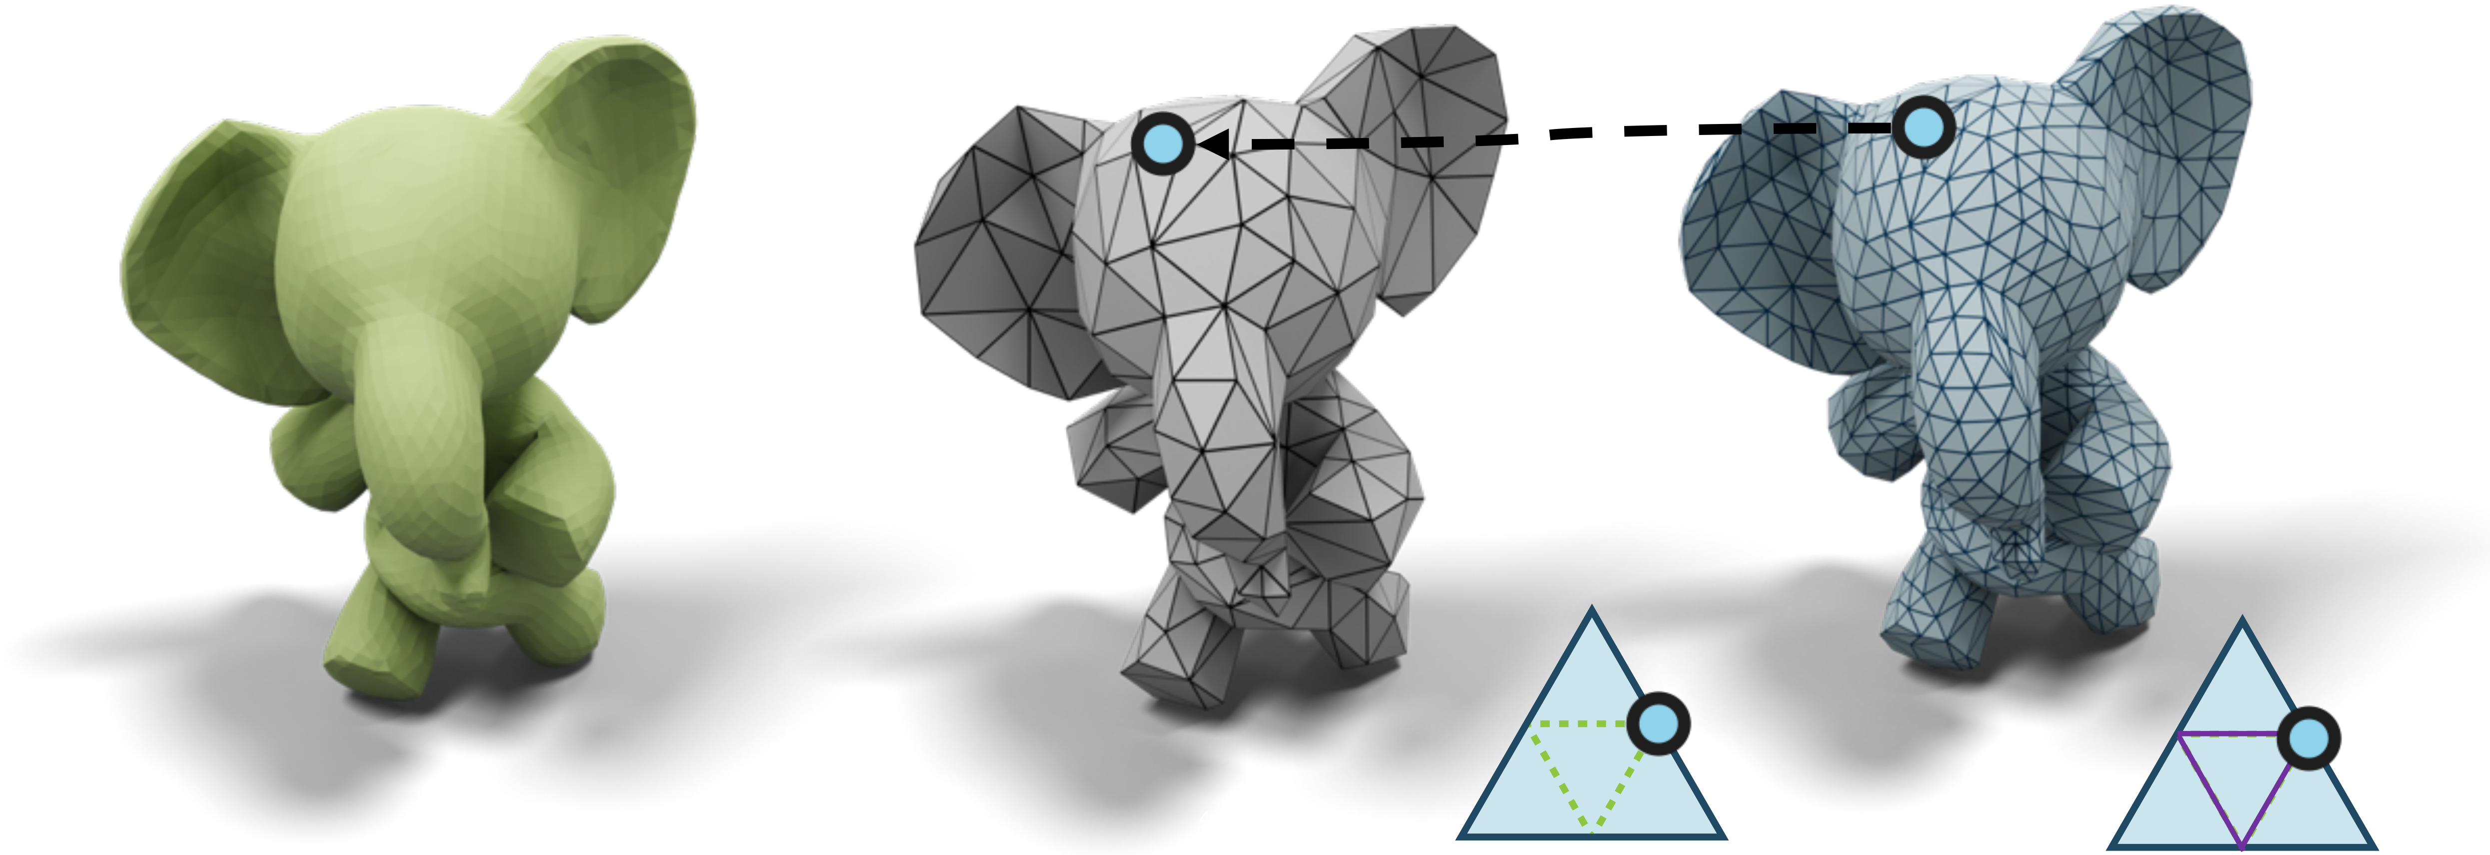

Maintaining Bijective Mapping

- Record barycentric coordinates during subdivision
- Match via parameterization during decimation

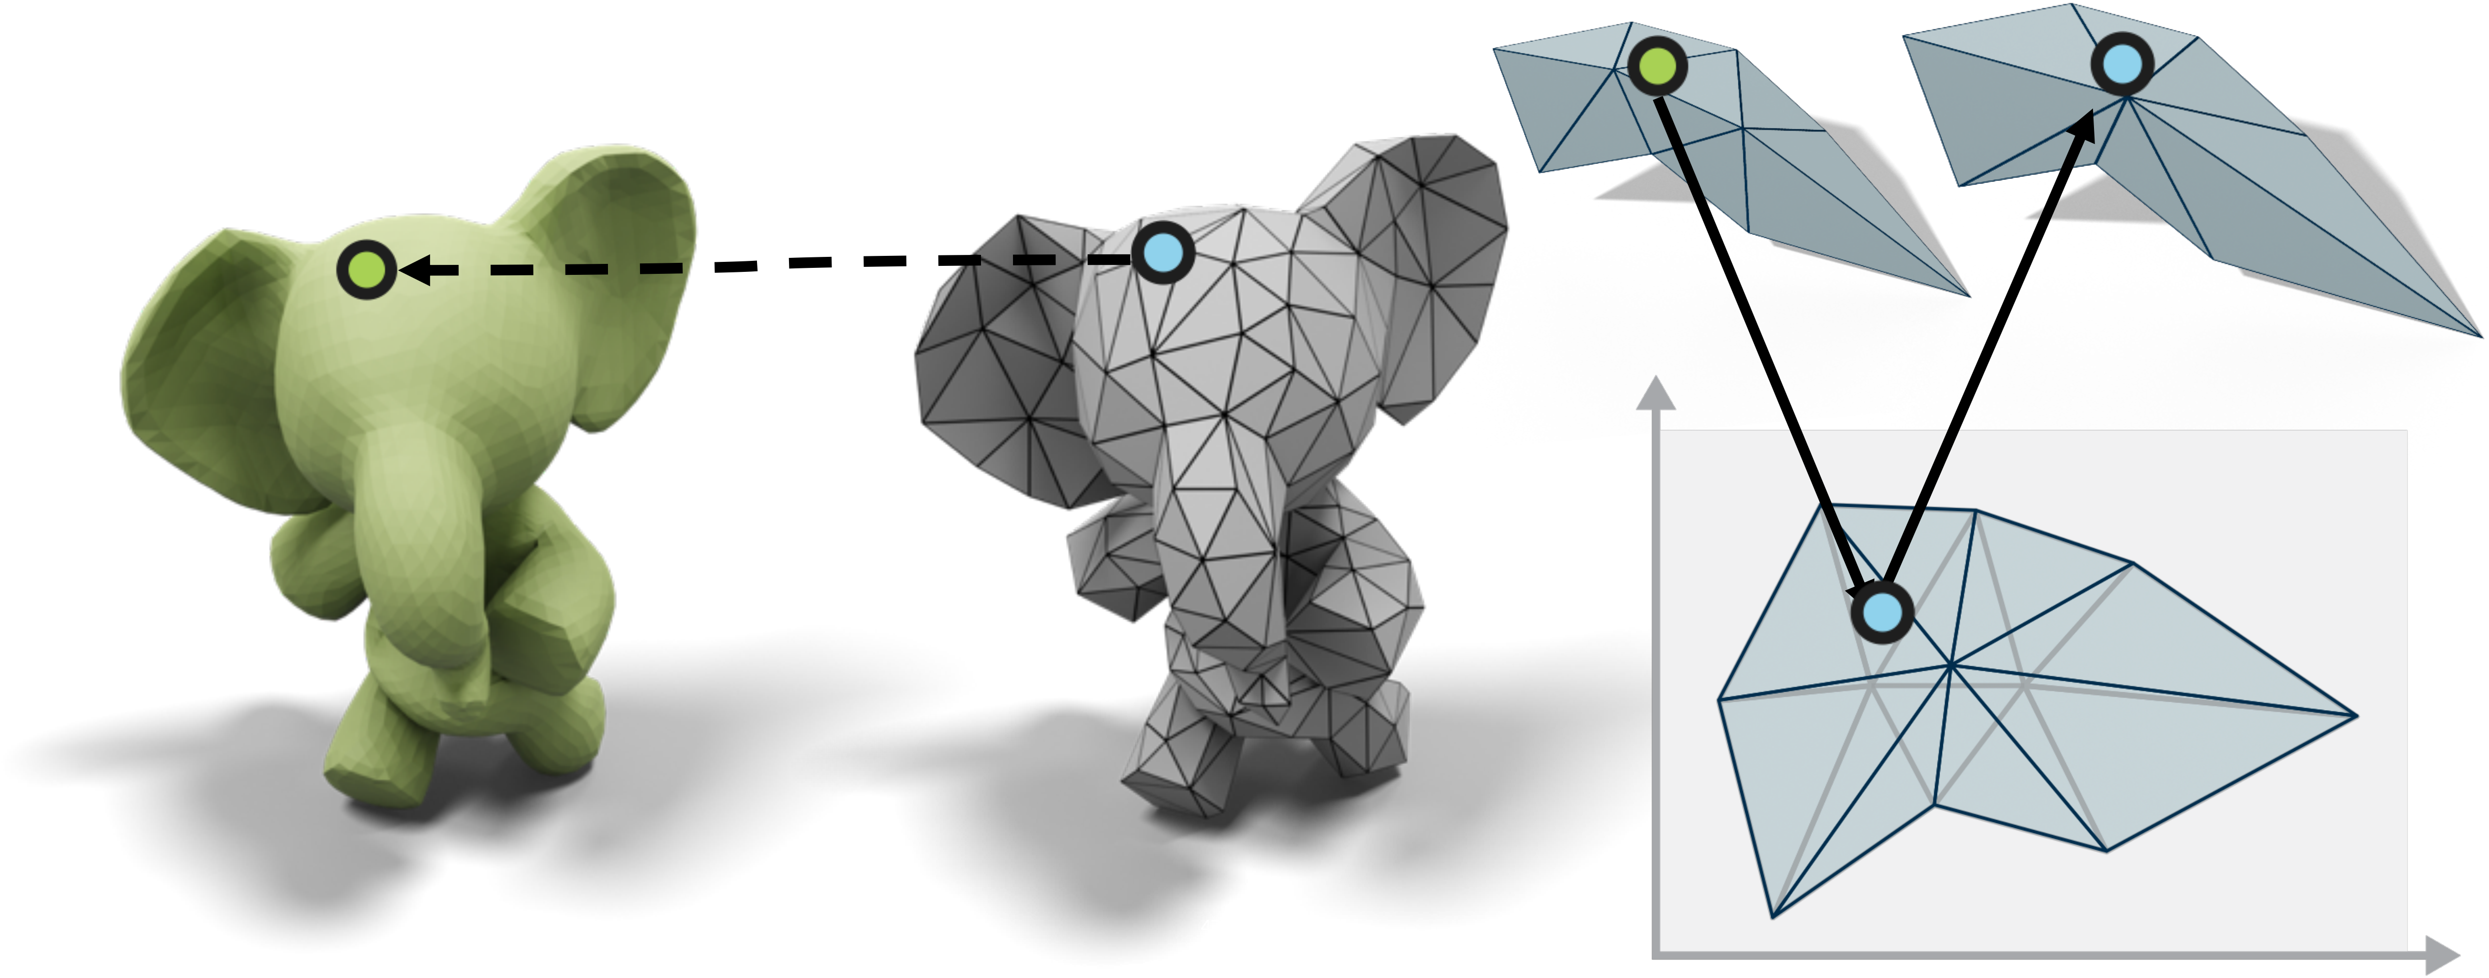

Neural Subdivision

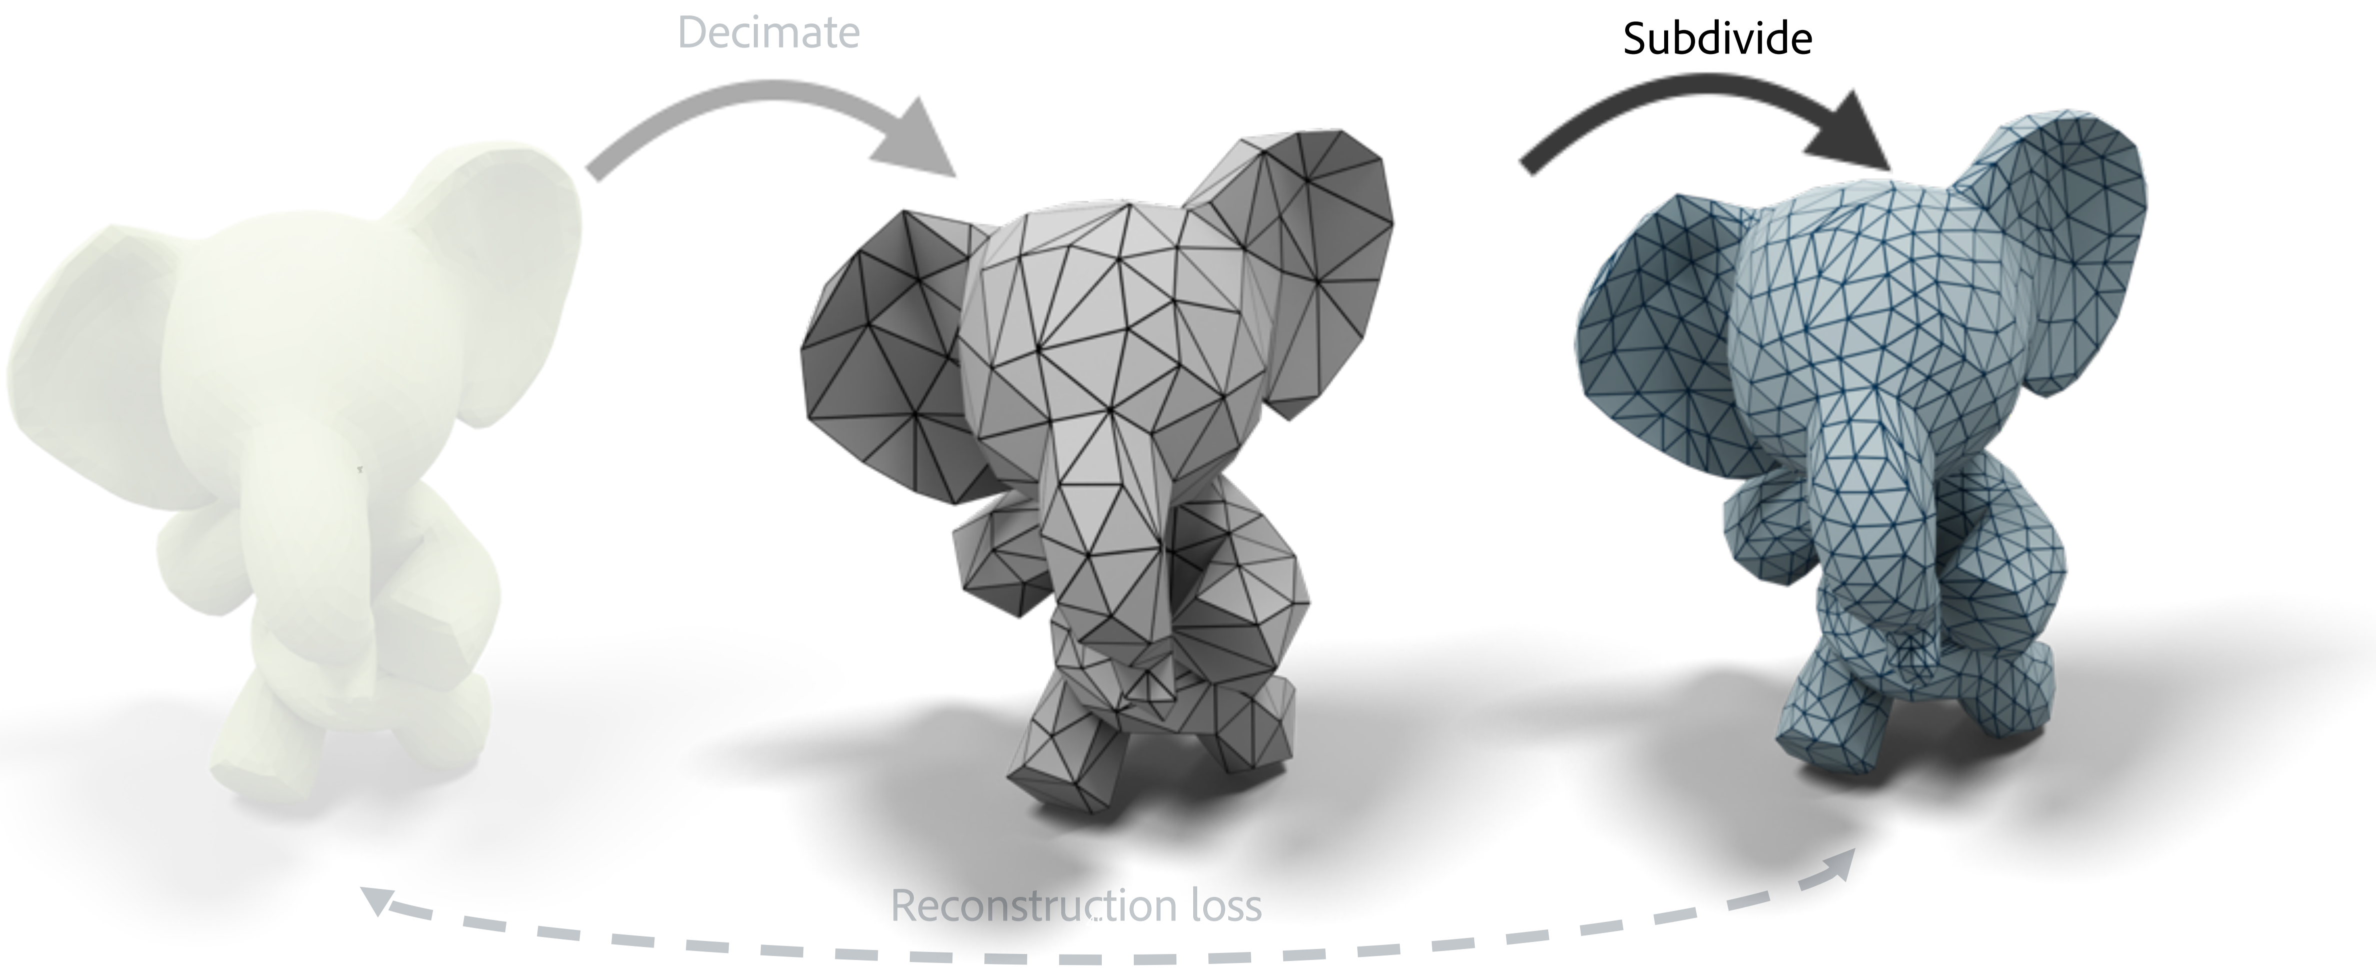

Neural Subdivision

- Triangle Split (mid-edge)

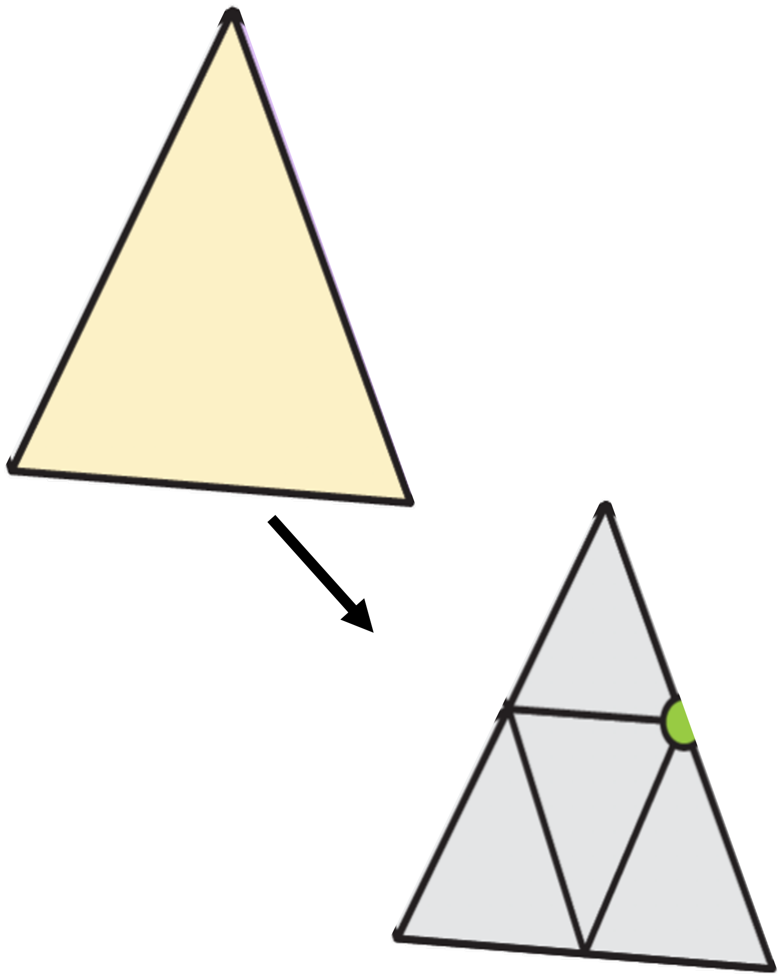

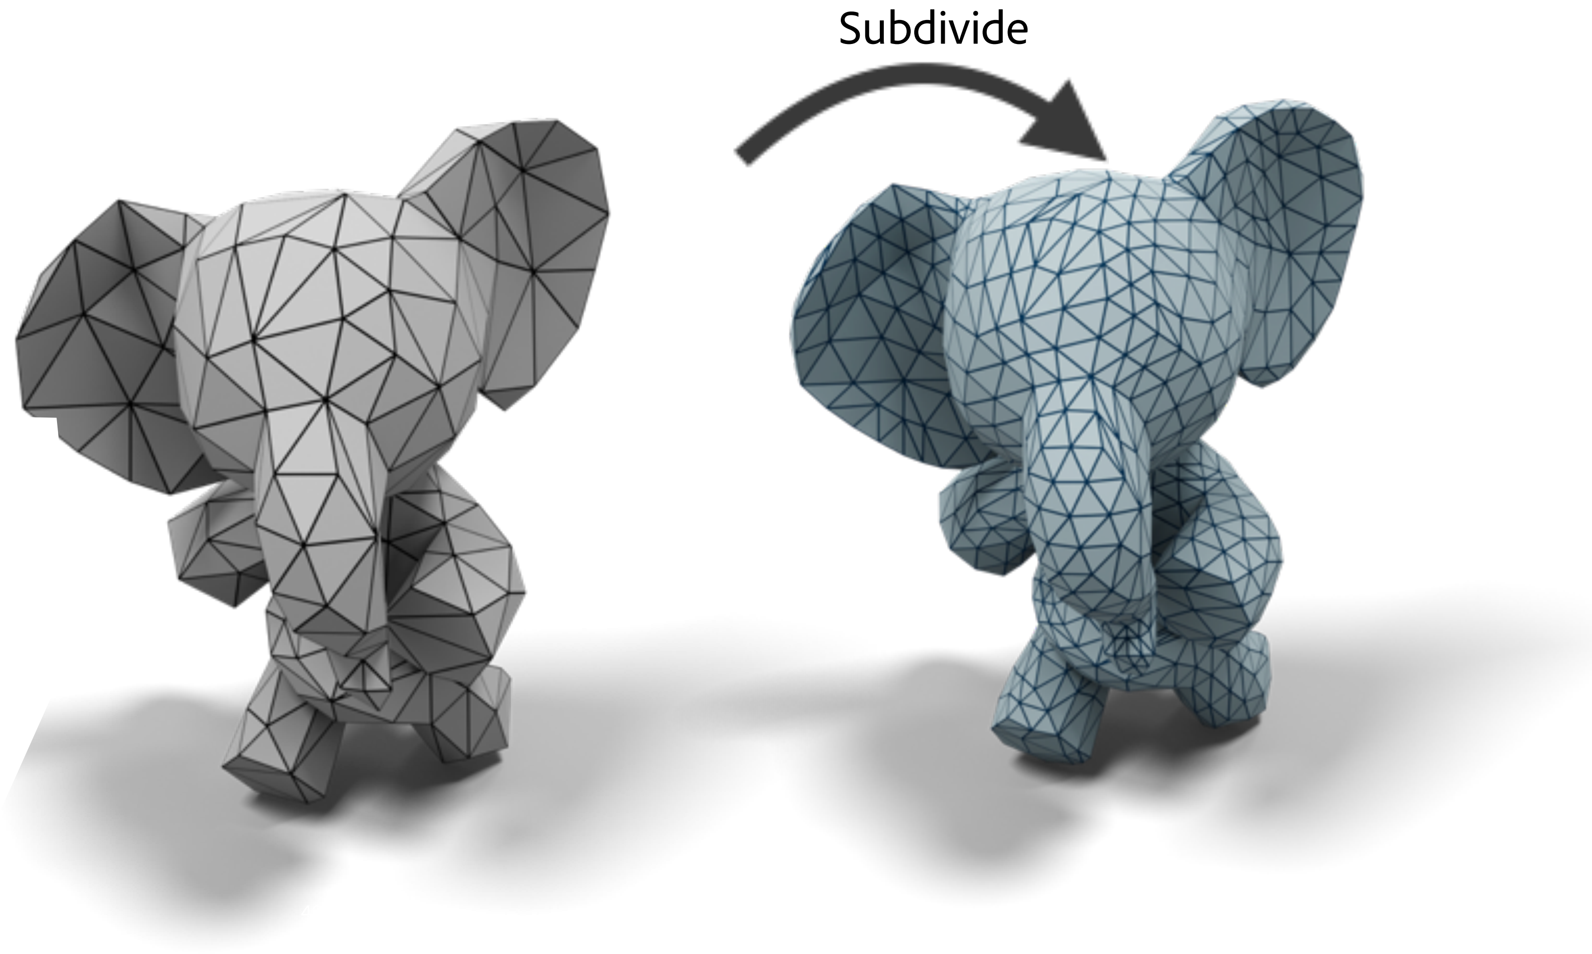

Neural Subdivision

- Triangle Split (mid-edge)

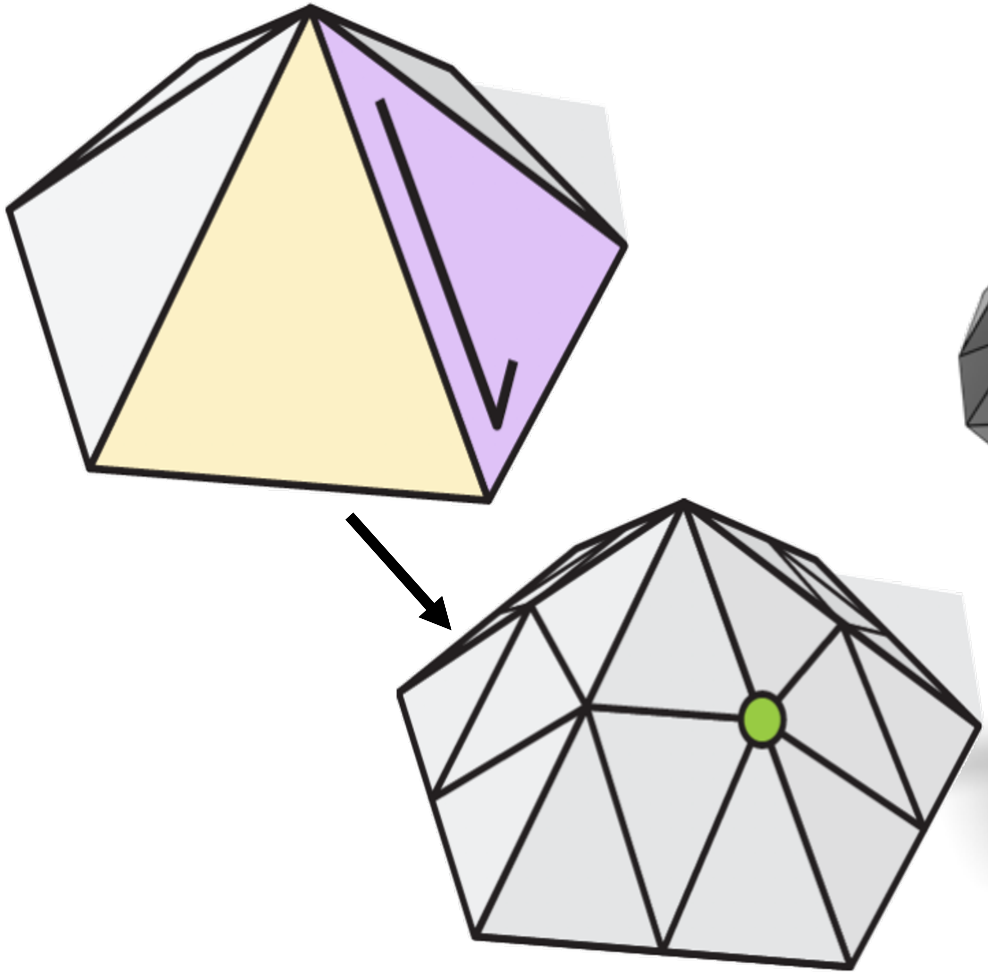

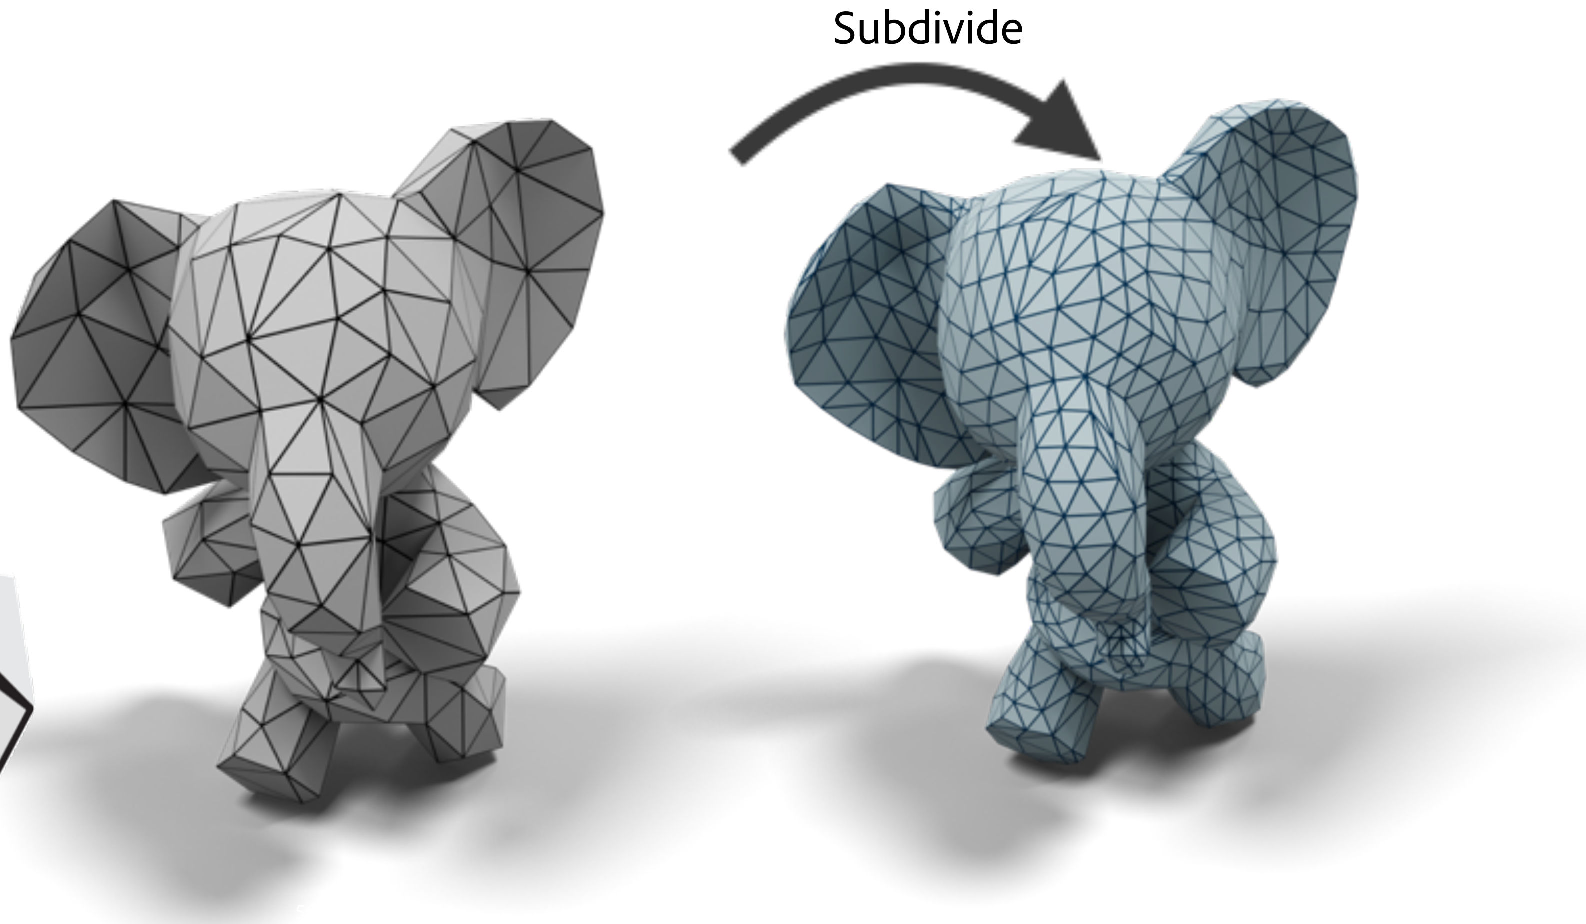

Neural Subdivision

- Triangle Split (mid-edge)
- Set vertex positions via neural network

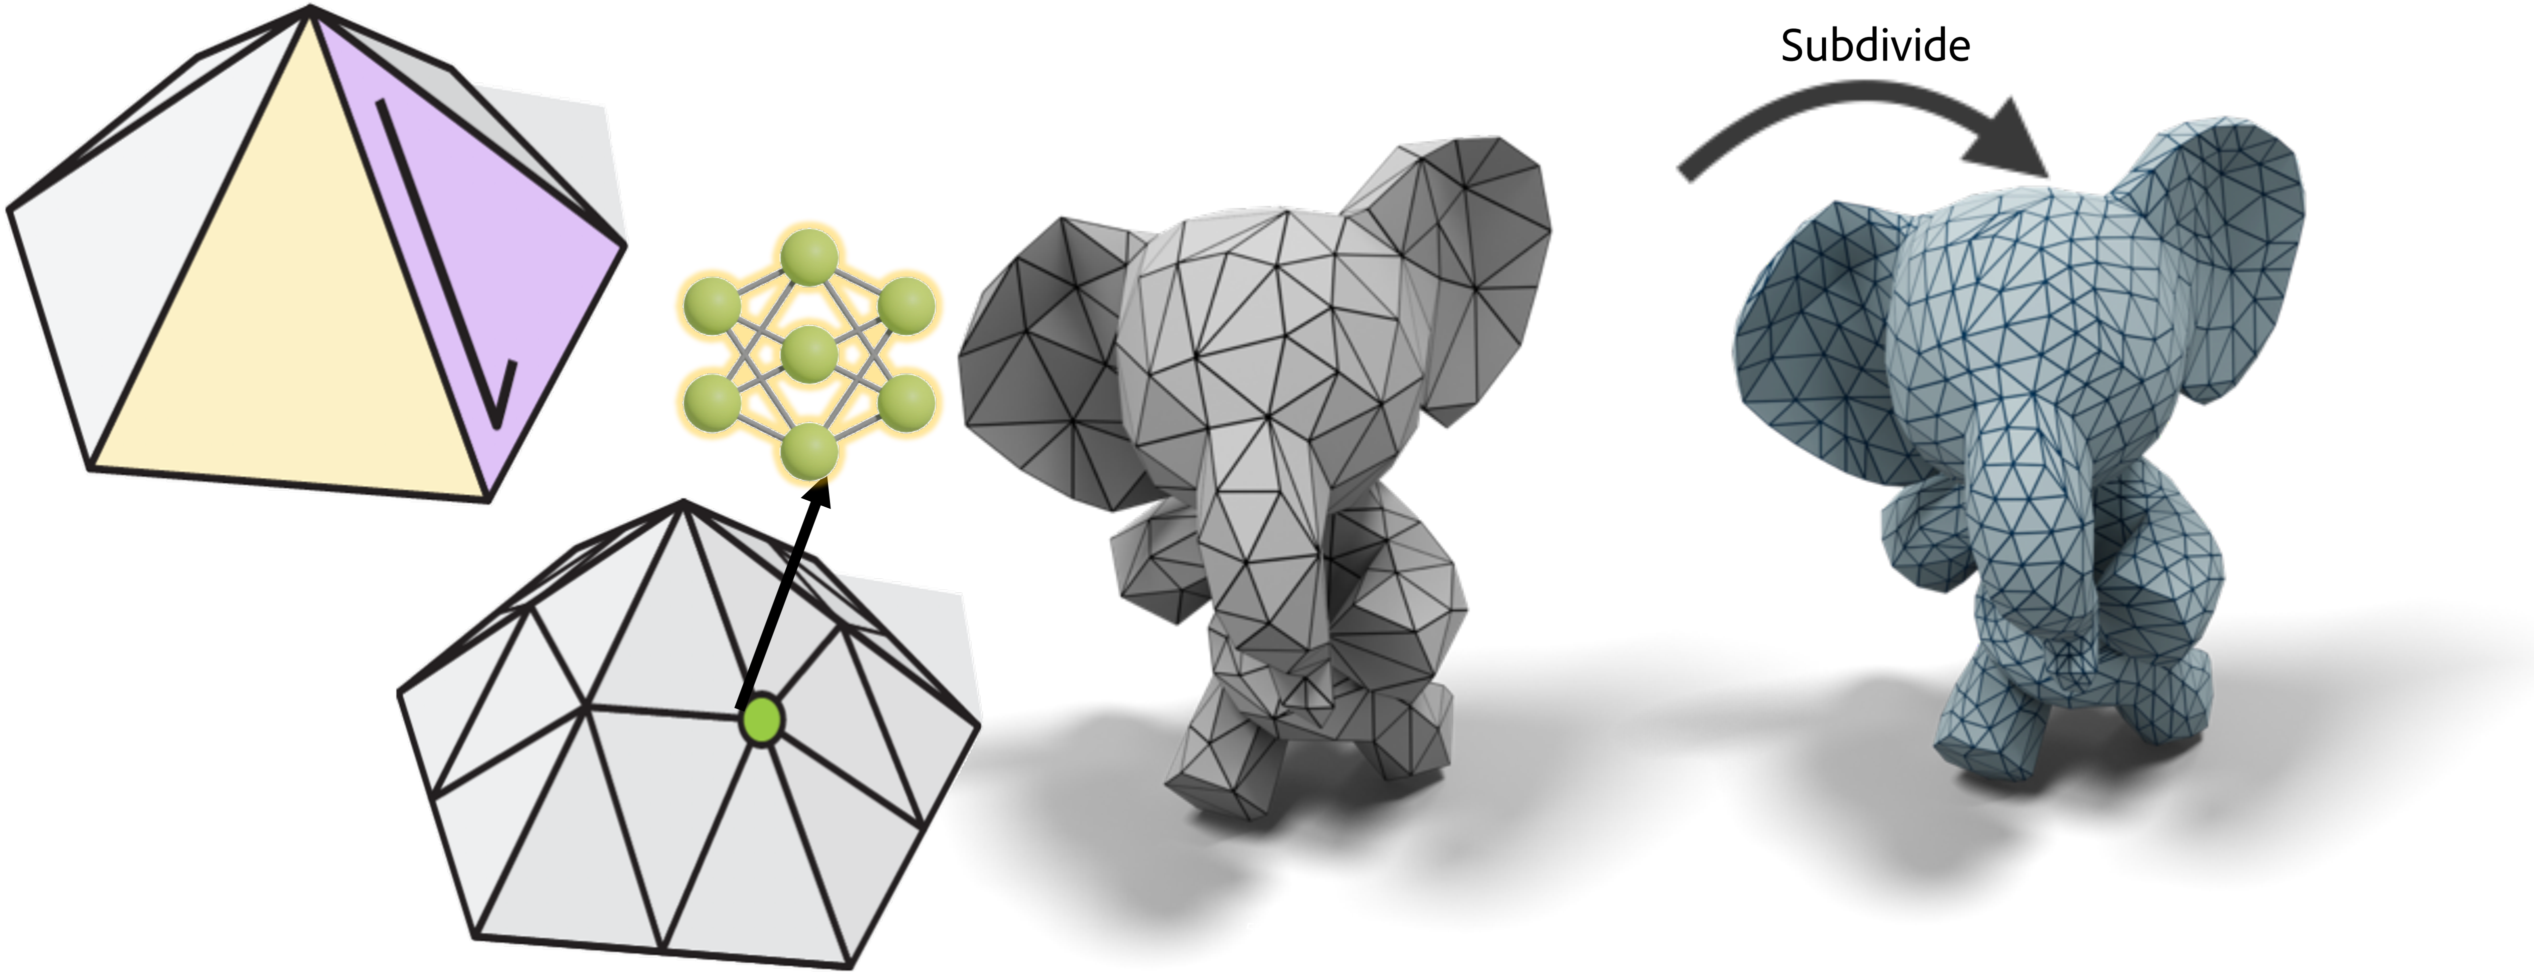

Architecture

- Half-flap: directed edges and two adjacent triangles
  - Fixed Dimensions
  - Canonical Ordering

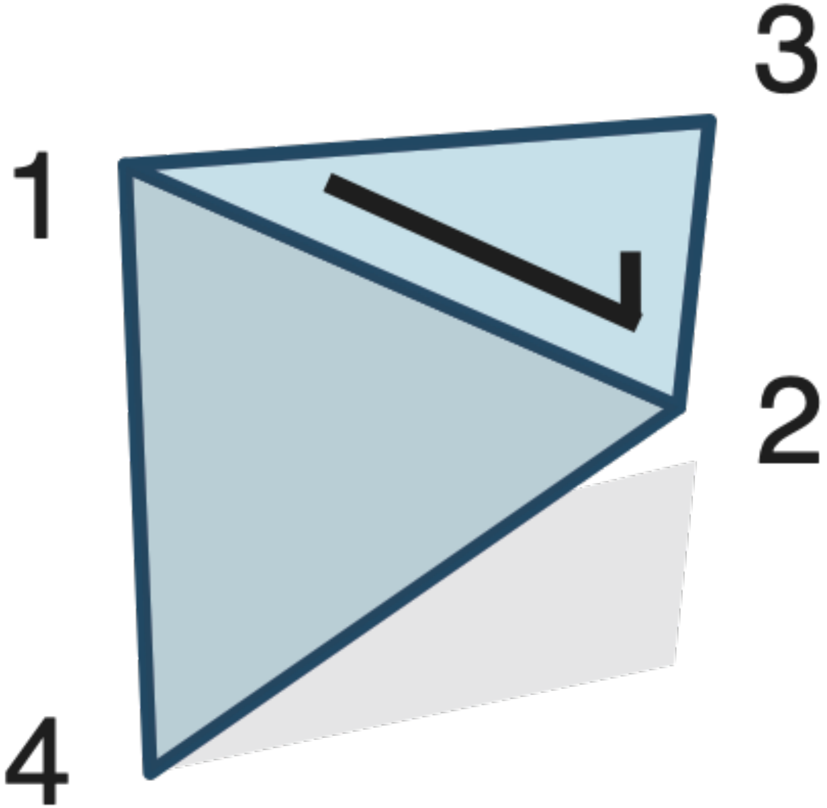

Architecture

- Half-flap: directed edges and two adjacent triangles
- Represent (differential) geometry in flap's local coordinate frame

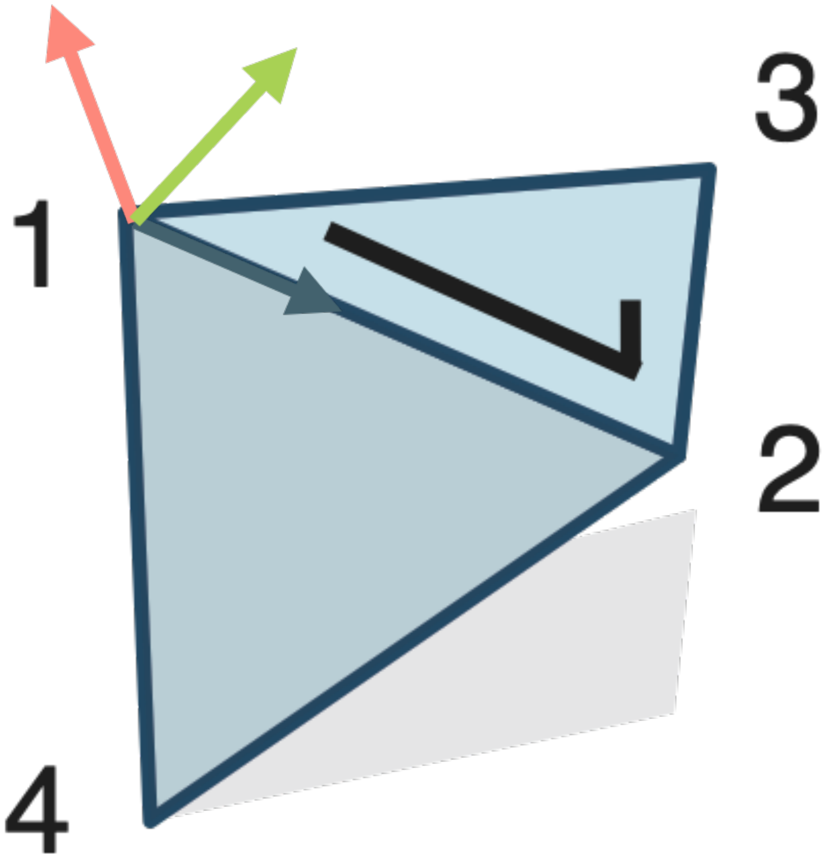

Pipeline

- Initialize per-vertex features

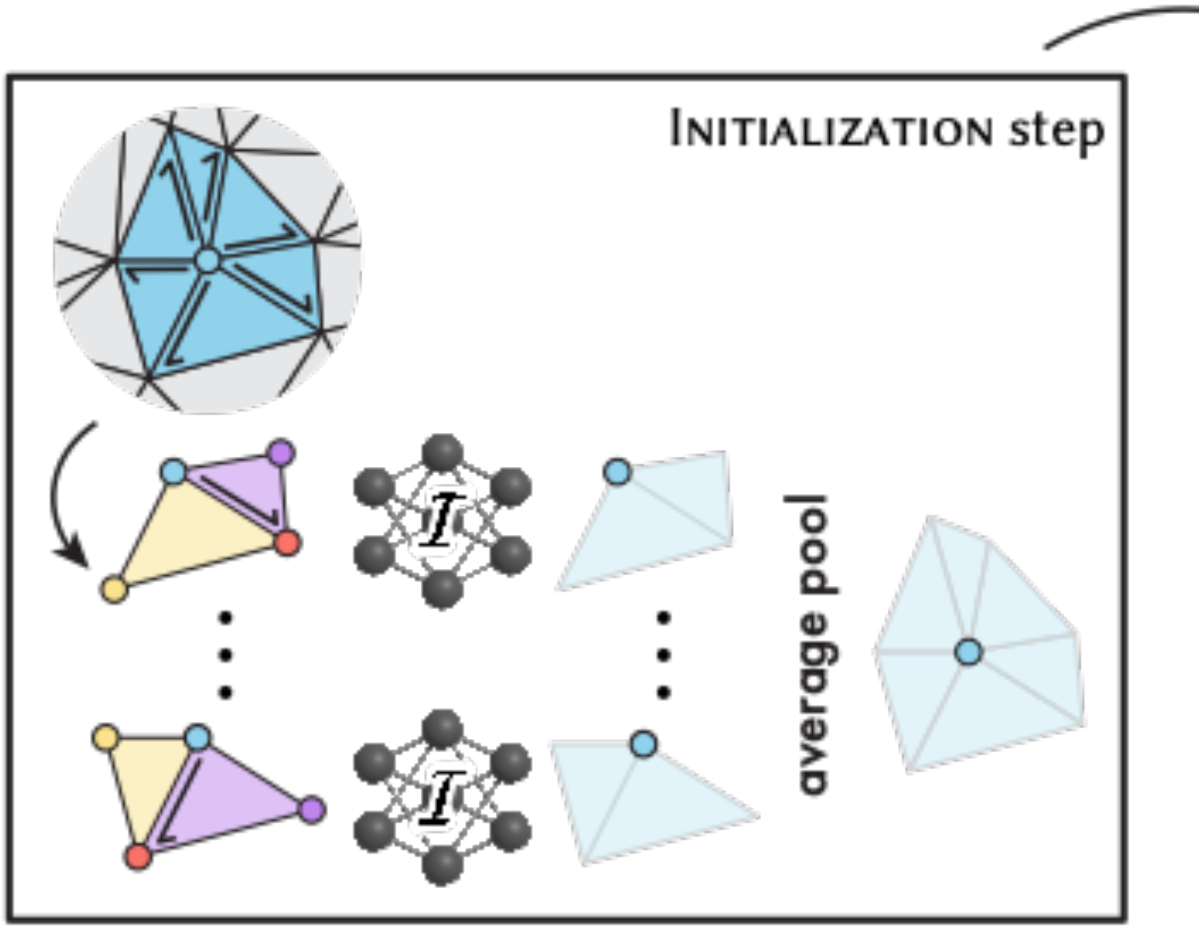

Pipeline

- Initialize per-vertex features
- Iteratively update features and geometry at old & new vertices

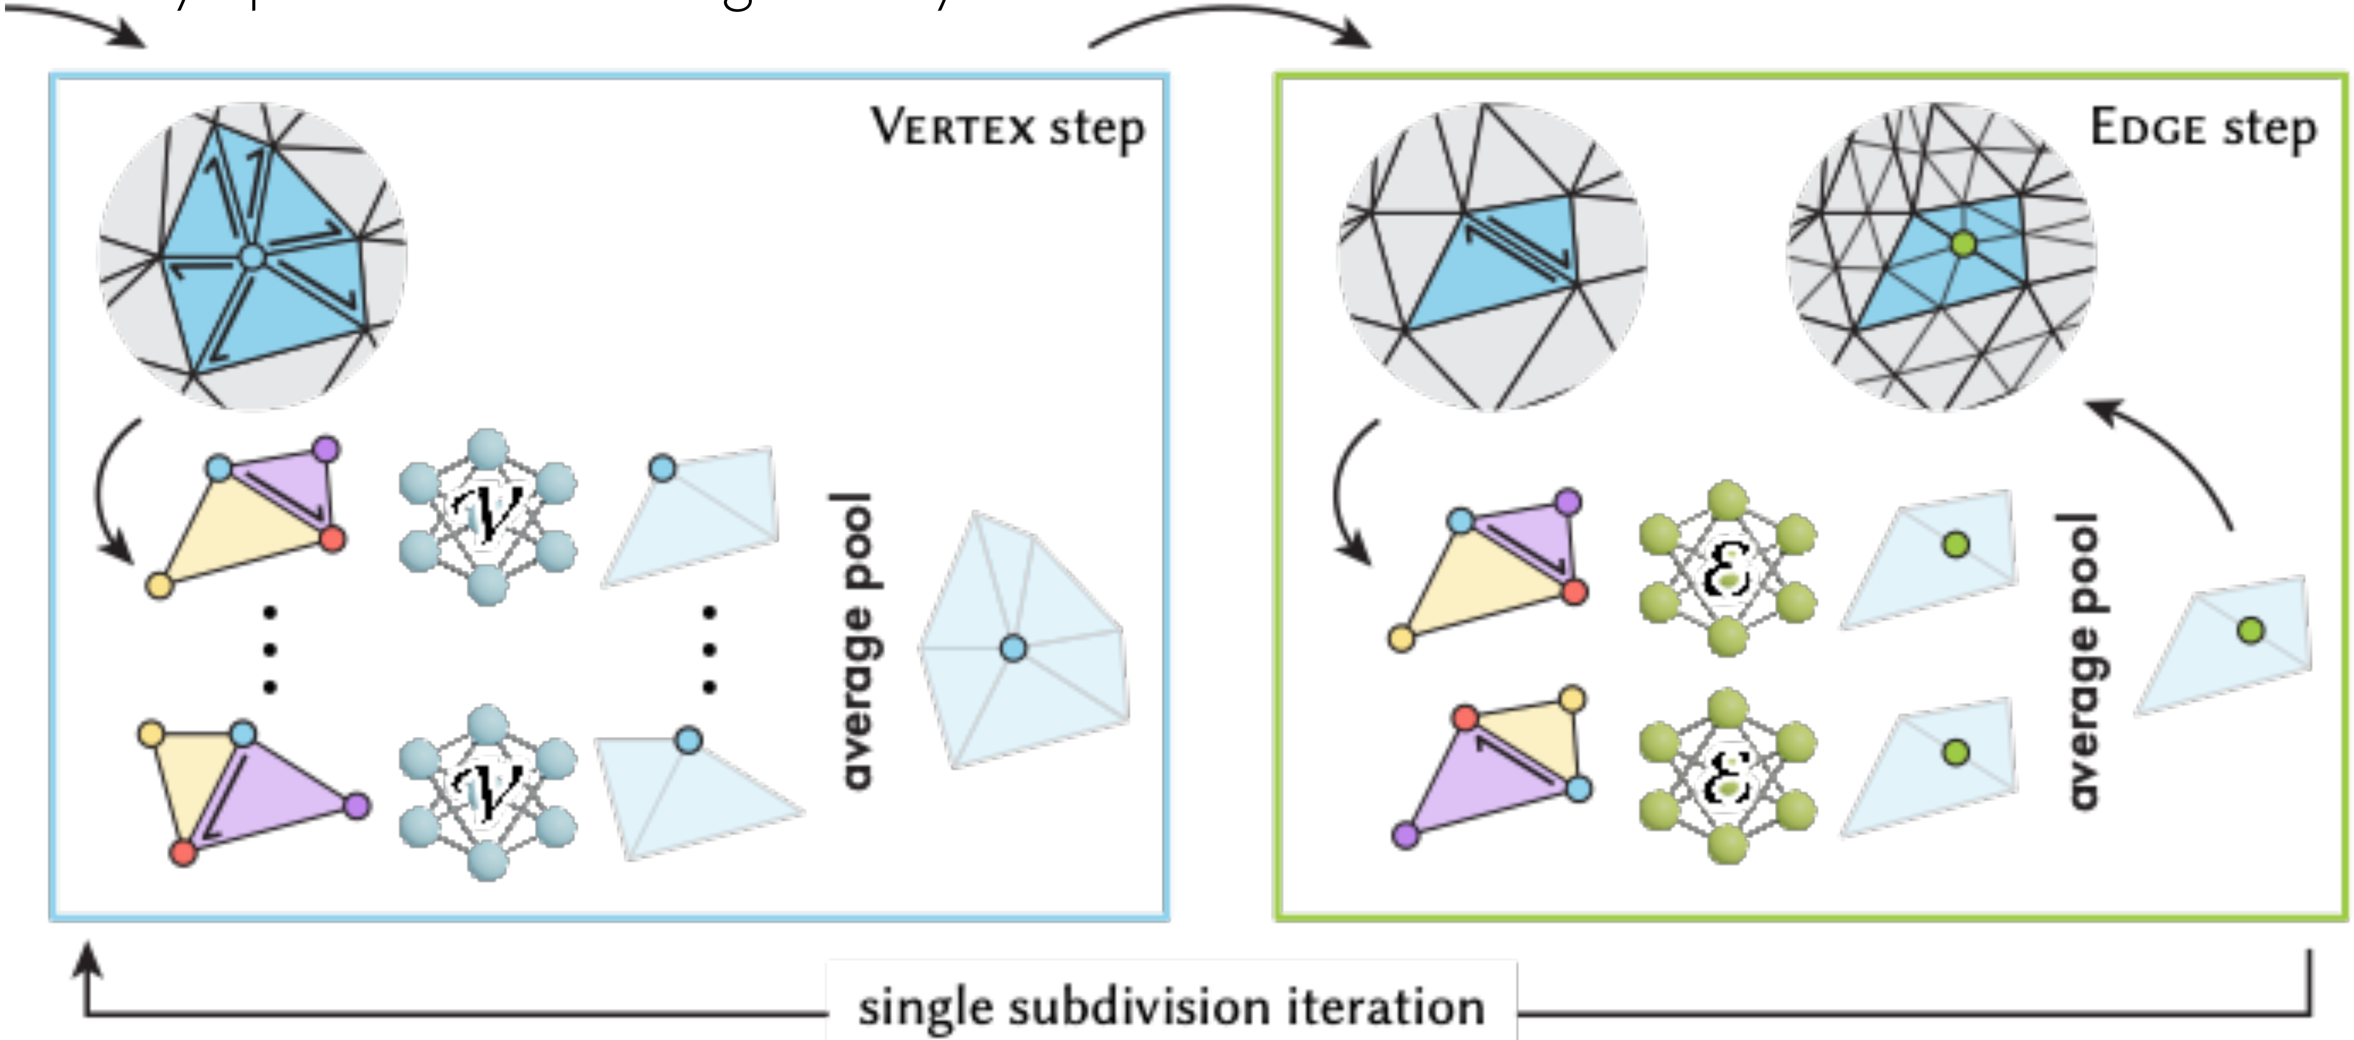

Results: Neural Detail Transfer

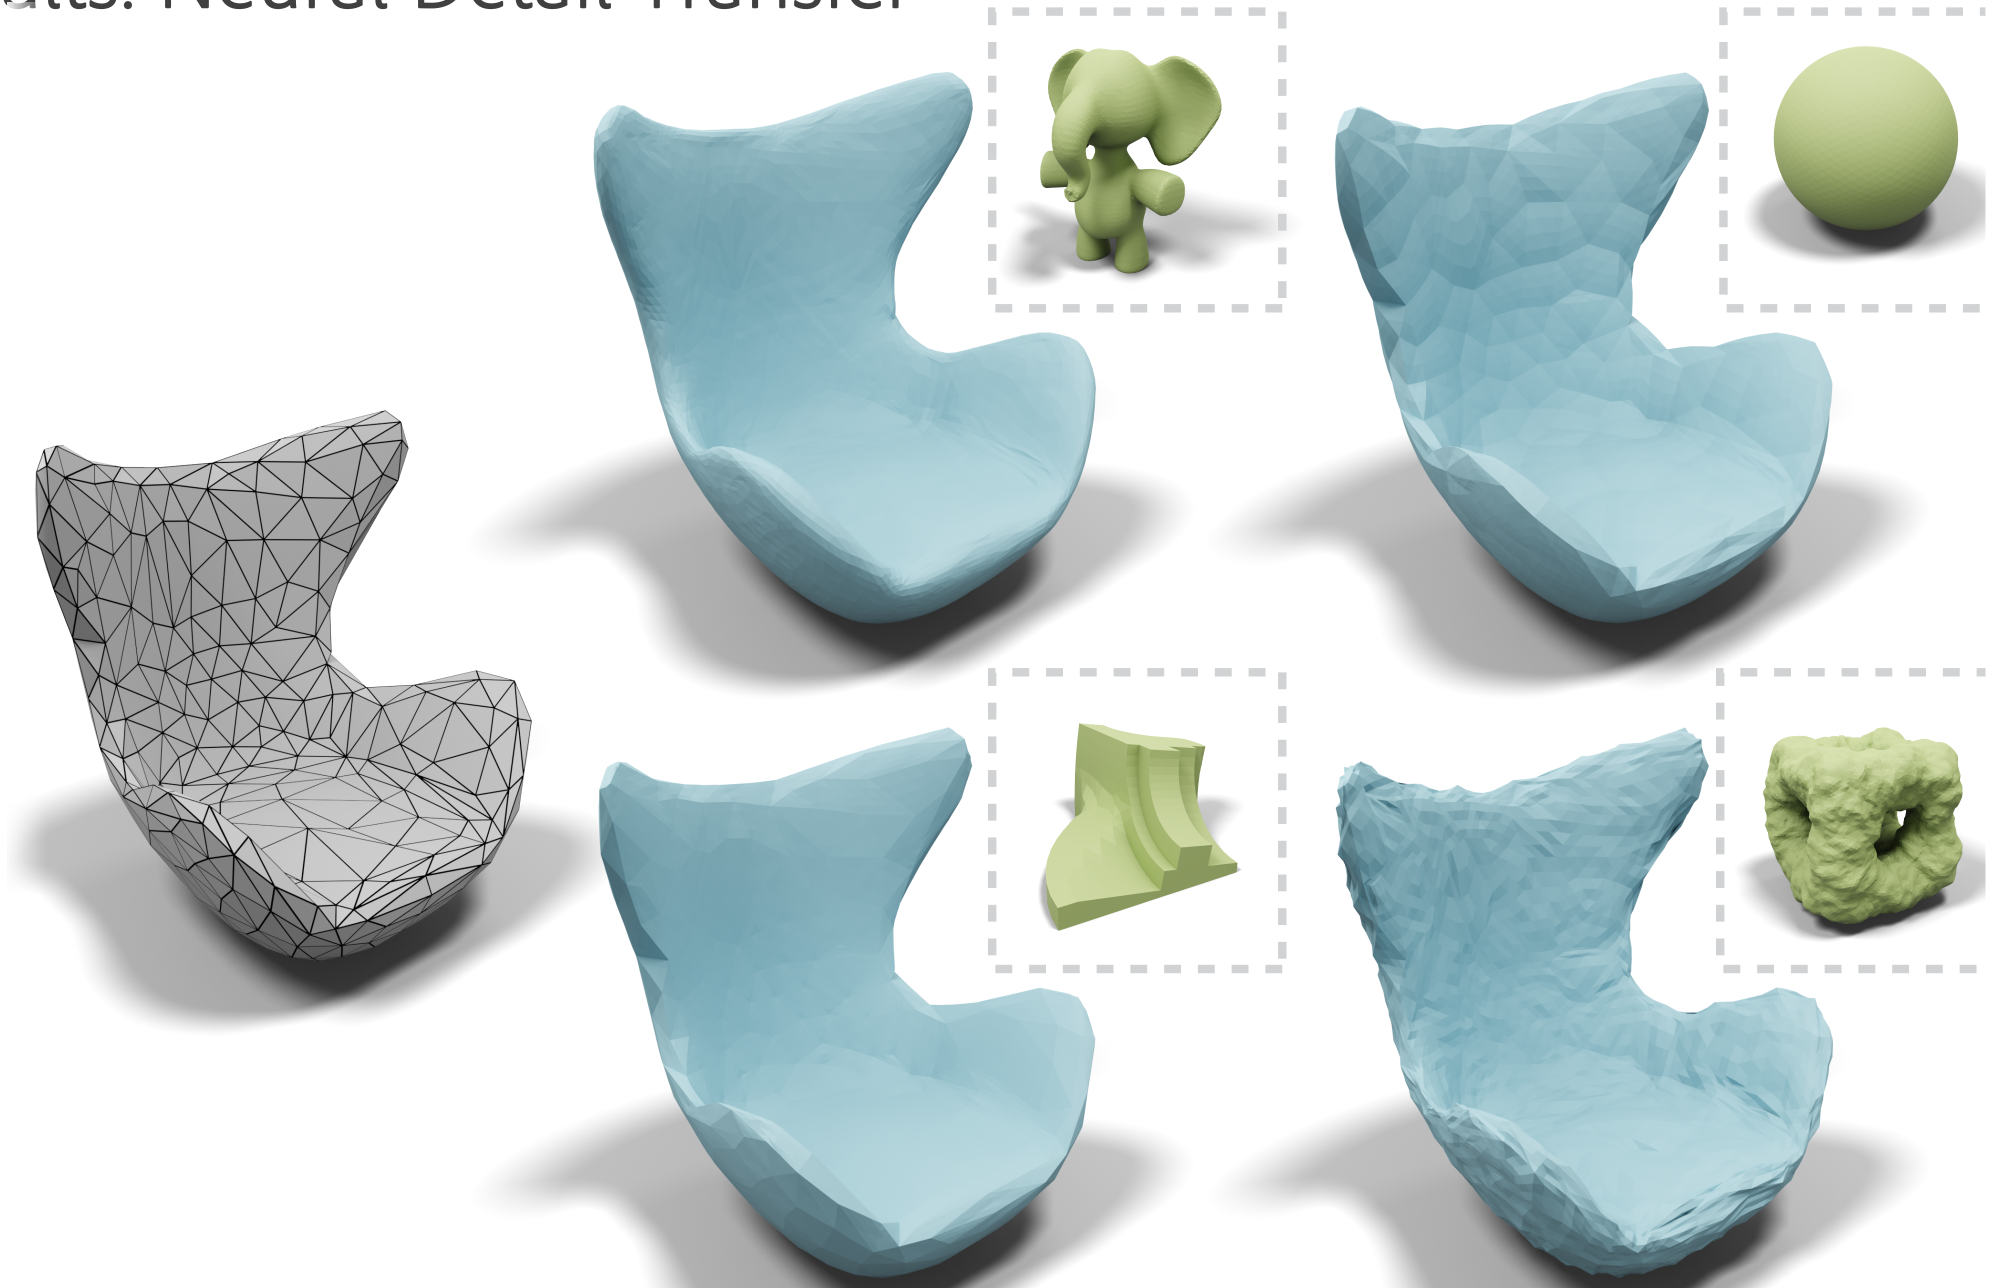

Results: Neural Detail Upsampling

- Neural subdivision trained on a single example

Training Example

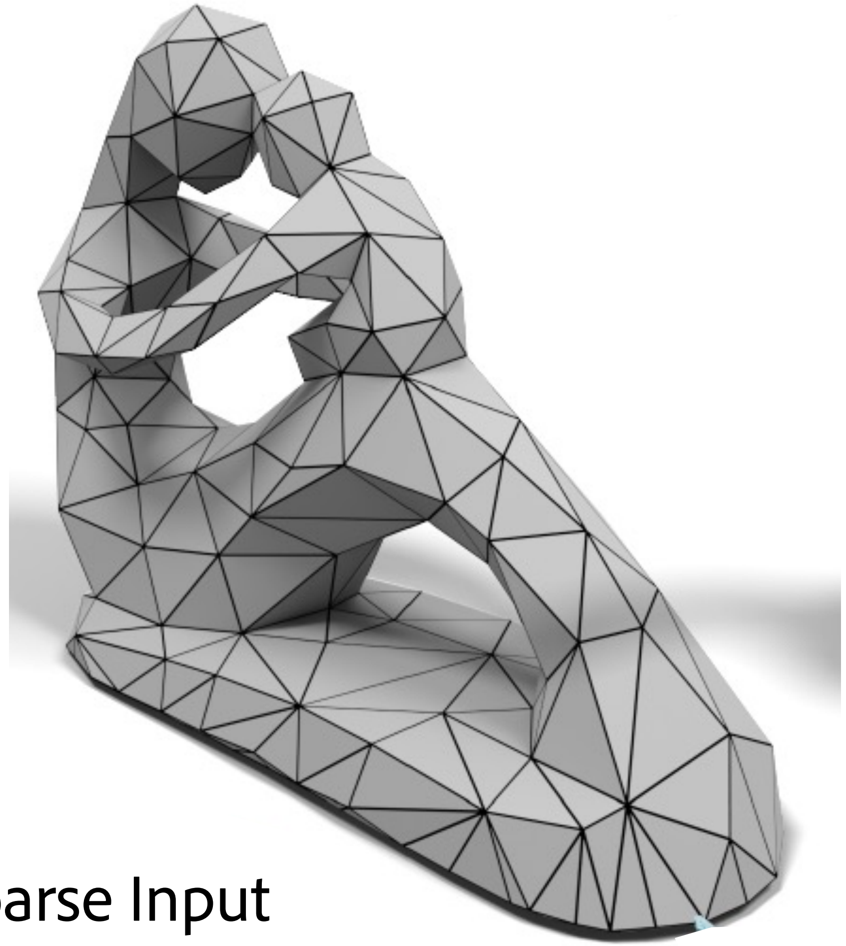

Coarse Input

Subdivided Output

Results: Neural Detail Upsampling

- Neural subdivision trained on a single example

Training Example

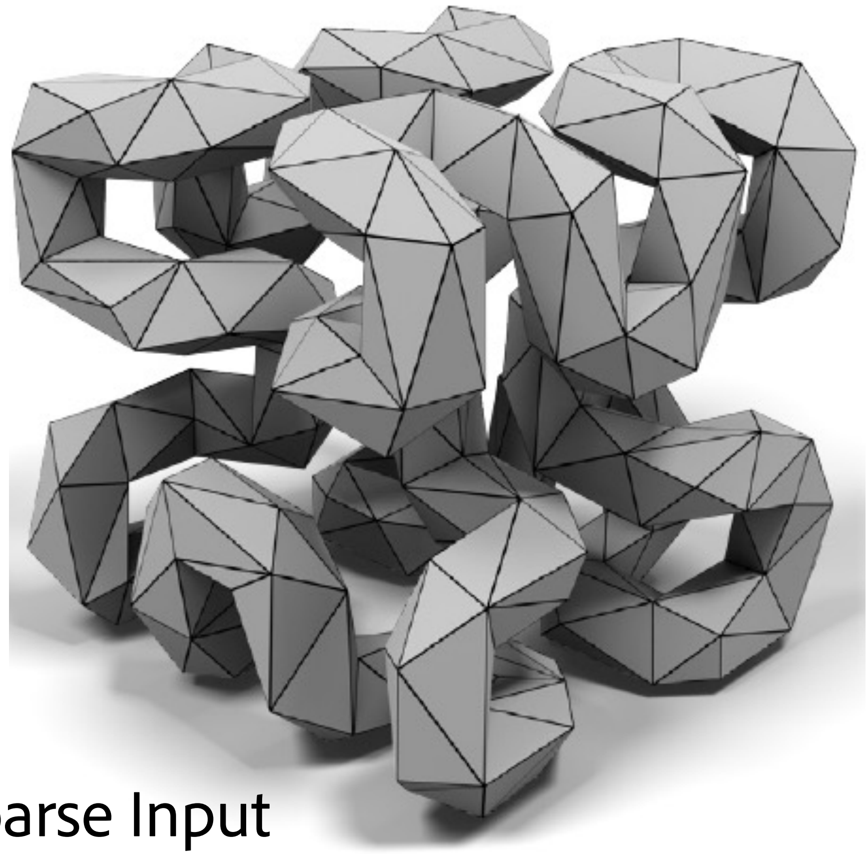

Coarse Input

Subdivided Output

Detail Transfer and Synthesis

- Hallucinating details with complex topology

Input

Target Style

Detail Transfer and Synthesis

- Hallucinating details with complex topology

Input

Target Style

Output

DÉCOR-GAN Neural Network

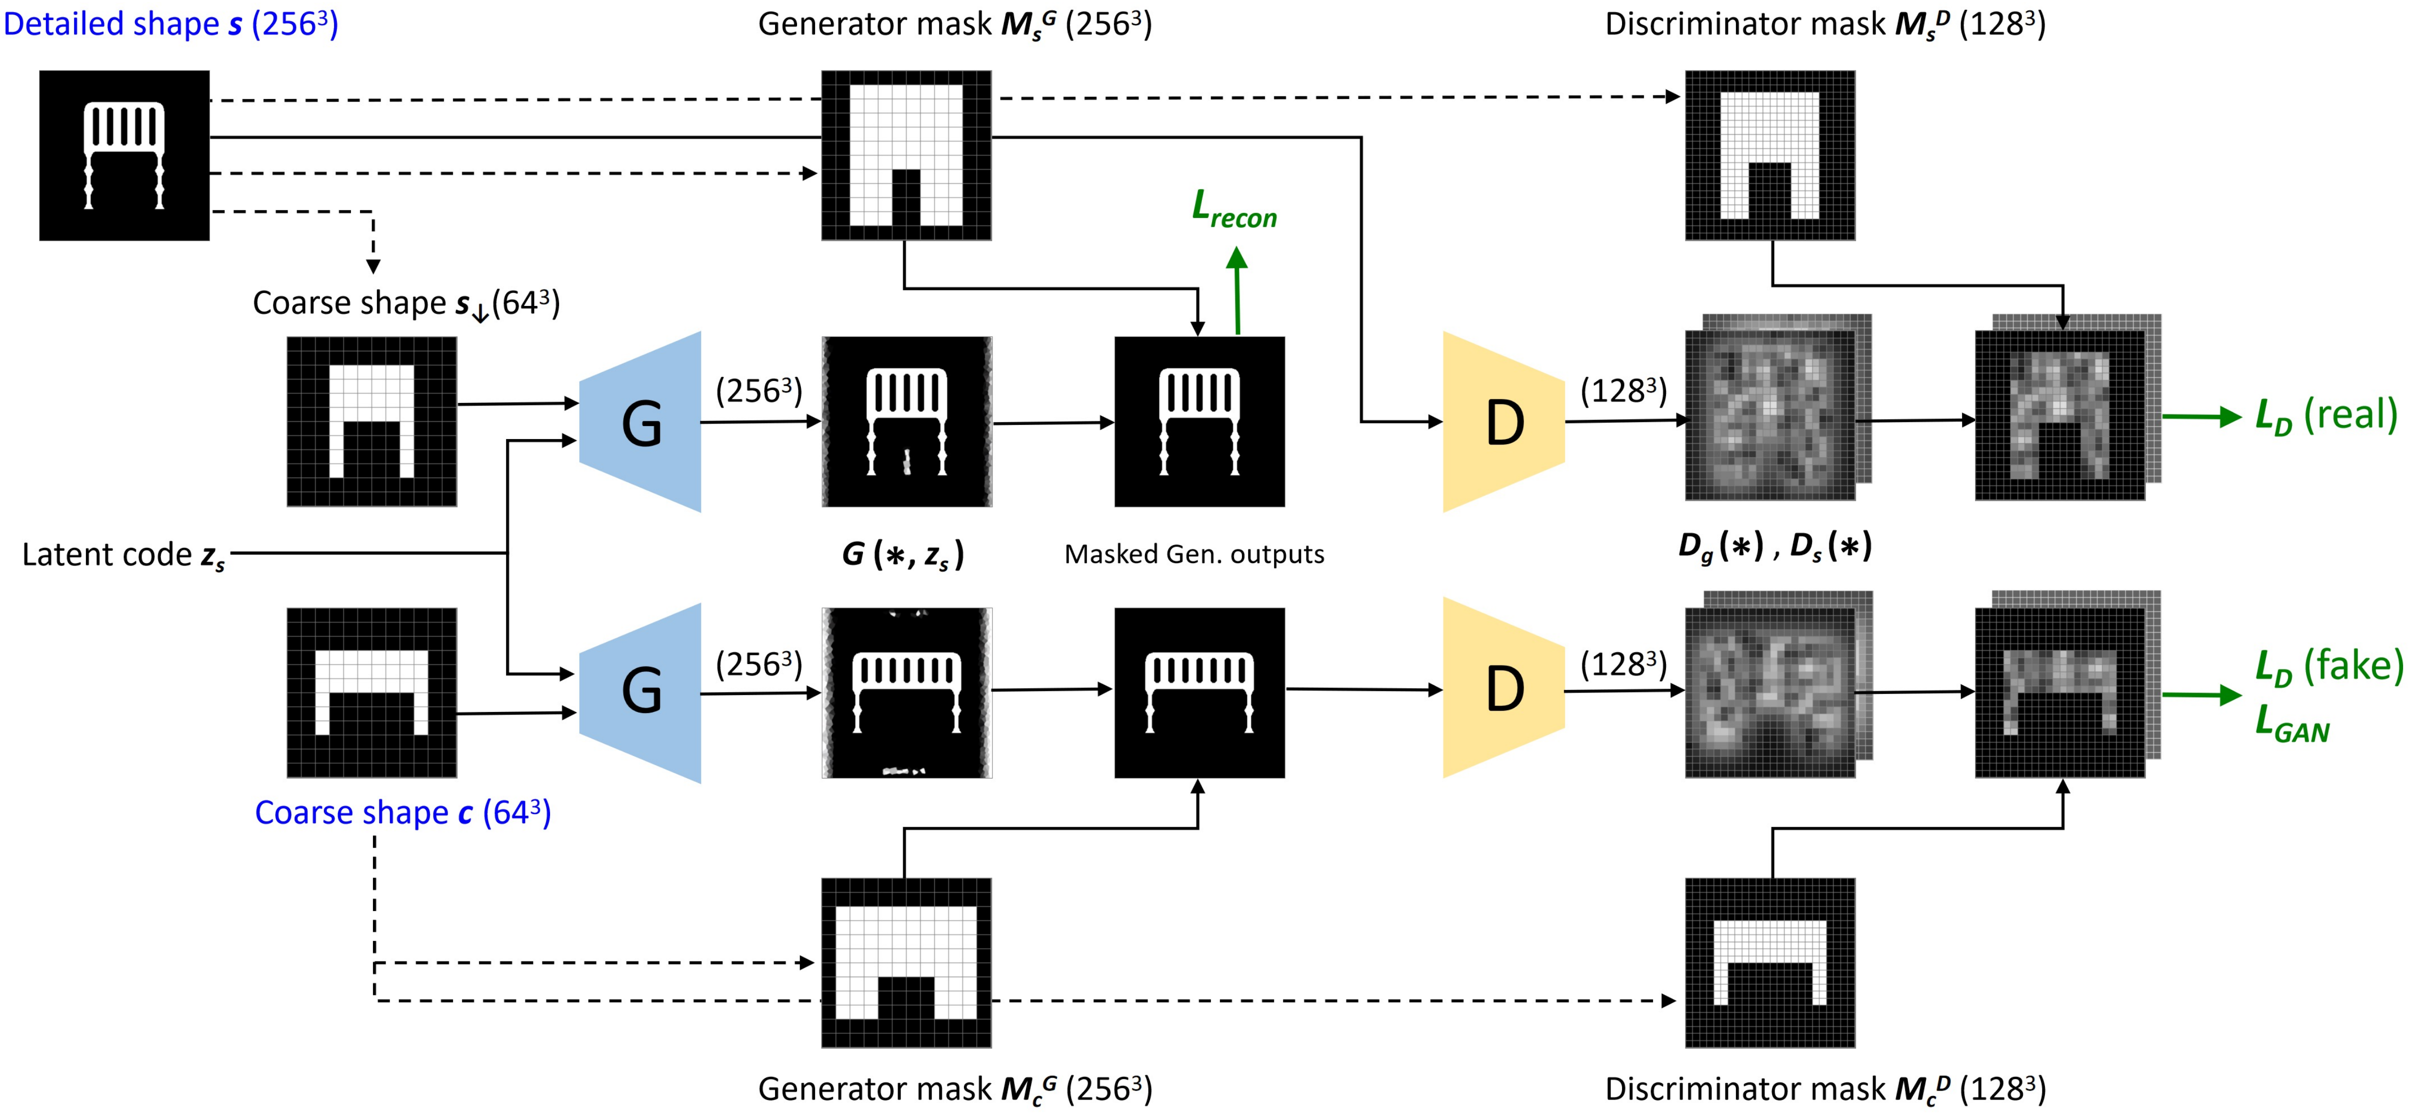

DÉCOR-GAN Neural Network

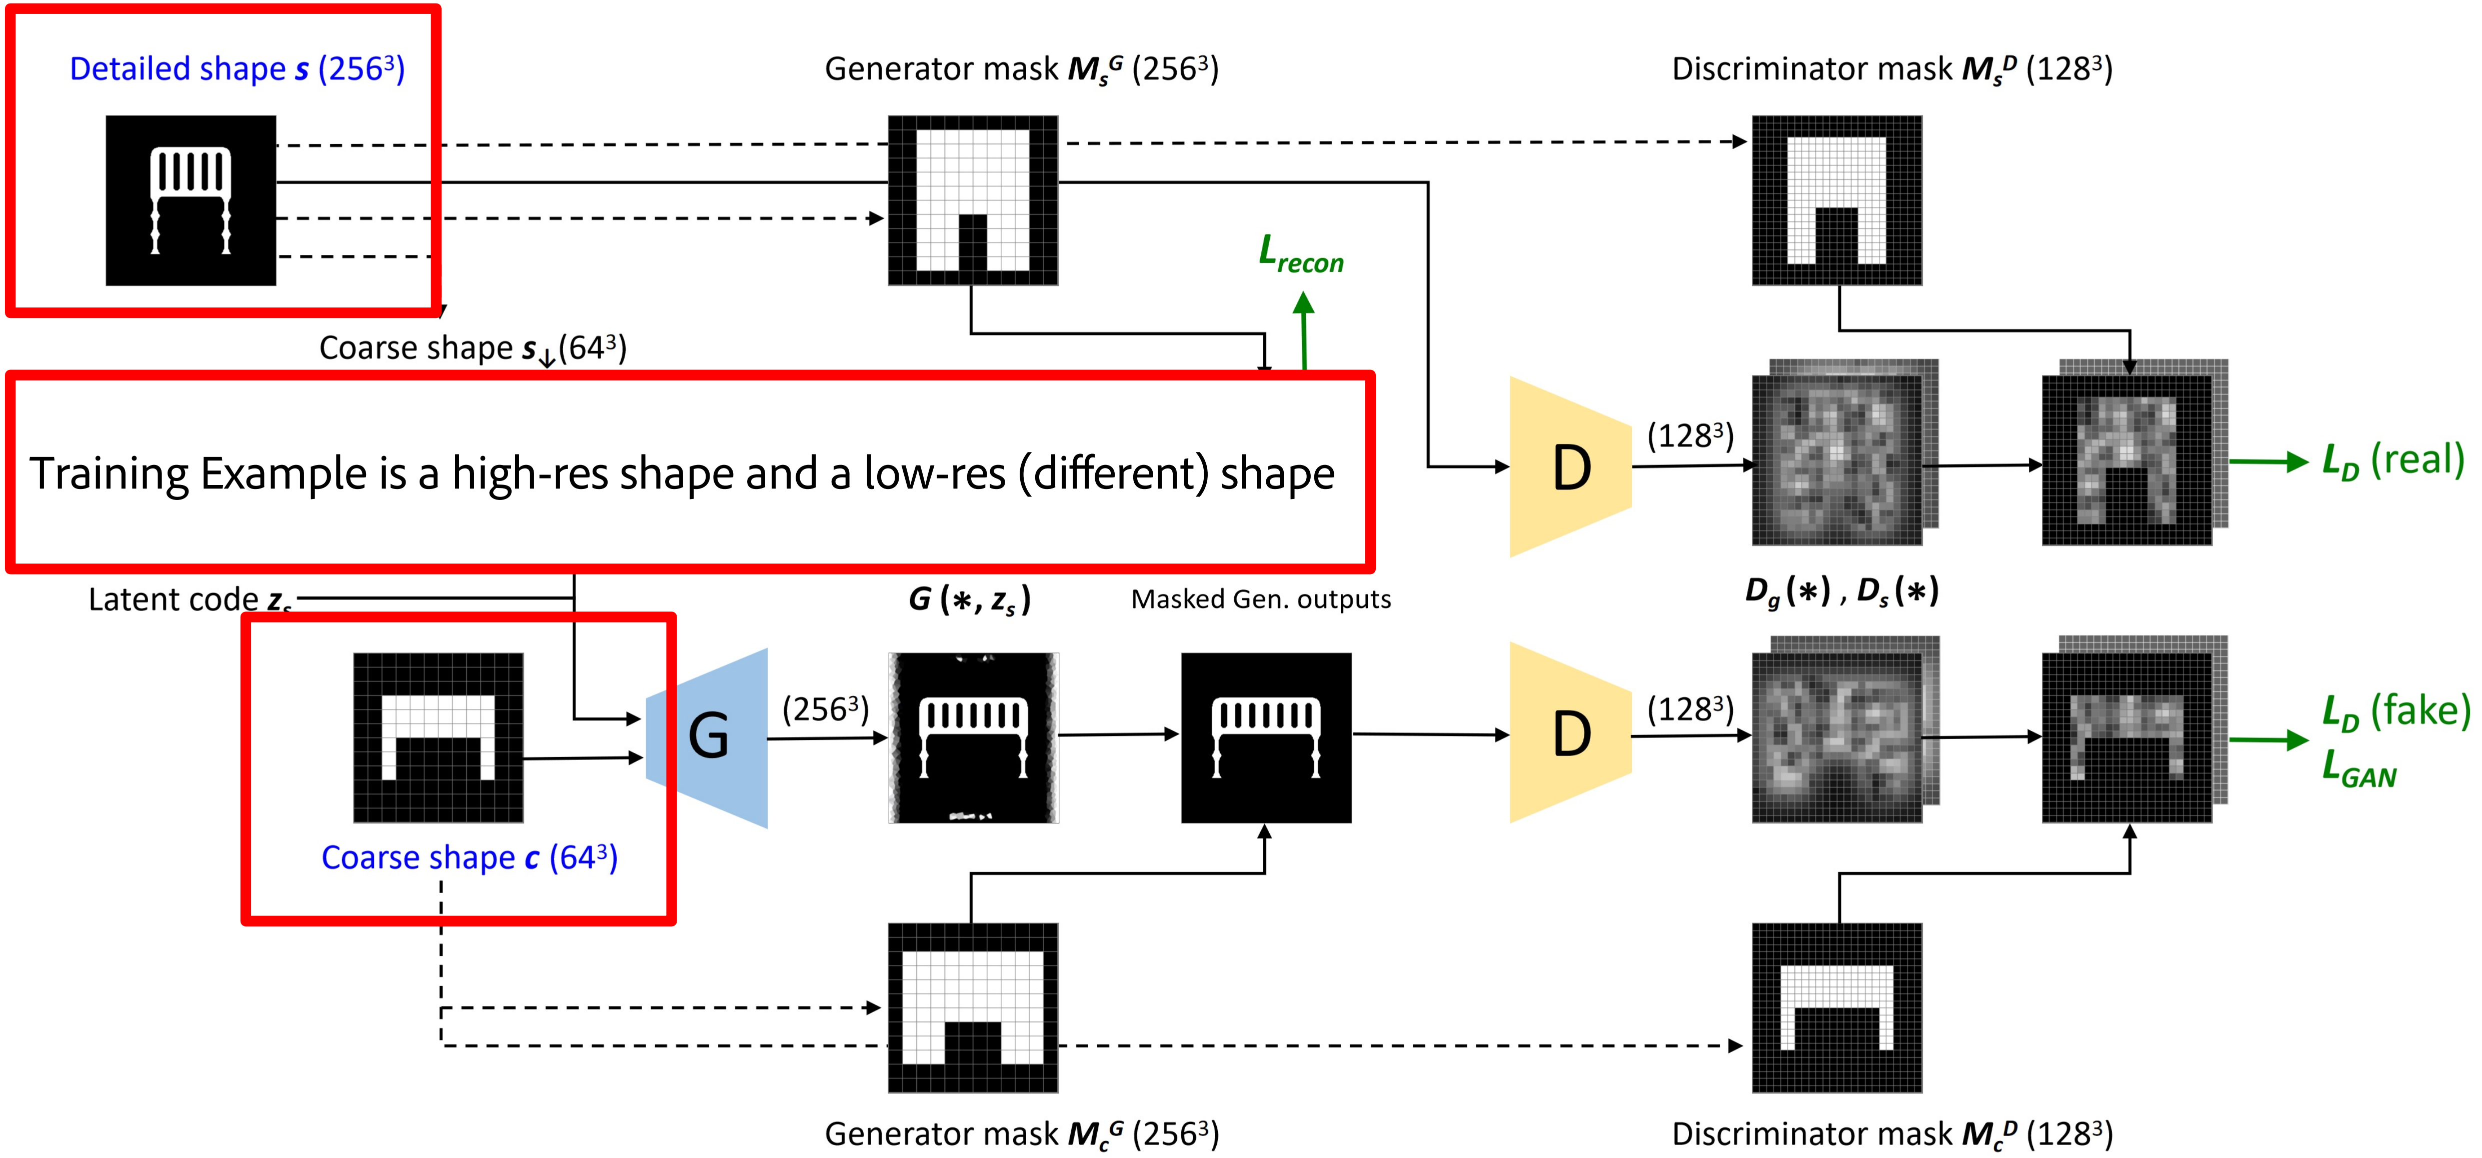

DÉCOR-GAN Neural Network

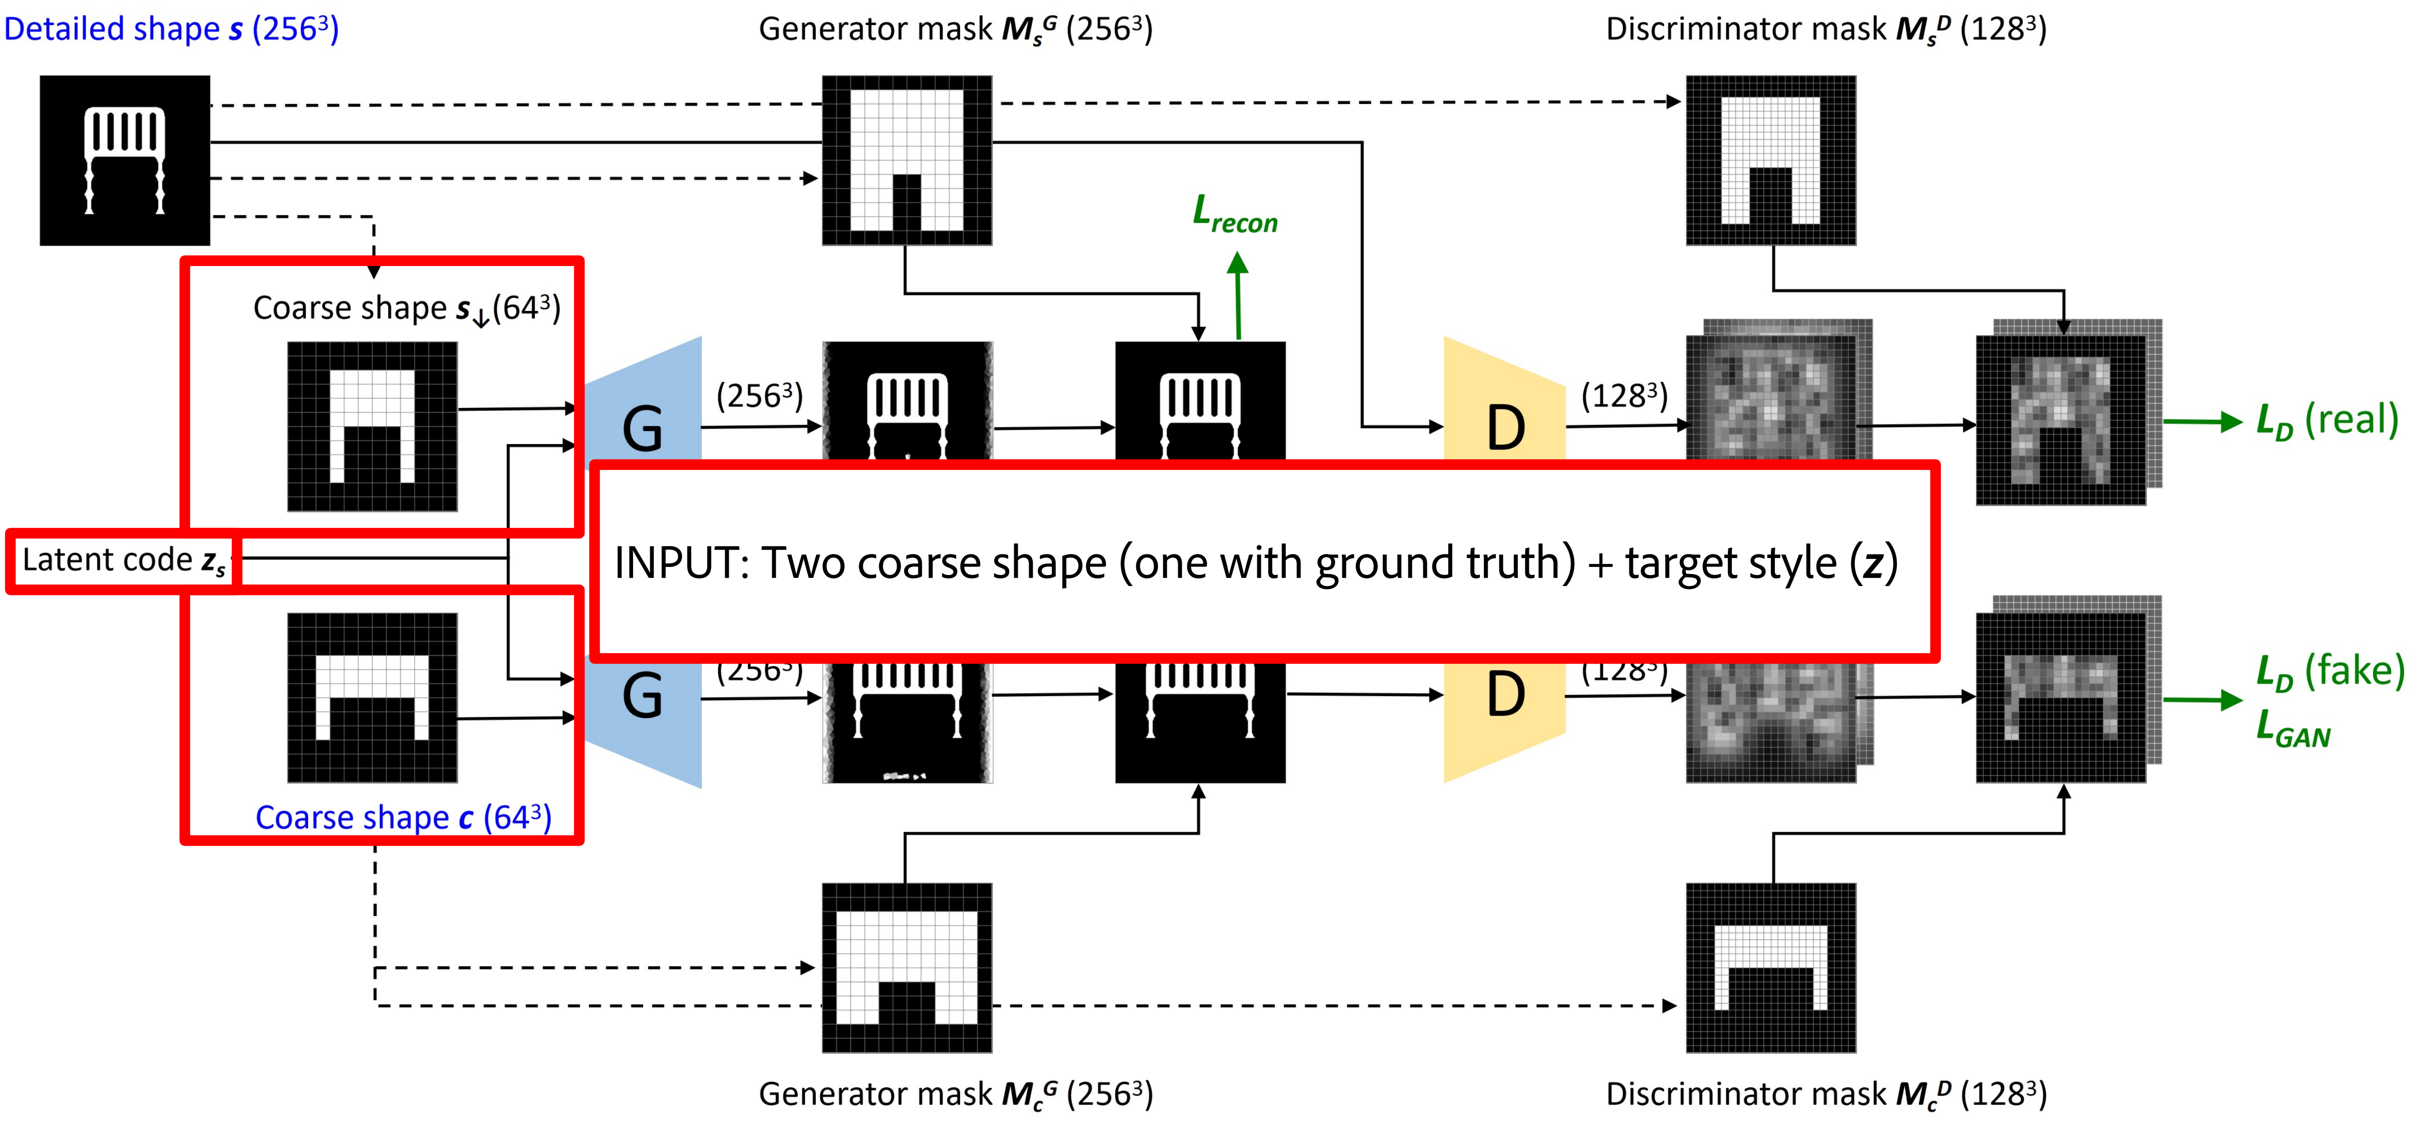

DÉCOR-GAN Neural Network

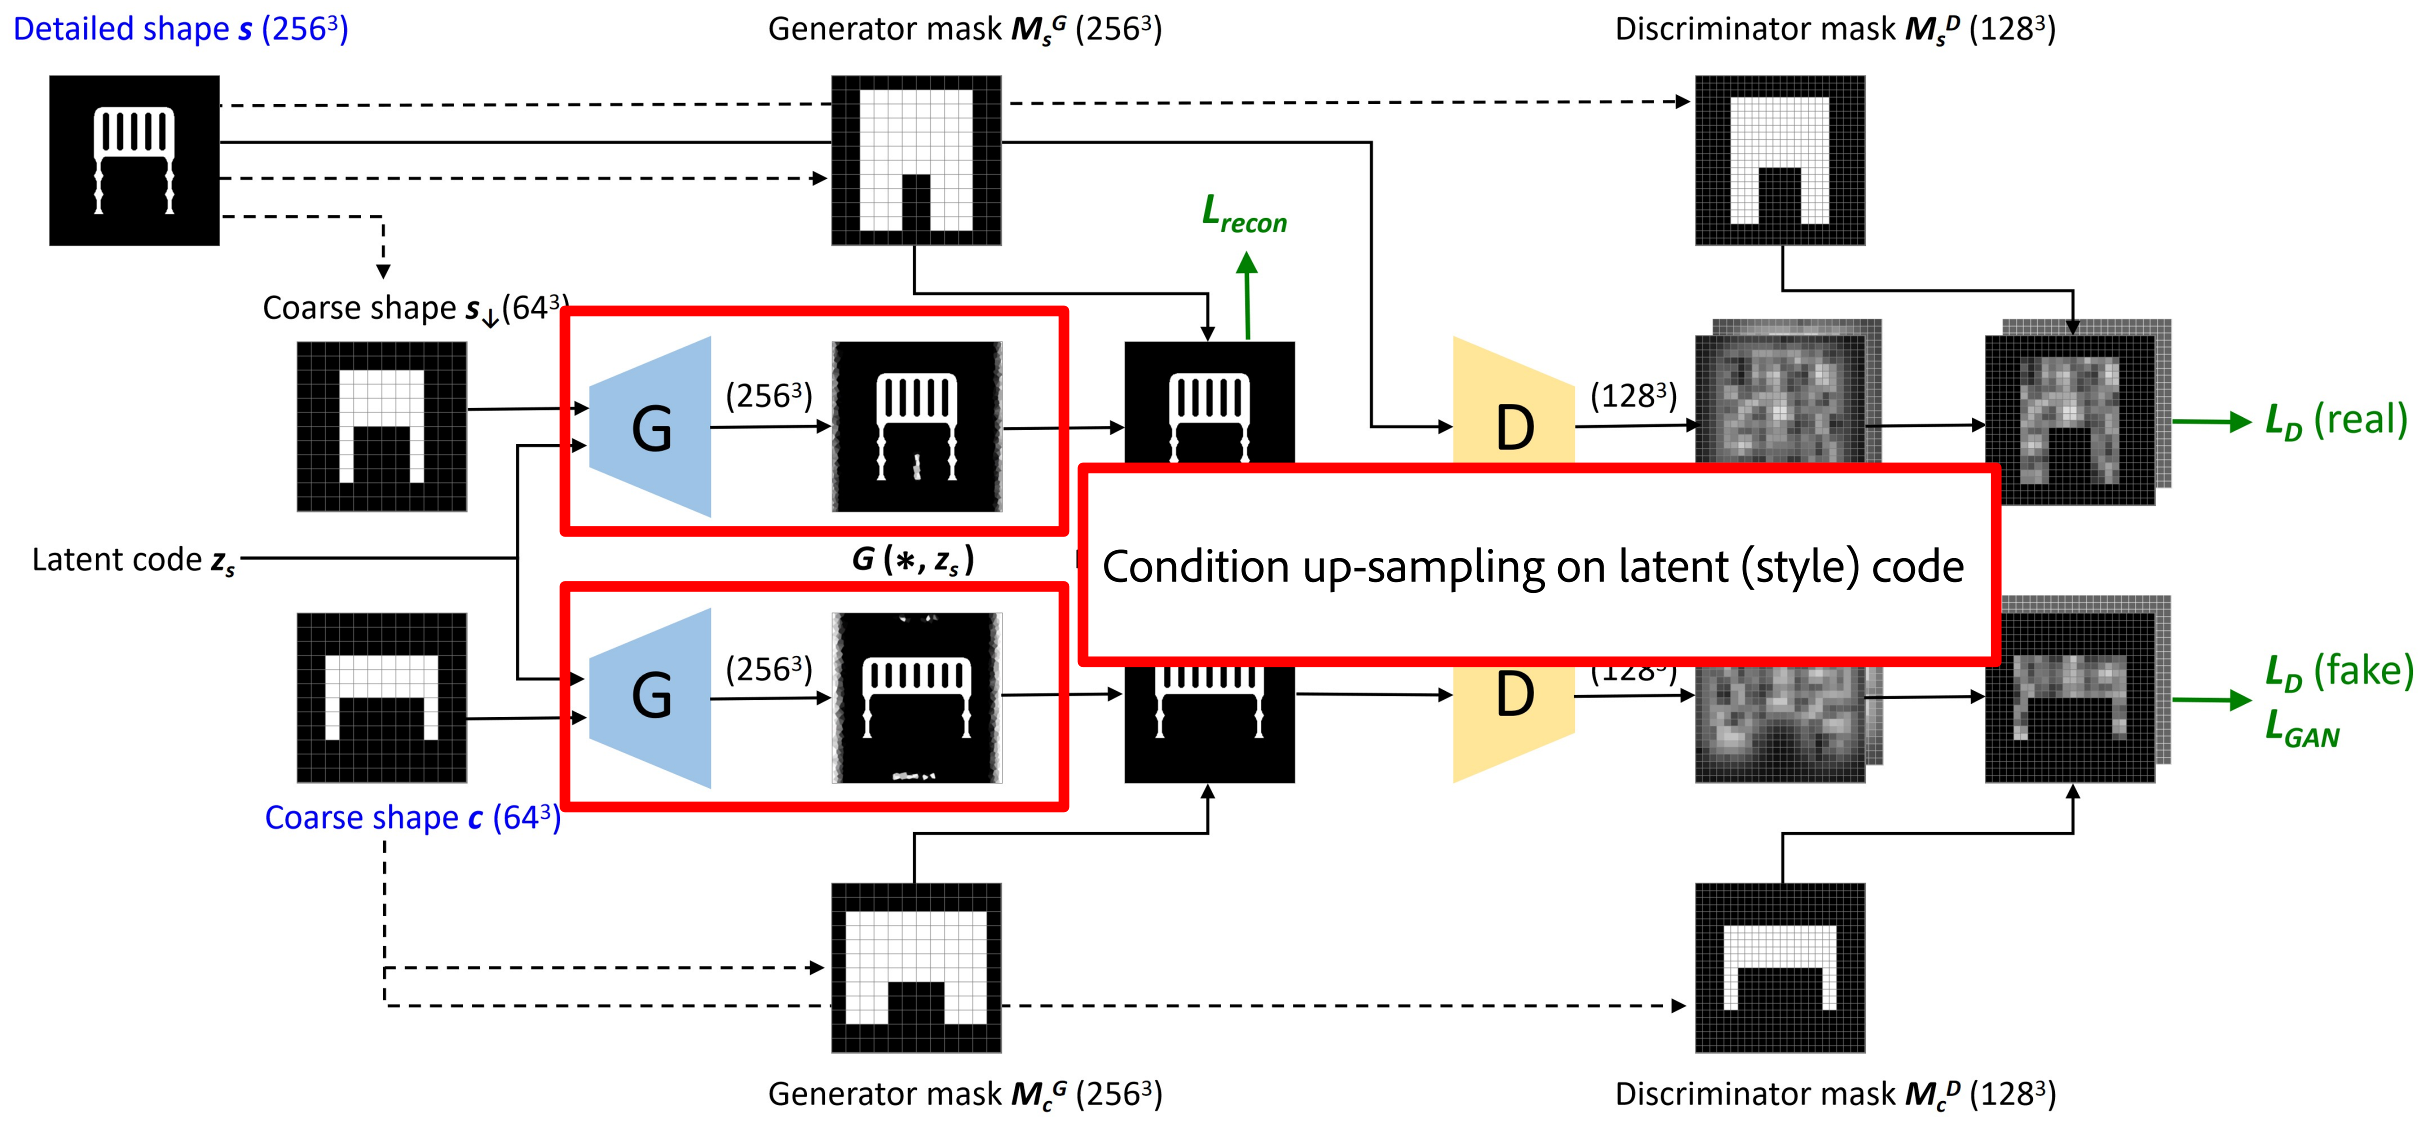

DÉCOR-GAN Neural Network

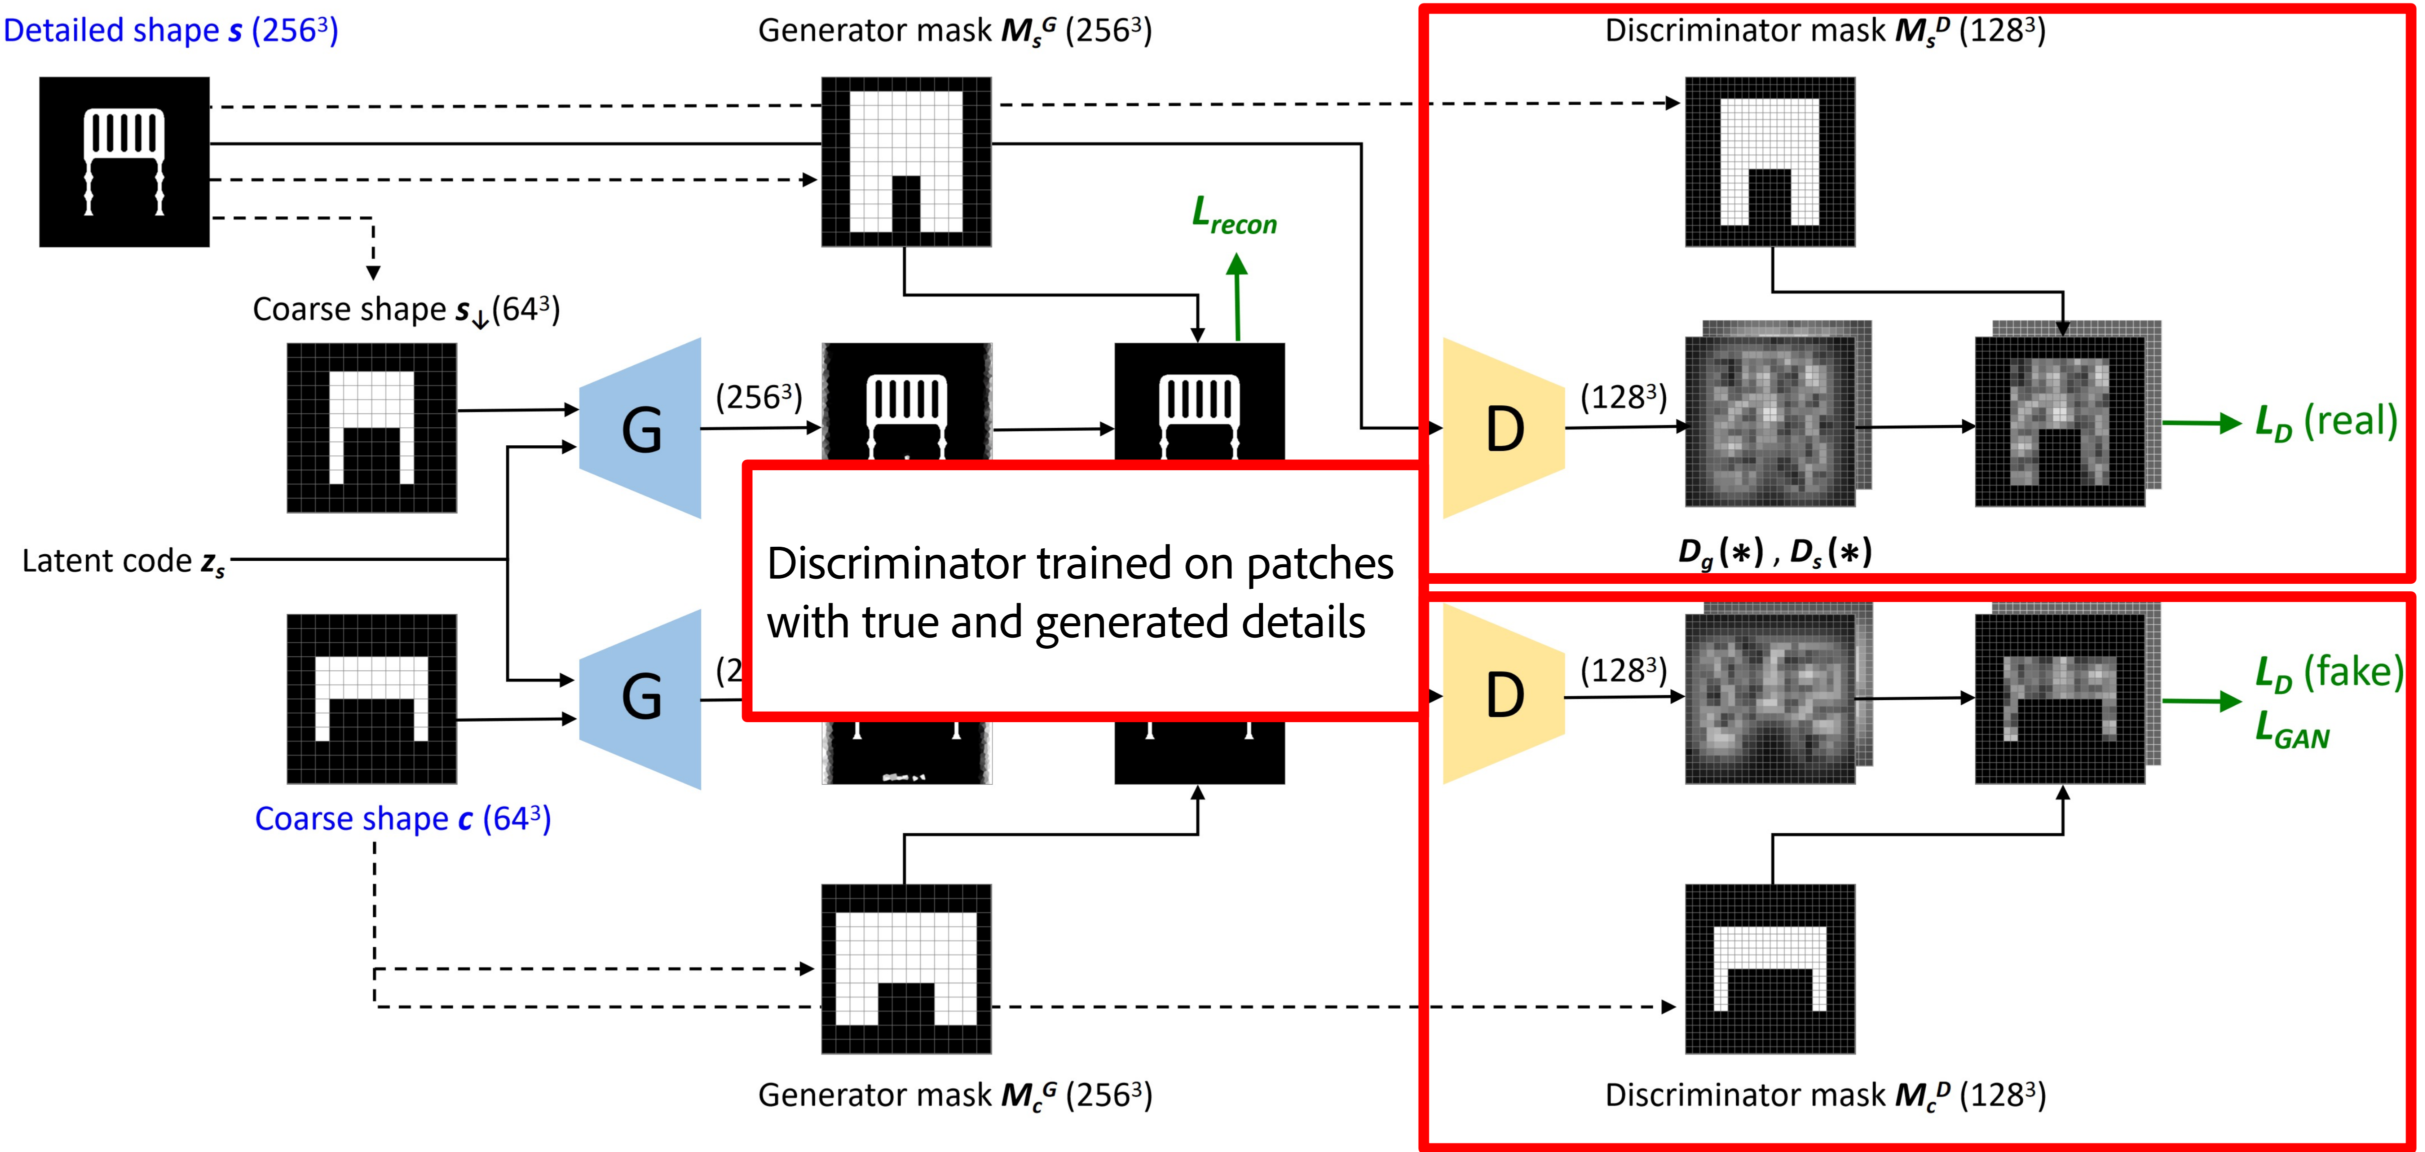

DÉCOR-GAN Neural Network

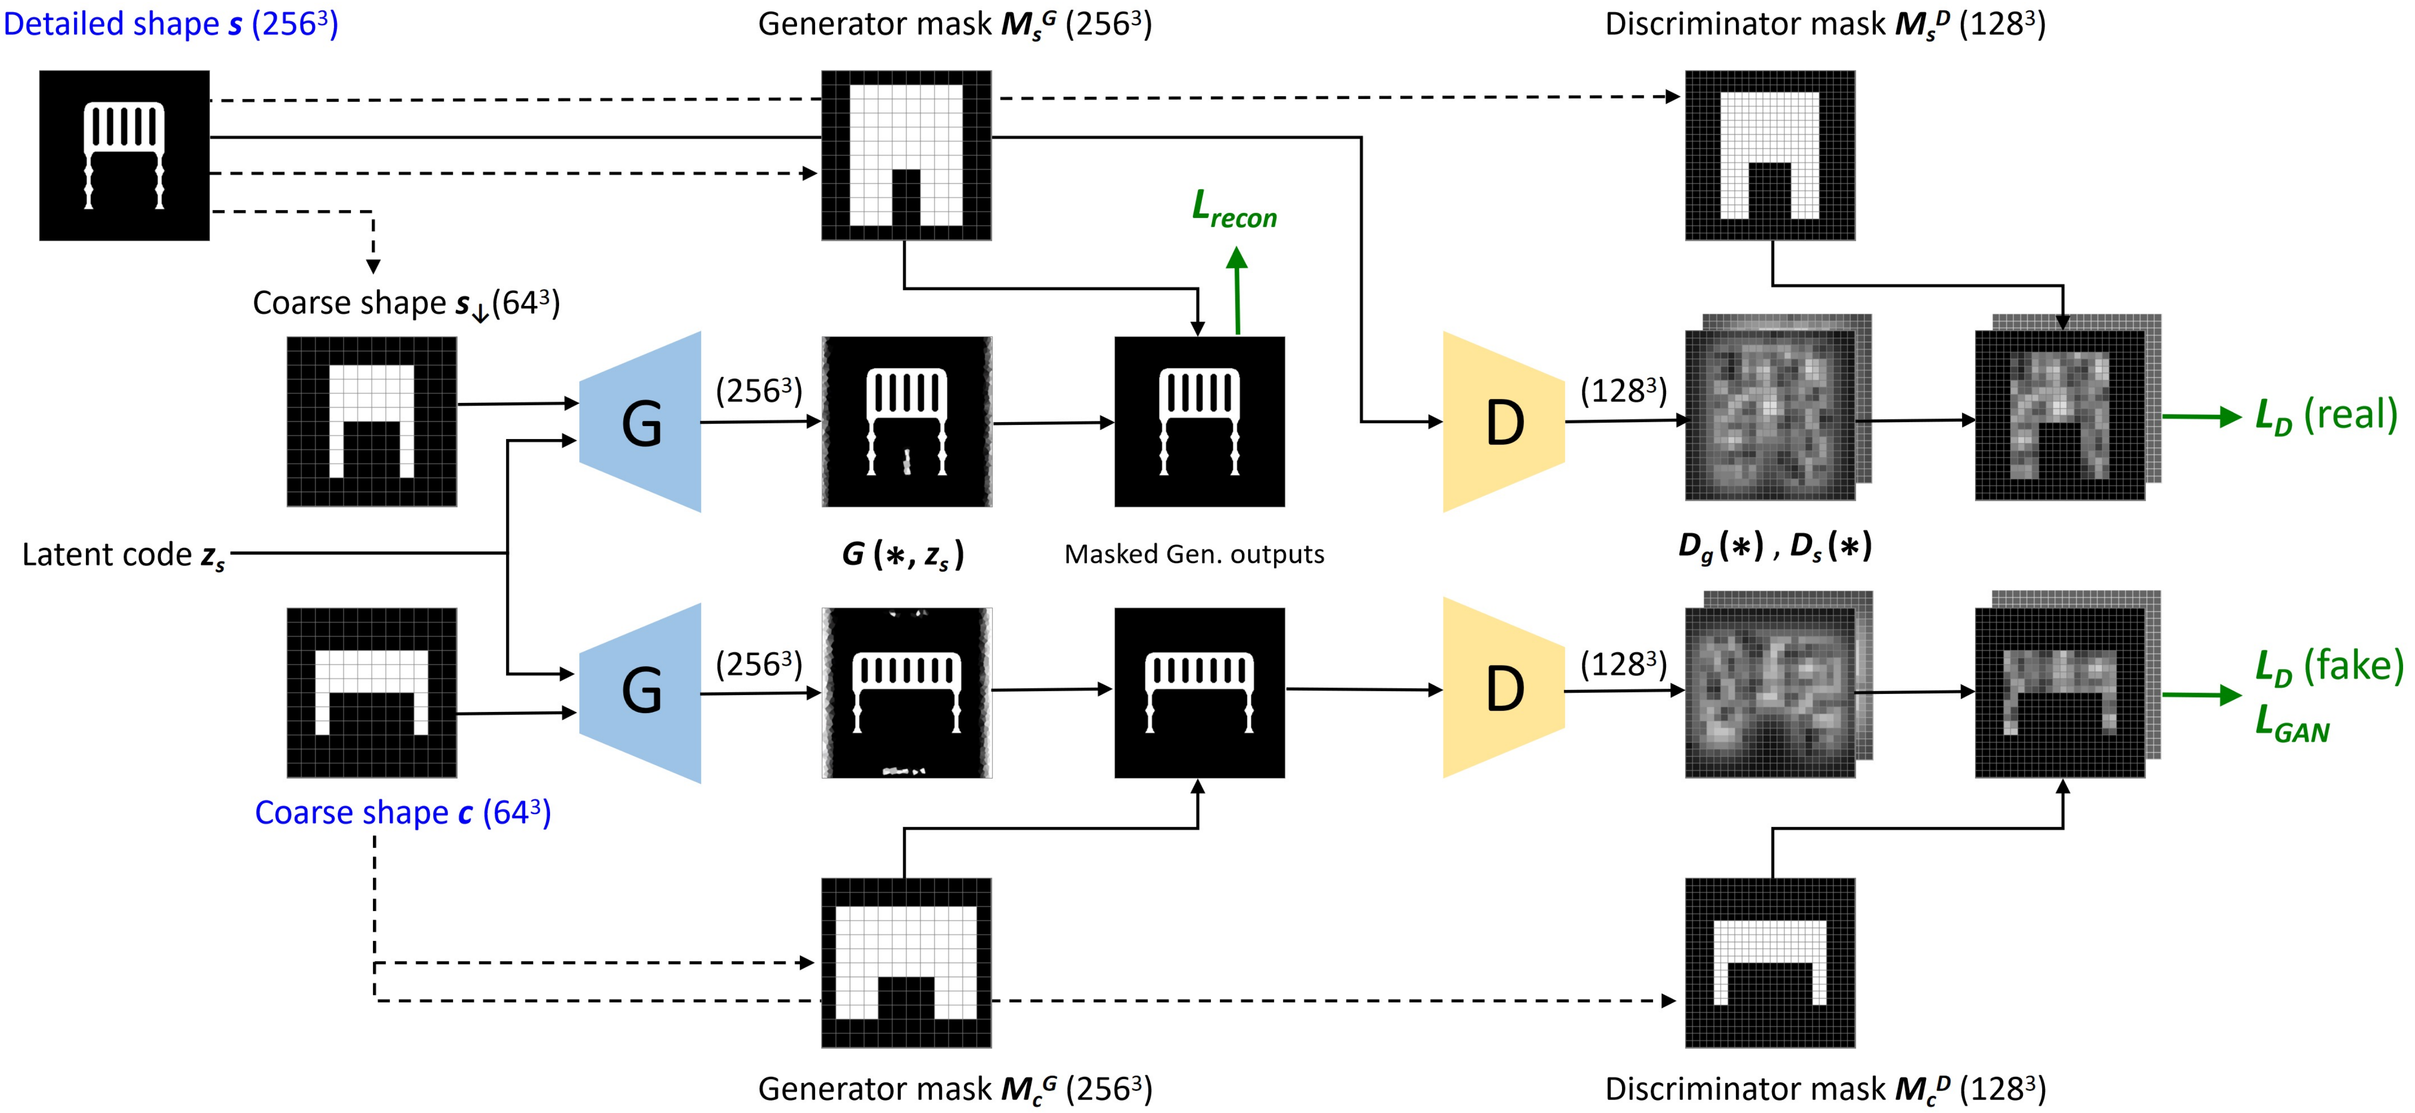

Results: Vases

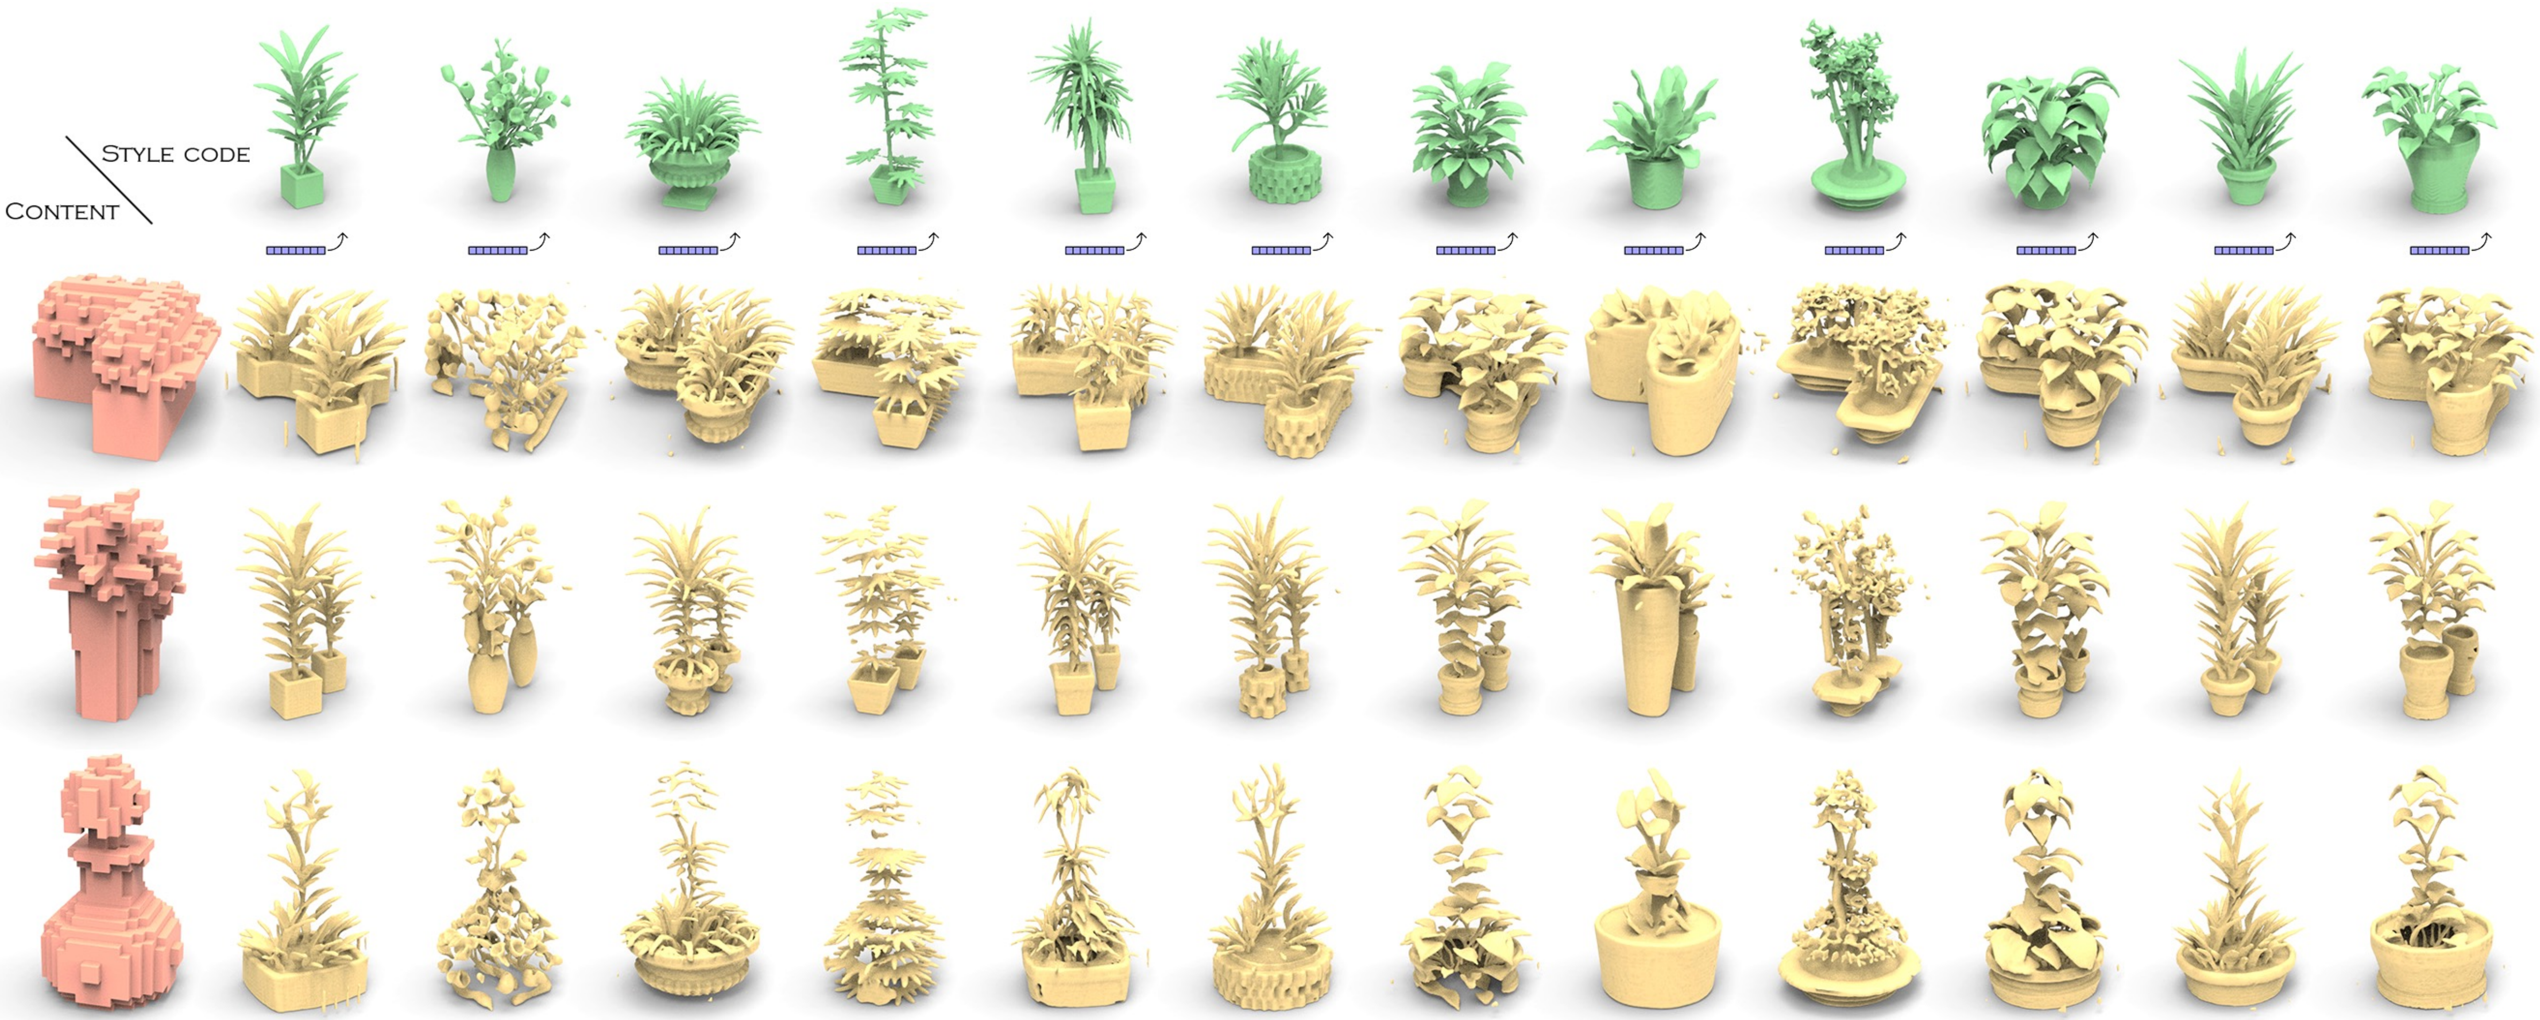

Results: Chairs

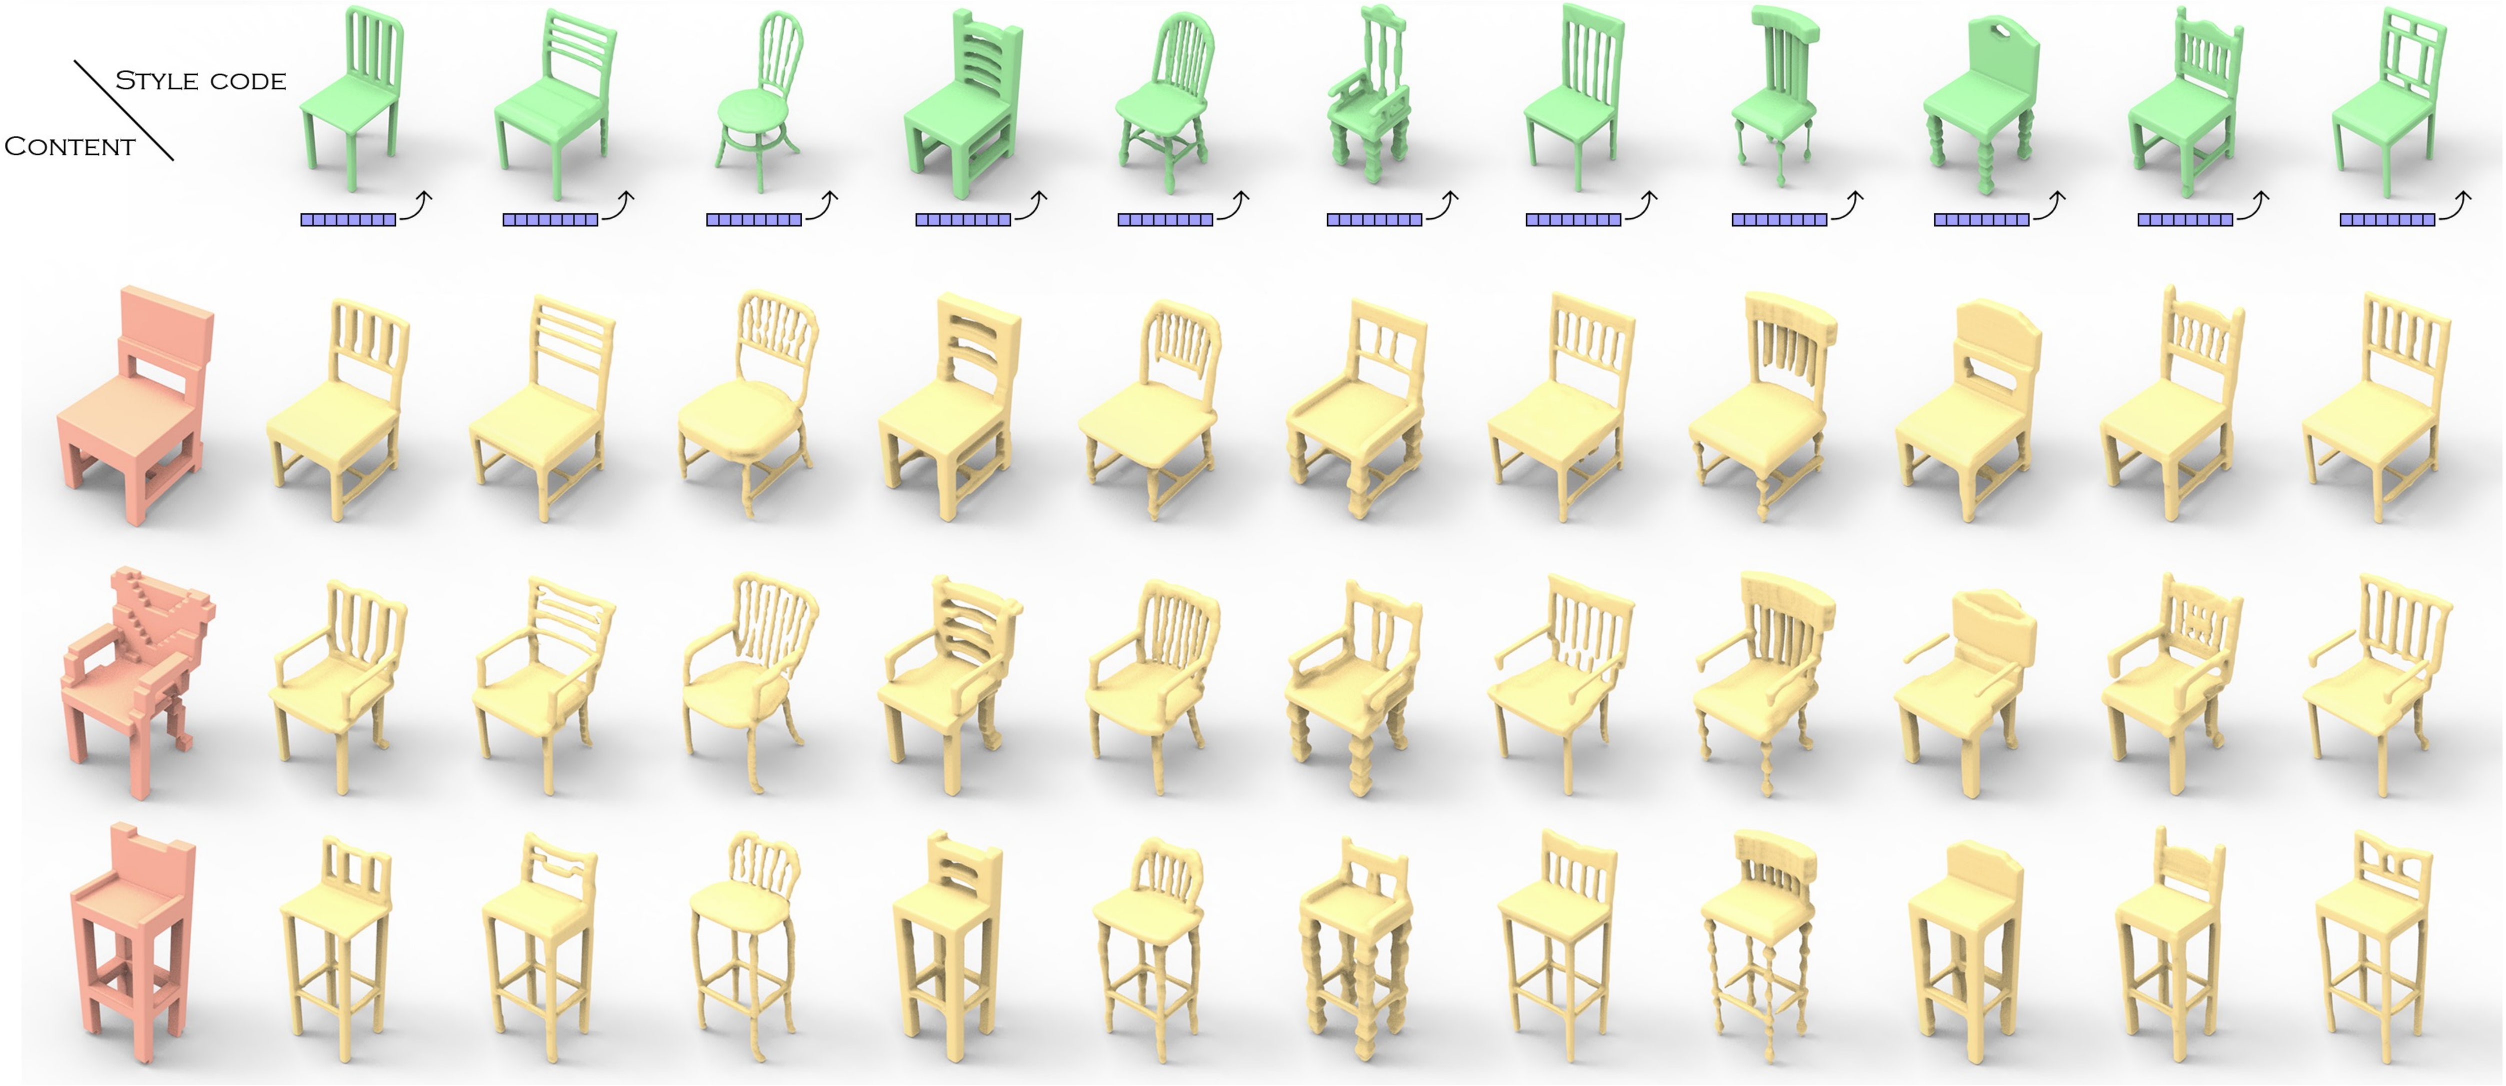

Results: Tables

Results: Airplanes

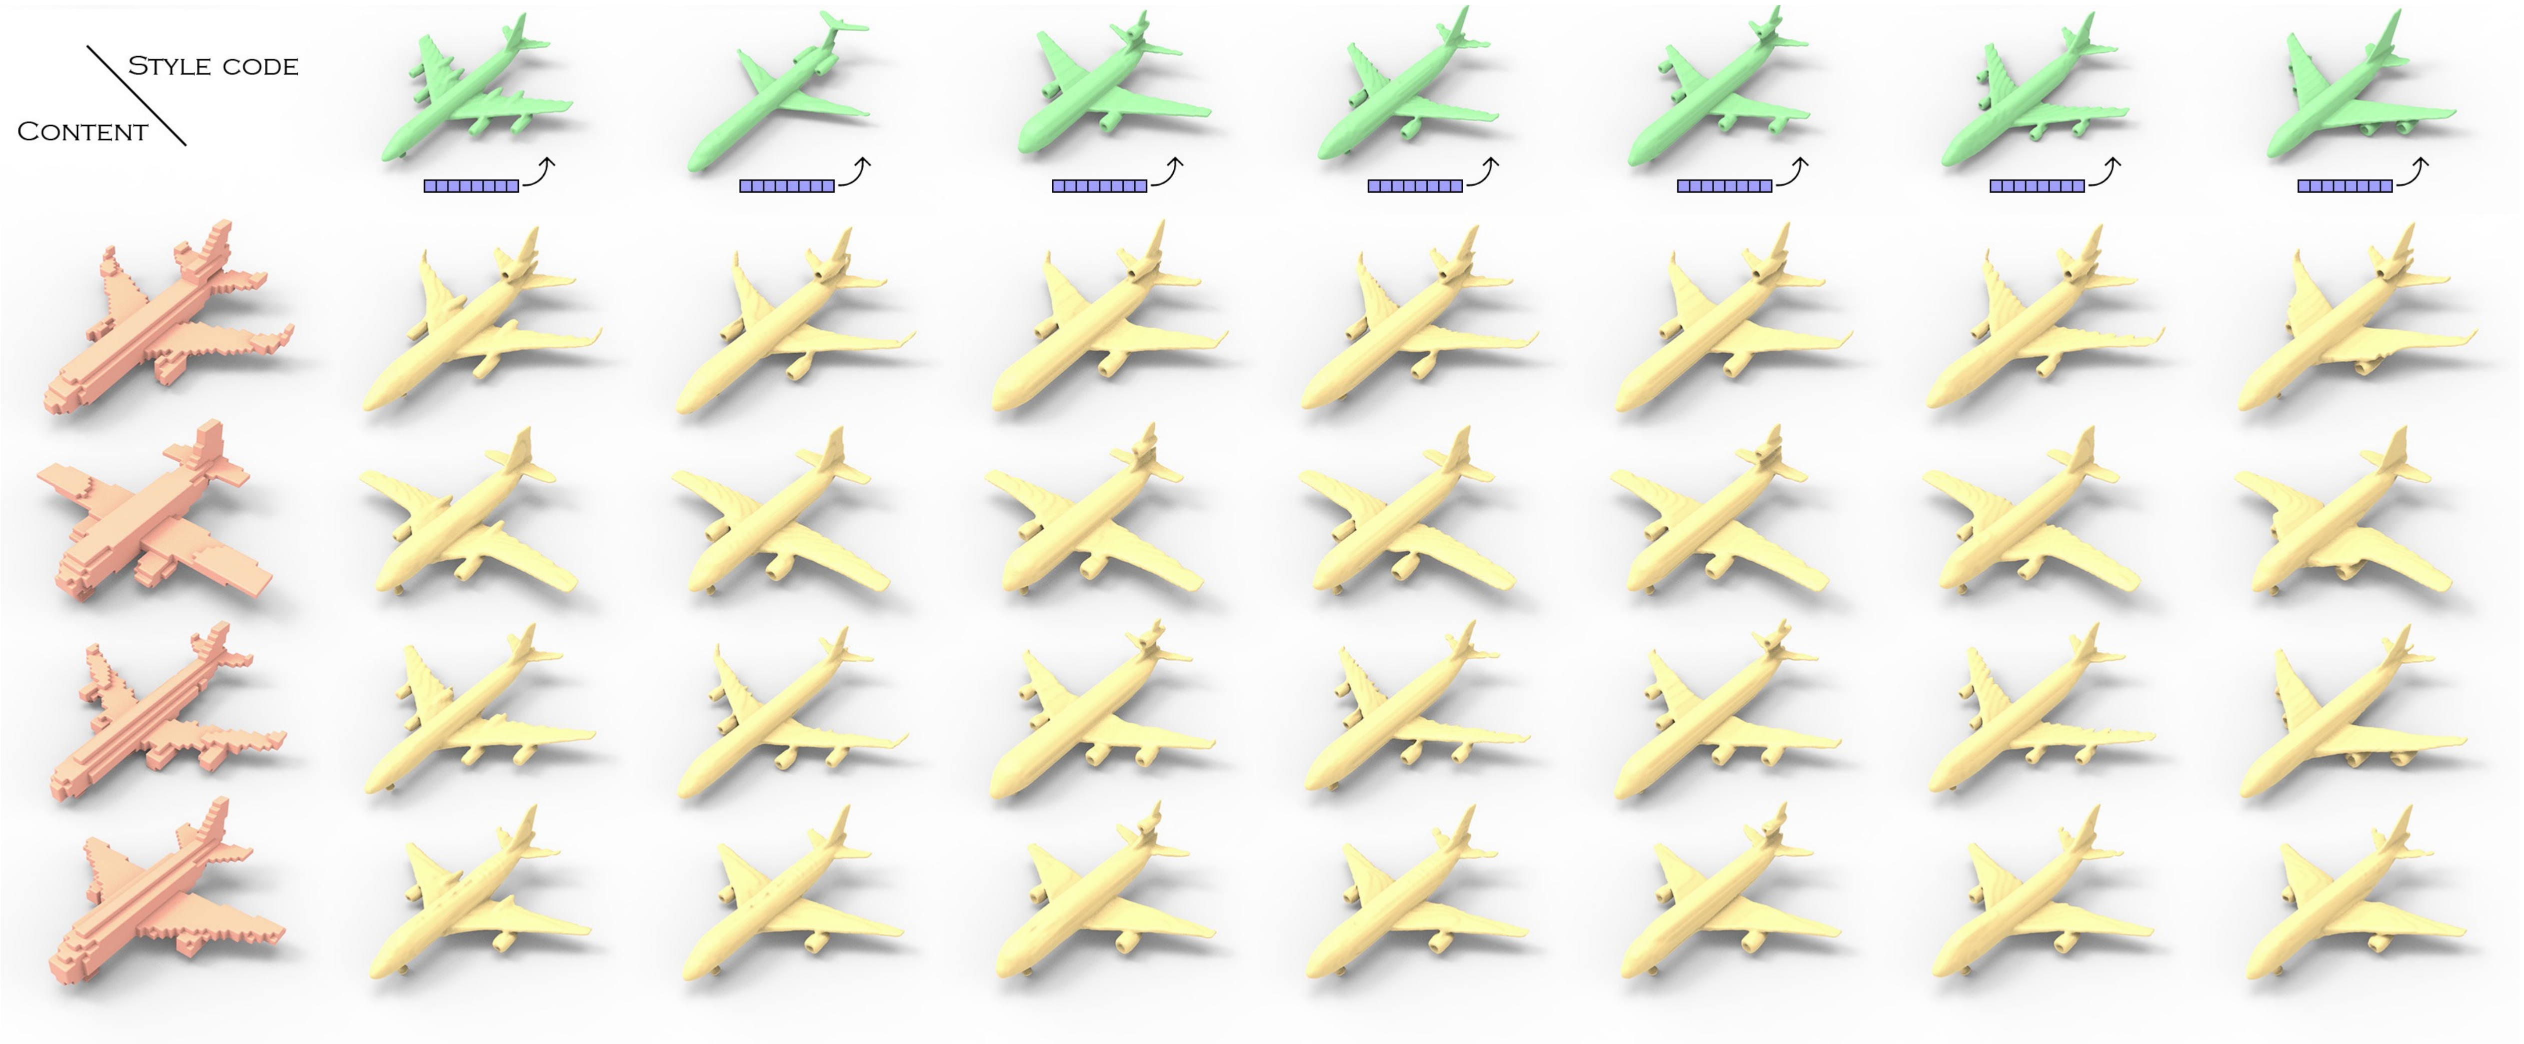

Results: Motorcycles

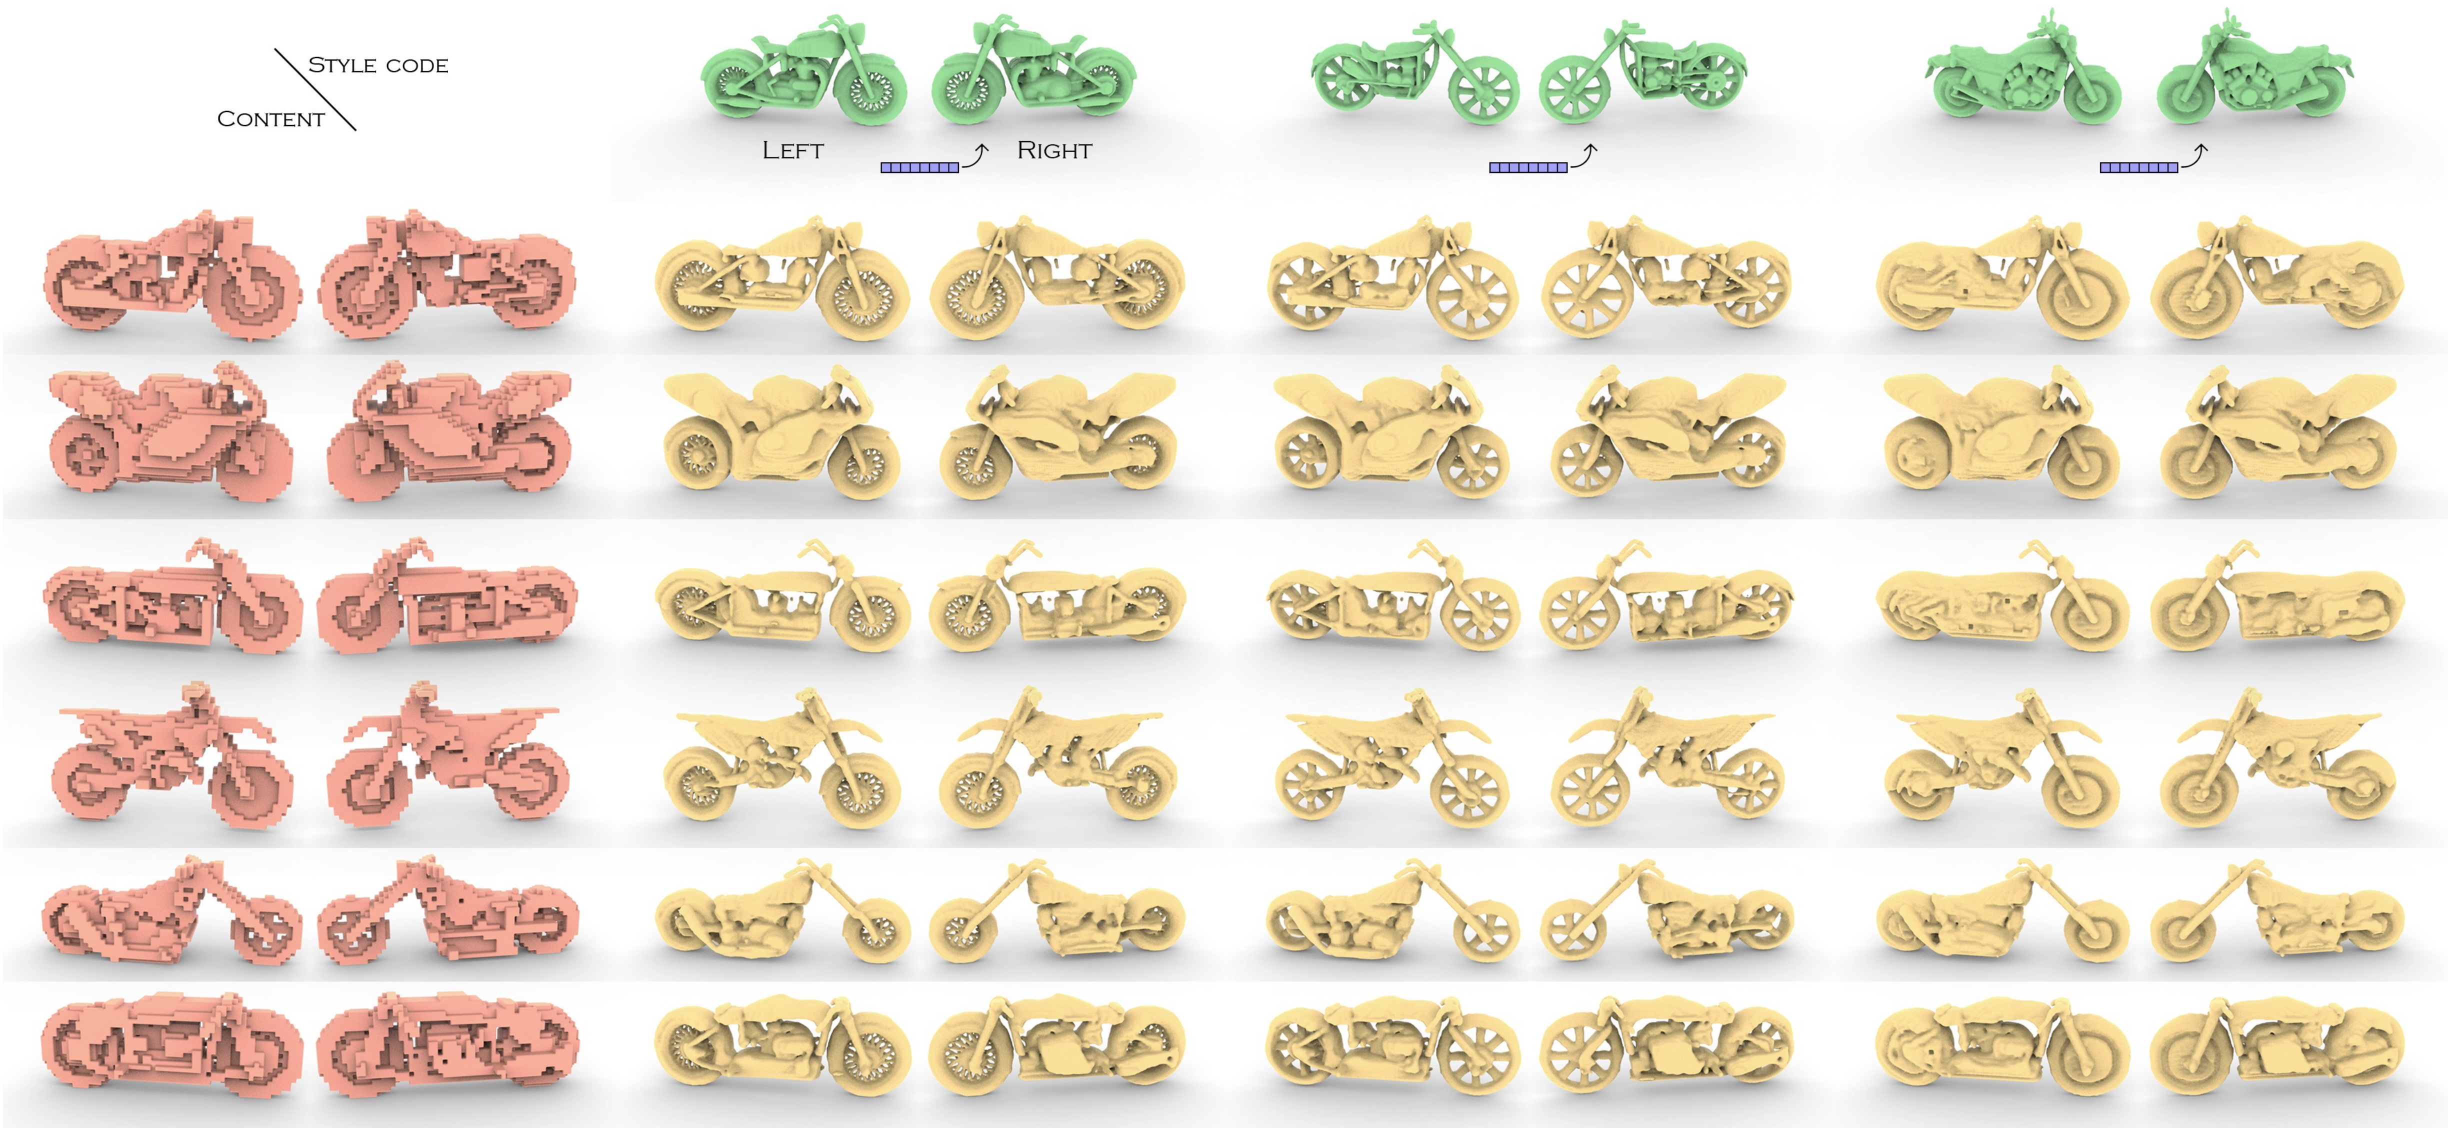

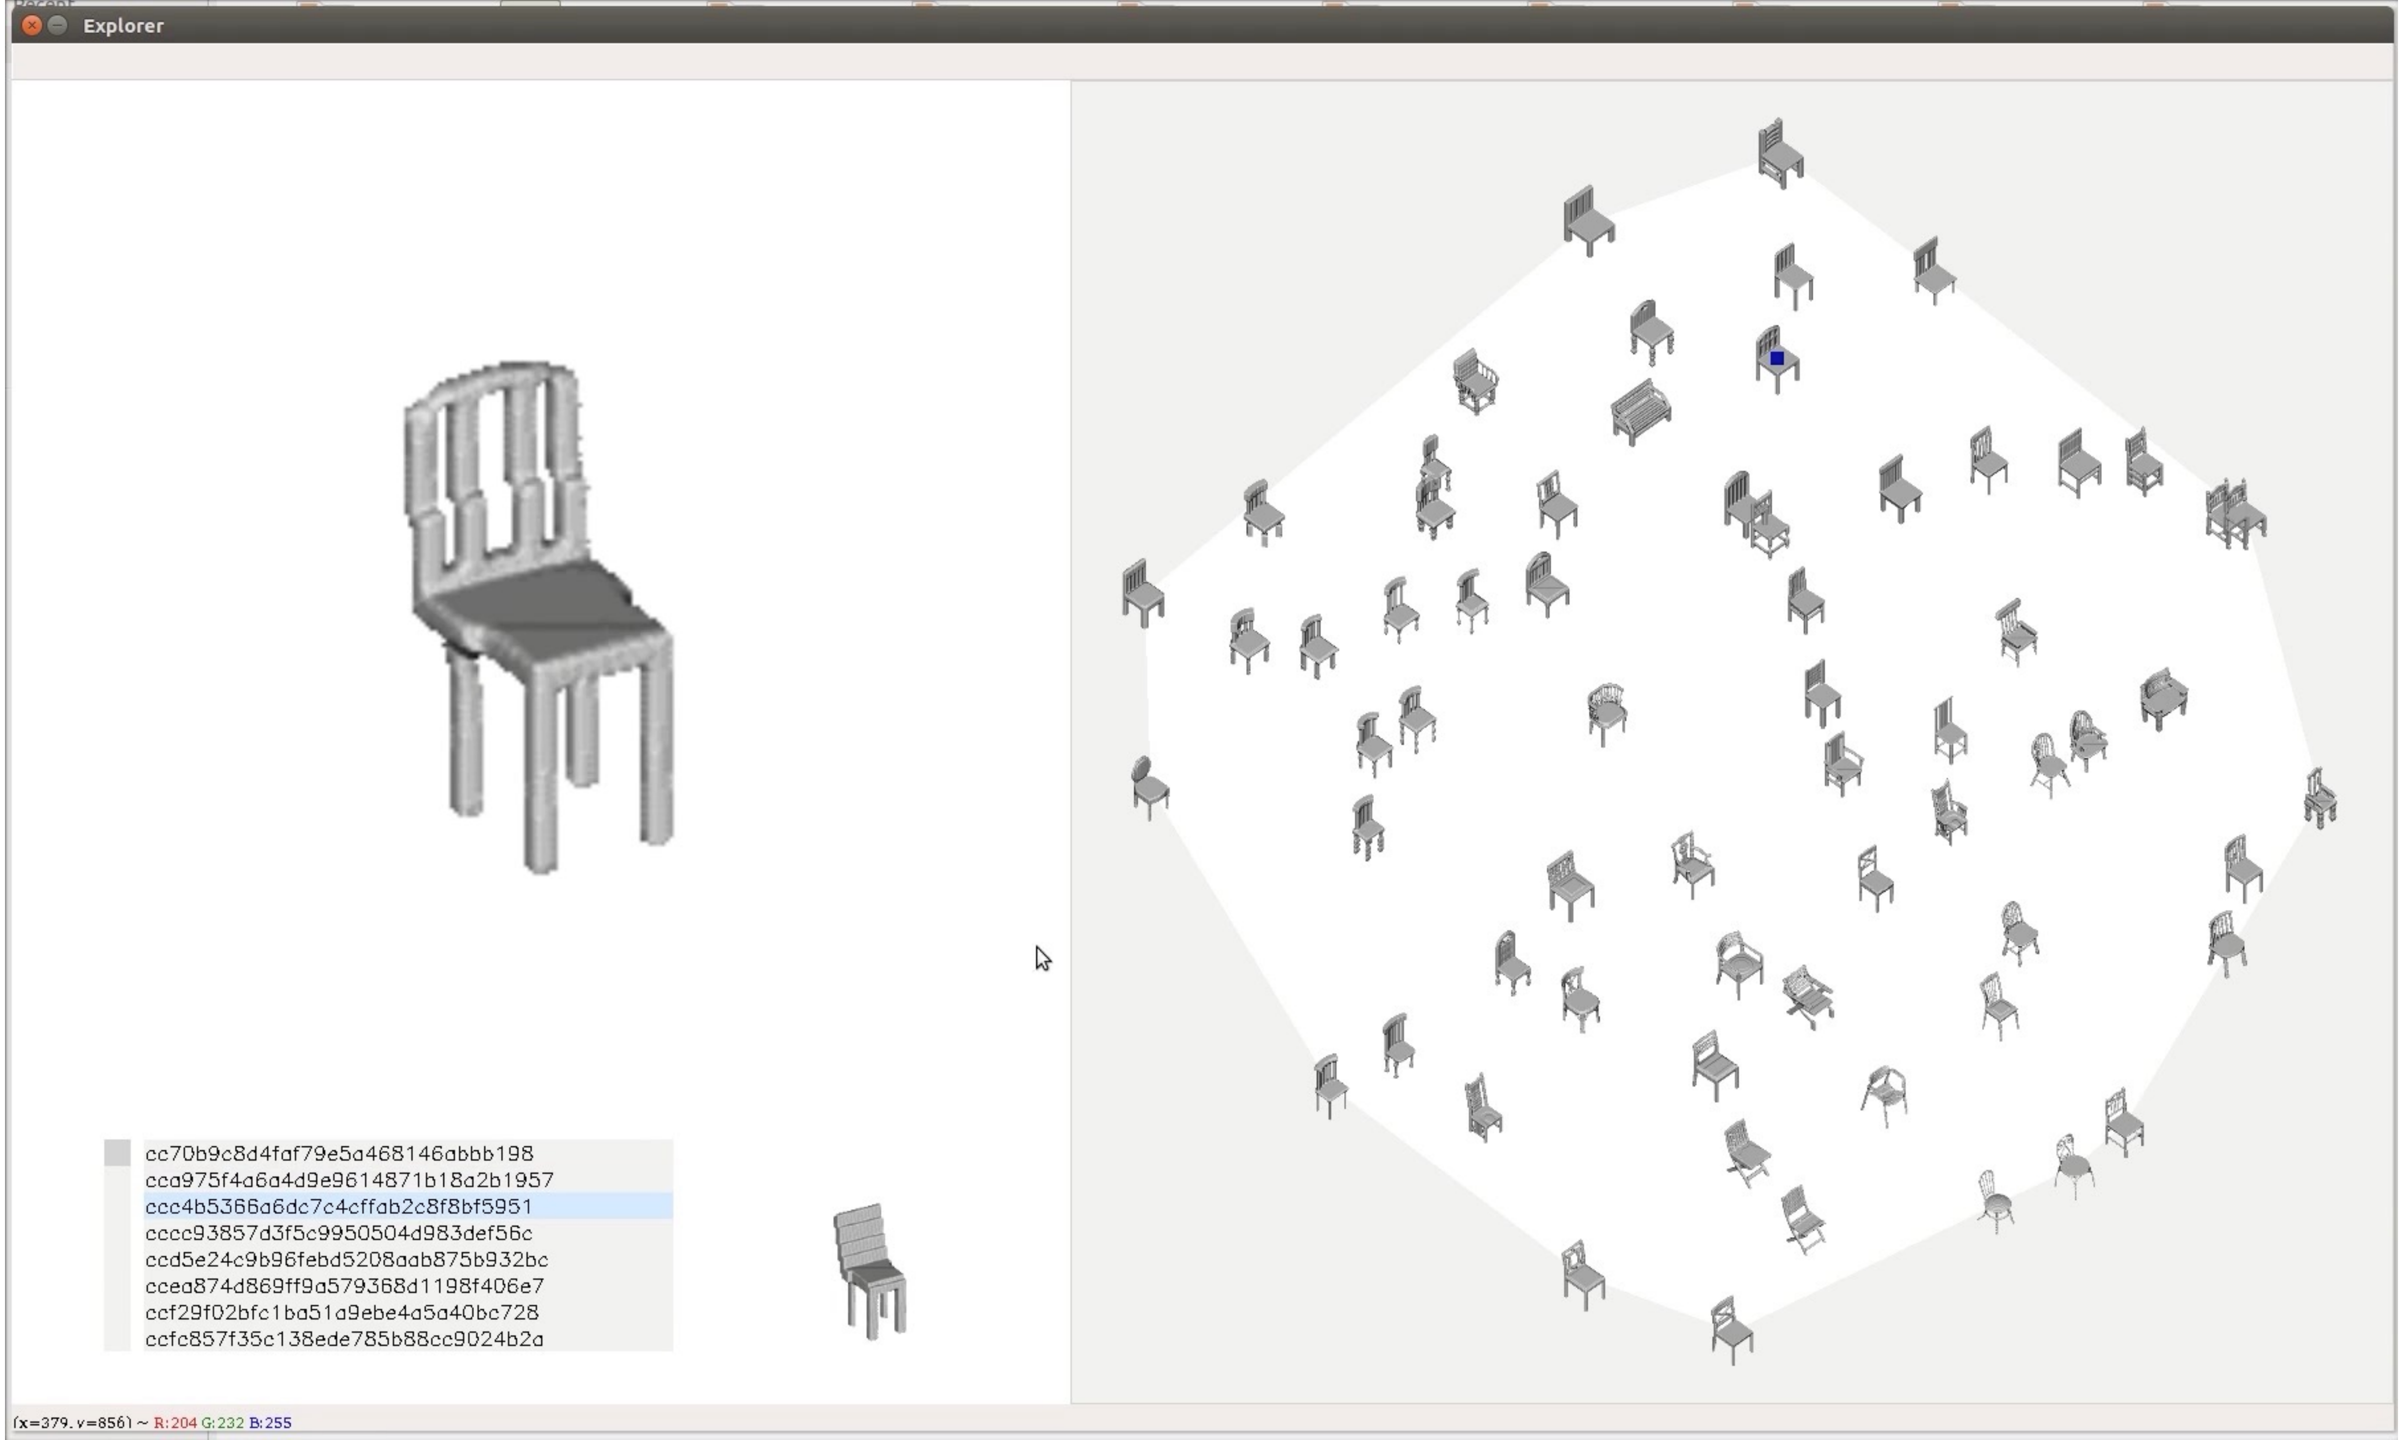
Explorer
cc70b9c8d4faf79e5a468146abbb198
cca975f4a6a4d9e9614871b18a2b1957
ccc4b5366a6dc7c4cffab2c8f8bf5951
cccc93857d3f5c9950504d983def56c
ccd5e24c9b96febd5208aab875b932bc
ccea874d869ff9a579368d1198f406e7
ccf29f02bfc1ba51a9ebe4a5a40bc728
ccfc857f35c138ede785b88cc9024b2a
(x=379, y=856) ~ R:204 G:232 B:255

Key Takeaways

- Retrieval and Deformation should be trained jointly

- Neural Deformation
  - MLP for entire map is too flexible (shape gets distorted)
  - Neural Cage-based deformation is too constrained (OK in some cases)
  - Neural Jacobian Fields (flexible and low-distortion)

- Neural Detailization
  - Neural Subdivision: effective for meshes, but input has to have the right topology
  - DÉCOR-GAN: voxel grids are good for learning details, classical image-based ideas are directly applicable

Future Work

- Geometry learning for production-quality assets
  - Diverse representation: different level of details and tessellation quality
  - Diverse content: few-shot learning
  - Model appearance: geometry, materials, and environment
- Neural Shape Processing: re-use and re-purpose the existing assets

Made with Adobe Stager

Collaborators

- Project Leads
  - Mikaela Uy, Stanford, Adobe Intern (Joint Learning of 3D Shape Retrieval and Deformation, **CVPR 2021**)
  - Thibault Groueix, ParisTech, Adobe (Unsupervised cycle-consistent deformation for shape matching, **SGP 2019**)
  - Yifan Wang, ETH Zurich, Adobe Intern (Neural Cages for Detail-Preserving 3D Deformations, **CVPR 2020, oral**)
  - Noam Aigerman, Adobe (Neural Jacobian Fields: Learning Intrinsic Mappings of Arbitrary Meshes, **SIGGRAPH 2022**)
  - Hsueh-Ti Liu, U. Toronto, Adobe Intern (Neural Subdivision, **SIGGRAPH 2020**)
  - Zhiqin Chen, SFU, Adobe Intern (Décor-GAN, **CVPR 2021, oral**)

- Collaborators
  - Siddhartha Chaudhuri, Matt Fisher, Bryan Russell, Alec Jacobson, Minhyuk Sung– **Adobe Research**
  - Leonidas Guibas – **Stanford University**
  - Olga Sorkine – **ETH Zurich**
  - Mathieu Aubry – **Ecole des Ponts ParisTech**
  - Richard Zhang – **Simon Fraser University**