NEURAL INFORMATION PROCESSING SYSTEMS

GuideFlow3D

Optimization-Guided Rectified Flow For Appearance Transfer

Sayan Deb Sarkar Sinisa Stekovic Vincent Lepetit Iro Armeni

Stanford University

Gradient Spaces

École des Ponts ParisTech

Why 3D Appearance Transfer?

- Bring real-world styles and materials into design
- Generate 3D assets by transferring texture and fine geometry while preserving shape

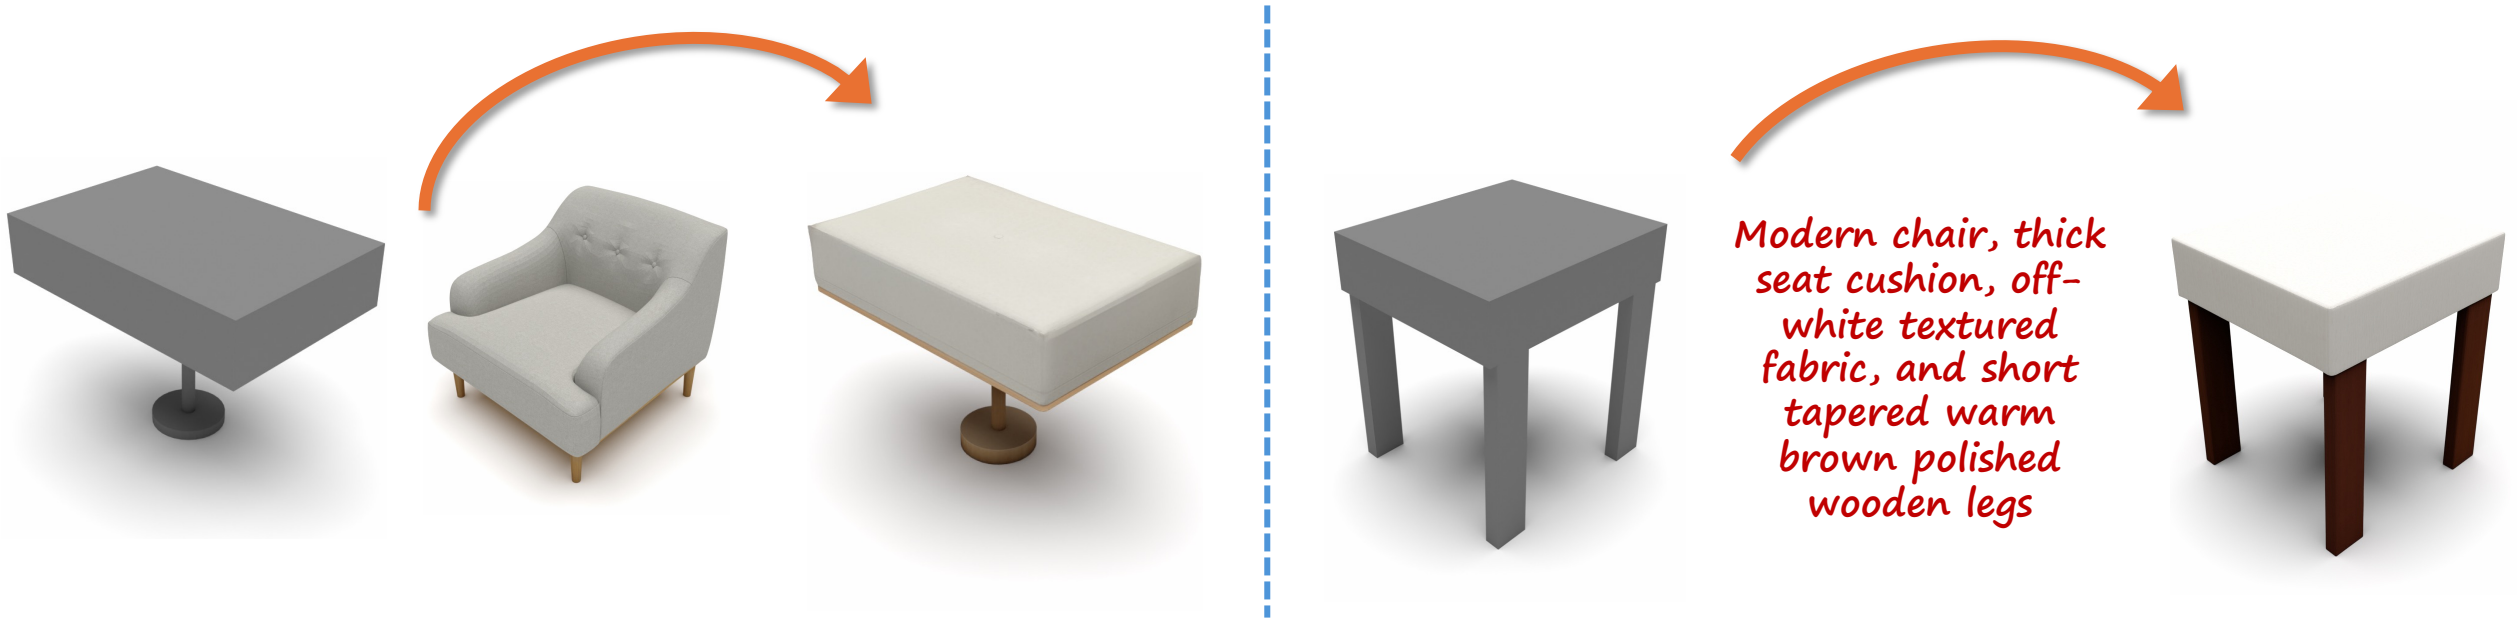

Goal: Accelerate stylized asset creation for gaming, AR/VR, and digital prototyping

Why 3D Appearance Transfer can be Hard?

1. Geometric irregularity and absence of part-aware grounding disrupt texture alignment and structural consistency.

Even similar shapes misalign with **no structural awareness**.

[1] Perla et al, *EASI-Tex: Edge-Aware Mesh Texturing from Single Image*, SIGGRAPH 2024.

Why 3D Appearance Transfer can be Hard?

2. Large semantic gaps across categories break correspondence, causing style leakage and textures that fail to align with object geometry.

Different shapes collapse when **style overpowers structure**.

[1] Xiang et al, *Structured 3D Latents for Scalable and Versatile 3D Generation*, CVPR 2025.

Our Solution

Guided Flow For Appearance Transfer

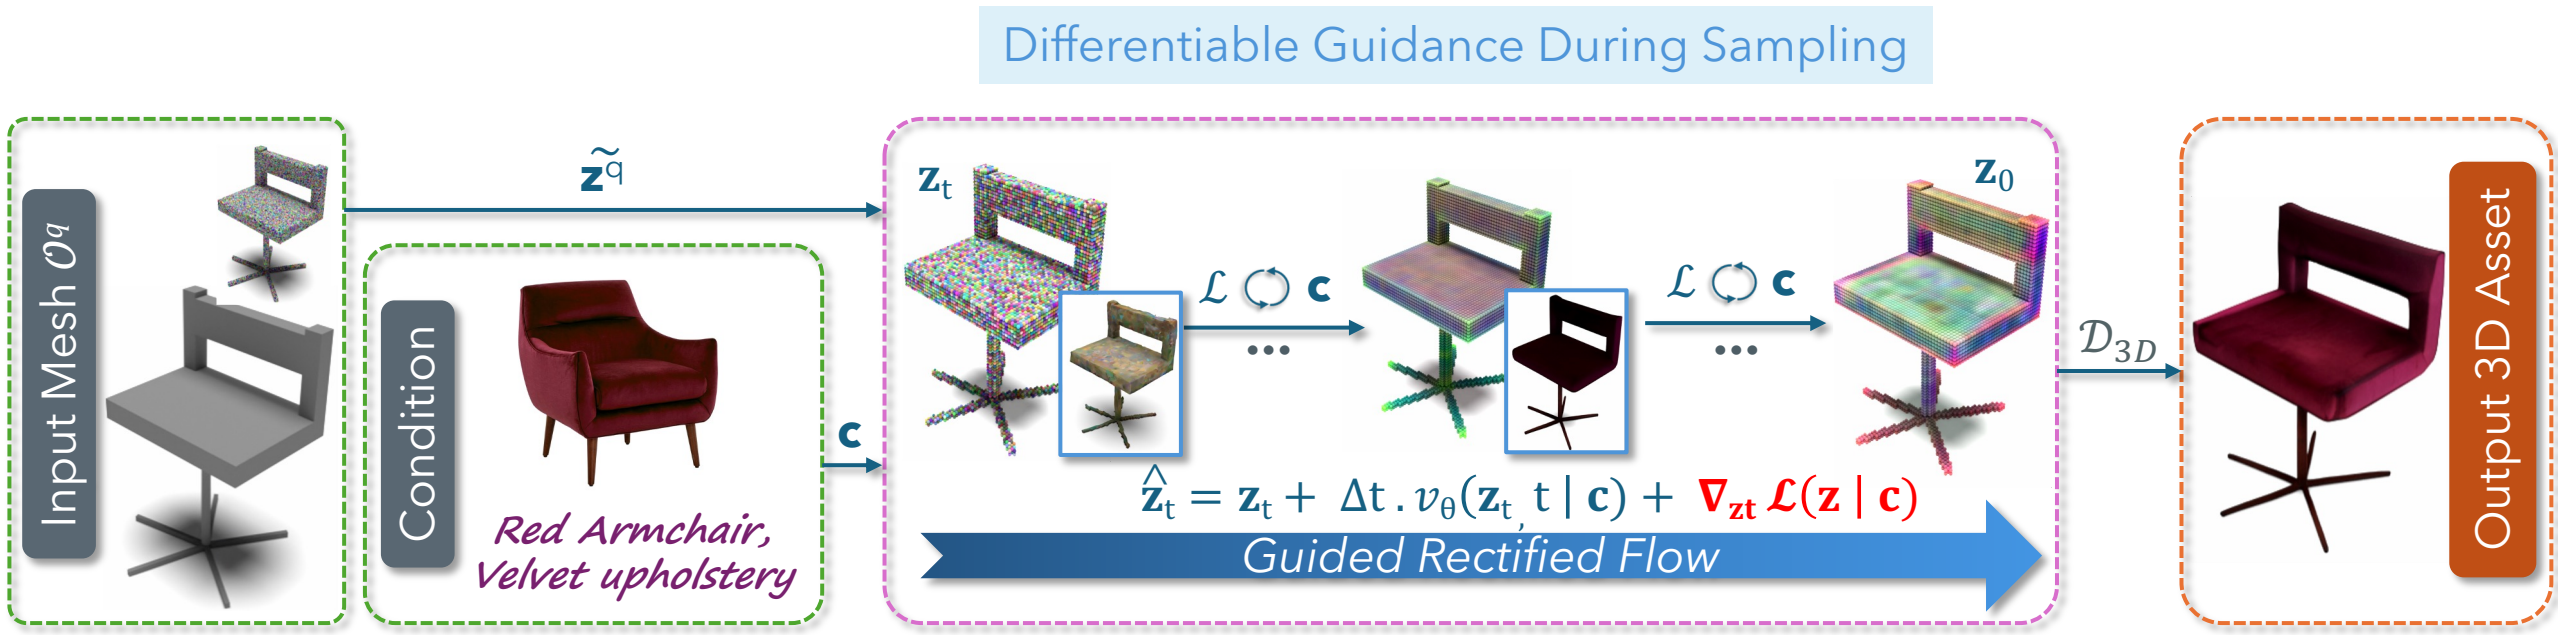

Interleave **rectified flow sampling** and **semantic + geometric** prior as guidance objective.

Our Solution

Guided Flow For Appearance Transfer

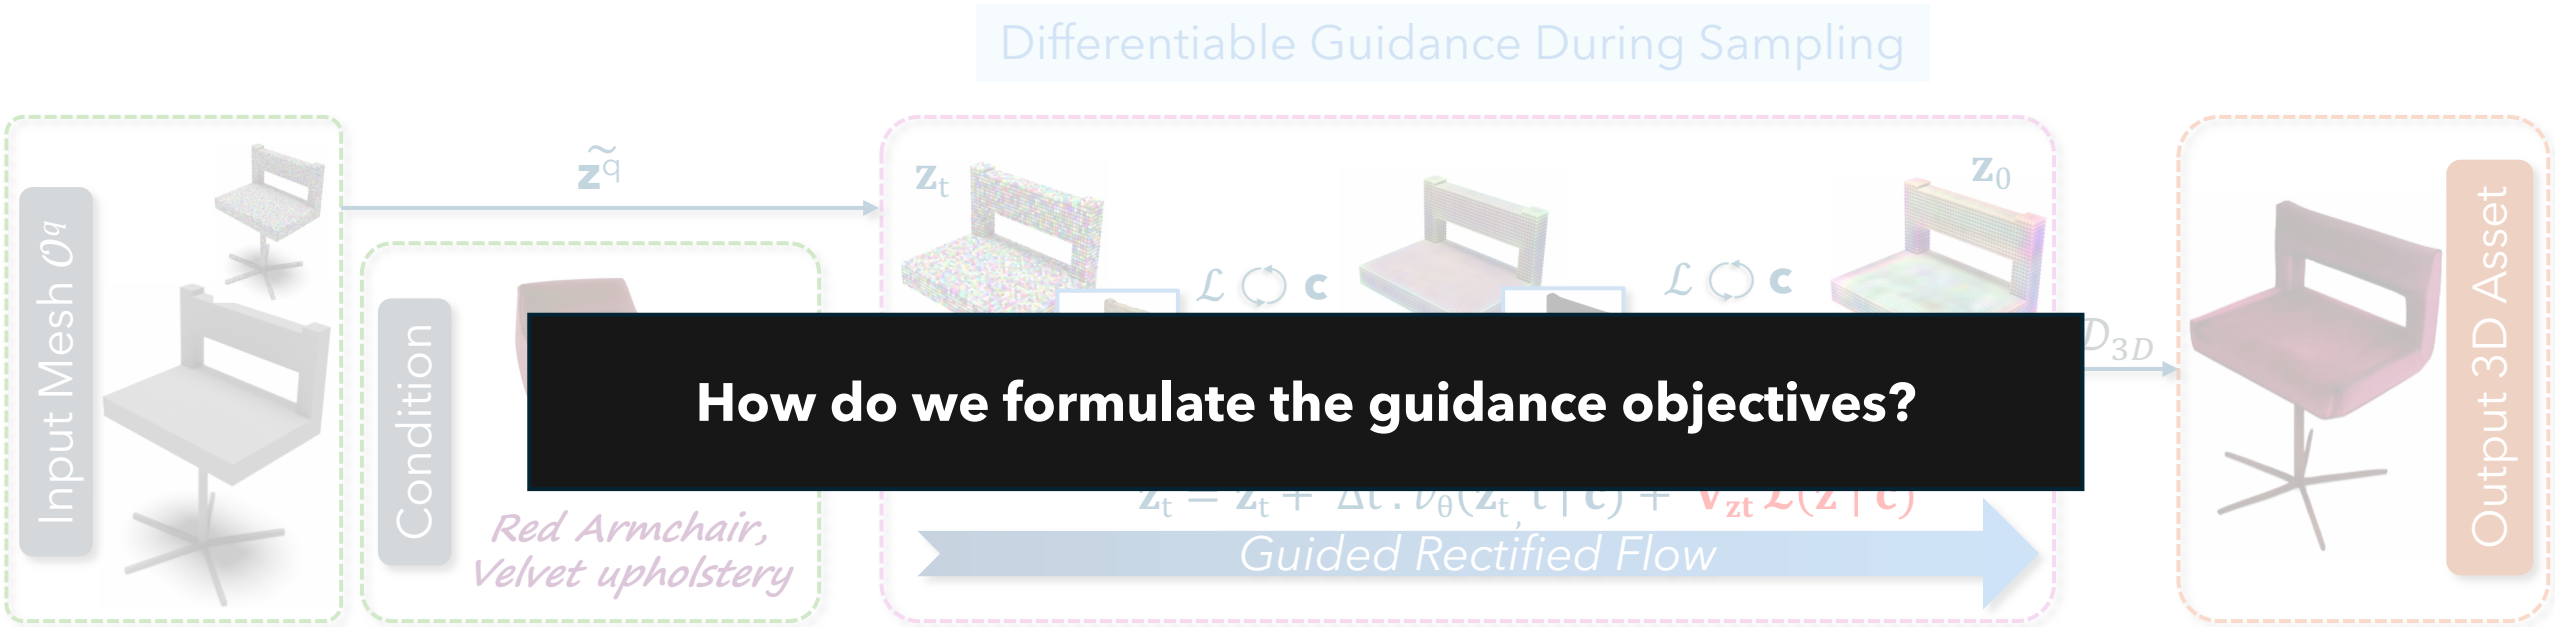

How do we formulate the guidance objectives?

Interleave **rectified flow sampling** and **semantic + geometric** prior as guidance objective.

Guided Flow For Appearance Transfer

Part-Aware Semantic Guidance

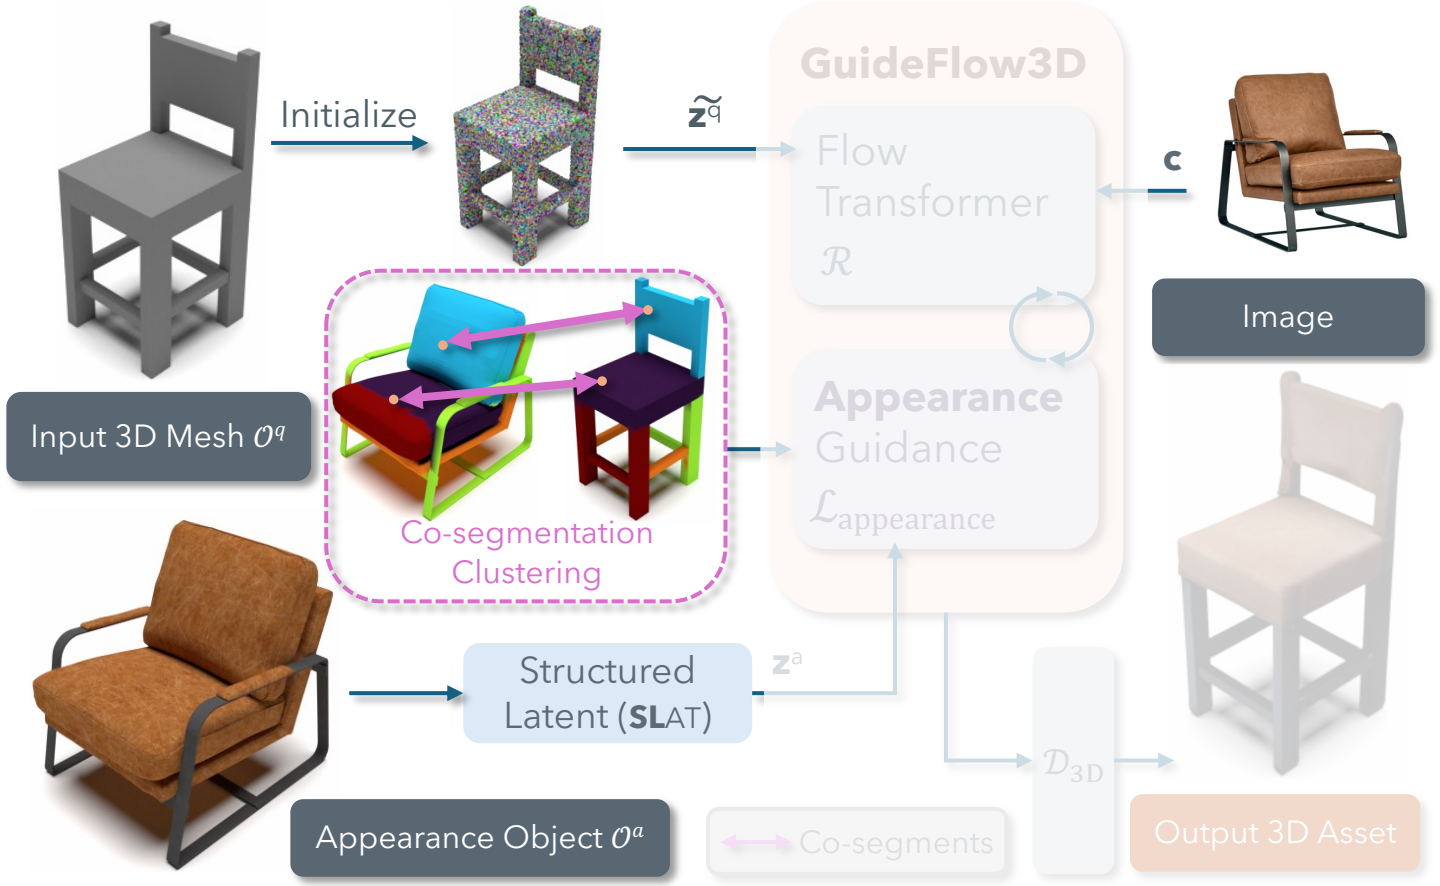

Matches input and appearance latents through **part-based co-segmentation**.

Guided Flow For Appearance Transfer

Self-Similarity Guidance

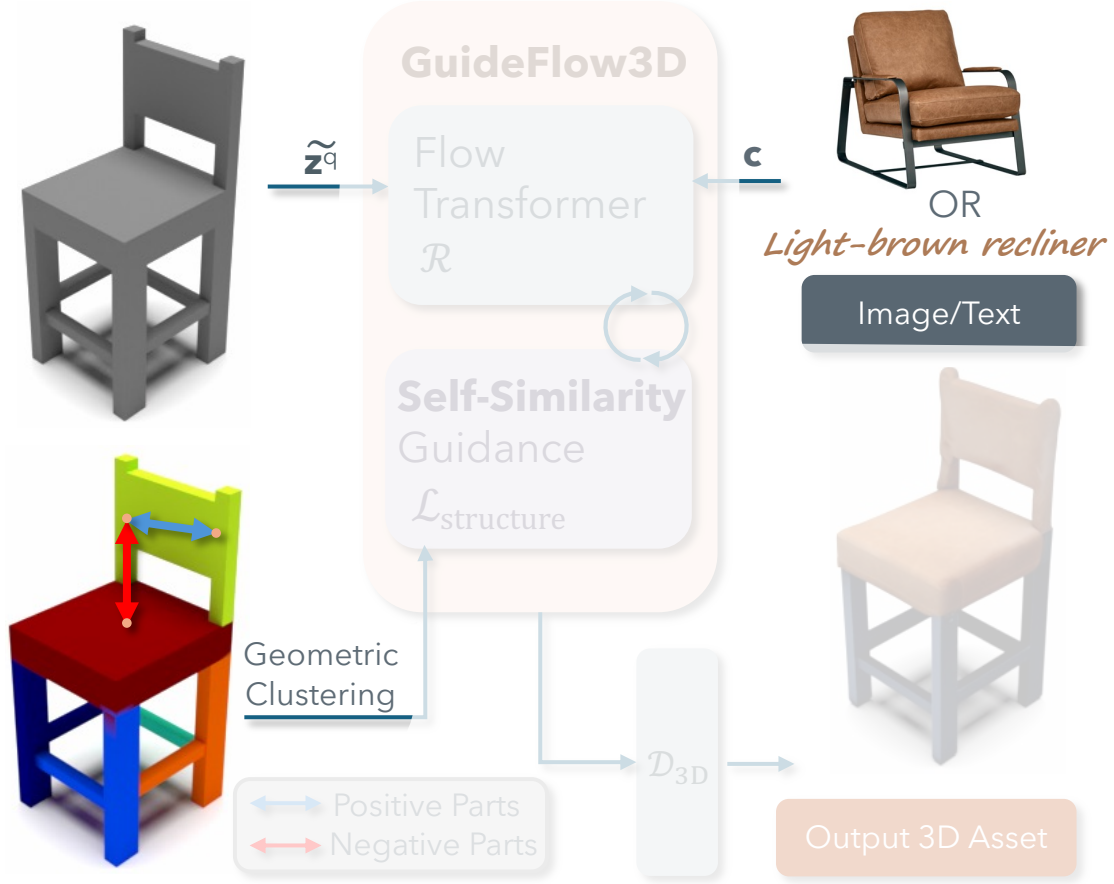

Promotes **local consistency** without **homogenizing appearance** globally.

Guided Flow For Appearance Transfer

Self-Similarity Guidance

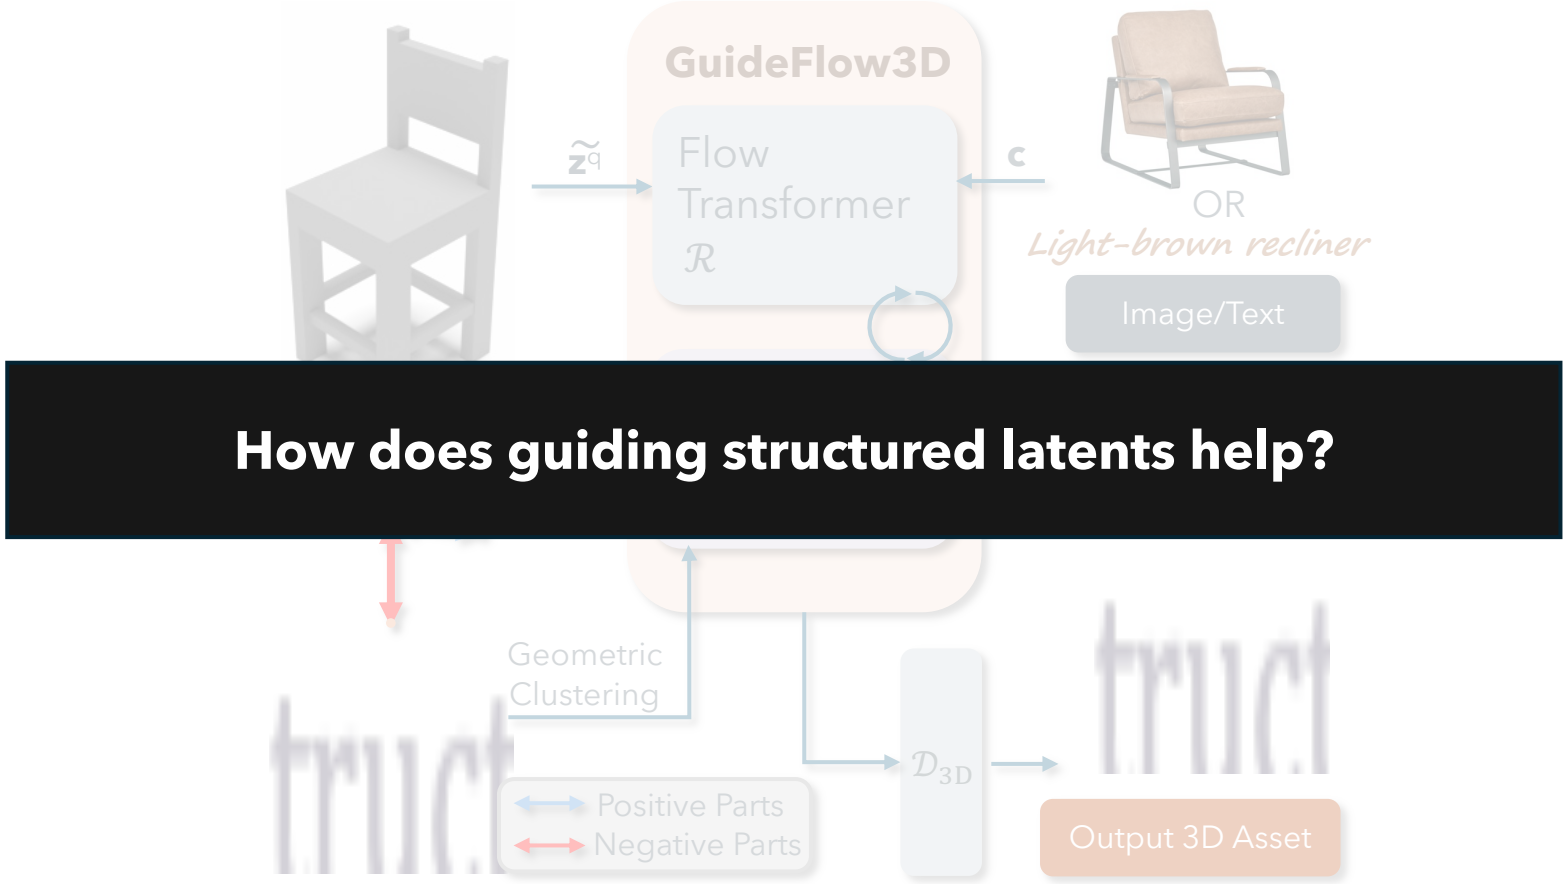

How does guiding structured latents help?

Promotes local consistency without homogenizing appearance globally.

Guided Flow For Appearance Transfer

Appearance Object

w/o Rectified Flow

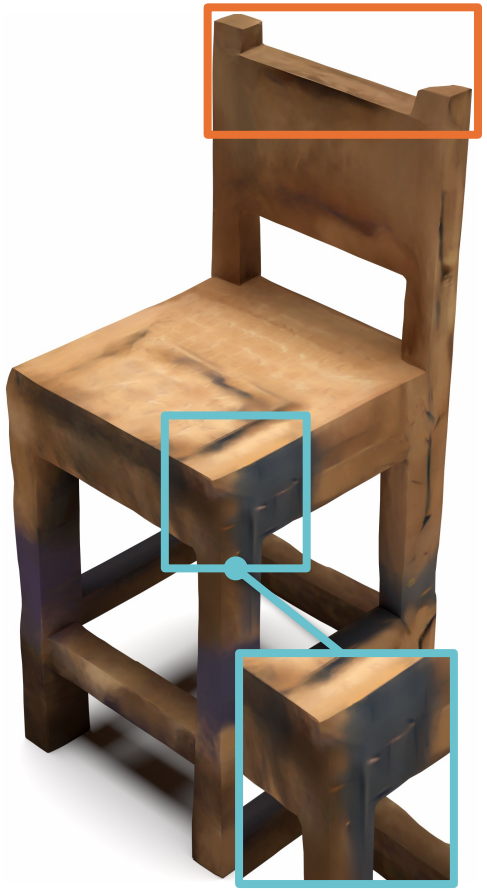

w/o Guidance

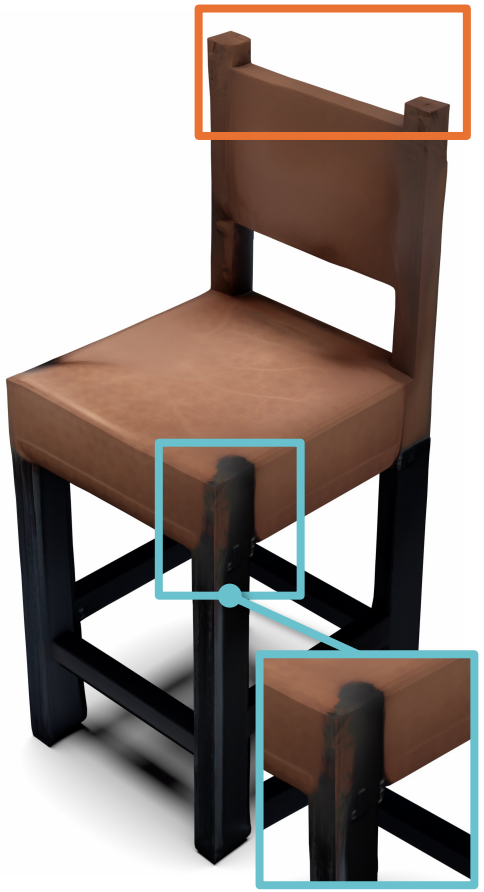

w/ GuideFlow3D

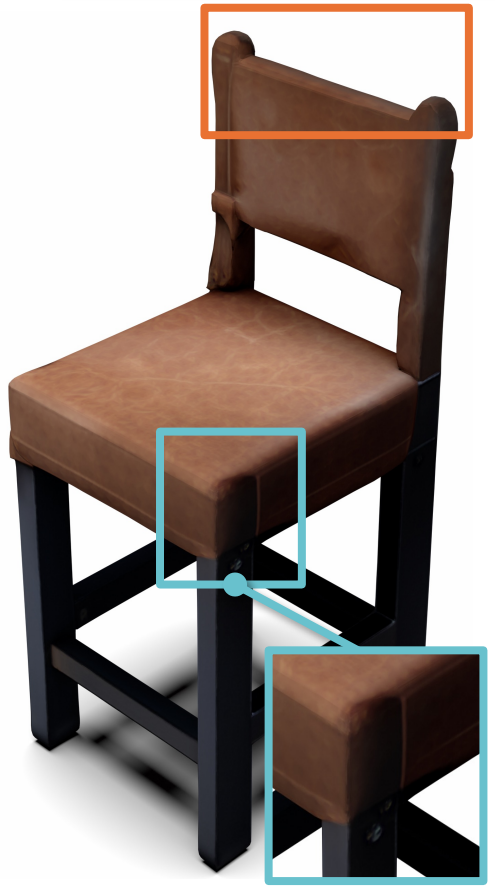

From plausible shapes to semantically grounded textures.

Results: Intra-Category with Image

Transfers **fine-grained texture** while preserving **part consistency** and **shape fidelity**.

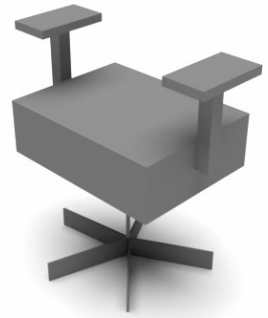

Input 3D Mesh

Appearance Image

UV Nearest Neighbor

Mamba-ST

Cross Image Attention

EASI-Tex

Trellis

GuideFlow3D *(Ours)*

Results: Intra-Category with Text

Aligns text-based appearance with object structure and semantics.

Input 3D Mesh

Reddish-brown rectangular wooden cabinet on short legs with two drawers, an open shelf, and a built-in power socket.

Appearance Text

UV Nearest Neighbor

SDXL + Cross Image Attention

Trellis

GuideFlow3D *(Ours)*

Results: Inter-Category with Image

Generalizes across category, retaining **realistic materials** and **geometric structure**.

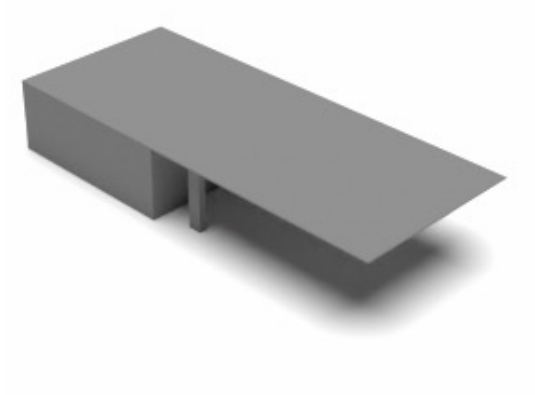

Input 3D Mesh

Appearance Image

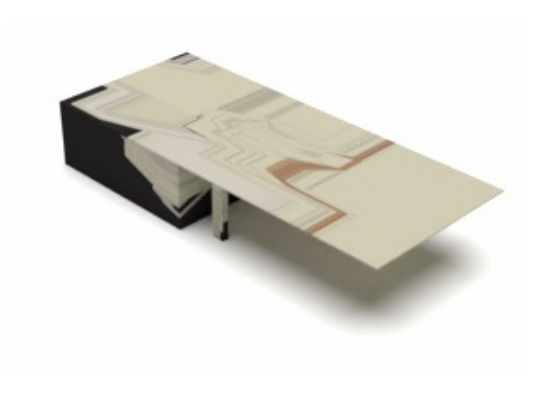

UV Nearest Neighbor

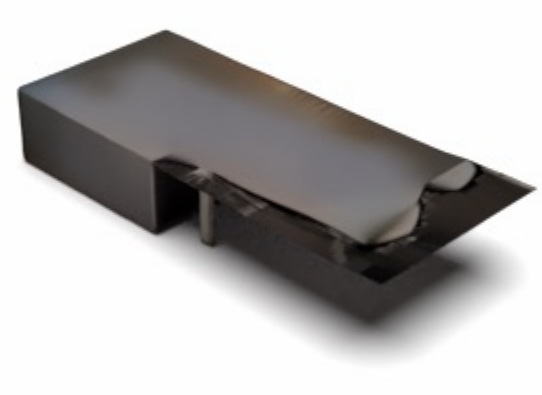

Mamba-ST

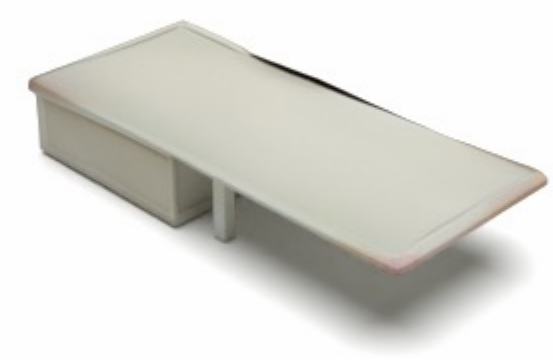

Cross Image Attention

EASI-Tex

Trellis

GuideFlow3D *(Ours)*

Results: Inter-Category with Text

Generalizes across categories from **abstract textual cues**.

Input 3D Mesh

Light rectangular coffee table with three-plank top, black metal accents, central drawer with ring handle, and thick black tapered legs.

Appearance Text

UV Nearest Neighbor

SDXL + Cross Image Attention

Trellis

GuideFlow3D *(Ours)*

Quantitative Evaluation

LLM Evaluation

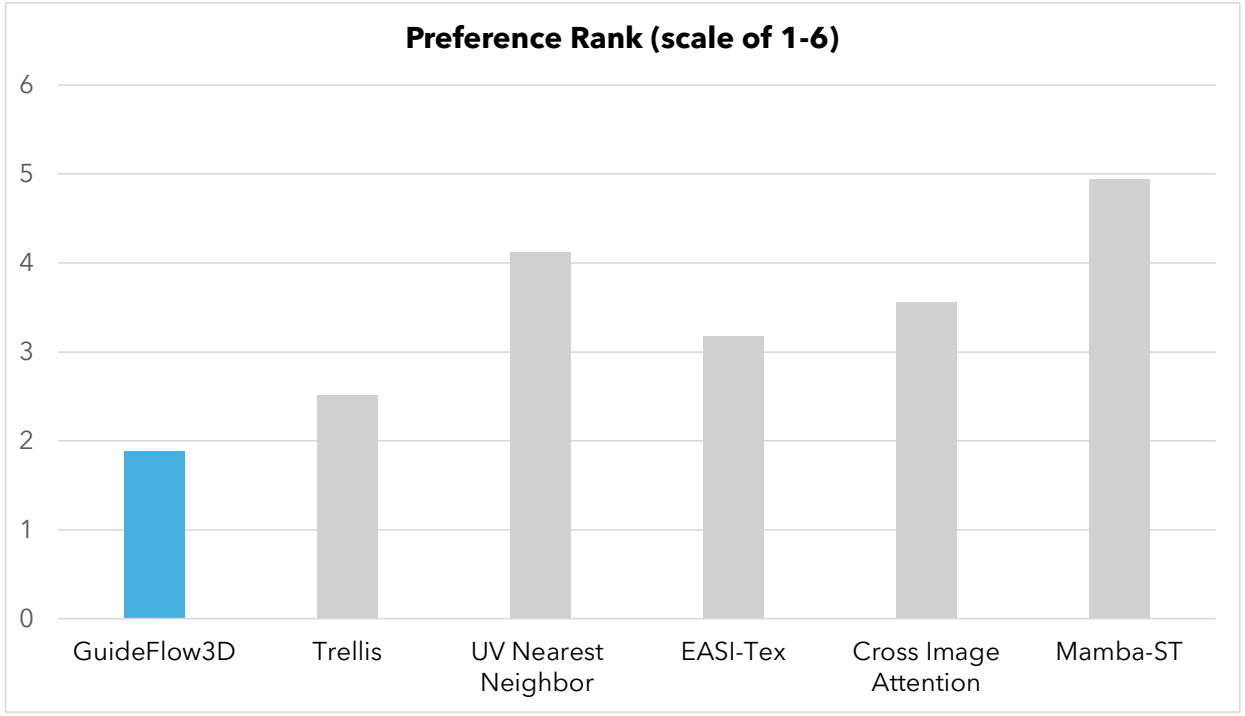

Human Evaluation

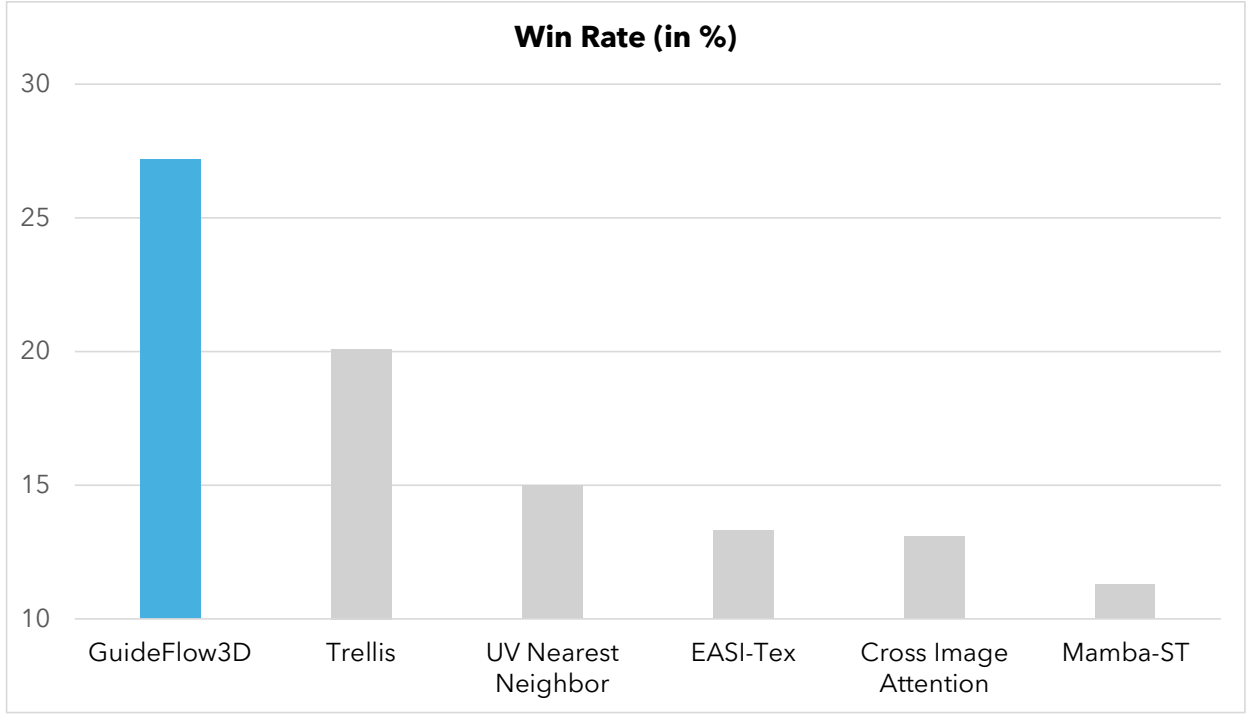

LLM rankings align with human preferences, confirming superior appearance transfer quality.

Application: In-the-Wild

Input 3D Mesh

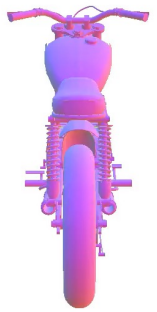
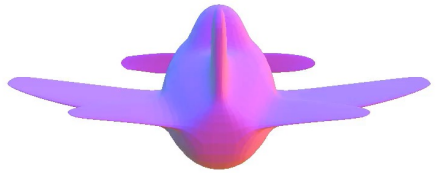


Appearance Object





Output 3D Asset





Application: Scene Editing

Input Scene

Appearance Objects

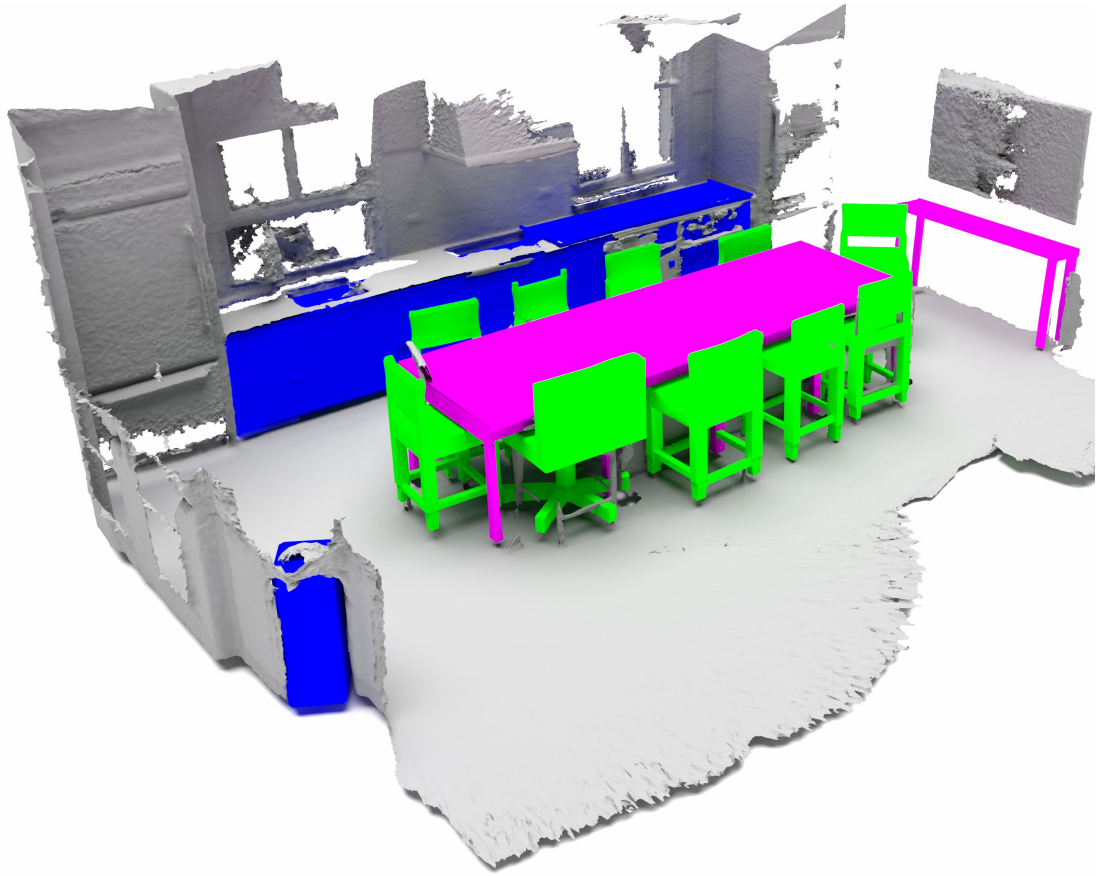

Application: Scene Editing

Input Scene

Appearance Objects

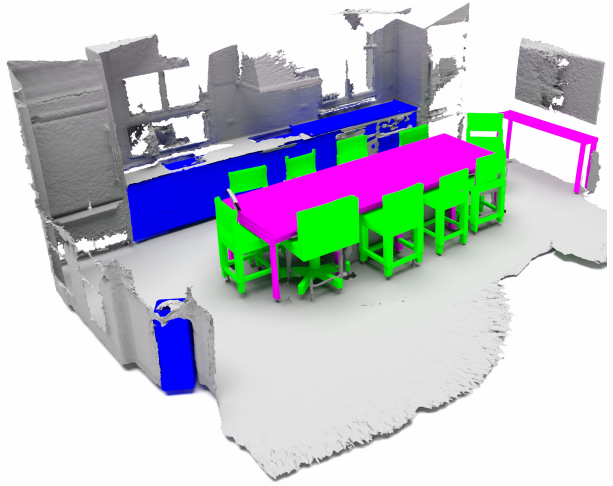

Output Scene

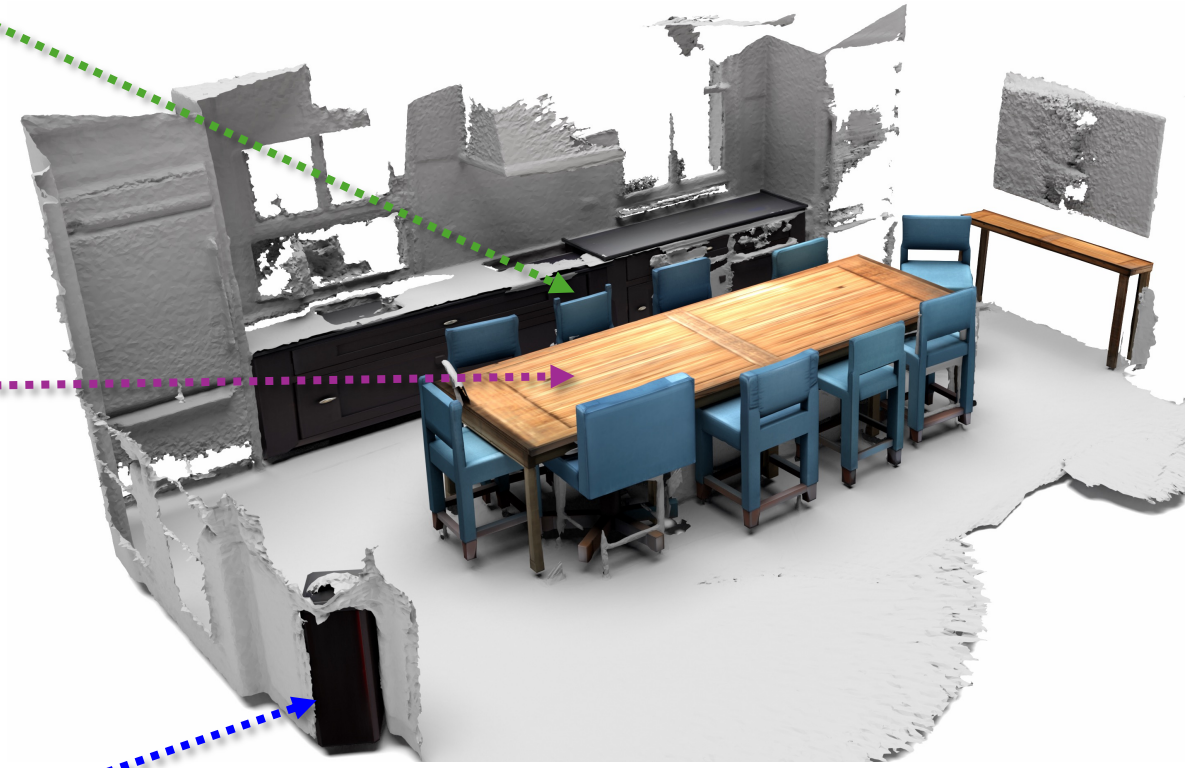

Scene Restyling with context-aware transfer

Limitations: Where Do We Still Fail?

Interpreting **abstract semantics without ambiguity** remains an open challenge!

Input 3D Mesh

Appearance Image

UV Nearest Neighbor

Mamba-ST

Cross Image Attention

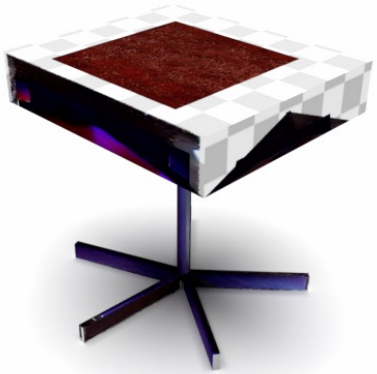

EASI-Tex

Trellis

GuideFlow3D *(Ours)*

Key Takeaways

- Novel framework for 3D appearance transfer that applies universal, differentiable guidance to a pretrained rectified flow model
- Training-free approach, generalizable to different appearance representations

Bring controllable creativity to 3D generative design!

Future Directions

- Can we train self-supervised model for fast and efficient inference?
- Can we, in principle, extend the guidance objectives to other tasks, eg, 3D reconstruction?

Stanford University

Gradient Spaces

École des Ponts ParisTech

GuideFlow3D

Optimization-Guided Rectified Flow For Appearance Transfer

Sayan Deb Sarkar [1] Sinisa Stekovic [2] Vincent Lepetit [2] Iro Armeni [1]

[1] Stanford University [2] ENPC, IP Paris

NeurIPS 2025

Paper | arXiv | Code | Poster

Project Page: https://sayands.github.io/guideflow3d/

Project Page

Poster Session

Fri 5 Dec
4:30-7:30pm