



TL;DR: MATLAT learns a material latent space that adapts pretrained image diffusion priors for fast, high-fidelity, multi-view-consistent PBR texturing of 3D meshes.

## Problem / Motivation

**Goal**  
Given a text prompt and a 3D mesh, generate relightable PBR textures: albedo, roughness, and metallic.



**Challenges in PBR Texturing**  
PBR texturing requires physical material maps beyond plausible colors, but high-quality PBR assets are scarce.



## Three Approaches for PBR Texture Generation

From Scratch (SC)	Score Distillation Sampling (SDS)	Multi-View Generation (MV)
<ul style="list-style-type: none"> <li>Train natively on PBR textures.</li> <li>Limited by scarce PBR data.</li> </ul>	<ul style="list-style-type: none"> <li>Uses pretrained image priors.</li> <li>Slow and often oversaturated.</li> </ul>	<ul style="list-style-type: none"> <li>Leverages pretrained image priors.</li> <li>Enables fast feed-forward inference.</li> </ul>

## What Multi-View Generation Must Solve

- How do we leverage pretrained RGB diffusion priors despite **extra material channels**?
- How do we preserve **multi-view consistency**?

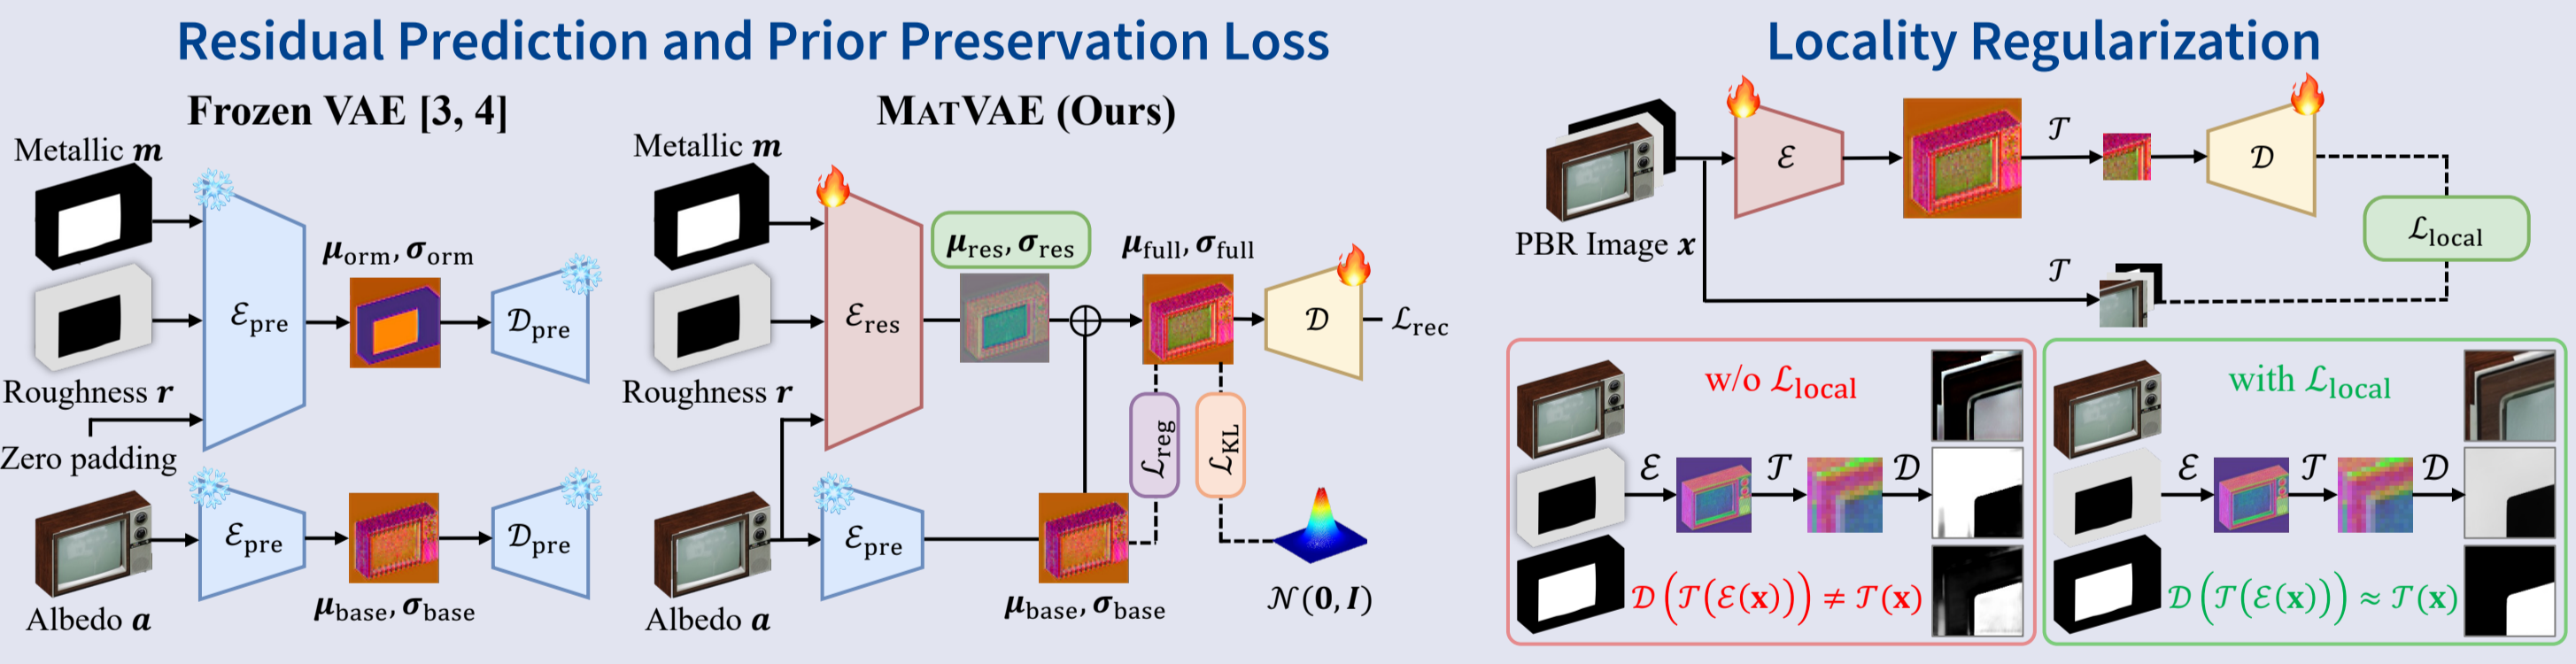
## High-Fidelity PBR Texture Generation



## Our Method: MATLAT

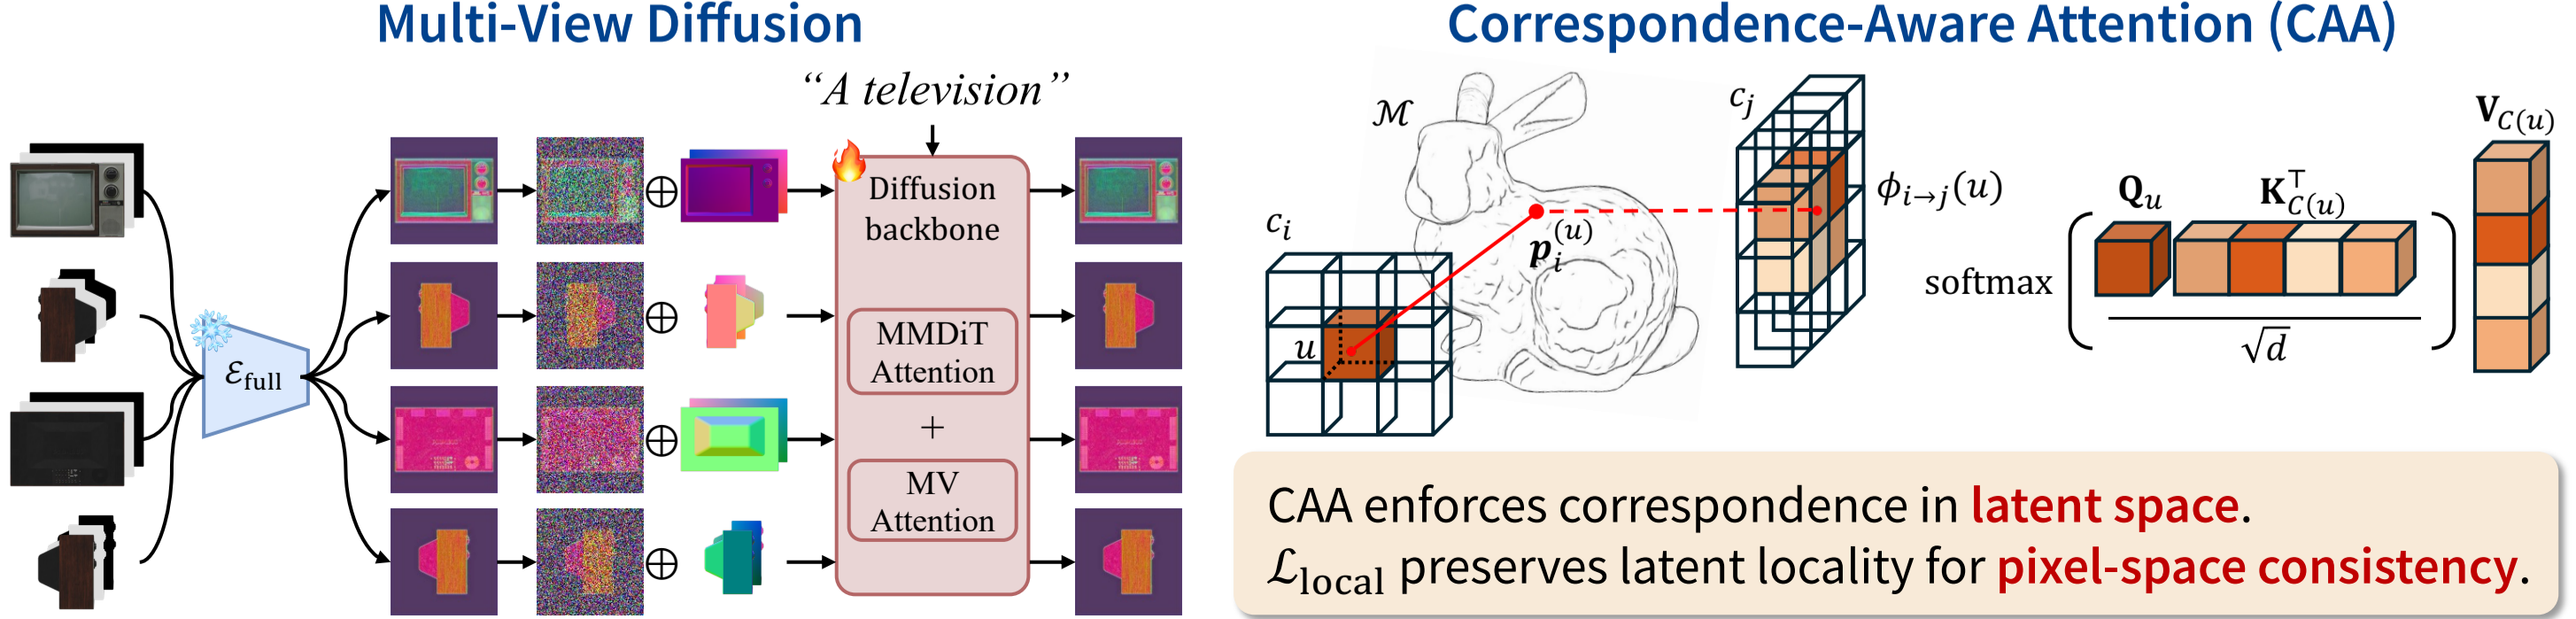
### Stage 1: Learning MATVAE

Unlike frozen-encoder approaches, we adapt the VAE to incorporate PBR channels with minimal latent shift.



### Stage 2: Finetuning Diffusion Prior

We fine-tune diffusion in the material latent space with CAA and  $\mathcal{L}_{local}$  to improve multi-view consistency.



## Results

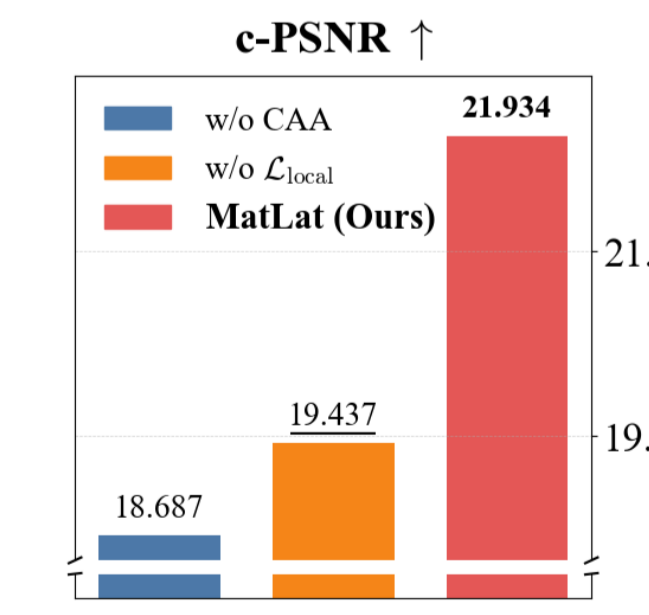
### Quantitative Comparison

Method	Type	Shaded		Albedo		Roughness	Metallic	Time↓
		FID <sub>CLIP</sub> ↓	CLIP↑	FID <sub>CLIP</sub> ↓	CLIP↑	RMSE↓	RMSE↓	
TexGaussian [1]	SC	6.025	0.301	12.119	0.299	<b>0.145</b>	0.243	73s
DreamMat [2]	SDS	<u>5.422</u>	0.311	<u>9.621</u>	0.311	0.167	0.165	2400s
MaterialAnything [3]	MV	6.582	<u>0.312</u>	12.691	<u>0.317</u>	0.233	0.200	500s
MaterialMVP [4]	MV	6.309	0.294	9.630	0.290	0.175	<b>0.133</b>	<u>35s</u>
<b>MATLAT (Ours)</b>	MV	<b>3.083</b>	<b>0.318</b>	<b>4.599</b>	<u>0.314</u>	<u>0.158</u>	<u>0.134</u>	<b>34s</b>

### Qualitative Results



### Ablation Study



### Relighting Examples



### References

- Xiong et al., TexGaussian: Octree-based 3DGS for PBR Material, CVPR 2025.
- Zhang et al., DreamMat: High-quality PBR Material Generation, ACM TOG 2024.
- Huang et al., Material Anything: Generating Materials for Any 3D Object via Diffusion, CVPR 2025.
- He et al., MaterialMVP: Illumination-Invariant Material Generation via Multi-view PBR Diffusion, ICCV 2025.
- Zhang and Agrawala, Transparent Image Layer Diffusion using Latent Transparency, ACM TOG 2024.
- Krishnan et al., Orchid: Image Latent Diffusion for Joint Appearance and Geometry Generation, ICCV 2025.