SUPRA-TEX: SUPER-ALIGNED 3D TEXTURE GENERATION

Anonymous authors
Paper under double-blind review



Figure 1: Super-aligned textured mesh generated by SUPRA-Tex. Our method preserves fine-grained appearance details from the reference image. ***(Best viewed with zoom.)***

ABSTRACT

Single-image 3D texture generation aims to dress a given 3D mesh with relightable materials while faithfully preserving the visual details of a reference image. Existing methods either introduce artifacts when fusing inconsistent multi-view predictions or struggle to preserve fine details through implicit image-feature conditioning. To resolve these challenges, we introduce **SUPRA-Tex**, built on the principle of *decoupling visible-material estimation, surface alignment, and observation-guided 3D completion*. This division of labor allows the 3D generator, trained on limited 3D data, to focus its capacity on preserving visible evidence and completing unobserved surfaces. Specifically, we first train an MLLM-based image editing model to predict lighting-separated PBR maps in the reference view. We then project these maps onto the visible surface and encode them with the pretrained material VAE encoder into *native observation latents* that share the generation target's latent space. An *interleaved flow transformer* holds these latents as fixed, read-only anchors while alternating unobserved-region and full-surface updates to generate complete material latents. This design achieves *super-alignment* with the reference, delivering *What You See Is What You Get* fidelity on visible regions and coherent completion on unobserved surfaces. Trained solely on public datasets, SUPRA-Tex achieves SOTA performance in both geometry-conditioned texturing and end-to-end asset generation, matching or even exceeding leading commercial systems on in-the-wild images. Code and models will be released.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1 INTRODUCTION

For applications such as product digitization and embodied AI, generating relightable materials for a given 3D mesh from a single reference image requires faithful preservation of the reference appearance. We refer to this requirement as *super-alignment*: accurately reproducing fine-grained text, logos, and patterns at their corresponding visible surface locations, while completing unobserved regions with materials consistent with the visible appearance.

Two prominent approaches to full-surface texturing are multi-view texturing and native 3D generation. Multi-view pipelines leverage geometry-conditioned 2D generative models to synthesize appearance across surrounding viewpoints, then project and fuse the resulting images onto the mesh (Tencent Hunyuan3D Team, 2025a). However, cross-view discrepancies can introduce seams and ghosting, while blending conflicting predictions blurs sharp reference details (Lai et al., 2026; Zeng et al., 2026).

Native 3D generation circumvents view-fusion artifacts by synthesizing appearance directly on 3D representations (Lai et al., 2026; Chen et al., 2026; Zeng et al., 2026; Xiang et al., 2026), yet preserving fine reference details remains challenging. For example, TRELLIS.2 employs a 3D material VAE to compress surface PBR properties into material latents, which are generated under geometry and image conditioning before being decoded into surface materials. Its image cross-attention provides no explicit correspondence between image details and 3D surface locations, leaving the generator to learn where those details belong. To improve spatial alignment, recent work lifts high-resolution image features into 3D through camera projection (Li et al., 2026a). Yet even with explicit spatial correspondence, these conditions remain visual features rather than material observations. Consequently, the 3D generator must simultaneously shoulder the responsibilities of inferring PBR properties from visual features, reconstructing fine reference details at their corresponding surface locations, and synthesizing materials for unobserved regions.

We argue that these responsibilities should not fall on a single 3D generator. Given the inherent scarcity of 3D training data, requiring the 3D model to learn these heterogeneous tasks jointly increases its learning burden. Instead, we offload visible-material estimation to 2D generative priors and establish explicit surface alignment through camera projection, separating both from the 3D generative task. Our key insight is to translate visible reference evidence into *native material observations* before 3D generation, then organize full-surface completion around this grounded evidence. Rather than forcing the 3D model to infer physical materials from abstract visual features, this formulation enables the generator to anchor directly on established material evidence, preserving reference details while focusing more of its capacity on completing unobserved surfaces.

Based on this principle, we introduce **SUPRA-Tex**, a framework for super-aligned, relightable 3D texture generation. First, we leverage strong 2D priors by training an MLLM-based image editing model to jointly predict lighting-separated PBR maps in the reference view in a single forward pass, while preserving fine visual details. Next, we project these maps onto valid visible surface locations and encode the resulting partial materials with the pretrained material VAE encoder into *native observation latents*. Sharing the target latent space and directly decodable by the pretrained material decoder, these latents provide spatially aligned material content rather than abstract visual features that the 3D backbone must translate into PBR properties. We also encode the spatial extent of these observations to guide completion.

To guide full-surface material generation, we design an *interleaved flow transformer* that reads the native observation latents as fixed, read-only reference memory at every block. Its blocks alternate between updating unobserved latent cells and the full surface, allocating more updates to cells without observations while retaining full-surface interactions for coherent completion. Since observations are indexed by 3D surface coordinates, the same interface incorporates additional views, including 2D-generated back views, and interactive 2D edits without retraining the 3D backbone. Through this division of labor, SUPRA-Tex delivers *What You See Is What You Get* fidelity in the reference view while completing the asset with globally coherent, relightable materials, as illustrated in Fig. 1.

Using only public datasets for training, SUPRA-Tex achieves SOTA results in both geometry-conditioned texturing and end-to-end asset generation. With synthetic back-view conditioning, SUPRA-Tex achieves reference fidelity and visual quality on real photographs comparable to or better

than those of the leading commercial systems in our comparisons. Code and models will be released for further research.

Our main contributions are as follows:

- Our key insight is to establish surface-aligned material observations before 3D generation, allowing the generator to focus on preserving reference details and completing unobserved surfaces. Built on this insight, SUPRA-Tex is trained solely on public datasets and, with synthetic back-view conditioning, achieves reference fidelity and visual quality comparable to or better than those of leading commercial systems in our comparisons.
- We present *native observation encoding*, which projects reference-view PBR estimates onto valid surface locations and encodes the resulting partial materials in the generation target's latent space, simplifying the learning of reference appearance preservation.
- We propose an *interleaved flow transformer* that uses fixed, read-only observation anchors and alternates unobserved-cell and full-surface updates for reference-faithful, coherent material completion.

2 Related Work

Texture generation has progressed from multi-view synthesis and fusion to native surface generation. Early methods such as TEXTure and Paint3D (Richardson et al., 2023; Zeng et al., 2024a) paint meshes view-by-view with diffusion models. MVPaint (Cheng et al., 2025) improves cross-view coherence, while Hunyuan3D (Tencent Hunyuan3D Team, 2025a; Lai et al., 2025) extends multi-view generation to PBR materials; fusing inconsistent views can nevertheless blur sharp reference details. TRELLIS and TRELLIS.2 (Xiang et al., 2025; 2026) advance native generation through structured 3D latents, with the latter modeling complete PBR materials. TEXGen (Yu et al., 2024) uses partial UV maps, NaTex (Lai et al., 2026) encodes initial textures for refinement, and TEXTRIX (Zeng et al., 2026) encodes projected appearance with its pretrained attribute VAE. Pixal3D (Li et al., 2026a) aligns image features with 3D geometry through camera-frame back-projection. Building on these advances, SUPRA-Tex encodes lighting-separated PBR estimates as native material latents and preserves them as read-only anchors for interleaved completion. Additional related work is discussed in Appendix H.

3 Method

Given a reference image I, a mesh $\mathcal{M}$, and its camera projection π, SUPRA-Tex generates full-surface, relightable PBR materials. Figure 2 summarizes the three stages: reference-view material estimation, surface-aligned native observation encoding, and observation-guided 3D completion. We first introduce the material representation and upstream geometry model, then describe these stages and their extension to multi-view inputs. Additional method, implementation, and training details are provided in the supplementary material.

3.1 Preliminaries

Surface PBR and material latents. We represent surface appearance through physically based rendering (PBR) properties over the mesh $\mathcal{M}$. At any surface location $\mathbf{p} \in \mathcal{M}$, these properties comprise albedo $\mathbf{a} \in \mathbb{R}^3$, metallic $m \in [0,1]$, and roughness $r \in [0,1]$. We build on the O-Voxel representation of TRELLIS.2 (Xiang et al., 2026), which additionally includes opacity $\alpha \in [0,1]$, giving the material tuple $T(\mathbf{p}) = (\mathbf{a}(\mathbf{p}), m(\mathbf{p}), r(\mathbf{p}), \alpha(\mathbf{p}))$. Separate sparse VAEs compress 3D geometry and surface materials into geometry latents $G = \{g_s\}_{s\in\mathcal{S}}$ and material latents $Z = \{z_s\}_{s\in\mathcal{S}}$ over active coordinates $\mathcal{S}$. For sparse surface attributes $\mathcal{P}$, the pretrained material VAE encoder E_{mat} yields Z; conditioned on G, the material generation stage models PBR materials directly in the native 3D space by generating Z, which the decoder D_{mat} restores onto the mesh surface.

Geometry foundation and camera projection. Our framework takes a mesh $\mathcal{M}$ and camera projection π, either provided as ground truth or predicted from the reference image. We use the open-source TRELLIS.2-based Pixal3D variant (Li et al., 2026a; Xiang et al., 2026), denoted Pixal3D*,

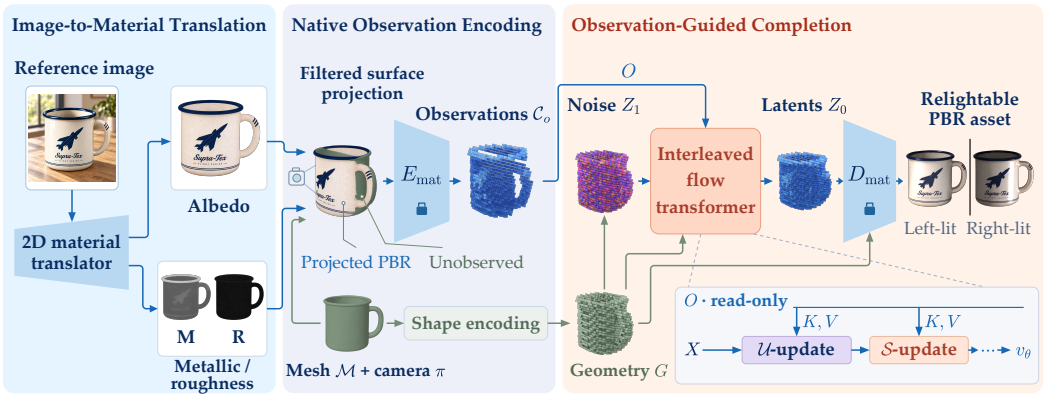

Figure 2: Overview of SUPRA-Tex: reference-view PBR predictions are projected onto the mesh and encoded as native observation latents, which guide an interleaved flow transformer to generate full-surface, relightable materials.

which supports geometry and material generation. Its geometry module predicts the reference-aligned mesh for end-to-end generation.

3.2 Image-to-Material Translation

To reduce the learning burden on the 3D generator, we separate reference-view material estimation from 3D surface completion. Specifically, we train an MLLM-based image editing model F_{2D} to infer reference-view albedo and metallic-roughness (MR) maps directly from the input photograph:

$$\widehat{T}_{\text{2D}} = F_{\text{2D}}(I) = (\widehat{\mathbf{a}}, \widehat{m}, \widehat{r}), \tag{1}$$

where $\widehat{\mathbf{a}}$ denotes the albedo map, and $\widehat{m}$ and $\widehat{r}$ denote the metallic and roughness maps. To reuse the 3-channel pretrained image VAE, we pack $\widehat{m}$ and $\widehat{r}$ with a constant channel into a 3-channel map, which also fulfills the input format expected by the pretrained material encoder.

Joint alignment and one-step inference. To coordinate predictions across channels while retaining pretrained image priors, F_{2D} adapts a multimodal DiT backbone (Super Intelligence Team, 2026) via *factorized modality-axis joint attention*, jointly predicting albedo, metallic–roughness, and auxiliary normal maps. This attention exchanges features across output modalities at corresponding spatial locations. We train the frontend on RGB images paired with ground-truth material and normal maps rendered from Objaverse, using a latent flow-matching objective at the fixed maximum-noise endpoint $\sigma = 1.0$. At inference, a single DiT forward pass at this endpoint (only 1.2 s on a single H800 GPU) estimates clean material latents, which the VAE decodes into 1024×1024 material maps without iterative denoising.

This design brings two key advantages. First, it leverages powerful 2D image priors to disentangle illumination from intrinsic materials at full reference resolution. Second, it shields the downstream 3D generator from abstract visual features, allowing the 3D model to operate purely on physical PBR observations and focus its capacity on reference preservation and unobserved-surface completion.

3.3 Native Observation Encoding

To supply physical observations to the 3D generator in a representation it natively understands, we bridge the estimated reference-view PBR maps $\widehat{T}_{\text{2D}}$ directly to the surface of $\mathcal{M}$ using the pretrained material encoder E_{mat}.

Filtered surface projection. Through the camera projection π, we associate each foreground pixel $\mathbf{u}$ of $\widehat{T}_{\text{2D}}$ with its first visible surface hit $\mathbf{p} \in \mathcal{M}$. At inference, foreground-mask erosion and local depth-discontinuity filtering further remove unreliable boundary correspondences. The retained correspondences form the observation set $\mathcal{H}_{\text{valid}}$, shared by all PBR channels. We voxelize the retained materials into a partial surface representation:

$$\mathcal{P}_o = \text{Voxelize}\Big(\Big\{\Big(\mathbf{p}, \widehat{T}_{\text{2D}}(\mathbf{u})\Big) \mid (\mathbf{u}, \mathbf{p}) \in \mathcal{H}_{\text{valid}}\Big\}\Big), \tag{2}$$

where $\mathbf{u}$ is a 2D pixel coordinate and $\mathbf{p}$ is its corresponding surface location on $\mathcal{M}$. Surface regions without valid observations are omitted from $\mathcal{P}_o$.

Sparse coordinate encoding. The partial surface attributes $\mathcal{P}_o$ match the input domain of the pretrained material VAE, allowing direct ingestion by E_{mat} without intermediate adapters:

$$E_{\text{mat}}(\mathcal{P}_o) = \{(s, z_s^o) \mid s \in \mathcal{S}_o\}, \qquad \mathcal{S}_o \subseteq \mathcal{S}, \tag{3}$$

using the posterior mean as output. We designate $\{z_s^o\}$ the *native observation latents*, where $\mathcal{S}_o$ contains the latent coordinates that received observations.

Rather than zero-padding unobserved regions, which misleads the generator by imposing concrete physical semantics such as black albedo, we exploit the sparse coordinate architecture of E_{mat} by physically pruning unobserved voxels entirely. Passing strictly the visible support $\mathcal{P}_o$ into E_{mat} prevents fictitious artifacts on unobserved surfaces while preserving authentic local observations. This encoding runs once before sampling and remains fixed throughout generation.

Spatial observation metadata. At the latent resolution, each cell aggregates multiple surface voxels across its spatial span. Consequently, an observed cell in $\mathcal{S}_o$ is rarely covered in its entirety from a single viewpoint. The generator therefore requires explicit quantification of how thoroughly each cell is supported by observation and its proximity to unobserved regions. Let $\mathcal{V}_s$ be the surface voxels present in cell s and $\mathcal{V}_o$ the filtered, deduplicated set of observed voxels. Coverage is

$$\rho_s = \frac{|\mathcal{V}_s \cap \mathcal{V}_o|}{|\mathcal{V}_s|}, \qquad s \in \mathcal{S}_o, \tag{4}$$

ensuring that multiple pixel projections onto a single surface voxel do not artificially inflate coverage. Let $\mathcal{U} = \mathcal{S} \setminus \mathcal{S}_o$ be the unobserved cells and $\mathcal{G}_{\mathcal{S}}$ the six-neighbor graph over $\mathcal{S}$; with $\mathcal{U}_s \subseteq \mathcal{U}$ the unobserved cells reachable from s, the boundary distance is

$$d_s = \begin{cases} \min\limits_{q \in \mathcal{U}_s} \text{dist}_{\mathcal{G}_{\mathcal{S}}}(s, q), & \mathcal{U}_s \neq \varnothing, \\ -1, & \mathcal{U}_s = \varnothing, \end{cases} \tag{5}$$

measured in graph steps, with -1 marking cells from which no unobserved region is reachable. Both quantities describe the extent of observation, not the accuracy of the 2D estimate. These native observation latents and their spatial extent form the native observation representation

$$\mathcal{C}_o = \{(s, z_s^o, \rho_s, d_s) \mid s \in \mathcal{S}_o\}, \tag{6}$$

which encapsulates native material observations alongside their physical extent without fictitious padding, ready to anchor full-surface generation. Because anchors are indexed in 3D coordinates, the same interface can combine observations from additional or edited views without retraining the generator (Appendix G.1).

3.4 Observation-Guided Completion

We generate complete material latents via flow matching over the mesh $\mathcal{M}$, predicting a velocity field at time t from the current noisy material latents Z_t, geometry latents G, and native observation representation $\mathcal{C}_o$ (Lipman et al., 2023). Native observation latents provide spatially aligned material evidence for reconstructing visible appearance. The generator must learn to use this evidence while completing unobserved surfaces coherently. Pixal3D* injects projected image features through additive residuals and updates all surface tokens at every layer. We introduce an interleaved flow transformer that retains observations as fixed, read-only memory and alternates unobserved-cell and full-surface updates. This schedule assigns more hidden-state updates to unobserved cells while allowing visible materials to be refined during full-surface generation.

Interleaved flow transformer. The generator operates across two distinct coordinate supports: the full active surface $\mathcal{S}$ and the observed subset $\mathcal{S}_o \subseteq \mathcal{S}$. We maintain two separate streams: a dynamic generation stream X tracking the hidden state of Z_t over $\mathcal{S}$, and a static observation stream O acting as immutable reference memory over $\mathcal{S}_o$. At observed coordinates, O_s stores partial evidence while X_s predicts the complete cell. Both streams use the same 3D coordinate frame and a shared input projection W_{in} over material and geometry latents:

$$\begin{aligned} X_s^{(0)} &= W_{\text{in}}([Z_{t,s}; g_s]), && s \in \mathcal{S}, \\ O_s &= W_{\text{in}}([\bar{z}_s^o; g_s]) + W_{\text{meta}}([\rho_s; d_s]) + b_o, && s \in \mathcal{S}_o, \end{aligned} \tag{7}$$

where $\bar{z}_s^o$ scales the native observation latents to match the normalization of Z_t, W_{meta} is a learnable affine projection of spatial metadata, and b_o is a learned observation-role bias initialized to zero.

Across the Transformer blocks, O produces no queries and receives no updates with $O^{(\ell+1)} = O$ across blocks, serving as read-only memory supplying keys and values to every layer. To concentrate generative capacity on unobserved surfaces, query coordinates alternate across blocks $\ell = 0, \ldots, L-1$: $\mathcal{J}_\ell = \mathcal{U}$ for even ℓ over unobserved cells $\mathcal{U} = \mathcal{S} \setminus \mathcal{S}_o$, and $\mathcal{J}_\ell = \mathcal{S}$ for odd ℓ over the full surface to maintain global coherence. Let $\widetilde{X}^{(\ell)} = N_\ell(X^{(\ell)}, t)$ and $\widetilde{O}^{(\ell)} = N_\ell(O, 0)$ denote the streams after normalization and time modulation, where observations carry a fixed clean-time indicator. The cross-stream attention is formulated as:

$$Q_\ell = q_\ell\Big(\widetilde{X}^{(\ell)}_{\mathcal{J}_\ell}\Big), \quad K_\ell = \Big[k_\ell\Big(\widetilde{O}^{(\ell)}\Big); k_\ell\Big(\widetilde{X}^{(\ell)}\Big)\Big], \quad V_\ell = \Big[v_\ell\Big(\widetilde{O}^{(\ell)}\Big); v_\ell\Big(\widetilde{X}^{(\ell)}\Big)\Big], \\ A_\ell = \mathrm{softmax}\Big(Q_\ell K_\ell^\top / \sqrt{d_h}\Big) V_\ell, \tag{8}$$

where $d_h = 128$ is the per-head dimension and q_ℓ, k_ℓ apply 3D RoPE to latent coordinates, enabling attention based on relative position and content. Residual updates affect only $\mathcal{J}_\ell$; other generation tokens are unchanged in that block, while O remains fixed throughout. An output projection maps the terminal state $X^{(L)}$ to predicted velocities for all of $\mathcal{S}$. We train the generator with a flow-matching loss over all valid surface tokens and channels, including observed coordinates. During training, observations encoded from visible ground-truth materials remain fixed conditioning inputs along the flow path, while the velocity targets are derived from complete ground-truth material latents.

4 Experiments

Dataset and implementation details, baseline configurations, evaluation protocols, and additional analyses are provided in the supplementary material.

Datasets. Our 3D texture generator is trained on approximately 570k unique 3D assets curated from public object repositories and human-scan datasets. We use a dual-track curation pipeline to filter web-scale noise while retaining high-value geometries. We evaluate on two datasets with no overlap with the training set: DTC-Tex, comprising 100 texture-rich product scans selected from the Digital Twin Catalog (Dong et al., 2025), and Complex-Tex, comprising 110 Sketchfab (Sketchfab, n.d.) meshes with denser inscriptions, finer patterns, and more varied shapes.

Implementation details. The 2D frontend adapts FireRed-Image-Edit-1.1 (Super Intelligence Team, 2026) at 1024×1024 resolution. The 3D generator uses 60 Transformer blocks, with self-attention initialized from TRELLIS.2 and the remaining components initialized randomly; the pretrained material encoder and decoder remain frozen. We train the generator with a global batch size of 576 using AdamW on 64 GPUs. Inference uses 25 Euler steps without classifier-free guidance.

4.1 Texture Generation with Given Geometry

Setting and baselines. Each method textures the ground-truth mesh from a single forward-lit reference image. We record the ground-truth lighting and camera poses and render each textured result under these same conditions at the reference view and 15 novel views, isolating material quality from geometry, pose, and illumination differences. Table 1 compares ten baselines spanning native 3D generation and dedicated texturing methods. PBR methods are additionally evaluated on albedo, metallic, and roughness.

Results. Table 1 shows that SUPRA-Tex leads across all shaded and PBR metrics on both datasets. Compared with feature-conditioned Pixal3D*, shaded PSNR improves by 1.7 dB and 1.2 dB, and albedo PSNR_{fg} by 2.0 dB and 0.9 dB on DTC-Tex and Complex-Tex, respectively. The gains over multi-view PBR methods are larger: LumiTex trails by more than 2 dB in shaded PSNR and by over 1.5 dB in albedo PSNR on both datasets. The albedo gains and lowest metallic/roughness MAE show that the improvements extend to material attributes. On DTC-Tex/Complex-Tex, reference-view shaded PSNR_{fg} improves by 2.56/0.95 dB over TEXGen, while the mean over 15 novel views improves by 1.76/1.09 dB over Pixal3D* (Table 7). These gains support reference fidelity together with improved appearance across viewpoints. Figure 3 shows that SUPRA-Tex better preserves the mug's fine lettering and the plate's intricate patterns than multi-view and feature-conditioned baselines, while novel-view renderings show coherent appearance on the completed surfaces.

Table 1: PBR texture generation with given geometry. Per-view metrics are averaged over 16 views; fg denotes the GT foreground. A.: albedo; Met./Rough.: metallic/roughness MAE.

Method	Shaded						PBR					
	PSNR ↑	PSNR_{fg} ↑	SSIM ↑	LPIPS ↓	CLIP-I ↑	$\text{FD}_{\text{DINOv3}}$ ↓	A. PSNR ↑	A. PSNR_{fg} ↑	A. SSIM ↑	A. LPIPS ↓	Met. ↓	Rough. ↓
DTC-Tex												
Step1X-3D	23.12	14.21	0.914	0.083	0.851	33.30	–	–	–	–	–	–
MV-Adapter	24.64	15.73	0.917	0.077	0.858	33.17	–	–	–	–	–	–
Ink3D	25.82	17.08	0.921	0.076	0.859	35.01	–	–	–	–	–	–
UniTEX	25.42	16.54	0.925	0.071	0.877	29.24	–	–	–	–	–	–
MVPainter	26.44	17.59	0.924	0.077	0.856	33.40	24.17	15.30	0.911	0.089	0.121	0.119
TEXGen	27.24	18.33	0.932	0.072	0.848	51.11	–	–	–	–	–	–
LumiTex	28.29	19.38	0.939	0.062	0.892	24.42	26.30	17.39	0.928	0.069	0.352	0.098
Hunyuan3D 2.1	26.80	17.89	0.926	0.070	0.877	28.18	24.00	15.09	0.910	0.079	0.142	0.079
TRELLIS.2	28.60	19.73	0.935	0.062	0.880	32.73	25.78	16.90	0.921	0.070	0.063	0.059
Pixal3D*	29.19	20.47	0.950	0.050	0.908	20.35	26.39	17.60	0.937	0.057	0.051	0.076
Ours	**30.89**	**22.34**	**0.961**	**0.042**	**0.927**	**16.51**	**28.20**	**19.55**	**0.951**	**0.048**	**0.043**	**0.041**
Complex-Tex												
Step1X-3D	23.61	13.97	0.924	0.066	0.848	30.93	–	–	–	–	–	–
MV-Adapter	24.59	14.95	0.926	0.064	0.849	31.19	–	–	–	–	–	–
Ink3D	26.15	16.69	0.935	0.060	0.864	30.86	–	–	–	–	–	–
UniTEX	25.78	16.19	0.940	0.056	0.872	27.42	–	–	–	–	–	–
MVPainter	25.72	16.19	0.933	0.067	0.846	33.87	23.05	13.51	0.922	0.079	0.228	0.188
TEXGen	27.90	18.25	0.945	0.053	0.871	32.89	–	–	–	–	–	–
LumiTex	27.77	18.21	0.947	0.051	0.887	23.29	25.26	15.68	0.938	0.060	0.256	0.181
Hunyuan3D 2.1	25.83	16.19	0.931	0.063	0.857	31.05	22.89	13.25	0.918	0.074	0.247	0.156
TRELLIS.2	27.97	18.41	0.945	0.051	0.876	27.42	25.77	16.17	0.935	0.058	0.131	0.111
Pixal3D*	28.84	19.44	0.950	0.046	0.899	19.50	26.62	17.25	0.935	0.056	0.145	0.122
Ours	**30.05**	**20.64**	**0.960**	**0.039**	**0.915**	**17.36**	**27.61**	**18.12**	**0.952**	**0.046**	**0.109**	**0.084**

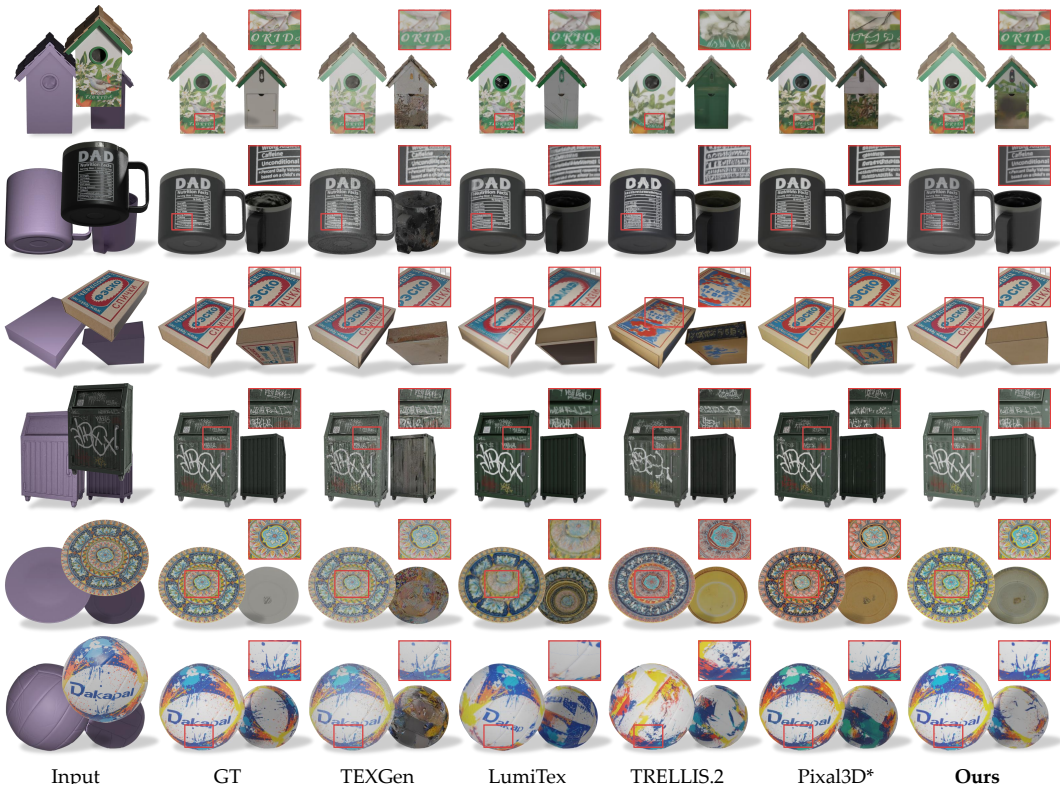

Figure 3: Qualitative comparison with given geometry. Enlarged crops compare reference-view details; novel views show the completed appearance.

4.2 End-to-End Single-Image-to-3D Generation

Setting and baselines. Each method generates both geometry and appearance end-to-end from a single reference image, without access to ground-truth geometry or cameras during generation.

Table 2: End-to-end single-image-to-3D generation. Entries report DTC-Tex / Complex-Tex. Reconstruction metrics use four views, except reference-view CLIP-I; distribution metrics compare 16-view rendered sets.

Method	Reconstruction					Distribution			
	PSNR ↑	PSNR_{fg} ↑	SSIM ↑	LPIPS ↓	CLIP-I ↑	$\text{FD}_{\text{DINOv3}}$ ↓	FD_{CLIP} ↓	FID ↓	CMMD ↓
Step1X-3D	19.70 / 20.66	12.00 / 12.36	0.894 / 0.900	0.115 / 0.099	0.844 / 0.882	33.4 / 26.3	19.6 / 17.9	33.5 / 31.4	14.40 / 8.73
Hunyuan3D 2.1	21.97 / 21.87	14.52 / 13.83	0.904 / 0.908	0.095 / 0.087	0.835 / 0.843	34.8 / 31.2	22.7 / 23.2	33.6 / 35.5	12.78 / 11.39
UniLat3D†	19.78 / 20.41	11.89 / 11.89	0.900 / 0.902	0.105 / 0.096	0.822 / 0.830	34.2 / 29.8	27.2 / 27.5	36.1 / 35.2	27.17 / 21.57
CUPID†	20.26 / 20.68	12.38 / 12.12	0.906 / 0.905	0.093 / 0.090	0.850 / 0.852	32.1 / 28.1	25.1 / 24.4	33.0 / 34.1	24.18 / 18.44
TRELLIS.2†	22.59 / 22.22	15.71 / 14.32	0.907 / 0.909	0.098 / 0.087	0.816 / 0.830	41.7 / 32.6	26.2 / 23.4	32.2 / 32.0	16.63 / 10.49
Pixal3D⋆	26.88 / 26.24	19.90 / 18.64	0.945 / 0.940	0.071 / 0.070	0.910 / 0.909	27.2 / 22.9	16.5 / 15.9	26.0 / 25.8	8.29 / 5.96
Ours	**28.65 / 28.73**	**20.79 / 20.24**	**0.967 / 0.963**	**0.025 / 0.029**	**0.971 / 0.968**	**21.1 / 19.4**	**9.9 / 9.3**	**23.0 / 23.3**	**4.97 / 4.00**

† Similarity-aligned to the ground-truth mesh.

We compare with TRELLIS.2, Pixal3D⋆, Hunyuan3D 2.1, Step1X-3D, UniLat3D, and CUPID (Xiang et al., 2026; Li et al., 2026a; Tencent Hunyuan3D Team, 2025a; Li et al., 2025a; Wu et al., 2026a; Huang et al., 2026). SUPRA-Tex and Pixal3D⋆ share geometry predicted by the latter, with separate appearance pipelines. Other baselines generate their own geometry. For scoring, outputs are brought into a common coordinate frame, with GT similarity alignment applied to methods marked † in Table 2; this alignment plays no role in generation.

Results. Table 2 shows that SUPRA-Tex leads across all 9 metrics on both datasets. Relative to Pixal3D⋆, four-view PSNR improves by 1.77 dB on DTC-Tex and 2.49 dB on Complex-Tex, foreground PSNR by 0.89 dB and 1.60 dB, and LPIPS decreases by 64.8% and 58.6%, respectively. Higher reference-view CLIP-I and lower 16-view distribution distances further indicate stronger input consistency and multi-view appearance quality. Since Pixal3D⋆ and SUPRA-Tex share predicted geometry, these results show that our appearance pipeline retains its advantage on predicted surfaces, complementing the given-geometry results in Table 1.

Generalization and commercial comparison. We evaluate SUPRA-Tex against five commercial systems on 45 real and synthetic images. We use two-view conditioning only for commercial comparisons, combining the reference image with a back view synthesized from it.



Figure 4: Comparison with commercial systems: each row shows the input and reference-view/back-view renderings from commercial systems and SUPRA-Tex.

Table 3: Comparison with commercial models. Preference reports our pairwise win rate (%); VLM reports mean scores out of 10.

System	Ours win rate (%) ↑			VLM (0–10) ↑		
	Align	PBR	Quality	Align	PBR	Quality
HY3.1	84.7	85.1	84.8	6.77	**8.67**	8.32
Rodin 2.5	75.8	72.3	75.5	6.81	8.64	**8.38**
Tripo H3.1	71.6	62.9	67.8	7.71	8.53	8.28
Meshy 7.1	60.8	54.7	57.6	7.74	8.61	8.34
Hi3D 3.0	51.1	44.9	51.1	**8.19**	**8.67**	8.37
Ours	–	–	–	8.08	**8.67**	8.34

Table 4: Observation ablations. Shad. and Alb. report shaded/albedo PSNR_{fg}; MR reports mean metallic/roughness MAE.

Variant	DTC-Tex			Complex-Tex		
	Shad. ↑	Alb. ↑	MR ↓	Shad. ↑	Alb. ↑	MR ↓
Residual injection	21.47	18.81	0.086	20.35	17.18	0.136
Joint-update attention	22.32	19.40	0.043	20.63	17.96	0.097
Fixed MR ($M=0$, $R=0.5$)	21.90	19.53	0.064	19.95	18.15	0.189
Dataset-mean MR	22.08	19.53	0.232	20.07	18.13	0.258
Ours	22.34	19.55	0.042	20.64	18.12	0.096
Oracle PBR (GT)	24.65	22.17	0.011	21.76	20.63	0.030

Without ground-truth geometry or cameras, paired human preferences and VLM ratings assess input alignment, PBR plausibility, and visual quality. Trained solely on public data, SUPRA-Tex receives 51.1%–84.8% preference in human visual-quality comparisons and ties for the highest VLM PBR score of 8.67 (Table 3). Figure 4 illustrates reference-detail preservation and coherent surface completion.

4.3 Ablation and Analysis

We evaluate the observation-injection architecture and material content under the given-geometry protocol of Sec. 4.1. The architecture variants share training data, initialization protocol, optimization settings, training budget, and checkpoint-selection rule. All comparisons use the same geometry, cameras, valid observation support, and sampling seed.

How observations are used. Table 4 compares three conditioning designs with the same predicted PBR observations. *Residual injection* adds encoded observations through layerwise residuals. *Joint-update attention* updates observation states in selected layers and all generation states in every layer. In contrast, SUPRA-Tex keeps observations read-only and alternates unobserved-cell and full-surface updates (Sec. 3.4). It improves shaded PSNR_{fg} over residual injection by 0.87 dB on DTC-Tex and 0.29 dB on Complex-Tex, and matches joint-update attention in shaded fidelity while improving albedo PSNR_{fg} by 0.15 and 0.16 dB with lower MR MAE. Thus, the design preserves or improves reconstruction quality without spending computation on observation-state updates or full-surface queries at every block. The complete design also samples $1.20\times$ and $1.18\times$ faster than joint-update attention on the two datasets (Appendix D.3).

Effectiveness of 2D material estimation. To test whether the FireRed-based frontend supplies useful physical material estimates, we replace its predicted metallic and roughness with fixed values ($M = 0, R = 0.5$) or dataset means, keeping predicted albedo and the 3D generator unchanged (Table 4). Relative to fixed and dataset-mean MR, predicted MR improves shaded PSNR_{fg} by 0.44 and 0.26 dB on DTC-Tex, and 0.69 and 0.57 dB on Complex-Tex, respectively, with lower MR MAE on both datasets. These gains show that the 2D estimates provide useful material information beyond albedo alone. Replacing all predicted PBR channels with ground truth (Oracle PBR) further improves shaded PSNR_{fg} by 2.31/1.12 dB and albedo PSNR_{fg} by 2.62/2.51 dB on DTC-Tex/Complex-Tex. The remaining gap indicates that improving 2D material estimation can further improve 3D texture quality without changing the completion model.

5 Conclusion

We presented SUPRA-Tex, which formulates single-image 3D texturing as surface completion guided by explicit material observations. By separating visible-material estimation and surface alignment from 3D generation and using native observation latents as fixed, read-only anchors, SUPRA-Tex preserves fine reference details while generating coherent, relightable materials across the full surface. Trained solely on public data, it achieves state-of-the-art results in geometry-conditioned texturing and end-to-end asset generation; with synthetic back views, it matches or exceeds leading commercial systems in reference fidelity and visual quality in our comparisons. These results demonstrate the value of spatially aligned, directly usable material observations for reconciling reference fidelity with full-surface generation.

Ethics Statement

Our models use public training datasets summarized in Appendix B. Use and redistribution of the source assets should respect their licenses and attribution requirements. Generated assets may reproduce copyrighted designs or personal likenesses, and unseen surfaces are model predictions rather than observed facts.

Reproducibility Statement

Section 3 describes the observation representation, generator, and training objective. Appendices A and B provide details on observation construction, inference, data curation, model architecture, and training. Appendices C–E document the evaluation protocols, baseline configurations, additional results, runtime measurements, and ablation settings. We use random seed 42 throughout our training and evaluation pipelines.

Use of AI Assistants

Generative AI tools assisted literature search and synthesis, manuscript organization and language editing, code development and debugging, and the writing of scripts for runtime measurement. In addition, some test input images used in comparisons with commercial models were generated using generative AI tools. The authors take full responsibility for the scientific design, experimental results, factual accuracy, and final content of the paper, including AI-assisted text and code.

References

Jingzhi Bao, Hongze Chen, Lingting Zhu, Chenyu Liu, Runze Zhang, Keyang Luo, Zeyu Hu, Weikai Chen, Yingda Yin, Xin Wang, Zehong Lin, Jun Zhang, and Xiaoguang Han. LumiTex: Towards high-fidelity PBR texture generation with illumination context. In *International Conference on Learning Representations*, 2026. URL https://arxiv.org/abs/2511.19437.

Raphael Bensadoun, Yanir Kleiman, Idan Azuri, Omri Harosh, Andrea Vedaldi, Natalia Neverova, and Oran Gafni. Meta 3D TextureGen: Fast and consistent texture generation for 3D objects. *arXiv preprint arXiv:2407.02430*, 2024. URL https://arxiv.org/abs/2407.02430.

Tianshi Cao, Karsten Kreis, Sanja Fidler, Nicholas Sharp, and Kangxue Yin. TexFusion: Synthesizing 3D textures with text-guided image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4169–4181, 2023. URL https://openaccess.thecvf.com/content/ICCV2023/html/Cao_TexFusion_Synthesizing_3D_Textures_with_Text-Guided_Image_Diffusion_Models_ICCV_2023_paper.html.

Chris Careaga and Yağız Aksoy. Colorful diffuse intrinsic image decomposition in the wild. *ACM Transactions on Graphics*, 43(6):178:1–178:12, 2024. doi: 10.1145/3687984. URL https://arxiv.org/abs/2409.13690.

Chia-Hao Chen, Yuan-Chen Guo, Zi-Xin Zou, Ze Yuan, Guan Luo, Xiaojuan Qi, Ding Liang, Yan-Pei Cao, and Song-Hai Zhang. Lafite: A generative latent field for 3D native texturing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19960–19971, 2026. URL https://openaccess.thecvf.com/content/CVPR2026/html/Chen_Lafite_A_Generative_Latent_Field_for_3D_Native_Texturing_CVPR_2026_paper.html.

Dave Zhenyu Chen, Yawar Siddiqui, Hsin-Ying Lee, Sergey Tulyakov, and Matthias Nießner. Text2Tex: Text-driven texture synthesis via diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 18558–18568, 2023. URL https://openaccess.thecvf.com/content/ICCV2023/html/Chen_Text2Tex_Text-driven_Texture_Synthesis_via_Diffusion_Models_ICCV_2023_paper.html.

Xi Chen, Sida Peng, Dongchen Yang, Yuan Liu, Bowen Pan, Chengfei Lv, and Xiaowei Zhou. IntrinsicAnything: Learning diffusion priors for inverse rendering under unknown illumination. In *European Conference on Computer Vision*, pp. 450–467, 2024. URL https://arxiv.org/abs/2404.11593.

Zhaoxi Chen, Jiaxiang Tang, Yuhao Dong, Ziang Cao, Fangzhou Hong, Yushi Lan, Tengfei Wang, Haozhe Xie, Tong Wu, Shunsuke Saito, Liang Pan, Dahua Lin, and Ziwei Liu. 3DTopia-XL: Scaling high-quality 3D asset generation via primitive diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26576–26586, 2025. URL https://openaccess.thecvf.com/content/CVPR2025/html/Chen_3DTopia-XL_Scaling_High-quality_3D_Asset_Generation_via_Primitive_Diffusion_CVPR_2025_paper.html.

Wei Cheng, Juncheng Mu, Xianfang Zeng, Xin Chen, Anqi Pang, Chi Zhang, Zhibin Wang, Bin Fu, Gang Yu, Ziwei Liu, and Liang Pan. MVPaint: Synchronized multi-view diffusion for painting anything 3D. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 585–594, 2025. URL https://openaccess.thecvf.com/content/CVPR2025/html/Cheng_MVPaint_Synchronized_Multi-View_Diffusion_for_Painting_Anything_3D_CVPR_2025_paper.html.

Jasmine Collins, Shubham Goel, Kenan Deng, Achleshwar Luthra, Leon Xu, Erhan Gundogdu, Xi Zhang, Tomas F. Yago Vicente, Thomas Dideriksen, Himanshu Arora, Matthieu Guillaumin, and Jitendra Malik. ABO: Dataset and benchmarks for real-world 3D object understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 21126–21136, 2022.

Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, Eli VanderBilt, Aniruddha Kembhavi, Carl Vondrick, Georgia Gkioxari, Kiana Ehsani, Ludwig Schmidt, and Ali Farhadi. Objaverse-XL: A universe of 10M+ 3D objects. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

Kangle Deng, Timothy Omernick, Alexander Weiss, Deva Ramanan, Jun-Yan Zhu, Tinghui Zhou, and Maneesh Agrawala. FlashTex: Fast relightable mesh texturing with LightControlNet. In *European Conference on Computer Vision*, 2024. URL https://eccv.ecva.net/virtual/2024/poster/537.

Zhao Dong, Ka Chen, Zhaoyang Lv, Hong-Xing Yu, Yunzhi Zhang, Cheng Zhang, Yufeng Zhu, Stephen Tian, Zhengqin Li, Geordie Moffatt, Sean Christofferson, James Fort, Xiaqing Pan, Mingfei Yan, Jiajun Wu, Carl Yuheng Ren, and Richard Newcombe. Digital twin catalog: A large-scale photorealistic 3D object digital twin dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025. URL https://openaccess.thecvf.com/content/CVPR2025/html/Dong_Digital_Twin_Catalog_A_Large-Scale_Photorealistic_3D_Object_Digital_Twin_CVPR_2025_paper.html.

Yifei Feng, Mingxin Yang, Shuhui Yang, Sheng Zhang, Jiaao Yu, Zibo Zhao, Yuhong Liu, Jie Jiang, and Chunchao Guo. RomanTex: Decoupling 3D-aware rotary positional embedded multi-attention network for texture synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 17203–17213, 2025. URL https://openaccess.thecvf.com/content/ICCV2025/html/Feng_RomanTex_Decoupling_3D-aware_Rotary_Positional_Embedded_Multi-Attention_Network_for_Texture_ICCV_2025_paper.html.

Yue Han, Chong Li, Zhening Liu, Cong Huang, Fang Deng, Yong Liu, Fangyun Wei, and Yan Lu. Ink3D: Sculpting 3D assets with extremely complex textures via video generative models. In *European Conference on Computer Vision*, 2026. URL https://arxiv.org/abs/2607.01222.

Zebin He, Mingxin Yang, Shuhui Yang, Yixuan Tang, Tao Wang, Kaihao Zhang, Guanying Chen, Yuhong Liu, Jie Jiang, Chunchao Guo, and Wenhan Luo. MaterialMVP: Illumination-invariant material generation via multi-view PBR diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 26294–26305, 2025a. URL https://openaccess.thecvf.

com/content/ICCV2025/html/He_MaterialMVP_Illumination-Invariant_Material_Generation_via_Multi-view_PBR_Diffusion_ICCV_2025_paper.html.

Zexin He, Tengfei Wang, Xin Huang, Xingang Pan, and Ziwei Liu. Neural LightRig: Unlocking accurate object normal and material estimation with multi-light diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26514–26524, 2025b. URL https://openaccess.thecvf.com/content/CVPR2025/html/He_Neural_LightRig_Unlocking_Accurate_Object_Normal_and_Material_Estimation_with_CVPR_2025_paper.html.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, 2017. URL https://arxiv.org/abs/1706.08500.

Binbin Huang, Haobin Duan, Yiqun Zhao, Zibo Zhao, Yi Ma, and Shenghua Gao. CUPID: Generative 3D reconstruction via joint object and pose modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12741–12752, 2026. URL https://arxiv.org/abs/2510.20776.

Xin Huang, Tengfei Wang, Ziwei Liu, and Qing Wang. Material anything: Generating materials for any 3D object via diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26556–26565, 2025a. URL https://openaccess.thecvf.com/content/CVPR2025/html/Huang_Material_Anything_Generating_Materials_for_Any_3D_Object_via_Diffusion_CVPR_2025_paper.html.

Zehuan Huang, Yuan-Chen Guo, Haoran Wang, Ran Yi, Lizhuang Ma, Yan-Pei Cao, and Lu Sheng. MV-Adapter: Multi-view consistent image generation made easy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 16377–16387, 2025b. URL https://arxiv.org/abs/2412.03632.

Sadeep Jayasumana, Srikumar Ramalingam, Andreas Veit, Daniel Glasner, Ayan Chakrabarti, and Sanjiv Kumar. Rethinking FID: Towards a better evaluation metric for image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9307–9315, 2024. URL https://arxiv.org/abs/2401.09603.

Bingxin Ke, Kevin Qu, Tianfu Wang, Nando Metzger, Shengyu Huang, Bo Li, Anton Obukhov, and Konrad Schindler. Marigold: Affordable adaptation of diffusion-based image generators for image analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. doi: 10.1109/TPAMI.2025.3591076. URL https://arxiv.org/abs/2505.09358.

Mukul Khanna, Yongsen Mao, Hanxiao Jiang, Sanjay Haresh, Brennan Shacklett, Dhruv Batra, Alexander Clegg, Eric Undersander, Angel X. Chang, and Manolis Savva. Habitat synthetic scenes dataset (HSSD-200): An analysis of 3D scene scale and realism tradeoffs for ObjectGoal navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 16384–16393, 2024.

Peter Kocsis, Vincent Sitzmann, and Matthias Nießner. Intrinsic image diffusion for indoor single-view material estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5198–5208, 2024. URL https://openaccess.thecvf.com/content/CVPR2024/html/Kocsis_Intrinsic_Image_Diffusion_for_Indoor_Single-view_Material_Estimation_CVPR_2024_paper.html.

Zeqiang Lai, Yunfei Zhao, Haolin Liu, Zibo Zhao, Qingxiang Lin, Huiwen Shi, Xianghui Yang, Mingxin Yang, Shuhui Yang, Yifei Feng, Sheng Zhang, Xin Huang, Di Luo, Fan Yang, Fang Yang, Lifu Wang, Sicong Liu, Yixuan Tang, Yulin Cai, Zebin He, Tian Liu, Yuhong Liu, Jie Jiang, Linus, Jingwei Huang, and Chunchao Guo. Hunyuan3D 2.5: Towards high-fidelity 3D assets generation with ultimate details. *arXiv preprint arXiv:2506.16504*, 2025. URL https://arxiv.org/abs/2506.16504.

Zeqiang Lai, Yunfei Zhao, Zibo Zhao, Xin Yang, Xin Huang, Jingwei Huang, Xiangyu Yue, and Chunchao Guo. NaTex: Seamless texture generation as latent color diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18619–18629, 2026. URL https://openaccess.thecvf.com/content/CVPR2026/html/Lai_NaTex_Seamless_Texture_Generation_as_Latent_Color_Diffusion_CVPR_2026_paper.html.

Dong-Yang Li, Wang Zhao, Yuxin Chen, Wenbo Hu, Meng-Hao Guo, Fang-Lue Zhang, Ying Shan, and Shi-Min Hu. Pixal3D: Pixel-aligned 3D generation from images. In *ACM SIGGRAPH 2026 Conference Papers*, 2026a. doi: 10.1145/3799902.3811175. URL https://arxiv.org/abs/2605.10922.

Sikuang Li, Chen Yang, Jiemin Fang, Taoran Yi, Jia Lu, Jiazhong Cen, Lingxi Xie, Wei Shen, and Qi Tian. WorldGrow: Generating infinite 3D world. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 6433–6441, 2026b. doi: 10.1609/aaai.v40i8.37571.

Weiyu Li, Xuanyang Zhang, Zheng Sun, Di Qi, Hao Li, Wei Cheng, Weiwei Cai, Shihao Wu, Jiarui Liu, Zihao Wang, Xiao Chen, Feipeng Tian, Jianxiong Pan, Zeming Li, Gang Yu, Xiangyu Zhang, Daxin Jiang, and Ping Tan. Step1X-3D: Towards high-fidelity and controllable generation of textured 3D assets. *arXiv preprint arXiv:2505.07747*, 2025a. URL https://arxiv.org/abs/2505.07747.

Yangguang Li, Zi-Xin Zou, Zexiang Liu, Dehu Wang, Yuan Liang, Zhipeng Yu, Xingchao Liu, Yuan-Chen Guo, Ding Liang, Wanli Ouyang, and Yan-Pei Cao. TripoSG: High-fidelity 3D shape synthesis using large-scale rectified flow models. *arXiv preprint arXiv:2502.06608*, 2025b. URL https://arxiv.org/abs/2502.06608.

Yixun Liang, Kunming Luo, Xiao Chen, Rui Chen, Hongyu Yan, Weiyu Li, Jiarui Liu, Fei-Peng Tian, and Ping Tan. UniTEX: Universal high fidelity generative texturing for 3D shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19917–19927, 2026. URL https://arxiv.org/abs/2505.23253.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Jialun Liu, Chenming Wu, Xinqi Liu, Xing Liu, Jinbo Wu, Haotian Peng, Chen Zhao, Haocheng Feng, Jingtuo Liu, and Errui Ding. TexOct: Generating textures of 3D models with octree-based diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4284–4293, 2024a. URL https://openaccess.thecvf.com/content/CVPR2024/html/Liu_TexOct_Generating_Textures_of_3D_Models_with_Octree-based_Diffusion_CVPR_2024_paper.html.

Yuxin Liu, Minshan Xie, Hanyuan Liu, and Tien-Tsin Wong. Text-guided texturing by synchronized multi-view diffusion. In *SIGGRAPH Asia 2024 Conference Papers*, 2024b. doi: 10.1145/3680528.3687621. URL https://doi.org/10.1145/3680528.3687621.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.

Sai Raj Kishore Perla, Yizhi Wang, Ali Mahdavi-Amiri, and Hao Zhang. EASI-Tex: Edge-aware mesh texturing from single image. *ACM Transactions on Graphics*, 43(4):40:1–40:11, 2024. doi: 10.1145/3658222.

Qwen Team. Qwen3.8-27B, 2026. URL https://huggingface.co/Qwen/Qwen3.8-27B. Model card.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pp. 8748–8763, 2021. URL https://proceedings.mlr.press/v139/radford21a.html.

Elad Richardson, Gal Metzer, Yuval Alaluf, Raja Giryes, and Daniel Cohen-Or. TEXTure: Text-guided texturing of 3D shapes. In *ACM SIGGRAPH 2023 Conference Proceedings*, pp. 1–11, 2023. doi: 10.1145/3588432.3591503. URL https://doi.org/10.1145/3588432.3591503.

Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2304–2314, 2019. URL https://openaccess.thecvf.com/content_ICCV_2019/html/Saito_PIFu_Pixel-Aligned_Implicit_Function_for_High-Resolution_Clothed_Human_Digitization_ICCV_2019_paper.html.

Mingqi Shao, Feng Xiong, Zhaoxu Sun, and Mu Xu. MVPainter: Accurate and detailed 3D texture generation via multi-view diffusion with geometric control. *arXiv preprint arXiv:2505.12635*, 2025. URL https://arxiv.org/abs/2505.12635.

Ruizhi Shao, Youxin Pang, Zerong Zheng, Jingxiang Sun, and Yebin Liu. 360-degree human video generation with 4D diffusion transformer. *ACM Transactions on Graphics (TOG)*, 43(6), 2024. doi: 10.1145/3687980.

Kaiyue Shen, Chen Guo, Manuel Kaufmann, Juan Jose Zarate, Julien Valentin, Jie Song, and Otmar Hilliges. X-Avatar: Expressive human avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 16911–16921, 2023.

Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. DINOv3. *arXiv preprint arXiv:2508.10104*, 2025. URL https://arxiv.org/abs/2508.10104.

Sketchfab. Sketchfab, n.d. URL https://sketchfab.com.

Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. RoFormer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. doi: 10.1016/j.neucom.2023.127063.

Super Intelligence Team. FireRed-Image-Edit-1.0 technical report. *arXiv preprint arXiv:2602.13344*, 2026. URL https://arxiv.org/abs/2602.13344.

Tencent Hunyuan3D Team. Hunyuan3D 2.1: From images to high-fidelity 3D assets with production-ready PBR material, 2025a. URL https://arxiv.org/abs/2506.15442.

Tencent Hunyuan3D Team. Hunyuan3D 2.0: Scaling diffusion models for high resolution textured 3D assets generation. *arXiv preprint arXiv:2501.12202*, 2025b. URL https://arxiv.org/abs/2501.12202.

Wenbo Wang, Hsuan-I Ho, Chen Guo, Boxiang Rong, Artur Grigorev, Jie Song, Juan Jose Zarate, and Otmar Hilliges. 4D-DRESS: A 4D dataset of real-world human clothing with semantic annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 550–560, 2024.

Zhou Wang, Alan Conrad Bovik, Hamid Rahim Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13 (4):600–612, 2004. URL https://www.cns.nyu.edu/pub/eero/wang03-reprint.pdf.

Guanjun Wu, Jiemin Fang, Chen Yang, Sikuang Li, Taoran Yi, Jia Lu, Zanwei Zhou, Jiazhong Cen, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Xinggang Wang, and Qi Tian. UniLat3D: Geometry-appearance unified latents for single-stage 3D generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (Findings)*, 2026a. URL https://arxiv.org/abs/2509.25079.

Xiuchao Wu, Pengfei Zhu, Jiangjing Lyu, Xinguo Liu, Jie Guo, Yanwen Guo, Weiwei Xu, and Chengfei Lyu. MatMart: Material reconstruction of 3D objects via diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2336–2345, 2026b. URL https://openaccess.thecvf.com/content/CVPR2026/papers/Wu_MatMart_Material_Reconstruction_of_3D_Objects_via_Diffusion_CVPR_2026_paper.pdf.

Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3D latents for scalable and versatile 3D generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 21469–21480, 2025. URL https://openaccess.thecvf.com/content/CVPR2025/html/Xiang_Structured_3D_Latents_for_Scalable_and_Versatile_3D_Generation_CVPR_2025_paper.html.

Jianfeng Xiang, Xiaoxue Chen, Sicheng Xu, Ruicheng Wang, Zelong Lv, Yu Deng, Hongyuan Zhu, Yue Dong, Hao Zhao, Nicholas Jing Yuan, and Jiaolong Yang. Native and compact structured latents for 3D generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14419–14429, 2026. URL https://openaccess.thecvf.com/content/CVPR2026/html/Xiang_Native_and_Compact_Structured_Latents_for_3D_Generation_CVPR_2026_paper.html.

Bojun Xiong, Jialun Liu, Jiakui Hu, Chenming Wu, Jinbo Wu, Xing Liu, Chen Zhao, Errui Ding, and Zhouhui Lian. TexGaussian: Generating high-quality PBR material via octree-based 3D gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 551–561, 2025. URL https://openaccess.thecvf.com/content/CVPR2025/html/Xiong_TexGaussian_Generating_High-quality_PBR_Material_via_Octree-based_3D_Gaussian_Splatting_CVPR_2025_paper.html.

Kyeongmin Yeo, Yunhong Min, Jaihoon Kim, and Minhyuk Sung. MatLat: Material latent space for PBR texture generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4602–4612, 2026. URL https://openaccess.thecvf.com/content/CVPR2026/html/Yeo_MatLat_Material_Latent_Space_for_PBR_Texture_Generation_CVPR_2026_paper.html.

Kim Youwang, Tae-Hyun Oh, and Gerard Pons-Moll. Paint-it: Text-to-texture synthesis via deep convolutional texture map optimization and physically-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4347–4356, 2024. URL https://openaccess.thecvf.com/content/CVPR2024/html/Youwang_Paint-it_Text-to-Texture_Synthesis_via_Deep_Convolutional_Texture_Map_Optimization_and_CVPR_2024_paper.html.

Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4578–4587, 2021a. URL https://openaccess.thecvf.com/content/CVPR2021/html/Yu_pixelNeRF_Neural_Radiance_Fields_From_One_or_Few_Images_CVPR_2021_paper.html.

Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. Function4D: Real-time human volumetric capture from very sparse consumer RGBD sensors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5746–5756, 2021b.

Xin Yu, Peng Dai, Wenbo Li, Lan Ma, Zhengzhe Liu, and Xiaojuan Qi. Texture generation on 3D meshes with point-UV diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4206–4216, 2023. URL https://arxiv.org/abs/2308.10490.

Xin Yu, Ze Yuan, Yuan-Chen Guo, Ying-Tian Liu, Jianhui Liu, Yangguang Li, Yan-Pei Cao, Ding Liang, and Xiaojuan Qi. TEXGen: a generative diffusion model for mesh textures. *ACM Transactions on Graphics*, 43(6):213:1–213:14, 2024. doi: 10.1145/3687909. URL https://arxiv.org/abs/2411.14740.

Xianfang Zeng, Xin Chen, Zhongqi Qi, Wen Liu, Zibo Zhao, Zhibin Wang, Bin Fu, Yong Liu, and Gang Yu. Paint3D: Paint anything 3D with lighting-less texture diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4252–4262, 2024a. URL https://openaccess.thecvf.com/content/CVPR2024/html/Zeng_Paint3D_Paint_Anything_3D_with_Lighting-Less_Texture_Diffusion_Models_CVPR_2024_paper.html.

Yifei Zeng, Yajie Bao, Jiachen Qian, Shuang Wu, Youtian Lin, Hao Zhu, Buyu Li, Feihu Zhang, Xun Cao, and Yao Yao. TEXTRIX: Latent attribute grid for native texture generation and beyond. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 27104–27113, 2026. URL https://arxiv.org/abs/2512.02993.

Zheng Zeng, Valentin Deschaintre, Iliyan Georgiev, Yannick Hold-Geoffroy, Yiwei Hu, Fujun Luan, Ling-Qi Yan, and Miloš Hašan. RGB↔X: Image decomposition and synthesis using material- and lighting-aware diffusion models. In *ACM SIGGRAPH 2024 Conference Papers*, 2024b. doi: 10.1145/3641519.3657445. URL https://arxiv.org/abs/2405.00666.

Longwen Zhang, Ziyu Wang, Qixuan Zhang, Qiwei Qiu, Anqi Pang, Haoran Jiang, Wei Yang, Lan Xu, and Jingyi Yu. CLAY: A controllable large-scale generative model for creating high-quality 3D assets. In *ACM SIGGRAPH 2024 Conference Papers*, 2024a. doi: 10.1145/3658146. URL https://doi.org/10.1145/3658146.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. URL https://arxiv.org/abs/1801.03924.

Yibo Zhang, Li Zhang, Rui Ma, and Nan Cao. Texverse: A universe of 3D objects with high-resolution textures. *arXiv preprint arXiv:2508.10868*, 2025a.

Yuqing Zhang, Yuan Liu, Zhiyu Xie, Lei Yang, Zhongyuan Liu, Mengzhou Yang, Runze Zhang, Qilong Kou, Cheng Lin, Wenping Wang, and Xiaogang Jin. DreamMat: High-quality PBR material generation with geometry- and light-aware diffusion models. *ACM Transactions on Graphics*, 43 (4):39:1–39:18, 2024b. doi: 10.1145/3658170.

Yuqing Zhang, Hao Xu, Yiqian Wu, Sirui Chen, Sirui Lin, Xiang Li, Xifeng Gao, and Xiaogang Jin. AlignTex: Pixel-precise texture generation from multi-view artwork. *ACM Transactions on Graphics*, 44(4):161:1–161:12, 2025b. doi: 10.1145/3731158.

Yuqing Zhang, Yan-Pei Cao, Hao Xu, Yiqian Wu, Sirui Lin, Yuqing Wang, Ding Liang, Yuan-Chen Guo, and Xiaogang Jin. PixTex: Consistent 3D texturing via pixel-space multi-view diffusion. In *ACM SIGGRAPH Conference Papers*, 2026. doi: 10.1145/3799902.3811120.

Zoey. Multiview upgrade: Multi-view geometry and texture in Meshy. Meshy Blog, 2026. URL https://www.meshy.ai/blog/multiview-upgrade. Published August 7, 2026. Accessed September 26, 2026.

SUPPLEMENTARY MATERIAL

This supplement provides method and training details, evaluation protocols, additional results, and extended discussion. The guide below links to each appendix and selected subsections; page numbers indicate where each appendix begins.

- **Appendix A: Method and Inference Details** (p. 18). Geometry and camera projection, visible-surface observations, and boundary filtering; the interleaved flow transformer, including token dimensions and its read-only update schedule; and the inference procedure and sampling settings.
- **Appendix B: Training Data and Optimization** (p. 20). Training sources and asset curation, the 2D material frontend, and 3D generator training, including observation construction, the full-surface flow-matching objective, optimization, and checkpoint selection.
- **Appendix C: Evaluation Protocols** (p. 22). Evaluation datasets, rendering cameras, alignment, and metrics, baseline configurations, and commercial comparisons. The complete VLM scoring prompts are provided for input alignment, PBR plausibility, and visual quality.
- **Appendix D: Additional Quantitative Results** (p. 26). Reference-view versus novel-view fidelity, 2D material prediction accuracy, and runtime and memory measurements, including the effects of observation KV caching.
- **Appendix E: Ablation and Observation Analyses** (p. 29). Ablation configurations, partial versus complete material encoding, observed-support encoding versus zero filling, and early learning with native PBR latents.
- **Appendix F: Additional Qualitative Comparisons** (p. 33). Extended visual comparisons for given-geometry texturing, end-to-end image-to-3D generation, and commercial systems, complementing the examples in the main paper.
- **Appendix G: Editing and Limitations** (p. 36). Synthetic back-view conditions and image-based texture editing, followed by limitations and failure cases involving geometry alignment, material estimation, and unseen-surface completion.
- **Appendix H: Additional Related Work** (p. 39). An expanded discussion of multi-view texture synthesis, reference-guided appearance preservation, PBR material generation, and native 3D representations and conditioning.
- **Appendix I: Notation Reference** (p. 40). A consolidated symbol table with dimensions and definitions for the physical material channels, surface and latent coordinates, observation metadata, and generator states.

A METHOD AND INFERENCE DETAILS

A.1 GEOMETRY AND SURFACE OBSERVATIONS

Geometry foundation and camera projection. Our framework operates on an underlying mesh $\mathcal{M}$ and camera projection π, provided as ground-truth inputs or estimated from the reference image. For end-to-end inference, we follow the default MoGe-2-based camera-estimation pipeline of Pixal3D*, whose geometry module generates the mesh. For the 3D position $\mathbf{x}_s$ of latent coordinate s, the model samples image features through camera projection:

$$f_s = \mathrm{Sample}(\Phi(I), \pi(\mathbf{x}_s)) \,, \tag{9}$$

where $\Phi(I)$ denotes a 2D feature map. In the geometry refinement stages, Pixal3D* projects both the original and spatially upsampled image features and concatenates the sampled features. Combined with global image conditioning, these features align the predicted mesh $\mathcal{M}$ and camera projection π with the reference view, serving as the geometric input for our material generation.

Surface observations and latent coordinates. The full surface defines the active latent-coordinate set $\mathcal{S}$. Valid visible surface voxels form the partial PBR input $\mathcal{P}_o$, which is encoded before any latent-level conditioning is constructed. The resulting coordinates $\mathcal{S}_o$ identify cells with observations, but a cell in $\mathcal{S}_o$ can also contain unseen surface voxels; $\mathcal{U} = \mathcal{S} \setminus \mathcal{S}_o$ therefore denotes cells without any observation, not all unseen surface. Coverage ρ_s is the observed fraction of the surface voxels actually present in a cell, and repeated image samples of one voxel do not increase it. Boundary distance d_s is defined on the six-neighbor graph of active latent coordinates, with -1 when no unobserved cell is reachable.

Inference-time boundary filtering. At inference, we additionally apply foreground-mask erosion and local depth-discontinuity filtering to reduce texture contamination from unreliable correspondences. Foreground erosion removes mixed pixels along the silhouette, and depth-discontinuity filtering rejects correspondences whose neighboring visible depths differ by more than a threshold, where a small projection error would move the observation onto a different surface. All material channels share the retained observation support, from which local material encoding, coverage, and boundary distance are computed.

A.2 INTERLEAVED FLOW TRANSFORMER

3D generator architecture. Material and geometry latents have 32 channels each on a sparse latent grid of resolution 64^3. The generator has 60 Transformer blocks with hidden dimension 1536 and 12 attention heads, 30 of each update type; its self-attention modules are initialized from TRELLIS.2 and the remaining components randomly.

Two states at an observed coordinate. At every $s \in \mathcal{S}_o$, the read-only observation token O_s represents a partial material observation, while the generation token X_s predicts the complete material of the cell. Geometry and material latents use fixed per-channel normalization; observation latents are mapped to the material-latent scale before the shared input projection. Projected spatial metadata and the learned role bias $b_o \in \mathbb{R}^{1536}$, initialized to zero, are added only to observation tokens. Table 5 summarizes the dimensions.

Table 5: Representation dimensions. $N = |\mathcal{S}|$ and $M = |\mathcal{S}_o|$ vary across objects.

Quantity	Dimensions	Meaning
Geometry latents G	$N \times 32$	Geometry conditions on the full surface
Noisy material latents Z_t	$N \times 32$	Full-surface material state at time t
Native observation latents $\{z_s^o\}_{s\in\mathcal{S}_o}$	$M \times 32$	Encoded partial PBR materials
Observation metadata	$M \times 2$	Coverage and boundary distance
Generation stream X	$N \times 1536$	Hidden generation states
Observation stream O	$M \times 1536$	Read-only observation states
Output velocity	$N \times 32$	Updates to all material latents

Attention and update schedule. Figure 5 illustrates the two streams and their alternating update schedule. Each of the 60 Transformer blocks reads keys and values from both streams. Even-indexed blocks use queries from $X_{\mathcal{U}}$, odd-indexed blocks from all of X; generation states in $\mathcal{U}$ are therefore updated in 60 blocks and those in $\mathcal{S}_o$ in 30. Observation tokens pass through each block's normalization and adaptive modulation with the clean-time indicator $t = 0$, produce no queries, and receive no residual or feed-forward updates. Queries from X, and keys from both O and X, carry rotary positional encoding (Su et al., 2024) extended to 3D latent coordinates; values are not rotated. These counts describe hidden-state updates within one forward pass and do not restrict which material latents evolve during sampling: the output velocity and Euler updates cover all of $\mathcal{S}$.

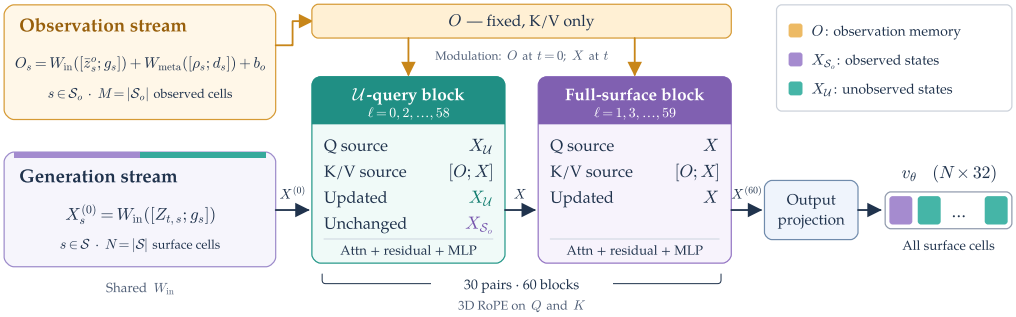

Figure 5: **Interleaved flow transformer.** One forward pass through 60 distinct blocks arranged in 30 alternating pairs. The fixed observation stream O (orange) supplies keys and values without queries or state updates. The generation stream X spans the full active latent-cell set $\mathcal{S}$, with states at observed cells $\mathcal{S}_o$ (purple) and cells without any observation $\mathcal{U} = \mathcal{S} \setminus \mathcal{S}_o$ (teal); the subscript o denotes observed. An observed cell may still contain unseen surface voxels. Even-indexed blocks query and update $X_{\mathcal{U}}$; odd-indexed blocks query and update all of X. Both block types read keys and values from all of O and X, and the output projection predicts velocities over all $N = |\mathcal{S}|$ cells. Source labels denote tokens before normalization, time modulation, and attention projections. At inference, 25 Euler sampling steps each invoke a full transformer forward pass.

A.3 Inference

Single-view inference procedure. Given dataset-provided or predicted geometry and cameras, we encode the mesh into G and $\mathcal{S}$ and the filtered observations into anchors O (Sec. 3.3). Starting from $Z_1 \sim \mathcal{N}(0, \mathbf{I})$ over $\mathcal{S}$, we integrate the learned velocity field with Euler steps, then de-normalize and decode the final latents with D_{mat} into full-surface PBR attributes for export. Optional KV caching reuses the observation stream's per-layer keys and values across sampling steps; Table 9 reports its runtime and memory effects.

Inference uses 25 Euler steps without classifier-free guidance on the time grid $t_j = 3u_j/(1 + 2u_j)$, $u_j = 1 - j/25$, $j = 0, \ldots, 25$.

B Training Data and Optimization

B.1 Training Data and Curation

The generator is trained on approximately 570k unique assets from TexVerse (Zhang et al., 2025a), Objaverse-XL (Deitke et al., 2023), ABO (Collins et al., 2022), HSSD (Khanna et al., 2024), and open synthetic 3D assets, together with human scans from 4D-Dress (Wang et al., 2024), Human4DiT (Shao et al., 2024), THuman2.1 (Yu et al., 2021b), and X-Humans (Shen et al., 2023). These 570k assets are a subset of a scored pool of 1,048,073 objects from public collections.

Figure 6 illustrates the VLM-based data filtering pipeline. For each asset, 6 RGB views are scored in one forward pass of Qwen3.8-27B (Qwen Team, 2026). The rater classifies structural defects, extraneous ground planes, untextured geometry, broken indoor scans, and outdoor solid terrain, and predicts geometric density s_{geo}, texture richness s_{tex}, and an aesthetic score $s_{\text{aes}} \in [0, 10]$. Extraneous ground planes and broken scans are discarded regardless of score; untextured objects are capped at $s \leq 2.99$. The remaining objects receive a composite quality score that combines geometric density, texture richness, aesthetics, and a class prior.

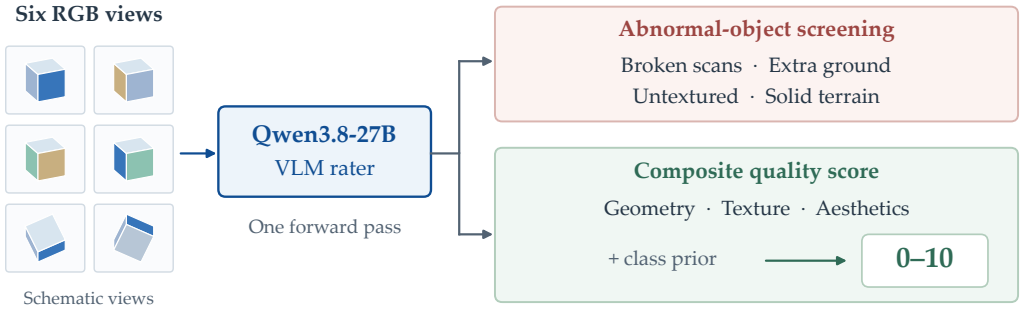

Figure 6: Six-view VLM rater used to screen training assets. The input tiles schematically represent six RGB views scored jointly by Qwen3.8-27B (Qwen Team, 2026) in one forward pass. The upper branch flags broken scans, untextured geometry, extraneous ground planes, and outdoor solid terrain. The lower branch combines geometry density, texture complexity, aesthetics, and a class prior into a quality score in $[0, 10]$.

The admitted corpus is the unique union of four disjoint slices (Table 6). An aesthetic track follows the TRELLIS.2 thumbnail cutoff (Xiang et al., 2026) and retains about 430k assets with an aesthetic score of at least 4.5. A VLM track recovers a disjoint set of about 90k TexVerse and Objaverse-XL Sketchfab assets whose composite score is at least 6. The remaining about 50k assets are coordinate-repaired TexVerse meshes (about 22k) and the four human collections (about 25k); neither slice uses the aesthetic threshold. The four slices do not overlap; their unique union is the reported 570k set.

Table 6: Composition of the admitted training set. The four slices are disjoint; the total is a unique-asset count, not the sum of rounded figures. This is not the scored pool of 1,048,073 assets.

Slice	Assets	Admission rule
Aesthetic track	∼430k	aesthetic score ≥ 4.5
VLM track	∼90k	composite $s \geq 6$; TexVerse and Objaverse-XL Sketchfab
Coord.-repaired TexVerse	∼22k	no aesthetic cutoff
Human collections	∼25k	no aesthetic cutoff
Unique total	∼570k	disjoint union

We use every source under its published license and only for non-commercial academic research. TexVerse and Complex-Tex retain per-asset Sketchfab terms. The Objaverse-XL collection is released under ODC-By 1.0, with object-level licenses unchanged. ABO is released under CC BY 4.0. HSSD and the Digital Twin Catalog under CC BY-NC 4.0. The human scans follow their institutional non-commercial agreements.

B.2 2D Material Frontend Training

We initialize the frontend from FireRed-Image-Edit-1.1 adapted to individual PBR and normal prediction tasks, and jointly fine-tune it on aligned RGB images and ground-truth material and normal maps rendered from Objaverse assets. Both the task-specific adaptation and joint fine-tuning use these rendered data. Modality-axis attention is inserted at Transformer blocks 29, 39, 49, and 59, with rank 256 and four attention heads. We optimize these adapters and backbone blocks 29–49 using a latent flow-matching objective and AdamW (Loshchilov & Hutter, 2019), with learning rates of 2×10^{-4} and 5×10^{-7}, respectively. The VAE and text encoder remain frozen.

By folding the batch dimension, the model processes albedo, metallic-roughness, and an auxiliary normal stream in parallel, inserting lightweight cross-attention adapters at specific transformer layers to exchange features across modality axes at identical spatial coordinates. We train F_{2D} with a flow-matching objective anchored at the fixed maximum-noise endpoint $\sigma = 1.0$ using an outlier-masked MSE loss. Elementwise latent-flow residuals with absolute magnitude greater than 50 contribute zero loss but remain in the averaging denominator. Losses are averaged equally across the albedo, metallic-roughness, and normal streams. At inference, we use a single prediction at this endpoint to estimate the clean material latents, which the image VAE decodes into 1024×1024 PBR maps.

B.3 3D Generator Training

Preprocessing for training. We build training observations from ground-truth PBR materials of public 3D assets using the visible-surface observation encoding pipeline. Foreground validity and first-hit visibility determine the observed surface voxels supplied to the sparse material encoder. We follow Pixal3D (Li et al., 2026a) for asset preprocessing and reference-view sampling. During training, we drop the entire observation condition with probability 0.03 while retaining the geometry condition.

Full-surface training objective. The generation stream X synthesizes material latents across the entire surface $\mathcal{S}$. During training, observations are constructed from ground-truth materials of public 3D assets using the visible-surface observation encoding pipeline, while complete material latents serve as the target. With Z_{gt} the normalized ground-truth latents and $\epsilon \sim \mathcal{N}(0, \mathbf{I})$,

$$Z_t = (1-t)Z_{\mathrm{gt}} + [\sigma_{\min} + (1-\sigma_{\min})t]\,\epsilon, \qquad v^* = (1-\sigma_{\min})\epsilon - Z_{\mathrm{gt}}, \tag{10}$$

where $\sigma_{\min}$ is a small residual noise scale, $t = 1$ corresponds to pure noise, and $t = 0$ to material latents with residual noise $\sigma_{\min}\epsilon$. The loss is

$$\mathcal{L}_{\mathrm{FM}} = \mathbb{E}[\mathrm{MSE}(v_\theta(Z_t, t; O, G), v^*)]\,, \tag{11}$$

averaged over all valid tokens and channels, including observed coordinates. The observation representation remains fixed and is not noised along the flow path; it serves as conditioning input to the generator. The generation stream is supervised against complete-material targets over the entire surface, learning to reconstruct visible appearance from the fixed observations and to complete unobserved regions.

Optimization. The material encoder and decoder are frozen. Generator self-attention is initialized from TRELLIS.2 and the remaining components randomly. We use the flow-matching objective in Eq. equation 11. We optimize the 3D generator for 30,000 steps with a global batch size of 576 using AdamW, with $\beta_1 = 0.9$, $\beta_2 = 0.95$, weight decay 0.01, and $\epsilon = 10^{-8}$. The initial learning-rate schedule uses 150 steps of linear warmup to 10^{-4}, followed by cosine decay with a 20,000-step horizon and a minimum learning rate of 10^{-5}. At step 15,000, we extend training with a cosine continuation from approximately 2.337×10^{-5} to 10^{-5} at step 30,000, without additional warmup. We use BF16 automatic mixed precision and adaptive gradient clipping with a maximum norm of 1.0 and a 95th-percentile threshold. The final training stage uses 64 GPUs and FlashAttention-3 acceleration. Evaluation uses the raw generator weights at step 30,000.

Training-time sampling. We use $\sigma_{\min} = 10^{-5}$ and sample training times as $t = \mathrm{sigmoid}(\xi)$ with $\xi \sim \mathcal{N}(1, 1)$. Observation dropout is applied as described in preprocessing above.

C Evaluation Protocols

C.1 Evaluation Datasets

DTC-Tex comprises 100 objects from the Digital Twin Catalog (Dong et al., 2025). We score the 2,399 catalog assets for surface-texture complexity from six views with Qwen3.8-27B, discard near-duplicates and broken meshes, and retain the 100 highest-scoring objects. Each object is evaluated from a front-view reference, together with its ground-truth mesh, PBR maps, and a 16-view camera rig. Typical cases are household products with printed labels or decorative motifs, such as canned goods, decorated ceramics, and books.

Complex-Tex complements this set with 110 Sketchfab (Sketchfab, n.d.) meshes that carry denser inscriptions and finer patterns on more varied shapes. This subset includes branded packaging and vehicles covered by logos and fine print, as well as buildings and characters with more irregular shapes. A method receives only the designated reference image, or that image together with the mesh when geometry is given. Figure 7 shows selected assets from DTC-Tex and Complex-Tex.



Figure 7: **Examples from DTC-Tex and Complex-Tex.** Ten selected assets from DTC-Tex (top) and Complex-Tex (bottom), shown in their designated reference views. The examples include dense text, fine patterns, and reflective surfaces.

Both DTC-Tex and Complex-Tex underwent strict duplicate checks during data collection to ensure no overlap with the training assets. We will release both evaluation benchmarks.

C.2 Rendering and Metrics

Protocol. Both settings evaluate the PBR asset each method exports, rendered under a shared protocol rather than compared in UV space, because latent decoding or remeshing changes the face layout. Predicted and ground-truth assets are rendered with the same cameras and forward lighting at 1024×1024 and downsampled to 512×512 with Lanczos interpolation; methods whose output lives in another coordinate frame are brought into the evaluation frame by their own transform or, where marked †, by a similarity transform to the ground-truth mesh. Per-view metrics are PSNR, SSIM (Wang et al., 2004), and VGG-LPIPS (Zhang et al., 2018) over the full image with a black background, so they penalize errors in silhouette as well as appearance, and CLIP-I, the cosine similarity of CLIP ViT-L/14 features (Radford et al., 2021) between corresponding views. We additionally report PSNR_{fg} on both shaded RGB and albedo, the same PSNR restricted to the ground-truth object mask, so that large black background regions do not dominate the score. Distribution metrics compare multi-view image sets within each dataset: $\text{FD}_{\text{DINOv3}}$ uses 1024-dimensional DINOv3 ViT-L/16 CLS features (Siméoni et al., 2025). In end-to-end evaluation, FID (Heusel et al., 2017) uses Inception-v3 features, while FD_{CLIP} and CMMD (Jayasumana et al., 2024) provide additional distribution measures. For methods that output PBR, we report mean absolute error (MAE)

of metallic and roughness over the intersection of the predicted and ground-truth object masks. Per-view metrics are averaged over views per object, then over valid objects. We report results separately on DTC-Tex and Complex-Tex. Texture generation with given geometry uses the ground-truth mesh and camera; end-to-end generation predicts geometry and applies similarity alignment to the ground truth only for scoring.

Rendering cameras and alignment. For given-geometry evaluation, we replay the 16 camera poses and front-light settings stored in each ground-truth asset's `transforms.json`; the designated input camera defines the reference view, and the other 15 cameras define novel views. Ground-truth and predicted assets use identical cameras and lighting, rendered with Blender 3.0.1 and Filmic color management. Where geometric alignment is required, we use iterative closest point (ICP) and silhouette-based refinement to estimate a similarity transform to the ground-truth geometry before rendering. This alignment is used only for evaluation.

DINOv3 Fréchet distance for given-geometry texturing. For Table 1, we extract the 1024-dimensional final-layer CLS token of DINOv3 ViT-L/16 without additional ℓ_2 normalization. We alpha-composite each shaded rendering onto black, resize its shorter side to 256 pixels using bilinear interpolation, center-crop to 224×224, and apply ImageNet mean and standard-deviation normalization. These features are computed in FP32 with TF32 disabled using Transformers 4.57.3. We compute the Fréchet distance between the feature means and unbiased covariances of the generated and corresponding ground-truth image sets, using float64 statistics. This is one set-level score per method and dataset, rather than an average of object-level distances.

C.3 Open-Source Baselines

Given-geometry comparison. Each object provides a ground-truth mesh and a forward-lit reference image from a specified view; every method textures the ground-truth mesh from this image through its own interface, and we render 16 views comprising the reference view and 15 novel views. We compare against the texturing stages of the native 3D generators TRELLIS.2 and Pixal3D* and of the pipelines Hunyuan3D 2.1 and Step1X-3D (Xiang et al., 2026; Li et al., 2026a; Tencent Hunyuan3D Team, 2025a; Li et al., 2025a), and against the dedicated texturing methods LumiTex, MVPainter, MV-Adapter, TEXGen, UniTEX, and Ink3D (Bao et al., 2026; Shao et al., 2025; Huang et al., 2025b; Yu et al., 2024; Liang et al., 2026; Han et al., 2026). TRELLIS.2, Pixal3D*, Hunyuan3D 2.1, LumiTex, and MVPainter output PBR materials and are evaluated on shaded renderings, albedo, metallic, and roughness; the remaining methods output RGB textures and are evaluated on shaded renderings only, with unavailable channels marked "–". We use the official implementations and released checkpoints for all baselines, providing the reference camera to every method that supports camera conditioning.

End-to-end comparison. When geometry is also predicted, reference fidelity depends on geometry reconstruction, surface correspondence, and material generation together. This is the setting the observation interface is most exposed to: the anchors now land on a predicted mesh through a predicted camera, so any geometry or correspondence error propagates into the materials. Each method takes the reference image alone. Pixal3D* and SUPRA-Tex share the same predicted geometry and apply their respective appearance-generation pipelines; the other baselines use their own geometry-generation pipelines. SUPRA-Tex textures this predicted mesh with the pipeline of Appendix A.3. We compare with TRELLIS.2, Pixal3D*, Hunyuan3D 2.1, Step1X-3D, UniLat3D, and CUPID (Xiang et al., 2026; Li et al., 2026a; Tencent Hunyuan3D Team, 2025a; Li et al., 2025a; Wu et al., 2026a; Huang et al., 2026), covering native latent generation, staged geometry-then-texture pipelines, and joint object–pose modeling. Since output material attributes differ across these systems, we compare complete assets through rendered appearance only. For scoring, outputs are brought into a common coordinate frame using their coordinate transforms or, for methods marked †, similarity alignment to the ground truth; this alignment plays no role in generation.

Reconstruction metrics take the reference view as anchor and compare four rendered views against the ground truth: the reference direction, its two sides, and the back. The side and back views thus test the completed surface, not only the reference view. CLIP-I compares the reference-view rendering with the input image. $\mathrm{FD}_{\mathrm{DINOv3}}$, $\mathrm{FD}_{\mathrm{CLIP}}$, FID, and CMMD compare the set of 16 rendered views of generated assets with the corresponding ground-truth set within each dataset. Reconstruction and

CLIP-I measure fidelity to a specific object; the distribution metrics measure the visual quality of the generated population.

Baselines. We use the authors' released checkpoints and implementations for the open-source baselines, following their inference recipes with method-specific sampling, resolution, and texture-export settings. Integration includes input-format conversion and repairs to mesh preprocessing, UV construction, and export. For given-geometry texturing, each method receives the designated reference image and GT mesh with its original materials removed, using the method's texturing interface; the reference camera is supplied where the interface supports it. No GT material maps or additional GT images are provided. For image-to-3D generation, each method receives the same single reference image. Pixal3D* and SUPRA-Tex share predicted geometry, as specified in Sec. 4.2.

Exported assets are converted from their method-specific coordinate frames for evaluation; similarity alignment is applied where indicated by † in the tables. Rendering and scoring follow the shared evaluation protocol. TEXTRIX (Zeng et al., 2026) was not included in the quantitative comparison because its official implementation and pretrained weights were unavailable when our experiments were conducted. The same availability constraint applied to NaTex, LaFiTe, and MatMart (Lai et al., 2026; Chen et al., 2026; Wu et al., 2026b). We discuss these methods in Appendix H, but do not mix their reported scores with our results because the datasets and evaluation protocols differ.

C.4 Commercial Comparisons

Commercial systems do not disclose their training data, so overlap with our ground-truth benchmarks cannot be ruled out. We therefore evaluate SUPRA-Tex without additional fine-tuning on a separate set of 45 images, comprising real photographs sourced online or contributed by users and synthetic images generated by image generation models such as GPT-Image-2. We compare end-to-end with the five commercial systems in Table 3: Hunyuan3D (v3.1), Hyper3D (Rodin-2.5 Extremely High), Tripo (H3.1), Meshy (v7.1), and Hi3D (v3.0), using their most advanced configurations and accessed after September 24, 2026. Hyper3D is listed as Rodin in the table.

We use two-view conditioning only for commercial comparisons, combining the reference image with a back view synthesized from it through the interface in Appendix G.1. We submit only the original reference image to each commercial system. Auxiliary-view synthesis is also supported by commercial workflows such as Meshy's multi-view mode, which can generate missing views from a single image (Zoey, 2026).

The 3D generator receives projected PBR observations rather than photographs, so generalization to real inputs depends on the fine-tuned 2D material estimator and accurate placement of the surface observations. In the absence of ground-truth geometry and cameras, paired reconstruction metrics are unavailable. We instead assess three criteria: alignment with the input image, including text, logos, and color layout; PBR plausibility; and overall visual quality.

We collected 512 valid pairwise evaluations, each comprising one participant's judgments on all three criteria for SUPRA-Tex versus one commercial system on a given input case. Each participant evaluated at most one model pair per case. For each of the five commercial systems and each criterion, Table 3 reports SUPRA-Tex's wins divided by the total number of non-tied votes, pooled across evaluations.

Qwen3.8-27B assigns each result a score from 0 to 10 using the rendered outputs, with criterion-specific inputs and prompts specified in Appendix C.4.1. Table 3 reports both human preferences and VLM ratings. Figure 4 illustrates how SUPRA-Tex preserves reference details and completes unseen surfaces consistently with the visible design while using only public training data.

C.4.1 VLM RATING PROMPTS

The VLM columns of Table 3 are one-shot ratings from Qwen3.8-27B, one call per criterion, on a 0–10 scale with half-point steps. Alignment compares the reference with the same-camera prediction. PBR and visual quality see only the prediction.

Alignment. The reference image is first. The prediction under the same light and camera is second.

```
Score alignment only.
If the reference has text, a face, or a complex pattern, the prediction
↪  must keep the same content in the same place. Skip a cue that is
↪  absent.
Then check identity, part layout, and main color.
Do not score sharpness, relighting, or mesh defects.
Illegible or wrong text, face, or pattern scores low. A clear match
↪  scores 9--10.
```

PBR. The VLM PBR criterion assesses relighting appearance, focusing on highlight motion and the appearance of illuminated regions. It uses relit views of the prediction without a reference image.

```
Score PBR only. No reference.
A moving highlight with a deep, saturated front light scores 9.
A pale but still recognizable front light scores 8.5.
A clipped white region or a highlight that does not move scores 7.
Do not score sharpness or the mesh.
```

Visual quality. Orbit views of the prediction are shown. No reference image is attached.

```
Score texture quality only. No reference.
Sharp weave, fur, leather, engraving, or print in every view scores 9.
The same details with mild softness score 8.5.
A uniform plastic or waxy skin scores 8, even when the face and clothes
↪  remain recognizable.
One view that loses the pattern scores 8. Two or more such views score 6.
Do not score similarity to a reference, relighting, or mesh holes.
```

Table 7: Reference-view and novel-view fidelity using the same per-method, per-metric valid objects as Table 1. PSNR_{fg} uses the GT object mask.

	Reference view (1)				Novel views (15)			
	Shaded		Albedo		Shaded		Albedo	
Method	PSNR_{fg} ↑	LPIPS↓	PSNR_{fg} ↑	LPIPS↓	PSNR_{fg} ↑	LPIPS↓	PSNR_{fg} ↑	LPIPS↓
DTC-Tex								
Step1X-3D	14.288	0.0861	–	–	14.207	0.0831	–	–
MV-Adapter	16.079	0.0806	–	–	15.706	0.0766	–	–
Ink3D	16.365	0.0795	–	–	17.128	0.0757	–	–
UniTEX	16.132	0.0711	–	–	16.572	0.0712	–	–
MVPainter	16.789	0.0815	15.530	0.0912	17.640	0.0770	15.286	0.0887
TEXGen	22.814	0.0380	–	–	18.029	0.0746	–	–
LumiTex	18.598	0.0594	18.083	0.0643	19.434	0.0620	17.342	0.0693
Hunyuan3D 2.1	16.553	0.0696	15.206	0.0768	17.979	0.0700	15.079	0.0789
TRELLIS.2	18.609	0.0660	17.093	0.0713	19.803	0.0621	16.885	0.0698
Pixal3D*	21.837	0.0304	19.884	0.0363	20.377	0.0510	17.450	0.0585
Ours	25.372	0.0132	23.870	0.0156	22.134	0.0444	19.264	0.0500
Complex-Tex								
Step1X-3D	13.431	0.0737	–	–	14.003	0.0651	–	–
MV-Adapter	14.748	0.0708	–	–	14.963	0.0630	–	–
Ink3D	15.801	0.0639	–	–	16.752	0.0595	–	–
UniTEX	15.261	0.0593	–	–	16.249	0.0562	–	–
MVPainter	14.927	0.0769	13.430	0.0881	16.272	0.0662	13.512	0.0783
TEXGen	21.974	0.0266	–	–	18.005	0.0551	–	–
LumiTex	17.419	0.0516	15.722	0.0607	18.264	0.0514	15.677	0.0604
Hunyuan3D 2.1	15.041	0.0694	13.145	0.0808	16.268	0.0624	13.262	0.0739
TRELLIS.2	17.123	0.0561	15.545	0.0623	18.493	0.0503	16.209	0.0577
Pixal3D*	20.062	0.0389	18.291	0.0453	19.402	0.0461	17.180	0.0566
Ours	22.927	0.0202	21.721	0.0254	20.488	0.0401	17.882	0.0474

D Additional Quantitative Results

D.1 Reference-View and Novel-View Fidelity

Table 7 separates the designated input view from the other 15 evaluation views, averaging novel-view scores per object before averaging over objects. For each method and metric, both splits use the same valid objects as Table 1. Novel views include both previously observed and unobserved surfaces; this split measures viewpoint generalization, rather than error exclusively on unseen surface regions.

SUPRA-Tex achieves the best mean in every reported metric on both view splits and datasets. Against the strongest baseline for shaded PSNR_{fg} in each split, reference-view gains are 2.56 dB on DTC-Tex and 0.95 dB on Complex-Tex (TEXGen), while novel-view gains are 1.76 and 1.09 dB (Pixal3D*). The split supports improved reference-view fidelity together with improved novel-view appearance.

D.2 2D Material Prediction

We evaluate the 2D material frontend directly in image space, before 3D projection, sparse encoding, or completion. The evaluation includes all 100 objects of DTC-Tex and all 110 objects of Complex-Tex, using each object's designated reference view. Predictions use one deterministic step at native 1024×1024 resolution, CFG 1, and seed 42. All 210 cases complete successfully, with no exclusions. The auxiliary normal output is generated but is not scored here.

Targets and metric domains. We reuse the benchmark's existing reference-view PNG targets without resizing or inverse display/gamma conversion. Each image is independently composited onto black using its own alpha, when present. Albedo PSNR and albedo, metallic, and roughness MAE use the GT foreground ($\alpha > 0.5$). Albedo full-image PSNR uses the mean squared error over all RGB pixels with unit data range. Albedo SSIM uses the full image with a 7×7 window and unit data range; VGG-LPIPS also uses the full image. None of the full-image metrics masks prediction pixels with the GT foreground. Metallic and roughness predictions are read from the red and green channels of the MRA output. Metrics are computed per object and averaged separately within each dataset.

Metrics are computed in the stored PNG encoding; metallic and roughness MAE therefore measure encoded-value differences rather than physical-linear material error.

Table 8: Standalone 2D material prediction. Metrics are averaged over objects in the benchmark PNG code domain; metallic and roughness MAE do not measure physical-linear error.

Dataset	n	Alb. PSNR_{fg} ↑	Alb. $\text{PSNR}_{\text{full}}$ ↑	$\text{SSIM}_{\text{full}}$ ↑	$\text{LPIPS}_{\text{full}}$ ↓	Alb. MAE_{fg} ↓	M MAE_{fg} ↓	R MAE_{fg} ↓	Frontend (s)↓
DTC-Tex	100	20.19	28.25	0.9678	0.0331	0.0872	0.0267	0.2540	1.682
Complex-Tex	110	19.44	28.19	0.9596	0.0427	0.1011	0.1167	0.2025	1.684

Results. The measured foreground PSNR is approximately 19–20 dB, so these results do not establish near-ground-truth observation accuracy. Full-image PSNR is 28.25 and 28.19 dB on DTC-Tex and Complex-Tex, respectively. Like full-image SSIM and LPIPS, it includes the shared background and should be interpreted separately from foreground PSNR. This experiment measures the standalone frontend; the effect of its errors on 3D generation is a separate comparison in Appendix E.

D.3 Runtime and Memory

D.3.1 3D Sampling

Table 9 compares Joint-update attention with SUPRA-Tex on all 210 given-geometry cases, using one H800 GPU per request. Both receive identical geometry and observation latents and use 25 Euler steps, seed 42, FP32 parameters, BF16 autocast, and FlashAttention 3 without classifier-free guidance. We report the mean of per-object median times from three runs after two warmups. These runtime measurements use the 10,000-step Joint-update attention checkpoint and the 30,000-step SUPRA-Tex checkpoint, separately from the matched-training quality ablations in Table 4; this timing comparison does not isolate interleaving alone.

SUPRA-Tex achieves sampling speedups of $1.20\times$ on DTC-Tex and $1.18\times$ on Complex-Tex. Observation KV caching adds a further $1.067\times$ and $1.064\times$ speedup, including cache construction, with identical latent outputs at every Euler step. It increases mean peak allocation by 1.94 and 1.81 GiB, respectively. The interleaved schedule also reduces query-token evaluations by 17.54% and 16.91% relative to querying all surface cells at every block; these counts are not total FLOP reductions.

Table 9: Runtime and sampling memory. Speedup is relative to Joint-update attention. Peak mean/max summarize sampling allocations; stage sum adds separately timed pipeline stages, not continuous end-to-end latency.

Dataset	Configuration	Sampling (s)↓	Speedup	Stage sum (s)↓	Peak mean (GiB)	Peak max (GiB)
DTC-Tex	Joint-update attention	10.103	1.000×	26.840	12.586	13.757
	Ours	8.440	1.197×	25.177	12.255	13.794
	Ours + KV cache	7.909	1.277×	24.647	14.191	16.978
Complex-Tex	Joint-update attention	9.970	1.000×	21.542	12.562	14.293
	Ours	8.440	1.181×	20.012	12.243	14.699
	Ours + KV cache	7.931	1.257×	19.503	14.049	17.340

Stage sums exclude model loading and file I/O. Geometry and observation preparation are timed once; frontend inference, sampling, and decoding use warm measurements. Peak memory covers the resident denoiser and sampling tensors, not the complete pipeline.

D.3.2 2D FRONTEND

Each request runs on one NVIDIA H800. After two warmup calls per worker, synchronized wall time covers input loading, conditioning, the single DiT call, and VAE/PIL decoding; model initialization and PNG writing are excluded. The mean isolated DiT latency is 1.205 s on both datasets, and peak allocated memory is 66.25 GiB. Four independent workers produce the cohort, with one GPU per request.

Repeated DiT forward timing. We additionally measure the single DiT forward pass with the model loaded once and kept resident on one NVIDIA H800. After 10 excluded warmup calls, we time 30 repetitions of one reference image and 30 calls cycling through six reference images (three each from DTC-Tex and Complex-Tex). CUDA-event mean forward times are 1.1990 s and 1.1987 s, respectively, rounded to 1.20 s in the main text. Each call jointly predicts albedo, metallic–roughness, and normal outputs at 1024×1024 with BF16, one inference step, CFG 1, and seed 42. The reported interval covers only the DiT forward pass; model loading, conditioning, VAE encoding/decoding, and file I/O are outside this interval.

E Ablation and Observation Analyses

E.1 Ablation Configurations

Table 4 reports the architecture and material-content ablations. Table 10 specifies the observation source and update rule for each row. The architecture variants use the same training data, initialization protocol, optimization settings, training budget, and checkpoint-selection rule. All variants use the given-geometry protocol of Sec. 4.1: GT geometry and cameras, matched valid observation support and sampling seed, and 25 Euler steps without classifier-free guidance. Scores average 16 views per object on each dataset; shaded and albedo PSNR_{fg} use the GT mask, and MR MAE averages metallic and roughness errors on the GT–prediction foreground intersection.

Table 10: Configurations for the six rows of Table 4. All observation conditions are encoded with the pretrained material VAE. $\widehat{\mathbf{a}}$, $\widehat{m}$, and $\widehat{r}$ denote predicted albedo, metallic, and roughness.

Variant	Material observations	Conditioning architecture
Residual injection	Predicted PBR	Observation latents added as layerwise residuals.
Joint-update attention	Predicted PBR	Observation states updated in selected layers; all generation states updated in every layer.
Fixed MR	$\widehat{\mathbf{a}}$; $m = 0$, $r = 0.5$	Read-only observations; interleaved generation updates.
Dataset-mean MR	$\widehat{\mathbf{a}}$; dataset-mean m, r	Same architecture as Ours.
Ours	Predicted PBR	Read-only observations; alternating unobserved-cell and full-surface updates.
Oracle PBR (GT)	GT PBR	Reported separately as a material-quality reference.

How observations are used. The architecture rows compare the conditioning designs as a whole.

What observations carry. The material-content comparison retains predicted albedo and replaces only metallic and roughness with fixed or dataset-mean constants. Ours uses all predicted PBR channels. Oracle PBR replaces the local observations with GT materials and is reported separately as a material-quality reference. The standalone image-space accuracy of the 2D frontend is reported in Appendix D.2.

E.2 Partial and Complete Material Encoding

Sharing the material encoder does not imply that a partial observation yields the corresponding slice of a complete-material latent. We isolate this distinction using GT geometry, cameras, and PBR values on Complex-Tex, bypassing both 2D prediction and flow generation. For each object, we encode complete materials and the filtered visible subset with the same frozen encoder using its posterior mean, then compare features at identical observed coordinates. We measure the RMS difference over coordinates and channels after dividing each channel by its fixed material-latent standard deviation, and average the resulting scores equally across objects. We retain objects that pass the coordinate and projection checks, excluding failed projection-IoU checks or incompatible geometry/material support. All retained objects have nonzero differences: the mean standardized RMS is 0.713 (95% object-bootstrap interval 0.684–0.742), with median 0.680, while repeated encoding of identical inputs gives zero difference (Fig. 8a). Thus, partial encoding changes the representation even where the input material values are unchanged.

Latent distance alone does not establish material degradation. We therefore decode the preselected cases nearest the 10th, 50th, and 90th percentiles of latent difference with the same frozen material decoder. Each arm retains its encoder's spatial correspondence cache, so this tests material reconstruction on supplied geometry rather than autonomous recovery of surface support. Figure 8b compares identical GT-selected crops; albedo PSNR is computed on the same retained GT surface voxels for both arms. Full/partial PSNR is 43.39/29.67, 40.68/26.52, and 27.41/24.37 dB for the cup, dartboard, and canister, respectively. Partial reconstructions retain the main patterns but introduce additional detail and color errors, including a pink shift in the dartboard's white regions. The canister also

shows appreciable full-input VAE error, which must be separated from the additional partial-input error. These three examples demonstrate decodable material content with imperfect fidelity; they are not a cohort-wide reconstruction estimate. Together, these results show that partial encoding supplies directly decodable material evidence while leaving room for refinement during full-surface generation.

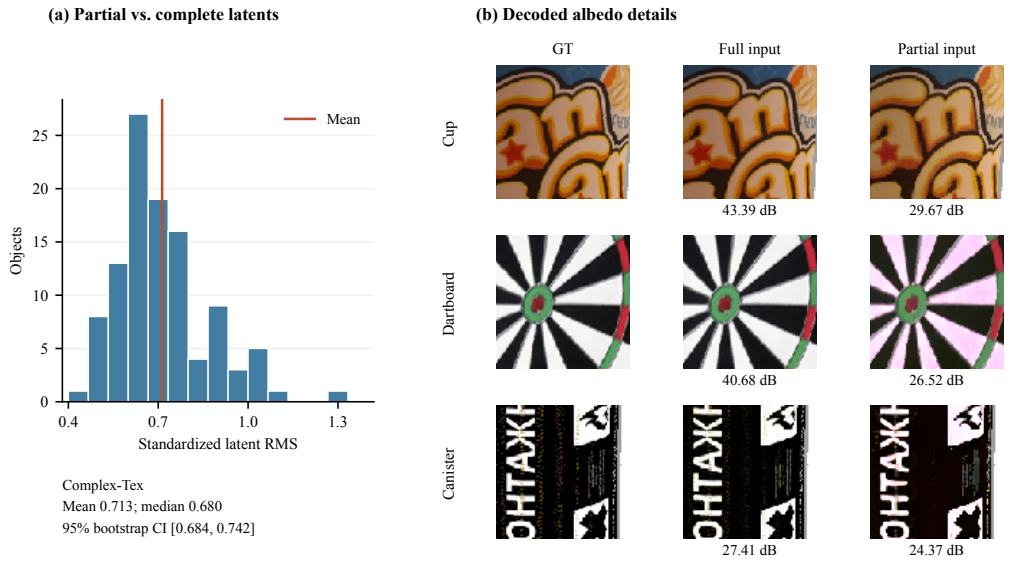

Figure 8: Native observation encoding with GT geometry and materials. (a) Partial-versus-complete latent differences across valid Complex-Tex objects. (b) Decoded albedo for three cases selected by latent-difference percentile, with shared GT-selected crops. PSNR covers all retained observed GT voxels, not only the displayed crops.

Diagnostic protocol. Figure 8 reports the paired GT partial-versus-complete encoding analysis and three illustrative material reconstructions. The diagnostic uses each object's designated reference view, a two-pixel foreground erosion, and a 3×3 depth-discontinuity filter with relative threshold 0.04; projection IoU must be at least 0.90. The latent summary weights objects equally and uses 10,000 object-bootstrap resamples. The decoded examples are the objects nearest the 10th, 50th, and 90th percentiles of object-level latent RMS. Their 128×128 crops maximize GT albedo gradient on a 16-pixel candidate grid with at least 80% valid pixels. Material values are converted to $[0, 1]$ and clamped as in the production decoder path; PSNR uses the same retained GT surface voxels for both reconstructions. Geometry support is supplied by each encoder's spatial correspondence cache, so coordinate coverage is not evidence of independent geometry recovery.

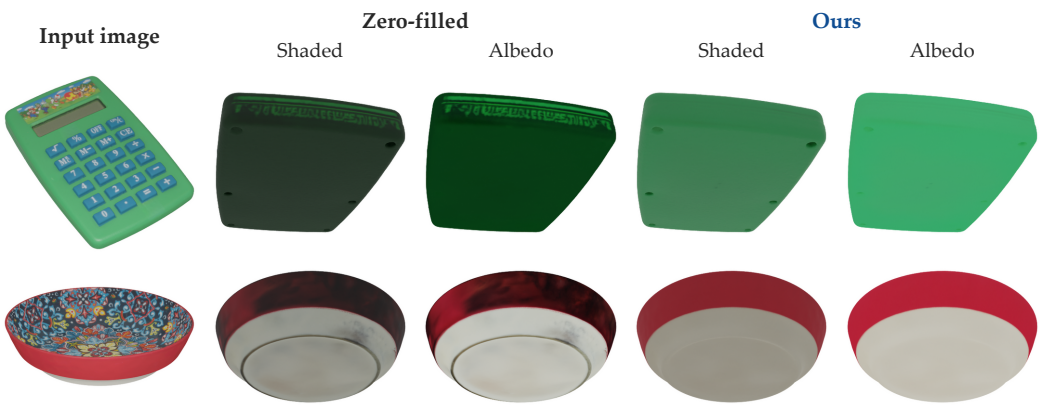

Figure 9: **Dark bias on unseen surfaces with zero-filled observations.** The calculator and dish are shown with their input images, followed by shaded and albedo renderings from the zero-filled model and SUPRA-Tex. The zero-filled results darken the unseen surfaces; SUPRA-Tex retains the corresponding green and red-and-white materials. Each comparison uses the same geometry and perspective view.

E.3 Observed-Support Encoding versus Zero Filling

As discussed in Sec. 3.3, zero-padding unobserved regions assigns the missing values concrete material semantics, such as black albedo. In the zero-filled formulation, we observed a pronounced dark bias on unseen surfaces. Figure 9 illustrates this behavior: the back of the calculator becomes nearly black and retains spurious front-side characters, while the dish develops dark mottling on its red side and white underside. With observed-support-only encoding, SUPRA-Tex instead produces a green calculator back and preserves the dish's red-and-white color regions.

E.4 Early Learning with Native PBR Latents

The native PBR latent is an *in-domain* observation representation: visible material evidence is encoded by the same pretrained material VAE that defines the generator's output space. We examine how quickly the generator learns to recover reference-view appearance from this representation, using the official Pixal3D* (Li et al., 2026a) implementation as an image-feature-conditioned comparison.

Both models are trained on the full collection of 571,051 training instances, and initialized from the same pretrained TRELLIS.2 denoiser with their respective architectural adaptations and fresh optimizers. TeaPot is used only for evaluation, not for single-object fitting. We evaluate all 16 raw checkpoints from step 0 to 1500 at intervals of 100 on its designated front view; there are no failed or excluded checkpoints. Our model receives ground-truth albedo, metallic, roughness, and opacity through native partial-material encoding, while Pixal3D* receives the corresponding front-lit RGB image through its official image-feature projection interface. This oracle-material setting examines the 3D model's ability to learn from native observations independently of frontend prediction errors. The two conditioning designs retain their respective architectures and optimization recipes.

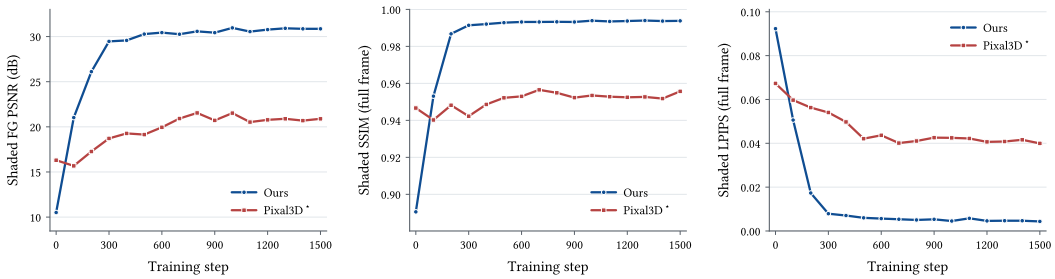

Figure 10: **Early reference-view learning on TeaPot during full-data training.** Ours uses GT PBR observations; Pixal3D⋆ uses the corresponding shaded RGB image. Each point is an evaluated checkpoint. Most of our model's improvement occurs within the first 300 steps.

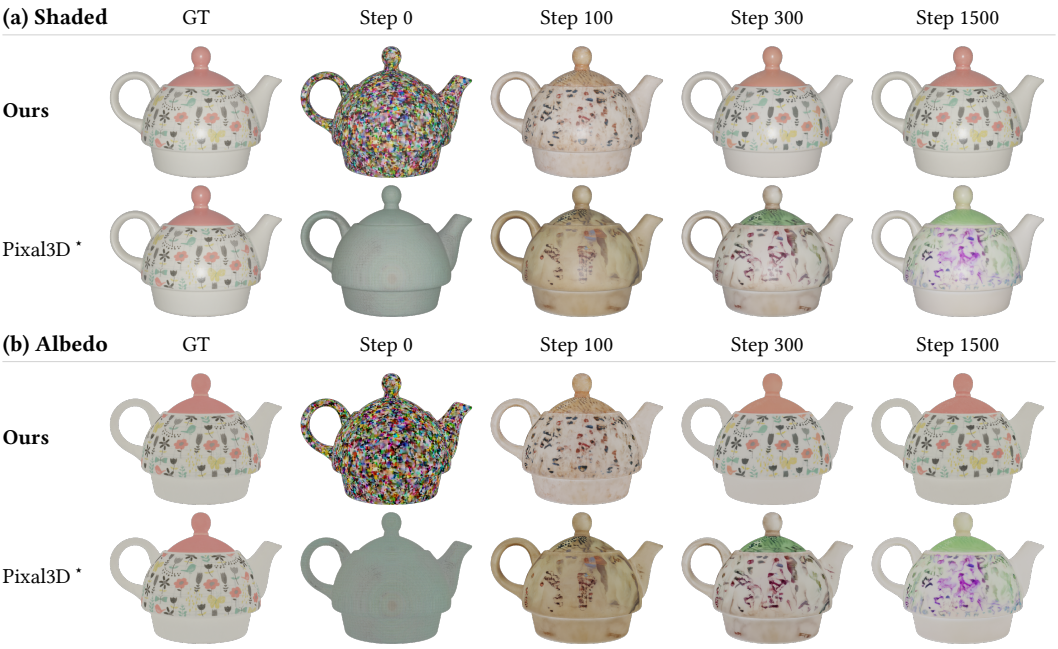

Figure 11: **Reference-view appearance during early training.** Renderings correspond to the experiment in Fig. 10; our model recovers the reference colors and floral patterns within 300 steps.

All checkpoints use the same GT geometry and camera, 25 Euler sampling steps, seed 42, and time rescaling factor 3, with no reprojection repainting or material postprocessing. We render with Blender 3.0.1 using Filmic color management and 32 samples at 1024×1024, then resize to 512×512 with Lanczos for scoring. Foreground PSNR uses the same 31,016 GT-mask pixels without erosion or intersection with the prediction mask. SSIM and VGG-LPIPS use the full image, including the background.

With GT PBR observations, our model reaches 29.46 dB shaded foreground PSNR on the TeaPot reference view after 300 steps, compared with 18.72 dB for Pixal3D⋆ under shaded RGB conditioning. At step 1500, the scores are 30.86 versus 20.89 dB, SSIM is 0.9938 versus 0.9557, and LPIPS is 0.00436 versus 0.03996. Figure 11 shows the early recovery of reference colors and floral patterns. This single-view diagnostic supports the learnability of native observations under oracle-material conditioning; it does not isolate conditioning from the models' architectural and optimization differences. Full-surface and dataset-level quality are evaluated separately in the main experiments.

F Additional Qualitative Comparisons

F.1 Given-Geometry Texturing

Figure 12 provides additional given-geometry comparisons across twelve cases.



Figure 12: Additional qualitative comparisons with given geometry across twelve cases. Novel views show the completed appearance.

F.2 End-to-End Generation

With predicted geometry, the question in the images is whether local details still land on the right surface. Figure 13 places each input image alongside reference-view and novel-view renderings from each method to compare how these details carry over to the generated asset. Reference views highlight text, logos, repeated patterns, and color boundaries; novel views show the continuity of the completed appearance.

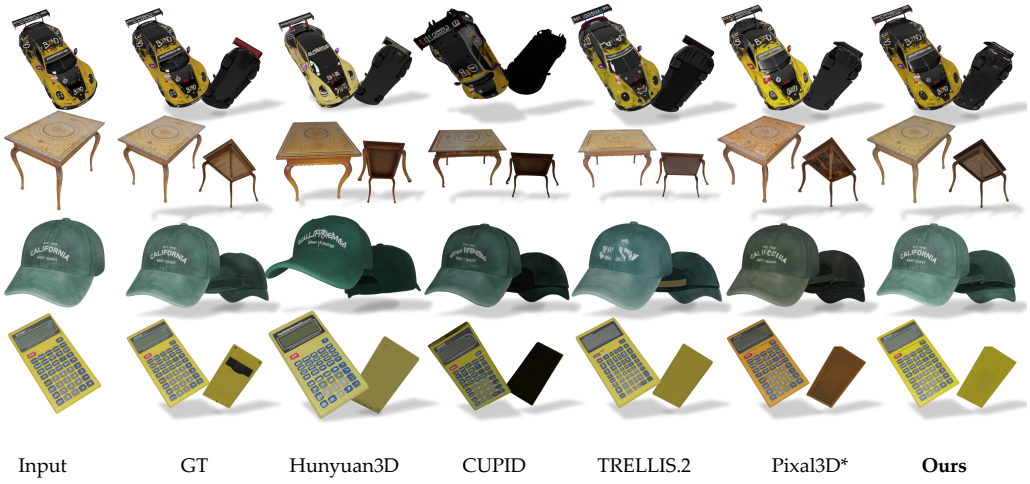

Figure 13: Qualitative comparison of end-to-end single-image asset generation (Table 2) across four cases. All methods use the same reference image. Pixal3D* and SUPRA-Tex share predicted geometry; other baselines use their own geometry pipelines. For each object we show the input, the reference-view rendering, and a novel view. Reference views assess text and repeated patterns; novel views assess the continuity of the completed appearance.

F.3 Commercial Comparisons

Figure 14 complements Figure 4 with eight additional cases under the protocol in Appendix C.4. Front views show the reproduction of text, patterns, and color layout from the input, while back views show the completion of unseen surfaces.



Figure 14: Additional comparisons with commercial models across eight cases. Each row shows the input image and front/back renderings from Hunyuan3D, Rodin, Tripo, Meshy, Hi3D, and SUPRA-Tex. Model versions and evaluation settings are given in Appendix C.4.

G EDITING AND LIMITATIONS

G.1 SINGLE-VIEW GENERATION AND BACK-VIEW EDITING

Because any image attached to the surface becomes an anchor (Sec. 3.3), the same generator accepts a second view or an edited view without retraining: a 2D editing model can supply appearance for surfaces the reference never shows, or change the appearance of surfaces it does. Both put the unseen surface under user control: the single-view results of Secs. 4.1 and 4.2 show what the generator completes on its own, and the extensions show what it produces when instructed. We demonstrate both with the same 3D generator and no additional 3D training.

Synthetic back-view conditions. From the same reference photograph, we synthesize a geometry-consistent back view and encode it as an additional material condition. Specifically, we select a viewpoint with minimal visual overlap relative to the reference, typically the opposite back side, render its surface normal map, and synthesize its shaded appearance via an MLLM, such as GPT-image, guided by the reference image, the rendered normal map, and optional user prompts. Both shaded views are translated into PBR maps, aggregated in 3D coordinate space via projection, and encoded into unified anchors O for full-surface completion. We use this two-view variant for comparisons with commercial systems; the benchmark comparisons in Tables 1 and 2 use the single-view model.

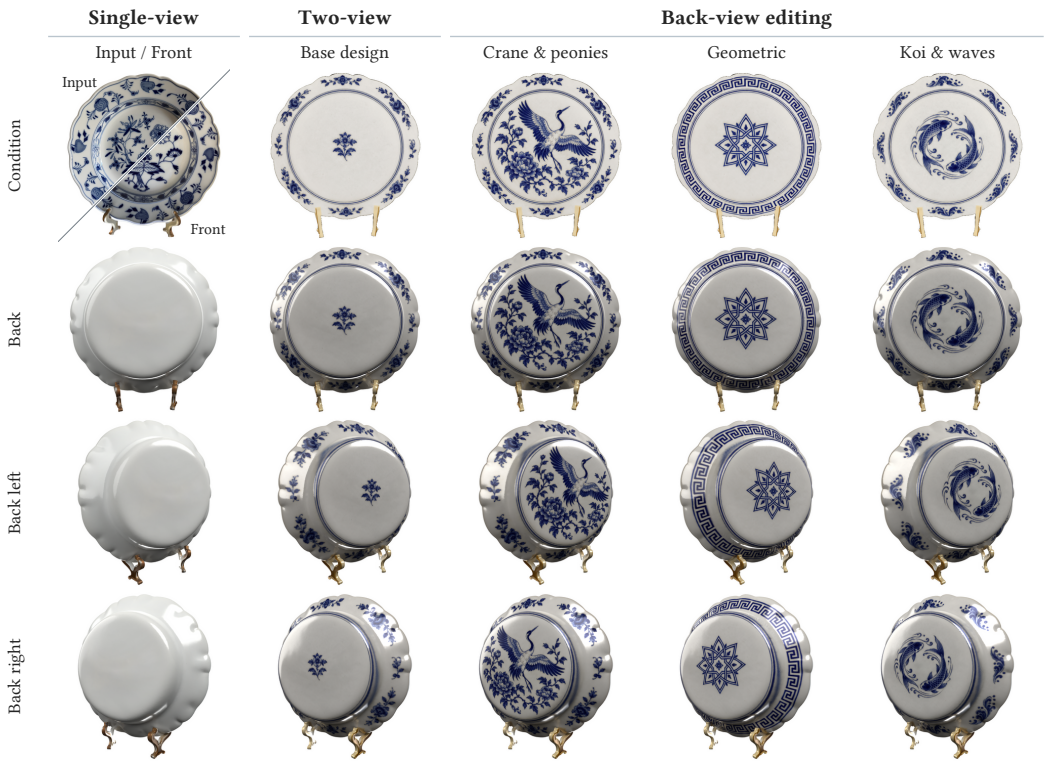

Figure 15: Single-view generation, two-view fusion, and back-view editing. The first condition panel diagonally compares the input photograph (upper left) and generated front (lower right); the other condition panels show designed back-view materials.

Texture editing. Editing the observation condition transfers image edits to the 3D surface. We use the image editing model to modify the front or back view, rebuild that view's observations, and regenerate the complete material. The visualizations show how the requested edit appears on the corresponding surface and how the edited region stays consistent with the rest of the asset.

Porcelain-plate example. Figure 15 demonstrates single-view generation, two-view fusion, and image-space material editing on a porcelain plate. Starting from the same front photograph and source

geometry, the single-view model completes the unobserved back. A geometry-guided 2D design supplies a second material observation; editing that design produces crane-and-peony, geometric, and koi variants without additional 3D training. The synthetic back is a design hypothesis rather than a second photograph of the real object.

We keep the source geometry, front material observation, model checkpoint, and sampling settings fixed. Front and back material observations are combined on the shared surface before native material encoding. The diagonal input–render comparison shows reference-view fidelity, while the back and opposite rear oblique views show how the edited motifs are carried onto the complete asset.

G.2 Limitations and Failure Cases

Figure 16 illustrates how imperfect geometry, uncertain material estimates, and missing observations limit the final asset quality.

Large image–geometry discrepancies. SUPRA-Tex can tolerate limited misalignment between the reference image and the input geometry, including geometry predicted by Pixal3D*. However, large discrepancies in shape or pose can place projected material observations on incorrect surface locations. Correspondence filtering removes some unreliable observations, but cannot correct the underlying geometry, which remains fixed during material generation. Severe discrepancies can therefore prevent faithful texture alignment. More accurate, better-aligned geometry—potentially supplied by systems such as Meshy 7.1 and Hi3D 3.0—could therefore translate directly into more precise material placement and finer reference-detail preservation, allowing SUPRA-Tex to benefit from advances in geometry generation without redesigning its texture model.

Hallucinated geometry on unseen surfaces. Upstream geometry generators such as Pixal3D* can introduce implausible structures on the unseen side of an object, even when the visible shape agrees with the reference. These structures may deviate from the shape configurations learned by the material generator, making coherent texture completion more difficult. Since SUPRA-Tex conditions on fixed geometry, it cannot remove these structures; their influence on texture generation can therefore reduce the visual quality of the final asset.

Inconsistent completion across visibility boundaries. Completing unseen surfaces is particularly challenging when text or patterns extend beyond the visible region. For example, lettering that wraps around a can may be only partially visible in the input. The model may preserve the visible fragment yet fail to continue its strokes or layout coherently onto the hidden side, producing inconsistent text or patterns across the visibility boundary.

Errors in 2D material estimation. The 2D frontend can misinterpret material properties, for example by assigning high metallic values to a non-metallic surface. These estimates condition the 3D generator, which may retain the error in the completed asset and produce implausible reflections under rendering or relighting. Accurate transfer of reference details thus does not guarantee correct material decomposition.

Failure cases. Figure 16 illustrates these limitations. Reference fidelity depends on usable image–surface correspondence and reliable material observations, while fine detail on unseen surfaces must be generated. The plate shows a printed motif reconstructed as protruding geometry, and the sign retains spurious structures after texture refinement. The owl separates recognizable appearance details from an ambiguous reflectance estimate; the teapot separates sharp observed patterns from less precise unseen completion. These are qualitative cases, without corrected-geometry controls or measured unseen-surface ground truth. Additional back-view conditioning may help the last case, but no paired two-view result is claimed here.



Figure 16: Limitations under imperfect inputs and incomplete observations. (a) The front output exhibits local motif distortion around the dragon head; the oblique view reveals a spurious protrusion in the underlying geometry. (b) Two front-oblique views show spurious protrusions in the upstream geometry that remain visible after texture refinement. (c) The glossy painted-wood appearance is assigned a high metallic value. Without material ground truth, this example illustrates ambiguity in single-image material inference rather than a verified estimation error. Grayscale maps range from black (0) to white (1). (d) Observed and unobserved base-color views show a loss of motif precision on the hidden side; boxes identify the enlarged details. The front output in (a) and both outputs in (d) use unlit base color; the oblique output in (a) and both outputs in (b) use PBR rendering under the same neutral lighting.

H Additional Related Work

Multi-view texture generation. Most texturing pipelines synthesize appearance in 2D views and fuse the predictions onto the mesh. Early approaches paint the mesh view-by-view with depth-guided diffusion and inpaint unobserved UV regions (Richardson et al., 2023; Chen et al., 2023; Zeng et al., 2024a). TexFusion aggregates view-wise denoising predictions in a shared latent texture map (Cao et al., 2023). Subsequent works improve cross-view coherence by synchronizing multi-view denoising (Liu et al., 2024b; Cheng et al., 2025; Bensadoun et al., 2024) or generating views jointly via geometry-conditioned diffusion (Huang et al., 2025b; Shao et al., 2025; Liang et al., 2026; Han et al., 2026; Zhang et al., 2026), a recipe also adopted by large image-to-3D systems (Zhang et al., 2024a; Li et al., 2025b; Tencent Hunyuan3D Team, 2025b). Reference-guided methods also explicitly target appearance preservation: EASI-Tex combines mesh-edge guidance with image features for single-image texturing (Perla et al., 2024), while AlignTex aligns multi-view artwork with the mesh before projection and texture refinement (Zhang et al., 2025b). Optimization-based approaches such as DreamMat distill geometry- and light-aware diffusion priors into PBR materials (Zhang et al., 2024b). Paint-it instead optimizes convolutionally parameterized PBR texture maps with score distillation (Youwang et al., 2024). To produce relightable assets, multi-view generation is extended to PBR channels using lighting-aware priors, material latent representations, or alternating view-space baking (Deng et al., 2024; Huang et al., 2025a; He et al., 2025a; Yeo et al., 2026; Feng et al., 2025; Tencent Hunyuan3D Team, 2025a; Lai et al., 2025; Bao et al., 2026; Wu et al., 2026b), often aided by single-image intrinsic delighting (Kocsis et al., 2024; Zeng et al., 2024b; Ke et al., 2025; Chen et al., 2024; He et al., 2025b; Careaga & Aksoy, 2024). Nevertheless, residual inconsistencies between synthesized views can blur or ghost sharp reference details during projection and blending. In contrast, SUPRA-Tex completes materials in native 3D space, using projected observations as conditioning anchors rather than directly fusing synthesized views into the final texture.

Native texture generation. Driven by the rise of large-scale 3D pretraining, native texture generation has emerged as an alternative paradigm. Existing approaches explore diffusion over point, octree, and UV representations (Yu et al., 2023; Liu et al., 2024a; Yu et al., 2024) and PBR prediction with octree-aligned Gaussians (Xiong et al., 2025). Latent texture generators, such as NaTex, LaFiTe, and TEXTRIX, learn structured latent spaces for surface colors or attributes (Lai et al., 2026; Chen et al., 2026; Zeng et al., 2026), while native 3D foundation models jointly represent geometry and appearance (Chen et al., 2025; Xiang et al., 2025; 2026). Beyond global conditioning, partial observations and spatially aligned features have also been explored. At scene scale, WorldGrow (Li et al., 2026b) extends TRELLIS with context-conditioned 3D block inpainting for progressive scene generation. TEXGen back-projects partial UV maps (Yu et al., 2024), NaTex encodes initial textures for refinement (Lai et al., 2026), and TEXTRIX encodes projected appearance with its pretrained attribute VAE in the generation target's latent space (Zeng et al., 2026). For spatial feature alignment, camera projection techniques originally developed for implicit reconstruction (Saito et al., 2019; Yu et al., 2021a) have been adapted to 3D generative models, such as CUPID's pose-aligned features and Pixal3D's camera-frame back-projected multi-scale features (Huang et al., 2026; Li et al., 2026a). Even with spatially aligned image features, the 3D generator must still infer material properties, preserve reference details, and complete unobserved surfaces.

I NOTATION REFERENCE

Table 11 summarizes the notation used in the main paper and supplement.

Table 11: Notation and symbol reference.

Symbol	Domain / Dimensions	Physical or Mathematical Interpretation
Inputs and Geometry Foundation		
I	$\mathbb{R}^{H\times W\times 3}$	Input reference photograph in RGB space
$\mathcal{M}$	3D Surface	Given or upstream-predicted object surface mesh
π	Projection operator	Camera intrinsic and extrinsic projection mapping 3D to 2D
$\mathbf{p}\in\mathcal{M}$	$\mathbb{R}^3$	Continuous 3D coordinate on the mesh surface
$\mathbf{x}_s$	$\mathbb{R}^3$	Continuous 3D centroid of sparse latent cell s
PBR Materials and 2D Translation Frontend		
$T(\mathbf{p})$	$\mathbb{R}^3\times[0,1]^3$	PBR material tuple $(\mathbf{a}(\mathbf{p}), m(\mathbf{p}), r(\mathbf{p}), \alpha(\mathbf{p}))$ at surface location $\mathbf{p}$
$\mathbf{a}, m, r, \alpha$	$\mathbb{R}^3$, $[0,1]$, $[0,1]$, $[0,1]$	Albedo, metallic, roughness, and opacity physical channels
F_{2D}	Modality-DiT	MLLM-based 2D image editing model for reference-view PBR estimation
$\widehat{T}_{\text{2D}}$	$\mathbb{R}^{1024\times1024\times5}$	Predicted physical maps $(\widehat{\mathbf{a}}, \widehat{m}, \widehat{r})$; MR is packed with a constant channel for the image VAE
$\mathbf{u}$	$\mathbb{Z}^2$	2D pixel coordinate in reference image and PBR maps
Sparse Coordinate Space and 3D VAE Latents		
$\mathcal{S}$	Coordinate set ($\lvert\mathcal{S}\rvert = N$)	Set of active sparse voxel coordinates on the surface of $\mathcal{M}$
$G = \{g_s\}_{s\in\mathcal{S}}$	$\mathbb{R}^{N\times32}$	Full-surface geometry latents produced by shape VAE
$Z = \{z_s\}_{s\in\mathcal{S}}$	$\mathbb{R}^{N\times32}$	Complete full-surface material latents (locked target representation)
$E_{\text{mat}}, D_{\text{mat}}$	Neural networks	Pretrained 3D sparse material VAE encoder and decoder
$\mathcal{P}$	Sparse attributes	Full-surface PBR attribute voxels (input to E_{mat} during ground-truth pretraining)

Continued on the next page.

Table 11 (continued).

Symbol	Domain / Dimensions	Physical or Mathematical Interpretation
Native Observation Encoding and Spatial Metadata		
$\mathcal{H}_{\text{valid}}$	$\{(\mathbf{u}, \mathbf{p})\}$	Valid pixel-surface correspondence pairs retained for observation encoding
$\mathcal{P}_o$	Sparse attributes	Voxelized partial PBR attributes supported strictly on visible surface
$\mathcal{S}_o \subseteq \mathcal{S}$	Coordinate subset ($\lvert\mathcal{S}_o\rvert = M$)	Coordinate indices of latent cells that received valid observations
z_s^o	$\mathbb{R}^{32}$	Native observation latent at cell $s \in \mathcal{S}_o$ encoded by E_{mat}
$\mathcal{U} = \mathcal{S} \setminus \mathcal{S}_o$	Coordinate subset	Indices of unobserved latent cells required to be completed
$\mathcal{V}_s, \mathcal{V}_o$	Discrete voxel sets	Surface voxels in latent cell s, and the deduplicated observed surface voxels
$\rho_s \in (0, 1]$	Scalar ratio	Observation coverage fraction $\lvert\mathcal{V}_s \cap \mathcal{V}_o\rvert / \lvert\mathcal{V}_s\rvert$ of cell s
$d_s \in \mathbb{Z}_{\geq -1}$	Graph distance	6-neighbor lattice geodesic distance from cell s to $\mathcal{U}$ (-1 if unreachable)
$\mathcal{C}_o$	Structured set	Native observation representation $\{(s, z_s^o, \rho_s, d_s) \mid s \in \mathcal{S}_o\}$
Observation-Guided Completion and Interleaved Flow Transformer		
$t \in [0, 1]$	Continuous scalar	Flow matching time ($t = 1$: Gaussian noise; $t = 0$: material endpoint with residual noise $\sigma_{\min}\epsilon$)
Z_t	$\mathbb{R}^{N \times 32}$	Noisy full-surface material latents along probability trajectory at time t
$X = \{X_s^{(\ell)}\}$	$\mathbb{R}^{N \times 1536}$	Dynamic full-surface generation stream hidden states across blocks ℓ
$O = \{O_s\}$	$\mathbb{R}^{M \times 1536}$	Static read-only observation stream memory (fixed anchor representation)
$\ell \in \{0, \ldots, L-1\}$	Integer ($L = 60$)	Transformer block index
$\mathcal{J}_\ell$	Coordinate subset	Query coordinates for block ℓ ($\mathcal{J}_\ell = \mathcal{U}$ for even ℓ, $\mathcal{J}_\ell = \mathcal{S}$ for odd ℓ)
W_{in}, W_{meta}, b_o	Learnable parameters	Shared material-geometry projection, metadata projection, and observation-role bias
Q_ℓ, K_ℓ, V_ℓ	Hidden tensors	Cross-stream attention queries (from $X_{\mathcal{J}_\ell}$), keys and values (from $[O; X]$)
$d_h = 128$	Per-head dimension	Attention head width ($1536/12$)
v^*, v_θ	$\mathbb{R}^{N \times 32}$	Ground-truth flow target velocity and predicted full-surface velocity field