

MatPedia: A Universal Generative Foundation for High-Fidelity Material Synthesis

Di Luo^{1,2*} Shuhui Yang^{2*} Mingxin Yang^{2*} Jiawei Lu¹ Yixuan Tang^{2,3}
Xintong Han² Zhuo Chen² Beibei Wang^{4†} Chunchao Guo^{2†}

¹Nankai University ²Tencent Hunyuan ³Xi’an Jiaotong University ⁴Nanjing University

{diluo, jsljiawei}@mail.nankai.edu.cn, beibei.wang@nju.edu.cn

{oakyang, maddoxyang, setskytang, pathan, zhuooochen, chunchaoguo}@tencent.com



Figure 1. We introduce **MatPedia**, a foundation model with a joint RGB-PBR representation that supports diverse material-related tasks. A vibrant ensemble of robots and the “MatPedia” typography, each adorned with distinct, photorealistic materials entirely synthesized by our method, showcasing our work’s capability in generating diverse, high-fidelity materials for 3D assets.

Abstract

*Physically-based rendering (PBR) materials are fundamental to photorealistic graphics, yet their creation remains labor-intensive and requires specialized expertise. While generative models have advanced material synthesis, existing methods lack a unified representation bridging natural image appearance and PBR properties, leading to fragmented task-specific pipelines and inability to leverage large-scale RGB image data. We present **MatPedia**, a foundation model built upon a novel joint RGB-*

*PBR representation that compactly encodes materials into two interdependent latents: one for RGB appearance and one for the four PBR maps encoding complementary physical properties. By formulating them as a 5-frame sequence and employing video diffusion architectures, MatPedia naturally captures their correlations while transferring visual priors from RGB generation models. This joint representation enables a unified framework handling multiple material tasks—text-to-material generation, image-to-material generation, and intrinsic decomposition—within a single architecture. Trained on **MatHybrid-410K**, a mixed corpus combining PBR datasets with large-scale RGB images, MatPedia achieves native 1024×1024 synthesis that substantially*

* Equal Contribution.

† Corresponding author.

surpasses existing approaches in both quality and diversity.

1. Introduction

Physically-based rendering materials are fundamental to modern computer graphics, enabling the creation of photo-realistic surfaces that respond accurately to diverse lighting and viewing conditions. They are essential in applications such as visual effects, video games, virtual reality, architectural visualization, and industrial design. As rendering technologies advance—supporting higher resolutions and more complex scenes—the demand for diverse, high-quality material libraries grows substantially. However, creating these materials remains a meticulous, labor-intensive process that requires specialized expertise and tools, thereby limiting the scalability of asset production.

The extension of generative models, notably GANs [9] and diffusion models [16, 38], to PBR material synthesis has enabled the generation of complex, interdependent texture maps (such as basecolor, metallic, roughness, and normal) from various inputs [28, 32, 45, 46, 56, 60]. However, current approaches lack a unified latent representation bridging natural image appearance (RGB) and PBR material properties. This fundamental gap leads to two critical limitations. First, existing methods remain fragmented into task-specific pipelines—such as intrinsic decomposition [14, 26, 56] versus direct material generation [32, 45, 46]—preventing universal architectures capable of handling diverse material-related tasks. Second, they are constrained to training on limited PBR datasets [8, 31, 43], unable to leverage large-scale, high-quality RGB data. Consequently, the quality and diversity of synthesized materials remain far below the potential demonstrated by modern RGB image generators.

To address these challenges, we introduce **MatPedia**, a foundation model based on a *joint RGB-PBR representation* that unifies visual appearance and physical material properties within a single encoding framework. This representation supports universal architectures for diverse material tasks and enables the use of large-scale RGB image data for material synthesis. Our key observation is that RGB images already contain rich cues about material appearance, while PBR maps provide the underlying physical explanations of that appearance. Leveraging this asymmetry, we encode PBR maps conditioned on the RGB image rather than treating them as independent modalities. The four PBR maps (basecolor, metallic, roughness, normal) collectively describe the physical properties manifested in the RGB image; thus, they can be compressed into a single latent vector that captures only the complementary physical information. To achieve this compact joint encoding, we draw inspiration from video compression, where 3D variational

autoencoders (VAEs) [24, 35, 49, 58] model dependencies across temporally coherent frames. We note a natural analogy: just as consecutive video frames share temporal coherence, RGB appearance and PBR maps are physically coupled through shared material properties.

We formulate the RGB and four PBR maps as a unified 5-frame input, enabling us to adapt video diffusion architectures to learn their joint distribution via a 3D VAE.

Building on this joint representation, **MatPedia** employs a video Diffusion Transformer (DiT) operating on the joint latent space to perform diverse material generation tasks—including text-to-material, image-to-material, and intrinsic decomposition—by conditioning on different input modalities. We initialize the DiT from pre-trained video generation models [24, 35, 49, 58] and adapt it via LoRA [19], transferring visual priors from large-scale RGB data. To train **MatPedia**, we construct **MatHybrid-410K**, a mixed corpus combining open-source PBR datasets [8, 31, 43] with a large-scale RGB image collection. Beyond standard training on RGB-PBR paired data, we incorporate abundant RGB-only images to enhance the model’s generalization capability and generation quality, enabling it to leverage diverse visual knowledge from the broader image domain. By jointly modeling multiple material representations and unifying diverse material-related tasks within a single architecture, **MatPedia**—named for its encyclopedic coverage of material understanding and generation—achieves native 1024×1024 generation across a wide range of material types.

In summary, our main contributions are:

- We propose a **joint RGB-PBR representation** that enables unified material modeling while leveraging large-scale RGB image data.
- We present **MatPedia**, a unified framework handling multiple material tasks—text-to-material, image-to-material, and intrinsic decomposition—achieving native 1024×1024 generation.
- We construct and will release **MatHybrid-410K**, a large-scale mixed dataset designed to facilitate a hybrid training strategy, enhancing generation quality and diversity.

2. Related Work

2.1. Material Estimation and Generation

Recovering physically-based material representations from images is a long-standing challenge in computer graphics [10]. Traditional methods often rely on strong priors or controlled capture setups, such as known illumination [5, 20, 37], self-similarity [1], or specific spatial relations between views [53], which limit applicability in uncontrolled environments. With the introduction of large synthetic datasets [8], deep learning methods are applied to material recovery from single or multiple images [11,

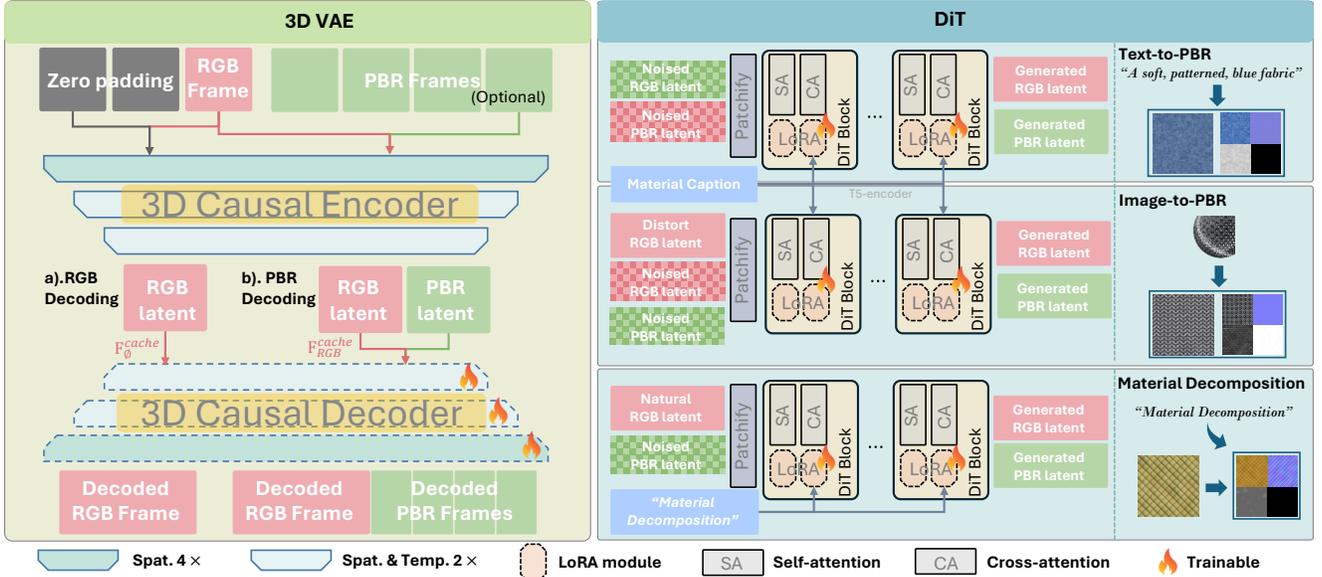


Figure 2. Pipeline of the proposed MatPedia framework. **Left:** The **3D VAE** encodes a shaded RGB frame together with optional PBR maps into a *joint RGB-PBR latent representation*, where PBR maps are conditioned on the RGB appearance. This compact representation supports both (a) shaded RGB decoding and (b) PBR decoding at native 1024×1024 resolution. **Right:** The **DiT**, initialized from large-scale video generation models and adapted via LoRA, operates on the joint latents to perform three tasks: *Text-to-PBR* (generate RGB/PBR from material captions), *Image-to-PBR* (generate planar RGB/PBR from distorted input images), and *Material Decomposition* (recover PBR maps from natural images). DiT blocks integrate self-attention (SA), cross-attention (CA), and LoRA modules to enable flexible conditioning across modalities.

30, 40, 44, 52, 59]. While these approaches relax some capture constraints, many still require near fronto-parallel views or specific lighting, making in-the-wild material extraction challenging.

Generative models offer an alternative via learned priors for material synthesis. GAN-based methods [12, 60, 61] enable unconditional generation and latent optimization, while diffusion models [38, 39] improve stability and controllability. Text2Mat [13] generates PBR materials from text using a Stable Diffusion-based model [38], and ReflectanceFusion [54] produces editable PBR maps via tandem diffusion, but both face “baked lighting” issues. ControlMat [45] synthesizes tileable high-res PBR materials from single photos via diffusion, but may have slight appearance deviations. MatFuse [46] unifies multi-modal material creation and editing but lacks diversity due to small training datasets. Material Palette [28] extracts tileable PBR materials from real images but struggles with uniform materials, illumination ambiguities, and strong geometric distortions including wrinkled or distorted surfaces. RGB \leftrightarrow X [56] synthesize RGB from specified channels and estimate intrinsic channels from images and via modality switching, but this incurs longer inference time during material decomposition. MaterialPicker [32] handles such challenging images via a video backbone but compresses frames independently (limiting resolution to 256×256). Regarding IntrinsicX [23]: they use per-channel LoRAs (one per PBR map) with cross-attention for consistency, while

ours uses per-task LoRAs on a shared joint latent space. We adopt a joint RGB-PBR representation that encodes PBR maps conditioned on the RGB appearance, achieving a high compression ratio while preserving material detail, and supporting native 1024×1024 resolution generation.

2.2. Video Diffusion Models

Video diffusion models extend denoising frameworks to spatio-temporal data, enabling coherent motion modeling and high-quality video synthesis. Early works such as Imagen Video [17] and Phenaki [47] introduced latent-space spatio-temporal attention and causal modeling for text-conditioned and long-form generation. Subsequent systems, including Video LDM [3] and Sora [4], improved multi-modal alignment, scalability, and fidelity, achieving high-resolution, long-duration videos with complex dynamics. Specialized designs like CogVideo [18] and Pyramid Video Diffusion [21] further enhanced efficiency and detail through hybrid attention and hierarchical latent spaces. Spatio-temporal VAEs are central to these pipelines, providing compact video representations while preserving temporal coherence, as demonstrated by CogVideoX [55], HunyuanVideo [24], VideoGen [25], and Improved ST-VAE [51]. Wan [49] introduces a 3D VAE with causal convolutions and high compression ratios, delivering efficient text-and-image-to-video synthesis with strong temporal consistency, which we adopt for our joint RGB-PBR representation.

3. Preliminaries

Spatio-Temporal Variational Autoencoder. Variational Autoencoders (VAEs) [3, 22, 49] map high-dimensional video data to a compact latent space for efficient generative modeling. The Wan2.2-VAE [49] adopts a 3D causal convolution design with high spatio-temporal compression: the first frame is only downsampled by $16\times$ spatially, while the remaining T frames are downsampled by $16\times$ spatially and $4\times$ temporally, producing a latent tensor of shape $[1 + T/4, H/16, W/16]$.

Material Representation. We represent each material as a Spatially Varying Bidirectional Reflectance Distribution Function parameterized by four maps: basecolor \mathbf{a} (diffuse albedo), normal \mathbf{n} (surface orientation), roughness \mathbf{r} (microfacet distribution width), and metallic \mathbf{m} (metallic factor). Together, these maps encode the material’s appearance under varying lighting and viewing conditions following the Cook-Torrance microfacet model [6, 48].

4. Method

4.1. Overview

Our goal is to develop a unified high-resolution generative model (Fig. 2) that integrates three PBR material tasks—text-to-material generation, image-to-material generation, and intrinsic decomposition—within a single architecture. At its core is a joint RGB-PBR representation (Sec. 4.2): a fine-tuned 3D VAE encodes materials into two interdependent latents, one for shaded RGB appearance and one jointly encoding four PBR maps. Building on this, we adopt a video DiT backbone with LoRA fine-tuning (Sec. 4.3) that handles all three tasks through flexible conditioning. We curate a large-scale hybrid dataset (Sec. 4.4) combining RGB-PBR pairs and RGB-only samples for training.

4.2. Joint RGB-PBR Representation

Our method builds on a *joint RGB-PBR representation*: one latent encodes the shaded RGB appearance, and another jointly encodes all four PBR maps. Crucially, RGB appearance already captures substantial visual information (texture, color, structure), while PBR maps primarily encode complementary physical properties (surface geometry, material type, reflectance). This complementary structure enables highly compact encoding—the PBR latent need only represent incremental physical attributes rather than redundantly encoding visual structure already present in RGB.

Formally, given an RGB image $\mathbf{I}_{\text{rgb}} \in \mathbb{R}^{H \times W \times 3}$ and four PBR maps (basecolor \mathbf{a} , normal \mathbf{n} , roughness \mathbf{r} , metallic \mathbf{m}), we employ a pretrained video VAE [49] that treats them as a 5-frame sequence along the temporal dimension. Video VAEs exploit redundancies across frames through 3D

convolutions, naturally capturing the structural dependencies between RGB and PBR. The encoder produces two interdependent latents:

$$\mathbf{z}_{\text{rgb}} = \mathcal{E}_{\text{rgb}}(\mathbf{I}_{\text{rgb}}), \quad \mathbf{z}_{\text{pbr}} = \mathcal{E}_{\text{pbr}}([\mathcal{F}_{\text{enc}}(\mathbf{z}_{\text{rgb}}), \mathbf{a}, \mathbf{n}, \mathbf{r}, \mathbf{m}]), \quad (1)$$

where RGB is encoded independently, while PBR maps are encoded using cached features \mathcal{F}_{enc} from the RGB encoding branch, exploiting their inherent correlation for compact representation.

The decoder mirrors this asymmetric structure:

$$\mathbf{I}'_{\text{rgb}} = \mathcal{D}_{\text{rgb}}(\mathbf{z}_{\text{rgb}}), \quad \{\mathbf{a}', \mathbf{n}', \mathbf{r}', \mathbf{m}'\} = \mathcal{D}_{\text{pbr}}([\mathcal{F}_{\text{dec}}(\mathbf{z}_{\text{rgb}}), \mathbf{z}_{\text{pbr}}]), \quad (2)$$

where RGB latents are decoded independently, while PBR latents are decoded using cached features \mathcal{F}_{dec} from the RGB decoder, producing high-fidelity PBR maps as incremental refinements. This symmetric encoder-decoder design, combined with compact dual-latent encoding, enables scalable high-resolution synthesis (Fig. 2).

While the pretrained video VAE provides strong visual priors from large-scale data, we fine-tune the decoder on PBR material data to achieve high-fidelity material reconstruction, keeping the encoder fixed to preserve the pretrained latent distribution. The decoder is optimized with pixel-wise and perceptual losses:

$$\mathcal{L}_{\text{VAE}} = \lambda_1 \|\hat{\mathbf{x}} - \mathbf{x}\|_1 + \lambda_2 \|\phi(\hat{\mathbf{x}}) - \phi(\mathbf{x})\|_2^2, \quad (3)$$

where \mathbf{x} denotes the ground-truth 5-frame sequence, $\hat{\mathbf{x}}$ is the reconstruction, and $\phi(\cdot)$ extracts features from a pretrained VGG network [41].

4.3. Universal Generative Material Model

Building on our joint RGB-PBR representation, we adopt a video DiT [49] as the generative backbone. The model treats materials as two interdependent latents: one for RGB appearance and one jointly encoding the four PBR maps. By leveraging video priors, the DiT naturally captures cross-map correlations while maintaining spatial alignment. We unify three PBR-related tasks within this single architecture through flexible conditioning and LoRA fine-tuning [19] (as shown in Fig. 2), enabling diverse material generation at 1024×1024 resolution.

Text-to-Material Generation. The DiT generates both RGB and PBR latents from noise conditioned on text prompts. The decoder then reconstructs the complete material: the RGB latent is decoded independently, and the PBR latent is decoded with feature caching from the RGB decoder, producing spatially aligned RGB appearance and PBR maps at 1024×1024 resolution. Crucially, we train this task using MatHybrid-410K, which includes both RGB-PBR paired sequences and RGB-only images. For RGB-only samples, we supervise only the RGB latent generation,

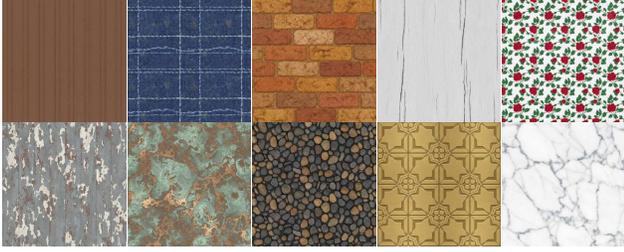


Figure 3. Examples of planar material images from the RGB Appearance Dataset, generated by the Gemini 2.5 Flash Image model [7].

allowing the model to learn rich visual priors from large-scale RGB data without requiring corresponding PBR annotations. Importantly, this strategy keeps the PBR latent distribution intact, as it continues to be learned from RGB-PBR paired data, while the RGB branch benefits from the additional diverse visual knowledge.

Image-to-Material Generation. Given a potentially distorted RGB photograph, we encode it through the VAE’s encoder as a conditioning latent. The DiT then generates two new latents—rectified planar RGB and PBR—conditioned on this encoding. During decoding, the generated RGB latent is reconstructed independently (rectifying geometric distortions), while the PBR latent is decoded with cached RGB features, producing flat material representations with accurate physical properties. This task is fine-tuned from the text-to-material checkpoint using LoRA, transferring the learned RGB priors to image-conditioned generation.

Intrinsic Decomposition. Given a planar RGB image, we encode it into an RGB latent and condition the DiT to generate only the corresponding PBR latent. The decoder then reconstructs the PBR maps using feature caching from the encoded RGB latent, decomposing the rendered appearance into underlying physical properties while ensuring spatial alignment. Similar to image-to-material generation, this task is fine-tuned from the text-to-material weights via LoRA, enabling the model to leverage RGB visual knowledge for accurate material decomposition.

Training Strategy. All tasks are optimized using the rectified flow objective [27]:

$$\mathcal{L}_{\text{RF}} = \mathbb{E}_{\mathbf{x}_0, \mathbf{x}_1, t} \left[\|v_\theta(\mathbf{x}_t, t, \mathbf{c}) - (\mathbf{x}_0 - \mathbf{x}_1)\|_2^2 \right], \quad (4)$$

where \mathbf{x}_0 denotes the ground-truth latent, $\mathbf{x}_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is sampled noise, $\mathbf{x}_t = (1 - t)\mathbf{x}_0 + t\mathbf{x}_1$ is the interpolated latent at timestep t , \mathbf{c} is the conditioning signal (text or image features), and v_θ is the velocity prediction network. For text-to-material generation, the loss is applied to both RGB and PBR latents for paired data, or solely to the RGB latent for RGB-only samples. Image-to-material and intrinsic decomposition are fine-tuned from this text-to-material checkpoint via LoRA.

Table 1. Quantitative comparison of text-to-material generation using CLIP and DINO-FID metrics. Best results are in **bold**.

	MatFuse	Ours	Ours(w/o mixed training)
CLIP↑	0.261	0.283	0.275
DINO-FID↓	1.90	1.31	1.62

4.4. Dataset

To support diverse material generation and robust material recovery, we construct MatHybrid-410K, a large-scale hybrid dataset with two complementary subsets.

RGB Appearance Dataset. We collect approximately 50,000 planar material images from two sources: (1) procedurally generated flat surfaces (as shown in Fig. 3) using Gemini 2.5 Flash Image [7], and (2) real-world planar material photographs from public repositories. Each image is paired with a text description generated by Qwen2.5-VL-72B-Instruct [2], enabling text-to-material generation. This RGB-only subset provides diverse appearance patterns for training without requiring paired PBR maps.

Complete PBR Material Dataset. We source approximately 6,000 complete PBR material sets from Matsynth [43] and additional collections. For each material, we render two types of training data using Blender’s Disney Principled BSDF [33]: (1) planar views under 32 HDR environment maps, yielding 192,000 pairs for intrinsic decomposition training; (2) distorted views rendered on geometric primitives (cubes, spheres, cylinders, cones, toruses) with varied lighting and camera angles, yielding approximately 168,000 pairs for image-to-material generation. Distorted images are cropped to ensure the material occupies at least 70% of the area.

4.5. Training and Inference

3D VAE Training. We fine-tune the decoder on RGB-PBR paired data (Sec. 4.4) at 1024×1024 for 10K steps using AdamW ($\text{lr}=5 \times 10^{-5}$, $\lambda_1 = 10$, $\lambda_2 = 1$).

DiT Training. We train the video DiT [49] with LoRA [19] (rank 128) on attention projections and FFN linear layers. Each task uses the hybrid dataset at 1024×1024 resolution, trained for 200K steps with batch size 16 and learning rate 1×10^{-4} .

Inference. Generation at 1024×1024 is upsampled to 4K using RealESRGAN [50]. Complete PBR map generation takes 20s with 50 sampling steps.

5. Experiments

We evaluate our method on the same test set as MaterialPicker [32]. Our experiments are organized into four parts: text-to-material generation, image-to-material generation, material decomposition, and ablation studies. We adopt the following evaluation metrics: **CLIP score** [36] measures semantic alignment, evaluating *text-image* sim-

Table 2. Quantitative comparison of image-to-material generation using CLIP score and DINO score metrics. Best results are in **bold**.

CLIP score \uparrow	basecolor	Normal	Roughness	Render
MatFuse	0.833	0.906	0.873	0.859
Material Palette	0.813	0.875	0.780	0.824
Ours	0.943	0.927	0.903	0.923
DINO score \uparrow	basecolor	Normal	Roughness	Render
MatFuse	0.649	0.755	0.717	0.677
Material Palette	0.543	0.579	0.448	0.541
Ours	0.907	0.762	0.752	0.843

ilarity for text-conditioned generation and *image-image* similarity for image-conditioned generation. **DINO score** [42] measures perceptual similarity between generated and reference images using DINOv2 [34] embeddings. **DINO-FID** extends the Fréchet Inception Distance [15] by replacing Inception features with DINOv2 embeddings, providing perceptually relevant distribution similarity. **MSE** measures pixel-wise reconstruction error. **LPIPS** [57] quantifies perceptual similarity using deep feature distances, correlating well with human judgments. For all metrics, higher CLIP and DINO scores indicate better alignment, while lower DINO-FID, MSE, and LPIPS indicate better fidelity. More results are provided in the supplementary.

5.1. Text-to-Material Generation

Since MaterialPicker [32] and ControlMat [45] are not publicly available, we perform qualitative comparisons with these methods using the cases presented in the MaterialPicker paper. Quantitative evaluation is conducted only against MatFuse [46], a unified diffusion-based framework supporting multimodal material creation and editing, using CLIP score (text-image similarity) and DINO-FID (distribution similarity).

As shown in Table 1, our method achieves higher CLIP score and lower DINO-FID than MatFuse, indicating better semantic consistency and visual fidelity. The qualitative comparison in Fig. 4 further demonstrates that our method produces materials with more accurate patterns, material attributes, and surface details. For instance, in the “Oak floor glossy” case, our method correctly captures both the wood grain structure and realistic gloss distribution, while baseline methods struggle with either texture clarity or material properties.

5.2. Image-to-Material Generation

We compare our method against MatFuse and Material Palette [28], which extract PBR materials from single real-world images via diffusion-based texture synthesis and PBR decomposition. Evaluation uses CLIP score (image-image similarity) and DINO score (perceptual similarity).

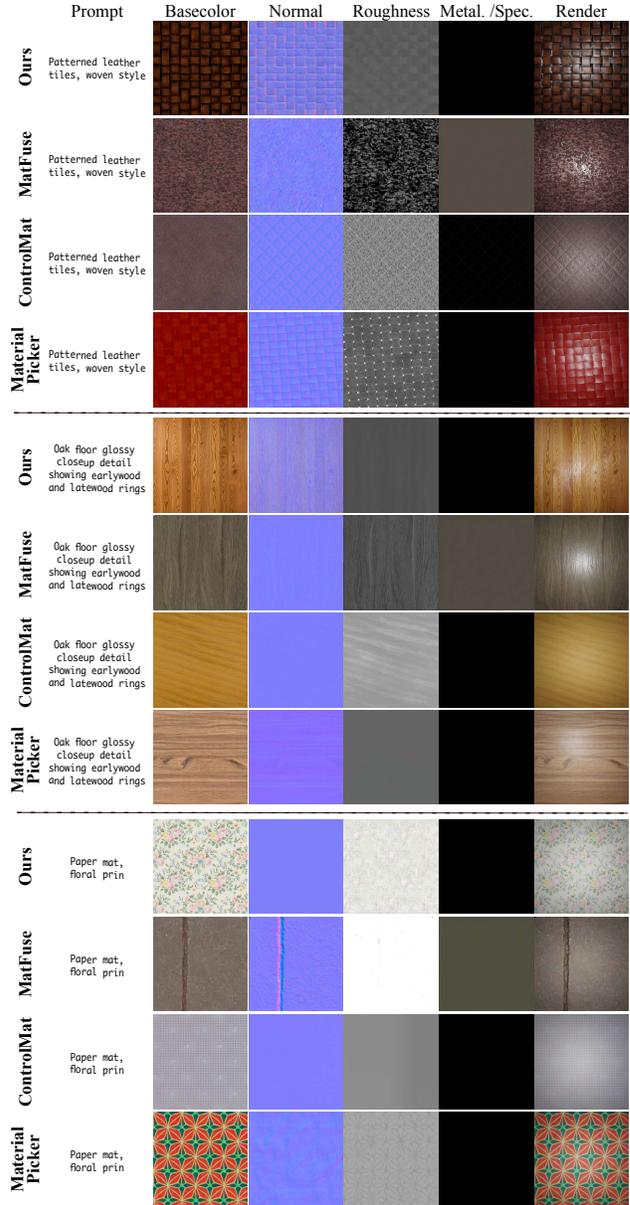


Figure 4. Qualitative comparison of text-conditioned PBR material generation among our method, MatFuse [46], ControlMat [45], and MaterialPicker [32]. For each prompt, we show the generated PBR maps (Basecolor, Normal, Roughness, Metallic) followed by a render view under point-light illumination. We note that MatFuse generates a specular map rather than a metallic map.

Our method achieves the highest CLIP and DINO scores across all channels (Table 2), with notable gains in *basecolor* (+0.11 CLIP, +0.26 DINO over MatFuse), indicating improved recovery of intrinsic colors under distortion. Qualitatively (Fig. 5), our framework generates flattened, artifact-free basecolor, normal, and roughness maps with fine detail and spatial consistency. In contrast, MatFuse ex-

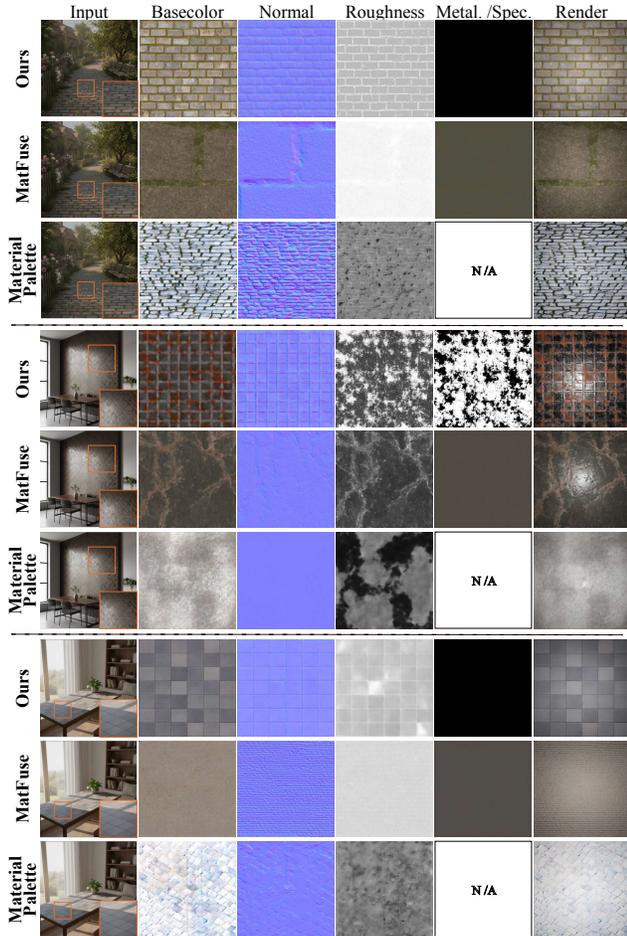


Figure 5. Qualitative comparison of image-conditioned PBR generation. For each sample, the first column shows the distorted input image (cropped from the scene), and the second to last columns present the generated material maps together with a rendering under point-light illumination. Our method produces geometrically flattened and artifact-free maps, while MatFuse shows reduced roughness fidelity and Material Palette retains geometric distortions from the input.

hibits muted texture detail and reduced roughness fidelity, while Material Palette fails to remove geometric distortions, leaving shading and perspective artifacts in the generated maps.

5.3. Material Decomposition

We compare our method against Material Palette [28] and RGB \leftrightarrow X [56]. Evaluation uses MSE (pixel-wise error) and LPIPS (deep feature perceptual distance) for reconstruction quality assessment.

Table 3 and Fig. 6 compare our method with Material Palette and RGB \leftrightarrow X for PBR material decomposition from single RGB inputs. Across diverse material categories, our approach consistently achieves the lowest MSE and LPIPS, indicating more accurate pixel-level reconstruction of PBR

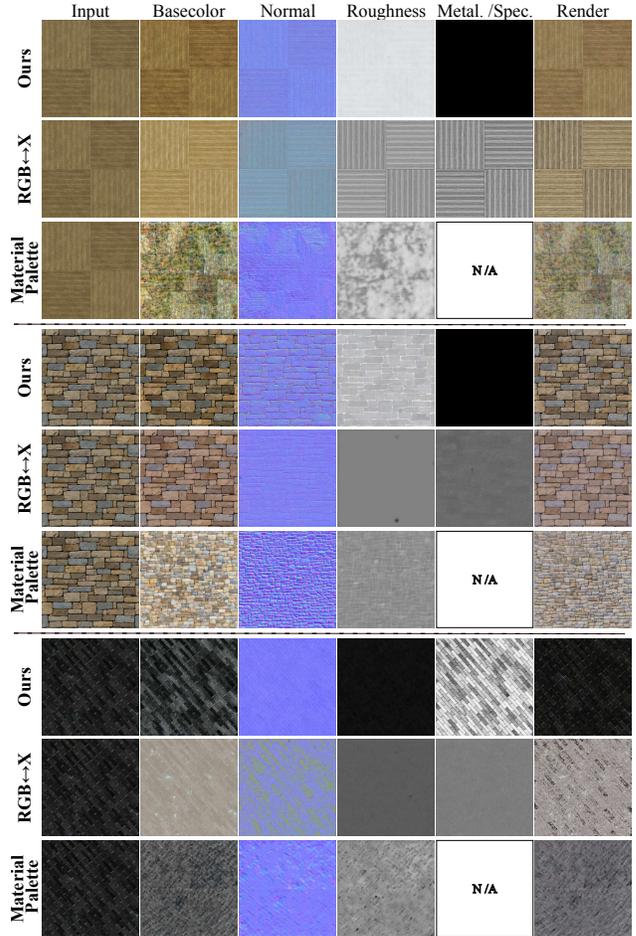


Figure 6. Qualitative comparison of material decomposition. For each sample, the first column shows the planar input image, and the second to last columns present the generated material maps together with a rendering under environment lighting. Our method produces consistent structural patterns, yielding rendered views that closely match the input appearance.

Table 3. Quantitative comparison of material decomposition using MSE and LPIPS metrics. Best results are in **bold**.

MSE↓	basecolor	Normal	Roughness	Render
Material Palette	0.058	0.014	0.103	0.034
RGB \leftrightarrow X	0.122	0.034	0.083	0.079
Ours	0.009	0.007	0.041	0.008
LPIPS↓	basecolor	Normal	Roughness	Render
Material Palette	0.716	0.513	0.761	0.644
RGB \leftrightarrow X	0.770	0.643	0.622	0.706
Ours	0.664	0.419	0.423	0.627

maps. Qualitatively, it produces basecolor maps with accurate colors and minimal noise, normal maps with sharp geometric details, and roughness maps with coherent structural patterns, yielding rendered views that closely match the inputs. In contrast, RGB \leftrightarrow X and Material Palette exhibit artifacts, loss of fine details, and tonal inconsistencies,

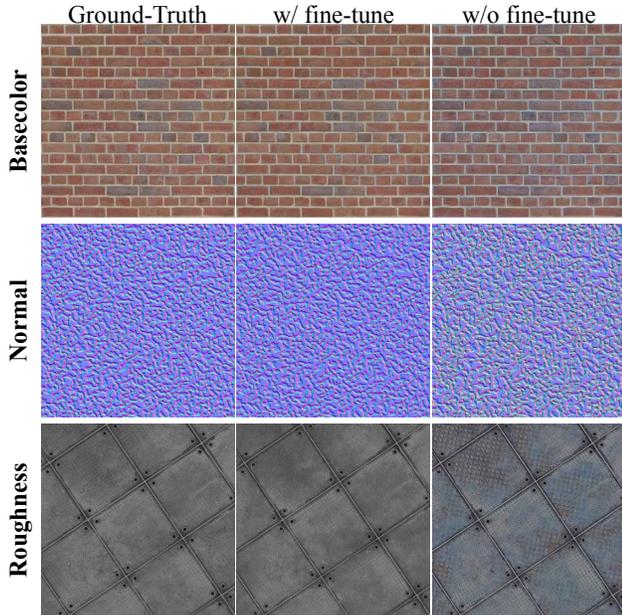


Figure 7. Visual comparison of VAE reconstruction quality before and after decoder fine-tuning. Fine-tuning significantly improves fidelity for basecolor, normal, and roughness maps, with sharper details and reduced artifacts.

especially in challenging textures such as woven surfaces and irregular stone.

5.4. Ablation Study

3D VAE Decoder Fine-tuning. Table 4 shows that decoder fine-tuning substantially improves reconstruction quality, with the largest gains in *Normal* (+3.55 dB) and *Roughness* (+5.20 dB), both critical for material appearance. Fig. 7 confirms sharper reconstructions with fewer artifacts.

Table 4. Reconstruction quality (PSNR) before and after decoder fine-tuning on our PBR dataset. Fine-tuning improves fidelity across all channels.

Method	Render	Basecolor	Normal	Roughness	Metallic
Before fine-tuning	33.12	29.27	27.29	31.36	42.45
After fine-tuning	34.25	31.30	30.84	36.56	45.59

Effect of Hybrid Training. We evaluate the impact of hybrid training by comparing our full model with a variant trained solely on the Complete PBR Material Dataset. Table 1 shows that excluding the RGB Appearance Dataset degrades performance: CLIP score drops from 0.283 to 0.275, and DINO-FID increases from 1.31 to 1.62. This confirms that incorporating the RGB Appearance Dataset enhances both semantic alignment and perceptual realism in text-conditioned generation.

Generalization to Point-Light Illumination. We evaluate material extraction from inputs captured under point-

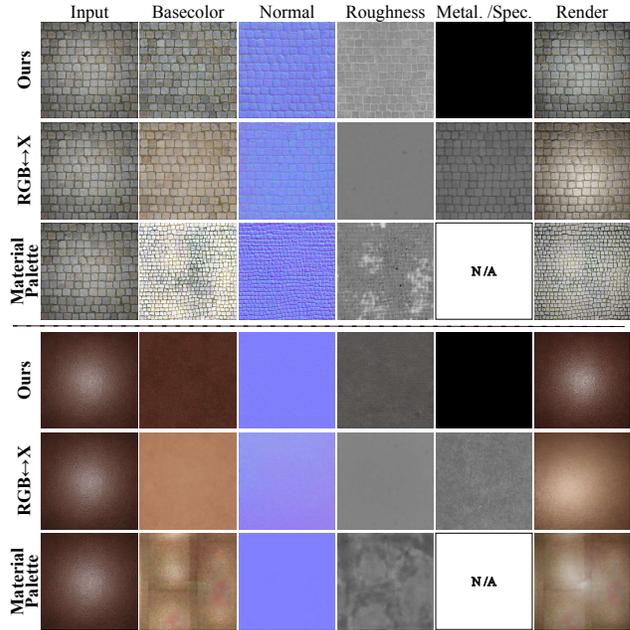


Figure 8. Qualitative comparison under point-light illumination, which introduces strong shading and non-uniform lighting absent from training data.

light illumination, which introduces strong shading and non-uniform lighting patterns. As shown in Fig. 8, our method produces basecolor, normal, and roughness maps whose rendered appearance closely matches the input image, indicating effective disentanglement even under challenging illumination. In contrast, RGB↔X shows color shifts and reduced detail fidelity, while Material Palette exhibits severe artifacts in the extracted maps.

6. Conclusion and Discussion

We introduce MatPedia, a unified PBR material generation framework built on a novel joint RGB-PBR representation. By treating materials as a 5-frame sequence and leveraging a video VAE, our approach captures cross-map correlations while enabling the use of abundant RGB data to overcome PBR data scarcity. MatPedia unifies three material-related tasks—text-to-material generation, image-to-material generation, and intrinsic decomposition—within a single architecture, producing high-fidelity results at 1024×1024 resolution that surpass existing methods in quality, diversity, and scalability. While the joint compression approach couples spatial features, limiting direct support for tileable generation via noise rolling, the native 1024×1024 resolution (upsampled to 4096×4096) provides sufficiently large textures for most production scenarios. Future work will focus on expanding to additional material channels (e.g., height, subsurface scattering) to enable more comprehensive physical modeling for advanced rendering pipelines.

Acknowledgment

We thank the reviewers for the valuable comments. This work has been partially supported by the National Natural Science Foundation of China under grant No. 62572230.

References

- [1] Miika Aittala, Tim Weyrich, Jaakko Lehtinen, et al. Two-shot svbrdf capture for stationary materials. *ACM Trans. Graph.*, 34(4):110–1, 2015. 2
- [2] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 5, 1
- [3] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22563–22575, 2023. 3, 4
- [4] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, et al. Video generation models as world simulators. *OpenAI Blog*, 1(8):1, 2024. 3
- [5] Manmohan Chandraker. On shape and material recovery from motion. In *European Conference on Computer Vision*, pages 202–217. Springer, 2014. 2
- [6] Robert L Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *ACM Transactions on Graphics (ToG)*, 1(1):7–24, 1982. 4
- [7] Google DeepMind. Gemini 2.5 flash image. <https://aistudio.google.com/models/gemini-2-5-flash-image>, 2025. 5, 1
- [8] Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. Single-image svbrdf capture with a rendering-aware deep network. *ACM Transactions on Graphics (ToG)*, 37(4):1–15, 2018. 2
- [9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 2
- [10] Darya Guarnera, Giuseppe Claudio Guarnera, Abhijeet Ghosh, Cornelia Denk, and Mashhuda Glencross. Brdf representation and acquisition. In *Computer graphics forum*, pages 625–650. Wiley Online Library, 2016. 2
- [11] Jie Guo, Shuichang Lai, Chengzhi Tao, Yuelong Cai, Lei Wang, Yanwen Guo, and Ling-Qi Yan. Highlight-aware two-stream network for single-image svbrdf acquisition. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021. 2
- [12] Yu Guo, Cameron Smith, Miloš Hašan, Kalyan Sunkavalli, and Shuang Zhao. Materialgan: Reflectance capture using a generative svbrdf model. *arXiv preprint arXiv:2010.00114*, 2020. 3
- [13] Zhen He, Jie Guo, Yan Zhang, Qinghao Tu, Mufan Chen, Yanwen Guo, Pengyu Wang, and Wei Dai. Text2mat: Generating materials from text. 2023. 3
- [14] Zebin He, Mingxin Yang, Shuhui Yang, Yixuan Tang, Tao Wang, Kaihao Zhang, Guanying Chen, Yuhong Liu, Jie Jiang, Chunchao Guo, et al. Materialmvp: Illumination-invariant material generation via multi-view pbr diffusion. *arXiv preprint arXiv:2503.10289*, 2025. 2
- [15] 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. *Advances in neural information processing systems*, 30, 2017. 6
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2
- [17] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. 3
- [18] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*, 2022. 3
- [19] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022. 2, 4, 5, 1
- [20] Zhuo Hui and Aswin C Sankaranarayanan. A dictionary-based approach for estimating shape and spatially-varying reflectance. In *2015 IEEE International Conference on Computational Photography (ICCP)*, pages 1–9. IEEE, 2015. 2
- [21] Yang Jin, Zhicheng Sun, Ningyuan Li, Kun Xu, Hao Jiang, Nan Zhuang, Quzhe Huang, Yang Song, Yadong Mu, and Zhouchen Lin. Pyramidal flow matching for efficient video generative modeling. *arXiv preprint arXiv:2410.05954*, 2024. 3
- [22] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 4
- [23] Peter Kocsis, Lukas Höllein, and Matthias Nießner. Intrinsic: High-quality pbr generation using image priors. *arXiv preprint arXiv:2504.01008*, 2025. 3
- [24] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*, 2024. 2, 3
- [25] Zongjian Li, Bin Lin, Yang Ye, Liuhan Chen, Xinhua Cheng, Shenghai Yuan, and Li Yuan. Wf-vae: Enhancing video vae by wavelet-driven energy flow for latent video diffusion model. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 17778–17788, 2025. 3
- [26] Ruofan Liang, Zan Gojcic, Huan Ling, Jacob Munkberg, Jon Hasselgren, Chih-Hao Lin, Jun Gao, Alexander Keller, Nandita Vijaykumar, Sanja Fidler, et al. Diffusion renderer: Neural inverse and forward rendering with video diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 26069–26080, 2025. 2

- [27] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. 5
- [28] Ivan Lopes, Fabio Pizzati, and Raoul de Charette. Material palette: Extraction of materials from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4379–4388, 2024. 2, 3, 6, 7
- [29] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 1
- [30] Di Luo, Hanxiao Sun, Lei Ma, Jian Yang, and Beibei Wang. Correlation-aware encoder-decoder with adapters for svbrdf acquisition. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–10, 2024. 3
- [31] Xiaohe Ma, Xianmin Xu, Leyao Zhang, Kun Zhou, and Hongzhi Wu. Opensvbrdf: A database of measured spatially-varying reflectance. *ACM Transactions on Graphics (TOG)*, 42(6):1–14, 2023. 2
- [32] Xiaohe Ma, Valentin Deschaintre, Miloš Hašan, Fujun Luan, Kun Zhou, Hongzhi Wu, and Yiwei Hu. Materialpicker: Multi-modal material generation with diffusion transformers. *arXiv preprint arXiv:2412.03225*, 2024. 2, 3, 5, 6
- [33] Stephen McAuley, Stephen Hill, Naty Hoffman, Yoshiharu Gotanda, Brian Smits, Brent Burley, and Adam Martinez. Practical physically-based shading in film and game production. In *ACM SIGGRAPH 2012 courses*, pages 1–7. 2012. 5
- [34] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafranec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 6
- [35] Xiangyu Peng, Zangwei Zheng, Chenhui Shen, Tom Young, Xinying Guo, Binluo Wang, Hang Xu, Hongxin Liu, Mingyan Jiang, Wenjun Li, et al. Open-sora 2.0: Training a commercial-level video generation model in 200 k. *arXiv preprint arXiv:2503.09642*, 2025. 2
- [36] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 5
- [37] Jérémy Riviere, Pieter Peers, and Abhijeet Ghosh. Mobile surface reflectometry. In *ACM SIGGRAPH 2014 Posters*, pages 1–1. 2014. 2
- [38] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2, 3
- [39] Sam Sartor and Pieter Peers. Matfusion: a generative diffusion model for svbrdf capture. In *SIGGRAPH Asia 2023 conference papers*, pages 1–10, 2023. 3
- [40] Liang Shi, Beichen Li, Miloš Hašan, Kalyan Sunkavalli, Tamy Boubekeur, Radomir Mech, and Wojciech Matusik. Match: Differentiable material graphs for procedural material capture. *ACM Transactions on Graphics (TOG)*, 39(6):1–15, 2020. 3
- [41] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 4
- [42] George Stein, Jesse Cresswell, Rasa Hosseinzadeh, Yi Sui, Brendan Ross, Valentin Villicroze, Zhaoyan Liu, Anthony L Caterini, Eric Taylor, and Gabriel Loaiza-Ganem. Exposing flaws of generative model evaluation metrics and their unfair treatment of diffusion models. *Advances in Neural Information Processing Systems*, 36:3732–3784, 2023. 6
- [43] Giuseppe Vecchio and Valentin Deschaintre. Matsynth: A modern pbr materials dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22109–22118, 2024. 2, 5
- [44] Giuseppe Vecchio, Simone Palazzo, and Concetto Spampinato. Surfacenet: Adversarial svbrdf estimation from a single image. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12840–12848, 2021. 3
- [45] Giuseppe Vecchio, Rosalie Martin, Arthur Roullier, Adrien Kaiser, Romain Rouffet, Valentin Deschaintre, and Tamy Boubekeur. Controlmat: a controlled generative approach to material capture. *ACM Transactions on Graphics*, 43(5):1–17, 2024. 2, 3, 6
- [46] Giuseppe Vecchio, Renato Sortino, Simone Palazzo, and Concetto Spampinato. Matfuse: controllable material generation with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4429–4438, 2024. 2, 3, 6, 1, 4, 5
- [47] Ruben Villegas, Mohammad Babaeizadeh, Pieter-Jan Kindermans, Hernan Moraldo, Han Zhang, Mohammad Taghi Saffar, Santiago Castro, Julius Kunze, and Dumitru Erhan. Phenaki: Variable length video generation from open domain textual description. *arXiv preprint arXiv:2210.02399*, 2022. 3
- [48] Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. Microfacet models for refraction through rough surfaces. *Rendering techniques*, 2007:18th, 2007. 4
- [49] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 2, 3, 4, 5, 1
- [50] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *International Conference on Computer Vision Workshops (ICCVW)*. 5
- [51] Pingyu Wu, Kai Zhu, Yu Liu, Liming Zhao, Wei Zhai, Yang Cao, and Zheng-Jun Zha. Improved video vae for latent video diffusion model. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 18124–18133, 2025. 3
- [52] YOUXIN XING, ZHENG ZENG, YOUYANG DU, LU WANG, and BEIBEI WANG. Diffusion-guided relighting for single-image svbrdf estimation. 2025. 3
- [53] Zexiang Xu, Jannik Boll Nielsen, Jiyang Yu, Henrik Wann Jensen, and Ravi Ramamoorthi. Minimal brdf sampling for two-shot near-field reflectance acquisition. *ACM Transactions on Graphics (TOG)*, 35(6):1–12, 2016. 2

- [54] Bowen Xue, Claudio Guarnera, Shuang Zhao, and Zahra Montazeri. Reflectancefusion: Diffusion-based text to svbrdf generation. In *Eurographics Symposium on Rendering*. Eurographics Association, 2024. 3
- [55] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 3
- [56] Zheng Zeng, Valentin Deschaintre, Iliyan Georgiev, Yannick Hold-Geoffroy, Yiwei Hu, Fujun Luan, Ling-Qi Yan, and Miloš Hašan. Rgb \leftrightarrow x: Image decomposition and synthesis using material- and lighting-aware diffusion models. In *ACM SIGGRAPH 2024 Conference Papers*, New York, NY, USA, 2024. 2, 3, 7
- [57] 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*, pages 586–595, 2018. 6
- [58] Sijie Zhao, Yong Zhang, Xiaodong Cun, Shaoshu Yang, Muyao Niu, Xiaoyu Li, Wenbo Hu, and Ying Shan. Cv-vae: A compatible video vae for latent generative video models. *Advances in Neural Information Processing Systems*, 37: 12847–12871, 2024. 2
- [59] Xilong Zhou and Nima Khademi Kalantari. Look-ahead training with learned reflectance loss for single-image svbrdf estimation. *ACM Transactions on Graphics (TOG)*, 41(6):1–12, 2022. 3
- [60] Xilong Zhou, Milos Hasan, Valentin Deschaintre, Paul Guerrero, Kalyan Sunkavalli, and Nima Khademi Kalantari. Tilegen: Tileable, controllable material generation and capture. In *SIGGRAPH Asia 2022 conference papers*, pages 1–9, 2022. 2, 3
- [61] Xilong Zhou, Milos Hasan, Valentin Deschaintre, Paul Guerrero, Yannick Hold-Geoffroy, Kalyan Sunkavalli, and Nima Khademi Kalantari. Photomat: A material generator learned from single flash photos. In *ACM SIGGRAPH 2023 conference proceedings*, pages 1–11, 2023. 3