

DualMat: PBR Material Estimation via Coherent Dual-Path Diffusion

Yifeng Huang*
Stony Brook University
Stony Brook, NY, USA

Zhang Chen
Meta
Pittsburgh, PA, USA

Yi Xu
Goertek Alpha Labs
Santa Clara, CA, USA

Minh Hoai
The University of Adelaide
Adelaide, SA, Australia

Zhong Li†
Apple Inc.
Cupertino, CA, USA

ABSTRACT

We present DualMat, a novel dual-path diffusion framework for estimating Physically Based Rendering (PBR) materials from single images under complex lighting conditions. Our approach operates in two distinct latent spaces: an albedo-optimized path leveraging pre-trained visual knowledge through RGB latent space, and a material-specialized path operating in a compact latent space designed for precise metallic and roughness estimation. To ensure coherent predictions between the albedo-optimized and material-specialized paths, we introduce feature distillation during training. We employ rectified flow to enhance efficiency by reducing inference steps while maintaining quality. Our framework extends to high-resolution and multi-view inputs through patch-based estimation and cross-view attention, enabling seamless integration into image-to-3D pipelines. DualMat achieves state-of-the-art performance on both Objaverse and real-world data, significantly outperforming existing methods with up to 28% improvement in albedo estimation and 39% reduction in metallic-roughness prediction errors. Our project can be found at [yifehuang97.github.io/DualMatProjPage/](https://github.com/yifehuang97/DualMatProjPage/).

CCS CONCEPTS

• **Computing methodologies** → **Computer vision**.

KEYWORDS

Computer Vision, Computer Graphics, Physically Based Rendering, Material Estimation, Texture Synthesis, Generative Models

ACM Reference Format:

Yifeng Huang, Zhang Chen, Yi Xu, Minh Hoai, and Zhong Li†. 2025. DualMat: PBR Material Estimation via Coherent Dual-Path Diffusion. In *Proceedings of the 33rd ACM International Conference on Multimedia (MM '25)*, October 27–31, 2025, Dublin, Ireland.

*Work initiated and partially completed during his internship at OPPO US Research Center. Email: yifehuang@cs.stonybrook.edu

†Corresponding author. Email: zhonglee323@gmail.com. The work was conducted when Zhong Li, Zhang Chen, Yi Xu was with OPPO.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

MM '25, October 27–31, 2025, Dublin, Ireland.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-2035-2/2025/10

<https://doi.org/10.1145/3746027.3755237>

27–31, 2025, Dublin, Ireland. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3746027.3755237>

1 INTRODUCTION

Physically-Based Rendering (PBR) materials are fundamental to realistic computer graphics, defining how rendered surfaces should interact with light through properties such as diffuse reflection, metallic behavior, and surface roughness. These materials form the cornerstone of modern graphics applications, from video games and digital content creation to augmented/virtual reality (AR/VR). Accurate PBR material estimation is crucial for faithfully replicating real-world materials in digital environments, yet recovering these properties from a single image remains challenging due to the complex interplay of material attributes, lighting conditions, and underlying geometry.

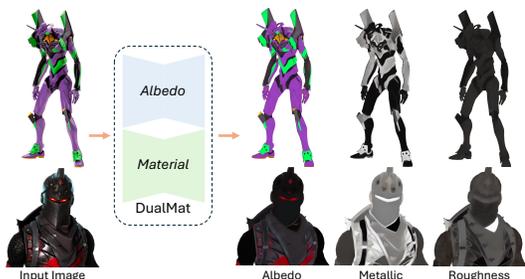


Figure 1: Given a single input image (left), DualMat’s dual-path design combines an albedo-optimized path and a material-specialized path to achieve accurate PBR material estimation.

The importance of accurate PBR material estimation has been further amplified by recent breakthroughs in image-to-3D generation. Modern image-to-3D pipelines aim to reconstruct complete 3D assets from images. Early works like DreamFusion [31] distill a 3D representation with NeRF [29] from a 2D image diffusion model with Score Distillation Sampling loss. Recent works either employ large transformer architectures [14, 18, 49, 53–55] to train a triplane representation in a data-driven manner, or explicitly generate multiview images and then reconstruct 3D models from the generated images [19, 25, 26, 28, 38, 39, 42, 48, 50]. However, these methods primarily focus on geometry and view-dependent appearance, often producing simplified or view-dependent textures

that lack explicit material properties. This limitation severely restricts their applications in scenarios requiring physically accurate rendering, such as relighting under different environment maps or realistic integration into virtual environments. A robust PBR material estimation method could bridge this gap, enabling these 3D generation pipelines to produce assets with physically accurate rendering capabilities.

To address these challenges, PBR material estimation is crucial and can be formulated as an image-to-2D map translation task. Traditionally, this task has been tackled with methods heavily reliant on deep convolutional networks [4, 5, 8–11, 15, 20–22, 24, 35, 46, 51, 57, 58, 61, 62]. While methods such as SurfaceNet have demonstrated the ability to estimate spatially-varying bidirectional reflectance distribution functions (SVBRDF) from a single image, these convolution-based architectures face significant limitations: they struggle to capture high-frequency details, resulting in blurred or oversimplified material maps; their performance degrades in multi-view scenarios where consistency across viewpoints is critical; and they lack the expressive power needed to generate coherent outputs under varying lighting conditions. The limitations of these conventional approaches suggest the need for a more powerful generative framework.

Recent advances in generative models, particularly diffusion models [12, 13, 16, 23, 32, 40, 52], offer promising solutions to these challenges. These models have demonstrated remarkable success in high-quality image synthesis [6, 30], material generation [36, 43–45, 47], super-resolution [7, 34, 59], and image translation [37, 41, 60]. Motivated by their ability to capture complex data distributions and generate high-fidelity outputs, we formulate PBR material estimation using a diffusion-based approach.

Building on these insights, we present DualMat, a novel dual-path diffusion framework for estimating PBR materials from single images, as shown in Fig. 1. Our approach operates in two distinct latent spaces to effectively capture material properties and ensure consistency with input images. The albedo-optimized path leverages pretrained visual knowledge through RGB latent space to generate high-quality albedo maps, while the material-specialized path operates in a compact latent space designed to precisely model metallic and roughness properties.

To ensure coherent predictions between these complementary paths, we introduce a feature distillation loss during training that aligns intermediate representations. We further enhance computational efficiency by incorporating rectified flow, significantly reducing the number of denoising steps needed in both training and inference stages while preserving high-quality results.

Our framework extends beyond single-image estimation through two crucial adaptations: a patch-based strategy for high-resolution inputs that preserves fine details while maintaining global consistency, and a cross-view attention mechanism for multi-view scenarios that ensures consistent material estimation across viewpoints. These capabilities position DualMat as a versatile PBR estimator that seamlessly integrates into existing image-to-3D generation pipelines, enabling applications from realistic relighting to textured mesh generation.

In short, the primary contributions of our work are: (1) we propose a novel dual-path diffusion framework for PBR material estimation that leverages complementary latent spaces, consisting of an

albedo-optimized path utilizing pre-trained visual knowledge and a material-specialized path dedicated to precise property estimation, combined with a feature distillation training strategy to ensure coherent outputs; (2) we design an efficient inference framework incorporating rectified flow sampling, patch-based high-resolution processing, and cross-view attention mechanisms, enabling high-quality material generation within just 2–4 steps and supporting seamless integration into image-to-3D pipelines; and (3) we demonstrate state-of-the-art performance across all evaluation metrics, achieving up to 28% improvement in albedo estimation and a 39% reduction in metallic-roughness prediction errors.

2 RELATED WORKS

This section reviews deep learning-based methods for material estimation (PBR or SVBRDF) and the application of diffusion models in material generation or estimation.

Material estimation. Estimating spatially-varying material parameters, such as PBR and SVBRDF maps, from images is a well-established yet challenging task. Traditional methods leverage convolutional neural networks (CNNs) to capture material properties from single images under unknown natural lighting. Early CNN-based methods [20, 58] demonstrated plausible SVBRDF estimation for specific material classes, such as metals and plastics. Subsequent works [8, 22] introduced methods that recover SVBRDFs from flash-lit or casually captured images by incorporating physical priors. Sang and Chandraker [35] extended these approaches by jointly estimating material and shape with relighting capabilities. The emergence of GAN-based methods, such as SurfaceNet [46] and Zhou and Kalantari [62], enhanced detail recovery through adversarial training while bridging the synthetic-to-real gap using real image pairs. MaterialGAN [10] and Henzler et al. [11] adapted GANs with learned latent spaces to serve as priors for inverse rendering and unsupervised SVBRDF reconstruction, respectively. Wen et al. [51] improved material diversity by generating non-repeating BRDF fields from latent codes. TileGen [61] and Hu et al. [15] further advanced GAN-based material modeling, generating tileable SVBRDFs with category-specific control and enabling texture transfer while preserving structural integrity.

Diffusion models in material estimation/generation. The emergence of diffusion models has opened new possibilities for material estimation and generation. MatFusion [36] estimates SVBRDFs using diffusion, replacing residual convolution blocks with ConvNeXt [27] blocks and injecting image conditions by modifying the first convolution layer to accept additional input channels. Vainer et al. [43] model PBR material distributions with cross-network communication, linking a new PBR model to a frozen RGB model to tackle data scarcity and high-dimensional outputs. MatFuse [47] offers a unified diffusion-based approach for 3D material creation and editing, using multiple conditioning sources and disentangled latent spaces for flexible map-level control. StableMaterials [44] applies semi-supervised latent diffusion, integrating adversarial training, a diffusion refiner, and a latent consistency model for fast, high-quality, and tileable material generation. Compared to prior CNN- and diffusion-based approaches, our dual-path diffusion framework captures high-quality material details while ensuring multi-view consistency through cross-view attention mechanisms,

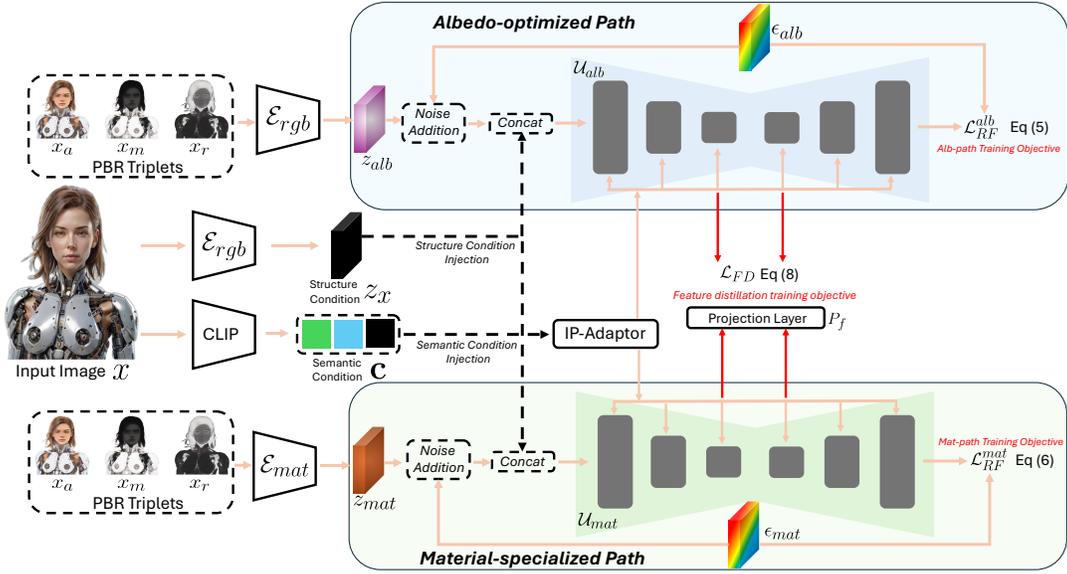


Figure 2: DualMat’s training architecture with dual paths and feature distillation. The albedo-optimized path (top) and material-specialized path (bottom) are trained with rectified flow objectives while maintaining consistency through feature distillation and dual conditioning mechanism.

addressing challenges related to diverse lighting conditions and viewpoints.

Image condition injection in diffusion. A simple approach to inject image conditions is to concatenate them with the sampled noise [17, 33]. IP-Adapter [56] refines this by extracting visual embeddings via a CLIP encoder, projecting them through a learned layer.

3 PROPOSED APPROACH

To introduce DualMat, we begin with preliminaries on diffusion models and rectified flow Sec. 3.1, followed by PBR latent space design Sec. 3.2, dual-path diffusion model Sec. 3.3, flow matching and feature transfer in training Sec. 3.4, and extensions to high-resolution and multi-view scenarios Sec. 3.5. Fig. 2 illustrate our training pipelines.

3.1 Preliminaries

Diffusion models. Diffusion models define a forward process that progressively corrupts data samples, transforming an initial data distribution $\mathbf{x}_0 \sim p_{\text{data}}(\mathbf{x})$ into a standard Gaussian distribution $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ over T time steps. The forward process follows a Markov chain, where at each step t , Gaussian noise is added according to: $\mathbf{x}_t \sim \mathcal{N}(\sqrt{\alpha_t}\mathbf{x}_{t-1}, (1 - \alpha_t)\mathbf{I})$, with variance schedule $\alpha_t \in (0, 1)$. To reverse this corruption, a neural network f_θ learns to predict the distribution $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ at each step. The training objective minimizes the prediction error of the added noise: $\mathbb{E}_{t, \mathbf{x}_0, \epsilon} [\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|^2]$, where ϵ represents the Gaussian noise from the forward process.

Rectified flow. Flow-based models offer an alternative approach by learning a continuous transformation between distributions

through an ordinary differential equation. The Rectified Flow framework defines this transformation using $\frac{d\mathbf{x}_t}{dt} = v_\theta(\mathbf{x}_t, t)$, where v_θ represents a learned velocity field and $t \in [0, 1]$. The model transforms data between the initial distribution $\mathbf{x}_0 \sim p_{\text{data}}(\mathbf{x})$ and a noise distribution $\mathbf{x}_1 \sim \pi_1(\mathbf{x})$ via linear interpolation: $\mathbf{x}_t = (1 - t)\mathbf{x}_0 + t\mathbf{x}_1$. The training objective minimizes the flow matching loss:

$$\min_{\theta} \mathbb{E}_{\mathbf{x}_1 \sim \pi_1, \mathbf{x}_0 \sim p_{\text{data}}} \left[\int_0^1 \|\mathbf{x}_1 - \mathbf{x}_0 - v_\theta(\mathbf{x}_t, t)\|^2 dt \right]. \quad (1)$$

This formulation enables high-quality generation with significantly fewer steps than traditional diffusion models.

3.2 PBR Latent Space Representation

The foundation of our approach lies in establishing effective latent representations for PBR materials within the diffusion framework. We investigate two complementary encoding strategies that, while both capable of encoding complete PBR triplets, are designed to excel at different aspects of material estimation: (1) an albedo-optimized strategy leveraging pretrained image-based diffusion models, particularly Stable Diffusion, whose rich visual priors prove essential for high-quality albedo estimation despite a less compact representation; and (2) a material-specialized strategy developing a compact and efficient latent space, specifically optimized for precise metallic and roughness property estimation. This complementary design ultimately allows us to combine the strengths of both paths—superior albedo predictions from the albedo-optimized path and precise material properties from the material-specialized path.

Albedo-optimized path via visual prior. To leverage rich visual understanding from pretrained diffusion models for high-quality albedo estimation, we design our first path based on an established RGB encoder \mathcal{E}_{rgb} . This path processes PBR material components:

albedo (x_a), metallic (x_m), and roughness (x_r) through independent encoding. The albedo component, being inherently an RGB image, is processed directly through \mathcal{E}_{rgb} . For metallic and roughness components, which are single-channel maps, we employ channel repetition to satisfy the encoder’s input specifications. The complete latent representation is formed by concatenating the individually encoded components:

$$z_{alb} = [\mathcal{E}_{rgb}(x_a), \mathcal{E}_{rgb}(R(x_m)), \mathcal{E}_{rgb}(R(x_r))], \quad (2)$$

where $R(\cdot)$ represents our channel repetition operation that transforms single-channel inputs into three-channel representations compatible with \mathcal{E}_{rgb} . Specifically, utilizing the VAE architecture from Stable Diffusion v2 (SD2) [32], this encoding process yields a 12-channel latent vector, with four channels dedicated to each material component. While this approach effectively leverages pre-trained visual understanding, particularly advantageous for high-quality albedo estimation, its independent processing of components may not fully capture the intricate interdependencies inherent in PBR materials.

Material-specialized path with unified encoding. To achieve precise metallic-roughness estimation through joint material understanding, we design our second path using a specialized vector quantized architecture. Unlike the independent processing in the albedo-optimized path, this path employs a unified encoder \mathcal{E}_{mat} that processes the complete PBR triplet (x_a, x_m, x_r) simultaneously:

$$z_{mat} = \mathcal{E}_{mat}([x_a, x_m, x_r]), \quad (3)$$

where $[x_a, x_m, x_r]$ represents the channel-wise concatenation of material components. Our architecture is specifically engineered to handle this 5-channel input (three for albedo, one each for metallic and roughness) and maps it into a compact 14-channel latent space. This unified encoding enables the capture of intricate relationships between material properties while maintaining a compact representation, crucial for precise metallic and roughness estimation. The encoder is trained with a comprehensive multi-objective loss:

$$\mathcal{L}_{mat} = \lambda_r \mathcal{L}_{rec} + \lambda_p \mathcal{L}_{perc} + \lambda_a \mathcal{L}_{adv} + \lambda_c \mathcal{L}_{code}. \quad (4)$$

This loss consists of reconstruction \mathcal{L}_{rec} , perceptual \mathcal{L}_{perc} , adversarial \mathcal{L}_{adv} , and codebook commitment \mathcal{L}_{code} terms to ensure high-quality material encoding.

3.3 Dual-Path Diffusion for PBR Estimation

To synergistically leverage the rich visual understanding from pre-trained diffusion models and specialized material modeling capabilities, we introduce DualMat, a novel dual-path diffusion framework illustrated in Fig. 2. While both paths are trained to predict complete PBR triplets (x_a, x_m, x_r), each is architecturally optimized for different components: an albedo-optimized path operating in the pre-trained visual latent space, and a material-specialized path working in our unified compact latent space.

Each path employs a dedicated denoising U-Net (U_{alb} and U_{mat}) carefully tailored to its respective latent representation. This dual-path design enables specialization where it matters most—the albedo-optimized path excels at generating high-fidelity albedo predictions by leveraging pre-trained visual knowledge, while the material-specialized path achieves precise metallic and roughness estimations through its unified material understanding. Although both

paths produce complete PBR predictions, we strategically combine their strengths in the final output: taking the albedo component from U_{alb} and the metallic-roughness components from U_{mat} , formally expressed as $(x_a^{alb}, x_m^{mat}, x_r^{mat})$. The consistency between these paths is ensured through feature distillation during training, as detailed in the following sections.

To ensure robust conditioning on the input image x , DualMat incorporates two complementary conditioning mechanisms that operate at different levels of abstraction:

1) **High-level semantic conditioning.** We leverage the rich semantic understanding of CLIP through IP-Adapter. Specifically, we extract deep visual features $\mathbf{c} = P_{ip}(\mathcal{E}_{clip}(x))$, where \mathcal{E}_{clip} is the CLIP vision encoder and P_{ip} represents the IP-Adapter projection. These semantic features are integrated via cross-attention layers throughout both U-Nets, enabling a global understanding of material properties for each path’s specialized focus.

2) **Low-level structure conditioning.** We further strengthen the connection to input image details through direct latent conditioning $z_x = \mathcal{E}_{rgb}(x)$. By concatenating this latent representation with the noise input at each denoising step, we provide fine-grained structural guidance that preserves local material details and ensures spatial coherence across predictions.

3.4 Flow Matching and Feature Transfer

To achieve both efficient training and consistent predictions across paths, we introduce two key mechanisms in our framework. First, we adopt rectified flow to enable rapid inference, modeling the continuous transformation between data and noise distributions with remarkably few sampling steps (2–4) compared to traditional diffusion approaches. Second, given that our final prediction combines albedo from the albedo-optimized path (x_a^{alb}) with material properties from the material-specialized path (x_m^{mat}, x_r^{mat}), we incorporate feature transfer strategies to ensure cross-path consistency. In our dual-path framework, each U-Net predicts a velocity field that guides the noise-to-data transformation:

$$\mathcal{L}_{RF}^{alb} = \mathbb{E}_{t,x,\epsilon_{alb}} \left[\left\| (z_{alb} - \epsilon_{alb}) - v_{\theta_{alb}}(z_t^{alb}, t, x) \right\|^2 \right], \quad (5)$$

$$\mathcal{L}_{RF}^{mat} = \mathbb{E}_{t,x,\epsilon_{mat}} \left[\left\| (z_{mat} - \epsilon_{mat}) - v_{\theta_{mat}}(z_t^{mat}, t, x) \right\|^2 \right], \quad (6)$$

where $z_{alb} = \mathcal{E}_{rgb}(x)$ and $z_{mat} = \mathcal{E}_{mat}(x)$ encode the PBR materials in their respective latent spaces, v_{θ} represents the predicted velocity field, z_t denotes the noisy latent at timestep t , and x is the input image. These velocity fields guide the progressive refinement from noise to structured material representations.

Our training pipeline consists of two strategically designed stages to ensure consistent predictions while maintaining path specialization:

1) **Albedo path training:** We initialize U_{alb} from a pre-trained Stable Diffusion v2 model and finetune it with \mathcal{L}_{RF}^{alb} . The model’s first and last convolutional layers are carefully adapted for PBR material estimation while preserving the rich visual understanding encoded in the pre-trained weights, essential for high-quality albedo prediction.

2) **Material path training:** With U_{alb} frozen, we train U_{mat} from scratch using both \mathcal{L}_{RF}^{mat} and feature distillation. This distillation transfers crucial visual knowledge from the albedo-optimized

path while ensuring consistency between material property predictions and albedo estimates.

For effective feature transfer, we align the intermediate representations across both paths through a learned projection P_f as shown in Fig. 2. At each network layer l , we minimize the L_2 distance between albedo path features and their projected material path counterparts:

$$\mathcal{L}_p^l = \left\| f_l^{alb} - P_f(f_l^{mat}) \right\|_2^2, \quad (7)$$

where f_l^{alb} and f_l^{mat} denote feature maps from corresponding layers. The total feature distillation loss aggregates these alignments across the entire U-Net hierarchy:

$$\mathcal{L}_{FD} = \sum_{l \in \mathcal{L}} \left\| f_l^{alb} - P_f(f_l^{mat}) \right\|_2^2, \quad (8)$$

where \mathcal{L} spans four down-sampling blocks, one middle block, and four up-sampling blocks. This comprehensive feature alignment ensures that while each path maintains its specialized capabilities, albedo prediction in U_{alb} and material property estimation in U_{mat} , their predictions remain coherent when combined in the final output $(x_a^{alb}, x_m^{mat}, x_r^{mat})$.

3.5 High-Resolution and Multi-View Support

High-resolution processing. While our model is trained at fixed resolutions (e.g., 256, 512), real-world applications often require processing higher-resolution images. Simply downsampling the input would result in loss of fine details, so we propose a coarse-to-fine strategy that maintains high-frequency details while ensuring global-local consistency.

First, we obtain a global image prediction I^g by applying our model to a down-sampled input. We then process the full-resolution image in overlapping patches that match the training resolution. During patch-wise inference, we blend global and local guidance:

$$\hat{z}_t = z_t + \gamma \nabla_{z_t} \left(\|D(z_0)_{\text{blur}} - I_p^g\|_2^2 + \sum_{n \in \{l, t\}} \|D(z_0)_{p_n} - I_{p_n}^n\|_2^2 \right), \quad (9)$$

where z_0 is the predicted clean latent, $D(\cdot)$ decodes a latent to RGB image space, I_p^g is the aligned patch from the global prediction, and $I_{p_n}^n$ are the overlapping regions from the previously generated left ($n = l$) and top ($n = t$) patches. This image-space formulation preserves fine details while ensuring seamless boundaries between windows.

Multi-view consistency. To ensure consistent PBR material estimation across multiple views, which is essential for effective integration into image-to-3D generation tasks, we enhance our single-view diffusion models by incorporating cross-view attention mechanisms. Specifically, we fine-tune both the albedo-optimized and material-specialized paths, initially trained for single-view predictions, using multi-view inputs (e.g., four-view images) during an additional fine-tuning stage.

During this process, each self-attention layer within our denoising U-Nets (U_{alb} and U_{mat}) is modified to perform attention operations across tokens from all viewpoints simultaneously. Given multiple input views, we first independently encode each view into its latent representation $z^{(v)}$. These latent tokens from all views are concatenated into a single sequence as: $z = [z^{(1)}, z^{(2)}, \dots, z^{(V)}]$,

where V denotes the total number of views. The concatenated tokens are then processed using a self-attention mechanism:

$$\text{Attention}(z) = \text{softmax} \left(\frac{Q(z)K(z)^T}{\sqrt{d}} \right) V(z),$$

where $Q(\cdot)$, $K(\cdot)$, and $V(\cdot)$ represent query, key, and value projections, respectively, and d is the dimension of attention heads. After this cross-view attention operation, the tokens are separated back into individual view representations: $[z^{(1)'}, z^{(2)'}, \dots, z^{(V)'}] = \text{Split}(\text{Attention}(z))$.

This cross-view attention enables each view’s latent representation to interact and integrate relevant information from other viewpoints, thereby promoting coherent and aligned material properties across all views. Fine-tuning our diffusion models with this mechanism significantly enhances multi-view consistency.

4 EXPERIMENTS

4.1 PBR Estimation

Dataset. We construct our training dataset from Objaverse by selecting 65,000 objects with valid PBR materials. Each object is rendered from 16 different viewpoints under 100 distinct HDR environment maps, yielding over one million image-PBR pairs for training.

Evaluation protocol. We evaluate our method primarily on the fundamental task of single-view material estimation. Using a test set of 654 objects from Objaverse (rendered under novel HDR environments), we assess each view independently, even when multiple views of the same object are available. For each single-view estimation, we employ multiple quality metrics: PSNR, SSIM, and LPIPS for albedo evaluation, and RMSE for metallic and roughness components. This single-view evaluation protocol provides a rigorous assessment of our method’s core material estimation capabilities, separate from its multi-view extensions.

Baselines. We compare DualMat against three representative methods: Intrinsic Anything [2] (trained on Objaverse, albedo estimation only), SurfaceNet [46] (fine-tuned with PBR prediction head), and Stable Diffusion (adapted with extended convolution layers). For a fair comparison, both SurfaceNet and Stable Diffusion are trained on our dataset with identical settings.

Implementation details. Our material-specialized encoder \mathcal{E}_{mat} is trained for 1.5M iterations with a batch size of 32. The loss weights are set as follows: reconstruction (1.0), LPIPS (0.001), commitment (0.1), and adversarial (0.01). The albedo-optimized U-Net (U_{alb}) is fine-tuned from Stable Diffusion v2 (SD2) for 500K iterations with a batch size of 64. Subsequently, we train the material-specialized U-Net (U_{mat}) from scratch for 500K iterations with a batch size of 48, employing both rectified flow and feature distillation objectives. Notably, the feature distillation loss is applied only during the initial 50K iterations. Both U-Nets share identical architectures except for their input and output layers. During inference, we achieve high-quality results using only three sampling steps. Additionally, we preserve the text-conditioning architecture with empty embeddings to facilitate potential future extensions. All experiments are conducted using four NVIDIA A100 GPUs.

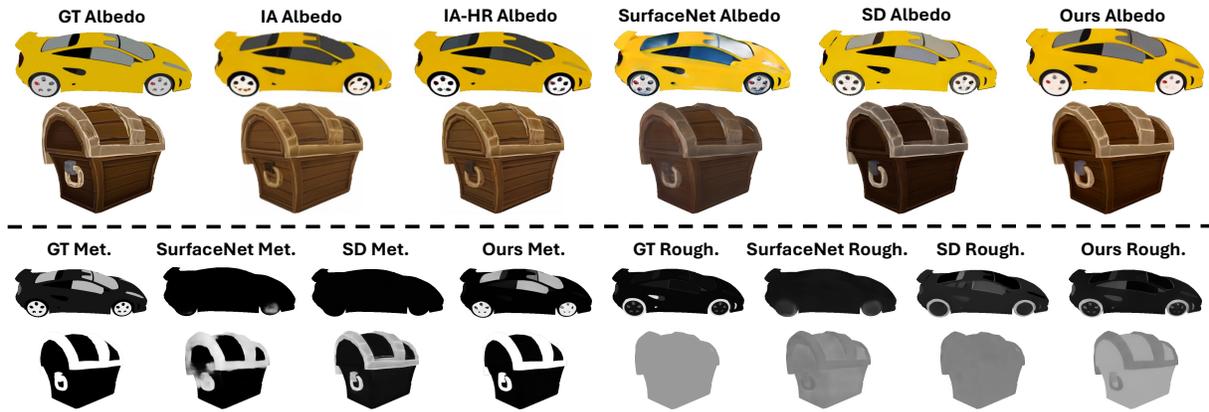


Figure 3: Comparison on Objaverse test objects, illustrating material decomposition capabilities of our method versus baseline approaches.

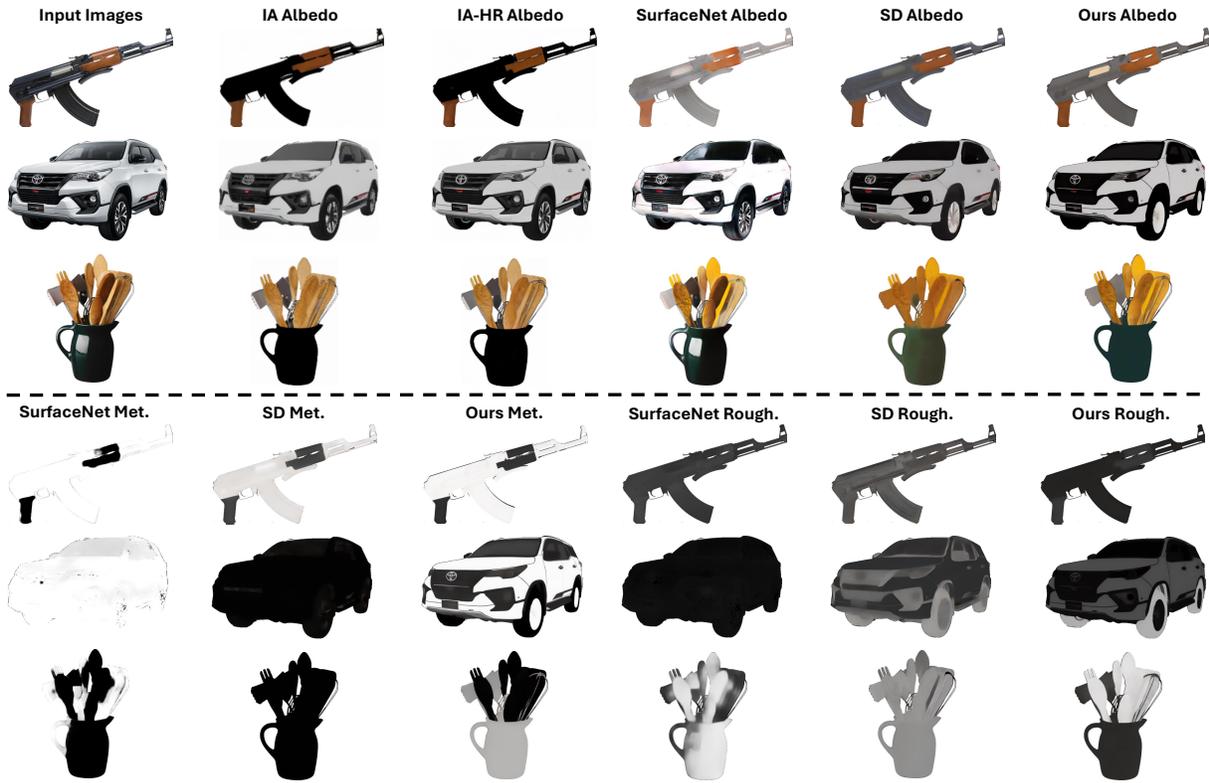


Figure 4: Results on real-world images, demonstrating our method’s superior generalization to diverse real-world materials with improved detail preservation and accurate material separation.

Results. The quantitative evaluation results are summarized in Table 1. DualMat consistently outperforms existing methods across all metrics. For albedo estimation, our method achieves significant improvements in both pixel-wise accuracy (PSNR: 28.6 dB, SSIM: 0.932) and perceptual quality (LPIPS: 0.047), surpassing Intrinsic Anything even with its high-resolution guidance enhancement

(PSNR: +4.4 dB, LPIPS: -0.026). The dual-path design proves particularly effective for metallic and roughness estimation, where we achieve RMSE values of 0.057 and 0.060, respectively, representing a 39% reduction in estimation error compared to previous methods.

| Method | Albedo | | | Met. | Rough. |
|------------------------|-------------|--------------|--------------|--------------|--------------|
| | PSNR | SSIM | LPIPS | | |
| IA [2] | 23.4 | 0.895 | 0.092 | - | - |
| IA-HR [†] [2] | 24.2 | 0.912 | 0.073 | - | - |
| SurfaceNet [46] | 26.1 | 0.931 | 0.052 | 0.093 | 0.092 |
| SD | 26.4 | 0.918 | 0.053 | 0.089 | 0.089 |
| DualMat(Ours) | 28.6 | 0.932 | 0.047 | 0.057 | 0.060 |

Table 1: Quantitative evaluation on single-view PBR material estimation. We report PSNR, SSIM, and LPIPS for albedo quality, and RMSE for metallic (Met.) and roughness (Rough.) components. Our method achieves superior performance across all metrics on the Objaverse test set. [†]IA-HR refers to Intrinsic Anything with high-resolution guidance enhancement.

4.2 Qualitative Results

Fig. 3 demonstrates our method’s qualitative performance on Objaverse test objects, where our method shows improvements in material decomposition over existing approaches. Most notably, in Fig. 4 with real-world objects, our method demonstrates significantly better generalization capability, preserving high-frequency details such as car logos and achieving cleaner albedo-metallic separation in challenging materials like car windows. These results validate our method’s effectiveness in handling complex real-world materials while maintaining intricate surface details. Additional qualitative results for highly specular objects are provided in Fig. 5.



Figure 5: Qualitative results on highly specular objects.

4.3 Ablation Study

4.3.1 Ablation on Dual-Path Architecture. We conduct ablation studies to validate the necessity of our dual-path architecture, as shown in Table 2, comparing three variants: using only the material-specialized path, only the albedo-optimized path, and our full dual-path model. The results demonstrate clear complementary benefits of both paths. The albedo-optimized path alone achieves strong albedo quality (matching our dual-path results in PSNR and LPIPS) but struggles with material property estimation. Conversely, the material-specialized path provides precise metallic-roughness estimates (matching our dual-path performance) but lower albedo quality. Our dual-path design successfully preserves the optimal components from each path, maintaining both high albedo quality and precise material property estimation.

| \mathcal{U}_{mat} | ✓ | ✗ | ✓ |
|---------------------|-------|-------|--------------|
| \mathcal{U}_{alb} | ✗ | ✓ | ✓ |
| Albedo PSNR | 27.8 | 28.6 | 28.6 |
| Albedo SSIM | 0.928 | 0.932 | 0.932 |
| Albedo LPIPS | 0.056 | 0.047 | 0.047 |
| Metallic RMSE | 0.057 | 0.065 | 0.057 |
| Roughness RMSE | 0.061 | 0.064 | 0.060 |

Table 2: Ablation study comparing single-path versus dual-path performance. Our dual-path design achieves optimal performance across all metrics.

4.3.2 Ablation on Multi-view Consistency. We present a qualitative ablation in Fig. 6, demonstrating the impact of our multi-view consistency mechanism. With multi-view consistency enabled, our model significantly enhances the coherence and alignment of predicted material maps across different viewpoints. As illustrated in the comparison, omitting multi-view consistency introduces noticeable artifacts and inconsistencies within the albedo, metallic, and roughness maps, particularly around intricate structures and boundaries. In contrast, incorporating our proposed cross-view attention mechanism ensures stable and coherent predictions, effectively preserving fine surface details and maintaining visual consistency across multiple viewpoints. These qualitative improvements clearly demonstrate the effectiveness and importance of our cross-view attention mechanism in achieving high-quality multi-view material estimation.

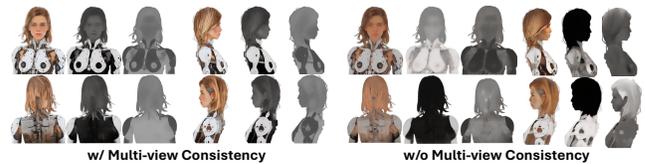


Figure 6: Qualitative ablation study demonstrating the effectiveness of our proposed multi-view consistency mechanism.

4.3.3 Ablation on Feature Distillation Loss. To evaluate the effectiveness of our feature distillation loss, we conduct an ablation study comparing model performance with and without feature distillation during the initial 50k training iterations. As shown in Table 3, incorporating the feature distillation loss consistently improves model performance across all evaluated metrics. Specifically, using feature distillation leads to significant gains in albedo estimation accuracy, increasing PSNR by 1.45 dB and SSIM by 0.007, while also improving metallic and roughness predictions by reducing their RMSE values. These results highlight the importance of feature distillation in effectively transferring visual knowledge between the albedo-optimized and material-specialized paths, ultimately enhancing overall material estimation quality.

4.3.4 Ablation on High-Resolution Processing. We present a qualitative ablation on our high-resolution processing strategy, as illustrated in Fig. 7. When high-resolution processing is enabled,

| Method | Alb. PSNR | Alb. SSIM | Met. | Rough. |
|-----------------------|-----------|-----------|-------|--------|
| W/ distillation loss | 26.64 | 0.897 | 0.070 | 0.073 |
| W/O distillation loss | 25.19 | 0.890 | 0.072 | 0.079 |

Table 3: Ablation study evaluating the effectiveness of the proposed feature distillation loss on PBR material estimation performance.

the model successfully preserves fine details and textures, closely matching the intricate features of the input image. In contrast, the version without high-resolution processing produces notably blurry and oversimplified results, losing essential facial and textural details. This ablation highlights the effectiveness and necessity of our high-resolution processing approach for maintaining detailed and realistic material predictions in practical high-resolution scenarios.



Figure 7: Qualitative ablation study illustrating the impact of our proposed high-resolution processing strategy.

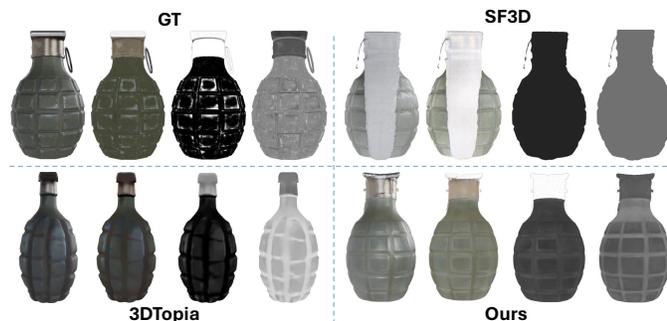


Figure 8: Comparison with image-to-3D methods. Each row shows rendered RGB followed by estimated albedo, metallic, and roughness maps.

4.4 Integration with Image-to-3D Pipelines

We demonstrate our method’s effectiveness in enhancing modern image-to-3D generation pipelines by integrating it with Era3D [19]. As shown in Figure 9, our method first processes Era3D’s multi-view outputs to obtain consistent PBR material properties across different viewpoints (front, left, right, back), which is crucial for high-quality 3D reconstruction. Using these consistent material estimates, we then generate a textured mesh following Era3D’s pipeline with an additional PBR reconstruction loss. The quality of our estimated materials is validated through relighting experiments, as shown

in Figure 10, where the reconstructed mesh maintains physically plausible appearance under various environment illuminations.

Furthermore, we compare our PBR estimation capabilities against other end-to-end image-to-3D methods that directly generate textured meshes with material properties, including SF3D [1] and 3DTopia [3], as shown in Fig. 8. While these methods can generate textured 3D meshes, they often struggle to accurately decompose materials into physically-based properties that enable realistic relighting. In contrast, our method produces more precise PBR estimations that better support downstream rendering applications. Additional results and comparisons across diverse objects are provided in the supplementary material.

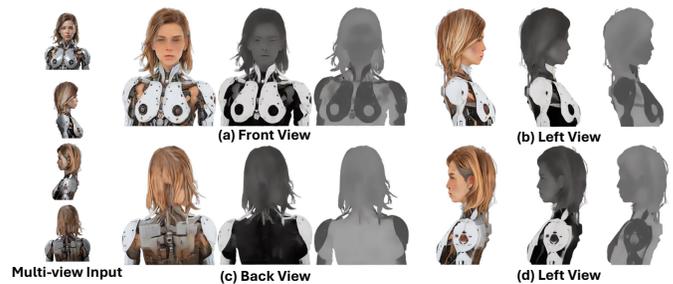


Figure 9: Multi-view material estimation results. Given multi-view input images (left), our method produces consistent PBR decomposition across different viewpoints, showing albedo, metallic, and roughness maps for each view (a-d). Note how material properties remain coherent across viewpoint changes.



Figure 10: Relighting results of our PBR-estimated textured mesh.

5 CONCLUSION

We introduce DualMat, a novel dual-path diffusion framework for PBR material estimation that combines pretrained visual knowledge with specialized material understanding. Operating in complementary latent spaces for albedo and metallic-roughness estimation, our approach achieves state-of-the-art performance in material recovery from single images. Through feature distillation and efficient rectified flow sampling, we enable high-quality generation in just three steps, with extensions to high-resolution and multi-view scenarios. While the dual-path design increases memory cost, future work could explore more efficient architectures and language-guided material editing.

REFERENCES

- [1] Mark Boss, Zixuan Huang, Aaryaman Vasishtha, and Varun Jampani. 2024. SF3D: Stable Fast 3D Mesh Reconstruction with UV-unwrapping and Illumination Disentanglement. *arXiv preprint* (2024).
- [2] Xi Chen, Sida Peng, Dongchen Yang, Yuan Liu, Bowen Pan, Chengfei Lv, and Xiaowei Zhou. 2024. IntrinsicAnything: Learning Diffusion Priors for Inverse Rendering Under Unknown Illumination. *arXiv preprint arXiv:2404.11593* (2024).
- [3] 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 Zifei Liu. 2024. 3DTopia-XL: High-Quality 3D PBR Asset Generation via Primitive Diffusion. *arXiv preprint arXiv:2409.12957* (2024).
- [4] Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. 2018. Single-image svbrdf capture with a rendering-aware deep network. *ACM Transactions on Graphics* (2018).
- [5] Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. 2019. Flexible svbrdf capture with a multi-image deep network. In *Computer graphics forum*.
- [6] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *International Conference on Machine Learning*.
- [7] Sicheng Gao, Xuhui Liu, Bohan Zeng, Sheng Xu, Yanjing Li, Xiaoyan Luo, Jianzhuang Liu, Xiantong Zhen, and Baochang Zhang. 2023. Implicit diffusion models for continuous super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- [8] Jie Guo, Shuichang Lai, Chengzhi Tao, Yuelong Cai, Lei Wang, Yanwen Guo, and Ling-Qi Yan. 2021. Highlight-aware two-stream network for single-image SVBRDF acquisition. *ACM Transactions on Graphics* (2021).
- [9] Yu Guo, Milovs Havsan, Lingqi Yan, and Shuang Zhao. 2020. A bayesian inference framework for procedural material parameter estimation. In *Computer Graphics Forum*.
- [10] Yu Guo, Cameron Smith, Milovs Havsan, Kalyan Sunkavalli, and Shuang Zhao. 2020. MaterialGAN: Reflectance capture using a generative SVBRDF model. *arXiv preprint arXiv:2010.00114* (2020).
- [11] Philipp Henzler, Valentin Deschaintre, Niloy J Mitra, and Tobias Ritschel. 2021. Generative modelling of BRDF textures from flash images. *arXiv preprint arXiv:2102.11861* (2021).
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. In *Advances in neural information processing systems*.
- [13] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. 2022. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research* (2022).
- [14] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. 2023. Lrm: Large reconstruction model for single image to 3d. *arXiv preprint arXiv:2311.04400* (2023).
- [15] Yiwei Hu, Milovs Havsan, Paul Guerrero, Holly Rushmeier, and Valentin Deschaintre. 2022. Controlling material appearance by examples. In *Computer graphics forum*.
- [16] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. 2022. Elucidating the design space of diffusion-based generative models. In *Advances in neural information processing systems*.
- [17] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. 2024. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [18] Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan Sunkavalli, Greg Shakhnarovich, and Sai Bi. 2023. Instant3d: Fast text-to-3d with sparse-view generation and large reconstruction model. *arXiv preprint arXiv:2311.06214* (2023).
- [19] Peng Li, Yuan Liu, Xiaoxiao Long, Feihu Zhang, Cheng Lin, Mengfei Li, Xingqun Qi, Shanghang Zhang, Wenhan Luo, Ping Tan, et al. 2024. Era3D: High-Resolution Multiview Diffusion using Efficient Row-wise Attention. *arXiv preprint arXiv:2405.11616* (2024).
- [20] Xiao Li, Yue Dong, Pieter Peers, and Xin Tong. 2017. Modeling surface appearance from a single photograph using self-augmented convolutional neural networks. *ACM Transactions on Graphics* (2017).
- [21] Zhong Li, Liangchen Song, Zhang Chen, Xiangyu Du, Lele Chen, Junsong Yuan, and Yi Xu. 2023. Relit-neufl: Efficient relighting and novel view synthesis via neural 4d light field. In *Proceedings of the 31st ACM International Conference on Multimedia*.
- [22] Zhengqin Li, Kalyan Sunkavalli, and Manmohan Chandraker. 2018. Materials for masses: SVBRDF acquisition with a single mobile phone image. In *Proceedings of the European conference on computer vision*.
- [23] Aoming Liu, Zhong Li, Zhang Chen, Nannan Li, Yi Xu, and Bryan A Plummer. 2024. Panofree: Tuning-free holistic multi-view image generation with cross-view self-guidance. In *European Conference on Computer Vision*.
- [24] Isabella Liu, Linghao Chen, Ziyang Fu, Liwen Wu, Haian Jin, Zhong Li, Chin Ming Ryan Wong, Yi Xu, Ravi Ramamoorthi, Zexiang Xu, et al. 2023. Openillumination: A multi-illumination dataset for inverse rendering evaluation on real objects. *Advances in Neural Information Processing Systems* (2023).
- [25] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. 2023. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF international conference on computer vision*.
- [26] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. 2023. Syncdreamer: Generating multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453* (2023).
- [27] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. 2022. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- [28] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobald, et al. 2024. Wonder3d: Single image to 3d using cross-domain diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [29] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM* (2021).
- [30] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952* (2023).
- [31] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. 2022. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988* (2022).
- [32] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- [33] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. 2022. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH*.
- [34] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. 2022. Image super-resolution via iterative refinement. *IEEE transactions on pattern analysis and machine intelligence* (2022).
- [35] Shen Sang and Manmohan Chandraker. 2020. Single-shot neural relighting and svbrdf estimation. In *Proceedings of the European conference on computer vision*.
- [36] Sam Sartor and Pieter Peers. 2023. Matfusion: a generative diffusion model for svbrdf capture. In *SIGGRAPH Asia 2023 Conference Papers*.
- [37] Hiroshi Sasaki, Chris G Willcocks, and Toby P Breckon. 2021. Unit-ddpm: Unpaired image translation with denoising diffusion probabilistic models. *arXiv preprint arXiv:2104.05358* (2021).
- [38] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. 2023. Zero123++: a single image to consistent multi-view diffusion base model. *arXiv preprint arXiv:2310.15110* (2023).
- [39] Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. 2023. Mvdream: Multi-view diffusion for 3d generation. *arXiv preprint arXiv:2308.16512* (2023).
- [40] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2021. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*.
- [41] Xuan Su, Jiamei Song, Chenlin Meng, and Stefano Ermon. 2022. Dual diffusion implicit bridges for image-to-image translation. *arXiv preprint arXiv:2203.08382* (2022).
- [42] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. 2023. Make-it-3d: High-fidelity 3d creation from a single image with diffusion prior. In *Proceedings of the IEEE/CVF international conference on computer vision*.
- [43] Shimon Vainer, Mark Boss, Mathias Parger, Konstantin Kutsy, Dante De Nigris, Ciara Rowles, Nicolas Perony, and Simon Donné. 2024. Collaborative Control for Geometry-Conditioned PBR Image Generation. *arXiv preprint arXiv:2402.05919* (2024).
- [44] Giuseppe Vecchio. 2024. StableMaterials: Enhancing Diversity in Material Generation via Semi-Supervised Learning. *arXiv preprint arXiv:2406.09293* (2024).
- [45] Giuseppe Vecchio, Rosalie Martin, Arthur Roullier, Adrien Kaiser, Romain Rouffet, Valentin Deschaintre, and Tamy Boubekeur. 2024. Controlmat: A controlled generative approach to material capture. *ACM Transactions on Graphics* (2024).
- [46] Giuseppe Vecchio, Simone Palazzo, and Concetto Spampinato. 2021. Surfacenet: Adversarial svbrdf estimation from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- [47] Giuseppe Vecchio, Renato Sortino, Simone Palazzo, and Concetto Spampinato. 2024. Matfuse: controllable material generation with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [48] Peng Wang and Yichun Shi. 2023. Imagedream: Image-prompt multi-view diffusion for 3d generation. *arXiv preprint arXiv:2312.02201* (2023).

- [49] Peng Wang, Hao Tan, Sai Bi, Yinghao Xu, Fujun Luan, Kalyan Sunkavalli, Wenping Wang, Zexiang Xu, and Kai Zhang. 2023. Pf-lrm: Pose-free large reconstruction model for joint pose and shape prediction. *arXiv preprint arXiv:2311.12024* (2023).
- [50] Zhengyi Wang, Yikai Wang, Yifei Chen, Chendong Xiang, Shuo Chen, Dajiang Yu, Chongxuan Li, Hang Su, and Jun Zhu. 2024. Crm: Single image to 3d textured mesh with convolutional reconstruction model. *arXiv preprint arXiv:2403.05034* (2024).
- [51] Tao Wen, Beibei Wang, Lei Zhang, Jie Guo, and Nicolas Holzschuch. 2022. SVBRDF Recovery from a Single Image with Highlights Using a Pre-trained Generative Adversarial Network. In *Computer Graphics Forum*.
- [52] Zhexiao Xiong, Zhang Chen, Zhong Li, Yi Xu, and Nathan Jacobs. 2025. PanoDreamer: Consistent Text to 360-Degree Scene Generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*.
- [53] Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. 2024. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191* (2024).
- [54] Yinghao Xu, Zifan Shi, Wang Yifan, Hansheng Chen, Ceyuan Yang, Sida Peng, Yujun Shen, and Gordon Wetzstein. 2024. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation. *arXiv preprint arXiv:2403.14621* (2024).
- [55] Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli, Gordon Wetzstein, Zexiang Xu, et al. 2023. Dmv3d: Denoising multi-view diffusion using 3d large reconstruction model. *arXiv preprint arXiv:2311.09217* (2023).
- [56] Hu Ye, Jun Zhang, Sibao Liu, Xiao Han, and Wei Yang. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721* (2023).
- [57] Wenjie Ye, Yue Dong, Pieter Peers, and Baining Guo. 2021. Deep Reflectance Scanning: Recovering Spatially-varying Material Appearance from a Flash-lit Video Sequence. In *Computer Graphics Forum*.
- [58] Wenjie Ye, Xiao Li, Yue Dong, Pieter Peers, and Xin Tong. 2018. Single image surface appearance modeling with self-augmented cnns and inexact supervision. In *Computer Graphics Forum*.
- [59] Zongsheng Yue, Jianyi Wang, and Chen Change Loy. 2024. Resshift: Efficient diffusion model for image super-resolution by residual shifting. In *Advances in Neural Information Processing Systems*.
- [60] Min Zhao, Fan Bao, Chongxuan Li, and Jun Zhu. 2022. Egsde: Unpaired image-to-image translation via energy-guided stochastic differential equations. *Advances in Neural Information Processing Systems* (2022).
- [61] Xilong Zhou, Milos Hasan, Valentin Deschaintre, Paul Guerrero, Kalyan Sunkavalli, and Nima Khademi Kalantari. 2022. Tilegen: Tileable, controllable material generation and capture. In *SIGGRAPH Asia*.
- [62] Xilong Zhou and Nima Khademi Kalantari. 2021. Adversarial Single-Image SVBRDF Estimation with Hybrid Training. In *Computer Graphics Forum*.