

# ARM: Appearance Reconstruction Model for Relightable 3D Generation

Xiang Feng<sup>1,2\*</sup> Chang Yu<sup>3\*</sup> Zoubin Bi<sup>2\*</sup> Yintong Shang<sup>1</sup> Feng Gao<sup>4†</sup>  
 Hongzhi Wu<sup>2</sup> Kun Zhou<sup>2</sup> Chenfanfu Jiang<sup>3</sup> Yin Yang<sup>1</sup>  
<sup>1</sup>University of Utah <sup>2</sup>Zhejiang University <sup>3</sup>UCLA <sup>4</sup>Amazon



Figure 1. ARM generates high-quality, relightable 3D content from a single image input. This figure presents sample results generated from different input images, demonstrating ARM’s ability to reconstruct a variety of objects with spatially-varying appearance. Please refer to our supplementary video for results under dynamic view and lighting.

## Abstract

Recent image-to-3D reconstruction models have greatly advanced geometry generation, but they still struggle to faithfully generate realistic appearance. To address this, we introduce ARM, a novel method that reconstructs high-quality 3D meshes and realistic appearance from sparse-view images. The core of ARM lies in decoupling geometry from appearance, processing appearance within the UV texture space. Unlike previous methods, ARM improves texture quality by explicitly back-projecting measurements onto the texture map and processing them in a UV space module with a global receptive field. To resolve ambiguities between material and illumination in input images, ARM introduces a material prior that encodes semantic appearance information, enhancing the robustness of appearance decomposition. Trained on

just 8 H100 GPUs, ARM outperforms existing methods both quantitatively and qualitatively. Our project page is available at <https://arm-aigc.github.io>.

## 1. Introduction

Obtaining high-quality 3D mesh models with realistic appearance from 2D images has become a critical task in computer vision and graphics, with applications spanning the metaverse, gaming, and e-commerce. Conventionally, 3D models with appearance are hand-crafted by skilled artists using specialized modeling software—a highly time-consuming process that can take hours or even days. Alternatively, 3D meshes can be reconstructed from multi-view input images using optimization-based approaches [26, 44, 45, 72], typically requiring over a hundred images from different views. Photometric devices are often needed to accurately capture high-quality appearance [25, 42, 47, 52, 56].

\* indicates equal contributions.

† This work is not related to F. Gao’s position at Amazon.

Recent advances in 3D generation and reconstruction models have brought new insights to the field, leading to several dominant approaches. One line of research [9, 31, 36, 49, 60, 74] leverages priors from pretrained 2D diffusion models to distill 3D shapes, typically by generating images from multiple viewpoints with a 2D generative model, followed by per-scene optimization. However, this distillation process is time-consuming, limiting its practicality for real-world applications. Another line [21, 65, 66, 77, 81] trains a feed-forward neural network directly on large-scale 3D datasets [12, 13] to learn 3D priors from single-view or sparse-view input images, an approach our method also adopts. This strategy can improve consistency across multi-views and achieves faster inference compared to time-intensive distillation techniques. However, existing methods in this area still struggle with several limitations: the reconstructed textures often appear blurred and lack fine details, leading to overall low-quality results. Moreover, most current methods represent object appearance only through per-vertex colors, which is even simpler than Lambertian shading and includes baked-in reflections, lighting, and shadows. This simplified shading model fails to capture realistic view-dependent and lighting-dependent effects, making the generated assets unsuitable for downstream applications such as gaming or metaverse, where dynamic lighting and viewpoint changes are essential for realism.

In this work, we introduce a framework called ARM for reconstructing high-quality 3D meshes with fine-detailed textures and realistic appearance. ARM builds on Large Reconstruction Models (LRMs) [21], using triplanes as its 3D representation. While LRMs offer strong geometry capabilities, we observed that reconstructed textures often appear overly blurred due to the limited resolution of triplanes and the relatively small decoding MLPs. ARM’s core innovation lies in decoupling geometry generation from appearance modeling by processing appearance directly within the UV space. Unlike previous methods that decode color from learned triplanes using MLPs, ARM enhances texture quality by explicitly back-projecting multi-view measurements onto the texture map and processing them with a UV-space module featuring a global receptive field. The UV texture space offers advantages over the triplane space by directly representing color variations on the object surface. ARM also introduces an approach to address the material and illumination ambiguities present in sparse-view input images—a fundamentally ill-posed problem. Previous methods [61] often attempt to tackle this issue using a rendering loss; however, these inverse rendering approaches have been shown to struggle with sparse inputs and may even fail with dense multi-view data. In contrast, ARM incorporates a material prior that encodes semantic appearance information and directly fits to ground-truth materials, enhancing the robustness of appearance decompo-

sition. Experimental results show that ARM outperforms recent image-to-3D methods both qualitatively and quantitatively, demonstrating its capability to generate versatile objects with realistic appearance.

## 2. Related work

### 2.1. 3D generation with 2D diffusion priors

The emergence of 2D diffusion models [20, 54, 55] has driven significant advancements in 3D generation. DreamFusion [49] was the first to use SDS [70] loss to iteratively distill 3D representations from a 2D diffusion model, inspiring numerous follow-up works [9, 31, 40, 41, 43, 57, 60, 63, 71, 78]. These methods typically convert multi-view images generated by diffusion models into 3D representations such as NeRF [44], NeuS [72], and 3DGS [26], using SDS-like losses and optimization-based techniques. Subsequent work has focused on enhancing cross-view consistency and generalization. For example, Zero-1-to-3 [36] leverages diffusion models to create consistent multi-view images for 3D reconstruction, while other methods [8, 22, 32, 37, 38, 51, 59, 85] further improve consistency through conditioning and distribution modeling. Wang et al. [74] proposed Variational Score Distillation (VSD) to address over-saturation in SDS and improve texture diversity. Recently, Voleti et al. [68] introduced an image-to-video diffusion model for enhanced generalization and multi-view consistency in novel view synthesis. Despite these advancements, per-shape optimization methods still face key limitations, with long runtimes as the main bottleneck to broader application. Issues with Janus inconsistencies and over-saturation also remain.

### 2.2. 3D generation with 3D feed-forward models

With the availability of high-quality, large-scale annotated 3D datasets [12, 13], recent works have focused on generating 3D objects using feed-forward models [21, 24, 53, 80], achieving significantly faster speeds than score-distillation methods. For example, One-2-3-45 [34] uses the generalizable NeRF method for 3D reconstruction, while One-2-3-45++ [33] improves quality with a two-stage, coarse-to-fine 3D diffusion model similar to LAS-Diffusion [94], though high-quality textures remain a challenge. The LRM model [21] and its extensions [29, 66, 73, 75, 77, 81, 83] enhance generation quality through transformer-based architectures and triplane representation [7]. While triplanes provide an effective hybrid representation, other works [65, 84, 90, 95] explore 3DGS [26] for potentially faster generation. Recently, MeshFormer [35] replaced triplanes with 3D sparse voxels to represent fine-grained shapes explicitly. Our method also adopts a feed-forward approach for efficient generation. Using LRM as the backbone, we address a key limitation: relying solely on triplanes restricts texture

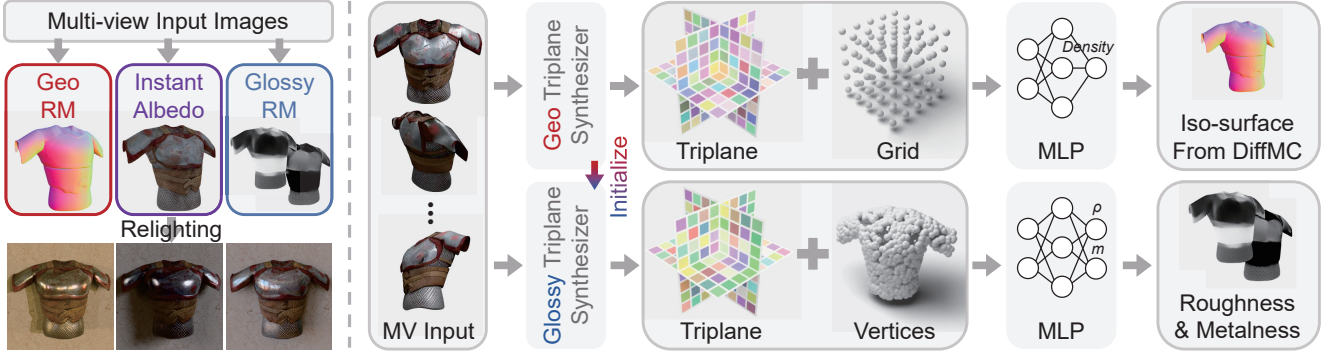


Figure 2. **Overview of our pipeline.** (left) Starting from sparse-view input images generated by a diffusion model [59], ARM separates shape and appearance generation into two stages. In the geometry stage, ARM uses GeoRM to predict a 3D shape from the input images. In the appearance stage, ARM employs InstantAlbedo and GlossyRM to reconstruct PBR maps, enabling realistic relighting under varied lighting conditions. (right) Both GeoRM and GlossyRM share the same architecture, consisting of a triplane synthesizer and a decoding MLP. GeoRM is trained to predict density and extracts an iso-surface from the density grid with DiffMC [76], while GlossyRM is trained to predict roughness and metalness. GlossyRM is trained after GeoRM and initializes with the weights of GeoRM at the start of training.

detail due to resolution constraints, and scaling up the network is impractical due to memory demands. Instead, we directly learn textures in UV space, significantly enhancing texture quality.

### 2.3. Material generation and decomposition

Estimating surface material properties remains a longstanding challenge in 3D reconstruction and generation, as joint optimization of unknown lighting and appearance makes this problem inherently ill-posed. Recent advances in multi-view reconstruction, including NeRF-based methods [2, 3, 15, 18, 46, 62, 87, 89, 92] and 3DGS-based approaches [1, 17, 23, 30], show promise for material estimation but still require dense multi-view input. SDS-based text-to-3D pipelines [9, 39, 82] offer a solution but remain time-intensive due to SDS optimization. Make-it-Real [16] uses LLMs to identify object semantics and retrieve materials from a library, while SF3D [4] adds UV unwrapping with material prediction in LRM. Other methods [67, 86, 88, 93] apply diffusion models for appearance generation on existing geometries with text prompts. Recently, CLAY [91] and 3DTopia-XL [10] leverage Diffusion Transformer models (DiT) [48] for 3D asset generation with appearance. Our method differs by being the first to perform material decomposition entirely in UV space, conditioned directly on in-the-wild image input. Since our network inherently learns texture-level priors through back-projection in the first step, it mitigates the inconsistent artifacts caused by generating multi-view appearance parameters first and then back-projecting them onto textures. Additionally, by introducing a material prior in texture-level appearance decomposition, our method more accurately disentangles illumination from materials, achieving greater fidelity than previous approaches.

### 3. Preliminaries

We model the spatially varying, view-dependent, and lighting-dependent appearance of objects by an Spatially Varying Bi-directional Reflectance Distribution Function (SVBRDF), where each surface point’s reflectance is modeled by a microfacet BRDF parameterized by its diffuse albedo  $c_d$ , roughness  $\rho$ , and metalness  $m$ . The BRDF  $f_r$  consists of a diffuse component and a glossy component:

$$f_r(\mathbf{l}, \mathbf{v}) = (1 - m) \frac{c_d}{\pi} + \frac{DFG}{4(\mathbf{n} \cdot \mathbf{l})(\mathbf{n} \cdot \mathbf{v})} \quad (1)$$

where  $\mathbf{l}$  and  $\mathbf{v}$  are the directions of the incoming light and the view,  $\mathbf{n}$  is the surface normal, and  $D$ ,  $F$ , and  $G$  represent the microfacet normal distribution, Fresnel, and geometry terms, respectively. These terms are determined by  $\rho$ ,  $m$ , and  $c_d$ , with detailed expressions provided in the supplementary material. In summary, ARM reconstructs spatially varying diffuse, roughness, and metalness maps—also known as Physically-Based Rendering (PBR) maps [5]—to efficiently capture diverse material appearances.

### 4. Overview

We propose ARM, a framework for simultaneously reconstructing high-quality 3D meshes and PBR texture maps, as illustrated in Fig. 2. Starting from sparse-view images generated from a single view using a diffusion model [59], ARM separates shape and appearance generation into two stages. In the geometry stage, ARM employs *GeoRM* (Sec. 5.1) to predict a 3D shape from input images, which is then unwrapped into atlas charts for latter processing in the UV texture space. In the appearance stage, ARM introduces two specialized models: *InstantAlbedo* (Sec. 5.2) and *GlossyRM* (Sec. 5.1). Both models take input as foremen-

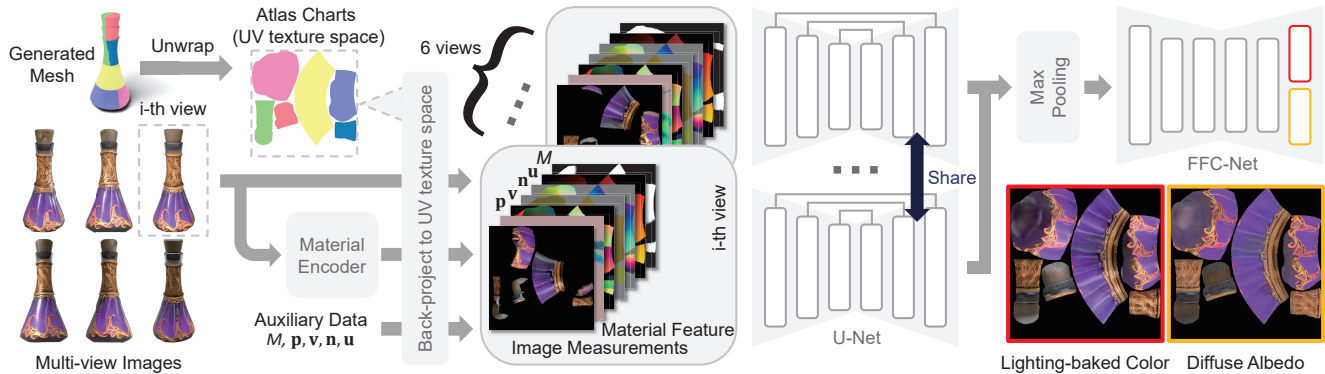


Figure 3. **Overview of InstantAlbedo.** InstantAlbedo operates in the texture UV space. This process begins by converting all necessary data to UV texture space. Given the unwrapped mesh from GeoRM, we back-project images, material encodings, and auxiliary data into UV texture space, resulting in six sets of inputs corresponding to the six input views. InstantAlbedo then processes these maps into a U-Net and an inpainting-specific FFC-Net to predict both the lighting-baked color and the decomposed diffuse albedo UV textures.

tioned sparse-view input images and the predicted 3D shape from GeoRM. GlossyRM outputs per-vertex roughness  $\rho$  and metalness  $m$ , capturing the glossy component of the appearance. InstantAlbedo, in contrast, operates in the UV texture space, producing two color maps: one with baked-in lighting and another with diffuse albedo, representing the diffuse component of appearance. To enhance the accuracy and robustness of albedo decomposition, ARM also incorporates a material prior (Sec. 5.3) into InstantAlbedo to encode semantic appearance information from the multi-view images. The final material is stored as PBR texture maps, enabling realistic relighting under novel lighting.

## 5. Method

In this section, we present the design of ARM and outline the key insights and motivations behind our approach.

### 5.1. Decoupling shape and appearance

To produce high-quality 3D meshes with high-fidelity reflective properties, ARM introduces two distinct LRM-based models: GeoRM and GlossyRM. Both models share the same architecture, including a transformer-based triplane synthesizer and a decoding MLP. The triplane, as an internal 3D representation, is generated by the synthesizer from sparse-view input images. Any queried 3D point is then projected onto the triplane to retrieve its corresponding feature, which is decoded by the MLP to produce the desired output. By decoupling shape and appearance, ARM designs the two models to serve different purposes, each trained with its own objective function.

GeoRM is trained to predict the density, with the isosurface extracted from a  $256^3$  density grid using differentiable marching cubes [76]. Optimized solely with geometry-related losses (mask, depth, and normal supervision), GeoRM’s purpose is exclusively to generate geometry, and its weights are frozen once trained.

GlossyRM, trained after GeoRM, uses the 3D mesh generated by GeoRM to query its own learned triplane and decode per-vertex roughness  $\rho$  and metalness  $m$ . Its training is focused solely on losses related to these parameters.

The core idea here is to decouple geometry from appearance, enabling (1) efficient processing in texture space, since unwrapping geometry into texture space is time-consuming and impractical during training, and (2) increased network capacity for both GeoRM and GlossyRM, with each model focused on a specific task. A naive approach would be to train a single LRM to predict all desired outputs—density, vertex color, and material parameters—but our pilot study (Fig. 7) showed that as the number of prediction targets increases, output quality degrades significantly, with results becoming blurred. This issue is particularly pronounced with material predictions, as these parameters are more challenging to infer than colors directly visible in the input images, increasing learning complexity.

Furthermore, we observe that the triplane resolution is directly correlated with shape quality and the presence of voxel artifacts, which can be mitigated by increasing the resolution. To address this, we introduce a super-resolution module that raises triplane resolution to  $256 \times 256$ , significantly enhancing shape reconstruction quality. By decoupling geometry and appearance, the memory footprint of both models is reduced, making this resolution increase feasible. Further details on the architecture of GeoRM and GlossyRM can be found in our supplementary material.

### 5.2. Appearance decomposition in texture space

The key insight of ARM is to reconstruct fine-detailed textures in UV texture space, which directly aligns with the spatial variations on the object surface, rather than relying on the triplane representation used in previous LRMs. There are two main reasons for moving away from triplanes: First, texture details are limited by the triplane res-

olution, which is constrained by memory usage of large transformer. Increasing triplane resolution to capture finer details quadratically increases the transformer complexity, reducing practicality for high-detail reconstructions. Second, triplanes function similarly to volumetric representations, where querying with decoding MLPs often results in blurriness. This is because triplanes store 3D information across three planes, but the spatial variation in these planes does not directly correspond to the texture variation on the object surface, leading to interpolation mismatches and degraded texture quality.

To overcome these limitations, ARM employs InstantAlbedo, a network specifically designed to operate in UV space, as illustrated in Fig. 3. By working directly in texture space, InstantAlbedo captures fine details without the resolution and memory constraints associated with triplanes.

The process begins by converting all necessary data to UV texture space. Given the 3D mesh generated by GeoRM, we unwrap it into atlas charts in the UV texture, where each texel corresponds to a specific surface point. Since unwrapping is time-consuming and impractical to perform during training, we synthesize a pre-unwrapped training dataset offline, which we detail later. Next, with multi-view input images generated at 6 known camera poses, we directly back-project these multi-view images onto the UV texture. Along with the images, we also back-project auxiliary data, including masks  $M$ , positions  $\mathbf{p}$ , texture coordinates  $\mathbf{u}$ , view directions  $\mathbf{v}$ , and surface normals  $\mathbf{n}$ . These inputs provide valid contextual information for material decomposition according to Eq. 1, effectively reducing the learning complexity for the network. The back-projection process generates six sets of input maps in UV texture space, corresponding to the six views.

InstantAlbedo takes above six sets of maps as input and outputs two color maps: one with baked-in lighting and another with the decomposed diffuse albedo. In the first step, a U-Net is employed to extract per-view features from each set, yielding six feature maps. Since each view captures only a partial region of the object’s surface, it is necessary to fuse information from multiple views to form a complete texture. Inspired by PointNet [50], we use max-pooling to aggregate information across the six feature maps. However, with only six views, some regions of the object surface remain unobserved. To address this, we incorporate a FFCNet [11, 64] with a global receptive field to extract information from other areas in the texture, inpaint unseen regions, and refine the fused result. As shown in Fig. 6, this design significantly improves the completeness of the reconstructed texture. Please note that InstantAlbedo focuses solely on albedo decomposition rather than predicting the entire material (including roughness and metalness) in a unified manner, as we found that this approach led to inaccuracies in decomposed roughness and metalness. Ad-

ditional details are provided in the supplementary material.

### 5.3. Appearance encoding with a material prior

ARM tackles the inherent ambiguity between material and illumination in sparse-view input, another challenging and inherently ill-posed problem. With fewer than ten input images, attempting to use inverse rendering with a simple rendering loss often results in flawed decompositions for glossy materials, where lighting effects may bake into the diffuse albedo to perfectly match the input images. To address this, ARM directly fits to ground-truth material instead of relying on a rendering loss to separate material and lighting. However, this approach is still challenging because any lighting in the input images can be hard for the network to ignore, often leaving traces of illumination in the final albedo output.

To improve robustness, ARM incorporates a material prior—a material-aware image encoder—into the back-projection process of InstantAlbedo, as illustrated in Fig. 3. By transforming multi-view input images into semantic, material-aware feature maps, InstantAlbedo back-projects these encoded features onto the UV texture. Combined with other auxiliary inputs, these features allow InstantAlbedo to produce a more accurate, decomposed result.

The image encoder, based on DINO’s ViT  $8 \times 8$  configuration [6], combines intermediate features at multiple scales into a unified feature map through upscaling convolutional networks, similar to the cascade architecture in [58]. We initialize the encoder with weights trained on a dataset with semantic material maps, making it suitable for recognizing materials. Integrated into the InstantAlbedo pipeline, the encoder is fine-tuned jointly with the rest of the model. As shown in Fig. 6, This design significantly enhances the accuracy of albedo decomposition, even when strong lighting and materials are tightly coupled in the input.

## 6. Training

To train GeoRM, GlossyRM, and InstantAlbedo, we synthesize two separate datasets: one for GeoRM and GlossyRM, and another for InstantAlbedo.

**GeoRM and GlossyRM** are trained on a 150K subset of the Objaverse dataset [12]. For each object, we render it from 32 random views, generating measurements, depth, normal, diffuse albedo, roughness, metalness, and mask maps for supervision. We start by training GeoRM, with a two-stage strategy similar to [81] (see supplementary material for more details). GeoRM is trained based on differentiable rendering, exclusively with geometry-related losses:

$$\mathcal{L}_{\text{geo}} = \frac{1}{N} \sum_i \lambda_z |z_i^{\text{gt}} - \hat{z}_i| + \lambda_M \mathcal{L}_{\text{mse}}(M_i^{\text{gt}}, \hat{M}_i) + \lambda_n \mathcal{L}_{\text{lpips}}(\mathbf{n}_i^{\text{gt}}, \hat{\mathbf{n}}_i), \quad (2)$$

where  $z_i^{\text{gt}}$ ,  $M_i^{\text{gt}}$ , and  $\mathbf{n}_i^{\text{gt}}$  are the ground-truth depth, mask, and surface normal for the  $i$ -th view, and  $\hat{z}_i$ ,  $\hat{M}_i$ , and  $\hat{\mathbf{n}}_i$  are their counterparts rendered from predicted mesh. We randomly select  $N$  views from the dataset, with 6 used as input. Next, we train GlossyRM with GeoRM fixed to provide a proxy shape, using the following loss:

$$\mathcal{L}_{\text{glossy}} = \frac{1}{N} \sum_i \mathcal{L}_0(\rho_i^{\text{gt}}, \hat{\rho}_i) + \mathcal{L}_0(m_i^{\text{gt}}, \hat{m}_i), \quad (3)$$

$$\mathcal{L}_0(x, y) = \lambda_1 \mathcal{L}_{\text{mse}}(x, y) + \lambda_2 \mathcal{L}_{\text{lpiPs}}(x, y) + \lambda_3 \mathcal{L}_{\text{ssim}}(x, y).$$

Here,  $\rho_i^{\text{gt}}$  and  $m_i^{\text{gt}}$  represent the ground-truth roughness and metalness for the  $i$ -th view, and  $\hat{\rho}_i$  and  $\hat{m}_i$  are the predicted roughness and metalness rendered from the predicted mesh.

**InstantAlbedo** is also trained after GeoRM, but independently from GlossyRM, allowing both models to be trained in parallel. Once GeoRM is trained, we use it to generate 55K shapes from the 150K subset mentioned above, each pre-unwrapped into atlas charts. During training, InstantAlbedo takes the unwrapped mesh and corresponding input images as input, and is trained to predict both the lighting-baked color and the decomposed diffuse albedo. The overall loss is defined as:

$$\mathcal{L}_{\text{albedo}} = \frac{1}{N} \sum_i \mathcal{L}_0(\mathbf{c}_i^{\text{gt}}, \hat{\mathbf{c}}_i) + \mathcal{L}_0(\mathbf{c}_{d_i}^{\text{gt}}, \hat{\mathbf{c}}_{d_i}), \quad (4)$$

where  $\mathbf{c}_i^{\text{gt}}$  and  $\mathbf{c}_{d_i}^{\text{gt}}$  denote the ground-truth lighting-baked color and decomposed diffuse albedo for the  $i$ -th view, and  $\hat{\mathbf{c}}_i$  and  $\hat{\mathbf{c}}_{d_i}$  are their predicted counterparts.

## 7. Experiments

### 7.1. Experimental settings

We trained ARM on 8 H100 GPUs with a batch size of 1 per GPU over approximately 5 days: two days for GeoRM, two for GlossyRM, and one for InstantAlbedo. We used the Adam optimizer with a learning rate of  $4 \times 10^{-5}$  for GeoRM and GlossyRM, and  $4 \times 10^{-4}$  for InstantAlbedo. The loss-balancing coefficients were set as  $\lambda_z = 0.5$ ,  $\lambda_M = 1.0$ ,  $\lambda_n = 0.2$ ,  $\lambda_1 = 0.7$ ,  $\lambda_2 = 0.3$ , and  $\lambda_3 = 0.1$ . We set  $N = 10$  when selecting views.

### 7.2. Evaluation settings

We assess the methods on three datasets: GSO [14], OmniObject3D [79], and a custom dataset specifically for relightable appearance evaluation. All datasets consist of 3D objects that were unseen during training. For the GSO dataset, we evaluate using all 1,030 available 3D shapes. From the OmniObject3D dataset, we randomly sample up to five shapes from each category, totaling 1,038 shapes for evaluation. The third dataset, introduced for relightable appearance evaluation, includes 100 objects with PBR materials, similar to [88]. For each object, we generate 144 images

under varied lighting conditions by rendering 24 random views across six different environmental lightings. To evaluate 3D reconstruction quality, we use both F-score (with a threshold of 0.1) and Chamfer Distance (CD), comparing the predicted meshes to ground truth meshes. For 2D appearance evaluation, we compute PSNR, SSIM, and LPIPS on rendered images. Since coordinate frames may differ across methods, we align each method’s predicted mesh to the ground truth before calculating metrics. Further alignment details are provided in the supplementary materials.

### 7.3. Comparisons

We first compare our method quantitatively to LGM [65], CRM [75], InstantMesh [81], MeshFormer [35], and SF3D [4] in Tab. 1. For a fair comparison, all methods are evaluated in a unified single-view to 3D setup. We use the first thumbnail image as input for the GSO dataset, while for the OmniObject3D dataset, we use a rendered image from a random view. For InstantMesh, MeshFormer, and our method, we employ the same Zero123++ [59] model to generate multi-view images from the single-view input. Since most methods only produce appearance with baked-in lighting, we compare based on lighting-baked color. For SF3D, which lacks lighting-baked vertex color, we render it under even environmental lighting. Our method outperforms others in both geometry and appearance accuracy.

In Fig. 4, we provide qualitative examples to visually demonstrate ARM’s superior performance over existing methods. For full results including CRM and LGM, please see our supplementary material. The reconstructed textures from ARM contain significantly richer details, owing to our design in UV texture space. While other methods suffer from blurriness, ARM accurately reconstructs complex and sharp patterns. Some methods, such as SF3D, struggle to generate plausible shape and texture in unseen areas due to training on single-view inputs.

In Fig. 5, we compare our reconstructed PBR maps and their relighted images under novel lighting conditions to those produced by SF3D [4], which also reconstructs PBR from single-view input. Our method outperforms SF3D in two key areas: First, when multiple materials are present in the input image, our method reconstructs spatially-varying roughness and metalness, while SF3D generates only constant values, resulting in a homogeneous appearance. Second, SF3D struggles with separating illumination from material properties in the input, leading to baked-in lighting effects. In the cup and ball example, lighting artifacts are embedded in SF3D’s reconstructed diffuse albedo, resulting in inaccurate relighting under novel conditions. In contrast, our method successfully decomposes illumination and material, yielding realistic results.

Finally, we present a quantitative comparison of our method and SF3D on the synthetic relighting dataset in



Figure 4. **Qualitative comparison.** We present examples of single-image 3D generation across different methods. While other methods exhibit blurriness, ARM reconstructs complex patterns with sharp details. Please zoom in to examine the texture quality. Full results, including comparisons with LGM [65] and CRM [75], are provided in the supplementary material.

Method	GSO [14]					OmniObject3D [79]				
	F-Score $\uparrow$	CD $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	F-Score $\uparrow$	CD $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
LGM [65]	0.784	0.132	18.173	0.848	0.207	0.801	0.127	17.979	0.843	0.229
CRM [75]	0.893	0.091	19.390	0.857	0.180	0.845	0.110	19.083	<b>0.852</b>	0.200
InstantMesh [81]	0.938	0.065	19.744	0.858	0.146	0.877	0.094	19.193	0.840	0.187
SF3D [4]	0.888	0.089	18.540	0.848	0.175	0.857	0.105	18.529	0.839	0.195
MeshFormer [35]	0.966	0.052	20.500	0.867	0.141	0.927	0.072	19.402	0.839	0.183
Ours	<b>0.968</b>	<b>0.049</b>	<b>21.692</b>	<b>0.880</b>	<b>0.137</b>	<b>0.936</b>	<b>0.067</b>	<b>20.874</b>	<b>0.850</b>	<b>0.165</b>

Table 1. **Quantitative results of single image to 3D.** We evaluate on the GSO [14] (1,030 shapes) and OmniObject3D [79] (1,038 shapes) datasets, reporting various metrics. CD denotes Chamfer Distance.

Tab. 2. Our method surpasses SF3D in relit rendering, demonstrating superior accuracy and robustness.

#### 7.4. Ablations

We perform both visual and quantitative ablation experiments on the synthetic dataset to evaluate the impact of components in our method. To exclude the effect of multi-

Setting	PSNR-A $\uparrow$	LPIPS-A $\downarrow$
SF3D [4]	18.592	0.200
Ours	<b>21.750</b>	<b>0.171</b>

Table 2. **Quantitative results of relighted appearance.** We evaluate relighted renderings under novel lighting conditions. -A denotes appearance.

view diffusion model, we use ground-truth multi-view images as input. The setups are as follows:

*w/o measurements* Back-projecting measurements directly visible in the input image provides crucial information for reconstructing textures. In this setup, we omit back-projected image measurements on texture maps.

*w/o material prior* Incorporating material-aware image en-

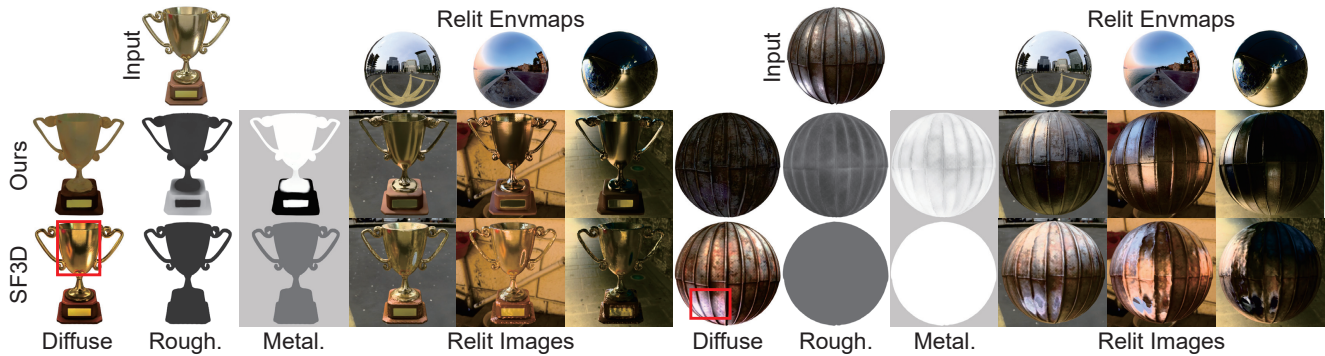


Figure 5. **PBR comparison.** We compare reconstructed PBR maps and relit images under novel lighting to SF3D [4]. While SF3D produces constant roughness and material with lighting baked into the diffuse color (highlighted in the figure), our method generates spatially-varying appearance, with well-separated illumination and materials. See supplementary material for full results.

Setting	PSNR-A $\uparrow$	LPIPS-A $\downarrow$	PSNR-D $\uparrow$	LPIPS-D $\downarrow$
wo/ Measurements	24.780	0.104	23.398	0.114
wo/ Material Prior	24.471	0.108	22.687	0.121
wo/ FFC-Net	24.612	0.110	23.360	0.123
Baseline	<b>25.074</b>	<b>0.096</b>	<b>24.116</b>	<b>0.098</b>

Table 3. **Ablation study.** We evaluate three alternative setups against our full method. -A indicates relighted appearance under novel lighting, and -D denotes predicted diffuse albedo.



Figure 6. **Ablation results for FFC-Net and material prior.** FFC-Net aids in inpainting unseen regions (**top**), while the material prior improves the diffuse decomposition (**bottom**), both for the InstantAlbedo stage.



Figure 7. **Quality degradation.** As the number of predicted targets increases (from  $c$  to  $c/c_d/\rho/m$ ), the quality deteriorates and the results become blurred, highlighting the necessity of our decoupling strategy.

coding adds semantic appearance information, helping to disentangle illumination from materials and improve the accuracy and robustness of appearance decomposition. In this setup, we exclude material features from back-projection.

*w/o FFC-Net* Since the multi-view images do not cover the full object surface, it is essential to inpaint unseen regions by extracting information from other regions. In this setup, we replace the FFC-Net with a U-Net that has a local receptive field, in contrast to the global scope of the FFC-Net. For more examples on FFC-Net and material prior, please refer to our supplementary material.

As shown in Tab. 3, both modifications degrade quality. Removing the material prior leads to a significant drop in the metrics for decomposed diffuse albedo, while removing the FFC-Net introduces artifacts in unseen areas, reducing perceptual quality notably. Fig. 6 further illustrates the visual impact of removing the material prior and FFC-Net, underscoring that the material prior and FFC-Net play crucial roles in disentangling illumination and material, as well as in inpainting, respectively. In Fig. 7, we further show the quality degradation caused by squeezing predictions into a single LRM. The results show pronounced blurring, highlighting the necessity of our decoupling strategy.

## 8. Conclusion and limitation

We present ARM, a novel method for reconstructing high-quality 3D meshes and PBR maps from sparse-view images, leveraging the advantages of UV space. ARM generates detailed meshes with high-quality textures and spatially-varying materials, outperforming existing methods both qualitatively and quantitatively. However, inconsistencies remain in multi-view images generated by upstream models, which can introduce artifacts in the reconstructed textures. Developing strategies to resolve conflicts across inconsistent views, such as weighting input views based on user specification, is a valuable direction for exploration.

## Acknowledgement

We thank anonymous reviewers and AC for their constructive comments. This work is partially supported by NSF (2153851), NSF China (62332015, 62227806 & 62421003), the XPLOER PRIZE and Information Technology Center and State Key Lab of CAD&CG, Zhejiang University.

## References

- [1] Zoubin Bi, Yixin Zeng, Chong Zeng, Fan Pei, Xiang Feng, Kun Zhou, and Hongzhi Wu. *Gs<sup>+</sup> 3: Efficient relighting with triple gaussian splatting*. *arXiv preprint arXiv:2410.11419*, 2024. 3
- [2] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. *Nerd: Neural reflectance decomposition from image collections*. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12684–12694, 2021. 3
- [3] Mark Boss, Andreas Engelhardt, Abhishek Kar, Yuanzhen Li, Deqing Sun, Jonathan Barron, Hendrik Lensch, and Varun Jampani. *Samurai: Shape and material from unconstrained real-world arbitrary image collections*. *Advances in Neural Information Processing Systems*, 35:26389–26403, 2022. 3
- [4] Mark Boss, Zixuan Huang, Aaryaman Vasishtha, and Varun Jampani. *Sf3d: Stable fast 3d mesh reconstruction with uv-unwrapping and illumination disentanglement*. *arXiv preprint arXiv:2408.00653*, 2024. 3, 6, 7, 8, 2, 4
- [5] Brent Burley and Walt Disney Animation Studios. *Physically-based shading at disney*. In *Acm Siggraph*, pages 1–7. vol. 2012, 2012. 3
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. *Emerging properties in self-supervised vision transformers*. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. 5, 1
- [7] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J Guibas, Jonathan Tremblay, Sameh Khamis, et al. *Efficient geometry-aware 3d generative adversarial networks*. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16123–16133, 2022. 2
- [8] Eric R Chan, Koki Nagano, Matthew A Chan, Alexander W Bergman, Jeong Joon Park, Axel Levy, Miika Aittala, Shalini De Mello, Tero Karras, and Gordon Wetzstein. *Generative novel view synthesis with 3d-aware diffusion models*. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4217–4229, 2023. 2
- [9] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. *Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation*. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 22246–22256, 2023. 2, 3
- [10] Zhaoxi Chen, Jiaxiang Tang, Yuhao Dong, Ziang Cao, Fangzhou Hong, Yushi Lan, Tengfei Wang, Haozhe Xie, Tong Wu, Shunsuke Saito, et al. *3dtopia-xl: Scaling high-quality 3d asset generation via primitive diffusion*. *arXiv preprint arXiv:2409.12957*, 2024. 3
- [11] Lu Chi, Borui Jiang, and Yadong Mu. *Fast fourier convolution*. *Advances in Neural Information Processing Systems*, 33:4479–4488, 2020. 5
- [12] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. *Objaverse: A universe of annotated 3d objects*. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023. 2, 5, 3
- [13] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. *Objaverse-xl: A universe of 10m+ 3d objects*. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [14] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. *Google scanned objects: A high-quality dataset of 3d scanned household items*. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2553–2560. IEEE, 2022. 6, 7
- [15] Andreas Engelhardt, Amit Raj, Mark Boss, Yunzhi Zhang, Abhishek Kar, Yuanzhen Li, Deqing Sun, Ricardo Martin Brualla, Jonathan T Barron, Hendrik Lensch, et al. *Shinobi: Shape and illumination using neural object decomposition via brdf optimization in-the-wild*. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19636–19646, 2024. 3
- [16] Ye Fang, Zeyi Sun, Tong Wu, Jiaqi Wang, Ziwei Liu, Gordon Wetzstein, and Dahua Lin. *Make-it-real: Unleashing large multimodal model’s ability for painting 3d objects with realistic materials*. *arXiv preprint arXiv:2404.16829*, 2024. 3
- [17] Jian Gao, Chun Gu, Youtian Lin, Hao Zhu, Xun Cao, Li Zhang, and Yao Yao. *Relightable 3d gaussian: Real-time point cloud relighting with brdf decomposition and ray tracing*. *arXiv preprint arXiv:2311.16043*, 2023. 3
- [18] Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. *Shape, light, and material decomposition from images using monte carlo rendering and denoising*. *Advances in Neural Information Processing Systems*, 35:22856–22869, 2022. 3
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. *Deep residual learning for image recognition*. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. *Denoising diffusion probabilistic models*. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2
- [21] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. *Lrm: Large reconstruction model for single image to 3d*. *arXiv preprint arXiv:2311.04400*, 2023. 2
- [22] Hanzhe Hu, Zhizhuo Zhou, Varun Jampani, and Shubham Tulsiani. *Mvd-fusion: Single-view 3d via depth-consistent multi-view generation*. In *Proceedings of the IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition, pages 9698–9707, 2024. 2
- [23] Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuexin Ma. Gaussianshader: 3d gaussian splatting with shading functions for reflective surfaces. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 5322–5332, 2024. 3
- [24] Heewoo Jun and Alex Nichol. Shap-e: Generating conditional 3d implicit functions. arXiv preprint arXiv:2305.02463, 2023. 2
- [25] Kaizhang Kang, Cihui Xie, Chengan He, Mingqi Yi, Minyi Gu, Zimin Chen, Kun Zhou, and Hongzhi Wu. Learning efficient illumination multiplexing for joint capture of reflectance and shape. ACM Trans. Graph., 38(6):165:1–165:12, 2019. 1
- [26] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. ACM Trans. Graph., 42(4):139–1, 2023. 1, 2
- [27] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. ACM Transactions on Graphics, 39(6), 2020. 2
- [28] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photorealistic single image super-resolution using a generative adversarial network. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 4681–4690, 2017. 1
- [29] Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan Sunkavalli, Greg Shakhnarovich, and Sai Bi. Instant3d: Fast text-to-3d with sparse-view generation and large reconstruction model. arXiv preprint arXiv:2311.06214, 2023. 2, 1
- [30] Zhihao Liang, Qi Zhang, Ying Feng, Ying Shan, and Kui Jia. Gs-ir: 3d gaussian splatting for inverse rendering. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 21644–21653, 2024. 3
- [31] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 300–309, 2023. 2
- [32] Yukang Lin, Haonan Han, Chaoqun Gong, Zunnan Xu, Yachao Zhang, and Xiu Li. Consistent123: One image to highly consistent 3d asset using case-aware diffusion priors. In Proceedings of the 32nd ACM International Conference on Multimedia, pages 6715–6724, 2024. 2
- [33] Minghua Liu, Ruoxi Shi, Linghao Chen, Zhuoyang Zhang, Chao Xu, Xinyue Wei, Hansheng Chen, Chong Zeng, Jiayuan Gu, and Hao Su. One-2-3-45++: Fast single image to 3d objects with consistent multi-view generation and 3d diffusion. arXiv preprint arXiv:2311.07885, 2023. 2
- [34] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. Advances in Neural Information Processing Systems, 36, 2024. 2
- [35] Minghua Liu, Chong Zeng, Xinyue Wei, Ruoxi Shi, Linghao Chen, Chao Xu, Mengqi Zhang, Zhaoning Wang, Xiaoshuai Zhang, Isabella Liu, et al. Meshformer: High-quality mesh generation with 3d-guided reconstruction model. arXiv preprint arXiv:2408.10198, 2024. 2, 6, 7, 4
- [36] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object, 2023. 2
- [37] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Generating multiview-consistent images from a single-view image. arXiv preprint arXiv:2309.03453, 2023. 2
- [38] Yuxin Liu, Minshan Xie, Hanyuan Liu, and Tien-Tsin Wong. Text-guided texturing by synchronized multi-view diffusion. arXiv preprint arXiv:2311.12891, 2023. 2
- [39] Zexiang Liu, Yangguang Li, Youtian Lin, Xin Yu, Sida Peng, Yan-Pei Cao, Xiaojuan Qi, Xiaoshui Huang, Ding Liang, and Wanli Ouyang. Unidream: Unifying diffusion priors for relightable text-to-3d generation. In European Conference on Computer Vision, pages 74–91. Springer, 2025. 3
- [40] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9970–9980, 2024. 2
- [41] Jonathan Lorraine, Kevin Xie, Xiaohui Zeng, Chen-Hsuan Lin, Towaki Takikawa, Nicholas Sharp, Tsung-Yi Lin, Ming-Yu Liu, Sanja Fidler, and James Lucas. Att3d: Amortized text-to-3d object synthesis. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 17946–17956, 2023. 2
- [42] Abhimitra Meka, Rohit Pandey, Christian Haene, Sergio Orts-Escolano, Peter Barnum, Philip David-Son, Daniel Erickson, Yinda Zhang, Jonathan Taylor, Sofien Bouaziz, et al. Deep relightable textures: volumetric performance capture with neural rendering. ACM Transactions on Graphics (TOG), 39(6):1–21, 2020. 1
- [43] Gal Metzger, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12663–12673, 2023. 2
- [44] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM, 65(1):99–106, 2021. 1, 2
- [45] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting Triangular 3D Models, Materials, and Lighting From Images. In Proceedings of the IEEE/CVF

- Conference on Computer Vision and Pattern Recognition (CVPR), pages 8280–8290, 2022. 1
- [46] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting triangular 3d models, materials, and lighting from images. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8280–8290, 2022. 3
- [47] Giljoo Nam, Joo Ho Lee, Diego Gutierrez, and Min H Kim. Practical svbrdf acquisition of 3d objects with unstructured flash photography. ACM Transactions on Graphics (ToG), 37(6):1–12, 2018. 1
- [48] William Peebles and Saining Xie. Scalable diffusion models with transformers. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 4195–4205, 2023. 3
- [49] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2209.14988, 2022. 2
- [50] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. arXiv preprint arXiv:1612.00593, 2016. 5
- [51] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. arXiv preprint arXiv:2306.17843, 2023. 2
- [52] Jieji Ren, Feishi Wang, Jiahao Zhang, Qian Zheng, Mingjun Ren, and Boxin Shi. Diligent102: A photometric stereo benchmark dataset with controlled shape and material variation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 12581–12590, 2022. 1
- [53] Xuanchi Ren, Jiahui Huang, Xiaohui Zeng, Ken Museth, Sanja Fidler, and Francis Williams. Xcube: Large-scale 3d generative modeling using sparse voxel hierarchies. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4209–4219, 2024. 2
- [54] 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
- [55] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. Advances in neural information processing systems, 35:36479–36494, 2022. 2
- [56] Shunsuke Saito, Gabriel Schwartz, Tomas Simon, Junxuan Li, and Giljoo Nam. Relightable gaussian codec avatars. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 130–141, 2024. 1
- [57] Junyoung Seo, Wooseok Jang, Min-Seop Kwak, Hyeonsu Kim, Jaehoon Ko, Junho Kim, Jin-Hwa Kim, Jiyoung Lee, and Seungryong Kim. Let 2d diffusion model know 3d-consistency for robust text-to-3d generation. arXiv preprint arXiv:2303.07937, 2023. 2
- [58] Prafull Sharma, Julien Philip, Michaël Gharbi, Bill Freeman, Fredo Durand, and Valentin Deschaintre. Materialistic: Selecting similar materials in images. ACM Trans. Graph., 42(4), 2023. 5, 2
- [59] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. Zero123++: a single image to consistent multi-view diffusion base model. arXiv preprint arXiv:2310.15110, 2023. 2, 3, 6
- [60] Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. arXiv preprint arXiv:2308.16512, 2023. 2
- [61] Yawar Siddiqui, Tom Monnier, Filippos Kokkinos, Mahendra Kariya, Yanir Kleiman, Emilien Garreau, Oran Gafni, Natalia Neverova, Andrea Vedaldi, Roman Shapovalov, et al. Meta 3d assetgen: Text-to-mesh generation with high-quality geometry, texture, and pbr materials. arXiv preprint arXiv:2407.02445, 2024. 2
- [62] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 7495–7504, 2021. 3
- [63] Jingxiang Sun, Bo Zhang, Ruizhi Shao, Lizhen Wang, Wen Liu, Zhenda Xie, and Yebin Liu. Dreamcraft3d: Hierarchical 3d generation with bootstrapped diffusion prior. arXiv preprint arXiv:2310.16818, 2023. 2
- [64] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In Proceedings of the IEEE/CVF winter conference on applications of computer vision, pages 2149–2159, 2022. 5
- [65] Jiayang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. arXiv preprint arXiv:2402.05054, 2024. 2, 6, 7
- [66] Dmitry Tochilkin, David Pankratz, Zexiang Liu, Zixuan Huang, Adam Letts, Yangguang Li, Ding Liang, Christian Laforte, Varun Jampani, and Yan-Pei Cao. Tripotr: Fast 3d object reconstruction from a single image. arXiv preprint arXiv:2403.02151, 2024. 2
- [67] Shimon Vainer, Mark Boss, Mathias Parger, Konstantin Kutsy, Dante De Nigris, Ciara Rowles, Nicolas Perony, and Simon Donné. Collaborative control for geometry-conditioned pbr image generation. In European Conference on Computer Vision, pages 127–145. Springer, 2025. 3
- [68] Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitry Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. Sv3d: Novel multi-view synthesis and 3d generation from a single image using latent video diffusion. In European Conference on Computer Vision, pages 439–457. Springer, 2025. 2

- [69] Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. Microfacet models for refraction through rough surfaces. *Rendering techniques*, 2007:18th, 2007. 1
- [70] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12619–12629, 2023. 2
- [71] Peng Wang and Yichun Shi. Imagedream: Image-prompt multi-view diffusion for 3d generation. *arXiv preprint arXiv:2312.02201*, 2023. 2
- [72] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021. 1, 2
- [73] Peng Wang, Hao Tan, Sai Bi, Yinghao Xu, Fujun Luan, Kalyan Sunkavalli, Wenping Wang, Zexiang Xu, and Kai Zhang. Pf-lrm: Pose-free large reconstruction model for joint pose and shape prediction. *arXiv preprint arXiv:2311.12024*, 2023. 2
- [74] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [75] Zhengyi Wang, Yikai Wang, Yifei Chen, Chendong Xiang, Shuo Chen, Dajiang Yu, Chongxuan Li, Hang Su, and Jun Zhu. Crm: Single image to 3d textured mesh with convolutional reconstruction model. *arXiv preprint arXiv:2403.05034*, 2024. 2, 6, 7
- [76] Xinyue Wei, Fanbo Xiang, Sai Bi, Anpei Chen, Kalyan Sunkavalli, Zexiang Xu, and Hao Su. Neumanifold: Neural watertight manifold reconstruction with efficient and high-quality rendering support. *arXiv preprint arXiv:2305.17134*, 2023. 3, 4
- [77] Xinyue Wei, Kai Zhang, Sai Bi, Hao Tan, Fujun Luan, Valentin Deschaintre, Kalyan Sunkavalli, Hao Su, and Zexiang Xu. Meshlrm: Large reconstruction model for high-quality mesh. *arXiv preprint arXiv:2404.12385*, 2024. 2
- [78] Kailu Wu, Fangfu Liu, Zhihan Cai, Runjie Yan, Hanyang Wang, Yating Hu, Yueqi Duan, and Kaisheng Ma. Unique3d: High-quality and efficient 3d mesh generation from a single image. *arXiv preprint arXiv:2405.20343*, 2024. 2
- [79] Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Liang Pan Jiawei Ren, Wayne Wu, Lei Yang, Jiaqi Wang, Chen Qian, Dahua Lin, and Ziwei Liu. Omniobject3d: Large-vocabulary 3d object dataset for realistic perception, reconstruction and generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 6, 7
- [80] Kevin Xie, Jonathan Lorraine, Tianshi Cao, Jun Gao, James Lucas, Antonio Torralba, Sanja Fidler, and Xiaohui Zeng. Latte3d: Large-scale amortized text-to-enhanced3d synthesis. *arXiv preprint arXiv:2403.15385*, 2024. 2
- [81] Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191*, 2024. 2, 5, 6, 7
- [82] Xudong Xu, Zhaoyang Lyu, Xingang Pan, and Bo Dai. Mat-lab: Material-aware text-to-3d via latent brdf auto-encoder. *arXiv preprint arXiv:2308.09278*, 2023. 3
- [83] Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli, Gordon Wetzstein, Zexiang Xu, et al. Dmv3d: Denoising multi-view diffusion using 3d large reconstruction model. *arXiv preprint arXiv:2311.09217*, 2023. 2
- [84] Yinghao Xu, Zifan Shi, Wang Yifan, Hansheng Chen, Ceyuan Yang, Sida Peng, Yujun Shen, and Gordon Wetzstein. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation. *arXiv preprint arXiv:2403.14621*, 2024. 2
- [85] Jianglong Ye, Peng Wang, Kejie Li, Yichun Shi, and Heng Wang. Consistent-1-to-3: Consistent image to 3d view synthesis via geometry-aware diffusion models. In *2024 International Conference on 3D Vision (3DV)*, pages 664–674. IEEE, 2024. 2
- [86] Kim Youwang, Tae-Hyun Oh, and Gerard Pons-Moll. Paint-it: Text-to-texture synthesis via deep convolutional texture map optimization and physically-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4347–4356, 2024. 3
- [87] Chong Zeng, Guojun Chen, Yue Dong, Pieter Peers, Hongzhi Wu, and Xin Tong. Relighting neural radiance fields with shadow and highlight hints. In *ACM SIGGRAPH 2023 Conference Proceedings*, 2023. 3
- [88] Chong Zeng, Yue Dong, Pieter Peers, Youkang Kong, Hongzhi Wu, and Xin Tong. Dilightnet: Fine-grained lighting control for diffusion-based image generation. In *ACM SIGGRAPH 2024 Conference Papers*, 2024. 3, 6
- [89] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5453–5462, 2021. 3
- [90] Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. Gs-lrm: Large reconstruction model for 3d gaussian splatting. In *European Conference on Computer Vision*, pages 1–19. Springer, 2025. 2
- [91] Longwen Zhang, Ziyu Wang, Qixuan Zhang, Qiwei Qiu, Anqi Pang, Haoran Jiang, Wei Yang, Lan Xu, and Jingyi Yu. Clay: A controllable large-scale generative model for creating high-quality 3d assets. *ACM Transactions on Graphics (TOG)*, 43(4):1–20, 2024. 3
- [92] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. Ner-factor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (ToG)*, 40(6):1–18, 2021. 3
- [93] Yuqing Zhang, Yuan Liu, Zhiyu Xie, Lei Yang, Zhongyuan Liu, Mengzhou Yang, Runze Zhang, Qilong Kou, Cheng Lin, Wenping Wang, et al. Dreammat: High-quality pbr material generation with geometry-and light-aware diffusion models. *ACM Transactions on Graphics (TOG)*, 43(4):1–18, 2024. 3

- [94] Xin-Yang Zheng, Hao Pan, Peng-Shuai Wang, Xin Tong, Yang Liu, and Heung-Yeung Shum. Locally attentional sdf diffusion for controllable 3d shape generation. ACM Transactions on Graphics (ToG), 42(4):1–13, 2023. [2](#)
- [95] Zi-Xin Zou, Zhipeng Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Yan-Pei Cao, and Song-Hai Zhang. Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10324–10335, 2024. [2](#)