

Progressive Rendering Distillation: Adapting Stable Diffusion for Instant Text-to-Mesh Generation without 3D Data

Zhiyuan Ma^{1,2} Xinyue Liang^{1,2} Rongyuan Wu¹ Xiangyu Zhu^{3,4} Zhen Lei^{1,2,3,4*} Lei Zhang^{1*}

¹The Hong Kong Polytechnic University

²Center for Artificial Intelligence and Robotics, HKISI CAS

³State Key Laboratory of Multimodal Artificial Intelligence Systems, CASIA

⁴School of Artificial Intelligence, University of Chinese Academy of Sciences, UCAS



Figure 1. Our method adapts Stable Diffusion [89] to generate high-fidelity textured meshes in 1.2 seconds.

Abstract

It is highly desirable to obtain a model that can generate high-quality 3D meshes from text prompts in just seconds. While recent attempts have adapted pre-trained text-to-image diffusion models, such as Stable Diffusion (SD), into generators of 3D representations (e.g., Triplane), they often suffer from poor quality due to the lack of sufficient high-quality 3D training data. Aiming at overcoming the data shortage, we propose a novel training scheme, termed as Progressive Rendering Distillation (PRD), eliminating the need for 3D ground-truths by distilling multi-view diffusion models and adapting SD into a native 3D generator. In each iteration of training, PRD uses the U-Net to progressively denoise the latent from random noise for a few steps, and in each step it decodes the denoised latent into 3D output. Multi-view diffusion models, including MVDream and RichDreamer, are used in joint with SD

to distill text-consistent textures and geometries into the 3D outputs through score distillation. Since PRD supports training without 3D ground-truths, we can easily scale up the training data and improve generation quality for challenging text prompts with creative concepts. Meanwhile, PRD can accelerate the inference speed of the generation model in just a few steps. With PRD, we train a Triplane generator, namely TriplaneTurbo, which adds only 2.5% trainable parameters to adapt SD for Triplane generation. TriplaneTurbo outperforms previous text-to-3D generators in both efficiency and quality. Specifically, it can produce high-quality 3D meshes in 1.2 seconds and generalize well for challenging text input. The code is available at github.com/theEricMa/TriplaneTurbo.

1. Introduction

Text-to-3D aims to create 3D content faithful to the given textual descriptions. The optimization-based text-to-3D

*Corresponding authors.

methods can achieve high generation fidelity by using image diffusion priors [53, 63, 68, 70, 73, 86, 96, 113] and score distillation techniques [43, 57, 75, 84, 114, 119, 148, 158] to optimize 3D representations [10, 30, 60, 61, 108]. However, these approaches encounter bottlenecks in computational efficiency since they take hours to generate 3D textured meshes. Recent researches have shifted towards learning-based methods [7, 25, 38, 50, 82], which generate 3D content through feedforward networks, reducing the latency to a few seconds per output. Unfortunately, existing 3D datasets [16–18, 83, 127] are much smaller compared to those used in training text-to-image generation models, while the 3D data therein suffer from texture quality and inconsistent object poses [62]. Consequently, these approaches struggle to produce high-quality 3D outputs. As a promising alternative to the above mentioned methods, one can adapt pretrained text-to-image models, such as Stable Diffusion (SD) [89], into generators of 3D representations [55, 67, 78]. Recent studies have enabled the use of SD to generate 2D planes for 3D representations such as Triplane [8]. PI3D [67] adapts SD to generate six 2D planes pre-optimized for 3D object reconstruction. HexaGen [78] first trains a VAE to compress 3D objects into latent space, then adapts SD to generate the desired latents from text prompts. These methods leverage SD’s prior knowledge for diverse text prompts while reducing training costs. However, their reliance on data-driven training limits the generalization performance.

In this paper, we propose to address the data shortage problem by using high-quality multi-view diffusion models [68, 86, 96] as teachers and performing multi-view distillation to adapt SD into an instant text-to-mesh generator without using 3D ground-truth data. Data-free distillation techniques [71, 159] have been applied in previous methods [52, 72, 75, 134] to train native 3D generators from scratch, whereas they struggle to produce high-quality 3D outputs. We argue that if we can adapt SD for native 3D generation, instead of training from scratch, not only the training efficiency can be improved but also the generation quality can be enhanced. To our best knowledge, however, no such attempt has been made before. The primary obstacle lies in the training scheme of SD, which requires 3D ground-truth data to supervise the denoising process. This requirement, unfortunately, conflicts with our goal to eliminate the dependency on 3D training data.

To address the above challenges, we propose **Progressive Rendering Distillation** (PRD), which enables 3D-data-free distillation. PRD achieves this goal by denoising the latent from random noise rather than 3D ground-truths. In each training iteration, PRD uses a few steps to progressively reverse noise to latent space using the adapted U-Net of SD. Then the denoised latent of each step is decoded to 3D content via the adapted VAE decoder. Multi-

view teacher models are used to distill high-quality renderings through efficient score distillation techniques such as ASD [75]. An additional benefit of PRD is that by adapting the SD to generate 3D content in just four steps, the overall generation process can be much accelerated.

PRD is flexible in the choices of multi-view teacher models and the types of 3D representations. Unlike existing approaches using SD alone as the teacher [60, 107, 119, 144, 158], which risk the multi-view inconsistency and the Janus problem [4], we employ multiple teachers in training. Specifically, we employ MVDream [96] and SD to suppress the Janus problem and ensure text consistency in 3D content supervision. Additionally, we use RichDreamer [86] for geometry supervision through normal and depth. With PRD, we adapt SD into a native 3D generator to produce geometry and texture Triplanes in 4 steps, which is named as **TriplaneTurbo** in the following context. To enhance TriplaneTurbo’s textured mesh quality, we employ SDF-based volumetric rendering [142] and mesh rasterization [122] for multi-view teacher supervision.

In addition, to address the high GPU memory usage caused by multiple teachers and multi-view renderings in the training process, we introduce the **Parameter-Efficient Triplane Adapter** (PETA), which adds only 2.5% trainable parameters to the frozen SD and effectively adapts it for 3D generation. Some results are shown in Fig. 1. To our knowledge, *this is the first parameter-efficient training method for direct 3D content generation from 2D diffusion models*, departing from full-parameter adaptation approaches. Our key contributions are summarized as follows:

- We make the first attempt to adapt the pretrained 2D SD model into a native 3D generator without 3D data. With the proposed PRD scheme, we use multi-view diffusion models as teachers and distill SD into a four-step native 3D generator, namely TriplaneTurbo.
- TriplaneTurbo adds only 2.5% additional trainable parameters to the frozen SD for Triplane adaptation. It marks the first use of parameter-efficient training to adapt 2D diffusion models for native 3D generation.
- TriplaneTurbo surpasses existing text-to-3D models not only in quality but also in speed, reducing text-to-mesh generation time to just 1.2 seconds. In addition, by scaling up the text training data, the model can generalize much better to complex text input.

2. Related Work

Data-Driven Models for 3D Generation. Employing 3D data to train generators has shown its effectiveness for single-category 3D generation such as human faces [2, 8, 32, 74, 101, 129, 132], body shapes [35, 58, 91, 138], everyday objects [3, 40, 66, 80], and structured room layouts [5, 21, 77, 88, 110, 130]. This is mainly because for these specific categories, the training data are relatively

easy to collect. However, text-to-3D generation requires open-category generation, bringing a rather different challenge. The difficulties in obtaining sufficient 3D training data largely limit the model generalizability across diverse text prompts [75], regardless of the type of used generation models (e.g., GAN [24, 81, 93] or diffusion models [29, 41, 99]) or 3D representations (e.g., NeRF [6, 79] or Gaussian Splatting [45, 149]). The publicly available 3D datasets [16–18, 83, 127] mostly contain only hundreds or thousands of samples, which are very hard, if not impossible, to train a generalized text-to-3D generation model.

2D Diffusion Models for 3D Generation. The difficulty of text-to-3D generation can be alleviated by leveraging 2D diffusion models as priors, thanks to their training on vast text-image pairs. Early approaches like SDS [107] and VSD [13] pioneered zero-shot text-to-3D generation by optimizing 3D representations (NeRF [60, 63, 84, 114], mesh [11, 60, 119, 140], 3DGS [57, 108, 144, 145]) through score distillation [1, 37, 42, 56, 57, 75, 84, 106, 117, 119, 121, 131, 147, 155, 157, 160], which serves as a bridge to transfer the generative capability of 2D diffusion models to 3D representations through rendered views. However, pre-trained text-to-image models like SD lack multi-view consistency, leading to the Janus problem [4]. MV-Dream [96] and subsequent works [36, 69] address this issue by camera-aware adaptations and synchronized multi-view generation, and they incorporate additional modalities such as normal [86], depth [22, 44], and CCM [65, 135] to enhance geometric quality. While achieving improved 3D generation results, these methods require computationally intensive score distillation [12, 86, 92, 96] or 3D reconstruction [33, 54, 64, 70, 95, 109, 120, 139, 151]. Native 3D generators [17, 26, 35, 50, 51, 59, 76, 111, 126, 133, 150, 153, 153, 156, 156] can reduce the generation time to seconds by directly producing 3D content without view rendering as proxies. For example, LN3Diff [50] and GVGGEN [26] compress NeRF [79] and 3DGS [45] into latent spaces using VAEs [48], then train text-conditioned latent diffusion models [29]. However, these methods show limited generalizability across text prompts, as their performance is constrained by the insufficient text-3D training pairs.

Some methods [67, 78, 86, 96] leverage SD as their backbone to transfer text-to-image knowledge into text-to-3D generation. SD provides strong generative prior and improves the model generalizability [9, 19, 59, 102–105, 124, 125, 141, 152]. Approaches like PI3D [67] adapt SD to generate multiple planes for constructing 3D space, yet their performance remains limited due to the insufficient 3D training data. Instead of adapting SD for multi-view generation, native 3D generation requires substantially more 3D data for effective adaptation. We propose to address the data insufficiency challenge by distilling knowledge from multi-view diffusion models into an SD-adapted

native 3D generator, eliminating the need for 3D training data. While previous works like ATT3D [72] and Scale-Dreamer [75] have investigated such a data-free training strategy, they employ multi-view distillation to train generators from scratch, and show limited performance due to the insufficient training scale. To overcome the huge training cost, we propose a cost-effective solution that combines multi-view distillation with SD-based native 3D generation. A fundamental challenge to achieve this goal is how to adapt SD for 3D generation without ground-truth data. We address this by proposing a novel Progressive Rendering Distillation scheme, which not only eliminates the need of 3D ground-truths but also enables few-step generation.

3. Method

3.1. Preliminary

Stable Diffusion (SD) performs diffusion in latent space for efficient text-to-image generation. Its VAE encoder $E_{\phi^{\text{SD}}}$ compresses an image x into a latent code z , while its decoder $D_{\phi^{\text{SD}}}$ reconstructs the image. Given text prompt y , a U-net $\epsilon_{\phi^{\text{SD}}}$ predicts noise ϵ , which is added to $z_t = \alpha_t z + \sigma_t \epsilon$, where timestep $t \in 1, \dots, T$ controls noise level via scalars α_t and σ_t . Generation proceeds by iteratively denoising from z_T to prompt-aligned z_0 . At each step, the U-net estimates noise $\hat{\epsilon} = \epsilon_{\phi^{\text{SD}}}(z_t; t, y)$ to compute $\hat{z}_0 = \frac{z_t - \sigma_t \hat{\epsilon}}{\alpha_t}$, denoted as $\hat{z}_0 = z_{\phi^{\text{SD}}}(z_t; t, y)$. Results can be refined through additional diffusion steps $t' < t$. The final latent is decoded to an image via $\hat{x} = D_{\phi^{\text{SD}}}(\hat{z}_0)$.

Score Distillation. 2D diffusion models can optimize 3D representations θ through differentiable rendering $x_\pi = g(\theta, \pi)$ [46, 94, 97], which produces images x_π from camera view π . Here 2D diffusion models serve as a metric $\mathcal{L}(x_\pi; \pi, y)$ that measures the consistency between x_π and the text conditions y . The 3D representation is optimized using gradient $\nabla_\theta \mathcal{L}(x_\pi; \pi, y) = \nabla x_\pi \mathcal{L}(x_\pi; \pi, y) \frac{\partial x_\pi}{\partial \theta}$, which also trains the native 3D generator. The computation of $\nabla x_\pi \mathcal{L}$ depends on the chosen score distillation method.

3.2. Progressive Rendering Distillation

We now detail our proposed training scheme for adapting SD as a native 3D generator. Traditional adaptation approaches require preparing ground-truth 3D representations θ and their corresponding latents z for each 3D sample in the dataset. The U-net is adapted to denoise diffused latents z_t by minimizing the noise prediction mean squared error (MSE). Fig. 2(a) illustrates this paradigm, which has been used by several native 3D generators [67, 118, 126, 153]. However, this paradigm faces limitations in both the quantity and quality of available 3D representations θ , as existing 3D datasets lack sufficient high-quality data for training text-to-3D generators. Actually, the pretrained SD models already possess denoising capabilities for image generation.

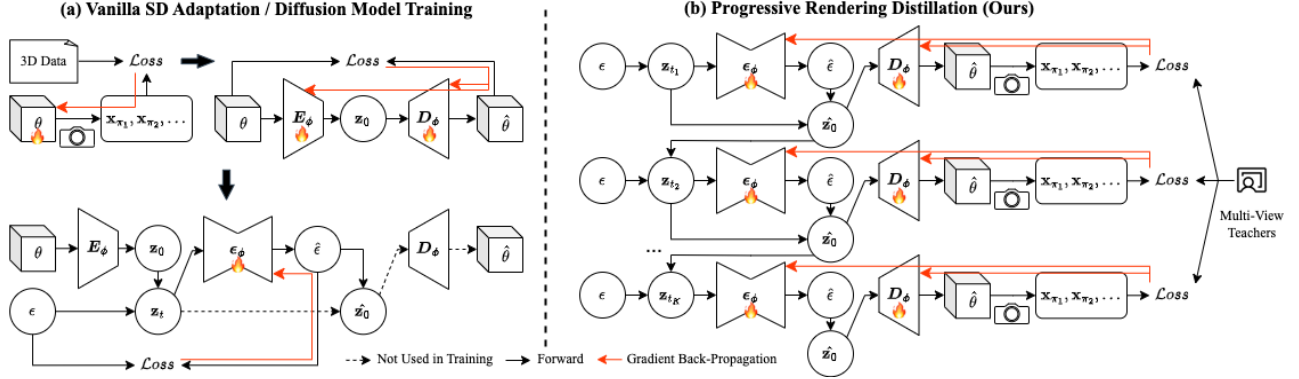


Figure 2. Comparison between (a) traditional SD adaptation and (b) our proposed progressive rendering distillation (PRD) for native 3D generation. Traditional approach requires ground-truth 3D representations θ and their latents z_0 for each 3D sample to generate z_0 . Our proposed PRD scheme progressively denoises latents z_t initialized from random noise into z_0 , which are decoded to θ , using multi-view diffusion models as teachers for distillation, eliminating the need for 3D data during adaptation and overcoming data scarcity.

In other words, pretrained SD is well-trained for a Markov Chain to reverse $z_T = \epsilon \sim \mathcal{N}(0, I)$ to z_0 with its U-net $z_{\phi^{\text{SD}}}$, and decode z_0 to image with its decoder $D_{\phi^{\text{SD}}}$. Our goal is to modify the Markov chain by transforming $z_{\phi^{\text{SD}}}$ and $D_{\phi^{\text{SD}}}$ into 3D generators $z_{\phi^{\text{3D}}}$ and $D_{\phi^{\text{3D}}}$, from which the 3D representations θ can be decoded. Note that our modification of the Markov chain differs from the traditional diffusion model adaptation objectives, as it requires neither ground-truth latents z_0 nor their noise-diffused variants z_t in the training process.

Specifically, at the beginning of the Markov chain, the network takes random noise ϵ as input to represent z_T . At each step, the current state z_t is used to estimate \hat{z}_0 through $z_{\phi^{\text{SD}}}$, which is then decoded to 3D output $\hat{\theta}$ with $D_{\phi^{\text{SD}}}$. The 3D output $\hat{\theta}$ is used to render images $x_{\pi_1}, x_{\pi_2}, \dots$ at camera views π_1, π_2, \dots and receive supervision from multi-view teachers via score distillation, while the estimated \hat{z}_0 is diffused to the next timestep t' as $z_{t'} = \alpha_{t'} \hat{z}_0 + \sigma_{t'} \epsilon$ for subsequent operations. We name this training scheme Progressive Rendering Distillation (PRD), as illustrated in Fig. 2(b). From the total T timesteps, we select a decreasing sequence of K timesteps $T = t_1 > t_2 > \dots > t_K = T/K$ with uniform spacing to perform score distillation from multi-view teachers. The gradient at each step is not back-propagated to previous steps; therefore, we can largely reduce the GPU memory usage and prevent gradient explosion [15, 85, 128, 137]. Since ϕ^{3D} is initialized from ϕ^{SD} , this specialized gradient detachment strategy still maintains good convergence. The pseudo code of our algorithm is provided in Algorithm 1. While the time cost of our training strategy increases with the increase of K , we can ensure the model generates high-quality results in just a few steps, thereby accelerating inference. We set $K = 4$ to balance quality and speed.

3.3. Parameter-Efficient Triplane Adaptation

While various 3D representations could be employed by our PRD scheme, in this paper we demonstrate an exemplar solution using Triplanes as the representation. We denote our adapted model as **TriplaneTurbo** with parameters ϕ^{3D} . Specifically, TriplaneTurbo adapts SD to generate a 3D representation consisting of two Triplanes [8] $\theta = (\theta_{\text{geo}}, \theta_{\text{tex}})$: a geometry Triplane θ_{geo} storing Signed Distance Function (SDF) and deformation values for mesh extraction, and a texture Triplane θ_{tex} containing RGB attributes for painting texture on the mesh. For each point in 3D space, we use a two-layer MLP to decode its SDF value. The same process applies to texture and mesh deformation [94, 122]. This separation of geometry and texture planes follows the work in [23, 100, 123]. We set the Triplane resolution as 256×256 , and replace the last convolution in SD's decoder to output 32 channels. The Triplanes' output has a dimension of $6 \times 256 \times 256 \times 32$ in feature space. Due to the $8 \times$ compression of the VAE, this corresponds to $6 \times 32 \times 32 \times 4$ in latent space, distinct from the $1 \times 64 \times 64 \times 4$ latent generated by pretrained SD. To enable interaction between the six feature planes, we follow existing approaches [70, 95, 96, 115] to adapt the U-net's self-attention [89, 112] to allow cross-plane attention. Unlike existing works that fully retrain SD [12, 55, 67, 78, 96], which can lead to catastrophic forgetting [67, 96], we propose a parameter-efficient adaptation approach. The core of our design lies in the fact that each of the six feature planes maintains its own unique feature distribution. Therefore, plane-specific characteristics must be incorporated into the adaptation process. As illustrated in Fig. 3, our adaptation modifies the convolution, self-attention, and cross-attention layers. We name our approach **Parameter-Efficient Triplane Adaptation (PETA)**.

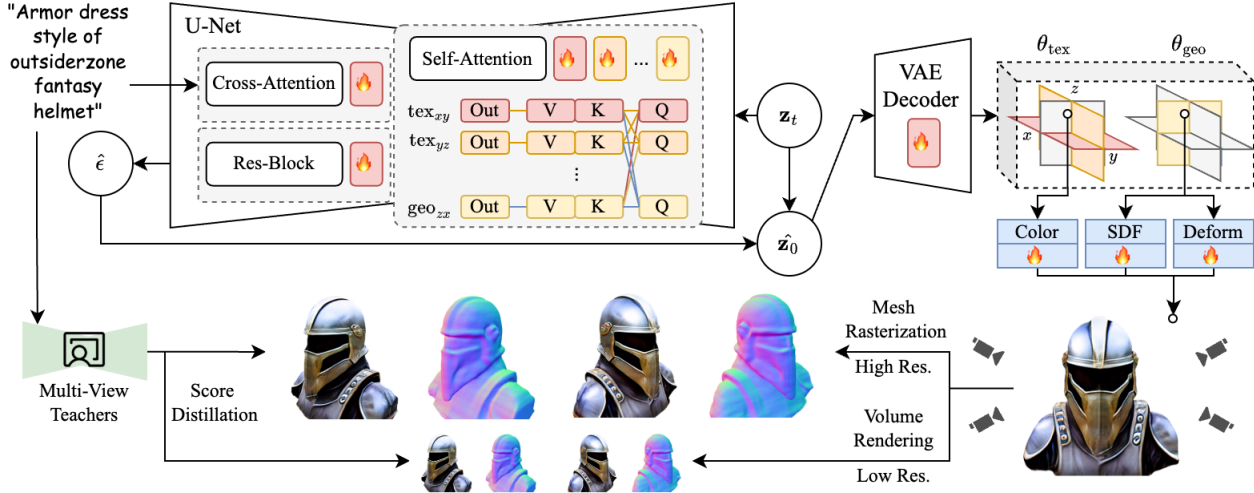


Figure 3. Illustration of TriplaneTurbo: an SD-adapted native 3D generator using our PRD scheme. Our model generates six feature planes comprising geometry Triplane θ_{geo} and texture Triplane θ_{tex} in 4 steps. We introduce Parameter Efficient Triplane Adaptation (PETA), which requires only 2.5% additional parameters for adaptation. The parameter arrangement is illustrated in the figure.

As shown in Fig. 3, for the convolution blocks (Res-Blocks) and cross-attention layers, we implement LoRA [34, 143] for parameter-efficient adaptation, and process the six planes uniformly. The plane-specific adaptations are then applied to the self-attention. For the self-attention blocks, we apply distinct LoRA layers [34] to the to_Q, to_K, to_V, and to_Out linear layers when processing each of the six feature planes. In each linear transformation within a self-attention block, the linear projection with multiple LoRAs is implemented in two steps. First, the frozen linear layer batch processes the features extracted from all the six planes together. Then, separate LoRA transformations are applied independently to the features of each plane. This adaptation maintains low computational overhead while effectively introducing plane-specific processing during attention calculations across the six feature planes. It can be applied with other techniques like AdaLoRA [154] and Vera [49]. We leave this for further exploration. We set the LoRA rank to 16 by default. While this adaptation adds only 2.5% of the parameters to the SD model, it effectively enables native 3D generation.

3.4. Distillation Details

Since PRD eliminates the need for 3D data by referring to multi-view teachers for distillation, using **multiple teachers** allows us to combine their strengths while mitigating individual biases. Most previous works [14, 60, 72, 84, 108] use SD model (parametrized by ϕ^{SD}) as the teacher for its ability to generate high-fidelity, text-consistent images. However, SD lacks camera-awareness, which can lead to the Janus problem [4]. MVDream [96] (MV, parametrized by ϕ^{MV}) addresses this by generating four

camera-conditioned views simultaneously, but at the cost of reduced prompt consistency [115]. While SD and MV complement each other, both of them focus on RGB rendering and provide no direct supervision on geometry. We further incorporate RichDreamer [86] (RD, parameterized by ϕ^{RD}), a model that generates four-view normal and depth maps based on text prompts. The score distillation guidance $\mathcal{L}(x_\pi; \pi, y)$ in our implementation thus integrates SD, MV, and RD. At each PRD step (see Sec. 3.2 and Fig. 2), given a text prompt y and generated 3D representation θ from $z_{\phi^{\text{3D}}}$ and $D_{\phi^{\text{3D}}}$, we sample four views π_1, \dots, π_4 at uniform azimuth intervals. We sample one view π from π_1, \dots, π_4 to compute $\mathcal{L}_{\text{SD}} = \mathcal{L}_{\phi^{\text{SD}}}(x_\pi; \pi, y)$. For MV, all the four views are used to compute $\mathcal{L}_{\text{MV}} = \mathcal{L}_{\phi^{\text{MV}}}(x_{\pi_1}, \dots, x_{\pi_4}; \pi_1, \dots, \pi_4, y)$. RD operates on concatenated normal and depth renderings $x'_{\pi_1}, \dots, x'_{\pi_4}$, and its objective is $\mathcal{L}_{\text{RD}} = \mathcal{L}_{\phi^{\text{RD}}}(x'_{\pi_1}, \dots, x'_{\pi_4}; \pi_1, \dots, \pi_4, y)$. Among existing score distillation methods [37, 42, 56, 57, 75, 84, 119, 121, 131, 147], we adopt Asynchronous Score Distillation (ASD) [75] for its pioneered efficiency in training deep text-to-3D generators. For 3D rendering, as illustrated in Fig. 3, we combine volumetric rendering [142] with mesh rasterization [122] to overcome the instability of pure mesh supervision [136]. See Sec. B for more details.

4. Experiments

4.1. Experimental Settings

Implementation Details. For a fair comparison with existing methods, we use captions of the 3D objects [17] provided by [31] to train our model, which comprises a total of 360K text prompts. We employ SD v2.1-base [90] as the



Figure 4. Qualitative comparison of text-to-mesh generation results by competing methods. Please refer to Sec. 4.2 for details.

base model. Our model is trained for 15K iterations with a learning rate of $2e-4$, costing 40 hours on 8 NVIDIA H20 GPUs. We further collect 1.6 million text prompts to evaluate the benefit of data scaling supported by our method; the detail is provided in Sec. D. Additional training details and loss weights are provided in Sec. B.

Compared Methods. Our proposed PRD aims for fast text-to-mesh generation. Therefore, we compare it against state-of-the-art approaches that can generate textured meshes within one minute, including Shape-E [39], Direct3D [62], 3DTopia [31], GVGEN [26], LN3Diff [50], LGM [109] and PI3D [67]. Note that we do not compare with some relevant works [78, 134] since their codes/models are not publicly available and they employ different evaluation protocols. In our experiments, the results of Shape-E, Direct3D, 3DTopia, GVGEN, LN3Diff and LGM are obtained by running their publicly available models. For PI3D, we can only perform quantitative com-

parisons with it by copying its results from the original paper, but we cannot perform visual comparisons with it since its code/model is not yet publicly available.

Evaluation Protocol. We employ the protocol used in our competing methods [26, 31, 39, 50, 67, 109] to evaluate our PRD model. Specifically, we use the ViT-B/34 CLIP model [87] to evaluate the test prompts from the DreamFusion gallery [84]. As in [67], we render the generated 3D results at 512 resolution from four viewpoints at 15° elevation across four azimuth angles: 0° , 90° , 180° and 270° . Under these views, we evaluate the performance of competing methods using CLIP Score [27] (C.S.) and CLIP R-precision (R@1). To ensure a fair comparison for instant text-to-mesh generation, we exclude the post-processing steps such as SDS refinement [84], as they will prolong the generation time by several minutes. For models [26, 109] using Gaussian Splatting representations, we use the conversion script from [109] to generate textured meshes.

4.2. Experimental Results

We showcase qualitative comparisons of the competing methods in Fig. 4 using challenging test prompts. One can see that most of the existing methods struggle to generate satisfactory outputs. Our method, in comparison, produces better quality results with complete and vivid meshes. The quantitative results are reported in Tab. 1. Again, our method demonstrates superior performance to its competitors in all the metrics. Existing methods fall short in the consistency of object poses and in the capability of handling complex prompts. Our PRD method addresses these challenges, enabling faster inference speed and improving the model performance by data scaling.

Handling Inconsistent Generation Poses. As shown in Fig. 4, existing models often generate objects with incorrect poses. For prompts like ‘*A teal moped*’, the competing methods like Direct3D [126], Shape-E [39] and GVGEN [26] generate objects with misaligned orientations, such as facing sideways or backwards. This issue stems from the inconsistent object orientations in 3D training datasets, which are difficult to detect and correct automatically. Current methods struggle with noisy training data, leading to pose ambiguity and geometric defects. Our approach addresses this challenge by learning from multi-view teachers. Though our teachers are trained on noisy 3D datasets and may occasionally provide incorrect directional guidance, our PRD scheme exposes the 3D outputs to multi-view teachers K times per iteration (see Sec. 3.2). Through the iterative distillation process, we substantially decrease the impact of incorrect directional guidance, ensuring consistent pose alignment in the 3D outputs.

Handling Complex Text Prompts. The failure of our competing methods mainly stems from their reliance on the existing text-3D paired datasets, which are not comprehensive enough compared to the diverse user inputs. Therefore, trained on these data, existing methods fail to produce good results when the input is complex. For example, existing methods might perform well when the input prompt is ‘*A robot*’ but fail when the prompt becomes ‘*A robot tiger*’. The quality of existing 3D datasets is also compromised by the absence of creative concepts. Creating 3D models for imaginative concepts like ‘*A robot tiger*’ demands significant time and expertise from 3D artists. As a result, existing 3D datasets are largely limited to common everyday objects. Models trained on these limited datasets fail to generalize effectively to imaginative or complex prompts. We solve this problem by adapting SD as a 3D generator and inheriting its generative power, and more crucially, by introducing a training scheme that completely eliminates the need of 3D data. Our training scheme not only improves the quality of results on the existing training corpus, as shown by the qualitative and quantitative results in Fig. 4 and Tab. 1, but also enhances the capability to handle complex prompts by

	C.S. \uparrow	R@1 \uparrow	Latency (/sec)
Shape-E [39]	55.1	27.1	13.0
Direct3D	60.8	4.33	16.0
3DTopia [31]	59.7	11.2	23.7
PI3D* [67]	65.9	25.2	3.00
GVGEN [26]	51.1	2.44	49.2
LN3Diff [50]	55.9	5.09	8.16
LGM [109]	67.4	28.3	56.1
Ours	68.2	32.3	1.23
+More Text Data	75.1	46.0	1.23

Table 1. Quality and speed comparison for text-to-mesh generation. *indicates that the values are quoted from the original papers.


expanding the available training data, as detailed below.

Fast Inference Speed. Our PRD method also demonstrates superior computational efficiency. As shown in Tab. 1, we evaluate the average inference latency from text input to textured mesh output across all test prompts on the H20 GPU device. While some methods [26, 109] require dozens of steps, our PRD approach enables the native 3D generator to produce quality results in just K steps. With the suggested setting of $K = 4$, as shown in Tab. 1, our model achieves sub-second latency for text-to-mesh generation, significantly outperforming previous methods.

Scaling Up Training Corpus. Since our method is free from the constraints of 3D training data, it can be easily scaled to accommodate more complex and creative text prompts during training. It preserves the SD model’s ability to handle creative concepts throughout the 3D adaptation process, generating 3D outputs that faithfully represent the input prompts. As can be seen in the **+More Text Data** column, when we scale up the training text data from 360K to 1.6M, the CLIP similarity improves by as much as 7%, leading to more consistent generation for challenging text prompts such as *A lion reading the newspaper* and *A tray of Sushi containing pugs*. This is because our collected text data covers a wider range of creative concepts than the existing 3D datasets [17, 127], thus providing more sufficient training and improving the generalizability of the model. More visual examples are partially presented in Fig. 1 and detailed in Fig. 5 and Fig. 8.

4.3. Ablation Study

The Effectiveness of PRD. We validate our PRD algorithm by testing a simplified configuration with $\mathbf{K=1}$, which reduces our method to a single-step generation process that is equivalent to a vanilla native 3D generator trained by score distillation [75]. As shown in Fig. 6 and quantified in Tab. 2, this configuration fails to generate proper 3D structures because it needs to simultaneously handle two complex challenges: adapting SD for 3D generation and performing single-step generation. In contrast, our PRD




A black Dragonborn Bard that plays an ocarina in a fantasy setting



Dinosaur in new york by Jean Dubuffet




20 year old Serbian with brown curly mullet in Naruto art form




A hamster wearing a top hat and suit imagining of kicing a football, award winning, realistic painting



Dragon ball supers goku, photorealistic ultra, detailed_8k



Ghost on skateboard cartoon style



A hobbit with silver hair planting raspberries in a cafeteria, graffiti art, highly detailed



Female beauty by the standards of 5th century Europe



Jared Leto's Joker in the style of The Batman Animated Series episode screencapture



Arnold schwarzenegger shirt suit shirtless muscle




Female halfelf rogue red hair slightly pointed earscanyon landscape



Dungeons and Dragons Bugbear Merchant fat many ring piercings creepy smile



DayZ videogame bear attack scene



Beautiful Elsa princess eating ice-cream in a snowy wonderland, fantasy style and hyperrealistic



Formula 1 view from the side off-road, wheels rugged feel white background



A police woman wearing gas mask



Cleopatra wearing the VR glasses and earphones typing on the labtop white dress animation style cyperbunk




A happy moment as a bearded man with a bald head finds a key



Medusa wearing a sunglass and shopping with a snake around her neck



Dungeons and Dragons dwarf paladin



Harry Potter riding a giraffe in a secret hideaway, realistic photograph, hyper-realistic

Figure 5. More results of our model trained with expanded corpus.



Figure 6. Visualizations on the ablation studies of PRD algorithm.

	C.S. \uparrow	R@1 \uparrow
K=1	50.9	14.4
K=2	62.6	22.4
K=4 (Proposed)	68.2	32.3

Table 2. Ablation study on the hyper-parameters of PRD.

scheme can use alternative step configurations such as **K=2**, which produces suboptimal yet acceptable results. We found that the configuration of **K=4** provides the best trade-off between output quality and computational efficiency, which is used as the default setting of PRD.

The Effectiveness of PETA. We first compare our proposed PETA method with conventional full parameter fine-tuning (shown as **Full Param Tuning** in Fig. 7 and Tab. 3). We see that full parameter fine-tuning exhibits training instability and catastrophic forgetting, resulting in text-inconsistent outputs. We then conduct additional ablation studies using vanilla LoRA tuning, maintaining an equivalent parameter size (22.6M), shown as the configuration of **w/ Basic LoRA**. We see that this vanilla adaptation produces degraded geometric structures and textures, and lacks the capability to handle the unique characteristics of each plane. Since different geometry and texture planes (see Fig. 3) exhibit distinct feature distributions, they require specialized consideration. Our solution (denoted as **w/ PETA**) considers plane dependency using multiple LoRAs in self-attention blocks, enabling each plane to maintain its unique representation while allowing cross-plane interactions during self-attention computation. As shown in Fig. 6 and Tab. 2, PETA achieves enhanced 3D generation quality with good text consistency.

We also perform ablation studies to investigate the impact of multiple multi-view teachers in training, the choice of LoRA rank, our hybrid rendering approach that combines volumetric rendering [116] and mesh rasterization [122] for multi-view distillation in PRD training scheme. The details




Figure 7. Visualizations on the ablation study of PETA.

	C.S. \uparrow	R@1 \uparrow
Full Param Tuning	35.8	0.35
w/ Basic LoRA	54.2	11.1
w/ PETA (Proposed)	68.2	32.3

Table 3. Ablation study on the effectiveness of PETA.

can be found in Sec. E.

5. Conclusion and Limitation

In this paper, we presented Progressive Rendering Distillation (PRD), a novel training scheme that adapts Stable Diffusion (SD) for instant text-to-mesh generation without relying on 3D data. We also introduced PETA (Parameter-Efficient Triplane Adaptation), a parameter-efficient method that introduces only 2.5% additional parameters to effectively enable SD for instant text-to-mesh generation. Our model, namely TriplaneTurbo, can produce text-consistent textured meshes in only 1.2 second. Through comprehensive experiments, we validated the effectiveness of our approach. Our methodology has the potential to be extended to 3D scene generation and image-to-3D tasks. While currently implemented with SD, the PRD approach can also be applied to other pre-trained models like DiT [20]. We hope our work can inspire new directions in 3D generation to overcome the dependency on 3D data.

Limitations. One limitation of our method lies in the generation of precise numbers of multiple 3D objects, which may require more sophisticated multi-view teachers, potentially enhanced with layout guidance. Besides, our results for full-body humans might exhibit limited facial and hand details, which can be improved by extending SD adaptation to more advanced 3D structures than Triplane.

Acknowledgment

This work is supported by the InnoHK program.

References

- [1] Thiemo Alldieck, Nikos Kolotouros, and Cristian Sminchisescu. Score distillation sampling with learned manifold corrective, 2024. 3
- [2] Sizhe An, Hongyi Xu, Yichun Shi, Guoxian Song, Umit Y Ogras, and Linjie Luo. Panohead: Geometry-aware 3d full-head synthesis in 360deg. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 20950–20959, 2023. 2
- [3] Titas Anciukevičius, Zexiang Xu, Matthew Fisher, Paul Henderson, Hakan Bilen, Niloy J Mitra, and Paul Guerrero. Renderdiffusion: Image diffusion for 3d reconstruction, inpainting and generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12608–12618, 2023. 2
- [4] Mohammadreza Armandpour, Ali Sadeghian, Huangjie Zheng, Amir Sadeghian, and Mingyuan Zhou. Re-imagine the negative prompt algorithm: Transform 2d diffusion into 3d, alleviate janus problem and beyond. *arXiv preprint arXiv:2304.04968*, 2023. 2, 3, 5
- [5] Sherwin Bahmani, Jeong Joon Park, Despoina Paschalidou, Xingguang Yan, Gordon Wetzstein, Leonidas Guibas, and Andrea Tagliasacchi. Cc3d: Layout-conditioned generation of compositional 3d scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7171–7181, 2023. 2
- [6] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5855–5864, 2021. 3
- [7] Ziang Cao, Fangzhou Hong, Tong Wu, Liang Pan, and Zhiwei Liu. Large-vocabulary 3d diffusion model with transformer. *arXiv preprint arXiv:2309.07920*, 2023. 2
- [8] 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, 4
- [9] Bin Chen, Gehui Li, Rongyuan Wu, Xindong Zhang, Jie Chen, Jian Zhang, and Lei Zhang. Adversarial diffusion compression for real-world image super-resolution. *arXiv preprint arXiv:2411.13383*, 2024. 3
- [10] 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 (ICCV)*, 2023. 2
- [11] 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. 3
- [12] Yabo Chen, Jiemin Fang, Yuyang Huang, Taoran Yi, Xiaopeng Zhang, Lingxi Xie, Xinggang Wang, Wenrui Dai, Hongkai Xiong, and Qi Tian. Cascade-zero123: One image to highly consistent 3d with self-prompted nearby views. *arXiv preprint arXiv:2312.04424*, 2023. 3, 4
- [13] Xinhua Cheng, Tianyu Yang, Jianan Wang, Yu Li, Lei Zhang, Jian Zhang, and Li Yuan. Progressive3d: Progressively local editing for text-to-3d content creation with complex semantic prompts, 2023. 3
- [14] Jaeyoung Chung, Suyoung Lee, Hyeongjin Nam, Jaerin Lee, and Kyoung Mu Lee. Luciddreamer: Domain-free generation of 3d gaussian splatting scenes. *arXiv preprint arXiv:2311.13384*, 2023. 5
- [15] Kevin Clark, Paul Vicol, Kevin Swersky, and David J Fleet. Directly fine-tuning diffusion models on differentiable rewards. *arXiv preprint arXiv:2309.17400*, 2023. 4
- [16] Jasmine Collins, Shubham Goel, Kenan Deng, Achleshwar Luthra, Leon Xu, Erhan Gundogdu, Xi Zhang, Tomas F Yago Vicente, Thomas Dideriksen, Himanshu Arora, et al. Abo: Dataset and benchmarks for real-world 3d object understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 21126–21136, 2022. 2, 3
- [17] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, Eli VanderBilt, Aniruddha Kembhavi, Carl Vondrick, Georgia Gkioxari, Kiana Ehsani, Ludwig Schmidt, and Ali Farhadi. Objaverse-XL: A universe of 10m+ 3d objects. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. 3, 5, 7, 19, 20, 22
- [18] 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. 2, 3
- [19] Slava Elizarov, Ciara Rowles, and Simon Donné. Geometry image diffusion: Fast and data-efficient text-to-3d with image-based surface representation. *arXiv preprint arXiv:2409.03718*, 2024. 3
- [20] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. 9
- [21] Chengzeng Feng, Jiacheng Wei, Cheng Chen, Yang Li, Pan Ji, Fayao Liu, Hongdong Li, and Guosheng Lin. Prim2room: Layout-controllable room mesh generation from primitives. *arXiv preprint arXiv:2409.05380*, 2024. 2
- [22] Xiao Fu, Wei Yin, Mu Hu, Kaixuan Wang, Yuexin Ma, Ping Tan, Shaojie Shen, Dahua Lin, and Xiaoxiao Long. Geowizard: Unleashing the diffusion priors for 3d geometry estimation from a single image. *arXiv preprint arXiv:2403.12013*, 2024. 3

- [23] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. Get3d: A generative model of high quality 3d textured shapes learned from images. In *Advances In Neural Information Processing Systems*, 2022. 4
- [24] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 3
- [25] Anchit Gupta, Wenhan Xiong, Yixin Nie, Ian Jones, and Barlas Ögüz. 3dgen: Triplane latent diffusion for textured mesh generation, 2023. 2
- [26] Xianglong He, Junyi Chen, Sida Peng, Di Huang, Yangguang Li, Xiaoshui Huang, Chun Yuan, Wanli Ouyang, and Tong He. Gvgen: Text-to-3d generation with volumetric representation. *arXiv preprint arXiv:2403.12957*, 2024. 3, 6, 7
- [27] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718*, 2021. 6
- [28] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 17
- [29] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3, 17, 19
- [30] Lukas Höllein, Ang Cao, Andrew Owens, Justin Johnson, and Matthias Nießner. Text2room: Extracting textured 3d meshes from 2d text-to-image models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7909–7920, 2023. 2
- [31] Fangzhou Hong, Jiuxiang Tang, Ziang Cao, Min Shi, Tong Wu, Zhaoxi Chen, Shuai Yang, Tengfei Wang, Liang Pan, Dahua Lin, et al. 3dtopia: Large text-to-3d generation model with hybrid diffusion priors. *arXiv preprint arXiv:2403.02234*, 2024. 5, 6, 7
- [32] Yang Hong, Bo Peng, Haiyao Xiao, Ligang Liu, and Juyong Zhang. Headnerf: A real-time nerf-based parametric head model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20374–20384, 2022. 2
- [33] 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, 2023. 3
- [34] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. 5
- [35] Tao Hu, Fangzhou Hong, and Ziwei Liu. Structldm: Structured latent diffusion for 3d human generation. *arXiv preprint arXiv:2404.01241*, 2024. 2, 3
- [36] Zhipeng Hu, Minda Zhao, Chaoyi Zhao, Xinyue Liang, Lincheng Li, Zeng Zhao, Changjie Fan, Xiaowei Zhou, and Xin Yu. Efficientdreamer: High-fidelity and robust 3d creation via orthogonal-view diffusion priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4949–4958, 2024. 3
- [37] Shuo Huang, Shikun Sun, Zixuan Wang, Xiaoyu Qin, Yanmin Xiong, Yuan Zhang, Pengfei Wan, Di Zhang, and Jia Jia. Placiddreamer: Advancing harmony in text-to-3d generation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 6880–6889, 2024. 3, 5
- [38] Tianyu Huang, Yihan Zeng, Bowen Dong, Hang Xu, Songcen Xu, Rynson WH Lau, and Wangmeng Zuo. Textfield3d: Towards enhancing open-vocabulary 3d generation with noisy text fields. *arXiv preprint arXiv:2309.17175*, 2023. 2
- [39] Heewoo Jun and Alex Nichol. Shap-e: Generating conditional 3d implicit functions, 2023. 6, 7
- [40] Animesh Karnewar, Andrea Vedaldi, David Novotny, and Niloy Mitra. Holodiffusion: Training a 3D diffusion model using 2D images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023. 2
- [41] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35:26565–26577, 2022. 3
- [42] Oren Katzir, Or Patashnik, Daniel Cohen-Or, and Dani Lischinski. Noise-free score distillation. *arXiv preprint arXiv:2310.17590*, 2023. 3, 5
- [43] Oren Katzir, Or Patashnik, Daniel Cohen-Or, and Dani Lischinski. Noise-free score distillation, 2023. 2
- [44] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9492–9502, 2024. 3
- [45] 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. 3
- [46] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. 3
- [47] Seung Wook Kim, Bradley Brown, Kangxue Yin, Karsten Kreis, Katja Schwarz, Daiqing Li, Robin Rombach, Antonio Torralba, and Sanja Fidler. Neuralfield-ldm: Scene generation with hierarchical latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8496–8506, 2023. 21
- [48] Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 3
- [49] Dawid J Kopiczko, Tijmen Blankevoort, and Yuki M Asano. Vera: Vector-based random matrix adaptation. *arXiv preprint arXiv:2310.11454*, 2023. 5
- [50] Yushi Lan, Fangzhou Hong, Shuai Yang, Shangchen Zhou, Xuyi Meng, Bo Dai, Xingang Pan, and Chen Change Loy. Ln3diff: Scalable latent neural fields diffusion for speedy 3d generation, 2024. 2, 3, 6, 7
- [51] Yushi Lan, Shangchen Zhou, Zhaoyang Lyu, Fangzhou Hong, Shuai Yang, Bo Dai, Xingang Pan, and Chen Change

- Loy. Gaussiananything: Interactive point cloud latent diffusion for 3d generation. *arXiv preprint arXiv:2411.08033*, 2024. 3
- [52] Ming Li, Pan Zhou, Jia-Wei Liu, Jussi Keppo, Min Lin, Shuicheng Yan, and Xiangyu Xu. Instant3d: Instant text-to-3d generation. *International Journal of Computer Vision*, pages 1–17, 2024. 2
- [53] Weiyu Li, Rui Chen, Xuelin Chen, and Ping Tan. Sweetdreamer: Aligning geometric priors in 2d diffusion for consistent text-to-3d. *arxiv:2310.02596*, 2023. 2
- [54] Weiyu Li, Jiarui Liu, Rui Chen, Yixun Liang, Xuelin Chen, Ping Tan, and Xiaoxiao Long. Craftsman: High-fidelity mesh generation with 3d native generation and interactive geometry refiner, 2024. 3
- [55] Xinyang Li, Zhangyu Lai, Linning Xu, Jianfei Guo, Liujuan Cao, Shengchuan Zhang, Bo Dai, and Rongrong Ji. Dual3d: Efficient and consistent text-to-3d generation with dual-mode multi-view latent diffusion. *arXiv preprint arXiv:2405.09874*, 2024. 2, 4
- [56] Zongrui Li, Minghui Hu, Qian Zheng, and Xudong Jiang. Connecting consistency distillation to score distillation for text-to-3d generation. In *European Conference on Computer Vision*, pages 274–291. Springer, 2025. 3, 5
- [57] Yixun Liang, Xin Yang, Jiantao Lin, Haodong Li, Xiaogang Xu, and Yingcong Chen. Luciddreamer: Towards high-fidelity text-to-3d generation via interval score matching, 2023. 2, 3, 5
- [58] Tingting Liao, Xiaomei Zhang, Yuliang Xiu, Hongwei Yi, Xudong Liu, Guo-Jun Qi, Yong Zhang, Xuan Wang, Xiangyu Zhu, and Zhen Lei. High-fidelity clothed avatar reconstruction from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8662–8672, 2023. 2
- [59] Chenguo Lin, Panwang Pan, Bangbang Yang, Zeming Li, and Yadong Mu. Diffsplat: Repurposing image diffusion models for scalable gaussian splat generation. *arXiv preprint arXiv:2501.16764*, 2025. 3
- [60] 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 3, 5
- [61] Hongyu Liu, Xuan Wang, Ziyu Wan, Yujun Shen, Yibing Song, Jing Liao, and Qifeng Chen. Headartist: Text-conditioned 3d head generation with self score distillation. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–12, 2024. 2
- [62] Qihao Liu, Yi Zhang, Song Bai, Adam Kortylewski, and Alan Yuille. Direct-3d: Learning direct text-to-3d generation on massive noisy 3d data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6881–6891, 2024. 2, 6
- [63] 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, 3
- [64] Xiangyu Liu, Xiaomei Zhang, Zhiyuan Ma, Xiangyu Zhu, and Zhen Lei. Mvboost: Boost 3d reconstruction with multi-view refinement. *arXiv preprint arXiv:2411.17772*, 2024. 3
- [65] 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. 3
- [66] Yibo Liu, Zheyuan Yang, Guile Wu, Yuan Ren, Kejian Lin, Bingbing Liu, Yang Liu, and Jinjun Shan. Vqa-diff: Exploiting vqa and diffusion for zero-shot image-to-3d vehicle asset generation in autonomous driving. *arXiv preprint arXiv:2407.06516*, 2024. 2
- [67] Ying-Tian Liu, Yuan-Chen Guo, Guan Luo, Heyi Sun, Wei Yin, and Song-Hai Zhang. Pi3d: Efficient text-to-3d generation with pseudo-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19915–19924, 2024. 2, 3, 4, 6, 7, 19, 20
- [68] 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, 2023. 2
- [69] 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
- [70] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, and Wenping Wang. Wonder3d: Single image to 3d using cross-domain diffusion, 2023. 2, 3, 4
- [71] Raphael Gontijo Lopes, Stefano Fenu, and Thad Starner. Data-free knowledge distillation for deep neural networks. *arXiv preprint arXiv:1710.07535*, 2017. 2
- [72] 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, 3, 5
- [73] Baorui Ma, Haoge Deng, Junsheng Zhou, Yu-Shen Liu, Tiejun Huang, and Xinlong Wang. Geodream: Disentangling 2d and geometric priors for high-fidelity and consistent 3d generation. *arXiv preprint arXiv:2311.17971*, 2023. 2
- [74] Zhiyuan Ma, Xiangyu Zhu, Guo-Jun Qi, Zhen Lei, and Lei Zhang. Otavatar: One-shot talking face avatar with controllable tri-plane rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16901–16910, 2023. 2
- [75] Zhiyuan Ma, Yuxiang Wei, Yabin Zhang, Xiangyu Zhu, Zhen Lei, and Lei Zhang. Scaledreamer: Scalable text-to-3d synthesis with asynchronous score distillation. *arXiv preprint arXiv:2407.02040*, 2024. 2, 3, 5, 7, 17
- [76] Zhiyuan Ma, Xiangyu Zhu, Guojun Qi, Chen Qian, Zhaoxiang Zhang, and Zhen Lei. Diffspeaker: Speech-driven 3d

- facial animation with diffusion transformer. *arXiv preprint arXiv:2402.05712*, 2024. 3
- [77] Quan Meng, Lei Li, Matthias Nießner, and Angela Dai. Lt3sd: Latent trees for 3d scene diffusion. *arXiv preprint arXiv:2409.08215*, 2024. 2
- [78] Antoine Mercier, Ramin Nakhli, Mahesh Reddy, Rajeev Yasarla, Hong Cai, Fatih Porikli, and Guillaume Berger. Hexagen3d: Stablediffusion is just one step away from fast and diverse text-to-3d generation, 2024. 2, 3, 4, 6, 19, 20
- [79] 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. 3
- [80] Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulò, Peter Kotschieder, and Matthias Nießner. Diffrr: Rendering-guided 3d radiance field diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4328–4338, 2023. 2
- [81] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11453–11464, 2021. 3
- [82] Evangelos Ntavelis, Aliaksandr Siarohin, Kyle Olszewski, Chaoyang Wang, Luc Van Gool, and Sergey Tulyakov. Autodecoding latent 3d diffusion models, 2023. 2
- [83] Xiaqing Pan, Nicholas Charron, Yongqian Yang, Scott Peters, Thomas Whelan, Chen Kong, Omkar Parkhi, Richard Newcombe, and Carl Yuheng Ren. Aria digital twin: A new benchmark dataset for egocentric 3d machine perception, 2023. 2, 3
- [84] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion, 2022. 2, 3, 5, 6, 17
- [85] Mihir Prabhudesai, Anirudh Goyal, Deepak Pathak, and Katerina Fragkiadaki. Aligning text-to-image diffusion models with reward backpropagation. *arXiv preprint arXiv:2310.03739*, 2023. 4
- [86] Lingteng Qiu, Guanying Chen, Xiaodong Gu, Qi zuo, Mutian Xu, Yushuang Wu, Weihao Yuan, Zilong Dong, Liefeng Bo, and Xiaoguang Han. Richdreamer: A generalizable normal-depth diffusion model for detail richness in text-to-3d. *arXiv preprint arXiv:2311.16918*, 2023. 2, 3, 5, 17, 20
- [87] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 6
- [88] Barbara Roessle, Norman Müller, Lorenzo Porzi, Samuel Rota Bulò, Peter Kotschieder, Angela Dai, and Matthias Nießner. L3dg: Latent 3d gaussian diffusion. *arXiv preprint arXiv:2410.13530*, 2024. 2
- [89] 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. 1, 2, 4, 17, 19, 20
- [90] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. Stable-diffusion-v2.1-base. <https://huggingface.co/stabilityai/stable-diffusion-2-1-base>, 2022. 5
- [91] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 84–93, 2020. 2
- [92] Kyle Sargent, Zizhang Li, Tanmay Shah, Charles Herrmann, Hong-Xing Yu, Yunzhi Zhang, Eric Ryan Chan, Dmitry Lagun, Li Fei-Fei, Deqing Sun, and Jiajun Wu. Zerovs: Zero-shot 360-degree view synthesis from a single real image, 2023. 3
- [93] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *Advances in Neural Information Processing Systems*, 33:20154–20166, 2020. 3
- [94] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. *Advances in Neural Information Processing Systems*, 34:6087–6101, 2021. 3, 4
- [95] 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, 2023. 3, 4
- [96] Yichun Shi, Peng Wang, Jiandong Ye, Long Mai, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. *arXiv:2308.16512*, 2023. 2, 3, 4, 5, 17, 20
- [97] Sanghyun Son, Matheus Gadelha, Yang Zhou, Zexiang Xu, Ming C. Lin, and Yi Zhou. Dmesh: A differentiable representation for general meshes, 2024. 3
- [98] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 19
- [99] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 3
- [100] Jingxiang Sun, Xuan Wang, Yichun Shi, Lizhen Wang, Jue Wang, and Yebin Liu. Ide-3d: Interactive disentangled editing for high-resolution 3d-aware portrait synthesis. *ACM Transactions on Graphics (ToG)*, 41(6):1–10, 2022. 4
- [101] Jingxiang Sun, Xuan Wang, Lizhen Wang, Xiaoyu Li, Yong Zhang, Hongwen Zhang, and Yebin Liu. Next3d: Generative neural texture rasterization for 3d-aware head avatars. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 20991–21002, 2023. 2
- [102] Lingchen Sun, Rongyuan Wu, Jie Liang, Zhengqiang Zhang, Hongwei Yong, and Lei Zhang. Improving the stability and efficiency of diffusion models for content consistent super-resolution. *arXiv preprint arXiv:2401.00877*, 2023. 3

- [103] Lingchen Sun, Rongyuan Wu, Zhiyuan Ma, Shuaizheng Liu, Qiaosi Yi, and Lei Zhang. Pixel-level and semantic-level adjustable super-resolution: A dual-lora approach. *arXiv preprint arXiv:2412.03017*, 2024. 17
- [104] Yichun Tai, Zhenzhen Huang, Tao Peng, and Zhijiang Zhang. Deffiller: Mask-conditioned diffusion for salient steel surface defect generation. *arXiv preprint arXiv:2412.15570*, 2024.
- [105] Yichun Tai, Kun Yang, Tao Peng, Zhenzhen Huang, and Zhijiang Zhang. Defect image sample generation with diffusion prior for steel surface defect recognition. *IEEE Transactions on Automation Science and Engineering*, 22: 8239–8251, 2025. 3
- [106] Boshi Tang, Jianan Wang, Zhiyong Wu, and Lei Zhang. Stable score distillation for high-quality 3d generation, 2023. 3
- [107] Jiaxiang Tang. Stable-dreamfusion: Text-to-3d with stable-diffusion, 2022. <https://github.com/ashawkey/stable-dreamfusion>. 2, 3
- [108] Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023. 2, 3, 5
- [109] Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm: Large multi-view gaussian model for high-resolution 3d content creation, 2024. 3, 6, 7
- [110] Jiapeng Tang, Yinyu Nie, Lev Markhasin, Angela Dai, Justus Thies, and Matthias Nießner. Diffuscene: Denoising diffusion models for generative indoor scene synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 20507–20518, 2024. 2
- [111] Zhicong Tang, Shuyang Gu, Chunyu Wang, Ting Zhang, Jianmin Bao, Dong Chen, and Baining Guo. Volumediffusion: Flexible text-to-3d generation with efficient volumetric encoder. *arXiv preprint arXiv:2312.11459*, 2023. 3
- [112] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 4
- [113] Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitrii Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. SV3D: Novel multi-view synthesis and 3D generation from a single image using latent video diffusion. *arXiv*, 2024. 2
- [114] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A. Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. *arXiv preprint arXiv:2212.00774*, 2022. 2, 3
- [115] Peng Wang and Yichun Shi. Imagedream: Image-prompt multi-view diffusion for 3d generation. *arXiv preprint arXiv:2312.02201*, 2023. 4, 5
- [116] 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. 9, 17
- [117] Peihao Wang, Zhiwen Fan, Dejia Xu, Dilin Wang, Sreyas Mohan, Forrest Iandola, Rakesh Ranjan, Yilei Li, Qiang Liu, Zhangyang Wang, et al. Steindreamer: Variance reduction for text-to-3d score distillation via stein identity. *arXiv preprint arXiv:2401.00604*, 2023. 3
- [118] Tengfei Wang, Bo Zhang, Ting Zhang, Shuyang Gu, Jianmin Bao, Tadas Baltrusaitis, Jingjing Shen, Dong Chen, Fang Wen, Qifeng Chen, et al. Rodin: A generative model for sculpting 3d digital avatars using diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4563–4573, 2023. 3
- [119] 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. *arXiv preprint arXiv:2305.16213*, 2023. 2, 3, 5
- [120] 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, 2024. 3
- [121] Min Wei, Jingkai Zhou, Junyao Sun, and Xuesong Zhang. Adversarial score distillation: When score distillation meets gan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8131–8141, 2024. 3, 5
- [122] 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. 2, 4, 5, 9, 17
- [123] Bin-Shih Wu, Hong-En Chen, Sheng-Yu Huang, and Yu-Chiang Frank Wang. Tpa3d: Triplane attention for fast text-to-3d generation. In *European Conference on Computer Vision*, pages 438–455. Springer, 2024. 4
- [124] Rongyuan Wu, Lingchen Sun, Zhiyuan Ma, and Lei Zhang. One-step effective diffusion network for real-world image super-resolution. *Advances in Neural Information Processing Systems*, 37:92529–92553, 2024. 3, 17
- [125] Rongyuan Wu, Tao Yang, Lingchen Sun, Zhengqiang Zhang, Shuai Li, and Lei Zhang. Seesr: Towards semantics-aware real-world image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 25456–25467, 2024. 3
- [126] Shuang Wu, Youtian Lin, Feihu Zhang, Yifei Zeng, Jingxi Xu, Philip Torr, Xun Cao, and Yao Yao. Direct3d: Scalable image-to-3d generation via 3d latent diffusion transformer. *arXiv preprint arXiv:2405.14832*, 2024. 3, 7
- [127] Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Jiawei Ren, Liang Pan, Wayne Wu, Lei Yang, Jiaqi Wang, Chen Qian, et al. Omniobject3d: Large-vocabulary 3d object dataset for realistic perception, reconstruction and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 803–814, 2023. 2, 3, 7
- [128] Xiaoshi Wu, Yiming Hao, Manyuan Zhang, Keqiang Sun, Zhaoyang Huang, Guanglu Song, Yu Liu, and Hongsheng Li. Deep reward supervisions for tuning text-to-image diffusion models. *arXiv preprint arXiv:2405.00760*, 2024. 4
- [129] Yue Wu, Sicheng Xu, Jianfeng Xiang, Fangyun Wei, Qifeng Chen, Jiaolong Yang, and Xin Tong. Aniporportraitgan:

- animatable 3d portrait generation from 2d image collections. In *SIGGRAPH Asia 2023 Conference Papers*, pages 1–9, 2023. 2
- [130] Zhennan Wu, Yang Li, Han Yan, Taizhang Shang, Weixuan Sun, Senbo Wang, Ruikai Cui, Weizhe Liu, Hiroyuki Sato, Hongdong Li, et al. Blockfusion: Expandable 3d scene generation using latent tri-plane extrapolation. *ACM Transactions on Graphics (TOG)*, 43(4):1–17, 2024. 2
- [131] Zike Wu, Pan Zhou, Xuanyu Yi, Xiaoding Yuan, and Hanwang Zhang. Consistent3d: Towards consistent high-fidelity text-to-3d generation with deterministic sampling prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9892–9902, 2024. 3, 5
- [132] Jianfeng Xiang, Jiaolong Yang, Yu Deng, and Xin Tong. Gram-hd: 3d-consistent image generation at high resolution with generative radiance manifolds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2195–2205, 2023. 2
- [133] Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. *arXiv preprint arXiv:2412.01506*, 2024. 3
- [134] 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, 2024. 2, 6, 19, 20
- [135] Chao Xu, Ang Li, Linghao Chen, Yulin Liu, Ruoxi Shi, Hao Su, and Minghua Liu. Sparp: Fast 3d object reconstruction and pose estimation from sparse views. In *European Conference on Computer Vision*, pages 143–163. Springer, 2025. 3
- [136] 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. 5, 17
- [137] Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36, 2024. 4
- [138] Yinghao Xu, Wang Yifan, Alexander W Bergman, Menglei Chai, Bolei Zhou, and Gordon Wetzstein. Efficient 3d articulated human generation with layered surface volumes. *arXiv preprint arXiv:2307.05462*, 2023. 2
- [139] Yinghao Xu, Zifan Shi, Wang Yifan, Sida Peng, Ceyuan Yang, Yujun Shen, and Wetzstein Gordon. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation. *arxiv: 2403.14621*, 2024. 3
- [140] Haibo Yang, Yang Chen, Yingwei Pan, Ting Yao, Zhiheng Chen, Zuxuan Wu, Yu-Gang Jiang, and Tao Mei. Dreammesh: Jointly manipulating and texturing triangle meshes for text-to-3d generation. *arXiv preprint arXiv:2409.07454*, 2024. 3
- [141] Tao Yang, Rongyuan Wu, Peiran Ren, Xuansong Xie, and Lei Zhang. Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization. In *European Conference on Computer Vision*, pages 74–91. Springer, 2024. 3
- [142] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. 2, 5, 17, 21
- [143] Shih-Ying Yeh, Yu-Guan Hsieh, Zhidong Gao, Bernard BW Yang, Giyeong Oh, and Yanmin Gong. Navigating text-to-image customization: From lycoris fine-tuning to model evaluation. In *The Twelfth International Conference on Learning Representations*, 2023. 5
- [144] Taoran Yi, Jiemin Fang, Guanjun Wu, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Qi Tian, and Xinggang Wang. Gaussiandreamer: Fast generation from text to 3d gaussian splatting with point cloud priors. *arxiv:2310.08529*, 2023. 2, 3
- [145] Taoran Yi, Jiemin Fang, Zanwei Zhou, Junjie Wang, Guanjun Wu, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Xinggang Wang, and Qi Tian. Gaussiandreamerpro: Text to manipulable 3d gaussians with highly enhanced quality. *arXiv:2406.18462*, 2024. 3
- [146] Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6613–6623, 2024. 17
- [147] Xin Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Song-Hai Zhang, and Xiaojuan Qi. Text-to-3d with classifier score distillation. *arXiv preprint arXiv:2310.19415*, 2023. 3, 5
- [148] Xin Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Song-Hai Zhang, and Xiaojuan Qi. Text-to-3d with classifier score distillation, 2023. 2
- [149] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19447–19456, 2024. 3
- [150] Bowen Zhang, Yiji Cheng, Jiaolong Yang, Chunyu Wang, Feng Zhao, Yansong Tang, Dong Chen, and Baining Guo. Gaussiancube: A structured and explicit radiance representation for 3d generative modeling. *arXiv preprint arXiv:2403.19655*, 2024. 3
- [151] Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. Gs-irm: Large reconstruction model for 3d gaussian splatting. *arXiv*, 2024. 3
- [152] Lvmin Zhang and Maneesh Agrawala. Transparent image layer diffusion using latent transparency. *arXiv preprint arXiv:2402.17113*, 2024. 3
- [153] 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. *arXiv preprint arXiv:2406.13897*, 2024. 3

- [154] Qingru Zhang, Minshuo Chen, Alexander Bukharin, Nikos Karampatziakis, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. Adalora: Adaptive budget allocation for parameter-efficient fine-tuning. *arXiv preprint arXiv:2303.10512*, 2023. 5
- [155] Chenxi Zheng, Yihong Lin, Bangzhen Liu, Xuemiao Xu, Yongwei Nie, and Shengfeng He. Recdreamer: Consistent text-to-3d generation via uniform score distillation. In *The Thirteenth International Conference on Learning Representations*, 2025. 3
- [156] Junsheng Zhou, Weiqi Zhang, and Yu-Shen Liu. Diffgs: Functional gaussian splatting diffusion. *Advances in Neural Information Processing Systems*, 37:37535–37560, 2024. 3
- [157] Mingyuan Zhou, Huangjie Zheng, Zhendong Wang, Mingzhang Yin, and Hai Huang. Score identity distillation: Exponentially fast distillation of pretrained diffusion models for one-step generation. In *Forty-first International Conference on Machine Learning*, 2024. 3
- [158] Junzhe Zhu and Peiye Zhuang. Hifa: High-fidelity text-to-3d with advanced diffusion guidance, 2023. 2
- [159] Zhuangdi Zhu, Junyuan Hong, and Jiayu Zhou. Data-free knowledge distillation for heterogeneous federated learning. In *International conference on machine learning*, pages 12878–12889. PMLR, 2021. 2
- [160] Wenjie Zhuo, Fan Ma, Hehe Fan, and Yi Yang. Vividreamer: invariant score distillation for hyper-realistic text-to-3d generation. In *European Conference on Computer Vision*, pages 122–139. Springer, 2024. 3

Supplementary Material to “Progressive Rendering Distillation: Adapting Stable Diffusion for Instant Text-to-Mesh Generation without 3D Data”

The contents of this supplementary file include:

- Progressive Rendering Distillation pseudo code (referring to Sec. 3.2 in the main paper).
- More implementation details (referring to Sec. 4.1 in the main paper).
- Additional qualitative comparisons (referring to Sec. 4.2 in the main paper).
- Additional results with expanded training corpus (referring to Sec. 4.2 in the main paper).
- Additional ablation experiments (referring to Sec. 4.3 in the main paper).

A. Pseudo Code

The pseudo-code of our Progressive Rendering Distillation (PRD) training scheme is appended in Algorithm 1.

B. More Implementation Details

Dual rendering. We integrate DiffMC [122] for mesh rasterization and NeuS [116] for volume rendering to supervise the generation of 3D outputs. Such a dual rendering approach can ensure the training stability: when SDF values are all positive or all negative throughout the 3D space and thus the mesh extraction fails, volume rendering can still guide the training process to optimize the 3D space. Due to memory constraints, volume rendering is limited to low resolution (128×128). We complement this with high-resolution (512×512) mesh rasterization. To handle mesh extraction failures caused by the uniformly distributed SDF signs, we implement the method proposed in [136] to artificially enforce the position of the zero-level set in the 3D space. We manually control gradient magnitudes during backpropagation. The gradient of volume rendered multi-views with respect to the texture decoding MLP starts at 1.0 and linearly decreases to 0.01 at the end of training, preventing blurry textures caused by low-resolution volume rendering supervision. The gradient of mesh rasterized multi-views with respect to both the SDF decoding MLP and deformation decoding MLP is fixed at 0.001 throughout training, which stabilizes training and improves generation performance.

Training objective. With the multi-view teacher [86, 89, 96], we decode the multi-views \mathbf{x}_π to latent \mathbf{z}_π , which are diffused by adding Gaussian noise at timestep t [29], denoted by $\mathbf{z}_{\pi,t}$. We write the diffusion module of the multi-view teacher as $\mathbf{z}_{\phi^{2D}}(\mathbf{z}_{\pi,t}; t, \pi, y)$ to represent the process of noise prediction and latent denoising, where y is the text prompt. With ASD [75], the derivative of the objective with respect to the 3D generator ϕ^{3D} is:

$$\nabla_{\phi^{3D}} \mathcal{L}_{\phi^{2D}}(\mathbf{x}_\pi; \pi, y) = \mathbb{E}_{t, \epsilon, \Delta t} \left[\omega(t) \left(\mathbf{z}_{\phi^{2D}}^{Cls}(\mathbf{z}_{\pi,t}; t, \pi, y) - \mathbf{z}_{\phi^{2D}}(\mathbf{z}_{\pi,t+\Delta t}; t + \Delta t, \pi, y) \right) \frac{\partial \mathbf{z}_\pi}{\partial \phi^{3D}} \right], \quad (1)$$

where ϕ^{2D} denotes the teacher model parameters, t is sampled from $\mathcal{U}[T_{\text{Min}}, T_{\text{Max}}]$ with $0 < T_{\text{Min}} < T_{\text{Max}} < T = 1000$, and Cls indicates classifier-free guidance (CFG) [28]. By introducing a timestep shift Δt [75] sampled from a uniform distribution $\mathcal{U}[0, \eta(t - T_{\text{Min}})]$, ASD achieves more effective training of the native 3D generator. We utilize the timestep-dependent weighting factor from DMD [146], as implemented in [103, 124]. We let

$$\omega(t) = \frac{1}{\text{NoGrad}(\text{Mean}(\mathbf{z}_\pi - \mathbf{z}_{\phi^{2D}}^{Cls}(\mathbf{z}_{\pi,t}; t, \pi, y))) + \delta}, \quad (2)$$

where NoGrad detaches gradients for loss back-propagation, and Mean applies L_1 -norm across height, width, channel dimensions and all rendered views. Unlike [103, 124, 146], we add constant $\delta = 0.1$ to the denominator, which stabilizes training and improves generation performance. We apply this objective function to supervise 3D outputs using three teacher models (SD, MV, RD) and two rendering pipelines (volume rendering and mesh rasterization). Regarding the sampling range of timestep t , $T_{\text{Max}} = 980$ throughout training, while T_{Min} starts at 500 and linearly decreases to 20. Teacher-specific hyperparameters vary: RD uses CFG=20 and $\eta = 0.1$; MV uses CFG decreasing from 20 to 10 and $\eta = 0$; SD uses CFG=5 and $\eta = 0$. Setting $\eta = 0$ for multi-view teachers that supervise RGB renderings aligns with the findings in PiSA-SR [103]. Additionally, we incorporate regularization terms during training, such as sparsity loss [84] and eikonal loss [142]. We linearly reduce the sparsity and eikonal loss weights from 1 to 0 throughout the training process.



Figure 8. More results of our model trained with expanded corpus.

Algorithm 1 Progressive Rendering Distillation (PRD)

Input: SD-based native 3D generator with $z_{\phi^{3D}}$ and $D_{\phi^{3D}}$; score distillation objective $\mathcal{L}_{\phi^{2D}}$ parameterized by multi-view diffusion model ϕ^{2D} ; prompt corpus \mathbb{S}_y ; number of rendered views N ; number of steps K

```
1 Initialize optimizer Opt for  $z_{\phi^{3D}}$  and  $D_{\phi^{3D}}$ 
2 Define fixed timesteps  $T = t_1 > t_2 > \dots > t_K > 0$ 
3 while not converged do
4   Sample text prompt  $y \in \mathbb{S}_y$ 
5   Sample  $\hat{z}_0 \sim \mathcal{N}(0, I)$ 
6   for  $t \leftarrow t_1$  to  $t_K$  do
7     Sample  $\epsilon \sim \mathcal{N}(0, I)$ 
8      $z_t \leftarrow \alpha_t \hat{z}_0 + \sigma_t \epsilon$ 
9      $\hat{z}_0 \leftarrow z_{\phi^{3D}}(z_t; t, y)$ 
10     $\hat{\theta} \leftarrow D_{\phi^{3D}}(\hat{z}_0)$ 
11    Sample  $K$  camera poses  $\pi_1, \dots, \pi_N$ 
12    for  $i \leftarrow 1$  to  $N$  do
13       $x_{\pi_i} \leftarrow g(\hat{\theta}, \pi_i)$ 
14    end
15     $L \leftarrow \mathcal{L}_{\phi^{2D}}(x_{\pi_1}, \dots, x_{\pi_N}; \pi_1, \dots, \pi_N, y)$ 
16    Save  $\frac{1}{K} \nabla_{\phi^{3D}} L$  in Opt
17  end
18  Update  $z_{\phi^{3D}}$  and  $D_{\phi^{3D}}$  with gradient saved in Opt
19 end
20 return  $z_{\phi^{3D}}$  and  $D_{\phi^{3D}}$ 
```




Figure 9. Qualitative comparison with LATTE3D [134].

Noise schedule. The PRD training incorporates progressive noise addition to the denoised latents (see Line 8 in Algorithm 1). Being adapted from SD [89], our native 3D generator follows the DDPM [29] noise schedule in training. During inference, we employ DDIM [98].

C. More Qualitative Comparison Results

Comparison with methods adapting SD as native 3D generators. Since the codes or trained models of current SD-based native 3D generators [67, 78] are not publicly available, we conduct our comparisons by using their visual results presented in the original publications. The qualitative comparisons with PI3D [67] and HexaGen3D [78] are presented in Fig. 10 and Fig. 11, respectively. As both the two compared methods employ data-driven training, they inherit pose inconsistencies existed in the 3D training datasets [17], leading to the issue of occasional pose misalignment. This can be clearly observed from PI3D’s result of ‘A dalmatian wearing a fireman’s hat’ shown in Fig. 10, where the dog is oriented sideways. The comparison results demonstrate our method’s superior visual fidelity with the input prompts. These improvements are attributed



Figure 10. Qualitative comparison with PI3D [67].




Figure 11. Qualitative comparison with HexaGen3D [78].

to our proposed Progressive Rendering Distillation (PRD) scheme, which utilizes multi-view teachers in training without requiring 3D training data.

Comparison with native 3D generators trained with score distillation. We further compare our approach with existing methods that employ score distillation for native 3D generator training. Specifically, we compare against the current state-of-the-art method, LATTE3D [134]. Since the code or model of LATTE3D is unavailable, we conduct qualitative comparisons using their published results. The visual comparisons are presented in Fig. 9. It can be seen that our method demonstrates notable improvements in both texture fidelity and geometric accuracy. For example, in ‘A blue tulip’, our model captures more natural flower textures, while in ‘A pile of dice on a green tabletop’, our model achieves more precise geometric structures. These improvements can be attributed to our strategic adaptation of SD as the backbone architecture, which allows our model to leverage its powerful generative capabilities.

D. Expanding Training Corpus

Since our proposed training scheme does not require 3D ground truth data, it can be easily up-scaled to a large amount of text prompts. We collect a total number of 1.7 million text prompts from HuggingFace that were used to generate images by DALL-E and Midjourney. This corpus has more unnatural prompts than the Objaverse [17], and it is more challenging. To the best of our knowledge, our work is the first that can process more than 1 million training data. Our model, trained on this expanded dataset, achieves enhanced visual quality, as demonstrated in Fig. 1 in the main paper and Fig. 5, Fig. 8 in this supplementary file.

E. More Ablation Studies

The necessity of multiple teachers. We employ SD [89], MV [96] and RD [86] as teachers for multi-view supervision of RGB, normal and depth maps. Here we perform ablation studies by systematically removing individual components.

First, as visualized by **w/o SD** in Fig. 12, when SD is removed, leaving only MV and RD as teachers, the model can collapse to generate results inconsistent with text prompts. For example, given the prompt ‘A DSLR photo of a cracked egg with the yolk spilling out on a wooden table’, the model collapses to generating a stack of discs. This occurs because training for multi-view generation may impair MV and SD’s text understanding capabilities, resulting in outputs that diverge from the specified text descriptions. SD can prevent from training collapse and improve the generation stability. Second, the importance of MV is demonstrated by the visualizations of **w/o MV** in Fig. 12. Without multi-view RGB supervision, the generated results tend to show repetitive and redundant contents across different viewpoints. For instance, multiple ‘egg yolks’ might appear in the results of ‘A DSLR photo of a cracked egg with the yolk spilling out on a wooden table’. Finally, the importance of RD is validated by the visualizations of **w/o RD**. We can see that adding normal and depth constraints enhances text consistency in the outputs, such as the generated ‘wooden table’ in the results of ‘A DSLR photo of a cracked egg with the yolk spilling out on a wooden table’. Overall, the combination of SD, MV, and RD as teachers achieves the best results, as validated by the visualization of **w/ All** and the metrics shown in Tab. 4.

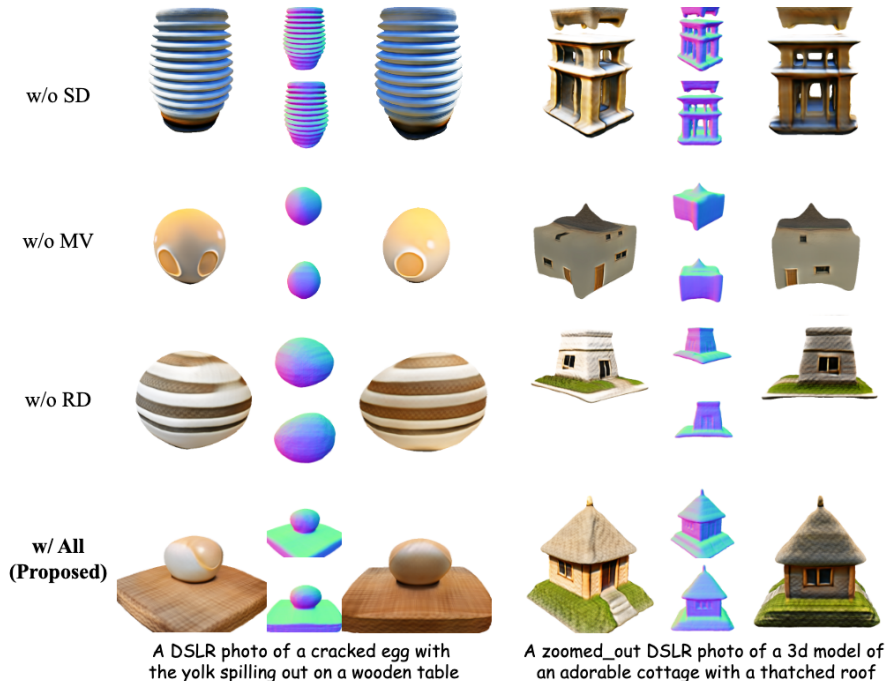


Figure 12. Visualizations for the ablation study on jointly using SD, MV and RD as multi-view teachers.

	C.S. \uparrow	R@1 \uparrow
w/o SD	63.0	20.1
w/o MV	67.4	25.9
w/o RD	41.5	11.4
w/ All (Proposed)	68.2	32.3

Table 4. Ablation study on jointly using SD, MV and RD as multi-view teachers.

The necessity of dual rendering. We use a dual rendering framework that integrates mesh rasterization [47] and volume rendering [142] for 3D output supervision, as detailed in Sec. B. The effectiveness of this dual approach is demonstrated through quantitative and qualitative evaluations in Tab. 5 and Fig. 13, respectively. Without volume rendering, relying solely on mesh rasterization leads to training collapse and invalid mesh extraction. The results labeled as **w/o Volume Rendering** in Fig. 13 demonstrate that training converges to a state where the SDF’s zero-level set vanishes, resulting in mesh extraction failure and empty space. Conversely, using only volume rendering, which is constrained to low-resolution training, fails to produce high-quality mesh geometry, leading to rough and coarse textural details, as shown by **w/o Mesh Rasterization** in Fig. 13. For example, it fails to produce the shining gold texture for the prompt ‘A DSLR photo of a toilet made out of gold’. Moreover, without direct mesh supervision, volume rendering-based methods may produce geometrically invalid structures. This limitation is evident in the result of ‘A DSLR photo of aerial view of a ruined castle’, where the extracted meshes exhibit incorrect structural features and poor textures, manifesting as gray regions in parts of the mesh. As shown by **w/ Both** in Fig. 13 and supported by the superior metrics in Tab. 5, our dual rendering approach enables stable training while producing meshes with detailed textures and well-defined geometric structures.

The choice of LoRA rank. We demonstrate the significance of using a LoRA rank of 16 in our Parameter-Efficient Triplane Adaption (PETA). With a lower rank of 8, shown as **Rank=8** in Fig. 14, the model exhibits insufficient learning capacity, as evidenced by its failure to generate the top hat structure for the prompt ‘A capybara wearing a top hat, low poly’. However, setting a higher rank, such as 32, can also lead to unreasonable geometric outputs. As shown in **Rank=32** in Fig. 14, some unwanted platform structures appear at the bottom of results of ‘A capybara wearing a top hat, low poly’ and ‘A zoomed out DSLR photo of a baby dragon’. Such artifacts stem from the inherent generation biases in both MV

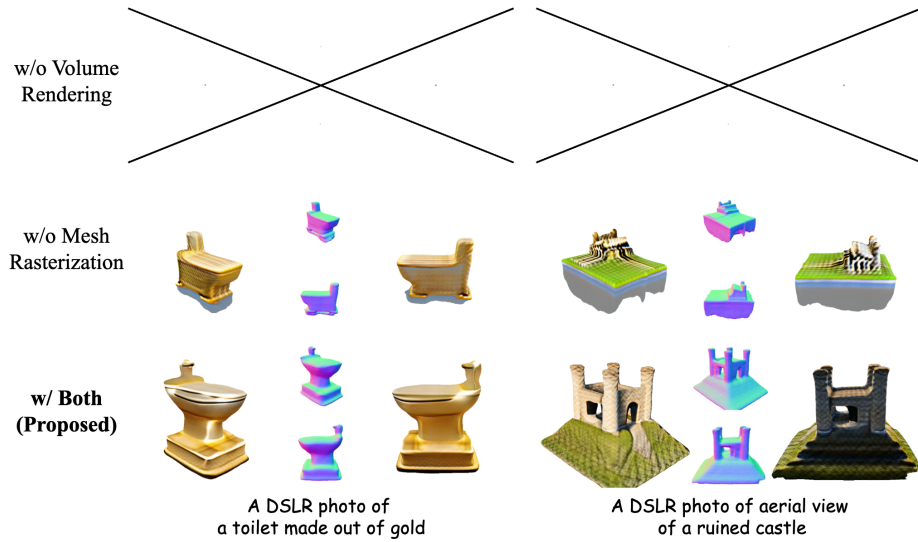


Figure 13. Ablation study on dual rendering. The cross mark means the model fails to generate mesh due to training instability.



Figure 14. Visualization for the ablation study on the LoRA rank in PETA.

and SD, as their training dataset [17] contains numerous examples where objects rest on square platforms. As a result, the multi-view teachers are fitted to generate outputs with similar structures. When the LoRA rank is set too high, the native 3D generator tends to learn and reproduce the biases from the teachers during the distillation. Setting the rank to a balanced value of 16 enables the model to generate text-aligned 3D results while avoiding the incorporation of undesirable biases into the 3D generation model. Denoted as **Rank=16**, both qualitative results in Fig. 14 and quantitative results in Fig. 14 show that a rank of 16 yields the best performance.

	C.S. \uparrow	R@1 \uparrow
w/o Volume Rendering	25.1	0.01
w/o Mesh Rasterization	67.4	25.9
Joint (Proposed)	68.2	32.3

Table 5. Ablation study on the dual renders.

	C.S. \uparrow	R@1 \uparrow
rank=8	62.9	15.6
rank=16 (Proposed)	68.2	32.3
rank=32	66.2	26.6

Table 6. Ablation study on the LoRA rank in PETA.