

Text-to-3D Generation with Bidirectional Diffusion using both 2D and 3D priors

Lihe Ding^{1,4*‡}, Shaocong Dong^{2*‡}, Zhanpeng Huang³, Zibin Wang^{3†},
Yiyuan Zhang¹, Kaixiong Gong¹, Dan Xu^{2†}, Tianfan Xue^{1†}

¹The Chinese University of Hong Kong ²Hong Kong University of Science and Technology
³SenseTime Research ⁴Shanghai AI Laboratory

{dl023, gk023, tfxue}@ie.cuhk.edu.hk, {sdongae, danxu}@cse.ust.hk
{wangzb02, yiyuanzhang.ai}@gmail.com, {huangzhanpeng}@sensetime.com

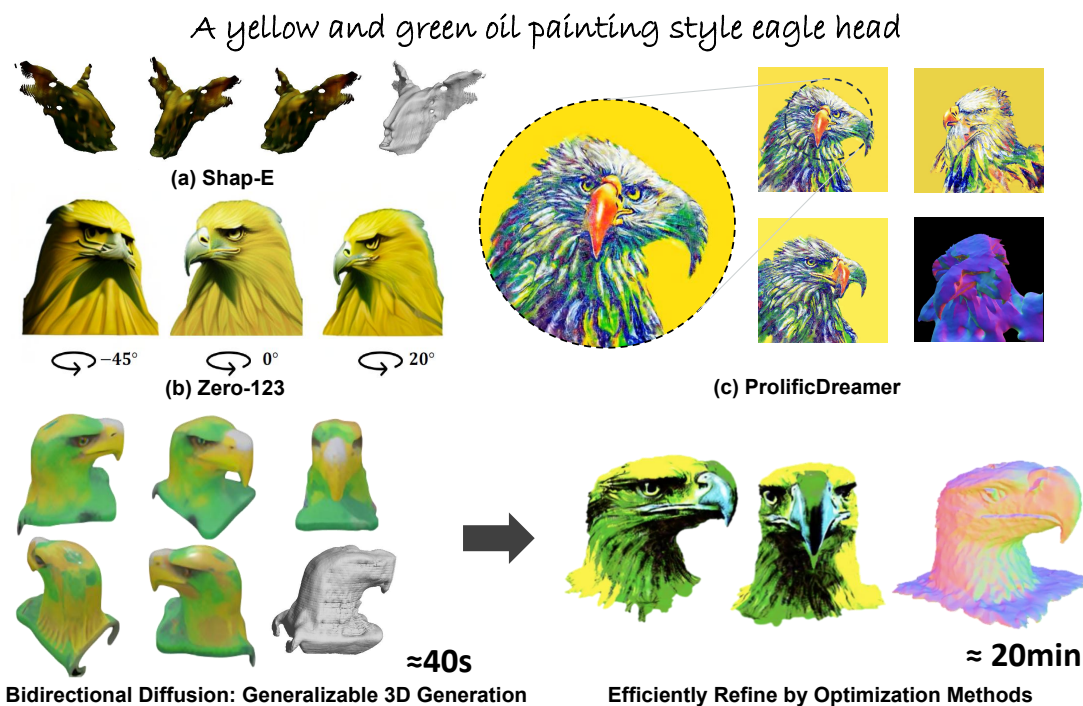


Figure 1. Our *BiDiff* can efficiently generate high-quality 3D objects. It alleviates all these issues in previous 3D generative models: (a) low-texture quality, (b) multi-view inconsistency, and (c) geometric incorrectness (e.g., multi-face Janus problem). The outputs of our model can be further combined with optimization-based methods (e.g., ProlificDreamer) to generate better 3D geometries with slightly longer processing time (bottom row).

Abstract

Most 3D generation research focuses on up-projecting 2D foundation models into the 3D space, either by minimizing 2D Score Distillation Sampling (SDS) loss or fine-tuning on multi-view datasets. Without explicit 3D priors, these methods often lead to geometric anomalies and multi-view

inconsistency. Recently, researchers have attempted to improve the genuineness of 3D objects by directly training on 3D datasets, albeit at the cost of low-quality texture generation due to the limited texture diversity in 3D datasets. To harness the advantages of both approaches, we propose Bidirectional Diffusion (*BiDiff*), a unified framework that incorporates both a 3D and a 2D diffusion process, to preserve both 3D fidelity and 2D texture richness, respectively. Moreover, as a simple combination may yield inconsistent generation results, we further bridge them with novel bidirectional guidance. In addition, our method can be used

*Equal contribution.

†Corresponding author.

‡Part of this work was done when Lihe Ding and Shaocong Dong interned at SenseTime.

as an initialization of optimization-based models to further improve the quality of 3D models and the efficiency of optimization, reducing the process from 3.4 hours to 20 minutes. Experimental results have shown that our model achieves high-quality, diverse, and scalable 3D generation. Project website <https://bidiff.github.io/>.

1. Introduction

Recent advancements in text-to-3D generation [24] mainly focus on lifting 2D foundation models into 3D space. One of the most popular solutions [15, 30] uses 2D Score Distillation Sampling (SDS) loss derived from a 2D diffusion model to supervise 3D generation. While these methods can generate high-quality textures, they often lead to geometric ambiguity, such as the multi-face Janus problem [25], due to the lack of 3D constraints (Fig. 1(c)). Moreover, these optimization methods are time-consuming, taking hours to generate one object. Zero-123[18] tries to alleviate the problem by fine-tuning the 2D diffusion models on multi-view datasets, but it still cannot guarantee geometric consistency (Fig. 1(b)).

To ensure better 3D consistency, another solution is to directly learn 3D structures from 3D datasets [12, 28]. However, many existing 3D datasets [1, 4] only contain hand-crafted objects or lack high-quality 3D geometries, with textures very different from real-world objects. Moreover, 3D datasets are often much smaller than, and also difficult to scale up to, their 2D counterparts. As a result, the 3D diffusion models (Fig. 1 (a)) normally cannot generate detailed textures and complicated geometries, even if they have better 3D consistency compared to up-projecting 2D diffusion models.

Therefore, a straightforward way to leverage the advantages of both methods is to combine both 2D and 3D diffusion models. However, a simple combination may result in inconsistent generative directions as they are learned in two independent diffusion processes. In addition, the two diffusion models are represented in separate 2D and 3D spaces without knowledge sharing.

To overcome these problems, we propose *Bidirectional Diffusion (BiDiff)*, a method to seamlessly integrate both 2D and 3D diffusion models within a unified framework. Specifically, we employ a hybrid representation in which a signed distance field (SDF) is used for 3D feature learning and multi-view images for 2D feature learning. The two representations are mutually transformable by rendering 3D feature volume into 2D features and back-projecting 2D features to 3D feature volume. Starting from pretrained 3D and 2D diffusion models, the two diffusion models are jointly finetuned to capture a joint 2D and 3D prior facilitating 3D generation.

However, correlating the 2D and 3D representations is

not enough to combine two diffusion processes, as they may deviate from each other in the following diffusion steps. To solve this problem, we further introduce bidirectional guidance to align the generative directions of the two diffusion models. At each diffusion step, the intermediate results from the 3D diffusion scheme are rendered into 2D images as guidance signals to the 2D diffusion model. Meanwhile, the multi-view intermediate results from the 2D diffusion process are also back-projected to 3D, guiding the 3D diffusion. The mutual guidance regularizes the two diffusion processes to learn in the same direction.

The proposed bidirectional diffusion poses several advantages over the previous 3D generation models. First, users can separately control the generation of 2D texture and 3D geometry, as shown in Fig. 2, because the 2D diffusion model focuses on texture generation and the 3D diffusion model focuses on geometry. This is impossible for previous 3D diffusion methods. Secondly, compared to 3D-only diffusion models [12], our method takes advantage of a 2D diffusion model trained on much larger datasets. Therefore, it can generate more diversified objects and create a completely new object like “A strong muscular chicken” illustrated in Fig 2. Thirdly, compared to previous optimization methods [30, 41] that often take several hours to generate one object, we utilize a fast feed-forward joint 2D-3D diffusion model for scalable generation, which only takes about 40 seconds to generate one object.

Moreover, because of the efficacy of *BiDiff*, we also propose an optional step to utilize its output as an initialization for the existing optimization-based methods (e.g., ProlificDreamer [41]). This optional step can further improve the quality of a 3D object, as demonstrated in the bottom row of Fig. 1. Also, the good initialization from *BiDiff* helps to reduce optimization time from around 3.4 hours to 20 minutes, and concurrently resolves geometrical inaccuracy issues, like multi-face anomalies. Moreover, this two-step generation enables creators to rapidly adjust prompts to obtain a satisfactory preliminary 3D model through a lightweight feed-forward generation process, subsequently refining it into high-fidelity results.

Through training on ShapeNet [1] and Objaverse 40K [4], our framework is shown to generate high-quality textured 3D objects with strong generalizability. In summary, our contributions are as follows: 1) We propose *BiDiff*, a joint 2D-3D diffusion model, that can generate high-quality, 3D-consistent, and diversified 3D objects; 2) We propose a novel training pipeline that utilizes both pretrained 2D and 3D generative foundation models; 3) We propose the first diffusion-based 3D generation model that allows independent control of texture and geometry; 4) We utilize the outputs from *BiDiff* as a strong initialization for the optimization-based methods, generating high-quality geometries while ensuring that users receive quick feedback for each prompt update.




Figure 2. **Texture Control (Top)**: we change the texture while maintaining the overall shape. **Shape Control (Bottom)**: we fix texture patterns and generate various shapes.

2. Related Work

Early 3D generative methods adopt various 3D representations, including 3D voxels [8, 37, 42], point clouds [29, 43], meshes [16, 26], and implicit functions [2, 11] for category-level 3D generations. These methods directly train the generative model on a small-scale 3D dataset, and, as a result, the generated objects may either miss tiny geometric structures or lose diversity. Although there are large-scale [4] or high-quality 3D datasets [38] recently, they are still much smaller than the datasets used for 2D image generation training.

With the powerful text-to-image synthesis models [32–34], a new paradigm emerges for 3D generation without large-scale 3D datasets by leveraging 2D generative model. One line of works utilizes 2D priors from pre-trained text-to-image model (known as CLIP) [10, 13] or 2D diffusion generative models [15, 24, 39] to guide the optimization of underlying 3D representations. However, these models could not guarantee cross-view 3D consistency and the per-instance optimization scheme suffers both high computational cost and over-saturated problems. Later on, researchers improve these models using textual codes or depth maps [5, 23, 35], and [41] directly model 3D distribution to improve diversity. These methods alleviate the visual artifacts but still cannot guarantee high-quality 3D results.

Another line of works learn 3D priors directly from 3D datasets. As the diffusion model has been the de-facto network backbone for most recent generative models, it has been adapted to learn 3D priors using implicit spaces such as point cloud features [28, 45], NeRF parameters [6, 12], or SDF spaces [3, 20]. The synthesized multi-view images rendered from 3D datasets were also utilized to provide cross-view 3D consistent knowledge, thereby achieving image-to-3D gener-

ation under given single image condition [18, 19, 22]. These methods normally highlight fast inference and 3D consistent results. However, due to inferior 3D dataset quality and size, these methods generally yield geometric lower-quality results with limited diversity. Recently a few methods [31, 36] explored to combine 2D priors and 3D priors from individual pre-trained diffusion models, but they often suffer from inconsistency between two generative processes.

3. Method

As many previous studies [18, 31] have illustrated, both 2D texture and 3D geometry are important for 3D object generation. However, incorporating 3D structural priors and 2D textural priors is challenging: i) combining both 3D and 2D generative models into a single cohesive framework is not trivial; ii) in both training and inference, two generative models may lead to opposite generative directions.

To tackle these problems, we propose *BiDiff*, a novel bidirectional diffusion model that marries a pretrained 3D diffusion model with another 2D one using bidirectional guidance. Fig. 3 illustrates the overall architecture of our framework. Details of each component will be discussed below. Specifically, in Sec. 3.1, we will introduce our novel hybrid representation that includes both 2D and 3D information, and the bidirectional diffusion model built on top of this hybrid representation. In Sec. 3.2 and Sec. 3.3, to ensure the two generative models lead to the same generative direction, we will introduce how to add bidirectional guidance to both 3D and 2D diffusion models. In Sec. 3.4, we discuss one advantage of *BiDiff*, which is independent control of texture and geometry generation, as shown in Fig. 2. Finally, in Sec. 3.5, we discuss another advantage of *BiDiff*, which is

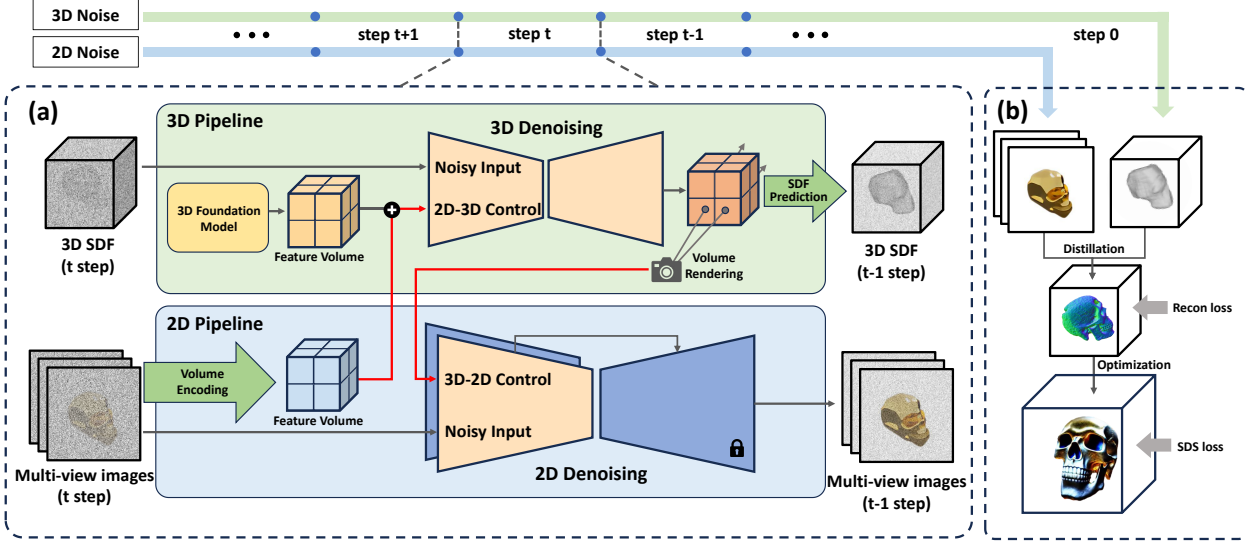


Figure 3. The *BiDiff* framework operates as follows: (a) At each step of diffusion, we render the 3D diffusion’s intermediate outputs into 2D images, which then guide the denoising of the 2D diffusion model. Simultaneously, the intermediate multi-view outputs from the 2D diffusion are re-projected to assist the denoising of the 3D diffusion model. Red arrows show the bidirectional guidance, which ensures that both diffusion processes evolve coherently. (b) We use the outcomes of the 2D-3D diffusion as a strong starting point for optimization methods, allowing for further refinement with fewer optimization steps.

to use the results from *BiDiff* as a strong initialization for optimization-based methods to obtain more delicate results efficiently.

3.1. Bidirectional Diffusion

To incorporate both 2D and 3D priors, we represent a 3D object using a hybrid combination of two formats: Signed Distance Field (SDF) \mathcal{F} and multi-view image set $\mathcal{V} = \{\mathcal{I}^i\}_{i=1}^M$, where \mathcal{F} is computed from signed distance values on an $N \times N \times N$ grid, and \mathcal{I}^i is the i -th image from a multi-view image set of size M . This hybrid representation is shown on the left side of Fig. 3.

With this representation, we learn a joint distribution $\{\mathcal{F}, \mathcal{V}\}$ utilizing two distinct diffusion models: a 3D diffusion model \mathcal{D}_{3d} in the SDF space (the green 3D denoising block in Fig. 3) and a 2D multi-view diffusion model \mathcal{D}_{2d} within the image domain (the blue 2D denoising block in Fig. 3). Specifically, given a timestep t , we add Gaussian noises to both SDF and multi-view images as

$$\begin{aligned} \mathcal{F}_t &= \sqrt{\bar{\alpha}_t} \mathcal{F}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon_{3d} \quad \text{and} \\ \mathcal{I}_t^i &= \sqrt{\bar{\alpha}_t} \mathcal{I}_0^i + \sqrt{1 - \bar{\alpha}_t} \epsilon_{2d}^i \quad \text{for } \forall i, \end{aligned} \quad (1)$$

where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ is random noise, and $\bar{\alpha}_t$ is a noise schedule which is different in 3D and 2D. Subsequently, the straightforward way is to separately train these two diffusion models by minimizing the following two objectives:

$$\begin{aligned} L_{simple3d} &= \mathbb{E}_{\mathcal{F}_t, \epsilon_{3d}, t} \|\epsilon_{3d} - \mathcal{D}_{3d}(\mathcal{F}_t, t)\|_2^2, \\ L_{simple2d} &= \frac{1}{N} \sum_{i=1}^N (\mathbb{E}_{\mathcal{I}_t^i, \epsilon_{2d}^i, t} \|\epsilon_{2d}^i - \mathcal{D}_{2d}(\mathcal{I}_t^i, t)\|_2^2), \end{aligned} \quad (2)$$

where ϵ_{3d} and ϵ_{2d} are Gaussian noises $\epsilon_{3d}, \epsilon_{2d}^i \sim \mathcal{N}(0, \mathbf{I})$, SDF and image set are sampled from forward diffusion processes $\mathcal{F}_t \sim q(\mathcal{F}_t), \mathcal{I}_t^i \sim q(\mathcal{I}_t^i)$, and timestep is uniformly sampled $t \sim U[1, T]$.

However, this simple combination does not consider the correlations between 3D and 2D diffusion, which may hinder the understanding of 2D and 3D consistency, leading to inconsistent generation between 3D geometry and 2D multi-view images.

We resolve this problem by a novel *Bidirectional Diffusion*. In this model, the consistency between 3D and 2D diffusion output is enforced through bidirectional guidance. First, we add guidance from the 2D diffusion process to the 3D generative process, which is the red arrow pointing to the “2D-3D Control”. Specifically, during each denoising step t , we feed the denoised multi-view images $\mathcal{V}'_{t+1} = \{\mathcal{I}_{t+1}^i\}_{i=1}^N$ in previous step into the 3D diffusion model as $\epsilon'_{3d} = \mathcal{D}_{3d}(\mathcal{F}_t, \mathcal{V}'_{t+1}, t)$. This guidance steers the current 3D denoising direction to ensure 2D-3D consistency. It’s worth mentioning that the denoised output \mathcal{V}'_{t+1} from the previous step $t + 1$ is inaccessible in training, therefore we directly substitute it with the ground truth \mathcal{V}_t . During inference, we utilize the denoised images from the previous step. Then we could obtain the denoised radiance field \mathcal{F}'_0 given the 3D guided noise prediction ϵ'_{3d} by $\mathcal{F}'_0 = \frac{1}{\sqrt{\bar{\alpha}_t}} (\mathcal{F}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon'_{3d})$.

Secondly, we also add guidance from the 3D diffusion process to the 2D generative process. Specifically, using the same camera poses, we render multi-view images \mathcal{H}_t^i derived from the radiance field \mathcal{F}'_0 by the 3D diffusion model: $\mathcal{H}_t^i = \mathcal{R}(\mathcal{F}'_0, \mathcal{P}^i), i = 1, \dots, M$, where \mathcal{P}^i is the i^{th} camera pose. These images are further used as guidance to the 2D multi-view denoising process \mathcal{D}_{2d} by $\epsilon'_{2d} = \mathcal{D}_{2d}(\mathcal{V}_t, \{\mathcal{H}_t^i\}_{i=1}^N, t)$. This guidance is the red arrow pointing to the “3D-2D control” in Fig. 3.

Our method can seamlessly integrate and synchronize both the 3D and 2D diffusion processes within a unified framework. In the following sections, we will delve into each component in detail.

3.2. 3D Diffusion Model with 2D Guidance

Our 3D diffusion model aims to generate a neural surface field (NeuS) [21], with novel 2D-to-3D guidance derived from the denoised 2D multi-view images. To train our 3D diffusion model, at each training timestep t , we add noise to a clean radiance field, yielding a noisy one \mathcal{F}_t . This field, combined with the timestep t embeddings and the text embeddings, is then passed through 3D sparse convolutions to generate a 3D feature volume \mathcal{M} as: $\mathcal{M} = \text{Sp3DConv}(\mathcal{F}_t, t, \text{text})$. Then we sample $N \times N \times N$ grid points from \mathcal{M} and project these points onto all denoised multi-view images \mathcal{V}'_{t+1} from the previous step of the 2D diffusion model. At each grid point p , we aggregate the interpolated 2D feature at its 2D projected location on each view, and calculate the mean and variance over all N interpolated features to obtain the image-conditioned feature volume \mathcal{N} :

$$\mathcal{N}(p) = [\text{Mean}(\mathcal{V}'_{t+1}(\pi(p))), \text{Var}(\mathcal{V}'_{t+1}(\pi(p)))], \quad (3)$$

where π denotes the projection operation from 3D to 2D image plane. We fuse these two feature volumes with further sparse convolutions for predicting the clean \mathcal{F}_0 .

One important design of our 3D diffusion model is that it incorporates geometry priors derived from the 3D foundation model, Shap-E [12]. Shap-E is a latent diffusion [24] model trained on several millions 3D objects, and thus ensures the genuineness of generated 3D objects. Still, we do not want Shap-E to limit the creativity of our 3D generative model, and try to preserve the capability of generating novel objects that Shap-E cannot.

To achieve this target, we design a feature volume \mathcal{G} to represent a radiance field converted from the Shap-E latent code \mathcal{C} . It is implemented using NeRF MLPs by setting their parameters to the latent code \mathcal{C} : $\mathcal{G}(p) = \text{MLP}(\lambda(p); \theta = \mathcal{C})$, where λ denotes the positional encoding operation.

One limitation of directly introducing Shap-E latent code is that the network is prone to shortcut the training process, effectively memorizing the radiance field derived from Shap-E. To generate 3D objects beyond Shap-E model, we add Gaussian noise at level t_0 to the clean latent code, resulting in

the noisy latent representation \mathcal{C}_{t_0} , where t_0 represents a pre-defined constant timestep. Subsequently, the noisy radiance field \mathcal{G}_{t_0} is decoded by substituting \mathcal{C} with \mathcal{C}_{t_0} . This design establishes a coarse-to-fine relationship between the 3D prior and the ground truth, prompting the 3D diffusion process to leverage the 3D prior without excessively depending on it.

In this way, we can get the fused feature volume as:

$$\mathcal{S} = \mathcal{U}([\mathcal{M}, \text{Sp3DConv}(\mathcal{N}), \text{Sp3DConv}(\mathcal{G}_{t_0})]), \quad (4)$$

where \mathcal{U} denotes 3D sparse U-Net. Then we can query features from \mathcal{S} for each grid point p and decode it to SDF values through several MLPs: $\mathcal{F}'_0(p) = \text{MLP}(\mathcal{S}(p), \lambda(p))$, where $\mathcal{S}(p)$ represents the interpolated features from \mathcal{S} at position p . In Sec. 4.2 and Fig. 4, our experiments also demonstrate that our model can generate 3D objects beyond Shap-E model.

3.3. 2D Diffusion Model with 3D Guidance

Our 2D diffusion model simultaneously generates multi-view images by jointly denoising multi-view noisy images $\mathcal{V}_t = \{\mathcal{I}_t^i\}_{i=1}^M$. To encourage 2D-3D consistency, the 2D diffusion model is also guided by the 3D radiance field output from 3D diffusion process mentioned above. Specifically, for better image quality, 2D multi-view diffusion model is built on the multiple independently frozen 2D foundation models (e.g., DeepFloyd [7]) to harness the potent 2D priors. Each of these frozen 2D foundation models (the blue network in Fig. 3) is modulated by view-specific 3D-consistent residual features and responsible for the denoising of a specific view, as described below.

First, to achieve 3D-to-2D guidance, we render multi-view images from the 3D denoised radiance field \mathcal{F}'_0 and feed them to 2D denoising model. Note that the radiance field consists of a density field and a color field. The density field is constructed from the signed distance field (SDF) generated by our 3D diffusion model using S-density introduced in NeuS [40]. To obtain the color field, we apply another color MLP to the feature volume in the 3D diffusion process.

Upon obtaining the color field c and density field σ , we conduct volumetric rendering on each ray $\mathbf{r}(m) = \mathbf{o} + m\mathbf{d}$ which extends from the camera origin \mathbf{o} along a direction \mathbf{d} to produce multi-view consistent images $\{\mathcal{H}^i\}_{i=1}^M$:

$$\hat{\mathcal{C}}(\mathbf{r}) = \int_0^\infty T(m)\sigma(\mathbf{r}(m))c(\mathbf{r}(m)), \mathbf{d}dm, \quad (5)$$

where $T(m) = \exp(-\int_0^m \sigma(\mathbf{r}(s))ds)$ handles occlusion.

Secondly, we use these rendered multi-view images as guidance for the 2D foundation model. We first use a shared feature extractor \mathcal{E} to extract hierarchical multi-view consistent features from these images. Then each extracted feature is added as residual to the decoder of its corresponding frozen 2D foundation denoising U-Net (the red arrow

pointing to “3D-2D Control” in Fig. 3), achieving multi-view modulation and joint denoising following ControlNet [46] as $\hat{f}_k^i = f_k^i + \text{ZeroConv}(\mathcal{E}(\mathcal{H}^i)[k])$, where f_k^i denotes the original feature maps of the k -th decoder layer in 2D foundation model, $\mathcal{E}(\mathcal{H}^i)[k]$ denotes the k -th residual features of the i -th view, and ZeroConv [46] is 1×1 convolution which is initialized by zeros and gradually updated during training. Experimental results show that this 3D-to-2D guidance helps to ensure multi-view consistency and facilitate geometry understanding.

3.4. Separate Control of Geometry and Texture

One advantage of *BiDiff* is that it naturally separates 2D texture generation using 2D diffusion model from 3D geometry generation using 3D diffusion model. Because of this, users can separately control geometry and texture generation, as shown in Fig. 2.

To achieve this, we first propose a prior enhancement strategy to empower a manual control of the strength of 3D and 2D priors independently. Inspired by the classifier-free guidance [9], during training, we randomly drop the information from 3D priors by setting condition feature volume from \mathcal{G} to zero and weaken the 2D priors by using empty text prompts. Consequently, upon completing the training, we can employ two guidance scales, γ_{3d} and γ_{2d} , to independently modulate the influence of these two priors.

Specifically, to adjust the strength of 3D prior, we calculate the difference between 3D diffusion outputs with and without conditional 3D feature volumes, and add them back to 3D diffusion output:

$$\hat{\epsilon}_{3d} = \mathcal{D}_{3d}(\mathcal{F}_t, \mathcal{V}'_{t+1}, t) + \gamma_{3d} \cdot ((\mathcal{D}_{3d}(\mathcal{F}_t, \mathcal{V}'_{t+1}, t | \mathcal{G}) - \mathcal{D}_{3d}(\mathcal{F}_t, \mathcal{V}'_{t+1}, t)). \quad (6)$$

Then we can control the strength of 3D prior by adjusting the weight γ_{3d} of this difference term. When $\gamma_{3d} = 0$, it will completely ignore 3D prior. When $\gamma_{3d} = 1$, this is just the previous model that uses both 3D prior and 2D prior. When $\gamma_{3d} > 1$, the model will produce geometries close to the conditional radiance field but with less diversity.

Also, we can similarly adjust the strength of 2D priors by adding differences between 2D diffusion outputs with and without conditional 2D text input:

$$\hat{\epsilon}_{2d} = \mathcal{D}_{2d}(\mathcal{V}_t, \{\mathcal{H}_t^i\}_{i=1}^M, t) + \gamma_{2d} \cdot ((\mathcal{D}_{2d}(\mathcal{V}_t, \{\mathcal{H}_t^i\}_{i=1}^M, t | \text{text}) - \mathcal{D}_{2d}(\mathcal{V}_t, \{\mathcal{H}_t^i\}_{i=1}^M, t)). \quad (7)$$

Increasing γ_{2d} results in more coherent textures with text, albeit at the expense of diversity. It is worth noting that while we adjust the 3D and 2D priors independently via Eq. (6) and Eq. (7), the influence inherently propagates to the other

domain due to the intertwined nature of our bidirectional diffusion process.

With these two guidance scales γ_{3d} and γ_{2d} , we can easily achieve a separate control of geometry and texture. First, to only change texture while keeping geometry untouched, we just fix the initial 3D noisy SDF grids and the conditional radiance field \mathcal{C}_{t_0} , while enlarge their influence by Eq. (6). On the other hand, to only change geometry while keeping texture style untouched, we can maintain keywords in text prompts and enlarge its influence by Eq. (7). By doing so, the shape will be adjusted by the 3D diffusion process.

3.5. Optimization with BiDiff Initialization

The generated radiance field \mathcal{F}_0 using *BiDiff* can be further used as a strong initialization of the optimization-based methods [41]. This additional step can further improve the quality of the 3D model, as shown in Fig. 1 and Fig. 5. Importantly, compared to the geometries directly generated by optimization, our *BiDiff* can output more diversified geometry which better aligns with users’ input text, and also has more accurate 3D structure. Therefore, the optimization started from this strong initialization can be rather efficient (≈ 20 min) and avoid incorrect geometries like multi-face and floaters.

Specifically, we first convert generated radiance field \mathcal{F}_0 by *BiDiff* into a higher resolution one $\bar{\mathcal{F}}_0$ that supports 512×512 resolution image rendering, as shown on the right of Fig. 3. This process is achieved by a fast NeRF distillation operation (≈ 2 min). The distillation first bounds the occupancy grids of $\bar{\mathcal{F}}_0$ with the estimated binary grids (transmittance > 0.01) from the original radiance field \mathcal{F}_0 , then overfits $\bar{\mathcal{F}}_0$ to \mathcal{F}_0 by minimizing both the L_1 distance between two density fields and L_1 distance between their renderings 2D images under random viewpoints. Thanks to this flexible and fast distillation operation, we can efficiently convert generated radiance field by *BiDiff* into any 3D representations an optimization-based method requires. In our experiments, since we are using ProlificDreamer [41], we use the InstantNGP [27] as the high-resolution radiance field.

After initialization, we optimize $\bar{\mathcal{F}}_0$ by SDS loss following the previous methods [30, 41]. It is noteworthy that since we already have a good initialized radiance field, we only need to apply a small noise level SDS loss. Specifically, we set the ratio range of denoise timestep t_{opt} to $[0.02, 0.5]$ during the entire optimization process.

4. Experiment

In this section, we described our experimental results. We train our framework on the ShapeNet-Chair [1] and Objaverse LVIS 40k datasets [4]. We use the pre-trained DeepFloyd-IF-XL [7] as our 2D foundation model and ShapE [12] as our 3D priors. We adopt the SparseNeuS [21] as the neural surface field presentation with $N = 128$. For the




Figure 4. Qualitative sampling results of Bidirectional Diffusion model, including multi-view images and 3D mesh from diffusion sampling. The top two rows are the results on the Shapenet-Chair, and the bottom two rows are the results on the Objaverse. We compared the results of Shap-E in the last column.

3D-to-2D guidance, We follow the setup of ControlNet [46] to render $M = 8$ multi-view images with 64×64 resolution using SparseNeuS. We train our framework on 4 NVIDIA A100 GPUs for both ShapeNet and Objaverse 40k experiments with a batch size of 4. During sampling, we set the 3D and 2D prior guidance scale to 3.0 and 7.5 respectively. More details on data processing and model architecture are included in supplementary material. We discuss the evaluation and ablation results below. Please refer to supplementary webpages and videos for more visual results.

4.1. Text-to-3D Results

ShapeNet-Chair results. The first and second rows of Fig. 4 present our results trained on the ShapeNet-Chair dataset. Although the chair category often contains complicated geometric details, our framework demonstrates the capability to capture those fine details. Concurrently, our approach exhibits a remarkable capability to produce rich and diverse textures by merely modulating the textual prompts, leading to compelling visual outcomes.

Objaverse results. Scaling to a much larger 3D dataset, Objaverse, our framework’s efficacy becomes increasingly pronounced. The bottom two rows of Fig. 4 are results from the Objaverse dataset. Compared to objects generated by Shap-E, our model closely adheres to the given textual prompts. This again shows that the proposed *BiDiff* learns to model both 2D textures and 3D geometries better compared with 3D-only solutions, and is capable of generating more diverse geometries.

Decouple geometry and texture control. Lastly, we illustrate that our *BiDiff* can separately control geometry and

texture generation. First, as illustrated in the first row of Fig. 2, when the 3D prior is fixed, we have the flexibility to manipulate the 2D diffusion model using varying textual prompts to guide the texture generation process. This capability enables the generation of a diverse range of textured objects while maintaining a consistent overall shape. Second, when we fix the textual prompt for the 2D priors (e.g., "a xxx with Van Gogh starry sky style"), we can adjust the 3D diffusion model by varying the conditional radiance field derived from the 3D priors. This procedure results in the generation of a variety of shapes, while maintaining a similar texture, as shown in the second row of Fig. 2.

4.2. Comparison with other Generation Models

Comparison with optimization methods. Our framework is capable of simultaneously generating multi-view consistent images alongside a 3D mesh in a scalable manner. In

Table 1. CLIP R-precision.

Method	R-P	time
DreamFusion	0.67	1.1h
ProlificDreamer	0.83	3.4h
Ours-sampling	0.79	40s
Ours-post	0.85	20min

contrast, the SDS-based methods [30, 41] utilize a one-by-one optimization approach. Tab. 1 reports the CLIP R-Precision [12] and inference time on 50 test prompts manually derived from the captioned untrained Objaverse to quantitatively evaluate these methods. Also, optimization methods, Dreamfusion [30] and ProlificDreamer [41], are expensive, taking several hours to generate an object. Moreover, these optimization methods may lead to more severe multi-face problems. In contrast, our method can produce realistic objects with reasonable geometry in only 40 seconds. Fur-

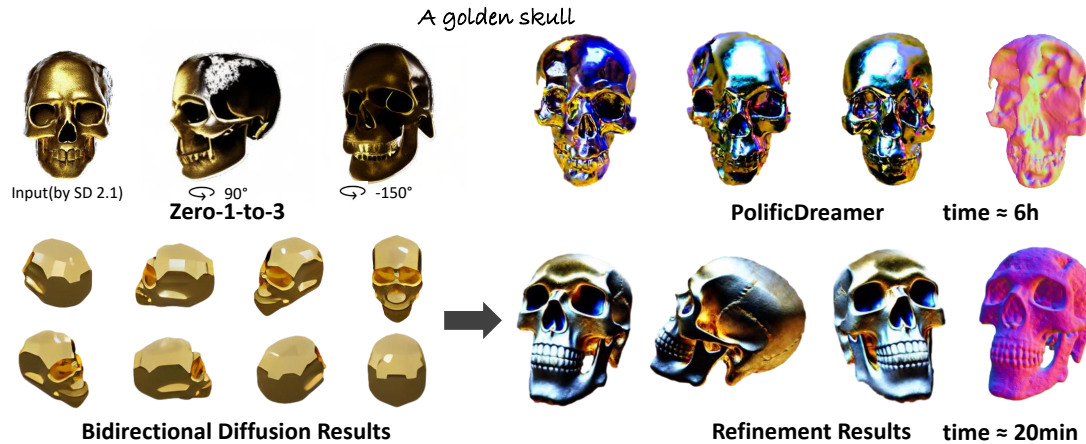


Figure 5. Comparison with other optimization or multi-view diffusion based works. We show both multi-view images (left) and 3D results (right). Zero-1-to-3 [18] is not good at predicting results from a large perspective, and PolificDreamer [41] suffers from the multi face problem. Our method has excellent robustness and can obtain high-quality results in a short period of time.




Figure 6. Ablation of prior and prior enhancement.

thermore, *BiDiff* can serve as a strong prior for optimization-based methods and significantly boost their performance. Initializing the radiance field in ProlificDreamer [41] with our outputs shows remarkable improvements in both quality and computational efficiency, as shown in Fig. 5.

Comparison with multi-view methods Given one reference image, the multi-view method Zero-1-to-3 [18] produces images from novel views by fine-tuning a pre-trained 2D diffusion model on multi-view datasets. However, this method employs cross-view attention to establish multi-view correspondence without an inherent understanding of 3D structures, inevitably leading to inconsistent multi-view images as shown in Fig. 5. Moreover, the Zero-1-to-3 series cannot directly generate the 3D mesh, requiring substantial post-processing to acquire the geometry. In contrast, our framework also incorporates 3D priors, in addition to 2D priors, and thus can generate more accurate 3D geometries.

4.3. Abalation Studies

We perform comprehensive ablation studies on the ShapeNet-Chair dataset [1] to evaluate the importance of each com-

ponent below. More ablation results can be found in the supplementary material.

3D priors. To assess the impact of 3D priors, we eliminate the conditional radiance field from Shap-E and train the 3D geometry generation from scratch. The experimental results in the second row of Fig. 6 demonstrate that in the absence of the 3D priors, our framework can only generate common objects in the training set.

2D priors. To delve into the impact of 2D priors, we randomly initiate the parameters of the 2D diffusion model, instead of fine-tuning on a pretrained model. The results in the first row of Fig. 6 show that in the absence of 2D priors, the textures generated tend to fit the stylistic attributes of the synthetic training data. Conversely, with 2D priors, we can produce more realistic textures.

Prior enhancement strategy. As discussed in Sec. 3.4, we can adjust the influence of both 3D and 2D priors by the prior enhancement strategy. Fig. 6 also shows the results without this strategy. The prior enhancement strategy plays a vital role in achieving diverse and flexible 3D generation.

5. Conclusion

In this paper, we propose Bidirectional Diffusion, which incorporates both 3D and 2D diffusion processes into a unified framework. Furthermore, Bidirectional Diffusion leverages the robust priors from 3D and 2D foundation models, achieving generalizable geometry and texture understanding.

Acknowledgments: This work was supported by CUHK Direct Grants (RCFUS) No. 4055189 and Research Grants Council (RGC) of the Hong Kong SAR under grant No. 26202321. And we gratefully acknowledge the support of SenseTime.

References

- [1] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. **2, 6, 8, 1**
- [2] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, 2019. **3**
- [3] Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. In *CVPR*, 2023. **3**
- [4] 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 *CVPR*, 2023. **2, 3, 6, 1**
- [5] C. Deng, C. Jiang, C. R. Qi, X. Yan, Y. Zhou, L. Guibas, and D. Anguelov. Nerdi: Single-view nerf synthesis with language-guided diffusion as general image priors. In *CVPR*, 2023. **3**
- [6] Ziya Erkoç, Fangchang Ma, Qi Shan, Matthias Nießner, and Angela Dai. Hyperdiffusion: Generating implicit neural fields with weight-space diffusion. In *ICCV*, 2023. **3**
- [7] Deep Floyd. If project. <https://github.com/deep-floyd/IF>, 2023. **5, 6**
- [8] Philipp Henzler, Niloy J. Mitra, and Tobias Ritschel. Escaping plato’s cave: 3d shape from adversarial rendering. In *ICCV*, 2019. **3**
- [9] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. **6**
- [10] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *CVPR*, 2022. **3**
- [11] Julian Straub Richard Newcombe Jeong Joon Park, Peter Florence and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. **3**
- [12] Heewoo Jun and Alex Nichol. Shap-e: Generating conditional 3d implicit functions. *arXiv preprint arXiv:2305.02463*, 2023. **2, 3, 5, 6, 7, 1**
- [13] Nasir Mohammad Khalid, Tianhao Xie, Eugene Belilovsky, and Popa Tiberiu. Clip-mesh: Generating textured meshes from text using pretrained image-text models. *SIGGRAPH Asia 2022 Conference Papers*, 2022. **3**
- [14] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 2023. **1**
- [15] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiao-hui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *CVPR*, 2023. **2, 3**
- [16] Tong Wu Yu-Jie Yuan Hongbo Fu Yu-Kun Lai Lin Gao, Jie Yang and Hao Zhang. Sdm-net: Deep generative network for structured deformable mesh. *ACM Transactions on Graphics*, 38:1–15, 2019. **3**
- [17] 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. *NeurIPS*, 2024. **4**
- [18] 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. In *ICCV*, 2023. **2, 3, 8**
- [19] 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**
- [20] Zhen Liu, Yao Feng, Michael J. Black, Derek Nowrouzezahrai, Liam Paull, and Weiyang Liu. Meshdiffusion: Score-based generative 3d mesh modeling. In *ICLR*, 2023. **3**
- [21] Xiaoxiao Long, Cheng Lin, Peng Wang, Taku Komura, and Wenping Wang. Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In *ECCV*, 2022. **5, 6**
- [22] 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. *arXiv preprint arXiv:2310.15008*, 2023. **3**
- [23] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Realfusion: 360 reconstruction of any object from a single image. In *CVPR*, 2023. **3**
- [24] Gal Metzger, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. In *CVPR*, 2023. **2, 3, 5**
- [25] Gal Metzger, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. In *CVPR*, 2023. **2**
- [26] Gregor Kobsik Moritz Ibing and Leif Kobbelt. Octree transformer: Autoregressive 3d shape generation on hierarchically structured sequences. *arXiv preprint arXiv:2111.12480*, 2021. **3**
- [27] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Transactions on Graphics*, 41(4):102:1–102:15, 2022. **6**
- [28] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022. **2, 3**
- [29] Ioannis Mitliagkas Panos Achlioptas, Olga Diamanti and Leonidas Guibas. Learning representations and generative models for 3d point clouds. In *ICML*, 2018. **3**
- [30] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *ICLR*, 2022. **2, 6, 7**
- [31] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, 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. **3**
- [32] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry,

- Amanda Aspell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 3
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022.
- [34] 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. *NeurIPS*, 2022. 3
- [35] Junyoung Seo, Wooseok Jang, Min-Seop Kwak, Jaehoon Ko, Hyeonsu Kim, 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. 3
- [36] Yichun Shi, Peng Wang, Jianglong Ye, Long Mai, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. *arXiv:2308.16512*, 2023. 3
- [37] Edward Smith and David Meger. Deep unsupervised learning using nonequilibrium thermodynamics. In *Conference on Robot Learning*, 2017. 3
- [38] Xiao Fu Yuxin Wang Jiawei Ren Liang Pan-Wayne Wu Lei Yang Jiaqi Wang Chen Qian Dahua Lin Ziwei Liu Tong Wu, Jiarui Zhang. Omniobject3d: Large-vocabulary 3d object dataset for realistic perception, reconstruction and generation. *CVPR*, 2023. 3
- [39] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *CVPR*, 2023. 3
- [40] 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. 5
- [41] 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. *NeurIPS*, 2024. 2, 3, 6, 7, 8, 5
- [42] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *NeurIPS*, 2016. 3
- [43] Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. Pointflow: 3d point cloud generation with continuous normalizing flows. In *ICCV*, 2019. 3
- [44] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *CVPR*, 2021. 1
- [45] Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. Lion: Latent point diffusion models for 3d shape generation. *arXiv preprint arXiv:2210.06978*, 2022. 3
- [46] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023. 6, 7, 3