

## Wonder3D: Single Image to 3D using Cross-Domain Diffusion

Xiaoxiao Long<sup>1,3\*</sup>, Yuan-Chen Guo<sup>2\*‡</sup>, Cheng Lin<sup>1†</sup>, Yuan Liu<sup>1</sup>, Zhiyang Dou<sup>1</sup>  
 Lingjie Liu<sup>4</sup>, Yuexin Ma<sup>5</sup>, Song-Hai Zhang<sup>2</sup>, Marc Habermann<sup>6</sup>, Christian Theobalt<sup>6</sup>, Wenping Wang<sup>7†</sup>

<sup>1</sup>The University of Hong Kong <sup>2</sup>Tsinghua University <sup>3</sup>VAST

<sup>4</sup>University of Pennsylvania <sup>5</sup>Shanghai Tech University <sup>6</sup>MPI Informatik <sup>7</sup>Texas A&M University

\*Core contributions †Corresponding authors ‡Intern at VAST

<https://www.xxlong.site/Wonder3D/>



Figure 1. *Wonder3D* reconstructs highly-detailed textured meshes from a single-view image in only 2 ~ 3 minutes. *Wonder3D* first generates consistent multi-view normal maps with corresponding color images via a cross-domain diffusion model, and then leverages a novel normal fusion method to achieve fast and high-quality reconstruction. Compared to the prior works, *Wonder3D* has achieved a leading-level of geometric details with high efficiency.

## Abstract

*In this work, we introduce **Wonder3D**, a novel method for efficiently generating high-fidelity textured meshes from single-view images. Recent methods based on Score Distillation Sampling (SDS) have shown the potential to recover 3D geometry from 2D diffusion priors, but they typically suffer from time-consuming per-shape optimization and inconsistent geometry. In contrast, certain works directly produce 3D information via fast network inferences, but their results are often of low quality and lack geometric details. To holistically improve the quality, consistency, and efficiency of single-view reconstruction tasks, we propose a cross-domain diffusion model that generates multi-view normal maps and the corresponding color images. To ensure the consistency of generation, we employ a multi-view cross-domain attention mechanism that facilitates information exchange across views and modalities. Lastly, we introduce a geometry-aware normal fusion algorithm that extracts high-quality surfaces from the multi-view 2D representations in only 2 ~ 3 minutes. Our extensive evaluations demonstrate that our method achieves high-quality reconstruction results, robust generalization, and good efficiency compared to prior works.*

## 1. Introduction

Reconstructing 3D geometry from a single image [13, 26, 33, 38, 40, 42, 45] stands as a fundamental task in computer graphics and 3D computer vision, benefiting a wide range of versatile applications such as novel view synthesis [7, 25, 35], 3D content creation [36, 48], and robotics grasping [29, 79]. However, this task is notably challenging since it is ill-posed and demands the ability to discern the 3D geometry of both visible and invisible parts. This ability requires extensive knowledge of the 3D world.

Recently, the field of 3D generation has experienced rapid and flourishing development with the introduction of diffusion models. A growing body of research [5, 31, 47, 63, 67], such as DreamField [24], DreamFusion [47], and Magic3D [31], resort to distilling prior knowledge of 2D image diffusion models or vision language models to create 3D models from text or images via Score Distillation Sampling (SDS) [47]. Despite their compelling results, these methods suffer from two main limitations: **efficiency** and **consistency**. The per-shape optimization process typically entails tens of thousands of iterations, involving full-image volume rendering and inferences of the diffusion models. Consequently, it often consumes tens of minutes or even hours on per-shape optimization. Moreover, the 2D prior model operates by considering only a single view at each iteration and strives to make every view resemble the input image. This often results in the generation of 3D shapes

exhibiting inconsistencies, thus, often leading to the generation of 3D shapes with inconsistencies such as multiple faces (i.e., the Janus problem [47]).

There exists another group of works that endeavor to directly produce 3D geometries like point clouds [41, 45, 75, 80], meshes [16, 37], neural fields [1, 4, 8, 15, 17, 21, 26–28, 44, 46, 65, 76] via network inference to avoid time-consuming per-shape optimization. Most of them attempt to train 3D generative diffusion models from scratch on 3D assets. However, due to the limited size of publicly available 3D datasets, these methods demonstrate poor **generalizability**, most of which can only generate shapes on specific categories.

More recently, several methods have emerged that directly generate multi-view 2D images, with representative works including SyncDreamer [36] and MVDream [55]. By enhancing the multi-view consistency of image generation, these methods can recover 3D shapes from the generated multi-view images. Our method also adopts a multi-view generation scheme to favor the flexibility and efficiency of 2D representations. However, due to only relying on color images, the **fidelity** of the generated shapes is not well-maintained, and they struggle to recover geometric details or come with enormous computational costs.

To better address the issues of fidelity, consistency, generalizability and efficiency in the aforementioned works, in this paper, we introduce a new approach to the task of single-view 3D reconstruction by generating multi-view consistent normal maps and their corresponding color images with a cross-domain diffusion model. The key idea is to extend diffusion frameworks to model the joint distribution of two different domains, i.e., normals and colors. The normal map characterizes the undulations and variations presented on the surface of the shape, thus encoding rich detailed geometric information. This allows for the high-fidelity extraction of 3D geometry from 2D normal maps. Adopting the 2D representations enables *Wonder3D* to be built on the pre-trained Stable Diffusion model [49], where strong priors facilitate zero-shot generalization ability.

The following technical designs of *Wonder3D* make it a robust and efficient tool to create 3D shapes from single images. 1) **Cross-domain switcher**. The introduced domain switcher allows the diffusion model to generate either normal maps or color images without significantly modifying the original model. 2) **Cross-domain attentions**. We further leverage cross-domain attention mechanisms to assist in the information exchange between the two domains, ultimately improving consistency and quality. This mechanism facilitates information perception across different domains, enabling our method to recover high-fidelity geometry. 3) **Geometry-aware normal fusion**. In order to stably extract surfaces from the generated views, we propose a geometry-aware normal fusion algorithm that is robust to subtle inac-




Figure 2. Overview of *Wonder3D*. Given a single image, *Wonder3D* takes the input image, the text embedding produced by CLIP model [49], the camera parameters of multiple views, and a domain switcher as conditioning to generate consistent multi-view normal maps and color images. Subsequently, *Wonder3D* employs an innovative normal fusion algorithm to robustly reconstruct high-quality 3D geometry from the 2D representations, yielding high-fidelity textured meshes.

curate generations and capable of reconstructing clean and high-quality geometries (see Figure 1).

We conduct extensive experiments on the Google Scanned Object dataset [14] and various 2D images with different styles. The experiments validate that *Wonder3D* achieves a leading level of geometric details with high efficiency and strong generalization among current zero-shot single-view reconstruction methods.

## 2. Related Works

### 2.1. 2D Diffusion Models for 3D Generation

Recent compelling successes in 2D diffusion models [9, 22, 51] and large vision language models (e.g., CLIP model [49]) provide new possibilities for generating 3D assets using the strong priors of 2D diffusion models. Pioneering works DreamFusion [47] and SJC [63] propose to distill a 2D text-to-image generation model to generate 3D shapes from texts, and many follow-up works follow such per-shape optimization scheme. For the task of text-to-3D [2, 5, 6, 23, 31, 52, 53, 61, 67, 69, 73, 82] or image-to-3D synthesis [42, 48, 50, 54, 58, 71], these methods typically optimize a 3D representation (i.e., NeRF, mesh, or SDF), and then leverage neural rendering to generate 2D images from various viewpoints. The images are then fed into the 2D diffusion models or CLIP model for calculating SDS [47] losses, which guide the 3D shape optimization. However, most of these methods always suffer from low efficiency and multi-face problem, where a per-shape optimization consumes tens of minutes and the optimized geometry tends to produce multiple faces due to the lack of explicit 3D supervision.

### 2.2. 3D Generative Models

Instead of performing a time-consuming per-shape optimization guided by 2D diffusion models, some works attempt to directly train 3D diffusion models based on various 3D representations, like point clouds [41, 45, 75, 80],

meshes [16, 37], neural fields [1, 4, 8, 15, 17, 21, 26–28, 44, 46, 65, 76] However, due to the limited size of public available 3D assets dataset, most of the works have only been validated on limited categories of shapes, and how to scale up on large datasets is still an open problem. On the contrary, our method adopts 2D representations and, thus, can be built upon the 2D diffusion models [51] whose pre-trained priors significantly facilitate zero-shot generalization ability.

### 2.3. Multi-view Diffusion Models

To generate consistent multi-view images, some efforts [3, 11, 18, 30, 34, 57, 59, 60, 62, 68, 70, 72, 74, 77, 81] are made to extend 2D diffusion models from single-view images to multi-view images. However, most of these methods focus on image generation and are not designed for 3D reconstruction. The recent works Viewset Diffusion [57], SyncDreamer [36], and MVDream [55] share a similar idea to produce consistent multi-view color images via attention layers. However, unlike that normal maps explicitly encode geometric information, reconstruction from color images always suffers from texture ambiguity, and, thus, they either struggle to recover geometric details or require huge computational costs. SyncDreamer [36] requires dense views for 3D reconstruction but still suffers from low-quality geometry and blurring textures. MVDream [55] still resorts to a time-consuming optimization using SDS loss for 3D reconstruction, and its multi-view distillation scheme requires 1.5 hours. In contrast, our method can reconstruct high-quality textured meshes in just 2 minutes.

## 3. Problem Formulation

### 3.1. Diffusion Models

Diffusion models [22, 56] are first proposed to gradually recover images from a specifically designed degradation process, where a forward Markov chain and a Reverse Markov chain are adopted. The forward process will be iteratively applied to the target image  $z$  until the image becomes com-



plete Gaussian noise at the end. On the contrary, the reverse chain then is employed to iteratively denoise the corrupted image, i.e., recovering  $z_{t-1}$  from  $z_t$  by predicting the added random noise  $\epsilon$ . The readers can refer to [22, 56] for more details about image diffusion models.

### 3.2. The Distribution of 3D Assets

Unlike that prior works adopt 3D representations like point clouds, tri-planes, or neural radiance fields, we propose that the distribution of 3D assets, denoted as  $p_a(\mathbf{z})$ , can be modeled as a joint distribution of its corresponding 2D multi-view normal maps and color images. Specifically, given a set of cameras  $\{\pi_1, \pi_2, \dots, \pi_K\}$  and a conditional input image  $y$ , we have

$$p_a(\mathbf{z}) = p_{nc}(n^{1:K}, x^{1:K}|y), \quad (1)$$

where  $p_{nc}$  is the distribution of the normal maps  $n^{1:K}$  and color images  $x^{1:K}$  observed from a 3D asset conditioned on an image  $y$ . For simplicity, we omit the symbol  $y$  for this equation in the following discussions. Therefore, our goal is to learn a model  $f$  that synthesizes multiple normal maps and color images of a set of camera poses denoted as

$$(n^{1:K}, x^{1:K}) = f(y, \pi_{1:K}). \quad (2)$$

Finally, we can formulate this cross-domain joint distribution as a Markov chain within the diffusion scheme:

$$p\left(n_T^{(1:K)}, x_T^{(1:K)}\right) \prod_t p_{nc}\left(n_{t-1}^{(1:K)}, x_{t-1}^{(1:K)} | n_t^{(1:K)}, x_t^{(1:K)}\right), \quad (3)$$

where  $p\left(n_T^{(1:K)}, x_T^{(1:K)}\right)$  are Gaussian noises. Our key problem is to characterize the distribution  $p_{nc}$ , so that we can sample from this Markov chain to generate normal maps and images.

## 4. Method

As per our problem formulation in Section 3.2, we propose a multi-view cross-domain diffusion scheme, which operates on two distinct domains to generate multi-view consistent normal maps and color images. The overview of our method is presented in Figure 2. First, our method adopts a multi-view diffusion scheme to generate multi-view normal maps and color images, and enforces the consistency across different views using multi-view attentions (see Section 4.1). Second, our proposed domain switcher allows the diffusion model to operate on more than one domain. A cross-domain attention is proposed to propagate information between the normal domain and color image domain ensuring geometric and visual coherence between the two domains (see Section 4.2). Finally, our novel geometry-aware normal fusion reconstructs the high-quality geometry and appearance from the multi-view 2D normal and color images (see Section 4.3).

### 4.1. Consistent Multi-view Generation

The prior 2D diffusion models [33, 49] generate each image separately, so that the resulting images are not geometrically and visually consistent across different views. To enhance consistency among different views, similar to prior works such as SyncDreamer [36] and MVDream [55], we utilize attention mechanism to facilitate information propagation across different views, implicitly encoding multi-view dependencies (as illustrated in Figure 3).

This is achieved by extending the original self-attention layers to be global-aware, allowing connections to other views within the attention layers. Keys and values from different views are connected to each other to facilitate the exchange of information. By sharing information across different views within the attention layers, the diffusion model perceives multi-view correlation and is capable of generating consistent multi-view color images and normal maps.

### 4.2. Cross-Domain Diffusion

Our model is built upon pre-trained stable diffusion models [49] to leverage its strong priors. Given that current 2D diffusion models [33, 49] are designed for a single domain, the main challenge lies in how to effectively extend stable diffusion models to support multi-domain operations.

**Naive Solutions.** To achieve this goal, we explore several possible designs. A straightforward solution is to add four more channels to the output of the UNet module representing the extra domain. Therefore, the diffusion model can simultaneously output normals and color image domains. However, we notice that such a design suffers from low convergence speed and poor generalization. This is because the channel expansion may perturb the pre-trained weights of stable diffusion models and therefore cause catastrophic model forgetting.

Revisiting Eq. 1, it is possible to factor the joint distribution into two conditional distributions:

$$q_a(\mathbf{z}) = q_n(n^{1:K}) \cdot q_c(x^{1:K} | n^{1:K}). \quad (4)$$

This equation suggests an alternative solution where we could initially train a diffusion model to generate normal maps and then train another diffusion model to produce color images, conditioning on the generated normal maps (or vice versa). Nonetheless, the implementation of this two-stage framework introduces certain complications. It not only substantially increases the computational cost but also results in performance degradation. Please refer to Section 5.6 for an in-depth discussion.

**Domain Switcher.** To overcome these difficulties mentioned above, we design a cross-domain diffusion scheme via a *domain switcher*, denoted as  $s$ . The switcher  $s$  is a one-dimensional vector serving as labels for distinct domains, and we further feed the switcher into the diffusion

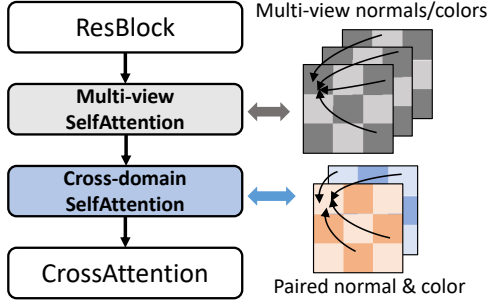


Figure 3. The illustration of the structure of the multi-view cross-domain transformer block.

model as an extra input. Therefore, the formulation of Eq. 2 can be extended as:

$$n^{1:K}, x^{1:K} = f(y, \pi_{1:K}, s_n), f(y, \pi_{1:K}, s_c). \quad (5)$$

The domain switcher  $s$  is first encoded via positional encoding [43] and subsequently concatenated with the time embedding. This combined representation is then injected into the UNet of the stable diffusion models. Interestingly, experiments show that this subtle modification does not significantly alter the pre-trained priors. As a result, it allows for fast convergence and robust generalization, without requiring substantial changes to the stable diffusion models.

**Cross-domain Attention.** Using the proposed domain switcher, the diffusion model can generate two different domains. However, it is important to note that for a single view, there is no guarantee that the generated color image and the normal map will be geometrically consistent. To address this issue and ensure the consistency between the generated normal maps and color images, we introduce a cross-domain attention mechanism to facilitate the exchange of information between the two domains. This mechanism aims to ensure that the generated outputs align well in terms of geometry and appearance.

The cross-domain attention layer maintains the same structure as the original self-attention layer and is integrated before the cross-attention layer in each transformer block of the UNet, as depicted in Figure 3. In the cross-domain attention layer, the keys and values from the normal and color image domains are combined and processed through attention operations. This design ensures that the generations of color images and normal maps are closely correlated, thus promoting geometric consistency between the two domains.

### 4.3. Textured Mesh Extraction

To extract explicit 3D geometry from 2D normal maps and color images, we optimize a neural implicit signed distance field (SDF) to amalgamate all 2D generated data. Nonetheless, adopting existing SDF-based reconstruction methods, such as NeuS [64], proves unviable. These methods were

tailored for real-captured images and necessitate dense input views. In contrast, our generated views are relatively sparse, and the generated normal maps and color images may exhibit subtle inaccurate predictions of some pixels. Regrettably, these errors accumulate during the geometry optimization, leading to distorted geometries, outliers, and incompleteness. To overcome the challenges above, we propose a novel geometric-aware optimization scheme.

**Optimization Objectives.** With the obtained normal maps  $G_{0:N}$  and color images  $H_{0:N}$ , we first leverage segmentation models to segment the object masks  $M_{0:N}$  from the normal maps or color images. Then we optimize SDF field with such an objective function:

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_{normal} + \mathcal{L}_{rgb} + \mathcal{L}_{mask} \\ & + \mathcal{R}_{eik} + \mathcal{R}_{sparse} + \mathcal{R}_{smooth}, \end{aligned} \quad (6)$$

where  $\mathcal{L}_{normal}$  denotes the normal loss term that will be discussed later,  $\mathcal{L}_{rgb}$  denotes a MSE loss term that calculates the errors between rendered colors  $\hat{h}_k$  and generated colors  $h_k$ ,  $\mathcal{L}_{mask}$  denotes a binary cross-entropy loss term that calculating errors between the rendered mask  $\hat{m}_k$  and the generated mask  $m_k$ ,  $\mathcal{R}_{eik}$  denotes eikonal regularization term that encourages the magnitude of the SDF gradients to be unit length,  $\mathcal{R}_{sparse}$  denotes a sparsity regularization term that avoid floaters of SDF, and  $\mathcal{R}_{smooth}$  denotes a 3D smoothness regularization term that enforces the SDF gradients to be smooth in 3D space (see details in the supp.).

**Geometry-aware Normal Loss.** Thanks to the differentiable nature of SDF representation, we can easily extract normal values  $\hat{g}$  of the optimized SDF via calculating the gradients of SDF. We maximize the similarity of the normal of SDF  $\hat{g}$  and our generated normal  $g$  to provide 3D geometric supervision. To tolerate trivial inaccuracies of the generated normals from different views, we introduce a geometry-aware normal loss:

$$\mathcal{L}_{normal} = \frac{1}{\sum w_k} \sum w_k \cdot (1 - \cos(\hat{g}_k, g_k)) \quad (7)$$

where we measure the error between the normal of SDF  $\hat{g}_k$  and the generated normal  $g_k$  for the  $k_{th}$  sampled ray,  $\cos(\cdot, \cdot)$  denotes cosine function, and  $w_k$  is a geometric-aware weight defined as

$$w_k = \begin{cases} 0, & \cos(\mathbf{v}_k, \mathbf{g}_k) > \epsilon \\ \exp(-|\cos(\mathbf{v}_k, \mathbf{g}_k)|), & \cos(\mathbf{v}_k, \mathbf{g}_k) \leq \epsilon \end{cases}. \quad (8)$$

Here  $\exp(\cdot)$  denotes exponential function,  $|\cdot|$  denotes absolute function,  $\epsilon$  is a negative threshold closing to zero, and we measure the cosine value of the angle between the generated normal  $g_k$  and the  $k_{th}$  ray's viewing direction  $\mathbf{v}_k$ .




Figure 4. The qualitative comparisons with baseline models on synthesized multi-view color images. We present the picked best results of Zero123 [33] after multiple runs.




Figure 5. The qualitative comparisons with baseline methods on GSO [14] dataset in terms of the reconstructed textured meshes.

The design rationale behind this approach lies in the orientation of normals, which are deliberately set to face outward, while the viewing direction is inward-facing. This configuration ensures that the angle between the normal vector and the viewing ray remains not less than  $90^\circ$ . A deviation from this criterion would imply inaccuracies in the generated normals. Furthermore, rather than treating normals from different views equally, we introduce a weighting mechanism. We assign higher weights to normals that form larger angles with the viewing rays. This prioritization enhances the accuracy of our geometric supervision process.

**Outlier Dropping.** Instead of directly summing up the color errors of all sampled rays at each iteration, we first sort

these errors in a descending order and then discard the top largest errors according to a predefined percentage. This enables the optimization process to circumvent the outlier regions with poor quality or inconsistency, and instead leverage the regularization capability of MLP to learn the geometry. This strategy effectively eliminates inaccurate isolated surfaces and distorted textures.

## 5. Experiments

### 5.1. Implementation Details

We train our model on the LVIS subset of the Objaverse dataset [10], which comprises approximately 30k+ objects



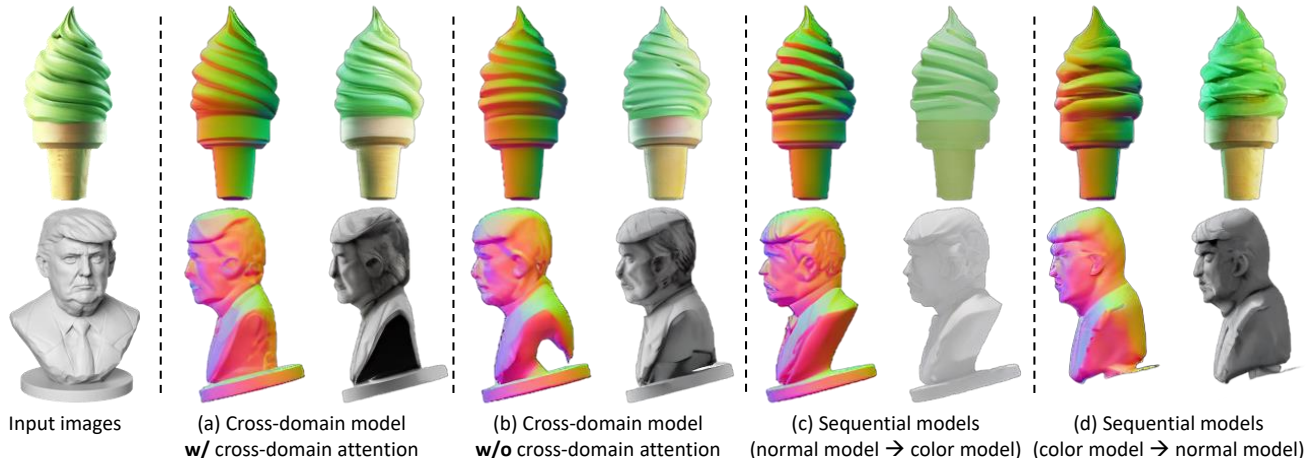


Figure 6. Ablation studies on different cross-domain diffusion schemes.

following a cleanup process. To create the rendered multi-view dataset, we first normalized each object to be centered and of unit scale. Then we render normal maps and color images from six views, including the front, back, left, right, front-right, and front-left views, using Blenderproc [12].

We fine-tune our model starting from the Stable Diffusion Image Variations Model, which has previously been fine-tuned with image conditions. During fine-tuning, we use a reduced image size of  $256 \times 256$  and a total batch size of 512 for training. The fine-tuning process involves training the model for 30,000 steps. This entire training procedure typically requires approximately 3 days on a cluster of 8 Nvidia Tesla A800 GPUs. To reconstruct 3D geometry from the 2D representations, our normal fusion method is built on the instant-NGP SDF reconstruction method [19].

## 5.2. Baselines

We adopt Zero123 [33], RealFusion [42], Magic123 [48], One-2-3-45 [32], Point-E [45], Shap-E [26] and a recent work SyncDreamer [36] as baseline methods. Given an input image, Zero123 [33] is capable of generating novel views of arbitrary viewpoints, and it can be incorporated with SDS loss [47] for 3D reconstruction (we adopt the implementation of ThreeStudio [20]). RealFusion [42] and Magic123 [48] leverage Stable Diffusion [51] and SDS loss for single-view reconstruction. One-2-3-45 [32] directly predict SDFs via SparseNeuS [39] by taking the generated multiple images of Zero123 [33]. Point-E [45] and Shap-E [26] are 3D generative models trained on a large internal OpenAI 3D dataset, both of which are able to convert a single-view image into a point cloud or an implicit representation. SyncDreamer[36] generates multi-view consistent images from a single image for deriving 3D geometry.

## 5.3. Evaluation Protocol

**Evaluation Datasets.** Following prior research [33, 36], we adopt the Google Scanned Object dataset [14] for our eval-

Method	Chamfer Dist.↓	Volume IoU↑
Realfusion [42]	0.0819	0.2741
Magic123 [48]	0.0516	0.4528
One-2-3-45 [32]	0.0629	0.4086
Point-E [45]	0.0426	0.2875
Shap-E [26]	0.0436	0.3584
Zero123 [33]	0.0339	0.5035
SyncDreamer [36]	0.0261	0.5421
Ours	<b>0.0199</b>	<b>0.6244</b>

Table 1. Quantitative comparison with baseline methods. We report Chamfer Distance and Volume IoU on the GSO [14] dataset.

uation, which includes a wide variety of common everyday objects. Our evaluation dataset matches that of SyncDreamer [36], comprising 30 objects that span from everyday items to animals. For each object in the evaluation set, we render an image with a size of  $256 \times 256$ , which serves as the input. Additionally, we include some images with diverse styles collected from the internet or text2image models in our evaluation.

**Metrics.** To evaluate the quality of the single-view reconstructions, we adopt two metrics: Chamfer Distances (CD) and Volume IoU between ground-truth shapes and reconstructed shapes. Since different methods adopt various canonical systems, we first align the generated shapes to the ground-truth shapes before calculating the two metrics. Moreover, we adopt the metrics PSNR, SSIM [66] and LPIPS [78] for evaluating the generated color images.

## 5.4. Single View Reconstruction

We evaluate the quality of the reconstructions of different methods. The quantitative results are summarized in Table 1, and the qualitative comparisons are presented in Fig. 5. Shap-E [26] tends to produce incomplete and distorted meshes. SyncDreamer [36] generates shapes that are roughly aligned with the input image but lack detailed geometries, and the texture quality is subpar. One-2-3-45 [32] attempts to reconstruct meshes from the multiview-




Figure 7. Ablation study on the strategies in the mesh extraction module: geometry-aware normal loss and outlier-dropping strategy.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Realfusion [42]	15.26	0.722	0.283
Zero123 [33]	18.93	0.779	0.166
SyncDreamer [36]	20.05	0.798	0.146
Ours	<b>26.07</b>	<b>0.924</b>	<b>0.065</b>

Table 2. The quantitative comparison in novel view synthesis. We report PSNR, SSIM [66], LPIPS [78] on the GSO [14] dataset.

inconsistent outputs of Zero123 [33]. While it can capture coarse geometries, it loses important details in the process. In comparison, our method stands out by achieving the highest quality, both in terms of geometry and textures.

### 5.5. Novel View Synthesis

We evaluate the quality of novel view synthesis for different methods. The quantitative results are presented in Table 2, and the qualitative results can be found in Figure 4. Zero123 [33] produces visually reasonable images, but they lack multi-view consistency since it operates on each view independently. Although SyncDreamer [33] introduces a volume attention scheme to enhance the consistency of multi-view images, their model is sensitive to the elevation degrees of the input images and tends to produce unreasonable results. In contrast, our method generates images that not only exhibit semantic consistency with the input image but also maintain high consistency across multiple views in terms of both colors and geometry.

### 5.6. Discussions

In this section, we perform a series of experiments to validate the effectiveness of the designs in our method.

**Cross-Domain Diffusion.** To validate the effectiveness of our proposed cross-domain diffusion scheme, we study the following settings: (a) cross-domain model with cross-domain attention; (b) cross-domain model without cross-domain attention; (c) sequential models rgb-to-normal: first train a multi-view color diffusion model then train a multi-view normal diffusion model conditioned on the previously generated color images; (d) sequential models normal-to-rgb: first train a multi-view normal diffusion model then train a multi-view color diffusion model conditioned on the previously generated normal images.

As shown in (a) and (b) of Figure 6, it’s evident that the cross-domain attentions significantly enhance the consistency between color images and normals, particularly in terms of the detailed geometries of objects like the ice-cream and the sculpture. From (c) and (d) of Figure 6, while the normals and color images generated by sequential models maintain some consistency, their results suffer from performance drops, like color shifts of (c) and wrong geometries of (d).

**Normal Fusion.** To assess the efficacy of our normal fusion algorithm, we conducted experiments using the complex lion model, which is rich in geometric details, as illustrated in Figure 7. The baseline model’s surfaces exhibited numerous holes and noises. Utilizing either the geometry-aware normal loss or the outlier-dropping loss helps mitigate the noisy surfaces. Finally, combining both strategies yields the best performance, resulting in clean surfaces while preserving detailed geometries.

## 6. Conclusions and Future Works

In this paper, we present *Wonder3D*, an innovative approach designed for efficiently generating high-fidelity textured meshes from single-view images. Experimental results demonstrate that our method upholds good efficiency and robust generalization, and delivers high-quality geometry. However, limited by computational resources, the current implementation of *Wonder3D* only produces normals and color images from six views. This limited number of views makes it challenging for our method to accurately reconstruct objects with very thin structures and severe occlusions. To address this issue, *Wonder3D* may benefit from leveraging more efficient multi-view attention mechanisms to handle a greater number of views effectively.

## Acknowledgements

This research is partially supported by the Innovation and Technology Commission of the HKSAR Government under the InnoHK initiative and Ref. T45-205/21-N of Hong Kong RGC. Song-Hai Zhang is supported by the National Key Research and Development Program of China (No. 2023YFF0905104), the Natural Science Foundation of China (No. 62132012), and Tsinghua-Tencent Joint Laboratory for Internet Innovation Technology.



## References

- [1] 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 *CVPR*, 2023. 2, 3
- [2] Mohammadreza Armandpour, Huangjie Zheng, Ali Sadeghian, Amir Sadeghian, and Mingyuan Zhou. Reimagine the negative prompt algorithm: Transform 2d diffusion into 3d, alleviate janus problem and beyond. *arXiv preprint arXiv:2304.04968*, 2023. 3
- [3] Eric R Chan, Koki Nagano, Matthew A Chan, Alexander W Bergman, Jeong Joon Park, Axel Levy, Miika Aittala, Shalini De Mello, Tero Karras, and Gordon Wetzstein. Generative novel view synthesis with 3d-aware diffusion models. In *ICCV*, 2023. 3
- [4] Hansheng Chen, Jiatao Gu, Anpei Chen, Wei Tian, Zhuowen Tu, Lingjie Liu, and Hao Su. Single-stage diffusion nerf: A unified approach to 3d generation and reconstruction. In *ICCV*, 2023. 2, 3
- [5] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. *arXiv preprint arXiv:2303.13873*, 2023. 2, 3
- [6] Yiwen Chen, Chi Zhang, Xiaofeng Yang, Zhongang Cai, Gang Yu, Lei Yang, and Guosheng Lin. It3d: Improved text-to-3d generation with explicit view synthesis. *arXiv preprint arXiv:2308.11473*, 2023. 3
- [7] Kai Cheng, Xiaoxiao Long, Kaizhi Yang, Yao Yao, Wei Yin, Yuexin Ma, Wenping Wang, and Xuejin Chen. Gaussianpro: 3d gaussian splatting with progressive propagation. *arXiv preprint arXiv:2402.14650*, 2024. 2
- [8] 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. 2, 3
- [9] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *T-PAMI*, 2023. 3
- [10] 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. 6
- [11] Congyue Deng, Chiyu Jiang, Charles R Qi, Xinchun Yan, Yin Zhou, Leonidas Guibas, Dragomir Anguelov, et al. Nerdi: Single-view nerf synthesis with language-guided diffusion as general image priors. In *CVPR*, 2023. 3
- [12] Maximilian Denninger, Dominik Winkelbauer, Martin Sundermeyer, Wout Boerdijk, Markus Knauer, Klaus H. Strobl, Matthias Humt, and Rudolph Triebel. Blenderproc2: A procedural pipeline for photorealistic rendering. *Journal of Open Source Software*, 8(82):4901, 2023. 7
- [13] Zhiyang Dou, Qingxuan Wu, Cheng Lin, Zeyu Cao, Qiangqiang Wu, Weilin Wan, Taku Komura, and Wenping Wang. Tore: Token reduction for efficient human mesh recovery with transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15143–15155, 2023. 2
- [14] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. In *ICRA*, 2022. 3, 6, 7, 8
- [15] Ziya Erkoç, Fangchang Ma, Qi Shan, Matthias Nießner, and Angela Dai. Hyperdiffusion: Generating implicit neural fields with weight-space diffusion. *arXiv preprint arXiv:2303.17015*, 2023. 2, 3
- [16] 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. *NeurIPS*, 2022. 2, 3
- [17] Jiatao Gu, Qingzhe Gao, Shuangfei Zhai, Baoquan Chen, Lingjie Liu, and Josh Susskind. Learning controllable 3d diffusion models from single-view images. *arXiv preprint arXiv:2304.06700*, 2023. 2, 3
- [18] Jiatao Gu, Alex Trevithick, Kai-En Lin, Joshua M Susskind, Christian Theobalt, Lingjie Liu, and Ravi Ramamoorthi. Nerfdiff: Single-image view synthesis with nerf-guided distillation from 3d-aware diffusion. In *ICML*, 2023. 3
- [19] Yuan-Chen Guo. Instant neural surface reconstruction, 2022. <https://github.com/bennyguo/instant-nsr-pl>. 7
- [20] Yuan-Chen Guo, Ying-Tian Liu, Ruizhi Shao, Christian Laforte, Vikram Voleti, Guan Luo, Chia-Hao Chen, Zi-Xin Zou, Chen Wang, Yan-Pei Cao, and Song-Hai Zhang. threestudio: A unified framework for 3d content generation. <https://github.com/threestudio-project/threestudio>, 2023. 7
- [21] Anchit Gupta, Wenhan Xiong, Yixin Nie, Ian Jones, and Barlas Ögüz. 3dgen: Triplane latent diffusion for textured mesh generation. *arXiv preprint arXiv:2303.05371*, 2023. 2, 3
- [22] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 3, 4
- [23] Yukun Huang, Jianan Wang, Yukai Shi, Xianbiao Qi, Zheng-Jun Zha, and Lei Zhang. Dreamtime: An improved optimization strategy for text-to-3d content creation. *arXiv preprint arXiv:2306.12422*, 2023. 3
- [24] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 867–876, 2022. 2
- [25] Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuexin Ma. Gaussianshader: 3d gaussian splatting with shading functions for reflective surfaces. *arXiv preprint arXiv:2311.17977*, 2023. 2
- [26] Heewoo Jun and Alex Nichol. Shap-e: Generating conditional 3d implicit functions. *arXiv preprint arXiv:2305.02463*, 2023. 2, 3, 7
- [27] Animesh Karnewar, Niloy J Mitra, Andrea Vedaldi, and David Novotny. Holofusion: Towards photo-realistic 3d generative modeling. In *ICCV*, 2023.
- [28] 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 *CVPR*, 2023. 2, 3

- [29] Kilian Kleeberger, Richard Bormann, Werner Kraus, and Marco F Huber. A survey on learning-based robotic grasping. *Current Robotics Reports*, 1:239–249, 2020. [2](#)
- [30] Jiabao Lei, Jiapeng Tang, and Kui Jia. Generative scene synthesis via incremental view inpainting using rgb-d diffusion models. In *CVPR*, 2022. [3](#)
- [31] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *CVPR*, 2023. [2](#), [3](#)
- [32] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Zexiang Xu, and Hao Su. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *arXiv preprint arXiv:2306.16928*, 2023. [7](#)
- [33] 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](#), [4](#), [6](#), [7](#), [8](#)
- [34] Xinhang Liu, Shiu-hong Kao, Jiaben Chen, Yu-Wing Tai, and Chi-Keung Tang. Deceptive-nerf: Enhancing nerf reconstruction using pseudo-observations from diffusion models. *arXiv preprint arXiv:2305.15171*, 2023. [3](#)
- [35] Yuan Liu, Sida Peng, Lingjie Liu, Qianqian Wang, Peng Wang, Christian Theobalt, Xiaowei Zhou, and Wenping Wang. Neural rays for occlusion-aware image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7824–7833, 2022. [2](#)
- [36] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Generating multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453*, 2023. [2](#), [3](#), [4](#), [7](#), [8](#)
- [37] Zhen Liu, Yao Feng, Michael J Black, Derek Nowrouzezahrai, Liam Paull, and Weiyang Liu. Meshdiffusion: Score-based generative 3d mesh modeling. In *ICLR*, 2023. [2](#), [3](#)
- [38] Xiaoxiao Long, Cheng Lin, Lingjie Liu, Wei Li, Christian Theobalt, Ruigang Yang, and Wenping Wang. Adaptive surface normal constraint for depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12849–12858, 2021. [2](#)
- [39] Xiaoxiao Long, Cheng Lin, Peng Wang, Taku Komura, and Wenping Wang. Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In *European Conference on Computer Vision*, pages 210–227. Springer, 2022. [7](#)
- [40] Xiaoxiao Long, Yuhang Zheng, Yupeng Zheng, Beiwen Tian, Cheng Lin, Lingjie Liu, Hao Zhao, Guyue Zhou, and Wenping Wang. Adaptive surface normal constraint for geometric estimation from monocular images. *arXiv preprint arXiv:2402.05869*, 2024. [2](#)
- [41] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2837–2845, 2021. [2](#), [3](#)
- [42] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Realfusion: 360deg reconstruction of any object from a single image. In *CVPR*, 2023. [2](#), [3](#), [7](#), [8](#)
- [43] 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. In *ECCV*, 2020. [5](#)
- [44] Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Buló, Peter Kotschieder, and Matthias Nießner. Diffrf: Rendering-guided 3d radiance field diffusion. In *CVPR*, 2023. [2](#), [3](#)
- [45] 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](#), [7](#)
- [46] Evangelos Ntavelis, Aliaksandr Siarohin, Kyle Olszewski, Chaoyang Wang, Luc Van Gool, and Sergey Tulyakov. Autodecoding latent 3d diffusion models. *arXiv preprint arXiv:2307.05445*, 2023. [2](#), [3](#)
- [47] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *ICLR*, 2023. [2](#), [3](#), [7](#)
- [48] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint arXiv:2306.17843*, 2023. [2](#), [3](#), [7](#)
- [49] 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 *ICML*, 2021. [2](#), [3](#), [4](#)
- [50] Amit Raj, Srinivas Kaza, Ben Poole, Michael Niemeyer, Nataniel Ruiz, Ben Mildenhall, Shiran Zada, Kfir Aberman, Michael Rubinstein, Jonathan Barron, et al. Dreambooth3d: Subject-driven text-to-3d generation. *arXiv preprint arXiv:2303.13508*, 2023. [3](#)
- [51] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. [3](#), [7](#)
- [52] Hoigi Seo, Hayeon Kim, Gwanghyun Kim, and Se Young Chun. Ditto-nerf: Diffusion-based iterative text to omnidirectional 3d model. *arXiv preprint arXiv:2304.02827*, 2023. [3](#)
- [53] 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](#)
- [54] Qihong Shen, Xingyi Yang, and Xinchao Wang. Anything-3d: Towards single-view anything reconstruction in the wild. *arXiv preprint arXiv:2304.10261*, 2023. [3](#)
- [55] Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. *arXiv preprint arXiv:2308.16512*, 2023. [2](#), [3](#), [4](#)
- [56] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. [3](#), [4](#)

- [57] Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. Viewset diffusion:(0-) image-conditioned 3d generative models from 2d data. *arXiv preprint arXiv:2306.07881*, 2023. 3
- [58] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-it-3d: High-fidelity 3d creation from a single image with diffusion prior. In *ICCV*, 2023. 3
- [59] Shitao Tang, Fuyang Zhang, Jiacheng Chen, Peng Wang, and Yasutaka Furukawa. Mvdiffusion: Enabling holistic multi-view image generation with correspondence-aware diffusion. *arXiv preprint arXiv:2307.01097*, 2023. 3
- [60] Ayush Tewari, Tianwei Yin, George Cazenavette, Semon Rezhikov, Joshua B Tenenbaum, Frédo Durand, William T Freeman, and Vincent Sitzmann. Diffusion with forward models: Solving stochastic inverse problems without direct supervision. *arXiv preprint arXiv:2306.11719*, 2023. 3
- [61] Christina Tsalicoglou, Fabian Manhardt, Alessio Tonioni, Michael Niemeyer, and Federico Tombari. Textmesh: Generation of realistic 3d meshes from text prompts. *arXiv preprint arXiv:2304.12439*, 2023. 3
- [62] Hung-Yu Tseng, Qinbo Li, Changil Kim, Suhub Alsisan, Jia-Bin Huang, and Johannes Kopf. Consistent view synthesis with pose-guided diffusion models. In *CVPR*, 2023. 3
- [63] 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. 2, 3
- [64] 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. In *NeurIPS*, 2021. 5
- [65] 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 *CVPR*, 2023. 2, 3
- [66] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 2004. 7, 8
- [67] 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
- [68] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel view synthesis with diffusion models. *arXiv preprint arXiv:2210.04628*, 2022. 3
- [69] Jinbo Wu, Xiaobo Gao, Xing Liu, Zhengyang Shen, Chen Zhao, Haocheng Feng, Jingtuo Liu, and Errui Ding. Hd-fusion: Detailed text-to-3d generation leveraging multiple noise estimation. *arXiv preprint arXiv:2307.16183*, 2023. 3
- [70] Jianfeng Xiang, Jiaolong Yang, Binbin Huang, and Xin Tong. 3d-aware image generation using 2d diffusion models. *arXiv preprint arXiv:2303.17905*, 2023. 3
- [71] Deja Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. Neurallift-360: Lifting an in-the-wild 2d photo to a 3d object with 360 views. *arXiv e-prints*, pages arXiv–2211, 2022. 3
- [72] Paul Yoo, Jiaxian Guo, Yutaka Matsuo, and Shixiang Shane Gu. Dreamsparse: Escaping from plato’s cave with 2d frozen diffusion model given sparse views. *CoRR*, 2023. 3
- [73] Chaohui Yu, Qiang Zhou, Jingliang Li, Zhe Zhang, Zhibin Wang, and Fan Wang. Points-to-3d: Bridging the gap between sparse points and shape-controllable text-to-3d generation. *arXiv preprint arXiv:2307.13908*, 2023. 3
- [74] Jason J. Yu, Fereshteh Forghani, Konstantinos G. Derpanis, and Marcus A. Brubaker. Long-term photometric consistent novel view synthesis with diffusion models. In *ICCV*, 2023. 3
- [75] Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. Lion: Latent point diffusion models for 3d shape generation. In *NeurIPS*, 2022. 2, 3
- [76] Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models. In *SIGGRAPH*, 2023. 2, 3
- [77] Jingbo Zhang, Xiaoyu Li, Ziyu Wan, Can Wang, and Jing Liao. Text2nerf: Text-driven 3d scene generation with neural radiance fields. *arXiv preprint arXiv:2305.11588*, 2023. 3
- [78] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 7, 8
- [79] Yuhang Zheng, Xiangyu Chen, Yupeng Zheng, Songen Gu, Runyi Yang, Bu Jin, Pengfei Li, Chengliang Zhong, Zengmao Wang, Lina Liu, et al. Gaussiangrasper: 3d language gaussian splatting for open-vocabulary robotic grasping. *arXiv preprint arXiv:2403.09637*, 2024. 2
- [80] Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5826–5835, 2021. 2, 3
- [81] Zhizhuo Zhou and Shubham Tulsiani. Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. In *CVPR*, 2023. 3
- [82] Joseph Zhu and Peiye Zhuang. Hifa: High-fidelity text-to-3d with advanced diffusion guidance. *arXiv preprint arXiv:2305.18766*, 2023. 3