

# Cascade-Zero123: One Image to Highly Consistent 3D with Self-Prompted Nearby Views

Yabo Chen<sup>1\*</sup>, Jiemin Fang<sup>2\*†</sup>, Yuyang Huang<sup>1</sup>, Taoran Yi<sup>3</sup>,  
Xiaopeng Zhang<sup>2✉</sup>, Lingxi Xie<sup>2</sup>, Xinggang Wang<sup>3</sup>, Wenrui Dai<sup>1✉</sup>,  
Hongkai Xiong<sup>1</sup>, and Qi Tian<sup>2</sup>

<sup>1</sup> Shanghai Jiao Tong University, Shanghai, China

<sup>2</sup> Huawei Inc., Shenzhen, China

<sup>3</sup> Huazhong University of Science and Technology, Wuhan, China  
{chenyabo, huangyuyang, daiwenrui, xionghongkai}@sjtu.edu.cn  
jaminfong@gmail.com, {taoranyi, xgwang}@hust.edu.cn  
zpxhistory@gmail.com, 198808xc@gmail.com, tian.qi1@huawei.com



**Fig. 1:** Rather than adopting limited input information, which Zero-1-to-3 [43] generation pipeline only has a single-view source image, Cascade-Zero123 progressively extracts the 3D information from more condition images by self-prompting. View-consistent images can be generated in a cascade manner. Cascade-Zero123 shows the strong capability on various complex objects, *e.g.* insects, robots, or multiple objects stacked.

**Abstract.** Synthesizing multi-view 3D from one single image is a significant but challenging task. Zero-1-to-3 methods have achieved great success by lifting a 2D latent diffusion model to the 3D scope. The target-view image is generated with a single-view source image and the camera

\*Equal contribution. †Project lead.

✉Correspondence to Xiaopeng Zhang and Wenrui Dai.

pose as condition information. However, due to the high sparsity of the single input image, Zero-1-to-3 tends to produce geometry and appearance inconsistency across views, especially for complex objects. To tackle this issue, we propose to supply more condition information for the generation model but in a self-prompt way. A cascade framework is constructed with two Zero-1-to-3 models, named Cascade-Zero123, which progressively extract 3D information from the source image. Specifically, several nearby views are first generated by the first model and then fed into the second-stage model along with the source image as generation conditions. With amplified self-prompted condition images, our Cascade-Zero123 generates more consistent novel-view images than Zero-1-to-3. Experiment results demonstrate remarkable promotion, especially for various complex and challenging scenes, involving insects, humans, transparent objects, and stacked multiple objects *etc.* More demos and code are available at <https://cascadezero123.github.io>.

**Keywords:** Single Image to 3D · Novel View Synthesis · Cascade Networks · Diffusion Models

## 1 Introduction

Generating the 3D object from a single image has become an appealing research topic, for its flexibility and convenience of creating 3D assets, which can be applied to various real-world applications including virtual reality, computer gaming, movie and animation production [10, 13, 31, 38, 44, 49, 56, 57, 64, 76, 77, 82, 97, 110]. However, due to the input information is too limited – only one image, it is quite challenging to reconstruct a high-fidelity 3D object. Many research works attempt to lift strong generation power from 2D large latent diffusion models [53, 63, 65, 67] to the 3D fields [2, 3, 8, 19–21, 29, 39, 45, 48, 54, 58, 72, 84, 87, 91, 92, 94, 98, 99, 106, 109], bringing possibilities to untractable 3D generation tasks.

As a representative work, Zero-1-to-3 [43] innovatively proposes to generate novel view images by conditioning the diffusion model with the source image and corresponding camera pose transition. The Zero-1-to-3 diffusion model is tuned on a large-scale multi-view 3D dataset, Objaverse [15, 16], and has achieved remarkable success and generalizability in synthesizing novel views from any single image. However, still suffering from the information sparsity of a single image, Zero-1-to-3 may inevitably fail to keep geometry or appearance consistency across different views, especially for objects with complex structures. As shown in Fig. 2, when the camera pose changes drastically, *e.g.* rotating more than  $45^\circ$ , it is hard to guarantee a high synthesis quality.

Providing more input information beyond a single image will undoubtedly ease the difficulty of generating 3D objects. We expect to improve the Zero-1-to-3 model by supplying more condition images from different views but in a self-prompting way. The main idea is to design a cascade framework to achieve this goal, which contains two Zero-1-to-3 models, *i.e.* Base-0123 and Refiner-0123. The Base-0123 model is conditioned to generate multiple images from



**Fig. 2:** The performance comparison of Zero-1-to-3 [43] and our methods on Google Scanned Object [18] with different camera pose rotation angles. When the camera pose changes drastically, the synthesis quality of Zero-1-to-3 will drop drastically. But our method can promote the synthesis quality in all the transition ranges of camera poses.

various views. These images are constrained in a small range, *e.g.*  $45^\circ$  in our implementation, to guarantee reliable quality as observed in Fig. 2. The Refiner-0123 model takes in all these generated images along with the source image to output the target-view image. We name the proposed framework as Cascade-Zero123. Rather than directly generating the target view as in Zero-1-to-3, Cascade-Zero123 lowers the synthesis difficulty by progressively digging out the 3D information starting from the input source image. To further enhance the Base-0123 model, we propose to update its parameters with a self-distillation manner from Refiner-0123. As shown in Fig. 2 and Fig. 1, our method can consistently promote the synthesis quality in all the transition ranges of camera poses.

Our contributions can be summarized as follows:

- We propose a cascade framework to progressively extract 3D information from a single image, preventing inconsistency caused by drastic view transitions.
- We design a self-prompting method that generates multiple nearby viewpoints, providing reliable condition information for the second-stage Refiner-0123 model.
- Our proposed Cascade-Zero123 has shown notable promotion on complex scenes that are challenging for Zero-1-to-3, such as insects, humans, robots, or multiple objects stacked together.

## 2 Related Work

### 2.1 Single Image to 3D

Many researchers have studied the tasks of generating 3D models and achieving novel view synthesis using only a single image [12, 32, 37, 41, 74, 75, 90, 93, 104, 107] and text [7, 14, 26, 34, 51, 52, 60, 68, 78, 83, 103, 105]. Some researchers directly train a 3D model on 3D data [6, 25, 30, 33, 54, 62, 95], but they tend to have good generation quality only on scenes similar to the training set.

Recently, by constructing a conditional latent diffusion model based on camera viewpoints, many works have made it possible to pre-train a single image-to-3D model [22, 24, 40, 42, 43, 46, 47, 50, 61, 69–71, 73, 79, 80, 88, 89, 96, 100, 101].

Zero-1-to-3 [43] learns from large-scale multi-view images [16] to build the geometric priors of large-scale diffusion models. Zero-1-to-3 can lift various images that training sets have never been seen before to 3D with good quality. After that, many works have utilized Zero-1-to-3 as a module to enhance the quality of meshes or 3D models. Magic123 [61] combines the capabilities of Zero-1-to-3 and stable diffusion together to generate 3D models. One-2-3-45 [42] also leverages Zero-1-to-3 to generate different views to assist in mesh generation. However, these methods simply use it as pre-trained 3D diffusion model tools with fixed checkpoints.

There have also been efforts to improve Zero-1-to-3. Approaches like Consistent1-to-3 [101], aim to enhance the consistency of view generation by introducing priors during the denoising process in view-conditioned diffusion models. However, they add models like additional Transformers to render new views which will result in poor generalization due to these additional model’s capability, they may not integrate well with generalized latent diffusion models. In addition, the training of additional models usually costs a lot.


Therefore, we argue that Zero-1-to-3 itself has the ability to provide 3D priors and establish 3D consistency, albeit requiring a progressive manner to achieve it. By using self-prompted condition images, Zero-1-to-3 can generate images of higher quality and higher consistency while also maintaining its strong open-set generalization capabilities brought by the large-scale latent diffusion model.

## 2.2 Multi-stage Diffusion Models

Cascade networks have been widely used in boosting the performance of models in the computer vision field, such as cascade RCNN [4], cascade DETR [102], and so on [35, 36]. After large-scale latent diffusion models become popular, many influential works also use cascade networks to improve the generation quality of models. For example, DeepFloyd-IF [66] achieves high-resolution and high-detail image generation by constructing three cascade pixel diffusion modules: a base model that generates a  $64 \times 64$  pixel image based on text prompts, and two super-resolution models. Furthermore, SDXL [59] uses a cascade structure, in which a base model aligned with a refiner in the latent space to generate higher-quality and higher-resolution images. I2VGen-XL [108] also employs cascade diffusion models to enhance the quality of video generation. Inspired by these cascade network designs, we attempt to progressively extract 3D information from a single image to improve the novel view synthesis quality.

## 3 Methods

We present the Cascade-Zero123 approach from the following aspects. First, we briefly review the framework of Zero-1-to-3 [43] in Sec. 3.1. Then, we explain how we construct the cascade structure of Zero-1-to-3 in Sec. 3.2. Next, we discuss the design of Base-0123 in Sec. 3.3 and Refiner-0123 to use self-prompted views in Sec. 3.4. Finally, we describe the self-distillation design in Sec. 3.5 and the inference process of the model in Sec. 3.6.



**Fig. 3:** The architecture of Cascade-Zero123. Cascade-Zero123 can be divided into two parts. The left part is Base-0123, which takes a set of R and T values as input to generate corresponding multi-view images. These output images are concatenated with the input condition image and its corresponding camera pose, forming a self-prompted input denoted as a set of  $c(x_c, \Delta R, \Delta T)$  for the right part Refiner-0123. The corresponding camera pose transition for each condition image to the target image needs to be recalculated as shown in detailed camera pose rotations. After each iteration of training, Base-0123 is updated through exponential moving average (EMA) using Refiner-0123.

### 3.1 Preliminary

We first give a brief introduction to diffusion models [17]. The diffusion model’s forward process adds the Gaussian noise  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ , which can be defined as:

$$q(x^t | x^{t-1}) = \mathcal{N}(x^t; \sqrt{\alpha_t}x^{t-1}, (1 - \alpha_t)\mathbf{I}), \quad (1)$$

where  $\alpha$  is a scheduling hyper-parameter and  $t \in [1, 1000]$  denotes the diffusion timestep.  $q(x^t | x^{t-1})$  estimate the probability of  $x^t$  using  $x^{t-1}$ , and  $\mathbf{I}$  is the normally distributed variance.

Zero-1-to-3 [43] proposes a latent diffusion model to learn the relationship between the source image  $x_c$  and target image  $x_t$  during the denoising process, which can be simply defined as:

$$x_t = f_{\phi}(x_c, \Delta R^{ct}, \Delta T^{ct}). \quad (2)$$

The Zero-1-to-3 model takes an image  $x_c$  and the relative camera pose  $(\Delta R^{ct}, \Delta T^{ct})$  as the condition. For simplicity, we use  $\Delta$  to represent the pose transition, including the relative angle rotation and translation from the condition view  $(R^c, T^c)$  to the target view  $(R^t, T^t)$ , *i.e.*  $(\Delta R^{ct}, \Delta T^{ct}) = (R^t, T^t) \ominus (R^c, T^c)$ , where  $\ominus$  is the pose transition in the corresponding world coordinate and  $R$  and  $T$  represent the rotation and translation matrix of the camera pose respectively. Specifically,

the pose transition computation of the Objaverse dataset coordinate [15,16] can be found in the appendix.

Using a latent diffusion model with an encoder  $f_\phi$ , a denoiser U-Net  $\epsilon_\theta$  and a decoder  $\mathcal{D}$ , the denoising process can be defined as follows:

$$\begin{aligned} p(x^{t-1} | x^t, c(x_c, \Delta R^{ct}, \Delta T^{ct})) \\ = \mathcal{N}(x^{t-1}; \mu_\theta(x^t, t, c(x_c, \Delta R^{ct}, \Delta T^{ct})), \\ \Sigma_\theta(x^t, t, c(x_c, \Delta R^{ct}, \Delta T^{ct}))). \end{aligned} \quad (3)$$

where  $c(x_c, \Delta R^{ct}, \Delta T^{ct})$  represents the embedding encoded from the condition image and relative camera pose and the mean distribution  $\mu_\theta$  and variance function  $\Sigma_\theta$  is modeled by the denoising U-Net.

At the diffusion time step  $t$ , Zero-1-to-3 encodes the embedding of the input view and relative camera pose as  $c(x_c, \Delta R^{ct}, \Delta T^{ct})$  and set the objective loss function as:

$$L(\theta) = \min_{\theta} \mathbb{E}_{z \sim \mathcal{E}(x), t, \epsilon \sim \mathcal{N}(0,1)} \|\epsilon - \epsilon_\theta(z_t, t, c(x_c, \Delta R^{ct}, \Delta T^{ct}))\|_2^2, \quad (4)$$

where  $\epsilon$  is the noise prediction corresponding to the distribution. With the model  $\epsilon_\theta$  trained, the inference model  $f_\phi$  can generate an image by denoising the Gaussian noise conditioned on the embedding of  $c(x_c, \Delta R^{ct}, \Delta T^{ct})$ .

### 3.2 Cascade-Zero123 Framework

Suffering from the information sparsity of a single image, it is highly challenging for Zero-1-to-3 [43] to generate geometry and appearance consistent novel views with a single source image. We ease this task with a cascade structure, which progressively moves the camera pose with slight changes.

As shown in Fig. 3. Cascade-Zero123 consists of two cascade Zero-1-to-3 models. The first Zero-1-to-3 is referred to as Base-0123, and the second one is referred to as Refiner-0123. In terms of the structure, Base-0123 is responsible for generating multi-view images. These images serve as rough condition inputs, which, along with their corresponding camera poses, are fed into the network of the Refiner-0123 model. The Refiner-0123 is aware of the multi-view input and it will compute the rotation and translation of different viewpoints to the final target view. Supplying more condition images from different views to Refiner-0123 can ease the difficulty of generating novel views. After that, the training framework incorporates the model parameters of the second Refiner-0123 model back into the Base-0123 through exponential moving average (EMA).

Compared with Eq. 2, we can formulate the proposed Cascade-Zero123 as:

$$x_T = f_{\phi_r}((x_c^0, f_{\phi_b}(x, \Delta R_b^{c\{1:P\}}, \Delta T_b^{c\{1:P\}})), \Delta R_r^{\{1:P\}t}, \Delta T_r^{\{1:P\}t}), \quad (5)$$

where  $f_{\phi_b}$  and  $f_{\phi_r}$  are the Base-0123 and Refiner-0123 respectively, and the parameters of  $f_{\phi_b}$  can be denoted as  $\phi_b$  and those parameters of  $f_{\phi_r}$  can be denoted as  $\phi_r$ .  $x_T$  is the final target image.  $P$  is the number of the prompt views,  $\Delta R_r^{it}, \Delta T_r^{it}$  is the prompted views camera pose of Refiner-0123. For easy understanding, we include Table 1 to list some symbols mentioned in the paper along with their meanings.

**Table 1:** Table of mentioned symbol representation of this paper.

Symbol	Meaning
$x_c, x_t$	source image and target image
$\Delta R^{ct}, \Delta T^{ct}$	rotation and translation from the condition pose to the target pose
$\Delta R_b, \Delta T_b$	rotation and translation condition to Base-0123
$\Delta R_r, \Delta T_r$	rotation and translation condition to Refiner-0123
$f_{\phi_b}, f_{\phi_r}$	Base-0123 model and Refiner-0123 model
$\ominus$	pose transitions in the world coordinate
$\hat{x}_t^1, \dots, \hat{x}_t^P$	self-prompted views

### 3.3 Base-0123 Framework

In particular, at the beginning of the pre-training process, a set of input images denoted as  $x_c^0$  is provided to Base-0123. Then we sample some rotation and translation matrices from the nearby viewpoints of the input view, which we denote as  $\{(R^i, T^i) | i \in [1, P]\}$ , where  $P$  is the total number of prompt nearby views. We set the poses of  $P$  viewpoints as constant ones to avoid the gap between the training and inference phases. Constant poses can also save computation resources due to that these poses can be reused across different targets.

Our ablation experiments have also demonstrated that generating prompt images with small angles produces good performance. Detailed view setting and hyper-parameter selection can be found in the implementation details of Sec. 4.2.

All the  $P$  prompted viewpoints attended with the target view pose  $(R_t, T_t)$  are concatenated with the same input embedding as  $c(x_c^0, \Delta R_b^{ci}, \Delta T_b^{ci})$ ,  $i \in [1, P]$ . Specifically, we calculate:

$$(\Delta R_b^{ci}, \Delta T_b^{ci}) = (R^i, T^i) \ominus (R^c, T^c) \quad i \in [1, P], \quad (6)$$

where  $\ominus$  are the pose transitions in the corresponding world coordinate. In this stage, all these poses are drawn as target views. They are fed into the first Base-0123  $f_\phi$  in parallel.

Compared with Eq. 2, we can formulate the Base-0123 framework as:

$$\begin{aligned} \{\hat{x}_t^1, \hat{x}_t^2, \dots, \hat{x}_t^P\} = & \{f_{\phi_b}(x_c^0, \Delta R_b^{c1}, \Delta T_b^{c1}), \\ & f_{\phi_b}(x_c^0, \Delta R_b^{c2}, \Delta T_b^{c2}), \dots \\ & f_{\phi_b}(x_c^0, \Delta R_b^{cP}, \Delta T_b^{cP})\}, \end{aligned} \quad (7)$$

where  $\{\hat{x}_t^1, \hat{x}_t^2, \dots, \hat{x}_t^P\}$  are exactly self-prompted views.

### 3.4 Refiner-0123 Framework

After obtaining the multi-view images generated by the first-stage Base-0123 and their corresponding camera poses, we proceed to compute the camera pose rotation changes for each of these images with respect to the final target image.

Specifically, we calculate the relative pose transition from self-prompted views to the target views as:

$$(\Delta R_r^{it}, \Delta T_r^{it}) = (R^t, T^t) \ominus (R^i, T^i) \quad i \in [1, P], \quad (8)$$

where  $\ominus$  are the pose transitions in the corresponding world coordinate, and  $(R_r, T_r)$  is the camera pose of the target image. Because the camera rotation and translation from the target view to itself are zero and the rotation of the input view remains the same, *i.e.*  $(\Delta R_r^{tt}, \Delta T_r^{tt}) = (0, 0)$ ,  $(\Delta R_r^{ct}, \Delta T_r^{ct}) = (\Delta R_c^{ct}, \Delta T_c^{ct})$ . Then as shown in Fig. 3. The second Refiner-0123 takes the input images and self-prompted nearby views  $\{\hat{x}_t^1, \hat{x}_t^2, \dots, \hat{x}_t^P\}$  as input. Compared with Eq. 2, we can formulate the Refiner-0123 framework as:

$$x_T = f_{\phi_r}((x_c^0, \Delta R_r^{ct}, \Delta T_r^{ct}), (\hat{x}_t^1, \Delta R_r^{1t}, \Delta T_r^{1t}), \dots, (\hat{x}_t^P, \Delta R_r^{Pt}, \Delta T_r^{Pt}), (\hat{x}_t^t, \Delta R_r^{tt}, \Delta T_r^{tt})). \quad (9)$$

Reviewing that the conditional denoising autoencoder can control the synthesis process through inputs context  $y$  such as text, semantic maps or images [65], the latent diffusion models use attention-based models. The context cross-attention can be formulated as:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right) \cdot V. \quad (10)$$

$$Q = W_Q^{(i)} \cdot \varphi_i(z_t), K = W_K^{(i)} \cdot \tau_{\phi_r}(y), V = W_V^{(i)} \cdot \tau_{\phi_r}(y). \quad (11)$$

Here,  $\varphi_i(z_t) \in \mathbb{R}^{N \times d_i^\epsilon}$  denotes a (flattened) intermediate representation of the UNet implementing  $\epsilon_{\phi_r}$  and  $W_V^{(i)} \in \mathbb{R}^{d \times d_i^\epsilon}$ ,  $W_Q^{(i)} \in \mathbb{R}^{d \times d_\tau}$  and  $W_K^{(i)} \in \mathbb{R}^{d \times d_\tau}$  are learnable projection matrices. All prompt views are concatenated through the token dimension so that all these conditions are set as the context of the input  $\tau_{\phi_r}(y)$ , all conditional embeddings can calculate cross attention to the flattened intermediate representation  $\epsilon_{\phi_r}$ .

Then, the objective loss function of the Refiner-0123 model can be set as:

$$\min_{\theta} \mathbb{E}_{z \sim \mathcal{E}(x), t, \epsilon \sim \mathcal{N}(0,1)} \|\epsilon - \epsilon_{\phi_r}(z_r, t, c(x_c^0, \hat{x}_t^{\{1:p\}}, \Delta R_r^{\{1:p\}}, \Delta T_r^{\{1:p\}}))\|_2^2. \quad (12)$$

After the model  $\epsilon_{\phi_r}$  is trained, we can get the Refiner-0123 model  $f_{\phi_r}$ .

### 3.5 Self-Distillation Design

Considering that the Refiner-0123 network  $f_{\phi_r}$  is trained via back-propagation to minimize the denoising loss. Inspired by the methods [5, 9, 11, 23] in the self-supervised learning field. Our ablation experiments demonstrated that even if only accepted a single image, Refiner-0123 still improves performance. To further enhance the Base-0123 model, the Base-0123 network is updated in a momentum update way using exponential moving average (EMA). Specifically, we have denoted the parameters of  $f_{\phi_b}$  as  $\phi_b$  and those of  $f_{\phi_r}$  as  $\phi_r$  we update  $\phi_b$  by

$$\phi_b \leftarrow \eta \cdot \phi_b + (1 - \eta) \cdot \phi_r, \quad (13)$$

where  $\eta \in [0, 1)$  is a momentum coefficient to control the magnitude decay of updates from the Refiner-0123 to the Base-0123.

### 3.6 Inference

During the inference stage, given an input image, two stages of the DDIM (Denoising Diffusion Probabilistic Model) [28] schedule are performed. The input image passes through both the Base-0123 and Refiner-0123 simultaneously.

These views along with their corresponding camera pose rotation and translation differences, are concatenated and input into the Refiner-0123. After computing cross-attention with the context and inputs, another DDIM sampling is performed, generating the final target image.

For the generation from a single image to 3D, the incurred cost is quite similar to the original Zero-1-to-3. Refiner-0123 leverages these condition images to calculate the SDS loss without repeatedly invoking the denoising process of Base-0123. The process is similar to novel view synthesis, but the Base-0123 only needs to be processed once. All the self-prompted nearby views along with the input views are used as condition images and fed into the Refiner-0123 network for computing Score Distillation Sampling (SDS) [60] loss.

## 4 Experiments

We assess our model’s performance on novel view synthesis and single-image to 3D reconstruction tasks. We introduce the datasets, implementation details, and metrics from Sec.4.1 to Sec.4.3. And show qualitative results, quantitative results on Sec.4.4. Ablation studies are introduced in supplementary materials.

### 4.1 Datasets

**Objaverse Dataset** is a large-scale dataset containing 800K+ annotated 3D mesh objects [55]. We use this dataset for training and validation. We render 12 images per object from uniformly distributed viewpoints followed by [43].

**Realfusion15** Realfusion15 is the dataset collected and released by RealFusion [50], consisting of 15 natural images.

**Google Scanned Object (GSO)** [18] is a high-quality scanned household items dataset. We adopt the same data split as in SyncDreamer [46].

**RTMV** [81] consists of complex scenes, and each scene is composed of 20 random objects. We adopt the same data split as in Zero-1-to-3 [43].

### 4.2 Implementation Details


We train Cascade-Zero123 on the Objaverse dataset [16] which contains about 800k objects. Following Zero-1-to-3, the number of viewpoints rendered is 12. Following the assumption of Zero-1-to-3, we also assume that the azimuth of



**Fig. 4:** Novel view synthesis compared with Zero123-XL [15], and SyncDreamer [46], where Zero123-XL is Zero-1-to-3 pre-trained on Objaverse-XL datasets [15], achieving higher performance. We selected some challenging scenes, including stacked objects, parallel objects, and objects with multiple branches. Zero123-XL exhibits good quality in image generation but lacks consistency in these complex scenes. SyncDreamer demonstrates good consistency but struggles to maintain good quality in image generation. Our model, however, maintains both quality and consistency in these scenarios.

both the input view and the first target view is  $0^\circ$ . We train the Cascade-Zero123 for 200k steps with 8 V100 GPUs using a total batch size of 96.

Both Base-0123 and Refiner-0123 load the pre-trained Zero123-XL model initially. In Base-0123, to avoid increasing the pretraining costs, we perform only 25 iterations of DDIM (Differentiable Diffusion Model) for inference on each input view of the object during the pretraining phase. Multiple nearby views share the same input view but concatenate with different camera poses. In Refiner-0123, we choose nearby views of azimuth rotations of  $45^\circ$  and  $-45^\circ$ , a view of elevation rotations of 30 degrees, and the target view generated in the first stage is also set as inputs to Refiner-0123. More details about the selection of nearby views can be found in the ablation studies.



**Fig. 5:** Single image to 3D reconstruction using SDS loss [60] compared with Zero123-XL. The first two rows illustrate that Cascade-Zero123 can correct the problem that Zero-1-to-3 sometimes learns inaccurate colors of the backside. The middle two lines describe how Cascade-Zero123 can rectify structural errors through multi-view self-prompting. The last two lines indicate that Cascade-Zero123 can address the problem of transparent or high-brightness objects being mistakenly learned as white clouds.

### 4.3 Metrics

We use the following evaluation metrics to quantitatively evaluate the performance of our model. We first report the Peak Signal-to-Noise Ratio (PSNR). Perceptual Loss (LPIPS) measures the perceptual distance between two images by comparing the deep features extracted by deep neural networks given each image as input. Structural Similarity (SSIM) measures the structural similarity between two images considering both color and texture information. The CLIP-score (CLIP) quantifies the average CLIP distance between the rendered image and the reference image, serving as a measure of 3D consistency by assessing appearance similarity across novel views and the reference view. We measured Chamfer Distance (C Dist.) to evaluate point-by-point shape similarity and volumetric IoU (Vol IoU) to quantify the overlap between the reconstructed and the ground truth shape.

### 4.4 Qualitative Results

**Novel View Synthesis** We show qualitative results generated by our Cascade-Zero123 in Fig. 1 and Fig. 4. Our model can generalize well to unseen data.

**Table 2:** Quantitative results on Objaverse. We evaluate our method on the test split of Objaverse [16].

Methods	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$
Zero123-XL [15, 43]	18.68	0.883	0.189	0.758
Magic123 [61]	18.95	0.882	0.167	0.778
Ours	<b>21.42</b>	<b>0.911</b>	<b>0.125</b>	<b>0.802</b>

We selected various images from real-world scenes or high-quality scenes generated by Stable Diffusion 2.1 [65] for the experiments on novel view synthesis. These scenes include different kinds of environments and objects. We also tested scenarios involving multiple object stacking (such as stacked donuts) and complex branching structures (such as ladybugs and peeled bananas). The selected example images were deliberately chosen to avoid central symmetry or left-right symmetry, which can pose challenges for generating novel views using conditional latent diffusion models. In the case of these difficult-to-maintain consistency scenes, our Cascade-Zero123 achieved better consistency compared to Zero123-XL and significantly improved image quality compared to SyncDreamer. Note that Zero123-XL is a model pre-trained on Objaverse-XL [15, 16] and has superior generation quality.

**Single Image to 3D Using SDS Loss** We show qualitative results of a single image to 3D using Dreamfusion’s Score Distillation Sampling (SDS) loss [60]. All samples are generated by our Cascade-Zero123 and baseline Zero123-XL in Fig. 5. Firstly, as shown in the upper two rows of Fig. 5, after using SDS loss, Zero-1-to-3 tends to learn the backside as the smooth color of the front side. This is due to that the model can only perceive information from the prompt of a single image. As a result, the ears of the fox, the back and hair of Lisa (the woman in yellow), the back of the dinosaur, and the glass tiger are all learned as the same color. With the correction provided by self-prompted views, we no longer have to speculate the color of the backside from scratch. Instead, we can progressively predict the backside based on the side views. This allows us to better learn the true color consistency of unseen views.

In addition, the proposal of self-prompted views has enhanced the ability of Zero-1-to-3 to model the shapes and textures of objects as shown in the middle of Fig. 5. The shapes of the ice cream are well preserved, and even the complex structure of the ninja’s back can be modeled. Moreover, Zero-1-to-3, when combined with SDS, has difficulty modeling objects with transparency. Zero-1-to-3 sometimes learns the backside of transparent or high-brightness objects as pure white mist-like clouds. However, the multi-view conditioning information prevents the model from lacking information about other sides, avoiding the learning of a pure white plane for the backside and providing 3D information for objects with transparency.

**Table 3:** Evaluation of novel view synthesis on GSO dataset [18].

Methods	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$
Realfusion [50]	15.26	0.722	0.283	-
Zero-1-to-3 [43]	18.93	0.779	0.166	-
SyncDreamer [46]	20.05	0.798	0.146	-
Zero123-XL [15]	18.81	0.828	0.142	0.813
Ours(Zero123-XL)	<b>20.35</b>	<b>0.850</b>	<b>0.113</b>	<b>0.846</b>
Ours(SyncDreamer)	<b>21.10</b>	<b>0.862</b>	<b>0.095</b>	<b>0.944</b>

**Table 4:** Evaluation of novel view synthesis on RTMV dataset [81].

Methods	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$
DietNeRF [27]	7.13	0.406	0.507	-
Image Variation [1]	6.56	0.442	0.564	-
SJC-I [85]	7.95	0.456	0.545	-
Zero123-XL [15]	10.97	0.573	0.465	0.702
Ours	<b>11.12</b>	<b>0.617</b>	<b>0.441</b>	<b>0.734</b>

**Table 5:** Quantitative results on RealFusion15. Evaluation of novel-view synthesis on the RealFusion15 dataset [50].

Methods	PSNR $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$
RealFusion [50]	20.216	0.197	0.735
Make-it-3D [79]	20.010	0.119	0.839
Zero-1-to-3 [43]	25.386	0.068	0.759
Zero123-XL [15]	25.220	0.050	0.897
Magic123 [61]	25.637	0.062	0.747
Consistent-123 [40]	25.682	0.056	0.844
Ours	<b>26.098</b>	<b>0.043</b>	<b>0.916</b>

**Table 6:** Evaluation of single image to 3D reconstruction on GSO dataset [18].

Methods	C Dist. $\downarrow$	Vol IoU $\uparrow$
Realfusion [50]	0.0819	0.2741
Magic123 [61]	0.0516	0.4528
One-2-3-45 [42]	0.0629	0.4086
Point-E [54]	0.0426	0.2875
Shap-E [30]	0.0436	0.3584
Zero123 [43]	0.0339	0.5035
SyncDreamer [46]	0.0261	0.5421
Ours	<b>0.0207</b>	<b>0.5792</b>

#### 4.5 Quantitative Results

**Novel view synthesis on Objaverse testset.** Following prior research [88], which randomly picked up 100 objects from the Objaverse testset, Since the entire Objaverse test set [16] is quite large, testing all samples would require an excessively long time. Following the approach of Consistent123 [88], we randomly selected a subset of samples from the Objaverse test set for testing. These selected samples have not been seen during the training process. Compared with Consistent123 [88], we take a larger number of samples for evaluation (Consistent123 only takes 100 samples). Specifically, we randomly selected 200 samples and then randomly selected an input view and a target view following the same setting as Zero-1-to-3. On the Objaverse test set, we measured the reconstruction performance and consistency using metrics of PSNR, SSIM, LPIPS, and CLIP-score (CLIP). We utilize the checkpoints from Zero123-XL [15, 43]. The index of the selected samples will be made available in the open-source release.

**Novel view synthesis on Google Scanned Object (GSO).** Following SyncDreamer [46] and Zero-1-to-3 [43], we adopt the Google Scanned Object [18] dataset as the evaluation dataset. We followed the setting of SyncDreamer [46] that chose the same 30 objects ranging from daily objects to animals of Google Scanned Object. We also apply the proposed cascade mechanism to SyncDreamer [46], named as Ours(SyncDreamer). The self-prompted views generated from the first-stage Base-SyncDreamer are fed into the Synchronized

Multiview Noise Predictor in the second-stage Refiner-SyncDreamer, as shown in Table 3. We replicate Zero123-XL [15] for fair comparison. We measured the reconstruction performance and consistency using metrics of PSNR, SSIM, LPIPS, and CLIP-score (CLIP). Our cascade method also shows notable improvement on both Zero123-XL and SyncDreamer.

**Novel view synthesis on RTMV.** In Table 4, we use the same experimental setup as Zero1-to-3 [43] on RTMV [81] and replicate the performance of the Zero123-XL [15]. We measured the reconstruction performance and consistency using metrics of PSNR, SSIM, LPIPS, and CLIP-score (CLIP).

**Novel view synthesis on Realfusion15 testset.** Following Magic123 [61] and Consistent123 [40] We evaluate Cascade-Zero123 against many related baselines, including RealFusion [50], Make-it-3D [79], Zero-1-to-3 [43] and Magic123 [61], on the RealFusion15 datasets. Like Magic123, we measure the PSNR, LPIPS, and CLIP-scores (CLIP), which assess the reconstruction quality and visual consistency. As shown in Table. 5.

**Single image to 3D reconstruction on GSO.** In Table 6, we compare our approach with various other methods following the setting of SyncDreamer [46]. Both our method and SDS-free methods utilize NeuS [86], a neural reconstruction method for converting multi-view images into 3D shapes. To achieve faithful reconstruction of 3D mesh that aligns well with ground truth, the generated multi-view images should be geometrically coherent. Cascade-Zero123 achieves the best results in both Chamfer distance (C Dist.) and volumetric IoU (Vol IoU) metrics, demonstrating proficiency in producing multi-view consistent images.

## 5 Limitation and Conclusion

*Limitation* Cascade-Zero123 is based on the pre-trained Zero-1-to-3 model [43]. For cases that are extremely difficult for Zero-1-to-3, Cascade-Zero123 has limited ability to handle. With 2D image input, it is challenging to figure out the exact overlap, so even with nearby views, the overlapping parts may still appear to be stuck together. The 3D depth information will be lost and a flat structure will be potentially learned. Additionally, Zero-1-to-3 is sensitive to camera pose elevation. Therefore, Cascade-Zero123 also struggles with input images that have a high elevation. While our model has already achieved performance and consistency improvement compared to previous models, there are still some potential future directions, *e.g.* enhancing Cascade-Zero123 by incorporating attention between multiple views as [72] and involving multi-modal conditioning information, *e.g.* depth or normal.

*Conclusion* We have witnessed Zero-1-to-3’s success in synthesizing novel views from a single image. However, inconsistency still appears as the input information is too limited. We propose a Cascade-Zero123 network to progressively dig out the multi-view information from the input image by self-prompting nearby views. A series of experimental proofs have demonstrated the effectiveness of this cascade design. We believe the cascade framework can be compatible with more Zero-1-to-3 variants and leave this as an important future work.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant 62125109, Grant 61931023, Grant 61932022, Grant 62371288, Grant 62320106003, Grant 62301299, Grant T2122024, Grant 62120106007.

## References

1. Stable diffusion image variations. - a hugging face space by lambdalabs (2023)
2. Alldieck, T., Kolotouros, N., Sminchisescu, C.: Score distillation sampling with learned manifold corrective. arXiv:2401.05293 (2024)
3. Armandpour, M., Zheng, H., Sadeghian, A., Sadeghian, A., Zhou, M.: Re-imagine the negative prompt algorithm: Transform 2d diffusion into 3d, alleviate janus problem and beyond. arXiv:2304.04968 (2023)
4. Cai, Z., Vasconcelos, N.: Cascade r-cnn: High quality object detection and instance segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019)
5. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: *Proceedings of the International Conference on Computer Vision (ICCV)* (2021)
6. Chan, E.R., Nagano, K., Chan, M.A., Bergman, A.W., Park, J.J., Levy, A., Aitala, M., Mello, S.D., Karras, T., Wetzstein, G.: GeNVS: Generative novel view synthesis with 3D-aware diffusion models. In: arXiv (2023)
7. Chen, M., Yuan, W., Wang, Y., Sheng, Z., He, Y., Dong, Z., Bo, L., Guo, Y.: Sketch2nerf: Multi-view sketch-guided text-to-3d generation. arXiv:2401.14257 (2024)
8. Chen, R., Chen, Y., Jiao, N., Jia, K.: Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. arXiv:2303.13873 (2023)
9. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv (2020)
10. Chen, X., Mihajlovic, M., Wang, S., Prokudin, S., Tang, S.: Morphable diffusion: 3d-consistent diffusion for single-image avatar creation. arXiv:2401.04728 (2024)
11. Chen, Y., Liu, Y., Jiang, D., Zhang, X., Dai, W., Xiong, H., Tian, Q.: Sdae: Self-distilled masked autoencoder. In: *ECCV* (2022)
12. Chen, Y., Ni, J., Jiang, N., Zhang, Y., Zhu, Y., Huang, S.: Single-view 3d scene reconstruction with high-fidelity shape and texture. arXiv:2311.00457 (2023)
13. Chen, Y., Xie, R., Ye, Q., Yang, S., Xie, Z., Chen, T., Li, R., Huo, Y.: 2l3: Lifting imperfect generated 2d images into accurate 3d. arXiv:2401.15841 (2024)
14. Chen, Z., Wang, F., Liu, H.: Text-to-3d using gaussian splatting. arXiv:2309.16585 (2023)
15. Deitke, M., Liu, R., Wallingford, M., Ngo, H., Michel, O., Kusupati, A., Fan, A., Laforte, C., Voleti, V., Gadre, S.Y., VanderBilt, E., Kembhavi, A., Vondrick, C., Gkioxari, G., Ehsani, K., Schmidt, L., Farhadi, A.: Objaverse-xl: A universe of 10m+ 3d objects. arXiv preprint arXiv:2307.05663 (2023)
16. Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., VanderBilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A universe of annotated 3d objects. In: *CVPR*. pp. 13142–13153 (2023)
17. Dhariwal, Prafulla, , Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* (2021)

18. Downs, L., Francis, A., Koenig, N., Kinman, B., Hickman, R., Reymann, K., McHugh, T.B., Vanhoucke, V.: Google scanned objects: A high-quality dataset of 3d scanned household items. In: 2022 International Conference on Robotics and Automation (ICRA). pp. 2553–2560 (2022). <https://doi.org/10.1109/ICRA46639.2022.9811809>
19. Fang, J., Wang, J., Zhang, X., Xie, L., Tian, Q.: Gaussianeditor: Editing 3d gaussians delicately with text instructions. arXiv preprint arXiv:2311.16037 (2023)
20. Gao, J., Shen, T., Wang, Z., Chen, W., Yin, K., Li, D., Litany, O., Gojcic, Z., Fidler, S.: Get3d: A generative model of high quality 3d textured shapes learned from images. *Advances In Neural Information Processing Systems* **35**, 31841–31854 (2022)
21. Gupta, A., Xiong, W., Nie, Y., Jones, I., Oğuz, B.: 3dgen: Triplane latent diffusion for textured mesh generation. arXiv:2303.05371 (2023)
22. Hamdi, A., Ghanem, B., Nießner, M.: Sparf: Large-scale learning of 3d sparse radiance fields from few input images. In: ICCV. pp. 2930–2940 (2023)
23. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR. pp. 9729–9738 (2020)
24. Hu, S., Hong, F., Hu, T., Pan, L., Mei, H., Xiao, W., Yang, L., Liu, Z.: Humanliff: Layer-wise 3d human generation with diffusion model. arXiv:2308.09712 (2023)
25. Huang, Z., Stojanov, S., Thai, A., Jampani, V., Rehg, J.M.: Zeroshape: Regression-based zero-shot shape reconstruction. arXiv:2312.14198 (2024)
26. Jain, A., Mildenhall, B., Barron, J.T., Abbeel, P., Poole, B.: Zero-shot text-guided object generation with dream fields. In: CVPR. pp. 867–876 (2022)
27. Jain, A., Tancik, M., Abbeel, P.: Putting nerf on a diet: Semantically consistent few-shot view synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 5885–5894 (October 2021)
28. Jiaming Song, Chenlin Meng, S.E.: Denoising diffusion implicit models. ICLR (2021)
29. Jun, H., Nichol, A.: Shap-e: Generating conditional 3d implicit functions. arXiv:2305.02463 (2023)
30. Jun, H., Nichol, A.e.a.: Shap-e: Generating conditional 3d implicit functions. arXiv:2305.02463 (2023)
31. Kant, Y., Wu, Z., Vasilkovsky, M., Qian, G., Ren, J., Guler, R.A., Ghanem, B., Tulyakov, S., Gilitschenski, I., Siarohin, A.: Spad : Spatially aware multiview diffusers. arXiv:2402.05235 (2024)
32. Kocsis, P., Sitzmann, V., Nießner, M.: Intrinsic image diffusion for single-view material estimation. arXiv:2312.12274 (2023)
33. Lee, D., Kim, C., Cho, M., Han, W.S.: Locality-aware generalizable implicit neural representation. In: arXiv:2310.05624 (2023)
34. Lei, J., Zhang, Y., Jia, K., et al.: Tango: Text-driven photorealistic and robust 3d stylization via lighting decomposition. *Advances in Neural Information Processing Systems* **35**, 30923–30936 (2022)
35. Li, H., Shi, B., Dai, W., Chen, Y., Wang, B., Sun, Y.: Hierarchical graph networks for 3d human pose estimation. arXiv:2111.11927 (2021)
36. Li, K., Wang, S., Zhang, X., Xu, Y., Xu, W., Tu, Z.: Pose recognition with cascade transformers
37. Li, S., Zanjani, F.G., Yahia, H.B., Asano, Y.M., Gall, J., Habibiyan, A.: Valid: Variable-length input diffusion for novel view synthesis. arXiv:2312.08892 (2023)
38. Li, Z., Litvak, D., Li, R., Zhang, Y., Jakab, T., Rupperecht, C., Wu, S., Vedaldi, A., Wu, J.: Learning the 3d fauna of the web. arXiv:2401.02400 (2024)

39. Lin, C.H., Gao, J., Tang, L., Takikawa, T., Zeng, X., Huang, X., Kreis, K., Fidler, S., Liu, M.Y., Lin, T.Y.: Magic3d: High-resolution text-to-3d content creation. In: CVPR. pp. 300–309 (2023)
40. Lin, Y., Han, H., Gong, C., Xu, Z., Zhang, Y., Li, X.: Consistent123: One image to highly consistent 3d asset using case-aware diffusion priors. arXiv:2309.17261 (2023)
41. Liu, M., Shi, R., Chen, L., Zhang, Z., Xu, C., Wei, X., Chen, H., Zeng, C., Gu, J., Su, H.: One-2-3-45++: Fast single image to 3d objects with consistent multi-view generation and 3d diffusion. arXiv:2311.07885 (2023)
42. Liu, M., Xu, C., Jin, H., Chen, L., Xu, Z., Su, H., et al.: One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. arXiv:2306.16928 (2023)
43. Liu, R., Wu, R., Hoorick, B.V., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. arXiv:2303.11328 (2023)
44. Liu, T., Zhao, H., Yu, Y., Zhou, G., Liu, M.: Car-studio: Learning car radiance fields from single-view and unlimited in-the-wild images. IEEE Robotics and Automation Letters pp. 2024–2031 (2024). <https://doi.org/10.1109/LRA.2024.3349949>
45. Liu, X., Kao, S.H., Chen, J., Tai, Y.W., Tang, C.K.: Deceptive-nerf: Enhancing nerf reconstruction using pseudo-observations from diffusion models. arXiv:2305.15171 (2023)
46. Liu, Y., Lin, C., Zeng, Z., Long, X., Liu, L., Komura, T., Wang, W.: Sync-dreamer: Generating multiview-consistent images from a single-view image. arXiv:2309.03453 (2023)
47. Long, X., Guo, Y.C., Lin, C., Liu, Y., Dou, Z., Liu, L., Ma, Y., Zhang, S.H., Habermann, M., Theobalt, C., et al.: Wonder3d: Single image to 3d using cross-domain diffusion. arXiv:2310.15008 (2023)
48. Luo, T., Rockwell, C., Lee, H., Johnson, J.: Scalable 3d captioning with pretrained models. arXiv:2306.07279 (2023)
49. Melas-Kyriazi, L., Laina, I., Rupprecht, C., Neverova, N., Vedaldi, A., Gafni, O., Kokkinos, F.: Im-3d: Iterative multiview diffusion and reconstruction for high-quality 3d generation. arXiv:2402.08682 (2024)
50. Melas-Kyriazi, L., Laina, I., Rupprecht, C., Vedaldi, A.: Realfusion: 360deg reconstruction of any object from a single image. In: CVPR. pp. 8446–8455 (2023)
51. Michel, O., Bar-On, R., Liu, R., Benaim, S., Hanocka, R.: Text2mesh: Text-driven neural stylization for meshes. In: CVPR. pp. 13492–13502 (2022)
52. Mohammad Khalid, N., Xie, T., Belilovsky, E., Popa, T.: Clip-mesh: Generating textured meshes from text using pretrained image-text models. In: SIGGRAPH Asia 2022 conference papers. pp. 1–8 (2022)
53. Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv:2112.10741 (2021)
54. Nichol, A., Jun, H., Dhariwal, P., Mishkin, P., Chen, M.: Point-e: A system for generating 3d point clouds from complex prompts. arXiv:2212.08751 (2022)
55. Ouyang, Y., Chai, W., Ye, J., Tao, D., Zhan, Y., Wang, G.: Chasing consistency in text-to-3d generation from a single image. arXiv:2309.03599 (2023)
56. Paliwal, A., Nguyen, B., Tsarov, A., Kalantari, N.K.: Reshader: View-dependent highlights for single image view-synthesis. arXiv:2309.10689 (2023)
57. Pan, X., Yang, Z., Bai, S., Yang, Y.: Gd<sup>2</sup>-nerf: Generative detail compensation via gan and diffusion for one-shot generalizable neural radiance fields. arXiv:2401.00616 (2024)

58. Pan, Z., Yang, Z., Zhu, X., Zhang, L.: Fast dynamic 3d object generation from a single-view video. arXiv:2401.08742 (2024)
59. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. arXiv:2307.01952 (2023)
60. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. arXiv (2022)
61. Qian, G., Mai, J., Hamdi, A., Ren, J., Siarohin, A., Li, B., Lee, H.Y., Skorokhodov, I., Wonka, P., Tulyakov, S., et al.: Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. arXiv:2306.17843 (2023)
62. Qian, X., Wang, Y., Luo, S., Zhang, Y., Tai, Y., Zhang, Z., Wang, C., Xue, X., Zhao, B., Huang, T., Wu, Y., Fu, Y.: Pushing auto-regressive models for 3d shape generation at capacity and scalability. arXiv:2402.12225 (2024)
63. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv:2204.06125 (2022)
64. Roessle, B., Müller, N., Porzi, L., Bulò, S.R., Kotschieder, P., Nießner, M.: Ganerf: Leveraging discriminators to optimize neural radiance fields. In: arXiv:2306.06044. (2023)
65. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022)
66. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L.: Photorealistic text-to-image diffusion models with deep language understanding. Advances in Neural Information Processing Systems (2022)
67. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. Advances in Neural Information Processing Systems (2022)
68. Sanghi, A., Chu, H., Lambourne, J.G., Wang, Y., Cheng, C.Y., Fumero, M., Malekshan, K.R.: Clip-forge: Towards zero-shot text-to-shape generation. In: CVPR. pp. 18603–18613 (2022)
69. Sargent, K., Li, Z., Shah, T., Herrmann, C., Yu, H.X., Zhang, Y., Chan, E.R., Lagun, D., Fei-Fei, L., Sun, D., et al.: Zeronvs: Zero-shot 360-degree view synthesis from a single real image. arXiv:2310.17994 (2023)
70. Shen, Q., Yang, X., Wang, X.: Anything-3d: Towards single-view anything reconstruction in the wild. arXiv:2304.10261 (2023)
71. Shi, R., Chen, H., Zhang, Z., Liu, M., Xu, C., Wei, X., Chen, L., Zeng, C., Su, H.: Zero123++: a single image to consistent multi-view diffusion base model. arXiv:2310.15110 (2023)
72. Shi, Y., Wang, P., Ye, J., Long, M., Li, K., Yang, X.: Mvdream: Multi-view diffusion for 3d generation. arXiv:2308.16512 (2023)
73. Shi, Y., Wang, J., Cao, H., Tang, B., Qi, X., Yang, T., Huang, Y., Liu, S., Zhang, L., Shum, H.Y.: Toss: High-quality text-guided novel view synthesis from a single image. arXiv:2310.10644 (2023)
74. Shi, Y., Wang, J., Cao, H., Tang, B., Qi, X., Yang, T., Huang, Y., Liu, S., Zhang, L., Shum, H.Y.: Toss:high-quality text-guided novel view synthesis from a single image. arXiv:2310.10644 (2023)
75. Simon, C., He, S., Perez-Rua, J.M., Xu, M., Benhalloum, A., Xiang, T.: Hypervoltran: Fast and generalizable one-shot image to 3d object structure via hypernetworks. arXiv:2312.16218 (2024)
76. Spiegl, B., Perin, A., Deny, S., Ilin, A.: Viewfusion: Learning composable diffusion models for novel view synthesis. arXiv:2402.02906 (2024)

77. Tang, J., Chen, Z., Chen, X., Wang, T., Zeng, G., Liu, Z.: Lgm: Large multi-view gaussian model for high-resolution 3d content creation. arXiv:2402.05054 (2024)
78. Tang, J., Ren, J., Zhou, H., Liu, Z., Zeng, G.: Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. arXiv:2309.16653 (2023)
79. Tang, J., Wang, T., Zhang, B., Zhang, T., Yi, R., Ma, L., Chen, D.: Make-it-3d: High-fidelity 3d creation from a single image with diffusion prior. arXiv:2303.14184 (2023)
80. Tang, S., Chen, J., Wang, D., Tang, C., Zhang, F., Fan, Y., Chandra, V., Furukawa, Y., Ranjan, R.: Mvdifffusion++: A dense high-resolution multi-view diffusion model for single or sparse-view 3d object reconstruction. arXiv:2402.12712 (2024)
81. Tremblay, J., Meshry, M., Evans, A., Kautz, J., Keller, A., Khamis, S., Müller, T., Loop, C., Morrical, N., Nagano, K., Takikawa, T., Birchfield, S.: Rtmv: A ray-traced multi-view synthetic dataset for novel view synthesis. arXiv:2205.07058 (2022)
82. Vainer, S., Boss, M., Parger, M., Kutsy, K., Nigris, D.D., Rowles, C., Perony, N., Donné, S.: Collaborative control for geometry-conditioned pbr image generation. arXiv:2402.05919 (2024)
83. Wang, C., Chai, M., He, M., Chen, D., Liao, J.: Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In: CVPR. pp. 3835–3844 (2022)
84. Wang, H., Du, X., Li, J., Yeh, R.A., Shakhnarovich, G.: Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In: CVPR. pp. 12619–12629 (2023)
85. Wang, H., Du, X., Li, J., Yeh, R.A., Shakhnarovich, G.: Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12619–12629 (June 2023)
86. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. arXiv:2106.10689 (2023)
87. Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. arXiv:2305.16213 (2023)
88. Weng, H., Yang, T., Wang, J., Li, Y., Zhang, T., Chen, C., Zhang, L.: Consistent123: Improve consistency for one image to 3d object synthesis. arXiv:2310.08092 (2023)
89. Weng, Z., Wang, Z., Yeung, S.: Zeroavatar: Zero-shot 3d avatar generation from a single image. arXiv:2305.16411 (2023)
90. Woo, S., Park, B., Go, H., Kim, J.Y., Kim, C.: Harmonyview: Harmonizing consistency and diversity in one-image-to-3d. arXiv:2312.15980 (2023)
91. Wu, C.H., Chen, Y.C., Solarte, B., Yuan, L., Sun, M.: ifusion: Inverting diffusion for pose-free reconstruction from sparse views. arXiv:2312.17250 (2023)
92. Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Xinggang, W.: 4d gaussian splatting for real-time dynamic scene rendering. arXiv preprint arXiv:2310.08528 (2023)
93. Wu, T., Li, Z., Yang, S., Zhang, P., Pan, X., Wang, J., Lin, D., Liu, Z.: Hyperdreamer: Hyper-realistic 3d content generation and editing from a single image. In: SIGGRAPH Asia 2023 Conference Papers (2023). <https://doi.org/10.1145/3610548.3618168>

94. Wu, Z., Li, Y., Yan, H., Shang, T., Sun, W., Wang, S., Cui, R., Liu, W., Sato, H., Li, H., Ji, P.: Blockfusion: Expandable 3d scene generation using latent tri-plane extrapolation. arXiv:2401.17053 (2024)
95. Xiang, J., Yang, J., Huang, B., Tong, X.: 3d-aware image generation using 2d diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2383–2393 (October 2023)
96. Xu, D., Jiang, Y., Wang, P., Fan, Z., Wang, Y., Wang, Z.: Neurallift-360: Lifting an in-the-wild 2d photo to a 3d object with 360deg views. In: CVPR. pp. 4479–4489 (2023)
97. Xu, D., Yuan, Y., Mardani, M., Liu, S., Song, J., Wang, Z., Vahdat, A.: Agg: Amortized generative 3d gaussians for single image to 3d. arXiv:2401.04099 (2024)
98. Xu, J., Wang, X., Cheng, W., Cao, Y.P., Shan, Y., Qie, X., Gao, S.: Dream3d: Zero-shot text-to-3d synthesis using 3d shape prior and text-to-image diffusion models. In: CVPR. pp. 20908–20918 (2023)
99. Yang, C., Li, S., Fang, J., Liang, R., Xie, L., Zhang, X., Shen, W., Tian, Q.: Gaussianobject: Just taking four images to get a high-quality 3d object with gaussian splatting. arXiv:2402.10259 (2024)
100. Yang, J., Cheng, Z., Duan, Y., Ji, P., Li, H.: Consistnet: Enforcing 3d consistency for multi-view images diffusion. arXiv:2310.10343 (2023)
101. Ye, J., Wang, P., Li, K., Shi, Y., Wang, H.: Consistent-1-to-3: Consistent image to 3d view synthesis via geometry-aware diffusion models. arXiv:2310.03020 (2023)
102. Ye, M., Ke, L., Li, S., Tai, Y.W., Tang, C.K., Danelljan, M., Yu, F.: Cascade-detr: Delving into high-quality universal object detection. In: ICCV (2023)
103. Yi, T., Fang, J., Wu, G., Xie, L., Zhang, X., Liu, W., Tian, Q., Wang, X.: Gaussiandreamer: Fast generation from text to 3d gaussian splatting with point cloud priors. arxiv:2310.08529 (2023)
104. Yu, K., Liu, J., Feng, M., Cui, M., Xie, X.: Boosting3d: High-fidelity image-to-3d by boosting 2d diffusion prior to 3d prior with progressive learning. arXiv:2311.13617 (2023)
105. Yu, Y., Zhu, S., Qin, H., Li, H.: Boostdream: Efficient refining for high-quality text-to-3d generation from multi-view diffusion. arXiv:2401.16764 (2024)
106. Zeng, X., Chen, X., Qi, Z., Liu, W., Zhao, Z., Wang, Z., Fu, B., Liu, Y., Yu, G.: Paint3d: Paint anything 3d with lighting-less texture diffusion models. arXiv:2312.13913 (2023)
107. Zhang, J., Tang, Z., Pang, Y., Cheng, X., Jin, P., Wei, Y., Ning, M., Yuan, L.: Repaint123: Fast and high-quality one image to 3d generation with progressive controllable 2d repainting. arXiv:2312.13271 (2023)
108. Zhang, S., Wang, J., Zhang, Y., Zhao, K., Yuan, H., Qing, Z., Wang, X., Zhao, D., Zhou, J.: I2vgen-xl: High-quality image-to-video synthesis via cascaded diffusion models (2023)
109. Zhao, M., Zhao, C., Liang, X., Li, L., Zhao, Z., Hu, Z., Fan, C., Yu, X.: Efficient-dreamer: High-fidelity and robust 3d creation via orthogonal-view diffusion prior. arXiv:2308.13223 (2023)
110. Zheng, X.Y., Pan, H., Guo, Y.X., Tong, X., Liu, Y.: Mvd<sup>2</sup>: Efficient multiview 3d reconstruction for multiview diffusion. arXiv:2402.14253 (2024)