

# SV3D: Novel Multi-view Synthesis and 3D Generation from a Single Image using Latent Video Diffusion

Vikram Voleti\* Chun-Han Yao\* Mark Boss\* Adam Letts David Pankratz  
Dmitry Tochilkin Christian Laforte Robin Rombach Varun Jampani\*

Stability AI




Figure 1. **Stable Video 3D (SV3D)**. From a single image, SV3D generates consistent novel multi-view images. We then optimize a 3D representation with SV3D generated views resulting in high-quality 3D meshes.

## Abstract

We present *Stable Video 3D (SV3D)* — a latent video diffusion model for high-resolution, image-to-multi-view generation of orbital videos around a 3D object. Recent work on 3D generation propose techniques to adapt 2D generative models for novel view synthesis (NVS) and 3D optimization. However, these methods have several disadvantages due to either limited views or inconsistent NVS, thereby affecting the performance of 3D object generation. In this work, we propose SV3D that adapts image-to-video diffusion model for novel multi-view synthesis and 3D generation, thereby leveraging the generalization and multi-view consistency of the video models, while further adding explicit camera control for NVS. We also propose improved 3D optimization techniques to use SV3D and its NVS outputs for image-to-3D generation. Extensive experimental results on multiple datasets with 2D and 3D metrics as well as user study demonstrate SV3D’s state-of-the-art performance on NVS as well as 3D reconstruction compared to prior works.

\* Core technical contribution.

## 1. Introduction

Single-image 3D object reconstruction is a long-standing problem in computer vision with a wide range of applications in game design, AR/VR, e-commerce, robotics, etc. It is a highly challenging and ill-posed problem as it requires lifting 2D pixels to 3D space while also reasoning about the unseen portions of the object in 3D.

Despite being a long-standing vision problem, it is only recently that practically useful results are being produced by leveraging advances in the generative AI. This is mainly made possible by the large-scale pretraining of generative models which enables sufficient generalization to various domains. A typical strategy is to use image-based 2D generative models (e.g., Imagen [39], Stable Diffusion (SD) [37]) to provide a 3D optimization loss function for unseen novel views of a given object [20, 28, 34]. In addition, several works repurpose these 2D generative models to perform novel view synthesis (NVS) from a single image [23, 24, 27, 50], and then use the synthesized novel views for 3D generation. Conceptually, these works mimic a typical photogrammetry-based 3D object capture pipeline, i.e., first photographing multi-view images of an object, followed by 3D optimization; except that the explicit multi-view capture is replaced by novel-view synthesis using a

generative model, either via text prompt or camera pose control.

A key issue in these generation-based reconstruction methods is the lack of multi-view consistency in the underlying generative model resulting in inconsistent novel views. Some works [6, 15, 24] try to address the multi-view consistency by jointly reasoning a 3D representation during NVS, but this comes at the cost of high computational and data requirements, often still resulting in unsatisfactory results with inconsistent geometric and texture details. In this work, we tackle these issues by adapting a high-resolution, image-conditioned video diffusion model for NVS followed by 3D generation.

**Novel Multi-view Synthesis.** We adapt a latent video diffusion model (Stable Video Diffusion - SVD [2]) to generate multiple novel views of a given object with *explicit camera pose conditioning*. SVD demonstrates excellent *multi-view consistency* for video generation, and we repurpose it for NVS. In addition, SVD also has good *generalization* capabilities as it is trained on large-scale image and video data that are more readily available than large-scale 3D data. In short, we adapt the video diffusion model for NVS from a single image with three useful properties for 3D object generation: *pose-controllable*, *multi-view consistent*, and *generalizable*. We call our resulting NVS network ‘SV3D’. To our knowledge, this is the first work that adapts a video diffusion model for explicit pose-controlled view synthesis. Some contemporary works such as [2, 26] demonstrate the use of video models for view synthesis, but they typically only generate orbital videos without any explicit camera control.

**3D Generation.** We then utilize our SV3D model for 3D object generation by optimizing a NeRF and DMTet mesh in a coarse-to-fine manner. Benefiting from the multi-view consistency in SV3D, we are able to produce high-quality 3D meshes directly from the SV3D novel view images. We also design a masked score distillation sampling (SDS) [34] loss to further enhance 3D quality in the regions that are not visible in the SV3D-predicted novel views. In addition, we propose to jointly optimize a disentangled illumination model along with 3D shape and texture, effectively reducing the issue of baked-in lighting.

We perform extensive comparisons of both our NVS and 3D generation results with respective state-of-the-art methods, demonstrating considerably better outputs with SV3D. For NVS, SV3D shows high-level of multi-view consistency and generalization to real-world images while being pose controllable. Our resulting 3D meshes are able to capture intricate geometric and texture details.

## 2. Background

### 2.1. Novel View Synthesis

We organize the related works along three crucial aspects of novel view synthesis (NVS): generalization, controllability, and multi-view (3D) consistency.

**Generalization.** Diffusion models [14, 45] have recently shown to be powerful generative models that can generate a wide variety of images [3, 37, 38] and videos [2, 12, 49] by iteratively denoising a noise sample. Among these models, the publicly available Stable Diffusion (SD) [37] and Stable Video Diffusion (SVD) [2] demonstrate strong generalization ability by being trained on extremely large datasets like LAION [40] and LVD [2]. Hence, they are commonly used as foundation models for various generation tasks, *e.g.* novel view synthesis.

**Controllability.** Ideally, a controllable NVS model allows us to generate an image corresponding to any arbitrary viewpoint. For this, Zero123 [23] repurposes an image diffusion model to a novel view synthesizer, conditioned on a single-view image and the pose difference between the input and target views. Follow-up works such as Zero123XL [8] and Stable Zero123 [46] advance the quality of diffusion-based NVS, as well as the trained NeRFs using SDS loss. However, these methods only generate one novel view at a time, and thus are not designed to be multi-view consistent inherently. Recent works such as EscherNet [18] and Free3D [60] are capable of generating with intelligent camera position embedding design encouraging better multi-view consistency. However, they only make use of image-based diffusion models, and generate images at  $256 \times 256$  resolution. We finetune a video diffusion model to generate novel views at  $576 \times 576$  resolution.

**Multi-view Consistency.** Multi-view consistency is the most critical requirement for high-quality NVS and 3D generation. To address the inconsistency issue in prior works, MVDream [43], SyncDreamer [24], HexGen3D [27], and Zero123++ [42] propose to generate multiple (specific) views of an object simultaneously. However, they are not controllable: they only generate specific views given a conditional image, not arbitrary viewpoints. Moreover, they were finetuned from image-based diffusion models, *i.e.* multi-view consistency was imposed on image-based diffusion by adding interaction among the multiple generated views through cross-attention layers. Hence, their output quality is limited to the generalization capability of the image-based model they finetuned from, as well as the 3D dataset they were finetuned on. Efficient-3DiM finetunes the SD model with a stronger vision transformer DINO v2 [33]. Consistent-1-to-3 [55] and SPAD leverage epipolar geometry to generate multiview consistent images. One-2-3-45 [22] and One-2-3-45++ [21] train additional 3D network using the outputs of 2D generator. MVDream [43],

Consistent123 [51], and Wonder3D [25] also train multi-view diffusion models, yet still require post-processing for video rendering. SyncDreamer [24] and ConsistNet [53] employ 3D representation into the latent diffusion model.

**Exploiting Video Diffusion Models.** To achieve better generalization and multi-view consistency, some contemporary works exploit the temporal priors in video diffusion models for NVS. For instance, Vivid-1-to-3 [19] combines a view-conditioned diffusion model and video diffusion model to generate consistent views. SVD-MV [2] and IM-3D [26] finetune a video diffusion model for NVS. However, they generate  $\leq 360^\circ$  views at the same elevation only. Unlike SV3D, none of them are capable of rendering any arbitrary view of the 3D object.

We argue that the existing NVS and 3D generation methods do not fully leverage the superior generalization capability, controllability, and consistency in video diffusion models. In this paper, we fill this important gap and train SV3D, a state-of-the-art novel multi-view synthesis video diffusion model at  $576 \times 576$  resolution, and leverage it for 3D generation.

## 2.2. 3D Generation

Recent advances in 3D representations and diffusion-based generative models have significantly improved the quality of image-to-3D generation. Here, we briefly summarize the related works in these two categories.

**3D Representation.** 3D generation has seen great progress since the advent of Neural Radiance Fields (NeRFs) [29] and its subsequent variants [1], which implicitly represents a 3D scene as a volumetric function, typically parameterized by a neural network. Notably, Instant-NGP [30] introduces a hash grid feature encoding that can be used as a NeRF backbone for fast inference and ability to recover complex geometry. On the other hand, several recent works improve from an explicit representation such as DMTet [41], which is capable of generating high-resolution 3D shapes due to its hybrid SDF-Mesh representation and high memory efficiency. Similar to [20, 35], we adopt coarse-to-fine training for 3D generation, by first learning a rough object with Instant-NGP NeRF and then refining it using the DMTet representation.

**Diffusion-Based 3D Generation.** Several recent works [16, 31] train a 3D diffusion model to learn these flexible 3D representations, which, however, lack generalization ability due to the scarcity of 3D data. To learn 3D generation without ground truth 3D data, image/multi-view diffusion models have been used as guidance for 3D generation. DreamFusion [34] and its follow-up works [20, 28] leverage a trained image diffusion model as a ‘scoring’ function and calculate the SDS loss for text-to-3D generation. However, they are prone to artifacts like Janus problem [28, 34] and over-

saturated texture. Inspired by Zero123 [23], several recent works [18, 21, 22, 24, 26, 35, 42, 46, 60] finetune image or video diffusion models to generate novel view images as a stronger guidance for 3D generation. Our method shares the same spirit as this line of work, but produces denser, controllable, and more consistent multi-view images, thus resulting in better 3D generation quality.

## 3. SV3D: Novel Multi-view Synthesis

Our main idea is to repurpose temporal consistency in a video diffusion model for spatial 3D consistency of an object. Specifically, we finetune SVD [2] to generate an orbital video around a 3D object, conditioned on a single-view image. This orbital video need not be at the same elevation, or at regularly spaced azimuth angles. SVD is well-suited for this task since it is trained to generate smooth and consistent videos on large-scale datasets of real and high-quality videos. The exposure to superior data quantity and quality makes it more *generalizable* and *multi-view consistent*, and the flexibility of the SVD architecture makes it amenable to be finetuned for *camera controllability*.

Some prior works attempt to leverage such properties by finetuning image diffusion models, training video diffusion models from scratch, or finetuning video diffusion models to generate pre-defined views at the same elevation (static orbit) around an object [2, 26]. However, we argue that these methods do not fully exploit the potential of video diffusion models. To the best of our knowledge, SV3D is the first video diffusion-based framework for *controllable* multi-view synthesis at  $576 \times 576$  resolution (and subsequently for 3D generation).

**Problem Setting.** Formally, given an image  $\mathbf{I} \in \mathbb{R}^{3 \times H \times W}$  of an object, our goal is to generate an orbital video  $\mathbf{J} \in \mathbb{R}^{K \times 3 \times H \times W}$  around the object consisting of  $K = 21$  multi-view images along a camera pose trajectory  $\boldsymbol{\pi} \in \mathbb{R}^{K \times 2} = \{(e_i, a_i)\}_{i=1}^K$  as a sequence of  $K$  tuples of elevation  $e$  and azimuth  $a$  angles. We assume that the camera always looks at the center of an object (origin of the world coordinates), and so any viewpoint can be specified by only two parameters: elevation and azimuth. We generate this orbital video by iteratively denoising samples from a learned conditional distribution  $p(\mathbf{J}|\mathbf{I}, \boldsymbol{\pi})$ , parameterized by a video diffusion model.

**SV3D Architecture.** As shown in Fig. 2, the architecture of SV3D builds on that of SVD consisting of a UNet with multiple layers, each layer containing sequences of one residual block with Conv3D layers, and two transformer blocks (spatial and temporal) with attention layers. (i) We remove the vector conditionings of ‘fps id’ and ‘motion bucket id’ since they are irrelevant for SV3D. (ii) The conditioning image is concatenated to the noisy latent state input  $\mathbf{z}_t$  at noise timestep  $t$  to the UNet, after being embedded into

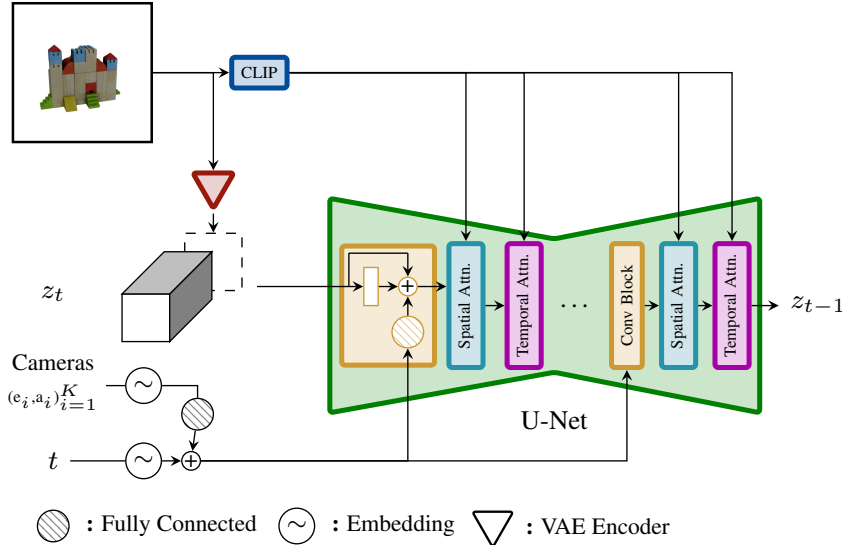


Figure 2. **SV3D Architecture.** We add the sinusoidal embedding of the camera orbit elevation and azimuth angles  $(e, a)$  to that of the noise step  $t$ , and feed the sum to the convolutional blocks in the UNet. We feed the single input image’s CLIP embedding to the attention blocks, and concatenate its latent embedding to the noisy state  $z_t$ .

latent space by the VAE encoder of SVD. (iii) The CLIP-embedding [36] matrix of the conditioning image is provided to the cross-attention layers of each transformer block as its key and value, the query being the feature at that layer. (iv) The camera trajectory is fed into the residual blocks along with the diffusion noise timestep. The camera pose angles  $e_i$  and  $a_i$  and the noise timestep  $t$  are first embedded into sinusoidal position embeddings. Then, the camera pose embeddings are concatenated together, linearly transformed, and added to the noise timestep embedding. This is fed to every residual block, where they are added to the block’s output feature (after being linearly transformed again to match the feature size).

**Static v.s. Dynamic Orbits.** As shown in Fig. 3, we design *static* and *dynamic* orbits to study the effects of camera pose conditioning. In a **static** orbit, the camera circles around an object at regularly-spaced azimuths at the same elevation angle as that in the conditioning image. The disadvantage with the static orbit is that we might not get any information about the top or bottom of the object depending on the conditioning elevation angle. In a **dynamic** orbit, the azimuths can be irregularly spaced, and the elevation can vary per view. To build a dynamic orbit, we sample a static orbit, add small random noise to the azimuth angles, and add a random weighted combination of sinusoids with different frequencies to its elevation. This provides temporal smoothness, and ensures that the camera trajectory loops around to end at the same azimuth and elevation as those of the conditioning image.

Thus, with this strategy, we are able to tackle all three




Figure 3. **Static vs. Dynamic Orbits.** We use two types of orbits for training the SV3D models.

aspects of *generalization*, *controllability*, and *consistency* in novel multi-view synthesis by leveraging video diffusion models, providing the camera trajectory as additional conditioning, and repurposing temporal consistency for spatial 3D consistency of the object, respectively.

**Triangular CFG Scaling.** SVD uses a linearly increasing scale for classifier-free guidance (CFG) from 1 to 4 across the generated frames. However, this scaling causes the last few frames in our generated orbits to be over-sharpened, as shown in Fig. 4 Frame 20. Since we generate videos looping back to the front-view image, we propose to use a triangle wave CFG scaling during inference: linearly increase CFG from 1 at the front view to 2.5 at the back view, then linearly decrease it back to 1 at the front view. Fig. 4 also demonstrates that our triangle CFG scaling produces more details in the back view (Frame 12).

**Models.** We train three image-to-3D-video models finetuned from SVD. First, we train a pose-unconditioned model,  $SV3D^u$ , which generates a video of static orbit around an object while only conditioned on a single-view image. Note that unlike SVD-MV [2], we do not provide the elevation angle to the pose-unconditioned model, as we find that the model is able to infer it from the conditioning image. Our second model, the pose-conditioned  $SV3D^c$  is conditioned on the input image as well as a sequence of camera elevation and azimuth angles in an orbit, trained on dynamic orbits. Following SVD’s [2] intuition to progressively increase the task difficulty during training, we train our third model,  $SV3D^p$ , by first finetuning SVD to produce static orbits unconditionally, then further fine-




Figure 4. **Linear vs. Triangle CFG Scaling.** Notice the increased oversharing in the penultimate frame in the linear scaling vs. our proposed triangle scaling.




Figure 5. **LPIPS vs. Frame Number.** We find that SV3D has the best reconstruction metric per frame.

tuning on dynamic orbits with camera pose condition.

**Training Details.** We train SV3D on the Objaverse dataset [9], which contains synthetic 3D objects covering a wide diversity. For each object, we render 21 frames around it on random color background at  $576 \times 576$  resolution, field-of-view of 33.8 degrees. We choose to finetune the SVD-xt model to output 21 frames. All three models (SV3D<sup>u</sup>, SV3D<sup>c</sup>, SV3D<sup>p</sup>) are trained for 105k iterations in total (SV3D<sup>p</sup> is trained unconditionally for 55k iterations and conditionally for 50k iterations), with an effective batch size of 64 on 4 nodes of 8 80GB A100 GPUs for around 6 days. For more training details, please see the appendix.

### 3.1. Experiments and Results

**Datasets.** We evaluate the SV3D synthesized multi-view images on static and dynamic orbits on the unseen GSO [10] and OmniObject3D [52] datasets. Since many GSO objects are the same items with slightly different colors, we filter 300 objects from GSO to reduce redundancy and maintain diversity. For each object, we render ground truth static and dynamic orbit videos and pick the last frame of each video as the conditioning image. We also conduct a user study on novel view videos from a dataset of 22 real-world images. See the appendix for more details on the datasets.

**Metrics.** We use the SV3D models to generate static and dynamic orbit videos corresponding to the ground truth camera trajectories in the evaluation datasets. We compare each generated frame with the corresponding ground truth frames, in terms of Learned Perceptual Similarity (LPIPS [59]), Peak Signal-to-Noise Ratio (PSNR), Structural SIMilarity (SSIM), Mean Squared-Error (MSE), and CLIP-score (CLIP-S). This range of metrics covers both pixel-level as well as semantic aspects. Note that testing on

dynamic orbits evaluates the controllability of SV3D models, and all the metrics evaluate multi-view consistency.

**Baselines.** We compare SV3D with several recent NVS methods capable of generating arbitrary views, including Zero123 [23], Zero123-XL [8], SyncDreamer [24], Stable Zero123 [46], Free3D [60], EscherNet [18].

**Results.** As shown in Tables 1 to 4, our SV3D models achieve state-of-the-art performance on novel multi-view synthesis. Tables 1 and 3 show results on static orbits, and include all our three models. We see that even our pose-unconditioned model SV3D<sup>u</sup> performs better than all prior methods. Tables 2 and 4 show results on dynamic orbits, and include our pose-conditioned models SV3D<sup>c</sup> and SV3D<sup>p</sup>.

**Ablative Analyses.** Interestingly, from Tables 1 and 3, we find that both SV3D<sup>c</sup> and SV3D<sup>p</sup> outperform SV3D<sup>u</sup> on generations of static orbits, even though SV3D<sup>u</sup> is trained specifically on static orbits. We also observe that SV3D<sup>p</sup> achieves better metrics than SV3D<sup>c</sup> on both static and dynamic orbits, making it the best performing SV3D model overall. This shows that progressive finetuning from easier (static) to harder (dynamic) tasks is indeed a favorable way to finetune a video diffusion model.

**Visual Comparisons** in Fig. 6 further demonstrate that SV3D-generated images are more detailed, faithful to the conditioning image, and multi-view consistent compared to the prior works.

**Quality Per Frame.** We also observe from Fig. 5 that SV3D produces better quality at every frame. We plot the average LPIPS value for each generated frame, across generated GSO static orbit videos. The quality is generally worse around the back view, and better at the beginning and the end (*i.e.* near the conditioning image), as expected.




Figure 6. **Visual Comparison of Novel Multi-view Synthesis.** SV3D is able to generate novel multi-views that are more detailed, faithful to the conditioning image, and multi-view consistent compared to the prior works.

Table 1. Evaluation of novel multi-view synthesis on GSO static orbits

Model	LPIPS↓	PSNR↑	SSIM↑	CLIP-S↑	MSE↓
SyncDreamer [24]	0.17	15.78	0.76	0.87	0.03
Zero123 [23]	0.13	17.29	0.79	0.85	0.04
Zero123XL [8]	0.14	17.11	0.78	0.85	0.04
Stable Zero123 [46]	0.13	18.34	0.78	0.85	0.05
Free3D [60]	0.15	16.18	0.79	0.84	0.04
EscherNet [18]	0.13	16.73	0.79	0.85	0.03
SV3D <sup>st</sup>	0.09	21.14	0.87	<b>0.89</b>	<b>0.02</b>
SV3D <sup>c</sup>	0.09	20.56	0.87	0.88	<b>0.02</b>
SV3D <sup>p</sup>	<b>0.08</b>	<b>21.26</b>	<b>0.88</b>	<b>0.89</b>	<b>0.02</b>

**User Study on Real-World Images.** We conducted a user survey to study human preference between static orbital videos generated by SV3D and by other methods. We asked 30 users to pick one between our SV3D-generated static video and other method-generated video as the best orbital video for the corresponding image, using 22 real-world images. We noted that users preferred SV3D-generated videos over Zero123XL, Stable Zero123, EscherNet, and Free3D, 96%, 99%, 96%, and 98% of the time, respectively.

Table 2. Evaluation of novel multi-view synthesis on GSO dynamic orbits

Model	LPIPS↓	PSNR↑	SSIM↑	CLIP-S↑	MSE↓
Zero123 [23]	0.14	16.99	0.79	0.84	0.04
Zero123XL [8]	0.14	16.73	0.78	0.84	0.04
Stable Zero123 [46]	0.13	18.04	0.78	0.85	0.05
Free3D [60]	0.18	14.93	0.77	0.83	0.05
EscherNet [18]	0.13	16.47	0.79	0.84	0.03
SV3D <sup>c</sup>	0.10	19.99	0.86	<b>0.87</b>	<b>0.02</b>
SV3D <sup>p</sup>	<b>0.09</b>	<b>20.38</b>	<b>0.87</b>	<b>0.87</b>	<b>0.02</b>

#### 4. 3D Generation from a Single Image Using SV3D

We then generate 3D meshes of objects from a single image by leveraging SV3D. One way is to use the generated static/dynamic orbital samples from SV3D as direct reconstruction targets. Another way is to use SV3D as diffusion guidance with Score Distillation Sampling (SDS) loss [34].

Since SV3D produces more consistent multi-views compared to existing works, we already observe higher-quality 3D reconstructions by only using SV3D outputs for recon-

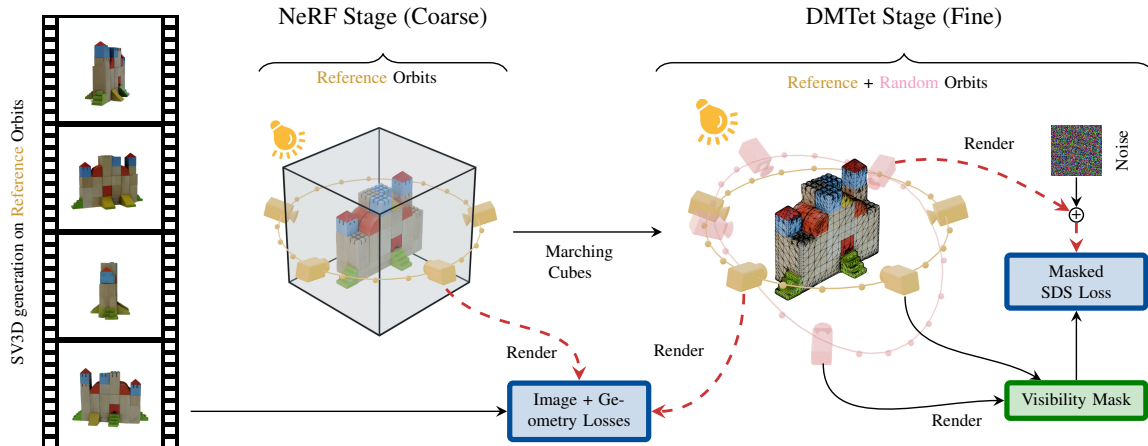


Figure 7. **3D Optimization Pipeline.** We perform a two-stage optimization. In the short first stage, we extract the general shape, texture and illumination from the SV3D generated multi-view images. In the second stage, we extract a mesh with marching cubes and use DMTet to further optimize the shape, texture and illumination. We not only use the SV3D-generated images but a soft-masked SDS loss for unseen areas. Dashed red lines represent backpropagation into the differentiable rendering pipeline.

Table 3. Evaluation of novel multi-view synthesis on OmniObject3D static orbits

Model	LPIPS↓	PSNR↑	SSIM↑	CLIP-S↑	MSE↓
Zero123 [23]	0.17	15.50	0.76	0.83	0.05
Zero123XL [8]	0.18	15.36	0.75	0.83	0.06
Stable Zero123 [46]	0.15	16.86	0.77	0.84	0.06
Free3D [60]	0.16	15.29	0.78	0.83	0.05
EscherNet [18]	0.17	14.63	0.74	0.83	0.05
SV3D <sup>u</sup>	<b>0.10</b>	19.68	<b>0.86</b>	<b>0.86</b>	<b>0.02</b>
SV3D <sup>c</sup>	<b>0.10</b>	19.50	0.85	0.85	<b>0.02</b>
SV3D <sup>p</sup>	<b>0.10</b>	<b>19.91</b>	<b>0.86</b>	<b>0.86</b>	<b>0.02</b>

Table 4. Evaluation of novel multi-view synthesis on OmniObject3D dynamic orbits

Model	LPIPS↓	PSNR↑	SSIM↑	CLIP-S↑	MSE↓
Zero123 [23]	0.16	15.78	0.77	0.82	0.05
Zero123XL [8]	0.17	15.49	0.76	0.83	0.05
Stable Zero123 [46]	0.14	16.74	0.77	0.83	0.05
Free3D [60]	0.19	14.28	0.76	0.82	0.06
EscherNet [18]	0.16	15.05	0.76	0.83	0.05
SV3D <sup>c</sup>	<b>0.10</b>	19.21	<b>0.85</b>	0.84	<b>0.02</b>
SV3D <sup>p</sup>	<b>0.10</b>	<b>19.28</b>	<b>0.85</b>	<b>0.85</b>	<b>0.02</b>

struction when compared to existing works. However, we observe that this naive approach often leads to artifacts like baked-in illumination, rough surfaces, and noisy texture, especially for the unseen regions in the reference orbit. Thus, we further propose several techniques to address these issues.

**Coarse-to-Fine Training.** We adopt a two-stage, coarse-to-fine training scheme to generate a 3D mesh from input images, similar to [20, 35]. Fig. 7 illustrate an overview

of our 3D optimization pipeline. In the coarse stage, we train an Instant-NGP [30] NeRF representation to reconstruct the SV3D-generated images (*i.e.* without SDS loss) at a lower resolution. In the fine stage, we extract a mesh from the trained NeRF using marching cubes [7], and adopt a DMTet [41] representation to finetune the 3D mesh using SDS-based diffusion guidance from SV3D at full-resolution. Finally, we use xatlas [56] to perform the UV unwrapping and export the output object mesh.



Figure 8. **Constant vs. SGs Illumination.** Notice that our SGs-based reconstructions do not exhibit darkening on the side of the bus, which enables easier and more convincing relighting for downstream applications.

**Disentangled Illumination Model.** Similar to other recent 3D object generation methods [20, 28, 34], our output target is a mesh with a diffuse texture. Typically, such SDS-based optimization techniques use random illuminations at every iteration. However, our SV3D-generated videos are under consistent illumination, *i.e.*, the lighting remains static while the camera circles around an object. Hence, to disentangle lighting effects and obtain a cleaner texture, we propose to fit a simple illumination model of 24 Spherical Gaussians (SG) inspired by prior decomposition methods [4, 58]. We model white light and hence only use a

scalar amplitude for the SGs. We only consider Lambertian shading, where the cosine shading term is approximated with another SG. We learn the parameters of the illumination SGs using a reconstruction loss between the rendered images and SV3D-generated images.




Figure 9. **Influence of Training Orbits.** We show that using a dynamic orbit is crucial to 3D generations that are complete from diverse views.




Figure 10. **Influence of SDS.** Using our masked SDS loss, we are able to fill in unseen surfaces in the training orbit, producing a cleaner result without oversaturation or blurry artifacts caused by naive SDS.

Inspired by [13, 34] we reduce baked-in illumination with a loss term that replicates the HSV-value component of the input image  $\mathbf{I}$  with the rendered illumination  $L$ :  $\mathcal{L}_{\text{illum}} = |V(\mathbf{I}) - L|^2$ , with  $V(\mathbf{c}) = \max(c_r, c_g, c_b)$ . Given these changes, our disentangled illumination model is able to express lighting variation properly and can severely reduce baked-in illumination. Fig. 8 shows sample reconstructions with and without our illumination modelling. From the re-

sults, it is clear that we are able to disentangle the illumination effects from the base color (e.g., dark side of the school bus).

#### 4.1. 3D Optimization Strategies and Losses

**Reconstruction via Photometric Losses.** Intuitively, we can treat the SV3D-generated images as multi-view pseudo-ground truth, and apply 2D reconstruction losses to train the 3D models. In this case, we apply photometric losses on the rendered images from NeRF or DMTet, including the pixel-level MSE loss, mask loss, and a perceptual LPIPS [59] loss. These photometric losses also optimize our illumination model through the differential rendering pipeline.

**Training Orbits.** For 3D generation, we use SV3D to generate multi-view images following a camera orbit  $\pi_{\text{ref}}$ , referred to as the *reference orbit* (also see Fig. 7 for the overview). Fig. 9 shows sample reconstruction with using static and dynamic orbital outputs from SV3D. Using a dynamic orbit for training is crucial to high-quality 3D outputs when viewed from various angles, since some top/bottom views are missing in the static orbit (fixed elevation). Hence, for SV3D<sup>c</sup> and SV3D<sup>p</sup>, we render images on a dynamic orbit whose elevation follows a sine function to ensure that top and bottom views are covered.

**SV3D-Based SDS Loss.** In addition to the reconstruction losses, we can also use SV3D via score-distillation sampling (SDS) [34, 54]. Fig. 10 shows sample reconstructions with and without using SDS losses. As shown in Fig. 10 (left), although training with a dynamic orbit improves overall visibility, we observe that sometimes the output texture is still noisy, perhaps due to partial visibility, self-occlusions, or inconsistent texture/shape between images. Hence, we handle those unseen areas using SDS loss [34] with SV3D as a diffusion guidance.

We sample a random camera orbit  $\pi_{\text{rand}}$ , and use our 3D NeRF/DMTet parameterized by  $\theta$  to render the views  $\hat{\mathbf{J}}$  of the object along  $\pi_{\text{rand}}$ . Then, noise  $\epsilon$  at level  $t$  is added to the latent embedding  $\mathbf{z}_t$  of  $\hat{\mathbf{J}}$ , and the following SDS loss (taken expectation over  $t$  and  $\epsilon$ ) is backpropagated through the differentiable rendering pipeline:  $\mathcal{L}_{\text{SDS}} = w(t)(\epsilon_{\phi}(\mathbf{z}_t; \mathbf{I}, \pi_{\text{rand}}, t) - \epsilon) \frac{\partial \hat{\mathbf{J}}}{\partial \theta}$ , where  $w$  is  $t$ -dependent weight,  $\epsilon$  and  $\epsilon_{\phi}$  are the added and predicted noise, and  $\phi$  and  $\theta$  are the parameters of SV3D and NeRF/DMTet, respectively. See Fig. 7 for an illustration of these loss functions in the overall pipeline.

**Masked SDS Loss.** As shown in Fig. 10 (middle), in our experiments we found that adding the SDS loss naively can cause unstable training and unfaithful texture to the input images such as oversaturation or blurry artifacts. Therefore, we design a soft masking mechanism to only apply SDS loss on the unseen/occluded areas, allowing it to inpaint the missing details while preserving the texture of clearly-



visible surfaces in the training orbit (as seen in Fig. 10 (right)). Also, we only apply the masked SDS loss in the final stage of DM Tet optimization, which greatly increased the convergence speed.

We apply SDS loss on only those pixels in the random orbit views that are not likely visible in the reference orbit views. For this, we first render the object from the random camera orbit  $\pi_{\text{rand}}$ . For each random camera view, we obtain the visible surface points  $\mathbf{p} \in \mathbb{R}^3$  and their corresponding surface normals  $\mathbf{n}$ . Then, for each reference camera  $i$ , we calculate the view directions  $\mathbf{v}_i$  of the surface  $\mathbf{p}$  towards its position  $\bar{\pi}_{\text{ref}}^i \in \mathbb{R}^3$  (calculated from  $\pi_{\text{ref}}^i \in \mathbb{R}^2$ ) as  $\mathbf{v}_i = \frac{\bar{\pi}_{\text{ref}}^i - \mathbf{p}}{\|\bar{\pi}_{\text{ref}}^i - \mathbf{p}\|}$ . We infer the visibility of this surface from the reference camera based on the dot product between  $\mathbf{v}_i$  and  $\mathbf{n}$  *i.e.*  $\mathbf{v}_i \cdot \mathbf{n}$ . Since higher values roughly indicate more visibility of the surface from the reference camera, we chose that reference camera  $c$  that has maximum likelihood of visibility:  $c = \max_i (\mathbf{v}_i \cdot \mathbf{n})$ . As we only want to apply SDS loss to unseen or grazing angle areas from  $c$ , we use the smoothstep function  $f_s$  to smoothly clip to  $c$ 's visibility range  $\mathbf{v}_c \cdot \mathbf{n}$ . In this way, we create a pseudo visibility mask  $M = 1 - f_s(\mathbf{v}_c \cdot \mathbf{n}, 0, 0.5)$ , where  $f_s(x; f_0, f_1) = \hat{x}^2(3 - 2\hat{x})$ , with  $\hat{x} = \frac{x - f_0}{f_1 - f_0}$ . Thus,  $M$  is calculated for each random camera render, and the combined visibility mask  $\mathbf{M}$  is applied to SDS loss:  $\mathcal{L}_{\text{mask-sds}} = \mathbf{M}\mathcal{L}_{\text{sds}}$ .

**Geometric Priors.** Since our rendering-based optimization operates at the image level, we adopt several geometric priors to regularize the output shapes. We add a smooth depth loss from RegNeRF [32] and a bilateral normal smoothness loss [5] to encourage smooth 3D surfaces where the projected image gradients are low. Moreover, we obtain normal estimates from Omnidata [11] and calculate a mono normal loss similar to MonoSDF [57], which can effectively reduce noisy surfaces in the output mesh. Further details about the training losses and optimization process are available in the appendix.

## 4.2. Experiments and Results

Due to the strong regularization, we only require 600 steps in the coarse stage and 1000 in the fine stage. Overall, the entire mesh extraction process takes  $\approx 8$  minutes without SDS loss, and  $\approx 20$  minutes with SDS loss. The coarse stage only takes  $\approx 2$  minutes and provides a full 3D representation of the object.

**Evaluation.** We evaluate our 3D generation framework on 50 randomly sampled objects from the GSO dataset as described in Sec. 3.1. We compute image-based reconstruction metrics (LPIPS, PSNR, SSIM, MSE, and CLIP-S) between the ground truth (GT) GSO images, and rendered images from the trained 3D meshes on the same static/dynamic orbits. In addition, we compute 3D reconstruction metrics of Chamfer distance (CD) and 3D IoU between

the GT and predicted meshes. We compare our SV3D-guided 3D generations with several prior methods including Point-E [31], Shap-E [16], One-2-3-45++ [21], DreamGaussian [47], SyncDreamer [24], EscherNet [18], Free3D [60], and Stable Zero123 [46].

**Visual Comparison.** In Fig. 11, we show visual comparison of our results with those from prior methods. Qualitatively, Point-E [31] and Shap-E [16] often produce incomplete 3D shapes. DreamGaussian [47], SyncDreamer [24], EscherNet [18], and Free3D [60] outputs tend to contain rough surfaces. One-2-3-45++ [21] and Stable Zero123 [46] are able to reconstruct meshes with smooth surface, but lack geometric details. Our mesh outputs are detailed, faithful to input image, and consistent in 3D (see appendix for more examples). We also show 3D mesh renders from real-world images in Fig. 12.

**Quantitative Comparison.** Tables 5 and 6 show the 2D and 3D metric comparisons respectively. All our 3D models achieve better 2D and 3D metrics compared to the prior and concurrent methods, showing the high-fidelity texture and geometry of our output meshes. We render all 3D meshes on the same dynamic orbits and compare them against the GT renders. Our best model, SV3D<sup>P</sup>, performs comparably to using GT renders for reconstruction in terms of the 3D metrics, which further demonstrates the 3D consistency of our generated images.

**Effects of Photometric and (Masked) SDS Loss.** As shown in Tables 5 and 6, the 3D outputs using both photometric and Masked SDS losses ('SV3D<sup>P</sup>') achieves the best results, while training without SDS ('SV3D<sup>P</sup> no SDS') leads to marginally lower performance. This demonstrates that the images generated by SV3D are high-quality reconstruction targets, and are often sufficient for 3D generation without the cumbersome SDS-based optimization. Nevertheless, adding SDS generally improves quality, as also shown in Fig. 10.

**Effects of SV3D Model and Training Orbit.** As shown in Tables 5 and 6, SV3D<sup>P</sup> achieves the best performance among the three SV3D models, indicating that its synthesized novel views are most faithful to the input image and consistent in 3D. On the other hand, SV3D<sup>u</sup> shares the same disadvantage as several prior works in that it can only generate views at the same elevation, which is insufficient to build a legible 3D object, as shown in Fig. 9. This also leads to the worse performance of 'SV3D<sup>P</sup> with static orbit' in Tables 5 and 6. Overall, SV3D<sup>P</sup> with dynamic orbit and masked SDS loss performs favorably against all other configurations since it can leverage more diverse views of the object.

**Limitations.** Our SV3D model is by design only capable of handling 2 degrees of freedom: elevation and azimuth; which is usually sufficient for 3D generation from a single image. One may want to tackle more degrees of

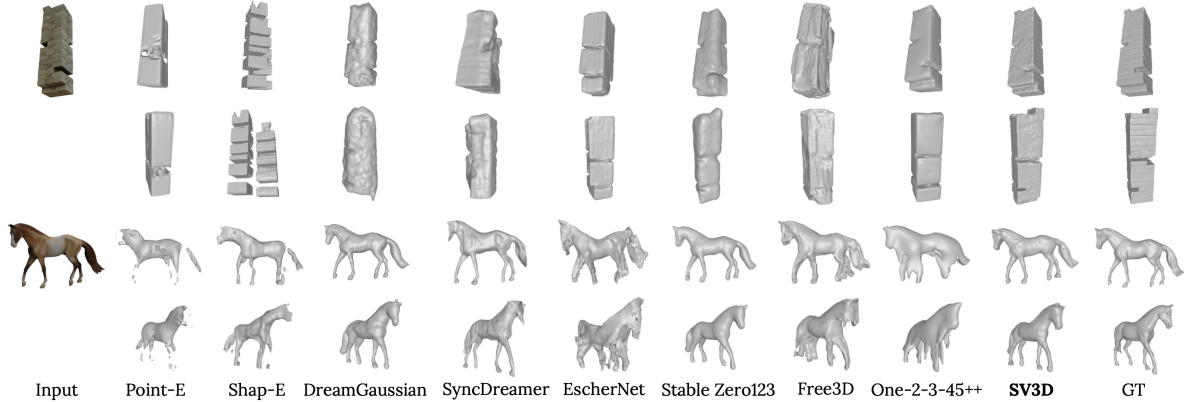


Figure 11. **Visual Comparison of 3D Meshes.** For each object, we show the conditioning image (left), and the output meshes rendered in two different views. Our generated meshes are more detailed, faithful to input images, and consistent in 3D. This demonstrates the quality of novel multi-view synthesis by our SV3D model.




Figure 12. **Real-World 3D Results.** Notice the accurate shape and details in our reconstructions even from the diverse images in-the-wild.

Table 5. 2D comparison of our 3D outputs against prior methods on the GSO dataset. Our best performing method uses **SV3D<sup>p</sup>** with dynamic (sine elevation) orbit and SDS guidance. Note that all our models achieve better 2D metrics than prior works.

Model	LPIPS↓	PSNR↑	SSIM↑	MSE↓	CLIP-S↑
GT renders	0.078	19.508	0.879	0.014	0.897
EscherNet [18]	0.178	14.438	0.804	0.041	0.835
Free3D [60]	0.197	14.202	0.799	0.043	0.809
Stable Zero123 [46]	0.166	14.635	0.813	0.040	0.805
SV3D <sup>u</sup>	0.133	15.957	0.834	0.031	0.871
SV3D <sup>c</sup>	0.132	16.373	0.834	0.027	0.870
SV3D <sup>p</sup> static orbit	0.125	16.821	0.848	0.025	0.864
SV3D <sup>p</sup> no SDS	0.124	16.864	0.841	0.024	0.875
<b>SV3D<sup>p</sup></b>	<b>0.119</b>	<b>17.405</b>	<b>0.849</b>	<b>0.021</b>	<b>0.877</b>

freedom in cameras for a generalized NVS system, which forms an interesting future work. We also notice that SV3D exhibits some view inconsistency for mirror-like reflective surfaces, which provide a challenge to 3D reconstruction. Lastly, such reflective surfaces are not representable by our Lambertian reflection-based shading model. Conditioning SV3D on the full camera matrix, and extending the shading model are interesting directions for future research.

Table 6. Comparison of 3D metrics. Our models perform favorably against prior works.

Model	CD↓	3D IoU↑
GT renders	0.021	0.689
Point-E [31]	0.074	0.162
Shap-E [16]	0.071	0.267
DreamGaussian [47]	0.055	0.411
One-2-3-45++ [21]	0.054	0.406
SyncDreamer [24]	0.053	0.451
EscherNet [18]	0.042	0.466
Free3D [60]	0.047	0.426
Stable Zero123 [46]	0.039	0.550
SV3D <sup>u</sup>	0.027	0.589
SV3D <sup>c</sup>	0.027	0.584
SV3D <sup>p</sup> static orbit	0.028	0.610
SV3D <sup>p</sup> no SDS	<b>0.024</b>	0.611
<b>SV3D<sup>p</sup></b>	<b>0.024</b>	<b>0.614</b>

## 5. Conclusion

We present SV3D, a latent video diffusion model for novel multi-view synthesis and 3D generation. In addition to leveraging the generalizability and view-consistent prior in SVD, SV3D enables controllability via camera pose con-

ditioning, and generates orbital videos of an object at high resolution on arbitrary camera orbits. We further propose several techniques for improved 3D generation from SV3D, including triangle CFG scaling, disentangled illumination, and masked SDS loss. We conduct extensive experiments to show that SV3D is controllable, multi-view consistent, and generalizable to the real-world, achieving state-of-the-art performance on novel multi-view synthesis and 3D generation. We believe SV3D provides a solid foundation model for further research on 3D object generation. We plan to publicly release SV3D models.

**Acknowledgments.** We thank Emad Mostaque, Bryce Wilson, Anel Islamovic, Savannah Martin, Ana Guillen, Josephine Parquet, Adam Chen and Ella Irwin for their help in various aspects of the model development and release. We thank Kyle Lacey and Yan Marek for their help with demos. We also thank Eric Courtemanche for his help with visual results.

## References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-NeRF: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 3
- [2] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable Video Diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023. 2, 3, 4, 15
- [3] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your Latents: High-Resolution Video Synthesis with Latent Diffusion Models. *arXiv:2304.08818*, 2023. 2
- [4] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P.A. Lensch. NeRD: Neural reflectance decomposition from image collections. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 7, 16
- [5] Mark Boss, Andreas Engelhardt, Abhishek Kar, Yuanzhen Li, Deqing Sun, Jonathan T. Barron, Hendrik P.A. Lensch, and Varun Jampani. SAMURAI: Shape And Material from Unconstrained Real-world Arbitrary Image collections. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 9, 16
- [6] 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. GeNVS: Generative novel view synthesis with 3D-aware diffusion models. *International Conference on Computer Vision*, 2023. 2
- [7] Daren BH Cline. Admissible kernel estimators of a multivariate density. *The Annals of Statistics*, 16(4):1421–1427, 1988. 7
- [8] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-XL: A universe of 10m+ 3D objects. *arXiv preprint arXiv:2307.05663*, 2023. 2, 5, 6, 7
- [9] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3D objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023. 5, 15
- [10] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. Google Scanned Objects: A high-quality dataset of 3D scanned household items. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2553–2560. IEEE, 2022. 5
- [11] Ainaz Eftekhari, Alexander Sax, Roman Bachmann, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3D scans. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 9, 16
- [12] Rohit Girdhar, Mannat Singh, Andrew Brown, Quentin Duval, Samaneh Azadi, Sai Saketh Rambhatla, Akbar Shah, Xi Yin, Devi Parikh, and Ishan Misra. EMU VIDEO: Factorizing Text-to-Video Generation by Explicit Image Conditioning, 2023. 2
- [13] Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. Shape, Light, and Material Decomposition from Images using Monte Carlo Rendering and Denoising. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 8
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020. 2
- [15] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. LRM: Large reconstruction model for single image to 3D. *International Conference on Learning Representations*, 2024. 2
- [16] Heewoo Jun and Alex Nichol. Shap-E: Generating conditional 3D implicit functions, 2023. 3, 9, 10
- [17] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the Design Space of Diffusion-Based Generative Models. *arXiv:2206.00364*, 2022. 15
- [18] Xin Kong, Shikun Liu, Xiaoyang Lyu, Marwan Taher, Xiaojuan Qi, and Andrew J Davison. Eschnet: A generative model for scalable view synthesis. *arXiv preprint arXiv:2402.03908*, 2024. 2, 3, 5, 6, 7, 9, 10
- [19] Jeong-gi Kwak, Erqun Dong, Yuhe Jin, Hanseok Ko, Shweta Mahajan, and Kwang Moo Yi. Vivid-1-to-3: Novel view synthesis with video diffusion models. *IEEE conference on computer vision and pattern recognition*, 2024. 3
- [20] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3D: High-resolution

- text-to-3D content creation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 300–309, 2023. 1, 3, 7
- [21] Minghua Liu, Ruoxi Shi, Linghao Chen, Zhuoyang Zhang, Chao Xu, Xinyue Wei, Hansheng Chen, Chong Zeng, Jiayuan Gu, and Hao Su. One-2-3-45++: Fast single image to 3D objects with consistent multi-view generation and 3D diffusion. *arXiv preprint arXiv:2311.07885*, 2023. 2, 3, 9, 10
- [22] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. One-2-3-45: Any single image to 3D mesh in 45 seconds without per-shape optimization. *Advances in Neural Information Processing Systems*, 36, 2023. 2, 3
- [23] 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. *International Conference on Computer Vision*, 2023. 1, 2, 3, 5, 6, 7
- [24] 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. 1, 2, 3, 5, 6, 9, 10
- [25] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3D: Single image to 3D using cross-domain diffusion. *arXiv preprint arXiv:2310.15008*, 2023. 3
- [26] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, Natalia Neverova, Andrea Vedaldi, Oran Gafni, and Filippos Kokkinos. IM-3D: Iterative multiview diffusion and reconstruction for high-quality 3D generation. *arXiv preprint arXiv:2402.08682*, 2024. 2, 3
- [27] Antoine Mercier, Ramin Nakhli, Mahesh Reddy, Rajeev Yasarla, Hong Cai, Fatih Porikli, and Guillaume Berger. HexaGen3D: Stablediffusion is just one step away from fast and diverse Text-to-3D generation. *arXiv preprint arXiv:2401.07727*, 2024. 1, 2
- [28] Gal Metzer, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-NeRF for shape-guided generation of 3D shapes and textures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12663–12673, 2023. 1, 3, 7
- [29] 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 *European Conference on Computer Vision*, pages 405–421, 2020. 3
- [30] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 2022. 3, 7
- [31] Alex Nichol, Heewoo Jun, Pratul Dhariwal, Pamela Mishkin, and Mark Chen. Point-E: A System for Generating 3D Point Clouds from Complex Prompts, 2022. 3, 9, 10
- [32] Michael Niemeyer, Jonathan T. Barron, Ben Mildenhall, Mehdi S. M. Sajjadi, Andreas Geiger, and Noha Radwan. RegNeRF: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 9, 16
- [33] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 2
- [34] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3D using 2D diffusion. *arXiv*, 2022. 1, 2, 3, 6, 7, 8
- [35] 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. 3, 7
- [36] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 4
- [37] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1, 2
- [38] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. DreamBooth: Fine tuning text-to-image diffusion models for subject-driven generation. *arXiv preprint arXiv:2208.12242*, 2022. 2
- [39] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. *arXiv:2205.11487*, 2022. 1
- [40] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. LAION-5B: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 2
- [41] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep Marching Tetrahedra: a hybrid representation for high-resolution 3D shape synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 3, 7
- [42] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. Zero123++: a single image to consistent multi-view diffusion base model. *arXiv preprint arXiv:2310.15110*, 2023. 2, 3
- [43] 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

- [44] Yang Song and Stefano Ermon. Improved Techniques for Training Score-Based Generative Models. *arXiv:2006.09011*, 2020. 16
- [45] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 2
- [46] StabilityAI. Stable Zero123, 2023. 2, 3, 5, 6, 7, 9, 10
- [47] Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. DreamGaussian: Generative Gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023. 9, 10
- [48] Yu-Ting Tsai and Zen-Chung Shih. All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *ACM Transactions on Graphics (ToG)*, 2006. 16
- [49] Vikram Voleti, Alexia Jolicoeur-Martineau, and Christopher Pal. MCVD: Masked conditional video diffusion for prediction, generation, and interpolation. In *(NeurIPS) Advances in Neural Information Processing Systems*, 2022. 2
- [50] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel view synthesis with diffusion models, 2022. 1
- [51] Haohan Weng, Tianyu Yang, Jianan Wang, Yu Li, Tong Zhang, C. L. Philip Chen, and Lei Zhang. Consistent123: Improve consistency for one image to 3D object synthesis, 2023. 3
- [52] Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Jiawei Ren, Liang Pan, Wayne Wu, Lei Yang, Jiaqi Wang, Chen Qian, et al. Omniobject3D: Large-vocabulary 3D object dataset for realistic perception, reconstruction and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 803–814, 2023. 5, 17, 18, 19
- [53] Jiayu Yang, Ziang Cheng, Yunfei Duan, Pan Ji, and Hongdong Li. ConsistNet: Enforcing 3D consistency for multi-view images diffusion. *arXiv preprint arXiv:2310.10343*, 2023. 3
- [54] Chun-Han Yao, Amit Raj, Wei-Chih Hung, Michael Rubinstein, Yuanzhen Li, Ming-Hsuan Yang, and Varun Jampani. ARTIC3D: Learning robust articulated 3D shapes from noisy web image collections. *Advances in Neural Information Processing Systems*, 36, 2024. 8
- [55] Jianglong Ye, Peng Wang, Kejie Li, Yichun Shi, and Heng Wang. Consistent-1-to-3: Consistent image to 3D view synthesis via geometry-aware diffusion models. In *3DV*, 2024. 2
- [56] Jonathan Young. xatlas. <https://github.com/jpcy/xatlas>, 2024. 7
- [57] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. MonoSDF: Exploring monocular geometric cues for neural implicit surface reconstruction. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 9
- [58] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. PhysSG: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 7, 16
- [59] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 5, 8, 16
- [60] Chuanxia Zheng and Andrea Vedaldi. Free3D: Consistent novel view synthesis without 3D representation. *arXiv preprint arXiv:2312.04551*, 2023. 2, 3, 5, 6, 7, 9, 10, 17

# Contents

<b>1. Introduction</b>	<b>1</b>
<b>2. Background</b>	<b>2</b>
2.1. Novel View Synthesis	2
2.2. 3D Generation	3
<b>3. SV3D: Novel Multi-view Synthesis</b>	<b>3</b>
3.1. Experiments and Results	5
<b>4. 3D Generation from a Single Image Using SV3D</b>	<b>6</b>
4.1. 3D Optimization Strategies and Losses	8
4.2. Experiments and Results	9
<b>5. Conclusion</b>	<b>10</b>
<b>A Broader Impact</b>	<b>14</b>
<b>B Data Details</b>	<b>15</b>
<b>C Training Details</b>	<b>15</b>
<b>D Inference Details</b>	<b>16</b>
<b>E Additional Details on Illumination Model</b>	<b>16</b>
<b>F. Losses and Optimization for 3D Generation</b>	<b>16</b>
<b>G Additional Ablative Analyses</b>	<b>16</b>
G.1. SV3D Models	16
G.2. Static v.s. Dynamic Orbits	16
G.3. Masked SDS Loss	17
<b>H Additional Visual Results</b>	<b>17</b>
H.1. Novel View Synthesis	17
H.2. 3D Generation	17

## Appendix

### A. Broader Impact

The advancement of generative models in different media forms is changing how we make and use content. These AI-powered models can create images, videos, 3D objects, and more, in ways we’ve never seen before. They offer huge potential for innovation in media production. But along with this potential, there are also risks. Before we start using these models widely, it is crucial to make sure we understand the possible downsides and have plans in place to deal with them effectively.

In the case of 3D object generation, the input provided by the user plays a crucial role in constraining the model’s output. By supplying a full front view of an object, users limit the model’s creative freedom to the visible or unoccluded regions, thus minimizing the potential for generating problematic imagery. Additionally, factors such as predicted depth values and lighting further influence the fidelity and realism of generated content.

Moreover, ensuring the integrity and appropriateness of training data is critical in mitigating risks associated with generative models. Platforms like Sketchfab, which serve as repositories for 3D models used in training, enforce strict content policies to prevent the dissemination of “Unacceptable Content” and disallows it on their platform: <https://help.sketchfab.com/hc/en-us/articles/214867883-What-is-Restricted-Content>. By adhering to these guidelines and actively monitoring dataset quality, developers can reduce the likelihood of biased or inappropriate outputs.

Sketchfab also has a tag for “Restricted Content” which is deemed to be similar to the PG-13 content rating used in the US (i.e. inappropriate for children under 13). We have confirmed that none of the objects that we use in training have this flag set to true. Thus we go the extra step of excluding even tagged PG-13 content from the training set.

There is a chance that certain content may not have been correctly labeled on Sketchfab. In cases where the uploader fails to tag an object appropriately, Sketchfab provides publicly accessible listings of objects and a mechanism for the community to report any content that may be deemed offensive.

In our regular utilization of Objaverse content, we haven’t observed any significant amount of questionable material. Nevertheless, there are occasional instances of doll-like nudity stemming from basic 3D models, which could be crucial for accurately depicting humanoid anatomy. Additionally, the training dataset contains some presence of drugs, drug paraphernalia, as well as a certain level of blood content and weaponry, resembling what might be encountered in a video game context. Should the model be provided with imagery featuring these categories of content, it possesses the capability to generate corresponding 3D models to some extent.

It is to be noted that SV3D mainly focuses on generating hidden details in the user’s input image. If the image is unclear or some parts are hidden, the model guesses what those parts might look like based on its training data. This means it might create details similar to what it has seen before. The training data generally follows the standards of 3D modeling and gaming. However, this could lead to criticisms about the models being too similar to existing trends. But the user’s input image limits how creative the model can be and reduces the chance of biases showing up in its creations, especially if the image is clear and straightforward.

## B. Data Details

Similar to SVD-MV [2], we render views of a subset of the Objaverse dataset [9] consisting of 150K curated CC-licensed 3D objects from the original dataset. Each loaded object is scaled such that the largest world-space XYZ extent of its bounding box is 1. The object is then repositioned such that it this bounding box is centered around the origin.

For both the static and dynamic orbits, we use Blender’s EEVEE renderer to render an 84-frame RGBA orbit at  $576 \times 576$  resolution. During training, any 21-frame orbit can be subsampled from this by picking any frame as the first frame, and then choosing every 4th frame after that.

We apply two background colors to each of these images: random RGB color, and white. This results in a doubling of the number of orbit samples for training. We then encode all of these images into latent space using SVD’s VAE, and using CLIP. We then store the latent and CLIP embeddings for all of these images along with the corresponding elevation and azimuth values.

For lighting, we randomly select from a set of 20 curated HDRI envmaps. Each orbit begins with the camera positioned at azimuth 0. Our camera uses a field-of-view of 33.8 degrees. For each object, we adaptively position the camera to a distance sufficient to ensure that the rendered object content makes good and consistent use of the image extents without being clipped in any view.

For static orbits, the camera is positioned at a randomly sampled elevation between  $[-5, 30]$  degrees. The azimuth steps by a constant  $\frac{360}{84}$  degree delta between each frame. For dynamic orbits, the sequence of camera elevations for each orbit are obtained from a random weighted combination of sinusoids with different frequencies. For each sinusoid, the whole number period is sampled from  $[1, 5]$ , the amplitude is sampled from  $[0.5, 10]$ , and a random phase shift is applied. The azimuth angles are sampled regularly, and then a small amount of noise is added to make them irregular. The elevation values are smoothed using a simple convolution kernel and then clamped to a maximum elevation of 89 degrees.

## C. Training Details

Our approach involves utilizing the widely used EDM [17] framework, incorporating a simplified diffusion loss for finetuning, as followed in SVD [2]. We eliminate the conditions of ‘fps\_id’, ‘motion\_bucket\_id’, etc. since they are irrelevant for SV3D. Furthermore, we adjust the loss function to assign lower weights to frames closer to the front-view conditioning image, ensuring that challenging back views receive equal training focus as the easier front views. To optimize training efficiency and conserve GPU VRAM, we preprocess and store the precomputed latent and CLIP embeddings for all video frames in advance. During training, these tensors are directly loaded rather than being computed in real-time. We choose to finetune the SVD-xt model to output 21 frames instead of 25 frames. We found that with 21 frames we were able to fit a batch size of 2 on each GPU, instead of 1 with 25 frames at  $576 \times 576$  resolution. All SV3D models are trained for 105k iterations with an effective batch size of 64 on 4 nodes of 8 80GB A100 GPUs for around 6 days.

## D. Inference Details

To generate an orbital video during inference, we use 50 steps of the deterministic DDIM sampler [44] with the triangular classifier-free guidance (CFG) scale described in the main paper. This takes  $\approx 40$  seconds for the SV3D model.

## E. Additional Details on Illumination Model

We base our rendering model on Spherical Gaussians (SG) [4, 58]. A SG at query location  $\mathbf{x} \in \mathbb{R}^3$  is defined as  $G(\mathbf{x}; \boldsymbol{\mu}, c, a) = ae^{s(\boldsymbol{\mu} \cdot \mathbf{x} - 1)}$ , where  $\boldsymbol{\mu} \in \mathbb{R}^3$  is the axis,  $s \in \mathbb{R}$  the sharpness of the lobe, and  $a \in \mathbb{R}$  the amplitude. Here, we point out that we only model white light and hence only use a scalar amplitude. One particularly interesting property of SGs is that the inner product of two SGs is the integral of the product of two SGs. The operation is defined as [48]:

$$G_1(\mathbf{x}) \cdot G_2(\mathbf{x}) = \int_{\Omega} G_1(\mathbf{x})G_2(\mathbf{x})d\mathbf{x} = \frac{1}{d_m} (2\pi a_1 a_2 e^{d_m - \lambda_m} (1.0 - e^{-2d_m}))$$

$$\lambda_m = \lambda_1 - \lambda_2$$

$$d_m = \|\lambda_1 \boldsymbol{\mu}_1 + \lambda_2 \boldsymbol{\mu}_2\|. \tag{1}$$

In our illumination model we only consider Lambertian shading. Here, the cosine shading term influences the output the most. This term can be approximated with another SG  $G_{\text{cosine}} = (\mathbf{x}; \mathbf{n}, 2.133, 1.17)$ , where  $\mathbf{n}$  defines the surface normal at  $\mathbf{x}$ . The lighting evaluation using SGs  $G_i$  then becomes:  $L = \sum_{i=1}^{24} \frac{1}{\pi} \max(G_i \cdot G_{\text{cosine}}, 0)$ . As defined previously this results in the full integration of incoming light for each SG and as light is additive evaluating and summing all SGs results in the complete environment illumination. This  $L$  is also used in the  $\mathcal{L}_{\text{illum}}$  loss described in the main paper. The rendered textured image is then defined as  $\hat{\mathbf{I}} = \mathbf{c}_d \mathbf{L}$ , where  $\mathbf{c}_d$  is the diffuse albedo. We learn  $\boldsymbol{\mu}, c, a$  for each SG  $G_i$  using reconstruction loss between these rendered images and SV3D-generated images.

## F. Losses and Optimization for 3D Generation

In addition to the masked SDS loss  $\mathcal{L}_{\text{mask-sds}}$  and illumination loss  $\mathcal{L}_{\text{illum}}$  detailed in the manuscript, we use several other losses for 3D reconstruction. Our main reconstruction losses are the pixel-level mean squared error  $\mathcal{L}_{\text{mse}} = \|\mathbf{I} - \hat{\mathbf{I}}\|^2$ , LPIPS [59] loss  $\mathcal{L}_{\text{lpips}}$ , and mask loss  $\mathcal{L}_{\text{mask}} = \|\mathbf{S} - \hat{\mathbf{S}}\|^2$ , where  $\mathbf{S}, \hat{\mathbf{S}}$  are the predicted and ground-truth masks. We further employ a normal loss using the estimated mono normal by Omnidata [11], which is defined as the cosine similarity between the rendered normal  $\mathbf{n}$  and estimated pseudo ground truths  $\bar{\mathbf{n}}$ :  $\mathcal{L}_{\text{normal}} = 1 - (\mathbf{n} \cdot \bar{\mathbf{n}})$ . To regularize the output geometry, we apply a smooth depth loss inspired by RegNeRF [32]:  $\mathcal{L}_{\text{depth}}(i, j) = (d(i, j) - d(i, j + 1))^2 + (d(i, j) - (d(i + 1, j)))^2$ , where  $i, j$  indicate the pixel coordinate. For surface normal we instead rely on a bilateral smoothness loss similar to [5]. We found that this is crucial to getting high-frequency details and avoiding over-smoothed surfaces. For this loss we compute the image gradients of the input image  $\nabla \mathbf{I}$  with a Sobel filter [?]. We then encourage the gradients of rendered normal  $\nabla \mathbf{n}$  to be smooth if (and only if) the input image gradients  $\nabla \mathbf{I}$  are smooth. The loss can be written as  $\mathcal{L}_{\text{bilateral}} = e^{-3\nabla \mathbf{I}} \sqrt{1 + \|\nabla \mathbf{n}\|}$ . We also found that adding a spatial smoothness regularization on the albedo is beneficial:  $\mathcal{L}_{\text{albedo}} = (\mathbf{c}_d(\mathbf{x}) - \mathbf{c}_d(\mathbf{x} + \boldsymbol{\epsilon}))^2$ , where  $\mathbf{c}_d$  denotes the albedo,  $\mathbf{x}$  is a 3D surface point, and  $\boldsymbol{\epsilon} \in \mathbb{R}^3$  is a normal distributed offset with a scale of 0.01. The overall objective is then defined as the weighted sum of these losses. All losses are applied in both coarse and fine stages, except that we only apply  $\mathcal{L}_{\text{mask-sds}}$  in the last 200 iterations of the fine stage. We use an Adam optimizer [?] with a learning rate of 0.01 for both stages.

## G. Additional Ablative Analyses

We conduct additional ablative analyses of our 3D generation pipeline in this section.

### G.1. SV3D Models

In Table 7, we compare the quantitative results using different SV3D models and training losses. Both 2D and 3D evaluation shows that SV3D<sup>P</sup> is our best performing model, either for pure photometric reconstruction or SDS-based optimization.

### G.2. Static v.s. Dynamic Orbits

We also compare the results using different camera orbits for 3D training in Table 8. The results show that using a dynamic orbit (sine-30) produces better 3D outputs compared to static orbit since it contains more information of the top and bottom



Table 7. **Ablative results of different SV3D models and training losses.** We show that our SV3D<sup>p</sup> model with Photo+SDS losses achieves the best 2D and 3D metrics.

Model	Training losses	LPIPS↓	PSNR↑	SSIM↑	MSE↓	CLIP-S↑	CD↓	3D IoU↑
SV3D <sup>u</sup>	Photo	0.132	15.951	0.827	0.032	0.873	0.028	0.583
SV3D <sup>u</sup>	Photo+SDS	0.133	15.957	0.834	0.031	0.871	0.027	0.589
SV3D <sup>c</sup>	Photo	0.135	15.826	0.832	0.033	0.871	0.029	0.579
SV3D <sup>c</sup>	Photo+SDS	0.132	16.373	0.834	0.027	0.870	0.027	0.584
SV3D <sup>p</sup>	Photo	0.124	16.864	0.841	0.024	0.875	<b>0.024</b>	0.611
SV3D <sup>p</sup>	Photo+SDS	<b>0.119</b>	<b>17.405</b>	<b>0.849</b>	<b>0.021</b>	<b>0.877</b>	<b>0.024</b>	<b>0.614</b>

Table 8. **Ablative results of different reference orbits for 3D generation.** We show that using a dynamic orbit (sine elevation) with moderate amplitude performs better than orbits with no or extreme elevation variations.

Training orbit	LPIPS↓	PSNR↑	SSIM↑	MSE↓	CLIP-S↑	CD↓	3D IoU↑
Static	0.125	16.821	0.848	0.025	0.864	0.028	0.610
Sine-30	<b>0.119</b>	<b>17.405</b>	0.849	<b>0.021</b>	<b>0.877</b>	<b>0.024</b>	<b>0.614</b>
Sine-50	0.123	17.057	<b>0.854</b>	0.025	0.873	0.026	0.609

Table 9. **Ablative analyses of Masked SDS loss.** Overall, our soft-masked SDS loss leads to higher-quality mesh outputs in terms of most 2D and 3D metrics.

Training losses	LPIPS↓	PSNR↑	SSIM↑	MSE↓	CLIP-S↑	CD↓	3D IoU↑
Photo	0.124	16.864	0.841	0.024	0.875	<b>0.024</b>	0.611
Photo+SDS (naive)	0.124	17.007	<b>0.850</b>	0.024	0.867	0.025	<b>0.615</b>
Photo+SDS (hard masked)	0.124	17.335	0.845	0.022	<b>0.877</b>	<b>0.024</b>	0.610
Photo+SDS (soft masked)	<b>0.119</b>	<b>17.405</b>	0.849	<b>0.021</b>	<b>0.877</b>	<b>0.024</b>	0.614

views of the object. However, higher elevation (sine-50) tends to increase inconsistency between multi-view images, and thus resulting in worse 3D reconstruction. In our experiments, we find that setting the elevation within  $\pm 30$  degree generally leads to desirable 3D outputs.

### G.3. Masked SDS Loss

Finally, we show the ablative results of our SDS loss in Table 9. We compare the results of 1) pure photometric losses, 2) with naive SDS loss (no masking), 3) with hard-masked SDS loss by thresholding the dot product of surface normal and camera viewing angle as visibility masks, and 4) with soft-masked SDS loss as described in the manuscript. Overall, adding SDS guidance from the SV3D model can improve the 2D metrics while maintaining similar 3D metrics. Our novel soft-masked SDS loss generally achieves the best results compared to other baselines.

## H. Additional Visual Results

In this section, we show more results of novel view synthesis and 3D generation.

### H.1. Novel View Synthesis

We show the additional NVS results on OmniObject3D [52] and real-world images in Fig. 13 and Fig. 14, respectively. The generated novel multi-view images by SV3D are more detailed and consistent compared to prior state-of-the-arts.

### H.2. 3D Generation

We show the additional 3D generation results on OmniObject3D [52] and real-world images in Fig. 15 and Fig. 16, respectively. Thanks to the consistent multi-view images by SV3D and the novel Masked SDS loss, our 3D generations are detailed, high-fidelity, and generalizable to a wide range of images. Since Free3D [60] does not include a 3D generation method, we run our 3D pipeline on its generated multi-view images for fair 3D comparison.



Figure 13. NVS Results on OmniObject3D [52]. Notice the consistent color, geometry, and pose in SV3D NVS outputs compared to prior works.



Figure 14. NVS Results on Real-World images. Notice the consistent color, geometry, and pose in SV3D NVS outputs compared to prior works.




Figure 15. Mesh Results on OmniObject3D [52]. Notice the accurate shape and details in our reconstructions even from the diverse images.




Figure 16. **Mesh Results on Real-World images.** Notice the accurate shape and details in our reconstructions even from the diverse images in-the-wild.