

---

# SV4D: Dynamic 3D Content Generation with Multi-Frame and Multi-View Consistency

---

Yiming Xie<sup>1,2\*</sup> Chun-Han Yao<sup>1\*</sup> Vikram Voleti<sup>1</sup> Huaizu Jiang<sup>2†</sup> Varun Jampani<sup>1†</sup>

<sup>1</sup> Stability AI <sup>2</sup> Northeastern University

\* Equal contribution † Equal advising

## Abstract

We present Stable Video 4D (SV4D) — a latent video diffusion model for multi-frame and multi-view consistent dynamic 3D content generation. Unlike previous methods that rely on separately trained generative models for video generation and novel view synthesis, we design a unified diffusion model to generate novel view videos of dynamic 3D objects. Specifically, given a monocular reference video, SV4D generates novel views for each video frame that are temporally consistent. We then use the generated novel view videos to optimize an implicit 4D representation (dynamic NeRF) efficiently, without the need for cumbersome SDS-based optimization used in most prior works. To train our unified novel-view video generation model, we curated a dynamic 3D object dataset from the existing Objaverse dataset. Extensive experimental results on multiple datasets and user studies demonstrate SV4D’s state-of-the-art performance on novel-view video synthesis as well as 4D generation compared to prior works. Project page: <https://sv4d.github.io>.

## 1 Introduction

The 3D world we live in is dynamic in nature with moving people, playing pets, bouncing balls, waving flags, etc. Dynamic 3D object generation, also known as 4D generation, is the task of generating not just the 3D shape and appearance (texture) of a 3D object, but also its motion in 3D space. In this work, we tackle the problem of generating a 4D (dynamic 3D) object from a single monocular video of that object. 4D generation enables the effortless creation of realistic visual experiences, such as for video games, movies, AR/VR, etc.

4D generation from a single video is highly challenging as this involves simultaneously reasoning both the object appearance and motion at unseen camera views around the object. This is also an ill-posed problem as a multitude of 4D results can plausibly explain a single given video. There are two main technical challenges in training a robust 4D generative model that can generalize to different object types and motions. First, there exists no large-scale dataset with 4D objects to train a robust generative model. Second, the higher dimensional nature of the problem requires a large number of parameters to represent the 3D shape, appearance, and motion of an object.

As a result, several recent techniques [48, 1, 27, 42, 68, 18, 66, 64] optimize 4D content by leveraging priors in pre-trained video and multi-view generative models via score-distillation sampling (SDS) loss [40] and its variants. However, they tend to produce unsatisfactory results due to independent modeling of object motion using video models, and novel view synthesis using multi-view generative models. In addition, they tend to take hours to generate a single 4D object due to time-consuming SDS-based optimization. A concurrent work [61] tries to partially address these issues by jointly sampling novel view videos (along both video frame and view axes) using both video and multi-view generative models; and then using the resulting novel view videos for 4D optimization. This results in a better novel view video synthesis with multi-view dynamic consistency, but several inconsistencies still remain due to the use of separate video and multi-view generative models.




Figure 1: **Stable Video 4D (SV4D)** framework overview and generated 4D assets. We adapt and train a video diffusion model to generate novel view videos, conditioned on a single-view video and a multi-view orbital video of the first frame. SV4D generated novel view videos are consistent across both view and motion axis, which we directly use to optimize dynamic 3D objects without the cumbersome SDS loss.

In this work, we propose Stable Video 4D (SV4D) model that takes as input a single video, such as of a dynamic object, along with a user-specified camera trajectory around the object, and outputs videos of the object along each of the specified camera views. That is, given a video with  $F$  number of frames and a camera trajectory with  $V$  number of camera views, SV4D outputs a  $V \times F$  grid of images, as illustrated in Fig. 1. Given a reference video, we first obtain the reference multi-views for the first frame using an off-the-shelf multi-view generator (SV3D [54]). SV4D then jointly outputs the remaining grid of images as highlighted by a yellow box in Fig. 1. In contrast to prior and concurrent works, SV4D jointly reasons along both view and motion axes, resulting in state-of-the-art multi-frame and multi-view consistency in the output novel view videos.

Specifically, we start with Stable Video Diffusion (SVD) [2], a state-of-the-art video generator, and equip it with two more attention blocks: view attention and frame attention. The view attention block aligns the multi-view images at each video frame, conditioning on the first view *i.e.*, in the reference video. Similarly, the frame attention block aligns the multi-frame images at each view, conditioning on the first frame at each view, *i.e.*, reference multi-view. This dual-attention design leads to significantly improved dynamic and multi-view consistency compared to prior art.

A key challenge here is that SV4D needs to simultaneously generate the  $V \times F$  grid of images, which can quickly become large with long input videos; making it infeasible to fit into memory even on modern GPUs. As a remedy, we propose techniques to sequentially process an interleaved subset of input frames while also retaining consistency in the output image grid. After generating the multiple novel view videos, we optimize a 4D representation of the dynamic 3D asset, as illustrated in Fig. 1.

Given the lack of large-scale 4D datasets, we carefully initialized SV4D weights with those from SVD [2] and SV3D [54] networks, thereby leveraging the priors learned in the existing video and multi-view diffusion models. To further train SV4D, we carefully curated a subject of Objaverse [9, 8] dataset with dynamic 3D objects, resulting in the ObjaverseDy dataset.

We perform extensive comparisons of both novel view video synthesis and 4D generation results with respective state-of-the-art methods on datasets with synthetic (ObjaverseDy, Consistent4D [18]), and real-world (DAVIS [38, 39, 4]) data. We modify the FVD metric [52] to evaluate both video frame and view consistency, validating the effectiveness and design choices of our approach. Fig. 1 shows some sample results of our approach.

To summarize, our contributions include:

- A novel SV4D network that can simultaneously reason across both frame and view axes. To our knowledge, this is the first work that trains a single novel view video synthesis network using 4D datasets, that can jointly perform novel view synthesis as well as video generation.
- A mixed sampling scheme that enables sequential processing of arbitrary long input videos while also retaining the multi-frame and multi-view consistency.
- State-of-the-art results on multiple benchmark datasets with both novel view video synthesis as well as 4D generation.

## 2 Related Work

**3D Generation.** Here, we refer to the works that generate static 3D content as 3D generation. DreamFusion [40] first proposed to distill priors from the 2D diffusion model via SDS loss to optimize the 3D content from text or image. Several subsequent works [63, 50, 45, 57, 24, 59, 36, 6, 49, 43, 26, 69, 11] try to solve issues caused by the original SDS loss, such as multi-face Janus, slow generation speed, and over-saturated/smoothed generations. Recent works [16, 17, 56, 70, 58, 51] try to directly predict the 3D model of an object via a large reconstruction model. Another approach [30, 31, 32, 54, 62, 19, 23, 46, 44, 55, 29, 28] to 3D generation is generating dense multi-view images with sufficient 3D consistency. 3D content is reconstructed based on the dense multi-view images. We follow this strategy, but generate consistent multi-view videos (instead of images) and then reconstruct the 4D object.

**Video Generation.** Recent video generation models [15, 53, 3, 2, 14, 47, 12] have shown very impressive performance with consistent geometry and realistic motions. Video generation models have good generalization capabilities, as they are trained on large-scale image and video data that are easier to collect than large-scale 3D or 4D data. Hence, they are commonly used as foundation models for various generation tasks. Well-trained video generative models have shown their potential to generate multi-view images as a 3D generator [54, 7, 13, 22, 33]. SV3D [54] adapts a latent video diffusion model (Stable Video Diffusion - SVD [2]) to generate novel multiple views. In this work, we leverage the pre-trained video generation model for 4D generation by adding an additional view attention layer to align the multiview images.

**4D Generation.** On one hand, recent optimization-based methods [48, 1, 27, 42, 68, 18, 66, 64] can generate 4D content by distilling pre-trained diffusion models in a 4D representation [5, 20, 34] via SDS loss [40]. MAV3D [48] distills a text-to-video diffusion model and several following works [1, 27, 42, 68] combine both video and multi-view diffusion models to leverage the motion and multi-view priors. However, they tend to take hours to generate 4D content due to time-consuming optimization. On the other hand, photogrammetry-based methods [61, 37] mimic a 3D object capture pipeline by directly generating multi-frame multi-view images of a 4D content with dynamic and multi-view consistency and then directly reconstructing 4D representations with them. Although these methods have much faster speeds, inference-only pipelines are adopted due to the paucity of the 4D data, thus making the spatial-temporal consistency still unsatisfactory. In this work, we proposed a monolithic model to generate a more consistent image grid, and we use a carefully curated 4D dataset to train the model.

## 3 Method

Our main idea is to build multi-frame and multi-view consistency in a 4D object by surgically combining the frame-consistency in a video diffusion model, with the multi-view consistency in a multi-view diffusion model. In our case, we choose SVD [2] and SV3D [54] as the video and multi-view diffusion models respectively, for the advantages reasoned below. However, it is to be noted that any choice of attention-based diffusion models should work.

### 3.1 Novel View Video Synthesis via SV4D

**Problem Setting.** Formally, we begin with a monocular input video  $\mathbf{J} \in \mathbb{R}^{F \times D}$  of a dynamic object of  $F$  frames and  $D := 3 \times H \times W$  dimensions, where  $H$  and  $W$  are the height and width of each frame. Our goal is to generate an image matrix  $\mathbf{M} \in \mathbb{R}^{V \times F \times D}$  of the 4D object consisting of  $F$  dynamic frames and  $V$  camera views of each frame. Similar to SV3D [54], the multi-view frames follow a camera pose trajectory  $\boldsymbol{\pi} \in \mathbb{R}^{V \times 2} = \{(e_v, a_v)\}_{v=1}^V$  as a sequence of  $V$  tuples of elevation  $e$  and azimuth angles  $a$ . We generate this image matrix by iteratively denoising samples from a learned conditional distribution  $p(\mathbf{M}|\mathbf{J}, \boldsymbol{\pi})$ , parameterized by a 4D diffusion model. Due to memory limitation in generating  $V \times F$  images simultaneously, we break down the full generation process into multiple submatrix generation steps.

**SV4D Network.** Our goal is to make the generated image matrix  $\mathbf{M}$  dynamically consistent in the “frame” axis, and multi-view consistent in the “view” axis (see Fig. 2 left). To achieve this, we condition the image matrix generation on the corresponding **frames** of the monocular input video  $\{\mathbf{M}_{0,f}\} = \mathbf{J}$  as well as the **views** from reference multi-view images of the first video frame  $\{\mathbf{M}_{v,0}\}$ , which provide motion and multi-view information of the 4D object, respectively. Without loss of generality, we obtain the reference multi-view images of the first input frame by sampling from a pre-

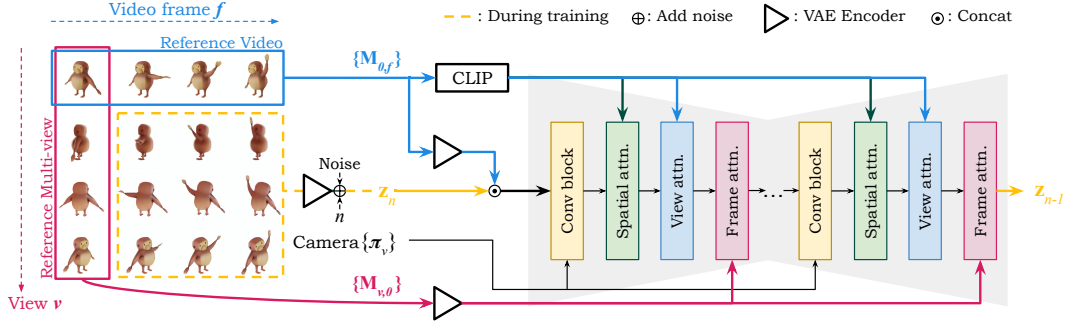


Figure 2: **SV4D Model Architecture.** For camera conditioning, we feed the sinusoidal embedding of camera viewpoints to the convolutional blocks in the UNet, and use the input video for cross-attention conditioning in the spatial and view attention blocks. To improve temporal consistency, we introduce an additional motion attention block, conditioned on the corresponding views of the first frame.

trained SV3D [54] model, which can be expressed as  $p(\{\mathbf{M}_{v,0}\}|\{\mathbf{M}_{0,0}\}, \boldsymbol{\pi})$ . The overall SV4D novel view video synthesis can be rewritten as sampling from the distribution  $p(\mathbf{M}|\{\mathbf{M}_{0,f}\}, \{\mathbf{M}_{v,0}\}, \boldsymbol{\pi})$

We build the SV4D network based on SVD [2] and SV3D [54] models to combine the advantages of both video and multi-view diffusion models. As illustrated in Fig. 2, SV4D consists of a UNet with multiple layers, where each layer contains a sequence of one residual block with Conv3D layers and three transformer blocks with attention layers: spatial, view, and frame attention. Similar to SV3D, the residual Conv block takes in the noisy latents of flattened image matrix as well as handles the incorporation of conditioning camera poses  $\{\boldsymbol{\pi}_v\}$ , and the spatial attention layer handles image-level details by performing attention across the image width and height axes. To better capture motion cues in the input monocular video  $\{\mathbf{M}_{0,f}\}$ , we concatenate its VAE latents to the noisy latents  $\mathbf{z}_n$  before feeding it to the UNet.

For multi-view consistency, the view attention block transposes the features and performs attention in the multi-view axis. By using the CLIP embedding of corresponding input frames  $\{\mathbf{M}_{0,f}\}$  as cross-attention conditioning, it allows the network to learn spatial consistency across novel views while maintaining the semantic context from the input video.

To further ensure dynamic consistency across video frames, we insert a frame attention layer in each UNet block, which applies the attention mechanism in the video frame dimension. The frame attention of each novel view video is conditioned on the corresponding reference view (of the first frame)  $\{\mathbf{M}_{v,0}\}$  via cross-attention, allowing the network to preserve dynamic coherence starting from the first frame. We initialize the weights of the frame attention layers from SVD and the rest of the network from SV3D<sub>p</sub>, to leverage the generalizability as well as rich dynamic and multi-view priors learned from large-scale video and 3D datasets.

**ObjaverseDy Dataset.** Considering that there exists no large-scale training datasets with dynamic 3D objects, we curate a new 4D dataset from the existing Objaverse dataset [9, 8], a massive dataset with annotated 3D objects. Objaverse includes animated 3D objects, however, several of these animated 3D objects are not suitable for training due to having too few animated frames or insufficient motion. In addition, in the rendering stage, the common rendering and sampling setting may cause some issues. For example, dynamic objects may be out of the image if the camera distance is fixed because they have global motion; the motion of objects may be too fast or too slow if the temporal sampling step is fixed.

We follow several steps to curate and clean the 4D objects for our training purposes. We first filter out the objects based on a review of licenses. Then, we remove the objects whose animated frames are too few. To further filter out objects with minimal motion, we subsample keyframes from each video and apply simple thresholding on the maximum  $L1$  distance between these frames as motion measurement. To render the training novel view videos, we flexibly choose the camera distance from the object. Starting from a base value, we increase the camera distance until the object fits within all frames of the rendered images. We also dynamically adjust the temporal sampling step. Starting from a base value, we increase the temporal sample step until the  $L1$  distance between consecutive




Figure 3: **SV4D Model Sampling** (7 frames and 3 views are for illustrative purposes only). To extend the generated multi-view videos while preserving temporal consistency, we propose a novel mixed-sampling strategy during inference. We first sample a sparse set of anchor frames, then use the anchor frames as new conditioning images to densely sample/interpolate the middle frames. To ensure a smooth transition between consecutive generations, we alternatively use the first (forward) or last (backward) frame within a time window for conditioning during dense sampling.

keyframes exceeds a certain threshold. These steps ensure a high-quality collection of 4D objects, with rendered multi-view videos that form our ObjaverseDy dataset.

**Training Details.** We train SV4D on our ObjaverseDy dataset. We choose to finetune from the SV3D<sub>p</sub> [54] model to output 40 frames ( $F = 5$  frames along each of the  $V = 8$  views), where the parameters in the frame attention layers are initialized from SVD-xt [2]. Similar to SV3D, we train SV4D progressively by first training on the static camera orbits for 40K iterations, then fine-tuning it for 20K iterations on the dynamic orbits. We use an effective batch size of 16 during training on 2 nodes of 8 80GB H100 GPUs. For more training details, please see the appendix.

**Inference Sampling Scheme.** Due to memory limitations, we cannot generate all the novel view frames at once when  $F > 5$ . To generate the full  $V \times F$  image matrix with arbitrary length videos during inference, one can naively run the submatrix generation in an auto-regressive manner. However, we observe that it often leads to severe artifacts due to error propagation during the generation process. Hence, we design a novel sampling scheme to mitigate the issue. The proposed sampling scheme is illustrated in Fig. 3. Note that we only show motion frame extensions here for illustrative purposes. We first generate a sparse set of anchor frames with SV4D (*interleaved sampling*) as shown in Fig. 3 (left). Then we use the anchor frames as new reference views to densely sample the remaining frames (*dense sampling*). To ensure a smooth transition between consecutive generations, we alternatively use the first (forward) or last (backward) anchor frame for conditioning at each diffusion step, as shown in Fig. 3 (middle). In our experiments (Fig. 8), we demonstrate that our sampling scheme can generate novel view videos that are more temporally consistent compared to using an off-the-shelf video interpolation model [25] to interpolate in between the SV4D-generated anchor frames.

**Spatio-temporal CFG Scaling.** During inference sampling, SVD uses a linearly increasing scale for classifier-free guidance (CFG) over time (video frame axis) to avoid losing context from the first reference frame. SV3D, on the other hand, adopts a triangular CFG scaling (in the view axis) when generating an orbital video around a static 3D object to avoid over-sharpened or over-saturated outputs. However, we observe that neither of these CFG scaling schemes suit the spatio-temporal nature in 4D image matrix generation and often cause undesirable artifacts. Hence, we propose a novel spatio-temporal CFG scaling which integrates a linearly-increasing scale in the frame axis and a triangular wave in the view axis (linearly increase when moving away from the input view, and linearly decrease when rotating back towards the input view). Details can be found in the appendix.

### 3.2 4D Optimization from SV4D Generated Novel View Videos

**4D Representation.** As shown in Fig. 4, given the novel view videos generated by SV4D, we optimize a 4D representation to reconstruct the dynamic 3D asset. Formally, we learn a neural representation  $\Psi_\theta : (\mathbf{x}, t) \mapsto (\sigma, c)$  that maps the sampled 3D points  $\mathbf{x} = (x, y, z)$  along a camera ray, and the time embedding  $t$  (i.e. continuous version of “frame” in the earlier discrete case) to its volumetric density  $\sigma \in \mathbb{R}_+$  and color  $c \in \mathbb{R}_+^3$ . Similar to D-NeRF [41], we represent a 4D object by the composition of a canonical NeRF that captures the static 3D appearance and a deformation field which handles the object motion across time. For each 3D point  $\mathbf{x}$  in the canonical space, we trilinearly interpolate the multi-resolution hash grid features following Instant-NGP [35], and decode them as density and color via  $\text{MLP}_\sigma$  and  $\text{MLP}_c$ , respectively. The deformation field is represented by

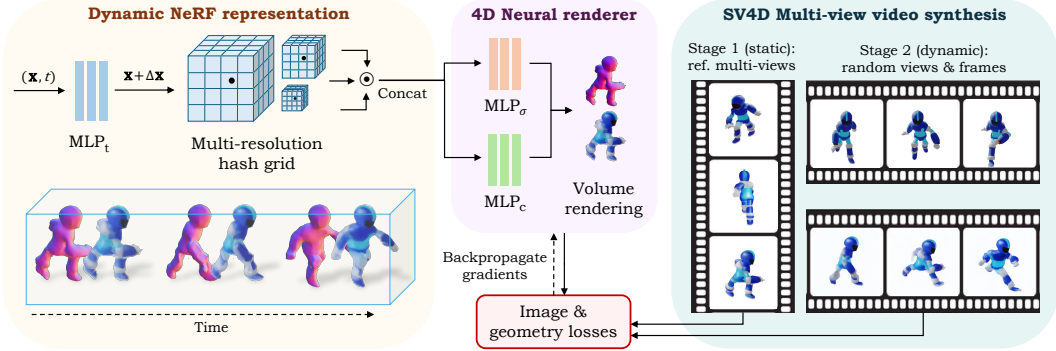


Figure 4: **Overview of Optimization Framework.** We first use the reference multi-view images (of the first frame) to optimize a static NeRF represented by a multi-resolution hash grid as well as density and color MLPs. Then, we unfreeze the temporal deformation MLP and optimize the dynamic NeRF with randomly sampled views and frames.

an MLP network  $MLP_t$  conditioned on time embedding  $t$ , mapping temporal deformation of  $\mathbf{x}$  to the common canonical space  $\mathbf{x} + \Delta\mathbf{x}$ . Overall, the dynamic NeRF parameters  $\theta$  include the canonical hash grid,  $MLP_t$ ,  $MLP_\sigma$ , and  $MLP_c$ . Since we do not exhaustively sample temporal timestamps or dense spatial views like in prior SDS-based approaches, we observe that this dynamic NeRF representation produces better 4D results compared to other representations such as 4D Gaussian Splatting [60], which suffers from flickering artifacts and does not interpolate well across time or views.

**Optimization Details.** By leveraging the consistent image matrix generated by SV4D as pseudo ground-truths, we adopt a simple photometry-based optimization without the cumbersome SDS losses. The reconstruction losses include a pixel-level MSE loss, mask loss, and a perceptual LPIPS [67] loss. Similar to SV3D [54], we also use several geometric priors to regularize the output shapes, such as a mono normal loss similar to MonoSDF [65] as well as bilateral depth and normal smoothness losses [54] to encourage smooth 3D surfaces where the projected image gradients are low.

For training efficiency and stability, we follow a coarse-to-fine, static-to-dynamic strategy to optimize a 4D representation. That is, we first freeze the deformation field  $MLP_t$  and only optimize the canonical NeRF on the multi-view images of the first frame, while gradually increasing the rendering resolution from  $128 \times 128$  to  $512 \times 512$ . Then, we unfreeze  $MLP_t$  and randomly sample  $4 \text{ frames} \times 4 \text{ views}$  for training. Following the static-to-dynamic strategy, we also gradually optimize the time embedding  $t$  from low to high temporal frequency. In our experiments, we find that sampling more timestamps in one batch and progressive optimization techniques are crucial to 4D output quality. We render the dynamic NeRF at  $512 \times 512$  resolution and use an Adam [21] optimizer to train all model parameters. The overall optimization takes roughly 15-20 minutes per object. More implementation details can be found in the appendix.

## 4 Experiments

**Datasets.** We evaluate SV4D-synthesized novel view videos and 4D optimization results on the synthetic datasets ObjaverseDy, Consistent4D [9], and the real-world videos dataset DAVIS [38, 39, 4]. Consistent4D dataset contains some dynamic objects collected from Objaverse [9, 8]. We excluded these objects from our training set to make a fair comparison. We used the same input video and evaluated views as Consistent4D. DAVIS dataset consists of single-view videos in the wild with corresponding foreground masks. For qualitative evaluation, we selected single-view videos with static cameras.

**Metrics.** We use the SV4D model to generate multiple novel view videos corresponding to the trajectories of the ground truth camera in the evaluation datasets. We compare each generated image with its corresponding ground-truth, in terms of Learned Perceptual Similarity (LPIPS [67]) and CLIP-score (CLIP-S) to evaluate visual quality. In addition, we evaluate the video coherence by

Table 1: **Evaluation of Novel View Video Synthesis on the Consistent4D Dataset.** SV4D can achieve better video frame consistency while maintaining comparable image quality. † Our reproduced version.

Model	LPIPS↓	CLIP-S↑	FVD-F↓
SV3D [54]	<b>0.129</b>	0.925	989.53
Diffusion <sup>2</sup> † [61]	0.171	0.919	1430.35
STAG4D† [66]	0.131	<b>0.929</b>	861.88
SV4D	<b>0.129</b>	<b>0.929</b>	<b>677.68</b>

Table 2: **Evaluation of 4D Outputs on the Consistent4D Dataset.** SV4D can achieve better visual quality and video frame smoothness.

Model	LPIPS↓	CLIP-S↑	FVD-F↓
Consistent4D [18]	0.160	0.87	1133.93
STAG4D [66]	0.126	0.91	992.21
Efficient4D [37]	0.140	<b>0.92</b>	-
4DGen [64]	0.130	0.89	-
DG4D [42]	0.160	0.87	-
GaussianFlow [10]	0.140	0.91	-
SV4D	<b>0.118</b>	<b>0.92</b>	<b>732.40</b>

Table 3: **Evaluation of Novel View Video Synthesis on the ObjaverseDy Dataset.** SV4D can get superior performance in video frame consistency and multi-view consistency. † Our reproduced version.

Model	LPIPS↓	CLIP-S↑	FVD-F↓	FVD-V↓	FVD-Diag↓	FV4D↓
SV3D [54]	<b>0.131</b>	<b>0.920</b>	729.67	375.49	526.78	690.49
Diffusion <sup>2</sup> † [61]	0.152	0.905	1067.78	442.18	612.82	977.19
STAG4D† [66]	0.133	0.917	652.43	469.07	636.83	546.56
SV4D	0.136	<b>0.920</b>	<b>585.09</b>	<b>331.94</b>	<b>503.02</b>	<b>470.46</b>

Table 4: **Evaluation of 4D outputs on the ObjaverseDy dataset.** SV4D consistently outperforms baselines in terms of all metrics, demonstrating superior performance in visual quality (*LPIPS* and *CLIP-S*), video frame consistency (*FVD-F*), multi-view consistency (*FVD-V*), and multi-view video consistency (*FVD-Diag* and *FV4D*).

Model	LPIPS↓	CLIP-S↑	FVD-F↓	FVD-V↓	FVD-Diag↓	FV4D↓
Consistent4D [18]	0.165	0.896	880.54	488.38	741.52	871.95
STAG4D [66]	0.158	0.860	929.10	453.62	663.50	1003.16
DreamGaussian4D [42]	0.152	0.897	697.80	450.58	615.68	638.15
SV4D	<b>0.131</b>	<b>0.905</b>	<b>659.66</b>	<b>368.53</b>	<b>525.65</b>	<b>614.35</b>

reporting FVD [52], a video-level metric commonly used in video generation tasks. We calculate FVD with different ways (see Fig. 7):

- *FVD-F*: calculate FVD over frames at each view.
- *FVD-V*: calculate FVD over views at each frame.
- *FVD-Diag*: calculate FVD over the diagonal images of the image matrix.
- *FV4D*: calculate FVD over all images by scanning them in a bidirectional raster order.

**Baselines.** For *novel view video synthesis*, we compare SV4D with several recent methods capable of generating multiple novel view videos from a single-view video, including SV3D [54], Diffusion<sup>2</sup> [61], STAG4D [61]. It is to be noted that all of these methods are inference-only techniques, and do not involve direct training in the 4D space like our method. We run SV3D to generate multi-view images at each video frame separately. Diffusion<sup>2</sup> has not released their code yet. STAG4D used Zero123++ [44] as the multi-view generator, which fixed the view angle, and hence the novel views generated from STAG4D cannot be changed to be consistent with the views evaluated. We reproduced Diffusion<sup>2</sup> and STAG4D with SV3D as the multi-view generator. SV3D has been shown to generate more consistent 3D results than Zero123++, so this serves as a better baseline. For *4D generation*, we compare SV4D with other methods that can generate 4D representations, including Consistent4D [18], STAG4D [66], DreamGaussian4D (DG4D) [42], GaussianFlow [10], 4DGen [64], Efficient4D [37].

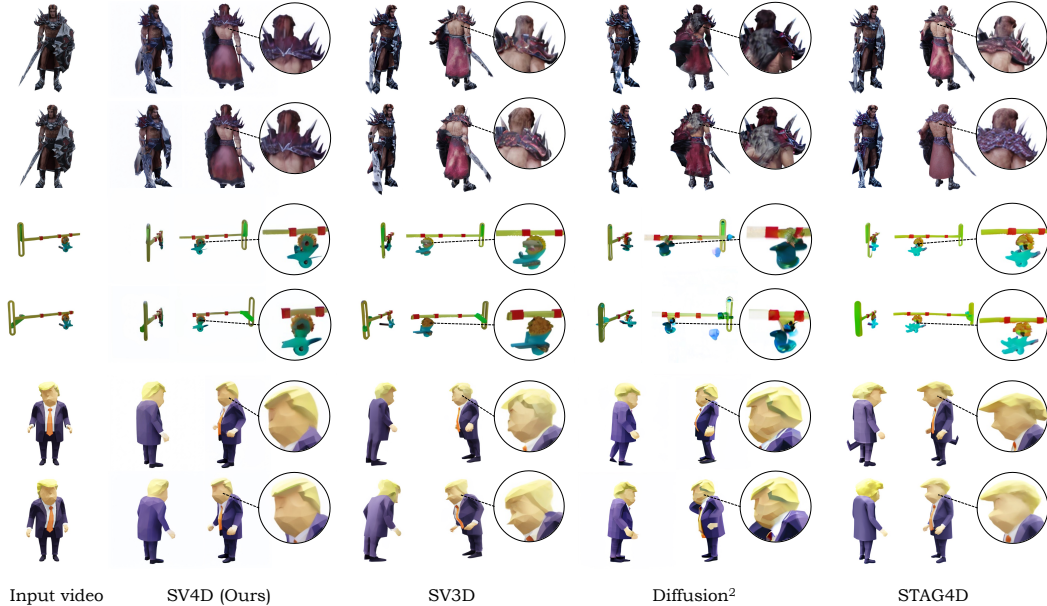


Figure 5: **Visual Comparison of Novel View Video Synthesis Results.** We show two frames in the input videos and two novel-view results of the corresponding frames. Compared to the baseline methods, SV4D outputs contain geometry and texture details that are more faithful to the input video and consistent across frames.

#### 4.1 Quantitative Comparison

**Novel View Video Synthesis.** We quantitatively compare our method with the baselines in terms of novel view video synthesis results. Table 1 reports the comparisons in the Consistent4D dataset. Due to the fact that the evaluated views are too sparse, we only report  $FVD-F$  in the Consistent4D dataset. Our method can achieve much better performance in terms of video frame consistency while maintaining comparable performance in terms of image quality. In particular, our approach has a significant reduction of 31.5% and 21.4% in  $FVD-F$  compared to SV3D and STAG4D, respectively. Table 3 reports the comparisons in the Objaverse dataset. Our method can achieve much better video frame consistency as well as multi-view consistency while maintaining comparable performance in terms of image quality. In particular, our approach has much lower  $FVD-F$  compared to baseline methods, demonstrating our generated videos are much smoother. In addition, SV4D achieves better  $FVD-V$  which shows better multi-view consistency. Our method also surpasses the previous state-of-the-art methods in terms of  $FVD-Diag$  and  $FV4D$ , proving that the synthesized novel view videos have better video frame and multi-view consistency.

**4D Generation.** We quantitatively compare our optimized 4D outputs with the baselines in Consistent4D and ObjaverseDy dataset, as shown in Tables 2 and 4, respectively. Our method consistently outperforms baselines in terms of all metrics, demonstrating superior performance in visual quality ( $LPIPS$  and  $CLIP-S$ ), motion smoothness ( $FVD-F$ ), multi-view smoothness ( $FVD-V$ ), and motion-multi-view joint smoothness ( $FVD-Diag$  and  $FV4D$ ).

#### 4.2 Visual Comparison

**Novel View Video Synthesis.** In Fig. 5 we show the visual comparison of our multi-view video synthesis results against other methods. We observe that applying SV3D frame-by-frame leads to inconsistent geometry and texture in novel view videos. Diffusion<sup>2</sup> slightly improves the temporal coherence but tends to produce blurry or flickering artifacts. STAG4D can produce smoother videos but often fails at capturing large motion. Compared to these methods, SV4D can generate high-quality multi-view videos that are detailed, faithful to input videos, and temporally consistent.

**4D Generation.** We compare our generated 4D results with the prior methods in Fig. 6. For each video, we render the 4D outputs at two different timestamps and two novel views. Consistent4D and



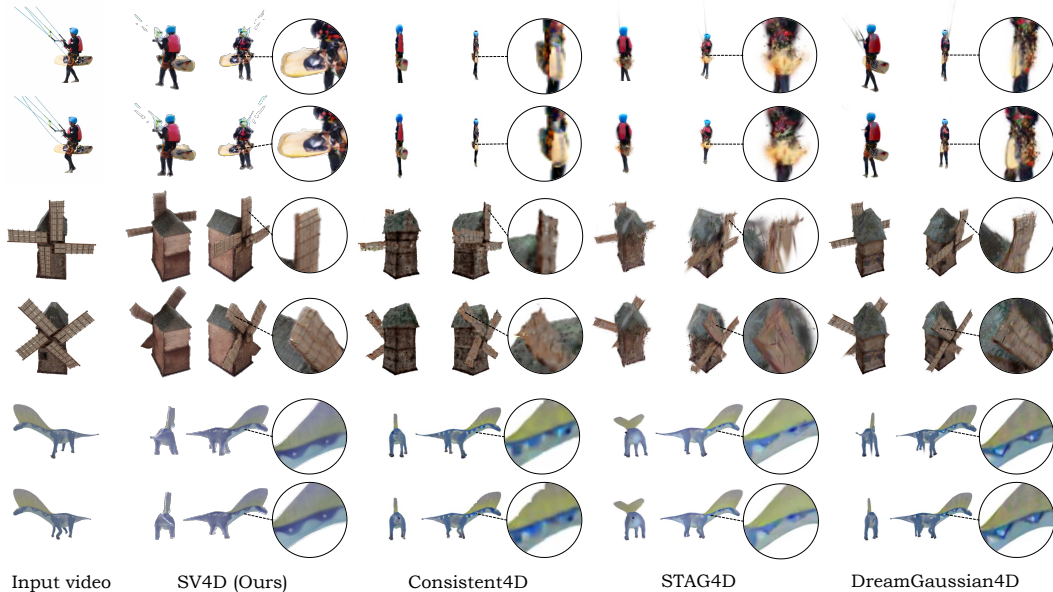


Figure 6: **Visual Comparison of Generated 4D Outputs.** By leveraging the consistent multi-view videos generated by SV4D, we can learn higher-quality 4D assets compared to the prior works with SDS-based losses. Our results are more detailed, consistent, and faithful to the input videos.

Table 5: **Evaluation of Novel View Video Synthesis (Anchor Frames Only).** SV4D can effectively sample frames with faithful consistency and visual details. † Our reproduced version.

Model	ObjaverseDy		Consistent4D	
	FVD-F↓	FV4D↓	FVD-F↓	FV4D↓
SV3D [54]	700.72	831.18	656.59	780.99
Diffusion <sup>2</sup> † [61]	896.98	890.70	801.99	1093.24
STAG4D† [66]	765.79	669.28	601.23	659.79
SV4D	<b>629.22</b>	<b>569.08</b>	<b>469.73</b>	<b>621.00</b>

Table 6: **Evaluation of Different Sampling Strategies on ObjaverseDy Dataset.** SV4D sampling can effectively generate full image matrices with faithful consistency and visual details.

Sampling	LPIPS↓	FVD-F↓	FVD-Diag↓	FV4D↓
Interleaved	0.136	731.28	567.15	717.21
Independent	0.136	663.41	512.14	488.31
AMT [25]	0.140	612.61	505.36	472.87
SV4D (anchor)	<b>0.130</b>	629.22	-	569.08
SV4D	0.136	<b>585.09</b>	<b>503.02</b>	<b>470.46</b>

STAG4D often produce blurry outputs with inconsistent geometry and texture. DreamGaussian4D can generate finer texture details but still suffers from flickering artifacts and sometimes creates inaccurate geometry. Moreover, all these methods rely on the computationally expensive SDS loss, which is prone to spatially incoherent results and over-saturated texture. On the contrary, SV4D optimizes 4D assets using purely photometric and geometric losses, resulting in smoother videos with realistic and faithful object appearance.

### 4.3 User Study

In addition to the quantitative evaluation, we also conduct two user studies on our multi-view video synthesis results and 4D outputs, respectively. Concretely, we randomly select 10 real-world videos from the DAVIS dataset and 10 synthetic videos from the Objaverse or Consistent4D datasets. For each video, we randomly choose a novel camera view and ask the user to compare the novel view videos generated by 4 different methods (SV4D and 3 baselines). The users are asked to select a video that “looks more stable, realistic, and closely resembles the reference subject”. For multi-view video synthesis, SV4D results are preferred 73.3% over per-frame SV3D (13.7%), Diffusion<sup>2</sup> (5.0%), and STAG4D (8.0%) among multiple participants. For the optimized 4D outputs, SV4D achieves 60% overall preference against Consistent4D (12.7%), STAG4D (9.7%), and DreamGaussian4D (17.6%) among multiple users.

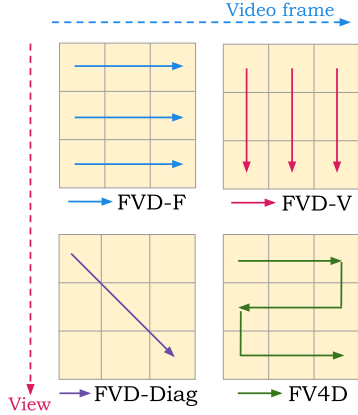


Figure 7: **Illustrations of Video and 4D Metrics.** FVD-F evaluates coherence between video frames from a fixed view. FVD-V captures multi-view consistency. We also design FVD-Diag and FV4D to evaluate 4D consistency by traversing the image matrix through different paths.




Figure 8: **Visual Comparison of SV4D Sampling and Video Interpolation.** Given the same sparse anchor frames, SV4D can effectively sample the middle frames with faithful motion and details, whereas an off-the-shelf video interpolation model creates blurry results or missing parts.

#### 4.4 Ablative Analyses

We conduct several ablation experiments to validate the effectiveness of our model’s design choices. We summarize the key findings below.

**SV4D Can Generate Anchor Frames with Better Consistency.** To validate the quality of anchor frames, we compare the anchor frames ( $8 \text{ views} \times 5 \text{ video frames}$ ) generated from SV4D and the baselines on consistent4D and ObjaverseDy dataset, as shown in Table 5. SV4D results have much lower *FVD-F* and *FV4D*, showing much better video frame and multi-view consistency while maintaining comparable performance in terms of visual quality.

**SV4D Sampling Is Better Than Off-the-shelf Interpolation Method.** Extending anchor frames to the full  $V \times F$  image matrix is not trivial. In Table 6, we compare our SV4D sampling (*ours*) and the off-the-shelf interpolation method AMT [25]. We also show the results of anchor frames as a reference (*SV4D (anchor)*). The model with SV4D sampling outperforms AMT across all metrics. Fig. 8 validates this observation, where we can see that the SV4D sampling can synthesize distinct images while the results of the off-the-shelf interpolation method have a lot of motion blur. Table 6 also shows two intuitive sampling strategies, independent and interleaved sampling, which generate the videos with worse consistency. See the Appendix for details on these two sampling strategies. These results validate that the SV4D sampling is effective.

## 5 Conclusion

We present SV4D, a latent video diffusion model for novel view video synthesis and 4D generation. Given an input video, SV4D can generate multiple novel view videos that are dynamically and spatially consistent, by leveraging the dynamic prior in SVD and multi-view prior in SV3D within a unified architecture. The generated novel view videos can then be used to efficiently optimize a 4D asset without SDS losses from one or multiple diffusion models. Our extensive experiments show that SV4D outputs are more multi-frame and multi-view consistent than existing methods and generalizable to real-world videos, achieving the state-of-the-art performance on novel view video synthesis and 4D generation. We believe that SV4D provides a solid foundation model for further research on dynamic 3D object generation.

**Acknowledgments.** We thank Prem Akkaraju, Christian Laforte, Adam Letts, Josephine Parquet, Savannah Martin, Cathy Lee, Amir Etemadieh, Patricia Wood, Nadiia Klokova, Shan Shan Wong, Ryan Holeman, Jerry Chi, Javad Azimi for their help in various aspects of the model development and release. We also thank Eric Courtemanche for his help with visual results.

## References

- [1] Sherwin Bahmani, Ivan Skorokhodov, Victor Rong, Gordon Wetzstein, Leonidas Guibas, Peter Wonka, Sergey Tulyakov, Jeong Joon Park, Andrea Tagliasacchi, and David B. Lindell. 4D-fy: Text-to-4D generation using hybrid score distillation sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- [2] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendeleevitch, 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.
- [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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [4] Sergi Caelles, Jordi Pont-Tuset, Federico Perazzi, Alberto Montes, Kevis-Kokitsi Maninis, and Luc Van Gool. The 2019 davis challenge on video object segmentation. *arXiv*, 2019.
- [5] Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [6] Zilong Chen, Feng Wang, and Huaping Liu. Text-to-3D using gaussian splatting. *arXiv preprint arXiv:2309.16585*, 2023.
- [7] Zilong Chen, Yikai Wang, Feng Wang, Zhengyi Wang, and Huaping Liu. V3D: Video diffusion models are effective 3D generators. *arXiv preprint arXiv:2403.06738*, 2024.
- [8] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, Eli VanderBilt, Aniruddha Kembhavi, Carl Vondrick, Georgia Gkioxari, Kiana Ehsani, Ludwig Schmidt, and Ali Farhadi. Objaverse-xl: A universe of 10m+ 3d objects. *arXiv preprint arXiv:2307.05663*, 2023.
- [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*, 2023.
- [10] Quankai Gao, Qiangeng Xu, Zhe Cao, Ben Mildenhall, Wenchao Ma, Le Chen, Danhang Tang, and Ulrich Neumann. GaussianFlow: Splatting gaussian dynamics for 4D content creation. *arXiv preprint arXiv:2403.12365*, 2024.
- [11] Pengsheng Guo, Hans Hao, Adam Caccavale, Zhongzheng Ren, Edward Zhang, Qi Shan, Aditya Sankar, Alexander G Schwing, Alex Colburn, and Fangchang Ma. StableDreamer: Taming noisy score distillation sampling for text-to-3d. *arXiv preprint arXiv:2312.02189*, 2023.
- [12] Yuwei Guo, Ceyuan Yang, Anyi Rao, Yaohui Wang, Yu Qiao, Dahua Lin, and Bo Dai. AnimateDiff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv preprint arXiv:2307.04725*, 2023.
- [13] Junlin Han, Filippos Kokkinos, and Philip Torr. VFusion3D: Learning scalable 3D generative models from video diffusion models. *arXiv preprint arXiv:2403.12034*, 2024.
- [14] Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. Latent video diffusion models for high-fidelity long video generation. *arXiv*, 2022.
- [15] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 2022.
- [16] 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. *arXiv preprint arXiv:2311.04400*, 2023.
- [17] Hanwen Jiang, Zhenyu Jiang, Yue Zhao, and Qixing Huang. LEAP: Liberate sparse-view 3D modeling from camera poses. *ArXiv*, 2310.01410, 2023.

- [18] Yanqin Jiang, Li Zhang, Jin Gao, Weimin Hu, and Yao Yao. Consistent4D: Consistent 360 $\{\backslash\text{deg}\}$  dynamic object generation from monocular video. *arXiv preprint arXiv:2311.02848*, 2023.
- [19] Animesh Karnear, Niloy J Mitra, Andrea Vedaldi, and David Novotny. HoloFusion: Towards photo-realistic 3D generative modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023.
- [20] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 2023.
- [21] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [22] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- [23] Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan Sunkavalli, Greg Shakhnarovich, and Sai Bi. Instant3d: Fast text-to-3d with sparse-view generation and large reconstruction model. *arXiv preprint arXiv:2311.06214*, 2023.
- [24] Weiyu Li, Rui Chen, Xuelin Chen, and Ping Tan. SweetDreamer: Aligning geometric priors in 2D diffusion for consistent text-to-3D. *arXiv preprint arXiv:2310.02596*, 2023.
- [25] Zhen Li, Zuo-Liang Zhu, Ling-Hao Han, Qibin Hou, Chun-Le Guo, and Ming-Ming Cheng. Amt: All-pairs multi-field transforms for efficient frame interpolation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2023.
- [26] Yixun Liang, Xin Yang, Jiantao Lin, Haodong Li, Xiaogang Xu, and Yingcong Chen. LucidDreamer: Towards high-fidelity text-to-3D generation via interval score matching. *arXiv*, 2023.
- [27] Huan Ling, Seung Wook Kim, Antonio Torralba, Sanja Fidler, and Karsten Kreis. Align your gaussians: Text-to-4D with dynamic 3D gaussians and composed diffusion models. *arXiv preprint arXiv:2312.13763*, 2023.
- [28] 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.
- [29] Minghua Liu, Chao Xu, Haiyan 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. In *Advances in Neural Information Processing Systems*, 2023.
- [30] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3D object. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023.
- [31] 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.
- [32] 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.
- [33] 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.
- [34] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 2021.

- [35] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 2022.
- [36] Zijie Pan, Jiachen Lu, Xiatian Zhu, and Li Zhang. Enhancing high-resolution 3D generation through pixel-wise gradient clipping. *arXiv preprint arXiv:2310.12474*, 2023.
- [37] Zijie Pan, Zeyu Yang, Xiatian Zhu, and Li Zhang. Fast dynamic 3D object generation from a single-view video. *arXiv preprint arXiv:2401.08742*, 2024.
- [38] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [39] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alexander Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv:1704.00675*, 2017.
- [40] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3D using 2D diffusion. *arXiv preprint arXiv:2209.14988*, 2022.
- [41] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [42] Jiawei Ren, Liang Pan, Jiaxiang Tang, Chi Zhang, Ang Cao, Gang Zeng, and Ziwei Liu. DreamGaussian4D: Generative 4D gaussian splatting. *arXiv preprint arXiv:2312.17142*, 2023.
- [43] Kyle Sargent, Zizhang Li, Tanmay Shah, Charles Herrmann, Hong-Xing Yu, Yunzhi Zhang, Eric Ryan Chan, Dmitry Lagun, Li Fei-Fei, Deqing Sun, et al. ZeroNVS: Zero-shot 360-degree view synthesis from a single real image. *arXiv preprint arXiv:2310.17994*, 2023.
- [44] 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.
- [45] 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.
- [46] Yukai Shi, Jianan Wang, He Cao, Boshi Tang, Xianbiao Qi, Tianyu Yang, Yukun Huang, Shilong Liu, Lei Zhang, and Heung-Yeung Shum. Toss: High-quality text-guided novel view synthesis from a single image. *arXiv preprint arXiv:2310.10644*, 2023.
- [47] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022.
- [48] Uriel Singer, Shelly Sheynin, Adam Polyak, Oron Ashual, Iurii Makarov, Filippos Kokkinos, Naman Goyal, Andrea Vedaldi, Devi Parikh, Justin Johnson, and Yaniv Taigman. Text-to-4D dynamic scene generation. *arXiv:2301.11280*, 2023.
- [49] Jingxiang Sun, Bo Zhang, Ruizhi Shao, Lizhen Wang, Wen Liu, Zhenda Xie, and Yebin Liu. DreamCraft3D: Hierarchical 3D generation with bootstrapped diffusion prior. *arXiv preprint arXiv:2310.16818*, 2023.
- [50] 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.
- [51] Dmitry Tochilkin, David Pankratz, Zexiang Liu, Zixuan Huang, Adam Letts, Yangguang Li, Ding Liang, Christian Laforte, Varun Jampani, and Yan-Pei Cao. TripoSR: Fast 3D object reconstruction from a single image. *arXiv preprint arXiv:2403.02151*, 2024.
- [52] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018.

- [53] Vikram Voleti, Alexia Jolicoeur-Martineau, and Christopher Pal. MCVD: Masked conditional video diffusion for prediction, generation, and interpolation. In *Advances in Neural Information Processing Systems*, 2022.
- [54] Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitrii Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. SV3D: Novel multi-view synthesis and 3D generation from a single image using latent video diffusion. In *European Conference on Computer Vision*, 2024.
- [55] Peng Wang and Yichun Shi. ImageDream: Image-prompt multi-view diffusion for 3D generation. *arXiv preprint arXiv:2312.02201*, 2023.
- [56] Peng Wang, Hao Tan, Sai Bi, Yinghao Xu, Fujun Luan, Kalyan Sunkavalli, Wenping Wang, Zexiang Xu, and Kai Zhang. PF-LRM: Pose-free large reconstruction model for joint pose and shape prediction. *arXiv preprint arXiv:2311.12024*, 2023.
- [57] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. ProlificDreamer: High-fidelity and diverse text-to-3D generation with variational score distillation. In *Advances in Neural Information Processing Systems*, 2024.
- [58] Xinyue Wei, Kai Zhang, Sai Bi, Hao Tan, Fujun Luan, Valentin Deschaintre, Kalyan Sunkavalli, Hao Su, and Zexiang Xu. MeshLRM: Large reconstruction model for high-quality mesh. *arXiv preprint arXiv:2404.12385*, 2024.
- [59] Haohan Weng, Tianyu Yang, Jianan Wang, Yu Li, Tong Zhang, CL Chen, and Lei Zhang. Consistent123: Improve consistency for one image to 3D object synthesis. *arXiv preprint arXiv:2310.08092*, 2023.
- [60] Guanjuan Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4D gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- [61] Zeyu Yang, Zijie Pan, Chun Gu, and Li Zhang. Diffusion2: Dynamic 3D content generation via score composition of orthogonal diffusion models. *arXiv preprint arXiv:2404.02148*, 2024.
- [62] 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. *arXiv preprint arXiv:2310.03020*, 2023.
- [63] Taoran Yi, Jiemin Fang, Guanjuan Wu, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Qi Tian, and Xinggang Wang. GaussianDreamer: Fast generation from text to 3D gaussian splatting with point cloud priors. *arXiv preprint arXiv:2310.08529*, 2023.
- [64] Yuyang Yin, Dejie Xu, Zhangyang Wang, Yao Zhao, and Yunchao Wei. 4DGen: Grounded 4D content generation with spatial-temporal consistency. *arXiv preprint arXiv:2312.17225*, 2023.
- [65] 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*, 2022.
- [66] Yifei Zeng, Yanqin Jiang, Siyu Zhu, Yuanxun Lu, Youtian Lin, Hao Zhu, Weiming Hu, Xun Cao, and Yao Yao. STAG4D: Spatial-temporal anchored generative 4D gaussians. *arXiv preprint arXiv:2403.14939*, 2024.
- [67] 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*, 2018.
- [68] Yuyang Zhao, Zhiwen Yan, Enze Xie, Lanqing Hong, Zhenguo Li, and Gim Hee Lee. Animate124: Animating one image to 4D dynamic scene. *arXiv preprint arXiv:2311.14603*, 2023.
- [69] Linqi Zhou, Andy Shih, Chenlin Meng, and Stefano Ermon. DreamPropeller: Supercharge text-to-3D generation with parallel sampling. *arXiv preprint arXiv:2311.17082*, 2023.

- [70] Zi-Xin Zou, Zhipeng Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Yan-Pei Cao, and Song-Hai Zhang. Triplane meets gaussian splatting: Fast and generalizable single-view 3D reconstruction with transformers. *arXiv preprint arXiv:2312.09147*, 2023.