

SV4D 2.0: Enhancing Spatio-Temporal Consistency in Multi-View Video Diffusion for High-Quality 4D Generation

Chun-Han Yao^{1*} Yiming Xie^{1,2*} Vikram Voleti¹ Huaizu Jiang^{2†} Varun Jampani^{1†}
¹ Stability AI ² Northeastern University
 * Equal contribution † Equal advising

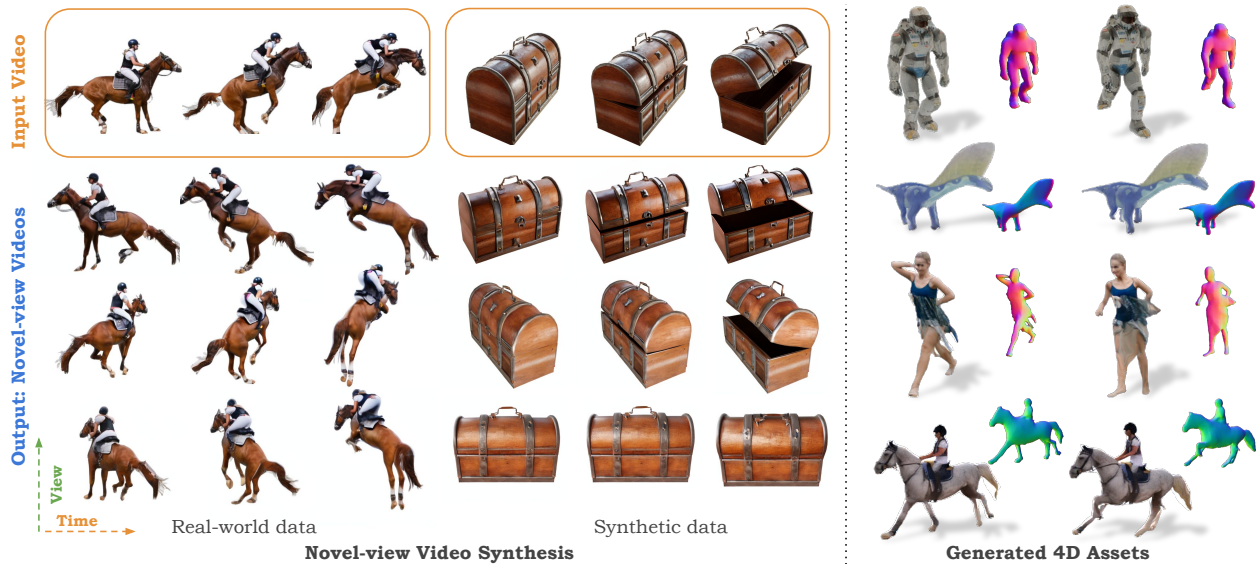


Figure 1. **SV4D 2.0** generates multiple novel-view videos from an input monocular video. The generated novel-view videos have high fidelity in terms of detail sharpness and consistency across view and time/frame axes, which can be used to optimize high-quality 4D assets. We also refer reviewers to the Supplemental Material for the video demo.

Abstract

We present *Stable Video 4D 2.0 (SV4D 2.0)*, a multi-view video diffusion model for dynamic 3D asset generation. Compared to its predecessor *SV4D* [69], *SV4D 2.0* is more robust to occlusions and large motion, generalizes better to real-world videos, and produces higher-quality outputs in terms of detail sharpness and spatio-temporal consistency. We achieve this by introducing key improvements in multiple aspects: 1) network architecture: eliminating the dependency of reference multi-views and designing blending mechanism for 3D and frame attention, 2) data: enhancing quality and quantity of training data, 3) training strategy: adopting progressive 3D-4D training for better generalization, and 4) 4D optimization: handling 3D inconsistency and large motion via 2-stage refinement and progressive frame sampling. Extensive experiments demonstrate significant performance gain by *SV4D 2.0* both visually and quantitatively, achiev-

ing better detail (-14% LPIPS) and 4D consistency (-44% FV4D) in novel-view video synthesis and 4D optimization (-12% LPIPS and -24% FV4D) compared to *SV4D*.

1. Introduction

Dynamic 3D generation, also known as 4D generation, aims to synthesize a moving object/scene in 3D space, capturing the dynamics of the subject across time. The generated 4D assets can be used in a wide range of applications such as video games, movie production, AR/VR experience, etc. Traditionally, such 4D asset generation requires a tedious process where 3D artists handcraft 3D geometry, texture, as well as rigging on the 3D object. With the recent advances in image [45, 46] and video [3, 4, 14, 60] diffusion models which show astonishing generation quality in the 2D domain, there has been a growing interest in extending them to 3D and 4D to speed up and/or automate this process.

Generating a dynamic 3D object from a single-view video

is a highly ill-posed problem due to the ambiguity of object appearance and motion at unseen camera views. Moreover, considering the lack of large-scale 4D datasets, and the requirement of high model capacity to handle the high-dimensional nature of 4D, it is challenging to train a native 4D generative model that can generalize to diverse object categories and motion. A common way to tackle this problem using diffusion models is to directly optimize a 4D asset using the priors in pre-trained video and multi-view generative models via a score-distillation sampling (SDS) loss [1, 22, 29, 43, 52, 73, 76, 79]. While proven to be effective in some cases, SDS loss can easily lead to artifacts such as spatio-temporal inconsistency or over-saturated color if not handled properly. Several recent works [54, 69, 70] have shown that reconstruction from synthesized videos is sufficient to produce high-quality 4D assets without SDS loss, by sampling videos of the subject at novel viewpoints via novel-view video synthesis (NVVS) techniques and using them as photogrammetry pseudo-ground truths.

The main challenge for NVVS is to generate multi-view videos that are consistent spatially (*i.e.*, across novel views), and temporally (*i.e.*, in dynamic motion). Earlier works [54, 70] on NVVS use separate models for novel-view synthesis and dynamic object motion, which often causes blurry details or poor spatio-temporal consistency, and thus can only handle small object motion. Some recent methods [24, 27, 69, 77] jointly learn spatial and temporal consistency in a single NVS model by repurposing and finetuning a video diffusion model on synthetic 4D data [9, 10]. However, their success is limited to low resolution [24, 27, 77], small object motion [27], or short video length [69].

In this work, we propose Stable Video 4D 2.0 (SV4D 2.0), a NVVS diffusion model that generates high-resolution multi-view videos of a dynamic 3D object, given a monocular video and a user-specified camera trajectory as input. We then use the NVVS outputs of SV4D 2.0 to optimize a 4D asset. In SV4D 2.0, we build upon the SV4D [69] framework and make several key modifications to improve 1) the ability to generate longer videos and sparser (*i.e.*, more distant) views, 2) robustness to self-occlusion and large motion, and 3) output quality in terms of details and spatio-temporal consistency. Our improvements over SV4D include:

- **Network architecture:** We make several modifications to the architecture, including 3D attention layers and disentangled α -blending of 3D and temporal information, enabling sparse novel-view synthesis and thus longer video generation in a single inference pass. Unlike SV4D which requires SV3D [61] to generate reference multi-views, we remove this dependency using a random masking strategy, resulting in robustness to self-occlusion in the first frame. This also enables more flexible inference sampling.
- **Data Curation:** We improve the quality of our synthetic training dataset by disentangling global transformation

from local motion, rectifying off-center objects, and filtering out objects with small motion or inconsistent scaling.

- **Training strategy:** Unlike most 4D generation methods [24, 27, 44, 70] which directly finetune a modified video or 3D generation model on 4D data, we adopt a progressive 3D-to-4D training. This smoothens the adaptation to new network architecture and bridges the 3D-4D domain gap, which helps preserve the rich video and 3D priors and thus generalizes better to in-the-wild videos.
- **4D optimization:** We propose a novel 2-stage optimization scheme for a 4D asset with visibility-weighted losses and progressive frame sampling to handle view inconsistencies in the synthesized novel-view videos and large object motion, respectively. This leads to improved 4D reconstruction quality compared to prior works.

We perform comprehensive evaluations of NVVS and 4D optimization on both synthetic datasets (ObjaverseDy [69] and Consistent4D [22]) and real-world videos (DAVIS [5, 38, 39]). As shown in Fig. 1, SV4D 2.0 outputs demonstrate high image quality, spatio-temporal consistency, and fidelity to the input videos. Despite being only trained on the synthetic 3D and 4D object datasets, SV4D 2.0 also demonstrates good generalization on real-world videos. Quantitatively, SV4D 2.0 achieves better details (-14% LPIPS) and 4D consistency (-44% FV4D) in video synthesis and -12% LPIPS and -24% FV4D in 4D optimization compared to SV4D, setting a new state-of-the-art for 4D generation.

2. Related Work

3D Generation. The SDS-based methods [8, 15, 26, 28, 36, 40, 47, 50, 53, 56, 65, 67, 72, 81] propose to distill priors from the 2D generative model via SDS loss to optimize the 3D content from text or image. Another line of works [19, 20, 58, 64, 66, 83] directly predict the 3D model of an static object via a large reconstruction model. Other approaches [23, 25, 30–34, 48, 49, 61, 63, 71] first generate dense consistent multi-view images, which are used for 3D content reconstruction. We follow this strategy, but generate consistent multi-view videos instead of images and then reconstruct the 4D object.

Video Generation. Recent advancements in video generation models [3, 4, 16–18, 51, 60] have shown very impressive performance with consistent geometry and realistic motions. Video generation models demonstrate strong generalization capabilities due to their training on large-scale image and video datasets, which are more readily available than 3D or 4D data. In this work, we adapt the pre-trained video generation model to generate multi-view videos for 4D generation.

4D Generation. Recent works generate 4D content by separately utilizing pre-trained video generative model and multi-view generative model, either through SDS-based optimization [1, 6, 13, 22, 29, 42, 43, 52, 73, 74, 76, 79] or some inference-only pipelines [37, 54, 55, 70]. However, SDS-

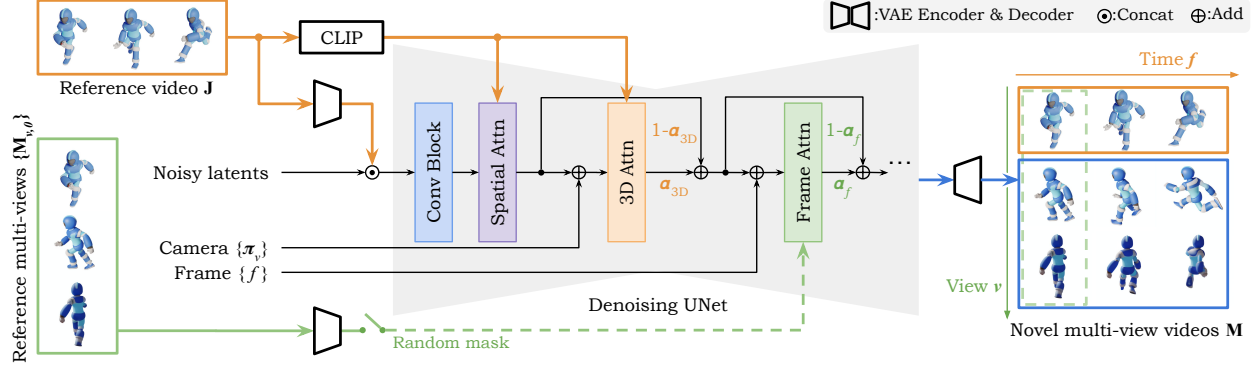


Figure 2. **Stable Video 4D 2.0 (SV4D 2.0) network architecture.** Our model is similar to SV4D [69] with several key differences: 1) we randomly mask the reference multi-view latents for cross-attention conditioning, allowing the model to generate multi-view videos without dependency of a separate multi-view diffusion model; 2) we replace view attention with 3D attention and condition it on the camera poses relative to the input view, making it more robust to arbitrary and sparse novel views; 3) we design an α -blending strategy for both 3D and frame attention layers to merge spatial and temporal information effectively while preserving priors from multi-view and video diffusion models, which also enables joint training on 3D and 4D data.

based methods tend to take hours to generate 4D content, while inference-only pipelines struggle to maintain satisfactory spatio-temporal consistency. A different line of approaches [2, 21, 24, 27, 44, 62, 68–70, 77, 80, 82] train a unified 4D generative model using synthetic 4D data. For instance, SV4D [69] incorporates view and frame attention layers to maintain spatio-temporal consistency. However, these works often produce blurry details and poor spatio-temporal consistency, particularly when handling large object motion and self-occlusion. Building upon the SV4D framework, we make several key improvements to the network architecture, training strategy, data curation, and 4D optimization. These enhance the output quality in terms of detail preservation and spatio-temporal consistency, while better handling self-occlusion, large motion, and longer video generation.

3. Multi-view Video Synthesis via SV4D 2.0

SV4D 2.0 generates dynamic 3D content by first synthesizing multiple novel-view videos given a monocular video as input then optimizing a 4D representation via photogrammetry. Similar to SV4D [69], we leverage the temporal prior in a video diffusion model as well as the 3D prior in a multi-view diffusion model, while introducing key improvements in both multi-view video synthesis and 4D optimization.

3.1. Problem Setting

Given a monocular video $\mathbf{J} \in \mathbb{R}^{F \times D}$ of a dynamic object as input, where F is the number of frames and $D := 3 \times H \times W$ is the merged image dimension, our goal is to generate novel multi-view videos $\mathbf{M} \in \mathbb{R}^{V \times F \times D}$ at V views that are consistent spatially (i.e. V 's dimension) and temporally (i.e. F 's dimension). The generated novel views must follow a conditioning camera trajectory $\boldsymbol{\pi} \in \mathbb{R}^{V \times 2}$ specified simply by a sequence of elevation e and azimuth angles a relative to the input view of the monocular video. Prior work SV4D [69]

generates novel multi-view videos \mathbf{M} , by learning the conditional distribution $p(\mathbf{M} | \{\mathbf{M}_{v,0}\}, \mathbf{J}, \boldsymbol{\pi})$ parameterized by a 4D diffusion model, where $\{\mathbf{M}_{v,0}\}$ denotes the multi-view images of the first frame obtained using SV3D [61]. In contrast, SV4D 2.0 directly learns $p(\mathbf{M} | \mathbf{J}, \boldsymbol{\pi})$ without additionally conditioning on the output of SV3D.

3.2. Network Architecture

The model architecture of SV4D 2.0, as shown in Fig. 2, consists of a denoising UNet, where each layer consists of a residual block with Conv3D layers, and three transformer blocks with spatial, 3D, and frame attention respectively. To leverage the rich motion and multi-view priors learned from large-scale video and 3D datasets, the weights of the frame attention layers are initialized from the corresponding layers in SVD [3], and the rest from SV3D_p [61]. While our architecture shares similarities with those of SV3D [61] and SV4D [69], we make important changes in SV4D 2.0 that lead to enhanced NVVS quality.

3D Attention for Sparse-view Synthesis. To better handle correspondence between sparse novel views i.e. views with significant distance between them, we utilize 3D attention layer [50, 63], while SV3D [61] and SV4D [69] use view attention layers. Let \mathbf{L} denote the latent features before being fed to each attention layer $\gamma_{(\cdot)}(\cdot)$. The view attention module γ_v in [61, 69] learns 3D consistency by attending to the same spatial location across views by first reshaping the latents as:

$$\gamma_v \left(\text{reshape} \left(\mathbf{L}, F \times V \times H \times W \times C \rightarrow (FHW) \times \underline{V} \times C \right) \right), \quad (1)$$

where C is the number of latent channels. However, the same spatial location in a neighboring view need not correspond to the same 3D region, especially when generating sparse novel views. Thus, we adopt 3D attention γ_{3D} by attending to merged view and spatial dimensions as:

$$\gamma_{3D} \left(\text{reshape}(\mathbf{L}, F V H W C \rightarrow F \underline{(V H W)} C) \right). \quad (2)$$

We show in our ablation study (Tab. 1) that this considerably improves 3D consistency across sparse novel views, which also allows us to generate longer videos under the same compute/memory constraints.

Blending 3D and Temporal Information. To combine the spatial and temporal information with flexible control, we use an α -blending mechanism. Specifically, we introduce learnable blending weights α_{3D} and α_f for the 3D and frame attention block respectively, to merge the attended latents with their skip-connected versions. Compared to SV4D [69] which uses the same blending weights after multi-view and frame attention (*i.e.*, $\alpha_{3D} = \alpha_f$), this design can more effectively preserve the 3D prior in a multi-view diffusion model and temporal prior in a video diffusion model, while learning the joint spatio-temporal consistency from 4D data. Furthermore, by setting $\alpha_f = 0$ to bypass the frame attention layers, we can train the model with more abundant 3D data for better generalization. In 4D training, we initialize α_f to a near-zero value to speed up convergence by allowing the model to maximally reuse the priors in 3D pre-training.

Camera and Frame Index Conditioning. To condition on the camera and frame index, we add the learned embeddings of camera trajectory $\{\pi_v\}$ and frames indices $\{f\}$ to the latents before 3D and frame attention layers, respectively. In contrast to SV4D [69] which uses view indices $\{v\}$ to condition view attention, camera conditioning allows the model to flexibly generate novel views following an arbitrary camera trajectory. Note that our conditioning camera trajectory is relative to the input view, whereas SV4D takes in world-space camera poses and thus requires additional pose estimation of the input video during inference.

Optional Conditioning of Reference Multi-views. We additionally make the dependency on reference multi-views $\mathbf{M}_{v,0}$ optional, by randomly masking their latents for cross-attention conditioning during training. Although the reference multi-views provide rich 3D guidance for NVVS, it requires either ground-truth 3D renders or synthesized novel views via a multi-view diffusion model like SV3D [61] during inference. We remove this requirement to make inference more practically feasible. More importantly, it encourages the model to focus on temporal correspondence and motion information in all input frames (instead of the first frame only), allowing NVVS that is consistent with the input video.

Note that many existing 4D generation methods [44, 54, 62, 69, 70, 73, 74] generate such reference multi-view images conditioned on a single anchor frame (typically frame 0), based on the underlying assumption that the appearance of a 3D object from novel views (*i.e.*, $\{\mathbf{M}_{v,0}\}$) is conditionally independent of its appearance in motion from a single view (*i.e.*, $\{\mathbf{M}_{0,f}\}$). We argue that this assumption is false for general articulated objects, and demonstrate that our

model is more robust to self-occlusions in a single frame as we generate the spatio-temporal images jointly. We also show several examples in the supplemental material where the reference multi-views conflict with the input videos and cause blurry artifacts.

Inference Sampling. During inference, we sample multi-view videos of $V = 4$ novel views and $F = 12$ frames from the learned distribution $p(\mathbf{M} | \mathbf{J}, \boldsymbol{\pi})$ without the reference multi-view condition of $\mathbf{M}_{v,0}$. To extend the videos while maintaining temporal consistency, we use the generated multi-views of the last generated frame as multi-view conditioning and sample longer videos from the distribution $p(\mathbf{M} | \mathbf{M}_{v,0}, \mathbf{J}, \boldsymbol{\pi})$. Note that this flexible inference sampling is only enabled by the proposed random masking during training. We linearly increase the scale of classifier-free guidance (CFG) in both view and time axes when deviating from the input view and first timestamp.

3.3. Data Curation

We curate a 4D dataset, ObjaverseDy++, to train our model, which consists of renders of dynamic objects from the Objaverse [10] and ObjaverseXL datasets [9]. It is a quality-enhanced version of the ObjaverseDy dataset in [69]. Particularly, we prevent objects from moving off-center to facilitate the learning of temporal consistency. To this end, we propose to disentangle the global and local motion of each dynamic 3D object, by 1) calculating the mean temporal offsets of each surface point, 2) identifying the static regions via simple thresholding, 3) computing the average offset within a most static 3D bounding box, then 4) subtracting it globally. We observe that such global motion calibration before rendering 4D objects leads to much higher spatio-temporal consistency in the multi-view video synthesis. Moreover, we filter out objects with small motion or inconsistent scaling and reduce baked-in lighting effect that sometimes causes dark back views. Lastly, we render denser spatial views and longer animation sequences for each object. More data curation details are shown in supplemental material.

3.4. Progressive 3D-to-4D Training

While most 4D generation methods [24, 27, 44, 70] directly finetune a modified diffusion model on 4D data, we argue that it is non-optimal to adapt to the new network architecture and bridge the 3D-4D or video-4D domain gap. Instead, we train the SV4D 2.0 model in a progressive 3D-to-4D manner, facilitating 4D learning while preserving rich 3D priors. Concretely, we first train the model on the Objaverse dataset (static 3D) in [61] while bypassing the frame attention layers by setting $\alpha_f = 0$. Then, we unfreeze the frame attention layers as well as α_f and finetune the model on the ObjaverseDy++ dataset to boost the learning of object motion. Although our model architecture allows joint 3D and 4D training from scratch, we show in our ablation study

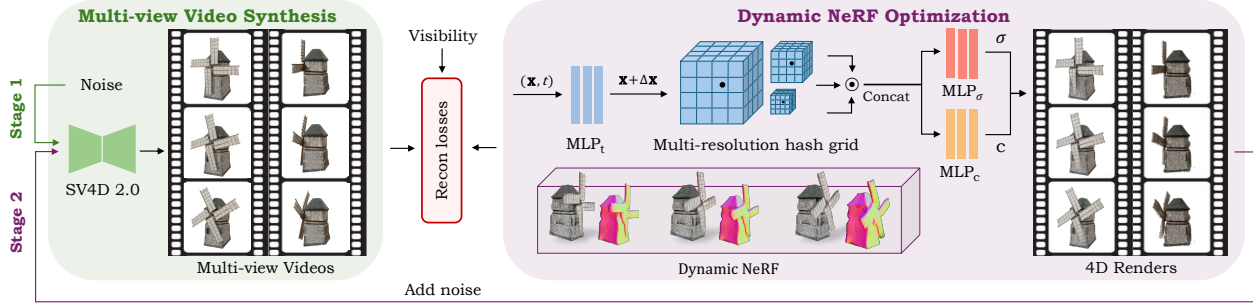


Figure 3. **4D optimization overview.** In the first stage, we use the initial synthesized multi-view videos as pseudo ground-truths to optimize a dynamic NeRF. To handle the 3D inconsistency and pose misalignment in novel view synthesis, we propose a second-stage refinement by noising and denoising the renders of dynamic NeRF as enhanced (3D consistent) photogrammetry targets. We also propose a visibility weighting scheme for the reconstruction losses to mitigate inconsistent texture across views.

(see supplemental material) that this progressive learning reaches the best tradeoff between training efficiency and model performance considering the imbalanced size of 3D and 4D datasets. The training objective of SV4D 2.0 can be expressed as: $\mathbb{E}_{z_n, n, \mathbf{J}, \pi, \epsilon} = \left[\|\epsilon - \epsilon_\theta(z_n, n, \mathbf{J}, \pi)\|_2^2 \right]$, where ϵ_θ is the denoising network and z_n is the noisy latents at denoising timestep n with added noise ϵ .

4. 4D Optimization from Multi-view Videos

4D Representation and Photogrammetry-based Optimization. As shown in Fig. 3, by using the multi-view videos generated by SV4D 2.0 as pseudo ground-truths, we adopt photogrammetry-based optimization to learn a dynamic NeRF representation for each 4D asset. Formally, the implicit NeRF representation can be expressed as $\Psi : (\mathbf{x}, \mathbf{t}) \mapsto (\sigma, \mathbf{c})$, mapping each sampled 3D point $\mathbf{x} = (x, y, z)$ and time embedding \mathbf{t} to the corresponding density $\sigma \in \mathbb{R}_+$ and color $\mathbf{c} \in \mathbb{R}_+^3$, which can then be used to render images/videos by densely sampling points along camera rays and timestamps. Similar to D-NeRF [41] and SV4D [69], we compose a 4D object with a canonical NeRF and a temporal deformation field, where the former is parametrized by a multi-resolution hash grid [35] with density and color MLPs (MLP_σ and MLP_c), and the later by a deformation MLP (MLP_t).

Instead of relying on the cumbersome SDS losses from image/video diffusion models, we follow SV4D [69] to optimize a dynamic NeRF via photogrammetry. The main reconstruction losses include the pixel-level mean squared error $\mathcal{L}_{\text{mse}} = \|\mathbf{M} - \hat{\mathbf{M}}\|^2$, perceptual LPIPS [78] loss $\mathcal{L}_{\text{lipis}}$, 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 silhouettes. We learn the 4D representation in a static-to-dynamic manner to improve training efficiency and stability. That is, we first optimize the canonical NeRF only on the multi-view images of the first frame, then unfreeze MLP_t while sampling multiple views and frames for joint optimization.

Improving 3D Consistency and Pose Alignment. Although SV4D [69] and other prior methods [27, 54, 70] demonstrate that the such photogrammetry-based optimization can efficiently create a 4D object without SDS losses, we observe that the slight 3D inconsistency or pose misalignment (generated view not aligned with input camera pose) in multi-view video synthesis can greatly affect the 4D output quality, resulting in undesirable artifacts or blurry details. To address this, we propose a 2-stage optimization framework and weighted reconstruction loss based on soft visibility estimation, as shown in Fig. 3.

In the first stage, we sample multi-view videos from noise as pseudo ground-truths and optimize a dynamic NeRF via the aforementioned reconstruction losses. In the second stage, we refine the pseudo ground-truth videos by noising and denoising the rendered views and frames of the stage-1 dynamic NeRF. Since these videos are rendered and denoised from an optimized 4D representation, they are by nature more consistent and pose-aligned in 3D compared to the initial video synthesis from noise, thus providing a better reconstruction target. Note that such iterative refinement can be applied multiple times [75] and at each denoising step [7, 57] with a tradeoff for optimization speed. In practice, we find that adding a second-stage optimization already achieves noticeably better output quality in terms of 3D coherence and detail sharpness, as shown in Fig. 4 (a).

To better deal with the inconsistent texture across views, we also propose a loss weighting scheme by treating visibility information as confidence for each generated view. Concretely, for each foreground pixel \mathbf{p} in view v , we estimate its soft visibility \mathbf{W}_p based on the dot product between camera ray direction \mathbf{v} (calculated from camera pose π_v) and surface normal \mathbf{n} (estimated from Omnidata [11]) as: $\mathbf{W}_p = \mathbf{v} \cdot \mathbf{n}$. Since higher values of \mathbf{W}_p indicate better visibility of the surface from the camera, we clip the range of \mathbf{W} to $(0, 1)$ and use it as confidence weights for the MSE loss: $\mathcal{L}_{\text{mse}} = \|\mathbf{W}(\mathbf{M} - \hat{\mathbf{M}})\|^2$. We show examples of the visibility maps in Fig. 4 (b).

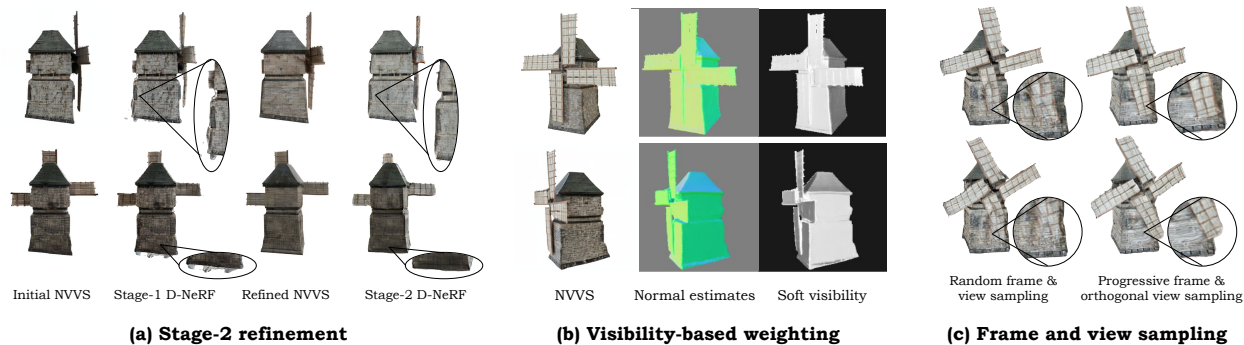


Figure 4. **Detailed analyses of our 4D optimization strategies.** (a) Our Stage-2 refinement can effectively reduce the artifacts in dynamic NeRF caused by inconsistent novel-view video synthesis (NVVS). (b) We compute soft visibility maps as view-dependent loss weights based on surface normal estimates to further mitigate texture inconsistency. (c) The proposed progressive frame and orthogonal view sampling are shown to facilitate the learning of temporal deformation and capture better details in motion.

View and Frame Sampling for Stable 4D Optimization.

Due to memory limitation, SV4D [69] randomly samples 4 views and 4 frames during 4D optimization. To further stabilize training, we propose orthogonal view sampling and progressive frame sampling to stabilize training. Specifically, we sample near-orthogonal views (*e.g.*, azimuth degree $a = 0, 90, 180, 270$) to minimize free deformation in unseen regions while progressively increasing the temporal window (range of frame indices) for frame sampling. We show in Fig. 4 (c) that this sampling strategy effectively improves the learning of temporal deformation field in presence of large motion.

5. Experiments

Datasets. We evaluate the results of NVVS and 4D optimization on two synthetic datasets: ObjaverseDy [10, 69] and Consistent4D [10]. For ObjaverseDy, we use the same validation objects as [69] and render 21 frames \times 8 views per object for evaluation. We run all baseline models on our validation set using their official code, conditioned on a near-front view video for each object. Consistent4D dataset contains multi-view renders of several dynamic objects from Objaverse [10]. We used the same input video and evaluate on the same 4 novel views as [22, 69]. Additionally, we use monocular videos from the real-world video dataset DAVIS [5, 38, 39] for the visual comparison and user study.

Metrics. We evaluate the visual quality of our video and 4D outputs using common image metrics like LPIPS [78], CLIP-S, PSNR, SSIM, and MSE. Following [69], we also evaluate spatio-temporal consistency by calculating FVD [59] in different ways: *FVD-F*: over frames at a fixed view. *FVD-V*: over views at a fixed frame. *FVD-Diag*: over the diagonal images of the image matrix. *FV4D*: over all images by scanning them in a bidirectional raster (zig-zag) order.

Baselines. For NVVS, we compare SV4D 2.0 with several recent methods capable of generating multiple novel-view

videos from a monocular video, including SV4D [69], per-frame SV3D [61], Diffusion² [70], 4Diffusion [77]. For *4D generation*, we compare SV4D 2.0 with other baselines that can generate 4D representations, including SV4D [69], Consistent4D [22], STAG4D [76], 4Diffusion [77], Dream-Gaussian4D (DG4D) [43], GaussianFlow [12], 4DGen [73], Efficient4D [37], L4GM [44]. We omit other prior methods since they either cannot perform video-to-4D generation [13, 21, 42] or do not release the code or model weights at the time of submission [2, 13, 24, 27, 55, 62, 68, 74, 80].

5.1. Visual Comparison

In Figs. 5 and 6, we show visual comparison against prior methods on multi-view video synthesis and 4D generation, respectively. We observe that applying SV3D frame-by-frame can preserve geometry and texture details across novel views but it fails to maintain temporal consistency. Diffusion² tends to produce over-smoothed results by sacrificing details for temporal coherence. SV4D can generate novel-view videos with decent spatio-temporal consistency, but it often produces blurry details during large motion. By comparison, SV4D 2.0 outputs are higher-quality in terms of details, sharpness and consistency across views and frames. In 4D generation (Fig. 6), STAG4D often produce artifacts caused by its Gaussian representation and inconsistent novel views. DG4D generates finer details but suffers from over-saturated texture, temporal flickering, and inaccurate geometry, due to SDS-based optimization. L4GM does not generalize well on real-world data (no video prior like ours) and struggles with videos at non-zero elevation (training data primarily at 0° elevation). SV4D improves with cleaner texture and geometry, though it fails to capture thin geometric structures or texture details. In contrast, SV4D 2.0 can generate high-fidelity 4D assets with realistic geometry, texture, and motion. We also show ablative results in Fig. 4 to justify the effectiveness of our 4D optimization strategies.

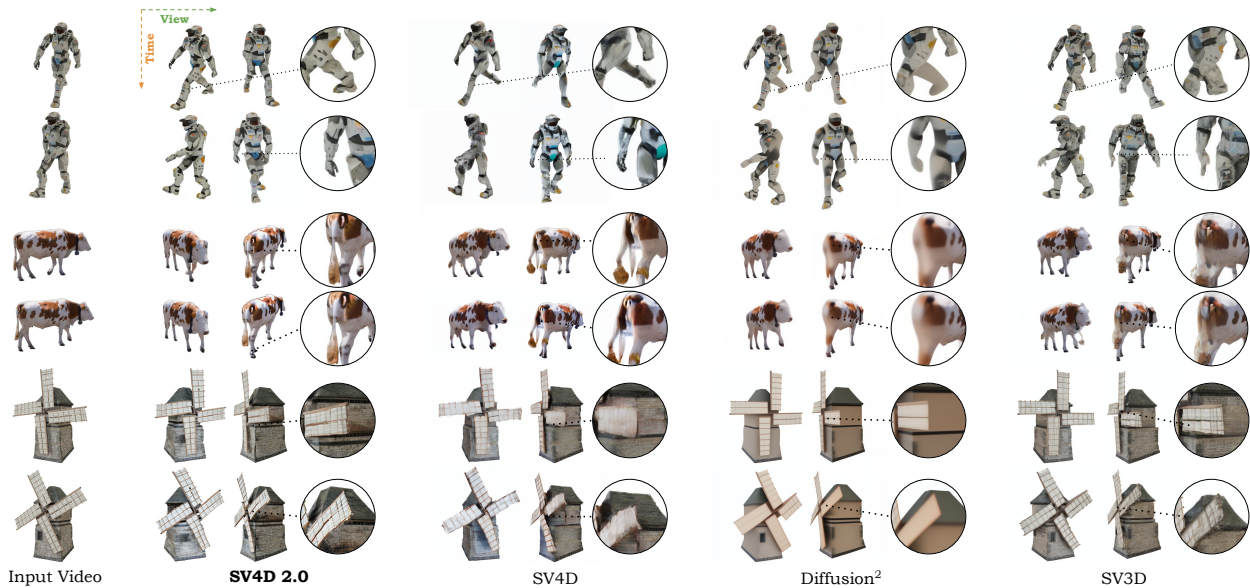


Figure 5. **Visual Comparison of Novel View Video Synthesis Results.** We show two frames in the input videos and two novel-view results of the corresponding frames. Compared to the baseline methods, SV4D 2.0 outputs contain geometry and texture details that are more faithful to the input video and consistent across frames. We also refer reviewers to the Supplemental Material for the video comparison.




Figure 6. **Visual Comparison of 4D Optimization Results.** We show two frames in the input videos and render a novel view of the corresponding frames. By leveraging the detailed and consistent multi-view videos generated by SV4D 2.0, we can optimize higher-quality 4D assets compared to prior works. In particular, our results are more detailed, consistent, and faithful to the input videos.

5.2. Quantitative Comparison

For NVVS, we show the quantitative comparison on the ObjaverseDy and Consistent4D datasets in Tabs. 1 and 3, respectively. The results on both datasets show a consistent performance gain by SV4D 2.0 over all baselines on all metrics. In particular, the pixel-level metrics (PSNR,

SSIM, MSE) demonstrate that SV4D 2.0 outputs have higher image-quality and align better with the ground-truths. The perceptual-level metrics (LPIPS, CLIP-S) show that our results are semantically more faithful to the objects. More importantly, SV4D 2.0 achieves much lower FVD metrics across frames, views, diagonal images, and full image ma-

Table 1. **Evaluation of NVVS on the ObjaverseDy dataset.** SV4D 2.0 achieves better visual quality (*LPIPS*, *CLIP-S*, *PSNR*, *SSIM*, *MSE*), video frame consistency (*FVD-F*), multi-view consistency (*FVD-V*), and spatio-temporal consistency (*FVD-Diag* and *FV4D*).

Model	LPIPS↓	CLIP-S↑	PSNR↑	SSIM↑	MSE↓	FVD-F↓	FVD-V↓	FVD-Diag↓	FV4D↓
SV3D [61]	0.131	0.922	17.60	0.873	0.0244	799.73	397.37	531.85	754.09
Diffusion ² [70]	0.133	0.892	18.08	0.886	0.0225	987.19	477.66	789.12	1020.15
SV4D [69]	0.122	0.919	17.83	0.881	0.0217	647.97	427.97	564.24	458.34
SV4D 2.0 w/o 3D attention	0.115	0.934	17.89	0.878	0.0214	548.14	368.68	463.59	283.10
SV4D 2.0 w/o improved data	0.112	0.931	18.19	0.884	0.0192	576.66	344.68	467.57	316.91
SV4D 2.0 (Ours)	0.105	0.939	18.67	0.886	0.0180	473.90	272.85	426.92	256.84

Table 2. **Evaluation of 4D outputs on the ObjaverseDy dataset.** SV4D 2.0 consistently outperforms baselines in all metrics.

Model	LPIPS↓	CLIP-S↑	PSNR↑	SSIM↑	MSE↓	FVD-F↓	FVD-V↓	FVD-Diag↓	FV4D↓
Consistent4D [22]	0.148	0.899	16.44	0.866	0.0297	781.38	510.04	782.79	658.31
STAG4D [76]	0.155	0.868	16.73	0.867	0.0260	848.83	539.96	709.52	833.08
DG4D [43]	0.156	0.874	16.02	0.860	0.0293	826.72	543.29	761.58	741.99
L4GM [44]	0.146	0.902	17.65	0.877	0.0227	805.87	537.46	782.92	666.48
SV4D [69]	0.122	0.902	18.47	0.884	0.0189	754.23	436.86	666.59	699.04
SV4D 2.0 (Ours)	0.107	0.916	19.37	0.890	0.0153	654.64	375.03	513.26	532.55

Table 3. **Evaluation of NVVS on Consistent4D.** SV4D 2.0 achieves a consistent performance gain on image and video metrics.

Model	LPIPS↓	CLIP-S↑	FVD-F↓
SV3D [61]	0.129	0.925	989.53
4Diffusion [77]	0.164	0.863	-
Diffusion ² [70]	0.189	0.907	1205.16
SV4D [69]	0.129	0.929	677.68
SV4D 2.0 (Ours)	0.115	0.946	553.43

trix, justifying its superior spatio-temporal consistency. We also show several ablative results (w/o 3D attention, w/o improved data curation in ObjaverseDy++) in Tab. 1 to validate the effectiveness of our model design. For *4D generation*, we evaluate the rendered videos on the ObjaverseDy and Consistent4D datasets in Tabs. 2 and 4. The 4D results also show that SV4D 2.0 consistently achieves higher quality in terms of both image and video quality. L4GM [44] performs comparably to our method on Consistent4D but struggles with videos at non-zero elevation in ObjaverseDy. We show more results and ablations in the supplemental material.

5.3. User Study

To further evaluate our model performance on in-the-wild videos, we also conduct user study to compare the NVVS results of SV4D 2.0 against 3 prior methods. Specifically, we select 20 real-world videos from the DAVIS dataset [5] with relatively steady camera and minimal truncation/occlusions. Following [69], we randomly choose a novel camera view for each input video and ask users to select a novel-view video (generated by 4 different methods) that “looks more stable, realistic, and closely resembles the reference subject”. As a result, SV4D 2.0 outputs are preferred 85.2% over Diffusion² (1.8%), STAG4D (5.1%), and SV4D (7.9%)

Table 4. **Evaluation of 4D outputs on the Consistent4D dataset.** SV4D 2.0 achieves state-of-the-art visual quality (LPIPS, CLIP-S) and temporal smoothness (FVD-F) compared to prior methods.

Model	LPIPS↓	CLIP-S↑	FVD-F↓
Consistent4D [22]	0.160	0.87	1133.93
STAG4D [76]	0.126	0.91	992.21
4Diffusion [77]	0.165	0.88	-
Efficient4D [37]	0.130	0.92	-
4DGen [73]	0.140	0.89	-
DG4D [43]	0.160	0.87	-
GaussianFlow [12]	0.140	0.91	-
L4GM [44]	0.120	0.94	691.87
SV4D [69]	0.118	0.92	732.40
SV4D 2.0 (Ours)	0.117	0.93	692.68

among all participants, showing much better generalization to these challenging real-world videos.

6. Conclusion

We present SV4D 2.0, an enhanced multi-view video diffusion model that can generate high-quality novel-view videos and 4D assets given a monocular video. Compared to our baseline [69], SV4D 2.0 can produce longer videos, from sparser camera views, with better spatio-temporal consistency, detail sharpness, and generalizability to real-world videos. Moreover, it does not require reference multi-views from a separate model during inference, making it more practical and robust to occlusions in anchor frame. We achieve these by introducing novel techniques in data curation, network architecture, progressive training, and 4D optimization. Our extensive evaluation and user study show that SV4D 2.0 outperforms prior methods in both NVVS and 4D generation, making it a favorable 4D foundation model.

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 *CVPR*, 2024. 2
- [2] Jianhong Bai, Menghan Xia, Xintao Wang, Ziyang Yuan, Xiao Fu, Zuozhu Liu, Haoji Hu, Pengfei Wan, and Di Zhang. SynCamMaster: Synchronizing multi-camera video generation from diverse viewpoints. *arXiv preprint arXiv:2412.07760*, 2024. 3, 6
- [3] 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. 1, 2, 3
- [4] 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 *CVPR*, 2023. 1, 2
- [5] Sergi Caelles, Jordi Pont-Tuset, Federico Perazzi, Alberto Montes, Kevis-Kokitsi Maninis, and Luc Van Gool. The 2019 davis challenge on vos: Unsupervised multi-object segmentation. *arXiv preprint arXiv:1905.00737*, 2019. 2, 6, 8
- [6] Ce Chen, Shaoli Huang, Xuelin Chen, Guangyi Chen, Xiaoguang Han, Kun Zhang, and Mingming Gong. CT4D: Consistent text-to-4D generation with animatable meshes. *arXiv preprint arXiv:2408.08342*, 2024. 2
- [7] Hansheng Chen, Bokui Shen, Yulin Liu, Ruoxi Shi, Linqi Zhou, Connor Z Lin, Jiayuan Gu, Hao Su, Gordon Wetzstein, and Leonidas Guibas. 3D-Adapter: Geometry-consistent multi-view diffusion for high-quality 3D generation. *arXiv preprint arXiv:2410.18974*, 2024. 5
- [8] Zilong Chen, Feng Wang, Yikai Wang, and Huaping Liu. Text-to-3D using gaussian splatting. In *CVPR*, 2024. 2
- [9] 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. In *NeurIPS*, 2023. 2, 4
- [10] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3D objects. In *CVPR*, 2023. 2, 4, 6
- [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 *ICCV*, 2021. 5
- [12] Quankai Gao, Qiangeng Xu, Zhe Cao, Ben Mildenhall, Wen-chao Ma, Le Chen, Danhang Tang, and Ulrich Neumann. GaussianFlow: Splatting gaussian dynamics for 4D content creation. *arXiv preprint arXiv:2403.12365*, 2024. 6, 8
- [13] Chen Geng, Yunzhi Zhang, Shangzhe Wu, and Jiajun Wu. Birth and death of a rose. *arXiv preprint arXiv:2412.05278*, 2024. 2, 6
- [14] Rohit Girdhar, Mannat Singh, Andrew Brown, Quentin Duval, Samaneh Azadi, Sai Saketh Rambhatla, Akbar Shah, Xi Yin, Devi Parikh, and Ishan Misra. Factorizing text-to-video generation by explicit image conditioning. In *ECCV*, 2024. 1
- [15] 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. 2
- [16] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yao-hui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. AnimateDiff: Animate your personalized text-to-image diffusion models without specific tuning. In *ICLR*, 2024. 2
- [17] Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. Latent video diffusion models for high-fidelity long video generation. *arXiv*, 2022.
- [18] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. In *NeurIPS*, 2022. 2
- [19] 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. In *ICLR*, 2024. 2
- [20] Hanwen Jiang, Zhenyu Jiang, Yue Zhao, and Qixing Huang. LEAP: Liberate sparse-view 3D modeling from camera poses. In *ICLR*, 2024. 2
- [21] Yanqin Jiang, Chaohui Yu, Chenjie Cao, Fan Wang, Weiming Hu, and Jin Gao. Animate3D: Animating any 3D model with multi-view video diffusion. In *NeurIPS*, 2024. 3, 6
- [22] Yanqin Jiang, Li Zhang, Jin Gao, Weiming Hu, and Yao Yao. Consistent4D: Consistent 360° dynamic object generation from monocular video. In *ICLR*, 2024. 2, 6, 8
- [23] Animesh Karnwar, Niloy J Mitra, Andrea Vedaldi, and David Novotny. HoloFusion: Towards photo-realistic 3D generative modeling. In *ICCV*, 2023. 2
- [24] Bing Li, Cheng Zheng, Wenxuan Zhu, Jinjie Mai, Biao Zhang, Peter Wonka, and Bernard Ghanem. Vivid-ZOO: Multi-view video generation with diffusion model. *NeurIPS*, 2024. 2, 3, 4, 6
- [25] 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. In *ICLR*, 2024. 2
- [26] Weiyu Li, Rui Chen, Xuelin Chen, and Ping Tan. SweetDreamer: Aligning geometric priors in 2D diffusion for consistent text-to-3D. In *ICLR*, 2024. 2
- [27] Hanwen Liang, Yuyang Yin, Dejia Xu, hanxue liang, Zhangyang Wang, Konstantinos N Plataniotis, Yao Zhao, and Yunchao Wei. Diffusion4D: Fast spatial-temporal consistent 4D generation via video diffusion models. In *NeurIPS*, 2024. 2, 3, 4, 5, 6
- [28] 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. In *CVPR*, 2024. 2
- [29] 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. In *CVPR*, 2024. 2

- [30] 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. In *NeurIPS*, 2023. 2
- [31] 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. In *CVPR*, 2024.
- [32] Ruoshi Liu, Rund Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3D object. In *ICCV*, 2023.
- [33] 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. In *ICLR*, 2024.
- [34] 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. In *CVPR*, 2024. 2
- [35] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 2022. 5
- [36] Zijie Pan, Jiachen Lu, Xiatian Zhu, and Li Zhang. Enhancing high-resolution 3D generation through pixel-wise gradient clipping. In *ICLR*, 2024. 2
- [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. 2, 6, 8
- [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 *CVPR*, 2016. 2, 6
- [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. 2, 6
- [40] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D diffusion. In *ICLR*, 2023. 2
- [41] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural radiance fields for dynamic scenes. In *CVPR*, 2021. 5
- [42] Ohad Rahamim, Ori Malca, Dvir Samuel, and Gal Chechik. Bringing objects to life: 4D generation from 3D objects. *arXiv preprint arXiv:2412.20422*, 2024. 2, 6
- [43] 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. 2, 6, 8
- [44] Jiawei Ren, Cheng Xie, Ashkan Mirzaei, Karsten Kreis, Ziwei Liu, Antonio Torralba, Sanja Fidler, Seung Wook Kim, Huan Ling, et al. L4GM: Large 4D gaussian reconstruction model. In *NeurIPS*, 2024. 2, 3, 4, 6, 8
- [45] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 1
- [46] 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. In *CVPR*, 2023. 1
- [47] 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 image. In *CVPR*, 2024. 2
- [48] 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
- [49] Yukai Shi, Jianan Wang, CAO He, 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. In *ICLR*, 2024. 2
- [50] Yichun Shi, Peng Wang, Jianglong Ye, Long Mai, Kejie Li, and Xiao Yang. MVDream: Multi-view diffusion for 3D generation. In *ICLR*, 2024. 2, 3
- [51] 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. In *ICLR*, 2023. 2
- [52] Uriel Singer, Shelly Sheynin, Adam Polyak, Oron Ashual, Iurii Makarov, Filippos Kokkinos, Naman Goyal, Andrea Vedaldi, Devi Parikh, Justin Johnson, et al. Text-to-4D dynamic scene generation. In *ICML*, 2023. 2
- [53] Jingxiang Sun, Bo Zhang, Ruizhi Shao, Lizhen Wang, Wen Liu, Zhenda Xie, and Yebin Liu. DreamCraft3D: Hierarchical 3D generation with bootstrapped diffusion prior. In *ICLR*, 2024. 2
- [54] Qi Sun, Zhiyang Guo, Ziyu Wan, Jing Nathan Yan, Shengming Yin, Wengang Zhou, Jing Liao, and Houqiang Li. EG4D: Explicit generation of 4D object without score distillation. In *ICLR*, 2025. 2, 4, 5
- [55] Wenqiang Sun, Shuo Chen, Fangfu Liu, Zilong Chen, Yueqi Duan, Jun Zhang, and Yikai Wang. DimensionX: Create any 3D and 4D scenes from a single image with controllable video diffusion. *arXiv preprint arXiv:2411.04928*, 2024. 2, 6
- [56] Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. DreamGaussian: Generative gaussian splatting for efficient 3D content creation. In *ICLR*, 2024. 2
- [57] Zhenyu Tang, Junwu Zhang, Xinhua Cheng, Wangbo Yu, Chaoran Feng, Yatian Pang, Bin Lin, and Li Yuan. Cycle3D: High-quality and consistent image-to-3D generation via generation-reconstruction cycle. *arXiv preprint arXiv:2407.19548*, 2024. 5
- [58] 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. 2
- [59] 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. 6

- [60] Vikram Voleti, Alexia Jolicoeur-Martineau, and Christopher Pal. MCVD: Masked conditional video diffusion for prediction, generation, and interpolation. In *NeurIPS*, 2022. 1, 2
- [61] 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 *ECCV*, 2024. 2, 3, 4, 6, 8
- [62] Chaoyang Wang, Peiye Zhuang, Tuan Duc Ngo, Willi Menapace, Aliaksandr Siarohin, Michael Vasilkovsky, Ivan Skokhodov, Sergey Tulyakov, Peter Wonka, and Hsin-Ying Lee. 4Real-Video: Learning generalizable photo-realistic 4D video diffusion. *arXiv preprint arXiv:2412.04462*, 2024. 3, 4, 6
- [63] Peng Wang and Yichun Shi. ImageDream: Image-prompt multi-view diffusion for 3D generation. *arXiv preprint arXiv:2312.02201*, 2023. 2, 3
- [64] 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. In *ICLR*, 2024. 2
- [65] 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 *NeurIPS*, 2024. 2
- [66] 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. 2
- [67] 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. 2
- [68] Rundi Wu, Ruiqi Gao, Ben Poole, Alex Trevithick, Changxi Zheng, Jonathan T Barron, and Aleksander Holynski. CAT4D: Create anything in 4D with multi-view video diffusion models. *arXiv preprint arXiv:2411.18613*, 2024. 3, 6
- [69] Yiming Xie, Chun-Han Yao, Vikram Voleti, Huaizu Jiang, and Varun Jampani. SV4D: Dynamic 3D content generation with multi-frame and multi-view consistency. In *ICLR*, 2025. 1, 2, 3, 4, 5, 6, 8
- [70] Zeyu Yang, Zijie Pan, Chun Gu, and Li Zhang. Diffusion $\$^2\$: Dynamic 3D content generation via score composition of video and multi-view diffusion models. In *ICLR*, 2025. 2, 3, 4, 5, 6, 8$
- [71] 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
- [72] Taoran Yi, Jiemin Fang, Junjie Wang, Guanjun Wu, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Qi Tian, and Xinggang Wang. GaussianDreamer: Fast generation from text to 3D gaussians by bridging 2D and 3D diffusion models. In *CVPR*, 2024. 2
- [73] 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. 2, 4, 6, 8
- [74] Heng Yu, Chaoyang Wang, Peiye Zhuang, Willi Menapace, Aliaksandr Siarohin, Junli Cao, László Jeni, Sergey Tulyakov, and Hsin-Ying Lee. 4Real: Towards photorealistic 4D scene generation via video diffusion models. In *NeurIPS*, 2024. 2, 4, 6
- [75] Wangbo Yu, Jinbo Xing, Li Yuan, Wenbo Hu, Xiaoyu Li, Zhipeng Huang, Xiangjun Gao, Tien-Tsin Wong, Ying Shan, and Yonghong Tian. ViewCrafter: Taming video diffusion models for high-fidelity novel view synthesis. *arXiv preprint arXiv:2409.02048*, 2024. 5
- [76] 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. In *ECCV*, 2024. 2, 6, 8
- [77] Haiyu Zhang, Xinyuan Chen, Yaohui Wang, Xihui Liu, Yunhong Wang, and Yu Qiao. 4Diffusion: Multi-view video diffusion model for 4D generation. In *NeurIPS*, 2024. 2, 3, 6, 8
- [78] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5, 6
- [79] Yuyang Zhao, Zhiwen Yan, Enze Xie, Lanqing Hong, Zhenqiu Li, and Gim Hee Lee. Animate124: Animating one image to 4D dynamic scene. *arXiv preprint arXiv:2311.14603*, 2023. 2
- [80] Yuyang Zhao, Chung-Ching Lin, Kevin Lin, Zhiwen Yan, Linjie Li, Zhengyuan Yang, Jianfeng Wang, Gim Hee Lee, and Lijuan Wang. GenXD: Generating any 3D and 4D scenes. In *ICLR*, 2025. 3, 6
- [81] Linqi Zhou, Andy Shih, Chenlin Meng, and Stefano Ermon. DreamPropeller: Supercharge text-to-3D generation with parallel sampling. In *CVPR*, 2024. 2
- [82] Hanxin Zhu, Tianyu He, Xiqian Yu, Junliang Guo, Zhibo Chen, and Jiang Bian. AR4D: Autoregressive 4D generation from monocular videos. *arXiv preprint arXiv:2501.01722*, 2025. 3
- [83] 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. In *CVPR*, 2024. 2