

ViVid-1-to-3: Novel View Synthesis with Video Diffusion Models

Jeong-gi Kwak^{1,2 *†} Erqun Dong^{1,3,4 *†} Yuhe Jin¹ Hanseok Ko²
 Shweta Mahajan^{1,5} Kwang Moo Yi^{1,4}

¹ University of British Columbia ² Korea University ³ McGill University
⁴ Haiper Ltd. ⁵ Vector Institute for AI

<https://ubc-vision.github.io/vivid123/>

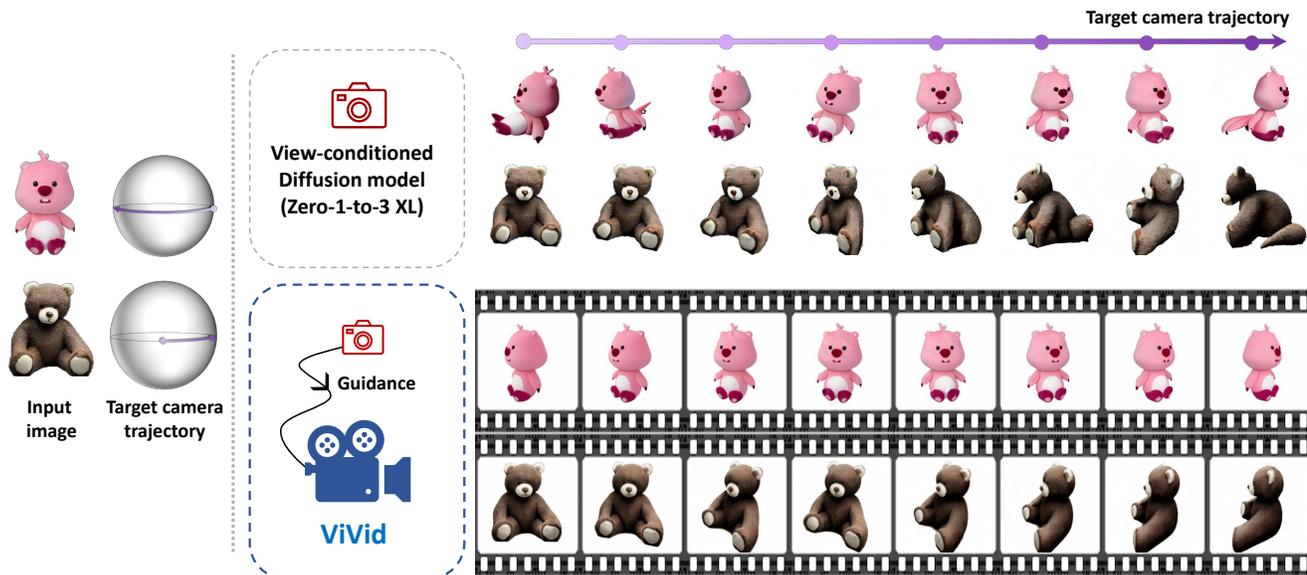


Figure 1. **Teaser** – we present a strikingly simple training-free method to make already available novel-view synthesis diffusion models more consistent both in terms of the desired viewing angle and the content—combining it with video diffusion. As shown in the example, the results of our method are more consistent with the input images and correspond more to the target views.

Abstract

Generating novel views of an object from a single image is a challenging task. It requires an understanding of the underlying 3D structure of the object from an image and rendering high-quality, spatially consistent new views. While recent methods for view synthesis based on diffusion have shown great progress, achieving consistency among various view estimates and at the same time abiding by the desired camera pose remains a critical problem yet to be solved. In this work, we demonstrate a strikingly simple method, where we utilize a pre-trained video diffusion model to solve this problem. Our key idea is that synthesizing a novel view could be reformulated as synthesizing a video of a camera going around the object of interest—a scanning video—

which then allows us to leverage the powerful priors that a video diffusion model would have learned. Thus, to perform novel-view synthesis, we create a smooth camera trajectory to the target view that we wish to render, and denoise using both a view-conditioned diffusion model and a video diffusion model. By doing so, we obtain a highly consistent novel view synthesis, outperforming the state of the art.

1. Introduction

Novel view synthesis from a single image is an interesting problem in Computer Vision as it requires an understanding of the 3D characteristics of an object or a scene, simply by looking at a 2D image. Recent methods have utilized Neural Radiance Fields (NeRF) [26, 78], and more recently diffusion models [7, 73], including those that expand upon Stable Diffusion [50] such as Zero-1-to-3 [34]. While recent

*Equal contribution

†Work done while visiting University of British Columbia

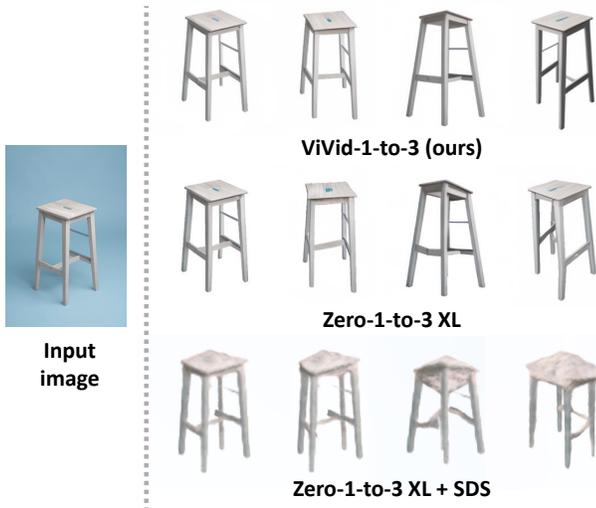


Figure 2. **Example rendering inconsistencies** – we show example novel views generated for a chair image as shown on the left. Our method provides improved consistency with the input image, compared to both Zero-1-to-3 XL [10, 34], a pure 2D novel-view diffusion model, and even when it is combined with the Score Distillation Sampling (SDS) [43] for improved 3D consistency. Note that the 3D distilled version exhibits blurriness due to the pose errors that Zero-1-to-3 XL makes.

methods show high-quality rendering from a given view, when examined more carefully, these models lack consistency [34], and their poses are also not accurate. This is because diffusion models are mainly geared towards generating images that look good. Moreover, diffusion models learn the 3D consistency implicitly, not explicitly. These limitations are not resolved even when models are trained with large datasets [10, 11]; see Fig. 1.

Very recently, researchers have thus focused their efforts on making diffusion model-based novel view synthesis more ‘consistent’. These include methods that train or finetune 2D diffusion models to make them more consistent [10, 35, 56, 57, 74, 77] or embed 3D constraints in the form of 3D representations [28, 42, 70] through score distillation techniques [43, 44, 62, 69, 72]. While they provide improved results, there are still several limitations. 2D view-conditioned models often require retraining [35, 56, 74] which is costly, and as shown in Fig. 2 (middle), even with this additional training they still have considerable room for improvement. In the case of methods that bring in explicit 3D representations such as NeRF [42] or Gaussian Splatting [28] on top [33, 44, 62], suffer from poor visual quality often due to the blurs that the pose inconsistencies of 2D models bring; see Fig. 2 (bottom).

In this work, we show that a strikingly simple solution that *does not require any new training or fine-tuning* exists for this problem—leveraging pre-trained video diffu-

sion models. Our key idea is that, given their recent improvement in quality, video diffusion models [1, 18] can be used as strong priors for any task that can be represented as a video—including the current task of novel view synthesis. For example, consider generating a video of someone scanning an object. If we were to generate such a video, where the first frame of the video corresponds to the object that we wish to generate a novel view of, and the last frame is from the target novel view, generating this video is effectively equivalent to the single image novel view synthesis task, except now with redundant intermediate frames. Importantly, we already have pre-trained public models for both generating these individual frames—view-conditioned diffusion models [34]—and the video as a whole—video diffusion models [1, 18].

To implement our method we use Zero-1-to-3 XL [10, 34] and ZeroScope [1], both of which are based on Stable Diffusion [50]. We generate a smooth camera trajectory from the frontal view to the desired view and provide Zero-1-to-3 XL with the individual camera locations for each frame, along with the input image. We then denoise a video, where the noise estimates for each frame are a combination of noise estimates from both diffusion models. We investigate multiple different strategies for combining the noise estimates, and find that starting to denoise with equal emphasis on both diffusion models and then reducing the weight on the video diffusion model to half throughout the denoising process is the best.

To evaluate our method we rely on 100 shapes from the Google Scanned Object (GSO) dataset [12]. As we are able to render these shapes from any view, we compare the novel-view synthesis results with the ground-truth rendering. We find that, due to the difficulty of the novel-view synthesis tasks, it is very easy to have *minor misalignments*, and the typical image quality metrics do not provide a complete view. We thus propose a novel metric based on optical flow, and measure the *spatial deviations* of the novel-view renderings. Our metric, together with the standard ones provides a holistic view of the performances of each method.

2. Related work

We first review novel-view synthesis methods based on explicit geometric constraints, then discuss the more recent trend of using diffusion models. As we rely on video diffusion, we also briefly discuss noteworthy works.

Novel view synthesis with geometric constraints. Early work on novel view synthesis recovers the 3D structure of a scene by incorporating geometric prior such as camera parameters [52, 60]. Since then, as in many other areas in computer vision, deep learning-based methods have been proposed [15, 48, 49], which often combine traditional 3D geometry aware multi-view synthesis [3, 17] with modern

deep learning. Other works leverage voxels [36, 59], depth maps [14, 66] or epipolar constraints [30] in their model.

More recently, since the introduction of Neural Radiance Fields (NeRF) [42], 3D constraints via volume rendering have become popular. While training a NeRF requires multiple views, PixelNeRF [78] utilizes convolutional features of pre-trained deep networks [19] to reduce the number of required views, as little as a single view. GRF [65] extends this idea to the concept of canonical space. GenVNS [7] further appends diffusion models into the pipeline for improved rendering. While they have shown impressive progress, these methods require per-dataset training, which limits their applicability.

Novel-view synthesis with diffusion models. Recently, like many other applications that involve image generation [20, 29, 41, 50], there has been a flurry of research for novel view synthesis based on diffusion models. These include those that directly aim to generate 2D novel views conditioned on the input image and the camera pose [34, 73]. Among them, Zero-1-to-3 [34], being based on Stable Diffusion [50], has demonstrated impressive results, benefiting from the original weights of Stable Diffusion that have been trained with a very large dataset [53]. This method has been further trained on a large-scale 3D dataset [10, 11], further improving its performance—we use this model.

While these direct 2D methods have shown impressive rendering quality, as shown in Fig. 2, they often fall short when it comes to the consistency of what they render, and also from where they are viewed. Various approaches, concurrently to our work [35, 56, 57, 74, 77] have thus been presented in an attempt to make them more ‘consistent’. They, however, require re-training or fine-tuning, and sometimes only provide restricted views [35]. Our method suffers from neither of these shortcomings.

Alternatively, to enforce consistency, methods that aim to *distill* what diffusion models have learned to 3D representations have also been suggested. DreamFusion [43] and Score Jacobian Chaining (SJC) [69] leverage a pre-trained diffusion model [51] to distill it into a NeRF representation [42], allowing text-to-3D, which can also be utilized for novel-view synthesis. Various followups [4, 8, 27, 32, 40, 54, 72, 75, 76] have since then been suggested to improve the rendering quality via additional constraints such as subject-driven diffusion guidance [45] and frontal camera position [32, 54].

More recent, concurrent works [33, 37, 44] have further integrated this idea with view-conditioned diffusion methods [34], but due to misalignments between desired camera pose and the outcome of 2D part of their pipelines, their results can become blurry (Fig. 2). Our method is complementary to these efforts in that we still rely on a pure 2D pipeline, yet allow more consistent renderings.

Video diffusion models. We also briefly review video dif-

fusion models since we utilize them. A recent trend in video diffusion models is to take advantage of the advancements in image diffusion models by factorizing space and time [6, 24, 25, 58, 68, 80]. They include methods that use cascaded diffusion models [24], and spatio-temporal interpolation [58]. Among them, methods that extend already trained test-to-image latent diffusion models such as Stable Diffusion [50] have become popular. For example, Blattmann et al. [6] choose to train separate temporal layers, Guo et al. [18] injects motion modules, ZeroScope [1] finetunes VideoFusion [39], which decomposes the diffusion process as a sum of a base noise and a residual noise. Among them, we use ZeroScope, as their models are publicly available, and they also build on top of Stable Diffusion, as is the case of view-synthesis diffusion model Zero-1-to-3 XL [10, 34].

3. Method

The key idea behind our method is strikingly simple—video diffusion models can be used as a strong prior in conjunction with novel-view diffusion models to improve consistency in synthesized views, *without* any training or fine-tuning. To do this, instead of directly generating a novel view, we propose to generate a camera trajectory towards the camera view that we wish to render.

As illustrated in Fig. 3, images generated by the view-conditioned diffusion model for generating novel views often do not faithfully follow the target camera path, and can also have (object) pose errors. As such renderings, when viewed as a video, would not look like a natural video due to abrupt motions, using a video diffusion model helps to smooth out this error. This results in renderings that abide by the requested novel pose, as well as the contents being consistent with the input view.

To explain our method, we first provide a brief overview of diffusion models to introduce the notations that we will use (Sec. 3.1) then detail our method (Sec. 3.2).

3.1. Preliminary: diffusion models

Specifically among the diffusion models, we rely on Latent Diffusion Models (LDM) [13, 50, 67]. Latent diffusion models utilize an encoder-decoder, $\mathcal{E}\text{-}\mathcal{D}$, pair—often pre-trained, *e.g.*, with Vector-Quantized Generative Adversarial Nets (VQ-GAN) [50]—to convert an input image \mathbf{x} into a latent code $\mathbf{z} = \mathcal{E}(\mathbf{x})$ and transform it into noise by iteratively adding Gaussian noise for T steps [50]. The diffusion model then learns to estimate the amount of noise at a given time step t , which is then used to reverse the diffusion process and denoise to the corresponding signal \mathbf{z} . Specifically, denoting the denoising model as $\epsilon(\cdot)$, to generate an image we write

$$\mathbf{z}_{t-1} = \Phi(\mathbf{z}_t, \epsilon(\mathbf{z}_t, t, y)), \quad (1)$$

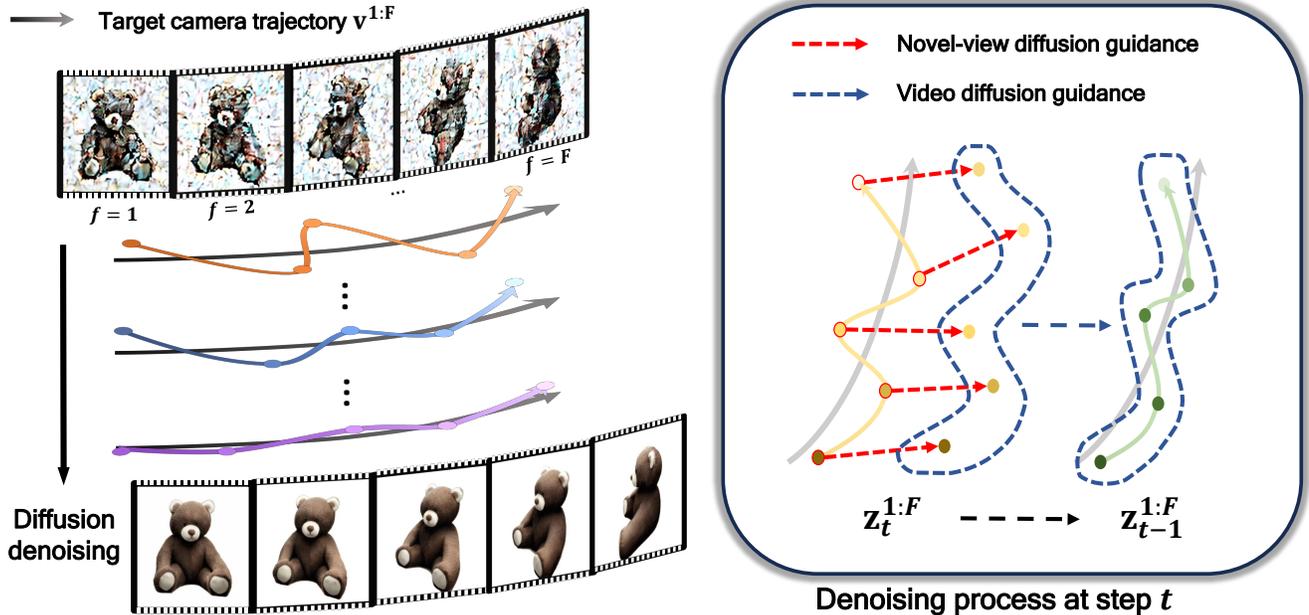


Figure 3. **Overview** – to synthesize a target novel view of a given object, we synthesize along a smooth camera trajectory, starting from an initial view and ending at the target. We then denoise to generate images with guidance from two diffusion models: a novel-view synthesis diffusion model; and a video diffusion model. The video diffusion model helps create a smooth camera trajectory and preserve consistency, which the novel-view synthesis diffusion model can lack. The two together, as shown, provide high-quality novel-view synthesis, that is consistent in both the content and the desired camera views.

where $\Phi(\cdot)$ is an update (sampling) rule for denoising such as Denoising Diffusion Probabilistic Models (DDPM) [23], and y is the conditioning text. Note that for clarity in notation, we drop text conditioning from Eq. (1).

We now formalize the two types of denoising diffusion models that we use, a model for novel-view synthesis, and one for generating videos.

Diffusion models for novel view synthesis. For novel view synthesis with diffusion models, instead of the text conditioning via y , it is often replaced with an encoding of the image and the desired camera pose [34, 57]. Thus, denoising with this model takes the form,

$$\mathbf{z}_{t-1} = \Phi(\mathbf{z}_t, \epsilon_{\text{view}}(\mathbf{z}_t, t, \mathcal{V}(\mathbf{x}^0, \mathbf{v}))), \quad (2)$$

where \mathbf{x}^0 is the input image, \mathbf{v} denotes the camera poses to synthesize, and \mathcal{V} is a mapping that encodes the input image \mathbf{x}^0 and the desired view \mathbf{v} to a conditioning signal.

Video diffusion models. While video diffusion models vary in their architectural designs and their sophisticated training schemes [1, 6, 18, 39], at a high level they diffuse *multiple* images (frames) together so that they form a video. At any timestep of the denoising process, the denoising model is applied across all the frames, given by,

$$\mathbf{z}_{t-1}^{1:F} = \Phi(\mathbf{z}_t^{1:F}, \epsilon_{\text{video}}(\mathbf{z}_t^{1:F}, t, y)), \quad (3)$$

where $\mathbf{z}^{1:F} = \{\mathbf{z}^1, \dots, \mathbf{z}^F\}$ denotes the F frames of a video at a denoising step t .

3.2. Video diffusion for novel view synthesis

To combine the two diffusion models for novel view synthesis, we generate a trajectory of camera poses for each view. views $\mathbf{v}^{1:F} = \{\mathbf{v}^1, \dots, \mathbf{v}^F\}$ where \mathbf{v}^F is the desired novel view. We generate these views by creating a smooth trajectory of views through Spherical Linear Interpolation (Slerp) [9], starting from the initial view \mathbf{v}^1 to the target view \mathbf{v}^F . For convenience, we set the initial view to be the same as \mathbf{v}^0 corresponding to the input image \mathbf{x}^0 (i.e., $\mathbf{v}^1 := \mathbf{v}^0$), but it can be any arbitrary view. We then initialize the latent for each view $\mathbf{z}_T^{1:F}$ by drawing them from a Gaussian distribution, that is, $\mathbf{z}_T^{1:F} \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$. With these, we iteratively denoise them according to the typical diffusion process with the two denoisers ϵ_{view} and ϵ_{video} together.

Specifically, simplifying the notation for the denoiser estimates in Eqs. (2) and (3) by dropping t , y , and \mathcal{V} , as $\epsilon_{\text{view}}(\mathbf{z}_t, \mathbf{x}^0, \mathbf{v})$ and $\epsilon_{\text{video}}(\mathbf{z}^{1:F})$, we write our denoising process as

$$\mathbf{z}_{t-1}^{1:F} = \Phi(\mathbf{z}_t^{1:F}, \epsilon_{\text{both}}^{1:F}), \quad (4)$$

where, denoting the view index as superscript f ,

$$\epsilon_{\text{both}}^f = \lambda_{\text{view}} \epsilon_{\text{view}}(\mathbf{z}^f, \mathbf{x}^0, \mathbf{v}^f) + \lambda_{\text{video}} \epsilon_{\text{video}}(\mathbf{z}^{1:F})^f, \quad (5)$$

where λ_{view} and λ_{video} denote the hyperparameters controlling the influence of each denoising model.

Note that combining two noise estimates together is similar to how conditional and unconditional models are used together for Classifier Free Guidance [22]. Here, we are combining two models, each with different conditioning and architecture to perform the *same task* of generating the *video* frames of the camera going about a specific trajectory. Ultimately, for the generated frames to be precise, they must satisfy both models, which is the core idea that provides improved synthesis results as shown in Fig. 3.

Null prompting. As shown in Eq. (5), $\epsilon_{\text{view}}(\cdot)$, the denoiser for the novel views is already conditioned on the input image x^0 . Because of this, there is no need for the user to provide the text prompt y for $\epsilon_{\text{video}}(\cdot)$. We thus set it to a null prompt, that is, $y = \emptyset$. Hence, the guidance of the content for the video diffusion model is purely reliant on $\epsilon_{\text{view}}(\cdot)$, the novel-view synthesis diffusion model.

Scheduling influence of each model. We found that the way the two noise estimates are combined in Eq. (5) has an impact on the behaviour of our method. Too strong emphasis—large λ_{video} —on the $\epsilon_{\text{video}}(\cdot)$ estimates in later stages of the denoising process causes the entire generation to be overly smooth, while having too small of an impact—small λ_{video} —in early stages restrict the impact of the video diffusion model, as the early stages determine the global structures [5, 21] often related to the camera view. We thus propose a linearly decaying setup, where $\lambda_{\text{view}}=1$ is set to a constant, and λ_{video} decay linearly from 1.0 at the 0-th step to 0.5 at the 50-th step. We ablate our choice in Sec. 4.3.

4. Results

4.1. Datasets and experimental setup

We will release the code and the experimental settings to make our results completely reproducible.

Datasets. To systematically evaluate the quality of the synthesized images, we rely on the 100 3D shapes from the Google Scanned Objects (GSO) dataset [12]. Among the 1,030 objects in the dataset, we manually select 100 shapes that look interesting—for example, we ignore shapes such as cubes and spheres. We render these images using the same lighting as in Zero-1-to-3 [34]. We render 25 views, with an elevation of 15 degrees and azimuth ranging from -45 to 45 degrees from manually selected views that capture the characteristics of the object.

Metrics. While results from diffusion-based novel-view synthesis models look good in general it can be misleading—synthesized images may not be of the desired view, and their contents may change as diffusion-based models often do not have the explicit notion of 3D space. Thus, with the synthetic dataset (described above), we systematically measure the quality of the generated images.

| | GT | Zero-1-to-3 XL | ViVid-1-to-3 (Ours) |
|---|----|---|---|
|  | |  |  |
| PSNR / SSIM / LPIPS / FOR ₈ | | 24.33 / 0.949 / 0.076 / 0.5337 | 23.97 / 0.951 / 0.049 / 0.1444 |
|  | |  |  |
| PSNR / SSIM / LPIPS / FOR ₈ | | 17.75 / 0.825 / 0.163 / 0.6415 | 17.71 / 0.843 / 0.188 / 0.5838 |
|  | |  |  |
| PSNR / SSIM / LPIPS / FOR ₈ | | 20.82 / 0.914 / 0.083 / 0.7219 | 20.15 / 0.908 / 0.102 / 0.6310 |

Figure 4. **Shortcoming of standard image metrics** – we show example (left) ground-truth renderings from the target view, and the novel-view images by (middle) Zero-1-to-3 XL and (right) our method. We show PSNR / SSIM / LPIPS / FOR₈ for each shape. As shown, standard image metrics—PSNR, SSIM, and LPIPS—may not correspond to how ‘accurate’ each image is due to misalignments. Arguably, our renderings are more consistent and faithful to the ground truth. However, each metric provides a different story. We propose an optical flow-based metric, FOR_k, to compensate for this shortcoming and provide a holistic view.

While we also report standard images for novel-view reconstruction, Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Metric (SSIM) [71], and Learned Perceptual Image Patch Similarity (LPIPS) [79], we notice that these metrics do not faithfully resemble the synthesis quality. It is because synthesizing a novel view from a single image is inherently an ill-posed problem, and even the slightest error in the implied 3D estimate can lead to a large error for the standard metric; see Fig. 4. This is expected, since these standard metrics are well known to be either highly sensitive to misalignment as in the case of PSNR and SSIM, or lead to measures that cannot distinguish between good and poor alignment due to insensitivity, *e.g.*, with LPIPS.

Optical flow metric – FOR_k. We thus propose to use a metric that measures how many of the rendered pixels are close enough to where they really should be – optical flow. To retrieve optical flow, given that the generated images and the ground truth are supposed to be similar, we use RAFT [64]. We then measure the ratio of flow estimates that deviate significantly (we use various thresholds) from RAFT estimates to account for the potential errors that RAFT itself may make. We denote our metric as FOR_k (flow outlier ratio), where k is the pixel threshold—we use 8 and 16 as RAFT is reported [64] to make an average pixel error of 5 pixels on the KITTI benchmark [16]. Note that our metric is similar to how the KITTI benchmark measures optical flow accuracy by counting the ratio of optical flow outliers. Ideally, if the rendering was perfect, this metric

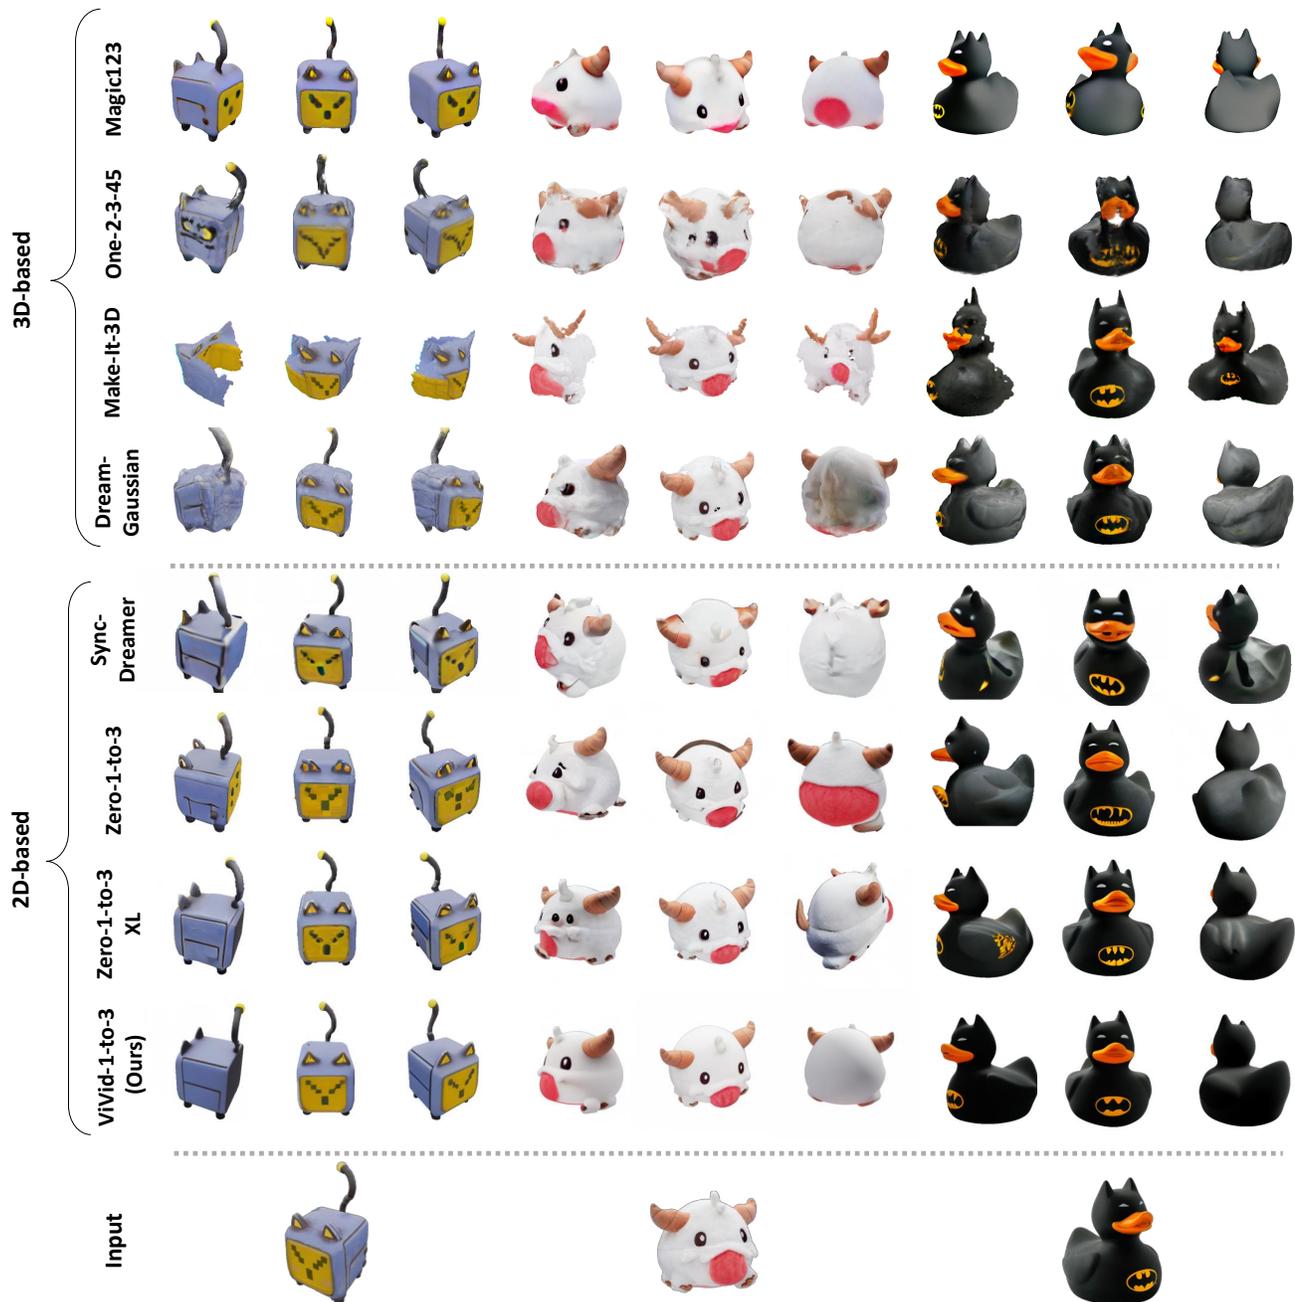


Figure 5. **Qualitative highlights** – we provide examples for multiple recent baselines. Our results are most consistent with the input image. Note the direction of the tail for the left object, the horns of the object in the middle, and the tail of the object on the right.

would be zero, and the lower the ratio of flow outliers the more accurate. Since we are looking at estimated optical flow, which relies on appearance, this metric takes into account both the faithfulness of the appearance and the alignment; see Fig. 4.

4.2. Qualitative comparison

Baselines. We compare our method against the following 2D and 3D baselines:

- **2D) Zero-1-to-3 (XL)** [10, 34]: the base model for our novel-view diffusion model. We compare against both the original version and the one fine-tuned on Objaverse-XL [10].

- **2D) SyncDreamer** [35]: a 2D novel-view diffusion model that improves the consistency of Zero-1-to-3 XL by finetuning it with pre-defined fixed views. For this method, we thus show the nearest view, in its pre-defined view set.
- **3D) Magic123** [44]: a method that consists of two stages of coarse-to-fine generation process for textured 3D mesh. Firstly it utilizes both 2D [50] and 3D [43] diffusion prior to train NeRF. It then converts the trained NeRF to DM Tet [55] to generate high-resolution 3D textured models.
- **3D) Make-It-3D** [63]: a method that first uses the reference image and its monocular depth estimate to optimize a NeRF. It then uses the Score Distillation Sampling (SDS) [43] loss with a text prompt obtained via BLIP-2 [31] and optimizes a textured point cloud, which is then rendered as novel views via deferred neural rendering.
- **3D) One-2-3-45** [33]: a method that trains a straightforward and efficient multi-view reconstruction model with a 3D convolutional neural network using images generated from Zero-1-to-3 (XL) [10, 34].
- **3D) DreamGaussian** [62]: a method that leverages the recently popularized 3D Gaussian Splatting [28] for efficient 3D textured mesh generation.

For all baselines, we use the official code provided by the authors and use default parameters. We visualize each baseline from a similar view, which we manually align for 3D methods due to the different coordinate conventions of each method, and the inconsistency between the pose estimates and the requested pose.

Discussion Fig. 5. We first demonstrate our results qualitatively in Fig. 5. As shown, our method provides results that are much more consistent with the input image. Note that many of the methods shown, qualitatively, look good. However, upon close inspection, these models are often not consistent with either the input image, especially as the target view deviates strongly from the input image. For example, the tail direction of the left object, the horns of the middle object, or the ears and the tails of the right object.

We notice that the methods that utilize explicit 3D representations show worse quality than the pure 2D ones. This is potentially due to their reliance on monocular depth estimates from off-the-shelf methods [46, 47], and the pose inconsistencies that 2D novel-view diffusion models bring. While their renderings are 3D consistent by construction, they tend to be blurrier and less faithful to the input image.

From text to novel views – Fig. 6. We further demonstrate rendering multiple views of objects purely from text. We generate an image of an object with Stable Diffusion, remove the background via Ranftl et al. [47], and then use our method to generate novel views. The generated views are not only consistent with the input view but also preserve

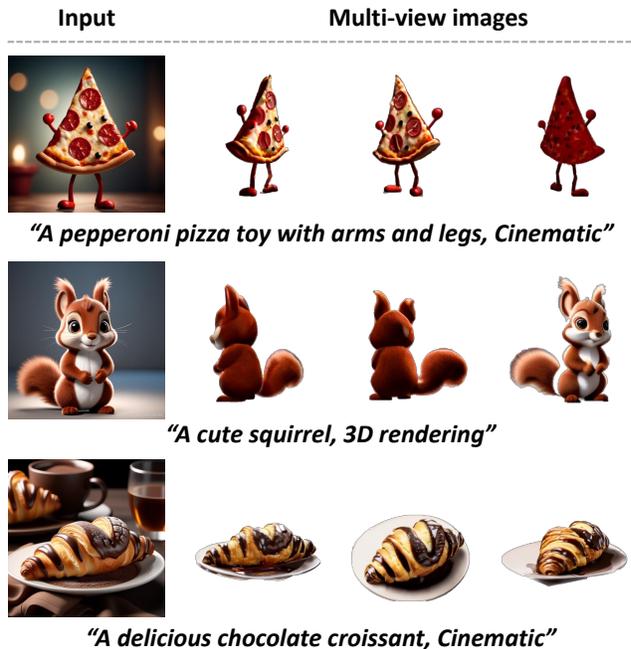


Figure 6. **From text to novel-view synthesis** – we show examples of multiple views of objects being generated purely from text prompts via Stable Diffusion [50] and our method. We show both the original view and the novel view.

the semantics of the input prompt. Our framework yields a high-quality text-to-novel view synthesis model when combined with the denoising pipeline of stable diffusion.

4.3. Quantitative results

Baselines. Due to the limited amount of computational resources, we focus our efforts to peer-reviewed baselines at the time of writing: Zero-1-to-3 (XL) [10, 34] and Make-It-3D [63]. As before, we use the official implementations, with their default configurations. For Make-It-3D, we render images at 800×800 which is the default rendering resolution, and then resize it to 512×512 with bilinear sampling.

Flow-based metric – Fig. 8. We report the flow-based metric at varying thresholds—8 and 16 pixels—in Fig. 8. Here, we exclude Make-It-3D as its results were significantly worse than others, as we will show via standard metrics in Tab. 1. Our method significantly improves over Zero-1-to-3 XL, especially when the requested viewing angle is distant from the original view. This demonstrates how video diffusion helps in providing more consistent renderings.

Image-based metrics – Tab. 1. For completeness, we further report the standard image-based metrics in Tab. 1. Our method also improves over the state of the art in terms of traditional image quality metrics. Interestingly, Make-It-3D, by focusing on building an explicit 3D representation,

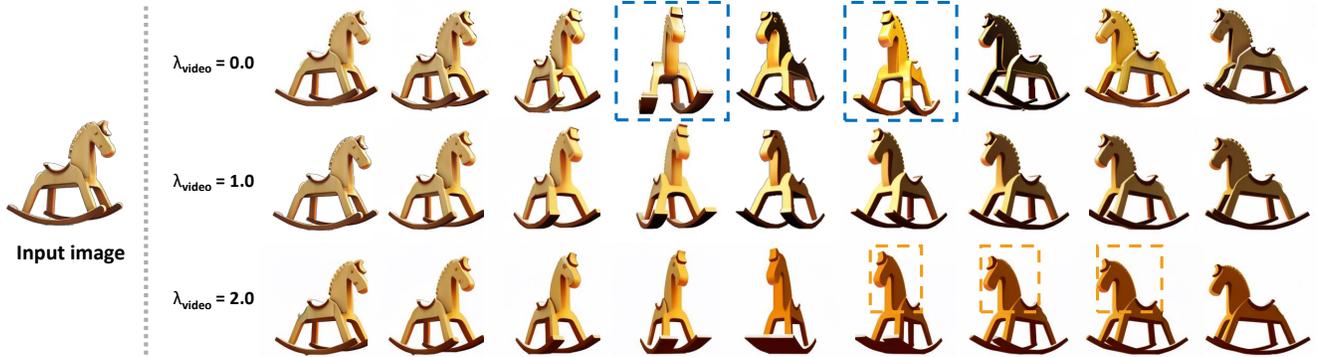


Figure 7. **Effect of hyperparameters** – we show an example of our method when (**top**) video diffusion is not used hence equivalent to Zero-1-to-3 XL [10, 34], (**middle**) with optimal hyper parameters, and (**bottom**) with too much influence from video diffusion model. Without video diffusion the poses are inconsistent—they change abruptly from one image to another (marked with blue boxes). With too much video diffusion the content smooths out, losing detail (marked with orange boxes).

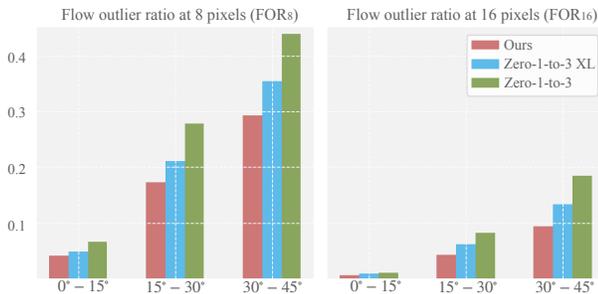


Figure 8. **Optical flow outlier ratio** – we report the optical flow outlier ratio with varying thresholds (8 and 16 pixels) for each method for novel view images generated for different viewing angles. Our method provides significant improvement over Zero-1-to-3 XL [10, 34] and outperforms all methods.

| | PSNR \uparrow | SSIM \uparrow | LPIPS \downarrow | FOR $_8\downarrow$ | FOR $_{16}\downarrow$ |
|-------------------------|-----------------|-----------------|--------------------|--------------------|-----------------------|
| Make-It-3D [63] | 17.13 | 0.841 | 0.135 | 0.876 | 0.501 |
| Zero-1-to-3 [34] | 22.66 | 0.902 | 0.106 | 0.255 | 0.091 |
| Zero-1-to-3 XL [10, 34] | 23.47 | 0.909 | 0.101 | 0.203 | 0.068 |
| ViVid-1-to-3 (Ours) | 24.05 | 0.917 | 0.099 | 0.178 | 0.060 |

Table 1. **Image metrics** – We report the PSNR, SSIM, and LPIPS for each method for all views. Our method performs best.

| λ_{video}^s | λ_{video}^e | PSNR \uparrow | SSIM \uparrow | LPIPS \downarrow | FOR $_8\downarrow$ | FOR $_{16}\downarrow$ |
|----------------------------|----------------------------|-----------------|-----------------|--------------------|--------------------|-----------------------|
| 1.0 | 0.5 | 22.55 | 0.905 | 0.105 | 0.2923 | 0.0944 |
| 1.0 | 1.0 | 22.55 | 0.905 | 0.106 | 0.2958 | 0.0949 |
| 1.0 | 0.0 | 22.48 | 0.905 | 0.105 | 0.3082 | 0.0995 |
| 1.5 | 0.5 | 22.44 | 0.904 | 0.106 | 0.2952 | 0.1072 |

Table 2. **Effect of hyperparameters** – we show all metrics for the azimuth range of 30–45 degrees for various scheduling of λ_{view} and λ_{video} . We keep $\lambda_{\text{view}}=1$ and schedule λ_{video} , where we denote the linear scheduling as $\lambda_{\text{video}}^s-\lambda_{\text{video}}^e-t$, where λ_{video}^s is the starting λ_{video} value, λ_{video}^e is the end value at timestep 50.

loses quality when it comes to actual 2D renderings.

Effect of hyperparameters λ_{view} and λ_{video} – Tab. 2 and

Fig. 7. As discussed earlier in Sec. 3.2, the choice of λ_{view} and λ_{video} matters. We investigate multiple parameter settings and report a subset of our search in Tab. 2. Note that our optical flow-based metric, FOR $_k$ is more distinctive. As shown, relying either too much or too little on the video diffusion model is suboptimal. A representative example of both cases is shown in Fig. 7.

5. Conclusion

We have presented a framework for novel-view synthesis, that poses the problem as a video generation problem, which allows combining novel-view diffusion models with video diffusion models. By utilizing the strong priors learned within video diffusion models, we achieve more consistent novel-view synthesis results. To compensate for the shortcomings of standard image-based metrics, we propose a novel metric based on optical flow. We compared our method existing methods, achieving the state of the art.

Limitations and future work. While our method delivers improved consistency with the input image, it still does not have an explicit 3D model and can be multi-view inconsistent. We note, however, that our method is complementary to other methods that focus on consistent novel-view synthesis including those that embed 3D models. With our improved consistency in a pure 2D pipeline, a promising future direction would be to incorporate explicit 3D pipelines.

Acknowledgment. This work was supported in part by the MOTIE (Ministry of Trade, Industry, and Energy) in Korea, under Human Resource Development Program for Industrial Innovation (Global) (P0017311) supervised by the Korea Institute for Advancement of Technology (KIAT), Brand Engagement Network (BEN) [Q2312881], the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant, Digital Research Alliance of Canada, and Advanced Research Computing at the University of British Columbia.

References

- [1] Zeroscope v2 XL model. https://huggingface.co/cerspense/zeroscope_v2_XL, . 2, 3, 4, 12
- [2] Zeroscope v2 576w model. https://huggingface.co/cerspense/zeroscope_v2_576w, . 12
- [3] Sameer Agarwal, Noah Snavely, Ian Simon, Steven M. Seitz, and Richard Szeliski. Building Rome in a day. In *ICCV*, 2009. 2
- [4] Mohammadreza Armandpour, Huangjie Zheng, Ali Sadeghian, Amir Sadeghian, and Mingyuan Zhou. Re-imagine the Negative Prompt Algorithm: Transform 2D Diffusion into 3D, alleviate Janus problem and Beyond. *arXiv*, 2023. 3
- [5] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. eDiff-I: Text-to-Image Diffusion Models with an Ensemble of Expert Denoisers. *arXiv*, 2022. 5
- [6] 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. 3, 4
- [7] Eric R. Chan, Koki Nagano, Matthew A. Chan, Alexander W. Bergman, Jeong Joon Park, Axel Levy, Miika Aittala, Shalini De Mello, Tero Karras, and Gordon Wetzstein. Generative Novel View Synthesis with 3D-Aware Diffusion Models. *CoRR*, 2023. 1, 3
- [8] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3D: Disentangling Geometry and Appearance for High-quality Text-to-3D Content Creation. In *ICCV*, 2023. 3
- [9] Erik B Dam, Martin Koch, and Martin Lillholm. *Quaternions, Interpolation and Animation*. Citeseer, 1998. 4
- [10] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-XL: A Universe of 10M+ 3D Objects. *arXiv*, 2023. 2, 3, 6, 7, 8, 12
- [11] 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, 3
- [12] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. Google Scanned Objects: A High-Quality Dataset of 3D Scanned Household Items. In *ICRA*, 2022. 2, 5, 12, 14, 15
- [13] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming Transformers for High-Resolution Image Synthesis. In *CVPR*, 2021. 3
- [14] John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. Deep Stereo: Learning to Predict New Views from the World’s Imagery. In *CVPR*, 2016. 3
- [15] Jonathan Freer, Kwang Moo Yi, Wei Jiang, Jongwon Choi, and Hyung Jin Chang. Novel-View Synthesis of Human Tourist Photos. In *WACV*, 2022. 2
- [16] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for Autonomous Driving? The Kitti Vision Benchmark Suite. In *CVPR*, 2012. 5
- [17] Michael Goesele, Noah Snavely, Brian Curless, Hugues Hoppe, and Steven M. Seitz. Multi-View Stereo for Community Photo Collections. In *ICCV*, 2007. 2
- [18] 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*, 2023. 2, 3, 4
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016. 3
- [20] Eric Hedlin, Gopal Sharma, Shweta Mahajan, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. Unsupervised Semantic Correspondence Using Stable Diffusion. *arXiv*, 2023. 3
- [21] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv*, 2022. 5
- [22] Jonathan Ho and Tim Salimans. Classifier-Free Diffusion Guidance. In *NeurIPS*, 2021. 5
- [23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In *NeurIPS*, 2020. 4
- [24] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey A. Gritsenko, Diederik P. Kingma, Ben Poole, Mohammad Norouzi, David J. Fleet, and Tim Salimans. Imagen Video: High Definition Video Generation with Diffusion Models. *CoRR*, 2022. 3
- [25] Jonathan Ho, Tim Salimans, Alexey A. Gritsenko, William Chan, Mohammad Norouzi, and David J. Fleet. Video diffusion models. In *NeurIPS*, 2022. 3
- [26] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting Nerf on a Diet: Semantically Consistent Few-Shot View Synthesis. In *ICCV*, 2021. 1
- [27] Ajay Jain, Ben Mildenhall, Jonathan T. Barron, Pieter Abbeel, and Ben Poole. Zero-Shot Text-Guided Object Generation with Dream Fields. In *CVPR*, 2022. 3
- [28] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ToG*, 2023. 2, 7
- [29] Aliasghar Khani, Saeid Asgari Taghanaki, Aditya Sanghi, Ali Mahdavi Amiri, and Ghassan Hamarneh. Slime: Segment like me. *arXiv*, 2023. 3
- [30] Gaëtan Landreau and Mohamed Tamaazousti. EpipolarNVS: leveraging on Epipolar geometry for single-image Novel View Synthesis. In *BMVC*, 2022. 3
- [31] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv*, 2023. 7
- [32] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3D: High-Resolution Text-to-3D Content Creation. In *CVPR*, 2023. 3
- [33] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Zexiang Xu, Hao Su, et al. One-2-3-45: Any Single Image to 3D Mesh in 45 Seconds without Per-Shape Optimization. *arXiv*, 2023. 2, 3, 7, 12

- [34] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot One Image to 3D Object. In *ICCV*, 2023. 1, 2, 3, 4, 5, 6, 7, 8, 12
- [35] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. SyncDreamer: Learning to Generate Multiview-consistent Images from a Single-view Image. *arXiv*, 2023. 2, 3, 7, 12
- [36] Stephen Lombardi, Tomas Simon, Jason M. Saragih, Gabriel Schwartz, Andreas M. Lehrmann, and Yaser Sheikh. Neural Volumes: Learning Dynamic Renderable Volumes from Images. *TOG*, 2019. 3
- [37] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, and Wenping Wang. Wonder3D: Single Image to 3D using Cross-Domain Diffusion. *arXiv*, 2023. 3
- [38] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. In *NeurIPS*, 2022. 12
- [39] Zhengxiong Luo, Dayou Chen, Yingya Zhang, Yan Huang, Liang Wang, Yujun Shen, Deli Zhao, Jingren Zhou, and Tieniu Tan. VideoFusion: Decomposed Diffusion Models for High-Quality Video Generation. In *CVPR*, 2023. 3, 4
- [40] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. RealFusion 360° Reconstruction of Any Object from a Single Image. In *CVPR*, 2023. 3
- [41] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations. *arXiv*, 2021. 3
- [42] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*, 2020. 2, 3
- [43] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D Diffusion. In *ICLR*, 2023. 2, 3, 7
- [44] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skokhodov, Peter Wonka, Sergey Tulyakov, and Bernard Ghanem. Magic123: One Image to High-Quality 3D Object Generation Using Both 2D and 3D Diffusion Priors. *CoRR*, 2023. 2, 3, 7, 12
- [45] Amit Raj, Srinivas Kaza, Ben Poole, Michael Niemeyer, Nataniel Ruiz, Ben Mildenhall, Shiran Zada, Kfir Aberman, Michael Rubinstein, Jonathan T. Barron, Yuanzhen Li, and Varun Jampani. DreamBooth3D: Subject-Driven Text-to-3D Generation. *CoRR*, 2023. 3
- [46] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards Robust Monocular Depth Estimation: Mixing Datasets for Zero-shot Cross-Dataset Transfer. *TPAMI*, 2020. 7
- [47] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision Transformers for Dense Prediction. In *ICCV*, 2021. 7
- [48] Gernot Riegler and Vladlen Koltun. Free View Synthesis. In *ECCV*, 2020. 2
- [49] Gernot Riegler and Vladlen Koltun. Stable View Synthesis. In *CVPR*, 2021. 2
- [50] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 1, 2, 3, 7
- [51] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. In *NeurIPS*, 2022. 3
- [52] Johannes L. Schönberger and Jan-Michael Frahm. Structure-from-Motion Revisited. In *CVPR*, 2016. 2
- [53] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. LAION-5B: An open large-scale dataset for training next generation image-text models. In *NeurIPS*, 2022. 3
- [54] Junyoung Seo, Wooseok Jang, Minseop Kwak, Jaehoon Ko, Hyeonsu Kim, Junho Kim, Jin-Hwa Kim, Jiyoung Lee, and Seungryong Kim. Let 2D Diffusion Model Know 3D-Consistency for Robust Text-to-3D Generation. *CoRR*, 2023. 3
- [55] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep Marching Tetrahedra: a Hybrid Representation for High-Resolution 3D Shape Synthesis. In *NeurIPS*, 2021. 7
- [56] 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, 2023. 2, 3
- [57] Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. MVDream: Multi-view Diffusion for 3D Generation. *arXiv*, 2023. 2, 3, 4
- [58] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, Devi Parikh, Sonal Gupta, and Yaniv Taigman. Make-A-Video: Text-to-Video Generation without Text-Video Data. In *ICLR*, 2023. 3
- [59] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. DeepVoxels: Learning Persistent 3D Feature Embeddings. In *CVPR*, 2019. 3
- [60] Noah Snavely, Steven M. Seitz, and Richard Szeliski. Photo Tourism: Exploring Photo Collections in 3D. *TOG*, 2006. 2
- [61] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021. 12
- [62] Jiayang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. DreamGaussian: Generative Gaussian Splatting for Efficient 3D Content Creation. *arXiv*, 2023. 2, 7, 12
- [63] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-It-3D: High-fidelity 3D Creation from A Single Image with Diffusion Prior. In *ICCV*, 2023. 7, 8
- [64] Zachary Teed and Jia Deng. RAFT: Recurrent All-Pairs Field Transforms for Optical Flow. In *ECCV*, 2020. 5

- [65] Alex Trevithick and Bo Yang. GRF: Learning a General Radiance Field for 3D Representation and Rendering. In *ICCV*, 2021. 3
- [66] Richard Tucker and Noah Snavely. Single-View View Synthesis With Multiplane Images. In *CVPR*, 2020. 3
- [67] Aaron Van Den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural Discrete Representation Learning. In *NeurIPS*, 2017. 3
- [68] Vikram Voleti, Alexia Jolicoeur-Martineau, and Chris Pal. MCVD: Masked Conditional Video Diffusion for Prediction, Generation, and Interpolation. In *NeurIPS*, 2022. 3
- [69] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A. Yeh, and Greg Shakhnarovich. Score Jacobian Chaining: Lifting Pretrained 2D Diffusion Models for 3D Generation. In *CVPR*, 2023. 2, 3
- [70] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. In *NeurIPS*, 2021. 2
- [71] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image Quality Assessment: from Error Visibility to Structural Similarity. *TIP*, 2004. 5
- [72] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. ProlificDreamer: High-Fidelity and Diverse Text-to-3D Generation with Variational Score Distillation. *arXiv*, 2023. 2, 3
- [73] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel View Synthesis with Diffusion Models. In *ICLR*, 2023. 1, 3
- [74] 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*, 2023. 2, 3
- [75] DeJia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. NeuralLift-360: Lifting an in-the-Wild 2D Photo to A 3D Object with 360° Views. In *CVPR*, 2023. 3
- [76] Jiale Xu, Xintao Wang, Weihao Cheng, Yan-Pei Cao, Ying Shan, Xiaohu Qie, and Shenghua Gao. Dream3D: Zero-Shot Text-to-3D Synthesis Using 3D Shape Prior and Text-to-Image Diffusion Models. In *CVPR*, 2023. 3
- [77] Jiayu Yang, Ziang Cheng, Yunfei Duan, Pan Ji, and Hongdong Li. ConsistNet: Enforcing 3D Consistency for Multi-view Images Diffusion. *arXiv*, 2023. 2, 3
- [78] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural Radiance Fields From One or Few Images. In *CVPR*, 2021. 1, 3
- [79] 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
- [80] Daquan Zhou, Weimin Wang, Hanshu Yan, Weiwei Lv, Yizhe Zhu, and Jiashi Feng. MagicVideo: Efficient Video Generation With Latent Diffusion Models. *CoRR*, 2022. 3