

4Real: Towards Photorealistic 4D Scene Generation via Video Diffusion Models

Heng Yu^{1,2,*†} Chaoyang Wang^{1*} Peiye Zhuang¹ Willi Menapace¹ Aliaksandr Siarohin¹
Junli Cao¹ László A Jeni² Sergey Tulyakov¹ Hsin-Ying Lee¹

¹Snap Inc. ²Carnegie Mellon University

Project page: <https://snap-research.github.io/4Real/>



Figure 1: **4Real** is a 4D generation framework that can generate near-*photorealistic* dynamic scenes from text prompts. We use deformable 3D Gaussian Splats (D-3DGS) to model the scene. After generation, We can view the generated dynamic scenes at any timestep from different camera poses.

Abstract

Existing dynamic scene generation methods mostly rely on distilling knowledge from pre-trained 3D generative models, which are typically fine-tuned on synthetic object datasets. As a result, the generated scenes are often object-centric and lack photorealism. To address these limitations, we introduce a novel pipeline designed for photorealistic text-to-4D scene generation, discarding the dependency on multi-view generative models and instead fully utilizing video generative models trained on diverse real-world datasets. Our method begins by generating a reference video using the video generation model. We then learn the canonical 3D representation of the video using a freeze-time video, delicately generated from the reference video. To handle inconsistencies in the freeze-time video, we jointly learn a per-frame deformation to model these imperfections. We then learn the temporal deformation based on the canonical representation to capture dynamic interactions in the reference video. The pipeline facilitates the generation of dynamic scenes with enhanced photorealism and structural integrity, viewable from multiple perspectives, thereby setting a new standard in 4D scene generation.

*Equal contribution.

†Work done during internship at Snap Inc.

1 Introduction

As industries ranging from film production to virtual reality seek increasingly immersive and interactive experiences, the ability to generate dynamic 3D scenes over time—essentially, 4D environments—promises to revolutionize how we interact with digital content. Recently, significant advances in image and video generation have been driven by large-scale text-image and text-video datasets, along with the development of diffusion models. Furthermore, image diffusion models have been adapted into 3D-aware multi-view generative models through fine-tuning with limited 3D data, serving as foundational priors for 3D and 4D generation.

In this work, we focus on photorealistic text-to-4D scene generation. Existing 4D generation pipelines, due to the lack of 4D data, typically employ image, multi-view, and video generation models as priors to synthesize 4D samples. However, the multi-view models, which provide critical 3D information, are fine-tuned on static and synthetic 3D assets. As a result, current generated 4D results are predominantly object-centric, lacking photorealism, and limited in their ability to capture complex interactions between objects and environments.

In response, we propose **4Real**, a novel pipeline designed for photorealistic dynamic scenes with dynamic objects within the environment. We discard the reliance on multi-view generative models and exploit video generative models trained on large-scale real-world videos, covering more diverse and general appearance, shape, motion, and the interaction between objects and environments. Compared to existing methods purely relying on score distillation sampling, the proposed pipeline provides more flexible use cases, more diverse results, and requires fewer computations.

The proposed method adopts deformable 3D Gaussian Splats (D-3DGS) as the representation for dynamic scenes. The pipeline unfolds in three steps. First, we begin by creating a reference video featuring a dynamic scene with a pre-trained text-to-video diffusion model. Next, we reconstruct a canonical 3D representation from a selected frame of the reference video. To achieve this, we generate a ‘freeze-time’ video with camera motion and minimal object movement by applying dataset context embedding and prompt engineering to the video diffusion model. However, practically, the generated video still may contain object motion and is prone to be geometrically inconsistent despite looking temporally smooth. We propose to recognize the inconsistencies as per-frame deformation with respect to the canonical representation and learn these deformations simultaneously with the canonical representation. Finally, we reconstruct temporal deformation from the reference video with the learned canonical representation. We choose a video score distillation sampling (SDS) strategy that renders two types of videos: a fixed-viewpoint video with a static camera, and a freeze-time video with a fixed time step. The fixed-viewpoint video helps learn temporally consistent motions, while the freeze-time video learns multi-view geometrical consistency.

4Real achieves text-driven dynamic scene generation with a near-photorealistic appearance and realistic 3D motions. The generated scenes are viewable from different viewing angles and over time, as shown in Figure 1. Our contributions are summarized as follows:

- We propose **4Real**, the first photorealistic text-to-4D scene generation pipeline. Without reliance on biased multi-view image generation models trained with specialized datasets, the proposed pipeline can generate more diverse and near-photorealistic results with dynamic objects within realistic environments.
- We leverage text-to-video diffusion to generate target reference videos and auxiliary freeze-time videos, thereby transforming the generation problem into a reconstruction problem and reducing reliance on time-consuming score distillation sampling steps. We learn from the freeze-time videos canonical 3DGS along with per-frame deformation modeling video inconsistency, and learn from the reference videos temporal deformation.
- The proposed pipeline provides users flexibility in selecting and editing videos that they want to lift to 4D, and can generate high-quality samples in a more reasonable computation budget, taking 1.5 hours on an A100 GPU compared to 10+ hours with competing methods [2, 72].

2 Related Works

Text-to-video generation. Our pipeline leveraged the advancements in video generation models, which have been spurred by the recent success of text-to-image diffusion models. Some approaches

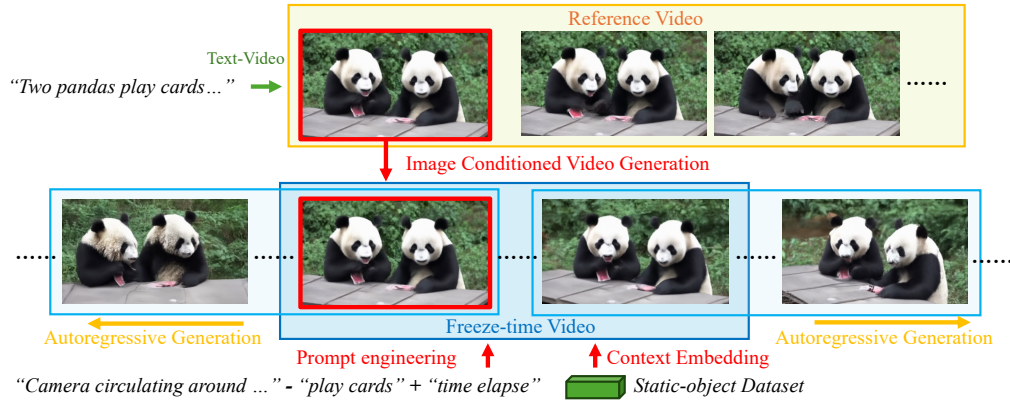


Figure 2: **Generate reference and freeze-time videos.** Given the input text prompt, we first generate a *reference video* using text-to-video diffusion model. The reference video will be our target to perform 4D reconstruction. In order to obtain canonical Gaussian Splats, we generate a *freeze-time video* by performing frame-conditioned video generation with prompt engineering and context embedding. We further perform auto-regressive generation to expand view angle coverage.

extend these capabilities by directly fine-tuning pre-trained image diffusion models [3, 4] or by integrating motion modules while maintaining fixed image models [17, 34, 71]. Others utilize both image and video data using 3D-Uet [20] and cascaded temporal and spatial upsampler [19, 48]. More recently, some efforts [5, 35] propose to replace the Unet architectures with transformer-based architectures [9, 41] for improved quality, scalability, and training efficiency.

Object-centric 3D and 4D generation. The challenges of generating 3D and 4D data mostly lie in the scarcity of available data, due to the costly nature of data capture. Conventional 3D generative models [1, 10–12, 22, 28, 38, 50, 57] that train on limited 3D data in various representation fall short in quality and diversity, while 2D diffusion models have shown incredible generalizability using large-scale data. To bridge the gap, Score Distillation Sampling (SDS) [42] and its variants [8, 27, 56, 58, 73] have been developed to utilize pre-trained text-to-image models as priors to generate 3D objects by optimizing parametric spaces for 3D representations. Recently, multiview images have emerged as a popular representation in 3D generation. Multiview image diffusion models [30–32, 47] are fine-tuned from image diffusion models using synthetic object dataset [13]. To generate 3D objects, the multiview image diffusion models are used as approximate 3D priors [44, 62], or as providing intermediate medium followed by 3D reconstruction methods [21, 25, 32, 51, 59, 65, 74]. Meanwhile, video diffusion models have also been explored as priors for some early 4D generation attempts [49] or as data generators to provide multiview supervision for 3D model [52]. More recently, image, multiview, and video diffusion models are jointly used as priors for 4D generation [2, 29, 45, 68, 72]. However, they typically produce results biased toward object-centric and non-photorealistic outputs as the 3D priors are obtained from synthetic object datasets.

4D Reconstruction. Neural Radiance Fields (NeRFs) [36] has revolutionized 3D scene reconstruction and novel view synthesis. The breakthroughs have been expanded to dynamic scenes. There are two major branches of methods for dynamic NeRFs: one views them as a 4D representation by extending the radiance fields with a time dimension [14, 16, 26, 53, 55, 64], while the other represents a dynamic scene by a canonical space coupled with a deformation field that maps local observation to this space [15, 39, 40, 43, 54, 66]. Recently, 3D Gaussian Splatting (3DGS) [23] has been introduced as a promising representation that facilitates fast neural rendering using explicit Gaussian particles. 3DGS has then been applied to model dynamic scenes with the similar idea of building a deformation field [33, 63, 67, 69]. We adopt deformable 3D Gaussian Splats (D-3DGS) as our dynamic scene representation.

3 Method

Given input text prompts, we aim to generate photorealistic dynamic scenes including animated objects and detailed backgrounds that exhibit plausible relative scales and realistic interactions with moving objects. We adopt deformable 3D Gaussian Splats (D-3DGS) as our representation. First, we

utilize a text-to-video diffusion model to create a reference video with dynamic scenes. Next, a frame from this reference video serves as the conditional input for the video diffusion model to produce a freeze-time video featuring circular camera motion and minimal object movement. Subsequently, we reconstruct the canonical 3D representation from the freeze-time video. Finally, we reconstruct temporal deformations to align with the object motions in the reference video.

In the following sections, we will first discuss the 4D Gaussian splats representation for dynamic scenes in Section 3.1, the generation of freeze-time videos in Section 3.2, the reconstruction of canonical 3DGS in Section 3.3 and the temporal reconstruction in Section 3.4.

3.1 Dynamic Scene Representation

We employ Deformable 3D Gaussian Splats (D-3DGS) to model dynamic scenes, a popular choice for dynamic novel view synthesis tasks due to its low rendering computation cost, state-of-the-art rendering fidelity, and compatibility with existing graphics rendering pipelines [33, 63, 67, 69]. D-3DGS consists of 3D Gaussian splats that represent the static 3D scene at a canonical frame and a deformation field that models dynamic motions.

Canonical 3D Gaussian splats. 3D Gaussian Splats [23] consist of a set of 3D Gaussian points, each associated with an opacity value α and RGB color $\mathbf{c} \in \mathbb{R}^3$. The position and shape of each 3D Gaussian point are represented by a Gaussian mean position $\mathbf{x} \in \mathbb{R}^3$ and a covariance matrix, which is factorized into orientation, represented by a quaternion $q \in \mathbb{R}^4$, and scaling, represented by $\mathbf{s} \in \mathbb{R}^3$. 3DGS are rendered by projecting Gaussian points onto the image plane and aggregating pixel values using NeRF-like volumetric rendering equations.

Deformation field. To represent motion in 4D scenes, we employ a deformation field $w_{t\text{-deform}}$ that represents the offset of 3DGS from the canonical frame to the frame at time t :

$$w_{t\text{-deform}}(\mathbf{x}, t) = (\Delta\mathbf{x}_t, \Delta\mathbf{q}_t). \quad (1)$$

The position and orientation of 3DGS at time t are simply obtained by $\mathbf{x}_t = \mathbf{x} + \Delta\mathbf{x}_t$ and $\mathbf{q}_t = \mathbf{q} + \Delta\mathbf{q}_t$. We implement w using a multi-layer perceptron (MLP), which serves as a general representation for arbitrary deformations. It is worth noting that this can be substituted by more specialized representations, such as linear blend skinning, which is more suitable for articulated motions. In this work, we prefer MLPs due to their simplicity and generality.

3.2 Generating Reference and Freeze-time Videos

Given a text prompt, we generate a reference video using video diffusion models. First, we would like to obtain the canonical 3D Gaussian splats for this reference video by creating a freeze-time video conditioned on a selected frame. This process requires a video generation model that can: 1) perform image-to-video generation, ensuring the output video frames remain consistent with the input image, and 2) generate videos with substantial camera motion while maintaining a relatively static scene.

We evaluated several recent video models accessible to the authors, including SVD [3], SV3D [52], VideoCrafter [7], and Snap Video Model [35]. SVD and VideoCrafter support image-to-video generation but have limited camera movement and introduce random motions in non-rigid objects. SV3D creates object-centric 360° videos with a white background, unsuitable for the scene-level generation. The Snap Video Model aligns most closely with our requirements, offering the required capabilities for our application.

Frames-conditioned pixel-space video diffusion. The Snap Video Model is a text-to-video diffusion model that diffuses pixel values instead of latents, resulting in less motion blur and fewer gridding effects, which are more common in latent diffusion models. It also tends to produce larger camera and object motions [35]. We use a variant trained to generate videos from both text prompts and an arbitrary set of video frame inputs. This feature enhances flexibility for autoregressive generation, particularly when creating freeze-time videos with extensive view coverage. Details of training and inference are in the Supplementary.

Context conditioning and prompt engineering. The Snap Video Model is trained on mixed video datasets, with a unique context embedding for each dataset to guide the model in generating videos that follow the specific distribution. One dataset consists of static real-world objects with circular camera trajectories, helping the model generate freeze-time videos with large camera motions and

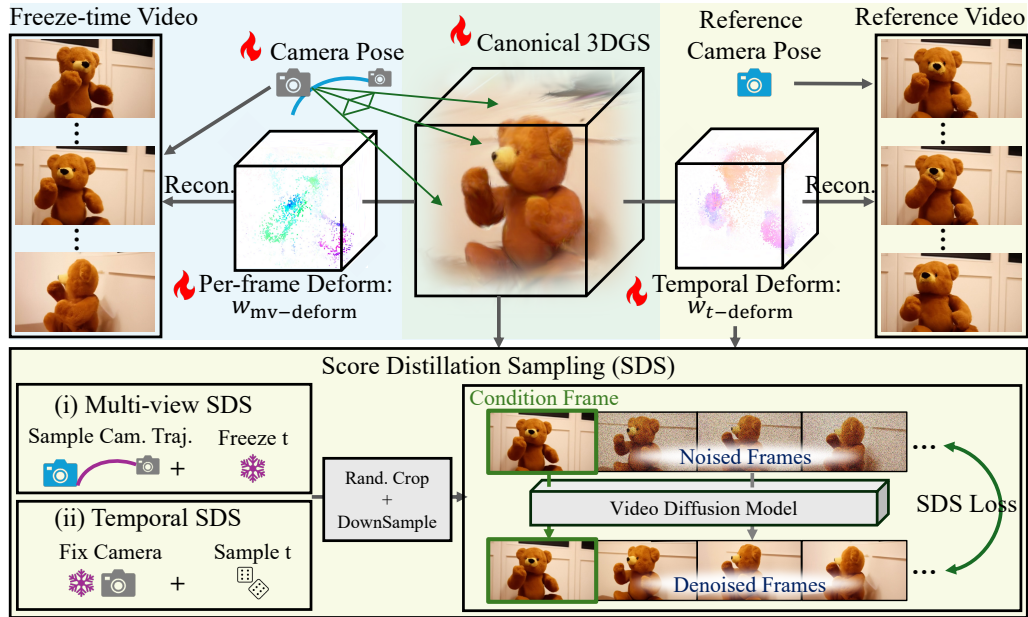


Figure 3: **Reconstructing Deformable 3DGS.** First, we reconstruct the canonical 3DGS using the freeze-time video. As the generated freeze-time video may still contain object motions and geometric inconsistencies, we propose to model these imperfections as per-frame deformations that are learned jointly with the canonical 3DGS. Then, we reconstruct the temporal deformation from the reference video with the learned canonical 3DGS. In addition to the reconstruction loss and motion regularization, we propose a video SDS strategy. The video SDS strategy includes a multi-view SDS, using videos of freeze time and sampled camera trajectory, and a temporal SDS, using videos of fixed camera and sampled time steps. 🔥 denotes learnable parameters.

minimal object movements when its context embedding is used. The success rate can be further improved by adding phrases such as "Camera circulating around" and "time elapse" to the beginning and end of the input prompt, while omitting action verbs. Despite this, limited by the current capability of the Snap Video Model, multi-view inconsistency and moderate object motion may still occur, which our reconstruction steps in Section 3.3 address.

Extending view coverage. The raw output from the Snap Video Model is a sequence of 16 frames, resulting in limited view coverage. To extend coverage, we perform iterative autoregressive generation. At each step, we condition on the 8 previously generated frames, guiding the model to generate 8 new frames that follow the camera trajectory of the prior frames. We extend the view coverage in both forward and backward directions until it reaches a sufficient amount, such as 180° . Covering 360° is currently challenging due to background clutter and the model's limitation in generating loop videos.

3.3 Robust Reconstruction of 3DGS from Noisy Freeze-time Videos

Constrained by the capability of current video models, generated freeze-time videos may still contain multi-view inconsistencies, preventing direct reconstruction of high-quality static 3DGS. These inconsistencies can be: 1) temporally consistent but geometrically incorrect 2D video, and 2) videos with object motions. To address these imperfections, we propose the following treatments.

Representing multi-view inconsistency as deformation. We treat multi-view inconsistency as per-frame deformation with respect to the canonical 3DGS. Specifically, we introduce another deformation field $w_{\text{multi-view}}(\mathbf{x}, k) = (\Delta\mathbf{x}_k, \Delta\mathbf{q}_k)$ similarly defined as in Eq. (1), which outputs changes in position and orientation of a Gaussian point at the canonical position \mathbf{x} with respect to its new position and orientation at the k -th frame. We choose the canonical frame to be the one from the reference video that is used as the conditional input to generate the freeze-time video. This implies zero-valued outputs for the deformation field at the canonical time frame.

We jointly optimize the deformation field $w_{\text{multi-view}}$ and the parameters for the 3DGS by minimizing a combination of the following losses.

Image reconstruction loss computes the difference between images \mathcal{I}_{GS} rendered by the deformed 3DGS and the frames \mathcal{I}_k from the freeze-time video.

$$\mathcal{L}_{\text{recon}} = \sum_k \|\mathcal{I}_{\text{GS}}(\mathbf{x} + \Delta\mathbf{x}_k, \mathbf{q} + \Delta\mathbf{q}_k, \mathbf{s}, \alpha, \mathbf{c}, \mathbf{P}_k) - \mathcal{I}_k\|_1, \quad (2)$$

where k denotes the frame index, \mathbf{P}_k represents the camera projection matrix which is initialized using Colmap [46], then jointly finetuned with the 3DGS, and with some abuse of notations, $(\mathbf{x} + \Delta\mathbf{x}_k, \mathbf{q} + \Delta\mathbf{q}_k, \mathbf{s}, \alpha, \mathbf{c})$ denotes parameters for the deformed 3DGS.

Small motion loss. Minimizing $\mathcal{L}_{\text{recon}}$ alone is insufficient to obtain plausible 3DGS, since there can be infinitely many pairs of deformation fields and 3DGS that produce identical image renderings. To address this under-constrained problem, we introduce the assumption that the deformations should be as small as possible. Specifically:

$$\mathcal{L}_{\text{small mo.}} = \sum_k \|\Delta\mathbf{x}_k\|_1 + \|\Delta\mathbf{q}_k\|_1. \quad (3)$$

Score distillation sampling (SDS) loss. Hand-crafted heuristics like $\mathcal{L}_{\text{small mo.}}$ may be insufficient for plausible canonical 3DGS, especially with significant multi-view inconsistencies in the freeze-time videos. For improved regularization, in the *later* stage of optimization, we employ score distillation sampling. This method ensures that the rendering of canonical 3DGS on randomly sampled camera trajectories aligns with the distribution dictated by the video diffusion model.

During each optimization step, we sample a camera trajectory by first randomly selecting a camera pose \mathbf{P}_k from the camera trajectories of the freeze-time video and then randomly perturbing it to create $\hat{\mathbf{P}}$. The trajectory $\{\hat{\mathbf{P}}_{k=1, \dots, 16}\}$ is uniformly interpolated between the canonical camera pose and $\hat{\mathbf{P}}$. We then render the canonical 3DGS using this sampled trajectory to generate video frames:

$$\mathcal{I}_{\text{can-GS}}^k := \mathcal{I}_{\text{GS}}(\mathbf{x}, \mathbf{q}, \mathbf{s}, \alpha, \mathbf{c}, \hat{\mathbf{P}}_k), \quad \forall k = 1, \dots, 16. \quad (4)$$

To perform score distillation sampling, we add Gaussian noise to the rendered images $\mathcal{I}_{\text{can-GS}}^k$ and then use the video diffusion model to produce one-step denoised images. Since the sampled camera trajectory originates from the canonical camera pose, the first frame $\mathcal{I}_{\text{can-GS}}^0$ of the rendered videos is identical to the corresponding frame \mathcal{I}_c from the freeze-time video, allowing us to substitute $\mathcal{I}_{\text{can-GS}}^0$ with \mathcal{I}_c , which can then serve as the conditional frame for the diffusion model. More specifically:

$$\{\hat{\mathcal{I}}_2, \hat{\mathcal{I}}_3, \dots, \hat{\mathcal{I}}_{16}\} = f_{\text{denoise}}(\mathcal{I}_{\text{can-GS}}^2 + \epsilon_2, \mathcal{I}_{\text{can-GS}}^3 + \epsilon_3, \dots, \mathcal{I}_{\text{can-GS}}^{16} + \epsilon_{16} | \mathcal{I}_c, \sigma), \quad (5)$$

where ϵ 's represent the noise applied to the video frames, σ denotes the noise level for the current diffusion step, and $\hat{\mathcal{I}}_k$ indicates the x_0 prediction for each frame. Using \mathcal{I}_c as the image condition input is crucial because it allows the diffusion model to provide more precise supervisions, and sharper details that are aligned with the reference video.

Then SDS is implemented by minimizing the sum-of-square difference between the rendered video frames $\mathcal{I}_{\text{can-GS}}^k$ and the denoised frames $\hat{\mathcal{I}}_k$, as follows,

$$\mathcal{L}_{\text{SDS}} = \|\mathcal{I}_{\text{can-GS}}^k - [\hat{\mathcal{I}}_k]\|^2. \quad (6)$$

In this context, we use $[\cdot]$ to indicate that the variables are detached from the automatic differentiation computation graph. It is straightforward to show that the gradient of \mathcal{L}_{SDS} is proportional to the original gradient formulation by Poole *et al.* [42].

Sampling with low-res pixel-space diffusion model. The video diffusion model we employ is a two-stage pixel-space model. It first generates a video at a low resolution of 36×64 , which is then upsampled to a resolution of 288×512 . For SDS, we employ only the first-stage low-resolution model, as it is significantly less computationally demanding and not straightforward to apply both stages simultaneously. Fortunately, utilizing just the low-resolution SDS is adequate to yield high-quality results. This effectiveness stems from our method relying more on reconstruction loss to enhance fine details, rather than depending primarily on SDS.

To compute the SDS loss, we first downsample the rendered video frames to low resolution. Before the downsampling, frames are randomly shifted and cropped to mitigate aliasing effects.

Remarks on computational efficiency. Previous methods employing SDS are typically computationally intensive, making them less practical for real-world applications. For instance, executing one step of SDS training using single precision on A100-80G GPUs with SVD can take ≈ 2 seconds, not including the time required for rendering. Moreover, GPU memory often becomes overloaded with SVD, which cannot compute gradients for more than four frames simultaneously due to substantial memory demands for backpropagation through the image encoder.

In contrast, our method implements SDS using the low-resolution, pixel-space Snap Video Model. This model leverages a compressed transformer architecture [9], which is significantly faster, requiring ≈ 200 ms per step. It is also more memory-efficient, allowing for gradient computation across all 16 frames in a single batch. This efficiency makes our approach more viable for practical applications.

3.4 Reconstruct Temporal Deformation

After obtaining canonical 3DGS representation, we proceed to generate temporal motion by fitting the deformation field (as detailed in Eq. (1)) to align with the reference video.

Image alignment loss computes the similarity between the rendered frames $\mathcal{I}_{\text{GS}}^t$ and the reference video frames \mathcal{I}_t , combining pixel-wise intensity loss and structural similarity index measure (SSIM) loss [60], i.e., $\mathcal{L}_{\text{align}} = \sum_t \|\mathcal{I}_{\text{GS}}^t - \mathcal{I}_t\| + \text{SSIM}(\mathcal{I}_{\text{GS}}^t, \mathcal{I}_t)$.

Deformation regularization. To ensure that the deformation field learns plausible motions, we incorporate losses that enforce both spatial and temporal smoothness in the motion trajectories of each Gaussian splat, following CoGS [69] and Dynamic 3DGS [33]. These as-rigid-as-possible losses encourage minimal variation in both translational and rotational motions of Gaussian splats within their spatial and temporal neighborhoods. We refer readers to the respective papers for detailed treatment. Implementation details specific to our adaptation are included in the supplementary.

Joint temporal and multi-view SDS. Relying solely on hand-crafted regularization often leads to noticeable artifacts in novel view rendering, especially when the camera position deviates significantly from that in the reference video. Therefore, in the later stages of optimization, we employ SDS to ensure that the rendered video frames adhere to the distributions of real videos. For evaluating the SDS loss, we render two types of videos:

1. *Fixed-Viewpoint Videos:* The camera remains static while we sample time steps to deform the 3D Gaussian Splats. To ensure that the first frame can be used as a conditional input for the video diffusion model, the sampled time steps always start from 0. This initial frame is rendered using the canonical 3D Gaussian Splats without any deformation.
2. *Freeze-Time Videos:* With a fixed time step, we sample a continuous sequence of camera poses starting from the reference video’s camera pose. This setup allows us to directly use the corresponding frame from the reference video as the conditional input to the video diffusion model.

Applying SDS loss to both fixed-viewpoint and freeze-time videos is crucial. The former regularizes D-3DGS for *temporally* consistent motions, and the latter ensures *multi-view* consistency, aligning D-3DGS motions with the reference video and avoiding multi-perspective optical illusions.

Improving details by re-upsampling We optionally perform an additional step to enhance the rendering details further. We found that downsampling and then re-upsampling the videos rendered by Gaussian splatting results in fewer artifacts and higher fidelity. Therefore, we render a video with continuously changing viewpoints and time, and then re-upsample it using the upsampler of the video diffusion model. The resulting video serves as a reconstruction target to fine-tune the Gaussian splats, thereby improving the overall quality.

4 Experiment

We evaluated our method using a set of 30 text prompts sourced from related works [2, 35, 72]. Please refer to the included webpage for the complete results. As illustrated in Figure 4, our method successfully generates scenes featuring complex illumination and semi-transparent objects such as water. Additionally, it demonstrates flexibility in creating diverse multi-object scenes.




Figure 4: **Generating challenging scenes.** 4Real is able to generate dynamic scenes with complex illumination and (semi)-transparent materials such as water. It is flexible to produce diverse content and generate multi-object scenes. Please refer to our supplementary webpage for the complete results.

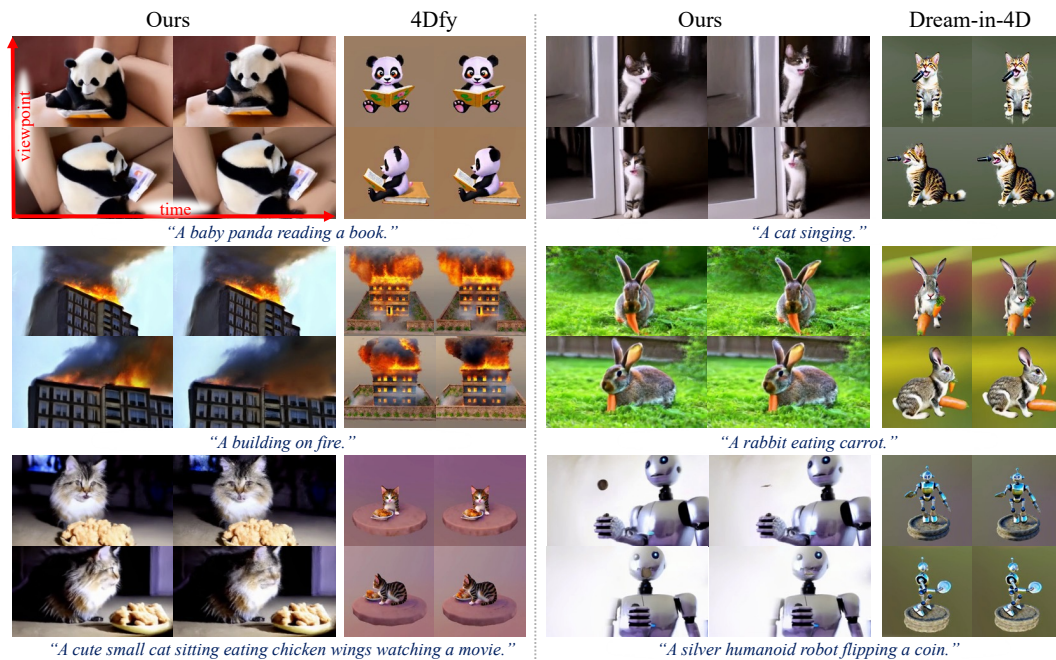


Figure 5: **Qualitative comparison** to state-of-the-arts *object-centric* 4D generation methods.

Implementation details. The optimization consists of two stages: first, reconstructing the canonical 3DGS, and then learning the temporal deformation. At each stage, we perform SDS only during the last 5k iterations, and we anneal the diffusion time steps similarly to the scheduler used in Hifa [73]. Throughout our experiments, we run 20k iterations per stage, totaling 1.5 hours on a single A100 GPU. This duration is significantly shorter than that required by recent state-of-the-art methods, which typically require over 10 hours [2, 72]. Details are included in the Supplementary.

Compare to object-centric text-4D generation. To the best of our knowledge, there are currently no existing text-to-4D scene generation methods, so we choose to evaluate our method against recent state-of-the-art object-centric text-to-4D generation methods. Specifically, we conduct a user study against 4Dfy [2] and Dream-in-4D [72], which are based on NeRF representation. These methods are known to achieve some of the best visual quality compared to other methods [29, 45, 49] using the same representations.

The user study was conducted using Amazon Turk, involving 30 evaluators per video pair. In each session, evaluators were presented with two anonymized videos. Each video depicted a dynamic object or scene, with the camera moving along a circular trajectory and stopping at four fixed poses to highlight object motions. We obtained 16 videos for 4Dfy and 14 videos for Dream-in-4D from their respective project web pages. Evaluators were tasked with selecting their preferences based on seven criteria: *motion realism*, *foreground/background photo-realism*, *3D shape realism*, *general realism*, *significance of motion*, and *video-text alignment*. As shown in Figure 6, our method outperformed the competition in every category. Sample frames are displayed in Figure 5, and all videos used in the study are available in the supplementary materials.




Figure 6: **Comparative user study** with state-of-the-art *object-centric* 4D generation methods.

	X-CLIP ↑	Visual Quality ↑	Temporal Consistency↑	Dynamic Degree ↑	T-V Alignment ↑	Factual Consistency ↑
4Dfy [2]	20.03	1.43	1.49	3.05	2.26	1.30
Ours	24.23	2.43	2.17	3.15	2.91	2.49
Dream-in-4D [72]	19.52	1.34	1.37	3.02	2.27	1.20
Ours	24.77	2.41	2.15	3.14	2.89	2.46
AYG [29]	19.87	2.49	2.09	3.15	2.80	2.47
Ours	23.09	2.44	2.16	3.16	2.90	2.50

Table 1: **Quantitative comparison against baselines** using X-CLIP [37] and VideoScore [18].




Figure 7: **Qualitative comparison against baseline of 3D scene generation + 4D object-centric generation.** Our result is more natural in terms of object placement, motion and lighting.

In addition to existing methods, we also implement and compare with a straightforward baseline: combining 3D scene generation with 4D object-centric generation. Although straightforward, implementation with high quality is challenging due to: (1) Inserting 3D objects with realistic and physically correct placement is not trivial. Common issues include floating objects, misaligned scales between objects and background, and unrealistic scene layout. (2) Generating object motions that align well with the backgrounds is difficult, such as ensuring a generated target subject sits on a sofa rather than stands on it. (3) Complex procedures are needed to relight objects to match environment lighting. We tried our best to create a few baseline results with heavy manual efforts, as shown in Figure 7. We create the 3D background by removing the object from our generated freeze-time video. We then employ 4Dfy [2] to generate the foreground objects. Finally, we manually insert the objects into the background with the most plausible position and scale we can find. Despite these efforts, we find the inserted object appears disconnected from the background, especially when compared to our results.

On top of user study, we also provide quantitative results in Table 1 using employ the X-CLIP score [37], a minimal extension of the CLIP score for videos. We also run VideoScore [18], a video quality evaluation model trained using human feedback, which provides five scores assessing visual quality, temporal consistency, text-video alignment, and factual consistency. Our method significantly outperformed other methods in these metrics. However, we remain conservative about the effectiveness of these metrics since they have not been thoroughly studied and evaluated yet.

Ablation studies. We diagnose our system by removing individual components from the pipeline.

Per-frame deformation is crucial for reconstructing high-quality canonical 3D representations from noisy freeze-time videos. As illustrated in Figure 8, the absence of per-frame deformation results in significant artifacts, such as on the faces of the panda and the cat, and in blurrier areas like the boundaries of doors in a kitchen.

Small motion regularization is employed in our pipeline. While view-dependent deformation can improve underfitting, it can also lead to overfitting to the reference video, resulting in noisy reconstruction. We thus adopt a regularization loss of the magnitude of the deformation to mitigate the issues. As shown in Figure 9.



Figure 8: **Ablation study** of the impact of removing each component from the proposed pipeline.



Figure 9: **Ablation on per-view deformation and small motion regularization.** Removing per-view deformation results in underfitting the regions which are not geometrically consistent. Removing small motion regularization of the per-view deformation field causes overfitting to the freeze-time field, leading to noisy results.

Multi-view SDS aids in noise reduction and enhances shape realism. For example, as illustrated in Figure 8, it improves the depiction of the human head in the background, the upper boundary of the sofa, and the limbs of the goat.

Freeze-time video: Removing the step of generating freeze-time videos and the corresponding reconstruction loss $\mathcal{L}_{\text{recon}}$ will result in systematic failure. This is partly because we apply SDS only in the later stages of training, which is insufficient to correct significant reconstruction errors.

Joint multi-view and temporal SDS helps improve the sharpness of the final result during the reconstruction of temporal deformation. It also helps prevent erroneous artifacts, such as the floating spot on the edge of the cat’s ear and the dark shadows above the cat’s back, as shown in Figure 8.

We have included additional samples from the ablation study in the supplementary web page.

5 Discussion and Conclusion

We propose **4Real**, which, to our knowledge, is the first method capable of generating near-photorealistic 4D scenes from text inputs.

However, our method has **limitations**: 1) It inherits limitations from the underlying video generation model, such as video resolution, blurriness or artifacts during fast motion, particularly with thin or small objects, and occasional text-video misalignment. 2) Reconstructing from videos with dynamic content is challenging, and our method may fail due to inaccuracies in camera pose estimation, rapid movements, the sudden appearance and disappearance of objects, and abrupt lighting changes. 3) Our method does not produce high-quality geometry such as meshes, due to the limitations of using 3DGS. 4) It still takes over an hour to generate a 2-second 4D scene.

For **future works**, we envision the availability of stronger video generation models with more accurate camera pose and object motion control [6, 61], incorporating cross-frame attention [70] when generating freeze-time videos and utilizing feedforward 3D reconstruction [51] would help address these limitations.

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3d point clouds. In *IMCL*, 2018. 3
- [2] 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*, 2023. 2, 3, 7, 8, 9
- [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. 3, 4
- [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. 3
- [5] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. Video generation models as world simulators. 2024. URL <https://openai.com/research/video-generation-models-as-world-simulators>. 3
- [6] Changgu Chen, Junwei Shu, Lianggangxun Chen, Gaoqi He, Changbo Wang, and Yang Li. Motion-zero: Zero-shot moving object control framework for diffusion-based video generation. *arXiv preprint arXiv:2401.10150*, 2024. 10
- [7] Haoxin Chen, Yong Zhang, Xiaodong Cun, Menghan Xia, Xintao Wang, Chao Weng, and Ying Shan. Videocrafter2: Overcoming data limitations for high-quality video diffusion models, 2024. 4
- [8] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. *arXiv preprint arXiv:2303.13873*, 2023. 3
- [9] Ting Chen and Lala Li. Fit: Far-reaching interleaved transformers. *arXiv preprint arXiv:2305.12689*, 2023. 3, 7
- [10] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, 2019. 3
- [11] Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. In *CVPR*, 2023.
- [12] Zezhou Cheng, Menglei Chai, Jian Ren, Hsin-Ying Lee, Kyle Olszewski, Zeng Huang, Subhansu Maji, and Sergey Tulyakov. Cross-modal 3d shape generation and manipulation. In *ECCV*, 2022. 3
- [13] 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. 3
- [14] Yilun Du, Yanan Zhang, Hong-Xing Yu, Joshua B Tenenbaum, and Jiajun Wu. Neural radiance flow for 4d view synthesis and video processing. In *ICCV*, 2021. 3
- [15] Jiemin Fang, Taoran Yi, Xinggang Wang, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Matthias Nießner, and Qi Tian. Fast dynamic radiance fields with time-aware neural voxels. In *SIGGRAPH Asia*, 2022. 3
- [16] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *ICCV*, 2021. 3
- [17] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui 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. 3
- [18] Xuan He, Dongfu Jiang, Ge Zhang, Max Ku, Achint Soni, Sherman Siu, Haonan Chen, Abhramil Chandra, Ziyang Jiang, Aaran Arulraj, et al. Mantisscore: Building automatic metrics to simulate fine-grained human feedback for video generation. *arXiv preprint arXiv:2406.15252*, 2024. 9
- [19] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. 3
- [20] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. In *NeurIPS*, 2022. 3

- [21] 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. 3
- [22] Heewoo Jun and Alex Nichol. Shap-e: Generating conditional 3d implicit functions. *arXiv preprint arXiv:2305.02463*, 2023. 3
- [23] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM TOG*, 2023. 3, 4, 16, 17
- [24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 17
- [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. 3
- [26] Zhengqi Li, Simon Niklaus, Noah Snively, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *CVPR*, 2021. 3
- [27] 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
- [28] Chieh Hubert Lin, Hsin-Ying Lee, Willi Menapace, Menglei Chai, Aliaksandr Siarohin, Ming-Hsuan Yang, and Sergey Tulyakov. Infinicity: Infinite-scale city synthesis. In *ICCV*, 2023. 3
- [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. *arXiv preprint arXiv:2312.13763*, 2023. 3, 8, 9
- [30] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *ICCV*, 2023. 3
- [31] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Generating multiview-consistent images from a single-view image. In *ICLR*, 2024.
- [32] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. *arXiv preprint arXiv:2310.15008*, 2023. 3
- [33] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. *arXiv preprint arXiv:2308.09713*, 2023. 3, 4, 7, 15
- [34] Aniruddha Mahapatra, Aliaksandr Siarohin, Hsin-Ying Lee, Sergey Tulyakov, and Jun-Yan Zhu. Text-guided synthesis of eulerian cinemagraphs. *ACM TOG*, 2023. 3
- [35] Willi Menapace, Aliaksandr Siarohin, Ivan Skorokhodov, Ekaterina Deyneka, Tsai-Shien Chen, Anil Kag, Yuwei Fang, Aleksei Stoliar, Elisa Ricci, Jian Ren, et al. Snap video: Scaled spatiotemporal transformers for text-to-video synthesis. In *CVPR*, 2024. 3, 4, 7
- [36] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 2021. 3
- [37] Bolin Ni, Houwen Peng, Minghao Chen, Songyang Zhang, Gaofeng Meng, Jianlong Fu, Shiming Xiang, and Haibin Ling. Expanding language-image pretrained models for general video recognition. In *ECCV*, 2022. 9
- [38] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022. 3
- [39] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *ICCV*, 2021. 3
- [40] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *ACM Trans. Graph.*, 2021. 3
- [41] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *ICCV*, 2023. 3

- [42] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *ICLR*, 2023. 3, 6
- [43] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *CVPR*, 2021. 3
- [44] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. In *ICLR*, 2024. 3
- [45] 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. 3, 8
- [46] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. 6, 17
- [47] Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. In *ICLR*, 2024. 3
- [48] 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. 3
- [49] 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. *arXiv preprint arXiv:2301.11280*, 2023. 3, 8
- [50] Edward J Smith and David Meger. Improved adversarial systems for 3d object generation and reconstruction. In *CoRL*, 2017. 3
- [51] Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. *arXiv preprint arXiv:2402.05054*, 2024. 3, 10
- [52] Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitry Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. Sv3d: Novel multi-view synthesis and 3d generation from a single image using latent video diffusion. *arXiv preprint arXiv:2403.12008*, 2024. 3, 4
- [53] Chaoyang Wang, Ben Eckart, Simon Lucey, and Orazio Gallo. Neural trajectory fields for dynamic novel view synthesis. *arXiv preprint arXiv:2105.05994*, 2021. 3
- [54] Chaoyang Wang, Lachlan Ewen MacDonald, László A. Jeni, and Simon Lucey. Flow supervision for deformable nerf. In *CVPR*, 2023. 3
- [55] Chaoyang Wang, Peiye Zhuang, Aliaksandr Siarohin, Junli Cao, Guocheng Qian, Hsin-Ying Lee, and Sergey Tulyakov. Diffusion priors for dynamic view synthesis from monocular videos. *arXiv preprint arXiv:2401.05583*, 2024. 3
- [56] 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. 3
- [57] Tengfei Wang, Bo Zhang, Ting Zhang, Shuyang Gu, Jianmin Bao, Tadas Baltrusaitis, Jingjing Shen, Dong Chen, Fang Wen, Qifeng Chen, et al. Rodin: A generative model for sculpting 3d digital avatars using diffusion. In *CVPR*, 2023. 3
- [58] 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 preprint arXiv:2305.16213*, 2023. 3
- [59] Zhengyi Wang, Yikai Wang, Yifei Chen, Chendong Xiang, Shuo Chen, Dajiang Yu, Chongxuan Li, Hang Su, and Jun Zhu. Crm: Single image to 3d textured mesh with convolutional reconstruction model. *arXiv preprint arXiv:2403.05034*, 2024. 3
- [60] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE TIP*, 2004. 7
- [61] Zhouxia Wang, Ziyang Yuan, Xintao Wang, Tianshui Chen, Menghan Xia, Ping Luo, and Yin Shan. Motionctrl: A unified and flexible motion controller for video generation. *arXiv preprint arXiv:2312.03641*, 2023. 10

- [62] 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. 3
- [63] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528*, 2023. 3, 4
- [64] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *CVPR*, 2021. 3
- [65] Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli, Gordon Wetzstein, Zexiang Xu, et al. Dmv3d: Denoising multi-view diffusion using 3d large reconstruction model. In *ICLR*, 2024. 3
- [66] Gengshan Yang, Chaoyang Wang, N. Dinesh Reddy, and Deva Ramanan. Reconstructing animatable categories from videos. In *CVPR*, 2023. 3
- [67] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. *arXiv preprint arXiv:2309.13101*, 2023. 3, 4, 17
- [68] Yuyang Yin, Dejia Xu, Zhangyang Wang, Yao Zhao, and Yunchao Wei. 4dgen: Grounded 4d content generation with spatial-temporal consistency. *arXiv preprint arXiv:2312.17225*, 2023. 3
- [69] Heng Yu, Joel Julin, Zoltán Á. Milacski, Koichiro Niinuma, and László A. Jeni. Cogs: Controllable gaussian splatting. In *CVPR*, 2024. 3, 4, 7
- [70] Yifei Zeng, Yanqin Jiang, Siyu Zhu, Yuanxun Lu, Youtian Lin, Hao Zhu, Weiming Hu, Xun Cao, and Yao Yao. Stag4d: Spatial-temporal anchored generative 4d gaussians. *arXiv preprint arXiv:2403.14939*, 2024. 10
- [71] Yabo Zhang, Yuxiang Wei, Dongsheng Jiang, Xiaopeng Zhang, Wangmeng Zuo, and Qi Tian. Controlvideo: Training-free controllable text-to-video generation. In *ICLR*, 2024. 3
- [72] Yufeng Zheng, Xueting Li, Koki Nagano, Sifei Liu, Otmar Hilliges, and Shalini De Mello. A unified approach for text- and image-guided 4d scene generation. In *CVPR*, 2024. 2, 3, 7, 8, 9
- [73] Joseph Zhu and Peiye Zhuang. Hifa: High-fidelity text-to-3d with advanced diffusion guidance. In *ICLR*, 2023. 3, 8
- [74] Peiye Zhuang, Songfang Han, Chaoyang Wang, Aliaksandr Siarohin, Jiaxu Zou, Michael Vasilkovsky, Vladislav Shakhrai, Sergey Korolev, Sergey Tulyakov, and Hsin-Ying Lee. Gtr: Improving large 3d reconstruction models through geometry and texture refinement. *arXiv preprint arXiv:2406.05649*, 2024. 3

A Masked Snap Video Model Details

We employ a masked variant of the Snap Video model that is trained to generate videos using both text and a set of known frames as conditioning. Given a set of known frames at arbitrary positions in the input sequence, the masked Snap Video model produces an output video whose frames at those positions match the known frames. This mechanism allows the model to perform various tasks such as image animation, frame interpolation and temporal extension of videos by providing known frames in the corresponding arrangements.

The masked model variant is obtained by extending the original Snap Video model's input channels by 3 to accommodate the conditioning frames input, and by training the model for an additional 200k steps using a masked video modeling objective, while keeping the other training parameters unchanged.

In particular, we adopt a set of masking strategies which define the set of known frames provided as conditioning during training, each of which is randomly applied to an input batch element with probability m :

- *Unconditional generation* $m = 0.6$. Mask the whole input video. This setting corresponds to the original Snap Video setup, where no known frame is present.
- *Bernoulli* $m = 0.3$. Mask each frames of the input video with probability $(1 - 1/16)$, leaving on average a single frame unmasked.
- *Frame interpolation* $m = 0.075$. Mask all video frames apart from a set of frames sampled at regular time intervals.
- *Video extension* $m = 0.075$. Mask the last N video frames.

B Deformation Regularization Details

Our motion deformation network predicts position offsets ($\Delta\mathbf{x}$), rotation offsets ($\Delta\mathbf{q}$), and scale offsets ($\Delta\mathbf{s}$) for each time step. To prevent unnecessary movements (e.g., in the background) and to ensure movement consistency, we apply an L_1 norm regularization as follows:

$$\mathcal{L}_{\text{norm}} = \frac{1}{N} \sum_{i=1}^N (\|\Delta\mathbf{x}_i\| + \|\Delta\mathbf{q}_i\| + \|\Delta\mathbf{s}_i\|). \quad (7)$$

We introduce a difference loss, denoted as $\mathcal{L}_{\text{diff}}$, to ensure that the movement or trajectory of each Gaussian is consistent and smooth over time. This loss penalizes abrupt changes in the trajectory, promoting smoother transitions. The difference loss is formulated as follows:

$$\mathcal{L}_{\text{diff}} = \frac{1}{N} \sum_{i=1}^N \|\Delta\mathbf{x}_{i,t} - \Delta\mathbf{x}_{i,t-1}\|. \quad (8)$$

Here, $\Delta\mathbf{x}_{i,t}$ represents the position offset of Gaussian i at time t . We also borrow the local-rigidity loss $\mathcal{L}_{\text{rigid}}$ and rotational loss \mathcal{L}_{rot} from [33] (For a more comprehensive understanding, please refer to their detailed paper). For the local-rigidity loss, a weighting scheme is applied that utilizes an unnormalized Gaussian weighting factor. This factor assigns different weights to each point based on its proximity to the center of the Gaussian distribution. The unnormalized Gaussian weighting factor is defined as follows:

$$w_{i,j} = \exp(-\lambda_w \|\mathbf{x}_{j,0} - \mathbf{x}_{i,0}\|_2^2) \quad (9)$$

Using this weighting scheme, a local-rigidity loss is employed, denoted as $\mathcal{L}_{\text{rigid}}$, defined as follows:

$$\mathcal{L}_{\text{rigid}}^{i,j} = w_{i,j} \|(\mathbf{x}_{j,t-1} - \mathbf{x}_{i,t-1}) - \mathbf{R}_{i,t-1} \mathbf{R}_{i,t}^{-1} (\mathbf{x}_{j,t} - \mathbf{x}_{i,t})\|_2, \quad (10)$$

$$\mathcal{L}_{\text{rigid}} = \frac{1}{k|G|} \sum_{i \in G} \sum_{j \in \text{knn}_{i,k}} \mathcal{L}_{\text{rigid}}^{i,j}. \quad (11)$$

This loss ensures that for each Gaussian i , neighboring Gaussians j move in a manner consistent with the rigid body transformation of the coordinate system over time. Specifically, this means that the relative positions and orientations of neighboring Gaussians are maintained as the coordinate system evolves, ensuring a coherent and physically plausible motion pattern.

Additionally, we incorporate a rotational loss \mathcal{L}^{rot} to ensure consistent rotations among neighboring Gaussians across different time steps. This loss function helps maintain the relative orientations of nearby Gaussians as they evolve, preserving the overall rotational coherence of the system. The rotational loss is expressed as:

$$\mathcal{L}_{\text{rot}} = \frac{1}{k|G|} \sum_{i \in G} \sum_{j \in \text{kn}_i, k} w_{i,j} \|\hat{\mathbf{q}}_{j,t} \hat{\mathbf{q}}_{j,t-1}^{-1} - \hat{\mathbf{q}}_{i,t} \hat{\mathbf{q}}_{i,t-1}^{-1}\|_2, \quad (12)$$

Here, $\hat{\mathbf{q}}$ represents the normalized quaternion rotation of each Gaussian. The k -nearest neighbors, identified in the same manner as in the preceding losses, are utilized to ensure consistent rotational behavior among neighboring Gaussians.

C Stage-wise Training

Our method consists of two main components: canonical space reconstruction and motion fitting.

1. Canonical Space Reconstruction

(a) Warm-up Stage:

- Train the network for 3,000 iterations without any per-frame deformations to establish a baseline for the canonical space. This stage follows the same method as described in [23].

(b) Learning Deformations and GS Optimization:

- Learn deformations for each frame and optimize the Gaussians over the course of 20,000 iterations.
- Implement GS growth (including splitting, cloning, and pruning) using the strategy described in [23] during the first 15,000 iterations.

(c) Multi-view SDS:

- Enable multi-view SDS in this stage from 15,000 iterations to help learn a good canonical space.

2. Motion Fitting

(a) Initial Motion Fitting:

- Fit the motion over the next 10,000 iterations using only reference views at different timesteps.

Note: Cease GS optimization to focus solely on learning motion deformations.

(b) GS Growth for Motion Compensation:

- Resume GS growth using the same strategy for an additional 5,000 iterations to address missing points in the canonical space caused by motion.

Note: Fix the learned canonical space from the previous step, and perform GS growth by focusing solely on the newly added Gaussians.

(c) Simultaneous Fine-Tuning:

- Fine-tune both the GS and the motion deformation network simultaneously using all training frames consisting of different camera views and timesteps.

(d) Joint Temporal and Multi-view SDS:

- Enable temporal and multi-view SDS jointly in this stage when doing fine-tuning from 35,000 iterations and last for 5,000 iterations to help learn a consistent motion.

D Implementation Details

D.1 Multi-view and Temporal SDS

During training with SDS, we alternately perform multi-view and temporal SDS. For each iteration, we randomly sample either a camera pose or a timestep. This sampling determines the specific spatial (view) dimension or the temporal (time) dimension to train. By alternating between these two types of SDS, we ensure comprehensive coverage of both spatial and temporal aspects. This strategy enhances the overall consistency and robustness of the model. For each SDS step, we randomly sample $M = 16$ frame poses and perform rendering, either following the sequence of camera movement or based on the order of time. It is important to note that we perform noisy virtual trajectory sampling during multi-view SDS. Specifically, we randomly sample a camera pose that is different from the reference pose used to generate the temporal video. Subsequently, we sample a virtual camera pose within a specified radius of a circle centered on the sampled pose. By interpolating between this virtual camera pose and the reference camera pose, we generate $M = 16$ frame poses for rendering. This approach ensures a diverse set of viewpoints, enhancing the robustness and accuracy of the multi-view SDS process. When inputting our high-resolution rendered images into the low-resolution model for SDS, we first apply one of two preprocessing steps with equal probability. The image is either randomly shifted by $S - 1$ pixels (where S is the downsampling scale) or randomly cropped to half its size. Following this, we downsample the image and input it into the low-resolution model. This preprocessing ensures that the low-resolution model receives varied inputs, enhancing the robustness of the training process.

D.2 GS Growth for Motion Fitting

During the GS growth for motion fitting, we fix the attributes (position, rotation, scale, opacity and color) of the learned canonical Gaussians by stopping the gradient flow to them. Instead, we focus on performing splitting, cloning, and pruning operations based on the gradients of the newly added Gaussians. These new Gaussians are introduced at the beginning of the growth stage, specifically according to the gradients derived from the canonical Gaussians. This method ensures that the canonical structure remains stable while allowing the newly added Gaussians to adapt and optimize the motion fitting process.

D.3 Deformation Network

For the training process, we adopted the differential Gaussian rasterization technique from 3D GS [23] and utilized the same deformation network structure as described in [67] for both per-frame deformation and motion deformation. Both deformation networks consist of an MLP \mathcal{F}_θ : $(\gamma(\mathbf{x}), \gamma(t)) \rightarrow (\Delta\mathbf{x}, \Delta\mathbf{q}, \Delta\mathbf{s})$, which applies position embedding γ on each coordinate of the 3D Gaussians and time (or per-frame deformation code) and maps them to their corresponding deviations in position, rotation, and scaling. The weights θ of the MLP are optimized during this mapping process. Each MLP \mathcal{F}_θ processes the input through eight fully connected layers, each employing ReLU activations and containing 256-dimensional hidden units, resulting in a 256-dimensional feature vector. This feature vector is then passed through three additional fully connected layers (without activation functions) to independently output the offsets for position, rotation, and scaling over time. It is worth noting that, similar to NeRF, the feature vector and the input are concatenated at the fourth layer, enhancing the network's ability to capture complex deformations.

D.4 Training Setup

The initialization of the Gaussians was based on Structure-from-Motion (SfM) results obtained from COLMAP [46]. We adhered to the training settings specified in [23] for the canonical GS training. The learning rate for the positions of the Gaussians was set to undergo exponential decay, starting from 1.6×10^{-4} and decreasing to 1.6×10^{-6} . Additionally, the learning rate for both scale and rotation parameters was maintained at 1×10^{-3} throughout the training process. The learning rate for the deformation network was set to exponentially decay from 1×10^{-3} to 1×10^{-5} . The optimization across these processes was conducted using the Adam optimizer [24] with β parameters set to (0.9, 0.999). All experiments were performed on single NVIDIA A100 GPUs with 80GB of memory.

E Key Hyperparameters

E.1 Loss Weighting

For the regularization terms mentioned above, we assign a weight of 0.01 to each within the overall loss function, ensuring they contribute appropriately to the optimization process. For the reconstruction loss, we use a weight of 1. Additionally, for the SDS loss, we differentiate the weights based on the type of SDS being applied: a weight of 20 for temporal SDS and a weight of 5 for multi-view SDS. These weights are chosen to effectively capture motion dynamics and stabilize the training process.

E.2 GS Growth Threshold

For the GS growth stage during the canonical space reconstruction, we set the opacity threshold to $\tau_\alpha = 5 \times 10^{-3}$ and the gradient threshold to $\tau_{grad} = 2 \times 10^{-4}$. These thresholds help control the addition of new Gaussians based on their opacity and gradient values. In contrast, for the GS growth stage during motion fitting, we use a relatively higher opacity threshold of $\tau_\alpha = 1 \times 10^{-2}$ to avoid introducing redundant Gaussians. This higher threshold ensures that only significant Gaussians are added, which aids in maintaining efficiency and relevance during the motion fitting process.

F User Study Details

The user study was conducted through a single survey comprising seven questions. Each question asked the evaluators to compare two videos: one rendered with our method and the other with a different method using the same text prompt as input. Evaluators indicated their preferred video based on various aspects, as illustrated in Fig. 10. The evaluators were given the following instructions for each metric.


- **Motion Realism:** Take into consideration of magnitude, smoothness, consistency of the motion. Pay close attention to deformed limbs of human and animals and unnatural deformations.
- **Foreground Photo-realism:** Assess realism of the appearance of the foreground object, do not consider the motion factor. That is the main object that looks more like taken with real video camera, pay attention to unnatural color or styles of the objects.
- **Background Photo-realism:** Assess realism of the appearance of the background. That is the background that looks more like taken with real video camera, pay attention to unnatural color or styles of the objects or bluriness.
- **3D Shape Realism:** That is the video in which main object has the most natural shape, again pay attention to deformed limbs of human and animals and unnatural deformations.
- **General Realism:** Please provide your subjective appraisal of which video, in your opinion, stands out as superior based on appearance quality, 3D structure quality, and motion quality.
- **Significance of Motion:** The video that contain the most amount of motion. Please exclude from consideration any random limbs deformations.
- **Video-text Alignment:** That is which video reflect all the aspects included in the text description.

You are shown a description of a video and two different videos generated by AI based on this description. Your task is to answer 7 questions regarding quality of these videos. Please pay close attention to instructions and answer as thoughtfully as you can. In each task one question is a test one on which acceptance of the task will be judged. Please exclude size of the video from any considerations and pay attention only to the video content.

- 1) Which video has more realistic motion? Take into consideration of magnitude, smoothness, consistency of the motion. Pay close attention to deformed limbs of human and animals and unnatural deformations.
- 2) Which video has the most photo-realistic foreground? Assess realism of the appearance of the foreground object, do not consider the motion factor. That is the main object that looks more like taken with real video camera, pay attention to unnatural color or styles of the objects.
- 3) Which video has the most photo-realistic background? Assess realism of the appearance of the background. That is the background that looks more like taken with real video camera, pay attention to unnatural color or styles of the objects or blurriness.
- 4) Which video has object of better more realist shape? That is the video in which main object has the most natural shape, again pay attention to deformed limbs of human and animals and unnatural deformations.
- 5) In general which video looks more realistic?
- 6) Which video is most dynamic? The video that contain the most amount of motion. Please exclude from consideration any random limbs deformations.
- 7) Which video is better following text description? That is which video reflect all the aspects included in the text description

Text Description: a humanoid robot playing the violin

Video A



Video B



Which video has more realistic motion?

☐ Option A

☐ Option B

Which video has the most photo-realistic foreground?

☐ Option A

☐ Option B

Which video has the most photo-realistic background?

☐ Option A

☐ Option B

Which video has object of better more realist shape?

☐ Option A

☐ Option B

In general which video looks more realistic?

☐ Option A

☐ Option B

Which video is most dynamic?

☐ Option A

☐ Option B

Which video is better following text description?

☐ Option A

☐ Option B

Figure 10: Screenshot of the user study webpage.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: We claim to propose a pipeline that facilitates the generation of dynamic scenes with enhanced photorealism and structural integrity, viewable from multiple perspectives. Such claim is empirically verified with our generated results.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: Limitations are thoroughly discussed in the discussion section.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This is not a theoretical paper.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We detailed each implementation steps in the main paper and appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: The implementation involves models that have not been open-sourced.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We specified the training details in the supplementary, and how the user study is conducted.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We conduct user study and report win rates between our method and competitions. The preference of our method ranges between 66.9 to 77.1 percentage across seven different categories.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We stated that our method runs with a single A100 GPU.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: There is no concern over code of ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Included at the end of conclusion section.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper does not release new models.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: The paper does not use existing assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [Yes]

Justification: We included detailed descriptions, and screenshot of instructions in the appendix.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [Yes]

Justification: The experiment is conducted with Amazon Turk.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.