
Vidu4D: Single Generated Video to High-Fidelity 4D Reconstruction with Dynamic Gaussian Surfels

Yikai Wang^{*1}, Xinzhou Wang^{*1,2,3}, Zilong Chen^{1,2}, Zhengyi Wang^{1,2}, Fuchun Sun¹, Jun Zhu^{†1,2}

¹Department of Computer Science and Technology, BNRist Center, Tsinghua University

²ShengShu ³College of Electronic and Information Engineering, Tongji University

yikaiw@outlook.com, wangxinzhou@tongji.edu.cn, dcszj@tsinghua.edu.cn

Abstract

Video generative models are receiving particular attention given their ability to generate realistic and imaginative frames. Besides, these models are also observed to exhibit strong 3D consistency, significantly enhancing their potential to act as world simulators. In this work, we present Vidu4D, a novel reconstruction model that excels in accurately reconstructing 4D (*i.e.*, sequential 3D) representations from single generated videos, addressing challenges associated with non-rigidity and frame distortion. This capability is pivotal for creating high-fidelity virtual contents that maintain both spatial and temporal coherence. At the core of Vidu4D is our proposed *Dynamic Gaussian Surfels* (DGS) technique. DGS optimizes time-varying warping functions to transform Gaussian surfels (surface elements) from a static state to a dynamically warped state. This transformation enables a precise depiction of motion and deformation over time. To preserve the structural integrity of surface-aligned Gaussian surfels, we design the warped-state geometric regularization based on continuous warping fields for estimating normals. Additionally, we learn refinements on rotation and scaling parameters of Gaussian surfels, which greatly alleviates texture flickering during the warping process and enhances the capture of fine-grained appearance details. Vidu4D also contains a novel initialization state that provides a proper start for the warping fields in DGS. Equipping Vidu4D with an existing video generative model, the overall framework demonstrates high-fidelity text-to-4D generation in both appearance and geometry. Project page: <https://vidu4d-dgs.github.io>.

1 Introduction


The field of multimodal generation exhibits significant advancements and holds great promise for various applications. Recently, video generative models have garnered attention for their remarkable capability to craft immersive and lifelike frames [4, 8]. These models produce visually stunning content while also exhibiting strong 3D consistency [15, 81], largely increasing their potential to simulate realistic environments.

Parallel to these developments, high-quality 4D reconstruction has made great strides [19, 58, 63, 94, 100]. This technique involves capturing and rendering detailed spatial and temporal information. When integrated with generative video technologies, 4D reconstruction potentially enables the creation of models that capture static scenes and dynamic sequences over time. This synthesis provides a more holistic representation of reality, which is crucial for applications such as virtual reality, scientific visualization, and embodied artificial intelligence.

*Equal contribution. †The corresponding author.




(a) Prompt: A portrait captures the dignified presence of an orange cat with striking blue eyes. The cat wears a single pearl earring. Her head tilts in contemplation, reminiscent of a Dutch cap.



(b) Prompt: A dragon with its hair blown by a strong wind. Devil enters the soul with ethereal landscapes.



(c) Prompt: Light painting photo of a cheetah, cinematic.



(d) Prompt: A goldfish seemingly swimming through the air.



(e) Prompt: A small, fluffy creature with an appearance reminiscent of a mythical being. The creature's fur texture is rendered in high detail. The monster's large eyes and open mouth express wonder and curiosity.



(f) Prompt: An isolated coloured abstract sculpture with a dali shape.

Figure 1: Text-(to-video)-to-4D samples generated by equipping Vidu4D with a pretrained video diffusion model [4]. For each sample, we exhibit per-frame 3D rendering for novel-view color, normal, and surfel feature. We observe that Vidu4D can reconstruct precisely detailed and photo-realistic 4D representation. See our accompanying videos in our [project page](#) for better visual quality.

However, achieving high-fidelity 4D reconstruction from generated videos poses great challenges. Non-rigidity and frame distortion are prevalent issues that can undermine the temporal and spatial coherence of the reconstructed content, thus complicating the creation of a seamless and coherent depiction of dynamic subjects.

In this work, we introduce Vidu4D, a novel reconstruction pipeline designed to accurately reconstruct 4D representations from single generated videos, facilitating the creation of 4D content with high precision in spatial and temporal coherence. Vidu4D contains two novel stages, namely, the initialization of non-rigid warping fields and Dynamic Gaussian Surfels (DGS), together enabling the reconstruction of high-fidelity 4D content with detailed appearance and accurate geometry.

Specifically, the proposed DGS optimizes non-rigid warping functions that transform Gaussian surfels from static to dynamically warped states. This dynamic transformation accurately represents motion and deformation over time, crucial for capturing realistic 4D representations. Besides, DGS demonstrates superior 4D reconstruction performance due to two other key aspects. Firstly, in terms of geometry, DGS adheres to Gaussian surfels principles [16, 28] to achieve precise geometric representation. Unlike existing methods, DGS incorporates warped-state normal consistency regularization to align surfels with actual surfaces with learnable continuous fields (*w.r.t.* spatial coordinate and time) to ensure smooth warping when estimating normals. Secondly, for appearance, DGS learns additional refinements on the rotation and scaling parameters of Gaussian surfels by a dual branch structure. This refinement reduces the flickering artifacts during warping and allows for the precise rendering of appearance details, resulting in high-quality reconstructed 4D representations.

By integrating Vidu4D with an existing powerful video generative model named Vidu [4], the overall framework demonstrates exceptional capabilities in text-to-4D generation. We provide 4D visualization results in Fig. 1. Extensive experiments based on the generated videos verify the effectiveness of our method compared to current state-of-the-art methods.

2 Related works

3D representation. Transforming 2D images into 3D representations has long been a central challenge in the field. Initially, triangle meshes were favored for their compactness and compatibility with rendering pipelines [9, 17, 67, 78, 82, 93]. However, the transition to more sophisticated volumetric methods was inevitable due to the limitations of surface-based approaches. Early volumetric representations included voxel grids [35, 48, 61, 72] and multi-plane images [20, 54, 74, 75, 80, 105], which, despite their straightforwardness, demanded intricate optimization strategies. The introduction of neural radiance fields (NeRF) [55] marked a great advancement, offering an implicit volumetric neural representation that could store and query the density and color of each point, leading to highly realistic reconstructions. The NeRF paradigm has since been improved upon in terms of reconstruction quality [5, 6, 33, 53, 92] and rendering [12, 23, 25, 27, 41, 45, 49, 64–66, 85, 102]. To address the limitations of NeRF, such as rendering speed and memory usage, recent work dubbed 3D Gaussian splatting (3DGS) [33] has proposed anisotropic Gaussian representations with GPU-optimized tile-based rasterization. This has opened up new avenues for surface extraction [24, 28], generation [14, 77, 96], and large-scale scene reconstruction [34, 46, 70], with 3DGS emerging as a universal representation for 3D scenes and objects. Gaussian surfels methods [16, 28] further exhibit advantages in modeling accurate geometry. While these methods have significantly advanced the field of static 3D representation, capturing the dynamic aspects of real-world scenes with non-rigid motion and deformation introduces a distinct set of challenges that demand innovative solutions.

Dynamic reconstruction and generation. The dynamic reconstruction of scenes from video captures presents a more complex challenge than static reconstruction, necessitating the capture of non-rigid motion and deformation over time [30, 37, 60, 76, 88]. Traditional methods have explored dynamic reconstruction using synchronized multi-view videos [1, 3, 11, 36, 48, 59, 73, 83, 84, 86, 89] or have focused on specific dynamic elements like humans or animals. More recently, there has been a shift towards reconstructing non-rigid objects from monocular videos, which is a more practical yet challenging scenario. One approach involves incorporating time as an additional input to the neural radiance field [11, 38, 68, 98], allowing for explicit querying of spatiotemporal information. Another line of research decomposes the spatiotemporal radiance field into a canonical space and a deformation field, representing spatial attributes and their temporal variations [18, 19, 19, 21, 22, 31, 39, 44, 47, 57, 58, 63, 69, 79, 95, 104]. With advancements in 3DGS, deformable-GS [100]




Figure 2: Illustration of the pipeline of Vidu4D, including the initialization stage and the DGS stage.

and 4DGS [94] have been developed, utilizing neural deformation fields with multi-layer perception (MLP) and triplane, respectively. SCGS [29] and dynamic 3D Gaussians [52] also advance the field by modeling time-varying scenes. Building on these advances, our work introduces dynamic Gaussian surfels, a novel extension of Gaussian representations that enhances the quality of both appearance and surface reconstruction under dynamic scenarios. A concurrent work DGM [43] builds time-consistent meshes from a monocular video with 3D Gaussian Splatting. In the realm of 3D or 4D generation, our approach diverges from recent progress in optimization-based [2, 13, 14, 40, 42, 62, 71, 88, 90], feed-forward [26, 91, 106], and multi-view reconstruction methods [15, 50, 51] by leveraging a video generative model to achieve generation capabilities. Our primary focus is on preserving high-quality appearance and geometrical integrity from generated videos. This results in a generation process that not only captures the nuances of motion and deformation but also maintains the high standards of realism and detail that are essential for creating immersive and lifelike virtual 3D representations.

3 Method

We start by introducing the problem definition for 4D reconstruction in Sec. 3.1. Following that, we introduce our Vidu4D which encompasses two novel stages. The first stage is designed to learn Dynamic Gaussian Surfels (DGS), ensuring precise representation of both visual appearance and geometric structure during the non-rigid reconstruction process, as detailed in Sec. 3.2. The second stage focuses on establishing the initial non-rigid warping fields of DGS, as detailed in Sec. 3.3.

3.1 Problem Definition

When given a single sequence of RGB video with T frames, the goal of 4D reconstruction is to determine a sequential 3D representation that could be rendered to fit each video frame as much as possible. Specifically, suppose the 3D representation for the t -th frame (termed as time t) is parameterized by θ_t , where $t = 1, \dots, T$. Given a differentiable rendering mapping \mathbf{g} , we could obtain the rendered color at the frame pixel $\bar{\mathbf{x}}^t \in \mathbb{R}^2$. We choose volume rendering as commonly adopted in NeRF [55], Gaussian Splatting [33], and Gaussian Surfels [16, 28]. The optimization of 4D reconstruction can be implemented by minimizing the empirical loss as

$$\min_{\theta} \frac{1}{T} \sum_{t=1}^T \sum_{\bar{\mathbf{x}}^t} \mathcal{L}(\mathbf{c}(\bar{\mathbf{x}}^t) = \mathbf{g}(\theta_t, \{\mathbf{x}_i^t\}_{i=1, \dots, N}), \hat{\mathbf{c}}(\bar{\mathbf{x}}^t)), \quad (1)$$

where $\mathbf{x}_i^t \in \mathbb{R}^3$ is the i -th 3D point sampled or intersected with Gaussian primitives along the ray that emanates from the frame pixel $\bar{\mathbf{x}}^t$; N is the number of sampled or intersected points per ray; $\mathbf{c}(\bar{\mathbf{x}}^t)$ and $\hat{\mathbf{c}}(\bar{\mathbf{x}}^t)$ are the rendered color and the observed color at $\bar{\mathbf{x}}^t$, respectively.

In the following, we detail the proposed **Vidu4D**, a reconstruction pipeline comprising two key stages as illustrated in Fig. 2, including a field initialization stage and a DGS stage.

3.2 Dynamic Gaussian Surfels

By optimizing Eq. (1), essentially our goal is to build a sequential 3D representation that could deform to be consistent with each 2D frame. We first start by considering an ideal video exhibiting different views of the same static object without object deformation, movement, or video distortion. To model the 3D representation with high appearance fidelity and geometry accuracy, we follow the method of using differentiable 2D Gaussian primitives as proposed by recent Gaussian Surfels advances [16, 28].

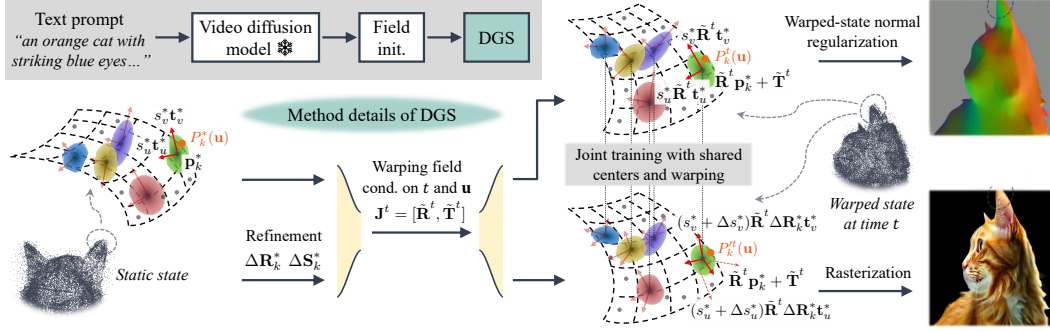


Figure 3: Illustration of the overall framework and our DGS in detail. For DGS, Gaussian surfels in the static state are transformed to the warped state by learning non-rigid warping functions conditioned on time t and coordinate \mathbf{u} . We incorporate warped-state normal regularization for accurate geometry, and refined rotation and scaling matrices of Gaussian surfels for detailed appearance. Both branches in the warped state, including with and without refinement, share the same centers of Gaussian surfels and the same warping functions. “Field init.” stands for field initialization as introduced in Sec. 3.3.

Specifically, the k -th Gaussian surfel (of the total K) is characterized by a central point $\mathbf{p}_k^* \in \mathbb{R}^3$ and a local coordinate system centered at \mathbf{p}_k^* with two principal tangential vectors $\mathbf{t}_u^* \in \mathbb{R}^{3 \times 1}$, $\mathbf{t}_v^* \in \mathbb{R}^{3 \times 1}$ and scaling factors $s_u^* \in \mathbb{R}$, $s_v^* \in \mathbb{R}$. Here, we use the notation “*” to represent parameters in the static state. A Gaussian surfel is computed as a 2D Gaussian defined in a local tangent plane in the world space. Following [28], for any point $\mathbf{u} = (u, v)$ located on the uv coordinate system centered at \mathbf{p}_k^* , its coordinate in the world space, denoted as $P_k^*(\mathbf{u}) \in \mathbb{R}^{3 \times 1}$, is computed by

$$P_k^*(\mathbf{u}) = \mathbf{p}_k^* + s_u^* \mathbf{t}_u^* u + s_v^* \mathbf{t}_v^* v = [\mathbf{R}_k^* \mathbf{S}_k^* \quad \mathbf{p}_k^*] (u, v, 1, 1)^\top, \quad (2)$$

where $\mathbf{R}_k^* = [\mathbf{t}_u^*, \mathbf{t}_v^*, \mathbf{t}_u^* \times \mathbf{t}_v^*] \in \text{SO}(3)$ denotes the rotation matrix, and the diagonal matrix $\mathbf{S}_k^* = \text{diag}(s_u^*, s_v^*, 0) \in \mathbb{R}^{3 \times 3}$ denotes the scaling matrix.

In this work, our focus is on 4D reconstruction from a single generated video, which may exhibit large non-rigidity, distortion, or illumination changes. We introduce **Dynamic Gaussian Surfels (DGS)**, a method designed to achieve precise 4D reconstruction while accommodating non-rigidity and other time-varying effects.

Motivated by recent advancements in non-rigid reconstruction methods [57, 88, 98], we aim to ensure that the target object maintains a consistent static state across different frames, thereby mitigating non-rigidity and distortion effects. To achieve this, we employ warping techniques on each Gaussian surfel represented by $P_k^*(\mathbf{u})$, transforming them into a corresponding Gaussian surfel $P_k^t(\mathbf{u})$ at time t , which is centered at $\mathbf{p}_k^t \in \mathbb{R}^3$ with a rotation matrix $\mathbf{R}_k^t \in \text{SO}(3)$ and a scaling matrix $\mathbf{S}_k^t \in \mathbb{R}^{3 \times 3}$.

Non-rigid warping for Gaussian surfels. We now build the warping process from the static state to the warped state. We leverage a non-rigid warping function with B bones as key points to ease the training of deformation. In the static state, the b -th bone is represented by 3D Gaussian ellipsoids [97], with more details provided in the Appendix. We let $\mathbf{J}_b^t \in \text{SE}(3)$ represent a rigid transformation that moves the b -th bone from its static state to the warped state at time t . In effect, \mathbf{J}_b^t is achieved by non-linear mappings using a multi-layer perception (MLP) with $\text{SE}(3)$ guaranteed, as will be given later in Eq. (5). The non-rigid warping function can be written as the weighted combination of \mathbf{J}_b^t , where we apply dual quaternion blend skinning (DQB) [32] to ensure valid $\text{SE}(3)$ after combination,

$$\mathbf{J}^t = \mathcal{R} \left(\sum_{b=1}^B w_b^t \mathcal{Q}(\mathbf{J}_b^t) \right), \quad (3)$$

where w_b^t is the b -th element of the skinning weight vector $\mathbf{w}^t \in \mathbb{R}^{B \times 1}$, as detailed in the Appendix; \mathcal{Q} and \mathcal{R} denote the quaternion process and the inverse quaternion process, respectively. In this case, there is $\mathbf{J}^t \in \text{SE}(3)$.

We therefore rewrite the warping as $\mathbf{J}^t = [\tilde{\mathbf{R}}^t, \tilde{\mathbf{T}}^t]$ with the rotation $\tilde{\mathbf{R}}^t \in \text{SO}(3)$ and translation $\tilde{\mathbf{T}}^t \in \mathbb{R}^3$, and apply the corresponding transformation to Eq. (2) by

$$P_k^t(\mathbf{u}) = \mathbf{J}^t P_k^*(\mathbf{u}) = [\tilde{\mathbf{R}}^t \mathbf{R}_k^* \mathbf{S}_k^* \quad \tilde{\mathbf{R}}^t \mathbf{p}_k^* + \tilde{\mathbf{T}}^t] (u, v, 1, 1)^\top. \quad (4)$$

Note that Eq. (4) holds for any given point $P_k^*(\mathbf{u})$ including the center point of the k -th Gaussian surfel (*i.e.*, \mathbf{p}_k^*) when $\mathbf{u} = (0, 0)$. By deriving Eq. (4), we enable connection of the warping function *w.r.t.* to any point $\mathbf{u} = (u, v)$ on the local coordinate system centered at \mathbf{p}_k^* , which is needed later in Eq. (8) where \mathbf{u} is an intersection with Gaussian surfels and a ray that emanates from the frame pixel.

Warped-state normal regularization. To accurately capture the geometric representation, we follow similar methods in Gaussian Surfels [16, 28] to add normal consistency regularization which encourages all Gaussian surfels to be locally aligned with the actual surfaces. Differently, unlike 3D reconstruction for static scenes, 4D reconstruction commonly faces non-rigidity and distortion. Thus simply performing regularization to promote surface-aligned Gaussian surfels like previous methods harms the structural integrity due to the non-rigid warping.

We therefore design a warped-state normal regularization. As mentioned, each point $P_k^t(\mathbf{u})$ in the warped state at time t is transformed from its corresponding static point $P_k^*(\mathbf{u})$ based on the warping function in Eq. (4), namely, $P_k^t(\mathbf{u}) = \mathbf{J}^t P_k^*(\mathbf{u})$ with \mathbf{J}^t composed by \mathbf{J}_b^t . To maintain the structural integrity to a large extent when regularizing normal, we design \mathbf{J}_b^t as a continuous field that takes both the point $P_k^*(\mathbf{u})$ (or equivalently, \mathbf{u} in the local coordinate system) and the time t as conditions. By this setting, \mathbf{J}_b^t is expected to change continuously with the change of \mathbf{u} or t . We implement the continuous field by using a NeRF-style MLP which directly outputs a 6-dimensional dual quaternion, and rely on the inverse quaternion process \mathcal{R} to guarantee $\text{SE}(3)$, *i.e.*,

$$\mathbf{J}_b^t = \mathcal{R}(\text{MLP}(\gamma_b^t; \mathbf{u}, t)), \quad (5)$$

where γ_b^t is a learnable latent code for encoding the b -th bone at time t ; both \mathbf{u} and t are sent to the MLP as conditions to obtain \mathbf{J}_b^t . Thus \mathbf{J}^t is also expected to be continuous *w.r.t.* \mathbf{u} and t .

Based on the above design, the normal consistency loss at time t is obtained similar to [28],

$$\mathcal{L}_n = \sum_k \omega_k (1 - \mathbf{n}_k^\top \mathbf{N}^t), \quad \mathbf{N}^t(x, y) = \frac{\nabla_x \mathbf{p}^t \times \nabla_y \mathbf{p}^t}{|\nabla_x \mathbf{p}^t \times \nabla_y \mathbf{p}^t|}, \quad (6)$$

where k indexes over intersected surfels along the ray that emanates from the frame pixel $\bar{\mathbf{x}}$; $\omega_k = \alpha_k \mathcal{G}_k(\mathbf{u}(\bar{\mathbf{x}})) \prod_{j=1}^{k-1} (1 - \alpha_j \mathcal{G}_j(\mathbf{u}(\bar{\mathbf{x}})))$ denotes the blending weight of the intersection point; \mathbf{n}_k represents the normal of the surfel that is oriented towards the camera; \mathbf{N}^t , computed with finite differences, is the surface normal estimated by the nearby depth point \mathbf{p}^t at warped state time t .

In summary, by learning a continuous warping field and aligning the surfel normal with the estimated surface normal in the warped state, we ensure that all Gaussian surfels locally approximate the actual object surface without being noticeably impaired by the non-rigid warping.

Dual branch structure with refinement. To further achieve fine-grained appearance and reduce the texture flickering during warping, we propose to learn refinement terms for adjusting the rotation matrices \mathbf{R}_k^* and scaling matrices \mathbf{S}_k^* (defined in Eq. (2)) in the static state. We suppose the refinement terms are $\Delta \mathbf{R}_k^* \in \text{SO}(3)$ and $\Delta \mathbf{S}_k^* \in \mathbb{R}^{3 \times 3}$, respectively. Note that the third-axis of $\Delta \mathbf{S}_k^*$ is no longer necessarily 0. During refinement, we remain the center points \mathbf{p}_k^* and the warping \mathbf{J}^t (*i.e.*, including both $\tilde{\mathbf{R}}^t$ and $\tilde{\mathbf{T}}^t$) to be unchanged. The new warped process is formulated as,

$$P_k^t(\mathbf{u}) = [\tilde{\mathbf{R}}^t (\Delta \mathbf{R}_k^* \mathbf{R}_k^*) (\mathbf{S}_k^* + \Delta \mathbf{S}_k^*) \quad \tilde{\mathbf{R}}^t \mathbf{p}_k^* + \tilde{\mathbf{T}}^t] (u, v, 1, 1)^\top. \quad (7)$$

During the training of DGS, we maintain two branches including one with refinement and one without. In the warped state, both branches are jointly trained with shared warping functions and centers of Gaussian primitives². Due to the involvement of $\Delta \mathbf{R}_k^*$ and $\Delta \mathbf{S}_k^*$, both branches have different rotation and scaling matrices of Gaussian primitives.

Rasterization. Given a frame pixel $\bar{\mathbf{x}}$ and a camera ray that emanates from $\bar{\mathbf{x}}$, following the static-state methods to calculate intersection coordinates with Gaussian primitives along the ray [28, 33], we could obtain warped-state intersection coordinates based on Eq. (4) and Eq. (7). We then perform the volume rendering process [28] that integrates alpha-weighted appearance along the ray by

$$\mathbf{c}(\bar{\mathbf{x}}) = \sum_k \mathbf{c}_k \alpha_k \mathcal{G}_k(\mathbf{u}(\bar{\mathbf{x}})) \prod_{j=1}^{k-1} (1 - \alpha_j \mathcal{G}_j(\mathbf{u}(\bar{\mathbf{x}}))), \quad (8)$$

²Here, since the third-axis of the refined scaling matrix is not necessarily 0, we adopt ‘‘Gaussian primitive’’ for commonly referring to both Gaussian surfel and the refined Gaussian.

where k indexes over intersected Gaussian primitives along the ray that emanates from the frame pixel $\bar{\mathbf{x}}$; α_k and c_k denote the opacity and view-dependent appearance parameterized with spherical harmonics of the k -th Gaussian surfel, respectively; $\mathcal{G}_k(\mathbf{u}(\bar{\mathbf{x}})) = \exp\left(-\frac{u^2+v^2}{2}\right)$ corresponds to the k -th intersection point $\mathbf{u}(\bar{\mathbf{x}})$ which could be directly calculated when given $P_k^t(\mathbf{u})$ or $P_k^{tt}(\mathbf{u})$ and the corresponding local coordinate system. During implementation, $\mathcal{G}_k(\mathbf{u}(\bar{\mathbf{x}}))$ is further applied a low-pass filter following [7, 28].

A detailed architecture of DGS is depicted in Fig. 3. Important symbols are summarized in our Appendix.

3.3 Field Initialization

Given that the camera trajectory of generated videos is unknown, SfM methods like COLMAP struggle to converge due to rigidity violations. Additionally, since the background of generated videos appears to exhibit soft deformation or flickering colors, proper estimation of camera/body poses through background SfM is hindered. These challenges often result in very few successful registrations, as demonstrated in previous monocular 4D reconstruction tasks [98].

To address this, we design an implicit field before performing DGS to initialize the camera poses and establish the continuous warping field in Eq. (5). In this part, we propose the **field initialization** as another key component of our pipeline to initialize the continuous warping field of DGS for fast and stable convergence, as detailed below.

Initially, we train a neural Signed Distance Function (SDF) model [87], leveraging the same warping structure with bones as utilized in DGS. While DGS transforms Gaussian surfels from the static state to the warped state for rasterization, the neural SDF reverses this process, mapping points along camera rays from the warped state back to the static state. For the neural SDF component, we optimize the reverse warping process and deduce the forward warping as its inverse by minimizing a cycle loss, inspired by [10, 98]. Subsequently, we initialize $\text{MLP}(\cdot)$ in Eq. (5) to obtain warping functions \mathbf{J}_b^f by the network weights learned by the neural SDF part.

During the rendering of the neural SDF, we perform backward warping on the warped-state sampling points to the static state,

$$\mathbf{J}^{f,-1} = \mathcal{R}\left(\sum_{b=1}^B w_b^f \mathcal{Q}(\mathbf{J}_b^f)^{-1}\right), \quad \mathbf{X}^f = \mathbf{J}^{f,-1} \mathbf{X}^*, \quad (9)$$

which is an inversion of Eq. (3). By querying the SDF with a sample point \mathbf{X}^* in the static state, we render RGB and compute the photometric loss to optimize the SDF and the warping field defined in Eq. (5).

Nevertheless, there are two discrepancies between the neural SDF warping and DGS warping. Firstly, sampling points of the neural SDF are distributed in the frustum of the camera, while sampling points of DGS are distributed on the object surface. Additionally, we train the inversion of during initialization, while we utilize the non-inverse ones in DGS. To resolve the distribution gap and ensure that faithfully models the forward warping, we add a cycle loss,

$$\mathcal{L}_{\text{cyc}} = \|\mathbf{J}^f(\mathbf{J}^{f,-1}(\mathbf{X}^f)) - \mathbf{X}^f\|_2^2, \quad (10)$$

where \mathbf{X}^f could be either a mesh surface point or a sample point on the camera ray.

After initialization, we extract the canonical space mesh using marching cubes and initialize Gaussian surfels on it. We set the spherical harmonic in 0-th order to the RGB value of the nearest vertices. The warping field and learned camera poses are retained.

With the field initialization before DGS, our Vidu4D is capable of performing a text-(to-video)-to-4D generation task with the integration of existing video diffusion models.

4 Experiment

In this section, we provide an extensive evaluation of our method DGS (Sec. 3.2) with the initialization in Sec. 3.3, comparing both appearance and geometry against previous state-of-the-art methods. Additionally, we analyze the contributions of each proposed component in detail.

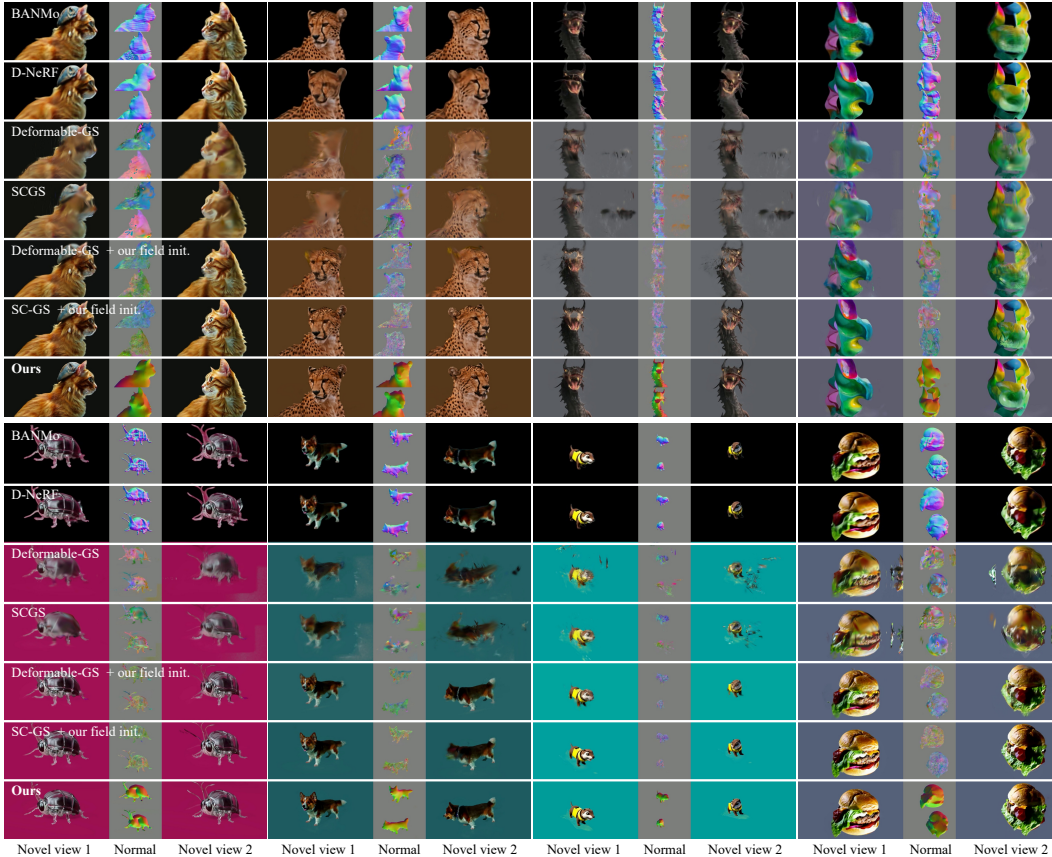


Figure 4: Novel-view qualitative evaluation compared with SOTA methods including NeRF-based methods (BANMo [98] and D-NeRF [63]) and Gaussian splatting-based methods (Deformable-GS [100] and SCGS [29]). We also provide our learned camera poses to baseline approaches for a fair comparison. These variants are denoted as “w. Poses”. Best view in color and zoom in.

4.1 Implementation

For all qualitative and quantitative experiments, we follow the standard pipeline for dynamic reconstruction [58], to construct our evaluation setup by selecting every fourth frame as a training frame and designating the middle frame between each pair of training frames as a validation frame.

Our model configuration involves several key parameters to balance reconstruction and regularization losses. For the field initialization stage, we use a similar architecture with 8 layers for volume rendering as in NeRF [55], and initialize MLP for predicting SDF as an approximate unit sphere [101]. We obtain a neural SDF, a warping field, and camera poses after this stage. For the DGS stage, we initialize centers of the Gaussian surfels with the sampled surface points extracted from the neural SDF, and initialize the warping field by the forward field from the first stage. The dimension of the latent code embedding γ_b^t is set as 128. Following BANMo [98], we adopt 25 bones to optimize skinning weights. For each reconstruction, the overall training takes over 1 hour on an A800 GPU.

4.2 Qualitative Evaluation

In the qualitative evaluation, we visually compare the novel-view reconstructions produced by our DGS against those generated by other state-of-the-art models, as illustrated in Fig. 4. Our evaluation focuses on several key aspects including detail preservation, texture quality, and geometric accuracy. Compared to methods based on implicit fields, the integration of Gaussian in our approach facilitates the rendering of highly detailed textures. Additionally, benefiting from a more geometry-aware representation, our method produces superior normal maps compared to those purely Gaussian-based methods. This also enhances the robustness of our method against artifacts of the generated videos

Table 1: Novel-view quantitative results on generated videos. Evaluation metrics are PSNR, SSIM, and LPIPS. We report results on three single videos and the averaged results over 30 single videos.

	Cat			Cheetah			Dragon			Average over 30 videos		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
BANMo [98]	15.10	0.6514	0.2575	13.15	0.5921	0.3241	18.48	0.6423	0.3500	13.62 \pm 2.99	0.6153 \pm 0.0714	0.3738 \pm 0.0665
D-NeRF [63]	15.15	0.6537	0.2657	13.21	0.5930	0.3344	18.53	0.6489	0.3527	21.01 \pm 2.86	0.8519 \pm 0.0717	0.1522 \pm 0.0754
Deformable-GS [100]	19.09	0.7815	0.2434	20.35	0.8039	0.1982	24.19	0.9100	0.0992	13.22 \pm 3.42	0.5934 \pm 0.0535	0.3749 \pm 0.0763
SCGS [29]	19.46	0.7867	0.2405	20.87	0.8123	0.1919	24.03	0.9083	0.1009	21.17 \pm 2.69	0.8547 \pm 0.0691	0.1504 \pm 0.0737
Deformable-GS + our field init.	21.94	0.8123	0.1816	22.41	0.8200	0.1687	26.05	0.9218	0.0894	22.63 \pm 2.14	0.8469 \pm 0.0438	0.1452 \pm 0.0354
SCGS + our field init.	23.25	0.8268	0.1574	23.70	0.8338	0.1497	28.40	0.9375	0.0686	24.75 \pm 2.11	0.8680 \pm 0.0440	0.1201 \pm 0.0359
Ours	24.63	0.8432	0.1559	25.68	0.8843	0.1117	28.58	0.9392	0.0618	27.30 \pm 2.66	0.9152 \pm 0.0602	0.0877 \pm 0.0564

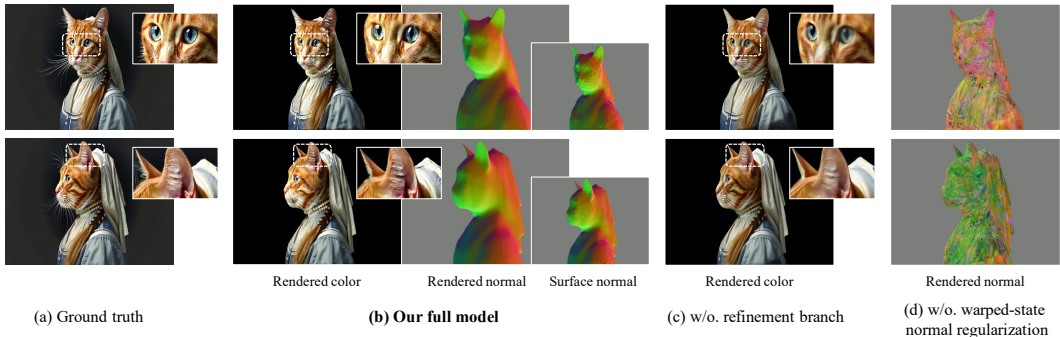


Figure 5: Ablation studies on the geometric regularization and refinement strategy. For our full model shown in (b), we provide our rendered color, rendered normal, and surface normal (estimated from the depth points for regularization). Additionally, for comparison, we visualize the rendered color for the case without refinements in (c) and the rendered normal for the case without warped-state normal regularization in (d), respectively. We showcase our model’s fidelity with close-ups.

like occlusions. For instance, in the third clip of the series, which features a dragon shrouded in fog, both SCGS and Deformable-GS methods tend to overfit and subsequently show a decline in performance. In contrast, our method consistently delivers superior results.

4.3 Quantitative Evaluation

We provide the quantitative evaluation comparing our method with state-of-the-art works in Table 1. Metrics include Peak Signal-to-Noise Ratio (PSNR) to evaluate the fidelity of the reconstructed textures, Structural Similarity Index (SSIM) for the quality evaluation, and LPIPS [103] as a perceptual metric. Our method exhibits superiority over all baseline methods, even with our learned poses, *e.g.*, ~ 2.5 PSNR increase over SCGS with poses for the averaged results.

4.4 Ablations

To understand the contributions of each component in Vidu4D, especially DGS, we conduct ablation studies in this section. We remove or alter specific elements of our model and observe the resulting performance changes in both appearance and geometry reconstruction.

Geometric regularization. We evaluate the impact of warped-state normal regularization by disabling it during training. From Fig. 5(b)(d), we observe that when removing the regularization, there is an obvious degradation in the structural integrity of surface-aligned Gaussian surfels, leading to noticeable inconsistency in the reconstructed 4D models.

Refinement strategy. We examine the effect of omitting refinements by keeping one branch (the concept of branches could be better visualized in Fig. 3) during training, shown in Fig. 5(b)(c). The performance indicates that removing refinements increases the loss of fine-grained appearance details. Additionally, we also find that refinements are crucial for mitigating the texture flickering issue.

Additional ablations. Please refer to the Appendix for additional ablation studies that detail the effectiveness of our refinement strategy and field initialization.

5 Conclusion

We introduce Vidu4D as a novel reconstruction model to achieve high-fidelity 4D representations from single generated videos. Vidu4D is powerful with our proposed DGS which builds the non-rigid warping field to transform Gaussian surfels, ensuring precise capture of motion and deformation over time. DGS also introduces key innovations that greatly enhance the accuracy and fidelity of 4D reconstruction, including dual branch refinement and warped-state geometric regularization. Our experiments demonstrate that Vidu4D outperforms existing methods in both quantitative and qualitative evaluations, highlighting its superiority in generating realistic and immersive 4D content.

Limitations and broader impact. While Vidu4D with DGS presents a significant performance in 4D reconstruction, currently there are still limitations such as the reliance on video quality, scalability challenges for large scenes, and computational difficulties in real-time applications. Additionally, when equipping Vidu4D with generative models, as with any generative technology, there is a risk of producing deceptive content which needs more caution.

6 Acknowledgement

This work was partly supported by the NSFC Projects (Nos. 62350080, 62306163, 62276149, 92370124, 92248303, U2341228, 62061136001, 62076147), BNRist (BNR2022RC01006), Tsinghua Institute for Guo Qiang, and the High Performance Computing Center, Tsinghua University. J. Zhu was also supported by the XPlorer Prize. Y.K. Wang was also supported by the China National Postdoctoral Program (No. 2023M741951) and Major Project of the New Generation of Artificial Intelligence (No. 2018AAA0102900).

References

- [1] Attal, B., Huang, J.B., Richardt, C., Zollhoefer, M., Kopf, J., O’Toole, M., Kim, C.: Hyperreel: High-fidelity 6-dof video with ray-conditioned sampling. arXiv preprint arXiv:2301.02238 (2023)
- [2] Bahmani, S., Skorokhodov, I., Rong, V., Wetzstein, G., Guibas, L., Wonka, P., Tulyakov, S., Park, J.J., Tagliasacchi, A., Lindell, D.B.: 4d-fy: Text-to-4d generation using hybrid score distillation sampling. In: CVPR (2024)
- [3] Bansal, A., Vo, M., Sheikh, Y., Ramanan, D., Narasimhan, S.: 4d visualization of dynamic events from unconstrained multi-view videos. In: CVPR (2020)
- [4] Bao, F., Xiang, C., Yue, G., He, G., Zhu, H., Zheng, K., Zhao, M., Liu, S., Wang, Y., Zhu, J.: Vidu: a highly consistent, dynamic and skilled text-to-video generator with diffusion models. arXiv preprint arXiv:2405.04233 (2024)
- [5] Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: ICCV (2021)
- [6] Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Zip-nerf: Anti-aliased grid-based neural radiance fields. arXiv preprint arXiv:2304.06706 (2023)
- [7] Botsch, M., Hornung, A., Zwicker, M., Kobbelt, L.: High-quality surface splatting on today’s gpus. In: Proceedings Eurographics/IEEE VGTC Symposium Point-Based Graphics (2005)
- [8] Brooks, T., Peebles, B., Holmes, C., DePue, W., Guo, Y., Jing, L., Schnurr, D., Taylor, J., Luhman, T., Luhman, E., Ng, C., Wang, R., Ramesh, A.: Video generation models as world simulators (2024), <https://openai.com/research/video-generation-models-as-world-simulators>
- [9] Buehler, C., Bosse, M., McMillan, L., Gortler, S., Cohen, M.: Unstructured lumigraph rendering. In: Proceedings of the 28th annual conference on Computer graphics and interactive techniques (2001)
- [10] Cai, H., Feng, W., Feng, X., Wang, Y., Zhang, J.: Neural surface reconstruction of dynamic scenes with monocular RGB-D camera. In: NeurIPS (2022)
- [11] Cao, A., Johnson, J.: Hexplane: a fast representation for dynamic scenes. arXiv preprint arXiv:2301.09632 (2023)
- [12] Cao, J., Wang, H., Chemerys, P., Shakhrai, V., Hu, J., Fu, Y., Makoviichuk, D., Tulyakov, S., Ren, J.: Real-time neural light field on mobile devices. arXiv preprint arXiv:2212.08057 (2022)
- [13] Chen, R., Chen, Y., Jiao, N., Jia, K.: Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. In: ICCV (2023)
- [14] Chen, Z., Wang, F., Wang, Y., Liu, H.: Text-to-3d using gaussian splatting. In: CVPR (2024)

- [15] Chen, Z., Wang, Y., Wang, F., Wang, Z., Liu, H.: V3d: Video diffusion models are effective 3d generators. arXiv preprint arXiv:2403.06738 (2024)
- [16] Dai, P., Xu, J., Xie, W., Liu, X., Wang, H., Xu, W.: High-quality surface reconstruction using gaussian surfels. In: SIGGRAPH (2024)
- [17] Debevec, P.E., Taylor, C.J., Malik, J.: Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques (1996)
- [18] Du, Y., Zhang, Y., Yu, H.X., Tenenbaum, J.B., Wu, J.: Neural radiance flow for 4d view synthesis and video processing. In: ICCV (2021)
- [19] Fang, J., Yi, T., Wang, X., Xie, L., Zhang, X., Liu, W., Nießner, M., Tian, Q.: Fast dynamic radiance fields with time-aware neural voxels. In: SIGGRAPH Asia (2022)
- [20] Flynn, J., Broxton, M., Debevec, P., DuVall, M., Fyffe, G., Overbeck, R., Snavely, N., Tucker, R.: Deepview: View synthesis with learned gradient descent. In: CVPR (2019)
- [21] Gao, C., Saraf, A., Kopf, J., Huang, J.B.: Dynamic view synthesis from dynamic monocular video. In: ICCV (2021)
- [22] Gao, H., Li, R., Tulsiani, S., Russell, B., Kanazawa, A.: Dynamic novel-view synthesis: A reality check. In: NeurIPS (2022)
- [23] Garbin, S.J., Kowalski, M., Johnson, M., Shotton, J., Valentin, J.: Fastnerf: High-fidelity neural rendering at 200fps. In: ICCV (2021)
- [24] Guédon, A., Lepetit, V.: Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. arXiv preprint arXiv:2311.12775 (2023)
- [25] Hedman, P., Srinivasan, P.P., Mildenhall, B., Barron, J.T., Debevec, P.: Baking neural radiance fields for real-time view synthesis. ICCV (2021)
- [26] Hong, Y., Zhang, K., Gu, J., Bi, S., Zhou, Y., Liu, D., Liu, F., Sunkavalli, K., Bui, T., Tan, H.: Lrm: Large reconstruction model for single image to 3d. arXiv preprint arXiv:2311.04400 (2023)
- [27] Hu, T., Liu, S., Chen, Y., Shen, T., Jia, J.: Efficientnerf efficient neural radiance fields. In: CVPR (2022)
- [28] Huang, B., Yu, Z., Chen, A., Geiger, A., Gao, S.: 2d gaussian splatting for geometrically accurate radiance fields. In: SIGGRAPH. Association for Computing Machinery (2024)
- [29] Huang, Y.H., Sun, Y.T., Yang, Z., Lyu, X., Cao, Y.P., Qi, X.: Sc-gs: Sparse-controlled gaussian splatting for editable dynamic scenes. arXiv preprint arXiv:2312.14937 (2023)
- [30] Jiakai, Z., Xinhang, L., Xinyi, Y., Fuqiang, Z., Yanshun, Z., Minye, W., Yingliang, Z., Lan, X., Jingyi, Y.: Editable free-viewpoint video using a layered neural representation. In: SIGGRAPH (2021)
- [31] Jiang, Y., Hedman, P., Mildenhall, B., Xu, D., Barron, J.T., Wang, Z., Xue, T.: Alignerf: High-fidelity neural radiance fields via alignment-aware training. arXiv preprint arXiv:2211.09682 (2022)
- [32] Kavan, L., Collins, S., Zára, J., O’Sullivan, C.: Skinning with dual quaternions. In: SI3D (2007)
- [33] Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Trans. Graph. (2023)
- [34] Kerbl, B., Meuleman, A., Kopanas, G., Wimmer, M., Lanvin, A., Drettakis, G.: A hierarchical 3d gaussian representation for real-time rendering of very large datasets. ACM Trans. Graph. (2024)
- [35] Kutulakos, K.N., Seitz, S.M.: A theory of shape by space carving. IJCV (2000)
- [36] Li, L., Shen, Z., Wang, Z., Shen, L., Tan, P.: Streaming radiance fields for 3d video synthesis. arXiv preprint arXiv:2210.14831 (2022)
- [37] Li, R., Tanke, J., Vo, M., Zollhofer, M., Gall, J., Kanazawa, A., Lassner, C.: Tava: Template-free animatable volumetric actors (2022)
- [38] Li, T., Slavcheva, M., Zollhofer, M., Green, S., Lassner, C., Kim, C., Schmidt, T., Lovegrove, S., Goesele, M., Newcombe, R., et al.: Neural 3d video synthesis from multi-view video. In: CVPR (2022)
- [39] Li, Z., Niklaus, S., Snavely, N., Wang, O.: Neural scene flow fields for space-time view synthesis of dynamic scenes. In: CVPR (2021)
- [40] Lin, C.H., Gao, J., Tang, L., Takikawa, T., Zeng, X., Huang, X., Kreis, K., Fidler, S., Liu, M.Y., Lin, T.Y.: Magic3d: High-resolution text-to-3d content creation. In: CVPR (2023)
- [41] Lindell, D.B., Martel, J.N., Wetzstein, G.: Autoint: Automatic integration for fast neural volume rendering. In: CVPR (2021)
- [42] Ling, H., Kim, S.W., Torralba, A., Fidler, S., Kreis, K.: Align your gaussians: Text-to-4d with dynamic 3d gaussians and composed diffusion models. arXiv preprint arXiv:2312.13763 (2023)

- [43] Liu, I., Su, H., Wang, X.: Dynamic gaussians mesh: Consistent mesh reconstruction from monocular videos. arXiv preprint arXiv:2404.12379 (2024)
- [44] Liu, J.W., Cao, Y.P., Mao, W., Zhang, W., Zhang, D.J., Keppo, J., Shan, Y., Qie, X., Shou, M.Z.: Devrf: Fast deformable voxel radiance fields for dynamic scenes. arXiv preprint arXiv:2205.15723 (2022)
- [45] Liu, L., Gu, J., Zaw Lin, K., Chua, T.S., Theobalt, C.: Neural sparse voxel fields. In: NeurIPS (2020)
- [46] Liu, Y., Guan, H., Luo, C., Fan, L., Peng, J., Zhang, Z.: Citygaussian: Real-time high-quality large-scale scene rendering with gaussians. arXiv preprint arXiv: 2404.01133 (2024)
- [47] Liu, Y.L., Gao, C., Meuleman, A., Tseng, H.Y., Saraf, A., Kim, C., Chuang, Y.Y., Kopf, J., Huang, J.B.: Robust dynamic radiance fields. In: CVPR (2023)
- [48] Lombardi, S., Simon, T., Saragih, J., Schwartz, G., Lehrmann, A., Sheikh, Y.: Neural volumes: Learning dynamic renderable volumes from images. arXiv preprint arXiv:1906.07751 (2019)
- [49] Lombardi, S., Simon, T., Schwartz, G., Zollhoefer, M., Sheikh, Y., Saragih, J.: Mixture of volumetric primitives for efficient neural rendering. arXiv preprint arXiv:2103.01954 (2021)
- [50] Long, X., Guo, Y.C., Lin, C., Liu, Y., Dou, Z., Liu, L., Ma, Y., Zhang, S.H., Habermann, M., Theobalt, C., et al.: Wonder3d: Single image to 3d using cross-domain diffusion. arXiv preprint arXiv:2310.15008 (2023)
- [51] Lu, Y., Zhang, J., Li, S., Fang, T., McKinnon, D., Tsin, Y., Quan, L., Cao, X., Yao, Y.: Direct2. 5: Diverse text-to-3d generation via multi-view 2.5 d diffusion. arXiv preprint arXiv:2311.15980 (2023)
- [52] Luiten, J., Kopanas, G., Leibe, B., Ramanan, D.: Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In: 3DV (2024)
- [53] Ma, L., Li, X., Liao, J., Zhang, Q., Wang, X., Wang, J., Sander, P.V.: Deblur-nerf: Neural radiance fields from blurry images. arXiv preprint arXiv:2111.14292 (2021)
- [54] Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. ACM Trans. Graph. (2019)
- [55] Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020)
- [56] Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023)
- [57] Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: ICCV (2021)
- [58] Park, K., Sinha, U., Hedman, P., Barron, J.T., Bouaziz, S., Goldman, D.B., Martin-Brualla, R., Seitz, S.M.: Hypernerf: a higher-dimensional representation for topologically varying neural radiance fields. ACM Trans. Graph. (2021)
- [59] Peng, S., Yan, Y., Shuai, Q., Bao, H., Zhou, X.: Representing volumetric videos as dynamic mlp maps. In: CVPR (2023)
- [60] Peng, S., Zhang, Y., Xu, Y., Wang, Q., Shuai, Q., Bao, H., Zhou, X.: Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In: CVPR (2021)
- [61] Penner, E., Zhang, L.: Soft 3d reconstruction for view synthesis. ACM Trans. Graph. (2017)
- [62] Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2209.14988 (2022)
- [63] Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: CVPR (2021)
- [64] Rebain, D., Jiang, W., Yazdani, S., Li, K., Yi, K.M., Tagliasacchi, A.: Derf: Decomposed radiance fields. In: CVPR (2021)
- [65] Reiser, C., Peng, S., Liao, Y., Geiger, A.: Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In: CVPR (2021)
- [66] Reiser, C., Szeliski, R., Verbin, D., Srinivasan, P.P., Mildenhall, B., Geiger, A., Barron, J.T., Hedman, P.: Merf: Memory-efficient radiance fields for real-time view synthesis in unbounded scenes. arXiv preprint arXiv:2302.12249 (2023)
- [67] Riegler, G., Koltun, V.: Free view synthesis. In: ECCV (2020)
- [68] Sara Fridovich-Keil and Giacomo Meanti, Warburg, F.R., Recht, B., Kanazawa, A.: K-planes: Explicit radiance fields in space, time, and appearance. In: CVPR (2023)

- [69] Shao, R., Zheng, Z., Tu, H., Liu, B., Zhang, H., Liu, Y.: Tensor4d: Efficient neural 4d decomposition for high-fidelity dynamic reconstruction and rendering. In: CVPR (2023)
- [70] Shuai, Q., Guo, H., Xu, Z., Lin, H., Peng, S., Bao, H., Zhou, X.: Real-time view synthesis for large scenes with millions of square meters (2024)
- [71] Singer, U., Sheynin, S., Polyak, A., Ashual, O., Makarov, I., Kokkinos, F., Goyal, N., Vedaldi, A., Parikh, D., Johnson, J., et al.: Text-to-4d dynamic scene generation. arXiv preprint arXiv:2301.11280 (2023)
- [72] Sitzmann, V., Thies, J., Heide, F., Nießner, M., Wetzstein, G., Zollhofer, M.: Deepvoxels: Learning persistent 3d feature embeddings. In: CVPR (2019)
- [73] Song, L., Chen, A., Li, Z., Chen, Z., Chen, L., Yuan, J., Xu, Y., Geiger, A.: Nerfplayer: A streamable dynamic scene representation with decomposed neural radiance fields. arXiv preprint arXiv:2210.15947 (2022)
- [74] Srinivasan, P.P., Mildenhall, B., Tancik, M., Barron, J.T., Tucker, R., Snavely, N.: Lighthouse: Predicting lighting volumes for spatially-coherent illumination. In: CVPR (2020)
- [75] Srinivasan, P.P., Tucker, R., Barron, J.T., Ramamoorthi, R., Ng, R., Snavely, N.: Pushing the boundaries of view extrapolation with multiplane images. In: CVPR (2019)
- [76] Su, S.Y., Yu, F., Zollhöfer, M., Rhodin, H.: A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. In: NeurIPS (2021)
- [77] Tang, J., Chen, Z., Chen, X., Wang, T., Zeng, G., Liu, Z.: Lgm: Large multi-view gaussian model for high-resolution 3d content creation. arXiv preprint arXiv:2402.05054 (2024)
- [78] Thies, J., Zollhöfer, M., Nießner, M.: Deferred neural rendering: Image synthesis using neural textures. ACM Trans. Graph. (2019)
- [79] Tretschk, E., Tewari, A., Golyanik, V., Zollhöfer, M., Lassner, C., Theobalt, C.: Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In: ICCV (2021)
- [80] Tucker, R., Snavely, N.: Single-view view synthesis with multiplane images. In: CVPR (2020)
- [81] Voleti, V., Yao, C.H., Boss, M., Letts, A., Pankratz, D., Tochilkin, D., Laforte, C., Rombach, R., Jampani, V.: Sv3d: Novel multi-view synthesis and 3d generation from a single image using latent video diffusion. arXiv preprint arXiv: 2403.12008 (2024)
- [82] Waechter, M., Moehrl, N., Goesele, M.: Let there be color! large-scale texturing of 3d reconstructions. In: ECCV (2014)
- [83] Wang, F., Chen, Z., Wang, G., Song, Y., Liu, H.: Masked space-time hash encoding for efficient dynamic scene reconstruction. In: NeurIPS (2023)
- [84] Wang, F., Tan, S., Li, X., Tian, Z., Liu, H.: Mixed neural voxels for fast multi-view video synthesis. arXiv preprint arXiv:2212.00190 (2022)
- [85] Wang, H., Ren, J., Huang, Z., Olszewski, K., Chai, M., Fu, Y., Tulyakov, S.: R2l: Distilling neural radiance field to neural light field for efficient novel view synthesis. In: ECCV (2022)
- [86] Wang, L., Zhang, J., Liu, X., Zhao, F., Zhang, Y., Zhang, Y., Wu, M., Yu, J., Xu, L.: Fourier plencotrees for dynamic radiance field rendering in real-time. In: CVPR (2022)
- [87] Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In: NeurIPS (2021)
- [88] Wang, X., Wang, Y., Ye, J., Wang, Z., Sun, F., Liu, P., Wang, L., Sun, K., Wang, X., He, B.: Animatable-dreamer: Text-guided non-rigid 3d model generation and reconstruction with canonical score distillation. arXiv preprint arXiv:2312.03795 (2023)
- [89] Wang, Y., Dong, Y., Sun, F., Yang, X.: Root pose decomposition towards generic non-rigid 3d reconstruction with monocular videos. In: ICCV (2023)
- [90] Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. In: NeurIPS (2023)
- [91] Wang, Z., Wang, Y., Chen, Y., Xiang, C., Chen, S., Yu, D., Li, C., Su, H., Zhu, J.: Crm: Single image to 3d textured mesh with convolutional reconstruction model. arXiv preprint arXiv:2403.05034 (2024)
- [92] Wang, Z., Li, L., Shen, Z., Shen, L., Bo, L.: 4k-nerf: High fidelity neural radiance fields at ultra high resolutions. arXiv preprint arXiv:2212.04701 (2022)
- [93] Wood, D.N., Azuma, D.I., Aldinger, K., Curless, B., Duchamp, T., Salesin, D.H., Stuetzle, W.: Surface light fields for 3d photography. In: Proceedings of the 27th annual conference on Computer graphics and interactive techniques (2000)

- [94] Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Xinggang, W.: 4d gaussian splatting for real-time dynamic scene rendering. arXiv preprint arXiv:2310.08528 (2023)
- [95] Xian, W., Huang, J.B., Kopf, J., Kim, C.: Space-time neural irradiance fields for free-viewpoint video. In: CVPR (2021)
- [96] Xu, Y., Shi, Z., Yifan, W., Peng, S., Yang, C., Shen, Y., Gordon, W.: Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation. arXiv preprint arXiv: 2403.14621 (2024)
- [97] Yang, G., Sun, D., Jampani, V., Vlastic, D., Cole, F., Chang, H., Ramanan, D., Freeman, W.T., Liu, C.: LASR: learning articulated shape reconstruction from a monocular video. In: ICCV (2021)
- [98] Yang, G., Vo, M., Neverova, N., Ramanan, D., Vedaldi, A., Joo, H.: Banmo: Building animatable 3d neural models from many casual videos. In: CVPR (2022)
- [99] Yang, G., Wang, C., Reddy, N.D., Ramanan, D.: Reconstructing animatable categories from videos. In: CVPR (2023)
- [100] Yang, Z., Gao, X., Zhou, W., Jiao, S., Zhang, Y., Jin, X.: Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. arXiv preprint arXiv:2309.13101 (2023)
- [101] Yariv, L., Kasten, Y., Moran, D., Galun, M., Atzmon, M., Basri, R., Lipman, Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. In: NeurIPS (2020)
- [102] Yu, A., Li, R., Tancik, M., Li, H., Ng, R., Kanazawa, A.: Plenotrees for real-time rendering of neural radiance fields. In: CVPR (2021)
- [103] Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018)
- [104] Zhao, F., Yang, W., Zhang, J., Lin, P., Zhang, Y., Yu, J., Xu, L.: Humannerf: Efficiently generated human radiance field from sparse inputs. In: CVPR (2022)
- [105] Zhou, T., Tucker, R., Flynn, J., Fyffe, G., Snavely, N.: Stereo magnification: Learning view synthesis using multiplane images. arXiv preprint arXiv:1805.09817 (2018)
- [106] Zou, Z.X., Yu, Z., Guo, Y.C., Li, Y., Liang, D., Cao, Y.P., Zhang, S.H.: Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. arXiv preprint arXiv:2312.09147 (2023)

A Appendix / Supplemental Material

Table 2: A summary of important symbols in DGS.

Symbol	Definition and Usage
$\mathbf{t}_u^* \in \mathbb{R}^{3 \times 1}, \mathbf{t}_v^* \in \mathbb{R}^{3 \times 1}$	Principal tangential vectors in the static state.
$s_u^* \in \mathbb{R}, s_v^* \in \mathbb{R}$	Scaling factors in the static state.
$\mathbf{p}_k^* \in \mathbb{R}^{3 \times 1}$	Center point coordinate (world space) of the k -th Gaussian surfel in the static state.
$P_k^*(\mathbf{u}) \in \mathbb{R}^{3 \times 1}$	Coordinate (world space) in the static state, given $\mathbf{u} = (u, v)$ on the local uv coordinate system centered at \mathbf{p}_k^* .
$\mathbf{R}_k^* = [\mathbf{t}_u^*, \mathbf{t}_v^*, \mathbf{t}_u^* \times \mathbf{t}_v^*] \in \text{SO}(3)$	Rotation matrix of the k -th Gaussian surfel in the static state.
$\mathbf{S}_k^* = \text{diag}(s_u^*, s_v^*, 0) \in \mathbb{R}^{3 \times 3}$	Scaling matrix of the k -th Gaussian surfel in the static state, a diagonal matrix.
$\mathbf{p}_k^t \in \mathbb{R}^{3 \times 1}$	Center point coordinate (world space) of the k -th Gaussian surfel in the warped state.
$P_k^t(\mathbf{u}) \in \mathbb{R}^{3 \times 1}$	Coordinate (world space) in the warped state, given $\mathbf{u} = (u, v)$ on the local uv coordinate system centered at \mathbf{p}_k^t .
$\mathbf{c}_b^* \in \mathbb{R}^{3 \times 1}, \mathbf{V}_b^* \in \mathbb{R}^{3 \times 3}, \mathbf{\Lambda}_b^* \in \mathbb{R}^{3 \times 3}$	Center, rotation matrix, and diagonal scaling matrix of the b -th Gaussian ellipsoid bone.
$\mathbf{w}^t \in \mathbb{R}^{B \times 1}$	Skinning weight vectors.
$\mathbf{J}_b^t \in \text{SE}(3)$	A rigid transformation that moves the b -th bone from its static state to the warped state at time t .
$\mathbf{J}^t = [\mathbf{R}^t, \mathbf{T}^t] \in \text{SE}(3)$	The warping function, a weighted combination of \mathbf{J}_b^t .
\mathcal{Q}, \mathcal{R}	The quaternion process and the inverse quaternion process.
$\omega_b^t \in \mathbb{R}^{128}$	A learnable latent code for representing the body pose at time t .
$\mathbf{n}_k \in \mathbb{R}^{3 \times 1}$	The normal of the k -intersected Gaussian surfel that is oriented towards the camera.
$\mathbf{N}^t \in \mathbb{R}^{3 \times 1}$	The surface normal estimated by the nearby depth point \mathbf{p}^t at warped state time t .
$\Delta \mathbf{R}_k^* \in \text{SO}(3)$	Learnable refinement term for adjusting \mathbf{R}_k^* .
$\Delta \mathbf{S}_k^* \in \text{SO}(3)$	Learnable refinement term for adjusting \mathbf{S}_k^* .

A.1 Details of Skinning Representation

As mentioned in the main paper, the warping process from the static state to the warped state is modelled as a time-varying non-rigid warping function with B bones to be key points. In the static state, the b -th bone is represented by 3D Gaussian ellipsoids [97] with the center $\mathbf{c}_b^* \in \mathbb{R}^{3 \times 1}$, rotation matrix $\mathbf{V}_b^* \in \mathbb{R}^{3 \times 3}$, and diagonal scaling matrix $\mathbf{\Lambda}_b^* \in \mathbb{R}^{3 \times 3}$. For a 3D point $P_k^*(\mathbf{u})$, the skinning weight vectors $\mathbf{w}^t \in \mathbb{R}^{B \times 1}$ at time t is calculated by the normalized Mahalanobis distance following [98]

$$m_b^t = (P_k^*(\mathbf{u}) - \mathbf{c}_b^t)^\top \mathbf{Q}_b^t (P_k^*(\mathbf{u}) - \mathbf{c}_b^t), \quad \mathbf{w}^t = \sigma_{\text{softmax}}(m_1^t, m_2^t, \dots, m_B^t)^\top, \quad (11)$$

where m_b^t denotes the squared distance between $P_k^*(\mathbf{u})$ and the b -th bone; $\mathbf{c}_b^t \in \mathbb{R}^{3 \times 1}$ is the center of the b -th bone at time t , and $\mathbf{Q}_b^t = \mathbf{V}_b^{t \top} \mathbf{\Lambda}_b^* \mathbf{V}_b^t$ is the precision matrix composed by the bone orientation matrix $\mathbf{V}_b^t \in \mathbb{R}^{3 \times 3}$ at time t and $\mathbf{\Lambda}_b^*$. Specifically, there is $(\mathbf{V}_b^t | \mathbf{c}^t) = \mathbf{J}_b^t (\mathbf{V}_b^* | \mathbf{c}^*)$ with \mathbf{c}_b^* , \mathbf{V}_b^* , and $\mathbf{\Lambda}_b^*$ being learnable parameters. σ_{softmax} is the softmax function.

A.2 Ablation Studies of Field Initialization and Refinement

In dynamic videos captured in the wild, one of the primary challenges is the initialization of camera poses. In synthetic videos, preserving temporal consistency in texture and geometry is problematic, which significantly complicates the task of camera registration. To address this, we utilize an implicit field to both initialize the camera poses and establish the warping field. Initially, we estimate the transformation for each frame, followed by the computation of coarse camera poses through an iterative process. Subsequently, we adopt the approach outlined in NeuS [87] for scene representation. Feature extraction is performed using DinoV2 [56], facilitating unsupervised registration. To enhance this process, we train an additional channel in NeuS specifically for rendering features, which are then employed for registration purposes as described in RAC [99]. The camera poses without initialization and refined camera poses are depicted in Fig. 6. Without field initialization, the performance of DGS will degrade, as shown in Table 3. Also, please refer to the quantitative ablation of refinement in Table 3.

A.3 Additional Qualitative Comparison

In this section, we present a detailed comparison of our results with previous works, as illustrated in Fig. 7-10. Our method consistently achieves high-quality texture details while maintaining smooth and realistic geometry.

Table 3: Quantitative ablation studies of the initialization and refinement.

	Cat			Cheetah			Dragon		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Ours w.o. init.	20.15	0.7961	0.2393	20.96	0.8194	0.1940	25.33	0.9146	0.0938
Ours w.o. refinement	24.19	0.8196	0.1797	24.10	0.8582	0.1242	27.71	0.9128	0.0687
Ours full	24.63	0.8432	0.1559	25.68	0.8843	0.1117	28.58	0.9392	0.0618

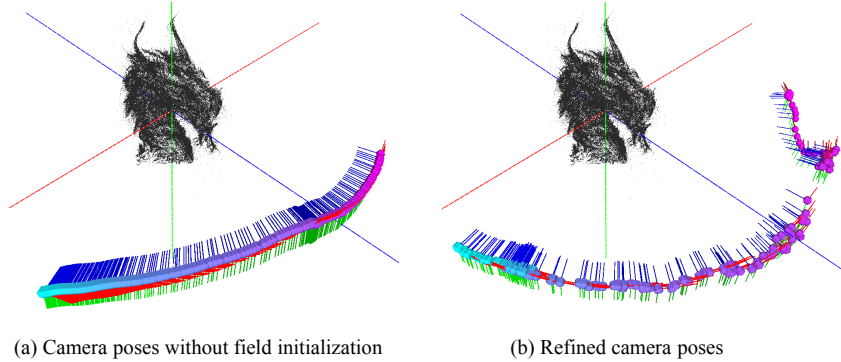


Figure 6: Coarse camera poses and refined camera poses.

A.4 Interpolation on Time and Views

We present results for interpolation on time and views, as illustrated in Fig. 11 and Fig. 12.

A.5 Broader Impact

Generative models used in video generation might pose risks, for example, the potential for creating deepfakes or other misleading content that could be used for harmful purposes like misinformation, privacy invasion, or defamation. To mitigate these risks, we have chosen to release only the reconstruction code, deliberately avoiding the release of components that could facilitate the generation of content with ethical concerns. This decision ensures that our contribution is focused on advancing reconstruction techniques without enabling the creation of new, potentially harmful video content.

Besides, we have carefully considered the potential ethical risks associated with generative models, particularly in video content creation. To address these concerns, our model includes robust safety mechanisms designed to screen and prevent any misuse. We believe these measures effectively mitigate potential ethical risks and align with the standards of the community.



Figure 7: Additional qualitative comparison with more novel views.



Figure 8: Additional qualitative comparison with more novel views.




Figure 9: Additional qualitative comparison with more novel views.



Figure 10: Additional qualitative comparison with more novel views.

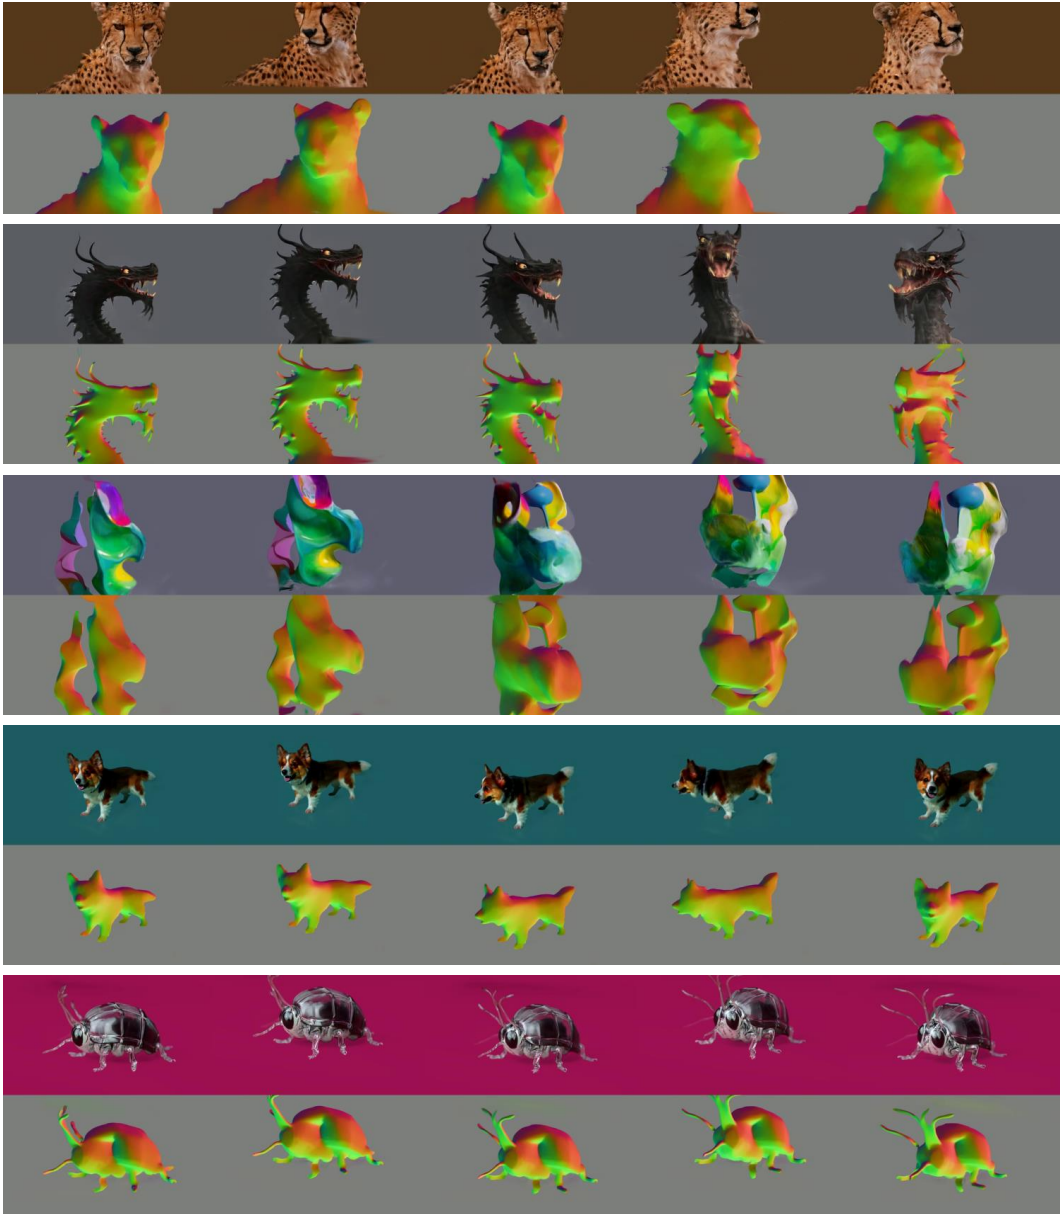


Figure 11: Interpolation on time and views.




Figure 12: Interpolation on time and views.

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: Our main claims in the abstract and introduction accurately reflect the contributions and scope.

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: We discuss it Sec. 5.

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: [Yes]

Justification: We build our proof upon papers cited.

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: All relevant details are provided within the text of our paper or through the references we have cited.

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: [Yes]

Justification: We will release it after acceptance.

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, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: All relevant details are provided within the text of our paper or through the references we have cited.

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: Please refer to our Tab. 1.

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: Please refer to Sec. 4.

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: The NeurIPS Code of Ethics is conformed.

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: We discuss it in Sec. 5.

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: No such risks.

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: [Yes]

Justification: We cite the assets properly.

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: [No]

Justification: No 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: [NA] .

Justification: No human subjects.

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: [NA] .

Justification: No human subjects.

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.