

Efficient Gaussian Splatting for Monocular Dynamic Scene Rendering via Sparse Time-Variant Attribute Modeling

Hanyang Kong, Xingyi Yang, Xinchao Wang*

National University of Singapore
hanyang.k@u.nus.edu, xyang@u.nus.edu, xinchao@nus.edu.sg

Abstract

Rendering dynamic scenes from monocular videos is a crucial yet challenging task. The recent deformable Gaussian Splatting has emerged as a robust solution to represent real-world dynamic scenes. However, it often leads to heavily redundant Gaussians, attempting to fit every training view at various time steps, leading to slower rendering speeds. Additionally, the attributes of Gaussians in static areas are time-invariant, making it unnecessary to model every Gaussian, which can cause jittering in static regions. In practice, the primary bottleneck in rendering speed for dynamic scenes is the number of Gaussians. In response, we introduce Efficient Dynamic Gaussian Splatting (EDGS), which represents dynamic scenes via sparse time-variant attribute modeling. Our approach formulates dynamic scenes using a sparse anchor-grid representation, with the motion flow of dense Gaussians calculated via a classical kernel representation. Furthermore, we propose an unsupervised strategy to efficiently filter out anchors corresponding to static areas. Only anchors associated with deformable objects are input into MLPs to query time-variant attributes. Experiments on two real-world datasets demonstrate that our EDGS significantly improves the rendering speed with superior rendering quality compared to previous state-of-the-art methods.

1 Introduction

Novel view synthesis (NVS) is a pivotal challenge in the field of 3D vision, essential for applications such as virtual reality, augmented reality, and film production. NVS involves generating images from arbitrary viewpoints or times within a scene, typically necessitating accurate reconstruction based on several 2D images. While recent advances in diffusion models (Kong et al. 2025; Yu et al. 2024; Ma, Fang, and Wang 2024) have shown promise in NVS, dynamic scenes remain challenging due to the need to model complex motions and the need for real-time processing.

Recent advancements in 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) provide new tools to tackle these challenges by enabling real-time rendering. Extensions of 3DGS, such as (Huang et al. 2023; Wu et al. 2023; Yang et al. 2023), have further enhanced its ability to handle dynamic monocular scenes. For example, Deformable

3DGS (Yang et al. 2023) incorporates a deformable network to capture motion fields, enhancing 3DGS’s adaptability to dynamic scenes. Similarly, 4DGS (Wu et al. 2023) utilizes a hexplane-based (Cao and Johnson 2023) encoder to optimize deformation queries, making the process more efficient.

While previous methods deliver high-quality renderings, they tend to generate an excessive number of Gaussians, resulting in considerable redundancy. We analyze the relationship between the rendering speed *w.r.t.* number of Gaussians based on (Yang et al. 2023) in Fig. 1(c). As a result, the rendering speed declines as the number of redundant Gaussians increases. Inspired by recent work (Lu et al. 2023) which improved efficiency in static scenes by sparsifying and reducing Gaussian points, we question: *Can we achieve a more efficient and compact representation of dynamic scenes without compromising rendering quality?*

In this paper, we affirmatively answer this question by introducing Efficient Dynamic Gaussian Splatting (EDGS). This method efficiently represents dynamic scenes as a set of sparse time-variant attributes. At its heart, EDGS separately processes time-invariant attributes, time-variant attributes, and motions. This design allows precise control over each aspect of the scene’s dynamics.

For time-invariant attributes, EDGS uses a sparse anchor-grid initialization to capture geometry and appearance. These attributes, such as color and opacity, are decoded by lightweight MLPs, enabling accurate deformation modeling.

Unlike prior methods that process all time-variant Gaussian attributes (e.g., location, scale, quaternion) indiscriminately through MLPs, EDGS adopts a selective approach. A tiny MLP filters static anchors, identifying those that are deformable. Only these deformable anchors are then processed by corresponding MLPs for their dynamic attributes. Notably, this selection process is trained in a fully unsupervised manner.

Moreover, we efficiently model the movements of deformable Gaussians in a sparse manner. Specifically, we track the movements of anchors linked to deformable objects at each time step. The offsets for corresponding Gaussians are computed using a radial basis function (RBF) kernel. This sparse approach enables precise and efficient rendering of dynamic changes.

By integrating these techniques, we carried out detailed

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: (a) Given a set of monocular multi-view images and camera poses, our method achieves real-time rendering for dynamic scenes while maintaining high rendering quality. (b) Our method achieves promising rendering quality with faster rendering speed and fewer Gaussians. The radius of the circle is the number of time-variant Gaussians whose attributes need to be queried by MLPs. (c) The bottleneck of the rendering speed for dynamic scenes is the number of Gaussians. The fewer the number of Gaussians, the faster the rendering speed.

experiments on two real-world datasets, NeRF-DS (Yan, Li, and Lee 2023) and HyperNeRF (Park et al. 2021b). We visualize the rendered novel views and show the rendering speed in Fig. 1. Our method achieves a much faster rendering speed with a higher PSNR score and fewer points that the time-variant attributes need to be queried by MLPs. In summary, our contributions are summarized as follows:

- We formulate deformable 3DGS through sparse, time-variant attribute modeling and introduce a novel unsupervised strategy to filter out Gaussians with time-invariant attributes.
- We employ a classical kernel representation to formulate motion flow sparsely for the deformable Gaussians.
- Our EDGS achieves faster rendering speed compared with other state-of-the-art methods with higher rendering quality.

2 Related Works

Dynamic 3D scene reconstruction. 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) offers a distinct representation for novel view synthesis with improved rendering quality and speed. Several concurrent works (Wang et al. 2025; Wu et al. 2023; Yang et al. 2023; Huang et al. 2023) have adapted 3DGS for reconstructing dynamic scenes. For instance, Deformable 3DGS (Yang et al. 2023) formulates the dynamic field based on multi-layer perceptron (MLP) with differentiable rendering. 4DGS (Wu et al. 2023) improves the rendering speed with a compact network. SCGS (Huang et al. 2023) explicitly decomposes the motion and dynamic scenes into sparse control points and the deformation of Gaussians is controlled by its k -nearest control points.

Though achieving promising rendered quality, the rendering speed for dynamic scenes is much slower than rendering static scenes (100 FPS *v.s.* 20 FPS at 1K resolution). The bottleneck of the rendering speed for dynamic scenes is the number of Gaussian points. As shown in Fig. 1, as more

Gaussians are densified, an enormous amount of Gaussians are fed into the deformation network, which leads to slower rendering speed. Moreover, (Huang et al. 2023; Yang et al. 2023; Wu et al. 2023) query the attributes of all Gaussians at each timestep, even though the static areas in real-world scenes are time-invariant and do not require querying by MLPs. As shown in Fig. 5, the static areas of (Yang et al. 2023; Wu et al. 2023)’s rendered results are still jittering across time.

Grid-based rendering for acceleration. Grid-based representations are based on a dense uniform grid of voxels (Fridovich-Keil et al. 2022, 2023; Cao and Johnson 2023; Lu et al. 2023) explicitly or implicitly. For instance, K-Planes (Fridovich-Keil et al. 2023) applies neural planes to parameterize a 3D scene, with an additional temporal plane to accommodate dynamics. HexPlane (Cao and Johnson 2023) further enhances the neural planes for time and space by factorizing space and time into compact neural representations. Scaffold-GS (Lu et al. 2023) introduces anchor-voxel-based strategy to achieve reduced storage parameters for static scenes. Scaffold-GS can render a similar speed (100 FPS at 1K resolution) as 3DGS and the storage requirements are significantly reduced as Scaffold-GS only needs to store anchor points and MLPs for each scene. However, Scaffold-GS is designed for real-world static scenes and we are the first to adapt anchor-voxel-based strategy to dynamic scene reconstruction.

3 Preliminaries

In this section, we simply review the representation and rendering process of 3DGS (Kerbl et al. 2023) with the formula of deformable 3DGS (Wu et al. 2023; Yang et al. 2023) and our baseline method (Lu et al. 2023).

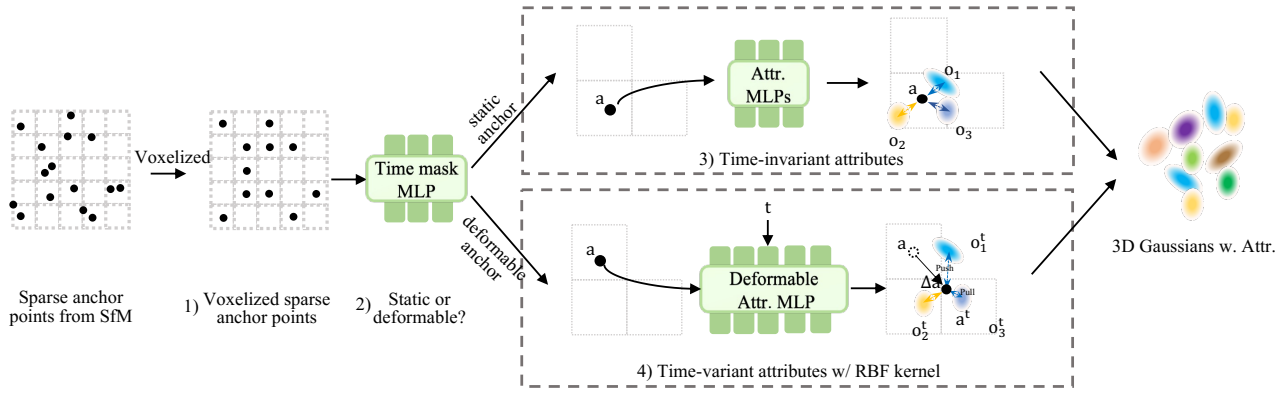


Figure 2: **The pipeline of our EDGS.** 1) We first initialize voxelized sparse anchor points from Structure from Motion (SfM) points derived from COLMAP. 2) A time-mask MLP is applied to classify if the anchor belongs to the static area or the deformable area. 3) k Gaussian offsets are initialized for each anchor a belonging to static area. The time-invariant attributes of each Gaussian, *i.e.*, opacity, quaternion, scale, and color are calculated from its feature by corresponding tiny MLPs. 4) Time-variant attributes for anchors from dynamic areas are decoded by a deformable attribute MLP. RBF kernel function is employed to compute the location of each Gaussian at timestep t by calculating the similarity between each Gaussian and its belonging anchor point. This pipeline is compact and efficient, featuring only a few tiny MLPs for the attributes of the Gaussians and a single network for deformations. Notably, the position of each anchor remains static and is not subject to updates.

3.1 Dynamic 3D Gaussian Splatting

3DGS (Kerbl et al. 2023) is an explicit 3D representation in the form of point clouds. The process starts with 3D point clouds generated through Structure-from-Motion (SfM). Each Gaussian $G(\mathbf{x})$ is defined by a mean position $\boldsymbol{\mu}$ and a covariance matrix $\boldsymbol{\Sigma}$

$$G(\mathbf{x}) = e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}, \quad (1)$$

where \mathbf{x} is a 3D point location within the 3D scene. $\boldsymbol{\Sigma}$ is formulated using a scaling matrix \mathbf{S} and rotation matrix \mathbf{R} :

$$\boldsymbol{\Sigma} = \mathbf{R} \mathbf{S} \mathbf{S}^T \mathbf{R}^T. \quad (2)$$

To render an image from a random viewpoint, 3D Gaussians are first splatted to 2D, and render the pixel value C by the following formula:

$$C = \sum_{i \in N} c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (3)$$

where c_i and α_i represent the density and view-dependent color of the point, and N is the number of sorted Gaussians contributing to the rendering process.

To model the dynamic 3D Gaussians that vary over time, (Huang et al. 2023; Yang et al. 2023; Wu et al. 2023) decouple the 3D Gaussians and the deformation field, separately. To be specific, given time t and the location \mathbf{x} of 3D Gaussians as input, deformable 3DGS methods (Huang et al. 2023; Yang et al. 2023; Wu et al. 2023) apply the trained deformation MLP to produce offsets, which subsequently transform the 3D Gaussians from canonical space to the deformed space:

$$\Delta \mathbf{x}, \Delta \mathbf{r}, \Delta \mathbf{s} = \mathcal{F}(\gamma(\mathbf{x}, t)), \quad (4)$$

where γ is the encoding of space and time, $\Delta \mathbf{x}$, $\Delta \mathbf{r}$, and $\Delta \mathbf{s}$ are the offsets of the Gaussian location \mathbf{x} , quaternion \mathbf{r} ,

and scaling \mathbf{s} , respectively. Subsequently, the deformed 3D Gaussians $G(\mathbf{x} + \Delta \mathbf{x}, \mathbf{r} + \Delta \mathbf{r}, \mathbf{s} + \Delta \mathbf{s})$ are fed into the differential Gaussian rasterization pipeline for rendering novel views with various times t . All attributes of 3D Gaussians and the deformable network \mathcal{F} are learnable and optimized end-to-end directly via training view reconstruction.

Though achieving promising results, the rendering speed of dynamic scenes remains significantly slower than that of static scenes (e.g., 20 FPS vs. 100 FPS for 1K resolution). This slowdown is due to the redundancy of Gaussians and the need to query the attributes of all Gaussians using the deformation network \mathcal{F} . A sparse Gaussian representation would accelerate the rendering process. Additionally, it is more efficient to query only the attributes of Gaussians associated with deformable objects, as the attributes of static areas remain unchanged over time.

3.2 Scaffold-GS

Scaffold-GS (Lu et al. 2023) adheres to the framework of 3DGS for static scenes and introduces a storage-friendly anchor-based strategy. Scaffold-GS derives Gaussians from the anchors and deduces the attributes from the attributes of the attached anchors through MLPs, rather than storing Gaussians directly. To be specific, given an anchor point located at \mathbf{x}^a and its attributes $\mathcal{A} = \{f \in \mathbb{R}^d, l \in \mathbb{R}^6, \mathbf{o} \in \mathbb{R}^{3K}\}$, where f , l , and \mathbf{o} are the anchor feature, attribute scaling, and Gaussian offsets of anchor \mathbf{x}^a , respectively. K is the number of Gaussian offsets belonging to anchor \mathbf{x}^a . During rendering, f^a are fed into MLPs to generate attributes \mathcal{A} for Gaussians. The location of each Gaussian are calculated by adding \mathbf{x} and Gaussian offsets \mathbf{o} and l is applied to scale the location and shape of the Gaussians.

Scaffold-GS achieves a similar rendering speed as the original 3DGS for the static scenes and the storage require-

ments are significantly reduced as Scaffold-GS only stores anchor points and MLPs for each scenes. The number of anchor points is much fewer than the number of Gaussians. As discussed earlier, the bottleneck of the rendering speed for dynamic scenes is the number of Gaussians. To this end, we extend the anchor-based strategy to the dynamic scene reconstruction for faster rendering speed.

4 Method

Our goal is to reconstruct a dynamic scene from a monocular video. We present the geometry and appearance of the dynamic scene with sparse anchor-grid initialization while modeling the deformation based on the sparse representation. The time-invariant attributes of Gaussians, *e.g.*, color and opacity, are decoded by the tiny MLPs. Other attributes of the deformable objects, *e.g.*, location, scale, and quaternion, vary across time, whereas those attributes of the static objects are consistent across time. We filter the anchors of the static objects with a tiny MLP and only the anchors of deformation objects are fed into corresponding MLPs for time-variant attributes. We formulate the movements of deformable Gaussians in a sparse semantic manner. Specifically, we query the movement of the deformed anchors at each time step and the movements of the assigned Gaussian offsets are calculated by the radial basis function (RBF) kernel. Please note that we do not need any additional 2D or 3D masks for supervision.

4.1 Anchor-grid Initialization

The static anchor-grid representation is initialized by the sparse point cloud from COLMAP (Schonberger and Frahm 2016). To be specific, given the 3D point cloud $\mathbf{P} \in \mathbb{R}^{M \times 3}$, we first voxelize the scene by:

$$\mathbf{A} = \left\lfloor \frac{\mathbf{P}}{\Delta d} \right\rfloor \cdot \Delta d, \quad (5)$$

where $\mathbf{A} = \{\mathbf{a}_0, \mathbf{a}_1, \dots, \mathbf{a}_{N-1}\} \in \mathbb{R}^{N \times 3}$ is the structured anchor-grid, $\lfloor \cdot \rfloor$ is the floor operation, and Δd is the voxel size. The floor operation removes the redundant and over-dense points \mathbf{P} .

4.2 Gaussian Attribute Derivation

We first introduce how to derive Gaussian’s attribute, *i.e.*, opacity, quaternion, scaling, and color from anchor points at various timestep t . The attributes of each Gaussian are decoded from its feature and time positional encoding γ by tiny MLPs $\Phi_*(\cdot)$, where $\Phi_*(\cdot)$ is the tiny MLP for the attribute $*$. Given one anchor \mathbf{a} , there are K attributes decoded by $\Phi_*(\cdot)$ where K is the number of Gaussians belonging to the anchor \mathbf{a} . Time-invariant attributes, *i.e.*, opacity and color, are decoded by $\Phi_\alpha(f_{\mathbf{a}})$ and $\Phi_c(f_{\mathbf{a}})$, where $f_{\mathbf{a}}$ are anchor’s feature. Quaternion, scale, and Gaussian offsets vary across time. For anchor \mathbf{a} with scale $s_{\mathbf{a}}$, the scale of its belonging k^{th} Gaussian offsets at time t is calculated by

$$s_{\mathbf{o}}^{k,t} = s_{\mathbf{a}} + s_{\mathbf{o}}^k + \Delta s_{\mathbf{o}}^{k,t}, \quad (6)$$

where $s_{\mathbf{o}}^k \in \{s_{\mathbf{o}}^0, s_{\mathbf{o}}^1, \dots, s_{\mathbf{o}}^{K-1}\}$ are K learnable scales belonging to anchor \mathbf{a} and $\Delta s_{\mathbf{o}}^{k,t}$ is the offset variation at time

t and $\Delta s_{\mathbf{o}}^{k,t}$ is calculated by

$$\Delta s_{\mathbf{o}}^{k,t} = \Phi_s(\text{cat}\{f_{\mathbf{a}}, \gamma(t)\}), \quad (7)$$

where cat is the concatenation operation. The quaternion $r_{\mathbf{o}}^{k,t}$ is calculated in a similar manner.

4.3 Gaussian Offsets Derivation with RBF Kernel

As formulated in Eq. (4), previous methods (Wu et al. 2023; Yang et al. 2023; Huang et al. 2023) query Gaussians attributes by the deformation network \mathcal{F} . Intuitively, the movements of anchor \mathbf{a} ’s K Gaussians should be semantically aligned. For rigid objects, the deformation of K Gaussians is exactly the same as its anchor \mathbf{a} ’s movement. The strict assumption seriously affects the performance for rendering novel views of deformed objects because there are so many non-rigid objects in the rendering scenes. Another SOTA solution is to formulate the deformation of offsets based on anchors is k -nearest neighborhood method (KNN). For instance, SCGS (Huang et al. 2023) calculates the deformation of Gaussians based on k -nearest control points. However, KNN is calculated by the distance between points and the performance is degraded when the moving objects are separated after collision.

In order to formulate the deformation of objects in a semantic manner, we calculate the location of anchor \mathbf{a} ’s k^{th} Gaussian offsets at timestep t with radial basis function (RBF) kernel. To be specific, we first apply a deformation MLP \mathcal{F} to calculate the movement of anchor \mathbf{a} across time t :

$$\Delta \mathbf{x}_{\mathbf{a}}^t = \mathcal{F}(\gamma(\mathbf{x}_{\mathbf{a}}, t)), \quad (8)$$

where $\mathbf{x}_{\mathbf{a}}$ is the location of anchor \mathbf{a} and γ is the position encoding of location and time. The movement of anchor \mathbf{a} ’s k^{th} Gaussians at time t is calculated by

$$\Delta \mathbf{x}_{\mathbf{o}}^{k,t} = \mathcal{K}(f_{\mathbf{a}}, f_{\mathbf{o}}^k) \cdot \Delta \mathbf{x}_{\mathbf{a}}^t, \quad (9)$$

where $f_{\mathbf{a}}$ and $f_{\mathbf{o}}^k$ are the features of anchor \mathbf{a} and its k^{th} belonging Gaussians. $\mathcal{K}(\cdot, \cdot)$ is the kernel basis function on anchor’s feature $f_{\mathbf{a}}$ and its offsets’ feature $f_{\mathbf{o}}^k$. Here we use the common radial basis function (RBF) kernel:

$$\mathcal{K}(f_{\mathbf{a}}, f_{\mathbf{o}}^k) = \exp\left(-\frac{\|f_{\mathbf{a}} - f_{\mathbf{o}}^k\|^2}{2\sigma^2}\right), \quad (10)$$

where σ is the covariance of the RBF kernel. We set $\sigma = 1$ in all experiments. Similar to calculating scale and quaternion, the location of anchor \mathbf{a} ’s k^{th} Gaussian is calculated as

$$\mathbf{x}_{\mathbf{o}}^{k,t} = \mathbf{x}_{\mathbf{a}} + \mathbf{x}_{\mathbf{o}}^k + \Delta \mathbf{x}_{\mathbf{o}}^{k,t}, \quad (11)$$

4.4 Time Mask

As discussed earlier, the scales, quaternion, and location of Gaussians belonging to static scenes are not changed over time. The rendering speed would be faster if we first filter the anchors belonging to static scenes and then only query time-variant attributes by MLPs for anchors that are responsible for deformation objects. To this end, we introduce a time mask MLP Φ_{mask} to filter the anchors responsible for the deformation objects. Φ_{mask} is a tiny binary-classifier. All anchors’ features $f_{\mathbf{a}} = \{f_{\mathbf{a}}^0, f_{\mathbf{a}}^1, \dots, f_{\mathbf{a}}^{N-1}\} \in \mathbb{R}^{N \times d}$

Method	Sieve			Plate			Bell			Press		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
3D-GS (Kerbl et al. 2023)	23.16	0.8203	0.2247	16.14	0.6970	0.4093	21.01	0.7885	0.2503	22.89	0.8163	0.2904
TiNeuVox (Fang et al. 2022)	21.49	0.8265	0.3176	20.58	0.8027	0.3317	23.08	0.8242	0.2568	24.47	0.8613	0.3001
HyperNeRF (Park et al. 2021b)	25.43	0.8798	0.1645	18.93	0.7709	0.2940	23.06	0.8097	0.2052	26.15	0.8897	0.1959
NeRF-DS (Yan, Li, and Lee 2023)	25.78	0.8900	0.1472	20.54	0.8042	0.1996	23.19	0.8212	0.1867	25.72	0.8618	0.2047
4DGS (Wu et al. 2023)	26.11	0.9193	0.1107	20.41	0.8311	0.2010	25.70	0.9088	0.1103	26.72	0.9031	0.1301
SCGS (Huang et al. 2023)	25.93	0.9187	0.1194	20.17	0.8257	0.2104	25.97	0.9172	0.1167	26.57	0.8971	0.1367
Deformable 3DGS (Yang et al. 2023)	25.70	0.8715	0.1504	20.48	0.8124	0.2224	25.74	0.8503	0.1537	26.01	0.8646	0.1905
Ours	27.12	0.9271	0.1151	21.21	0.8957	0.1873	26.01	0.9203	0.1204	26.61	0.9054	0.1313
Method	Cup			As			Basin			Mean		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
3D-GS (Kerbl et al. 2023)	21.71	0.8304	0.2548	22.69	0.8017	0.2994	18.42	0.7170	0.3153	20.29	0.7816	0.2920
TiNeuVox (Fang et al. 2022)	19.71	0.8109	0.3643	21.26	0.8289	0.3967	20.66	0.8145	0.2690	21.61	0.8234	0.2766
HyperNeRF (Park et al. 2021b)	24.59	0.8770	0.1650	25.58	0.8949	0.1777	20.41	0.8199	0.1911	23.45	0.8488	0.1990
NeRF-DS (Yan, Li, and Lee 2023)	24.91	0.8741	0.1737	25.13	0.8778	0.1741	19.96	0.8166	0.1855	23.60	0.8494	0.1816
4DGS (Wu et al. 2023)	24.57	0.9102	0.1185	26.30	0.8917	0.1499	19.01	0.8277	0.1631	24.18	0.8845	0.1405
SCGS (Huang et al. 2023)	24.32	0.9121	0.1207	26.17	0.8851	0.1491	19.23	0.8379	0.1514	24.05	0.8848	0.1439
Deformable 3DGS (Yang et al. 2023)	24.86	0.8908	0.1532	26.31	0.8842	0.1783	19.67	0.7934	0.1901	24.11	0.8524	0.1769
Ours	25.08	0.9132	0.1225	26.65	0.9015	0.1472	19.91	0.8351	0.1640	24.65	0.8998	0.1411

Table 1: **Quantitative comparison on NeRF-DS dataset per-scene.** We color each cell as **best**, **second best**, and **third best**.

are fed into Φ_{mask} . Only the anchors with output label 1 are further fed into corresponding MLPs for time-variant attributes. We further introduce a time mask regularization term \mathcal{L}_{t-mask} to optimize all the learnable parameters with few anchors that are responsible for deformation objects:

$$\mathcal{L}_{t-mask} = \sum_{n=0}^{N-1} \Phi_{mask}(f_{\mathbf{a}}^n)/N, \quad (12)$$

where N is the number of anchors.

4.5 Densify and Pruning for Anchors

Though applying sparse anchor points improves the rendering speed greatly, we still need to densify the 3D Gaussians for higher rendering performance. Similar to (Lu et al. 2023; Kerbl et al. 2023), we compute the gradients of the anchor points over 100 iterations and gradients greater than threshold ϵ will be densified. The voxel of the densified anchor points are divided into m sub-voxels and the anchor points are placed to each sub-voxel. The anchor feature of each sub-voxel are copied from the densified anchor and K offsets are further initialized for the anchor of each sub-voxel. For pruning redundant anchor points, we calculate the average opacity of the Gaussian offsets for each anchor point. The anchor is removed if the average opacity of its belonging Gaussian offsets is lower than the pre-defined threshold.

4.6 Loss Function

Similar to previous methods (Kerbl et al. 2023; Huang et al. 2023; Yang et al. 2023; Wu et al. 2023), we optimize the feature f each anchor \mathbf{a} and its offsets $\mathbf{o}_{\mathbf{a}}$, tiny MLPs F_{α} , F_r , F_s , and F_c for the attributes of Gaussians, and anchor’s deformable network with respect to the \mathcal{L}_1 loss and SSIM loss \mathcal{L}_{SSIM} with a time mask regularization term \mathcal{L}_{t-mask} over the rendered RGB images to supervise the training process. Our final loss function is defined as:

$$\mathcal{L} = (1 - \lambda)\mathcal{L}_1 + \lambda\mathcal{L}_{SSIM} + \lambda_t\mathcal{L}_{t-mask}, \quad (13)$$

where $\lambda = 0.2$ and $\lambda_t = 0.2$ in our all experiments. It should be noted that during the training process, the position of each

anchor $\mathbf{a} \in \mathbf{A}$ remains frozen and does not undergo any updates.

Model	PSNR(dB)↑	MS-SSIM↑	FPS↑
Nerfies (Park et al. 2021a)	22.2	0.803	< 1
HyperNeRF (Park et al. 2021b)	22.4	0.814	< 1
TiNeuVox-B (Fang et al. 2022)	24.3	0.836	1
3D-GS (Kerbl et al. 2023)	19.7	0.680	32
FFDNeRF (Guo et al. 2023)	24.2	0.842	0.05
Deformable 3DGS (Yang et al. 2023)	25.0	0.822	13
4DGS (Wu et al. 2023)	25.2	0.845	34
SCGS (Huang et al. 2023)	24.6	0.813	12
Ours	25.7	0.860	117

Table 2: **Quantitative results on HyperNeRF’s (Park et al. 2021b) vrig dataset.** The rendering resolution is set to 960×540 . We color each cell as **best**, **second best**, and **third best**.

5 Experiments

5.1 Experimental Setup

Dataset and metrics. We conducted extensive experiments on two real-world datasets: the HyperNeRF dataset (Park et al. 2021b) and the NeRF-DS dataset (Yan, Li, and Lee 2023). The division on the training and testing subsets and other experimental protocols are perfectly aligned with the original papers. The metrics we apply to evaluate the performance are Peak Signal-to-Noise Ratio (RSNR), Structural Similarity (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al. 2018). Apart from these commonly used metrics, we additionally report the training time and the rendering speed (FPS) for model compactness and efficiency. We report the metrics per scene on the NeRF-DS dataset (Yan, Li, and Lee 2023) and the averaged metrics over all scenes on the HyperNeRF (Park et al. 2021b) dataset.

Baselines and implementation. To evaluate the performance of novel view synthesis for real-world dynamic scenes, we conducted benchmarks against several state-of-the-art methods in the field, including NeRF-based methods (Fang et al. 2022; Park et al. 2021b; Yan, Li, and Lee

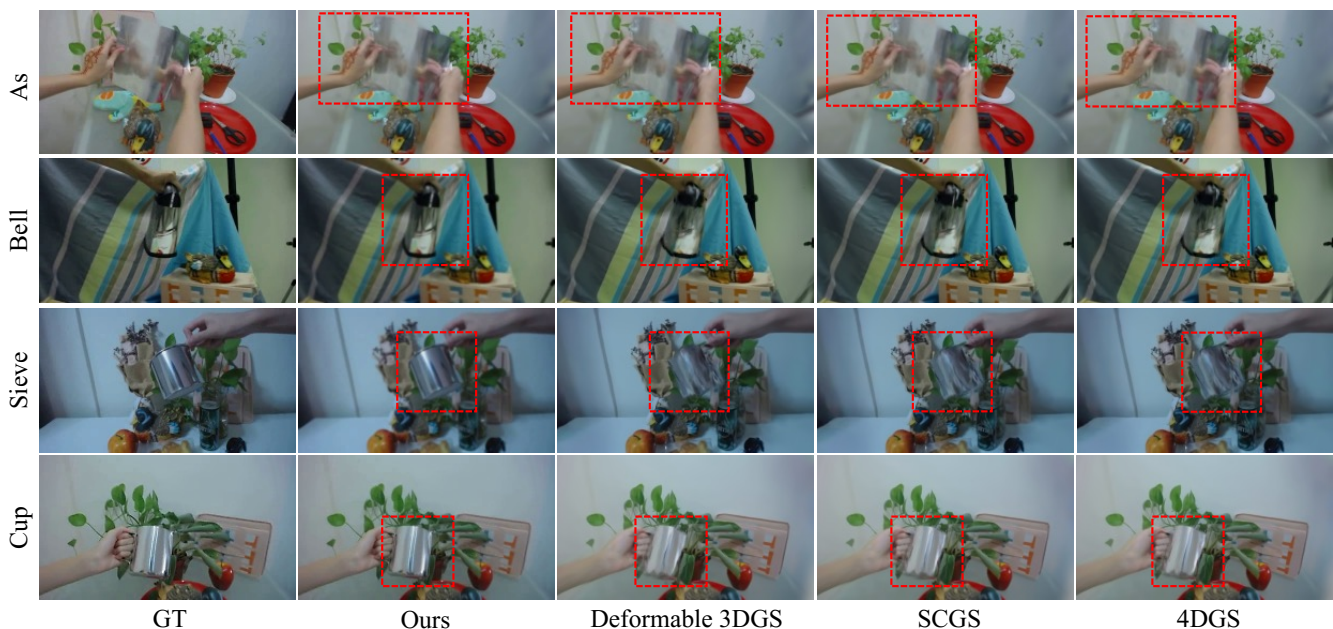


Figure 3: **Qualitative comparison on the NeRF-DS dataset (Yan, Li, and Lee 2023)**. Compared with other SOTA methods, our method reconstructs finer details and produces a structured rendering of the moving objects, *e.g.*, the cup on human’s hand.

2023; Park et al. 2021a; Song et al. 2023; Attal et al. 2023; Fridovich-Keil et al. 2023; Lin et al. 2023; Wang et al. 2023) and 3DGS-based methods (Kerbl et al. 2023; Huang et al. 2023; Yang et al. 2023; Wu et al. 2023). Our implementation is primarily based on PyTorch (Paszke et al. 2019) framework and evaluated by a Nvidia V100 GPU. Most of our hyper-parameters follow 3DGS (Kerbl et al. 2023). Our pipeline is trained for 30k iterations. We set $k = 10$ offset Gaussians for each anchor and the dimension of the anchor features f_a and offset features f_o is 8. The tiny MLPs Φ_* for Gaussian attributes are two-layer MLPs and the dimension of the middle layer is 64. The anchor’s deformation MLP is four-layer fully connected layers that employ ReLU activation and the dimension of intermediate layers is 128. The voxel size for the initialization of anchors is 0.6.

5.2 Quantitative Comparisons

NeRF-DS dataset. We first compare our method with baselines using monocular real-world NeRF-DS (Yan, Li, and Lee 2023) dataset. The experimental results on the NeRF-DS dataset, presented in Tab. 1, clearly demonstrate the superiority of our proposed method across multiple evaluation metrics and scenes. our method achieves the highest mean PSNR (24.39) and SSIM (0.8873) across all scenes, indicating its superior performance and consistency. Moreover, with the help of the structured anchor and the formulation between anchor and Gaussian offsets, our method achieves high rendering performance for moving objects. Please refer to Fig. 3 for visualization.

HyperNeRF dataset. We compare our method with other baselines on the HyperNeRF dataset (Park et al. 2021b). The experimental results on HyperNeRF’s VRIG dataset, shown




Figure 4: **Qualitative comparison on the HyperNeRF dataset (Park et al. 2021b)**. Our EDGS reconstructs detailed texture and reliable structure compared with other SOTA methods.

in Tab. 2, highlight the superior performance of our proposed method. Achieving the highest PSNR of 25.7 dB and MS-SSIM of 0.860, our approach ensures exceptional image quality and structural fidelity. Furthermore, it demonstrates remarkable efficiency with a rendering time of just 20 minutes and an impressive 117 FPS, significantly outperforming other methods. Additionally, our method requires only 7K Gaussians which need to query time-variant attributes by MLP, the lowest among all compared techniques, underscoring its computational efficiency. These results validate the robustness and practicality of our EDGS for dynamic scene reconstruction.




Figure 5: **Visualizaition of the difference map (diff.) and the optical flow with fixed camera views.** We synthesis fixed novel view across time for (Yang et al. 2023; Wu et al. 2023) and ours. The 1st row is the rendered frames at various time steps. The 2nd and 3rd rows are the difference map between t^{th} frame and the 1st frame and the optical flow, respectively. The response in the highlighted red area indicates that the static area rendered by deformable GS and 4DGS is jittering. Our method achieves better quality for static and dynamic objects.

5.3 Qualitative Comparisons

We conduct qualitative comparisons to illustrate the advantages of our method over other SOTA methods. The comparisons on the NeRF-DS dataset are shown in Fig. 3. Please zoom in to the highlighted red region for rendering comparisons of the moving objects. Compared with other SOTA methods, our method reconstructs finer details and produces a more structured rendering of the moving objects, *e.g.*, the cup in hand.

We also visualize the rendered results with fixed cameras and show the difference map between each frame and the 1st frame and the optical flow of the rendered video in Fig. 5. As highlighted in the red box in Fig. 5, deformable 3DGS and 4DGS fail to render the static area. The difference map and the optical flow at the static area indicates that the static area is jittering for those methods. We render the static and

dynamic objects more accurate with fewer jittering issues because of the time-mask MLP.

We also visualize the mask predicted by the time-mask MLP in Fig. 6. The visualization comparisons on the HyperNeRF dataset (Park et al. 2021b) are shown in Fig. 4. Compared to other SOTA methods, our EDGS reconstructs finer details (*e.g.*, the red chicken toy and the banana in hand) and produces a more structured rendering of moving objects.

5.4 Ablation Studies

Efficacy of anchor-voxel strategy and time-mask MLP.

We assess the efficiency of the anchor-voxel strategy and time-mask MLP on three scenes from the NeRF-DS (Yan, Li, and Lee 2023) dataset. Without the anchor-voxel strategy, rendering speed drops to around 20 FPS due to redundant densified Gaussians, showing that the bottleneck in ren-

dering deformation scenes is the number of Gaussians. Introducing the time-mask MLP further improves rendering speed by filtering out anchors corresponding to static areas, allowing only those in dynamic areas to be queried for time-variant attributes. This also enhances rendering quality (e.g. PSNR increases from 26.18 to 27.12 on the sieve scene), likely because the MLPs are optimized with fewer anchor points, making the optimization process more effective. See Fig. 6 for the time mask visualization.

scene	As		Bell		Sieve	
	PSNR	FPS	PSNR	FPS	PSNR	FPS
w/o anchor-grid init	24.57	27	24.31	23	24.74	17
w/o time-mask	26.02	127	25.78	132	26.18	122
full model	26.65	152	26.01	147	27.12	151

Table 3: **Effects of anchor-voxel strategy and time-mask MLP.** The anchor-grid strategy is crucial for rendering speed due to the reduced number of Gaussians. The time-mask MLP filters time-variant anchors, further reducing the number of anchors required for time-variant attribute queries.

Efficacy of time-variant Gaussian offsets derivation with various strategies. Tab. 4 compares time-variant Gaussian offsets derived using different strategies across three scenes on the NeRF-DS (Yan, Li, and Lee 2023) dataset. Regarding the rigid transformation strategy, we assume that the movements of the anchors and their corresponding Gaussians same. The lowest performance shows that simply regarding all deformation objects as rigid ones degrades the performance. We further calculate the similarity between the feature of anchor and Gaussian offsets using KNN method and cosine similarity. The KNN strategy performs well for the reconstruction of single deformation objects (Huang et al. 2023) and the performance degrades on the real-world scenes. Cosine similarity achieves similar performance compared with RBF kernel but is VRAM-consuming because the similarity is calculated by a $N \times (KN)$ huge matrix.

scene	As		Bell		Sieve	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
rigid transformation	23.51	0.8259	24.21	0.8571	23.95	0.8750
KNN	25.75	0.8954	25.53	0.8816	25.68	0.8907
cosine similarity	26.31	0.8928	25.74	0.8807	26.57	0.9155
RBF kernel	26.65	0.9015	26.01	0.9203	27.12	0.9271

Table 4: **Gaussian offsets derivation with various strategies.** We conduct four different strategies for formulating the variation of the Gaussian offsets across time.

Visualization of anchor features and time mask. We perform an analysis of the learnable anchor features to assess their effectiveness in dynamic scene reconstruction. As shown in Fig. 6, the clustered patterns indicate that the anchor features learn and encode similar semantic meanings across both static and dynamic objects of the scene. Additionally, the time-mask MLP Φ_{mask} effectively classifies time-variant anchors without any mask supervision, demonstrating its ability to adaptively distinguish between static and dynamic regions.

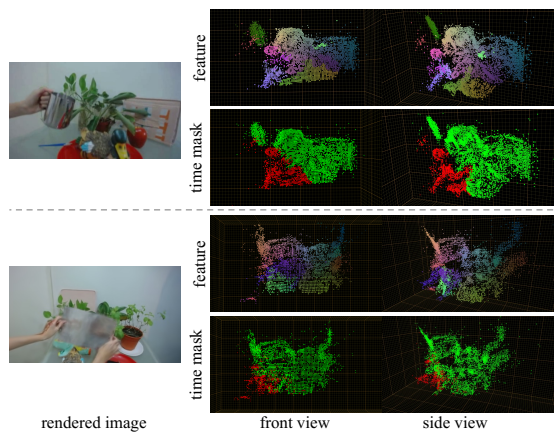


Figure 6: **Visualization of anchor features and time mask.** We visualize the rendered images, anchor features, and time masks for two deformation scenes. Anchor features are visualized using the UMAP function. In the time-mask visualization, the time masks are predicted by time-mask MLP Φ_{mask} . **Green** anchors represent static scenes, while **red** anchors indicate deformation scenes.

6 Conclusion

In conclusion, we develop Efficient Dynamic Gaussian Splatting (EDGS). It efficiently models dynamic scenes by emphasizing sparse, time-variant attributes and selectively processing static objects. This strategy significantly reduces computational complexity while preserving high rendering quality. Additionally, the incorporation of a classical kernel for motion flow optimization further enhances this process. Our evaluations using the NeRF-DS and HyperNeRF datasets show that EDGS not only achieves faster rendering speeds but also higher Peak Signal-to-Noise Ratio (PSNR) scores, surpassing current state-of-the-art methods.

Acknowledgements

This project is supported by the Ministry of Education, Singapore, under its Academic Research Fund Tier 2 (Award Number: MOE-T2EP20122-0006), and the National Research Foundation, Singapore, under its Medium Sized Center for Advanced Robotics Technology Innovation.

References

- Attal, B.; Huang, J.-B.; Richardt, C.; Zollhoefer, M.; Kopf, J.; O’Toole, M.; and Kim, C. 2023. HyperReel: High-fidelity 6-DoF video with ray-conditioned sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16610–16620.
- Cao, A.; and Johnson, J. 2023. Hexplane: A fast representation for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 130–141.
- Fang, J.; Yi, T.; Wang, X.; Xie, L.; Zhang, X.; Liu, W.; Nießner, M.; and Tian, Q. 2022. Fast dynamic radiance fields with time-aware neural voxels. In *SIGGRAPH Asia 2022 Conference Papers*, 1–9.

- Fridovich-Keil, S.; Meanti, G.; Warburg, F. R.; Recht, B.; and Kanazawa, A. 2023. K-planes: Explicit radiance fields in space, time, and appearance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12479–12488.
- Fridovich-Keil, S.; Yu, A.; Tancik, M.; Chen, Q.; Recht, B.; and Kanazawa, A. 2022. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5501–5510.
- Guo, X.; Sun, J.; Dai, Y.; Chen, G.; Ye, X.; Tan, X.; Ding, E.; Zhang, Y.; and Wang, J. 2023. Forward Flow for Novel View Synthesis of Dynamic Scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 16022–16033.
- Huang, Y.-H.; Sun, Y.-T.; Yang, Z.; Lyu, X.; Cao, Y.-P.; and Qi, X. 2023. SC-GS: Sparse-Controlled Gaussian Splatting for Editable Dynamic Scenes. *arXiv preprint arXiv:2312.14937*.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4): 1–14.
- Kong, H.; Lian, D.; Mi, M. B.; and Wang, X. 2025. Dream-Drone: Text-to-Image Diffusion Models Are Zero-Shot Perpetual View Generators. In *European Conference on Computer Vision*, 324–341. Springer.
- Lin, H.; Peng, S.; Xu, Z.; Xie, T.; He, X.; Bao, H.; and Zhou, X. 2023. High-Fidelity and Real-Time Novel View Synthesis for Dynamic Scenes. In *SIGGRAPH Asia Conference Proceedings*.
- Lu, T.; Yu, M.; Xu, L.; Xiangli, Y.; Wang, L.; Lin, D.; and Dai, B. 2023. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. *arXiv preprint arXiv:2312.00109*.
- Ma, X.; Fang, G.; and Wang, X. 2024. Deepcache: Accelerating diffusion models for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15762–15772.
- Park, K.; Sinha, U.; Barron, J. T.; Bouaziz, S.; Goldman, D. B.; Seitz, S. M.; and Martin-Brualla, R. 2021a. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5865–5874.
- Park, K.; Sinha, U.; Hedman, P.; Barron, J. T.; Bouaziz, S.; Goldman, D. B.; Martin-Brualla, R.; and Seitz, S. M. 2021b. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *arXiv preprint arXiv:2106.13228*.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Schonberger, J. L.; and Frahm, J.-M. 2016. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4104–4113.
- Song, L.; Chen, A.; Li, Z.; Chen, Z.; Chen, L.; Yuan, J.; Xu, Y.; and Geiger, A. 2023. Nerfplayer: A streamable dynamic scene representation with decomposed neural radiance fields. *IEEE Transactions on Visualization and Computer Graphics*, 29(5): 2732–2742.
- Wang, F.; Chen, Z.; Wang, G.; Song, Y.; and Liu, H. 2023. Masked Space-Time Hash Encoding for Efficient Dynamic Scene Reconstruction. *Advances in neural information processing systems*.
- Wang, S.; Yang, X.; Shen, Q.; Jiang, Z.; and Wang, X. 2025. GFlow: Recovering 4D World from Monocular Video. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Wu, G.; Yi, T.; Fang, J.; Xie, L.; Zhang, X.; Wei, W.; Liu, W.; Tian, Q.; and Wang, X. 2023. 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528*.
- Yan, Z.; Li, C.; and Lee, G. H. 2023. Nerf-ds: Neural radiance fields for dynamic specular objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8285–8295.
- Yang, Z.; Gao, X.; Zhou, W.; Jiao, S.; Zhang, Y.; and Jin, X. 2023. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. *arXiv preprint arXiv:2309.13101*.
- Yu, H.-X.; Duan, H.; Hur, J.; Sargent, K.; Rubinstein, M.; Freeman, W. T.; Cole, F.; Sun, D.; Snively, N.; Wu, J.; et al. 2024. Wonderjourney: Going from anywhere to everywhere. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6658–6667.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.