

MEGA: Memory-Efficient 4D Gaussian Splatting for Dynamic Scenes

Xinjie Zhang^{1*} Zhening Liu¹ Yifan Zhang^{2†} Xingtong Ge¹ Dailan He³
Tongda Xu⁵ Yan Wang⁵ Zehong Lin¹ Shuicheng Yan^{4,2†} Jun Zhang^{1†}

¹iComAI Lab, The Hong Kong University of Science and Technology ²Skywork AI

³The Chinese University of Hong Kong ⁴National University of Singapore

⁵Institute for AI Industry Research (AIR), Tsinghua University

{xzhangga, zhening.liu}@connect.ust.hk, yifan.zhang7@kunlun-inc.com, xingtong.ge@gmail.com

hedailan@link.cuhk.edu.hk, x.tongda@nyu.edu, wangyan202199@163.com, eezhlin@ust.hk

yansc@comp.nus.edu.sg, eejzhang@ust.hk

Abstract

4D Gaussian Splatting (4DGS) has recently emerged as a promising technique for capturing complex dynamic 3D scenes with high fidelity. It utilizes a 4D Gaussian representation and a GPU-friendly rasterizer, enabling rapid rendering speeds. Despite its advantages, 4DGS faces significant challenges, notably the requirement of millions of 4D Gaussians, each with extensive associated attributes, leading to substantial memory and storage cost. This paper introduces a memory-efficient framework for 4DGS. We streamline the color attribute by decomposing it into a per-Gaussian direct color component with only 3 parameters and a shared lightweight alternating current color predictor. This approach eliminates the need for spherical harmonics coefficients, which typically involve up to 144 parameters in classic 4DGS, thereby creating a memory-efficient 4D Gaussian representation. Furthermore, we introduce an entropy-constrained Gaussian deformation technique that uses a deformation field to expand the action range of each Gaussian and integrates an opacity-based entropy loss to limit the number of Gaussians, thus forcing our model to use as few Gaussians as possible to fit a dynamic scene well. With simple half-precision storage and zip compression, our framework achieves a storage reduction by approximately 190× and 125× on the Technicolor and Neural 3D Video datasets, respectively, compared to the original 4DGS. Meanwhile, it maintains comparable rendering speeds and scene representation quality, setting a new standard in the field. Code is available at <https://github.com/Xinjie-Q/MEGA>.

1. Introduction

Dynamic scene reconstruction from multi-view videos is gaining widespread interest in computer vision and graphics due to its broad applications in virtual reality (VR), augmented reality (AR), and 3D content production. The emergence of neural radiance field (NeRF) [31] enables high-quality novel view synthesis from multi-view image inputs. It has been further extended to represent dynamic scenes by modeling a direct mapping from spatio-temporal coordinates to color and density [3, 22, 36]. Despite the impressive visual quality of NeRF-based methods, they require dense sampling along rays, leading to slow rendering speeds that hinder practical applications.

The recent introduction of 3D Gaussian Splatting (3DGS) [17] marks a significant shift in the field of novel view synthesis. This approach incorporates the explicit 3D Gaussian representation and differentiable tile-based rasterization to enable real-time rendering speeds that significantly outperform NeRF-based methods. Built on this framework, subsequent studies have developed 4D Gaussian Splatting (4DGS) [10, 46], which conceptualizes scene variations across different timestamps as 4D spatio-temporal Gaussian hyper-cylinder. As shown in Fig. 2, when depicting a 3D scene at a given timestamp, these 4D Gaussians will first be sliced into 3D Gaussians with time-varying positions and opacity. Then, the 3D Gaussians with the temporal decay opacity below a specific threshold are filtered out. This filtering operation helps 4DGS to describe the transient content such as emerging or vanishing objects. Finally, following 3DGS, the remaining 3D Gaussians are projected onto 2D screens through fast rasterization. By directly optimizing a collection of 4D Gaussians, 4DGS effectively captures both static and dynamic scene elements, thereby achieving photorealistic visual quality.

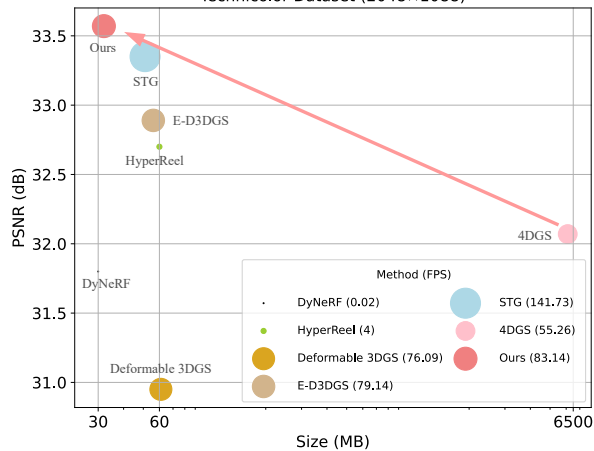
However, 4DGS requires millions of Gaussians to ad-

*This work was partially performed when Xinjie Zhang was an Intern at Skywork AI.

†Corresponding Authors.



(a) High performance at the *Birthday* scene. Technicolor Dataset (2048×1088)



(b) Comparison on quality, size, and speed.

Figure 1. Our approach significantly reduces storage requirements while maintaining comparable photorealistic quality and real-time rendering speed with 4D Gaussian Splatting (4DGS) [46]. The core idea is to develop a memory-efficient 4D Gaussian representation and use as few Gaussians as possible to fit dynamic scenes well. (a) 4DGS requires up to 13 million Gaussians to render the *Birthday* scene, whereas our method only needs 0.91 million Gaussians. (b) Quantitative comparisons of rendering quality, storage size, and speed against various competitive baselines on the Technicolor dataset.

equately represent dynamic scenes with high fidelity. As depicted in Fig. 1 (a), rendering the *Birthday* scene necessitates up to 13 million Gaussian points, leading to a storage overhead of approximately 7.79GB. This substantial storage and transmission challenge can severely limit the practical applications of 4DGS, particularly on resource-constrained devices. For example, the significant memory requirements may make it impractical to store, transmit, and render various scenes on AR/VR headsets. Consequently, it is of critical importance to compress 4D Gaussians to minimize the memory footprint of 4DGS while preserving high-quality scene representation and reconstruction.

To address the significant memory and storage challenges associated with 4DGS, we propose a **Memory-Efficient 4D Gaussian Splatting (MEGA)** framework. In the original 4D Gaussian representation, 144 out of the to-




Figure 2. Illustration of temporal slicing in 4DGS, with the z-axis omitted for simplicity. A 4D Gaussian can be conceptualized as a hyper-cylinder in 4D space. Given the specific time query, a corresponding 3D Gaussian ellipsoid is extracted from this hyper-cylinder. The depth of color in the 3D Gaussian ellipsoid represents its temporal opacity. Those 3D Gaussian ellipsoids with temporal opacity below a predefined threshold are excluded from the scene rendering.

tal 161 parameters are 4D spherical harmonics (SH) coefficients, which occupy the majority of the storage space and exhibit considerable redundancy. To develop a memory-efficient 4D Gaussian representation, we draw inspiration from the concepts of Direct Current (DC) and Alternating Current (AC) in electrical engineering, which symbolize the steady and varying components, respectively. Specifically, we decouple the color attribute into a per-Gaussian DC color component and a shared temporal-viewpoint aware AC color predictor. This predictor is capable of accurately estimating the color variations of a Gaussian at given times and viewing angles, thereby effectively preserving visual quality. It is noteworthy that our DC color component requires only 3 parameters, while the predictor utilizes a lightweight multi-layer perceptron (MLP) with three linear layers. Consequently, this modification achieves a compression ratio of approximately $8\times$ relative to the original 4D Gaussians with equivalent Gaussian points, substantially reducing the storage demands of the Gaussian representation.

Nevertheless, compacting the properties of the 4D Gaussian alone cannot effectively alleviate the problem of excessive number of Gaussians required. Existing 4DGS baselines [10, 46] assume that each sliced 4D Gaussian exhibits only linear movement over time while maintaining constant covariance, which means that the complex motion in the scene has to be modeled by a combination of multiple Gaussians. Moreover, as illustrated in Fig. 4 (a), only about 6% of Gaussians actively participate in rendering at any given time, because the temporal decay opacity forces each Gaussian to be visible only near its mean time center and invisible at other times. These inherent properties significantly limit the effective utilization of each Gaussian, thereby increasing the number of Gaussians needed for adequate scene rendering. To overcome this limitation, we introduce an efficient entropy-constrained Gaussian deformation field designed to expand the operational range of 4D Gaussians. This deformation model leverages both tempo-

ral and viewpoint information to accurately represent Gaussian motion, shape, and transience changes. Meanwhile, a spatial opacity-based entropy loss is introduced to push the spatial opacity of each Gaussian towards binary states (either one or zero). This adjustment aids in identifying and eliminating non-essential Gaussians that contribute minimally to the overall performance. In this way, our proposed strategy not only effectively reduces the number of Gaussians, but also improves the utilization rate of the Gaussians involved in rendering given the time and viewing angle. Finally, to store the parameters of our streamlined 4DGS, we employ 16-bit floating-point (FP16) precision with zip delta compression algorithm to achieve further reductions in memory footprint. In summary, our main contributions are three-fold:

- To the best of our knowledge, we are among the first to develop a memory-efficient framework for 4D Gaussian Splatting. By decomposing the color attribute into a per-Gaussian DC color component and a lightweight, temporal-viewpoint aware AC color predictor, we successfully eliminate the need for redundant spherical harmonics coefficients.
- We introduce an entropy-constrained Gaussian deformation technique to enhance the potential of each 4D Gaussian for depicting complex scene motion. This approach not only substantially reduces the number of Gaussians but also improves their utilization rate. Moreover, we integrate straightforward post-processing techniques, such as FP16 precision and zip delta compression, to further decrease storage overhead.
- Extensive experimental results demonstrate that our proposed method achieves significant storage reductions—approximately 190× and 125× on the Technicolor and Neural 3D Video datasets, respectively—while maintaining comparable quality of scene representation and rendering speed relative to the original 4DGS.

2. Related Works

Neural Rendering for Static Scenes. Recently, the advent of neural rendering has attracted increasing interest in 3D scene representation and reconstruction. NeRF, pioneered by [31], represents the volume density and view-dependent emitted radiance of a 3D scene as a function of 5D coordinates (3D position and 2D viewing direction) using an MLP. However, the vanilla NeRF relies solely on a large MLP to store scene information, significantly limiting its training and rendering efficiency. Subsequent works have explored explicit grid-based representations [4, 12, 32, 40] to enhance training efficiency. Nonetheless, these NeRF-based methods still face challenges of slow rendering due to dense sampling for each ray. In contrast, the work [17] introduce 3D Gaussian Splatting, a novel explicit representation framework that employs a highly optimized custom

CUDA rasterizer to achieve unparalleled rendering speeds with high-fidelity novel view synthesis for complex scenes. Subsequent studies have further improved the representation efficiency [7, 15, 20, 27, 48, 50] and have been extended to various vision understanding and editing applications [5, 26, 35, 38, 43, 49].

Neural Rendering for Dynamic Scenes. Synthesizing new views of dynamic scenes from a series of 2D images captured at different times presents a significant challenge. Recent advancements have extended NeRF to handle monocular or multi-object dynamic scenes by learning a mapping from spatio-temporal coordinates to color and density [1, 3, 13, 21, 22, 24, 25, 30, 36, 39, 41]. Unfortunately, these methods suffer from low rendering efficiency. To address this issue, some recent studies [2, 9, 14, 28, 44, 45] have developed deformable 3D GS, which decouples dynamic scenes into a static canonical 3DGS and a deformation motion field to account for temporal variations in the 3D Gaussian parameters. Concurrently, a series of recent studies [10, 16, 18, 23, 46] directly learn a set of spatio-temporal Gaussians to model static, dynamic, and transient content within a scene. However, these methods require a large number of Gaussians to achieve high-quality scene modeling, which brings expensive storage overhead. To this end, our work focuses on developing effective compression techniques for 4DGS [46].

3D Gaussian Splatting Compression. Since optimized scenes in 3DGS typically comprise millions of 3D Gaussians and require up to several gigabytes of storage, various compression strategies have been proposed to reduce the size, including redundant Gaussian pruning [11, 20], spherical harmonics distillation or compactness [11, 20, 34, 42], vector quantization [11, 20, 33, 42], and entropy models [6]. However, due to the differences between 3DGS for static scene representation and 4DGS for dynamic scene representation, existing methods may be inapplicable to or unsuitable for 4DGS. In this paper, we aim to develop a more compact color representation and reduce the number of 4D Gaussians by considering temporal and viewpoint factors, thereby achieving a more efficient memory footprint.

3. Method

In Section 3.1, we first review the technique of 4DGS [46], which serves as the foundation of our method. Subsequently, in Section 3.2, we introduce how to develop our memory-efficient 4D Gaussian Splatting for modeling dynamic scenes. Finally, we detail the training process and describe how to store our compact 4DGS in Section 3.3.

3.1. Preliminary: 4D Gaussian Splatting

4D Gaussian Splatting [46] optimizes a set of anisotropic 4D Gaussians via differentiable rasterization to effectively represent a dynamic scene. With a highly efficient ras-

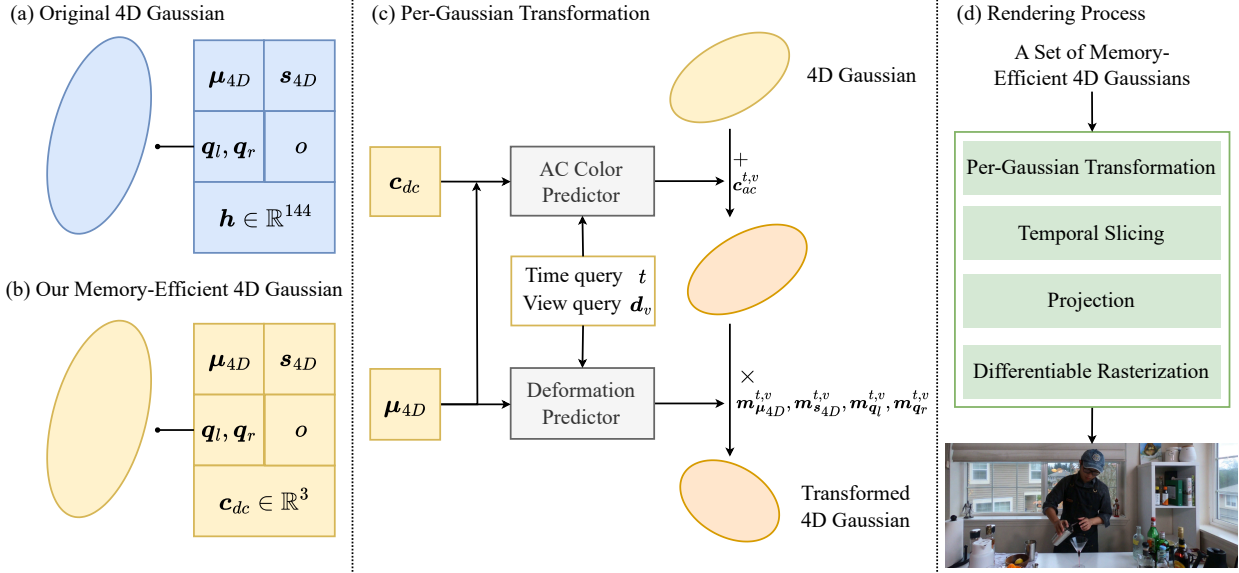


Figure 3. Overview of our proposed memory-efficient Gaussian Splatting framework. (a) The original 4D Gaussian uses 4D spherical harmonics \mathbf{h} to represent color, which is highly redundant and consumes substantial memory. (b) Our memory-efficient 4D Gaussian replaces \mathbf{h} with a compact, view-independent, and time-independent color component \mathbf{c}_{dc} , achieving an about $8\times$ reduction in storage overhead. (c) In the per-Gaussian transformation, a lightweight AC color predictor compensates for the absent viewpoint and temporal information in \mathbf{c}_{dc} , and a deformation predictor expands the action range of each Gaussian. (d) Our rendering process consists of four steps: per-Gaussian transformation, temporal slicing, projection, and differentiable rasterization.

terizer, the optimized model facilitates real-time rendering of high-fidelity novel views. Each 4D Gaussian is characterized by the following attributes: (i) 4D center $\boldsymbol{\mu}_{4D} = (\mu_x, \mu_y, \mu_z, \mu_t)^T \in \mathbb{R}^4$; (ii) 4D rotation \mathbf{R}_{4D} represented by a pair of left quaternion $\mathbf{q}_l \in \mathbb{R}^4$ and right quaternion $\mathbf{q}_r \in \mathbb{R}^4$; (iii) 4D scaling factor $\mathbf{s}_{4D} = (s_x, s_y, s_z, s_t)^T \in \mathbb{R}^4$; (iv) time- and view-dependent RGB color represented by 4D spherical harmonics coefficients $\mathbf{h} \in \mathbb{R}^{3(k_v+1)^2(k_t+1)}$ with the view degrees of freedom k_v and time degree of freedom k_t ; (v) spatial opacity $o \in [0, 1]$.

Given 4D scaling matrix $\mathbf{S}_{4D} = \text{diag}(\mathbf{s}_{4D})$ and 4D rotation matrix \mathbf{R}_{4D} , we parameterize 4D Gaussian’s covariance matrix as:

$$\boldsymbol{\Sigma}_{4D} = \mathbf{R}_{4D} \mathbf{S}_{4D} \mathbf{S}_{4D}^T \mathbf{R}_{4D}^T = \begin{pmatrix} \mathbf{U} & \mathbf{V} \\ \mathbf{V}^T & \mathbf{W} \end{pmatrix}, \quad (1)$$

where $\mathbf{U} \in \mathbb{R}^{3 \times 3}$. When rendering the scene at time t , each 4D Gaussian is sliced into 3D space. The density of the sliced 3D Gaussian at the spatial point \mathbf{x} is expressed as:

$$G_{3D}(\mathbf{x}, t) = \sigma(t) e^{-\frac{1}{2}[\mathbf{x} - \boldsymbol{\mu}_{3D}(t)]^T \boldsymbol{\Sigma}_{3D}^{-1} [\mathbf{x} - \boldsymbol{\mu}_{3D}(t)]}, \quad (2)$$

where $\boldsymbol{\Sigma}_{3D} = \mathbf{U} - \frac{\mathbf{V}\mathbf{V}^T}{\mathbf{W}}$ represents the time-invariant 3D covariance matrix. The temporal decay opacity, $\sigma(t) = e^{-\frac{(t-\mu_t)^2}{2\mathbf{W}}}$, utilizes a 1D Gaussian function to modulate the contribution of each Gaussian to the t -th scene. The time-variant 3D center, $\boldsymbol{\mu}_{3D}(t) = \boldsymbol{\mu}_{3D} + (t - \mu_t) \frac{\mathbf{V}}{\mathbf{W}}$, introduces a linear motion term to the 3D center position

$\boldsymbol{\mu}_{3D} = (\mu_x, \mu_y, \mu_z)^T$, assuming that all motions can be approximated as linear motion within a very small time range. After temporal slicing, the following process involves projecting sliced 3D Gaussians onto the 2D image plane based on depth order from specific view direction, and executing the fast differentiable rasterization to render the final image. Although this paradigm provides high-quality novel view synthesis, it necessitates large amount of Gaussians to fully reconstruct a dynamic scene, which brings unbearable storage overhead. This challenge drives our memory-efficient 4D Gaussian Splatting design.

3.2. Memory-efficient 4D Gaussian Splatting

Overview. As illustrated in Fig. 3, we develop our memory-efficient 4D Gaussian framework to significantly reduce the number of per-Gaussian stored parameters and drive the model to reconstruct dynamic scene with fewer 4D Gaussians. During the rendering process, we utilize a set of optimized 4D Gaussians and initially transform each Gaussian based on specific time and view direction. This transformation procedure involves Gaussian color prediction and geometry deformation. By modifying the geometric structure of each Gaussian, we effectively broaden its action range. This expansion not only reduces the total number of Gaussians required but also increases the rendering participation rate of each Gaussian. Following the per-Gaussian transformation, we adhere to the established protocols of the original 4DGS [46] to carry out temporal slicing, projection, and




Figure 4. (a) The ratio of Gaussians involved in rendering the *Birthday* scene at different time steps. The blue line shows how many Gaussians are involved in rendering in our MEGA model if we do not use per-Gaussian transformation. (b) Visualization of the varying number of Gaussians on the *Birthday* scene during training.

differentiable rasterization, all critical for rendering high-quality frames.

Memory-efficient 4D Gaussian. 4DGS introduces 4D spherical harmonics \mathbf{h} to model the temporal evolution of view-dependent color in dynamic scenes, which typically requires 144 of the total 161 parameters and contributes to the main storage overhead. While [20] have explored the use of a grid-based neural field to replace SH coefficients \mathbf{h} , we find that directly applying this method results in severe performance loss compared to 4DGS (see Table 3).

To overcome this issue, we propose a compact DC-AC color (DAC) representation. Specifically, we decouple the color attribute as a per-Gaussian DC color component $\mathbf{c}_{dc} \in \mathbb{R}^3$ and a temporal-viewpoint aware AC color predictor \mathcal{F}_ϕ . To predict the final color $\mathbf{c}_{t,v}$ of each Gaussian, we first compute the normalized view direction $\mathbf{d}_v = \frac{\boldsymbol{\mu}_{3D} - \mathbf{p}_v}{\|\boldsymbol{\mu}_{3D} - \mathbf{p}_v\|_2}$ for each Gaussian according to the camera center point $\mathbf{p}_v \in \mathbb{R}^3$ at the viewpoint v . Then, we concatenate the 3D position $\boldsymbol{\mu}_{3D}$, view direction \mathbf{d}_i , time t , and DC color \mathbf{c}_{dc} and input them to a lightweight MLP network \mathcal{F}_ϕ :

$$\mathbf{c}_{t,v} = \text{sigmoid}(\mathbf{c}_{dc} + \mathcal{F}_\phi(\text{sg}(\boldsymbol{\mu}_{3D}), \text{sg}(\mathbf{d}_v), t, \mathbf{c}_{dc})), \quad (3)$$

where $\text{sg}(\cdot)$ indicates a stop-gradient operation. This hybrid color composition method not only effectively preserves the individual information using DC component and supplements the missing viewpoint and time information using the AC predictor to maintain high rendering quality, but also reduces the storage overhead by up to $8\times$ compared to the original 4DGS [46].

Entropy-constrained Gaussian Deformation. For a specific time t , 4DGS [46] presupposes that the sliced 4D Gaussians exhibit linear movement while their rotation and scale remain constant. This strict assumption simplifies

the movement representation and forces the model to combine multiple extra Gaussians to present any complex non-linear motions. Moreover, the sliced 4D Gaussian introduces the temporal decay opacity σ_t . From its definition, it is analyzed that a Gaussian gradually appears as time t approaches its temporal position μ_t , peaks in opacity at $t = \mu_t$, and gradually diminishes in density as t moves away from μ_t . As shown in Fig. 4 (a), this limited temporal operation range results in more than 90% of Gaussians being excluded at each time, causing the model to densify a large amount of Gaussians for rendering high-quality scene.

To address these limitations, we advocate for improving flexibility in the motion representation and geometric structure of each 4D Gaussian. Specifically, we introduce a temporal-viewpoint aware deformation predictor to enlarge the action range of Gaussians. The 4D Gaussian center $\boldsymbol{\mu}_{4D}$, view direction \mathbf{d}_i , and time t are mapped to a high-dimensional space using a regular frequency positional encoding function γ [31], and then processed through a lightweight MLP network \mathcal{F}_θ to predict the position deformation $\mathbf{m}_{\boldsymbol{\mu}_{4D}}^{t,v} \in \mathbb{R}^4$, scale deformation $\mathbf{m}_{\mathbf{s}_{4D}}^{t,v} \in \mathbb{R}^4$, and rotation deformations $\mathbf{m}_{\mathbf{q}_l}^{t,v} \in \mathbb{R}^4, \mathbf{m}_{\mathbf{q}_r}^{t,v} \in \mathbb{R}^4$ as:

$$(\mathbf{m}_{\boldsymbol{\mu}_{4D}}^{t,v}, \mathbf{m}_{\mathbf{s}_{4D}}^{t,v}, \mathbf{m}_{\mathbf{q}_l}^{t,v}, \mathbf{m}_{\mathbf{q}_r}^{t,v}) = \mathcal{F}_\theta(\gamma(\text{sg}(\boldsymbol{\mu}_{4D})), \gamma(\text{sg}(\mathbf{d}_v)), \gamma(t)), \quad (4)$$

where γ is defined as $(\sin(2^l \pi p), \cos(2^l \pi p))_{l=0}^{L-1}$. Based on the estimated deformation for time t and viewpoint v , we transform the original 4D Gaussian to a temporal-viewpoint aware 4D Gaussian:

$$\begin{aligned} \boldsymbol{\mu}_{4D}^{t,v} &= \boldsymbol{\mu}_{4D} \times \mathbf{m}_{\boldsymbol{\mu}_{4D}}^{t,v}, & \mathbf{s}_{4D}^{t,v} &= \mathbf{s}_{4D} \times \mathbf{m}_{\mathbf{s}_{4D}}^{t,v}, \\ \mathbf{q}_l^{t,v} &= \mathbf{q}_l \times \mathbf{m}_{\mathbf{q}_l}^{t,v}, & \mathbf{q}_r^{t,v} &= \mathbf{q}_r \times \mathbf{m}_{\mathbf{q}_r}^{t,v}. \end{aligned} \quad (5)$$

Table 1. Quantitative comparison with various competitive baselines on the Technicolor Dataset. “Storage” refers to the total model size for 50 frames.

Method	PSNR \uparrow	DSSIM $_1\downarrow$	DSSIM $_2\downarrow$	LPIPS \downarrow	FPS \uparrow	Storage \downarrow
DyNeRF [22]	31.80	-	0.0210	0.1400	0.02	30.00MB
HyperReel [1]	32.70	0.0470	-	0.1090	4.00	60.00MB
Deformable 3DGS [44]	30.95	0.0696	0.0353	0.1553	76.09	61.36MB
STG [23]	33.35	0.0404	0.0187	0.0846	141.73	51.35MB
E-D3DGS [2]	32.89	0.0494	0.0231	0.1114	79.14	56.07MB
4DGS [46]	32.07	0.0535	0.0263	0.1189	55.26	6107.07MB
Ours	33.57	0.0442	0.0204	0.1014	83.14	32.45MB

Nonetheless, as depicted in Fig. 4 (b), without constraints on the number of Gaussians, a significant proliferation occurs where Gaussians are continuously split and cloned during the densification process. To force the model to use fewer Gaussians while accurately simulating complex scene changes, we introduce a spatial opacity-based entropy loss \mathcal{L}_{opa} that encourages the spatial opacity o of each Gaussian to approach one or zero:

$$\mathcal{L}_{opa} = \frac{1}{N} \sum_{j=1}^N (-o_j \log(o_j)), \quad (6)$$

where N denotes the number of Gaussians. During optimization, we actively prune Gaussians that exhibit near-zero opacity at every K iterations, which ensures efficient computation and maintains a low storage footprint throughout the training phase. Furthermore, as shown in Fig. 4 (a), with the opacity-based entropy loss \mathcal{L}_{opa} , our deformation field successfully enlarges the action range of each Gaussian, increasing the Gaussian participation ratio from less than 50% to about 75% under the same Gaussian points.

3.3. Training and Compression Pipeline

Loss Function. Following the original 4DGS [46], we adopt the photometric loss \mathcal{L}_{photo} , consisting of \mathcal{L}_1 loss and structural similarity loss \mathcal{L}_{ssim} , to measure the distortion between the rendered image and ground truth image. By adding the loss for opacity regularization \mathcal{L}_{opa} , the overall loss \mathcal{L} is defined as:

$$\mathcal{L} = \mathcal{L}_{photo} + \kappa \mathcal{L}_{opa} = (1 - \lambda) \mathcal{L}_1 + \lambda \mathcal{L}_{ssim} + \kappa \mathcal{L}_{opa}, \quad (7)$$

where both λ and κ are trade-off parameters to balance the components.

Compression Pipeline. During the optimization phase, we adopt half-precision training. After obtaining the optimized MEGA representation, we store these learnable parameters in the FP16 format, then apply the zip delta compression algorithm. This lossless compression technique typically reduces storage overhead by approximately 10%.

4. Experiments

4.1. Experimental Setup

Datasets. We evaluate the effectiveness of our method using two real-world benchmarks that are representative of various challenges in dynamic scene rendering: (1) **Technicolor Light Field Dataset** [37]: This dataset consists of multi-view video data captured by a time-synchronized 4×4 camera rig. Following HyperReel [1], we exclude the camera at the second row, second column and evaluate on five scenes (*Birthday*, *Fabien*, *Painter*, *Theater*, and *Trains*) at 2048×1088 full resolution. (2) **Neural 3D Video Dataset** (Neu3DV) [22]: This dataset includes six indoor multi-view video scenes captured by 18 to 21 cameras, each at a resolution of 2704×2028 pixels. The scenes (*Coffee Martini*, *Cook Spinach*, *Cut Roasted Beef*, *Flame Salmon*, *Flame Steak*, *Sear Steak*) vary in duration and feature dynamic movements, some with multiple objects in motion. Consistent with existing practices [22, 46], evaluations are conducted at half resolution of 300-frame scenes.

Evaluation Metrics. To assess the quality of rendered videos, we utilize three popular image quality assessment metrics: Peak Signal-to-Noise Ratio (PSNR), Dissimilarity Structural Similarity Index Measure (DSSIM), and Learned Perceptual Image Patch Similarity (LPIPS) [47]. PSNR quantifies the pixel color error between the rendered and original frames. DSSIM evaluates the perceived dissimilarity of the rendered image, while LPIPS measures the higher-level perceptual similarity using an AlexNet backbone [19]. Given the inconsistency in DSSIM implementation noted across different methods [1, 13], we follow [23] to distinguish DSSIM results into two categories: DSSIM $_1$ and DSSIM $_2$. DSSIM $_1$ is calculated with a data range set to 1.0, based on the structural similarity function from the *scikit-image* library, whereas DSSIM uses a data range of 2.0. For rendering speed, we measure the performance in frames per second (FPS).

Baselines. As we introduce MEGA, a novel method for compressing 4DGS [46], our primary comparison focuses on the baseline 4DGS method. Additionally, we benchmark MEGA against a range of NeRF-based baselines, including DyNeRF [22], HyperReel [1], Neural Volume

Table 2. Quantitative comparisons with various competitive baselines on the Neural 3D Video Dataset. “Storage” refers to the total model size for 300 frames. ¹: Only report the result on the *Flame Salmon* scene. ²: Exclude the *Coffee Martini* scene. ³: These methods train each model with a 50-frame video sequence to prevent memory overflow, requiring six models to complete the overall evaluation.

Method	PSNR \uparrow	DSSIM $_1\downarrow$	DSSIM $_2\downarrow$	LPIPS \downarrow	FPS \uparrow	Storage \downarrow
Neural Volume ¹ [25]	22.80	-	0.0620	0.2950	-	-
LLFF ¹ [30]	23.24	-	0.0200	0.2350	-	-
DyNeRF ¹ [22]	29.58	-	0.0200	0.0830	0.015	28.00MB
HexPlane ^{2,3} [3]	31.71	-	0.0140	0.0750	0.56	200.00MB
StreamRF [21]	28.26	-	-	-	10.90	5310.00MB
NeRFPlayer ³ [39]	30.69	0.0340	-	0.1110	0.05	5130.00MB
HyperReel [1]	31.10	0.0360	-	0.0960	2.00	360.00MB
K-Planes [13]	31.63	-	0.0180	-	0.30	311.00MB
MixVoxels-L [41]	31.34	-	0.0170	0.0960	37.70	500.00MB
MixVoxels-X [41]	31.73	-	0.0150	0.0640	4.60	500.00MB
Dynamic 3DGS [29]	30.46	0.0350	0.0190	0.0990	460.00	2772.00MB
C-D3DGS [16]	30.46	-	-	0.1500	118.00	338.00MB
Deformable 3DGS [44]	30.98	0.0331	0.0191	0.0594	29.62	32.64MB
E-D3DGS [2]	31.20	0.0259	0.0151	0.0304	69.70	40.20MB
STG ³ [23]	32.04	0.0261	0.0145	0.0440	273.47	175.35MB
4DGS [46]	31.57	0.0290	0.0164	0.0573	96.69	3128.00MB
Ours	31.49	0.0290	0.0165	0.0568	77.42	25.05MB

[25], LLFF [30], HexPlane [3], StreamRF [21], NeRF-Player [39], MixVoxels [41], and K-Planes [13]. Other recent competitive Gaussian-based methods are also considered in our comparisons, including Dynamic 3DGS [29], C-D3DGS [16], Deformable 3DGS [44], E-D3DGS [2], and STG [23]. The numerical results of Deformable 3DGS, E-D3DGS, STG, and 4DGS are produced by running their released codes on a single NVIDIA A800 GPU, while results for other baselines are from their original papers.

Implementation Details. We train our MEGA model over 30k iterations and stop densification at the midpoint. We use the Adam optimizer with a batch size of one, adopting the hyperparameter settings from the original 4DGS [46] framework, including loss weight, learning rate, and threshold parameters. When rendering the view at time t , we filter out those Gaussians with $\sigma(t) \leq 0.05$. To ensure stable training of our deformation predictor, we introduce weight regularization and set it at $1e^{-6}$. The learning rate of the deformation predictor undergoes exponential decay, starting from $8e^{-4}$ and reducing to $1.6e^{-6}$. For the AC color predictor, we start with an initial learning rate of 0.01, incorporating a 100-step warm-up phase. Subsequently, its learning rate is decreased by a factor of three at the 5k, 15k, and 25k steps. Regarding the hyper-parameters in the loss function, we set λ and κ as 0.2 and 0.0005, respectively, to balance the contributions of different components.

4.2. Experimental Results

Table 1 details a quantitative evaluation of our MEGA method on the Technicolor dataset. Notably, our method surpasses the main baseline 4DGS [46], with PSNR,

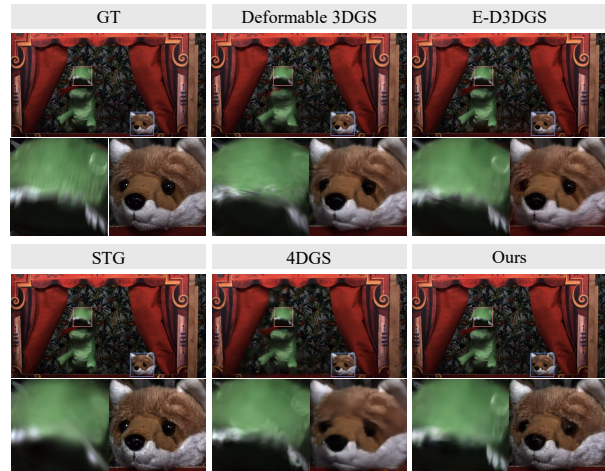


Figure 5. Subjective comparison of various methods on *Theater* scene from the Technicolor Dataset.

DSSIM $_1$, DSSIM $_2$, and LPIPS improvements by 1.2dB, 0.01, 0.006, and 0.018, respectively. Meanwhile, it significantly reduces storage requirements, achieving a 190 \times compactness and improving rendering speed by 50%. When compared with the NeRF-based method HyperReel [1], MEGA achieves a substantial improvement in representation, with an increase of about 0.87dB in PSNR and a 20 \times faster rendering speed, while halving the storage overhead. Moreover, our MEGA records a 0.22dB gain in visual fidelity over the state-of-the-art Gaussian-based method STG [23], and reduces storage overhead by 40%. Fig. 5 offers qualitative comparisons for the *Theater* scene, demonstrating that our results contain more vivid details and provide artifact-less rendering.

Table 3. Ablation study of the proposed components. N denotes the number of Gaussians. The last row represents our final solution.

(a) Technicolor Dataset								
Variants	<i>Birthday</i>				<i>Fabien</i>			
	PSNR \uparrow	DSSIM $_1\downarrow$	$N\downarrow$	Params \downarrow	PSNR \uparrow	DSSIM $_1\downarrow$	$N\downarrow$	Params \downarrow
4DGS [46]	31.00	0.0383	13.00M	2093.56M	33.57	0.0582	5.43M	874.14M
w/ grid [20]	30.49	0.0410	16.33M	293.07M	32.99	0.0620	4.61M	93.77M
w/ DAC	31.60	0.0355	15.43M	308.65M	34.21	0.0587	4.57M	91.48M
w/ DAC+Deformation	31.35	0.0368	15.75M	315.36M	33.02	0.0604	11.56M	231.53M
w/ DAC+ \mathcal{L}_{opa}	31.46	0.0370	9.15M	183.23M	33.96	0.0603	2.32M	46.40M
w/ DAC+Deformation+ \mathcal{L}_{opa}	32.02	0.0309	0.91M	18.48M	34.89	0.0597	0.31M	6.43M

(b) Neural 3D Video Dataset								
Variants	<i>Flame Steak</i>				<i>Sear Steak</i>			
	PSNR \uparrow	DSSIM $_1\downarrow$	$N\downarrow$	Params \downarrow	PSNR \uparrow	DSSIM $_1\downarrow$	$N\downarrow$	Params \downarrow
4DGS [46]	33.19	0.0204	5.17M	831.88M	33.44	0.0204	3.52M	567.30M
w/ grid [20]	31.07	0.0279	4.82M	97.35M	31.313	0.0281	3.25M	70.76M
w/ DAC	33.34	0.0210	5.31M	106.33M	33.67	0.0206	3.61M	72.18M
w/ DAC+Deformation	33.47	0.0209	6.34M	127.16M	33.46	0.0208	4.17M	83.78M
w/ DAC+ \mathcal{L}_{opa}	33.45	0.0208	2.76M	55.22M	33.58	0.0215	1.99M	39.74M
w/ DAC+Deformation+ \mathcal{L}_{opa}	32.27	0.0242	0.87M	17.79M	33.67	0.0200	0.56M	11.50M

Besides, we report the quantitative comparisons on the Neu3DV dataset in Table 2. Relative to 4DGS, our method achieves up to a $125\times$ compression ratio while preserving similar visual quality and rendering speed. It is observed that compared to the SOTA NeRF-based baseline MixVoxels [41], our method achieves a $20\times$ storage reduction and a $16\times$ inference speed improvement, maintaining comparable rendering quality. Furthermore, our approach exhibits higher rendering quality and smaller storage overhead compared to most Gaussian-based methods.

4.3. Ablation Study

To validate the effectiveness of various components within our proposed method, we conduct ablation experiments on scenes from the Technicolor dataset (*Birthday*, *Fabien*) and the Neu3DV dataset (*Flame Steak*, *Sear Steak*). Detailed results are presented in Table 3.

Compact DC-AC Color Representation. Building on the original 4DGS, we substitute the 4D SH coefficients with a grid-based neural field representation [20], and our proposed DAC representation, respectively. While the grid-based approach, referred to as “w/ grid,” achieves a reduction of approximately $10\times$ in parameters, it leads to a significant performance degradation compared to 4DGS. This performance loss may be attributed to the grid’s inability to retain sufficient detail, thereby discarding critical information. To address this issue, we use a DC component to preserve essential color information inherently present in the scene, and an AC predictor to encode the temporal-viewpoint variations in color. This method allows us to achieve a comparable reduction in storage as the grid-based approach while maintaining high-quality rendering consistent with 4DGS.

Entropy-constrained Gaussian Deformation. This part of our ablation study evaluates the impact of Gaussian deformation and opacity-based entropy loss \mathcal{L}_{opa} . Starting

from the configuration “w/ DAC”, we observe that implementing a deformation predictor alone (referred to as “w/ DAC+Deformation”) leads to an increased number of Gaussians. Conversely, employing \mathcal{L}_{opa} without the deformation predictor (referred to as “w/ DAC+ \mathcal{L}_{opa} ”) limits the action range of each Gaussian, inhibiting their efficacy. However, when combining our deformation predictor with \mathcal{L}_{opa} , this strategy significantly reduces the number of Gaussians needed while maintaining rendering quality comparable to that of 4DGS.

5. Conclusion

In this paper, we develop a novel, memory-efficient framework tailored for 4D Gaussian Splatting. By decomposing the color attribute into a per-Gaussian direct current component and a shared, lightweight alternating current color predictor, our approach significantly reduces the per-Gaussian parameters without compromising performance. Furthermore, to reduce redundancy among the 4D Gaussians, we introduce entropy-constrained Gaussian deformation. This technique expands the action range of each Gaussian to enhance the effective utilization rate, thereby enabling the model to render high-quality scenes with as few Gaussians as possible. Extensive experimental results underscore the efficacy of our approach, demonstrating more than a hundredfold reduction in storage requirements while maintaining high-quality reconstruction and real-time rendering speeds in comparison to the original 4D Gaussian Splatting. These advancements establish a new benchmark in the field, combining high performance, compactness, and real-time rendering capabilities.

Acknowledgments. This work was supported by the Hong Kong Research Grants Council under the Areas of Excellence scheme grant AoE/E-601/22-R and NSFC/RGC Collaborative Research Scheme grant CRS_HKUST603/22.

References

- [1] Benjamin Attal, Jia-Bin Huang, Christian Richardt, Michael Zollhoefer, Johannes Kopf, Matthew O’Toole, and Changil Kim. HyperReel: High-fidelity 6-DoF video with ray-conditioned sampling. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3, 6, 7, 13, 14
- [2] Jeongmin Bae, Seoha Kim, Youngsik Yun, Hahyun Lee, Gun Bang, and Youngjung Uh. Per-gaussian embedding-based deformation for deformable 3d gaussian splatting. In *European Conference on Computer Vision*, 2024. 3, 6, 7, 13, 14, 15
- [3] Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 130–141, 2023. 1, 3, 7, 14
- [4] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision*, pages 333–350. Springer, 2022. 3
- [5] Yiwen Chen, Zilong Chen, Chi Zhang, Feng Wang, Xiaofeng Yang, Yikai Wang, Zhongang Cai, Lei Yang, Huaping Liu, and Guosheng Lin. Gaussianeditor: Swift and controllable 3d editing with gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21476–21485, 2024. 3
- [6] Yihang Chen, Qianyi Wu, Jianfei Cai, Mehrtash Harandi, and Weiyao Lin. Hac: Hash-grid assisted context for 3d gaussian splatting compression. In *European Conference on Computer Vision*, 2024. 3
- [7] Yuedong Chen, Haoifei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. In *European Conference on Computer Vision*, pages 370–386. Springer, 2024. 3
- [8] François Darmon and et al. Robust gaussian splatting. *arXiv*, 2024. [available online](#). 15
- [9] Devikalyan Das, Christopher Wewer, Raza Yunus, Eddy Ilg, and Jan Eric Lenssen. Neural parametric gaussians for monocular non-rigid object reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10715–10725, 2024. 3
- [10] Yuanxing Duan, Fangyin Wei, Qiyu Dai, Yuhang He, Wenzheng Chen, and Baoquan Chen. 4d-rotor gaussian splatting: Towards efficient novel view synthesis for dynamic scenes. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 1, 2, 3
- [11] Zhiwen Fan, Kevin Wang, Kairun Wen, Zehao Zhu, Dejie Xu, and Zhangyang Wang. Lightgaussian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps. In *European Conference on Computer Vision*, 2024. 3
- [12] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. 3
- [13] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12479–12488, 2023. 3, 6, 7, 14
- [14] Zhiyang Guo, Wengang Zhou, Li Li, Min Wang, and Houqiang Li. Motion-aware 3d gaussian splatting for efficient dynamic scene reconstruction. In *European Conference on Computer Vision*, 2024. 3
- [15] Alex Hanson, Allen Tu, Geng Lin, Vasu Singla, Matthias Zwicker, and Tom Goldstein. Speedy-splat: Fast 3d gaussian splatting with sparse pixels and sparse primitives. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 21537–21546, 2025. 3
- [16] Kai Katsumata, Duc Minh Vo, and Hideki Nakayama. A compact dynamic 3d gaussian representation for real-time dynamic view synthesis. In *European Conference on Computer Vision*, 2024. 3, 7, 14
- [17] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42 (4):139–1, 2023. 1, 3
- [18] Agelos Kratimenos, Jiahui Lei, and Kostas Daniilidis. Dynmf: Neural motion factorization for real-time dynamic view synthesis with 3d gaussian splatting. In *European Conference on Computer Vision*, 2024. 3
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2012. 6
- [20] Joo Chan Lee, Daniel Rho, Xiangyu Sun, Jong Hwan Ko, and Eunbyung Park. Compact 3d gaussian representation for radiance field. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21719–21728, 2024. 3, 5, 8
- [21] Lingzhi Li, Zhen Shen, zhongshu wang, Li Shen, and Ping Tan. Streaming radiance fields for 3d video synthesis. In *Advances in Neural Information Processing Systems*, 2022. 3, 7, 14
- [22] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5521–5531, 2022. 1, 3, 6, 7, 13, 14
- [23] Zhan Li, Zhang Chen, Zhong Li, and Yi Xu. Spacetime gaussian feature splatting for real-time dynamic view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8508–8520, 2024. 3, 6, 7, 13, 14, 15
- [24] Zhening Liu, Yingdong Hu, Xinjie Zhang, Jiawei Shao, Zehong Lin, and Jun Zhang. Dynamics-aware gaussian splatting streaming towards fast on-the-fly training for 4d reconstruction. *arXiv preprint arXiv:2411.14847*, 2024. 3
- [25] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images.

- ACM Transactions on Graphics*, 38(4):65:1–65:14, 2019. 3, 7, 14
- [26] Jiahao Lu, Yifan Zhang, Qihong Shen, Xinchao Wang, and Shuicheng YAN. Poison-splat: Computation cost attack on 3d gaussian splatting. In *International Conference on Learning Representations*, 2025. 3
- [27] Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20654–20664, 2024. 3
- [28] Zhicheng Lu, Xiang Guo, Le Hui, Tianrui Chen, Min Yang, Xiao Tang, Feng Zhu, and Yuchao Dai. 3d geometry-aware deformable gaussian splatting for dynamic view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8900–8910, 2024. 3
- [29] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *International Conference on 3D Vision*, pages 800–809. IEEE, 2024. 7, 14
- [30] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics*, 2019. 3, 7, 14
- [31] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 1, 3, 5
- [32] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 41(4):1–15, 2022. 3
- [33] KL Navaneet, Kossar Pourahmadi Meibodi, Soroush Abbasi Koohpayegani, and Hamed Pirsiavash. Compgs: Smaller and faster gaussian splatting with vector quantization. In *European Conference on Computer Vision*, 2024. 3
- [34] Simon Niedermayr, Josef Stumpegger, and Rüdiger Westermann. Compressed 3d gaussian splatting for accelerated novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10349–10358, 2024. 3
- [35] Long Peng, Anran Wu, Wenbo Li, Peizhe Xia, Xueyuan Dai, Xinjie Zhang, Xin Di, Haoze Sun, Renjing Pei, Yang Wang, et al. Pixel to gaussian: Ultra-fast continuous super-resolution with 2d gaussian modeling. *arXiv preprint arXiv:2503.06617*, 2025. 3
- [36] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. 1, 3
- [37] Neus Sabater, Guillaume Boisson, Benoit Vandame, Paul Kerbiriou, Frederic Babon, Matthieu Hog, Remy Gendrot, Tristan Langlois, Olivier Bureller, Arno Schubert, et al. Dataset and pipeline for multi-view light-field video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 30–40, 2017. 6
- [38] Qihong Shen, Xingyi Yang, and Xinchao Wang. Flashsplat: 2d to 3d gaussian splatting segmentation solved optimally. In *European Conference on Computer Vision*, pages 456–472. Springer, 2024. 3
- [39] Liangchen Song, Anpei Chen, Zhong Li, Zhang Chen, Lele Chen, Junsong Yuan, Yi Xu, and Andreas Geiger. Nerf-player: A streamable dynamic scene representation with decomposed neural radiance fields. *IEEE Transactions on Visualization and Computer Graphics*, 29(5):2732–2742, 2023. 3, 7, 14
- [40] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022. 3
- [41] Feng Wang, Sinan Tan, Xinghang Li, Zeyue Tian, Yafei Song, and Huaping Liu. Mixed neural voxels for fast multi-view video synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19706–19716, 2023. 3, 7, 8, 14
- [42] Henan Wang, Hanxin Zhu, Tianyu He, Runsen Feng, Jiajun Deng, Jiang Bian, and Zhibo Chen. End-to-end rate-distortion optimized 3d gaussian representation. In *European Conference on Computer Vision*, 2024. 3
- [43] Yuxin Wang, Qianyi Wu, Guofeng Zhang, and Dan Xu. Learning 3d geometry and feature consistent gaussian splatting for object removal. In *European Conference on Computer Vision*, pages 1–17. Springer, 2024. 3
- [44] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20310–20320, 2024. 3, 6, 7, 13, 14
- [45] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20331–20341, 2024. 3
- [46] Zeyu Yang, Hongye Yang, Zijie Pan, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. In *International Conference on Learning Representations*, 2024. 1, 2, 3, 4, 5, 6, 7, 8, 13, 14, 15
- [47] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. 6
- [48] Xinjie Zhang, Xingtong Ge, Tongda Xu, Dailan He, Yan Wang, Hongwei Qin, Guo Lu, Jing Geng, and Jun Zhang. Gaussianimage: 1000 fps image representation and compression by 2d gaussian splatting. In *European Conference on Computer Vision*, pages 327–345. Springer, 2024. 3
- [49] Shijie Zhou, Haoran Chang, Sicheng Jiang, Zhiwen Fan, Zehao Zhu, Dejia Xu, Pradyumna Chari, Suyu You, Zhangyang

Wang, and Achuta Kadambi. Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21676–21685, 2024. [3](#)

- [50] Lingting Zhu, Guying Lin, Jinnan Chen, Xinjie Zhang, Zhenchao Jin, Zhao Wang, and Lequan Yu. Large images are gaussians: High-quality large image representation with levels of 2d gaussian splatting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10977–10985, 2025. [3](#)