

D-FCGS: Feedforward Compression of Dynamic Gaussian Splatting for Free-Viewpoint Videos

Wenkang Zhang^{1,2}, Yan Zhao¹, Qiang Wang³, Zhixin Xu¹, Li Song¹, Zhengxue Cheng^{1*}

¹Shanghai Jiao Tong University

²Zhiyuan College, Shanghai Jiao Tong University

³Visionstar Information Technology (Shanghai) Co., Ltd

{conquer.wkzhang, zhaoyanzy, zhixin.xu, song_li, zxcheng}@sjtu.edu.cn, wq@sightp.com

Abstract

Free-Viewpoint Video (FVV) enables immersive 3D experiences, but efficient compression of dynamic 3D representation remains a major challenge. Existing dynamic 3D Gaussian Splatting methods couple reconstruction with optimization-dependent compression and customized motion formats, limiting generalization and standardization. To address this, we propose D-FCGS, a novel **Feedforward Compression** framework for **Dynamic Gaussian Splatting**. Key innovations include: (1) a standardized Group-of-Frames (GoF) structure with I-P coding, leveraging sparse control points to extract inter-frame motion tensors; (2) a dual prior-aware entropy model that fuses hyperprior and spatial-temporal priors for accurate rate estimation; (3) a control-point-guided motion compensation mechanism and refinement network to enhance view-consistent fidelity. Trained on Gaussian frames derived from multi-view videos, D-FCGS generalizes across diverse scenes in a zero-shot fashion. Experiments show that it matches the rate-distortion performance of optimization-based methods, achieving over 17× compression compared to the baseline while preserving visual quality across viewpoints. This work advances feedforward compression of dynamic 3DGS, facilitating scalable FVV transmission and storage for immersive applications.

Code — <https://github.com/Mr-Zwkid/D-FCGS>

1 Introduction

Our world is inherently dynamic, with 3D scenes evolving over time and observable from arbitrary viewpoints. Capturing and representing such 4D dynamics has long been a fundamental challenge in computer vision and graphics. Free-viewpoint video (FVV), which enables immersive 6-DoF experiences, has emerged as a promising solution with applications in virtual reality, telepresence, and remote education. However, realizing practical FVV systems demands efficient solutions for reconstruction, compression, transmission, and rendering. In this work, we focus on the critical problem of compression for dynamic 3D scenes.

Recent advances of 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) have revolutionized 3D scene representation, of-

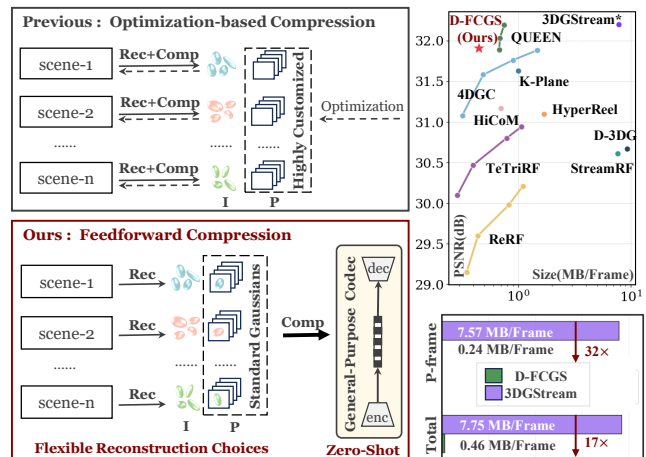


Figure 1: **Left:** Illustration of differences between previous optimization-based methods and our D-FCGS. **Right:** R-D curve and storage size comparison on N3V dataset.

fering unparalleled rendering quality and real-time performance. Extending to 4D, dynamic forms of 3DGS (Yang et al. 2024; Wu et al. 2024a; Li et al. 2024; Yang et al. 2023; Sun et al. 2024) gradually garner attention. Analogous to the temporal expansion of images facilitated by videos, frame-by-frame 3D Gaussians can naturally serve as a temporal expansion of 3DGS, forming the basis for FVV. Based on this per-frame idea, 3DGSStream (Sun et al. 2024) pioneers dynamic scene reconstruction via on-the-fly training, while subsequent works improve compression efficiency through techniques like rate-aware training (Hu et al. 2025), vector-quantized residuals (Girish et al. 2024), and static-dynamic decomposition (Wu et al. 2025).

Despite these advances, existing methods remain constrained by their coupled optimization of reconstruction and compression. By representing frame-wise Gaussian motions through specialized formats (e.g., neural networks), these approaches require per-scene optimization and meticulous hyperparameter tuning. This optimization-dependent paradigm creates two critical barriers to practical deployment: (1) it severely limits generalization to unseen scenes; (2) it hinders the development of standardized compression

*Corresponding Author

schemes that could enable widespread adoption of FVV.

To address these challenges, we propose **D-FCGS**, a novel feedforward compression framework for dynamic Gaussian Splatting. Our key insight is that temporal coherence in Gaussian point clouds can be efficiently modeled via the Group-of-Frames (GoF) representation, where inter-frame motions can be compressed in a feedforward and scene-agnostic manner. Specifically, we adopt the standard GS format as input and leverage sparse control points to efficiently extract motion tensors, optimal for both compression efficiency and computational performance. These motion tensors are then processed through our feedforward motion compression pipeline, which incorporates a tailored dual prior-aware entropy model to enhance probability estimation accuracy. Following decompression, the sparse motions are propagated across the entire Gaussian frame under the guidance of the control points. Finally, lightweight color refinement is applied to improve view-consistent fidelity.

Our D-FCGS is trained end-to-end on frame-wise GS sequences constructed from both real-world and synthetic multi-view video datasets. Once trained, it serves as a **general-purpose inter-frame compression codec**, requiring no scene-specific optimization or access to multi-view images during zero-shot inference. Extensive experiments show that D-FCGS exhibits strong generalization across diverse dynamic scenes and achieves state-of-the-art rate-distortion performance, surpassing $17\times$ compression over 3DGStream while maintaining comparable fidelity.

Our contributions can be summarized as follows:

- We present D-FCGS, a novel feedforward compression framework for dynamic 3DGS that enables zero-shot inter-frame coding of Gaussian sequences.
- We propose a sparse motion representation with control points for I-P coding, coupled with a dual prior-aware entropy model for efficient compression, and develop a decoder with motion compensation and refinement for enhanced fidelity.
- Experiments across various scenes show the effectiveness and robustness of D-FCGS, achieving over $17\times$ compression over 3DGStream while preserving fidelity, outperforming most optimization-based methods.

2 Related Work

2.1 3D Gaussian Splatting Compression

High storage demands of 3DGS have motivated compression efforts. Optimization-based compression methods include value-based and structure-based ones. The former involve the use of pruning (Girish, Gupta, and Shrivastava 2024; Ali et al. 2024), masking (Lee et al. 2024b; Wang et al. 2024a), vector quantization (Niedermayr, Stumpfegger, and Westermann 2024; Navaneet et al. 2023) and distillation (Fan et al. 2024) to reduce insignificant Gaussians or redundancy of Gaussian parameters. The latter utilize structural modeling such as anchors (Lu et al. 2024), tri-planes (Lee et al. 2025) and 2D grids (Morgenstern et al. 2024) to address the sparsity and unorganized nature of Gaussian Splatting. When combined with learned entropy models (Ballé,

Laparra, and Simoncelli 2016; Cheng et al. 2020), they achieve excellent rate-distortion performance (Chen et al. 2024b, 2025; Zhan et al. 2025).

However, these optimization-based methods require per-scene fine-tuning, limiting their generalization. Recent works have introduced feedforward pipelines (Chen et al. 2024a; Huang et al. 2025; Yang et al. 2025b) that decouples reconstruction and compression. These pre-trained codecs can directly compress arbitrary Gaussian clouds without multi-view supervision, offering greater practicality. We aim to extend this paradigm to dynamic GS compression.

2.2 Dynamic GS and its Compression

Dynamic GS methods can be categorized by their motion representation approaches. Implicit/explicit motion fields (Yang et al. 2024; Wu et al. 2024a; Li et al. 2024; Lin et al. 2024) and 4D Gaussian formulations (Duan et al. 2024; Lee et al. 2024a) often face challenges in real-time streaming, variable resolutions, or long durations. In contrast, per-frame approaches (Luiten et al. 2024; Sun et al. 2024; Wang et al. 2024b; Yan et al. 2025) demonstrate superior practicality through incremental frame updates, making them ideal for free-viewpoint video streaming. Our work focuses on this line of *Per-Frame Gaussian* methods.

For compression of *Per-Frame Gaussian*, HiCoM (Gao et al. 2024) employs hierarchical grid-wise motion representation and QUEEN (Girish et al. 2024) leverages a latent-decoder for quantization of attribute residuals. 4DGC (Hu et al. 2025) instills rate-aware training into frame reconstruction, achieving decent compression. While effective, they fundamentally couple reconstruction with compression, requiring scene-specific tuning. Our work breaks this constraint by introducing a novel optimization-free compression framework for dynamic GS, enabling zero-shot compression to arbitrary 4D scenes.

3 Preliminary

Our feedforward compression method takes frame-wise standard 3DGS as input, denoted as Gaussian frames, and each frame is pre-optimized from multi-view images at the respective timestamp. Within a given Gaussian frame, each Gaussian primitive is characterized by: (1) **geometry parameters** including position $\mu \in \mathbb{R}^3$ and covariance matrix $\Sigma \in \mathbb{R}^{3 \times 3}$; (2) **appearance parameters** including opacity $o \in \mathbb{R}^1$ and SH-based color $c_{SH} \in \mathbb{R}^{48}$. The covariance matrix can be further represented as $\Sigma = RSS^T R^T$, where $R \in \mathbb{R}^{3 \times 3}$ is the rotation matrix parameterized by the quaternion $q \in \mathbb{R}^4$, and the scale matrix $S \in \mathbb{R}^{3 \times 3}$ is a diagonal matrix with elements $s \in \mathbb{R}^3$. The geometry of a Gaussian primitive can be formulated as:

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

where $\mathbf{x} \in \mathbb{R}^3$ is any random 3D location within the scene.

Given a viewpoint, 3D Gaussians are projected into a 2D plane, and the color of a pixel $C \in \mathbb{R}^3$ is derived by alpha-blending of overlapping 2D Gaussians:

$$C = \sum_i c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (2)$$

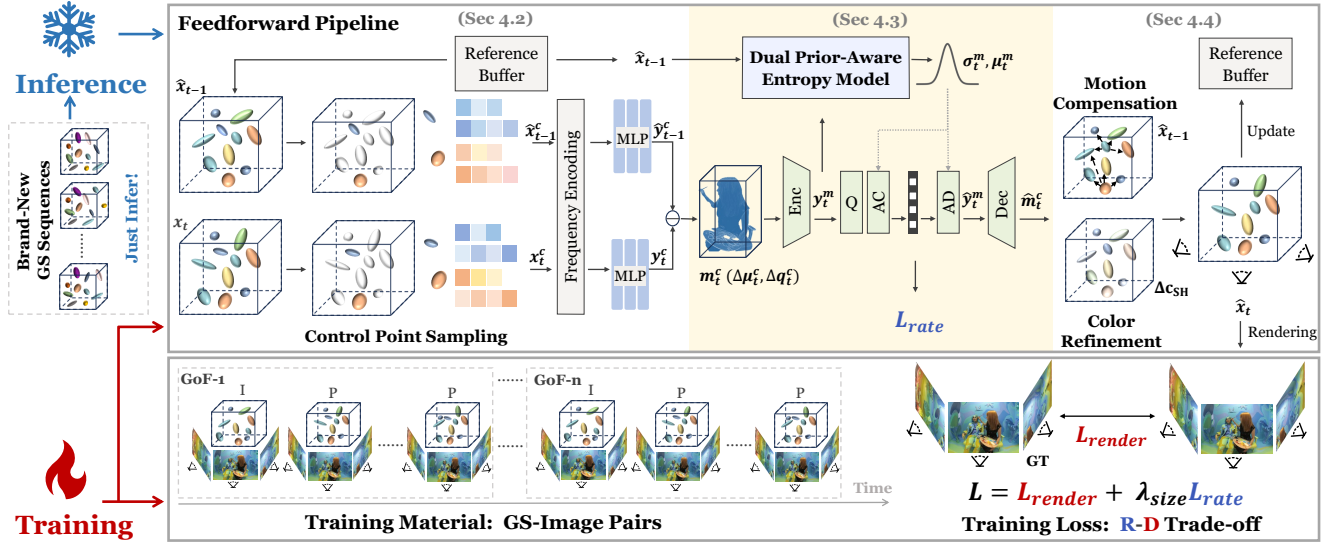


Figure 2: **Overview of D-FCGS framework.** Our feedforward pipeline processes sequential Gaussian frames in GoF through three stages: (1) sparse motion extraction (Section 4.2), (2) feedforward motion compression (Section 4.3), and (3) motion compensation and refinement (Section 4.4). Once trained with rate-distortion loss, D-FCGS can infer on brand-new GS sequences in a zero-shot manner.

where $c_i \in \mathbb{R}^3$ is the view-dependent color calculated from SH-based color c_{SH} , and $\alpha_i \in \mathbb{R}^1$ is the blending weight derived from opacity o .

With the differentiable rasterizer, training can be supervised by 2D images from varied views in an end-to-end way. The rendering loss of vanilla 3DGS is:

$$L_{\text{render}} = \lambda L_{\text{D-SSIM}} + (1 - \lambda) L_1. \quad (3)$$

4 Method

4.1 Overview

As shown in Fig. 2, our D-FCGS framework comprises three key components: (1) sparse motion extraction (Section 4.2), (2) feedforward motion compression (Section 4.3) and (3) motion compensation and refinement (Section 4.4). Training and inference procedures will be mentioned in Section 4.5.

Given two adjacent Gaussian frames x_t and \hat{x}_{t-1} , we first sample sparse control points from dense Gaussians and extract motion features. These motion tensors are compressed using our dual prior-aware entropy model for efficient rate estimation. During decoding, we reconstruct \hat{x}_t via control-point-guided motion compensation and color refinement, while storing it into the buffer for future reference.

4.2 Sparse Motion Extraction via Control Points

Building on observations that most Gaussians exhibit local motion coherence, we propose an efficient sparse representation. Inspired by previous optimization-based 4D reconstruction methods (He et al. 2024; Yan et al. 2025; Huang et al. 2024), we employ Farthest Point Sampling (FPS) to select $N^c = \frac{N}{M}$ control points from the N Gaussian primitives, and derive corresponding geometry parameters, for-

ulated as:

$$\mu^c = FPS(\{\mu_i\}_{i \in N}, \frac{N}{M}), \quad (4)$$

$$x^c = \text{Index}(x, \mu^c), \quad (5)$$

where M denotes the downscale factor, and x^c ($x^c = \{\mu^c, q^c\}$) represent geometry parameters of control points. This process efficiently reduces storage demand and processing costs, while preserving motion characteristics.

After control point sampling, we derive the native sparse geometry attributes of the current frame and reference Gaussian frame, denoted as x_t^c and \hat{x}_{t-1}^c . We encode these attributes using frequency encoding (Mildenhall et al. 2021) and MLP projection:

$$y_t^c = MLP(\text{FreqEnc}\{x_t^c\}), \quad (6)$$

$$\hat{y}_{t-1}^c = MLP(\text{FreqEnc}\{\hat{x}_{t-1}^c\}). \quad (7)$$

Motion tensors are then derived in the feature domain:

$$m_t^c = \text{Converter}(y_t^c - \hat{y}_{t-1}^c), \quad (8)$$

where m_t^c denotes the motion tensors at time t , which keeps the same dimension as x_t^c and \hat{x}_{t-1}^c . *Converter* is realized by MLP at the feature level.

4.3 Feedforward Motion Compression

End-to-End Motion Compression. The obtained sparse motion tensors m_t^c are then fed into our end-to-end compression module. This process begins with data encoding, followed by differentiable quantization simulated by additive uniform noise (Ballé, Laparra, and Simoncelli 2016) to




Figure 3: **Illustration of the dual prior-aware entropy model.** Spatial-temporal context priors extracted via multi-scale hashgrids and hyperpriors generated through the factorized model are fused by a lightweight fusion network for bitrate estimation.

enable gradient backpropagation :

$$\begin{aligned} \hat{\mathbf{y}}_t^m &= Q(\mathbf{y}_t^m) = \mathbf{y}_t^m + \mathcal{U}\left(-\frac{q'}{2}, \frac{q'}{2}\right), \text{ for training} \\ &= \text{Round}\left(\frac{\mathbf{y}_t^m}{q'}\right) \cdot q', \text{ for testing} \end{aligned} \quad (9)$$

where $q' \in \mathbb{R}^1$ is the quantization step size, and \mathbf{y}_t^m , $\hat{\mathbf{y}}_t^m$ denote the encoded latent motion before and after quantization, respectively. Next to that, arithmetic coding (AC) converts the quantized data into a compact bitstream for efficient transmission and storage. On the decoder side, the bitstream is decompressed back into motion tensors using arithmetic decoding (AD) for frame reconstruction.

Dual Prior-Aware Entropy Model. According to Shannon’s theory (Shannon 1948), the cross-entropy between the estimated and true latent distributions provides a tight lower bound on the achievable bitrate:

$$R(\hat{\mathbf{y}}_t^m) \geq \mathbb{E}_{\hat{\mathbf{y}}_t^m \sim q_{\hat{\mathbf{y}}_t^m}} [-\log_2 p_{\hat{\mathbf{y}}_t^m}(\hat{\mathbf{y}}_t^m)], \quad (10)$$

where $p_{\hat{\mathbf{y}}_t^m}$ and $q_{\hat{\mathbf{y}}_t^m}$ are respectively estimated and true probability mass functions (PMFs) of the quantized latent codes $\hat{\mathbf{y}}_t^m$. Since arithmetic coding can achieve a bitrate close to this bound, our goal is to devise an entropy model that accurately estimates $p_{\hat{\mathbf{y}}_t^m}$.

Fig. 3 shows the proposed dual prior-aware entropy model that combines hyperprior and spatial-temporal context prior for precise distribution estimation. Following (Ballé et al. 2018), we use a factorized model to learn the hyperprior and estimate its PMF $p(\hat{\mathbf{z}}_t^m)$, which is common in deep image and video compression. However, for GS-based FVV, the latent codes also exhibit strong 3D-spatial and temporal correlations. Assuming adjacent Gaussian frames share similar features, we strive to extract spatial-temporal priors from the reference frame $\hat{\mathbf{x}}_{t-1}$ through a multi-resolution hash grid encoding scheme. Specifically, at each grid intersection, we store learnable feature vectors that capture position-specific information. For each Gaussian positioned at $\boldsymbol{\mu}_t$, we retrieve multi-scale features by performing tri-linear interpolation across different levels G^l of the voxel grid, where l indicates the resolution level. The interpolated features from all levels are then concatenated and processed by a lightweight

MLP to generate comprehensive positional contexts:

$$\text{context} = \text{MLP}\left(\bigcup_{l=1}^L \text{Interp}(\boldsymbol{\mu}_t, G^l)\right), \quad (11)$$

where $\text{Interp}(\cdot)$ denotes the grid interpolation operation. These positional contexts are subsequently combined with the transformed appearance parameters as the final spatial-temporal priors. The prior fusion network is then to integrate these spatial-temporal priors with the hyperpriors, estimating the mean $\boldsymbol{\mu}_t^m$ and scale σ_t^m of the latent code distribution (assumed as normal distribution). The probability mass for each quantized latent value $\hat{\mathbf{y}}_{t,i}^m$ is computed by integrating the PMF over the quantization bin:

$$p(\hat{\mathbf{y}}_{t,i}^m) = \int_{\hat{\mathbf{y}}_{t,i}^m - \frac{q'}{2}}^{\hat{\mathbf{y}}_{t,i}^m + \frac{q'}{2}} \mathcal{N}(y | \boldsymbol{\mu}_{t,i}^m, \sigma_{t,i}^m) dy, \quad (12)$$

where i corresponds to the index of a certain control point. The rate loss combines contributions from both the motion and hyper latent distributions:

$$L_{\text{rate}} = \frac{1}{N^c} \sum_{i=1}^{N^c} (-\log_2(p(\hat{\mathbf{y}}_{t,i}^m)) - \log_2(p(\hat{\mathbf{z}}_{t,i}^m))), \quad (13)$$

where $N^c = \frac{N}{M}$ is the total number of control points.

4.4 Motion Compensation and Refinement

Control Point Guided Motion Compensation. The decoded control point motions $\hat{\mathbf{m}}_t^c$ are propagated to the entire Gaussian set in a distance-aware compensation way. For the i^{th} control point, we first identify its K -nearest Gaussians $\mathcal{K}(i)$ via KNN search. The motion vectors are then distributed to neighboring Gaussians using an exponentially decaying weight function based on spatial distance. The closer the distance is, the more influence we consider the control point can exert on the motion of this neighboring point. The motion that i^{th} control point assigns to its j^{th} neighbor can be written as:

$$m_{i,j} = \frac{e^{-d_{i,j}} m_i}{\sum_{k \in \mathcal{K}(i)} e^{-d_{i,k}}}, \quad (14)$$

where $d_{i,j}$ represents the Euclidean distance between the i^{th} control point and its j^{th} neighbor. The aggregated motion vectors are finally applied to adjust the geometry parameters of the reference frame via addition.

This approach provides two key benefits: (1) enhanced motion estimation accuracy through localized correlation modeling that suppresses error propagation, and (2) inherent parallelizability due to the independent processing of control points, ensuring computational efficiency. More discussions on frame matching during motion compensation can be seen in the supplementary.

Color Refinement. For image and video compression, post-compression refinement plays a critical role in mitigating visual artifacts such as color banding and blurring. In our GS-based framework, we specifically target refinement

at SH coefficients c_{SH} while preserving sensitive geometry and opacity parameters (Chen et al. 2024a; Girish et al. 2024; Papantonakis et al. 2024). The spatial-temporal priors from the entropy model module are repurposed to predict color residuals Δc_{SH} , which are dynamically added to c_{SH} during decoding without additional storage. This **on-the-fly** refinement is fully differentiable, allowing gradient backpropagation through the entropy model for joint rate-distortion optimization.

4.5 Training and Inference Pipeline of D-FCGS

Training Process and Loss. Following established practices (Wang et al. 2023; Zheng et al. 2024a,b), we adopt a Group-of-Frames (GoF) paradigm with the structure:

$$\underbrace{(I - P - \dots - P)}_{\text{GoF}_1(L)} \quad \underbrace{(I - P - \dots - P)}_{\text{GoF}_2(L)} \quad \dots \quad \underbrace{(I - P - \dots - P)}_{\text{GoF}_k(L)}, \quad (15)$$

where $\text{GoF}_k(L)$ denotes the k^{th} group containing one intra-coded (I) frame followed by $L - 1$ predictively-coded (P) frames. The model is trained end-to-end with a composite loss function:

$$L_{\text{total}} = L_{\text{render}} + \lambda_{\text{size}} L_{\text{rate}}, \quad (16)$$

where L_{render} replicates the original 3D Gaussian Splatting objective, and λ_{size} controls the rate-distortion trade-off.

Encoding and Decoding Process. The encoder extracts sparse motion from control points and compresses motion tensors \hat{y}_t^m and hyperpriors \hat{z}_t^m via arithmetic coding. During decoding, the system first reconstructs \hat{z}_t^m , then combines it with spatial-temporal priors to estimate distribution parameters (σ_t^m, μ_t^m) for decoding \hat{y}_t^m . These decoded motion features enable subsequent motion compensation.

5 Experiments

We rigorously evaluate our D-FCGS model by addressing three critical research questions:

- **Generalization and Effectiveness:** Can D-FCGS generalize across new scenes from widely used datasets while achieving competitive rate-distortion performance compared to optimization-based methods? (Section 5.2)
- **Robustness and Stability:** How does D-FCGS perform under diverse and high-dynamic scenes, and is the system stable under varying hyperparameters? (Section 5.3)
- **Module Efficacy:** Do individual modules (e.g. control points) contribute meaningfully? (Section 5.4)

5.1 Experimental Setup

Datasets and Implementation Details. We derive sequential Gaussian frames from six multi-view video datasets: (1) *N3V* (Li et al. 2022b) (2) *MeetRoom* (Li et al. 2022a) (3) *WideRange4D* (Yang et al. 2025a) (4) *Google Immersive* (Broxton et al. 2020) (5) *Self-Cap* (Xu et al. 2024) (6) *VRU* (Wu et al. 2025). For training, we use 3 scenes from *MeetRoom* and 28 scenes from *WideRange4D*. For evaluation, we reserve the "discussion" scene from *MeetRoom* for in-domain testing and six scenes from *N3V* for out-of-domain benchmarking. Additional robustness tests are

Method	PSNR (dB) ↑	SSIM ↑	Size (MB) ↓	Render (FPS) ↑	Feedforward Compression
K-Planes	31.63	0.920	1.0	0.15	✗
HyperReel	31.10	0.931	1.7	16.7	✗
NeRFPlayer	30.69	0.931	18.4	0.05	✗
StreamRF	30.61	0.930	7.6	8.3	✗
ReRF	29.71	0.918	0.77	2.0	✗
TeTriRF	30.65	0.931	0.76	2.7	✗
D-3DG	30.67	0.931	9.2	460	✗
3DGStream*	32.20	0.953	7.75	215	✗
HiCoM	31.17	-	0.70	<u>274</u>	✗
QUEEN	<u>32.19</u>	0.946	0.75	248	✗
4DGC	31.58	0.943	<u>0.49</u>	168	✗
D-FCGS (ours)	31.91	<u>0.952</u>	0.46	215	✓

Table 1: Quantitative results on N3V (Li et al. 2022b) dataset, averaged over 300 frames across six scenes. * denotes results reproduced by our implementation. **Bold** and underlined values indicate the best and second-best performance, respectively. Detailed per-scene results are reported in supplementary material.

Method	PSNR (dB) ↑	SSIM ↑	Size (MB) ↓	Render (FPS) ↑
StreamRF	26.71	0.913	8.23	10
ReRF	26.43	0.911	0.63	2.9
TeTriRF	27.37	0.917	0.61	3.8
3DGStream*	31.74	0.957	7.66	288
HiCoM	29.61	-	<u>0.40</u>	<u>284</u>
4DGC	28.08	0.922	0.42	213
D-FCGS (ours)	<u>30.97</u>	<u>0.950</u>	0.38	288

Table 2: Quantitative results on MeetRoom (Li et al. 2022a) dataset, averaged over 300 frames.

performed on diverse and high-dynamic scenes (see Section 5.3). All experiments are conducted on NVIDIA RTX 4090 GPU. Key hyperparameters include: (1) the downscale factor $M = 70$, (2) $K = 30$ for KNN, (3) quantization step $q' = 1$, (4) GoF size $L = 10$ during inference, and (4) $\lambda_{\text{size}} = 1e-3$. Further details of datasets and implementation are provided in the supplementary material.

Evaluation Metrics. We evaluate D-FCGS using three kinds of metrics: (1) PSNR and SSIM (Wang et al. 2004) for reconstruction quality, (2) compressed size (MB/frame) for rate efficiency, and (3) rendering speed (FPS) plus encoding/decoding time (sec) for computational performance.

5.2 Benchmark Comparison

Benchmark Methods. As the first feedforward inter-frame codec for Gaussian point clouds (to our knowledge), D-FCGS faces unique evaluation challenges due to the absence of directly comparable optimization-free approaches. Our analysis therefore encompasses two categories of benchmark methods in optimization-based dynamic scene reconstruction: (1) NeRF-based techniques including K-Planes (Fridovich-Keil et al. 2023), HyperReel (Attal et al. 2023), NeRFPlayer (Song et al. 2023),

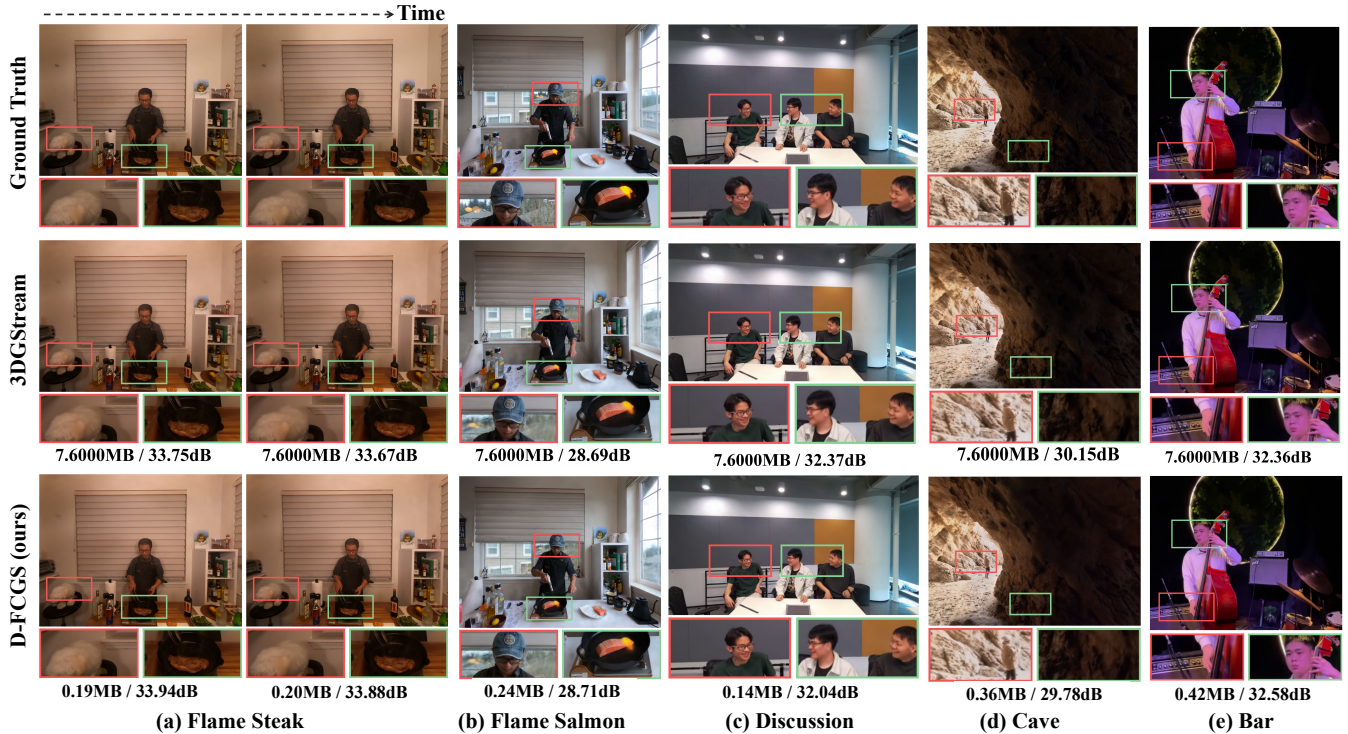


Figure 4: **Qualitative comparison of rendering results.** Visual comparisons between (1) Ground Truth, (2) 3DGStream, and (3) our D-FCGS. D-FCGS significantly reduces storage size while maintaining comparable high fidelity to 3DGStream.

Method	Google Immersive			Self-Cap			VRU		
	PSNR(dB)↑	SSIM↑	Size(MB)↓	PSNR(dB)↑	SSIM↑	Size(MB)↓	PSNR(dB)↑	SSIM↑	Size(MB)↓
3DGStream*	25.84	0.872	7.6	26.43	0.858	7.6	26.18	0.892	7.6
D-FCGS(ours)	25.58	0.879	0.34	26.00	0.854	0.90	25.13	0.876	0.43

Table 3: Robustness test on diverse scenes from extensive datasets. Averaged PSNR, SSIM and P-frame size are reported. Per-scene results are provided in the appendix.

K	30	60	120	30	30	30	30
M	70	70	70	140	280	70	70
L	300	300	300	300	300	32	4
Size↓	8.9	9.0	9.0	4.0	2.0	9.0	9.2
PSNR↑	33.596	33.596	33.596	33.599	33.599	33.604	33.598

Table 4: Robustness test on hyperparameter settings (K for KNN, downscale factor M , and GoF size L). Averaged P-frame size ($\times 0.1$ MB) and PSNR (dB) are reported on "cut.beef" scene.

StreamRF (Li et al. 2022a), ReRF (Wang et al. 2023), and TeTriRF (Wu et al. 2024b); and (2) *Per-Frame Gaussian* methods including D-3DG (Luiten et al. 2024), 3DGStream (Sun et al. 2024), HiCoM (Gao et al. 2024), QUEEN (Girish et al. 2024), and 4DGC (Hu et al. 2025). While this comparison inherently favors optimization-based approaches that benefit from scene-specific tuning, it provides critical insights into D-FCGS's performance relative to cur-

Method	Encoding(sec)	Decoding(sec)	Total(sec)
proposed	0.61	0.72	1.33
w/o control points	1.33	2.88	4.21

Table 5: Average encoding and decoding time for P-frames.

rent paradigms in dynamic scene compression.

Quantitative Results. Our compression results demonstrate significant improvements over existing methods. As shown in Section 5.1 and Section 5.1, D-FCGS achieves remarkable average sizes of **0.46MB** ($N3V$) and **0.38MB** ($MeetRoom$) per frame, namely a **17 \times** reduction compared to 3DGStream's 7.75MB and 7.66MB. Note that our codec is designed for inter-frame compression, and our P-frame compression ratio exceeds 32 \times in most cases (Fig. 1). While I-frames could be further compressed using existing static methods (e.g., FCGS (Chen et al. 2024a)), we maintain their original size for fair comparison. Additionally, the entire en-

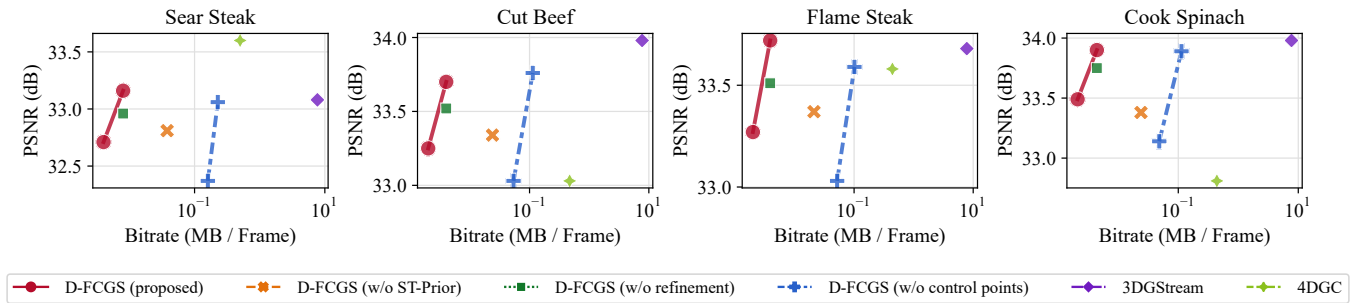


Figure 5: Rate-Distortion comparison of the proposed method and ablations.

coding/decoding pipeline operates efficiently, with both processes completing in under 1 second (Section 5.1).

Qualitative Results. We visualize the qualitative comparisons with 3DGStream (Sun et al. 2024) in Fig. 4, presenting results on “flame steak”, “flame salmon” (*N3V* dataset) and “discussion” (*MeetRoom* dataset). From scene details, we can tell that D-FCGS achieves comparable fidelity to 3DGStream, even better in some cases.

5.3 Robustness Test

Robustness to Diverse Scenes. We evaluate D-FCGS on three additional datasets: *Google Immersive*, *Self-Cap*, and *VRU*, covering scenarios such as indoor painting, cave exploration, and basketball games. As summarized in Section 5.1, for Gaussian sequences reconstructed by 3DGStream, our approach achieves unprecedented P-frame compression while maintaining near-identical SSIM and small PSNR drop (< 0.5 dB for *Google Immersive* and *Self-Cap*). While PSNR gaps exist for *VRU* (full of blurry motions), our method offers a more practical and efficient pipeline for real-world applications, sacrificing acceptable visual fidelity for great improvements in bitrates. Fig. 4 highlights visual results on the outdoor “cave” (*Google Immersive*) and high-dynamic “bar” (*Self-Cap*).

Robustness to Hyperparameter Settings. Here, we apply varied parameter configurations (KNN neighborhood size K , downscale factor M , and GoF size L) to the “cut beef” scene (*N3V*). As presented in Section 5.1, D-FCGS maintains stable rendering quality across most parameter combinations, showing strong parametric robustness. Notably, increasing the downscale factor M effectively reduces control point counts, leading to a linear decrease in P-frame compression size, which aligns with expectations.

5.4 Ablation Study

In this section, we evaluate key components of D-FCGS through systematic ablation studies.

Effect of Control Points. The sparse motion representation via control points forms the foundation of our efficient compression pipeline. Removing control points and predicting motions for all Gaussians increases storage substantially (Fig. 5) and slows encoding/decoding by 3.2 \times (Section 5.1).

This validates our sparse motion representation’s efficiency in rate saving and computational cost.

Effect of Dual Prior-Aware Entropy Model. Our entropy model employs a novel dual-prior architecture. While hyperpriors effectively capture global dependencies (Ballé et al. 2018), our key innovation lies in the hash-grid-based spatial-temporal prior that models local correlations. Removing the spatial-temporal prior branch leads to noticeable R-D performance degradation (Fig. 5), confirming its importance to optimal compression.

Effect of Color Refinement. Our online color refinement module improves rendering quality (PSNR) by 0.1 \sim 0.5 dB (Fig. 5), while requiring no additional storage and negligible decoding overhead. Per-scene results on *N3V* dataset are shown in the supplementary.

6 Conclusion

In this paper, we propose Feedforward Compression of Dynamic Gaussian Splatting (D-FCGS), a novel feedforward framework for zero-shot dynamic Gaussian sequence compression. Our contributions are threefold. First, we adopt the I-P coding profile for standard GS compression and introduce sparse inter-frame motion extraction via control points. Second, we present an end-to-end motion compression framework with a dual prior-aware entropy model, fully leveraging hyperpriors and spatial-temporal context to improve rate estimation. Third, control-point-guided motion compensation is combined with a color refinement network to guarantee high-fidelity and view-consistent reconstruction. Experimental results show that *D-FCGS* achieves superior compression efficiency (over 17 \times compression) on two benchmark datasets (*MeetRoom* and *N3V*) and remarkable robustness across diverse scenes from extensive datasets, significantly enhancing transmission and storage efficiency for free-viewpoint video applications.

Acknowledgments

This work was partly supported by the NSFC62431015, Science and Technology Commission of Shanghai Municipality No.24511106200, the Fundamental Research Funds for the Central Universities, Shanghai Key Laboratory of Digital Media Processing and Transmission under Grant 22DZ2229005, 111 project BP0719010, Okawa Research Grant and Explore-X Research Fund.

References

- Ali, M. S.; Qamar, M.; Bae, S.-H.; and Tartaglione, E. 2024. Trimming the fat: Efficient compression of 3d gaussian splats through pruning. *arXiv preprint arXiv:2406.18214*.
- 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.
- Ballé, J.; Laparra, V.; and Simoncelli, E. P. 2016. End-to-end optimized image compression. *arXiv preprint arXiv:1611.01704*.
- Ballé, J.; Minnen, D.; Singh, S.; Hwang, S. J.; and Johnston, N. 2018. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*.
- Broxton, M.; Flynn, J.; Overbeck, R.; Erickson, D.; Hedman, P.; Duvall, M.; Dourgarian, J.; Busch, J.; Whalen, M.; and Debevec, P. 2020. Immersive light field video with a layered mesh representation. *ACM Transactions on Graphics (TOG)*, 39(4): 86–1.
- Chen, Y.; Wu, Q.; Li, M.; Lin, W.; Harandi, M.; and Cai, J. 2024a. Fast feedforward 3d gaussian splatting compression. *arXiv preprint arXiv:2410.08017*.
- Chen, Y.; Wu, Q.; Lin, W.; Harandi, M.; and Cai, J. 2024b. Hac: Hash-grid assisted context for 3d gaussian splatting compression. In *European Conference on Computer Vision*, 422–438. Springer.
- Chen, Y.; Wu, Q.; Lin, W.; Harandi, M.; and Cai, J. 2025. HAC++: Towards 100X Compression of 3D Gaussian Splatting. *arXiv preprint arXiv:2501.12255*.
- Cheng, Z.; Sun, H.; Takeuchi, M.; and Katto, J. 2020. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7939–7948.
- Duan, Y.; Wei, F.; Dai, Q.; He, Y.; Chen, W.; and Chen, B. 2024. 4d-rotor gaussian splatting: towards efficient novel view synthesis for dynamic scenes. In *ACM SIGGRAPH 2024 Conference Papers*, 1–11.
- Fan, Z.; Wang, K.; Wen, K.; Zhu, Z.; Xu, D.; Wang, Z.; et al. 2024. Lightgaussian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps. *Advances in neural information processing systems*, 37: 140138–140158.
- 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.
- Gao, Q.; Meng, J.; Wen, C.; Chen, J.; and Zhang, J. 2024. Hicom: Hierarchical coherent motion for dynamic streamable scenes with 3d gaussian splatting. *Advances in Neural Information Processing Systems*, 37: 80609–80633.
- Girish, S.; Gupta, K.; and Shrivastava, A. 2024. Eagles: Efficient accelerated 3d gaussians with lightweight encodings. In *European Conference on Computer Vision*, 54–71. Springer.
- Girish, S.; Li, T.; Mazumdar, A.; Shrivastava, A.; De Mello, S.; et al. 2024. QUEEN: QUantized Efficient ENcoding of Dynamic Gaussians for Streaming Free-viewpoint Videos. *Advances in Neural Information Processing Systems*, 37: 43435–43467.
- He, B.; Chen, Y.; Lu, G.; Wang, Q.; Gu, Q.; Xie, R.; Song, L.; and Zhang, W. 2024. S4d: Streaming 4d real-world reconstruction with gaussians and 3d control points. *arXiv preprint arXiv:2408.13036*.
- Hu, Q.; Zheng, Z.; Zhong, H.; Fu, S.; Song, L.; Xiaoyun-Zhang; Zhai, G.; and Wang, Y. 2025. 4DGC: Rate-Aware 4D Gaussian Compression for Efficient Streamable Free-Viewpoint Video. *arXiv:2503.18421*.
- Huang, H.; Huang, W.; Yang, Q.; Xu, Y.; and Li, Z. 2025. A hierarchical compression technique for 3d gaussian splatting compression. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Huang, Y.-H.; Sun, Y.-T.; Yang, Z.; Lyu, X.; Cao, Y.-P.; and Qi, X. 2024. Sc-gs: Sparse-controlled gaussian splatting for editable dynamic scenes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4220–4230.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4): 139–1.
- Lee, J.; Won, C.; Jung, H.; Bae, I.; and Jeon, H.-G. 2024a. Fully explicit dynamic gaussian splatting. *Advances in Neural Information Processing Systems*, 37: 5384–5409.
- Lee, J. C.; Rho, D.; Sun, X.; Ko, J. H.; and Park, E. 2024b. Compact 3d gaussian representation for radiance field. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21719–21728.
- Lee, S.; Shu, F.; Sanchez, Y.; Schierl, T.; and Hellge, C. 2025. Compression of 3D Gaussian Splatting with Optimized Feature Planes and Standard Video Codecs. *arXiv preprint arXiv:2501.03399*.
- Li, L.; Shen, Z.; Wang, Z.; Shen, L.; and Tan, P. 2022a. Streaming radiance fields for 3d video synthesis. *Advances in Neural Information Processing Systems*, 35: 13485–13498.
- Li, T.; Slavcheva, M.; Zollhoefer, M.; Green, S.; Lassner, C.; Kim, C.; Schmidt, T.; Lovegrove, S.; Goesele, M.; Newcombe, R.; et al. 2022b. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5521–5531.
- Li, Z.; Chen, Z.; Li, Z.; and Xu, Y. 2024. Spacetime gaussian feature splatting for real-time dynamic view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8508–8520.
- Lin, Y.; Dai, Z.; Zhu, S.; and Yao, Y. 2024. Gaussian-flow: 4d reconstruction with dynamic 3d gaussian particle. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21136–21145.

- Lu, T.; Yu, M.; Xu, L.; Xiangli, Y.; Wang, L.; Lin, D.; and Dai, B. 2024. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20654–20664.
- Luiten, J.; Kopanas, G.; Leibe, B.; and Ramanan, D. 2024. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *2024 International Conference on 3D Vision (3DV)*, 800–809. IEEE.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Morgenstern, W.; Barthel, F.; Hilsmann, A.; and Eisert, P. 2024. Compact 3d scene representation via self-organizing gaussian grids. In *European Conference on Computer Vision*, 18–34. Springer.
- Navaneet, K.; Meibodi, K. P.; Koohpayegani, S. A.; and Pirsivash, H. 2023. Compact3d: Compressing gaussian splat radiance field models with vector quantization. *arXiv preprint arXiv:2311.18159*, 4.
- Niedermayr, S.; Stumpfegger, J.; and Westermann, R. 2024. Compressed 3d gaussian splatting for accelerated novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10349–10358.
- Papantonakis, P.; Kopanas, G.; Kerbl, B.; Lanvin, A.; and Drettakis, G. 2024. Reducing the memory footprint of 3d gaussian splatting. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 7(1): 1–17.
- Shannon, C. E. 1948. A mathematical theory of communication. *The Bell system technical journal*, 27(3): 379–423.
- 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.
- Sun, J.; Jiao, H.; Li, G.; Zhang, Z.; Zhao, L.; and Xing, W. 2024. 3dstream: On-the-fly training of 3d gaussians for efficient streaming of photo-realistic free-viewpoint videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20675–20685.
- Wang, H.; Zhu, H.; He, T.; Feng, R.; Deng, J.; Bian, J.; and Chen, Z. 2024a. End-to-end rate-distortion optimized 3d gaussian representation. In *European Conference on Computer Vision*, 76–92. Springer.
- Wang, L.; Hu, Q.; He, Q.; Wang, Z.; Yu, J.; Tuytelaars, T.; Xu, L.; and Wu, M. 2023. Neural residual radiance fields for streamably free-viewpoint videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 76–87.
- Wang, P.; Zhang, Z.; Wang, L.; Yao, K.; Xie, S.; Yu, J.; Wu, M.; and Xu, L. 2024b. V³: Viewing Volumetric Videos on Mobiles via Streamable 2D Dynamic Gaussians. *ACM Transactions on Graphics (TOG)*, 43(6): 1–13.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Wu, G.; Yi, T.; Fang, J.; Xie, L.; Zhang, X.; Wei, W.; Liu, W.; Tian, Q.; and Wang, X. 2024a. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 20310–20320.
- Wu, J.; Peng, R.; Wang, Z.; Xiao, L.; Tang, L.; Yan, J.; Xiong, K.; and Wang, R. 2025. Swift4D: Adaptive divide-and-conquer Gaussian Splatting for compact and efficient reconstruction of dynamic scene. *arXiv preprint arXiv:2503.12307*.
- Wu, M.; Wang, Z.; Kouros, G.; and Tuytelaars, T. 2024b. Tetrirf: Temporal tri-plane radiance fields for efficient free-viewpoint video. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6487–6496.
- Xu, Z.; Xu, Y.; Yu, Z.; Peng, S.; Sun, J.; Bao, H.; and Zhou, X. 2024. Representing Long Volumetric Video with Temporal Gaussian Hierarchy. *ACM Transactions on Graphics*, 43(6).
- Yan, J.; Peng, R.; Wang, Z.; Tang, L.; Yang, J.; Liang, J.; Wu, J.; and Wang, R. 2025. Instant gaussian stream: Fast and generalizable streaming of dynamic scene reconstruction via gaussian splatting. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 16520–16531.
- Yang, L.; Zhu, K.; Tian, J.; Zeng, B.; Lin, M.; Pei, H.; Zhang, W.; and Yan, S. 2025a. WideRange4D: Enabling High-Quality 4D Reconstruction with Wide-Range Movements and Scenes. *arXiv preprint arXiv:2503.13435*.
- Yang, Q.; Yang, L.; Van Der Auwera, G.; and Li, Z. 2025b. HybridGS: High-Efficiency Gaussian Splatting Data Compression using Dual-Channel Sparse Representation and Point Cloud Encoder. *arXiv preprint arXiv:2505.01938*.
- Yang, Z.; Gao, X.; Zhou, W.; Jiao, S.; Zhang, Y.; and Jin, X. 2024. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 20331–20341.
- Yang, Z.; Yang, H.; Pan, Z.; and Zhang, L. 2023. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. *arXiv preprint arXiv:2310.10642*.
- Zhan, Y.-T.; Ho, C.-Y.; Yang, H.; Chen, Y.-H.; Chiang, J. C.; Liu, Y.-L.; and Peng, W.-H. 2025. CAT-3DGS: A Context-Adaptive Triplane Approach to Rate-Distortion-Optimized 3DGS Compression. *arXiv preprint arXiv:2503.00357*.
- Zheng, Z.; Zhong, H.; Hu, Q.; Zhang, X.; Song, L.; Zhang, Y.; and Wang, Y. 2024a. HPC: Hierarchical Progressive Coding Framework for Volumetric Video. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 7937–7946.
- Zheng, Z.; Zhong, H.; Hu, Q.; Zhang, X.; Song, L.; Zhang, Y.; and Wang, Y. 2024b. JointRF: End-to-End Joint Optimization for Dynamic Neural Radiance Field Representation and Compression. In *2024 IEEE International Conference on Image Processing (ICIP)*, 3292–3298. IEEE.