

Dynamic Gaussian Scene Reconstruction from Unsynchronized Videos

Zhixin Xu^{1*}, Hengyu Zhou^{1*}, Yuan Liu², Wenhan Xue¹, Hao Pan¹, Wenping Wang³, Bin Wang^{1,4†}

¹Tsinghua University

²Hong Kong University of Science and Technology

³Texas A&M University

⁴Beijing National Research Center for Information Science and Technology

{xuzx23, zhouhy22, xuewh22}@mails.tsinghua.edu.cn, yuanly@ust.hk, wenping@tamu.edu,

{haopan, wangbins}@tsinghua.edu.cn

Abstract

Multi-view video reconstruction plays a vital role in computer vision, enabling applications in film production, virtual reality, and motion analysis. While recent advances such as 4D Gaussian Splatting (4DGS) have demonstrated impressive capabilities in dynamic scene reconstruction, they typically rely on the assumption that input video streams are temporally synchronized. However, in real-world scenarios, this assumption often fails due to factors like camera trigger delays or independent recording setups, leading to temporal misalignment across views and reduced reconstruction quality. To address this challenge, a novel temporal alignment strategy is proposed for high-quality 4DGS reconstruction from unsynchronized multi-view videos. Our method features a coarse-to-fine alignment module that estimates and compensates for each camera’s time shift. The method first determines a coarse, frame-level offset and then refines it to achieve sub-frame accuracy. This strategy can be integrated as a readily integrable module into existing 4DGS frameworks, enhancing their robustness when handling asynchronous data. Experiments show that our approach effectively processes temporally misaligned videos and significantly enhances baseline methods.

Introduction

The reconstruction of dynamic scenes in four dimensions (4D) is a frontier topic in computer graphics and computer vision, aiming to capture and reconstruct time-evolving 3D scenes from multi-view videos. This technology is fundamental to enabling applications such as free-viewpoint video, immersive virtual and augmented reality experiences, digital twins, and visual effects production. In recent years, 4D Gaussian Splatting (4DGS) has emerged as a leading paradigm in dynamic scene reconstruction, distinguished by its exceptional rendering quality and unprecedented real-time rendering speeds. By representing the scene with explicit Gaussian primitives, it achieves highly efficient and high-fidelity modeling of complex dynamic details.

However, a critical yet demanding assumption underpinning most state-of-the-art 4DGS methods (Huang et al.

*These authors contributed equally.

†Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

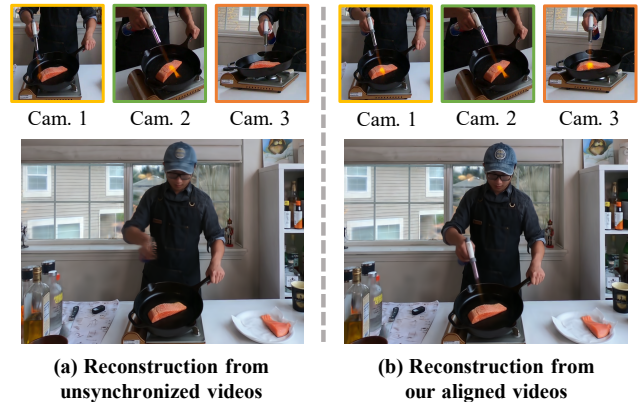


Figure 1: We introduce a novel readily integrable module that significantly enhances existing 4D dynamic scene reconstruction methods (e.g., 4D Gaussians and SC-GS). Compared to the baseline result (left), our method (right) demonstrates superior quality in capturing intricate details and complex motion.

2024b; Yang et al. 2024c; Wu et al. 2024) is the requirement for strict temporal synchronization across all capture cameras. This implies that for any given timestamp t , the camera shutters are triggered simultaneously. While this condition can be met in professional studios or laboratory settings using expensive hardware like Genlock signal generators, this assumption is often violated in more general, real-world capture scenarios. For instance, when using a set of independent consumer-grade cameras (e.g., smartphones, GoPros) or a distributed camera system controlled over a wireless network, the lack of a centralized clock signal, coupled with network latency and manual start-and-stop operations, almost inevitably introduces temporal misalignment ranging from milliseconds to even seconds between video streams. This temporal asynchronicity poses a critical challenge to 4D reconstruction. When a model attempts to fuse views captured at physically distinct moments to reconstruct the scene at a single logical timestamp, severe visual artifacts arise. For objects in fast motion, even minor time discrepancies can cause significant positional differences between

views, leading to ghosting artifacts and motion blur in the final reconstruction. This inconsistent observational data misguides the 4DGS optimization process, causing it to incorrectly attribute temporal errors to flaws in spatial geometry or appearance. This ultimately leads to a sharp degradation in reconstruction quality, or even complete failure. Consequently, the problem of temporal misalignment has become a key bottleneck preventing the widespread adoption of high-fidelity 4D reconstruction technology beyond controlled environments.

To address this challenge, we propose a novel and general temporal alignment strategy for high-quality 4DGS reconstruction from unsynchronized multi-view videos. Instead of passively accepting data with temporal errors, we explicitly incorporate the unknown temporal offset of each camera into the optimization objective. Specifically, we design a coarse-to-fine optimization framework to estimate and compensate for each camera’s time shift precisely. The framework first operates on a coarse-grained temporal scale to rapidly identify an approximate offset range, effectively preventing the optimization from converging to a poor local minimum. Subsequently, it performs fine-tuning at a more granular scale to achieve sub-frame alignment accuracy. Our strategy can be seamlessly integrated as a readily integrable module into various existing 4DGS frameworks, significantly enhancing their robustness and reconstruction quality when handling asynchronous data.

The main contributions of this paper are summarized as follows:

- We systematically identify and analyze the problem of temporal misalignment, which is prevalent in real-world capture, and reveal its severely detrimental effects on existing 4DGS methods.
- We propose, for the first time, a coarse-to-fine optimization framework that jointly solves for the unknown temporal offsets of each camera stream concurrently with the 4DGS reconstruction process.
- Extensive experiments demonstrate that our method effectively handles temporally asynchronous videos, significantly outperforms baseline approaches on multi-view datasets, and successfully reconstructs high-quality, artifact-free dynamic scenes. This work substantially expands the application boundaries of dynamic Gaussian reconstruction, enabling its use with lower-cost and more flexible capture setups.

Related Work

Novel View Synthesis

Novel view synthesis (NVS) for dynamic scenes has advanced considerably in recent years, fueled by breakthroughs in Neural Radiance Fields (NeRF) and 3D Gaussian Splatting (3DGS). The foundational work, NeRF (Mildenhall et al. 2021), introduced a method to represent 3D scenes as continuous volumetric functions parameterized by neural networks. Various strategies have been proposed to reduce memory consumption and speed up the reconstruction process in neural rendering. These include

techniques such as hash-based encoding (Müller et al. 2022), mobile-optimized architectures (Chen et al. 2023), scalable training frameworks (Wu et al. 2022), highly parallelized networks (Reiser et al. 2021), and efficient factorized representations (Garbin et al. 2021), all of which contribute to making real-time or large-scale neural rendering more feasible.

The advent of 3DGS marked a shift towards rasterization-based explicit methods. 3DGS (Kerbl et al. 2023) modeled scenes as collections of 3D Gaussians, enabling real-time rendering. Following the introduction of 3DGS, extensive research has focused on improving its storage efficiency, rendering speed, and visual quality. Scaffold-GS (Lu et al. 2024) improves 3DGS by replacing unstructured Gaussians with a structured scaffold of anchors that generate view-dependent Gaussians, effectively reducing redundancy and enhancing view consistency. The COLMAP-free approach (Fu et al. 2024) enables the progressive reconstruction of a 3D Gaussian scene by simultaneously estimating camera poses from video input. GaussianPro (Cheng et al. 2024) introduces a progressive propagation strategy that efficiently densifies the Gaussian cloud from a sparse initialization to capture finer details. To optimize rendering performance and consistency, Octree-GS (Ren et al. 2024) organizes Gaussians into a level-of-detail (LOD) hierarchy using an octree structure, ensuring real-time rates at different viewing distances. Departing from the standard primitive, GES (Hamdi et al. 2024) proposes a generalized exponential function as a more flexible alternative to the 3D Gaussian, improving rendering efficiency and quality. In parallel, extensions of the original method have also been explored, including approaches (Huang et al. 2024a) that leverage 2D Gaussian primitives to achieve smoother surface representations.

Dynamic Gaussian Reconstruction

Recent advancements have extended 3D Gaussian Splatting to dynamic scenes. Deformable 3DGS (Yang et al. 2024a) introduces a deformable formulation that learns scene representations in a canonical space, using a deformation field to capture monocular dynamics. 4DGaussians (Wu et al. 2024) proposes a hybrid explicit representation that combines 3D Gaussians with 4D neural voxels, employing decomposed voxel encoding and a lightweight MLP to predict Gaussian deformations over time. Building on these foundations, subsequent works (Duan et al. 2024; Yan et al. 2024; Yang et al. 2024b) have improved dynamic modeling through enhancements in primitive representation, motion optimization, and sampling strategies. SplineGS (Park et al. 2024) further models smooth Gaussian trajectories using cubic Hermite splines with motion-adaptive control point pruning, enabling efficient and expressive deformation under varying motion patterns.

Recent methods for dynamic reconstruction from monocular real-world videos can be categorized into reconstruction-based approaches, which explicitly recover scene geometry and motion over time (LIU et al. 2025; Wang et al. 2024a; Stearns et al. 2024; Lei et al. 2024), and prediction-based approaches, which directly estimate mo-

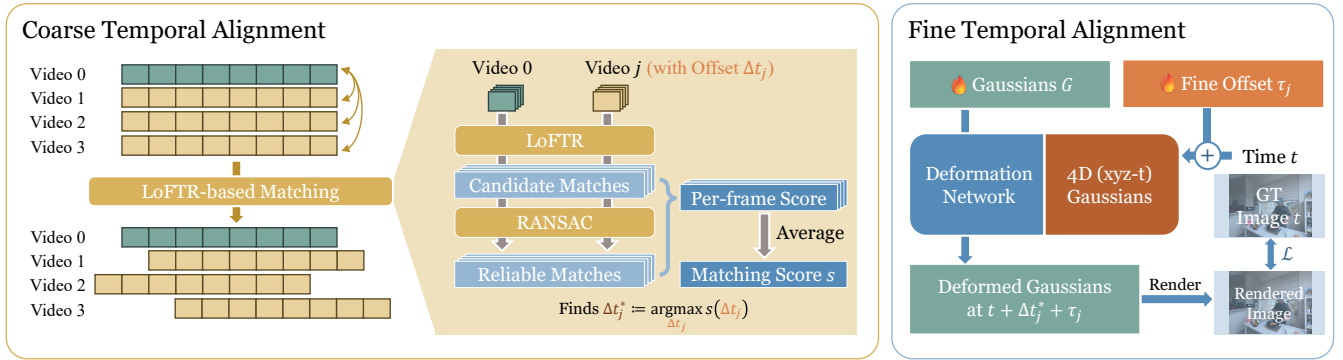


Figure 2: Overview of our two-stage temporal alignment pipeline. **Left:** Coarse Temporal Alignment estimates integer offsets Δt_j^* by matching frames across videos using LoFTR and RANSAC. **Right:** Fine Temporal Alignment refines offsets with a learnable τ_j . The result is supervised by a photometric reconstruction loss.

tion or deformation using priors or flow-based representations (Wang et al. 2024b; Liang et al. 2024). These methods collectively advance the accuracy, coherence, and efficiency of dynamic scene understanding in real-world settings. GaussianFlow (Lin et al. 2024) and MotionGS (Zhu et al. 2024) incorporate optical flow information to introduce additional motion constraints in dynamic motion modeling, significantly improving temporal consistency and geometric accuracy.

Preliminary

3D Gaussian Splatting In the 3D Gaussian Splatting (3DGS) framework, a scene is modeled as a collection of anisotropic 3D Gaussians $\{G_i\}$. Each Gaussian G is parameterized by its center $\mu \in \mathbb{R}^3$, covariance matrix $\Sigma \in \mathbb{R}^{3 \times 3}$, opacity $\alpha \in \mathbb{R}$, and color $c \in \mathbb{R}^3$. After scaling by opacity, its density at any point $\mathbf{x} \in \mathbb{R}^3$ is given by

$$G(\mathbf{x}) = \alpha \exp \left[-\frac{1}{2} (\mathbf{x} - \mu)^\top \Sigma^{-1} (\mathbf{x} - \mu) \right]. \quad (1)$$

Its covariance matrix Σ is parameterized as

$$\Sigma = R S S^\top R^\top \quad (2)$$

where $R \in SO(3)$ is the rotation matrix and $S = \text{diag}(s)$ is a diagonal scale matrix.

The input for 3DGS scene reconstruction generally involves a set of posed images, and the 3DGS model is optimized to minimize the difference between its rendered images and the given images at the given poses.

4D Gaussian Splatting 4DGS extends the static scene of 3DGS with an extra temporal dimension. This is typically achieved by associating each 3D Gaussian primitive with the time dimension, i.e. $G(\mathbf{x})$ becomes $G(\mathbf{x}, t)$, where the parameters of a Gaussian primitive can vary with t (Wu et al. 2024; Yang et al. 2024c).

The input for 4DGS scene reconstruction involves a set of videos rather than static images, where the videos do not have to be synchronized in the time dimension in practice, and the 4DGS model is reconstructed by minimizing the difference between its rendered images from the given videos.

The unsynchronized videos pose additional challenge for 4DGS reconstruction, because for the 4DGS model $\{G_i(\mathbf{x}, t)\}$ at the moment t , it's not clear to which frame $I_{j,t'}$ of the j -th video the GS model should compare. As we can see in Fig. 1, a naive treatment of $t' = t$ leads to severe reconstruction failures, especially for highly dynamic regions. Our goal in this paper is to present an approach such that an optimal temporal alignment $t' = f_j(t)$ is found for each video, so that the total reconstruction error is minimized. In this paper, we assume the temporal alignment function $f_j(t) = t + \delta t_j$ is an offset transformation, which complies with practical scenarios where the videos are unsynchronized mostly by a time shift.

Method

We present a time alignment module in our method, designed to address the challenge of reconstructing dynamic scenes from asynchronous multi-view videos. Rather than proposing an entirely new reconstruction pipeline, our key contribution lies in a modular and pluggable component that can be seamlessly integrated into existing state-of-the-art dynamic Gaussian methods. The module is intended to automatically estimate and compensate for temporal misalignments caused by hardware limitations or capture inconsistencies in multi-camera systems.

In the following sections, we first introduce our coarse-to-fine time alignment module, where the time shift of each camera $\delta t = \Delta t + \tau$ is decomposed into a coarse level shift Δt and a fine level shift τ to assist globally optimal alignment. Then we show how our module can be easily plugged into state-of-the-art dynamic representations implementing 4DGS.

Coarse Temporal Alignment

To address temporal asynchrony among multi-view video streams, we first perform a coarse, frame-level synchronization. Our approach is based on the insight that dense feature matchers like LoFTR (Sun et al. 2021), while designed for static scenes, can be powerfully repurposed for temporal alignment. Specifically, when video frames from differ-




Figure 3: Illustration of our coarse temporal alignment. For each candidate temporal offset Δt between two camera views, we evaluate the number of geometrically consistent feature matches, quantified by the number of RANSAC inliers. For a given view pair, we first extract putative matches using LoFTR (red dots), then apply RANSAC to robustly find inliers (green lines), which serve as our alignment score.

ent views are captured at the exact same moment, a dynamic foreground effectively becomes static from multiple perspectives. For LoFTR, this "instantaneously static" object provides a rich source of stable features, leading to a surge in the number of high-confidence matches, as illustrated by Fig. 3. We leverage this principle by identifying the frame pairings that maximize these feature matches, thereby establishing a robust initial synchronization.

Building on this insight, we cast temporal alignment as a frame-pair selection problem: for each pair of camera views, we search for the time offset Δt that matches the views best, as measured by the number of reliable feature correspondences.




Figure 4: Corresponding frames for matching score calculation under different Δt_j s. We find the offset Δt_j for each video j relative to the reference video that maximizes the matching score.

As a preliminary step, we process all video streams through a pre-trained video segmentation model (Yang et al. 2023) to obtain a binary foreground mask, M , for every frame. This mask, M^t , isolates the pixels corresponding to the dynamic subject in frame I^t . By focusing exclusively on the dynamic subject, this foreground-centric strategy provides a more direct and reliable signal for synchronization.

For a set of multi-view videos as the input, we select one as the reference, from which reference frames $\{I_{\text{ref}}^{t_i}\}$ are sampled. For each other video j , we search for the best Δt_j in the range $[-k, k]$, where k is a hyperparameter to account for the maximum potential misalignment.

For every candidate Δt_j , we calculate the average matching score between each reference frame $I_{\text{ref}}^{t_i}$ and its corresponding frame $I_j^{t_i+\Delta t_j}$. The correspondence is illustrated in Fig. 4. The score is computed in two stages: generating candidate correspondences, and then estimating robust geometric matches.

In the first stage, we use LoFTR to search for candidate feature correspondences between each frame pair $(I_{\text{ref}}^t, I_j^{t+\Delta t_j})$. Critically, we filter those correspondences using the foreground masks, retaining only those matches where both points lie within their respective masks, M_{ref}^t and $M_j^{t+\Delta t_j}$.

In the second stage, we apply the RANSAC (Fischler and Bolles 1981) algorithm to this set of foreground-only correspondences. RANSAC robustly fits a geometric model (in

our case, a fundamental matrix) to find the largest set of self-consistent inliers. This step effectively tests the hypothesis that the foreground object captured in the two views is a rigid projection from a single moment in time.

The number of foreground inliers, $N_{\text{inlier-fg}}(I_{\text{ref}}^{t_i}, I_j^{t_i+\Delta t})$, serves as the alignment score. A high score indicates that the subject’s pose is geometrically consistent between the two views, providing strong evidence of correct temporal alignment. The optimal offset Δt_j^* for video j is the one that maximizes this score:

$$\Delta t_j^* = \arg \max_{\Delta t_j \in [-k, k]} \sum_{t_i} N_{\text{inlier-fg}}(I_{\text{ref}}^{t_i}, I_j^{t_i+\Delta t}). \quad (3)$$

Fine Temporal Refinement

The coarse alignment provides a robust, integer-frame synchronization for all video streams. These offsets are added to the frame indices of each source camera stream j . To achieve sub-frame precision, we introduce a small, learnable temporal refinement parameter, τ_j , for each camera j . This parameter represents a continuous, residual time shift that is learned during training. τ_j , in addition to Δt_j^* , is added to the timestamp before calculating the parameters of Gaussians at time t , resulting in the queried t' being

$$t' = t + \Delta t_j^* + \tau_j. \quad (4)$$

These refinement parameters $\{\tau_j\}$ are learnable and are optimized jointly with the 4DGS model. The gradients are derived from the end-to-end photometric reconstruction loss. This allows the optimization process to dynamically discover and correct any residual, sub-frame temporal discrepancies, which is essential for minimizing motion artifacts and achieving high-fidelity dynamic scene reconstruction.

Integration with 4D Representations

Our coarse alignment module is an individual part that does not require modifying the baseline implementation. The fine temporal alignment needs to be implemented into existing methods.

Neural 4D Representations The proposed module can be integrated with 4D Gaussian Splatting systems that represent scene dynamics through neural networks—a paradigm widely adopted by recent high-performance dynamic 4DGS methods (Wu et al. 2024; Huang et al. 2024b). These systems model dynamics by using a neural network to deform a static set of canonical 3D Gaussians.

The neural deformation network \mathcal{D}_θ , parameterized by θ and typically implemented as a multi-layer perceptron (MLP), models the temporal evolution of each Gaussian. Specifically, it takes as input the positionally encoded canonical location $\gamma(\mu_k)$ and time $\gamma(t)$, and outputs deformations, such as position offsets $\Delta\mu_k$:

$$\mathcal{D}_\theta(\gamma(\mu_k), \gamma(t)) \rightarrow (\Delta\mu_k, \dots) \quad (5)$$

where $\gamma(\cdot)$ denotes a positional encoding function.

Our module interfaces with the deformation network via its temporal input. Instead of feeding the global time t , we add the coarse offset Δ_j^* and a per-camera learnable offset

τ_j to each frame. Since the output of \mathcal{D}_θ is differentiable with respect to its temporal input, gradients with respect to the per-camera offset τ_j can be naturally computed during backpropagation, thereby enabling end-to-end optimization.

Direct 4D Representations An alternative approach to representing dynamic scenes is to extend 3D Gaussians directly into a 4D (xyz-t) space. The RT4DGS method (Yang et al. 2024c) adapts the 3DGS methodology by parameterizing 4D covariance matrices using 4D rotations and 4D scales.

Similar to 3DGS, rendered images from this representation are fully differentiable to their inputs. However, the original RT4DGS implementation does not include gradient computation concerning the input timestamp t . Therefore, we compute the derivative of the reconstruction loss \mathcal{L} with respect to t using a finite difference approximation:

$$\frac{\partial \mathcal{L}}{\partial t} \approx \frac{\mathcal{L}(t+h) - \mathcal{L}(t)}{h}. \quad (6)$$

The selection of an appropriate step size, h , is crucial for numerical stability. In our setup, we find that values for h ranging from one-hundredth to one-tenth of the frame interval provide stable results.

Experiments

Datasets and Metrics

To evaluate our method, we utilize the Dyrnerf dataset (Li et al. 2022), which contains six challenging dynamic scenes, each captured from approximately 20 viewpoints in 9-second video clips. To heighten the challenge and better simulate real-world conditions, we create a more demanding benchmark from this data. The videos are subsampled to 15 FPS, which introduces more significant motion between consecutive frames. We apply a random temporal offset of up to 10 frames to each video. This process yields a multi-view, temporally misaligned dataset that mimics asynchronous capture scenarios. On this benchmark, we assess our method’s performance on the task of novel view synthesis. We report three standard quantitative metrics: PSNR, SSIM, and LPIPS (Zhang et al. 2018).

Implementation Details

Experiments based on 4DGaussians and SC-GS were conducted on an RTX 3090, while those using the RT4DGS method were run on an RTX A6000. Initial point clouds were generated using the COLMAP SfM pipeline.

For fair comparison, the hyperparameters from baseline methods are untouched. We choose h to be one-thirtieth of the frame interval for derivative calculation. Hyperparameters for correspondence estimation follow standard practice. We adopt the official Outdoor configuration of LoFTR, and set the RANSAC inlier threshold to 0.2 px. The temporal search window k depends on estimated asynchrony and frame rate: larger asynchrony requires a wider range, while lower FPS reduces it. In typical cases where the asynchrony is below 0.5 s and the frame rate lies within 15–30 FPS, choosing $k = 10$ –20 works well.

Method	Coffee Martini			Cook Spinach			Cut Roasted Beef		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Sync-Nerf	27.64	0.893	0.147	28.90	0.918	0.132	29.96	0.925	0.125
SC-GS (Huang et al. 2024b)	26.81	0.906	0.113	30.53	0.939	0.104	31.13	0.943	0.095
SC-GS+Ours	25.68	0.898	0.125	31.54	0.947	0.089	31.37	0.949	0.084
4DGaussians (Wu et al. 2024)	26.44	0.905	0.120	31.44	0.946	0.098	31.37	0.941	0.096
4DGaussians+Ours	28.01	0.918	0.108	32.57	0.951	0.089	31.71	0.948	0.093
RT4DGS* (Yang et al. 2024c)	27.92	0.919	0.085	31.15	0.948	0.077	31.43	0.953	0.072
RT4DGS*+Ours	28.35	0.924	0.079	33.15	0.962	0.059	32.94	0.963	0.056

Method	Flame Salmon			Flame Steak			Sear Steak		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Sync-Nerf	27.00	0.890	0.150	30.66	0.942	0.093	30.50	0.938	0.110
SC-GS (Huang et al. 2024b)	26.83	0.912	0.105	30.34	0.947	0.096	30.92	0.950	0.085
SC-GS+Ours	27.10	0.915	0.102	31.46	0.953	0.081	31.20	0.949	0.092
4DGaussians (Wu et al. 2024)	28.01	0.917	0.109	30.68	0.952	0.087	29.67	0.947	0.082
4DGaussians+Ours	29.53	0.923	0.103	32.63	0.955	0.077	32.51	0.959	0.080
RT4DGS* (Yang et al. 2024c)	27.78	0.923	0.080	31.13	0.956	0.068	32.94	0.965	0.058
RT4DGS*+Ours	28.79	0.929	0.070	33.34	0.968	0.050	33.51	0.968	0.055

Table 1: Quantitative comparison of our method against baselines on the Dynerf dataset. Our method consistently enhances RT4DGS, 4DGaussians, and SC-GS frameworks. RT4DGS*: We follow RT4DGS using half the resolution for evaluation.



Figure 5: **Visual comparison of reconstruction results from unsynchronized inputs.** We compare novel view synthesis results from the original 4DGaussians and RT4DGS methods against versions enhanced by our approach (+Ours). Images are from four scenes of the DyNeRF dataset: coffee_martini, cut_roasted_beef, cook_spinach, and flame_salmon.

Comparisons

We integrate our module into three state-of-the-art 4D scene reconstruction pipelines: 4DGaussians (Wu et al. 2024), SC-GS (Huang et al. 2024b), and RT4DGS (Yang et al. 2024c). Direct comparisons with the original baselines demonstrate the effectiveness of our approach. As shown in Tab. 1, our method achieves significant quantitative improvements. Qualitative results in Fig. 5 further illustrate reduced visual artifacts, particularly those caused by unsynchronized reconstruction inputs.

Ablation Study

To analyze the contribution of each component within our proposed temporal refinement module, we conduct a series of ablation studies. These experiments are designed to validate two key aspects of our method: (1) the necessity of our coarse-to-fine strategy for predicting time offsets, and (2) the robustness of our model against varying degrees of random temporal perturbations.

Ablation on Coarse and Fine Temporal Alignment We investigate the effectiveness of our proposed two-stage (Coarse + Fine) strategy for time offset prediction. To this end, we compare our **Full** method against three variants: a baseline without any temporal correction, a **Coarse**-only version, and a **Fine**-only version. Quantitative results are provided in Tab. 2.

As shown in Fig. 6, both the Coarse-only and Fine-only models significantly reduce visual artifacts such as blurring and ghosting compared to the baseline. This demonstrates that each stage independently contributes to mitigating temporal misalignment. However, the Coarse-only or the Fine-only results still suffer from residual misalignment in dynamic scenes. In contrast, the full two-stage approach achieves the best visual and quantitative results. Both the quantitative metrics in Tab. 2 and the qualitative comparisons in Fig. 6 confirm that the combined Coarse + Fine strategy is essential for achieving optimal reconstruction quality.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
4DGaussians	29.56	0.935	0.099
4DGaussians+Coarse	30.92	0.943	0.092
4DGaussians+Fine	30.87	0.941	0.091
4DGaussians+Full	31.16	0.942	0.091

Table 2: Our full two-stage method (4DGaussians+Full) achieves the best performance, highlighting the complementary roles of coarse and fine temporal refinement.

Ablation on Random Time Offsets To assess the robustness and performance of our method under varying levels of input asynchronicity, we conduct an ablation study by introducing different magnitudes of random time offsets. During training, we add a temporal jitter Δt to the timestamp of each input view, where Δt is sampled uniformly from the range $[0, \tau_{\max}]$. We evaluate our model’s performance by

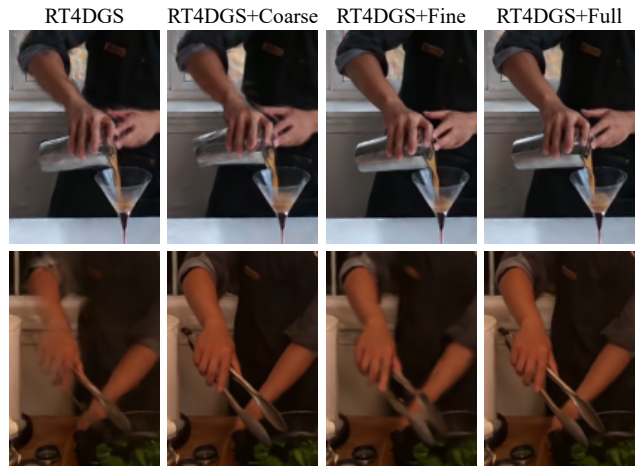


Figure 6: Ablation study results. The full method yields the best reconstruction quality, while removing components leads to visible performance drops.

progressively increasing the upper bound of this jitter, setting the maximum offset τ_{\max} to 3, 5, and 10 frames, respectively.

The experimental results in Tab. 3 clearly show that the performance of the Baseline model degrades dramatically as the maximum offset τ_{\max} increases. In stark contrast, the method augmented with our temporal refinement module maintains high reconstruction fidelity even under the most severe perturbation.

τ_{\max}	4DGaussians			+Ours		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
3	30.69	0.938	0.097	31.25	0.943	0.092
5	30.31	0.938	0.097	31.29	0.943	0.091
10	29.60	0.935	0.099	31.16	0.942	0.091

Table 3: Ablation study on random time offsets.

Conclusion

This paper addresses a prevalent yet challenging problem in computer vision and graphics: high-quality dynamic scene reconstruction from temporally unsynchronized multi-view videos. We introduce a novel temporal alignment module that, when integrated into existing 4D reconstruction frameworks, yields significant improvements in reconstruction quality. This was demonstrated on a custom benchmark created from the DyNeRF dataset, which was modified to include random temporal offsets to simulate real-world conditions. Ultimately, our work resolves the critical issue of temporal asynchrony, greatly expanding the application boundaries of high-quality 4D dynamic scene reconstruction. It enables high-fidelity dynamic capture using lower-cost, more flexible camera systems, thereby promoting the technology’s widespread adoption for real-world applications beyond controlled environments.

References

- Chen, Z.; Funkhouser, T.; Hedman, P.; and Tagliasacchi, A. 2023. MobileNeRF: Exploiting the Polygon Rasterization Pipeline for Efficient Neural Field Rendering on Mobile Architectures. In *The Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Cheng, K.; Long, X.; Yang, K.; Yao, Y.; Yin, W.; Ma, Y.; Wang, W.; and Chen, X. 2024. Gaussianpro: 3d gaussian splatting with progressive propagation. In *Forty-first International Conference on Machine Learning*.
- Duan, Y.; Wei, F.; Dai, Q.; He, Y.; Chen, W.; and Chen, B. 2024. 4d gaussian splatting: Towards efficient novel view synthesis for dynamic scenes. *arXiv e-prints*, arXiv-2402.
- Fischler, M. A.; and Bolles, R. C. 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6): 381–395.
- Fu, Y.; Liu, S.; Kulkarni, A.; Kautz, J.; Efros, A. A.; and Wang, X. 2024. COLMAP-Free 3D Gaussian Splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 20796–20805.
- Garbin, S. J.; Kowalski, M.; Johnson, M.; Shotton, J.; and Valentin, J. 2021. Fastnerf: High-fidelity neural rendering at 200fps. In *Proceedings of the IEEE/CVF international conference on computer vision*, 14346–14355.
- Hamdi, A.; Melas-Kyriazi, L.; Mai, J.; Qian, G.; Liu, R.; Vondrick, C.; Ghanem, B.; and Vedaldi, A. 2024. Ges: Generalized exponential splatting for efficient radiance field rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19812–19822.
- Huang, B.; Yu, Z.; Chen, A.; Geiger, A.; and Gao, S. 2024a. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 conference papers*, 1–11.
- Huang, Y.-H.; Sun, Y.-T.; Yang, Z.; Lyu, X.; Cao, Y.-P.; and Qi, X. 2024b. 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.
- Lei, J.; Weng, Y.; Harley, A.; Guibas, L.; and Daniilidis, K. 2024. Mosca: Dynamic gaussian fusion from casual videos via 4d motion scaffolds. *arXiv preprint arXiv:2405.17421*.
- Li, T.; Slavcheva, M.; Zollhoefer, M.; Green, S.; Lassner, C.; Kim, C.; Schmidt, T.; Lovegrove, S.; Goesele, M.; Newcombe, R.; et al. 2022. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5521–5531.
- Liang, H.; Ren, J.; Mirzaei, A.; Torralba, A.; Liu, Z.; Gilitschenski, I.; Fidler, S.; Oztireli, C.; Ling, H.; Gojcic, Z.; et al. 2024. Feed-Forward Bullet-Time Reconstruction of Dynamic Scenes from Monocular Videos. *arXiv preprint arXiv:2412.03526*.
- 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.
- LIU, Q.; Liu, Y.; Wang, J.; Lyu, X.; Wang, P.; Wang, W.; and Hou, J. 2025. MoDGS: Dynamic Gaussian Splatting from Casually-captured Monocular Videos with Depth Priors. In *The Thirteenth International Conference on Learning Representations*.
- 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.
- 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.
- Müller, T.; Evans, A.; Schied, C.; and Keller, A. 2022. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4): 1–15.
- Park, J.; Bui, M.-Q. V.; Bello, J. L. G.; Moon, J.; Oh, J.; and Kim, M. 2024. SplineGS: Robust Motion-Adaptive Spline for Real-Time Dynamic 3D Gaussians from Monocular Video. *arXiv preprint arXiv:2412.09982*.
- Reiser, C.; Peng, S.; Liao, Y.; and Geiger, A. 2021. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF international conference on computer vision*, 14335–14345.
- Ren, K.; Jiang, L.; Lu, T.; Yu, M.; Xu, L.; Ni, Z.; and Dai, B. 2024. Octree-gs: Towards consistent real-time rendering with lod-structured 3d gaussians. *arXiv preprint arXiv:2403.17898*.
- Stearns, C.; Harley, A.; Uy, M.; Dubost, F.; Tombari, F.; Wetstein, G.; and Guibas, L. 2024. Dynamic gaussian marbles for novel view synthesis of casual monocular videos. In *SIGGRAPH Asia 2024 Conference Papers*, 1–11.
- Sun, J.; Shen, Z.; Wang, Y.; Bao, H.; and Zhou, X. 2021. LoFTR: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8922–8931.
- Wang, Q.; Ye, V.; Gao, H.; Austin, J.; Li, Z.; and Kanazawa, A. 2024a. Shape of motion: 4d reconstruction from a single video. *arXiv preprint arXiv:2407.13764*.
- Wang, S.; Yang, X.; Shen, Q.; Jiang, Z.; and Wang, X. 2024b. Gflow: Recovering 4d world from monocular video. *arXiv preprint arXiv:2405.18426*.
- Wu, G.; Yi, T.; Fang, J.; Xie, L.; Zhang, X.; Wei, W.; Liu, W.; Tian, Q.; and Wang, X. 2024. 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, X.; Xu, J.; Zhu, Z.; Bao, H.; Huang, Q.; Tompkin, J.; and Xu, W. 2022. Scalable neural indoor scene rendering. *ACM transactions on graphics*, 41(4).

- Yan, J.; Peng, R.; Tang, L.; and Wang, R. 2024. 4D Gaussian Splatting with Scale-aware Residual Field and Adaptive Optimization for Real-time rendering of temporally complex dynamic scenes. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 7871–7880.
- Yang, J.; Gao, M.; Li, Z.; Gao, S.; Wang, F.; and Zheng, F. 2023. Track anything: Segment anything meets videos. *arXiv preprint arXiv:2304.11968*.
- Yang, Z.; Gao, X.; Zhou, W.; Jiao, S.; Zhang, Y.; and Jin, X. 2024a. 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.; Pan, Z.; Zhu, X.; Zhang, L.; Jiang, Y.-G.; and Torr, P. H. 2024b. 4D Gaussian Splatting: Modeling Dynamic Scenes with Native 4D Primitives. *arXiv preprint arXiv:2412.20720*.
- Yang, Z.; Yang, H.; Pan, Z.; and Zhang, L. 2024c. Real-time Photorealistic Dynamic Scene Representation and Rendering with 4D Gaussian Splatting. In *International Conference on Learning Representations (ICLR)*.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhu, R.; Liang, Y.; Chang, H.; Deng, J.; Lu, J.; Yang, W.; Zhang, T.; and Zhang, Y. 2024. Motiongs: Exploring explicit motion guidance for deformable 3d gaussian splatting. *Advances in Neural Information Processing Systems*, 37: 101790–101817.