

# CTRL-GS: Cascaded Temporal Residue Learning for 4D Gaussian Splatting

Karly Hou, Wanhua Li, Hanspeter Pfister  
Harvard University

karly@alumni.harvard.edu, wanhuali@g.harvard.edu, pfister@seas.harvard.edu

## Abstract

*Recently, Gaussian Splatting methods have emerged as a desirable substitute for prior Radiance Field methods for novel-view synthesis of scenes captured with multi-view images or videos. In this work, we propose a novel extension to 4D Gaussian Splatting for dynamic scenes. Drawing on ideas from residual learning, we hierarchically decompose the dynamic scene into a "video-segment-frame" structure, with segments dynamically adjusted by optical flow. Then, instead of directly predicting the time-dependent signals, we model the signal as the sum of video-constant values, segment-constant values, and frame-specific residuals, as inspired by the success of residual learning. This approach allows more flexible models that adapt to highly variable scenes. We demonstrate state-of-the-art visual quality and real-time rendering on several established datasets, with the greatest improvements on complex scenes with large movements, occlusions, and fine details, where current methods degrade most.*

## 1. Introduction

Novel view synthesis is an important task for applications including medical imaging, autonomous driving, robotics, AR/VR, animation and gaming, and more. However, representing and rendering dynamic scenes with high visual accuracy and efficiency is challenging, especially for scenes with complex geometry and motion.

In recent years, Neural Radiance Fields (NeRFs) [26] and their extensions have achieved high-quality novel view synthesis results by using implicit radiance fields to represent light distribution without explicit scene geometry definition. Point radiances are computed on-the-fly by prompting the neural network rather than explicitly stored, allowing differentiable and compact representations of complex scenes, but training and rendering can be slow due to volumetric ray marching.

Kerbl and Kopanas et al. [14] introduced 3D Gaussian Splatting (3D-GS) as an alternative to prior radiance fields, aiming to achieve high visual quality for unbounded

and complete static scenes (as opposed to isolated objects) and render high-resolution images in real time. Explicit scene representation with 3D Gaussians allows flexible manipulation and differentiable splatting significantly boosts rendering speed from volumetric rendering. Splatting [16, 22, 27, 44] involves projecting geometric primitives like points or discs onto an image plane to simulate the appearance of 3D objects.

4D Gaussian Splatting (4D-GS) [45] extends this work to dynamic scenes with moving objects by leveraging a spatial-temporal structure encoder, Gaussian deformation field, and deformation decoder. Instead of constructing 3D Gaussians at each timestamp, which is costly especially at scale, 4D-GS performs transformations on a canonical set of 3D Gaussians. 4D-GS achieves high-quality reconstructions while maintaining high rendering speed, but it often produces errors on scenes with large movements, occlusions, and fine details.

In this paper, we propose a novel extension to 4D Gaussian Splatting to address these challenges. We evaluate CTRL-GS on synthetic and real-world scenes with a range of types of motion. CTRL-GS achieves higher quantitative and qualitative reconstruction accuracy, with the largest improvements for challenging scenes with large movements. In summary, our main contributions include:

1. Three approaches to constructing temporal windows for spatial-temporal features.
2. Temporally-local dynamic MLPs that model deformation fields over a subsegment of time.
3. A hierarchical time signal modeling scheme that decomposes the dynamic scene into a "video-segment-frame residual" structure, offering both built-in scene continuity and flexibility for complex scenes.

## 2. Related Work

### 2.1. NeRF Methods

Nerfies [29] extends NeRF to optimize a deformation field per observation to handle non-rigidly deforming objects. HyperNeRF [30] models more complex scene dynamics through a higher-dimension canonical space and a learned



Figure 1. Our method achieves improved reconstruction accuracy for dynamic scenes at high image resolutions across settings. Current models tend to perform poorly on scenes with high motion, occluded areas, and fine details. **Left:** CTRL-GS correctly constructs the broom handle and more accurate shading details, where the 4D-GS reconstruction is blurry and omits the handle. **Right:** CTRL-GS reproduces fine details with accurate structure and depth, where the 4D-GS reconstruction omits fine textures and lighting changes.

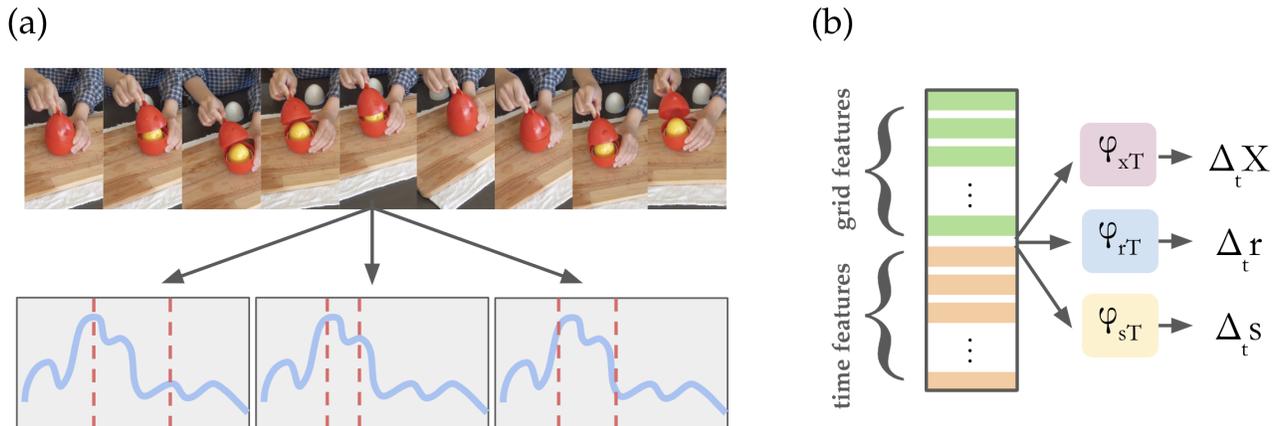


Figure 2. Two key components of our method. a) We divide the frames of the dynamic scene into temporal windows. We assess three different methods: 1) equal divisions, 2) divisions at the  $N$  frame-pairs of highest optical flow, 3) divisions based on greedy optical flow thresholds. b) We construct MLPs  $\varphi_{xT}$ ,  $\varphi_{rT}$ ,  $\varphi_{sT}$ , which take encoded spatial-temporal features and produce intermediate mean components for position, rotation, and scaling for each temporal window individually. These components are combined with the original 3D Gaussians and point-in-time deformations to construct our final features.

warp field to handle both geometric and temporal variations. FFDNeRF [10] uses differentiable forward flow motion modeling for dynamic view synthesis. TiNeuVox-B [5] represents scenes with time-aware voxel features and uses a multi-distance interpolation method to model both small and large motions. V4D [9] uses 3D voxels to directly model 4D neural radiance fields without the need for a canonical space. K-Planes [7] uses  $d$  choose 2 planes to represent a  $d$ -dimension scene, allowing easy addition of dimension-specific priors. A variety of other works [2–4, 6, 19, 24, 25, 48] explore related ideas.

## 2.2. Gaussian Splatting

A variety of work builds on 3D-GS [14] to achieve higher accuracy, faster rendering, more efficient representation, flexible manipulation, and diverse applications. [8] demonstrates COLMAP-free 3D-GS. GauFRe [20] constructs deformable models and separate static and dynamic components. DynMF [15] decomposes per-point motions in a dynamic scene into a small set of explicit trajectories. SC-GS [12] uses sparse control points and dense Gaussians to enable user-controlled motion editing. PhysGaussian [46] integrates physical dynamics, combining 3D-GS with kinematic deformations and mechanical stress attributes. SG-

Splatting [42] accelerates rendering speed with Spherical Gaussians. SwinGS [38] uses temporally-local windows for dynamic 3D Gaussians. DreamGaussian4D [35] refines generated 4D Gaussians with a pre-trained image-to-video diffusion model. GAGS [31] incorporates CLIP features for scene understanding.

Many Gaussian Splatting applications [1, 17, 18, 23, 28, 33, 34] have been demonstrated, such as [36] for human and face models, [50] for urban and autonomous driving scenes, [13] for SLAM, [47] for reflective surfaces, and more.

### 2.3. Residual Learning

The original deep residual learning for image recognition paper [11] introduces the idea that instead of approximating a target function  $H(x)$  directly, a network can be designed to instead approximate a residual function  $F(x)$  such that  $H(x) = F(x) + x$ . Though both forms may be able to asymptotically approximate the desired functions, the reformulation may be easier to learn. If additional layers can be constructed as identity mappings, then a deeper model should be able to perform at least as well as a shallower one; with the residual learning reformulation, the model can drive weights of nonlinear layers toward zero to approach identity mappings. If the optimal function is closer to an identity mapping than to zero mapping, it should be easier for the model to find perturbations with reference to an identity mapping than to learn the entire function. This idea inspires our hierarchical decomposition in a temporal setting.

## 3. Method

### 3.1. Preliminaries: 3D Gaussian Splatting

Starting from a sparse set of SfM [37] points, each 3D Gaussian is defined by covariance matrix  $\Sigma$  in world space [51] centered at point  $\mathcal{X}$ :

$$G(\mathcal{X}) = e^{-\frac{1}{2}\mathcal{X}^T\Sigma^{-1}\mathcal{X}} \quad (1)$$

For rendering,  $\Sigma$  is converted to a  $\Sigma'$  covariance matrix represented in camera coordinates:

$$\Sigma' = JW\Sigma W^T J^T \quad (2)$$

where  $W$  is a viewing transformation matrix and  $J$  is the Jacobian of the affine approximation of the projective transformation.

This covariance  $\Sigma$  can be decomposed into scaling matrix  $S$  and rotation matrix  $R$  for optimization:

$$\Sigma = RSS^T R^T \quad (3)$$

Next, optimization of positions  $\mathcal{X} \in \mathbb{R}^3$ , color as defined by spherical harmonic (SH) coefficients  $\mathcal{C} \in \mathbb{R}^k$  (where  $k$  represents the number of SH functions), opacities  $\alpha \in \mathbb{R}$ ,

rotations  $r \in \mathbb{R}^4$ , and scaling factors  $s \in \mathbb{R}^3$  is interleaved with steps that adaptively control density of the Gaussians.

Finally, a tile-based rasterizer filters out Gaussians unlikely to be visible in the viewing area, new instances of the Gaussians are sorted, and the  $N$  points overlapping a pixel are blended using

$$C = \sum_{i \in N} c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j) \quad (4)$$

where  $c_i$  and  $\alpha_i$  represent the color and opacity of each point.

### 3.2. Preliminaries: 4D Gaussian Splatting

4D-GS [45] leverages a Gaussian deformation field network  $\mathcal{F}$ , spatial-temporal structure encoder  $\mathcal{H}$ , and multi-head deformation decoder  $\mathcal{D}$  to produce deformed 3D Gaussians  $\mathcal{G}'$  from the original 3D Gaussians  $\mathcal{G}$ .

First,  $\mathcal{H}$  including a multi-resolution Hex-Plane  $R(i, j)$  and a tiny MLP  $\phi_d$  encode spatial-temporal features:  $\mathcal{H}(\mathcal{G}, t) = \{R_l(i, j), \phi_d(i, j) \in \{(x, y), (x, z), (y, z), (x, t), (y, t), (z, t)\}, l \in \{1, 2\}$ , where  $\mu = (x, y, z)$  is the mean of 3D Gaussians  $\mathcal{G}$ .

A multi-head Gaussian deformation decoder  $\mathcal{D} = \{\varphi_x, \varphi_r, \varphi_s\}$  computes deformation components for position, rotation, and scaling, respectively, where each  $\varphi$  is a separate MLP, leading to deformed features

$$(\mathcal{X}', r', s') = (\mathcal{X} + \Delta\mathcal{X}, r + \Delta r, s + \Delta s) \quad (5)$$

where  $\Delta\mathcal{X} = \varphi_x(\mathcal{X})$ ,  $\Delta r = \varphi_r(r)$ ,  $\Delta s = \varphi_s(s)$ . So the deformed 3D Gaussians are  $\mathcal{G}' = \mathcal{G} + \Delta\mathcal{G} = \{\mathcal{X}', s', r', \sigma, \mathcal{C}\}$ . Finally, a Gaussian conversion process based on timestamps  $t_i$  maintains differential splatting efficiencies.

### 3.3. CTRL-GS: Overview

Key components of our method are depicted in Figure 2. We build on the 4D-GS [45] method for dynamic scenes.

At a high level, our cascaded decomposition is motivated by the real-world observation that dynamic scenes often exhibit hierarchical temporal structure: global geometry and lighting remain largely stable across the video, mid-level motion patterns (e.g. limb movement, deforming surfaces) evolve at an intermediate rate, and fine-grained changes (e.g. shading, occlusion boundaries) fluctuate on a per-frame basis.

We first divide the frames of the dynamic scene into temporal windows. We experiment with three window construction methods: equal length, N-highest-flow, and dynamic thresholds. Next, we create reusable intermediate MLPs for position, rotation, and scaling, which are applied to the temporal windows and learn the deformation field from canonical 3D Gaussians for the frames within that

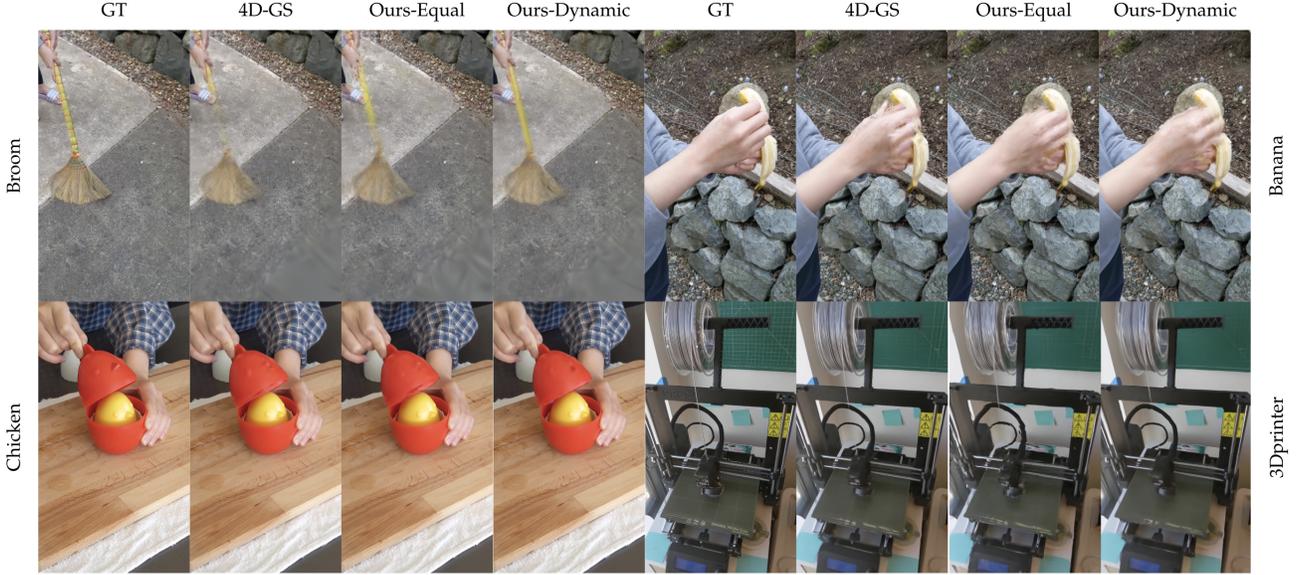


Figure 3. Comparisons between ground truth (GT), baseline 4D-GS [45] result, and our method with equal and dynamic threshold-based windows on HyperNeRF [30] validation rig scenes. **Broom:** CTRL-GS correctly constructs the broom handle and offers better depth shading. **Banana:** CTRL-GS correctly places fingers. **Chicken:** CTRL-GS preserves edge structural details and beak shading, and reduces blur. **3D Printer:** CTRL-GS does not produce improved results on this example.

window. Finally, our deformed features are computed as the sum of video-constant features, segment-specific means, and point-in-time deformations.

This ”cascaded” approach allows us to capture changing scene information with greater flexibility and accuracy. We take advantage of information shared throughout the entire scene, while affording flexibility for appearing and disappearing elements and regions, large movements, and other temporal scene shifts.

### 3.4. Equal Temporal Window Construction

Let timestamps  $t_i$  within a scene fall within the range  $[0,1]$ , and  $N$  be the number of periods we select for temporal segmentation. Then the interval length  $I = \frac{1}{N}$ . We perform the following temporal feature transformation:

$$\mathcal{T} = \lfloor \frac{\mathcal{T}}{I + 10^{-9}} \rfloor * I + q * I \quad (6)$$

where  $\mathcal{T}$  is an array containing the point timestamp of the same width as the other grid features and  $q \in \mathbb{Q}^+$  is the quantization coefficient, a performance hyperparameter that adjusts the interpolation between segment-level predictions and frame-specific residuals. We include  $10^{-9}$  to ensure the floor function performs timestamp quantization as expected. Finally, we append this constructed temporal feature to the other existing grid features (representing  $\mathcal{X}, r, s$ ) to produce features with quantized temporal information  $\mathcal{X}_{\mathcal{T}}, r_{\mathcal{T}}, s_{\mathcal{T}}$ , as input to the segment-specific MLPs.

### 3.5. Dynamic Window Construction: N-Highest-Flow

Instead of simply dividing the frames into equal temporal windows, we can observe video characteristics and construct windows adaptively. Using a pretrained RAFT [39] ’raft-things’ model and one camera perspective, we calculate the optical flow  $f_{ij}$  between each pair  $i, j$  of consecutive frames in a dynamic scene.

In each scene, we notice several ’peaks’ of highest motion. These high-motion periods tend to divide a scene into different visual settings. Because models may have trouble learning over these high-motion periods, we use separate intermediate MLPs to learn before/after the moment of high motion. To yield  $N$  temporal windows, we select the top  $N - 1$  frame-pairs  $i, j$  with highest optical flow  $f_{ij}$  as window divisions, so frame  $i$  ends the previous window and  $j$  begins the next.

### 3.6. Dynamic Window Construction: Greedy Thresholds

Extending this idea, we can try to construct more, shorter windows during periods of high motion where additional granularity can help learn significant geometry changes. We construct fewer, longer windows for periods of low motion where geometry is more temporally consistent. This approach allows us to incorporate more flexibility, leading to better reconstructions for high-motion scenes where current

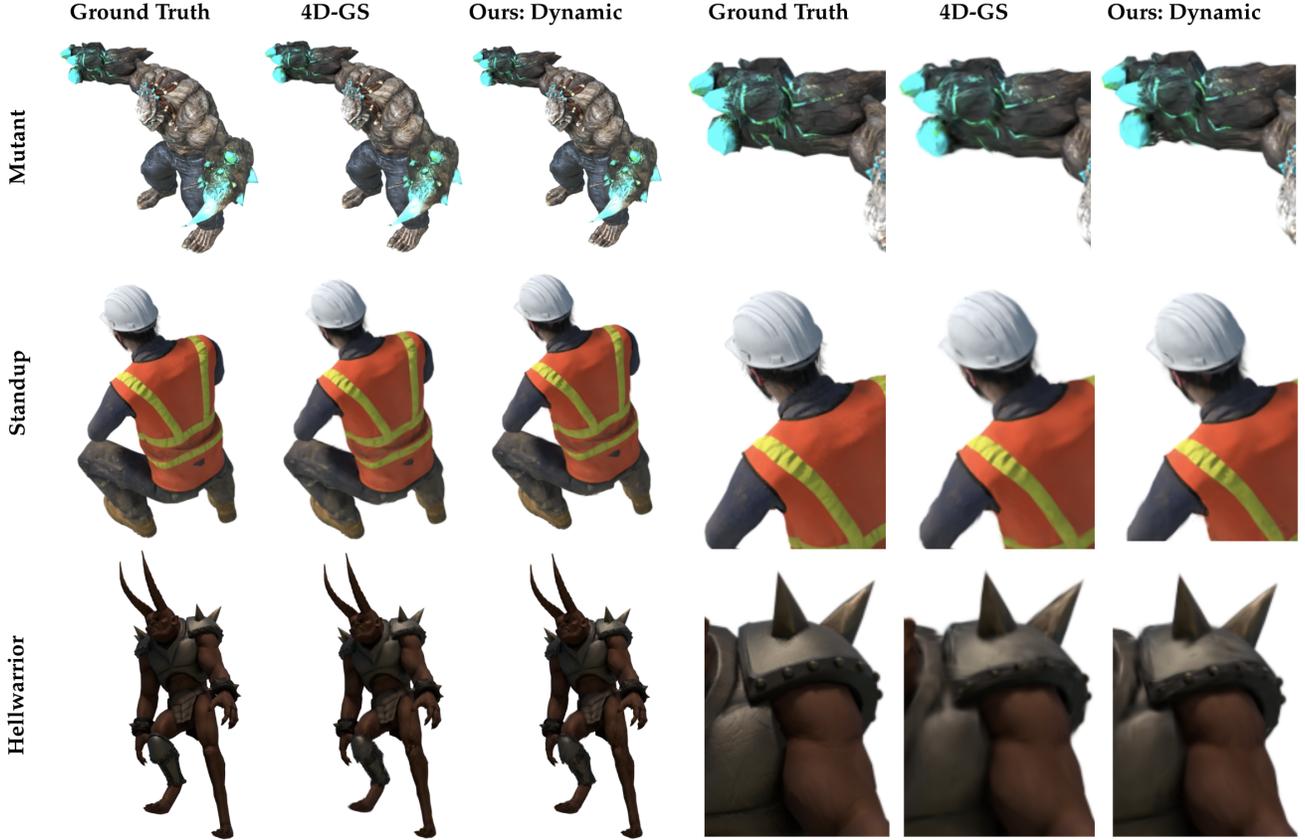


Figure 4. Comparisons between ground truth (GT), 4D-GS, and our method with dynamic threshold-based windows on D-NeRF [30] synthetic scenes. CTRL-GS shows greater reconstruction fidelity and less blur, especially on fine details and sharp edges. **Mutant**: See blue claw structure, clarity of small blue veins, and placement of cracks. **Standup**: See helmet edges, helmet ridges, and vest wrinkles. **Hellwarrior**: See shoulder spheres, chestplate ridges, and luminosity of collar highlight.

models tend to degrade most.

After selecting  $N$  time periods for segmentation, we set threshold:

$$T = \frac{\sum f_{ij}}{N} \quad (7)$$

Iterating over the frames, we follow a greedy approach, keeping a sum of total flow so far. When we reach the threshold, we begin a new window and reset the flow sum.

### 3.7. Cascaded Temporal Residue Deformation

After constructing temporal windows from our dynamic scene with one of the three prior methods, we introduce three new MLPs  $\varphi_{xT}, \varphi_{rT}, \varphi_{sT}$ , which produce intermediate mean components for position, rotation, and scaling for each temporal segment individually:  $\Delta_t \mathcal{X} = \varphi_{xT}(\mathcal{X}_T), \Delta_t r = \varphi_{rT}(r_T), \Delta_t s = \varphi_{sT}(s_T)$ .

Now our deformed features are:

$$(\mathcal{X}', r', s') = (\mathcal{X} + \Delta_t \mathcal{X} + \Delta \mathcal{X}, r + \Delta_t r + \Delta r, s + \Delta_t s + \Delta s) \quad (8)$$

Finally, our deformed 4D Gaussians  $\mathcal{G}'$  are defined by the updated features and the original opacity and spherical harmonic features:  $\{\mathcal{X}', s', r', \alpha, \mathcal{C}\}$ .

By using a "cascaded" approach, we model spatial-temporal features as the sum of video-constant features, segment-specific means, and point-in-time deformations. This approach offers greater reconstructive control over changing scenes.

## 4. Experiments

### 4.1. Setting and Implementation Details

We build on the implementation of 4D-GS [45] using PyTorch. We use COLMAP [37] to generate point clouds to initialize 3D Gaussians. To improve stability and convergence, we "warm up" optimization using smaller image resolution for 3000 initial iterations, which are then upsampled. We perform densification every 100 iterations until 10000 iterations, removing Gaussians with opacity  $\alpha$  below threshold  $\delta_\alpha$ , which are essentially transparent. We observe

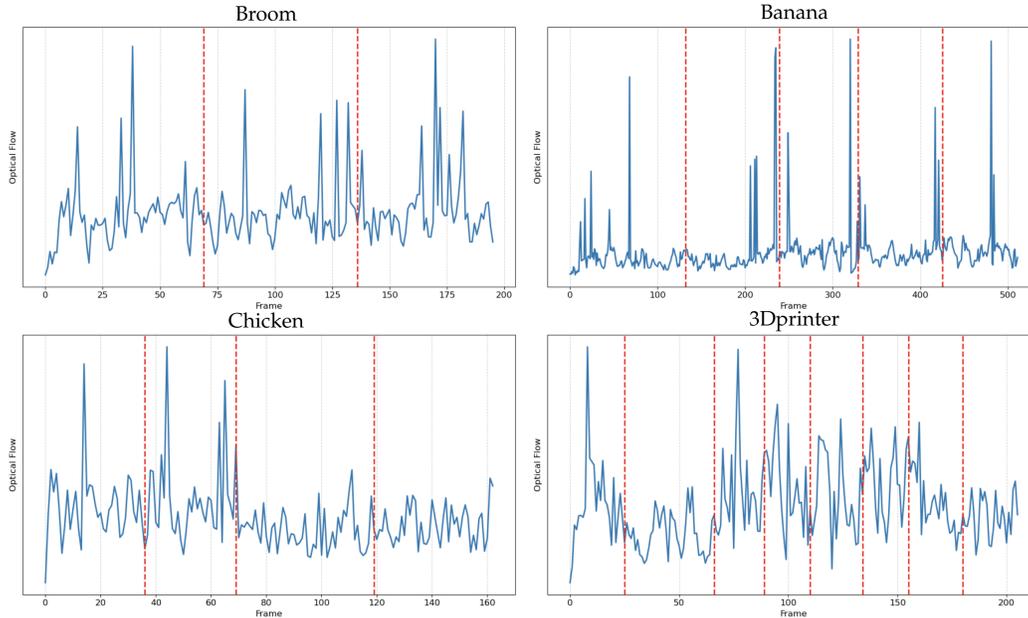


Figure 5. Magnitude of optical flow  $f_{ij}$  over each of the HyperNeRF [30] datasets a) BROOM, b) BANANA, c) CHICKEN, 3DPRINTER. Red vertical lines denote optimal temporal window segmentation when constructed with the dynamic thresholds method.

Table 1. Quantitative results on HyperNeRF [30] vrig datasets with rendering resolution 960 x 540. First, second, and third best results are highlighted in yellow, blue, and pink. CTRL-GS achieves the most improvements with dynamic threshold windows.

Model	PSNR (dB) $\uparrow$	MS-SSIM $\uparrow$	FPS $\downarrow$
3D-GS [14]	19.7	0.680	55
Nerfies [29]	22.2	0.803	<1
HyperNeRF [30]	22.3	0.814	<1
FFDNeRF [10]	24.2	0.842	0.05
TiNeuVox-B [5]	24.3	0.836	1
V4D [9]	24.8	0.832	0.29
4D-GS [45]	25.5	0.845	34
Equal Windows	25.8	0.831	27
Dyn. N-highest	25.9	0.860	25
Dyn. Threshold	26.0	0.863	22

30,000 iterations for each scene and report best results for our models against baselines. All training and testing is performed on a single RTX A4000 GPU.

## 4.2. Real-World Datasets

We primarily assess the performance of our models on the 4 HyperNeRF [30] vrig datasets, which contain real-world scenarios captured by two cameras. Each video, originally 30-60s long, has been sub-sampled to 15fps. The videos

range approximately 300-1000 frames; the datasets are constructed by using every 4th frame for training, and the middle frame between each training pair for validation.

## 4.3. Synthetic Datasets

We also evaluate our model on the 8 D-NeRF [32] datasets, which contain monocular synthetic dynamic scenes without backgrounds. Each video contains 50-200 frames.

## 4.4. Evaluation

We primarily assess experimental results using metrics like peak-signal-to-noise ratio (PSNR), perceptual quality measure LPIPS [49], structural similarity index (SSIM) [43], and its extension multiscale SSIM (MS-SSIM). We also compare our results with several other state-of-the-art methods as described previously, including Nerfies [29], HyperNeRF [30], FFDNeRF [10], TiNeuVox [5], V4D [9], K-Planes [7], HexPlane [4], MSTH [41], 3D Gaussian Splatting [14], and 4D Gaussian Splatting [45]. Results for comparison methods are from their papers, code, or otherwise provided by their authors. Rendering speed is borrowed from estimates by [45] on official implementations.

## 4.5. Quantitative Results

Results on the real-world datasets are summarized in Tab. 1, and enumerated per scene in Tab. 3. Results on the synthetic datasets are summarized in Tab. 2, and enumerated per scene in Tab. 4. CTRL-GS shows improved accuracy across scenes, with the greatest improvements occurring in scenes

Table 2. Quantitative results on D-NeRF [32] synthetic datasets with rendering resolution 1352 x 1014. First, second, and third best results are highlighted in yellow, blue, and pink. CTRL-GS with dynamic threshold-based temporal windows achieves the highest PSNR and SSIM metrics, but renders more slowly than 4D-GS.

Model	PSNR (dB) $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FPS $\uparrow$
TiNeuVox-B [5]	32.67	0.97	0.04	1.5
K-Planes [7]	31.61	0.97	0.04	0.97
HexPlane-Slim [4]	31.04	0.97	0.04	2.5
3D-GS [14]	23.19	0.93	0.12	170
FFDNeRF [10]	32.68	0.97	0.04	<1
MSTH [41]	31.34	0.98	0.02	6
V4D [9]	33.72	0.98	0.02	82
4D-GS [45]	34.05	0.98	0.02	82
Ours: Equal Windows	32.87	0.97	0.04	73
Ours: Dynamic N-highest	33.90	0.98	0.03	69
Ours: Dynamic Threshold	34.34	0.98	0.03	63

Table 3. Per-scene results on the HyperNeRF validation rig datasets [30] across models. First, second, and third best results are highlighted in yellow, blue, and pink, respectively. CTRL-GS achieves improved results across most datasets across multiple metrics.

Method	3D Printer		Chicken		Broom		Banana	
	PSNR	MS-SSIM	PSNR	MS-SSIM	PSNR	MS-SSIM	PSNR	MS-SSIM
Nerfies [29]	20.6	0.83	26.7	0.94	19.2	0.56	22.4	0.87
HyperNeRF [30]	20.0	0.59	26.9	0.94	19.3	0.59	23.3	0.90
TiNeuVox-B [5]	22.8	0.84	28.3	0.95	21.5	0.69	24.4	0.87
FFDNeRF [10]	22.8	0.84	28.0	0.94	21.9	0.71	24.3	0.86
3D-GS [14]	18.3	0.60	19.7	0.70	20.6	0.63	24.0	0.80
4D-GS [45]	22.1	0.81	28.7	0.93	22.0	0.70	28.0	0.94
Ours: Equal Windows	22.2	0.74	29.1	0.93	22.3	0.72	29.4	0.92
Ours: Dynamic N-highest	21.9	0.82	29.5	0.93	22.9	0.75	29.4	0.94
Ours: Dynamic Threshold	21.9	0.82	29.6	0.94	22.9	0.75	29.5	0.94

with high motion, occluded areas, or fine detail that is challenging for current models to reconstruct. For example, on the BROOM sequence—characterized by all of the aforementioned challenges, CTRL-GS achieves 22.9 PSNR, outperforming all cited methods that presented scene-level results. These include (in addition to the results collated in the tables) DynMF [15] (22.1), Gaussian-Flow [21] (22.8), and SP-GS [40] (22.8).

Like other Gaussian Splatting models, CTRL-GS enjoys much higher rendering speed than NeRF-based methods. We observed some slowdown from 4D-GS, which is primarily due to repeated segment identification during rendering, which was introduced to support our dynamic segment-based querying. Specifically, in the forward-dynamic method, we perform per-frame checks to determine segment membership. The underlying issue is not the segment strategy itself, but rather the inefficiency of dy-

namically querying segment info on-the-fly. This can be fixed by precomputing segment indices during training and caching them for efficient lookup during rendering. We plan to implement this improvement for the official release.

#### 4.6. Qualitative Results

Qualitative results are presented in Fig. 3 for real-world datasets and 4 for synthetic datasets. Our method can more accurately reconstruct structure, handle edges more smoothly, recover more accurate color and lighting, and preserve fine details. For example, in HyperNeRF-BANANA, CTRL-GS correctly places fingers where 4D-GS constructs a thumb not present in the ground truth. In CHICKEN, CTRL-GS correctly constructs a smooth lid and clear fingernail contours, where 4D-GS produces incorrect lid holes and fingertip splintering. In D-NeRF-MUTANT, CTRL-GS captures the fine details of dark cracks, blue

Table 4. Per-scene results on the synthetic datasets across models. First, second, and third best results are highlighted in yellow, blue, and pink, respectively. CTRL-GS achieves improved results across most datasets across multiple metrics.

Method	Bouncing Balls			Hellwarrior			Hook			Jumpingjacks		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
3D-GS	23.20	0.9591	0.0600	24.53	0.9336	0.0580	21.71	0.8876	0.1034	23.20	0.9591	0.0600
K-Planes	40.05	0.9934	0.0322	24.58	0.9520	0.0824	28.12	0.9489	0.0662	31.11	0.9708	0.0468
HexPlane	39.86	0.9915	0.0323	25.00	0.9443	0.0732	28.63	0.9433	0.0636	33.49	0.9772	0.0398
TiNeuVox	40.23	0.9926	0.0416	27.10	0.9638	0.0786	28.63	0.9433	0.0636	34.39	0.9771	0.0408
4D-GS	40.62	0.9942	0.0155	28.71	0.9733	0.0369	32.73	0.9760	0.0272	35.42	0.9857	0.0128
Ours: Threshold	40.84	0.9946	0.0142	29.33	0.9742	0.0354	33.15	0.9765	0.0263	35.90	0.9862	0.0182

Method	Lego			Mutant			Standup			Trex		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
3D-GS	23.06	0.9290	0.0642	20.64	0.9297	0.0828	21.91	0.9301	0.0785	21.93	0.9539	0.0487
K-Planes	25.49	0.9483	0.0331	32.50	0.9713	0.0362	33.10	0.9793	0.0310	30.43	0.9737	0.0343
HexPlane	25.10	0.9388	0.0437	33.67	0.9802	0.0261	34.40	0.9839	0.0204	36.07	0.9749	0.0273
TiNeuVox	24.65	0.9063	0.0648	30.87	0.9607	0.0474	34.61	0.9797	0.0326	31.25	0.9666	0.0478
4D-GS	25.03	0.9376	0.0382	37.59	0.9880	0.0167	38.11	0.9898	0.0074	34.23	0.9850	0.0131
Ours: Threshold	25.23	0.9362	0.0611	38.16	0.9897	0.0143	38.15	0.9906	0.0126	33.94	0.9858	0.0203

veins, and surface textures, also preserving original luminosity. Note that 4D-GS shows significant blurring both within and on the edges of the structures. In STANDUP, CTRL-GS replicates the crisp and precise helmet contours, more accurate stripe texture, and beard shape, where 4D-GS performs incorrect texture and shape blurring throughout. In HELLWARRIOR, our method correctly constructs the spheres on the armor, upper breastplate segmentation, highlight luminosity and shadow depth, which are missed by 4D-GS. These improvements demonstrate an improved ability to capture realism, especially in scenes with fine details and textures. Additional comparison frames can be found in the appendix.

#### 4.7. Ablation Studies

We compare our three temporal window construction methods and find that the greedy threshold-based dynamic method tends to perform best across settings. However, even simpler methods like equal temporal windows and N-highest windows offer improvements from baseline models.

#### 4.8. Limitations

While CTRL-GS does not uniformly surpass all existing methods across every dataset, it delivers notable and consistent improvements on several key benchmarks and scenes (for example, HyperNeRF’s BROOM). In general, CTRL-GS offers the most meaningful advances on challenging dynamic scenes where range of motion is high and/or complex occlusions are present, cases where existing approaches degrade most (often producing significantly inaccurate recon-

structions with holes, visual sharding, temporal discontinuity, and ghosting). The benchmark datasets offer a range of scenes with different characteristics. In scenes where motion is limited and temporal variation is low, existing models already perform reasonably well, and improvements offered by CTRL-GS are not as significant.

### 5. Conclusion

This paper proposes Cascaded Temporal Residue Learning for 4D-GS (CTRL-GS) to improve visual fidelity in reconstruction tasks. We evaluate our models across a wide set of real-world and synthetic data, with different kinds of motion, lighting, and occlusions, and find that CTRL-GS produces the greatest improvements on challenging settings where there is high motion throughout the temporal sequence, where current models struggle most, often producing significantly inaccurate reconstructions with holes, visual sharding, temporal discontinuity, and ghosting. Future work may explore 1) further reducing aberrations, 2) static-dynamic decomposition, 3) adaptive opacity controls for objects entering and leaving the scene, and 4) improving rendering speed.

### Acknowledgments

This work is supported in part by NIH grant R01HD104969.

### References

- [1] Rameen Abdal, Wang Yifan, Zifan Shi, Yinghao Xu, Ryan Po, Zhengfei Kuang, Qifeng Chen, Dit-Yan Yeung, and Gor-

- don Wetzstein. Gaussian shell maps for efficient 3d human generation. In *CVPR*, 2024. 3
- [2] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [3] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [4] Ang Cao and Justin Johnson. Hexplane: a fast representation for dynamic scenes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 6, 7
- [5] Jiemin Fang, Taoran Yi, Xinggang Wang, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Matthias Nießner, and Qi Tian. Fast dynamic radiance fields with time-aware neural voxels. In *SIGGRAPH Asia Conference Papers*, 2022. 2, 6, 7
- [6] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: radiance fields without neural networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [7] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbaek Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: explicit radiance fields in space, time, and appearance. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 6, 7
- [8] Yang Fu, Sifei Liu, Amey Kulkarni, Jan Kautz, Alexei A. Efros, and Xiaolong Wang. Colmap-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20796–20805, 2024. 2
- [9] Wnhui Gan, Hongbin Xu, Yi Huang, Shifeng Chen, and Naoto Yokoya. V4d: Voxel for 4d novel view synthesis. In *IEEE Transactions on Visualization and Computer Graphics*, 2023. 2, 6, 7
- [10] X. Guo, J. Sun, Y. Dai, G. Chen, X. Ye, X. Tan, E. Ding, Y. Zhang, and J. Wang. Forward flow for novel view synthesis of dynamic scenes. In *International Conference on Computer Vision (ICCV)*, 2023. 2, 6, 7
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 3
- [12] Yi-Hua Huang, Yang-Tian Sun, Ziyi Yang, Xiaoyang Lyu, Yan-Pei Cao, and Xiaojuan Qi. Sc-gs: Sparse-controlled gaussian splatting for editable dynamic scenes. In *CVPR*, 2024. 2
- [13] Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer, Deva Ramanan, and Jonathon Luiten. Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 3
- [14] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42(4), 2023. 1, 2, 6, 7
- [15] Agelos Kratimenos, Jiahui Lei, and Kostas Daniilidis. Dynmf: Neural motion factorization for real-time dynamic view synthesis with 3d gaussian splatting. In *ECCV*, 2024. 2, 7
- [16] D. Laur and P. Hanrahan. Hierarchical splatting: a progressive refinement algorithm for volume rendering. In *Computer Graphics SIGGRAPH*, 1991. 1
- [17] Mingwei Li, Jiachen Tao, Zongxin Yang, and Yi Yang. Human101: Training 100+fps human gaussians in 100s from 1 view, 2023. 3
- [18] Wanhua Li, Renping Zhou, Jiawei Zhou, Yingwei Song, Johannes Herter, Minghan Qin, Gao Huang, and Hanspeter Pfister. 4d langsplat: 4d language gaussian splatting via multimodal large language models. *arXiv preprint arXiv:2503.10437*, 2025. 3
- [19] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [20] Yiqing Liang, Numair Khan, Zhengqin Li, Thu Nguyen-Phuoc, Douglas Lanman, James Tompkin, and Lei Xiao. Gafre: Gaussian deformation fields for real-time dynamic novel view synthesis. In *Proc. IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2025. 2
- [21] Youtian Lin, Zuozhuo Dai, Siyu Zhu, and Yao Yao. Gaussian-flow: 4d reconstruction with dynamic 3d gaussian particle. In *CVPR*, 2024. 7
- [22] L. Lippert and M.H. Gross. Fast wavelet based volume rendering by accumulation of transparent texture maps. *Computer Graphics Forum*, 14(3), 1995. 1
- [23] Xian Liu, Xiaohang Zhan, Jiayang Tang, Ying Shan, Gang Zeng, Dahua Lin, Xihui Liu, and Ziwei Liu. Humangaussian: Text-driven 3d human generation with gaussian splatting. *arXiv preprint arXiv:2311.17061*, 2023. 3
- [24] Yu-Lun Liu, Chen Gao, Andreas Meuleman, Hung-Yu Tseng, Ayush Saraf, Changil Kim, Yung-Yu Chuang, Johannes Kopf, and Jia-Bin Huang. Robust dynamic radiance fields. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [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), 2019. 2
- [26] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Monathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision (ECCV)*, 2020. 1
- [27] Klaus Mueller and Roni Yagel. Fast perspective volume rendering with splatting by utilizing a ray-driven approach. In *IEEE Visualization*, 1996. 1
- [28] Panwang Pan, Zhuo Su, Chenguo Lin, Zhen Fan, Yongjie Zhang, Zeming Li, Tingting Shen, Yadong Mu, and Yebin Liu. Humansplat: Generalizable single-image human gaussian splatting with structure priors. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 3

- [29] K. Park, U. Sinha, J. T. Barron, S. Bouaziz, D. B. Goldman, S. M. Seitz, and R. Martin-Brualla. Nerfies: deformable neural radiance fields. In *International Conference on Computer Vision (ICCV)*, 2021. 1, 6, 7
- [30] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *ACM Transactions on Graphics*, 40(6), 2021. 1, 4, 5, 6, 7
- [31] Yuning Peng, Haiping Wang, Yuan Liu, Chenglu Wen, Zhen Dong, and Bisheng Yang. Gags: Granularity-aware 3d feature distillation for gaussian splatting. *arXiv preprint arXiv:2412.13654*, 2024. 3
- [32] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: neural radiance fields for dynamic scenes. In *CVPR*, 2021. 6, 7, 1
- [33] Zhiyin Qian, Shaofei Wang, Marko Mihajlovic, Andreas Geiger, and Siyu Tang. 3dgs-avatar: Animatable avatars via deformable 3d gaussian splatting. In *CVPR*, 2024. 3
- [34] Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. Langsplat: 3d language gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20051–20060, 2024. 3
- [35] Jiawei Ren, Liang Pan, Jiaxiang Tang, Chi Zhang, Ang Cao, Gang Zeng, and Ziwei Liu. Dreamgaussian4d: Generative 4d gaussian splatting. *arXiv preprint arXiv:2312.17142*, 2023. 3
- [36] Shunsuke Saito, Gabriel Schwartz, Tomas Simon, Junxuan Li, and Giljoo Nam. Relightable gaussian codec avatars. In *CVPR*, 2024. 3
- [37] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3, 5
- [38] Richard Shaw, Michal Nazarczuk, Jifei Song, Arthur Moreau, Sibi Catley-Chandar, Helisa Dharmo, and Eduardo Perez-Pellitero. Swings: Sliding windows for dynamic 3d gaussian splatting. *arXiv preprint arXiv:2312.13308*, 2023. 3
- [39] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *European Conference on Computer Vision*, 2020. 4
- [40] Diwen Wan, Ruijie Lu, and Gang Zeng. Superpoint gaussian splatting for real-time high-fidelity dynamic scene reconstruction. In *ICML*, 2024. 7
- [41] Feng Wang, Zilong Chen, Guokang Wang, Yafei Song, and Huaping Liu. Masked space-time hash encoding for efficient dynamic scene reconstruction. In *NeurIPS*, 2023. 6, 7
- [42] Yiwen Wang, Siyuan Chen, and Ran Yi. Sg-splatting: accelerating 3d gaussian splatting with spherical gaussians. *arXiv preprint arXiv:2501.00342*, 2024. 3
- [43] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. In *IEEE Transactions on Image Processing*, 2004. 6
- [44] Lee Westover. Interactive volume rendering. In *Proceedings of the Chapel Hill workshop on Volume Visualization*, 1989. 1
- [45] GuanJun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Wang Xinggang. 4D Gaussian Splatting for Real-Time Dynamic Scene Rendering. *arXiv preprint arXiv:2310.08528*, 2023. 1, 3, 4, 5, 6, 7
- [46] Tianyi Xie, Zeshun Zong, Yuxing Qiu, Xuan Li, Yutao Feng, Yin Yang, and Chenfanfu Jiang. Physgaussian: Physics-integrated 3d gaussians for generative dynamics. *arXiv preprint arXiv:2311.12198*, 2023. 2
- [47] Yuxuan Yao, Zixuan Zeng, Chun Gu, Xiatian Zhu, and Li Zhang. Reflective gaussian splatting. *arXiv preprint arXiv:2412.19282*, 2024. 3
- [48] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492v2*, 2020. 2
- [49] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [50] Hongyu Zhou, Jiahao Shao, Lu Xu, Dongfeng Bai, Weichao Qiu, Bingbing Liu, Yue Wang, Andreas Geiger, and Yiyi Liao. Hugs: Holistic urban 3d scene understanding via gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21336–21345, 2024. 3
- [51] Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Ewa volume splatting. In *IEEE Proceedings Visualization*, 2001. 3

# CTRL-GS: Cascaded Temporal Residue Learning for 4D Gaussian Splatting

## Supplementary Material



Figure 6. Additional frame comparisons for the HyperNeRF [30] BROOM dataset across ground truth (GT), baseline 4D-GS, CTRL-GS with equal temporal windows, and CTRL-GS with dynamic threshold-based windows. In frame 177, our model shows a full handle reconstruction. In frame 171, our model shows more accurate red/green coloring, depth, and fiber detail.

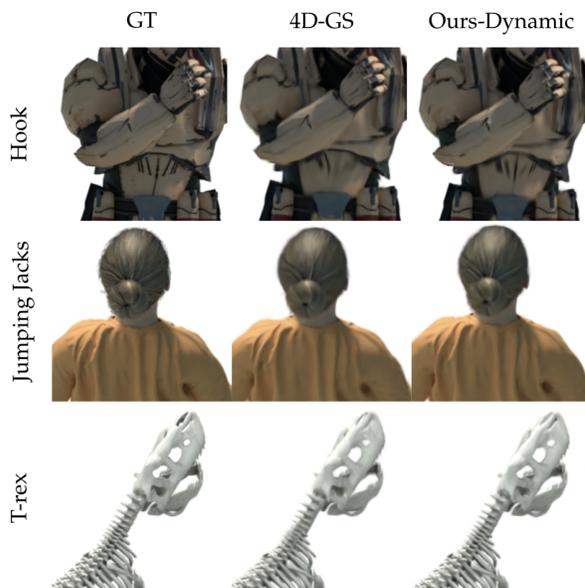


Figure 7. Frame comparisons for the D-NeRF [32] HOOK, JUMPING JACKS, and T-REX datasets across ground truth (GT), baseline 4D-GS, and CTRL-GS with dynamic threshold-based windows. CTRL-GS more accurately reconstructs fine details, like armor lines, hair texture, and T-rex teeth.

## Optimization Details

We mostly follow the same settings as used in 4D-GS [45]. We observe 30,000 iterations for each scene and report best results for our models against baselines.

To test each temporal window construction method, we perform a hyperparameter sweep over number of windows  $N \in \{2, 3, 4, 5, 6, 7, 8, 9\}$  for all methods, and quantization coefficient  $q \in \{0.0, 0.1, 0.2, 0.3, 0.5, 0.7, 0.9\}$  for the equal-length method, reporting best results. This set of values for  $q$  was selected to test a range of interpolation weights between segment and residual predictions. We did not choose to test higher numbers of  $N$ , because we tended to note performance degradation at the upper range of this set (which makes sense, because for short videos, too many temporal segments can lead to the consequences of temporal discontinuity outweighing gains from temporal fitting).