

# TurboGS: Accelerating 3D Gaussian Splatting via Error-Guided Sparse Pixel Sampling and Optimization

Zheng Dong<sup>†</sup><sup>1</sup> Daifei Qiu<sup>2</sup> Pinxuan Dai<sup>3</sup> Ke Xu<sup>4</sup> Jiamin Xu<sup>5</sup> Lili He<sup>1</sup> Rynson W.H. Lau<sup>4</sup> Weiwei Xu<sup>†</sup><sup>3</sup>

## Abstract

Consumer-level applications require fast optimization of 3D Gaussian Splatting (3DGS) with high-fidelity novel view rendering. However, existing 3DGS acceleration approaches still incur substantial computation on redundant pixels while sacrificing fine details. In this paper, we present TurboGS, an error-guided training framework that accelerates 3DGS by concentrating optimization on perceptually informative pixels. TurboGS is built upon four core components: (1) a tile-wise sparse pixel sampling, which, driven by multi-view reconstruction errors during training, prioritizes challenging regions and skips well-reconstructed ones to avoid redundant gradient computation; (2) a tile-wise structure-aware loss with sparse Normalized Cross-Correlation, which provides sparse yet effective supervision to preserve fine details and stabilize training; (3) an error-driven Gaussian density control strategy, which dynamically allocates model capacity and removes redundant primitives; and (4) a tailored hybrid optimizer that couples Hessian-informed updates with Adam moment damping to stabilize and improve convergence under sparse supervision. Experiments on standard benchmarks demonstrate that TurboGS can deliver on par or superior rendering quality within 100 seconds (up to  $\sim 10\times$  training speedup over vanilla 3DGS).

## 1. Introduction

Novel view synthesis (NVS) is a long-standing and challenging research topic with downstream applications such as VR/AR. In recent years, Neural Radiance Fields (NeRF) (Mildenhall et al., 2020) and their variants demon-

<sup>1</sup>Zhejiang Sci-Tech University <sup>2</sup>Malanshan Audio&Video Laboratory <sup>3</sup>CAD&CG Laboratory, Zhejiang University <sup>4</sup>City University of Hong Kong <sup>5</sup>Hangzhou Dianzi University. Correspondence to: Zheng Dong<sup>†</sup>, Weiwei Xu<sup>†</sup> <zhengdong@zstu.edu.cn, xww@cad.zju.edu.cn>.

Proceedings of the 43<sup>rd</sup> International Conference on Machine Learning, Seoul, South Korea. PMLR 306, 2026. Copyright 2026 by the author(s).

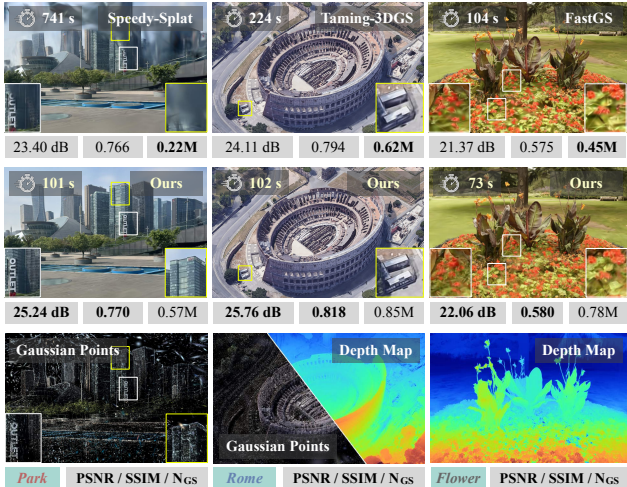


Figure 1. We propose **TurboGS**, a novel framework that *accelerates 3DGS optimization* while delivering high-quality novel view rendering results with notable details across scene scales.

strate impressive rendering fidelity, yet require hours to days of per-scene optimization. Despite substantial progress on acceleration (Müller et al., 2022; Hedman et al., 2021), NeRF-style approaches still struggle to simultaneously balance optimization time, memory footprint, and rendering quality of high-frequency details. On the other hand, by representing scenes with explicit Gaussian primitives, 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) offers fast optimization and rendering that achieves NeRF-comparable quality. Nonetheless, to satisfy consumer-facing or time-sensitive real-world applications, accelerating 3DGS (e.g., with tight computation budgets or under large-scale/fast-turnaround settings (Liu et al., 2024; Zhu et al., 2026)) remains highly desirable.

Existing 3DGS acceleration methods largely fall into two categories. A group of works accelerates the rendering and backpropagation procedures via improved rasterization (Hanson et al., 2025a; Feng et al., 2025) or near-second-order optimizers (Höllein et al., 2025; Lan et al., 2025). The other group of methods focuses on controlling the Gaussian primitive complexity (i.e., adaptive density control) by scheduling the primitive growth or pruning redundant Gaussians (Mallick et al., 2024; Chen et al., 2025; Fang & Wang, 2024; Hanson et al., 2025b; Papantonakis et al., 2024b; Wang et al., 2024; Baranowski et al., 2026). De-

spite their success, as shown in Fig. 1 (first row), existing approaches typically suffer from (1) substantial redundant computation in well-reconstructed regions due to their uniform optimization of pixels: rasterizing and backpropagating dense pixels/tiles per iteration during training; and (2) loss of fine details due to their aggressive primitive pruning.

To address these problems, we draw inspiration from the sparse ray sampling for optimizing NeRF (Mildenhall et al., 2020), where focusing on a set of informative rays improves efficiency. Analogously, we observe that during 3DGS optimization, residual errors are typically distributed across a small subset of difficult pixels (e.g., edges, thin structures, and other high-frequency regions), while a large portion of pixels quickly becomes low-error and contributes marginal gradients. Repeated rasterization in such well-reconstructed regions may hinder convergence under a fixed time budget. To this end, we ask: *can we allocate optimization capacity adaptively according to the reconstruction difficulty, instead of uniformly over the whole image plane?*

In this paper, we propose **TurboGS**, a novel error-guided training framework that accelerates 3DGS by focusing optimization on perceptually informative pixels. We first propose tile-wise sparse pixel sampling (guided by online-maintained multi-view error maps) to prioritize per-iteration rasterization and gradient backpropagation on challenging regions. Based on the per-tile sparsely sampled pixels, we introduce a local structure-aware loss that stabilizes the optimization while preserving fine details, by measuring the sparse Normalized Cross-Correlation (NCC). We also propose an error-driven density control mechanism for adaptive Gaussian densification and pruning, which improves both efficiency and rendering quality by accumulating and leveraging per-Gaussian statistics corresponding to the sparse pixel errors. Last, we optimize TurboGS with a new, hybrid optimizer that blends Hessian-informed updates with Adam moment damping, to achieve stable and fast convergence.

Extensive experiments on standard benchmarks show that TurboGS delivers high-quality rendering results with fine details (see Fig. 1 (second row)) and converges fast (see Fig. 2). In sum, we present the following main contributions:

- TurboGS, a novel error-guided training framework for 3DGS that accelerates optimization by dynamically focusing computation on challenging pixels.
- A tile-wise structure-aware loss based on sparse Normalized Cross-Correlation (NCC) to provide effective supervision from sparse samples, preserving fine details and local texture consistency.
- An error-driven density control strategy that enables targeted densification of challenging regions and pruning of redundant Gaussians.
- A tailored hybrid optimizer that integrates Adam moments with Hessian-informed updates, to enable faster and more

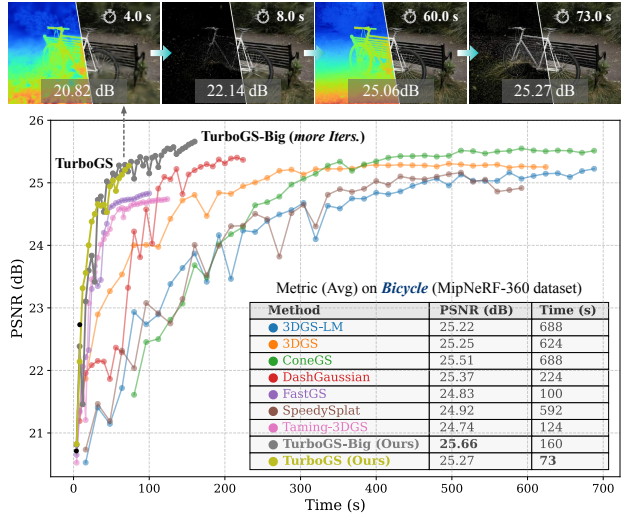


Figure 2. We visualize our intermediate test view rendering results (upper part) and training time-PSNR trajectories of existing methods and ours on the *Bicycle* scene (Yu et al., 2024) (lower part). Our method converges faster to better results.

stable convergence under challenging conditions of sparse, non-uniform gradients.

## 2. Related Work

**Novel View Synthesis (NVS)** attracts significant research interest in computer vision and graphics. Neural Radiance Fields (NeRF) (Mildenhall et al., 2020) significantly facilitates NVS by implicitly parameterizing scene information via an MLP network, and achieving high-quality image synthesis through volumetric rendering. While NeRFs suffer from training overheads, follow-up works combine NeRF with explicit or hybrid representations, such as voxels (Fridovich-Keil et al., 2022; Sun et al., 2022), hash grids (Müller et al., 2022), point-based formulations (Xu et al., 2022), to improve the training or rendering efficiency.

Alternatively, 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) represents scenes with explicit Gaussian primitives and renders images through GPU-friendly rasterization and  $\alpha$ -compositing, enabling much faster optimization and real-time rendering. Building upon 3DGS, prior works improve rendering quality (Hamdi et al., 2024; Lu et al., 2024; Yu et al., 2024), enhance surface reconstruction (Guédon & Lepetit, 2024; Huang et al., 2024), scale to large scenes or high-resolution rendering (Kerbl et al., 2024; Feng et al., 2025), and reduce reliance on SfM (Snively et al., 2006) initialization (Ji & Yao, 2025; Pan et al., 2025; Fan et al., 2024b). Our TurboGS is inspired by NeRF-style sparse ray sampling, where we introduce a tailored principle to 3DGS by focusing supervision on informative pixels.

**Efficient 3DGS.** Recent methods aim to pursue higher efficiency in both training (e.g., density control in Taming-3DGS (Mallick et al., 2024)) and rendering (e.g., LoD

structures in Octree-GS (Ren et al., 2025)), while reducing the storage overhead induced by large numbers of primitives (Papantonakis et al., 2024a). In contrast, our TurboGS allocates the optimization budget over pixels based on reconstruction difficulty and training history.

*Primitive Management (Densification and Pruning).* A key bottleneck of 3DGS (Kerbl et al., 2023) is the rapid growth of Gaussian primitive count, yielding redundancy and slow optimization. Recent works control density by refining this schedule (Mallick et al., 2024; Hanson et al., 2025a; Chen et al., 2025; Kheradmand et al., 2024; Baranowski et al., 2026; Ren et al., 2026), developing stronger pruning criteria (Fan et al., 2024a; Rota Bulò et al., 2024; Fang & Wang, 2024), and combining compression with pruning for efficiency (Girish et al., 2024). For instance, Dash-Gaussian (Chen et al., 2025) schedules rendering resolution with primitive growth, FastGS (Ren et al., 2026) guides densification and pruning with multi-view reconstruction consistency, HGS (Xu et al., 2026) mines hard Gaussians from multi-view positional gradients, 3DGS-MCMC (Kheradmand et al., 2024) formulates MCMC with SGLD-style updates (Brosse et al., 2018; Shakiba Kheradmand, 2024), while Mini-GS (Fang & Wang, 2024) combines depth and blurriness cues with splitting and importance-aware pruning. In TurboGS, we back-project sparse multi-view pixel errors onto Gaussians during rasterization, where the aggregated error statistics guide targeted densification and pruning.

*Optimizer and Rasterization Acceleration.* Another group of recent works improves the efficiency of optimization and rasterization. 3DGS-LM (Hölllein et al., 2025) introduces a two-stage Adam-to-LM optimization, while stochastic Newton-style prioritized per-kernel updates are used in 3DGS<sup>2</sup> (Lan et al., 2025) for fast convergence. To accelerate training, per-splat backward is used in Taming-3DGS (Mallick et al., 2024), while StopThePop (Radl et al., 2024), FlashGS (Feng et al., 2025), and Speedy-Splat (Hanson et al., 2025a) focus on Gaussian-tile pairs with accurate tile intersection. In TurboGS, we design an error-guided sparse pixel supervision with a lightweight hybrid solver, which enables more efficient optimization and better time-quality tradeoffs.

## 3. Our Proposed Method : TurboGS

### 3.1. Preliminary

**3D Gaussian Splatting (3DGS).** Given calibrated multi-view images  $\{\mathbf{I}_v\}_{v=1}^N$  with corresponding cameras  $\{\mathbf{C}_v\}$ , and the point cloud derived from the structure-from-motion (SfM) (Snavely et al., 2006) algorithm, 3DGS represents a scene as a set of  $M$  anisotropic Gaussian primitives  $\mathcal{G} = \{g_i\}_{i=1}^M$ . Each primitive  $g_i$  is parameterized by geometry and appearance: a 3D center position  $\boldsymbol{\mu}_i \in \mathbb{R}^3$ , a 3D covariance  $\boldsymbol{\Sigma}_i \in \mathbb{R}^{3 \times 3}$  (typically via scale  $s_i$  and rotation  $\mathbf{R}_i$ ), an opacity  $\alpha_i \in (0, 1)$ , and color coefficients  $\mathbf{c}_i$ .

For a pixel  $\mathbf{p}$  in view  $v$ , 3DGS projects Gaussians to the image plane and rasterizes them as 2D splats. Let  $\mathcal{K}(\mathbf{p})$  be the ordered set of Gaussians intersecting pixel  $\mathbf{p}$  (sorted by depth). The rendered color follows alpha compositing:

$$\hat{\mathbf{I}}_v(\mathbf{p}) = \sum_{k \in \mathcal{K}(\mathbf{p})} T_k(\mathbf{p}) \alpha_k(\mathbf{p}) \mathbf{c}_k(\mathbf{p}), \quad T_k(\mathbf{p}) = \prod_{j < k} (1 - \alpha_j(\mathbf{p})), \quad (1)$$

where  $\alpha_k(\mathbf{p})$  depends on the projected 2D Gaussian footprint and its opacity. This pipeline is differentiable, enabling gradient-based optimization of  $\mathcal{G}$ .

**Sparse-Pixel Training Setting.** Different from vanilla 3DGS (Kerbl et al., 2023) that rasterizes a dense image (or full tiles) for every selected view at each iteration, our TurboGS optimizes Gaussian primitives  $\mathcal{G}$  using a sparse set of pixels to reduce redundant computation.

At iteration  $t$ , we first sample a small set of views  $\mathcal{V}_t$  (Sec. 3.4). For each view  $v \in \mathcal{V}_t$ , we then sample a tile-wise sparse pixel set  $\mathcal{P}_{t,v}$  (Sec. 3.5), where the sampled pixels are ordered by their tile indices to match the underlying rasterization schedule. TurboGS executes the standard 3DGS training loop only on these sampled pixels, including sparse forward rasterization, loss evaluation, and backpropagation.

To make sparse-pixel training efficient in practice, we introduce modifications to the original rasterization kernels via *tile-wise pixel mapping* and the *per-Gaussian backward mechanism* (Mallick et al., 2024), which restricts computation to the sampled pixel set, *i.e.*,  $\{(v, \mathbf{p}) \mid v \in \mathcal{V}_t, \mathbf{p} \in \mathcal{P}_{t,v}\}$ , while preserving tile-parallel forward execution and efficient splat-parallel computation in the backward pass.

### 3.2. Overview of TurboGS

Motivated by the observation that, as training progresses, only a small subset of pixels remains difficult to reconstruct while most become well explained and yield marginal gradients (see the decreasing pixel errors in Fig. 5), TurboGS maintains per-pixel statistics to allocate computation to informative pixels throughout training. As shown in Fig. 3 and Alg. 1, each iteration performs (1) *view selection* and (2) *tile-wise sparse pixel sampling* guided by pixel error and age, followed by (3) *sparse pixel rasterization*, (4) *tile-level sparse supervisions*, (5) *online EMA updates of error and age maps*, (6) *error-driven Gaussian density control*, and (7) *a lightweight hybrid solver* for stable sparse training.

### 3.3. Persistent Pixel-wise Informative Map

To support error-guided pixel sampling throughout training, TurboGS maintains *pixel-wise maps* for each training view. For each view  $v$ , we store (i) an **error map**  $\mathbf{E}_v \in \mathbb{R}^{H \times W}$  that records an online estimate of per-pixel reconstruction difficulty, and (ii) an **age map**  $\mathbf{A}_v \in \mathbb{N}^{H \times W}$  that measures how long a pixel has not been sampled. The error map

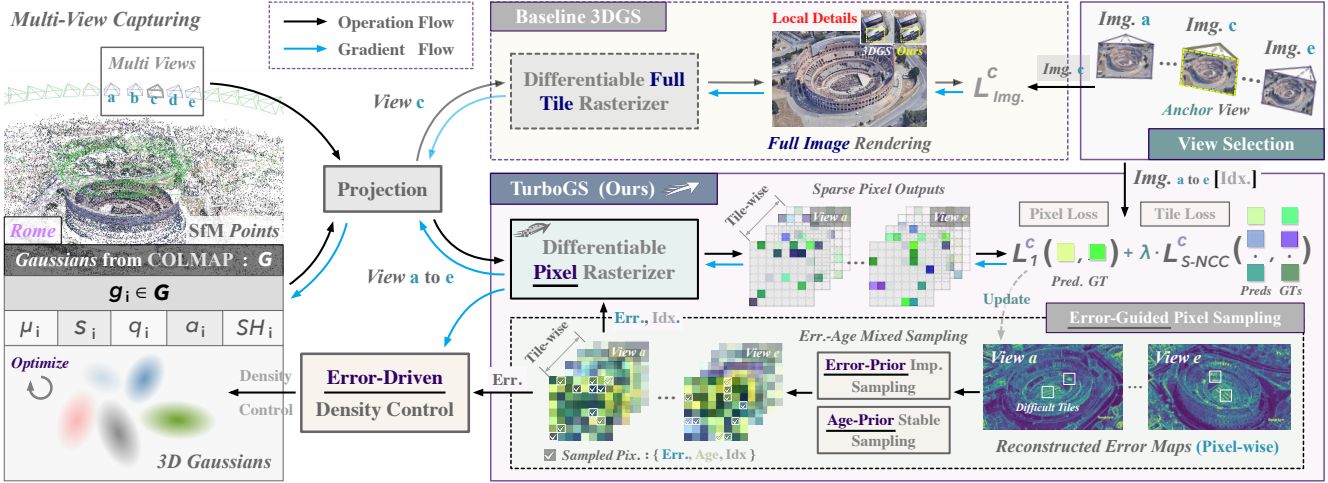


Figure 3. **TurboGS framework overview.** Unlike existing 3DGS (Kerbl et al., 2023; Mallick et al., 2024; Ren et al., 2026) methods that rasterize the entire tile in each iteration, TurboGS allocates more attention to pixels that are difficult to reconstruct and reduces gradient backpropagation in well-interpreted regions, thereby lowering the rasterization cost per iteration. To achieve this, TurboGS records multi-view pixel reconstruction errors and performs tile-wise sparse pixel sampling based on pixel error and age. Built upon this sparse sampling scheme, we further incorporate a local structure-aware loss and error-driven density control to preserve fine details.

promotes exploitation of difficult regions, while the age map encourages periodic revisiting of neglected pixels to refresh outdated estimates and to avoid over-focusing on a small set of persistently high-error locations.

**Sparse Pixel Loss as Error Signal.** For sampled pixels, we compute a per-pixel reconstruction error, as:

$$e_{t,v}(\mathbf{p}) = \|\hat{\mathbf{I}}_{t,v}(\mathbf{p}) - \mathbf{I}_v(\mathbf{p})\|_1, \quad \mathbf{p} \in \mathcal{P}_{t,v}, \quad (2)$$

and we initialize error map  $\mathbf{E}_v$  by densely evaluating  $e_{0,v}(\mathbf{p})$  for each pixel rendered by initial Gaussian  $\mathcal{G}_0$ .

**EMA Update of Pixel Errors.** We update  $\mathbf{E}_v$  only at sampled locations using an exponential moving average:

$$\mathbf{E}_v(\mathbf{p}) \leftarrow (1 - \beta) \mathbf{E}_v(\mathbf{p}) + \beta e_{t,v}(\mathbf{p}), \quad \mathbf{p} \in \mathcal{P}_{t,v}, \quad (3)$$

and keep other pixels unchanged. This avoids re-estimating errors over all pixels at every iteration. Fig. 5 illustrates the error map  $\mathbf{E}_v$  update process during training.

**Pixel Age Update.** We increment the age of all pixels globally and reset the age of sampled pixels:

$$\mathbf{A}_v(\mathbf{p}) \leftarrow \begin{cases} 0, & \mathbf{p} \in \mathcal{P}_{t,v}, \\ \mathbf{A}_v(\mathbf{p}) + 1, & \text{otherwise.} \end{cases} \quad (4)$$

In practice, we store  $\mathbf{E}_v$  and  $\mathbf{A}_v$  in contiguous GPU memories, and use half precision for  $\mathbf{E}_v$  and uint16 for  $\mathbf{A}_v$ .

### 3.4. Geometry-Aware View Sampling

At iteration  $t$ , we sample  $K$  views to form a view set  $\mathcal{V}_t$ . As shown in Fig. 3 and Fig. 4, we first choose an *anchor view* according to the average error over its full pixels,

$$\bar{E}_v = \frac{1}{|\Omega|} \sum_{\mathbf{p} \in \Omega} \mathbf{E}_v(\mathbf{p}), \quad v^* = \arg \max_v \bar{E}_v, \quad (5)$$

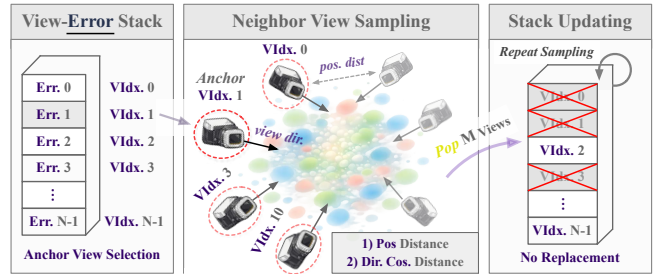


Figure 4. **View selection in TurboGS.** An anchor view is chosen based on view errors, followed by geometry-aware sampling of  $K-1$  nearby views from a candidate stack. Selected views are then removed from the stack to avoid repeated sampling.

where  $\Omega$  denotes the image domain. We then sample the remaining  $K-1$  views without repetition according to a geometry-aware distribution that favors nearby cameras with similar viewing directions to the selected anchor view:

$$\pi(u | v^*) \propto \exp\left(-\lambda_d \|\mathbf{c}_u - \mathbf{c}_{v^*}\|_2 - \lambda_\theta (1 - \langle \mathbf{d}_u, \mathbf{d}_{v^*} \rangle)\right), \quad (6)$$

where  $\mathbf{c}_u$  and  $\mathbf{d}_u$  denote the camera center and unit viewing direction of the view  $u$ . In practice, we maintain a shuffled view stack and draw the  $K$  views *without replacement* by popping sampled indices from the stack, which will be refilled and reshuffled once fewer than  $K$  views remain.

### 3.5. Tile-Wise Error-Guided Sparse Pixel Sampling

We partition each image into several tiles of size  $S \times S$  (e.g.,  $16 \times 16$ ). For each selected view  $v$  at iteration  $t$ , TurboGS samples a fixed number, i.e.,  $N_{\text{pix}} = \max(1, \lfloor r_t \cdot S^2 \rfloor)$  of pixels *within each tile* to form  $\mathcal{P}_{t,v}$ , where  $r_t$  denotes the pixel sampling rate. We further split  $N_{\text{pix}}$  into a hard set




Figure 5. **Visualization of error maps during training.** Pixel-wise error maps and corresponding renderings over training (upper: *Amsterdam*; lower: *Street*). As optimization proceeds, pixel errors gradually weaken as reconstruction quality improves, and the error maps flatten toward predominantly low-frequency residuals.

and a stable set, with  $N_{\text{hard}} = \lfloor r_{\text{hard}} \cdot N_{\text{pix}} \rfloor$ , and  $N_{\text{stable}} = \lfloor r_{\text{stable}} \cdot N_{\text{pix}} \rfloor$ , satisfying  $N_{\text{hard}} + N_{\text{stable}} \leq N_{\text{pix}}$ .

**Error-Prior Importance Sampling.** For a tile  $\tau$ , we define an error-based sampling distribution over its pixels:

$$\pi_E(\mathbf{p} \mid \tau, v) = \frac{\mathbf{E}_v(\mathbf{p})}{\sum_{\mathbf{q} \in \tau} \mathbf{E}_v(\mathbf{q}) + \epsilon}, \quad \mathbf{p} \in \tau, \quad (7)$$

where  $\mathbf{E}_v(\mathbf{p})$  is the current error-map value (Sec. 3.3). We then draw  $N_{\text{hard}}$  hard pixels from the tile  $\tau$  by multinomial sampling based on  $\pi_E(\cdot \mid \tau, v)$  without replacement.

**Age-Prior Stable Sampling.** Let  $\mathcal{H}_{t,v}(\tau)$  be the selected hard pixels in tile  $\tau$ . We draw age-prior stable pixels from the remaining pixel set,  $\tau \setminus \mathcal{H}_{t,v}(\tau)$ , as:

$$\mathcal{S}_{t,v}(\tau) = \text{TopK}_{\mathbf{p} \in \tau \setminus \mathcal{H}_{t,v}(\tau)}(\mathbf{A}_v(\mathbf{p}), N_{\text{stable}}), \quad (8)$$

where  $\mathbf{A}_v(\mathbf{p})$  is the age map and  $\text{TopK}(\cdot, N_{\text{stable}})$  returns the  $N_{\text{stable}}$  pixels with the largest ages.

Finally, the sampled pixels in tile  $\tau$  are  $\mathcal{P}_{t,v}(\tau) = \mathcal{H}_{t,v}(\tau) \cup \mathcal{S}_{t,v}(\tau)$ , and the full sampled set is  $\mathcal{P}_{t,v} = \bigcup_{\tau} \mathcal{P}_{t,v}(\tau)$ .

### 3.6. Differentiable Sparse Pixel Rasterization

We output sampled pixels in tile order for GPU tile-parallel execution and perform differentiable rasterization on  $\mathcal{P}_{t,v}$ , to obtain rendered colors  $\hat{\mathbf{I}}_{t,v}(\mathbf{p})$  and Gaussian gradients.

As illustrated in Fig. 6, the forward pass maintains valid sampled tiles and computes tile-wise pixel offsets via prefix-sum, enabling each CUDA block (tile) to efficiently access

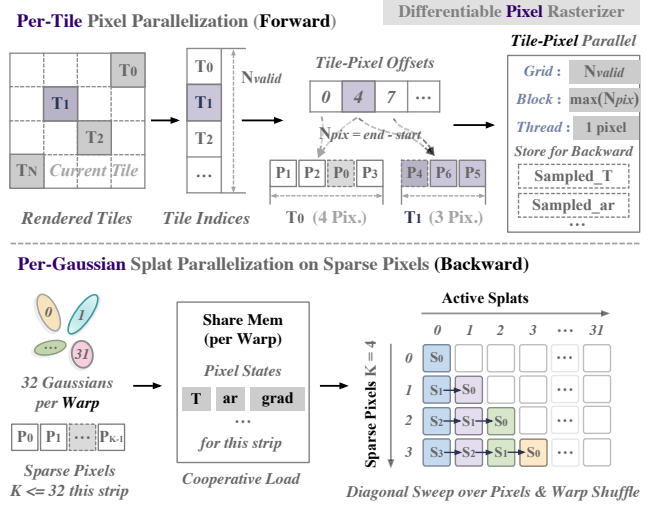


Figure 6. **Pixel rasterization in TurboGS.** Forward rasterization follows tile-parallel execution by organizing sampled pixels in tile order, while the backward pass adopts per-Gaussian splat parallelization to propagate gradients only over sparse sampled pixels.

sampled pixel indices and rasterize only sparse pixels. For the backward pass, following the splat-parallel strategy of Taming-3DGS (Mallick et al., 2024), we retrieve sparse pixels within each tile through tile-pixel indexing and propagate gradients via warp shuffle only over sampled pixels, avoiding dense rasterization replay while preserving efficient gradient accumulation.

### 3.7. Sparse Supervision with Tile-Level NCC

Sparse supervision can be unreliable when relying only on pixel-wise losses (only  $\ell_1$  loss in Tab. 4 and w/o S-NCC in Fig. 11). To better preserve local structures, we introduce a tile-level sparse NCC regularization on sampled pixels.

For a tile  $\tau$  in view  $v$ , we compute a sparse Normalized Cross-Correlation (NCC) similarity over the sampled pixels  $\mathcal{P}_{t,v}(\tau)$  and define the tile-wise structure loss:

$$\mathcal{L}_{\text{S-NCC}}^\tau = 1 - \text{NCC}(\hat{\mathbf{I}}_{t,v}(\mathcal{P}_{t,v}(\tau)), \mathbf{I}_v(\mathcal{P}_{t,v}(\tau))). \quad (9)$$

**Gradient-Balanced Weighting.** Instead of using a fixed NCC weight, we balance the loss by the ratio of the average gradient norms of the  $\ell_1$  and NCC terms within each tile:

$$\lambda_{\text{ncc}}^\tau = \text{clip}\left(\frac{\bar{g}_{\ell_1}^\tau}{\bar{g}_{\text{S-NCC}}^\tau + \epsilon}, 0, 1\right),$$

$$\bar{g}_k^\tau = \mathbb{E}_{\mathbf{p} \in \mathcal{P}_{t,v}(\tau)} \left[ \|\nabla_{\hat{\mathbf{I}}(\mathbf{p})} \mathcal{L}_k^\tau(\mathbf{p})\|_1 \right], \quad k \in \{\ell_1, \text{S-NCC}\}. \quad (10)$$

In practice, this is computed in a CUDA kernel, where each tile corresponds to one thread block, and  $\bar{g}_k^\tau$  is obtained via block-level gradient reduction of pixel gradient norm.

**Sparse Pixel Supervision.** Our sparse pixel supervision is

Table 1. **Quantitative comparisons with existing improved 3DGS optimization methods.** TurboGS completes training **within 80 seconds** while achieving competitive rendering quality compared with prior methods. Best, second-best, and third-best results are highlighted by **best score**, **second best score**, and **third best score**, respectively. Time is reported in seconds.

Method	<i>Mip-NeRF 360 (Barron et al., 2022)</i>						<i>Tanks &amp; Temples (Knapitsch et al., 2017)</i>						<i>Deep Blending (Hedman et al., 2018)</i>					
	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑
3DGS	891	27.55	0.816	0.215	2.73M	199	583	23.84	<b>0.853</b>	0.169	1.58M	252	921	29.85	0.907	<b>0.238</b>	2.48M	208
Taming-3DGS	148	27.26	0.795	0.259	0.67M	413	97	23.77	0.836	0.211	0.32M	612	100	29.92	0.903	0.271	0.29M	601
Mini-Splatting	674	27.32	<b>0.822</b>	0.217	0.49M	279	479	23.24	0.836	0.202	<b>0.20M</b>	405	586	29.95	<b>0.907</b>	0.254	0.35M	364
Speedy-Splat	533	26.95	0.786	0.288	<b>0.32M</b>	<b>983</b>	292	23.41	0.820	0.239	<b>0.19M</b>	<b>1138</b>	497	29.55	0.903	0.268	<b>0.25M</b>	<b>1121</b>
3DGS-LM	622	27.41	0.814	0.220	3.35M	218	369	23.57	<b>0.844</b>	0.185	1.80M	285	572	29.58	<b>0.906</b>	<b>0.246</b>	2.89M	218
DashGaussian	169	<b>27.69</b>	<b>0.817</b>	0.220	2.24M	259	141	24.04	<b>0.851</b>	0.181	1.20M	334	112	<b>30.12</b>	<b>0.907</b>	<b>0.248</b>	1.95M	289
FastGS	111	27.51	0.805	0.261	<b>0.38M</b>	<b>1002</b>	91	<b>24.06</b>	0.842	0.209	<b>0.24M</b>	<b>1087</b>	81	30.06	0.905	0.267	<b>0.22M</b>	<b>1144</b>
ConeGS	778	<b>27.74</b>	<b>0.823</b>	<b>0.202</b>	0.87M	427	749	23.88	<b>0.853</b>	<b>0.160</b>	0.55M	564	805	<b>30.28</b>	<b>0.909</b>	<b>0.238</b>	<b>0.52M</b>	<b>655</b>
<b>TurboGS (Ours)</b>	<b>77</b>	27.57	0.794	0.256	0.64M	153	<b>73</b>	23.79	0.832	0.200	0.58M	170	<b>62</b>	<b>30.00</b>	0.900	0.277	0.49M	231
<b>TurboGS-Big (Ours)</b>	168	<b>27.99</b>	0.815	0.224	1.34M	134	146	<b>24.14</b>	0.841	<b>0.181</b>	0.91M	151	132	<b>30.26</b>	<b>0.909</b>	0.259	0.69M	216

formulated only by the sampled pixels for each tile:

$$\mathcal{L} = \mathbb{E}_{v \in \mathcal{V}_t, \mathbf{p} \in \mathcal{P}_{t,v}} [\ell_1(\mathbf{p})] + \mathbb{E}_{v \in \mathcal{V}_t, \tau} [\lambda_{\text{ncc}}^\tau \mathcal{L}_{\text{S-NCC}}^\tau]. \quad (11)$$

In practice, loss evaluation and per-tile gradient statistics are efficiently fused in an independent CUDA kernel for fast execution with low overhead (see runtime in Fig. 7).

### 3.8. Error-Driven Density Control

Besides sparse supervision, TurboGS further adapts the Gaussian density using multi-view errors  $e_{t,v}(\mathbf{p})$  to guide densification and pruning. When a Gaussian  $g_i$  contributes to the pixel  $\mathbf{p}$  with compositing weight  $w_{t,i}(\mathbf{p}) = T_{t,i}(\mathbf{p}) \alpha_{t,i}(\mathbf{p})$ , we accumulate per-Gaussian statistics:

$$\mathcal{E}_i^{(t)} = \sum_{v, \mathbf{p}} e_{t,v}(\mathbf{p}) w_{t,i}(\mathbf{p}), \quad \mathcal{W}_i^{(t)} = \sum_{v, \mathbf{p}} w_{t,i}(\mathbf{p}). \quad (12)$$

Similarly, we accumulate a distance term  $\mathcal{D}_i^{(t)}$  by replacing  $e_{t,v}(\mathbf{p})$  with a pixel-splat distance term  $D_{t,i}(\mathbf{p})$  (i.e., the conic Mahalanobis distance computed from the offset between 2D Gaussian projection and the pixel coordinate).

We normalize the accumulated statistics by visibility:  $e_i^{(t)} = \mathcal{E}_i^{(t)} / (\mathcal{W}_i^{(t)} + \epsilon)$  and  $d_i^{(t)} = \mathcal{D}_i^{(t)} / (\mathcal{W}_i^{(t)} + \epsilon)$ , and define the Gaussian difficulty score as:

$$s_i^{(t)} = e_i^{(t)} + \lambda_{\text{eff}} d_i^{(t)}, \quad \lambda_{\text{eff}} = \lambda_{\text{base}}(1 - \rho_t), \quad (13)$$

where  $\rho_t \in [0, 1)$  denotes the training progress. Normalization by  $\mathcal{W}_i^{(t)}$  removes bias toward frequently visible Gaussians, while the distance term emphasizes boundary regions that often indicate under-fitting.

Here, different from the density control in FastGS (Ren et al., 2026), which relies on post-hoc dense error maps for densification and pruning, TurboGS uses Gaussian scores  $s_i^{(t)}$  and  $e_i^{(t)}$  obtained from online sparse pixel statistics, along with their percentile thresholds for adaptive density control, reducing computational complexity for faster training.




Figure 7. **Per-step runtime breakdown of TurboGS.** We report the average execution time (ms) of each optimization step on the *Bicycle* and large-scale *Rome* scenes. Step abbreviations denote the corresponding operations in Alg. 1 (Appendix).

### 3.9. Moment-Damped LM Solver for Sparse Training

Sparse supervision alters the optimization landscape and can amplify gradient variance across views and tiles. To improve convergence while keeping the update cost lightweight, TurboGS adopts a moment-damped Levenberg-Marquardt (LM) solver implemented as per-Gaussian CUDA kernels.

At each iteration, we form a small local equation for each Gaussian  $g$  and solve it independently. Let  $\mathbf{m}_g$  and  $\mathbf{v}_g$  denote the Adam first and second moments for a parameter block, and let  $\mathbf{J}_g$  be the local Jacobian estimated from the accumulated gradients. We compute an LM-style update:

$$(\mathbf{J}_g^\top \mathbf{J}_g + \text{diag}(\sqrt{\mathbf{v}_g} + \epsilon)) \Delta_g = -\mathbf{m}_g, \quad (14)$$

where the moment  $\mathbf{v}_g$  serves as a diagonal damping to stabilize sparse gradients. Each local formulation is solved via *Cholesky decomposition*, resulting in a fully parallel and lightweight solver tailored for sparse-pixel training.

## 4. Experiments

We evaluate TurboGS through comparisons with representative improved 3DGS optimization methods and ablations of key components. To analyze computational efficiency, we further report the runtime breakdown of each optimization step in Fig. 7. Experimental details (GPU and training settings) are provided in Sec. B in the appendix.

**Datasets and Metrics.** We conduct experiments on three real-world datasets (i.e., Mip-NeRF 360 (Barron et al.,

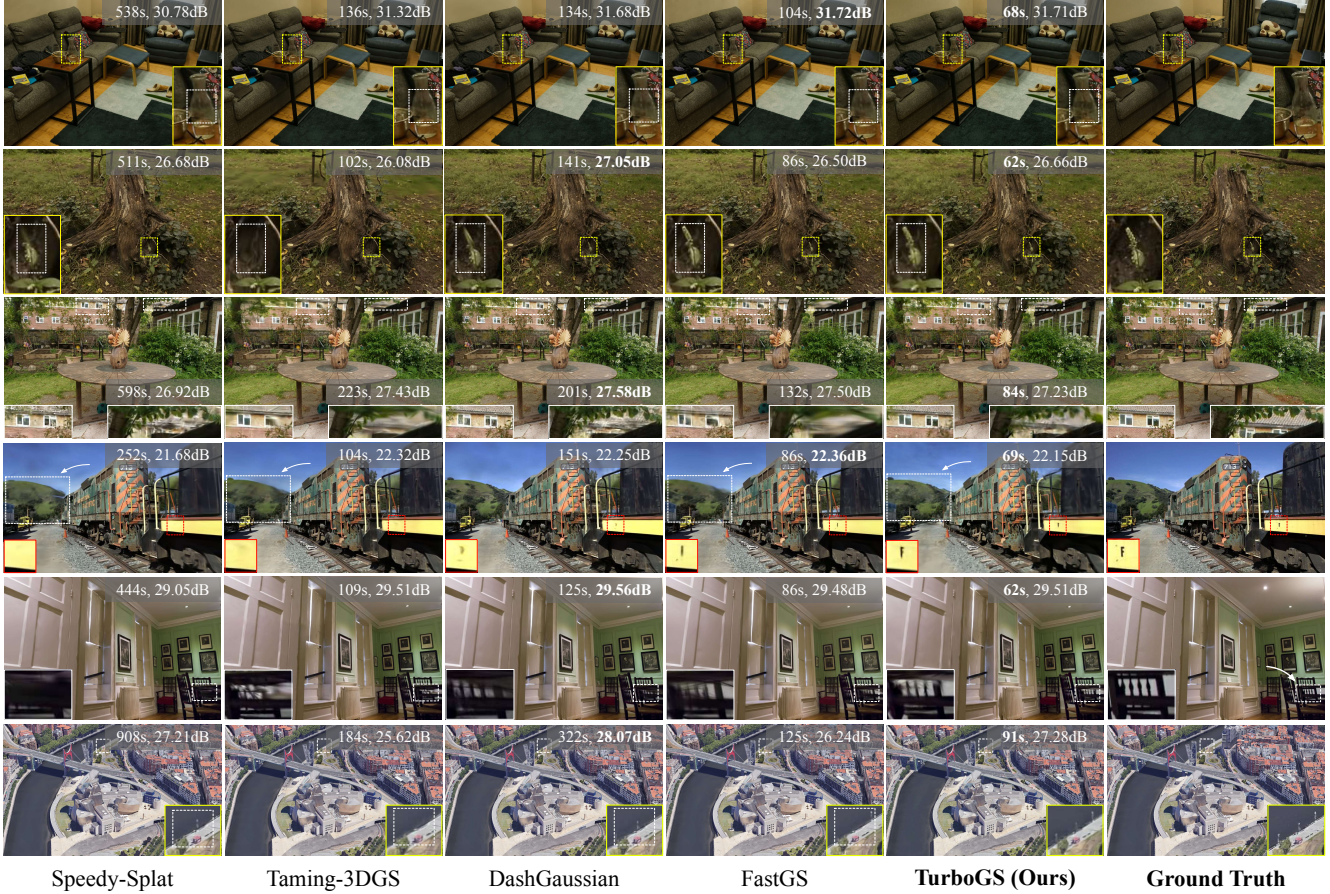


Figure 8. **Qualitative comparisons with fast 3DGS optimization methods.** Results are shown on Mip-NeRF 360 (Barron et al., 2022) (Room, Stump and Garden in the first three rows), Tanks & Temples (Knapitsch et al., 2017) (Train in the 4th row), Deep Blending (Hedman et al., 2018) (Drjohnson in the 5th row), and BungeeNeRF (Xiangli et al., 2022) (Bilbao in the last row). TurboGS better preserves some fine structures and details while achieving the **fastest training speed within 100s** among the compared methods.

Table 2. **Quantitative comparisons with existing fast 3DGS optimization methods.** Time is reported in seconds.

Method	BungeeNeRF (Xiangli et al., 2022)					
	Time↓	PSNR↑	SSIM↑	LPIPS↓	$N_{GS}$ ↓	FPS↑
Taming-3DGS	222	24.54	0.792	0.279	0.62M	243
DashGaussian	332	26.55	0.870	0.172	3.05M	133
FastGS	136	24.97	0.812	0.256	0.61M	700
FastGS-Big	227	25.76	0.853	0.194	1.38M	496
<b>TurboGS (Ours)</b>	<b>101</b>	25.69	0.808	0.252	0.84M	136
<b>TurboGS-Big</b>	230	<b>26.68</b>	<b>0.863</b>	<b>0.180</b>	2.25M	90

2022), Tanks & Temples (Knapitsch et al., 2017), and Deep Blending (Hedman et al., 2018)), one large-scale dataset from Google Earth imagery used in BungeeNeRF (Xiangli et al., 2022), as well as validations on the OMMO (Lu et al., 2023) dataset and the 4K sub-regions of Rubble (Turki et al., 2022) dataset. Novel view quality is evaluated using average PSNR, SSIM (Wang et al., 2004), and LPIPS (Zhang et al., 2018). Training efficiency and compactness are assessed by total training time (Time, in seconds), the final number of Gaussians ( $N_{GS}$ ), and rendering speed (FPS).

#### 4.1. Comparison with 3DGS Optimization Methods

**Baselines.** We compare TurboGS with representative 3DGS optimization methods that target acceleration and primitive control, including Taming-3DGS (Mallick et al., 2024), Speedy-Splat (Hanson et al., 2025a), Mini-Splatting (Fang & Wang, 2024), FastGS (Ren et al., 2026), ConeGS (Baranowski et al., 2026), DashGaussian (Chen et al., 2025), and 3DGS-LM (Höllein et al., 2025), together with vanilla 3DGS (Kerbl et al., 2023) as a reference.

**Quantitative Results.** Tab. 1 and Tab. 2 report quantitative comparisons. TurboGS consistently achieves the fastest training speed across datasets, completing optimization within 80 seconds on three real-world benchmarks, corresponding to a  $10\times \sim 14\times$  speedup over vanilla 3DGS and a  $1.2\times \sim 2\times$  speedup over recent fast methods. Despite the reduced training cost, TurboGS maintains competitive rendering quality, while our TurboGS-Big (with more iterations; see Sec. B) further achieves the best or near-best PSNR on multiple benchmarks. We attribute this to the error-guided sparse-pixel training scheme, which focuses

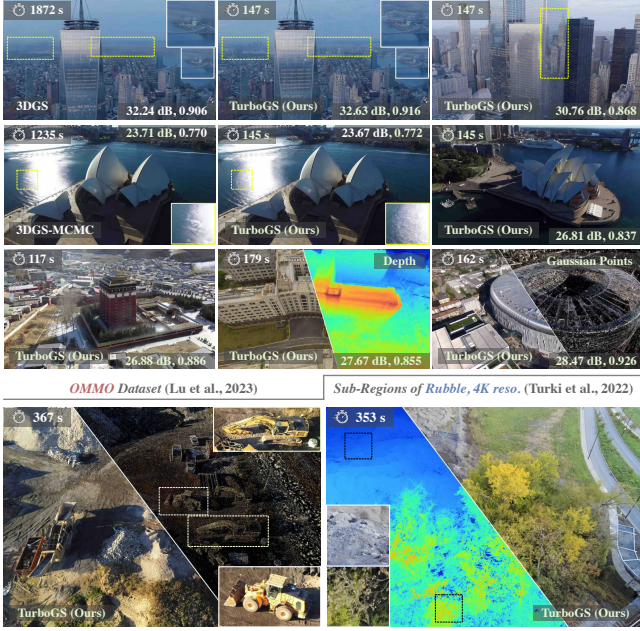


Figure 9. Large-scale and high-resolution scene reconstruction. Results on the large-scale OMMO scene (Lu et al., 2023) and representative 4K Rubble (Turki et al., 2022) sub-regions.

Table 3. Quantitative comparisons on the 4K sub-region Rubble (Turki et al., 2022) dataset. Time is reported in seconds.

Method	4K Sub-Regions of Rubble (Turki et al., 2022)					
	Time↓	PSNR↑	SSIM↑	LPIPS↓	$N_{GS}$ ↓	FPS↑
Taming-3DGS	898	25.45	0.733	0.400	<b>0.94M</b>	51
FastGS	614	26.16	<b>0.788</b>	0.309	2.90M	<b>125</b>
<b>TurboGS (Ours)</b>	<b>360</b>	<b>26.32</b>	<u>0.771</u>	<b>0.284</b>	2.73M	<b>51</b>

computation on difficult pixels. Although SSIM and LPIPS are slightly lower in some cases, likely because our pixel-wise training aligns more closely with PSNR, TurboGS achieves a favorable trade-off between speed, quality, and model compactness. Unlike 3DGS-LM, which requires over 32 GB GPU memory, our hybrid optimizer incurs substantially lower computational and memory overhead.

**Qualitative Results.** Fig. 8 presents qualitative comparisons across datasets. Despite the shorter training time, TurboGS preserves some fine structures and local details comparable to or better than prior fast optimization methods. This mainly benefits from error-guided sparse sampling and density control, which help stabilize training in challenging regions. More results are provided in the appendix.

**Large-Scale and High-Resolution Scenes.** Fig. 9 demonstrates the effectiveness of TurboGS on the large-scale OMMO scenes (Lu et al., 2023), showing both training speedup and improved local details compared with vanilla 3DGS and 3DGS-MCMC (Kheradmand et al., 2024). Fig. 10 and Tab. 3 further validate the advantage of TurboGS for high-resolution (e.g., 4K) training, mainly benefiting from

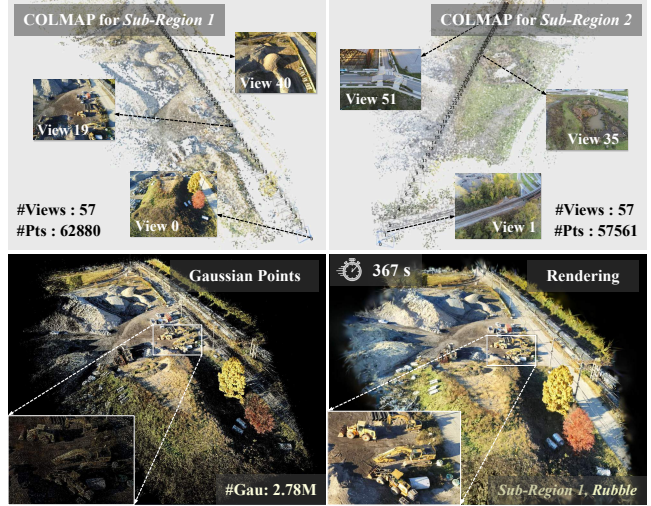


Figure 10. 4K Rubble sub-region reconstruction. Independent COLMAP initialization (cameras and sparse points), and our efficient 3DGS optimization with high-quality rendering results.

Table 4. Ablation studies of TurboGS. Experiments are conducted on Mip-NeRF 360 (Barron et al., 2022) dataset with vanilla 3DGS (Kerbl et al., 2023) as the baseline. We progressively add core components and replace key modules to isolate their effects.

Ablated Method	Time↓	PSNR↑	SSIM↑	LPIPS↓	$N_{GS}$ ↓
3DGS	891	27.55	<b>0.816</b>	<b>0.215</b>	2.73M
+ Err-SPS.	85	26.58	0.737	0.331	0.76M
+ Geo-VS.	89	26.73	0.743	0.321	1.05M
+ Sparse-NCC.	95	27.14	0.773	0.279	0.98M
+ Err-DP. (Full)	77	<b>27.57</b>	0.794	<u>0.256</u>	0.64M
Err-SPS. → RPS.	<b>75</b>	27.10	0.780	0.276	<u>0.56M</u>
Geo-VS. → RVS.	<u>76</u>	27.26	0.780	0.274	<b>0.47M</b>
Err-DP. → FastGS	91	<u>27.55</u>	<u>0.801</u>	0.258	0.60M
$\ell_1$ loss (only)	92	27.09	0.761	0.302	0.76M
S-NCC. → S-SSIM.	86	27.21	0.773	0.287	0.74M

sparse pixel sampling that reduces computation proportionally to image resolution. For the 4K Rubble benchmark, we extract three representative 4K sub-region image sets and independently run COLMAP for fair comparison. For higher-resolution scenes at extreme view scales, storing error/age maps may introduce additional overhead, which could be alleviated by scalable strategies such as multi-GPU partitioning or periodic GPU-CPU memory offloading.

## 4.2. Ablation Study

We conduct ablation studies on the Mip-NeRF 360 (Barron et al., 2022) dataset with vanilla 3DGS (Kerbl et al., 2023) as the baseline to analyze the contribution of each component in TurboGS. In particular, our ablation design jointly considers optimization efficiency and rendering quality, aiming to analyze how different components contribute to a *balanced trade-off between training speed and visual fidelity*. Quantitative results are reported in Tab. 4 and Tab. 5, with qualitative comparisons shown in Fig. 11.




Figure 11. **Qualitative ablation results.** Our full model retains finer details than ablated variants while reducing training time. Results are shown on *Park* (Wu et al., 2023) dataset (first and third rows) and Mip-NeRF 360 dataset (second and 4th rows).




Figure 12. **Limitations of TurboGS.** Examples of subtle rippling artifacts in low-frequency regions despite low error-map responses.

**Effect of Adding Key Components.** As shown in the first five rows of Tab. 4, progressively adding our components substantially reduces training time while improving rendering quality. Error-guided sparse pixel sampling (“Err-SPS.”) provides most of the acceleration, reducing training time from 891 to 85 seconds and  $N_{GS}$  from 2.73M to 0.76M. However, relying only on sparse-pixel  $\ell_1$  supervision degrades reconstruction quality. Geometry-aware view sampling (“Geo-VS.”) and sparse NCC regularization gradually restore rendering quality and enhance local high-frequency details. After introducing error-driven density control (“Err-DP.”), the full TurboGS model achieves the best trade-off, completing training in 77 seconds while improving PSNR by +0.02 dB over the baseline, and reducing  $N_{GS}$  to 0.64M. As shown in the first two rows of Fig. 11, the full model removes distant background artifacts and progressively improves fine structural details with faster convergence.

**Sampling Strategies.** Replacing “Err-SPS.” with random pixel sampling (“RPS.”), and “Geo-VS.” with random view sampling (“RVS.”), both lead to consistent drops in PSNR and perceptual metrics (the sixth and seventh rows in Tab. 4). Fig. 11 (the third row) also shows that random sampling fails to preserve high-frequency details, with “RPS.” causing

Table 5. **Optimizer ablation in TurboGS.** Results on Mip-NeRF 360 (Barron et al., 2022), comparing our Moment-LM optimizer against the sparse Adam-only version under identical settings.

Optimizer	Time↓	PSNR↑	SSIM↑	LPIPS↓	$N_{GS}$ ↓
Sparse Adam (only)	80	27.52	0.793	0.261	<b>0.64M</b>
Moment-LM (Ours)	<b>77</b>	<b>27.57</b>	<b>0.794</b>	<b>0.256</b>	<b>0.64M</b>

more severe degradation since error-dominant pixels are less likely to be prioritized, resulting in less informative supervision. These results highlight the importance of our pixel-and-view sampling under sparse supervision.

**Gaussian Density Control Strategy.** Replacing our density control (“Err-DP.”) with the FastGS-style strategy increases training time and slightly reduces PSNR (8th row in Tab. 4), likely due to its reliance on post-hoc full-image error accumulation, which is less efficient in our online sparse setting.

**Objective Functions.** Using only the  $\ell_1$  loss degrades reconstruction quality. Replacing sparse NCC (“S-NCC.”) with sparse SSIM (“S-SSIM.”) slightly increases training time and sometimes yields less sharp local structures (the 4th row in Fig. 11), suggesting that NCC offers more effective structural supervision under sparse training.

**Optimization Strategy.** As shown in Tab. 5, replacing our moment-damped LM with sparse Adam yields slightly lower PSNR/SSIM and higher LPIPS under similar training time. Qualitatively, the Adam-only variant produces less stable high-frequency details (the 4th row in Fig. 11), suggesting that our hybrid optimizer improves stability and reconstruction quality under sparse supervision.

## 5. Conclusion

In this paper, we have proposed TurboGS, an error-guided sparse optimization framework for accelerating 3DGS. TurboGS adaptively reallocates computation from well-reconstructed regions to informative pixels by incorporating tile-wise sparse pixel sampling with persistent error and age maps, geometry-aware view sampling, structure-aware sparse supervision, error-driven density control, and a hybrid moment-damped optimizer. Extensive experiments demonstrate that TurboGS achieves rapid convergence with competitive rendering quality across diverse scenes.

Despite these advantages, TurboGS still has limitations. As shown in Fig. 12, sparse supervision may occasionally introduce subtle rippling artifacts in low-frequency regions with weak texture variations (e.g., walls or ceilings), even when the error-map response is already low. A possible remedy is to incorporate adaptive dense supervision for tiles in such regions to better regularize structural consistency. Moreover, extending TurboGS to dynamic scenes further requires spatiotemporal sampling and cross-frame error aggregation, which we consider an interesting future work.

## Acknowledgements

We thank all the anonymous reviewers for their professional and constructive comments. Weiwei Xu is partially supported by the National Key Research and Development Program of China (No. 2024YFE0216600), and the NSFC grant (No. 62421003). Zheng Dong is supported by the grant of Zhejiang Provincial Natural Science Foundation, China (No. LQN26F020035).

## Impact Statement

This paper focuses on accelerating 3DGS to reduce training overhead while improving the trade-off between computational efficiency and novel view synthesis quality. Although our work may have broader societal implications, we do not foresee any requiring specific discussion at this stage.

## References

- Baranowski, B., Esposito, S., Gschößmann, P., Chen, A., and Geiger, A. Conegs: Error-guided densification using pixel cones for improved reconstruction with fewer primitives. In *3DV*, 2026.
- Barron, J. T., Mildenhall, B., Verbin, D., Srinivasan, P. P., and Hedman, P. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *CVPR*, 2022.
- Brosse, N., Durmus, A., and Moulines, E. The promises and pitfalls of stochastic gradient langevin dynamics. In *NeurIPS*, 2018.
- Chen, Y., Jiang, J., Jiang, K., Tang, X., Li, Z., Liu, X., and Nie, Y. Dashgaussian: Optimizing 3d gaussian splatting in 200 seconds. In *CVPR*, 2025.
- Fan, Z., Wang, K., Wen, K., Zhu, Z., Xu, D., and Wang, Z. Lightgaussian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps. In *NeurIPS*, 2024a.
- Fan, Z., Wen, K., Cong, W., Wang, K., Zhang, J., Ding, X., Xu, D., Ivanovic, B., Pavone, M., Pavlakos, G., Wang, Z., and Wang, Y. Instantsplat: Sparse-view gaussian splatting in seconds. *arXiv:2403.20309*, 2024b.
- Fang, G. and Wang, B. Mini-splatting: Representing scenes with a constrained number of gaussians. In *ECCV*, 2024.
- Feng, G., Chen, S., Fu, R., Liao, Z., Wang, Y., Liu, T., Pei, Z., Li, H., Zhang, X., and Dai, B. Flashgs: Efficient 3d gaussian splatting for large-scale and high-resolution rendering. In *CVPR*, 2025.
- Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., and Kanazawa, A. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022.
- Girish, S., Gupta, K., and Shrivastava, A. Eagles: Efficient accelerated 3d gaussians with lightweight encodings. In *ECCV*, 2024.
- Guédon, A. and Lepetit, V. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *CVPR*, 2024.
- Hamdi, A., Melas-Kyriazi, L., Mai, J., Qian, G., Liu, R., Vondrick, C., Ghanem, B., and Vedaldi, A. Ges : Generalized exponential splatting for efficient radiance field rendering. In *CVPR*, 2024.
- Hanson, A., Tu, A., Lin, G., Singla, V., Zwicker, M., and Goldstein, T. Speedy-splat: Fast 3d gaussian splatting with sparse pixels and sparse primitives. In *CVPR*, 2025a.
- Hanson, A., Tu, A., Singla, V., Jayawardhana, M., Zwicker, M., and Goldstein, T. Pup 3d-gs: Principled uncertainty pruning for 3d gaussian splatting. In *CVPR*, 2025b.
- Hedman, P., Philip, J., Price, T., Frahm, J.-M., Drettakis, G., and Brostow, G. Deep blending for free-viewpoint image-based rendering. *ACM TOG*, 2018.
- Hedman, P., Srinivasan, P. P., Mildenhall, B., Barron, J. T., and Debevec, P. Baking neural radiance fields for real-time view synthesis. In *ICCV*, 2021.
- Höllein, L., Božič, A., Zollhöfer, M., and Nießner, M. 3dgs-1m: Faster gaussian-splatting optimization with levenberg-marquardt. In *ICCV*, 2025.
- Huang, B., Yu, Z., Chen, A., Geiger, A., and Gao, S. 2d gaussian splatting for geometrically accurate radiance fields. In *SIGGRAPH*, 2024.
- Ji, B. and Yao, A. Sfm-free 3d gaussian splatting via hierarchical training. In *CVPR*, 2025.
- Kerbl, B., Kopanas, G., Leimkühler, T., and Drettakis, G. 3d gaussian splatting for real-time radiance field rendering. *ACM TOG*, 2023.
- Kerbl, B., Meuleman, A., Kopanas, G., Wimmer, M., Lanvin, A., and Drettakis, G. A hierarchical 3d gaussian representation for real-time rendering of very large datasets. *ACM TOG*, 2024.
- Kheradmand, S., Rebain, D., Sharma, G., Sun, W., Tseng, Y.-C., Isack, H., Kar, A., Tagliasacchi, A., and Yi, K. M. 3d gaussian splatting as markov chain monte carlo. In *NeurIPS*, 2024.
- Knapitsch, A., Park, J., Zhou, Q.-Y., and Koltun, V. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM TOG*, 2017.

- Lan, L., Shao, T., Lu, Z., Zhang, Y., Jiang, C., and Yang, Y. 3dgs<sup>2</sup>: Near second-order converging 3d gaussian splatting. In *SIGGRAPH*, 2025.
- Liu, Y., Luo, C., Fan, L., Wang, N., Peng, J., and Zhang, Z. Citygaussian: Real-time high-quality large-scale scene rendering with gaussians. In *ECCV*, 2024.
- Lu, C., Yin, F., Chen, X., Chen, T., Yu, G., and Fan, J. A large-scale outdoor multi-modal dataset and benchmark for novel view synthesis and implicit scene reconstruction. In *ICCV*, 2023.
- Lu, T., Yu, M., Xu, L., Xiangli, Y., Wang, L., Lin, D., and Dai, B. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In *CVPR*, 2024.
- Mallick, S. S., Goel, R., Kerbl, B., Steinberger, M., Carrasco, F. V., and De La Torre, F. Taming 3dgs: High-quality radiance fields with limited resources. In *SIGGRAPH Asia*, 2024.
- Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., and Ng, R. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.
- Müller, T., Evans, A., Schied, C., and Keller, A. Instant neural graphics primitives with a multiresolution hash encoding. *ACM TOG*, 2022.
- Pan, W., Zhang, X., Zhai, H., Xiang, X. a. H., and Zhang, G. Liberated-gs: 3d gaussian splatting independent from sfm point clouds. In *ICCV*, 2025.
- Papantonakis, P., Kopanas, G., Kerbl, B., Lanvin, A., and Drettakis, G. Reducing the memory footprint of 3d gaussian splatting. *ACM CGIT*, 2024a.
- Papantonakis, P., Kopanas, G., Kerbl, B., Lanvin, A., and Drettakis, G. Reducing the memory footprint of 3d gaussian splatting. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 2024b.
- Radl, L., Steiner, M., Parger, M., Weinrauch, A., Kerbl, B., and Steinberger, M. Stopthepop: Sorted gaussian splatting for view-consistent real-time rendering. *ACM TOG*, 2024.
- Ren, K., Jiang, L., Lu, T., Yu, M., Xu, L., Ni, Z., and Dai, B. Octree-gs: Towards consistent real-time rendering with lod-structured 3d gaussians. *IEEE TPAMI*, 2025.
- Ren, S., Wen, T., Fang, Y., and Lu, B. Fastgs: Training 3d gaussian splatting in 100 seconds. In *CVPR*, 2026.
- Rota Bulò, S., Porzi, L., and Kotschieder, P. Revising densification in gaussian splatting. In *ECCV*, 2024.
- Schönberger, J. L. and Frahm, J.-M. Structure-from-motion revisited. In *CVPR*, 2016.
- Shakiba Kheradmand, Daniel Rebain, G. S. H. I. A. K. A. T. K. M. Y. Accelerating neural field training via soft mining. In *CVPR*, 2024.
- Snaveley, N., Seitz, S. M., and Szeliski, R. Photo tourism: exploring photo collections in 3d. In *SIGGRAPH*. 2006.
- Sun, C., Sun, M., and Chen, H. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *CVPR*, 2022.
- Turki, H., Ramanan, D., and Satyanarayanan, M. Meganerf: Scalable construction of large-scale nerfs for virtual fly-throughs. In *CVPR*, 2022.
- Wang, X., Yi, R., and Ma, L. Adr-gaussian: Accelerating gaussian splatting with adaptive radius. In *SIGGRAPH Asia*, 2024.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 2004.
- Wu, X., Xu, J., Zhang, X., Bao, H., Huang, Q., Shen, Y., Tompkin, J., and Xu, W. Scannerf: Scalable bundle-adjusting neural radiance fields for large-scale scene rendering. *ACM TOG*, 2023.
- Xiangli, Y., Xu, L., Pan, X., Zhao, N., Rao, A., Theobalt, C., Dai, B., and Lin, D. Bungeenerf: Progressive neural radiance field for extreme multi-scale scene rendering. In *ECCV*, 2022.
- Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., and Neumann, U. Point-nerf: Point-based neural radiance fields. In *CVPR*, 2022.
- Xu, Q., Cui, J., Yi, X., Wang, Y., Zhou, Y., Ong, Y.-S., and Zhang, H. Pushing rendering boundaries: Hard gaussian splatting. In *AAAI*, 2026.
- Yu, Z., Chen, A., Huang, B., Sattler, T., and Geiger, A. Mip-splatting: Alias-free 3d gaussian splatting. In *CVPR*, 2024.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- Zhu, H., Liu, Z., Li, X., Wu, A., Zhao, J., Liu, F., Gan, Y., Leng, J., and Feng, Y. Nebula: Enable city-scale 3d gaussian splatting in virtual reality via collaborative rendering and accelerated stereo rasterization. In *ASPLOS*, 2026.

## Appendix

### A. TurboGS Optimization Procedure

**Algorithm 1 : TurboGS Training Framework.** Steps marked in blue are implemented as custom CUDA kernels / operators / buffers; Gray denotes lightweight CPU bookkeeping / value.

---

**Input** : Images & cameras  $(\mathcal{I}, \mathcal{C})$ , initial Gaussians  $\mathcal{G}$ , total training iterations  $T$ , sampling view budget range  $[K_{\min}, K_{\max}]$ , other hyper-params  $\Theta$

**Output** : Optimized Gaussians  $\mathcal{G}$

**Initialization** :  $\bar{\mathcal{E}}_{\text{init}}, \mathcal{E} \leftarrow \text{Initialize\_error\_maps}(\mathcal{G}, \mathcal{C}, \mathcal{I})$  // per-pixel error buffers (GPU)

**for**  $t \leftarrow 0$  **to**  $T - 1$  **do**

$\rho \leftarrow \min(1, t/\tau_{\text{until}})$  // training progress

**1) Sampling Parameters :** // non-linear query view budget & pixel sampling rate  
 $(K_t, r_t) \leftarrow \text{query\_sampling\_schedule}(K_{\max}, K_{\min}, \bar{\mathcal{E}}_{\text{init}}, \Theta, t)$

**2) View Selection (GVS) :** // anchor-by-error, then sampling-by pose/dir distance  
 $\mathcal{V}_t \leftarrow \text{geometry\_aware\_view\_sampling}(K_t, \rho)$

**3) Pixel Sampling (EPS) :** // tile-wise error-guided sparse pixel sampling  
 $\mathcal{P}_t \leftarrow \text{tile\_wise\_pixel\_sampling}(\mathcal{V}_t, r_t, \Theta)$

**4) Gather GT Color & Error :** // colors and current errors on sampled pixels  
 $(\mathbf{c}^*, \mathbf{e}^*) \leftarrow \text{get\_sampling\_pixel\_gt\_err}(\mathcal{V}_t, \mathcal{P}_t, \mathcal{I}, \mathcal{E})$

**5) Sparse Forward Rendering (FWD) :**  $\mathbf{o}, \mathbf{d} \leftarrow \text{rasterize\_forward\_pixels}(\mathcal{V}_t, \mathcal{P}_t, \mathbf{e}^*)$

**6) Sparse Loss & Grads. ( $\mathcal{L}$ ) :**  $(\nabla \mathbf{c}, \nabla \mathbf{d}, \mathcal{L}_{\text{pix}}) \leftarrow \text{calculate\_ll\_sncc\_grad\_loss}(\mathbf{o}, \mathbf{c}^*, \mathbf{d}, \mathbf{d}^*(\text{optional}))$

**7) EMA Error-Map Update ( $\mathcal{E}$ ) :**  $\mathcal{E} \leftarrow \text{update\_err\_statistics\_EMA}(\mathcal{E}, \mathcal{L}_{\text{pix}}, \mathcal{V}_t, \mathcal{P}_t)$

**8) Gaussian Backward & Screen Properties (BWD) :** // gradients of Gaussian properties  
 $\nabla \mathcal{G} \leftarrow \text{rasterize\_backward\_pixels}(\nabla \mathbf{c}, \nabla \mathbf{d}, \mathcal{G})$   
 $(\nabla \mathcal{G}_{\text{screen}}, \mathbf{r}) \leftarrow \text{get\_gaussian\_grad\_screen\_radii}(\nabla \mathcal{G}, \mathcal{G})$

**9) Solve Gaussian Updates ( $\Delta \mathcal{G}$ ) :**  $\Delta \mathcal{G} \leftarrow \text{solve\_gaussian\_delta\_properties}(\nabla \mathcal{G}, t, \rho)$

**10) Apply Gaussian Update :**  $\mathcal{G} \leftarrow \text{update\_gaussian\_properties}(\mathcal{G}, \Delta \mathcal{G})$

**11) Error-Driven Density Control (EDC) :** // adaptive Gaussian densification & pruning  
**Back-Projection Pixel Properties to Gaussians :** // accumulate (err, dist,  $\alpha$ ) per Gaussian  
 $e_{\mathcal{G}}, D_{\mathcal{G}}, \alpha_{\mathcal{G}} \leftarrow \text{rasterize\_forward\_pixels}(\mathcal{V}_t, \mathcal{P}_t, \mathcal{L}_{\text{pix}})$   
 $\mathcal{G} \leftarrow \text{error\_driven\_densify\_and\_prune}(\mathcal{G}, \{\nabla \mathcal{G}_{\text{screen}}, \mathbf{r}\}, \{e_{\mathcal{G}}, D_{\mathcal{G}}, \alpha_{\mathcal{G}}\}, t, \Theta)$

**return**  $\mathcal{G}$

---

Alg. 1 outlines the complete optimization pipeline of TurboGS, illustrating how error-guided sampling and sparse supervision are integrated into the 3DGS training loop. We highlight the input-output variable relationships between our *sparse sampling*, *sparse supervision*, *Gaussian property optimization*, and *error-driven density control*. In addition to the core algorithmic details already explained in the main paper, we further clarify the details of several steps:

**Adaptive Scheduling of View Budget and Pixel Sampling.** TurboGS dynamically adjusts both the sampled view budget and tile-wise pixel sampling rate according to reconstruction difficulty and training progress. At iteration  $t$ , we first compute a normalized error ratio:

$$r_t = \text{clip}\left(\frac{\bar{E}_t}{\bar{E}_{\text{init}} + \epsilon}, 0, 1\right), \quad \phi(r_t) = \sin\left(\frac{\pi}{2}\sqrt{r_t}\right), \quad (15)$$

where  $\bar{E}_t$  and  $\bar{E}_{\text{init}}$  denote the current and initial mean view errors. The sinusoidal modulation is motivated by the observation that reconstruction error drops rapidly in early training and then remains within a narrow range for most iterations (see Fig. 2 and Fig. 5). It amplifies sampling variation in this long low-error regime, improving sensitivity to subtle reconstruction differences.

We further define a progress coefficient  $\rho_t = 1 - \text{clip}(t/\tau_{\text{until}}, 0, 1)$  to gradually anneal sampling after densification. The tile-wise pixel sampling rate and sampled view number are jointly scheduled as:

$$r_{\text{tile}}^{(t)} = r_{\min} + (r_{\max} - r_{\min})(1 - \rho_t \phi(r_t)), \quad K_t = \text{round}\left(K_{\min} + (K_{\max} - K_{\min})e^{-\alpha(1 - \rho_t \phi(r_t))}\right), \quad (16)$$

where  $r_{\min} = \lambda/K_{\max}$ ,  $r_{\max} = \lambda/K_{\min}$ ,  $\lambda$  denotes the tile sampling scale, and  $\alpha$  controls the transition speed. This design allocates more views and denser pixel sampling during difficult stages, while progressively reducing computation and shifting from global to more local view exploration as optimization converges.

**Gradient Computation under Sparse Supervision.** Under sparse pixel sampling, TurboGS computes gradients by combining a pixel-wise Charbonnier  $\ell_1$  term and a tile-level sparse NCC term. For a sampled pixel with prediction  $x$  and ground truth  $y$ , the gradients are:

$$\frac{\partial \ell_1}{\partial x} = \frac{x - y}{\sqrt{(x - y)^2 + \epsilon^2}}, \quad \frac{\partial \ell_{\text{NCC}}}{\partial x} = - \left( \frac{y - \mu_y}{\sqrt{\sigma_x^2 \sigma_y^2 + \epsilon}} - \text{NCC} \frac{x - \mu_x}{\sigma_x^2 + \epsilon} \right), \quad (17)$$

where  $\mu_x, \mu_y$  and  $\sigma_x^2, \sigma_y^2$  denote the tile-wise mean and variance computed from sampled pixels.

To stabilize sparse optimization, we adaptively balance the two terms using tile-level gradient statistics:

$$\lambda_{\text{ncc}}^\tau = \text{clip} \left( \frac{\mathbb{E}[|\nabla \ell_1|]}{\mathbb{E}[|\nabla \ell_{\text{NCC}}|] + \epsilon}, 0, 1 \right), \quad \frac{\partial \mathcal{L}}{\partial x} = \frac{\partial \ell_1}{\partial x} - \lambda_{\text{ncc}}^\tau \frac{\partial \ell_{\text{NCC}}}{\partial x}. \quad (18)$$

**Details of Our Moment-Damped LM Optimizer.** TurboGS adopts a hybrid optimization strategy that combines lightweight Levenberg–Marquardt (LM) updates with Adam-style moment damping. For each Gaussian, we partition the parameters into two compact blocks: geometry parameters  $\mathbf{g} \in \mathbb{R}^{10}$  (position, scale, rotation) and base appearance parameters  $\mathbf{a} \in \mathbb{R}^4$  (RGB DC and opacity), which are solved independently through small LM systems:

$$(\mathbf{J}^\top \mathbf{J} + \text{diag}(\sqrt{\mathbf{v}} + \epsilon)) \Delta = -\mathbf{m}, \quad \Delta \theta_{\text{SH}} = -\eta \frac{\hat{\mathbf{m}}}{\sqrt{\hat{\mathbf{v}} + \epsilon}}, \quad (19)$$

where  $\mathbf{m}$  and  $\mathbf{v}$  denote Adam first and second moments, and  $\mathbf{v}$  serves as an adaptive diagonal damping term. Each LM system is efficiently solved via Cholesky decomposition.

Higher-order spherical harmonic (SH) coefficients are updated separately using Adam to avoid enlarging the LM system size. Ignoring SH cross-parameter correlations is a practical trade-off between accuracy and efficiency, since joint optimization would substantially increase Jacobian size and memory overhead. Empirically, this hybrid strategy improves stability under sparse gradients without noticeable convergence degradation. In practice, LM-based updates and SH optimization are executed in two dedicated CUDA kernels for efficient parallelism.

**Convergence Stability under Sparse Optimization.** TurboGS adopts a structured sparse optimization scheme whose stochastic gradients remain aligned with the dense objective. Let the dense loss be  $\mathcal{L}(\theta) = \sum_{i=1}^N \ell_i(\theta)$  over image pixels  $i$ . Under sparse sampling set  $\mathcal{S}$ , the gradient estimator is:

$$\hat{\mathbf{g}} = \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} w(i) \nabla \ell_i(\theta), \quad \mathbb{E}[\hat{\mathbf{g}}] = \sum_{i=1}^N p(i) w(i) \nabla \ell_i(\theta), \quad (20)$$

where  $p(i)$  denotes the sampling probability and  $w(i)$  is the importance weight. Under ideal importance sampling with  $w(i) = 1/p(i)$ , the estimator recovers the dense gradient in expectation, *i.e.*,  $\mathbb{E}[\hat{\mathbf{g}}] \propto \nabla \mathcal{L}(\theta)$ , providing convergence consistency while reducing computation.

In practice, TurboGS combines error-guided hard sampling and age-aware stable sampling, prioritizing informative pixels while periodically revisiting under-sampled regions to avoid overly localized supervision. Moreover, the moment-damped LM update improves optimization stability by introducing adaptive diagonal damping through the Adam second moment,  $\mathbf{J}^\top \mathbf{J} + \text{diag}(\sqrt{\mathbf{v}} + \epsilon)$ , which improves local conditioning under sparse gradients. Together, these designs lead to stable empirical convergence in sparse-pixel optimization.

## B. Implementation Details

All experiments are conducted on an NVIDIA RTX 5090 GPU, except for 3DGS-LM (Höllerle et al., 2025), which is evaluated on an NVIDIA RTX PRO 6000 GPU due to its memory requirement exceeding 32 GB. The RTX PRO 6000

Table 6. Definition and values of the hyper-parameters  $\Theta$  in TurboGS.

H.P.	Definition	Value
$T$	Total training iterations	5000 (TurboGS, fast), 8000 (TurboGS-Big)
$\tau_{\text{until}}$	Last iteration for densification and pruning	3500 (TurboGS, fast), 6000 (TurboGS-Big)
$\tau_{\text{from}}$	First iteration for densification	200
$K_{\text{max}}$	Maximum sampled views per iteration	10
$K_{\text{min}}$	Minimum sampled views per iteration	3
$\lambda$	Tile pixel sampling scale	(0, 1]
$\alpha$	Transition speed of adaptive view scheduling	3.0
$\Delta_{\text{dens}}$	Densification interval (iterations)	100 ~ 350
$\Delta_{\text{prune}}$	Final pruning interval (iterations)	500
$\eta_{\text{topk}}$	Fraction for top- $k$ densification	0.5 ~ 0.95
$\eta_{\text{final}}$	Fraction retained by final pruning	0.8 ~ 0.95
$\tau_{\text{grad}}$	Gradient threshold for densification	$1.0 \times 10^{-4} \sim 2.0 \times 10^{-4}$
$\tau_{\text{grad}}^{\text{abs}}$	Absolute-gradient threshold	$2.0 \times 10^{-4} \sim 4.0 \times 10^{-4}$
$\rho_{\text{dense}}$	Percent dense	0.001 ~ 0.02
$r_{\text{hard}}$	Sampling ratio for difficult pixels	0.4
$r_{\text{stable}}$	Sampling ratio for long-untrained pixels	0.3
-	Remaining ratio for mixed sampling	0.3
$\beta$	EMA blending ratio for pixel errors	0.4
$\lambda_{\text{base}}$	Base weight of distance-aware Gaussian difficulty	0.8

provides comparable compute throughput (TFLOPS) to the RTX 5090 while offering 96 GB GPU memory. All baselines are evaluated using their official implementations under the original 3DGS setting. TurboGS is implemented in CUDA and PyTorch, with sparse rasterization, loss and gradient evaluation, and optimization executed as custom CUDA kernels. Training hyper-parameters are summarized in Tab. 6.

### C. Ablation on Sampling Hyper-parameters

Table 7. Ablations on adaptive sampling hyper-parameters. We evaluate the effects of the adaptive view budget ( $K_{\text{max}} \rightarrow K_{\text{min}}$ ), hard-pixel sampling ratio ( $r_{\text{hard}}$ ), and error-ratio scheduling function  $\phi(r_t)$  on Mip-NeRF 360 (Barron et al., 2022) under the same setup.

$K$	Time↓	PSNR↑	SSIM↑	LPIPS↓	$N_{\text{GS}} \downarrow$	$r_{\text{hard}}$	Time↓	PSNR↑	SSIM↑	LPIPS↓	$N_{\text{GS}} \downarrow$
1 (fixed)	81s	26.50	0.768	0.282	1.24M	0.2	<b>77s</b>	27.29	0.785	0.266	<b>0.59M</b>
3 (fixed)	<u>78s</u>	27.23	<u>0.787</u>	<u>0.263</u>	0.74M	<b>0.4 (Ours)</b>	<b>77s</b>	<b>27.57</b>	<b>0.794</b>	<b>0.256</b>	0.64M
10 $\Rightarrow$ 3 (Ours)	<b>77s</b>	<b>27.57</b>	<b>0.794</b>	<b>0.256</b>	<u>0.64M</u>	0.6	82s	27.33	<u>0.787</u>	<u>0.262</u>	0.65M
10 (fixed)	121	<u>27.41</u>	0.783	<u>0.263</u>	<b>0.51M</b>	0.8	<u>80s</u>	<u>27.41</u>	0.785	0.264	<u>0.60M</u>

(a) Ablation on adaptive view budget

(b) Ablation on hard-pixel sampling ratio

Schedule	Time↓	PSNR↑	SSIM↑	LPIPS↓	$N_{\text{GS}} \downarrow$
<i>Linear</i>	81s	<u>27.35</u>	<u>0.788</u>	<u>0.261</u>	<u>0.63M</u>
<i>Exponential</i>	<u>79s</u>	27.27	0.783	0.268	<b>0.55M</b>
<i>Sinusoidal</i> (Eq. 15, Ours)	<b>77s</b>	<b>27.57</b>	<b>0.794</b>	<b>0.256</b>	0.64M

(c) Ablation on error-ratio scheduling

We further analyze the sensitivity of TurboGS to key sampling hyper-parameters, including the adaptive view budget, hard-pixel sampling ratio ( $r_{\text{hard}}$ ), and error-ratio scheduling function  $\phi(r_t)$ . Results are reported in Tab. 7.

**Effect of Adaptive View Budget.** As shown in Tab. 7(a), increasing the number of sampled views improves reconstruction quality but incurs higher computation. A very small view budget ( $K=1$ ) degrades performance due to insufficient multi-view constraints, while a large fixed budget ( $K=10$ ) increases training time without gains. Our adaptive strategy (10 $\rightarrow$ 3) achieves the best trade-off by gradually shifting from global multi-view consistency to local refinement.

**Effect of Hard-Pixel Sampling Ratio.** Tab. 7(b) shows that a small  $r_{\text{hard}}$  under-samples informative regions, whereas a large value reduces sampling diversity and increases cost. We find  $r_{\text{hard}}=0.4$  provides a balance between speed and quality.




Figure 13. Interactive GUI of TurboGS. Real-time rendering and free-viewpoint exploration of reconstructed scenes.

**Effect of Error-Ratio Scheduling.** Tab. 7(c) compares different scheduling functions for adaptive sampling. The proposed sinusoidal schedule consistently outperforms linear and exponential variants, likely because it amplifies sampling variation in the long low-error regime, improving sensitivity to subtle reconstruction differences.

## D. Additional Scene-wise Comparisons and Visualizations

We provide additional scene-wise evaluations and visualizations. Tab. 8, Tab. 9, Tab. 10, Tab. 11, and Tab. 12 report detailed scene-wise quantitative comparisons on Mip-NeRF 360, Tanks & Temples, Deep Blending, BungeeNeRF, and 4K Rubble sub-regions, respectively. Fig. 13 presents the interactive GUI of our reconstruction system. Fig. 14 and Fig. 15 visualize intermediate optimization results on the large-scale OMMO dataset and 4K Rubble sub-regions, respectively. In addition, Fig. 16 shows further qualitative comparisons on representative scenes under matched training time.

Table 8. Scene-wise quantitative results on Mip-NeRF 360 (Barron et al., 2022) dataset. We report the per-scene comparisons of TurboGS and other accelerated or improved 3DGS optimization methods.

Method	Bicycle						Flowers						Garden					
	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑
3DGS	1069	25.27	0.767	0.208	4.91M	105	706	21.57	0.605	0.336	2.90M	207	1105	27.52	0.868	0.107	4.14M	138
Taming-3DGS	139	24.82	0.717	0.294	0.81M	407	118	21.10	0.555	0.407	0.57M	526	223	27.43	0.858	0.126	1.94M	320
Mini-Splatting	524	25.20	0.772	0.225	0.53M	241	569	21.42	0.625	0.327	0.57M	224	566	26.84	0.847	0.150	0.56M	233
Speedy-Splat	584	24.90	0.724	0.303	0.61M	777	470	21.39	0.579	0.395	0.36M	982	598	26.92	0.834	0.183	0.53M	843
3DGS-LM	727	25.21	0.764	0.219	6.06M	105	524	21.49	0.600	0.342	3.60M	196	762	27.33	0.865	0.110	5.93M	116
DashGaussian	225	25.38	0.771	0.213	4.26M	197	168	21.94	0.617	0.327	2.50M	302	201	27.58	0.863	0.121	2.82M	270
FastGS	100	24.82	0.722	0.296	0.48M	1057	104	21.37	0.575	0.388	0.45M	970	132	27.50	0.847	0.157	0.66M	958
ConeGS	694	25.50	0.782	0.188	1.27M	354	646	21.55	0.625	0.303	1.00M	393	671	27.65	0.870	0.106	1.32M	362
<b>TurboGS (Ours)</b>	<b>73</b>	25.27	0.716	0.291	0.85M	187	<b>73</b>	22.06	0.580	0.387	0.78M	131	<b>84</b>	27.23	0.823	0.208	0.75M	166
<b>TurboGS-Big (Ours)</b>	160	<b>25.66</b>	0.747	0.257	1.87M	147	180	<b>22.61</b>	<b>0.630</b>	0.327	2.36M	118	154	<b>27.62</b>	0.858	0.155	1.06M	148

Method	Stump						Treehill						Room					
	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑
3DGS	898	26.66	0.771	0.216	4.29M	156	826	22.59	0.634	0.326	3.27M	175	880	31.62	0.921	0.217	1.30M	234
Taming-3DGS	102	26.08	0.736	0.292	0.48M	568	134	23.03	0.625	0.384	0.78M	420	136	31.32	0.909	0.249	0.23M	389
Mini-Splatting	504	27.17	<b>0.806</b>	0.198	0.61M	217	564	22.72	0.654	<b>0.314</b>	0.57M	218	780	31.27	0.921	0.212	0.40M	440
Speedy-Splat	511	26.68	0.770	0.259	0.51M	876	460	22.50	0.592	0.443	<b>0.36M</b>	<b>1066</b>	538	30.78	0.896	0.278	<b>0.12M</b>	<b>1088</b>
3DGS-LM	994	26.70	0.776	0.218	4.86M	154	519	22.55	0.634	0.332	3.80M	179	520	31.31	0.915	0.222	1.54M	283
DashGaussian	141	27.05	0.783	0.216	3.42M	313	188	23.06	0.640	<b>0.319</b>	3.20M	257	134	31.68	0.916	0.230	1.04M	245
FastGS	<b>86</b>	26.50	0.751	0.283	<b>0.35M</b>	<b>1081</b>	90	22.67	<b>0.674</b>	0.416	<b>0.36M</b>	1040	104	31.72	0.913	0.239	<b>0.21M</b>	<b>1045</b>
ConeGS	675	27.19	0.804	<b>0.182</b>	1.24M	375	604	22.65	0.605	0.351	0.76M	473	941	32.13	<b>0.927</b>	<b>0.200</b>	0.47M	524
<b>TurboGS (Ours)</b>	<b>62</b>	26.66	0.742	0.279	0.65M	179	<b>69</b>	<b>23.18</b>	0.609	0.392	0.69M	164	<b>68</b>	31.71	0.915	0.215	0.45M	155
<b>TurboGS-Big (Ours)</b>	171	<b>27.37</b>	0.784	0.227	2.04M	137	114	<b>23.22</b>	0.614	0.368	1.17M	130	173	<b>32.15</b>	<b>0.925</b>	<b>0.191</b>	0.95M	137

Method	Counter						Kitchen						Bonsai					
	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑
3DGS	836	29.07	0.909	0.200	1.08M	242	1006	31.28	0.928	0.126	1.59M	202	692	32.35	0.943	0.203	1.07M	332
Taming-3DGS	156	28.61	0.898	0.223	0.31M	341	179	31.15	0.922	0.141	0.48M	335	147	31.78	0.935	0.219	0.41M	412
Mini-Splatting	896	28.61	0.905	0.198	0.41M	358	900	31.24	0.926	0.129	0.43M	272	764	31.41	0.939	0.200	0.36M	310
Speedy-Splat	538	28.18	0.869	0.274	<b>0.10M</b>	<b>1058</b>	589	30.01	0.890	0.202	<b>0.11M</b>	<b>1055</b>	508	31.18	0.920	0.260	<b>0.13M</b>	<b>1102</b>
3DGS-LM	504	28.83	0.907	0.204	1.21M	300	596	31.30	0.928	0.127	1.82M	239	453	31.97	0.941	0.206	1.25M	386
DashGaussian	139	28.95	0.903	0.209	0.79M	250	197	31.44	0.921	0.141	1.26M	199	131	32.09	0.940	0.206	0.84M	295
FastGS	116	29.10	0.900	0.219	0.21M	988	153	31.75	0.925	0.136	0.39M	880	118	32.15	0.937	0.212	0.28M	998
ConeGS	991	28.97	<b>0.910</b>	0.189	0.56M	410	828	31.68	<b>0.934</b>	<b>0.117</b>	0.74M	410	953	<b>32.37</b>	<b>0.947</b>	<b>0.185</b>	0.46M	540
<b>TurboGS (Ours)</b>	<b>85</b>	28.63	0.898	0.204	0.49M	125	<b>92</b>	31.55	0.923	0.131	0.59M	148	<b>85</b>	31.82	0.939	0.202	0.55M	122
<b>TurboGS-Big (Ours)</b>	207	<b>29.14</b>	0.906	<b>0.186</b>	0.80M	123	191	<b>31.81</b>	0.926	0.127	0.73M	135	162	32.31	<b>0.945</b>	<b>0.176</b>	1.07M	131

Table 9. Scene-wise quantitative results over the Tanks & Temples (Knapitsch et al., 2017) dataset.

Method	Truck						Train					
	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑
3DGS	625	25.47	0.884	0.142	2.07M	231	540	22.22	<b>0.822</b>	<b>0.196</b>	1.09M	272
Taming-3DGS	<b>90</b>	25.22	0.868	0.184	0.27M	<b>659</b>	104	22.32	0.804	0.238	0.37M	565
Mini-Splatting	527	25.05	0.874	0.160	<b>0.19M</b>	329	431	21.42	0.799	0.244	<b>0.21M</b>	482
Speedy-Splat	332	25.13	0.868	0.189	0.26M	<b>1122</b>	252	21.68	0.773	0.289	<b>0.11M</b>	<b>1153</b>
3DGS-LM	424	25.37	0.881	0.152	2.51M	235	314	21.77	0.807	0.218	1.09M	335
DashGaussian	130	<b>25.84</b>	<b>0.885</b>	0.152	1.40M	359	151	22.25	<b>0.818</b>	0.209	1.01M	309
FastGS	<b>96</b>	25.76	0.877	0.177	<b>0.25M</b>	<b>1134</b>	<b>86</b>	<b>22.36</b>	0.807	0.240	<b>0.23M</b>	<b>1040</b>
ConeGS	740	25.75	<b>0.890</b>	<b>0.116</b>	0.65M	511	758	22.01	<b>0.816</b>	<b>0.204</b>	0.46M	617
<b>TurboGS (Ours)</b>	<b>76</b>	25.43	0.873	0.166	0.75M	159	<b>69</b>	22.15	0.791	0.234	0.41M	182
<b>TurboGS-Big (Ours)</b>	148	<b>25.82</b>	0.883	0.146	1.16M	144	144	<b>22.46</b>	0.799	0.216	0.65M	158

Table 10. Scene-wise quantitative results over the Deep Blending (Hedman et al., 2018) dataset.

Method	Drjohnson						Playroom					
	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑
3DGS	1046	29.45	<b>0.905</b>	<b>0.236</b>	1.84M	163	795	30.25	0.910	<b>0.240</b>	3.11M	252
Taming-3DGS	<b>109</b>	29.51	0.903	0.267	0.40M	561	<b>91</b>	30.33	0.903	0.276	<b>0.18M</b>	640
Mini-Splatting	632	29.40	<b>0.904</b>	0.257	<b>0.38M</b>	340	539	30.49	<b>0.911</b>	0.250	0.32M	387
Speedy-Splat	444	29.05	0.900	0.267	<b>0.31M</b>	<b>1094</b>	549	30.05	0.907	0.270	<b>0.19M</b>	<b>1149</b>
3DGS-LM	645	28.97	0.903	0.247	3.48M	172	498	30.19	0.908	<b>0.246</b>	2.30M	264
DashGaussian	125	<b>29.56</b>	<b>0.905</b>	0.247	2.53M	248	99	<b>30.67</b>	0.910	0.249	1.37M	329
FastGS	<b>86</b>	29.48	0.900	0.272	<b>0.25M</b>	<b>1137</b>	<b>76</b>	30.64	<b>0.911</b>	0.262	<b>0.13M</b>	<b>1152</b>
ConeGS	826	<b>29.84</b>	<b>0.905</b>	<b>0.241</b>	0.61M	610	783	<b>30.72</b>	<b>0.913</b>	<b>0.235</b>	0.44M	700
<b>TurboGS (Ours)</b>	<b>62</b>	29.53	0.892	0.285	0.42M	273	<b>61</b>	30.47	0.908	0.270	0.55M	188
<b>TurboGS-Big (Ours)</b>	139	<b>29.78</b>	<b>0.904</b>	0.260	0.67M	204	125	<b>30.73</b>	<b>0.913</b>	0.258	0.72M	227

Table 11. Scene-wise quantitative results over the BungeeNeRF (Xiangli et al., 2022) dataset.

Method	Amsterdam						Bilbao						Hollywood					
	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑
Taming-3DGS	234	24.45	0.804	0.257	<b>0.65M</b>	243	184	25.62	0.825	0.246	<b>0.44M</b>	278	198	23.72	0.713	0.363	<b>0.55M</b>	248
DashGaussian	350	<b>26.68</b>	<b>0.880</b>	<b>0.159</b>	3.17M	125	322	<b>28.07</b>	<b>0.900</b>	<b>0.135</b>	2.75M	126	347	<b>25.83</b>	<b>0.843</b>	<b>0.192</b>	3.45M	143
FastGS	<b>129</b>	24.82	0.816	0.247	<b>0.56M</b>	<b>628</b>	125	26.24	0.846	0.220	<b>0.55M</b>	<b>766</b>	<b>143</b>	24.22	0.754	0.317	<b>0.72M</b>	<b>743</b>
FastGS-Big	212	25.81	0.857	0.191	1.30M	<b>509</b>	207	27.11	0.875	0.170	1.28M	<b>549</b>	232	24.68	0.793	0.256	1.41M	<b>476</b>
<b>TurboGS (Ours)</b>	<b>96</b>	25.40	0.803	0.260	0.69M	150	<b>91</b>	27.28	0.843	0.219	0.76M	145	<b>105</b>	24.75	0.745	0.314	0.93M	132
<b>TurboGS-Big (Ours)</b>	236	<b>26.57</b>	<b>0.863</b>	<b>0.184</b>	1.86M	99	227	<b>27.94</b>	<b>0.880</b>	<b>0.168</b>	1.86M	95	258	<b>25.78</b>	<b>0.817</b>	<b>0.221</b>	2.51M	85

Method	Pompidou						Rome					
	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑
Taming-3DGS	230	24.04	0.830	0.234	<b>0.68M</b>	271	224	24.11	0.794	0.278	<b>0.62M</b>	242
DashGaussian	409	<b>26.44</b>	<b>0.903</b>	<b>0.124</b>	4.20M	127	337	<b>26.69</b>	<b>0.890</b>	<b>0.151</b>	3.12M	130
FastGS	<b>150</b>	24.56	0.848	0.210	<b>0.73M</b>	<b>609</b>	<b>133</b>	24.71	0.818	0.252	<b>0.61M</b>	<b>578</b>
FastGS-Big	227	25.41	0.877	0.166	1.52M	<b>511</b>	256	25.81	0.861	0.190	1.41M	<b>433</b>
<b>TurboGS (Ours)</b>	<b>112</b>	25.27	0.830	0.225	0.99M	128	<b>102</b>	25.77	0.818	0.244	<b>0.85M</b>	126
<b>TurboGS-Big (Ours)</b>	184	<b>26.24</b>	<b>0.883</b>	<b>0.159</b>	2.61M	86	245	<b>26.85</b>	<b>0.874</b>	<b>0.170</b>	2.39M	88

Table 12. Scene-wise quantitative results on representative sub-regions of the Rubble (Turki et al., 2022) dataset at 4K resolution, where COLMAP (Schönberger & Frahm, 2016) is independently performed for each sub-region to ensure fair evaluation.

Method	Sub-Region #1 (view 172 to 228)						Sub-Region #2 (view 58 to 114)						Sub-Region #3 (view 690 to 746)					
	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑
Taming-3DGS	858	25.34	0.737	0.393	<b>1.00M</b>	53	925	24.81	0.736	0.399	<b>0.94M</b>	51	911	26.18	0.728	0.409	<b>0.86M</b>	50
FastGS	<b>639</b>	<b>26.17</b>	<b>0.791</b>	<b>0.322</b>	<b>1.90M</b>	<b>87</b>	<b>630</b>	<b>25.79</b>	<b>0.799</b>	<b>0.290</b>	3.73M	<b>141</b>	<b>574</b>	<b>26.51</b>	<b>0.775</b>	<b>0.315</b>	3.05M	<b>149</b>
<b>TurboGS (Ours)</b>	<b>361</b>	<b>26.57</b>	<b>0.776</b>	<b>0.295</b>	2.83M	50	<b>367</b>	<b>25.88</b>	<b>0.775</b>	<b>0.273</b>	<b>2.78M</b>	47	<b>353</b>	<b>26.73</b>	<b>0.760</b>	<b>0.285</b>	<b>2.44M</b>	48

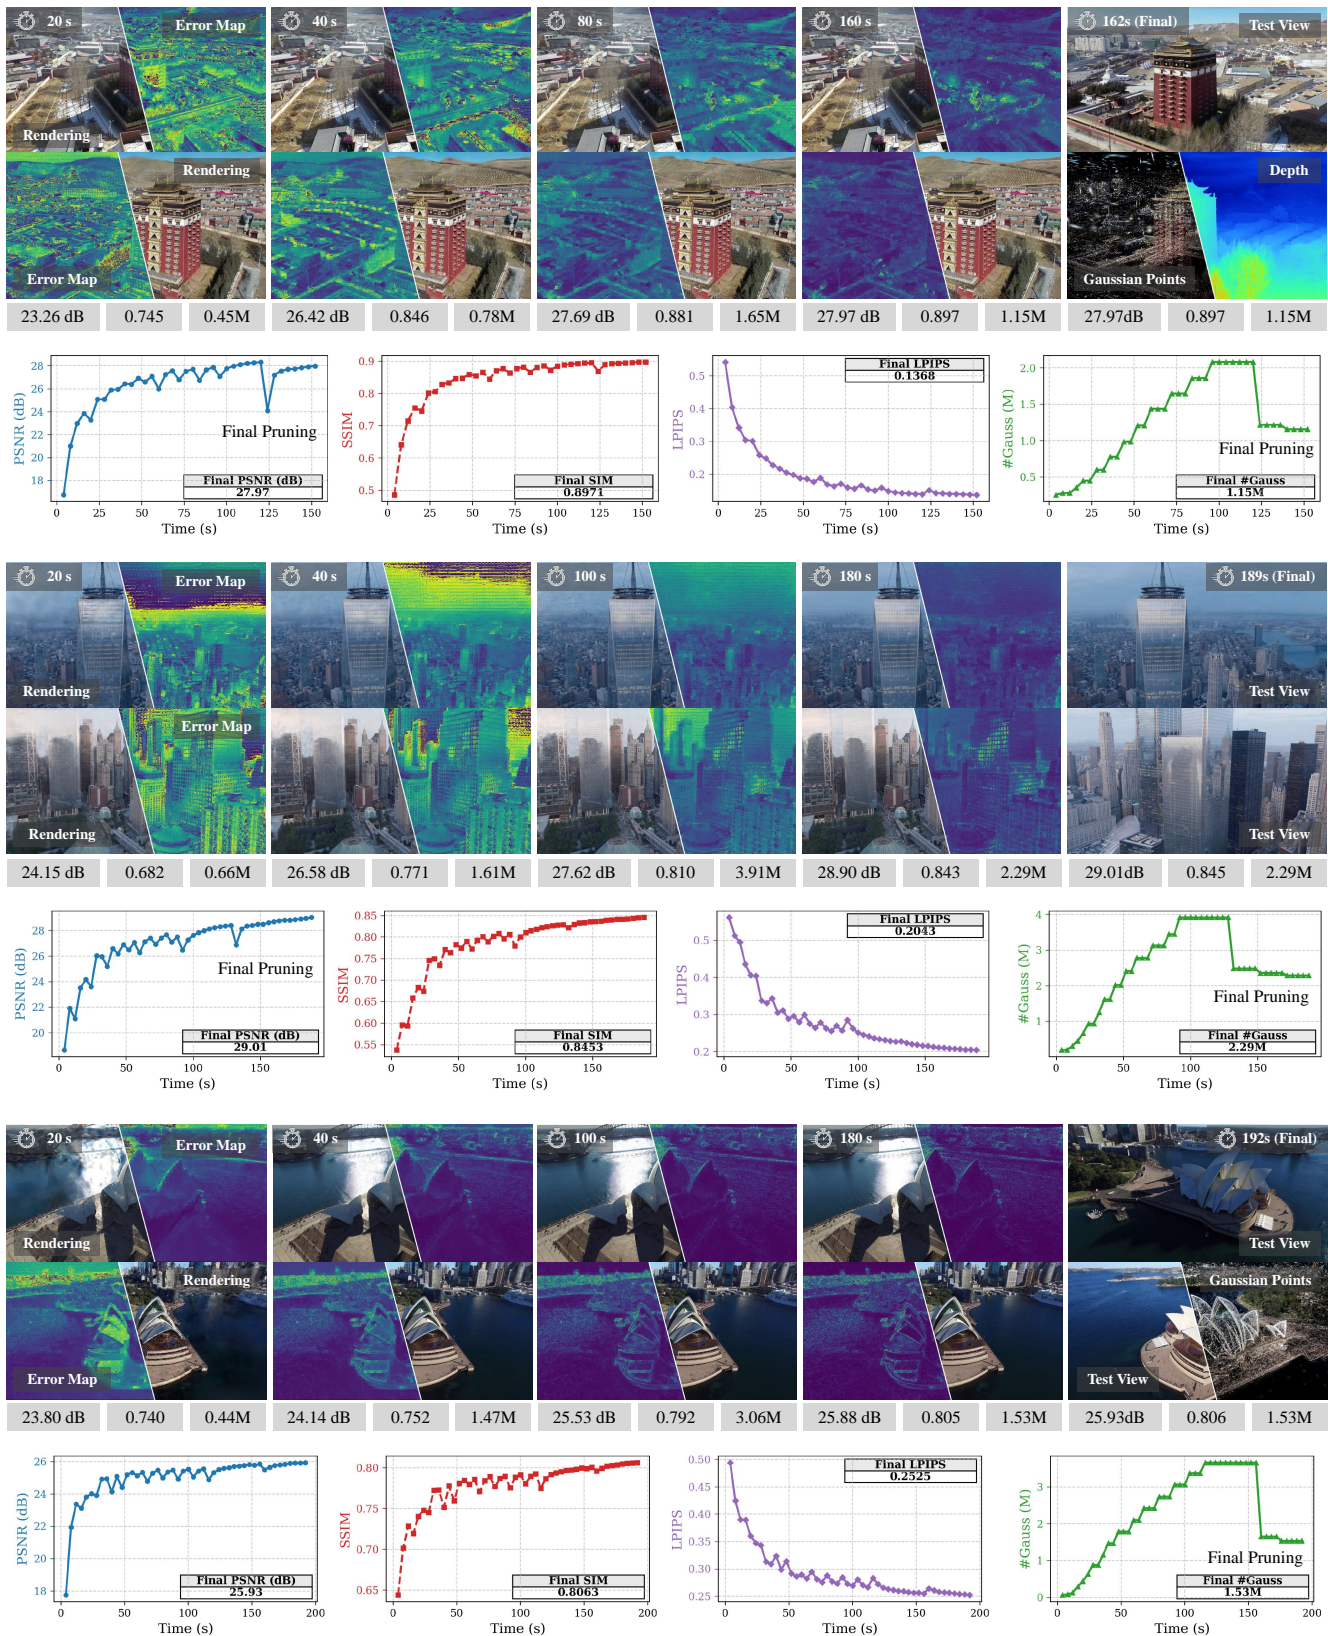


Figure 14. Optimization progress on large-scale OMMO scenes. We visualize rendering results, error maps, and training curves at representative iterations (e.g., 20s, 40s, 80s, and final), showing progressive reconstruction improvement and convergence behavior.

# TurboGS: Accelerating 3D Gaussian Splatting via Error-Guided Sparse Pixel Sampling and Optimization

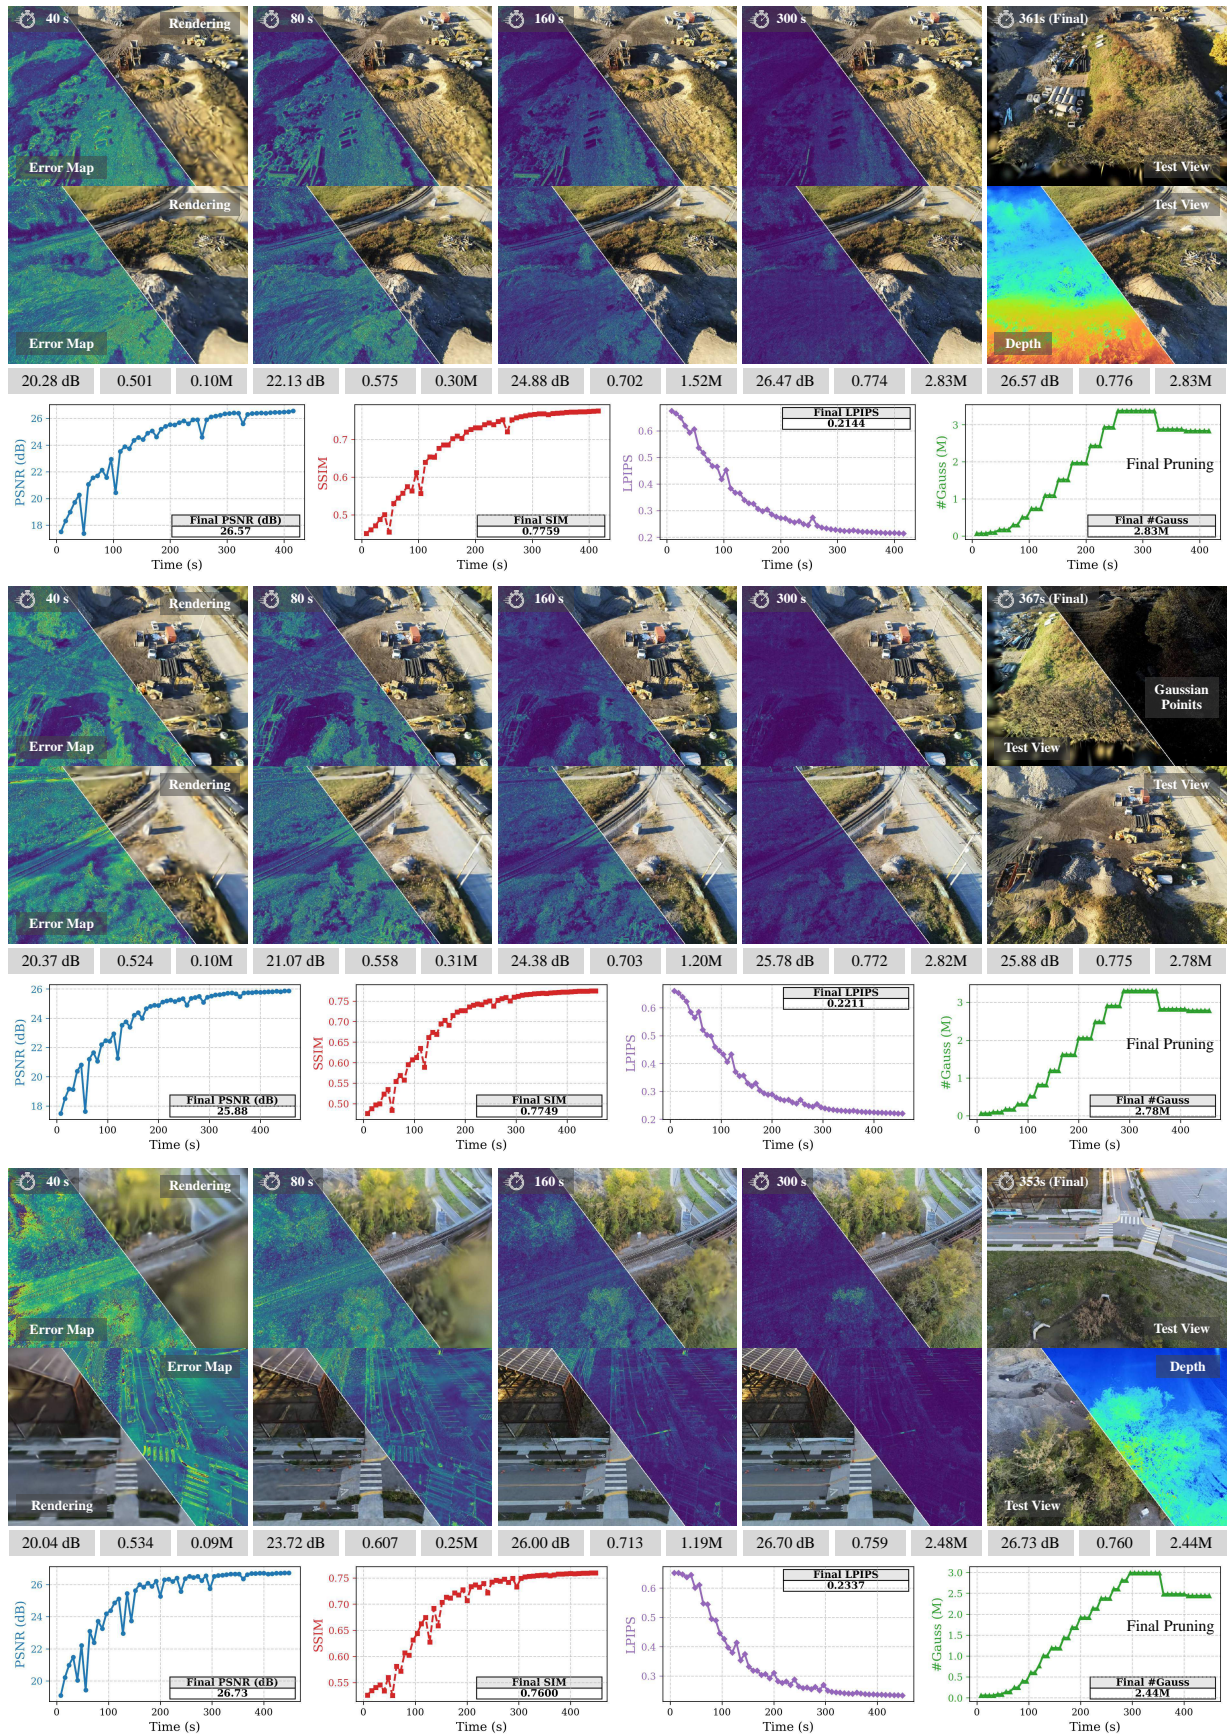


Figure 15. Optimization progress on 4K Rubble sub-regions. Representative rendering results and training curves during optimization.

**TurboGS: Accelerating 3D Gaussian Splatting via Error-Guided Sparse Pixel Sampling and Optimization**




Figure 16. **Additional qualitative comparisons under the same training time budget.** We visualize representative results from existing fast 3DGS optimization methods at the time when TurboGS finishes optimization, on the Mip-NeRF 360 (Barron et al., 2022) dataset (*Bicycle*, *Counter*, *Flowers*) and the Deep Blending (Hedman et al., 2018) dataset (*Drjohnson*). Cropped patches highlight differences in detail reconstruction and visual fidelity.