

Gaussian On-the-Fly Splatting: A Progressive Framework for Robust Near Real-Time 3DGS Optimization

Yiwei Xu, Yifei Yu, Wentian Gan, Tengfei Wang, Zongqian Zhan, Hao Cheng and Xin Wang

Abstract—3D Gaussian Splatting (3DGS) achieves high-fidelity rendering with fast real-time performance, but existing methods rely on offline training after full Structure-from-Motion (SfM) processing. In contrast, this work introduces On-the-Fly GS, a progressive framework enabling near real-time 3DGS optimization during image capture. As each image arrives, its pose and sparse points are updated via on-the-fly SfM, and newly optimized Gaussians are immediately integrated into the 3DGS field. We propose a progressive local optimization strategy to prioritize new images and their neighbors by their corresponding overlapping relationship, allowing the new image and its overlapping images to get more training. To further stabilize training across old and new images, an adaptive learning rate schedule balances the iterations and the learning rate. Moreover, to maintain overall quality of the 3DGS field, an efficient global optimization scheme prevents overfitting to the newly added images. Experiments on multiple benchmark datasets show that our On-the-Fly GS reduces training time significantly, optimizing each new image in seconds with minimal rendering loss, offering the first practical step toward rapid, progressive 3DGS reconstruction. https://xywjohn.github.io/GS_On-the-Fly.github.io/

I. INTRODUCTION

Novel View Synthesis (NVS) aims to generate realistic images of novel views based on a set of input images [1]. Recently, Kerbl et al. [2] introduced a novel 3D representation technique known as 3D Gaussian Splatting (3DGS), which employs a large number of 3D Gaussian kernels to represent a 3D scene. Subsequent research has made notable advancements in improving the 3DGS field in terms of training efficiency [3]–[11], rendering speed [12], [13], and output quality [14]–[18]. However, most existing approaches rely on an offline training solution, which means that 3DGS can only be optimized after all images are processed by Colmap [19], preventing immediate synthesis of novel views upon completion of data acquisition [20]. While some recent works [20]–[24] have made progress towards real-time/near real-time training of 3DGS by integrating SLAM, such as GS-SLAM [21], MonoGS++ [24]. They take indoor video frames with spatial and temporal continuity as input, which might degenerate in some practical applications [25], [26].

To address these limitations, we propose On-the-Fly GS, a progressive framework for near real-time 3DGS field training

*This work was supported by the National Natural Science Foundation of China (No.42301507) and Natural Science Foundation of Hubei Province, China (No. 2022CFB727).

¹Yiwei Xu, Yifei Yu, Wentian Gan, Tengfei Wang, Zongqian Zhan, Xin Wang are with the School of Geodesy and Geomatics, Wuhan University, China PR. (Corresponding Author: Xin Wang, xwang@sgg.whu.edu.cn)

²Hao Cheng is with the Department of Earth Observation Science at ITC Faculty Geo-Information Science and Earth Observation, University of Twente, the Netherlands.

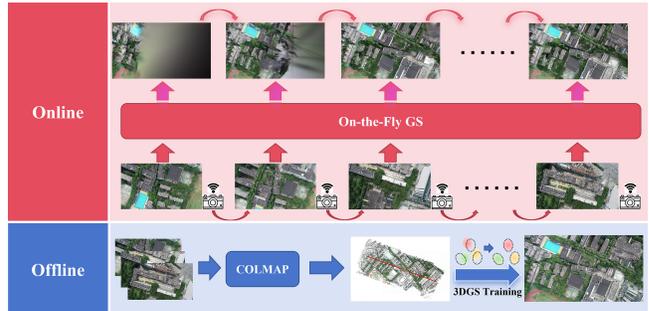


Fig. 1. Our online On-the-Fly GS vs. Offline Training Method

that enables 3D Gaussian training during discrete image capturing. As shown in Fig. 1, in our work, a new image can be acquired in an arbitrary manner. Its pose and relevant new point cloud are estimated online by the on-the-fly SfM [26]¹, which are then inputted into our on-the-fly GS to optimize the corresponding local 3DGS field in near real time and update the entire 3DGS field.

In contrast to the original 3DGS that focuses on reconstructing a static scenario using a fixed set of images and their corresponding poses, our On-the-Fly GS targets a dynamically expanding scenario, where images are progressively captured over time. This shift necessitates a different training strategy tailored for the incremental expansion of the 3DGS fields. A key challenge in training dynamic scenes lies in the inconsistency of local quality within the 3DGS fields. Images captured earlier undergo more training iterations, resulting in higher rendering quality, while newly acquired images, having undergone fewer training iterations, exhibit lower rendering quality. Consequently, the learning rate decay strategy of the original 3DGS based on complete 3DGS field quality is no longer applicable. Instead, On-the-Fly GS calculates the learning rate for each image based on the number of iterations it has undergone. Images with fewer training iterations are assigned higher learning rates, while images with more iterations receive lower learning rates.

Furthermore, in the local optimization, based on the overlapping relationships, each local connected image is assigned a weight that determines the number of training iterations it will undergo. This weight encourages prioritizing training the newly added images and their local overlapping images, thereby accelerating the rendering quality in regions captured more recently. To further enhance training efficiency and near real-time performance for each image, two strategies have been integrated into On-the-Fly GS: inspired by Adr-

¹See more at <https://yifeiyu225.github.io/on-the-flySfMv2.github.io/>

Gaussian [4], a load-balancing strategy is incorporated to mitigate the over-accumulation of 3D Gaussians, while simultaneously ensuring fast training of newly added regions; Drawing from insights from Taming-3DGS [8], we reimplement the back-propagation by transitioning from pixel-based parallelism to splatting-based parallel computation. In summary, this work makes the following contributions:

- We propose On-the-Fly GS, a novel progressive solution for near real-time 3DGS field optimization that enables robust online training during image acquisition.
- To achieve robust and near real-time progressive 3DGS training, self-adaptive methods are presented to determine the training iteration number of each image and to update the learning rate during each iteration.
- Our On-the-Fly GS reduces training time significantly on several benchmark datasets, offering the first practical step toward rapid 3DGS optimization during image capturing.

II. RELATED WORKS

Novel View Synthesis (NVS). Novel View Synthesis has been a central research topic in both computer vision and computer graphics for many years [27]. Traditionally, various approaches have been developed to achieve NVS. Light Field Rendering employs light field cameras to capture light field data, enabling novel view synthesis by interpolation [28], [29]. The Lumigraph method [30], [31] improves Light Field Rendering with greater flexibility and broader applicability. View Interpolation [32] utilizes multi-view images to interpolate among existing images, while it still employs a traditional photogrammetry method to generate a mesh model of the scene and renders images from arbitrary viewpoints using back-projection [33]. In recent years, the advent of Neural Radiance Fields (NeRF) [34] and 3D Gaussian Splatting (3DGS) [2] has introduced novel paradigms for NVS. NeRF, based on volumetric rendering, leverages a multi-layer perceptron (MLP) to implicitly represent a 3D scene. The original NeRF demonstrated remarkable performance in small-scale scenes (e.g., desk-size objects), subsequent NeRF variants further improved its performance by reducing training time [35], enhancing rendering efficiency [36] and quality [37], enabling large-scale scene modeling [38], etc. However, contemporary NeRF methods still face significant challenges in complex, unbounded outdoor scenes where achieving high-quality, high-resolution rendering often requires extensive training time, making real-time rendering infeasible. To address these limitations, Kerbl et al. [2] introduced 3DGS to explicitly represent 3D scenes.

3DGS fast training. Although 3DGS delivers high-quality and real-time rendering, its training process remains computationally intensive. Consequently, extensive research has been dedicated to accelerating the 3D Gaussians optimization. Several studies have sought to expedite 3DGS training by reducing the number of Gaussians in the scene representation [3]–[7]. Panagiotis et al. [3] proposed an efficient and resolution-aware Gaussians pruning strategy along with an adaptive adjustment method for the degree of spherical

harmonic coefficients. It reduces the number of Gaussians by more than 50% and saves the memory. ADR-Gaussian [4] introduced a load balancing strategy to optimize the distribution of Gaussians and prevent their excessive accumulation in specific regions, thereby reducing the number of Gaussians. EfficientGS [6] selectively restricted the increase in key Gaussians and implemented a pruning mechanism to remove redundant ones, retaining critical Gaussians while removing unnecessary ones. Similarly, LightGaussian [7] introduced a global importance-based pruning strategy, retaining critical and removing redundant Gaussians. Since 3DGS training requires optimizing the parameters of a large number of Gaussians, reducing 3D Gaussians can significantly enhance training speed and decrease memory overhead.

Some studies also focused on refining training strategies to improve computational efficiency. ADR-Gaussian [4] adopts an early culling strategy during forward-propagation, refining the influence range of each Gaussian. Taming 3DGS [8] implements a Gaussian-based parallel strategy for gradient accumulation and optimizes the computation of spherical harmonics and differentiable loss functions, thereby enhancing back-propagation efficiency. RetinaGS [9] and DOGS [10] address the challenge of large-scale scenes and high-resolution image data. They propose a parallel training method that divides the whole scene into several overlapping blocks, enabling parallel and accelerating training for large-scale scenes. Unlike methods that focus on reducing the number of Gaussians, these approaches optimize specific computational steps in the training pipeline, achieving significant speed improvements.

3DGS real-time training. Recent studies have explored real-time training of 3DGS. Before 3DGS, through integration with SLAM, NICE-SLAM [22] and Vox-Fusion [23] successfully achieve real-time NeRF training. In addition, Gaussian Splatting SLAM [21], MonoGS++ [24] and RTG-SLAM [20] have further improved real-time 3DGS training, delivering higher rendering quality. While these methods enable real-time 3DGS training, they struggle to handle discrete images that are not with spatially and temporally continuous. In contrast, our On-the-Fly GS, which can deal with images captured in an arbitrary way, employs a progressive framework for near real-time 3DGS training. It addresses the limitations of existing SLAM-based methods and extends the applicability of 3DGS training to more general use cases.

III. PRELIMINARIES

In this section, two key preliminaries of 3DGS [2] and On-the-fly SfM [26] are introduced.

3D Gaussian Splatting. 3DGS employs a large number of Gaussians to represent a real 3D scene [2]. The geometric properties of each Gaussian are defined using a 3D Gaussian distribution:

$$G(X) = e^{-\frac{1}{2}(X-X_0)^T \Sigma^{-1}(X-X_0)}, \quad (1)$$

Where X_0 represents the center position of a Gaussian, X represents the position of a sampled point, and Σ corresponds

to the covariance matrix of the Gaussian distribution. For each Gaussian, its covariance matrix Σ can be formulated by a rotation matrix R and a scaling matrix S , as follows:

$$\Sigma = RSS^T R^T. \quad (2)$$

The above parameters define the geometric properties of a Gaussian. Besides, 3DGS uses spherical harmonics (SH) and opacity σ to reveal the material properties. Before calculating the color of one specific pixel, 3DGS estimates the splatting opacity α based on the Gaussian-pixel 2D distance x and the projected covariance matrix Σ' :

$$\alpha = \sigma e^{-\frac{1}{2}x^T \Sigma'^{-1}x}. \quad (3)$$

Given the splatting opacities α_i and color c_i calculated from SH of N Gaussians associated with the current pixel, 3DGS calculates the color of the current pixel c_p by alpha blending, as follows:

$$c_{\text{pixel}} = \sum_{i=1}^N \left(c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j) \right) \quad (4)$$

During a novel view rendering, 3DGS assigns a dedicated thread to each pixel, enabling parallel computation of the color of each pixel to facilitate high-quality, real-time image rendering.

On-the-Fly SfM. Before training the original 3DGS field, COLMAP [19] is typically used for pose estimation and sparse point cloud generation. However, COLMAP requires feature matching across all images, making it impractical in scenarios where image acquisition and 3DGS training occur simultaneously, such as our On-the-Fly GS training. To overcome this limitation, instead of COLMAP, we use the On-the-Fly SfM [26] for our subsequent experiments. On-the-Fly SfM employs a Hierarchical Weighted Local Bundle Adjustment (BA), in which only neighboring connected and registered images are incorporated into BA after acquiring a new image. Unlike COLMAP’s BA, which incurs a substantial computational cost as the number of images increases, Hierarchical Weighted Local BA significantly reduces time complexity, ensuring efficient processing. Additionally, On-the-Fly SfM can achieve online feature matching via global feature retrieval on HNSW (Hierarchical Navigable Small World) [39], enabling efficient dynamic updates for the matching matrix. On-the-fly SfM allows for near real-time pose estimation and sparse point cloud generation immediately after capturing a new image, ensuring a continuous and timely input for our approach.

IV. METHODOLOGY

A. Overview of on-the-fly GS

To achieve near real-time 3DGS field optimization, the 3DGS field is supposed to be trained concurrently with image acquisition. Therefore, rather than relying on the original offline training solution of 3DGS, we propose a progressive training framework. As illustrated in Fig. 2, our training process is structured into three distinct phases:

Phase 1: Initial 3DGS Training. Training the 3DGS field with a limited number of images risks overfitting to the initial few images (e.g., 2 or 3 images) [40]. To mitigate this, we begin by training an initial 3DGS field after On-the-Fly SfM has estimated the poses and generated the sparse point cloud for N_0 initial images. This initial 3DGS field serves as the foundation for subsequent progressive training. In this phase, we adopt the original 3DGS training method but set a more severe number for training iterations to prevent overfitting to these initial images.

Phase 2: Progressive 3DGS Training. As a new image is acquired, On-the-Fly SfM continuously updates the camera poses of both the newly acquired and previously registered images, while refining the sparse point cloud. Based on the updated sparse point cloud, some new Gaussians are appended to the previous 3DGS field. The updated model will be then optimized. The optimization process contains local optimization and global optimization. The local optimization focuses on fast optimization of the newly added regions of the 3DGS field, while the global optimization ensures the overall quality and consistency of the 3DGS field and prevents overfitting to the newly added images.

Phase 3: Final Optimization. Once all images have been acquired and the progressive training is complete, a final optimization is performed on the current 3DGS field. This step further refines the model and enhances its overall quality, ensuring the best reconstruction result.

B. 3D Gaussian field Update with newly fly-in images

For the newly fly-in image, On-the-Fly SfM estimates the pose of the new image and optimizes the pose of its neighboring connected images via an improved local BA [26], simultaneously updating the sparse point cloud. Before updating the previously trained 3DGS field, a search tree is constructed for the existing sparse point cloud. The sparse points from the new image are then queried against this search tree using a predefined threshold. Points in the new image whose shortest distance from the existing points exceeds the threshold are classified as newly added points. These newly added points are then initialized as newly added Gaussians and directly incorporated into the current 3DGS field, facilitating the 3DGS field training on the newly expanded region.

C. Self-adaptive Training Iteration Allocation via Hierarchical Image Weighting

Newly added images often introduce novel unexplored regions or correspond to areas where the current 3DGS field exhibits poor rendering quality. To ensure that more training is allocated to the new image and its neighboring connected images during local optimization, we propose an image hierarchical weighting strategy, which can be used to adaptively allocate training iterations for each image. The weight is determined by the overlap relationship between the newly fly-in image and the already registered images, as depicted in Fig. 3.

Assuming that there are $n - 1$ registered images, when the image n (I_{11} , $l = 1$) is newly added, the weight of

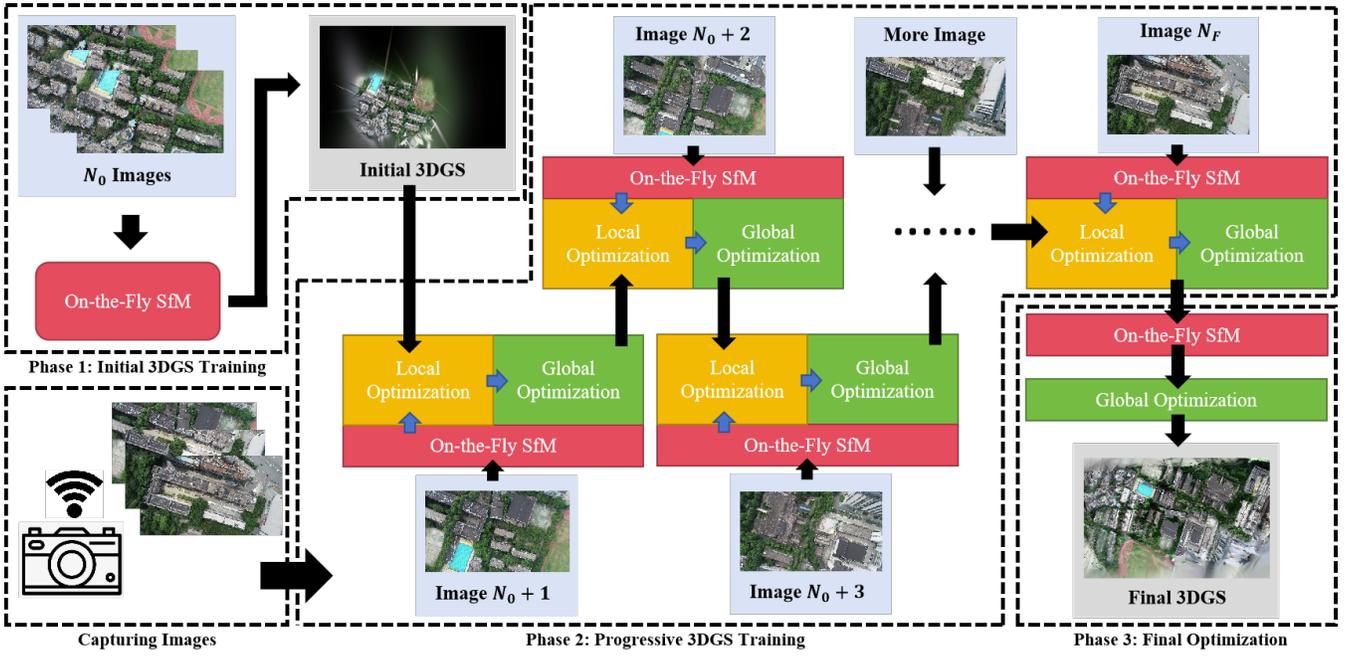


Fig. 2. Workflow of the proposed On-the-Fly GS.

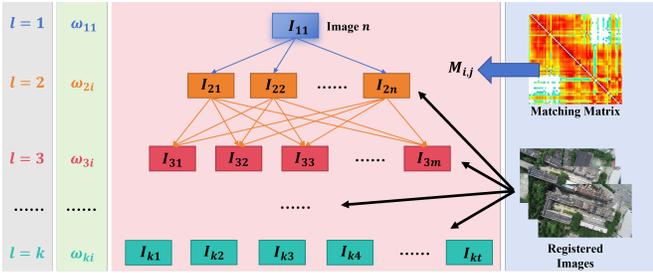


Fig. 3. The Hierarchical Image Weighting.

the newly added image is fixed as $\omega_{11} = 1$. The weights of the remaining registered images are calculated based on the match matrix online estimated by On-the-Fly SfM. If the overlapping degree of any two images is $M_{i,j}$, for images I_{2i} that have a direct overlapping relationship ($M_{i,n} \neq 0, l = 2$) with the newly added image, their weights ω_{2i} are calculated using the following equation:

$$\omega_{2i} = \omega_{11} \times M_{i,n} = M_{i,n}. \quad (5)$$

For images I_{li} that have an indirect overlapping relationship ($M_{i,n} = 0, l \geq 3$) with the newly added image, their weights are determined by their overlapping relationship with images in the previous layer and the weights of those images. If the layer $k - 1$ contains D images, the weight of the image j in the layer $k - 1$ is $\omega_{(k-1)j}$, the weight of the corresponding image i in layer k will be calculated as follows:

$$\omega_{ki} = \frac{\sum_{j=1}^D \omega_{(k-1)j} \times M_{(k-1)j,ki}}{D} \quad (6)$$

Based on the hierarchical weighting method, we can adaptively assign a weight to all registered images including the newly added image. Subsequently, the training iterations for local neighboring images will be allocated based on their

weights. For N images, with given weight ω_i of image i , the training iterations T_i can be calculated as follows:

$$T_i = \frac{\omega_i}{\sum_{j=1}^N \omega_j} \times \text{LOT} \quad (7)$$

LOT represents the number of training iterations required during each local optimization. This hierarchical weighting method assigns higher weights for the newly added image and its neighboring connected images while the remaining images are given lower weights, ensuring the newly added image and its neighboring connected images receive more training in the current local optimization.

D. Adaptively Learning Rate Updating

The progressive training in On-the-Fly GS may result in inconsistency in local quality within the 3DGS field. To address this issue, we dynamically update the learning rate based on the rendering quality of each image, instead of updating the learning rate based on the overall quality of the 3DGS field (this was applied in the original 3DGS). Specifically, for image j , if it has been trained for IT_j iterations, with the initial learning rate ILr and the final learning rate FLr , the current learning rate LogLr_j for the image j can be calculated as follows:

$$\text{LogLr}_j = e^{\ln(\text{ILr}) \times (1-t_i) + \ln(\text{FLr}) \times t_i}, \quad (8)$$

$$t_i = \begin{cases} 1, & \text{IT}_j > \text{Mean} \\ \frac{\text{IT}_j}{\text{Mean}}, & \text{IT}_j \leq \text{Mean} \end{cases}, \quad (9)$$

where Mean is a predefined parameter that represents the number of training iterations required for an image to achieve satisfactory rendering quality. However, determining the learning rate based on the already trained iterations of a single image is not rigorous. This is because other

images overlapping with image j might already achieve satisfactory rendering quality. Thus, when determining the current learning rate, it makes sense to also consider the overlapping images among the registered images. If the set of overlapping images O_I for the current image j contains RN images (i.e., for any $i \in O_I, M_{i,j} > 0$). The learning rate Lr_j for image j can be calculated as follows:

$$Lr_j = \frac{\text{Log}Lr_j + \sum_{i=1}^{RN} M_{i,j} \times \text{Log}Lr_i}{RN + 1}. \quad (10)$$

The proposed learning rate updating method considers the rendering quality of the 3DGS field within the local scene. If the local scene corresponding to the current image j has achieved high rendering quality, the learning rate Lr_j is supposed to be lower. This is because multiple images overlapping with image j have been trained for more iterations (even if image j is newly added), as shown in Fig. 4 (left). If the local scene corresponding to the image j performs poor rendering quality, which indicates that both image j and its neighboring connected images are only trained for fewer iterations, the learning rate Lr_j should be higher, as shown in Fig. 4 (right).

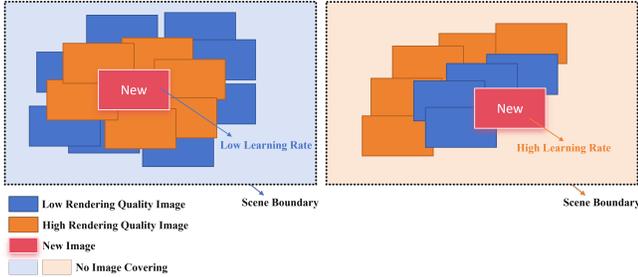


Fig. 4. Adaptively Learning Rate Updating

E. Load Balancing and 3DGS Runtime Optimization

Some local areas may undergo excessive Gaussian accumulation due to being visible across multiple images, while newly introduced scenes may suffer from insufficient Gaussian kernels. To address this uneven distribution of Gaussians, inspired by ADR-Gaussian [4], we employ a load balancing training. For the current training image i , we record the number of Gaussians GN_p associated with each pixel p in image i . Using the standard deviation of Gaussian distribution within the visible range of the image, we define a Load-Balancing-Loss L_{load} as follows:

$$L_{load} = \text{std}(GN_p), p \in i. \quad (11)$$

With L1-Loss and SSIM-Loss, the total loss is as follows:

$$\text{LOSS} = \lambda_1 L_1 + \lambda_{\text{ssim}} L_{\text{ssim}} + \lambda_{\text{load}} L_{\text{load}}. \quad (12)$$

Subsequent experiments demonstrate that the load balancing strategy not only prevents the abnormal accumulation of Gaussians in some specific regions but also effectively reduces the total number of Gaussians in the 3DGS field,

thereby accelerating the overall training process. Furthermore, it facilitates the densification of Gaussians in newly added regions, promoting faster 3DGS training.

To further accelerate the training of the 3DGS field, we integrate the strategy proposed by Taming 3DGS [8]. Taming 3DGS identified that the per-pixel gradient accumulation method used in the original 3DGS leads to multiple threads contending for accessing the same locations and thus serialized atomic operations. To address this, we implemented the corresponding Splatting-based parallelism in the gradient accumulation method. In this implementation, gradients are accumulated directly for each Gaussian, significantly reducing training time by eliminating contention and enhancing parallel processing efficiency.

V. EXPERIMENT

A. Experimental Settings

Datasets and Metrics. To validate the efficacy of our On-the-Fly GS, we utilize several datasets from the 3DGS [2], Mip-NeRF 360 [41], On-the-Fly SfM [25], [26], and TUM-RGBD [42]. This contains three indoor scenes, two small-scale outdoor scenes, and one UAV image, which are dynamically processed as simulated sequential inputs according to the original stored order. Based on these datasets, the corresponding performance is evaluated from three aspects: **1) Training efficiency.** Since On-the-Fly GS is capable of simultaneous estimation of image poses and 3DGS training, we record the time consumption of different methods starting from the moment image pose estimation begins until the final 3DGS field is obtained. **2) Rendering quality.** We use the same common metrics as those used in 3DGS, including PNSR, SSIM and LPIPS. **3) Video Memory.** We analyse the trained 3DGS field memory consumption of each method. In addition, for On-the-Fly GS, we record the time of the complete training process and evaluate the rendering quality. To compare with other offline-training methods, we first investigate the training time required to achieve the same rendering quality as our On-the-Fly GS, followed by exploring their default training iteration settings. All experiments are conducted under the same hardware configuration (Single NVIDIA RTX 4090 GPU).

Implementation details. The optimization process of our On-the-Fly GS includes three stages (see Fig. 2). In phase 1 of initialization, after On-the-Fly SfM completes pose estimation and sparse point cloud generation for the first $N_0 = 16$ images, we perform 2,000 training iterations on these images. In phase 2, for each newly acquired image, we first run 100 iterations of local optimization, followed by 100 iterations of global optimization. Additionally, the preset parameter Mean is 195, ILr is $1.6e - 4$ and FLr is $1.6e - 6$. In the last phase 3, an additional 1,000 iterations of fast global optimization are performed to further improve the overall model quality. The weight of Load-Balancing-Loss λ_{load} , L1-Loss λ_1 and SSIM-Loss λ_{ssim} are set to 0.41, 0.47 and 0.12, respectively.

TABLE I

COMPARISON WITH OFFLINE SOLUTIONS AT SAME RENDERING PERFORMANCE AS ON-THE-FLY GS. BEST AND SECOND-BEST ARE HIGHLIGHTED IN BOLD AND UNDERLINE.

Dataset	ImageNum	Methods	Video Memory (MB)	PSNR	COLMAP/ On-the-Fly SfM (min)	3DGS Training (min)	Total Time (min)
Counter	240	3DGS	218.07	27.68	16.92	4.91	21.83
		Taming-3DGS	221.38			1.38	<u>18.20</u>
		Adr-Gaussian	121.56			3.76	20.68
		Ours-Offline GS	119.06			3.14	20.06
		On-the-Fly GS	118.36			\	14.85+0.48
Train	301	3DGS	128.15	21.01	16.39	2.82	19.21
		Taming-3DGS	135.44			1.31	17.70
		Adr-Gaussian	69.73			2.18	18.57
		Ours-Offline GS	75.61			1.44	<u>17.83</u>
		On-the-Fly GS	159.68			19.21	\
YD	125	3DGS	2124.37	22.04	12.55	13.58	26.13
		Taming-3DGS	2183.01			8.31	20.85
		Adr-Gaussian	1384.139			15.27	27.82
		Ours-Offline GS	1084.87			6.48	<u>19.03</u>
		On-the-Fly GS	699.34			7.70	\
Palace	210	3DGS	100.95	21.41	14.17	10.49	24.66
		Taming-3DGS	118.36			4.02	18.19
		Adr-Gaussian	48.68			9.65	23.82
		Ours-Offline GS	62.01			3.46	<u>17.63</u>
		On-the-Fly GS	173.94			16.54	\

B. Comparisons with other SOTA methods

We compare the proposed On-the-Fly GS with four offline-training solutions (3DGS [2], Adr-Gaussian [4], Taming-3DGS [8] and Ours-Offline GS training). Ours-Offline GS refers to an offline training pipeline that incorporates Load Balancing and Splatting-based parallelism in backpropagation strategies, while keeping all other settings identical to 3DGS. COLMAP is employed for pose estimation and sparse point cloud generation for the offline-training methods, while On-the-Fly SfM is employed to estimate pose and generate sparse point cloud for our On-the-Fly GS.

Table I reports the corresponding metrics when the four offline training solutions achieve the same rendering performance as that can be generated by our On-the-Fly GS. It can be seen that On-the-Fly GS shows a significant advantage in terms of cost time. Apart from the time (before the +) required for On-the-Fly SfM to do online pose estimation and sparse point cloud generation, only a small amount of extra overload (after the +) is needed to complete the 3DGS training. Generally, it takes around 2-3 seconds for a newly fly-in image to finish the 3DGS optimization, which is sufficient for near real-time performance and demonstrates the capability of simultaneous image acquisition, pose estimation, and 3DGS training. On the other hand, the offline training solutions need to finish the Colmap processing before training. They also typically take longer for both pose estimation and 3DGS training, further leading to more total time for offline training. Furthermore, the model obtained through On-the-Fly GS has the smallest model size in some scenarios, which means that the proposed method is promising to enable lightweight 3DGS training.

To further investigate the discrepancy between our work and these offline solutions with default setting, the number

TABLE II

COMPARISON WITH OFFLINE SOLUTIONS OF DEFAULT SETTING. BEST AND SECOND-BEST ARE HIGHLIGHTED IN BOLD AND UNDERLINE.

Methods	Total Time (min)				PSNR			
	Counter	Train	YD	Palace	Counter	Train	YD	Palace
3DGS	37.35	26.06	25.03	39.16	30.37	24.64	22.09	23.98
Taming-3DGS	24.66	21.39	20.85	23.89	30.25	24.31	22.15	23.71
Adr-Gaussian	28.68	21.71	45.19	28.09	29.57	23.07	23.57	22.19
Ours-Offline GS	<u>24.16</u>	<u>19.98</u>	<u>34.71</u>	<u>21.93</u>	29.75	22.79	24.13	23.83
On-the-Fly GS	15.33	19.57	8.25	16.92	27.68	21.01	22.04	21.41

of iterations for these offline methods is set to 30,000. As shown in Tab. II and Fig. 5, in general, the overall rendering quality of our On-the-Fly GS is around 80%–90% of that generated by these offline solutions. Nevertheless, our On-the-Fly GS demonstrates a significant advantage in terms of time efficiency, while also producing superior results in certain rendered images. Furthermore, as YD is a UAV image dataset, the 3DGS and Taming-3DGS can only run 11,000 and 10,500 iterations due to excessive memory consumption (the results presented in Tab. II and Fig. 5 correspond to this specific training setting), which also prove the effectiveness of our method in handling model size.

On-the-fly SfM [25] can also process sequential video frames, thus, we further tested two SLAM benchmarks and compared with several real-time NVS methods based on SLAM, including NICE-SLAM [22], Vox-Fusion [23], GS-SLAM [21], and MonoGS++ [24]. Quantitative comparisons are shown in Table III. On-the-Fly GS typically performs the best or second best on the evaluation metrics.

C. Ablation Studies

To more explicitly assess the efficacy of each component implemented in our method, we conduct an extensive

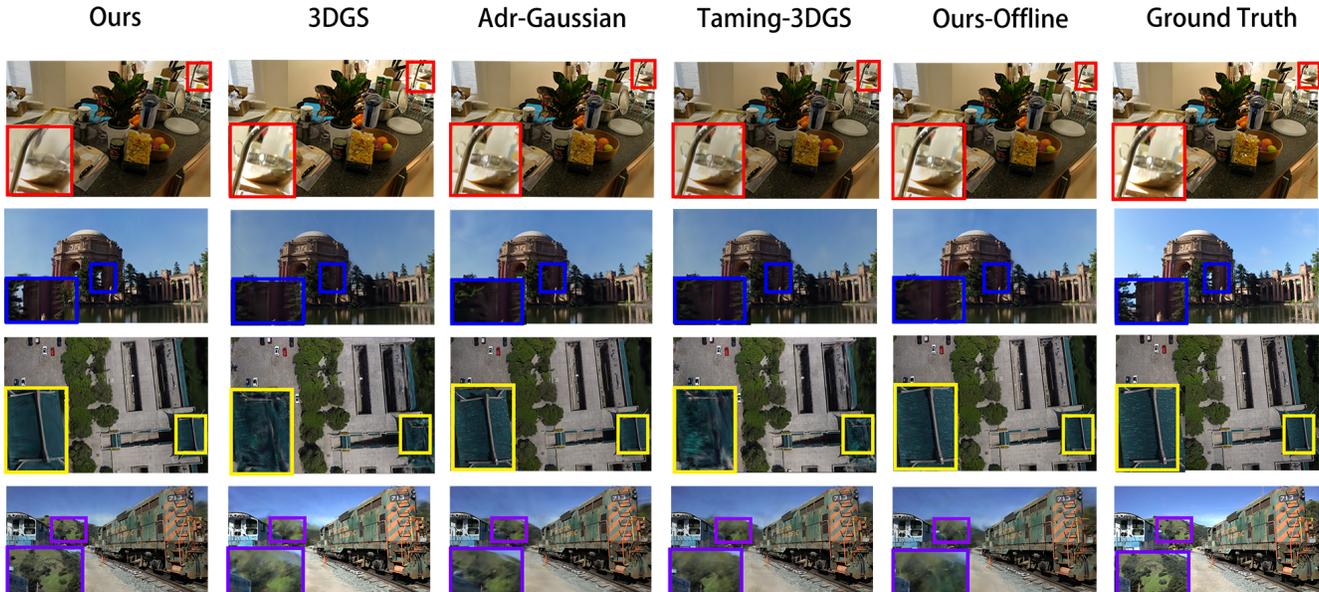


Fig. 5. Rendering Results of On-the-Fly GS and Offline solutions with default settings

TABLE III

COMPARISON WITH SLAM-BASED METHOD. BEST AND SECOND-BEST ARE HIGHLIGHTED IN BOLD AND UNDERLINE.

Methods	fr1_room			fr1_desk		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
NICE-SLAM	11.39	0.37	0.62	13.83	0.56	0.48
Vox-Fusion	14.20	0.56	0.55	15.79	0.64	0.52
MonoGS	14.76	0.52	0.60	17.31	0.65	<u>0.38</u>
MonoGS++	22.16	0.77	0.31	22.85	0.82	0.60
Ours	22.17	0.84	<u>0.32</u>	24.23	0.86	0.26

TABLE IV

ABLATION EXPERIMENT OF THE PROPOSAL COMPONENTS IN THE ON-THE-FLY GS METHOD.

Methods	Time Cost (min)	Video Memory (MB)	PSNR	SSIM	LPIPS
No_IW	11.21	123.7	21.04	0.83	0.36
No_GO	12.96	271.13	16.62	0.69	0.48
No_Load	14.02	400.47	21.03	0.84	0.33
No_SPB	18.29	182.96	<u>21.21</u>	0.86	0.36
Ours	11.19	<u>173.94</u>	21.24	0.86	0.36

ablation study as follows. Based on the *Palace* dataset, we first process it with our full On-the-Fly GS and then sequentially switch off one of the following training strategies: 1) Image Weighting and Training Iteration Allocation (No_IW). 2) Global Optimization (No_GO). 3) Load Balancing (No_Load). 4) Splatting-based Parallelism in Backpropagation (No_SPB). In particular, the cost time in this context refers to the total time required for 3DGS model training after the On-the-Fly SfM completes the pose estimation for each newly fly-in image. The corresponding quantitative results are presented in Table IV and Fig. 6.

Based on Table IV and Fig. 6, switching off any of the proposed components can negatively impact the 3DGS training, validating their corresponding effectiveness: 1) Image Weighting and Training Iteration Allocation improves the final rendering quality and also slightly reduces time consumption. 2) Global Optimization reduces the memory

and time consumption while enhancing model quality. 3) Load Balancing primarily works on memory saving, whereas Splatting-based Parallelism in Backpropagation accelerates training speed. Overall, these ablation study results demonstrate the effectiveness of all the proposed strategies.



Fig. 6. Rendering Results of Ablation Studies

VI. CONCLUSION

In this paper, we propose On-the-Fly GS, a novel near real-time 3DGS training framework via a progressive optimization manner, which can dynamically output 3DGS field during image capturing. First, we introduce a method for weighting images and allocating different training iterations to each newly captured image during local optimization, enabling fast training for the newly expanded region. Second, we present a new learning rate updating solution based on image rendering quality, adaptively calculating the learning rate for each training iteration based on the training progress of the current image and its neighboring connected images. Finally, inspired by Adr-Gaussian and Taming 3DGS, we further control the number of Gaussians and improve training efficiency. The conducted experiments demonstrate that our On-the-Fly GS successfully achieves near real-time 3DGS training and considerable rendering performance, even when the input images lack spatial and temporal continuity.

REFERENCES

- [1] Y. Chen, Z. Chen, C. Zhang, *et al.*, “Gaussianeditor: Swift and controllable 3d editing with gaussian splatting,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024, pp. 21 476–21 485.
- [2] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, “3d gaussian splatting for real-time radiance field rendering,” *ACM Trans. Graph.*, vol. 42, no. 4, pp. 139–1, 2023.
- [3] P. Papantonakis, G. Kopanas, B. Kerbl, A. Lanvin, and G. Drettakis, “Reducing the memory footprint of 3d gaussian splatting,” *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, vol. 7, no. 1, pp. 1–17, 2024.
- [4] X. Wang, R. Yi, and L. Ma, “Adr-gaussian: Accelerating gaussian splatting with adaptive radius,” in *SIGGRAPH Asia 2024 Conference Papers*, 2024, pp. 1–10.
- [5] W. Liu, T. Guan, B. Zhu, *et al.*, “Efficientgs: Streamlining gaussian splatting for large-scale high-resolution scene representation,” *IEEE MultiMedia*, 2025.
- [6] Z. Fan, K. Wang, K. Wen, Z. Zhu, D. Xu, Z. Wang, *et al.*, “Light-gaussian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps,” *Advances in neural information processing systems*, vol. 37, pp. 140 138–140 158, 2025.
- [7] J. C. Lee, D. Rho, X. Sun, J. H. Ko, and E. Park, “Compact 3d gaussian representation for radiance field,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 21 719–21 728.
- [8] S. S. Mallick, R. Goel, B. Kerbl, *et al.*, “Taming 3dgs: High-quality radiance fields with limited resources,” in *SIGGRAPH Asia 2024 Conference Papers*, 2024, pp. 1–11.
- [9] B. Li, S. Chen, L. Wang, *et al.*, “Retinags: Scalable training for dense scene rendering with billion-scale 3d gaussians,” *arXiv preprint arXiv:2406.11836*, 2024.
- [10] Y. Chen and G. H. Lee, “Dogs: Distributed-oriented gaussian splatting for large-scale 3d reconstruction via gaussian consensus,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 34 487–34 512, 2025.
- [11] S. Durvasula, A. Zhao, F. Chen, *et al.*, “Distwar: Fast differentiable rendering on raster-based rendering pipelines,” *arXiv preprint arXiv:2401.05345*, 2023.
- [12] J. Jo, H. Kim, and J. Park, “Identifying unnecessary 3d gaussians using clustering for fast rendering of 3d gaussian splatting,” *arXiv preprint arXiv:2402.13827*, 2024.
- [13] J. Lee, S. Lee, J. Lee, J. Park, and J. Sim, “Gscore: Efficient radiance field rendering via architectural support for 3d gaussian splatting,” in *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*, 2024, pp. 497–511.
- [14] K. Cheng, X. Long, K. Yang, *et al.*, “Gaussianpro: 3d gaussian splatting with progressive propagation,” in *Forty-first International Conference on Machine Learning*, 2024.
- [15] J. Zhang, F. Zhan, M. Xu, S. Lu, and E. Xing, “Fregs: 3d gaussian splatting with progressive frequency regularization,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 21 424–21 433.
- [16] Z. Yu, A. Chen, B. Huang, T. Sattler, and A. Geiger, “Mip-splatting: Alias-free 3d gaussian splatting,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024, pp. 19 447–19 456.
- [17] Y. Jiang, J. Tu, Y. Liu, *et al.*, “Gaussianshader: 3d gaussian splatting with shading functions for reflective surfaces,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 5322–5332.
- [18] X. Wu, J. Xu, C. Wang, *et al.*, “Local gaussian density mixtures for unstructured lumigraph rendering,” in *SIGGRAPH Asia 2024 Conference Papers*, 2024, pp. 1–11.
- [19] J. L. Schonberger and J.-M. Frahm, “Structure-from-motion revisited,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4104–4113.
- [20] Z. Peng, T. Shao, Y. Liu, *et al.*, “Rtg-slam: Real-time 3d reconstruction at scale using gaussian splatting,” in *ACM SIGGRAPH 2024 Conference Papers*, 2024, pp. 1–11.
- [21] H. Matsuki, R. Murai, P. H. Kelly, and A. J. Davison, “Gaussian splatting slam,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 18 039–18 048.
- [22] Z. Zhu, S. Peng, V. Larsson, *et al.*, “Nice-slam: Neural implicit scalable encoding for slam,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 12 786–12 796.
- [23] X. Yang, H. Li, H. Zhai, *et al.*, “Vox-fusion: Dense tracking and mapping with voxel-based neural implicit representation,” in *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 2022, pp. 499–507.
- [24] R.-W. Li, W. Ke, D. Li, L. Tian, and E. Barsoum, “Monogs++: Fast and accurate monocular rgb gaussian slam,” in *35th British Machine Vision Conference 2024, BMVC 2024, Glasgow, UK, November 25-28, 2024*. BMVA, 2024.
- [25] Z. Zhan, R. Xia, Y. Yu, *et al.*, “On-the-fly sfm: What you capture is what you get,” *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. X-1-2024, pp. 297–304, 2024.
- [26] Z. Zhan, Y. Yu, R. Xia, *et al.*, “Sfm on-the-fly: Get better 3d from what you capture,” *arXiv preprint arXiv:2407.03939*, 2024.
- [27] Y. Bao, T. Ding, J. Huo, *et al.*, “3d gaussian splatting: Survey, technologies, challenges, and opportunities,” *IEEE Transactions on Circuits and Systems for Video Technology*, 2025.
- [28] M. Levoy and P. Hanrahan, “Light field rendering,” in *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, 2023, pp. 441–452.
- [29] A. Davis, M. Levoy, and F. Durand, “Unstructured light fields,” in *Computer Graphics Forum*, vol. 31, no. 2pt1. Wiley Online Library, 2012, pp. 305–314.
- [30] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen, “The lumigraph,” in *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, 2023, pp. 453–464.
- [31] C. Buehler, M. Bosse, L. McMillan, S. Gortler, and M. Cohen, “Unstructured lumigraph rendering,” in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 2001, pp. 425–432.
- [32] S. E. Chen and L. Williams, “View interpolation for image synthesis,” in *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, 2023, pp. 423–432.
- [33] Z. Ma and S. Liu, “A review of 3d reconstruction techniques in civil engineering and their applications,” *Advanced Engineering Informatics*, vol. 37, pp. 163–174, 2018.
- [34] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “Nerf: Representing scenes as neural radiance fields for view synthesis,” *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [35] T. Müller, A. Evans, C. Schied, and A. Keller, “Instant neural graphics primitives with a multiresolution hash encoding,” *ACM transactions on graphics (TOG)*, vol. 41, no. 4, pp. 1–15, 2022.
- [36] Z. Shu, R. Yi, Y. Meng, Y. Wu, and L. Ma, “Rt-octree: accelerate plenotree rendering with batched regular tracking and neural denoising for real-time neural radiance fields,” in *SIGGRAPH Asia 2023 Conference Papers*, 2023, pp. 1–11.
- [37] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman, “Mip-nerf 360: Unbounded anti-aliased neural radiance fields,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 5470–5479.
- [38] H. Turki, D. Ramanan, and M. Satyanarayanan, “Mega-nerf: Scalable construction of large-scale nerfs for virtual fly-throughs,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 12 922–12 931.
- [39] Y. A. Malkov and D. A. Yashunin, “Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 4, pp. 824–836, 2018.
- [40] H. Xiong, *SparseGS: Real-time 360° sparse view synthesis using Gaussian splatting*. University of California, Los Angeles, 2024.
- [41] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman, “Mip-nerf 360: Unbounded anti-aliased neural radiance fields,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 5470–5479.
- [42] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, “A benchmark for the evaluation of rgb-d slam systems,” in *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2012, pp. 573–580.