

Mini-Splatting: Representing Scenes with a Constrained Number of Gaussians

Guangchi Fang¹ and Bing Wang^{1*}

Spatial Intelligence Group, The Hong Kong Polytechnic University
guangchi.fang@gmail.com, bingwang@polyu.edu.hk




Fig. 1: By reorganizing the spatial distribution of 3D Gaussians [17], our Mini-Splatting reduces the number of Gaussians (in millions) while enhancing model performance in terms of rendering speed, training time, and rendering quality. The Gaussian centers of a foreground object *bicycle* are projected onto the rendered image as blue points, demonstrating the more uniform spatial distribution achieved by our algorithm.

Abstract. In this study, we explore the challenge of efficiently representing scenes with a constrained number of Gaussians. Our analysis shifts from traditional graphics and 2D computer vision to the perspective of point clouds, highlighting the inefficient spatial distribution of Gaussian representation as a key limitation in model performance. To address this, we introduce strategies for densification including blur split and depth reinitialization, and simplification through intersection preserving and sampling. These techniques reorganize the spatial positions of the Gaussians, resulting in significant improvements across various datasets and benchmarks in terms of rendering quality, resource consumption, and storage compression. Our Mini-Splatting integrates seamlessly with the original rasterization pipeline, providing a strong baseline for future research in Gaussian-Splatting-based works. Code is available.

Keywords: Gaussian Splatting · Point Clouds · Scene Representation

1 Introduction

Recent advancements in 3D Gaussian Splatting (3DGS) [17] have showcased its significant potential across a range of applications, such as immersive rendering and 3D reconstruction, attributed to its capabilities for real-time and high-quality rendering. Unlike traditional radiance field-based algorithms [1, 2, 24],

* Corresponding author

3DGS employs millions of elliptical Gaussians for scene modeling (*e.g.*, 1 to 6 million Gaussians for each scene in the Mip-NeRF360 dataset [2]). The detailed characterization of these Gaussians, including their 3D position, opacity, covariance, and spherical harmonic (SH) coefficients, allows for a remarkably detailed and realistic scene representation. However, the extensive use of Gaussians also leads to inefficiencies, as depicted in Fig. 1, where Gaussian centers are prone to clustering together. This suboptimal spatial distribution can constrain the rendering quality and speed. In this regard, our study aims to address the challenge of efficiently representing a scene with a constrained number of Gaussians.

The optimization of minimal Gaussian representation remains a relatively underexplored topic, with only a handful of compression-focused studies venturing into this area. Drawing inspiration from voxel pruning techniques used in grid-based neural representations [21, 40], Niedermayr *et al.* [27] proposed pruning Gaussians based on parameter sensitivity derived from image gradients. Similarly, Lee *et al.* [20] and LightGaussian [11] have implemented pruning strategies based on Gaussian opacity and scale. However, these methods, which primarily aim at storage compression, often neglect the inefficient spatial distribution of input Gaussians. As a result, direct pruning, as practiced in these approaches, tends to yield suboptimal simplification outcomes. Moreover, most existing 3D data simplification algorithms focus on preserving feature points [16, 19] or geometric structures [23, 29], making them less suitable for explicit representations with densely optimizable parameters like 3DGS. Therefore, adapting these techniques to the 3DGS model presents a non-trivial challenge.

In this paper, we start with a comprehensive analysis of Gaussian representation. By visualizing the projected Gaussian centers and rendered results, we find that 3D Gaussians are not uniformly distributed around the object surface, which limits the rendering speed and quality of the model. To construct an efficient Gaussian representation, we pivot the task of directly simplifying 3D Gaussians to reorganizing their spatial positions employing our Gaussian densification and simplification method. Specifically, our proposed densification algorithm includes two key components: blur split, which involves splitting Gaussians in the blurring areas, and depth reinitialization, which uses merged depth points to reinitialize the scene. Our simplification approach comprises intersection preserving, which preserve Gaussians with maximal impact on the rendered image, and sampling, designed to preserve both the geometric structure and rendering quality of the original model. By stacking our densification and simplification techniques, we achieve a more uniform and efficient spatial distribution of Gaussians, as illustrated in Fig. 1. Our main contributions are as follows:

- Through a deep analysis of the spatial distribution of Gaussians, we observe that the vanilla 3DGS presents two significant phenomena: ‘overlapping’ and ‘under-reconstruction’. We identify that this inefficient distribution limits the rendering quality and speed, and it is non-trivial to achieve a minimal Gaussian representation while maintaining rendering quality.
- We propose a Gaussian densification and simplification algorithm to reorganize the spatial positions of Gaussians rather than directly pruning them.

The proposed densification (blur split and depth reinitialization) and simplification (intersection preserving and Gaussian sampling) strategies, leveraging both screen-space and world-space information, effectively achieve a dense distribution of Gaussians around the object.

- By integrating densification and simplification with corresponding further processing, our proposed Mini-Splatting achieves a balanced trade-off between rendering quality, resource consumption, and storage. Extensive experimental results on multiple benchmarks and datasets demonstrate the potential and scalability of our method.

2 Related Work

2.1 Gaussian Splatting

3D Gaussian Splatting (3DGS) [17] represents a recent advancement in novel view synthesis, enabling real-time high-quality rendering. Over the past year, numerous follow-up works have emerged. The concept of 3DGS has been applied to domains such as autonomous driving [39], human avatar [31], content generation [34,35] and dynamic scene [22], resulting in advancements within each domain [4]. While several works aim to enhance 3DGS from the perspectives of graphics [12], image processing [41], and system performance [10], few studies focus on delving into the nature of Gaussians as the scene representation.

The Gaussian representation constitutes a key distinction between 3DGS and previous standard techniques for novel view rendering. While most prior works employ implicit [1, 2, 24] or explicit [33, 38, 40] neural representations with volume rendering based on ray marching, the primary challenge lies in the dense ray point sampling required by ray marching, which consequently limits both training and rendering speed. In contrast, 3DGS utilizes a large collection of elliptical Gaussians [42] for scene representation and rasterize these Gaussians onto the image plane for alpha blending. This carefully designed rasterization pipeline circumvents the issues associated with ray marching. However, the performance of 3DGS models is constrained by the large number of Gaussians due to the rasterization process.

2.2 3D Data Simplification

For point cloud simplification purposes, conventional techniques such as farthest point sampling [30], random sampling [15], and grid sampling [13] are commonly utilized in point cloud processing tasks. In contrast, most works oriented towards point cloud simplification focus on preserving either feature points [5, 16, 19] or geometric structures [23, 25, 29]. Several learning-based algorithms [9, 19, 36] have been proposed to sample points during training, integrating downstream tasks in an end-to-end manner. However, it is non-trivial to directly apply these sampling techniques to 3DGS due to differences in representations.

With recent advancements in neural radiance fields, several works have been proposed to simplify neural representations, primarily focusing on feature grids

and introducing voxel pruning [7,21] or voxel masking [32,37] strategies. Inspired by these developments, similar pruning techniques have been adapted in recent 3DGS compression-oriented works [11,20,27]. For instance, Niedermayr et al. [27] proposed to prune Gaussians based on image gradients, while Lee et al. [20] and LightGaussian [11] performed pruning according to Gaussian opacity and scale. However, these works overlook the issue of inefficient spatial distribution of Gaussians, which can easily lead to suboptimal simplification results after Gaussian pruning. One concurrent work, RadSplat [28], also share a similar technical pipeline with our algorithm.

3 Analyzing 3DGS from a Perspective of Point Clouds

3.1 Preliminaries of Gaussian Splatting

Gaussian Representation. 3D Gaussian Splatting (3DGS) [17] models the scene geometry as a set of 3D Gaussians $\{\mathcal{G}_i | i = 1, \dots, N\}$. Each Gaussian \mathcal{G}_i is characterized by its opacity $\alpha_i \in [0, 1]$, center $\mathbf{p}_i \in \mathbb{R}^{3 \times 1}$ and covariance matrix in world space $\Sigma_i \in \mathbb{R}^{3 \times 3}$ [42] as $\mathcal{G}_i(\mathbf{x}) = e^{-\frac{1}{2}(\mathbf{x}-\mathbf{p}_i)^T \Sigma_i^{-1}(\mathbf{x}-\mathbf{p}_i)}$. To render an image from these Gaussians, 3DGS sorts them in an approximate depth order based on the distance from their centers \mathbf{p}_i to the image plane, and further applies alpha-blending [24] as follows:

$$\mathbf{c}(\mathbf{x}) = \sum_{i=1}^N w_i \cdot \mathbf{c}_i, \quad \text{where } w_i = T_i \cdot \alpha_i \cdot \mathcal{G}_i^{2D}(\mathbf{x}), \quad T_i = \prod_{j=1}^{i-1} (1 - w_j), \quad (1)$$

where \mathbf{c}_i represents the view-dependent color modeled by spherical harmonic (SH) coefficients, and \mathcal{G}_i^{2D} is the projected 2D Gaussian distribution of \mathcal{G}_i through a local affine transformation [42].

Adaptive Density Control of Gaussians. In 3DGS, sparse points from Structure-from-Motion (SfM) are first initialized as Gaussians. During the optimization, these Gaussians undergo an adaptive density control strategy, allowing the dynamic addition and removal of individual Gaussians. Specifically, Gaussians with a significant gradient are identified and are either split or cloned depending on their scales. Additionally, Gaussians with low opacity or an overlarge screen space size are pruned from the representation. Ablation studies presented in [17] demonstrate that this clone and split strategy can control the density of Gaussians, which in turn benefits scene recovery.

3.2 Challenges in Gaussian Simplification

Visual Analysis of Gaussians. In 3DGS, scene geometry is represented using elliptical Gaussian primitives. Ideally, these Gaussians should be distributed around the geometric surface to accurately model the scene, particularly for the foreground Gaussians with relatively small 3D scales. Under this assumption, the




Fig. 2: Visual analysis of Gaussian centers. (a): Visualization of the original 3DGS model’s projected Gaussian centers, along with its corresponding rendering quality (PSNR in dB) and the number of Gaussians (in millions). Notably, we observe the phenomena of ‘overlapping’ and ‘under-reconstruction’. (b), (c), (d), and (e): Visualization of Gaussian centers after applying different sampling techniques, including pruning, random sampling, grid sampling, and density-preserved sampling, respectively.

Gaussian centers \mathbf{p}_i , akin to point clouds, can be approximated as the geometric structure with their scales as the corresponding surface segments.

Following the experimental setup of 3DGS, we utilize sparse points as input initialization in the main paper. As depicted in Fig. 2 (a), we project Gaussian centers of the vanilla 3DGS model onto the rendered image as blue points. Typically, points captured by depth sensors or reconstructed by multiview stereo methods tend to uniformly adhere to the object surface. However, the first row illustrates the ‘overlapping’ phenomenon, where most Gaussians are clustered in certain areas but sparse in neighboring parts. The second row highlights ‘under-reconstruction’, wherein 3DGS fails to model the details of a specific area and presents artifacts in the form of Gaussian blurs. These two phenomena collectively indicate the inefficient spatial distribution of Gaussians, which underscores the inherent redundancy of the Gaussian representation. Additional analysis of 3DGS from initialization with dense point cloud is provided in the appendix.

Gaussian Pruning and Sampling. The most straightforward approach to reduce the number of Gaussians is to perform pruning and sampling. Similar to the importance calculation in [21], we accumulate all blending weights w_{ij} of Gaussian \mathcal{G}_i to obtain its importance score $I_i = \sum_{j=1}^K w_{ij}$, where K is the total number of rays intersected with \mathcal{G}_i . Then, we prune Gaussians with importance lower than a certain threshold. The Gaussian representation is further optimized until convergence and is then shown in Fig. 2 (b). The pruning strategy can reduce the number of Gaussians from 6.1 million to 0.6 million, but it also degrades PSNR from 25.2 dB to 24.9 dB, and the phenomena of ‘overlapping’ and ‘under-reconstruction’ still exist. We also explore combining pruning with various point cloud sampling techniques, such as random sampling, grid sampling, and




Fig. 3: Visual analysis of blur split. (a) Ground truth image. (b) Rendered image using the original 3DGS model [17]. (c) Rendered index of Gaussians with the maximum contribution to each pixel. (d) Rendered image after applying the blur split. (e) Rendered Gaussian index after applying the blur split.

density-preserved sampling, shown in Figs. 2 (c), (d), and (e), respectively. These results suggest that addressing the inefficient spatial distribution of Gaussians and achieving a minimal Gaussian representation while maintaining rendering quality requires further efforts.

4 Methodology

In this section, we introduce our densification and simplification algorithms aimed at reorganizing the spatial distribution of Gaussians, thereby constraining the number of Gaussians while preserving rendering quality.

4.1 Densification

Our densification approach comprises two key components: blur split and depth initialization, to obtain Gaussians with a dense and uniform spatial distribution.

Blur Split. The blur split strategy aims to address the Gaussian blur artifacts, as shown in Fig. 3 (b), which are also discussed in Sec. 3.2. The gradient-based split and clone strategy [17] may fail in areas with smooth color transitions, as illustrated in Fig. 3 (a). Consequently, the corresponding oversized Gaussians \mathcal{G}_i tend to be preserved during optimization. One intuitive approach adopted by 3DGS is to prune Gaussians with an excessively large scale in screen space based on their projected 2D distribution \mathcal{G}_i^{2D} . However, this pruning method does not address the issue of ‘under-reconstruction’.

As illustrated in Fig. 3 (c), we render indexes of Gaussians with the maximum contribution to a pixel as $i_{max} = \arg \max_i w_i$ (*i.e.*, the Gaussian with the maximum weight in alpha-blending) onto the image. The rendered indexes $i_{max}(\mathbf{x})$




Fig. 4: Visual analysis of depth map and depth points from 3DGS [17]. (a): Rendered depth map from alpha blending. (b): Reconstructed depth points from alpha blending with red boxes highlighting significant artifacts. (c): Reconstructed depth points from our depth map using Gaussian centers. (d): Reconstructed depth points from our depth map using mid-points.

of each pixel \mathbf{x} are mapped to random colors. To distinguish with the rendered indexes from all Gaussians $i_{max}(\mathbf{x})$, we denote the projected index of i th Gaussian \mathcal{G}_i as $i(\mathbf{x})$. A key observation is that, compared to the projected index $i(\mathbf{x})$ (*i.e.*, the projected area of Gaussian \mathcal{G}_i), the blurry artifacts are more directly related to the rendered index $i_{max}(\mathbf{x})$. This indicates that Gaussians with large maximum contribution areas lead to the problem of ‘under-reconstruction’.

Based on this observation, for each image, we identify Gaussians with large blurry areas, denoted as \mathcal{G}^{blur} , using the following criterion:

$$\mathcal{G}^{blur} = \{\mathcal{G}_i | S_i > \mathcal{T}_{blur} \wedge i \in [1, N]\}, \quad \mathcal{T}_{blur} = \theta_{blur} \cdot H \cdot W. \quad (2)$$

Here, $S_i = \sum_{\mathbf{x}=(1,1)}^{(H,W)} \delta(i(\mathbf{x}) = i_{max}(\mathbf{x}))$ represents the maximum contribution area of Gaussian \mathcal{G}_i , where δ is the indicator function, and (H, W) denotes the image resolution. The value of S_i can be computed during the forward pass of rasterization. The hyperparameter θ_{blur} is empirically set as 2×10^{-4} to control the extent of the blurry area. Gaussians that meet the criteria \mathcal{G}^{blur} are then split according to the split strategy in [17] during the optimization process. The results of this deblurring strategy are illustrated in Fig. 3 (d) and (e), showcasing the effectiveness of the blur split approach in enhancing image clarity.

Depth Reinitialization. Our depth reinitialization leverages dense depth information to alleviate both the ‘overlapping’ and ‘under-reconstruction’ phenomena. The concept of utilizing depth to enhance neural rendering is both straightforward and effective, as evidenced by its application in several recent NeRF-based algorithms [8, 38]. However, in Gaussian-Splatting-based works, the primary challenge lies not in how to utilize the reconstructed depth, but in how to generate accurate and practical depth data.

Similar to NeRF [24], 3DGS can render blending depth by directly replacing the Gaussian color \mathbf{c}_i with the depth of its center d_i , as $d^{blend} = \sum_{i=1}^N w_i \cdot d_i$. This approach is popular among many subsequent works, and the rendered depth map is shown in Fig. 4 (a). However, we observe that this seemingly decent 2D depth map can be misleading, often resulting in a depth point cloud with significant artifacts. In Fig. 4 (b), we reproject the blending depth to world space and




Fig. 5: Depth points. (a): Depth points from our depth map using Gaussian centers. (b): Depth points from our depth map using mid-points. (c): Merged depth points from our depth map using mid-points.

also visualize Gaussian centers with their colors. Three types of artifacts (depth collapse, object misalignment, and blending boundary) are highlighted with red boxes for better illustration. These failures can be attributed to the alpha blending of multiple Gaussians, and more analysis is provided in the appendix.

Here, we present our formulation of the rendered depth. Given one elliptical Gaussian defined by its scale $\mathbf{s} = (s_x, s_y, s_z)$, we model the Gaussian as an ellipsoid in 3D space as $g(x, y, z) = \frac{x^2}{s_x^2} + \frac{y^2}{s_y^2} + \frac{z^2}{s_z^2} = 1$. With an input ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$, where \mathbf{o} and \mathbf{d} are the ray origin and the ray direction transferred in the coordinate system of the ellipsoid, we can simply calculate the position of the mid-point between two intersection points of the ray and the ellipsoid. We set the depth of this mid-point as our Gaussian depth d_i^{mid} . Another intuitive formulation is to calculate the optimized t where the maximum density of the Gaussian distribution along the ray is reached. We adopt the former formulation because these two share the same numerical result, but the former one can provide an auxiliary discriminant (*i.e.*, the discriminant to distinguish the intersection point of the ray and the ellipsoid exists or not). Details of the calculation are provided in the appendix.

For multiple Gaussians, we only collect the Gaussian with the maximum contribution to the pixel to avoid the artifacts from alpha blending. The final depth d^{mid} can be formulated as $d^{mid} = d_{i_{max}}^{mid}$, where $i_{max} = \arg \max_i w_i$. In practice, we observe that setting the depth based on the Gaussian center, as $d^{center} = d_{i_{max}}$, can also achieve a similar rendering quality as d^{mid} after further optimization. However, considering that d^{mid} shows a better reconstruction result as illustrated in Figs. 5 (a) and (b), we chose to include it in our main pipeline of Mini-Splatting for potential future usage.

After generating the depth map, we reproject all depth points into the world space. For each image, we randomly select a certain number of points as initial points and assign the ground truth colors as their corresponding colors. The total number of sampled points is empirically set at around 3.5 million for each scene. During optimization, we repeatedly reinitialize our Gaussian representation with these points to reorganize the spatial distribution of Gaussians. The resulting

merged point cloud is shown in Fig. 5 (c). These dense and uniform depth points assist our model in initializing a dense and uniform point distribution.

4.2 Simplification

After the densification process, we achieve a denser Gaussian representation of the scene. To further refine this representation, we propose simplification techniques, including intersection preserving and Gaussian sampling, to attain an efficient Gaussian representation while still maintaining high-quality rendering.

Intersection Preserving. Drawing inspiration from the concept of ray-mesh intersection, our intersection preserving technique is designed to discard Gaussians that do not directly intersect with the ray. However, we avoid strictly converting the smooth opacity or blending weights into binary values of 0, 1, as this could potentially compromise the rendering quality.

In the depth reinitialization step, we model 3D Gaussians as ellipsoids and determine their depths based on the mid-points. We render the depth of the Gaussian with the maximum contribution, representing it as the depth of the scene. These mid-points, indicating the intersection of the rays with the scene representation, serve as the basis for our intersected Gaussians \mathcal{G}^{int} . Specifically, for each image, we only retain Gaussians that provide these intersection points, and formulate these Gaussians as the intersected Gaussians:

$$\mathcal{G}^{int} = \{\mathcal{G}_i | i \in \mathcal{I}_{max} \wedge i \in [1, N]\}, \quad (3)$$

where \mathcal{I}_{max} is the collection of all indexes in the rendered indexes $i_{max}(\mathbf{x})$.

Importance-weighted Sampling. In this section, we address the limitations of the direct pruning strategy and introduce our importance-weighted sampling approach. This method considers the geometric structure of the scene, thereby enhancing rendering quality while maintaining a sparse set of Gaussians.

The concept of pruning in neural representation is initially introduced for grid-based neural radiance fields [21, 40], aimed at removing empty or low-information voxels, and this technique can also be extended to Gaussian representations. By assigning each Gaussian a pre-defined importance value I_i , we can simply prune Gaussians with low importance. While this strategy is effective for removing a limited number of Gaussians, it may lead to a collapse in rendering quality when the pruning ratio is relatively large. The main issue lies in the fact that Gaussian importance only reflects its significance before simplification. Direct pruning can easily disrupt local geometry, which in turn degrades the final rendering quality after further optimization. Specifically, neighboring Gaussians often exhibit similar importance in a given area, causing them to be either




Fig. 6: The relation between the preserving ratio and the geometry.

removed or preserved simultaneously, thus risking the destruction of local geometry. To illustrate this point, we conduct Gaussian pruning on our intersected Gaussians and measure the geometry quality using the chamfer distance between the centers of pruned and non-simplified Gaussians. The significant chamfer distance observed after pruning, especially at low preserving ratios, suggests that relying solely on Gaussian importance leads to suboptimal simplification results.

Compared to deterministic pruning, a stochastic sampling strategy can better maintain overall geometry. Working under this observation, we incorporate Gaussian importance I_i into the sampling probability of each Gaussian, formulated as $\mathcal{P}_i = \frac{I_i}{\sum_{i=1}^N I_i}$. Thus, Gaussians can be sampled based on their sampling probability. As illustrated in Fig. 6, the results demonstrate that this importance-weighted sampling strategy exhibits a stronger capability to preserve geometry quality, thereby enhancing the final rendering result. In our experiments, we find that using blending weight-based importance yields better performance in both pruning and sampling. Additionally, different combinations of blending weight and other information (*e.g.*, the scale of projected Gaussians and the rendered index) exhibit varying performance on indoor and outdoor scans. Further implementation details of the importance metric are provided in the appendix.

5 Implementation Variants

Based on our densification and simplification algorithms, we construct three implementation variants. Details of each variant are provided in the appendix.

Mini-Splatting. Our Mini-Splatting is specifically designed for resource-efficient training and rendering. Following a similar optimization schedule as 3DGS [17], we set the total number of optimization steps to $30K$. Densification is enabled until the $15K$ th iteration, after which simplification is conducted at the $15K$ th and $20K$ th iterations. Additionally, we observe that incorporating view-dependent colors barely enhance densification. Therefore, we only enable SH coefficients and increase the SH level after simplification ($15K$).

Mini-Splatting-D. Our Mini-Splatting-D is tailored for quality-prioritized rendering. While its overall pipeline is similar to that of Mini-Splatting, the main distinction lies in the omission of the simplification process.

Mini-Splatting-C. Our Mini-Splatting-C is designed for scenarios requiring storage compression. This pipeline is directly based on our Mini-Splatting. Given a pre-trained Mini-Splatting model, we store the Gaussian centers as `float32`, while other attributes, such as scales and SH coefficients, are processed using transform coding [6] and a standard zip compression.

6 Experiments

Datasets. We perform experiments on three real-world datasets: Mip-NeRF360 [2], Tanks&Temples [18] and Deep Blending [14]. To ensure consistency and fairness in our evaluations, we adopt identical processing details for these datasets,

Table 1: Quantitative evaluation of our Mini-Splatting and previous works. 3DGS* indicates the retrained model from the official implementation.

Dataset Method Metric	Mip-NeRF 360				Tanks&Temples				Deep Blending			
	SSIM ↑	PSNR ↑	LPIPS ↓	Num	SSIM ↑	PSNR ↑	LPIPS ↓	Num	SSIM ↑	PSNR ↑	LPIPS ↓	Num
Plenoxels [40]	0.626	23.08	0.463	-	0.719	21.08	0.379	-	0.795	23.06	0.510	-
INGP-Big [26]	0.699	25.59	0.331	-	0.745	21.92	0.305	-	0.817	24.96	0.390	-
mip-NeRF 360 [2]	0.792	27.69	0.237	-	0.759	22.22	0.257	-	0.901	29.40	0.245	-
Zip-NeRF [3]	0.828	28.54	0.189	-	-	-	-	-	-	-	-	-
3DGS [17]	0.815	27.21	0.214	3.36	0.841	23.14	0.183	1.78	0.903	29.41	0.243	2.98
3DGS [17]*	0.815	27.47	0.216	3.35	0.848	23.66	0.176	1.84	0.904	29.54	0.244	2.82
Mini-Splatting-D	0.831	27.51	0.176	4.69	0.853	23.23	0.140	4.28	0.906	29.88	0.211	4.63
Mini-Splatting	0.822	27.34	0.217	0.49	0.835	23.18	0.202	0.20	0.908	29.98	0.253	0.35

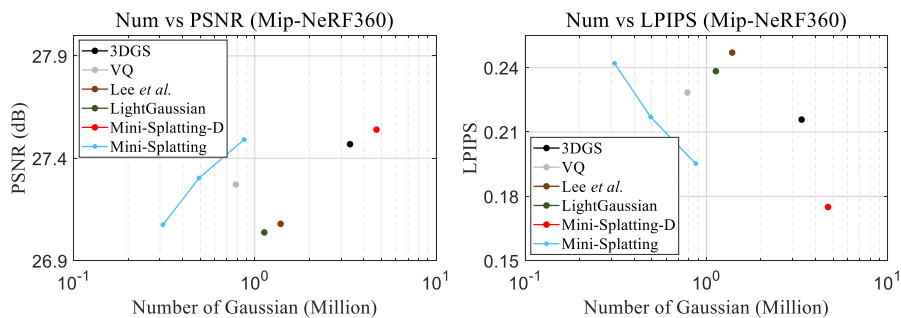


Fig. 7: Number-Quality curves of our algorithms and other baselines on the Mip-NeRF360 dataset.

including scene selection, train/test split, and image resolution, as specified in the official implementation of 3DGS [17].

Implementation Details. We implement our densification and simplification algorithms using PyTorch and integrate them into the optimization pipeline of 3DGS [17]. We modify the Gaussian rasterization module of 3DGS to render Gaussian indexes and depth points. Further implementation details of our Mini-Splatting, Mini-Splatting-D, and Mini-Splatting-C are provided in the appendix.

6.1 Experimental Results

Rendering Quality. The quantitative results of various methods are presented in Table 1. We compare our Mini-Splatting and Mini-Splatting-D against the baseline 3DGS and other NeRF-based algorithms, utilizing standard metrics such as PSNR, SSIM, and LPIPS. The table demonstrates that, in comparison to the 3DGS baseline, our densification method, Mini-Splatting-D, exhibits superior performance across most metrics, highlighting the significant impact of Gaussian spatial distribution in rendering. Notably, our Mini-Splatting-D even surpasses the state-of-the-art neural rendering algorithm, Zip-NeRF [3], in terms of SSIM and LPIPS on the Mip-NeRF360 dataset. Moreover, our Mini-Splatting is comparable to 3DGS, despite employing only a fraction (7× fewer) of the Gaussians. Our algorithms exhibit a decreased PSNR on the Tanks&Temples




Fig. 8: Qualitative results (*bicycle* and *truck*) of our Mini-Splatting and 3DGS. LPIPS and number of Gaussians in million are displayed as (LPIPS, Number).

dataset. This can be attributed to large sky areas in the scans (*i.e.*, *train*) within Tanks&Temples, which our depth-based strategy does not accurately model. Further analysis is provided in the appendix.

As depicted in Fig. 7, we present the number-quality curves of our algorithms alongside 3DGS. The curve for our Mini-Splatting is generated by controlling the sampling ratio during the importance-weighted sampling process. Additionally, we compare against several recent Gaussian-Splatting-based algorithms, incorporating Gaussian pruning within their pipelines. Specifically, we apply the pruning technique proposed by [21] to 3DGS, denoted as VQ, and re-implement the prune and recovery stage outlined by [11] based on their official repository, referred to as LightGaussian. The results reported by Lee *et al.* [20] are collected from their paper. Notably, with a larger number of Gaussians, our Mini-Splatting-D demonstrates superior rendering quality, and it is clear that our Mini-Splatting outperforms all its counterparts at similar numbers of Gaussians.

The rendered images are presented in Fig. 8. These qualitative outcomes are in line with the quantitative results provided in Table 1 and Fig. 7. Notably, our Mini-Splatting-D exhibits superior rendering quality in terms of finer details, while our Mini-Splatting maintains an acceptable rendering quality despite employing a limited number of Gaussians. These results indicate the effectiveness of our densification and simplification algorithm.

Resource Consumption. We also present resource consumption metrics, including training time, training peak GPU memory usage, rendering FPS, and rendering peak GPU memory usage, in Table 2, where peak memory consumption is measured using `torch.cuda.max_memory_allocated()`. In outdoor scenes, our Mini-Splatting-D, despite employing more Gaussians, exhibits similar training time and memory consumption to 3DGS due to our strategy of constraining SH coefficients during densification. Furthermore, our efficient representation significantly accelerates both training and rendering processes while reducing peak memory usage. In indoor scenes, both our Mini-Splatting-D and Mini-Splatting learn denser Gaussian distributions, yet the results still show an acceptable re-

Table 2: Resource consumption of our algorithms and 3DGS on Mip-NeRF360. Mini-Splatting* indicates Mini-Splatting trained on a GTX 1060 6G GPU, and all other models are trained with a RTX 3090 GPU.

Method Metric	Mip-NeRF360 (Outdoor)				Mip-NeRF360 (Indoor)					
	Num	Training		Rendering		Num	Training		Rendering	
		Time	Memory	FPS	Memory		Time	Memory	FPS	Memory
3DGS [17]	4.86	30m8s	7.45GB	98	2.79GB	1.46	24m41s	2.75GB	151	1.07GB
Mini-Splatting-D	5.40	31m48s	7.45GB	83	3.12GB	3.80	40m13s	5.55GB	83	2.46GB
Mini-Splatting	0.57	17m56s	2.61GB	410	0.40GB	0.40	27m2s	2.77GB	362	0.35GB
Mini-Splatting*	0.57	101m11s	2.61GB	64	0.40GB	0.40	154m	2.82GB	40	0.35GB




Fig. 9: Rate-distribution curves for our Mini-Splatting-C and recent compression-oriented works [20, 27] applied to 3DGS.

source consumption of our algorithms. Additionally, we successfully train and evaluate our Mini-Splatting* on a low-cost graphics card (a GTX 1060 6G GPU). This highlights the potential for on-device neural rendering and reconstruction on consumer-grade machines.

Storage Compression. As depicted in Fig. 9, we provide rate-distribution curves comparing our Mini-Splatting-C with two compression-oriented methods [20, 27] applied to 3DGS. Results for these baseline approaches are collected from their respective papers. Due to the optimized and simplified Gaussian representation, our Mini-Splatting-C, which solely integrates basic post-processing techniques, demonstrates superior performance compared to other methods primarily tailored for storage compression.

6.2 Ablation Studies

Densification. We begin by presenting the ablation study of our densification in Table 3. Starting from the 3DGS baseline, we incrementally introduce blur splitting and depth reinitialization steps. The results demonstrate a consistent improvement in rendering quality with the increasing number of Gaussians. Notably, we observe a high correlation between the number of Gaussians and the LPIPS metric, indicating that a dense and uniform spatial distribution of Gaussians positively influences human perception metrics.

We extend our analysis by comparing blending depth with depth obtained from Gaussian center and mid-point in Table 4. As discussed in Sec. 4.1, reinitial-

Table 3: Ablation study of densification.


	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	Num
Baseline	0.815	27.47	0.216	3.35
+ Blur Split	0.819	27.47	0.195	3.74
+ Depth Reinit	0.832	27.54	0.175	4.32

Table 4: Ablation study of depth reinitialization.

	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	Num
Blending	0.513	17.67	0.475	4.01
Center	0.832	27.57	0.176	4.70
Mid	0.832	27.54	0.175	4.69

izing Gaussians with blending depth may degrade rendering quality due to the presence of noise points resulting from depth artifacts. Although depth reinitialization using Gaussian center yields results comparable to those obtained from the mid-point, depth derived from the mid-point exhibits superior dense point cloud reconstruction, as illustrated in Fig. 5, and holds potential for extension to normal estimation tasks.

Simplification. The ablations of simplification on Mip-NeRF360 are depicted in Fig. 10. Starting from the 3DGS baseline, we first add densification to obtain our Mini-Splatting-D. Subsequently, we integrate direct pruning into Mini-Splatting-D and tune the pruning ratio to establish the baseline curve, Add Pruning, for our simplification algorithm. It is worth noting that this baseline encompasses the combination of Mini-Splatting-D and direct pruning, surpassing 3DGS with pruning (*i.e.*, VQ) as depicted in Fig. 7. To underscore the effectiveness of our intersection preserving and sampling techniques, we add intersection preserving and then sampling to Mini-Splatting-D (Add Intersection and Add Sampling). The results clearly demonstrate that our simplification algorithm outperforms direct pruning, achieving approximately a halved reduction in the number of Gaussians while maintaining comparable rendering quality. This result indicates the effectiveness of the proposed simplification method.

**Fig. 10:** Ablations of simplification.

7 Conclusion

We introduced Mini-Splatting, which aimed at representing scene with a constrained number of Gaussians, through our proposed densification and simplification algorithm. Our Gaussian densification process, involving blur splitting and depth initialization, facilitates the construction of a dense Gaussian representation, while our simplification algorithm, featuring intersection preserving and Gaussian sampling, effectively limits the number of Gaussians while preserving high-quality rendering. Our experimental results show the effectiveness of Mini-Splatting and its variants across benchmarks, encompassing not only rendering quality but also resource consumption and storage compression.

Acknowledgements

This work was jointly supported by the Young Scientists Fund of the National Natural Science Foundation of China (42301520), the Research Grants Council of Hong Kong (25206524), the Platform Project of Unmanned Autonomous Systems Research Centre (P0049516), the Seed Project of Smart Cities Research Institute (P0051028), the Research Team Cultivation Program of ShenZhen University (2023JCT003), and the Pearl River Talent Program (2021JC02G046).

References

1. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields. *ICCV* (2021)
2. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. *CVPR* (2022)
3. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Zip-NeRF: Anti-Aliased Grid-Based Neural Radiance Fields. *ICCV* (2023)
4. Chen, G., Wang, W.: A Survey on 3D Gaussian Splatting. *arXiv preprint arXiv:2401.03890* (2024)
5. Chen, S., Tian, D., Feng, C., Vetro, A., Kovačević, J.: Fast Resampling of 3D Point Clouds via Graphs. *IEEE Transactions on Signal Processing* **66**(3), 666–681 (2017)
6. De Queiroz, R.L., Chou, P.A.: Compression of 3D Point Clouds Using a Region-Adaptive Hierarchical Transform. *IEEE TIP* **25**(8), 3947–3956 (2016)
7. Deng, C.L., Tartaglione, E.: Compressing Explicit Voxel Grid Representations: fast NeRFs become also small. In: *WACV*. pp. 1236–1245 (2023)
8. Deng, K., Liu, A., Zhu, J.Y., Ramanan, D.: Depth-supervised NeRF: Fewer Views and Faster Training for Free. In: *CVPR*. pp. 12882–12891 (2022)
9. Dovrat, O., Lang, I., Avidan, S.: Learning to sample. In: *CVPR*. pp. 2760–2769 (2019)
10. Durvasula, S., Zhao, A., Chen, F., Liang, R., Sanjaya, P.K., Vijaykumar, N.: DIST-WAR: Fast Differentiable Rendering on Raster-based Rendering Pipelines. *arXiv preprint arXiv:2401.05345* (2023)
11. Fan, Z., Wang, K., Wen, K., Zhu, Z., Xu, D., Wang, Z.: LightGaussian: Unbounded 3D Gaussian Compression with 15x Reduction and 200+ FPS. *arXiv preprint arXiv:2311.17245* (2023)
12. Franke, L., Rückert, D., Fink, L., Stamminger, M.: TRIPS: Trilinear Point Splatting for Real-Time Radiance Field Rendering. *arXiv preprint arXiv:2401.06003* (2024)
13. Graham, B., Van der Maaten, L.: Submanifold Sparse Convolutional Networks. *arXiv preprint arXiv:1706.01307* (2017)
14. Hedman, P., Philip, J., Price, T., Frahm, J.M., Drettakis, G., Brostow, G.: Deep Blending for Free-Viewpoint Image-Based Rendering. *SIGGRAPH Asia* (2018)
15. Hu, Q., Yang, B., Xie, L., Rosa, S., Guo, Y., Wang, Z., Trigoni, N., Markham, A.: RandLA-Net: Efficient Semantic Segmentation of Large-Scale Point Clouds. In: *CVPR*. pp. 11108–11117 (2020)
16. Katz, S., Tal, A.: Improving the Visual Comprehension of Point Sets. In: *CVPR*. pp. 121–128 (2013)

17. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *TOG* **42**(4) (2023)
18. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and Temples: Benchmarking Large-Scale Scene Reconstruction. *TOG* **36**(4) (2017)
19. Lang, I., Manor, A., Avidan, S.: SampleNet: Differentiable Point Cloud Sampling. In: *CVPR*. pp. 7578–7588 (2020)
20. Lee, J.C., Rho, D., Sun, X., Ko, J.H., Park, E.: Compact 3D Gaussian Representation for Radiance Field. *arXiv preprint arXiv:2311.13681* (2023)
21. Li, L., Shen, Z., Wang, Z., Shen, L., Bo, L.: Compressing Volumetric Radiance Fields to 1 MB. In: *CVPR*. pp. 4222–4231 (2023)
22. Luiten, J., Kopanas, G., Leibe, B., Ramanan, D.: Dynamic 3D Gaussians: Tracking by Persistent Dynamic View Synthesis. *3DV* (2024)
23. Lv, C., Lin, W., Zhao, B.: Intrinsic and Isotropic Resampling for 3D Point Clouds. *PAMI* **45**(3), 3274–3291 (2022)
24. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. *ECCV* (2020)
25. Moening, C., Dodgson, N.A.: Intrinsic Point Cloud Simplification. *Proc. 14th GrahCon* **14**(23), 2 (2004)
26. Müller, T., Evans, A., Schied, C., Keller, A.: Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *SIGGRAPH* (2022)
27. Niedermayr, S., Stumpfegger, J., Westermann, R.: Compressed 3D Gaussian Splatting for Accelerated Novel View Synthesis. *arXiv preprint arXiv:2401.02436* (2023)
28. Niemeyer, M., Manhardt, F., Rakotosaona, M.J., Oechsle, M., Duckworth, D., Gosula, R., Tateno, K., Bates, J., Kaeser, D., Tombari, F.: RadSplat: Radiance Field-Informed Gaussian Splatting for Robust Real-Time Rendering with 900+ FPS. *arXiv preprint arXiv:2403.13806* (2024)
29. Pauly, M., Gross, M., Kobbelt, L.P.: Efficient Simplification of Point-Sampled Surfaces. In: *VIS*. pp. 163–170. *IEEE* (2002)
30. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In: *CVPR*. pp. 652–660 (2017)
31. Qian, S., Kirschstein, T., Schoneveld, L., Davoli, D., Giebenhain, S., Nießner, M.: GaussianAvatars: Photorealistic Head Avatars with Rigged 3D Gaussians. *CVPR* (2024)
32. Rho, D., Lee, B., Nam, S., Lee, J.C., Ko, J.H., Park, E.: Masked Wavelet Representation for Compact Neural Radiance Fields. In: *CVPR*. pp. 20680–20690 (2023)
33. Sun, C., Sun, M., Chen, H.T.: Direct Voxel Grid Optimization: Super-fast Convergence for Radiance Fields Reconstruction. *CVPR* (2022)
34. Sun, J.M., Wu, T., Gao, L.: Recent Advances in Implicit Representation-based 3D Shape Generation. *Visual Intelligence* **2**(1), 9 (2024)
35. Tang, J., Ren, J., Zhou, H., Liu, Z., Zeng, G.: DreamGaussian: Generative Gaussian Splatting for Efficient 3D Content Creation. *ICLR* (2024)
36. Wen, C., Yu, B., Tao, D.: Learnable Skeleton-Aware 3D Point Cloud Sampling. In: *CVPR*. pp. 17671–17681 (2023)
37. Xie, X., Gherardi, R., Pan, Z., Huang, S.: HollowNeRF: Pruning Hashgrid-Based NeRFs with Trainable Collision Mitigation. In: *ICCV*. pp. 3480–3490 (2023)
38. Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., Neumann, U.: Point-NeRF: Point-based Neural Radiance Fields. In: *CVPR*. pp. 5438–5448 (2022)
39. Yan, Y., Lin, H., Zhou, C., Wang, W., Sun, H., Zhan, K., Lang, X., Zhou, X., Peng, S.: Street Gaussians for Modeling Dynamic Urban Scenes. *arXiv preprint arXiv:2401.01339* (2024)

40. Yu, A., Fridovich-Keil, S., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance Fields without Neural Networks. CVPR (2022)
41. Yu, Z., Chen, A., Huang, B., Sattler, T., Geiger, A.: Mip-Splatting: Alias-free 3D Gaussian Splatting. CVPR (2024)
42. Zwicker, M., Pfister, H., Van Baar, J., Gross, M.: EWA Splatting. TVCG **8**(3), 223–238 (2002)