

Improving Gaussian Splatting with Localized Points Management

Haosen Yang^{1*} Chenhao Zhang^{1*} Wenqing Wang^{1*} Marco Volino¹ Adrian Hilton¹
 Li Zhang² Xiatian Zhu¹

¹University of Surrey ²Fudan University
 *Equal contribution

<https://happy-hsy.github.io/projects/LPM/>




Figure 1. Visualization of points behavior. 3DGS produces ill-conditioned Gaussians (red box) that occlude other valid points, resulting in noticeably incorrect depth estimation. **LPM** handles these ill-conditioned points to reduce negative impacts and further calibrate the geometry.

Abstract

Point management is critical for optimizing 3D Gaussian Splatting models, as point initiation (e.g., via structure from motion) is often distributionally inappropriate. Typically, Adaptive Density Control (ADC) algorithm is adopted, leveraging view-averaged gradient magnitude thresholding for point densification, opacity thresholding for pruning, and regular all-points opacity reset. We reveal that this strategy is limited in tackling intricate/special image regions (e.g., transparent) due to inability of identifying all 3D zones requiring point densification, and lacking an appropriate mechanism to handle ill-conditioned points with negative impacts (e.g., occlusion due to false high opacity). To address these limitations, we propose a **Localized Point Management (LPM)** strategy, capable of identify-

ing those error-contributing zones in greatest need for both point addition and geometry calibration. Zone identification is achieved by leveraging the underlying multiview geometry constraints, subject to image rendering errors. We apply point densification in the identified zones and then reset the opacity of the points in front of these regions, creating a new opportunity to correct poorly conditioned points. Serving as a versatile plugin, LPM can be seamlessly integrated into existing static 3D and dynamic 4D Gaussian Splatting models with minimal additional cost. Experimental evaluations validate the efficacy of our LPM in boosting a variety of existing 3D/4D models both quantitatively and qualitatively. Notably, LPM improves both static 3DGS and dynamic SpaceTimeGS to achieve state-of-the-art rendering quality while retaining real-time speeds, excelling on challenging datasets such as Tanks & Temples and the Neural

1. Introduction

Neural rendering has emerged as a generalizable, flexible, and powerful approach for photorealistic novel view synthesis (NVS) of any camera poses [23], underpinning a wide variety of applications in augmented/virtual/mixed reality [9], robotics [37], and generation [25], among more others. For example, taking a learning-based parametric idea, Neural Radiance Fields (NeRFs) [23] implicitly represent the scene radiance of any complexity using neural networks (e.g., MLPs), without the tedious requirements of model handcrafting for accounting the scene variations in geometry, texture, illumination. However, their view rendering is inefficient computationally due to heavy ray sampling, thus suffer in scaling to high-resolution content applications and large scale scene modeling [29, 32].

Recently, 3D Gaussian Splatting (3DGS) [14] has come as an alternative with explicit representation, much faster model optimization and real-time neural rendering. The process begins by initializing a set of 3D Gaussian points using Structure from Motion (SfM) [28]. This is followed by optimizing the parameters of these points through view reconstruction loss, resulting in a view output generated with differentiable splatting-based rasterization. However, the point initialization is often distributionally non-optimal, leading to issues such as under-population (e.g., insufficient points) or over-population (e.g., excessive points) in the 3D space. Consequently, a point management mechanism, such as Adaptive Density Control (ADC), is necessary during optimization. However, we identify several limitations with ADC: (i) Thresholding the average gradient to determine regions for point densification often overlooks under-optimized points. For instance, larger Gaussian points typically have lower average gradients and may frequently appear across various views in screen space. (ii) Point sparsity complicates the addition of sufficient and reliable points needed to comprehensively cover the scene. (iii) Mis-optimized Gaussian points can have detrimental effects, such as occluding other valuable points and leading to incorrect depth estimates (see erroneous placements on windows in Fig. 1). While GaussianPro [7] and PixelGS [41] try to solve (1) and (2) through multi-view stereo [2] and additional gradient propagation, respectively, these methods significantly increase the training budget, as shown in the Table 4.

To overcome the aforementioned limitations, in this paper we propose a novel and efficient *Localized Point Management* (LPM) approach. Our idea is intuitive – identifying those 3D Gaussian points leading to rendering errors. Thus we start with an image rendering error map of a specific view. To obtain the error contributing 3D points, we leverage the region correspondence between different

views via feature mapping, subject to the multiview geometry constraint. For each pair of corresponded regions, we cast the rays through them at their respective camera views in the cone shape, and consider their *intersection* as the error source zone. Within each such zone, we consider two situations: (1) At presence of points, we further apply point densification at a lower threshold to complement the original counterpart locally; (2) In case no point due to point sparsity, we add new Gaussian points. Concurrently, we reset the opacity of points with high opacity estimates that are located in front of these zones, as they can significantly affect view rendering. This provides an opportunity to correct potentially ill-conditioned points while tuning the newly added ones in the subsequent optimization. To minimize model expansion, we prune the points by opacity in a density-aware manner.

We summarize the *contributions* below: (1) Through in-depth analysis, we have identified several limitations in the standard point management mechanism used in Gaussian Splatting that impede model optimization. (2) We present *Localized Point Management* (LPM) for these issues by identifying error-contributing 3D zones and implementing appropriate operations for point densification and opacity reset. (3) Extensive experiments validate the benefits of our LPM in improving a diversity of existing 3D and 4D Gaussian Splatting models in novel view synthesis on both static and dynamic scenes.

2. Related Work

Neural Scene Representations has always been an important direction in novel view synthesis. Previous methods allocate neural features to structures such as volume [21, 27], texture [31], and point cloud [1]. The pioneering work of NeRF [23] proposes integrating neural networks with 3D volumetric representations to convert a 3D scene into a learnable density field, enabling high-quality novel view synthesis without requiring explicit modeling of the 3D scene and illumination. Later on, numerous works emerge to boost the quality and efficiency of volume rendering, [3, 5, 35] refine the point sampling strategy in ray marching, some advanced works [4, 33] reparameterize the scene to produce a more compact representation. Additionally, regularization terms [8, 39] can be incorporated to constrain the scene representation, resulting in a closer approximation to real-world geometry. Despite their high-quality representational performance, these methods are typically computationally inefficient for view rendering due to the extensive ray sampling required and the use of Multi-Layer Perceptrons (MLPs) to represent the scene, complicating the computation and optimization of any point within the scene. To address this, several works have proposed novel scene representations aimed at accelerating the rendering process. These representations replace MLPs with sparse

voxels [20], hash tables [24], or triplanes [6], significantly enhancing rendering speed. However, real-time rendering remains challenging due to the inherent complexity of the ray marching strategy in volume rendering.

Gaussian Splatting represents a recent advancement in novel view synthesis, enabling real-time high-quality rendering. It contributes to splatting-based rasterization by computing pixel colors through depth sorting and α -blending of projected 2D Gaussians, thereby avoiding the complex sampling strategies of ray marching and achieving real-time performance. It is precisely due to its real-time high-quality rendering capabilities that 3DGS has been applied to various domains, including autonomous driving, content generation [30], and 4D dynamic scenes [18, 34, 38], among others. Despite these advancements, 3DGS still has some drawbacks, such as the storage of Gaussians and handling multi-resolution, and so on. Several works have enhanced 3DGS by improving Gaussian representation, including techniques such as low-pass filtering [40], multiscale Gaussian representations [36], and interpolating Gaussian attributes from structured grid features [22]. However, these works often overlook the importance of point management, specifically Adaptive Density Control, which is typically applied during optimization to address issues like under-population or over-population in the 3D space. Only a few works have focused on point management. For example, GaussianPro [7] directly tackles densification limitations, bridging gaps from SfM-based initialization. Pixel-GS [41] proposes a gradient scaling strategy to suppress artifacts near the camera. Additionally, [26] introduces an auxiliary per-pixel error function to implicitly supervise point contributions.

Although these methods improve densification, they are still unable to identify all 3D zones that require point densification and lack a proper mechanism to handle ill-conditioned points with negative impacts. Here, we propose a novel approach, Localized Point Management, capable of identifying error-contributing zones with greatest demand for both point addition and geometry calibration.

3. Method

3.1. Preliminaries: 3D Gaussian Splatting

Gaussian Splatting builds upon concepts from EWA [42] splatting and proposes modeling a 3D scene as a collection of 3D Gaussian points $\{G_i \mid i = 1, \dots, K\}$, rendered through volume splatting. Each 3D Gaussian G is defined by the equation:

$$G(x) = e^{-\frac{1}{2}(x-\mu)^T \Sigma^{-1}(x-\mu)},$$

where $\mu \in \mathbb{R}^{3 \times 1}$ represents the mean vector, and $\Sigma \in \mathbb{R}^{3 \times 3}$ denotes its covariance matrix. To maintain the positive semi-definite nature of Σ during optimization, it is repre-

sented as $\Sigma = RSS^T R^T$, with the orthogonal rotation matrix $R \in \mathbb{R}^{3 \times 3}$ and the diagonal scale matrix $S \in \mathbb{R}^{3 \times 3}$.

To render an image from a specific viewpoint, the color of each pixel p is determined by blending N ordered Gaussians $\{G_i \mid i = 1, \dots, N\}$ that overlap p , using the formula:

$$c(p) = \sum_{i=1}^N c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j),$$

where α_i is derived by evaluating a projected 2D Gaussian from G_i at pixel p combined with a learned opacity for G_i , and c_i is the learnable, view-dependent color modeled using spherical harmonics in 3DGS. Gaussians that influence p are arranged in ascending order based on their depth from the current viewpoint. Employing differentiable rendering techniques allows for the end-to-end optimization of all Gaussian attributes through training view reconstruction.

Point management Since existing 3DGS variants start by initializing 3D Gaussian points using Structure from Motion (SfM), the points are often coarse and non-optimal in space. During optimization, a point management mechanism, Adaptive Density Control (ADC), is typically applied to manage point distribution issues. Specifically, thresholding the average gradient is used to decide on point densification. For each Gaussian point G_i , 3DGS tracks the magnitude of the positional gradient $\frac{\partial L_p}{\partial \mu_i}$ across all rendered views, which is then averaged to a quantity T_i . During each training iteration, if the gradient T_i surpasses a predefined threshold, it considers this point as inadequately representing the corresponding 3D region. With the scale of the Gaussian as the size measure, a large Gaussian will be split into two, while a small one leads to point cloning.

However, this commonly used ADC strategy is unable to identify all the 3D zones with the underlying need for point densification. This is because, often the local complexity of scene geometry varies significantly, which beyond the reach of any single-value based thresholding. Besides, there is lacking of a proper mechanism to handle ill-conditioned points with negative impacts (e.g., wrong opacity values estimated during training with points distributed here and there).

3.2. Localized Gaussian Point Management

To address the aforementioned issues, we introduce a novel model agnostic point management approach, *Localized Point Management* (LPM), which leverages multiview geometry constraints to identify error contributing 3D points, with the guidance of image rendering errors. This approach can be seamlessly integrated with existing 3DGS models without the need for architectural modification. As illustrated in Figure 2, we begin with an image rendering error map for a specific view. Under the multiview geometry constraint, the corresponding regions in the referred view are

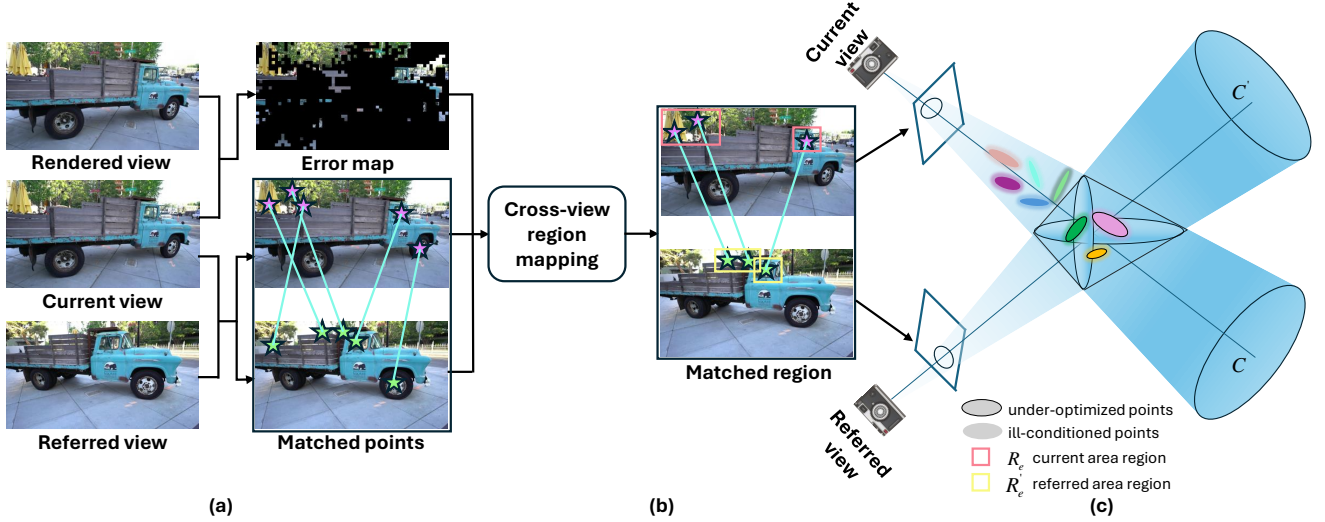


Figure 2. Overview of our Localized Point Management (LPM). (a) We start with an image rendering error map *versus* the current view (the ground-truth). Concurrently, matching points are identified between the current view and a referred view sampled as an adjacent view via off-the-shelf feature mapping. (b) Subsequently, cross-view region mapping is then employed to locate the correspondence region in the referred view. (c) For each pair of corresponded regions, we cast the rays through them at their respective camera views in the cone shape, and consider their *intersection* as the error source zone. The final step involves identifying under-optimized or ill-conditioned points within these zones, where under-optimized/empty places are densified, and ill-conditioned points are reset.

matched via feature mapping. For each pair of corresponding regions, we then cast rays through them from their respective camera views in a cone and identify their intersection as the error source zone. Within each zone, we perform localized point manipulation.

Error map generation To accurately localize those zones in the 3D space that require point densification and geometry calibration, we initiate our process by rendering the current view image through the splatting of 3D Gaussians. This is followed by generating an error map (Figure 2(a)) for this specific view against the ground-truth image using an error function [18].

Error contributing 3D zone identification To project this rendering error back to the 3D space, we leverage the region correspondence between different views under multiview geometry constraints. This involves the following two key steps.

(i) **Cross-view region mapping** We select a neighboring view as the referred image. Following LightGlue [19], which predicts a partial assignment between two sets of local features extracted from images A and B from different views. Each feature consists of sets of 2D features position $\{F_i \mid (x_i, y_i) \in [0, 1]^2\}$, normalized by the image size. The images A and B contain M and N local features. LightGlue outputs a set of correspondences $\mathcal{M} = \{(i, j)\} \subseteq A \times B$, where i and j denote the indices of matched points from sets A and B , respectively. Since the 2D rendering error regions in the current view may not all appear in the referred image, we select the paired regions (R_e, R'_e) (Fig-

ure 2(b)) based on the matching points, where R_e represents the region in the current view and R'_e corresponds to the region in the reference view. Additionally, this paired region undergoes multiview adaptive adjustments based on the error map throughout the optimization process.

(ii) **2D-to-3D projection** After obtaining the paired regions with render errors, we project each 2D error region to the 3D space via multiview geometry constraints. Specifically, we cast the rays \mathcal{C} in cone shape for region R_e from the camera's center of projection o along the direction d , which aligns with the pixel's center (Figure 2(c)). The apex of this cone is located at o , and its radius at the image plane. Hence, $o + d$ is parameterized as \mathcal{C} . The radius r_{Cone} is set to match the radius of the smallest circumscribed circle of the 2D plane error region, creating a cone on the 3D space that can trace the Gaussian points contributing to the 2D error region. Concurrently, a corresponding cone, denoted as C' , belong to region R'_e is similarly projected. Subsequently, we compute the intersection points of these rays. In order to regionalize these points, we directly use a smallest sphere that can contain these points as error source 3D zone R_{zone} .

Points manipulation Recall that in existing 3DGS, points management only relies on the view-averaged gradient magnitude τ to determine point densification *globally*. In addition to this, we *further* perform localized points addition and geometry calibration within the identified error source 3D zone R_{zone} . For the point addition, we consider two common situations: (1) In the presence of points, we



Figure 3. Qualitative evaluation of our LPM across diverse static datasets [4, 12]. Our LPM improves 2DGS [13] and 3DGS [14] on these challenging scenarios, e.g. (a) **Light artifacts**, (b) **Completeness in the distance**, (c) **Depth structure** and (d) **Mesh details**. See red patches for highlighted visual differences.

apply point densification to locally complement the original counterparts. We set a lower threshold to select the points that need densification, aiming to enhance the geometric details. The densification rule is consistent with 3DGS, but it focuses on local 3D zones that need it most. Specifically, for small Gaussians, our strategy involves cloning the Gaussians while maintaining their size and repositioning them along the positional gradient to better capture emerging geometrical features. Conversely, larger Gaussians situated in areas of high variance are split into smaller points to more accurately represent the underlying geometry. (2) In cases

of point sparsity, we add new Gaussian points at the center of the 3D zone. In the context of α -blending in 3DGS, if the points at the forefront of the identified 3D zone R_{zone} have the highest opacity, they may occlude valid points, leading to incorrect depth estimation, as shown in Figure 1. To deal with such issues, we treat these points as potentially ill-conditioned. We reset these points to provide an opportunity for correction, further calibrating the geometry. To minimize model expansion, we adaptively prune points based on their opacity values, starting from low to high. The number of points pruned is determined by the density of

Method	Mip-NeRF 360			Tanks&Temples			Deep Blending		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
Plenoxels	23.08	0.625	0.463	21.08	0.719	0.379	23.06	0.795	0.510
INGP-Big	25.59	0.699	0.331	21.92	0.745	0.305	24.96	0.817	0.390
Mip-NeRF 360	27.69	0.792	0.237	22.22	0.759	0.257	29.40	0.901	0.245
3DGS	27.21	0.815	0.214	23.14	0.841	0.183	29.41	0.903	0.243
3DGS*	27.47	0.816	0.216	23.67	0.849	0.177	29.55	0.904	0.245
3DGS* + LPM	27.59	0.820	0.216	23.83	0.850	0.181	29.76	0.908	0.241
2DGS*	27.15	0.808	0.246	23.58	0.832	0.185	29.35	0.899	0.262
2DGS* + LPM	27.42	0.817	0.228	23.65	0.848	0.180	29.52	0.903	0.240
MipGS*	27.51	0.817	0.210	23.69	0.852	0.173	29.58	0.910	0.242
MipGS* + LPM	27.70	0.821	0.210	23.82	0.851	0.180	29.61	0.910	0.241
PixelGS*	27.54	0.819	0.203	23.75	0.850	0.175	29.58	0.920	0.220
PixelGS* + LPM	27.80	0.830	0.190	24.02	0.856	0.173	29.65	0.910	0.196

Table 1. Comparison of various methods across different scenes on the Mip-NeRF 360 dataset, Tanks&Temples and Deep Blending. * indicates the retrained model from the official implementation. Bold represents best, underline indicates second best.

points in the zone. This strategic reduction ensures that our point management remains cost efficient and adaptive to the evolving needs of the scene representation.

4. Experiment

Datasets and metrics We conducted an extensive evaluation using both static and dynamic scenes derived from publicly datasets. For static scenes, our approach was applied to a total of 13 scenes as specified in the 3DGS framework [14], which includes nine scenes from Mip-NeRF360 [3], two from Tanks&Temples [15], and two from DeepBlending [12]. In the context of dynamic scenes, our approach was tested across six scenes from the Neural 3D Video Dataset [17].

To evaluate novel view synthesis performance, we followed standard protocols by selecting one out of every eight images as test images, with the remaining used for training in static scenes. For each dynamic scene within the Neural 3D Video Dataset, one view was designated for testing while the others were allocated for training purposes. Evaluation metrics included the peak signal-to-noise ratio (PSNR), structural similarity index measure (SSIM), and the learned perceptual image patch similarity (LPIPS), which are broadly recognized standards in the field.

Baselines and implementation Vanilla 3D Gaussian Splatting (3DGS) [14], 2D Gaussian Splatting (2DGS) [13], Mip Gaussian Splatting (MipGS) [40], PixelGS [41] and SpacetimeGS (STGS) [18] were selected as our main baselines for their established art performance in novel view synthesis. For the static 3D benchmark, we also recorded the results of Mip-NeRF360 [3], iNGP [24] and Plenoxels [10] as in [14]. For the Dynamic 4D benchmark, we performed system comparison, such as DyNeRF [16], K-planes [11] and so on. In alignment with the approach

described in 3DGS an STGS, our models were trained for 30k iterations across all scenes, following the same training schedule and hyperparameters. In addition to the original Gaussian densification strategies used in 3DGS and Space-Time Gaussian, we also performed localized points management, including addition, reset, and pruning. We maintained the same thresholds for splitting and cloning points as in the original 3DGS and SpaceTime Gaussian. For point matching, we performed offline extraction to save computational cost. All experiments were conducted on an RTX 3090 GPU with 24GB of memory.

4.1. Main Results

Results on static 3D datasets The quantitative results (PSNR, SSIM, and LPIPS) on the Mip-NeRF 360 and Tanks & Temples datasets are presented in Tables 1. We retrained the 3DGS model (referred to as 3DGS*) as it yields better performance compared to the vanilla 3DGS and its variants. Our approach achieves results comparable to the state-of-the-art on the Mip-NeRF360 dataset and further enhances all 3DGS based method using our point management technique. Additionally, LPM improve vanilla 3DGS and PixelGS to set new state-of-the-art results on the Mip-NeRF 360, Tanks & Temples datasets and Dep Blending, effectively capturing more challenging environments (*e.g.*, light effects, transparency). These results quantitatively validate the effectiveness of our method in improving the quality of reconstruction.

In Figures 3, we present a comparison between 3DGS [14] and 3DGS* + LPM, focusing on both appearance and depth. A variety of improvements can be observed, particularly in challenging cases such as light effects, completeness at a distance. Our LPM significantly reduces artifacts in specific regions on top of 3DGS, partic-

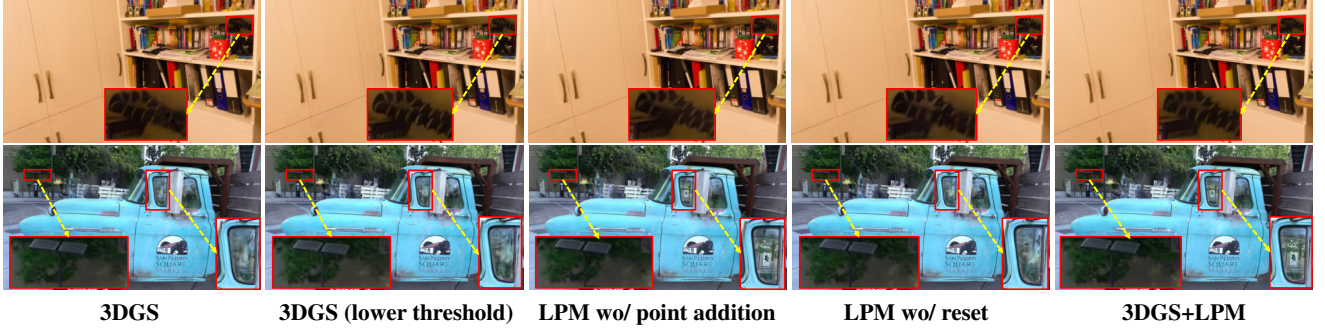


Figure 4. Effect of key operations of LPM. We show that the *point addition* operation effectively captures the geometric details in the scene; The *point reset* operation based on the error map further calibrate the geometry.

Method	PSNR	DSSIM	LPIPS	FPS
LLFF ¹	23.24	0.076	0.235	-
DyNeRF ¹	29.58	0.083	0.063	0.015
Dynamic-4DGS ¹	-	-	-	30
4DGS ¹	29.38	-	-	114
STGS ¹	<u>29.58</u>	0.022	<u>0.063</u>	103
STGS* ¹	29.48	<u>0.023</u>	0.066	<u>110</u>
STGS*¹ + LPM	29.84	0.022	0.062	105
StreamRF	28.26	-	0.039	10.9
NeRFPlayer	30.69	-	0.111	0.05
HyperReal	31.10	-	0.096	2
K-planes	31.63	-	0.31	3
MixVoxels-X	31.73	-	0.064	4.6
Dynamic-4DGS	31.15	0.016	0.049	30
4DGS	32.01	-	0.055	114
STGS	<u>32.05</u>	0.014	<u>0.044</u>	140
STGS*	31.99	<u>0.015</u>	0.045	145
STGS*+ LPM	32.40	0.014	0.045	<u>140</u>

Table 2. Quantitative comparisons on the Neural 3D Video dataset. “FPS” is measured at a resolution of 1352×1014 . Some methods only report results for a subset of scenes. For a fair comparison, we report LPM’s results under two pre-existing settings. ¹ Only includes the *Flame Salmon* scene. **Bold** represents best, underline indicates second best.

ularly in the tree at the second. These regions require more points for accurate population, leading to a more precise and detailed reconstruction. Additionally, the tablecloth in the first row is affected by ill-conditioned points. Furthermore, we provide depth and mesh comparisons in the third and final rows. All these observations demonstrate that our geometry calibration with LPM offers an opportunity to correct these potentially ill-conditioned points, thereby enhancing the overall reconstruction accuracy.

Results on dynamic 4D datasets Table 2 presents a quantitative evaluation on the Neural 3D Video Dataset.



Figure 5. Qualitative evaluation on dynamic Neural 3D Video dataset [17]. LPM improves STGS [18] for both scenes **Transparent** (e.g., window) and **Dynamic movements** (e.g., dog’s tongue).

Following established practices, training and evaluation are conducted at half resolution, with the first camera held out for evaluation [16]. Integrating our LPM into Space-TimeGS yields the best performance across all comparisons. Notably, our method demonstrates significant improvements in the challenging *Flame Salmon* scene compared to SpaceTimeGS [18]. Our approach not only surpasses previous methods in rendering quality but also maintains comparable rendering speed.

In addition to the quantitative assessment, we provide qualitative comparisons on the *Flame Salmon* and *Flame Steak* scenes, as illustrated in Figure 5. The quality of synthesis in both static and dynamic regions markedly outperforms STGS. Several intricate details, including the tree behind the window and the fine features like the dog’s tongue, are faithfully reproduced with higher accuracy compared to STGS [18]. Both examples indicate that LPM improves upon STGS for superior scene modeling.

4.2. Ablation study

We conducted ablation studies on the more challenging scene: PlayRoom from Deep Blending [12] and Truck from Tanks&Temples [15].

Effectiveness and cost of LPM We hypothesize that the Adaptive Density Control (ADC) tends to overlook under-

Method	PlayRoom			Truck		
	PSNR	LPIPS	SSIM	PSNR	LPIPS	SSIM
Full LPM	30.22	0.241	0.910	25.61	0.154	0.883
wo/ point addition	30.10	0.241	0.910	25.43	0.153	0.883
wo/ reset	30.07	0.243	0.908	25.52	0.144	0.883

Table 3. Performance comparison for different configurations

Scene	Method	PSNR	LPIPS	Gaussians	Training time
PlayRoom	3DGS*	30.03	0.244	232k	22min
	3DG* (lower threshold)	29.69	0.240	523k	36min
	GaussianPro	Out of Memory			
	PiexlGS	30.09	0.241	186k	35min
	3DGS + LPM	30.22	0.241	186k	23min
Truck	3DGS*	25.42	0.146	257k	19 min
	3DGS* (lower threshold)	25.45	0.127	635k	35min
	GaussianPro	25.40	0.164	312k	36min
	PiexlGS	25.51	0.121	518k	37min
	3DGS + LPM	25.61	0.154	265k	21min

Table 4. Cost-effectiveness analysis. Rendering speed of both methods are measured on our machine. **Note:** For 3DGS+LPM, training time includes the feature matching process.

optimized points due to its simplistic approach of thresholding the average gradient. The straight way to identify the all points is lowering threshold to densification process. Although this solution can reduce blurring in specific regions, such as the toy (red box) illustrated in Figure 4, it still has limitations. As shown in Table 4, lowering the threshold for 3DGS significantly increases the number of Gaussian points and decreases rendering speed. Additionally, the PSNR of the quantitative results decreases due to the introduction of unnecessary points in already dense areas. In contrast, **LPM** effectively generates points in areas indicated by the error map, leading to more accurate and detailed reconstructions while maintaining real-time rendering speed. As demonstrated by the qualitative comparison in Figure 4, 3DGS with LPM achieves superior qualitative results. We further compare our method with other recent methods that also focus on adaptive density control (ADC). While PixelGS and GaussianPro achieve improvements in rendering quality, their training times increase substantially as they only consider point addition and extra gradient propagation. In contrast, LPM achieves a noticeable improvement with only a slight increase in training time due to (1) point matching [19] is much faster (2) considering model expansion to dynamically prune the points by their addition number and (3) selecting points in error-contributing zones 3D zone using the parallel matrix operations.

Individual points manipulation We study the effect of individual points manipulation of LPM, including the *point addition* and *reset ill-conditional points*. The results in Ta-

ble 3 show that, (1) each manipulation is useful with positive gain, suggesting that the LPM is meaningful. (2) The *point addition* operation densify the under-optimized points which may be overlook in the 3DGS, further captures the geometry details (*e.g.*, detail of toy and leaf of the tree, see Fig. 4). (3) Reset points in ceratin zone provide the opportunity of correct the ill-conditioned points to achieve geometry calibration, (*e.g.*, window of the trunk, see Fig. 4).

5. Conclusion

We propose Localized Point Management (LPM), a novel point management approach to address the limitations of the Adaptive Density Control (ADC) mechanism in 3D Gaussian Splatting (3DGS). The core idea of LPM is identifying the error-contributing 3D zones that require both point addition and geometry calibration under multiview geometry constraints, guided by image rendering errors. We implement appropriate operations for point densification and opacity reset. As a versatile plugin, LPM can be seamlessly integrated into existing 3DGS-based rendering methods. Extensive experiments across both static 3D and dynamic 4D scenes validate the efficacy of LPM in enhancing existing ADC mechanisms both quantitatively and qualitatively. While our method identifies the 3D Gaussian points that lead to rendering errors, it still follows the densification rules of 3DGS [14]. This approach may not be optimal for under-optimized points, and we leave this aspect for further investigation.

References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*, pages 696–712. Springer, 2020. 2
- [2] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3):24, 2009. 2
- [3] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 2, 6, 1
- [4] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. 2, 5
- [5] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19697–19705, 2023. 2
- [6] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision*, pages 333–350. Springer, 2022. 3
- [7] Kai Cheng, Xiaoxiao Long, Kaizhi Yang, Yao Yao, Wei Yin, Yuexin Ma, Wenping Wang, and Xuejin Chen. Gaussianpro: 3d gaussian splatting with progressive propagation. *arXiv preprint arXiv:2402.14650*, 2024. 2, 3
- [8] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12882–12891, 2022. 2
- [9] Nianchen Deng, Zhenyi He, Jiannan Ye, Budmonde Duinkharjav, Praneeth Chakravarthula, Xubo Yang, and Qi Sun. Fov-nerf: Foveated neural radiance fields for virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 28(11):3854–3864, 2022. 2
- [10] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinlong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. 6
- [11] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12479–12488, 2023. 6
- [12] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (ToG)*, 37(6):1–15, 2018. 5, 6, 7, 1
- [13] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 5, 6
- [14] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):1–14, 2023. 2, 5, 6, 8, 1
- [15] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017. 6, 7, 1
- [16] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, and Zhaoyang Lv. Neural 3d video synthesis from multi-view video. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5511–5521, 2022. 6, 7
- [17] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5521–5531, 2022. 6, 7
- [18] Zhan Li, Zhang Chen, Zhong Li, and Yi Xu. Spacetime gaussian feature splatting for real-time dynamic view synthesis. *arXiv preprint arXiv:2312.16812*, 2023. 3, 4, 6, 7, 1
- [19] Philipp Lindenberger, Paul-Edouard Sarlin, and Marc Pollefeys. Lightglue: Local feature matching at light speed. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17627–17638, 2023. 4, 8
- [20] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020. 3
- [21] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751*, 2019. 2
- [22] Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. *arXiv preprint arXiv:2312.00109*, 2023. 3
- [23] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2
- [24] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4):1–15, 2022. 3, 6

- [25] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 2
- [26] Samuel Rota Bulò, Lorenzo Porzi, and Peter Kotschieder. Revising densification in gaussian splatting. *arXiv e-prints*, pages arXiv–2404, 2024. 3
- [27] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2437–2446, 2019. 2
- [28] Noah Snavely, Steven M Seitz, and Richard Szeliski. Photo tourism: exploring photo collections in 3d. In *ACM siggraph 2006 papers*, pages 835–846. 2006. 2
- [29] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8248–8258, 2022. 2
- [30] Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023. 3
- [31] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *Acm Transactions on Graphics (TOG)*, 38(4):1–12, 2019. 2
- [32] Haithem Turki, Deva Ramanan, and Mahadev 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*, pages 12922–12931, 2022. 2
- [33] Peng Wang, Yuan Liu, Zhaoxi Chen, Lingjie Liu, Ziwei Liu, Taku Komura, Christian Theobalt, and Wenping Wang. F2-nerf: Fast neural radiance field training with free camera trajectories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4150–4159, 2023. 2
- [34] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528*, 2023. 3
- [35] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5438–5448, 2022. 2
- [36] Zhiwen Yan, Weng Fei Low, Yu Chen, and Gim Hee Lee. Multi-scale 3d gaussian splatting for anti-aliased rendering. *arXiv preprint arXiv:2311.17089*, 2023. 3
- [37] Ze Yang, Yun Chen, Jingkan Wang, Sivabalan Manivasagam, Wei-Chiu Ma, Anqi Joyce Yang, and Raquel Urtasun. Unisim: A neural closed-loop sensor simulator. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1389–1399, 2023. 2
- [38] Zeyu Yang, Hongye Yang, Zijie Pan, Xiatian Zhu, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. In *ICLR*, 2024. 3
- [39] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *Advances in neural information processing systems*, 35:25018–25032, 2022. 2
- [40] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. *arXiv preprint arXiv:2311.16493*, 2023. 3, 6
- [41] Zheng Zhang, Wenbo Hu, Yixing Lao, Tong He, and Hengshuang Zhao. Pixel-gs: Density control with pixel-aware gradient for 3d gaussian splatting. *arXiv preprint arXiv:2403.15530*, 2024. 2, 3, 6
- [42] Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Ewa volume splatting. In *Proceedings Visualization, 2001. VIS'01.*, pages 29–538. IEEE, 2001. 3