

CityGaussian: Real-time High-quality Large-Scale Scene Rendering with Gaussians

Yang Liu^{1,2}, Chuanchen Luo⁴, Lue Fan^{1,2}, Naiyan Wang⁵,
Junran Peng⁶✉, and Zhaoxiang Zhang^{1,2,3} ✉

¹ NLPR, MAIS, Institute of Automation, Chinese Academy of Sciences
² University of Chinese Academy of Sciences
³ Centre for Artificial Intelligence and Robotics ⁴ Shandong University
⁵ TuSimple ⁶ University of Science and Technology Beijing
{liuyang2022, lue.fan, zhaoxiang.zhang}@ia.ac.cn,
chuanchen.luo@sdu.edu.cn, winsty@gmail.com, jrpeng4ever@126.com

Abstract. The advancement of real-time 3D scene reconstruction and novel view synthesis has been significantly propelled by 3D Gaussian Splatting (3DGS). However, effectively training large-scale 3DGS and rendering it in real-time across various scales remains challenging. This paper introduces CityGaussian (CityGS), which employs a novel divide-and-conquer training approach and Level-of-Detail (LoD) strategy for efficient large-scale 3DGS training and rendering. Specifically, the global scene prior and adaptive training data selection enables efficient training and seamless fusion. Based on fused Gaussian primitives, we generate different detail levels through compression, and realize fast rendering across various scales through the proposed block-wise detail levels selection and aggregation strategy. Extensive experimental results on large-scale scenes demonstrate that our approach attains state-of-the-art rendering quality, enabling consistent real-time rendering of large-scale scenes across vastly different scales. Our project page is available at <https://dekulitesla.github.io/citygs/>.

Keywords: Large-Scale Scene Reconstruction · Novel View Synthesis · 3D Gaussian Splatting

1 Introduction

3D large-scale scene reconstruction, as a pivotal component in AR/VR [6, 11], aerial surveying [36], smart city [4, 8], and autonomous driving [34], has drawn extensive attention from academia and industry in recent decades. Such a task pursues high-fidelity reconstruction and real-time rendering at different scales for large areas that typically span over 1.5 km^2 [36]. In the past few years, this field has been dominated by neural radiance fields (NeRF) [21] based methods. Representative works include Block-NeRF [34], BungeeNeRF [41], and ScaNeRF [40]. But they still lack fidelity in details or exhibit sluggish performance.

Recently, 3D Gaussian Splatting (3DGS) [12] emerged as a promising alternative solution. In contrast to NeRF, it employs explicit 3D Gaussians as

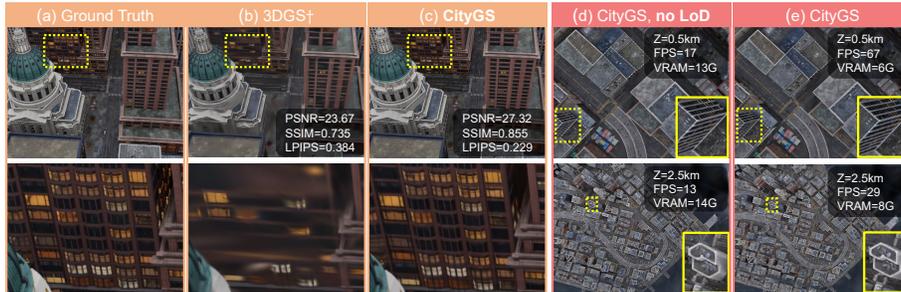


Fig. 1: (a, b, c) Our proposed CityGS achieves the SOTA rendering fidelity on Small City scene (5620 training images, 740 test images) of *MatrixCity* dataset. The setting of baseline 3DGS[†] is discussed in Sec. 4.1. (d, e) Here Z denotes camera height. Without LoD, CityGS would render over 20 million points, leading to considerable VRAM and time costs. The LoD saves VRAM and enables real-time performance under various scales. Note that the FPS is tested with CUDA synchronization for objective evaluation.

primitives to represent the scene. Thanks to highly efficient rasterization algorithm, 3DGS achieves high-quality visual effects at real-time rendering speed. Most existing academic exploration around 3DGS mainly focuses on objects or small scenes. However, devils emerge when 3DGS is applied to large-scale scene reconstruction. *On the one hand*, directly deploying 3DGS to large-scale scenes results in prohibitive overhead in GPU memory during training. For instance, the 24G RTX3090 raises out-of-memory errors when the Gaussian number grows above 11 million. But to reconstruct over 1.5 km^2 city area with high visual quality from aerial view, over 20 million Gaussians might be required. 3DGS of such capacity can't be directly trained even on 40G A100. *On the other hand*, the rendering speed bottleneck lies in depth sorting. As the number of Gaussians rises to the order of millions, the rasterization becomes extremely slow. For instance, a small *Train* scene of Tanks&Temples [13] dataset of 1.1 million Gaussians is rendered with an average visible Gaussians number of around 0.65 million and a speed of 103 FPS. But the 2.7 km^2 *MatrixCity* scene of 23 million Gaussians can only be rendered at the speed of 21 FPS even though the average visible Gaussians number is also around 0.65 million. And how to free unnecessary Gaussians from rasterization is the key to real-time large-scale scene rendering.

To address the problems mentioned above, we propose the **CityGaussian (CityGS)**. Inspired by MegaNeRF [36], we adopt a divide-and-conquer strategy. The whole scene is first partitioned into spatially adjacent blocks and parallelly trained. Each block is represented by much fewer Gaussians and trained with less data, the memory cost of a single GPU is well attenuated. For Gaussian partitioning, we contract unbounded regions to normalized bounded cubic and apply uniform grid division, to better balance workloads across different blocks. For training data partitioning, a pose is kept only if it is inside the considered block, or if the considered block has a considerable contribution to the rendering result. This novel strategy efficiently avoids distraction from irrelative data, while

enabling higher fidelity with less Gaussian consumption, as depicted in Tab. 4. To align trained Gaussians from adjacent blocks, we guide the training of each block with a coarse global Gaussian prior. This strategy can efficiently avoid interactive interference among blocks and enable seamless fusion.

To alleviate the computation burden when rendering large-scale Gaussians, we propose a block-wise Level-of-Detail (LoD) strategy. The key idea is to only feed necessary Gaussians to the rasterizer while eliminating extra computation costs. Specifically, we take the previously divided blocks as units to quickly decide which Gaussians are likely to be contained in the frustum. Furthermore, due to the perspective effect, distant regions occupy a small area of screen space and contain fewer details. Thus the Gaussian blocks that are far from the camera can be replaced with the compressed version, using fewer points and features. In this way, the processing burden of the rasterizer is significantly reduced, while the introduced extra computation remains acceptable. As illustrated in (d,e) of Fig. 1, our CityGS can maintain real-time large-scale scene rendering even under drastically large field of view.

In a nutshell, this work has three-fold contributions:

- We propose an effective divide-and-conquer strategy to reconstruct large-scale 3D Gaussian Splatting in parallel manner.
- With the proposed LoD strategy, we realize real-time large-scale scene rendering under drastically different scales with minimal quality loss.
- Our method, termed as CityGS, performs favorably against current state-of-the-art methods in public benchmarks.

2 Related Works

2.1 Neural Rendering

Neural Radiance Field is an instrumental technique for 3D scene reconstruction and novel view synthesis. As an implicit neural scene representation, it employs Multilayer Perceptrons (MLPs) as the mapping function between query positions and corresponding radiances. Volumetric rendering is then applied to render such a representation to 2D images. The success of NeRF has spawned a wide range of follow-up works [2, 3, 19, 20, 25–27, 35, 42] that improve upon various aspects of the original method. Among them, Mip-NeRF360 [3] serves as a recent milestone with outstanding rendering quality. However, NeRFs suffer from intensive sampling along emitted rays, resulting in relatively high training and inference latency. A series of methods [5, 10, 23, 32, 46] have been proposed to alleviate this problem. And the most representatives include InstantNGP [23] and Plenoxels [10]. Combined with a multiresolution hash grid and a small neural network, InstantNGP achieves speedup of several orders of magnitude while maintaining high image quality. On the other hand, Plenoxels represents the continuous density field with a sparse voxel grid, to get considerable speedup and outstanding performance together.

Point-based Rendering Another parallel line of works renders the scene with point-based representation. Such explicit geometry primitives enable fast rendering speed and high editability. Pioneering works include [14, 39, 44, 45, 51], but discontinuity in rendered images remains a problem. The recently proposed 3D Gaussian Splatting (3DGS) [12] solved this problem by using 3D Gaussians as primitives. Combined with a highly optimized rasterizer, the 3DGS can achieve superior rendering speed over the NeRF-based paradigm with no loss of visual fidelity. Nevertheless, the explicit millions of Gaussians with high dimensional features (RGB, spherical harmonics, etc) lead to significant memory and storage footprint. To mitigate this burden, methodologies such as [9, 15, 22] are proposed. Apart from vector quantization [24], [9] and [22] further combine distillation and 2D grid decomposition respectively to compress storage while improving speed. Joo et al. [15] realize similar performance by exploiting the geometrical similarity and local feature similarity of Gaussians. Despite the success in compression and speedup, these researches concentrate on small scenes or single objects. Under large-scale scenes, the excessive memory cost and computation burden lead to difficulty in high-quality reconstruction and real-time rendering. And our CityGS serves as an efficient solution for these problems.

2.2 Large Scale Scene Reconstruction

3D reconstruction from large image collections has been an aspiration for many researchers and engineers for decades. Photo Tourism [30] and Building Rome in a Day [1] are two representatives of early exploration in robust and parallel large-scale scene reconstruction. With the advance of NeRF [21], the paradigm has shifted. In Block-NeRF [34] and Mega-NeRF [36], the divide-and-conquer strategy is adopted and each divided block is represented by a small MLP. Switch-NeRF [49] further improves the performance by introducing learnable scene decomposition strategy. Urban Radiance Fields [28] and SUDS [37] further explores the large scene reconstruction with modalities beyond RGB images, such as LiDAR and 2D optical flow. To better balance the model storage and performance, Grid-NeRF [43] introduces guidance from the multiresolution feature planes, while GP-NeRF [48] and City-on-Web [31] utilizes hash grid and multiresolution tri-plane representations. To realize real-time rendering, UE4-NeRF [11] transforms the divided sub-NeRF to polygonal meshes and combines the matured rasterization pipeline in Unreal Engine 4. ScaNeRF [40] further counteracts the crux of inaccurate camera poses in realistic scenes. VastGaussian [17] explores the application of 3DGS under large-scale scenes and deals with appearance variation. However, the low rendering speed of such a large scene remains a problem. In contrast, our approach combines 3D Gaussian primitives with well-designed training and LoD methodology, elevating real-time render quality by a large margin.

2.3 Level of Detail

In computer graphics, Level of Detail (LoD) techniques regulate the amount of detail used to represent the virtual world so as to bridge complexity and performance [18]. Typically, LoD decreases the workload of objects that are becoming less important (e.g. moving away from the viewer). In recent years, the incorporation of LoD and the neural radiance field has received extensive concern. BungeeNeRF [41] equips NeRF with a progressive growing strategy, where each residual block is responsible for a finer detail level. On the other hand, NGLoD [33] represents different detail levels of neural signed distance functions (SDFs) with multi-resolution sparse voxel octree, while VQ-AD [32] uses hierarchical feature grid to compactly represent LoD of 3D signals. Motivated by Mip-NeRF [21], both Tri-MipRF and LoD-NeuS [50] apply cone casting with multiresolution tri-plane representation for anti-aliased LoD. City-on-Web [31] generates coarser detail levels by training feature grid of lower resolution and corresponding deferred MLP decoder. By realizing LoD on explicit 3D Gaussian representation, we efficiently improve its real-time performance under large-scale scenes.

3 Method

Overview The training and rendering pipelines are respectively shown in Fig. 2 and Fig. 3. We first generate a 3DGS that offers global scene depiction with normal 3DGS training strategy described in Sec. 3.1. Building upon this global prior, we employ the training strategy presented in Sec. 3.2 to adaptively divide the Gaussian primitives and data for further parallel training. Based on the fused large-scale Gaussians, our Level of Detail (LoD) algorithm discussed in Sec. 3.3 dynamically selects required Gaussians for fast rendering.

3.1 Preliminary

We begin with a brief introduction of 3DGS [12]. The 3DGS represents the scene with discrete 3D Gaussians $\mathbf{G}_{\mathbf{K}} = \{G_k | k = 1, \dots, K\}$, each of them is equipped with learnable properties including 3D position $\mathbf{p}_{\mathbf{k}} \in \mathbb{R}^{3 \times 1}$, opacity $\alpha_k \in [0, 1]$, geometry (i.e. scaling and rotation used to construct Gaussian covariance), spherical harmonics (SH) features $\mathbf{f}_{\mathbf{k}} \in \mathbb{R}^{3 \times 16}$ for view-dependent color $\mathbf{c}_{\mathbf{k}} \in \mathbb{R}^{3 \times 1}$. In rendering, given the intrinsics κ and pose τ_i of i -th image, the Gaussians are splatted to screen space, sorted in depth order, and rendered via alpha blending:

$$c_i(\mathbf{x}) = \sum_{k=1}^K \alpha_k \mathbf{c}_{\mathbf{k}} G_k^{2D}(\mathbf{x}) \prod_{t=1}^{k-1} (1 - \alpha_t G_t^{2D}(\mathbf{x})), G_k^{2D} = \mathbf{proj}(G_k, \kappa, \tau_i), \quad (1)$$

where \mathbf{proj} is projection operation, $c_i(\mathbf{x})$ is the color at pixel position \mathbf{x} , G_k^{2D} is projected Gaussian distribution. We refer the readers to the original paper [12]

for details. The final rendered image is denoted as $I_{\mathbf{G}_K}(\tau_i)$. During training, a typical initialization choice is the point cloud generated with the Structure-from-Motion (SfM), such as COLMAP [29]. Then based on gradients derived from differentiable rendering, the Gaussians would be cloned, densified, pruned, and consistently refined. However, the depiction of large-scale scenes can consume over 20 million primitives, which easily causes out-of-memory errors in training process, and rendering time is slowed down as well.

3.2 Training CityGS

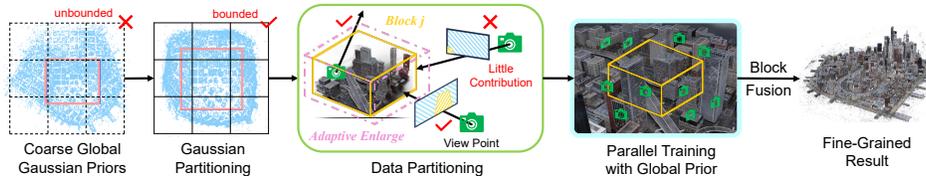


Fig. 2: The training process of CityGS. The pink square bounds the foreground area, facilitating subsequent unbounded space contraction and Gaussian partitioning. Then for a specific block, a pose is assigned to training set if it is inside the block or if the block has a considerable contribution to the rendering result. These blocks are then parallelly trained and merged together to depict the whole scene.

In this section, we first illustrate the necessity of a coarse global Gaussian prior and how to generate it. Based on this prior, we describe the Gaussian and data primitives division strategy. At the end of this section, we present the training and post-processing details. The pipeline is shown in Fig. 2.

Global Gaussian Prior Generation This part serves as the basis for further Gaussian and data division. An intuitive large-scale scene training strategy involves applying a divide-and-conquer strategy to the COLMAP points. However, due to a lack of depth constraints and global awareness, many geometrically inaccurate floaters will be generated to overfit regions outside the block, making a reliable fusion of different blocks difficult. Additionally, the rendered images from COLMAP points tend to be blurred and inaccurate, making it difficult to assess whether a specific view is important to training the block. To this end, we propose a simple yet effective way to solve the problems. Specifically, we first train the COLMAP points with all observations for 30,000 iterations, yielding a coarse description of the overall geometry distribution. The resulting set of Gaussian primitives are denoted as $\mathbf{G}_K = \{G_k | k = 1, \dots, K\}$, where K is the total amount. In further block-wise finetuning, such a strong global geometry prior leads point to geometrically appropriate positions, eliminating severe interference in fusion. As proved in Tab. 4, this strategy can efficiently improve

rendering fidelity. Moreover, this coarse Gaussian provides more accurate geometry distribution and cleaner rendered images, facilitating subsequent primitives and data division.

Primitives and Data Division Considering most real-world scenes are unbounded [35], the optimized Gaussians could extend indefinitely. Directly applying uniform grid division on the original 3D space would lead to many almost empty blocks, thus making workloads severely unbalanced. To alleviate this problem, we first contract the global Gaussian prior to a bounded cubic region.

For contraction, the inner foreground region, i.e. the pink square in Fig. 2, contains a linear space mapping, while the outer background region contains a non-linear mapping. Specifically, we depict the foreground region with the minimum and maximum positions of its corners \mathbf{p}_{\min} and \mathbf{p}_{\max} . Then we normalize Gaussian positions as $\hat{\mathbf{p}}_k = 2(\mathbf{p}_k - \mathbf{p}_{\min}) / (\mathbf{p}_{\max} - \mathbf{p}_{\min}) - 1$. Consequently, the positions of foreground Gaussians fall within the range $[-1, 1]$. The subsequent contraction is executed using the following function [40]:

$$\mathbf{contract}(\hat{\mathbf{p}}_k) = \begin{cases} \hat{\mathbf{p}}_k, & \text{if } \|\hat{\mathbf{p}}_k\|_\infty \leq 1, \\ \left(2 - \frac{1}{\|\hat{\mathbf{p}}_k\|_\infty}\right) \frac{\hat{\mathbf{p}}_k}{\|\hat{\mathbf{p}}_k\|_\infty}, & \text{if } \|\hat{\mathbf{p}}_k\|_\infty > 1. \end{cases} \quad (2)$$

By evenly partitioning this contracted cubic space, a more balanced Gaussian partitioning is derived.

In the finetuning phase, we hope each block is sufficiently trained. To be specific, the assigned data should keep training progress focused on refining details within the block. So one pose needs to be retained only when the considered block projects a significant amount of visible content to the rendering results. Distractive cases like severe occlusion or minor content contribution should be occluded. Since SSIM loss can efficiently capture the structural difference and is to some extent insensitive to brightness changes [38], we take it as the basis of our data partition strategy.

Specifically, for the j -th block, the containing coarse global Gaussians are denoted as $\mathbf{G}_{\mathbf{K}_j} = \{G_k | \mathbf{b}_{j,\min} \leq \mathbf{contract}(\hat{\mathbf{p}}_k) < \mathbf{b}_{j,\max}, k = 1, \dots, K_j\}$, where $\mathbf{b}_{j,\min}$ and $\mathbf{b}_{j,\max}$ defines the x,y,z bound of block j , and K_j is the number of contained Gaussians. Then whether the i -th pose τ_i is assigned to j -th block is determined by:

$$\mathbf{B}_1(\tau_i, \mathbf{G}_{\mathbf{K}_j}) = \begin{cases} 1, & L_{\text{SSIM}}(I_{\mathbf{G}_{\mathbf{K}}}(\tau_i), I_{\mathbf{G}_{\mathbf{K}} \setminus \mathbf{G}_{\mathbf{K}_j}}(\tau_i)) > \varepsilon, \\ 0, & \text{otherwise,} \end{cases} \quad (3)$$

where $\mathbf{G}_{\mathbf{K}} \setminus \mathbf{G}_{\mathbf{K}_j}$ defines difference set of $\mathbf{G}_{\mathbf{K}}$ and $\mathbf{G}_{\mathbf{K}_j}$. The SSIM loss L_{SSIM} larger than threshold ε means a considerable contribution of block j to the rendered image and thus leads to an assignment.

However, solely relying on the first principle can lead to artifacts when viewing outside at the edge of the block. Because these cases rarely involve the projection of considered block, they will not be sufficiently trained under the

first principle. Therefore we also include poses that fall into considered blocks, i.e.

$$B_2(\tau_i, \mathbf{G}_{\mathbf{K}_j}) = \begin{cases} 1, & \mathbf{b}_{j,\min} \leq \mathbf{contract}(\hat{\mathbf{p}}_{\tau_i}) < \mathbf{b}_{j,\max}, \\ 0, & \text{otherwise,} \end{cases} \quad (4)$$

where $\hat{\mathbf{p}}_{\tau_i}$ is the position under world coordinate of pose i . And the final assignment is:

$$\mathbf{B}(\tau_i, \mathbf{G}_{\mathbf{K}_j}) = \mathbf{B}_1(\tau_i, \mathbf{G}_{\mathbf{K}_j}) + \mathbf{B}_2(\tau_i, \mathbf{G}_{\mathbf{K}_j}). \quad (5)$$

Despite having the above strategies in place, empty blocks may still exist in cases of extremely uneven distributions or high block dimensions. To prevent overfitting, we enlarge the bound $\mathbf{b}_{j,\min}$ and $\mathbf{b}_{j,\max}$ until K_j exceeds certain threshold. This process is only used in data assignment to ensure enough training data for each block.

Finetuning and Post-processing After data and primitives division, we proceed to train each block in parallel. **It is worth noting this fine-tuning stage is under original uncontracted space.** Specifically, we utilize the coarse global prior generated in Sec. 3.2 to initialize the finetuning of each block. The training loss follows the approach outlined in the original 3DGS paper [12], comprising a weighted sum of $L1$ loss and SSIM loss. Then for each block, we filter out the finetuned Gaussians contained within its spatial bound. Thanks to the global geometric prior, interference among blocks is significantly mitigated. Thus a high-quality overall model can be derived through direct concatenation. Additional qualitative validation can be found in the Appendix. Further refinement is left in the LoD part.

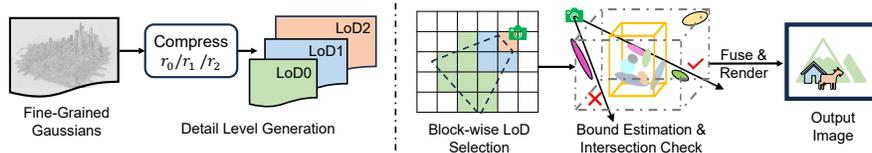


Fig. 3: Rendering of CityGS. Based on trained dense Gaussians, we generate detail levels with different compression rates r_0 , r_1 and r_2 . When rendering, all the Gaussians in the same block will share the same detail level, which is determined by the block’s distance to the camera. Since the contraction-based block partition leads to some irregular block shapes, we estimate their bounding boxes after removing floaters. The frustum intersection with the estimated block shape determines whether the block will be fed to rasterizer.

3.3 Level-of-Detail on CityGS

As discussed in Sec. 1, to eliminate the computation burden brought by unnecessary Gaussians to rasterizer, our CityGS involves multiple detail levels gener-

ation and block-wise visible Gaussian selection. We introduce these two parts respectively in the following two subsections.

Detail Level Generation As objects move away from the camera, they occupy less area on the screen, while contributing less high-frequency information. Thus distant, low detail-level regions can be well represented by models of low capacity, i.e. fewer points, less feature dimension, and lower data precision. In practice, we generate different detail levels with the advancing compression strategy LightGaussian [9], which operates directly on trained Gaussians, achieving a substantial compression rate with minimal performance degradation. Consequently, the memory and computation demand for required Gaussians is significantly alleviated, while the rendering fidelity is still well maintained.

Detail Level Selection and Fusion A baseline of detail level selection is to fill frustum regions between different distance intervals with Gaussians from the corresponding detail level. However, this method necessitates per-point distance calculation and assignment, resulting in significant computational overhead, as confirmed in Sec. 4.4. Therefore, we adopt a block-wise strategy, considering spatially adjacent blocks as the unit, as depicted in the right part of Fig. 3. Each block is considered as a cubic with eight corners for calculation of frustum intersection. All the Gaussians contained in certain blocks will share the same detail level, which is determined by the minimum distance from eight corners to the camera center. However, in practice, we found the minimum and maximum coordinates of Gaussians are usually determined by floaters. The resulting volume would be unreasonably enlarged, leading to many fake intersections. To avoid the influence of these floaters, we take the Median Absolute Deviation (MAD) [7] algorithm. The bounds of the j -th block, denoted as \mathbf{p}_{\min}^j and \mathbf{p}_{\max}^j , are determined by:

$$\begin{aligned} MAD_j &= \text{median} \left(\left| \mathbf{p}_k^j - \text{median} \left(\mathbf{p}_k^j \right) \right| \right), \\ \mathbf{p}_{\min}^j &= \max \left(\min \left(\mathbf{p}_k^j \right), \text{median} \left(\mathbf{p}_k^j \right) - n_{MAD} \times MAD_j \right), \\ \mathbf{p}_{\max}^j &= \min \left(\max \left(\mathbf{p}_k^j \right), \text{median} \left(\mathbf{p}_k^j \right) + n_{MAD} \times MAD_j \right), \end{aligned} \quad (6)$$

where n_{MAD} is the hyper-parameter. By choosing the appropriate n_{MAD} , this method can capture the bound of the block more accurately.

After that, all the corners in front of the camera will be projected into screen space. The minimum and maximum of these projected points composite a bounding box. By calculating its Intersection-over-Union (IoU) with screen area, we can check if the block has an intersection with the frustum. Along with the block where the camera is in, all visible blocks of corresponding detail level will be used for rendering.

In the fusion step, different detail levels are still comprehended via direct concatenation, which generates negligible discontinuity.

4 Experiments

4.1 Experiments setup

Dataset and Metrics Our algorithm is benchmarked on five scenes with various scales and environments. Specifically, we adopt the $2.7km^2$ Small City scene of synthetic city-scale dataset *MatrixCity* [16], as done in [31]. However, instead of solely training and evaluating on partial city area, we construct the whole city and compare performances. And we rescale the image width to 1600 pixels. We also carried out experiments on public real-world scene datasets, including *Residence*, *Sci-Art*, *Rubble*, and *Building* [36]. Following the approach in [36, 48, 49], the image resolution for these datasets is reduced by a factor of $4\times$. To further validate the generalization ability of our CityGS, we also test it on the street view scene Block_A of *MatrixCity* [16]. This challenging scene contains 4076 training images and 495 test images. To comprehensively measure the reconstruction quality of different methods, we take standard **SSIM**, **PSNR**, and **LPIPS** as our metrics [47]. We also compare **FPS** to evaluate the rendering speed. It is noteworthy that we synchronize all CUDA streams before measuring time, ensuring an objective evaluation of render times for each frame.

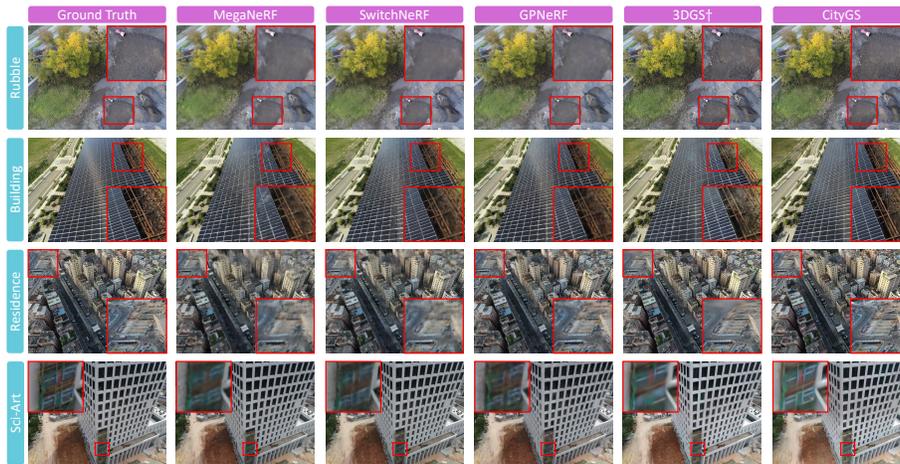


Fig. 4: Qualitative comparison with SOTA methods on real-scene datasets.

Implementations and Baselines Our method first trains coarse global Gaussian priors for further block-wise refinement. The training of priors for *Sci-Art*, *Residence* and *Rubble* adheres to the parameter settings outlined in 3DGS [12]. But for *Building* and *MatrixCity*, we halve the learning rate of position and scaling to prevent underfitting caused by aggressive optimization. In the fine-tuning

phase, we train each block for another 30,000 iterations with coarse global Gaussian prior as initialization. Furthermore, the learning rate of position is reduced by 60%, while that of scaling is empirically reduced by 20%, compared to setting in 3DGS [12]. The foreground area of contraction function is chosen as the central 1/3 area, and detailed block dimensions can be found in the Appendix. Our method is benchmarked against Mega-NeRF [36], Switch-NeRF [49], GP-NeRF [48], and 3DGS [12]. Since the datasets contain thousands of images, which are much larger than that used in the original 3DGS paper [12], we adjust the total iteration to 60,000, while densifying Gaussians from iteration 1,000 to 30,000 with an iteration interval of 200. This strong baseline is named as 3DGS[†]. In LoD, we evaluate on the *MatrixCity* dataset. We use 3 detail levels, where LoD 2 is the finest and LoD 0 is the coarsest. n_{MAD} mentioned in Sec. 3.3 is set to 4. The blocks within 0m to 200m are represented by LoD 2, blocks within 200m to 400m are represented by LoD 1, while others are represented by LoD 0.

4.2 Comparison with SOTA



Fig. 5: Qualitative comparison with SOTA methods on *MatrixCity* dataset.

For an apple-to-apple rendering quality comparison with the SOTA large-scale scene reconstruction strategy, we use the performance of the CityGS with no LoD. The quantitative result is shown in Tab. 1. Due to page limits, we put results of *Sci-Art* in the Appendix. It can be observed that our method outperforms the NeRF-based baselines by a large margin. On the one hand, as far as we know, our method is the first of all attempts that successfully reconstructs the whole *MatrixCity* with highly variable camera altitude, ranging from 150m to 500m [16]. The **PSNR** arrives at 27.46 and qualitative results presented in Fig. 5 also validate the high fidelity of our renderings. Compared with the baseline 3DGS[†], our CityGS can capture much richer details. More visualizations on *MatrixCity* can be found in the Appendix. On the other three realistic scenes, our method enables much higher **SSIM** and **LPIPS**, which indicates outstanding visual quality. As shown in Fig. 4, thin structures like girder steel and window frames can be well reconstructed. The details of complex

structures like grass and rubble are well recovered as well. In addition, results in Tab. 3 validates that our training strategy can generalize well to street scenes.

Table 1: Quantitative Comparison on four large-scale scene datasets. The ‘-’ symbol indicates Mega-NeRF [36] and Switch-NeRF [49] were not evaluated on *MatrixCity* due to difficulties in adjusting its training configurations beyond the provided, resulting in poor performance on this dataset. ‘n/a’ means the number of Gaussians ‘#GS’ does not apply to NeRF-based methods. The best results of each metric are in **bold**.

Metrics	MatrixCity				Residence				Rubble				Building			
	SSIM↑	PSNR↑	LPIPS↓	#GS	SSIM↑	PSNR↑	LPIPS↓	#GS	SSIM↑	PSNR↑	LPIPS↓	#GS	SSIM↑	PSNR↑	LPIPS↓	#GS
MegaNeRF [36]	-	-	-	n/a	0.628	22.08	0.489	n/a	0.553	24.06	0.516	n/a	0.547	20.93	0.504	n/a
Switch-NeRF [49]	-	-	-	n/a	0.654	22.57	0.457	n/a	0.562	24.31	0.496	n/a	0.579	21.54	0.474	n/a
GP-NeRF [43]	0.611	23.56	0.630	n/a	0.661	22.31	0.448	n/a	0.565	24.06	0.496	n/a	0.566	21.03	0.486	n/a
3DGS [†] [12]	0.735	23.67	0.384	9.7M	0.791	21.44	0.236	6.6M	0.777	25.47	0.2774	6.1M	0.720	20.46	0.305	6.4M
Ours	0.865	27.46	0.204	23.7M	0.813	22.00	0.211	10.8M	0.813	25.77	0.228	9.7M	0.778	21.55	0.246	13.2M

4.3 Level of Detail

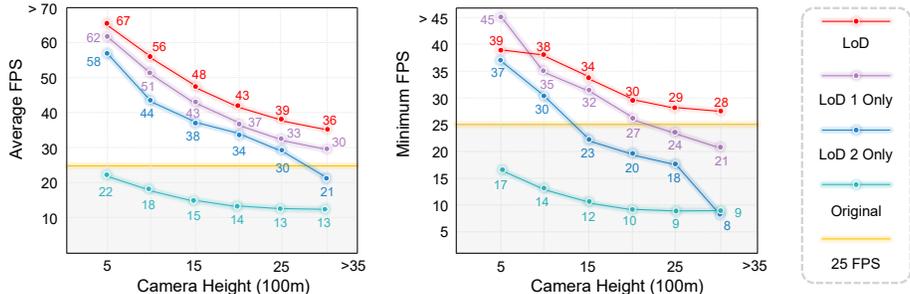


Fig. 6: Validation of LoD. We test rendering speed under looking-down cameras of different heights. The left part denotes **average** FPS, while the right part denotes **minimum** FPS. The shadowed area under 25 FPS line is the non-real-time zone.

Considering that *MatrixCity* has test split of over 700 images and various altitudes, we take it as our benchmark to evaluate effectiveness of our LoD strategy. Specifically, we first generate three detail levels with compression rates 50%, 34%, and 25%, namely LoD 2, LoD 1, and LoD 0. Our CityGS then applies the proposed LoD technique to combine all these detail levels. As shown in Tab. 2, the most fine-grained LoD 2 gains the best rendering quality, while the coarsest LoD 0 has the fastest rendering speed. Compared with the three detail levels, the version with LoD technique obtains SSIM and PSNR only second to LoD 2, while the speed is very close to LoD 1.

Table 2: Validation of LoD. Without considering the result with no LoD, the best performances are in **bold**, while the second best results are in **red**. Only LoD i means only using i -th detail level for rendering. And LoD means applying all three detail levels.

Models	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	FPS \uparrow
no-LoD	0.865	27.46	0.204	21.6
Only LoD 2	0.863	27.54	0.215	45.6
Only LoD 1	0.848	27.20	0.244	57.2
Only LoD 0	0.825	26.57	0.279	69.4
LoD	0.855	27.32	0.229	53.7

However, the camera altitude of *MatrixCity* contest split is bounded below 500m. To validate performance under extreme scale variance, we adjust the pitch to view straightly downside with appointed altitude, while keeping other pose attributes unchanged. The rendering results from different altitudes can be found in parts (d) and (e) of Fig. 1. The visual disparity between versions with and without LoD can be observed only when zooming deeply in.

In Fig. 6, we present how the mean and minimum rendering speed changes with different camera heights. LoD 0 is not considered for its inferior rendering quality. Comparing the mean FPS, both results from LoD and single detail level greatly outperform vanilla by a large margin. And the LoD version wins the highest speed at all heights. Comparing the minimum FPS, only the LoD version gains consistent real-time performance under worst cases of various heights, while the worst speed of single detail level drops dramatically as the camera lifts up. In other words, the proposed LoD helps smooth real-time transitions among drastically varying scales.

4.4 Ablation

In this section, we analyze the impact of training strategies as shown in Tab. 4. As a baseline, we trained using cameras in a $1.5\times$ larger area of each block. The results in the first and last lines of Tab. 4 indicate that this method significantly underperforms compared to our CityGS. The main issue is floaters caused by overfitting areas outside the block, but this is mitigated with guidance from the global model, improving performance as seen in the second line of Tab. 4. However, this improved baseline uses $1.5\times$ more Gaussians than CityGS, making it inefficient. The last three lines demonstrate that Eq.(3) is crucial for data partitioning, and we found Eq.(4) helps prevent floaters at block edges. Further hyper-parameter ablations are detailed in the Appendix.

Table 3: Quantitative Comparison on *MatrixCity* street scene Block_A. The best results of each metric are in **bold**.

Method	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow
MipNeRF360 [3]	0.717	22.00	0.488
3DGS \dagger [12]	0.701	21.14	0.441
CityGS	0.808	22.98	0.301

Table 4: Ablation on training strategy. The experiment is conducted on the *Rubble* dataset. The baseline denotes simply selecting cameras within the $1.5\times$ larger area of the block for training. The best performances are in **bold**.

	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	#GS
baseline	0.779	23.98	0.251	12.2
w global, baseline	0.801	25.01	0.227	15.4
w/o Eq.(3), Ours	0.749	23.43	0.252	19.7
w/o Eq.(4), Ours	0.810	25.45	0.233	9.6
Full, Ours	0.813	25.77	0.231	9.7

Table 5: Ablation on LoD Strategy. In-Frustum denotes the strategy used to select Gaussians within frustum. The block-wise and point-wise strategy are discussed in Sec. 3.3. Distance Interval here denotes the distance interval used by different detail levels. The area within the near interval would be represented by a higher detail level. The best results are in **bold**.

In-Frustum	Distance Interval (m)	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	FPS \uparrow
block-wise	[0,200],[200,400],[400, ∞]	0.855	27.32	0.229	53.7
point-wise	[0,200],[200,400],[400, ∞]	0.849	27.18	0.239	30.3
block-wise	[0,150],[150,300],[300, ∞]	0.848	27.16	0.242	57.6
block-wise	[0,250],[250,500],[500, ∞]	0.858	27.39	0.223	46.4

The quantitative ablation on LoD strategy is shown in Tab. 5. Comparing the first and the second row, the computation burden introduced by the point-wise strategy leads to the considerably worse real-time performance. Its rendering quality is on par with the smaller interval setting, i.e. third row. Comparing the first row with the third and the fourth row, it can be observed that with larger intervals for LoD 2 and LoD 1, the rendered results would contain richer details, but downgraded speed. And first row arrives at a balance, which is adopted as the standard setting.

5 Conclusions

In this work, we present CityGS, which successfully realizes real-time large-scale scene reconstruction with high fidelity. Through the blocking and LoD strategy tailored for Gaussian geometry, we obtain state-of-the-art rendering fidelity on mainstream benchmarks, while significantly reducing time costs when rendering drastically different scales of the same scene. However, the hidden static scene assumption limits its generalization ability. Training with the combination of drastically different views such as aerial and street views also degrades instead of boosting the performance of CityGS. The inner mechanism deserves to be further explored and well-resolved.

Acknowledgements

This work was supported in part by the National Key R&D Program of China (No. 2022ZD0116500), the National Natural Science Foundation of China (No. U21B2042, No. 62320106010), and in part by the 2035 Innovation Program of CAS, and the InnoHK program.

References

1. Agarwal, S., Furukawa, Y., Snavely, N., Simon, I., Curless, B., Seitz, S.M., Szeliski, R.: Building rome in a day. *Communications of the ACM* **54**(10), 105–112 (2011)
2. 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. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5855–5864 (2021)
3. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5470–5479 (2022)
4. Chaturvedi, K., Kolbe, T.H.: Integrating dynamic data and sensors with semantic 3d city models in the context of smart cities. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **4**, 31–38 (2016)
5. Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: Tensorf: Tensorial radiance fields. In: *European Conference on Computer Vision*. pp. 333–350. Springer (2022)
6. Chen, G., Wang, W.: A survey on 3d gaussian splatting. *arXiv preprint arXiv:2401.03890* (2024)
7. Dodge, Y.: *The concise encyclopedia of statistics*. Springer Science & Business Media (2008)
8. Dong, Q., Shu, M., Cui, H., Xu, H., Hu, Z.: Learning stratified 3d reconstruction. *Science China Information Sciences* **61**, 1–16 (2018)
9. 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)
10. Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5501–5510 (2022)
11. Gu, J., Jiang, M., Li, H., Lu, X., Zhu, G., Shah, S.A.A., Zhang, L., Bennamoun, M.: Ue4-nerf: Neural radiance field for real-time rendering of large-scale scene. *arXiv preprint arXiv:2310.13263* (2023)
12. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4) (2023)
13. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)* **36**(4), 1–13 (2017)
14. Lassner, C., Zollhofer, M.: Pulsar: Efficient sphere-based neural rendering. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1440–1449 (2021)
15. 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)

16. Li, Y., Jiang, L., Xu, L., Xiangli, Y., Wang, Z., Lin, D., Dai, B.: Matrixcity: A large-scale city dataset for city-scale neural rendering and beyond. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3205–3215 (2023)
17. Lin, J., Li, Z., Tang, X., Liu, J., Liu, S., Liu, J., Lu, Y., Wu, X., Xu, S., Yan, Y., Yang, W.: Vastgaussian: Vast 3d gaussians for large scene reconstruction. In: CVPR (2024)
18. Luebke, D.: Level of detail for 3D graphics. Morgan Kaufmann (2003)
19. Martin-Brualla, R., Radwan, N., Sajjadi, M.S., Barron, J.T., Dosovitskiy, A., Duckworth, D.: Nerf in the wild: Neural radiance fields for unconstrained photo collections. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7210–7219 (2021)
20. Mildenhall, B., Hedman, P., Martin-Brualla, R., Srinivasan, P.P., Barron, J.T.: Nerf in the dark: High dynamic range view synthesis from noisy raw images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16190–16199 (2022)
21. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
22. Morgenstern, W., Barthel, F., Hilsmann, A., Eisert, P.: Compact 3d scene representation via self-organizing gaussian grids. *arXiv preprint arXiv:2312.13299* (2023)
23. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)* **41**(4), 1–15 (2022)
24. Navaneet, K., Meibodi, K.P., Koochpayegani, S.A., Pirsiavash, H.: Compact3d: Compressing gaussian splat radiance field models with vector quantization. *arXiv preprint arXiv:2311.18159* (2023)
25. Niemeyer, M., Barron, J.T., Mildenhall, B., Sajjadi, M.S., Geiger, A., Radwan, N.: Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5480–5490 (2022)
26. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10318–10327 (2021)
27. Reiser, C., Szeliski, R., Verbin, D., Srinivasan, P., Mildenhall, B., Geiger, A., Barron, J., Hedman, P.: Merf: Memory-efficient radiance fields for real-time view synthesis in unbounded scenes. *ACM Transactions on Graphics (TOG)* **42**(4), 1–12 (2023)
28. Rematas, K., Liu, A., Srinivasan, P.P., Barron, J.T., Tagliasacchi, A., Funkhouser, T., Ferrari, V.: Urban radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12932–12942 (2022)
29. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
30. Snavely, N., Seitz, S.M., Szeliski, R.: Photo tourism: exploring photo collections in 3d. In: *ACM siggraph 2006 papers*, pp. 835–846 (2006)
31. Song, K., Zhang, J.: City-on-web: Real-time neural rendering of large-scale scenes on the web. *arXiv preprint arXiv:2312.16457* (2023)
32. Takikawa, T., Evans, A., Tremblay, J., Müller, T., McGuire, M., Jacobson, A., Fidler, S.: Variable bitrate neural fields. In: *ACM SIGGRAPH 2022 Conference Proceedings*. pp. 1–9 (2022)

33. Takikawa, T., Litalien, J., Yin, K., Kreis, K., Loop, C., Nowrouzezahrai, D., Jacobson, A., McGuire, M., Fidler, S.: Neural geometric level of detail: Real-time rendering with implicit 3d shapes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11358–11367 (2021)
34. Tancik, M., Casser, V., Yan, X., Pradhan, S., Mildenhall, B., Srinivasan, P.P., Barron, J.T., Kretschmar, H.: Block-nerf: Scalable large scene neural view synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8248–8258 (2022)
35. Tancik, M., Weber, E., Ng, E., Li, R., Yi, B., Wang, T., Kristoffersen, A., Austin, J., Salahi, K., Ahuja, A., et al.: Nerfstudio: A modular framework for neural radiance field development. In: ACM SIGGRAPH 2023 Conference Proceedings. pp. 1–12 (2023)
36. Turki, H., Ramanan, D., Satyanarayanan, M.: 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. pp. 12922–12931 (2022)
37. Turki, H., Zhang, J.Y., Ferroni, F., Ramanan, D.: Suds: Scalable urban dynamic scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12375–12385 (2023)
38. Wang, Z., Simoncelli, E.P., Bovik, A.C.: Multiscale structural similarity for image quality assessment. In: The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003. vol. 2, pp. 1398–1402. Ieee (2003)
39. Wiles, O., Gkioxari, G., Szeliski, R., Johnson, J.: Synsin: End-to-end view synthesis from a single image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7467–7477 (2020)
40. Wu, X., Xu, J., Zhang, X., Bao, H., Huang, Q., Shen, Y., Tompkin, J., Xu, W.: Scannerf: Scalable bundle-adjusting neural radiance fields for large-scale scene rendering. *ACM Transactions on Graphics (TOG)* **42**(6), 1–18 (2023)
41. Xiangli, Y., Xu, L., Pan, X., Zhao, N., Rao, A., Theobalt, C., Dai, B., Lin, D.: Bungeenerf: Progressive neural radiance field for extreme multi-scale scene rendering. In: European conference on computer vision. pp. 106–122. Springer (2022)
42. Xu, D., Jiang, Y., Wang, P., Fan, Z., Shi, H., Wang, Z.: Sinnerf: Training neural radiance fields on complex scenes from a single image. In: European Conference on Computer Vision. pp. 736–753. Springer (2022)
43. Xu, L., Xiangli, Y., Peng, S., Pan, X., Zhao, N., Theobalt, C., Dai, B., Lin, D.: Grid-guided neural radiance fields for large urban scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8296–8306 (2023)
44. Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., Neumann, U.: Point-nerf: Point-based neural radiance fields. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (Jun 2022). <https://doi.org/10.1109/cvpr52688.2022.00536>, <http://dx.doi.org/10.1109/cvpr52688.2022.00536>
45. Yifan, W., Serena, F., Wu, S., Öztireli, C., Sorkine-Hornung, O.: Differentiable surface splatting for point-based geometry processing. *ACM Transactions on Graphics* p. 1–14 (Dec 2019). <https://doi.org/10.1145/3355089.3356513>, <http://dx.doi.org/10.1145/3355089.3356513>
46. Yu, A., Li, R., Tancik, M., Li, H., Ng, R., Kanazawa, A.: Plenotrees for real-time rendering of neural radiance fields. In: 2021 IEEE/CVF International Conference on Computer Vision (ICCV) (Oct 2021). <https://doi.org/10.1109/iccv48922.2021.00570>, <http://dx.doi.org/10.1109/iccv48922.2021.00570>

47. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)
48. Zhang, Y., Chen, G., Cui, S.: Efficient large-scale scene representation with a hybrid of high-resolution grid and plane features. arXiv preprint arXiv:2303.03003 (2023)
49. Zhenxing, M., Xu, D.: Switch-nerf: Learning scene decomposition with mixture of experts for large-scale neural radiance fields. In: The Eleventh International Conference on Learning Representations (2022)
50. Zhuang, Y., Zhang, Q., Feng, Y., Zhu, H., Yao, Y., Li, X., Cao, Y.P., Shan, Y., Cao, X.: Anti-aliased neural implicit surfaces with encoding level of detail. In: SIGGRAPH Asia 2023 Conference Papers. pp. 1–10 (2023)
51. Zwicker, M., Pfister, H., Van Baar, J., Gross, M.: Ewa volume splatting. In: Proceedings Visualization, 2001. VIS'01. pp. 29–538. IEEE (2001)