

A LoD of Gaussians: Out-of-Core Training and Rendering for Seamless Ultra-Large Scene Reconstruction

FELIX WINDISCH, Graz University of Technology, Austria
 THOMAS KÖHLER, Graz University of Technology, Austria
 LUKAS RADL, Graz University of Technology, Austria
 MATTIA D'URSO, Graz University of Technology, Austria
 MICHAEL STEINER, Graz University of Technology, Austria
 DIETER SCHMALSTIEG, Graz University of Technology, Austria and University of Stuttgart, Germany
 MARKUS STEINBERGER, Graz University of Technology, Austria and Huawei Technologies, Austria




Fig. 1. **Teaser:** We introduce a fully hierarchical 3D Gaussian representation trained directly across unstructured, multi-scale image sets—including street-level and far aerial views—without scene partitioning. Our method maintains a consistent global scene model, eliminating boundary artifacts typical of chunked approaches. A hybrid Level-of-Detail system combines Gaussian hierarchies with Sequential Point Trees, enabling dynamic, view-dependent streaming and LoD selection. The entire model resides in external memory, with only a small, adaptive subset loaded on demand, allowing seamless training and interactive rendering of scenes with 150M+ Gaussians on a single consumer GPU ($\leq 24\text{GB}$ VRAM).

Gaussian Splatting has emerged as a high-performance technique for novel view synthesis, enabling real-time rendering and high-quality reconstruction of small scenes. However, scaling to larger environments has so far relied on partitioning the scene into chunks—a strategy that introduces artifacts at chunk boundaries, complicates training across varying scales, and is poorly suited to unstructured scenarios such as city-scale flyovers combined with street-level views. Moreover, rendering remains fundamentally limited by GPU memory, as all visible chunks must reside in VRAM simultaneously. We introduce *A LoD of Gaussians*, a framework for training and rendering ultra-large-scale Gaussian scenes on a single consumer-grade

GPU without partitioning. Our method stores the full scene out-of-core (e.g., in CPU memory) and trains a Level-of-Detail (LoD) representation directly, dynamically streaming only the relevant Gaussians. A hybrid data structure combining Gaussian hierarchies with Sequential Point Trees enables efficient, view-dependent LoD selection, while a lightweight caching and view scheduling system exploits temporal coherence to minimize the loading overhead. Together, these innovations enable seamless multi-scale reconstruction and interactive visualization of complex scenes—from broad aerial views to fine-grained ground-level details.

Authors' Contact Information: Felix Windisch, felix.windisch@tugraz.at, Graz University of Technology, Graz, Austria; Thomas Köhler, Graz University of Technology, Austria, t.koehler@tugraz.at; Lukas Radl, Graz University of Technology, Austria, lukas.radl@tugraz.at; Mattia D'Urso, mattia.durso@tugraz.at, Graz University of Technology, Austria; Michael Steiner, michael.steiner@tugraz.at, Graz University of Technology, Austria; Dieter Schmalstieg, Graz University of Technology, Austria and University of Stuttgart, Germany, dieter.schmalstieg@visus.uni-stuttgart.de; Markus Steinberger, Graz University of Technology, Austria and Huawei Technologies, Austria.

CCS Concepts: • **Computing methodologies** → *Rasterization*; *Visibility*; *Machine learning*.

Additional Key Words and Phrases: Level of Detail, Gaussian Splatting, Large-Scale Reconstruction

1 Introduction

Given a set of posed images of a 3D scene, the task of novel view synthesis (NVS) is to generate plausible images of the scene from unseen viewpoints. Early approaches achieved this via image-based blending [Zhang et al. 2020], but the introduction of Neural Radiance Fields (NeRF) [Mildenhall et al. 2021] marked a breakthrough, enabling high-quality results by optimizing an implicit volumetric scene representation through a multi-layer perceptron. More recently, 3D Gaussian Splatting (3DGS) [Kerbl et al. 2023] extended this paradigm to an explicit representation: a set of Gaussian primitives that are efficiently rasterized using splatting techniques [Zwicker et al. 2001], replacing costly ray marching and allowing real-time rendering with fast convergence.

Despite these advances, both NeRF and 3DGS remain constrained by memory bottlenecks when applied to large-scale environments. Prior methods address this by dividing scenes into smaller chunks [Chen et al. 2025; Kerbl et al. 2024; Lin et al. 2024; Liu et al. 2024a; Tancik et al. 2022; Xu et al. 2023], training each independently before merging results. While chunking strategies mitigate memory usage during training, they introduce several key limitations:

- (1) **View-chunk misalignment:** Camera views often span multiple chunks, especially in open or multi-scale datasets (e.g., combining aerial and street-level images). As a result, chunk boundaries become arbitrary with respect to images, complicating scene partitioning and training, resulting in artifacts: *Chunk bleeding* occurs when Gaussians extend out of their assigned chunk and obscure neighbouring chunks after merging. *Chunk ghosting* occurs when an occluder present in a training image is not part of the current chunk, training its ‘ghostly’ outline into the wrong chunk. See Figures 7 and 12.
- (2) **Redundant overlap:** To avoid artifacts at chunk boundaries, regions are typically trained with significant overlap, which duplicates parameters and optimizer state, increasing memory usage and prolonging training time.
- (3) **Asymmetric hardware demands:** Although chunking reduces memory requirements during training, rendering may require all visible chunks in memory simultaneously—often exceeding the capacity of the original training setup and undermining the practical benefit of partitioning.

The simplest and most robust alternative to chunking is to avoid splitting altogether. With *A LoD of Gaussians*, we introduce a seamless pipeline that enables training and rendering of ultra-large-scale scenes directly on a single consumer-grade GPU, without any form of scene partitioning (see Figure 6). To handle scenes that exceed available VRAM, we store all Gaussian data in CPU RAM and dynamically stream only those visible from the current training view into GPU memory. Still, a single distant view could require access to the full scene. To address this, we construct a hierarchical Level-of-Detail (LoD) model inspired by Kerbl et al. [2024], loading detail proportional to view distance. Crucially, this hierarchy must be maintained throughout training, as Gaussian parameters and their spatial distribution evolve dynamically. We propose a novel hierarchy densification strategy, inspired by MCMC-style spawning [Kheradmand et al. 2024], to support stable and progressive refinement. Efficient view-dependent selection from the hierarchy

is challenging for large models, particularly when the hierarchy structure changes during optimization. Instead of full tree traversal, we adopt Sequential Point Trees (SPTs) [Dachsbacher et al. 2003], originally developed for point cloud rendering. Our Hierarchical SPT version allows us to compute the correct LoD cut efficiently for individual views and camera paths. Finally, to reduce CPU-GPU data transfer, we introduce a lightweight caching system that tracks recently used Gaussians and reuses them across training iterations. In summary, we make the following contributions:

- (1) **Seamless, non-partitioned training of ultra-large Gaussian scenes.** We present the first 3D Gaussian Splatting framework that enables training and interactive rendering of city-scale scenes from arbitrary views without spatial partitioning, using out-of-core memory and view-dependent streaming on a single consumer GPU.
- (2) **Dynamic LoD hierarchy densification during training.** We propose a novel coarse-to-fine hierarchy densification strategy that allows Gaussian LoD hierarchies to evolve continuously during optimization, supporting stable refinement and restructuring without post-training hierarchy construction.
- (3) **Hierarchical Sequential Point Trees for Gaussian Splatting.** We adapt Sequential Point Trees (SPTs) to Gaussian splatting and introduce the Hierarchical SPT (HSPT), a hybrid data structure enabling efficient, parallelizable LoD selection while remaining robust to hierarchy updates during training.
- (4) **Efficient out-of-core execution via cache-aware streaming and view scheduling.** We design a lightweight GPU caching and view selection strategy that exploits temporal coherence to substantially reduce CPU-GPU transfer overhead during both training and rendering.
- (5) **Large-scale evaluation and new benchmarks.** We introduce the Uni10k dataset and an expanded version of the MatrixCity Small City scene, and demonstrate state-of-the-art performance on large-scale, multi-view datasets spanning aerial, street-level, and indoor environments.

2 Related Work

Large Scale Reconstruction. Reconstructing large-scale scenes from images has long been a central challenge in visual computing. Traditional approaches relied on Structure-from-Motion (SfM) pipelines to recover geometry from unordered photo collections [Agarwal et al. 2009; Schönberger and Frahm 2016]. Differentiable rendering techniques, notably NeRF [Mildenhall et al. 2021] and 3DGS [Kerbl et al. 2023], marked a paradigm shift by optimizing volumetric scene representations. Extensions of NeRF to large scenes typically employ scene partitioning [Tancik et al. 2022; Xu et al. 2023] or multi-GPU training strategies [Li et al. 2024b]. Similarly, most large-scale 3DGS pipelines adopt chunk-based training: *Hierarchical-3DGS (H-3DGS)* [Kerbl et al. 2024] trains chunks independently and then merges them into a global LoD hierarchy; *CityGaussian* [Liu et al. 2024a] combines chunked training with per-chunk LoD selection using *LightGaussian* [Fan et al. 2024]; *OccluGaussian* [Liu et al. 2025] partitions the scene to maximize camera correlation in each chunk; and

VastGaussian [Lin et al. 2024] introduces decoupled appearance modeling and progressive partitioning. *Horizon-GS* [Jiang et al. 2025] integrates divide-and-conquer strategies with LoD mechanisms from Ren et al. [2024], specifically targeting hybrid aerial/street-view datasets. *GrendelGS* [Zhao et al. 2024] avoids spatial chunking by distributing training images across GPUs, such that each device renders a disjoint screen region. Another research direction focuses on extracting geometric proxies from large-scale 3DGS scenes [Chen et al. 2025; Li et al. 2025; Liu et al. 2024b]. These methods leverage TSDF fusion and geometric losses to generate multiple meshes, which are fused and rendered efficiently using traditional rasterization. Other approaches [Lee et al. 2025; Zhao et al. 2026] make use of out-of-core memory to scale training, but neglect LoD, which limits their memory reduction during training to frustum culling.

Level-of-Detail Rendering. Level-of-detail techniques reduce geometric complexity of distant scene content to accelerate rendering. In 3DGS, LoD approaches have mainly targeted efficient rendering on memory-constrained or mobile devices. Compression-based strategies include attribute quantization via codebooks, pruning low-impact Gaussians, and adapting the degree of spherical harmonics per primitive [Fan et al. 2024; Fang and Wang 2024; Huang et al. 2025; Niedermayr et al. 2024; Niemeyer et al. 2025; Papantonakis et al. 2024; Seo et al. 2024]. *Scaffold-GS* [Lu et al. 2024] uses latent vectors anchored to reference Gaussians, with an MLP generating associated Gaussians at render time. *Octree-GS* [Ren et al. 2024] extends this to hierarchical LoD rendering via spatial subdivision. *Virtualized 3D Gaussians* [Yang et al. 2025] targets composed scenes of reconstructed objects using an LoD scheme inspired by Unreal Engine 5’s *Nanite* [Karis et al. 2021].

3 Preliminaries of Hierarchical 3D Gaussian Splatting

3DGS [Kerbl et al. 2023] models a radiance field using a set of spatially distributed Gaussians, each with mean $\boldsymbol{\mu}_i \in \mathbb{R}^3$, RGB base colors $\mathbf{b}_i \in \mathbb{R}^3$ and covariance matrices $\Sigma_i = \mathbf{R}_i \mathbf{S}_i \mathbf{S}_i^\top \mathbf{R}_i^\top$, which are parameterized via a diagonal scaling matrix $\mathbf{S}_i = \text{diag}(s_i^1, s_i^2, s_i^3)$ and an orthonormal rotation matrix \mathbf{R}_i . Each Gaussian also stores an opacity σ_i and view-dependent color, modeled using spherical harmonics (SH) coefficients \mathbf{f}_i^d . The SH degree d controls expressiveness, with each Gaussian requiring $\sum_{j=1}^d 3 \cdot (2j + 1)$ parameters. For rendering, all N Gaussians are sorted by distance to the camera and a discrete approximation of the volume rendering equation is evaluated for every pixel \mathbf{x} with corresponding view direction \mathbf{v} :

$$\mathbf{C}(\mathbf{x}) = \sum_{i=1}^N \mathbf{c}_i(\mathbf{v}) \alpha_i(\mathbf{x}) \prod_{j=1}^{i-1} (1 - \alpha_j(\mathbf{x})), \quad (1)$$

where α_j is the opacity of the j -th Gaussian along the view ray:

$$\alpha_j(\mathbf{x}) = \sigma_j e^{-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}'_j) \boldsymbol{\Sigma}'_j (\mathbf{x} - \boldsymbol{\mu}'_j)^T}. \quad (2)$$

Here $\boldsymbol{\mu}'$ and $\boldsymbol{\Sigma}'$ denote the projected 2D mean and covariance on the image plane, obtained by applying an affine approximation of the projective transform [Zwicker et al. 2001].

3DGS Memory. Standard 3DGS pipelines store the full set of per-Gaussian attributes, training images, and optimizer state in GPU

memory (VRAM); see Figure 9 for a detailed per-Gaussian breakdown. Additional temporary allocations occur during forward and backward passes (e.g., for sorting and gradient accumulation). This overhead varies with the scale of the Gaussians and effectiveness of culling strategies, but can be roughly upper-bounded by around 800 bytes per Gaussian in practice. This limits typical training to roughly 500 000 Gaussians per GB of GPU memory, imposing strict constraints on the detail and extent of reconstructions.

Gaussian Hierarchies. As introduced in H-3DGS [Kerbl et al. 2024], Gaussian hierarchies recursively merge nearby Gaussians into a tree, where each non-leaf node approximates its children, and leaves correspond to the original Gaussians. A cut is defined by a condition $c_{\text{hier}}(i, \text{cam})$ evaluated in a breadth-first search (BFS). If a node satisfies this condition, it is added to the cut set and its children are skipped; otherwise, the BFS continues. A *proper cut set* includes no parents or children of any included node and thus provides a view-adaptive LoD representation.

The cut condition used in [Kerbl et al. 2024] is a simple cut-off according camera distance:

$$c_{\text{hier}}(i, \text{cam}) = \|\boldsymbol{\mu}_i - \mathbf{p}_{\text{cam}}\|_2 \geq m_d(i), \quad m_d(i) = \frac{T}{\max_j s_i^j}, \quad (3)$$

where \mathbf{p}_{cam} is the camera position, T is a global LoD threshold, and $m_d(i)$ is the minimum acceptable distance for viewing Gaussian i . The BFS ensures that if i is in the cut set, $\text{parent}(i)$ failed the condition: $m_d(\text{parent}(i)) > \|\boldsymbol{\mu}_i - \mathbf{p}_{\text{cam}}\|_2 \geq m_d(i)$.

4 Method

To train models that exceed GPU memory limits, all Gaussian attributes are stored in CPU RAM and streamed to the GPU on demand for each training view according to the LoD hierarchy. To accelerate the hierarchy cut, we store a copy of only the tree structure in VRAM, where larger subtrees are replaced by *Sequential Point Trees* (SPTs), forming a *hierarchical SPT* (HSPT). To minimize costly transfers between RAM and VRAM, we track which SPTs currently reside in GPU memory and at which detail. Only if an SPT is not present in this GPU cache at a sufficiently similar level of detail, will it be loaded from RAM. Densification is performed on the CPU by adding new leaf nodes to the hierarchy and respawning low-opacity leaf nodes. Following densification, the updated hierarchy is converted back to an HSPT and transferred to the GPU for a new round of training iterations. Suppl. A.2 includes more details on initialization and training. An overview of our training and densification process can be found in Figure 8 and pseudocode in Suppl. A.7.

4.1 Sequential Point Trees for Gaussian Splatting

Sequential Point Trees [Dachsbacher et al. 2003] were originally developed for point cloud LoD rendering, but can be trivially extended to ellipsoids and Gaussians. They enforce a more constrained but more efficient cut condition than Kerbl et al. [2024]:

$$c_{\text{SPT}}(i, \text{cam}) = m_d(\text{parent}(i)) > \|\boldsymbol{\mu}_{\text{root}} - \mathbf{p}_{\text{cam}}\|_2 \geq m_d(i). \quad (4)$$

This condition is evaluated for all Gaussians in parallel, using the shared root-camera distance $\|\boldsymbol{\mu}_{\text{root}} - \mathbf{p}_{\text{cam}}\|_2$. It requires storing only sorted pairs $(m_d(i), m_d(\text{parent}(i)))$, significantly reducing memory compared to full Gaussian hierarchies. To optimize cuts, Gaussians




Fig. 2. **Visualization of Hierarchy and Densification.** (a) An SPT and Gaussian hierarchy show the same 5 Gaussians at varying levels of detail (red lines indicate cuts). Vertical lines show binary search results; horizontal lines show distance cuts. (b) Leaf node densification and respawning.

are sorted by $m_d(\text{parent}(i))$ in descending order. A binary search determines the cutoff index N , above which Gaussians are too fine to be rendered. Note that cuts are guaranteed to be proper, with nodes where $m_d(\text{parent}(i)) > m_d(i)$ never being selected for the cut. Evidently, the level of detail of all Gaussians in the SPT is dictated by the camera’s distance to its root node. This can lead to Gaussians with $m_d(i) > \|\mu_i - \mathbf{p}_{\text{cam}}\|_2$ being rendered, even though they would be too coarse for the current view. To counteract this issue, we define: $M_d(i) = m_d(i) + \|\mu_i - \mathbf{p}_{\text{cam}}\|_2$, as a conservative minimum distance function. By the triangle inequality, selecting Gaussians satisfying $M_d(i) \leq \|\mu_{\text{root}} - \mathbf{p}_{\text{cam}}\|_2$ guarantees $m_d(i) \leq \|\mu_i - \mathbf{p}_{\text{cam}}\|_2$. In turn, this means that Gaussians that are further away from the camera than the root node will be selected at a higher level of detail. SPTs are best suited for tightly grouped Gaussians observed from distances greater than their mutual spacing. Their compact memory footprint and parallel evaluation make them well-suited for large-scale scenes. Figure 2 visualizes both hierarchy types and their LoD cuts.

4.2 Densification

Densifying an LoD representation presents a unique challenge, as the hierarchical structure must evolve continuously during training. Prior works circumvent this issue by constructing LoD hierarchies only after chunk-level training and densification are complete.

We take inspiration from *3DGS-MCMC* [Kheradmand et al. 2024], which ‘splits’ Gaussians, replacing them with two new Gaussians which together should appear similarly to the original. Notably, this mirrors how a parent node in a Gaussian hierarchy approximates its children. Motivated by this correspondence, we adopt this approach and instead ‘spawn’ two new child nodes for a Gaussian, increasing the size of the hierarchy with minimal artifacts. This procedure leads to gradual increase in detail during densification, thereby avoiding the instability on large scenes observed in [Kerbl et al. 2024].

Instead of pruning, *3DGS-MCMC* declares Gaussians below a certain opacity threshold as ‘dead’, and respawns them at the position of a high-opacity Gaussian. We propose a similar strategy: when a leaf node dies, its parent is replaced by its sibling node; the dead leaf node and its parent are then respawnd as children to another node, which is selected to be densified. See Figure 2b for an overview of the two hierarchy densification operations. Together, they ensure




Fig. 3. A Gaussian hierarchy is converted to an HSPT by cutting according to Gaussian volume and converting sufficiently large subtrees to SPTs. The HSPT can then be cut in a 2-step process.

that the hierarchy can be expanded during training in a stable and valid manner while being rebalanced as required.

While *3DGS-MCMC* choose Gaussians to densify using a random selection weighted by opacity, we employ the strategy from Kerbl et al. [2024], which selects Gaussians for densification based on their maximal screen-space gradient. This criterion better aligns densification with view-dependent reconstruction error in large-scale scenes. For further details on densification, see Suppl. A.6.

4.3 The Hierarchical SPT Datastructure

We first review previous LoD selection approaches and motivate the need for a new datastructure—the hierarchical SPT—for robust and efficient training.

BFS. Computing the cut set of a large Gaussian hierarchy is costly and must be done for every frame. A straightforward solution is a BFS from the root, which guarantees a proper cut and enables early pruning of large subtrees (e.g., via frustum culling). However, graph traversal poorly suited to parallel execution on the GPU, making this approach prohibitively expensive at scale.

Parallel Cut. To enable GPU-accelerated cuts, Kerbl et al. [2024] evaluate the cut condition in parallel for each Gaussian:

$$\left(m_d(i) < \|\mu_i - \mathbf{p}_{\text{cam}}\|_2\right) \wedge \left(m_d(\text{parent}(i)) \geq \|\mu_{\text{parent}(i)} - \mathbf{p}_{\text{cam}}\|_2\right), \quad (5)$$

where any Gaussian that is sufficiently small at its current camera distance and whose parent is too large to be rendered, should be part of the cut set. This produces a proper cut under the assumption that child Gaussians always have a smaller minimal distance than their parents (i.e. the heap condition is fulfilled): $\forall i : m_d(i) < m_d(\text{parent}(i))$. This is generally valid when hierarchies are constructed after training, since parent Gaussians represent coarser approximations of their children. However, when the hierarchy is modified during training and densification, optimization can break the heap condition—leading to invalid cut sets and degenerate hierarchies that worsen over time.

Hierarchical SPT (HSPT). Our HSPT data structure combines the benefits of both approaches. To construct it from a Gaussian hierarchy, we cut it using a BFS on the condition $c_{\text{HSPT}}(i) = s_i^1 \cdot s_i^2 \cdot s_i^3 < \text{size}$ with volume threshold size. The resulting cut set \mathbb{C}_{HSPT} partitions the hierarchy into the *upper hierarchy*, which includes all

Gaussians with volume greater than size, and the *lower hierarchy*, consisting of the subtrees rooted at the nodes in the cut set.

The volume of each root node in the lower hierarchy is now bounded by size, which also roughly bounds the extent of all Gaussians in the subtree. This provides an upper bound on the error introduced if the subtree is converted into an SPT. Consequently, each subtree of sufficient size in \mathbb{C}_{HSPT} can be transformed into an SPT to accelerate cut computation without violating cut correctness at higher levels. The HSPT-based cutting process then proceeds in two steps: first, a BFS on the upper hierarchy selects the required nodes and leaf/SPT subtrees for the current view. Second, each selected SPT is cut according to the camera’s distance to its root node. Together, these yield the full set of Gaussians needed for rendering the current frame. The construction and cutting process of an HSPT is illustrated in Figure 3. Figure 11 and 14 show example frames with respectively SPTs highlighted and different levels of detail.

Rebuilding the HSPT every training iteration would eliminate its performance benefits. Instead, we exploit the fact that the minimum distance m_d evolves slowly during optimization and thus requires only infrequent updates. In practice, we rebuild the HSPT only after each densification step. This infrequent recomputation allows us to use a more accurate—albeit more expensive—minimum distance metric than the inverse of maximal scale. Specifically, we define:

$$m'_d(i) = \frac{T}{\sqrt{s_i^1 \cdot s_i^2 + s_i^1 \cdot s_i^3 + s_i^2 \cdot s_i^3}}, \quad (6)$$

which corresponds to the inverse square root of the surface area of the Gaussian ellipsoid (up to a constant factor). This better captures the perceived size of anisotropic Gaussians, especially those that are significantly elongated in one or more directions.

Frustum Culling. The main benefits of using BFS to cut the upper hierarchy are the guarantee of a proper cut and early culling of subtrees. We therefore frustum cull every node considered in the BFS by checking if a sphere around the Gaussians with radius $3 \cdot \max_j s_i^j$ intersects the view frustum. While using the Gaussian scale as a conservative proxy for the entire subtree extent is not perfectly accurate, we observed no discernable difference compared to a full bounding sphere hierarchy in our experiments. Although Gaussians are implicitly frustum culled during rasterization, this early culling accelerates cut computation and significantly reduces the number of Gaussians loaded from RAM (see Figure 10).

4.4 Caching on the GPU

Loading Gaussian data from RAM is a costly operation that can become a significant bottleneck during large-scale training. To mitigate this, we maintain a GPU-resident cache of Gaussians that are likely to be reused across consecutive training views. However, checking the cache for every individual Gaussian would introduce non-trivial overhead. Once again, SPTs offer an efficient alternative.

Rather than caching individual Gaussians, we store the Gaussians from SPT cuts along with the cached distance from the camera to the root of each SPT, denoted \bar{d}^j for the j -th SPT. During rendering, when the upper hierarchy is cut and the required SPTs identified, we compute $d^j = \|\mu_{\text{root}(j)} - \mathbf{p}_{\text{cam}}\|_2$ and check whether a matching cut is cached, using a simple distance ratio tolerance:




Fig. 4. Gaussians required for the current training view are assembled from three sources: the upper tree, newly loaded SPT cuts from RAM, and cache hits. After optimization, newly accessed SPTs are added to the GPU cache.

$$D_{\min} \leq \frac{d^j}{\bar{d}^j} \leq D_{\max}. \quad (7)$$

Here, D_{\max} defines the allowable range for using coarser-than-optimal detail, while D_{\min} limits how much finer detail can be tolerated. If this condition is met, the cached SPT cut is reused, avoiding a costly RAM-to-GPU transfer.

While this heuristic introduces slight variability in rendered detail—since the LoD may depend on the cache state—we find that this stochasticity actually improves training robustness. In particular, subtle variations in detail across views discourage overfitting to fixed camera distances and promotes generalization across scales.

For each training view, visible Gaussians are assembled from the upper hierarchy, the cached SPTs, and the skybox Gaussians (which remains in VRAM). Uncached SPTs are streamed from RAM. After each training iteration, newly loaded SPTs are added to the cache.

To bound VRAM usage, we use a least-recently-used (LRU) write-back policy. When a memory threshold is exceeded, entries are written back to RAM. Additionally, to prevent overfitting to persistent cache entries, the entire cache is flushed every 1 000 iterations. Figure 4 illustrates the caching process across two frames.

View Selection. In large-scale scenes, the GPU cache typically covers only a small fraction of the overall geometry, leading to sparse cache hits. To improve cache utilization, we prioritize spatial locality by selecting successive training views close to the current one, maximizing Gaussian reuse.

To this end, we precompute a k -nearest-neighbour graph over all training view positions, where edge weights w_{ij} correspond to the Euclidean distance between views i and j . The next training view j is then sampled from the k -nearest neighbours of the current view i according to the distribution: $\mathbb{P}(j | i) \propto \frac{1}{w_{ij} + W}$, where W is a normalization constant that also controls the degree of exploration. However, care must be taken when deviating from a uniformly sampled training view selection as this may introduce bias. To counteract this, we inject a randomly selected view every 128 iterations, which we find sufficient to preserve generalization performance.

4.5 Memory Layout

Figure 5 illustrates the peak memory usage for a single training iteration of a 60-million-Gaussian hierarchy on the MC-smaller-city+ dataset. The majority of RAM usage is consumed by per-Gaussian properties and their corresponding ADAM optimizer states. In contrast, the hierarchy structure itself accounts for less than 10% of

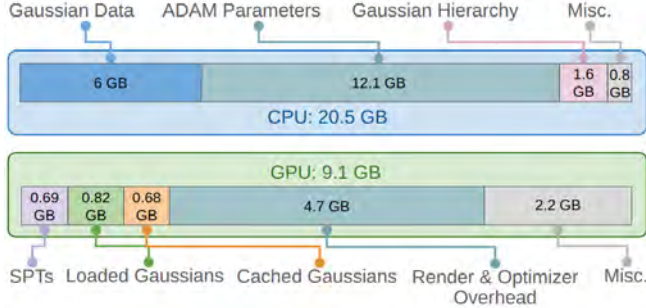


Fig. 5. Peak memory consumption of CPU and GPU for a training iteration on MC-smaller-city+ with 60 million Gaussians.

Table 1. **Novel view synthesis results.** Results with [†] require suboptimal COLMAP initialization instead of the provided point cloud. VRAM usage is measured while rendering the test images. Methods are considered out-of-memory (OOM) if they exceed 141GB of VRAM during training or rendering.

Method	PSNR [†]	SSIM [†]	LPIPS [†]	VRAM (render) [†]	VRAM (train) [†]	#iterations [†]	#Gaussians
MC-small-city+ (42.2k images)							
CityGaussian	19.78	0.650	0.475	2.4GB	4.1GB	1.11M	2.7M
HorizonGS	12.06	0.521	0.544	20.6GB	25.3GB	2.5M	N/A
OctreeGS	OOM	OOM	OOM	OOM	OOM	OOM	OOM
H-3DGS [†]	OOM	OOM	OOM	OOM	OOM	OOM	OOM
CLM	15.68	0.584	0.523	18.7GB	25.1GB	600k	24.2M
Ours	21.59	0.711	0.396	16.4GB	18.0GB	600k	136.7M
MC-smaller-city+ (15.1k images)							
HorizonGS	15.28	0.601	0.435	13.9GB	23.6GB	1.2M	34.1M
OctreeGS	15.93	0.560	0.556	26.6GB	98.7GB	310k	N/A
Ours	21.71	0.710	0.370	15.7GB	17.2GB	310k	60M
H-3DGS [†]	14.57	0.516	0.563	97.0GB	15.1GB	4M	83.1M
Ours [†]	20.59	0.658	0.468	20.7GB	22.1GB	310k	24.7M
Uni10k (10.1k images)							
CityGaussian	19.78	0.642	0.414	3.7GB	14.6GB	300k	2.2M
H-3DGS	19.96	0.705	0.319	39.2GB	26.5GB	255k	32.5M
OctreeGS	20.06	0.602	0.341	25.3GB	22.9GB	160k	N/A
CLM-GS	19.89	0.608	0.455	46.7GB	39.7GB	160k	12.6M
HorizonGS	21.72	0.722	0.326	10.3GB	44.3GB	400k	N/A
Ours	21.32	0.718	0.336	22.6GB	31.3GB	160k	28.6M

the total RAM footprint. On the GPU, the SPT metadata for all 60 million Gaussians occupies just 680 MB of VRAM and the upper hierarchy negligible 24 MB.

Even in wide-angle aerial views, only a subset of the scene is actively loaded into GPU memory. In the example shown, 2.2 million Gaussians are rendered directly, while an additional 2.4 million are retained in the cache for future use. The bulk of GPU memory is instead consumed by temporary allocations for rasterization and optimization, which scale with the number of Gaussians rendered. Therefore, minimizing the number of active Gaussians is critical for staying within the VRAM budget.

The remaining GPU memory usage consists of auxiliary data, including cache management, hierarchy cut tracking, training and ground-truth images, as well as general PyTorch overhead. Figure 9 provides a breakdown of all data associated with a single Gaussian and how they are distributed across CPU and GPU.

5 Evaluation

Our method is designed to enable seamless training and rendering on ultra-large-scale scenes comprising tens of thousands of views captured at vastly different scales. Unfortunately, most datasets of sufficient size contain either street-level or aerial views—but not both. To address this gap, we captured the campus of Udine University with over 10 000 images at 4k resolution from varying aerial heights and street views in the Uni10k scene. We also introduce MC-small-city+, which expands the small-city scene from the MatrixCity [Li et al. 2023] dataset with new high-altitude aerial views covering the entire scene, resulting in 42.2k images spanning hundreds of buildings. Since some methods run out-of-memory on scenes of this scale, we additionally construct a subset of MC-small-city+ covering about a third of the area (15.1k images, MC-smaller-city+). Additionally, we present results on large-scale street-view and indoor scenes from the Hierarchical 3DGS [Kerbl et al. 2024] and OccluGaussian [Liu et al. 2025] datasets. Results on the aerial datasets UrbanScene3D [Lin et al. 2022] and Mill19 [Turki et al. 2022], as well as the smaller-scale H-3DGS single chunk scene are included in Suppl. A.3 along with additional details on all scenes.

We choose recent divide-and-conquer based 3DGS methods *CityGaussian* [Liu et al. 2024a], *H-3DGS* [Kerbl et al. 2024], the large-scale neural Gaussian method *OctreeGS* [Ren et al. 2024] and out-of-core training method *CLM-GS* [Zhao et al. 2026] as baselines. Because *OccluGaussian* [Liu et al. 2025] has not released code yet, we compare only against their self-reported results on their dataset. To enable training on scenes of this scale, it was necessary to modify *CityGaussian*, *HorizonGS* and *OctreeGS* to load images from disk instead of caching them fully in RAM or VRAM. These changes affect training throughput but do not alter optimization behavior or final reconstruction quality. In general, the hyperparameters suggested for large-scale scenes were used for the experiments. Exact details on the training setting and configuration files are included in the supplemental material. Evaluations were run on a single H200 GPU with 141 GB VRAM to allow baselines methods to complete training. Where applicable, we report rendering performance on consumer GPUs to demonstrate practical deployability.

5.1 Results

Qualitative comparisons are shown in Figure 12 and quantitative comparisons in Tables 1 and 2. The outputs of *H-3DGS* show significant floating and ghosting artifacts. While *H-3DGS* performs well on individual chunks of MC-smaller-city+, oversized floaters survive the merging procedure and obscure the majority of test images (cf. Figure 7). Floaters and ghosting artifacts can also be found in the results on other scenes. The merging procedure of *CityGaussian*, which is designed for aerial-only datasets, discards most of the trained Gaussians to avoid chunk artifacts, resulting in mostly artifact-free views and low memory requirements, but noticeably blurry results. A *LoD of Gaussians* reconstructs novel views for joined street- and aerial data with a high degree of detail while avoiding chunk-based artifacts, resulting in a consistent improvement in quality metrics across scenes. Seamless training also enables faster convergence, substantially reducing the number of training iterations required compared to divide-and-conquer based methods.

Results on the OccluGaussian dataset further demonstrate generalizability to indoor environments, which are particularly difficult for divide-and-conquer based methods. All experiments of our method were run on an RTX 3090 GPU, except Uni10k, which required more VRAM for the 4k training images. Lowering the resolution to HD reduces peak VRAM to 20GB for the same parameters.

CLM-GS also employs out-of-core memory during training, but relies solely on frustum culling to limit memory pressure. On Uni10k and MC-small-city+, this results in substantially higher VRAM usage per Gaussian than our method. Further, unlike *OcTreeGS* [Ren et al. 2024] and our approach, their densification is not tailored to large-scale scene reconstruction from sparse views and point clouds. On MatrixCity, *CLM-GS* instead initializes from a fully dense synthetic point cloud; larger non-aerial scenes (e.g. Campus) fail to reconstruct under all configurations provided with their code.

HorizonGS [Jiang et al. 2025] is not competitive on the MC-city+ tests due to the large scale of the dataset and its diversity of view-points. While the authors report results for MatrixCity, only a single block is used in their evaluation. However, it performs competitively on Uni10k, for which it is explicitly designed: the method targets mixed street- and aerial-view scenarios and relies on ground-truth labels to primarily target street-level views for densification. Our universal approach—without additional supervision—closely matches *HorizonGS* on Uni10k and drastically outperforms it on both MC-city+ scenes. Furthermore, rendering the neural *HorizonGS* representation is impractical due to a costly preprocessing step: while enabling low VRAM usage, it results in render times of up to 10 seconds per frame.

Table 2. **H-3DGS and OccluGaussian dataset novel view synthesis results.** Results with \dagger are taken from [Liu et al. 2025], parentheses contain the number of images per-scene. We strongly suspect they evaluated LPIPS using the AlexNet backbone, so we report the same for our results as LPIPS_A.

Dataset	Hierarchical 3DGS						OccluGaussian					
	Small City (5.8k)			Campus (22.0k)			Canteen (8.6k)			Classbuilding (9.1k)		
Method	PSNR \dagger	SSIM \dagger	LPIPS \dagger	PSNR \dagger	SSIM \dagger	LPIPS \dagger	PSNR \dagger	SSIM \dagger	LPIPS \dagger	PSNR \dagger	SSIM \dagger	LPIPS \dagger
OcTreeGS	19.20	0.601	0.498	20.54	0.643	0.544	21.49	0.826	0.293	23.08	0.889	0.197
CityGS	21.40	0.692	0.383	19.47	0.645	0.531	20.41 \dagger	0.794 \dagger	0.275 \dagger	20.48 \dagger	0.840 \dagger	0.244 \dagger
H-3DGS	24.35	0.788	0.273	17.84	0.610	0.467	22.71 \dagger	0.825 \dagger	0.178 \dagger	23.87 \dagger	0.881 \dagger	0.128 \dagger
CLM-GS	22.54	0.699	0.392	16.60	0.613	0.558	23.13	0.864	0.168	22.89	0.892	0.159
OccluGS	N/A	N/A	N/A	N/A	N/A	N/A	25.25 \dagger	0.900 \dagger	0.100 \dagger	25.33 \dagger	0.921 \dagger	0.083 \dagger
Ours	24.61	0.774	0.297	21.85	0.707	0.456	23.33	0.867	0.155	24.14	0.915	0.116

Rendering. Our level-of-detail and caching strategy can also be applied to efficiently render the trained models. As supplemental material, we include fly-through videos of the evaluated scenes rendered on an RTX 3090. Table 1 demonstrates effective VRAM reduction during rendering of test images compared to other 3DGS methods. Figure 13 compares rendering VRAM consumption and image quality of our LoD method with full-detail 3DGS [Kerbl et al. 2023] and gsplat [Ye et al. 2025]. Our approach achieves visual fidelity comparable to the baselines while significantly reducing VRAM usage, highlighting the effectiveness of our LoD scheme.

Ablations. We assess the contribution of key components through an ablation study using recorded camera paths across a selection of scenes (see supplemental videos). Table 3 reports average frame

times over these paths. Caching significantly improves rendering performance, roughly doubling the framerate across scenes by reducing the average number of Gaussians loaded from RAM by 93% on Campus and 86% on MC-small-city+. The effectiveness of frustum culling scales with scene size: On the full MC-small-city+ scene, 24.5 million Gaussians are frustum culled on average (88% reduction), while for Campus and Small City the corresponding values are 9.8 million and 7.9 million (74% and 65% reduction), respectively. As expected, the overhead of frustum culling amortizes with scene size, making it essential for the largest datasets.

To evaluate cut efficiency, we compare the time required to compute the visible set using either full hierarchy BFS or our HSPT-based approach. HSPT consistently yields faster cut times due to improved parallelization. Moreover, the BFS approach requires positions and scales for all Gaussians to reside in memory, causing it to exceed 24GB of VRAM on MC-small-city+, whereas the HSPT method peaks at 21GB. For training ablations, we measure average iteration durations over 1 000 steps. Here, frustum culling and caching prove essential, substantially reducing the number of Gaussians loaded and rendered per view.

Table 3. **Ablations.** Average frame times for rendering camera paths and average iteration times during training with and without caching Gaussians. The final results show the average timings of the hierarchy cut during rendering using our HSPT and the baseline BFS approach. For Campus, we evaluate two different models with 38M and 80M Gaussians respectively.

	MC-smaller-city+ 60M	MC-small-city+ 136M	Campus 38M	Campus 80M
Render Full	48.1 ms	68.3 ms	47.1 ms	83.2 ms
Render w/o Cache	119.4 ms	127.4 ms	92.6 ms	222.3 ms
Render w/o Frustum Culling	52.3 ms	89.3 ms	38.3 ms	110.2 ms
Training Full	156 ms	174 ms	205 ms	
Training w/o Cache	471 ms	257 ms	244 ms	
Training w/o Frustum Culling	685 ms	626 ms	312 ms	
Render Cut (HSPT)	31.9 ms	45.7ms	31.3 ms	36.5 ms
Render Cut (BFS)	47.8 ms	OOM	40.0 ms	53.7 ms

6 Discussion and Outlook

A LoD of Gaussians enables seamless training and rendering of ultra-large 3DGS models on consumer hardware. By storing Gaussian data in external memory and streaming it on demand, our method avoids the pitfalls of chunk-based pipelines. The HSPT datastructure accelerates LoD selection and remains robust to ongoing training changes. Combined with caching and view selection, our approach significantly reduces out-of-core overhead. These components enable efficient reconstruction and rendering at scale, as demonstrated on challenging multi-scale scenes such as MC-small-city+.

Limitations and Future Work. Our method represents an informed trade-off between performance and memory. While it greatly reduces the necessary number of training iterations for large-scale scenes, individual iterations take longer than standard 3DGS training due to data loading and hierarchy management overhead.

Similarly, although rendering framerates are a significant improvement over neural-based and other out-of-core methods, they do not yet reach the peak performance of fully in-core 3DGS.

Our approach requires roughly 1 GB of RAM per million Gaussians, which—while more efficient than prior methods—still constrains scalability for extremely large scenes. Loading from disk is feasible in our experiments, but at the cost of about a 10× slowdown, making fast secondary storage highly desirable.

The level-of-detail system makes our method robust to large variations in view distance. However, when such variation is absent—for example in single-height aerial datasets as evaluated in Suppl. A.3—the LoD machinery introduces unnecessary overhead, making more straightforward training a competitive alternative. Interactive rendering performance could be further improved by avoiding per-frame hierarchy cut recomputation and asynchronous streaming.

While frustum culling effectively reduces memory load in most views, it is ineffective when the entire scene lies inside the frustum. Occlusion culling could address this limitation by skipping entire SPTs before they are loaded into memory.

Overall, we believe that out-of-core 3D Gaussian Splatting is a promising direction for scaling radiance field methods to city-scale scenes and beyond on consumer-grade hardware.

References

- Sameer Agarwal, Noah Snavely, Ian Simon, Steven M. Seitz, and Richard Szeliski. 2009. Building Rome in a day. In *2009 IEEE 12th International Conference on Computer Vision*. 72–79. doi:10.1109/ICCV.2009.5459148
- Junyi Chen, Weicai Ye, Yifan Wang, Dapeng Chen, Di Huang, Wanli Ouyang, Guofeng Zhang, Yu Qiao, and Tong He. 2025. GigaGS: 3D Gaussian Based Planar Representation for Large-Scene Surface Reconstruction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 2088–2096.
- Carsten Dachsbacher, Christian Vogelgsang, and Marc Stamminger. 2003. Sequential point trees. *ACM Trans. Graph.* 22, 3 (July 2003), 657–662. doi:10.1145/882262.882321
- Zhiwen Fan, Kevin Wang, Kairun Wen, Zehao Zhu, Dejie Xu, Zhangyang Wang, et al. 2024. Lightgaussian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps. *Advances in neural information processing systems* 37 (2024), 140138–140158.
- Guangchi Fang and Bing Wang. 2024. Mini-splatting: Representing scenes with a constrained number of gaussians. In *European Conference on Computer Vision*. Springer, 165–181.
- He Huang, Wenjie Huang, Qi Yang, Yiling Xu, and Zhu Li. 2025. A hierarchical compression technique for 3d gaussian splatting compression. In *ICASSP 2025–2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1–5.
- Lihan Jiang, Kerui Ren, Mulin Yu, Linning Xu, Junting Dong, Tao Lu, Feng Zhao, Dahua Lin, and Bo Dai. 2025. Horizon-GS: Unified 3D Gaussian Splatting for Large-Scale Aerial-to-Ground Scenes. In *Proceedings of the Computer Vision and Pattern Recognition Conference*. 26789–26799.
- Brian Karis, Rune Stubbe, and Graham Wihlidal. 2021. A Deep Dive into Nanite Virtualized Geometry, Advances in Real-Time Rendering in Games: Part I. In *ACM SIGGRAPH 2021 Courses (SIGGRAPH '21)*. Association for Computing Machinery. https://advances.realtimereading.com/s2021/Karis_Nanite_SIGGRAPH_Advances_2021_final.pdf
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)* 42, 4 (July 2023). <http://www.sop.inria.fr/revues/Basilic/2023/KKLD23>
- Bernhard Kerbl, Andreas Meuleman, Georgios Kopanas, Michael Wimmer, Alexandre Lanvin, and George Drettakis. 2024. A Hierarchical 3D Gaussian Representation for Real-Time Rendering of Very Large Datasets. *ACM Trans. Graph.* 43, 4, Article 62 (July 2024), 15 pages. doi:10.1145/3658160
- Shakiba Kheradmand, Daniel Rebaïn, Gopal Sharma, Weiwei Sun, Yang-Che Tseng, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 2024. 3D Gaussian Splatting as Markov Chain Monte Carlo. In *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Eds.), Vol. 37. Curran Associates, Inc., 80965–80986. https://proceedings.neurips.cc/paper_files/paper/2024/file/93be245fce00a9bb2333c17ceae4b732-Paper-Conference.pdf
- Donghyun Lee, Dawoon Jeong, Jae W. Lee, and Hongil Yoon. 2025. GS-Scale: Unlocking Large-Scale 3D Gaussian Splatting Training via Host Offloading. arXiv:2509.15645 [cs.CV] <https://arxiv.org/abs/2509.15645>
- Bingling Li, Shengyi Chen, Luchao Wang, Kaimin Liao, Sijie Yan, and Yuanjun Xiong. 2024a. RetinaGS: Scalable Training for Dense Scene Rendering with Billion-Scale 3D Gaussians. arXiv:2406.11836 [cs.CV] <https://arxiv.org/abs/2406.11836>
- Ruilong Li, Sanja Fidler, Angjoo Kanazawa, and Francis Williams. 2024b. Nerf-xl: Scaling nerfs with multiple gpus. In *European Conference on Computer Vision*. Springer, 92–107.
- Yixuan Li, Lihan Jiang, Linning Xu, Yuanbo Xiangli, Zhenzhi Wang, Dahua Lin, and Bo Dai. 2023. Matrixcity: A large-scale city dataset for city-scale neural rendering and beyond. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 3205–3215.
- Zhuoxiao Li, Shanliang Yao, Yong Yue, Wufan Zhao, Rongjun Qin, Angel F. Garcia-Fernandez, Andrew Levers, and Xiaohui Zhu. 2025. ULSR-GS: Ultra Large-scale Surface Reconstruction Gaussian Splatting with Multi-View Geometric Consistency. arXiv:2412.01402 [cs.CV] <https://arxiv.org/abs/2412.01402>
- Jiaqi Lin, Zhihao Li, Xiao Tang, Jianzhuang Liu, Shiyong Liu, Jiayue Liu, Yangdi Lu, Xiaofei Wu, Songcen Xu, Youliang Yan, et al. 2024. Vastgaussian: Vast 3d gaussians for large scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5166–5175.
- Liqiang Lin, Yilin Liu, Yue Hu, Xingguang Yan, Ke Xie, and Hui Huang. 2022. Capturing, Reconstructing, and Simulating: the UrbanScene3D Dataset. In *European Conference on Computer Vision*. 93–109.
- Shiyong Liu, Xiao Tang, Zhihao Li, Yingfan He, Chongjie Ye, Jianzhuang Liu, Binxiao Huang, Shunbo Zhou, and Xiaofei Wu. 2025. OccluGaussian: Occlusion-Aware Gaussian Splatting for Large Scene Reconstruction and Rendering. (2025).
- Yang Liu, Chuanchen Luo, Lue Fan, Naiyan Wang, Junran Peng, and Zhaoxiang Zhang. 2024a. Citygaussian: Real-time high-quality large-scale scene rendering with gaussians. In *European Conference on Computer Vision*. Springer, 265–282.
- Yang Liu, Chuanchen Luo, Zhongkai Mao, Junran Peng, and Zhaoxiang Zhang. 2024b. Citygaussianv2: Efficient and geometrically accurate reconstruction for large-scale scenes. *arXiv preprint arXiv:2411.00771* (2024).
- Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. 2024. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 20654–20664.
- Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65, 1 (2021), 99–106.
- Simon Niedermayr, Josef Stumpfegger, and Rüdiger Westermann. 2024. Compressed 3d gaussian splatting for accelerated novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10349–10358.
- Michael Niemeyer, Fabian Manhardt, Marie-Julie Rakotosaona, Michael Oechsle, Daniel Duckworth, Rama Gosula, Keisuke Tateno, John Bates, Dominik Kaeser, and Federico Tombari. 2025. Radsplat: Radiance field-informed gaussian splatting for robust real-time rendering with 900+ fps. In *2025 International Conference on 3D Vision (3DV)*. IEEE, 134–144.
- Panagiotis Papanotakis, Georgios Kopanas, Bernhard Kerbl, Alexandre Lanvin, and George Drettakis. 2024. Reducing the Memory Footprint of 3D Gaussian Splatting. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 7, 1 (May 2024), 1–17. doi:10.1145/3651282
- Kerui Ren, Lihan Jiang, Tao Lu, Mulin Yu, Linning Xu, Zhangkai Ni, and Bo Dai. 2024. Octree-gs: Towards consistent real-time rendering with lod-structured 3d gaussians. *arXiv preprint arXiv:2403.17898* (2024).
- Johannes L. Schönberger and Jan-Michael Frahm. 2016. Structure-from-Motion Revisited. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 4104–4113. doi:10.1109/CVPR.2016.445
- Yunji Seo, Young Sun Choi, Hyun Seung Son, and Youngjung Uh. 2024. Flod: Integrating flexible level of detail into 3d gaussian splatting for customizable rendering. *arXiv preprint arXiv:2408.12894* (2024).
- Matthew Tancik, Vincent Casser, Xinchen Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. 2022. Block-nerf: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 8248–8258.
- Haitthem Turki, Deva Ramanan, and Mahadev Satyanarayanan. 2022. 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*. 12922–12931.
- Linning Xu, Yuanbo Xiangli, Sida Peng, Xingang Pan, Nanxuan Zhao, Christian Theobalt, Bo Dai, and Dahua Lin. 2023. Grid-guided Neural Radiance Fields for Large Urban Scenes. arXiv:2303.14001 [cs.CV] <https://arxiv.org/abs/2303.14001>
- Xijie Yang, Linning Xu, Lihan Jiang, Dahua Lin, and Bo Dai. 2025. Virtualized 3D Gaussians: Flexible Cluster-based Level-of-Detail System for Real-Time Rendering of Composed Scenes. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Papers (SIGGRAPH Conference Papers '25)*. Association for Computing Machinery, New York, NY, USA, Article 100, 11 pages. doi:10.1145/3721238.3730602
- Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari, Jianbo Ye, Jeffrey Hu, Matthew Tancik, et al. 2025. gsplat: An open-source library for Gaussian splatting. *Journal of Machine Learning Research* 26, 34 (2025), 1–17.
- Lingzhi Zhang, Tarmily Wen, and Jianbo Shi. 2020. Deep image blending. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. 231–240.
- Saining Zhang, Baijun Ye, Xiaoxue Chen, Yuantao Chen, Zongzheng Zhang, Cheng Peng, Yongliang Shi, and Hao Zhao. 2024. Drone-assisted Road Gaussian Splatting with Cross-view Uncertainty. arXiv:2408.15242 [cs.CV] <https://arxiv.org/abs/2408.15242>
- Hexu Zhao, Xiwen Min, Xiaoteng Liu, Moonjun Gong, Yiming Li, Ang Li, Saining Xie, Jinyang Li, and Aurojit Panda. 2026. CLM: Removing the GPU Memory Barrier for 3D Gaussian Splatting. In *Proceedings of the 2026 International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'26)*. Pittsburgh, PA, USA.
- Hexu Zhao, Haoyang Weng, Daohan Lu, Ang Li, Jinyang Li, Aurojit Panda, and Saining Xie. 2024. On scaling up 3d gaussian splatting training. In *European Conference on Computer Vision*. Springer, 14–36.
- Mathias Zwicker, Hanspeter Pfister, Jeroen van Baar, and Markus Gross. 2001. EWA Volume Splatting. In *IEEE Visualization*.

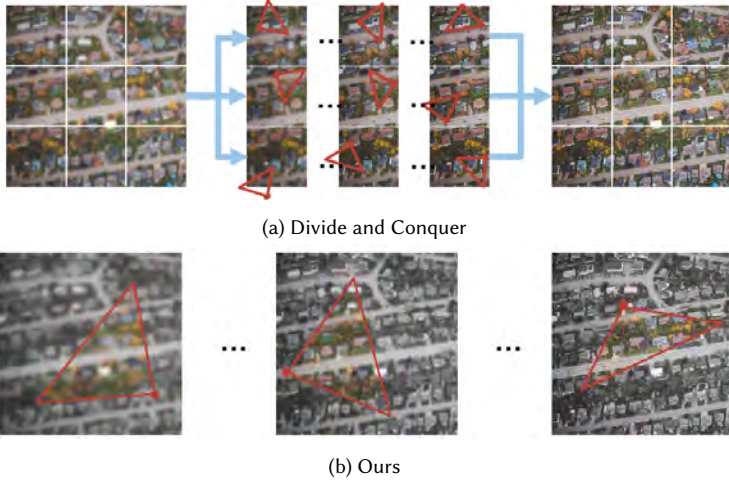


Fig. 6. Comparison between the divide and conquer training process as used in Kerbl et al. [2024]; Lin et al. [2024]; Liu et al. [2024a] to our training process. Colored regions are present in VRAM, training views are drawn in red.



(a) Chunk Bleeding in *H-3DGS* [Kerbl et al. 2024].



(b) Chunk Ghosting in *CityGaussian* [Liu et al. 2024a].

Fig. 7. Artifacts caused by the divide-and-conquer strategy on MC-small-city+.

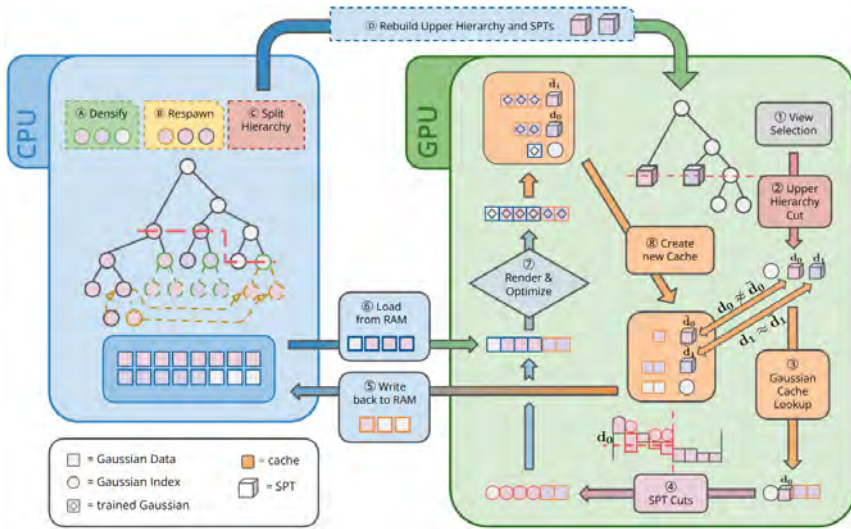


Fig. 8. Method Overview: Steps ① to ⑧ show the process of a single training iteration, while ⑨ and ⑩ show a densification step.

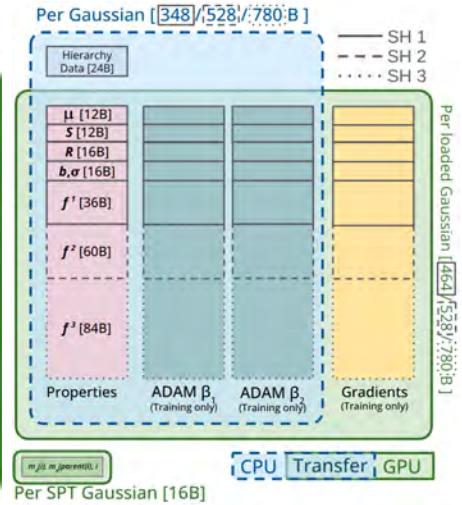


Fig. 9. Memory Layout: Only the currently loaded Gaussians and slim SPT information needs to be stored on GPU.




Fig. 10. Frustum Culling and LoD selection (left) greatly reduces the number of Gaussians required to render a view (right).



Fig. 11. Hierarchical SPTs enable smooth transitions between detailed and coarse representation.



Fig. 12. Qualitative comparison of our method and SOTA methods [Kerbl et al. 2024; Liu et al. 2024a; Ren et al. 2024; Zhao et al. 2026] on the MC-small-city+, Campus and Uni10k scenes.

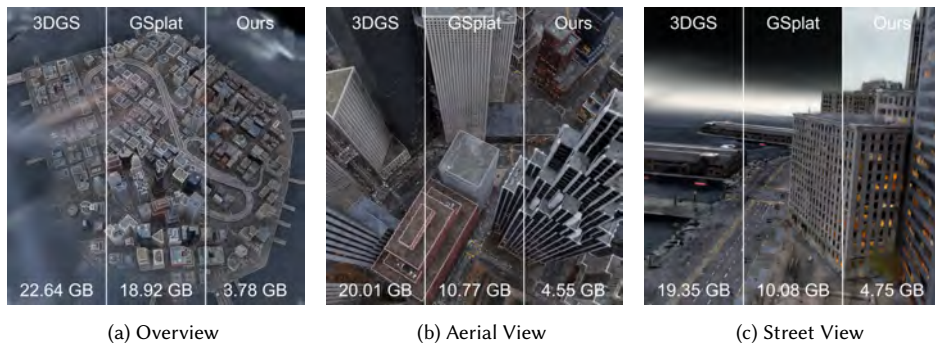


Fig. 13. Comparison to standard 3DGS [Kerbl et al. 2023] and gsplat [Ye et al. 2025] in terms of rendering quality and VRAM usage on the MC-small-city+ scene.




Fig. 14. SPTs for a frame of MatrixCity rendered in different colors.

A Supplementary Material

Our full supplementary material consists of:

- (1) This document, which includes additional ablation studies, implementation details and experiments, which we could not include in the main paper due to page restrictions
- (2) A short **video presentation** about this manuscript
- (3) An html-page with more qualitative comparisons and short videos
- (4) Our full source code contained in the Code directory
- (5) Training configurations detailing all hyperparameters for our and baseline methods on each experiment

A.1 Additional Ablations

We present additional experiments to demonstrate the effectiveness of each component of our method.

View Selection. To verify the effectiveness of caching and view selection, we conduct an experiment by training the MC-smaller-city+ dataset (100k iterations, up to 30M Gaussians) with and without view selection and caching. The results in Table 4 show that while quality metrics are barely affected (with a maximal difference of 0.006 PSNR), the number of Gaussians that need to be loaded into VRAM decrease by a factor of almost 20× with view selection and caching enabled.

Effectiveness of LoD. To evaluate the effectiveness of our level-of-detail system, we choose 5 random images from each dataset and compare the resulting quality metrics to the number of Gaussians rendered at 50 different levels of details. The results are visualized in Figure 16 in addition to average FPS for rendering the view 10 times with pre-warmed cache on an H200 GPU. As expected, FPS scales closely with the inverse of the number of Gaussians. In the vast majority of cases, increasing the LoD level consistently improves rendering quality. Overall, we only notice a small dip in quality when reducing the Gaussian count by 50%, verifying the effectiveness of our LoD structure. Further, using continuous instead of discrete levels of detail leads to the smoothness of the curves in Figure 16, indicating a smooth transition between LoD levels, which reduces popping.

Effectiveness of Caching. In Figure 15, we provide sensitivity curves for cache size vs. FPS during rendering. In each iteration up to *cache size* Gaussians from the previous frames may be reused. The results once again demonstrate the importance of the caching system to rendering efficiency. The graphs show that performance increases steeply with increasing cache size until it exceeds the typical number of Gaussians required per iteration and starts to plateau.

A.2 Initialization and Training Details

Following Kerbl et al. [2024], we initialize the Gaussian model from a sparse point cloud, augmented with skybox points. This initial representation is small enough to fully reside in GPU memory and is trained for 100k iterations without densification. The goal of this phase is to establish a stable global scene structure before constructing the hierarchy. After this initial optimization, we build a binary Gaussian hierarchy, where the trained Gaussians act as leaf nodes

Table 4. Ablation results on the MC-smaller-city+ dataset. #Loaded refers to the average number of Gaussians loaded from RAM over 100 000 iterations.

	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	#Loaded
w/o caching	21.831	0.701	0.425	1 629 187
w/ caching	21.829	0.700	0.429	290 572
w/ view selection	21.825	0.697	0.432	89 451




Fig. 15. Average FPS for rendering 1000 frames of the camera paths for the MC-small-city+ and Campus scenes with various cache sizes.

and the parent nodes represent merged approximations of their two children.

The Gaussian properties—including base color, SH coefficients, position, and covariance—are optimized using the standard loss propagation introduced by Kerbl et al. [2023]. To fulfill our tight memory requirements, we update only the parameters of the cut set chosen for each training view. Thus, gradients may, in general, be propagated to Gaussians in the middle of the hierarchy. Assuming good coverage and diversity of training views, these updates will diffuse over the entire hierarchy over the course of training, leading to smooth transitions from the highest to the lowest LoD.

A.3 Dataset Details and Further Results

We include more details on the datasets used in the evaluation and provide additional results for smaller-scale scenes.

MC-small-city+. For our main benchmark, MC-small-city+, we aggregate 33 006 street-view images and 7 672 aerial views from the small-city scene of the MatrixCity dataset [Li et al. 2023] and generate 533 additional high-altitude views. We evaluate reconstruction quality on a separate set of 4 228 test views, according to the test split provided by the dataset. The scene is extremely challenging due to its enormous scale, sparse views and wide variation in scale: As such, we also construct a subset of the entire dataset (15.1k images, covering about a third of the area), denoted as MC-smaller-city+ for baselines that are not able to reconstruct the full dataset. For MC-small-city+, we use the camera poses provided by the MatrixCity dataset and convert them to the COLMAP format for Gaussian splatting. We merge the provided street and aerial sparse point clouds and randomly downsample them by a factor of 5 to get more realistic initialization conditions. We do not make use of the ground truth depth images provided by the dataset, as

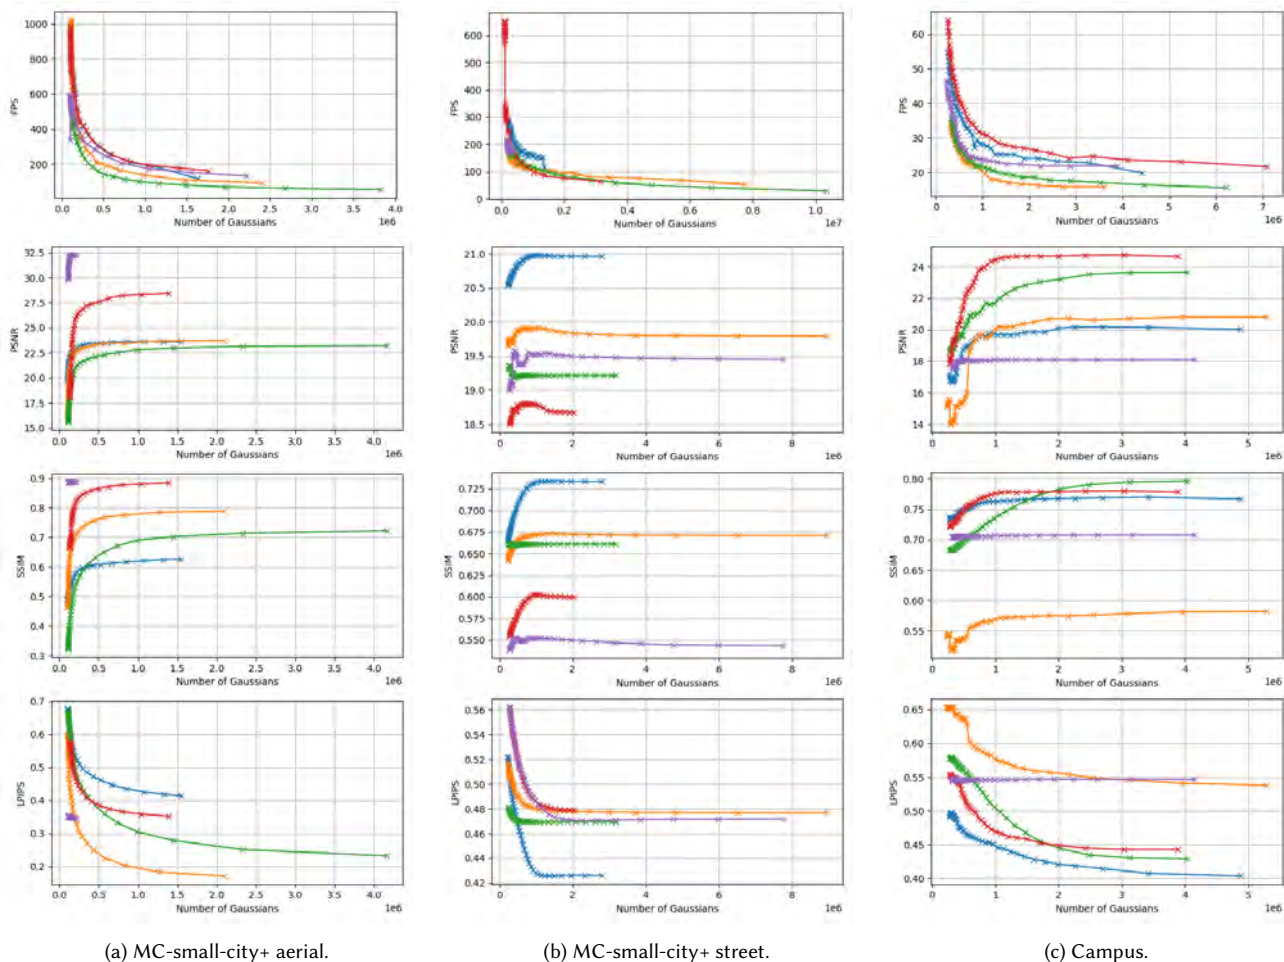


Fig. 16. Qualitative evaluation of our LoD system for 5 randomly chosen views and 50 different levels of detail.

we consider this to be an unrealistic advantage. Figure 17 shows the scene along with the camera distribution. While some methods have successfully reconstructed only aerial [Liu et al. 2024a,b; Ren et al. 2024] or only street-level views [Li et al. 2024b; Liu et al. 2024b], training a model that holds up to scrutiny from both perspectives presents a particular challenge. Training on close and far views simultaneously—without a proper LoD system like ours in place—significantly degrades visual quality for both, as noted by Jiang et al. [2025]; Zhang et al. [2024]. Moreover, such a scenario greatly complicates partitioning the scene into independent chunks.

Why are the quality metrics on MC-small-city+ so low? When comparing our evaluation on the MatrixCity dataset with those in other works like Jiang et al. [2025]; Liu et al. [2024a]; Ren et al. [2024]; Zhao et al. [2026], it might seem surprising that they achieve significantly lower quality metrics on our constructed MC-small-city+ scene. There are several reasons for this:

- (1) These works only evaluate on the aerial views, which make up less than 20% of the scene images and are generally easier to reconstruct.
- (2) Reconstructing street, aerial and the newly generated high-aerial views in the same scene is much more difficult than reconstructing just aerial views.
- (3) Some competing works [Liu et al. 2024a; Zhao et al. 2026] use the ground-truth depth maps or dense point clouds from the MatrixCity dataset. We would consider this to be an unrealistic advantage, since they contain almost the entire ground truth geometry of the scene. Zhao et al. [2026] even starts the reconstruction from the full dense point cloud and never densifies, whereas our evaluation starts from the same point cloud downsampled by a factor of 5.

Uni10K. In order to address the gap of real-world large-scale datasets that are captured across multiple scales, we present Uni10k. The scene consists of an outdoor campus of approximately 100 000 m², captured from both ground-level and aerial perspectives. Standard




Fig. 17. The MC-small-city+ dataset spans hundreds of buildings, which are supervised by tens of thousands of views, varying widely in scale.

reconstruction via COLMAP at this scale would typically require weeks of computation due to the complexity of the image-matching and mapping stages. To mitigate this, we leverage spatial and temporal priors in conjunction with a coarse-to-fine scheme. Specifically, we utilize GPS data to reduce the matching complexity from quadratic to near-linear by limiting image comparisons to a predefined spatial radius. Regarding temporal priors, since frames are sampled from video sequences, we initially reconstruct a baseline model starting from a set of frames uniformly sampled every second of video. Then, we densify the camera coverage by incrementally registering, triangulating, and refining new images using only local BA. The process concludes with several rounds of Global bundle adjustment. The final reconstruction comprises more than 10 000 images at 4K resolution and 6.2M sparse points, with an overall MRE of approximately 0.63 pixels. We train on the full resolution images and hold back every 8th frame alphabetically for the test set. Figure 18 shows the sparse point cloud and camera distribution.

Hierarchical 3DGS dataset. Since H-3DGS Kerbl et al. [2024] did not introduce a test split for their dataset, we hold back every 100th image alphabetically for testing. We reevaluate H-3DGS for this scene with the new test split, using the camera calibrations and chunk splits provided on their website. In accordance with the instructions, we disable exposure optimization for evaluation. On this dataset only, we use depth supervision identical to Kerbl et al. [2024]. The distribution of street views in the Campus scene is visualized in Figure 19b.

Hierarchical 3DGS single-chunk dataset. To demonstrate our ability to reconstruct small scenes, we also evaluate our method on the smaller, single-chunk versions of the Campus and Small City scenes (cf. Table 5). The quality metrics indicate that training on the LoD structure (which would not be required for scenes of this scales)



Fig. 18. Uni10K’s sparse point cloud. The scene is reconstructed from more than 10 thousand 4k images at widely varying heights, both from aerial and ground perspective.

Table 5. H-3DGS single chunk view synthesis results. Results with \dagger are taken from [Kerbl et al. 2024].

Method	Small City Chunk (1.1k images)			Campus Chunk (1.5k images)		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Mip-NeRF 360 \dagger	24.70	0.765	0.348	20.95	0.731	0.442
INGP-big \dagger	23.47	0.715	0.426	20.37	0.700	0.476
F2-NeRF-big \dagger	24.53	0.762	0.342	19.25	0.681	0.478
3DGS \dagger	25.34	0.776	0.337	23.87	0.785	0.378
H-3DGS \dagger	26.62	0.820	0.259	24.61	0.807	0.331
Ours	25.94	0.808	0.276	24.86	0.796	0.360

during training only leads to marginal reductions in visual quality.

OccluGaussian Dataset. The OccluGaussian dataset features three large-scale indoor environments. Unfortunately, the download link for the Gallery scene on the official website leads to a broken zip

file. As we have not received an answer from the authors, we were unable to evaluate on this scene. Additionally, some training images that are referenced in the camera pose files were not included in the download of the Canteen scene. We ran the evaluation on all images that were provided using full resolution images.

Mill19 and Urbanscene3D dataset. Table 6 and 7 show results on the Mill19 [Turki et al. 2022] and UrbanScene3D [Lin et al. 2022] datasets. We use the camera poses and sparse point cloud provided by Liu et al. [2024a]. In accordance with baselines, we downscale all images by a factor of 4. We have included these dataset, because they are widely used in large-scale novel-view synthesis. At the same time, they represents a worst-case scenario for our method, as the regular, same-height aerial views (cf. Figure 19a) negate any benefit of our level-of-detail method and can be split into independent chunks trivially.

Table 6. **Mill19 novel view synthesis results** Results with \dagger are taken from Liu et al. [2024a]. Results of *H-3DGS* are taken from Kerbl et al. [2024].

Method	Rubble (1.6k images)				Building (1.9k images)			
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Size	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Size
MegaNeRF \dagger	24.06	0.553	0.516	n/a	20.93	0.547	0.504	n/a
3DGS \dagger	25.47	0.777	0.277	6.1M	20.46	0.720	0.305	6.4M
CityGaussian \dagger	25.77	0.813	0.228	9.7M	21.55	0.778	0.246	13.2M
H-3DGS	24.64	0.755	0.284	n/a	21.52	0.723	0.297	n/a
Ours	23.75	0.704	0.330	25M	20.67	0.682	0.318	25M

Table 7. **UrbanScene3D novel view synthesis Results.** Results with \dagger are taken from Liu et al. [2024a].

Method	Residence (2.5k images)			Sci-Art (3k images)		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
CityGaussian \dagger	21.90	0.805	0.217	21.34	0.833	0.232
3DGS \dagger	21.44	0.791	0.236	21.05	0.830	0.242
H-3DGS	17.52	0.416	0.391	18.42	0.652	0.340
OctreeGS	18.76	0.523	0.539	18.97	0.620	0.535
Ours	20.25	0.705	0.319	20.70	0.794	0.272

A.4 Evaluation Details

Finding baselines for the MC-small-city+ scene proved challenging, as the only methods we are aware of that have successfully trained on all fused aerial and street-level views—Li et al. [2024a] and Li et al. [2024b]—both require 64 GPUs running in parallel. We were not able to meet these hardware requirements for evaluation and hold that these methods and similar multi-GPU works like Zhao et al. [2024] are orthogonal to our method. H-3DGS [Kerbl et al. 2024] has demonstrated the ability to reconstruct similar sized datasets, but its evaluation was restricted to street-level views. We choose Kerbl et al. [2024] and other popular large-scale Gaussian Splatting methods [Jiang et al. 2025; Liu et al. 2024a; Ren et al. 2024; Zhao et al. 2026] as baselines. Table 8 compares the supported features of our method and all baselines.

Most of the large-scale 3DGS frameworks have complicated multi-stage training processes and are very sensitive to hyperparameters,

which we detail in the following section. We consider it a benefit of our training pipeline that the user is not confronted with tuning the error-prone partitioning process, being more similar to the original Gaussian Splatting training.

OctreeGS. For *OctreeGS* [Ren et al. 2024], we follow the provided instructions on training custom datasets. We trained using the suggested hyperparameters for standard scenes and those used for the MatrixCity dataset. The latter achieved higher quality metrics, which is what we report. Note that OctreeGS does not perform any scene division and uses the memory-intensive ScaffoldGS [Lu et al. 2024] method, making it unsuitable for ultra-large scale scenes.

CityGaussian. We choose *CityGaussian* [Liu et al. 2024a] (specifically the 1.2 version of the repository) as a baseline instead of *CityGaussianV2* [Liu et al. 2024b], as the code release for *CityGaussianV2* does not yet support level-of-detail rendering. We follow the instructions on training large datasets and use the parameters and chunk split from the provided configuration file for the MatrixCity-Aerial dataset (which covers the same region as the MC-small-city+ dataset, but without the street views), as they produced better metrics than the default parameters. Note that the low number of Gaussians in the final model is due to *CityGS* discarding most (about 90%) of the trained Gaussians after chunk training in order to avoid artifacts. For Campus and Small City, we used the same parameters, but a 4×4 and 2×2 chunk split to match the number of chunks as closely as possible with the 12 and 4 chunks respectively used by *H-3DGS* [Kerbl et al. 2024] and a 3×3 chunk split for the Uni10k dataset.

CLM-GS. Significant modifications were necessary to evaluate CLM-GS [Zhao et al. 2026] on our large-scale datasets. In the paper, the issue of large-scale densification is sidestepped, by starting from a dense (102 million point) ground truth point cloud and disabling densification all-together for the aerial matrix-city scene. All other evaluated scenes are sufficiently small that densification is not an issue. Running the unmodified method on large-scale datasets from sparse point clouds leads to either collapsing training or pruning of all Gaussians. Thus, we disable pruning to avoid large parts of the scene disappearing and disable opacity resets, which make the training process unstable for large-scale scenes. Experimentally we found that the high densification parameters for the 28M mode of Rubble produced the best results and used it in our evaluation. We further implemented the suggestion of Kerbl et al. [2024] to densify according to maximal instead of average screen-space gradients, but did not notice significant improvements.

As suggested by the authors for large datasets, we use the clm-offload strategy with a batch size of 4, and match the number of training iterations with our method for each experiment. We enable the option to load images from disk in order to avoid going out of memory. As of writing the CLM codebase contains no possibility of rendering test images or evaluating performance metrics, so we implemented the necessary functions ourselves.

Hierarchical-3DGS. We follow the instructions of Kerbl et al. [2024] for running the method on large scenes. Note that we do not make use of their particular COLMAP pipeline for MC-smaller-city+, as it only works with unbroken camera paths. We disable exposure




Fig. 19. Overview of the aerial and street datasets.

	LoD Training	LoD Rendering	Out-of-Core Training	Interactive Rendering	No Chunking Artifacts	Gaussians
CityGaussian	✗	✓	✗	✓	✗	Explicit
HorizonGS	✓	✓	✗	✗	✗	Neural
H-3DGS	✗	✓	✗	✓	✗	Explicit
OctreeGS	✓	✓	✗	✓	✓	Neural
CLM-GS	✗	✗	✓	✗	✓	Explicit
VastGaussian	✗	✗	✗	✓	✗	Explicit
OccluGaussian	✗	✗	✗	✓	✗	Explicit
Ours	✓	✓	✓	✓	✓	Explicit

Table 8. Comparison of features.

optimization for evaluation in accordance with the instructions. The partitioning step of H-3DGS requires point correspondences not provided by the dataset. To substitute, we use the camera poses provided by the dataset and generate the sparse point cloud and correspondences from scratch using COLMAP. On scenes of this scale, COLMAP output contains significant noise, which also leads to reduced quality metrics for our method.

H-3DGS achieves drastically worse PSNR and SSIM results on the full Campus scene, compared to the single-chunk results, due to a slight perspective distortion that occurs specifically on this dataset (it did not occur on Small City). This can be reproduced by evaluating the trained model on their Campus scene, both of which are available on their website. As such, we include the results, but encourage visual comparison in Figure 12 of the main paper.

HorizonGS. We expand the configuration provided for a single chunk of the MatrixCity small-scene dataset (named *ours/large_scene/block_A*) with a chunk split of 3×4 and 5×5 for MC-smaller-city+ and MC-small-city+, respectively, and 2×2 for Uni10k. Since the method throws an error if it is not given both a set of street- and aerial images, we limit evaluation to those two scenes. Each chunk is trained for 60000 iterations and then fine-tuned for another 40000.

The chunk partitions are shown in Figure 20; as we can see, the complexity of these partitions is vastly different. For Uni10k, the

simpler chunk partitioning leads to more stable training, and, as can be seen in Table 1, improved results.




Fig. 20. Partitions of HorizonGS on MC-small-city+ (left) and Uni10k (right).

VastGaussian. *VastGaussian* [Lin et al. 2024] would have presented a natural comparison point, but there is currently no official code release available.

OccluGaussian. As of writing, the code for [Liu et al. 2025] has not been released, so we are only able to compare to their reported results on their own dataset. Because the method is specifically tailored to indoor scenes, they never published results on any outdoor street-level datasets.

Ours. In general, we perform 60k iterations of coarse training (except for 100k on MC-small-city+) and then 150k iterations of fine

training (except 100k on Uni10k, 250k on Campus and MC-smaller-city+ and 500k on MC-small-city+). We choose SH degree 1 and halve the densification gradient threshold as opposed to *H-3DGS*. We pick a cache size of 15 million Gaussians. For each scene, we pick a maximal number of Gaussians and desired SPT volume according to the scene scale. We deactivate the use of pinned CPU memory, which would accelerate memory transfers, because it would not be available in sufficient quantities on the consumer devices we are targeting. We provide detailed configuration files containing all hyperparameters for all of the scenes as supplemental material.

A.5 Implementation Details

Training models on this unprecedented scale on consumer hardware required solving many technical issues. In this section, we elaborate on the most important of these design decisions.

Codebase. Our general code structure is based on *H-3DGS* [Kerbl et al. 2024]. In particular, we reuse the hierarchy creator, but replace the cut procedure and chunk based training, as well as the rasterizer for the one used in Kerbl et al. [2023]. We also use the densification implementation of Kheradmand et al. [2024]. All of the code is implemented in PyTorch and C++/CUDA. The entire source code is included as supplemental material.

Varying level of detail. We introduce noise into the hierarchy cuts in order to prevent overfitting of Gaussians of a certain scale to views from a particular distance. In particular, we multiply the distances to the SPT centers each iteration with $1 + 5r^4$, with a uniformly random variable $r \in U(0, 1)$. This factor is designed such that most iterations will train near the highest level of detail, but coarser levels will also be trained occasionally. While this does not improve image metrics on the test set, it shows significant improvement to out-of-distribution views and reduces LoD popping artifacts.

Varying focal length. The required level of detail is not only dependent on a camera’s distance d to the Gaussian, but also on its focal length, which we need to account for as our datasets contain views with differing focal lengths. Therefore, we choose a base focal length f_b and use a relative distance metric \hat{d} for our HSPT cuts, which is calculated for the current camera’s focal length f_i as $\hat{d} = \frac{f_b}{f_i} d$. For example, given a camera with double the focal length, the cut distance should be halved to account for the larger projection footprint on the image plane.

SH Degree. For the most part, experiments indicate that SH degrees higher than 1 do not significantly contribute to image quality in the tested scenes. For training higher degrees n of SH, we have found that increasing the degree from 0 to 1 during the course stage and then gradually increasing the degree from 1 to n during fine training (with each increase happening after 10% of total training iterations) yields the best results.

Gaussian Order. To exploit spatial coherency and improve training performance, we store the Gaussians on CPU in Morton Z-order. As this order can change during training, we resort the Gaussians at every densification iteration.

Training Images. Conventional Gaussian Splatting stores the entire training dataset in VRAM to achieve their impressive training speed, however, this approach is infeasible for training ultra-large-scale datasets on consumer-grade hardware. Instead, we load the ground truth images from disk every iteration, saving VRAM at the cost of training performance.

Unreachable Gaussians. During training, it can occur that a child Gaussian becomes larger than its parent. This can lead to cases where $m_d(\text{parent}(i)) < m_d(i)$, making it impossible for the parent Gaussian to fulfil the SPT cut condition. While we are aware of this inefficiency, we found that in practice this only happens to a small number of Gaussians (< 10%). They still occupy a portion of RAM, but are never rendered or transferred. We have experimented with rebalancing the hierarchy during training, but found it too costly for little benefit. In general, we find this behaviour to be preferable over generating improper cuts, which can derail training.

Respawning, Densifying, and Pruning Gaussians. For improved performance, all respawn- and densify-operations are performed in parallel during the densification step. This can lead to difficult edge cases that need to be handled to prevent the hierarchy from degenerating: When two sibling nodes both need to be respawned, we only respawn the right node. This will cause the left node to become a new dead leaf, which will be respawned in the next densification iteration. If a node needs to be respawned whose sibling is not a leaf node, its entire subtree will replace the parent node. To minimize the number of unnecessary Gaussians, we apply a simple pruning strategy where we zero the opacity of Gaussians that have not contributed for a number of iterations equal to twice the size of the training set. This will cause them to be respawned in the next densification iteration.

Performance Optimization. The performance of the HSPT cut and Gaussian loading is particularly important, as they happen every iteration. We store the SPT properties ($m_d(i)$, $m_d(\text{parent}(i))$, i) for all Gaussians in a single continuous GPU memory buffer and perform the cuts in parallel using optimized CUDA kernels. Similarly, the Gaussian properties are stored in a single PyTorch tensor in RAM as an array of structures (each Gaussian’s properties concurrently). The properties of all required Gaussians are then transferred to the GPU via a single copy operation, reorganized on the GPU in a structure of arrays layout (as required by the rasterizer), and appended to the current render set.

A.6 Additional Densification Details

Our densification method combines the split operation from Kheradmand et al. [2024] and the selection method from Kerbl et al. [2024] with coarse-to-fine expansion of the LoD structure. Relying purely on the densification strategy of Kheradmand et al. [2024]—applying a noise to every Gaussian’s position and an opacity/scaling loss—becomes problematic for large scenes, where a majority of Gaussians are not visible from any one view: The result is a gradual disappearance of density in the scene or extreme “stringing” artifacts (cf. Figure 21). However, without these losses to encourage respawning of Gaussians, the densification becomes aimless and random. We implemented another strategy that only applies the losses and noise

Algorithm 2 SPT Cut

```


procedure CUT(SPT  $\mathcal{S}$ , distance  $d^j$ )
   $high \leftarrow \text{BINARYSEARCH}(\mathcal{S}_{max}, d^j)$ 
   $cut \leftarrow \{i \in [0, high] | \mathcal{S}_{min}^i < d^j\}$ 
  return  $cut$ 
end procedure

```

to Gaussians that affected at least one pixel in the current iteration. While this significantly improved results, we found that the MCMC densification strategy in general performs poorly when Gaussian density is low, which is necessitated by ultra-large-scale scenes. On small scenes such as the single-chunk datasets, our modified MCMC strategy slightly outperforms the presented strategy given sufficient Gaussian budget, but it lags behind in the larger scenes. Avoiding destructive opacity resets, as advocated by Kerbl et al. [2023], is particularly important in our setting, as many Gaussians will no longer conform to the hierarchical structure after opacity is restored to normal.



(a) Campus



(b) Small City

Fig. 21. Without regularization, the scale- and opacity-losses of 3DGS-as-MCMC lead to either stringing or disappearance of the entire scene.

A.7 Pseudocode

This section provides pseudocode for our core algorithms. Algorithm 1 contains a procedure to ”cut” a binary tree hierarchy along a certain condition. Algorithm 2 shows how to cut a single SPT at a particular distance d^j . The entire training procedure is sketched in Algorithm 3.

Algorithm 1 BFS Hierarchy Cut

```

procedure CUT(Hierarchy  $\mathcal{H}$ , condition)
   $cut \leftarrow \{\}$ 
   $Q \leftarrow \text{QUEUE}(\text{root}(\mathcal{H}))$ 
  while  $Q$  not empty do
     $node \leftarrow Q.\text{DEQUEUE}()$ 
    if condition( $node$ ) then
       $cut \leftarrow cut \cup node$ 
    else
       $Q.\text{ENQUEUE}(node \rightarrow left)$ 
       $Q.\text{ENQUEUE}(node \rightarrow right)$ 
    end if
  end while
  return  $cut$ 
end procedure

```

Algorithm 3 Full train procedure

```

procedure TRAIN(upper Hierarchy  $\mathcal{U}$ , cache, skybox)
  view  $\leftarrow$  0
  Cache_Distances  $\leftarrow$  {}
  Cache_SPTs  $\leftarrow$  {}
  while True do
    view  $\leftarrow$  SAMPLE(knn_graph, view) ▷ Find a new nearby view
    cut_condition  $\leftarrow$  [ $\lambda(i) = \neg \text{is\_in\_frustum}(\mu_i) \vee m_d(i) < T$ ]
    upper_cut  $\leftarrow$  CUT( $\mathcal{U}$ , cut_condition)
    SPT_roots  $\leftarrow$  { $u \in \text{upper\_cut} \mid u$  contains an SPT}
    upper_nodes_to_render  $\leftarrow$  upper_cut  $\setminus$  SPT_roots
    SPT_distances  $\leftarrow$  { $\|\mu_r - p_{\text{view}}\|_2^2$  for  $r \in \text{SPT\_roots}$ }
    Cache_Reuse_SPTs  $\leftarrow$  { $s_i \in (\text{Cache\_SPTs} \cap \text{SPT\_roots}) \mid D_{\min} \leq \left\| \frac{\mu_i - p_{\text{cam}}}{\text{Cache\_Distances}_i} \right\|_2 \leq D_{\max}$ }
    load_SPTs  $\leftarrow$  (SPT_roots  $\setminus$  Cache_Reuse_SPTs)
    load_Gaussian_indices  $\leftarrow$  CUT_SPTs(load_SPTs, SPT_distances[load_SPTs])
    load_Gaussians  $\leftarrow$  LOAD(load_Gaussian_indices)
    Gaussians  $\leftarrow$  skybox  $\cup$  upper_nodes_to_render  $\cup$  load_Gaussians
    Gaussians  $\leftarrow$  Gaussians  $\cup$  LOAD_CACHE(Cache_Reuse_SPTs)
    RENDER_AND_OPTIMIZE(Gaussians)
    Cache_SPTs  $\leftarrow$  Cache_SPTs  $\cup$  load_SPTs
    if len(Cache_Gaussians) > max_cache_size then
      WRITE_BACK_LRU(Cache_Gaussians) ▷ Write cached Gaussians to RAM
    end if
  end while
end procedure

```
