

URBANGS: A SCALABLE AND EFFICIENT ARCHITECTURE FOR GEOMETRICALLY ACCURATE LARGE-SCENE RECONSTRUCTION

Anonymous authors

Paper under double-blind review

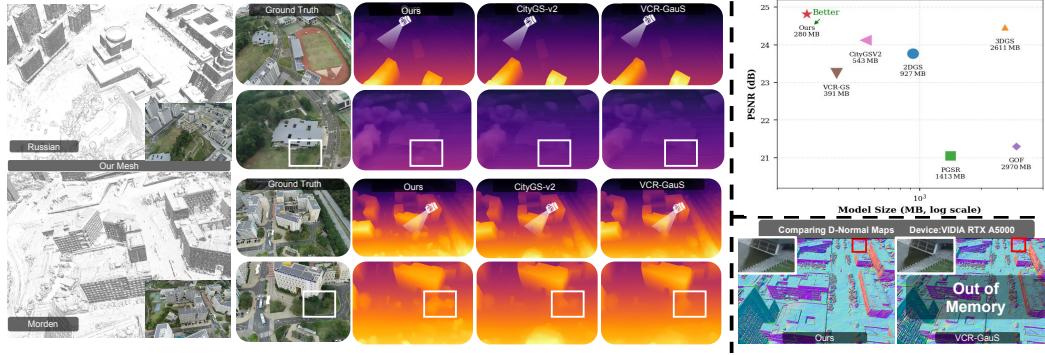


Figure 1: We propose UrbanGS, a scalable framework for high-fidelity large-scale scene reconstruction. Left: It reconstructs complex urban environments from multi-view RGB images, capturing fine details like trees, buildings, and roads. Middle: Compared with CityGS-v2 (Liu et al., 2024b) and VCR-Gaus (Chen et al., 2024b), by comparing rendered depth maps, our method can intuitively demonstrate its geometric advantages in terms of the surface smoothness of objects. Top-right: Our Spatially Adaptive Gaussian Pruning enables significant model compression while preserving quality. Bottom-right: UrbanGS efficiently reconstructs large scenes on an A5000 GPU, whereas VCR-Gaus (Chen et al., 2024b) fails due to out-of-memory issues.

ABSTRACT

While 3D Gaussian Splatting (3DGS) enables high-quality, real-time rendering for bounded scenes, its extension to large-scale urban environments gives rise to critical challenges in terms of geometric consistency, memory efficiency, and computational scalability. To address these issues, we present UrbanGS, a scalable reconstruction framework that effectively tackles these challenges for city-scale applications. First, we propose a Depth-Consistent D-Normal Regularization module. Unlike existing approaches that rely solely on monocular normal estimators, which can effectively update rotation parameters yet struggle to update position parameters, our method integrates D-Normal constraints with external depth supervision. This allows for comprehensive updates of all geometric parameters. By further incorporating an adaptive confidence weighting mechanism based on gradient consistency and inverse depth deviation, our approach significantly enhances multi-view depth alignment and geometric coherence, which effectively resolves the issue of geometric accuracy in complex large-scale scenes. To improve scalability, we introduce a Spatially Adaptive Gaussian Pruning (SAGP) strategy, which dynamically adjusts Gaussian density based on local geometric complexity and visibility to reduce redundancy. Additionally, a unified partitioning and view assignment scheme is designed to eliminate boundary artifacts and optimize computational load. Extensive experiments on multiple urban datasets demonstrate that UrbanGS achieves superior performance in rendering quality, geometric accuracy, and memory efficiency, providing a systematic solution for high-fidelity large-scale scene reconstruction.

054
055
056
1 INTRODUCTION

057 3D scene reconstruction is a long-standing research topic in computer vision and computer graphics,
 058 with its core objective of achieving photorealistic rendering and accurate geometric reconstruction.
 059 Following the introduction of Neural Radiance Fields (NeRF) (Mildenhall et al., 2021), 3D Gaus-
 060 sian Splatting (3DGS) (Kerbl et al., 2023) has emerged as a mainstream technique in this field,
 061 thanks to its advantages in training convergence and rendering efficiency. 3DGS represents scenes
 062 using a set of discrete Gaussian ellipsoids and leverages a highly optimized rasterizer for render-
 063 ing. However, due to the unstructured nature of 3DGS, accurately representing surfaces—especially
 064 in large-scale complex scenes—remains a significant challenge. In recent years, numerous promi-
 065 nent studies (Huang et al., 2024a; Chen et al., 2024b;a) have been proposed to address this issue.
 066 While these methods have achieved remarkable success in single-object or small-scale scene recon-
 067 struction, directly extending them to complex large-scale scenes reveals several critical limitations.
 068 For instance, vanilla 3DGS suffers from inadequate geometric modeling accuracy and incomplete
 069 parameter updates when applied to city-scale environments, failing to meet the high-fidelity recon-
 070 struction requirements of complex urban scenes.

070 To tackle the challenges of urban-scale modeling, various technical solutions have been developed.
 071 Methods such as CityGaussian (Liu et al., 2024a) and VastGaussian (Lin et al., 2024) have proposed
 072 block-wise partitioning strategies; although these strategies improve rendering efficiency, they still
 073 suffer from geometric inconsistencies, low geometric accuracy, and fail to reduce memory require-
 074 ments during training. CityGaussianv2 (Liu et al., 2024b) adopts a hybrid approach integrating 2D
 075 Gaussian Splatting (Huang et al., 2024a), while this accelerates training and enhances geometric
 076 accuracy, it comes at the cost of degraded rendering quality. Furthermore, vanilla 3DGS generates
 077 excessive redundant Gaussian primitives in homogeneous regions (e.g., skies, distant building fa-
 078 cades), and naive pruning heuristics often sacrifice fine-grained details (Fan et al., 2023). Existing
 079 partitioning schemes also introduce computational inefficiencies by processing irrelevant views and
 080 generating boundary discontinuities (Liu et al., 2024a). These limitations underscore the urgent
 081 need for a unified framework that balances geometric precision, memory efficiency, and seamless
 082 scalability.

083 We propose **UrbanGS**, a strategy that achieves high geometric accuracy, fidelity, and efficiency in
 084 large-scale scene reconstruction. To enhance geometric fidelity in large-scale settings, we directly
 085 supervise the rendered normal maps of 3D Gaussians with external pseudo-normal priors. However,
 086 this form of supervision alone is insufficient for updating the position parameters of Gaussians,
 087 which is critical for accurate surface reconstruction (Chen et al., 2024b). To overcome this limita-
 088 tion, we introduce a Depth-Consistent D-Normal Regularization framework. Instead of supervising
 089 the rendered normals directly, we first derive depth-normal (D-Normal) from the spatial gradient
 090 of the rendered depth maps, which are then supervised by the pseudo-normal priors. This estab-
 091 lishes a geometric constraint intrinsically linked to depth, thereby enabling comprehensive updates
 092 of both rotation and position parameters of the Gaussians. Furthermore, considering the limitation
 093 that supervision based on D-Normal relies on the accuracy of rendered depth maps, we introduce
 094 a depth estimator (Pseudo Depth) (Hu et al., 2024) to directly supervise the rendered depth maps,
 095 thereby constructing the “Pseudo Depth & D-Normal Dual Supervision Mechanism” (with theore-
 096 matical proofs provided in the supplementary materials). To ensure the reliability of depth alignment
 097 across multiple views, we propose an adaptive confidence weighting strategy that dynamically ad-
 098 justs supervision weights for different regions, thus reducing the impact of depth errors on surface
 099 reconstruction results.

100 To meet the memory and computational demands of urban-scale reconstruction, we propose a Spat-
 101 ially Adaptive Gaussian Pruning (SAGP) method. Traditional pruning approaches, designed for
 102 small-scale or object-level scenes, rely on global metrics or fixed thresholds (Fan et al., 2023). When
 103 applied to city-scale scenes with high spatial heterogeneity and numerous Gaussian primitives, such
 104 strategies often oversimplify local structures or lose fine details (see Table 4). To our knowledge,
 105 this is the first pruning framework specifically designed for city-scale 3D Gaussian Splatting. SAGP
 106 operates within local voxel cells, integrating local geometric complexity, ray-intersection frequency,
 107 and visibility-aware importance scores to decide which Gaussians to prune. This adaptively removes
 108 redundant primitives—especially in uniform or distant regions—while preserving perceptually and
 109 geometrically critical structures. Applied progressively during training, SAGP significantly reduces
 110 model complexity and memory usage while maintaining high rendering and geometric quality (Ta-

108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
ble 4, Fig. 4). We also incorporate a partitioning strategy (Liu et al., 2024a) to enable parallel processing, supporting efficient and scalable reconstruction of large-scale urban scenes. Our main contributions are summarized below:

- We propose a Depth-Consistent D-Normal Regularizer that enables holistic optimization of all Gaussian parameters (position, rotation), addressing the limitation of incomplete geometric updates in methods that supervise only rendered normals.
- We introduce an adaptive confidence term to enhance robustness, which suppresses unreliable depth predictions and strengthens multi-view geometric alignment.
- To address Gaussian redundancy and memory explosion in city-scale scenes, we design a Spatially Adaptive Gaussian Pruning (SAGP) algorithm that is aware of local geometric complexity.
- Extensive experiments demonstrate that our method outperforms existing large-scale scene reconstruction techniques, thus laying a solid foundation for future further research in this field.

2 RELATED WORK

Neural Rendering. Novel view synthesis and multi-view surface reconstruction are interconnected tasks in 3D scene reconstruction. Traditional reconstruction pipelines relied on Structure-from-Motion (SfM) (Duisterhof et al., 2024; Wang et al., 2023; He et al., 2024) and Multi-View Stereo (MVS) (Furukawa et al., 2015; Tang et al., 2024; Yao et al., 2018) with feature matching (Wang & Shen, 2018), but suffered from artifacts and noise sensitivity (Leroy et al., 2024). Early synthesis methods like Soft3D (Penner & Zhang, 2017) used volumetric ray-marching with high computational costs. The neural revolution began with NeRF (Mildenhall et al., 2021), which improved quality through positional encoding yet remained slow due to MLPs; variants like Mip-NeRF (Barron et al., 2022), InstantNGP (Müller et al., 2022), and Plenoxels (Fridovich-Keil et al., 2022) balanced efficiency but struggled with empty spaces. For reconstruction, implicit methods like NeuS (Wang et al., 2021) and Neuralangelo (Li et al., 2023) integrated SDFs (Park et al., 2019; Yu et al., 2022b) for detailed surfaces at the cost of lengthy training. The paradigm shifted with 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023), enabling real-time synthesis via unstructured Gaussians, though its explicit form caused reconstruction issues like depth ambiguities (Zhang et al., 2024; Chen et al., 2024a). Subsequent optimizations addressed both domains: synthesis-focused improvements included Mip-splatting (Yu et al., 2024a) and HRGS (Li et al., 2025), while reconstruction enhancements featured SuGaR’s mesh binding (Guédon & Lepetit, 2024) (despite scalability limits (Chen et al., 2024a)), 2DGS’s surfel-based normal alignment (Huang et al., 2024b), VCR-GauS’s depth-normal regularizers (Chen et al., 2024b), and GOF’s ray-tracing for unbounded scenes (Yu et al., 2024b). However, when reconstructing complex large-scale scenes, 3DGS faces considerable challenges in terms of rendering quality and geometric accuracy. Furthermore, it also has the problems of a surge in video memory usage and excessively long training times, all of which limit the further expansion of 3DGS in large-scale scenes.

Large-Scale Scene Reconstruction. Reconstructing large-scale scenes (e.g., urban areas, expansive landscapes) faces significant challenges, including computational inefficiency, memory constraints, and geometric inconsistencies across sub-scenes processed in a block-wise manner (Tancik et al., 2022). Early NeRF-based methods partitioned scenes into blocks for parallel training (Turki et al., 2022a; Zhang et al., 2025), but due to the limitations of multi-layer perceptrons (MLPs), these approaches suffered from slow rendering speeds and poor scalability (Kerbl et al., 2023). Although 3DGS-based methods improved efficiency, they introduced new issues: partition-and-merge strategies such as VastGaussian (Lin et al., 2024) often lead to boundary inconsistencies due to insufficient multi-view constraints; methods like CityGaussian (Liu et al., 2024a) require time-consuming post-processing for pruning or distillation; and while these methods improve rendering quality, they still struggle with geometric accuracy, training cost, and efficiency (Chen & Lee, 2024). More recently, CityGS-X (Gao et al., 2025) revisits large-scale 3DGS from a systems perspective, introducing a parallel hierarchical representation with multi-task supervision and progressive optimization that eliminates the partition-and-merge pipeline and improves geometric consistency under scalable multi-GPU training, but its surface reconstruction quality for high-fidelity urban details remains limited. Optimization-focused solutions like CityGaussianV2 (Liu et al., 2024b), despite employing techniques to control Gaussian proliferation, sacrifice rendering quality to some extent. To address these limitations, we propose the UrbanGS framework, which establishes a unified depth-normal regularizer for holistic geometric optimization, incorporates confidence-aware weighting to enhance

162 robustness, introduces spatially adaptive pruning to manage redundancy, and designs a seamless par-
 163 titioning scheme, collectively achieving high-fidelity, efficient, and geometrically consistent large-
 164 scale reconstruction.

166 3 METHODOLOGY

168 3.1 PRELIMINARIES

170 **3D Gaussian Splatting.** 3D Gaussian Splatting models a scene using a collection of anisotropic 3D
 171 Gaussians $G = \{G_i \mid i \in \mathbb{N}\}$. Each 3D Gaussian unit G_i is characterized by a center $u \in \mathbb{R}^3$ and a
 172 covariance matrix $\Sigma \in \mathbb{R}^{3 \times 3}$, and can be mathematically expressed as:

$$174 \quad G_i(p) = \exp \left\{ -\frac{1}{2} (p - u_i)^\top \Sigma_i^{-1} (p - u_i) \right\}. \quad (1)$$

176 During the training process, the covariance matrix is decomposed into a rotation matrix $R \in \mathbb{R}^{3 \times 3}$
 177 and a diagonal scaling matrix $S \in \mathbb{R}^{3 \times 3}$, that is,

$$179 \quad \Sigma_i = RSS^\top R^\top, \quad (2)$$

180 to ensure the covariance matrix is positive semi-definite. For rendering the color of a pixel p , the 3D
 181 Gaussians are projected into the image space for alpha blending:

$$183 \quad C = \sum_i c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (3)$$

186 where c_i and $\alpha_i = o_i G(x_i)$ denote the color and density of a point, respectively.

188 3.2 DEPTH-CONSISTENT D-NORMAL REGULARIZATION

190 **D-Normal Regularization.** To reconstruct scene surfaces, we enforce normal priors N predicted
 191 by a pretrained monocular deep neural network (Bae & Davison, 2024) to supervise the rendered
 192 normal map \hat{N} using L_1 and cosine losses:

$$193 \quad \mathcal{L}_n = \|\hat{N} - N\|_1 + (1 - \hat{N} \cdot N). \quad (4)$$

195 In our method, the depth map is rendered by performing a weighted sum of depths (Bae & Davison,
 196 2024; Chen et al., 2024b; Yu et al., 2022a), with the formula given as follows:

$$199 \quad \hat{D} = \frac{\sum_{i \in M} d_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j)}{\sum_{i \in M} \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j)}, \quad (5)$$

202 where d_i denotes the intersection depth (Chen et al., 2023; 2024b) and is distinct from the depth
 203 estimation in conventional 3D Gaussian Splatting (3DGS). Specifically, it refers to the distance from
 204 the camera to the intersection point calculated along the z-axis of the camera coordinate system; this
 205 intersection point is formed between the ray emitted from the camera center and the elliptical plane
 206 obtained by compressing the ellipsoid of 3DGS (Further details are provided in the supplementary
 207 material in the Appendix).

208 Additionally, to effectively update Gaussian positions, we utilize the predicted normal N from the
 209 pretrained model to supervise the D-Normal \bar{N}_d . The derivation of the D-Normal from the rendered
 210 depth involves two sequential steps. First, the rendered depth map is back-projected into point
 211 clouds $\{\mathbf{d}_k(n, p)\}$, using the camera intrinsic matrix. Subsequently, the horizontal and vertical finite
 212 differences are computed between adjacent points in this back-projected point cloud; the D-Normal
 213 is then obtained by calculating the cross-product of these two sets of finite differences.

$$214 \quad \bar{N}_d(n, p) = \frac{\nabla_v d(n, p) \times \nabla_h d(n, p)}{|\nabla_v d \times \nabla_h d|}, \quad (6)$$

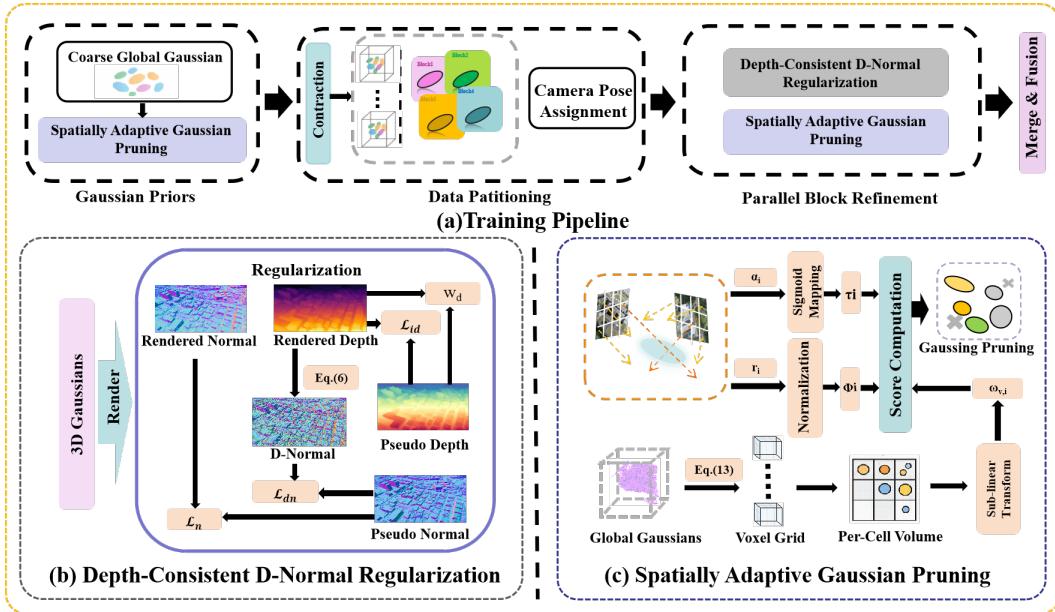


Figure 2: **UrbanGS training pipeline and core components.** (a) **Training Pipeline:** Starting from coarse global Gaussians, we apply spatially adaptive Gaussian pruning to obtain compact priors, contract and partition the scene into blocks, assign camera views using geometric and SSIM-based criteria, and refine all blocks in parallel before merging them into a unified large-scale 3D Gaussian scene. (b) **Depth-Consistent D-Normal Regularization:** 3D Gaussians are rendered to depth and normal maps, depth is converted to D-normals and jointly supervised with pseudo-depth and pseudo-normal priors from pretrained models via the loss $\mathcal{L}_n + \mathcal{L}_{dn} + w_d \mathcal{L}_{id}$, yielding stable and globally consistent geometry. (c) **Spatially Adaptive Gaussian Pruning:** Global Gaussians are discretized into a voxel grid, where per-cell importance $\omega_{v,i}$ and view-dependent cues are fused into pruning scores to remove redundant Gaussians and obtain an efficient yet accurate representation.

where d represents the 3D coordinates of a pixel obtained via back-projection from the depth map. We then apply the D-Normal regularization:

$$\mathcal{L}_{dn} = \|\bar{N}_d - N\|_1 + (1 - \bar{N}_d \cdot N), \quad (7)$$

Depth Consistency Regularization. In urban-scale scenes, D-Normal regularization optimizes geometry through normal-depth associations but lacks explicit cross-view depth constraints, frequently causing building misalignment and street distortion—especially in distant/complex areas. To resolve inconsistent multi-view depth predictions, we propose a depth consistency framework integrating inverse depth constraints with geometry-aware confidence. This extends normal-based regularization by incorporating robust priors from monocular depth estimators, where depth anchors D_{ext} (Hu et al., 2024) ensure cross-view consistency during optimization. Specifically, we derive dense relative depth anchors by processing training images with a pre-trained DepthAnything-v2 model (Hu et al., 2024). To align these monocular predictions with the unified metric scale of the 3D reconstruction, we leverage sparse 3D points from COLMAP’s Structure-from-Motion (SfM). Specifically, we compute per-view scale and shift parameters by robustly fitting the monocular depth maps to the sparse COLMAP depth values at valid 2D-3D correspondences. This process brings the relative depth estimates into alignment with the scale of the multi-view geometry. We define an inverse depth loss \mathcal{L}_{id} that operates on reciprocal depths to balance optimization sensitivity across distance ranges (Kerbl et al., 2024):

$$\mathcal{L}_{id}(u, v) = \left| \hat{D}^{-1}(u, v) - D_{\text{ext}}^{-1}(u, v) \right|. \quad (8)$$

where $\hat{D}^{-1} \equiv 1/\hat{D}$ is the reciprocal of the rendered depth map. This formulation minimizes relative depth errors per pixel while enhancing distant surface accuracy where linear depth gradients diminish.

270 Complementing this loss, we define a geometry-aware confidence measure w_d based on two geo-
 271 metric cues. First, the cosine similarity of depth gradients:
 272

$$273 \quad \cos \phi = \frac{\nabla \hat{D} \cdot \nabla D_{\text{ext}}}{\|\nabla \hat{D}\|_2 \|\nabla D_{\text{ext}}\|_2}, \quad (9)$$

274 quantifies gradient reliability by measuring local surface orientation consistency. Second, we mea-
 275 sure error sensitivity via normalized inverse depth deviation to suppress high-discrepancy regions:
 276

$$277 \quad \epsilon_d(u, v) = \frac{\mathcal{L}_{\text{id}}(u, v)}{\text{median}(\hat{D}^{-1})}. \quad (10)$$

278 The unified confidence w_d combines both cues through exponential decay:
 279

$$280 \quad w_d = \exp\left(\frac{\cos \phi - 1}{0.01}\right) \cdot \exp\left(-\frac{\epsilon_d}{0.1}\right), \quad (11)$$

281 where hyperparameters $\gamma_d = 0.01$ and $\tau = 0.1$ balance directional and magnitude sensitivity. The
 282 total optimization objective is consequently augmented to:
 283

$$284 \quad \mathcal{L}_{\text{total}} = \mathcal{L}_{\text{RGB}} + \lambda_1 \mathcal{L}_{\text{n}} + \lambda_2 \mathcal{L}_{\text{dn}} + \lambda_3 (w_d \cdot \mathcal{L}_{\text{id}}), \quad (12)$$

285 where $\lambda_i (i = 1, 2, 3)$ balancing the individual components. \mathcal{L}_{RGB} includes \mathcal{L}_1 and D-SSIM
 286 losses (Kerbl et al., 2023).
 287

288 3.3 SPATIALLY ADAPTIVE GAUSSIAN PRUNING (SAGP)

289 Large-scale 3D scenes exhibit pronounced spatial heterogeneity. Specifically, detailed foreground
 290 geometries necessitate dense Gaussian representations to accurately capture fine structures, while
 291 distant regions often suffer from excessive Gaussian proliferation—which in turn leads to pro-
 292hibitive memory consumption and degraded rendering performance. Existing pruning strategies,
 293 grounded in global significance metrics or fixed opacity thresholds (Kerbl et al., 2023; Fan et al.,
 294 2023), frequently oversimplify local features or inadvertently remove crucial far-field Gaussians,
 295 this ultimately results in incomplete reconstructions and visual artifacts.
 296

297 To overcome these limitations, we propose a unified, spatially adaptive pruning framework. The
 298 scene is first partitioned into volumetric cells whose characteristic length ℓ scales with the overall
 299 Gaussian density:
 300

$$301 \quad \ell = \lambda \left(\frac{\mathcal{V}_{\text{scene}}}{\mathcal{N}} \right)^{1/3}, \quad (13)$$

302 where $\mathcal{V}_{\text{scene}}$ denotes the bounding-box volume and \mathcal{N} the total number of Gaussians. We set
 303 $\lambda = 1.2$ to slightly enlarge the cell size for more stable local statistics.
 304

305 Within each cell, we compute the t -th percentile Gaussian volume $\vartheta_{\text{local}}^{(t)}$ and normalize individual
 306 volumes via a sub-linear transform:
 307

$$308 \quad w_{v,i} = \left(\min\left(\frac{v_i}{\vartheta_{\text{local}}^{(t)}}, 1\right) \right)^\kappa. \quad (14)$$

309 We use $t = 90\%$ to represent the typical volume in each cell while mitigating outlier influence.
 310 The sub-linear exponent $\kappa = 0.5$ (i.e., a square root) is applied to compress the dynamic range of
 311 volume ratios, thereby amplifying the importance of fine-scale structures while suppressing overly
 312 large Gaussians. This operation attenuates oversized background Gaussians while amplifying fine-
 313 scale structures, thereby establishing a context-aware basis for importance estimation.
 314

315 Building on these localized volume weights, we define each Gaussian’s importance score S_i as
 316 a weighted combination of three normalized attributes. The first attribute is the normalized ray-
 317 intersection frequency (Fan et al., 2023), given by $\phi_i = \frac{r_i}{\max_{j \in \mathcal{G}(i)} r_j}$, where r_i denotes the total
 318 number of intersections between the i -th Gaussian and all sampled rendering rays during the current
 319 training iteration. The second attribute is the Sigmoid-mapped opacity (Kerbl et al., 2023), expressed
 320

324 as $\tau_i = \sigma(a_i) = \frac{1}{1+e^{-a_i}}$, with a_i representing the raw opacity parameter of the i -th Gaussian
 325 (a learnable parameter in 3DGS that controls the Gaussian’s contribution to scene occlusion and
 326 visibility). The third attribute is the sub-linear volume weight, defined as $w_{v,i}$. The composite
 327 importance score is then given by
 328

$$S_i = \alpha \cdot \phi_i + \beta \cdot \tau_i + \gamma \cdot w_{v,i}, \quad (15)$$

331 where α, β, γ are tunable hyperparameters that balance visual salience, occlusion relevance, and
 332 geometric scale awareness. In our experiments, we use $\alpha = 1.2$, $\beta = 1.0$, and $\gamma = 0.8$ as default
 333 values to emphasize multi-view consistency while maintaining a balance with visual quality and
 334 geometric detail(See Table. F).

335 3.4 PARTITIONING STRATEGY

337 Our partitioning strategy is improved based on CityGS (Liu et al., 2024a), as illustrated in part (a)
 338 of Fig. 2. First, when obtaining the global coarse 3DGS model, we first eliminate redundant Gaus-
 339 sians through SAGP pruning to prevent these redundant Gaussians from attracting non-contributing
 340 views and amplifying the computational load during subsequent block-wise training. Then, in the
 341 partitioning phase, we retain common Gaussian primitives at the boundaries of each sub-block to
 342 avoid introducing visible fusion artifacts caused by geometric discontinuities between blocks. All
 343 other modules follow the methodologies of CityGS, and the specific formulas are referred to in the
 344 supplementary materials C.

345 4 EXPERIMENTS

348 Table 1: Quantitative comparisons on the Mill19 (Turki et al., 2022c) and UrbanScene3D (Lin et al.,
 349 2022) datasets for novel view synthesis. \uparrow indicates higher is better, while \downarrow indicates lower is better.
 350 The top three results are highlighted with red, orange, and yellow backgrounds, respectively. \dagger
 351 denotes results obtained without the decoupled appearance encoding.

	Building			Rubble			Residence			Sci-Art		
	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow
w/o Geometric Optimization												
Mega-NeRF	0.547	20.92	0.454	0.553	24.06	0.508	0.628	22.08	0.401	0.770	25.60	0.312
Switch-NeRF	0.579	21.54	0.397	0.562	24.31	0.478	0.654	22.57	0.352	0.795	26.51	0.271
VastGaussian \dagger	0.728	21.80	0.225	0.742	25.20	0.264	0.699	21.01	0.261	0.761	22.64	0.261
3DGS	0.738	22.53	0.214	0.725	25.51	0.316	0.745	22.36	0.247	0.791	24.13	0.262
DoGaussian	0.759	22.73	0.204	0.765	25.78	0.257	0.740	21.94	0.244	0.804	24.42	0.219
CityGaussian	0.778	21.55	0.246	0.813	25.77	0.228	0.813	22.00	0.211	0.837	21.39	0.230
w/ Geometric Optimization												
SuGaR	0.507	17.76	0.455	0.577	20.69	0.453	0.603	18.74	0.406	0.698	18.60	0.349
NeuS	0.463	18.01	0.611	0.480	20.46	0.618	0.503	17.85	0.533	0.633	18.62	0.472
Neuralangelo	0.582	17.89	0.322	0.625	20.18	0.314	0.644	18.03	0.263	0.769	19.10	0.231
PGSR	0.480	16.12	0.573	0.728	23.09	0.334	0.746	20.57	0.289	0.799	19.72	0.275
VCR-Gaus	0.502	19.56	0.502	0.541	21.34	0.428	0.623	20.59	0.359	0.665	19.31	0.465
CityGaussianV2	0.650	19.07	0.397	0.720	23.75	0.322	0.769	21.15	0.234	0.810	20.66	0.266
Ours	0.802	22.82	0.208	0.791	26.25	0.210	0.823	22.48	0.205	0.824	22.62	0.279

365 4.1 EXPERIMENTAL SETUP

368 Our experiments cover seven representative scenes drawn from four datasets: Building and Rubble
 369 from Mill-19 (Turki et al., 2022c); Residence and Sci-Art from UrbanScene3D (Lin et al., 2022);
 370 and Residence, Russian Building, and Modern Building from GauU-Scene (Xiong et al., 2024).
 371 Unless otherwise noted, competing methods were evaluated on RTXA800 GPUs, while UrbanGS
 372 was trained on eight RTXA5000 GPUs. Additional details on training protocols and evaluation
 373 settings are provided in the supplementary material.

375 4.2 MAIN RESULTS

377 **Novel View Synthesis.** As shown in Table 1 and Fig. 3, we present quantitative and qualitative eval-
 378 uations of large-scale scene reconstruction methods with and without geometric optimization (de-

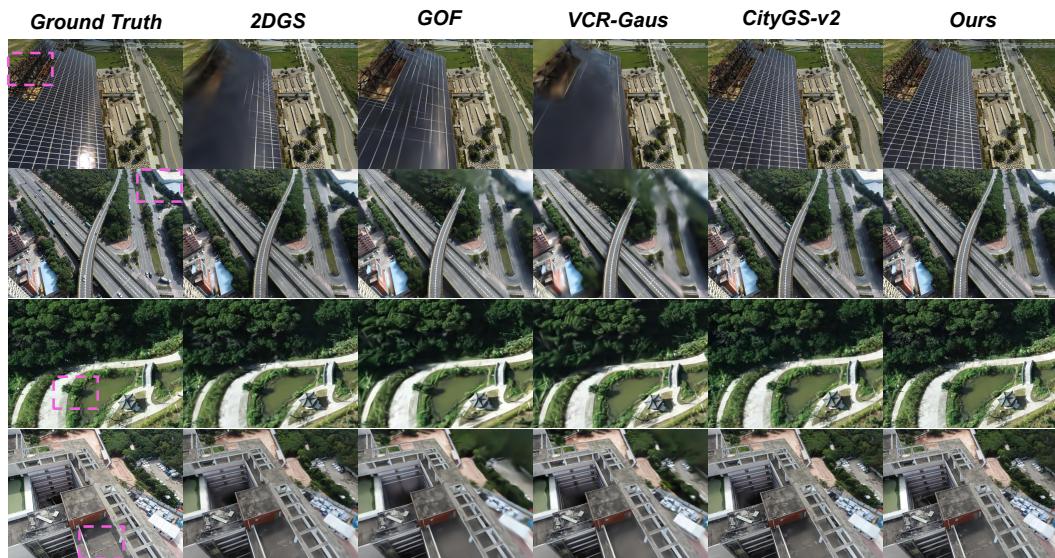


Figure 3: Qualitative results of ours and other methods in image rendering on Mill-19 (Turki et al., 2022c) and Urbanscene3D (Lin et al., 2022).

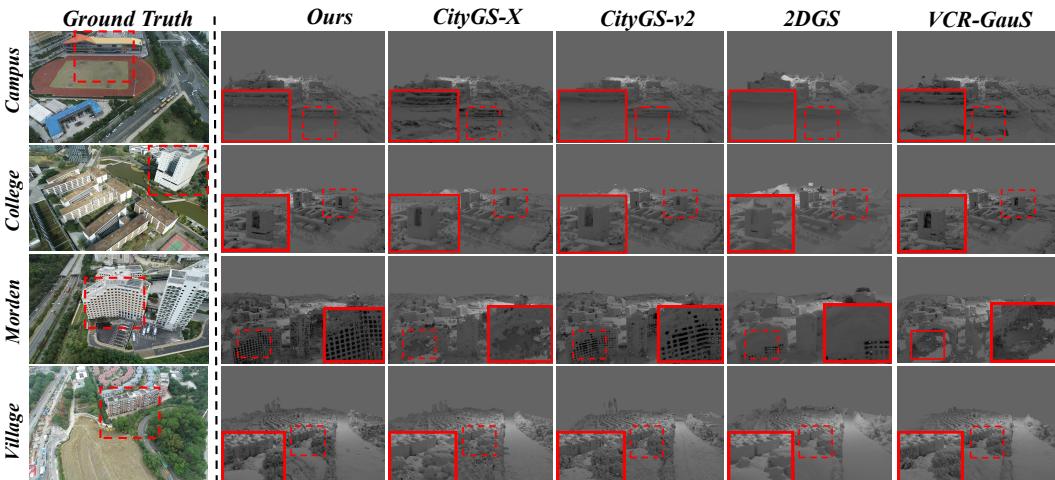


Figure 4: Qualitative mesh and texture comparison between SOTA and our method on [GauU-Scene dataset](#) (Xiong et al., 2024).

noted as w/ Geometric Optimization and w/o Geometric Optimization, respectively). Our method, UrbanGS, consistently achieves state-of-the-art performance—even when compared to methods without geometric constraints. Specifically, it attains the highest PSNR and SSIM in building scenes; in residential scenes, it reduces the LPIPS by 0.006 compared to CityGS (Liu et al., 2024a), demonstrating the robustness of our method. Qualitative comparisons in Fig. 3 show that our method can effectively mitigate the impacts of floating artifacts and lighting inconsistencies, resulting in superior rendering quality. These results highlight the robustness of UrbanGS in maintaining multi-view consistency and faithfully preserving appearance information, thereby enabling accurate large-scale scene reconstruction.

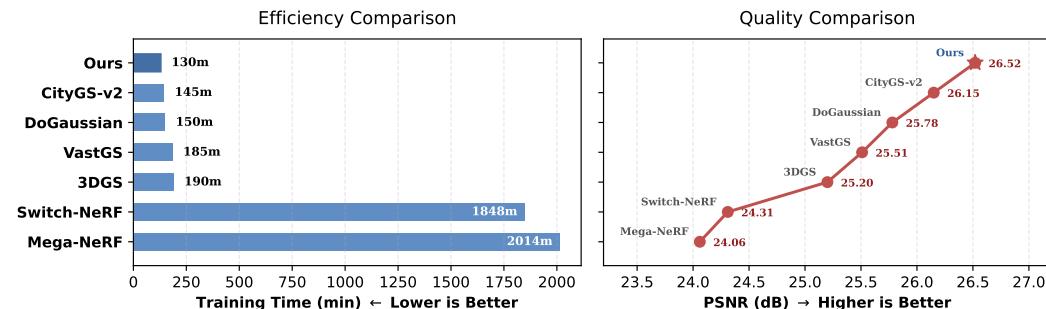
Surface Reconstruction. We compare our method with existing surface reconstruction approaches on the GauU-Scene datasets (Xiong et al., 2024). As shown in Table 2, our method achieves state-of-the-art performance among both neural implicit baselines and recent 3DGS-based city-scale methods. In particular, compared with CityGS-X, our approach attains higher F1 scores across all scenes

432 Table 2: Detailed geometry evaluation on the GauU-Scene dataset (Xiong et al., 2024). “NaN”
 433 indicates that the method produced invalid numerical results, while “FAIL” denotes a failure to
 434 extract a valid mesh. For all metrics, \uparrow indicates that higher values are better.

Methods	Residence			Russian Building			Modern Building		
	P \uparrow	R \uparrow	F1 \uparrow	P \uparrow	R \uparrow	F1 \uparrow	P \uparrow	R \uparrow	F1 \uparrow
NeuS	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL
Neuralangelo	NaN	NaN	NaN	FAIL	FAIL	FAIL	NaN	NaN	NaN
SuGaR	0.579	0.287	0.384	0.480	0.369	0.417	0.650	0.220	0.329
GOF	0.404	0.418	0.411	0.294	0.394	0.330	0.411	0.357	0.382
VCR-Gaus	0.498	0.402	0.445	0.538	0.454	0.492	0.591	0.401	0.478
2DGS	0.526	0.406	0.458	0.544	0.519	0.531	0.588	0.413	0.485
CityGS-X	0.512	0.411	0.456	0.572	0.516	0.542	0.653	0.389	0.487
CityGaussianV2	0.524	0.421	0.467	0.560	0.530	0.544	0.643	0.398	0.492
Ours	0.529	0.461	0.493	0.568	0.525	0.546	0.662	0.408	0.503

448 Table 3: Under the GauU-Scene dataset (Lin et al., 2022), comparison of Large-Scale Scene Mod-
 449 eling Methods, the best result for specific metrics under each scene is highlighted in **bold**.

Scene	Method	PSNR \uparrow	F1 \uparrow	#GS(M) \downarrow	Size(G) \downarrow	Mem.(G) \downarrow
Residence	CityGS	23.17	0.453	8.05	0.44	31.5
	CityGS-v2	23.46	0.465	8.07	0.44	14.2
	Ours	23.78	0.493	7.78	0.37	13.2
Russia	CityGS	24.19	0.455	7.00	0.38	27.4
	CityGS-v2	23.89	0.537	6.97	0.38	15.0
	Ours	24.53	0.546	6.56	0.35	11.4
Modern	CityGS	26.22	0.462	7.90	0.43	29.2
	CityGS-v2	25.53	0.489	7.90	0.42	16.1
	Ours	26.44	0.503	7.45	0.39	15.0



473 Figure 5: Experimental results on the Rubble dataset (Turki et al., 2022c) demonstrate that the
 474 proposed method outperforms comparative approaches in terms of PSNR while achieving superior
 475 training efficiency.

477 by improving recall while maintaining comparable precision. It also surpasses CityGS-v2 (Liu et al.,
 478 2024b) on most metrics. Qualitative comparisons in Figure 4 further show that our method produces
 479 more detailed and clearer surface structures. These results indicate that the proposed framework ef-
 480 fectively learns high-fidelity geometric representations and enables accurate surface reconstruction.
 481 Additional mesh visualizations are provided in Appendix B.

482 **Efficiency Comparison.** We compare the training time of our method with that of existing methods.
 483 As shown in Fig. 5, our method only takes 2 hours and 10 minutes to complete the training on
 484 the Rubble (Lin et al., 2022), which is significantly faster than competing methods. As presented
 485 in Table. 3, when compared with other large-scale scene algorithms, our method requires lower
 486 computational costs while achieving better rendering quality and geometric accuracy.

486 Table 4: Ablation Results on Russian dataset (Xiong et al., 2024). **Bold** indicates best performance.
 487 Note that OOM denotes Out Of Memory.

Method	Rendering Quality			Geometric Quality			Training Statistics			
	PSNR↑	SSIM↑	LPIPS↓	P↑	R↑	F1↑	GS (M)↓	Time↓	Size↓	Mem↓
Baseline	22.54	0.778	0.231	0.532	0.501	0.516	6.43	235	1102.23	OOM
+ST	24.68	0.816	0.188	0.571	0.518	0.543	6.37	188	1035.02	26.3
+LP	24.53	0.785	0.195	0.556	0.502	0.528	3.02	134	467.47	17.1
+SAGP (Ours)	24.66	0.813	0.184	0.568	0.525	0.546	2.45	122	314.24	14.4
STPG	24.57	0.801	0.201	0.563	0.511	0.536	2.73	119	320.12	13.9

496 Table 5: Ablation study on the effects of D-Normal Regularization and Depth Consistency Reg-
 497 ularization, conducted on the Morden Building dataset (Xiong et al., 2024). **Bold** indicates best
 498 performance.

Method	PSNR↑	SSIM↑	LPIPS↓	F1↑
w/o D-Normal	25.02	0.743	0.215	0.463
w/o Depth Consistency	24.59	0.792	0.201	0.453
w/o Geometry-Aware Confidence	26.02	0.795	0.163	0.493
Full	26.44	0.805	0.157	0.503

506 4.3 ABLATION STUDIES

507 To validate the effectiveness of individual components in our method, we conduct a series of ablation
 508 studies on the GauU-Scene dataset. Specifically, we evaluate the impact of the following compo-
 509 nents: Spatially Adaptive Gaussian Pruning (SAGP), Depth-Consistent D-Normal Regularization,
 510 and the partitioning strategy.

511 **Ablation of SAGP & Gaussian Partitioning**. As summarized in Tab. 4, we conduct a systematic
 512 ablation to evaluate the individual contributions of our proposed SAGP and partitioning strategy. We
 513 first establish a Baseline that employs neither our SAGP nor any partitioning strategy.

514 **Ablation on SAGP**. We compare our SAGP pruning against LP, the pruning method from LightGaus-
 515 sian (Fan et al., 2023). The results demonstrate that our SAGP is more effective at preserving the
 516 original geometric quality (higher F1 score) while significantly reducing the number of Gaussians,
 517 training time, and memory consumption, with only a minor impact on rendering quality.

518 **Ablation on Partitioning Strategy**. Our full method (Ours) integrates the proposed partitioning strat-
 519 egy (ST) with SAGP. We further compare it against STPG, which uses the partitioning strategy from
 520 CityGaussian (Liu et al., 2024a) with our SAGP. The comparison validates the superior effective-
 521 ness of our partitioning strategy, as it achieves better rendering and geometric quality under the same
 522 pruning method, demonstrating its ability to better preserve structural consistency across blocks.

523 **Ablation of Depth-Consistent D-Normal Regularization**. As shown in Tab. 5, we conduct abla-
 524 tion studies on each component of the Depth-Consistent D-Normal Regularization, demon-
 525 strating that its introduction significantly enhances both rendering quality and geometric accuracy for large-
 526 scale scenes. Quantitative results reveal consistent improvements across all evaluation metrics, with
 527 notable gains in F1-score (from 0.453 to 0.503) and PSNR (from 24.59 to 26.44), validating the
 528 critical importance of this component for high-quality large-scale reconstruction. Furthermore, as
 529 illustrated in Fig. D, the Geometric Regularization substantially improves the details in rendered
 530 images, as well as the quality of rendered normal and depth maps.

532 5 CONCLUSIONS

533 This paper presents UrbanGS, a scalable framework for urban-scale scene reconstruction. It in-
 534 troduces a depth-consistent D-Normal regularizer that enables comprehensive optimization of all
 535 Gaussian geometric parameters by fusing depth and normal cues. A spatial pruning strategy and
 536 seamless partitioning further enhance efficiency and avoid artifacts. Experiments show UrbanGS
 537 outperforms existing methods in rendering, geometry, and training speed, offering a practical solu-
 538 tion for large-scale 3D reconstruction.

540 Ethics Statement

541
 542 This work presents UrbanGS, a scalable framework for high-fidelity large-scale scene reconstruc-
 543 tion. The research focuses on methodological innovation to address challenges in geometric con-
 544 sistency, memory efficiency, and computational scalability of 3D Gaussian Splatting in urban-scale
 545 applications. All experiments use publicly available benchmark datasets (Mill-19, UrbanScene3D,
 546 GauU-Scene) in line with academic practices, involving no human subjects, personal data, or so-
 547 cial risk assessment. The authors encourage ethical and legal use of this technology and declare no
 548 potential conflicts of interest.

549 Reproducibility Statement

550 To ensure reproducibility of UrbanGS's results, we provide key details: the proposed components
 551 (Depth-Consistent D-Normal Regularization, Spatially Adaptive Gaussian Pruning, partitioning
 552 strategy) are detailed in the methodology section with mathematical formulations . Experimental
 553 setups include training on 8 NVIDIA RTX A5000 GPUs (baselines on RTX A800), using PyTorch
 554 2.0+, Open3D 0.18.0+, and pretrained models (DepthAnything-v2, Dsine) . Dataset preparation
 555 follows the image downsampling strategy (resizing images wider than 1600 pixels) and original
 556 train/validation splits . We will make the complete code and training scripts publicly available on
 557 GitHub upon the final revision and acceptance of this paper.

558 REFERENCES

559
 560 Gwangbin Bae and Andrew J. Davison. Rethinking inductive biases for surface normal estimation.
 561 In *CVPR*, 2024.

562
 563 Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf
 564 360: Unbounded anti-aliased neural radiance fields. In *CVPR*, 2022.

565
 566 Danpeng Chen, Hai Li, Weicai Ye, Yifan Wang, Weijian Xie, Shangjin Zhai, Nan Wang, Haomin
 567 Liu, Hujun Bao, and Guofeng Zhang. PGSR: planar-based gaussian splatting for efficient and
 568 high-fidelity surface reconstruction. *CoRR*, abs/2406.06521, 2024a.

569
 570 Hanlin Chen, Chen Li, and Gim Hee Lee. Neusg: Neural implicit surface reconstruction with 3d
 571 gaussian splatting guidance. *CoRR*, abs/2312.00846, 2023.

572
 573 Hanlin Chen, Fangyin Wei, Chen Li, Tianxin Huang, Yunsong Wang, and Gim Hee Lee. Vcr-gaus:
 574 View consistent depth-normal regularizer for gaussian surface reconstruction. In Amir Globersons,
 575 Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and
 576 Cheng Zhang (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference
 577 on Neural Information Processing Systems 2024, Vancouver, BC, Canada, December 10 - 15,
 578 2024*, 2024b.

579
 580 Yu Chen and Gim Hee Lee. DOGS: distributed-oriented gaussian splatting for large-scale 3d recon-
 581 struction via gaussian consensus. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela
 582 Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural Infor-
 583 mation Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024,
 584 NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024.

585
 586 Bardienus Duisterhof, Lojze Zust, Philippe Weinzaepfel, Vincent Leroy, Yohann Cabon, and Jerome
 587 Revaud. Mast3r-sfm: a fully-integrated solution for unconstrained structure-from-motion. *arXiv
 588 preprint arXiv:2409.19152*, 2024.

589
 590 Zhiwen Fan, Kevin Wang, Kairun Wen, Zehao Zhu, Dejia Xu, and Zhangyang Wang. Lightgaus-
 591 sian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps. *arXiv preprint
 592 arXiv:2310.01534*, 2023.

593
 594 Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo
 595 Kanazawa. Plenoxels: Radiance fields without neural networks. 2022.

596
 597 Xiao Fu, Wei Yin, Mu Hu, Kaixuan Wang, Yuexin Ma, Ping Tan, Shaojie Shen, Dahua Lin, and
 598 Xiaoxiao Long. Geowizard: Unleashing the diffusion priors for 3d geometry estimation from a
 599 single image. In *European Conference on Computer Vision*, pp. 241–258. Springer, 2024.

594 Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and*
 595 *Trends® in Computer Graphics and Vision*, 2015.

596

597 Yuanyuan Gao, Hao Li, Jiaqi Chen, Zhengyu Zou, Zhihang Zhong, Dingwen Zhang, Xiao Sun, and
 598 Junwei Han. Citygs-x: A scalable architecture for efficient and geometrically accurate large-scale
 599 scene reconstruction. *CoRR*, abs/2503.23044, 2025.

600 Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d
 601 mesh reconstruction and high-quality mesh rendering. In *IEEE/CVF Conference on Computer*
 602 *Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pp. 5354–
 603 5363, 2024.

604 Xingyi He, Jiaming Sun, Yifan Wang, Sida Peng, Qixing Huang, Hujun Bao, and Xiaowei Zhou.
 605 Detector-free structure from motion. In *CVPR*, 2024.

606

607 Zao Hu, Haocheng Wang, Xuan Rong, Danni Zhang, Zhang Xu, and Fangli Guan. Integrating
 608 depth-anything-v2 depth estimation and sobel operator matrix for UAV landing-site detection.
 609 In *International Conference on Cyberworlds, CW 2024, Kofu, Japan, October 29-31, 2024*, pp.
 610 342–343. IEEE, 2024.

611 Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting
 612 for geometrically accurate radiance fields. In *SIGGRAPH*, 2024a.

613

614 Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for
 615 geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers, SIGGRAPH*
 616 *2024, Denver, CO, USA, 27 July 2024- 1 August 2024*, pp. 32, 2024b.

617 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splat-
 618 ting for real-time radiance field rendering. *ACM Trans. Graph.*, 2023.

619 Bernhard Kerbl, Andreas Meuleman, Georgios Kopanas, Michael Wimmer, Alexandre Lanvin, and
 620 George Drettakis. A hierarchical 3d gaussian representation for real-time rendering of very large
 621 datasets. *ACM Trans. Graph.*, 43(4):62:1–62:15, 2024.

622

623 Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r.
 624 *arXiv preprint arXiv:2406.09756*, 2024.

625 Changbai Li, Haodong Zhu, Hanlin Chen, Juan Zhang, Tongfei Chen, Shuo Yang, Shuwei Shao,
 626 Wenhao Dong, and Baochang Zhang. HRGS: hierarchical gaussian splatting for memory-efficient
 627 high-resolution 3d reconstruction. *CoRR*, abs/2506.14229, 2025.

628

629 Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H. Taylor, Mathias Unberath, Ming-Yu Liu, and
 630 Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *CVPR*, 2023.

631 Jiaqi Lin, Zhihao Li, Xiao Tang, Jianzhuang Liu, Shiyong Liu, Jiayue Liu, Yangdi Lu, Xiaofei Wu,
 632 Songcen Xu, Youliang Yan, and Wenming Yang. Vastgaussian: Vast 3d gaussians for large scene
 633 reconstruction. In *CVPR*, 2024.

634

635 Liqiang Lin, Yilin Liu, Yue Hu, Xingguang Yan, Ke Xie, and Hui Huang. Capturing, reconstructing,
 636 and simulating: The urbanscene3d dataset. In Shai Avidan, Gabriel J. Brostow, Moustapha Cissé,
 637 Giovanni Maria Farinella, and Tal Hassner (eds.), *Computer Vision - ECCV 2022 - 17th Euro-
 638 pean Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part VIII*, volume 13668 of
 639 *Lecture Notes in Computer Science*, pp. 93–109. Springer, 2022.

640

641 Yang Liu, Chuanchen Luo, Lue Fan, Naiyan Wang, Junran Peng, and Zhaoxiang Zhang. Citygaus-
 642 sian: Real-time high-quality large-scale scene rendering with gaussians. In Ales Leonardis, Elisa
 643 Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gü̈l Varol (eds.), *Computer Vision
 644 - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Pro-
 ceedings, Part XVI*, volume 15074 of *Lecture Notes in Computer Science*, pp. 265–282. Springer,
 645 2024a.

646 Yang Liu, Chuanchen Luo, Zhongkai Mao, Junran Peng, and Zhaoxiang Zhang. Citygaussianv2: Ef-
 647 ficient and geometrically accurate reconstruction for large-scale scenes. *CoRR*, abs/2411.00771,
 648 2024b.

648 Zhenxing Mi and Dan Xu. Switch-nerf: Learning scene decomposition with mixture of experts
 649 for large-scale neural radiance fields. In *The Eleventh International Conference on Learning*
 650 *Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023.

651

652 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and
 653 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications*
 654 *of the ACM*, 2021.

655 Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics prim-
 656 itives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, 2022.

657

658 Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove.
 659 DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*,
 660 2019.

661 Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *ACM Transactions on Graph-*
 662 *ics*, 2017.

663 René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust
 664 monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transac-*
 665 *tions on pattern analysis and machine intelligence*, 44(3):1623–1637, 2020.

666

667 Matthew Tancik, Vincent Casser, Xincheng Yan, Sabeek Pradhan, Ben P. Mildenhall, Pratul P. Sriniv-
 668 asan, Jonathan T. Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view
 669 synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022,*
 670 *New Orleans, LA, USA, June 18-24, 2022*, pp. 8238–8248. IEEE, 2022.

671 Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative
 672 gaussian splatting for efficient 3d content creation. 2023.

673

674 Zhenggang Tang, Yuchen Fan, Dilin Wang, Hongyu Xu, Rakesh Ranjan, Alexander G. Schwing, and
 675 Zhicheng Yan. Mv-dust3r+: Single-stage scene reconstruction from sparse views in 2 seconds.
 676 *CoRR*, 2024.

677 Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-nerf: Scalable construction
 678 of large-scale nerfs for virtual fly- throughs. In *IEEE/CVF Conference on Computer Vision and*
 679 *Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pp. 12912–12921.
 680 IEEE, 2022a.

681 Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-nerf: Scalable construction of
 682 large-scale nerfs for virtual fly-throughs. In *CVPR*, 2022b.

683

684 Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-nerf: Scalable construction
 685 of large-scale nerfs for virtual fly- throughs. In *IEEE/CVF Conference on Computer Vision and*
 686 *Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pp. 12912–12921.
 687 IEEE, 2022c.

688 Jianyuan Wang, Nikita Karaev, Christian Rupprecht, and David Novotny. Visual geometry grounded
 689 deep structure from motion. *arXiv preprint arXiv:2312.04563*, 2023.

690

691 Kaixuan Wang and Shaojie Shen. Mvdepthnet: Real-time multiview depth estimation neural net-
 692 work. In *3DV*, 2018.

693

694 Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus:
 695 Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv*
 696 *preprint arXiv:2106.10689*, 2021.

697

698 Xiuchao Wu, Jiamin Xu, Xin Zhang, Hujun Bao, Qixing Huang, Yujun Shen, James Tompkin,
 699 and Weiwei Xu. Scanerf: Scalable bundle-adjusting neural radiance fields for large-scale scene
 700 rendering. *ACM Trans. Graph.*, 42(6):261:1–261:18, 2023.

701

Butian Xiong, Nanjun Zheng, Junhua Liu, and Zhen Li. Gauu-scene V2: assessing the reli-
 702 ability of image-based metrics with expansive lidar image dataset using 3dgs and nerf. *CoRR*,
 703 abs/2404.04880, 2024.

702 Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstruc-
703 tured multi-view stereo. In *ECCV*, 2018.

704

705 Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Ex-
706 ploring monocular geometric cues for neural implicit surface reconstruction. In *NeurIPS 2022*,
707 2022a.

708 Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Ex-
709 ploring monocular geometric cues for neural implicit surface reconstruction. 2022b.

710

711 Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-
712 free 3d gaussian splatting. In *CVPR*, 2024a.

713 Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian opacity fields: Efficient adaptive surface
714 reconstruction in unbounded scenes. *ACM Trans. Graph.*, 43(6):271:1–271:13, 2024b.

715

716 Wenyuan Zhang, Yu-Shen Liu, and Zhizhong Han. Neural signed distance function inference
717 through splatting 3d gaussians pulled on zero-level set. In Amir Globersons, Lester Mackey,
718 Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Ad-
719 vances in Neural Information Processing Systems 38: Annual Conference on Neural Information
720 Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024.

721 Yuqi Zhang, Guanying Chen, and Shuguang Cui. Efficient large-scale scene representation with a
722 hybrid of high-resolution grid and plane features. *Pattern Recognit.*, 158:111001, 2025.

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 **A IMPLEMENTATION DETAILS**
757758 **Training Setup.** UrbanGS are trained using NVIDIA A5000 GPUs, while all baseline methods are
759 trained on NVIDIA A800 GPUs. Since the Mill-19 (Turki et al., 2022c), UrbanScene3D (Lin et al.,
760 2022), and GauU-Scene (Xiong et al., 2024) datasets contain thousands of high-resolution images,
761 we follow the image downsampling strategy proposed in 3DGS: any image with a width exceeding
762 1600 pixels is resized proportionally during both training and validation. For geometric priors, we
763 utilize the DepthAnything-v2 model (Hu et al., 2024) for depth prediction and the pre-trained Dsine
764 model (Bae & Davison, 2024) for surface normal estimation.765 Regarding the pruning schedule, our design follows the training dynamics of 3DGS and prior prac-
766 tice. As shown in the pipeline 2, we use two stages of pruning. When constructing the coarse global
767 Gaussian model, we apply an initial, simple pruning rule to remove obviously redundant Gaussians,
768 reduce memory, and obtain a compact global prior for subsequent block-wise training. During block
769 refinement, we prune at 7k, 15k, and 25k iterations (out of 30k). The 7k step is applied after the
770 scene has roughly formed and the Gaussian distribution starts to stabilize, consistent with the behav-
771 ior observed in 3DGS (Kerbl et al., 2023), and removes early exploratory Gaussians that no longer
772 contribute to the final geometry. The 15k step follows the original 3DGS setting, occurring at the
773 end of densification when the Gaussian count peaks, and is most effective for controlling model
774 complexity and overfitting. The final pruning at 25k, inspired by LightGaussian (Fan et al., 2023),
775 acts as a consolidation step near convergence, further eliminating residual redundancy and ensuring
776 a good balance between high-fidelity reconstruction and compact, efficient rendering.777 **Mesh Extraction.** To obtain the final mesh, we employ Open3D’s volumetric TSDF fusion method,
778 which integrates rendered depth maps and corresponding camera poses to construct a continuous
779 Signed Distance Field (SDF). The surface is then extracted using the Marching Cubes algorithm at
780 the zero-level isosurface, enabling direct reconstruction of 3D geometry without relying on interme-
781 diate point cloud representations.782 **B PROOF ON A DEPTH-CONSISTENT D-NORMAL REGULARIZER**
783784 These propositions systematically validate the evolutionary process from traditional rendered normal
785 supervision to our proposed depth-normal regularizer. Proposition 1.1 reveals the limitation of su-
786 pervising only rendered normals in updating Gaussian positions; Proposition 1.2 demonstrates that
787 the depth-normal regularizer can effectively optimize Gaussian positions; Proposition 2.1 further
788 proves that the depth-consistent regularizer significantly improves geometric accuracy, highlighting
789 the enhanced effectiveness of our method.790 **B.1 GEOMETRIC PROPERTIES**
791792 To reconstruct the 3D surface, we focus on the geometric properties of Gaussians that enable accu-
793 rate intersection depth calculation, as detailed below.794 **Normal Vector** Following NeuSG (Chen et al., 2023), the Gaussian’s normal vector $\mathbf{n} \in \mathbb{R}^3$ is
795 defined as the direction of its minimized scaling factor:

796
$$\mathbf{n} = \mathbf{R}[k, :], \quad k = \arg \min ([s_1, s_2, s_3]). \quad (16)$$

797

798 Both \mathbf{n} and Gaussian center \mathbf{p} are transformed to the camera coordinate system (default unless stated
799 otherwise).800 **Intersection Depth** Existing work (Tang et al., 2023) uses \mathbf{p} for depth calculation, which is in-
801 accurate (depth unrelated to \mathbf{n}). We instead compute the ray-Gaussian intersection depth via the
802 following steps:803 Gaussian Flattening with Scale Regularization: To simplify intersection computation, we adopt
804 NeuSG’s (Chen et al., 2023) scale loss to flatten 3D Gaussian ellipsoids into planes (\mathbf{p}, \mathbf{n}) :

805
$$\mathcal{L}_s = \|\min (s_1, s_2, s_3)\|_1. \quad (17)$$

806

807 This loss constrains the minimum scaling factor component to approach zero.

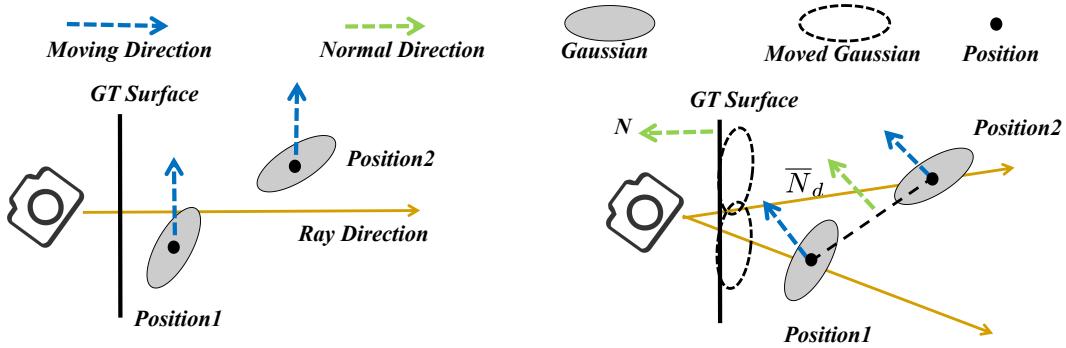


Figure A: Illustration of Proof of the Proposition on Comprehensive Update of Gaussian Parameters. (a) After back-propagation through alpha-blending Eq. 1, the rendered normal supervision loss \mathcal{L}_n moves Gaussians either closer to (corresponding to *Position*₁) or farther from (corresponding to *Position*₂) the intersecting ray. When the normal of a Gaussian is closer to the ground truth (GT) surface normal, this supervision mechanism pushes the Gaussian (e.g., *Position*₁) toward the ray to increase its weight in the rendering equation; conversely, if there is a significant deviation between the two normals, it pushes the Gaussian (e.g., *Position*₂) away from the ray. (b) In contrast, the D-Normal regularizer loss \mathcal{L}_{dn} can move Gaussians either closer to or farther from the GT surface. Here, *Position*₁ and *Position*₂ are 3D positions corresponding to the mean depth of two adjacent pixels (rays), computed via Eq. 5; the D-Normal \bar{N}_d is derived from *Position*₁ and *Position*₂ using Eq. 6. Notably, \mathcal{L}_{dn} relies on the intersection depth, related to Gaussian position *Position* and normal \mathbf{n}) to encourage \bar{N}_d alignment with the GT normal N , ultimately enabling Gaussians to move toward or away from the (GT) surface.

For the plane constraint, any point o_p on plane (\mathbf{p}, \mathbf{n}) satisfies $\mathbf{n} \cdot (o_p - \mathbf{p}) = 0$. For the ray representation, a ray originating from the origin is expressed as $o_l = \mathbf{r}t$, where \mathbf{r} denotes the ray direction and t is the distance from the origin. At the intersection ($o_l = o_p$), solving for the depth along the camera z -axis yields:

$$d(\mathbf{n}, \mathbf{p}) = r_z \cdot \frac{\mathbf{n} \cdot \mathbf{p}}{\mathbf{n} \cdot \mathbf{r}}, \quad (18)$$

where r_z is the z -component of \mathbf{r} .

This $d(\mathbf{n}, \mathbf{p})$ is correlated with both \mathbf{p} and \mathbf{n} , ensuring accuracy and enabling D-Normal regularization to backpropagate loss to Gaussian parameters.

B.2 PROOF PROPOSITIONS

Proposition 1.1 Supervising the rendered normals cannot effectively influence the positions of Gaussians. The rendered normal \hat{N} is defined as the opacity-weighted average of Gaussian normals. Considering the normal loss \mathcal{L}_n , its gradient with respect to the Gaussian position p_i can be expressed via the chain rule as:

$$\frac{\partial \mathcal{L}_n}{\partial p_i} = \frac{\partial \mathcal{L}_n}{\partial \hat{N}} \cdot \frac{\partial \hat{N}}{\partial p_i}. \quad (19)$$

Since each Gaussian normal n_i is determined solely by the rotation parameters, $\frac{\partial n_i}{\partial p_i} = 0$. Thus, the dependency of \hat{N} on p_i originates only from the opacity weights α_i :

$$\frac{\partial \hat{N}}{\partial p_i} = \frac{\partial \hat{N}}{\partial \alpha_i} \cdot \frac{\partial \alpha_i}{\partial G(x)} \cdot \frac{\partial G(x)}{\partial p_i}. \quad (20)$$

For a Gaussian distribution

$$G(x) = \exp\left(-\frac{1}{2}(x - p_i)^\top \Sigma^{-1}(x - p_i)\right). \quad (21)$$

864 we obtain
 865

866
 867
$$\frac{\partial G(x)}{\partial p_i} = -G(x) \Sigma^{-1}(x - p_i). \quad (22)$$

 868

869 In our implementation, following scale regularization, each Gaussian is flattened into an approxi-
 870 mate plane, so we approximate Σ^{-1} by the identity matrix to emphasize directionality. Hence,
 871

872
 873
$$\frac{\partial G(x)}{\partial p_i} \approx -G(x) (x - p_i). \quad (23)$$

 874

875 Substituting into Eq.(16), the resulting position gradient is
 876

877
 878
$$\frac{\partial \mathcal{L}_n}{\partial p_i} \propto (x - p_i). \quad (24)$$

 879

880 This indicates that the position update depends only on the spatial offset between the pixel-aligned
 881 point x and the Gaussian center p_i , without involving the surface normal n_i . Consequently, con-
 882 ventional normal supervision can only adjust opacities but fails to drive positions toward the true
 883 surface along its normal direction. This explains why rendered-normal supervision alone leads to
 884 incomplete geometric optimization.

885 **Proposition 1.2** Supervising our proposed Depth-Normal (D-Normal) regularizer can effectively
 886 influence the positions of Gaussians.

887 We now consider our proposed D-Normal loss \mathcal{L}_{dn} . By definition, the D-Normal \bar{N}_d is computed
 888 from the gradients of rendered depth maps. The gradient of \mathcal{L}_{dn} with respect to the Gaussian position
 889 p_i follows a three-stage chain rule:
 890

891
 892
$$\frac{\partial \mathcal{L}_{dn}}{\partial p_i} = \frac{\partial \mathcal{L}_{dn}}{\partial \bar{N}_d} \cdot \frac{\partial \bar{N}_d}{\partial \hat{D}} \cdot \frac{\partial \hat{D}}{\partial p_i}. \quad (25)$$

 893

894 where \hat{D} denotes the rendered depth.
 895

896 Since \hat{D} is the opacity-weighted average of Gaussian intersection depths d_i , its derivative can be
 897 decomposed as:
 898

899
 900
$$\frac{\partial \hat{D}}{\partial p_i} = \underbrace{\frac{\partial \hat{D}}{\partial \alpha_i} \cdot \frac{\partial \alpha_i}{\partial p_i}}_{\text{(A) Conventional weight term}} + \underbrace{\frac{\partial \hat{D}}{\partial d_i} \cdot \frac{\partial d_i}{\partial p_i}}_{\text{(B) Depth term (new)}}, \quad (26)$$

 901
 902

903 The depth of a Gaussian-ray intersection is given by:
 904

905
 906
$$d_i = r_z \cdot \frac{n_i \cdot p_i}{n_i \cdot r}, \quad (27)$$

 907

908 where r is the viewing ray and r_z its z -component. Differentiating with respect to p_i yields:
 909

910
 911
$$\frac{\partial d_i}{\partial p_i} = r_z \cdot \frac{n_i}{n_i \cdot r}. \quad (28)$$

 912

913 Thus, the second term (B) in $\partial \hat{D} / \partial p_i$ explicitly involves the surface normal n_i . we obtain:
 914

915
 916
$$\frac{\partial \mathcal{L}_{dn}}{\partial p_i} = \underbrace{\frac{\partial \mathcal{L}_{dn}}{\partial \bar{N}_d} \cdot \frac{\partial \bar{N}_d}{\partial \hat{D}} \cdot \frac{\partial \hat{D}}{\partial \alpha_i} \cdot \frac{\partial \alpha_i}{\partial G(x)} \cdot \frac{\partial G(x)}{\partial p_i}}_{\text{traditional weight-dependent term}} + \underbrace{\frac{\partial \mathcal{L}_{dn}}{\partial \bar{N}_d} \cdot \frac{\partial \bar{N}_d}{\partial \hat{D}} \cdot \frac{\partial \hat{D}}{\partial d_i} \cdot r_z \cdot \frac{n_i}{n_i \cdot r}}_{\text{new term proportional to } n_i}. \quad (29)$$

 917

The new term proportional to n_i provides a direct mechanism to update the position p_i along the normal direction. As a result, D-Normal supervision not only influences Gaussian rotations (as conventional normal supervision does) but also effectively aligns Gaussian positions with the underlying surface geometry. This theoretical insight explains the substantial geometric improvements observed in our experiments.

Proposition 2.1

Supervising our proposed Depth-Consistent D-Normal Regularizer, which incorporates the Pseudo Depth & D-Normal Dual Supervision Mechanism by utilizing both D-Normal maps and pseudo depth maps, can effectively and stably influence Gaussian positions along the normal direction, thereby achieving comprehensive updates of geometric parameters (rotation and position) and significantly improving geometric and reconstruction accuracy.

From Proposition 1.1, conventional rendered normal supervision provides gradients $\frac{\partial \mathcal{L}_n}{\partial p_i} \propto (x - p_i)$, which are independent of Gaussian normals n_i and thus fail to guide positions toward the true surface.

From Proposition 1.2, D-Normal supervision introduces an additional term

$$\frac{\partial \mathcal{L}_{dn}}{\partial p_i} \supset \frac{\partial \mathcal{L}_{dn}}{\partial \bar{N}_d} \cdot \frac{\partial \bar{N}_d}{\partial \hat{D}} \cdot \frac{\partial \hat{D}}{\partial d_i} \cdot r_z \cdot \frac{n_i}{n_i \cdot r}, \quad (30)$$

which is explicitly proportional to the normal n_i . This enables position updates along the surface normal direction, thus coupling position and rotation optimization.

However, the reliability of this update depends on the accuracy of rendered depth \hat{D} . To further enhance stability, we introduce pseudo depth supervision $\mathcal{L}_{pD}(\hat{D}, D^{\text{pseudo}})$. Its gradient contributes

$$\frac{\partial \mathcal{L}_{pD}}{\partial p_i} = \frac{\partial \mathcal{L}_{pD}}{\partial \hat{D}} \cdot \frac{\partial \hat{D}}{\partial p_i}, \quad (31)$$

which shares the same structural dependence on $\frac{\partial \hat{D}}{\partial p_i}$ as the D-Normal term, and therefore reinforces the normal-dependent component introduced above.

Combining these two complementary signals, the dual supervision mechanism (i) stabilizes depth estimation via pseudo depth, and (ii) ensures normal-consistent position updates via D-Normal. As a result, both rotation and position parameters of Gaussians are comprehensively optimized, yielding improved geometric accuracy in reconstruction.

C SUPPLEMENTATION TO THE PARTITIONING STRATEGY

Existing large-scale 3DGS frameworks exhibit two critical limitations: geometric discontinuities at block boundaries introduce visible fusion artifacts, while redundant Gaussians attract non-contributing views that inflate computational loads during block-wise training. Our unified approach addresses both issues through integrated redundancy reduction and geometric continuity enforcement.

The pipeline begins with global pruning of low-impact Gaussians using spatially adaptive scoring (Eq. 15):

$$\mathbf{G}^{\text{pruned}} = G_k \in \mathbf{G} \mid S_k > \theta_{\text{prune}}, \quad (32)$$

This operation targets background and low-contribution primitives that attract irrelevant views. Pre-partition pruning eliminates redundancy propagation to local blocks, significantly reducing computational load.

To enable spatially balanced partitioning, the pruned Gaussians are contracted into a normalized cube $[-1, 1]^3$ using a hybrid contraction function (as in (Wu et al., 2023)), which applies a linear mapping to the foreground region and a nonlinear scaling to the unbounded background. This contraction yields a compact representation of the full scene and facilitates uniform space division.

Within this contracted space, the scene is partitioned into regular blocks. To preserve geometric continuity across adjacent partitions, boundary Gaussians are explicitly duplicated:

$$\mathbf{G}_{\text{shared}}^j = \{G_k \mid \text{dist}(G_k, \partial \mathcal{B}_j) < \delta_{\text{share}}\}. \quad (33)$$

972 This duplication enforces overlapping geometric constraints near block interfaces, thereby suppressing
 973 boundary artifacts during block-wise fusion.
 974

975 Camera pose assignment for each block \mathcal{B}_j integrates geometric proximity and perceptual contribution
 976 through dual evaluation criteria. The geometric criterion assesses physical containment by
 977 checking if the contracted camera position $\mathbf{p}_{\tau_i}^{\text{ctr}} = \text{contract}(\hat{\mathbf{p}}_{\tau_i})$ falls within the block's spatial
 978 extent $[\mathbf{b}_{j,\min}, \mathbf{b}_{j,\max}]$, formalized as:
 979

$$B_{\text{geo}}(\tau_i) = \begin{cases} 1 & \mathbf{p}_{\tau_i}^{\text{ctr}} \in [\mathbf{b}_{j,\min}, \mathbf{b}_{j,\max}] \\ 0 & \text{otherwise} \end{cases}, \quad (34)$$

981 where the contraction operator follows (Liu et al., 2024a).
 982

983 The perceptual criterion quantifies visual degradation when removing Gaussians $\mathbf{G}_{\mathcal{B}_j}$. By comparing
 984 renders $I_{\tau_i}^{\text{full}}$ (full model) and $I_{\tau_i}^{\text{excl-}j}$ (excluding $\mathbf{G}_{\mathcal{B}_j}$) in original space, it computes:
 985

$$B_{\text{vis}}(\tau_i) = \begin{cases} 1 & \text{SSIM}(I_{\tau_i}^{\text{full}}, I_{\tau_i}^{\text{excl-}j}) < 1 - \varepsilon_j \\ 0 & \text{otherwise} \end{cases}, \quad (35)$$

988 with ε_j controlling sensitivity to structural loss, identifying perceptually dependent poses.
 989

990 The final assignment combines both criteria:
 991

$$B(\tau_i) = B_{\text{geo}}(\tau_i) \vee B_{\text{vis}}(\tau_i), \quad (36)$$

992 ensuring each pose is assigned to blocks it physically occupies or visually relies upon. This establishes
 993 efficient view-block correspondence while maintaining rendering consistency.
 994

995 Table A: Comparison of training times across multiple state-of-the-art methods on the Mill-19 (Turki
 996 et al., 2022c) and UrbanScene3D (Lin et al., 2022). **Bold** indicates best performance..
 997

Models	Building	Rubble	Residence	Sci-Art
	Time ↓	Time ↓	Time ↓	Time ↓
Mega-NeRF	19:49	30:48	27:20	27:39
Switch-NeRF	24:46	38:30	35:11	34:34
VastGS †	03:26	02:30	03:12	03:13
DOGS	03:51	02:25	04:33	04:23
CityGS-v2	04:25	03:05	04:45	04:38
Ours	03:13	02:10	02:45	03:40

1008 Table B: Novel View Synthesis Performance Evaluation on the GauU-Scene datasets (Xiong et al.,
 1009 2024). **Bold** indicates best performance.
 1010

Methods	Residence			Russian Building			Modern Building		
	SSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	PSNR ↑	LPIPS ↓
NeuS	0.244	15.16	0.674	0.202	13.65	0.694	0.236	14.58	0.694
Neuralangelo	NaN	NaN	NaN	0.328	12.48	0.698	NaN	NaN	NaN
SuGaR	0.612	21.95	0.452	0.738	23.62	0.332	0.700	24.92	0.381
GOF	0.652	20.68	0.391	0.713	21.30	0.322	0.749	25.01	0.286
VCR-Gaus	0.663	22.69	0.404	0.724	22.89	0.273	0.726	25.19	0.230
2DGS	0.703	22.24	0.306	0.788	23.77	0.189	0.776	25.77	0.202
CityGS-v2	0.742	23.57	0.243	0.784	24.12	0.196	0.770	25.84	0.207
Ours	0.762	23.78	0.206	0.810	24.53	0.158	0.805	26.44	0.157

1021
 1022 **D MORE EXPERIMENTS**
 1023

1024 The experimental section of this paper focuses on evaluating the performance of UrbanGS in large-
 1025 scale scene reconstruction. Through comprehensive comparisons with a variety of baseline methods,
 1026

1026 Table C: Detailed geometry evaluation on GauU-Scene datasets (Xiong et al., 2024). “NaN” indicates
 1027 invalid numerical results, while “FAIL” denotes failure to extract valid mesh. For all metrics,
 1028 ↑ indicates higher values are better.

Methods	Campus			Village			College		
	P ↑	R ↑	F1 ↑	P ↑	R ↑	F1 ↑	P ↑	R ↑	F1 ↑
NeuS	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL
Neuralangelo	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN
SuGaR	0.321	0.272	0.294	0.354	0.253	0.295	0.409	0.271	0.326
VCR-Gaus	0.478	0.312	0.379	0.492	0.412	0.448	0.456	0.361	0.401
2DGS	0.389	0.304	0.341	0.442	0.283	0.345	0.340	0.182	0.237
GOF	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL
PGSR	0.464	0.355	0.403	0.535	0.445	0.486	0.349	0.349	0.354
CityGS-X	0.505	0.361	0.421	0.545	0.443	0.489	0.559	0.371	0.446
CityGaussianV2	0.486	0.383	0.428	0.580	0.503	0.543	0.577	0.373	0.453
Ours	0.492	0.388	0.435	0.567	0.512	0.538	0.564	0.381	0.456

Table D: Ablation study of different priors on the Modern Building dataset.

Priors				Rendering Quality			Geometric Quality		
Dav2	MiDaS	Dsine	GeoWizard	SSIM ↑	PSNR ↑	LPIPS ↓	P ↑	R ↑	F1 ↑
✓		✓		0.805	26.44	0.157	0.663	0.404	0.503
✓			✓	0.802	26.48	0.163	0.665	0.399	0.498
	✓	✓		0.798	26.33	0.161	0.645	0.410	0.501
	✓		✓	0.785	26.12	0.166	0.658	0.392	0.491

1055 the effectiveness of UrbanGS is validated across three key aspects: training efficiency, novel view
 1056 synthesis quality, and geometric accuracy.

D.1 TRAINING EFFICIENCY

1060 As shown in table A, in training time comparison experiments conducted on diverse scenes such as
 1061 Building, Rubble, Residence, and Sci-Art, UrbanGS demonstrates significant efficiency gains. For
 1062 instance, in the Building scene (Turki et al., 2022c), UrbanGS completes training in only 3 hours and
 1063 13 minutes, substantially outperforming Mega-NeRF (19 hours 49 minutes) (Turki et al., 2022b) and
 1064 Switch-NeRF (24 hours 46 minutes) (Mi & Xu, 2023). Even when compared with more efficient
 1065 baselines such as VastGS† (Lin et al., 2024), UrbanGS consistently achieves competitive or superior
 1066 training times across most scenes.

1067 Table E: Effect of block partitioning on the *Russian* dataset. Memory (Mem) in GB, Time in minutes. **Bold**
 1068 indicates best performance.

Block/GPU	PSNR↑	SSIM↑	LPIPS↓	F1↑	Mem↓	Time↓
2/2	23.43	0.779	0.215	0.518	25.2	170
4/4	24.55	0.804	0.201	0.539	20.1	140
8/8	24.66	0.813	0.184	0.546	14.4	122

D.2 NOVEL VIEW SYNTHESIS

1077 As shown in table B, the novel view synthesis performance is also evaluated on the GauU-Scene
 1078 dataset (Xiong et al., 2024), UrbanGS outperforms competing methods in all tested scenes using
 1079 SSIM, PSNR (higher is better), and LPIPS (lower is better) as evaluation metrics. Specifically,
 in the Residence scene, it achieves SSIM 0.762, PSNR 23.78, and LPIPS 0.206; in the Russian

1080 Building scene, SSIM 0.810, PSNR 24.53, and LPIPS 0.158; and in the Modern Building scene,
 1081 SSIM 0.805, PSNR 26.44, and LPIPS 0.157. These results consistently surpass baselines such as
 1082 SuGaR (Guédon & Lepetit, 2024) and GOF (Yu et al., 2024b).
 1083

1084 D.3 GEOMETRIC ACCURACY

1085 To further assess the generalization of our method to large-scale urban scenes, we additionally eval-
 1086 uate geometry quality on the GauU-Scene dataset (Xiong et al., 2024). As summarized in Table C,
 1087 we report precision (P), recall (R), and F1 score across three subsets (Campus, Village, and College).
 1088 Several NeRF- and 3DGS-based baselines either produce invalid numerical results (“NaN”) or fail to
 1089 extract a valid mesh (“FAIL”), highlighting the difficulty of this benchmark. In contrast, our method
 1090 consistently reconstructs valid meshes and achieves the best or highly competitive performance on
 1091 all metrics. As a qualitative complement to the numerical results, Figure B and Figure C compares
 1092 the reconstructed meshes of representative methods on multiple GauU-Scene subsets. Baseline ap-
 1093 proaches often suffer from over-smoothed surfaces, broken structures, or missing fine-scale details,
 1094 particularly around building facades and road layouts, whereas our method produces more complete
 1095 and coherent geometry with sharper boundaries.
 1096

1097 D.4 MORE ABLATIONS

1098 As shown in Figure D, we conducted ablation studies on the Campus dataset. When the Depth-
 1099 Consistent D-Normal Regularization module and Partitioning Strategy module are ablated, signif-
 1100 icant differences are observed in both rendered normal maps and rendered depth maps compared
 1101 with our full method, demonstrating the substantial effectiveness of these modules in the proposed
 1102 approach. To more clearly demonstrate the effectiveness of our module in pushing 3D points along
 1103 the normal direction, we conduct an experiment on a small scene. As shown in Figure E, when the
 1104 Depth-Consistent D-Normal Regularization module is removed, the point cloud on object surfaces
 1105 becomes highly scattered. With our regularizer enabled, the points are driven toward the underlying
 1106 surfaces, resulting in a much more compact point cloud and a significantly cleaner reconstruction.
 1107

1108
 1109 Table D studies the influence of using different depth estimators (Dav2 (Hu et al., 2024), MiDaS
 1110 (Ranftl et al., 2020)) and the normal prior (GeoWizard (Fu et al., 2024), Dsine (Bae & Davison,
 1111 2024)) on the Modern Building dataset. Across all combinations, the rendering metrics remain very
 1112 close (SSIM \approx 0.79–0.81, PSNR \approx 26.1–26.5, LPIPS \approx 0.157–0.166) and the geometric quality
 1113 (P/R/F1) only fluctuates within a small range. This indicates that our framework is not sensitive
 1114 to the particular choice of depth or normal prior. Thanks to the explicit depth-consistency and
 1115 normal-consistency constraints, the optimization can effectively correct the bias of different priors
 1116 and consistently recover high-quality geometry.
 1117

1118 As shown in Table E, under the experimental scenario of the Russian dataset, this table presents
 1119 the impacts of different “Block/GPU count” configurations on model performance, memory con-
 1120 sumption, and training time. As the configuration is scaled up from 2/2 (2 blocks, 2 GPUs) to 8/8
 1121 (8 blocks, 8 GPUs), the model performance is gradually optimized: PSNR increases from 23.43
 1122 to 24.66, SSIM rises from 0.779 to 0.813, LPIPS decreases from 0.215 to 0.184, and F1 improves
 1123 from 0.518 to 0.546. Meanwhile, resource consumption is significantly reduced, with memory usage
 1124 dropping from 25.2 GB to 14.4 GB and training time shortening from 170 minutes to 122 minutes.
 1125 This demonstrates the positive role of the parallel training strategy—where the number of blocks
 1126 matches the number of GPUs—in balancing “performance-efficiency”.

1127 As shown in Table F, through systematic weight ablation experiments, we identified the optimal
 1128 configuration ($\alpha=1.2$, $\beta=1.0$, $\gamma=0.8$), which serves as the default setting for all subsequent experi-
 1129 ments. This configuration achieves the best results across all metrics - PSNR (26.44), SSIM (0.805),
 1130 LPIPS (0.157), and F1-score (0.503) - demonstrating its comprehensive advantages in both ren-
 1131 dering quality and geometric accuracy. Compared to the equal-weight baseline (1.0,1.0,1.0), our
 1132 configuration improves PSNR and F1-score by 0.55 dB and 0.016, respectively. Univariate analysis
 1133 further validates the design rationale: the ray intersection frequency weight (α) is crucial for multi-
 view consistency (14.1% F1-score drop when $\alpha=0.0$), the opacity weight (β) directly affects visual
 quality (26.1% LPIPS increase when $\beta=0.0$), while the volume weight ($\gamma=0.8$) preserves details

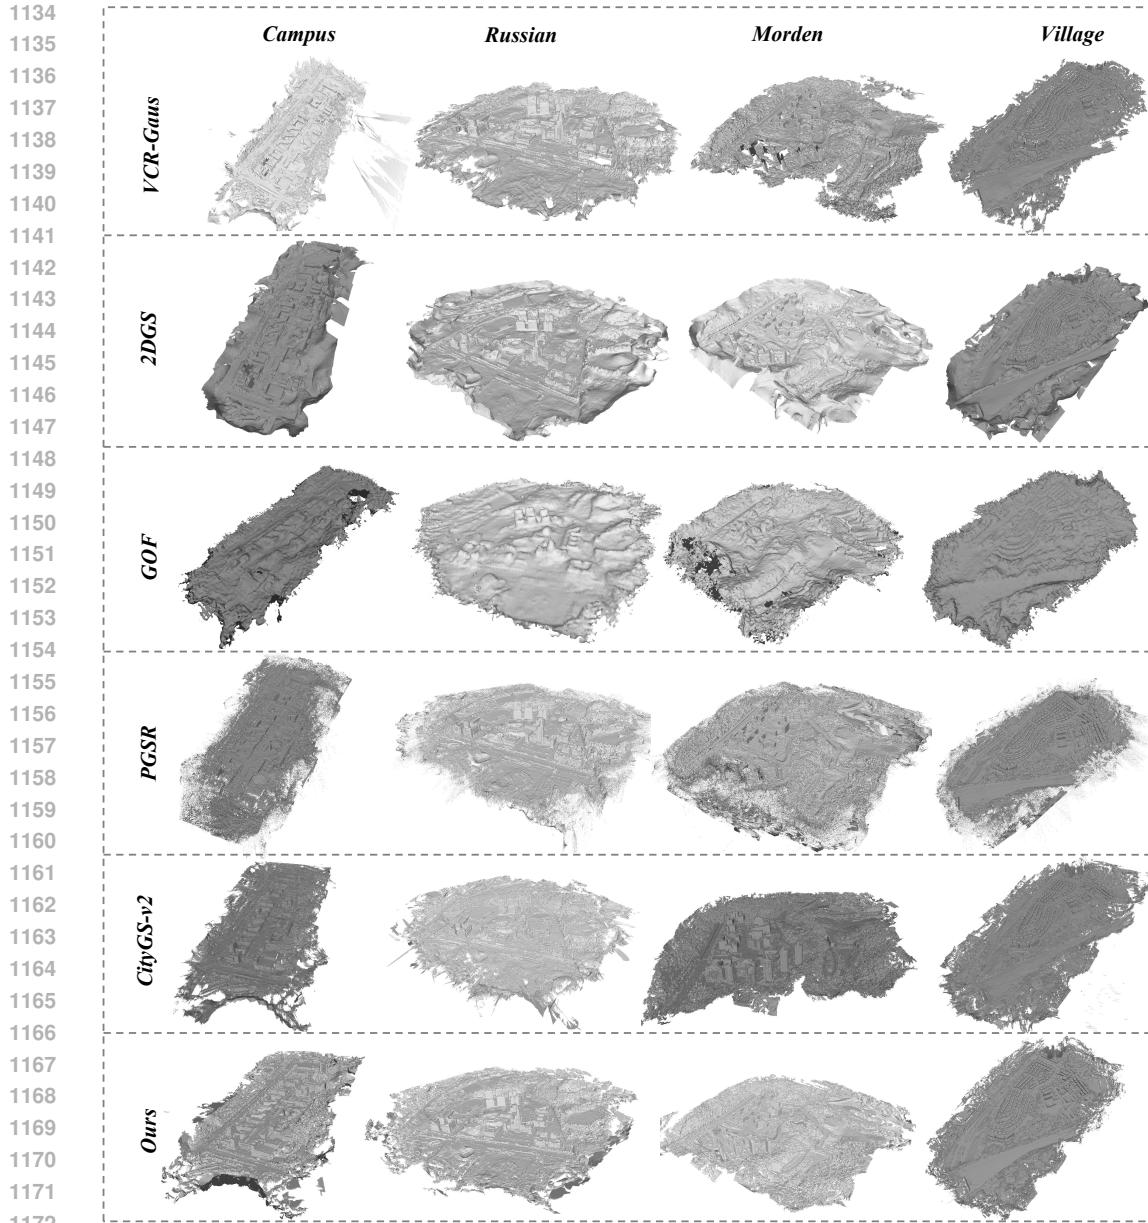
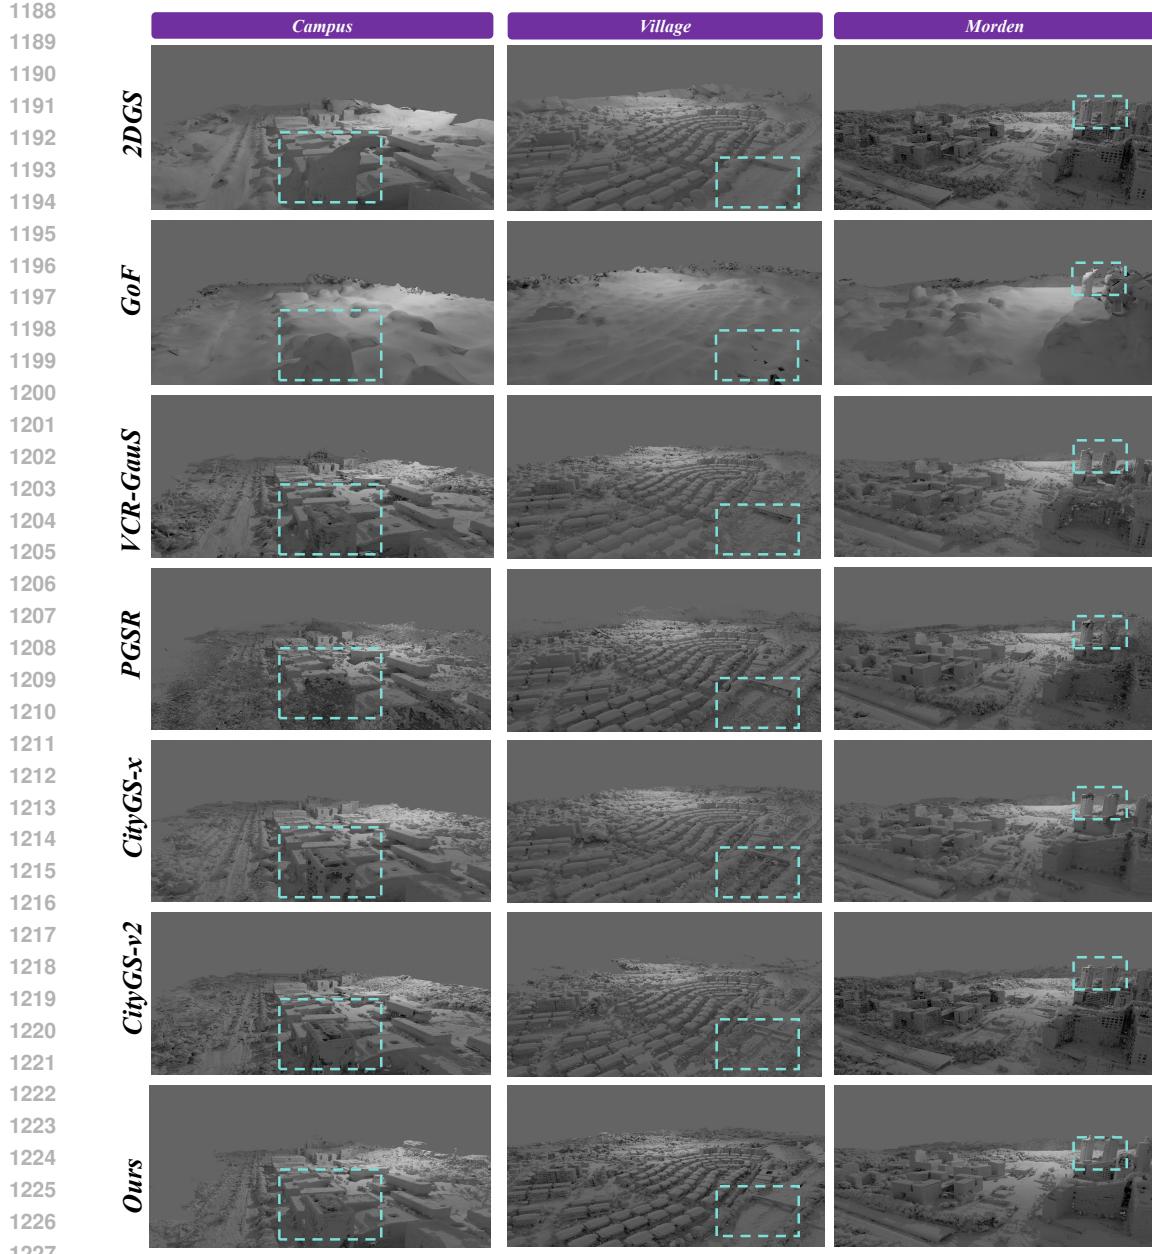


Figure B: Visual comparison of meshes from state-of-the-art (SOTA) methods.

while avoiding excessive redundancy (performance degradation across all metrics when $\gamma=1.0$). This weight combination thus achieves the optimal balance between quality and efficiency. In addition, Figure F visualizes rendered views under different weight combinations. The reconstructions are visually very similar across settings, with no catastrophic degradation, even for suboptimal weights. This qualitative evidence further confirms that our method is robust and only weakly sensitive to the choice of (α, β, γ) , as long as they remain within a reasonable range.

Table G reports the effect of removing individual components in our block partition strategy on the Russian scene of the GauU-Scene dataset (Xiong et al., 2024). Starting from our full model, discarding the global pruning term in Eq. 32 leads to a clear increase in the number of Gaussians (2.45M \rightarrow 3.01M), longer training time and higher memory usage, together with a slight drop in both rendering and geometric quality. This confirms that the spatially adaptive pruning not only reduces redundancy but also facilitates optimization. Removing the boundary-duplication rule in Eq. 33 also degrades SSIM, LPIPS, and F1, indicating that sharing Gaussians across neighboring



1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

Figure C: Qualitative mesh and texture comparison between SOTA and our method on the Campus, Village, and Morden Buliding scenes (Xiong et al., 2024).

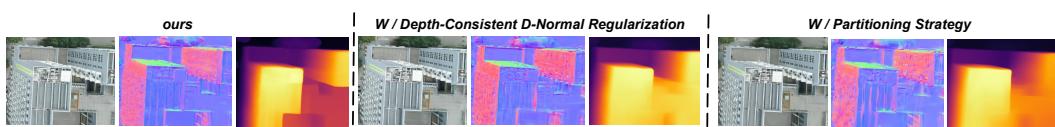


Figure D: Ablation experiments on the Campus Dataset Lin et al. (2022)

blocks is important for suppressing block-boundary artifacts. When the geometric pose-assignment criterion in Eq. 34 is disabled, the performance drops most significantly (e.g., PSNR and F1 both decrease considerably) while the computational cost increases, showing that aligning camera frusta with physically relevant blocks is crucial for both fidelity and efficiency. Finally, removing the perceptual criterion in Eq. 35 causes a moderate decline in rendering quality and F1, demonstrating that perceptual filtering helps retain poses that are visually important for each block. Taken together,

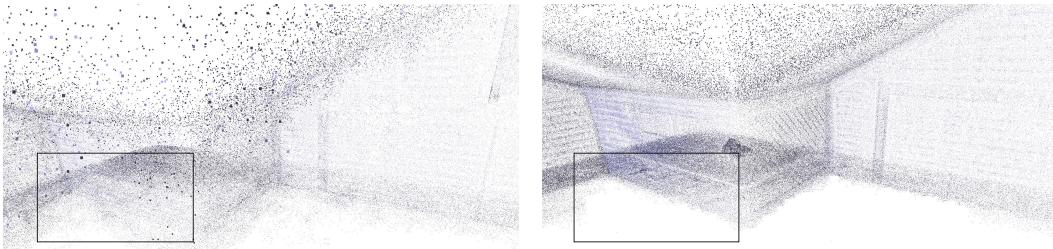


Figure E: Qualitative ablation for the Depth-Consistent D-Normal Regularizer. We visualized the centers of Gaussian ellipsoids in a 3D scene. In the left figure, the Depth-Consistent D-Normal Regularizer is disabled, while the right figure demonstrates the results with our proposed regularization. In comparison, the left figure exhibits a notable number of Gaussian ellipsoids floating off the surface. Our proposed Depth-Consistent D-Normal Regularizer effectively pushes the 3D Gaussians toward the surface, resulting in a cleaner reconstruction.

Table F: Quantitative Analysis of Weight Configuration Ablation Study

Weight Configuration	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	F1 \uparrow
(1.2, 1.0, 0.8)	26.44	0.805	0.157	0.503
(1.0, 1.0, 1.0)	26.19	0.791	0.172	0.487
(0.0, 1.0, 1.0)	24.32	0.763	0.215	0.432
(1.0, 0.0, 1.0)	25.12	0.778	0.198	0.468
(1.0, 1.0, 0.0)	25.87	0.793	0.169	0.485
(1.2, 1.0, 1.0)	25.95	0.782	0.185	0.451

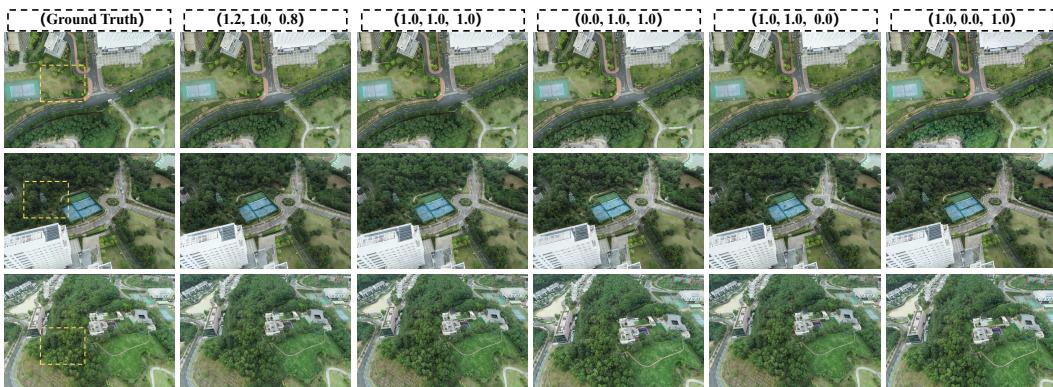


Figure F: Ablation experiments on the Morden Building Dataset Xiong et al. (2024)

these results show that all components of our partition strategy contribute to the overall trade-off, and the full design achieves the best balance between reconstruction quality and resource consumption.

As shown in Figure G, on the Sci-Art scenes (Lin et al., 2022) we observe that 3DGS-based methods with explicit geometry optimization often yield lower rendering quality than the original 3DGS. These scenes contain many aerial-style images dominated by distant sky regions with weak or ambiguous geometry. In such backgrounds, geometry-optimized variants tend to degrade the sky appearance, producing coarse color blotches and unnatural boundaries in the rendered views. In contrast, although the original 3DGS is also imperfect in sky modeling, its results still vaguely preserve cloud layers and building silhouettes. This discrepancy highlights a limitation of current geometry optimization objectives when applied to background regions lacking clear geometric structure.

Table H presents the ablation study results of the two key hyperparameters γ_d and τ in the geometry-aware confidence mechanism, conducted on the Modern Building (Lin et al., 2022). scene. The baseline configuration ($\gamma_d = 0.1$, $\tau = 0.01$) achieves the optimal performance across all evaluation metrics, with PSNR of 26.44 and F1-score of 0.503. Decreasing γ_d to 0.05 alone leads to noticeable performance degradation (PSNR drops by 0.32, F1-score drops by 0.018), indicating that excessive sensitivity to depth gradient consistency suppresses valid geometric signals. Similarly, increasing τ

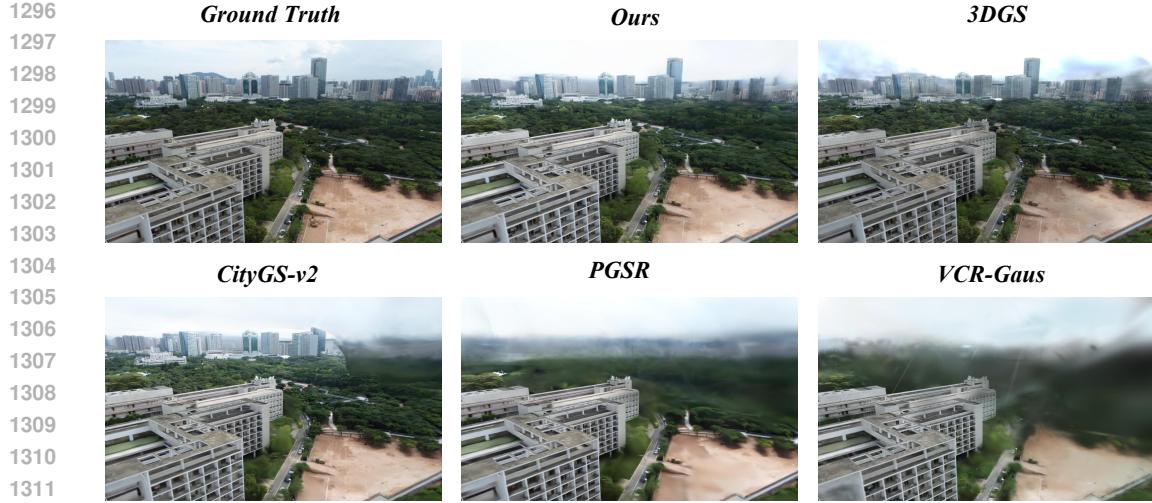


Figure G: Qualitative mesh and texture comparison between SOTA and our method on Art-Sci Scene (Lin et al., 2022).

to 0.02 alone also causes performance deterioration (PSNR drops by 0.16, F1-score drops by 0.011), suggesting insufficient suppression of depth errors adversely affects reconstruction quality. The worst performance occurs when both parameters are modified ($\gamma_d = 0.15, \tau = 0.005$), with PSNR and F1-score decreasing by 0.55 and 0.025 respectively, validating the coupling relationship between the two hyperparameters and the rationality of the baseline selection. These results comprehensively demonstrate that our chosen hyperparameter combination achieves the optimal balance between geometric consistency and error suppression.

Table G: Ablation Results of Block Partition Strategy on Russian Scene Dataset (Xiong et al., 2024). **Bold** indicates best performance.

Method	Rendering Quality			Geometric Quality			Training Statistics			
	PSNR↑	SSIM↑	LPIPS↓	P↑	R↑	F1↑	GS (millions)↓	Time (min)↓	Size (MB)↓	Mem (GB)↓
Effect of Removing Individual Components										
baseline (ours)	24.66	0.813	0.184	0.568	0.525	0.546	2.45	122	314.24	14.4
baseline w/o Eq. 32	24.43	0.797	0.201	0.562	0.518	0.539	3.01	142	429.41	17.5
baseline w/o Eq. 33	24.51	0.802	0.198	0.564	0.513	0.537	2.44	129	314.24	15.1
baseline w/o Eq. 34	22.32	0.764	0.231	0.531	0.498	0.513	2.56	157	334.31	20.3
baseline w/o Eq. 35	24.42	0.808	0.188	0.566	0.521	0.543	2.46	125	314.31	14.7

Table H: Ablation Study of Geometry-Aware Confidence Hyperparameters

Hyperparameter Settings	PSNR ↑	F1-score ↑
Baseline ($\gamma_d = 0.1, \tau = 0.01$)	26.44	0.503
Only $\gamma_d = 0.05$	26.12	0.485
Only $\tau = 0.02$	26.28	0.492
Both modified ($\gamma_d = 0.15, \tau = 0.005$)	25.89	0.478

E MULTI-VIEW GEOMETRIC CONSISTENCY

Our method employs multi-view geometric consistency principles to conduct quantitative assessment of depth map quality. By back-projecting the depth map of a reference view to generate a 3D point cloud, we transform this point cloud into the coordinate system of adjacent views using camera poses and perform reprojection. The relative error between the reprojected depth values and the ac-

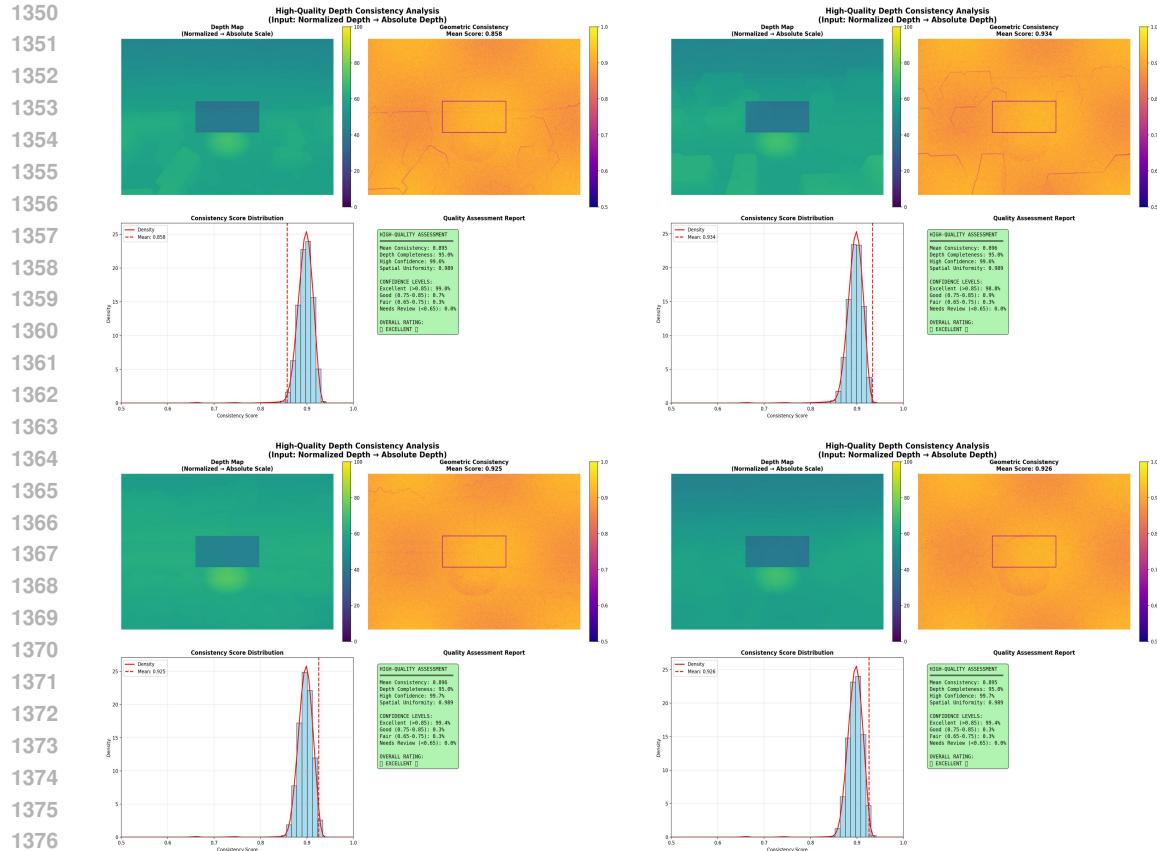


Figure H: Experimental results on the Morden Building (Xiong et al., 2024).

tual depth values in the target view is then computed to verify the geometric correctness of the depth estimation. The evaluation process utilizes absolute depth values (in meters), dynamically adapts to scale variations across depth maps of different resolutions and camera parameters, and applies a 10% relative error threshold to determine consistency. The final output is a normalized consistency score in the range of 0-1, where scores above 0.8 indicate high consistency across multiple views, 0.6-0.8 represents good consistency, 0.4-0.6 indicates moderate consistency, and scores below 0.4 suggest significant geometric inconsistencies. This approach provides a reliable geometric verification metric for evaluating the performance of depth estimation models. As shown in Table I, our evaluation of fifty depth maps demonstrates strong geometric consistency, with an average score of 0.87 confirming reliable depth estimation across multiple viewpoints. The 83% consistency pass rate and 78% check coverage validate the robustness of our approach for 3D reconstruction applications. Furthermore, as illustrated in Figure H, we present the depth consistency test results for four sample images.

Table I: Depth Map Geometric Consistency Evaluation Results

Evaluation Metric	Value	Quality Level
Average Consistency Score	0.87 ± 0.08	Excellent
Check Coverage	78%	Good
Consistency Pass Rate	83%	Excellent

F LIMITATIONS

Although UrbanGS demonstrates advantages in large-scale reconstruction, it still exhibits certain limitations. Its geometric regularization relies on monocular depth/normal priors derived from pre-trained networks, which may propagate estimation errors into the reconstruction—particularly in

1404 regions with weak textures or extreme lighting conditions. Additionally, the method primarily fo-
1405 cuses on static environments and does not explicitly model dynamic objects commonly found in
1406 urban scenes. Future work will aim to mitigate dependency on monocular priors through multi-
1407 view geometric consensus and extend the framework to dynamic urban objects via explicit motion
1408 modeling.
1409

1410 G USE OF LARGE LANGUAGE MODELS 1411

1412 A large language model (LLM) was used solely for language-level assistance, such as improving
1413 readability, fluency of the text and formatting \LaTeX tables and retrieve related works. The research
1414 ideas, experiments, and results are entirely the work of the authors, who bear full responsibility for
1415 the content of this submission.
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457