
D²GS: Dense Depth Regularization for LiDAR-free Urban Scene Reconstruction

Kejing Xia^{1*} Jidong Jia^{2*} Ke Jin³ Yucai Bai⁴ Li Sun⁴ Dacheng Tao⁵ Youjian Zhang^{4†}

¹Wuhan University ²Shanghai Jiaotong University ³Tongji University ⁴Bosch

⁵Nanyang Technological University

xiakejing.lesia@whu.edu.cn, jjd1123@sjtu.edu.cn, 2330379@tongji.edu.cn
dacheng.tao@ntu.edu.sg, {Yucai.BAI, Kevin.SUN, Youjian.ZHANG}@cn.bosch.com

Abstract

Recently, Gaussian Splatting (GS) has shown great potential for urban scene reconstruction in the field of autonomous driving. However, current urban scene reconstruction methods often depend on multimodal sensors as inputs, *i.e.* LiDAR and images. Though the geometry prior provided by LiDAR point clouds can largely mitigate ill-posedness in reconstruction, acquiring such accurate LiDAR data is still challenging in practice: i) precise spatiotemporal calibration between LiDAR and other sensors is required, as they may not capture data simultaneously; ii) reprojection errors arise from spatial misalignment when LiDAR and cameras are mounted at different locations. To avoid the difficulty of acquiring accurate LiDAR depth, we propose D²GS, a LiDAR-free urban scene reconstruction framework. In this work, we obtain geometry priors that are as effective as LiDAR while being denser and more accurate. **First**, we initialize a dense point cloud by back-projecting multi-view metric depth predictions. This point cloud is then optimized by a Progressive Pruning strategy to improve the global consistency. **Second**, we jointly refine Gaussian geometry and predicted dense metric depth via a Depth Enhancer. Specifically, we leverage diffusion priors from a depth foundation model to enhance the depth maps rendered by Gaussians. In turn, the enhanced depths provide stronger geometric constraints during Gaussian training. **Finally**, we improve the accuracy of ground geometry by constraining the shape and normal attributes of Gaussians within road regions. Extensive experiments on the Waymo dataset demonstrate that our method consistently outperforms state-of-the-art methods, producing more accurate geometry even when compared with those using ground-truth LiDAR data.

1 Introduction

Modeling urban street scenes is fundamental to autonomous driving [2]. Accurate 3D reconstructions of roads, buildings, and street furniture enable the creation of high-fidelity virtual environments, which are crucial for the closed-loop simulation and testing of perception, planning, and control systems [3]. The ability to reuse and edit these digital urban scenes, rather than repeatedly collecting fresh data for every test scenario, significantly improves the efficiency of data collection.

Recently, methods leveraging Neural Radiance Field (NeRF) [4] and 3D Gaussian Splatting (3DGS) [5] have become popular in this field for their photorealistic results. However, NeRF-based methods are computationally intensive while 3DGS is favored due to its explicit representation, high

*This work was completed during an internship at Bosch. Equal contribution.

†Corresponding author.



Figure 1: Examples of common LiDAR acquisition issues: a) reprojection error, b) calibration misalignments, and c) LiDAR missing problem. OmniRe [1] yields poor performance in both image reconstruction and depth estimation when using these inaccurate LiDAR measurements.

efficiency, and rapid rendering speed [6]. In fact, 3DGS has demonstrated its powerful reconstruction ability for dynamic urban scenes [7, 8, 1, 9, 10, 11].

Most of the 3DGS urban street scene reconstruction methods rely on LiDAR data to provide an accurate geometry prior. Usually, this prior can be applied in two aspects: i) LiDAR point cloud can be utilized to initialize Gaussian points, ii) LiDAR point cloud can provide a sparse depth supervision during the training phase. However, there are several notable drawbacks when acquiring LiDAR data in practice. First, the acquisition of LiDAR data requires specialized vehicles and costly equipment [12], which makes the data collection process challenging and less scalable. Second, the calibration between sensors (LiDARs and cameras) is highly demanding. Since these sensors do not capture data simultaneously, both spatial and temporal calibration are required to accurately align the LiDAR points with images. Furthermore, in practical data collection systems, LiDAR sensors and cameras are inevitably mounted at different positions and heights, causing significant viewpoint disparities and resulting in reprojection errors when LiDAR points are projected onto images (Fig. 1). However, most existing methods use this LiDAR projection result as ground-truth depth map. In preliminary experiments, we found that misalignment between LiDAR projections and pixels significantly affects the reconstruction results (as shown in Fig. 1). Finally, LiDAR provides only sparse depth supervision during the GS training phase, limiting the model’s ability to learn fine-grained geometric details.

In the meantime, many recent studies [13, 14, 15, 16, 17, 18, 19, 20, 21] have also explored ways to incorporate geometric constraints into the optimization of sparse-view scene reconstruction. However, these methods either rely on monocular depth estimation [13, 14, 15, 16, 17, 18], which only provides relative depth, or obtain metric depth through multi-view depth estimation [19, 20, 21], which is unsuitable for dynamic scenes and still suffers from large absolute errors.

To overcome the challenges of acquiring accurate geometry prior using LiDAR points, while obtaining a more robust and dense depth for regularization, we introduce D²GS – a LiDAR-free dynamic urban street scene reconstruction framework. First, to obtain an initial point cloud as a replacement for LiDAR point clouds, we leverage depth predictions from a pre-trained multi-view estimation model [21]. Re-projecting these depths yields an initial, dense point cloud. However, this point cloud is per-pixel, resulting in an excessive number of points. To manage the computational cost, we introduce a progressive pruning strategy. This strategy starts training the Gaussian field with minimal scene graph components, then gradually prunes the less important Gaussians to retain a compact yet representative point set that effectively captures the global scene geometry. Second, for the absence of

depth supervision provided by LiDAR data, we propose a novel joint optimization strategy for depth and the Gaussian representation. Inspired by incorporating image-based diffusion priors [22] used in Gaussian training, we introduce a Depth Enhancer module that leverages priors from the depth foundation model [23] to refine the depth rendered from the current Gaussian state. Specifically, we iteratively refine and update the densely rendered depth by using the current geometric information as guidance within the diffusion framework, and using refined depth for Gaussian training. The refined depth maps can offer a dense and accurate depth supervision that is semantically consistent with the scene, making them a superior alternative to LiDAR point clouds. Finally, since the ground plane possesses strong geometric priors in urban scenes [10], we initialize and optimize a dedicated road node in the scene graph for a better geometry and depth estimation of the road. Specifically, we constrain each Gaussian in the road node to be a planar parallel to a precomputed ground plane by regularizing their shape and normal attributes. In summary, our main contributions are as follows:

- We introduce D²GS, a framework for reconstructing dynamic urban street scenes using only camera inputs, eliminating the need for LiDAR points acquisition and the error brought by calibration and reprojection.
- We propose a progressive pruning strategy to efficiently manage the dense point cloud derived from initial depth estimates, yielding a compact Gaussian representation that captures global scene geometry.
- We develop a joint optimization strategy using a diffusion-based Depth Enhancer that iteratively refines estimated depths and the Gaussian representation, providing robust geometric supervision with dense metric depths.
- We further enhance the reconstruction accuracy on roads by creating a Road Node in the scene graph, explicitly modeling the ground plane using strong geometric priors.

2 Related Works

Urban scene reconstruction. Neural Radiance Field (NeRF) [4] and 3D Gaussian Splatting (3DGS) [5] have recently become the leading solutions for urban scene reconstruction. Early methods reconstruct entire city blocks and improve scalability to larger environments [24, 25], but all assumed static scenes. To handle dynamic elements, methods such as SUDS [26] and EmerN-erF [27] employ self-supervised decompositions into static and dynamic components, while neural scene-graph approaches explicitly disentangle motion and structure [28] and then leverage instance-aware frameworks to decompose each instance and model each moving object independently [29]. Nowadays, S3GS [7] introduces deformable fields for Gaussians to capture scene dynamics, and PVG [8] assigns each Gaussian a lifespan attribute for temporal modeling. Then scene graphs are also used in GS-based method, StreetGaussians [9] represents static and dynamic components as individual Gaussians within a Gaussian scene graph and optimizes them jointly, and OmniRe [1] integrates multiple Gaussian types to better reconstruct dynamic elements.

Depths in 3D Gaussian. Depth information provides essential geometric priors for resolving ambiguities in 3D reconstruction [19] and has been widely used in sparse-view reconstruction fields. In scenarios lacking ground-truth depth, many works adopt monocular depth estimates as weak supervision, leveraging scale-invariant loss functions [13, 14, 15] and depth ranking constraints [16, 17]. Recent advances such as DNGaussian [18] further improve this paradigm with patch-wise depth normalization. However, these monocular supervision methods inherently suffer from scale ambiguities due to the unknown shift and scale factors in predicted depth [21]. While some feed-forward 3DGS methods like MVSplat [20] and DepthSplat [21] address this by incorporating multi-view depth estimation, but they cannot handle dynamic scenes, and feed-forward depths are imperfect for further training. In contrast, our method leverages metric depth supervision while progressively refining depth predictions with generative priors through an iterative training scheme.

Diffusion priors in 3D Gaussian. Generative priors have recently been incorporated into 3DGS to reduce artifacts arising from geometric inaccuracies in novel-view synthesis [30, 22]. These approaches leverage diffusion models to enforce semantic coherence, hallucinate photometric details, and extrapolate occluded regions. DriveDreamer4D [31] further adapts this framework to dynamic driving scenarios. However, all of these methods concentrate on image-space priors and overlook the orthogonal role of depth-space generative priors. Diffusion-based monocular depth estimators such as Marigold [32] establish a foundation for learned depth priors, and Marigold-DC [23] shows




Figure 2: Pipeline of D^2GS . We first employ a Progressive Pruning strategy to obtain a robust global Gaussian initialization. A Road Node is incorporated into the scene graph structure to regularize the road region using strong geometric priors. During training, Gaussian optimization and depth refinement are performed iteratively, allowing depth to be learned jointly from Gaussian supervision and enhanced by diffusion priors from a pretrained depth foundation model.

that classifier-guided sampling can propagate sparse metric depths into dense predictions. However, our approach diverges by directly integrating Marigold’s generative depth prior into the 3DGS optimization loop to refine the dense imperfect depths.

3 Method

Our proposed D^2GS (Fig. 2) addresses the challenge of LiDAR-free dynamic urban scene reconstruction through three key components. First, we propose an initialization strategy that leverages estimated depth as input and applies progressive pruning to obtain a compact Gaussian representation (Sec. 3.1). Second, during Gaussian optimization, we propose a multi-view Depth Enhancer module to iteratively refine the rendering depth, providing dense geometric supervision (Sec. 3.2). Third, we incorporate a specialized road node into the scene graph representation to explicitly model the ground plane using strong geometric priors (Sec. 3.3). We illustrate these components in the following subsections.

3.1 Initialization via Progressive Pruning

Effective initialization of the 3D Gaussians is crucial for achieving high-quality reconstruction results with Gaussian Splatting. While most current methods [9, 1] rely on LiDAR point clouds for this purpose, our LiDAR-free setting requires an alternative solution. We begin by predicting an initial depth map for each input image using a pre-trained multi-view depth estimation network [21], denoted as D_i for i_{th} image. Given the camera pose and D_i , we unproject each depth map into a point cloud. By aggregating these point clouds, we obtain a unified point cloud $\mathbf{P}$.

The point cloud $\mathbf{P}$ contains as many points as image pixels, resulting in an excessively large size (about 100 million) that is computationally impractical for Gaussian training. To address this challenge, one common solution is to de-densify the points by voxel grids. However, this can lead to the loss of fine geometric details. To mitigate this, we propose a progressive pruning strategy that learns to de-densify the points in a learnable manner. Specifically, we initialize Gaussians with a non-deformable and low-resolution training setting to optimize these points and progressively eliminate less important Gaussians based on their opacity, *i.e.* the significance of the Gaussian in the scene representation. After each round of progressive pruning, we reset the opacity values and continue the next round. The progressive pruning process can be formulated as:

$$\mathbf{G}^{(N+1)} = \left\{ g_i \in \mathbf{G}^{(N)} \mid o_i^{(N)} \geq \tau \right\},$$

where $\mathbf{G}^{(N)}$ represents the set of Gaussians before pruning at round N , $o_i^{(N)}$ is the opacity of one Gaussian point g_i , and τ is the opacity threshold. This progressive strategy produces a manageable set of Gaussians that capture the essential scene geometry, effectively balancing computational efficiency and the quality of the initial geometric representation.

3.2 Joint Optimization of Depth and Gaussian Splatting

In the absence of LiDAR depth, obtaining reliable metric depth as a geometric constraint becomes challenging. Our experiments found that the estimated metric depths [21] are not sufficiently accurate for training, as directly applying these noisy depth maps as supervision leads to significant degradation in both reconstruction quality and geometric accuracy. (refer to OmniRe+DS results in Tab. 2).

On the one hand, Gaussian training can infer depth leveraging spatial consistency, and correct inconsistent depth estimates [21]. On the other hand, metric depth estimation, which aligns closely with image semantics, can provide smooth and fine-grained depth supervision, especially when training views are insufficient. To marry the advantages of both sides, we propose a joint optimization strategy that iteratively refines both the depth predictions and the Gaussian parameters. We perform iterative depth optimization as part of the Gaussian optimization process. For each depth optimization iteration, we first identify the high-confidence region from the previously estimated depth map, then refine the dense depth with the geometric information guided by high-confidence regions, using current Gaussian geometry to condition the *Depth Enhancer*. We then update the supervision (Multiview) depth pool to guide the Gaussian learning process.

Identify the High-confidence Region. Inspired by methods handling uncertain inputs [33, 34], we introduce a confidence measure based on the consistency between the previously estimated depth map and the currently rendered depth. The high-confidence region reflects the consistency between Gaussian training and depth estimation, indicating high accuracy. Specifically, for each pixel in view i , we compute the relative difference between the current depth $\mathbf{D}_i^{\text{cur}}$ and the previous iteration's depth $\mathbf{D}_i^{\text{prev}}$. The confidence map $\mathbf{C}_i$ is formulated as:

$$\mathbf{C}_i = \mathbb{I} \left\{ \max(\mathbf{D}_i^{\text{cur}} \oslash \mathbf{D}_i^{\text{prev}}, \mathbf{D}_i^{\text{prev}} \oslash \mathbf{D}_i^{\text{cur}}) < \lambda_C \right\}, \quad (1)$$

where $\oslash$ denotes element-wise division, $\mathbb{I}$ is an element-wise indicator function that evaluates to 1 if the enclosed condition is true, and 0 otherwise.

Multi-view Depth Enhancer for Depth Refinement. We employ a Multi-view Depth Enhancer module to refine the depth map, especially the low-confidence region. Inspired by the Marigold-DC [23], the proposed module leverages generative priors from the pre-trained monocular depth prediction model, *i.e.*, Marigold [32], by optimizing latent depth representations and scale-shift parameters through classifier-free guidance. However, it differs from Marigold-DC in the guidance accuracy: instead of relying on sparse ground-truth depth, our Depth Enhancer's diffusion steps are guided by the potentially noisy depth estimates $\mathbf{D}_i$ which are rendered by the Gaussian at the current state. Therefore, we design three losses to suppress the potential noise and enforce multi-view geometric consistency to obtain more accurate estimates.

i) Reference loss L_{ref} . Similar to Marigold-DC that use sparse ground-truth depth to inject metric depth information, we also use the rendered depth as a reference signal. However, we do not use it directly as supervision since the rendered depth $\mathbf{D}_i$ is not reliable. Instead, we first downsample both the predicted depth $\hat{\mathbf{D}}_i^{(k)}$ (the predicted depth map for view i at diffusion step k) and the reference depth $\mathbf{D}_i$ to a patch level (size $p \times p$) with average pooling, denoted by $\Downarrow$. This reduces the influence of local fluctuations. The loss then computes a combination of L1 and L2 distances between the downsampled maps, modulated by the confidence map $\mathbf{C}_i$ from Eq. 1:

$$L_{\text{ref}} = \Downarrow \mathbf{C}_i \odot \left(\left\| \Downarrow \hat{\mathbf{D}}_i^{(k)} - \Downarrow \mathbf{D}_i \right\|_1 + \left\| \Downarrow \hat{\mathbf{D}}_i^{(k)} - \Downarrow \mathbf{D}_i \right\|_2^2 \right). \quad (2)$$

The confidence-weighted loss ensures that the refinement closely follows the reference depth in high-confidence regions, while in uncertain areas, it guides the generation of metric depth based on both the reference depth and the image semantics.

ii) Multi-view warping loss L_w . It is also essential to ensure that the depth remains consistent across multiple viewpoints. Thus, we employ a projection mechanism to further guide the depth generation. Specifically, the depth warping operation $\mathcal{W}$ transforms the predicted depth map from another view j to the view of the reference frame i using the given camera poses. Then we penalize cross-view inconsistency via an L1 loss between warped and predicted depths:

$$L_w = \sum_{j \neq i} \left\| \mathcal{W}_{j \rightarrow i}(\mathbf{D}_j) - \hat{\mathbf{D}}_i^{(k)} \right\|_1, \quad (3)$$

The gradient of the warping loss enforces the geometry consistency between the predicted depth and the depth of the adjacent frames.

iii) Smooth loss L_{smooth} . To promote smoothness in the refined depth map, we incorporate an edge-aware smoothness constraint on the predicted depth $\hat{\mathbf{D}}_i^{(k)}$, conditioned on the ground-truth input image $\mathbf{I}_i^{gt}$:

$$L_{\text{smooth}} = \frac{1}{HW} \sum_{x,y} \left(\left| \partial_x \hat{\mathbf{D}}_i^{(k)}(x,y) \right| e^{-|\partial_x \mathbf{I}_i(x,y)|} + \left| \partial_y \hat{\mathbf{D}}_i^{(k)}(x,y) \right| e^{-|\partial_y \mathbf{I}_i(x,y)|} \right), \quad (4)$$

where ∂_x and ∂_y denote the spatial gradients in the horizontal and vertical directions, respectively. The weight $e^{-|\partial \mathbf{I}_i(x,y)|}$ reduces the influence of the smooth loss in areas with strong edge.

Finally, the total loss function guiding the diffusion optimization at each step is a weighted sum:

$$L_{\text{diffuse}} = \lambda_{\text{ref}} L_{\text{ref}} + \lambda_{\text{smooth}} L_{\text{smooth}} + \lambda_w L_w. \quad (5)$$

During the Gaussian training phase, the Depth Enhancer iteratively refines the depth $\mathbf{D}_i^{\text{cur}}$ to $\mathbf{D}_i^{\text{next}}$ every N iterations in order to balance time costs and effectiveness. More details of the iterative training scheme will be illustrated in the supplementary.

3.3 Optimization of Road Node in Gaussian Scene Graph

Roads are a key component of urban scenes and exhibit strong geometric priors, *i.e.* a specific height that can be precomputed from the point cloud and a known planar structure. Similar to [35, 10, 36], we also take advantage of the prior knowledge of roads to obtain a better depth/geometry on road region. Specifically, we introduce a specialized Road Node within our Gaussian scene graph structure, initializing it with all points identified as belonging to the road through segmentation. Then two geometry regularization are applied:

Position Constraint. We pre-compute the average height of the initial road point cloud and constrain the positions of the road Gaussians to stay close to this reference height.

$$L_{\text{plane}} = \frac{1}{|\mathbf{G}_{\text{road}}|} \sum_{g_i \in \mathbf{G}_{\text{road}}} (\mu_{z,i} - \bar{\mu}_z)^2, \quad (6)$$

where $\mathbf{G}_{\text{road}}$ is the set of Gaussians assigned to the Road Node, $\mu_{z,i}$ denotes the z-coordinate of the center of Gaussian g_i , and $\bar{\mu}_z$ is the mean road height of the initial points on the road.

Planar Constraint. Also, the shape and orientation of individual Gaussians within the Road Node should be constrained as a flat planar. Specifically, we encourage the Gaussian to be flat along the ground. Similar to [36], this constraint is formulated as,

$$L_{\text{shape}} = \frac{1}{|\mathbf{G}_{\text{road}}|} \sum_{g_i \in \mathbf{G}_{\text{road}}} \left(\lambda_{\text{normal}} \|\mathbf{n}_i - \hat{\mathbf{z}}\|_2 + \lambda_{\text{flat}} \|s_{\text{min},i}\|_1 \right), \quad (7)$$

where λ_{normal} and λ_{flat} are hyperparameters. For each Gaussian, $s_{\text{min},i}$ denote the smallest scale and the normal $\mathbf{n}_i$ derives from the the axis aligns with smallest scale. $\hat{\mathbf{z}}$ is the normal of the ground. The overall loss for the Road Node is the sum of these two components: $L_{\text{road}} = L_{\text{plane}} + L_{\text{shape}}$.

3.4 Loss functions for Gaussian Training

In Gaussian training iterations, we follow the training loss of OmniRe [1], except for the road loss:

$$L = (1 - \lambda_r) L_1 + \lambda_r L_{\text{SSIM}} + \lambda_{\text{road}} L_{\text{road}} + \lambda_{\text{depth}} L_{\text{depth}} + \lambda_{\text{opacity}} L_{\text{opacity}} + L_{\text{reg}}, \quad (8)$$

Table 1: Comparison of image reconstruction performances on Waymo NOTR Dynamic32 Dataset. OmniRe+DS refers to OmniRe using DepthSplat’s estimated depth for depth regularization.

Inputs	Methods	Image reconstruction			Novel view synthesis		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Img+LiDAR	EmerNeRF	28.16	0.806	0.228	25.14	0.747	0.313
	3DGS	28.47	0.876	0.136	25.14	0.813	0.165
	S3GS	31.35	0.911	0.106	27.44	0.857	0.137
	PVG	34.25	0.936	0.111	28.83	0.855	0.159
	OmniRe	33.35	0.945	0.067	29.02	0.871	0.115
Img	OmniRe w/o LiDAR	34.07	0.947	0.079	27.69	0.853	0.146
	OmniRe+DS	31.26	0.932	0.089	27.37	0.854	0.136
	Ours	34.34	0.950	0.074	28.80	0.867	0.129

Table 2: Comparison of depth estimation performances among multiple LiDAR-free methods and a multi-view depth estimation method DepthSplat on Waymo NOTR Dynamic32 Dataset. All metrics are computed with sparse GT depth from LiDAR.

Inputs	Methods	L1 $\downarrow$	Abs. Rel. $\downarrow$	RMSE $\downarrow$	$\delta < 1.25 \uparrow$
Img	DepthSplat	8.7088	0.5412	14.6494	0.6020
	OmniRe w/o LiDAR	4.2197	0.1977	8.3008	0.7212
	OmniRe+DS	5.5345	0.3137	8.5426	0.6422
	Ours	2.6043	0.1313	4.6723	0.8404

where L_1 and L_{SSIM} are the L1 and SSIM losses on rendered images, L_{depth} compares the rendered depth of Gaussians with dense reference depth, $L_{opacity}$ encourages the opacity of the Gaussians to align with the non-sky mask, and L_{reg} represents various regularization terms applied to different Gaussian representations.

4 Experiments

4.1 Experimental Setup

Datasets. We conducted a detailed experimental analysis using the Waymo dataset [37], a widely-used dataset for real-world driving scenarios. To evaluate our method’s performance in dynamic urban environments, we selected the challenging Waymo dataset subset, NOTR Dynamic32, proposed by [27]. We utilize data from the three frontal cameras (FRONT, FRONT RIGHT, FRONT LEFT), resizing the images to a resolution of 640×960 .

Metrics. We adopt PSNR, SSIM, and LPIPS as metrics to evaluate the quality of image reconstruction. Since our method does not use ground-truth point cloud depth for supervision, we employ L1, RMSE, Abs.Rel., and Threshold accuracy(δ) [21] to assess the depth estimation performance. Specifically, these metrics are evaluated with the sparse ground-truth depth obtained from LiDAR.

Baselines. We compare our method against several Gaussian Splatting approaches: 3DGS [5], S3GS [7], PVG [8], and OmniRe [1], and NeRF-based approach EmerNeRF [27]. However, these methods all require a LiDAR point cloud for optimization. For a fairer comparison, we include two additional baselines, (1) **OmniRe w/o LiDAR**, which trains OmniRe without any depth regularization, and (2) **OmniRe + DS**, which trains OmniRe using depth estimates from DepthSplat [21] in place of LiDAR depths. For more implementation details, please refer to the supplemental material.

4.2 Evaluation on Urban Driving Dataset

4.2.1 Image Reconstruction Performance

Quantitative results for image reconstruction and novel view synthesis, presented in Tab. 1, demonstrate that our D²GS method consistently outperforms all other methods and is even comparable to methods that use ground-truth (GT) LiDAR. Notably, OmniRe [1] without depth supervision shows competitive image quality on training views. We infer that the absence of explicit geometric




Figure 3: Comparison of image reconstruction and depth estimation performances on Waymo NOTR Dynamic32 Dataset with S3GS [7], PVG [8], OmniRe [1], and LiDAR-free baselines. Zoom in for better visual comparison.

constraints allows its Gaussian representation to overfit to the training views, resulting in poor generalization to novel views (PSNR drops from 34.07 to 27.69). The results of OmniRe+DS also validates that using the inaccurate depth prediction as supervision will in turn degrade reconstruction quality.

As shown in Fig. 3, D²GS renders more coherent images with superior reconstruction of fine details like vehicle contours and building textures. In contrast, former methods with LiDAR supervision may suffer from a quality degradation where the LiDAR is missing; the results of OmniRe [1] PVG [8] and S3GS [7] become blurry in the upper region. We can also observe the inaccurate depth rendering in the corresponding regions, S3GS [7] shows an evident dividing line in the depth map. As for LiDAR-free methods, although Gaussian training can overfit to images in the training views, it suffers from distortion and floater artifacts in novel views. Visualizations of the corresponding depth maps further confirm that it is difficult to recover accurate depth either by relying solely on Gaussian training (OmniRe w/o LiDAR) or on noisy depth predictions (OmniRe + DS).

4.2.2 Depth Estimation Performance

As shown in Tab. 2, D²GS achieves significant improvement on depth estimation accuracy compared to methods without GT LiDAR supervision. First, DepthSplat, which provides our initial depth estimates, shows inferior performance. This may be due to its inability to handle dynamic scenes. Moreover, compared to DepthSplat [21], the OmniRe + DS method shows an improvement, suggesting that Gaussian training has some capacity to correct noisy depth predictions. However, it still performs worse than OmniRe w/o LiDAR. We infer that, while depth supervision can serve as an effective regularizer for Gaussian training, inaccurate depth predictions may mislead the optimization process. In contrast, the depth accuracy of our method achieves approximately a 53% improvement over OmniRe + DS, validating the effectiveness of our proposed modules, given that the same initialization is employed.

A similar conclusion can be drawn from the Qualitative results (Fig. 3), as we have provided some analysis in Sec. 4.2.1. Our proposed method generates accurate dense depth maps, which offer

Table 3: Ablation study of our proposed key components and training strategy.

Components			Settings of D.E.				L1 ↓	Abs. Rel. ↓	RMSE ↓	$\delta < 1.25 \uparrow$
R.N.	P.P.	D.E.	L_{ref}	L_w	L_{smooth}	Update				
×	✓	✓	✓	✓	✓	Real-time	2.7580	0.1535	4.7748	0.8080
✓	×	✓	✓	✓	✓	Real-time	2.9121	0.1491	5.1675	0.8041
✓	✓	×	×	×	×	Real-time	2.8190	0.1595	5.3086	0.8107
✓	✓	✓	✓	×	×	Real-time	2.6477	0.1488	4.6977	0.8163
✓	✓	✓	✓	✓	✓	Periodic	2.9168	0.1700	5.0385	0.7701
✓	✓	✓	✓	✓	✓	Real-time	2.3847	0.1303	4.2854	0.8526




Figure 4: Visualization comparison of the proposed modules and training strategy. Zoom in for better comparison.

effective regularization and suppress floaters in regions where the depth data is missing or unreliable. It is worth noting that, apart from the issue of missing upper LiDAR points, the Waymo dataset exhibits minimal reprojection and calibration errors, as it is a meticulously calibrated dataset. However, such errors are common in real-world, industrial-scale data, where our method will demonstrate more significant improvements. More visual comparisons are provided in Supplementary.

4.3 Ablation Studies

Building upon the previous results, we perform ablation studies on a Waymo subset (12 sequences, detailed in supplementary material) to assess the impact of D²GS’s key components: Progressive Pruning (P.P.), Depth Enhancer (D.E.), and Road Node (R.N.). We evaluate variants by removing each component individually. We further analyze the loss function and update strategy of the Depth Enhancer.

4.3.1 Network Components

Effectiveness of Progressive Pruning. Without progressive pruning (D²GS w/o P.P.), Gaussians are initialized directly from a dense point cloud. The results (Tab. 3) are worse because initial noise and outliers in the dense points lead to flawed Gaussian representations. Floaters arise when training from this setting (Fig. 4). Our P.P. strategy successfully filters these initial points, creating a robust global geometry for the following reconstruction.

Effectiveness of Depth Enhancer. Removing the Depth Enhancer (D²GS w/o D.E.) degrades depth estimation accuracy and quality (Tab. 3). The D.E. iteratively refines depth rendered from current Gaussians using diffusion priors, establishing a crucial loop for geometric accuracy. Without it, the model struggles with insufficient geometric constraints, failing to capture fine geometric details and exhibiting depth errors and distortions.

Effectiveness of Road Node Constraint. Excluding the Road Node constraint (D²GS w/o R.N.) removes the strong ground plane prior, leading to a decline in depth metrics (Tab. 3) and qualitative artifacts such as uneven ground surfaces (Fig. 4). Inferring geometry for large, textureless surfaces like roads from multi-view images alone is challenging. Road Nodes enforce shape and normal constraints on Gaussians in detected road areas, promoting smooth and continuous ground geometry.

4.3.2 Settings of Depth Enhancer

Loss Function. We ablate the D.E.’s loss function by removing its projection and smooth loss terms. As shown in Tab. 3, this degrades D.E. performance and overall reconstruction quality. These loss terms provide geometric information from the Gaussian with less error, enabling the D.E. to generate more consistent and accurate depth maps, which in turn provide more effective geometric supervision during iterative optimization, validating our design.

Update Strategy. We compare our adopted *Real-time Update* strategy for the D.E. with a *Periodic Update* strategy. In *Periodic Update*, depth maps for all training views are updated at the same iteration. In our *Real-time Update*, only a single depth map of one training view is rendered, enhanced, and used for depth regularization during a depth enhancement iteration. Experimental results (Tab. 3) show that *Real-time Update* significantly outperforms *Periodic Update*. The reason is that *Periodic Update* suffers from using outdated depth for supervision, which leads to a conflict between Gaussian optimization and depth regularization. *Real-time Update* ensures the D.E. takes the latest Gaussian geometry as a conditioning signal, providing timely geometric constraints for Gaussian training.

5 Conclusion

In this paper, we presented D²GS, a novel LiDAR-free framework for reconstructing dynamic urban street scenes. By effectively leveraging geometric priors derived from images, D²GS eliminates the dependency on LiDAR data, thus significantly simplifying data acquisition for autonomous driving applications. The core ideas of D²GS include: 1) An initialization strategy using multi-view depth estimation followed by a progressive training and pruning process, which efficiently generates a compact and representative initial Gaussian set. 2) A joint optimization approach employing a diffusion-based depth enhancer, which iteratively refines both the depth estimates and the Gaussian representation, providing robust geometric supervision throughout training. 3) The integration of a dedicated road node into the scene graph GS representation, explicitly modeling the ground plane using strong geometric priors. Extensive experimental results on the Waymo dataset validate the efficacy of our LiDAR-free reconstruction pipeline. While D²GS successfully eliminates the need for LiDAR, it still relies on accurate camera poses. In future work, we will investigate pose-free implementations to further simplify the pipeline and broaden its applicability.

References

- [1] Z. Chen, J. Yang, J. Huang, R. de Lutio, J. M. Esturo, B. Ivanovic, O. Litany, Z. Gojcic, S. Fidler, M. Pavone *et al.*, “Omnire: Omni urban scene reconstruction,” *arXiv preprint arXiv:2408.16760*, 2024.
- [2] H. Zhu, Z. Zhang, J. Zhao, H. Duan, Y. Ding, X. Xiao, and J. Yuan, “Scene reconstruction techniques for autonomous driving: a review of 3d gaussian splatting,” *Artificial Intelligence Review*, vol. 58, no. 1, p. 30, 2024.
- [3] Z. Yang, Y. Chen, J. Wang, S. Manivasagam, W.-C. Ma, A. J. Yang, and R. Urtasun, “Unisim: A neural closed-loop sensor simulator,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 1389–1399.
- [4] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “Nerf: Representing scenes as neural radiance fields for view synthesis,” *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [5] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, “3d gaussian splatting for real-time radiance field rendering,” *ACM Trans. Graph.*, vol. 42, no. 4, pp. 139–1, 2023.
- [6] B. Fei, J. Xu, R. Zhang, Q. Zhou, W. Yang, and Y. He, “3d gaussian splatting as new era: A survey,” *IEEE Transactions on Visualization and Computer Graphics*, 2024.
- [7] N. Huang, X. Wei, W. Zheng, P. An, M. Lu, W. Zhan, M. Tomizuka, K. Keutzer, and S. Zhang, “S3gaussian: Self-supervised street gaussians for autonomous driving,” *arXiv preprint arXiv:2405.20323*, 2024.
- [8] Y. Chen, C. Gu, J. Jiang, X. Zhu, and L. Zhang, “Periodic vibration gaussian: Dynamic urban scene reconstruction and real-time rendering,” *arXiv preprint arXiv:2311.18561*, 2023.
- [9] Y. Yan, H. Lin, C. Zhou, W. Wang, H. Sun, K. Zhan, X. Lang, X. Zhou, and S. Peng, “Street gaussians: Modeling dynamic urban scenes with gaussian splatting,” in *European Conference on Computer Vision*. Springer, 2024, pp. 156–173.

- [10] M. Khan, H. Fazlali, D. Sharma, T. Cao, D. Bai, Y. Ren, and B. Liu, "Autosplat: Constrained gaussian splatting for autonomous driving scene reconstruction," *arXiv preprint arXiv:2407.02598*, 2024.
- [11] X. Zhou, Z. Lin, X. Shan, Y. Wang, D. Sun, and M.-H. Yang, "Drivinggaussian: Composite gaussian splatting for surrounding dynamic autonomous driving scenes," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024, pp. 21 634–21 643.
- [12] Y. Li and J. Ibanez-Guzman, "Lidar for autonomous driving: The principles, challenges, and trends for automotive lidar and perception systems," *IEEE Signal Processing Magazine*, vol. 37, no. 4, pp. 50–61, 2020.
- [13] S. Kim, "Därf: Boosting radiance fields from sparse inputs with monocular depth adaptation," in *Neural Information Processing Systems (NeurIPS)*, 2023.
- [14] H. Xiong, S. Muttukuru, R. Upadhyay, P. Chari, and A. Kadambi, "Sparsegs: Real-time 360° sparse view synthesis using gaussian splatting," *arXiv preprint arXiv:2312.00206*, 2023.
- [15] Z. Zhu, Z. Fan, Y. Jiang, and Z. Wang, "Fsgs: Real-time few-shot view synthesis using gaussian splatting," in *European conference on computer vision*. Springer, 2024, pp. 145–163.
- [16] G. Wang, Z. Chen, C. C. Loy, and Z. Liu, "Sparsenerf: Distilling depth ranking for few-shot novel view synthesis," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 9065–9076.
- [17] D. Xu, Y. Jiang, P. Wang, Z. Fan, Y. Wang, and Z. Wang, "Neurallift-360: Lifting an in-the-wild 2d photo to a 3d object with 360deg views," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 4479–4489.
- [18] J. Li, J. Zhang, X. Bai, J. Zheng, X. Ning, J. Zhou, and L. Gu, "Dngaussian: Optimizing sparse-view 3d gaussian radiance fields with global-local depth normalization," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024, pp. 20 775–20 785.
- [19] K. Deng, A. Liu, J.-Y. Zhu, and D. Ramanan, "Depth-supervised nerf: Fewer views and faster training for free," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 12 882–12 891.
- [20] Y. Chen, H. Xu, C. Zheng, B. Zhuang, M. Pollefeys, A. Geiger, T.-J. Cham, and J. Cai, "Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images," in *European Conference on Computer Vision*. Springer, 2024, pp. 370–386.
- [21] H. Xu, S. Peng, F. Wang, H. Blum, D. Barath, A. Geiger, and M. Pollefeys, "Depthsplat: Connecting gaussian splatting and depth," *arXiv preprint arXiv:2410.13862*, 2024.
- [22] X. Liu, C. Zhou, and S. Huang, "3dgs-enhancer: Enhancing unbounded 3d gaussian splatting with view-consistent 2d diffusion priors," *Advances in Neural Information Processing Systems*, vol. 37, pp. 133 305–133 327, 2024.
- [23] M. Viola, K. Qu, N. Metzger, B. Ke, A. Becker, K. Schindler, and A. Obukhov, "Marigold-dc: Zero-shot monocular depth completion with guided diffusion," *arXiv preprint arXiv:2412.13389*, 2024.
- [24] H. Turki, D. Ramanan, and M. Satyanarayanan, "Mega-nerf: Scalable construction of large-scale nerfs for virtual fly-throughs," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 12 922–12 931.
- [25] M. Tancik, V. Casser, X. Yan, S. Pradhan, B. Mildenhall, P. P. Srinivasan, J. T. Barron, and H. Kretzschmar, "Block-nerf: Scalable large scene neural view synthesis," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 8248–8258.
- [26] H. Turki, J. Y. Zhang, F. Ferroni, and D. Ramanan, "Suds: Scalable urban dynamic scenes," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 12 375–12 385.
- [27] J. Yang, B. Ivanovic, O. Litany, X. Weng, S. W. Kim, B. Li, T. Che, D. Xu, S. Fidler, M. Pavone *et al.*, "Emernerf: Emergent spatial-temporal scene decomposition via self-supervision," *arXiv preprint arXiv:2311.02077*, 2023.
- [28] J. Ost, F. Mannan, N. Thuerey, J. Knodt, and F. Heide, "Neural scene graphs for dynamic scenes," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 2856–2865.
- [29] Z. Wu, T. Liu, L. Luo, Z. Zhong, J. Chen, H. Xiao, C. Hou, H. Lou, Y. Chen, R. Yang *et al.*, "Mars: An instance-aware, modular and realistic simulator for autonomous driving," in *CAAI International Conference on Artificial Intelligence*. Springer, 2023, pp. 3–15.
- [30] F. Liu, W. Sun, H. Wang, Y. Wang, H. Sun, J. Ye, J. Zhang, and Y. Duan, "Reconx: Reconstruct any scene from sparse views with video diffusion model," *arXiv preprint arXiv:2408.16767*, 2024.

- [31] G. Zhao, C. Ni, X. Wang, Z. Zhu, X. Zhang, Y. Wang, G. Huang, X. Chen, B. Wang, Y. Zhang *et al.*, “Drivedreamer4d: World models are effective data machines for 4d driving scene representation,” *arXiv preprint arXiv:2410.13571*, 2024.
- [32] B. Ke, A. Obukhov, S. Huang, N. Metzger, R. C. Daudt, and K. Schindler, “Repurposing diffusion-based image generators for monocular depth estimation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 9492–9502.
- [33] R. Martin-Brualla, N. Radwan, M. S. Sajjadi, J. T. Barron, A. Dosovitskiy, and D. Duckworth, “Nerf in the wild: Neural radiance fields for unconstrained photo collections,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 7210–7219.
- [34] Z. Xie, J. Zhang, W. Li, F. Zhang, and L. Zhang, “S-nerf: Neural radiance fields for street views,” *arXiv preprint arXiv:2303.00749*, 2023.
- [35] H. Zhou, L. Lin, J. Wang, Y. Lu, D. Bai, B. Liu, Y. Wang, A. Geiger, and Y. Liao, “Hugsim: A real-time, photo-realistic and closed-loop simulator for autonomous driving,” *arXiv preprint arXiv:2412.01718*, 2024.
- [36] C. Peng, C. Zhang, Y. Wang, C. Xu, Y. Xie, W. Zheng, K. Keutzer, M. Tomizuka, and W. Zhan, “Desire-gs: 4d street gaussians for static-dynamic decomposition and surface reconstruction for urban driving scenes,” *arXiv preprint arXiv:2411.11921*, 2024.
- [37] P. Sun, H. Kretschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine *et al.*, “Scalability in perception for autonomous driving: Waymo open dataset,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 2446–2454.
- [38] L. Piccinelli, C. Sakaridis, Y.-H. Yang, M. Segu, S. Li, W. Abbeloos, and L. Van Gool, “Unidepthv2: Universal monocular metric depth estimation made simpler,” *arXiv preprint arXiv:2502.20110*, 2025.
- [39] W. Yin, C. Zhang, H. Chen, Z. Cai, G. Yu, K. Wang, X. Chen, and C. Shen, “Metric3d: Towards zero-shot metric 3d prediction from a single image,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 9043–9053.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction can clearly reflect our contribution, including mentioning our motivation, introduction to our method, and a brief summary of the experiment results.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We have mentioned it in the conclusion and supplementary material.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All of the formulas have been checked by the authors.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All the information needed for reproduction is provided in the Method, Experiments, and Supplemental Material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The code will be open-sourced after the company's review process.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide all the information in the supplemental material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We have tested our method on 32 sequences, which are commonly used in the field.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We have mentioned it in the Supplemental Material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have carefully checked the Code of Ethics, and our paper conforms it.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: There is no societal impact of the work performed.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We have cited all the papers we referenced. And there are no Licenses problem.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

Supplemental Material

A Implementation Details

Initialization: Our pipeline begins by generating initial depth estimates using DepthSplat [21] (weight name: "depthsplat-gs-base-dl3dv-256x448-randview2-6-d94d996f.pth"), which processes input images and camera parameters to produce per-view depth maps. These depth maps are then unprojected into 3D space using camera intrinsics and extrinsics to form a dense point cloud. We initialize static Gaussians from this point cloud, and then perform low-resolution training at a resolution of 240×160 for three 3,000-iteration pruning epochs. Through our progressive pruning strategy, we gradually filter out less significant Gaussians based on opacity values (the threshold is computed from percentiles of all Gaussians), ultimately retaining a subset of 2×10^6 Gaussians. To mitigate the impact of initial depth inaccuracies, we defer the activation of the Depth Enhancer module until the 1,000th iteration after the first epoch. The final rendered depth maps from this initialization phase (denoted as $\mathbf{D}^{\text{prev}}$) serve as the foundation for subsequent full training processes. For Road Node initialization, we use the road mask to identify the corresponding Gaussians and assign them to the Road Node.

Training Details: The model contains 60,000 iterations of training, divided into 40,000 iterations for RGB-only training and 20,000 iterations for iterative Gaussian and depth enhancement training. The reason why we do not perform depth enhancement in the first 40,000 iteration is that, the early training stages typically produce meaningless depth and image outputs. There are severe floaters and artifacts resulting from the aggressive densification process. We follow the default learning rate settings for Gaussian properties as established in OmniRe [1]. Our proposed Road Node's learning rate is 0.05. The Depth Enhancer module is activated halfway through training, with the diffusion process consisting of 80 steps and an early stop mechanism that prevents the smoothness loss from increasing for 8 continuous steps. For the warping loss, we utilize the views from 6 closest timestamp of the same camera ID. For quantitative comparisons, results for 3DGS, EmerNeRF, and S3GS are adopted from the S3GS [7] paper. All other methods were evaluated by us using their official codebases on a single NVIDIA H20 GPU. Key hyperparameters for our method are detailed in Tab. 4.

Table 4: Value of the key hyperparameters in our method.

Hyperparameter	Value	Description
$p \times p$	$\frac{H}{8} \times \frac{H}{8}$	Patch size for depth downsampling
λ_c	2.0	Confidence threshold in Depth Enhancer
λ_{ref}	1.0	Weight for reference depth loss in Depth Enhancer
λ_{smooth}	$\frac{1}{8}$	Weight for depth smoothness loss in Depth Enhancer
λ_w	$\frac{1}{16}$	Weight for warping loss in Depth Enhancer
N	20	Depth Enhancer update frequency
λ_{normal}	10.0	Weight for road normal alignment loss in Road Node
λ_{flat}	1.0	Weight for road flatness loss in Road Node
λ_{road}	0.1	Weight for Road Node loss

B Datasets and Baselines:

Our main experiments are conducted on the Waymo NOTR Dynamic32 [27] dataset, comparing our method against S3GS [7], PVG [8], OmniRe [1], and LiDAR-free baselines. The specific segment IDs are listed in Tab. 5. Our ablation studies are performed on a 12-sequence subset of the Waymo NOTR dataset, with segment IDs provided in Tab. 6.

Table 5: Segment IDs of Waymo NOTR Dynamic32 Dataset.

seg102319...	seg103913...	seg104444...	seg104980...	seg105887...	seg106250...	seg106648...	seg109636...
seg110170...	seg117188...	seg118463...	seg119178...	seg119284...	seg120278...	seg121618...	seg122514...
seg123392...	seg148106...	seg168016...	seg181118...	seg191876...	seg225932...	seg254789...	seg441423...
seg508351...	seg522233...	seg583504...	seg624282...	seg767010...	seg882250...	seg990779...	seg990914...

Table 6: Segment IDs of a 12 sequences subset used in our ablation experiments.

seg102319...	seg104980...	seg105887...	seg109636...	seg110170...	seg117188...
seg119284...	seg122514...	seg123392...	seg168016...	seg191876...	seg441423...

C Additional Results

In this section, we present more qualitative comparisons of our methods with others. Same as the main submission, we show the results of the LiDAR supervised methods (S3GS [7], PVG [8], OmniRe [1], and LiDAR-free methods (OmniRe w/o LiDAR, OmniRe + DepthSplat, ours).




Figure 5: Comparison of our method with S3GS, PVG, OmniRe, and LiDAR-free baselines.




Figure 6: Comparison of our method with S3GS, PVG, OmniRe, and LiDAR-free baselines.




Figure 7: Comparison of our method with S3GS, PVG, OmniRe, and LiDAR-free baselines.




Figure 8: Comparison of our method with S3GS, PVG, OmniRe, and LiDAR-free baselines.




Figure 9: Comparison of our method with S3GS, PVG, OmniRe, and LiDAR-free baselines.




Figure 10: Comparison of our method with S3GS, PVG, OmniRe, and LiDAR-free baselines.




Figure 11: Comparison of our method with S3GS, PVG, OmniRe, and LiDAR-free baselines.




Figure 12: Comparison of our method with S3GS, PVG, OmniRe, and LiDAR-free baselines.

D Additional Experiments and Analyses

Summary. This section expands the computational analysis, clarifies assumptions about camera poses and our handling of dynamic objects, details sky modeling, strengthens ablations (including image-level metrics and road-node losses), visualizes depth evolution during training, and analyzes the initialization choices (SfM vs. MVS vs. monocular metric depth).

Notation. We use the main-paper notation throughout. “P.P.” denotes *Progressive Pruning*, “D.E.” denotes the *Depth Enhancer*, and “R.N.” denotes the *Road Node*. NVS denotes novel-view synthesis.

D.1 Computational Analysis of Progressive Pruning

Dense depth initialization yields an extremely large set of back-projected points. Training Gaussians on these points is both redundant and error-prone, because abundant low-accuracy points often produce artifacts. P.P. therefore serves two goals: (i) reduce computation and memory by early removal of inconsistent Gaussians and (ii) deliver a higher-quality initial point cloud that improves reconstruction.

Computational Analysis. Tab. 7 profiles memory, runtime, and the evolution from initial to retained training points over the first 30k iterations. The results show that when the initial input point cloud contains as many as 15M points, the Gaussian model retains around 6M points that are not effectively pruned. In contrast, our P.P. method prevents a significant waste of computational resources and avoids efficiency bottlenecks. Even with very large initial sets, P.P. can avoid out-of-memory errors and shorten training by eliminating redundant, inconsistent Gaussians before full optimization. In contrast, the original pruning method prunes primarily during rendering optimization, which defers these savings.

Table 7: Analysis of computational cost, comparing our P.P. method with a baseline using a manageable number of points to demonstrate the efficiency gains.

Method	Initial Points	Training Points [†]	GPU RAM	Runtime
w/o P.P.	15M	6M - 8M	34 GB	10 h
P.P.	2M	2M - 3M	20 GB	8 h

[†]Retained points used in optimization within the first 30k iterations.

D.2 The Dependency on Camera Pose Accuracy

Despite the LiDAR calibration issue which we aim to solve in our method, we assume accurate camera poses, consistent with common urban driving scene reconstruction settings and dataset priors.

To further address inaccurate camera poses, a common issue in data calibration, we propose the following possible solutions for future work. Specifically, when precise camera poses are unavailable, pose estimation algorithms, such as learning-based methods or traditional SfM, can provide initial estimates. Subsequently, a joint optimization module for camera poses can be incorporated during reconstruction, enabling simultaneous refinement of both the poses and the reconstructed scene.

D.3 Modeling Dynamic Objects

Our method can handle dynamic objects properly through the following aspects:

Scene graph decomposition. We model and decompose dynamic objects using a Scene Graph, an approach that follows our baseline, OmniRe [1]. Each dynamic object is represented as an independent node within this graph, equipped with its own set of Gaussian models (for dynamic or rigid representations).

Depth refinement for dynamics. Initial MVS depths inevitably contain errors on moving objects. First, since our method uses given ground-truth camera poses from Waymo, dynamic objects do not affect pose estimation. Depth errors are therefore localized to these dynamic areas, with minimal

impact on the static background. Furthermore, during the iterative optimization of the Gaussians and the D.E., the depth of dynamic objects is further refined (as shown in Fig. 13). Specifically, this refinement is driven by both the photometric loss from Gaussian splatting and the accurate geometric supervision provided by our diffusion-based D.E.. This joint optimization loop is effective for both static and dynamic scene components, enabling correction of initial errors and more accurate depth prediction.

D.4 Detailed Ablations with Image-Level Metrics

We complement geometric ablations with image reconstruction and NVS metrics, with results shown in Tab. 8. The full model performs best across settings, corroborating Tab. 3 in the main paper: degraded depth quality typically yields poorer NVS. Notably, the Periodic Update variant attains high reconstruction quality but markedly worse NVS, suggesting overfitting to training-view photometry with stale depth. This also underscores the importance of our real-time iterative refinement.

Table 8: Image-level metrics for ablations (image reconstruction and NVS). “Update” denotes D.E. update scheme.

Components			Settings of D.E.				Image reconstruction			Novel view synthesis		
R.N.	P.P.	D.E.	L_{ref}	L_w	L_{smooth}	Update	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
×	✓	✓	✓	✓	✓	Real-time	32.67	0.939	0.087	27.67	0.854	0.139
✓	×	✓	✓	✓	✓	Real-time	32.25	0.932	0.099	27.17	0.847	0.149
✓	✓	×	×	×	×	Real-time	32.77	0.939	0.087	27.52	0.851	0.143
✓	✓	✓	✓	×	×	Real-time	32.57	0.940	0.086	27.29	0.852	0.140
✓	✓	✓	✓	✓	✓	Periodic	33.95	0.951	0.075	26.90	0.823	0.128
✓	✓	✓	✓	✓	✓	Real-time	33.63	0.947	0.079	28.04	0.863	0.130

D.5 Road Node: Loss Analysis

We ablate the two R.N. loss terms independently, and conduct experiments on a subset of Waymo dataset (sequence id: ‘031’, ‘035’, ‘084’, ‘089’, ‘102’, ‘111’, ‘222’, ‘382’). As shown in Tab. 9, both the loss L_{plane} and the loss L_{shape} contribute positively to a more accurate geometry (evaluated with depth results). The best performance is achieved when both are used together.

Table 9: Ablation of road-node losses on Waymo subset.

Method	L1 $\downarrow$	Abs. Rel. $\downarrow$	RMSE $\downarrow$	$\delta < 1.25 \uparrow$
w/o L_{plane}	2.7879	0.1395	5.0075	0.9270
w/o L_{shape}	2.6502	0.1211	5.0192	0.9320
Full loss	2.4245	0.1161	4.5535	0.9459

D.6 Depth Evolution Across Training

Fig. 13 visualizes the online refinement process by showing: (i) the Gaussian-rendered depth (D.E. input), (ii) the refined depth (D.E. output), and (iii) visualizations of this process at three checkpoints (initial, mid, and final training stages). We observe a clear removal of floating fragments, stabilization of large planar regions, and improved cross-view geometric consistency as the refinement progresses.

D.7 Initialization: SfM vs. MVS vs. Monocular

In this section, we give a more comprehensive analysis of different initialization methods. Existing depth estimation methods are inevitably imperfect and unsuitable to use as strict ground-truth for high-fidelity reconstruction. As discussed in the main paper, monocular methods often yield relative depth with unknown scale, while MVS methods can struggle in dynamic scenes. Our core contribution is a framework that *starts* from an imperfect *metric* depth prior and jointly optimizes geometry and depth. Metric initialization is crucial for scale-correct point-cloud seeding and for strong, multi-view-consistent supervision.




Figure 13: Depth evolution with online updates in three different stages: initial → final.

Recent monocular depth estimation foundation models (e.g., UniDepthV2 [38] and Metric3D [39]) also report strong metric depth prediction performance. We adopt a DepthSplat-style MVS depth estimator in our main paper. Nevertheless, our pipeline is orthogonal to the depth source: any reasonable *metric* initializer can be used.

Monocular metric depth (UniDepthV2) in place of MVS. We replace the DepthSplat MVS method with UniDepthV2 [38] to initialize our pipeline. We conduct experiments on a subset of Waymo dataset (sequence id: '031', '035', '084', '089', '102', '111', '222', '382'), and Tab. 10 shows that our method still improves over baselines, confirming orthogonality to the initializer. Gains are smaller than with DS, likely because UniDepthV2 provides stronger initial depths and thus less headroom for refinement.

Why not SfM: a sparse point cloud initialization. We define dense initialization as back-projecting *pixelwise* metric depth maps into a redundant but coverage-complete point cloud. Sparse initialization denotes sampled point clouds (e.g., LiDAR, COLMAP), which provide only sparse supervision and may miss fine-scale or specific parts of the geometry.

In urban driving scenes, SfM still tends to be sparse due to limited overlap, textureless regions, and occlusions. Our contributions align naturally with dense initialization: P.P. compresses dense

Table 10: Monocular metric depth estimation as an initialization alternative on Waymo subset.

Methods	Image reconstruction			Novel view synthesis		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
OmniRe	33.96	0.955	0.057	28.50	0.874	0.117
OmniRe+DS	31.30	0.936	0.078	28.11	0.876	0.110
OmniRe+UniDepthV2	32.73	0.944	0.068	29.67	0.885	0.098
Ours+DS	34.35	0.954	0.068	29.75	0.883	0.107
Ours+UnidepthV2	34.39	0.954	0.069	29.92	0.891	0.096

initialization priors before full optimization, and D.E. requires dense targets for per-pixel supervision. However, with minor changes (e.g., drop P.P. and seed D.E. with reprojected sparse depths), the pipeline still applies to SfM sparse initialization.