

EGSRAL: An Enhanced 3D Gaussian Splatting Based Renderer with Automated Labeling for Large-Scale Driving Scene

Yixiong Huo*, Guangfeng Jiang*, Hongyang Wei, Ji Liu[†], Song Zhang, Han Liu, Xingliang Huang, Mingjie Lu, Jinzhang Peng, Dong Li, Lu Tian, Emad Barsoum

Advanced Micro Devices, Inc., Beijing, China

{yixiong.huo, guangfeng.jiang, ji.liu, mingjiel, jinz.peng, d.li, lu.tian, emad.barsoum}@amd.com

Abstract

3D Gaussian Splatting (3D GS) has gained popularity due to its faster rendering speed and high-quality novel view synthesis. Some researchers have explored using 3D GS for reconstructing driving scenes. However, these methods often rely on various data types, such as depth maps, 3D boxes, and trajectories of moving objects. Additionally, the lack of annotations for synthesized images limits their direct application in downstream tasks. To address these issues, we propose EGSRAL, a 3D GS-based method that relies solely on training images without extra annotations. EGSRAL enhances 3D GS's capability to model both dynamic objects and static backgrounds and introduces a novel adaptor for auto labeling, generating corresponding annotations based on existing annotations. We also propose a grouping strategy for vanilla 3D GS to address perspective issues in rendering large-scale, complex scenes. Our method achieves state-of-the-art performance on multiple datasets without any extra annotation. For example, the PSNR metric reaches 29.04 on the nuScenes dataset. Moreover, our automated labeling can significantly improve the performance of 2D/3D detection tasks.

1 Introduction

Synthesizing novel photorealistic views represents a complex and crucial challenge within the fields of computer vision and graphics. With the rapid development of neural radiance fields (NeRFs) (Mildenhall et al. 2021), free-view synthesis has gradually transferred to the domain of large-scale view synthesis, particularly in synthesizing streetscapes critical for autonomous driving (Chen 2019). However, Simulating outdoor environments is challenging due to the complexity of geographical locations, diverse surroundings, and varying road conditions. Image-to-image translation methods (Wang et al. 2018; Isola et al. 2017) have been proposed to synthesize semantically labeled streetscapes by learning the mapping between source and target images. While these methods generate visually impressive street views, they often exhibit noticeable artifacts and inconsistent textures in local details. Additionally, the relatively uniform viewpoints of the synthesized

images present challenges for their application in complex autonomous driving scenes.

To address these challenges, Drive-3DAu (Tong et al. 2023) introduces a 3D data augmentation approach using NeRF, designed to augment driving scenes in the 3D space. DGNR (Li et al. 2023) presents a novel framework that learns a density space from scenes to guide the construction of a point-based renderer. Meanwhile, READ (Li, Li, and Zhu 2023) offers a large-scale driving simulation environment for generating realistic data for advanced driver assistance systems. 3D GS-based methods (Zhou et al. 2024; Yan et al. 2024) have been employed to synthesize driving scenes due to their superior generation capabilities. While these methods generate realistic images suitable for autonomous driving, they cannot simultaneously synthesize novel views and provide corresponding 2D/3D annotation boxes, which are crucial for supervised model training. **Consequently, enhancing novel view synthesis for large-scale scenes and achieving automatic annotations for new views remain key challenges in autonomous driving.**

To overcome these challenges, we introduce a novel framework called EGSRAL based on an enhanced 3D Gaussian Splatting (3D GS) technique. This framework improves the quality of novel view synthesis while simultaneously generating corresponding annotations. Specifically, we propose a deformation enhancement module to refine the Gaussian deformation field enhancing the modeling of both dynamic objects and static backgrounds. Additionally, we introduce an opacity enhancement module that leverages a neural network, replacing the original learnable parameter, to significantly boost the modeling capacity of complex driving scenes. Furthermore, to address the issue of unrealistic viewpoints in rendering large-scale complex scenes, where occluded distant Gaussians should not be included, we also propose a grouping strategy for vanilla 3D GS.

In short, our contributions are as follows: (1) We propose an enhanced 3D GS-based renderer called EGSRAL, which can synthesize novel view images with corresponding annotations based on existing dataset annotations. EGSRAL introduces a deformation enhancement module and an opacity enhancement module, both of which enhance the modeling capability of 3D GS for complex scenes. (2) Additionally, to address the issue of unreasonable perspectives in rendering large-scale, complex scenes, we introduce a grouping

*Equal contribution.

[†]Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

strategy for vanilla 3D GS. (3) Unlike previous methods that focus solely on novel view synthesis, we propose an adaptor with three constraints to transform neighboring annotation boxes into novel annotation boxes in the domain of autonomous driving. (4) Experimental results demonstrate that our method outperforms existing rendering methods for large-scale scenes. Additionally, the novel view images with corresponding annotations effectively improve the performance of 2D/3D detection models.

2 Related Work

2.1 3D Gaussian Splatting

Recently 3D GS has achieved exciting results in novel view synthesis and real-time rendering. Unlike NeRF, which uses implicit functions and volumetric rendering, 3D GS leverages a set of 3D Gaussians to model the scene, allowing for fast rendering through a tile-based rasterizer.

The original 3D GS (Kerbl et al. 2023) is designed to model static scenes, and subsequent work has tried to extend it to dynamic scenes. (Yang et al. 2024) introduce a deformation field to model monocular dynamic scenes, which use encoded Gaussian position and timestep to predict the position, rotation, and scaling offsets of 3D Gaussian. 4D GS (Wu et al. 2024) uses HexPlane (Cao and Johnson 2023) to encode temporal and spatial information for 3D Gaussian, building Gaussian features from 4D neural voxels, and then predicting the deformation of 3D Gaussians. However, compared to indoor scenes, outdoor scenes present greater challenges due to their larger spatial extent and the wider range of motion of dynamic objects. The above 3D GS methods often do not scale well to large-scale outdoor autonomous driving scenes. In contrast, our approach is more effective at reconstructing autonomous driving scenes.

2.2 Novel-View Synthesis for Driving Scene

Autonomous driving scene reconstruction methods based on NeRF can be roughly divided into two categories: perception-based (Fu et al. 2022; Xu et al. 2023; Kundu et al. 2022; Zhi et al. 2021) and simulation-based (Wu et al. 2023; Zhang et al. 2023; Li, Li, and Zhu 2023). Perception-based methods leverage capabilities of NeRF to capture the semantic information and geometrical representations. For simulation, MARS (Wu et al. 2023) models the foreground objects and background environments separately based on NeRF, making it flexible for scene controlling in autonomous driving simulation. Furthermore, READ (Li, Li, and Zhu 2023) explores diverse sampling strategies aimed at enabling the synthesis of expansive driving scenarios.

Some recent works have begun to bring 3D GS into the reconstruction of driving scenes. DrivingGaussian (Zhou et al. 2024) uses incremental static Gaussians and composite dynamic Gaussians to reconstruct static scenes and dynamic objects, respectively. Street Gaussians (Yan et al. 2024) represents the scene as a set of point-based background and foreground objects where an optimization strategy is introduced to deal with dynamic foreground vehicles. Although the above methods introduce a novel view synthesis paradigm into the autonomous driving domain, mainstream

methods still heavily rely on annotated data. Simply synthesizing novel views without annotations can not be directly applied to the downstream perception tasks. Therefore, to address this challenge, we first propose a new framework for synthesizing new views based on the 3D GS while enabling automatic annotation of downstream perception tasks.

2.3 Auto Labeling With Scene Reconstruction

Currently, some works partially explored the potential of NeRFs to serve as data factories at a level to achieve auto labeling in other task domains (e.g., depth estimation (Chang et al. 2022), object detection (Ge et al. 2022), semantic segmentation (Zhi et al. 2021)) to further improve domain performance. For instance, NS (Tosi et al. 2023) introduces a NeRF-supervised training procedure which exploits rendered stereo triplets to compensate for occlusions and depth maps as proxy labels to improve the depth stereo performance. Semantic-NeRF (Zhi et al. 2021) extends NeRF to jointly encode semantics with appearance and geometry to complete and accurate 2D semantic labels which can be achieved using a small amount of in-place annotations specific to the scene. Neural-Sim (Ge et al. 2022) proposes a novel bilevel optimization algorithm to automatically optimize rendering parameters (pose, zoom, illumination) to generate optimal data for downstream tasks using NeRF, which substitutes the traditional graphics pipeline and synthesizes useful images. However, there is currently no work to introduce auto labeling into 3D GS scene reconstruction. Inspired by the above methods, we consider incorporating auto labeling in the autonomous driving domain to further advance downstream perception tasks.

3 Methodology

3.1 Overview

Given input image sequences of a driving scene and point clouds estimated through structure-from-motion (SfM) methods (Agisoft 2019; Schonberger and Frahm 2016), our EGSRAL framework can synthesize realistic driving scenes from multiple perspectives while automatically labeling the corresponding novel synthetic views. We also propose a grouping strategy to address perspective issues in large-scale driving scenes. Our framework is divided into two parts: an Enhanced 3D GS Rendering and a Novel View Auto Labeling, as illustrated in Figure 1. The 3D GS rendering is based on the Deformable 3D GS (Yang et al. 2024), which we have extended with innovative modules (Sec. 3.2) to improve novel view synthesis. For auto labeling, we introduce an adaptor to transform camera poses and bounding boxes, generating corresponding annotations for the novel views (Sec. 3.3).

3.2 Enhanced 3D GS Rendering

3D GS is typically used for modeling static scenes and its ability to handle dynamic scenes is limited. Recent research (Yang et al. 2024; Sun et al. 2024; Wu et al. 2024; Cao and Johnson 2023) has focused on improving dynamic scene modeling. We use Deformable 3D GS (Yang et al. 2024) as

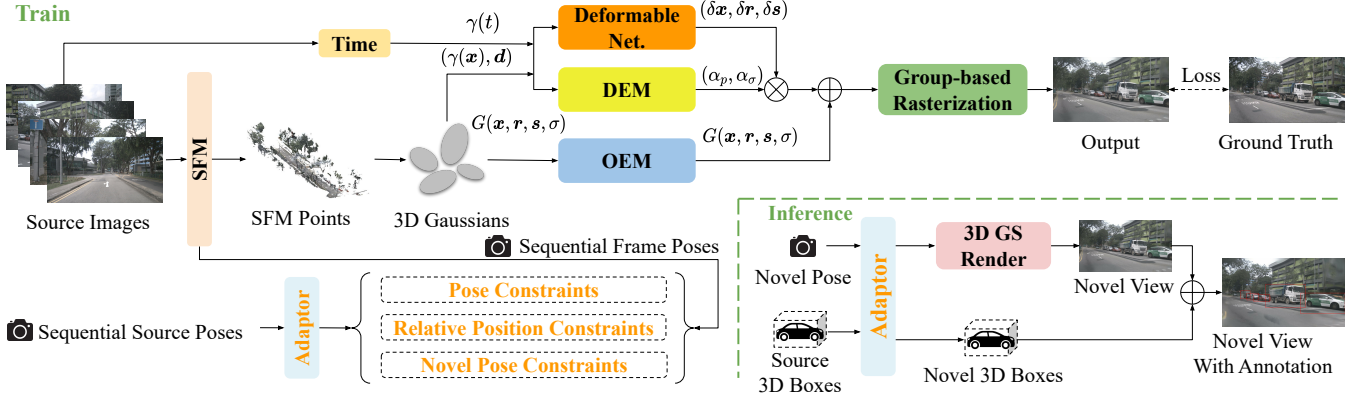


Figure 1: Illustration of overall EGSRAL. The EGSRAL framework begins by aligning the input image, followed by initializing the 3D Gaussian using the point cloud generated by the SfM. A deformable network (orange block) constructs the 3D Gaussian deformation field while the deformation enhancement module (DEM) (yellow blocks) refines this field. The opacity enhancement module (OEM) (blue blocks) optimizes opacity. To address perspective issues in large-scale, complex scenes, a group-based training and rendering strategy (green block) is employed (Section 3.2). Additionally, the adaptor is trained using three constraints (orange) to enhance its coordinate relationship modeling. During inference, these modules render and synthesize novel view images with corresponding annotations (Section 3.3).

our baseline and apply it to automatic driving scene reconstruction. This approach employs a unified deformation field for both static backgrounds and dynamic objects. To better represent each Gaussian primitive’s state, we introduce a state attribute $\mathbf{d} \in \mathbb{R}^{d \times 1}$ which implicitly indicates whether the primitive is static or dynamic. We further enhance the deformation field by incorporating a deformation enhancement module. This module decodes each Gaussian primitive’s state attribute \mathbf{d} and time encoding $\gamma(t)$ to determine the adjustment factor for the deformation field. Recognizing the importance of opacity in rendering, we also introduce an opacity enhancement module to improve the capacity.

Deformable 3D GS Network. To reduce data dependency, we use only image data for driving scene reconstruction. We begin by initializing a set of 3D Gaussians $G_0(\mathbf{x}, \mathbf{r}, \mathbf{s}, \sigma)$ through SfM, where \mathbf{x} , \mathbf{r} , \mathbf{s} , and σ represent the position of the Gaussian primitive, quaternion, scaling, and opacity, respectively. To better model dynamic 3D Gaussians, the Deformable 3D GS introduces a deformation field for the 3D Gaussians. The deformation field takes the position \mathbf{x} and time t as inputs to predict offsets $\delta\mathbf{x}$, $\delta\mathbf{r}$, and $\delta\mathbf{s}$ for \mathbf{x} , \mathbf{r} , and \mathbf{s} , respectively. Subsequently, these deformed 3D Gaussians $G_1(\mathbf{x} + \delta\mathbf{x}, \mathbf{r} + \delta\mathbf{r}, \mathbf{s} + \delta\mathbf{s}, \sigma)$ are then passed through a differentiable tile rasterizer (Kerbl et al. 2023) to render the novel image:

$$C_{\{u,v\}} = \sum_{i \in N} T_i \alpha_i c_i \quad (1)$$

where $C_{\{u,v\}}$ denotes the rendered color of the (u, v) pixel, T_i is the transmittance defined by $\prod_{j=1}^{i-1} (1 - \alpha_j)$, c_i is the color of each Gaussian primitive, and α_i is calculated by evaluating a 2D Gaussian with covariance matrix (Yifan et al. 2019) multiplied by a learned per-primitive opacity.

However, Deformable 3D GS struggles to effectively model dynamic objects and static backgrounds. We believe

this limitation stems from the inability of 3D Gaussians to adequately represent both dynamic and static elements. To address this, we introduce a state attribute \mathbf{d} for each Gaussian primitive and incorporate it as an input to the deformable network. This enhancement improves the network’s ability to capture the state of the Gaussian primitives:

$$(\delta\mathbf{x}, \delta\mathbf{r}, \delta\mathbf{s}) = \mathcal{F}_\theta(\gamma(\text{sg}(\mathbf{x})), \gamma(t), \mathbf{d}) \quad (2)$$

where \mathcal{F}_θ represents the deformable network’s parameters, $\text{sg}(\cdot)$ denotes a stop-gradient operation and γ is the positional encoding.

Deformation Enhancement Module (DEM). To further enhance the modeling of dynamic objects, we propose a deformation enhancement module that fine-tunes the deformation field based on the time t and the state attribute \mathbf{d} of the Gaussian primitive. Specifically, we use the state attribute \mathbf{d} and time encoding $\gamma(t)$ as inputs to output an adjustment factor α_p for each Gaussian primitive’s deformation field:

$$\alpha_p = \text{sigmoid}(MLP_1(\mathcal{F}_\beta(\mathbf{d}, \gamma(t), \mathbf{x} - \mathbf{c}_{view}))) \quad (3)$$

where \mathcal{F}_β represents the dynamic encoding network’s parameters, \mathbf{c}_{view} is the current camera’s coordinate position, and $\text{sigmoid}(\cdot)$ is the activation function.

Additionally, in driving scenes, we dynamically adjust the opacity of dynamic objects using an opacity adjustment factor α_σ to more effectively model their appearance and disappearance:

$$\alpha_\sigma = \tanh(MLP_2(\mathcal{F}_\beta(\mathbf{d}, \gamma(t), \mathbf{x} - \mathbf{c}_{view}))) \quad (4)$$

where $\tanh(\cdot)$ is the activation function.

Using the adjustment factors α_p and α_σ , we modify the attributes of deformed 3D Gaussians G_1 to $G_2(\mathbf{x} + \alpha_p \delta\mathbf{x}, \mathbf{r} + \delta\mathbf{r}, \mathbf{s} + \delta\mathbf{s}, \alpha_\sigma \sigma)$, enabling the deformation field to model dynamic objects and static backgrounds with greater detail.

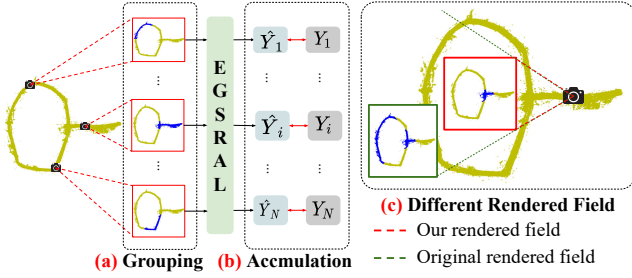


Figure 2: Illustration of the grouping strategy.

Opacity Enhancement Module (OEM). As seen in the image rendering formula (Equation 1), the rendering of the current pixel depends on the color and opacity of the Gaussian primitive. To enhance the capacity of opacity prediction, we initialize the Gaussian primitive by initializing opacity to a trainable parameter $\sigma' \in R^{16 \times 1}$ and introduce a lightweight network to accelerate opacity optimization:

$$\sigma = \left(\mathcal{F}_\eta \left(\sigma' \right) \right) \quad (5)$$

where \mathcal{F}_η is the opacity enhancement module’s parameters.

Grouping Strategy (GPS). Previous methods (Zhou et al. 2024; Yan et al. 2024) for reconstructing driving scenes using 3D GS have not accounted for large-scale scenes. To address this, we propose a grouping strategy. As illustrated in Figure 2(a), we divide the scene into N groups using a fixed image interval and assign each Gaussian primitive a group identification id , which allows us to perform subsequent cloning, splitting, and rendering based on these group ids . By grouping strategy, we solve the problem of unreasonable fields of view. As shown in Figure 2(c), the original 3D GS rendering field of view includes distant Gaussian primitives (green rectangle), which is impractical. By grouping (red rectangle), we limit the field of view to a specific range, thus reducing the optimization burden by excluding Gaussian primitives outside this range.

Unlike the static incremental training mode in Driving-Gaussian (Zhou et al. 2024), our method avoids sequential group training and the need for separate deformable networks for each group, leading to fewer models and shorter training times. However, the varying positional distribution of Gaussian primitives can affect convergence. To address this, we introduce a multi-group joint optimization strategy. As shown in Figure 2(b), each group performs forward propagation to accumulate gradients independently. After all groups have completed their forward passes, we perform gradient backpropagation to optimize the network parameters, stabilizing the training of the deformable network.

Additionally, to address the poor reconstruction quality of 3D GS in the initial frames, we propose an overlap training strategy. Specifically, for each group, we use N images from the previous group to train the current group, significantly improving reconstruction quality.

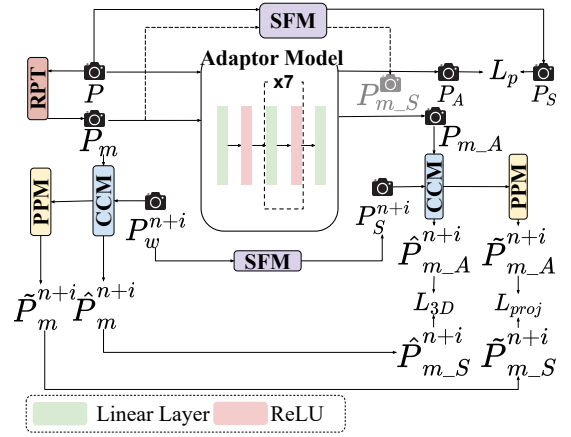


Figure 3: Illustration of the adaptor module including model and transformation modules. Camera pose $P_{m,S}$ in SfM coordinate system is the corresponding pose of novel camera pose P_m , where m denotes n_{nov} for better visualization.

3.3 Novel View Auto Labeling

Adaptor Requirement. Our method leverages image sequences (dataset) to construct scene point clouds and estimate camera poses and parameters using SfM methods. The generated point clouds and camera poses are defined in a new coordinate system generated by the SfM methods. However, the 3D annotations of the image sequences and corresponding camera poses are defined in the original world coordinate system, as in nuScenes (Caesar et al. 2020). Consequently, there exist two different coordinate systems: the original world coordinate system (OWCS) and the estimated world coordinate system (EWCS) from SfM methods.

Our renderer is trained based on the EWCS, as it takes the estimated points and camera poses in the EWCS as input. The novel camera poses used for generating novel view images are also based on the EWCS, which complicates the utilization of the dataset’s 3D annotations in the OWCS. A transformation adaptor is necessary to establish the relationship between the two coordinate systems, allowing us to utilize the 3D annotations effectively. This adaptor transforms coordinates from the OWCS to the EWCS. As a result, novel view camera poses from the OWCS can be converted into the EWCS and fed into the renderer to generate novel view images.

Our Adaptor. We take a neural network to model the transformation relationship of camera poses from the two coordinate systems, which takes camera poses in the OWCS as input and predicts the corresponding camera poses in the EWCS. We represent the camera pose as a 3×4 matrix that includes both rotation and translation information. Therefore the shapes of input and output are both 3×4 . We take MLPs to build our adaptor, as shown in Figure 3, the backbone of our adaptor is composed of 8 layers of MLP, and the output head is a simple linear layer, which is used to predict camera poses in the EWCS. To optimize our adaptor, we introduce three constraints during the training process.

Constraints for Adaptor. To effectively train the adaptor,

we need camera poses from the dataset in both coordinate systems. It is straightforward to obtain the existing camera pose pairs from these systems, where each pair consists of camera poses of the same frame from both the OWCS and the EWCS. We adopt the $smooth_{L_1}$ loss (Girshick 2015) to ensure that the adaptor’s predictions closely match the corresponding SfM camera poses in the EWCS, given the original camera poses in the OWCS. The loss function for camera pose constraints is:

$$L_p = smooth_{L_1}(P_A, P_S) \quad (6)$$

where L_p is the loss of the existing camera pose constraint, P_A is the pose predicted by the adaptor given camera pose P in the OWCS, and P_S is the corresponding pose generated by SfM in the EWCS. After applying the initial constraint, the adaptor can transform existing camera poses between the two coordinate systems. However, due to the limited number of camera pose pairs, the adaptor’s generalization ability is restricted, resulting in poor performance with novel camera poses not included in the existing pairs. To address this issue, we introduce new camera pose constraints.

The novel camera pose constraints aim to enhance the adaptor’s generalization ability by utilizing existing data to constrain new camera poses. Based on the projection consistency rule of the same objects in two camera coordinate systems, we can construct projection constraints for new camera poses. Specifically, for a camera pose P in the dataset, we use a Random Position Transformation (RPT) module to sample a nearby novel pose P_{nov} . This pose has a corresponding camera pose P_{nov-S} in the EWCS, which is implicit and not directly available. Next, we take the following N camera pose pairs after the current camera pose and project them onto the corresponding novel camera pose planes in both coordinate systems. The corresponding projected points should have the same pixel coordinates. As shown in Figure 3, the coordinate conversion module (CCM) is used to convert points (the N camera poses) in the OWCS to the camera coordinate system firstly by leveraging coordinate transformation formulas shown in Equation 7. In the coordinate system of the novel camera pose, the relative pose of P_{n+i} is represented as $\hat{P}_{n_{nov}^{n+i}}$.

$$\hat{P}_{n_{nov}^{n+i}} = (P_w^{n_{nov}})^{-1} * P_w^{n+i} \quad (7)$$

where P_w^{n+i} is the camera pose of the $frame_{n+i}$ in the OWCS, $P_w^{n_{nov}}$ is the novel camera pose, sampled nearby the current frame $frame_n$ in the OWCS. Then, we project these points to the novel camera pose plane. As shown in Figure 3, the point project module (PPM) is used to project points from the camera coordinate system to the pixel coordinate system. Next, we calculate the coordinate of $\hat{P}_{n_{nov}^{n+i}}$ in the pixel coordinate system of the novel camera pose plane by utilizing the camera intrinsic parameters, as illustrated in Equation 8 and 9 below:

$$\dot{P}_{n_{nov}^{n+i}} = K_{intr} * \begin{bmatrix} T_{\hat{P}_{n_{nov}^{n+i}}} \\ 1 \end{bmatrix} \quad (8)$$

$$\tilde{P}_{n_{nov}^{n+i}} = \begin{bmatrix} x_{\dot{P}_{n_{nov}^{n+i}}} / z_{\dot{P}_{n_{nov}^{n+i}}} \\ y_{\dot{P}_{n_{nov}^{n+i}}} / z_{\dot{P}_{n_{nov}^{n+i}}} \end{bmatrix} \quad (9)$$

where K_{intr} is camera intrinsic parameter matrix, and $T_{\hat{P}_{n_{nov}^{n+i}}}$ is the translation of $\hat{P}_{n_{nov}^{n+i}}$. Then the projection constraint can be summarized as shown in Equation 10.

$$L_{proj} = \sum_{i=1}^N MSE(\tilde{P}_{n_{nov-A}^{n+i}}, \tilde{P}_{n_{nov-S}^{n+i}}) \quad (10)$$

where $\tilde{P}_{n_{nov-A}^{n+i}}$ is the pixel coordinate of the relative camera pose $\hat{P}_{n_{nov-A}^{n+i}}$ in the EWCS projected to the novel camera pose $P_{n_{nov-A}}$ plane. $P_{n_{nov-A}}$ is predicted by our adaptor fed with novel camera pose $P_{n_{nov}}$ as shown in Figure 3. Our adaptor aims to ensure that $P_{n_{nov-A}}$ is equal to $P_{n_{nov-S}}$ by making $\tilde{P}_{n_{nov-A}^{n+i}}$ equal to $\tilde{P}_{n_{nov-S}^{n+i}}$. $\tilde{P}_{n_{nov-S}^{n+i}}$ is estimated by using the projected pixel coordinates from the corresponding camera poses in the OWCS as shown in Figure 3 based on the projection consistency rule.

Due to the limitations of pixel coordinate constraints (as in Equation 10), which only restrict the relationship between x and z and between y and z without determining their specific values, the adaptor fails to converge to the desired outcome. Including the constraints on the position in the camera coordinate system (as in Equation 11) into L_{All} effectively resolves the aforementioned issue. Given the similarity between real camera intrinsic parameters and those estimated by SfM, it is possible to assume identical position information in the camera coordinate system during the initial stages of adaptor training. Subsequently, the weight on the camera coordinate system constraints can be appropriately reduced. The overall loss is expressed as in Equation 12 below:

$$L_{3D} = \sum_{i=1}^N smooth_{L_1}(\hat{P}_{n_{nov-A}^{n+i}}, \hat{P}_{n_{nov-S}^{n+i}}) \quad (11)$$

$$L_{all} = w_1 * L_p + w_2 * L_{3D} + w_3 * L_{proj} \quad (12)$$

where $\hat{P}_{n_{nov-A}^{n+i}}$ is the 3D coordinate of the relative camera pose of camera pose P_S^{n+i} in the EWCS relative to novel camera pose $P_{n_{nov-A}}$, as shown in Figure 3. $\hat{P}_{n_{nov-S}^{n+i}}$ is the estimated 3D coordinate of the relative camera pose of the camera pose P_S^{n+i} in the EWCS, relative to novel camera pose $P_{n_{nov-S}}$, obtained using the 3D coordinates from the corresponding camera poses in the OWCS. L_{3D} is the loss of novel camera pose constraints in 3D coordinates, L_{proj} is the loss of novel camera pose constraints in project pixel coordinates, w_1 is the weight of L_p , w_2 is the weight of L_{3D} , w_3 is the weight of L_{proj} .

Annotation Generation. During the inference phase, as illustrated in Figure 1, we start with the original camera pose P_{ori} in the OWCS from the dataset. This pose undergoes an affine random transformation to obtain P_{nov} by the RPT module shown in Figure 3. Then we feed P_{nov} into the adaptor to produce the corresponding pose in the EWCS, denoted as P_{nov-A} . Feeding P_{nov-A} into the renderer generates the corresponding novel view image, called $image_{nov}$. Similarly, applying the same affine transformation to the original dataset’s annotations $anns_{ori}$ results in $anns_{nov}$, which corresponds to $image_{nov}$. Combining $image_{nov}$ with $anns_{nov}$ results in a novel view image with annotations, effectively accomplishing auto labeling.

Methods	Supervision	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow
NPBG	RGB + Points	19.58	0.248	0.628
ADOP	RGB + Points	20.08	0.183	0.623
Deepview	RGB	17.28	0.196	0.716
MPI	RGB	19.54	0.158	0.733
READ	RGB	23.48	0.132	0.787
3D GS	RGB	20.37	0.357	0.679
Mip-splatting	RGB	20.86	0.341	0.696
Deformable 3D GS	RGB	21.33	0.306	0.713
EGSRAL(Ours)	RGB	23.60	0.192	0.787

Table 1: Comparisons with existing methods on KITTI City dataset.

Methods	Input	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow
S-NeRF	RGB+LiDAR	25.43	0.302	0.730
SUDS	RGB+LiDAR	21.26	0.466	0.603
EmerNeRF	RGB+LiDAR	26.75	0.311	0.760
3D GS	RGB+SfM	26.08	0.298	0.717
DrivingGaussian-L	RGB+LiDAR	28.74	0.237	0.865
EGSRAL(Ours)	RGB+LiDAR	29.04	0.162	0.883

Table 2: Comparisons with existing methods on the nuScenes dataset used in DrivingGaussian.

4 Experiment

4.1 Datasets

KITTI Dataset. KITTI dataset (Geiger, Lenz, and Urtasun 2012) contains a wide range of driving scenes. Our experiments focus on a single distinct scene: City, which consists of 1424 train images and 159 test images.

NuScenes Dataset. NuScenes (Caesar et al. 2020) stands as a comprehensive autonomous driving dataset, featuring 3D bounding boxes for 1000 scenes. We follow the different scenes selected by S-NeRF (Xie et al. 2023) and DrivingGaussian (Zhou et al. 2024) for our experiments.

4.2 Experimental Settings

Implementation Details. Our implementation is primarily based on Deformable 3D GS (Yang et al. 2024), maintaining consistency with its training iterations and optimizer settings. Additionally, the DEM and the OEM are optimized using Adam with a learning rate of 0.001.

Adaptor Model Training Settings. Our training dataset is nuScenes poses P and poses P_S generated by SfM as data pairs. We need to train on 17 different scenes, each with approximately 230 pose pairs. Each scene requires training for 1000 epochs. We use a Adam optimizer with a learning rate of $2e-4$ and a batch size of 16. It takes about 2 hours to train using AMD MI100. Adaptor model loss weights w_1 , w_2 and w_3 in Equation 12 are 50, 0.1 and 1, respectively, and the N camera pose is 15.

2D/3D Detection Settings. To verify the effectiveness of synthetic data, we use mainstream detection model to perform 2D and 3D detection tasks on four categories: car, bus, truck, and trailer. 2D detection tasks use Co-DETR (Zong, Song, and Liu 2023) and 3D tasks use MonoLSS (Li, Jia, and Shi 2024). For the test dataset, we select the first fifteen scene data in the nuScenes validation set.

Methods	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow
Mip-NeRF	18.22	0.421	0.655
Mip-NeRF360	24.37	0.240	0.795
Urban-NeRF	21.49	0.491	0.661
S-NeRF	26.21	0.228	0.831
3D GS	32.82	0.225	0.925
Mip-splatting	32.22	0.224	0.928
Deformable 3D GS	33.43	0.224	0.932
EGSRAL(Ours)	34.43	0.205	0.939

Table 3: Comparisons with existing methods on the nuScenes dataset used in S-NeRF.

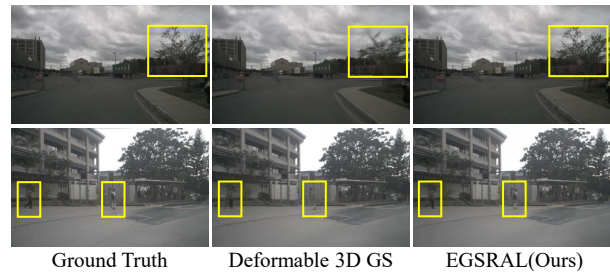


Figure 4: Qualitative comparison of novel view synthesis on the nuScenes dataset.

4.3 Performance of Novel View Synthesis

Evaluation on KITTI. We compare our method with state-of-the-art (SOTA) methods, as shown in Table 1. Our approach demonstrates superior performance compared to other 3D GS methods, with PSNR and SSIM improvements of 2.27 and 0.074, respectively, and 0.114 reduction in LPIPS over Deformable 3D GS. Additionally, compared to other methods like READ, our method enhances the PSNR metric to 23.60 while achieving comparable SSIM results. However, the 3D GS-based method has a slight deficiency in modeling high-frequency information, such as ground texture, which leads to a higher LPIPS. For detailed visualization results, please refer to the supplementary materials.

Evaluation on nuScenes. Following the scenes used by DrivingGaussian, as shown in Table 2, our method achieves the highest performance, with an increase of 0.3 and 0.018 in PSNR and SSIM, respectively, and a decrease of 0.075 in the LPIPS. It is worth noting that DrivingGaussian trains the dynamic objects and the static backgrounds separately, training an individual 3D GS model for each dynamic object then merging them. In contrast, our method requires only a single training process and does not need any annotations, resulting in better performance and reducing the dependence on data annotations.

Following the scenes used by S-NeRF (Xie et al. 2023), as shown in Table 3, we contrast our method with the SOTA method. In addition, we replicated 3D GS (Kerbl et al. 2023), Deformable 3D GS (Yang et al. 2024), and mip-splatting (Yu et al. 2024) on this dataset. The 3D GS-based method has much higher metrics on these scenes than the NeRF-based method. Our method achieves the best performance, with PSNR improved by 1.00, SSIM improved by

Model	Dataset	Total_amount	mAP	AP50	AP75
Co-DETR	sample	1×	23.7	44.6	22.9
		2×	25.9	45.9	23.1
		3×	26.8	46.1	23.2
	all	1×	28.1	46.5	24.2
		2×	29.9	47.7	24.8
		3×	31.3	48.2	25.3

Table 4: Augmentation results of 2D detection on nuScenes.

Model	Dataset	Total_amount	AP3d	Mod.	Hard
MonoLSS	sample	1×	17.21	12.55	11.30
		2×	19.78	13.32	12.81
		3×	20.31	13.85	12.97
	all	1×	21.69	14.27	13.95
		2×	22.15	14.93	14.22
		3×	22.87	15.25	14.40

Table 5: Augmentation results of 3D detection on nuScenes.

0.007, and LPIPS reduced by 0.019 compared to the Deformable 3D GS.

In Figure 4, we compare the view synthesis results of our method with those of other 3D GS methods. The visualization results show superior synthesis capabilities, especially achieving more realistic rendering results on static backgrounds (trees) and dynamic objects (pedestrians).

4.4 Performance of Auto Labeling

We select 17 different scenes from the nuScenes dataset for the 2D/3D auto-labeling task. Using the adaptor, we generate novel view images and automatically label them simultaneously. The results are presented in Figure 5. The three columns represent three different scenes: the first row displays the original images, while the last two rows show the novel view images with auto labeling. In these images, the yellow boxes denote 2D annotations, and the orange boxes denote 3D annotations. The results demonstrate that our adaptor performs effectively across different scenes.

Performance of 2D/3D Object Detection. We select 17 different scenes containing 3,898 images (674 of which are sample images with annotations) from the nuScenes dataset as the baseline dataset for the detection task. We produce the annotations for the non-sample images based on the annotations of sample images by camera coordinate system transformation. To augment the dataset, we synthesize additional images using an adaptor, increasing the total number of images to two and three times that of the baseline dataset. The 2D detection results are presented in Table 4, and the 3D detection results are presented in Table 5. In these tables, *sample* refers to the 674 sample images, while *all* refers to the 3,898 images, including both sample and non-sample images. The results demonstrate that, without any additional input, our data augmentation method improves the accuracy of the detection model in both the *sample* and the *all* dataset.



Figure 5: Visualizing 2D/3D auto labeling on nuScenes.

DEM	OEM	GPS	PSNR↑	LPIPS↓	SSIM↑
-	-	-	21.33	0.306	0.713
✓	-	-	22.37	0.279	0.743
-	✓	-	22.32	0.277	0.738
✓	✓	-	22.84	0.253	0.760
✓	✓	✓	23.60	0.192	0.787

Table 6: Effect of each module in our method.

4.5 Ablation Study and Analysis

In this section, we provide ablation experiments for different modules of our method on the KITTI City dataset, as shown in Table 6. For more experiments and analyses, please see Appendix Section 3. Comparing rows 2 and 3, we observe that using the DEM alone to refine the deformation field improves PSNR by 1.04 and SSIM by 0.03. In contrast, using the OEM alone results in improvements of 0.99 in PSNR and 0.025 in SSIM, with a reduction of 0.029 in LPIPS. When both DEM and OEM are combined, the evaluation metrics show significant improvements over Deformable 3D GS, demonstrating the enhanced synthesis capabilities of our method for large-scale scene modeling. Additionally, we validate our proposed grouping strategy. As shown in the last row, our approach further enhances PSNR by 0.76, improves SSIM by 0.027, and reduces LPIPS by 0.061 by incorporating the grouping strategy. In conclusion, the above ablation studies fully validate the effectiveness of all proposed modules in our method.

5 Conclusion

In this paper, we present EGSRAL, a novel 3D GS-based renderer with an automated labeling framework capable of synthesizing novel view images with corresponding annotations. For novel view rendering, we introduce two effective modules to improve the ability of 3D GS to model complex scenes and propose a grouping strategy to solve the problem of unreasonable view of large-scale scenes. For novel view auto labeling, we propose an adaptor to generate new annotations for novel views. Experimental results show that EGSRAL significantly outperforms existing methods in novel view synthesis and achieves superior object detection performance on annotated images.

References

- Agisoft. 2019. Agisoft: Metashape software]. retrieved 20.05.2019 (2019).
- Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2020. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, 11621–11631.
- Cao, A.; and Johnson, J. 2023. Hexplane: A fast representation for dynamic scenes. In *CVPR*, 130–141.
- Chang, D.; Božič, A.; Zhang, T.; Yan, Q.; Chen, Y.; Süsstrunk, S.; and Nießner, M. 2022. RC-MVSNet: unsupervised multi-view stereo with neural rendering. In *ECCV*, 665–680. Springer.
- Chen, S.-C. 2019. Multimedia for autonomous driving. *IEEE MultiMedia*, 5–8.
- Fu, X.; Zhang, S.; Chen, T.; Lu, Y.; Zhu, L.; Zhou, X.; Geiger, A.; and Liao, Y. 2022. Panoptic nerf: 3d-to-2d label transfer for panoptic urban scene segmentation. In *3DV*, 1–11. IEEE.
- Ge, Y.; Behl, H.; Xu, J.; Gunasekar, S.; Joshi, N.; Song, Y.; Wang, X.; Itti, L.; and Vineet, V. 2022. Neural-sim: Learning to generate training data with nerf. In *ECCV*, 477–493. Springer.
- Geiger, A.; Lenz, P.; and Urtasun, R. 2012. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 3354–3361. IEEE.
- Girshick, R. 2015. Fast r-cnn. In *ICCV*, 1440–1448.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *CVPR*, 1125–1134.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.*, 139–1.
- Kundu, A.; Genova, K.; Yin, X.; Fathi, A.; Pantofaru, C.; Guibas, L. J.; Tagliasacchi, A.; Dellaert, F.; and Funkhouser, T. 2022. Panoptic neural fields: A semantic object-aware neural scene representation. In *CVPR*, 12871–12881.
- Li, Z.; Jia, J.; and Shi, Y. 2024. MonoLSS: Learnable Sample Selection For Monocular 3D Detection. In *3DV*, 1125–1135. IEEE.
- Li, Z.; Li, L.; and Zhu, J. 2023. Read: Large-scale neural scene rendering for autonomous driving. In *AAAI*, 1522–1529.
- Li, Z.; Wu, C.; Zhang, L.; and Zhu, J. 2023. DGNR: Density-Guided Neural Point Rendering of Large Driving Scenes.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 99–106.
- Schonberger, J. L.; and Frahm, J.-M. 2016. Structure-from-motion revisited. In *CVPR*, 4104–4113.
- Sun, J.; Jiao, H.; Li, G.; Zhang, Z.; Zhao, L.; and Xing, W. 2024. 3dstream: On-the-fly training of 3d gaussians for efficient streaming of photo-realistic free-viewpoint videos. In *CVPR*, 20675–20685.
- Tong, W.; Xie, J.; Li, T.; Deng, H.; Geng, X.; Zhou, R.; Yang, D.; Dai, B.; Lu, L.; and Li, H. 2023. 3D Data Augmentation for Driving Scenes on Camera.
- Tosi, F.; Tonioni, A.; De Gregorio, D.; and Poggi, M. 2023. NeRF-Supervised Deep Stereo. In *CVPR*, 855–866.
- Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018. High-resolution image synthesis and semantic manipulation with conditional gans. In *CVPR*, 8798–8807.
- Wu, G.; Yi, T.; Fang, J.; Xie, L.; Zhang, X.; Wei, W.; Liu, W.; Tian, Q.; and Wang, X. 2024. 4d gaussian splatting for real-time dynamic scene rendering. In *CVPR*, 20310–20320.
- Wu, Z.; Liu, T.; Luo, L.; Zhong, Z.; Chen, J.; Xiao, H.; Hou, C.; Lou, H.; Chen, Y.; Yang, R.; et al. 2023. Mars: An instance-aware, modular and realistic simulator for autonomous driving.
- Xie, Z.; Zhang, J.; Li, W.; Zhang, F.; and Zhang, L. 2023. S-nerf: Neural radiance fields for street views.
- Xu, C.; Wu, B.; Hou, J.; Tsai, S.; Li, R.; Wang, J.; Zhan, W.; He, Z.; Vajda, P.; Keutzer, K.; et al. 2023. Nerf-det: Learning geometry-aware volumetric representation for multi-view 3d object detection. In *ICCV*, 23320–23330.
- Yan, Y.; Lin, H.; Zhou, C.; Wang, W.; Sun, H.; Zhan, K.; Lang, X.; Zhou, X.; and Peng, S. 2024. Street gaussians for modeling dynamic urban scenes.
- Yang, Z.; Gao, X.; Zhou, W.; Jiao, S.; Zhang, Y.; and Jin, X. 2024. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In *CVPR*, 20331–20341.
- Yifan, W.; Serena, F.; Wu, S.; Öztireli, C.; and Sorkine-Hornung, O. 2019. Differentiable surface splatting for point-based geometry processing. *TOG*, 1–14.
- Yu, Z.; Chen, A.; Huang, B.; Sattler, T.; and Geiger, A. 2024. Mip-Splatting: Alias-free 3D Gaussian Splatting. In *CVPR*, 19447–19456.
- Zhang, J.; Zhang, F.; Kuang, S.; and Zhang, L. 2023. NeRF-LiDAR: Generating Realistic LiDAR Point Clouds with Neural Radiance Fields.
- Zhi, S.; Laidlow, T.; Leutenegger, S.; and Davison, A. J. 2021. In-place scene labelling and understanding with implicit scene representation. In *ICCV*, 15838–15847.
- Zhou, X.; Lin, Z.; Shan, X.; Wang, Y.; Sun, D.; and Yang, M.-H. 2024. Drivinggaussian: Composite gaussian splatting for surrounding dynamic autonomous driving scenes. In *CVPR*, 21634–21643.
- Zong, Z.; Song, G.; and Liu, Y. 2023. DETRs with Collaborative Hybrid Assignments Training. In *ICCV*, 6748–6758.