

Nighttime Autonomous Driving Scene Reconstruction with Physically-Based Gaussian Splatting

Tae-Kyeong Kim^{1,2}, Xingxin Chen¹, Guile Wu¹, Chengjie Huang¹, Dongfeng Bai¹, and Bingbing Liu¹

Abstract—This paper focuses on scene reconstruction under nighttime conditions in autonomous driving simulation. Recent methods based on Neural Radiance Fields (NeRFs) and 3D Gaussian Splatting (3DGS) have achieved photorealistic modeling in autonomous driving scene reconstruction, but they primarily focus on normal-light conditions. Low-light driving scenes are more challenging to model due to their complex lighting and appearance conditions, which often causes performance degradation of existing methods. To address this problem, this work presents a novel approach that integrates physically based rendering into 3DGS to enhance nighttime scene reconstruction for autonomous driving. Specifically, our approach integrates physically based rendering into composite scene Gaussian representations and jointly optimizes Bidirectional Reflectance Distribution Function (BRDF) based material properties. We explicitly model diffuse components through a global illumination module and specular components by anisotropic spherical Gaussians. As a result, our approach improves reconstruction quality for outdoor nighttime driving scenes, while maintaining real-time rendering. Extensive experiments across diverse nighttime scenarios on two real-world autonomous driving datasets, including nuScenes and Waymo, demonstrate that our approach outperforms the state-of-the-art methods both quantitatively and qualitatively.

I. INTRODUCTION

With the rapid development of autonomous driving, the demand for safety continues to increase. Creating digital twins of driving scenes has gained increasing attention because these digital twins can be used to simulate safety-critical corner cases that are difficult and costly to capture in the real world. By performing closed-loop evaluation with these reconstructed driving scenes, the reliability of autonomous driving systems can be continuously improved.

Contemporary driving scene reconstruction methods can be categorized into Neural Radiance Fields (NeRF) based methods and 3D Gaussian Splatting (3DGS) based methods. NeRF-based methods implicitly model driving scenes with neural scene graphs and multi-layer perceptrons (MLPs). Although NeRF-based methods have achieved photorealistic scene modeling [1]–[4], they inherently suffer from slow training and rendering speeds. On the other hand, 3DGS-based methods explicitly employ composite Gaussian primitives and tile-based differentiable rasterization for driving scene reconstruction [5]–[8]. They have shown promising results for modeling dynamic driving scenes, while maintaining real-time rendering performance. However, existing

methods primarily focus on driving scene reconstruction under normal-light conditions, neglecting more challenging nighttime driving scene reconstruction. Consequently, existing methods often suffer from performance degradation for nighttime scene reconstruction due to complex lighting and appearance conditions under low-light conditions. Although scene reconstruction under low-light conditions [9]–[11] has been explored in some other fields, these methods are mostly designed for static scenes and cannot be directly used for dynamic driving scene modeling. In addition, their reliance on environment maps may lead to failures in urban driving scenes due to complex lighting and appearance conditions.

To address these problems, this work presents a novel approach that integrates physically based rendering into composite scene Gaussian representations for autonomous driving scene reconstruction. Specifically, our approach disentangles diffuse and specular component modeling and jointly optimizes Bidirectional Reflectance Distribution Function (BRDF) based material properties to enhance nighttime scene reconstruction. Unlike prior work that relies on environment map sampling, we propose a global illumination module to model the diffuse component. For specular component, we employ anisotropic spherical Gaussians (ASGs) [12] to better capture high-frequency specular effects and equip each Gaussian with BRDF-based material properties. Then, we solve a physically based rendering function to obtain color contribution of each Gaussian and map High Dynamic Range (HDR) colors to Low Dynamic Range (LDR) colors before differentiable rasterization. In this way, our approach is capable of improving outdoor nighttime driving scene reconstruction while maintaining real-time rendering. Our experiments on two real-world autonomous driving datasets, including nuScenes [13] and Waymo [14], demonstrate that our approach achieves the state-of-the-art performance compared with existing methods for nighttime driving scene reconstruction. In summary, our contributions are:

- We propose a novel method to model nighttime driving scene by integrating physically based rendering into composite scene Gaussian representations, which fills a gap left by existing works.
- We design a global illumination module to model diffuse component which does not require per-ray environment map sampling and is suitable for dynamic scene modeling.
- We combine ASGs with BRDF-constrained rendering for specular component modeling to capture high-frequency effects while ensuring physical plausibility.

¹The authors are with Huawei Noah's Ark Lab. ²Tae-Kyeong Kim is with the University of Toronto and contributed to this work during an internship at Huawei Noah's Ark Lab, Canada. taekyeong.kim@mail.utoronto.ca. {xingxin.chen, baidongfeng, liu.bingbing}@huawei.com, {guile.wu, chengjie.huang}@outlook.com.

II. RELATED WORKS

A. Gaussian Splatting for Driving Scene Reconstruction

Recently, Gaussian splatting has been widely used for driving scene reconstruction in autonomous driving simulation due to its faster training and real-time rendering efficiency. To model dynamic driving scenes, researchers often employ composite scene Gaussian representations, which decompose scenes into dynamic object nodes, static background nodes and distant region nodes in Gaussian scene graphs [5]–[8], [15], [16]. OmniRe [5] constructs a Gaussian scene graph of diverse Gaussian node types to unify reconstruction of rigid vehicles, deformable agents, and non-rigid backgrounds, which achieves state-of-the-art dynamic reconstruction for autonomous driving. StreetGS [6] introduces decomposition of urban environments into background, foreground actors, and sky, using tracked poses and cubemaps for large-scale rendering. HUGS [17] and ArmGS [7] integrate appearance modeling into composite scene Gaussian representations to model camera exposure and dynamic scene conditions. UniGaussian [16] introduces multi-modal multi-sensor simulation using composite scene Gaussian representations. However, these works primarily focus on driving scene reconstruction under normal-light conditions, leaving low-light (nighttime) scenarios largely underexplored. This motivates us to develop a novel approach that remains effective under challenging nighttime conditions. To this end, our work adopts a Gaussian scene graph for dynamic driving scene reconstruction and integrates physically based rendering into composite scene Gaussian representations, which enables effective modeling of complex illumination effects in nighttime environments.

B. Lighting Modeling for Scene Reconstruction

Lighting modeling for scene reconstruction has been actively explored to improve reconstruction quality under complex illumination. One line of work incorporates physically based rendering (PBR), which is widely used in computer graphics for physically plausible relighting. NeILF [18] first introduces per-scene material parameters and enables relighting via the Simplified Disney BRDF model [19]. Building on this idea, RelightableGS [11] extends PBR to 3DGS, assigning material attributes (albedo, roughness, metallic) to each Gaussian, and use a point-based rasterizer to produce the final image from the shaded Gaussians. However, these relighting methods are designed mainly for static scenes. They usually require per-ray sampling of environment maps or incident lighting at each rendering step, which slows down inference compared to standard SH-based 3DGS. In addition, they often rely on extended shading pipelines rather than the lightweight rasterization used in vanilla 3DGS. More recently, Spec-Gaussian [20] employs per-Gaussian ASGs for modeling sharper highlights. DarkGS [21] and LL-Gaussian [22] enhance the robustness of 3DGS in low-light scenarios. However, these methods still focus on static scenes and cannot be directly used for dynamic driving scene modeling. Different from these works, we integrate

physically based rendering into composite scene Gaussian representations and present a global illumination module to model diffuse component and BRDF-constrained rendering with ASGs to model specular component for autonomous driving simulation.

III. METHODS

A. Preliminaries

1) *Gaussian Splatting*: 3D Gaussian Splatting (3DGS) [23] represents a scene as a collection of colorized 3D Gaussian primitives. Each Gaussian is parameterized by opacity $o \in [0, 1]$, position $\mu \in \mathbb{R}^3$, covariance matrix Σ (derived from a quaternion rotation $q \in \mathbb{R}^4$ and scaling vector $s \in \mathbb{R}^3$), and color $c \in \mathbb{R}^F$, often expressed as spherical harmonics (SH) coefficients. These parameters define the Gaussian in 3D space as the following:

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

During rendering stage, the 3D Gaussians are projected onto the image plane as 2D Gaussians. The color of each pixel is then computed by alpha-blending N-ordered 2D Gaussians:

$$C = \sum_{i \in N} c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (2)$$

where c_i is the color derived from SH coefficients and α_i is determined by the opacity of the Gaussian.

2) *Composite Scene Graph Structure*: Following previous work [5]–[8], [15], [16], we construct a 3DGS scene graph that consists of three components: (i) static background, (ii) dynamic actors, and (iii) sky map. Background Gaussians are initialized from LiDAR points. Rigid objects (e.g., vehicles) are represented by Gaussian nodes whose positions are updated according to the vehicle trajectory $T_v \in \text{SE}(3)$. Non-rigid objects are modeled as SMPL Gaussians for pedestrians [5] and deformable Gaussians for cyclists and distant pedestrians [24]. The sky region is modeled using an explicit cubemap representation, which is sampled on given viewing direction. While this scene structure provides a robust representation of autonomous driving scenes, our method further integrates a lighting module to achieve photorealistic reconstructions under challenging low-light conditions.

B. Approach Overview

An overview of our proposed framework for nighttime autonomous driving scene reconstruction is shown in Figure 1. Our method reconstructs dynamic nighttime driving scenes by integrating physically based rendering with a composite 3DGS scene graph. Specifically, for a given camera and timestep, we decompose the Gaussian appearance attributes into diffuse and specular components: a global lighting module predicts SH coefficients to model the diffuse illumination, while per-Gaussian ASGs capture the incident specular lighting. The two components are synergically combined with additional material attributes to produce HDR radiance, which is tone-mapped to LDR via Reinhard and rendered with the standard differentiable 3DGS rasterizer.

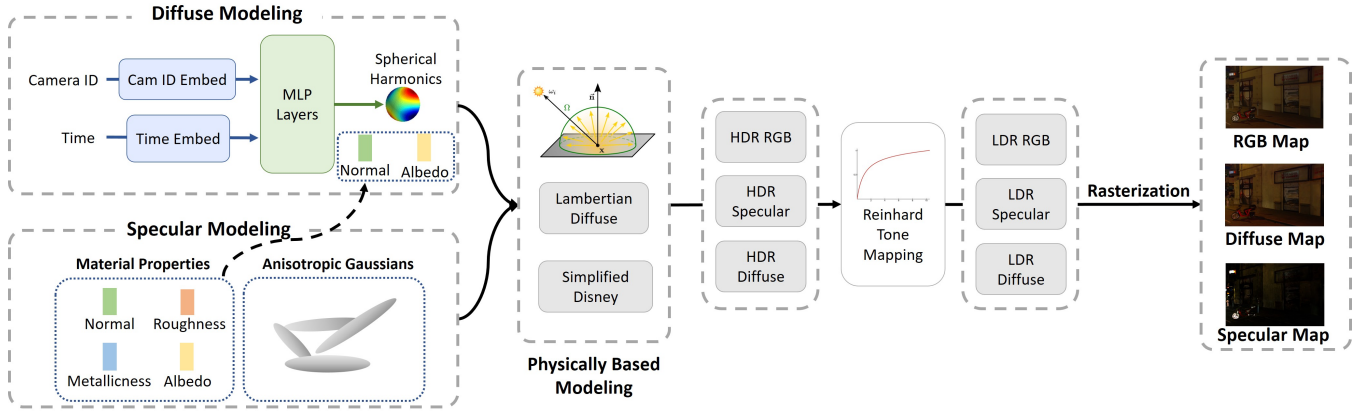


Fig. 1. An overview of our approach. Our framework decomposes lighting into specular and diffuse component. Each of the respective component is modeled using the per-Gaussian ASGs and global SH module that are constrained by the BRDF

C. BRDF Modeling

Our lighting method utilizes the physically based rendering equation [25] to compute the interaction between the lighting of the scene and the Gaussian surfaces. The rendering equation is defined as follows:

$$L_o(\omega_o, \mathbf{x}) = \int_{\Omega} f(\omega_o, \omega_i, \mathbf{x}) L_i(\omega_i, \mathbf{x}) (\omega_i \cdot \mathbf{n}) d\omega_i, \quad (3)$$

where x denotes a surface point with normal vector n , f is the BRDF function, and L_i and L_o represent the incoming and outgoing radiance in directions ω_i and ω_o , respectively. The domain Ω corresponds to the hemisphere above the surface. Following [11], each Gaussian primitive is parameterized as $g(\mu, q, s; b, r, m, n)$, where in addition to the mean μ and covariance (represented by rotation quaternion q and scale s), each Gaussian primitive is assigned material attributes and a surface normal:

- $b \in [0, 1]^3$: albedo (base color),
- $r \in [0, 1]$: roughness,
- $m \in [0, 1]$: metallicness,
- $n \in \mathbb{R}^3$: surface normal.

D. Global Lighting Module

Previous work [11], [18] models global incident light with environment maps, which requires sampling for every rendering step. Additionally, they are primarily tested for static scenes, and building a compact robust environment map for dynamic scenes is challenging. To address these limitations, we propose a global lighting module that predicts scene illumination conditioned on normalized timestep and camera embeddings. Specifically, we use an MLP to produce a latent representation, which is passed through a detection head for each SH level. In our design, we employ second-degree SH, producing three sets of coefficients. These SH coefficients are combined with the normal n and albedo b properties of the Gaussians, which are then evaluated using the Lambertian diffuse equation [26]. Following [27], the diffuse contribution can be approximated as a weighted

summation over SH coefficients:

$$L_d = \frac{b}{\pi} \sum_{l=0}^2 A_l \sum_{m=-l}^l c_m^l Y_m^l(n), \quad (4)$$

where A_l is the cosine-kernel convolution factors that convert incident radiance SH into irradiance SH for a Lambertian, Y_l^m represents the real SH basis evaluated at Gaussian normal, and c_m^l is the SH coefficients produced by the environment lighting module.

E. Specular Light Modeling

Nighttime driving scenes often contain artificial lighting and intense reflections from retro reflectors. To better handle these situations, we employ ASGs [12] within our rendering pipeline, as they have been shown to capture specular highlights more effectively than SH-based approaches [11]. Following [12], an ASG is defined as:

$$G(v; [x, y, z], [\lambda, \mu], c) = c \cdot S(v; z) \cdot e^{-\lambda(v \cdot x)^2 - \mu(v \cdot y)^2}, \quad (5)$$

where x, y, z represents the lobe direction, λ, μ denotes the lobe sharpness in the respective x and y coordinates, and S is the saturation between the viewing direction v and the lobe z axis. ASGs can capture specular highlights more effectively than previous approaches based on SH coefficients to approximate incident specular lighting. Instead of using ASGs as a latent representation passed to an MLP for the final RGB contribution like [20], our framework uses ASGs to model incident specular lighting. Each Gaussian is assigned a small set of ASG lobes (four in our case), and their specular contribution is evaluated during rendering via a BRDF-constrained PBR. Following [11], [18], we use a simplified Disney BRDF representation:

$$f_s(w_o, w_i) = \frac{D(h; r) \cdot F(w_o, h; b, m) \cdot G(w_i, w_o, h; r)}{(n \cdot w_i) \cdot (n \cdot w_o)}, \quad (6)$$

where the h represents the half vector between incoming and outgoing radiance, and D , F and G indicate the normal distribution function, the Fresnel term, and the geometry term respectively. As the normal distribution can be modeled as

a spherical Gaussian (SG) [11], [18], the convolution of the SG and ASG product inside the PBR can be simplified [12]:

$$L_s = G(\omega_i, \omega_0, h; r) \cdot F(\omega_o, h; b, m) \cdot \sum_i A_i \cdot s_i \cdot \text{ASG}_i \left(w_r, [x_i, y_i, z_i], \left[\frac{\nu \lambda_i}{\nu + \lambda_i}, \frac{\nu \mu_i}{\nu + \mu_i} \right], a_{\text{ndf}} \frac{\pi}{\sqrt{(\nu + \lambda_i)(\nu + \mu_i)}} \right), \quad (7)$$

where ν and a_{ndf} represent concentration parameter and amplitude of the normal distribution, and w_r denotes the reflection direction between incoming radiance and the Gaussian normal. With this simplification, our approach can efficiently apply specular lighting to every Gaussian without any additional sampling.

F. Final Lighting and Tone Mapping

The diffuse and specular lighting is summed together to produce the final relighted value for each Gaussian in HDR:

$$L_{\text{HDR}} = L_d + L_s. \quad (8)$$

Since the image pixel values are expressed in LDR values, we apply Reinhard tone mapping [28] to produce LDR values:

$$L_{\text{LDR}} = \frac{L_{\text{HDR}}}{1 + L_{\text{HDR}}}. \quad (9)$$

The resulting L_{LDR} serves as the per-Gaussian color c_i used for the standard rasterization.

G. Optimization

In order to obtain realistic lighting effects, accurate surface normals are crucial for each Gaussian. We employ an off-the-shelf model [29] to obtain a normal map prior N for the input image. During rasterization, the per-Gaussian normals n_i are accumulated to produce the rendered normal map:

$$\hat{N} = \frac{\sum_{i \in M} n_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j)}{\sum_{i \in M} \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j)}, \quad (10)$$

where α_i denotes the opacity contribution of the i -th Gaussian along the ray, and M is the set of overlapping Gaussians. To improve normal supervision with these priors, we follow the depth-normal loss regularizer from [30]:

$$L_n = \|\hat{N} - N\|_1 + (1 - \hat{N} \cdot N), \quad (11)$$

Additionally, we derive depth-based normals \hat{N}_d from the rendered depth map and compute a confidence weight to measure the consistency between \hat{N}_d and the prior normal N :

$$w = \exp \left(\frac{\hat{N}_d \cdot N - 1}{\gamma} \right), \quad (12)$$

where γ is a hyperparameter controlling the sharpness of the weighting. Using this, we define the view-consistent D-Normal regularizer [30] as:

$$L_{dn} = w \cdot \left(\|\hat{N}_d - N\|_1 + (1 - \hat{N}_d \cdot N) \right). \quad (13)$$

The overall training objective combines the depth-normal loss with the standard RGB pixel and D-SSIM loss [5], [6], [17], which can be formulated as:

$$L_{\text{total}} = w_{\text{rgb}} \cdot L_{\text{rgb}} + w_{\text{D-SSIM}} \cdot L_{\text{D-SSIM}} + w_{dn} \cdot (L_d + L_n). \quad (14)$$

IV. EXPERIMENT

A. Datasets

We conducted our experiments on two challenging real-world autonomous driving datasets, including nuScenes [13] and the Waymo Open Dataset [14]. The Waymo Open Dataset contains high-resolution images with challenging nighttime scenes, while the nuScenes dataset has lower resolution but provides diverse urban scenes with low ambient lighting and sharp highlights. In this work, we consider *low-light dynamic scenes* as those that exhibit low ambient illumination from the sky and contain sharp localized lighting sources (e.g., headlights, buildings, street lamps), which reflects real-world challenges for autonomous driving in low-light conditions. We therefore select 11 low-light driving scenes from each datasets for our experiments. Specifically, we use scenes 754, 763, 774, 776, 780, 781, 782, 783, 784, 785, and 790 for nuScenes, and scenes 007, 012, 015, 018, 030, 038, 051, 099, 106, 129, and 166 for Waymo. For evaluation, every eighth frame is reserved for novel view testing, while the remaining frames are used for training.

B. Implementation Details

In our experiments, each driving scene is initialized with LiDAR point clouds, supplemented by uniformly sampled points following [15]. The model is trained for 40,000 iterations using the Adam optimizer [31]. The learning rate for normals, roughness, and metallicness is set to 1×10^{-3} , while the learning rate for Gaussian lobe axes, stretch parameters along the x and y coordinates, and amplitude is set to 1×10^{-5} . The global SH illumination module consists of 8 linear layers, and each latent representation is passed through a single linear head to produce the SH coefficients at each level. For training, the loss weights are empirically set as $\omega_{\text{rgb}} = 0.8$, $\omega_{\text{D-SSIM}} = 0.2$, and $\omega_{\text{LPIPS}} = 0.025$.

C. Baselines

We compare our framework against recent state-of-the-art 3DGS methods for dynamic urban scene reconstruction: OmniRe [5] and StreetGaussians [6]. We also include 3DGS [23] as a baseline for comparison. Note that, since previous work [5], [6] has demonstrated the superiority of 3DGS-based method over NeRF-based methods, we do not include comparison with NeRF-based methods in this work. For evaluation metrics, following [5], [6], we evaluate PSNR for pixel-level, SSIM for structural-level [32], and LPIPS [33] for perceptual similarity.

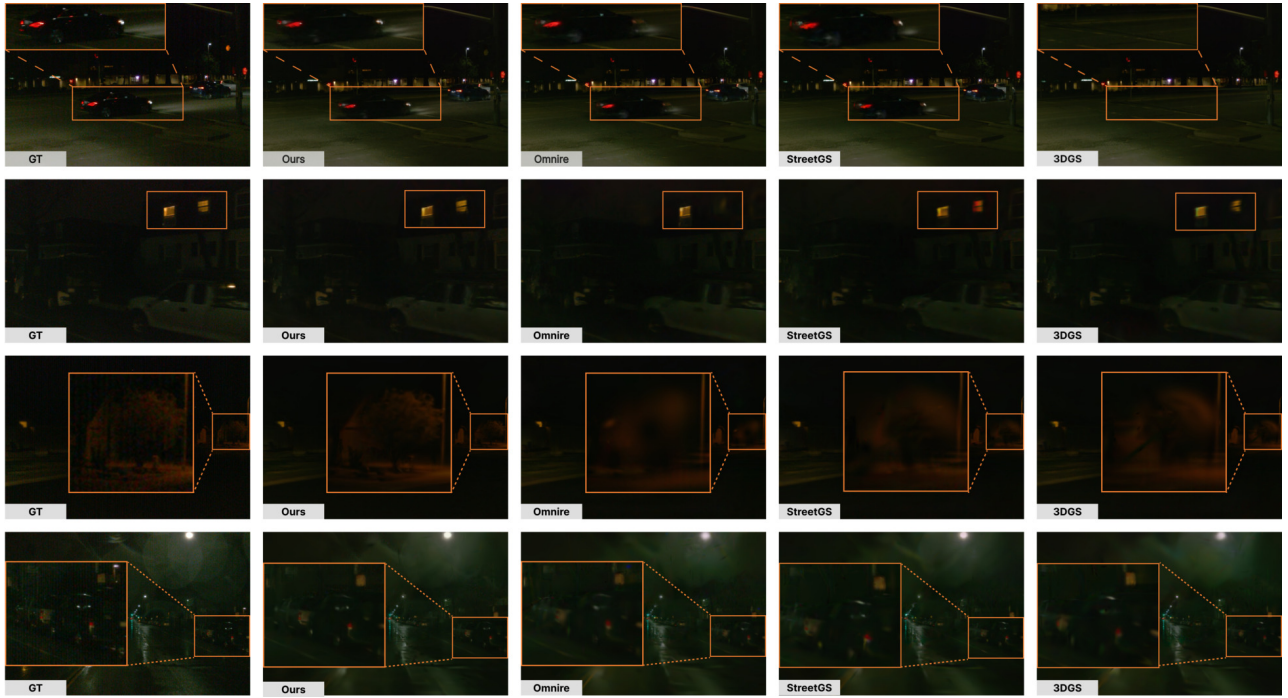


Fig. 2. Qualitative results from the Waymo Open Dataset. Columns show results from GT, Ours, OmniRe, StreetGS, and 3DGS, respectively. Fine-grained details such as traffic lights, vehicles, and trees are highlighted to illustrate differences in reconstruction quality.

TABLE I

QUANTITATIVE COMPARISON ON THE WAYMO DATASET FOR SCENE RECONSTRUCTION AND NOVEL VIEW SYNTHESIS TASKS.

Method	Reconstruction			Novel View		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
3DGS [23]	30.1	0.755	0.487	28.3	0.720	0.503
StreetGaussians [6]	30.5	0.757	0.482	28.5	0.722	0.499
OmniRe [5]	31.0	0.768	0.455	28.5	0.718	0.481
Ours	31.9	0.781	0.441	28.8	0.720	0.467

TABLE II

QUANTITATIVE COMPARISON ON THE NUSCENES DATASET FOR SCENE RECONSTRUCTION AND NOVEL VIEW SYNTHESIS TASKS.

Method	Reconstruction			Novel View		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
3DGS [23]	27.6	0.741	0.378	26.4	0.704	0.382
StreetGaussians [6]	28.3	0.749	0.364	26.9	0.709	0.376
OmniRe [5]	28.7	0.760	0.335	26.9	0.705	0.351
Ours	29.7	0.775	0.319	27.7	0.718	0.338

D. Results on Waymo

The quantitative results on Waymo are reported in Table I. Overall, our approach consistently outperforms prior state-of-the-art methods on both novel view synthesis and scene reconstruction. Specifically, for scene reconstruction, our model achieves a PSNR of 31.9 dB, SSIM of 0.781, and LPIPS of 0.441, surpassing any other existing models. We attribute these improvements to our decomposed lighting module: SH-based global illumination enhances stability across viewpoints and timestamps, while per-Gaussian ASGs effectively capture localized highlights such as headlights and reflective surfaces. For novel view synthesis, our model achieves the highest overall perceptual quality, with a PSNR of 28.8 dB, SSIM of 0.720, and LPIPS of 0.467. Although the PSNR and SSIM improvements are moderate, the consistent reduction in LPIPS underscores the ability of our lighting decomposition to produce sharper and more visually realistic renderings under low-light conditions. Importantly, despite the inherent challenge of low-light driving scenarios,

our method demonstrates consistent performance across 11 diverse scenes.

A qualitative comparison on Waymo nighttime sequences is shown in Figure 2. We can see that our method produces clearer vehicle renderings and accurately reconstructs the effect of headlights on the road surface. Additionally, it recovers more fine-grained details, such as illuminated building windows and reflective surfaces, which other baselines tend to oversmooth or fail to capture. For instance, vehicles rendered by other methods are overly smoothed and fail to model the car headlights and their interaction with the road, while our method generates sharper vehicles and captures localized lighting effects on the background. These results highlight the superiority of our framework in modeling low-light driving scenes. Furthermore, our approach also reconstructs clearer obstacles, such as trees, even under extremely low-light conditions, which is crucial for autonomous driving applications.



Fig. 3. Qualitative results from the nuScenes Dataset. Columns show results from GT, Ours, OmniRe, StreetGS, and 3DGS, respectively. Fine-grained details such as traffic lights, vehicles, and trees are highlighted to illustrate differences in reconstruction quality.



Fig. 4. Comparison of our decomposed lighting modules. Without ASGs, fine details on vehicles and scene elements are lost due to oversmoothing. Without global SH illumination, the model fails to capture the interaction between light sources (e.g., vehicle headlights) and the surrounding environment, including shadows cast on the ground.

E. Results on nuScenes

The quantitative results on nuScenes are presented in Table II. Despite the lower resolution, our model exhibits robust performance across all three metrics for both novel view synthesis and scene reconstruction. Specifically, our framework achieves a PSNR of 29.7 dB, SSIM of 0.775, and LPIPS of 0.319 for scene reconstruction, outperforming all prior state-of-the-art frameworks. Compared to the second-best baseline, OmniRe, this corresponds to an approximately 1.0 dB improvement in PSNR and a notable reduction in LPIPS, indicating sharper and robust reconstructions. For novel view synthesis, our model also outperforms the state-of-the-art models, achieving a PSNR of 27.7 dB, SSIM of 0.718, and LPIPS of 0.338. Although the improvements are slightly less than those for scene reconstruction, our method consistently surpasses other methods, highlighting its ability to generalize across challenging low-light driving scenes.

The qualitative results in Figure 3 on the nuScenes dataset further demonstrate the advantages of our approach. From Figure 3, we can see that our method produces clearer vehicle renderings and accurately reconstructs background elements, such as buildings, where other methods struggle. Additionally, our approach effectively reconstructs road lane signs, which are important for navigation in low-light conditions, while other methods fail to recover these features. These highlight the superiority of our framework in modeling low-light driving scenes and its ability to preserve critical

TABLE III
ABLATION STUDY ON THE EFFECTIVENESS OF OUR DECOMPOSED LIGHTING MECHANISM FOR SCENE RECONSTRUCTION AND NOVEL VIEW SYNTHESIS TASKS.

Variant	Reconstruction			Novel View		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
w/o Global Diffuse Modeling	30.1	0.765	0.397	27.9	0.716	0.417
w/o Specular Modeling	29.8	0.767	0.390	27.9	0.713	0.416
Full Model	30.8	0.778	0.380	28.2	0.720	0.403

details for autonomous driving tasks. These results confirm that our decomposition generalizes well across various scenes from different datasets, ensuring the fidelity of the our framework.

F. Ablation Study

1) *Effectiveness of Lighting Decomposition:* As shown in Table III, removing either the global diffuse SH-based illumination or the per-Gaussian ASG-based lighting degrades performance in both novel view and reconstruction tasks. Without global illumination, the model struggles to maintain consistency across the scene, leading to reduced PSNR and SSIM. Without ASG-based specular modeling, localized highlights such as vehicle headlights and reflective building windows are lost, causing a drop in perceptual quality (higher LPIPS). Furthermore, as shown in Figure 4, the module without ASGs fails to render fine details and

TABLE IV

ABLATION STUDY ON THE EFFECTIVENESS OF DIFFERENT SPECULAR MODELING TECHNIQUES.

Variant	Reconstruction			Novel View		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
	Specular Modeling w/ SH	30.1	0.768	0.391	28.1	0.717
Specular Modeling w/ ASG	30.8	0.778	0.380	28.2	0.720	0.403




Fig. 5. Comparison of specular maps produced using SH and ASGs. SH-based incident specular lighting tends to create exaggerated highlights that deviate from realistic appearance, whereas ASGs capture sharper and more physically plausible specular effects.

oversmooths the vehicles, while removing the global SH illumination results in incorrect modeling of the interaction between vehicle headlights and the ground and the shadows cast by the lighting. These results confirm that both global and local lighting are complementary and necessary for robust rendering.

2) *Analysis of ASG vs. SH for Specular Modeling:* As shown in Table IV, replacing ASGs with SH for incident lighting yields similar PSNR and SSIM, but increases in LPIPS, reflecting poorer perceptual sharpness. Qualitatively, as shown in Figure 5, we observe that SH fails to capture sharp specular highlights and instead produces smoother, overly diffuse reflections. ASGs, with their anisotropic nature, more accurately reproduce high-frequency specular effects, making them better suited for specular modeling.

3) *Analysis of BRDF Constraint:* As shown in Table V, when ASGs are directly evaluated without BRDF constraints, the performance degrades dramatically in all metrics compared to BRDF-constraint ASGs. Specifically, PSNR drops by 1.0 dB with a significant increase in LPIPS. This aligns with the result reported by [20]. Furthermore, a comparison of BRDF-constrained and direct evaluation of ASGs is shown in Figure 6. This confirms that constraining specular component with a physically based BRDF framework is crucial for producing realistic and physically consistent lighting, rather than arbitrary highlight patterns as in Figure 6.

4) *A Visualization of Lighting Component Decomposition:* We show a visualization of lighting component decomposition in Figure 7. We can see that the albedo represents the base color of the scene without any lighting or shading. The diffuse component models non-directional light and accounts for shading caused by environmental illumination, while the specular component captures sharp highlights from reflective surfaces or self-illuminating objects, such as neon signs or traffic signals. These results demonstrate the robustness of our framework and its ability to produce physically plausible and realistic lighting decomposition.

TABLE V

ABLATION STUDY ON THE EFFECTIVENESS OF BRDF CONSTRAINT.

Variant	Reconstruction			Novel View		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Ours w/o BRDF Constraint	29.8	0.760	0.426	27.6	0.712	0.446
Ours w/ BRDF Constraint	30.8	0.778	0.380	28.2	0.720	0.403



Fig. 6. Comparison of BRDF-constrained and direct evaluation of ASGs. BRDF-constrained ASGs produces more physically realistic lighting and construct overall better quality rendering than direct evaluation.



Fig. 7. Visualization of the outputs. From top to bottom and left to right: the final rendered image, albedo, diffuse, and specular. This illustrates the decomposition and the contribution of each lighting component to the final rendering.

V. CONCLUSION

This work proposes a novel framework to effectively reconstruct low-light autonomous driving scenes. Our key idea is to integrate physically based rendering into composite scene Gaussian representations, which jointly optimizes BRDF-based material properties, for nighttime driving scene reconstruction. Experiments on two challenging autonomous driving datasets demonstrate the superiority of our framework both quantitatively and qualitatively.

Limitation and Future Work. Since our approach adopts a per-sequence reconstruction paradigm, our approach has scalability issues when applied to large-scale driving scene simulation in digital twin environments. In addition, the current lighting modeling strategy may not fully capture complex lighting conditions in large-scale driving scenes. Our future work aims to address these limitations by exploring feed-forward reconstruction paradigms, incorporating richer priors beyond surface normals to better capture material properties, exploring alternative BRDF models other than the simplified Disney model, and integrating more advanced tone-mapping module to further enhance reconstruction quality.

REFERENCES

- [1] 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.
- [2] J. Cao, Z. Li, N. Wang, and C. Ma, "Lightning nerf: Efficient hybrid scene representation for autonomous driving," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 16 803–16 809.
- [3] Z. Yang, Y. Chen, J. Wang, S. Manivasagam, W.-C. Ma, A. J. Yang, and R. Urtaşun, "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] 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.
- [5] Z. Chen, J. Yang, J. Huang, R. de Lutió, J. M. Esturo, B. Ivanovic, O. Litany, Z. Gojčić, S. Fidler, M. Pavone, L. Song, and Y. Wang, "Omnire: Omni urban scene reconstruction," in *The Thirteenth International Conference on Learning Representations*, 2025.
- [6] 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.
- [7] G. Wu, D. Bai, and B. Liu, "Arngs: Composite gaussian appearance refinement for modeling dynamic urban environments," *arXiv preprint arXiv:2507.03886*, 2025.
- [8] 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.
- [9] Z. Qu, K. Xu, G. P. Hancke, and R. W. Lau, "Lush-nerf: lighting up and sharpening nerfs for low-light scenes," in *Proceedings of the 38th International Conference on Neural Information Processing Systems*, 2024, pp. 109 871–109 893.
- [10] S. Ye, Z.-H. Dong, Y. Hu, Y.-H. Wen, and Y.-J. Liu, "Gaussian in the dark: Real-time view synthesis from inconsistent dark images using gaussian splatting," in *Computer Graphics Forum*, vol. 43, no. 7. Wiley Online Library, 2024, p. e15213.
- [11] J. Gao, C. Gu, Y. Lin, Z. Li, H. Zhu, X. Cao, L. Zhang, and Y. Yao, "Relightable 3d gaussians: Realistic point cloud relighting with brdf decomposition and ray tracing," in *European Conference on Computer Vision*. Springer, 2024, pp. 73–89.
- [12] K. Xu, W.-L. Sun, Z. Dong, D.-Y. Zhao, R.-D. Wu, and S.-M. Hu, "Anisotropic spherical gaussians," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 6, pp. 1–11, 2013.
- [13] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.
- [14] P. Sun, H. Kretschmar, X. Dotiwala, 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.
- [15] 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.
- [16] Y. Ren, G. Wu, R. Li, Z. Yang, Y. Liu, X. Chen, T. Cao, and B. Liu, "Unigaussian: Driving scene reconstruction from multiple camera models via unified gaussian representations," *arXiv preprint arXiv:2411.15355*, 2024.
- [17] H. Zhou, J. Shao, L. Xu, D. Bai, W. Qiu, B. Liu, Y. Wang, A. Geiger, and Y. Liao, "Hugs: Holistic urban 3d scene understanding via gaussian splatting," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 21 336–21 345.
- [18] Y. Yao, J. Zhang, J. Liu, Y. Qu, T. Fang, D. McKinnon, Y. Tsin, and L. Quan, "Neilf: Neural incident light field for physically-based material estimation," in *European conference on computer vision*. Springer, 2022, pp. 700–716.
- [19] B. Burley and W. D. A. Studios, "Physically-based shading at disney," in *Acm siggraph*, vol. 2012, no. 2012. vol. 2012, 2012, pp. 1–7.
- [20] Z. Yang, X. Gao, Y.-T. Sun, Y. Huang, X. Lyu, W. Zhou, S. Jiao, X. Qi, and X. Jin, "Spec-gaussian: Anisotropic view-dependent appearance for 3d gaussian splatting," *Advances in Neural Information Processing Systems*, vol. 37, pp. 61 192–61 216, 2024.
- [21] T. Zhang, K. Huang, W. Zhi, and M. Johnson-Roberson, "Darkgs: Learning neural illumination and 3d gaussians relighting for robotic exploration in the dark," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 12 864–12 871.
- [22] H. Sun, F. Yu, H. Xu, T. Zhang, and C. Zou, "Li-gaussian: Low-light scene reconstruction and enhancement via gaussian splatting for novel view synthesis," *arXiv preprint arXiv:2504.10331*, 2025.
- [23] 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.
- [24] Z. Yang, X. Gao, W. Zhou, S. Jiao, Y. Zhang, and X. Jin, "Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024, pp. 20 331–20 341.
- [25] J. T. Kajiya, "The rendering equation," in *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, 1986, pp. 143–150.
- [26] M. F. Cohen and J. R. Wallace, *Radiosity and realistic image synthesis*. Morgan Kaufmann, 1993.
- [27] R. Green, "Spherical harmonic lighting: The gritty details," in *Archives of the game developers conference*, vol. 56, 2003, p. 4.
- [28] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda, "Photographic tone reproduction for digital images," in *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, 2023, pp. 661–670.
- [29] R. Wang, S. Xu, Y. Dong, Y. Deng, J. Xiang, Z. Lv, G. Sun, X. Tong, and J. Yang, "Moge-2: Accurate monocular geometry with metric scale and sharp details," *arXiv preprint arXiv:2507.02546*, 2025.
- [30] H. Chen, F. Wei, C. Li, T. Huang, Y. Wang, and G. H. Lee, "Vcr-gaus: View consistent depth-normal regularizer for gaussian surface reconstruction," *Advances in Neural Information Processing Systems*, vol. 37, pp. 139 725–139 750, 2024.
- [31] D. P. Kingma, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [32] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 586–595.
- [33] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.