

Lightning NeRF: Efficient Hybrid Scene Representation for Autonomous Driving

Junyi Cao¹, Zhichao Li², Naiyan Wang², and Chao Ma^{1,*}

Abstract—Recent studies have highlighted the promising application of NeRF in autonomous driving contexts. However, the complexity of outdoor environments, combined with the restricted viewpoints in driving scenarios, complicates the task of precisely reconstructing scene geometry. Such challenges often lead to diminished quality in reconstructions and extended durations for both training and rendering. To tackle these challenges, we present Lightning NeRF. It uses an efficient hybrid scene representation that effectively utilizes the geometry prior from LiDAR in autonomous driving scenarios. Lightning NeRF significantly improves the novel view synthesis performance of NeRF and reduces computational overheads. Through evaluations on real-world datasets, such as KITTI-360, Argoverse2, and our private dataset, we demonstrate that our approach not only exceeds the current state-of-the-art in novel view synthesis quality but also achieves a five-fold increase in training speed and a ten-fold improvement in rendering speed. Codes are available at <https://github.com/VISION-SJTU/Lightning-NeRF>.

I. INTRODUCTION

Neural Radiance Fields (NeRFs) [1] have paved a novel pathway for novel view synthesis, exhibiting remarkable results across various datasets. When dealing with outdoor scenes, the vastness, complex structures, and fluctuating lighting conditions considerably amplify the complexity of scene reconstruction and substantially heighten the computational demands. Recently, several methods have emerged to address the challenges of applying NeRF to outdoor environments. Specifically, NeRF-W [2] incorporates a learnable appearance embedding to address lighting variations. In the context of autonomous driving, some techniques [3], [4], [5], [6] integrate point clouds to deliver enhanced geometric information, addressing the issue of representing complex structures. However, these methods often overlook the efficiency and computational overhead associated with training and rendering. More complex modeling and larger scenes invariably lead to extended model training times.

In this paper, we propose an efficient hybrid scene representation. We separately model density and color in NeRF, using explicit and implicit approaches respectively. For density, point clouds offer an effective initialization, substantially reducing representational challenges. This allows us to explicitly model density using a limited-resolution voxel grid, eliminating the need for a Multi-Layer Perceptron (MLP). For rendering image details, we retain the implicit

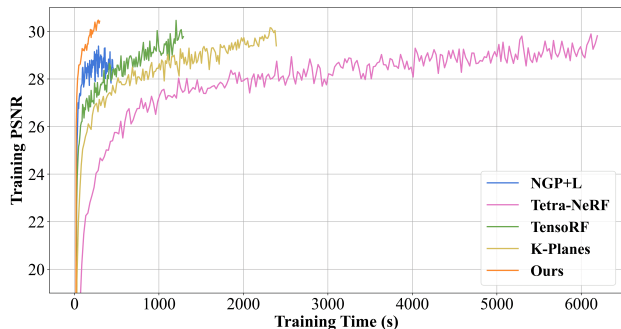


Fig. 1. Training efficiency. These curves reflect the trend of training PSNR with time. The values are obtained on sequence 133e2e0b of Argoverse2 [8].

modeled color MLP to ensure the capacity to accommodate the highly variable real world. Moreover, we present a more realistic model of the outdoor scene background and color decomposition, further elevating the quality of novel view synthesis and rendering efficiency. Comparative studies on real-world autonomous driving datasets, including KITTI-360 [7], Argoverse2 [8] and a private dataset, show that our method not only surpasses current state-of-the-arts in performance for novel view synthesis but also achieves a five-fold improvement in training speed and a ten-fold boost in rendering speed.

II. RELATED WORK

A. Neural Radiance Fields (NeRF)

NeRF [1] emerged as a groundbreaking approach to synthesize novel views of a scene by fitting a neural radiance field given calibrated RGB images, offering unprecedented quality in scene representation. Recent works [9], [10], [11], [12], [13], [14], [15], [16], [17] pursuit of more detailed modeling, coupled with architectural innovations to push the boundaries of NeRF. However, a significant challenge associated with NeRF is its computational intensity. Specifically, NeRF requires querying a deep MLP millions of times, which results in slow training and rendering. Many studies [18], [19] attempt to accelerate the process through more efficient sampling [20], [21], [22], [23] or enhanced data structures [24], [25], [26], [27], [28], [29], [30], [31], [32], essentially trading space for time. However, when applied to outdoor unbounded scenes, they often encounter problems such as slow convergence, insufficient capacity, or excessive memory overhead.

B. NeRF in Outdoor Scenes

Outdoor environments, with their expansive scales and complex lighting dynamics, pose fresh challenges for NeRF.

¹ Junyi Cao and Chao Ma are with MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, Shanghai, China. {junyicao, chaoma}@sjtu.edu.cn

² Zhichao Li and Naiyan Wang are with Tusimple, Beijing, China. {leeisabug, winsty}@gmail.com

* Chao Ma is the corresponding author.




Fig. 2. Overview of the proposed framework. The red and green boxes represent the foreground and background in our proposed scene representation, respectively. Given point cloud data from LiDAR observations, we first use LiDAR Initialization to initialize the scene geometry (see Sec. III-C). Then, we query the volume density σ and the color embedding feature \mathbf{f} of each sample point along a ray from the voxel grids (see Sec. III-B). We adopt separate MLPs for modeling view-dependent (with the viewing direction \mathbf{d} as an additional input) and view-independent colors. Combining the two components achieves the final rendered image (see Sec. III-D).

To break the limit of single model capacity, Block NeRF [33] strategically segments vast scenes into smaller subsections for more manageable reconstruction. NeRF-W [2] introduces an appearance embedding for each camera view to learn the photometric differences caused by variations in lighting and viewpoint. To enhance the reconstruction fidelity of large scene geometry, some researches consider additional supervisions, *e.g.*, point cloud data [3], [5], [6], [34] or semantic labels [35], [36]. However, the expansion of scene doesn't merely demand augmented model capacity; it also precipitates a decline in both model convergence and rendering speed. To this end, we harness the geometry cues from point clouds and integrate them with efficient hybrid scene representation. Our method not only facilitates high-fidelity reconstructions, but also effectively boosts the model's convergence and rendering pace.

III. METHOD

A. Preliminaries

NeRF [1] represents a scene with an implicit function $\mathcal{F}(\mathbf{x}, \mathbf{d})$, often parameterized by MLPs, that returns a color value \mathbf{c} and a volume density prediction σ for a random 3D point \mathbf{x} in the scene observed from a viewing direction \mathbf{d} :

$$(\sigma, \mathbf{f}) = \mathcal{M}^{(\text{den})}(\mathbf{x}), \quad (1)$$

$$\mathbf{c} = \mathcal{M}^{(\text{rgb})}(\mathbf{f}, \mathbf{d}), \quad (2)$$

where $\mathcal{M}^{(\text{den})}$, $\mathcal{M}^{(\text{rgb})}$ are MLPs and \mathbf{f} is a latent embedding to assist the much shallower $\mathcal{M}^{(\text{rgb})}$ to learn \mathbf{c} (see NeRF++ [9] for detailed discussions). Volume rendering [37] is used to synthesize novel views. Concretely, to render a pixel, NeRF leverages hierarchical volume sampling to generate N points along the ray \mathbf{r} . The predicted density $\{\sigma_i\}_{i=1}^N$ and color feature $\{\mathbf{c}_i\}_{i=1}^N$ at these positions are combined by accumulation:

$$\hat{\mathbf{C}}(\mathbf{r}) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, T_i = \exp(-\sum_{j=1}^{i-1} \sigma_j \delta_j), \quad (3)$$

where δ_i is the distance between adjacent samples. This process is fully differentiable. Hence, the mean squared error between the predicted and ground-truth colors corresponding to each ray is used to supervise the optimization of NeRF.

While NeRF exhibits superior performance in novel view synthesis, it suffers long training time and slow rendering speed, which is in part because of the inefficiency of its sampling strategy. To this end, [19], [20], [27] maintain a coarse occupancy grid during training and only sample locations inside occupied volume. We use a similar sampling strategy as these works to boost the efficiency of our model.

B. Hybrid Scene Representation

Hybrid volumetric representations [20], [27], [30] have achieved fast optimization and rendering with a compact model. In view of that, we adopt a hybrid voxel-grid representation to model the radiance field for efficiency. Briefly, we explicitly model the volume density by storing σ at grid vertices while using shallow MLPs to decode color embeddings \mathbf{f} into final color \mathbf{c} in an implicit way. To handle the unbounded nature of outdoor environments, we split our scene representation into two portions for foreground and background separately, as shown in Fig. 2. Concretely, we inspect the camera frustum in each frame from a trajectory sequence and define the foreground bounding box so that it tightly wraps all frustums in the aligned coordinate system. The background box is obtained by enlarging the foreground box proportionally along each dimension.

Voxel-grid Representation. A voxel-grid representation explicitly stores the scene properties (*e.g.*, density, RGB color, or feature) in its grid vertices to support efficient feature queries. In this way, for a given 3D location $\mathbf{x} \in \mathbb{R}^3$, we can decode the corresponding property by tri-linear interpolation:

$$\mathcal{G}(\mathbf{x}) \triangleq \psi(\mathbf{x}, \mathbf{G}) : (\mathbb{R}^3, \mathbb{R}^{C \times N_G \times N_G \times N_G}) \rightarrow \mathbb{R}^C, \quad (4)$$

where $\mathcal{G}(\cdot)$ denotes the decoding operation, $\psi(\cdot)$ is the tri-linear interpolation, and \mathbf{G} is the voxel grid. C indicates the dimension of the outputs and N_G gives the grid resolution.

Foreground. We set up two separate feature grids for modeling density and color embedding in the foreground region. Specifically, the density grid mapping $\mathcal{G}_{\text{fg}}^{(\text{den})}$ maps positions into density scalars σ for volume rendering. For the color-embedding grid mapping $\mathcal{G}_{\text{fg}}^{(\text{rgb})}$, we instantiate multiple voxel-grid at different resolution backup by a hash-table, as in [19], to achieve finer details with affordable memory overhead. The final color embedding \mathbf{f} is obtained by concatenating outputs over L resolution levels.

Background. Although the aforementioned foreground modeling works well for radiance fields at an object level, extending it to unbounded outdoor scenes is non-trivial. Some related arts like NGP [19] directly expands its scene bounding box so that the background region can be included, while GANcraft [38] and URF [5] introduces a spherical background radiance to handle this issue. However, the former attempt leads to a waste of its capability as most regions inside its scene box are for the background scene. For the latter proposal, it may fail in handling sophisticated panorama in urban scenes (*e.g.*, undulating buildings or complex landscapes) since it simply assumes the background radiance is only conditioned on view directions.

To this end, we set up an additional background grid model to keep the resolution in our foreground portion unchanged. We adopt the scene parameterization in [9] for the background with deliberate designs. First, unlike its inverse sphere modeling, we use inverse cubic modeling that replaces the ℓ_2 norm with the ℓ_∞ norm since we use a voxel-grid representation. Second, we do not instantiate additional MLPs for querying background color to save memory. Specifically, we warp 3D background points into 4D by

$$\mathbf{x} = (x_1, x_2, x_3) \rightarrow (x'_1, x'_2, x'_3, \frac{1}{r}) = \mathbf{x}', \quad (5)$$

where $r = \|\mathbf{x}\|_\infty$ and $x'_k = x_k/r$ for $k = \{1, 2, 3\}$. Then, similar to the grid mapping used in the foreground, we query 4D grid $\mathcal{G}_{\text{bg}}^{(\text{den})}$ and $\mathcal{G}_{\text{bg}}^{(\text{rgb})}$ given \mathbf{x}' to obtain σ and \mathbf{f} for background locations. Note that $\mathcal{G}_{\text{bg}}^{(\text{rgb})}$ shares the same numbers of levels L and output dimensions C with $\mathcal{G}_{\text{fg}}^{(\text{rgb})}$.

In a nutshell, our hybrid scene representation (HSR) outputs density scalars and color embeddings for any 3D locations. The foreground and background decomposition is based on the points' relative positions to the foreground bounding box. Therefore, for simplicity, we will omit the subscript for the grid mapping operation and use $\mathcal{G}^{(\text{den})}$, $\mathcal{G}^{(\text{rgb})}$ in the following sections.

C. LiDAR Initialization

With our hybrid scene representation, the model can save computation and memory as we query density value directly from the efficient voxel-grid representation rather than the computationally intensive MLP $\mathcal{M}^{(\text{den})}$. However, given the large-scale nature and complexity of urban scenes, this lightweight representation easily gets trapped into local minima in optimization due to the limited resolution of the density grid. Fortunately, in autonomous driving scenarios,




Fig. 3. Visual results in the extrapolation setting with or without the proposed color decomposition (CD). Best viewed in color with zoomed-in.

most Self-Driving Vehicles (SDVs) are equipped with LiDAR sensors, which provide coarse geometry priors for scene reconstruction. To this end, we propose to use the LiDAR point clouds to initialize our density grid to alleviate the handicap of the joint optimization of scene geometry as well as the radiance. Specifically, we first load LiDAR point clouds $P_{\text{pc}} = \{\mathbf{p}_i | \mathbf{p}_i \in \mathbb{R}^3\}_i^{N_{\text{pc}}}$ and additionally create a set of points $P_{\text{bg}} = \{\mathbf{p}_j | \mathbf{p}_j \in \mathbb{R}^3\}_j^{N_{\text{bg}}}$ scattering in the top, front, left and right surfaces of the background box. Then, we manually initialize the density value σ at the corresponding location $\mathbf{p} \in P_{\text{pc}} \cup P_{\text{bg}}$ to be a positive constant σ_0 by adjusting the values in the density grid. Meanwhile, we also update the occupancy grid used for efficient sampling based on the output of the initialized density mapping. This initialization process ensures the sample locations are sparse in the scene and cluster around the objects' surfaces. We find that this step is crucial to the rendering quality and efficiency of our hybrid scene representation (see ablation study in Tab. V).

D. Color Decomposition

The original NeRF [1] uses a view-dependent MLP $\mathcal{M}^{(\text{rgb})}$ to model the color in the radiance field, which is a simplification of the physical world where radiance is composed of diffuse (view-independent) color and specular (view-dependent) color [39], [40], [41], [42]. Besides, since the final output color \mathbf{c} is fully entangled with the viewing direction \mathbf{d} , it poses difficulties to render high-fidelity images in unseen views (*i.e.*, extrapolation), as $\mathcal{M}^{(\text{rgb})}$ is not optimized for viewing directions out of training domains. However, this is the common case in autonomous driving since single trip usually can only scan limited space. Thus, we propose to decompose the color into view-dependent factor \mathbf{c}_{vd} and view-independent factor \mathbf{c}_{vi} to better simulate the physical world. In this way, when the model faces large shifts in viewing directions, the view-independent factor can still give a reasonable prediction for the final color. As shown in Fig. 3, our method trained without color decomposition (CD) fails in novel view synthesis in the extrapolation setting (*i.e.*, shift the viewing direction 2 meters left based on training views) while ours with color decomposition gives a plausible rendering result.

We introduce two lightweight MLPs, denoted as $\mathcal{M}^{(\text{rgb-vd})}$, $\mathcal{M}^{(\text{rgb-vi})}$, for modeling view-dependent color \mathbf{c}_{vd} and view-independent color \mathbf{c}_{vi} , respectively:

$$\mathbf{c}_{\text{vd}} = \mathcal{M}^{(\text{rgb-vd})}(\mathbf{f}, \mathbf{d}), \quad (6)$$

$$\mathbf{c}_{\text{vi}} = \mathcal{M}^{(\text{rgb-vi})}(\mathbf{f}). \quad (7)$$

The final color at the sample location is the sum of these two factors:

$$\mathbf{c} = \mathbf{c}_{\text{vd}} + \mathbf{c}_{\text{vi}}. \quad (8)$$

We then use the conventional volume rendering to obtain the color for a pixel as in Eqn. (3).

E. Training Loss

We modify the photometric loss with a re-scaling weight w_i to optimize our model to focus on the hard samples for fast convergence. The weight coefficient is defined as:

$$w_i = \text{clamp}(\tilde{w}_i), \quad (9)$$

where $\text{clamp}(\cdot)$ means to clip the value into $[1, 10]$ and

$$\tilde{w}_i = \frac{\|\hat{C}(\mathbf{r}_i) - C(\mathbf{r}_i)\|_2^2}{\min_{i \in \{1, \dots, N\}} (\|\hat{C}(\mathbf{r}_i) - C(\mathbf{r}_i)\|_2^2)}. \quad (10)$$

Our weighted photometric loss is written as:

$$\mathcal{L}_p = \frac{1}{N} \sum_{i=1}^N \langle w_i \rangle \cdot \|\hat{C}(\mathbf{r}_i) - C(\mathbf{r}_i)\|_2^2, \quad (11)$$

where $\hat{C}(\mathbf{r}), C(\mathbf{r})$ are the predicted and the ground-truth color, $\langle \cdot \rangle$ denotes the stop-gradient [43] operation. We additionally regularize the view-dependent color \mathbf{c}_{vd} with a ℓ_1 regularizer during training:

$$\mathcal{L}_r = \frac{1}{M} \sum_{j=1}^M \|\mathbf{c}_{\text{vd},j}\|_1. \quad (12)$$

The total loss function is:

$$\mathcal{L} = \mathcal{L}_p + \lambda \mathcal{L}_r, \quad (13)$$

where λ is a weight parameter for balancing the two losses.

IV. EXPERIMENTS

A. Experimental Setup

Datasets. Three urban datasets are considered in our experiments. Two of them are public-available, *i.e.*, KITTI-360 [7] and Argoverse2 [8], while the remaining one is a private dataset used for evaluation on long trajectories. For KITTI-360, we follow [3] to select 5 scenes mainly containing static objects. Note that LiDAR sensors used in KITTI-360 only capture near-ground objects due to a limited vertical field of view. The incomplete LiDAR data inevitably hampers the optimization of our model. To test our method with complete LiDAR data, we use Argoverse2 and the private dataset. We select 6 sequences from Argoverse2 *Sensor Dataset* and 2 trajectories from the private dataset. We select every 10th image in the sequences as the test data for these two datasets and take the rest as the training data.

Implementation Details. Our model is based on the NeRF-Studio framework [46]. During training, we use RAdam [47] optimizer for the density and color embedding grids with an initial learning rate of 1.0. The MLP networks are optimized via Adam [48] optimizer with an initial learning rate of 0.01. λ in Eqn. (13) is set to be 0.01. We use a batch size of 65,536

TABLE I
QUANTITATIVE RESULTS ON THE KITTI-360.

Method	Input	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
NeRF-W [2]	I	22.77	79.40	44.60
NGP [19]+L [5]	I+L	23.36	85.05	35.63
TensorRF [30]	I	23.55	<u>85.14</u>	36.37
K-Planes [44]	I	<u>23.63</u>	84.94	37.00
Point-NeRF [23]	I+L	21.54	79.30	40.60
Tetra-NeRF [45]	I+L	21.73	80.22	43.60
Mip-NeRF 360 [13]	I	23.27	83.60	35.50
DNMP [3]	I+L	23.41	84.60	30.50
Ours	I+L	23.71	85.21	<u>34.27</u>

TABLE II
QUANTITATIVE RESULTS ON THE ARGOVERSE2.

Method	Input	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
NGP [19]	I	28.40	83.74	40.71
NGP [19]+L [5]	I+L	28.64	84.07	<u>40.37</u>
TensorRF [30]	I	<u>30.05</u>	84.09	42.31
K-Planes [44]	I	29.26	83.56	41.86
Tetra-NeRF [45]	I+L	29.59	<u>84.33</u>	42.68
Ours	I+L	30.62	85.59	38.32

rays for training 30k iterations unless otherwise specified. All experiments are conducted on an NVIDIA RTX 3090 GPU.

Evaluation Metrics. Following previous works [1], [3], [13], we evaluate the novel view synthesis performance based on three commonly-used metrics, *i.e.*, peak signal-to-noise ratio (PSNR), structural similarity index measure (SSIM) [49], and the learned perceptual image patch similarity (LPIPS) [50].

B. Comparisons with the State-of-the-art Methods

Quality of Novel View Synthesis. To analyze the rendering quality of our model, we present the novel view synthesis metrics on KITTI-360 and Argoverse2 in Tab. I and Tab. II. Due to the truncation of height in KITTI-360’s point clouds, there’s a significant portion of missing coverage. This results in many foreground areas remaining uninitialized after LiDAR initialization. Despite the challenge, we still surpass the state-of-the-art methods in terms of PSNR and SSIM. It should be noted that training DNMP [3] is far less efficient than ours as it requires a mesh optimization stage before training the NeRF. In contrast to the KITTI-360, the point cloud coverage of Argoverse2 is much more comprehensive. Consequently, our method exhibits clear superiority across all metrics. We also display the rendered images and depth maps in Fig. 4 for a qualitative comparison. Our depth maps offer richer details with more consistent transitions, effectively capitalizing on the prior information provided by the point clouds. Furthermore, in Fig. 5, we showcase the extrapolation results of our method. Extrapolation is essential for closed-loop simulation systems. When SDVs exhibit behaviors different from the recorded data, the simulator should be capable of generating observations from corresponding views. Thanks to the refined depth and color decomposition

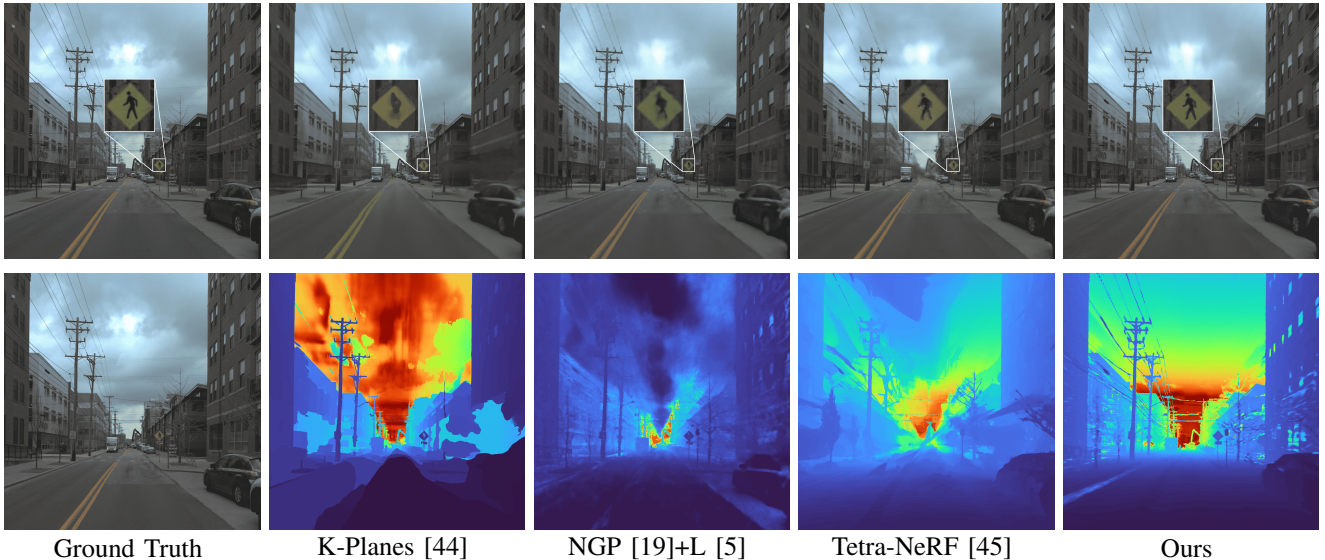


Fig. 4. Qualitative results on the Argoverse2. The first row shows rendered images and the second row shows depth maps. Best viewed in color.

TABLE III

COMPARISON OF TRAINING AND RENDERING SPEED WITH OTHER METHODS. IN THE TABLE, “TT” REFERS TO THE TRAINING TIME, AND “R-FPS” REFERS TO THE RENDERING FPS.

Method	133e2e0b (156m)		2aea7bd1 (139m)		b1a98ad6 (106m)	
	TT(m)	R-FPS	TT(m)	R-FPS	TT(m)	R-FPS
NGP [19]	2.62	0.19	4.32	0.19	3.61	0.17
NGP [19]+L [5]	2.27	0.22	3.88	0.23	3.17	0.19
TensoRF [30]	7.65	0.15	7.10	0.15	8.64	0.15
K-Planes [44]	13.13	0.23	12.92	0.23	10.74	0.23
Tetra-NeRF [45]	44.64	0.02	129.98	0.02	60.01	0.02
Ours	0.76	2.30	1.07	2.56	0.88	2.18

achieved by our method, our extrapolation results are considerably superior compared to other techniques.

Training & Rendering Cost. Tab. III showcases the advantages of our proposed method in terms of training and rendering speed. For a fair comparison of training speed, we report the training time for each method to converge to the same PSNR as NGP [19]. Specifically, we choose 3 sequences from Argoverse2 [8] in this section for comparisons. For sequences 133e2e0b, 2aea7bd1, and b1a98ad6, we use PSNR values of 28.5, 27.0, and 28.0, respectively, to report the training time. We achieve significant speed improvements in all test scenes. As can be seen from the table, the combination of our proposed LiDAR Initialization (LI) and Hybrid Scene Representation (HSR) yields more than 4 times boost in terms of convergence speed of our method compared with NGP [19]. Additionally, because we model the occupancy information explicitly, our instructive sampling reduces the number of sample points during rendering and saves time on MLP inference, thereby achieving a rendering speed that is 10 times faster. We will discuss in detail the impact of each component on speed in the following Sec. IV-C. We also show the results in the private dataset in Tab. IV to demonstrate the advantage of training speed on large-scale scenes. For the 600m scene, the proposed method trained within 1 minute exceeds the performance of NGP+L

TABLE IV

RESULTS ON THE PRIVATE DATASET. WE PRESENTED METRICS FOR 1-MINUTE AND 5-MINUTE TRAINING.

v1 (600m)	1min		5min	
	PSNR \uparrow	LPIPS \downarrow	PSNR \uparrow	LPIPS \downarrow
NGP [19]+L [5]	27.95	41.68	29.03	38.43
Ours	30.48	37.14	31.69	33.14
v2 (1000m)	PSNR \uparrow		LPIPS \downarrow	
	PSNR \uparrow	LPIPS \downarrow	PSNR \uparrow	LPIPS \downarrow
NGP [19]+L [5]	27.34	45.06	29.27	41.33
Ours	28.79	40.65	30.04	36.83

trained for 5 minutes. These results convincingly verify the efficiency of our proposed method.

C. Ablation Study

Analysis on Proposed Components. To analyze the roles of different components, we present the novel view synthesis quality and efficiency in Tab V. The PSNR values in the table are obtained under the standard training setting on test sets, while the training time reflects how long the models converge to the same PSNR as NGP [19], same as in Tab. III. Based on our method, we respectively remove corresponding modules to analyze the differences in outcomes, including Hybrid Scene Representation (HSR), LiDAR Initialization (LI), and Color Decomposition (CD). From the table, it’s evident that the most significant impacts on training and rendering speed come from HSR and LI. Specifically, when HSR is ablated, the model uses NGP’s multi-scale hash grids to store features and uses separate MLPs to implicitly model both volume density and color. In large-scale scenes, high-resolution hash grids cause severe hash collisions. NGP typically relies on MLPs to handle these collisions, making the spatial representation ambiguous and leading to declines in novel view synthesis performance. Additionally, the density MLP introduces prediction overheads, resulting in a significant drop in training and rendering speeds. On the other hand,



Reference View K-Planes [44] NGP [19]+L [5] Tetra-NeRF [45] Ours

Fig. 5. Extrapolation results on the Argoverse2. Based on the reference view, we moved the camera 2 meters to the left.

TABLE V
ABLATION STUDIES FOR PROPOSED COMPONENTS ON THE ARGOVERSE2.

Variant	133e2e0b				2aea7bd1				b1a98ad6			
	PSNR	#Sample/Ray	TT(m)	R-FPS	PSNR	#Sample/Ray	TT(m)	R-FPS	PSNR	#Sample/Ray	TT(m)	R-FPS
w/o HSR	29.23	77.12	3.93	0.34	26.36	51.73	8.71	0.56	28.89	64.01	4.56	0.40
w/o LI	25.65	39.42	Inf.	0.72	25.86	34.48	Inf.	0.76	24.61	47.19	Inf.	0.60
w/o CD	31.86	27.07	0.57	1.29	30.42	21.40	0.70	1.53	31.42	24.06	0.82	1.27
Ours	31.74	12.06	0.76	2.30	30.04	11.26	1.07	2.56	31.55	12.61	0.88	2.18

TABLE VI
ABLATION STUDIES FOR BACKGROUND MODELING IN TERMS OF PSNR.

Background Modeling	133e2e0b	2aea7bd1	b1a98ad6
Enlarged Scene Box [19]	31.43	24.46	30.02
Sphere Background [5]	30.94	26.34	29.21
Ours	31.74	30.04	31.55

TABLE VII
ABLATION STUDIES FOR COLOR MODELING.

Color Modeling	133e2e0b		2aea7bd1		b1a98ad6	
	PSNR	R-FPS	PSNR	R-FPS	PSNR	R-FPS
SH (degree=0)	27.42	3.83	22.80	3.91	28.25	3.64
SH (degree=1)	29.06	3.33	24.33	3.43	28.98	2.96
SH (degree=2)	29.14	2.64	24.68	2.79	29.29	2.43
Ours	31.74	2.30	30.04	2.56	31.55	2.18

without LI, plus the absence of an MLP in our model to fit density, it cannot converge to NGP’s PSNR. This results in the training time being shown as “Inf.” Furthermore, due to the absence of initialization and fitting capabilities, the model struggles to represent space sparsely, leading to a substantial number of samples per ray. This, in turn, leads to a noticeable decline in both rendering speed and quality. Finally, CD enables the model to focus its representation on most non-translucent and non-reflective areas without requiring excessive sample points for description. However, unsupervised decomposition slightly increases the fitting complexity, leading to a minor drop in PSNR. This represents a trade-off between performance and efficiency.

Analysis on Background Modeling. Here, we compare our background modeling strategy with the aforementioned proposals, *i.e.*, “Enlarged Scene Box” used by [19] and “Sphere Background” adopted by [5]. The comparative results are shown in Tab. VI. From the table, it is observed that our background modeling approach consistently outperforms the others by a clear margin, which verifies the advantage of our proposed strategy in autonomous driving scenarios.

Analysis on Color Modeling. Some recent work [26], [28] propose to use spherical harmonic (SH) to model scene appearance explicitly for efficiency, avoiding the need for a color MLP. To compare the effectiveness of our proposed color modeling with SH representations, we conduct ablation studies and show the results in Tab. VII. It is observed that SH representations achieve fast rendering speed thanks to

explicit color modeling. However, due to the complexity inherited from large-scale scenarios, SH with limited grid resolution cannot achieve fine rendering quality compared with our hybrid representations. In summary, our method enjoys efficient training and rendering via explicit modeling of scene geometry while maintaining rich expression ability for scene appearance via implicit color modeling.

V. CONCLUSION

This paper introduces Lightning NeRF, an efficient novel view synthesis framework for outdoor scenes that integrates point clouds and images. The proposed method utilizes point clouds for a swift initialization of a sparse representation of the scene, achieving significant performance and speed enhancements. By modeling the background more effectively, we reduce the representational strain on the foreground. Finally, through color decomposition, we model view-dependent and view-independent colors separately, enhancing the model’s extrapolation capability. Extensive experiments on various autonomous driving datasets demonstrate that our method outperforms the previous state-of-the-art techniques in both performance and efficiency.

Acknowledgments. This work was supported in part by NSFC (62322113, 62376156), Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), and the Fundamental Research Funds for the Central Universities.

REFERENCES

- [1] 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," in *ECCV*, 2020.
- [2] 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 *CVPR*, 2021.
- [3] F. Lu, Y. Xu, G. Chen, H. Li, K.-Y. Lin, and C. Jiang, "Urban radiance field representation with deformable neural mesh primitives," in *ICCV*, 2023.
- [4] J. Guo, N. Deng, X. Li, Y. Bai, B. Shi, C. Wang, C. Ding, D. Wang, and Y. Li, "StreetSurf: Extending multi-view implicit surface reconstruction to street views," *arXiv preprint arXiv:2306.04988*, 2023.
- [5] K. Rematas, A. Liu, P. P. Srinivasan, J. T. Barron, A. Tagliasacchi, T. Funkhouser, and V. Ferrari, "Urban radiance fields," in *CVPR*, 2022.
- [6] 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 *CVPR*, 2023.
- [7] Y. Liao, J. Xie, and A. Geiger, "KITTI-360: A novel dataset and benchmarks for urban scene understanding in 2D and 3D," *TPAMI*, vol. 45, no. 3, pp. 3292–3310, 2023.
- [8] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes, *et al.*, "Argoverse 2: Next generation datasets for self-driving perception and forecasting," *arXiv preprint arXiv:2301.00493*, 2023.
- [9] K. Zhang, G. Riegler, N. Snively, and V. Koltun, "NeRF++: Analyzing and improving neural radiance fields," *arXiv preprint arXiv:2010.07492*, 2020.
- [10] Z. Wang, S. Wu, W. Xie, M. Chen, and V. A. Prisacariu, "NeRF-: Neural radiance fields without known camera parameters," *arXiv preprint arXiv:2102.07064*, 2021.
- [11] C.-H. Lin, W.-C. Ma, A. Torralba, and S. Lucey, "BARF: Bundle-adjusting neural radiance fields," in *ICCV*, 2021.
- [12] J. T. Barron, B. Mildenhall, M. Tancik, P. Hedman, R. Martin-Brualla, and P. P. Srinivasan, "Mip-NeRF: A multiscale representation for anti-aliasing neural radiance fields," in *CVPR*, 2021.
- [13] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman, "Mip-NeRF 360: Unbounded anti-aliased neural radiance fields," in *CVPR*, 2022.
- [14] A. Pumarola, E. Corona, G. Pons-Moll, and F. Moreno-Noguer, "D-NeRF: Neural radiance fields for dynamic scenes," in *CVPR*, 2021.
- [15] X. Fu, S. Zhang, T. Chen, Y. Lu, L. Zhu, X. Zhou, A. Geiger, and Y. Liao, "Panoptic NeRF: 3D-to-2D label transfer for panoptic urban scene segmentation," in *3DV*, 2022.
- [16] J. Ye, N. Wang, and X. Wang, "FeatureNeRF: Learning generalizable NeRFs by distilling foundation models," in *ICCV*, 2023.
- [17] H. Chen, C. Li, M. Guo, Z. Yan, and G. H. Lee, "GNeSF: Generalizable neural semantic fields," in *NeurIPS*, 2023.
- [18] P. Hedman, P. P. Srinivasan, B. Mildenhall, J. T. Barron, and P. Debevec, "Baking neural radiance fields for real-time view synthesis," in *ICCV*, 2021.
- [19] T. Müller, A. Evans, C. Schied, and A. Keller, "Instant neural graphics primitives with a multiresolution hash encoding," *ToG*, vol. 41, no. 4, pp. 1–15, 2022.
- [20] L. Liu, J. Gu, K. Zaw Lin, T.-S. Chua, and C. Theobalt, "Neural sparse voxel fields," *NIPS*, 2020.
- [21] M. Píala and R. Clark, "TerminNeRF: Ray termination prediction for efficient neural rendering," in *3DV*, 2021.
- [22] T. Neff, P. Stadlbauer, M. Parger, A. Kurz, J. H. Mueller, C. R. A. Chaitanya, A. Kaplanyan, and M. Steinberger, "DONeRF: Towards real-time rendering of compact neural radiance fields using depth oracle networks," in *Computer Graphics Forum*, 2021.
- [23] Q. Xu, Z. Xu, J. Philip, S. Bi, Z. Shu, K. Sunkavalli, and U. Neumann, "Point-NeRF: Point-based neural radiance fields," in *CVPR*, 2022.
- [24] C. Reiser, S. Peng, Y. Liao, and A. Geiger, "KiloNeRF: Speeding up neural radiance fields with thousands of tiny mlps," in *ICCV*, 2021.
- [25] S. J. Garbin, M. Kowalski, M. Johnson, J. Shotton, and J. Valentin, "FastNeRF: High-fidelity neural rendering at 200fps," in *ICCV*, 2021.
- [26] A. Yu, R. Li, M. Tancik, H. Li, R. Ng, and A. Kanazawa, "PlenOctrees for real-time rendering of neural radiance fields," in *ICCV*, 2021.
- [27] C. Sun, M. Sun, and H.-T. Chen, "Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction," in *CVPR*, 2022.
- [28] S. Fridovich-Keil, A. Yu, M. Tancik, Q. Chen, B. Recht, and A. Kanazawa, "Plenoxels: Radiance fields without neural networks," in *CVPR*, 2022.
- [29] Z. Chen, T. Funkhouser, P. Hedman, and A. Tagliasacchi, "MobileNeRF: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures," in *CVPR*, 2023.
- [30] A. Chen, Z. Xu, A. Geiger, J. Yu, and H. Su, "Tensorf: Tensorial radiance fields," in *ECCV*, 2022.
- [31] H. Yan, C. Liu, C. Ma, and X. Mei, "PlenVDB: Memory efficient VDB-based radiance fields for fast training and rendering," in *CVPR*, 2023.
- [32] P. Wang, Y. Liu, Z. Chen, L. Liu, Z. Liu, T. Komura, C. Theobalt, and W. Wang, "F²-NeRF: Fast neural radiance field training with free camera trajectories," in *CVPR*, 2023.
- [33] 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 *CVPR*, 2022.
- [34] J. Ost, I. Laradji, A. Newell, Y. Bahat, and F. Heide, "Neural point light fields," in *CVPR*, 2022.
- [35] A. Kundu, K. Genova, X. Yin, A. Fathi, C. Pantofaru, L. J. Guibas, A. Tagliasacchi, F. Dellaert, and T. Funkhouser, "Panoptic neural fields: A semantic object-aware neural scene representation," in *CVPR*, 2022.
- [36] X. Zhang, A. Kundu, T. Funkhouser, L. Guibas, H. Su, and K. Genova, "Nerflits: Local radiance fields for efficient structure-aware 3D scene representation from 2D supervision," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [37] N. Max, "Optical models for direct volume rendering," *IEEE Transactions on Visualization and Computer Graphics*, vol. 1, no. 2, pp. 99–108, 1995.
- [38] Z. Hao, A. Mallya, S. Belongie, and M.-Y. Liu, "GANcraft: Unsupervised 3D neural rendering of minecraft worlds," in *ICCV*, 2021.
- [39] B. T. Phong, "Illumination for computer generated pictures," *Communications of ACM*, pp. 95–101, 1998.
- [40] D. P. Greenberg, K. E. Torrance, P. Shirley, J. Arvo, E. Lafortune, J. A. Ferwerda, B. Walter, B. Trumbore, S. Pattanaik, and S.-C. Foo, "A framework for realistic image synthesis," in *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, 1997.
- [41] X. Zhang, P. P. Srinivasan, B. Deng, P. Debevec, W. T. Freeman, and J. T. Barron, "NeRFactor: Neural factorization of shape and reflectance under an unknown illumination," *TOG*, vol. 40, no. 6, pp. 1–18, 2021.
- [42] M. Pharr, W. Jakob, and G. Humphreys, *Physically based rendering: From theory to implementation*. MIT Press, 2023.
- [43] X. Chen and K. He, "Exploring simple siamese representation learning," in *CVPR*, 2021.
- [44] S. Fridovich-Keil, G. Meanti, F. R. Warburg, B. Recht, and A. Kanazawa, "K-planes: Explicit radiance fields in space, time, and appearance," in *CVPR*, 2023.
- [45] J. Kulhanek and T. Sattler, "Tetra-NeRF: Representing neural radiance fields using tetrahedra," in *ICCV*, 2023.
- [46] M. Tancik, E. Weber, E. Ng, R. Li, B. Yi, T. Wang, A. Kristoffersen, J. Austin, K. Salahi, A. Ahuja, *et al.*, "NeRFstudio: A modular framework for neural radiance field development," in *SIGGRAPH*, 2023.
- [47] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han, "On the variance of the adaptive learning rate and beyond," *arXiv preprint arXiv:1908.03265*, 2019.
- [48] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [49] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *TIP*, vol. 13, no. 4, pp. 600–612, 2004.
- [50] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *CVPR*, 2018.