

Reconstructing Objects in-the-wild for Realistic Sensor Simulation

Ze Yang^{1,2}, Sivabalan Manivasagam^{1,2}, Yun Chen^{1,2}, Jingkang Wang^{1,2}, Rui Hu¹, Raquel Urtasun^{1,2}
 Waabi¹, University of Toronto²
 {zeyang, manivasagam, yun, wangjk, urtasun}@cs.toronto.edu, rhu@waabi.ai

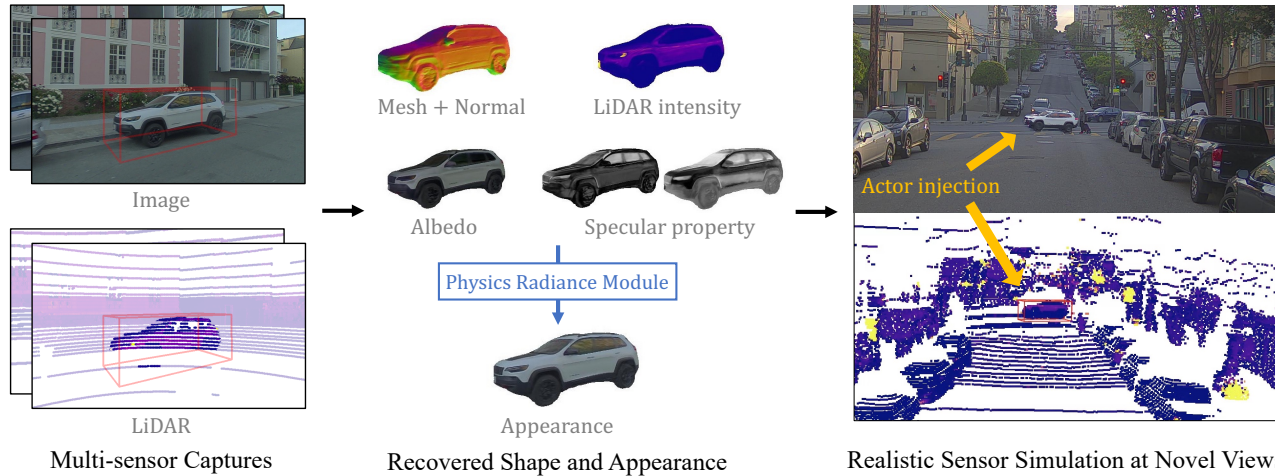


Fig. 1. Given a camera video and LiDAR sweeps as input, our model reconstructs accurate geometry and surface properties, which can be used to synthesize realistic appearance under novel viewpoints using our physics-based radiance module, enabling realistic sensor simulation for self-driving.

Abstract—Reconstructing objects from real world data and rendering them at novel views is critical to bringing realism, diversity and scale to simulation for robotics training and testing. In this work, we present NeuSim, a novel approach that estimates accurate geometry and realistic appearance from sparse in-the-wild data captured at distance and at limited viewpoints. Towards this goal, we represent the object surface as a neural signed distance function and leverage both LiDAR and camera sensor data to reconstruct smooth and accurate geometry and normals. We model the object appearance with a robust physics-inspired reflectance representation effective for in-the-wild data. Our experiments show that NeuSim has strong view synthesis performance on challenging scenarios with sparse training views. Furthermore, we showcase composing NeuSim assets into a virtual world and generating realistic multi-sensor data for evaluating self-driving perception models. The supplementary material can be found at the project website: <https://waabi.ai/research/neusim/>

I. INTRODUCTION

Simulation is key for testing at scale self-driving systems. To allow for end-to-end closed loop testing, the simulator should produce in real-time realistic sensor data, which are rendered views of a 3D virtual world. Most simulators employ CAD models [1], [2], which have unrealistic appearance, require costly manual construction and cannot scale to represent the diversity and complexity of the real world. To address this issue, we focus on automatically reconstructing high-quality 3D assets cost-efficiently from sparse, in-the-wild, multi-sensor data captured by a moving platform along constrained trajectories. The reconstructed assets should have accurate shape and appearance, and should render efficiently.

A recent promising approach to asset shape reconstruct-

tion and novel view synthesis is Neural Radiance Fields (NeRF) [3], which represents the scene as a continuous density and radiance field parameterized by neural networks, and leverages volume rendering [4] to render the scene. NeRF produces renderings that match the observed images; however, the resulting asset representation is not suitable for scalable virtual world creation from in-the-wild data due to shape radiance ambiguity [5], noisy geometry [6], [7], incomplete reconstruction (see Fig. 2) and significant artifacts and distortion at extrapolated viewpoints (see Fig. 3). Several works have aimed to improve NeRF [5]–[8] and adapt it to outdoor scenes [5], [9], but they focus on either synthetic data or small scenes densely captured in controlled environments. In contrast, we focus on reconstruction from data captured by moving platforms (*e.g.*, self-driving vehicle) in outdoor environments, which are more challenging due to the sparsity and limited range of sensor viewpoints, varying resolution and distance from the sensor, and sensor noise.

In this paper, we propose NeuSim, a novel neural volume rendering approach that takes sparse, in-the-wild image and LiDAR data and learns an asset’s shape and surface properties for robust and realistic multi-sensor simulation. NeuSim is composed of a neural geometry representation that generates precise surfaces and models appearance via a physics-based radiance model, which accurately learns texture and reflectance. NeuSim also incorporates structural symmetry priors for common traffic actors (primarily cars, trucks) to learn the surface properties of the unobserved side (see Fig. 2), enabling robust novel view synthesis and seamless simulation of new scenarios. NeuSim not only

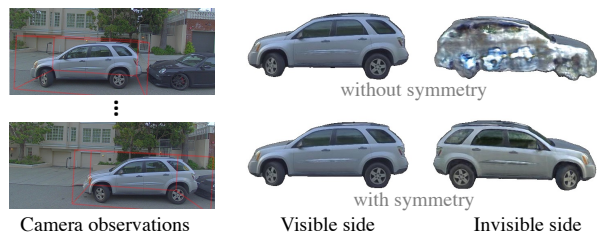


Fig. 2. The structural symmetry priors helps to reconstruct shape and appearance on unseen regions for common traffic objects (e.g., cars, trucks).




Fig. 3. Prior MLP radiance model [3], [6], [7] cannot generalize well to large viewpoint change though the underlying geometry is correct. NeuSim’s Physics Radiance Module models the appearance robustly and realistically.

renders and learns images, but also predicts LiDAR depths and intensities, resulting in better geometry than assets generated from images alone, and enabling consistent multi-sensor simulation, which is key for self-driving testing, as modern perception systems exploit multiple sensors for robustness. Our factorized geometry and radiance model easily bakes into an explicit mesh allowing for easy modification, (1000x) faster rendering, and integration into existing simulators.

Our experiments show that NeuSim reconstructs high-quality assets from in-the-wild data that render efficiently in existing graphics engines. Finally, we leverage NeuSim assets to generate sensor data for testing perception systems.

II. RELATED WORK

1) *Assets for Sensor Simulation*: Self-driving simulators such as CARLA [1] and AirSim [2] leverage artist-created 3D CAD models to build virtual worlds. While these assets have clean geometry and are easily modifiable, they are expensive to create, have limited diversity and scale, and lack realism [10]. Several works leverage in-the-wild data from driving scenes to build assets at scale. LiDARSim and SurfGAN [11], [12] aggregate LiDAR across frames to build diverse textured surfel representations, but have noisy geometry and appearance. Other works perform shape completion and texture estimation using a learned asset representation. Several leverage synthetic CAD data [13]–[17] as a prior, and then generate the final asset either through feed-forward prediction or optimization on real data. These approaches are limited by the set of CAD models used and are often not photorealistic. GeoSim [18] deforms a template mesh and leverages image warping for photorealistic video simulation, but the assets do not model reflectance and have overly smooth geometry due to its limited mesh representation. DriveGAN [19] represents assets as disentangled latent codes and generates video from control inputs, allowing for fully differentiable simulation, but is limited in its realism and is not temporally consistent. In contrast, we build realistic

assets that have high-fidelity geometry, can be inserted into new and diverse scenarios at scale, and are temporally and multi-sensor consistent.

2) *Neural Volume Rendering*: Neural implicit representation [20]–[28] and volumetric rendering approaches [4], [29], such as Neural Radiance Fields (NeRF) [3] have achieved high quality image rendering given dense and uniform training views, but can have noisy geometry and artifacts, especially when trained on sparse and limited viewpoints. To learn a better geometry, recent works have assumed solid closed surface objects and combined SDFs with radiance fields [6], [7], [30]–[32] and leverage 3D information such as RGB-D measurements [31], [33], resulting in improved geometry. We also leverage an SDF representation, and incorporate LiDAR data as supervision. Our work combines these enhancements with a learned reflectance model to learn more robust appearance (see Fig. 3) for in-the-wild data.

Several works learn a reflectance model and estimate scene lighting [34]–[38], but they focus on improving controllability of NeRF and learn on synthetic data or controlled environments. We leverage a simple yet effective Phong illumination model to model reflectance [39], [40], and supervise with multi-sensor data (RGB and LiDAR intensity), to better render objects at novel views from sparse in-the-wild data. We also incorporate structural priors such as symmetry over the asset shape and surface properties to complete the missing information due to limited sensor observations.

Other works have focused on the composability of neural radiance fields [41]–[46], but do not model the reflectance and do not generate 3D geometry for multi-sensor simulation. In contrast, our approach models reflectance and can be baked into an explicit textured mesh for fast rendering.

3) *Inverse Graphics*: Several recent works generate assets from in-the-wild data with decomposed shape, material, and illumination via differentiable mesh rendering [38], [47], [48]. These approaches build on intrinsic decomposition, the task of decomposing the image into albedo, illumination, etc. based on image formation [49]–[51]. Our work is most similar to NeRS [48], which estimates shape through a deformable implicit surface representation and appearance with a learned Phong reflectance model. However, our use of volume rendering and an SDF representation results in more accurate assets, as shown in our experiments.

III. METHOD

Our goal is to automatically build rigid object assets from in-the-wild data for sensor simulation. Given camera images and LiDAR point clouds captured by a moving platform, we want to learn the object’s shape and appearance. Towards this goal, we propose NeuSim, a novel approach that is composed of a structured neural surface representation and a physics-based reflectance model. This decomposed representation enables generalization to new views from sparse in-the-wild viewpoints. Fig. 4 provides an overview of NeuSim. First, we briefly review NeRF. Next, we introduce Neural Surface Modeling for accurate surface and normal estimation. We then introduce our Physics-based Radiance Module for ro-

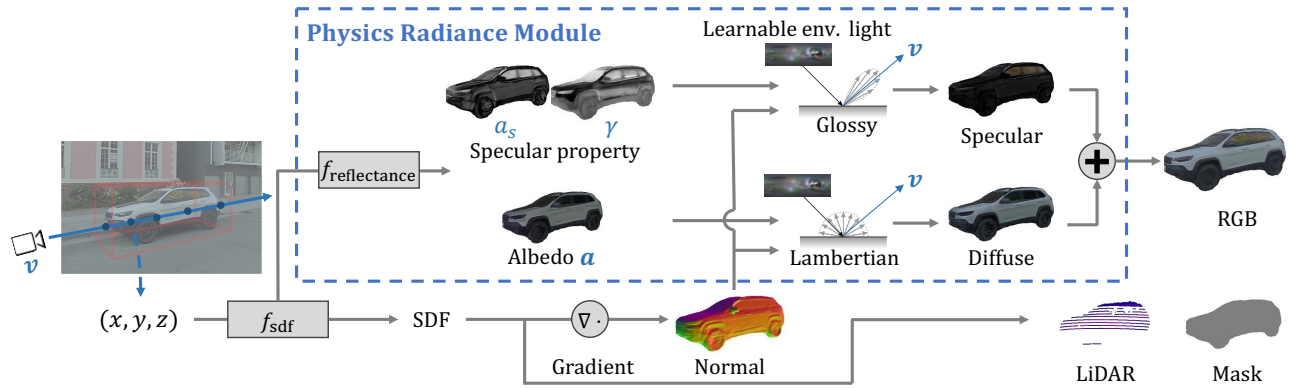


Fig. 4. Given a continuous 3D location, NeuSim outputs the signed distance value of the point to object surface, the albedo, and the specular property. The signed distance value is used to derive the surface normal, which is then used to shade the diffuse and specular components to obtain the final RGB color. We also render the LiDAR depth and intensity, as well as object mask from the learned representation.

bust and realistic appearance modelling. Finally we present how our model performs rendering and learns from data.

A. Review of NeRF Representations

NeRF [3] represents the scene with a Multi-Layer Perceptron (MLP) that maps a point location $\mathbf{x} \in \mathbb{R}^3$ and viewing direction $\mathbf{d} \in \mathbb{R}^3$ to a volume density $\sigma(\mathbf{x}) \in \mathbb{R}$ and RGB color radiance $\mathbf{c}(\mathbf{x}, \mathbf{d}) \in \mathbb{R}^3$. The color is conditioned on the viewing direction to model view-dependent effects. Given a posed camera, NeRF performs volume rendering by evaluating the MLP at points along a camera ray to compute pixel color. We define each pixel’s ray as $\{\mathbf{r}(t) = \mathbf{o} + t\mathbf{d} \mid t > 0\}$, where \mathbf{o} is camera center and \mathbf{d} is the viewing direction originating from the camera center. Given N sample points $\{\mathbf{x}_i\}_{i=1}^N$ along the camera ray, the pixel color is

$$\mathbf{C}(\mathbf{r}) = \sum_{i=1}^N \alpha(\mathbf{x}_i) T(\mathbf{x}_i) \mathbf{c}(\mathbf{x}_i, \mathbf{d}) \quad (1)$$

$$\alpha(\mathbf{x}_i) = 1 - \exp(-\sigma(\mathbf{x}_i) \delta_i), \quad T(\mathbf{x}_i) = \prod_{j=1}^{i-1} (1 - \alpha(\mathbf{x}_j))$$

where $\alpha(\mathbf{x}_i)$ is the alpha compositing computed using the density prediction $\sigma(\mathbf{x})$ and the distance between adjacent samples $\delta_i = \|\mathbf{x}_{i+1} - \mathbf{x}_i\|_2$, and $T(\mathbf{x}_i)$ is the accumulated transmittance along the ray.

B. Neural Surface Modeling

NeRF-based approaches [3], [5], [52] generate photorealistic renderings, but are unable to reconstruct accurate geometry, especially when trained on sparse in-the-wild views. Rather than model the space as a heterogenous soft volume as in NeRF [3], we take inspiration from [7], [20], [30], [31] and assume the object of interest has a topologically closed surface that we represent as the zero-level set of a signed distance function (SDF), parameterized by an MLP. This representation is effective for most objects of interest in self-driving simulations (e.g., vehicles, construction elements).

We define an SDF MLP network $f_{\text{SDF}} : \mathbf{x} \rightarrow s$ mapping a point in 3D space $\mathbf{x} \in \mathbb{R}^3$ to its signed distance $s(\mathbf{x}) \in \mathbb{R}$ from the object surface. The surface \mathcal{S} of the object can then be defined as the zero-level set of the SDF function:

$$\mathcal{S} = \{\mathbf{x} \in \mathbb{R}^3 \mid s(\mathbf{x}) = 0\}. \quad (2)$$

And the normal is derived as the SDF’s gradient, $\mathbf{n} = \nabla s(\mathbf{x})$.

To perform volume rendering (Eq. (1)), we need to convert the signed distance to an alpha value $\alpha(\mathbf{x}) \in [0, 1]$, where $\alpha(\mathbf{x}) = 1$ if \mathbf{x} is inside the object, and $\alpha(\mathbf{x}) = 0$ otherwise. Towards this goal, we use a sigmoid-like function:

$$\alpha(\mathbf{x}) = \frac{1}{1 + \exp(\beta \cdot s(\mathbf{x}))}, \quad (3)$$

where $\alpha(\mathbf{x})$ transitions from 0 to 1 when $s(\mathbf{x})$ transitions from positive (outside) to negative (inside). To reconstruct solid objects, the converting function should model a step function at the object’s surface, i.e., $\beta \rightarrow \infty$. To prevent vanishing gradients, we make β a learnable parameter.

C. Robust and Realistic Appearance Modelling

Given our surface geometry representation, we now discuss how we model object appearance. We represent the appearance of a 3D point \mathbf{x} viewed from direction \mathbf{d} as a radiance field $\mathbf{c}(\mathbf{x}, \mathbf{d})$, a function of the material properties f_r , the radiance from the environment $L(\mathbf{x}, \omega)$ (with ω the incoming radiance direction), and the surface normals $\mathbf{n}(\mathbf{x})$. $\mathbf{c}(\mathbf{x}, \mathbf{d})$ is estimated via rendering equation [53], [54]:

$$\mathbf{c}(\mathbf{x}, \mathbf{d}) = \int_{\Omega} f_r(\mathbf{x}, \omega, \mathbf{d}) L(\mathbf{x}, \omega) (\omega \cdot \mathbf{n}) d\omega \quad (4)$$

where $f_r(\mathbf{x}, \omega, \mathbf{d})$ is the bidirectional reflectance distribution function (BRDF), which defines the proportion of light reflected from the incoming light direction ω to the viewing direction \mathbf{d} at position \mathbf{x} . The integral is computed over the hemisphere Ω centered at location \mathbf{x} and oriented towards the normal \mathbf{n} . Since in outdoor settings the light mainly comes from the sun, we model the scene lighting as a set of infinitely far-away directional lights, and ignore attenuation effects to the incoming light intensity. Thus the light function is independent of the point location, i.e., $L(\mathbf{x}, \omega) = L(\omega)$.

NeRF models the radiance equation using an MLP that directly maps a point \mathbf{x} and view direction \mathbf{d} to the emitted radiance $\mathbf{c}(\mathbf{x}, \mathbf{d})$. This does not generalize well to larger viewpoint variation (see Fig. 3), especially when trained on a sparse set of images. The MLP radiance lacks knowledge of the underlying light transport physics and overfits. To address this, we model the radiance function using a simple




Fig. 5. Rendering of normal derived from density and SDF representation.

yet effective Phong illumination model [39], [40] with spatial varying albedo and shininess. Our framework also supports other PBR models (e.g., microfacet [55]). The Phong model consists of diffuse and specular components:

$$\mathbf{c}_{\text{phy}}(\mathbf{x}, \mathbf{d}) = \int_{\Omega} \left(\underbrace{\mathbf{a}(\mathbf{x})(\omega \cdot \mathbf{n})}_{\text{Diffuse term}} + \underbrace{f_s(\mathbf{x}, \mathbf{d})}_{\text{Specular term}} \right) \underbrace{L(\omega)}_{\text{Light}} d\omega \quad (5)$$

The diffuse component depends on the albedo $\mathbf{a}(\mathbf{x}) \in \mathbb{R}^3$. The specular component $f_s(\mathbf{x}, \mathbf{d})$ is view-dependant:

$$f_s(\mathbf{x}, \mathbf{d}) = a_s(\mathbf{x}) \frac{\gamma(\mathbf{x}) + 1}{2\pi} (\mathbf{r}_{\omega} \cdot \mathbf{d})^{\gamma(\mathbf{x})} \quad (6)$$

where $a_s(\mathbf{x}) \in \mathbb{R}$ is the specular albedo measuring the specular highlight intensity, $\mathbf{r}_{\omega} = 2(\omega \cdot \mathbf{n})\mathbf{n} - \omega$ is the reflected light direction, and $\gamma(\mathbf{x})$ is the shininess.

We predict the diffuse albedo $\mathbf{a}(\mathbf{x})$, specular albedo $a_s(\mathbf{x})$ and shininess $\gamma(\mathbf{x})$ using reflectance MLPs $f_{\text{reflectance}}$. We use a learnable 2D environment map $\mathcal{E} \in \mathbb{R}^{A \times E \times 3}$ to parameterize the environment light $L(\omega)$, where A and E are the azimuth and elevation resolution respectively. \mathcal{E} stores the light intensity for each discrete incoming light direction ω . We numerically estimate integral in Eq. (5) as the sum of radiance values from discrete incoming lights $\mathcal{E}_{\omega} = L(\omega)$:

$$\mathbf{c}_{\text{phy}} \approx \mathbf{a} \sum_{\omega} (\omega \cdot \mathbf{n}) \mathcal{E}_{\omega} + \frac{a_s(\gamma + 1)}{2\pi} \sum_{\omega} (\mathbf{r}_{\omega} \cdot \mathbf{d})^{\gamma} \mathcal{E}_{\omega} \quad (7)$$

$\mathbf{c}_{\text{phy}}(\mathbf{x}, \mathbf{d})$ relies on accurate normal estimation \mathbf{n} . Prior NeRF-based methods estimates the normal [34], [35], [37] by taking the normalized gradient of volume density w.r.t. input 3D location, but they are noisy (see Fig. 5). Other methods use a MLP to predict the normal vector and enforce smoothness via regularization [34], [36]. These predicted normals tend to be smoother, but they are decoupled from the actual shape and do not reflect the exact geometry. In contrast, our neural SDF representation ensures smooth and accurate normal estimated as the SDF's gradient, $\mathbf{n} = \nabla_s(\mathbf{x})$.

D. Rendering with NeuSim

1) *Scene Representation*: As we focus on object reconstruction, we assume that the object of interest in the scene is bounded by a cuboid and we only render in the frustum generated between the viewing plane and the projected cuboid [41]. We assume the rendered ray $r(t)$ intersects the cuboid at t_{near} and t_{far} . We divide the traversed space into foreground ($\{t_{\text{near}} < t < t_{\text{far}}\}$) and background. To decouple the object sensor observations from the background's, we take inspiration from NeRF++ [5] and model them with separate networks. We use our proposed model to represent the foreground and use NeRF with an inverted sphere parameterization [5] to represent the background. For rendering only the asset, the background network is discarded.

2) *Rendering Camera*: To render the RGB observations, we first draw stratified samples on the ray and query the SDF MLPs to compute $\alpha(\mathbf{x})$ as in Eq. (3). For each sampled point, we query the reflectance MLPs to compute the albedo $\mathbf{a}(\mathbf{x})$, $a_s(\mathbf{x})$ and shininess $\gamma(\mathbf{x})$. We then compute the foreground radiance $\mathbf{c}(\mathbf{x}, \mathbf{d})$ for each sampled point \mathbf{x} with view direction \mathbf{d} , using Eq. (7) and the learned environment map \mathcal{E} . To compute the background scene radiance, we sample the ray's intersections with Multiple-Sphere Images (MSI) surrounding the object of interest. We generate the radii for MSI by linearly interpolating inverse depths. The rendered RGB is computed by alpha compositing using Eq. (1).

3) *Rendering LiDAR Depth and Intensity*: Given a query LiDAR ray, we sample points along the ray similar to camera rendering, and then render the LiDAR depth and intensity as:

$$D(\mathbf{r}) = \sum_{i=1}^N \alpha(\mathbf{x}_i) T(\mathbf{x}_i) d_i, \quad I(\mathbf{r}) = \sum_{i=1}^N \alpha(\mathbf{x}_i) T(\mathbf{x}_i) i(\mathbf{x}). \quad (8)$$

where $d_i = \|\mathbf{x}_i - \mathbf{o}\|_2$ is the depth value of sample point \mathbf{x}_i on the ray originating from \mathbf{o} . We add another branch on the reflectance MLP to predict the intensity value $i(\mathbf{x}) \in \mathbb{R}^+$.

4) *Rendering Object Mask*: We render the object mask to help provide additional signal on object boundaries. We estimate the rendered foreground probability as the aggregated weights in the foreground intervals:

$$m(\mathbf{r}) = \sum_{t_{\text{near}} < t < t_{\text{far}}} \alpha(\mathbf{o} + t\mathbf{d}) T(\mathbf{o} + t\mathbf{d}). \quad (9)$$

E. Learning NeuSim

To learn the model, we minimize the difference between the sensor observations and our rendered outputs. We leverage RGB color ($\mathcal{L}_{\text{color}}$), LiDAR point clouds ($\mathcal{L}_{\text{lidar}}$) and object foreground masks ($\mathcal{L}_{\text{mask}}$). We also add an Eikonal term (\mathcal{L}_{Eik}) to regularize the predicted SDF, and a symmetry term (\mathcal{L}_{sym}) for vehicle objects. The full training loss is:

$$\mathcal{L} = \mathcal{L}_{\text{color}} + \lambda_{\text{lidar}} \mathcal{L}_{\text{lidar}} + \lambda_{\text{mask}} \mathcal{L}_{\text{mask}} + \lambda_{\text{Eik}} \mathcal{L}_{\text{Eik}} + \lambda_{\text{sym}} \mathcal{L}_{\text{sym}} \quad (10)$$

For each asset, we train the shape and reflectance networks and environment lighting map jointly via gradient descent over ray batches of size N randomly sampled from the sensor data. We now review each loss term.

1) *RGB supervision*: Similar to NeRF, we want to ensure that the rendered pixels match the observed ones. The camera image loss $\mathcal{L}_{\text{color}}$ is defined as:

$$\mathcal{L}_{\text{color}} = \frac{1}{N} \sum_{i=1}^N \left\| \mathbf{C}(\mathbf{r}_i) - \hat{\mathbf{C}}_i \right\|_1 \quad (11)$$

where $\mathbf{C}(\mathbf{r}_i)$ is rendered color and $\hat{\mathbf{C}}_i$ is observed color.

2) *LiDAR supervision*: We leverage LiDAR depth measurements to supervise the SDF field for more accurate geometry, as well as intensity to learn better shape and surface properties. The LiDAR loss $\mathcal{L}_{\text{lidar}}$ is defined as:

$$\mathcal{L}_{\text{lidar}} = \frac{1}{N} \sum_{i=1}^N \left(\left\| D(\mathbf{r}_i) - \hat{D}_i \right\|_2 + \lambda_{\text{int}} \left\| I(\mathbf{r}_i) - \hat{I}_i \right\|_2 \right) \quad (12)$$

Method	MSE↓	PSNR↑	SSIM↑	LPIPS↓
SI-ViewWarp [57]	0.0233	17.51	0.514	0.371
SAMP [13]	0.0144	19.52	0.628	0.283
NeRS [48]	0.0176	18.49	0.562	0.265
NVDiffRec [58]	0.0114	20.46	0.593	0.396
NeRF++ [5]	0.0138	20.86	0.611	0.300
NeuS [6]	0.0115	21.37	0.640	0.247
Ours	0.0081	22.44	0.692	0.202

TABLE I

QUANTITATIVE COMPARISONS ON PANDAVEHICLE FOR NVS TASK.

where $D(\mathbf{r}_i)$ is the rendered LiDAR depth and $I(\mathbf{r}_i)$ is the rendered LiDAR intensity. \hat{D}_i and \hat{I}_i are the LiDAR depth and intensity observations, respectively. We also penalize large weight predictions that are far from the observations:

$$\mathcal{L}_{\text{lidar}} = \mathcal{L}_{\text{lidar}} + \frac{1}{N} \sum_{i=1}^N \left(\sum_{\|d_{ij} - \hat{D}_i\|_2 > \epsilon} w_{ij} \right) \quad (13)$$

3) *Object mask supervision*: Foreground masks identify the object in the image. The mask loss $\mathcal{L}_{\text{mask}}$ is defined as:

$$\mathcal{L}_{\text{mask}} = \frac{1}{N} \sum_{i=1}^N \|m(\mathbf{r}_i) - \hat{m}_i\|_2, \quad (14)$$

where $m(\mathbf{r}_i)$ is the rendered foreground probability and \hat{m}_i is the estimated mask from an off-the-shelf algorithm [56].

4) *Eikonal Regularizer*: This term encourages the SDF to satisfy the Eikonal equation and generate unit normals [28]:

$$\mathcal{L}_{\text{Eik}} = \frac{1}{M} \sum_{i=1}^M (\|\nabla s(\mathbf{x}_i)\|_2 - 1)^2 \quad (15)$$

where $\{s(\mathbf{x}_i)\}_{i=1}^M$ is the predicted signed distance in a batch.

5) *Structural Symmetry Prior*: To reconstruct unseen regions, we incorporate symmetry priors for common traffic objects (e.g., cars, trucks). Although the radiance is not symmetric due to diffuse and specular shading with lighting, the surface geometry $s(\mathbf{x})$ and material properties ($\mathbf{a}(\mathbf{x})$, $a_s(\mathbf{x})$, $\gamma(\mathbf{x})$) are approximately symmetric. We denote the transform from world coordinate to the canonical object coordinate (Front-Left-Up) as $\mathbf{T} = \begin{pmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0} & 1 \end{pmatrix} \in SE(3)$. For each query point \mathbf{x} with normal \mathbf{n} in world coordinate, the symmetrized point \mathbf{x}' and normal \mathbf{n}' are:

$$\mathbf{x}' = \mathbf{T}^{-1} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \mathbf{T}\mathbf{x}, \quad \mathbf{n}' = \mathbf{R}^{-1} \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \mathbf{R}\mathbf{n} \quad (16)$$

We jointly optimize \mathbf{T} during training, the symmetry loss is:

$$\mathcal{L}_{\text{sym}} = \frac{1}{M} \sum_{i=1}^M \|s(\mathbf{x}_i) - s(\mathbf{x}'_i)\|_2 + \|\mathbf{n}'(\mathbf{x}_i) - \mathbf{n}(\mathbf{x}'_i)\|_2 + \|\mathbf{a}(\mathbf{x}_i) - \mathbf{a}(\mathbf{x}'_i)\|_2 + \|a_s(\mathbf{x}_i) - a_s(\mathbf{x}'_i)\|_2 + \|\gamma(\mathbf{x}_i) - \gamma(\mathbf{x}'_i)\|_2 \quad (17)$$

IV. EXPERIMENTAL EVALUATION

In this section we demonstrate our model performance on in-the-wild data. We first introduce our experimental settings. We then compare our model against state-of-the-art methods, and also ablate our design choices. Finally, we apply NeuSim for fast and realistic sensor simulation for self-driving.

Supervision	MSE↓	PSNR↑	SSIM↑	LPIPS↓
img	0.0131	21.07	0.646	0.269
img + lidar	0.0089	22.07	0.677	0.211
img + lidar + mask	0.0091	22.13	0.679	0.198
img + lidar + mask + sym	0.0081	22.44	0.692	0.202

TABLE II

ABLATION ON LEARNING SUPERVISION FOR NEUSIM.

Radiance Model	MSE↓	PSNR↑	SSIM↑	LPIPS↓
MLP Radiance	0.0118	21.31	0.636	0.232
Physics-based Radiance	0.0091	22.13	0.679	0.198

TABLE III

ABLATION ON RADIANCE MODEL FOR APPEARANCE PREDICTION.

A. Experimental Setting

We focus on recovering the shape and appearance of vehicles, the most common actor in self-driving scenes, and evaluate on the task of novel view synthesis. We curated 10 vehicles with diverse shape and appearance under complex illumination from the PandaSet [59] to derive the PandaVehicle dataset. This dataset has calibrated LiDARs and multiple cameras. We use the left camera for training and the front-left camera for evaluation. Each asset has on average ~ 24 training views. PandaVehicle is more challenging than existing novel view synthesis datasets due to limited range and number of viewpoints available for training, as well as the complex illumination. We evaluate the NVS performance using Mean-Square Error (MSE), Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) [60]. Since we focus on assets, we use the predicted segmentation mask [56] to only evaluate the foreground pixels.

B. Novel View Synthesis

1) *State-of-the-art (SoTA) Comparison*: We compare our model with state-of-the-art Neural Radiance Fields based approach NeRF++ [5], NeuS [6] and inverse graphics model NeRS [48], NVDiffRec [58]. We choose these baselines as they model view-dependent appearance and work well in our outdoor setting. We also compare against non-learning based LiDAR guided single view warping model [57] and CAD-based model SAMP [13]. As shown in Table I, our model achieves the best performance across all metrics. A qualitative comparison is depicted by Fig. 6. Our model generalizes better to large viewpoint changes when compared to NeRF-based method [5], [6], demonstrating the value of our physics-based reflectance module. Our model also captures more fine-grained details than inverse graphics model [48] due to the expressive implicit representation.

2) *Ablation on Learning Supervision*: We study the effect of LiDAR, mask and symmetry supervision in Table II. Incorporating LiDAR improves performance since the additional depth and intensity measurements help learn better geometry and reflectance. Mask and symmetry supervision does not significantly improve metrics but it helps separate objects from the ground and complete unseen regions.



Fig. 6. Qualitative comparisons on PandaVehicle for novel view synthesis. For each vehicle, we show results at novel views with large view variation.

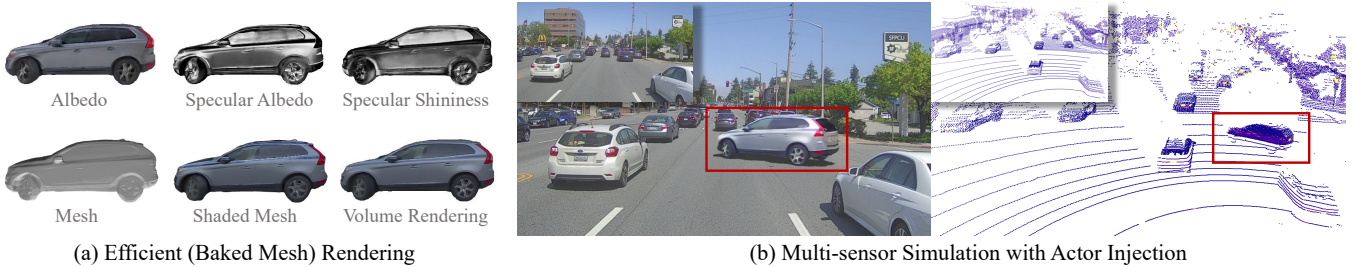


Fig. 7. NeuSim reconstructs disentangled shape and appearance, enabling (a) efficient mesh rendering by baking diffuse and specular properties into an explicit mesh, which is 1000x faster than volume rendering, (b) and multi-sensor simulation for self-driving via actor insertion.

3) *Ablation on Radiance Model*: We study the effect of the radiance model in Table III. For the MLP setting, we replace the reflectance model with a 4-layer viewpoint-dependant MLP predicting the color similar to [6]. Our reflectance model achieves better performance, demonstrating the benefits of a physics-based decomposed representation.

C. Downstream Applications

1) *Efficient Rendering*: Our approach recovers geometry, surface texture and specular material that can be baked into an explicit mesh for efficient rendering. We perform marching cubes to generate the mesh from the SDF representation [61], and evaluate on each vertex to compute the per-vertex albedo, specular albedo, and shininess mappings. We render the explicit mesh with OpenGL using a customized shader implementing Eq. (7). In Fig. 7 (a), we show the reconstructed albedo and specular material, followed by a visual comparison between using mesh rendering and volume rendering. Mesh rendering is 1000x faster than volume rendering (~ 77 FPS vs. ~ 0.03 FPS on consumer GPU GTX 1080 Ti and Ubuntu OS) and still provides good visual quality. This enables efficient simulation.

2) *Downstream Evaluation on Camera Simulation*: We evaluate the object detection and instance segmentation algorithm on the simulated camera rendering at novel viewpoints. Specifically, we insert and blend the actors into background image and replace the existing actors. Then we compute the instance-level IoU of predicted bounding box and segmentation mask between real image and simulated image. This detection/segmentation agreement metric indicates how well we can use the sensor simulation to test existing perception systems. As shown in Tab. IV, NeuSim achieves highest agreement for both detection and instance segmentation.

3) *Realistic Sensor Simulation*: Using our reconstructed asset from NeuSim, we can create consistent multi-sensor simulations for self-driving. For camera simulation, we render the asset to the target view and then apply a post-composition network [18] to seamlessly blend the actor to

Method	Detection (IoU) \uparrow	Inst. Segmentation (IoU) \uparrow
SAMP [13]	90.39	89.58
NVDiffRec [58]	85.26	85.88
NeRF++ [5]	92.81	93.22
NeuS [6]	93.97	94.22
Ours	94.82	95.48

TABLE IV

EVALUATE DOWNSTREAM PERCEPTION TASKS ON CAMERA SIMULATION

the background. For LiDAR simulation, we use an approach similar to [62], [63] and perform actor injection by raycasting the asset according to the LiDAR calibration and removing points in the real LiDAR sweep that are occluded by the added actor. Fig. 7 (b) shows that we can generate realistic camera and LiDAR simulations for the added actor, enabling diverse data generation and end to end autonomy testing.

V. CONCLUSION

In this paper, we propose NeuSim, a novel approach for 3D object reconstruction and novel view synthesis from in-the-wild camera and LiDAR data. NeuSim represents the object geometry as a neural SDF, and the appearance with a physics-based reflectance model. With this decomposed representation, we can realistically and efficiently render assets at novel views. We demonstrated that NeuSim assets can be inserted into new scenarios, generate realistic multi-sensor data, and can be used to evaluate autonomy perception, enabling scalable and diverse simulation for self-driving. Future work involves explicitly modelling scene lighting [34], [64], large-scale and dynamic scene [41], [65], [66], efficient training and real-time rendering [67], [68], and dealing with inaccurate sensor poses [69], [70].

VI. ACKNOWLEDGEMENTS

We are grateful to Anqi Joyce Yang, Ioan-Andrei Barsan and Wei-Chiu Ma for profound discussion, constructive feedback and proofreading. We also thank the Waabi team for their valuable assistance and support.

REFERENCES

- [1] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in *CoRL*, 2017. 1, 2
- [2] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "AirSim: High-fidelity visual and physical simulation for autonomous vehicles," in *Field and service robotics*, 2018. 1, 2
- [3] 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. 1, 2, 3
- [4] J. T. Kajiya and B. P. Von Herzen, "Ray tracing volume densities," 1984. 1, 2
- [5] K. Zhang, G. Riegler, N. Snavely, and V. Koltun, "NeRF++: Analyzing and improving neural radiance fields," *arXiv*, 2020. 1, 3, 4, 5, 6
- [6] P. Wang, L. Liu, Y. Liu, C. Theobalt, T. Komura, and W. Wang, "Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction," 2021. 1, 2, 5, 6
- [7] M. Oechsle, S. Peng, and A. Geiger, "UNISURF: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction," in *ICCV*, 2021. 1, 2, 3
- [8] A. Jain, M. Tancik, and P. Abbeel, "Putting nerf on a diet: Semantically consistent few-shot view synthesis," in *ICCV*, 2021. 1
- [9] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman, "Mip-nerf 360: Unbounded anti-aliased neural radiance fields," *arXiv*, 2021. 1
- [10] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. Efros, and T. Darrell, "CyCADA: Cycle-consistent adversarial domain adaptation," in *ICML*, 2018. 2
- [11] S. Manivasagam, S. Wang, K. Wong, W. Zeng, M. Sazanovich, S. Tan, B. Yang, W.-C. Ma, and R. Urtaun, "LiDARsim: Realistic lidar simulation by leveraging the real world," in *CVPR*, 2020. 2
- [12] Z. Yang, Y. Chai, D. Anguelov, Y. Zhou, P. Sun, D. Erhan, S. Rafferty, and H. Kretschmar, "Surfelgan: Synthesizing realistic sensor data for autonomous driving," in *CVPR*, 2020. 2
- [13] F. Engelmann, J. Stückler, and B. Leibe, "Samp: shape and motion priors for 4d vehicle reconstruction," in *WACV*. IEEE, 2017, pp. 400–408. 2, 5, 6
- [14] F. Lu, Z. Liu, X. Song, D. Zhou, W. Li, H. Miao, M. Liao, L. Zhang, B. Zhou, R. Yang, *et al.*, "PerMO: Perceiving more at once from a single image for autonomous driving," *arXiv*, 2020. 2
- [15] Z. Yang, S. Manivasagam, M. Liang, B. Yang, W.-C. Ma, and R. Urtaun, "Recovering and simulating pedestrians in the wild," in *CoRL*, 2020. 2
- [16] S. Duggal, Z. Wang, W.-C. Ma, S. Manivasagam, J. Liang, S. Wang, and R. Urtaun, "Secrets of 3d implicit object shape reconstruction in the wild," *arXiv*, 2021. 2
- [17] X. Zhao, Z. Zheng, C. Ji, Z. Liu, S. Lin, T. Yu, J. Suo, and Y. Liu, "Vehicle reconstruction and texture estimation using deep implicit semantic template mapping," *arXiv*, 2020. 2
- [18] Y. Chen, F. Rong, S. Duggal, S. Wang, X. Yan, S. Manivasagam, S. Xue, E. Yumer, and R. Urtaun, "Geosim: Realistic video simulation via geometry-aware composition for self-driving," in *CVPR*, 2021. 2, 6
- [19] S. W. Kim, J. Philion, A. Torralba, and S. Fidler, "DriveGAN: Towards a controllable high-quality neural simulation," in *CVPR*, 2021. 2
- [20] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "DeepSDF: Learning continuous signed distance functions for shape representation," in *CVPR*, 2019. 2, 3
- [21] Z. Chen and H. Zhang, "Learning implicit fields for generative shape modeling," in *CVPR*, 2019. 2
- [22] Z. Yang, S. Wang, S. Manivasagam, Z. Huang, W.-C. Ma, X. Yan, E. Yumer, and R. Urtaun, "S3: Neural shape, skeleton, and skinning fields for 3d human modeling," in *CVPR*, 2021, pp. 13 284–13 293. 2
- [23] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, "Occupancy networks: Learning 3d reconstruction in function space," in *CVPR*, 2019. 2
- [24] M. Atzmon and Y. Lipman, "SAL: Sign agnostic learning of shapes from raw data," in *CVPR*, 2020. 2
- [25] T. Davies, D. Nowrouzezahrai, and A. Jacobson, "On the effectiveness of weight-encoded neural implicit 3d shapes," *arXiv*, 2020. 2
- [26] V. Sitzmann, "Awesome implicit representations - a curated list of resources on implicit neural representations." 2
- [27] V. Sitzmann, M. Zollhöfer, and G. Wetzstein, "Scene representation networks: Continuous 3d-structure-aware neural scene representations," 2019. 2
- [28] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," *NeurIPS*, 2020. 2, 5
- [29] M. Niemeyer, L. Mescheder, M. Oechsle, and A. Geiger, "Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision," in *CVPR*, 2020. 2
- [30] L. Yariv, J. Gu, Y. Kasten, and Y. Lipman, "Volume rendering of neural implicit surfaces," in *NeurIPS*, 2021. 2, 3
- [31] D. Azinović, R. Martin-Brualla, D. B. Goldman, M. Nießner, and J. Thies, "Neural RGB-D surface reconstruction," *arXiv*, 2021. 2, 3
- [32] E. Insafutdinov, D. Campbell, J. F. Henriques, and A. Vedaldi, "Snes: Learning probably symmetric neural surfaces from incomplete data," in *ECCV*, 2022. 2
- [33] K. Deng, A. Liu, J.-Y. Zhu, and D. Ramanan, "Depth-supervised nerf: Fewer views and faster training for free," *arXiv*, 2021. 2
- [34] 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," *arXiv*, 2021. 2, 4, 6
- [35] P. P. Srinivasan, B. Deng, X. Zhang, M. Tancik, B. Mildenhall, and J. T. Barron, "NeRV: Neural reflectance and visibility fields for relighting and view synthesis," in *CVPR*, 2021. 2, 4
- [36] S. Bi, Z. Xu, P. Srinivasan, B. Mildenhall, K. Sunkavalli, M. Hašan, Y. Hold-Geoffroy, D. Kriegman, and R. Ramamoorthi, "Neural reflectance fields for appearance acquisition," *arXiv*, 2020. 2, 4
- [37] M. Boss, R. Braun, V. Jampani, J. T. Barron, C. Liu, and H. Lensch, "NeRD: Neural reflectance decomposition from image collections," in *ICCV*, 2021. 2, 4
- [38] K. Zhang, F. Luan, Q. Wang, K. Bala, and N. Snavely, "Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting," in *CVPR*, 2021. 2
- [39] B. T. Phong, "Illumination for computer generated pictures," *Communications of the ACM*, 1975. 2, 4
- [40] J. Arvo, "Applications of irradiance tensors to the simulation of non-lambertian phenomena," in *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, 1995. 2, 4
- [41] J. Ost, F. Mannan, N. Thuerey, J. Knodt, and F. Heide, "Neural scene graphs for dynamic scenes," in *CVPR*, 2021. 2, 4, 6
- [42] M. Niemeyer and A. Geiger, "GIRAFFE: Representing scenes as compositional generative neural feature fields," in *CVPR*, 2021. 2
- [43] S. Liu, X. Zhang, Z. Zhang, R. Zhang, J.-Y. Zhu, and B. Russell, "Editing conditional radiance fields," *arXiv*, 2021. 2
- [44] B. Yang, Y. Zhang, Y. Xu, Y. Li, H. Zhou, H. Bao, G. Zhang, and Z. Cui, "Learning object-compositional neural radiance field for editable scene rendering," in *ICCV*, 2021. 2
- [45] K. Schwarz, Y. Liao, M. Niemeyer, and A. Geiger, "GRaF: Generative radiance fields for 3d-aware image synthesis," 2020. 2
- [46] M. Guo, A. Fathi, J. Wu, and T. Funkhouser, "Object-centric neural scene rendering," 2021. 2
- [47] W. Chen, J. Litalien, J. Gao, Z. Wang, C. F. Tsang, S. Khamis, O. Litany, and S. Fidler, "DIB-R++: Learning to predict lighting and material with a hybrid differentiable renderer," *arXiv*, 2021. 2
- [48] J. Y. Zhang, G. Yang, S. Tulsiani, and D. Ramanan, "Ners: Neural reflectance surfaces for sparse-view 3d reconstruction in the wild," in *NeurIPS*, 2021. 2, 5
- [49] H. Barrow, J. Tenenbaum, A. Hanson, and E. Riseman, "Recovering intrinsic scene characteristics," *Comput. Vis. Syst.*, vol. 2, no. 3-26, p. 2, 1978. 2
- [50] J. T. Barron and J. Malik, "Intrinsic scene properties from a single RGB-D image," in *CVPR*, 2013. 2
- [51] W.-C. Ma, H. Chu, B. Zhou, R. Urtaun, and A. Torralba, "Single image intrinsic decomposition without a single intrinsic image," in *ECCV*, 2018. 2
- [52] 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," 2021. 3
- [53] D. S. Immel, M. F. Cohen, and D. P. Greenberg, "A radiosity method for non-diffuse environments," 1986. 3
- [54] J. T. Kajiya, "The rendering equation," in *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, 1986. 3
- [55] B. Walter, S. R. Marschner, H. Li, and K. E. Torrance, "Microfacet models for refraction through rough surfaces," *Rendering techniques*, vol. 2007, p. 18th, 2007. 4

- [56] A. Kirillov, Y. Wu, K. He, and R. Girshick, "Pointrend: Image segmentation as rendering," in *CVPR*, 2020. 5
- [57] S. Tulsiani, R. Tucker, and N. Snavely, "Layer-structured 3d scene inference via view synthesis," in *ECCV*, 2018, pp. 302–317. 5
- [58] J. Munkberg, J. Hasselgren, T. Shen, J. Gao, W. Chen, A. Evans, T. Müller, and S. Fidler, "Extracting triangular 3d models, materials, and lighting from images," in *CVPR*, 2022, pp. 8280–8290. 5, 6
- [59] P. Xiao, Z. Shao, S. Hao, Z. Zhang, X. Chai, J. Jiao, Z. Li, J. Wu, K. Sun, K. Jiang, *et al.*, "Pandaset: Advanced sensor suite dataset for autonomous driving," in *ITSC*, 2021. 5
- [60] 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. 5
- [61] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3d surface construction algorithm," *ACM siggraph computer graphics*, vol. 21, no. 4, pp. 163–169, 1987. 6
- [62] J. Fang, X. Zuo, D. Zhou, S. Jin, S. Wang, and L. Zhang, "Lidar-aug: A general rendering-based augmentation framework for 3d object detection," in *CVPR*, 2021. 6
- [63] J. Wang, A. Pun, J. Tu, S. Manivasagam, A. Sadat, S. Casas, M. Ren, and R. Urtasun, "Advsim: Generating safety-critical scenarios for self-driving vehicles," in *CVPR*, 2021. 6
- [64] Z. Wang, W. Chen, D. Acuna, J. Kautz, and S. Fidler, "Neural light field estimation for street scenes with differentiable virtual object insertion," in *ECCV*. Springer, 2022, pp. 380–397. 6
- [65] M. Tancik, V. Casser, X. Yan, S. Pradhan, B. Mildenhall, P. Srinivasan, J. T. Barron, and H. Kretzschmar, "Block-NeRF: Scalable large scene neural view synthesis," *arXiv*, 2022. 6
- [66] K. Rematas, A. Liu, P. P. Srinivasan, J. T. Barron, A. Tagliasacchi, T. Funkhouser, and V. Ferrari, "Urban radiance fields," *arXiv preprint arXiv:2111.14643*, 2021. 6
- [67] S. J. Garbin, M. Kowalski, M. Johnson, J. Shotton, and J. Valentin, "Fastnerf: High-fidelity neural rendering at 200fps," in *ICCV*, 2021, pp. 14 346–14 355. 6
- [68] T. Müller, A. Evans, C. Schied, and A. Keller, "Instant neural graphics primitives with a multiresolution hash encoding," *ACM Transactions on Graphics (ToG)*, vol. 41, no. 4, pp. 1–15, 2022. 6
- [69] C.-H. Lin, W.-C. Ma, A. Torralba, and S. Lucey, "Barf: Bundle-adjusting neural radiance fields," in *ICCV*, 2021, pp. 5741–5751. 6
- [70] L. Yen-Chen, P. Florence, J. T. Barron, A. Rodriguez, P. Isola, and T.-Y. Lin, "inerf: Inverting neural radiance fields for pose estimation," in *IROS*. IEEE, 2021, pp. 1323–1330. 6