

# Efficient Depth-Guided Urban View Synthesis

Sheng Miao<sup>1\*</sup>, Jiaxin Huang<sup>1\*</sup>, Dongfeng Bai<sup>2</sup>, Weichao Qiu<sup>2</sup>, Bingbing Liu<sup>2</sup>,  
Andreas Geiger<sup>3,4</sup>, and Yiyi Liao<sup>1✉</sup>

<sup>1</sup>Zhejiang University    <sup>2</sup>Huawei Noah’s Ark Lab    <sup>3</sup>University of Tübingen  
<sup>4</sup>Tübingen AI Center  
<https://xdimlab.github.io/EDUS/>

**Abstract.** Recent advances in implicit scene representation enable high-fidelity street view novel view synthesis. However, existing methods optimize a neural radiance field for each scene, relying heavily on dense training images and extensive computation resources. To mitigate this shortcoming, we introduce a new method called Efficient Depth-Guided Urban View Synthesis (EDUS) for fast feed-forward inference and efficient per-scene fine-tuning. Different from prior generalizable methods that infer geometry based on feature matching, EDUS leverages noisy predicted geometric priors as guidance to enable generalizable urban view synthesis from sparse input images. The geometric priors allow us to apply our generalizable model directly in the 3D space, gaining robustness across various sparsity levels. Through comprehensive experiments on the KITTI-360 and Waymo datasets, we demonstrate promising generalization abilities on novel street scenes. Moreover, our results indicate that EDUS achieves state-of-the-art performance in sparse view settings when combined with fast test-time optimization.

**Keywords:** urban view synthesis, generalizable NeRF, sparse view

## 1 Introduction


Novel View Synthesis (NVS) of street scenes is a key issue in autonomous driving and robotics. Recently, Neural Radiance Fields (NeRF) [25] have evolved as a popular scene representation for NVS. Several works [30, 43, 46, 50, 54] subsequently demonstrate promising NVS performance for street scenes.

However, these techniques are not applicable to sparse image settings, as they are typical for autonomous driving. When vehicles move at high speeds, a substantial portion of the content is captured from merely two or three view-points. This insufficient overlap between consecutive views significantly hinders the performance of existing methods. Furthermore, as the cameras are in forward motion, the parallax angle between frames is typically small. This reduction in parallax angle further amplifies the uncertainty of reconstruction.

Several recent methods address the sparse view setting: One line of works [6, 33, 41] exploits external geometric priors or regularization terms, such as depth

---

\* Equally contributed. ✉ Corresponding author.



**Fig. 1: Illustration.** Left: Most existing generalizable NeRF methods rely on feature matching for recovering the geometry, e.g., by constructing local cost volumes, potentially overfitting certain reference camera pose distributions. Right: Our method lifts geometric priors to the 3D space and fuses them into a global volume to be processed by a generalizable network. This enhances robustness as the geometric priors are unaffected by the reference image poses. We show synthesized images and depth maps through a feed-forward inference in the middle.

maps, geometric continuity, and color cues, to enhance model performance. However, these approaches require per-scene optimization, which may take up to a few hours to converge on a single scene.

Another line of research [4, 14, 44, 57], known as generalizable NeRF, resorts to pre-training on large-scale datasets to gain domain-specific prior knowledge. In new scenarios, these data-driven methods are able to efficiently infer novel views from sparse inputs in a feed-forward manner, with the option for additional test-time optimization. Most of these methods rely on feature matching to learn the geometry, which suffers from insufficient overlap and textureless regions given sparse street views shown in Fig. 1. Furthermore, these methods typically retrieve the nearest reference images for feature matching, potentially leading to overfitting to specific camera pose configurations. This hampers their generalization capability to unseen sparsity levels and various lengths of camera baselines, especially in highly sparse settings. We hence ask the following question: *Can we develop an efficient and generalizable urban view synthesis method capable of robustly handling diverse sparsity levels?*

To tackle this challenge, we propose Efficient Depth-Guided Urban View Synthesis (EDUS) which mitigates overfitting to a certain reference image pose configuration. This is achieved by lifting the reference images to 3D space utilizing approximate geometric priors, which are then processed by a feed-forward network directly applied in 3D space. In contrast to existing generalizable NeRF methods relying on feature matching for geometry prediction, whose performance is dependent on the relative poses between reference images and the target frame [4, 44], our method explicitly leverages geometric priors, e.g., monocular depth. These priors remain unaffected by the reference image poses, making our method more robust to various sparsity levels. In addition, our method allows for accumulating geometric priors of reference frames within a large window (60 meters) to learn a generalizable model in this global volume. This further enables efficient per-scene fine-tuning by directly updating the global volume,

circumventing the need to fine-tune numerous local volumes [4] or the entire network [44, 47].

In order to leverage the geometric prior within a limited distance while being able to perform feed-forward inference for the full unbounded scene, we further decompose the scene into three components: foreground regions, background, and the sky, with a separate generalizable module designed for each. Specifically, assuming a noisy point cloud of the foreground regions is provided by accumulating monocular/stereo depth estimations of the reference views, we train a generalizable 3D CNN to predict a feature volume. This 3D-to-3D generalizable module enhances robustness against various reference pose configurations as its input is the point cloud represented in world coordinates. While the 3D CNN leads to good geometry prediction, it encounters challenges in capturing high-frequency appearance details due to its inherent smoothness bias. To address this limitation, we augment the model with 2D features retrieved from nearby input views, enhancing the model’s ability to predict detailed appearance. Finally, we leverage image-based rendering for the background and sky.

Our contributions in this paper can be summarized as follows: 1) We investigate the utility of noisy geometric priors and 3D-to-3D refinement for generalizable NeRF models, achieving promising feed-forward inference on various sparse urban settings. Our global volume-based representation further allows for efficient per-scene optimization which converges within five minutes. 2) We build our generalizable model tailored for unbounded street scenes by utilizing three separate generalizable components for the foreground, background, and the sky. This allows us to leverage the geometric priors within limited distances while being able to represent the full unbounded scene. 3) Our proposed approach outperforms baselines for sparse novel view synthesis on KITTI-360 dataset, including geometric regularization-based methods and generalization-based methods. We further demonstrate that our model trained on KITTI-360 shows convincing generalization performance on the Waymo dataset.

## 2 Related Work

**Novel View Synthesis:** Neural Radiance Fields [25] revolutionized the task of Novel View Synthesis (NVS). Many follow-up works improve NeRF to achieve real-time rendering, fast training [7, 22, 35, 36], scene editing [11], etc. In this work, we focus on generalizable NeRF on street views which has many applications in autonomous driving.

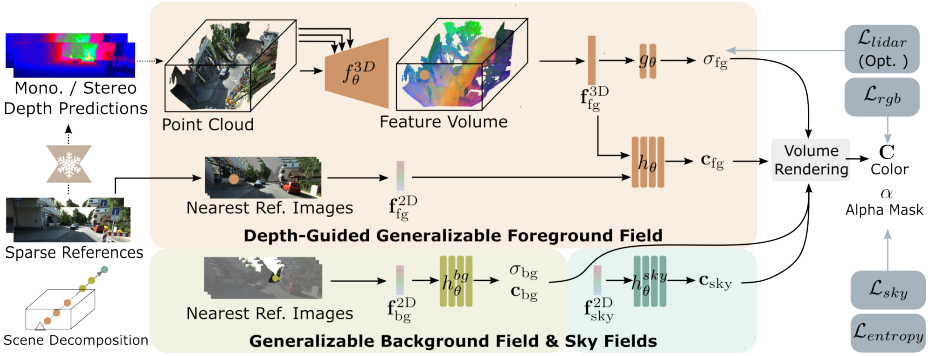
**Street Scene NeRF:** There are a few NeRF-based methods carefully designed to deal with street scenes [19, 27, 30, 46]. One line of methods [30, 46, 54] combines LiDAR and RGB data for street scene reconstruction. In this work, we optionally leverage LiDAR in the generalizable training stage and consider an RGB-only setting for test scenes. Another line of methods consider the street scenes based on RGB images [8, 9, 19, 27, 45]. All aforementioned methods takes relatively dense images as input and consider per-scene optimization. Our method, in contrast, focuses on generalizable street scenes NeRF which is applicable to various

sparsity levels. Methods for unbounded scene representation are also relevant for street scenes [2, 3, 24, 43]. Our method takes inspiration from these per-scene optimization methods and decomposes the scene into foreground, background, and sky with a generalizable module for each.

**Sparse View NeRF:** One line of works leverages depth or appearance regularization [6, 15, 26, 29, 33, 41, 51] to address the challenging task of sparse view synthesis. While leading to improvement in performance, these methods suffer from long training time. In contrast, our method takes the geometric priors as input to train a generalizable NeRF, enabling feed-forward inference and fast per-scene optimization. Generalizable NeRF approaches are also intensively studied in the literature [4, 16, 23, 44, 47, 48, 57]. Many works are built on image-based rendering [44, 57] or local cost volume [4, 16, 23]. As suggested in the current work MuRF [47], these methods struggle to generalize well to large baselines. MuRF tackles this issue by constructing a cost volume at the target view and achieves strong generalization to multiple baselines. However, it requires a long time for per-scene optimization, as the full computationally expensive model needs to be fine-tuned. Our method is more closely related to PointNeRF [48] and NeO 360 [14], both reconstructing a scene representation in a global world coordinate. PointNeRF leverages point clouds from a multi-view stereo algorithm and builds a point-based radiance field. We instead apply a 3D CNN to the input point cloud for refinement, allowing for better toleration of noisy point cloud input. NeO 360 [14] proposes to learn a tri-plane-based global scene representation from sparse RGB images, enabling 360° on the synthetic dataset. However, when it comes to complex street scenes with forward-moving cameras, NeO 360 struggles to directly predict the global scene representation. Our method tackles this challenge by incorporating explicit geometric priors.

**NVS with Point Cloud:** There are many existing NVS methods that take explicit geometric priors into account. [17, 48] learn a scene starting from a coarse point cloud, which can be obtained from multi-view stereo techniques [32] or learning-based methods [52, 55, 56] by unprojecting depth maps. These point-based works combine explicit geometric priors and implicit networks to regress color and other attributes of each point in the point cloud, and then render 2D images through volume rendering or rasterization. Other methods [20, 31] project the point cloud into feature images first and then learn a model to fill holes during neural rendering. Instead, we only take as input the point cloud, extracting geometry and appearance guidance for the generalizable reconstruction.

Some methods utilize encoders to process point clouds, followed by a volume-based neural rendering procedure. [13, 58] mainly focus on pre-trained point cloud encoder for downstream tasks, and do not consider cases like sparse street views into consideration. [49] reconstructs a driving scene well with either 2D images or 3D point cloud inputs, but it is limited by per-scene optimization. In contrast, our work is a generalizable method designed for autonomous driving scenes, which has a fast convergence speed and can obtain good reconstruction results even in extremely sparse viewpoints with fine-tuning.



**Fig. 2: EDUS.** Our model takes as input sparse reference images and renders a color image at a target view point. This is achieved by decomposing the scene into three generalizable modules: 1) depth-guided generalizable foreground fields to model the scene within a foreground volume; 2) generalizable background fields to model the background objects and stuff; and 3) generalizable sky fields. Our model is trained on multiple street scenes using RGB supervision  $\mathcal{L}_{rgb}$  and optionally LiDAR supervision  $\mathcal{L}_{lidar}$ . We further apply  $\mathcal{L}_{sky}$  to decompose the sky from the other regions and  $\mathcal{L}_{entropy}$  to penalize semi-transparent reconstructions.

### 3 Method

This paper explores the learning of a generalizable Neural Radiance Field (NeRF) for rendering novel views from sparse reference images in unbounded street scenes, with the option of test-time optimization. An overview of our EDUS is presented in Fig. 2. Initially, we partition the scene into foreground, background, and sky components. This partitioning facilitates the characterization of foreground areas with restricted extent utilizing a 3D-to-3D adaptable framework, while the remaining distant regions are modeled through image rendering techniques. Throughout this paper, the term **foreground** denotes the region confined within a predetermined volume, while **background** encompasses substances and objects outside of this volume, excluding the **sky**. We develop distinct generalizable modules for each component. The core of our method is the depth-guided generalizable foreground field, effectively combining noisy geometric priors and image-based rendering for representing close-range regions (Sec. 3.1). We further exploit image-based rendering to model the background regions and the sky (Sec. 3.2). Finally, these three parts are composed to represent the unbounded street scenes (Sec. 3.3). Our method is trained on multiple street scenes and enables feed-forward NVS on unseen validation scenes, with optional fine-tuning for further improvement (Sec. 3.4).

#### 3.1 Depth-Guided Generalizable Foreground Fields

**Depth Estimation:** We explore geometric priors of existing powerful depth estimation methods to enhance our generalizable NeRF for foreground regions,

i.e., regions within a close-range volume. Given  $N$  input images  $\{I_i\}_{i=0}^N$  for a single scene, we exploit off-the-shelf depth estimators to predict metric depth maps  $\{D_i\}_{i=0}^N$ . The depth estimator can be a stereo depth model [34] when the input contains stereo pairs. Otherwise, we adopt a metric monocular depth estimator [56] which provides less accurate but plausible metric depth estimations.

**Point Cloud Accumulation:** With the provided camera intrinsic matrices  $\{K_i\}_{i=0}^N$  and camera poses  $\{\mathbf{T}_i\}_{i=0}^N$ , we unproject the predicted depth maps  $\{D_i\}_{i=0}^N$  into 3D space, and accumulate them in the world coordinate system to form a scene point cloud  $\mathcal{P} \in \mathbb{R}^{N_p \times 3}$ . Specifically, for each pixel with homogeneous coordinate  $\mathbf{u}$  from a given frame, its world coordinate  $\mathbf{x}$  is computed as:  $\mathbf{x} = (d\mathbf{R}_i\mathbf{K}^{-1}\mathbf{u} + \mathbf{t}_i)$ , where  $d$  denotes the pixel’s estimated depth and  $(\mathbf{R}_i, \mathbf{t}_i)$  is derived from the  $i$ th frame’s cam-to-world transformation  $\mathbf{T}_i$ . Additionally, we assign a 3-channel color value to each unprojected point by retrieving the color from the RGB images. Note that we keep points that fall within the foreground volume and omit the rest, as depth predictions are usually untrustworthy at far regions.

We also utilize depth consistency check to filter noise. Specifically, we unproject a depth map  $\mathbf{D}_i$  of  $i$ th frame to 3D and reproject it to a nearby view  $j$ , obtaining a projected depth  $\mathbf{D}_{i \rightarrow j}$ . Next, we compare  $\mathbf{D}_{i \rightarrow j}$  and  $\mathbf{D}_j$ , and mask out pixels if the difference exceeds an empirical threshold  $\sigma = 0.2m$ .

**Modulation-based 3D Feature Extraction:** The resulting noisy foreground point cloud serves as input for our generalizable feature extraction network  $f_\theta^{3D}$ . We first discretize the point cloud  $\mathcal{P}$  into a volume  $\mathbf{P} \in \mathbb{R}^{H \times W \times D \times 3}$ , where  $H \times W \times D$  denotes the spatial resolution. The discretized point cloud is mapped to a feature volume  $\mathbf{F} \in \mathbb{R}^{H \times W \times D \times F}$ , where  $F$  denotes the feature dimension:

$$\mathbf{F} = f_\theta^{3D}(\mathbf{P}). \quad (1)$$

Notably, the design choice of  $f_\theta^{3D}$  affects the generalization performance. We experimentally observe that a conventional 3D U-Net with encoder-decoder architecture results in blurry artifacts when generalizes to new scenes, as shown in supplementary material. Instead, we draw inspiration from [28, 53] to employ a 3D Spatially-Adaptive Normalization convolutional neural network (SPADE CNN) to modulate feature volume generation at multiple scales, which is a popular strategy in semantic image synthesis. More specifically, our SPADE CNN discards the encoder part and comprises 3 SPADE residual blocks and upsampling layers. For each block, SPADE CNN downsamples the input volume  $\mathbf{P}$  to the corresponding spatial resolution and uses it to modulate layer activations with learned scale and bias parameters. We hypothesize that SPADE CNN is more effective in preserving the appearance information encoded in  $\mathbf{P}$  thanks to its multi-resolution modulation. Given a sampled point  $\mathbf{x} \in \mathbb{R}^3$  in the 3D space, we retrieve its feature  $\mathbf{f}_{\text{bg}}^{3D} \in \mathbb{R}^F$  from  $\mathbf{F}$  using trilinear interpolation. More network details are provided in the supplementary materials.

**Image-based 2D Feature Retrieval:** While the SPADE CNN takes as input the noisy depth predictions and achieves reasonable generalization, we observe that using  $\mathbf{f}_{\text{bg}}^{3D}$  alone struggles to provide high-frequency details in appearance.

This is due to two reasons: 1) the point cloud discretization limits the spatial resolution; 2) the 3D CNN suffers from an inductive smoothness bias. Therefore, we propose to additionally exploit the idea of image-based rendering to achieve better generalization of the appearance. Inspired by [44], we first select  $n$  nearest neighboring views from the target view based on the distances of the camera poses. This forms a group of nearest reference images  $\hat{\mathcal{I}} = \{I_k\}_{k=0}^K$ . For each sample point  $\mathbf{x}$ , we project it to the reference frame in  $\hat{\mathcal{I}}$ . Next, we retrieve the colors from the reference frames based on bilinear interpolation, and concatenate them to form a 2D feature vector  $\mathbf{f}_{\text{fg}}^{2D} \in \mathbb{R}^{3K}$ . We set  $K = 3$  in this paper.

**Color and Density Decoder:** Given a 3D point confined to the foreground volume, we predict its density  $\sigma_{\text{fg}}$  and color  $\mathbf{c}_{\text{fg}}$  based on the volumetric feature  $\mathbf{f}_{\text{fg}}^{3D}$  and the 2D feature  $\mathbf{f}_{\text{fg}}^{2D}$ . More specifically, the density  $\sigma_{\text{fg}}$  is predicted by applying a tiny MLP  $g_\theta$  to the volumetric feature  $\mathbf{f}_{\text{fg}}^{3D}$ .

For prediction of the color  $\mathbf{c}_{\text{fg}}$ , we concatenate  $\mathbf{f}_{\text{fg}}^{3D}$  and  $\mathbf{f}_{\text{fg}}^{2D}$  and feed them into a six-layer MLP. The 3D location  $\mathbf{x}$  and the viewing direction  $\mathbf{d}$  are further considered as input. This process can be expressed as:

$$\sigma_{\text{fg}} = g_\theta(\mathbf{f}_{\text{fg}}^{3D}), \quad \mathbf{c}_{\text{fg}} = h_\theta(\mathbf{f}_{\text{fg}}^{3D}, \mathbf{f}_{\text{fg}}^{2D}, \gamma(\mathbf{x}), \mathbf{d}) \quad (2)$$

Different from the original NeRF [25], we apply positional encoding  $\gamma(\cdot)$  to the location  $\mathbf{x}$  but not to the viewing direction  $\mathbf{d}$  as we expect non-Lambertian objects to be mainly modeled by the 2D feature  $\mathbf{f}_{\text{fg}}^{2D}$ . Please refer to the supplementary material for more details.

### 3.2 Generalizable Background and Sky Fields

As mentioned above, the foreground volume only covers a limited street scene and falls short in representing the distant landscape (often hundreds of meters away) and the sky. Therefore, we model the background region and the sky separately.

**Image-based Background Modelling:** An object placed outside the foreground volume is regarded as a background landscape. Due to perspective projection, distant objects occupy only a small portion of the image and have less detailed appearance. In this case, we observe that image-based rendering is sufficient to faithfully reconstruct the background regions as relative depth changes are small. Thus, we employ an MLP  $h_\theta^{bg}$  to jointly predict the background density  $\sigma_{\text{bg}}$  and the color  $\mathbf{c}_{\text{bg}}$  based on 2D image-based feature  $\mathbf{f}_{\text{bg}}^{2D}$ :

$$\sigma_{\text{bg}}, \mathbf{c}_{\text{bg}} = h_\theta^{bg}(\mathbf{f}_{\text{bg}}^{2D}, \gamma(\mathbf{x}), \mathbf{d}) \quad (3)$$

where  $\mathbf{x}$  is spatial location and  $\mathbf{d}$  is the view direction. Same to the foreground, we retrieve color from nearby images and concatenate them to form vector  $\mathbf{f}_{\text{bg}}^{2D}$ .

**Sky Modelling:** Street scenes invariably contain the infinite sky region where rays do not collide with any physical objects, leading to minimal changes in appearance while moving forward. Thus, we omit the positional influence and represent the sky as a view-dependent environment map. Given that the sky’s



**Fig. 3: Scene Decomposition Visual Results.**

appearance between consecutive frames is remarkably similar, we retrieve the 2D image feature  $\mathbf{f}_{\text{sky}}^{2\text{D}}$  from reference frames, as we discussed in Sec. 3.1, and then blend the sky color using a single-layer MLP  $h_{\theta}^{\text{sky}}$  conditioned on view directions:

$$\mathbf{c}_{\text{sky}} = h_{\theta}^{\text{sky}}(\mathbf{f}_{\text{sky}}^{2\text{D}}, \mathbf{d}). \quad (4)$$

### 3.3 Scene Decomposition

The composition of the foreground, background, and sky leads to our full model for the unbounded street scenes, as shown in Fig. 3. Here, we elaborate on the compositional volume rendering for each ray. To render a pixel, we cast a ray  $\mathbf{r} = \mathbf{r}_o + t\mathbf{d}$  from the camera center  $\mathbf{r}_o$  along its view direction  $\mathbf{d}$ . We sample points along the ray to query their color and density. The color and density of points falling within the foreground volume are queried through Eq. (2), otherwise the background model in Eq. (3) is applied.

We accumulate all foreground and background points to obtain the color and the accumulated alpha mask:

$$\mathbf{C}^{(\text{fg}+\text{bg})} = \sum_{i=1}^M T_k (1 - \exp(-\sigma_i)) \mathbf{c}_i, \quad \alpha^{(\text{fg}+\text{bg})} = \sum_{i=1}^M T_k (1 - \exp(-\sigma_i)) \quad (5)$$

where  $T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j\right)$  denotes the accumulated transmittance. Finally, we add the sky color of the ray assuming it is located at an infinite distance:

$$\mathbf{C} = \mathbf{C}^{(\text{fg}+\text{bg})} + (1 - \alpha^{(\text{fg}+\text{bg})})\mathbf{c}_{\text{sky}} \quad (6)$$

### 3.4 Training and Fine-tuning

we use multiple training losses to train on different scenes and (optional) test-time optimization on a specific scene. The most common sensors used in self-driving cars are cameras and LiDAR. Compared to cameras, LiDAR data are more expensive to acquire and may not be deployed in unseen test scenarios. As a result, we use LiDAR information only for training, to enhance geometric understanding of the model. This leads to the training loss function and fine-tuning loss functions as follows:

$$\mathcal{L}_{\text{training}} = \mathcal{L}_{\text{rgb}} + \lambda_1 \mathcal{L}_{\text{lidar}} + \lambda_2 \mathcal{L}_{\text{sky}} + \lambda_3 \mathcal{L}_{\text{entropy}} \quad (7)$$

$$\mathcal{L}_{\text{fine-tuning}} = \mathcal{L}_{\text{rgb}} + \lambda_2 \mathcal{L}_{\text{sky}} + \lambda_3 \mathcal{L}_{\text{entropy}} \quad (8)$$



We now elaborate on the loss terms.

**Photometric loss:** We use the  $\mathcal{L}_2$  loss to minimize the photometric loss between the rendered colors  $\mathbf{C}$  and ground truth pixel color  $\mathbf{C}$ :

$$\mathcal{L}_{\text{rgb}} = \|\mathbf{C} - \mathbf{C}_{gt}\|_2^2 \quad (9)$$

**Sky loss:** In order to separate the sky from the foreground and background solid structures, we leverage a pre-trained segmentation model [40] to provide pseudo sky mask  $M^{\text{sky}}$ , and supervise the rendered sky mask with binary cross entropy (BCE) loss:

$$\mathcal{L}_{\text{sky}} = \text{BCE}\left(1 - \alpha^{(\text{fg}+\text{bg})}, M^{\text{sky}}\right) \quad (10)$$

**LiDAR loss:** During training, we optionally exploit LiDAR as an extra sensor to improve the reconstruction of our method. Given a set of collected lidar samples  $\mathbf{r}(z) = \mathbf{o}_\ell + z\mathbf{d}_\ell$ , Urban Radiance Fields [30] proposes line-of-sight loss which encourages the ray distribution to approximate the Gaussian Distribution  $\mathcal{N}(0, \epsilon^2)$ , with  $\epsilon$  being the bound width of the solid structure. As the assumption of Gaussian distribution may not hold for all rays, we experimentally observe that the loss in [30] leads to unstable training in our setting. Hence, we propose to soft the strict constraints as follows:

$$\mathcal{L}_{\text{empty}} = \int_{t_n}^{z-\epsilon} w(t)^2 dt, \quad \mathcal{L}_{\text{near}} = 1 - \int_{z-\epsilon}^{z+\epsilon} w(t) dt, \quad \mathcal{L}_{\text{dist}} = \int_{z+\epsilon}^{t_f} w(t)^2 dt \quad (11)$$

where the bound width  $\epsilon$  is set to 0.5 meters as initialization and we apply exponential decay to the minimum 0.1 meters, which enables gradually decreasing ambiguities along the training step. Note that the second term  $\mathcal{L}_{\text{near}}$  encourages our model to increase volumetric density within a certain range but without specifying its distribution, whereas the first and third terms are required to keep the space empty for the rest of the regions.

**Entropy regularization loss:** To penalize our model representing the distant landscape as semi-transparent, we introduce the entropy regularization loss to encourage opaque rendering followed by StreetSurf [10]:

$$\mathcal{L}_{\text{entropy}} = -(\alpha^{\text{fg}} \ln \alpha^{\text{fg}} + (1 - \alpha^{\text{fg}}) \ln(1 - \alpha^{\text{fg}})) \quad (12)$$

## 4 Implementation details

**Input Volume Masking:** In the training stage, we randomly mask small regions of the input volume  $\mathbf{P}$ . Similar to Masked AutoEncoders [12], this trick improves the completion capability when given a sparse and incomplete point cloud. More details and ablation results are shown in the supplementary material.

**Hierarchical Sampling:** For rendering a ray, we follow NeuS [42] to perform hierarchical sampling during training and inference. This is achieved by distributing the initial half of these samples uniformly, followed by iterative importance sampling for the rest.

**Appearance Embedding:** The images collected from urban scenarios frequently present variable illumination or other environmental changes. Inspired by [24, 38], we add per-frame appearance embedding to account for exposure variations. For feed-forward inference on unseen scenes, we set the appearance embeddings as the mean embedding averaged across all training frames. Once fine-tuning is considered, we learn the appearance embedding of the reference images and interpolate them for the test views.

**Training Details:** We train our neural network end-to-end using Adam optimizer [18] with initial learning rate  $5 \times 10^{-3}$ . We set  $\lambda_1 = 0.1$ ,  $\lambda_2 = 1$  and  $\lambda_3 = 0.002$  as loss weights in our work. During training, we randomly select a single image and randomly sample 4096 pixels as a batch. Our EDUS is trained on one RTX4090 GPU with 500k steps, which takes about 2 days.

## 5 Experiment

**Dataset:** We conduct generalizable training on the KITTI-360 dataset [21]. Specifically, we gather 80 sequences encompassing varied street perspectives for training purposes, with each sequence comprising 60 stereo images alongside their respective camera poses. In the case of training scenes, we optionally utilize accumulated LiDAR points to supervise the geometry. Across all scenes, we standardize the pose by centering the translation around the origin and establish the axis-aligned bounding box (AABB) to decompose the street scene. The foreground range is set to  $[-12.6, 12.6 \text{ m}]$  for X axes,  $[-3, 9.8 \text{ m}]$  for Y axes,  $[-20, 31.2 \text{ m}]$  for Z axes, and dimension of the voxelized pointcloud is  $128 \times 64 \times 256$  since the voxel size is set to  $(0.2 \text{ m}, 0.2 \text{ m}, 0.2 \text{ m})$ .


We use five public validation sets from the KITTI-360 dataset and five scenes sourced from Waymo [37] as our test scenarios, with the Waymo dataset serving as an out-of-domain test set. It is important to note that the test sets from KITTI-360 have no overlap with the training scenes. Each test sequence comprises 60 images covering approximately 60 meters. To further assess our robustness against various levels of sparsity, we consider three sparsity settings: 50%, 80%, and 90% drop rates, where a higher drop rate indicates a more sparse set of reference images. To ensure fairness in comparison, we employ the same test images across all drop rates.

**Depth:** In this paper, we use both monocular and binocular depth to conduct our experiment. For test sequences in KITTI-360, depth used as both supervision and input is stereo depth generated by MobileStereoNet [34]. We only supervise the regularization-based methods by left-eye depth. For the geometry-guided methods, we use the left-eye depth values of corresponding training images and aggregate them into a point cloud after depth consistency check [5]. For experiments on Waymo, we employ Metric3D [56] to get metric monocular depth maps and unproject them to form point clouds for evaluation.

**Baselines:** We conduct a comparative analysis of our method against two lines of approaches. Firstly, we compare our model with generalizable approaches: IBRNet [44], MVSNeRF [4], NeO 360 [14] and MuRF [47]. We train each model

Methods	Setting	KITTI360 <sub>drop50%</sub>			KITTI360 <sub>drop80%</sub>			Waymo <sub>drop50%</sub>		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
IBR-Net		19.99	0.624	0.217	15.96	0.469	0.354	21.28	<b>0.777</b>	0.199
MVSNeRF	No	17.73	0.618	0.328	16.50	0.577	0.365	19.58	0.662	0.278
Neo360	per-scene	13.73	0.394	0.624	12.98	0.357	0.659	14.07	0.541	0.708
MuRF	opt.	<b>22.19</b>	0.741	0.264	18.69	0.639	0.353	23.12	0.779	0.318
<b>Ours</b>		21.93	<b>0.745</b>	<b>0.178</b>	<b>19.63</b>	<b>0.668</b>	<b>0.244</b>	<b>23.16</b>	0.761	<b>0.189</b>
IBR-Net		21.17	0.657	0.199	17.98	0.529	0.279	23.39	0.825	0.163
MVSNeRF	Per-scene	19.47	0.647	0.310	18.06	0.602	0.353	24.28	0.759	0.207
Neo360		17.92	0.489	0.566	17.51	0.445	0.581	22.59	0.670	0.522
MuRF	opt.	23.71	0.762	0.233	19.70	0.666	0.321	28.30	0.846	0.175
<b>Ours</b>		<b>24.43</b>	<b>0.793</b>	<b>0.136</b>	<b>20.91</b>	<b>0.712</b>	<b>0.220</b>	<b>28.45</b>	<b>0.834</b>	<b>0.132</b>

**Table 1: Quantitative Comparison** on five test scenes among generalizable methods. All models are trained on the KITTI-360 dataset using drop 50% sparsity level.



**Fig. 4: Qualitative Comparison** with generalizable baselines.

from scratch on the 80 training sequences of KITTI-360, and assess their performance both with and without per-scene fine-tuning. For all generalizable methods, we use three reference frames for training and inference. Secondly, we compare with test-time optimization methods, including methods for sparse view settings: MixNeRF [33], SparseNeRF [41], DS-NeRF [6], and one taking point clouds as input: 3DGS [17]. To maintain parity in our evaluation, we furnish the depth maps utilized in our method to train depth-supervision approaches [6, 41], and employ our accumulated point cloud as the initialization for 3DGS [17].

## 5.1 Comparison with Generalizable NeRFs

**Feed-Foward Inference:** We train all generalizable methods under 50% drop rate setting on KITTI-360 and assess their performance using both the 50% and 80% drop rate sparsity levels on KITTI-360. Tab. 1 and Fig. 4 show the qualitative and quantitative comparisons respectively. The results show that our proposed EDUS achieves photorealistic novel view synthesis on the complex street




Fig. 5: Zero-shot Inference on Waymo.




Fig. 6: Training Time.

view using only feed-forward inference. More importantly, our method outperforms the baselines with a larger gap when evaluated at a higher unseen sparsity level (drop 80%), as our method is less sensitive to the camera pose distribution. IBR-Net and MVSNeRF are mainly suitable for small-scale, object-level scenes and have difficulty handling unbounded street views, resulting in blurry images. NeO 360 directly predicts a tri-plane-based radiance field from RGB images, making it struggle to reconstruct the street scene of 60 meters driving distance. Our concurrent work MuRF [47] achieves comparable PSNR and SSIM in drop 50% but sometimes shows blocky artifacts, resulting in worse LPIPS. Furthermore, our methods outperform all baselines on the unseen Waymo datasets, demonstrating the generalization capability empowered by the geometric priors, see Fig. 5.

**Per-Scene Fine-Tuning:** We further fine-tune all generalizable methods under different sparsity levels, as shown in Tab. 1. For each novel street sequence, we freeze our 3D CNN and directly fine-tune the feature volume to enable efficient convergence. Note that our feature volume is initialized by running a single feed-forward pass of the 3D CNN, yielding a good initialization and subsequently yielding improved performance than baselines after fine-tuning.

Note that MuRF constructs local volume at each target view and requires complete propagation of gradients, resulting in slower convergence during per-scene optimization. In contrast, our global volume-based approach demonstrates faster convergence speed, see Fig. 6, making it more suitable for urban scenes. Besides, EDUS shows superior memory efficiency when handling full-resolution ( $376 \times 1408$ ) inference with 6GB while MuRF requires splitting the same size image into 4 patches, each consuming 16.2GB on a single RTX4090 (24GB).

## 5.2 Comparison with Test-Time Optimization NeRFs

We proceed to compare our method with test-time optimization approaches across various sparsity levels. Tab. 2 and Fig. 7 demonstrate that our method achieves state-of-the-art performance at high sparsity levels, attributable to the robust initialization provided by our generalizable model. In the drop 50% setting, the recent state-of-the-art method 3DGS [17] achieves the best SSIM and LPIPS, but their performance experiences significant degradation in the 80% and 90% settings. Our method achieves superior performance than other geometric regularization [33] or depth-supervision [6, 41] methods. Although they

Setting	Methods	Drop Rate	Cost Time	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Geometric Reg.	MixNeRF	50%	$\sim 51min$	21.50	0.625	0.357
		80%	$\sim 38min$	18.89	0.558	0.404
		90%	$\sim 31min$	17.89	0.529	0.416
	SparseNeRF	50%	$\sim 35min$	21.34	0.642	0.552
		80%	$\sim 27min$	19.18	0.607	0.559
		90%	$\sim 21min$	17.94	0.559	0.608
	DS-NeRF	50%	$\sim 36min$	20.40	0.604	0.601
		80%	$\sim 28min$	18.91	0.581	0.620
		90%	$\sim 22min$	17.91	0.564	0.632
Depth Input	3D GS	50%	$\sim 29min$	24.37	<b>0.804</b>	<b>0.109</b>
		80%	$\sim 18min$	19.80	0.679	0.234
		90%	$\sim 10min$	17.46	0.597	0.332
	Ours	50%	$\sim 5min$	<b>24.43</b>	0.793	0.136
		80%	$< 5min$	<b>20.91</b>	<b>0.712</b>	<b>0.220</b>
		90%	$< 5min$	<b>19.16</b>	<b>0.657</b>	<b>0.271</b>

Table 2: Test-Time Optimization Comparison.



Fig. 7: Qualitative Comparison with Test-Time Optimization baselines.

show good performance on small-scale datasets (e.g. DTU, LLFF datasets), they struggle with urban scenes which contain more complex geometry layouts. In a sparser setting, we maintain our advantages and achieve a PSNR of 20.91dB in 80% drop rate and 19.16dB PSNR in 90% drop rate, respectively. Another advantage of our work is that we achieve good quality with fast fine-tuning on novel scenes. Given the same computation recourse, our training time (**5 minutes**) outperforms other baselines by a large margin,  $5\times$  faster than DS-Nerf [6] and SparseNerf [41] and  $10\times$  faster than MixNeRF [33].

### 5.3 Ablation Study

In Tab. 6, we ablate the design choices of our method. We report the results on KITTI-360 validation dataset with 50% drop rate. First, we replace the modulated SPADE CNN with the traditional 3D U-Net, yielding the results (3D U-Net). This worsens feed-forward inference results, indicating that the modulated convolutions better preserve the information of the input point cloud.

Ablation	Feed-forward NVS			Finetuning NVS		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
3D U-Net	21.49	0.716	0.194	24.14	0.789	0.145
w/o $\mathbf{f}_{\text{fg}}^{3\text{D}}$	15.71	0.532	0.441	18.91	0.609	0.343
w/o $\mathbf{f}_{\text{fg}}^{2\text{D}}$	18.28	0.593	0.342	21.41	0.648	0.315
w/o decomposition	20.65	0.684	0.195	22.54	0.754	0.175
w/o $\mathcal{L}_{\text{lidar}}$	21.75	0.742	0.189	23.77	0.788	0.139
Full model	21.93	0.745	0.178	24.43	0.793	0.136

**Table 3: Ablation study.** Metrics are averaged over the 5 test scenes.

Next, we remove the 3D feature from the full model, yielding a baseline similar to IBRNet but augmented with scene decomposition (w/o  $\mathbf{f}_{\text{fg}}^{3\text{D}}$ ). This leads to a significant drop in performance, indicating the importance of the geometric priors. We also study the effect of the 2D image-based feature, and observe that removing it leads to blurry rendering (w/o  $\mathbf{f}_{\text{fg}}^{2\text{D}}$ ). This demonstrates that the 3D feature  $\mathbf{f}_{\text{fg}}^{3\text{D}}$  alone is insufficient due to the inherent inductive bias of the 3D CNN. When removing the street decomposition and representing the entire scene as foreground (w/o decomposition), it leads to artifacts in background regions due to the incorrect geometric modeling. Finally, we demonstrate that the LiDAR supervision of the training sequences leads to slight improvements (w/o  $\mathcal{L}_{\text{lidar}}$ ). We observe that the LiDAR supervision improves details, leading to a larger improvement in LPIPS compared to PSNR and SSIM. Note that LiDAR supervision is not applied during test, but only applied to the training sequences to enhance our generalizable geometry prediction against noisy geometric priors. Please refer to our supplementary material for qualitative comparisons.

## 6 Conclusion

The paper presents EDUS, a generalizable and efficient method for sparse urban view synthesis. By integrating geometric priors into a generalizable model, we demonstrate that our method achieves robust performance across various density levels and datasets. It further enables efficient fine-tuning and outperforms existing sparse view methods based on good initialization. We will investigate handling dynamic objects in street views in future work, which poses challenges to generalizable approaches.

## Acknowledgements

We would like to thank Dzmitry Tsishkou for his useful suggestions and tips. This work was supported by the National Key R&D Program of China under Grant 2021ZD0114500, and NSFC under grant 62202418 and U21B2004. A. Geiger was supported by the ERC Starting Grant LEGO-3D (850533) and the DFG EXC number 2064/1 - project number 390727645.

## References

1. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5855–5864 (2021)
2. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5470–5479 (2022)
3. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Zip-nerf: Anti-aliased grid-based neural radiance fields. arXiv preprint arXiv:2304.06706 (2023)
4. Chen, A., Xu, Z., Zhao, F., Zhang, X., Xiang, F., Yu, J., Su, H.: Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14124–14133 (2021)
5. Cheng, K., Long, X., Yin, W., Wang, J., Wu, Z., Ma, Y., Wang, K., Chen, X., Chen, X.: Uc-nerf: Neural radiance field for under-calibrated multi-view cameras. In: The Twelfth International Conference on Learning Representations (2023)
6. Deng, K., Liu, A., Zhu, J.Y., Ramanan, D.: Depth-supervised nerf: Fewer views and faster training for free. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12882–12891 (2022)
7. Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: CVPR (2022)
8. Fu, X., Zhang, S., Chen, T., Lu, Y., Zhou, X., Geiger, A., Liao, Y.: Panopticnerf-360: Panoramic 3d-to-2d label transfer in urban scenes. arXiv preprint arXiv:2309.10815 (2023)
9. Fu, X., Zhang, S., Chen, T., Lu, Y., Zhu, L., Zhou, X., Geiger, A., Liao, Y.: Panoptic nerf: 3d-to-2d label transfer for panoptic urban scene segmentation. In: 2022 International Conference on 3D Vision (3DV). pp. 1–11. IEEE (2022)
10. Guo, J., Deng, N., Li, X., Bai, Y., Shi, B., Wang, C., Ding, C., Wang, D., Li, Y.: Streetsurf: Extending multi-view implicit surface reconstruction to street views. arXiv preprint arXiv:2306.04988 (2023)
11. Haque, A., Tancik, M., Efros, A.A., Holynski, A., Kanazawa, A.: Instruct-nerf2nerf: Editing 3d scenes with instructions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 19740–19750 (2023)
12. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022)
13. Huang, D., Peng, S., He, T., Yang, H., Zhou, X., Ouyang, W.: Ponder: Point cloud pre-training via neural rendering. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16089–16098 (2023)
14. Irshad, M.Z., Zakharov, S., Liu, K., Guizilini, V., Kollar, T., Gaidon, A., Kira, Z., Ambrus, R.: Neo 360: Neural fields for sparse view synthesis of outdoor scenes. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9187–9198 (2023)
15. Jain, A., Tancik, M., Abbeel, P.: Putting nerf on a diet: Semantically consistent few-shot view synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5885–5894 (2021)

16. Johari, M.M., Lepoittevin, Y., Fleuret, F.: Geonerf: Generalizing nerf with geometry priors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18365–18375 (2022)
17. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4) (2023)
18. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
19. Kundu, A., Genova, K., Yin, X., Fathi, A., Pantofaru, C., Guibas, L.J., Tagliasacchi, A., Dellaert, F., Funkhouser, T.: Panoptic neural fields: A semantic object-aware neural scene representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12871–12881 (2022)
20. Li, Z., Li, L., Zhu, J.: Read: Large-scale neural scene rendering for autonomous driving. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 1522–1529 (2023)
21. Liao, Y., Xie, J., Geiger, A.: Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(3), 3292–3310 (2022)
22. Liu, L., Gu, J., Zaw Lin, K., Chua, T.S., Theobalt, C.: Neural sparse voxel fields. *Advances in Neural Information Processing Systems* **33**, 15651–15663 (2020)
23. Liu, Y., Peng, S., Liu, L., Wang, Q., Wang, P., Theobalt, C., Zhou, X., Wang, W.: Neural rays for occlusion-aware image-based rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7824–7833 (2022)
24. Martin-Brualla, R., Radwan, N., Sajjadi, M.S., Barron, J.T., Dosovitskiy, A., Duckworth, D.: Nerf in the wild: Neural radiance fields for unconstrained photo collections. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7210–7219 (2021)
25. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
26. Niemeyer, M., Barron, J.T., Mildenhall, B., Sajjadi, M.S., Geiger, A., Radwan, N.: Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5480–5490 (2022)
27. Ost, J., Mannan, F., Thuerey, N., Knodt, J., Heide, F.: Neural scene graphs for dynamic scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2856–2865 (2021)
28. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2337–2346 (2019)
29. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
30. Rematas, K., Liu, A., Srinivasan, P.P., Barron, J.T., Tagliasacchi, A., Funkhouser, T., Ferrari, V.: Urban radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12932–12942 (2022)
31. Rückert, D., Franke, L., Stamminger, M.: Adop: Approximate differentiable one-pixel point rendering. *ACM Transactions on Graphics (ToG)* **41**(4), 1–14 (2022)



32. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4104–4113 (2016)
33. Seo, S., Han, D., Chang, Y., Kwak, N.: Mixnerf: Modeling a ray with mixture density for novel view synthesis from sparse inputs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20659–20668 (2023)
34. Shamsafar, F., Woerz, S., Rahim, R., Zell, A.: Mobilestereonet: Towards lightweight deep networks for stereo matching. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2417–2426 (2022)
35. Sun, C., Sun, M., Chen, H.T.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction (2022)
36. Sun, C., Sun, M., Chen, H.T.: Improved direct voxel grid optimization for radiance fields reconstruction. arXiv preprint arXiv:2206.05085 (2022)
37. Sun, P., Kretzschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2446–2454 (2020)
38. Tancik, M., Casser, V., Yan, X., Pradhan, S., Mildenhall, B., Srinivasan, P.P., Barron, J.T., Kretzschmar, H.: Block-nerf: Scalable large scene neural view synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8248–8258 (2022)
39. Tancik, M., Weber, E., Ng, E., Li, R., Yi, B., Wang, T., Kristoffersen, A., Austin, J., Salahi, K., Ahuja, A., et al.: Nerfstudio: A modular framework for neural radiance field development. In: ACM SIGGRAPH 2023 Conference Proceedings. pp. 1–12 (2023)
40. Tao, A., Sapra, K., Catanzaro, B.: Hierarchical multi-scale attention for semantic segmentation. arXiv preprint arXiv:2005.10821 (2020)
41. Wang, G., Chen, Z., Loy, C.C., Liu, Z.: Sparsenerf: Distilling depth ranking for few-shot novel view synthesis. arXiv preprint arXiv:2303.16196 (2023)
42. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. arXiv preprint arXiv:2106.10689 (2021)
43. Wang, P., Liu, Y., Chen, Z., Liu, L., Liu, Z., Komura, T., Theobalt, C., Wang, W.: F2-nerf: Fast neural radiance field training with free camera trajectories. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4150–4159 (2023)
44. Wang, Q., Wang, Z., Genova, K., Srinivasan, P.P., Zhou, H., Barron, J.T., Martin-Brualla, R., Snavely, N., Funkhouser, T.: Ibrnet: Learning multi-view image-based rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4690–4699 (2021)
45. Wu, Z., Liu, T., Luo, L., Zhong, Z., Chen, J., Xiao, H., Hou, C., Lou, H., Chen, Y., Yang, R., et al.: Mars: An instance-aware, modular and realistic simulator for autonomous driving. In: CAAI International Conference on Artificial Intelligence. pp. 3–15. Springer (2023)
46. Xie, Z., Zhang, J., Li, W., Zhang, F., Zhang, L.: S-nerf: Neural radiance fields for street views. arXiv preprint arXiv:2303.00749 (2023)
47. Xu, H., Chen, A., Chen, Y., Sakaridis, C., Zhang, Y., Pollefeys, M., Geiger, A., Yu, F.: Murf: Multi-baseline radiance fields. arXiv preprint arXiv:2312.04565 (2023)
48. Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., Neumann, U.: Point-nerf: Point-based neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5438–5448 (2022)

49. Yang, H., Zhang, S., Huang, D., Wu, X., Zhu, H., He, T., Tang, S., Zhao, H., Qiu, Q., Lin, B., et al.: Unipad: A universal pre-training paradigm for autonomous driving. arXiv preprint arXiv:2310.08370 (2023)
50. Yang, J., Ivanovic, B., Litany, O., Weng, X., Kim, S.W., Li, B., Che, T., Xu, D., Fidler, S., Pavone, M., et al.: Emernerf: Emergent spatial-temporal scene decomposition via self-supervision. arXiv preprint arXiv:2311.02077 (2023)
51. Yang, J., Pavone, M., Wang, Y.: Freenerf: Improving few-shot neural rendering with free frequency regularization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8254–8263 (2023)
52. Yang, L., Kang, B., Huang, Z., Xu, X., Feng, J., Zhao, H.: Depth anything: Unleashing the power of large-scale unlabeled data. arXiv preprint arXiv:2401.10891 (2024)
53. Yang, Y., Yang, Y., Guo, H., Xiong, R., Wang, Y., Liao, Y.: Urbangiraffe: Representing urban scenes as compositional generative neural feature fields. arXiv preprint arXiv:2303.14167 (2023)
54. Yang, Z., Chen, Y., Wang, J., Manivasagam, S., Ma, W.C., Yang, A.J., Urtasun, R.: Unisim: A neural closed-loop sensor simulator. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1389–1399 (2023)
55. Yao, Y., Luo, Z., Li, S., Fang, T., Quan, L.: Mvsnet: Depth inference for unstructured multi-view stereo. In: Proceedings of the European conference on computer vision (ECCV). pp. 767–783 (2018)
56. Yin, W., Zhang, C., Chen, H., Cai, Z., Yu, G., Wang, K., Chen, X., Shen, C.: Metric3d: Towards zero-shot metric 3d prediction from a single image. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9043–9053 (2023)
57. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelnerf: Neural radiance fields from one or few images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4578–4587 (2021)
58. Zhu, H., Yang, H., Wu, X., Huang, D., Zhang, S., He, X., He, T., Zhao, H., Shen, C., Qiao, Y., et al.: Ponderv2: Pave the way for 3d foundataion model with a universal pre-training paradigm. arXiv preprint arXiv:2310.08586 (2023)


# Supplementary Material for Efficient Depth-Guided Urban View Synthesis

**Abstract.** In this **supplementary document**, we initially present an in-depth overview of our network architecture in Appendix A. Subsequently, we elaborate on our implementation details in Appendix B, encompassing the random masking strategy, the sampling strategy, and the point cloud accumulation configurations. Following this, we introduce the baselines for comparison in Appendix C, and discuss the datasets utilized in our experiments in Appendix D. Finally, we present additional experimental results in Appendix E. Besides, the **supplementary video** showcases our feed-forward and fine-tuned results across various scenes, along with novel view renderings compared against multiple baselines.

## A Network Architecture

### A.1 Modulation-based CNN

Our 3D Spatially-Adaptive Normalization convolutional neural network (SPADE CNN) modulates the feature volume generation at multiple resolutions. It takes a single voxelized point cloud with dimension  $64 \times 128 \times 256$  as input. The appearance information of the input point cloud is injected into the feature volume at each resolution. Fig. 8 presents the detailed network architecture.



**Fig. 8: Architecture of our 3D SPADE CNN.**

## A.2 Foreground Network Architecture

The feature volume generated by the 3D SPADE CNN provides both geometric priors and appearance information for foreground contents. Based on a feature vector  $\mathbf{f}_{\text{fg}}^{3\text{D}} \in \mathbb{R}^{16}$  extracted from this feature volume, our density decoder  $g_\theta$  predicts the density value. In order to capture high-frequency appearance during feedforward inference, the color decoder  $h_\theta$  combines the 3D feature  $\mathbf{f}_{\text{fg}}^{3\text{D}} \in \mathbb{R}^{16}$  and 2D features queried from nearest views  $\mathbf{f}_{\text{fg}}^{2\text{D}} \in \mathbb{R}^9$  for color prediction. In addition, the color decoder further takes the positional encoding of  $\mathbf{x}$  as input, setting the number of frequencies to 10 following [25]. Furthermore, we inject viewing direction  $\mathbf{d} \in \mathbb{R}^3$  and a per-frame appearance embedding  $\mathbf{w} \in \mathbb{R}^{32}$ . We present the detailed network architecture in Tab. 4.

Density Decoder			
Layer	in channels	out channels	description
LinearRelu <sub>0</sub>	16	64	in: $\mathbf{f}_{\text{fg}}^{3\text{D}}$
LinearRelu <sub>1</sub>	64	1	out: $\sigma_{\text{fg}}$
Color Decoder			
Layer	in channels	out channels	description
LinearRelu <sub>0</sub>	16+9+63	128	in: $\mathbf{f}_{\text{fg}}^{3\text{D}}, \mathbf{f}_{\text{fg}}^{2\text{D}}, \gamma(\mathbf{x})$
LinearRelu <sub><math>i=1,2</math></sub>	128	128	/
LinearRelu <sub>3</sub>	128+32+3	128	in: + $\mathbf{d}, \mathbf{w}$
LinearRelu <sub>4</sub>	128	64	/
LinearRelu <sub>5</sub>	64	3	out: $\mathbf{c}_{\text{fg}}$

**Table 4: Density and Color Decoder Architectures.** From top to bottom: two-layer MLP for density prediction and six-layer MLP for color prediction.

## A.3 Background and Sky Network Architectures

We also build the generalizable model for the background and sky region via the image-based rendering paradigm as we mentioned in the main paper. We now show the details of the background and sky network details in Tab. 5.

## B Implementation Details

### B.1 Input Volume Random Masking

Our target is to perform novel view synthesis from sparse urban images. When provided with sparse images, the accumulated point clouds are often incomplete and contain holes due to occlusions or insufficient overlap across the input frames. To equip EDUS with the capability to handle sparse and incomplete point clouds by filling in the missing parts, we randomly mask portions of the point cloud

<b>Background network architecture</b>			
<b>Layer</b>	<b>in channels</b>	<b>out channels</b>	<b>description</b>
LinearRelu <sub>0</sub>	63+9	128	in: $\mathbf{f}_{\text{fg}}^{2D}, \gamma(\mathbf{x})$
LinearRelu <sub><i>i=1,2,3</i></sub>	128	128	/
LinearRelu <sub>4</sub>	128	16	out: $\sigma_{\text{bg}}, \mathbf{f}_{\text{e}}$
LinearRelu <sub>5</sub>	15+3	64	/
LinearRelu <sub>6</sub>	64	64	/
LinearRelu <sub>7</sub>	64	3	out: $\mathbf{c}_{\text{bg}}$
<b>Sky network architecture</b>			
<b>Layer</b>	<b>in channels</b>	<b>out channels</b>	<b>description</b>
LinearRelu <sub>0</sub>	9+3	3	in: $\mathbf{f}_{\text{fg}}^{2D}, \mathbf{d}$ ; out: $\mathbf{c}_{\text{sky}}$

**Table 5: Background and Sky Network Architectures.** The  $\mathbf{f}_{\text{e}} \in \mathbb{R}^{15}$  represents the geometric embedding vector and serves as the input for background color prediction.

during the training stage. Specifically, we randomly select an  $8\text{m} \times 8\text{m} \times 12\text{m}$  cuboid in each iteration, corresponding to a volume of size  $[40 \times 40 \times 60]$ , and eliminate the point cloud within this area. Note that this is not applied in feed-forward inference nor per-scene fine-tuning. This simple yet effective strategy enhances the generalizable ability on novel scenes, as shown in our ablation study in Appendix E.2.

## B.2 Hierarchical Sampling

The total number of samples located at the ray is 160 with 128 sampled from the foreground and 32 from the background. Specifically, we first select 80 points along the ray within the foreground volume, distributing one-half of these samples uniformly, while the other half is allocated linearly based on disparity spacing, following Nerfacto [39]. Next, we iteratively perform importance sampling to sample valid regions containing solid contents based on the coarse density prediction following NeuS [42], with three iterations and 16 samples for each iteration. For the background, we additionally uniformly sample 32 points out of the foreground volume.

## B.3 Accumulated Point Cloud

We design our accumulated point cloud to encourage the model to learn completion. Considering that the sparsity level in real-world driving data is unknown in advance, we aim to train a *single* generalizable model and evaluate its performance at different sparsity levels. In our main experiments, all generalizable methods are trained at the sparsity level of the drop50 setting, i.e., 50% images are available for supervision. Instead of directly using the 50% reference images to create the accumulated point cloud, we create a more sparse point cloud for training to encourage the model to learn scene completion. Specifically, the accumulated point cloud we use for training is from every fifth of the stereo pairs.

Note that during the feed-forward inference or per-scene optimization, we always use the given reference images of different sparsity levels (50%, 80%, or 90%) to generate the corresponding accumulated point cloud. This ensures fair comparison to the baselines and further verifies our method’s robustness to the density of the input point cloud.

## C Baselines

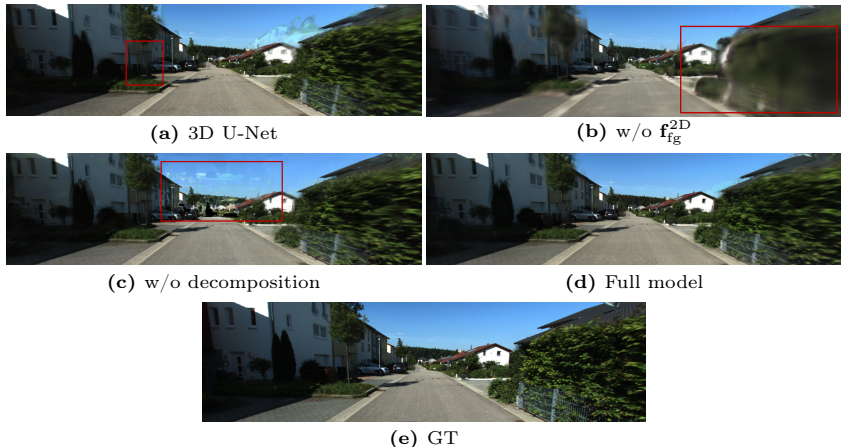
**Generalizable NeRFs** We utilize the official implementations of IBR-Net, MVSNerF, Neo360, and MuRF. For each method, we pretrain the model using 80 scenes from the KITTI360 dataset [21] under a 50% drop rate. In the case of IBR-Net, MVSNerF, and MuRF, we select the nearest three training frames of the target view as reference frames. As MVSNerF constructs local cost volumes defined at reference views, we adjust the fine-tuning strategy of MVSNerF to dynamically use the nearest reference frames to construct the cost volume, similar to the pre-training phase, rather than employing a fixed set of frames as in their original implementation.

**Test-Time Optimization Methods** The test-time optimization approaches are individually trained for each scene under three sparsity levels. We employ the official codebase of MixNeRF [33] and implement SparseNeRF [41] and DS-NeRF [6] by integrating stereo depth supervision onto Nerfstudio’s Mip-NeRF model [1, 39]. Additionally, we utilize the splatfacto model of Nerfstudio to train our 3D GS [17]. To ensure a fair comparison, we provide the depth maps used in our approach to train the depth-supervision methods [6, 41]. When it comes to point cloud-based methods [17], we find that the missing of sky may lead to a decline of performance, so we keep the sky in the point cloud and initialize the model with it.

## D Datasets

### D.1 KITTI-360

We follow the 50% drop rate (Drop50) and 90% drop rate (Drop90) settings of the KITTI-360 dataset and expand them into 80% drop rate (Drop80). To be more specific, we take a pair of left and right eye image pairs as the unit. In the Drop50 setting, we use odd-numbered image pairs as the training set. Similarly, in the Drop80 and Drop90 settings, we use every fifth or every tenth pair of images for training, respectively. Note that all test frames are the same to allow comparison across different sparsity levels. For example, considering ten consecutive image pairs, Drop50 refers to using pairs of id 0, 2, 4, 6, 8 for training, whereas Drop80 and Drop90 refer to using 0, 5 and 0, respectively. The test frames are selected as the left-eye image of 1, 3, 7, 9 pairs to avoid overlapping with the training frames.



**Fig. 9: Qualitative Results of Ablation Study.** The experiments are conducted on KITTI-360 under Drop50 sparsity level.

## D.2 Waymo

We use the same drop rate setting and point cloud accumulation scheme as the KITTI-360 dataset for Waymo. Note that we only use the front camera of the Waymo dataset. Therefore, every odd-numbered image is used for training for the Drop50 setting, and every fifth or tenth image is used for Drop80 and Drop90, respectively. The test views are configured to encompass frames at the 3rd and 7th positions, as well as frames with id equivalent to 3 or 7 plus the multiples of 10, which has no overlapping with the training set.

## E Additional Experimental Results

### E.1 Qualitative Results of Ablation Study

Fig. 9 demonstrates the qualitative outcomes of our ablation study. These results are all derived from feed-forward inference images under the Drop50 sparsity level, which align seamlessly with the corresponding quantitative results in Table 3 of the main paper. The substitution of the modulated SPADE CNN with a conventional 3D U-Net leads to less detailed prediction and yields blurriness in some regions. The elimination of the 2D image-based feature results in the high-frequency contents losing their sharpness and becoming noticeably blurry. Moreover, if we solely represent the scene as foreground, it results in artifacts in the distant region, due to the wrong modeling of the scene geometry. Our fully integrated model exhibits the highest quality of reconstruction, effectively validating our design choices.

Ablation	KITTI360 <sub>drop50</sub>			KITTI360 <sub>drop80</sub>		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
w/o random masking	20.62	0.694	0.228	18.54	0.641	0.288
Full model	21.93	0.745	0.178	19.63	0.668	0.244

**Table 6: Ablation Study of Input Volume Random Masking.** The results show that the random masking strategy in the training stage improves the feed-forward inference performance using different sparse settings.

## E.2 Additional Ablation of Input Volume Random Masking

As shown in Tab. 6, the random mask strategy helps improve our feed-forward performance, which means our model is equipped with filling capacity to process the incomplete point cloud by adding random masking.

## E.3 Additional Comparison on Waymo Dataset

We conduct zero-shot NVS under a drop rate of 50% using the model trained on the KITTI-360 dataset, and the rendered images are shown in Fig. 10. To further verify our method’s performance given different training scenes, we train all generalizable methods on the Waymo dataset and evaluate them on Waymo and KITTI-360, see Tab. 7. Specifically, all models are trained with the Drop50 setting on the Waymo dataset. Our model has better performance under 80% drop rate thanks to the depth-guided global volume and achieves the best metrics after fine-tuning. Note that our depth guidance of the Waymo dataset is obtained from a monocular depth estimation method [56], demonstrating that we are able to learn an effective generalizable model guided by noisy monocular depth priors. Fig. 11 shows the feed-forward and fine-tuned qualitative results of our method trained and evaluated on the Waymo dataset.

## E.4 Additional Qualitative Analysis of Geometry Prediction

Fig. 12 shows the novel view synthesis and depth comparison of our method to other optimization-based methods under Drop90 setting. The quantitative comparisons are reported in Table 2 of the main paper. Note that all methods except for Mix-NeRF use the same depth maps as supervision (DS-NeRF, SparseNeRF) or as input (3DGS, Ours). Although 3DGS [17] renders relatively good depth maps, it shows obvious flaws in appearance under sparse settings. Our model shows the best geometry thanks to our generalizable prior. As shown in Fig. 13, we illustrate the geometry extracted from the feed-forward inference, demonstrating that our EDUS is able to refine the noisy input point cloud and fill holes with generalizable 3D CNN.





**Fig. 10: Zero-shot Inference** on Waymo dataset. Our model is trained on the KITTI-360 dataset using Drop50 sparsity level.



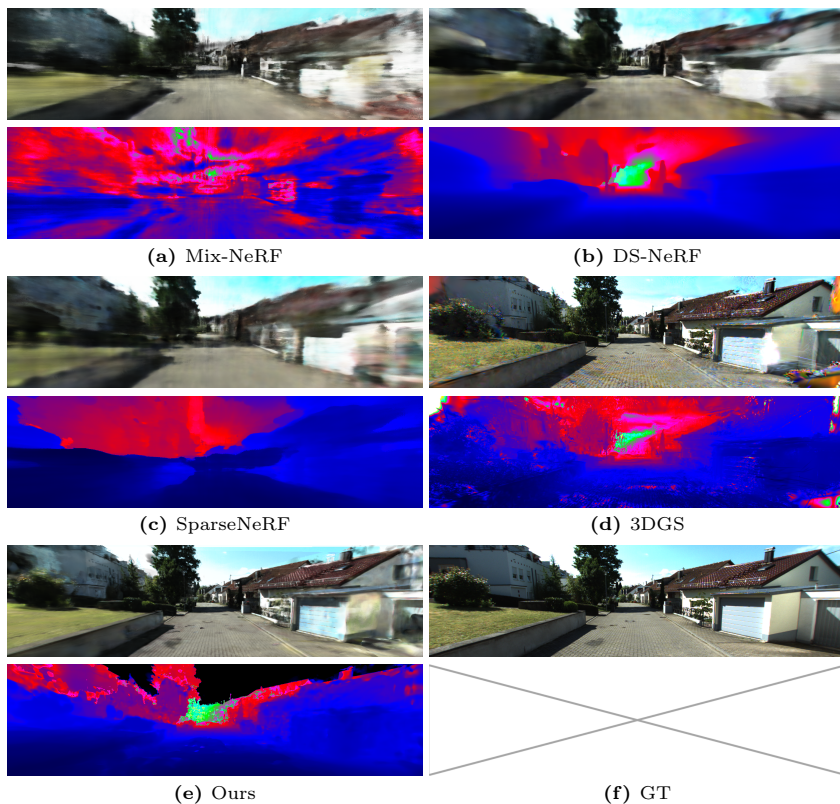
**Fig. 11: Qualitative Results** trained and evaluated on Waymo dataset. Our model is trained on Waymo dataset using Drop50 sparsity level. The feed-forward inference and fine-tuning are performed under the Drop80 setting.

## E.5 Number of Reference Views for Generalizable NeRFs


As we have mentioned in the main paper, our method takes input as a point cloud, which accumulates multi-frame information in addition to the reference frames. To verify fairness and our advancement compared with other reference-based methods, we evaluate the state-of-the-art approach MuRF [47] using different numbers of reference images. Specifically, we gradually increase the number of reference frames which is used for feature extraction and conduct feed-forward inference on our test sequence of the KITTI-360 dataset under Drop80 setting.

Methods	Setting	Waymo <sub>drop50</sub>			Waymo <sub>drop80</sub>			KITTI360 <sub>drop50</sub>		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
IBR-Net		22.90	0.739	0.148	17.90	0.619	0.242	18.88	0.606	0.272
MVSNeRF	No	18.05	0.595	0.378	17.41	0.581	0.404	15.35	0.524	0.389
Neo360	per-scene	16.28	0.505	0.569	15.83	0.505	0.584	10.99	0.204	0.687
MuRF	opt.	<b>23.66</b>	0.746	0.256	21.25	0.664	0.327	19.83	<b>0.669</b>	0.340
<b>Ours</b>		23.41	<b>0.769</b>	<b>0.147</b>	<b>21.64</b>	<b>0.703</b>	<b>0.204</b>	<b>20.13</b>	0.659	<b>0.257</b>
IBR-Net		23.71	<b>0.861</b>	0.121	18.58	0.741	0.222	20.78	0.650	0.205
MVSNeRF	Per-scene	22.42	0.699	0.309	20.81	0.657	0.348	18.80	0.611	0.354
Neo360	opt.	21.76	0.647	0.517	20.66	0.631	0.533	17.89	0.480	0.571
MuRF		24.06	0.753	0.224	22.85	0.711	0.260	21.24	0.710	0.265
<b>Ours</b>		<b>26.88</b>	0.839	<b>0.109</b>	<b>24.33</b>	<b>0.776</b>	<b>0.164</b>	<b>24.35</b>	<b>0.787</b>	<b>0.145</b>


**Table 7: Quantitative Comparison** on five test scenes among generalizable methods. All models are trained on the Waymo dataset using Drop50 sparsity level.



**Fig. 12: Qualitative Depth Comparison** with test-time optimization methods.



(a) Filling holes in input point clouds.



(b) Removing noise from input point clouds.

**Fig. 13: Qualitative Results** of density volume extracted from our feed-forward inference.


We find that increasing the number of reference frames does not improve the model’s performance. This may be due to the fact that further reference frames have stronger occlusions, hence increasing the difficulty of feature matching. On the other hand, the usage of more input frames also increases the burden on computing resources and causes out-of-memory when the number of views achieves 10. On the contrary, our method utilizes a fixed size of point cloud and manages to combine multi-frame guidance to predict accurate scene geometry. The results of MuRF using different numbers of reference views are shown in Tab. 8.

## E.6 Limitations

Although our approach achieves superior performance on NVS from sparse urban images, it still suffers degeneration when there is less overlap between the target image and reference images. In this case, 3D sample points will not query accurate appearance information from nearby images, leading to blurry rendering at the bottom of the test images, as illustrated in Fig. 14. This problem may be mitigated in future work by improving the 3D CNN’s capacity to directly predict high-frequency appearance instead of querying color from reference images.

Num. of Views	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
2	18.06	0.606	0.374
3	18.69	0.639	0.353
4	18.47	0.633	0.361
6	16.61	0.589	0.410
8	15.01	0.544	0.459
10	/	/	/

**Table 8: Number of Reference Views for MuRF.** We evaluate MuRF using various numbers of reference views. Note that the performance of MuRF does not increase wrt. the number of reference views.



**Fig. 14: Limitations in Drop90 Feed-Foward Inference.** Our performance degenerates in regions with insufficient overlap between reference and target frames, e.g., the bottom of the images.