

Improving Neural Indoor Surface Reconstruction with Mask-Guided Adaptive Consistency Constraints

Xinyi Yu¹, Liqin Lu¹, Jintao Rong¹, Guangkai Xu^{2,*} and Linlin Ou¹

Abstract—3D scene reconstruction from 2D images has been a long-standing task. Instead of estimating per-frame depth maps and fusing them in 3D, recent researches leverage the neural implicit surface as a global representation for 3D reconstruction. Equipped with data-driven pre-trained geometric cues, these methods have demonstrated promising performance. However, the inevitable inaccurate estimation of priors can lead to suboptimal reconstruction quality, particularly in some geometrically complex regions. In this paper, we propose a two-stage training process to further improve the reconstruction quality. It decouples the view-dependent and view-independent colors, and leverages two novel consistency constraints to enhance detail reconstruction performance without requiring extra priors. Additionally, we introduce an essential mask scheme to adaptively influence the selection of supervision constraints, thereby improving performance in a self-supervised paradigm. Experiments on synthetic and real-world datasets show the capability of reducing the side effects of inaccurately estimated priors and achieving high-quality scene reconstruction with rich geometric details.

I. INTRODUCTION

3D scene reconstruction from multiple images is a fundamental vision task with diverse applications, including robotics [1]–[3], virtual reality, augmented reality, etc. In robotics, reconstructions are widely used in trajectory planning [2] and mapping [4]. Given posed images, traditional algorithms usually estimate depth maps and lift them into 3D space, which can be categorized into multi-view stereo methods and monocular depth estimation methods. Multi-view stereo (MVS [5]–[7]) leverages accurate feature correspondences between keyframes to recover the 3D structure, and monocular depth estimation relies on large-scale training datasets to improve the generalization to diverse scenes. While feature matching lacks confidence in lighting changes, occlusion, and low-texture regions, robust monocular depth estimation also suffers from the unknown scale. They both struggle to handle the multi-frame inconsistency issues and are less satisfactory. Although some RGB-D fusion algorithms [8], [9] and post-processing optimization methods [10]–[12] are committed to ensuring consistency, it is found that they have difficulty handling some inaccurate depth predictions.

In order to tackle the consistency problem, there is an urgent need for a global 3D representation instead of per-frame

2D depth maps. Some learning-based methods [13], [14] project 2D features to spatial and directly predict TSDF value in 3D position. On the other hand, armed with the volume rendering theory, optimization-based methods usually encode a specific scene with a neural implicit scene representation by overfitting the pair of the input 3D position and the output color and geometry. However, unlike object-centric cases, the sparse views and texture-less areas of indoor scenes may lead to limited surface quality and local minima in optimization. To address the issue, some approaches integrate affine-invariant depth [15]–[19] and predicted normal priors [20], [21] as supervision. Although promising results have been achieved, they struggle to handle both the unknown scale-shift values of depth priors and inaccurate prior estimations, which results in poor reconstruction quality of complex geometric details.

In this study, we propose mask-guided adaptive consistency constraints to improve the detail reconstruction performance of neural surface representation. Similar to previous works [20]–[22], we also use two kinds of MLPs to predict signed distance and color information and employ the surface normal predicted by a pre-trained model [23] as priors. Specifically, we divide our training process into two stages. In the first stage, we focus on optimizing the RGB image reconstruction constraint as our primary objective while incorporating the estimated normal vector cues as additional supervision signals. This allows us to obtain an initial scene geometry. In the second stage, based on the principle that excellent reconstruction quality is closely associated with multi-view consistency, we categorize the reconstructed components into two groups: the accurate part and the inaccurate part, based on the difference in the rendered normal vectors at different viewpoints. During training, except for the color reconstruction constraint, to the sample rays passing through the inaccurate part, we employ a geometric consistency constraint to improve the accuracy of rendered depth; For the accurate part, we continue to apply normal priors to supervision. Besides, we decouple the rendering color into a view-dependent one and a view-independent one, and leverage another photometric consistency constraint to supervise the view-independent color. This two-stage training strategy adaptively distinguishes the accuracy of surface normal priors and adopts different supervision paradigms separately. Experimental results on both synthetic and real-world datasets demonstrate that our method achieves high-quality reconstructions with rich geometric details, outperforming other existing methods. Our main contributions are summarized as follows:

¹Xinyi Yu, Liqin Lu, Jintao Rong, and Linlin Ou are with College of Information Engineering, Zhejiang University of Technology, Hangzhou, China. {yuxy, liqinlu, 2111903071, linlinou}@zjut.edu.cn

²Guangkai Xu is with College of Computer Science and Technology, Zhejiang University, Hangzhou, China. guangkai.xu@zju.edu.cn

* Corresponding author.

- We propose a two-stage training process for neural surface reconstruction, which decouples the view-dependent and view-independent colors and leverages two novel consistency constraints to improve the quality.
- Based on the essential mask scheme, our model can reduce the side effects of inaccurate surface normal priors and enhance performance in a self-supervised paradigm during training.
- Experimental results on both synthetic and real-world datasets show that we can achieve high-quality scene reconstruction performance with rich geometric details.

II. RELATED WORK

A. Multi View Stereo

For years, 3D reconstruction from multi-view perspectives has remained a challenging yet significant task. Multi-view stereo (MVS) is a traditional reconstruction approach that leverages feature matching and triangulation methods to estimate 3D positions corresponding to pixels or features across multi-views. Using estimated positions and bundle adjustment, MVS recovered depth or normal to recover geometry [5], [7], [24], [25]. While MVS has achieved significant success, the estimated depth suffers from inaccurate and scale-inconsistent in abundant specular reflection or untextured regions, such as indoor scenes. Although some RGB-D fusion [8], [9] and post-processing optimization methods [10]–[12] have handled scale consistency issues, they still fail in inaccurate estimations. With the advancements in deep learning, learning-based methods have witnessed significant development in recent years. These methods employ neural networks to estimate depth or TSDF (truncated signed distance function) end-to-end. Depth-based methods estimate depth maps and utilize fusion procedures to reconstruct, which encounter challenges in noisy surfaces, and scale ambiguities. TSDF-based methods like NeuralRecon [14] propose a novel framework to lift the 2D features, fuse them spatially and temporally in 3D space, and predict the TSDF volume directly. Those methods always produce overly smooth reconstructions.

B. Implicit Representation of Geometry

Recently, with the success of volume rendering theory, some methods leverage Multi-Layer Perceptrons (MLPs) to implicitly represent geometry. These approaches are supervised by RGB images, overfit the pair of the 3D coordinates, and the corresponding color and geometric properties, like volume density or occupancy [26]–[30]. One notable advancement in this area is the Neural Radiance Field (NeRF) technique, which has demonstrated remarkable results in both novel view synthesis [31], [32] and implicit surface reconstruction [22], [30], [33]. These approaches utilize volume rendering methods to supervise implicit scene representation through a 2D photometric loss. However, volume density representations often fall short in geometric details due to a lack of sufficient constraints. Some methods have attempted to improve the volume rendering framework, such as VolSDF [33] and NeuS [30], which have achieved better

surface reconstructions but still face challenges in reconstructing geometric details. Consequently, certain methods aim to integrate geometric cues acquired from sensors [34], [35] or predicted by models [20], [21], [36] to strengthen geometric constraints. MonoSDF [20] integrates estimated monocular geometric clues into a neural volume rendering framework to enhance the overall quality of the reconstruction. NeuRIS [21] adaptively utilizes predicted normal cues by patch matching between neighboring images. While these methods have achieved accurate reconstruction results, they are sensitive to the accuracy in priors, especially when applied to real scenes. Compared to these methods, our method utilizes normal vector cues more efficiently and performs better in real-world scenarios.

III. METHOD

Aiming at 3D indoor scene reconstruction from posed images $\{I_k\}_{k=0\dots M}$, we use a neural surface representation optimized through the supervision of rendered RGB images. To enhance the robustness, normal priors obtained from pre-trained models are adopted. However, directly supervising the rendered normal may disturb the training process if the normal priors are inaccurate. Therefore, We optimize it with color and normal constraints to get the initial shape in the first stage. Then, we introduce geometric and photometric constraints to further improve reconstruction quality. What’s more, all the training constraints except the color one are guided by our proposed mask scheme in stage two. The overall pipeline of the second stage is shown in Fig. 1.

Concretely, for each selected image I_k , we sample m rays $\{\mathbf{r}_j\}_{j=0\dots m}$ passing through it. Randomly generating a virtual ray \mathbf{r}_j^v to pair with the current sample ray \mathbf{r}_j . Using MLPs and volume rendering framework, we render color $\hat{\mathbf{C}}$, decomposed color $\hat{\mathbf{C}}^{vi}$, depth \hat{D} , and normal $\hat{\mathbf{N}}$ for both \mathbf{r}_j and \mathbf{r}_j^v . To train our model, we minimize the color and normal differences, between $\hat{\mathbf{C}}(\mathbf{r}_j)$ and the given color $\mathbf{C}(\mathbf{r}_j)$, $\hat{\mathbf{N}}(\mathbf{r}_j)$ and the normal prior $\bar{\mathbf{N}}(\mathbf{r}_j)$ predicted by pre-trained models [23], respectively. Furthermore, we introduce a multi-view consistency constraint to guide the training process. During optimization, we use a masking scheme to adaptively select different training constraints for different sample rays.

A. Neural Surface Representation

As shown in Fig. 2, the neural surface representation is composed of two kinds of MLPs: a geometric geometry network f_g and two color networks f_c . The geometry network maps a 3D coordinate $\mathbf{x} \in \mathbb{R}^3$ to a feature vector \mathbf{z} and a signed distance function (SDF) value $\hat{s} \in \mathbb{R}$, which indicates the shortest distance to the closest geometry surface, and its sign shows whether the point is inside or outside the object.

$$[\hat{s}, \mathbf{z}] = f_g(\mathbf{x}; \theta_g) \quad (1)$$

where θ_g is the trainable parameters. The surface S is defined as the set of \mathbf{x} where the corresponding \hat{s} is equal to 0.

The color networks f_c generate view-dependent color $\hat{\mathbf{c}}^{vd} \in \mathbb{R}^3$ and view-independent color $\hat{\mathbf{c}}^{vi} \in \mathbb{R}^3$ from the

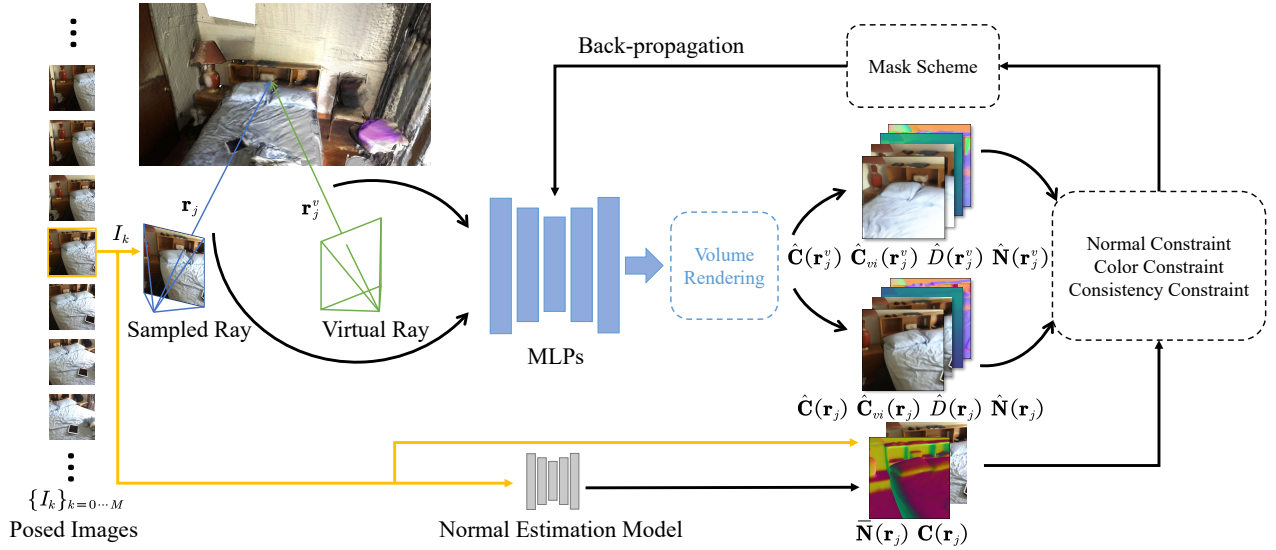


Fig. 1. We sample rays pass through selected image I_k . Randomly generating virtual ray \mathbf{r}_j^v corresponding to each sample ray \mathbf{r}_j . Rendering the color $\hat{\mathbf{C}}$, view-independent color $\hat{\mathbf{C}}_{vi}$, depth \hat{D} and normal $\hat{\mathbf{N}}$ along these rays. Using a pre-trained model to estimate normal priors $\tilde{\mathbf{N}}$ of \mathbf{r}_j . To learn MLPs' weights, we minimize the difference between $\hat{\mathbf{C}}(\mathbf{r}_j)$ and the given color $\mathbf{C}(\mathbf{r}_j)$. Besides, we utilize the mask-guided consistency and normal constraints.

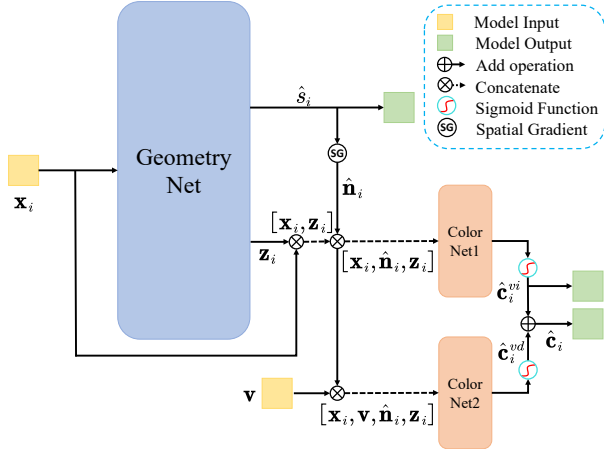


Fig. 2. Network architecture of our method.

input feature vectors. The final color $\hat{\mathbf{c}} \in \mathbb{R}^3$ is obtained by adding $\hat{\mathbf{c}}^{vd}$ with $\hat{\mathbf{c}}^{vi}$. This setting helps us to decompose color which we will discuss in Section III-B.3.

To recover the surface under the supervision of input RGB images, we render the colors of sample rays. Taking render in the ray \mathbf{r} as an example, we feed n sampling points $\mathbf{x}_i = \mathbf{o} + t_i \mathbf{v}$ along the ray into the geometric network f_g to obtain the corresponding SDF values \hat{s}_i and geometric feature \mathbf{z}_i . Here, $\mathbf{o} \in \mathbb{R}^3$ and $\mathbf{v} \in \mathbb{R}^3$ represents the camera position and ray direction with $|\mathbf{v}| = 1$ and $t_i \geq 0$, respectively. Subsequently, we concatenate \mathbf{x}_i , \mathbf{v} , $\hat{\mathbf{n}}_i$ and \mathbf{z}_i into two feature vectors and obtain the corresponding view-dependent color $\hat{\mathbf{c}}_i^{vd} = f_c(\mathbf{x}_i, \hat{\mathbf{n}}_i, \mathbf{v}, \mathbf{z}_i; \theta_c)$ and view-independent color $\hat{\mathbf{c}}_i^{vi} = f_c(\mathbf{x}_i, \hat{\mathbf{n}}_i, \mathbf{z}_i; \theta_c)$. The normal vector $\hat{\mathbf{n}}_i \in \mathbb{R}^3$ is the analytical gradient of the corresponding SDF value \hat{s}_i , and the final color $\hat{\mathbf{c}}_i$ is obtained by adding two color components together:

$$[\hat{s}_i, \mathbf{z}_i] = f_g(\mathbf{x}_i), \quad \hat{\mathbf{n}}_i = \frac{\partial \hat{s}_i}{\partial \mathbf{x}_i} \quad (2)$$

$$\hat{\mathbf{c}}_i = \hat{\mathbf{c}}_i^{vd} + \hat{\mathbf{c}}_i^{vi} \quad (3)$$

Following the volume rendering framework [22], [26], the color $\mathbf{C}(\mathbf{r})$ is accumulated along the ray.

$$\hat{\mathbf{C}}(\mathbf{r}) = \sum_{i=1}^n T_i \alpha_i \hat{\mathbf{c}}_i \quad (4)$$

where $T_i = \prod_{j=1}^{i-1} (1 - \alpha_j)$ and $\alpha_i = 1 - \exp(-\sigma_i \delta_i)$ denote the transmittance and alpha value, respectively. δ_i is the distance between neighbouring points, and σ_i is the density value corresponding to \mathbf{x}_i . To improve the geometric representation, we compute density values σ_i from \hat{s}_i [33]:

$$\sigma_i(\hat{s}_i) = \begin{cases} \frac{1}{2\beta} \exp(\frac{-\hat{s}_i}{\beta}), & \text{if } \hat{s}_i > 0. \\ \frac{1}{\beta} (1 - \frac{1}{2} \exp(\frac{\hat{s}_i}{\beta})), & \text{if } \hat{s}_i \leq 0. \end{cases} \quad (5)$$

where β is trainable. As β approach 0, the sensitivity of $\sigma_i(\hat{s}_i)$ to \hat{s}_i increases, contributing to edge reconstruction.

B. Supervision Constraints

1) *Color and Normal Constraints*: Since we have obtained each sample rays' rendering color $\hat{\mathbf{C}}(\mathbf{r})$, the parameters of f_g and f_c can be updated by minimizing the difference between $\hat{\mathbf{C}}(\mathbf{r})$ and the given color $\mathbf{C}(\mathbf{r})$:

$$\mathcal{L}_{rgb} = \sum_{\mathbf{r} \in \mathcal{R}} \left\| \hat{\mathbf{C}}(\mathbf{r}) - \mathbf{C}(\mathbf{r}) \right\|_1 \quad (6)$$

where \mathcal{R} represents the sample rays in a batch.

Geometric properties like normal vectors $\hat{\mathbf{N}}(\mathbf{r})$ and depth $\hat{D}(\mathbf{r})$ can be rendered by accumulating sample points' features along the ray, similar to rendering colors. We utilize normal constraints to guide the training process:

$$\hat{D}(\mathbf{r}) = \sum_{i=1}^n T_i \alpha_i t_i, \quad \hat{\mathbf{N}}(\mathbf{r}) = \sum_{i=1}^n T_i \alpha_i \hat{\mathbf{n}}_i \quad (7)$$

$$\mathcal{L}_{normal} = \frac{1}{|\mathcal{M}_r|} \sum_{\mathbf{r} \in \mathcal{M}_r} \left\| \hat{\mathbf{N}}(\mathbf{r}) - \bar{\mathbf{N}}(\mathbf{r}) \right\|_1 + \left\| 1 - \hat{\mathbf{N}}(\mathbf{r})^T \bar{\mathbf{N}}(\mathbf{r}) \right\|_1 \quad (8)$$

where \mathcal{M}_r denotes the ray mask. We will describe the details of \mathcal{M}_r in Section III-C. Besides, we also utilize the Eikonal term to regularize each sample point's SDF value in space [20], [21].

2) *Geometric Consistency Constraint*: The proposed geometric consistency is based on the principle that geometric properties of the surface, such as depth or normal, should be consistent among different viewpoints in unobstructed regions. We utilize these consistencies, visualized in Fig. 3, to constrain the optimization process.

Specifically, for each ray \mathbf{r} passing through the sample pixel, we calculate the corresponding rendered depth $\hat{D}(\mathbf{r})$ and normal $\hat{\mathbf{N}}(\mathbf{r})$ (Eq. 7). Using render depth $\hat{D}(\mathbf{r})$, we compute the target point \mathbf{x}_t (Eq. 9), which serves a similar purpose as the feature point in MVS but in 3D form. It is worth noting that this setting helps to avoid inaccuracy in feature extraction and matching. After that, we randomly generate a virtual viewpoint \mathbf{o}^v , which contributes to the robustness of the proposed constraint. Based on target point \mathbf{x}_t and \mathbf{o}^v , we calculate the virtual ray's direction \mathbf{v}^v . Consequently, we obtain a virtual ray \mathbf{r}^v originating from \mathbf{o}^v in direction \mathbf{v}^v , and virtual sample points $\mathbf{x}_i^v = \mathbf{o}^v + t_i^v \mathbf{v}^v$, where $t_i^v \geq 0$, positioned along this ray.

$$\mathbf{x}_t = \mathbf{o} + \hat{D}(\mathbf{r}) \mathbf{v}, \quad \mathbf{v}^v = \frac{\mathbf{x}_t - \mathbf{o}^v}{\|\mathbf{x}_t - \mathbf{o}^v\|_2} \quad (9)$$

Using the volume rendering framework, we render the depth $\bar{D}(\mathbf{r}_v)$ and normal $\hat{\mathbf{N}}(\mathbf{r}_v)$ of \mathbf{r}_v . Due to the geometric consistency between the depth of both rays, we propose a novel optimization target:

$$\mathcal{L}_{gc} = \frac{1}{2|\mathcal{M}_v|} \sum_{\mathbf{r}_v \in \mathcal{M}_v} |\hat{D}(\mathbf{r}_v) - \bar{D}(\mathbf{r}_v)|^2 \quad (10)$$

where $\bar{D}(\mathbf{r}_v) = \|\mathbf{x}_t - \mathbf{o}^v\|_2$, and \mathcal{M}_v denotes the mask for valid sample rays which fail in multi-view normal consistency. We will describe \mathcal{M}_v in Section III-C.

3) *Photometric Consistency Constraint*: Similar to the geometric consistency in views, the appearance of the scene also exhibits consistency. Due to the illumination changes or material properties, colors may appear different from various viewpoints. Inspired by [37], we decompose the render color of each sample point. Concretely, we leverage two color networks to predict view-dependent color $\hat{\mathbf{c}}_i^v$ and view-independent color $\hat{\mathbf{c}}_i^v$ of sample points, and the final rendering color $\hat{\mathbf{c}}_i$ is obtained by summing these two terms, as shown in Fig. 2. Cause the portion of view-dependent color is usually small in reality, we use the L1 loss to constrain $\hat{\mathbf{c}}_i^v$ and improve the color decomposition.

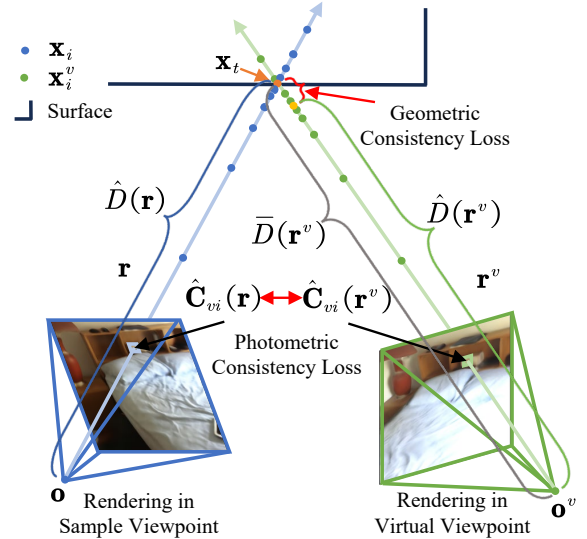


Fig. 3. Illustration of Consistency Constraints

The view-independent colors $\hat{\mathbf{C}}_{vi}$ of two kinds of rays are accumulated (Eq. 4). We propose an additional photometric consistency:

$$\mathcal{L}_{pc} = \frac{1}{|\mathcal{M}_r|} \sum_{\mathbf{r} \in \mathcal{M}_r} \left\| \hat{\mathbf{C}}_{vi}(\mathbf{r}) - \hat{\mathbf{C}}_{vi}(\mathbf{r}^v) \right\|_1 \quad (11)$$

C. Mask Scheme

In this section, we will introduce our mask scheme applied in the second training stage and utilize the AND operation to combine masks into \mathcal{M}_v and \mathcal{M}_r . It is worth noting that we present the valid rays as 1 in all of our masks.

1) *Sample Mask*: To enforce multi-view consistency, a virtual viewpoint \mathbf{o}^v is randomly generated for each sample ray. However, this setting may result in \mathbf{o}^v being positioned outside the scene or inside objects. To address this issue, we propose a sample mask \mathcal{M}_s to select valid virtual viewpoints. Specifically, we primarily utilize the SDF value in \mathbf{o}^v . Our reconstruction starts with a sphere that encloses all the given camera poses. As the training goes on, this sphere gradually approaches our target. Consequently, the outside part can be considered as the interior of the object, which means that if \mathbf{o}^v is valid, we will get a positive corresponding SDF value $\hat{s}(\mathbf{o}_v)$. Our sample mask is as follows:

$$\mathcal{M}_s = \begin{cases} 1, & \text{if } \hat{s}(\mathbf{o}_v) > 0. \\ 0, & \text{otherwise.} \end{cases} \quad (12)$$

2) *Occlusion Mask*: To address the problem of errors in multi-view consistency caused by occlusion along both rays, we propose an occlusion mask \mathcal{M}_o . Following the sampling algorithm in [33], our sample points are concentrated near the surfaces where the rays pass through. Hence, we can identify the presence of occlusion by analyzing the sign change in the SDF values associated with the sampling points along the ray.

$$\mathcal{M}_o^s = \begin{cases} 1, & \text{if } \|\text{diff}(\text{sgn}(\hat{\mathbf{s}}))\|_1 \leq 2. \\ 0, & \text{otherwise.} \end{cases}$$

$$\mathcal{M}_o^v = \begin{cases} 1, & \text{if } \|\text{diff}(\text{sgn}(\hat{\mathbf{s}}^v))\|_1 \leq 2. \\ 0, & \text{otherwise.} \end{cases}$$

$$\mathcal{M}_o = \mathcal{M}_o^s \& \mathcal{M}_o^v \quad (13)$$

where $\text{diff}(\cdot)$ computes the n -th forward difference along the given vector’s dimension, and $\text{sgn}(\cdot)$ is the sign function. $\hat{\mathbf{s}}$ and \mathcal{M}_o^s denote the vector of SDF values along the sample ray and the occlusion mask of this ray. Similarly, $\hat{\mathbf{s}}^v$ and \mathcal{M}_o^v represent the corresponding values for the virtual rays. Finally, the final occlusion mask \mathcal{M}_o is obtained through the AND operation between \mathcal{M}_o^s and \mathcal{M}_o^v .

3) *Adaptive Check Mask*: As described in Section III-B.2, high-quality reconstruction conforms to geometric consistency in multi-views. Therefore, we utilize the consistency of the rendered normal in multi-views as an adaptive check. Specifically, we use the normal Cosine Similarity (Eq. 14) to compute the difference between the sample ray’s render normal $\hat{\mathbf{N}}(\mathbf{r})$ and the virtual ray’s render normal $\hat{\mathbf{N}}(\mathbf{r}_v)$. We compare this difference to a certain threshold value ϵ . Rays with significant differences are identified by \mathcal{M}_a :

$$\cos(\hat{\mathbf{N}}(\mathbf{r}), \hat{\mathbf{N}}(\mathbf{r}_v)) = \frac{\hat{\mathbf{N}}(\mathbf{r}) \cdot \hat{\mathbf{N}}(\mathbf{r}_v)}{\|\hat{\mathbf{N}}(\mathbf{r})\|_2 \|\hat{\mathbf{N}}(\mathbf{r}_v)\|_2} \quad (14)$$

$$\mathcal{M}_a = \begin{cases} 1, & \text{if } \cos(\hat{\mathbf{N}}(\mathbf{r}), \hat{\mathbf{N}}(\mathbf{r}_v)) < \epsilon. \\ 0, & \text{otherwise.} \end{cases} \quad (15)$$

4) *Mask integration*: To better utilize the estimated normal cues and multi-view consistency, we organize the sample ray masks by AND operations:

$$\begin{aligned} \mathcal{M}_v &= \mathcal{M}_s \& \mathcal{M}_o \& \mathcal{M}_a \\ \mathcal{M}_r &= \mathcal{M}_s \& \mathcal{M}_o \& (1 - \mathcal{M}_a) \end{aligned} \quad (16)$$

Rays selected by \mathcal{M}_v have valid virtual viewpoints and no occlusion issues but fail to check in multi-view normal consistency. We put the geometric consistency constraint on these rays’ training process. \mathcal{M}_r chooses rays that conform to the normal consistency check. In those rays, predicted normal cues contribute to the reconstruction and we continue to apply them. In addition, we incorporate photometric consistency to further improve the quality.

IV. EXPERIMENTS

A. Implementation Detail

We implement our method with PyTorch and the network training is performed on one NVIDIA RTX 3090 GPU. The normal priors in our method are predicted by Omnidata model [23]. The network is trained over 25k iterations and each batch consists of 1024 sample rays. We first optimize the model directly guided with normal priors and RGB images over 25k iterations. In the second stage, in addition to color constraint, our model is trained under mask-guided normal and geometric consistency constraints until 75k iterations. After that, we add photometric consistency to our

optimization target. Following the optimization, we discrete our implicit function into voxel grids with a resolution of 512 and extract the mesh using Marching Cubes [38].

B. Experimental Settings

1) *Datasets*: Since our method primarily focuses on indoor scenes, we conduct evaluations on the Replica dataset [39] and the ScanNet dataset [40]. Replica consists of high-quality reconstructions of various indoor spaces. Each scene in Replica offers clean dense geometry and high-resolution images captured from multiple viewpoints. ScanNet is an RGB-D video dataset that comprises over 1500 indoor scenes with 2.5 million views. It is annotated with ground-truth camera poses, surface reconstructions, and instance-level semantic segmentations.

2) *Baselines*: We conduct a comparative analysis of our method with other methods. (1) COLMAP [5]: Traditional MVS reconstruction method, using screened Poisson Surface reconstruction (sPSR) to reconstruct mesh from point clouds. (2) NeuralRecon [14]: A learning-based TSDF fusion module. (3) MonoSDF(MLP version) [20]: Implicit method using predicted normal and depth priors directly. (4) NeuRIS [21]: Implicit method adaptive using normal priors.

3) *Evaluation Metrics*: To evaluate the quality of reconstruction, following [13], [20], [41], we mainly report *Accuracy*, *Completeness*, *Chamfer Distance*, *Precision*, *Recall* and *F-score* with the threshold of 0.05 meter. We further report *Normal Consistency* measure [20] to better evaluate reconstructions under the synthetic dataset.

TABLE I
QUANTITATIVE RESULTS ON THE SCANNET DATASET

Method	Acc.↓	Comp.↓	Chamfer- L_1 ↓	Prec.↑	Recall↑	F-score↑
COLMAP	0.047	0.235	0.141	0.711	0.441	0.537
NeuralRecon	0.044	0.123	0.084	0.741	0.502	0.595
NeuRIS	0.050	0.049	0.050	0.717	0.669	0.692
MonoSDF(MLP)	0.035	0.048	0.042	0.799	0.681	0.733
Ours	0.036	0.045	0.040	0.837	0.734	0.780

C. Results in Realistic Dataset

We conducted experiments using ScanNet dataset, which provides real-world data. We selected four scenarios and chose every 10th image from the original sets (about 2k-4k images). After the whole training process, we saved the reconstructed mesh and evaluated the final trained model. In Fig. 4, we compare our generated mesh with the original reconstructions from baselines. It shows that ours has more geometric detail and fixes missing caused by inaccurate estimation of normals. Additionally, we quantitatively evaluated our method compared to others in Table I. It can be seen that our method achieves more accurate results.

COLMAP [5] exhibits limitations in low-texture regions. NeuralRecon [14] relies on TSDF values for supervision and suffers from limitations in unseen scenarios. The performance of MonoSDF [20] is dependent on the quality of predicted geometric cues, which may lead to inaccuracy. NeuRIS [21] incorporates normal priors adaptively but is susceptible to noise in real-world datasets, resulting in artifacts in reconstructions.

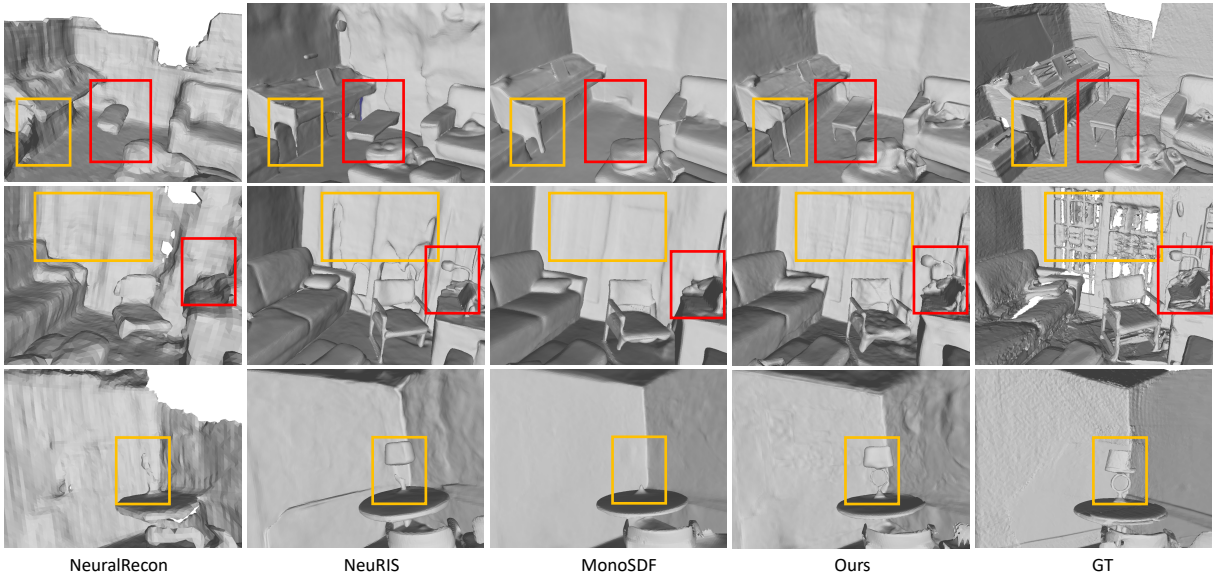


Fig. 4. Qualitative comparison of indoor 3D surface reconstruction on the ScanNet dataset.

D. Results in Synthetic Datasets

We conducted a quantitative evaluation of five scenarios from the Replica and averaged the results for comparison. It is worth noting that Replica is a synthetic dataset with little noise, which makes the predicted cues have minimal errors. To demonstrate the effectiveness of our approach in utilizing normal priors, we compare it with modified MonoSDF* that directly utilizes priors for supervision and outperforms other baselines. We conducted a comparison in Table II, which indicates ours achieves more accurate results, despite the majority of normal priors being accurate.

TABLE II

QUANTITATIVE COMPARISON ON THE REPLICA DATASET

Scan	Chamfer- L_1 ↓		F-score ↑		Normal. C ↑	
	monosdf*	Ours	monosdf*	Ours	monosdf*	Ours
scan1	4.104	2.593	76.945	89.589	88.133	91.619
scan2	2.415	2.996	92.316	89.045	93.945	94.304
scan3	6.292	3.718	85.153	85.843	93.851	93.376
scan4	3.119	2.562	86.367	92.031	91.290	92.329
scan5	4.627	3.900	84.976	85.604	92.244	93.833
mean	4.111	3.154	85.151	88.423	91.893	93.029

E. Ablation Studies

We conducted an ablation study to analyze the impact of consistency constraints and the adaptive mask in our method. To do this, we removed each setting from our framework and evaluated the results. We randomly selected three scenarios from the ScanNet and averaged the final evaluated metrics to perform the ablation experiment. The comprehensive results of the ablation study can be found in Table III.

TABLE III

QUANTITATIVE RESULTS OF ABLATION STUDIES

Norm. Mask	Geo. Pho.	Acc. ↓	Comp. ↓	Chamfer- L_1 ↓	Prec. ↑	Recall ↑	F-score ↑
✓		3.539	5.042	4.290	80.051	65.162	71.752
✓	✓	3.519	4.990	4.254	80.392	65.652	72.204
✓	✓	3.526	4.864	4.195	80.834	66.832	73.094
✓	✓	3.546	4.949	4.248	80.428	65.561	72.170
✓	✓	3.247	4.667	3.957	84.079	69.696	76.149

1) *Effectiveness of Geometric Consistency:* In regions that fail in normal consistency checks, we add geometric constraints to improve geometric detail reconstruction. In the

first two rows of Table III, geometric constraint improves quantitative results in terms of F-score and Chamfer- L_1 metrics, compared with directly utilizing normal priors. This demonstrates the effectiveness of the geometric consistency.

2) *Effectiveness of Photometric Consistency:* We utilize photometric consistency for further improvement in areas that conform to normal consistency. Cause the majority of normal priors are accurate, most optimization will be constrained by photometric consistency rather than geometric consistency. Therefore, this setting is more effective than using geometric consistency, as shown in Table III.

3) *Effectiveness of Adaptive Mask:* To enhance the effectiveness of constraints, we employ an adaptive mask for selection. As shown in the last two rows of Table III, applying all constraints without selection has limited improvement, due to the conflict between the inaccurate normal priors constraint and consistency constraints.

TABLE IV

THE EFFECTIVENESS OF COLOR DECOMPOSITION

	Acc. ↓	Comp. ↓	Chamfer- L_1 ↓	Prec. ↑	Recall ↑	F-score ↑
full(w/o CD.)	3.576	4.940	4.258	80.323	66.094	72.445
full	3.247	4.667	3.957	84.079	69.696	76.149

Table III shows that all of the settings contribute to enhancing reconstruction results, and lead to higher-quality reconstructions when employed in conjunction. Furthermore, to demonstrate the effectiveness of color decomposition, we conducted a comparative experiment on color. As shown in Table IV, color decomposition also contributes to reconstruction.

V. CONCLUSION

In this paper, we propose a 3D indoor scene reconstruction method using a two-stage training process. We decompose the color by perspective dependency and utilize two novel consistency constraints in both geometry and photography to improve the reconstruction quality. Besides, we introduce an essential mask to effectively select constraints in the training process. The experiments show that our approach achieves more geometric detail and high-quality reconstruction.

REFERENCES

- [1] S. Sajjan, M. Moore, M. Pan, G. Nagaraja, J. Lee, A. Zeng, and S. Song, "Clear grasp: 3d shape estimation of transparent objects for manipulation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 3634–3642.
- [2] D. Yang, T. Tosun, B. Eisner, V. Isler, and D. Lee, "Robotic grasping through combined image-based grasp proposal and 3d reconstruction," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6350–6356.
- [3] K. A. Skinner, E. Iscar, and M. Johnson-Roberson, "Automatic color correction for 3d reconstruction of underwater scenes," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 5140–5147.
- [4] W. Wang, B. Joshi, N. Burgdorfer, K. Batsosc, A. Q. Lid, P. Mordohaia, and I. Rekleitish, "Real-time dense 3d mapping of underwater environments," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5184–5191.
- [5] J. L. Schönberger, E. Zheng, J.-M. Frahm, and M. Pollefeys, "Pixel-wise view selection for unstructured multi-view stereo," in *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*. Springer, 2016, pp. 501–518.
- [6] J. L. Schönberger and J.-M. Frahm, "Structure-from-motion revisited," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4104–4113.
- [7] S. M. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski, "A comparison and evaluation of multi-view stereo reconstruction algorithms," in *2006 IEEE computer society conference on computer vision and pattern recognition (CVPR'06)*, vol. 1. IEEE, 2006, pp. 519–528.
- [8] T. Whelan, M. Kaess, H. Johannsson, M. Fallon, J. J. Leonard, and J. McDonald, "Real-time large-scale dense rgb-d slam with volumetric fusion," *The International Journal of Robotics Research*, vol. 34, no. 4–5, pp. 598–626, 2015.
- [9] H. Wang, M. Wang, Z. Che, Z. Xu, X. Qiao, M. Qi, F. Feng, and J. Tang, "Rgb-depth fusion gan for indoor depth completion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6209–6218.
- [10] G. Xu, W. Yin, H. Chen, C. Shen, K. Cheng, and F. Zhao, "Frozenrecon: Pose-free 3d scene reconstruction with frozen depth models," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023.
- [11] J.-W. Bian, H. Zhan, N. Wang, Z. Li, L. Zhang, C. Shen, M.-M. Cheng, and I. Reid, "Unsupervised scale-consistent depth learning from video," *International Journal of Computer Vision*, vol. 129, no. 9, pp. 2548–2564, 2021.
- [12] J.-W. Bian, H. Zhan, N. Wang, T.-J. Chin, C. Shen, and I. Reid, "Auto-rectify network for unsupervised indoor depth estimation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 12, pp. 9802–9813, 2021.
- [13] Z. Murez, T. van As, J. Bartolozzi, A. Sinha, V. Badrinarayanan, and A. Rabinovich, "Atlas: End-to-end 3d scene reconstruction from posed images," in *Proceedings of the European Conference on Computer Vision*, 2020, pp. 414–431.
- [14] J. Sun, Y. Xie, L. Chen, X. Zhou, and H. Bao, "Neuralrecon: Real-time coherent 3d reconstruction from monocular video," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 15 598–15 607.
- [15] G. Xu, W. Yin, H. Chen, C. Shen, K. Cheng, F. Wu, and F. Zhao, "Towards 3d scene reconstruction from locally scale-aligned monocular video depth," *arXiv preprint arXiv:2202.01470*, 2022.
- [16] W. Yin, J. Zhang, O. Wang, S. Niklaus, L. Mai, S. Chen, and C. Shen, "Learning to recover 3d scene shape from a single image," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 204–213.
- [17] R. Ranftl, K. Lasinger, D. Hafner, K. Schindler, and V. Koltun, "Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 3, pp. 1623–1637, 2022.
- [18] B. Roessle, J. T. Barron, B. Mildenhall, P. P. Srinivasan, and M. Nießner, "Dense depth priors for neural radiance fields from sparse input views," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 12 892–12 901.
- [19] T. Neff, P. Stadlbauer, M. Parger, A. Kurz, J. H. Mueller, C. R. A. Chaitanya, A. Kaplanyan, and M. Steinberger, "Donerf: Towards real-time rendering of compact neural radiance fields using depth oracle networks," in *Computer Graphics Forum*, vol. 40, no. 4. Wiley Online Library, 2021, pp. 45–59.
- [20] Z. Yu, S. Peng, M. Niemeyer, T. Sattler, and A. Geiger, "Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction," *Advances in neural information processing systems*, vol. 35, pp. 25 018–25 032, 2022.
- [21] J. Wang, P. Wang, X. Long, C. Theobalt, T. Komura, L. Liu, and W. Wang, "Neuris: Neural reconstruction of indoor scenes using normal priors," in *Proceedings of the European Conference on Computer Vision*. Springer, 2022, pp. 139–155.
- [22] L. Yariv, Y. Kasten, D. Moran, M. Galun, M. Atzmon, B. Ronen, and Y. Lipman, "Multiview neural surface reconstruction by disentangling geometry and appearance," *Advances in Neural Information Processing Systems*, vol. 33, pp. 2492–2502, 2020.
- [23] A. Eftekhar, A. Sax, J. Malik, and A. Zamir, "Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10 786–10 796.
- [24] Y. Furukawa, C. Hernández, et al., "Multi-view stereo: A tutorial," *Foundations and Trends® in Computer Graphics and Vision*, vol. 9, no. 1–2, pp. 1–148, 2015.
- [25] P. Lindenberger, P.-E. Sarlin, V. Larsson, and M. Pollefeys, "Pixel-perfect structure-from-motion with featuremetric refinement," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 5987–5997.
- [26] 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," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [27] Z. Zhu, S. Peng, V. Larsson, W. Xu, H. Bao, Z. Cui, M. R. Oswald, and M. Pollefeys, "Nice-slam: Neural implicit scalable encoding for slam," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 12 786–12 796.
- [28] A. Tewari, O. Fried, J. Thies, V. Sitzmann, S. Lombardi, K. Sunkavalli, R. Martin-Brualla, T. Simon, J. Saragih, M. Nießner, et al., "State of the art on neural rendering," in *Computer Graphics Forum*, vol. 39, no. 2. Wiley Online Library, 2020, pp. 701–727.
- [29] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, "Occupancy networks: Learning 3d reconstruction in function space," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4460–4470.
- [30] 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," *Advances in Neural Information Processing Systems*, vol. 34, pp. 27 171–27 183, 2021.
- [31] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman, "Mip-nerf 360: Unbounded anti-aliased neural radiance fields," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5470–5479.
- [32] M. M. Johari, Y. Lepoittevin, and F. Fleuret, "Geonerf: Generalizing nerf with geometry priors," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 18 365–18 375.
- [33] L. Yariv, J. Gu, Y. Kasten, and Y. Lipman, "Volume rendering of neural implicit surfaces," *Advances in Neural Information Processing Systems*, vol. 34, pp. 4805–4815, 2021.
- [34] D. Yan, X. Lyu, J. Shi, and Y. Lin, "Efficient implicit neural reconstruction using lidar," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 8407–8414.
- [35] D. Azinović, R. Martin-Brualla, D. B. Goldman, M. Nießner, and J. Thies, "Neural rgb-d surface reconstruction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6290–6301.
- [36] Y. Wang, Z. Li, Y. Jiang, K. Zhou, T. Cao, Y. Fu, and C. Xiao, "Neuralroom: Geometry-constrained neural implicit surfaces for indoor scene reconstruction," *ACM Transactions on Graphics (TOG)*, vol. 41, no. 6, pp. 1–15, 2022.
- [37] J. Tang, H. Zhou, X. Chen, T. Hu, E. Ding, J. Wang, and G. Zeng, "Delicate textured mesh recovery from nerf via adaptive surface refinement," *arXiv preprint arXiv:2303.02091*, 2022.
- [38] W. E. Lorenson and H. E. Cline, "Marching cubes: A high resolution

- 3d surface construction algorithm,” in *Seminal graphics: pioneering efforts that shaped the field*, 1998, pp. 347–353.
- [39] J. Straub, T. Whelan, L. Ma, Y. Chen, E. Wijmans, S. Green, J. J. Engel, R. Mur-Artal, C. Ren, S. Verma, *et al.*, “The replica dataset: A digital replica of indoor spaces,” *arXiv preprint arXiv:1906.05797*, 2019.
- [40] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, “Scannet: Richly-annotated 3d reconstructions of indoor scenes,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 5828–5839.
- [41] H. Guo, S. Peng, H. Lin, Q. Wang, G. Zhang, H. Bao, and X. Zhou, “Neural 3d scene reconstruction with the manhattan-world assumption,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5511–5520.