

DebSDF: Delving into the Details and Bias of Neural Indoor Scene Reconstruction

Yuting Xiao*, Jingwei Xu*, Zehao Yu, Shenghua Gao

Abstract—In recent years, the neural implicit surface has emerged as a powerful representation for multi-view surface reconstruction due to its simplicity and state-of-the-art performance. However, reconstructing smooth and detailed surfaces in indoor scenes from multi-view images presents unique challenges. Indoor scenes typically contain large texture-less regions, making the photometric loss unreliable for optimizing the implicit surface. Previous work utilizes monocular geometry priors to improve the reconstruction in indoor scenes. However, monocular priors often contain substantial errors in thin structure regions due to domain gaps and the inherent inconsistencies when derived independently from different views. This paper presents **DebSDF** to address these challenges, focusing on the utilization of uncertainty in monocular priors and the bias in SDF-based volume rendering. We propose an uncertainty modeling technique that associates larger uncertainties with larger errors in the monocular priors. High-uncertainty priors are then excluded from optimization to prevent bias. This uncertainty measure also informs an importance-guided ray sampling and adaptive smoothness regularization, enhancing the learning of fine structures. We further introduce a bias-aware signed distance function to density transformation that takes into account the curvature and the angle between the view direction and the SDF normals to reconstruct fine details better. Our approach has been validated through extensive experiments on several challenging datasets, demonstrating improved qualitative and quantitative results in reconstructing thin structures in indoor scenes, thereby outperforming previous work.

Index Terms—Multi-view Reconstruction, Implicit Representation, Indoor Scenes Reconstruction, Uncertainty Learning, Bias-aware SDF to Density Transformation.

1 INTRODUCTION

SURFACE reconstruction from multiple calibrated RGB images is a long-standing goal in computer vision and graphics with various applications ranging from robotics to virtual reality. Traditional methods tackle this problem with a multi-step pipeline. It first estimates dense depth maps for each RGB image by the multi-view stereo, fuses the depth maps to unstructured point clouds, and then converts it to a triangle mesh with Poisson surface reconstruction [3]. Recently, the neural implicit surface [2], [4], [5], [6], [7] has emerged as a powerful representation for multi-view surface reconstruction due to its simplicity. The key ideas are the utilization of coordinate-based networks for scene representation that map 3D coordinates to different scene properties such as signed distance and the use of differentiable volume rendering techniques that project the implicit surface to 2D observations such as RGBs, depths, and normals. The neural implicit surface can be optimized with multi-view images in an end-to-end manner and can be further converted to triangle mesh easily with the marching cubes algorithm at arbitrary resolution.

Despite the impressive reconstruction performance that has been shown for object-centric scenes with implicit surfaces, it is still challenging to reconstruct smooth and detailed surfaces from multi-view images for indoor scenes. First, multi-view reconstruction is an inherently under-

constrained problem as there exists an infinite number of plausible implicit surfaces that match the input images after rendering. Second, indoor scenes usually contain large texture-less regions and the photometric loss used for optimizing the implicit surface is unreliable. One possible way to circumvent these challenges is to leverage the priors about the indoor scenes. For example, by assuming a Manhattan world [8] where the wall and floor regions are orthogonal, Manhattan-SDF [9] achieves better performance than methods that directly optimize SDF from multi-view images. MonoSDF [1] and NeuRIS [10] further improve the reconstruction quality by incorporating large pre-trained models to infer geometric prior from a single view, e.g. monocular depth and normal maps.

While the overall structures such as walls and floors of indoor scenes can be faithfully reconstructed, these methods still struggle to recover fine details, e.g. chair legs and bracket of the desk, as shown in Figure 1. These failures result from several factors: 1) Monocular priors have large errors in these regions due to domain gaps between the training data used for the pre-trained model and the scenes we want to reconstruct. 2) As the monocular priors of each input image are predicted independently, they are unlikely to be multi-view consistent, especially in the thin and detailed structure regions. 3) Thin structure regions occupy a small area in the input images. Existing methods sample training rays uniformly across all training images, resulting in low sampling probabilities of these regions in comparison to walls and floors. 4) Smoothness regularization is applied uniformly to the whole space, which suppresses the learning of thin structure surfaces. These combined factors together make fine and detailed surfaces hard to form.

In this paper, we address these challenges of utilizing

- Yuting Xiao and Jingwei Xu contributed equally to this work;
- Corresponding Author: Shenghua Gao;
E-mail: gaoshh@shanghaitech.edu.cn
- Yuting Xiao, Jingwei Xu, and Shenghua Gao are with the School of Information Science and Technology, ShanghaiTech University, Shanghai 201210, China;
- Zehao Yu is with the Department of Computer Science, University of Tübingen, 72076, Germany;

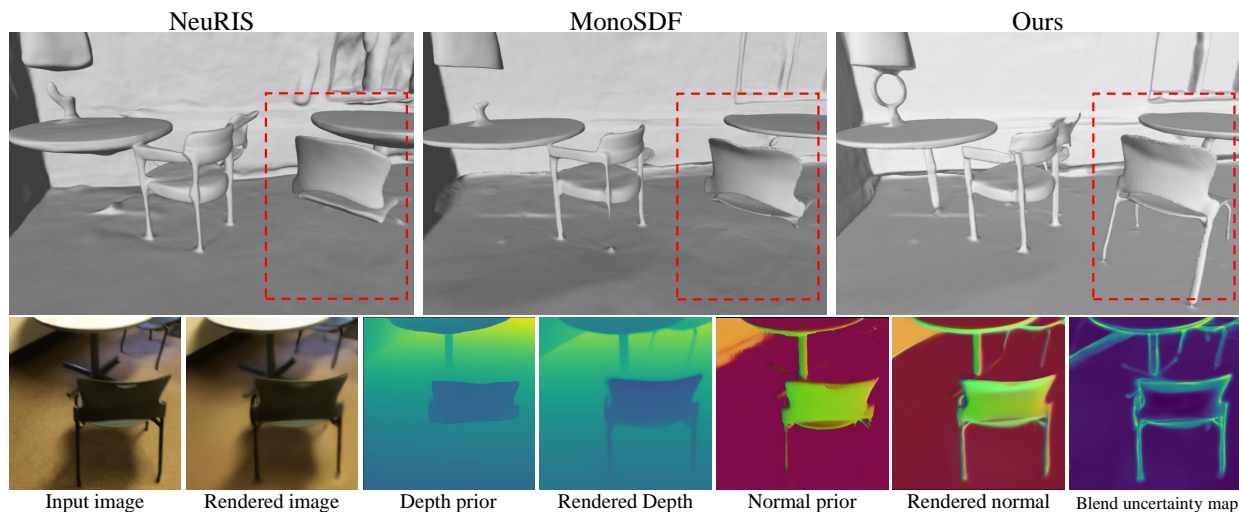


Fig. 1: We can observe that our method can reconstruct the indoor scene with more detailed structures, such as the chair legs and bracket of the desk lamp. Previous works such as MonoSDF [1], which is based on the VolSDF [2], can not reconstruct the thin and detailed surface due to the inaccurate geometry prior at these regions. Our method can accurately generate the uncertainty map, which can localize the inaccurate priors and reduce the bias in SDF-based rendering with a proposed bias-aware SDF to density transformation approach so that our method can reconstruct the indoor scene significantly better than previous works.

monocular geometric priors in a unified perspective. Our key insight is based on the assumption that a prior is correct if it is consistent with other priors, e.g. consistent with other views or other modalities (depth normal consistency in our case). While manually detecting this consistency or inconsistency can be challenging, we observe that the optimization of implicit surface implicitly aggregates information/priors from multi-views and multi-modality (depth and normal). Geometric priors that are consistent with others will have low errors, while priors that deviate from others will have large errors in comparison to learned implicit surfaces. Therefore, we achieve this with uncertainty modeling where large uncertainty reflects the large error of the priors. We introduce the uncertainty into the monocular prior guided optimization to avoid the effect of wrong priors in surface reconstruction. We observe that the filtered regions usually correspond to thin structure regions, as shown in the *Importance map* in Figure 1. Therefore, we propose an importance-guided ray sampling that samples more rays from these regions and an adaptive smoothness regularization strategy that applies smaller regularization on these regions to facilitate the learning of fine structures.

Further, we find that the commonly used SDF to density transformation has a non-negligible bias when a ray is casting around the surface where multiple peaks exist in the volume rendering weights. Optimization with such an ambiguous formula suppresses the reconstruction of foreground objects, e.g. thin structures in indoor scenes. We propose to transform the SDF to density with the curvature radius and the angle between the view direction and the SDF normals. As computing the Hessian matrix to obtain the analytical solution is computationally expensive, we approximate the curvature radius with a triangle from adjacent points. With the proposed bias-aware SDF to density transformation, our method reconstructs fine details faithfully for indoor scenes,

as shown in Figure 1.

In summary, we make the following contributions:

- We identify the key reasons for previous methods failing to recover thin and detailed structures in indoor scenes. Consequently, we introduce **DebSDF**, which utilizes uncertainty modeling for filtering large error monocular priors, guiding the ray sampling, and applying smoothness regularization adaptively for reconstructing detailed indoor scenes.
- We propose a novel bias-aware SDF to density transformation for volume rendering, enabling the reconstruction of thin and detailed structures.
- Extensive experiments on four challenging datasets, i.e., ScanNet [11], ICL-NUIM [12], Replica [13], and Tanks and Temples [14], validate the effectiveness of our method quantitatively and qualitatively.

The rest of the paper is organized as follows: In Sec. 2, we introduce some related works and contrast them with our neural indoor reconstruction methods to highlight our contribution. In Sec. 3, we introduce the details of our method, including uncertainty-guided prior localization, filtering, ray sampling, regularization, and our bias-aware SDF to density transformation. We conduct extensive experiments to validate their effectiveness on multiple challenging datasets in Sec. 4 and conclude our work in Sec. 5.

2 RELATED WORKS

Multi-view stereo. Multi-view stereo is a technique in computer vision that involves reconstructing a 3D scene from multiple 2D images captured from different viewpoints. Due to its wide-ranging applications in various fields, such as robotics, augmented reality, and computer graphics, it has been extensively studied within the field of computer vision. The classic multi-view stereo algorithms [15], [16], [17], [18], [19] employ patch matching across multiple images to

estimate the depth [15], [19] of each pixel. Some methods [20], [21], [22] utilize voxel representation to model shapes. However, these methods often encounter difficulties with the texture-less regions in the indoor scene since depth estimation and patch matching are difficult in these areas.

Recent advancement of multi-view stereo focusing on the usage of deep learning approaches. Some learning-based multi-view stereo methods replace some steps in the classic MVS pipeline. [23], [24], [25], [26] leverage a 3D CNN for depth estimation, yielding superior results. Most of these methods apply the MVSNet [24] as the backbone, which is computationally expensive due to the large number of parameters involved in 3D CNN. Some works also apply the pyramid structure [27], [28], [29] to reduce the parameter quantity, while other works incorporate patch matching [30], attention mechanism [31], and semantic segmentation [32] for performance improvement. However, these approaches often suffer from geometry inconsistency since the depth maps are estimated individually for each view, resulting in noisy surfaces with many hollows.

Neural scene representations. Instead of the multi-stage pipeline, some learning-based approaches propose an end-to-end framework to reconstruct 3D scenes based on the truncated signed distance function (TSDF) representation. Atlas [33] extracts the image features by 2D CNN and back projects features from each view to the same 3D voxel volume. A 3D CNN is applied to infer the TSDF and labels of each point. NeuralRecon [34] utilizes a recurrent network for feature fusion from previous fragments sequentially. This can enhance the ability to both represent global and local details and reconstruct 3D scenes in real-time. TransformerFusion [35] applies the transformer for more efficient feature fusion from multi-view images, which benefits from the attention mechanism.

Differing from treating the images as input to extract features by CNN, some methods [2], [4], [5], [6], [36], [37], [38], [39], [40] represent the scene as a coordinate-based implicit neural network and apply the differentiable neural rendering technique for supervision. Benefiting from the contiguous representation, these methods can generate arbitrary reconstruction results with only a simple fully-connected deep network. Most of these methods have been integrated into open-source frameworks such as NeRFStudio [41] and SDFStudio [42] to promote the development of neural fields. Among them, NeRF maps the 5D coordinate to the volume density, and view-dependent emitted radiance, which achieves outstanding novel view synthesis performance, but the geometry is noisy. IDR [37] represents the geometry as the zero level-set of a neural network and utilizes the surface rendering with the masks of each view to simultaneously optimize the geometry, view-dependent color, and camera poses. VolSDF [2], NeuS [4], and HF-NeuS [43] apply SDF to represent the geometry and transform SDF to volumetric density such that the volume rendering can be used to render the image from each viewpoint. However, these methods usually fail in the reconstruction of the texture-less regions, such as the pure white wall in the indoor scene. This would lead the optimization to the local optimum, which indicates incorrect geometry results.

Indoor scene reconstruction with priors. To tackle the aforementioned local minimization issue caused by the

texture-less regions, some NeRF-based methods incorporate prior information such as depth smoothness [44], semantic mask [45], [46], and depth prior [36], [47], [48]. [36], [47], [48] integrate the depth information from the Structure-from-motion (Sfm) technique to regularize the geometry optimization. DSNeRF [47] and NerfingMVS [36] apply the sparse depth map from COLMAP [15] to regularize the optimization. DDPNeRF [48] utilizes a pre-trained depth completion network to predict the dense depth map and a 2D uncertainty map to regularize the estimated depth and standard deviation of rendering weights. However, the predicted 2D uncertainty map from the pre-trained network suffers from the domain gap problem and cannot keep consistency in 3D space. And the standard deviation of rendering weights is not suitable to apply to the SDF-based methods [2], [4], [5] since the SDF-to-density transformation constrains the density distribution.

Although the above NeRF-based methods can generate high-quality novel view synthesis results, the geometry reconstruction results are still noisy and incomplete. Some previous works [1], [8], [10] focus on generating high-fidelity geometry by applying monocular priors such as depth and normal estimated by the pre-trained network. Manhattan-SDF [9] regularize the surface normal at the wall and floor regions under the Manhattan world assumption [8]. NeuRIS [10] adopts the normal prior for regularization and uses an adaptive manner to impose the normal prior based on the assumption that the normal priors are not faithful at the pixels with rich visual features. MonoSDF [1] utilizes both the depth and normal priors predicted by the Omnidata model [49] for 3D indoor scene reconstruction. The aforementioned previous works use the geometry prior for regularization and [10] filter the normal prior according to the image feature. This can significantly improve the reconstruction quality in smooth and simple regions but still can not reconstruct the complex detailed surfaces well. The TUVr [39] utilizes the MVS network [24], [27] for depth prior prediction and applies the cosine of the normal and ray direction to reduce the bias. The [50] applies the self-supervised super-plane constraint by exploring the indoor reconstruction without the geometry cues. Similar to [50], the HelixSurf [51] uses an unsupervised MVS approach to predict the depth point for intertwined regularization. The [52] directly uses the signed distance function in sparse voxel block grids for fast reconstruction without MLPs.

These methods [1], [10], [39] cannot robustly filter the inaccurate monocular priors, which is important to reconstruct the thin and detailed surfaces. Besides, other methods [2], [4], [39] ignore the bias in SDF-based rendering caused by the curvature of SDF. In this work, we focus on these problems to boost the 3D reconstruction of thin structures.

3 OUR METHOD

Previous works [1], [9], [10] have shown that regularizing the optimization with geometry priors can significantly improve the reconstruction quality at texture-less areas such as the wall and floor within the implicit neural surface representation framework and SDF-based volume rendering. However, it is still difficult to reconstruct the complex and detailed surface, especially when it is less observed in the

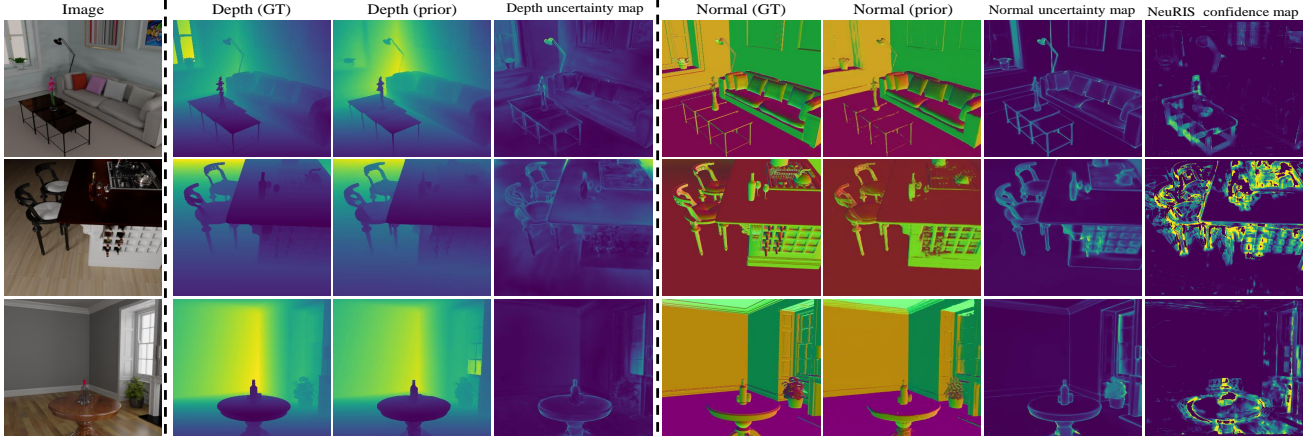


Fig. 2: The rendered uncertainty map can localize the regions where the monocular prior is inaccurate, which usually corresponds to thin structures with high levels of detail. Compared with NeuRIS [10], our model is capable of generating more reasonable results.

indoor scene, such as the legs of the chair. We analyze that there are some reasons for this problem: (i) The obtained geometry priors have significantly larger errors in these regions than in other planar regions. (ii) Areas of fine detail occupy a small area in the indoor scene, so the unbalanced sampling harms the reconstruction quality. (iii) Applying the smooth regularization indiscriminately degrades the reconstruction of high-frequency signals. (iv) The SDF-based volume rendering has geometry bias resulting from the curvature of SDF, which leads to the elimination of fine and detailed thin geometry structures, especially with regularization from monocular geometry prior. To tackle these problems, we propose **DebSDF**, which filters the inaccurate monocular priors and uses a bias-aware transformation from SDF to density to reduce the ambiguity of density representation such that the elimination of the fine structure problem can be solved.

3.1 Preliminaries

Following previous works [1], [2], we apply the implicit neural network to represent the geometry and radiance field and optimize this by differentiable volume rendering. Suppose a ray \mathbf{r} is cast from the camera location \mathbf{o} and passes through the pixel along the ray direction \mathbf{v} . N points are sampled on the ray and the i -th point is defined as $\mathbf{r}(t_i) = \mathbf{o} + t_i\mathbf{v}$ where the t_i is the distance to camera. The SDF s_i and color value \mathbf{c}_i corresponding to the $\mathbf{r}(t_i)$ is predicted by the implicit network. To apply the volume rendering technique, the transformation from SDF s to density can be defined as the Laplace CDF [2]:

$$\sigma(s) = \begin{cases} \frac{1}{2\beta} \exp(-\frac{s}{\beta}) & \text{if } s > 0 \\ \frac{1}{\beta} - \frac{1}{2\beta} \exp(\frac{s}{\beta}) & \text{if } s \leq 0 \end{cases} \quad (1)$$

or the Logisitic CDF [43], which proved with less bias than the former:

$$\sigma(s) = \frac{1}{\beta} \cdot \frac{1}{1 + \exp(-\frac{s}{\beta})} \quad (2)$$

where the β is a learnable parameter.

The rendered color [6] of the ray \mathbf{r} in the space is:

$$\hat{\mathbf{C}}(\mathbf{r}) = \sum_{i=1}^N T_i \alpha_i \mathbf{c}_i, \quad (3)$$

where T_i and α_i denote the transparency and alpha value at the i^{th} point on the ray \mathbf{r} , respectively [6]:

$$T_i = \prod_{j=1}^{i-1} (1 - \alpha_j), \quad \alpha_i = 1 - \exp(-\sigma_i \delta_i) \quad (4)$$

The rendered depth $\hat{D}(\mathbf{r})$ and normal $\hat{\mathbf{N}}(\mathbf{r})$ corresponding to the surface intersecting the ray \mathbf{r} are [1], [36], [48]:

$$\hat{D}(\mathbf{r}) = \sum_{i=1}^N T_i \alpha_i t_i, \quad \hat{\mathbf{N}}(\mathbf{r}) = \sum_{i=1}^N T_i \alpha_i \mathbf{n}_i \quad (5)$$

where the \mathbf{n}_i is the SDF gradient of i -th point on ray \mathbf{r} .

The depth estimation network predicts the depth only up to scale, so the scale w and the shift q computed by the least-squares method [53] is applied to normalize the depth prior, which is denoted as $D(\mathbf{r}) = wD'(\mathbf{r}) + q$. The $D'(\mathbf{r})$ is the depth prior predicted by the pre-trained depth estimation network. The monocular depth and normal loss function for regularization of previous works [1], [10] are:

$$\begin{aligned} \mathcal{L}_{\text{depth}} &= \sum_{\mathbf{r} \in \mathcal{R}} \|w\hat{D}(\mathbf{r}) + q - D(\mathbf{r})\|^2 \\ \mathcal{L}_{\text{normal}} &= \sum_{\mathbf{r} \in \mathcal{R}} \|\hat{\mathbf{N}}(\mathbf{r}) - \mathbf{N}(\mathbf{r})\|_1 + \|1 - \hat{\mathbf{N}}(\mathbf{r})^T \mathbf{N}(\mathbf{r})\|_1 \end{aligned} \quad (6)$$

where the $D(\mathbf{r})$ and $\mathbf{N}(\mathbf{r})$ are the depth and normal prior obtained from the pre-trained Omnidata model [49].

3.2 Uncertainty Guided Prior Filtering

Since the monocular priors provided by the pre-trained models, such as Omnidata [49] and SNU [54], are not perfectly accurate, it is necessary to apply an adaptive strategy to filter the monocular prior. NeuRIS [10] filters the monocular prior based on the assumption that the regions where the monocular priors are not faithful typically consist of high-frequency features or irregular shapes with relatively

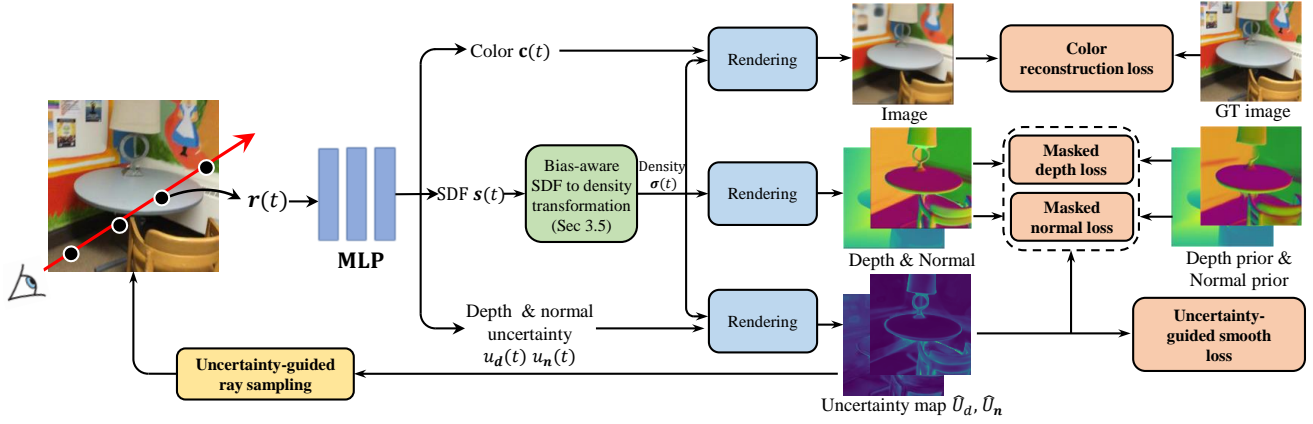


Fig. 3: The overview of our method. We propose the masked uncertainty learning to adaptively filter the geometry prior and localize the detailed and thin region in 5D space so that the small and thin structure would not be lost due to the wrong prior. Then, the localized uncertainty maps are utilized to guide ray sampling and smooth regularization to improve the reconstruction details of geometry. Besides, we analyze the bias in volume rendering caused by the transformation from SDF to density, which has a significant negative impact on the small and thin structure with geometry prior. A bias-aware SDF to density transformation is proposed to significantly reduce the bias for reconstructing the small and thin objects.

rich visual features in the input images. However, this assumption does not generalize well since the monocular prior could be faithful at the simple planar regions, and some planar surfaces also have high-frequency appearance features, such as the wall with lots of texture details. Instead of applying the image feature for monocular prior filtering, we utilize the prior uncertainty from multi-view to filter the faithful prior. Specifically, the prior from a viewpoint is considered to be inaccurate if it has a large variance from other viewpoints. Besides, this also cannot guarantee the occlusion-aware property since whether a point on the ray is visible to other views is unknown. Based on the observation that the inaccurate priors usually have a large variance from multiple viewpoints, we introduce the masked uncertainty learning loss function to model this variance.

Following the [1], [2], the SDF values are predicted by an coordinate-based implicit network f_g :

$$s, \mathbf{z} = f_g(\mathbf{x}) \quad (7)$$

where the \mathbf{z} is the feature vector and the \mathbf{x} is the point coordinate on the ray \mathbf{r} .

To model the prior uncertainty of a pixel, a straightforward approach is to apply the variance of geometry prior from different viewpoints at the same point on the surface as the prior uncertainty. However, this approach models the uncertainty view-independent, which is the main drawback since only the viewpoints with faithful priors need to be filtered. Besides, it is still unknown whether the points on a queried ray are visible to the other views due to occlusion. The uncertainty computed by the pre-trained depth or normal estimation network [48], [54] has low accuracy because of the domain gap. Due to this reason, we model the uncertainty of the prior as a view-dependent representation by the volume rendering [55], [56].

Suppose the uncertainty scores at each point for modeling the variance corresponding to the monocular prior are $u_d \in \mathbb{R}$ and $\mathbf{v}_n \in \mathbb{R}^3$, respectively. We apply the view-

dependent color network to predict the uncertainty scores:

$$\mathbf{c}, u_d, \mathbf{v}_n = f_c(\mathbf{x}, \mathbf{v}, \mathbf{n}, \mathbf{z}) \quad (8)$$

We compute the depth uncertainty score $\hat{U}_d(\mathbf{r})$ and normal uncertainty vector $\hat{\mathbf{V}}_n(\mathbf{r})$ of the pixel corresponds to the ray \mathbf{r} based on the volume rendering:

$$\hat{U}_d(\mathbf{r}) = \sum_{i=1}^N T_i \alpha_i u_{di}, \quad \hat{\mathbf{V}}_n(\mathbf{r}) = \sum_{i=1}^N T_i \alpha_i \mathbf{v}_{ni}. \quad (9)$$

We compute the normal uncertainty score $\hat{U}_n(\mathbf{r})$ as the mean value of the normal uncertainty vector $\hat{\mathbf{V}}_n(\mathbf{r})$ with 3 dimensions.

Moreover, we design the loss function based on the uncertainty learning perception with the mask to optimize the implicitly represented uncertainty field. The mask is applied to filter the negative impact from monocular prior with large uncertainty, which indicates a large probability of being inaccurate.

Masked Depth loss. The $\hat{D}(\mathbf{r})$ is the predicted geometry prior. We design the loss function as:

$$\mathcal{L}_{\text{Mdepth}} = \ln |\hat{U}_d(\mathbf{r})| + \frac{\Omega(\hat{U}_d(\mathbf{r}), \tau_d) \odot |\hat{D}(\mathbf{r}) - D(\mathbf{r})|}{|\hat{U}_d(\mathbf{r})|} \quad (10)$$

where $\Omega(U, \tau) \odot F$ is an adaptive gradient detach operation. The gradients from loss function would be detached if $U > \tau$ while the gradients are not detached if $U \leq \tau$. This indicates that the geometry prior is not utilized in the regions with high uncertainty while the uncertainty score will still be optimized.

Masked Normal loss. Similar to the depth loss, the normal prior is transformed into the world coordinate space for regularization and computing the uncertainty.

$$\mathcal{L}_{\text{Mnormal}} = \ln \hat{U}_n^2(\mathbf{r}) + \frac{\Omega(\hat{U}_n(\mathbf{r}), \tau_n) \odot \|\hat{\mathbf{N}}(\mathbf{r}) - \mathbf{N}(\mathbf{r})\|_2}{\hat{U}_n^2(\mathbf{r})} \quad (11)$$

Specifically, the inconsistency can be considered as the adaptive weight to adjust the monocular prior to each pixel. Low weight is applied to the prior with a large inconsistency.

Color reconstruction loss. Optimize the scene representation though the observation in 2D space.

$$\mathcal{L}_{\text{rgb}} = \sum_{\mathbf{r} \in \mathcal{R}} \|\hat{\mathbf{C}}(\mathbf{r}) - \mathbf{C}(\mathbf{r})\|_1 \quad (12)$$

Eikonal loss. Following the [57], the Eikonal loss is applied such that the property of SDF can be satisfied.

$$\mathcal{L}_{\text{eik}} = \sum_{\mathbf{x} \in \mathcal{X}} (\|\nabla f_g(\mathbf{x})\|_2 - 1)^2 \quad (13)$$

where the \mathcal{X} is the set of points uniformly sampled in the 3D space and the regions near the surface.

The uncertainty map can estimate the multi-view consistency of depth and normal prior. If the most geometry prior to a region is inaccurate, 2 edge cases may occur:

- The geometry prior from a viewpoint is inaccurate, but the corresponding uncertainty is low.
- The geometry prior from a viewpoint is accurate, but the corresponding uncertainty is high.

Since the geometry prior is predicted by a pre-trained model and the real prior is unavailable, it is impossible to filter the inaccurate prior perfectly. However, our uncertainty-guided prior filtering can still filter the vast majority of incorrect prior. Since the inaccurate prior is unlikely to be multi-view consistent. After the negative influence of inaccurate prior is reduced, the color reconstruction loss would gradually dominate. Based on the this analysis, these 2 edge cases would become increasingly rare along with the training.

3.3 Uncertainty-Guided Ray Sampling

Though applying the geometry prior with the uncertainty-based filter can improve the 3D reconstruction quality since inconsistency prior for some pixels can be filtered out, the fine and detailed structures are still hard to reconstruct. We observe that the thin object structures only occupy a small area in the image, so the probability of being sampled is low. In contrast, the texture-less and planar regions occupy most of the room's area and can already be reconstructed with high fidelity by a small number of ray samples, which benefits from the geometry prior supervision. Sampling more rays on these simple geometry surfaces than complex and fine geometry causes computational waste.

Localizing the thin geometry surface, which usually corresponds to the high-frequency surface, is perceptual. A straightforward approach is to apply a high pass filter or keypoint extractor to localize the location of high-frequency signal for sampling, which generates incorrect results at the planar surfaces with high-frequency color appearance. Even with the auxiliary prior information such as the geometry prior [49], directly utilizing the high-frequency part from geometry prior is still inaccurate since the geometry prior would miss some detailed structure, which indicates predicting these regions as the smooth and planar surfaces. This can be observed in Fig. 1 where the chair legs are lost.

Based on this analysis, we compute the blend uncertainty score $A(\mathbf{r})$ of each pixel by combining both inconsistency representations for depth and normal:

$$A(\mathbf{r}) = \hat{U}_d^{1-\lambda}(\mathbf{r}) \hat{U}_n^\lambda(\mathbf{r}) \quad (14)$$




Fig. 4: As shown in sub-figure (i), the TUVR [39] applies the cosine between the ray direction and normal $\cos \theta$ to reduce the bias. Specifically, it transforms the SDF s to the depth $s/|\cos \theta|$ ($of \rightarrow od$) by assuming the ray intersects with the planar surface. However, TUVR [39], NeuS [4], and HF-NeuS [43] only consider the planar surface, which ignores the curvature of the surface. As shown in sub-figure (ii), the dd' indicates the error between the rendered depth and the real depth, which causes the biased volume rendering.

where the λ is a hyper-parameter. The blend uncertainty score is utilized as the guidance for the ray sampling.

Compared with inferring the confident maps for monocular prior filtering and ray sampling by only detecting the high-frequency image features, our method infers the uncertainty maps by using the information from multiple viewpoints implicitly. For each ray \mathbf{r} , the probability to be sampled is calculated as:

$$p(\mathbf{r}_i) = \frac{A(\mathbf{r}_i)}{\sum_i A(\mathbf{r}_i)} \quad (15)$$

3.4 Uncertainty-Guided Smooth Regularization

To avoid the reconstructed surfaces being too noisy, smooth regularization [7] is widely applied to reduce the floater on the surfaces. The smooth loss [7] requires the gradients of SDF to be the same in a local region, which can reduce the floaters near the surface.

However, this regularization would also damage the fine details, which indicates that not all the surfaces in the indoor scene need to be smooth. According to this analysis, we utilize the smooth regularization term in an adaptive manner, which can not only keep the simple surfaces smooth but also preserve the fine details of complex geometry.

For each sampled ray \mathbf{r} , the $\mathcal{S}(\mathbf{r})$ denotes the points sampled near the surface along this ray. The smoothness loss term is:

$$\mathcal{L}_{\text{smooth}} = \mathbb{M}(A(\mathbf{r}), \tau_s) \sum_{\mathbf{x} \in \mathcal{S}(\mathbf{r})} \|\nabla f_g(\mathbf{x}) - \nabla f_g(\mathbf{x} + \epsilon)\| \quad (16)$$

where the ϵ is a random offset sampled on a Gaussian distribution $\mathcal{N}(0, \xi)$ whose ξ is a small variance. The $\mathbb{M}(A(\mathbf{r}), \tau_s)$ is a mask function:

$$\mathbb{M}(A(\mathbf{r}), \tau_s) = \begin{cases} 1 & \text{if } A(\mathbf{r}) \leq \tau_s \\ 0 & \text{if } A(\mathbf{r}) > \tau_s \end{cases} \quad (17)$$

Apart from the points sampled on the rays, a small batch of points is uniformly randomly sampled in the indoor space, and the aforementioned smoothness regularization is also applied. Since the probability of sampling near the

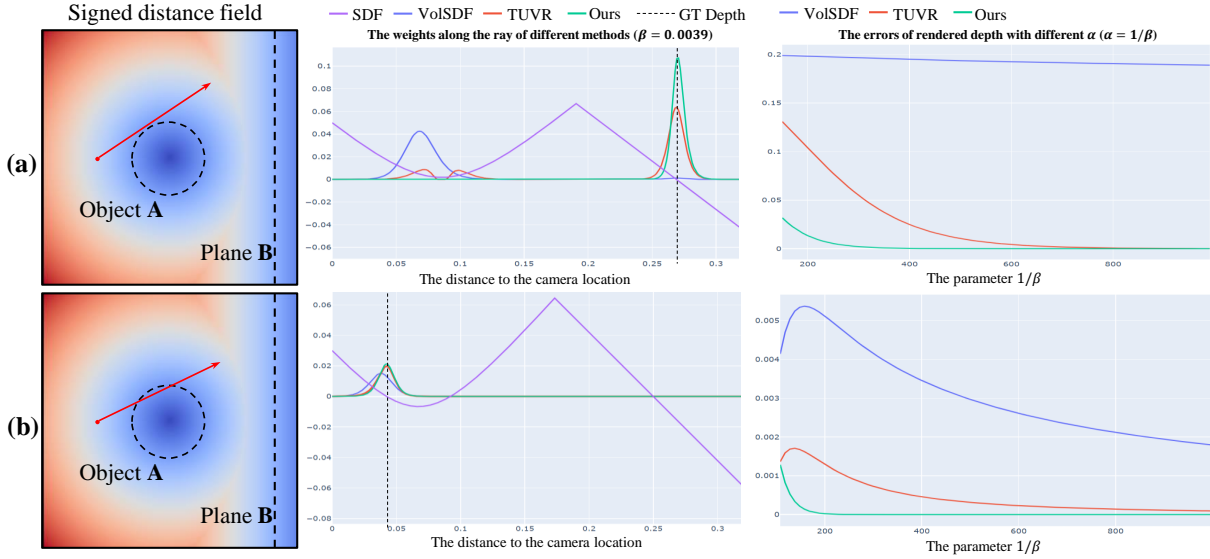


Fig. 5: We demonstrate 2 toy cases to simulate the SDF-based rendering by applying Logistic CDF when a ray brushes past and intersects with a small object. The weight function along the ray demonstrates that our method does not have a peak if there is no intersection with the surface. The sub-figures of the right column show the rendered depth corresponding to different α values. Our method can achieve smaller errors than previous works such as the TUVR [39] and the VolSDF [2]. We do not demonstrate NeuS [4] since the TUVR [39] demonstrates that its bias is smaller than NeuS [4].

surface is very low, the uncertainty-based adaptive manner is not used. Besides, we do not apply the adaptive manner on the Eikonal loss since the SDF around the small and thin structure still satisfies this property.

3.5 Bias-aware SDF to Density Transformation

Previous works [2], [4], [5] utilize a transformation from SDF to volumetric density and volume rendering technique for modeling the connection between the 3D geometry and the rendered images. All these methods are designed to model a weight function that is required to be unbiased for the zero-level set of SDF.

However, these methods, such as VolSDF [2] and NeuS [4], still suffer from biased volume rendering. The TUVR [39] analyzes the problem of these methods and proposes to scale the SDF with the angle between the normal and the ray direction, but it still ignore the bias caused by the curvature of SDF. Based on this, we think that the curvature of SDF should be considered for unbiased volume rendering. In this section, we analyze the weight function from volume rendering near the “small object” and design a bias-aware transformation from SDF to density to improve the reconstruction quality of the small and fine structures.

3.5.1 The Analysis of Bias in the SDF to Density Transformation

We can divide the design of ideal unbiased volume rendering into 2 steps. Firstly, design a transformation function from SDF to density which can ensure unbiased when the ray intersects with the surface vertically, which corresponds to the ray A in Fig. 4. (a), the SDF value equal to the depth. This has been discussed in some previous works [2], [4], [5], [43] while the Logistic transformation applied by HF-NeuS [43] can achieve unbiased transformation. Secondly,

since the ray may not intersect with the surface vertically, a transformation from the SDF to the depth is needed. The transformation proposed by TUVR [39] is shown in Fig. 4. It applies the cosine of the angle between the normal and the ray direction $\cos \theta$ for SDF transformation, which corresponds to ray B in Fig. 4 (i). Besides, the NeuS [4] multiplies $\cos \theta$ by the interval between points on the ray to adjust this situation. However, these methods assume the intersection with a planar surface, which ignores the curvature of the surface. As shown in Fig. 4, the bias still exists for the curved surface. Based on the aforementioned analysis, we propose a transformation from SDF to density, considering the curvature of SDF for reducing bias. Specifically, we assume that the surface is composed of multiple arcs and design a transformation from SDF to depth for the second step corresponding to the ($of \rightarrow od$) in Fig. 4 (ii).

We set 2 simple toy cases in Fig. 5 to demonstrate the benefits of considering the curvature. There is a small object A and a plane B in space. Case (a): A ray emits on the plane B without intersecting with object A but being close to the object A. Case (b): A ray intersects with the object A. In sub-figure (a), though the ray does not intersect with object A, the weight functions along the ray of VolSDF [2] and TUVR [39] are still influenced by object A. Multiple peaks can be observed on the weight function when near object A. In sub-figure (b), it can be observed that the previous works have larger depth errors than our method. Especially the case shown in sub-figure (a) leads to multiple peaks, which have significant ambiguity for optimization.

Because of this, the optimization of the masked depth loss, which minimizes the error between depth prior and rendered depth, would lead to the reduction of the density value around the peaks near the object A. And the masked normal loss suffers from the same problem as depth. Though the points sampled around the multiple peaks along




Fig. 6: The demonstration of our bias-aware transformation from SDF to density. The curvature radius $a > 0$ indicates the intersection to the inwardly curved surface, while the curvature radius $a < 0$ indicates the intersection with the outward curved surface. If the $|a| < |a + s| \sin \theta$ is satisfied, this indicates that the ray and the surface nearby have no intersection point. The fourth sub-figure illustrates the curvature radius estimation by difference method. We apply 2 adjacent points on the ray to estimate the curvature radius of the previous point.

the ray are not inside the object A , it would still influence the SDF representation of object A due to the local contiguity of the implicit neural network and the regularization of Eikonal loss. This ambiguity causes the inefficient optimization of SDF for the fine and small objects in the indoor scene. Specifically, this problem causes the small and fine structures to be smaller or even disappear.

Our model applies a better transformation from SDF to density, considering the curvature of the surface can have a smaller bias than other methods, which can have better reconstruction for the fine and detailed regions.

3.5.2 SDF to Density Mapping for Bias Reduction

To tackle the aforementioned problem, we design a SDF to density transformation based on the analysis of the SDF curvature and the extreme point of weight function. Specifically, we consider a simple scenario: a ray \mathbf{r} intersects a circle of radius a at point L . As shown in Fig. 6, the SDF of point $\mathbf{r}(t)$ is s , but the ray reaches the circle surface through a distance of y . Our target is to design an SDF mapping function to replace the $s(t)$ in $\sigma(s(t))$ (Eq. 1) as $y(t)$. We propose a function to map the SDF value s to the distance y in the transformation to density, which aims to reduce the negative influence from the curvature of SDF and the ray direction. This function can be divided into 2 situations:

- Ray intersects with the surface.
- Ray does not intersect the surface.

For the first situation, which is shown in Fig. 6 (i, ii), the SDF mapping function is denoted as:

$$y(t) = (a + s(t)) |\cos \theta(t)| - \text{sign}(a) \cdot l \quad (18)$$

where the $\theta(t)$ is the angle between the view direction and the normal of SDF. The a is the curvature radius of the point D on the surface of the object A . Through Pythagorean theorem, the distance l is denoted as:

$$l = \sqrt{a^2 - (a + s)^2 \sin^2 \theta} \quad (19)$$

Like the SDF, we define the distance function $y(t)$ as a negative number, so the above mathematical equations still hold when the $\mathbf{r}(t)$ is inside the circle. This indicates that we assume the small local area on the surface can be approximated as a circular arc. For the planar surface,

it can be considered the corresponding absolute value of the curvature radius is a very large number. The distance function $y(t)$ can be simplified to $y(t) = s(t) / |\cos \theta(t)|$, which is the same as TUVR [39]. We apply the distance $y(t)$ as a kind of calibration to SDF $s(t)$ for volume rendering, so the density field of each point is computed as $\sigma(y(t))$.

Further, for the second situation which indicates the ray does not intersect with the surface, we design the mapping function based on the consideration that the ray tangent to another circle whose center is O_2 as shown in Fig. 6 (iii). Specifically, if the number within the square root of Eq. 19 is less than 0, which corresponds to the situation $|a| < |a + s| \sin \theta$, the ray does not intersect with the surface. For this situation, a naive solution is setting the $y(t)$ to infinity, which causes the discontinuity of the density field. To keep the density field contiguous, we design the distance $y(t)$ in this case as:

$$y(t) = \frac{s(t)}{|\cos \theta(t)|} + s(t) |\tan \theta(t)| \quad (20)$$

which indicates the tangency of another circle whose radius is larger than a .

Further, we apply the normal curvature radius of SDF at $\mathbf{r}(t)$ to estimate the aforementioned curvature radius a . The curvature radius $a(t)$ corresponding to the point $\mathbf{r}(t)$ is:

$$a(t) = R(\mathbf{r}(t), \mathbf{v}) - s(t) \quad (21)$$

where the $R(\mathbf{r}(t), \mathbf{v})$ [58], [59] is the normal curvature radius of SDF at point $\mathbf{r}(t)$ toward the ray direction \mathbf{v} .

3.5.3 Curvature Radius Estimation

Estimating the analytical solution of curvature radius requires computing the Hessian matrix, which is computationally unfriendly. We apply an approach to estimate the numerical solution of the curvature radius. As shown in the fourth sub-figure in Fig. 6, suppose there are 2 points A and B with distance d on the ray \mathbf{r} . The normals of point A and B are \mathbf{n}_A and \mathbf{n}_B respectively, and the \mathbf{n}'_B is the projection of \mathbf{n}_B on the plane determined by normal \mathbf{n}_A and ray direction \mathbf{v} . The α is the angle between \mathbf{n}_A and \mathbf{n}'_B .

The curvature radius R of point A can be estimated based on the Law of Sines:

$$R = \chi(\mathbf{n}_A, \mathbf{n}'_B) \cdot d \cdot \frac{\sin \theta_B}{\sin \alpha} \quad (22)$$

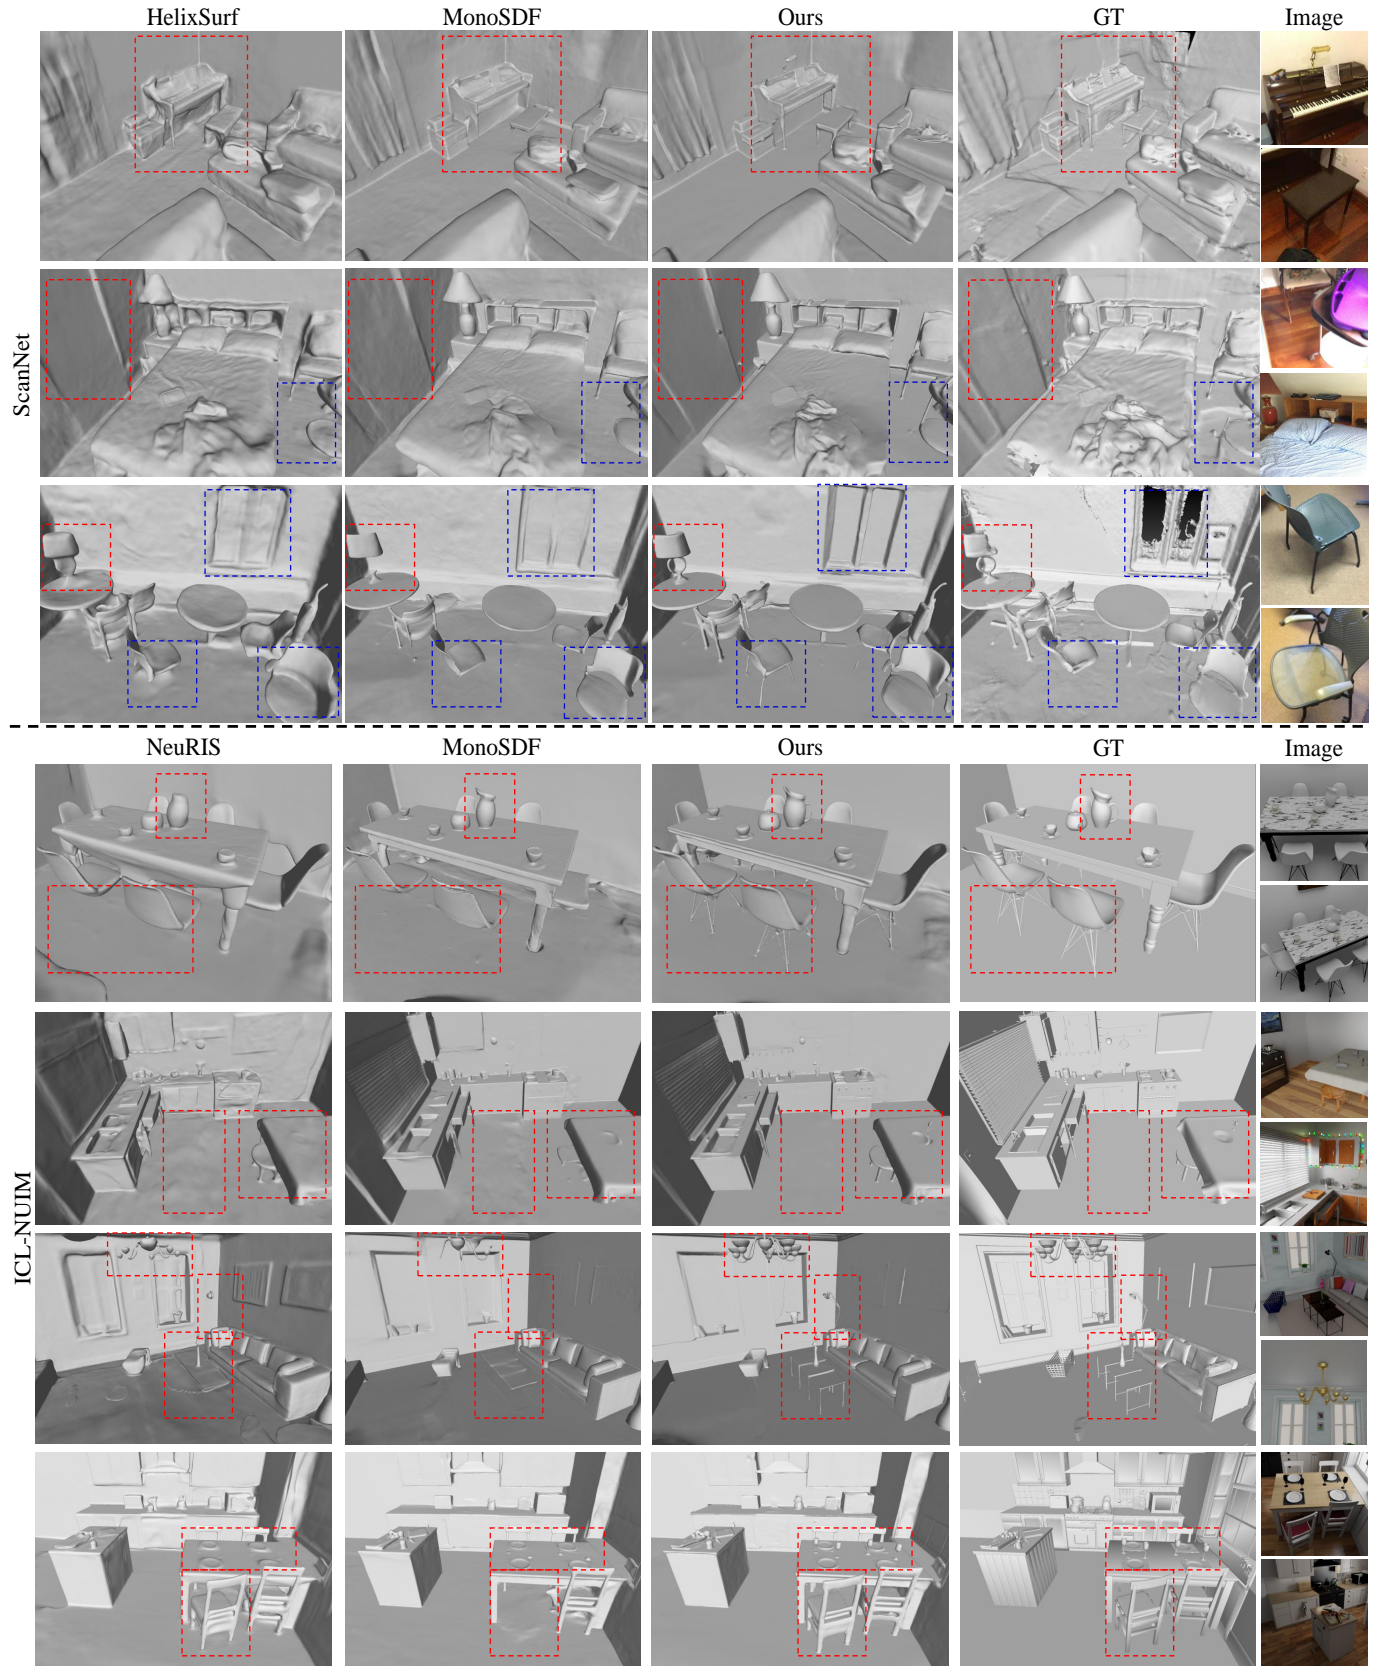


Fig. 7: The qualitative results on the ScanNet and ICL-NUIM dataset. We compare our method with previous state-of-the-art methods. Since the ground truth meshes of the ScanNet [11] dataset are generated by the RGBD sensor, it is still not the complete and accurate reconstruction results, especially for the thin and detailed structures such as the leg and handrail of chairs and the lamp on the piano, which are shown in the blue box. These can be confirmed by the column of images. We can observe that our method can reconstruct the thin and detailed surfaces significantly better than other methods.

where the $\chi(\mathbf{n}_A, \mathbf{n}'_B)$ is the indicator function:

$$\chi(\mathbf{n}_A, \mathbf{n}'_B) = \begin{cases} 1 & \text{if } \cos(\theta_A) \leq \cos(\theta_B) \\ -1 & \text{if } \cos(\theta_A) > \cos(\theta_B) \end{cases} \quad (23)$$

3.5.4 Progressive Warm-up

We design a progressive warm-up strategy to stabilize the training phase since the numerical estimation of the curvature radius is not stable at the beginning of training. Specifically, we replace the $\cos \theta(t)$ to $\cos^{p(t)} \theta(t)$, where the $p(t)$ is a parameter progressive growing with training iterations from 0 to 1. Then correspondingly, the $\sin \theta$ is adjusted according to $\sin^2 \theta + \cos^2 \theta = 1$. Moreover, since the sampling probability of the texture-less region is low, the growing $p(t)$ is designed as the product of a progressively increasing number of p' and the point-wise uncertainty score, such that the p of texture-less and planar regions increase slower than the detailed and important regions. The $p(t)$ corresponds to the point $\mathbf{r}(t)$ is denoted as:

$$p(t) = \min(p' \cdot u_d^{1-\lambda}(t) \cdot u_n^\lambda(t), 1) \quad (24)$$

where the $u_n(t)$ is the mean of the normal uncertainty score vector $\mathbf{v}_n(t)$ corresponds to each point instead of the $\hat{\mathbf{v}}_n$.

Optimization The total loss function to optimize the neural implicit geometry and color appearance network is:

$$\mathcal{L} = \mathcal{L}_{\text{rgb}} + \lambda_1 \mathcal{L}_{\text{eik}} + \lambda_2 \mathcal{L}_{\text{smooth}} + \lambda_3 \mathcal{L}_{\text{Mdepth}} + \lambda_4 \mathcal{L}_{\text{Mnormal}} \quad (25)$$

where the λ_i is the hyper-parameters for weighting each term of the loss function. We set $\lambda_1 = 0.05$, $\lambda_2 = 0.005$, $\lambda_3 = 0.006$, and $\lambda_4 = 0.0025$.

4 EXPERIMENTS

Datasets. We evaluate our method on the following 5 datasets: ScanNet [11], ICL-NUIM [12], Replica [13], Tanks and Temples [14] and DTU dataset [60]. Among them, ScanNet, Replica, and Tanks and Temples are real-world indoor scene datasets, and the Tanks and Temples dataset has primary large-scale indoor scenes. We select 4 scenes from ScanNet for performance evaluation by following the setting of Manhattan-SDF [9]. The ICL-NUIM [12] is a synthetic indoor scene dataset with the ground truth mesh.

Baselines. We compare our method with the following methods: (i) Neural volume rendering methods with prior, including MonoSDF [1], NeuRIS [10], Manhattan-SDF [9]; (ii) Neural volume rendering methods without prior, including VolSDF [2], NeuS [4], and Unisurf [5]; and (iii) Classical MVS reconstruction method: COLMAP [15].

Metrics. By following the evaluation protocol of [1], [10], we use the following evaluation metrics: Chamfer Distance (L_2), F-score with 5cm threshold, and Normal Consistency.

4.1 Implementation Details

By following the MonoSDF [1], we apply an MLP with 8 hidden layers as the geometry network to predict SDF and an MLP with 2 layers as the color network to predict the color field. Each layer has 256 hidden nodes. The Adam optimizer with a beginning learning rate of 5×10^{-4} is utilized to optimize the network. The learning rate exponentially decays in each iteration. The size of the input image

TABLE 1: The quantitative results on the ScanNet [11] dataset. Our method achieves state-of-the-art performance.

Method	Acc↓	Comp↓	Pre↑	Recall↑	Chamfer↓	F-score↑
COLMAP [15]	0.047	0.235	71.1	44.1	0.141	53.7
Unisurf [5]	0.554	0.164	21.2	36.2	0.359	26.7
VolSDF [2]	0.414	0.120	32.1	39.4	0.267	34.6
NeuS [4]	0.179	0.208	31.3	27.5	0.194	29.1
M-SDF [9]	0.072	0.068	62.1	56.8	0.070	60.2
NeuRIS [10]	0.050	0.049	71.7	66.9	0.050	69.2
MonoSDF [1]	0.035	0.048	79.9	68.1	0.042	73.3
TUVR [39]	0.038	0.042	76.6	72.9	0.040	74.7
HelixSurf [51]	0.038	0.044	78.6	72.7	0.042	75.5
Ours	0.036	0.040	80.7	76.5	0.038	78.5

TABLE 2: The quantitative results on the ICL-NUIM [12] and Replica [13] dataset. Our method achieves state-of-the-art performance on both datasets. It is noted that the official MonoSDF [1] implementation on the Replica dataset utilizes manually selected masks to disable the monocular priors of some images due to serious inaccuracy. In this paper, to validate the ability of our model to filter the inaccurate prior automatically, we do not apply this operation to MonoSDF [1] and our method for fair comparisons.

		Normal C↑	Chamfer↓	F-score↑
Replica	Unisurf [5]	90.96	0.049	78.99
	MonoSDF [1]	92.69	0.045	81.46
	Ours	93.23	0.029	88.36
ICL-NUIM	NeuRIS [10]	85.54	0.106	67.31
	MonoSDF [1]	87.51	0.104	67.39
	TUVR [39]	87.61	0.101	68.75
	Ours	88.34	0.092	77.38

is 384×384 , and the pre-trained Omnidata model [49] is applied to generate the geometry prior. We use PyTorch [61] to implement our model and train our model on one NVIDIA GeForce RTX 2080Ti GPU. We train our model for 200,000 iterations with 1024 rays sampled in each iteration. We set the $\lambda = 0.9$ to balance the uncertainty map localized from the depth and normal prior. We do not apply the estimated uncertainty regions to guide ray sampling and smoothing for the first 40,000 iterations since the uncertainty localization is not stable at the initial stage. To filter out the wrong prior, we set $\tau_d = 0.25$, $\tau_n = 0.4$, and $\tau_s = 0.3$.

4.2 Performance Comparison with Other Baselines

Regarding the baselines, COLMAP [15] reconstructs the mesh from the depth point cloud by the Poisson Surface reconstruction algorithm [3]. NeRF [6], Unisurf [5], NeuS [4], and VolSDF [2] apply the neural volume rendering technique for 3D reconstruction. The Manhattan-SDF [9], NeuRIS [10], and MonoSDF [1] utilize the auxiliary data from pre-trained models. As shown in Table 1 and Table 2, our method achieves state-of-the-art performance on multiple datasets. We also visualize the results of NeuRIS and MonoSDF, which are two state-of-the-art methods for indoor 3D reconstruction. The quantitative results are shown in Table 1 and Table 2; our method achieves significant improvement compared with other methods, which indicates the importance of the uncertainty-guided prior filtering and bias-aware SDF to density transformation. As shown in Fig. 7, our method well reconstructs the small and thin

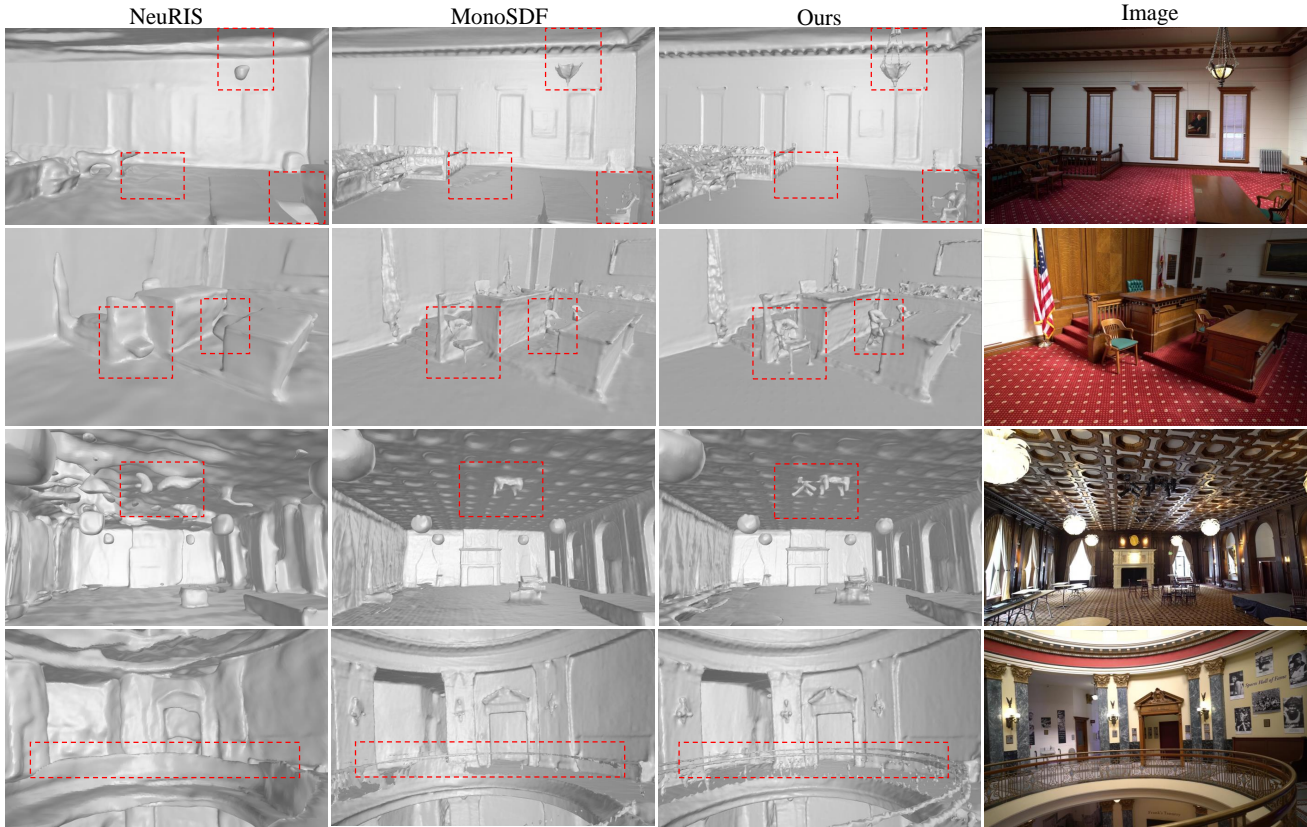


Fig. 8: We demonstrate the qualitative results on the Tanks and Temple [14] dataset. Our method reconstructs small objects better than MonoSDF [1], such as the pendant, chairs, and handrails of stairs. As shown in the first column, our method can also adapt to the inaccurate prior on the floor, so there is not a large pit on the floor. The Ground truth meshes of the Tanks and Temples [14] dataset are not released to the public.

TABLE 3: The quantitative results on Tanks and Temples dataset. * indicates using the hash encoding feature grids. The evaluation metric is the F-score, which is computed by the official website of the Tanks and Temple dataset.

Method	Auditorium	Ballroom	Courtroom	Museum	Mean
NeuRIS [10]	0.79	0.84	2.78	1.34	1.13
MonoSDF [1]	3.09	2.47	10.00	5.10	5.16
Ours	5.41	6.09	18.53	7.11	9.43
MonoSDF* [1]	3.17	3.70	13.75	5.68	6.58
Ours*	7.45	4.21	19.33	7.64	9.66

structures, such as the leg and the handrail of chairs, the bracket of the lamp, and the cups on the Table. Further, the reconstruction of the texture-less surfaces is also improved for our method. As shown in the first and second rows of Fig. 7, the reconstructed floor and wall of our method are smoother than other methods.

We also apply our methods on the Tanks and Temples dataset with/without hash encoding feature grids [1], [62]. As shown in Table 3, our method can achieve large improvement for both settings. The qualitative results are shown in Fig. 8. It can be observed that our method prevents reconstruction results from being corrupted by the wrong geometry priors and simultaneously improves the performance at thin and detailed regions.

4.3 Performance Evaluation of Different Regions

The aforementioned blend uncertainty map can be utilized to localize the thin and detailed structures in the image, so we divide each image into the masked and unmasked regions by setting different threshold τ_s 's for the blend uncertainty mask obtained from Eq. 17, and the masked regions and unmasked regions correspond to the detailed surface and simple planar surface in the indoor scene, respectively. The visualizations of both parts are shown in Fig. 10. Then, we use the depth abs rel, normal cos similarity, and normal L_1 similarity metrics to evaluate the improvements of different regions for each viewpoint. we generate the predicted depth and normal map from reconstructed meshes obtained from Marching Cubes [63] and volume rendering.

As shown in Table 5, the reconstruction quality improvements of our method on the complex detailed surfaces are more significant than those on the simple planar surfaces. For the bias-aware SDF to density transformation, the improvements of the masked regions are also more significant than the unmasked regions. This further validates the effectiveness of our method for the reconstruction of small objects and detailed regions.

4.4 Ablation Studies

4.4.1 Effectiveness of Different Modules

To validate the effectiveness of each proposed module, we conduct ablation studies on the ScanNet and ICL-NUIM




Fig. 9: The ablation studies of our method. "UF" denotes the uncertainty-guided prior filtering with the masked depth and normal loss; "US" denotes the uncertainty-guided ray sampling; "S" denotes the uncertainty-guided smooth. Compared with the baseline, the reconstruction of the thin and detailed structure benefits from our proposed modules.

TABLE 4: The ablation studies conducted on the ScanNet [11] and ICL-NUIM [12] datasets.

Prior Filtering	Uncertainty-Guided Ray Sampling	Uncertainty-Guided Smooth	Bias-Aware SDF to Density Transformation	ScanNet			ICL-NUIM		
				Normal C \uparrow	Chamfer \downarrow	F-score \uparrow	Normal C \uparrow	Chamfer \downarrow	F-score \uparrow
✓				87.85	0.0429	73.30	87.51	0.1034	67.39
✓				89.57	0.0402	76.82	88.12	0.0981	73.97
✓	✓			89.60	0.0399	77.12	87.99	0.0978	76.14
✓	✓	✓		89.98	0.0399	77.30	88.30	0.0953	76.75
✓	✓	✓	✓	90.21	0.0382	78.54	88.34	0.0928	77.59

TABLE 5: We divide each image in the ICL-NUIM dataset [12] into the masked region and unmasked region with the mask from the blend uncertainty map and compute the metrics for depth and normal map evaluation. The masked regions approximately correspond to the thin and detailed surface, while the unmasked regions correspond to the simple planar surface. "Mesh" indicates extracting the normal and depth map from the reconstructed mesh, while "Volume Rendering" indicates extracting the normal and depth map from volume rendering of the SDF.

τ_s	Method	Mesh						Volume Rendering					
		Masked region			Unmasked region			Masked region			Unmasked region		
		depth abs_rel \downarrow	normal cos \uparrow	normal $L_1 \downarrow$	depth abs_rel \downarrow	normal cos \uparrow	normal $L_1 \downarrow$	depth abs_rel \downarrow	normal cos \uparrow	normal $L_1 \downarrow$	depth abs_rel \downarrow	normal cos \uparrow	normal $L_1 \downarrow$
0.3	NeuRIS [10]	0.160	0.637	0.339	0.030	0.940	0.083	0.199	0.662	0.336	0.057	0.943	0.085
	MonoSDF [1]	0.167	0.664	0.323	0.026	0.962	0.054	0.180	0.650	0.331	0.035	0.959	0.057
	No Bias-aware	0.136	0.691	0.291	0.017	0.965	0.049	0.155	0.693	0.295	0.03	0.963	0.053
	Ours	0.117	0.702	0.287	0.016	0.967	0.049	0.133	0.700	0.289	0.027	0.964	0.052
0.4	NeuRIS [10]	0.185	0.623	0.345	0.032	0.934	0.089	0.243	0.646	0.344	0.059	0.938	0.089
	MonoSDF [1]	0.189	0.643	0.335	0.029	0.956	0.060	0.224	0.622	0.348	0.037	0.954	0.061
	No Bias-aware	0.166	0.674	0.298	0.019	0.959	0.055	0.199	0.660	0.312	0.031	0.959	0.056
	Ours	0.142	0.676	0.300	0.018	0.962	0.054	0.176	0.661	0.308	0.028	0.960	0.055

datasets. Five different configurations are investigated to evaluate the following modules: (1) Uncertainty-Guided Prior Filtering: applying the masked uncertainty learning to localize and filter out the noisy prior; (2) Uncertainty-Guided Ray sampling: sampling more rays from the localized uncertain regions; (3) Uncertainty-Guided Smooth: prohibiting the smooth regularization to the uncertain regions; (4) Bias-aware SDF to density transformation: utilizing curvature-based SDF mapping function to eliminate the bias in volume rendering. The quantitative results are shown in Table 4. The top row of Table 4 corresponds to the baseline MonoSDF [1].

Compared with MonoSDF [1], our uncertainty-guided prior filtering improves the F-score by 3.53 and 6.58 on ScanNet [11] and ICL-NUIM datasets, respectively. Based on this performance, the uncertainty-guided ray sampling and smooth regularization further improve the F-score by 1 to

TABLE 6: We evaluate the performance of our methods with different λ in Eq. 14. We set the λ is equal to 0.9.

λ	1	0.9	0.5	0.1	0
Chamfer \downarrow	0.095	0.093	0.095	0.096	0.097
F-score \uparrow	77.16	77.39	76.81	75.71	75.67
Normal C \uparrow	88.19	88.34	88.12	88.24	88.21

3 on both datasets. Further, the final performance increases by 1.24 and 0.84 in terms of F-score on these two datasets by introducing the bias-aware transformation from SDF to density module. Among these proposed modules, the performance improvement benefiting from the uncertainty-guided prior filtering, uncertainty-guided ray sampling, and the transformation from SDF to density than that from the uncertainty-guided smooth regularization.

As shown in Fig. 9, our method achieves a more stable

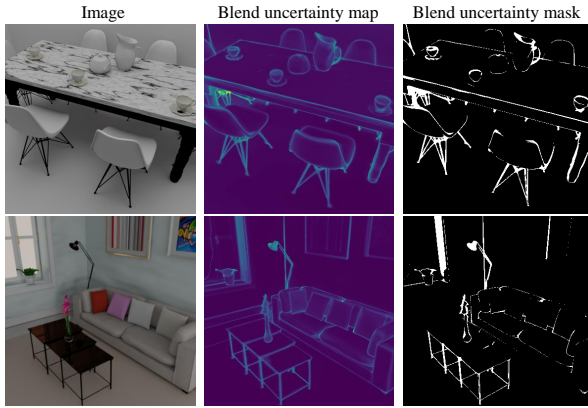


Fig. 10: The visualization of blend uncertainty map and mask. The mask can localize the fine and detailed regions in the indoor scene which would facilitate the ray sampling and smooth regularization.

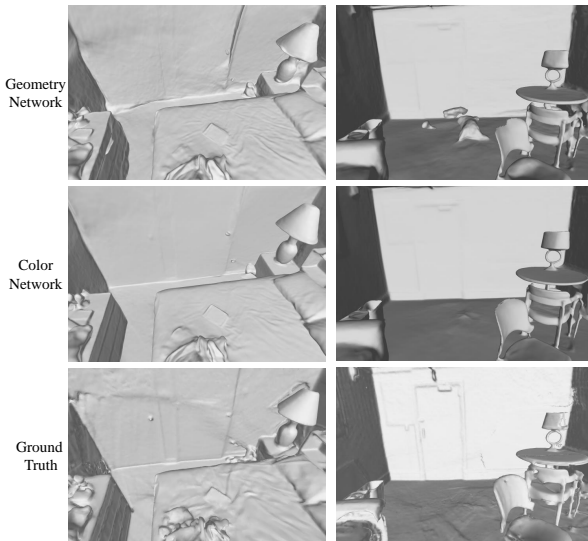


Fig. 11: The "Geometry Network" indicates predicting the uncertainty by the geometry network, which models the uncertainty view independently. The "color Network" indicates predicting the uncertainty by the color network which models the uncertainty view dependently.

reconstruction performance than MonoSDF. Some thin and detailed structures, such as the chair backrest and lamp, cannot be well reconstructed by MonoSDF, while these regions can be well reconstructed when applying prior filtering, uncertainty-guided ray sampling, and uncertainty-guided smooth regularization but still have structural damages. However, for those more tiny structures, such as the chair legs, still cannot be reconstructed, only applying these three modules is still not enough. But when introducing the bias-aware SDF to density transformation, these regions can be well reconstructed. This indicates that only detecting and filtering out the inaccurate geometry prior is not enough. A bias-aware SDF to density transformation is also indispensable for the reconstruction of the regions with finer details.

TABLE 7: We conduct experiments to evaluate the reconstruction performance of estimating the uncertainty by geometry network and color network on ScanNet [11] dataset. Evaluating the uncertainty by the color network is better than the geometry network.

Method	Acc↓	Comp↓	Pre↑	Recall↑	Chamfer↓	F-score ↑
GeoNet	0.039	0.044	75.6	71.2	0.042	73.2
ColorNet	0.036	0.040	80.7	76.5	0.038	78.5

4.4.2 Trade-off Between Different Uncertainties

We introduce a hyper-parameter λ in Eq. 14 to fuse the localized results from the depth and normal prior. As shown in Table 6, we set the λ with different values to study the effect of different priors in localizing the detailed and important regions. It can be observed that as λ gradually decreases, the performance decreases. This indicates that applying the normal prior for uncertainty-guided prior filtering can lead to better performance than utilizing the depth prior. The reason is that the obtained depth prior from the pre-trained model is defined only up to scale, such that the optimization of uncertainty would be misled by the normalization of depth prior. The normal prior can provide more accurate information than the depth prior since it is not related to the distance scale problem. These experiments show that the normal prior is more important for reconstructing the details of the indoor scene. We set $\lambda = 0.9$ in our implementation.

4.4.3 Modeling the Uncertainty

The uncertainty scores can be predicted by the geometry network or color network. The input of the geometry network is only the 3D space coordinate while the color network additionally considers the normal and ray direction. We conduct ablation studies to prove that modeling the prior uncertainty by the color network is better than a geometry network. The quantitative results are shown in Table 7. This proves that the modeling of uncertainty should be view-dependent. Suppose a point in 3D space can be viewed from different viewpoints, view-dependent modeling can predict different uncertainty scores for prior from different views such that the views with inaccurate prior can be filtered and the rest can be preserved.

More qualitative results are shown in Fig. 11. It can be observed that predicting the uncertainty score by the color network can reconstruct the scene more robustly than by the geometry network. The reason is that a region can be viewed from multiple viewpoints and only the viewpoints with inaccurate prior need to be filtered. Such that the uncertainty of region should be view dependent.

4.4.4 Efficiency of the SDF to Density Transformation

To evaluate the efficiency of the proposed SDF to density transformation, we conduct the experiments by only applying this sub-module to the baseline on the DTU dataset [60].

The results are shown in Table. 8. Our method achieves better performance than other baselines [4], [39]. This proves the efficiency of the proposed SDF to density transformation approach. The qualitative results are shown in Fig. 12, it can be observed that our method can reconstruct the detailed curved regions better than other baselines. These

TABLE 8: The comparison experiments on the DTU dataset [60]. We only modify the SDF to density transformation approach to the TUVr [39] and Ours.

Method	24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	Mean
NeuS [4]	1.00	1.37	0.93	0.43	1.10	0.65	0.57	1.48	1.09	0.83	0.52	1.2	0.35	0.49	0.54	0.84
TUVr [39]	0.83	1.06	0.57	0.40	1.00	0.62	0.62	1.41	1.32	0.94	0.57	1.07	0.35	0.49	0.51	0.79
Ours	0.71	0.94	0.46	0.39	1.05	0.61	0.59	1.49	1.20	0.88	0.61	1.05	0.34	0.47	0.49	0.75




Fig. 12: We conduct ablation studies on the DTU dataset [60] to prove the efficiency of the proposed SDF to density transformation. For our model and TUVr [39], we do not apply the geometry prior for reconstruction.

experiments prove that our proposed curvature-aware SDF to density transformation is efficient for the detailed and curved regions. Considering the curvature of the surface in the transformation of the SDF to the density can benefit the reconstruction because the bias is reduced.

5 CONCLUSION

We introduce DebSDF, which improved the detail and quality of indoor 3D reconstructions by localizing uncertainty regions and introducing a bias-aware SDF-to-density transformation for volume rendering of SDF. Based on the observation that a prior is correct if it is consistent with other priors, we propose an uncertainty modeling approach that effectively identifies large error regions in monocular geometric priors, which usually correspond to fine-detailed regions in the indoor scene. Accordingly, we selectively filter out geometry priors in these regions to avoid their potentially negative effect. We also assign a higher sampling probability to these regions and apply adaptive smooth regularization, further improving reconstruction quality. Furthermore, we found that the volume rendering technique of neural implicit surface used in previous work has a strong bias in

eliminating fine-detailed surfaces. Consequently, we propose a progressively growing, bias-aware SDF-to-density transformation method to reduce the impact of these biases, enhancing the reconstruction of thin, detailed structures in indoor environments. Our DebSDF demonstrates improved reconstruction compared to previous work, evidenced by experiments across five challenging datasets.

Limitation While DebSDF improves the reconstruction quality significantly over previous work, several limitations still exist. First, our method depends on the quality of monocular priors and could potentially benefit from future developments of monocular priors. Second, we use images and monocular priors with resolution 384×384 because Omnidata [49] model is trained on the images with this resolution, and their performances when using images with high resolution are poor. An alternative way is to resize the monocular priors to a high resolution, which we empirically found ineffective as it generates wrong priors to many pixels. More advanced methods of using high-resolution priors [1] are left as our future work.


REFERENCES

- [1] Z. Yu, S. Peng, M. Niemeyer, T. Sattler, and A. Geiger, "Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction," in *Advances in Neural Information Processing Systems*, 1, 2, 3, 4, 5, 10, 11, 12, 14.
- [2] 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. 1, 2, 3, 4, 5, 7, 10
- [3] M. Kazhdan and H. Hoppe, "Screened poisson surface reconstruction," *ACM Transactions on Graphics (ToG)*, vol. 32, no. 3, pp. 1–13, 2013. 1, 10
- [4] 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. 1, 3, 6, 7, 10, 13, 14
- [5] M. Oechsle, S. Peng, and A. Geiger, "Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5589–5599. 1, 3, 7, 10
- [6] 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. 1, 3, 4, 10
- [7] M. Niemeyer, L. Mescheder, M. Oechsle, and A. Geiger, "Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 3504–3515. 1, 6
- [8] J. M. Coughlan and A. L. Yuille, "Manhattan world: Compass direction from a single image by bayesian inference," in *Proceedings of the seventh IEEE international conference on computer vision*, vol. 2. IEEE, 1999, pp. 941–947. 1, 3
- [9] 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. 1, 3, 10
- [10] 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 *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*. Springer, 2022, pp. 139–155. 1, 3, 4, 10, 11, 12
- [11] 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. 2, 9, 10, 12, 13
- [12] A. Handa, T. Whelan, J. McDonald, and A. J. Davison, "A benchmark for rgb-d visual odometry, 3d reconstruction and slam," in *2014 IEEE international conference on Robotics and automation (ICRA)*. IEEE, 2014, pp. 1524–1531. 2, 10, 12
- [13] 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. 2, 10
- [14] A. Knapitsch, J. Park, Q.-Y. Zhou, and V. Koltun, "Tanks and temples: Benchmarking large-scale scene reconstruction," *ACM Transactions on Graphics*, vol. 36, no. 4, 2017. 2, 10, 11
- [15] 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. 2, 3, 10
- [16] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon, "Kinectfusion: Real-time dense surface mapping and tracking," in *2011 10th IEEE international symposium on mixed and augmented reality*. Ieee, 2011, pp. 127–136. 2
- [17] P. Merrell, A. Akbarzadeh, L. Wang, P. Mordohai, J.-M. Frahm, R. Yang, D. Nistér, and M. Pollefeys, "Real-time visibility-based fusion of depth maps," in *2007 IEEE 11th International Conference on Computer Vision*. Ieee, 2007, pp. 1–8. 2
- [18] 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. 2
- [19] M. Bleyer, C. Rhemann, and C. Rother, "Patchmatch stereo-stereo matching with slanted support windows," in *Bmvc*, vol. 11, 2011, pp. 1–11. 2, 3
- [20] D. Paschalidou, O. Ulusoy, C. Schmitt, L. Van Gool, and A. Geiger, "Raynet: Learning volumetric 3d reconstruction with ray potentials," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3897–3906. 3
- [21] A. O. Ulusoy, A. Geiger, and M. J. Black, "Towards probabilistic volumetric reconstruction using ray potentials," in *2015 International Conference on 3D Vision*. IEEE, 2015, pp. 10–18. 3
- [22] S. M. Seitz and C. R. Dyer, "Photorealistic scene reconstruction by voxel coloring," *International journal of computer vision*, vol. 35, pp. 151–173, 1999. 3
- [23] Z. Yu and S. Gao, "Fast-mvsnet: Sparse-to-dense multi-view stereo with learned propagation and gauss-newton refinement," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 1949–1958. 3
- [24] Y. Yao, Z. Luo, S. Li, T. Fang, and L. Quan, "Mvsnet: Depth inference for unstructured multi-view stereo," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 767–783. 3
- [25] P.-H. Huang, K. Matzen, J. Kopf, N. Ahuja, and J.-B. Huang, "Deepmvs: Learning multi-view stereopsis," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2821–2830. 3
- [26] Y. Yao, Z. Luo, S. Li, T. Shen, T. Fang, and L. Quan, "Recurrent mvsnet for high-resolution multi-view stereo depth inference," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 5525–5534. 3
- [27] X. Gu, Z. Fan, S. Zhu, Z. Dai, F. Tan, and P. Tan, "Cascade cost volume for high-resolution multi-view stereo and stereo matching," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 2495–2504. 3
- [28] J. Liao, Y. Fu, Q. Yan, F. Luo, and C. Xiao, "Adaptive depth estimation for pyramid multi-view stereo," *Computers & Graphics*, vol. 97, pp. 268–278, 2021. 3
- [29] S. Cheng, Z. Xu, S. Zhu, Z. Li, L. E. Li, R. Ramamoorthi, and H. Su, "Deep stereo using adaptive thin volume representation with uncertainty awareness," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 2524–2534. 3
- [30] F. Wang, S. Galliani, C. Vogel, P. Speciale, and M. Pollefeys, "Patchmatchnet: Learned multi-view patchmatch stereo," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 14 194–14 203. 3
- [31] Y. Ding, W. Yuan, Q. Zhu, H. Zhang, X. Liu, Y. Wang, and X. Liu, "Transmvsnet: Global context-aware multi-view stereo network with transformers," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 8585–8594. 3
- [32] H. Xu, Z. Zhou, Y. Qiao, W. Kang, and Q. Wu, "Self-supervised multi-view stereo via effective co-segmentation and data-augmentation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 4, 2021, pp. 3030–3038. 3
- [33] 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 *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16*. Springer, 2020, pp. 414–431. 3
- [34] 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. 3
- [35] A. Bozic, P. Palafox, J. Thies, A. Dai, and M. Nießner, "Transformerfusion: Monocular rgb scene reconstruction using transformers," *Advances in Neural Information Processing Systems*, vol. 34, pp. 1403–1414, 2021. 3
- [36] Y. Wei, S. Liu, Y. Rao, W. Zhao, J. Lu, and J. Zhou, "Nerfingmvs: Guided optimization of neural radiance fields for indoor multi-view stereo," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5610–5619. 3, 4
- [37] 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. 3
- [38] J. T. Barron, B. Mildenhall, M. Tancik, P. Hedman, R. Martin-Brualla, and P. P. Srinivasan, "Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5855–5864. 3
- [39] Y. Zhang, Z. Hu, H. Wu, M. Zhao, L. Li, Z. Zou, and C. Fan, "Towards unbiased volume rendering of neural implicit surfaces with geometry priors," in *Proceedings of the IEEE/CVF Conference*


- on *Computer Vision and Pattern Recognition*, 2023, pp. 4359–4368. **3, 6, 7, 8, 10, 13, 14**
- [40] Z. Li, T. Müller, A. Evans, R. H. Taylor, M. Unberath, M.-Y. Liu, and C.-H. Lin, “Neuralangelo: High-fidelity neural surface reconstruction,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 8456–8465. **3**
- [41] M. Tancik, E. Weber, E. Ng, R. Li, B. Yi, J. Kerr, T. Wang, A. Kristoffersen, J. Austin, K. Salahi, A. Ahuja, D. McAllister, and A. Kanazawa, “Nerfstudio: A modular framework for neural radiance field development,” in *ACM SIGGRAPH 2023 Conference Proceedings*, ser. SIGGRAPH ’23, 2023. **3**
- [42] Z. Yu, A. Chen, B. Antic, S. Peng, A. Bhattacharyya, M. Niemeyer, S. Tang, T. Sattler, and A. Geiger, “Sdfstudio: A unified framework for surface reconstruction,” 2022. [Online]. Available: <https://github.com/autonomousvision/sdfstudio> **3**
- [43] Y. Wang, I. Skorokhodov, and P. Wonka, “Hf-neus: Improved surface reconstruction using high-frequency details,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 1966–1978, 2022. **3, 4, 6, 7**
- [44] M. Niemeyer, J. T. Barron, B. Mildenhall, M. S. Sajjadi, A. Geiger, and N. Radwan, “Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5480–5490. **3**
- [45] S. Zhi, T. Laidlow, S. Leutenegger, and A. J. Davison, “In-place scene labelling and understanding with implicit scene representation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 838–15 847. **3**
- [46] A. Jain, M. Tancik, and P. Abbeel, “Putting nerf on a diet: Semantically consistent few-shot view synthesis,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5885–5894. **3**
- [47] K. Deng, A. Liu, J.-Y. Zhu, and D. Ramanan, “Depth-supervised nerf: Fewer views and faster training for free,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 12 882–12 891. **3**
- [48] 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. **3, 4, 5**
- [49] 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. **3, 4, 6, 10, 14**
- [50] B. Ye, S. Liu, X. Li, and M.-H. Yang, “Self-supervised super-plane for neural 3d reconstruction,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 21 415–21 424. **3**
- [51] Z. Liang, Z. Huang, C. Ding, and K. Jia, “Helixsurf: A robust and efficient neural implicit surface learning of indoor scenes with iterative intertwined regularization,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 13 165–13 174. **3, 10**
- [52] W. Dong, C. Choy, C. Loop, O. Litany, Y. Zhu, and A. Anandkumar, “Fast monocular scene reconstruction with global-sparse local-dense grids,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 4263–4272. **3**
- [53] D. Eigen, C. Puhrsch, and R. Fergus, “Depth map prediction from a single image using a multi-scale deep network,” *Advances in neural information processing systems*, vol. 27, 2014. **4**
- [54] G. Bae, I. Budvytis, and R. Cipolla, “Estimating and exploiting the aleatoric uncertainty in surface normal estimation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 13 137–13 146. **4, 5**
- [55] J. Shen, A. Agudo, F. Moreno-Noguer, and A. Ruiz, “Conditional-flow nerf: Accurate 3d modelling with reliable uncertainty quantification,” in *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part III*. Springer, 2022, pp. 540–557. **5**
- [56] X. Pan, Z. Lai, S. Song, and G. Huang, “Activenerf: Learning where to see with uncertainty estimation,” in *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII*. Springer, 2022, pp. 230–246. **5**
- [57] A. Gropp, L. Yariv, N. Haim, M. Atzmon, and Y. Lipman, “Implicit geometric regularization for learning shapes,” in *Proceedings of the 37th International Conference on Machine Learning*, 2020, pp. 3789–3799. **6**
- [58] X. Z. Wujun Che, Jean-Claude Paul, “Lines of curvature and umbilical points for implicit surfaces,” *Computer Aided Geometric Design*, pp. 395–409, 2007. **8**
- [59] T. Novello, G. Schardong, L. Schirmer, V. da Silva, H. Lopes, and L. Velho, “Exploring differential geometry in neural implicits,” *Computers & Graphics*, vol. 108, pp. 49–60, 2022. **8**
- [60] H. Aanaes, R. R. Jensen, G. Vogiatzis, E. Tola, and A. B. Dahl, “Large-scale data for multiple-view stereopsis,” *International Journal of Computer Vision*, vol. 120, pp. 153–168, 2016. **10, 13, 14**
- [61] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, “Pytorch: An imperative style, high-performance deep learning library,” *Advances in neural information processing systems*, vol. 32, 2019. **10**
- [62] T. Müller, A. Evans, C. Schied, and A. Keller, “Instant neural graphics primitives with a multiresolution hash encoding,” *ACM Transactions on Graphics (ToG)*, vol. 41, no. 4, pp. 1–15, 2022. **11**
- [63] W. E. Lorensen and H. E. Cline, “Marching cubes: A high resolution 3d surface construction algorithm,” *ACM siggraph computer graphics*, vol. 21, no. 4, pp. 163–169, 1987. **11**



Yuting Xiao received a bachelor's degree from Hunan University, Changsha, China, in 2018. He is currently pursuing a Ph.D. degree at ShanghaiTech University under the supervision of Prof. Shenghua Gao. His research interests include 3D reconstruction, image segmentation, and anomaly detection.



Jingwei Xu is currently a master student of ShanghaiTech University and will formally begin pursuing his M.S. in Computer Science from ShanghaiTech University under the supervision of Prof. Shenghua Gao in the near future. He has just achieved his bachelor's degree recently from Shanghai University.



Zehao Yu is a Ph.D. student with the Autonomous Vision Group led by Prof. Andreas Geiger, at the University of Tübingen, Germany. He obtained his M.S. in Computer Science in 2021 from ShanghaiTech University under the supervision of Prof. Shenghua Gao and his B.Sc. degree in Software Engineering from Xi'an University.



Shenghua Gao is an associate professor, PI in ShanghaiTech University, China. He received the B.E. degree from the University of Science and Technology of China in 2008 (outstanding graduates), and received the Ph.D. degree from the Nanyang Technological University in 2012. From Jun 2012 to Jul 2014, he worked as a postdoctoral fellow at Advanced Digital Sciences Center, in Singapore. His research interests include computer vision and machine learning.