

NLOS-NeuS: Non-line-of-sight Neural Implicit Surface

Yuki Fujimura¹

Takahiro Kushida²

Takuya Funatomi¹

Yasuhiro Mukaigawa¹

¹Nara Institute of Science and Technology

²Ritsumeikan University

{fujimura.yuki, funatomi, mukaigawa}@is.naist.jp, tkushida@fc.ritsumeikan.ac.jp

Abstract

Non-line-of-sight (NLOS) imaging is conducted to infer invisible scenes from indirect light on visible objects. The neural transient field (NeTF) was proposed for representing scenes as neural radiance fields in NLOS scenes. We propose NLOS neural implicit surface (NLOS-NeuS), which extends the NeTF to neural implicit surfaces with a signed distance function (SDF) for reconstructing three-dimensional surfaces in NLOS scenes. We introduce two constraints as loss functions for correctly learning an SDF to avoid non-zero level-set surfaces. We also introduce a lower bound constraint of an SDF based on the geometry of the first-returning photons. The experimental results indicate that these constraints are essential for learning a correct SDF in NLOS scenes. Compared with previous methods with discretized representation, NLOS-NeuS with the neural continuous representation enables us to reconstruct smooth surfaces while preserving fine details in NLOS scenes. To the best of our knowledge, this is the first study on neural implicit surfaces with volume rendering in NLOS scenes. Project page: <https://yfujimura.github.io/nlos-neus/>

1. Introduction

Many computer-vision applications are being integrated into society. However, most are intended for visible scenes from cameras. Non-line-of-sight (NLOS) imaging [16, 38] is used to infer invisible scenes occluded from the cameras. Figure 1(a) shows a typical confocal NLOS setup [31], in which an ultra-fast light source and time-resolved detector are collocated and the detector can only see a diffuse wall. Pulsed light from the light source is emitted into the visible wall, and reflected light on the wall reaches an NLOS scene. Reflected light in the NLOS scene then bounces back to the detector via the relay wall. A time-resolved detectors for such NLOS imaging is a single photon avalanche diode (SPAD) [5] (top in Fig 1(b)). A SPAD is capable of time-correlated single photon counting with

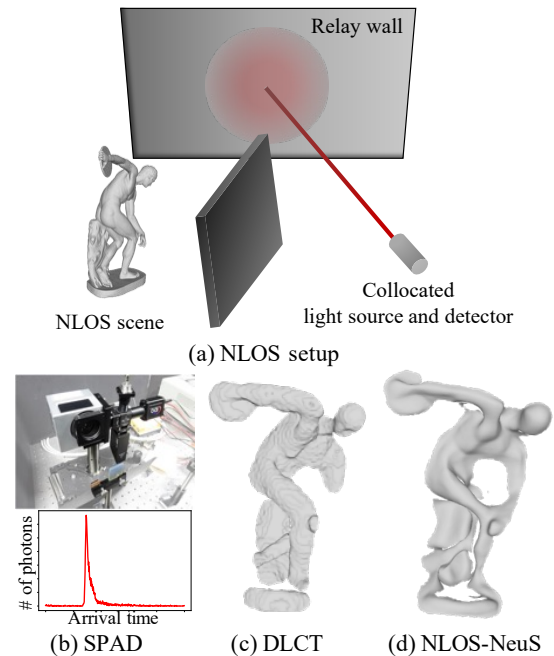


Figure 1. (a) Confocal NLOS setup, (b) SPAD system and measured transients, (c) NLOS surface reconstruction with DLCT [45] and (d) with our NLOS-NeuS. Neural implicit surface representation can enhance quality of reconstructed surfaces in NLOS scenes.

pico-second resolution [40]. Measured data with such a time-resolved detector are called transients, in which the number of counted photons or intensity is recorded at each arrival time (bottom in Fig. 1(b)). One approach for NLOS imaging is to use the transients on the relay wall as input. We aim to reconstruct three-dimensional (3D) surfaces in NLOS scenes from the transients, similar to previous studies [13, 34, 36, 37, 42, 45].

We extend the neural transient field (NeTF) [35] for neural implicit surface representation in NLOS scenes. The neural radiance field (NeRF) [26] is a powerful paradigm for scene representation, in which density and view-dependent color at each scene position are continuously pa-

parameterized with a multi-layer perceptron (MLP), which is optimized by reconstructing input multi-view images with volume rendering. The NeTF is the extended version of the NeRF for transient measurement in an NLOS setup. However, the NeTF (and also the NeRF) does not support geometric representation, which is necessary for accurate 3D reconstruction. We incorporate the framework of a neural implicit surface with volume rendering such as NeuS [39] and VolSDF [43] into the NeTF for NLOS neural implicit surface reconstruction, i.e., *NLOS-NeuS*.

We argue that a simple extension of the NeTF to neural implicit surface representation with a signed distance function (SDF) causes a non-zero level-set surface due to an under-constrained NLOS setup. We introduce effective training losses to avoid a non-zero level-set surface. Figure 1(d) shows an example of the surface reconstruction of an occluded object with NLOS-NeuS. Compared with the results of the state-of-the-art NLOS surface-reconstruction method [45] (Fig. 1(c)), the fine details of the target object can be reconstructed because of the nature of its continuous representation. To the best of our knowledge, this is the first study on a neural implicit surface with volume rendering in NLOS scenes.

2. Related work

Non-line-of-sight imaging NLOS imaging is attracting much more attention with the development of computational imaging devices [16, 38]. While the typical inputs of NLOS-imaging methods are transients on a diffuse relay wall, output scene representations differ depending on applications and methodology.

A voxel grid is one of the most commonly used representations [1, 19, 20, 31, 38]. For example, with back-projection-based methods [1, 17, 38], the measured intensity is back-projected to each voxel to estimate the probabilities of object existence. O’Toole et al. [31] proposed the light-corn transform (LCT), with which a closed-form solution is derived under volumetric albedo representation. Lindell et al. [19] formulated the NLOS problem as f-k migration in seismology, where electromagnetic radiation at each scene-grid point is recovered with the Fast Fourier Transform. However, such discrete representations are limited for representing scenes with fine details due to memory cost. Shen et al. [35] proposed the NeTF, with which volumetric density and reflectance are implicitly modeled by a continuous MLP with arbitrary resolution. Mu et al. [27] extended the NeTF to feedforward inference, which enables fast NLOS imaging with a non-confocal setup.

Although these methods model object existence or volumetric albedo in NLOS scenes, accurate 3D reconstruction requires explicit geometric representations, e.g., some methods incorporate surface normals into voxel representations [45], directly estimate a point cloud [42], or optimize

object surface with a differentiable renderer [13, 34, 37]. Our method extends continuous volumetric representation [35] with an SDF for surface reconstruction with arbitrary resolution.

Neural implicit surface Park et al. [32] and Mescheder et al. [25] respectively proposed the DeepSDF and Occupancy Networks, with which a scene is represented using a SDF and occupancy field parameterized with an MLP. In contrast to the traditional discretized representations such as voxels, these implicit representations are memory-efficient, and an object surface can be extracted as level-set at any resolution, which enables dense surface reconstruction from a coarse voxel grid or sparse point cloud with feedforward inference [7, 33] or test-time optimization [2, 6, 23, 41].

Other research directions include neural implicit surfaces from multi-view 2D images, where the implicit function is optimized by minimizing reconstruction loss between the input and rendered images. The key to optimization is how to connect the implicit function and surface rendering in a differentiable manner [28, 44]. In contrast to surface-rendering-based methods, Yariv et al. [43] and Wang et al. [39] respectively proposed VolSDF and NeuS, with which images are rendered with volume rendering similar to the NeRF [26]. Differing from the surface-rendering approach, images are rendered with multiple points on a ray with α -compositions, which enables back-propagation from not only the surface points but also points far from the surface. NLOS-NeuS uses a neural implicit surface with volume rendering for NLOS surface reconstruction.

3. Method

This section describes the key points for NLOS-NeuS to correctly learn an SDF in an NLOS scene. We first explain our problem setting and review the NeTF as preliminary for neural scene representation in an NLOS scene. We then describe the difficulty in learning an SDF in the under-constrained NLOS setup followed by additional constraints and loss functions for learning an SDF. At the end of this section, we introduce background rendering in NLOS scenes.

3.1. Problem setting

Our goal is to learn two MLPs, $d : \mathbb{R}^3 \rightarrow \mathbb{R}$ and $\rho : \mathbb{R}^6 \rightarrow \mathbb{R}$ from transients. The input of d is a scene position $\mathbf{p} = (x, y, z) \in \mathbb{R}^3$ in an NLOS scene to estimate a signed distance, and the inputs of ρ are \mathbf{p} and direction \mathbf{v} from \mathbf{p} to the position of a visible relay wall $\mathbf{p}' = (x', y', z') \in \mathbb{R}^3$ to estimate view-dependent reflectance. The parameters of these two MLPs are optimized by reconstructing measured transients on the relay wall with volume rendering.

The transient at each position on the relay wall is represented with a 1D vector $\tau(\mathbf{p}') \in \mathbb{R}^B$, where B is the

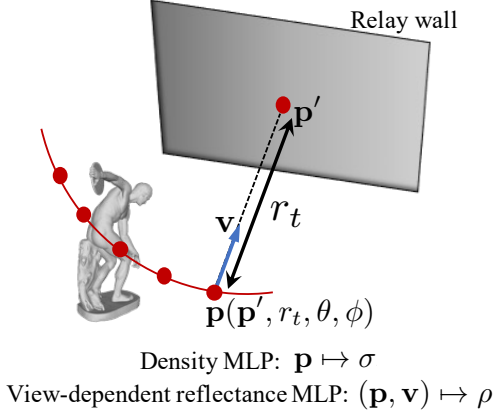


Figure 2. Spherical sampling in NeTF [35] for rendering transients. Scan sphere with radius r_t corresponds to transient at time bin t .

number of time bins. For example, in SPAD measurement, $\tau(\mathbf{p}', t)$ is the number of counted photons, the arrival time of which is t .

3.2. Preliminary

We first review the NeTF [35] for NLOS scene representation. The NeTF also learns two MLPs; one is $\sigma : \mathbb{R}^3 \rightarrow \mathbb{R}$ taking as input a position (x, y, z) to estimate density, and the other is ρ , which is essentially the same with our MLP to estimate view-dependent reflectance¹. Differing from the NeRF [26], where scene points are sampled on a ray for volume rendering, the NeTF spherically samples points from the relay wall for rendering the transients as shown in Fig. 2:

$$\mathbf{p}(\mathbf{p}', r_t, \theta, \phi) = \begin{bmatrix} r_t \sin \theta \cos \phi + x' \\ r_t \sin \theta \sin \phi + y' \\ r_t \cos \theta + z' \end{bmatrix}, \quad (1)$$

where r_t is a radius corresponding to a transient time bin t , and θ and ϕ are elevation and azimuth angles. A transient $\tau(\mathbf{p}', t)$ is computed as the sum of reflected intensities at sampled points on the scan sphere:

$$\tau(\mathbf{p}', t) = \iint A(r_t, \theta) T(\mathbf{p}, \mathbf{v}) \sigma(\mathbf{p}) \rho(\mathbf{p}, \mathbf{v}) d\theta d\phi, \quad (2)$$

where $A(r_t, \theta) = \sin \theta / r_t^2$ is an attenuation factor, and $T(\mathbf{p}, \mathbf{v})$ ² is the transparency between the object and relay wall. For simplicity, the arguments of \mathbf{p} and \mathbf{v} are omitted, i.e., $\mathbf{p}(\mathbf{p}', r_t, \theta, \phi)$ and $\mathbf{v}(\mathbf{p}(\mathbf{p}', r_t, \theta, \phi), \mathbf{p}')$.

The parameters of σ and ρ are optimized by minimizing the error between measured and rendered transients as

¹In fact, ρ is a 5D function in the NeTF, i.e., the inputs are (x, y, z) and (θ, ϕ) .

²This transparency is omitted in the implementation of the NeTF.

follows:

$$\mathcal{L}_\tau = \frac{1}{MB} \sum_{\mathbf{p}', t} (\tau_m(\mathbf{p}', t) - \tau(\mathbf{p}', t))^2, \quad (3)$$

where τ_m is the measured transient and M is the number of measured positions on the relay wall.

3.3. NLOS-NeuS

Following previous studies [39, 43], NLOS-NeuS uses volume rendering with an SDF. The MLP d first takes \mathbf{p} as input to estimate a signed distance, and ρ takes \mathbf{p} and \mathbf{v} to estimate view-dependent reflectance. The estimated signed distance is then converted to a density for volume rendering. We use the transformation proposed in StyleSDF [30] for computational efficiency with one learnable parameter and without gradient evaluation:

$$\sigma(\mathbf{p}) = \frac{1}{\alpha} \text{Sigmoid}\left(-\frac{d(\mathbf{p})}{\alpha}\right), \quad (4)$$

where $\text{Sigmoid}(\cdot)$ is a sigmoid function and $\alpha > 0$ is a learnable parameter that controls the tightness of the density around the object surface. After obtaining the density, we can render a transient by using Eq. (2). In the implementation, we use the following discrete counterparts [24, 26, 39]:

$$\tau(\mathbf{p}', t) = \sum_{\theta, \phi} A(r_t, \theta) w(\mathbf{p}, \mathbf{v}) \rho(\mathbf{p}, \mathbf{v}) \Delta\theta \Delta\phi, \quad (5)$$

$$w(\mathbf{p}, \mathbf{v}) = \sum_{s=\ell_{min}}^t T_s \left(1 - \exp\{-\sigma(\mathbf{p}(\mathbf{p}', r_s, \theta, \phi)) \Delta r_s\}\right), \quad (6)$$

$$T_s = \exp\left\{-\sum_{u=\ell_{min}}^{s-1} \sigma(\mathbf{p}(\mathbf{p}', r_u, \theta, \phi)) \Delta r_u\right\}, \quad (7)$$

where $w(\mathbf{p}, \mathbf{v})$ corresponds to the blending weights of volumetric albedos on the ray from the relay wall.

3.4. Learning zero level-set surface in NLOS setup

Many methods for learning neural implicit surfaces have been proposed under a multi-view setup [28, 29, 39, 43, 44]. One of the differences between these methods and our NLOS-NeuS is the configuration of the target scene. In a multi-view setup, the target object is surrounded by multiple cameras. In an NLOS setup, however, only one side of the target object is observed from the relay wall. We found that such an under-constrained setup causes a non-zero level-set surface in an SDF, as shown in Fig. 3(a). The back side of the object cannot be observed, which leads to the highest weight at the point with non-zero signed distance on the ray. Figure 3(b) shows the estimated signed distance for the synthetic bunny scene with the simple extension of the

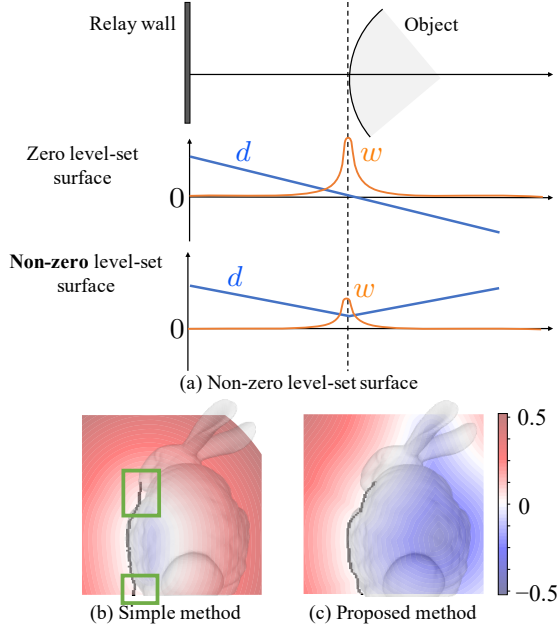


Figure 3. (a) Non-zero level-set surface. Under-constrained NLOS setup causes non-zero level-set surface because back side of object cannot be observed from relay wall. Point with non-zero signed distance incorrectly has highest weight on ray from relay wall. (b) Learned SDF of synthetic bunny scene, where points with highest weights (black curve) do not coincide with zero level-set in green boxes. (c) Our method enabled us to extract plausible zero level-set surfaces.

NeTF. The black curve indicates an extracted object surface as the points with the highest weight along the ray from the relay wall and does not coincide with the zero level-set in the green boxes. Figure 9(a) is the estimated depth from this SDF with the sphere tracing algorithm [12], where most parts of the bunny are missing.

To avoid such a non-zero level-set surface, the following two requirements should be satisfied. (1) At the object surface, the signed distance values should be zero. To achieve this, we have to detect the object surface and force explicit geometric supervision. (2) Points with non-zero signed distance should not contribute to the object surface. To achieve this, α in Eq. (4) decreases correctly during training. As explained in Sec. 3.3, $\alpha \rightarrow 0$ means a perfectly sharp surface, i.e., only points with zero signed distance contributes to object surfaces. We introduce two constraints as training losses for satisfying these two requirements, which enabled us to estimate a plausible SDF, as shown in Fig. 3(c).

Zero signed-distance self-supervision We supervise the MLP to have zero signed distances at the object surface. Recent studies have shown that such explicit geometric supervision enhances the quality of the learned SDF [3, 9]. Differing from these studies in which geometric information

was obtained with additional sensors, we have no explicit geometric information, thus automatically detect points to be supervised during training.

To detect surface points, we compute the probabilistic density function (PDF) as normalized $w(\mathbf{p}, \mathbf{v})\rho(\mathbf{p}, \mathbf{v})$ on each scan sphere during training and sample N_z points on the sphere on the basis of the PDF. These points are likely to be located on the object surface; thus, we force these points to have a zero signed distance with the following loss function:

$$\mathcal{L}_z = \frac{1}{MBN_z} \sum_{t, \mathbf{p}'} \sum_{n=1}^{N_z} m(\mathbf{p}', t) |d(\mathbf{p}(t, \theta_n, \phi_n))|, \quad (8)$$

where $m(\mathbf{p}', t) \in \{0, 1\}$ is a mask; $m(\mathbf{p}', t) = 1$ means that object points exist on the sphere with a radius r_t centered at \mathbf{p}' . This mask can be computed easily by thresholding the measured transient τ_m . If we model background effects (Sec. 3.6), we compute the mask with an object component $\tau_m - \xi\tau_b$ during training.

Constraint on volume-rendering weight Wang et al. [39] mentioned that a learnable parameter connecting signed distance and density (α in Eq. (4) in our case) converges to zero during training. However, we found that an additional constraint is necessary to correctly decrease the parameter for the NLOS setup.

Before introducing the constraint, we discuss a mask loss proposed in multi-view settings [39, 44]. In NeuS [39], the input images are masked, and the masks are used for an additional loss, where the sum of the volume rendering weights (Eq. (6) in our case) on each camera ray is forced to be equal to the value at the corresponding pixel of the mask, i.e., the mask loss is defined as

$$\mathcal{L}_{mask} = E(m_k, \hat{\delta}_k), \quad (9)$$

where m_k is the value of the mask at the k -th pixel, $\hat{\delta}_k$ is the accumulated weight, and E is an error metric such as the binary cross entropy. We provide the following observation.

Proposition 1. *Minimizing the mask loss (Eq. (9)) leads to the convergence of α (in Eq. (4)) to 0.*

We provide the proof in the supplementary material. Intuitively, all densities in the empty space should be 0, which results in the convergence of α to 0.

Therefore, the mask loss is effective for learning α , while the difficulty in the NLOS setup is that we cannot obtain an object mask. Instead of using such explicit mask loss, we use the following loss:

$$\mathcal{L}_{en} = \frac{1}{M|\theta||\phi|} \sum_{\mathbf{p}', \theta, \phi} -\hat{\delta} \log_2 \hat{\delta} - (1 - \hat{\delta}) \log_2 (1 - \hat{\delta}), \quad (10)$$




Figure 4. Plots of α during training. For learning sharp surface, α in Eq. (4) should decrease during training. This plot shows that \mathcal{L}_{en} can make α converge to very small value (orange and red plots)

where $\hat{o} = \sum_{t=t_{min}}^{t_{max}} w(\mathbf{p}, \mathbf{v})$ is the accumulated weight for each view line from the relay wall, and $|\cdot|$ is the number of sampled variables. This loss function is similar to the beta-distribution constraint of accumulated opacity used in the previous studies [4, 21, 27] where accumulated opacity from a camera should be 0 or 1. We use entropy as this loss function that bounds the loss value between 0 and 1.

Figure 4 shows the plots of α during training for the synthetic bunny scene, where α diverges without both \mathcal{L}_z and \mathcal{L}_{en} (blue plot). Although α does not diverge if we use \mathcal{L}_z (green plot), it begins to increase in the early stage and remains relatively large at the end of the training. On the other hand, α successfully decreases and converges to a very small value with \mathcal{L}_{en} (orange and red plots).

3.5. Training loss

In summary, we use the following loss function for training NLOS-NeuS:

$$\mathcal{L} = \lambda_\tau \mathcal{L}_\tau + \lambda_{ei} \mathcal{L}_{ei} + \lambda_z \mathcal{L}_z + \lambda_{en} \mathcal{L}_{en} + \lambda_f \mathcal{L}_f, \quad (11)$$

where \mathcal{L}_τ , \mathcal{L}_z , and \mathcal{L}_{en} are those described in the previous sections. λ_* are hyperparameters to control the contribution of each loss. \mathcal{L}_{ei} is the Eikonal loss [11] for the regularization of the SDF:

$$\mathcal{L}_{ei} = \frac{1}{|\Omega_{nlos}|} \sum_{\mathbf{p} \in \Omega_{nlos}} (\|\nabla d(\mathbf{p})\| - 1)^2, \quad (12)$$

where Ω_{nlos} is the target NLOS space.

We also introduce SDF regularization based on the geometry of first-returning photons [36]. From the traveling distances of first-returning photons, we can apply space carving to the NLOS scene to obtain the free space and rough shape. In the free space, the signed distances to the rough shape form the lower bounds of the true signed dis-

tances. We can leverage these lower bounds as the additional training loss:

$$\mathcal{L}_f = \frac{1}{|\Omega_{free}|} \sum_{\mathbf{p} \in \Omega_{free}} \max(0, b(\mathbf{p}) - d(\mathbf{p})), \quad (13)$$

where Ω_{free} is the free space and $b(\mathbf{p})$ is the lower bound of the signed distance at \mathbf{p} in the free space. This loss is inspired by the SDF lower bound from silhouette [18], while we derive the lower bounds from the geometry of first-returning photons. In the supplementary material, we provide the details of the detection of the first-returning photons and robust space carving algorithm.

3.6. Rendering background

In practical scenes, background effects, such as inter-reflection or reflected light from the floor, are not negligible. Plack et al. [34] proposed the background network, which takes as input a point on a relay wall and temporal bin for estimating components from the background. The scaling parameter between the components from the target object and background is also optimized with an additional constraint. To avoid an SDF fitting the background components, we follow the same approach for modeling the background but simply adjust the scaling parameter with the measured transient. Formally, the final transient τ is the weighted sum of the transients from the object τ_o (same with Eq. (5)) and from the background τ_b as follows:

$$\tau(\mathbf{p}') = \tau_o(\mathbf{p}') + \xi \tau_b(\mathbf{p}'), \quad (14)$$

where ξ is a scaling parameter computed as

$$\xi = \frac{\sum_t \tau_m(\mathbf{p}', t) - \sum_t \tau_o(\mathbf{p}', t)}{\sum_t \tau_b(\mathbf{p}', t)}. \quad (15)$$

The background network is jointly trained with ρ and d . In the supplementary material, we provide examples of the rendered background.

4. Experiments

We evaluated the effectiveness of NLOS-NeuS on publicly available NLOS datasets. We compared it with the NeTF [35] and directional LCT (DLCT) [45], which is a state-of-the-art surface-reconstruction method for NLOS scenes. Note that we retrained the NeTF with the author's code.

4.1. Implementation

Network architecture We use the same networks with that of NeuS [39] with geometric initialization [2] for d and ρ . The inputs of the networks are a scene position and direction to the relay wall with positional encoding [26]. Differing from NeuS, we do not use a surface normal for the




Figure 5. Results from NeTF [35] and NLOS-NeuS on synthetic bunny scene. (a) Rendered directional albedo, (b) depth reconstruction, and (c) surface-normals reconstruction. Neural implicit surface representation achieved geometrically consistent reconstruction.




Figure 6. Surface reconstruction of bunny. (a) Ground-truth mesh, (b) mesh from NeTF [35], (c) mesh from DLCT [45], and (d) mesh from NLOS-NeuS.

input of ρ due to heavy computational cost. For the background network, we use a three-layers MLP with positional encoding for both position and time.

Surface reconstruction After training NLOS-NeuS, we use the sphere tracing algorithm [12] to extract a point cloud instead of using the marching cubes algorithm [22] because an optimized SDF is not correctly closed due to the NLOS setup. Note that we first create an object mask with estimated directional albedo then apply the sphere tracing algorithm to only the object region. After extracting the point cloud with this algorithm, we can obtain a surface normal at each point by evaluating the gradient of d . We then use the Poisson surface reconstruction [15] to obtain the object surface.

Dataset We use two datasets preprocessed in the NeTF [35]. The first one is the ZNLOS dataset [10], which was synthetically created with the transient renderer [14]. The other one is the f-k dataset [19], which is a real capture using gated SPADs avoiding strong direct reflection from a relay wall. The temporal resolution of the SPAD system is approximately 70 pico seconds.

4.2. Results on synthetic data

Figure 5 shows the comparison of the NeTF and NLOS-NeuS on the synthetic bunny scene. Figure 5(a) visual-

Table 1. Quantitative comparison of NLOS-NeuS with NeTF [35] and DLCT [45] on bunny scene. We computed end-point-error for surface normals. Errors dramatically reduced with NLOS-NeuS compared with NeTF, and NLOS-NeuS is also comparable with DLCT.

	Depth [cm]		Normal [cm]	
	RMSE	MAE	RMSE	MAE
NeTF [35]	8.26	3.67	0.81	0.63
DLCT [45]	5.27	1.59	0.40	0.30
NLOS-NeuS	4.63	1.84	0.39	0.30

izes the rendered directional albedos with the NeTF and NLOS-NeuS. The directional albedo of the NeTF was simply extracted as the highest $\sigma\rho$ on a ray from the relay wall. The directional albedo of NLOS-NeuS was generated with volume rendering. In the supplementary material, we present our directional albedos rendered from different views, which is similar to novel-view synthesis with the NeRF [26]. Figure 5(b) shows the results of depth reconstruction. Note that the results are masked with the ground-truth depth map. The neural implicit surface representation with NLOS-NeuS dramatically improved the accuracy of the depth reconstruction. We also show the reconstructed surface normals in (c), where the NeTF normal is computed as the normalized gradient of the density network, demonstrating the geometrically consistent reconstruction of NLOS-NeuS. Table 1 shows the quantitative comparison with the NeTF and DLCT [45] on depth and surface normals. The errors dramatically reduced with NLOS-NeuS compared with the NeTF, and NLOS-NeuS is also comparable with the DLCT, which is the state-of-the-art surface reconstruction in NLOS scenes.

Figure 6 shows the qualitative comparison of reconstructed meshes with the NeTF, DLCT and NLOS-NeuS. Due to the lack of geometric representation, the NeTF mesh contains large errors. On the other hand, NLOS-NeuS enables high quality 3D surface reconstruction. Although the




Figure 7. Results on Lucy and Indonesian with (b) NeTF [35] and (c) NLOS-NeuS. Results contain directional albedos and reconstructed surface. Note that each object is located on floor and measured transient is highly affected by reflected light from floor.




Figure 8. Results of surface reconstruction on real dataset with (a) NeTF [35], (b) DLCT [45], and (c) NLOS-NeuS.

result with the DLCT is comparable with that with NLOS-NeuS on this synthetic data, its discretized representation is limited in real data, as shown in Fig. 1(c).

Figure 7 summarize the comparison with the NeTF for the Lucy and Indonesian scenes. Note that each object is located on a floor and the measured transient is highly affected by the reflected light from the floor. In the supplementary material, we present the learned transients of the floor reflection. We present the directional albedos and meshes of both methods. NLOS-NeuS can reconstruct the complicated structure in the Lucy. Although there are some missing parts in our results of the Indonesian, NLOS-NeuS can obtain geometrically plausible results compared with the NeTF.

4.3. Results on real capture

Figure 8 shows the results with the real data captured using a SPAD. We present the directional albedos and reconstructed meshes for both scenes. The NeTF fails to reconstruct the correct geometry in these scenes. DLCT can reconstruct the object structures, while the quality of the reconstructed geometry is limited due to its discrete representation. NLOS-NeuS, on the other hand, can reconstruct the fine details and smooth surface compared with DLCT because of its continuous representation of the neural implicit surface. These results indicate the effectiveness of NLOS-NeuS for real scenes.

4.4. Ablation study on loss functions

We introduce several losses for training an SDF. Figure 9 shows the comparison between the several patterns of the loss functions on the bunny. As discussed in Sec.




Figure 9. Ablation study of loss functions for learning zero level-set surface. By using both \mathcal{L}_z and \mathcal{L}_{en} , MLP can correctly learn zero level-set surface. We also show effectiveness of \mathcal{L}_f , which improves overall accuracy of reconstructed geometry.



Figure 10. Challenging scenes for SDF. (a) This scene consists of two planes, one of which is partially occluded. SDF is not suitable for such non-closed surface. (b) This scene consists of thin structures, which are also difficult for SDF to correctly represent geometry.

3.4, the simple extension of the NeTF with the neural implicit surface incorrectly estimates non-zero level-set surfaces. When we train the MLP without both \mathcal{L}_z and \mathcal{L}_{en} (a), most parts of the shape are missing due to such non-zero level-set surfaces. When we train the MLP only with \mathcal{L}_z (b) or \mathcal{L}_{en} (c), the results dramatically improve, while the shape near the object boundary cannot be reconstructed. By using both loss functions (e), we can correctly estimate the shape near the object boundary. We also show the effectiveness of \mathcal{L}_f , which improves the overall accuracy of the reconstructed geometry (d,e).

4.5. Limitations

We finally discuss the limitations due to our geometric representation. Figure 10 shows an example of a challenging scene. In Fig. 10(a), two planes are located and one is partially occluded. The figure also shows learned SDF, where the black lines indicate the positions of the ground-truth planes. Although the zero level-set is extracted on the front plane, the SDF around the back plane is not correct. Such a non-closed surface is difficult for an SDF to represent the geometry. Figure 10(b) is another example of a challenging scene for an SDF, where the target object con-

sists of thin structures. One approach for these scenes is to use more flexible geometric representation such as unsigned distance field [8].

5. Conclusion

We proposed NLOS-NeuS, which is an extension of the NeTF for representing 3D surfaces with an SDF in NLOS scenes. To the best of our knowledge, this is the first study on neural implicit surfaces with volume rendering in NLOS scenes. The experimental results indicate that the introduced constraints are essential for learning a correct SDF to avoid non-zero level-set surfaces in an under-constrained NLOS setup, and NLOS-NeuS with continuous representation enables high quality 3D surface reconstruction in NLOS scenes. Future research directions include using other geometric representations and an extension for a non-confocal NLOS setup.

Acknowledgement

This work was supported by the Japan Society for the Promotion of Science KAKENHI Grant Number 21K21317, 20K20629.

References

- [1] Victor Arellano, Diego Gutierrez, and Adrian Jarabo. Fast back-projection for non-line of sight reconstruction. *Opt. Express*, 25(10):11574–11583, 2017. 2
- [2] Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *CVPR*, 2020. 2, 5
- [3] Dejan Azinović, Ricardo Martin-Brualla, Dan B Goldman, Matthias Nießner, and Justus Thies. Neural rgb-d surface reconstruction. In *CVPR*, pages 6290–6301, 2022. 4
- [4] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Deep reflectance volumes: Relightable reconstructions from multi-view photometric images. In *ECCV*, 2020. 5
- [5] Mauro Buttafava, Jessica Zeman, Alberto Tosi, Kevin Elceiri, and Andreas Velten. Non-line-of-sight imaging using a time-gated single photon avalanche diode. *Opt. Express*, 23(16):20997–21011, 2015. 1
- [6] Rohan Chabra, Jan E. Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In *ECCV*, 2020. 2
- [7] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3d shape reconstruction and completion. In *CVPR*, 2020. 2
- [8] Julian Chibane, Aymen Mir, and Gerard Pons-Moll. Neural unsigned distance fields for implicit function learning. In *NeurIPS*, 2020. 8
- [9] Qiancheng Fu, Qingshan Xu, Yew-Soon Ong, and Wenbing Tao. Geo-neus: Geometry-consistent neural implicit surfaces learning for multi-view reconstruction. In *NeurIPS*, 2022. 4
- [10] Miguel Galindo, Julio Marco, Matthew O’Toole, Gordon Wetzstein, Diego Gutierrez, and Adrian Jarabo. A dataset for benchmarking time-resolved non-line-of-sight imaging. In *ICCP*, 2019. 6
- [11] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *ICML*, pages 3569–3579, 2020. 5
- [12] John C. Hart. Sphere tracing: a geometric method for the antialiased ray tracing of implicit surfaces. *The Visual Computer*, 12(10):527–545, 1996. 4, 6
- [13] Julian Iseringhausen and Matthias B. Hullin. Non-line-of-sight reconstruction using efficient transient rendering. *ACM TOG*, 39(1), 2020. 1, 2
- [14] Adrian Jarabo, Julio Marco, Adolfo Muñoz, Raul Buisan, Wojciech Jarosz, and Diego Gutierrez. A framework for transient rendering. *ACM Transactions on Graphics (SIGGRAPH Asia 2014)*, 33(6):177, 2014. 6
- [15] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM TOG*, 32(3):29, 2013. 6
- [16] Ahmed Kirmani, Tyler Hutchison, James Davis, and Ramesh Raskar. Looking around the corner using ultrafast transient imaging. *IJCV*, 95(1):13–28, 2011. 1, 2
- [17] Marco La Manna, Fiona Kine, Eric Breitbach, Jonathan Jackson, Talha Sultan, and Andreas Velten. Error back-projection algorithms for non-line-of-sight imaging. *IEEE TPAMI*, 41(7):1615–1626, 2019. 2
- [18] Chen-Hsuan Lin, Chaoyang Wang, and Simon Lucey. Sdfs: Learning signed distance 3d object reconstruction from static images. In *NeurIPS*, 2020. 5
- [19] David B. Lindell, Gordon Wetzstein, and Matthew O’Toole. Wave-based non-line-of-sight imaging using fast f-k migration. *ACM TOG*, 38(4):116, 2019. 2, 6
- [20] Xiaochun Liu, Ibón Guillén, Marco La Manna, Ji Hyun Nam, Syed Azer Reza, Toan Huu Le, Adrian Jarabo, Diego Gutierrez, and Andreas Velten. Non-line-of-sight imaging using phasor-field virtual wave optics. *Nature*, 572(7771):620–630, 2019. 2
- [21] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM TOG*, 38(4):65:1–65:14, 2019. 5
- [22] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM SIGGRAPH*, 21(4):163–169, 1987. 6
- [23] Baorui Ma, Yu-Shen Liu, and Zhizhong Han. Reconstructing surfaces for sparse point clouds with on-surface priors. In *CVPR*, pages 6315–6325, 2022. 2
- [24] N. Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995. 3
- [25] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 2019. 2
- [26] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 2, 3, 5, 6
- [27] Fangzhou Mu, Sicheng Mo, Jiayong Peng, Xiaochun Liu, Ji Hyun Nam, Siddeshwar Raghavan, Andreas Velten, and Yin Li. Physics to the rescue: Deep non-line-of-sight reconstruction for high-speed imaging. In *ICCP*, 2022. 2, 5
- [28] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *CVPR*, 2020. 2, 3
- [29] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *ICCV*, pages 5589–5599, 2021. 3
- [30] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. Stylesdf: High-resolution 3d-consistent image and geometry generation. In *CVPR*, pages 13503–13513, 2022. 3
- [31] Matthew O’Toole, David B. Lindell, and Gordon Wetzstein. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature*, 555(7696):338–341, 2018. 1, 2
- [32] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. 2
- [33] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *ECCV*, 2020. 2

- [34] Markus Plack, Clara Callenberg, Monika Schneider, and Matthias B. Hullin. Fast differentiable transient rendering for non-line-of-sight reconstruction. In *WACV*, pages 3067–3076, 2023. [1](#), [2](#), [5](#)
- [35] Siyuan Shen, Zi Wang, Ping Liu, Zhengqing Pan, Ruiqian Li, Tian Gao, Shiyang Li, and Jingyi Yu. Non-line-of-sight imaging via neural transient fields. *IEEE TPAMI*, 43(7):2257–2268, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [36] Chia-Yin Tsai, Kiriakos N. Kutulakos, Srinivasa G. Narasimhan, and Aswin C. Sankaranarayanan. The geometry of first-returning photons for non-line-of-sight imaging. In *CVPR*, 2017. [1](#), [5](#)
- [37] Chia-Yin Tsai, Aswin C. Sankaranarayanan, and Ioannis Gkioulekas. Beyond volumetric albedo – a surface optimization framework for non-line-of-sight imaging. In *CVPR*, 2019. [1](#), [2](#)
- [38] Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Mounsi G. Bawendi, and Ramesh Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communications*, 3(1):745, 2012. [1](#), [2](#)
- [39] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *NeurIPS*, 2021. [2](#), [3](#), [4](#), [5](#)
- [40] Ryan Warburton, Constantin Aniculaesei, Matteo Clerici, Yoann Altmann, Genevieve Garipey, Richard McCracken, Derryck Reid, Steve McLaughlin, Marco Petrovich, David Richardson, Robert Henderson, Daniele Faccio, and Jonathan Leach. Observation of laser pulse propagation in optical fibers with a spad camera. In *Imaging and Applied Optics 2016*, 2016. [1](#)
- [41] Francis Williams, Zan Gojcic, Sameh Khamis, Denis Zorin, Joan Bruna, Sanja Fidler, and Or Litany. Neural fields as learnable kernels for 3d reconstruction. In *CVPR*, pages 18500–18510, 2022. [2](#)
- [42] Shumian Xin, Sotiris Nousias, Kiriakos N. Kutulakos, Aswin C. Sankaranarayanan, Srinivasa G. Narasimhan, and Ioannis Gkioulekas. A theory of fermat paths for non-line-of-sight shape reconstruction. In *CVPR*, 2019. [1](#), [2](#)
- [43] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *NeurIPS*, 2021. [2](#), [3](#)
- [44] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. In *NeurIPS*, 2020. [2](#), [3](#), [4](#)
- [45] Sean I. Young, David B. Lindell, Bernd Girod, David Taubman, and Gordon Wetzstein. Non-line-of-sight surface reconstruction using the directional light-cone transform. In *CVPR*, 2020. [1](#), [2](#), [5](#), [6](#), [7](#)