

Few-shot Non-line-of-sight Imaging with Signal-surface Collaborative Regularization

Xintong Liu¹, Jianyu Wang¹, Leping Xiao¹, Xing Fu¹, Lingyun Qiu^{1,2}, and Zuoqiang Shi^{1,2}
¹Tsinghua University
²Yanqi Lake Beijing Institute of Mathematical Sciences and Applications

Abstract

The non-line-of-sight imaging technique aims to reconstruct targets from multiply reflected light. For most existing methods, dense points on the relay surface are raster scanned to obtain high-quality reconstructions, which requires a long acquisition time. In this work, we propose a signal-surface collaborative regularization (SSCR) framework that provides noise-robust reconstructions with a minimal number of measurements. Using Bayesian inference, we design joint regularizations of the estimated signal, the 3D voxel-based representation of the objects, and the 2D surface-based description of the targets. To our best knowledge, this is the first work that combines regularizations in mixed dimensions for hidden targets. Experiments on synthetic and experimental datasets illustrated the efficiency of the proposed method under both confocal and non-confocal settings. We report the reconstruction of the hidden targets with complex geometric structures with only 5×5 confocal measurements from public datasets, indicating an acceleration of the conventional measurement process by a factor of 10,000. Besides, the proposed method enjoys low time and memory complexity with sparse measurements. Our approach has great potential in real-time non-line-of-sight imaging applications such as rescue operations and autonomous driving.

1. Introduction

The non-line-of-sight (NLOS) imaging technique enables reconstructions of targets out of the direct line of sight, which is attractive in various applications such as autonomous driving, remote sensing, rescue operations and medical imaging [1, 5, 6, 10, 15, 16, 19, 21, 26, 33–35, 38–40]. A typical scenario of NLOS imaging is shown in Figure 1. Several points on the visible surface are illuminated by a laser and the back-scattered light from the target is detected to reconstruct the target. The NLOS detection system is confocal if each illumination point is the same with

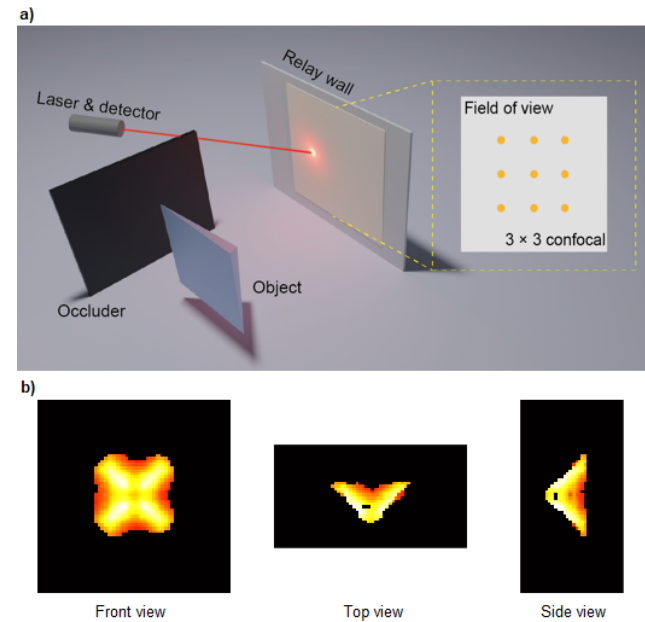


Figure 1. A typical non-line-of-sight imaging scenario. a) The time resolved signals are measured at only 3×3 focal points. b) The three views of the reconstructed target obtained with the proposed SSCR method.

the detection point, and non-confocal otherwise. The time-correlated single-photon counting (TCSPC) technique is applied in the detection process due to the extremely low photon intensity after multiple diffuse reflections. In practice, a single-photon avalanche diode (SPAD) in the Geiger-mode can be used to record the photon events with time-of-flight (TOF) information [3]. The first experimental demonstration of NLOS imaging dates back to 2012, where the targets are reconstructed with the back-projection (BP) method [37]. Extensions of this approach include its fast implementation [2], the filtering technique for reconstruction quality enhancement [17], and weighting factors for noise reduction [11].

A number of efficient methods have been designed

for fast reconstructions. The light cone transform (LCT) method [30] formulates the physical model as a convolution operator, so that the reconstructions can be obtained using the Wiener deconvolution method with the fast Fourier transform. The directional light cone transform (D-LCT) [42] generalizes the LCT and reconstructs the albedo and surface normal simultaneously. The method of frequency wavenumber migration (F-K) [20] formulates the propagation of light using the wave equation, and also provides a fast inversion algorithm with the frequency-domain interpolation technique. Whereas the LCT, D-LCT and F-K methods only work directly in confocal measurement scenarios, the phasor field (PF) method [23,24,32] converts the NLOS imaging scenarios to LOS cases and works for the general non-confocal setting with low computation complexity. For high-quality and noise-robust reconstructions, the signal-object collaborative regularization (SOCR) method can be applied, but brings additional computational cost. In recent years, deep learning-based methods are also introduced to the field of NLOS imaging [7,8,27,43]. Besides, advances in hardware enhance the distance of NLOS detection to kilometers [39], or make it possible to reconstruct targets on the scale of millimeters [38].

Despite these breakthroughs, the trade off between the acquisition time and the imaging quality is inevitable. In the raster scanning mode, the acquisition time is proportional to the number of measurement points with fixed scanning speed. Due to the intrinsic ill-posedness of the NLOS reconstruction problem [22] and heavy measurement noise [11], dense measurements are necessary for high quality reconstructions [20,23,30]. The measurement process may take from seconds to hours, which poses a great challenge for applications such as autonomous driving, where real-time reconstruction of the video stream is needed. The acquisition process can be accelerated by reducing the number of pulses used for each illumination point. In the work [18], the pulse number that record the first returning photon is used to reconstruct the target. Another way to reduce the acquisition time is to design array detectors for non-confocal measurements. For example, the implementation of the phasor field method with SPAD arrays realizes low-latency real-time video imaging of the hidden scenes [28]. A third way to accelerate the NLOS detection process is to reduce the number of measurement points. It is shown that 16×16 confocal measurements are enough to reconstruct the hidden target by incorporating the compressed sensing technique [41].

In this paper, we study the randomness in the photon detection process of NLOS scenarios and propose an imaging method that deals with a very limited number of spatial measurements, which we term the few-shot NLOS detection scenarios. We design joint regularizations of the estimated signal, the 3D voxel-based representation of the ob-




Figure 2. The least-squares solution of the statue with 3×3 confocal measurements [20]. The target cannot be identified even though its simulated signal matches the measurements well (see Fig. 3). Strong regularizations are needed to reconstruct the target. See also Fig. 6 for a comparison.




Figure 3. Comparisons of the measured data and the simulated data of the least-squares solution for the instance of the statue [20]. The measured signals are shown in black. The simulated data of the least-squares solution are shown in red. The shapes of the signals are very close to each other.

jects, and the 2D surface-based description of the targets, which leads to faithful reconstruction results. The main contributions of this work are as follows.

- We propose a signal-surface collaborative regularization (SSCR) framework for few-shot non-line-of-sight reconstructions, which works under both confocal and non-confocal settings.
- We report the reconstruction of the hidden targets with complex geometric structures with only 5×5 confocal measurements from public datasets, indicating an acceleration of the conventional measurement process by a factor of 10,000.

2. Related work

From a mathematical point of view, the non-line-of-sight imaging task belongs to the category of inverse problem. The goal is to reconstruct the surface of the hidden target with the measured signal. The NLOS inverse problem is ill-posed due to the intrinsic structure of the physical model [22] and heavy measurement noise [11]. When the number of measurement points is small, the lack of data leads to rank-deficiency of the measurement matrix, making the reconstruction task even harder [41]. As an example, Fig. 2 shows the reconstruction results of the instance of statue with 3×3 confocal measurements from the Stanford dataset [20] using the traditional least-squares method. The target cannot be identified in the reconstruction. The measured signal and the simulated signal of the least-squares solution are illustrated in Fig. 3, which shows that the shapes of these two signals are very similar. The experimental setup and more quantitative details are provided in the supplement. In such few-shot NLOS detection scenarios, regularization becomes a vital tool for high-quality reconstructions. Here, we briefly review the signal-object collaborative regularization method [25]. The proposed method is closely related to this method.

The signal-object collaborative regularization (SOCR) method [25] represents the hidden target with voxels and considers prior terms for both the reconstructed target and the signal. The joint regularization term concerns sparseness of the reconstruction, non-local self-similarity of the target and the smoothness of the signal. The method employs L_1 regularization, orthogonal optimization and empirical Wiener filter to provide noiseless reconstructions with sharp boundaries, but at the expense of additional computational cost. The SOCR method provides high-quality reconstructions with dense measurements, but may fail in few-shot NLOS detection scenarios. In such cases, the ill-posedness of the NLOS inverse problem is so strong that the method may result in biased estimations of the targets and artifacts in the background.

In this paper, we propose a signal-surface collaborative regularization method that provides faithful reconstructions of the targets in few-shot NLOS detection scenarios. The proposed approach differs from the SOCR method in two aspects. Firstly, the SOCR method uses the quadratic data misfit as the data fidelity term, which assumes noise of the Gaussian type. In the proposed SSCR method, we study the randomness of photon event stamping, and use logarithmic data fidelity terms derived from Bernoulli photon event assumptions. Secondly, the SOCR method considers regularizations for three-dimensional voxels, while the proposed method introduces two-dimensional regularizations for the surface of the hidden objects and incorporates priors in both 2D and 3D representations of the targets.

3. The physical model

In general, the reconstruction quality and the computation complexity of an NLOS imaging algorithm rely heavily on the forward model that simulates the physical measurement process. Fine physical models lead to reconstructions with clear geometric structures, but at the expense of high computational cost [13, 36]. Instead of putting forward a novel physical model, we aim at overcoming the ill-posedness of the NLOS imaging task with an extremely small number of measurements. We adopt a linear physical model which only considers the square fall-off of the photon intensity. Let x'_i and x'_d be the illumination and detection points on the visible surface, the photon intensity detected at time t is modeled as

$$\tau(x'_i, x'_d, t) = \int_{\Omega} \frac{u(x)}{\|x'_i - x\|^2 \|x'_d - x\|^2} \delta(\|x'_i - x\| + \|x'_d - x\| - ct) dx, \quad (1)$$

in which c is the speed of light, x is a point in the three-dimensional reconstruction domain Ω . The albedo value of the point x is represented by $u(x)$. The δ function describes the intrinsic domain of integration as the set of points with optical path length ct . When $x'_i \neq x'_d$, the domain of integration is a half ellipsoid with foci x'_i and x'_d . When $x'_i = x'_d = x'$, the domain of integration is a half sphere with center x' and the model reduces to the one used in the work [30], which writes

$$\tau(x', t) = \int_{\Omega} \frac{u(x)}{\|x' - x\|^4} \cdot \delta(2\|x' - x\| - ct) dx. \quad (2)$$

4. The SSCR method

In this section, we study the randomness in the measurement process and propose a signal-surface collaborative regularization framework for few-shot NLOS imaging scenarios.

Notation The L_2 norm of a vector \mathbf{x} is denoted by $\|\mathbf{x}\|$. We use $[N]$ as an abbreviation of the set $\{1, 2, \dots, N\}$. We use \otimes to denote the Kronecker product of matrices. The three-dimensional reconstruction domain Ω is discretized with voxels $V = \{v_{ijk} = (x_i, y_j, z_k) | i \in [I], j \in [J], k \in [K]\}$, where x_i , y_j , and z_k are the coordinates of the point v_{ijk} in the horizontal, vertical and depth directions. Each point in V represents a cubic voxel of the same size. The grid function \mathbf{u} is used to denote the discrete albedo values, with its components denoted by $u_{ijk} = \mathbf{u}(v_{ijk})$. To reconstruct the hidden targets, the signal is measured at P pairs of points $\{(x'_p, y'_p)\}_{p=1}^P$, in which x'_p and y'_p are the coordinates of the p^{th} illumination and detection points, respectively. For each measurement pair (x'_p, y'_p) , the time resolved signal contains Q time bins. The length of each time bin is a constant, usually at the scale of picoseconds in

real applications. We denote by $\tau_{p,q}$ the photon intensity of the p^{th} measurement pair and the q^{th} time bin. The set of photon intensities is denoted by $\boldsymbol{\tau}$.

The reconstructed surface We use $\mathbf{u} \in \mathbb{R}^{I \times J \times K}$ to represent the albedo of the NLOS scene, which is a three-dimensional tensor. However, it is only possible to reconstruct the portion of the hidden surface where photons are bounced back, which is a two dimensional geometric object. With this observation, we define a subset $\mathcal{G} \subsetneq \mathbb{R}^{I \times J \times K}$ as follows

$$\mathcal{G} = \{ \mathbf{g} = (g_{ijk}) \in \mathbb{R}^{I \times J \times K} \mid \forall (i, j) \in [I] \times [J], \\ \exists \text{ at most one } k = k_{ij} \in [K], \text{ s.t. } g_{i,j,k_{ij}} \neq 0 \}. \quad (3)$$

Each element $\mathbf{g} \in \mathcal{G}$ yields a trival two-dimensional parameterization. For each pixel $(i, j) \in [I] \times [J]$, only one of the following cases holds.

- Case 1: $g_{ijk} = 0$ for all $k \in [K]$. In this case, the line $x = x_i, y = y_j$ does not intersect with the target and we call (i, j) a *background pixel*.
- Case 2: There exists only one $k = k_{ij}$, such that $g_{i,j,k_{ij}} \neq 0$. In this case, we call (i, j) a *foreground pixel* and the corresponding *depth* is z_{ij} .

To express elements of \mathcal{G} with matrices, we show that there is a bijection from the set \mathcal{G} to the following set

$$\mathcal{G}' = \{ (\mathbf{e}, \mathbf{d}, \boldsymbol{\alpha}) \mid \mathbf{e} = (e_{ij})_{I \times J}, e_{ij} \in \{0, 1\}, \\ \mathbf{d} = (d_{ij})_{I \times J}, d_{ij} \in [K] \cup \{\text{NaN}\}, \\ \boldsymbol{\alpha} = (\alpha_{ij})_{I \times J}, \alpha_{ij} \in \mathbb{R} \cup \{\text{NaN}\}, \\ \forall (i, j), e_{ij} = 0 \iff d_{ij} = \text{NaN} \iff \alpha_{ij} = \text{NaN} \}. \quad (4)$$

In the definition of \mathcal{G}' , the placeholder NaN represents a background pixel and does not operate with real numbers. To construct a bijection from \mathcal{G} to \mathcal{G}' , for each element $\mathbf{g} \in \mathcal{G}$, let \mathbf{e} be the indicator function of the set of foreground pixels, \mathbf{d} be the depths of the foreground pixels and $\boldsymbol{\alpha}$ be the corresponding albedo values. Then, fill the matrices \mathbf{d} and $\boldsymbol{\alpha}$ with NaNs where necessary. It is easy to check that this map: $\mathcal{G} \rightarrow \mathcal{G}'$ is one to one and onto. We call \mathbf{e} , \mathbf{d} and $\boldsymbol{\alpha}$ the foreground indication matrix, the depth matrix and the albedo matrix of \mathbf{g} , respectively. We will design regularizations for elements of \mathcal{G} , in which the matrix representations bring remarkable convenience.

The data of photon event stamping In NLOS detections, the intensity of the back scattered light is extremely weak after multiple diffuse reflections. For each measurement pair, a total of N laser pulses are emitted to the illumination point. We use the binary variable $d_{p,q,n}$ to denote the recorded photon event

$$d_{p,q,n} = \begin{cases} 1, & \text{record a photon event} \\ 0, & \text{otherwise} \end{cases}, \quad (5)$$

in which $p \in [P]$, $q \in [Q]$, and $n \in [N]$ are indices of the measurement pair, the time bin and the pulse number. For each measurement pair, the detector can record at most one photon event for each pulse, which means that $\sum_q d_{p,q,n} \leq 1$. However, the case of recording more than one photon event in a single pulse can be neglected [18]. The collection of photon event stamping is denoted by \mathbf{d} .

The Bayesian framework We propose a unified Bayesian framework that reconstructs the hidden targets with the measured data of photon event stamping. For each measurement pair p , time bin q , and pulse number n , it is assumed that the detection of a photon event $e_{p,q,n}$ follows the Bernoulli distribution with probability $\mathbb{P}\{e_{p,q,n} = 1\} = 1 - e^{-\eta\tau_{p,q}}$, in which $\eta > 0$ is the detection efficiency [30]. The collection of random variables $e_{p,q,n}$ is denoted by \mathbf{e} . In NLOS detection scenarios, the probability of detecting a photon event is extremely small. The first order approximation of the exponent is adopted and we assume

$$\mathbb{P}\{e_{p,q,n} = 1\} = \eta\tau_{p,q}. \quad (6)$$

In Eq. (1), the photon intensity is linear with respect to the albedo, we choose $\eta = 1$ without loss of generality. We also assume that the detection of different photon events is independent. Let $\mathbf{g} \in \mathcal{G}$ be the albedo of the hidden surface. We view \mathbf{g} and $\boldsymbol{\tau}$ as random vectors and find them simultaneously by maximizing the posterior probability $\mathbb{P}(\mathbf{g}, \boldsymbol{\tau} \mid \mathbf{e} = \mathbf{d})$, where \mathbf{e} is related with $\boldsymbol{\tau}$ by Eq. (6). Noting that the set \mathcal{G} is not a convex subset of $\mathbb{R}^{I \times J \times K}$, the resulting optimization problem is non-convex and hard to solve. To tackle this problem, we introduce the random vector $\mathbf{u} \in \mathbb{R}^{I \times J \times K}$ as an approximation of the surface \mathbf{g} and maximize $\mathbb{P}(\mathbf{g}, \mathbf{u}, \boldsymbol{\tau} \mid \mathbf{e} = \mathbf{d})$. Besides, we assume that the conditional probability of \mathbf{e} only depends on $\boldsymbol{\tau}$, which means

$$\mathbb{P}(\mathbf{e} = \mathbf{d} \mid \mathbf{g}, \mathbf{u}, \boldsymbol{\tau}) = \mathbb{P}(\mathbf{e} = \mathbf{d} \mid \boldsymbol{\tau}). \quad (7)$$

Using the Bayesian formula, we have

$$\begin{aligned} & \arg \max_{\mathbf{g}, \mathbf{u}, \boldsymbol{\tau}} \mathbb{P}(\mathbf{g}, \mathbf{u}, \boldsymbol{\tau} \mid \mathbf{e} = \mathbf{d}) \\ &= \arg \max_{\mathbf{g}, \mathbf{u}, \boldsymbol{\tau}} \mathbb{P}(\mathbf{e} = \mathbf{d} \mid \mathbf{g}, \mathbf{u}, \boldsymbol{\tau}) \mathbb{P}(\mathbf{g}, \mathbf{u}, \boldsymbol{\tau}) \\ &= \arg \max_{\mathbf{g}, \mathbf{u}, \boldsymbol{\tau}} \mathbb{P}(\mathbf{e} = \mathbf{d} \mid \boldsymbol{\tau}) \mathbb{P}(\mathbf{g}, \mathbf{u}, \boldsymbol{\tau}) \\ &= \arg \max_{\mathbf{g}, \mathbf{u}, \boldsymbol{\tau}} \prod_{p,q,n} \mathbb{P}(e_{p,q,n} = d_{p,q,n} \mid \tau_{p,q}) \mathbb{P}(\mathbf{g}, \mathbf{u}, \boldsymbol{\tau}) \\ &= \arg \max_{\mathbf{g}, \mathbf{u}, \boldsymbol{\tau}} \prod_{p,q} (\tau_{p,q})^{d_{p,q}} (1 - \tau_{p,q})^{N - d_{p,q}} \mathbb{P}(\mathbf{g}, \mathbf{u}, \boldsymbol{\tau}) \\ &= \arg \min_{\mathbf{g}, \mathbf{u}, \boldsymbol{\tau}} \sum_{p,q} [(d_{p,q} - N) \ln(1 - \tau_{p,q}) \\ & \quad - d_{p,q} \ln(\tau_{p,q})] + \Gamma(\mathbf{g}, \mathbf{u}, \boldsymbol{\tau}), \end{aligned} \quad (8)$$

where $d_{p,q} = \sum_{n=1}^N d_{p,q,n}$ is the data of photon event histogram. $\Gamma(\mathbf{g}, \mathbf{u}, \boldsymbol{\tau})$ is the joint regularization term of \mathbf{g} , \mathbf{u} and $\boldsymbol{\tau}$.

The joint regularization term $\Gamma(\mathbf{g}, \mathbf{u}, \boldsymbol{\tau})$ plays a crucial role in the process of reconstruction. An ingenious design of this term not only results in faithful reconstructions, but also admits low-cost algorithms to solve the optimization problem (8). In this paper, we assume

$$\Gamma(\mathbf{g}, \mathbf{u}, \boldsymbol{\tau}) = \lambda \|\boldsymbol{\tau} - A\mathbf{u}\|^2 + J_1(\mathbf{u}) + J_2(\mathbf{u}, \boldsymbol{\tau}) + J_3(\mathbf{u}, \mathbf{g}), \quad (9)$$

in which A is the forward operator defined by Eq. (1). J_1 describes the prior distribution of \mathbf{u} , J_2 is the joint prior of \mathbf{u} and $\boldsymbol{\tau}$, J_3 is the joint prior of \mathbf{u} and \mathbf{g} . λ is a fixed parameter.

The prior $J_1(\mathbf{u})$ The priors of the voxel based representation of the targets have been widely used in existing works [14, 20, 30]. Two efficient priors are the sparseness and non-local self-similarity of the objects. In the work [25], these two priors are considered for the 4D tensor of the directional albedo. Here, we simplify the approach and directly use L_1 norm of the albedo to impose the sparseness of the target. For the non-local self-similarity prior, we directly follow the block-matching and sparse representation method in the work [25], and set

$$J_1(\mathbf{u}) = s_u \|\mathbf{u}\|_1 + \lambda_u \sum_i [\|\mathcal{B}\mathbf{u}_i - D_s C_i D_n^T\|^2 + \lambda_{pu} |C_i|_0], \quad (10)$$

in which s_u , λ_u , and λ_{pu} are fixed parameters. $|\cdot|_0$ denotes the number of nonzero elements. The summation is made over all possible local 3D albedo blocks. The matrix $\mathcal{B}\mathbf{u}_i$ is constructed by putting the vectorizations of the i^{th} local block \mathbf{u}_i and its neighbors column by column. The orthogonal matrices D_s and D_n capture the local structure and non-local self-similarity of the albedo \mathbf{u} . C_i contains the transform-domain coefficients of the i^{th} block, whose sparseness is imposed by the term $|C_i|_0$. For more details, we refer the readers to [4, 9, 25].

The prior $J_2(\mathbf{u}, \boldsymbol{\tau})$ We seek for a joint local sparse representation scheme for the estimated signal $\boldsymbol{\tau}$ and the simulated signal $A\mathbf{u}$. It is assumed that $\boldsymbol{\tau}$ is a three dimensional tensor of size $N_x \times N_y \times Q$, in which N_x and N_y are the number of measurement points in the horizontal and vertical directions. Q is the number of time bins. We call a three-dimensional sub-tensor of $\boldsymbol{\tau}$ a local patch. Consider the set of all possible patches of size $r_x \times r_y \times r_q$. We obtain the patch dataset $\mathcal{P}(\boldsymbol{\tau})$ by generating the vectorization of each patch and put them together column by column. We use the orthogonal dictionary $D = D_q \otimes D_y \otimes D_x$ as the transform basis, where D_q , D_y and D_x are matrices of the discrete

consine transform of orders q , y and x , respectively. The joint prior of \mathbf{u} and $\boldsymbol{\tau}$ is given by

$$J_2(\mathbf{u}, \boldsymbol{\tau}) = \lambda_t \|\mathcal{P}(\boldsymbol{\tau}) - DS\|^2 + \lambda_{ut} \|\mathcal{P}(A\mathbf{u}) - DS\|^2 + \lambda_{pt} |S|_0, \quad (11)$$

in which λ_t , λ_{ut} and λ_{pt} are fixed parameters. $\mathcal{P}(A\mathbf{u})$ represents the patch dataset generated by the simulated signal of \mathbf{u} . S contains the public transform domain coefficients of $\mathcal{P}(\boldsymbol{\tau})$ and $\mathcal{P}(A\mathbf{u})$, whose sparseness is imposed by the L_0 term.

The prior $J_3(\mathbf{u}, \mathbf{g})$ We express the joint regularization of \mathbf{u} and \mathbf{g} as

$$J_3(\mathbf{u}, \mathbf{g}) = \lambda_g [\|\mathbf{u} - \mathbf{g}\|^2 + \Upsilon(\mathbf{g})], \quad (12)$$

in which $\Upsilon(\mathbf{g})$ describes the prior distribution of $\mathbf{g} \in \mathcal{G}$. The set \mathcal{G} is not convex, making it difficult to design $\Upsilon(\mathbf{g})$ explicitly. In fact, it suffices to update \mathbf{g} with fixed \mathbf{u} and vice versa in the final optimization problem. With fixed \mathbf{g} , \mathbf{u} can be easily updated with the term $\|\mathbf{u} - \mathbf{g}\|^2$. To update \mathbf{g} with any fixed $\mathbf{u} \in \mathbb{R}^{I \times J \times K}$, we choose an element from the set \mathcal{G} which not only lies in the neighborhood of \mathbf{u} in the L_2 sense, but also acts like the surface of some real-world object. With this motivation in mind, we construct a map $\mathcal{S} : \mathbb{R}^{I \times J \times K} \rightarrow \mathcal{G}$ and view $\mathcal{S}(\mathbf{u})$ as the solution to the following optimization problem.

$$\mathcal{S}(\mathbf{u}) = \arg \min_{\mathbf{g}} \|\mathbf{u} - \mathbf{g}\|^2 + \Upsilon(\mathbf{g}). \quad (13)$$

We call $\mathcal{S}(\mathbf{u})$ a *surfaiation* of \mathbf{u} . To construct the map \mathcal{S} , we assign for each $\mathbf{u} \in \mathbb{R}^{I \times J \times K}$ an indicator matrix \mathbf{e} , a depth matrix \mathbf{d} and an albedo matrix $\boldsymbol{\alpha}$ (recall the definition of \mathcal{G}'). Different methods that generate the surfaiation of \mathbf{u} lead to different surface regularizations. For clarity, we state a basic method here and provide more technical tricks in the supplement.

For each pixel (i, j) of \mathbf{u} , let $u_{i,j,k_{ij}^1}, \dots, u_{i,j,k_{ij}^{n_{ij}}}$ be all n_{ij} non-zero albedo values in the depth direction. Define

$$\tilde{e}_{ij} = \begin{cases} 1, & n_{ij} > 0 \\ 0, & n_{ij} = 0 \end{cases} \quad (14)$$

to be the indicator function of the set of foreground pixels. There could be many mislabeled pixels due to heavy background noise in \mathbf{u} . To provide a noise-robust estimation of the indicator matrix \mathbf{e} , we consider the correlations of all pixels and solve

$$\begin{aligned} \mathbf{e}^* &= (\mathbf{e}_{ij}^*)_{I \times J} \\ &= \arg \min_{\{e_{ij}\}} \sum_{i=1}^I \sum_{j=1}^J \gamma_{ij} (e_{ij} - \tilde{e}_{ij})^2 \\ &\quad + \sum_{p=1}^I \sum_{q=1}^J \sum_{r=1}^I \sum_{s=1}^J w_{pq,rs}^e (e_{pq} - e_{rs})^2, \end{aligned} \quad (15)$$

in which $w_{pq,rs}^e$ describes the weight of the pixels (p, q) and (r, s) . The parameter γ_{ij} describes the confidence of indication of the original pixel \tilde{e}_{ij} . This least-squares problem has a unique solution and can be solved using the standard LSQR method. The foreground indication matrix $\mathbf{e} = (e_{ij})$ is determined by

$$e_{ij} = \begin{cases} 1, & e_{ij}^* \geq 0.5 \\ 0, & e_{ij}^* < 0.5 \end{cases} . \quad (16)$$

To obtain the depth matrix \mathbf{d} , we solve the following least-squares problem

$$\begin{aligned} \mathbf{d}^* &= (d_{ij}^*)_{I \times J} \\ &= \arg \min_{\{d_{ij}\}} \sum_{i=1}^I \sum_{j=1}^J \sum_{n=1}^{n_{ij}} \lambda_{ijn} (d_{ij} - z_{k_{ij}^n})^2 \\ &\quad + \sum_{p=1}^I \sum_{q=1}^J \sum_{r=1}^I \sum_{s=1}^J w_{pq,rs}^d (d_{pq} - d_{rs})^2, \end{aligned} \quad (17)$$

in which $z_{k_{ij}^n}$ is the depth of the voxel (i, j, k_{ij}^n) . λ_{ijn} and $w_{pq,rs}^d$ are fixed parameters that control the weight of the corresponding terms. The depth matrix $\mathbf{d} = (d_{ij})$ is then determined by

$$d_{ij} = \begin{cases} \arg \min_k \|z_k - d_{ij}^*\|, & e_{ij} = 1 \\ \text{NaN}, & e_{ij} = 0 \end{cases} . \quad (18)$$

If there are two different values of k that minimizes $\|z_k - d_{ij}^*\|$, the smaller one is used.

The albedo matrix $\alpha = (\alpha_{ij}^*)$ is obtained by solving the following optimization problem

$$\begin{aligned} \alpha &= (\alpha_{ij}^*)_{I \times J} \\ &= \arg \min_{\{\alpha_{ij}\}} \sum_{i=1}^I \sum_{j=1}^J \sum_{n=1}^{n_{ij}} \lambda_{ijn} (\alpha_{ij} - u_{ijk_{ij}^n})^2 \\ &\quad + \sum_{p=1}^I \sum_{q=1}^J \sum_{r=1}^I \sum_{s=1}^J w_{pq,rs}^\alpha (\alpha_{pq} - \alpha_{rs})^2, \end{aligned} \quad (19)$$

in which λ_{ijn} and $w_{pq,rs}^\alpha$ are fixed parameters. Finally, the element α_{ij}^* is reset as NaN if $e_{ij} = 0$.

Finally we obtain the optimization problem of the proposed signal-surface collaborative regularization (SSCR) framework as follows

$$\begin{aligned} &\arg \min_{\substack{\tau, \mathbf{u}, \mathbf{g}, \\ D_s, D_n, \mathcal{G}, S}} \sum_{p,q} [(d_{p,q} - N) \ln(1 - \tau_{p,q}) - d_{p,q} \ln(\tau_{p,q})] \\ &+ \lambda_t \|\mathcal{P}(\tau) - DS\|^2 + \lambda_{ut} \|\mathcal{P}(A\mathbf{u}) - DS\|^2 + \lambda_{pt} |S|_0 \\ &+ \lambda \|\tau - A\mathbf{u}\|^2 + s_u \|\mathbf{u}\|_1 + \lambda_g [\|\mathbf{u} - \mathbf{g}\|^2 + \Upsilon(\mathbf{g})] \\ &+ \lambda_u \sum_i [\|\mathcal{B}\mathbf{u}_i - D_s C_i D_n^T\|^2 + \lambda_{pu} |C_i|_0] \\ \text{s.t. } &\mathbf{g} \in \mathcal{G}, \quad D_s D_s^T = I_x, \quad D_n D_n^T = I_y. \end{aligned} \quad (20)$$




Figure 4. Reconstruction results of the pyramid with confocal synthetic signal. a) - f) The front view of the reconstructions obtained with the LOG-BP, F-K, LCT, D-LCT, SOCR, and the proposed SSCR methods.

in which x is the product of the block size in three directions, y is the number of neighbors for each block. This optimization problem is solved using the alternating iteration method, as summarized in Algorithm 1. In the supplement, we provide a detailed discussion of the solution to each sub-problem and the choices of parameters. The surface and voxel representations of the target are given by \mathbf{u} and \mathbf{g} , respectively.

5. Results

To validate the capability of the proposed method in reconstructing the hidden targets with sparse measurements, we compare our reconstruction results with the Laplacian of Gaussian filtered back-projection (LOG-BP) [17] and SOCR [25] methods. We also bring the F-K [20], LCT [30], D-LCT [42] and the PF [23] methods into comparisons by constructing dense measurements with the linear interpolation technique. In the supplement, we provide a gallery of reconstruction results of these methods with other signal interpolation techniques, where similar reconstruction results are shown.

Confocal experiments We use the synthetic signal of the instance of the pyramid [25] to test the proposed method. Only 3×3 of the original 64×64 synthetic signals are chosen. The physical model used to generate the data considers cosine attenuation of the photon intensity, and is finer than Eq. (2). The base length and height of the pyramid are 1 m and 0.2 m, respectively. The central axis of the regular quadrangular pyramid is vertical to the planer relay surface. The pyramid is 0.5 m away from the relay surface and the time resolution is 32 ps. The experimental setup and the three views of our reconstruction are shown in Fig. 1. The

Algorithm 1 The SSCR algorithm

Require: \mathbf{d}, N .

Ensure: \mathbf{u}, \mathbf{g} .

$$\tau_{p,q}^0 = d_{p,q}/N$$

$$\mathbf{u}^0 = \arg \min_{\mathbf{u}} \lambda \|\boldsymbol{\tau}^0 - \mathbf{A}\mathbf{u}\|^2 + s_u \|\mathbf{u}\|_1$$

for $k = 1$ to $K - 1$ **do**

$$\mathbf{g}^{k+1} = \arg \min_{\mathbf{g}} \|\mathbf{u}^k - \mathbf{g}\|^2 + \Upsilon(\mathbf{g})$$

$$(D_s^{k+1}, \mathbf{C}^{k+1}, D_n^{k+1}) = \arg \min_{D_s, \mathbf{C}, D_n} \|\mathcal{B}\mathbf{u}_i^k - D_s C_i D_n^T\|^2 + \lambda_{pu} |C_i|_0$$

$$S^{k+1} = \arg \min_S \lambda_t \|\mathcal{P}(\boldsymbol{\tau}^k) - DS\|^2 + \lambda_{ut} \|\mathcal{P}(\mathbf{A}\mathbf{u}^k) - DS\|^2 + \lambda_{pt} |S|_0$$

$$\boldsymbol{\tau}^{k+1} = \arg \min_{\boldsymbol{\tau}} \sum_{p,q} [(d_{p,q} - N) \ln(1 - \tau_{p,q}) - d_{p,q} \ln(\tau_{p,q})] + \lambda_t \|\mathcal{P}(\boldsymbol{\tau}) - DS^{k+1}\|^2 + \lambda \|\boldsymbol{\tau} - \mathbf{A}\mathbf{u}^k\|^2$$

$$\mathbf{u}^{k+1} = \arg \min_{\mathbf{u}} \lambda \|\boldsymbol{\tau}^{k+1} - \mathbf{A}\mathbf{u}\|^2 + s_u \|\mathbf{u}\|_1 + \lambda_{ut} \|\mathcal{P}(\mathbf{A}\mathbf{u}) - DS^{k+1}\|^2 + \lambda_u \sum_i \|\mathcal{B}\mathbf{u}_i - D_s^{k+1} C_i^{k+1} (D_n^{k+1})^T\|^2 + \lambda_g \|\mathbf{u} - \mathbf{g}^{k+1}\|^2$$

end for

$$\mathbf{u} = \mathbf{u}^K$$

$$\mathbf{g} = \mathbf{g}^K$$
















Alg	Oracle	FK	LCT	SOCR	SSCR
5×5 0.34 s					
PSNR	-	14.39	13.68	20.45	23.16
5×5 0.05 s					
PSNR	-	15.40	12.96	20.31	21.67
10×10 1.37 s					
PSNR	-	15.03	12.81	16.42	20.78

Figure 5. Reconstruction results for the instances of the statue and the dragon. The number of focal points and exposure time are listed in the first column. The oracle is shown in the second column. Reconstructions of the FK, LCT, SOCR and SSCR algorithms are shown in the third to sixth columns.

coordinates of the focal points are provided in the supplementary code. Figure 4 shows the front views of the reconstruction results of different methods. The LOG-BP method fails to locate the hidden object. The F-K, LCT, D-LCT and SOCR reconstructions contain artifacts. The proposed SSCR method provides a faithful estimation of the hidden target. Supplementary Figure S4 shows the three views of these results, revealing a significant depth error in the F-K reconstruction.

We use the measured data of the instances of the statue and the dragon in the Stanford dataset [20] to test the proposed method in real-world applications. The original dataset contains 512×512 confocal measurements over a square region of 2×2 m². For the instance of the statue,

the distance to the relay is 1 m and the exposure time of the original 512×512 measurements are 60 min and 10 min, respectively. We use 5×5 sub-sampled signals for the reconstruction, which would only take 0.34 s and 0.05 s to measure the sub-sampled signals. For the instance of the dragon, the distance is 1.3 m to the relay surface and the total exposure time is 60 min. We use 10×10 sub-sampled signals for the reconstructions, which would only take 1.37 s. For both instances, the time resolution is 32 ps. The focal points are provided in the supplementary code. Reconstruction results are compared in Fig. 5, where the oracle is generated with the SOCR method using 64×64 measurements. It is shown that the F-K and LCT reconstructions are blurry. The D-LCT reconstructions are provided in the supplement, and are similar to those obtained with the LCT method. The SOCR reconstructions are discontinuous and noisy. The SSCR reconstructions have the highest PSNR values and does not contain background noise. For the instance of the dragon, the specularity of the material leads to a large bias of the physical model and the SSCR method reconstructs a portion of the target.

For the instance of the statue with original exposure time of 60 min of the 512×512 measurements, we show three views of the SSCR reconstructions with different number of illumination points in Fig. 6. The target can be clearly reconstructed with 7×7 confocal measurements, which is only 0.01% of the original dataset. With 4×4 illuminations, the SSCR method still provides a reasonable estimation of the hidden target, which demonstrates the robustness of the proposed algorithm. More comparisons with existing methods are provided in the supplement.

Non-confocal experiments We use one simulated dataset and one measured dataset to test the method for non-confocal reconstructions. For the simulated experiment, we use the signal of the instance of letter ‘K’ from the NLoS




Figure 6. Reconstruction results of the statue with different illumination settings. The illumination points are shown in yellow in the first column. The front view, top view and side view of the reconstructions are shown in the second, third and fourth columns.

Benchmark dataset [12]. The oracle is generated with the SOCR method using measurements from 64×64 illumination points in a region of $0.512 \times 0.512 \text{ m}^2$. The detection point locates at the center of the illumination region. The photon travel distance is 0.001 m per second. A sub-sampled 6×6 signal is used for the reconstruction. For the measured dataset, we use the data of the instance of the figure ‘4’ provided by the work [23]. The original dataset contains 130×180 measurements in 390 min and we use a 6×6 sub-sampled dataset for the reconstruction. It would take 36 s to measure the sub-sampled signal. The coordinates of the illumination and detection points are provided in the supplementary code. Comparisons of the results with the LOG-BP [17], PF [23] and the SOCR methods [25] are shown in Fig. 7. The SSCR method yields the highest PSNR of the albedo values and does not contain background noise.

Alg	Oracle	LOG-BP	PF	SOCR	SSCR
6×6 Simu					
PSNR	-	10.28	16.68	16.08	23.88
6×6 36 s					
PSNR	-	10.42	15.63	17.59	19.90

Figure 7. Reconstruction results of non-confocal experiments. The SSCR method yields the highest PSNR values.

6. Discussion

When the reconstruction domain is discretized with $L \times L \times L$ voxels and the signal is detected at P measurement pairs ($P \leq L^2$), the time and memory complexity of the proposed SSCR method is $\mathcal{O}(PL^3)$ and $\mathcal{O}(L^3)$ respectively. In the supplement, we provide a detailed analysis of the complexity. Notably, when $P = \mathcal{O}(1)$, the time complexity of the method is $\mathcal{O}(L^3)$. This is smaller than the F-K, LCT, and D-LCT methods, which cost $\mathcal{O}(L^3 \log L)$. In this case, the LOG-BP and SOCR methods also yield a time complexity of $\mathcal{O}(L^3)$, but may provide biased estimations of the targets. This indicates the significance of the proposed two-dimensional regularization of the hidden surface. For the instance of the pyramid, the runtime of LOG-BP, F-K, LCT, D-LCT, SOCR and the proposed method are 0.1 s, 1.7 s, 1.1 s, 5.6 s, 142.0 s and 15.9 s on a laptop. In SOCR and SSCR methods, the sub-problems containing L_1 regularization take 20 iterations.

We discuss two aspects that could help improve the proposed method. In Eq. (6), we did not consider the dark count of the detector and the background noise. Modeling the noise distribution within the Bayesian framework would lead to better reconstruction results. Besides, for the two-dimensional regularization of \mathbf{g} , we simply use pixel level similarity in the SSCR method. Incorporating other regularizations [9, 29, 31] would lead to better estimations of the hidden surface.

7. Conclusion

We conclude that the two-dimensional surface regularization plays an important role in few-shot NLOS imaging scenarios. The joint regularization of the estimated signal, the voxel and surface representations of the target makes it possible to reconstruct the hidden object with certain complex geometric structures with only 5×5 confocal measurements, even in cases with measurement noise.

Acknowledgments. This work was supported by the National Natural Science Foundation of China (61975087, 12071244, 11971258).

References

- [1] Byeongjoo Ahn, Akshat Dave, Ashok Veeraraghavan, Ioannis Gkioulekas, and Aswin C. Sankaranarayanan. Convolutional approximations to the general non-line-of-sight imaging operator. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. **1**
- [2] Victor Arellano, Diego Gutierrez, and Adrian Jarabo. Fast back-projection for non-line of sight reconstruction. *Opt. Express*, 25(10):11574–11583, May 2017. **1**
- [3] Mauro Buttafava, Jessica Zeman, Alberto Tosi, Kevin Eliceri, and Andreas Velten. Non-line-of-sight imaging using a time-gated single photon avalanche diode. *Opt. Express*, 23(16):20997–21011, Aug 2015. **1**
- [4] Jian-Feng Cai, Hui Ji, Zuwei Shen, and Gui-Bo Ye. Data-driven tight frame construction and image denoising. *Applied and Computational Harmonic Analysis*, 37(1):89–105, 2014. **5**
- [5] Susan Chan, Ryan E. Warburton, Genevieve Gariepy, Jonathan Leach, and Daniele Faccio. Non-line-of-sight tracking of people at long range. *Opt. Express*, 25(9):10109–10117, May 2017. **1**
- [6] Sreenithy Chandran and Suren Jayasuriya. Adaptive lighting for data-driven non-line-of-sight 3d localization and object identification. *arXiv preprint arXiv:1905.11595*, 2019. **1**
- [7] Wenzheng Chen, Simon Daneau, Fahim Mannan, and Felix Heide. Steady-state non-line-of-sight imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. **2**
- [8] Javier Grau Chopite, Matthias B. Hullin, Michael Wand, and Julian Iseringhausen. Deep non-line-of-sight reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. **2**
- [9] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising with block-matching and 3d filtering. In Nasser M. Nasrabadi, Syed A. Rizvi, Edward R. Dougherty, Jaakko T. Astola, and Karen O. Egiazarian, editors, *Image Processing: Algorithms and Systems, Neural Networks, and Machine Learning*, volume 6064, page 606414. International Society for Optics and Photonics, SPIE, 2006. **5, 8**
- [10] Daniele Faccio, Andreas Velten, and Gordon Wetzstein. Non-line-of-sight imaging. *Nature Reviews Physics*, 2(6):318–327, Jun 2020. **1**
- [11] Xiaohua Feng and Liang Gao. Improving non-line-of-sight image reconstruction with weighting factors. *Opt. Lett.*, 45(14):3921–3924, Jul 2020. **1, 2, 3**
- [12] 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 *ACM SIGGRAPH 2019 Posters*, SIGGRAPH ’19, New York, NY, USA, 2019. Association for Computing Machinery. **8**
- [13] Felix Heide, Matthew O’Toole, Kai Zang, David B Lindell, Steven Diamond, and Gordon Wetzstein. Non-line-of-sight imaging with partial occluders and surface normals. *ACM Transactions on Graphics (ToG)*, 38(3):1–10, 2019. **3**
- [14] Mariko Isogawa, Dorian Chan, Ye Yuan, Kris M. Kitani, and Matthew O’Toole. Efficient non-line-of-sight imaging from transient sinograms. In *16th European Conference on Computer Vision (ECCV)*, 2020. **5**
- [15] Mariko Isogawa, Ye Yuan, Matthew O’Toole, and Kris M. Kitani. Optical non-line-of-sight physics-based 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. **1**
- [16] Masaki Kaga, Takahiro Kushida, Tsuyoshi Takatani, Kenichiro Tanaka, Takuya Funatomi, and Yasuhiro Mukaigawa. Thermal non-line-of-sight imaging from specular and diffuse reflections. *IPSJ Transactions on Computer Vision and Applications*, 11(1):8, Nov 2019. **1**
- [17] Martin Laurenzis and Andreas Velten. Feature selection and back-projection algorithms for nonline-of-sight laser-gated viewing. *Journal of Electronic Imaging*, 23(6):063003, 2014. **1, 6, 8**
- [18] Zhupeng Li, Xintong Liu, Jianyu Wang, Zuoqiang Shi, Lingyun Qiu, and Xing Fu. Fast non-line-of-sight imaging based on first photon event stamping. *Opt. Lett.*, 47(8):1928–1931, Apr 2022. **2, 4**
- [19] David B. Lindell, Gordon Wetzstein, and Vladlen Koltun. Acoustic non-line-of-sight imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. **1**
- [20] David B. Lindell, Gordon Wetzstein, and Matthew O’Toole. Wave-based non-line-of-sight imaging using fast f-k migration. *ACM Trans. Graph. (SIGGRAPH)*, 38(4):116, 2019. **2, 3, 5, 6, 7**
- [21] Jianjiang Liu, Yijun Zhou, Xin Huang, Zheng-Ping Li, and Feihu Xu. Photon-efficient non-line-of-sight imaging. *IEEE Transactions on Computational Imaging*, 8:639–650, 2022. **1**
- [22] Xiaochun Liu, Sebastian Bauer, and Andreas Velten. Analysis of feature visibility in non-line-of-sight measurements. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. **2, 3**
- [23] Xiaochun Liu, Sebastian Bauer, and Andreas Velten. Phasor field diffraction based reconstruction for fast non-line-of-sight imaging systems. *Nature Communications*, 11(1):1645, Apr 2020. **2, 6, 8**
- [24] 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 waveoptics. *Nature*, 572(7771):620–623, Aug 2019. **2**
- [25] Xintong Liu, Jianyu Wang, Zhupeng Li, Zuoqiang Shi, Xing Fu, and Lingyun Qiu. Non-line-of-sight reconstruction with signal–object collaborative regularization. *Light: Science & Applications*, 10(1):198, Sep 2021. **3, 5, 6, 8**
- [26] Tomohiro Maeda, Yiqin Wang, Ramesh Raskar, and Achuta Kadambi. Thermal non-line-of-sight imaging. In *2019 IEEE International Conference on Computational Photography (ICCP)*, pages 1–11, 2019. **1**
- [27] Christopher A. Metzler, Felix Heide, Prasana Rangarajan, Muralidhar Madabhushi Balaji, Aparna Viswanath, Ashok

- Veeraraghavan, and Richard G. Baraniuk. Deep-inverse correlative: towards real-time high-resolution non-line-of-sight imaging. *Optica*, 7(1):63–71, Jan 2020. 2
- [28] Ji Hyun Nam, Eric Brandt, Sebastian Bauer, Xiaochun Liu, Marco Renna, Alberto Tosi, Eftychios Sifakis, and Andreas Velten. Low-latency time-of-flight non-line-of-sight imaging at 5 frames per second. *Nature Communications*, 12(1):6526, Nov 2021. 2
- [29] Stanley Osher, Zuoqiang Shi, and Wei Zhu. Low dimensional manifold model for image processing. *SIAM Journal on Imaging Sciences*, 10(4):1669–1690, 2017. 8
- [30] 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, Mar 2018. 2, 3, 4, 5, 6
- [31] A. Rajwade, A. Rangarajan, and A. Banerjee. Image denoising using the higher order singular value decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(4), 2012. 8
- [32] Syed Azer Reza, Marco La Manna, Sebastian Bauer, and Andreas Velten. Phasor field waves: A Huygens-like light transport model for non-line-of-sight imaging applications. *Opt. Express*, 27(20):29380–29400, Sep 2019. 2
- [33] Nicolas Scheiner, Florian Kraus, Fangyin Wei, Buu Phan, Fahim Mannan, Nils Appenrodt, Werner Ritter, Jurgen Dickmann, Klaus Dietmayer, Bernhard Sick, and Felix Heide. Seeing around street corners: Non-line-of-sight detection and tracking in-the-wild using doppler radar. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1
- [34] Sheila W. Seidel, John Murray-Bruce, Yanting Ma, Christopher Yu, William T. Freeman, and Vivek K Goyal. Two-dimensional non-line-of-sight scene estimation from a single edge occluder. *IEEE Transactions on Computational Imaging*, 7:58–72, 2021. 1
- [35] Kenichiro Tanaka, Yasuhiro Mukaigawa, and Achuta Kadambi. Polarized non-line-of-sight imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1
- [36] Chia-Yin Tsai, Aswin C. Sankaranarayanan, and Ioannis Gkioulekas. Beyond volumetric albedo – a surface optimization framework for non-line-of-sight imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 3
- [37] 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, Mar 2012. 1
- [38] Bin Wang, Ming-Yang Zheng, Jin-Jian Han, Xin Huang, Xiu-Ping Xie, Feihu Xu, Qiang Zhang, and Jian-Wei Pan. Non-line-of-sight imaging with picosecond temporal resolution. *Phys. Rev. Lett.*, 127:053602, Jul 2021. 1, 2
- [39] Cheng Wu, Jianjiang Liu, Xin Huang, Zheng-Ping Li, Chao Yu, Jun-Tian Ye, Jun Zhang, Qiang Zhang, Xiankang Dou, Vivek K. Goyal, Feihu Xu, and Jian-Wei Pan. Non-line-of-sight imaging over 1.43 km. *Proceedings of the National Academy of Sciences*, 118(10), 2021. 1, 2
- [40] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 1
- [41] Jun-Tian Ye, Xin Huang, Zheng-Ping Li, and Feihu Xu. Compressed sensing for active non-line-of-sight imaging. *Opt. Express*, 29(2):1749–1763, Jan 2021. 2, 3
- [42] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2, 6
- [43] Dayu Zhu and Wenshan Cai. Fast non-line-of-sight imaging with two-step deep remapping. *ACS Photonics*, 9(6):2046–2055, 2022. 2