UC Berkeley
UC Berkeley Previously Published Works

Title
Fast Non-line-of-sight Imaging with Non-planar Relay Surfaces

Permalink
https://escholarship.org/uc/item/6cb4z24f

Authors
Gu, Chaoying
Sultan, Talha
Masumnia-Bisheh, Khadijeh
et al.

Publication Date
2023-07-30

DOI
10.1109/iccp56744.2023.10233262

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Fast Non-line-of-sight Imaging with Non-planar Relay Surfaces

Chaoying Gu, Talha Sultan, Khadijeh Masumnia-Bisheh, Laura Waller, and Andreas Velten

2023 IEEE International Conference on Computational Photography (ICCP) | 979-8-3503-1676-6/23/$31.00 ©2023 IEEE | DOI: 10.1109/ICCP56744.2023.10233262

Abstract—Non-line-of-sight imaging methods reconstruct images from light captured off a relay surface. In most prior work this relay surface is a diffuse plane. It has been shown that even small deviations from a planar relay wall geometry quickly degrade reconstruction quality. Although existing methods can account for relay surface geometry in a straightforward way, they typically have high computational complexity and take orders of magnitude longer time to compute than state-of-the-art planar methods. In this work, we propose a fast algorithm that can perform non-line-of-sight reconstruction on arbitrary non-planar relay surfaces. Our algorithm has the same computational and memory complexity as the fastest existing algorithms, yet it achieves comparable reconstruction quality to the widely-used slower algorithms.

Index Terms—Computational Photography, Time-of-flight Imaging, Non-line-of-sight Imaging

1 Introduction

Traditional LiDAR 3D imaging systems use pulsed lasers and Time-of-Flight (ToF) detectors to measure the ToF of the pulse from the laser to the object to the detector and generate a 3D image of the scene using this directly reflected signal. To generate 3D images of scenes that are hidden from the Line-of-Sight (LOS) of the observer, the directly visible LOS surface becomes a relay surface that reflects light from the hidden scene [1], [2]. In a ToF Non-Line-of-Sight (NLOS) measurement, a short laser pulse is reflected off points on the relay surface and the light returning from a hidden scene is measured to generate 3D reconstructions of the hidden scene. This NLOS imaging system has uses in diverse applications, such as autonomous navigation, disaster response, and infrastructure assessment.

Different approaches have been proposed to tackle the challenging problem of NLOS imaging [3]–[14]. Active ToF-based methods that utilize a pulsed laser and fast Single Photon Avalanche Diode (SPAD) detectors have been demonstrated to perform robust high quality 3D reconstructions of arbitrary room-sized hidden scenes [15]–[18].

However, most existing ToF-based reconstruction algorithms are limited to datasets acquired from planar relay surfaces with uniform sampling grids. Recent work has used data acquisition with multi-pixel SPAD detectors to improve light collection efficiency and reduce acquisition times [18]. While live video reconstructions have been demonstrated using this system and video rate reconstruction speeds are possible with multiple methods, these algorithms require the use of planar relay surfaces sampled with uniform grids.

The few algorithms that can handle general relay surfaces and sampling patterns while producing state-of-the-art reconstructions are computationally expensive [19]. Reconstruction times ranging from minutes to hours make them inappropriate for most real-world applications and inconvenient for research. Data-driven algorithms can provide fast high-quality reconstructions, but they are trained with a specific relay wall geometry and won't easily generalize to other geometries. Other deep methods learned to invert the light transport use a known forward model [20], [21]. These methods may be extendable to non-planar relay surfaces, but the necessity to retrain for each new relay surface typically results in a reconstruction time of hours.

The ability to reconstruct scenes from surfaces other than perfect planes, while maintaining speeds and quality comparable to existing state-of-the-art algorithms, is an important requirement for practical NLOS imaging systems. However, this task has received relatively little attention. In this paper, we aim to address these challenges by proposing a new 3D Rayleigh-Sommerfeld Diffraction(RSD) algorithm for NLOS imaging, which:

- can handle arbitrary non-planar relay surfaces
- can handle datasets collected with irregular sampling grids and non-confocal acquisition schemes on the relay surface
- has the same computational and memory complexity and has similar runtime as existing fast algorithms for planar relay surfaces
- yields similar reconstruction quality to the state-of-of-art (slow) algorithms

The experimental data and implementation code can be found in [22]. Our method can be applied to both confocal and non-confocal capture schemes. In a confocal measurement, an illumination laser and a single pixel detector scan together to match the illuminated and imaged points on the relay surface. This scheme simplifies the reconstruction process, but also limits light collection to a single detector pixel and a small area on the relay surface, resulting in light inefficiency [15]. Algorithms that can handle non-confocal datasets where the location of the detector pixel and laser

• *C. Gu and L. Waller are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, Berkeley, CA 94720. T. Sultan, K. Masumnia-Bisheh, and A. Velten are with the Department of Electrical and Computer Engineering, the University of Wisconsin-Madison, Madison, WI 53706 USA.*
E-mail: chaoying_gu@berkeley.edu

979-8-3503-1676-6/23/$31.00 ©2023 IEEE
Authorized licensed use limited to: Univ of Calif Berkeley. Downloaded on October 25,2023 at 23:49:16 UTC from IEEE Xplore. Restrictions apply.

illumination does not match can be extended to efficient capture systems with SPAD arrays [18].

The primary limitation of our algorithm is its inability to model non-linear light transport effects such as partial occlusion within the relay surface and partial reflection of the illumination laser at the edge of objects or at partially transmissive surfaces.

2 Related Work

2.1 Time of Flight based NLOS

The first method to successfully demonstrate ToF-based NLOS imaging used expensive Streak cameras and an ellipsoidal back-projection algorithm that operated on the time domain light transport of the measured signal [2]. Subsequent work sought to use cheaper SPAD detectors [23], decrease computation times by choosing to reconstruct along ellipsoidal surfaces in the hidden scene [24], and improve reconstruction quality by employing slow iterative algorithms [25], [26]. However time-domain reconstruction algorithms are computationally expensive, limiting the use of such methods "in the wild".

The light cone transform (LCT) [27] solves this problem of high computational complexity by utilizing a confocal acquisition scheme to formulate the forward model as a shift-invariant convolution. Then the inverse can be solved as a computationally efficient frequency-based deconvolution. However, this algorithm is restricted to singe pixel confocal data acquisition on planar relay wall with regular sampling grids.

2.2 Wave-Based NLOS

To enhance the efficiency and accuracy of LCT, new wave-based NLOS methods have been proposed. These methods utilize fast Fourier transforms and operate on the frequency domain of the measured signal [16], [17]. Individual frequency components of the measured signal are treated as waves, employing wave propagation techniques to model and propagate these virtual waves back into the hidden scene.

For the confocal acquisition, the measurement can be treated as a spatial boundary condition and the hidden scene as a temporal boundary condition. Then Frequency-Wavenumber (FK) migration algorithm from seismic imaging can be utilized to generate high-quality reconstructions [16]. While this algorithm is computationally efficient, the inefficiency of the confocal scan necessitates long acquisition times or the use of retroreflective objects.

In contrast, non-confocal schemes with independent laser and detector positions on relay surfaces can utilize multi-pixel ToF detectors to reduce data acquisition time [18]. Wave imaging methods can be used to generate state-of-the-art reconstructions of the hidden scene for both confocal and non-confocal acquisition schemes [15] using Rayleigh Sommerfeld Diffraction (RSD) propagator. Among the wave-based methods, the phasor field light transport formalism [28]–[35] can be understood as converting the visible relay surface into a virtual LOS camera. Fast frequency-domain reconstruction methods [17] paired with non-confocal acquisition schemes can generate dynamic reconstructions of the hidden scene [18]. However, aforementioned FK-migration and phasor field algorithms all require the use of a regular sampling grid on a planar "relay wall".

2.3 Non-Planar NLOS

In principle, time-domain ellipsoidal back projection methods [2], [15] can be used to reconstruct the hidden sene via arbitrary relay surfaces and sampling patterns, as long as the LOS geometry is known. The first method to experimentally validate this capability utilized a second SPAD detector to form a LiDAR system that dynamically updated the laser positions on the relay surface [36]. This setup enables high-quality reconstructions even when the relay surface is a moving curtain. However, the ellipsoidal back projection algorithm used [15] has high computational complexity resulting in hours of compute time for a single reconstruction.

Among wave-based methods, FK migration is the only approach that attempts to handle non-planar relay surfaces [16]. However, the two data pre-processing methods proposed for handling non-planarity have their limitations. One method increases the computational complexity to be greater than the ellipsoidal back projection, while the other is restricted to handling only minor deviations from planar walls as we will see in the comparisons below.

2.4 Computer-generated holography

Computer-generated holography (CGH) computes holograms from 3D description on an object, utilizing wave propagation similar to that in the phasor field NLOS reconstruction. However it generally relies on the Fresnel Transform, a simplified form of the RSD propagator used in our wave-based NLOS methods. The fast methods to perform volume-to-plane wave propagation and thus generate digital holograms are well-studied in point-based CGH [37]. For example, the wavefront recording plane (WRP) method accelerates this process, introducing an intermediate WRP in close proximity to the object. This method assumes that the thickness of the object space is small such that all the object points are contributing to a limited zone on WRP. Our methodology deviates from WRP, relaxing the surface thickness assumption and ensuring rapid propagation via 3D convolution. With the same derivation introduced in the following sections, our method is adaptable to the field of CGH. However, it is worth noting that digital holograms often have much higher resolution than what we process, so computational cost may still be an issue. Exploring this extension would be interesting future research.

3 Theoretical Background

We base our method on the phasor field formalism. This method converts captured NLOS data into a wavefront and employs diffraction theory that is typically used for LOS imaging systems to perform the reconstruction. The captured temporal data is encoded into a phasor field wavefront at the relay surface, then the reconstruction operator propagates this wavefront backward into the hidden scene to reconstruct the NLOS objects. We start with a brief introduction to the phasor field algorithm based on RSD [15], a fast NLOS reconstruction limited to a planar relay

Authorized licensed use limited to: Univ of Calif Berkeley. Downloaded on October 25,2023 at 23:49:16 UTC from IEEE Xplore. Restrictions apply.

wall. Then we introduce a time domain equivalent of RSD, which can deal with any given surface and sampling, but is very slow.

3.1 RSD

The phasor field RSD algorithm provides fast non-approximative scene reconstructions of room-sized scenes from non-confocal measurements with less computational and memory requirements than the prior art [17]. With its fast data acquisition and reconstruction capabilities, RSD has been applied for real-time NLOS imaging applications [18].

The core idea of RSD algorithm is that it codes the measured temporal signal into phasors in the Fourier domain, and accelerates the calculation of diffraction by Fourier-based convolution. However, it should be noted that the fast convolution requires uniform grid sampling of the relay surface. As a result, the fast RSD algorithm is naturally limited to scenarios that involve plane-to-plane propagation with uniform sampling.

We show the notation of our NLOS measurement in Fig.1. Because light transport at each spatial position is linear and time-invariant, we can characterize the temporal measurement as an impulse response $H(\mathbf{x_p}, \mathbf{x_c}, t)$, where $\mathbf{x_p}$ is the laser illumination point on the relay surface, and $\mathbf{x_c}$ is the sensor detection point on the relay surface.

To perform the reconstruction, the RSD algorithm picks a virtual illumination function $\mathcal{P}(\mathbf{x}_\mathrm{p}, t)$ first, and then computes the corresponding scene response. The temporal measurement denotes the impulse response at each $\mathbf{x_c}$, so we can use convolution to obtain the virtual captured wavefront on the relay wall as:

$$\mathcal{P}(\mathbf{x}_\mathrm{c}, t) = \int \left[\mathcal{P}(\mathbf{x}_\mathrm{p}, t) * H(\mathbf{x}_\mathrm{p}, \mathbf{x}_\mathrm{c}, t)\right] \mathrm{d}\mathbf{x}_\mathrm{p}. \tag{1}$$

The resulted phasors $\mathcal{P}(\mathbf{x}_\mathrm{c}, t)$ are converted to the frequency space quantity $\mathcal{P}_\mathcal{F}(\mathbf{x}_\mathrm{c}, \omega)$ where each frequency component at a fixed ω is a monochromatic wave front of a phasor reflected from the scene and captured at the relay surface. This wavefront is then propagated by the RSD kernel $R_{\mathbf{x}_\mathrm{v}}$ back into the hidden scene as:

$$I(\mathbf{x}_\mathrm{v}) = \left| \int_{-\infty}^{+\infty} \mathcal{R}_{\mathbf{x}_\mathrm{v}}\left(\mathcal{P}_\mathcal{F}(\mathbf{x}_\mathrm{c}, \omega)\right) \frac{\mathrm{d}\omega}{2\pi} \right| \tag{2}$$

where $\mathbf{x}_\mathrm{v}$ is voxel in the reconstructed hidden scene.

Given that all the detection points on the relay wall $\mathbf{x}_\mathrm{c}$ are distributed on a regular plane, the fast RSD algorithm computes the diffraction integral by fast Fourier transform (FFT) convolution:

$$R_{\mathbf{x}_\mathrm{v}}\left(\mathcal{P}_\mathcal{F}(\mathbf{x}_\mathrm{c}, \omega)\right) = \mathcal{P}_\mathcal{F}(x_c, y_c, 0, \omega) \underset{\text{2D}}{*} G(x_c, y_c, z_v, \omega) \tag{3}$$

where $G(x_c, y_c, z_v, \omega) = e^{i\frac{\omega}{c}\sqrt{{x_c}^2+y_c^2+z_v^2}} / \sqrt{x_c^2+y_c^2+z_v^2}$.

Thus, all the $\mathbf{x}_\mathrm{v}$ values with the same depth z_v can be computed simultaneously using a single 2D FFT convolution. Typically, the reconstruction volume has the same sampling distance and aperture as the relay surface. Following the convention in the literature [17], we denote N as the number of sampling points along each spatial dimension for both the relay surface and the reconstruction scene. Consequently, the computational complexity of RSD is $\mathcal{O}\left(N^3 log(N)\right)$, as it involves calculating N convolutions of size N^2 for a fixed number of frequency components.

3.2 FBP

Back-projection algorithm, together with the filtered back-projection (FBP) methods, are standard solutions to linear inverse problems. They are some of the most popular methods for NLOS 3D reconstruction [38]. Besides, the FBP algorithm is very similar to the phasor field algorithm with the time-domain implementation of RSD in terms of the underlying math.

While there are various kinds of filters that can be applied to the measured signal, the basic reconstruction of FBP at $\mathbf{x}_\mathrm{v}$ is represented by:

$$I(\mathbf{x}_\mathrm{v}) = \iiint H(\mathbf{x}_\mathrm{p}, \mathbf{x}_\mathrm{c}, t)\, \delta\left(t - \frac{|\mathbf{x}_\mathrm{p} - \mathbf{x}_\mathrm{v}| + |\mathbf{x}_\mathrm{c} - \mathbf{x}_\mathrm{v}|}{c}\right) d\mathbf{x}_\mathrm{p} d\mathbf{x}_\mathrm{c} dt. \tag{4}$$

The back projection specifies a voxel $\mathbf{x}_\mathrm{v}$ in the hidden scene, enumerates each pair of the illumination point $\mathbf{x}_\mathrm{p}$ and the camera point $\mathbf{x}_\mathrm{c}$ on the relay wall, and searches for the photons which have the correct time of flight to be scattered by the specific voxel $\mathbf{x}_\mathrm{v}$.

Some filters like Laplacian are often applied to the measured signal to accentuate the high-frequency components of reconstruction and enhance its resolution. When filtered by the illumination function $\mathcal{P}(\mathbf{x}_\mathrm{p}, t)$ in RSD, FBP becomes a time domain implementation of RSD. To ensure consistency across various methods, we apply the same filter used in RSD for all the subsequent FBP implementations.

In non-confocal experiments, we adopt one of the two approaches: either fix the SPAD while scanning the laser or fix the laser while employing a SPAD array. Owing to Helmholtz reciprocity [39], the lights and cameras are interchangeable, rendering these two measurement schemes mathematically equivalent. For the subsequent derivation, we fix the laser point $\mathbf{x}_\mathrm{p}$ for each measurement, then Eq.4 is simplified to:

$$I(\mathbf{x}_\mathrm{v}) = \iint \mathcal{P}(\mathbf{x}_\mathrm{c}, t)\, \delta\left(t - \frac{|\mathbf{x}_\mathrm{p} - \mathbf{x}_\mathrm{v}| + |\mathbf{x}_\mathrm{c} - \mathbf{x}_\mathrm{v}|}{c}\right) \mathrm{d}\mathbf{x}_\mathrm{c} \mathrm{d}t \tag{5}$$

where $\mathcal{P}(\mathbf{x}_\mathrm{c}, t)$ is defined in Eq.1.

The computational complexity of FBP is $\mathcal{O}\left(N^5\right)$, as it has to enumerate every possible pairs of N^3 voxels in the reconstruction volume and N^2 sampling points on the relay surface.

4 Our Method

The most fundamental initiative of our proposed 3D RSD algorithm is to extend the applicability of fast RSD to non-planar scenes. Our method is based on the well-known separability of the RSD [40], thus we can propagate our wavefront to an intermediate plane in the hidden scene and reconstruct the volume from that intermediate plane using existing methods. We first introduce this intermediate plane, and then propose a fast 3D convolutional version of FBP to deal with non-planar surface and non-uniform sampling.

Authorized licensed use limited to: Univ of Calif Berkeley. Downloaded on October 25,2023 at 23:49:16 UTC from IEEE Xplore. Restrictions apply.

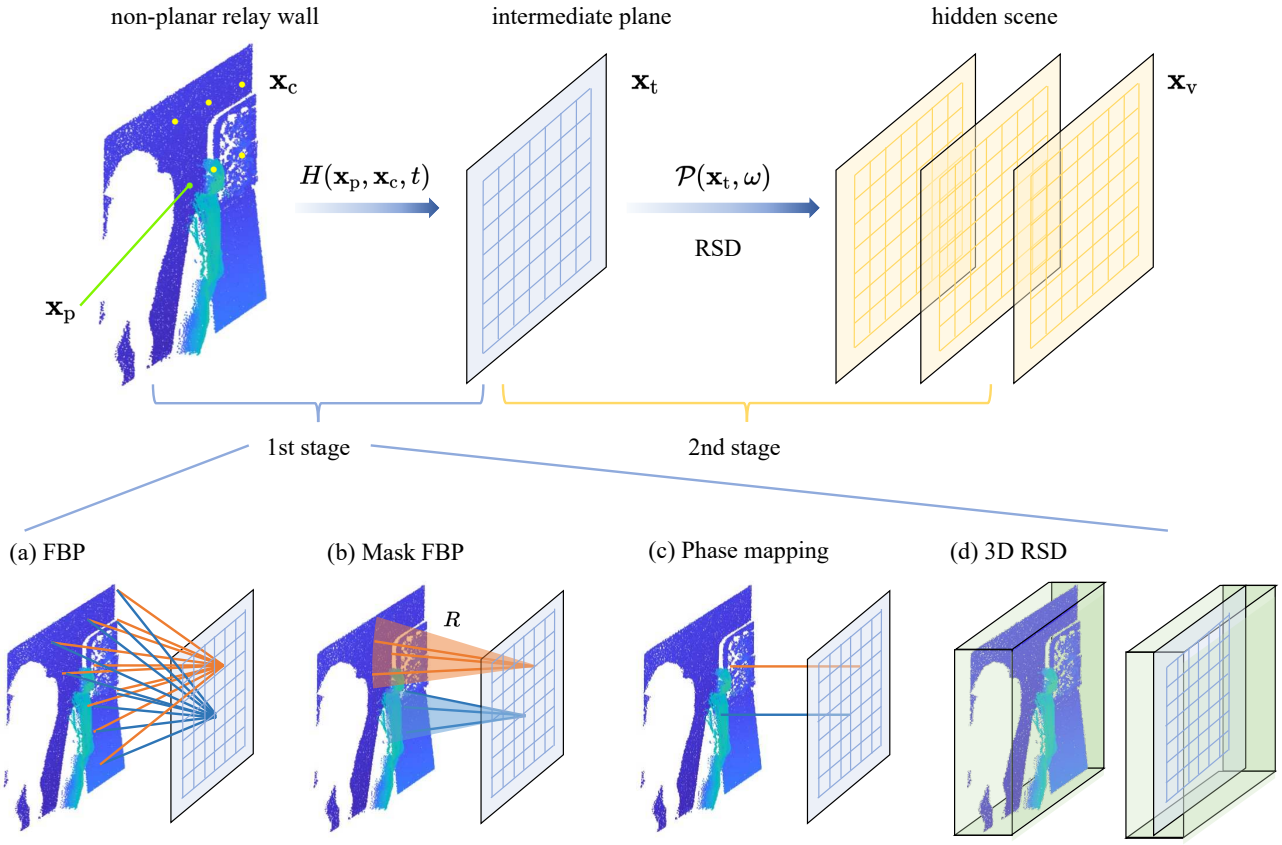

Fig. 1. 2-stage algorithm schematic. The top row shows the overall procedure for 2-stage non-planar reconstruction, and the bottom row illustrates four possible algorithms to implement the first stage propagation from a non-planar relay surface to an intermediate plane. The second stage involves the segmentation of the hidden scene into planar slices, and it is a standard planar reconstruction based on the RSD algorithm. **(a)** FBP: all the points inside the camera aperture ($\mathrm{x_c}$) are contributing to a specific $\mathrm{x_t}$ on the intermediate plane. **(b)** Mask FBP: only $\mathrm{x_c}$ inside a certain distance range are contributing to a specific $\mathrm{x_t}$. **(c)** Phase mapping: the intermediate plane is directly projected from the non-planar relay surface through phase shift, where $\mathrm{x_c}$ and $\mathrm{x_t}$ are one-to-one mapped. **(d)** 3D RSD: $\mathrm{x_c}$ is first fit into a 3D grid and then propagated by 3D convolution to retrieve the intermediate plane.

4.1 2-stage decomposition

The RSD is derived as an operator that propagates a wave from any closed surface (denoted by points $\mathbf{x}_c$ on the relay wall) to a different surface (denoted by points $\mathbf{x}_v$ in the hidden scene). For a closed or infinitely large surface, the operator is deterministic and inevitable. It can also be broken up into multiple steps by propagating to intermediate planes (denoted by $\mathbf{x}_t$ on the intermediate plane).

To utilize the fast RSD for our main reconstruction, we add an extra intermediate plane. This modification allows the slow algorithm, which involves non-planar geometries, to solely focus on the propagation to a single planar slice, rather than the entire hidden volume. Formally we define our 2-stage reconstruction algorithm as follows: the first stage propagates the measured field on the non-planar relay surface to an intermediate plane, and the second stage propagates from the intermediate plane to the hidden scene space through the familiar plane-to-plane RSD. The implementation of the first stage propagator will be discussed in more detail in the following two subsections.

4.2 Non-uniform sampling

In typical NLOS imaging setups, access to the occluded regions of a scene is infeasible, so the calibration of the relay surface must be angled. In situations where the relay surface is planar, it is possible to determine the angle between the wall and the camera axis, and subsequently, account for irregular sampling by adjusting the scanning pattern of galvos. However, in realistic non-planar scenarios, the deformation of relay surfaces makes it exceedingly arduous to mechanically compensate for uniform sampling. It would also require multiple iterative scans of the relay wall. Since the grid has to be uniform when viewed from the direction of the hidden scene, surfaces with self occlusions can create shadows that make complete uniform sampling as assumed by existing reconstruction algorithms impossible. As such, it is essential for reconstruction algorithms to be able to handle non-uniform sampling.

However, in most of the previous work capture systems are carefully calibrated to collect a sampling grid that appears uniform from the location of the hidden target [16], [17]. The 2D RSD approach achieves fast reconstruction owing to the uniform sampling grid in the x and y dimensions, which allows convolutional operations using FFT. Therefore, a regular grid is desired and a natural remedy for the non-planar and non-uniform sampling could be to interpolate the data. However, given that the points on the relay surface are scattered and have no structure between

Authorized licensed use limited to: Univ of Calif Berkeley. Downloaded on October 25,2023 at 23:49:16 UTC from IEEE Xplore. Restrictions apply.

their relative locations, the computational overhead can be extremely high when scaling. For example, generating the Delaunay triangulation using Bowyer-Watson algorithm has the complexity of $\mathcal{O}(N^4)$ for N^2 points on the relay surface [41].

With our intention to accelerate the existing FBP algorithm that can deal with irregular sampling, we opt for a more efficient way to interpolate. Specifically, we construct a uniform 3D grid and map the measurement on $\mathbf{x}_\mathrm{c}$ to the nearest grid point. The spacing of the resampled 3D grid is determined by the average sampling spacing on the relay surface, and the range is determined by the aperture size on the relay resurface.

The proposed simple interpolation method is particularly valid in the context of phasor field reconstruction. This is because, when viewing the object in the hidden scene as sources of spherical waves, the relay surface is almost tangential to the wavefront for the finite area inside the aperture. As a result, the gradient of the data along the relay surface is rather small at the scale of non-planar surface displacement. In our experimental datasets, surface gradients are typically small relative to the virtual wavelength and steps are limited to a few edges. The effectiveness of our approach, and how it measures up against more complex interpolation methods, is thoroughly discussed in Sec. 5.3.

However, it is worth noting that sharp discontinuities in the relay surface may result in substantial errors due to large gradients, especially parallel to wave propagation. However, the slow FBP might also yield unsatisfactory results due to the ringing artifacts introduced by sharp cutoffs in the sampling scheme. While we examined gradient-adapted windowing techniques for aliasing mitigation on more discontinuous surfaces, these often resulted in the loss of valuable high-frequency signals in reconstruction. Thus, addressing sharp discontinuities in non-planar relay surfaces without artifacts remains a compelling future research direction.

4.3 3D RSD

After the interpolation, we can derive the fast algorithm using tricks similar to 2D RSD. The 3D RSD algorithm first decomposes the measured signal into distinct frequency components and propagates them separately. Starting from Eq.5, we define the 4D reconstruction $I\left(\mathbf{x}_\mathrm{v}, t\right)$ as below and calculate its Fourier transform with expanded convolution:

$$\begin{aligned} I\left(\mathbf{x}_\mathrm{v}, t\right) &\triangleq \int H\left(\mathbf{x}_\mathrm{c}, t\right) \delta\left(t-\frac{\left|\mathbf{x}_\mathrm{p}-\mathbf{x}_\mathrm{v}\right|+\left|\mathbf{x}_\mathrm{c}-\mathbf{x}_\mathrm{v}\right|}{c}\right) d\mathbf{x}_\mathrm{c} \\ I_{\mathcal{F}}\left(\mathbf{x}_\mathrm{v}, \Omega\right) &= \int H\left(\mathbf{x}_\mathrm{c}, \Omega\right) * e^{-j\frac{\Omega}{c}\left(\left|\mathbf{x}_\mathrm{p}-\mathbf{x}_\mathrm{v}\right|+\left|\mathbf{x}_\mathrm{c}-\mathbf{x}_\mathrm{v}\right|\right)} d\mathbf{x}_\mathrm{c} \\ &= \iint H\left(\mathbf{x}_\mathrm{c}, \omega\right) e^{-j\frac{\Omega-\omega}{c}\left(\left|\mathbf{x}_\mathrm{p}-\mathbf{x}_\mathrm{v}\right|+\left|\mathbf{x}_\mathrm{c}-\mathbf{x}_\mathrm{v}\right|\right)} d\omega d\mathbf{x}_\mathrm{c}. \end{aligned} \tag{6}$$

Eq. 5 indicates that the integration of the 4D reconstruction over time is essentially the reconstruction intensity at $\mathbf{x}_\mathrm{v}$. Thus, we can rewrite the reconstruction as the DC component of the pre-defined 4D function:

$$I\left(\mathbf{x}_\mathrm{v}\right) = \int \mathcal{F}^{-1}\left\{I_{\mathcal{F}}\left(\mathbf{x}_\mathrm{v}, \Omega\right)\right\} dt = I_{\mathcal{F}}\left(\mathbf{x}_\mathrm{v}, 0\right). \tag{7}$$

We further use Eq. 6 to evaluate $I_{\mathcal{F}}\left(\mathbf{x}_\mathrm{v}, \Omega\right)$ at $\Omega = 0$. The sorted expression becomes a diffraction integral weighted by a phase mask depending on $\mathbf{x}_\mathrm{p}$:

$$\begin{aligned} I\left(\mathbf{x}_\mathrm{v}\right) &= \iint H\left(\mathbf{x}_\mathrm{c}, \omega\right) e^{-j\frac{\omega}{c}\left(\left|\mathbf{x}_\mathrm{p}-\mathbf{x}_\mathrm{v}\right|+\left|\mathbf{x}_\mathrm{c}-\mathbf{x}_\mathrm{v}\right|\right)} d\omega d\mathbf{x}_\mathrm{c} \\ &= \int e^{-j\frac{\omega}{c}\left|\mathbf{x}_\mathrm{p}-\mathbf{x}_\mathrm{v}\right|}\left[\int H\left(\mathbf{x}_\mathrm{c}, \omega\right) e^{-j\frac{\omega}{c}\left|\mathbf{x}_\mathrm{v}-\mathbf{x}_\mathrm{c}\right|} d\mathbf{x}_\mathrm{c}\right] d\omega \end{aligned} \tag{8}$$

where the inner integral can be accelerated through 3D FFT.

Let the grid-discretized $\mathbf{x}_\mathrm{c} = (x_c, y_c, z_c), \mathbf{x}_\mathrm{v} = (x_v, y_v, z_v)$, and we fix $z_v = z_0$ and consider one specific plane, which is exactly the intermediate plane in our specific case. Thus, the input field $\mathbf{x}_\mathrm{c}$ is centered at $z_c = 0$ and the propagated volume is centered at $z_v = z_0$. With these assumptions the inner integral simplifies into:

$$\begin{aligned} &\int H\left(\mathbf{x}_\mathrm{c}, \omega\right) e^{-j\frac{\omega}{c}\left|\mathbf{x}_\mathrm{v}-\mathbf{x}_\mathrm{c}\right|} d\mathbf{x}_\mathrm{c} \\ &= \sum_{x_c}\sum_{y_c}\sum_{z_c} H\left(x_c, y_c, z_c, \omega\right) e^{-j\frac{\omega}{c}\sqrt{(x_v-x_c)^2+(y_v-y_c)^2+(z_v-z_c)^2}} \\ &= H\left(x_v, y_v, z_v - z_0, \omega\right) \underset{\mathrm{3D}}{*} G(x_v, y_v, z_v, \omega)\Big|_{z_v=z_0} \end{aligned} \tag{9}$$

where $G(x_v, y_v, z_v, \omega) = e^{-j\frac{\omega}{c}\sqrt{(x_v)^2+(y_v)^2+(z_v)^2}}$.

It should be noted that while our input field, $H\left(x_v, y_v, z_v, \omega\right)$, has finite support, the diffraction kernel $G(x_v, y_v, z_v)$ does not. Consequently, 3D convolution cannot be used to calculate propagation to an entire volume. We observe that the 3D convolution relationship remains valid if the shifted kernel overlaps with the input, yielding the correct calculation for the central plane where $z_v = z_0$. Despite not being as efficient as whole volume-to-volume 3D convolution, this volume-to-plane propagation method still presents superior computational complexity when compared with the FBP integral.

Furthermore, the RSD operator describes propagation between 2D surfaces [42], so the extension to a 3D integral performed is not necessarily physically valid. In fact, since different planes within a volume are related through an RSD propagator, most 3D volumes are not valid inputs for this operator. Although Eq. 9 accurately represents the operator without necessitating any additional assumptions or approximations, regardless of the non-planar relay surface, there are limitations to our algorithm. Specifically, it does not account for occlusion or partial illumination, and operates under the presumption that the input data constitutes a surface. It is therefore necessary to have a valid input volume, which is one describing a 2D surface without occlusions and partially transparent voxels, with remaining voxels at zero.

The first stage computational complexity is a 3D convolution of N^3 size, and the second stage is 2D convolution of N^2 size repeated for N times. Therefore, the overall computational complexity can be expressed as $\mathcal{O}(N^3 \log(N))$. This complexity is the same as that of the existing fast planar RSD algorithm and is promising for real-time applications.

Authorized licensed use limited to: Univ of Calif Berkeley. Downloaded on October 25,2023 at 23:49:16 UTC from IEEE Xplore. Restrictions apply.

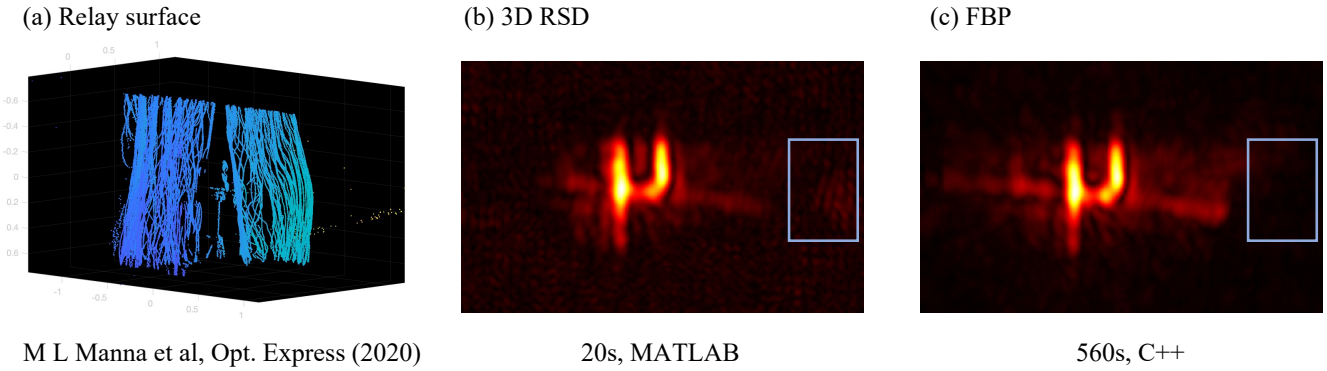

Fig. 2. NLOS reconstruction on existing non-uniform and non-planar datasets [36]. (a) Point cloud measured on the relay surface, representing a curtain adjacent to an air fan set on high. (b) Reconstruction using our 3D RSD algorithm, taking 20 seconds in MATLAB. (c) Reconstruction using standard FBP, as employed in the original paper, taking 560 seconds even with the accelerated C++ implementation.

5 Experimental Results

5.1 Performance analysis

In this section, we evaluate our algorithm on an existing dataset from a published paper [36]. The relay surface is a curtain blown by an air fan set to 'high' mode located beside it. Fig. 2(a) displays the measured point cloud on the relay surface, which is highly non-planar with non-uniform sampling. The original paper showed that the reconstruction assuming planar geometry failed to provide any information about the hidden scene. As a result, it is crucial to measure the non-planar relay surface and incorporate this information into the reconstruction process.

Fig. 2 demonstrates that the 3D RSD algorithm's reconstruction is slightly less sharp compared to the FBP method used in the original paper. As indicated by the blue frame, the 3D RSD algorithm also introduces more artifacts. This occurs because the intermediate plane cannot be infinitely large; we set it to the same size as the aperture on the relay surface to conserve computational resources. Additionally, errors are introduced by the nearest neighbor interpolation. However, most of the artifacts are common to both 3D RSD and FBP, likely resulting from partial occlusion that the linear light propagation cannot model. In general, the two reconstructions are visually similar, while our proposed method requires significantly less time.

While the FBP algorithm takes 560s with an efficient C++ version implementation, our 3D RSD reconstruction takes only 20s in MATLAB. In particular, fitting the measured data to a uniform 3D grid takes 14 seconds, while the actual 2-stage propagation requires only 6 seconds in total. This duration is essentially the same as that for the fast 2D RSD algorithm as introduced in Sec. 3.1.

5.2 Segmentation validation

To investigate whether the non-planar part indeed provides useful information in the reconstruction, we segmented the measured point cloud and perform NLOS reconstruction separately on planar and non-planar subsets of the relay surface. Fig.3 shows the results, including a side-by-side comparison with the FBP reconstruction. Our algorithm achieves comparable reconstruction quality to FBP with significantly less computational overhead on subsets reconstruction.

For the mannequin scene in Fig. 3(a), the non-planar part is contributing a noisy but significant part to the reconstruction. However, our method breaks for the tree scene in Fig. 3(b), where the segmented non-planar points fail to generate useful information of the hidden scene. We also demonstrate that the baseline FBP method is unable to get better reconstruction.

We conclude that the inverse problem becomes ill-posed when the relay surface has too much inter-reflection and occlusion. For example, in the tree dataset, one single laser pulse may illuminate multiple small leaves at varying depths, consequently producing multiple bounces. Additionally, the tree has complicated self-occlusion, eliminating some bounces which could have been propagated back. These factors remain unaccounted for in the linear wave-based light transport model. Learning-based methods or advanced neural rendering algorithms have potential to correct the forward model for extremely complex relay surfaces and are interesting topics for further research.

5.3 Interpolation comparison

Fig. 4 offers a detailed and quantitative analysis of interpolation methods. The analysis is performed on the most challenging non-planar dataset in Fig. 5 (d). The out-of-the-box application of MATLAB function to interpolate the whole 3D volume takes hours, and the fully-interpolated input volume is no longer valid as we defined in Sec. 4.3 , which accounts for the reconstruction in Fig. 4 (d).

When interpolating onto a regular 3D grid, we can identify the nearest grid points by simply dividing point positions by grid spacing. Therefore, during interpolation, we confine query points to these nearest grid points. This optimized approach makes the runtime of linear interpolation similar to nearest neighbor interpolation and marginally improves performance as Fig. 4 (e) shows.

In our method, we use optimized grid resolution as an alternative to more sophisticated interpolation, mainly serving the purpose of maintaining the $\mathcal{O}(N^3 logN)$ complexity. Increasing resolution is subject to the same $\mathcal{O}(N^3 logN)$ complexity as 3D convolution. Even the optimized interpolation method necessitates initial triangulation, which raises

Authorized licensed use limited to: Univ of Calif Berkeley. Downloaded on October 25,2023 at 23:49:16 UTC from IEEE Xplore. Restrictions apply.

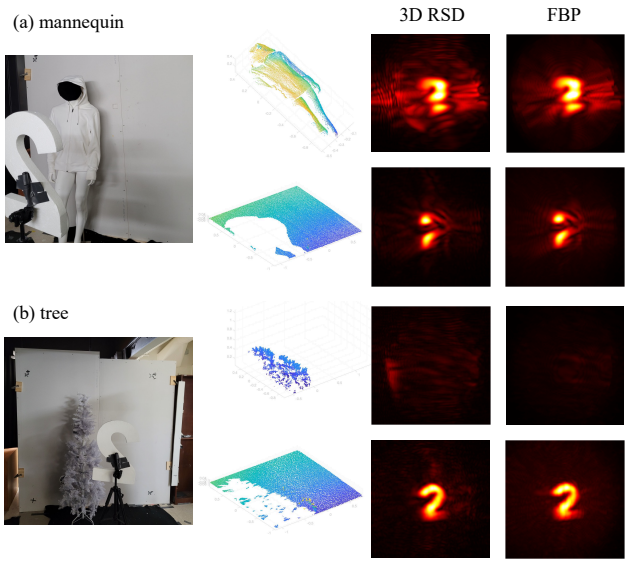

Fig. 3. NLOS reconstruction on segmented relay-surfaces. The sub-reconstructions using the same method are adjusted into the same intensity scale.

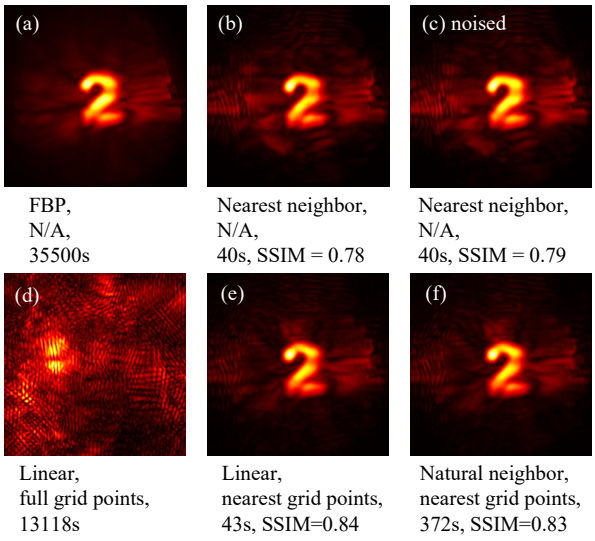

Fig. 4. The first line shows the interpolation method, the second line displays the query scheme, and the third line is the runtime alongside structure similarity with the FBP generated reconstruction.

computational complexity of the entire method to $\mathcal{O}(N^4)$, counteracting our aim. We demonstrate the sufficiency of our grid resolution by introducing random noise, uniformly distributed between $-d/2$ and $d/2$ (grid spacing d=1cm), to the measured positions. The reconstruction is unaffected as shown in Fig. 4 (c). Based on this analysis and our intuitive understanding we also optimize the grid size to have higher resolution in the depth dimension than lateral dimensions. Should an efficient and accurate interpolation algorithm emerge, it could be easily incorporated and allow reducing grid resolution.

6 Comparison to Alternative Methods and Discussion

The existing literature on NLOS reconstruction for non-uniformly sampled and non-planar relay surfaces is limited. To the best of our knowledge, only FBP is inherently capable of handling such situations. Moreover, data-driven methods require a time-consuming training stage for each specific relay surface, rendering them unsuitable for fast reconstruction purposes.

To establish the effectiveness and necessity of our 3D RSD algorithm, we conduct a comparative analysis with the two background algorithms in the literature, along with three other possibly intuitive implementations of the first stage propagation, as depicted in the bottom row of Fig. 1.

6.1 FBP-based methods

FBP can be employed either as the initial stage operator to propagate from the non-planar surface to an intermediate plane, or to propagate to the entire volume directly. In the former case, the computational complexity of the algorithm is $\mathcal{O}\left(N^4\right)$, as it has to examine every feasible pairing of two planes, each containing N^2 elements. The schematic is shown in Fig.1 (a). In the latter case, the computational complexity is $\mathcal{O}\left(N^5\right)$ because it has to repeat this for every reconstructed depth slice.

An intuitive observation is that when calculating the propagation kernel, the closest points contribute the most. So we introduce a modified version of mask FBP, as Fig. 1(b) depicts. We manually specify a radius threshold R, and only calculate the contribution from the source points inside the given radius ball. Eq. 4 is modified to:

$$I\left(\mathbf{x}_{\mathrm{v}}\right)=\iiint H\left(\mathbf{x}_{\mathrm{p}}, \mathbf{x}_{\mathrm{c}}, t\right) \delta\left(t-\frac{\left|\mathbf{x}_{\mathrm{p}}-\mathbf{x}_{\mathrm{v}}\right|+\left|\mathbf{x}_{\mathrm{c}}-\mathbf{x}_{\mathrm{v}}\right|}{c}\right) M\left(\left|\mathbf{x}_{\mathrm{c}}-\mathbf{x}_{\mathrm{v}}\right|\right) d \mathbf{x}_{\mathrm{p}} d \mathbf{x}_{\mathrm{c}} d t. \tag{10}$$

where the mask function is defined as:

$$M(x)=\begin{cases} 1, x \leqslant R \\ 0, x>R \end{cases} \tag{11}$$

The method has superior computational complexity to FBP, as it collapses the N^2 factor in the first stage to a constant parameter. Consequently, in theory, the computational complexity is solely dependent on the second stage of plane-to-plane RSD, and is scaled by $\mathcal{O}\left(N^3 log(N)\right)$.

6.2 FK-based methods

The FK migration paper has attempted to extend for the non-planar relay surfaces [16]. However, there are two significant challenges to its practical implementation: first, the FK migration method is designed for confocal measurement. The performance on non-confocal datasets significantly degrades after data conversion and pre-processing [15], [18], [19]. Therefore, it is not suitable for fast non-confocal data acquisition. Secondly, the proposed extrapolation technique for non-planar data pre-processing is computationally intensive. It requires N^2 times calculation of 3D FFT, which results in a computational complexity of $N^5 log(N)$. This complexity is worse than all the other non-planar algorithms,

Authorized licensed use limited to: Univ of Calif Berkeley. Downloaded on October 25,2023 at 23:49:16 UTC from IEEE Xplore. Restrictions apply.

making the FK migration with their proposed extrapolation method inapplicable for fast NLOS reconstruction.

In the FK paper [16], the authors present a more efficient approximation for non-planar data pre-processing, which is simply shifting the acquired histogram by the travel time corresponding to the depth offset of the relay surface. We adopt this idea and integrate it into the phasor algorithm context, resulting in a first stage propagator executed through phase mapping. Specifically, the intermediate plane is generated via a direct projection from the non-planar relay surface. This procedure involves the mapping of each $\mathrm{x_c}$ to its corresponding $\mathrm{x_t}$, which is achieved by introducing a phase shift proportional to the distance between the two points.

We map the non-planar points $\mathrm{x_c}$ to planar grid points $\mathrm{x_t}$ by:

$$\mathcal{P}\left(\mathbf{x}_{\mathrm{t}}, \omega\right) = \mathcal{P}\left(\mathbf{x}_{\mathrm{c}}, \omega\right) e^{i\frac{\omega}{c}|\mathbf{x}_{\mathrm{t}} - \mathbf{x}_{\mathrm{c}}|}. \tag{12}$$

The computational complexity is $\mathcal{O}(N^3 log(N))$, which is the same as RSD. However, this approximation is only valid if the relay wall has a modest level of deformation and the intermediate plane can be very close to the relay wall. Besides, it can not deal with non-uniform sampling in x and y axis caused by off-center calibration of non-planar geometry.

6.3 Performance comparison

To evaluate the performance of our proposed algorithm and compare it with alternative methods, we fix the hidden scene to be a single number '2' and set up four relay surfaces with varying degrees of non-planarity: (a) planar wall, (b) (c) wall with several big boards in front, and (d) wall with a walking mannequin in front. To avoid refocusing the SPAD array between scans, we fix the SPAD array to look at the same planar part of the relay wall for all four scans. This is done mostly for convenience and is not a requirement for the algorithm.

The comparison results for all the experiments are shown in Fig. 5. The first row in Fig. 5 shows the photo of relay surfaces. We apply the two one-stage algorithms in the literature as introduced in Sec. 3, three alternative two-stage algorithms, and our proposed 3D RSD algorithm to each of the dataset. We displayed the 3D reconstruction results using maximum intensity projection and labeled each reconstruction's runtime under the result. Providing quantitative metrics of reconstruction quality is difficult due to challenges acquiring hidden scene ground truth. However, as our algorithm approximates FBP, we list the structural similarity (SSIM) comparisons with FBP reconstructions for each algorithm, measured on normalized and projected reconstructions for visual perception alignment.

The comparison of time complexity and reconstruction performance is shown in Tab. 1. The memory complexity for all the algorithms is $O(N^3)$, although in practice previous work has demonstrated that the phasor field methods based on RSD require significantly less memory than other wave-based methods like F-K migration [17]. We put a question mark next to the performance of F-K migration method with full extrapolation as this algorithm was not executed on our datasets, primarily due to its high computational complexity. Furthermore, previous research [15] indicates that the F-K migration method yields inferior performance compared to FBP when applied to non-confocal datasets. Consequently, the performance of the F-K migration method is upper-bounded by the FBP method.

The comparison experiments on our datasets lead to the following conclusions: the direct application of the fast RSD algorithm to highly non-planar relay surfaces results in distorted reconstruction, indicating the knowledge of non-planar geometry is necessary. In contrast, the FBP method achieves faithful reconstruction with a substantial runtime of up to 10 hours. Separating the reconstruction process into two stages, with RSD applied in the second stage, leads to limited improvement. While using FBP solely for the first stage reduces the runtime by a factor of N and still produces clean results, it remains orders of magnitude slower than the second stage. Additionally, the mask FBP can only reduce runtime further by at most two times. The mask FBP method requires a sufficiently large mask radius to maintain reconstruction quality for non-planar geometries and keeping narrowing down the radius leads to severe artifacts. The FK-based phase mapping method is ineffective in correcting for non-planar geometry, since the depth offsets in our measured datasets are all significantly larger than the virtual wavelength. This method fails to correct for the non-planar geometry and yields outcomes that are almost identical to those produced by directly applying the planar RSD algorithm.

In general, the comparison in Fig. 5 reveals that our 3D RSD algorithm has superior performance compared with other alternative methods in two-stage reconstruction. Also, it has comparable runtime with planar fast RSD and generates faithful reconstruction comparable with the most universal FBP solver. The proposed algorithm has stable performance on non-planar datasets and non-uniform sampling and is orders of magnitude faster than the FBP, making it a promising technique for practical NLOS imaging applications.

Additionally, to prove our algorithm's versatility across various scenes, we demonstrate it on a fixed relay surface with varied scenes in Fig. 6.

TABLE 1
Complexity and performance comparison

Algorithm		Time complexity	Non-planar reconstruction quality
1-stage	FBP	$\mathcal{O}(N^5)$	high
	RSD	$\mathcal{O}(N^3 log(N))$	low
	FK	$\mathcal{O}(N^5 log(N))$	medium?
2-stage	FBP	$\mathcal{O}(N^4)$	high
	Mask FBP	$\mathcal{O}(N^3 log(N))$	medium
	Phase mapping	$\mathcal{O}(N^3 log(N))$	low
	3D RSD	$\mathcal{O}(N^3 log(N))$	high

Authorized licensed use limited to: Univ of Calif Berkeley. Downloaded on October 25,2023 at 23:49:16 UTC from IEEE Xplore. Restrictions apply.

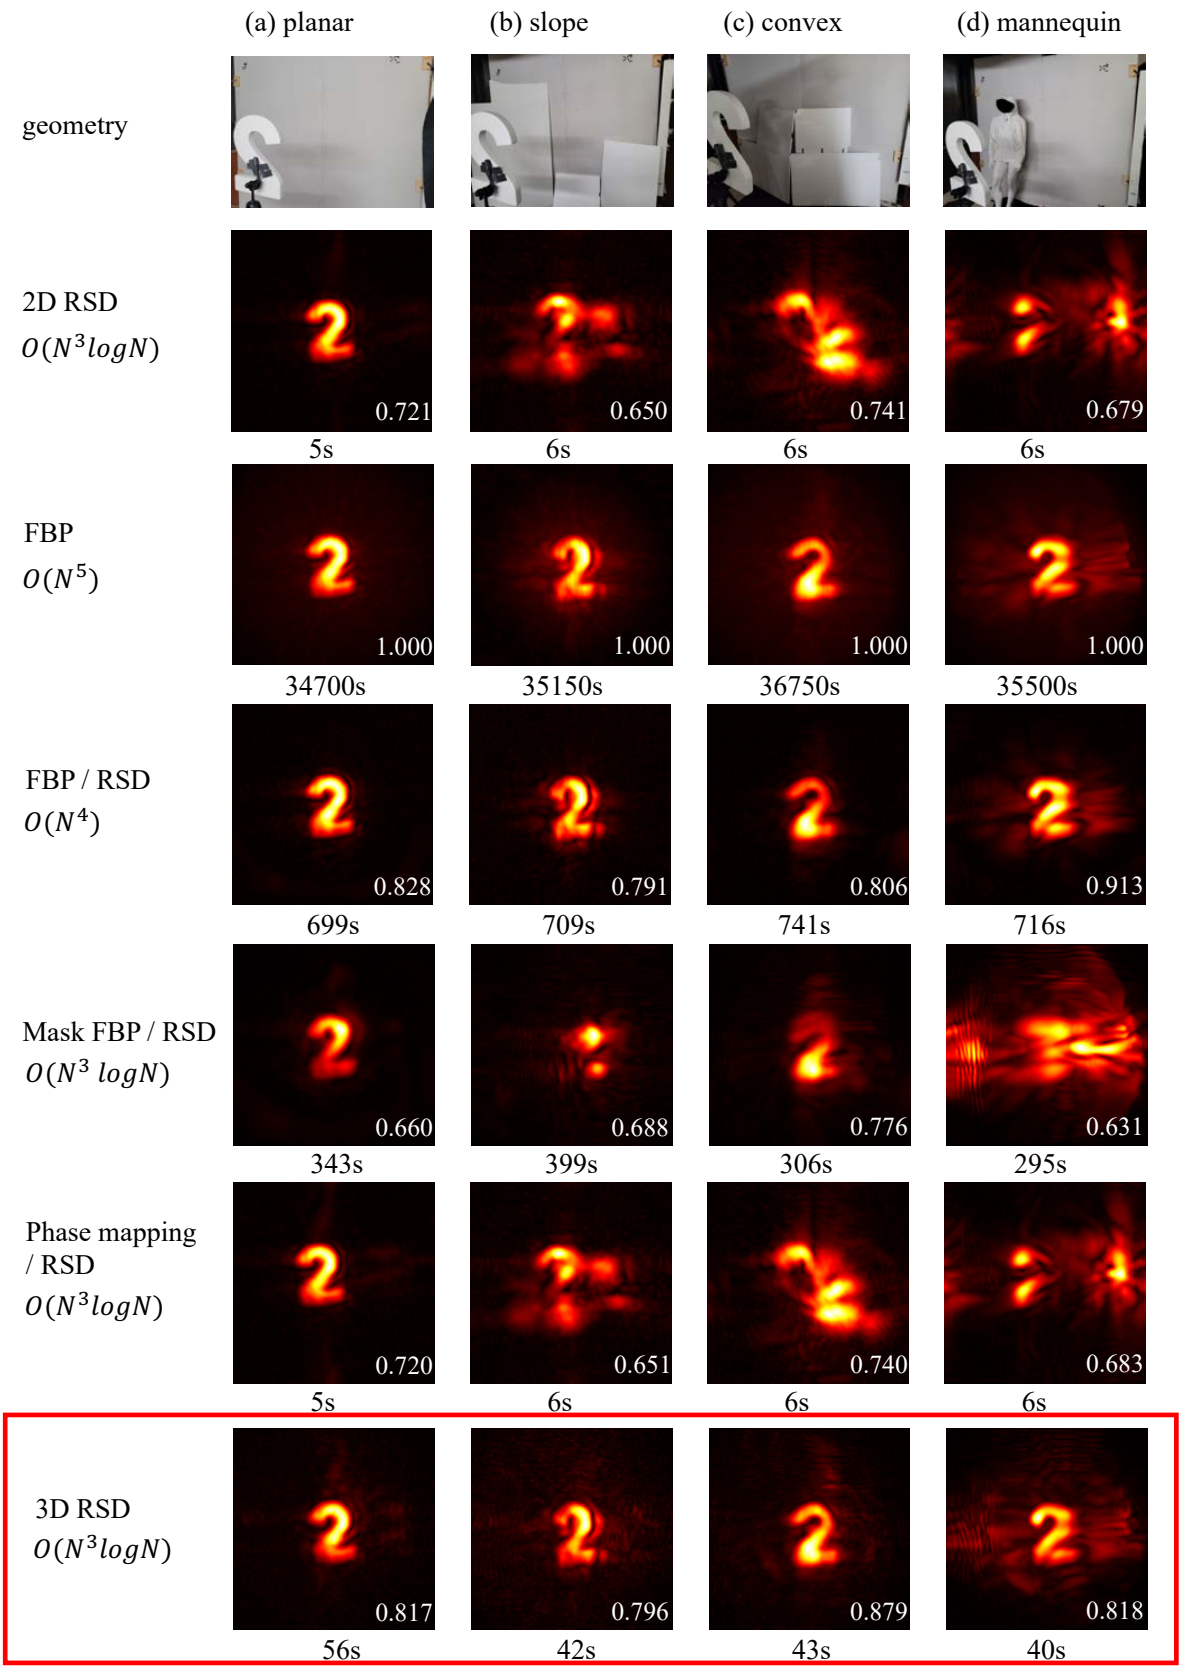

Fig. 5. Comparison between two one-stage and four two-stage algorithms on experimental non-confocal datasets with different relay surfaces. The real-world runtime is listed below, and the structure similarity (SSIM) with the FBP generated reconstruction is indicated in the bottom right corner of each reconstruction.

Authorized licensed use limited to: Univ of Calif Berkeley. Downloaded on October 25,2023 at 23:49:16 UTC from IEEE Xplore. Restrictions apply.

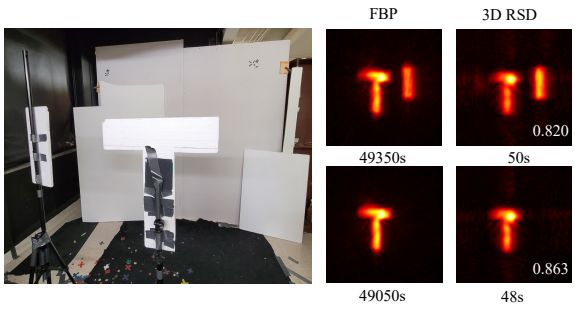

Fig. 6. 3D RSD reconstruction demonstrated on varied hidden scenes. The relay surface and an exemplary hidden scene featuring two letters at varying depths are displayed on the left.

7 Details on Data Acquisition and Processing

7.1 Set up hardware details

Our NLOS imaging system comprises an illumination source and detector, as shown in Fig. 7(a). Specifically, we employ a PM-1.03-25^{TM} laser from Polar Laser Laboratories as our illumination source, which is supplemented by a frequency doubler to generate 515 nm pulses with a maximum pulse width of 35 ps. The average power of the laser output is 375 mW, and it operates at an average repetition rate of 5 MHz. To scan the different relay surfaces, we use a two-mirror Thorlabs galvanometer (Thorlabs GVS012). While the total scan area varies, the total scan time is set to 36 seconds for each experiment. Our detector is a SPAD array [43] that is focused on a 7.1 cm by 4.7 cm area of the relay surface using a Canon EF 85 mm f/1.8 USM Lens. The temporal resolution of the SPAD array is characterized by a Full-Width at Half Maximum (FWHM) of approximately 60 ps and a deadtime of less than 100 ns.

To calibrate the 3D geometry of various relay surfaces, a secondary single-pixel, fast-gated SPAD module from MPD [44] is used to form a confocal LiDAR system with our PM-1.03-25^{TM} laser. A Nikon AF NIKKOR 50mm f/1.8D lens is used to focus the single pixel spad onto the relay surface, as shown in Fig. 2(b). Since the direct reflection from the relay surface is much brighter than the indirect reflection from the hidden scene, this confocal system can sufficiently infer the 3D coordinates of the relay surface with a 7s scan.

7.2 Algorithm details

The relay-surface has a aperture of roughly $2m \times 2m$, and is sampled at the distance of $1cm$. Since the sampling grid is projected to a non-planar surface, the scanning generates a 190×190 non-uniform sampling grid. Since the span of the 16×16 SPAD pixels on the relay surface is much smaller than the aperture, we simply add up all the data measured by all SPAD pixels to boost SNR. This is equivalent to treating the SPAD array as a defocused single pixel SPAD detector.

The sampling theorem gives that the least virtual wavelength for reconstruction is 2cm. To account for SPAD array span and improve SNR, we set the virtual wavelength to 6cm. The virtual illumination function $\mathcal{P}(\mathbf{x}_p, t)$ in Eq.1 is a wave packet of given frequency multiplied by a Gaussian window:

$$\mathcal{P}(\mathbf{x}_p, t) = e^{i\frac{2\pi}{\lambda_C}t} e^{-\frac{t^2}{2\sigma^2}} \tag{13}$$

where $\lambda_C = 6cm$ and $\sigma = 5 * \lambda_C/c$. The same filter is applied in FBP and mask FBP to keep consistency.

The hidden space is set to a 2 m $\times$ 2 m $\times$ 1 m cube sampled with 1 cm spacing. All the algorithms are implemented with MATLAB and run on an Apple M1 Pro CPU with 32 GB memory. The runtime of each algorithm is measured using MATLAB's timing function, with the exception of FBP, due to its high computational complexity. In the case of FBP, each reconstruction would have taken up to 10 hours with the same conditions. Therefore, we turn to a more efficient implementation using C++ [15] to generate the reconstruction results. To enable fair comparison, we calculated the theoretical time for FBP reconstruction under the same conditions, which is the product of the number of reconstructed depth slices and the runtime of the first-stage FBP. This theoretical time is labeled as the FBP reconstruction time.

8 Conclusion

Fast and accurate reconstruction of non-confocal measurements on non-planar relay surfaces is essential for practical and real-time deployment of NLOS imaging systems in the wild. By decomposing the reconstruction into two stages, this separation allows for a combination of non-planar and planar propagation algorithms. We leverage the fast RSD algorithm in the second stage and reduce the computational complexity of FBP by one order of magnitude without any quality loss. Furthermore, our proposed 3D RSD algorithm reduces the computational complexity even further to the same as fast RSD for planar relay surfaces. When FBP takes hours and fast RSD fails to generate useful reconstruction under highly non-planar scenarios, our proposed 3D RSD algorithm achieves comparable reconstruction quality to state-of-the-art FBP algorithms in comparable time to the existing fast algorithm for planar surfaces.

Acknowledgments

This work was supported by the National Science Foundation (1846884) and the Air Force Office for Scientific Research (FA9550-21-1-0341). The authors would like to thank ChatGPT for helping rephrase section 7.1 on Hardware details.

References

[1] A. Kirmani, T. Hutchison, J. Davis, and R. Raskar, "Looking around the corner using transient imaging," in *2009 IEEE 12th International Conference on Computer Vision*. Kyoto: IEEE, Sep. 2009, pp. 159–166. [Online]. Available: http://ieeexplore.ieee.org/document/5459160/

[2] A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, M. G. Bawendi, and R. Raskar, "Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging," *Nature Communications*, vol. 3, p. 745, Mar. 2012. [Online]. Available: https://www.nature.com/articles/ncomms1747

[3] K. L. Bouman, V. Ye, A. B. Yedidia, F. Durand, G. W. Wornell, A. Torralba, and W. T. Freeman, "Turning Corners into Cameras: Principles and Methods," in *2017 IEEE International Conference on Computer Vision (ICCV)*. Venice: IEEE, Oct. 2017, pp. 2289–2297.

[4] C. Saunders, J. Murray-Bruce, and V. K. Goyal, "Computational periscopy with an ordinary digital camera," *Nature*, vol. 565, no. 7740, pp. 472–475, Jan. 2019. [Online]. Available: http://www.nature.com/articles/s41586-018-0868-6

Authorized licensed use limited to: Univ of Calif Berkeley. Downloaded on October 25,2023 at 23:49:16 UTC from IEEE Xplore. Restrictions apply.

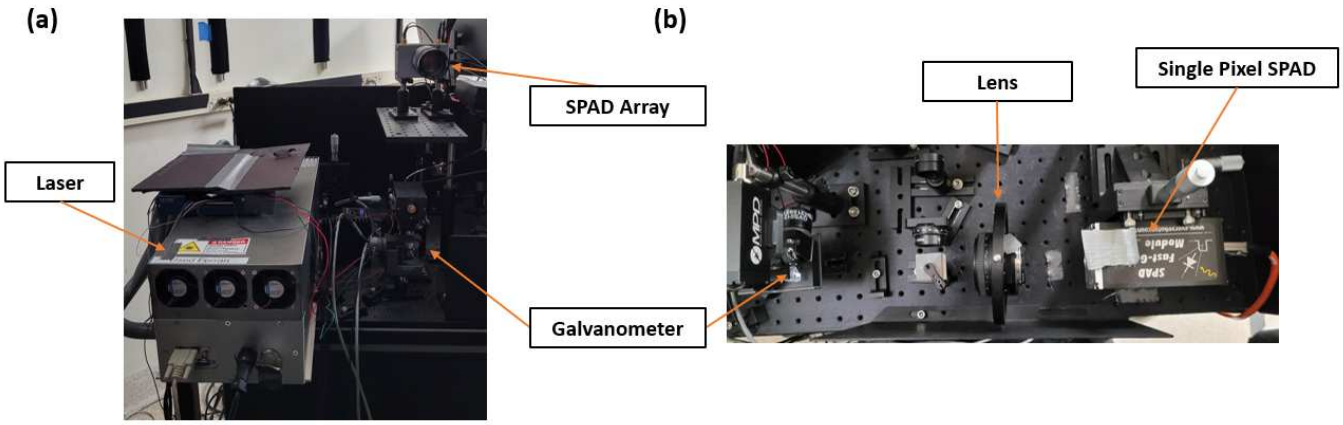

Fig. 7. Hardware set up used to validate our algorithm on experimental data. **(a)** Our NLOS imaging system consists of a picosecond pulsed laser, a galvanometer to scan the relay surface, and a 16 by 16 pixel SPAD array that uses a lens to focus on the relay wall. **(b)** Our confocal LiDAR system calibrates the coordinates of different relay surfaces. A single pixel SPAD uses a lens to focus on the laser position at the relay surface.

[5] J. Boger-Lombard and O. Katz, "Passive optical time-of-flight for non line-of-sight localization," *Nature Communications*, vol. 10, no. 1, p. 3343, Jul. 2019, number: 1 Publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/s41467-019-11279-6

[6] C. Thrampoulidis, G. Shulkind, F. Xu, W. T. Freeman, J. H. Shapiro, A. Torralba, F. N. C. Wong, and G. W. Wornell, "Exploiting Occlusion in Non-Line-of-Sight Active Imaging," *IEEE Transactions on Computational Imaging*, vol. 4, no. 3, pp. 419–431, Sep. 2018. [Online]. Available: https://ieeexplore.ieee.org/document/8345763/

[7] F. Xu, G. Shulkind, C. Thrampoulidis, J. H. Shapiro, A. Torralba, F. N. C. Wong, and G. W. Wornell, "Revealing hidden scenes by photon-efficient occlusion-based opportunistic active imaging," *Optics Express*, vol. 26, no. 8, pp. 9945–9962, Apr. 2018, tex.ids= xu_revealing_2018-1, xu_revealing_2018-2 publisher: Optical Society of America. [Online]. Available: https://www.osapublishing.org/oe/abstract.cfm?uri=oe-26-8-9945

[8] S. W. Seidel, Y. Ma, J. Murray-Bruce, C. Saunders, W. T. Freeman, C. C. Yu, and V. K. Goyal, "Corner Occluder Computational Periscopy: Estimating a Hidden Scene from a Single Photograph," in *2019 IEEE International Conference on Computational Photography (ICCP)*, May 2019, pp. 1–9, iSSN: 2472-7636.

[9] S. W. Seidel, J. Murray-Bruce, Y. Ma, C. Yu, W. T. Freeman, and V. K. Goyal, "Two-Dimensional Non-Line-of-Sight Scene Estimation From a Single Edge Occluder," *IEEE Transactions on Computational Imaging*, vol. 7, pp. 58–72, 2021, conference Name: IEEE Transactions on Computational Imaging.

[10] P. Rangarajan, F. Willomitzer, O. Cossairt, and M. P. Christensen, "Spatially resolved indirect imaging of objects beyond the line of sight," vol. 11135, p. 111350I, Sep. 2019, conference Name: Unconventional and Indirect Imaging, Image Reconstruction, and Wavefront Sensing 2019 ADS Bibcode: 2019SPIE11135E..0IR. [Online]. Available: https://ui.adsabs.harvard.edu/abs/2019SPIE11135E..0IR

[11] M. Baradad, V. Ye, A. B. Yedidia, F. Durand, W. T. Freeman, G. W. Wornell, and A. Torralba, "Inferring Light Fields from Shadows," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Jun. 2018, pp. 6267–6275, iSSN: 2575-7075.

[12] W. Chen, S. Daneau, F. Mannan, and F. Heide, "Steady-State Non-Line-Of-Sight Imaging," 2019, pp. 6790–6799. [Online]. Available: https://openaccess.thecvf.com/content_CVPR_2019/html/Chen_Steady-State_Non-Line-Of-Sight_Imaging_CVPR_2019_paper.html

[13] O. Katz, P. Heidmann, M. Fink, and S. Gigan, "Non-invasive single-shot imaging through scattering layers and around corners via speckle correlations," *Nature Photonics*, vol. 8, no. 10, pp. 784–790, Oct. 2014, number: 10 Publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/nphoton.2014.189

[14] W. Krska, S. W. Seidel, C. Saunders, R. Czajkowski, C. Yu, J. Murray-Bruce, and V. Goyal, "Double Your Corners, Double Your Fun: The Doorway Camera," in *2022 IEEE International Conference on Computational Photography (ICCP)*, Aug. 2022, pp. 1–12, iSSN: 2472-7636.

[15] X. Liu, I. Guillén, M. La Manna, J. H. Nam, S. A. Reza, T. Huu Le, A. Jarabo, D. Gutierrez, and A. Velten, "Non-line-of-sight imaging using phasor-field virtual wave optics," *Nature*, vol. 572, no. 7771, pp. 620–623, Aug. 2019, number: 7771 Publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/s41586-019-1461-3

[16] D. B. Lindell, G. Wetzstein, and M. O'Toole, "Wave-based non-line-of-sight imaging using fast *f-k* migration," *ACM Transactions on Graphics*, vol. 38, no. 4, pp. 1–13, Aug. 2019. [Online]. Available: https://dl.acm.org/doi/10.1145/3306346.3322937

[17] X. Liu, S. Bauer, and A. Velten, "Phasor field diffraction based reconstruction for fast non-line-of-sight imaging systems," *Nature Communications*, vol. 11, no. 1, pp. 1–13, Apr. 2020, number: 1 Publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/s41467-020-15157-4

[18] J. H. Nam, E. Brandt, S. Bauer, X. Liu, M. Renna, A. Tosi, E. Sifakis, and A. Velten, "Low-latency time-of-flight non-line-of-sight imaging at 5 frames per second," *Nature Communications*, vol. 12, no. 1, p. 6526, Dec. 2021. [Online]. Available: https://www.nature.com/articles/s41467-021-26721-x

[19] X. Liu, J. Wang, Z. Li, Z. Shi, X. Fu, and L. Qiu, "Non-line-of-sight reconstruction with signal–object collaborative regularization," *Light: Science & Applications*, vol. 10, no. 1, p. 198, Dec. 2021. [Online]. Available: https://www.nature.com/articles/s41377-021-00633-3

[20] J. Grau Chopite, M. B. Hullin, M. Wand, and J. Iseringhausen, "Deep Non-Line-of-Sight Reconstruction," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Seattle, WA, USA: IEEE, Jun. 2020, pp. 957–966.

[21] W. Chen, F. Wei, K. N. Kutulakos, S. Rusinkiewicz, and F. Heide, "Learned feature embeddings for non-line-of-sight imaging and recognition," *ACM Transactions on Graphics*, vol. 39, no. 6, pp. 1–18, Dec. 2020.

[22] https://github.com/ArianaGu/3D-RSD.

[23] M. Buttafava, J. Zeman, A. Tosi, K. Eliceiri, and A. Velten, "Non-line-of-sight imaging using a time-gated single photon avalanche diode," *Optics Express*, vol. 23, no. 16, pp. 20 997–21 011, Aug. 2015. [Online]. Available: https://www.osapublishing.org/oe/abstract.cfm?uri=oe-23-16-20997

[24] V. Arellano, D. Gutierrez, and A. Jarabo, "Fast back-projection for non-line of sight reconstruction," *Optics Express*, vol. 25, no. 10, pp. 11 574–11 583, May 2017. [Online]. Available: https://www.osapublishing.org/oe/abstract.cfm?uri=oe-25-10-11574

[25] M. L. Manna, F. Kine, E. Breitbach, J. Jackson, T. Sultan, and A. Velten, "Error Backprojection Algorithms for Non-Line-of-Sight Imaging," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2018.

Authorized licensed use limited to: Univ of Calif Berkeley. Downloaded on October 25,2023 at 23:49:16 UTC from IEEE Xplore. Restrictions apply.

[26] J. Iseringhausen and M. B. Hullin, "Non-line-of-sight Reconstruction Using Efficient Transient Rendering," *ACM Transactions on Graphics*, vol. 39, no. 1, pp. 8:1–8:14, Jan. 2020. [Online]. Available: https://doi.org/10.1145/3368314
[27] M. O'Toole, D. B. Lindell, and G. Wetzstein, "Confocal non-line-of-sight imaging based on the light-cone transform," *Nature*, vol. 555, no. 7696, pp. 338–341, Mar. 2018. [Online]. Available: https://www.nature.com/articles/nature25489
[28] S. A. Reza, M. L. Manna, S. Bauer, A. Velten, and A. Velten, "Phasor field waves: A Huygens-like light transport model for non-line-of-sight imaging applications," *Optics Express*, vol. 27, no. 20, pp. 29 380–29 400, Sep. 2019, publisher: Optical Society of America. [Online]. Available: https://www.osapublishing.org/oe/abstract.cfm?uri=oe-27-20-29380
[29] S. A. Reza, M. La Manna, S. Bauer, and A. Velten, "Phasor field waves: experimental demonstrations of wave-like properties," *Optics Express*, vol. 27, no. 22, p. 32587, Oct. 2019. [Online]. Available: https://www.osapublishing.org/abstract.cfm?URI=oe-27-22-32587
[30] J. A. Teichman, "Phasor field waves: a mathematical treatment," *Optics Express*, vol. 27, no. 20, pp. 27 500–27 506, Sep. 2019, publisher: Optical Society of America. [Online]. Available: https://www.osapublishing.org/oe/abstract.cfm?uri=oe-27-20-27500
[31] J. Dove and J. H. Shapiro, "Paraxial theory of phasor-field imaging," *Optics Express*, vol. 27, no. 13, pp. 18 016–18 037, Jun. 2019. [Online]. Available: https://www.osapublishing.org/oe/abstract.cfm?uri=oe-27-13-18016
[32] ——, "Nonparaxial phasor-field propagation," *Opt. Express*, vol. 28, no. 20, pp. 29 212–29 229, Sep 2020. [Online]. Available: https://opg.optica.org/oe/abstract.cfm?URI=oe-28-20-29212
[33] ——, "Paraxial phasor-field physical optics," *Optics Express*, vol. 28, no. 14, pp. 21 095–21 109, Jul. 2020, publisher: Optical Society of America. [Online]. Available: https://www.osapublishing.org/oe/abstract.cfm?uri=oe-28-14-21095
[34] ——, "Speckled speckled speckle," *Optics Express*, vol. 28, no. 15, pp. 22 105–22 120, Jul. 2020, publisher: Optical Society of America. [Online]. Available: https://www.osapublishing.org/oe/abstract.cfm?uri=oe-28-15-22105
[35] J. Marco, A. Jarabo, J. H. Nam, X. Liu, M. Á. Cosculluela, A. Velten, and D. Gutierrez, "Virtual Light Transport Matrices for Non-Line-of-Sight Imaging," in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct. 2021, pp. 2420–2429, iSSN: 2380-7504.
[36] M. L. Manna, J.-H. Nam, S. A. Reza, A. Velten, and A. Velten, "Non-line-of-sight-imaging using dynamic relay surfaces," *Optics Express*, vol. 28, no. 4, pp. 5331–5339, Feb. 2020, publisher: Optica Publishing Group. [Online]. Available: https://opg.optica.org/oe/abstract.cfm?uri=oe-28-4-5331
[37] P. W. M. Tsang, T.-C. Poon, and Y. M. Wu, "Review of fast methods for point-based computer-generated holography," *Photon. Res.*, vol. 6, no. 9, pp. 837–846, Sep 2018. [Online]. Available: https://opg.optica.org/prj/abstract.cfm?URI=prj-6-9-837
[38] D. Faccio, A. Velten, and G. Wetzstein, "Non-line-of-sight imaging," *Nature Reviews Physics*, vol. 2, no. 6, pp. 318–327, Jun. 2020.
[39] P. Sen, B. Chen, G. Garg, S. R. Marschner, M. Horowitz, M. Levoy, and H. P. A. Lensch, "Dual photography," in *ACM SIGGRAPH 2005 Papers*, ser. SIGGRAPH '05. New York, NY, USA: Association for Computing Machinery, 2005, p. 745–755. [Online]. Available: https://doi.org/10.1145/1186822.1073257
[40] M. Born and E. Wolf, *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Elsevier, 2013.
[41] D. Attali, J.-D. Boissonnat, and A. Lieutier, "Complexity of the delaunay triangulation of points on surfaces the smooth case," in *Proceedings of the nineteenth annual symposium on Computational Geometry*, 2003, pp. 201–210.
[42] J. W. Goodman, *Introduction to Fourier Optics*, 2017.
[43] S. Riccardo, E, Conca, V. Sesta, A. Velten, and A. Tosi, "Fast-Gated 16 × 16 SPAD Array with 16 on-chip 6 ps Time-to-Digital Converters for Non-Line-of-Sight Imaging," *IEEE Sensors Journal*, pp. 1–1, 2022, conference Name: IEEE Sensors Journal.
[44] M. Buttafava, G. Boso, A. Ruggeri, A. Dalla Mora, and A. Tosi, "Time-gated single-photon detection module with 110 ps transition time and up to 80 MHz repetition rate," *Review of Scientific Instruments*, vol. 85, no. 8, p. 083114, Aug. 2014, publisher: American Institute of Physics. [Online]. Available: https://aip.scitation.org/doi/10.1063/1.4893385

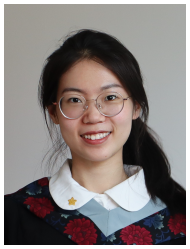

Chaoying Gu is a Ph.D. student in the Department of Electrical Engineering and Computer Sciences at UC Berkeley. She received B.S. in Electrical Engineering from Peking University. Her research interest lies in computational optics and imaging, especially the robust and efficient reconstruction for the inverse problems.

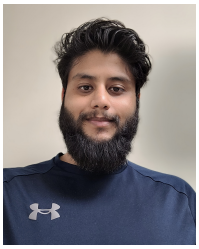

Talha Sultan (Member, IEEE) is working toward the Ph.D. degree in Electrical and Computer Engineering Department University of Wisconsin–Madison, Madison, WI. He received his BS degree in chemical engineering & applied mathematics from the University of Wisconsin–Madison, Madison, WI. His current research interests include computational optics and computational imaging. He specifically focuses on utilizing emerging sensing technologies to build non-line-of-sight imaging systems.

Khadijeh Masumnia-Bisheh (Member, IEEE) received the B.S. and M.S. degrees in electrical engineering from Babol Noshirvani University of Technology, Babol, Iran, in 2011 and 2015, and the Ph.D. degree in electrical engineering from Tarbiat Modares University, Tehran, Iran, in 2019. In 2018, she was a Visiting Scholar with the University of Utah. From 2020 to 2022, she was a Lecturer with the University of Science and Technology of Mazandaran, Behshahr, Iran. In April 2022, she joined the Computational Optics Group, University of Wisconsin-Madison, as a Postdoctoral Researcher. Her research interests include applied computational optics and imaging, computational electromagnetics, bioelectromagnetics, wave propagation and scattering in random media.

Laura Waller is a Professor of Electrical Engineering and Computer Sciences at UC Berkeley. She received B.S., M.Eng. and Ph.D. degrees from MIT and was a Postdoc/Lecturer of Physics at Princeton University. She is a Packard Fellow, Moore Foundation Data-driven Investigator, Bakar Fellow, OSA Fellow, AIMBE Fellow and Chan-Zuckerberg Biohub Investigator. She received the Carol D.Soc Distinguished Graduate Mentoring Award, OSA Adolph Lomb Medal, Ted Van Duzer Endowed Professorship, NSF CAREER Award and SPIE Early Career Achievement Award.

Andreas Velten received the Ph.D. degree in physics from the University of New Mexico, Albuquerque. He was a Postdoctoral Associate of the MIT Media Laboratory, Camera Culture Group. He is an Associate Professor at the Department of Biostatistics and Medical Informatics and the Department of Electrical and Computer Engineering, University of Wisconsin-Madison and directs the Computational Optics Group. He has included in the MIT TR35 list of the world's top innovators under the age of 35 and is a Senior Member of NAI, OSA, and SPIE and a member of Sigma Xi. He is the Co-Founder of Onlume, a company that develops surgical imaging systems, and Ubicept, a company developing single photon imaging solutions.

Authorized licensed use limited to: Univ of Calif Berkeley. Downloaded on October 25,2023 at 23:49:16 UTC from IEEE Xplore. Restrictions apply.