

Exploiting Occlusion in Non-Line-of-Sight Active Imaging

Christos Thrampoulidis , Gal Shulkind , Feihu Xu, William T. Freeman , Jeffrey H. Shapiro , Antonio Torralba, Franco N. C. Wong, and Gregory W. Wornell 

Abstract—Active non-line-of-sight imaging systems are of growing interest for diverse applications. The most commonly proposed approaches to date rely on exploiting time-resolved measurements, i.e., measuring the time it takes for short-duration light pulses to transit the scene. This typically requires expensive, specialized, ultrafast lasers, and detectors that must be carefully calibrated. We develop an alternative approach that exploits the valuable role that natural occluders in a scene play in enabling accurate and practical image formation in such settings without such hardware complexity. In particular, we demonstrate that the presence of occluders in the hidden scene can obviate the need for collecting time-resolved measurements, and develop an accompanying analysis for such systems and their generalizations. Ultimately, the results suggest the potential to develop increasingly sophisticated future systems that are able to identify and exploit diverse structural features of the environment to reconstruct scenes hidden from view.

Index Terms—Computational imaging, computer vision, LIDAR, non-line-of-sight imaging, time-of-flight cameras.

I. INTRODUCTION

IN CONTRAST to classical photography, where the scene of interest is in the observer's direct line of sight, non-line-of-sight (NLOS) imaging systems only have indirect access to a scene of interest via reflections from intermediary surfaces.

Manuscript received October 30, 2017; revised February 23, 2018; accepted April 7, 2018. Date of publication April 24, 2018; date of current version August 13, 2018. This work was supported in part by the DARPA REVEAL program under Contract HR0011-16-C-0030 and in part by the NSF under Grant CCF-1717610. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Orazio Gallo. (*Christos Thrampoulidis and Gal Shulkind contributed equally to this work.*) (*Corresponding author: Christos Thrampoulidis.*)

C. Thrampoulidis, G. Shulkind, J. H. Shapiro, F. N. C. Wong, and G. W. Wornell are with the Department of Electrical Engineering and Computer Science and the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: cthrampo@mit.edu; shulkind@mit.edu; jhs@mit.edu; ncw@mit.edu; gww@mit.edu).

F. Xu is with the Department of Electrical Engineering and Computer Science and the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA, and also with the Hefei National Laboratory for Physical Sciences at the Microscale, University of Science and Technology of China, Shanghai 201315, China (e-mail: fxu@mit.edu).

W. T. Freeman is with the Department of Electrical Engineering and Computer Science and the Computer Science and Artificial Intelligence Laboratory Massachusetts Institute of Technology, Cambridge, MA 02139 USA, and also with the Google Research, Cambridge, MA 02139 USA (e-mail: billf@mit.edu).

A. Torralba is with the Department of Electrical Engineering and Computer Science and the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: torralba@mit.edu).

Digital Object Identifier 10.1109/TCL.2018.2829599

Such systems are of considerable interest for applications spanning a wide variety of fields including medicine, manufacturing, transportation, public safety, and basic science.

Despite their obvious appeal, there are inherent challenges in the design of NLOS systems. In particular, typical surfaces (e.g., walls, floors, etc.) diffusely reflect light, effectively removing beam orientation information, and rendering the problem of scene reconstruction poorly conditioned. In order to compensate for the losses induced by diffuse reflections, initial demonstrations of NLOS imaging used ultrafast transient-imaging modalities [1], [2] that involved a laser source to send optical pulses of sub-picosecond duration, and a streak camera exhibiting temporal resolution in the picosecond range. A computational algorithm then used the fine time-resolved light intensity measurements to form a three-dimensional reconstruction of the hidden scene.

The system requirements posed by these systems, for transmission of very short-duration, high power optical pulses on the transmitter side, and for very high temporal resolution on the receiver side, inevitably imply high system complexity and cost. Thus, much of the follow-up work has focused on developing reduced cost and power implementations. For example, [3] uses a single-pixel, single-photon avalanche diode (SPAD) detector for reduced power consumption and cost; [4] uses a multi-pixel SPAD camera to demonstrate tracking of hidden moving objects; and [5] uses modulated illumination and low temporal-resolution CMOS time-of-flight sensors, including photonic mixer devices, to substantially reduce overall system cost, albeit at the expense of impairing the spatial resolution of the reconstruction.

A. Our Contribution

To address the limitations of such existing approaches, we introduce a rather different imaging modality for such problems. In particular, we develop the beneficial role that natural occlusions—which would traditionally be viewed as an impediment to imaging—can play in facilitating robust image reconstruction in NLOS settings. In fact, we demonstrate—analytically and experimentally—that in some cases the presence of occluders in the hidden scene can obviate the need for collecting time-resolved (TR) measurements, enabling imaging systems of significantly reduced cost. In turn, and in contrast to existing methods, this means our approach is compatible with wide field-of-view detectors, enabling the collection of

more photons per measurement and thus accelerating acquisition times so as to facilitate real-time operation.

We introduce the key concepts and principles in the context of imaging a hidden wall of unknown reflectivity. For this problem, we develop a framework of analysis that involves a mathematical formulation, as well as numerical and experimental illustrations. We further study diverse features of the proposed occlusion-based imaging system, such as robustness to modeling errors, and optimal selection of measurements. More generally, the ideas that we introduce open opportunities in designing more accurate, robust, and cost-effective NLOS imaging systems that relax the stringent temporal resolution requirements for optical measurements in the presence of occluders. We envision that our results will motivate further research towards the development of NLOS imaging systems that opportunistically exploit known structural features in the environment, such as occluders.

B. Related Work

To the best of our knowledge, this paper is the first to exploit the presence of occluders for high-resolution reconstruction of hidden-surface reflectivity from measurements of diffuse reflections. However, there is a variety of related work in computational imaging that investigates exploiting physical structure in the space between the scene of interest and the measurement system. Perhaps the best known is what is referred to as “coded-aperture imaging,” in which occlusion in the optical path takes the form of a carefully designed physical mask that modulates the light transferred from the scene of interest to a detector array. Among the earliest and simplest instances of coded-aperture imaging are those based on pinhole structure [6] and pinspeck (anti-pinhole) structure [7], though more complex structure is commonly used. Such methods are of particular interest in applications where lens fabrication is infeasible or impractical, such as in x-ray and gamma-ray imaging. More generally, a number of rich extensions to the basic methodology have been developed; see, e.g., [8] and the references therein.

In other developments, the value of using a mask in conjunction with a lens has been investigated in computational photography for motion deblurring [9], depth estimation [10], and digital refocusing and recovery of 4D light-fields [11]. More recently, there has been an increased interest in using masks with appropriate computational techniques, instead of traditional lens-based cameras, to build cameras that have fewer pixels, need not be focused [12], and/or meet physical constraints [13]. All these methods are passive imagers; only very recently has the addition of an active illumination source and time-resolved sensing been proposed to reduce acquisition time in lensless systems [14].

Perhaps the work most closely related to the present paper is that demonstrating how information about a scene outside the direct field of view can be revealed via “accidental” pinhole or anti-pinhole camera images [15]. The accidental camera is based on the use of video sequences obtained only with ambient illumination, and requires a reference frame without the occluder present. Another similar study has very recently demonstrated the ability to use an occluding wall edge to deduce a hidden subject’s pattern of motion [16]. While also relying on the

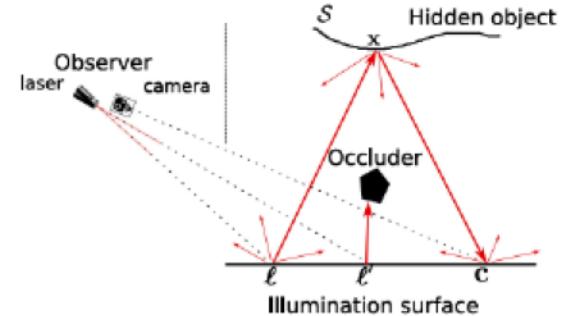


Fig. 1. Red lines trace beam paths reflecting from the virtual laser points ℓ , ℓ' , where a laser beam hits the illumination surface towards point x on the hidden object. The beam emanating from ℓ' is blocked by the occluder. Upon hitting the point x light reflects back towards a virtual camera position c , where a focused camera is steered.

presence of occlusions (albeit, specific occluding patterns induced by wall edges) for NLOS imaging, the results in [16] are limited to one-dimensional tracking of a hidden object that is assumed to be *moving*. In effect, the present paper can be viewed as quantifying the high-resolution imaging performance achievable without the aforementioned limitations, and in particular when we actively illuminate the scene with a scanning laser.

Finally, there have been recent demonstrations of a method for tracking moving objects in NLOS scenes via non-time-resolved intensity images of a visible wall [17]. By contrast, our framework emphasizes imaging without requiring the presence—and exploitation—of scene motion, so it can be applied much more broadly.

C. Paper Organization

The paper is organized as follows. Section II introduces a forward propagation model for NLOS imaging that accounts for sources of occlusion, and Section III introduces an analysis framework for NLOS imaging in the presence of such occlusion. Section IV then establishes the limitations of time-resolved measurements with respect to the temporal resolution of the detector, and Section V shows how to transcend these limitations by opportunistically exploiting occluders in the hidden scene. An experimental demonstration of the methodology is presented in Section VI. Finally, Section VII contains a discussion of extensions and opportunities for future research.

II. FORWARD MODEL FOR NLOS IMAGING

The goal of NLOS imaging systems is to process reflected light-intensity measurements and perform joint estimation of both the geometry and reflectivity properties of a hidden three-dimensional scene, as illustrated in Fig. 1. A focused laser beam is steered towards a visible *illumination surface* and reflects back towards a hidden object. Upon hitting the object light is reflected back towards the *illumination surface* and is measured by a focused camera. This forms a three-bounce problem in which light beams follow paths of the form

$$\text{Laser} \rightarrow \ell \rightarrow x \rightarrow c \rightarrow \text{Camera},$$

where ℓ, c lie on the illumination surface and x lies on the hidden object surface. By raster scanning the laser and/or changing the focal point of the camera, we retrieve multiple measurements corresponding to a set of K parameters $\mathcal{P} = \{(\ell_i, c_i)\}_{i=1,\dots,K}$.

In this section we formulate a forward propagation model that determines the irradiance waveform $y_{\ell,c}(t)$ measured at point c on the illumination surface in response to a single optical laser pulse $p(t)$ fired towards position ℓ . We let \mathcal{S} be a parametrization of the hidden object surface, and $f(x)$, $x \in \mathcal{S}$ denote the spatially varying reflectivity function (or, albedo). The model assumes that the illumination and hidden object surfaces are both ideal Lambertian reflectors.

In order to account for the presence of occluders in the scene (as illustrated in Fig. 1), we introduce a binary *visibility function* $V(x, z)$ which determines whether point x on the hidden object surface \mathcal{S} and point z on the illumination surface are visible to each other:

$$V(x, z) = \begin{cases} 1, & \text{clear line of sight between } x \text{ and } z, \\ 0, & \text{no line of sight between } x \text{ and } z. \end{cases} \quad (1)$$

With these, the forward model is given as follows¹:

$$y_{\ell,c}(t) = \int_{\mathcal{S}} f(x) \frac{V(x, \ell) V(x, c)}{\|x - \ell\|^2 \|x - c\|^2} G(x, \ell, c) \cdot p\left(t - \frac{\|x - \ell\| + \|x - c\|}{c}\right) dx. \quad (2)$$

Here, G is the Lambertian Bidirectional Reflectance Distribution Function (BRDF):

$$G(x, \ell, c) \equiv \cos(x - \ell, n_{\ell}) \cos(x - \ell, n_x) \cdot \cos(x - c, n_x) \cos(x - c, n_c),$$

n_x, n_c, n_{ℓ} are the surface normals at x, c, ℓ , respectively and c is the speed of light. The model can easily be generalized to account for non-Lambertian BRDFs for the illumination wall and the hidden object by appropriately adjusting G .

Several remarks are germane with respect to (2):

Virtual laser and camera positions: For simplicity in the exposition we have excluded from the model the attenuation, delay, and BRDF contributions accrued along the path from the laser to ℓ as well as those accrued from c to the camera. Note that those quantities are fixed and known to the observer, hence can be easily accounted for. In general, it is useful for our exposition to think of ℓ and c as *virtual unfocused illumination and camera positions* (Fig. 1), and (2) is consistent with that interpretation.

Visibility functions: The visibility functions in (2) account for obstructions of light beams in the imaging process, identifying hidden-object patches that are either not reached by the virtual illumination from ℓ or are not observable by the virtual camera at c . Implicit in this description is the partition of the objects occupying the space facing the illumination wall into: (a) the *hidden objects*, which are objects of interest in the reconstruction

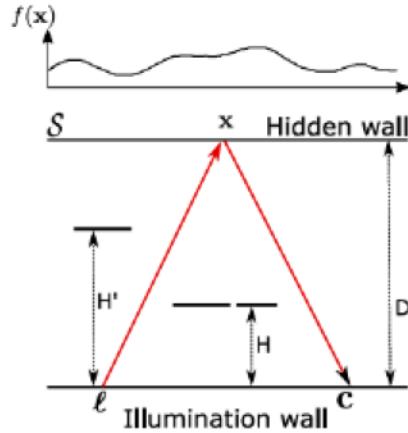


Fig. 2. The proposed imaging setting in which the objective is to reconstruct the reflectivity $f(x)$ of a flat hidden wall that is parallel to the illumination wall at known distance D . The positions and sizes of the fully absorbing occluders are known.

process; (b) the *occluders*, which are not of immediate interest (in fact, we usually assume that they are known), blocking at least some light paths between the illumination and hidden-object surfaces.

Third-bounces: The model (2) accounts for the contributions in the measurements resulting from three bounces (at ℓ, x, c) that are informative about the hidden objects. Higher-order bounces are neglected, since they typically experience high attenuation in the setting considered. Also, in deriving (2) we model the occluders as fully absorbing objects.²

Temporal resolution of the camera: The camera averages the incident irradiance at c with a finite temporal resolution Δt resulting in measurements $y_{\ell,c,\tau}, \tau = 1, 2, \dots, T$,

$$y_{\ell,c,\tau} = \int_{(\tau-1)\Delta t}^{\tau\Delta t} y_{\ell,c}(t) dt. \quad (3)$$

Since only third-bounce reflections involving the hidden object are of interest to us, with some abuse of notation we shift the time axis such that time $t = 0$ is the first instant when third bounce reflections reach the camera and $T\Delta t$ is chosen such that all relevant third-bounce reflections from the hidden object are included in the interval $[0, T\Delta t]$.

III. SCENE AND SYSTEM MODEL

To develop the key principles of approach, we turn to a specific instance of the general NLOS imaging scenario described in the previous section (also, Fig. 1), which we now describe. Extensions are discussed in Section VII.

A. Representative NLOS Imaging Setting

Our setup is illustrated in Fig. 2. It includes a planar hidden object and a parallel planar illumination surface, which we refer to as the *hidden wall* and the *illumination wall*, respectively. These two surfaces of known geometry are placed distance D apart.

¹A similar forward model is used in [5], and is based on well-known principles, namely quadratically decaying power with propagation distance for optical beams, and Lambert's cosine law for diffuse reflection. Equation (2) further accounts for possible occlusions in the scene through the visibility function.

²This model also applies for reflective occluders of known reflectivity pattern since their contribution in the measurements can be accounted for.

In between the illumination and the hidden walls lie occluders, whose effects on the imaging process are captured through the visibility function defined in (1). The occluders are fully absorbing objects of known geometries and locations. Hence, their visibility functions are known. The NLOS imaging objective under this setting is then to reconstruct the unknown reflectivity function $f(x)$ of the hidden wall from the measurements.

Under the aforementioned setting, the measurements $y_{\ell,c,\tau}$ in (3) are *linear* in the unknown reflectivity function $f(x)$. Let x_1, \dots, x_N be a discretization of the hidden wall, then, according to (2), each measurement $y_{\ell,c,\tau}$ corresponds to a measurement vector $a_{\ell,c,\tau} \in \mathbb{R}^N$ such that $y_{\ell,c,\tau} = a_{\ell,c,\tau}^\top f$, where $f := [f(x_1), \dots, f(x_N)]^\top$. Repeating the measurements for a total of K (ℓ, c) pairs, obtaining T time samples per each pair, and collecting them in a vector y of dimension $M = K \cdot T$, this gives rise to the linear system of equations $y = Af$ where A is an $M \times N$ measurement matrix whose rows are the vectors $a_{\ell,c,\tau}^\top$ that correspond to the chosen (ℓ, c) pairs and temporal resolution Δt . In this study we consider measurements that are contaminated by additive noise ϵ :

$$y = Af + \epsilon. \quad (4)$$

The noise term can be thought of as a simple means to capture system modeling errors, camera quantization errors, background noise, etc. We introduce the physics-based noise model for low-photon-count operation with a SPAD detector in [18].

B. Bayesian Scene Model

The idea of imposing Bayesian priors is well-established in image processing [19], [20]: past studies have considered various forms of Gaussian prior distributions on the unknown target scene, including variations promoting sparse derivatives [10], and natural image statistics [21]. Such priors offer enough flexibility and at the same time are amenable to analysis and intuitive interpretation. In this work, we let³:

$$f \sim \mathcal{N}(0, \Sigma_f), \quad (5)$$

with a smoothness promoting kernel function such that the entries of the covariance matrix are $[\Sigma_f]_{ij} = \exp(-\frac{1}{2\pi\sigma_f^2} \|x_i - x_j\|^2)$ and the spatial variance σ_f^2 controls the extent of smoothness. Additionally, we consider an i.i.d. Gaussian distribution for the measurement noise $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$ such that the Signal-to-Noise Ratio in our problem is given by $\text{SNR} = \text{tr}(A\Sigma_f A^\top) / (M\sigma^2)$, where M denotes the total number of measurements. For the reconstruction, we consider the minimum mean-squared error (MMSE) estimator, which under the Gaussian framework is explicitly computable as

$$\hat{f} = \Sigma_f A^\top (A\Sigma_f A^\top + \sigma^2 I)^{-1} y. \quad (6)$$

³The zero-mean assumption is somewhat simplistic, but not particularly restrictive. Strictly speaking, in order to respect the nonnegative nature of the reflectivity function, a positive mean should be added in all the scenarios considered in this paper. Similarly, a global scaling can be applied to ensure reflectivity values that are not greater than 1. However, these additions have no effect on the qualitative conclusions drawn from our results. This is further validated by the successful use of the Gaussian prior in the experimental demonstrations of Section VI.

We measure and compare reconstruction performance in different settings using the normalized mean squared error $\text{NMSE} = E\|\hat{f} - f\|_2^2 / E\|f\|_2^2$, which equals the (normalized) trace of the posterior covariance matrix

$$\text{NMSE} = \frac{1}{M} \text{tr}(\Sigma_f - \Sigma_f A^\top (A\Sigma_f A^\top + \sigma^2 I)^{-1} A\Sigma_f).$$

Note that the NMSE can be evaluated before collecting measurements y . Also, the reconstruction in (6) remains the optimal linear estimator under given first and second order statistics for f , even beyond Gaussian priors.

IV. TIME-RESOLVED MEASUREMENTS

In this section we study the limits of traditional NLOS imaging that is based on collecting fine time-resolved (TR) measurements, and thus set up a reference against which to compare the newly proposed imaging modality that uses occlusions and no TR information, which we formally introduce in Section V.

A. Virtues of Time-Resolved Measurements

Assuming an ideal pulse $p(t) = \delta(t)$, and considering the propagation of optical pulses at the speed of light c , the measurement $y_{\ell,c,\tau}$ taken at time step τ forms a linear combination of the reflectivity values of only those scene patches x_i whose sum distance to ℓ and c corresponds to a propagation time around $\tau\Delta t$. These patches fall within the elliptical annulus with focal points ℓ and c described by the following inequalities:

$$(\tau - 1) \cdot c\Delta t \leq \|x_i - \ell\| + \|x_i - c\| \leq \tau \cdot c\Delta t.$$

The thinner the annulus (eqv. the lower Δt), the more informative the measurements are about the reflectivity values of these patches. Furthermore, by scanning the laser and camera positions (ℓ, c) , different sets of light paths are probed, each generating a different set of elliptical annuli. For a total of K (ℓ, c) -pairs, this forms the linear system of (4).

We performed a simple numerical simulation to demonstrate scene reconstruction performance in a TR setup. For the purposes of illustration the simulations presented here and in the rest of the paper are in a two-dimensional world. This allows for easy visualization of important concepts such as the visibility function and the forward measurement operator, and it enables useful insights, but is otherwise non-restrictive. The room size was set such that the width of the walls is 1 m, the distance between the walls is $D = 2$ m and the temporal resolution was set at $\Delta t = 100$ ps. $K = 8$ (ℓ, c) pairs were randomly chosen, f was drawn according to the Gaussian prior with $\sigma_f^2 = 0.1$, and we set $\text{SNR} = 13.7$ dB. The results are summarized in Fig. 3, where we plot the measurement matrix A , the true reflectivity f and the estimated \hat{f} with the corresponding reconstruction uncertainty depicted in shaded color around the MMSE estimator. The reconstruction uncertainty for our purposes is the square-root of the diagonal entries in the posterior covariance matrix corresponding to the standard deviation of $\hat{f}_i - f_i$ for the individual points i on the wall. For this setup and resolution we collect $T = 16$ temporal samples per (ℓ, c) pair such that the total number of measurements is $M = 8 \cdot 16 = 128$. These

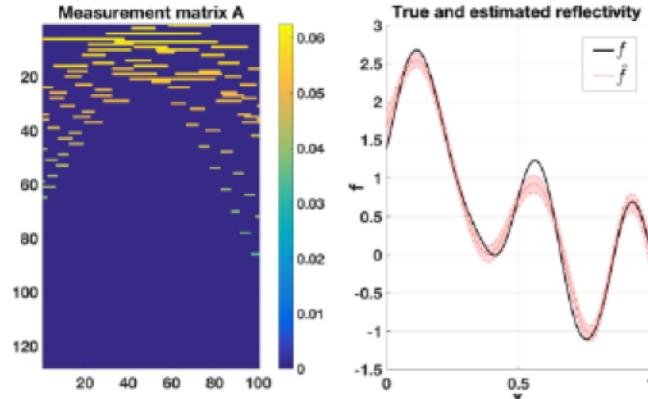


Fig. 3. Reflectivity reconstruction from TR measurements. (left) Measurement matrix, where each row corresponds to a specific choice for the (ℓ, c) pair and time index τ . The columns correspond to a discretization of the hidden wall to $N = 100$ points. (right) True reflectivity function versus the MMSE estimate \hat{f} .

are the rows of \mathbf{A} depicted in the figure, where each block of 8 consecutive rows corresponds to the measurements collected at a single time instant and for all (ℓ, c) pairs. Notice that the last few blocks are zero as at that time no patch on the hidden wall contributes to the measurements.

B. Performance Dependence on Temporal Resolution

The simulation results shown in Fig. 3 demonstrate high-fidelity reflectivity reconstruction when the available temporal resolution is fine ($\Delta t = 100$ ps). We show next that reconstruction fidelity depends on having access to measurements with high enough temporal resolution, and that it deteriorates substantially with lower temporal resolution of the measurements. As such, when we only have access to low temporal-resolution measurements, reconstruction fidelity may be severely limited.

Let us first consider an extreme situation where the temporal resolution is so low that the distance that light travels during a single resolution window of the detector is longer than the entire spatial extent of the scene. As an example, for the setup in Fig. 3 this happens when $\Delta t \gtrsim 1.5$ ns. In this extreme case, which is essentially equivalent to collecting *non-time-resolved measurements*, each (ℓ, c) -pair effectively generates just a single *scalar* measurement which we denote $\mathbf{y}_{\ell, c}$ and which is a linear combination of *all* the entries of \mathbf{f} . The combination coefficients are determined by the decay and cosine factors in (2). Focusing on the distance factors $\|\mathbf{x} - \ell\|^{-2} \|\mathbf{x} - \mathbf{c}\|^{-2}$ for intuition, the range of values that these can take is clearly determined by the geometry of the problem, and can be very limited; for example, if the two walls are far apart. This weak variation can result in poor conditioning of the measurement matrix \mathbf{A} and subsequently poor reconstruction fidelity.

This ill-conditioning is illustrated in Fig. 4(a). Here, for the Fig. 3 setup, we plot the NMSE versus the temporal resolution for $K = 30$ measurements and various SNR values, where for each data point we average over 10 random draws for (ℓ, c) . Observe that as Δt deteriorates, reconstruction fidelity decreases. Considering finite SNR for the purpose of this evaluation is key

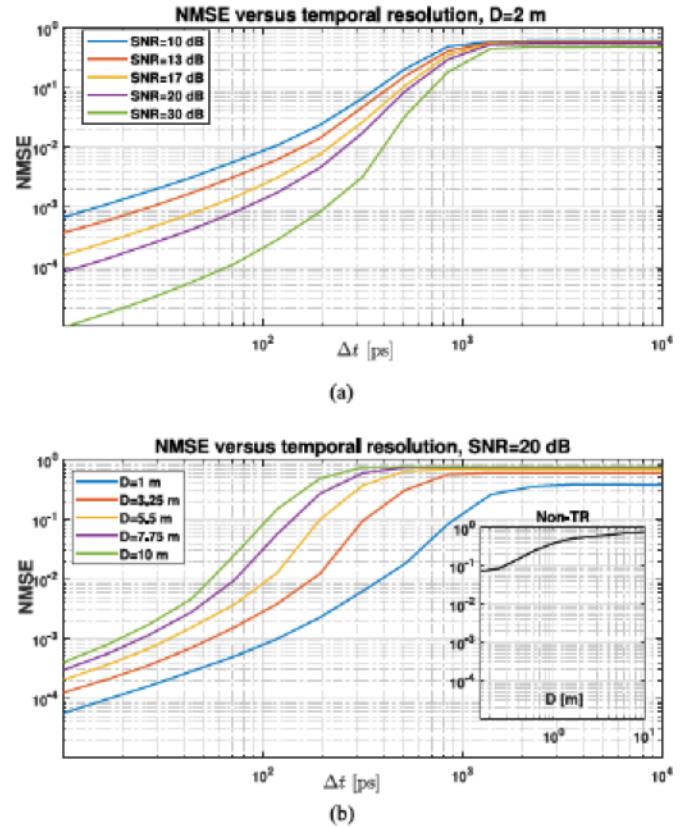


Fig. 4. Study of the reconstruction error as a function of the available temporal resolution Δt of the detector in TR sensing. (a) Normalized mean-squared error in reconstruction versus temporal resolution for different SNR values. (b) Normalized mean-squared error in reconstruction versus temporal resolution for different wall separations D .

as reconstruction in an ideal noise-free experiment could result in high-fidelity reconstruction even if \mathbf{A} is ill conditioned.⁴

When imaging more distant walls, the poor conditioning of \mathbf{A} further deteriorates as the distance decay factors become less varied and approach constants $\|\mathbf{x} - \ell\| \approx \|\mathbf{x} - \mathbf{c}\| \approx D$, as illustrated in Fig. 4(b) where reconstruction performance is parametrized against D for a fixed SNR in a setup with otherwise identical parameters as those of the first subfigure. In particular, notice in this plot the limit of non-time-resolved measurements, $\Delta t > 1.5$ ns, where the NMSE is always poor but is especially bad for larger D . This limit is separately summarized in the inset, which reveals that unless the room size is particularly small (i.e., just a few cm) high-fidelity reconstruction is impossible.

Summarizing, we see that unless very fine time-resolved measurements are available, NLOS scene reconstruction becomes ill-posed and reconstruction is not robust. To be more specific, notice that reconstruction from non-TR measurements in this NLOS setting fails, despite the fact that we are considering a simplified imaging problem with known geometry.

⁴Each curve in Fig. 4(a) corresponds to a different SNR. In practice, when comparing setups of different temporal resolutions, the equipment involved will be technologically different, so that a fair comparison does not necessarily entail assuming a fixed SNR common to all setups. Notice however the general trend of worsening reconstruction performance with diminishing temporal resolutions, which holds for all SNR levels.

In the next section we show how occluders can enable high-fidelity reconstruction in non-time-resolved and practical room-size settings.

V. IMAGING WITH OCCLUDERS

The inversion problem in the poor temporal-resolution regime is inherently difficult as rows of the linear forward operator \mathbf{A} are smooth functions over the spatial target coordinate \mathbf{x} , resulting in bad conditioning of the operator. The situation changes drastically when the line of sight between ℓ (and \mathbf{c}) and the hidden wall is partially obstructed by an occluder: for each (ℓ, \mathbf{c}) pair, certain segments of the hidden wall (that are different for different pairs) are occluded from ℓ or from \mathbf{c} . The occlusions are encoded in the linear forward operator \mathbf{A} via zero entries on the corresponding spatial target coordinates \mathbf{x} , such that its rows are choppy and varied. Consequently, the inverse problem (4) becomes significantly better conditioned. This section builds on this idea and studies situations in which high-fidelity reconstruction becomes possible without TR measurements.

A. Informative Measurements Through Occlusions

Non-TR measurements $y_{\ell, \mathbf{c}}$ correspond to integrating (2) over time. Assuming $\int p(t) dt = 1$, we get

$$y_{\ell, \mathbf{c}} = \int_S f(\mathbf{x}) \frac{V(\mathbf{x}, \ell) V(\mathbf{x}, \mathbf{c})}{\|\mathbf{x} - \ell\|^2 \|\mathbf{x} - \mathbf{c}\|^2} G(\mathbf{x}, \ell, \mathbf{c}) d\mathbf{x}. \quad (7)$$

Let L be the number of distinct occluders \mathcal{O}_t , $i = 1, \dots, L$ that are present in the scene. We associate a distinct (binary) visibility function $V_i(\mathbf{x}, \mathbf{z})$ to each of them. The overall visibility function $V(\mathbf{x}, \mathbf{z})$ becomes $V(\mathbf{x}, \mathbf{z}) = \prod_i V_i(\mathbf{x}, \mathbf{z})$, such that:

$$\mathbf{A} = \mathbf{A}_0 \circ (\mathbf{V}_1 \circ \dots \circ \mathbf{V}_L). \quad (8)$$

Here, \mathbf{A}_0 is the operator corresponding to a scene with no occluders, and \mathbf{V}_i is the (binary) *visibility matrix*, which has K rows (as many as the number of (ℓ, \mathbf{c}) pairs), N columns, and each of its entries takes values as follows:

$$(\mathbf{V}_i)_{(\ell, \mathbf{c}), \mathbf{x}} = V_i(\mathbf{x}, \ell) V_i(\mathbf{x}, \mathbf{c}). \quad (9)$$

Lastly, \circ denotes the Hadamard entry-wise product of matrices. On the one hand, the operator \mathbf{A}_0 is generally badly conditioned: successive entries of any of its rows exhibit small and smooth variations due only to the quadratic distance attenuation and the BRDF factors G in (7). On the other hand, the Hadamard multiplication with nontrivial binary visibility matrices results in an operator with much better conditioning.

This behavior is demonstrated through an example in Fig. 5, which compares reconstruction performance in the presence and absence of occluders. The setup, illustrated in Fig. 5(a), is as reported in previous simulations, with the addition of occluders as depicted. We collect $K = 30$ measurements with randomly drawn ℓ, \mathbf{c} parameters and noise variance such that SNR = 25 dB. The occluded measurement matrix \mathbf{A} and the non-occluded matrix \mathbf{A}_0 are depicted in Fig. 5(b) alongside their corresponding singular values. Observe that the singular values of \mathbf{A}_0 decay substantially faster than those of \mathbf{A} , which exhibits a much flatter spectrum. As expected, this

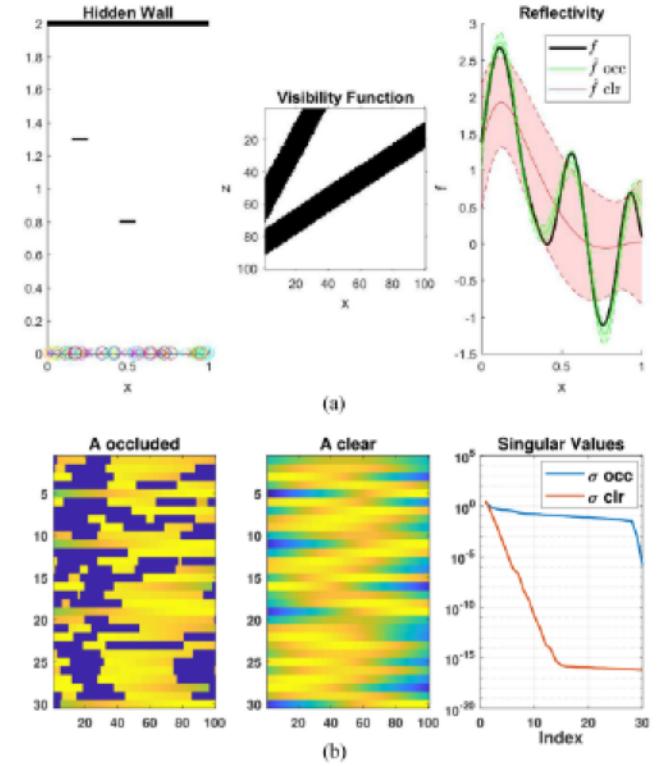


Fig. 5. Illustrating the beneficial role of occluders, by comparing imaging in their absence and presence. (a) (left) Room setup. On the illumination wall, positions that are marked with 'x' (resp. 'o') indicate virtual laser (resp. camera) points. (middle) Binary visibility matrix, with 0 (1) depicted in black (white). (right) Reflectivity reconstruction with (in green) and without occluders (in red). (b) (left) Measurement matrix when occluders are present in the room. The values of its entries are depicted in Matlab's jet colormap as in Fig. 3. (middle) Measurement matrix in the absence of occluders. (right) Singular values of the two matrices in decreasing order.

better conditioning translates to better image reconstruction, as illustrated in the rightmost panel of Fig. 5(a): in solid red is the poor reconstruction without the occluder (NMSE = 54%), and in solid green is the successful reconstruction with the occluder (NMSE = 2.4%). The dashed lines indicate the standard deviations of the error $\hat{f}_i - f_i$ for each spatial coordinate x_i , which correspond to the square-root of the diagonal entries of the posterior covariance matrix.

B. Measurement Schemes

So far we have considered a generic setting in which a focused laser source and a focused camera generate measurements corresponding to some given set of (ℓ, \mathbf{c}) pairs on the illumination wall. In principle, all possible such ℓ and \mathbf{c} combinations are allowed. In this section we discuss the following special instances of this general scheme: (i) selection of most informative subset of (ℓ, \mathbf{c}) pairs under a budget constraint on the number of allowed measurements; (ii) measurement collection with a wide field-of-view camera and (iii) specific measurement sets that are favorable from an analysis viewpoint.

Optimal measurement configuration:. We consider a situation where collection of at most K measurements is allowed, e.g., in order to limit the acquisition time of the imaging system.

Under such budget constraint, we suggest an efficient strategy to choose an optimal set \mathcal{P} of (ℓ, c) pairs and we study the imaging performance as a function of the number of allowed measurements.

Let \mathcal{D} be a (uniform) discretization of the illumination wall, and $(\ell, c) \in \mathcal{D} \times \mathcal{D}$. The idea is to choose a subset \mathcal{P} such that the corresponding measurement vector $\mathbf{y}_{\mathcal{P}} := \{\mathbf{y}_{\ell,c} \mid (\ell, c) \in \mathcal{P}\}$ is the most informative about the unknown reflectivity \mathbf{f} . Using $I(\cdot; \cdot)$ to denote the mutual information between two random vectors, this amounts to solving

$$\mathcal{P}^* = \underset{\mathcal{P}: |\mathcal{P}| \leq K}{\operatorname{argmax}} \Phi(\mathcal{P}), \quad \Phi(\mathcal{P}) \equiv I(\mathbf{y}_{\mathcal{P}}; \mathbf{f}). \quad (10)$$

The optimization problem in (10) is NP-hard in general. However, it turns out that under the framework of Section III the objective function $\Phi(\mathcal{P})$ is monotonic and submodular (see for example [22], [23] for similar derivations). The theory of submodular optimization then suggests that an efficient greedy solver obtains near-optimal solutions \mathcal{P}^{gr} satisfying: $\Phi(\mathcal{P}^{\text{gr}}) \geq (1 - \frac{1}{e})\Phi(\mathcal{P}^*)$ [24]. The greedy algorithm augments the set \mathcal{P} with an additional choice (ℓ, c) per iteration, for a total of K iterations. The solution has the property $\mathcal{P}_K^{\text{gr}} \subset \mathcal{P}_{K+1}^{\text{gr}}$, where we have used subscript notation for the budget constraint on the allowable size of \mathcal{P} . The algorithm picks the next element myopically given the solution set built so far, i.e., the algorithm picks the next element as the one which maximizes the marginal information gain. Submodular set functions are well studied and have many desirable properties that allow for efficient minimization and maximization with approximation guarantees, e.g., [24].

We illustrate the efficacy of this approach via numerical simulations. For the purpose of clearly illustrating the solution in a simple setting our setup is similar to that of Fig. 5(a), except we only position one of the two occluders (the one centered around 0.5 m). The noise variance is kept constant at $\sigma^2 = 0.1$, and we seek an optimal set \mathcal{P} of measurements under a budget constraint $|\mathcal{P}| \leq K$. Fig. 6(a) shows the output of the greedy algorithm for the most informative (ℓ, c) pairs for values of K up to 30. The selected parameters, marked with red crosses are accompanied by a number indicating the iteration cycle at which they were retrieved. Notice how the first two measurement configurations are selected one to the left and the other to the right of the occluder, thus casting effective shadows on different parts of the hidden wall. Fig. 6(b) validates the optimality features of the output \mathcal{P}^{gr} of the greedy algorithm by comparing it to an equal size subset of measurements chosen uniformly at random. For a fixed desired NMSE the number of measurements required when randomly picking can be as large as double the number required with optimal choice. On the other hand, observe that under both schemes the NMSE drops significantly for the first few added measurements and the marginal benefit degrades as more measurements are added.

Single-pixel camera with a wide field of view: An additional benefit from exploiting occlusion for scene reconstruction with non-TR measurements is the ability to use a single-pixel camera with a wide field-of-view in lieu of the focused detector that is typically required for TR imaging techniques. This camera change offers several advantages, such as reduced

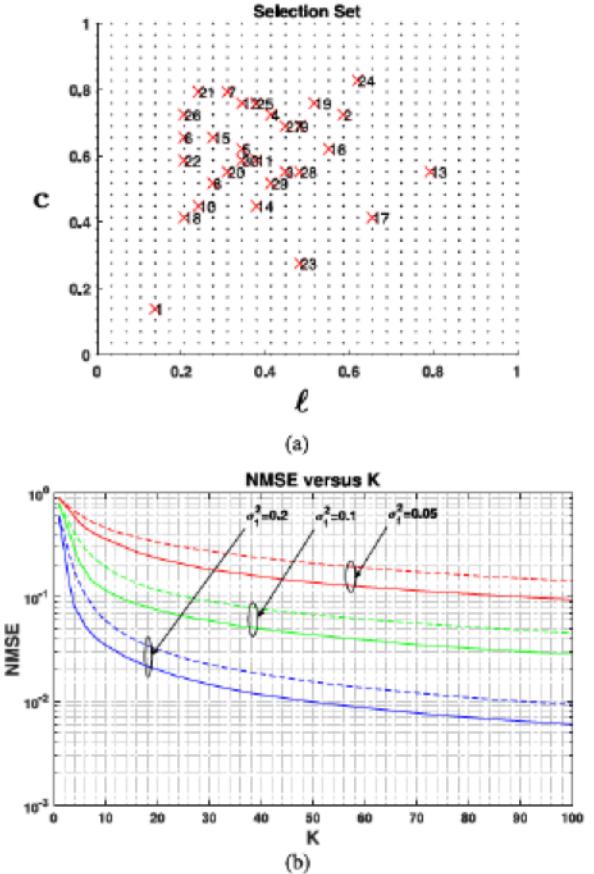


Fig. 6. Illustration of the efficient greedy selection algorithm for choosing informative measurements under a budget constraint. (a) Coordinates of virtual laser (ℓ , on the horizontal axis) and camera (c , on the vertical axis) positions. The set $\mathcal{D} \times \mathcal{D}$ of all possible locations is marked with black dots. The set \mathcal{P} selected by the greedy algorithm for a budget constraint $K = 30$ is marked with red crosses. The numbers indicate the order of selection. (b) Reconstruction performance versus total number of measurements for the random (dashed lines) and optimized by the greedy algorithm (solid lines) configurations for various values of the spatial correlation σ_f^2 parameter.

equipment cost (no lens required) and a dramatically increased SNR as more photons are collected per measurement. To the best of our knowledge, this is the first demonstration of NLOS imaging with a wide field-of-view detector. A camera that is configured for wide field-of-view operation detects light reflected from multiple positions c on the illumination wall. Thus, it captures more of the backscattered photons from the hidden scene and modifies the forward measurement model as explained next. Let \mathcal{C} represent the surface of the illumination wall that is in the camera's fixed field of view, while the laser source raster scans the illumination wall as before. This procedure yields measurements that are now parametrized only by ℓ , as follows:

$$\begin{aligned} y_{\ell} &= \int_{\mathcal{C}} \frac{y_{\ell,c}}{\|\mathbf{c} - \mathbf{\Gamma}\|^2} \cos(\mathbf{\Gamma} - \mathbf{c}, \mathbf{n}_c) d\mathbf{c} \\ &= \int_{\mathcal{S}} f(\mathbf{x}) \frac{V(\mathbf{x}, \ell)}{\|\mathbf{x} - \ell\|^2} \\ &\quad \cdot \left[\int_{\mathcal{C}} \frac{V(\mathbf{x}, \mathbf{c}) G(\mathbf{x}, \ell, \mathbf{c}) \cos(\mathbf{\Gamma} - \mathbf{c}, \mathbf{n}_c)}{\|\mathbf{x} - \mathbf{c}\|^2 \|\mathbf{c} - \mathbf{\Gamma}\|^2} d\mathbf{c} \right] d\mathbf{x}. \quad (11) \end{aligned}$$

In deriving (11), we used (7) and we further explicitly accounted for the quadratic power decay from the illumination wall to the position of the camera that is denoted by Γ . The measurements are again linear in the unknown reflectivity, hence, the same reconstruction techniques can be used. In the presence of occluders, the nontrivial visibility function $V(\mathbf{x}, \mathbf{z})$ results in a better-conditioned measurement operator and a successful image reconstruction. In particular, our experimental demonstration in Section VI is based on the forward model in (11). We mention in passing that the dual setting, where a wide field-of-view light projector is utilized instead of a focused laser illumination, with measurements collected at multiple locations \mathbf{c} on the illumination wall, might also be of interest.

Other measurement configurations: Lastly, we mention a specific configuration that reduces the dimensionality of the parameter space by imposing the restriction $\ell = \mathbf{c}$ on the measurements.⁵ This results in a strict subset of the entire measurement set $\mathcal{D} \times \mathcal{D}$ that is convenient for analysis purposes and for drawing insights about the features of the imaging system, and will be useful for our analysis in Section V-C.

C. Robustness

Here, we study in more detail the structural properties of the visibility function, which we use in turn to study the robustness of reconstruction with respect to a misspecified occluder location.

More on the visibility function: Henceforth, we focus on a simple, yet insightful, case of *flat horizontal* occluders, i.e. occluders aligned horizontally at some fixed distance from the illumination wall (see Fig. 2). This family of occluders is useful as any occluder that is small compared to the size of the room may be well approximated as being flat and horizontal. We show that the visibility function V associated with a flat horizontal occluder has a simple structure. To be concrete, suppose that the occluder \mathcal{O} lies on a horizontal plane at distance $H = \alpha D$ from the visible wall for some (known) $0 < \alpha < 1$. Further define the occupancy function $s(\mathbf{x})$ such that for all points \mathbf{x} on that plane $s(\mathbf{x}) = 0$ if \mathcal{O} occupies \mathbf{x} and $s(\mathbf{x}) = 1$ otherwise.⁶ A point \mathbf{x} on the hidden wall is *not* visible from a point \mathbf{z} on the illumination wall if and only if the line that connects them intersects with the occluder, or equivalently, if at the point of intersection it holds that $s(\alpha\mathbf{x} + (1 - \alpha)\mathbf{z}) = 0$. This translates to:

$$V(\mathbf{x}, \mathbf{z}) = s(\alpha\mathbf{x} + (1 - \alpha)\mathbf{z}). \quad (12)$$

In particular, when $\ell = \mathbf{c}$, it follows from (9) and (12) that

$$(V)_{(\ell, \mathbf{c}), \mathbf{x}} = s(\alpha\mathbf{x} + (1 - \alpha)\ell),$$

and the visibility function V_ℓ has a band-like structure. Ignoring edge-effects, its discretization corresponds to a convolution

⁵Strictly speaking, when $\ell = \mathbf{c}$, the camera focused at \mathbf{c} sees a first-bounce response in addition to the informative third-bounce. We assume here that the dimensions of the entire scene are such that it is possible to use time-gating to reject that first-bounce. Note that this is possible with mild temporal resolution requirements.

⁶Here, occluder \mathcal{O} is allowed to be composed of several patches as long as they all lie on the same plane. Equivalently, the set of values for which $s(\mathbf{x}) = 0$ need not be connected.

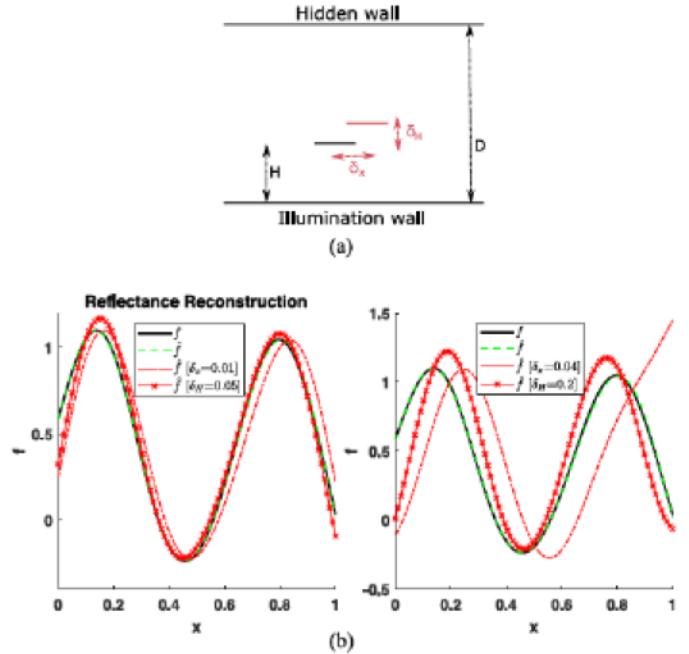


Fig. 7. Illustrating the effect of modeling mismatches on reconstruction. (a) A shifted occluder setup. The occluder appears in its actual position in black. We perform reconstruction under imperfect knowledge of its position; taken to be as appears in red. (b) Reconstruction with a mispositioned occluder. (left) Small and (right) large vertical and horizontal shifts in a far field setup.

matrix, which is favorable since the convolution structure makes possible deriving analytic conclusions regarding the effect of the occluder's parameters on the image reconstruction as shown next.

The effect of modeling mismatches: We study scene reconstruction under a mismatched model for the position of the occluders to evaluate the robustness of our imaging method with respect to such modeling errors. Fig. 7(a) illustrates our setup, where the true position of the occluder appears in black, and our mismatched model assumes the occluder is positioned, as appears in red, with δ_x and δ_H vertical and horizontal shifts, respectively. We study the resulting reconstruction under the following simplifications: (i) measurements are noiseless, (ii) measurements are taken with parameters satisfying $\ell = \mathbf{c}$, (iii) continuous measurements are collected, i.e., y_ℓ is available for all points ℓ on the visible wall, and (iv) we assume that the hidden wall is far from the illumination wall such that $\|\mathbf{x} - \ell\|^2 \|\mathbf{x} - \mathbf{c}\|^2$ and $G(\mathbf{x}, \ell, \mathbf{c})$ are approximately constant.

Under these assumptions, the measurements y_ℓ are expressed (up to a constant) as

$$y_\ell = \int f(\mathbf{x}) s(\alpha\mathbf{x} + (1 - \alpha)\ell) d\mathbf{x}, \quad (13)$$

where we have used (12), and $f(\mathbf{x})$ is the true reflectivity of the hidden wall.

In the presence of errors δ_x, δ_H , the misspecified visibility function can be expressed as $\tilde{V}(\mathbf{x}, \mathbf{z}) = s(\alpha'(\mathbf{x} - \delta_x) + (1 - \alpha')(\ell - \delta_x))$, where $\alpha' := \frac{H + \delta_H}{D} = \alpha + \frac{\delta_H}{D}$. This results in a

misspecified model:

$$y_\ell = \int \hat{f}(x) s(\alpha'(x - \delta_x) + (1 - \alpha')(\ell - \delta_x)) dx. \quad (14)$$

In order to study how $\hat{f}(x)$ relates to $f(x)$ it is convenient to work in the Fourier domain.⁷ Manipulating (13) and (14) accordingly, we show in the Appendix that

$$\hat{F}(\omega) = \frac{1 - \alpha'}{1 - \alpha} \frac{S(-\frac{1 - \alpha'}{1 - \alpha} \frac{\omega}{\alpha'})}{S(-\frac{\omega}{\alpha'})} e^{-j\omega \frac{\ell}{\alpha'}} F\left(\frac{\alpha}{\alpha'} \frac{1 - \alpha'}{1 - \alpha} \omega\right), \quad (15)$$

where $H(\omega)$ denotes the Fourier transform of a function $h(x)$. Of course, this holds for spatial frequencies at which $S(\omega)$ is non-vanishing.

The following conclusions regarding reconstruction distortion under mismatched occluder position are drawn from (15). (a) In the absence of errors ($\delta_x, \delta_H = 0$), the reflectivity function is perfectly reconstructed for those frequencies for which the occluder's occupancy function is non-zero. (b) Horizontal occluder translation errors ($\delta_x \neq 0, \delta_H = 0$) result in simple shifts of the true reflectivity. (c) Vertical occluder translation errors ($\delta_x = 0, \delta_H \neq 0$) result in two kinds of distortion. The first is a scaling effect, while the other is a distortion that depends on the shape of the occluder through the term $S(-\frac{1 - \alpha'}{1 - \alpha} \frac{\omega}{\alpha'})/S(-\frac{\omega}{\alpha'})$. For this latter term, observe that its effect diminishes for a spectrum $S(\omega)$ that is mostly flat over a large range of spatial frequencies. This property is approximately (due to the finite support⁷ of $s(x)$) satisfied by a very narrow occluder.

Recall that the above conclusions hold analytically in the limit of a distant hidden scene and a continuum of noiseless measurements. However, the conclusions are also suggestive and insightful for practical scenarios, as illustrated by the numerical study shown in Fig. 7, where we illustrate high SNR (35 dB) reconstruction with a mispositioned occluder. The room setup is $D = 5$ m, with a single occluder of width 0.25 m positioned at $[0.5, 2]$ m. Measurements are collected with random ℓ and random $c \neq \ell$. Black solid lines show the true reflectivity $f(x)$. The dashed green line depicts reconstruction under perfect occluder knowledge. The red curves show reconstructions with horizontally and vertically mispositioned occluders. The mispositioning is larger in the right subplot. It is evident from the images that horizontal mispositioning mostly results in a shifted reconstruction, whereas vertical mispositioning results in axis-scaling of the reconstructed scene. Our analysis-based conclusions seem to be valid for the middle section of the reflectivity function, whereas edge effects appearing close to the boundaries $x = 0, 1$ are not captured by the analysis.

The robustness of our imaging method with respect to occluder positioning errors is further supported by the experimental demonstration in Section VI, where such occluder modeling inaccuracies are unavoidable, yet the reconstruction results we demonstrate are satisfactory.

⁷The variable of integration x in (13) and (14) ranges over the finite surface of the hidden wall. Correspondingly, $f(x)$ and $s(x)$ are only defined over this region. Formally, when it comes to taking Fourier transforms, we extend the functions on the rest of the space by zero-padding.

D. Reconstruction of Reflectivity With Unknown Distance

Thus far, we have demonstrated the use of occluders to reconstruct the unknown reflectivity of a hidden wall when its geometry is known. Here, we develop a simple algorithm for reflectivity reconstruction with the aid of occluders when the distance D between the visible and hidden walls is unknown.

In line with the Bayesian approach in Section III-B, we associate some distribution with the unknown depth D , and attempt joint estimation of both D and f by solving the maximum a posteriori probability (MAP) problem:

$$(\hat{D}, \hat{f}) = \arg \max_{D', f'} p(D', f' | y), \quad (16)$$

where $y = A_D f + \epsilon$ as in (4). Observe that the distance D enters the measurement equations via the forward operator, which we have parametrized as A_D . For a fixed D' the maximization in (16) with respect to f has already been studied in terms of (efficient) implementation and performance. Namely, under a Gaussian prior assumption on f , each maximizer $\hat{f}_{D'}$ coincides with the MMSE estimator of Section III-B. Based on this observation, a simple and effective strategy for solving the joint optimization in (16) is as follows. Start with a range of candidate distance values, D_1, D_2, \dots, D_N . For each candidate, form the measurement matrix A_{D_i} and solve for the corresponding reflectivity vector \hat{f}_{D_i} . Then, for $i = 1, 2, \dots, N$, compute i_* that maximizes the likelihood (we assume here a uniform prior among the D_i 's):

$$i_* = \arg \max_i p(y | \hat{f}_{D_i}, D_i).$$

Finally, return $(\hat{D}, \hat{f}) = (D_{i_*}, \hat{f}_{i_*})$. In particular, under the Gaussian prior assumption, it can be shown that

$$-\log p(y | \hat{f}_{D_i}, D_i) = y^\top (A_{D_i} \Sigma_f A_{D_i}^\top + \sigma^2 I)^{-1} y. \quad (17)$$

Note however that the algorithm can be readily adapted to different priors on f .

Fig. 8 includes an illustration of the algorithm and a numerical demonstration of its performance for different values of parameters such as SNR and number of measurements. The room setup is the same as in Fig. 5. In particular, the true distance of the hidden wall is $D = 2$ m and the reflectivity is drawn from a Gaussian prior with $\sigma_f^2 = 0.05$. A total number of K randomly selected (ℓ, c) -measurements are collected. Observe in Fig. 8(a) that the negative log-likelihood in (17) shows a valley in the neighborhood of the true distance D . Higher values of SNR result in sharper valleys and the minimum occurs at the true distance (here $D = 2$ m) provided that enough measurements are available (see Fig. 8(b)). The plots shown are averages over 200 realizations drawn from the Gaussian prior with each instance measured by 30 randomly selected (ℓ, c) pairs.

E. Collecting TR-Measurements in Occluded Settings

Thus far we have focused on imaging systems that use either TR measurements or non-TR measurements and occlusions. It is natural to attempt combining the best of both worlds. A full study of this topic is beyond the scope of the paper, but we present numerical simulations to illustrate its promise. Consider the familiar setting of Fig. 5(a) and a detector with a nontrivial

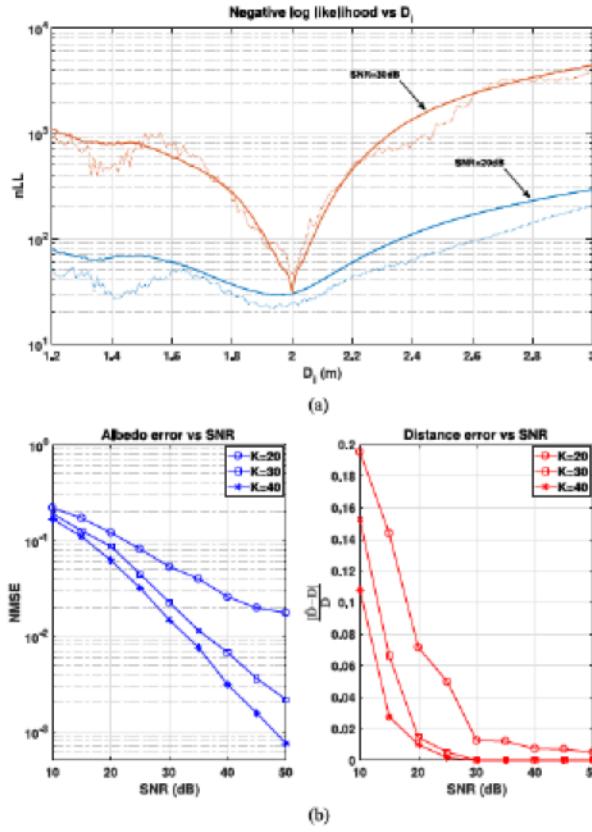


Fig. 8. Illustration of the proposed algorithm for reflectivity estimation when the distance D is unknown. (a) Plots of the negative log-likelihood (nLL) function versus each candidate distance value D_t (see (17)) for two different values of the SNR and for $K = 30$ measurements. The solid lines represent averages over 200 realizations of the reflectivity and of the measurement positions. The dashed lines show the nLL for a specific such realization. (b) Plots of (normalized) reconstruction error for the reflectivity and the distance as a function of the SNR for different numbers of measurements.

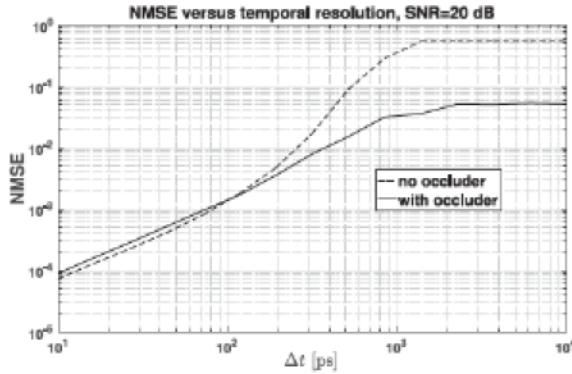


Fig. 9. Comparing reconstruction performance versus temporal resolution in the presence and absence of occluders.

temporal resolution Δt . We sweep Δt over a range of values, and plot the resulting NMSE in Fig. 9 (solid curve). For comparison, we also plot in dashed line the NMSE performance in the absence of an occluder (this corresponds exactly to the plot in Fig. 4(a)). For a large range of temporal resolutions (here, $\Delta t \gtrsim 150$ ps) the presence of occlusions leads to a substantial improvement in reconstruction performance, allowing the same level of performance to be maintained at inferior temporal resolution levels. When very fine temporal resolution is

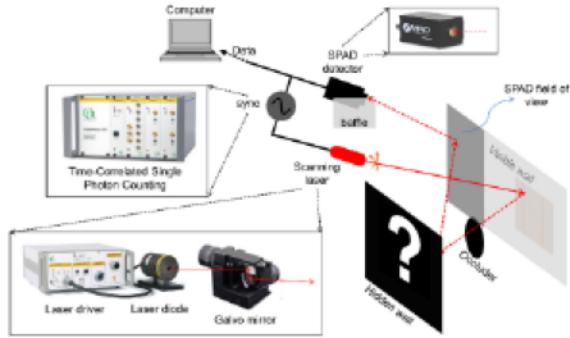


Fig. 10. Experimental setup. Visible wall to hidden wall: ~ 106 cm; visible wall to SPAD: ~ 156 cm; visible wall to occluder: ~ 37 cm; Diameter of the circular occluder is 3.4 cm. Scene reconstruction was done with non-TR measurements by discarding the time stamps obtained from the SPAD.

available, the reconstruction performance is almost identical whether occluders are present or not. Note here that TR measurements can be further utilized to improve other aspects of the system. For instance, one might imagine using (coarse) TR measurements to find the position of the occluder more effectively than could otherwise be possible. We comment more on this in Section VII.

VI. EXPERIMENTAL ILLUSTRATION

We experimentally demonstrate an instance of opportunistic exploitation of occluders to perform NLOS active imaging with non-TR measurements. More extensive experiments can be found in [18], where our methods are extended to the low-photon-count regime by means of a physics-based noise model for a SPAD detector.

Experimental Setup: The schematic setup of our experiment is shown in Fig. 10. A 640 nm laser source (Picoquant LDH-640B) operating at average output power ~ 1 mW transmits optical pulses with ~ 350 ps pulse width (full-width half-maximum), at a 5 MHz repetition rate towards a nearly Lambertian-surface visible wall (1st bounce). The scattered light travels to the hidden wall, which scatters the light back (2nd bounce). Finally, the backscattered light from the visible wall (3rd bounce) is collected by a SPAD detector (Micro Photon Devices PDM) with 35% quantum efficiency and ~ 50 ps timing resolution. This process is repeated multiple times as the laser illumination is raster scanned along a uniform grid of illumination points ℓ on the illumination wall, and for each ℓ we record the total number of photons detected by the SPAD over a fixed dwell duration.

The SPAD is capable of providing time-resolved measurements. However, for the purpose of this experiment we operate the SPAD as a regular camera, discarding the temporal information by integrating its response over time,⁸ yielding a single

⁸To be precise, we only use the SPAD's time-resolution capability to gate-out the first-bounce response from the illumination wall. Beyond that, no TR information is employed in our scene reconstructions as they employ just the sum total of post-first-bounce photons that were detected. Notice that the illumination wall is in the direct line of sight of the imaging equipment, thus its location can be well-estimated based on standard imaging techniques. With this information, the time window that corresponds to the first-bounce response is a-priori known. Hence, the same operation achieved here with a SPAD camera can in principle be performed using a time-gated camera collecting non-TR intensity measurements.

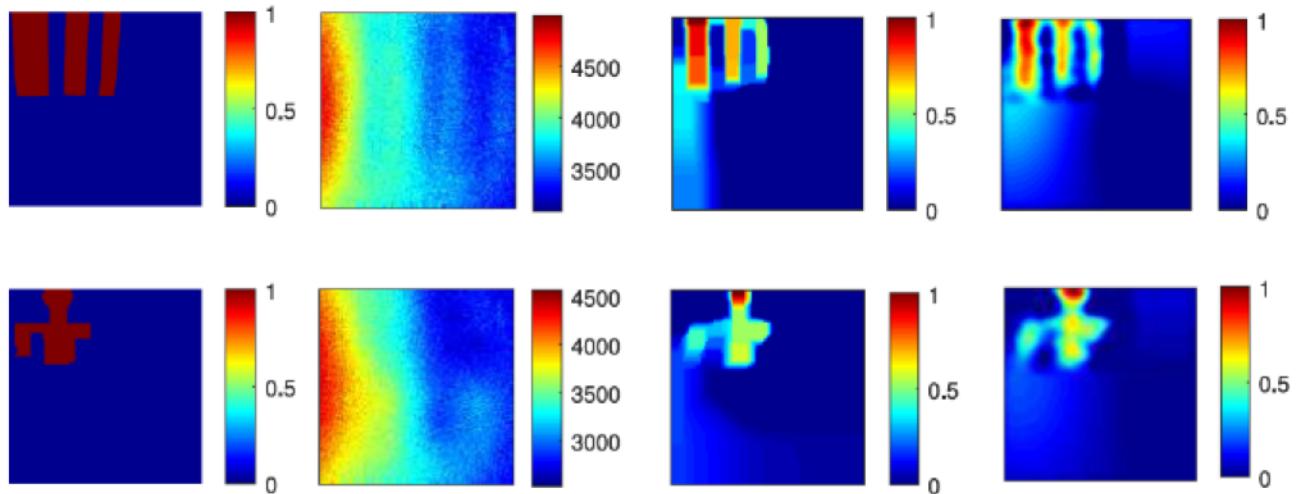


Fig. 11. (first col.) Ground truth of the tested scene patterns on the hidden wall. The patterns are placed in the upper-left corner of the hidden wall. (second col.) Raw measurement counts for 100×100 raster-scanning laser points. At each laser point, we turn on the SPAD for a fixed dwell time such that ~ 3500 photon counts are recorded on average. (third col.) Reconstruction results from (18). (fourth col.) Reconstruction results for the linear method in (6) that is based on the Gaussian prior model.

scalar intensity measurements per each ℓ configuration. In front of the SPAD, an interference filter (Andover, 2 nm bandwidth) centered at 640 nm is used to remove most of the background light. In the experiment, the SPAD is lensless and configured for wide-field-of-view observation of the left side of the visible wall to minimize first bounce light detection. On average, for each raster scanning laser point, the SPAD detects approximately one third-bounce photon per 3000 illumination pulses. The occluder is a black-surface circular patch without any back reflections. During the experiment, we turned off all ambient room light to minimize background noise.

Computational processing: Based on the forward model in (11) we obtain an estimate \hat{f} of the true reflectivity from the measurements by applying one of two computational methods.

The first computational method we apply is solving the following non-smooth convex optimization program:

$$\hat{f} = \arg \min_{f \geq 0} \frac{1}{2} \|y - Af\|_2^2 + \lambda \|f\|_{TV}, \quad (18)$$

where $\|\cdot\|_{TV}$ is the Total-Variation (TV)-norm and $\lambda > 0$ is a regularization parameter. To solve (18) we use an efficient dedicated iterative first-order solver [25], which is based on the popular FISTA algorithm [26]. TV-norm penalization is a standard technique that has been successfully applied in other imaging tasks (e.g., image restoration [10], [26], [27]). Its use is motivated by the observation that the derivatives of natural images have heavy-tailed prior distributions [10], [27].

The second computational method we use is obtaining an estimate of the scene \hat{f} by positing a Gaussian prior (GP) and performing Bayesian inference from the measurements to the scene using the linear methodology of (6).

Results: Our experimental results are summarized in Fig. 11. Two different reflectivity patterns on the hidden wall were tested (first column). The laser light was raster scanned on a 100×100 grid and, at each point, the SPAD detector was turned on for a fixed dwell time such that a total number of ~ 9 million laser pulses were emitted and ~ 3500 back-reflected third-bounce

photons were recorded on average. The laser's raster-scanning area is such that the hidden pattern, which is placed in the top-left quadrant of the hidden wall, is completely scanned by the occluder's shadow pattern (i.e., the projection of the circular object on the hidden wall). The raw measurement counts for each of the hidden patterns are shown in the second column of the figure: each one of the 100×100 entries corresponds to a measurement y_ℓ for the corresponding virtual laser position ℓ .

A pre-imaging measurement of background light, in the absence of a target pattern on the hidden wall, was made over a long observation time and used to subtract the *average* background-count level from the raw counts collected (over a much shorter measurement interval) when there was a target present. The background-corrected raw counts were then used to reconstruct hidden-wall reflectivity using the TV and GP methods.⁹

Reconstruction results using the optimization method in (18) are shown in the third column. The regularization parameter λ was tuned independently for each pattern to yield a reconstruction that is empirically closest to the ground truth. Tuning in this manner is convenient for such demonstrations, but in the absence of ground truth, one typically resorts to a cross-validation procedure.

Finally we performed reconstruction according to the linear scheme in (6) that assumes a Gaussian prior on $f(x)$ (see Section III-B) with $\sigma_f^2 = 0.02$ and σ^2 tuned to achieve good results. These results are shown in the fourth column of Fig. 11, where we threshold the reconstruction to only keep the positive values of \hat{f} , and scale such that the maximum is 1.

Comparing the two processing methods, we note that TV regularization is more accurate and emphasizes the edges in the scene, as expected. The linear reconstruction is blurry but satisfactory and yields a reconstruction that is easily interpretable by

⁹See [18] for a more effective background-suppression technique that relies on the binomial-likelihood which models SPAD operation in the low-photon-count regime. Ref. [18] also points out that operation at 1550 nm wavelength, instead of the 640 nm wavelength employed there and here, would greatly aid in reducing background-light detection and its accompanying shot noise.

human eyes. One should also note that the linear reconstruction is more computationally efficient. Both methods require tuning of the involved parameters: σ_f^2 and σ^2 for GP, and λ for TV.¹⁰

These results validate the forward model and the performance of the reconstruction algorithm.

VII. DISCUSSION AND FUTURE WORK

In this paper we introduced and explored the benefits of exploiting occlusions in NLOS imaging. We focused on the problem of reconstructing the reflectivity of a hidden surface of known geometry from diffuse reflections, further assuming that the occluders in this setup were absorbing and of known geometry. This served as a useful testing ground for demonstrating basic principles in occluder-assisted NLOS imaging. At the same time, our promising results suggest that it is of interest to extend the study to more complicated system models. It further suggests exploring the premises of opportunistic NLOS imaging under even broader settings. In what follows, we elaborate on relevant directions of future research.

Beyond the problem of reflectivity estimation, it remains to explore extensions towards full 3D reconstruction of more complicated scenes. While much of our focus has been on identifying scenarios where the use of occluders can alleviate the need to collect TR measurements, we speculate that combined use of both TR measurements and occluders can assist in approaching more complicated problems such as the aforementioned.

Another interesting extension is as follows. Rather than using known occluders to reconstruct the reflectivity function, one can imagine scenarios where the reflectivity function of a back wall is known, thus it can be exploited to identify the position of unknown objects in the hidden room. In terms of the forward model in (2) this essentially asks for an estimate of the visibility function given the measurements and the reflectivity $f(x)$, since the visibility function is in turn informative about the shape of the occluders.

Continuing along the same lines it is natural to consider the fully blind problem, in which both the reflectivity function and the occluder shape are unknown. A natural approach to solving this problem is an iterative alternating-optimization method, which iterates between the two subproblems that were previously discussed: solve for $f(x)$ given the occupancy function, and vice versa. For each subproblem, we can use convex-optimization with appropriate regularization to promote the statistical or structural properties of the desired quantities (e.g., Gaussian prior on $f(x)$ and a low total-variation assumption on the occupancy function). Analyzing the convergence proper-

¹⁰In [18], we replace the additive white-Gaussian noise model employed herein with a physics-based model for SPAD operation in the low-photon-count regime. This model combines the Poisson statistics of photodetection shot noise with the effect of SPAD detection's dead time to arrive at a binomial likelihood function. Experiments reported in [18], whose reconstructions used the binomial likelihood function together with TV regularization, showed a 16× reduction in the number of detected photons per pixel needed to achieve satisfactory NMSE's as compared to those obtained from the Gaussian likelihood function and TV regularization. Use of the Gaussian likelihood function in the experiments reported here, however, suffices for the present purpose, which is to provide validation of the paper's forward model and reconstruction algorithm in the high-photon-count regime.

ties of such procedures and further understanding the extent to which different priors are sufficient to identify the true underlying quantities are compelling research questions.

Similar to the use of occluders as a form of opportunistic imaging, it is possible that exploiting other structural features of the environment results in enhancement of NLOS imaging. As already discussed, one such example involves exploiting the possibly known reflectivity pattern on back walls. Another example is utilizing coincidental bumps or edges on the illumination wall itself, and the occlusions that those introduce. Finally, it is natural to attempt extensions of the discussed methods to dynamic environments. For instance, moving occluders will generate measurements with additional diversity that can be exploited towards more accurate and robust reconstructions.

APPENDIX

Here, we provide a detailed derivation of (15).

First, we make the substitutions $x' \equiv \alpha x$, $f'(x) \equiv f(\frac{x}{\alpha})$, $\ell' \equiv -(1 - \alpha)\ell$ in (13) to arrive at the following,

$$y_{-\frac{\ell'}{1-\alpha}} = \frac{1}{\alpha} \int f'(x')s(x' - \ell')dx'.$$

Taking the Fourier transform¹¹ of the expressions in both sides above, with $0 < \alpha < 1$ we have:

$$(1 - \alpha)Y(-(1 - \alpha)\omega) = \frac{1}{\alpha}F'(\omega)S(-\omega) = F(\alpha\omega)S(-\omega),$$

which can be written as:

$$Y(\omega) = \frac{1}{1 - \alpha}F\left(-\frac{\alpha}{1 - \alpha}\omega\right)S\left(\frac{1}{1 - \alpha}\omega\right). \quad (19)$$

Next, in (14) we make the substitutions $x' \equiv \alpha'x$, $\hat{f}'(x) \equiv \hat{f}(\frac{x}{\alpha'})$, $\ell' = \delta_x - (1 - \alpha')\ell$ to reach the following,

$$y_{-\frac{1}{1-\alpha'}\ell+\frac{\ell_x}{1-\alpha'}} = \frac{1}{\alpha'} \int \hat{f}'(x')s(x' - \ell')dx'.$$

By taking the Fourier transform we find that

$$(1 - \alpha')e^{-j\omega\delta_x}Y(-(1 - \alpha')\omega) = \hat{F}(\alpha'\omega)S(-\omega),$$

which can be written as:

$$Y(\omega) = \frac{e^{-j\omega\frac{\delta_x}{1-\alpha'}}}{1 - \alpha'}\hat{F}\left(-\frac{\alpha'}{1 - \alpha'}\omega\right)S\left(\frac{1}{1 - \alpha'}\omega\right). \quad (20)$$

Equating (19) and (20) and solving for $\hat{F}(\omega)$ we arrive at (15), as desired.

REFERENCES

- [1] A. Kirmani, T. Hutchison, J. Davis, and R. Raskar, "Looking around the corner using ultrafast transient imaging," *Int. J. Comput. Vis.*, vol. 95, no. 1, pp. 13–28, 2011.
- [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 Commun.*, vol. 3, p. 745, 2012.
- [3] M. Buttafava, J. Zeman, A. Tosi, K. Eliceiri, and A. Velten, "Non-line-of-sight imaging using a time-gated single photon avalanche diode," *Opt. Express*, vol. 23, no. 16, pp. 20 997–21 011, 2015.

¹¹Recall $\mathcal{F}[f(t)] = F(\omega) \rightarrow \mathcal{F}[f(at + b)] = \frac{1}{|a|}e^{j\omega\frac{b}{a}}F(\frac{\omega}{a})$.

- [4] G. Gariepy, F. Tonolini, R. Henderson, J. Leach, and D. Faccio, "Detection and tracking of moving objects hidden from view," *Nature Photon.*, vol. 10, no. 1, pp. 23–26, 2016.
- [5] F. Heide, L. Xiao, W. Heidrich, and M. B. Hullin, "Diffuse mirrors: 3-D reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 3222–3229.
- [6] E. E. Fenimore and T. Cannon, "Coded aperture imaging with uniformly redundant arrays," *Appl. Opt.*, vol. 17, no. 3, pp. 337–347, 1978.
- [7] A. L. Cohen, "Anti-pinhole imaging," *J. Modern Opt.*, vol. 29, no. 1, pp. 63–67, 1982.
- [8] D. J. Brady, N. P. Pitsianis, and X. Sun, "Reference structure tomography," *JOSA A*, vol. 21, no. 7, pp. 1140–1147, 2004.
- [9] R. Raskar, A. Agrawal, and J. Tumblin, "Coded exposure photography: Motion deblurring using fluttered shutter," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 795–804, 2006.
- [10] A. Levin, R. Fergus, F. Durand, and W. T. Freeman, "Image and depth from a conventional camera with a coded aperture," *ACM Trans. Graph.*, vol. 26, no. 3, p. 70, 2007.
- [11] A. Veeraraghavan, R. Raskar, A. Agrawal, A. Mohan, and J. Tumblin, "Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing," *ACM Trans. Graph.*, vol. 26, no. 3, p. 69, 2007.
- [12] M. F. Duarte, M. A. Davenport, D. Takbar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk, "Single-pixel imaging via compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 83–91, Mar. 2008.
- [13] M. S. Asif, A. Ayremliou, A. Sankaranarayanan, A. Veeraraghavan, and R. Baraniuk, "FlatCam: Thin, lensless cameras using coded aperture and computation," *IEEE Trans. Comput. Imag.*, vol. 3, no. 3, pp. 384–397, 2017.
- [14] G. Satat, M. Tancik, and R. Raskar, "Lensless imaging with compressive ultrafast sensing," *IEEE Trans. Comput. Imag.*, vol. 3, no. 3, pp. 398–407, 2017.
- [15] A. Torralba and W. T. Freeman, "Accidental pinhole and pinhole cameras: Revealing the scene outside the picture," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2012, pp. 374–381.
- [16] 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 *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 2270–2278.
- [17] J. Klein, C. Peters, J. Martin, M. Laurenzis, and M. B. Hullin, "Tracking objects outside the line of sight using 2-D intensity images," *Sci. Rep.*, vol. 6, 2016, Art. no. 32491.
- [18] 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," *Opt. Express*, vol. 26, pp. 9945–9962, 2018.
- [19] J. Besag, J. York, and A. Mollié, "Bayesian image restoration, with two applications in spatial statistics," *Ann. Inst. Stat. Math.*, vol. 43, no. 1, pp. 1–20, 1991.
- [20] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-6, no. 6, pp. 721–741, Nov. 1984.
- [21] M. K. Mihcak, I. Kozintsev, K. Ramchandran, and P. Moulin, "Low-complexity image denoising based on statistical modeling of wavelet coefficients," *IEEE Signal Process. Lett.*, vol. 6, no. 12, pp. 300–303, Dec. 1999.
- [22] G. Shulkind, S. Jegelka, and G. W. Wornell, "Sensor array design through submodular optimization," arXiv:1705.06616, 2017.
- [23] G. Shulkind, L. Horesh, and H. Avron, "Experimental design for non-parametric correction of misspecified dynamical models," arXiv:1705.00956, 2017.
- [24] S. Fujishige, *Submodular Functions and Optimization*. vol. 58. Amsterdam, The Netherlands: Elsevier, 2005.
- [25] Z. T. Harmany, R. F. Marcia, and R. M. Willett, "This is spiral-tap: Sparse poisson intensity reconstruction algorithms theory and practice," *IEEE Trans. Image Process.*, vol. 21, no. 3, pp. 1084–1096, Mar. 2012.
- [26] A. Beck and M. Teboulle, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems," *IEEE Trans. Image Process.*, vol. 18, no. 11, pp. 2419–2434, Nov. 2009.
- [27] D. Krishnan and R. Fergus, "Fast image deconvolution using hyper-Laplacian priors," in *Proc. Adv. Neural Inf. Process. Syst.*, 2009, pp. 1033–1041.