



Deep-inverse correlography: towards real-time high-resolution non-line-of-sight imaging

CHRISTOPHER A. METZLER,^{1,4,*} FELIX HEIDE,^{2,†} PRASANA RANGARAJAN,^{3,†} MURALIDHAR MADABHUSHI BALAJI,³ APARNA VISWANATH,³ ASHOK VEERARAGHAVAN,⁴ AND RICHARD G. BARANIUK⁴

¹Department of Electrical Engineering, Stanford University, 450 Serra Mall, Stanford, California 94305, USA

²Department of Computer Science, Princeton University, Princeton, New Jersey 08544, USA

³Department of Electrical Engineering, Southern Methodist University, 6425 Boaz Lane, Dallas, Texas 75205, USA

⁴Department of Electrical and Computer Engineering, Rice University, 6100 Main Street, Houston, Texas 77005, USA

*Corresponding author: cmetzler@stanford.edu

Received 29 July 2019; revised 2 December 2019; accepted 9 December 2019 (Doc. ID 374026); published 16 January 2020

Low signal-to-noise ratio (SNR) measurements, primarily due to the quartic attenuation of intensity with distance, are arguably the *fundamental barrier* to real-time, high-resolution, non-line-of-sight (NLoS) imaging at long standoffs. To better model, characterize, and exploit these low SNR measurements, we use spectral estimation theory to derive a noise model for NLoS correlography. We use this model to develop a speckle correlation-based technique for recovering occluded objects from indirect reflections. Then, using only synthetic data sampled from the proposed noise model, and without knowledge of the experimental scenes nor their geometry, we train a deep convolutional neural network to solve the noisy phase retrieval problem associated with correlography. We validate that the resulting *deep-inverse correlography* approach is exceptionally robust to noise, far exceeding the capabilities of existing NLoS systems both in terms of spatial resolution achieved and in terms of total capture time. We use the proposed technique to demonstrate NLoS imaging with 300 μm resolution at a 1 m standoff, using just two 1/8th s exposure-length images from a standard complementary metal oxide semiconductor detector. © 2020 Optical Society of America under the terms of the OSA Open Access Publishing Agreement

<https://doi.org/10.1364/OPTICA.374026>

<https://doi.org/10.1364/OPTICA.374026>

1. INTRODUCTION

Non-line-of-sight (NLoS) imaging recovers hidden objects from light scattered off these objects onto other surfaces in the scene; in essence, it lets us use a rough wall as a mirror. Although the majority of NLoS imaging methods have exploited time-of-travel information [1–12], recent work has demonstrated that spatial correlations in scattered light (speckle) contain sufficient information to image such hidden objects [13–15]. These methods recover the latent image of the hidden object by solving variations of a phase retrieval (PR) problem. While speckle-based methods achieve unprecedented results in high-flux illumination, such approaches struggle to recover a latent image in photon-starved environments, due in large part to the poorly understood noise characteristics of the correlation process. This shortcoming limits existing methods to long acquisition times and small standoff distances d .

In this paper, we examine the NLoS imaging problem illustrated in Fig. 1. The objective is to recover a spatially resolved image of a target hidden behind the corner. To this end, we indirectly illuminate the target using continuous-wave (CW) laser light scattered off a section of the visible wall (dubbed the virtual source), and record the object returns incident on another section of the visible

wall (dubbed the virtual detector). This configuration causes the hidden object's albedo, r , to be encoded in the distribution of the speckle pattern incident on the virtual detector. By analyzing spatial correlations in the speckle image, we can estimate the albedo's autocorrelation, $r \star r$. With this estimate in hand, we use a PR algorithm to recover r using the following relationship:

$$\mathcal{F}(r \star r) = |\mathcal{F}(r)|^2, \quad (1)$$

where \mathcal{F} denotes the Fourier transformation and the square is taken elementwise.

Although details differ in how $r \star r$ is encoded and estimated, this basic Fourier relationship is what underlies nearly all correlation imaging techniques [13,15–19], including NLoS correlography [14]. NLoS correlography is described in detail in Section 3.

While correlation imaging techniques provide a novel approach to the challenging NLoS problem, their performance is limited by the same d^4 intensity falloff that affects all NLoS methods. Note that because of eye safety limitations, this falloff cannot be overcome simply by increasing the laser power. Similarly, reducing the falloff to d^2 by collocating the illumination source and the hidden object, as was done in previous correlation-based NLoS




Fig. 1. NLoS imaging setup. A camera uses light scattered off of a rough wall, known as a virtual detector, to reconstruct an image of the hidden object. When using a continuous-wave laser, the camera records speckle. Inset, NLoS correlography estimates an object's autocorrelation using speckle images. It then uses this autocorrelation estimate to recover the object's shape by solving a PR problem.

experiments [13,15], is impractical for the vast majority of NLoS use cases. Thus, real-time NLoS imaging at even moderate stand-off distances is fundamentally a low-flux problem. As we will demonstrate in Section 6, existing PR methods [20–28] are too slow and sensitive to successfully operate in these photon-starved environments.

Deep-learning has recently achieved state-of-the-art performance on a range of challenging imaging inverse problems, such as super-resolution microscopy [29,30], lensless imaging [31], ghost imaging [32], and imaging through scattering media [33–36]. However, these existing approaches depend on large supervised training data and generally use shift-variant loss functions that do not translate to the NLoS correlography problem, Eq. (1). While we tackle the latter by introducing translation-invariant loss functions, acquiring training data experimentally is infeasible (consider the large combinatorial space of potential hidden scenes). To lift this limitation, we derive an accurate noise model for the NLoS correlography PR problem, which enables us to generate accurate training data synthetically. With the training data in hand, we propose and validate a learned reconstruction approach to NLoS correlography in the low-flux regime. In doing so, we make the following contributions.

- We use results from spectral density estimation to analyze the distribution of the noise associated with NLoS correlography.
- We propose a new approach for generating PR training data, without the need of experimental acquisition or modeling scene semantics.
- We propose a new mapping for the PR problem and propose and analyze several new, translation-invariant loss functions for learning-based PR.
- We validate our CNN on experimental NLoS imaging data, where it proves to be far faster and more robust than traditional methods; it enables reconstructions at a 300 μm resolution at 1 m standoff, using just two $1/8$ s exposure-length images.

A. Strengths and Limitations

In contrast to most NLoS imaging methods, our approach requires only standard CW laser sources and complementary metal oxide semiconductor (CMOS) sensors, solves the reconstruction problem in a fraction of a second, and does not need to know the location of the virtual source and detector. This significantly enhances the utility and practicality of our approach, bypassing the need for ultrafast sources and detectors, computationally expensive techniques for recovering a latent image from measurements, and impractical calibration steps. On the flip side, like other correlation-based methods and unlike time-of-travel-based techniques, our technique is best suited for imaging small isolated objects within the hidden volume; large objects lie outside the range of the memory effect [37] and so do not cause the self-interference upon which correlation-based techniques rely. Likewise, because of the translation invariance of the PR problem, our system is unable to localize the position of the objects within the hidden volume.

2. RELATED WORK

A. Impulse Non-Line-of-Sight Imaging

Kirmani *et al.* [1] first described the concept of imaging occluded objects using temporally coded measurement in which short pulses of light are captured propagating through the scene at the speed of light. These transient “light in flight” measurements are the temporal impulse response of light transport, and Abramson [38] first demonstrated a capture system for transient imaging. Velten *et al.* [2] showed the first experimental NLoS imaging results using a femtosecond laser and streak camera system to capture transient images. Building on these seminal works, a large body of work has explored impulse NLoS imaging [3–12], much of which is reviewed in Ref. [39]. These methods require detectors capable of high temporal resolution sampling to allow for impulse probing of the temporal light transport in the scene. Although the streak camera setup from Velten *et al.* [2] is essentially a decade-old technology, it allows for temporal precision of <10 ps. However, the high instrumentation cost and sensitivity of these experimental capture systems has sparked an interest in single photon avalanche diodes (SPADs) as a detector alternative [7,40]. Although SPAD detectors can offer resolution <10 ps [41], comparable to streak camera setups, they typically suffer from low fill factors typically around a few percent [42] and spatial resolution in the kilopixel range [43]. Compared to ubiquitous intensity sensors with pixel arrays of >10 megapixel and more, SPAD detectors are orders of magnitude less photon-efficient and more expensive.

Recently, a combination of noise robust algorithms [11,12] and extremely powerful illumination have moved these systems closer to real-time rates; using an illumination source with a peak power nearly $10000\times$ our own, [11] produced room-sized reconstructions in under 30 s. Additionally, Ref. [9] demonstrated real-time reconstructions of retroreflective objects; retroreflectors experience only a d^2 intensity falloff with confocal measurements [7].

B. Correlation Non-Line-of-Sight Imaging

Instead of directly acquiring transient transport measurements, a further line of research explores indirect coding using time-of-flight sensors [44–47]. Time-of-flight cameras capture correlation measurements of amplitude-modulated light, which encodes travel

time via the phase shift of the amplitude-modulated illumination. Although these cameras are readily available as consumer products, such as Microsoft's Kinect One, existing pixel technology offers only limited modulation bandwidths of around 100 MHz, truncating the effective temporal resolution to the nanosecond range.

Recently, an exciting line of work [13,15,19], loosely based off of ideas first developed in [37,48], has explored using correlations in the carrier wave itself, instead of amplitude modulation. This approach enables the use of conventional intensity sensors, while offering high modulation bandwidths in the THz range. Although seemingly a solution to the bandwidth and detector limitations of previous methods, existing approaches have been limited to scenes at microscopic scales [15] and lab setups with ambient light completely eliminated and have colocated the hidden object and the illumination source.

In this paper, we demonstrate a method that overcomes these issues by relying on temporally and spatially coherent measurements and a robust reconstruction framework, together allowing us to achieve photon-efficient NLoS imaging in the presence of strong ambient light and at large distances. A comparison of the two methods is provided in Section 3 of Supplement 1.

C. Non-Line-of-Sight Tracking

Recently, a number of methods for the orthogonal task of tracking and/or classifying occluded objects have been demonstrated using intensity measurements only [49–56]. Although these methods also rely on conventional intensity sensors, most are restricted to coarse localization tasks and often assume known object classes and scene geometries or require colocating the hidden object and the illumination source.

D. Traditional Phase Retrieval Algorithms

PR solves the problem of recovering a signal from a measurement of its Fourier magnitude (modulus). As the phase is lost in the measurement, this inverse problem is ill-posed in general. However, if the measurements are oversampled sufficiently, in theory, phase can be perfectly recovered by solving a set of nonlinear equations [57]. Together with an assumption on a nonzero support of the real-valued signal, practical error reduction methods have been designed [20] for a plethora of applications in optics, crystallography, biology, and physics. A popular extension of this algorithm is the hybrid input–output (HIO) method [21] and its various relaxations [22,58]. Recently, two major lines of research have explored alternating direction methods for PR [23,59], and overcoming the nonconvex nature of PR through convex relaxations [60].

E. CNN-Based Phase Retrieval Algorithms

Only very recently, deep neural networks have been explored for solving PR. Most previous attempts to apply convolutional neural networks (CNNs) to PR have been application specific. CNNs have been applied to ptychography [61,62], holography [63], quantitative phase microscopy [64,65], and coherent diffraction imaging (CDI) [66,67]. Among these works, the CDI approaches are the closest to our own.

There have also been attempts to use CNNs as regularizers within a PR optimization problem [68]. Unfortunately, these techniques require accurate initializations to succeed, which are

not available in the very low signal-to-noise ratio (SNR) regimes characteristic of NLoS imaging.

3. NON-LINE-OF-SIGHT CORRELOGRAPHY

A. Principles of Operation

The canonical NLoS imaging geometry of Fig. 1 will be used to introduce and develop the mathematical concepts underlying NLoS imaging. In this setup, a quasi-monochromatic laser source illuminates a portion of the visible wall, which we dub the “virtual source.” Laser light scattered by the optically rough virtual source surface propagates towards the hidden object. Due to the coherence of the laser source, the hidden object is illuminated by a fully developed speckle pattern, characterized by randomized constructive and destructive interference. A fraction of the light incident on the object is redirected (reflected and scattered) towards a second section of the visible wall, which we dub the “virtual detector.” A camera observing the virtual detector surface records the light reflected and scattering by the object. By scanning the position of the virtual source, a second statistically independent speckle realization is used to indirectly illuminate the hidden object. The corresponding camera image of the virtual detector surface, which is itself a low-contrast speckle pattern, is recorded. Using two or more measurements of this form, the autocorrelation of the hidden object's albedo is estimated, and the albedo is reconstructed by solving a PR problem.

B. Measurement Model

In this section, we provide details of the NLoS correlography measurement process, seeking to identify the (idealized) probability distribution associated with each measurement. We denote the complex-valued optical fields incident and emitted by the virtual source, hidden object, and virtual detector with the variables $E_{VS_{in}}$, $E_{VS_{out}}$, $E_{O_{in}}$, $E_{O_{out}}$, and $E_{VD_{in}}$, respectively. We index spatial locations on the virtual source, the hidden object, and the virtual detector using x_{VS} , x_O , and x_{VD} , respectively. In the interest of mathematical simplicity, we assume that the distances between the virtual source, object, and virtual detector are such that propagation between them can be modeled by the Fraunhofer approximation, i.e., proportional to a Fourier transformation. The propagation operator may be generalized to Fresnel propagation by absorbing the quadratic phase factors intrinsic to propagation within the fields being propagated, without affecting the outcome of our analysis. We additionally assume that the virtual source surface is illuminated by a collimated beam at normal incidence so that

$$E_{VS_{in}}(x_{VS}) = 1. \quad (2)$$

The optically rough virtual source scrambles the phase of the light emerging from the virtual source surface, so that

$$E_{VS_{out}}(x_{VS}) = e^{j\theta_{VS_{out}}(x_{VS})}, \quad (3)$$

where the phase term $\theta_{VS_{out}}$ is uniformly distributed over the interval $[0, 2\pi]$ and where the autocorrelation function of the phase term is a Dirac delta function. Thus, for $x_{VS,1} \neq x_{VS,2}$, $E_{VS_{out}}(x_{VS,1})$ and $E_{VS_{out}}(x_{VS,2})$ are uncorrelated with respect to the distribution of $\theta_{VS_{out}}$.

The field emerging from the virtual source then undergoes far-field propagation on its way to the object, which can be modeled

by a Fourier transformation. Accordingly, the field incident on the hidden object is

$$E_{O_{in}} \propto \mathcal{F}(E_{VS_{out}}), \quad (4)$$

where \mathcal{F} denotes the Fourier transform operator.

From the central limit theorem, the independence of the fields at different locations of $E_{VS_{out}}$, and the orthonormality of the Fourier transform, we have that, for all x_O , $E_{O_{in}}(x_O)$ follows a circular Gaussian distribution with autocorrelation function $\sigma^2 \delta(\Delta x_O)$ for some constant σ^2 .

Each location on the object modulates the incoming field according to the albedo of the hidden object. Thus,

$$E_{O_{out}}(x_0) = E_{O_{in}}(x_0)r(x_0), \quad (5)$$

and $E_{O_{out}}(x_0)$ follows a circular Gaussian distribution with autocorrelation function $r\sigma^2\delta(\Delta x_O)$.

The field emerging from the object propagates towards the virtual detector, in accordance with the far-field propagator. Thus,

$$E_{VD_{in}} \propto \mathcal{F}(E_{O_{out}}). \quad (6)$$

The camera then images the virtual detector. If we assume that the camera has infinite aperture and the virtual detector has uniform albedo, we have

$$I \propto |\mathcal{F}(E_{O_{out}})|^2. \quad (7)$$

One can illuminate different locations on the virtual source to capture independent measurements that follow Eq. (7). Images can additionally be subdivided into nonoverlapping patches and treated independently, e.g., a 2000×2000 image can be subdivided into 25 independent 400×400 patches [18,69].

C. Autocorrelation Estimation

Correlography takes measurements that follow the distribution specified by Eq. (7) and recovers the hidden object's albedo using the following two equalities. These equalities are derived in Section 1 in Supplement 1 and are based on the law of large numbers,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N |\mathcal{F}^{-1}(I_n)|^2 = r \star r(\Delta x_O) + \delta(\Delta x_O) \left[\int_{x_1=-\infty}^{\infty} r(x_1) dx_1 \right]^2, \quad (8)$$

and

$$\lim_{N \rightarrow \infty} \left| \frac{1}{N} \sum_{n=1}^N \mathcal{F}^{-1}(I_n) \right|^2 = \delta(\Delta x_O) \left[\int_{x_1=-\infty}^{\infty} r(x_1) dx_1 \right]^2. \quad (9)$$

By subtracting Eq. (9) from Eq. (8), we are left with an unbiased estimate of $r \star r$, which we denote with $\widehat{r \star r}$. An example of such an estimate is shown in Fig. 2(c).

The autocorrelation is related to the Fourier magnitudes of r through Eq. (1). Thus, by taking the Fourier transform of our estimate of $r \star r$ and then applying a PR algorithm to the result, we can recover an estimate of the albedo r . An example of one such




Fig. 2. Long-exposure NLoS correlography example. 25 nonoverlapping 400×400 speckle subimages were drawn from each of 50 distinct, 1 s exposure-length speckle images. These subimages were then used to estimate the hidden object's autocorrelation (middle) using Eqs. (8) and (9). HIO [21] was then used to reconstruct the object's albedo (right). (a) Hidden object; (b) speckle subimage; (c) estimate of $r \star r$; (d) estimate of r .




Fig. 3. Short-exposure NLoS correlography example. 25 nonoverlapping 400×400 speckle subimages were drawn from each of 2 distinct, 1/8 s exposure-length speckle images. These subimages were then used to estimate the hidden object's autocorrelation (middle) using Eqs. (8) and (9). HIO [21] was then used to reconstruct the object's albedo (right). (a) Hidden object; (b) speckle subimage; (c) estimate of $r \star r$; (d) estimate of r .

reconstruction, recovered using the HIO algorithm [21], is shown in Fig. 2(d).

Figures 3(c) and 3(d) repeat this experiment using two short exposure measurements, which produce a much noisier autocorrelation estimate: HIO fails to recover any structure in this context. Understanding and overcoming this noise is the key to enabling real-time NLoS imaging with correlography.

4. CORRELOGRAPHY NOISE MODEL

This section describes the fluctuations of $\widehat{r \star r}$ due to various sources of noise in the measurement process.

A. Distribution of the Autocorrelation Estimate

In practice, for locations $\Delta x \neq 0$, Eq. (8) is much greater than Eq. (9). Thus, for $\Delta x \neq 0$,

$$\widehat{r \star r}(\Delta x) \approx \frac{1}{N} \sum_{n=1}^N |\mathcal{F}^{-1} I_n|^2(\Delta x).$$

The expression $\frac{1}{N} \sum_{n=1}^N |\mathcal{F}^{-1} I_n|^2$ is the average of N i.i.d. random variables. From the central limit theorem, the elements of $\frac{1}{N} \sum_{n=1}^N |\mathcal{F}^{-1} I_n|^2$ follow a Gaussian distribution. Therefore, for $\Delta x \neq 0$

$$\widehat{r \star r}(\Delta x) \sim \mathcal{N}(\mu(\Delta x), \sigma^2(\Delta x))$$

for some mean $\mu(\Delta x)$ and variance $\sigma^2(\Delta x)$.

Next, we note that $\frac{1}{N} \sum_{n=1}^N |\mathcal{F}^{-1} I_n|^2$ is the average of N periodograms of the speckle images (because I_n is real, $|\mathcal{F}^{-1} I_n|^2 = |\mathcal{F} I_n|^2$). In the context of power spectral density

(PSD) estimation, averaging together multiple periodograms is known as Bartlett's procedure. Recognizing this fact allows us to rely upon existing theory to analyze $\widehat{r \star r}$.

Bartlett's procedure produces an unbiased estimate of the true PSD, $S(\Delta x)$. Moreover, if we assume I_n follows a Gaussian distribution, the pointwise variance of this estimate is proportional to $\frac{1}{N} S^2(\Delta x)$ [70]. The assumption is justified because I_n follows a noncentral chi-squared distribution with M degrees of freedom, where M denotes the dimension of the signal and, for large M , this can be approximated by a Gaussian distribution [71]. In summary, for $\Delta x \neq 0$, we have

$$\widehat{r \star r}(\Delta x) \sim \mathcal{N}(S(\Delta x), \gamma \frac{1}{N} S^2(\Delta x)), \quad (10)$$

where $S(\Delta x)$ is the PSD of the speckle at Δx and γ is some constant.

B. Sources of Noise

Multiple sources of noise and bias influence the NLoS correlography measurement process, the most important of which are

1. **Finite-sample approximation error:** Using a few samples (small N) will increase the variance of $\widehat{r \star r}$. It has no effect on the PSD of the speckle.
2. **Photon noise:** When dealing with weak, third bounce signals, Poisson shot noise shows up on the measurement I_n . This shot noise is white and will add a uniform offset to $S(\Delta x)$.
3. **Ambient illumination:** The measurements capture light not only from the hidden object, but also from walls, floors, and clutter in the scene. This shows up as both a diffuse background and uncorrelated shot noise in the speckle the measurements. The diffuse background will add a peak at $S(\Delta x = 0)$, and the shot noise adds a uniform offset to $S(\Delta x)$.
4. **Finite apertures:** The finite apertures of both the camera and the virtual detector mean that our measurements are low-pass filtered. This band limits the PSD of the speckle.

Combining these sources of noise, we get

$$S(\Delta x) = H(r \star r(\Delta x) + b),$$

where $H(\cdot)$ is the aperture's low-pass transfer function and b is an offset accounting for the shot noise (both from the hidden object and the background). Assuming our camera has a sufficiently large aperture, this model simplifies to

$$S(\Delta x) = r \star r(\Delta x) + b.$$

Combining this result with Eq. (10), with some slight abuse of notation, we have

$$\begin{aligned} \widehat{r \star r}(\Delta x) &\approx \mathcal{N}(r \star r(\Delta x) + b, \gamma \frac{1}{N} (r \star r(\Delta x) + b)^2) \\ &\approx r \star r(\Delta x) + \mathcal{N}(0, \frac{\gamma}{N} (r \star r(\Delta x))^2) + \mathcal{N}(b, \frac{\gamma}{N} b^2). \end{aligned} \quad (11)$$

In the above equation, we eliminate the variance's cross terms by assuming that, on the support of $r \star r(\Delta x)$, $r \star r(\Delta x) \gg b$.

In this expression, the signal-dependent, space-varying noise (second term) is due primarily to finite-sample approximation error. Figure 4(d) illustrates what happens when too few high SNR




Fig. 4. Spending the photon budget. Simulated autocorrelation estimates of a Lambertian 1 cm \times 1 cm letter "V" using (b) 100, 0.1 s measurements; (c) 10, 1 s measurements; and (d) 1, 10 s measurement. With many short-exposure measurements, the spatially invariant noise dominates. With too few measurements, shot-noise-like signal-dependent noise from the finite-sample error dominates. (a) Ground truth; (b) $100 \times .1$; (c) 10×1 ; (d) 1×10 .

measurements are used to estimate $\widehat{r \star r}$; error-type 1 dominates, and shot-noise-like, strongly signal-dependent noise shows up in the estimate.

In contrast, the offset and the spatially invariant noise (third term) are due primarily to the shot noise. Figure 4(b) illustrates the case when many low SNR measurements with significant shot noise components are used to estimate $\widehat{r \star r}$; error type 2 dominates, and an offset and uniform Gaussian noise appears in the estimate.

Finally, Fig. 4(c) demonstrates that for a large-enough photon budget, there is a Goldilocks zone wherein the shot-noise in the measurements is reduced but there are still enough samples to avoid finite-sample approximation error. Photons budgets should be spent so as to operate in this regime. In this paper, we found our photon budget was best spent capturing just two high-resolution speckle images (each broken into 64 smaller image patches, resulting in $N = 128$).

C. Validating the Noise Model

To further validate the proposed noise model, we inspect a series of experimental autocorrelation estimates, which were formed with a varying number of speckle images and with a variety of exposure times. For each autocorrelation estimate, we consider the statistics of a 20×20 patch in the top-left corner of the estimate, which, by visual inspection, contains no signal component. Our model predicts that, across estimates, such regions should be distributed according to $\mathcal{N}(b, \frac{\gamma}{N} b^2)$. Assuming these regions are ergodic, they should exhibit a similar distribution across the pixels of a single estimate.




Fig. 5. Distribution of experimental autocorrelation estimates. Variance versus mean with varying exposure and fixed $N = 128$ (left) and variance over mean with fixed exposure and a varying N (right). As predicted by our model Eq. (11), the variance of our $r \star r$ estimate grows quadratically with respect to its mean, and the ratio between the variance and mean grows linearly with respect to $\frac{1}{N}$.




Fig. 6. Unstructured training data. Examples of images formed with a Canny edge detector (top) and their associated noisy autocorrelations (bottom).

In Fig. 5, we plot the variance and mean of these patches (across pixels). We observe that as the mean increases the variance grows quadratically, as predicted. Furthermore, we see that as we reduce the number of speckle images used to form the estimates (N), the ratio between the variance and mean of the patches grows proportionally to $\frac{1}{N}$, as predicted.

5. LEARNING PHASE RETRIEVAL

As mentioned before, existing PR algorithms are not up to the task of solving the noisy PR problem associated with low-light NLoS correlography. In this section, we describe how we applied a CNN to the problem.

A. Training Datasets

Deep learning is a powerful tool for solving computational imaging problems, but requires vast quantities of training data to succeed. In the context of NLoS imaging, this training data is very hard to come by experimentally. Therefore, we leverage the noise model developed in the previous section to synthesize training data.

We generated training data consisting of $r, r \star \hat{r}$ pairs, where the $r \star \hat{r}$ examples were synthesized according to Eq. (11) with $b = 70$ and $\frac{\gamma}{N} = 0.015$, where elements of r are scaled such that $\max(r \star r) = 255$. These parameters were chosen by fitting the noise model to the mean (b) and variance ($\frac{\gamma}{N}b^2$) of a 20×20 patch from the top left corner of the autocorrelation estimate formed by a $1/8$ s exposure measurement.

The dataset used for r determines what priors the network learned about the problem—and how well it generalizes to different problems. In this paper, we train a CNN using a dataset of sparse, “unstructured” images (at the SNRs we are interested in, reconstructing dense, “natural image” scenes is infeasible). This dataset was formed by passing the Berkeley Segmentation Dataset 500 [72] through a Canny edge detector and then cropping to form a dataset of roughly 20000064×64 images with sparse edges. The images in this dataset are connected and sparse, but otherwise lack much structure. See Fig. 6.

B. Loss Function

One challenge in learning PR is that with Fourier measurement operators, PR is invariant to translations. Thus, training a neural network using a loss that is not invariant, such as the ℓ_1 or ℓ_2 distance between r and \hat{r} , would force the network to not only solve the PR problem but also memorize the locations of all the training data.




Fig. 7. Noisy estimates of (a) $r \star r$, (b) $|\mathcal{F}(r)|$ and (c) the associated r . $r \star r$ and r share similar features whereas $|\mathcal{F}(r)|$ and r , for the most part, do not.

To avoid this, we experimented with four different translation invariant losses: $\|\hat{r} \star \hat{r} - r \star r\|_1$, $\|\hat{r} \star \hat{r} - r \star r\|_2$, $\||\mathcal{F}(\hat{r})| - |\mathcal{F}(r)|\|_1$, and $\||\mathcal{F}(\hat{r})| - |\mathcal{F}(r)|\|_2$, where \hat{r} denotes the network's estimate of r . We found that the loss $\|\hat{r} \star \hat{r} - r \star r\|_1$ converged quicker than the others but that all four losses eventually produced similar solutions (the Pearson correlation coefficient may also have been effective [66]). See Figures S1 and S2 in Supplement 1. We use $\|\hat{r} \star \hat{r} - r \star r\|_1$ as the loss function throughout the rest of the paper.

C. Choosing the Mapping

Although in principle a CNN can learn almost arbitrary mappings between a measurement of r and the signal r , certain mappings are easier to learn than others. Neural networks with skip or residual connections excel at learning identity-like mappings [73]. As illustrated in Fig. 7, the mapping from $r \star r$ to r is much closer to an identity mapping ($r \star r$ and r share similar features) than the mapping from $|\mathcal{F}(r)|$ to r is. As such, we found that networks trained to recover r from an estimate of its autocorrelation did much better than those trained to reconstruct it from an estimate of its Fourier magnitudes.

D. CNN Architecture

We used the well-known U-net architecture as our CNN [74]. Our U-net consists of 12 convolutional layers, each with between 64 and 512 channels. The U-net is essentially a convolutional autoencoder with a large number of skip connections between its layers. While originally designed for segmentation, it has been applied successfully to a range of imaging inverse problems, such as the reconstruction of medical images [75] and low-light denoising [76].

E. Implementation and Training Details

Our network was implemented in PyTorch. We used a batch size of 32. We trained for 400 epochs at a learning rate starting at 0.002 and decaying to 0.000001, using the ADAM optimizer [77]. It took a little over a day to train the network using an Nvidia Titan RTX GPU. (Our code is available at [78].)

6. LOW-LIGHT NLoS CORRELOGRAPHY

In this section, we use simulations and experiments to validate the proposed approach and answer the following question: does the increased noise robustness afforded by the proposed learned PR enable real-time NLoS correlography?




Fig. 8. Experimental setup. Light passes from the laser, to the virtual source, to the hidden object, to the virtual detector, and finally to the camera.

A. Experimental Setup

Our NLoS correlography experimental setup is illustrated in Fig. 8. A steerable, 500 mW, 532 nm CW laser source (Azur Light Systems ALS-532) illuminates the virtual source. We operated the laser at 300 mW. A Canon telephoto macrolens with 180 mm focal length is used to image the virtual detector surface. With this lens, the image sensor (2056 × 2464 pixel Sony IMX 264 monochrome) has a magnification of about 0.5, a pixel size of 3.45 μm, and an active area of 8.47 mm × 7 mm. (We removed the camera’s cover glass to reduce internal reflections.)

We imaged the 1 cm hidden figures from Figs. 2(a) and 3(a), which had an average fill rate of about 0.2 (they occupy about 0.2 cm²). Our virtual source was 0.5 m from the hidden object, the hidden object was 1 m from the virtual detector, and the virtual detector was 0.8 meters from the camera.

B. Radiometric Throughput

Assuming the walls and target have albedo 1 and are perfectly Lambertian (they are not), the radiometric throughput of our system works out to −182 dB. With a 300 mW, 532 nm laser source, this translates to 2.7 million third bounce photons per second reaching the detector. These calculations can be found in Section D in Supplement 1.

C. Phase Retrieval Algorithms

Here we compare our CNN to the projection-based HIO PR algorithm [21] and the alternating minimization PR algorithm (Alt-Min) from [24]. Additional results with median truncated Wirtinger flow (MTWF) [25], truncated amplitude flow (TAF) [26], truncated Wirtinger flow (TWF) [27], and alternating direction method of multipliers with a total variation prior (ADMM-TV) [28] can be found in Section E of Supplement 1. These methods perform no better than HIO.

HIO was run for 2000 iterations with a step size $\beta = .9$. The support was assumed to be 64 × 64, out of 128 × 128 pixels total. Alternating minimization was run for 1000 iterations.

D. Recovery Times

With 2 speckle images, estimating the autocorrelation, $r \hat{\star} r$, took $\frac{1}{10}$ s. From there, HIO took just under 3 s to reconstruct r . The CNN took a few hundredths of a second. Exposure times, not processing, are the bottleneck in a CNN-based system.

E. Low-Light Imaging Simulations

Using the throughput estimates from the previous section, we simulated NLoS correlography with exposure times between 1/4 and 1/256 s using a 300 mW CW laser with hidden objects with a reflective area of roughly 0.2 cm². For each of the exposure times, we captured 2 images, each consisting of 64 patches.




Fig. 9. Simulated and experimental reconstructions with varying exposure lengths. Because it is more robust to noise, the CNN-based method can operate with far less light and thus at higher frame rates than a system relying on traditional PR algorithms like HIO [21] or Alt-Min [24]. (The vertical/horizontal lines that can be observed in the experimental short-exposure autocorrelation estimates are the result of correlated, fixed pattern read noise between pixels. The 7 and F were measured at different orientations.)

Our simulation results, presented in Fig. 9(a), demonstrate that for a given laser power, because it is more robust to noise, the CNN-based method can operate with reduced exposure times, and thus at higher frame rates.

F. Low-Light Imaging Experiments

We applied the CNN-based PR method to experimental NLoS imaging data consisting of multiple low exposure measurements of the objects from Figs. 2(a) and 3(a). Figure 9(b) demonstrates that the CNN is significantly more robust to noise and offers improved reconstructions across all operating regimes. The CNN offers recognizable reconstructions starting around exposure lengths of $\frac{1}{16}$ s. Additional results can be found in Supplement 1.

7. CONCLUSION

NLoS correlography promises, in theory, to enable real-time NLoS imaging at sub-mm resolutions. However, the limitations of existing PR methods, particularly their sensitivity to noise, prohibit this. More broadly, because of the quartic attenuation of intensity with distance, handling low-flux regimes is arguably the *fundamental barrier* to real-time NLoS imaging. This paper makes a step towards lifting these limitations.

Specifically, we first analyzed the NLoS correlography noise model. This analysis makes it possible to simulate adequate training data for learning NLoS correlography problems. Using the proposed dataset and new loss function, we then trained a CNN to solve the noisy PR problem. In simulation, we confirmed that the resulting CNN is computationally efficient and exceptionally robust to multiple forms of noise, far exceeding the capabilities of existing algorithms. We validated our approach on experimental NLoS imaging data and successfully reconstructed the shape of small hidden objects from a standoff distance of one meter away using just two $\frac{1}{8}$ s exposure-length images captured by a conventional CMOS detector, representing a significant step towards real-time high-resolution NLoS imaging.

Funding. Defense Advanced Research Projects Agency (Reveal: HR0011-16-C-0028.).

See Supplement 1 for supporting content.

†These authors contributed equally to this work.

REFERENCES

1. A. Kirmani, T. Hutchison, J. Davis, and R. Raskar, "Looking around the corner using transient imaging," in *Proceedings of IEEE International Conference on Computer Vision* (2009), pp. 159–166.
2. A. Velten, D. Wu, A. Jarabo, B. Masia, C. Barsi, C. Joshi, E. Lawson, M. Bawendi, D. Gutierrez, and R. Raskar, "Femto-photography: capturing and visualizing the propagation of light," *ACM Trans. Graphics* **32**, 44 (2013).
3. R. Pandharkar, A. Velten, A. Bardagiy, E. Lawson, M. Bawendi, and R. Raskar, "Estimating motion and size of moving non-line-of-sight objects in cluttered environments," in *Proc. of IEEE International Conference on Computer Vision and Pattern Recognition* (2011), pp. 265–272.
4. A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, M. Bawendi, and R. Raskar, "Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging," *Nat. Commun.* **3**, 745 (2012).
5. O. Gupta, T. Willwacher, A. Velten, A. Veeraraghavan, and R. Raskar, "Reconstruction of hidden 3D shapes using diffuse reflections," *Opt. Express* **20**, 19096–19108 (2012).
6. A. K. Pediredla, M. Buttafava, A. Tosi, O. Cossairt, and A. Veeraraghavan, "Reconstructing rooms using photon echoes: a plane based model and reconstruction algorithm for looking around the corner," in *Proceedings of IEEE International Conference on Computational Photography* (IEEE, 2017).
7. M. O'Toole, D. B. Lindell, and G. Wetzstein, "Confocal non-line-of-sight imaging based on the light cone transform," *Nature* **555**, 338–341 (2018).
8. 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* **26**, 9945–9962 (2018).
9. M. O'Toole, D. B. Lindell, and G. Wetzstein, "Real-time non-line-of-sight imaging," in *ACM SIGGRAPH 2018 Emerging Technologies* (ACM, 2018), paper 14.
10. S. Xin, S. Nousias, K. N. Kutulakos, A. C. Sankaranarayanan, S. G. Narasimhan, and I. Gkioulekas, "A theory of Fermat paths for non-line-of-sight shape reconstruction," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (2019), pp. 6800–6809.
11. X. Liu, I. Guillén, M. La Manna, J. H. Nam, S. A. Reza, T. H. Le, A. Jarabo, D. Gutierrez, and A. Velten, "Non-line-of-sight imaging using phasor-field virtual wave optics," *Nature* **572**, 620–623 (2019).
12. D. B. Lindell, G. Wetzstein, and M. O'Toole, "Wave-based non-line-of-sight imaging using fast f-k migration," *ACM Trans. Graphics* **38**, 116 (2019).
13. O. Katz, E. Small, and Y. Silberberg, "Looking around corners and through thin turbid layers in real time with scattered incoherent light," *Nat. Photonics* **6**, 549–553 (2012).
14. A. Viswanath, P. Rangarajan, D. MacFarlane, and M. P. Christensen, "Indirect imaging using correlography," in *Computational Optical Sensing and Imaging* (Optical Society of America, 2018), paper CM2E–3.
15. O. Katz, P. Heidmann, M. Fink, and S. Gigan, "Non-invasive single-shot imaging through scattering layers and around corners via speckle correlations," *Nat. Photonics* **8**, 784–790 (2014).
16. P. S. Idell, J. R. Fienup, and R. S. Goodman, "Image synthesis from non-imaged laser-speckle patterns," *Opt. Lett.* **12**, 858–860 (1987).
17. J. R. Fienup and P. S. Idell, "Imaging correlography with sparse arrays of detectors," *Opt. Eng.* **27**, 279778 (1988).
18. P. S. Idell, J. D. Gonglewski, D. G. Voelz, and J. Knopp, "Image synthesis from nonimaged laser-speckle patterns: experimental verification," *Opt. Lett.* **14**, 154–156 (1989).
19. J. Bertolotti, E. G. van Putten, C. Blum, A. Lagendijk, W. L. Vos, and A. P. Mosk, "Non-invasive imaging through opaque scattering layers," *Nature* **491**, 232–234 (2012).
20. R. W. Gerchberg, "A practical algorithm for the determination of phase from image and diffraction plane pictures," *Optik* **35**, 237–246 (1972).
21. J. R. Fienup, "Phase retrieval algorithms: a comparison," *Appl. Opt.* **21**, 2758–2769 (1982).
22. H. H. Bauschke, P. L. Combettes, and D. R. Luke, "Hybrid projection-reflection method for phase retrieval," *J. Opt. Soc. Am. A* **20**, 1025–1034 (2003).
23. S. Marchesini, Y.-C. Tu, and H.-T. Wu, "Alternating projection, Ptychographic imaging and phase synchronization," *Appl. Comput. Harmon. Anal.* **41**, 815–851 (2016).
24. P. Netrapalli, P. Jain, and S. Sanghavi, "Phase retrieval using alternating minimization," in *Advances in Neural Information Processing Systems* (2013), pp. 2796–2804.
25. H. Zhang, Y. Chi, and Y. Liang, "Provable non-convex phase retrieval with outliers: median truncated Wirtinger flow," in *Proc. International Conference on Machine Learning* (2016), pp. 1022–1031.
26. G. Wang, G. B. Giannakis, and Y. C. Eldar, "Solving systems of random quadratic equations via truncated amplitude flow," *IEEE Trans. Inf. Theory* **64**, 773–794 (2018).
27. Y. Chen and E. Candes, "Solving random quadratic systems of equations is nearly as easy as solving linear systems," in *Advances in Neural Information Processing Systems* (2015), pp. 739–747.
28. F. Heide, S. Diamond, M. Nießner, J. Ragan-Kelley, W. Heidrich, and G. Wetzstein, "Proximal: efficient image optimization using proximal algorithms," *ACM Trans. Graphics* **35**, 84 (2016).
29. Y. Rivenson, Z. Göröcs, H. Günaydin, Y. Zhang, H. Wang, and A. Ozcan, "Deep learning microscopy," *Optica* **4**, 1437–1443 (2017).

30. E. Nehme, L. E. Weiss, T. Michaeli, and Y. Shechtman, "Deep-storm: super-resolution single-molecule microscopy by deep learning," *Optica* **5**, 458–464 (2018).
31. A. Sinha, J. Lee, S. Li, and G. Barbastathis, "Lensless computational imaging through deep learning," *Optica* **4**, 1117–1125 (2017).
32. F. Wang, H. Wang, H. Wang, G. Li, and G. Situ, "Learning from simulation: an end-to-end deep-learning approach for computational ghost imaging," *Opt. Express* **27**, 25560–25572 (2019).
33. Y. Li, Y. Xue, and L. Tian, "Deep speckle correlation: a deep learning approach towards scalable imaging through scattering media," *Optica* **5**, 1181–1190 (2018).
34. S. Li, M. Deng, J. Lee, A. Sinha, and G. Barbastathis, "Imaging through glass diffusers using densely connected convolutional networks," *Optica* **5**, 803–813 (2018).
35. Y. Sun, Z. Xia, and U. S. Kamilov, "Efficient and accurate inversion of multiple scattering with deep learning," *Opt. Express* **26**, 14678–14688 (2018).
36. M. Lyu, H. Wang, G. Li, S. Zheng, and G. Situ, "Learning-based lensless imaging through optically thick scattering media," *Adv. Photon.* **1**, 036002 (2019).
37. I. Freund, M. Rosenbluh, and S. Feng, "Memory effects in propagation of optical waves through disordered media," *Phys. Rev. Lett.* **61**, 2328–2331 (1988).
38. N. Abramson, "Light-in-flight recording by holography," *Opt. Lett.* **3**, 121–123 (1978).
39. T. Maeda, G. Satat, T. Swedish, L. Sinha, and R. Raskar, "Recent advances in imaging around corners," arXiv:1910.05613 (2019).
40. 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* **23**, 20997–21011 (2015).
41. F. Nolet, S. Parent, N. Roy, M.-O. Mercier, S. Charlebois, R. Fontaine, and J.-F. Pratte, "Quenching circuit and SPAD integrated in CMOS 65 nm with 7.8 ps FWHM single photon timing resolution," *Instruments* **2**, 19 (2018).
42. L. Parmesan, N. A. Dutton, N. J. Calder, A. J. Holmes, L. A. Grant, and R. K. Henderson, "A 9.8 μm sample and hold time to amplitude converter CMOS SPAD pixel," in *44th European Solid State Device Research Conference (ESSDERC)* (IEEE, 2014), pp. 290–293.
43. Y. Maruyama and E. Charbon, "A time-gated 128 \times 128 CMOS SPAD array for on-chip fluorescence detection," in *Proceedings International Image Sensor Workshop (IISW)* (2011).
44. F. Heide, M. B. Hullin, J. Gregson, and W. Heidrich, "Low-budget transient imaging using photonic mixer devices," *ACM Trans. Graphics* **32**, 45 (2013).
45. A. Kadambi, R. Whyte, A. Bhandari, L. Streeter, C. Barsi, A. Dorrington, and R. Raskar, "Coded time of flight cameras: sparse deconvolution to address multipath interference and recover time profiles," *ACM Trans. Graphics* **32**, 167 (2013).
46. F. Heide, L. Xiao, W. Heidrich, and M. B. Hullin, "Diffuse mirrors: 3D reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (2014), pp. 3222–3229.
47. A. Kadambi, H. Zhao, B. Shi, and R. Raskar, "Occluded imaging with time-of-flight sensors," *ACM Trans. Graphics* **35**, 15 (2016).
48. I. Freund, "Looking through walls and around corners," *Phys. A* **168**, 49–65 (1990).
49. J. Klein, C. Peters, J. Martín, M. Laurenzis, and M. B. Hullin, "Tracking objects outside the line of sight using 2D intensity images," *Sci. Rep.* **6**, 32491 (2016).
50. P. Caramazza, A. Boccolini, D. Buschek, M. Hullin, C. F. Higham, R. Henderson, R. Murray-Smith, and D. Faccio, "Neural network identification of people hidden from view with a single-pixel, single-photon detector," *Sci. Rep.* **8**, 11945 (2018).
51. S. Chan, R. E. Warburton, G. Garipey, J. Leach, and D. Faccio, "Non-line-of-sight tracking of people at long range," *Opt. Express* **25**, 10109–10117 (2017).
52. 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 *Proceedings of IEEE International Conference on Computer Vision* (2017), Vol. 1, pp. 8.
53. B. M. Smith, M. O'Toole, and M. Gupta, "Tracking multiple objects outside the line of sight using speckle imaging," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (2018), pp. 6258–6266.
54. M. Tancik, G. Satat, and R. Raskar, "Flash photography for data-driven hidden scene recovery," arXiv:1810.11710 (2018).
55. C. Saunders, J. Murray-Bruce, and V. K. Goyal, "Computational periscope with an ordinary digital camera," *Nature* **565**, 472–475 (2019).
56. M. Batarseh, S. Sukhov, Z. Shen, H. Gemar, R. Rezvani, and A. Dogariu, "Passive sensing around the corner using spatial coherence," *Nat. Commun.* **9**, 3629 (2018).
57. R. Bates, "Fourier phase problems are uniquely solvable in mute than one dimension. I: underlying theory," *Optik (Stuttgart)* **61**, 247–262 (1982).
58. D. R. Luke, "Relaxed averaged alternating reflections for diffraction imaging," *Inverse Probl.* **21**, 37 (2004).
59. Z. Wen, C. Yang, X. Liu, and S. Marchesini, "Alternating direction methods for classical and ptychographic phase retrieval," *Inverse Probl.* **28**, 115010 (2012).
60. E. J. Candes, T. Strohmer, and V. Voroninski, "Phaselift: exact and stable signal recovery from magnitude measurements via convex programming," *Commun. Pure Appl. Math.* **66**, 1241–1274 (2013).
61. A. Kappeler, S. Ghosh, J. Holloway, O. Cossairt, and A. Katsaggelos, "Ptychnet: CNN based Fourier ptychography," in *Proceedings of IEEE International Conference on Image Processing (IEEE, 2017)*, pp. 1712–1716.
62. L. Boominathan, M. Maniparambil, H. Gupta, R. Baburajan, and K. Mitra, "Phase retrieval for Fourier ptychography under varying amount of measurements," arXiv:1805.03593 (2018).
63. Y. Rivenson, Y. Zhang, H. Günaydin, D. Teng, and A. Ozcan, "Phase recovery and holographic image reconstruction using deep learning in neural networks," *Light Sci. Appl.* **7**, 17141 (2018).
64. Z. Kemp, "Propagation based phase retrieval of simulated intensity measurements using artificial neural networks," *J. Opt.* **20**, 045606 (2018).
65. M. R. Kellman, E. Bostan, N. Repina, M. Lustig, and L. Waller, "Physics-based learned design: optimized coded-illumination for quantitative phase imaging," arXiv:1808.03571 (2018).
66. A. Goy, K. Arthur, S. Li, and G. Barbastathis, "Low photon count phase retrieval using deep learning," *Phys. Rev. Lett.* **121**, 243902 (2018).
67. M. J. Cherukara, Y. S. Nashed, and R. J. Harder, "Real-time coherent diffraction inversion using deep generative networks," *Sci. Rep.* **8**, 16520 (2018).
68. C. A. Metzler, P. Schniter, A. Veeraraghavan, and R. G. Baraniuk, "prDeep: robust phase retrieval with a flexible deep network," in *Proceedings International Conference on Machine Learning* (2018), pp. 3498–3507.
69. J. Fienup (private communication, 2017).
70. P. Welch, "The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms," *IEEE Trans. Audio Electroacoust.* **15**, 70–73 (1967).
71. R. J. Muirhead, *Aspects of Multivariate Statistical Theory* (Wiley, 2009), Vol. **197**.
72. D. Martin, C. Fowlkes, D. Tal, and J. Malik, "A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics," in *Proceedings of IEEE International Conference on Computer Vision* (2001), Vol. **2**, pp. 416–423.
73. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition* (2016), pp. 770–778.
74. O. Ronneberger, P. Fischer, and T. Brox, "U-net: convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention* (Springer, 2015), pp. 234–241.
75. K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Trans. Image Process.* **26**, 4509–4522 (2017).
76. C. Chen, Q. Chen, J. Xu, and V. Koltun, "Learning to see in the dark," in *Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition* (2018).
77. D. P. Kingma and J. Ba, "Adam: a method for stochastic optimization," arXiv:1412.6980 (2014).
78. https://github.com/ricedsp/Deep_Inverse_Correlography