

Imaging with Local Speckle Intensity Correlations: Theory and Practice

MARINA ALTERMAN and CHEN BAR, Department of Electrical Engineering, Technion

IOANNIS GKIΟULEKAS, Robotics Institute, Carnegie Mellon University

ANAT LEVIN, Department of Electrical Engineering, Technion

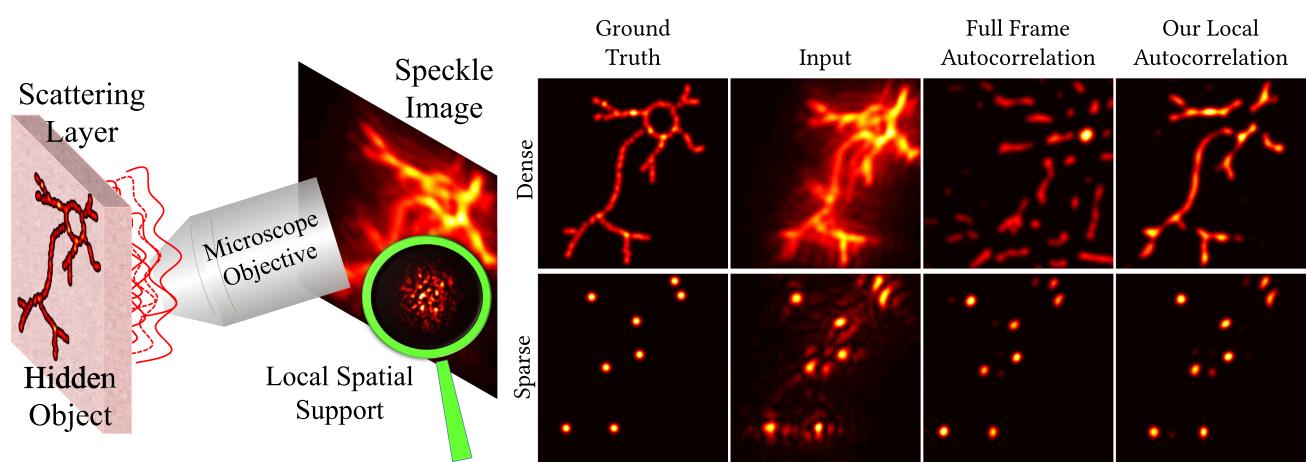


Fig. 1. Near-field imaging through scattering. A latent object, comprising mutually incoherent light sources, is seen through a scattering sample (e.g., fluorescent particles inside tissue). The image measured by the camera is degraded due to scattering. Previous approaches suggest that due to speckle statistics, the latent image can be recovered from the auto-correlation of the speckle image. Despite the potential of this idea, many previous experimental demonstrations considered sources a few centimeters behind the scattering layer rather than inside it, as would be desired in a real biological application. Here, we attempt to bring the algorithm from the far into the near field regime, taking advantage of a special characteristic of near field speckles: their local support. In the right part, both the classic full-frame auto-correlation and our approach can successfully recover the latent object when it is composed only of a small number of illuminating points (lower right). However, our approach can also recover a significantly denser object, while the classic algorithm fails (top right).

Recent advances in computational imaging have significantly expanded our ability to image through scattering layers such as biological tissues by exploiting the auto-correlation properties of captured speckle intensity patterns. However, most experimental demonstrations of this capability focus on the far-field imaging setting, where obscured light sources are very far from the scattering layer. By contrast, medical imaging applications such as fluorescent imaging operate in the near-field imaging setting, where sources are inside the scattering layer. We provide a theoretical and experimental study of the similarities and differences between the two settings, highlighting the increased challenges posed by the near-field

setting. We then draw insights from this analysis to develop a new algorithm for imaging through scattering that is tailored to the near-field setting by taking advantage of unique properties of speckle patterns formed under this setting, such as their local support. We present a theoretical analysis of the advantages of our algorithm and perform real experiments in both far-field and near-field configurations, showing an order-of-magnitude expansion in both the range and the density of the obscured patterns that can be recovered.

CCS Concepts: • Computing methodologies → Computational photography;

Additional Key Words and Phrases: Speckle, memory effect, scattering

ACM Reference format:

Marina Alterman, Chen Bar, Ioannis Gkioulekas, and Anat Levin. 2021. Imaging with Local Speckle Intensity Correlations: Theory and Practice. *ACM Trans. Graph.* 40, 3, Article 30 (July 2021), 22 pages. <https://doi.org/10.1145/3447392>

This work was supported by European Research Council Horizon 2020 635537, ISF 1947, the Ollendorff Minerva Center of the Technion, Technion System Eng Center, NSF-BSF award 2008123/2019758 and NSF Expeditions award 1730147. Authors' addresses: M. Alterman, C. Bar, and A. Levin, Department of Electrical Engineering, Technion; Haifa, 32000, Israel; emails: amarinago@gmail.com, chen.bar@campus.technion.ac.il, anat.levin@ee.technion.ac.il; I. Gkioulekas, Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue Pittsburgh, PA 15213; email: igkioule@andrew.cmu.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

0730-0301/2021/07-ART30 \$15.00

<https://doi.org/10.1145/3447392>

1 INTRODUCTION

Developing techniques for imaging through scattering layers, and in particular through layers of biological tissue, is a core challenge of modern imaging. The fundamental difficulty in achieving this objective is the fact that, when an incident wave propagates through such a layer, it interacts with its microstructure multiple times. For example, incident light arising from a single coherent

source will, after such a scattering process, result in images showing strong speckle patterns, spread over multiple pixels.

Despite their noise-like appearance, these speckle images have strong statistical properties relating to both the incident wave and the scattering material that produced them. One such property, which is the focus of this article, is the **memory effect (ME)**: This refers to the fact that speckle patterns are correlated and approximately shift-invariant with respect to small tilts in the illumination or viewing angles [Akkermans and Montambaux 2007; Baydoun et al. 2016; Dougherty et al. 1994; Osnabrugge et al. 2017; Freund and Eliyahu 1992; Berkovits and Feng 1994; Fried 1982; Feng et al. 1988]. This property underlies many recent techniques for producing clean images through scattering media [Katz et al. 2014; Bertolotti et al. 2012; Takasaki and Fleischer 2014; Edrei and Scarselli 2016a, 2016b; Hofer et al. 2018; Wu et al. 2017, 2020; Wang et al. 2020]. In particular, Katz et al. [2014] have demonstrated that it is possible to recover a clean, scattering-free image of a sparse set of mutually incoherent latent illuminators, observed through a thick scattering layer. Their algorithm works by simply computing the auto-correlation of the observed speckle image and performing phase retrieval [Fienup 1982]. This remarkable imaging capability has strong potential for applications in medical imaging, e.g., for imaging fluorescent cells beneath tissue and for performing non-invasive blood flow analysis.

Unfortunately, realizing this potential in practice remains difficult. One reason for this lies in the fundamental limitations restricting the applicability of memory effect algorithms: For instance, it is well-documented that the strength of the memory effect decays fast as the displacement between the latent illuminators increases [Wang et al. 2019; Schott et al. 2015]; and we show that it is also negatively affected as the density of the latent illuminators increases. Another reason is that nearly all previous demonstrations of imaging through scattering with the memory effect use illuminators and sensors that are in the *far field* of the scattering layer, i.e., placed at a large distance from it. (A notable exception is the work of Chang and Wetzstein [2018], which we discuss in detail later in the article.) This setting is reasonable for applications such as non-line-of-sight imaging [Smith et al. 2018; Freund 1990; Katz et al. 2012; Batarseh et al. 2018; Metzler et al. 2020; Viswanath et al. 2018; Boger-Lombard and Katz 2019]. However, it is unrealistic for tissue imaging applications, which typically require *near-field* imaging conditions. For example, in fluorescent imaging, the fluorescing particles are inside the scattering layer, rather than at a distance from it. Consequently, the current experimental protocol used in research papers for evaluating theory and algorithms about the memory effect is incompatible with how these theory and algorithms would be applied in medical imaging practice.

Our goal in this work is to draw attention to this incompatibility, show that its implications are significant, and propose ways to align research and practice. To this end, we begin with a detailed study of the memory effect in the near-field and far-field settings and highlight the differences between the two settings. In particular, we introduce a new theorem that allows us to draw direct analogies between the two settings and use physically accurate simulations [Bar et al. 2019, 2020] to both validate and generalize conclusions drawn from this theorem. The findings of our study suggest that, in the near-field setting, memory effect techniques

are only practical for scattering layers of modest optical depth, namely, scattering layers whose thickness is only a few mean free paths. In such layers, mid-order scattering is dominant, meaning that light undergoes a small number of scattering events. In tissue, this mid-order scattering regime corresponds to layers that are still well-beyond the maximum penetration depth of standard microscopes. Therefore, memory-effect techniques operating in this regime can be of great practical importance for medical imaging.

Based on this observation, we proceed to investigate how to improve imaging-through-scattering techniques in the mid-order scattering regime. Specifically, we document a property characteristic of speckle patterns arising due to mid-order scattering: The speckle pattern formed on a sensor due to a single latent illuminator is typically much smaller than the sensor. This *local support property* has not been studied in the past and, as we show, is key for enhancing the performance of imaging-through-scattering techniques in both the far-field and near-field settings.

In particular, we first derive an analytical expression for the signal-to-noise ratio that can be achieved when recovering scattering-free images using the memory effect. Our analysis suggests that there exists an optimal matched filter, corresponding to the local spatial support of speckle patterns, that maximizes this ratio. We then use this theoretical result to motivate and develop a new algorithm for using the memory effect to image through scattering, taking advantage of the local support property. Inspired by ptychography techniques [Rodenburg et al. 2007], our algorithm optimizes for the auto-correlation of overlapping local windows, instead of the full-frame auto-correlation of the entire sensor as in previous algorithms. Our algorithm can be used to improve imaging-through-scattering performance in both the near-field and far-field settings, so long as they operate in the mid-order scattering regime. We demonstrate this improved performance through experiments we perform using both near-field and far-field imaging prototypes. Our experiments show that, compared with previous auto-correlation approaches [Katz et al. 2014; Chang and Wetzstein 2018], our algorithm results in an order-of-magnitude expansion of both the range and density of independent illuminators that can be recovered.

Implications and future outlook. Together, our theory, simulations, and experiments shed light on a fundamental limit on the performance of memory effect algorithms for imaging through scattering: namely, their ability to recover a clean image of obscured incoherent illuminators deteriorates as the *density* of the illuminators increases. This is due to the fact that speckle contrast decays quickly when summing speckle patterns from multiple illuminators, as shown, e.g., when comparing the sparse and dense input images in Figure 1. The density of illuminators is a fundamental limit that has not previously received much attention in the literature and is distinct from the better-studied fundamental limit imposed by the memory effect's finite range. That is, recovering a clean image can be unsuccessful in the presence of a large number of independent illuminators, even if all of them are within the memory effect's range.

Our article additionally comprehensively catalogues similarities and differences between the far-field and near-field variants of the imaging-through-scattering problem. In particular, our article demonstrates that the near-field variant of the problem is harder in

two ways. First, ME correlation applies to much shorter displacements. Second, as we analyze in this article, due to restrictions on illuminators density, the latent patterns that can be recovered in the near-field setting are sparser and a lot more constrained in terms of their spatial layout than in the far-field setting. Despite these difficulties, our experimental results demonstrate that our algorithm can reconstruct latent near-field patterns of considerable size and density, both an order-of-magnitude larger than what was possible using previous algorithms. These latent patterns already have potential for applications in medical imaging, e.g., sparse blinking fluorescent sources used for STORM localization [Betzig et al. 2006] or sparse cell nuclei observed through microscopes. At the same time, the challenges and limitations analyzed in our article point towards the development of fully robust near-field memory effect algorithms as an important future research direction.

2 RELATED WORK

Imaging with speckle correlations. Several imaging techniques leverage *spatial* speckle correlations, often termed the *memory effect*. Example applications include motion tracking [Jacquot and RasTrans. Graphi 1979; Jakobsen et al. 2012; Jo et al. 2015; Smith et al. 2017], looking around the corner [Smith et al. 2018; Freund 1990; Katz et al. 2012; Batarseh et al. 2018; Metzler et al. 2020], super-resolution [Dertinger et al. 2009; Chaigne et al. 2016, 2017; Judkewitz et al. 2013], and seeing through [Katz et al. 2014; Bertolotti et al. 2012; Boniface et al. 2020] or focusing through [Mosk et al. 2013; Nixon et al. 2013; Horstmeyer et al. 2015; Osnabrugge et al. 2017; Vellekoop and Aegeerter 2010; Papadopoulos et al. 2016] tissue and other scattering layers.

Other techniques use *temporal* correlations due to scatterer motion, e.g., in liquid dispersions [Dougherty et al. 1994; Berne and Pecora 2000]. Temporal correlations provide information about liquid flow (e.g., blood flow [Durduran et al. 2010]) and composition. Example techniques include diffusing wave spectroscopy [Pine et al. 1988], laser speckle contrast imaging [Boas and Yodh 1997], and dynamic light scattering [Goldburg 1999].

Imaging through scattering. Other techniques for imaging through scattering with coherent illumination use adaptive optics to focus at specific points inside the scattering sample [Rueckel et al. 2006; Vellekoop and Mosk 2007; Yaqoob et al. 2008; Katz et al. 2010; Vellekoop et al. 2010; van Putten et al. 2011; Choi et al. 2011; Katz et al. 2012; Vellekoop et al. 2012; Lai et al. 2015; Horstmeyer et al. 2015; Boniface et al. 2019, 2020]. The main challenge for these techniques is the non-invasive recovery of the aberration correction pattern that the adaptive optics need to apply to achieve focusing. The memory effect can help alleviate this challenge by allowing to adapt a previously recovered pattern to focus at different nearby points [Osnabrugge et al. 2017].

Other approaches for imaging through scattering use incoherent illumination and rely on incoherent intensity models for scattering [Durduran et al. 2010]. Many of these techniques take advantage of additional information available in time-resolved measurements, captured using so-called *transient imaging* systems [Satat et al. 2015, 2016, 2017; Xin et al. 2019]. Noteworthy within this category are *diffuse optical tomography* techniques [Boas et al. 2001; Liu et al. 2020], which use diffusion

theory to achieve larger depth penetration, at the cost of reduced resolution compared to coherent techniques.

Improving speckle correlation algorithms. Speckle auto-correlation algorithms for imaging through scattering have recently received increased attention, with several works focusing on improving depth penetration, angular extent, and overall robustness. For example, Wang et al. [2019] proposed to decompose the auto-correlation as a superposition of multiple local auto-correlations, resulting in a threefold improvement in angular extent. Li et al. [2018a] extract spatially varying point spread functions from the speckle correlation using a sequence of illumination patterns. Other techniques [Liao et al. 2019; Chang and Wetzstein 2018] improve robustness by adding sparsity priors on the latent image. Complementary to these techniques are works [Li et al. 2018b; Guo et al. 2020] that use learning-based approaches to allow recovering illuminator patterns wider than the memory effect range. However, these come at the cost of reduced generality—only patterns similar to those available in constrained training datasets (e.g., handwritten digits) can be recovered.

Finally, related to our work are techniques that use ptychography [Rodenburg et al. 2007] to increase the memory effect range [Zhou et al. 2020; Gardner et al. 2019; Li et al. 2019a, 2019b]. These techniques take *multiple images* as input, each corresponding to illuminating a different area on the scattering sample. By contrast, our algorithm works using just a *single image* as input.

3 PROBLEM SETTING AND BACKGROUND

In this section, we formalize the imaging through scattering problem and clarify the distinction between near-field and far-field imaging conditions. We additionally provide background on speckle statistics, the memory effect, and its use for imaging through scattering.

Imaging geometry. We consider the setup in Figure 2(a). Without loss of generality, we assume that the optical axis of the system is aligned with the z axis. A scattering sample (e.g., tissue layer) of thickness L is positioned between depth planes $z_{\min} = -L/2$, and $z_{\max} = L/2$. We assume that the scattering sample has a width in the x, y dimensions that is much larger than the depth L .

The sample is illuminated by multiple co-planar sources located at depth z_i . The light propagates through the scattering sample and generates a speckle pattern, measured by a 2D sensor at depth z_v . This can be either a lensless sensor physically located at z_v (Figure 2(a)), or an imaging system focused at z_v (Figure 2(b)). We denote by \mathbf{i}, \mathbf{v} the 3D position of illumination and viewing points, and by $\mathbf{i}_{x,y}, \mathbf{v}_{x,y}$ their $x-y$ restriction to the $z=z_i$ and $z=z_v$ planes, respectively.

We restrict our discussion to the *transmissive* setting, where the illuminators and the sensor (or imaging lens) are placed at opposite sides of the scattering sample. Within this setting, selecting $z_i = z_{\min}$ corresponds to cases where the illuminators are immediately at the back of the sample (e.g., isotropic fluorescent sources inside the sample or confocal illumination focused at that depth). We refer to such cases as *near-field* configurations (Figure 2(c)). By contrast, selecting $z_i \ll z_{\min}$ corresponds to cases where the illuminators are placed at a large distance from the sample, as is

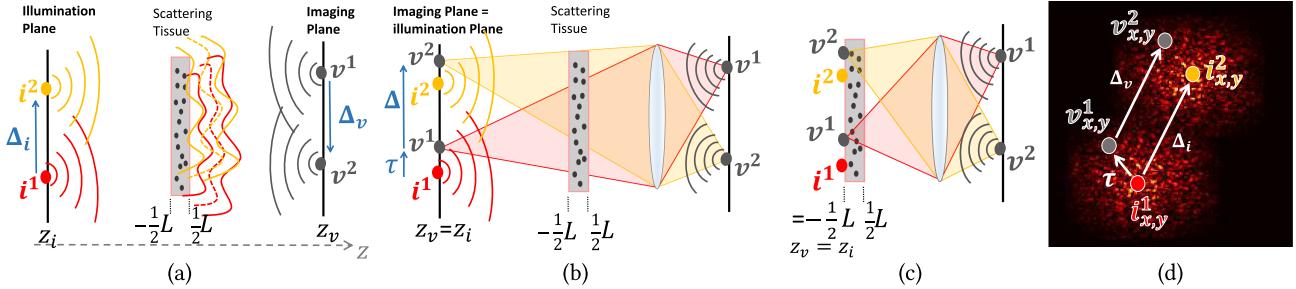


Fig. 2. **Notation for memory effect parameters.** (a) A sample is illuminated by sources at distance z_i behind it. Light propagates through the sample to generate a speckle pattern on a sensor plane at depth z_v . (b) The same scene is imaged using a lens focused at the illuminator plane so $z_v = z_i$. (c) In near-field configurations, the light is located inside (or at the back face of) the sample rather than far behind it. (d) A typical speckle image obtained on the sensor, as the superposition of scattering from two sources. We mark the illumination and viewing points and the displacements Δ_i , Δ_v , τ between them.

common in prior experimental realizations of imaging through scattering using the memory effect [Katz et al. 2014; Bertolotti et al. 2012; Edrei and Scarelli 2016a; Hofer et al. 2018; Li et al. 2018b]. We refer to such cases as *far-field* configurations. At the extreme case, the illuminators can be located at (negative) infinity, which corresponds to illuminating the scene with directional plane waves. When not clear from context, we denote these directional far-field sources and sensors using vectors with a circumflex, $\hat{\mathbf{i}}$, $\hat{\mathbf{v}}$, corresponding to their (unit-norm) directions.

Speckle statistics. We denote by $u^i(\mathbf{v})$ the complex *speckle field* generated when light from source i propagates through the scattering sample and is observed at viewing point \mathbf{v} . We denote by $I^i(\mathbf{v}) = |u^i(\mathbf{v})|^2$ the corresponding measured intensity.

Consider a scattering sample illuminated by two mutually coherent sources at $\mathbf{i}^1, \mathbf{i}^2$, and measured at two sensor positions $\mathbf{v}^1, \mathbf{v}^2$. Then, we define the *speckle covariance* as:

$$C_f(u^i(\mathbf{v}^1), u^i(\mathbf{v}^2)) \equiv E[u^i(\mathbf{v}^1)u^i(\mathbf{v}^2)^*] - E[u^i(\mathbf{v}^1)]E[u^i(\mathbf{v}^2)^*], \quad (1)$$

where $*$ denotes complex conjugation, and expectation is taken with respect to multiple realizations of random media with the same statistical properties (e.g., multiple tissue layers of the same type and thickness). Similarly, we can define the *intensity covariance*¹:

$$C_I(I^i(\mathbf{v}^1), I^i(\mathbf{v}^2)) \equiv E[I^i(\mathbf{v}^1)I^i(\mathbf{v}^2)] - E[I^i(\mathbf{v}^1)]E[I^i(\mathbf{v}^2)]. \quad (2)$$

Using classical statistics, it is easy to show that for zero mean fields:

$$C_I(I^i(\mathbf{v}^1), I^i(\mathbf{v}^2)) = |C_f(u^i(\mathbf{v}^1), u^i(\mathbf{v}^2))|^2. \quad (3)$$

We can now use these quantities to formally describe the *memory effect* (ME) property of speckle fields. We consider the speckle fields u^i , u^j generated by two *nearby* illuminators $\mathbf{i}^1, \mathbf{i}^2$, displaced relative to each other by a vector $\Delta_i \equiv \mathbf{i}_{x,y}^2 - \mathbf{i}_{x,y}^1$. The ME refers to the fact that u^i, u^j will be correlated shifted versions of each other. That is, there exists a displacement vector Δ_v in the *view plane* such that

$$u^i(\mathbf{v}) \approx u^j(\mathbf{v} + \Delta_v). \quad (4)$$

¹Intensity correlations are independent of the phase of the signal, and hence they do not require the sources to be mutually coherent.

In Section 4, we prove that for any given illuminator displacement Δ_i , we can compute an optimal view-plane displacement $\Delta_v^{\text{opt}}(\Delta_i)$ that maximizes the above correlation.

We show examples of the ME property in Figure 5: Speckle patterns are similar when generated by nearby illuminators, but become different as the illuminator displacement increases. This points to the fact that the ME property holds only for small displacements $|\Delta_i|$ between light sources, with the correlation decreasing as the displacement increases. To quantify how the ME correlation decays as a function of displacement, it is common to measure the sum of speckle correlations at all sensor pixels. For applications considering field correlations (e.g., adaptive optics for focusing through scattering [Judkewitz et al. 2014; Osnabrugge et al. 2017; Horstmeyer et al. 2015; Papadopoulos et al. 2016]), this corresponds to:

$$C_f(\Delta_i) = \sum_{\mathbf{v}_{x,y}} e^{ik\theta^{\text{opt}}(\Delta_i)\mathbf{v}_{x,y}} C_f(u^1(\mathbf{v}_{x,y}), u^2(\mathbf{v}_{x,y} + \Delta_v^{\text{opt}}(\Delta_i))), \quad (5)$$

where $e^{ik\theta^{\text{opt}}(\Delta_i)\mathbf{v}_{x,y}}$ is a phase correction, which is required to make $C_f(\Delta_i)$ meaningful, since the correlation at individual pixels are complex with possibly different phases. We derive the optimal frequency $\theta^{\text{opt}}(\Delta_i)$ in Appendix A.1. For applications considering intensity correlations, such as the imaging-through-scattering task we study in this article, the ME is instead quantified using:

$$C_I(\Delta_i) = \sum_{\mathbf{v}_{x,y}} C_I(I^1(\mathbf{v}_{x,y}), I^2(\mathbf{v}_{x,y} + \Delta_v^{\text{opt}}(\Delta_i))). \quad (6)$$

In both Equations (5) and (6), the correlation is evaluated with the optimal view-plane displacement corresponding to the illuminator displacement. We assume that, for a wide homogeneous sample, this correlation depends only on the displacement Δ_i , rather than on the exact spatial positions $\mathbf{i}_{x,y}^1, \mathbf{i}_{x,y}^2$ of the illumination sources. We will call the *ME range* the maximum displacement Δ_i for which the correlation $C_I(\Delta_i)$ remains significant.

Using ME to image through scattering. We now briefly review the method of Katz et al. [2014] for imaging through scattering using the ME. We will be using this method as our baseline throughout the article. Consider a speckle image I generated when a scattering sample is illuminated simultaneously by K *mutually incoherent* sources \mathbf{i}_k on one side and imaged by a camera on the other side. The camera measures the incoherent summation of speckle

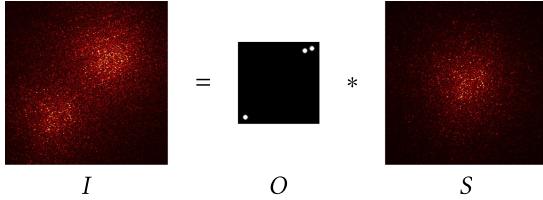


Fig. 3. **Image formation model within the memory effect** [Katz et al. 2014]. S is the speckle pattern from a single illuminator, O is a binary image of illuminator locations, and I is the image measured by the camera.

intensities,

$$I(v) = \sum_{k=1}^K I^k(v), \quad (7)$$

where $I^k(v)$ denotes the intensity image from the k th source.² We denote by $S \equiv I^{(0,0)}(v_{x,y})$ the speckle image generated by a source $i_{x,y} = (0, 0)$. If all K sources are within the ME range, then all $I^k(v)$ images are shifted versions of S . Denoting by O a binary *latent* image corresponding to the locations of all K sources (i.e., the image the camera would capture in the absence of the scattering sample), we can write $I = O \star S$, where \star denotes correlation, as visualized in Figure 3. We note that, if we replace I and S by

$$\bar{I} \equiv I - \mu(I), \quad \bar{S} \equiv S - \mu(S), \quad (8)$$

$\mu(I), \mu(S)$ being signal means, then the same correlation relation holds.

The image \bar{I} typically exhibits noise-like speckle, making it difficult to discern O . We can, however, consider its auto-correlation,

$$\bar{I} \star \bar{I} = (O \star \bar{S}) \star (O \star \bar{S}) = (O \star O) \star (\bar{S} \star \bar{S}). \quad (9)$$

As the intensity values in \bar{S} are approximately zero-mean independent noise, its auto-correlation is approximately an impulse, $\bar{S} \star \bar{S} \approx \delta$. Thus, the auto-correlation of \bar{I} is approximately equal to the auto-correlation of O ,

$$\bar{I} \star \bar{I} \approx O \star O. \quad (10)$$

Therefore, we can recover O from $\bar{I} \star \bar{I}$ using phase retrieval algorithms, e.g., the classical algorithm by Fienup [1982] or more robust strategies [Li et al. 2018b; Guo et al. 2020]. We refer to this procedure as the *full-frame auto-correlation algorithm* in the rest of the article.

A serious shortcoming of the full-frame auto-correlation algorithm is that the range of illuminators it can recover is small, as the maximal displacement between the illuminators $|i_{k_1} - i_{k_2}|$ needs to be within the ME range, which is typically very small. Additionally, our article highlights another shortcoming that has received less attention in the literature, namely, that recovery is only possible when the number of illuminators K contributing to Equation (7) is sufficiently small. Our goal in this article is to quantify and compare these constraints in the far-field and near-field settings; as well as to propose a new algorithm for imaging through scattering that can significantly relax these constraints.

²As the sources are incoherent, the camera measures $\sum_{k=1}^K I^k(v)$ rather than $|\sum_{k=1}^K u^k(v)|^2$.

4 COMPARING NEAR-FIELD AND FAR-FIELD SETTINGS

We start by exploring some properties of speckle statistics, with a focus towards understanding the relationship between speckle statistics in the near-field and far-field settings. In particular, we explore the effect of the parameters z_i and z_v , shown in Figure 2, on speckle correlations. The depth z_i of the illuminators is relative to the scattering sample controls whether we are in the near-field and far-field settings, and thus investigating how speckle correlations vary as a function of this parameter can help us understand the differences between the two settings. The depth z_v of the viewing plane corresponds to experimental choices such as deciding whether to measure speckle using a bare sensor on the front face of the sample, versus using a lens to focus the sensor at a different plane. We will use the observations we make in this section to develop better imaging-through-scattering algorithms in subsequent sections.

4.1 Analytic Field Correlation Relationship

In this section, we focus on correlation of complex fields (Equation (5)); we will discuss the correlation of intensity images later. We derive a new technical result that allows converting between field correlations in the near-field and far-field settings.

CLAIM 1. *Let $C_f(\Delta_i|z_i^j, z_v^j)$ denote the field correlation in Equation (5), with illuminators and sensors placed at planes z_i^j, z_v^j , respectively. Then, correlations measured at different illuminator placements can be related through a displacement scaling as:*

$$C_f(\Delta_i|z_i^1, z_v^1) = C_f\left(\frac{z_i^1}{z_i^2}\Delta_i|z_i^2, z_v^2\right), \quad (11)$$

where in both cases correlation is evaluated at the optimal view-plane displacement given by

$$\Delta_{\mathbf{v}}^{opt}(\Delta_i) = \frac{z_v}{z_i}\Delta_i \quad \text{or equivalently} \quad \frac{\Delta_{\mathbf{v}}^{opt}(\Delta_i)}{z_v} = \frac{\Delta_i}{z_i}. \quad (12)$$

We provide the proof in Appendix A.1 relying on ideas in Bar et al. [2021]. We will be using

$$\hat{\Delta}_i \equiv \frac{\Delta_i}{z_i}, \quad \hat{\Delta}_{\mathbf{v}} \equiv \frac{\Delta_{\mathbf{v}}}{z_v} \quad (13)$$

to denote normalized displacements. For small angles, the normalized displacement $\hat{\Delta}_i$ is equal to a first-order approximation to the angle two illuminators displaced by Δ_i at depth z_i form with their midpoint on the plane $z = 0$ in the middle of the sample; and similarly for $\hat{\Delta}_{\mathbf{v}}$. We visualize these angles in Figure 4(a). Thus, we refer to $\hat{\Delta}_i, \hat{\Delta}_{\mathbf{v}}$ as the *angular displacements*.

Claim 1 has three important implications: First, it states that field correlation is a function of only the *angular* displacement between the illuminators, and not their actual distance. Second, it states that field correlation is invariant to the viewing plane z_v , as long as the view-plane displacement is scaled as in Equation (12) to maintain a fixed angle with the scattering sample. Third, it states that we can convert field correlations between the far-field (large z_i values) and near-field (small z_i values) settings. This may be useful for translating knowledge about the far-field setting into the near-field setting.

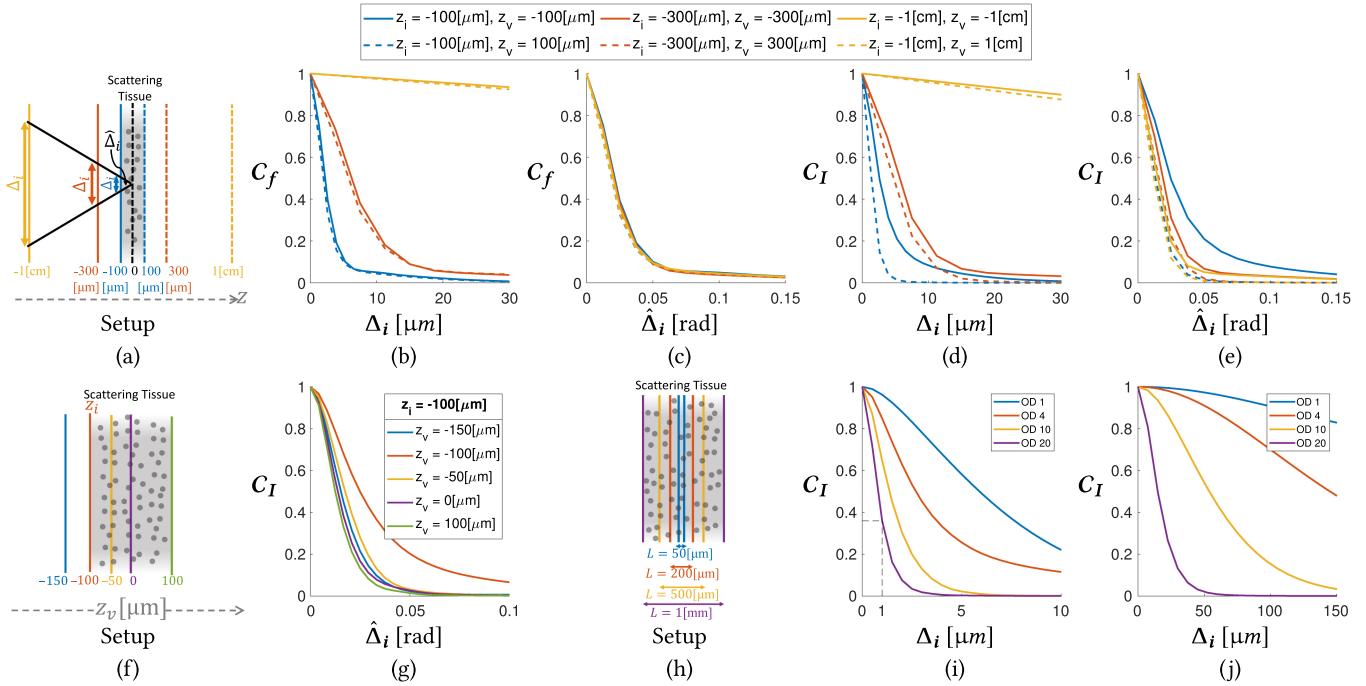


Fig. 4. Dependence of correlation on angular displacement. (a) Setup for simulation in (b)–(e), highlighting the angular displacement $\hat{\Delta}_i$ corresponding to spatial displacements Δ_i at different illumination planes. Plots (b)–(e) evaluate ME decay for three distances of the illumination source, and two focus settings. (b) Field correlation, plotted as a function of absolute displacement Δ_i . Far sources have much wider ME extant. (c) When plotted as a function of angular displacement $\hat{\Delta}_i$ rather than actual spatial displacements. That is, field correlation is a function of angular displacement $\hat{\Delta}_i$ rather than actual spatial displacements. (d) Intensity correlations as a function of spatial displacement are also much wider for far sources. (e) Intensity correlations as a function of angular displacement are not fully invariant to source distance. (f)–(g) Varying the view plane z_v for a fixed illumination plane z_i shows that the widest intensity correlation is obtained when the sensor focuses at the illuminator plane. (h) Setup for (i)–(j). (i) Near field intensity correlation as a function of absolute displacement, evaluated for a few material thickness. As the optical depth grows, ME correlation shrinks. For thick slices, the largest displacement at which correlation holds can be too small for any practical usage. (j) Same as (i) but for far sources. In this case, correlation at non zero displacements can be found even for thicker materials.

Claim 1 applies to the correlations of *complex fields*. However, imaging-through-scattering algorithms use correlations of speckle *intensities*. Unfortunately, we have not found a similar closed-form relationship for intensity correlations. However, the intuition that such correlations depend mostly on angular displacements still holds. In the following, we explore intensity correlations using numerical simulations as well as real tissue measurements.

4.2 Simulation-based Exploration

In this section, we use the physically accurate speckle rendering algorithms of Bar et al. (2019, 2020) to perform simulated experiments with two objectives in mind: First, we want to validate the predictions of Claim 1 for field correlations. Second, we want to explore properties of intensity correlations. In our simulations, we use scattering material parameters commonly used in medical imaging to model tissue [Igarashi et al. 2007]. In particular, we simulate a scattering sample with a Henyey-Greenstein phase function of anisotropy parameter $g = 0.99$, and mean free path $MFP = 50\mu\text{m}$ at wavelength $\lambda = 0.5\mu\text{m}$. Except where noted otherwise, we set sample thickness to $L = 200\mu\text{m}$, resulting in an optical depth $OD = 4$. We consider three illumination plane settings: (i) $z_i = z_{\min} = -100\mu\text{m}$, a near-field configuration where

illuminators are placed at the back face of the sample; (ii) $z_i = -300\mu\text{m}$, representing a small gap between the illuminator plane and the sample; (iii) $z_i = -1\text{cm}$, which is large enough to correspond to a far-field configuration. The second case is representative of experimental near-field realizations where fluorescent patterns are not really attached to the sample, but rather are projected onto it using a—potentially misfocused—relay lens [Chang and Wetzstein 2018]; or where fluorescent particles need to remain separated from the sample by a thin cover glass. These situations motivate our experiments to understand the effect of the resulting gap on speckle correlations. For each z_i setting, we consider two viewing plane settings: (i) $z_v = z_i$, corresponding to a sensor that is focused, through an imaging lens, at the illuminator plane; and (ii) $z_v = -z_i$, corresponding to a sensor placed at some distance from the front face of the sample. This case is representative of how far-field ME is often measured in practice, by placing a bare sensor far enough from the sample.

Dependence of correlation on angular displacement. In Figure 4(b)–(c), we simulate field correlation values C_f as a function of either absolute displacement Δ_i , or angular displacement $\hat{\Delta}_i$, for the six different configurations of illuminator and view planes z_i, z_v described above. When parameterizing C_f

by Δ_i , increasing the distance between the sample and illuminator plane increases the ME range. This is due to the fact that increasing this distance while keeping Δ_i constant reduces the angular displacement of the sources. As suggested by Claim 1, parameterizing C_f by $\hat{\Delta}_i$ makes all six configurations identical, showing that field correlations are invariant to the illuminator and view plane locations z_i, z_v .

In Figure 4(d)–(e), we repeat the above experiments, but this time simulating *intensity* correlation values C_I , which are the actual input to the imaging-through-scattering algorithms we develop later. As in the field case, when we parameterize C_I by absolute displacement Δ_i , the correlation significantly increases as the distance of the illuminator plane z_i increases. Parameterizing C_I as a function of angular displacement $\hat{\Delta}_i$ brings the six different configurations closer to each other, suggesting that, similarly to field correlation, intensity correlation also depends more strongly on angular displacement $\hat{\Delta}_i$ than on absolute displacement Δ_i . However, in contrast to the field correlation case, the six configurations are not equivalent, indicating that intensity correlations are not completely independent of absolute displacements. We explain this difference between field and intensity correlations in Appendix A.1; recall in particular that from Equations (5) and (6), $C_I(\hat{\Delta}_i)$ is *not* the squared amplitude of $C_f(\hat{\Delta}_i)$.

Dependence of intensity correlation on view plane. Figure 4(g) demonstrates that, unlike field correlations C_f , intensity correlations C_I are sensitive to the selection of the viewing plane z_v . In particular, we note that intensity correlation is maximized when the viewing plane coincides with the illuminator plane, $z_v = z_i$. This is an important practical observation. Often, such experiments use a bare sensor placed on the front face of the sample to measure speckle correlations. Our results suggest that we can increase the measured ME through a simple change in the imaging setup, namely, using a lens positioned in such a way that if there were no scattering sample, the sensor would image the illuminator plane z_i . We refer to this imaging configuration as the *focused configuration*. We validate this observation below using real tissue measurements.

ME ranges in practical near-field and far-field configurations. As ME range depends mostly on the angular displacement $\hat{\Delta}_i$, the size of the latent pattern one can handle using ME techniques increases when this pattern is placed further away from the scattering sample. This observation explains why imaging-through-scattering experiments are easier to perform in the far-field than in the near-field. In particular, even though the intensity correlation decay as a function of angular displacement is similar for near-field and far-field configurations, in the near-field case, the corresponding maximal absolute displacement Δ_i can become smaller than the wavelength.

To demonstrate this, in Figure 4(i), we simulate intensity correlation values in a near-field configuration, for scattering samples of progressively larger depths, with all other sample parameters remaining the same. This corresponds to increasing the optical depth of the sample, and thus increasing the average number of light scattering events; in turn, this results in a faster decay of intensity correlation as a function of absolute displacement Δ_i . The evaluation shows indeed that, for thick samples, near-field

ME vanishes for any realistic displacement. By contrast, when repeating the same simulations for a far-field configuration, as in Figure 4(j), we observe that the ME range remains non-negligible even for thick scattering samples. By moving the illuminator plane further away from the sample, we can scale the ME range to cover latent patterns of any size.

These simulations suggest that ME techniques in the near-field setting are only applicable for scattering samples of modest thickness, where mid-order scattering is dominant. Consequently, in the next section, we focus on exploring properties of speckle patterns formed under these conditions that can facilitate the development of imaging-through-scattering algorithms. We note here that the exact sample depth at which near-field ME is non-negligible will vary for different types of scattering materials. In particular, there are significant variations in the material parameters reported as representative of tissue in the literature [Cheong et al. 1990; Tuchin 2000; Igarashi et al. 2007]. We used one set of such parameters for our simulations, but the exact correlation values will be different for other tissue parameters. Therefore, our simulations, and in particular the near-field correlation plots in Figure 4(i), are not intended to precisely predict a tissue depth at which near-field ME vanishes, but rather to support our observation that near-field ME is only non-negligible under modest thicknesses corresponding to mid-order scattering. In practice, for chicken breast tissue, we detect near-field correlations for samples up to 200 μm thick. Despite the modest thickness of the samples, images captured through them still contain considerable degradation that can benefit from speckle correlation techniques.

4.3 Qualitative Validation Using Real Measurements

Before concluding this section, we present results from real measurements of chicken breast tissue. Data was captured using a near-field experimental imaging setup, described in Section 8.1. These results lend support to the observations presented earlier in this section.

Empirical correlation decay. In Figure 5, we use near field speckle images to demonstrate how intensity correlation C_I decays as a function of absolute displacement Δ_i . We compute $C_I(\Delta_i)$ by empirically correlating the captured speckle images. We perform measurements for different placements of the view and illuminator planes. We first compare the intensity correlation plots $C_I(\Delta_i)$ measured when the sensor is focused at the illuminator plane, versus when it is focused at a different plane. The ME range is wider in the former case, as predicted by our simulations. We then compare the intensity correlations $C_I(\Delta_i)$ measured for two placements of the illuminator plane: one where it is exactly at the back plane of the sample and another where it is at a distance of 200 μm from the back plane. In both of these cases, the sensor is focused at the illuminator plane. In agreement with our earlier observations, our measurements show that placing the illuminators at a distance from the sample results in an increased ME range, when measured as a function of absolute displacement Δ_i rather than angular displacement $\hat{\Delta}_i$. That is, even a small distance of 200 μm , which can occur in experiments due to misfocusing, can significantly impact the ME range.

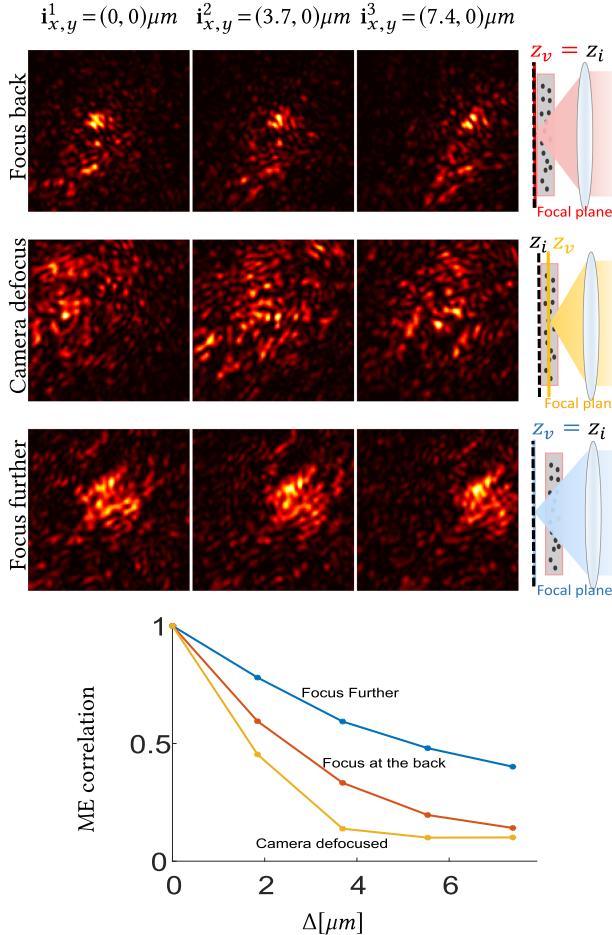


Fig. 5. Near-field ME. We present speckle images captured when placing an illuminator at three nearby positions; speckle patterns are shifted versions of each other, demonstrating the ME. The speckle spread is smallest when the illumination source is located exactly at the back plane of the sample and the objective distance is set to focus on that plane. Focusing the objective on a closer plane (2nd row) results in wider speckles. Computing ME correlation empirically from the captured speckles (lower part), we see that ME correlation holds for larger displacements when the camera is properly focused. We also test the option of moving the light source $200\mu\text{m}$ further than the back layer while correctly focusing on the illumination plane. In this configuration, higher correlation is measured at wider displacements.

Summary of important observations. To conclude this section, we summarize three important observations we presented: (i) We showed that, in near-field settings, the ME range is non-negligible only for scattering samples of modest optical depth, with a thickness of only a few mean free paths. Such layers are dominated by mid-order scattering. (ii) We showed that it is important to use a focused configuration where the view plane coincides with the illuminator plane to maximize the ME ranges. (iii) We showed that care needs to be taken when placing the illuminators behind the scattering sample, as even a small gap between the two can artificially increase the ME range, even if such a gap is unrealistic for applications such as fluorescent imaging. In the rest of the article,

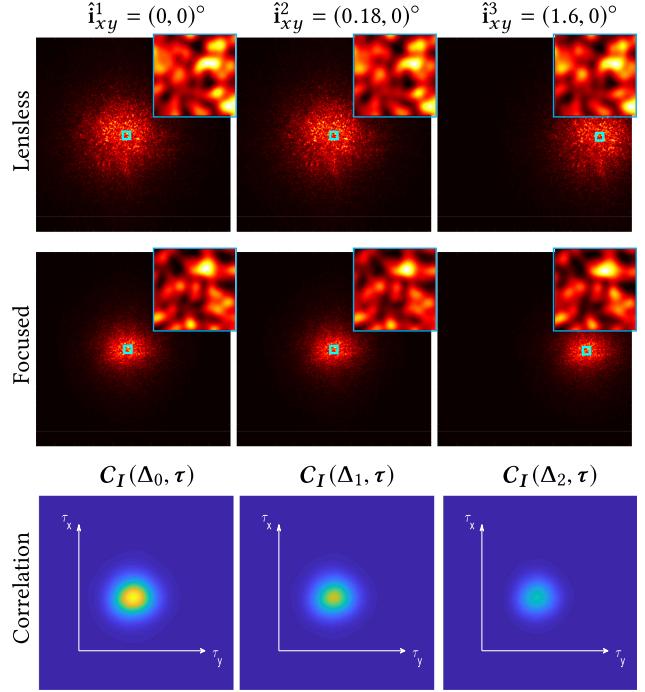


Fig. 6. Local support property. Images of far-field illumination scattering through a chicken breast slice of thickness $250\mu\text{m}$, for a focused and a bare sensor. Imaging with a focused sensor reduces the speckle support. We show speckle images for three different positions of the illumination source. The insets demonstrate the ME, namely, that the speckle patterns generated by different illuminators are shifted versions of each other. Lower panel: empirical speckle correlation $C_I(\Delta, \tau)$ (as defined in Equation (17)) evaluated from the focused data at the second panel. The correlation is displayed as a function of the 2D displacement vector τ for three different choices of illuminator displacements $\hat{\Delta}_0 = (0, 0)^\circ$, $\hat{\Delta}_1 = (0.18, 0)^\circ$, $\hat{\Delta}_2 = (1.6, 0)^\circ$. Due to the modest sample thickness, speckle spread is local, and so is the correlation.

we focus on experimental settings using samples of modest thickness, properly focused sensor, and carefully placed illuminators, as informed by the above three observations.

5 THE LOCAL SUPPORT PROPERTY

In this section, we document and characterize a property of speckle patterns formed under conditions where mid-order scattering is dominant. In particular, in Figures 5 and 6, we show speckle patterns measured through a tissue layer of modest thickness, using our near-field and far-field imaging setups. We observe that the speckle patterns have *local support*, much smaller than the full extent of the sensor. Local supports are prevalent in biological tissue samples of modest thickness, as the phase functions characterizing these samples are strongly forward-scattering—their average cosine is typically $g > 0.95$. Given such phase functions, light entering the sample with direction \hat{i} will, after undergoing a small number of scattering events, spread primarily towards outgoing directions $\hat{v} \approx \hat{i}$. As we will show in the next section, the local support property of speckle patterns due to mid-order scattering is key for improving imaging-through-scattering algorithms based

on the ME. Given its importance, we use this section to study this property in more detail.

Effect of focusing on speckle support. In Section 4, we showed that using a lens to focus the sensor at the same plane as the illuminators is important for increasing the ME range. Figures 5 and 6 demonstrate an additional advantage of using this focused configuration: In both the far-field and near-field settings, focusing decreases the speckle spread. Intuitively, the wider spread of non-focused configurations can be explained by the fact that the support of the scattered field is convolved with a defocus blur kernel.

Far-field versus near-field speckle patterns. Speckle patterns formed under far-field and near-field settings both exhibit the local support property. However, far-field patterns, such as those shown in Figure 6, include many more speckle features compared to near-field ones, such as those shown in Figure 5. We will see in subsequent sections that this difference makes imaging through scattering in the near-field setting more challenging than in the far-field setting.

Extended parameterization of intensity correlation. For the rest of this article, we restrict the discussion to focused configurations, where $z_v = z_i$. In this case, we get from Equation (12) that the displacement on the illuminator plane and the corresponding optimal displacement on the view plane are equal. Thus, we simplify notation using $\Delta \equiv \Delta_i = \Delta_v^{\text{opt}}(\Delta_i)$ for both displacements, leading to:

$$i_{x,y}^2 - i_{x,y}^1 = v_{x,y}^2 - v_{x,y}^1 = \Delta. \quad (14)$$

We note that the definition of the intensity correlation $C_I(\Delta)$ in Equation (6) treats all pixels $v_{x,y}^1$ equally and does not consider the location of the pixel $v_{x,y}^1$ relative to the illuminator location $i_{x,y}^1$. This is due to the fact that most prior literature on the ME focuses on settings where speckle patterns cover the entire sensor plane (e.g., cases where high-order scattering is dominant). Consequently, correlation does not vary significantly at different locations $v_{x,y}^1$ on the sensor, and it is sufficient to analyze how correlation C_I decays as a function of the displacement Δ alone. By contrast, in our setting, the local support property implies that speckle patterns generated by an illuminator at location $i_{x,y}^1$ are concentrated at pixels in locations $v_{x,y}^1$ adjacent to $i_{x,y}^1$. This suggests that correlation can vary at different locations $v_{x,y}^1$ on the sensor: For example, as we move away from $i_{x,y}^1$, less light is measured, and we expect correlation to be reduced. To characterize this effect of the local support property, we will modify the definition of the intensity correlation C_I so it takes as input the displacement between the illuminator and the pixel rather than only the displacement between the two sources.³

To this end, we denote the 2D displacement between the illuminator and pixel locations as:

$$\tau \equiv v_{x,y}^1 - i_{x,y}^1. \quad (15)$$

³We note that Judkewitz et al. [2014] have observed that materials with forward-scattering phase functions produce speckle patterns with local support, and also that this results in these materials having a larger ME range. In particular, they showed that the local support of the speckle patterns is related to a larger ME range, as $C_f(\Delta)$ is the Fourier transform of the angular speckle spread. However, their study still weighs all sensor pixels equally when computing correlation values.

We visualize both τ and Δ in Figure 2(d). For illuminator and pixel pairs satisfying Equation (14), it follows that:

$$v_{x,y}^1 - i_{x,y}^1 = v_{x,y}^2 - i_{x,y}^2 = \tau. \quad (16)$$

Then, we define the intensity correlation for illumination and pixel pairs satisfying both the illuminators displacement relation of Equation (14) and the illuminator-pixel displacement relation of Equation (16):

$$C_I(\Delta, \tau) \equiv C_I \left(I^{i_{x,y}^1}(i_{x,y}^1 + \tau), I^{i_{x,y}^1 + \Delta}(i_{x,y}^1 + \tau + \Delta) \right), \quad (17)$$

where the intensity covariance $C_I(I^{i^1}(v^1), I^{i^2}(v^2))$ was defined in Equation (2). We note that we can relate this definition to the definition of intensity correlation as a function of Δ alone in Equation (6) through the equation $C_I(\Delta) = \sum_{\tau} C_I(\Delta, \tau)$.

To demonstrate empirically the importance of parameterizing the intensity correlation C_I as a function of both Δ and τ , we use speckle images captured from chicken breast tissue samples with a far-field experimental imaging setup described in Section 8.1 below. In Figure 6, we show the intensity correlation $C_I(\Delta, \tau)$ computed from the image measurements. We can observe that C_I varies significantly as a function of both Δ and τ , and in particular that it quickly decays as the distance τ between the illuminator and pixel location increases. A schematic of the displacements Δ, τ is visualized in Figure 2(d). To our knowledge, the local support property and its effect on intensity correlations have not previously been used for imaging-through-scattering applications. In the next section, we provide a theoretical justification for using this property; then, in Section 7, we use it to develop an improved algorithm for imaging through scattering.

6 SIGNAL-TO-NOISE RATIO ANALYSIS

Previous studies of the full-frame speckle auto-correlation algorithm that we described in Section 3, for example by Wang et al. [2019], have focused on how the limited ME range constrains the size of the latent illuminator pattern that can be recovered. In this section, we study a second constraint on the recoverable latent illuminator pattern that has received little attention in the literature (see limited discussion in the supplement of Katz et al. [2014]): the fact that reconstruction is usually successful only when the number of different illuminators K in Equation (7) is sufficiently small. When a large number of incoherent sources contribute to the measured intensity image, speckle contrast decays and correlation becomes noisier. For example, this difference in speckle contrast is noticeable when comparing the sparse and dense inputs of Figure 1. We show that, by taking advantage of the local support property we described in the previous section, we can significantly increase the **signal-to-noise ratio (SNR)** of the correlation, and consequently, the density of illuminators we can recover. We note that the reconstruction algorithm by Katz et al. [2014] has two parts: first computing speckle correlation, and then performing phase retrieval. The focus of our analysis is on the first part, the SNR at which correlation can be computed. Even though we expect that the performance of the phase retrieval part will also improve as the noise characteristics of its input improve, a detailed analysis of phase retrieval convergence is beyond the scope of this work.

We are given a speckle image I formed as in Equation (7) and want to examine whether illuminators i^1, i^2 contributed to its formation. Denoting the illuminator displacement $\Delta = i_{x,y}^2 - i_{x,y}^1$ as in Equation (14), we can multiply the zero-mean speckle image \bar{I} with its shifted copy, then form a correlation estimate using *weighted* pixel averaging:

$$c_{\text{emp}}(i_{x,y}^1, \Delta) = \sum_{\tau} w(\Delta, \tau) \bar{I}(i_{x,y}^1 + \tau) \bar{I}(i_{x,y}^1 + \tau + \Delta). \quad (18)$$

We expect c_{emp} to have a large value when an illuminator pair $i_{x,y}^1, i_{x,y}^1 + \Delta$ exists and a value close to zero otherwise. When computing full-frame auto-correlation, as in Equation (10), the spatial weights w are uniform over the entire image I . However, if we know that the speckle patterns have local support, we can consider setting non-zero weights w only in a window around $i_{x,y}^1$, rather than in the entire image. We state a new technical result showing this can drastically improve SNR, and derive the optimal weighting strategy.

To formulate this result, we denote by P the number of sensor pixels and by F the number of speckle features in the image, where a feature refers to a diffraction-limited speckle spot. We have $F \leq P$, where a gap $F < P$ happens for two possible reasons: First, depending on the aperture, a diffraction limited feature can spread over more than a single pixel; and second, even for single pixel features, the combined speckles from all illuminators may not cover the entire sensor. Additionally, we denote by K the number of illuminators.⁴ Using these notations, we define the *density* of independent illuminators as

$$\alpha \equiv K/F. \quad (19)$$

Using the density definition, we state the following claim:

CLAIM 2. *The signal-to-noise ratio of the estimator of Equation (18) is*

$$\text{SNR} = \frac{E[c_{\text{emp}}(i_{x,y}^1, \Delta)]^2}{\text{Var}[c_{\text{emp}}(i_{x,y}^1, \Delta)]} = \frac{|\sum_{\tau} w(\Delta, \tau) C_I(\Delta, \tau)|^2}{\alpha^2 \cdot \sum_{\tau} |w(\Delta, \tau)|^2 \cdot |\sum_{\tau} C_I(0, \tau)|^2}, \quad (20)$$

and is maximized by the matched filter $w(\Delta, \tau) \equiv C_I(\Delta, \tau)$, reaching

$$\text{SNR}_{\text{matched}} = \frac{\sum_{\tau} |C_I(\Delta, \tau)|^2}{\alpha^2 \cdot |\sum_{\tau} C_I(0, \tau)|^2}. \quad (21)$$

We provide the proof in Appendix A.2. Algorithmically, using the matched filter requires averaging only within the local image window where we expect to have speckle from illuminator i^1 , and not within the entire sensor as in the full-frame auto-correlation algorithm. We provide algorithmic details in Section 7.

Implications. Claim 2 suggests that using the matched filter instead of uniform summation over the image can significantly improve SNR. To qualitatively characterize this improvement, we assume for simplicity that the size of a speckle feature is one pixel, the support of the speckle pattern due to one illuminator is N pixels, and all sensor pixels receive light from at least one illuminator, so $F = P$. Suppose also that $C_I(\Delta, \tau) = 1$ inside the support and 0 otherwise. From Claim 2, the matched and uniform filters achieve SNRs of

$$\text{SNR}_{\text{matched}} = 1/(\alpha^2 N), \quad \text{SNR}_{\text{uniform}} = 1/(\alpha^2 P). \quad (22)$$

⁴In the supplement of Katz et al. [2014], density is defined as the area of high emission in the target, divided by the area of a diffraction limited spot.

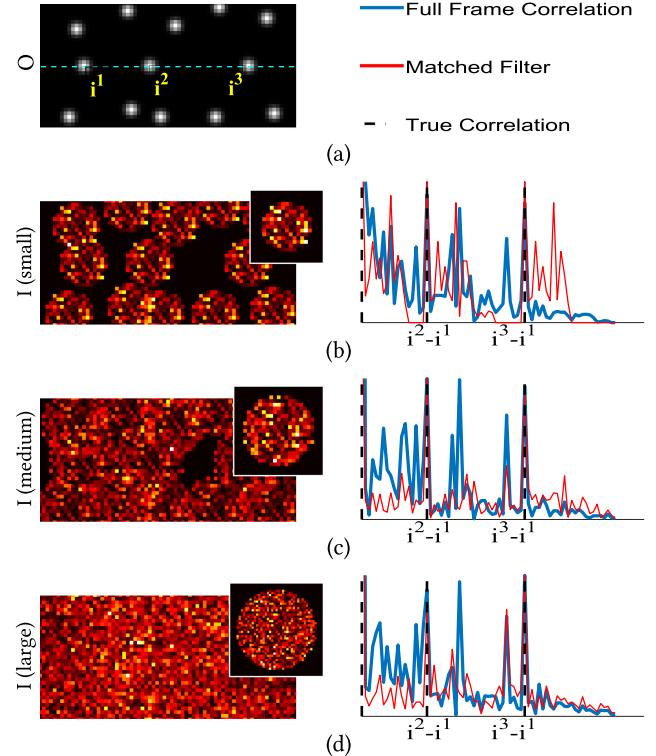


Fig. 7. SNR gain. We visualize the correlation $c_{\text{emp}}(i_{x,y}^1, \Delta)$ over a 1D line (highlighted on O). As this line includes three different illuminators, we expect to detect high correlations at three displacements. Correlations resulting from the matched filter (red curves) are less noisy than correlations from full-frame averaging (blue curves). We simulated observed images I due to the same illuminator arrangement O and three different speckle support sizes, visualized as insets at the top right corner of the corresponding observed images. As predicted by our theory, SNR improves for medium speckle support size, but decays for very small and very large support sizes.

Therefore, using the matched filter versus full-frame averaging improves SNR by N/P . When the sensor size is a few megapixels and the speckle support of each illuminator is only 100×100 pixels, this translates into an SNR improvement of two orders of magnitude.

We can also use Claim 2 to understand what illuminator density α we can expect to reliably detect. Suppose that for good detection we seek an SNR larger than a threshold R . As before, assume speckle features are single-pixel wide and the full sensor is covered by speckles. The matched and uniform filters lead to different upper bounds on the recoverable density:

$$\alpha_{\text{matched}} < \frac{1}{\sqrt{NR}}, \quad \alpha_{\text{uniform}} < \frac{1}{\sqrt{PR}}. \quad (23)$$

Selecting, e.g., $R = 100$ as a threshold for reliable detection, and given $P = F = 10^6$, we find that a uniform filter can reliably detect only one illuminator per 10^4 pixels. By contrast, if the speckle support from one illuminator includes $N = 10^4$ pixels (e.g., a 100×100 support), then the matched filter can reliably detect one

illuminator per 10^3 pixels; for $N=10^2$, this becomes one illuminator per 10^2 pixels.

We note that the above limits on SNR as a function of illuminator density hold even if all illuminators are within the ME range. Therefore, the constrained density is a fundamental limitation of the full-frame auto-correlation method for which there is limited discussion in the literature.

Visualizing the SNR gain. In Figure 7, we use a synthetic example to visualize the SNR gain achieved using the matched filter. We generate speckle images using the idealized formula $I = S \star O$, so all pixels are inside the ME range. We use three speckle patterns S of different support for the same latent image O . We then compute $c_{\text{emp}}(i_{x,y}^1, \Delta)$ as in Equation (18), for the $i_{x,y}^1$ point marked in Figure 7(a). For simplicity, we vary Δ only over one horizontal line marked in the figure. As the line contains only three illuminators, we ideally expect high correlation only for three translation values. We observe that, in agreement with Claim 2, uniform averaging produces significantly noisier correlations than the matched filter (compare the red and blue curves in Figure 7). In Figure 7(c), where we use a medium support N , the matched filter produces sharp correlation peaks at the correct displacements. In Figure 7(d), we increase the support N , and the matched-filter correlation becomes noisier; this agrees with Equation (22), which states that $\text{SNR}_{\text{matched}}$ decays as the support N increases.

We now consider Figure 7(b), where we use a small support N , and some sensor pixels do not receive light. The correlation becomes worse than that obtained with the medium support in Figure 7(c). To understand this, we note that when transitioning from the medium support in Figure 7(c) to the large support in Figure 7(d), the illuminator density α as defined in Equation (19) remains the same; as in both cases, the entire sensor is covered by speckle features so the feature number F remains the same. By contrast, in the case of small support in Figure 7(b), the density α is *higher*; this is because the number of sensor pixels covered by speckle, and thus the number of speckle features F , are both reduced. Consequently, from Equation (22), the SNR is also reduced. This example is important for understanding the near-field setting, where typically speckle support sizes are small and the speckles from all illuminators do not cover the entire sensor.

7 OPTIMIZING USING LOCAL SUPPORT

The previous section provides theoretical justification for using the matched filter, rather than uniform weights. In this section, we develop a new algorithm for imaging through scattering that explicitly takes into account the local support.

We begin by noting that, in many practical cases, we cannot measure the exact correlation $C_I(\Delta, \tau)$, and thus cannot compute the exact matched filter of Equation (21). Instead, our algorithm will approximate it using two binary thresholds T_τ, T_Δ , assuming that speckles from one illuminator are spread over pixels in a window of size T_τ around it, and that ME correlation holds for displacements $|\Delta| < T_\Delta$. The thresholds T_τ, T_Δ are free parameters that we can fine-tune to improve reconstruction quality. We show in Appendix A.5 that performance is robust to their exact values. As we discuss below, our algorithm offers improved performance compared to the baseline full-frame auto-correlation algorithm in

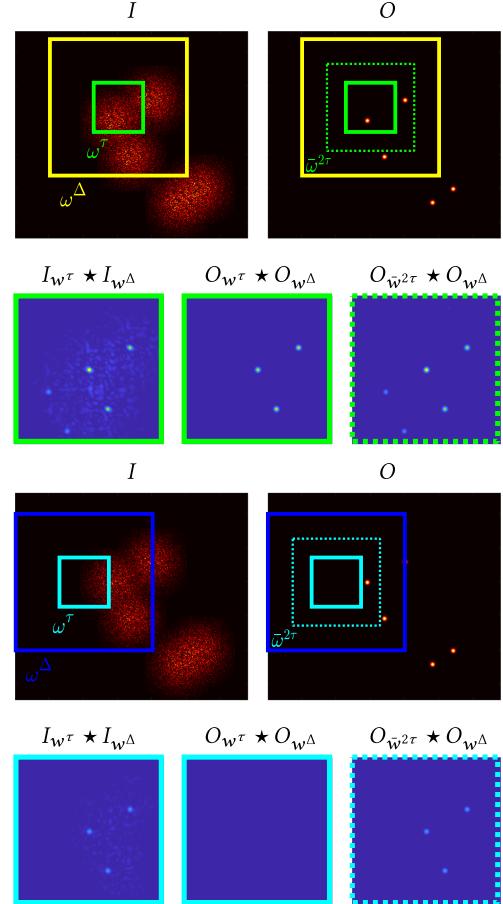


Fig. 8. **Local window selection for optimization.** We consider local subwindows w^τ (light green and cyan frames) whose support is equivalent to the speckle support size. Each such window is correlated with a wider window w^Δ (yellow and blue frames) around it, whose support is equivalent to the ME range. As speckle inside window w^τ can arise from a source outside w^τ , $O_{w^\tau} \star O_{w^\Delta}$ may not match $I_{w^\tau} \star I_{w^\Delta}$. To overcome this, we use an extended non-binary sub-window $\bar{w}^{2\tau} = w^\tau \star w^\tau$ for O , whose support is indicated by dashed lines.

situations where $T_\tau < T_\Delta$, namely, when the support from one illuminator is lower than the ME range. For thick scattering slices, where high-order scattering is dominant, this relationship does not hold, and our approach reduces to the baseline full-frame auto-correlation algorithm of Equation (10).

Our algorithm searches for a latent image O such that the auto-correlation in its local windows will match the auto-correlation in the local windows of the input image I . We define w^Δ and w^τ to be binary windows with support T_Δ, T_τ , respectively, and $\bar{w}^{2\tau} = w^\tau \star w^\tau$ —note that, from its definition, $\bar{w}^{2\tau}$ is non-binary. Then, we recover O by solving the optimization problem:

$$\min_O \sum_j \|\bar{I}_{w_j^\tau} \star \bar{I}_{w_j^\Delta} - O_{\bar{w}_j^{2\tau}} \star O_{w_j^\Delta}\|^2, \quad (24)$$

where $\bar{I}_{w_j^\tau}, \bar{I}_{w_j^\Delta}, O_{\bar{w}_j^{2\tau}}, O_{w_j^\Delta}$ denote windows of a given size cropped from the input and latent images, centered around the j th pixel.

Equation (24) uses windows of three different sizes, and we use Figure 8 to visualize their different roles: Each w_j^τ is a small window whose support is equivalent to the expected support size of the speckle pattern due to a single illuminator. w_j^Δ is a larger window around it, corresponding to the maximal displacement T_Δ for which we expect to find correlation, as dictated by the ME range. If the windows w_j^τ, w_j^Δ are centered around pixel $i_{x,y}$, then the Δ entry of the correlation $\bar{I}_{w_j^\tau} \star \bar{I}_{w_j^\Delta}$ is equal to $c_{\text{emp}}(i_{x,y}, \Delta)$ from Equation (18), where the matched filter is approximated by the binary window w^τ .

We note, additionally, that the window cropped from O should be wider than that from I . This is because speckle at a certain pixel can arise from an illuminator within a window around it. For example, in Figure 8, no illuminator is located inside the cyan subwindow of O , but a part of the speckle pattern is contained within the corresponding cyan subwindow of I . As a result $O_{w_j^\tau} \star O_{w_j^\Delta}$ is a zero image, even though $\bar{I}_{w_j^\tau} \star \bar{I}_{w_j^\Delta}$ detects three impulses. It is easy to prove that this can be addressed using the larger, non-binary window $\bar{w}^{2\tau}$ in the latent image, indicated in Figure 8 using dashed lines: In this case, $O_{\bar{w}^{2\tau}} \star O_{w_j^\Delta}$ correctly detects the same three impulses as $\bar{I}_{w_j^\tau} \star \bar{I}_{w_j^\Delta}$.

The motivation for the cost of Equation (24) is that, even if two illuminators in the latent pattern O are at a distance larger than the ME range T_Δ , they can be recovered if there exists a sequence of illuminators between them, where each two consecutive illuminators in the sequence are separated by a distance smaller than T_Δ . For example, in Figure 8, the illuminators outside the yellow and cyan w^Δ windows are recovered, thanks to the intermediate illuminators.

As pre-processing for our optimization procedure, we form an approximation for the zero mean speckle signal defined in Equation (8) by subtracting the local mean of each window:

$$\bar{I} \approx I - G * I, \quad (25)$$

where G is a Gaussian blur filter. The optimization problem in Equation (24) is no longer a phase retrieval problem as in standard full-frame auto-correlation algorithms. We minimize it using the ADAM gradient-based optimizer [Kingma and Ba 2014]. Gradient evaluation is described in Appendix A.3 and reduces to a sequence of convolution operations that can be performed efficiently, e.g., using a GPU-based fast Fourier transform. For initialization, we set the latent image to random noise; we have observed empirically that the optimization is fairly insensitive to initialization. Finally, we note that even though we could place a window w_j around every pixel of I , the empirical correlation is insensitive to small displacements of the central pixel j . Therefore, in practice, we consider windows only at strides $T_\tau/2$, which helps reduce computational complexity.

To conclude this section, we note that the optimization problem of Equation (24) is similar to ptychography algorithms [Rodenburg et al. 2007]. However, we emphasize that previous ptychographic approaches for extending the ME range recover the latent illuminators from *multiple* image measurements, captured by illuminating different areas on the scattering sample [Zhou et al. 2020; Gardner et al. 2019; Li et al. 2019a, 2019b; Shekel and Katz

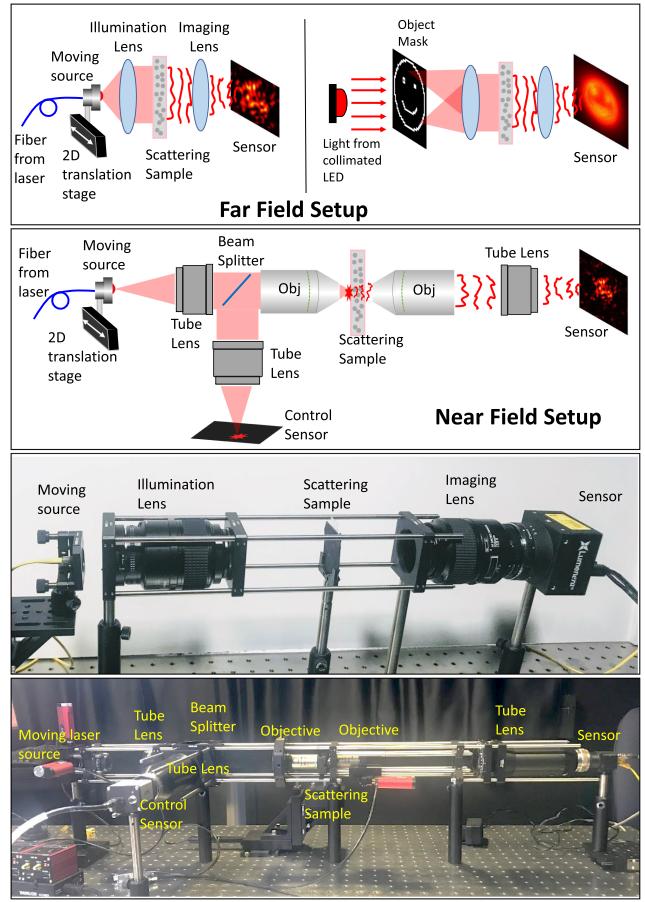


Fig. 9. **Hardware setup.** Top row: schematic of far-field setup, demonstrating two illumination configurations used in experimental setup. The first translates a single point source (fibered laser). The second uses an LED source with masked area of emission. Second row: schematic of near-field setup. Lower panels visualize our hardware lab setup.

2020]. By contrast, our algorithm recovers the latent illuminators from a *single* shot.

8 EXPERIMENTS

We begin by evaluating our algorithm in the far-field setting, demonstrating that even in this setting it provides an order-of-magnitude extension of both range and density of illuminators that can be recovered, compared to the full-frame auto-correlation algorithm. We then proceed to show experiments in near-field setting, demonstrating again significant improvement over previous approaches. We discuss the challenges of the near-field setting and show that they are in agreement with our theoretical analysis.

8.1 Experimental Setup

We built two hardware setups shown in Figure 9, implementing near-field and far-field imaging configurations.

For the near-field setup, we use a tube lens and an objective lens to focus a point source (the output of a single mode fiber connected to a 632nm laser) into a point source at the back side of a

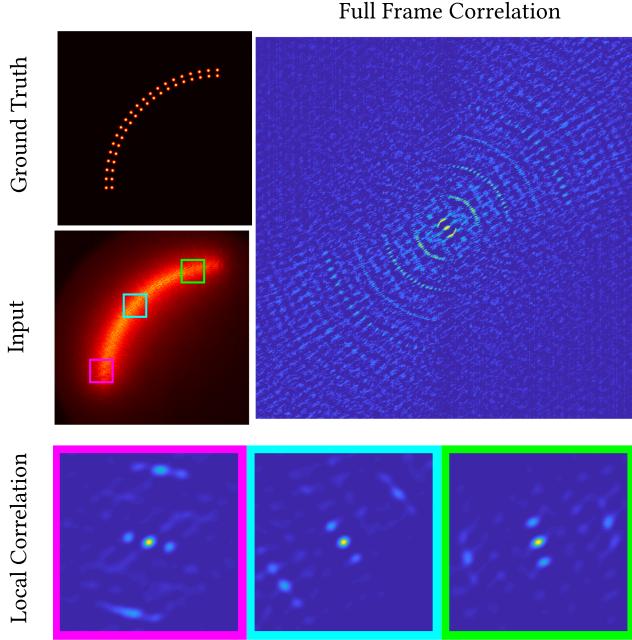


Fig. 10. **Local versus global auto-correlation.** The orientation of the auto-correlation evaluated in three different local windows of the image matches the orientation of the arc in the corresponding region of the latent image. By contrast, the auto-correlation of the full frame is much noisier and decays for large displacements due to limited ME.

scattering sample. To image the sample, we use a camera placed at the opposite side of the sample, similarly equipped with a tube lens and an objective lens. As we discussed in Section 4, verifying that the point source is focused exactly on the back face of the sample is important for ensuring that our experiments are representative of realistic scenarios where illuminators are located inside the tissue rather than beyond it. We confirmed that using a second, control camera. Additionally, as discussed in Section 4, we use a focused configuration where the camera is also focused at the same plane as the point source, at the back face of the sample. We verified the camera focusing by scanning its objective lens along the \hat{z} axis, capturing a focal stack, and selecting the position where speckle support size is smallest and ME range is largest. We used Nikon N20X-PF objectives with $NA = 0.5$ and $\times 20$ magnification, and Thorlabs TTL200 tube lenses.

For the far-field setup, instead of placing the source at a large distance from the sample, we placed it at the Fourier plane of a lens, creating fully directional illumination. This configuration is equivalent to a point source infinitely far away from the sample and allows for better light efficiency. Likewise, we use a camera on the opposite side of the sample, focused at infinity. This setup corresponds to a 4F system around the sample, which we implement using two macro lenses (Nikon 105mm f/2.8D).

For scattering samples, we use slices of chicken breast of thicknesses ranging between $100 - 400\mu\text{m}$. We measure the thickness of samples by placing them between two microscope slides of known dimensions and using a caliper to measure the total thickness.

In both the near-field and far-field setups, we translate the laser point source at different locations behind the sample, capturing different images I^k at each location. We then sum these intensity images to form $I = \sum_k I^k$, simulating the input from multiple mutually incoherent sources. Having access to the individual I^k images is useful for analyzing various algorithmic tradeoffs. We also use a second, single-shot setup, consisting of a binary mask illuminated by a wide-area, spatially incoherent 625nm LED.

8.2 Far-field Experiments

Local auto-correlation. In Figure 10, we visualize the different structure of local and global auto-correlations. Computing auto-correlation at small subwindows of the speckle image reveals the local orientation of the arc in the latent image. By contrast, when computing the auto-correlation of the full frame, the correlation is considerably noisier even for small displacements. Correlations between far illuminators are even harder to detect due to the limited ME range.

Range and density. As discussed in Section 6, a fundamental limitation of imaging-through scattering algorithms is the density of illuminators they can recover. To demonstrate this, in Figure 11, we compare recovery results for illuminator patterns of the same range and layout, but at different densities. In each case, we display the densest subset at which the full-frame and our local auto-correlation algorithms successfully recovered the latent pattern, with our local approach often handling order-of-magnitude larger densities. We captured the data by imaging speckle patterns created by individual point sources placed at different locations and summing the speckle images in post-processing, allowing us to form test images at any density of interest. Details on the full-frame phase-retrieval algorithm we used, as well as a comparison to the sparse approach of Chang and Wetzstein [2018] are provided in Section 8.4.

As we have access to the speckle images generated at each illuminator location, we can compute the decay of ME across the frame. Denoting by I^1, I^2 the individual speckle images from illuminator locations i^1, i^2 , and setting $\Delta_{1,2} = i_{x,y}^2 - i_{x,y}^1$, we evaluate:

$$C(\Delta_{1,2}) = \frac{\sum_{xy} \bar{I}^1(x,y) \cdot \bar{I}^2((x,y) + \Delta_{1,2})}{\sqrt{\sum_{xy} \bar{I}^1(x,y)^2} \sqrt{\sum_{xy} \bar{I}^2(x,y)^2}}. \quad (26)$$

We plot this correlation at the right of Figure 12, as a function of $|\Delta|$. We note that for the smile pattern, which was captured with a thin tissue layer, the ME range covers the entire frame (empirical correlation does not decrease below 0.8 even for the widest displacement). Even under these favorable conditions for the full-frame auto-correlation algorithm, our local algorithm recovers a denser set of illuminators.

In Figure 12 and additionally in Figure 24 of Appendix A.4, we demonstrate the increased range that our algorithm provides. To achieve this, we select a few local subwindows from the patterns in Figure 11 and display the maximal window for which the full-frame auto-correlation was successful—each pair of subwindows demonstrates a small window with reasonable reconstruction and a slightly bigger one where reconstruction already failed. Overall, in Figure 11 our local algorithm successfully handles patterns that

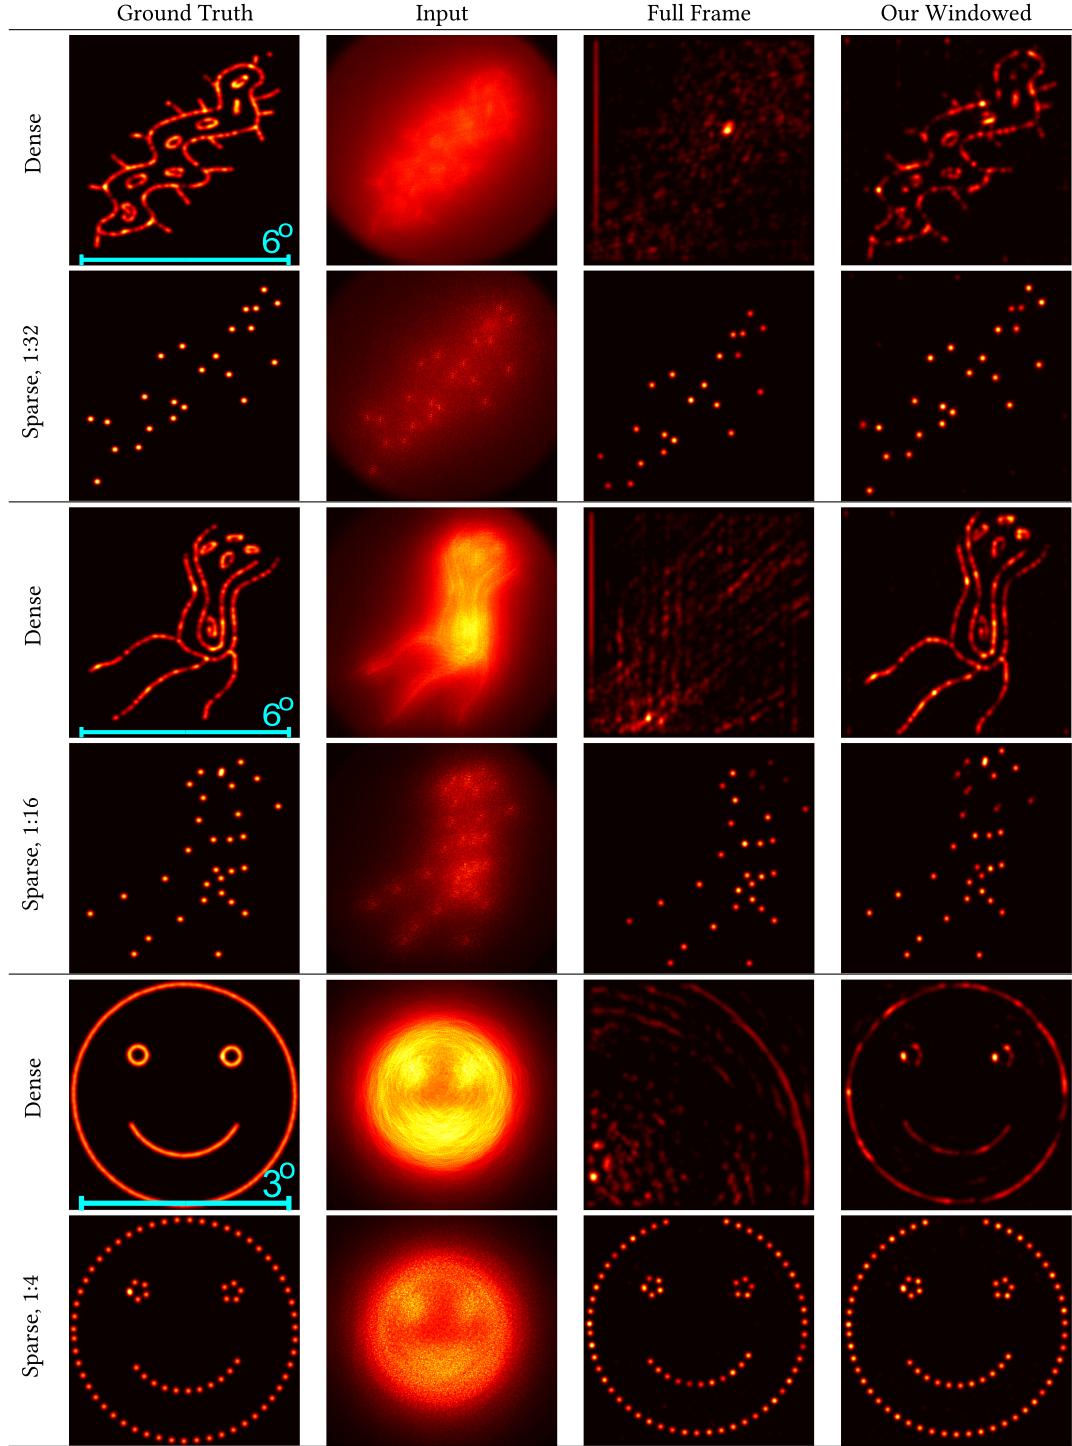


Fig. 11. **Comparison of our local and the full-frame auto-correlation algorithms.** For each example, we show the densest arrangement of illuminators for which the full-frame auto-correlation algorithm succeeded. In the top example, our algorithm successfully recovered $\times 32$ more illuminators. Even in the lower example where the ME extends over the entire frame (see correlation plots in Figure 12), our local approach outperforms the full-frame one. The tissue thickness of each example, from top to bottom, is $330\mu m$, $340\mu m$, and $200\mu m$.

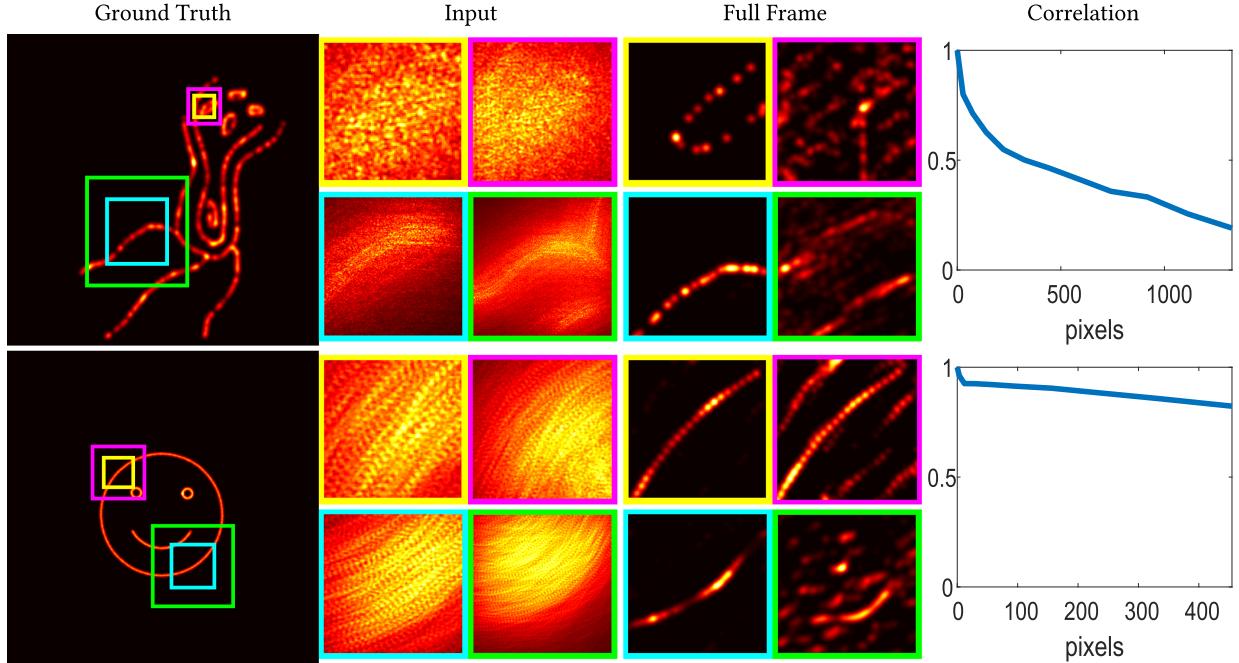


Fig. 12. **Full-frame auto-correlation algorithm applied to small crops of the patterns in Figure 11.** The yellow and cyan sub-windows demonstrate areas where reconstruction roughly succeeds, and the magenta and green ones a slightly larger window where reconstruction fails. To the right, we plot correlation as a function of displacement length $|\Delta|$, as measured for the corresponding tissue slice. Tissue thickness from top to bottom, are $340\mu\text{m}$ and $200\mu\text{m}$.

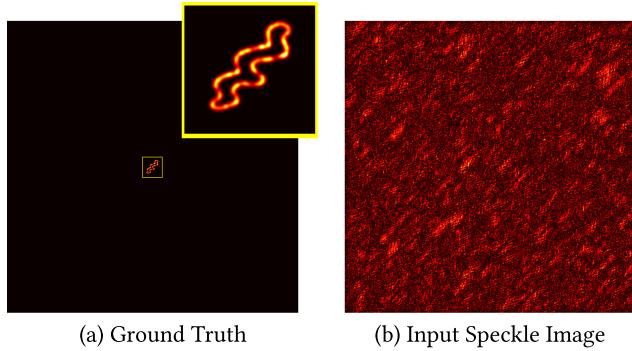


Fig. 13. **Classical setup illustration.** In the classical full-frame auto-correlation setup, the latent image is usually significantly smaller than the sensor width, and the number of speckle features it includes is about $\times 10^4$ higher than the number of independent illuminators.

are an order of magnitude wider than the maximal patterns recovered by the full-frame approach in Figure 12. There is no inherent limit preventing us from handling an even larger range, except that in the specific experimental setup we used, increasing the range would exceed the aperture width.

For the small images of Figure 12, our algorithm is equivalent to the full-frame auto-correlation algorithm, as the images are small enough that they do not fit more than one w^τ window. While both approaches fail on small images (Figure 12), the local approach is successful when applied to larger images, where the full-frame

algorithm still fails (Figure 11). This is because, when considering a larger image, our algorithm computes correlation with speckle patterns at other parts of the frame, providing additional constraints.

Contrasting with classical setup. The patterns recovered in our implementation are very different from the ones used in previous full-frame auto-correlation implementations [Katz et al. 2014]. The patterns in Figure 11 included about 10^3 illuminators spread near-uniformly across the area of a 2-megapixel sensor. By contrast, Figure 13 shows a typical input for previous full-frame implementations, where the target pattern is concentrated within a small area of about 100×100 pixels. Yet, the speckle support is much larger, covering the entire sensor. To achieve this wide speckle spread, previous implementations either imaged the sample with a lensless sensor rather than a focused one or used scattering layers that are thicker or have wider phase functions. In synthetic simulations of such a full-frame setup, the phase retrieval algorithm by Fienup [1982] usually fails if more than 100 sources are included; The number of sources can slightly increase with a better phase retrieval approach.

Finally, it is worth noting that, as illuminator density increases, the local approach eventually fails as well. We include an example of such a failure case in Figure 22 of Appendix A.4.

LED illumination. In Figure 14 and in Appendix A.4, we show reconstructions from the single-shot setup of Figure 9, where the entire area of a target mask is illuminated by spatially incoherent LED light. The main challenge in this case arises from the fact that the illumination is no longer purely monochromatic: Different wavelengths are diffracted in slightly different angles, blurring

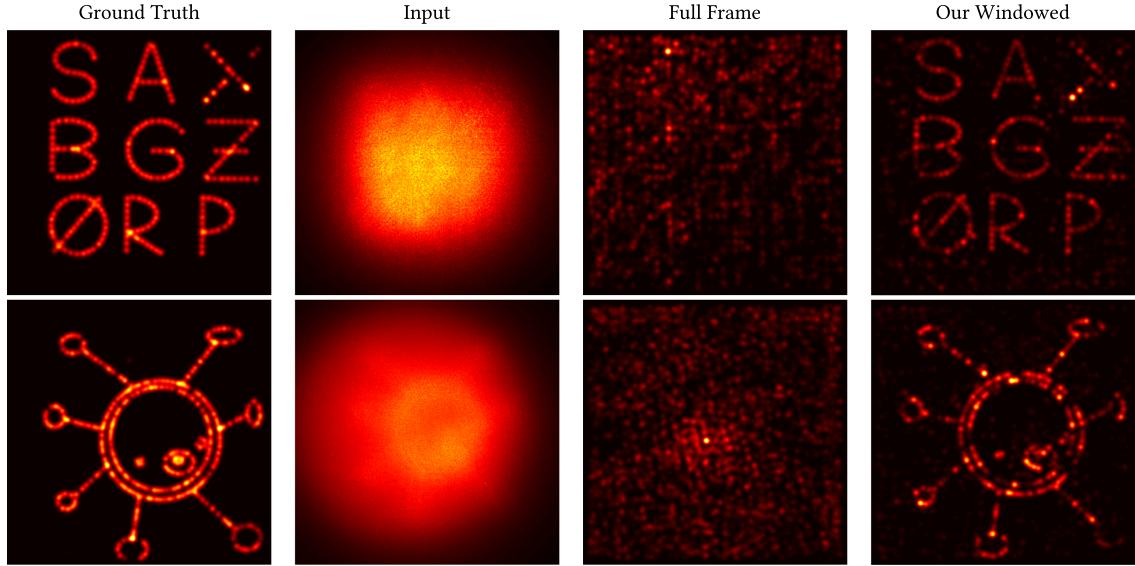


Fig. 14. **Far-field reconstruction using a single shot acquisition setup.** Our local correlation approach outperform the classical full-frame auto-correlation.

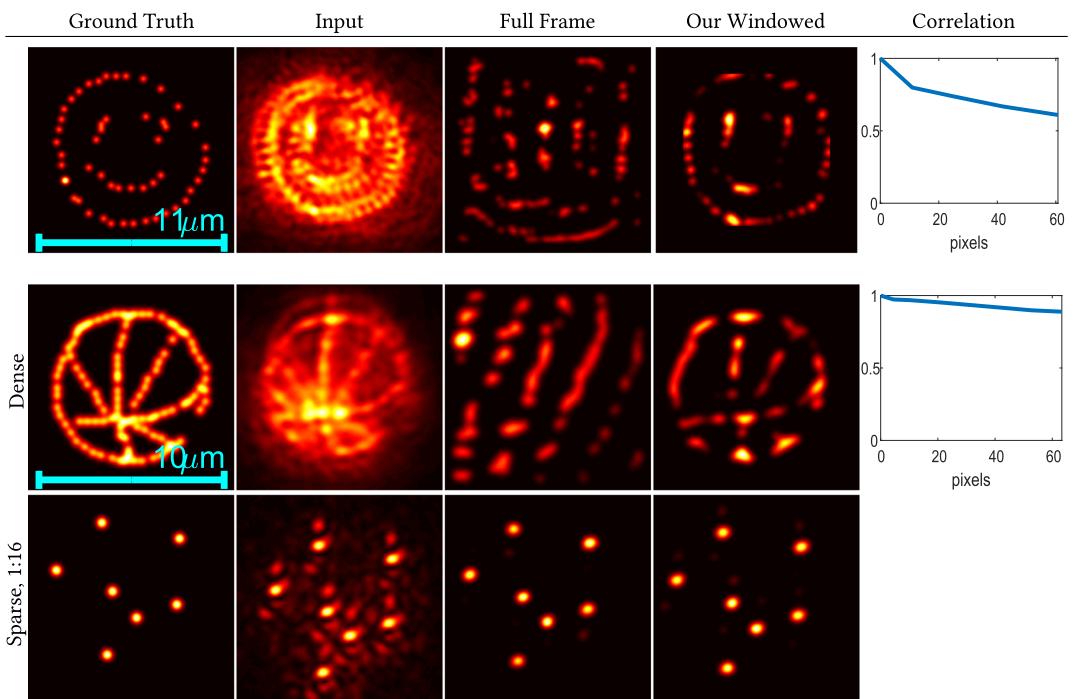


Fig. 15. **Near-field comparisons.** We compare our local and the full-frame auto-correlation algorithms on thin tissue examples ($L = 100 - 140\mu\text{m}$). Despite the seemingly small degradation, the full-frame approach fails unless provided an input composed of a considerably sparser set of illuminators. To the right, we plot correlation decay as measured using images of individual sources through the corresponding tissue slice.

speckle contrast. In the mid-order scattering examples that we are considering, this is mitigated by the fact that the speckle support size is limited, meaning that speckle patterns are less affected by blur. To reduce this effect, we placed a 10nm band-pass filter at the LED output.

8.3 Near-field Experiments

In Figure 1 and Figures 15 and 16, we show reconstruction results from our near-field setup. The first set of examples (Figures 1 and 15) come from thinner tissue layers ($L = 100 - 150\mu\text{m}$), for which scattering is modest and the latent pattern may be

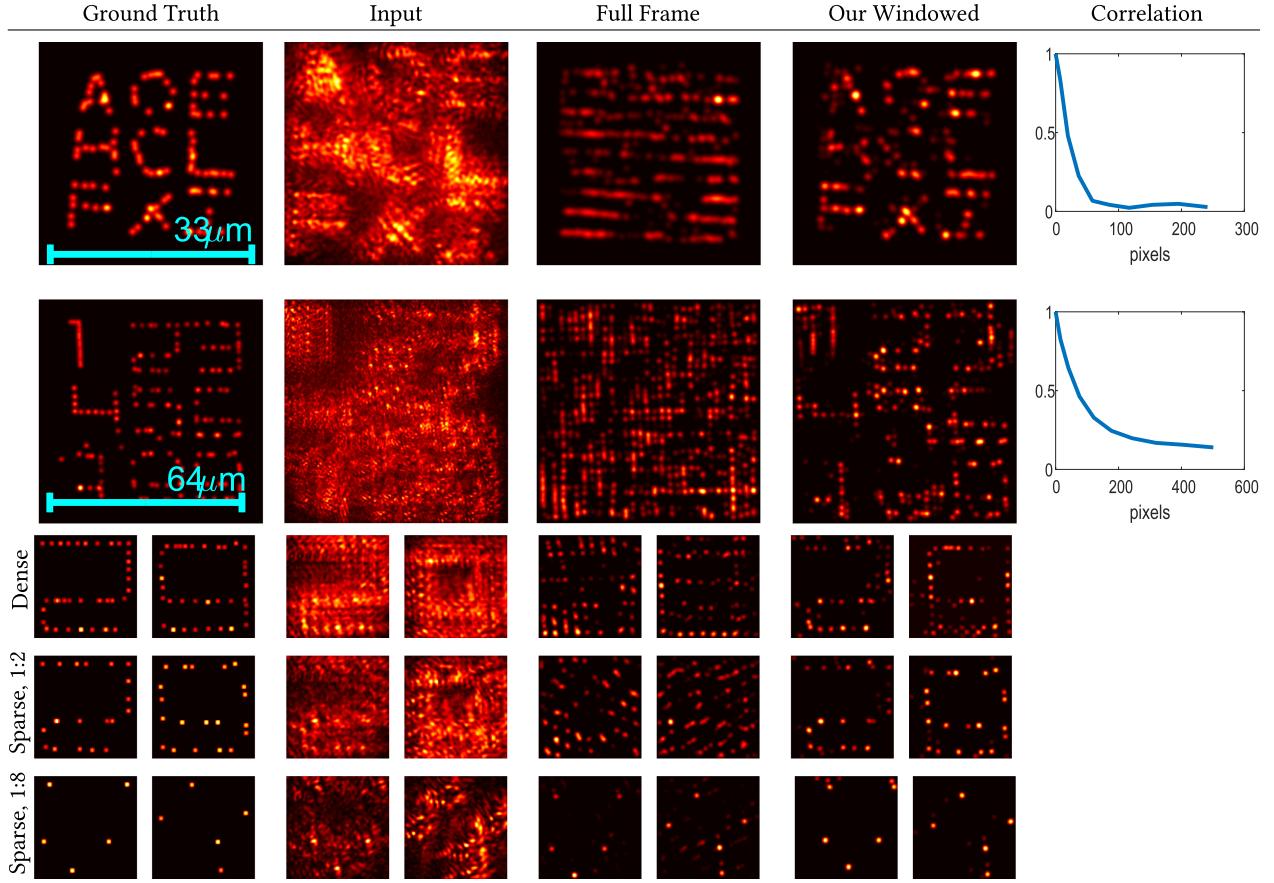


Fig. 16. More near-field comparisons. We now use a thicker tissue slice ($L = 200\mu\text{m}$). The lower rows zoom on two of the digits in the second row, demonstrating reconstruction at a few different densities. Reconstruction degrades as density increases. The full-frame approach was successful at a considerably lower density than our local approach.

recognizable from the degraded input image. Our algorithm still improves the pattern quality significantly and reconstructs fine details obscured by the speckle. However, even this modest degradation is challenging for the full-frame approach, which fails unless applied on a significantly sparser set of illuminators (last rows of Figures 1 and 15).

For the second set of examples in Figure 16, the sample thickness is larger at $L \approx 200\mu\text{m}$, and thus the degradation is stronger. This reconstruction task is more challenging for two reasons: First, the ME range is limited, as seen by the correlation curve at the right of Figure 16. Second, the illuminator density is large. Our local algorithm outperforms the full-frame auto-correlation approach, but its reconstruction is not free of artifacts either. In the lower rows, we zoom on two of the digit patterns, showing reconstructions for different illuminator densities. Our local approach performs worse as density increases, but still outperforms the full-frame approach, which is successful only at considerably smaller densities.

To demonstrate the difference between the far-field and near-field settings, in Figure 17, we compare reconstructions of a line pattern of the exact same length, from far-field and near-field measurements. We can see that the far-field speckles cover a much larger area of the sensor than the near-field ones. These additional speckle-covered pixels help improve the SNR of the far-field

correlation estimates. The explanation for this difference is the same as for the difference between the small and medium supports in Figure 7(b) and (c): The near-field image in Figure 17, due to the many pixels that do not receive light, has a larger effective illuminator density than the far-field image. This larger density results in worse correlation SNR, and thus reduced reconstruction quality. As another way to see this, we show in Figure 18 speckle patterns from a single point source captured under near-field and far-field conditions and the corresponding auto-correlation $\bar{S} \star \bar{S}$. Both auto-correlations resemble an impulse, up to noise. However, we observe that the near-field speckle image includes much fewer speckle features than the far-field one. Consequently, there is more noise in the near-field auto-correlation image than in the far-field one. This increased noise results in the reduced quality of the near-field reconstructions.

Previous near-field implementations. The only reported attempt to apply speckle auto-correlation techniques in the near-field setting that we are aware of is by Chang and Wetzstein [2018]. A direct comparison between our results and theirs is not possible, as the two sets of experiments use very different scattering samples. In particular, the speckle images captured by Chang and Wetzstein [2018] are similar to the ones used by Katz et al. [2014], comprising a target pattern concentrated with a small sensor area,

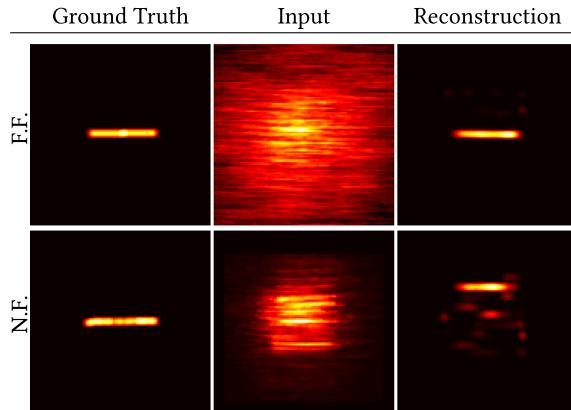


Fig. 17. **Comparison of near-field and far-field settings for a line pattern.** Even though the ground truth pattern is similar, the near-field speckle pattern has a much smaller support size. As more speckle pixels are provided, the far-field correlation is less noisy, improving reconstruction quality.

producing non-localized speckle patterns covering the entire sensor. The ME range reported by Chang and Wetzstein [2018] was approximately equal to a $10\mu\text{m}$ displacement. Accordingly, their experiments recovered illuminator patterns of size $10\mu\text{m} \times 10\mu\text{m}$. By contrast, the size of the near-field illuminator patterns we recovered scales up to $65\mu\text{m} \times 65\mu\text{m}$. Below, we additionally compare with their robust phase-retrieval algorithm, using measurements captured with our imaging setups.

8.4 Comparison to Alternative Algorithms and Limitations

In all of the previous figures, the reconstructions for the full-frame approach are achieved using an ADAM optimization procedure [Kingma and Ba 2014] with non-negativity constraints. This is analogous to Wirtinger flow optimization for phase retrieval [Candes et al. 2014; Chakravarthula et al. 2019]. We found that this approach works better than the classical optimization algorithm by Fienup [1982]. We show comparisons between these two algorithms in Figure 19, using sparse and dense sets of illuminators for both the far-field and near-field examples from Figures 11 and 16. For these examples the Fienup-based full-frame variant (Figure 19(b)) did not converge, whereas the ADAM-based full frame variant converged on the sparse set and failed to converge on the dense one. We additionally attempted to optimize the full-frame approach with the ADMM-based phase retrieval algorithm of Chang and Wetzstein [2018], which uses an L_1 regularization term. As shown in Figure 19(c), this performed better than the Fienup-based variant, but provided results very similar to the ADAM-based variant for most examples. We believe this is because the non-negativity constraints we enforce during ADAM optimization already leads to sparse solutions. We note in Figure 19 that all variants of the full-frame approach fail as we further increase the illuminator density. By contrast, our local correlation approach is successful in the higher density case, as shown in Figures 11 and 16.

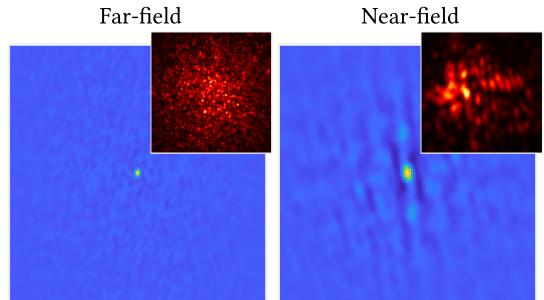


Fig. 18. **Single-point-source auto-correlation.** The auto-correlation of speckle images due to a single source $\bar{S} \star \bar{S}$ resembles an impulse plus noise, in both the near-field and far-field cases. However, in the near-field case, this auto-correlation is noisier, as fewer speckle features are averaged. The single-source speckle images \bar{S} are shown in the insets.

Given the local extent of the ME, another option one may consider is cropping local windows from the full speckle image, running the full-frame auto-correlation approach on each local window, and then seaming the individual local solutions. However, as shown in Figure 19(e), the independent solutions are rarely consistent and the seamed result has strong artifacts. We also note that the solution of the full-frame approach in each local window is only defined up to an arbitrary flip or shift. For the result in Figure 19(e), we favored this algorithm by flipping and shifting each window to best match the groundtruth. Even under this simplification, this algorithm is inferior to our approach that jointly optimizes all local windows.

Finally, in Figure 19(f), we compare against the approach of Wang et al. [2019] for extending the range of imaging-through-scattering algorithms. Their model assumes the latent image O can be decomposed into two parts O_1, O_2 of a smaller extent, and the ME applies in each window separately. Mathematically, they model the image formation as $I = O_1 \star S_1 + O_2 \star S_2$, where S_1 and S_2 are the speckle patterns from a single illuminator in each region, and which are assumed to be decorrelated, $S_1 \star S_2 = 0$. They then try to simultaneously solve for two smaller support images O_1, O_2 satisfying $O_1 \star O_1 + O_2 \star O_2 = I \star I$. As seen in Figure 19, this approach was successful on the far-field sparse examples, but failed on the denser and near-field examples. The reason for this is that, in our examples, the correlation decays gradually, and thus the assumption by Wang et al. [2019] that speckle patterns at different parts of the image are completely decorrelated does not hold. Additionally, the range of the latent illuminators is much larger than twice the ME range.

Runtime. Compared to the full-frame auto-correlation approach, one disadvantage of our algorithm is increased computational cost. In particular, our unoptimized Matlab implementation, when running on an NVIDIA Quadro RTX 8000 GPU, requires a few hours to converge for each of the results in this section. Gradient evaluations make up the bulk of this runtime. Gradient evaluations are essentially a sequence of convolution operations, which we implement using Matlab's GPU-based fast Fourier transform function. These gradient evaluations can potentially be accelerated using more sophisticated GPU-based convolution libraries. We also note that our optimization procedure only needs to be run once, as its

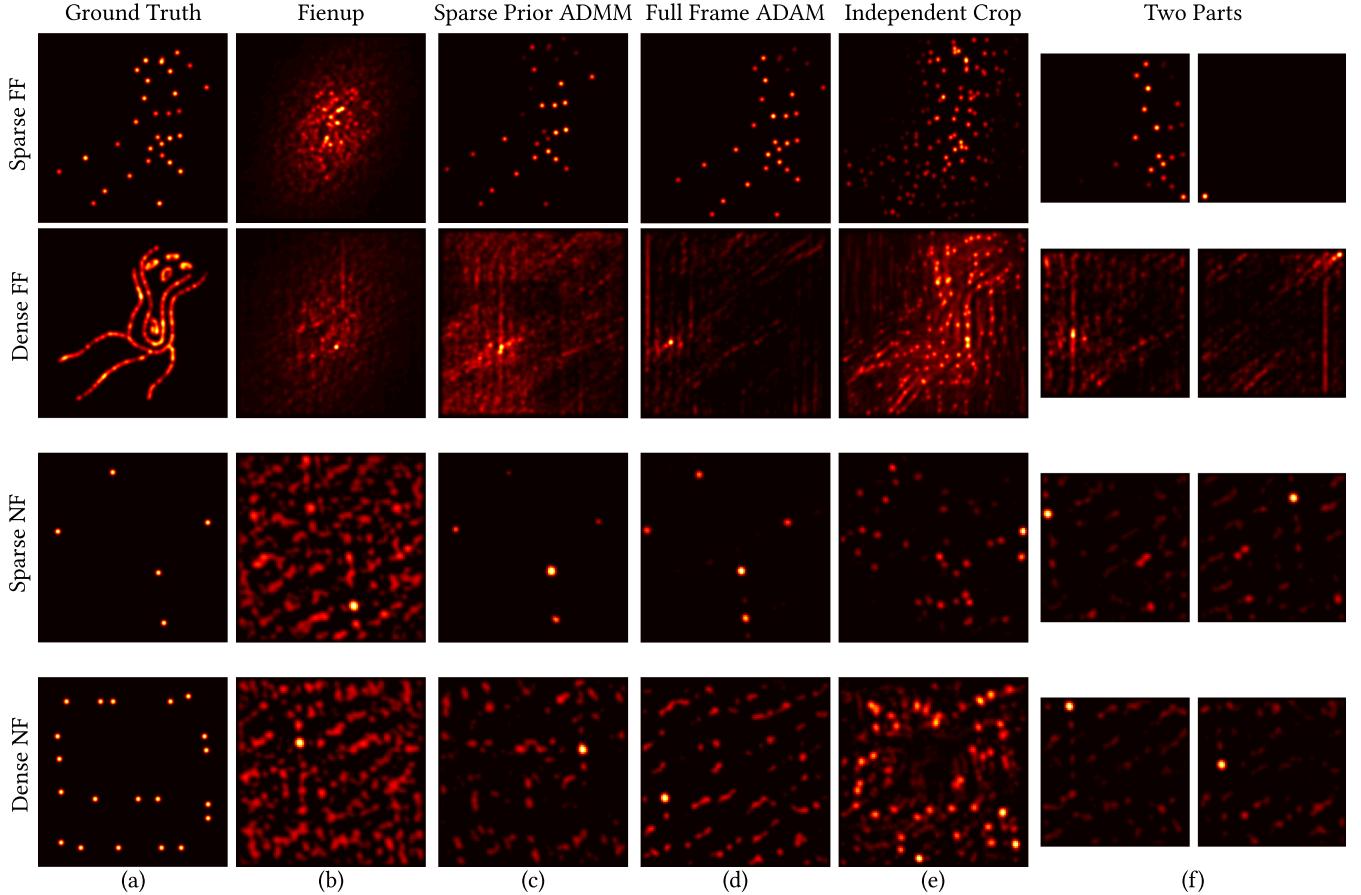


Fig. 19. **Comparison with alternative reconstruction algorithms.** Considering far- and near-field examples visualized in Figures 11 and 16, we evaluate a few alternative strategies. (b)–(d) show different full-frame phase retrieval approaches. Classical Fienup optimization [Fienup 1982] is rather noise-sensitive. Chang and Wetzstein [2018] proposed a better algorithm introducing a sparse prior and ADMM optimization. In this work, we used gradient decent update with ADAM step selection size, with similar results. These algorithms can solve the full frame phase retrieval on a sparse subset of the sources, but fail on a denser one. In contrast, our local cost led to much better results, as presented in Figures 11 and 16. (e) Solving the standard phase retrieval problem on independent local windows and seaming the results in post-processing leads to noticeable artifact. (f) Seeing beyond the ME range [Wang et al. 2019] by decomposing the speckle auto-correlation into two independent parts is sometimes successful on the sparse data, yet fails on the dense one.

results are insensitive to initialization. By contrast, iterative phase retrieval algorithms such as the algorithm by Fienup [1982] typically require multiple runs with different initializations.

9 DISCUSSION

We provided a comprehensive study of algorithms using the speckle ME to image through scattering. Using theory, simulations, and real experiments, we investigated the inherent limits of these approaches. In particular, we explored whether these approaches can be applied to practical biomedical imaging scenarios, where illumination sources are located inside, rather than far behind, a scattering sample such as a tissue layer.

We reported the following important findings of our study: First, we showed that the ME is affected by the angular difference between illumination sources and not by their actual displacement. Second, we showed that the correlation of speckle intensities can be improved through simple design choices in the imaging setup;

in particular, we found that the ME range is maximized when using a lens focused at the illuminator plane. Third, we showed that it is important to closely replicate near-field imaging conditions by ensuring that illuminators are accurately placed exactly behind the scattering sample, as doing otherwise can artificially increase the ME range. Fourth, we found that, in the near-field setting, the angular displacement for which significant correlations exist can correspond to actual displacements smaller than the illumination wavelength. As a consequence, ME approaches are only applicable to the near-field setting when considering scattering samples of modest thickness, where mid-order scattering is dominant. This thickness range still corresponds to penetration depths considerably beyond those achievable by a standard microscope. Therefore, analyzing and developing new ME approaches for imaging through scattering in this range can benefit biomedical imaging applications.

Our study additionally highlighted an important property of speckle intensity patterns formed due to samples where mid-order

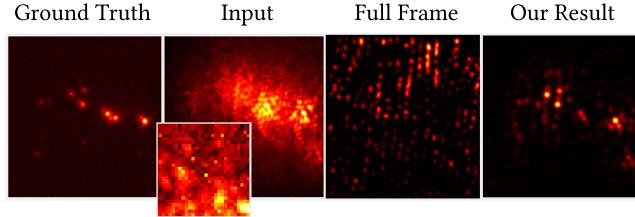


Fig. 20. **Reconstructing fluorescent beads.** We demonstrate single-shot near-field imaging of fluorescent beads attached at the back of a $[100]\mu\text{m}$ chicken breast slice and the reconstruction produced by the full-frame and our local auto-correlation approach. Speckle contrast is low due to the weak emission of fluorescent beads (see a close-up on imaging noise in the inset) and due to the fact that the emission is not fully monochromatic.

scattering is dominant when imaged with a focused lens. These patterns have a small support size, typically much smaller than the sensor size. We showed theoretically and experimentally that using this local support to create a matched filter when computing speckle correlations can boost the SNR of latent illuminator detection by orders of magnitude. We additionally developed an algorithm that takes advantage of this property by operating as a local version of classical full-frame auto-correlation techniques. Through experiments on real measurements captured using both far-field and near-field imaging configurations, we showed that our algorithm provides an order-of-magnitude improvement in terms of both the range and the density of recoverable illuminators.

Furthermore, our study shed light on two fundamental challenges associated with the near-field case. The first challenge is the fact that the ME holds for very small displacements. The second challenge relates to the fact that, even after exploiting the local support, only sparse latent patterns are recoverable. These challenges still leave ample room for applications in medical imaging settings where sparse targets are considered, for example STORM imaging of blinking fluorescent molecules, sparse nuclei, or other cell components.

Additional challenges can arise due to reduced speckle contrast and signal-to-noise ratio in measurements captured under real near-field fluorescent imaging conditions. To highlight these challenges, in Figure 20, we show captured speckled patterns generated by a sparse set of fluorescent beads placed at the back of a tissue sample, as well as the reconstructions produced by the full-frame and our local auto-correlation techniques. Details on the experimental setup are provided in Appendix A.6. The quality of the input speckle images is severely affected by two factors. First, as in Figure 14, the light emitted by the fluorescent beads is spectrally broadband, reducing speckle contrast. Second, the fluorescent emission is very weak, leading to noisy images. Despite these challenges, we observe that our local auto-correlation algorithm significantly improves reconstruction quality compared to the full-frame baseline. Therefore, our results showcase both the strong potential of local auto-correlation techniques compared to full-frame variants, and the need for further research towards ME techniques that are fully robust to real experimental conditions in applications such as fluorescence microscopy.

Last but not least, by drawing attention to the local support characteristics of near-field speckle images, our results open the door for future research on using different image processing approaches for imaging through scattering, such as local deconvolution [Wu et al. 2020] and sharpening operations. A particularly promising direction is adapting the large array of mature blind deconvolution techniques to the imaging-through-scattering setting by developing appropriate prior models for the scatter-free images and (spatially varying) speckle blur kernels. Our study on the statistics of speckle patterns, and their dependence on the scattering layer geometry and material properties, can help inform the development of such priors.

ACKNOWLEDGMENT

We thank Lucien Weiss for help with the fluorescent samples.

REFERENCES

- Eric Akkermans and Gilles Montambaux. 2007. *Mesoscopic Physics of Electrons and Photons*. Cambridge University Press.
- Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. 2019. A Monte Carlo framework for rendering speckle statistics in scattering media. *ACM Trans. Graph* 38, 4 (2019).
- C. Bar, M. Alterman, I. Gkioulekas, and A. Levin. 2021. Single scattering modeling of speckle correlation. In *ICCP*.
- Chen Bar, Ioannis Gkioulekas, and Anat Levin. 2020. Rendering near-field speckle statistics in scattering media. *ACM Trans. Graph* 39, 6 (2020).
- Mahed Batarseh, Sergey Sukhov, Zhean Shen, Heath Gemar, Roxana Rezvani, and Aristide Dogariu. 2018. Passive sensing around the corner using spatial coherence. *Nat. Commun.* 9, 3629 (2018).
- Ibrahim Baydoun, Diego Baresch, Romain Pierrat, and Arnaud Derode. 2016. Radiative transfer of acoustic waves in continuous complex media: Beyond the Helmholtz equation. *Phys. Rev. E* 94, 5 (2016).
- Richard Berkovits and Sheehao Feng. 1994. Correlations in coherent multiple scattering. *Phys. Rep.* 238, 3 (1994), 135–172.
- Bruce J. Berne and Robert Pecora. 2000. *Dynamic Light Scattering: With Applications to Chemistry, Biology, and Physics*. Courier Corporation.
- Jacopo Bertolotti, Elbert G. van Putten, Christian Blum, Ad Lagendijk, Willem L. Vos, and Allard P. Mosk. 2012. Non-invasive imaging through opaque scattering layers. *Nature* 491, 7423 (2012), 232–234.
- Eric Betzig, George H. Patterson, Rachid Sougrat, O. Wolf Lindwasser, Scott Olenych, Juan S. Bonifacino, Michael W. Davidson, Jennifer Lippincott-Schwartz, and Harald F. Hess. 2006. Imaging intracellular fluorescent proteins at nanometer resolution. *Science* 313, 5793 (2006), 1642–1645.
- David A. Boas, Dana H. Brooks, Eric L. Miller, Charles A. DiMarzio, Misha Kilmer, Richard J. Gaudette, and Quan Zhang. 2001. Imaging the body with diffuse optical tomography. *IEEE Sig. Proc. Mag.* 18, 6 (2001), 57–75.
- David A. Boas and Arjun G. Yodh. 1997. Spatially varying dynamical properties of turbid media probed with diffusing temporal light correlation. *J. Opt. Soc. Am. A* 14, 1 (1997), 192–215.
- Jeremy Boger-Lombard and Ori Katz. 2019. Passive optical time-of-flight for non line-of-sight localization. *Nat. Commun.* 10, 3343 (2019).
- Antoine Boniface, Baptiste Blochet, Jonathan Dong, and Sylvain Gigan. 2019. Non-invasive light focusing in scattering media using speckle variance optimization. *Optica* 6, 11 (2019), 1381–1385.
- Antoine Boniface, Jonathan Dong, and Sylvain Gigan. 2020. Non-invasive focusing and imaging in scattering media with a fluorescence-based transmission matrix. *Nature communications* 11, 6154 (2020).
- Emmanuel Candes, Xiaodong Li, and Mahdi Soltanolkotabi. 2014. Phase retrieval via wirtzinger flow: Theory and algorithms. *IEEE Trans. Inf. Theory* 61, 4 (2014), 1985–2007.
- Thomas Chaigne, Bastien Arnal, Sergey Vilov, Emmanuel Bossy, and Ori Katz. 2017. Super-resolution photoacoustic imaging via flow-induced absorption fluctuations. *Optica* 4, 11 (2017), 1397–1404.
- Thomas Chaigne, Jérôme Gateau, Marc Allain, Ori Katz, Sylvain Gigan, Anne Sentenac, and Emmanuel Bossy. 2016. Super-resolution photoacoustic fluctuation imaging with multiple speckle illumination. *Optica* 3, 1 (2016), 54–57.
- Praneeth Chakravarthula, Yifan Peng, Joel Kollin, Henry Fuchs, and Felix Heide. 2019. Wirtzinger holography for near-eye displays. *ACM Trans. Graph* 38, 6 (2019), 1–13.
- Julie Chang and Gordon Wetzstein. 2018. Single-shot speckle correlation fluorescence microscopy in thick scattering tissue with image reconstruction priors. *J. Biophot.* 11, 3 (2018), e201700224.

Wai-Fung Cheong, Scott A. Prahl, and Ashley J. Welch. 1990. A review of the optical properties of biological tissues. *IEEE J. Quant. Electron.* 26, 12 (1990), 2166–2185.

Youngwoon Choi, Taeseok Daniel Yang, Christopher Fang-Yen, Pilsung Kang, Kyoung Jin Lee, Ramachandra R. Dasari, Michael S. Feld, and Wonshik Choi. 2011. Overcoming the diffraction limit using multiple light scattering in a highly disordered medium. *Phys. Rev. Lett.* 107, 2 (2011), 023902.

Thomas Dertinger, Ryan Colyer, Gopal Iyer, Shimon Weiss, and Jorg Enderlein. 2009. Fast, background-free, 3D super-resolution optical fluctuation imaging (SOFI). *PNAS* 106, 52 (2009), 22287–22292.

Ronald L. Dougherty, Bruce J. Ackerson, Nafaa M. Reguigui, F. Dorri-Nowkoorani, and Ulf Nibbmann. 1994. Correlation transfer: Development and application. *JQSRT* 52, 6 (1994), 713–727.

Turgut Durduran, Regine Choe, Wesley B. Baker, and Arjun G. Yodh. 2010. Diffuse optics for tissue monitoring and tomography. *Rep. Prog. Phys.* 73, 7 (2010), 076701.

Eitan Edrei and Giuliano Scarselli. 2016b. Memory-effect based deconvolution microscopy for super-resolution imaging through scattering media. *Sci. Rep.* 6, 33558 (2016).

Eitan Edrei and Giuliano Scarselli. 2016a. Optical imaging through dynamic turbid media using the Fourier-domain shower-curtain effect. *Optica* 3, 1 (2016), 71–74.

Shechao Feng, Charles Kane, Patrick A. Lee, and A. Douglas Stone. 1988. Correlations and fluctuations of coherent wave transmission through disordered media. *Phys. Rev. Lett.* 61, 7 (1988), 834.

James R. Fienup. 1982. Phase retrieval algorithms: A comparison. *Appl. Opt.* 21, 15 (1982), 2758–2769.

Isaac Freund. 1990. Looking through walls and around corners. *Physica A* 168, 1 (1990), 49–65.

Isaac Freund and Danny Eliyahu. 1992. Surface correlations in multiple-scattering media. *Phys. Rev. A* 45, 8 (1992), 6133.

David L. Fried. 1982. Anisoplanatism in adaptive optics. *J. Opt. Soc. Am.* 72, 1 (1982), 52–61.

Dennis F. Gardner, Shawn Divitt, and Abbie T. Watnik. 2019. Ptychographic imaging of incoherently illuminated extended objects using speckle correlations. *Appl. Opt.* 58, 13 (2019), 3564–3569.

Walter I. Goldburg. 1999. Dynamic light scattering. *Am. J. Phys.* 67, 12 (1999), 1152–1160.

Enlai Guo, Shuo Zhu, Yan Sun, Lianfa Bai, Chao Zuo, and Jing Han. 2020. Learning-based method to reconstruct complex targets through scattering medium beyond the memory effect. *Opt. Express* 28, 2 (2020), 2433–2446.

Matthias Hofer, Christian Soeller, Sophie Brasselet, and Jacopo Bertolotti. 2018. Wide field fluorescence epi-microscopy behind a scattering medium enabled by speckle correlations. *Opt. Express* 26, 8 (2018), 9866–9881.

Roarke Horstmeyer, Haowen Ruan, and Changhuei Yang. 2015. Guidestar-assisted wavefront-shaping methods for focusing light into biological tissue. *Nat. Photon.* 9, 9 (2015), 563–571.

Takanori Igarashi, Ko Nishino, and Shree K. Nayar. 2007. The appearance of human skin: A survey. *Found. Trends Comput. Graph. Vis.* (2007).

Pierre Jacquot and Pramod K. Rastogi. 1979. Speckle motions induced by rigid-body movements in free-space geometry: An explicit investigation and extension to new cases. *Appl. Opt.* 18, 12 (1979), 2022–2032.

Michael L. Jakobsen, Hal T. Yura, and Steen G. Hanson. 2012. Spatial filtering velocimetry of objective speckles for measuring out-of-plane motion. *Appl. Opt.* (2012).

Kensei Jo, Mohit Gupta, and Shree Nayar. 2015. SpeDo: 6 DOF ego-motion sensor using speckle defocus imaging. *ICCV* (2015), 4319–4327.

Benjamin Judkewitz, Roarke Horstmeyer, Ivo Vellekoop, and Changhuei Yang. 2014. Translation correlations in anisotropically scattering media. *Nat. Phys.* 11, 8 (2015), 684–689.

Benjamin Judkewitz, Ying Wang, Roarke Horstmeyer, Alexandre Mathy, and Changhuei Yang. 2013. Speckle-scale focusing in the diffusive regime with time-reversal of variance-encoded light (TROVE). *Nat. Photon.* 7, 4 (2013), 300–305.

Ori Katz, Yaron Bromberg, Eran Small, and Yaron Silberberg. 2010. Focusing and compression of ultrashort pulses through scattering media. *Nat. Photon.* 5, 6 (2010), 372–377.

Ori Katz, Pierre Heidmann, Mathias Fink, and Sylvain Gigan. 2014. Non-invasive single-shot imaging through scattering layers and around corners via speckle correlation. *Nat. Photon.* 8, 10 (2014), 784–790.

Ori Katz, Eran Small, and Yaron Silberberg. 2012. Looking around corners and through thin turbid layers in real time with scattered incoherent light. *Nat. Photon.* 6, 8 (2012), 549–553.

Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *ICLR* (2015).

Puxiang Lai, Lidai Wang, Jian Wei Tay, and Lihong V. Wang. 2015. Photoacoustically guided wavefront shaping for enhanced optical focusing in scattering media. *Nat. Photon.* 9, 2 (2015), 126–132.

Guowei Li, Yang Wanqin, Wang Haichao, and Guohai Situ. 2019a. Image transmission through scattering media using ptychographic iterative engine. *Appl. Sci.* 9, 5 (2019), 849.

Long Li, Quan Li, Shuai Sun, Hui-Zu Lin, Wei-Tao Liu, and Ping-Xing Chen. 2018a. Imaging through scattering layers exceeding memory effect range with spatial-correlation-achieved point-spread-function. *Opt. Lett.* 43, 8 (2018), 1670–1673.

Lingling Li, Xinyu Wen, Rui Song, Ji-Tong Jiang, Hui-Li Zhang, Xiao-Bo Liu, and Liqing Wei. 2019b. Imaging corereography using ptychography. *Appl. Sci.* 9, 20 (2019), 4377.

Yunzhe Li, Yujia Xue, and Lei Tian. 2018b. Deep speckle correlation: A deep learning approach toward scalable imaging through scattering media. *Optica* 5, 10 (2018), 1181–1190.

Meihua Liao, Dajiang Lu, Wenqi He, Giancarlo Pedrini, Wolfgang Osten, and Xiang Peng. 2019. Improving reconstruction of speckle correlation imaging by using a modified phase retrieval algorithm with the number of nonzero-pixels constraint. *Appl. Opt.* 58, 2 (2019), 473–478.

Chao Liu, Akash K. Maity, Artur W. Dubrawski, Ashutosh Sabharwal, and Srinivasa G. Narasimhan. 2020. High resolution diffuse optical tomography using short range indirect subsurface imaging. *ICCP* (2020).

Christopher A. Metzler, Felix Heide, Prasanna Rangarajan, Muralidhar Madabhushi Balaji, Aparna Viswanath, Ashok Veeraraghavan, and Richard G. Baraniuk. 2020. Deep-inverse corereography: Towards real-time high-resolution non-line-of-sight imaging. *Optica* 7, 1 (2020), 63–71.

Allard P. Mosk, Ad Lagendijk, Geoffroy Lerosey, and Mathias Fink. 2013. Controlling waves in space and time for imaging and focusing in complex media. *Nat. Photon.* 6, 5 (2013), 283.

Micha Nixon, Ori Katz, Eran Small, Yaron Bromberg, Asher A. Friesem, Yaron Silberberg, and Nir Davidson. 2013. Real-time wavefront shaping through scattering media by all-optical feedback. *Nat. Photon.* 7, 11 (2013), 919–924.

Gerwin Osnabrugge, Roarke Horstmeyer, Ioannis N. Papadopoulos, Benjamin Judkewitz, and Ivo M. Vellekoop. 2017. Generalized optical memory effect. *Optica* 4, 8 (2017), 886–892.

Ioannis Papadopoulos, Jean-Sébastien Jouhanneau, James Poulet, and Benjamin Judkewitz. 2016. Scattering compensation by focus scanning holographic aberration probing (F-SHARP). *Nat. Photon.* 11, 2 (2016), 116–123.

David J. Pine, David A. Weitz, Paul M. Chaikin, and Eric Herbolzheimer. 1988. Diffusing wave spectroscopy. *Phys. Rev. Lett.* 60, 12:1134 (1988).

John M. Rodenburg, A. C. Hurst, Anthony G. Cullis, Barry R. Dobson, Franz Pfeiffer, Oliver Bunk, Christian David, Konstantinos Jefimovs, and I. Johnson. 2007. Hard-X-ray lensless imaging of extended objects. *Phys. Rev. Lett.* 98, 3 (2007).

Markus Rueckel, Julia A. Mack-Bucher, and Winfried Denk. 2006. Adaptive wavefront correction in two-photon microscopy using coherence-gated wavefront sensing. *PNAS* 103, 46 (2006), 17137–17142.

Guy Satat, Barmak Heshmat, Christopher Barsi, Dan Raviv, Ou Chen, Moungi Bawendi, and Ramesh Raskar. 2015. Locating and classifying fluorescent tags behind turbid layers using time-resolved inversion. *Nat. Commun.* 6, 6796 (2015).

Guy Satat, Barmak Heshmat, Dan Raviv, and Ramesh Raskar. 2016. All photons imaging through volumetric scattering. *Sci. Rep.* 6, 33946 (2016).

Guy Satat, Matthew Tancik, Otkrist Gupta, Barmak Heshmat, and Ramesh Raskar. 2017. Object classification through scattering media with deep learning on time resolved measurement. *Opt. Express* 25, 15 (2017), 17466–17479.

Sam Schott, Jacopo Bertolotti, Jean-Francois Léger, Laurent Bourdieu, and Sylvain Gigan. 2015. Characterization of the angular memory effect of scattered light in biological tissues. *Opt. Express* 23, 10 (2015), 13505–13516.

Noam Shekel and Ori Katz. 2020. Using fiber-bending-generated speckles for improved working distance and background rejection in lensless micro-endoscopy. *Opt. Lett.* 45, 15 (2020), 4288–4291.

Brandon M. Smith, Pratham Desai, Vishal Agarwal, and Mohit Gupta. 2017. CoLux: Multi-object 3D micro-motion analysis using speckle imaging. *ACM Trans. Graph.* 36, 4 (2017), 1–12.

Brandon M. Smith, Matthew O’Toole, and Mohit Gupta. 2018. Tracking multiple objects outside the line of sight using speckle imaging. *CVPR* (2018), 6258–6266.

Kevin T. Takasaki and Jason W. Fleischer. 2014. Phase-space measurement for depth-resolved memory-effect imaging. *Opt. Express* 22, 25 (2014), 31426–31433.

Valery Tuchin. 2000. *Tissue Optics: Light Scattering Methods and Instruments for Medical Diagnosis*. SPIE.

Elbert G. van Putten, Duygu Albulut, Jacopo Bertolotti, Willem. L. Vos, Ad Lagendijk, and Allard. P. Mosk. 2011. Scattering lens resolves sub-100 nm structures with visible light. *Phys. Rev. Lett.* 106, 19 (2011).

Ivo M. Vellekoop and Christof M. Aegerter. 2010. Scattered light fluorescence microscopy: Imaging through turbid layers. *Opt. Lett.* 35, 8 (2010), 1245–1247.

Ivo M. Vellekoop, Meng Cui, and Changhuei Yang. 2012. Digital optical phase conjugation of fluorescence in turbid tissue. *Appl. Phys. Lett.* 101, 8 (2012).

Ivo M. Vellekoop, Aart Lagendijk, and Allard P. Mosk. 2010. Exploiting disorder for perfect focusing. *Nat. Photon.* 4, 5 (2010), 320–322.

Ivo M. Vellekoop and Allard P. Mosk. 2007. Focusing coherent light through opaque strongly scattering media. *Opt. Lett.* 32, 16 (2007), 2309–2311.

Aparna Viswanath, Prasanna Rangarajan, Duncan Macfarlane, and Marc Christensen. 2018. Indirect imaging using corereography. *COSI CM2E.3* (2018).

Xiaoyu Wang, Xin Jin, and Junqi Li. 2020. Blind position detection for large field-of-view scattering imaging. *Photon. Res.* 8, 6 (2020), 920–928.

Xiaoyu Wang, Xin Jin, Junqi Li, Xiaocong Lian, Xiangyang Ji, and Qionghai Dai. 2019. Prior-information-free single-shot scattering imaging beyond the memory effect. *Opt. Lett.* 44, 6 (2019), 1423–1426.

Tengfei Wu, Jonathan Dong, and Sylvain Gigan. 2020. Non-invasive single-shot recovery of a point-spread function of a memory effect based scattering imaging system. *Opt. Lett.* 45, 19 (2020), 5397–5400.

Tengfei Wu, Jonathan Dong, Xiaopeng Shao, and Sylvain Gigan. 2017. Imaging through a thin scattering layer and jointly retrieving the point-spread-function using phase-diversity. *Opt. Express* 25, 22 (2017), 27182–27194.

Shumian Xin, Sotiris Nousias, Kyros Kutulakos, Aswin C. Sankaranarayanan, Sriniwas Narasimhan, and Ioannis Gkioulekas. 2019. A theory of fermat paths for non-line-of-sight shape reconstruction. *CVPR* (2019), 6800–6809.

Zahid Yaqoob, Demetri Psaltis, Michael Feld, and Changhuei Yang. 2008. Optical phase conjugation for turbidity suppression in biological samples. *Nat. Photon.* 2, 2 (2008), 110–115.

M. Zhou, A. Pan, R. Li, Y. Liang, J. Min, T. Peng, C. Bai, and B. Yao. 2020. Retrieval of non-sparse object through scattering media beyond the memory effect. *J. Opt.* 22, 8 (2020).

Received October 2020; revised January 2021; accepted January 2021