

Interferometric Transmission Probing with Coded Mutual Intensity

ALANKAR KOTWAL, Carnegie Mellon University

ANAT LEVIN, Technion

IOANNIS GKIOULEKAS, Carnegie Mellon University

We introduce a new interferometric imaging methodology that we term interferometry with coded mutual intensity, which allows selectively imaging photon paths based on attributes such as their length and endpoints. At the core of our methodology is a new technical result that shows that manipulating the spatial coherence properties of the light source used in an interferometric system is equivalent, through a Fourier transform, to implementing light path probing patterns. These patterns can be applied to either the coherent transmission matrix, or the incoherent light transport matrix describing the propagation of light in a scene. We test our theory by building a prototype inspired by the Michelson interferometer, extended to allow for programmable phase and amplitude modulation of the illumination injected in the interferometer. We use our prototype to perform experiments such as visualizing complex fields, capturing direct and global transport components, acquiring light transport matrices, and performing anisotropic descattering, both in steady-state imaging and, by combining our technique with optical coherence tomography, in transient imaging.

CCS Concepts: • **Computing methodologies** → **Computational photography**.

Additional Key Words and Phrases: spatial coherence, mutual intensity, interferometry, light transport matrix, transmission matrix

ACM Reference Format:

Alankar Kotwal, Anat Levin, and Ioannis Gkioulekas. 2020. Interferometric Transmission Probing with Coded Mutual Intensity. *ACM Trans. Graph.* 39, 4, Article 74 (July 2020), 16 pages. <https://doi.org/10.1145/3386569.3392384>

1 INTRODUCTION

Light propagation is an inherently multi-path phenomenon: When we look at our surroundings, we observe light that has interacted with one or multiple objects, either by reflecting on their surfaces, or by scattering in their interior. Imaging systems typically accumulate contributions from photons traveling along all of these paths, indiscriminately of characteristics such as the paths' origins and lengths. This accumulation process confounds the information that is available in imaging measurements about scene properties of interest, such as the shape and material of objects of interest.

Computational light transport techniques attempt to overcome this confounding effect, by measuring only light that has traveled along specific subsets of all the possible paths in a scene. These

subsets can be specified based on characteristics such as the endpoints and length of the paths, or combinations thereof. There is a growing number of systems that can perform this type of *selective probing* of the many light components that make up an image, including systems based on projector-camera combinations, continuous-wave amplitude-modulated sensors, streak cameras, single-photon avalanche diodes, as well as interferometry.

Our focus is on the latter type of systems: Interferometric systems operate by simultaneously measuring two light waves that originated at the same light source, and have traveled along different paths in an optical system. When superimposed on an optical sensor, the two waves will produce some measurable *interference*. Depending on the wave-optics properties of the original illumination, only light paths that satisfy certain characteristics will contribute to the interference. Therefore, by controlling the coherence properties of the source the waves originate from, and then measuring their interference after propagation, it is possible to isolate contributions from only specific light paths. This type of interferometric probing has been demonstrated in the past [Gkioulekas et al. 2015], using a system similar to optical coherence tomography [Huang et al. 1991]. However, these previously systems are severely limited in terms of the types of probing they can perform, compared to probing capabilities possible using other imaging technologies.

Our goal on this paper is to significantly expand the probing capabilities that can be realized using interferometry. To this end, we develop a new imaging technique that we term *interferometry with coded mutual intensity*. Our technique is based on a setup similar to the classical Michelson interferometer, augmented with optical components for amplitude and phase modulation. These components enable programmatic control of the spatial coherence properties of the illumination injected in the setup.

We perform a detailed theoretical analysis of our technique, and explain how the underlying wave-optics models relate to the incoherent models of light propagation typically used in computer graphics. Through this analysis, we show that our technique provides several probing capabilities: First, it enables probing the coherent *transmission matrix* of a scene, using arbitrary convolutional probing patterns. Second, it allows probing the incoherent *light transport matrix* of a scene, using probing patterns that are challenging to implement with alternative techniques. Third, it facilitates incorporating these probing capabilities within other interferometric techniques, for example optical coherence tomography.

Our paper begins with background on the Michelson interferometer and the notions of spatial and temporal coherence. We then use this background to relate interferometry to measurements of the transmission matrix characterizing coherent propagation of light. In particular, we explain how, by modulating the spatial coherence properties of the illumination used for interferometry, we can control which elements of the transmission matrix contribute to image

Authors' addresses: Alankar Kotwal, Carnegie Mellon University; Anat Levin, Technion; Ioannis Gkioulekas, Carnegie Mellon University.

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.

© 2020 Association for Computing Machinery.

0730-0301/2020/7-ART74 \$15.00

<https://doi.org/10.1145/3386569.3392384>

measurements. We continue to show how these interferometric measurements can be converted to measurements of the light transport matrix, characterizing radiometric propagation of light in a scene. Then, we design and build a physical prototype implementing our interferometric technique. Finally, we use this prototype to perform experiments such as visualizing complex fields, capturing direct and global transport components, acquiring light transport matrices, and performing anisotropic descattering, both in steady-state and transient imaging. We provide implementation details in the supplement, and our code and data in the project website [Kotwal et al. 2020], to facilitate reproducibility and follow-up research.

1.1 Why use interferometry with coded mutual intensity?

Our paper focuses on closing the gap between the probing capabilities that can be implemented using interferometric systems, versus those that are already possible using projector-camera systems [O'Toole et al. 2012]. This is important for making probing capabilities available in microscopic imaging applications, where projector-camera systems can be difficult to deploy and interferometric setups are widely-employed. For example, as we demonstrate in Section 7, our paper makes it possible to integrate very general probing functionalities with optical coherence tomography systems [Huang et al. 1991], commonly-used for tissue imaging and retinal imaging application. Another application includes interferometric shape-measurement devices [Li et al. 2018, 2017; Maeda et al. 2018], which are common in industry and fabrication where it is necessary to obtain measurements at micrometer scales: As these devices already use interferometric measurements, it is relatively straightforward to extend them to additionally incorporate probing using our technique. By contrast, it is unclear how one could implement probing in these systems using a projector-camera technique. In Section 8, we compare in more detail interferometric and projector-camera systems for probing.

Our paper builds upon the foundation laid down by Gkioulekas et al. [Gkioulekas et al. 2015], who first showed how interferometry can be used to implement probing of different light transport components. However, the technique by Gkioulekas et al. [Gkioulekas et al. 2015] is severely constrained in terms of the types of probing that are practically-implementable, with the paper's focus being almost exclusively on *diagonal probing* (corresponding to direct-only imaging). Other probing types require designing and fabricating complex mirror configurations, and for general probing types it is unclear whether such configurations even exist. By contrast, our technique makes it possible to implement arbitrary convolutional probing types. As we demonstrate in Section 7, our technique enables using interferometric systems to implement probing types that are particularly relevant for critical applications such as tissue imaging. These include probing with non-binary anisotropic patterns for descattering (Figure 8), as well as probing off-diagonal light transport components (Figure 9) that have been shown to be of great value for tomography (for example, using short-range-indirect imaging [Kubo et al. 2018], diffuse optical tomography [Liu et al. 2020], or inverse rendering [Gkioulekas et al. 2016]). Additionally, our technique allows implementing different probing types in a programmable way, without the need for any hardware changes

when changing from one probing type to another. This opens up interesting new possibilities, such as data-driven learning of optimal probing matrices, optical matched filtering for pattern detection, or adaptive probing through closed-loop control of the interferometric system. Overall, our paper makes the full potential of the probing framework, originally introduced by [Gkioulekas et al. 2015], available for the first time to a large range of critical applications in medicine and industry that rely on interferometry.

Last but not least, our paper introduces for the first time the ability to perform coherent probing of transmission matrices, in addition to probing incoherent light transport matrices. Transmission matrix formulations have emerged for several applications, including focusing and imaging through scattering layers [Popoff et al. 2010a; Vellekoop and Mosk 2007], as well as using their memory-effect structure for fluorescent imaging [Judkewitz et al. 2015; Osnabrugge et al. 2017]. The ability to probe these transmission matrices opens up several directions for future investigation: For example, it seems plausible that probing only the parts of the transmission matrix that contribute to the memory effect could improve descattering. Alternatively, probing may allow for more robust focusing inside tissue. We hope that, by introducing coherent probing, our paper will motivate the investigation of different components of the transmission matrix and the discovery of useful coherent probing patterns.

2 RELATED WORK

Computational light transport. The last two decades have seen the development of a large number of techniques for acquiring, decomposing, and selectively capturing different components of the light transport in an arbitrary scene. Our focus here will be on techniques that can achieve this without acquiring the entire light transport matrix of the scene [O'Toole and Kutulakos 2010; Peers et al. 2009; Sen et al. 2005; Wang et al. 2009]. We can broadly classify these techniques into three categories, based on the criteria they use to decompose light transport. First, *bounce decomposition* techniques capture only transport due to paths that bounce a specific number of times in a scene. Typically, this involves separating a direct component (single-bounce paths) from an indirect component (paths that bounce two or more times) [Gupta et al. 2011; Nayar et al. 2006; Reddy et al. 2012]. In Lambertian scenes, it is also possible to separate transport into discrete bounce components (only single-bounce, only two bounces, and so on) [Seitz et al. 2005]. These techniques are typically implemented using projector-camera systems.

Second, *transient imaging* techniques capture only light that has traveled along paths of a specific length (or equivalently, light that has a specific time-of-flight) within the scene [Jarabo et al. 2017]. This capability has been demonstrated using various sensing technologies, including ultrafast photodiodes [Kirmani et al. 2011], streak cameras [Velten et al. 2013; Wu et al. 2014a,b], continuous-wave time-of-flight cameras [Heide et al. 2013; Kadambi et al. 2013; Peters et al. 2015], single-photon avalanche diodes [Garipey et al. 2015; O'Toole et al. 2017], and interferometry [Abramson 1983]. Third, *probing* techniques isolate light following paths whose end-points satisfy specific correspondences between a two-dimensional source and a two-dimensional sensor [O'Toole et al. 2012]. These can be epipolar [O'Toole et al. 2014b], disparity [O'Toole et al. 2012;

Wang et al. 2018], or plane-ray [Kubo et al. 2018; Liu et al. 2020] correspondences, as well as their logical complements [O'Toole et al. 2015]. Finally, there exist techniques that implement hybrids between different types of decompositions, for example by combining probing and transient imaging [Achar et al. 2017; Gkioulekas et al. 2015; O'Toole et al. 2014a; Wang et al. 2018].

Optical interferometry. Interferometric techniques broadly operate by exploiting the interference between one or more light waves [Hariharan 2003]. These techniques have a long history in optics and other disciplines: for example, optical coherence tomography [Huang et al. 1991] is commonly employed for tissue and retinal imaging. More recently, an increased number of interferometric techniques have been introduced in the computer vision and graphics literature. For example, interferometric techniques have been used for refocusing [Cossairt et al. 2014], and high-resolution depth sensing [Li et al. 2018, 2017; Maeda et al. 2018]. Most closely related to our paper is the work of Gkioulekas et al. [2015], who use interferometry to perform transient imaging, optionally combined with different types of probing. Our technique significantly extends these capabilities, by using interferometry to implement arbitrary convolutional forms of probing, either on their own or jointly with transient imaging, in a programmable and light-efficient manner.

At the core of our technique is the ability to programmatically control the spatial coherence properties, and in particular the mutual intensity function [Zernike 1938], of the illumination injected in the two arms of a Michelson interferometer. This relates our technique to the extensive literature on visualizing, analyzing, and synthesizing the spatial coherence properties of test beams using shearing and Sagnac interferometry [Carter 1977; Koivurova et al. 2017; Mendlovic et al. 1999, 1998; Naik et al. 2009; Pan et al. 2019; Wax and Thomas 1996]. Instead of analyzing test beams, we use interferometry with illumination of controlled spatial coherence properties to analyze light transmission in arbitrary scenes.

Coherent imaging. Imaging with coherent effects has received increased attention within computer vision and graphics in recent years. In particular, techniques using speckle properties such as the memory effect [Goodman 2007] have been used for motion tracking [Jo et al. 2015; Smith et al. 2017], non-line-of-sight of imaging [Smith et al. 2018], and imaging through scattering layers [Chang and Wetzstein 2018], from just a few images under coherent illumination. Other techniques use exhaustive or partial measurements of the complex *transmission matrix* [Popoff et al. 2010a], in order to perform tasks such as focusing through scattering [Metzler et al. 2017; Sharma et al. 2019]. Many of these techniques have long prior history in optics and biomedical imaging [Bertolotti et al. 2012; Katz et al. 2014; Popoff et al. 2010b; Vellekoop et al. 2012; Vellekoop and Mosk 2007]. In particular, the structure of the transmission matrix for different types of scenes has been studied extensively using analytical approximations to coherent light propagation [Judkewitz et al. 2015; Osnabrugge et al. 2017], physically-accurate simulations [Bar et al. 2019], as well as measurements that allow directly visualizing both its spatial-domain and frequency-domain forms [Waller et al. 2012]. Both the memory effect and the transmission matrix are highly-relevant to our work: We rely on properties of the former for some of our derivations, and we show how our technique can




Fig. 1. **The Michelson interferometer.** (a) We use two different light sources to provide illumination to a Michelson interferometer. (b) A beam-splitter separates the incident illumination into two beams. After reflection at the reference arm (mirror) and target arm (retroreflector), the two beams recombine and are measured by the camera. (c) We can use the captured images to visualize the spatial coherence properties of the two light sources.

be used to perform probing of transmission matrices analogously to the probing of incoherent light transport matrices.

3 BACKGROUND ON INTERFEROMETRY AND SPATIAL COHERENCE

We begin by introducing some concepts from wave optics, necessary for understanding interferometry and spatial coherence. Our discussion largely follows Goodman [2000], Levin et al. [2013], and Gkioulekas et al. [2015]. The analysis in this section will serve as a warm-up for our technical results in Section 4.

Michelson interferometer. Throughout the paper, we use variants of the classical free-space full-field Michelson interferometer setup, with a two-dimensional illumination source plane and a two-dimensional sensor. This setup, shown in Figure 1(b), uses a beamsplitter to separate a light wave emitted from a source into two parts of (generally) equal intensity. One part propagates towards the so-called *target arm* of the interferometer, which contains the scene to be imaged. The other part propagates towards the *reference arm* of the interferometer, which contains some fixed, application-specific, optical configuration: most commonly, this is a planar mirror, but we will discuss later more general reference arm configurations. After reflection, the waves from the two arms are recombined by the beamsplitter, and finally imaged by a camera.

In this section, we focus on analyzing a simple interferometric experiment, which we use to provide insights into our main results in Section 4. We place a retroreflector (in the form of a corner reflector mirror) in the target arm of the Michelson interferometer. We inject light into this setup through a light source consisting of a lens and two types of (approximately) *monochromatic* emitters: a laser diode, and a gas-discharge lamp (Figure 1(a)). Both emitters are placed at the focal length of the lens, which collimates their output into a beam that propagates towards the beamsplitter. We use a camera to capture images when the interferometer receives light from one or the other source, and show the results in Figure 1(c).

We observe that both images show a fringe pattern, which is the result of *interference* between the two waves arriving at the camera, from the reference and target arm. In the laser case, this fringe pattern is present in the entire frame, whereas in the lamp case, the pattern is limited to a small area at the center of the image.




Fig. 2. **Notation and coordinate system.** We use the same coordinate system throughout Sections 3 and 4. (a) The parts of the Michelson interferometer corresponding to the target and reference arms. (b-e) The target and reference arms, shown unfolded in the same coordinate system. (b) and (c) show the target arm, with lighting and camera respectively. (d) and (e) show the reference arm, with lighting and camera respectively.

Qualitatively, the reason for this difference is the area of emission of the two sources: The laser diode is effectively a point emitter, having a circular area of emission with diameter approximately $5\text{ }\mu\text{m}$. By contrast, the gas-discharge lamp has a circular area of emission that is 10,000 times larger, with a diameter of approximately $500\text{ }\mu\text{m}$. The emission area discrepancy results in illumination of different *spatial coherence* properties, as we discuss in detail next.

3.1 Spatial coherence

Problem setup. To simplify notation, we derive our results in two dimensions, with the extension to the three-dimensional case being straightforward. We use an (x, z) coordinate system, shown in Figure 2, where z is the optical axis of the illumination lens and the camera. We restrict our derivations to the scalar wave theory, which is sufficient for describing spatial coherence effects. Additionally, we focus on the case of purely monochromatic light of wavelength λ (though note the discussions of temporal coherence at the end of this section, and of optical coherence tomography in Section 5). With these assumptions, we can describe electromagnetic fields as complex phasors that are independent of time, and arise as solutions of the Helmholtz equation [Goodman 2005, Section 3.3.1].

Incoherent area emitters. We denote by f the focal length of the illumination lens, and use s to parameterize its focal plane. When we place an ideal point emitter at some location s on the focal plane, as shown in Figure 1(a), the lens produces a *plane wave* propagating at an angle θ relative to the optical axis z . We can express the corresponding electromagnetic field at every location (x, z) as:

$$u^\theta(x, z) = \exp(-ik(x \sin(\theta) + z \cos(\theta))), \quad (1)$$

where $\kappa = 2\pi/\lambda$ is the *wavenumber*, and the direction θ satisfies $\sin(\theta) = s/f$. We use the *paraxial approximation*, requiring that θ is small, and therefore $\sin(\theta) = s/f \approx \theta$ and $\cos(\theta) \approx 1$. We can thus rewrite the plane wave as:

$$u^\theta(x, z) = \exp(-ik(x\theta + z)). \quad (2)$$

We now consider an area emitter of a non-zero emission area, characterized by the *complex emission* $E(s)$: For every point s on the focal plane, $|E(s)|^2$ and $\angle E(s)$ are, respectively, the energy and the phase of the emitter at that point. From Equation (2), $|E(s)|$ and $\angle E(s)$ are also the amplitude and phase of the plane wave at direction $\theta = s/f$; considering the linear relationship between s and θ , we abuse notation and write the complex emission $E(\theta)$ as a function

of θ , with the scaling by the focal length f being implicit. A laser diode of practically infinitesimal width and a uniform area emitter of width W , when centered on the focal plane, have respectively $|E_{\text{laser}}(\theta)|^2 \propto \delta(\theta)$ and $|E_{\text{area}}(\theta)|^2 \propto \text{rect}\left(\frac{\theta}{W}\right)$, where δ and rect are the Dirac delta and rectangular functions, respectively. For emitters with non-infinitesimal area, we assume that each point θ with non-zero emission power $|E(\theta)|^2$ is *statistically independent* with respect of every other such point. We precisely characterize the notion of statistical independence later in this section (Equation (7)). We refer to such emitters as *incoherent area emitters*.

With these assumptions, we can express the illumination produced by an incoherent area emitter with complex emission $E(\theta)$ as a collection of plane waves:¹

$$u(x, z) = \int_{\theta} E(\theta) u^\theta(x, z) d\theta. \quad (3)$$

Scattered and reference fields. We use the source described in Equation (3) as illumination in the Michelson interferometer of Figure 1(b). The beamsplitter creates two copies of the *incident field* $u_i(x, z) \equiv u(x, z)$. After interacting with the target and reference arms, these produce, respectively, the *scattered field* $u_s(x, z)$ and *reference field* $u_r(x, z)$, which we express analogously to Equation (3):

$$u_s(x, z) = \int_{\theta} E(\theta) u_s^\theta(x, z) d\theta, \quad (4)$$

$$u_r(x, z) = \int_{\theta} E(\theta) u_r^\theta(x, z) d\theta. \quad (5)$$

where $u_s^\theta(x, z)$ is the field that the target arm would produce if the incident field consisted of a single plane wave $u^\theta(x, z)$; and likewise for $u_r^\theta(x, z)$ and the reference arm. We note that $u_s^\theta(x, z)$ and $u_r^\theta(x, z)$ are not necessarily plane waves, and their exact form depend on the two arms. We analyze a specific example later in this section, and discuss the general case in Section 4.

Interference. We now consider an intensity sensor focused at location (x, z) . The sensor will capture a measurement $I(x, z)$ due to the superposition of the scattered and reference fields, combined by the beamsplitter. We can express this measurement as:²

$$I(x, z) = \left| \int_{\theta} E(\theta) u_s^\theta(x, z) + E(\theta) u_r^\theta(x, z) d\theta \right|^2. \quad (6)$$

The *statistical independence* assumption we introduced earlier can now be concretely defined as follows: When superimposing the scattered and reference fields, we first *coherently* sum the complex scattered and reference fields for each θ separately, and then we *incoherently* integrate the resulting intensities for all θ values. In practice, this statistical independence property is satisfied only when the plane waves u^θ that make up the incident illumination are the result of thermal emission from different points on an area emitter (for example, gas-discharge lamp or LED). However, such waves are typically not monochromatic. In Section 5, we show how to use

¹This equation assumes that the lens has an infinite aperture, and we can ignore diffraction and vignetting artifacts.

²We note that, typically, the intensity of a wave is defined as the *temporal average* of its squared amplitude [Goodman 2000, Equation (4.2-5)]. Because we use time-independent phasors to represent monochromatic waves, we can *equivalently* express intensity while omitting temporal averaging. The same holds for our expressions for correlation and mutual intensity later in this section.

temporal multiplexing to realize monochromatic illumination that satisfies this statistical independence property.

With the assumption of statistical independence, we can exchange the order of integration and squaring in Equation (6) to obtain:

$$I(x, z) = \int_{\theta} \left| E(\theta) u_s^{\theta}(x, z) + E(\theta) u_r^{\theta}(x, z) \right|^2 d\theta \quad (7)$$

$$= \int_{\theta} \left| E(\theta) u_s^{\theta}(x, z) \right|^2 d\theta + \int_{\theta} \left| E(\theta) u_r^{\theta}(x, z) \right|^2 d\theta + 2 \operatorname{Re} \left\{ \int_{\theta} \left(E(\theta) u_s^{\theta}(x, z) \right)^* \cdot \left(E(\theta) u_r^{\theta}(x, z) \right) d\theta \right\} \quad (8)$$

$$= \underbrace{I_s(x, z) + I_r(x, z)}_{\text{DC intensity}} + \underbrace{2 \operatorname{Re} \{ \operatorname{corr}(x, z) \}}_{\text{interference}}, \quad (9)$$

where we have used:

$$I_s(x, z) \equiv \int_{\theta} \left| E(\theta) u_s^{\theta}(x, z) \right|^2 d\theta, \quad (10)$$

$$I_r(x, z) \equiv \int_{\theta} \left| E(\theta) u_r^{\theta}(x, z) \right|^2 d\theta, \quad (11)$$

$$\operatorname{corr}(x, z) \equiv \int_{\theta} \left(E(\theta) u_s^{\theta}(x, z) \right)^* \cdot \left(E(\theta) u_r^{\theta}(x, z) \right) d\theta. \quad (12)$$

We observe that the sensor measurement in Equation (9) has two components: First, the *DC intensity* component is the sum of the intensity measurements $I_s(x, z)$ and $I_r(x, z)$ the sensor would capture if it were observing the target and reference arms separately. Second, the *interference* component is due to the *correlation* $\operatorname{corr}(x, z)$ of the scattered and reference fields. Our focus throughout the paper will be on characterizing this correlation, for different types of incident, scattered, and reference fields.

Equation (9) additionally suggests a straightforward way for isolating the interference component: We can use the camera to capture two additional images, one with the reference arm blocked (e.g., using a neutral density filter of very high optical density), and one with the target arm blocked. These two images will be equal to $I_s(x, z)$ and $I_r(x, z)$, respectively. By subtracting them from the superposition image $I(x, z)$, we are left with an estimate (up to noise) of the interference component. In Section 6, we describe in more detail how we can estimate the interference and correlation components.

Spatial coherence. Finally, we will derive the correlation for the specific target and reference arm configurations of Figure 1(b). For this, we first need to determine the fields $u_s^{\theta}(x, z)$ and $u_r^{\theta}(x, z)$ for all values of θ . A retroreflector placed at the in-focus plane and centered at $x = 0$ has the effect of flipping the x coordinate of an incoming plane wave with respect to the origin:

$$u_s^{\theta}(x, z) = u^{\theta}(-x, z). \quad (13)$$

A planar mirror, orthogonal to the optical axis and at distance τ from the in-focus plane (Figure 1(a)), introduces a phase delay:

$$u_r^{\theta}(x, z) = u^{\theta}(x, z) \exp(-ik\tau). \quad (14)$$

Then, we can prove the following proposition.

PROPOSITION 1. *Given a light source with complex emission $E(\theta)$, a reference arm containing a planar mirror, and a target arm containing a centered retroreflector, the correlation of Equation (12) equals:*

$$\operatorname{corr}(x, z) = \exp(-ik\tau) \mathcal{S}^c(2x), \quad (15)$$

where the mutual intensity kernel $\mathcal{S}^c$ is defined as:

$$\mathcal{S}^c(x) \equiv \mathcal{F} \{ |E(\theta)|^2 \} [x], \quad (16)$$

and $\mathcal{F}$ is the Fourier transform.

PROOF. Given Equations (13) and (14), we can write:

$$u_s(x, z) = \int_{\theta} E(\theta) \exp(-ik(-x\theta + z)) d\theta, \quad (17)$$

$$u_r(x, z) = \exp(-ik\tau) \int_{\theta} E(\theta) \exp(-ik(x\theta + z)) d\theta. \quad (18)$$

From the statistical independence property of different θ ,

$$\operatorname{corr}(x, z) = \exp(-ik\tau) \underbrace{\int_{\theta} |E(\theta)|^2 \exp(ik(2x\theta)) d\theta}_{\equiv \mathcal{S}^c(2x)}, \quad (19)$$

which concludes the proof. $\square$

The name *mutual intensity kernel* is motivated as follows: Given two points (x, z) and (x', z) at the in-focus plane, $\mathcal{S}^c(|x - x'|)$ equals the *mutual intensity* of the incident field at those points,

$$\mathcal{S}^c(|x - x'|) = u_i(x, z)^* u_i(x', z) \equiv J((x, z), (x', z)), \quad (20)$$

where J is the *mutual intensity function* of the incident field [Goodman 2000][Equation (5.2-32)]. Effectively, the mutual intensity kernel describes how the mutual intensity of the incident illumination varies as a function of the distance between points on the same plane. The mutual intensity kernel $\mathcal{S}^c$ is a function exclusively of the complex emission, and thus characteristic of the source; it describes the *spatial coherence* properties of the illumination generated by the source [Goodman 2000][Section 5.2.2]. For an ideal laser diode, $\mathcal{S}_{\text{laser}}^c(x) \propto 1$, and we refer to such illumination as *spatially coherent*. By contrast, for a uniform incoherent area emitter of width W , $\mathcal{S}_{\text{area}}^c(x) \propto \operatorname{sinc}(W \cdot x)$. As the width W of the emitter increases, the main lobe of the sinc function in the mutual intensity kernel decreases. Conversely, at the limit $W \rightarrow 0$, the mutual intensity kernel $\mathcal{S}_{\text{area}}^c$ becomes identical to that of the ideal laser diode.

Visualizing the mutual intensity kernel. Concluding our analysis of the experiment in Figure 1, we can use Equation (15) to interpret the images captured under the two light sources. In particular, the correlation captured with a retroreflector on the reference arm is exactly equal to (a scaled version of) the mutual intensity kernel $\mathcal{S}^c$: In the case of the ideal laser diode, the mutual intensity kernel is constant everywhere, resulting in a frame that is fully covered with fringes due to interference. In the case of the incoherent area emitter, the mutual intensity kernel is constrained primarily within the main lobe of a sinc function, with a width inversely proportional to that of the light source. In the corresponding captured image, we see that the fringes due to interference have a limited spatial extent, corresponding to the main lobe of the mutual intensity kernel. We note that the setup of Figure 1 provides us with an alternative to shearing interferometry techniques for visualizing the spatial coherence properties of a source [Mendlovic et al. 1998].

A different way to interpret the fringes on the retroreflector is by adopting a geometric optics viewpoint. We consider light paths starting at a plane immediately in front of the lens, and propagating

in the interferometer. From the paraxial approximation, the intensity the sensor measures at location x will be due to two paths, both traveling parallel to the optical axis, as shown in Figure 1(b). The first path arrives at x after reflection on the planar mirror of the reference arm, and from Equation (14) its origin on the lens is also at x . The second path arrives at x after reflection on the retroreflector of the target arm, and from Equation (13) its origin on the lens is at $-x$. Therefore, the origins of the two paths are separated by a distance $d = 2x$. The mutual intensity kernel determines the correlation of the two paths as a function of this separation d : In the case of the laser diode, two paths will always interfere, independently of d . In the case of the incoherent area emitter, they will only produce significant correlation if $d < W$. In the rest of the paper, we generalize this idea, and show how to selectively decide which paths will contribute to a sensor measurement based on their origins—that is, to *probe* light paths analogously to O’Toole et al. [2012].

Temporal coherence. Before we conclude this section, we discuss our assumption that the illumination we are using is purely monochromatic. When this is not the case, we need to consider the *temporal coherence* properties of the light source. As discussed by Gkioulekas et al. [2015], using a non-monochromatic source in the Michelson interferometer experiment of Figure 1 means that the correlation we observe in the captured images will additionally be a function of the difference τ between the distances of the reference and target arms from the beamsplitter (Figure 2(e)):

$$\text{corr}(x, z) = \exp(-i\kappa\tau) \mathcal{G}^c(\tau) S^c(2x), \quad (21)$$

where the *temporal coherence kernel* $\mathcal{G}^c$ is typically a positive real function that monotonically decreases as $|\tau|$ increases. The product $\mathcal{G}^c(\tau) S^c(x)$ characterizes the *mutual coherence* of the incident illumination generated by the light source [Goodman 2000][Equation (5.2-7)], analogously to Equation (20) for mutual intensity. The exact form of $\mathcal{G}^c$ is determined by the spectrum of the light source: when the light source is monochromatic, $\mathcal{G}^c(\tau) \propto 1$ and the light source is called *temporally coherent*; otherwise, the source is *temporally incoherent*. In practice, light sources are characterized by a *temporal coherence length*, which is the maximum separation distance τ within which they can be considered temporally coherent. We assume a temporally coherent source (infinite temporal coherence length) for deriving our theoretical results, and most of our experiments use a source of temporal coherence length significantly longer than the size of our test scenes. However, as we discuss in Section 5, our technique can be combined with temporally incoherent illumination, as the two types of coherence are orthogonal to each other; we demonstrate this experimentally in Section 7.

4 PROBING WITH CODED MUTUAL INTENSITY

In this section, we develop our main technical results, which take the form of two propositions: In Proposition 2, we show how to relate interferometric measurements of some arbitrary scene, captured as correlation measurements from a Michelson interferometer setup, to measurements of the scene’s *transmission matrix* combined with a *coherent probing matrix*. In Proposition 4, we additionally show that the squared amplitude of this correlation is approximately equal to measurements of the same scene’s radiometric *light transport*




Fig. 3. **Interferometry with coded mutual intensity.** (a) We consider a generalized version of the Michelson interferometer, where the illumination and reference arm are modified to perform, respectively, amplitude and phase modulation in the Fourier domain. (b) This allows us to implement probing of coherent transmission and incoherent light transport matrices.

matrix, combined with an associated *incoherent probing matrix*. We use *probing* in the same way as O’Toole et al. [2012], to refer to element-wise weights applied on the transmission or light transport matrix before imaging. Figure 3 summarizes these contributions. Finally, in the second part of Proposition 2, we show how to implement arbitrary convolutional probing matrices, by using appropriate amplitude and phase modulation of an incoherent area emitter.

4.1 Probing the transmission matrix

Throughout this section, we analyze a more general version of the Michelson interferometer setup, shown in Figure 3. We assume that we capture images using a camera that is focused at the plane $z = z_0$ (Figure 2). To simplify notation, we omit the z coordinate from fields. We begin by describing the main properties of the illumination, reference arm, and target arm of the interferometer.

Illumination. As in the previous section, we assume that the interferometer receives illumination that, after collimation by a lens, can be expressed as the superposition of purely monochromatic, independent plane waves, described by a complex emission function $E(\theta)$. This results in an incident field as in Equation (3):

$$u_i(x) = \int_{\theta} E(\theta) u^{\theta}(x) d\theta. \quad (22)$$

We refer to the squared amplitude of the complex emission function $A(\theta) \equiv |E(\theta)|^2$ as the *amplitude modulation*.

Reference arm. We assume that the reference arm contains an optical configuration that introduces a, potentially θ -dependent, phase delay to each incident plane wave:

$$u_r^{\theta}(x) = u^{\theta}(x) \exp(-i\kappa\Phi(\theta)), \quad (23)$$

We refer to $\Phi(\theta)$ as the *phase modulation* induced by the reference arm. Combining Equations (22)-(23) with Equation (5), we can write:

$$u_r(x) = \int_{\theta} E(\theta) \exp(-i\kappa\Phi(\theta)) u^{\theta}(x) d\theta. \quad (24)$$

Target arm. The target arm consists of a test scene that we are interested in imaging. In order to express the scattered field resulting from this scene, we use the scene’s *coherent Green’s function* $\mathcal{T}^c(x, x')$: This is a scalar complex function that describes the field

that is generated at point x on the plane $z = z_o$, by solving the Helmholtz equation with an impulse boundary condition at point x' also on the plane $z = z_o$. Adopting a geometric optics perspective, we can think of the coherent Green's function as the accumulation of complex contributions from all valid light paths starting at point x' and ending at point x on the plane $z = z_o$; the contribution of each path is the result of free space propagation, and potentially one or more scattering events on intermediate surfaces and volumes. We use the symbol $\mathcal{T}^c$ because of the relationship of this function to the *transmission matrix*, as we discuss later in this section.

Using the Green's function, we can express the scattered field as:

$$u_s(x) = \int_{x'} \mathcal{T}^c(x, x') u_i(x') dx'. \quad (25)$$

Combining Equations (22), (25) with Equation (4), we can write:

$$u_s(x) = \int_{\theta} E(\theta) \underbrace{\int_{x'} \mathcal{T}^c(x, x') u^{\theta}(x') dx'}_{\equiv u_s^{\theta}(x)} d\theta, \quad (26)$$

where u^{θ} is used as in Equation (2).

Interference of reference and scattered fields. We now describe the measurement captured by the camera when imaging this optical setup. As in Equation (9), this image can be decomposed into a DC intensity component, and an interference component due to the correlation between the reference and scattered field (Equation (12)). We can then prove the following proposition.

PROPOSITION 2. *The correlation of Equation (12) for the generalized Michelson interferometer of Figure (3) equals:*

$$\text{corr}(x; \mathcal{T}^c, \mathcal{P}^c) = \int_{\epsilon} \mathcal{T}^c(x, x + \epsilon) \mathcal{P}^c(\epsilon) d\epsilon, \quad (27)$$

where the coherent probing kernel $\mathcal{P}^c$ is defined as:

$$\mathcal{P}^c(x) \equiv \mathcal{F} \{A(\theta) \exp(-ik\Phi(\theta))\} [x], \quad (28)$$

and $\mathcal{F}$ is the Fourier transform.

PROOF. We use Equation (26) and express the scattered field as:

$$u_s(x) = \int_{\theta} E(\theta) \int_{x'} \mathcal{T}^c(x, x') \exp(-ikx'\theta) dx' d\theta. \quad (29)$$

Using Equation (24), we express the reference field as:

$$u_r(x) = \int_{\theta} E(\theta) \exp(-ik\Phi(\theta)) \exp(-ikx\theta) d\theta. \quad (30)$$

From the statistical independence property of different θ ,

$$\begin{aligned} \text{corr}(x; \mathcal{T}^c, \mathcal{P}^c) &= \int_{\theta} |E(\theta)|^2 \exp(-ik\Phi(\theta)) \int_{\epsilon} \mathcal{T}^c(x, x + \epsilon) \exp(-ik\epsilon\theta) d\epsilon d\theta \\ &= \int_{\epsilon} \mathcal{T}^c(x, x + \epsilon) \underbrace{\int_{\theta} A(\theta) \exp(-ik\Phi(\theta) - ik\epsilon\theta) d\theta}_{\equiv \mathcal{P}^c(\epsilon)} d\epsilon, \end{aligned} \quad (31)$$

where we used the fact that $A(\theta) \equiv |E(\theta)|^2$. $\square$

Discussion. We discuss the implications of the two parts of Proposition 2, Equations (27) and (28). To understand the first part, it is helpful to consider discretized versions of the coherent Green's function $\mathcal{T}^c$ and coherent probing kernel $\mathcal{P}^c$: We discretize the

focus plane $z = z_o$ into N points. Using the same discretization on the two arguments of the coherent Green's function $\mathcal{T}^c$ results in an $N \times N$ complex matrix $\mathbf{T}^c$ that is often referred to in optics as the *transmission matrix* [Judkewitz et al. 2015; Popoff et al. 2010a]. Likewise, using this discretization on the argument of the coherent probing kernel $\mathcal{P}^c$ and correlation function $\text{corr}(x; \mathcal{T}^c, \mathcal{P}^c)$ results in the $N \times 1$ complex *coherent probing vector* $\mathbf{p}^c$ and *correlation vector* $\mathbf{c}$, respectively. We will additionally introduce the complex *coherent probing matrix* $\mathbf{\Pi}^c \equiv \text{Toeplitz}(\mathbf{p}^c)$, which is an $N \times N$ Toeplitz matrix formed using the coherent probing vector as its generator. Then, Equation (27) can be written in discrete form:

$$\mathbf{c} = (\mathbf{\Pi}^c \odot \mathbf{T}^c) \mathbf{1}, \quad (32)$$

where $\mathbf{1}$ is an $N \times 1$ vector of ones, and $\odot$ denotes element-wise multiplication. This is the exact analogue of the *transport probing equation* of O'Toole et al. [2012], applied on the transmission matrix instead of the incoherent light transport matrix, and using a (possibly complex) convolutional probing matrix.

The second part of Proposition 2, Equation (28), explains how to implement an arbitrary coherent probing kernel $\mathcal{P}^c$: We can compute the inverse Fourier transform of $\mathcal{P}^c$, then use the amplitude $A(\theta)$ to modulate the amplitude of the plane waves induced into the interferometer by an incoherent area emitter, and its phase $\Phi(\theta)$ to modulate the phase of the response of the reference arm to each such plane wave. We discuss more details about how to implement arbitrary functions $A(\theta)$ and $\Phi(\theta)$ in Section 5.

Finally, we note that when the phase modulation function $\Phi(\theta)$ is independent of θ , and the amplitude modulation function $A(\theta)$ is a rectangular function, then the coherent probing kernel exactly equals the mutual intensity kernel $\mathcal{S}^c$ (Equation (16)) of an incoherent area emitter. Likewise, the coherent probing matrix $\mathbf{\Pi}^c$ is equivalent to the *mutual intensity matrix* [Ozaktas et al. 2016, 2002] of the incident field at the discrete set of locations x on the plane $z = z_o$. In this case, Equation (28) corresponds to the case of "spatial probing with a planar mirror" described by Gkioulekas et al. [2015, Figure 6(a)]. Proposition 2 generalizes the types of coherent probing that can be achieved through interferometric measurements: By appropriately modulating the amplitude of the source and the phase response of the reference arm, we code the mutual intensity function of the incident illumination. Then, our analysis shows that performing interferometry is equivalent to probing the scene's transmission matrix with a probing matrix equal to the corresponding coded mutual intensity matrix. We thus refer to imaging based on Proposition 2 as *interferometry with coded mutual intensity*.

4.2 Probing the light transport matrix

We now discuss how the coherent probing procedure we derived in the previous section relates to probing the incoherent light transport matrix, as introduced by O'Toole et al. [2012]. For this, we first prove the following proposition.

PROPOSITION 3. *When using sufficiently incoherent illumination, the intensity image I_s of Equation (10) of target arm of the generalized Michelson interferometer in Figure 3 is approximately equal to:*

$$I_s(x) \approx \int_{x'} |\mathcal{T}^c(x, x')|^2 dx'. \quad (33)$$

PROOF. We use Equation (26) and express the scattered field as:

$$u_s(x) = \int_{\theta} E(\theta) \int_{x'} \mathcal{T}^c(x, x') \exp(-ikx'\theta) dx' d\theta. \quad (34)$$

From the statistical independence property of different θ ,

$$I_s(x) = |u_s(x)|^2 = \iint_{x', x''} \mathcal{T}^c(x, x') \mathcal{T}^{c*}(x, x'') \underbrace{\int_{\theta} |E(\theta)|^2 \exp(-ik(x' - x'')\theta) d\theta}_{\equiv S^c(x' - x'')} dx' dx''. \quad (35)$$

Imaging in the radiometric light transport regime requires that the illumination is sufficient incoherent, or equivalently that the width W of the area of the source with non-zero emission $|E(\theta)|^2$ is large. Then, Equation (16) suggests that:

$$S^c(x' - x'') \propto \text{sinc}(W \cdot (x' - x'')) \approx \delta(x' - x''), \quad (36)$$

resulting in:

$$I_s(x) \approx \int_{x'} |\mathcal{T}^c(x, x')|^2 dx'. \quad (37)$$

This concludes the proof. $\square$

Equation (33) suggests that the squared amplitude of the coherent Green's function $|\mathcal{T}^c(x, x')|^2$ is approximately equal to the *incoherent Green's function* $\mathcal{T}^i(x, x')$: This is a scalar positive-valued function that radiometrically models light transport through the scene from a point source at y to a point sensor at x , by solving the radiative transfer and rendering equations with appropriate boundary conditions [Case and Zweifel 1967; Gkioulekas et al. 2016]. By discretizing this function in the same way we discretized the coherent Green's function, we obtain the $N \times N$ positive-valued *light transport matrix* $\mathbf{T}^i$ used by O'Toole et al. [2012]. We note that, in more precise terms, Equation (35) suggests that the equivalence between the incoherent Green's function $\mathcal{T}^i(x, x')$ and $|\mathcal{T}^c(x, x')|^2$ requires first spatially-blurring the coherent Green's function $\mathcal{T}^c(x, y)$ with a blur kernel corresponding to the sinc spatial coherence function in Equation (36). The width of the main lobe of this blur kernel represents an inherent resolution limit when working with spatially incoherent light [Levin et al. 2013; Zhang and Levoy 2009]. We refer to Levin et al. [2013] for a detailed discussion of this resolution limit.

The above discussion motivates the question: Can we relate the coherent probing Equation (27) and coherent probing kernel $\mathcal{P}^c$, to probing of the incoherent Green's function? For scenes with highly-diffuse transport, we can answer this question affirmatively.

PROPOSITION 4. *In scenes with sufficiently randomizing transmission, the squared amplitude of the correlation corr of Equation (10) captured from the target arm the generalized Michelson interferometer of Figure 3 is approximately equal to:*

$$|\text{corr}(x; \mathcal{T}^c, \mathcal{P}^c)|^2 \approx \int_{\epsilon} |\mathcal{T}^c(x, x + \epsilon)|^2 |\mathcal{P}^c(\epsilon)|^2 d\epsilon. \quad (38)$$

PROOF. Using Equation (27), we can write:

$$|\text{corr}(x; \mathcal{T}^c, \mathcal{P}^c)|^2 = \iint_{\epsilon, \zeta} \underbrace{(\mathcal{T}^c(x, x + \epsilon))^* \mathcal{T}^c(x, x + \zeta)}_{\equiv C_{x,x}^c(x + \epsilon, x + \zeta)} (\mathcal{P}^c(\epsilon))^* \mathcal{P}^c(\zeta) d\epsilon d\zeta, \quad (39)$$

where the function $C_{x,x}^c(x + \epsilon, x + \zeta)$ is the correlation between the scattered fields that would be observed at point x when exciting the target arm with point sources at points $x + \epsilon$ and $x + \zeta$. For scenes with highly-randomizing transmission, this function is approximately band-diagonal, with a very thin non-zero band region [Judkewitz et al. 2015]. Then, we can write:

$$C_{x,x}^c(x + \epsilon, x + \zeta) \approx \delta(\zeta - \epsilon), \quad (40)$$

and we can reduce Equation (39) to:

$$|\text{corr}(x; \mathcal{T}^c, \mathcal{P}^c)|^2 \approx \int_{\epsilon} |\mathcal{T}^c(x, x + \epsilon)|^2 |\mathcal{P}^c(\epsilon)|^2 d\epsilon. \quad (41)$$

This concludes the proof. $\square$

The condition of highly-randomizing transmission is typically satisfied by scenes with near-Lambertian surfaces, very rough reflection and transmission, or subsurface multiple scattering. In such scenes, and together with Proposition 3, Equation (38) shows that the squared amplitude of the correlation is equal to probing the incoherent Green's function $\mathcal{T}^i$ with an *incoherent probing kernel* that equals the squared amplitude of the coherent one, $\mathcal{P}^i(x) \equiv |\mathcal{P}^c(x)|^2$. Using discretization, we can define the $N \times 1$ positive-valued *incoherent probing vector* $\mathbf{p}^i$, and the $N \times N$ positive-valued *incoherent probing matrix* $\mathbf{\Pi}^i \equiv \text{Toeplitz}(\mathbf{p}^i)$, exactly analogously to their coherent counterparts. Then, Equation (38) can be rewritten as:

$$\mathbf{i} \approx (\mathbf{\Pi}^i \odot \mathbf{T}^i) \mathbf{1}, \quad (42)$$

where we use $\mathbf{i}$ to denote the discretization of $|\text{corr}(x; \mathcal{T}^c, \mathcal{P}^c)|^2$, to emphasize its relation to intensity. Equation (42) is exactly the light transport probing equation of O'Toole et al. [2012].

Before we conclude this section, it is worth revisiting Equation (39): Analogously to Proposition 3, this equation suggests that the equivalence with incoherence probing requires first spatially-blurring the coherent Green's function $\mathcal{T}^c(x, x')$ with a blur kernel corresponding to the diagonal band of the correlation function $C_{x,x}^c(x + \epsilon, x + \zeta)$ in Equation (40). This diagonal band corresponds to the *memory effect* of speckle [Goodman 2007, Section 5]; therefore, the extent of the memory effect introduces a limit on the resolution at which we can perform incoherent probing. The memory effect of speckle has recently been used in computer vision and graphics for motion tracking and non-line-of-sight imaging [Smith et al. 2017, 2018], and has been extensively modeled using Monte Carlo rendering algorithms [Bar et al. 2019]. We note that these works focus on *far-field* speckle correlations (directional sources and sensors), whereas the correlation and memory effect relevant for Proposition 4 are for the *near-field* case (point sources and sensors) [Judkewitz et al. 2015; Osnabrugge et al. 2017].

5 IMPLEMENTATION

In this section, we discuss how to design and optimize an imaging system for interferometry with coded mutual intensity as in Figure 3. Figure 4(a) shows a schematic of our design, which we use throughout this section for reference. Figure 4(b) shows a photograph of a physical prototype. In the supplement we provide details about constructing, aligning, and operating the setup.




Fig. 4. **Setup for interferometry with coded mutual intensity.** (a) Schematic of the setup. The setup is a modified Michelson interferometer that includes an illumination component with amplitude modulation (green), and a reference arm with phase modulation (blue). The dotted thick lines indicate the Fourier planes where amplitude and phase modulation take place. (b) Photograph of implemented physical prototype. Colors indicate the implementations of the corresponding components in the schematic in (a).

Amplitude modulation. We first design an optical configuration for creating *monochromatic* illumination with complex emission E such that its squared amplitude matches the desired amplitude illumination, $|E(\theta)|^2 = A(\theta)$. Theoretically, a straightforward way to generate such illumination would be as follows: first, we use a monochromatic area emitter whose emission area is larger than the support of A ; second, we image this emission area on an amplitude *spatial light modulator* (SLM) applying amplitude modulation $\sqrt{A}$; finally, we collimate the output with a lens. Unfortunately, such an implementation is challenging to realize for two reasons: First, area emitters are typically not sufficiently monochromatic. Second, for sparse amplitude modulation functions A , this implementation would be very light-inefficient, as most of the source’s power would be blocked by the amplitude SLM.

We overcome both challenges by using the setup of Figure 4(a): We use a two-dimensional MEMS mirror to steer a collimated coherent laser beam, which is then focused by a scan lens at the focal plane of the main illumination lens. As the direction of the beam incident on the scan lens changes, the focus spot scans the focal plane in a programmable manner, and this scan can take place *within exposure*. Effectively, this scanning scheme corresponds to using time-multiplexing to implement the integration over θ in Equation (7). To ensure temporal coherence, we use a *single-longitudinal-mode laser*, with an estimated temporal coherence length of 20 m. Theoretically, this optical configuration achieves optimal light efficiency, redirecting all of the laser source’s power towards the scene: The steering mirror can be programmed so that the focused point scans only locations θ of the focal plane where $A(\theta)$ is non-zero, and stays at each location for an amount of time proportional to $A(\theta)$.

In practice, not all scanning patterns are realizable, both because of acceleration and speed limits imposed by the MEMS mirror, and because the function $A(\theta)$ can be spatially discontinuous, requiring the focused point to instantaneously “jump” from one location θ to another. To address this issue, we place an amplitude *electro-optic modulator* (EOM) between the laser source and the steering mirror, which we synchronize with the mirror: Mirror steering is used to scan only locations θ within the support of $A(\theta)$ (or some superset

of this support, as dictated by speed and acceleration limits), and the EOM is used to attenuate the beam at each such location so that the effective overall modulation matches $A(\theta)$. As both the EOM and mirror support MHz operation, this scanning process can take place within exposure. Even though it does not achieve theoretical optimality, the resulting configuration remains significantly more light efficient than the alternative based on an amplitude SLM, and at the same time ensures temporal coherence.

We conclude this discussion with two remarks. First, the configuration we use for amplitude modulation is equivalent to using a laser projector coupled with a scan lens. Unfortunately, the laser diodes in commercial laser projectors have temporal coherence lengths of a few millimeters, making it necessary for us to implement a custom system incorporating a single-longitudinal-mode laser. Second, we can place the above light efficiency considerations within the framework of O’Toole et al. [2015]: Using their terminology, when $A(\theta)$ can be realized without an EOM through scanning patterns of the MEMS mirror, the amplitude modulation configuration of Figure 4(a) is equivalent to an *impulse projector*. When the EOM is necessary, the configuration is a *redistributive projector*. In both cases, the projector is used to project not directly the probing pattern, but instead its inverse Fourier transform (Equation (28)).

Phase modulation. We additionally need to design an optical configuration for implementing the phase modulation Φ . Unlike with amplitude modulation, which can be applied directly on the illumination incident on both arms of the interferometer, the phase modulation needs to be applied only on the reference arm. We achieve this using the optical configuration shown in Figure 4(a): A phase SLM is placed at the focal plane of a lens in the reference arm, and projects a phase modulation pattern equal to Φ . Unlike with amplitude modulation, the use of the phase SLM does not result in light loss, as phase SLMs reflect (most of) the energy incident on them. We note that the combination of the collimating lens with the phase SLM acts as a retroreflector, introducing a reflection of the x coordinate of the beam (Equation (13)). We use two additional lenses to cancel this flip (“flip compensation”).




Fig. 5. **Calibration and acquisition pipeline.** (a) To implement a Toeplitz probing matrix, we first compute the inverse Fourier transform of the probing kernel generating the matrix (Equation (28)). We then use the resulting amplitude and phase to program the amplitude and phase modulation components of our system. (b) To capture probing measurements of a scene, we first capture images at multiple sub-wavelength phase shifts, implemented using phase modulation. We then process these measurements to estimate the phase and amplitude of the interference component (Equations (44)-(45)).

Using temporally incoherent light. As we discussed in Section 3, our technique can be used with temporally incoherent light, to combine probing with transient imaging. This requires making two modifications to the setup of Figure 4: First, we replace the single-longitudinal-mode laser with a broadband spatially-coherent source. For this, we use a supercontinuum laser, though we could also use a superluminescent diode. Second, we mount the phase SLM and its lens on a translation stage, to be able to scan it at long distances. With these modifications, we capture transient sequences by applying, for each position of the translation stage, the acquisition pipeline described in Section 6 for the temporally-coherent case.

6 ACQUISITION PIPELINE

We use this section to provide details about how to use the optical setup of Section 5, in order to capture coherent and incoherent probing measurements. Figure 5 summarizes our acquisition pipeline, including both programming the modulation components of the setup and post-processing captured measurements. All of our code for this section is available at the project website [Kotwal et al. 2020].

Implementing a probing pattern. Before capturing measurements with our setup, we need to program it with the desired probing kernel $\mathcal{P}^c$. For this, we compute the inverse Fourier transform of $\mathcal{P}^c$ to obtain the required amplitude modulation $A(\theta)$ and phase modulation $\Phi(\theta)$ functions, as in Equation (28). We note that it is possible that the amplitude modulation function $A(\theta)$ will have negative values, which are not physically realizable. We address this by using $\tilde{A}(\theta) = |A(\theta)|$ and $\tilde{\Phi}(\theta) = \Phi(\theta) + \frac{\lambda}{2}$; from Euler's formula, this results in the same probing kernel. The amplitude and phase modulation functions are loaded programmatically on the EOM and MEMS mirror (amplitude) and phase SLM (phase). This calibration process only needs to be performed once per probing kernel.

Capturing probing measurements. Acquiring probing measurements is equivalent to measuring either the complex correlation $\text{corr}(x; \mathcal{T}^c, \mathcal{P}^c)$ (Equation (27)) for coherent probing; or its squared amplitude $|\text{corr}(x; \mathcal{T}^c, \mathcal{P}^c)|^2$ (Equation (38)) for incoherent probing. From Equation (9), we note that an intensity measurement captured at a camera pixel focused at location x on the plane $z = z_o$ equals:

$$I(x) = \underbrace{I_s(x) + I_r(x)}_{\equiv I_{DC}(x)} + \text{Re}\{\text{corr}(x; \mathcal{T}^c, \mathcal{P}^c)\}. \quad (43)$$

Analogously to continuous-wave time-of-flight cameras [Gupta et al. 2015] and phase-shifting interferometry [Creath 1985], we estimate the correlation term in a two-step procedure: First, we capture $N \geq 3$ intensity measurements, where for each of them we use the phase SLM to shift the phase modulation by a sub-wavelength amount, $\Phi(\theta) + n2\pi/N$, $n = 0, \dots, N$. Second, at each pixel, we fit a sinusoid to these N measurements, to obtain the per-pixel amplitude and phase of the correlation term $\text{corr}(x; \mathcal{T}^c, \mathcal{P}^c)$. For incoherent probing, we found empirically that it is more robust to directly estimate the squared amplitude $|\text{corr}(x; \mathcal{T}^c, \mathcal{P}^c)|^2$ as:

$$\langle I_{DC}(x) \rangle = \frac{1}{N} \sum_{n=1}^N I_n(x), \quad (44)$$

$$\left\langle |\text{corr}(x; \mathcal{T}^c, \mathcal{P}^c)|^2 \right\rangle = \frac{1}{N} \frac{1}{\pi} \sum (I_n(x, z) - \langle I_{DC}(x) \rangle)^2. \quad (45)$$

In practice, we found that we can reliably estimate the correlation term or its squared amplitude using measurements at $N = 10$ sub-wavelength phase shifts. We note that, as phase shifting is performed by the phase SLM and no mechanical parts are involved, these measurements can be captured at a frame rate of 6 Hz. This is limited by the maximum refresh rate of the phase SLM (60 Hz).

Dealing with speckle. In real scenes, the correlation measurements $\langle \text{corr}(x; \mathcal{T}^c, \mathcal{P}^c) \rangle$ will contain significant pseudo-random speckle noise. When doing incoherent probing, to eliminate these speckle artifacts, we blur our estimates with a small blur kernel $B(x)$. Therefore, our final estimate for incoherent probing is:

$$\left\langle |\text{corr}(x; \mathcal{T}^c, \mathcal{P}^c)|^2 \right\rangle * B(x). \quad (46)$$

For a detailed analysis of speckle in interferometry, we refer to Gkioulekas et al. [2015]. We follow their suggestions on how to set magnification and aperture size to maximize interference contrast. We also use a neutral density filter in the reference arm of our setup (Figure 4), to equalize the intensity of the reference and target arms.

7 EXPERIMENTS

In this section, we show results of experiments performed using the prototype described in Section 5. Our experiments emphasize probing types that cannot be implemented using previous interferometric techniques [Gkioulekas et al. 2015], or do not have light-efficient




Fig. 6. **Coherent probing of a retroreflector scene.** The images show the real part of transmission matrices probed using the amplitude modulation functions A shown in the inset. From Proposition 1, each image equals the real part of the corresponding probing kernel, and the support of the kernel corresponds to the part of the image where fringes are present. (a) Probing with horizontal line amplitude modulation functions of different widths. As we increase the size of A , the extent of the corresponding probing kernel decreases. (b) Probing with line amplitude modulation functions of different orientations. As we rotate A , the probing kernel rotates as well. (c) Probing with horizontal line amplitude modulation functions of different profiles. As we change the profile of A , the profile of the probing kernel also changes. The absolute values of the profiles of A and the probing kernel are shown in (d) and (e), respectively. We use functions A to implement probing kernels shaped (from left to right) as sinc, Laplacian-of-Gaussian, and Gaussian functions. Please see the project website [Kotwal et al. 2020] for the complete visualizations of the captured fields.

implementations using projector-camera techniques [O’Toole et al. 2012]. The project website [Kotwal et al. 2020] shows additional results, including full video sequences for the experiments combining probing with transient imaging. To facilitate reproducibility, we provide all of our data at the project website.

Coherent probing. In Figure 6, we perform experiments to coherently probe the transmission matrix of a retroreflector scene, similar to Figure 1. As we showed in Proposition 1, the resulting correlation will equal the probing kernel, which we can then visualize by plotting, for instance, its real part. To demonstrate this, we show experiments using a variety of one-dimensional probing kernels corresponding to amplitude modulation functions A of different sizes




Fig. 7. **Direct-only and global-only images.** We use three scenes: cup (top), bead (middle), and gummy bear (bottom). For all scenes, the left column shows a regular photograph of the scene, and the middle and right columns show direct-only and global-only images captured using our incoherent probing technique. The cup and gummy bear appear metallic in the direct-only images, as the subsurface scattering effects dominating their appearance are only measured in the global-only images. The direct-only image of the bead is missing the strong specular reflections on the bead’s faces and the on the wall, which are prominent in the global-only image.

(Figure 6(a)), orientations (Figure 6(b)), and profiles (Figure 6(c)). In the latter case, we implement one-dimensional sources with horizontal emission profiles corresponding to amplitude modulation that is: approximately constant, $A(\theta) \propto 1$; Gaussian, $A(\theta) \propto \exp(-\theta^2/\sigma^2)$ for some standard deviation σ ; and the inverse Fourier transform of the Laplacian-of-Gaussian function, $A(\theta) \propto \theta^2 \exp(-\theta^2/\sigma^2)$. In all cases, the resulting probing kernel is constant along the direction that is orthogonal to the orientation of A —for example, when A is horizontal, the probing kernel is constant along the vertical direction. In Figure 6(a), we observe that as we make A narrower, the support of the probing kernel becomes wider. In Figure 6(b), we observe that as we rotate A , the probing kernel rotates as well. Both of these observations correspond to standard properties of the Fourier transform. Finally, in Figure 6(c), we observe that the profile of the probing kernel matches the Fourier transform of A . These results are consistent with Proposition 1, which predicts that A and the probing kernel are related through a Fourier transform.

Direct-global separation. From the Fourier transform relationship of Proposition 2, the probing kernel due to a two-dimensional square amplitude modulation function A is a two-dimensional sinc. As the size of A increases, the central lobe of the probing kernel shrinks and the secondary lobes become negligible—that is, the probing kernel approximates a Dirac delta function. From Propositions 2 and 4, we see that this is equivalent to probing just the diagonal of the transmission and transport matrices, respectively. As explained by O’Toole et al. [2012], this is approximately equivalent to measuring only the direct and retroreflecting components of light propagation in a scene. Conversely, we can measure only the global component

by using both amplitude modulation A and phase modulation Φ to create a probing kernel that is constant everywhere, except for a small area around $x = 0$ where it is zero.

Figure 7 shows three example scenes where we use the two probing kernels described above to separately measure the direct and global components of the incoherent light transport matrix. The first scene (top row) uses a small toy cup, resulting in three types of light paths: direct reflections; retroreflective specular, which show up as the strong specularities at the upper and lower parts of the cup; and indirect paths due to diffuse interreflections and sub-surface scattering. The direct-only measurement removes the last two types of paths, making the cup appear metallic. The global-only measurement, on the other hand, includes the caustics, interreflections, and translucent effects due to the indirect paths.

The second scene consists of a plastic transparent bead placed between two diffuse walls. The bead faces reflect light and create strong specularities on the walls. Additionally, some of the light from the walls reflects specularly on the bead and towards the camera, resulting in specular reflections on some of the bead faces. All these specular effects are removed from the direct-only measurement, and are prominent in the global-only measurement. The global-only component additionally shows effects due to the interreflections between the diffuse walls and inside the bead.

The third scene is a gummy bear placed between two diffuse walls. The majority of the light received from the gummy bear is due to subsurface scattering inside it. This light is removed from the direct-only component, resulting in the gummy bear appearing metallic due to surface-only reflections. By contrast, the global-only component is dominated by this scattered light.

We note that most previous techniques estimate the global component by first capturing a direct-only image, and then subtracting it from a regular image. By contrast, our technique allows us to directly measure the global-only image. For all scenes in Figure 7, we have confirmed that the global-only image we capture using probing matches the image we obtain using the subtraction procedure.

Anisotropic descattering. As noted in O’Toole et al. [2012], when imaging a target through a scatterer, it is possible to suppress the effect of scattering by probing with a pattern that emphasizes the diagonal of the light transport matrix, and subtracts the first few off-diagonals. Using coherent probing, we can perform this type of descattering with a probing kernel that is shaped like a Laplacian-of-Gaussian function—the positive central lobe emphasizes the main diagonal of the transmission matrix, and the off-center negative values subtract the off-diagonals. We can additionally selectively perform descattering along only certain directions by using probing kernels that are shaped like *anisotropic* Laplacian-of-Gaussian functions. This is equivalent to a form of optical matched filtering that emphasizes features similar to the probing kernel.

Figure 8 shows an example of this optical matched filtering procedure. We use a template consisting of two vertical white stripes (width 100 μm) against a black background (Figure 8(a)). We place this template behind a semi-transparent scatterer of thickness 50 mm, created by immersing titanium-dioxide microspheres in silicone rubber. In a regular image of this scene (Figure 8(b)), the two stripe




Fig. 8. **Anisotropic descattering.** (a) We place a pattern consisting of two vertical stripes behind a thick semi-transparent scatterer. (b) In a regular image of the scene, the stripe features cannot be resolved because of the scattering. (c) Probing the scene with a horizontal Laplacian-of-Gaussian kernel that matches the orientation of the two stripes enhances their contrast and makes them clearly visible. (d) By contrast, probing the scene with a vertical kernel results in the two stripes remaining obscured.

features are obscured by the scattering and are not discernible. Probing with a horizontal Laplacian-of-Gaussian kernel enhances vertical features, making it possible to resolve the stripe patterns (Figure 8(c)). By contrast, the stripes remain obscured when we probe with a vertical kernel (Figure 8(d)). We note that, in contrast to our technique, these anisotropic probing kernels cannot be implemented in a light-efficient way using primal-dual coding.

Measuring the entire light transport matrix. We create a two-dimensional scene consisting of a diffuse surface and a mirror forming a right-angle corner (Figure 9(a)). The two-dimensional light transport matrix of this scene is dominated by the main diagonal and anti-diagonal (Figure 9(c)): The diagonal of the light transport matrix corresponds to contributions from direct and retroreflecting paths, whereas the anti-diagonal corresponds to contributions from two-bounce reflection paths (Figure 9(b)).

We use our technique to capture the two-dimensional light transport matrix of this scene (Figure 9(d)). We achieve this by successively probing with shifted diagonal kernels (Figure 9(e)), implemented using phase modulation Φ equal to ramp functions of varying slopes. The correlation measurements captured for different slopes (Figure 9(f)) are equal to different diagonals of the light transport matrix, and stacking them together forms the entire light transport matrix. We provide a video visualization of this procedure at the project website [Kotwal et al. 2020]. We note that the probing kernels in this experiment correspond to the short-range-indirect probing patterns that have recently been successfully used for tissue imaging [Gkioulekas et al. 2016; Kubo et al. 2018; Liu et al. 2020].

Combining probing and transient imaging. We perform experiments where we use temporally coherent light to combine our probing technique with transient imaging. We use our system to capture




Fig. 9. **Measuring the light transport matrix of a mirror-diffuser corner scene.** (a) Photograph of the scene for visualization. (b) Light paths in the scene: blue are direct paths, purple and yellow are two-bounce reflection paths, and green are retroreflecting paths. Below the schematic is a regular image of the scene. (c) The two-dimensional light transport matrix for this scene: direct and retroreflecting paths show up on the diagonal, and the two-bounce reflection paths show up on the antidiagonal. (d) Measured light transport matrix, with regions corresponding to different matrix parts in (c) marked in color. (e) Example probing patterns used to measure the light transport matrix. (f) Measurements using the probing patterns in (e), where the top and bottom images are due to two-bounce reflection paths, and the middle image is due to direct and retroreflecting paths. The dashed lines in (b) and (f) indicate the corner location. Please see the project website [Kotwal et al. 2020] for the complete set of probing measurements.

transient sequences without any probing, as well as sequences combined with anisotropic probing at different orientations. Figure 10 shows examples for two scenes, where we isolate the same frame in three transient sequences captured with different probing patterns. We observe that, when using anisotropic probing, only features that match the probing kernel are preserved, while other features are suppressed relative to their intensity when no probing is used. We provide video visualizations of the entire transient sequences at the project website [Kotwal et al. 2020].

8 DISCUSSION AND LIMITATIONS

We discuss limitations and directions for future exploration relating to interferometry with coded mutual intensity.

Comparison with primal-dual coding. Interferometry with coded mutual intensity and primal-dual coding [O’Toole et al. 2012] both implement probing of incoherent light transport, albeit using very different optical configurations and operating principles. Therefore, it is worthwhile discussing their relative strengths and weaknesses.

For this, we first contrast how the two methodologies implement different probing matrices. Given a desired probing matrix, primal-dual coding decomposes it into a sequence of illumination and sensing masks, which are then applied respectively on a projector source and sensor, successively within a single exposure. As discussed by O’Toole et al. [2015], this procedure is light-efficient only when the probing matrix decomposition is *homogeneous*; for arbitrary probing matrices, a homogeneous decomposition may not exist or may require a very large number of mask patterns. On the other hand, interferometry with coded mutual intensity uses an inverse Fourier transform to convert to a probing matrix to individual amplitude and phase modulation functions, which are then applied to the Fourier-domain of the illumination and reference arm, respectively, of an interferometer. From a hardware complexity perspective, the two techniques introduce a trade-off between sensor masking (primal-dual coding) versus interferometry with phase modulation (interferometry with coded mutual intensity). In terms of the types of probing that can be realized, by using a Fourier-domain redistributive projector for amplitude modulation (Section 5), interferometry with coded mutual intensity enables light-efficient implementations of probing matrices that would be challenging to implement with primal-dual coding (e.g., the anisotropic probing patterns shown in Section 7). However, this comes at the cost of only being able to implement Toeplitz probing matrices, a limitation not shared by primal-dual coding. Additionally, the advantage of better light efficiency can be outweighed by other noise performance issues typical in interferometry, as we discuss later in this section.

As probing the spatial dimensions of light transport is orthogonal to probing its temporal dimension, both primal-dual coding and interferometry with coded mutual intensity can be combined with transient imaging techniques. Primal-dual coding is better suited for macroscopic transient imaging, as the optical components it requires can be readily combined with continuous-wave time-of-flight cameras [Achar et al. 2017; O’Toole et al. 2014a]. By contrast, interferometry with coded mutual intensity is more readily combined with microscopic transient imaging, as both techniques rely on interferometric setups [Gkioulekas et al. 2015].

Finally, there are capabilities that have been demonstrated using only one or the other methodology: Primal-dual coding has been used in non-confocal setups to implement epipolar and disparity probing patterns [O’Toole et al. 2014b]. By contrast, we have only used interferometry with coded mutual intensity in a confocal setting, though our theory can also be applied to non-confocal settings (for example, by incorporating our amplitude and phase modulation systems within a Mach-Zehnder interferometer). Conversely, we have shown that interferometry with coded mutual intensity can be used to perform coherent probing of transmission matrices. This




Fig. 10. **Combining probing and transient imaging.** We use two scenes: cup (top), and bead (bottom). From left to right: regular image, and isolated frames of the transient sequence at the same time instance when using no probing, anisotropic probing at $+45^\circ$, and anisotropic probing at -45° . In the cup scene, when using probing, only the parts of the caustic that match the probing kernel are maintained, and the rest of the caustic is suppressed. Likewise, in the bead scene, different reflections on wall and on the bead faces are maintained and suppressed, depending on the probing kernel used. Please see the project website [Kotwal et al. 2020] for the complete transient sequences and additional probing patterns.

capability has not been demonstrated using primal-dual coding, though O’Toole [2016] discusses possible extensions to this end.

Noise and speckle. The signal-to-noise ratio of measurements captured using interferometry with coded mutual intensity can significantly deteriorate because of two factors inherent in interferometric systems. The first factor relates to the fact that the measurements captured by cameras in interferometric systems include not only the interference component, but additionally the DC intensity images of the reference and scattered arms (Equation (43)). The Poisson noise in these measurements is proportional to the total intensity [Hasi-noff et al. 2010], and thus it can potentially dominate the interference component when that is significantly weaker than the DC component. We refer to Takada [1998] for a detailed analysis of the noise characteristics of interferometric measurements.

The second factor relates to the presence of speckle due to the use of (partially-)coherent illumination. As Gkioulekas et al. [2015] have shown, speckle can negatively impact interference contrast, and careful optical design is needed to minimize its effect. The impact speckle has on the quality of interferometric measurements can be assessed from the effect speckle has on an image of the target arm under the same illumination as that used for interferometry.

In practice, we note that both of these issues are counteracted to a significant extent by the fact that we estimate the phase and amplitude of the interference component using multiple image measurements, captured with sub-wavelength phase shifts (Section 6). However, as we discuss immediately below, this comes at the cost of slower operation of the imaging system.

Towards real-time operation. Our current implementation captures images at a relatively low frame rate of 6 Hz. The primary

factor preventing us from achieving higher frame rates is our use of a phase SLM to implement not only phase modulation (which remains constant throughout capture), but also sub-wavelength phase shifts for estimating the interference component. Consequently, the maximum refresh rate of the phase SLM (60 Hz for commercially-available units) creates an upper bound on the maximum frame rate that can be achieved by our system.

To overcome this constraint, we can draw inspiration from other applications of phase-shifting interferometric. For example, we could potentially use the phase SLM to implement space-division multiplexing [Kakue et al. 2011], allowing us to capture multiple phase shifts in a single measurement, at the cost of reduced spatial resolution. Alternatively, we could apply phase retrieval algorithms [Fienup 1982], in order to recover the phase and amplitude of the interference component from measurements at just one or two phase shifts. A potential downside of these approaches is that, by reducing the number of measurements, they can result in significant deterioration of signal-to-noise ratio, as discussed above.

A different approach would be to accelerate measurements at multiple phase-shifts, by replacing the phase SLM with a different optical component. In particular, we could potentially use a combination of a phase EOM with a MEMS mirror at the reference arm of our interferometer to implement phase modulation at MHz rates, analogously to our implementation of amplitude modulation.

9 CONCLUSION

We introduced a new imaging methodology, interferometry with coded mutual intensity, that uses illumination with coded spatial coherence properties, in order to perform both coherent probing of

transmission matrices, and incoherent probing of light transport matrices. Our theory and optical design allow us to programmatically implement arbitrary convolutional probing patterns, by applying appropriate amplitude and phase modulation at the Fourier domain (i.e., the focal plane of the illumination lens) of a Michelson interferometer, without the need for hardware changes. These capabilities can additionally be incorporated into systems that use temporally-incoherent illumination for optical coherence tomography. We have developed a physical prototype, and demonstrated the validity of our theory and utility of our designs in experiments such as visualizing complex fields, capturing direct and global transport components, acquiring light transport matrices, and performing anisotropic descattering, both in steady-state and transient imaging. We hope that, by significantly expanding the types of probing that can be achieved using interferometric techniques, and by providing the design and engineering details for developing physical prototypes implementing these techniques, our paper will enable further research on interferometric light transport techniques.

ACKNOWLEDGMENTS

We thank Todd Zickler and Matthew O'Toole for helpful discussions. This work was supported by NSF Expeditions award 1730147, DARPA REVEAL contract HR0011-16-C-0028, ONR DURIP award N00014-16-1-2906, ERC 635537, ISF 1046-14, and the Ollendorff Minerva Center of the Technion.

REFERENCES

- Nils Abramson. 1983. Light-in-flight recording: high-speed holographic motion pictures of ultrafast phenomena. *Applied Optics* (1983).
- Supreeth Achar, Joseph R Bartels, William L Whittaker, Kiriakos N Kutulakos, and Srinivasa G Narasimhan. 2017. Epipolar time-of-flight imaging. *ACM TOG* (2017).
- Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. 2019. A Monte Carlo framework for rendering speckle statistics in scattering media. *ACM TOG* (2019).
- 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* (2012).
- William H Carter. 1977. Measurement of second-order coherence in a light beam using a microscope and a grating. *Applied Optics* (1977).
- Kenneth M Case and Paul Frederick Zweifel. 1967. *Linear Transport Theory*. Addison-Wesley.
- Julie Chang and Gordon Wetzstein. 2018. Single-shot speckle correlation fluorescence microscopy in thick scattering tissue with image reconstruction priors. *Journal of Biophotonics* (2018).
- Oliver Cossairt, Nathan Matsuda, and Mohit Gupta. 2014. Digital refocusing with incoherent holography. *IEEE ICCP* (2014).
- Katherine Creath. 1985. Phase-shifting speckle interferometry. *Applied Optics* (1985).
- James R Fienup. 1982. Phase retrieval algorithms: a comparison. *Applied Optics* (1982).
- Genevieve Garipey, Nikola Krstajić, Robert Henderson, Chunyong Li, Robert R Thomson, Gerald S Buller, Barmak Heshmat, Ramesh Raskar, Jonathan Leach, and Daniele Faccio. 2015. Single-photon sensitive light-in-flight imaging. *Nature communications* (2015).
- Ioannis Gkioulekas, Anat Levin, Frédo Durand, and Todd Zickler. 2015. Micron-scale light transport decomposition using interferometry. *ACM TOG* (2015).
- Ioannis Gkioulekas, Anat Levin, and Todd Zickler. 2016. An Evaluation of Computational Imaging Techniques for Heterogeneous Inverse Scattering. *ECCV* (2016).
- Joseph W Goodman. 2000. *Statistical Optics*. Wiley-Interscience.
- Joseph W Goodman. 2005. *Introduction to Fourier optics*. Roberts and Company Publishers.
- Joseph W Goodman. 2007. *Speckle phenomena in optics: theory and applications*. Roberts and Company Publishers.
- Mohit Gupta, Agrawal Agrawal, Ashok Veeraraghavan, and Srinivasa G Narasimhan. 2011. Structured light 3D scanning in the presence of global illumination. *CVPR* (2011).
- Mohit Gupta, Shree K Nayar, Matthias B Hullin, and Jaime Martin. 2015. Phasor imaging: A generalization of correlation-based time-of-flight imaging. *ACM TOG* (2015).
- Parameswaran Hariharan. 2003. *Optical interferometry*. Elsevier.
- Samuel W Hasinoff, Frédo Durand, and William T Freeman. 2010. Noise-optimal capture for high dynamic range photography. *CVPR* (2010).
- Felix Heide, Matthias B. Hullin, James Gregson, and Wolfgang Heidrich. 2013. Low-budget Transient Imaging Using Photonic Mixer Devices. *ACM Trans. Graph.* 32, 4, Article 45 (July 2013), 10 pages.
- David Huang, Eric A Swanson, Charles P Lin, Joel S Schuman, William G Stinson, Warren Chang, Michael R Hee, Thomas Flotte, Kenton Gregory, Carmen A Puliafito, and James G Fujimoto. 1991. Optical coherence tomography. *Science* (1991).
- Adrian Jarabo, Belen Masia, Julio Marco, and Diego Gutierrez. 2017. Recent advances in transient imaging: A computer graphics and vision perspective. *Visual Informatics* (2017).
- Kensei Jo, Mohit Gupta, and Shree K Nayar. 2015. Spedo: 6 dof ego-motion sensor using speckle defocus imaging. *ICCV* (2015).
- Benjamin Judkewitz, Roarke Horstmeyer, Ivo M Vellekoop, Ioannis N Papadopoulos, and Changhui Yang. 2015. Translation correlations in anisotropically scattering media. *Nature physics* (2015).
- Achuta Kadambi, Refael Whyte, Ayush Bhandari, Lee Streeter, Christopher Barsi, Adrian Dorrington, and Ramesh Raskar. 2013. Coded Time of Flight Cameras: Sparse Deconvolution to Address Multipath Interference and Recover Time Profiles. *ACM TOG* (2013).
- Takashi Kakue, Ryosuke Yonesaka, Tatsuki Tahara, Yasuhiro Awatsuji, Kenzo Nishio, Shogo Ura, Toshihiro Kubota, and Osamu Matoba. 2011. High-speed phase imaging by parallel phase-shifting digital holography. *Optics letters* (2011).
- Ori Katz, Pierre Heidmann, Mathias Fink, and Sylvain Gigan. 2014. Non-invasive single-shot imaging through scattering layers and around corners via speckle correlations. *Nature photonics* (2014).
- Ahmed Kirmani, Tyler Hutchison, James Davis, and Ramesh Raskar. 2011. Looking around the corner using ultrafast transient imaging. *IJCV* (2011).
- Matias Koivurova, Henri Partanen, Jari Turunen, and Ari T Friberg. 2017. Grating interferometer for light-efficient spatial coherence measurement of arbitrary sources. *Applied Optics* (2017).
- Alankar Kotwal, Anat Levin, and Ioannis Gkioulekas. 2020. Project Website. https://imaging.cs.cmu.edu/coded_mutual_intensity.
- Hirooyuki Kubo, Suren Jayasuriya, Takafumi Iwaguchi, Takuya Funatomi, Yasuhiro Mukaigawa, and Srinivasa G Narasimhan. 2018. Acquiring and characterizing plane-to-ray indirect light transport. *IEEE ICCP* (2018).
- Anat Levin, Daniel Glasner, Ying Xiong, Frédo Durand, William Freeman, Wojciech Matusik, and Todd Zickler. 2013. Fabricating BRDFs at high spatial resolution using wave optics. *ACM TOG* (2013).
- Fengqiang Li, Florian Willomitzer, Prasanna Rangarajan, Mohit Gupta, Andreas Velten, and Oliver Cossairt. 2018. Sh-tof: Micro resolution time-of-flight imaging with superheterodyne interferometry. *IEEE ICCP* (2018).
- Fengqiang Li, Joshua Yablon, Andreas Velten, Mohit Gupta, and Oliver Cossairt. 2017. High-depth-resolution range imaging with multiple-wavelength superheterodyne interferometry using 1550-nm lasers. *Applied Optics* (2017).
- Chao Liu, Akash Maity, Artur W Dubrawski, and Srinivasa G Sabharwal, Ashutosh Narasimhan. 2020. High Resolution Diffuse Optical Tomography using Short Range Indirect Subsurface Imaging. *IEEE ICCP* (2020).
- Tomohiro Maeda, Achuta Kadambi, Yoav Y Schechner, and Ramesh Raskar. 2018. Dynamic heterodyne interferometry. *IEEE ICCP* (2018).
- David Mendlovic, Gal Shabtay, and Adolf W Lohmann. 1999. Synthesis of spatial coherence. *Optics letters* (1999).
- David Mendlovic, Gal Shabtay, Adolf W Lohmann, and Naim Konforti. 1998. Display of spatial coherence. *Optics letters* (1998).
- Christopher A Metzler, Manoj K Sharma, Sudarshan Nagesh, Richard G Baraniuk, Oliver Cossairt, and Ashok Veeraraghavan. 2017. Coherent inverse scattering via transmission matrices: Efficient phase retrieval algorithms and a public dataset. *IEEE ICCP* (2017).
- Dinesh N Naik, Takahiro Ezawa, Yoko Miyamoto, and Mitsuo Takeda. 2009. 3-D coherence holography using a modified Sagnac radial shearing interferometer with geometric phase shift. *Optics express* (2009).
- Shree K Nayar, Gurunandan Krishnan, Michael D Grossberg, and Ramesh Raskar. 2006. Fast Separation of Direct and Global Components of a Scene Using High Frequency Illumination. *ACM TOG* (2006).
- Gerwin Osnabrugge, Roarke Horstmeyer, Ioannis N Papadopoulos, Benjamin Judkewitz, and Ivo M Vellekoop. 2017. Generalized optical memory effect. *Optica* (2017).
- Matthew O'Toole. 2016. *Optical Linear Algebra for Computational Light Transport*. Ph.D. Dissertation. University of Toronto (Canada).
- Matthew O'Toole, Supreeth Achar, Srinivasa G Narasimhan, and Kiriakos N Kutulakos. 2015. Homogeneous codes for energy-efficient illumination and imaging. *ACM TOG* (2015).
- Matthew O'Toole, Felix Heide, David B Lindell, Kai Zang, Steven Diamond, and Gordon Wetzstein. 2017. Reconstructing Transient Images from Single-Photon Sensors. *CVPR* (2017).
- Matthew O'Toole, Felix Heide, Lei Xiao, Matthias B Hullin, Wolfgang Heidrich, and Kiriakos N Kutulakos. 2014a. Temporal Frequency Probing for 5D Transient Analysis

- of Global Light Transport. *ACM TOG* (2014).
- Matthew O'Toole and Kiriakos N Kutulakos. 2010. Optical Computing for Fast Light Transport Analysis. *ACM TOG* (2010).
- Matthew O'Toole, John Mather, and Kiriakos N Kutulakos. 2014b. 3D Shape and Indirect Appearance by Structured Light Transport. *CVPR* (2014).
- Matthew O'Toole, Ramesh Raskar, and Kiriakos N Kutulakos. 2012. Primal-dual Coding to Probe Light Transport. *ACM TOG* (2012).
- Haldun M Ozaktas, Talha Cihad Gulcu, and M Alper Kutay. 2016. Linear algebraic theory of partial coherence: continuous fields and measures of partial coherence. *JOSA A* (2016).
- Haldun M Ozaktas, Serdar Yüksel, and M Alper Kutay. 2002. Linear algebraic theory of partial coherence: discrete fields and measures of partial coherence. *JOSA A* (2002).
- Luping Pan, Xingbing Chao, Zhi-Cheng Ren, Hui-Tian Wang, and Jianping Ding. 2019. Measuring spatial coherence by using a lateral shearing interferometry. *Applied optics* (2019).
- Pieter Peers, Dhruv K. Mahajan, Bruce Lamond, Abhijeet Ghosh, Wojciech Matusik, Ravi Ramamoorthi, and Paul Debevec. 2009. Compressive Light Transport Sensing. *ACM TOG* (2009).
- Christoph Peters, Jonathan Klein, Matthias B Hullin, and Reinhard Klein. 2015. Solving trigonometric moment problems for fast transient imaging. *ACM Trans. Graph.* 34, 6 (2015), 220–1.
- SM Popoff, G Leroosey, R Carminati, M Fink, AC Boccara, and S Gigan. 2010a. Measuring the transmission matrix in optics: an approach to the study and control of light propagation in disordered media. *Physical review letters* 104, 10 (2010), 100601.
- Sébastien Popoff, Geoffroy Leroosey, Mathias Fink, Albert Claude Boccara, and Sylvain Gigan. 2010b. Image transmission through an opaque material. *Nature communications* 1, 1 (2010), 1–5.
- Dikpal Reddy, Ravi Ramamoorthi, and Brian Curless. 2012. Frequency-space Decomposition and Acquisition of Light Transport Under Spatially Varying Illumination. *ECCV* (2012).
- Steve M Seitz, Yasuyuki Matsushita, and Kiriakos N Kutulakos. 2005. A theory of inverse light transport. *ICCV* (2005).
- Pradeep Sen, Billy Chen, Gaurav Garg, Stephen R. Marschner, Mark Horowitz, Marc Levoy, and Hendrik P. A. Lensch. 2005. Dual Photography. *ACM TOG* (2005).
- Manoj Sharma, Christopher A Metzler, Sudarshan Nagesh, Oliver Cossairt, Richard G Baraniuk, and Ashok Veeraraghavan. 2019. Inverse scattering via transmission matrices: Broadband illumination and fast phase retrieval algorithms. *IEEE TCI* (2019).
- Brandon M Smith, Pratham Desai, Vishal Agarwal, and Mohit Gupta. 2017. CoLux: Multi-object 3d micro-motion analysis using speckle imaging. *ACM TOG* (2017).
- Brandon M Smith, Matthew O'Toole, and Mohit Gupta. 2018. Tracking multiple objects outside the line of sight using speckle imaging. *CVPR* (2018).
- Kazumasa Takada. 1998. Noise in optical low-coherence reflectometry. *IEEE Journal of Quantum Electronics* (1998).
- Ivo M Vellekoop, Meng Cui, and Changhui Yang. 2012. Digital optical phase conjugation of fluorescence in turbid tissue. *Applied Physics Letters* (2012).
- Ivo M Vellekoop and AP Mosk. 2007. Focusing coherent light through opaque strongly scattering media. *Optics letters* (2007).
- Andreas Velten, Di Wu, Adrian Jarabo, Belen Masia, Christopher Barsi, Chinmaya Joshi, Everett Lawson, Moungi Bawendi, Diego Gutierrez, and Ramesh Raskar. 2013. Femto-photography: capturing and visualizing the propagation of light. *ACM TOG* (2013).
- Laura Waller, Guohai Situ, and Jason W Fleischer. 2012. Phase-space measurement and coherence synthesis of optical beams. *Nature Photonics* (2012).
- Jian Wang, Joseph Bartels, William Whittaker, Aswin C Sankaranarayanan, and Srinivasa G Narasimhan. 2018. Programmable triangulation light curtains. *ECCV* (2018).
- Jiaping Wang, Yue Dong, Xin Tong, Zhouchen Lin, and Baining Guo. 2009. Kernel Nyström Method for Light Transport. *ACM TOG* (2009).
- A Wax and JE Thomas. 1996. Optical heterodyne imaging and Wigner phase space distributions. *Optics letters* (1996).
- Di Wu, Andreas Velten, Matthew O'Toole, Belen Masia, Amit Agrawal, Qionghai Dai, and Ramesh Raskar. 2014a. Decomposing Global Light Transport Using Time of Flight Imaging. *IJCV* (2014).
- Di Wu, Gordon Wetzstein, Christopher Barsi, Thomas Willwacher, Qionghai Dai, and Ramesh Raskar. 2014b. Ultra-fast Lensless Computational Imaging through 5D Frequency Analysis of Time-resolved Light Transport. *IJCV* (2014).
- Frederik Zernike. 1938. The concept of degree of coherence and its application to optical problems. *Physica* (1938).
- Zhengyun Zhang and Marc Levoy. 2009. Wigner distributions and how they relate to the light field. *IEEE ICCP* (2009).