

Rendering Near-Field Speckle Statistics in Scattering Media

CHEN BAR, Department of Electrical Engineering, Technion, Israel

IOANNIS GKIOULEKAS, Robotics Institute, Carnegie Mellon University, USA

ANAT LEVIN, Department of Electrical Engineering, Technion, Israel

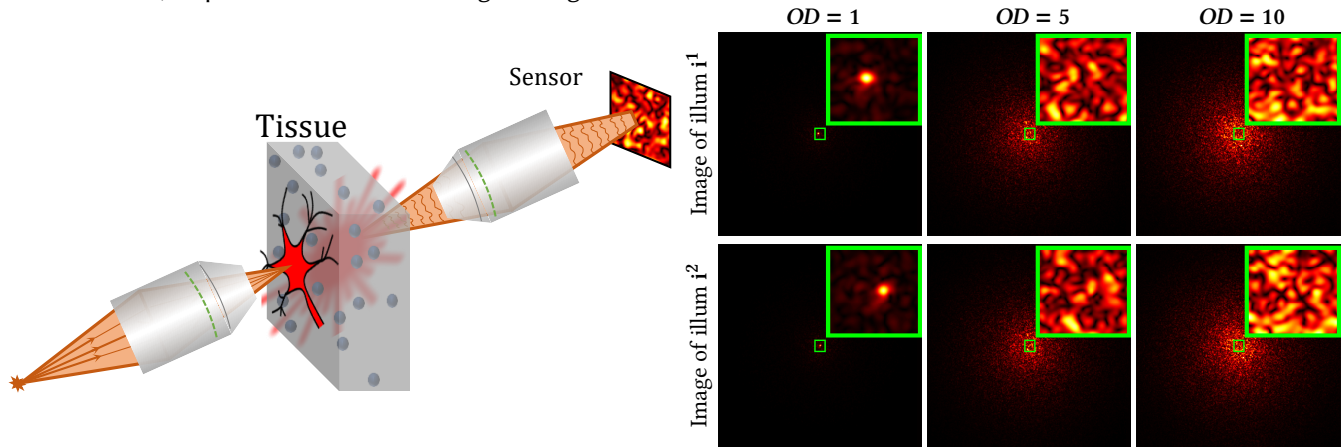


Fig. 1. **Rendering near-field speckle patterns.** We propose an efficient and physically-accurate algorithm that can simulate speckle patterns produced by coherent illumination sources located or focused very close to the material. One of the important properties of speckle is the *memory effect*: small translations of the illuminator produce shifted, highly-correlated speckle patterns. In the figure, we used our algorithm to simulate light scattering in a medium with realistic tissue parameters, at thicknesses of $50\text{ }\mu\text{m}$, $250\text{ }\mu\text{m}$, $500\text{ }\mu\text{m}$, and a mean free path of $50\text{ }\mu\text{m}$, equivalent to optical depths (OD) of 1, 5 and 10, respectively. In each case the figure presents two speckle images obtained with a small shift of the input illumination, leading to correlated shifted speckle patterns (e.g., the shift is visible in the insets of the first two columns). As expected, when the thickness increases, more scattering is present, and thus the memory-effect correlation becomes weaker (the correlation is less visible in the inset of the third column).

We introduce rendering algorithms for the simulation of speckle statistics observed in scattering media under coherent near-field imaging conditions. Our work is motivated by the recent proliferation of techniques that use speckle correlations for tissue imaging applications: The ability to simulate the image measurements used by these speckle imaging techniques in a physically-accurate and computationally-efficient way can facilitate the widespread adoption and improvement of these techniques. To this end, we draw inspiration from recently-introduced Monte Carlo algorithms for rendering speckle statistics under far-field conditions (collimated sensor and illumination). We derive variants of these algorithms that are better suited to the near-field conditions (focused sensor and illumination) required by tissue imaging applications. Our approach is based on using Gaussian apodization to approximate the sensor and illumination aperture, as well as von Mises-Fisher functions to approximate the phase function of the scattering material. We show that these approximations allow us to derive closed-form expressions for the focusing operations involved in simulating near-field speckle patterns. As we demonstrate in our experiments, these

approximations accelerate speckle rendering simulations by a few orders of magnitude compared to previous techniques, at the cost of negligible bias. We validate the accuracy of our algorithms by reproducing ground truth speckle statistics simulated using wave-optics solvers, and real-material measurements available in the literature. Finally, we use our algorithms to simulate biomedical imaging techniques for focusing through tissue.

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

Additional Key Words and Phrases: Monte Carlo rendering, speckle, memory effect

ACM Reference Format:

Chen Bar, Ioannis Gkioulekas, and Anat Levin. 2020. Rendering Near-Field Speckle Statistics in Scattering Media. *ACM Trans. Graph.* 39, 6, Article 187 (December 2020), 26 pages. <https://doi.org/10.1145/3414685.3417813>

1 INTRODUCTION

A core challenge in medicine is the development of technologies for imaging deep inside biological tissues at high spatial resolutions. What makes this type of imaging possible is the fact that, when a light source illuminates tissue, a significant amount of light enters, travels inside, and re-emerges out of the tissue. What makes this type of imaging difficult is the fact that, when inside the tissue, light *scatters* multiple times. Thus, the fundamental challenge that needs to be solved to enable imaging inside tissue is inverting the multiple scattering process, in order to extract the information that light carries about the tissue it interacted with.

Authors' addresses: Chen Bar, Department of Electrical Engineering, Technion, Haifa, Israel; Ioannis Gkioulekas, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA; Anat Levin, Department of Electrical Engineering, Technion, Haifa, Israel.

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.

0730-0301/2020/12-ART187 \$15.00

<https://doi.org/10.1145/3414685.3417813>

In recent years, several imaging techniques have emerged that address this challenge by taking advantage of the fact that images of tissue under coherent (e.g., laser) illumination contain significant *speckles*: These are pseudo-random, high-frequency spatial variations in the intensity of the captured images. The statistical properties of these speckle patterns is a classical research area within optics, having been the subject of several textbooks [Erf 1978; Goodman 2007; Jacquot and Fournier 2000; Kaufmann 2011]. This research has revealed that, despite its random appearance, a speckle has strong statistical properties that provide rich information about the underlying scattering material (e.g., tissue). Perhaps best known among them is the *memory effect* (ME) property, illustrated in Fig. 1, which describes how speckle fields remain correlated under small changes in imaging conditions. The memory effect is at the core of speckle-based techniques for tissue imaging applications such as fluorescence imaging and adaptive optics focusing inside tissue.

Unfortunately, most previous studies of speckle statistical properties and of the memory effect, have been performed under imaging conditions that are not suitable for tissue imaging. Typically, most studies assume that both the light sources and the sensors are outside and at a large distance from the scattering volume, a set of conditions referred to as *far-field* imaging. By contrast, tissue imaging applications require both sources and sensors to be focused very close to the tissue (e.g., confocal microscopy), or even located inside it (e.g., fluorescent particles), a set of conditions known as *near-field* imaging. Far-field imaging simplifies analysis, simulation, and experiments relating to speckles. However, inferences drawn for the far-field case do not necessarily generalize to the near-field one. Our goal is to develop physically-accurate and efficient rendering algorithms that can help improve our understanding of speckle statistical properties under near-field conditions.

For this, we draw inspiration from Bar et al. [2019], who introduced a Monte Carlo framework for simulating speckle correlations in a way that combines physical accuracy and computational efficiency. Unfortunately, despite offering orders-of-magnitude acceleration compared to previous physically-accurate simulation techniques (e.g., wave equation solvers), the algorithms of Bar et al. [2019] are primarily intended for simulating far-field speckle statistics, and remain impractical for the near-field case. For example, as we show in our experiments, using their framework to simulate current techniques for focusing inside scattering [Judkewitz et al. 2014] can take several days on a large cluster. These performance characteristics significantly constrain the scope of investigations that can be performed using these algorithms (e.g., evaluation of effectiveness of existing techniques under different imaging parameters or for tissue samples of different optical parameters).

With these considerations in mind, we develop a computationally-efficient algorithm for simulating near-field (focused) speckle statistics inside scattering media. To this end, we extend the Monte Carlo rendering algorithm of Bar et al. [2019] in several ways that make it better-suited for near-field simulations. Our innovations are three-fold: First, we derive a path-integral expression for near-field speckle statistics. Second, we approximate optical apertures and material phase functions using von Mises-Fisher functions, to obtain analytical expressions for connecting paths traced inside a volume to near-field sources and sensors. Third, we develop an

importance sampling scheme for starting the volume path tracing process that takes into account the focused beams inside the medium. In our experiments, we validate the physical accuracy of our algorithms by showing that they can reproduce speckle correlation statistics simulated using wave-equation solvers, as well as speckle correlation measurements of real materials that are publicly available in the literature [Osnabrugge et al. 2017]. Our algorithms match this synthetic and measured groundtruth more accurately than existing simulation techniques (e.g., multi-slice layer based algorithms [Schott et al. 2015]) and analytical models (e.g., tilt-shift memory effect [Osnabrugge et al. 2017]) from optics. Additionally, we show that our algorithms allow us to simulate focusing-inside-tissue applications, which require near-field conditions. For the small volumes we simulate, our algorithm is already 3000× faster than an adaptation of the far-field algorithm [Bar et al. 2019]; and this performance difference will only increase for real-sized volumes. We believe that the ability to simulate near-field speckle statistics outside the lab will accelerate ongoing research on speckle-based techniques for tissue imaging applications. To fortify this effort, we have made our implementation publicly available [Bar et al. 2020].

1.1 Why render near-field speckle statistics?

Speckle statistics have strong potential for applications in the context of tissue imaging, where scattering by cells and other variations of the local index of refraction in the tissue drastically degrade image contrast. For example, several papers have suggested using speckle correlations to detect incoherent fluorescence sources inside the tissue. As Katz et al. [2014] have observed, due to the memory effect, the auto-correlation of random speckle images, together with a phase retrieval algorithm, can help remove the effect of scattering and reveal the location of the sources under the skin. Unfortunately, this idea has been successfully demonstrated mostly in the far-field setting, with the sources located at a large distance outside the scattering medium, rather than inside it. The only successful demonstration of this idea in the near-field we are aware of is by Chang et al. [2018], whose experiments were able to recover fluorescent particles spanning a small spatial range of 10 μm . We argue that this state of affairs is due to the limited exploration of near-field speckle statistics, an issue we hope our paper will help address.

Another important application of speckle techniques in tissue imaging is the use of adaptive optics [Mertz et al. 2015] to focus light at points deep inside tissue. Achieving this type of focusing requires using a coherent wavefront of a shape specific to the tissue sample being imaged. Determining the exact wavefront is challenging, and typically involves using external information or a guiding star [Horstmeyer et al. 2015]. Once this wavefront is found, the memory effect can be used to scan an area inside tissue, e.g., by shifting and tilting the wavefront to focus at neighboring points [Osnabrugge et al. 2017]. In our experiments, we simulate this approach, and show preliminary investigations on the effectiveness of using the tilt-shift memory effect. By enabling researchers to perform such investigations in simulation, without the need for lab experiments, we hope that our paper can help expand the scanning range and operational capabilities of techniques for focusing inside tissue.

Furthermore, the ability to efficiently render speckle patterns can facilitate the widespread adoption of data-driven approaches in

tissue imaging. The use of such approaches is in part motivated by analogous successes in the far-field case, where it has been demonstrated that machine learning algorithms can improve the performance of memory-effect-based imaging around the corner and through scattering [Li et al. 2018; Metzler et al. 2020, 2018]. The successful deployment of machine learning algorithms requires large, physically-accurate datasets. Previously, the lack of physically-accurate simulation tools meant that datasets had to be collected using lab measurements, an approach that is not scalable. We hope that our rendering tools can help reduce the data collection overhead, making machine learning approaches tractable.

Last but not least, accurate speckle rendering algorithms can be useful for *inverse rendering* problems involving speckle measurements. In particular, prior work has shown that differentiable rendering techniques can be used to recover accurate scattering parameters of real-world materials from incoherent intensity measurements [Gkioulekas et al. 2013]. We expect that our algorithms can be combined with modern differentiable rendering techniques [Nimier-David et al. 2020; Zhang et al. 2020, 2019]. In turn, this has the potential to enable recovering accurate high-resolution models of important material classes, such as biological tissue, from measurements of speckle fields and speckle correlations.

1.2 Limitations

Our algorithms are subject to a few limitations that suggest important directions for future research. First, they assume that the simulated scattering medium is homogeneous, meaning that its optical scattering parameters are the same at all spatial locations inside the volume. Our algorithms additionally assume that the medium is exponential, meaning that it comprises uncorrelated scatterers [Bitterli et al. 2018; d'Eon 2018; Jarabo et al. 2018]; and scatters light isotropically, meaning that its phase function is rotationally-invariant [Jakob et al. 2010]. These assumptions are commonly used in biomedical imaging research and applications to approximate how biological tissues scatter light at optical and infrared wavelengths. Therefore, our algorithms can be used to simulate imaging applications involving such materials.

Second, our algorithms are primarily geared towards *transmission mode* imaging configurations, where illumination and sensing happen at opposite sides of a scattering volume. This is a consequence of our use of von Mises-Fisher functions to approximate the imaging aperture and material phase function. When using these approximations to simulate *reflection mode* imaging configurations, where illumination and sensing are on the same side of a scattering volume, accuracy will depend on the exact material parameters. In particular, when simulating materials with phase functions that have very little back-scattering, accuracy suffers because of numerical issues. We provide a detailed analysis in App. A.4. In practical terms, this means that our algorithms can be used to simulate tissue imaging applications such as fluorescence imaging and focusing through tissue, both of which correspond to transmissive configurations; but should be used with caution for applications such as confocal microscopy, which corresponds to a reflective configuration.

Third, our algorithms become inefficient for very optically-deep (e.g., more than 10 times the mean free path) volumes, where light is expected to perform a large number of scattering events. This

limitation is due to the fact that our algorithms use volume path tracing to sample light paths, and is thus shared with standard Monte Carlo volume rendering algorithms for simulating incoherent intensity [Novak et al. 2018]. As in the intensity case, techniques based on the diffusion approximation [Jensen et al. 2001] would be better-suited for simulating speckle statistics at volumes of larger optical depths. However, we note that near-field speckle correlations decay as a function of optical depth much faster than intensity does; and are essentially non-existent for volumes of optical depths significantly larger than what we simulate. Consequently, biomedical imaging techniques based on speckle correlations are typically only applicable for optical depths below the diffusive regime, coinciding with the range of optical depths we emphasize in our experiments. Considering that neither single-scattering [Narasimhan et al. 2006] nor diffusion approximations [Feng et al. 1988] are effective in this range, our algorithms provide a suitable simulation tool for researchers investigating these biomedical imaging techniques.

2 RELATED WORK

Speckle in computational imaging. Speckle statistics have found wide applicability in computational imaging. Example applications include motion tracking [Jacquot and Rastogi 1979; Jakobsen et al. 2012; Smith et al. 2017], looking around the corner [Batarseh et al. 2018; Freund 1990; Katz et al. 2012], and seeing through [Abookasis and Rosen 2004; Bertolotti et al. 2012; Katz et al. 2014; Rosen and Abookasis 2003; Takasaki and Fleischer 2014] or focusing through [Choi et al. 2011; Edrei and Scarcelli 2016; Katz et al. 2010, 2012; Lai et al. 2015; Mosk et al. 2012; Nixon et al. 2013; Rueckel et al. 2006; van Putten et al. 2011; Vellekoop and Aegerter 2010; Vellekoop et al. 2012, 2010; Vellekoop and Mosk 2007; Yaqoob et al. 2008] tissue and other scattering layers. Most of these imaging techniques rely on the *memory effect* of speckles, and therefore are based on *spatial correlations* between speckle images. Alternatively, imaging techniques such as diffusing wave spectroscopy [Pine et al. 1988], laser speckle contrast imaging [Boas and Yodh 1997], and dynamic light scattering [Goldburg 1999] use *temporal speckle correlations* [Berne and Pecora 2000; Dougherty et al. 1994] to estimate flow (e.g., blood flow [Durduran et al. 2010]) and liquid composition parameters.

Analytical models for speckle statistics. Quantifying differences in speckle characteristics between the near-field and far-field cases can be done using analytical tools that approximate speckle statistics with closed-form mathematical expressions. Most available such tools are for the far-field case [Akkermans and Montambaux 2007; Baydoun et al. 2016; Berkovits and Feng 1994; Dougherty et al. 1994; Feng et al. 1988; Freund and Eliyahu 1992; Fried 1982], though recently tools have been introduced for the near-field case [Judkewitz et al. 2014; Osnabrugge et al. 2017]. For example, Osnabrugge et al. [2017] derived a tilt-shift model for the memory effect, assuming that the underlying scattering material is optically thin and very forward-scattering. Deriving closed-form expressions requires restrictive assumptions (single scattering, diffusion, or Fokker-Planck limits), which typically do not apply to the predominantly-turbid

tissue scattering. Developing efficient and physically-accurate rendering tools for near-field speckle statistics can help assess the accuracy of these approximations, and facilitate their broader adoption in application regimes where they are accurate.

Rendering wave-optics effects. Monte Carlo rendering techniques within computer graphics have typically focused on simulating incoherent light transport. More recently, there have been a few works on simulating wave optics effects, such as diffraction and speckle due to rough surface geometry [Bergmann et al. 2016; Cuypers et al. 2012; Sadeghi et al. 2012; Stam 1999; Sur et al. 2018; Werner et al. 2017; Yan et al. 2018; Yeh et al. 2013]. Variants of volumetric path tracing for simulating speckle in scattering have appeared in optics [Lu et al. 2004; Mout et al. 2016; Pan et al. 1995; Sawicki et al. 2008; Schmitt and Knüttel 1997; Xu 2004], though these typically cannot estimate second-order statistics (e.g., the memory effect). Bar et al. [2019] addressed this shortcoming by introducing a Monte Carlo algorithm that uses simultaneous path connections to multiple light sources and sensors, to correctly model speckle correlations between them. Unfortunately, their algorithm is primarily tailored to far-field imaging and becomes inefficient for the near-field case, as discussed in Sec. 3. The difference between the two cases is reminiscent of the challenges in rendering depth-of-field effects in incoherent light transport, where special rendering algorithms are required due to the need to sample many rays on the aperture plane [Barsky and Kosloff 2008; Kolb et al. 1995; Soler et al. 2009].

3 BACKGROUND ON SPECKLE STATISTICS

We begin by providing background on speckle statistics and on their Monte Carlo modeling for the far-field case. We use this background in Sec. 4 to contrast the far-field and near-field cases, which helps highlight the challenges involved in rendering the latter case, and positions our work relative to the prior art of Bar et al. [2019].

Notation and setting. We use bold letters for vectors (e.g., points $\mathbf{o}, \mathbf{i}, \mathbf{v}$), with a circumflex for unit vectors (e.g., directions $\hat{\omega}, \hat{\mathbf{i}}, \hat{\mathbf{v}}$).

We consider scattering volumes $\mathcal{V} \in \mathbb{R}^3$ that satisfy the assumptions underlying classical radiative transfer for isotropic [Jakob et al. 2010] exponential media [Bitterli et al. 2018; d'Eon 2018; Jarabo et al. 2018]: Each volume comprises a set of scatterers, whose locations in the volume are statistically independent. These scatterers are assumed to be small enough relative to the wavelength of light to be considered infinitesimal points. They are also assumed to be spherically symmetric, and thus scatter incident light waves in a rotationally-invariant way. We model speckle fields arising in such volumes due to incident illumination that we assume to be *monochromatic, fully-coherent* and *unpolarized*. These fields are a function of the volumes' *bulk properties*, which we describe next.

Bulk material properties. We use a *statistical description* of the optical properties of scattering volumes. In particular, the *scattering and absorption coefficients* σ_s and σ_a model, respectively, the portion of energy that is scattered and absorbed upon interaction with a scatterer. Their sum is the *extinction coefficient* $\sigma_t \equiv \sigma_a + \sigma_s$, and its inverse is the *mean free path*, $MFP = 1/\sigma_t$, which is the average distance in the volume light travels between two scattering events. Given a volume \mathcal{V} , it is common to express its geometric dimensions

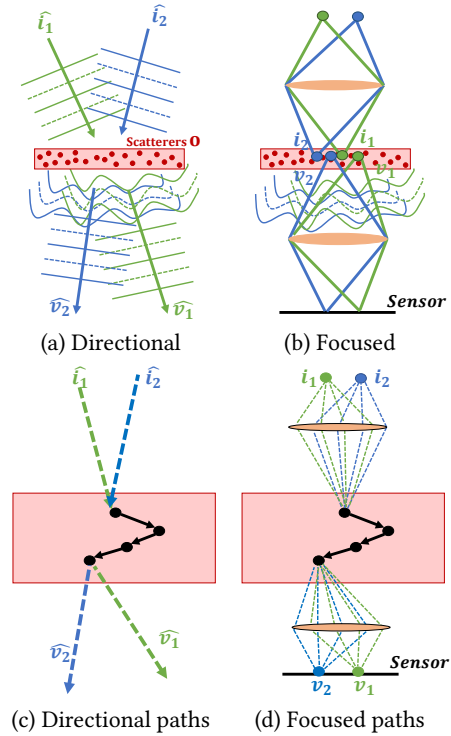


Fig. 2. **Transmissive imaging configurations.** Previous work [Bar et al. 2019] simulated speckle statistics by directional sources and sensors, as in (a). In contrast this work considers speckles by focused camera and sources, as in (b). (c-d) demonstrate what is implied when evaluating such speckle correlations using a Monte Carlo algorithm. In the directional formulation, one draws 4 directional connections from each path toward the illumination and viewing directions. In contrast, to simulate speckles through a focused lens, one needs to trace all paths via the aperture.

relative to *MFP*. For example, a volume has *optical depth* $OD = 2$ if its thickness is equal to $2 \cdot MFP$, meaning that light travelling through the volume undergoes on average two scattering events.

The *scattering amplitude function* $s(\cos \theta)$ describes how a field interacts with a scatterer: if a scatterer is illuminated from direction $\hat{\mathbf{i}}$, the complex scattered field u at direction $\hat{\mathbf{v}}$ is $u_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}} = s(\hat{\mathbf{i}} \cdot \hat{\mathbf{v}})$. The *phase function* is defined as $\rho(\cos \theta) \equiv |s(\cos \theta)|^2$. It is commonly characterized by an anisotropy parameter $-1 \leq g \leq 1$, equal to the average $\cos \theta$: $g = 0$ corresponds to scattering equally in all directions, and $g = 1$ to fully forward scattering. Tissue is characterized by very forward scattering ($g > 0.9$) [Cheong et al. 1990].

These parameters are a function of wavelength, and the scatterers' shape, size, and refractive index. For spherical scatterers, they can be computed using Mie theory [Bohren and Huffman 1983; Frisvad et al. 2007]. The three coefficients also depend on the *density* ζ , equal to the expected number of scatterers in a unit volume. We assume that scattering volumes are spatially *homogeneous*, meaning that scatterers are uniformly distributed, or equivalently, that the bulk parameters are the same everywhere inside a volume.

3.1 Modeling and rendering far-field speckle statistics

Transmissive far-field imaging. We focus on the geometry illustrated in Fig. 2(a): Scatterers are placed at a *configuration* of locations

$O = \{\mathbf{o}_1, \mathbf{o}_2, \dots\}$ inside the volume \mathcal{V} , each sampled *independently* from the others, using the volume density ς . This configuration is imaged using light sources and sensors that are on opposite sides of the volume, a setting we refer to as *transmission mode imaging*. Additionally, in this background section, we assume that the volume is illuminated by a *directional* plane wave $\hat{\mathbf{i}}$, and imaged with a *directional* sensor $\hat{\mathbf{v}}$. We refer to these conditions together as the *far-field* imaging conditions. If we know the exact scatterer locations, and incoming and outgoing directions, we can solve the wave equation to obtain the complex-valued *scattered field* $u_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}, O}$, arising from the interaction of the incident illumination with the scattering volume.

Defining speckle statistics. For any volume with a given scatterer configuration O , the scattered field typically contains large fluctuations with a semi-random noise structure known as *speckle* (see, e.g., Fig. 1). We can characterize speckle using the first and second-order statistics of fields due to different volumes *with the same bulk material properties*. In particular, we can define the *speckle mean*,

$$m_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}} \equiv E_O \left[u_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}, O} \right], \quad (1)$$

and the *speckle covariance*,

$$C_{\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2} \equiv E_O \left[u_{\hat{\mathbf{v}}^1}^{\hat{\mathbf{i}}^1, O} \cdot u_{\hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^2, O*} \right] - m_{\hat{\mathbf{v}}^1}^{\hat{\mathbf{i}}^1} \cdot m_{\hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^2*}, \quad (2)$$

where $(\cdot)^*$ is complex conjugation. $u_{\hat{\mathbf{v}}^1}^{\hat{\mathbf{i}}^1, O}$, $u_{\hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^2, O}$ are two speckle fields generated by the *same* scatterer configuration O , when illuminated by two *monochromatic, mutually-coherent* incident waves from $\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2$, and measured at two sensors $\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2$. The expectation E_O is taken with respect to all scatterer configurations O sampled from the same density ς . As we detail in App. A.1, the speckle mean $m_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}}$ can be computed using a closed-form expression and is typically zero. Therefore, we focus on modeling the speckle covariance $C_{\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2}$.

The definition of Eq. (2) suggests a straightforward approach for computing this covariance: randomly sample many scatterer configurations O from the material bulk parameters, solve the wave equation numerically to compute $u_{\hat{\mathbf{v}}^1}^{\hat{\mathbf{i}}^1, O}$, $u_{\hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^2, O}$, and use averaging to approximate the expectation in Eq. (2). Unfortunately, while exact wave-equation solvers exist [Thierry et al. 2015; Treeby and Cox. 2010; Yee 1966], their computational complexity is prohibitive, typically making them intractable for volumes of width larger than a few dozen wavelengths. This computational cost is further exacerbated by the need to use the solvers multiple times for averaging.

For an alternative, note first that, when $\hat{\mathbf{i}}^1 = \hat{\mathbf{i}}^2 = \hat{\mathbf{i}}$, $\hat{\mathbf{v}}^1 = \hat{\mathbf{v}}^2 = \hat{\mathbf{v}}$ and $m_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}} = 0$, $C_{\hat{\mathbf{v}}, \hat{\mathbf{v}}}^{\hat{\mathbf{i}}, \hat{\mathbf{i}}}$ reduces to the *intensity* $I_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}}$. This intensity is typically modeled in computer graphics using the radiative transfer equation, or its integral form, the volume rendering equation. The latter gives rise to Monte Carlo volume rendering algorithms, which compute intensity using as input only the volume's bulk material properties [Novak et al. 2018]. Bar et al. [2019] derived analogous Monte Carlo volume rendering algorithms for computing the speckle covariance $C_{\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2}$ for any directions $\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2, \hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2$. Like their intensity counterparts, these algorithms take as input bulk material parameters, and not particle positions. Bar et al. [2019] showed that their approach is physically accurate, orders-of-magnitude faster

than wave-equation solvers, and scalable to much larger volumes. We proceed to review this Monte Carlo rendering approach.

Covariance rendering. Bar et al. [2019] derive their algorithm from a *path integral* expression for speckle covariance, obtained by considering the correlation of fields that travel along all possible *pairs* of paths from $\hat{\mathbf{i}}^1$ to $\hat{\mathbf{v}}^1$, and from $\hat{\mathbf{i}}^2$ to $\hat{\mathbf{v}}^2$. They showed that this expression can be simplified to use only pairs of paths that coincide everywhere, except for their connections to $\hat{\mathbf{i}}^1, \hat{\mathbf{v}}^1, \hat{\mathbf{i}}^2, \hat{\mathbf{v}}^2$. We review this simplified formulation. Consider the space \mathbb{P} of *sub-paths* $\vec{\mathbf{x}}^s = \mathbf{o}_1 \rightarrow \dots \rightarrow \mathbf{o}_B$, $B \geq 1$, where each vertex $\mathbf{o}_b \in \mathcal{V}$; we denote by $\hat{\omega}_b \equiv \overline{\mathbf{o}_b, \mathbf{o}_{b+1}}$ the direction of the b -th edge of the sub-path. These vertices correspond to the shared part of two full paths $\vec{\mathbf{x}}^1 = \hat{\mathbf{i}}^1 \rightarrow \mathbf{o}_1 \rightarrow \dots \rightarrow \mathbf{o}_B \rightarrow \hat{\mathbf{v}}^1$, $\vec{\mathbf{x}}^2 = \hat{\mathbf{i}}^2 \rightarrow \mathbf{o}_1 \rightarrow \dots \rightarrow \mathbf{o}_B \rightarrow \hat{\mathbf{v}}^2$, formed by connecting the sub-path to $\hat{\mathbf{i}}^1, \hat{\mathbf{v}}^1$ and $\hat{\mathbf{i}}^2, \hat{\mathbf{v}}^2$. Then, the speckle covariance of Eq. (2) can be expressed as:

$$C_{\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2} = \int_{\mathbb{P}} c_{\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2}(\vec{\mathbf{x}}^s) d\vec{\mathbf{x}}^s, \quad (3)$$

where the far-field *path contribution function* $c_{\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2}$ equals the correlation of the fields that travel along $\vec{\mathbf{x}}^1, \vec{\mathbf{x}}^2$. For $B \geq 2$, this equals:

$$\begin{aligned} c_{\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2}(\vec{\mathbf{x}}^s) = & f(\vec{\mathbf{x}}^s) \cdot v(\hat{\mathbf{i}}^1 \rightarrow \mathbf{o}_1) s(\hat{\mathbf{i}}^1 \cdot \hat{\omega}_1) \\ & \cdot v(\mathbf{o}_B \rightarrow \hat{\mathbf{v}}^1) s(\hat{\omega}_{B-1} \cdot \hat{\mathbf{v}}^1) \\ & \cdot v(\hat{\mathbf{i}}^2 \rightarrow \mathbf{o}_1)^* s(\hat{\mathbf{i}}^2 \cdot \hat{\omega}_1)^* \\ & \cdot v(\mathbf{o}_B \rightarrow \hat{\mathbf{v}}^2)^* s(\hat{\omega}_{B-1} \cdot \hat{\mathbf{v}}^2)^*, \end{aligned} \quad (4)$$

and for $B = 1$:

$$\begin{aligned} c_{\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2}(\vec{\mathbf{x}}^s) = & \sigma_s \cdot v(\hat{\mathbf{i}}^1 \rightarrow \mathbf{o}_1) v(\mathbf{o}_1 \rightarrow \hat{\mathbf{v}}^1) s(\hat{\mathbf{i}}^1 \cdot \hat{\mathbf{v}}^1) \\ & \cdot v(\hat{\mathbf{i}}^2 \rightarrow \mathbf{o}_1)^* v(\mathbf{o}_1 \rightarrow \hat{\mathbf{v}}^2)^* s(\hat{\mathbf{i}}^2 \cdot \hat{\mathbf{v}}^2)^*. \end{aligned} \quad (5)$$

In the above, $f(\vec{\mathbf{x}}^s)$ is the standard *radiometric throughput* of $\vec{\mathbf{x}}^s$, augmented by scattering coefficients at the first and last vertex,

$$f(\vec{\mathbf{x}}^s) = (\sigma_s)^B \prod_{b=1}^{B-1} \frac{e^{-\sigma_t \|\mathbf{o}_{b+1} - \mathbf{o}_b\|}}{\|\mathbf{o}_{b+1} - \mathbf{o}_b\|^2} \rho(\hat{\omega}_{b-1} \cdot \hat{\omega}_b). \quad (6)$$

Finally, $v(\cdot)$ is the *complex volumetric throughput*, defined as:

$$v(\hat{\omega} \rightarrow \mathbf{o}) = e^{-\frac{1}{2} \sigma_t d(\hat{\omega} \rightarrow \mathbf{o})} e^{ik(\hat{\omega} \cdot \mathbf{o})}, \quad (7)$$

$$v(\mathbf{o} \rightarrow \hat{\omega}) = e^{-\frac{1}{2} \sigma_t d(\mathbf{o} \rightarrow \hat{\omega})} e^{-ik(\hat{\omega} \cdot \mathbf{o})}, \quad (8)$$

where $k \equiv 2\pi/\lambda$ is the wavenumber and λ the wavelength of the illumination; and $d(\hat{\omega} \rightarrow \mathbf{o})$, $d(\mathbf{o} \rightarrow \hat{\omega})$ denote the distance a ray entering or leaving, respectively, \mathbf{o} at direction $\hat{\omega}$, travels inside the scattering volume \mathcal{V} . Fig. 2(c) visualizes these terms.

The covariance rendering algorithm of Bar et al. [2019] uses a Monte Carlo path sampling approach to evaluate the speckle covariance integral of Eq. (3). This algorithm takes advantage of the presence of the radiometric throughput term in Eq. (4), and samples sub-paths $\vec{\mathbf{x}}^s$ using standard volumetric path tracing. Then, for each sampled sub-path, the endpoints $\mathbf{o}_1, \mathbf{o}_B$ are connected to the far-field illuminations $\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2$ and sensors $\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2$, to compute the complex volumetric throughput terms in Eqs. (4) and (5).

We note that when $\hat{\mathbf{i}}^1 = \hat{\mathbf{i}}^2 = \hat{\mathbf{i}}$, $\hat{\mathbf{v}}^1 = \hat{\mathbf{v}}^2 = \hat{\mathbf{v}}$, the above algorithm becomes equivalent to the standard volumetric path tracing algorithm used in computer graphics to render intensity $I_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}}$ [Novak et al. 2018]. A key observation by Bar et al. [2019] is that, in the case of covariance $C_{\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2}$, each sampled sub-path $\bar{\mathbf{x}}^s$ needs to contribute to two pairs of input-output directions (blue and green connections in Fig. 2(c)). Using the same sub-paths for both $(\hat{\mathbf{i}}^1, \hat{\mathbf{v}}^1)$, $(\hat{\mathbf{i}}^2, \hat{\mathbf{v}}^2)$ pairs is necessary to account for the correlation of fields from different illuminator-viewpoint combinations, and this correlation cannot be modeled using paths sampled independently for each pair.

Before concluding this section, we mention that Bar et al. [2019] present a variant of the above-described algorithm that simulates scattered fields $u_{\hat{\mathbf{v}}^1}^{\hat{\mathbf{i}}^1}$, $u_{\hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^2}$ in a way that accurately reproduces their correlations. In the rest of the paper, we show how to adapt both of these rendering algorithms to the near-field case. We present most of our theory in Secs. 4–6 in the context of covariance rendering, then adapt it in Sec. 7 for field rendering.

4 NEAR-FIELD SPECKLE STATISTICS

As mentioned in the introduction, many important applications require imaging speckle fields using a sensor that is focused at some point, rather than being at infinity; and using illumination that is better modeled as a point than as a directional source. Both of these points can be near or even inside the scattering volume. To distinguish them from the far-field case of the previous section, we refer to these conditions together as the *near-field* imaging conditions. Scattered fields formed under far-field and near-field conditions often have very different characteristics, as we discuss in the rest of the section. As in Sec. 3.1, we focus on transmission mode imaging configurations, where illumination and sensing are on opposite sides of the volume. The imaging configuration is shown in Fig. 2(b).

Modeling near-field speckle statistics. We use the notation $u_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}}$, $u_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}}$ to indicate scattered fields due to far-field and near-field imaging conditions, respectively, with the circumflex distinguishing between directional to point sources/sensors. If we know the scattered field $u_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}}$ for all illumination and sensing directions $\hat{\mathbf{i}}$, $\hat{\mathbf{v}}$, we can compute the scattered field $u_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}}$ by integrating over the sensor and illumination apertures (Fig. 2(d)) [Goodman 1968; Mertz 2019]. Concretely:

$$u_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}} = \int_{\hat{\mathbf{i}} \in \mathbb{S}^2} \int_{\hat{\mathbf{v}} \in \mathbb{S}^2} a_{\hat{\mathbf{i}}}(\hat{\mathbf{i}}) a_{\hat{\mathbf{v}}}(\hat{\mathbf{v}}) u_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}} d\hat{\mathbf{v}} d\hat{\mathbf{i}}, \quad (9)$$

where \mathbb{S}^2 is the unit sphere and, assuming an ideal lens,

$$a_{\hat{\mathbf{v}}}(\hat{\mathbf{v}}) \equiv m(\hat{\mathbf{v}}) e^{ik(\hat{\mathbf{v}} \cdot \mathbf{v})}, \quad a_{\hat{\mathbf{i}}}(\hat{\mathbf{i}}) \equiv m(\hat{\mathbf{i}}) e^{-ik(\hat{\mathbf{i}} \cdot \mathbf{i})}. \quad (10)$$

The functions $m(\hat{\mathbf{i}})$ and $m(\hat{\mathbf{v}})$ denote illumination and viewing aperture amplitudes. Typically, these are binary functions indicating which directions pass through an aperture of a finite extent. We indicate the width of these masks using the angle Θ_{\max} between the optical axis and the propagating direction that most deviates from this axis, and we refer to $\sin \Theta_{\max}$ as the *numerical aperture* (NA) (Fig. 3(a)). The term $\exp(ik(\hat{\mathbf{v}} \cdot \mathbf{v}))$ is the phase accumulated in direction $\hat{\mathbf{v}}$ when focusing at the point \mathbf{v} ; and analogously for the term $\exp(-ik(\hat{\mathbf{i}} \cdot \mathbf{i}))$. Using paraxial optics approximations, it is also possible to express the two integrals of Eq. (9) as a double Fourier

transform with respect to the variables $\hat{\mathbf{i}}$ and $\hat{\mathbf{v}}$, and thus $u_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}}$ is often referred to as the *Fourier field* [Goodman 1968; Mertz 2019].

The near-field speckle covariance $C_{\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2}$ can be defined analogously to the far-field covariance in Eq. (2). Combining this definition with Eq. (9), we can relate near-field and far-field covariances:

$$C_{\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2} = \iiint a_{\hat{\mathbf{v}}^1}(\hat{\mathbf{v}}^1) a_{\hat{\mathbf{i}}^1}(\hat{\mathbf{i}}^1) a_{\hat{\mathbf{v}}^2}(\hat{\mathbf{v}}^2)^* a_{\hat{\mathbf{i}}^2}(\hat{\mathbf{i}}^2)^* C_{\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2} d\hat{\mathbf{i}}^1 d\hat{\mathbf{i}}^2 d\hat{\mathbf{v}}^1 d\hat{\mathbf{v}}^2. \quad (11)$$

Why is rendering near-field covariance difficult? Eq. (11) provides a conceptually simple way to compute near-field covariance statistics: We sample, or discretize, the domains of $\hat{\mathbf{v}}^1$, $\hat{\mathbf{v}}^2$, $\hat{\mathbf{i}}^1$, $\hat{\mathbf{i}}^2$, estimate the corresponding far-field covariance values as in Sec. 3.1, and form their weighted average using the weights in Eq. (11).

In practice, this approach is computationally impractical, because of the very large number of samples of the far-field covariance $C_{\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2}^{\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2}$ we need to compute. This number is determined by the width W of the volume \mathcal{V} and the numerical aperture Θ_{\max} , and not by the size of the sensor. As we show in App. A.2, Nyquist sampling rate implies the number of samples per axis of integration should be:

$$N = \frac{2W \sin(\Theta_{\max})}{\lambda}. \quad (12)$$

For some representative numbers, we consider simulating a relatively thin volume of thickness $400 \mu\text{m}$. We assume that we image a spatial area of size $100 \mu\text{m} \times 100 \mu\text{m}$ using a numerical aperture $\sin(\Theta_{\max}) = 0.5$ and illumination wavelength $\lambda = 0.5 \mu\text{m}$. We need to simulate a volume that is at least as wide as the defocused beam, as light contributing to the scattered field can be due to scattering anywhere inside the beam (see Fig. 3). For the dimensions mentioned, the defocus blur has size $400 \mu\text{m}$ at the far edge of the target, suggesting that we need to simulate a volume of width $W = 100 \mu\text{m} + 400 \mu\text{m} = 500 \mu\text{m}$. Then Eq. (12) suggests using $N = 1000$ samples on each coordinate axis of the four directions we integrate in Eq. (11). Therefore, to simulate the near-field covariance, we need $(1000)^8$ samples, which is prohibitively large. We note that real near-field scenes used in, e.g., Osnabrügge et al. [2017] are bigger and would require an even more far-field samples.

Our discussion so far has been about covariance rendering, but similar arguments apply for field rendering, where using Eq. (9) would require N^4 samples, or $(1000)^4$ for the above example. We discuss the field rendering case in detail in Sec. 7, but for demonstration, we compare in Fig. 3 speckle images rendered by our proposed field rendering algorithm and using Eq. (9). Due to memory constraints, the far-field directions were sampled at 10% of the Nyquist sampling rate; this aliased sampling results in clear replica artifacts in the corresponding images. Even under these aliased sampling conditions, rendering using Eq. (9) was $100\times$ slower than our algorithm, while requiring 30 GB GPU memory. These challenges cannot be alleviated by using a large number of samples for single scattering and fewer samples for multiple scattering, as is often done in intensity rendering [Belcour et al. 2014]: even though intensity images from multiple scattering have low spatial frequency, coherent multiple scattering still creates high-frequency speckle patterns.

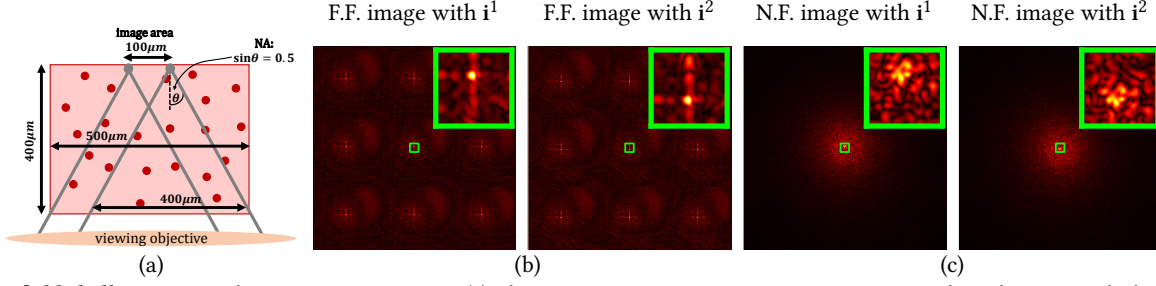


Fig. 3. **Near-field challenges.** Consider an imaging setup as in (a) where one wants to image a $100 \mu\text{m} \times 100 \mu\text{m}$ area through a $400 \mu\text{m}$ thick tissue with a $0.5NA$ objective. The resulting imaging cone is $400 \mu\text{m}$ wide. As scattering can arise from anywhere inside the imaging cone, for realistic simulation the medium should be at least as wide as the defocus cone, requiring us to simulate a slab of wider than $500 \mu\text{m}$. Assuming e.g. $\lambda = 0.5 \mu\text{m}$, this results in $N = 1000$ samples in each axis. (b) Two speckle images generated by two nearby illuminators, rendered using far-field covariances, sampled at 10% of the Nyquist limit revealing aliasing. The simulation still runs $100\times$ slower than our near-field approach, and occupies as much as 30 GB GPU global memory. (c) Aliasing-free speckle images by our suggested near field approach. Note the ME shift demonstrated in the insets.

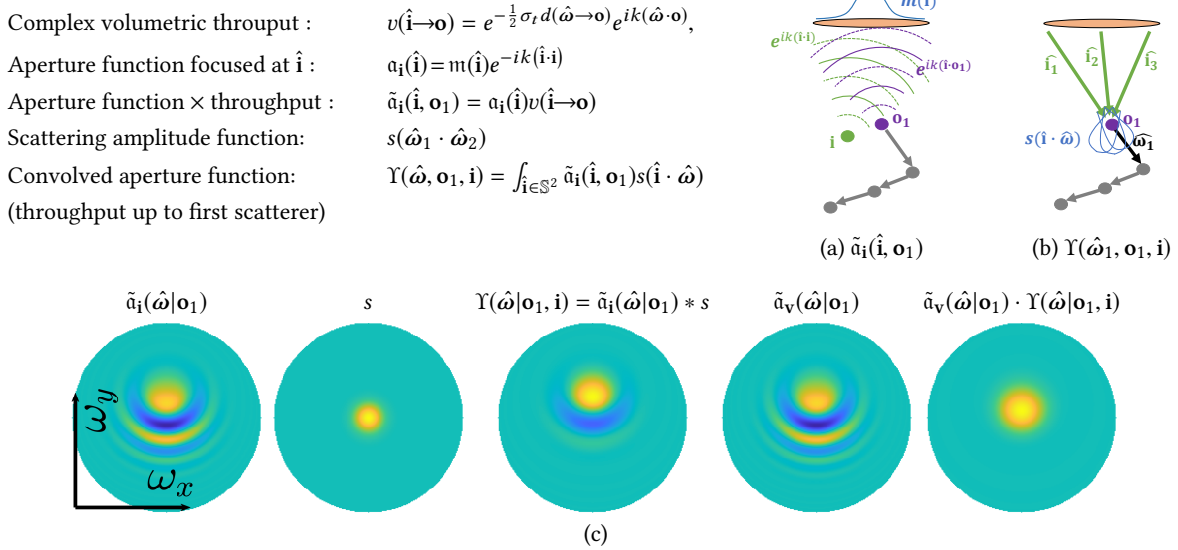


Fig. 4. **Notation summary and visualization.** (a) The aperture function $\tilde{a}_i(\hat{\mathbf{i}}, \mathbf{o}_1)$ includes an apodization mask $m(\hat{\mathbf{i}})$, a complex wave focusing at \mathbf{i} and a conjugate wave focusing at \mathbf{o}_1 , the first point on the sampled path (attenuation $\exp(-1/2\sigma_t d(\hat{\mathbf{o}} \rightarrow \mathbf{o}))$ is not visualized). (b) The aperture function convolved angularly with the scattering amplitude function leads to $\Upsilon(\hat{\omega}_1, \mathbf{o}_1, \mathbf{i})$, the throughput up to the first point and direction. That is, to compute the field propagating from \mathbf{o}_1 at direction $\hat{\omega}_1$ we integrate over all directions $\hat{\mathbf{i}}$ in the aperture. For each direction we consider the value of the aperture function in this direction, times the amount of energy scattered from $\hat{\mathbf{i}}$ to $\hat{\omega}_1$. (c) Visualizing the pipeline of the single-scattering covariance in Eq. (18). Illumination aperture \tilde{a}_i is convolved with scattering function s to generate $\Upsilon(\hat{\omega}_1, \mathbf{o}_1, \mathbf{i})$, which is then multiplied with the viewing aperture \tilde{a}_v . For visualization the real component of the involved spherical functions is projected onto the 2D $\hat{\omega}_x, \hat{\omega}_y$ plane (north hemisphere only).

4.1 Near-field covariance path integral

To overcome the computational challenges of evaluating near-field speckle covariance, we first derive for it a path-integral expression, which will absorb the directional integrations of Eq. (11) into the path contribution function. Then, in Sec. 5, we introduce an approximation that allows us to compute this path contribution function *analytically*. This completely removes the need for directional integration, drastically reducing computational complexity.

By combining Eq. (3) and Eq. (11), we can express the near-field covariance as a path integral on the same space of sub-paths \mathbb{P} :

$$C_{\mathbf{v}^1, \mathbf{v}^2}^{i^1, i^2} = \int_{\mathbb{P}} c_{\mathbf{v}^1, \mathbf{v}^2}^{i^1, i^2}(\vec{\mathbf{x}}^s) d\vec{\mathbf{x}}^s, \quad (13)$$

where now the contribution of each sub-path is determined by the *near-field* path contribution function $c_{\mathbf{v}^1, \mathbf{v}^2}^{i^1, i^2}$. For $B \geq 2$, this equals:

$$c_{\mathbf{v}^1, \mathbf{v}^2}^{i^1, i^2}(\vec{\mathbf{x}}^s) = f(\vec{\mathbf{x}}^s) \cdot \Upsilon(\hat{\omega}_1, \mathbf{o}_1, \mathbf{i}^1) \Upsilon(\hat{\omega}_{B-1}, \mathbf{o}_B, \mathbf{v}^1) \cdot \Upsilon(\hat{\omega}_1, \mathbf{o}_1, \mathbf{i}^2)^* \Upsilon(\hat{\omega}_{B-1}, \mathbf{o}_B, \mathbf{v}^2)^*, \quad (14)$$

where Υ denotes integration over the aperture of terms in Eq. (4):

$$\Upsilon(\hat{\omega}_1, \mathbf{o}_1, \mathbf{i}) \equiv \int_{\hat{\mathbf{i}} \in \mathbb{S}^2} \tilde{a}_i(\hat{\mathbf{i}}, \mathbf{o}_1) s(\hat{\mathbf{i}} \cdot \hat{\omega}_1) d\hat{\mathbf{i}}, \quad (15)$$

$$\Upsilon(\hat{\omega}_{B-1}, \mathbf{o}_B, \mathbf{v}) \equiv \int_{\hat{\mathbf{v}} \in \mathbb{S}^2} \tilde{a}_v(\hat{\mathbf{v}}, \mathbf{o}_B) s(\hat{\omega}_{B-1} \cdot \hat{\mathbf{v}}) d\hat{\mathbf{v}}, \quad (16)$$

and \tilde{a} denotes weighed aperture functions α :

$$\tilde{a}_i(\hat{\mathbf{i}}, \mathbf{o}_1) \equiv \alpha_i(\hat{\mathbf{i}}) v(\mathbf{i} \rightarrow \mathbf{o}_1), \quad \tilde{a}_v(\hat{\mathbf{v}}, \mathbf{o}_B) \equiv \alpha_v(\hat{\mathbf{v}}) v(\mathbf{o}_B \rightarrow \hat{\mathbf{v}}). \quad (17)$$

Similarly, for $B = 1$, the contribution function equals:

$$c_{\mathbf{v}^1, \mathbf{v}^2}^{i^1, i^2}(\vec{\mathbf{x}}^s) = \sigma_s(\mathbf{o}_1) \cdot \int_{\hat{\mathbf{v}}^1 \in \mathbb{S}^2} \tilde{a}_{\mathbf{v}^1}(\hat{\mathbf{v}}^1, \mathbf{o}_1) Y(\hat{\mathbf{v}}^1, \mathbf{o}_1, i^1) d\hat{\mathbf{v}}^1 \cdot \left(\int_{\hat{\mathbf{v}}^2 \in \mathbb{S}^2} \tilde{a}_{\mathbf{v}^2}(\hat{\mathbf{v}}^2, \mathbf{o}_1) Y(\hat{\mathbf{v}}^2, \mathbf{o}_1, i^2) d\hat{\mathbf{v}}^2 \right)^* \quad (18)$$

We use $\tilde{a}_i(\hat{\mathbf{i}}|\mathbf{o}_1)$, $Y(\hat{\omega}_1|\mathbf{o}_1, \mathbf{i})$ to denote versions of these functions with respect to only their first argument, conditioned on fixed values for their other arguments. These are *complex* functions on the unit sphere \mathbb{S}^2 . In particular, $Y(\hat{\omega}_1|\mathbf{o}_1, \mathbf{i})$ can be thought of as a *convolution* of the aperture function $\tilde{a}_i(\hat{\mathbf{i}}|\mathbf{o}_1)$ with the scattering amplitude function $s(\hat{\mathbf{i}} \cdot \hat{\omega})$. We summarize and visualize these terms in Fig. 4.

With Eqs. (14) and (18) at hand, conceptually we can compute the near-field covariance $c_{\mathbf{v}^1, \mathbf{v}^2}^{i^1, i^2}$ using a Monte Carlo rendering algorithm exactly analogous to the one proposed by Bar et al. [2019] for the far-field case: First, we sample sub-paths $\vec{\mathbf{x}}^s$ using standard volumetric path tracing. Second, we compute the path contribution function for each sampled path, and accumulate the results. We note however that, unlike the far-field case, in the near-field case the path contribution function cannot be computed analytically, as it requires spherical integration for evaluating Y . Theoretically, this could be done using a second-stage Monte Carlo integration procedure, by importance sampling one or more directions in the aperture plane. However, as we show experimentally in Sec. 8, because these integrals have complex integrands with highly-varying phases, Monte Carlo estimates have very high variance. Considering that this second-stage Monte Carlo integration procedure needs to be performed separately for each sampled sub-path $\vec{\mathbf{x}}^s$, it quickly results in an overwhelming computational overhead, making this overall rendering procedure intractable. In the next section, we sidestep this overhead by deriving closed-form approximations to these integration and convolution operations, which can be computed analytically without the need for Monte Carlo integration.

5 NEAR-FIELD USING VON MISES-FISHER FUNCTIONS

In this section, we present two main technical results. First, we show how to use mixtures of von Mises-Fisher functions to approximate the various spherical functions that appear in the near-field path contribution function $c_{\mathbf{v}^1, \mathbf{v}^2}^{i^1, i^2}(\vec{\mathbf{x}}^s)$ of Eqs. (14) and (18). Second, we show how this approximation allows us to derive closed-form expressions for the convolution and other integral terms in $c_{\mathbf{v}^1, \mathbf{v}^2}^{i^1, i^2}(\vec{\mathbf{x}}^s)$. Overall, this allows us to compute covariance contribution analytically for each sampled sub-path $\vec{\mathbf{x}}^s$, avoiding Monte Carlo integration.

Von Mises-Fisher functions. We begin with background on *complex* von Mises-Fisher functions, defined as [Mardia and Jupp 2000]:

$$h(\hat{\omega}) = \eta \cdot e^{(\boldsymbol{\mu} \cdot \hat{\omega})}, \quad (19)$$

where $\hat{\omega} \in \mathbb{S}^2$, and $\boldsymbol{\mu} = \boldsymbol{\mu}_r + i\boldsymbol{\mu}_i$ is a complex three-dimensional vector parameterizing the von Mises-Fisher function. We define:

$$\gamma_r = \|\boldsymbol{\mu}_r\|, \quad \gamma_i = \|\boldsymbol{\mu}_i\|, \quad \gamma = \|\boldsymbol{\mu}\| = \gamma_r - \gamma_i + 2i(\boldsymbol{\mu}_r \cdot \boldsymbol{\mu}_i), \quad (20)$$

$$\hat{\boldsymbol{\mu}} = \boldsymbol{\mu}/\gamma, \quad \hat{\boldsymbol{\mu}}_r = \boldsymbol{\mu}_r/\gamma_r, \quad \hat{\boldsymbol{\mu}}_i = \boldsymbol{\mu}_i/\gamma_i. \quad (21)$$

The value $|h(\hat{\omega})|$ is maximized when $\hat{\omega} = \hat{\boldsymbol{\mu}}$. The scale γ_r is inversely proportional to the support of the function: Large γ_r values result

in a narrow function that is sharply peaked around $\boldsymbol{\mu}_r$, whereas $\gamma_r = 0$ results in a uniform function over the sphere. Von Mises-Fisher functions have found use in various applications in computer graphics [Han et al. 2007], though we emphasize that, compared to this prior work, we use complex functions ($\gamma_i \neq 0$).

Von Mises-Fisher functions can be thought of as a generalization of the Gaussian distribution to the unit sphere, with variance $\sigma^2 = \frac{1}{\gamma}$ [Mardia and Jupp 2000]. Intuitively, we can see that this is the case by considering that, when we constrain $\|\hat{\boldsymbol{\mu}}\| = \|\hat{\omega}\| = 1$,

$$e^{-\frac{\gamma}{2} \|\hat{\omega} - \hat{\boldsymbol{\mu}}\|^2} \propto e^{-\gamma(\hat{\boldsymbol{\mu}} \cdot \hat{\omega})}, \quad (22)$$

where \propto denotes equality up to a multiplicative scale.

5.1 Working with von Mises-Fisher functions

Fitting with von Mises-Fisher functions. We aim to approximate the spherical functions that appear in the near-field path contribution function $c_{\mathbf{v}^1, \mathbf{v}^2}^{i^1, i^2}(\vec{\mathbf{x}}^s)$ using von Mises-Fisher functions.

We first consider the aperture function $\tilde{a}_i(\hat{\mathbf{i}})$, defined in Eqs. (10) and (17). We reproduce the definition here for convenience:

$$\tilde{a}_i(\hat{\mathbf{i}}) = m(\hat{\mathbf{i}}) e^{-ik(\hat{\mathbf{i}} \cdot \mathbf{i})} v(\mathbf{i} \rightarrow \mathbf{o}_1) = m(\hat{\mathbf{i}}) e^{-ik(\hat{\mathbf{i}} \cdot \mathbf{i}) - \frac{1}{2} \sigma_t d(\hat{\omega} \rightarrow \mathbf{o}) + ik(\hat{\omega} \cdot \mathbf{o})}. \quad (23)$$

We want to express $\tilde{a}_i(\hat{\mathbf{i}})$ as a von Mises-Fisher function. To achieve this, we first choose to approximate the aperture mask $m(\cdot)$ as a *real* von Mises-Fisher function. Denoting the optical axis of the system by $\hat{\boldsymbol{\mu}}_a$ (usually this is the z-axis $\hat{\boldsymbol{\mu}}_a = \hat{\mathbf{z}} = [0, 0, 1]$), we have,

$$m(\hat{\omega}) \approx e^{-\gamma_a} e^{\gamma_a(\hat{\boldsymbol{\mu}}_a \cdot \hat{\omega})}. \quad (24)$$

This approximation is a form of *apodization*: a binary aperture, which completely blocks or transmits fields propagating in different directions, is replaced by a non-binary mask, which attenuates the amplitude of transmitted fields by an increasing amount at larger propagation angles. Such non-binary apertures are typical of systems that use short-focal-length lenses with strong aberrations. On the illumination side, these non-binary apertures are also representative of the Gaussian profiles of laser beams. As we show in Sec. 8, even when the underlying aperture is binary, using the apodization of Eq. (24) produces accurate speckle statistics for transmissive imaging configurations. We note that the mean width of the non-binary aperture equals $\sigma_a = 1/\sqrt{\gamma_a}$, and should be set to match the width of the true binary aperture. Additionally, we can use $\gamma_a = 0$ to model isotropic point sources (e.g., fluorescent particles).

The phase terms in Eq. (23) already form a complex von Mises-Fisher function. Thus, to complete our treatment of $\tilde{a}_i(\hat{\mathbf{i}})$, we need to add the attenuation term, for which we assume that the attenuation is approximately constant over the aperture:

$$e^{-0.5\sigma_t d(\hat{\omega} \rightarrow \mathbf{o})} \approx e^{-0.5\sigma_t d(\hat{\boldsymbol{\mu}}_a \rightarrow \mathbf{o})}. \quad (25)$$

Putting things together, we approximate the aperture function as:

$$\tilde{a}_i(\hat{\omega}, \mathbf{o}) \approx \eta \cdot e^{\boldsymbol{\mu} \cdot \hat{\omega}}, \quad (26)$$

with

$$\eta = e^{-\gamma_a - \frac{1}{2} \sigma_t d(\hat{\boldsymbol{\mu}}_a \rightarrow \mathbf{o}_1)}, \quad \boldsymbol{\mu} = \gamma_a \hat{\boldsymbol{\mu}}_a + ik(\mathbf{o}_1 - \mathbf{i}). \quad (27)$$

We approximate the viewing aperture function similarly.

We now turn our attention to the scattering amplitude function s in Eqs. (14) and (18). We use the expectation-maximization algorithm




Fig. 5. **Phase function fitting.** Visualizing the von Mises-Fisher mixture fit of some Henyey Greenstein phase functions as well as real world phase functions measured by [Gkioulekas et al. 2013]. A small number of mixture components provides a good fit even when the phase function includes a back lobe (e.g. liquid clay). The top row shows a full $[-\pi, \pi]$ range, and the lower row zooms around the central forward/backward lobe of the function for better visualization.

of Banerjee et al. [2005] to approximate s as a *mixture* of von Mises-Fisher functions centered at $\hat{\mathbf{i}}$:

$$s(\hat{\mathbf{i}}, \hat{\mathbf{v}}) \approx \sum_m \pi_m e^{\gamma_{s,m}(\hat{\mathbf{i}} \cdot \hat{\mathbf{v}})}. \quad (28)$$

In Fig. 5, we show fits of this kind for Henyey-Greenstein phase functions, as well as real-world phase functions measured by Gkioulekas et al. [2013]; in all cases, the phase function is accurately approximated using a small number of mixture components.

Integration and convolution of von Mises-Fisher functions. The approximations we derived facilitate computing the covariance integrals Eqs. (14) and (18). In particular, these evaluation can now be done *analytically*, without the need for Monte Carlo integration, using the properties of von Mises-Fisher functions.

We consider first the spherical integration in Eq. (18). By approximating the integrand as a von Mises-Fisher function, as described above, we can compute this integral analytically using:

$$\int_{\hat{\omega} \in \mathbb{S}^2} \eta e^{\mu \cdot \hat{\omega}} = \eta \cdot 4\pi \frac{\sinh(\sqrt{\mu})}{\sqrt{\mu}} = \eta \cdot 2\pi \frac{e^{\sqrt{\mu}} - e^{-\sqrt{\mu}}}{\sqrt{\mu}}. \quad (29)$$

In this equation, $\sqrt{\mu}$ equals:

$$\sqrt{\mu} \equiv \sqrt{\mu_x^2 + \mu_y^2 + \mu_z^2}. \quad (30)$$

where each term in the summation involves complex square power rather than squared amplitude, and thus $\sqrt{\mu}$ is a *complex* number.

We now consider the spherical convolution in Eq. (14). We use the same approach as for integration, and replace the two functions that are being convolved with their approximation in terms of von Mises-Fisher functions we derived above. The resulting convolution of two von Mises-Fisher functions is also available through a simple analytical form. We note that, even though the exact convolution result is not itself a von Mises-Fisher function, it can be accurately approximated as one, as is necessary to facilitate subsequent computation steps in our rendering algorithm. In particular, in App. A.3, we show that the convolution can be approximated as:

$$\int_{\hat{\psi} \in \mathbb{S}^2} e^{(\mu \cdot \hat{\psi})} \cdot e^{\gamma_s(\hat{\omega} \cdot \hat{\psi})} \approx \frac{2\pi}{\beta_o} e^{\frac{\gamma_s}{\beta_o}(\mu \cdot \hat{\omega}) + c_o}, \quad (31)$$

where

$$\beta_o = \sqrt{\mu + \gamma_s \hat{\omega}_o}, \quad c_o = \beta_o - \frac{\gamma_s}{\beta_o}(\mu \cdot \hat{\omega}_o). \quad (32)$$

We discuss the selection of $\hat{\omega}_o$ in App. A.3, but the simplest strategy is to select it as the direction at the center of the viewing aperture. We compare approximated and exact convolutions in App. A.4. Note that, since the scattering amplitude function is approximated by a mixture of von Mises-Fisher functions, we need to compute the convolution with each mixture element separately.

5.2 Visualizing convolution

The function $Y(\hat{\omega}, \mathbf{o}_1, \mathbf{i})$, defined in Eq. (15) as the convolution of the aperture function with the scattering function, equals the throughput of a path up until the first scattering event. This function encodes the contribution of a path starting at node \mathbf{o}_1 and emerging from it at direction $\hat{\omega}$. This function is a fundamental building block of the near-field correlation, and as such we study and visualize its structure. In the next section, we use this to devise importance sampling schemes for accelerating our Monte Carlo rendering algorithms.

Spatial structure. The aperture function $\tilde{a}_i(\hat{\mathbf{i}}; \mathbf{o}_1)$ is complex, and its phase depends on the distance between the focus point \mathbf{i} and location \mathbf{o}_1 of the first scatterer in the path (Fig. 4(a)). When the phase variation is rapid, blurring this complex function will reduce the magnitude to zero. We want to understand for which \mathbf{o}_1 positions the throughput contribution $Y(\hat{\omega}, \mathbf{o}_1, \mathbf{i})$ is not zero. This will be valuable for defining an importance sampling strategy that avoids sampling \mathbf{o}_1 in areas receiving no energy. To this end we define

$$e(\mathbf{o}|\mathbf{i}) \equiv \int_{\hat{\omega} \in \mathbb{S}^2} |Y(\hat{\omega}, \mathbf{o}, \mathbf{i})|. \quad (33)$$

Fig. 6 visualizes the shape of $e(\mathbf{o}|\mathbf{i})$ for an $x-z$ volume slice. This is similar to the optics concept of a *Gaussian beam* [Yariv 1997], focused at \mathbf{i} . Considering that our illumination is a beam focused at \mathbf{i} , it is expected that scattering points \mathbf{o}_1 that are not located in the area of the illumination beam will not receive light. However, while the shape of a focused beam depends only on the aperture width γ_a , $e(\mathbf{o}|\mathbf{i})$ also depends on the width of the scattering function γ_s . In particular, assuming for ease of notation that the aperture axis is

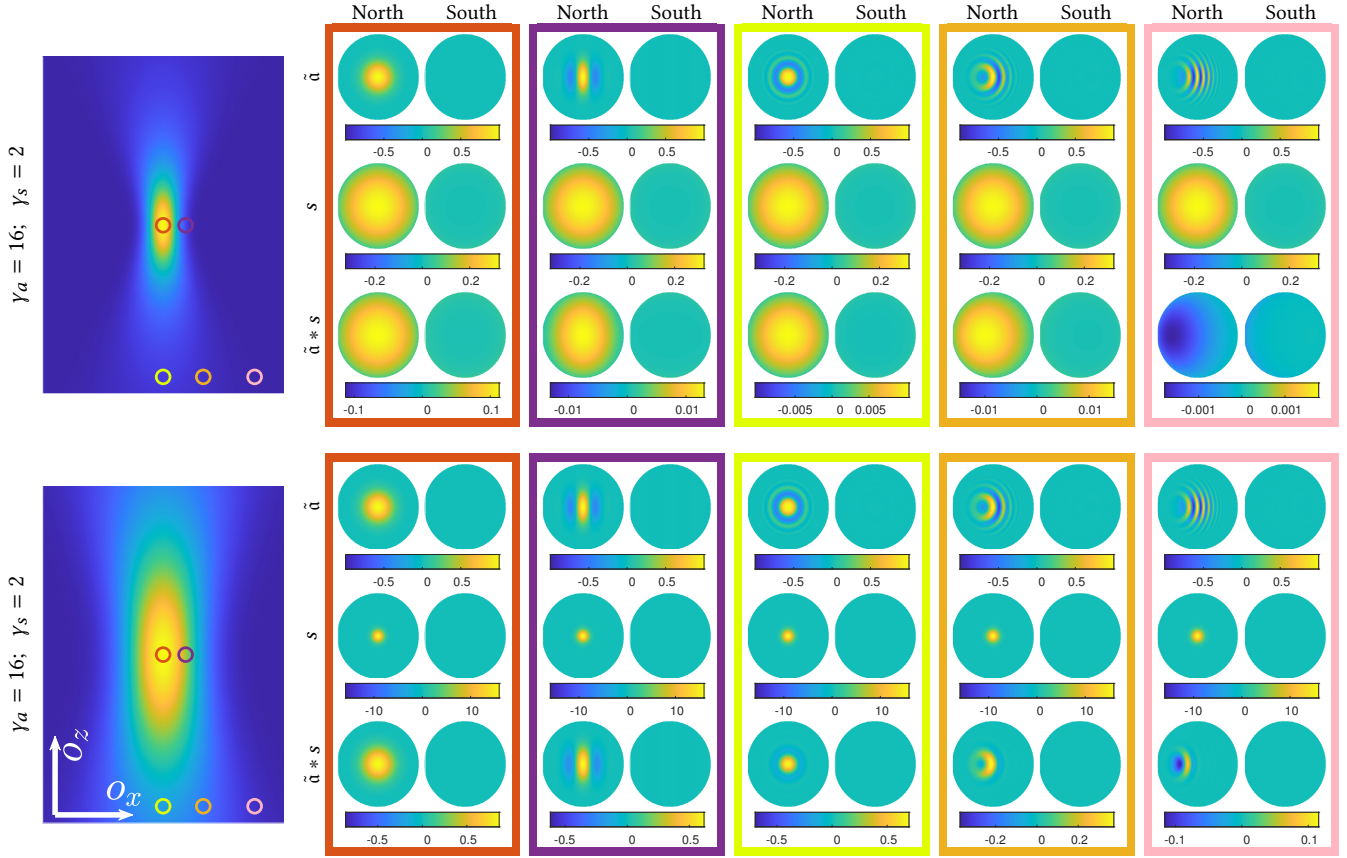


Fig. 6. **Visualizing spatial and angular throughput components.** Left: $e(\mathbf{o}|\mathbf{i})$, the energy of path starting points, as an $x - z$ slice through the volume. This resembles the shape of a *Gaussian beam*, which is narrow at the focus plane and wide at out-of-focus depths. The panels visualize two scattering functions defined by $\gamma_s = 2$ (wide) and $\gamma_s = 100$ (narrow), showing that the beam waist is wider for narrow forward-scattering functions. Right: the angular part of the throughput, $\Upsilon(\hat{\omega}|\mathbf{o}, \mathbf{i}) = \tilde{\mathbf{a}} * \mathbf{s}$ as a function of direction $\hat{\omega}$, for a subset of \mathbf{o} positions marked with corresponding colors on the Gaussian beams. These are functions on the 3D sphere, and we show the north (forward) and south (backward) hemispheres projected on the $\hat{\omega}_x - \hat{\omega}_y$ plane. The directions with high throughput shift for points \mathbf{o} at the periphery of the beam (e.g., orange and pink points). We display only the real part of these complex spherical functions.

aligned with the north pole $\hat{\mu}_a = \hat{\mathbf{z}}$, we show in App. A.5 that:

$$e(\mathbf{o}|\mathbf{i}) = G(\mathbf{o}_{xy}|\mathbf{o}_z, \mathbf{i}) = \beta(z)e^{-\frac{\|\mathbf{o}_{xy} - \mathbf{i}_{xy}\|^2}{2w(z)^2}}, \quad (34)$$

with

$$w(z; \gamma_s, \gamma_a) = \sqrt{\frac{\gamma_a + \gamma_s}{k^2} + \frac{z^2}{\gamma_a}}, \quad (35)$$

$$\beta(z) = e^{-\frac{1}{2}\sigma_t d(\hat{\mathbf{z}} \rightarrow \mathbf{o}_z)} \pi \frac{k^2 \gamma_a^2 (\gamma_s + \gamma_a)}{2w(z; \gamma_s, \gamma_a)^2} e^{\gamma_a + \gamma_s}. \quad (36)$$

and $z \equiv \mathbf{o}_z - \mathbf{i}_z$. For every z plane, $e(\mathbf{o}|\mathbf{i})$ is a planar Gaussian with standard deviation equal to $w(\mathbf{o}_z - \mathbf{i}_z)$. The Gaussian is narrowest when $\mathbf{o}_z - \mathbf{i}_z = 0$, that is, when \mathbf{o}_1 is at the same depth as \mathbf{i} . The beam expands at depths away from the focus depth.

Angular structure. Fig. 6 also provides a visualization of the angular part of the throughput, namely the variation of $\Upsilon(\hat{\omega}|\mathbf{o}, \mathbf{i})$ as a function of $\hat{\omega}$, in a few positions of the first point \mathbf{o} . For that, we display the aperture function $\tilde{\mathbf{a}}_i(\hat{\mathbf{i}}|\mathbf{o})$ as a spherical function (i.e., a function of direction), before and after convolution with \mathbf{s} . We compare the exact convolution against the approximation based on von

Mises-Fisher functions, showing good agreement. The directions with highest power after convolution can shift, and may not be located at the center of the sphere (e.g., the orange and pink points in lower panel). This happens at the periphery of the beam. Below we use the directional density $\Upsilon(\hat{\omega}|\mathbf{i}, \mathbf{o})$ for importance sampling.

6 IMPORTANCE SAMPLING

In Sec. 4.1 the covariance is expressed as an integral over path space, where each path contributes a term $c(\vec{\mathbf{x}}^s)$. For convenience we repeat here the definition for paths of length $B \geq 2$:

$$c(\vec{\mathbf{x}}^s) = f(\vec{\mathbf{x}}^s) \cdot \Upsilon(\hat{\omega}_1, \mathbf{o}_1, \mathbf{i}^1) \Upsilon(\hat{\omega}_{B-1}, \mathbf{o}_B, \mathbf{v}^1) \cdot \Upsilon(\hat{\omega}_1, \mathbf{o}_1, \mathbf{i}^2)^* \Upsilon(\hat{\omega}_{B-1}, \mathbf{o}_B, \mathbf{v}^2)^*. \quad (37)$$

We use Monte Carlo approximation of this integral, by sampling N sub-paths $\vec{\mathbf{x}}^{s,n}$ from a distribution $p(\vec{\mathbf{x}}^{s,n})$ and computing

$$C = \frac{1}{N} \sum_n \frac{c(\vec{\mathbf{x}}^{s,n})}{p(\vec{\mathbf{x}}^{s,n})}. \quad (38)$$

The quality of this estimator depends on the sampling distribution p , and estimation variance reduces when p closely approximates c .

ALGORITHM 1: Monte Carlo rendering of covariance $C_{v_1, v_2}^{i_1, i_2}$.

▷Initialize covariance estimate.

Set $C = 0$.

for iteration = 1 : N **do**

▷Sample a subpath:

▷Sample first vertex from the sum of Gaussian beams.

Sample beam index $j \in \{1, 2\}$

Sample point $\mathbf{o}_1 \sim e(\mathbf{o}_1 | i^j)$.

▷Update covariance with single scattering path.

Update $C += \frac{1}{p(\mathbf{o}_1)} \int_{\hat{\mathbf{v}}^1 \in \mathbb{S}^2} \tilde{a}_{v^1}(\hat{\mathbf{v}}^1, \mathbf{o}_1) \Upsilon(\hat{\mathbf{v}}^1, \mathbf{o}_1, i^1) \cdot \left(\int_{\hat{\mathbf{v}}^2 \in \mathbb{S}^2} \tilde{a}_{v^2}(\hat{\mathbf{v}}^2, \mathbf{o}_1) \Upsilon(\hat{\mathbf{v}}^2, \mathbf{o}_1, i^2) \right)^*$.

▷Sample first direction from the angular throughput.

Sample direction $\hat{\omega}_1 \sim |\Upsilon(\hat{\omega}_1 | \mathbf{o}_1, i^1)|^2$.

▷Continue tracing the subpath:

▷Sample second vertex of subpath.

Sample distance $d \sim \sigma_t e^{-\sigma_t d}$.

Set point $\mathbf{o}_2 = \mathbf{o}_1 + d \cdot \hat{\omega}_1$.

Set $b = 2$.

while \mathbf{o}_b inside medium **do**

▷Update covariance with next-event estimation.

Update $C += \frac{1}{p(\mathbf{o}_1)p(\hat{\omega}_1 | \mathbf{o}_1)} \Upsilon(\hat{\omega}_1, \mathbf{o}_1, i^1) \Upsilon(\hat{\omega}_{b-1}, \mathbf{o}_b, v^1) \Upsilon(\hat{\omega}_1, \mathbf{o}_1, i^2)^* \Upsilon(\hat{\omega}_{b-1}, \mathbf{o}_b, v^2)^*$

▷Sample next vertex of subpath:

▷Sample direction from phase function.

Sample direction $\hat{\omega}_b \sim |s(\hat{\omega}_{b-1} \cdot \hat{\omega}_b)|^2$.

▷Sample free path.

Sample distance $d \sim \sigma_t e^{-\sigma_t d}$.

▷Create next vertex of subpath.

Set point $\mathbf{o}_{b+1} = \mathbf{o}_b + d \cdot \hat{\omega}_b$.

▷Account for absorption.

Sample scalar $a \sim \text{Unif}[0, 1]$.

if $a > \sigma_s / \sigma_t$ **then**

▷Terminate subpath at absorption event.

break

end

Set $b = b + 1$.

end

▷Produce final covariance estimate.

Update $C = \frac{1}{N} C$.

return C .

Bar et al. [2019] sample sub-paths from a distribution $p(\vec{\mathbf{x}}^s) \propto f(\vec{\mathbf{x}}^s)$, where the first node \mathbf{o}_1 and direction $\hat{\omega}_1$ are sampled uniformly.

This uniform sampling strategy can be problematic, because the term c (Eq. (37)) includes the throughput Υ of the start and end segments. In particular, as discussed in Sec. 5.2, the spatial part of Υ has a shape similar to a Gaussian beam (Fig. 6). As the beam assigns zero weight to most points in space, uniform sampling of the starting node \mathbf{o}_1 will produce many paths with zero contribution. When the phase function is narrow, a similar argument holds for the first direction $\hat{\omega}_1$. We address this issue by deriving importance sampling strategies for the start node and direction.

To this end, we sample the path from a distribution

$$p(\mathbf{o}_1 \rightarrow \dots \rightarrow \mathbf{o}_B) = p(\mathbf{o}_2 \rightarrow \dots \rightarrow \mathbf{o}_B | \hat{\omega}_1, \mathbf{o}_1) p(\hat{\omega}_1 | \mathbf{o}_1) p(\mathbf{o}_1). \quad (39)$$

The first scattering point is sampled from

$$p(\mathbf{o}_1) \propto \frac{1}{2} \left(|e(\mathbf{o}_1 | i^1)|^2 + |e(\mathbf{o}_1 | i^2)|^2 \right). \quad (40)$$

We implement this sampling by first uniformly sampling one of the two beams in the summand, and then sampling \mathbf{o}_1 from the Gaussian beam $G(\mathbf{o}_{xy} | \mathbf{o}_z, i^j)$ of Eq. (34), where $j \in 1, 2$ denotes the index of the sampled beam. The first direction is sampled from

$$p(\hat{\omega}_1 | \mathbf{o}_1) \propto |\Upsilon(\hat{\omega}_1 | \mathbf{o}_1, i^j)|^2. \quad (41)$$

The rest of the path is sampled from

$$p(\mathbf{o}_2 \rightarrow \dots \rightarrow \mathbf{o}_B | \mathbf{o}_1, \hat{\omega}_1) \propto f(\mathbf{o}_1 \rightarrow \dots \rightarrow \mathbf{o}_B), \quad (42)$$

as in standard volumetric path tracing. The sampling scheme is summarized in Alg. 1, and we provides details in App. A.6. We note that we choose to importance sample the first segment of the path rather than the last one because, for most imaging configurations in this paper, we had a small number of illuminators and a large number of camera pixels. We can importance sample the last segment when illuminators are more than camera pixels, or importance sample both the first and last segment in bidirectional algorithms.

7 FIELD RENDERING

So far we focused on evaluating speckle covariance. A complementary question is how to directly render speckle images. Suppose, for example, that we want to generate N_i images from N_i different illumination points, where each image includes N_v sensor points. This requires sampling $N_{i,v} = N_i \cdot N_v$ complex numbers. One approach for this is to evaluate the corresponding covariance matrix of size $N_{i,v} \times N_{i,v}$, and then use it to sample values. However, for large $N_{i,v}$, the covariance matrix can be impractically large. Instead, we seek an algorithm with complexity $O(N_{i,v})$ rather than $O(N_{i,v}^2)$.

For this, we follow Bar et al. [2019], who note that the covariance matrix as given by Eq. (3), (14) and (18) is essentially an infinite summation of rank-1 matrices, decomposed over the path space

$$C_{v^k, v^m}^{i^k, i^m} = \int_{\mathbb{P}} f(\vec{\mathbf{x}}^s) b(i^k, v^k) b(i^m, v^m)^* d\vec{\mathbf{x}}^s, \quad (43)$$

for all (k, m) pairs $k, m \in \{1, \dots, N_{i,v}\}$, with

$$b(i^k, v^k) = \begin{cases} \int_{\hat{\mathbf{v}} \in \mathbb{S}^2} \tilde{a}_{v^k}(\hat{\mathbf{v}}, \mathbf{o}_1) \Upsilon(\hat{\mathbf{v}}, \mathbf{o}_1, i^k), & B = 1, \\ \Upsilon(\hat{\omega}_1, \mathbf{o}_1, i^k) \Upsilon(\hat{\omega}_B, \mathbf{o}_B, v^k), & B \geq 2. \end{cases} \quad (44)$$

Despite the fact that this rank-one decomposition is over-complete, we can use it to sample from the covariance matrix. We sample sub-paths $\vec{\mathbf{x}}^s \sim p(\vec{\mathbf{x}}^s)$, using the same importance function as in Sec. 6. The field is updated using $b(i^k, v^k)$, resulting in

$$u(i^k, v^k) = \frac{1}{\sqrt{N}} \sum_n b(i^k, v^k | \vec{\mathbf{x}}^{s,n}) \sqrt{\frac{f(\vec{\mathbf{x}}^{s,n})}{p(\vec{\mathbf{x}}^{s,n})}} e^{2\pi i \zeta_n}, \quad (45)$$

where $\zeta_n \in [0, 1]$ is a random phase ensuring

$$E[b(i^k, v^k | \vec{\mathbf{x}}^{s,n}) e^{2\pi i \zeta_n}] = 0. \quad (46)$$

As different paths are sampled independently, and given the zero-mean property, contributions from different sub-paths are uncorrelated. Namely, for $n_1 \neq n_2$:

$$E \left[b(i^k, v^k | \vec{\mathbf{x}}^{s,n_1}) e^{2\pi i \zeta_{n_1}} \cdot b(i^m, v^m | \vec{\mathbf{x}}^{s,n_2})^* e^{-2\pi i \zeta_{n_2}} \right] = 0, \quad (47)$$

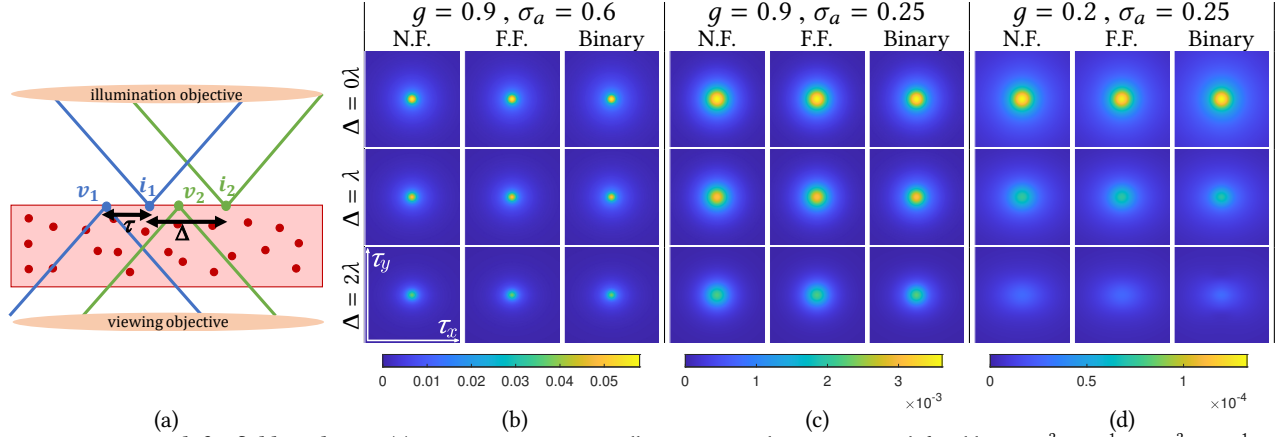


Fig. 7. **Comparisons with far field rendering.** (a) Imaging setup. We use illumination and viewing pairs defined by $\Delta = \mathbf{i}_{x,y}^2 - \mathbf{i}_{x,y}^1 = \mathbf{v}_{x,y}^2 - \mathbf{v}_{x,y}^1$ and $\tau = \mathbf{v}_{x,y}^1 - \mathbf{i}_{x,y}^1 = \mathbf{v}_{x,y}^2 - \mathbf{i}_{x,y}^2$. (This setup is also used in Fig. 8-11.) (b-d) We compare covariances rendered using our near-field (N.F.) algorithm against those obtained by rendering far-field (F.F.) correlations and applying focusing. We also compare with covariances rendered with a binary aperture of equivalent width to our apodized aperture. As binary apertures cannot be well-approximated using von Mises-Fisher functions, we only render them using the far-field algorithm. The volume has size $50\lambda \times 50\lambda \times 20\lambda$. We evaluate two aperture widths in (b,c) and two phase functions in (c,d). Each square visualizes correlation as a function of τ , and different rows correspond to different x -axis illuminator shifts $\Delta = \mathbf{i}_x^2 - \mathbf{i}_x^1$. We see close agreement in all cases.

for all (k, m) pairs. Therefore, we see that $E \left[u(\mathbf{i}^k, \mathbf{v}^k) \cdot u(\mathbf{i}^m, \mathbf{v}^m)^* \right]$ equals the desired covariance of Eq. (43).

Fig. 1 shows speckle images rendered with this algorithm. We note that in Eq. (45), the *same* set of paths is used to update *all* illuminators and sensors. As a result, the speckle images due to different illuminators in Fig. 1 are correlated, shifted versions of each other, corresponding to the memory effect property. Rendering each of these images independently, using a standard Monte Carlo volume rendering approach, would fail to reproduce these correlations.

8 EVALUATION

We now evaluate the efficiency and accuracy of our proposed rendering algorithms. We compare our algorithms against three alternatives: First, we compare with an approach that uses the Monte Carlo rendering algorithm of Bar et al. [2019] to produce far-field estimates, and then converts them to near-field estimates using Eqs. (2) and (9). We also consider a few variants of this approach that benefit from various acceleration techniques. Second, we compare with groundtruth estimates produced by a wave-equation solver. Third, we compare with estimates produced using a layered propagation approach popular in the optics literature. Additionally, we show experiments evaluating performance improvements due to the importance sampling scheme of Sec. 6. In Sec. 9, we show experiments relating to the tilt-shift memory effect [Osnabrugge et al. 2017], including experiments validating our algorithms against groundtruth measurements of real materials.

Experimental configurations. Our experiments focus on illumination and imaging configurations that are known from literature to produce strongly-correlated fields. We describe these configurations using Fig. 7(a) as reference: As we focus on transmission mode imaging, the focused illumination is placed at the back of the sample, and the focused sensor is placed at the front. We refer to App. A.4 for an evaluation of our technique under reflective imaging conditions. We consider the case where the illumination focus points $\mathbf{i}^1, \mathbf{i}^2$ and

sensing focus points $\mathbf{v}^1, \mathbf{v}^2$ are all located at the same depth plane, at the back plane of the volume. This configuration corresponds, for example, to the fluorescent imaging setting, where a fluorescent particle is deep inside a medium such as tissue, and a camera attempts to observe it by focusing through the scattering at the illuminator's plane.

We denote by $\mathbf{i}_{x,y}^1, \mathbf{i}_{x,y}^2, \mathbf{v}_{x,y}^1, \mathbf{v}_{x,y}^2$ the $x - y$ coordinates of the corresponding 3D points on the focus plane. We denote by $\Delta \equiv \mathbf{i}_{x,y}^2 - \mathbf{i}_{x,y}^1$ the 2D displacement between the illumination directions, and by $\tau \equiv \mathbf{v}_{x,y}^1 - \mathbf{i}_{x,y}^1$ the 2D displacement between the illumination and viewing directions (Fig. 7(a)). Classical memory effect theory [Feng et al. 1988] states that strong correlations exist between fields for illumination and viewing pairs satisfying $\mathbf{i}_{x,y}^2 - \mathbf{i}_{x,y}^1 = \mathbf{v}_{x,y}^2 - \mathbf{v}_{x,y}^1 = \Delta$, for small Δ values. With this in mind, in our experiments, we evaluate and visualize correlations of the form:

$$C(\Delta, \tau) = C_{\mathbf{i}_{x,y}^1, \mathbf{i}_{x,y}^1 + \Delta}^{\mathbf{i}_{x,y}^1, \mathbf{i}_{x,y}^1 + \Delta}, \quad (48)$$

for different displacements Δ and τ . For sufficiently wide volumes, $C(\Delta, \tau)$ is approximately invariant to \mathbf{i}^1 .

Comparison with far-field approach. We compare first with the simulation approach based on the far-field rendering algorithms of Bar et al. [2019]. To reduce the computational burden, we compare with an approach based on a combination of Eqs. (2) and (9), rather than Eq. (11): We first discretize the illumination and viewing apertures into a set of directions $\hat{\mathbf{i}}$ and $\hat{\mathbf{v}}$. For each set of $(\hat{\mathbf{i}}, \hat{\mathbf{v}})$ values, we use the implementation provided by Bar et al. [2019] to sample far-field scattered waves $u_{\hat{\mathbf{v}}}^{\hat{\mathbf{i}}}$, which we subsequently convert to the near-field scattered waves using Eq. (9). By running the field rendering algorithm of Bar et al. [2019] multiple times, we end up with multiple samples of near-field scattered waves, corresponding to different scatterer configurations. Finally, we use these near-field scattered waves to approximate the covariance as in Eq. (2). We

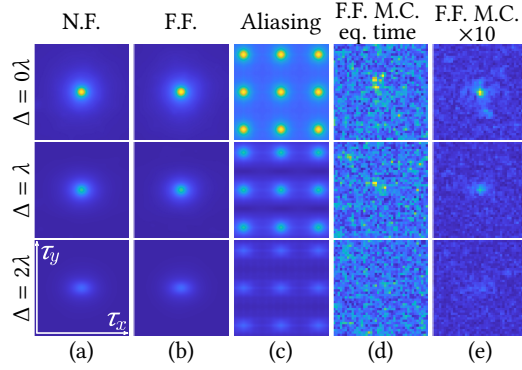


Fig. 8. **Acceleration of far-field rendering.** (a) N.F. (1× runtime). (b) F.F. with dense discretization of aperture integral (400× runtime). (c) F.F. with sparser discretization of aperture integral (100× runtime) (d, e) F.F. with Monte Carlo sampling of aperture integral (1× and 10× runtime).

Table 1. Runtimes of different algorithms for achieving $RMSE < 0.01$.

	$10\lambda \times 10\lambda \times 4\lambda$	$25\lambda \times 25\lambda \times 10\lambda$	$50\lambda \times 50\lambda \times 20\lambda$
N.F.	0.46 s	2.12 s	6 s
F.F. Tab.	31 s	221 s	2.4×10^3 s
F.F. M.C.	199 s	1.47×10^3 s	2.04×10^4 s

note that, in Sec. 4, we presented the formal approach for estimating near-field covariance using far-field rendering, based on the far-field covariance Eq. (11) rather than fields. In these experiments, we opt for the field-based approach, because the number of far-field samples it requires scales more favorably with the width W of the simulated volume (W^4 scaling for the field-based approach, W^8 scaling for the covariance-based approach).

In our experiments, to keep the number of far-field samples manageable, we simulate a relatively small volume of size $50\lambda \times 50\lambda \times 20\lambda$ with mean free path $MFP = 10\lambda$, leading to an optical depth $OD = 2$. We consider two choices of aperture width, $\sigma_a = 0.6$ and $\sigma_a = 0.25$, and both wide and narrow forward-scattering Henyey-Greenstein phase functions (low and high g values, respectively). When using the far-field approach, we simulate both von Mises-Fisher-apodized and binary aperture masks $m(\cdot)$ of the same width, to quantify the effect of our apodization approximation on accuracy (binary apertures can only be evaluated at the far-field, at increased computational complexity, as the von Mises-Fisher approximation does not apply). We show the results in Fig. 7. We observe that, in all cases, our near-field rendering algorithm produces very similar results to the far-field approaches. We also note that the far-field approach produces very close results when using apodization and binary masks, indicating that our apodization approximation does not introduce significant bias. At the same time, our near-field approach is orders-of-magnitude faster, as we quantify in detail below.

In Fig. 8, we additionally compare against two accelerated variants of the far-field approach. We reuse the configuration of Fig. 7(d), this time with a wider τ grid. First, we render far-field covariance only for a (fixed) regular subset of directions. As the number of samples is significantly lower than that required by the Nyquist sampling rate in Eq. (12), aliased replicas appear. Second, we use a Monte Carlo strategy that randomly samples the directions at which far-field

Table 2. Equal-time comparisons of different algorithms.

	$10\lambda \times 10\lambda \times 4\lambda$		$25\lambda \times 25\lambda \times 10\lambda$		$50\lambda \times 50\lambda \times 20\lambda$	
	1 s	60 s	5 s	500 s	15 s	6000 s
N.F.	0.0089	0.0076	0.0096	0.0093	0.0135	0.0110
N.F. biased ref.	0.0053	0.0021	0.0050	0.0018	0.0080	0.0045
F.F. Tab.	0.0396	0.0078	0.0672	0.0056	0.0915	0.0051
F.F. M.C.	0.1140	0.0198	0.2024	0.0207	0.3317	0.0632

covariance is computed when evaluating the integral of Eq. (11). The number of random samples is selected to achieve the lowest possible error for equal runtime. As the integrand is complex, using Monte Carlo approximation results in very high variance. At the cost of increased runtime, Fig. 8(e) demonstrates a better estimate.

Runtime and bias evaluation. We compare the runtime of our near-field approach against two versions of the far-field approach: a tabulated version, computing all far field directions before applying near field transformation, as in Fig. 8(b); and a Monte Carlo version, randomly sampling directions, as in Fig. 8(d-e). Table 1 shows the runtime each approach requires to converge to a root-mean-square-error (RMSE) relative to a reference rendering below 0.01. In all cases, our near-field approach has the best performance, and the performance improvement increases as the width W of the simulated volume increases. For the volume size we use for the results in Fig. 7, our near-field approach is 400× faster than the tabulated far-field approach, and 3000× faster than the Monte Carlo approach. The performance advantage will become even larger for volumes with sizes corresponding to realistic lab experiments. All algorithms were implemented and run on a V100 NVIDIA GPU.

Table 2 reports RMSE from equal-time comparisons of the three approaches using volumes of different sizes. Error was measured against a reference rendering produced by running the far-field approach till convergence. For each volume size, we report RMSE for both a small and a long rendering time. The latter allows us to quantify the bias of the near-field approach. While any such bias is too small to be noticeable in the visual comparisons of Fig. 7, the numerical difference between the near-field and far-field covariances does not fully converge to zero. We anticipate that most of the bias is due to the assumption that volumetric attenuation is constant for all directions through the aperture (see Eq. (25)). In the second row of Table 2 we compare our near-field approach against a reference evaluated using the far-field approach but with an equivalent constant directional attenuation, showing smaller error.

The benefit of importance sampling. To evaluate the effect of the importance sampling scheme we introduced in Sec. 6, we compare in Fig. 9 covariance estimates produced by our Monte Carlo algorithm, using three different sampling strategies: first, uniform sampling of both the first path vertex and direction; second, importance sampling of the first vertex and uniform sampling of the first direction; and third, importance sampling of both the first vertex and direction. We observe first that, when run till convergence (e.g., in the top panel where the sample is small and the phase function has significant side-scattering), all three combinations converge to the same result. This confirms that importance sampling does not introduce any additional bias. We additionally observe that the estimate using importance sampling of both the first vertex and




Fig. 9. **The benefit of importance sampling.** The figure compares uniform sampling vs. importance sampling (I.S.) of the first scatterer on the path \mathbf{o}_1 , vs. importance sampling of both the first scatterer and the first direction $\mathbf{o}_1, \hat{\omega}_1$. Top panel: a small target whose size is only $20\lambda \times 20\lambda \times 20\lambda$. For such small targets uniform sampling converges as well, although requiring a larger number of path samples (compare results with 10^3 samples to 10^5 samples), demonstrating that our importance sampling strategy is consistent. Middle panel: When expanding the size of the target volume, convergence of uniform sampling is very slow. As the phase function is rather wide, importance sampling of the first direction does not improve much. Lower panel: If we also use a very forward scattering phase function, convergence is much accelerated by importance sampling the first direction.

direction results in reduced noise in all cases. The improvement becomes more pronounced as the width of the volume increases (second panel in Fig. 9); this is because, as the size of the Gaussian beam relative to the volume decreases, uniform vertex sampling will result in more paths starting in points of the volume that do not receive any light. Finally, the improvement achieved by using importance sampling increases even further as the phase function becomes more forward-scattering (third panel in Fig. 9); in this case, it is necessary to importance sample the first direction as well, otherwise the majority of path-starting directions will have near-zero contribution. Fig. 10 additionally shows convergence plots for the first and third volumes in Fig. 9.

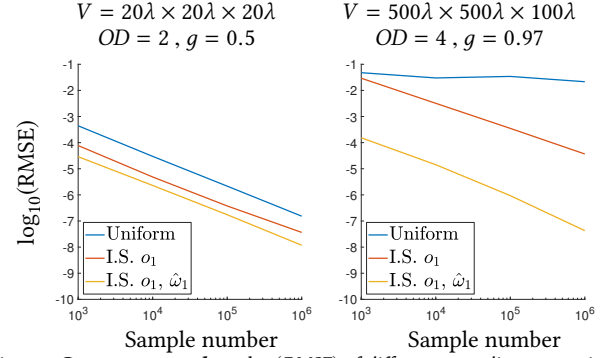


Fig. 10. **Convergence plots.** $\log(\text{RMSE})$ of different sampling strategies as a function of sample number, for the first and third volumes in Fig. 9.

Comparison with a wave-equation solver. Bar et al. [2019] demonstrated the accuracy of their far-field covariance rendering algorithm by comparing against numerical wave-equation solvers. For additional validation, we also compare directly with the μ -diff solver [Thierry et al. 2015] they use. The solver takes as input a configuration of scatterer locations, and uses numerical techniques to solve the wave equation and output an estimate of the complex scattered field. To compute covariance, we run the solver for multiple scatterer instantiations sampled from the same distribution, then use the results to estimate the speckle covariance of Eq. (2). We note that this simulator only works in 2D; thus, for these comparisons only, we restrict our Monte Carlo rendering algorithm to 2D as well.

The left part of Fig. 11 shows correlations values $C(\Delta, \tau)$ (Eq. (48)), as a function of τ for a few Δ and g values, simulated using μ -diff and our technique. Note that, as these simulations are in 2D, τ is a scalar, and thus $C(\Delta, \tau)$ is a 1D curve. We observe that the plots computed with our technique closely match those produced by the wave-equation solver, demonstrating the accuracy of our technique. At the same time, for this small example our technique is three orders of magnitude faster, and can scale to much larger volumes.

Comparison with multi-slice layered propagation. We compare additionally with the multi-slice beam propagation method [Schott et al. 2015], which is a popular numerical approach in optics for simulating wave propagation. This approach accounts for multiple-scattering effects by approximating the simulated volume as a sequence of planar slices orthogonal to the optical axis; layers are modeled as infinitesimally-thin 2D phase masks, separated by free space. The phase mask of each layer is selected to scatter light with an angular spread matching the phase function of the simulated volume. This technique has high computational efficiency compared to exact wave-equation solvers, but cannot model back-scattering, and has worse accuracy for wide scattering angles.

The right part of Fig. 11 shows correlation estimates from this technique, using the same experimental settings as in our comparisons with the wave-equation solver. We observe that, for very forward-scattering phase functions, both our Monte Carlo algorithm and the layered propagation technique closely match the groundtruth produced by the solver. However, for phase functions with significant side-scattering, the accuracy of the layered propagation technique is significantly worse than that of ours. Additionally, our Monte Carlo algorithm is significantly faster than the layered

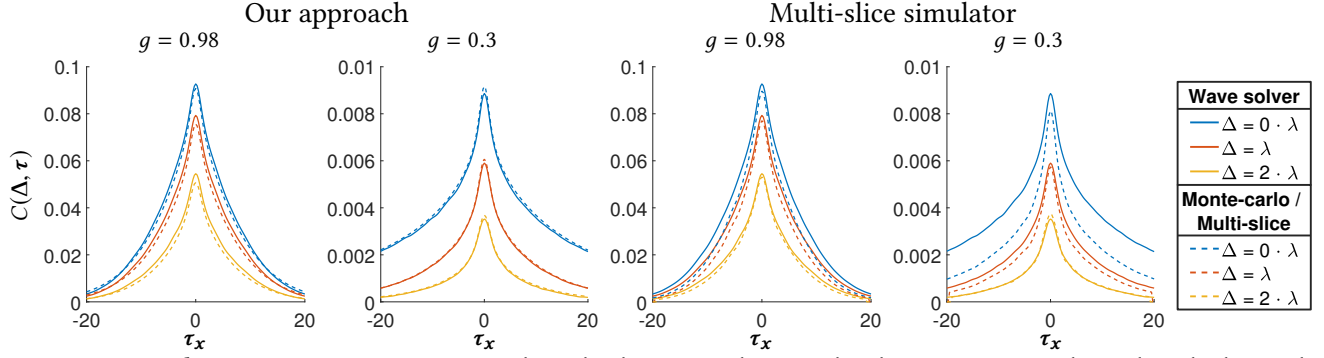


Fig. 11. **Wave optics alternatives.** We compare covariances obtained with our approach, against those by a wave-equation solver, and a multi-slice simulator from optics. Our simulator is physically accurate, producing results in close agreement with those of the solver, while being orders of magnitude faster. The multi-slice approach is valid for a very forward-scattering phase function ($g = 0.98$) and small optical depth. For a wider phase function ($g = 0.3$), the multi-slice approach produces inaccurate results, due to incorrect modeling of back and side scattering angles. The comparison is performed in 2D due to the limitations of the solver. The simulated volume is $200\lambda \times 70\lambda$ wide with $OD = 3$.

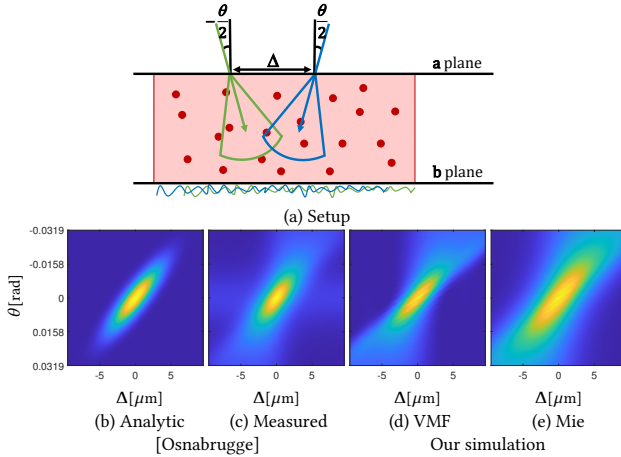


Fig. 12. **The tilt-shift memory effect.** (a) Setup: two beams with spatial displacement Δ_a and angular displacement θ enter the volume. We compute the correlation between the speckle patterns measured at the b plane as a function of both tilt and shift, assuming that the input displacement Δ_a is selected to maximize Eq. (50). The bottom images visualize the results. (b,c) Analytic and measured correlation, reproduced from Osnabrugge et al. [2017]. (d-e) Correlation computed by our algorithm for two different phase functions with the same average cosine $g = 0.98$.

propagation technique, as the latter requires for each layer a high-resolution discretization of the wave, resulting in large dense arrays that need to be convolved to model propagation between layers.

9 THE TILT-SHIFT MEMORY EFFECT

In this section, we use our near-field rendering technique to study the *tilt-shift memory effect* property introduced by Osnabrugge et al. [2017]. We first briefly review this property, using Fig. 12(a) as a reference: A scattering volume of thickness L is illuminated by two input beams from its top surface, denoted in Fig. 12(a) as the a plane. The scattered fields u^1, u^2 due to the two beams are imaged by a camera focused at the lower plane of the volume, denoted in Fig. 12(a) as the b plane. The two illumination beams have a displacement relative to each other equal to $\mathbf{i}^2 - \mathbf{i}^1 = \Delta_a$. The key observation of Osnabrugge et al. [2017] is that we can increase the correlation

of the speckle fields due to the two beams if we additionally *tilt* the beams at angles $-\theta/2, \theta/2$ respectively. Intuitively, appropriately selecting the tilting angle as a function of the displacement Δ_a helps increase the overlap between the defocused beams inside the medium, and thus increases the correlation of the resulting speckle fields.

This observation motivates evaluating how speckle field correlation varies as a function of tilt angle and shift displacements at the input and output planes. Concretely, we can write this as a function:

$$C(\Delta_a, \Delta_b, \theta) \equiv E \left[\sum_{\tau} u^1(\tau - \Delta_b/2) \cdot u^2(\tau + \Delta_b/2)^* e^{ik\tau\theta} \right], \quad (49)$$

where the expectation is taken over all fields with the same material parameters (e.g., fields generated by different scatterer instantiations O sampled from the same density). Osnabrugge et al. [2017] derived an analytic approximation for this function that takes the form:

$$C(\Delta_a, \Delta_b, \theta) \approx \delta(\Delta_b - \Delta_a - L\theta) e^{-\frac{L^3 k^2}{2\ell_{tr}} \left(\frac{\theta^2}{12} + \left(\frac{\theta}{2} + \frac{\Delta_b}{L} \right)^2 \right)}, \quad (50)$$

where $\delta(\cdot)$ is the Dirac delta function, and ℓ_{tr} is the *transport mean free path* $\ell_{tr} \equiv MFP/(1 - g)$. Their derivation is based on three simplifying assumptions: it uses a layered representation similar to that of the multi-slice layered propagation technique in Sec. 8; at each layer, it assumes forward-only propagation; and it uses a differential equation to integrate over multiple scattering planes.

Additionally, Osnabrugge et al. [2017] show measurements of the function $C(\Delta_a, \Delta_b, \theta)$ for a tissue phantom of thickness $L = 258 \mu\text{m}$, made of silica microspheres immersed in agarose gel. By combining Mie theory [Frisvad et al. 2007] with the dispersion and sizing properties of the materials used for fabrication, the authors estimate for the phantom an anisotropy parameter $g = 0.98$ and mean free path $MFP = 296 \mu\text{m}$. In Fig. 12(b,c), we replicate from their paper (using data provided by the authors) the measured correlation $C(\Delta_a, \Delta_b, \theta)$, as well as the analytical prediction using the model of Eq. (50). The correlation is displayed as a function of (Δ_b, θ) alone, with Δ_a selected according to the Dirac delta relationship in Eq. (50).

To evaluate the accuracy of our near-field covariance rendering algorithm, we use it to estimate the function $C(\Delta_a, \Delta_b, \theta)$ for the

same tissue phantom. We refer to App. A.7 for details on how to simulate this with our framework. We perform simulations using both the exact Mie-theory phase function describing the phantom, as well as a von Mises-Fisher phase function with the same $g = 0.98$ average cosine. Comparing with the measured data and the analytical approximation in Fig. 12(b,c), we make the following observations: Both our renderings and the analytical approximation produce a correlation function with a dominant lobe that has the same orientation as the one in the measured data. However, our simulations match the dimensions of this lobe more closely than the analytical model. Overall, our simulations reproduce the important qualitative features of the measurements, confirming the accuracy of our algorithms. Differences between our renderings and the measurements are likely due to inaccurate modeling of the true material properties of the phantom—as seen in Fig. 12(d-e), replacing the exact reported phase function with an approximation results in a better match to the measurements—and due to aberrations in the imaging optics. We also note that Osnabrugge et al. [2017] measured correlation with an interferometric setup, which typically produces very noisy estimates of small signals such as weak speckle correlations.

9.1 Focusing through turbid media

Having shown that our rendering algorithms can accurately reproduce the tilt-shift memory effect, we now use them to analyze this effect in the context of a specific biomedical imaging application: we simulate focusing through turbid media with adaptive optics. This involves using, e.g., a spatial light modulator to produce a coherent wavefront whose shape is specific to the tissue sample being imaged. Finding the exact shape of this wavefront is challenging and usually requires having external information or a guiding star [Horstmeyer et al. 2015]. Once we know the wavefront needed to focus at a specific point inside the volume, an important practical consideration is whether we can use this information to *refocus* at other points in a neighborhood around the original point [Judkewitz et al. 2014].

The memory effect property of scattered fields provides a way to achieve this refocusing task. The way this works is that we first use a guiding star to measure the scattered field that is created due to emission from a single point $\mathbf{i}_{x,y}$ at the a plane (Fig. 12(a)). Measuring this field $u(\mathbf{v}_{x,y})$ at all points $\mathbf{v}_{x,y}$ at the b plane provides us with exactly the wavefront shape we need to focus at point $\mathbf{i}_{x,y}$. Then, the memory effect property suggests that a shifted wave

$$u(\mathbf{v}_{x,y} + \Delta), \quad (51)$$

can be used to roughly focus at $\mathbf{i}_{x,y} + \Delta$. Osnabrugge et al. [2017] improve upon this idea by recommending to refocus using a wave that is both shifted and tilted; that is, a wave of the form

$$u(\mathbf{v}_{x,y} + \Delta)e^{ik\theta(\Delta)\mathbf{v}_{x,y}}, \quad (52)$$

where $\theta(\Delta)$ is selected according to Eq. (50). For example, the strength of the yellow spot at the fourth column of Fig. 13 is slightly higher than at the second column. Using Eq. (50), Osnabrugge et al. [2017] also predict the range of shifts for which this refocusing is effective; that is, they estimate the scanning range for which sufficient memory effect correlations exists.

We use our rendering algorithms to evaluate these refocusing techniques, as shown in Fig. 13. The top part of the figure compares




Fig. 13. **Adaptive optics focusing.** Illuminating a scattering slab with the complex wave visualized at the first and third columns of the top row leads to a sharp focused point at the other edge of the random media slab. Due to the ME, a small shift of the same pattern can focus at a nearby point, but focusing power degrades with displacement (note the weak power at the third row). Applying both *shift* and *tilt* to the input pattern (rightmost columns) leads to a stronger power at the same displacement (compare the power for the non-zero displacements at rows 2 and 3). Lower panel: Using our approach to evaluate the expected power one can achieve with the shift only and shift+tilt approaches, as a function of displacement (that is, the average power of the yellow dot at the three top rows). The scan range predicted by our accurate simulator is wider than the analytical prediction. Note that configurations (b) and (c) have the same transport mean free path and should be equivalent according to the simplified analytic model of Eq. (50), yet they are very different according to an accurate MC simulator.

refocusing simulations using only shifting (Eq. (51)) versus using both shifting and tilting (Eq. (52)). We observe that using the tilted-shifted wave improves refocusing, making the yellow focused spots of the second column in Fig. 13 stronger than the fourth one, and thus validating the observation of Osnabrugge et al. [2017]. We refer to App. A.8 for more details regarding this simulation.

We can additionally use our rendering algorithms to more accurately evaluate the scanning range over which this refocusing technique remains effective. In the lower panel of Fig. 13 we plot the expected power we can measure at focus points for different displacements Δ , comparing the prediction by our model with the analytical prediction of Osnabrugge et al. [2017]. We do simulations for a material with a Henyey-Greenstein phase function of $g = 0.98$ and mean free path of $MFP = 75 \mu\text{m}$ at wavelength $\lambda = 0.5 \mu\text{m}$. We test two material thicknesses $L = 75 \mu\text{m}$ and $L = 225 \mu\text{m}$, corresponding to optical depths of $OD = 1$ and $OD = 3$. We observe that the Monte Carlo simulations predict that the effective scanning

range is larger than the range predicted by the analytical model of Eq. (50). We additionally observe that, even though the analytical model depends only on the transport mean free path ℓ_{tr} , our simulations suggest that the scanning range varies significantly for materials with the same transport mean free path, but different phase function and actual mean free path (compare Fig. 13(b,c)).

These results demonstrate that our rendering algorithms can be used to evaluate how the performance of existing imaging techniques depend on exact material parameters, in ways that cannot be predicted using existing analytical models.

10 CONCLUSION

We presented computationally-efficient algorithms for simulating physically-accurate speckle fields and statistics under focused camera and light sources. The key element of our algorithms is the use of closed-form expression for transforming far-field correlations to the near-field ones, for individual Monte Carlo paths. We also derive efficient importance sampling strategies for path generation.

The closed-form expressions are made possible through the use of von Mises-Fisher functions to approximate all spherical functions appearing in the path integral expression of near-field speckle covariance. We make four such approximations: (i) We use apodization in the aperture plane to convert the aperture mask into a spherical von Mises-Fisher function. (ii) We assume the exponential attenuation is constant through all aperture directions. (iii) We approximate the phase function as a mixture of von Mises-Fisher functions. (iv) We approximate the analytical convolution of von Mises-Fisher functions, which is not in general a von Mises-Fisher function, as such a function. We have shown through simulations that: (i) Even though apodization cannot express an exactly binary aperture, if one matches its variance to the desired NA, a similar depth of field and similar speckle statistics are produced. (ii) The assumption of constant directional attenuation introduces negligible bias. (iii) Real-world and common parametric phase functions can be well-approximated using a small number of von Mises-Fisher functions. (iv) The convolution of von Mises-Fisher functions can be closely approximated as a von Mises-Fisher function. For narrow phase functions emitting little energy in the backward direction, numerical problems arise in back directions. As a result, our approach mostly applies for transmission mode imaging, and further research is required for reflection mode imaging.

As an application of our algorithms, we evaluated different performance metrics for memory-effect-based adaptive optics scanning, and showed that in practice performance can be better than what predicted by previous approximate analytical models. We hope that our algorithms will open the door for better understanding of speckle statistics under near-field imaging conditions, in turn allowing researchers to push the boundary of what is possible using speckle-based techniques in deep tissue imaging.

ACKNOWLEDGMENTS

This work was supported by European Research Council Horizon 2020 635537, ISF 1947-20, Ollendorff Minerva Center of the Technion, Gordon center for system Engineering, Porat award, NSF

Expeditions award 1730147, DARPA REVEAL grant HR0011-16-C-0028.

REFERENCES

- David Abookasis and Joseph Rosen. 2004. NOISE 2 imaging system: seeing through scattering tissue with a reference point. *Opt. Lett.* (2004).
- Eric Akkermans and Gilles Montambaux. 2007. *Mesoscopic Physics of Electrons and Photons*. Cambridge University Press.
- Arindam Banerjee, Inderjit S. Dhillon, Joydeep Ghosh, and Suvrit Sra. 2005. Clustering on the unit hypersphere using von Mises-Fisher distributions. *JMLR* (2005).
- Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. 2019. A Monte Carlo Framework for Rendering Speckle Statistics in Scattering Media. *ACM TOG* (2019).
- Chen Bar, Ioannis Gkioulekas, and Anat Levin. 2020. Project Website. <https://github.com/chabner/gaussianBeam-field>.
- Brian A. Barsky and Todd J. Kosloff. 2008. Algorithms for Rendering Depth of Field Effects in Computer Graphics. *ICCOMP* (2008).
- Mahed Batarseh, Sergey Sukhov, Zhean Shen, H. Gemar, Roxana Rezvani, and Aristide Dogariu. 2018. Passive sensing around the corner using spatial coherence. *Nature Communications* (2018).
- Ibrahim Baydoun, Diego Baresch, Romain Pierrat, and Arnaud Derode. 2016. Radiative transfer of acoustic waves in continuous complex media: Beyond the Helmholtz equation. *Physical Review E* (2016).
- Laurent Belcour, Kavita Bala, and Cyril Soler. 2014. A local frequency analysis of light scattering and absorption. *ACM TOG* (2014).
- Stephan Bergmann, Mahsa Mohammadikaji, Stephan Irgenfried, Heinz Worn, Jürgen Beyerer, and Carsten Dachsbaecher. 2016. A Phenomenological Approach to Integrating Gaussian Beam Properties and Speckle into a Physically-Based Renderer. *VMV* (2016).
- Richard Berkovits and Shechao Feng. 1994. Correlations in coherent multiple scattering. *Physics Reports* (1994).
- Bruce J. Berne and Robert Pecora. 2000. *Dynamic light scattering: with applications to chemistry, biology, and physics*. Courier Corporation.
- Jacopo Bertolotti, Elbert G. Van Putten, Christian Blum, Ad Legendijk, Willem L. Vos, and Allard P. Mosk. 2012. Non-invasive imaging through opaque scattering layers. *Nature* (2012).
- Benedikt Bitterli, Srinath Ravichandran, Thomas Müller, Magnus Wrenninge, Jan Novák, Steve Marschner, and Wojciech Jarosz. 2018. A radiative transfer framework for non-exponential media. *ACM TOG* (2018).
- David A. Boas and Arjun G. Yodh. 1997. Spatially varying dynamical properties of turbid media probed with diffusing temporal light correlation. *JOSA A* (1997).
- Craig F. Bohren and Donald R. Huffman. 1983. *Absorption and scattering of light by small particle*. John Wiley & Sons.
- Julie Chang and Gordon Wetzstein. 2018. Single-shot speckle correlation fluorescence microscopy in thick scattering tissue with image reconstruction priors. *Journal of Biophotonics* (2018).
- Wai-Fung Cheong, Scott A. Prael, and Ashley J. Welch. 1990. A review of the optical properties of biological tissues. *IEEE JQE* (1990).
- Youngwoon Choi, Taeseok Daniel Yang, Christopher Fang-Yen, Pilsung Kang, Kyoung Jin Lee, Ramachandra R. Dasari, Michael S. Feld, and Wonshik Choi. 2011. Overcoming the Diffraction Limit Using Multiple Light Scattering in a Highly Disordered Medium. *Phys. Rev. Lett.* (2011).
- Tom Cuypers, Tom Haber, Philippe Bekaert, Se Baek Oh, and Ramesh Raskar. 2012. Reflectance Model for Diffraction. *ACM TOG* (2012).
- Eugene d'Eon. 2018. A reciprocal formulation of nonexponential radiative transfer. 1: Sketch and motivation. *Journal of Computational and Theoretical Transport* (2018).
- Ronald L. Dougherty, Bruce J. Ackerson, Nafaa M. Reguigui, F. Dorri-Nowkhorani, and Ulf Nöbbmann. 1994. Correlation transfer: Development and application. *JQSRT* (1994).
- Turgut Durduran, Regine Choe, Wesley B. Baker, and Arjun G. Yodh. 2010. Diffuse optics for tissue monitoring and tomography. *Reports on Progress in Physics* (2010).
- Eitan Edrei and Giuliano Scarcelli. 2016. Memory-effect based deconvolution microscopy for super-resolution imaging through scattering media. *Scientific Reports* (2016).
- Robert K. Erf. 1978. Speckle Metrology. Elsevier.
- Shechao Feng, Charles Kane, Patrick A. Lee, and A. Douglas Stone. 1988. Correlations and fluctuations of coherent wave transmission through disordered media. *Phys. Rev. Lett.* (1988).
- Isaac Freund. 1990. Looking through walls and around corners. *Physica A* (1990).
- Isaac Freund and Danny Eliyahu. 1992. Surface correlations in multiple-scattering media. *Phys. Rev. A* (1992).
- David L. Fried. 1982. Anisoplanatism in adaptive optics. *JOSA* (1982).
- Jeppe Revall Frisvad, Niels Jørgen Christensen, and Henrik Wann Jensen. 2007. Computing the scattering properties of participating media using Lorenz-Mie theory. *ACM TOG* (2007).

- Ioannis Gkioulekas, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin. 2013. Inverse Volume Rendering with Material Dictionaries. *ACM TOG* (2013).
- Walter I. Goldburg. 1999. Dynamic light scattering. *American Journal of Physics* (1999).
- Joseph W. Goodman. 1968. *Introduction to Fourier Optics*. McGraw-Hill.
- Joseph W. Goodman. 2007. *Speckle Phenomena in Optics: Theory and Applications*. Roberts and Company Pub.
- Charles Han, Bo Sun, Ravi Ramamoorthi, and Eitan Grinspun. 2007. Frequency domain normal map filtering. *ACM TOG* (2007).
- Roarke Horstmeyer, Haowen Ruan, and Changhui Yang. 2015. Guidestar-assisted wavefront-shaping methods for focusing light into biological tissue. *Nature Photonics* (2015).
- Pierre Jacquot and Jean-Marc Fournier. 2000. *Interferometry in Speckle Light*. Springer.
- Pierre Jacquot and Pramod K. Rastogi. 1979. Speckle motions induced by rigid-body movements in free-space geometry: an explicit investigation and extension to new cases. *Appl. Opt.* (1979).
- Wenzel Jakob, Adam Arbree, Jonathan T. Moon, Kavita Bala, and Steve Marschner. 2010. A radiative transfer framework for rendering materials with anisotropic structure. *ACM TOG* (2010).
- Michael L. Jakobsen, Hal T. Yura, and Steen G. Hanson. 2012. Spatial filtering velocimetry of objective speckles for measuring out-of-plane motion. *Appl. Opt.* (2012).
- Adrian Jarabo, Carlos Aliaga, and Diego Gutierrez. 2018. A radiative transfer framework for spatially-correlated materials. *ACM TOG* (2018).
- Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. 2001. A practical model for subsurface light transport. *ACM TOG* (2001).
- Benjamin Judkewitz, Roarke Horstmeyer, Ivo Vellekoop, and Changhui Yang. 2014. Translation correlations in anisotropically scattering media. *Nature Physics* (2014).
- Ori Katz, Yaron Bromberg, Eran Small, and Yaron Silberberg. 2010. Focusing and Compression of Ultrashort Pulses through Scattering Media. *Nature Photonics* (2010).
- Ori Katz, Pierre Heidmann, Mathias Fink, and Sylvain Gigan. 2014. Non-invasive single-shot imaging through scattering layers and around corners via speckle correlation. *Nature Photonics* (2014).
- Ori Katz, Eran Small, and Yaron Silberberg. 2012. Looking around corners and through thin turbid layers in real time with scattered incoherent light. *Nature Photonics* (2012).
- Guillermo H. Kaufmann. 2011. *Advances in Speckle Metrology and Related Techniques*. Wiley.
- Craig Kolb, Don Mitchell, and Pat Hanrahan. 1995. A Realistic Camera Model for Computer Graphics. *ACM TOG* (1995).
- Puxiang Lai, Lidai Wang, Jian Wei Tay, and Lihong V. Wang. 2015. Photoacoustically guided wavefront shaping for enhanced optical focusing in scattering media. *Nature Photonics* (2015).
- Yunzhe Li, Yujia Xue, and Lei Tian. 2018. Deep speckle correlation: a deep learning approach toward scalable imaging through scattering media. *Optica* (2018).
- Qiang Lu, Xiaosong Gan, Min Gu, and Qingming Luo. 2004. Monte Carlo modeling of optical coherence tomography imaging through turbid media. *Applied optics* (2004).
- Kanti Mardia and Peter Jupp. 2000. *Directional statistics*. John Wiley & Sons.
- Jerome Mertz. 2019. *Introduction to Optical Microscopy*. Cambridge University Press.
- Jerome Mertz, Hari Paudel, and Thomas G. Bifano. 2015. Field of view advantage of conjugate adaptive optics in microscopy applications. *Applied Optics* (2015).
- Christopher A. Metzler, Felix Heide, Prasana Rangarajan, Muralidhar Madabhushi Balaji, Aparna Viswanath, Ashok Veeraraghavan, and Richard G. Baraniuk. 2020. Deep-inverse correlative: towards real-time high-resolution non-line-of-sight imaging. *Optica* (2020).
- Christopher A. Metzler, Philip Schniter, Ashok Veeraraghavan, and Richard G. Baraniuk. 2018. prDeep: Robust phase retrieval with a flexible deep network. *ICML* (2018).
- Allard P. Mosk, Ad Lagendijk, Geoffroy Lerosey, and Mathias Fink. 2012. Controlling waves in space and time for imaging and focusing in complex media. *Nature Photonics* (2012).
- Marco Mout, Michael Wick, F. Bociort, Joerg Petschulat, and Paul Urbach. 2016. Simulating multiple diffraction in imaging systems using a path integration method. *Applied Optics* (2016).
- Srinivasa G. Narasimhan, Mohit Gupta, Craig Donner, Ravi Ramamoorthi, Shree K. Nayar, and Henrik Wann Jensen. 2006. Acquiring scattering properties of participating media by dilution. *ACM TOG* (2006).
- Merlin Nimier-David, Sébastien Speierer, Benoît Ruiz, and Wenzel Jakob. 2020. Radiative backpropagation: an adjoint method for lightning-fast differentiable rendering. *ACM TOG* (2020).
- Micha Nixon, Ori Katz, Eran Small, Yaron Bromberg, Asher A. Friesem, Yaron Silberberg, and Nir Davidson. 2013. Real-time wavefront shaping through scattering media by all-optical feedback. *Nature Photonics* (2013).
- Jan Novak, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. 2018. Monte Carlo Methods for Volumetric Light Transport Simulation. *Computer Graphics Forum* (2018).
- Gerwin Osnabrugge, Roarke Horstmeyer, Ioannis N. Papadopoulos, Benjamin Judkewitz, and Ivo M. Vellekoop. 2017. Generalized optical memory effect. *Optica* (2017).
- Yingtian Pan, Reginald Birngruber, Jürgen Rosperich, and Ralf Engelhardt. 1995. Low-coherence optical tomography in turbid tissue: theoretical analysis. *Applied optics* (1995).
- David J. Pine, David A. Weitz, Paul M. Chaikin, and Eric Herbolzheimer. 1988. Diffusing wave spectroscopy. *Physical review letters* (1988).
- Joseph Rosen and David Abookasis. 2003. Seeing through biological tissues using the fly eye principle. *Optics Express* (2003).
- Markus Rueckel, Julia A. Mack-Bucher, and Winfried Denk. 2006. Adaptive wavefront correction in two-photon microscopy using coherence-gated wavefront sensing. *Proceedings of the National Academy of Sciences* (2006).
- Iman Sadeghi, Adolfo Munoz, Philip Laven, Wojciech Jarosz, Francisco Seron, Diego Gutierrez, and Henrik Jensen. 2012. Physically-Based Simulation of Rainbows. *ACM TOG* (2012).
- John Sawicki, Nikolas Kastor, and Min Xu. 2008. Electric field Monte Carlo simulation of coherent backscattering of polarized light by a turbid medium containing Mie scatterers. *Optical Express* (2008).
- Joseph M. Schmitt and A. Knüttel. 1997. Model of optical coherence tomography of heterogeneous tissue. *JOSA A* (1997).
- Sam Schott, Jacopo Bertolotti, Jean-Francois Léger, Laurent Bourdieu, and Sylvain Gigan. 2015. Characterization of the angular memory effect of scattered light in biological tissues. *Optics Express* (2015).
- Brandon M. Smith, Pratham Desai, Vishal Agarwal, and Mohit Gupta. 2017. CoLux: Multi-object 3D Micro-motion Analysis Using Speckle Imaging. *ACM TOG* (2017).
- Cyril Soler, Kartic Subr, Frédo Durand, Nicolas Holzschuch, and François Sillion. 2009. Fourier Depth of Field. *ACM TOG* (2009).
- Jos Stam. 1999. Diffraction shaders. *ACM TOG* (1999).
- Frédéric Sur, Benoît Blaysat, and Michel Grédiac. 2018. Rendering deformed speckle images with a Boolean model. *JMIV* (2018).
- Kevin T. Takasaki and Jason W. Fleischer. 2014. Phase-space measurement for depth-resolved memory-effect imaging. *Optical Express* (2014).
- Bertrand Thierry, Xavier Antoine, Chokri Chniti, and Hasan Alzubaidi. 2015. μ -diff: An open-source Matlab toolbox for computing multiple scattering problems by disks. *Computer Physics Communications* (2015).
- Bradley E. Treeby and Ben T. Cox. 2010. k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave-fields. *JBO* (2010).
- Elbert G. van Putten, Duygu Akbulut, Jacopo Bertolotti, Willem. L. Vos, Ad Lagendijk, and Allard P. Mosk. 2011. Scattering Lens Resolves Sub-100 nm Structures with Visible Light. *Phys. Rev. Lett.* (2011).
- Ivo M. Vellekoop and Christof M. Aegerter. 2010. Scattered light fluorescence microscopy: imaging through turbid layers. *Opt. Lett.* (2010).
- 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, Aart Lagendijk, and Allard P. Mosk. 2010. Exploiting disorder for perfect focusing. *Nature Photonics* (2010).
- Ivo M. Vellekoop and Allard P. Mosk. 2007. Focusing coherent light through opaque strongly scattering media. *Opt. Lett.* (2007).
- Sebastian Werner, Zdravko Velinov, Wenzel Jakob, and Matthias B. Hullin. 2017. Scratch Iridescence: Wave-Optical Rendering of Diffractive Surface Structure. *ACM TOG* (2017).
- Min Xu. 2004. Electric field Monte Carlo simulation of polarized light propagation in turbid media. *Optical Express* (2004).
- Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering specular microgeometry with wave optics. *ACM TOG* (2018).
- Zahid Yaqoob, Demetri Psaltis, Michael Feld, and Changhui Yang. 2008. Optical phase conjugation for turbidity suppression in biological samples. *Nature photonics* (2008).
- Amnon Yariv. 1997. *Optical electronics in modern communications*. New York : Oxford University Press.
- Kane Yee. 1966. Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media. *IEEE TAP* (1966).
- Hengchin Yeh, Ravish Mehra, Zhimin Ren, Lakulish Antani, Dinesh Manocha, and Ming Lin. 2013. Wave-ray Coupling for Interactive Sound Propagation in Large Complex Scenes. *ACM TOG* (2013).
- Cheng Zhang, Bailey Miller, Kai Yan, Ioannis Gkioulekas, and Shuang Zhao. 2020. Path-space differentiable rendering. *ACM TOG* (2020).
- Cheng Zhang, Lifan Wu, Changxi Zheng, Ioannis Gkioulekas, Ravi Ramamoorthi, and Shuang Zhao. 2019. A differential theory of radiative transfer. *ACM TOG* (2019).