A Monte Carlo Framework for Rendering Speckle Statistics in Scattering Media

CHEN BAR, Department of Electrical Engineering, Technion, Israel MARINA ALTERMAN, Department of Electrical Engineering, Technion, Israel IOANNIS GKIOULEKAS, Robotics Institute, Carnegie Mellon University, USA ANAT LEVIN, Department of Electrical Engineering, Technion, Israel

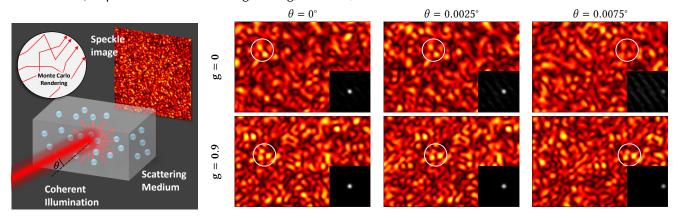


Fig. 1. **Simulation of memory effect in scattering.** Coherent images of translucent materials typically involve highly-fluctuating speckle structure. Despite their semi-random structure, speckles have strong statistical properties. For example, the memory effect property states that, as one tilts the illumination direction (setup at left), the resulting speckles shift. This property is at the core of multiple computational imaging applications. The memory effect is valid over a limited angular range that depends on material properties. Due to the absence of analytical formulas, it is generally necessary to measure this angular range for materials of interest empirically in the lab. We present a Monte Carlo rendering approach for simulating physically-accurate speckle images, as well as their statistics, as a function of material parameters. The figure shows speckle images rendered by our algorithm for a few illumination directions, as well as their auto-correlation (black insets), demonstrating the speckle shift property. As the angle difference increases, the correlation decays, and the decay rate is different material parameters—in this case, materials with Henyey-Greenstein (HG) phase functions of different parameters g. For the isotropic scattering case, g = 0, the pattern similarity is lost at the third column, whereas for the forward scattering case, g = 0.9, correlation is preserved. We verify the accuracy of our algorithm against an exact, yet computationally heavy, wave solver, as well as against analytical formulas derived under limiting assumptions.

We present a Monte Carlo rendering framework for the physically-accurate simulation of speckle patterns arising from volumetric scattering of coherent waves. These noise-like patterns are characterized by strong statistical properties, such as the so-called memory effect. These properties are at the core of imaging techniques for applications as diverse as tissue imaging, motion tracking, and non-line-of-sight imaging. Our rendering framework can replicate these properties computationally, in a way that is orders of magnitude more efficient than alternatives based on directly solving the wave equations. At the core of our framework is a path-space formulation for the covariance of speckle patterns arising from a scattering volume, which

Authors' addresses: Chen Bar, Department of Electrical Engineering, Technion, Haifa, Israel; Marina Alterman, 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 the author(s) 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.

@ 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2019/7-ART39 \$15.00

https://doi.org/10.1145/3306346.3322950

we derive from first principles. We use this formulation to develop two Monte Carlo rendering algorithms, for computing speckle covariance as well as directly speckle fields. While approaches based on wave equation solvers require knowing the microscopic position of wavelength-sized scatterers, our approach takes as input only bulk parameters describing the statistical distribution of these scatterers inside a volume. We validate the accuracy of our framework by comparing against speckle patterns simulated using wave equation solvers, use it to simulate memory effect observations that were previously only possible through lab measurements, and demonstrate its applicability for computational imaging tasks.

CCS Concepts: • Computing methodologies → Ray tracing.

Additional Key Words and Phrases: Scattering, Speckle statistics.

ACM Reference Format:

Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. 2019. A Monte Carlo Framework for Rendering Speckle Statistics in Scattering Media. *ACM Trans. Graph.* 38, 4, Article 39 (July 2019), 22 pages. https://doi.org/10.1145/3306346.3322950

1 INTRODUCTION

Scattering refers to the propagation of radiation (for instance, light or sound) in non-uniform media, composed of small discrete scatterers, usually particles of varying refractive properties: As an incident wave propagates through the medium, it will interact with scatterers multiple times, and each such interaction will change the wave's shape. Scattering is commonly encountered when visible light interacts with a large variety of materials, for instance biological tissues, minerals, the atmosphere and clouds, cosmetics, and many industrial chemicals. As a result of the ubiquity of scattering, its study has attracted numerous research efforts in computer graphics and vision, and much more broadly in medical imaging, remote sensing, seismic imaging, and almost any field of natural science.

The appearance of scattering materials is qualitatively very different, depending on whether they are imaged under incoherent or coherent conditions. In the incoherent case, scaterring results in images with smoothly-varying intensity distributions, often referred to as translucent appearance. By contrast, under coherent imaging conditions, the appearance of scattering materials is characterized by speckles, that is, pseudo-random high variations in the output waves and captured intensity images. Speckles have been the subject of multiple textbooks [Erf 1978; Goodman 2007; Jacquot and Fournier 2000; Kaufmann 2011], as despite their random structure, they have strong statistical properties that are characteristic of the underlying material. For example, a remarkable property of speckles is the memory effect: speckle fields produced under small perturbations in imaging parameters (e.g., change in illumination direction) are highly correlated shifted versions of each other (see Fig. 1). These speckle statistics have received strong attention since the invention of coherent laser illumination [Berkovits and Feng 1994; Feng et al. 1988; Freund et al. 1988; Li and Genack 1994], and are at the core of a large array of imaging techniques, with applications as diverse as motion tracking, estimating blood flow, looking around the corner, and seeing through scattering layers.

Unfortunately, and in stark contrast with the incoherent case, our ability to accurately simulate scattering in the coherent case is severely limited. Available algorithms generally fall into two categories. The first category consists of algorithms that compute output waves by numerically solving Maxwell's equations [Thierry et al. 2015; Treeby and Cox. 2010; Yee 1966]. These algorithms are physically accurate, but require as input the microscopic structure of the scattering medium, that is, knowledge of the exact (at subwavelength accuracy) locations of all scatterers in the medium. Even when such a microscopic characterization is available (e.g., specific samples examined with a microscope, or volumes with hallucinated scatterer locations), the high computational complexity of wave equation solvers makes them inapplicable for volumes larger than a few hundred cubic wavelengths, or containing more than a few hundred scatterers. The second category consists of approximate Monte Carlo rendering algorithms [Sawicki et al. 2008; Xu 2004], which accumulate the complex throughput (amplitude and phase) of paths sampled using standard volumetric path tracing. These algorithms are efficient, but cannot reproduce statistical properties of real speckles such as the memory effect. The lack of speckle

rendering algorithms that are both *physically accurate* and *computationally efficient* is a significant obstacle in the wide range of fields interested in coherent imaging of scattering volumes. Symptomatic of these shortcomings of existing rendering tools is the fact that the only reliable way for estimating the memory effect has been by conducting painstaking optical lab experiments [Schott et al. 2015].

In this paper, we change this state of affairs by developing a Monte Carlo framework for rendering speckles in volumetric scattering. Our framework builds on the following insight: Due to the central limit theorem, speckles are instances of a multivariate Gaussian distribution [Goodman 2007]. Therefore, it is sufficient to model their (scene and material-dependent) mean and covariance. To achieve this, we draw inspiration from Monte Carlo volume rendering algorithms for the incoherent case: These algorithms treat the scattering medium as a continuous volume, inside which light can scatter randomly at any location. Given bulk parameters characterizing the statistical distribution of scatterers in the medium, Monte Carlo algorithms synthesize images corresponding to the average distribution of scattered light across all scatterer instantiations that can be generated from the bulk parameters [Moon et al. 2007]. This macroscopic view of the medium enables efficient rendering, without the need to know and simulate the medium's microscopic structure.

To extend this approach to the coherent case, we begin by deriving a new path-integral formulation [Veach 1997] for the propagation of coherent light inside a scattering medium, which accurately encapsulates the first-order and second-order statistics of resulting speckle patterns. From this formulation, we derive two Monte Carlo rendering algorithms. The first algorithm estimates speckle covariance, which, together with an estimate of speckle mean obtained using a closed-form expression, can be subsequently used to sample multiple speckle images. The second algorithm directly simulates a physically-accurate speckle image, and operates by having sampled paths contribute to multiple pixels in a way that produces accurate speckle statistics. Both algorithms take as input only bulk macroscopic scattering parameters, as in the incoherent case. We validate our theory and algorithms in a few ways: First, we show that our approach can closely match "groundtruth" speckle estimates, obtained by averaging solutions of the wave equation across multiple particle instantiations, while also being orders of magnitude faster. Second, we show that our approach agrees with analytical formulas for speckle correlations derived for specific cases (e.g., diffusion). Finally, we show that our approach can accurately reproduce well-documented properties of speckles, such as the memory effect and coherent backscattering. We show example applications of our framework, including simulating speckle-based computational imaging techniques, and evaluating the extent of their applicability.

1.1 Why render speckle patterns?

There exist several imaging techniques that directly leverage second-order speckle statistics. 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 [Bertolotti et al. 2012; Katz et al. 2014] or focusing through [Mosk et al. 2013; Nixon et al. 2013; Osnabrugge et al. 2017; Vellekoop and Aegerter 2010] tissue and other scattering layers. Most of these imaging techniques rely on

the memory effect of speckles, a fact that has motivated significant research on quantifying this effect for different materials. Existing computational approaches generally attempt to derive closed-form expressions for the memory effect [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; Osnabrugge et al. 2017]. Unfortunately, these expressions only hold under assumptions such as diffusion or the Fokker-Planck limits, restricting their applicability. As a result, it has generally been necessary to measure the memory effect empirically using involved optical setups [Mesradi et al. 2013; Schott et al. 2015; Yang et al. 2014]. Our algorithm allows quantifying the memory effect for arbitrary scattering materials computationally, through accurate yet efficient simulations. This can significantly enhance our understanding of the applicability of memory effect techniques to different materials. Additionally, this new simulation capability can save considerable lab effort for tasks such as discovering optimal settings for computational imaging systems, and evaluating new imaging configurations.

The ability to efficiently render speckle patterns can facilitate the widespread adoption of data-driven approaches in fields where coherent imaging of scattering is common, such as tissue imaging and material science. Previously, the lack of physically-accurate simulation tools meant that training datasets had to be collected using lab measurements, an approach that is not scalable.

Finally, speckle statistics can be beneficial for *inverse rendering*, that is, retrieving material parameters from image measurements. While previous approaches use intensity measurements [Gkioulekas et al. 2016, 2013; Holodovski et al. 2016; Levis et al. 2015], measurements of speckle statistics may capture additional information and allow inverse rendering techniques to be applied in finer scales, where it is not possible to image without coherent effects.

2 RELATED WORK

Monte Carlo rendering of wave optics effects has recently attracted increased attention in computer graphics. A primary focus has been on rendering diffraction and speckle effects generated by surface microgeometry [Bergmann et al. 2016; Cuypers et al. 2012; Stam 1999; Sur et al. 2018; Werner et al. 2017; Yan et al. 2018; Yeh et al. 2013], without tackling volumetric scattering. Some approaches focusing on scattering and speckle effects can be found in the optics literature [Lu et al. 2004; Pan et al. 1995; Schmitt and Knüttel 1997]. For instance, Xu et al. [2008; 2004] modify volumetric path tracing, by tracking complex phase as a path is traced through the volume. By aggregating complex contributions from paths on the sensor, this technique produces images that resemble speckle patterns. However, because every pixel is rendered independently, this approach cannot reproduce spatial correlations between pixels. Additionally, it is impossible to use these approaches to reproduce correlations that exist across multiple illumination directions as in the memory effect.

There have been attempts to use Monte Carlo algorithms to evaluate various properties of coherence and partial coherence of light after propagating through a scattering tissue [Pierrat et al. 2005; Shen et al. 2017]. Often these are based on using the radiative transfer equation (RTE) and intensity-based Monte Carlo rendering, then applying a Fourier transform on its result. Such approaches can be justified as a special case of our algorithm.

An important result in the study of speckle statistics, which can be used to derive Monte Carlo rendering algorithms, is the correlation transfer equation (CTE) [Dougherty et al. 1994; Ishimaru 1999; Twersky 1964]. This integral equation extends the RTE, by modeling correlation of fields at different space points. As we show in Sec. 6.1, there are physical phenomena that are not accounted for by the CTE, such as coherent backscattering. While there exist some Monte Carlo rendering algorithms that take this effect into account [Ilyushin 2012; Sawicki et al. 2008], they only simulate intensity and not general covariance. We revisit the derivation of the CTE and its underlying assumptions, aiming to derive a more general rendering framework that accurately models both covariance and coherent backscattering.

Our derivation is fundamentally based on supplanting the true scattering volume, consisting of multiple discrete scatterers at fixed locations, with a continuous volume where scattering can happen randomly at any location. This macroscopic treatment of scattering underlies all current Monte Carlo volume rendering algorithms, and has also been used to accelerate rendering of so-called discrete random media, where the scatterers can be arbitrarily large or dense [Meng et al. 2015; Moon et al. 2007; Müller et al. 2016]. More recently, a number of works have used this approach to derive generalized versions of the RTE and Monte Carlo rendering algorithms, for media where the distribution of scatterer locations has spatial correlations, so-called non-exponential media [Bitterli et al. 2018; d'Eon 2018a,b; Jarabo et al. 2018]. Even though we focus exclusively on exponential media, our work provides the foundations for future investigations of Monte Carlo rendering of speckles in non-exponential media.

Finally, there is also research on temporal correlations in the presence of scatterer motion, e.g., in liquid dispersions [Berne and Pecora 2000; Dougherty et al. 1994]. Many established techniques use these temporal speckle correlations to estimate flow (e.g., blood flow [Durduran et al. 2010]) and liquid composition parameters. Example techniques include diffusing wave spectroscopy [Pine et al. 1988], laser speckle contrast imaging [Boas and Yodh 1997], and dynamic light scattering [Goldburg 1999]. Here we focus on spatial speckle correlations leaving these temporal effects for future work.

3 MODELING SPECKLE STATISTICS

Setting and notation. We use bold letters for three-dimensional vectors (e.g., points x, i, v), with a circumflex for unit vectors (e.g., directions ω , i, v). We also use \widehat{xy} for the unit vector from x to y. We assume fully-coherent and unpolarized illumination, which can be at either the near or the far field: Near-field illumination is an isotropic source at point i, whereas far-field illumination is a directional planewave source at direction i. Likewise, imaging is done with sensors at either near-field points v or far-field directions v. We often abuse the point notation i, v for both the far-field and near-field cases, except where context requires otherwise.

We consider scattering volumes $\mathcal{V} \in \mathbb{R}^3$ that satisfy four assumptions: First, they consist of scatterers with size comparable to the illumination wavelength, and which can therefore be considered infinitesimal. Second, the scatterers are far from each other, with an average pairwise distance (the mean free path) an order of magnitude



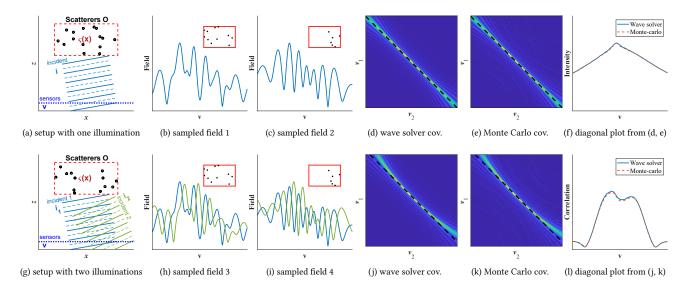


Fig. 2. Simulating speckle and their statistics. (a) Consider a rectangular scattering volume illuminated by a plane wave and a scattered field sensed by collinear sensors. For each scatterer instantiation we solve the wave equation using the package of Thierry et al. [2015] and compute the scattered field, shown in (b, c). Different scatterer positions lead to different high-fluctuation speckle fields. The empirical covariance of multiple fields obtained with the wave solver is demonstrated in (d), and is closely matched by the covariance computed directly by our Monte Carlo algorithm (e). To demonstrate the good agreement we overlay a diagonal plot (f). The diagonal of the speckle covariance is equivalent to intensity images from standard incoherent Monte Carlo algorithms. In the lower row (g) we consider a situation where the same scatterers instantiation is illuminated by two different incident directions highlighting that despite their semi-random structure speckles have strong statistical properties. In particular, the memory effect of speckles: when the same set of scatterers is illuminated by two incident directions the resulting speckle patterns are shifted versions of each other (h-i). This also implies that the covariance of the speckle fields (j) generated by two illumination directions has a shifted diagonal, where the diagonal offset corresponds to fields shift. Our Monte Carlo algorithm is physically correct and captures all such statistics, while having a computational complexity several orders of magnitude smaller than the wave equation solver.

larger than the wavelength. Third, the locations of scatterers are statistically independent. Fourth, scatterers scatter incident waves in a way that is invariant to rotation. These assumptions underly classical radiative transfer [Bitterli et al. 2018]. To simplify notation, in the main paper we derive results assuming scatterers of a single type (same shape, size, and refractive index), and extend to the case of multiple types in App. A.1. We denote by $\varsigma(\mathbf{x}), \mathbf{x} \in \mathcal{V}$ the, possibly spatially-varying, density describing the distribution of scatterers in the medium. Finally, we do not model the interface of volume V, ignoring interface events (reflection and refraction).

The scattered field. An incident wave of wavelength λ interacting with scatterers stimulates a scattered field (or scattered wave) u, which can be computed by solving the Helmholtz equation. When a single particle at location o is illuminated from direction i, the scattered field u at distance $|\mathbf{x} - \mathbf{o}| \gg \lambda$ is

$$u(\mathbf{x}) = \sqrt{c_s} \cdot s\left(\mathbf{i} \cdot \widehat{\mathbf{ox}}\right) \cdot \frac{\exp\left\{ik|\mathbf{x} - \mathbf{o}|\right\}}{|\mathbf{x} - \mathbf{o}|}, \ k = \frac{2\pi}{\lambda}.$$
 (1)

The real scalar c_s is the scattering cross-section, and accounts for the total energy scattered by the scatterer. The complex function $s(\cos \theta)$ is the scattering amplitude function, describing scattering at different angles. We can derive from it the positive probability function $\rho(\cos \theta) = |s(\cos \theta)|^2$, known in computer graphics as the phase function. All three quantities are functions of wavelength and the scatterer's shape, size, interior and exterior refractive index. For spherical scatterers, they can be computed using Mie theory [Bohren

and Huffman 1983; Frisvad et al. 2007]. We note that the scattering amplitude function is often defined as the product $\sqrt{c_s} \cdot s(\cos \theta)$. We separate the two terms and assume that $\rho(\cos \theta)$ integrates to 1.

We now consider the geometry illustrated in Fig. 2a: Scatterers are placed at locations $O = \{\mathbf{o}_1, \mathbf{o}_2, \ldots\}$, each sampled *independently* from the others, from the density $\varsigma(\mathbf{x})$. This configuration is illuminated from a source i, and imaged with a sensor v. Knowing the exact scatterer locations, we can solve the wave equation to obtain the complex-valued scattered field $u_{\mathbf{v}}^{\mathbf{i},O}$, which typically contains large fluctuations with a semi-random noise structure known as speckles (see flatland speckles in Fig. 2b, c).

Speckle statistics. Images modeled with the radiative transfer equation equal the expected intensity of the scattered field, averaged over all particle instantiations O sampled from $\varsigma(\mathbf{x})$, as in Fig. 2f:

$$I_{\mathbf{v}}^{\mathbf{i}} = E_O \left[\left| u_{\mathbf{v}}^{\mathbf{i},O} \right|^2 \right]. \tag{2}$$

These intensity images are typically smooth, without speckles. This is because of the *incoherent addition* in Eq. (2): The expectation is formed by averaging intensities of waves, whereas speckles are the result of coherent addition of complex valued waves. To capture speckle statistics, we can begin with the speckle mean,

$$m_{\mathbf{v}}^{\mathbf{i}} = E_O \left[u_{\mathbf{v}}^{\mathbf{i},O} \right]. \tag{3}$$

We can similarly define higher-order statistics of speckles. Of particular importance will be the speckle covariance,

$$C_{\mathbf{v}_{1},\mathbf{v}_{2}}^{\mathbf{i}_{1},\mathbf{i}_{2}} = E_{O} \left[u_{\mathbf{v}_{1}}^{\mathbf{i}_{1},O} \cdot u_{\mathbf{v}_{2}}^{\mathbf{i}_{2},O^{*}} \right] - m_{\mathbf{v}_{1}}^{\mathbf{i}_{1}} \cdot m_{\mathbf{v}_{2}}^{\mathbf{i}_{2}}^{*}, \tag{4}$$

where $(\cdot)^*$ denotes complex conjugation. In this case, $u_{v_1}^{i_1,O}, u_{v_2}^{i_2,O}$ are two speckle fields generated by the same scatterer configuration O, when illuminated by two incident waves from i_1 , i_2 , and measured at two sensors v_1, v_2 . When $i_1 = i_2 = i, v_1 = v_2 = v$ the term $C_{\mathbf{v},\mathbf{v}}^{\mathbf{i},\mathbf{i}} + |m_{\mathbf{v}}^{\mathbf{i}}|^2$ from Eqs. (3) and (4) reduces to the intensity $I_{\mathbf{v}}^{\mathbf{i}}$ of Eq. (2). As we discuss in Sec. 4.1, the speckle mean can be computed using a closed-form expression; in fact, because the speckle mean is the aggregate of complex numbers of essentially randomly-varying phase, it is typically zero. Therefore, when characterizing speckle statistics, the most challenging part is computing the covariance.

Gaussianity of speckles. Before we discuss ways to compute the speckle mean and covariance, one may wonder whether it is necessary to consider higher-order speckle statistics. The answer, in general, is negative: Classical results in optics [Goodman 2007] state that the space of solutions $u_{\mathbf{v}}^{\mathbf{i},O}$ of the wave equation, for all particle configurations O sampled from $\varsigma(\mathbf{x})$, follows a multivariate Gaussian distribution with scene-dependent mean and covariance. The Gaussianity results from the central limit theorem, as the particle locations are independent random variables. Consequently, the multivariate mean and covariance of Eqs. (3) and (4) provide sufficient statistics for speckle, and can be used to sample speckle patterns that are indistinguisable from patterns generated by specifying exact particle positions and solving the wave equation.

Computing speckle statistics. A straightforward approach for computing the speckle mean and covariance is to sample N different scatterer configurations $O^1, \dots O^N$, solve the wave equation for each configuration, and then compute the empirical moments:

$$m_{\mathbf{v}}^{\mathbf{i}} \approx \frac{1}{N} \sum_{v=1}^{N} u_{\mathbf{v}}^{\mathbf{i},O^n},$$
 (5)

$$C_{\mathbf{v}_{1},\mathbf{v}_{2}}^{\mathbf{i}_{1},\mathbf{i}_{2}} \approx \frac{1}{N} \sum_{n=1}^{N} u_{\mathbf{v}_{1}}^{\mathbf{i}_{1},O^{n}} \cdot u_{\mathbf{v}_{2}}^{\mathbf{i}_{2},O^{n}*} - m_{\mathbf{v}_{1}}^{\mathbf{i}_{1}} \cdot m_{\mathbf{v}_{2}}^{\mathbf{i}_{2}*}.$$
 (6)

Fig. 2(d,k) shows speckle covariances evaluated with this procedure. Solving the wave equation is only tractable for very small number of particles (a few thousands), and this computational cost is further exacerbated by the need to repeat this process multiple times. Our goal is to devise Monte Carlo algorithms that can compute speckle covariance directly and much more efficiently.

Bulk parameters. Unlike wave equation solvers, our algorithms are not tied to a specific configuration of scatterers. Instead, they rely only on the distribution of scatterers in the medium, as well as their size, shape, and refractive properties. As in the radiative transfer literature, we describe these using the phase function $\rho(\cos\theta)$ defined previously, and the scattering, absorption, and extinction coefficients, σ_s , σ_a , σ_t respectively, defined as

$$\sigma_s(\mathbf{x}) = \bar{N}(\mathbf{x})c_s, \quad \sigma_a(\mathbf{x}) = \bar{N}(\mathbf{x})c_a + \sigma_a^{\text{med}}, \quad \bar{N}(\mathbf{x}) = \frac{\varsigma(\mathbf{x})}{4/3\pi r^3}, \quad (7)$$

$$\sigma_t(\mathbf{x}) = \sigma_s(\mathbf{x}) + \sigma_a(\mathbf{x}),$$
 (8)

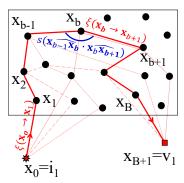


Fig. 3. Fields as path sums. The scattered field can be expressed as the sum of complex throughput contributions $\mu(\vec{x})$ from all possible paths passing through scatterers in a configuration O.

where: c_s , c_a are the scattering and absorption cross-sections, corresponding to the energy scattered or absorbed upon interaction with one particle; r is the radius of the particles; $\bar{N}(\mathbf{x})$ is the expected number of particles in the unit volume at location x; and σ_a^{med} is the absorption coefficient of the containing medium, determined by the imaginary part of the medium's refractive index. We also often reference the mean free path MFP = $1/\sigma_t$. The above definitions consider only particles of a single type, but in App. A.1 we extend them to multiple particle types.

PATH-SPACE VIEW OF SPECKLE STATISTICS

In this section, we derive path-space expressions for the speckle mean and covariance. These expressions will form the basis for the Monte Carlo rendering algorithms of Sec. 5. We note that, traditionally in computer graphics, path-space expressions are derived by recursively expanding integral equations such as the surface and volume rendering equations. Here, we start directly with a pathspace view, and discuss the relationship with an integral equation known as the correlation transfer equation (CTE) in App. A.4.

Fields as path sums. Our starting point is the classical theory of Twersky [1964]: Given a configuration O of scatterers, we can approximate the solution to the Helmholtz equation as the sum of contributions over all paths \vec{x} through O. That is, consider the (enumerable) set $\mathbb{P}_{\mathbf{V}}^{\mathbf{i},O}$ of all ordered sequences:

$$\vec{\mathbf{x}} = \mathbf{x}_0 \to \ldots \to \mathbf{x}_{B+1}$$
, with $\mathbf{x}_0 = \mathbf{i}$, $\mathbf{x}_{B+1} = \mathbf{v}$, $\mathbf{x}_1, \ldots, \mathbf{x}_B \in O$, (9)

where $B = 0, ..., \infty$. Then, the scattered field can be expressed as

$$u_{\mathbf{v}}^{\mathbf{i},O} = \sum_{\vec{\mathbf{x}} \in \mathbb{P}_{\mathbf{v}}^{\mathbf{i},O}} \mu(\vec{\mathbf{x}}) = \sum_{\vec{\mathbf{x}} \in \mathbb{P}_{\mathbf{v}}^{\mathbf{i},O}} \mu(\mathbf{x}_0 \to \mathbf{x}_1) \prod_{b=1}^{B} \mu(\mathbf{x}_{b-1} \to \mathbf{x}_b \to \mathbf{x}_{b+1}). \quad (10)$$

These paths are shown in Fig. 3. The *complex throughput* terms $\mu(\cdot)$ describe the amplitude and phase changes at each path segment, accounting for the scattering amplitude s and traveled length:

$$\mu(\mathbf{x}_{b-1} \to \mathbf{x}_b \to \mathbf{x}_{b+1}) = \xi(\mathbf{x}_b \to \mathbf{x}_{b+1}) s(\widehat{\mathbf{x}_{b-1}} \widehat{\mathbf{x}_b} \cdot \widehat{\mathbf{x}_b} \widehat{\mathbf{x}_{b+1}}), \quad (11)$$

$$\mu(\mathbf{x}_0 \to \mathbf{x}_1) = \xi(\mathbf{x}_0 \to \mathbf{x}_1). \tag{12}$$

The *complex transmission* terms $\xi(\cdot)$ account for phase change and radial decay between path vertices x_b, x_{b+1} , defined for points at the near field and far field, respectively, as

$$\xi(\mathbf{x}_1 \rightarrow \mathbf{x}_2) = \frac{e^{ik|\mathbf{x}_1 - \mathbf{x}_2|}}{|\mathbf{x}_1 - \mathbf{x}_2|}, \ \xi(i \rightarrow \mathbf{x}) = e^{ik(i \cdot \mathbf{x})}, \ \xi((\mathbf{x} \rightarrow \mathbf{v}) = e^{-ik(\mathbf{v} \cdot \mathbf{x})}.$$
 (13)

We note thatm for a fixed configuration O of scatterers, the complex transmission $\xi(\mathbf{x}_1 \rightarrow \mathbf{x}_2)$ does not have an attenuation term. As we see below, volumetric attenuation comes into play only once we start considering multiple random scatterer configurations.

Speckle statistics as path integrals. Using Eq. (10), we can now express the mean and covariance by averaging over all particle configurations O that can be sampled from the density ς :

$$m_{\mathbf{v}}^{\mathbf{i}} = E_O \left[\sum_{\mathbf{x} \in \mathbb{P}^{\mathbf{i}, O}} \mu(\mathbf{x}) \right],$$
 (14)

$$C_{\mathbf{v}_{1},\mathbf{v}_{2}}^{\mathbf{i}_{1},\mathbf{i}_{2}} = E_{O} \left[\sum_{\vec{\mathbf{x}}^{1} \in \mathbb{P}_{\mathbf{v}_{1}}^{\mathbf{i}_{1},O}, \vec{\mathbf{x}}^{2} \in \mathbb{P}_{\mathbf{v}_{2}}^{\mathbf{i}_{2},O}} \mu(\vec{\mathbf{x}}^{1}) \cdot \mu(\vec{\mathbf{x}}^{2})^{*} \right] - m_{\mathbf{v}_{1}}^{\mathbf{i}_{1}} \cdot m_{\mathbf{v}_{2}}^{\mathbf{i}_{2},*}.$$
(15)

Note that, within the expectation, the summation is over paths $\vec{\mathbf{x}}^1$, $\vec{\mathbf{x}}^2$ through the *same* particle instantiation *O*. By exchanging the order of expectation and summation in Eqs. (14) and (15), we have:

$$m_{\rm V}^{\rm i} = \int_{\mathbb{P}^{\rm i}} p(\vec{\mathbf{x}}) \mu(\vec{\mathbf{x}}) \, \mathrm{d}\vec{\mathbf{x}},\tag{16}$$

$$C_{\mathbf{v}_{1},\mathbf{v}_{2}}^{\mathbf{i}_{1},\mathbf{i}_{2}} = \iint_{\mathbb{P}_{\mathbf{v}_{1}}^{\mathbf{i}_{1}},\mathbb{P}_{\mathbf{v}_{2}}^{\mathbf{i}_{2}}} p(\vec{\mathbf{x}}^{1},\vec{\mathbf{x}}^{2}) \mu(\vec{\mathbf{x}}^{1}) \mu(\vec{\mathbf{x}}^{2})^{*} d\vec{\mathbf{x}}^{1} d\vec{\mathbf{x}}^{2} - m_{\mathbf{v}_{1}}^{\mathbf{i}_{1}} m_{\mathbf{v}_{2}}^{\mathbf{i}_{2}}^{*}, \quad (17)$$

where now the space $\mathbb{P}_{\mathbf{v}}^{\mathbf{i}}$ includes paths with vertices $\mathbf{x}_1,\dots,\mathbf{x}_B$ that can be *anywhere* in the volume \mathcal{V} , not only on fixed particle locations. Unlike $\mathbb{P}_{\mathbf{v}}^{\mathbf{i},O}$, $\mathbb{P}_{\mathbf{v}}^{\mathbf{i}}$ is not an enumerable space, thus summation is replaced with integration. The term $p(\vec{\mathbf{x}})$ is the probability that the path $\vec{\mathbf{x}}$ is included in the enumerable path space $\mathbb{P}_{\mathbf{v}}^{\mathbf{i},O}$ for some particle configuration O sampled from S. Similarly $P(\vec{\mathbf{x}}^1,\vec{\mathbf{x}}^2)$ is the probability that all vertices on both $\vec{\mathbf{x}}^1,\vec{\mathbf{x}}^2$ are included in the *same* sampled particle configuration O.

In the following sections, we show that $m_{\mathbf{v}}^{\mathbf{i}}$ can be computed in closed form, and we greatly simplify the path integral for $C_{\mathbf{v}_1,\mathbf{v}_2}^{\mathbf{i}_1,\mathbf{i}_2}$ by characterizing the pairs of paths that have non-zero contributions.

4.1 The speckle mean

Evaluating the speckle mean is addressed by standard textbooks on scattering [Ishimaru 1999; Mishchenko et al. 2006]. We present these results here, starting from a more general case, which subsumes the computation of speckle mean. The general case will also be useful for computing speckle covariance in the next section.

We consider a particle at \mathbf{x}_1 , illuminated from an incident wave with incident direction $\boldsymbol{\omega}$. As this wave scatters, we want to evaluate the average contribution of all paths $\vec{\mathbf{x}}$ starting at \mathbf{x}_1 and arriving at a second point \mathbf{x}_2 . This average equals [Mishchenko et al. 2006]

$$\int_{\mathbb{P}_{\mathbf{v}^2}^{\mathbf{x}_2}} p(\vec{\mathbf{x}}) \mu(\vec{\mathbf{x}}) \, d\vec{\mathbf{x}} = \tau(\mathbf{x}_1, \mathbf{x}_2) \cdot \mu(\boldsymbol{\omega} \rightarrow \mathbf{x}_1 \rightarrow \mathbf{x}_2), \tag{18}$$

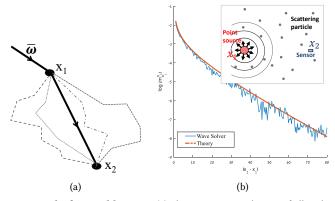


Fig. 4. **Paths for speckle mean.** (a) The average contribution of all paths connecting x_1 and x_2 (dashed lines) reduces to the contribution of the direct path (solid line). (b) We numerically simulate the speckle mean for the setup in the inset. We sample multiple particle configurations, use a wave equation solver to compute the field scattered from a source at point x_1 to a sensor at point x_2 , and average the solutions. The empirical mean of the scattered fields agrees with the speckle mean computed using Eq. (18).

where μ is defined as in Eq. (11). The *volumetric attenuation* τ is the probability of getting from \mathbf{x}_1 to \mathbf{x}_2 without encountering other particles, and equals for the near-field and far-field cases, respectively:

$$\tau(\mathbf{x}_1, \mathbf{x}_2) = e^{-\frac{1}{2} \int_0^1 \sigma_t(\alpha \mathbf{x}_1 + (1 - \alpha) \mathbf{x}_2) \, d\alpha}, \ \tau(\mathbf{i}, \mathbf{x}) = e^{-\frac{1}{2} \int_0^\infty \sigma_t(\mathbf{x}_1 - \alpha \mathbf{i}) \, d\alpha}, \ (19)$$

For a homogeneous medium, $\tau(\mathbf{x}_1, \mathbf{x}_2) = \exp(-\frac{1}{2}\sigma_t|\mathbf{x}_2 - \mathbf{x}_1|)$. The factor 1/2 in the exponent of Eq. (19) makes τ the square root of the volumetric attenuation term in standard radiative transfer. Intuitively, this is because we deal with the field rather than intensity.

The main intuition behind Eq. (18) is that, as most paths contribute essentially random complex phases, they cancel each other out. Therefore, the total field from \mathbf{x}_1 to \mathbf{x}_2 equals the field that travels only along the *direct path* between the two points, attenuated by the exponentially decaying probability $\tau(\mathbf{x}_1, \mathbf{x}_2)$, see Fig. 4(a).

Computing the speckle mean. We can now adapt this result for the speckle mean $m_{\mathbf{v}}^{\mathbf{i}}$ of Eq. (16), which is a special case of Eq. (18). Being the mean of paths from \mathbf{i} to \mathbf{v} without conditioning on an incoming direction $\boldsymbol{\omega}$, we can omit the s term due to scattering, thus

$$m_{\mathbf{v}}^{\mathbf{i}} = \int_{\mathbb{P}_{\mathbf{v}}^{\mathbf{i}}} p(\vec{\mathbf{x}}) \mu(\vec{\mathbf{x}}) \, d\vec{\mathbf{x}} = \tau(\mathbf{i}, \mathbf{v}) \cdot \mu(\mathbf{i} \to \mathbf{v}), \tag{20}$$

where now μ is defined as in Eq. (12) instead of Eq. (11).

The main consequence of this section is that computing the speckle mean becomes a *direct illumination* problem, which can be solved analytically without the need for path integration. In Fig. 4(b), we numerically evaluate the speckle mean by averaging multiple solutions of the wave equation as in Eq. (5), showing a good agreement with the analytic formula of Eq. (20). We note that, as the speckle mean decays exponentially with the distance, in most cases it is negligible, making the computation of covariance the main challenge in simulating speckle. We discuss this next.

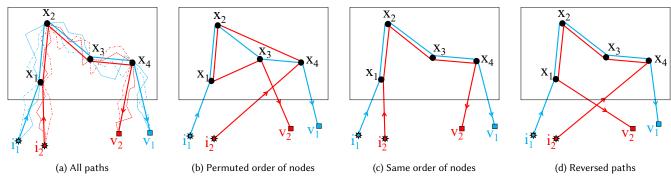


Fig. 5. Path pairs for speckle covariance. (a) When averaging over path pairs, we need only consider pairs that share the same vertices, as path segments from \mathbf{x}_j to \mathbf{x}_{j+1} with arbitrary length, and thus phase, cancel each other on average. (b) Pairs sharing the same vertices in different order also cancel out on average, due to length differences. Only pairs with the same vertices in the same order $x_1, ..., x_B$ (c) or reversed order $x_B, ..., x_1$ contribute to the average.

4.2 The speckle covariance

We have shown in Eq. (17) that the speckle covariance can be expressed as an integral over pairs of paths \vec{x}^1 from i_1 to v_1 and \vec{x}^2 from i_2 to v_2 . Unlike the mean, there is no closed-form expression for this integral. However, we can considerably simplify Eq. (17) by characterizing the pairs of paths $\vec{\mathbf{x}}^1$, $\vec{\mathbf{x}}^2$ for which its integrand

$$c_{\vec{\mathbf{x}}^1 \ \vec{\mathbf{x}}^2} = p(\vec{\mathbf{x}}^1, \vec{\mathbf{x}}^2) \cdot \mu(\vec{\mathbf{x}}^1) \cdot \mu(\vec{\mathbf{x}}^2)^*,$$
 (21)

is non-zero, as well as deriving a simple formula for $c_{\vec{\mathbf{x}}^1,\vec{\mathbf{x}}^2}$ for those pairs. Some of the arguments we use are discussed in Mishchenko et al. [2006]. We formalize these arguments and extend them to accurately account for both speckle covariance and, as we see below, coherent back-scattering. Our end result is a path-integral expression for covariance that lends itself to Monte Carlo integration.

Valid pairs of paths. Intuitively, if we aggregate complex contributions $c_{\vec{\mathbf{v}}^1 \vec{\mathbf{v}}^2}$ from different pairs of paths with very different phases, they will likely average to zero. The exception to this argument is cases where $c_{\vec{\mathbf{x}}^1,\vec{\mathbf{x}}^2}$ is not complex; this happens when every segment $\mathbf{x}_b \to \mathbf{x}_{b+1}$ that appears in $\vec{\mathbf{x}}^1$ also appears in $\vec{\mathbf{x}}^2$.

Consider, as in Fig. 5(a), the set of path pairs $\vec{\mathbf{x}}^1$, $\vec{\mathbf{x}}^2$ that have an arbitrary number of vertices, but share only vertices x_1, \ldots, x_B (in any order). Then, as in Sec. 4.1, we expect all the different path segments from \mathbf{x}_h to \mathbf{x}_{h+1} to average to the direct path between these points. In App. A.2, we prove that indeed all path pairs with disjoint vertices collapse to their joint vertices, and the average contribution of all pairs of paths sharing vertices x_1, \ldots, x_R is

$$c_{\vec{\mathbf{x}}^1.\vec{\mathbf{x}}^2} = v(\vec{\mathbf{x}}^1) \cdot v(\vec{\mathbf{x}}^2)^* \cdot \Pi_{b=1}^B \sigma_s(\mathbf{x}_b), \tag{22}$$

where
$$v(\vec{\mathbf{x}}) = v(\mathbf{x}_0 \to \mathbf{x}_1) \prod_{b=1}^B v(\mathbf{x}_{b-1} \to \mathbf{x}_b \to \mathbf{x}_{b+1}).$$
 (23)

The *complex volumetric throughput* terms $v(\cdot)$ combine the volumetric attenuation of Eq. (19) with the complex throughput of Eqs. (11) and (12). They can be defined as

$$\upsilon(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1}) = \tau(\mathbf{x}_b, \mathbf{x}_{b+1}) \cdot \mu(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1}), \ \ (24)$$

$$v(\mathbf{x}_0 \to \mathbf{x}_1) = \tau(\mathbf{x}_0, \mathbf{x}_1) \cdot \mu(\mathbf{x}_0 \to \mathbf{x}_1). \tag{25}$$

To recap, the complex volumetric throughput is the product of three factors: (i) the volumetric attenuation τ ; (ii) the complex transmission ξ , whose phase is proportional to the path segment length; and

(iii) the scattering amplitude function s due to change of direction (for paths with B > 1). The different terms are summarized in Fig. 6.

We can therefore restrict the integration space of Eq. (17) to only pairs of paths that share all vertices except, perhaps, their endpoints. The contribution of such pairs, given by Eq. (22), is Markovian and can be computed analytically. Next, we further constrain the integration space, by examining when pairs of paths sharing the same vertices but in *different order* have non-zero contribution.

Vertex permutations. We now consider the contribution of a pair of paths sharing the same vertices $\mathbf{x}_1, \dots, \mathbf{x}_B$, but in different permutations. The phase of the segment $\mathbf{x}_b \to \mathbf{x}_{b+1}$ is proportional to the length of that segment. Permutations that do not trace the vertices in the same order have segments of different lengths (Fig. 5b), and thus different phases. Intuitively, as in Sec. 4.1, they are likely to average to zero. However, for each *ordered* set of vertices $\mathbf{x}_1 \to \cdots \to \mathbf{x}_B$, there is one important permutation for which this argument does not apply, as path segments have the same length: the reversed permutation (Fig. 5c and 5d). Therefore, we need to consider contributions from pairs involving four paths [Mishchenko et al. 2006],

$$\vec{\mathbf{x}}^1 = \mathbf{i}_1 \rightarrow \mathbf{x}_1 \rightarrow \dots \rightarrow \mathbf{x}_B \rightarrow \mathbf{v}_1, \qquad \vec{\mathbf{x}}^2 = \mathbf{i}_2 \rightarrow \mathbf{x}_1 \rightarrow \dots \rightarrow \mathbf{x}_B \rightarrow \mathbf{v}_2, \vec{\mathbf{x}}^{1,r} = \mathbf{i}_1 \rightarrow \mathbf{x}_B \rightarrow \dots \rightarrow \mathbf{x}_1 \rightarrow \mathbf{v}_1, \qquad \vec{\mathbf{x}}^{2,r} = \mathbf{i}_2 \rightarrow \mathbf{x}_B \rightarrow \dots \rightarrow \mathbf{x}_1 \rightarrow \mathbf{v}_2.$$
(26)

The reversed paths are the cause of the well-documented phenomenon of coherent backscattering, which occurs when measuring backscattering from a dense scattering volume, with far-field coherent illumination and sensing. When the illumination and sensing directions are exactly equal, the scattered intensity is increased.

For intuition behind this effect, we first note that every particle instantiation O that contains the path $\mathbf{x}_1 \to \cdots \to \mathbf{x}_B$, also contains the reversed path $\mathbf{x}_B \to \cdots \to \mathbf{x}_1$; that is, the forward and reversed paths are not independent events. Consequently, their contribution in Eq. (15) is $(\mu(\vec{x}^1) + \mu(\vec{x}^{1,r})) \cdot (\mu(\vec{x}^2) + \mu(\vec{x}^{2,r}))^*$ rather than $\mu(\vec{\mathbf{x}}^1)\mu(\vec{\mathbf{x}}^2)^* + \mu(\vec{\mathbf{x}}^{1,r})\mu(\vec{\mathbf{x}}^{1,r})^*$. To appreciate the difference between these two terms, we consider the case $i_1 = i_2 = i$, $v_2 = v_2 = v$. Neglecting the scattering amplitude s for simplicity, the contribution of the forward and reversed paths becomes

$$\left|\mu(\vec{\mathbf{x}}) + \mu(\vec{\mathbf{x}}^r)\right|^2 = \left|\xi(\mathbf{i} \rightarrow \mathbf{x}_1)\xi(\mathbf{x}_B \rightarrow \mathbf{v}) + \xi(\mathbf{i} \rightarrow \mathbf{x}_B)\xi(\mathbf{x}_1 \rightarrow \mathbf{v})\right|^2$$
$$\cdot \left|\Pi_{b=1}^{B-1}\xi(\mathbf{x}_b \rightarrow \mathbf{x}_{b+1})\right|^2. \tag{27}$$

The shared intermediate segments have the same phase, therefore,

$$\left| \mu(\vec{\mathbf{x}}) + \mu(\vec{\mathbf{x}}^r) \right|^2 = \left| \xi(\mathbf{i} \rightarrow \mathbf{x}_1) \xi(\mathbf{x}_B \rightarrow \mathbf{v}) + \xi(\mathbf{i} \rightarrow \mathbf{x}_B) \xi(\mathbf{x}_1 \rightarrow \mathbf{v}) \right|^2$$

$$= \left| e^{ik(\mathbf{i}^T \mathbf{x}_1 - \mathbf{v}^T \mathbf{x}_B)} + e^{ik(\mathbf{i}^T \mathbf{x}_B - \mathbf{v}^T \mathbf{x}_1)} \right|^2$$

$$= 2 + 2Re \left(e^{ik(\mathbf{i} + \mathbf{v})^T (\mathbf{x}_1 - \mathbf{x}_B)} \right). \tag{28}$$

When i + v is large, the average of the real term in Eq. (28) over all space points is low. However, when $i \sim -v$, as in coherent backscattering, the real term approaches unity, and therefore the total contribution is doubled. In other words, we get *constructive interference* between the forward and reversed paths.

Covariance path integral. We can now state concretely our path integral formulation for speckle covariance. Consider the space \mathbb{P} of *sub-paths* $\vec{\mathbf{x}}^s = \mathbf{x}_1 \to \cdots \to \mathbf{x}_B$, where each vertex can be everywhere in \mathcal{V} , and $B = 0, \ldots, \infty$. Then, we can write:

$$C_{\mathbf{v}_1,\mathbf{v}_2}^{\mathbf{i}_1,\mathbf{i}_2} = \int_{\mathbb{P}} c(\vec{\mathbf{x}}^s) \, d\vec{\mathbf{x}}^s - m_{\mathbf{v}_1}^{\mathbf{i}_1} \cdot m_{\mathbf{v}_2}^{\mathbf{i}_2}^*.$$
 (29)

To define the integrand $c(\vec{\mathbf{x}}^s)$, we first form the four complete paths of Eq. (26), by connecting the forward and reversed versions of $\vec{\mathbf{x}}^s$ to the illumination and sensing conditions i_1 , v_1 and i_2 , v_2 . Then,

$$c(\vec{\mathbf{x}}^{s}) = c_{\vec{\mathbf{x}}^{1}, \vec{\mathbf{x}}^{2}} + c_{\vec{\mathbf{x}}^{1}, \vec{\mathbf{x}}^{2}, r} + c_{\vec{\mathbf{x}}^{1}, r, \vec{\mathbf{x}}^{2}} + c_{\vec{\mathbf{x}}^{1}, r, \vec{\mathbf{x}}^{2}, r}, \tag{30}$$

where the summands are defined in Eq. (22). By expanding the equations, and considering that now the pairs of paths have identical intermediate segments, we can rewrite this sum as

$$c(\vec{\mathbf{x}}^{s}) = f(\vec{\mathbf{x}}^{s}) \cdot \left(v(\mathbf{x}_{2} \to \mathbf{x}_{1} \to \mathbf{i}_{1}) v(\mathbf{x}_{B-1} \to \mathbf{x}_{B} \to \mathbf{v}_{1}) + v(\mathbf{x}_{B-1} \to \mathbf{x}_{B} \to \mathbf{i}_{1}) v(\mathbf{x}_{2} \to \mathbf{x}_{1} \to \mathbf{v}_{1}) \right) \cdot \left(v(\mathbf{x}_{2} \to \mathbf{x}_{1} \to \mathbf{i}_{2}) v(\mathbf{x}_{B-1} \to \mathbf{x}_{B} \to \mathbf{v}_{2}) + v(\mathbf{x}_{B-1} \to \mathbf{x}_{B} \to \mathbf{i}_{2}) v(\mathbf{x}_{2} \to \mathbf{x}_{1} \to \mathbf{v}_{2}) \right)^{*}, \quad (31)$$

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

$$f(\vec{\mathbf{x}}^{S}) = |v(\mathbf{x}_{1} \rightarrow \mathbf{x}_{2})|^{2} \prod_{b=2}^{B-1} |v(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_{b} \rightarrow \mathbf{x}_{b+1})|^{2} \prod_{b=1}^{B} \sigma_{S}(\mathbf{x}_{b})$$

$$= \sigma_{S}(\mathbf{x}_{1}) \sigma_{S}(\mathbf{x}_{B}) \tau^{2}(\mathbf{x}_{1}, \mathbf{x}_{2})$$

$$\prod_{b=2}^{B-1} \rho(\widehat{\mathbf{x}_{b-1}} \widehat{\mathbf{x}_{b}} \cdot \widehat{\mathbf{x}_{b}} \widehat{\mathbf{x}_{b+1}}) \tau^{2}(\mathbf{x}_{b}, \mathbf{x}_{b+1}) \sigma_{S}(\mathbf{x}_{b}). \tag{32}$$

The connections corresponding to the four complex volumetric throughput terms v in Eq. (31) are illustrated in Fig. 6, whereas $f(\vec{\mathbf{x}}^s)$ is the radiometric throughput of the central segments (gray path in Fig. 6). As we see in the next section, the radiometric throughput term in Eq. (32) allows us to reuse path sampling algorithms from intensity rendering also for covariance rendering.

We note that Eq. (29) allows for B = 0. This corresponds to an empty sub-path $\vec{\mathbf{x}}^s$, and therefore to complete paths that go directly from \mathbf{i}_1 to \mathbf{v}_1 and from \mathbf{i}_2 to \mathbf{v}_2 , without shared vertices. As a final simplification to Eq. (29), we note from Eq. (20) that the term $m_{\mathbf{v}_1}^{\mathbf{i}_1} \cdot m_{\mathbf{v}_2}^{\mathbf{i}_2}$ is exactly equal to the contributions of these direct paths. Therefore, we can remove this term from Eq. (29) by restricting integration to the space $\mathbb{P}_{B\geq 1}$ of sub-paths of length $B\geq 1$:

$$C_{\mathbf{v}_{1},\mathbf{v}_{2}}^{\mathbf{i}_{1},\mathbf{i}_{2}} = \int_{\mathbb{P}_{B>1}} c(\vec{\mathbf{x}}^{s}) \, d\vec{\mathbf{x}}^{s}. \tag{33}$$

We make three observations about the path integral formulation of Eq. (33). First, if we ignore the reversed paths, then the resulting path-integral formulation is equivalent to what can be obtained from the *correlation transfer equation* (CTE). We discuss this in App. A.4, and we also discuss how the Monte Carlo algorithms we derive in the next section compare to Monte Carlo algorithms derived from the CTE. In the evaluation of Sec. 6.1 we show that considering only forward paths can provide a good approximation in many cases; however, in cases where the sensor is close to collocated with the source, we should consider reversed paths as well.

Second, at the start of this section, we argued informally that pairs of paths with different permutations of $\mathbf{x}_1, \ldots, \mathbf{x}_B$ do not contribute to covariance. In App. A.3, we discuss this in more detail, and additionally show empirical evidence for ignoring these pairs. Likewise, the results in Sec. 6.1 show that accounting for only the forward and reversed path is accurate enough.

Third, it is worth considering the case of $i_1 = i_2 = i$ and $v_1 = i$ $v_2 = v$. Then, the sum of the covariance with the product of means, $C_{\mathbf{v}_1,\mathbf{v}_2}^{\mathbf{i}_1,\mathbf{i}_2} + m_{\mathbf{v}_1}^{\mathbf{i}_1} \cdot m_{\mathbf{v}_2}^{\mathbf{i}_2}$, becomes equal to the intensity $I_{\mathbf{v}}^{\mathbf{i}}$ of Eq. (2). If we ignore reversed paths, then the path contribution $c(\vec{\mathbf{x}}^s)$ of Eq. (31) reduces to the standard radiometric throughput $f(\vec{x})$ of the complete path $\vec{\mathbf{x}} = \mathbf{i} \rightarrow \vec{\mathbf{x}}^s \rightarrow \mathbf{v}$. Likewise, after adding the product of means to Eq. (33), we obtain for I_{v}^{i} a path-integral expression that exactly matches the one derived from the volume rendering equation [Dutré et al. 2006; Novak et al. 2018; Veach 1997]. When we consider reversed paths, the resulting path-integral expression for $I_{\mathbf{v}}^{\mathbf{i}}$ will be different from the one obtained by the volume rendering equation; the difference corresponds to the observation, previously reported in the literature [Mishchenko et al. 2006], that the radiative transfer and volume rendering equations cannot explain coherent back-scattering. Our derivation suggests a straightforward way to incorporate coherent back-scattering into existing volume rendering algorithms for intensity, by also considering the radiometric throughput of reversed paths. Finally, the fact that our formulation is consistent (up to coherent back-scattering effects) with the radiative transfer equation provides further justification for ignoring pairs of paths with different permutations of the same vertices.

5 MONTE CARLO RENDERING ALGORITHMS

We use the results of the previous section, to derive two Monte Carlo rendering algorithms. The first algorithm directly computes the speckle covariance, which we can use, together with an estimate of the speckle mean, to sample multiple speckle patterns. The second algorithm directly renders a speckle pattern, so that the empirical mean and covariance of multiple renderings is accurate.

 $\begin{array}{l} \xi(\mathbf{x}_b {\rightarrow} \mathbf{x}_{b+1}) = \frac{e^{ik|\mathbf{x}_b - \mathbf{x}_{b+1}|}}{|\mathbf{x}_b - \mathbf{x}_{b+1}|} \\ s(\widehat{\mathbf{x}_{b-1}} \mathbf{x}_{b} \cdot \widehat{\mathbf{x}_b} \mathbf{x}_{b+1}) \end{array}$ Complex transmission: Scattering amplitude function: $\mu(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1}) = \xi(\mathbf{x}_b \rightarrow \mathbf{x}_{b+1}) s(\widehat{\mathbf{x}_{b-1} \mathbf{x}_b} \cdot \widehat{\mathbf{x}_b \mathbf{x}_{b+1}})$ Complex throughput: $\tau(\mathbf{x}_{b}, \mathbf{x}_{b+1}) = e^{-\frac{1}{2}\sigma_{t}|\mathbf{x}_{b}-\mathbf{x}_{b+1}|}$ Volumetric attenuation: $v(\mathbf{x}_{b-1} \!\!\to\!\! \mathbf{x}_b \!\!\to\!\! \mathbf{x}_{b+1}) = \tau(\mathbf{x}_b, \mathbf{x}_{b+1}) \mu(\mathbf{x}_{b-1} \!\!\to\!\! \mathbf{x}_b \!\!\to\!\! \mathbf{x}_{b+1})$ Complex volumetric throughput: $f(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1}) = \sigma_s(\mathbf{x}_b)|v(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1})|^2$ Radiometric throughput:

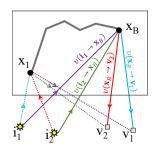


Fig. 6. Types of path contributions. Summary of notation and relationships between different throughput terms used in our Monte Carlo algorithms.

```
ALGORITHM 1: Monte Carlo rendering of covariance C_{\mathbf{v}_1,\mathbf{v}_2}^{\mathbf{i}_1,\mathbf{i}_2}
                                                                           ⊳Initialize covariance estimate.
    Set C = 0.
    for iteration = 1 : N do
                                                                          Sample first vertex of subpath.
            Sample point \mathbf{x}_1 \sim q_o(\mathbf{x}_1).
            Sample uniformly direction \omega_1.
                                            >Update covariance with single scattering path.
            Update
               C += V \cdot v(\mathbf{i}_1 \rightarrow \mathbf{x}_1)v(\mathbf{i}_1 \rightarrow \mathbf{x}_1 \rightarrow \mathbf{v}_1)v(\mathbf{i}_2 \rightarrow \mathbf{x}_1)^*v(\mathbf{i}_2 \rightarrow \mathbf{x}_1 \rightarrow \mathbf{v}_2)^*.
                                                                           ⊳Continue tracing the subpath.
                                                                     Sample second vertex of subpath.
            Sample distance d \sim \sigma_t(\mathbf{x}_1) |\tau(\mathbf{x}_1, \mathbf{x}_1 + \mathbf{d}\omega_1)|^2.
            Set point \mathbf{x}_2 = \mathbf{x}_1 + d \cdot \boldsymbol{\omega}_1.
            Set b = 2.
            while x_b inside medium do
                                            \triangleright Update covariance with next-event estimation.
                    Update C += \frac{V}{2} \left( v(\mathbf{x}_2 \rightarrow \mathbf{x}_1 \rightarrow \mathbf{i}_1) v(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{v}_1) \right)
                                                          +v(\mathbf{x}_{b-1}\rightarrow\mathbf{x}_b\rightarrow\mathbf{i}_1)v(\mathbf{x}_2\rightarrow\mathbf{x}_1\rightarrow\mathbf{v}_1)
                                                      \cdot \left( v(\mathbf{x}_2 \rightarrow \mathbf{x}_1 \rightarrow \mathbf{i}_2) v(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{v}_2) \right)
                                                          +v(\mathbf{x}_{b-1}\rightarrow\mathbf{x}_b\rightarrow\mathbf{i}_2)v(\mathbf{x}_2\rightarrow\mathbf{x}_1\rightarrow\mathbf{v}_2)
                                                                         Sample next vertex of subpath.
                    Sample direction \omega_b \sim \rho(\omega_{b-1} \cdot \omega_b).
                    Sample distance d \sim \sigma_t(\mathbf{x}_b) |\tau(\mathbf{x}_b, \mathbf{x}_b + d\omega_b)|^2.
                    Set point \mathbf{x}_{b+1} = \mathbf{x}_b + d \cdot \boldsymbol{\omega}_b.
                                                                                      \triangleright Account for absorption.
                    Sample scalar a \sim \text{Unif}[0, 1].
                    if a > \sigma_s(\mathbf{x}_{b+1})/\sigma_t(\mathbf{x}_{b+1}) then
                                                          ⊳ Terminate subpath at absorption event.
                           break
                    end
                    Set b = b + 1.
            end
    end
                                                                   ⊳Produce final covariance estimate.
     Update C = \frac{1}{N}C.
```

Rendering speckle covariance

return C.

To approximate the covariance integral of Eq. (33), we use a strategy that samples *sub-paths* $\vec{\mathbf{x}}^{s,n}$ from a distribution $q(\vec{\mathbf{x}}^{s,n})$ defined below. We then form a Monte Carlo estimate of the covariance as

$$C_{\mathbf{v}_1,\mathbf{v}_2}^{\mathbf{i}_1,\mathbf{i}_2} \approx \frac{1}{N} \sum_{n=1}^{N} \frac{c(\vec{\mathbf{x}}^{s,n})}{q(\vec{\mathbf{x}}^{s,n}) + q(\vec{\mathbf{x}}^{s,r,n})}$$
 (34)

ALGORITHM 2: Monte Carlo rendering of $J \times 1$ field **u** for $\{(i, v)_j\}_{j=1}^J$.

⊳Initialize field estimate.

```
Set \mathbf{u} = \mathbf{0}.
for iteration = 1 : N do
       Sample random phase \zeta \sim \text{Unif}[0, 1].
       Set z = e^{2\pi i \zeta}.
                                                                Sample first vertex of subpath.
       Sample point \mathbf{x}_1 \sim q_o(\mathbf{x}_1).
                                                Dupdate field with single scattering path. □
       Update \forall j, \mathbf{u}_j := z \cdot \sqrt{\frac{V}{2}} \cdot v(\mathbf{i}_j \rightarrow \mathbf{x}_1) v(\mathbf{i}_j \rightarrow \mathbf{x}_1 \rightarrow \mathbf{v}_j).
                                                                 ⊳Continue tracing the subpath.
                                                            Sample second vertex of subpath.
       Sample uniformly direction \omega_1.
       Sample distance d \sim \sigma_t(\mathbf{x}_1) |\tau(\mathbf{x}_1, \mathbf{x}_1 + d\omega_1)|^2
       Set point \mathbf{x}_2 = \mathbf{x}_1 + d \cdot \boldsymbol{\omega}_1.
       Set b = 2.
       while xb inside medium do
               Sample random phase \zeta \sim \text{Unif}[0, 1].
              Set z = e^{2\pi i \zeta}.
                                               ⊳Update field with next-event estimation.
                 \forall j, \mathbf{u}_j \mathrel{+=} z \cdot \sqrt{\frac{V}{2}} \cdot \left( v(\mathbf{x}_2 \rightarrow \mathbf{x}_1 \rightarrow \mathbf{i}_j) v(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{v}_j) \right)
                                                      +v(\mathbf{x}_2 \rightarrow \mathbf{x}_1 \rightarrow \mathbf{v}_j)v(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{i}_j).
                                                               Sample next vertex of subpath.
               Sample direction \omega_k \sim \rho(\omega_{b-1} \cdot \omega_b).
               Sample distance d \sim \sigma_t(\mathbf{x}_b) |\tau(\mathbf{x}_b, \mathbf{x}_b + d\omega_b)|^2.
              Set point \mathbf{x}_{b+1} = \mathbf{x}_b + d \cdot \boldsymbol{\omega}_b.
                                                                            ⊳ Account for absorption.
               Sample scalar a \sim \text{Unif}[0, 1].
               if a > \sigma_s(\mathbf{x}_{b+1})/\sigma_t(\mathbf{x}_{b+1}) then
                                                 ⊳ Terminate subpath at absorption event.
                     break
               end
              Set b = b + 1.
       end
                                                    \triangleright Produce final field with correct mean.
```

The denominator of Eq. (34) is the sampling probability. As it is possible to independently sample both the forward and reserved version of a subpath, the total probability is $q(\vec{\mathbf{x}}^s) + q(\vec{\mathbf{x}}^{s,r})$.

Update $\forall j$, $\mathbf{u}_j = m_{\mathbf{v}_j}^{\mathbf{i}_j} + \sqrt{\frac{1}{N}} \mathbf{u}_j$.

return u

The variance of the estimator in Eq. (34) reduces when $q(\vec{\mathbf{x}}^s)$ approximates $c(\vec{\mathbf{x}}^s)$. As $c(\vec{\mathbf{x}}^s)$ in Eq. (31) is Markovian, that is, expressed as a product of the contributions of individual segments, it lends itself to local sampling procedures. The sampling algorithm we use operates as follows: We sample the first vertex \mathbf{x}_1 according to the volume density, using the probability distribution q_o defined as:

$$q_o(\mathbf{x}) = \frac{\sigma_s(\mathbf{x})}{V}$$
 with $V = \int \sigma_s(\mathbf{x}) d\mathbf{x}$. (35)

For a homogeneous volume, q_o reduces to the uniform density. Then, taking advantage of the fact that $c(\vec{\mathbf{x}}^s)$ includes the radiometric throughput $f(\vec{\mathbf{x}}^s)$, we sample all other vertices of $\vec{\mathbf{x}}^s$ using volume path tracing [Dutré et al. 2006; Novak et al. 2018]. Finally, as we trace $\vec{\mathbf{x}}^s$, we perform *next event estimation*, connecting each vertex to the endpoints of the forward and reversed paths of Eq. (26), as shown in Fig. 6. This process is summarized in Alg. 1, which also details how to handle single-scattering subpaths consisting of only one vertex. We note that, in a heterogeneous volume, the exponential sampling of distances d in Alg. 1 is replaced with a tracking algorithm such as Woodcock tracking [Kutz et al. 2017]. When $\mathbf{i}_1 = \mathbf{i}_2, \mathbf{v}_1 = \mathbf{v}_2$, this algorithm reduces to the standard volume path tracing algorithm for rendering intensity, except for the sampling of the first vertex and the addition of reversed paths for coherent back-scattering.

The probability of a sampled sub-path \vec{x}^s sampled as above, and its contribution in Eq. (34), become:

$$q(\vec{\mathbf{x}}^s) = \frac{1}{V} f(\vec{\mathbf{x}}^s), \quad \text{and} \quad \frac{c(\vec{\mathbf{x}}^s)}{q(\vec{\mathbf{x}}^s) + q(\vec{\mathbf{x}}^{s,r})} = \frac{V}{2} \frac{c(\vec{\mathbf{x}}^s)}{f(\vec{\mathbf{x}}^s)}. \tag{36}$$

After term cancellations, we end up with only the terms $v(\cdot)$ in Eq. (31), for the four next event estimation connections in Fig. 6.

5.2 Rendering speckle fields

As discussed in Sec. 3, the space of speckle images follows a multivariate Gaussian distribution. Thus the mean and covariance provide sufficient statistics, which we can use to sample physically-correct speckle images, statistically indistinguishable from ones generated through an exact solution to the wave equation. However, with this approach, sampling an image of J pixels, such that the statistics of all pixels are consistent with each other, requires that we first render an $J \times J$ covariance matrix. While this is significantly more efficient than solving the wave equation, for large J values this can still be costly. To address this, we present a second rendering algorithm that can synthesize speckle images directly.

Our starting point is the following observation: Let C be the $J \times J$ covariance matrix corresponding to all pairwise combinations of J illumination and sensing conditions $\{(\mathbf{i},\mathbf{v})_j\}_{j=1}^J$. Then, from Eqs. (33) and (31), we can write C as an integral of rank-1 matrices,

$$C = \int_{\mathbb{P}} f(\vec{\mathbf{x}}^s) \cdot \mathbf{a}(\vec{\mathbf{x}}^s) \cdot \mathbf{a}^*(\vec{\mathbf{x}}^s) \, d\vec{\mathbf{x}}^s, \tag{37}$$

where: $f(\vec{\mathbf{x}}^s)$ is defined in Eq. (32), and $\mathbf{a}(\vec{\mathbf{x}}^s)$ is a $J \times 1$ vector with j-th entry equal to the $v(\cdot)$ terms in Eq. (31) applied to \mathbf{i}_j and \mathbf{v}_j ,

$$\mathbf{a}_{j}(\vec{\mathbf{x}}^{s}) = \left(v(\mathbf{x}_{2} \to \mathbf{x}_{1} \to \mathbf{i}_{j}) v(\mathbf{x}_{B-1} \to \mathbf{x}_{B} \to \mathbf{v}_{j}) + v(\mathbf{x}_{B-1} \to \mathbf{x}_{B} \to \mathbf{i}_{j}) v(\mathbf{x}_{2} \to \mathbf{x}_{1} \to \mathbf{v}_{j}) \right).$$
(38)

Sampling a $J \times 1$ field **u** from a multivariate Gaussian with a covariance as in Eq. (37) can be done by first initializing **u** to the zero vector, then repeating the following: (i) Sample a subpath $\vec{\mathbf{x}}^s$

as in Alg. 1. (ii) Sample a complex number z of unit magnitude and random phase. (iii) Increment \mathbf{u} by $\sqrt{\frac{1}{N}} \cdot \sqrt{\frac{V}{2}} \cdot z \cdot \mathbf{a}(\vec{\mathbf{x}}^s)$ (where $\sqrt{V/2}$ is the square root of the scale in Eq. (36)). This is summarized in Alg. 2, which also shows how to handle single-scattering subpaths.

We elaborate on two details of the above procedure: First, a single sample drawn according to Alg. 2 has the right covariance, but may not follow a Gaussian distribution. By averaging multiple samples, the central limit theorem implies that their average will converge to a Gaussian distribution. To keep the total variance of the average independent of the number of samples N, we scale each sample by $\sqrt{1/N}$. Second, we draw the random variable z to ensure that the mean of the samples is zero; we subsequently add the desired mean (computed as described in Sec. 4.1) to the final estimate.

Relationship to path tracing. Like Alg. 1, Alg. 2 is also closely related to volumetric path tracing for rendering intensity: The weight v is a complex square root of the next-event-estimation weight used for intensity. We can see this from Eq. (24), where v is defined as the product of: (i) the amplitude function s, which is the complex square root of the phase function ρ ; (ii) the volumetric attenuation τ , which is the square root of the volumetric attenuation for intensity; and (iii) the unit-magnitude transmission ξ .

We note, however, a critical difference: In Alg. 2, every sampled subpath is used to update *all sources and sensors*. This is the key for generating speckle images with accurate second-order statistics, and is the fundamental difference with previous speckle rendering algorithms [Sawicki et al. 2008; Xu 2004]. As those algorithms update different pixels *independently*, they cannot reproduce correlations between pixels or across different illumination conditions. We demonstrate this in Sec. 6.1.

6 EXPERIMENTS AND APPLICATIONS

We perform three sets of experiments. First, we validate the accuracy of our algorithms by comparing with a wave equation solver. Second, we use our algorithms to quantify the memory effect. Third, we reproduce computational imaging techniques based on that effect.

6.1 Validation against a wave-solver

To validate our rendering algorithms, we compare their outputs with "groundtruth" obtained as in Eq. (6), by first solving the Helmholtz equation for multiple scatterer configurations, and then computing the empirical statistics of the resulting scattered fields.

Wave equation solvers. We experimented with two types of solvers, both of which are well-established in the optics literature as accurate simulators that can be used to validate experimental measurements, albeit with very high computational cost. The first type of solvers are based on finite-difference time-domain (FDTD) methods [Treeby and Cox. 2010; Yee 1966], which voxelize the simulated volume at a sub-wavelength resolution, and have memory and computation complexity that scales with the size of the resulting grid. This makes FDTD solvers unsuitable for our experiments, which require volumes of linear dimension at least an order of magnitude larger than the wavelength, due to the far-field assumption for scatterers and our desire to model multiple scattering.

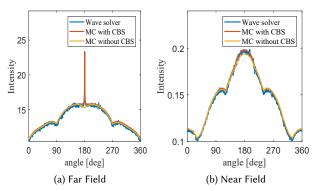


Fig. 7. Coherent back-scattering (CBS). We used a $100\lambda \times 100\lambda$ target with OD = 2 to validate CBS, and simulate Intensity as a function of sensor angle. (a) In the far-field case, when including both forward and reversed paths (red), our algorithm closely matches the intensity of the wave solver (blue). Neglecting reversed paths (orange) results in a mismatch at the exact back-scattering direction. (b) In the near-field case, due to the absence of CBS, both versions of our algorithm agree with the wave solver.

The second type of solvers use the integral version of the Helmholtz equation, and have a complexity that scales with the number of scatterers rather than volume size. This is significantly more efficient for our setting, though the complexity is still cubic in the number of scatterers. Therefore, even integral solvers become impractical for volumes with more than a few thousand scatterers. As an additional advantage, integral solvers produce higher-quality solutions, because of better boundary-handling properties. For our validation experiments, we use the μ -diff toolbox [Thierry et al. 2015], which is restricted to 2D volumes. Consequently, the experiments of this sub-section are all performed in 2D. We emphasize that this is for this sub-section only, and that the experiments of Secs. 6.2 and 6.3 are performed using full 3D simulations.

Coherent backscattering. Fig. 7 demonstrates coherent backscattering intensity, rendered using our algorithm with $i_1 = i_2$, $v_1 = v_2$. We use a target of size $100\lambda \times 100\lambda$, with a mean free path of 50λ , leading to an *optical depth* (that is, average number of scattering events) OD = 2. We simulate far-field sensors through all 360° around the target, and near-field sensors located on a 360° circle of diameter 200λ around the target. We compare the mean speckle intensity obtained from the electromagnetic solver with our Monte Carlo algorithm, considering forward and reversed paths, and with a simpler algorithm considering only forward paths derived in App. A.4. For far-field sensors, we see that when the viewing direction approaches the inverse of the illumination direction, a narrow peak in brightness occurs, which is the manifestation of coherent backscattering. This peak is not predicted when using forward-only paths, but is indeed explained when using both forward and reversed paths. For near-field sensors, coherent backscattering is less pronounced and the outputs of two Monte Carlo algorithms are closer to each other.

Memory effect. In Fig. 8, we show simulated covariance matrices for a target of size $20\lambda \times 20\lambda$ at OD = 2 and 0.5. The particles have a

radius of 0.48λ , and their phase function is computed using Mie theory. We visualize covariance matrices of a target illuminated by two plane waves, measured at the far-field over 360° viewing directions. In the covariance matrices, the memory effect is evident by the fact that, for small angles (e.g., $i = 1^{\circ}$ in Fig. 8), the strongest correlation is obtained at a diagonal that is offset from the main diagonal, and the offset increases with the illumination angle difference. When the angle difference is large (e.g., $i = 20^{\circ}$ in Fig. 8), the classical version of the memory effect no longer holds and the covariance is no longer a shifted diagonal. However, we can still observe some correlation along a curved set of viewing directions. To the best of our knowledge, such correlations have not yet been explored, and provide an exciting direction of future research. In particular they may allow for expanding the angular range of existing computational imaging techniques relying on the memory effect. We note also that, while the shape of the correlation curve is consistent, its exact value is a function of density, as seen from the two optical depths simulated in Fig. 8.

Runtime comparison. Figs. 7 and 8 show that our Monte Carlo algorithms provide accurate predictions of speckle correlations, while being orders of magnitudes more efficient than the wave solver. To quantify the performance difference, in the example of Fig. 8, simulating the covariance with the wave solver approach took six hours on a 50-core cluster, using the μ -diff solver [Thierry et al. 2015]. By contrast, our Monte Carlo algorithm produced the same estimate in 45 minutes on a single core, using an unoptimized Matlab implementation. The difference in performance becomes even more pronounced as the number of scatterers increases.

Field samples. We use Alg. 2 to sample multiple speckle fields, shown in Fig. 1, for target size and densities equivalent to the setup of Fig. 8 at OD = 2. We use these to compute an empirical covariance, shown n Fig. 9, which is in agreement with the covariance rendered directly using Alg. 1. We also compare with the "electric field Monte Carlo" (EMC) speckle rendering algorithm [Sawicki et al. 2008; Xu 2004]. This approach extends standard volumetric path tracing, by using the length of traced paths as complex phase. Its main difference with our algorithms is that each sampled path is used to update only one sensor point, and therefore different illumination and viewing directions are updated independently. As a consequence, while this approach can accurately render intensity and even simulate coherent backscattering, it cannot reproduce spatial correlation. We note, though, that while EMC also focuses on modeling polarization correctly, we do not account for polarization.

6.2 Quantifying the memory effect of speckles

As discussed in Sec. 1.1, the memory effect of speckles has been at the core of imaging techniques for a diverse range of applications, including seeing through tissue and around corners. Because of its wide applicability, understanding the range of illumination, viewing, and material conditions which result in high correlation between speckles is an active research area in optics.

There have been multiple attempts [Berkovits and Feng 1994; Fried 1982; Osnabrugge et al. 2017] to derive closed-form expressions for speckle correlation. The complexity of multiple scattering means that this is only possible under various assumptions, which limit

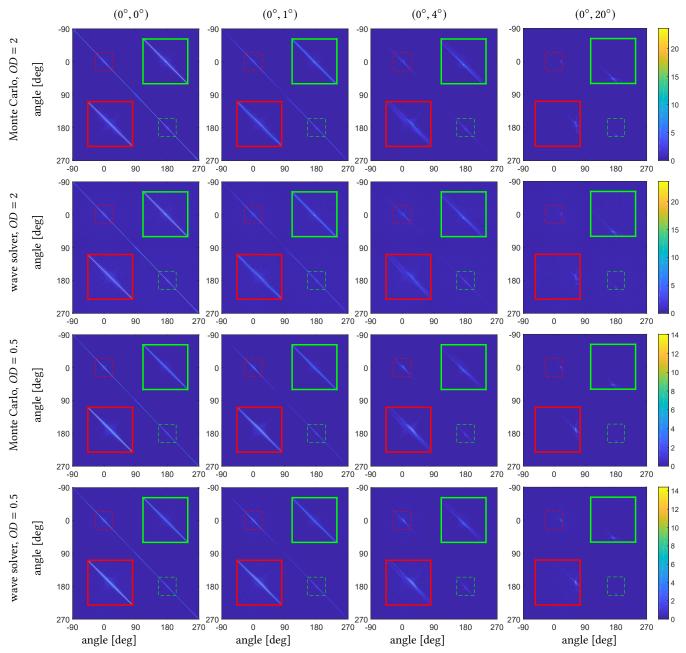


Fig. 8. Validating covariance rendering. Covariance matrices for illumination with two different plane waves, imaged at the far-field over viewing directions spanning 360°. Matrices rendered with Alg. 1 (rows 1 and 3) closely match those from a wave equation solver (rows 2 and 4). The covariance matrices demonstrate the memory effect for four different pairs of illumination angles and two different optical depths.

the approximation accuracy and the applicability of the resulting expressions. We state below a commonly-used result [Akkermans and Montambaux 2007; Feng et al. 1988] that is derived under a diffusion (that is, high-order scattering) assumption:

$$C(\theta) \approx \frac{(k\theta L)^2}{\sinh^2(k\theta L)},$$
 (39)

where θ is the angle between illumination and viewing directions, L is the material thickness, and $C(\theta)$ is the correlation between intensity images (rather than complex fields) $I_{\rm v}^{\rm i}$ and $I_{{\rm v}+\theta}^{{\rm i}+\theta}$. The correlation of Eq. (39) decays to zero exponentially fast as soon as $k\theta L>1$, hence the angular range at which the memory effect is valid is proportional to 1/(kL). We discuss later in this section how the Monte Carlo formulation can help understand this result.

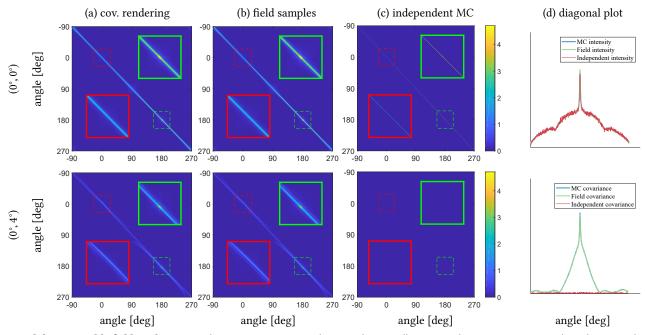


Fig. 9. **Validating speckle field rendering.** We show covariance matrices between the same illumination directions $i_1 = i_2 = 0^\circ$, where the matrix diagonal is intensity (top row), and different illumination directions $i_1 = 0^\circ$, $i_2 = 4^\circ$ (bottom row). The empirical covariance of speckle fields sampled using Alg. 2 in (b) closely matches the covariance rendered directly using Alg. 1 in (a). By contrast, speckle fields rendered with EMC [Sawicki et al. 2008] exhibit no spatial correlations (c), but can reproduce the correct intensity in the case of the top row, as seen from the plots of the matrix diagonals in (d).

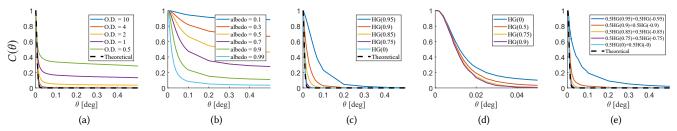


Fig. 10. Quantifying the memory effect. We compare analytical (dash) and numerical (solid) calculations of the correlation $C(\theta)$ as a function of angle θ . (a) Varying OD in a forward-scattering configuration for a material with isotropic phase function g=0. For high OD, the computed correlation agrees with diffusion theory. As OD decreases, the range of the memory effect increases. (b) Varying albedo in a back-scattering configuration with a fixed g and g an

The diffusion assumption used to derive Eq. (39) means that the formula applies only when the average number of scattering events on a path is large. However, empirical observations suggest that, in practice, the memory effect is valid through a much wider range. A few scenarios that have been observed to increase this range are (i) an average number of scattering events that is lower than the diffusive regime, (ii) absorption, (iii) forward scattering phase functions [Schott et al. 2015]. Forward scattering is particularly important in practice, as tissue is known to be highly forward scattering and is usually described by an Henyey-Greenstein (HG) phase function

with anisotropy parameter $g \in [0.85-0.95]$. Given the lack of analytic formulas and the practical importance of the problem, there have been multiple attempts to empirically measure the range of the memory effect of materials of interest in the lab [Mesradi et al. 2013; Schott et al. 2015; Yang et al. 2014].

Our Monte Carlo algorithm can compute the expected correlations directly, without the need for approximations or lab measurements. We note first that our Monte Carlo algorithm computes correlations of complex fields while Eq. (39) evaluates intensity correlations. However, field correlations can be easily converted to intensity correlations using $2|C_{\mathbf{i}_1,\mathbf{v}_2}^{\mathbf{i}_1,\mathbf{i}_2}|^2$. Additionally, correlations are

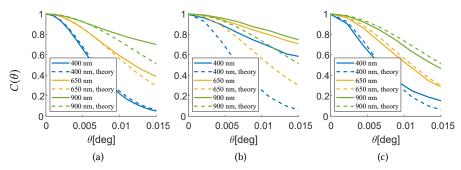


Fig. 11. Wavelength dependency of the memory effect. We compare analytical (dash) and numerical (solid) calculations of the correlation $C(\theta)$ as a function of angle θ , for varying illumination wavelengths. Wavelength-dependent scattering properties were computed using Mie theory. (a) Small particles of radius 10 nm, with corresponding OD = 20 for 400 nm, OD = 3 for 650 nm and OD = 1 for 900 nm. (b) Large particles of radius 10 μ m with average OD = 20. (c) Large particles as in (b), with average OD = 50.

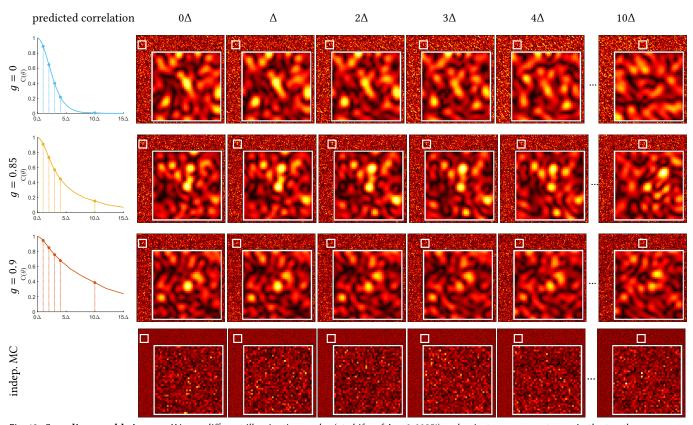


Fig. 12. **Sampling speckle images.** We use different illumination angles (at shifts of $\Delta=0.0025^{\circ}$) and anisotropy parameters g. In the top three rows, we show images rendered with Alg. 2, where we can observe the memory effect: For g=0, the correlation is lost at a shift of 3Δ , whereas for g=0.9, correlation remains even at shift of 10Δ . In the bottom row, we show images rendered with the EMC algorithm [Xu 2004], where each view is sampled independently.

computed as a function of simple parameters such as sample thickness, and material σ_t , σ_s and phase function. In Fig. 10, we show numerical simulations of the expected correlation as a function of angle θ . In Fig. 10a we use a forward scattering configuration, a sample of thickness L=1 mm at illumination wavelength $\lambda=500$ nm, $\sigma_a=0$, isotropic phase function g=0, and varying mean free path

(*MFP*) values. For a high optical depth, the correlation computed by our algorithm agrees with the theoretical formula of Eq. (39), and as the optical depth decreases, the range of the memory effect increases. In Fig. 10b we simulate a backscattering configuration for fixed g=0, *MFP* = 0.1 mm, $\sigma_t=1/MFP$, and varying albedo

 σ_s/σ_t . As expected, the memory effect is stronger as absorption increases (albedo decreases), as absorption attenuates the contribution of longer paths to the measured fields.

In Fig. 10c we keep the thickness and mean free path fixed to L = 1 mm, MFP = 0.1 mm, and vary the anisotropy parameter q of the phase function. In agreement with previous empirical observations [Schott et al. 2015], increasing *g* increases the *transport* mean free path, and thus the memory effect range expands. In Fig. 10d we investigate another common analytical approximation, the socalled similarity theory [Wyman et al. 1989; Zhao et al. 2014], which states that scattering coefficients and phase functions satisfying $\sigma_s^*(1-g^*) = \sigma_s(1-g)$ should produce similar scattering measurements. Using $L=1\,\mathrm{mm},\ \sigma_a=0,\ \mathrm{we\ set}\ \mathrm{at}\ g=0\ \mathrm{a}\ \mathrm{mean\ free}$ path of MFP = $250 \,\mu\text{m}$ (leading to OD = 4), and then vary g and $\sigma_s = \sigma_t = 1/MFP$ while maintaining the similarity relation. The graphs in Fig. 10d show that similarity theory is reasonably accurate, though low q values have a somewhat heavier tail. Finally in Fig. 10e we simulate a mixture of two HG phase functions whose mean cosine is always set to 0. We can see that the exact shape of the phase function influences the memory effect, and two phase functions with the same mean cosine can lead to very different decay graphs.

In Fig. 11 we study how the memory effect changes as a function of wavelength. In Fig. 11a we used particles of radius 10 nm, for which the cross-section varies with wavelength, and as a result the mean free path varies with wavelength as well. In Fig. 11(b-c) we used bigger particles of radius 10 μ m and set the density to achieve OD = 20 and OD = 50 respectively. As the phase function of such large particles is very forward scattering, the transport mean free path is much larger than the mean free path, and the thickness of the material is only two transport mean free paths for Fig. 11b and five transport mean free paths for Fig. 11c. We can see that the memory effect range varies with wavelength. When the transport mean free path is sufficently smaller than the target thickness, the graphs approach the theoretical prediction by the diffusion approximation.

Sampling speckle images. In Fig. 12 we use the sampling algorithm of Sec. 5.2 to sample speckle images as seen from a sensor at infinity over a viewing range of 0.1°, when the illumination direction is shifting (from 0° to 0.025° , at $\Delta = 0.0025^{\circ}$ intervals). As can be seen, these images reproduce the memory effect: For small changes in illumination angle the speckles appear as shifted versions of each other. When the illumination angle difference increases, the correlation decays. We show this simulation for a few anisotropy parameters q and as illustrated in Fig. 10c, when the anisotropy increases the memory effect can be observed over a wider angular range. In the last row of Fig. 12 we show simulations using the electric field Monte Carlo approach [Sawicki et al. 2008; Xu 2004], which updates different viewing and illumination directions independently. We observe that no joint speckle statistics are produced and the resulting images appear as independent noise.

Understanding the memory effect bounds. Before concluding this section, it is worth mentioning that our path integral formulation can provide an intuitive way to understand the memory effect range derived in Eq. (39). Consider two pairs of illumination and viewing directions i_1 , i_2 , v_1 , v_2 s.t. $i_1 - i_2 = v_1 - v_2 = \omega$, and consider a path starting at x_1 and ending at x_B . Dropping attenuation, the phase

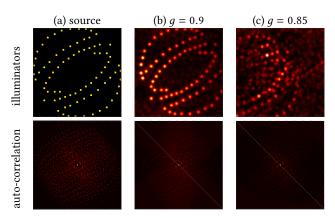


Fig. 13. Reproducing the seeing-through-scattering algorithm of Katz et al. [2014]. A set of illuminators with the arrangement at the top of (a) generates a semi-random speckle image, yet the auto-correlation of the speckle image is similar to the auto-correlation of the original illuminators and hence the illuminators can be recovered from the speckle image using phase retrieval algorithms. In (b,c) we show the auto-correlation and the corresponding reconstruction for different material parameters simulated with our speckle renderer. The success of the algorithm depends on the validity of the memory effect in this angular range for each type of material.

contributed by this path to the correlation is

$$e^{ik((i_1-i_2)\mathbf{x}_1-(\mathbf{v}_1-\mathbf{v}_2)\mathbf{x}_B)} = e^{ik\omega(\mathbf{x}_1-\mathbf{x}_B)}.$$
 (40)

If this complex number can have highly varying phases, than summing over multiple random paths averages to zero. The different paths interfere constructively only if the phase difference is negligible, roughly when $k|\omega||\mathbf{x}_1 - \mathbf{x}_n| < 1$. Intuitively, the average distance between an entrance point and an exit point on the target scales with the target depth, and it is reasonable to expect that $E[|\mathbf{x}_1 - \mathbf{x}_n|]$ is proportional to L. This implies that the memory effect holds when $k|\omega|L < 1$, in agreement with Eq. (39).

6.3 Seeing-through-scattering application

To demonstrate an application of speckle correlations, we reproduced the algorithm of Katz et al. [2014]. This algorithm attempts to recover a set of incoherent light sources located behind a scattering layer. Remarkably, due to the memory effect, the auto-correlation of the speckle image should be equivalent to the auto-correlation of light sources positions. Thus, given the seemingly random speckle image, one can recover the position of light sources behind it by applying an iterative phase retrieval algorithm [Fienup 1982]. In Fig. 13 we show the result of this reconstruction applied on speckle images rendered with Alg. 2. We use two of the materials in Fig. 10c, with anisotropy parameters q = 0.85, q = 0.9. The hidden source is placed over an angular range of $0.0125^{\circ} = 5\Delta$. As evaluated in Figs. 10c and 12, for this angular range the correlation for q = 0.9 is high, but for q = 0.85 we are already outside the memory effect range. Indeed the q = 0.9 speckle auto-correlation at the bottom of Fig. 13b is almost equivalent to the source auto-correlation (Fig. 13a[bottom]), while the auto-correlation of speckles rendered with q = 0.85 is darker due to the lower correlation (Fig. 13c[bottom]). As a result, phase retrieval with the g=0.9 speckles provides a good reconstruction of the original illuminator arrangement (Fig. 13b[top]). For g=0.85 (Fig. 13c[top]) only a cropped version of the illuminator pattern is recovered (along with background noise), as within this subset of illuminators the angular differences are smaller and the correlation is stronger. Experiments of this kind can be used to evaluate the applicability of the imaging technique of Katz et al. [2014] under different conditions, and to select optimal values for various parameters involved in an optical implementation of the technique.

7 SINGLE-SCATTERING APPROXIMATION

Before we conclude, we report an interesting property of speckle covariance, which can be used to accelerate its estimation under certain illumination and imaging conditions.

When simulating covariance using Monte Carlo rendering, we can separate contributions from paths of different numbers of bounces B. For example, in Fig. 14, we show simulations for a cube volume $\mathcal V$ of dimensions $100\lambda\times 100\lambda\times 100\lambda$, and with OD=5, resulting in strong multiple scattering. We simulate the covariance for multiple pairs of illumination and imaging sets satisfying $i_1-v_1-(i_2-v_2)=\omega$, for some target 3D vector ω . In each simulation, we decompose the rendered speckle covariance into two components, one accounting for contributions from paths that scattered once (B=1), and another accounting for paths that scattered two or more times $(B\geq 2)$. Within each rendered covariance matrix, the bottom left corner corresponds to rendering intensity.

We observe that, for the intensity case, the multiple-scattering component is dominant. By contrast, for cases where the difference between the two illumination or the two viewing directions is more than some small amount, the multiple-scattering component becomes negligible. This happens because, as the angle difference becomes large enough to bring us outside the range of the memory effect, multiply-scattered paths have complex contributions with randomly-varying phase, and therefore average to zero.

We conclude that, when the imaging and illumination conditions are such that we are outside the memory effect range, speckle covariance can be computed using only single scattering. Namely, from a short derivation we can obtain the formula:

$$C_{v_1,v_2}^{i_1,i_2} = s(i_1 \cdot v_1)s(i_2 \cdot v_2)^* \int_{\mathcal{V}} \sigma_s(\mathbf{x})e^{ik((i_1-v_1)-(i_2-v_2))\cdot\mathbf{x}} \eta(\mathbf{x}) d\mathbf{x}, \quad (41)$$

where $\eta(\mathbf{x}) = \tau(\mathbf{x}, i_1)\tau(\mathbf{x}, v_1)^*\tau(\mathbf{x}, i_2)^*\tau(\mathbf{x}, v_2)$. This integral can be evaluated much more efficiently than that of Eq. (33), without the need for expensive path tracing algorithms.

The above discussion indicates that, whenever we are outside the memory effect range, we can accelerate the computation of speckle covariance by using the single-scattering approximation, without significant loss in accuracy. This is analogous to the use of the single-scattering approximation for accelerating intensity rendering [Sun et al. 2005; Walter et al. 2009], with an important difference: In the case of intensity the single-scattering approximation is valid only for very optically-thin volumes [Narasimhan et al. 2006]. By contrast, in the case of covariance, the approximation can be accurate even for optically thick materials, given appropriate illumination and viewing conditions, making it more broadly applicable.

8 DISCUSSION

We presented a path-integral formulation for the covariance of speckle fields generated by the interaction of coherent light with scattering volumes. Using this formulation, we derived two Monte Carlo rendering algorithms, one for directly estimating covariance, and another for directly generating speckle patterns. As we demonstrated in Sec. 6, our algorithms provide a unique combination of physical accuracy (closely matching solutions of the wave equation, reproducing known physical phenomena such as memory effect and coherent backscattering), computational efficiency (outperforming wave equation solvers by orders of magnitude), and parsimony (using only bulk macroscopic parameters of a volume, instead of requiring knowledge of its microscopic structure). We conclude the paper with a discussion of limitations and possible future directions.

Both of our Monte Carlo rendering algorithms share strong similarities with Monte Carlo volume rendering algorithms for intensity, and in particular volumetric path tracing. This facilitates integration into popular physically-accurate rendering engines [Jakob 2010; Pharr et al. 2016], and reusing existing technology for efficient implementations. Likewise, existing results about the computational and theoretical properties of volumetric path tracing should be straightforward to apply to our setting. We should highlight, however, an important difference with the intensity case: The endpoints of each sampled path are connected to multiple light sources and sensors. This makes it challenging to importance-sample the first vertex and direction when starting to trace a path, which in turn can result in high variance in the resulting renderings. For example, in a material with a very forward-scattering phase function, we can select the first vertex and direction by importance-sampling at most one of the sensor or source connections; consequently, most other connections will produce paths of close-to-zero contributions.

A way to ameliorate this issue can be to apply to the speckle rendering setting variance reduction techniques that have been successful for intensity rendering. For example, multiple importance sampling [Veach and Guibas 1995b] can be used to reduce variance when the start of a path is sampled based on only one out of many source and sensor connections. More generally, it will be important to investigate path sampling techniques algorithms better suited to speckle rendering. To this end, we hope that our path-integral formulation for speckle covariance can provide the theoretical foundation for these investigations, analogous to how the path-integral formulation for intensity spurred the invention of algorithms such as bidirectional path tracing [Veach and Guibas 1995a] and Metropolis light transport [Veach and Guibas 1997]. In particular, we observe that our formulation is reciprocal, and therefore lends itself to the development of bidirectional, or even multidirectional for multiple sources and sensors, path sampling algorithms.

Despite their ability to accurately model first and second-order statistics of volumetric speckle, our theory and algorithms currently do not take into account wave effects that are likely to be important in real-world imaging situations. Namely, our path integral formulation ignores refraction and reflection events at the scattering volume's interface. While perfectly smooth interfaces can be incorporated into our formulation exactly analogously to standard Monte Carlo volume rendering, handling rough interfaces is

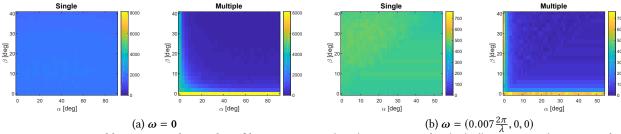


Fig. 14. **Decomposing speckle covariance by number of bounces.** We simulate the covariance of multiple illumination and viewing configurations satisfying $i_1 - v_1 - (i_2 - v_2) = \omega$, and plot contribution of single and multiple-scattering paths to the covariance as a function of two angles, $\alpha = \min(\angle(i_1, i_2), \angle(v_1, v_2))$, $\beta = \min(\angle(i_1, v_1), \angle(i_2, v_2))$. The multiple scattering term reduces to zero as soon as one of the above angular differences increases.

non-trivial due to the need to account for diffraction and speckle from surface events [Bergmann et al. 2016; Cuypers et al. 2012; Werner et al. 2017; Yan et al. 2018]. Additionally, our theory and algorithms cannot operate in scattering volumes that do not satisfy the assumptions of classical radiative transfer (Sec. 3). There is, by now, considerable literature extending Monte Carlo volume rendering to, for example, discrete random media with large and dense scatterers [Moon et al. 2007]; anisotropic media where scattering is not rotation-invariant [Jakob et al. 2010]; and non-exponential media where the locations of scatterers are not independent of each other [Bitterli et al. 2018]. All of these works focus exclusively on intensity rendering, and it would be interesting to investigate how to extend their techniques to the speckle rendering setting. Finally, our framework assumes unpolarized and fully-coherent illumination. Incorporating polarization effects into our framework can be done using existing techniques for both intensity [Jarabo and Arellano 2018] and speckle [Sawicki et al. 2008; Xu 2004] rendering. Accounting for partial coherence has received some attention [Pierrat et al. 2005; Shen et al. 2017], and these works can provide inspiration for developing similar extensions of our framework.

In this paper we chose to focus on spatial speckle correlations. However, this is only one of many other classes of second-order speckle statistics. For example, when the same volume is imaged under coherent illumination of different wavelengths, the resulting speckle patterns exhibit cross-wavelength correlations. This is a correlation property that so far remains relatively unexplored. As another example, speckle patterns in images of the same volume captured over time exhibit temporal correlations, due to moving scatterers in the volume [Dougherty et al. 1994]. These temporal correlations have found widespread use in techniques such as dynamic light scattering [Berne and Pecora 2000] and diffusing wave spectroscopy [Pine et al. 1988]. Even outside of volumes, speckle patterns due to surface microgeometry exhibit correlation properties analogous to the memory effect in volumes [Goodman 2007], with applications in non-line-of-sight imaging [Katz et al. 2012]. We hope that our results will motivate the development of analogous theoretical and simulation tools for these types of correlations. We expect that such tools can help broaden our understanding of speckle correlation effects, and extend their applicability to imaging applications. For instance, our rendering algorithms allow us to study the spatial memory effect and related applications in cases

where common assumptions (diffusion, Fokker-Planck limit [Osnabrugge et al. 2017]) do not hold. Likewise, rendering algorithms for temporal correlations can allow extending related applications to cases where the common assumption of Brownian motion of scatterers is invalid [Duncan and Kirkpatrick 2008]. Finally, the ability to render physically-accurate speckles can facilitate incorporating machine learning techniques into related imaging applications, where the collection of training data has been a major burden.

Last but not least, the findings of Sec. 7 suggest that measuring and rendering speckle covariance holds promise for *inverse rendering* applications. The fact that speckle covariance measurements are dominated by single scattering for a much larger class of materials than intensity measurements can potentially drastically simplify the volumetric inverse rendering problem, e.g., by potentially allowing us to replace the complex differentiable rendering of Gkioulekas et al. [2016; 2013] with simple analytic algorithms of Narasimhan et al. [2006]. In addition to simplifying computation, it will be interesting to examine whether speckle covariance measurements can be used to relax previously reported ambiguities between scattering parameters [Wyman et al. 1989; Zhao et al. 2014].

ACKNOWLEDGMENTS

We thank Prof. Amir Rosenthal for helpful discussions. This work was supported by European Research Council Horizon 2020 635537, ISF 1046-14, Ollendorff Minerva Center of the Technion, DARPA REVEAL HR0011-16-C-0028 and NSF Expeditions CCF-1730147.

REFERENCES

Eric Akkermans and Gilles Montambaux. 2007. Mesoscopic Physics of Electrons and Photons. Cambridge University Press.

M. Batarseh, S. Sukhov, Z. Shen, H. Gemar, R. Rezvani, and A. Dogariu. 2018. Passive sensing around the corner using spatial coherence. *Nature Communications*.

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*.

Stephan Bergmann, Mahsa Mohammadikaji, Stephan Irgenfried, Heinz Worn, Jürgen Beyerer, and Carsten Dachsbacher. 2016. A Phenomenological Approach to Integrating Gaussian Beam Properties and Speckle into a Physically-Based Renderer. In Vision, Modeling & Visualization.

Richard Berkovits and Shechao Feng. 1994. Correlations in coherent multiple scattering. Physics Reports (1994). https://doi.org/10.1016/0370-1573(94)90079-5

Bruce J Berne and Robert Pecora. 2000. Dynamic light scattering: with applications to chemistry, biology, and physics. Courier Corporation.

Jacopo Bertolotti, Elbert G. van Putten, Christian Blum, Ad Lagendijk, Willem L. Vos, and Allard P. Mosk. 2012. Non-invasive imaging through opaque scattering layers. *Nature* 491(7423), 232.

Benedikt Bitterli, Srinath Ravichandran, Thomas Müller, Magnus Wrenninge, Jan Novák, Steve Marschner, and Wojciech Jarosz. 2018. A radiative transfer framework for

- non-exponential media. In SIGGRAPH Asia. ACM, 225.
- David A Boas and Arjun G. Yodh. 1997. Spatially varying dynamical properties of turbid media probed with diffusing temporal light correlation. J. Opt. Soc. Am. A.
- Craig F. Bohren and Donald R. Huffman. 1983. Absorption and scattering of light by small particle. John Wiley & Sons.
- Tom Cuypers, Tom Haber, Philippe Bekaert, Se Baek Oh, and Ramesh Raskar. 2012. Reflectance Model for Diffraction. ACM Trans. Graph. 31, 5, Article 122, 11 pages. https://doi.org/10.1145/2231816.2231820
- Eugene d'Eon. 2018a. A Reciprocal Formulation of Non-Exponential Radiative Transfer. 2: Monte Carlo Estimation and Diffusion Approximation. arXiv preprint arXiv:1809.05881.
- Eugene d'Eon. 2018b. A reciprocal formulation of non-exponential radiative transfer with uncorrelated sources, detectors and boundaries. 1: Sketch and motivation. arXiv preprint arXiv:1803.03259.
- Ronald L. Dougherty, Bruce J. Ackerson, N.M. Reguigui, F. Dorri-Nowkoorani, and Ulf Nobbmann. 1994. Correlation transfer: Development and application. J. of Quantitative Spectroscopy and Radiative Transfer.
- Donald D. Duncan and Sean J. Kirkpatrick. 2008. Can laser speckle flowmetry be made a quantitative tool? J. Opt. Soc. Am. A 25, 8, 2088–2094. https://doi.org/10.1364/JOSAA.25.002088
- Turgut Durduran, Regine Choe, Wesley B. Baker, and Arjun G. Yodh. 2010. Diffuse optics for tissue monitoring and tomography. *Reports on Progress in Physics*.
- Philip Dutré, Kavita Bala, and Philippe Bekaert. 2006. Advanced global illumination. AK Peters, Ltd.
- Robert 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. *Physical review letters* 61, 7, 834.
- James R. Fienup. 1982. Phase retrieval algorithms: a comparison. Appl. Opt. 21, 15, 2758–2769.
- Isaac Freund. 1990. Looking through walls and around corners. *Physica: Statistical Mechanics and its App.*
- Isaac Freund and Danny Eliyahu. 1992. Surface correlations in multiple-scattering media. *Phys Rev A* (1992).
- Isaac Freund, Michael Rosenbluh, and Shechao. Feng. 1988. Memory Effects in Propagation of Optical Waves through Disordered Media. Phys. Rev. Lett. 61, 2328–2331.
 Issue 20. https://doi.org/10.1103/PhysRevLett.61.2328
- David L. Fried. 1982. Anisoplanatism in adaptive optics. J. Opt. Soc. Am. 72, 1, 52–61. https://doi.org/10.1364/JOSA.72.000052
- Jeppe Revall Frisvad, Niels Jørgen Christensen, and Henrik Wann Jensen. 2007. Computing the scattering properties of participating media using Lorenz-Mie theory. SIGGRAPH.
- Ioannis Gkioulekas, Anat Levin, and Todd Zickler. 2016. An Evaluation of Computational Imaging Techniques for Heterogeneous Inverse Scattering.
- I. Gkioulekas, S. Zhao, K. Bala, T. Zickler, and A. Levin. 2013. Inverse Volume Rendering with Material Dictionaries. ACM Transactions on Graphics (Proc. ACM SIGGRAPH Asia) (2013).
- W. I. Goldburg. 1999. Dynamic light scattering. American Journal of Physics (1999).
- Goodman. 2007. Speckle Phenomena in Optics: Theory and Applications. Roberts and Company Pub.
- Vadim Holodovski, Yoav Y. Schechner, Anat Levin, Aviad Levis, and Amit Aides. 2016. In-situ multi-view multi-scattering stochastic tomography. In ICCP.
- Y.A. Ilyushin. 2012. Coherent backscattering enhancement in highly anisotropically scattering media: Numerical solution. Journal of Quantitative Spectroscopy and Radiative Transfer (2012).
- Akira Ishimaru. 1999. Wave propagation and scattering in random media. Vol. 12. John Wiley & Sons.
- P. Jacquot and J. M. 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. 2010. Mitsuba renderer. http://www.mitsuba-renderer.org.
- Wenzel Jakob, Adam Arbree, Jonathan T Moon, Kavita Bala, and Steve Marschner. 2010. A radiative transfer framework for rendering materials with anisotropic structure. In ACM Transactions on Graphics (TOG), Vol. 29. ACM, 53.
- M. L. Jakobsen, H. T. Yura, and S. 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 Transactions on Graphics 37, 4 (2018).
 Adrian Jarabo and Victor Arellano. 2018. Bidirectional rendering of vector light trans-
- port. In Computer Graphics Forum, Vol. 37. Wiley Online Library, 96–105.

 O. Katz, P. Heidmann, M. Fink, and S. Gigan, 2014. Non-invasive single-shot imaging
- O. Katz, P. Heidmann, M. Fink, and S. Gigan. 2014. Non-invasive single-shot imaging through scattering layers and around corners via speckle correlation. *Nat. Photonics* (2014).
- O. Katz, E. Small, and Y. Silberberg. 2012. Looking around corners and through thin turbid layers in real time with scattered incoherent light. *Nature* (2012).

- Guillermo H. Kaufmann. 2011. Advances in Speckle Metrology and Related Techniques.
 Wiley
- Peter Kutz, Ralf Habel, Yining Karl Li, and Jan Novák. 2017. Spectral and decomposition tracking for rendering heterogeneous volumes. ACM Transactions on Graphics (TOG) 36, 4 (2017), 111.
- Aviad Levis, Yoav Y. Schechner, Amit Aides, and Anthony B. Davis. 2015. Airborne three-dimensional cloud tomography. In *ICCV*.
- J. H. Li and A. Z. Genack. 1994. Correlation in laser speckle. Phys. Rev. E 49 (May 1994), 4530–4533. Issue 5. https://doi.org/10.1103/PhysRevE.49.4530
- Qiang Lu, Xiaosong Gan, Min Gu, and Qingming Luo. 2004. Monte Carlo modeling of optical coherence tomography imaging through turbid media. Applied optics 43, 8 (2004), 1628–1637.
- Johannes Meng, Marios Papas, Ralf Habel, Carsten Dachsbacher, Steve Marschner, Markus H Gross, and Wojciech Jarosz. 2015. Multi-scale modeling and rendering of granular materials. ACM Trans. Graph. 34, 4 (2015), 49–1.
- M. Mesradi, A. Genoux, V. Cuplov, D. Abi Haidar, S. Jan, I. Buvat, and F. Pain. 2013. Experimental and analytical comparative study of optical coefficient of fresh and frozen rat tissues. *Journal of Biomedical Optics* (Nov. 2013). https://doi.org/10.1117/ 1.JBO.18.11.117010.
- M.I. Mishchenko, L.D. Travis, and A.A. Lacis. 2006. Multiple scattering of light by particles: radiative transfer and coherent backscattering. Cambridge Univ Pr.
- Jonathan T Moon, Bruce Walter, and Stephen R Marschner. 2007. Rendering discrete random media using precomputed scattering solutions. In Proceedings of the 18th Eurographics conference on Rendering Techniques. Eurographics Association, 231– 242.
- Allard P. Mosk, Ad Lagendijk, Geoffroy Lerosey, and Mathias Fink. 2013. Controlling waves in space and time for imaging and focusing in complex media. *Nat. Photonics* (2013).
- Thomas Müller, Marios Papas, Markus H Gross, Wojciech Jarosz, and Jan Novák. 2016. Efficient rendering of heterogeneous polydisperse granular media. ACM Trans. Graph. 35, 6 (2016), 168–1.
- S.G. Narasimhan, M. Gupta, C. Donner, R. Ramamoorthi, S.K. Nayar, and H.W. Jensen. 2006. Acquiring scattering properties of participating media by dilution. ACM Trans. Graph. 25, 3 (2006).
- Micha Nixon, Ori Katz, Eran Small, Yaron Bromberg, Asher A. Friesem, Yaron Silberberg, and Nir Davidson. 2013. Real-time wavefront shaping through scattering media by all-optical feedback. Nat. 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 4, 8 (2017), 886–892.
- Yingtian Pan, Reginald Birngruber, Jürgen Rosperich, and Ralf Engelhardt. 1995. Low-coherence optical tomography in turbid tissue: theoretical analysis. *Applied optics* 34, 28 (1995), 6564–6574.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. Physically based rendering: From theory to implementation. Morgan Kaufmann.
- Romain Pierrat, Jean-Jacques Greffet, Rémi Carminati, and Rachid Elaloufi. 2005. Spatial coherence in strongly scattering media. J. Opt. Soc. Am. A 22, 11 (Nov 2005), 2329– 2337.
- DJ Pine, DA Weitz, PM Chaikin, and E Herbolzheimer. 1988. Diffusing wave spectroscopy. *Physical review letters* 60, 12 (1988), 1134.
- 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. Opt. Express 16, 8 (Apr 2008), 5728–5738.
- Joseph M Schmitt and A Knüttel. 1997. Model of optical coherence tomography of heterogeneous tissue. JOSA A 14, 6 (1997), 1231–1242.
- Schott, Bertolotti, Léger, Bourdieu, and Gigan. 2015. Characterization of the angular memory effect of scattered light in biological tissues. *Opt. Express* (2015).
- Zhean Shen, Sergey Sukhov, and Aristide Dogariu. 2017. Monte Carlo method to model optical coherence propagation in random media. J. Opt. Soc. Am. A 34, 12 (Dec 2017), 2189–2193. https://doi.org/10.1364/JOSAA.34.002189
- Brandon M. Smith, Pratham Desai, Vishal Agarwal, and Mohit Gupta. 2017. CoLux: Multi-object 3D Micro-motion Analysis Using Speckle Imaging. ACM Trans. Graph. (2017).
- J. Stam. 1999. Diffraction shaders. In SIGGRAPH.
- Bo Sun, Ravi Ramamoorthi, Srinivasa G Narasimhan, and Shree K Nayar. 2005. A practical analytic single scattering model for real time rendering. In ACM Transactions on Graphics (TOG), Vol. 24. ACM, 1040–1049.
- F. Sur, B. Blaysat, and M. Grediac. 2018. Rendering deformed speckle images with a Boolean model. *Journal of Mathematical Imaging and Vision* (2018).
- B. Thierry, X. Antoine, C. Chniti, and H. Alzubaidi. 2015. µ-diff: An open-source Matlab toolbox for computing multiple scattering problems by disks. Computer Physics Communications 192 (2015), 348 362. https://doi.org/10.1016/j.cpc.2015.03.013

- B. E. Treeby and B. T. Cox. 2010. k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave-fields,. J. Biomed. Opt. (2010).
- V. Twersky. 1964. On propagation in random media of discrete scatterers. Am. Math. Sot. Symp. Stochastic Processes in Mathematical Physics and Engineering, Vol. 16, p. 84
- Veach. 1997. Robust Monte Carlo methods for light transport simulation. Ph.D. Dissertation. PhD thesis, Stanford University.
- Eric Veach and Leonidas Guibas. 1995a. Bidirectional estimators for light transport. In Photorealistic Rendering Techniques. Springer, 145-167.
- Eric Veach and Leonidas J Guibas. 1995b. Optimally combining sampling techniques for Monte Carlo rendering. In Proceedings of the 22nd annual conference on Computer graphics and interactive techniques. ACM, 419-428
- Eric Veach and Leonidas J Guibas. 1997. Metropolis light transport. In Proceedings of the 24th annual conference on Computer graphics and interactive techniques. ACM Press/Addison-Wesley Publishing Co., 65-76.
- Ivo M. Vellekoop and Christof M. Aegerter. 2010. Scattered light fluorescence microscopy: imaging through turbid layers. Opt. Lett. 35, 8 (Apr 2010), 1245-1247. https://doi.org/10.1364/OL.35.001245
- Bruce Walter, Shuang Zhao, Nicolas Holzschuch, and Kavita Bala. 2009. Single scattering in refractive media with triangle mesh boundaries. In ACM Transactions on Graphics (TOG), Vol. 28, ACM, 92,
- Sebastian Werner, Zdravko Velinov, Wenzel Jakob, and Matthias B. Hullin. 2017. Scratch Iridescence: Wave-Optical Rendering of Diffractive Surface Structure. ACM SIG-GRAPH Asia (2017).
- Douglas R Wyman, Michael S Patterson, and Brian C Wilson. 1989. Similarity relations for anisotropic scattering in Monte Carlo simulations of deeply penetrating neutral particles. J. Comput. Phys. 81, 1 (1989), 137-150.
- Min Xu. 2004. Electric field Monte Carlo simulation of polarized light propagation in turbid media. Opt. Express 12, 26 (Dec 2004), 6530-6539.
- Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering specular microgeometry with wave optics. ACM Transactions on Graphics (TOG) 37, 4 (2018), 75,
- Xin Yang, Ye Pu, and Demetri Psaltis. 2014. Imaging blood cells through scattering biological tissue using speckle scanning microscopy. Opt. Express 22, 3 (Feb 2014), 3405-3413. https://doi.org/10.1364/OE.22.003405
- K. Yee. 1966. Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media. EEE Trans. on Antennas and Propagation (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 Trans. Graph. (2013)
- Shuang Zhao, Ravi Ramamoorthi, and Kavita Bala. 2014. High-order similarity relations in radiative transfer. ACM Transactions on Graphics (TOG) 33, 4 (2014), 104.

APPENDIX

Dealing with multiple particle types A.1

In the main paper, we assumed, for simplicity, particles of a single type. We now show that, with small adjustments, our formulation can be extended to mixtures of particle types. We index particles of each type with a subscript ι . For each particle type ι , we assume we are given a, possibly spatially varying, density $\varsigma_t(\mathbf{x})$. We also denote its scattering and absorption cross-sections by $c_{s,\iota}, c_{a,\iota}$, its radius by r_t , and its (normalized) scattering amplitude function by s_t .

We start by defining the bulk scattering parameters, resulting from the mixture of particle types. Then, we revisit some equations in the main text that need to be adjusted for the case of multiple particle types. We denote by $\bar{N}_t(\mathbf{x})$ the mean number of particles of type ι in the unit volume at **x**. For spherical particles, this equals

$$\bar{N}_{l}(\mathbf{x}) = \frac{\varsigma_{l}(\mathbf{x})}{4/3\pi r_{s}^{3}}.$$
(42)

To account for the fact that we use normalized amplitude and phase functions, we define the total scattering and absorption coefficients σ_s , σ_a as the expected energy scattered or absorbed from all particles in a unit volume, which becomes,

$$\sigma_s(\mathbf{x}) = \sum_{l} \bar{N}_l(\mathbf{x}) c_{s,l}, \quad \sigma_a(\mathbf{x}) = \sum_{l} \bar{N}_l(\mathbf{x}) c_{a,l} + \sigma_a^{\text{med}}, \quad (43)$$

where σ_a^{med} is the attenuation coefficient of the containing medium. The extinction coefficient σ_t is $\sigma_t = \sigma_s + \sigma_a$. We will need to consider the bulk outer product of amplitude functions, defined as:

$$S(\theta_1, \theta_2) = \sum_{t} \beta_t s_t(\theta_1) s_t(\theta_2)^* \quad \text{with } \beta_t = \frac{c_{s,t} \bar{N}_t}{\sum_{t} c_{s,t} \bar{N}_t}. \tag{44}$$

The phase function can be defined as a special case of the above, $\rho(\theta) = S(\theta, \theta)$. The reason why the product of scattering amplitude functions at two angles is defined via Eq. (44) and not as $s(\theta_1)s(\theta_2)$ is because, when two paths scatter at the same point, their scattering is also due to the same particle of a single type.

To see how $S(\theta_1, \theta_2)$ becomes relevant for speckle rendering, we first introduce the notation

$$\Upsilon((\boldsymbol{\omega}_{1}, \boldsymbol{\omega}_{2}) \to \mathbf{x}_{o} \to (\mathbf{x}_{1}, \mathbf{x}_{2})) =$$

$$\sum_{l} \beta_{l} v_{l}(\boldsymbol{\omega}_{1} \to \mathbf{x}_{o} \to \mathbf{x}_{1}) v_{l}(\boldsymbol{\omega}_{2} \to \mathbf{x}_{o} \to \mathbf{x}_{2})^{*},$$

$$(45)$$

where $v_{\iota}(\cdot)$ is the equivalent of the complex volumetric throughput $v(\cdot)$ of Eq. (24), but with the amplitude function of particle type ι :

$$v_{i}(\boldsymbol{\omega}_{i} \rightarrow \mathbf{x}_{o} \rightarrow \mathbf{x}_{i}) = \tau(\mathbf{x}_{o}, \mathbf{x}_{i})\xi(\mathbf{x}_{o} \rightarrow \mathbf{x}_{i})s_{i}(\boldsymbol{\omega}_{i} \cdot \widehat{\mathbf{x}_{o}\mathbf{x}_{i}}). \tag{46}$$

The term $\Upsilon((\omega_1, \omega_2) \to x_o \to (x_1, x_2))$ replaces all terms of the form $v(\omega_1 \to \mathbf{x}_o \to \mathbf{x}_1)v(\omega_2 \to \mathbf{x}_o \to \mathbf{x}_2)^*$ in the definition of the speckle covariance in Sec. 4.2. The resulting changes to the corresponding covariance rendering algorithm of Sec. 5.1 are summarized in Alg. 3. Effectively, $\Upsilon(\cdot)$ accounts for the fact that, when two paths scatter at the same location, they interact with the same particle of a single type, so the same s_t should apply to both paths.

Finally, we discuss how to adjust Alg. 2 for directly sampling a speckle image, for the case of multiple types of particles. The covariance in Eq. (37) for the multiple type case becomes:

$$C = \int_{\mathbb{P}} f(\vec{\mathbf{x}}^s) \sum_{l_1, l_2} \beta_{l_1} \beta_{l_2} \cdot \mathbf{a}_{l_1, l_2} (\vec{\mathbf{x}}^s) \cdot \mathbf{a}_{l_1, l_2}^* (\vec{\mathbf{x}}^s) \, d\vec{\mathbf{x}}^s, \tag{47}$$

and Eq. (38) becomes:

$$\mathbf{a}_{t,j}(\vec{\mathbf{x}}^s) = \left(v_t(\mathbf{x}_2 \to \mathbf{x}_1 \to \mathbf{i}_j)v_t(\mathbf{x}_{B-1} \to \mathbf{x}_B \to \mathbf{v}_j) + v_t(\mathbf{x}_{B-1} \to \mathbf{x}_B \to \mathbf{i}_j)v_t(\mathbf{x}_2 \to \mathbf{x}_1 \to \mathbf{v}_j)\right). \tag{48}$$

Therefore, for every vertex sample, we should also sample a particle type $\iota \sim \beta_{\iota}$. We summarize the changes in Alg. 4.

A.2 Integrals in path space

Our goal in this section is to derive expectations of path contributions and justify Eq. (22). To analyze the path contributions, we will divide the space of all path pairs $\vec{\mathbf{x}}_1$, $\vec{\mathbf{x}}_2$ from \mathbf{i}_1 , \mathbf{i}_2 to \mathbf{v}_1 , \mathbf{v}_2 into sets defined by the vertices they have in common. Let $\vec{\mathbf{x}}^s = {\mathbf{x}_1, \dots, \mathbf{x}_B}$ denote a set of vertices and $\vec{\mathbf{x}}^{s,P_1}, \vec{\mathbf{x}}^{s,P_2}$ two possibly different permutations of these vertices. We look at the set of all paths that share exactly the vertices in $\vec{\mathbf{x}}^s$ in orders P_1, P_2 orders:

$$L(\vec{\mathbf{x}}^{s,P_{1}}, \vec{\mathbf{x}}^{s,P_{2}}) = \begin{cases} (\vec{\mathbf{x}}^{1}, \vec{\mathbf{x}}^{2}) & | \vec{\mathbf{x}}^{1} = \{\mathbf{i}_{1} \rightarrow \dots \rightarrow \mathbf{x}_{P_{1}(1)} \rightarrow \dots \rightarrow \mathbf{x}_{P_{1}(B)} \rightarrow \dots \rightarrow \mathbf{v}_{1}\} \\ \vec{\mathbf{x}}^{2} = \{\mathbf{i}_{2} \rightarrow \dots \rightarrow \mathbf{x}_{P_{2}(1)} \rightarrow \dots \rightarrow \mathbf{x}_{P_{2}(B)} \rightarrow \dots \rightarrow \mathbf{v}_{2}\} \end{cases}, \quad (49)$$

where any occurrence of . . . in Eq. (49) can be replaced with any sequence of vertices, as long as they are different from each other.

ALGORITHM 3: Monte Carlo rendering of covariance $C_{\mathbf{v_1},\mathbf{v_2}}^{\mathbf{i_1},\mathbf{i_2}}$

```
⊳Initialize covariance estimate.
Set C = 0.
for iteration = 1 : N do
                                                                         Sample first vertex of subpath.
        Sample point \mathbf{x}_1 \sim q_o(\mathbf{x}_1) inside medium.
        Sample uniform direction \omega_1.
                                              ⊳Update covariance by single scattering path
        Update C += Vv(\mathbf{i}_1 \rightarrow \mathbf{x}_1)v(\mathbf{i}_2 \rightarrow \mathbf{x}_1)^* \Upsilon((\mathbf{i}_1, \mathbf{i}_1) \rightarrow \mathbf{x}_1 \rightarrow (\mathbf{v}_1, \mathbf{v}_2)).
                                                                          \triangleright Continue tracing the subpath.
                                                                    Sample second vertex of subpath.
        Sample distance d \sim \sigma_t(\mathbf{x}_1) |\tau(\mathbf{x}_1, \mathbf{x}_1 + d\omega_1)|^2.
        Set point \mathbf{x}_2 = \mathbf{x}_1 + d \cdot \boldsymbol{\omega}_1.
        Set b = 2.
        while x_b inside medium do
                                          ⊳Update covariance with next-event estimation.
                 Forward-only version:
                 Update C += \frac{V}{2}\Upsilon(\mathbf{x}_2 \rightarrow \mathbf{x}_1 \rightarrow (\mathbf{i}_1, \mathbf{i}_2))\Upsilon(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow (\mathbf{v}_1, \mathbf{v}_2)).
                 Or, forward and reversed version:
                 Update C += \frac{V}{2} \Big( \Upsilon(\mathbf{x}_2 \rightarrow \mathbf{x}_1 \rightarrow (\mathbf{i}_1, \mathbf{i}_2)) \Upsilon(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow (\mathbf{v}_1, \mathbf{v}_2)) \Big)
                                     +\Upsilon(\mathbf{x}_2 \rightarrow \mathbf{x}_1 \rightarrow (\mathbf{i}_1, \mathbf{v}_2))\Upsilon(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow (\mathbf{v}_1, \mathbf{i}_2))
                                     + \Upsilon(\mathbf{x}_2 {\rightarrow} \mathbf{x}_1 {\rightarrow} (\mathbf{v}_1, \mathbf{i}_2)) \Upsilon(\mathbf{x}_{b-1} {\rightarrow} \mathbf{x}_b {\rightarrow} (\mathbf{i}_1, \mathbf{v}_2))
                                     +\Upsilon(\mathbf{x}_2 \rightarrow \mathbf{x}_1 \rightarrow (\mathbf{v}_1, \mathbf{v}_2))\Upsilon(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow (\mathbf{i}_1, \mathbf{i}_2))
                                                                        ⊳Sample next vertex of subpath.
                 Sample direction \omega_b \sim \rho(\omega_{b-1} \cdot \omega_b)
                 Sample distance d \sim \sigma_t(\mathbf{x}_b) |\tau(\mathbf{x}_b, \mathbf{x}_b + d\omega_b)|^2.
                 Set point \mathbf{x}_{b+1} = \mathbf{x}_b + d \cdot \boldsymbol{\omega}_b.
                                                                                      \triangleright Account for absorption.
                 Sample scalar a \sim \text{Unif}[0, 1].
                 if a > \sigma_s(\mathbf{x}_{b+1})/\sigma_t(\mathbf{x}_{b+1}) 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.
```

With this definition, we can divide the space of all paths $\vec{\mathbf{x}}^1, \vec{\mathbf{x}}^2$ into disjoint sets. We now argue that the throughput contribution from each set $L(\vec{\mathbf{x}}^{s,P_1}, \vec{\mathbf{x}}^{s,P_2})$ averages to the volumetric throughput contribution of the direct paths $\vec{\mathbf{x}}^{s,P_1}, \vec{\mathbf{x}}^{s,P_2}$. To show this, we first use the notation $b_1^- = P_1^{-1}(b) - 1$, $b_1^+ = P_1^{-1}(b) + 1$ for the vertices before and after \mathbf{x}_b in the permuted sequence P_1 , and similarly $b_2^- = P_2^{-1}(b) - 1$, $b_2^+ = P_2^{-1}(b) + 1$ for P_2 .

CLAIM 1.

$$\int_{(\vec{\mathbf{x}}^{1}, \vec{\mathbf{x}}^{2}) \in L(\vec{\mathbf{x}}^{s, P_{1}, \vec{\mathbf{x}}^{s, P_{2}})} p(\vec{\mathbf{x}}^{1}, \vec{\mathbf{x}}^{2}) \mu(\vec{\mathbf{x}}^{1}) \mu(\vec{\mathbf{x}}^{2})^{*} = \prod_{b=0}^{B} \Upsilon_{b}(\vec{\mathbf{x}}^{1, P_{1}}, \vec{\mathbf{x}}^{1, P_{2}}) \prod_{b=1}^{B} \sigma_{s}(\mathbf{x}_{b}), \quad (50)$$

with
$$\Upsilon_b(\vec{\mathbf{x}}^{1,P_1}, \vec{\mathbf{x}}^{2,P_2}) = \Upsilon((\mathbf{x}_{b_1^-}, \mathbf{x}_{b_2^-}) \to \mathbf{x}_b \to (\mathbf{x}_{b_1^+}, \mathbf{x}_{b_2^+})).$$
 (51)

PROOF. Let us start by drawing an independent set of *B* vertices $\mathbf{x}_1, \dots, \mathbf{x}_B$. According to the target density, the probability for these

```
ALGORITHM 4: Monte Carlo rendering of J \times 1 field u for \{(i, v)_j\}_{j=1}^J.
```

```
⊳Initialize field estimate.
Set \mathbf{u} = \mathbf{0}.
for iteration = 1 : N do
        Sample random phase \zeta \sim \text{Unif}[0, 1].
        Set z = e^{2\pi i \zeta}
                                                                   Sample first vertex of subpath.
        Sample point \mathbf{x}_1 \sim q_o(\mathbf{x}_1).
        Sample particle type \iota \sim \beta_{\iota}.
                                                   ⊳Update field with single scattering path.
        Update \forall j, \mathbf{u}_j += z \cdot \sqrt{\frac{V}{2}} \cdot v_i(\mathbf{i}_j \rightarrow \mathbf{x}_1) v_i(\mathbf{i}_j \rightarrow \mathbf{x}_1 \rightarrow \mathbf{v}_j).
                                                                     ⊳Continue tracing the subpath.
                                                               Sample second vertex of subpath.
        Sample uniformly direction \omega_1.
        Sample distance d \sim \sigma_t(\mathbf{x}_1)|\tau(\mathbf{x}_1,\mathbf{x}_1+d\omega_1)|^2.
        Set point \mathbf{x}_2 = \mathbf{x}_1 + d \cdot \boldsymbol{\omega}_1.
        Set b = 2
        while xb inside medium do
                Sample random phase \zeta \sim \text{Unif}[0, 1].
                Set z = e^{2\pi i \zeta}.
                Sample particle type \iota \sim \beta_{\iota}.
                                                   ⊳Update field with next-event estimation
                Forward-only version:
                \text{Update } \forall j, \ \mathbf{u}_j \mathrel{+}= z \cdot \sqrt{\tfrac{V}{2}} v_i (\mathbf{x}_2 {\rightarrow} \mathbf{x}_1 {\rightarrow} \mathbf{i}_j) v_i (\mathbf{x}_{b-1} {\rightarrow} \mathbf{x}_k {\rightarrow} \mathbf{v}_j). 
                Or, forward and reversed version:
                 \forall j, \ \mathbf{u}_j \mathrel{+=} z \cdot \sqrt{\frac{V}{2}} \Big( v_i(\mathbf{x}_2 \rightarrow \mathbf{x}_1 \rightarrow \mathbf{i}_j) v_i(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{v}_j)
                                                        +v_{i}(\mathbf{x}_{2}\rightarrow\mathbf{x}_{1}\rightarrow\mathbf{v}_{i})v_{i}(\mathbf{x}_{b-1}\rightarrow\mathbf{x}_{b}\rightarrow\mathbf{i}_{i})
                                                                  Sample next vertex of subpath.
                Sample direction \omega_b \sim \rho(\omega_{b-1} \cdot \omega_b).
                Sample distance d \sim \sigma_t(\mathbf{x}_h) |\tau(\mathbf{x}_h, \mathbf{x}_h + d\omega_h)|^2.
               Set point \mathbf{x}_{b+1} = \mathbf{x}_b + d \cdot \boldsymbol{\omega}_b.
                                                                                \triangleright Account for absorption.
                Sample scalar a \sim \text{Unif}[0, 1].
                if a > \sigma_s(\mathbf{x}_{b+1})/\sigma_t(\mathbf{x}_{b+1}) then
                                                    ⊳ Terminate subpath at absorption event.
                       break
                end
               Set b = b + 1.
        end
                                                       ⊳Produce final field with correct mean.
Update \forall j, \mathbf{u}_j = m_{\mathbf{v}_j}^{\mathbf{i}_j} + \sqrt{\frac{1}{N}} \mathbf{u}_j
```

particles is the last term of Eq. (50), $\prod_{b=1}^{B} \sigma_s(\mathbf{x}_b)$. For each position \mathbf{x}_b , we draw a particle type $\iota(b) \sim \beta_\iota$. Given the type of all particles on the paths, we decompose the path probabilities.

Let L_b denote the set of all disjoint paths $(\vec{\mathbf{x}}^{1,b}, \vec{\mathbf{x}}^{2,b})$ from $\mathbf{x}_{P_1(b)}$ to $\mathbf{x}_{P_1(b+1)}$ and from $\mathbf{x}_{P_2(b)}$ to $\mathbf{x}_{P_2(b+1)}$, and let ω_1^b, ω_2^b denote the end direction of $\vec{\mathbf{x}}^{1,b}, \vec{\mathbf{x}}^{2,b}$ (i.e. the direction at which the last segment is entering $\mathbf{x}_{P_1(b+1)}$ or $\mathbf{x}_{P_2(b+1)}$). While the only constraint on $\vec{\mathbf{x}}^{1,b}, \vec{\mathbf{x}}^{2,b}$ is that they are disjoint, we will make the approximation that they are independent. Mishchenko et al. [2006] show that the error introduced by this approximation is $o(1/\bar{N})$ where \bar{N} is the

expected number of particles in the medium. Thus we can write

$$p(\vec{\mathbf{x}}^1, \vec{\mathbf{x}}^2) = p(\vec{\mathbf{x}}^{1,0})p(\vec{\mathbf{x}}^{2,0}) \prod_{b=1}^{B} p(\vec{\mathbf{x}}^{1,b} | \omega^{1,b-1})p(\vec{\mathbf{x}}^{2,b} | \omega_1^{b-1}).$$
 (52)

Using Eq. (18):

$$\int_{L_0} p(\vec{\mathbf{x}}^{1,0}, \vec{\mathbf{x}}^{2,0}) \mu(\vec{\mathbf{x}}^{1,0}) \mu(\vec{\mathbf{x}}^{2,0})^*
= \int_{\vec{\mathbf{x}}^{1,0}} p(\vec{\mathbf{x}}^{1,0}) \mu(\vec{\mathbf{x}}^{1,0}) \cdot \int_{\vec{\mathbf{x}}^{1,0}} p(\vec{\mathbf{x}}^{2,0}) \mu(\vec{\mathbf{x}}^{2,0})^*
= v(\mathbf{i}_1 \to \mathbf{x}_{P_1(1)}) \cdot v(\mathbf{i}_2 \to \mathbf{x}_{P_2(1)})^*.$$
(53)

As all paths in L_0 integrate to the direct path, we know that the end directions when entering $\mathbf{x}_{P_1(1)}, \mathbf{x}_{P_2(1)}$ are $\boldsymbol{\omega}_0^1 = \overline{\mathbf{i}_1} \mathbf{x}_{P_1(1)}, \boldsymbol{\omega}_0^2 = \overline{\mathbf{i}_2} \mathbf{x}_{P_1(1)}$ $\widehat{i_2x_{P_2(1)}}$. Given the first segment's end direction, we can apply Eq. (18) to the second segment, and likewise to all successive segments:

$$\begin{split} \int_{L_{b}} p(\vec{\mathbf{x}}^{1,b}, \vec{\mathbf{x}}^{2,b} | \boldsymbol{\omega}_{b-1}^{1}, \boldsymbol{\omega}_{b-1}^{2}) \mu(\vec{\mathbf{x}}^{1,b}) \mu(\vec{\mathbf{x}}^{2,b})^{*} &= \\ v_{l(P_{1}(b))}(\boldsymbol{\omega}_{b-1}^{1} \to \mathbf{x}_{P_{1}(b)} \to \mathbf{x}_{P_{1}(b+1)}) \\ v_{l(P_{2}(b))}(\boldsymbol{\omega}_{b-1}^{2} \to \mathbf{x}_{P_{2}(b)} \to \mathbf{x}_{P_{2}(b+1)})^{*}. \end{split} \tag{54}$$

Concatenating Eqs. (53) and (54) assuming given particle position and type, we obtain

$$\prod_{b=0}^{B} v_{l(P_{1}(b))}(\mathbf{x}_{P_{1}(b-1)} \rightarrow \mathbf{x}_{P_{1}(b)} \rightarrow \mathbf{x}_{P_{1}(b+1)})$$

$$v_{l(P_{2}(b))}(\mathbf{x}_{P_{2}(b-1)} \rightarrow \mathbf{x}_{P_{2}(b)} \rightarrow \mathbf{x}_{P_{2}(b+1)})^{*}. \quad (55)$$

Eq. (55) is sorted by the order of vertices in the two permutations. We now rewrite it using a generic vertex order $1, \ldots, B$ using the notation $b_i^- = P_i^{-1}(b) - 1$, $b_i^+ = P_i^{-1}(b) + 1$, resulting in

$$\prod_{b=0}^{B} v_{\iota(b)}(\mathbf{x}_{b_{1}^{-}} \rightarrow \mathbf{x}_{b} \rightarrow \mathbf{x}_{b_{1}^{+}}) v_{\iota(b)}(\mathbf{x}_{b_{2}^{-}} \rightarrow \mathbf{x}_{b} \rightarrow \mathbf{x}_{b_{2}^{+}})^{*}$$
(56)

If we now sum Eq. (56) for all possible particle type assignments, and consider also the vertex sampling probability, we get Eq. (50). \square

A.3 Path permutations

As mentioned in Sec. 4.2, Claim 1 significantly simplifies the integration space of the path integral formulation for covariance, by restricting it to path pairs with only shared vertices. However, Claim 1 does not imply that these vertices appear at the same order.

The reason most permutations can be ignored stems from the fact that the phase of the path throughput is proportional to the path length. For permutations that do not trace the vertices in the same order, the segment lengths are different (Fig. 5b), resulting in path contributions of different path phases, whose sum quickly reduces to zero. To demonstrate this empirically, in Fig. 15 we consider paths of three vertices x_1, x_2, x_3 , which is the smallest path length with non trivial permutations. We use $i_1 = i_2, v_1 = v_2$, and we fix the vertices $\mathbf{x}_1, \mathbf{x}_2$ while varying the third vertex \mathbf{x}_3 over a $20\lambda \times 20\lambda$ area. We evaluate the path throughput contributions $\mu(\vec{\mathbf{x}}^1)\mu(\vec{\mathbf{x}}^{2,P})^*$ for various permutations P of these vertices. When P is the identity permutation, $\vec{\mathbf{x}}^1$ and $\vec{\mathbf{x}}^{2,P}$ have the same length, thus $\mu(\vec{\mathbf{x}}^1)\mu(\vec{\mathbf{x}}^{2,P})^*$ is always a positive number. If *P* is the reversed permutation, leading to the path \vec{x}_2 = i \rightarrow x_3 \rightarrow x_2 \rightarrow x_1 \rightarrow v, we get a fixed phase only for the backscattering direction $v \sim -i$. For other directions, we see in the second row of Fig. 15b that perturbing the position of x_3 changes the phase, and thus the average of the pairwise path throughput over all positions of x_3 becomes zero,

$$\int \mu(\vec{\mathbf{x}}^1)\mu(\vec{\mathbf{x}}^{2,P})^* d\mathbf{x}_3 \approx 0.$$
 (57)

For all other permutations, there is no configuration of illumination and viewing directions that leads to a fixed phase. Then, as can be seen in Fig. 15c, varying the position of one of the vertices locally quickly changes the phase, thus the average of different path contributions over a local window is zero. There are some rare path selections leading to a locally stationary phase, as can be seen in Fig. 15d. However, the probability of selecting such paths is low, and therefore the contribution to the overall covariance is negligible.

As an additional experiment, in Fig. 16 we numerically evaluate the integral of all six permutations of three numbers when varying two of the scatterers positions within a 2D square area,

$$\iint \mu(\vec{\mathbf{x}}^1)\mu(\vec{\mathbf{x}}^{2,P})^* d\mathbf{x}_2 d\mathbf{x}_3, \tag{58}$$

We see that, except for the forward and reversed permutations, the throughput of other permutations integrates to a contribution about two orders of magnitude smaller than the forward contribution.

The correlation transfer equation

Monte Carlo rendering algorithms for intensity were historically derived in computer graphics from the volume rendering equation. This is the integral form of the radiative transfer equation (RTE), expressing radiance at one point as a function of radiance at other points in space. Monte Carlo rendering algorithms are, then, derived by recursively unrolling the volume rendering equation [Dutré et al. 2006; Novak et al. 2018]. We can follow the same direction to derive Monte Carlo speckle rendering algorithms, starting from a key result in the speckle correlation literature, namely the correlation transfer equation (CTE) [Ishimaru 1999; Twersky 1964].

The CTE provides an expression for the second-order moment of speckles at different spatial points under the same illumination direction. Using our notation, we can write this as: $E[u_{\mathbf{v}_1}^{\mathbf{i},O} \cdot u_{\mathbf{v}_2}^{\mathbf{i},O^*}] =$ $C_{\mathbf{V}_1,\mathbf{V}_2}^{\mathbf{i}} + m_{\mathbf{V}_1}^{\mathbf{i}} m_{\mathbf{V}_2}^{\mathbf{i}}$. The CTE then states that:

$$E[u_{\mathbf{v}_{1}}^{\mathbf{i},O} \cdot u_{\mathbf{v}_{2}}^{\mathbf{i},O^{*}}] = m_{\mathbf{v}_{1}}^{\mathbf{i}} m_{\mathbf{v}_{2}}^{\mathbf{i}} + \int_{\mathbf{x}} \sigma_{s}(\mathbf{x}) \int_{\omega} v(\omega \to \mathbf{x} \to \mathbf{v}_{1}) \cdot v(\omega \to \mathbf{x} \to \mathbf{v}_{2})^{*} L_{\mathbf{x},\omega}^{\mathbf{i}}, \quad (59)$$

where $L_{\mathbf{x},\boldsymbol{\omega}}^{\mathbf{i}}$ is the "light field" as used traditionally in computer graphics, namely the radiance arriving at point x from direction ω .

The important observation made by the CTE is that, to compute correlations between the fields at sensor points v_1 , v_2 , we need to integrate only radiance from other space points, without the need to memorize any other correlations. The radiance at other points is weighted by the complex volumetric throughput *v*, namely the probability and phase of making a "single scattering" step from \boldsymbol{x} to \mathbf{v}_1 and \mathbf{v}_2 . For the case $\mathbf{v}_1 = \mathbf{v}_2$, the covariance reduces to intensity, and indeed Eq. (59) reduces to the volume rendering equation, assuming zero emission inside the volume.

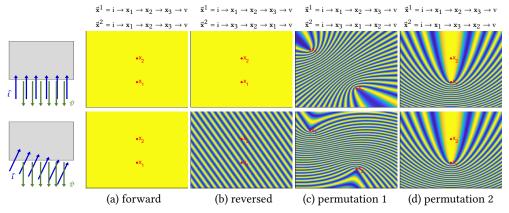


Fig. 15. **Phase of path permutations.** We show the phase of pairwise path throughput as a function of the position of one of the shared vertices, for different vertex permutations. For the forward permutation the phase is constant. For the reversed permutation a constant phase is achieved only at the backscattering direction. Other permutations result in spatially varying phase, thus the corresponding pairwise path throughputs cancel out after spatial integration.

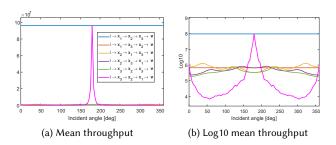


Fig. 16. Contributions from path permutations. We show the mean permutation throughput after integrating over spatial shifts of \mathbf{x}_2 , \mathbf{x}_3 . The mean contribution of the neglected permutations is two orders of magnitude smaller than the contributions of the forward and reversed permutations.

It is not hard to show that, for the case $i_1 = i_2$, a version of Alg. 3 considering only forward paths is equivalent to a recursive evaluation of the CTE. This version of the algorithm is derived by approximating the covariance as

$$C_{\mathbf{v}_{1},\mathbf{v}_{2}}^{\mathbf{i}_{1},\mathbf{i}_{2}} \approx \frac{1}{N} \sum_{\vec{\mathbf{x}}_{n}^{n} \sim q} \frac{c_{\vec{\mathbf{x}}^{1,n},\vec{\mathbf{x}}^{2,n}}}{q(\vec{\mathbf{x}}^{s,n})},$$
 (60)

rather than as

$$C_{\mathbf{v}_{1},\mathbf{v}_{2}}^{\mathbf{i}_{1},\mathbf{i}_{2}} \approx \frac{1}{N} \sum_{\vec{\mathbf{x}}^{n} \sim q} \frac{c_{\vec{\mathbf{x}}^{1},n},\vec{\mathbf{x}}^{2,n} + c_{\vec{\mathbf{x}}^{1},n},\vec{\mathbf{x}}^{2,n} + c_{\vec{\mathbf{x}}^{1},n},r,\vec{\mathbf{x}}^{2,n} + c_{\vec{\mathbf{x}}^{1},n},r,\vec{\mathbf{x}}^{2,n} + c_{\vec{\mathbf{x}}^{1},n},r,\vec{\mathbf{x}}^{2,n},r}{q(\vec{\mathbf{x}}^{s},r) + q(\vec{\mathbf{x}}^{s},r,n)}.$$
(61)

We highlight the forward-only version in Alg. 3 with blue font. A similar forward-only version can be derived for Alg. 4, also shown there in blue font.

The derivation of the CTE in, say, Mishchenko et al. [2006] follows from an expression of the solution to the wave equation as a sum of path contributions, analogous to what we presented in Sec. 4. Making the simplifying assumption that only forward path pairs need to be considered, the derivation reorganizes all the paths in the summation in a more compact recursive formula which is essentially the CTE, or the volume rendering equation in the $v_1 = v_2$ case. The fact that only forward paths are considered is an inherent assumption necessary for the compactness of the CTE, as it is equivalent to the fact that we only need to memorize the last vertex on a path and ignore the rest of its history. However, this compactness comes at a severe cost, namely the inability to explain coherent backscattering, which is an interference effect generated by the full path and not only by the last event. Due to this shortcoming, we chose to derive our Monte Carlo rendering algorithms directly from a path space formulation, and not from an integral equation.