Direct acquisition of volumetric scattering phase function using speckle correlations

Marina Alterman Evgeniia Saiko Anat Levin
Department of Electrical & Computer Engineering, Technion
Israel

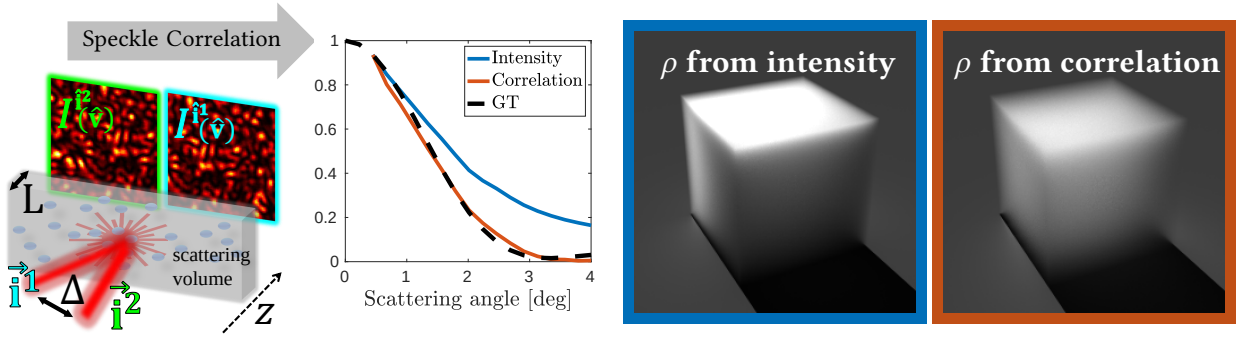

(a) Speckle acquisition (b) Estimated phase function (c) Rendered

Figure 1: Material acquisition: (a) Our setup captures the speckle patterns produced by a scattering volume under coherent laser illumination. We capture two such speckle images following a small tilt of the illumination beam. From the correlations between these two speckle images we can directly estimate the phase function of the material, bypassing costly optimization. (b) We use a validation material whose phase function is precisely known and compare it to the one estimated with our algorithm, showing good agreement. We compare that against a naive estimate, which assumes the material is thin enough so that light paths scatter only once, and uses the total intensity of light scattered at different angles as an estimate for the phase function. (c) To visually assess the difference between the naive intensity-based estimate of the phase function ρ and our accurate reconstruction, we used both phase functions to render a simple test scene of a scattering cube. The true phase function is narrower than the naive estimate, and the cube appears more translucent.

ABSTRACT

In material acquisition we want to infer the internal properties of materials from the way they scatter light. In particular, we are interested in measuring the phase function of the material, governing the amount of energy scattered towards different directions. This phase function has been shown to carry a lot of information about the type and size of particles dispersed in the medium, and is therefore essential for its characterization.

Previous approaches to this task have relied on computationally costly inverse rendering optimization. Alternatively, if the material can be made optically thin enough so that most light paths scatter only once, this optimization can be avoided and the phase function can be directly read from the profile of light scattering at different angles. However, in many realistic applications, it is not easy to slice or dilute the material so that it is thin enough for such a single scattering model to hold.

In this work we suggest a simple closed-form approach for acquiring material parameters from thick samples, avoiding costly optimization. Our approach is based on imaging the material of interest under coherent laser light and capturing speckle patterns. We show that memory-effect correlations between speckle patterns produced under nearby illumination directions provide a gating mechanism, allowing us to measure the singly scattered component of the light, even when observing thick samples where most light is scattered multiple times.

We have built an experimental prototype capable of measuring phase functions over a narrow angular cone. We test the accuracy of our approach using validation materials whose ground truth phase function is known; and we use it to capture a set of everyday materials.

CCS CONCEPTS

• **Computing methodologies** → **3D imaging**; **Computational photography**.

KEYWORDS

Material acquisition, scattering, speckles, memory effect.

ACM Reference Format:
Marina Alterman Evgeniia Saiko Anat Levin . 2022. Direct acquisition of volumetric scattering phase function using speckle correlations . In *SIGGRAPH Asia 2022 Conference Papers (SA '22 Conference Papers), December 6–9, 2022, Daegu, Republic of Korea.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3550469.3555379

SA '22 Conference Papers, December 6–9, 2022, Daegu, Republic of Korea

ACM ISBN 978-1-4503-9470-3/22/12…$15.00
https://doi.org/10.1145/3550469.3555379

1 INTRODUCTION

Scattering refers to the propagation of light in media composed of small discrete scatterers, usually particles of varying refractive properties. As an incident illumination propagates through the medium, it will interact with scatterers multiple times, and each such interaction will change its direction. 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.

Material acquisition refers to the task of recovering the intrinsic properties of such materials from their appearance, that is, given some measurement of the light they scatter. The task is of great value in many application settings. For example, in tissue imaging it can be used to detect tumors and classify them as malignant or non-malignant [Boas et al. 2001]; and in blood analysis, to recover diagnostically important parameters such as red and white blood counts [Berne and Pecora 2000; Durduran et al. 2010]; in material science and fabrication applications, to validate the fidelity of manufactured material samples [Sumin et al. 2019]; and in flow cytometry and particle sizing applications, to infer the chemical composition of nanodispersions [Pine et al. 1990].

Inverse radiative transport [Bal 2009] is studied in graphics, physics, chemistry, and biomedical sciences. Existing algorithms for acquiring scattering materials can be roughly classified into three categories. Methods based on the *diffusion* approximation consider optically thick media where high-order scattering is dominant. This allows for simpler inference and has been used for the acquisition of both homogeneous [Farrell et al. 1992; Jensen et al. 2001; Munoz et al. 2011; Papas et al. 2013] and heterogeneous materials [Boas et al. 2001; Cheong et al. 1990; Donner et al. 2008; Liu et al. 2020; Tuchin 2000; Wang et al. 2008]. However, it introduces ambiguities between different parameters [Wyman et al. 1989; Zhao et al. 2014].

At the other extreme, methods based on the *single scattering* approximation assume that the unknown medium is so optically thin that all photons scatter only once. This allows directly measuring scattering parameters of media such as smoke and thin or dilutable liquids [Fuchs et al. 2007; Gu et al. 2008; Hawkins et al. 2005; Narasimhan et al. 2006]. While the single scattering model makes parameter extraction direct and simple, not every material of interest is thin enough for this assumption to hold.

A third class of methods attempts to bridge the gap between the single scattering and the diffusion extremes, and seeks to use scattered paths of all lengths. They search for material parameters that explain input images of interest, when plugged into a Monte-Carlo simulator [Dutre et al. 2006; Novak et al. 2018]. These methods use various optimization schemes [Antyufeev 2000; Levis et al. 2015, 2017; Leyre et al. 2014; Prahl et al. 1993], where the modern approaches include differentiable rendering [Che et al. 2020; Gkioulekas et al. 2016, 2013; Khungurn et al. 2015; Nimier-David et al. 2020, 2019; Nindel et al. 2021; Velinov et al. 2018; Zhang et al. 2019] and machine learning [Che et al. 2020; Sde-Chen et al. 2021; Zheng et al. 2021]. While the approach can handle thicker materials, the need to simulate multiple scattering events for every candidate parameter set results in heavy computation.

In this work we relax this computational burden. For the first time, we propose a technique that can extract the phase function of a homogeneous material *in closed-form*. This direct extraction does not require thin, single scattering samples, and applies to materials of medium optical thickness. This is made possible by exploiting speckle statistics formed under coherent laser illumination.

Traditional computer graphics approaches have studied material acquisition under *incoherent illumination*, which is typically produced under natural illumination, LED, and in general when imaging resolution is not very high. Incoherent illumination usually results in smoothly-varying intensity distributions, often referred to as translucent appearance, and is well modeled by classical geometric optics models of light. In contrast, under *coherent imaging* conditions, obtained using laser light imaged in high magnification, the wave nature of light leads to *speckles*, which appear as pseudo-random high variations in the captured intensity images.

Despite their random appearance, speckles have strong statistical properties such as *the memory effect* (ME), implying that speckles formed under nearby illuminations are correlated shifted versions of each other, see Fig. 1(a). These fascinating properties have drawn a lot of attention in the literature [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] and have been the subject of multiple textbooks [Erf 1978; Goodman 2007; Jacquot and Fournier 2000; Kaufmann 2011]. The ME property has been exploited in multiple settings, in particular for seeing through scattering layers, bypassing very heavy aberration [Alterman et al. 2021; Bertolotti et al. 2012; Chang and Wetzstein 2018; Edrei and Scarcelli 2016a,b; Katz et al. 2014; Takasaki and Fleischer 2014]. In this work we explore a very different, novel application of the memory effect.

To this end, we start from a recent theoretical result by Bar *et al.* [2021], showing that under certain imaging conditions speckle correlations act as a gating mechanism, allowing an isolation of single-scattering light paths. By exploiting this theoretical result, we propose that as single-scattering light energy is proportional to the phase function, speckle correlations can produce a closed-form estimate of such phase functions. While the usage of single scattering models under previous incoherent imaging conditions requires that the material would indeed be thin enough so that most light paths scatter only once, with our coherent approach, we can isolate singly-scattered light from thicker materials samples, where most light paths scatter multiple times.

Our approach takes as input two speckle images of the same material target, captured following a small tilt of the illumination beam, and uses them to estimate local correlations. We demonstrate that our framework can recover phase functions, even from material samples thick enough to exhibit 8 scattering events in a typical light-path. We have built a proof-of-concept lab prototype, which can acquire phase functions within a narrow cone of 8°. We validate our approach against ground truth materials with known scattering parameters, and use it to measure a set of everyday materials.

Our approach applies mostly to viscous or solid materials, because it requires two sequential images of the same material, captured without any Brownian motion of internal scattering particles. We note, however, that solids have formed so far the hardest test-case for material acquisition, as most liquids can be easily diluted to match a single-scattering model [Narasimhan et al. 2006].

2 BACKGROUND

Notations: we use bold letters for vectors (e.g., 3D points $\mathbf{o}$). We use a top arrow to denote direction vectors (unit norm 3D vectors), e.g. $\vec{\mathbf{v}}$. We use circumflex to denote the first two coordinates of a unit norm 3D vector, namely its projection to the xy plane. e.g. $\hat{\mathbf{v}}$ denotes the xy coordinates of the 3D vector $\vec{\mathbf{v}}$. These are also known as the xy direction cosines.

2.1 Material representation and acquisition

Representation. Scattering materials are usually composed of small particles with different refractive properties. We use a statistical description of the bulk optical properties of the medium. In particular, the extinction coefficient σ_t describes the extinction cross section of the scattering particles per unit volume, which is proportional to the density of scattering particles in the material. It is the sum of the scattering and absorption coefficients $\sigma_t = \sigma_s + \sigma_a$, which model, respectively, the portion of energy that is scattered and absorbed per unit length along the path. The phase function of the material $\rho(\text{acos}(\vec{\mathbf{i}} \cdot \vec{\mathbf{v}}))$ describes the amount of energy scattered by a scatterer illuminated from direction $\vec{\mathbf{i}}$ toward direction $\vec{\mathbf{v}}$. ρ is a function of the scatterers' shape and refractive index. For spherical particles it can be computed using Mie theory formulas [Bohren and Huffman 1983; Frisvad et al. 2007; van de Hulst 1981]. It is often assumed that ρ is rotationally invariant in the sense that it depends only on the inner product $(\vec{\mathbf{i}} \cdot \vec{\mathbf{v}})$ of the illumination and viewing directions rather than on absolute directions. It is common to characterize the phase function by an anisotropy parameter $-1 \leq g \leq 1$, where $g = 0$ corresponds to scattering equally in all directions, and $g = 1$ to fully forward scattering. The mean free path (MFP) of the material is defined as the average distance that light travels in the volume between two successive scattering events, and can be shown to equal the inverse of the extinction coefficient, $MFP = 1/\sigma_t$. Given a scattering volume, it is common to express its geometric dimensions relative to the MFP. For example, a volume has optical depth $OD = 2$ if its thickness is $2 \cdot MFP$, meaning that light traveling through the volume undergoes on average two scattering events.

In this work we are interested in measuring the phase function ρ, and will not address $\sigma_t, \sigma_s, \sigma_a$. We model ρ as a general function and do not assume low-order parameterizations, such as a Henyey-Greenstein phase function [Henyey and Greenstein. 1941]

Phase function using single scattering models. A basic strategy for measuring phase functions [Narasimhan et al. 2006] is using optically thin samples where most light paths scatter only once. In this case, if we illuminate the material from direction $\vec{\mathbf{i}}$ and image from direction $\vec{\mathbf{v}}$, the amount of measured energy is equivalent to the phase function at angle $\text{acos}(\vec{\mathbf{i}} \cdot \vec{\mathbf{v}})$ (after accounting for foreshortening and other geometric factors). In the paraxial regime we can approximate this angle as the length of the displacement vector $\boldsymbol{\tau} = \hat{\mathbf{v}} - \hat{\mathbf{i}}$, and $\text{acos}(\vec{\mathbf{i}} \cdot \vec{\mathbf{v}}) \approx |\boldsymbol{\tau}|$, see Fig. 3(a). This model fails as material thickness increases and light scatters multiple times.

2.2 Speckle correlations

Consider a volume with scattering particles as in Fig. 1(a). We denote by $I^{\hat{\mathbf{i}}}(\hat{\mathbf{v}})$ the intensity scattered by a material illuminated from direction $\vec{\mathbf{i}}$ toward viewing direction $\vec{\mathbf{v}}$, when parameterized by their 2D projections $\hat{\mathbf{i}}$ and $\hat{\mathbf{v}}$. Assuming the same particle instantiation is illuminated from directions $\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2$, and considering two viewing directions $\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2$, we define the *speckle covariance* as:

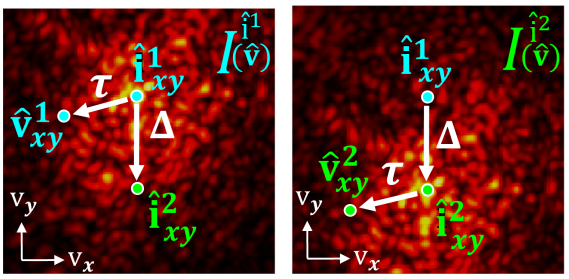

Figure 2: Notation: Speckle images obtained on the sensor plane, from two illumination directions $\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2$. Notice ME correlations between the two speckle patterns. We mark one pair of viewing directions $\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2$ satisfying the ME conditions of Eq. (2), and the corresponding displacements $\Delta, \boldsymbol{\tau}$.

$$\boldsymbol{C}^{\hat{\mathbf{i}}^1,\hat{\mathbf{i}}^2}_{\hat{\mathbf{v}}^1,\hat{\mathbf{v}}^2} \equiv E\left[I^{\hat{\mathbf{i}}^1}(\hat{\mathbf{v}}^1) \cdot I^{\hat{\mathbf{i}}^2}(\hat{\mathbf{v}}^2)\right] - E\left[I^{\hat{\mathbf{i}}^1}(\hat{\mathbf{v}}^1)\right] \cdot E\left[I^{\hat{\mathbf{i}}^2}(\hat{\mathbf{v}}^2)\right], \quad (1)$$

where the expectation is taken with respect to multiple realizations of random media with the same statistical properties (multiple samples from the same bulk material parameters).

The memory effect of speckles states that two speckle fields generated by nearby illumination directions are correlated shifted versions of each other. That is, for *small* displacements $\boldsymbol{\Delta} = \hat{\mathbf{i}}^2 - \hat{\mathbf{i}}^1$, we have $I^{\hat{\mathbf{i}}^1}(\hat{\mathbf{v}}) \approx I^{\hat{\mathbf{i}}^2}(\hat{\mathbf{v}} + \boldsymbol{\Delta})$, as illustrated in Fig. 2. This correlation decays as material thickness or illuminators' displacement increase.

The accurate way to calculate the covariance defined in Eq. (1) is by numerically solving wave equations [Thierry et al. 2015; Treeby and Cox. 2010; Yee 1966]. However, in this work we are interested in direct expressions relating them to material parameters.

Classical speckle theory developed by Twersky [1964] suggests that the covariance can be expressed as an analytic integral over the space of path-pairs p^1 from $\vec{\mathbf{i}}^1$ to $\vec{\mathbf{v}}^1$ and p^2 from $\vec{\mathbf{i}}^2$ to $\vec{\mathbf{v}}^2$, as illustrated in Fig. 3(b). As in classical ray tracing formulation, each path has a throughput, as the energy along the path is attenuated in proportion to its length, the extinction coefficient and scattering phase function. However, in a coherent setting the throughput is *a complex number*, whose phase is proportional to the path length. Due to the large variation of path lengths, most path pairs contribute complex numbers with rather random phases, and the average of such complex numbers is zero. Indeed, [Bar et al. 2019] prove analytically that many path pairs can be omitted. Based on this observation they can derive an *efficient* path sampling scheme, focusing only on path pairs that share their nodes, as illustrated in Fig. 3(c). The reduced path-integral is evaluated using Monte-Carlo ray tracing algorithms [Dutre et al. 2006; Novak et al. 2018].

In [Bar et al. 2021], it was further suggested that many of the longer path pairs still contribute random phases, and in fact, most correlation can be attributed to single scattering paths, of the form $p^1 = \vec{\mathbf{i}}^2 \to \mathbf{o} \to \vec{\mathbf{v}}^2$, $p^2 = \vec{\mathbf{i}}^2 \to \mathbf{o} \to \vec{\mathbf{v}}^2$ (Fig. 3(d)). This simplification is the ray analogous of the first Born approximation [Newton 2002]. Its main advantage is that unlike the full set of path pairs, which

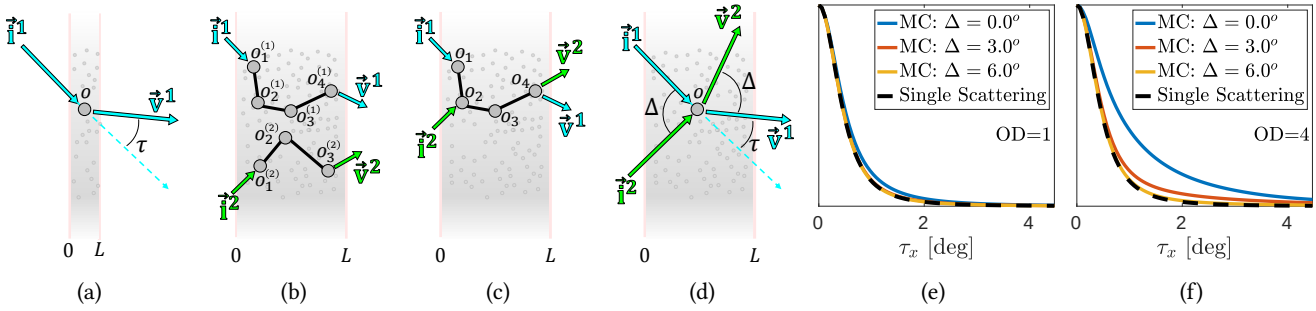

Figure 3: Single scattering correlations: (a) For a thin sample, most light paths do not scatter more than once. The intensity scattering at angle $\boldsymbol{\tau} = \hat{\mathbf{v}}^1 - \hat{\mathbf{i}}^1$ is proportional to the phase function $\rho(|\boldsymbol{\tau}|)$. (b-c) Attempting to explain speckle correlation using path pairs. (b) The complete description includes all pairs of paths from $\vec{\mathbf{i}}^1$ to $\vec{\mathbf{v}}^1$ and from $\vec{\mathbf{i}}^2$ to $\vec{\mathbf{v}}^2$. (c) Most such pairs are uncorrelated, and MC rendering algorithms only integrate path pairs sharing all intermediate nodes. (d) In practice, many of the path pairs in (b) are still uncorrelated and [Bar et al. 2021] suggests to approximate correlation by considering only path pairs that scatter at a single location. (e-f) Comparing between the exact correlation and the single scattering approximation of Claim 1. As we increase illuminator displacement Δ, single scattering becomes exact even at the thicker optical depth demonstrated in (f). For ease of visualization all curves were normalized to have the same maximum.

should be evaluated using Monte-Carlo sampling, single scattering paths can be evaluated in *closed-form*, to be summarized below.

While the result can be generalized, to simplify notation we consider illumination and viewing directions in the paraxial regime, forming a small angle with the optical z-axis. As reported in the literature, non-zero correlation is present mostly when the displacement between (the xy projections of) the illumination directions equals the displacement between viewing directions $\hat{\mathbf{i}}^2 - \hat{\mathbf{i}}^1 = \hat{\mathbf{v}}^2 - \hat{\mathbf{v}}^1$. We parameterize these illumination and viewing pairs using the 2D displacement vectors $\boldsymbol{\Delta}, \boldsymbol{\tau}$ (see Figs. 2 and 3(d)):

$$\boldsymbol{\Delta} \equiv \hat{\mathbf{i}}^2 - \hat{\mathbf{i}}^1 = \hat{\mathbf{v}}^2 - \hat{\mathbf{v}}^1, \quad \boldsymbol{\tau} \equiv \hat{\mathbf{v}}^1 - \hat{\mathbf{i}}^1 = \hat{\mathbf{v}}^2 - \hat{\mathbf{i}}^2, \tag{2}$$

encoding the angle between the two illumination vectors and the angle between the illumination and viewing directions, respectively. With this notation we can consider the phase function $\rho(|\boldsymbol{\tau}|)$ at angle $\boldsymbol{\tau}$. Finally, we denote the illumination wavelength by λ and the wave-number as $k = 2\pi/\lambda$. With this notation, [Bar et al. 2021] derive an analytic expression for the single scattering correlation.

Claim 1. *Consider a scattering slab of thickness L. A single scattering approximation to the speckle correlation (Eq. (1)) reduces to:*

$$\boldsymbol{C}(\boldsymbol{\tau}, \boldsymbol{\Delta}) = \left(\rho(|\boldsymbol{\tau}|) L \sigma_s e^{-\sigma_t L} \operatorname{sinc}\left(\frac{kL}{2}(\boldsymbol{\Delta} \cdot \boldsymbol{\tau})\right)\right)^2. \tag{3}$$

In Sec. 3 we propose that based on this formula, the phase function of the material can be read from correlation measurements.

Evaluating the single scattering approximation. In Fig. 3(e-f) we numerically evaluate the accuracy of this single-scattering approximation. We compare correlations produced by the closed-form expression in Claim 1 to those produced by a full Monte-Carlo simulator [Bar et al. 2019], which we treat as exact as it was exhaustively compared against an accurate wave optics solver [Thierry et al. 2015]. We simulate a forward-scattering phase function, which is common in many materials, such as tissue [Igarashi et al. 2007; Tuchin 2000].

We plot correlation as a function of the scattering angle $\boldsymbol{\tau} = (\tau_x, 0)$ (see definition in Eq. (2)), for a few values of illuminator displacements and two material thicknesses. For reasons that will be explained next, we consider illuminators displacement in the orthogonal direction, of the form $\boldsymbol{\Delta} = (0, \Delta_y)$. In the first case (Fig. 3(e)), $OD = 1$ and thus most light paths scatter only once. Indeed the closed-form single scattering expression agrees with the exact MC covariance for all Δ_y values. In the second case $OD = 4$ (Fig. 3(f)), and most light paths scatter multiple times. When the illuminators displacement is $\Delta_y = 0$, the two images we compare are identical, and it can be shown [Alterman et al. 2021] that correlation equals the total, multi-path intensity scattered to that direction. If the optical depth is larger than 1, long paths contribute a lot to the intensity, leading to a clear gap with the single scattering approximation. However, as the displacements Δ_y between the illuminators increase, the single scattering expression becomes an accurate approximation to the covariance, because long path pairs originated at two different directions $\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2$ exhibit larger phase variations, and hence they decorrelate and do not contribute to the covariance.

We acknowledge, however, that Eq. (3) also shows that correlation decays exponentially with $\sigma_t L$, which is essentially the optical depth of the material. As a result, the single scattering approximation is only useful in low to medium optical depths. For very large ODs, very limited correlation is actually present between the images, and cannot be reliably measured.

3 PHASE FUNCTION FROM SPECKLES

In this work we explore an important practical application of the theoretical result of [Bar et al. 2021]. For intuition, note that in a single scattering event, when light scatters from direction $\hat{\mathbf{i}}^1$ toward direction $\hat{\mathbf{v}}^1$ in direction $\boldsymbol{\tau} = \hat{\mathbf{v}}^1 - \hat{\mathbf{i}}^1$, the amount of scattered light is determined by the phase function at the angle $\rho(|\boldsymbol{\tau}|)$. If correlation can isolate single scattering paths, then as Eq. (3) states, correlation measured at different scattering angles is proportional to the phase function. As a result, *by detecting correlation between speckle patterns at different scattering angles we can directly read the phase function of the material.* This has a big practical advantage, because even in scattering volumes of medium thickness, one can extract the phase function without a computationally demanding inverse rendering optimization [Gkioulekas et al. 2013].

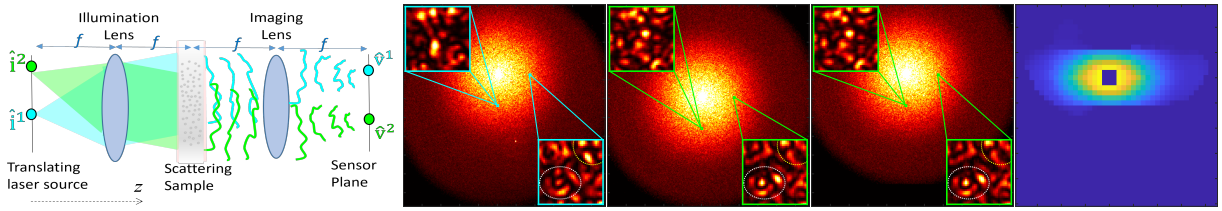

(a) Setup (b) $I^{\hat{\mathbf{i}}^1}(\hat{\mathbf{v}})$ (c) $I^{\hat{\mathbf{i}}^2}(\hat{\mathbf{v}})$ (d) $I^{\hat{\mathbf{i}}^2}(\hat{\mathbf{v}})$ shifted (e) Local correlations

Figure 4: Imaging setup and correlation estimation: (a) A 4f system with two lenses and a sample between them. An illumination source is varied at the focal plane of the first lens to generate a directional beam. Scattered waves are captured by a sensor at the focal plane of the 2nd lens. (b) Speckle image $I^{\hat{\mathbf{i}}^1}(\hat{\mathbf{v}})$ from the first illumination, the saturated pixels correspond to the direction $\hat{\mathbf{i}}^1$. (c) Speckle image $I^{\hat{\mathbf{i}}^2}(\hat{\mathbf{v}})$ from the second illumination. We zoom on the speckle pattern in two local windows, one of them displaced horizontally relative to the saturated region, while the other is displaced vertically. The horizontal window exhibits better correlation than the vertical window. (d) Image (c) shifted to match (b). (e) Local correlations between (b) and (d), evaluated as in Eq. (4), demonstrating the anisotropic shape of the correlation. Correlation is stronger along the horizontal axis than along the vertical axis. Correlation at the very central angles is corrupted by the ballistic term.

Acquisition setup. Our imaging system is visualized in Fig. 4(a). We use a 4f system consisting of two lenses of focal length f separated by distance $2f$, and place the scattering slab halfway between them. We illuminate the system with the output of a fibered laser (632nm), acting as a coherent point light source, placed at the focal plane of the first lens (i.e. the plane at distance f before it). This generates a collimated beam illuminating the sample. The point source is placed on a 2D translation stage, allowing us to translate it laterally along the focal plane. By translating the point source we can tilt the illumination beam, illuminating the sample in different angles. Similarly, we place a 2D sensor at distance f after the second lens, so that each sensor pixel measures a directional view of the scattering target, (as all rays emerging from the scene at a certain direction merge into one sensor point on the focal plane). Therefore, for each illumination angle $\hat{\mathbf{i}}$, we can capture on the sensor plane an image $I^{\hat{\mathbf{i}}}(\hat{\mathbf{v}})$ corresponding to a dense set of viewing directions $\hat{\mathbf{v}}$. For simplicity, we assume that the optical axis of the imaging system is aligned with the z axis.

Empirical correlations. In Fig. 4 we demonstrate a typical acquisition. We move the illumination source vertically, and capture two speckle images $I^{\hat{\mathbf{i}}^1}(\hat{\mathbf{v}}), I^{\hat{\mathbf{i}}^2}(\hat{\mathbf{v}})$ from two illumination directions. The speckle spread in Fig. 4(b-c) is limited to a circle corresponding to the aperture of the acquisition lens. We shift $I^{\hat{\mathbf{i}}^2}$ by $\boldsymbol{\Delta}$ to get $I^{\hat{\mathbf{i}}^2}(\hat{\mathbf{v}}+\boldsymbol{\Delta})$ (Fig. 4(d)), so that the speckle features are aligned.

To compute the correlation at scattering angle $\boldsymbol{\tau}$, we want to look at the product of the two images in Fig. 4(b,d). However, looking at the product at a single pixel would provide a very noisy estimate of the correlation, due to the large variation among speckle patterns. Moreover, when a lot of multiply scattered light is present, the correlation can be rather weak, and its detection is even more sensitive to noise. To eliminate the noise we need to average multiple measurements. For that we use spatial averaging and estimate

$$\begin{aligned}\boldsymbol{C}(\boldsymbol{\tau}) \quad = \quad & \frac{1}{N}\sum_{\boldsymbol{\zeta}\in W} I^{\hat{\mathbf{i}}^1}(\hat{\mathbf{i}}^1+\boldsymbol{\tau}+\boldsymbol{\zeta})\cdot I^{\hat{\mathbf{i}}^2}(\hat{\mathbf{i}}^1+\boldsymbol{\tau}+\boldsymbol{\zeta}+\boldsymbol{\Delta}) \\ & -\frac{1}{N}\sum_{\boldsymbol{\zeta}\in W} I^{\hat{\mathbf{i}}^1}(\hat{\mathbf{i}}^1+\boldsymbol{\tau}+\boldsymbol{\zeta})\cdot \frac{1}{N}\sum_{\boldsymbol{\zeta}\in W} I^{\hat{\mathbf{i}}^2}(\hat{\mathbf{i}}^1+\boldsymbol{\tau}+\boldsymbol{\zeta}+\boldsymbol{\Delta}),\end{aligned} \tag{4}$$

where W denotes a local window around $\boldsymbol{\tau}$, and N is the number of pixels in that window. Larger windows reduce noise but also reduce the resolution at which we measure the phase function. In practice, in our implementation we used 100×100 pixel windows, where the grain of individual speckle features is about 2 pixels. The total aperture diameter is about 4000 pixels. In Appendix Fig. 7 we test the effect of this window size. To further reject noise, we imaged multiple samples of the same material and averaged the correlations produced by such pairs as in Eq. (1). We discuss this and other implementation details in App. A.

In App. C we derive an analytic expression for the signal to noise ratio at which $\boldsymbol{C}(\boldsymbol{\tau})$ can be detected:

$$\mathrm{SNR} = \frac{N|\boldsymbol{C}(\boldsymbol{\tau})|^2}{E[I^{\hat{\mathbf{i}}^1}(\hat{\mathbf{i}}^1+\boldsymbol{\tau})]^4}, \tag{5}$$

where the denominator corresponds to the averaged intensity in the image region. The SNR improves when the number of averaged pixels N increases. More importantly, Eq. (5) implies that the SNR decreases when the ratio between the actual correlation $\boldsymbol{C}(\boldsymbol{\tau})$ and the overall measured intensity is lower. Thus, an estimate of $\boldsymbol{C}(\boldsymbol{\tau})$ is noisier as the optical depth of the material increases, as according to Eq. (3), correlation decays with OD. However, below we explain that for a given OD, some selections of the illumination and viewing directions lead to better correlation estimates.

Selecting measurement directions. As stated in Sec. 2, we follow previous work and consider rotationally invariant phase functions, which depend only on $|\boldsymbol{\tau}|$, the magnitude of the angle between the illumination and viewing directions, rather than on its exact orientation. We note, however, that even if the scattering is isotropic and ρ is rotationally invariant, the correlation in Eq. (3) is *not* rotationally invariant. According to Eq. (3), some selections of $\boldsymbol{\tau}$ would lead to a lower correlation than others, due to the *sinc* term. Thus, out of all $\boldsymbol{\tau}$ directions with the same magnitude $|\boldsymbol{\tau}|$, we can select directions which are expected to produce higher correlations in Eq. (3), by selecting the illuminator to viewpoint displacement $\boldsymbol{\tau}$ to be orthogonal to the illuminators displacement $\boldsymbol{\Delta}$:

$$(\boldsymbol{\Delta}\cdot\boldsymbol{\tau}) = 0. \tag{6}$$

Fig. 4(e) visualizes the windowed correlations of Eq. (4). We see high correlation along the horizontal axis, and correlation decays when moving along the vertical axis. As we shift the source vertically, this agrees with Eq. (6), stating that the correlation is maximal when $\boldsymbol{\Delta}, \boldsymbol{\tau}$ are orthogonal. As another way to see this, the insets of

Fig. 4(b,c) zoom on two speckle windows. For a window displaced horizontally, speckle patterns under the two illuminations resemble each other, while for a vertical displacement, no visual similarity is observed. According to Claim 1, if we move along the horizontal axis, the correlation values we read are proportional to the phase function of the material.

Ballistic light. When light propagates through a scattering volume, a significant part of the energy is scattered, but depending on the OD, some amount of ballistic light continues in the original direction. Since our acquisition setup is based on a 4f relay, the ballistic light is focused in a small number of pixels. Despite the fact that this is an exponentially small portion of the overall energy, as this energy is focused at a small number of pixels, these ballistic pixels are highly saturated, as can be seen in Fig. 4(b,c). As the energy of the ballistic term at these pixels is much higher than the energy of the scattered light, it heavily corrupts scattered values, even when HDR capture is used. We elaborate on this problem in App. A. The ballistic pixels correspond to small τ values, and as a result, the smaller angles of the phase function cannot be measured.

At the other end, the lens aperture limits the maximal scattering angle we could measure to about 8°. Thus, our current setup only measures a limited portion of the phase function. One can remove this restriction with a more complicated acquisition setup that would allow rotating the camera around the sample, but we did not implement it for this proof-of-concept demonstration.

4 RESULTS

4.1 Validation materials

To validate our approach we generated two scattering materials, by dispersing two types of micro-spheres in agarose: $10\mu m$ diameter SiO2, and $4\mu m$ polystyrene. To avoid scattering at the surface interface, we sandwich the material between two microscope glass slides to generate flat boundaries. We used micron-precision spacing stickers (ThermoFisher AB0578) between the two slides to ensure they are parallel, and the sample thickness is uniform. These stickers are of thickness $250\mu m$ and we used $1-2$ layers depending on the material of interest. We use Snell's law to correct the ingoing and outgoing ray directions, accounting for the refractive indices of the agar, glass and air.

We estimate the optical depth of the material by measuring the ratio between the intensity of an unscattered light (without a sample), and the ballistic term. This is an approximation since the ballistic term is somewhat aberrated. For the two validation materials, absorption is negligible.

As particle diameter and refractive index are known, we can compute the phase function using Mie theory [Bohren and Huffman 1983; Frisvad et al. 2007]. In Fig. 5 we plot the estimated correlations for a few illuminator displacements. At $\Delta_y = 0$ the correlation involves all multiple scattering paths, and for thick samples this is very different from the ground truth phase function. However, as the displacement Δ_y increases the correlation can be attributed to single scattering paths alone, and *empirical correlation agrees well with the ground truth.* We have repeated the measurement with multiple slices of the same material, at various thicknesses, all showing good agreement with the ground truth. We could estimate the correlation even at $OD = 7$, where most paths scatter much more than once. As discussed above, when material thickness increases the correlation estimate is noisier. Appendix Fig. 11 demonstrates a failure case.

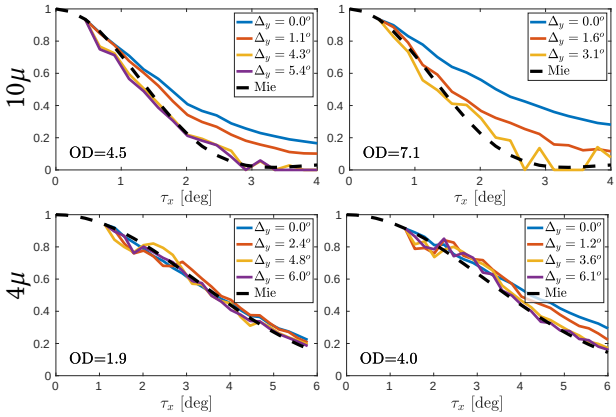

Figure 5: Validation results: We compare speckle correlations acquired with our setup to the ground truth phase function computed using Mie theory. For small illumination displacements Δ, the correlation contains multiple scattering. As Δ increases, correlation is only attributed to single scattering and is well matched with the Mie function. Correlation can isolate single scattering even in thick samples with $OD = 7$. Top row: $10\mu m$ Si02 particles, and lower row $4\mu m$ polystyrene, both mixed with agarose. For each material we demonstrate two samples at different thicknesses.

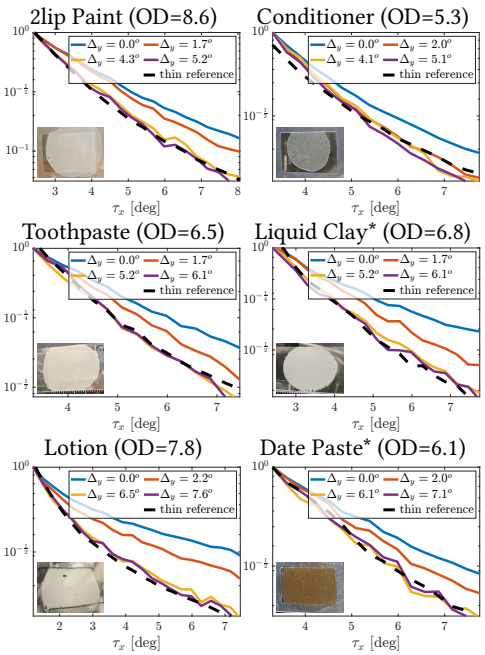

Figure 6: Everyday materials: Speckle correlation measured from various everyday materials. While we do not have a ground truth reference, we measured both a thin and a thick sample of the same material. Isolating single scattering correlation on the thinner sample is simpler, and we compare it with the correlations extracted from the thicker samples. Materials where the thinner reference is not a real single scattering one are marked with asterisk. vertical axis uses log scale.

4.2 Common materials

After validation, we used our approach for acquiring the phase functions of various everyday materials. We plot some results in Fig. 6, and Appendix Fig. 10. While we do not have ground truth, we prepared both a thin and a thicker sample of the same material. The thinner samples allow for a better correlation estimate, and we use them as a reference, while plotting correlation curves measured from the thicker samples. As before, at small illuminator displacements Δ_y the correlation is not free of multiple scattering paths, but as Δ_y increases, single scattering dominates and the correlation matches the one measured from an optically thinner sample.

For some of the tested materials, the thin samples are single-scattering (measured OD around 1). However, not all materials could be made sufficiently thin. Figs. 6 and 10 mark with asterisks materials where a single scattering reference was not achieved.

As another way of assessing the quality of the estimate, we note that for $\boldsymbol{\tau}$ values orthogonal to $\boldsymbol{\Delta}$, the single scattering correlation formula in Claim 1 does not depend on $\boldsymbol{\Delta}$. In practice, for small Δ_y values the correlations estimated from pairs at different Δ_y displacements do vary with Δ_y, because they include multiply scattered light. However, as Δ_y increases, correlations for different Δ_y values are consistent, see e.g. the yellow and purple curves in the figure. This agreement implies that at such Δ_y values, we have managed to isolate the single scattering component of the light, and we can use these correlations as an estimate of the true phase function.

In App. B we fitted the correlation-based phase functions of Fig. 6 with a Henyey-Greenstein model [Henyey and Greenstein. 1941], and provide a table summarizing anisotropy parameters. We compare these with a simpler, intensity-based estimate of the phase function as we explain next.

Code and data. are available for public usage at this URL.

4.3 Rendering visualization

In Fig. 1(c) and Appendix Fig. 12 we visualize the difference between the phase function estimated from speckle correlations and the naive estimation from intensity profiles. Intensity estimation simply assumes the material is thin enough to be treated as singly scattered, and hence the intensity scattered to different angles is the phase function. As explained by [Alterman et al. 2021], this is equivalent to the $\Delta_y = 0$ curve in the previous figures. We fitted both estimates with a Henyey-Greenstein phase function [Henyey and Greenstein. 1941], and used it in a Mitsuba renderer to render a simple scene, consisting of a scattering cube and a diffused basis plane. An incoherent illumination source is located behind the cube, and light scattering through this cube casts some shadow on the basis plane. While the difference between the correlation-based and intensity-based phase estimate may seem small when plotted as a function of angle, the small gap translates into a significant appearance difference.

5 DISCUSSION

We present a simple approach for reading the phase function of a volumetric-scattering material from the correlation between speckle patterns, produced when the material is imaged under coherent laser illumination. We show that speckle correlations serve as a gating mechanism, which can isolate single-scattering light paths. Thus, even in thick samples where most light paths scatter multiple times, the correlated portion can be explained by single scattering paths. This singly-scattered correlation is proportional to the phase function. We have validated our approach against ground truth materials, demonstrating accurate agreement.

Below we discuss remaining limitations.

Optical thickness. Our approach can estimate phase functions from material samples of modest optical depth. This restriction results from the fact that as OD increases, correlation is weaker, and its estimation is noise sensitive. In our current implementation we managed to measure materials of OD below 8.

Angular range. Our current proof-of-concept implementation only covers a limited range of angles. The largest scattering angle is limited by the camera aperture, and in our setup it is about 8°. Our attempts to increase the aperture resulted in too many aberrations. Capturing larger angles would require a more complex setup that can rotate the camera around the sample, as in [Gkioulekas et al. 2016]. From the other end, the very forward scattering angles are precluded by ballistic light.

Estimating material density. Another important component of material acquisition is the estimation of $\sigma_t, \sigma_s, \sigma_a$, or equivalently, the estimation of particle density and absorption. Ideally, σ_t can be computed by measuring the sample thickness L and the ratio between the maximal intensity of the ballistic light and the actual power of the laser (measured without any material). However, as the ballistic light is aberrated, this provides a coarse estimate.

Solid samples. Our approach requires two successive images of the same material imaged without any Brownian motion of internal particles, as such motion will reduce the correlation between the measured speckles. This can effectively be achieved with viscous or solid materials. However, these materials are actually the main open challenge, as liquids can easily be diluted and measured as in [Narasimhan et al. 2006].

Tissue. One of the more interesting applications of material acquisition is distinguishing between different types of tissue, e.g. for tumor detection. However, tissue is not yet covered by our model. The first issue is that we assume a flat interface and ignore any scattering and refraction at the surface boundary. In practice, tissue surface is not flat and a lot of scattering happens at the boundary. Additionally, in contrast to our isotropic scattering assumption, tissue exhibits significant anisotropy due to orientated tissue fibers.

ACKNOWLEDGMENTS

We thank Ioannis Gkioulekas for many helpful discussions, and Onit Alalouf for her help preparing experimental data. This project was funded by H2020 European Research Council (635537); United States-Israel Binational Science Foundation (2008123/2019758); Israel Science Foundation (1947/20).

REFERENCES

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

M. Alterman, C. Bar, I. Gkioulekas, and A. Levin. 2021. Imaging with local speckle intensity correlations: theory and practice. *ACM TOG* (2021).

V.S. Antyufeev. 2000. *Monte Carlo method for solving inverse problems of radiation transfer.* Vol. 20. Inverse and Ill-Posed Problems Series, V.S.P. International Science.

G. Bal. 2009. Inverse transport theory and applications. *Inverse Problems* 25, 5 (2009).

Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. 2019. A Monte Carlo Framework for Rendering Speckle Statistics in Scattering Media. *ACM TOG* (2019).

C. Bar, M. Alterman, I. Gkioulekas, and A. Levin. 2021. Single scattering modeling of speckle correlation. In *ICCP*.

Ibrahim Baydoun, Diego Baresch, Romain Pierrat, and Arnaud Derode. 2016. Radiative transfer of acoustic waves in continuous complex media: Beyond the Helmholtz equation. *Physical Review E* (2016).

Richard Berkovits and Shechao Feng. 1994. Correlations in coherent multiple scattering. *Physics Reports* (1994).

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

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

David A Boas, Dana H Brooks, Eric L Miller, Charles A DiMarzio, Misha Kilmer, Richard J Gaudette, and Quan Zhang. 2001. Imaging the body with diffuse optical tomography. *IEEE Signal Processing Magazine* (2001).

Craig F. Bohren and Donald R. Huffman. 1983. *Absorption and scattering of light by small particle.* John Wiley & Sons.

Julie Chang and Gordon Wetzstein. 2018. Single-shot speckle correlation fluorescence microscopy in thick scattering tissue with image reconstruction priors. *Journal of Biophotonics* (2018).

Chengqian Che, Fujun Luan, Shuang Zhao, Kavita Bala, and Ioannis Gkioulekas. 2020. Towards Learning-based Inverse Subsurface Scattering. In *2020 IEEE International Conference on Computational Photography (ICCP)*. 1–12.

Wai-Fung Cheong, Scott A. Prahl, and Ashley J. Welch. 1990. A review of the optical properties of biological tissues. *IEEE JQE* (1990).

Craig Donner, Tim Weyrich, Eugene d'Eon, Ravi Ramamoorthi, and Szymon Rusinkiewicz. 2008. A Layered, Heterogeneous Reflectance Model for Acquiring and Rendering Human Skin. *ACM Trans. Graph.* 27, 5, Article 140 (dec 2008), 12 pages.

Ronald L. Dougherty, Bruce J. Ackerson, Nafaa M. Reguigui, F. Dorri-Nowkoorani, and Ulf Nobbmann. 1994. Correlation transfer: Development and application. *JQSRT* (1994).

Turgut Durduran, Regine Choe, Wesley B. Baker, and Arjun G. Yodh. 2010. Diffuse optics for tissue monitoring and tomography. *Reports on Progress in Physics* (2010).

Phil Dutre, Kavita Bala, and Philippe Bekaert. 2006. *Advanced Global Illumination.* A K Peters, Natick, MA.

Eitan Edrei and Giuliano Scarcelli. 2016a. Memory-effect based deconvolution microscopy for super-resolution imaging through scattering media. *Scientific Reports* (2016).

Eitan Edrei and Giuliano Scarcelli. 2016b. Optical imaging through dynamic turbid media using the Fourier-domain shower-curtain effect. *Optica* (2016).

Robert K. Erf. 1978. Speckle Metrology. Elsevier.

T.J. Farrell, M.S. Patterson, and B. Wilson. 1992. A diffusion theory model of spatially resolved, steady-state diffuse reflectance for the noninvasive determination of tissue optical properties in vivo. *Med Phys.* (1992).

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

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

David L. Fried. 1982. Anisoplanatism in adaptive optics. *JOSA* (1982).

Jeppe R. Frisvad, Niels J. Christensen, and Henrik Wann Jensen. 2007. Computing the scattering properties of participating media using Lorenz-Mie theory. *ACM TOG* (2007).

C. Fuchs, T. Chen, M. Goesele, H. Theisel, and H.P. Seidel. 2007. Density estimation for dynamic volumes. *Computers & Graphics* 31, 2 (2007).

I. Gkioulekas, A. Levin, and T. Zickler. 2016. An Evaluation of Computational Imaging Techniques for Heterogeneous Inverse Scattering. *European Conference on Computer Vision* (2016).

Ioannis Gkioulekas, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin. 2013. Inverse Volume Rendering with Material Dictionaries. *ACM TOG* (2013).

Joseph W. Goodman. 2007. *Speckle Phenomena in Optics: Theory and Applications.* Roberts and Company Pub.

J. Gu, S. Nayar, E. Grinspun, P. Belhumeur, and R. Ramamoorthi. 2008. Compressive structured light for recovering inhomogeneous participating media. *ECCV* (2008).

T. Hawkins, P. Einarsson, and P. Debevec. 2005. Acquisition of time-varying participating media. *ACM Trans. Graph.* 24, 3 (2005).

L. C. Henyey and J.L. Greenstein. 1941. Diffuse radiation in the galaxy. *Astrophys.* 93 (1941), 70–83.

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

Pierre Jacquot and Jean-Marc Fournier. 2000. *Interferometry in Speckle Light.* Springer.

H.W. Jensen, S.R. Marschner, M. Levoy, and P. Hanrahan. 2001. A practical model for subsurface light transport. In *Proceedings of SIGGRAPH 2001, Annual Conference Series.*

Ori Katz, Pierre Heidmann, Mathias Fink, and Sylvain Gigan. 2014. Non-invasive single-shot imaging through scattering layers and around corners via speckle correlation. *Nature Photonics* (2014).

Guillermo H. Kaufmann. 2011. *Advances in Speckle Metrology and Related Techniques.* Wiley.

Pramook Khungurn, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner. 2015. Matching Real Fabrics with Micro-Appearance Models. *ACM Trans. Graph.* 35, 1 (2015), 1:1–1:26.

A. Levis, Y.Y. Schechner, A. Aides, and A.B. Davis. 2015. Airborne Three-Dimensional Cloud Tomography. *IEEE International Conference on Computer Vision* (2015).

Aviad Levis, Yoav Y Schechner, and Anthony B Davis. 2017. Multiple-scattering microphysics tomography. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.* 6740–6749.

Sven Leyre, Youri Meuret, Guy Durinck, Johan Hofkens, Geert Deconinck, and Peter Hanselaer. 2014. Estimation of the effective phase function of bulk diffusing materials with the inverse adding-doubling method. *Appl. Opt.* 53, 10 (Apr 2014), 2117–2125.

Chao Liu, Akash K. Maity, Artur W. Dubrawski, Ashutosh Sabharwal, and Srinivasa G. Narasimhan. 2020. High Resolution Diffuse Optical Tomography using Short Range Indirect Subsurface Imaging. *ICCP* (2020).

Adolfo Munoz, Jose I. Echevarria, Francisco J. Seron, Jorge Lopez-Moreno, Mashhuda Glencross, and Diego Gutierrez. 2011. BSSRDF Estimation from Single Images. *Computer Graphics Forum* 30, 2 (2011), 455–464.

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. In *ACM Transactions on Graphics (TOG)*.

Roger Newton. 2002. *Scattering Theory of Waves and Particles.* Dover Publications.

Merlin Nimier-David, Sébastien Speierer, Benoît Ruiz, and Wenzel Jakob. 2020. Radiative Backpropagation: An Adjoint Method for Lightning-Fast Differentiable Rendering. *ACM Trans. Graph.* 39, 4, Article 146 (jul 2020), 15 pages.

Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. 2019. Mitsuba 2: A Retargetable Forward and Inverse Renderer. *ACM Trans. Graph.* 38, 6, Article 203 (nov 2019), 17 pages.

Thomas Klaus Nindel, Tomáš Iser, Tobias Rittig, Alexander Wilkie, and Jaroslav Křivánek. 2021. A Gradient-Based Framework for 3D Print Appearance Optimization. 40, 4, Article 178 (jul 2021), 15 pages.

Jan Novak, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. 2018. Monte Carlo Methods for Volumetric Light Transport Simulation. *Computer Graphics Forum* (2018).

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

Marios Papas, Christian Regg, Wojciech Jarosz, Bernd Bickel, Philip Jackson, Wojciech Matusik, Steve Marschner, and Markus Gross. 2013. Fabricating Translucent Materials Using Continuous Pigment Mixtures. *ACM Trans. Graph.* 32, 4, Article 146 (July 2013), 12 pages.

D.J. Pine, D.A. Weitz, J.X. Zhu, and E. Herbolzheimer. 1990. Diffusing-wave spectroscopy: dynamic light scattering in the multiple scattering limit. *Journal de Physique* 51, 18 (1990).

S.A. Prahl, M.J.C. van Gemert, and A.J. Welch. 1993. Determining the optical properties of turbid media by using the adding–doubling method. *Applied optics* 32, 4 (1993).

Yael Sde-Chen, Yoav Y. Schechner, Vadim Holodovsky, and Eshkol Eytan. 2021. 3DeepCT: Learning Volumetric Scattering Tomography of Clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 5671–5682.

Denis Sumin, Tobias Rittig, Vahid Babaei, Thomas Nindel, Alexander Wilkie, Piotr Didyk, Bernd Bickel, J KŘivánek, Karol Myszkowski, and Tim Weyrich. 2019. Geometry-Aware Scattering Compensation for 3D Printing. *ACM Transactions on Graphics* 38 (2019).

Kevin T. Takasaki and Jason W. Fleischer. 2014. Phase-space measurement for depth-resolved memory-effect imaging. *Optical Express* (2014).

Bertrand Thierry, Xavier Antoine, Chokri Chniti, and Hasan Alzubaidi. 2015. μ-diff: An open-source Matlab toolbox for computing multiple scattering problems by disks. *Computer Physics Communications* (2015).

Bradley E. Treeby and Ben T. Cox. 2010. k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave-fields. *JBO* (2010).

Valery Tuchin. 2000. Tissue Optics Light Scattering Methods and Instruments for Medial Diagnosis. *SPIE* 13 (01 2000).

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* (1964).

H.C. van de Hulst. 1981. *Light scattering by small particles.* Dover.

Zdravko Velinov, Marios Papas, Derek Bradley, Paulo Gotardo, Parsa Mirdehghan, Steve Marschner, Jan Novák, and Thabo Beeler. 2018. Appearance Capture and Modeling of Human Teeth. *ACM Trans. Graph.* 37, 6, Article 207 (dec 2018), 13 pages.

J. Wang, S. Zhao, X. Tong, S. Lin, Z. Lin, Y. Dong, B. Guo, and H.Y. Shum. 2008. Modeling and rendering of heterogeneous translucent materials using the diffusion equation. *ACM Trans. Graph.* 27, 1 (2008).

D.R. Wyman, M.S. Patterson, and B.C. Wilson. 1989. Similarity relations for the interaction parameters in radiation transport. *Applied optics* 28, 24 (1989).

Kane Yee. 1966. Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media. *IEEE TAP* (1966).

Cheng Zhang, Lifan Wu, Changxi Zheng, Ioannis Gkioulekas, Ravi Ramamoorthi, and Shuang Zhao. 2019. A Differential Theory of Radiative Transfer. *ACM Trans. Graph.* 38, 6, Article 227 (nov 2019), 16 pages.

Shuang Zhao, Ravi Ramamoorthi, and Kavita Bala. 2014. High-order similarity relations in radiative transfer. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 104.

Quan Zheng, Gurprit Singh, and Hans-peter Seidel. 2021. Neural Relightable Participating Media Rendering. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 15203–15215.

Direct material acquisition using speckle correlations

Supplementary Appendix

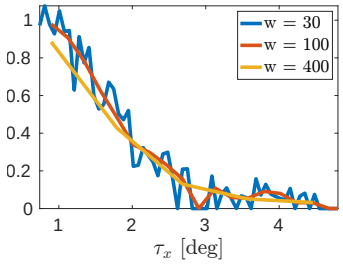

Figure 7: Window size: Correlations computed, with different window sizes w for spatial averaging. For smaller window size the correlation is noisy, and for very large ones the phase function is over smoothed. This result corresponds to the top left sample of Fig. 5, demonstrating $10\mu m$ beads at $OD = 4.5$, correlation for $\Delta_y = 5.4^o$.

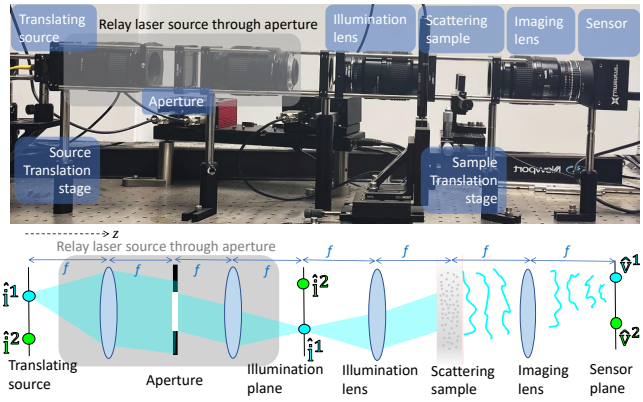

Figure 8: Full hardware setup: Top row: Our hardware lab setup. The setup consists of two 4f systems, which we implement using 4 macro lenses (Nikon 105 mm f/2.8D). Lower row: A schematic of our setup. The first 4f system relays the laser source through the aperture to the illumination plane. The scattering sample is located at the Fourier plane of a the imaging lens. We use a camera on the opposite side of the sample, focused at infinity. This part corresponds to a 4f system around the sample.

A IMPLEMENTATION DETAILS

Setting illumination diameter. To build a speckle acquisition system, one should adjust the resolution of individual speckle features (namely the speckle grain) to be about 1-2 pixels wide. This is a combination of two conflicting goals. On the one hand, if the speckles are smaller then a pixel, their contrast is blurred. On the other hand, if speckle features are much larger than a pixel, this is also sub-optimal as it reduces the number of different speckle samples we collect on the sensor. As the speckles we measure are in the Fourier plane, their size is inversely controlled by the diameter of the beam illuminating the scattering sample. In our implementation, with a $10.5cm$ focal length, a $5.5\mu m$ pixel pitch and a $632nm$ illumination, we used a $3mm$ beam diameter.

Spacing glass slides. In our implementation we cast a sample of the target material sandwiched between two microscope glass slides. We used spacing stickers (gene frames from ThermoFisher scientific AB0578) to ensure the two slides are actually parallel and the sample thickness is uniform. The stickers have a thickness of $250\mu n$ and an open area of 28×17 mm. For some of the optically thick materials we tested, the $250\mu m$ spacing translated into a very large OD. As we also tried to get real single scattering references, in some cases we just used tape layers whose thickness was measured at about $40\mu m$. The result, however, is less uniform than the spacing stickers.

Capturing multiple images. If we capture two images of the same sample and compute their empirical product as in Eq. (4), we often still get a noisy estimate to the speckle correlation, because the number of samples we have in a 100×100 pixel window is limited. This is especially a problem when the OD is high and correlation is weak relative to the overall scattered energy.

In Fig. 7 we test the effect of the window size on the estimated covariance. As the number of pixels N in the window increase, less noise is present. However, if the window is too wide, the structure of the phase function is blurred out.

To improve correlation estimation, we want more than one sample of the same material so that we can average the empirical speckle products from multiple samples as in Eq. (1). With an illumination beam diameter of $3mm$ we could illuminate multiple locations on the 28×17 mm sample and capture multiple speckle images. To this end we place our sample on a motorized translation stage, shift it, and capture multiple images. In practice not all translated images are good—the sample may contain some air bubbles, and also, despite the usage of spacing stickers, its thickness is not fully uniform. To filter out the bad images, we look at the maximal energy of the ballistic term, and reject images in which the ballistic term is too far from the median.

To allow translating the sample independently of the $3mm$ aperture, we effectively built two relay systems and placed the aperture at a plane conjugate to the sample plane, see Fig. 8.

Registering speckle images. Another issue we had to address is that the lens introduced some aberration. Thus, converting a pixel at position xy to direction is not a simple inverse tan function. Without accounting for this distortion, the correlations cannot be detected. Our goal is to compute correlations between images $I^{\hat{i}^1}(\hat{\mathbf{v}})$, $I^{\hat{i}^2}(\hat{\mathbf{v}})$ taken under different illumination directions. We start by detecting the strongest pixel in each image, which corresponds to the directions $\hat{\mathbf{i}}^1, \hat{\mathbf{i}}^2$ and compute the shift $\Delta = \hat{\mathbf{i}}^2 - \hat{\mathbf{i}}^1$. We globally shift $I^{\hat{i}^2}(\hat{\mathbf{v}})$ by Δ as in Fig. 4(d), but due to optical distortions, this is usually not enough to actually align the images. As a result, in different image positions we still see some local shifts between the

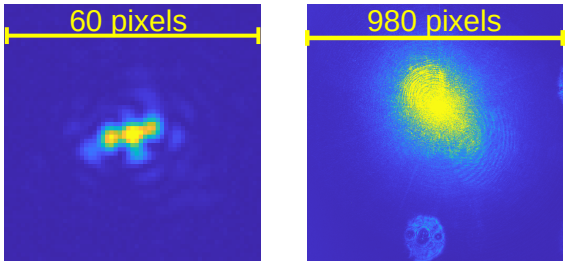

Figure 9: Aberrations of the ballistic light (a) Aberrations due to surface imperfections of the sample. The center of the ballistic term is not a symmetric Jinc. (b) Reflections and flare in a direct image of the laser (no scattering sample). These may not be negligible when added to the scattered speckle values. Note that the two images in (a,b) zoom on a different part of the sensor.

two speckle images. To account for this, for every 100×100 window we average, we first try to locally shift the windows by a few pixels to maximize the correlation.

Validation-sample preparation. To prepare validation materials we mixed agarose powder with distilled water, and added precision micro spheres. We mixed the hot blend using the Vortex mixer that blends fluids very thoroughly. After the blend cools down, this results in micro spheres embedded uniformly in an agarose gel. We used silica (SiO2) micro spheres of $10\mu m$ sold in powder and $4\mu m$ polystyrene latex particles in liquid suspension. Both were provided by Merck.

Ballistic problems. Despite the fact that the ballistic term is an exponentially small portion of the overall energy, in our 4f relay all ballistic light is focused in a small set of pixels. Thus, pixels corresponding to ballistic light are highly saturated, as can be seen in Fig. 4(b,c). In theory, for a circular illumination beam the ballistic term can be modeled as a Jinc function, and we have tried to use HDR capture and subtract it from the speckle image. This is not practical for two reasons, visualized in Fig. 9. The first problem is demonstrated in Fig. 9(a). As our aperture is circular we expect the ballistic term to look like a radially symmetric Jinc function. However, as demonstrated in Fig. 9(a), often we do not achieve this exact shape. This is due to the fact that despite our attempts to generate flat surface boundaries using microscope glass-slides, the sample is not sufficiently uniform, either because the glass is not of sufficiently high precision, or because the leading medium is not truly uniform. For the validation materials, the precision of the agarose is high, yet some small volumetric variations in the index of refraction are present. This is even more of a problem with various everyday materials, which are not designed to be uniform in any scientific standard. A second problem is demonstrated in Fig. 9(b), showing a long exposure image of the laser light going directly to the camera, without any sample. Some flare and other inter-reflections through the optics are visible. While these are orders of magnitude weaker with respect to the maximum of the ballistic term, they are not always negligible with respect to the scattered wavefront, because the scattered light is weaker than the maximal value of the ballistic term. As a result, a small amount of flare can heavily corrupt the speckle measurements.

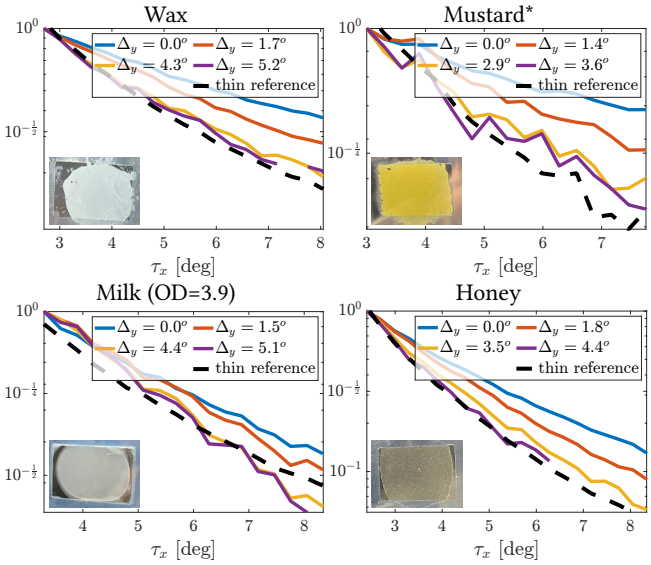

Figure 10: Additional materials: Speckle correlation measured from various common materials. While we do not have a ground truth reference, we measured both a thin and a thick sample of the same material. Isolating the single scattering correlation on the thinner sample is easier. We compare it with the correlations extracted from the thicker samples. Small Δ displacements involve multiple scattering, but with larger ones, correlation is mostly due to single scattering light, and matches the reference from the thinner sample. * stands for materials where the thinner reference is not a real single scattering one. y-axis is plotted using log scale.

B ADDITIONAL RESULTS

Fig. 10 follows on Fig. 6, demonstrating acquired phase function from additional materials. Note that the milk sample in Fig. 10 is mixed with agarose powder to solidify it, as we cannot measure liquids.

In our validation materials there is not much absorption and whenever a light path interacts with a scattering particle it actually scatters to a new direction. In contrast, some of the common materials in Figs. 6 and 10 do absorb light. Hence the multiply scattered light is attenuated compared to the singly scattered light. This absorption makes the detection of single scattering correlation less noisy, even in larger ODs. In Fig. 11 we demonstrate failure case, where we attempt to estimate the correlation from a very thick sample, of $OD = 11$. This correlation is too noisy to be of any practical usage.

For the phase functions in Figs. 6 and 10, we use the measured 8° angular cone to fit a parametric Henyey-Greenstein model [Henyey and Greenstein. 1941]. We summarize the anisotropy parameters of these functions in Table 1. We compare these values with the anisotropy parameter of a naive estimate from intensity profiles.

In Fig. 12 we visualize the effect of phase function accuracy when used for rendering, see discussion in Sec. 4.3 of the main paper.

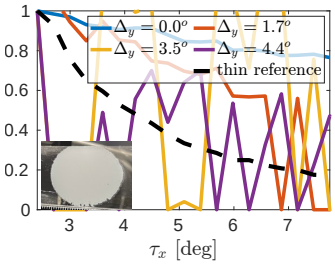

Figure 11: Noisy correlation result: Thick clay with $OD = 11.1$, at this thickness the correlation computed is too noisy to be of practice usage.

Material name	g (intensity)	g (correlation)
10μm SiO2	0.94	0.97
4μm polystyrene	0.91	0.93
Lotion	0.85	0.90
Honey	0.93	0.96
Wax	0.88	0.93
Toothpaste	0.88	0.92
2lip	0.94	0.97
Milk	0.87	0.92
Mustard	0.79	0.87
Date	0.91	0.95
Conditioner	0.93	0.96
Clay	0.87	0.92

Table 1: Summarizing the anisotropy parameter g for measured phase functions. We fit the measured phase functions in Figs. 6 and 10 with a parametric Henyey-Greenstein model [Henyey and Greenstein. 1941] and summarize the numerical values.

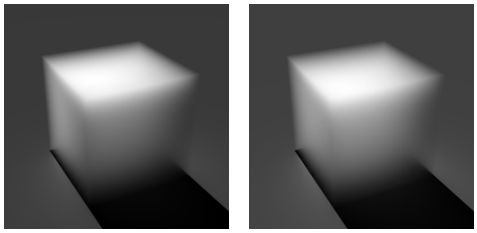

Lotion Intensity, $g = 0.85$ Lotion Correlation, $g = 0.9$

Figure 12: Rendering with estimated phase function: Visualizing appearance differences, for a phase function estimated from speckle correlation vs. a naive estimate from intensity profile. We use a lotion sample from Fig. 6 of the main paper.

C SIGNAL TO NOISE RATIO IN COVARIANCE ESTIMATION

Here we explain Eq. (5) of the main paper, and derive the signal to noise ratio we expect to achieve, when estimating covariance empirically from speckle images. For that, let us consider two complex random variables u_1, u_2, representing the complex speckle fields scattered from the same medium under two different illumination directions. It has been previously noted [Goodman 2007] that due to the central limit theorem, speckle fields follow a multi-variate Gaussian distribution. We thus assume that u_1, u_2 are sampled from a zero mean Gaussian distribution with covariance

$$\begin{pmatrix} C_{11} & C_{12} \\ C_{12}^* & C_{11} \end{pmatrix} \tag{7}$$

where we assume that the two images have the same mean intensity $E[|u_1|^2] = E[|u_2|^2] = C_{11}$. Note that C_{11} is a real positive number. In contrast, $C_{12} = E[u_1 \cdot u_2^*]$ can be complex.

A camera measures only the intensity of the speckle fields $|u_1|^2$, $|u_2|^2$. From these intensities we want to estimate the covariance, and we thus define a new random variable

$$X = (|u_1|^2 - C_{11})(|u_2|^2 - C_{11}). \tag{8}$$

For simplicity we assume here that the mean intensity C_{11} is known, through in practice we also estimate it from the speckle image. The SNR in the estimation of X is defined as

$$SNR(X) = \frac{E[X]^2}{Var[X]}. \tag{9}$$

Some derivation shows that $E[X] = |C_{12}|^2$, and

$$Var[X] = \frac{1}{16}\left(62C_{11}^2|C_{12}|^2 + 63|C_{12}|^4 + 15C_{11}^4\right) \tag{10}$$

Usually the correlation between the two speckle fields is not exact. Hence, covariance is smaller than the mean intensity and $|C_{12}| < C_{11}$. Thus, the largest term in Eq. (10) is C_{11}^4, and we simplify it by neglecting the first two terms. We thus express the SNR as

$$SNR(X) = \frac{|C_{12}|^4}{C_{11}^4}. \tag{11}$$

To relate this to Eq. (5) of the main paper, we note that the correlation between speckle intensity we aim to detect is essentially $\boldsymbol{C}(\boldsymbol{\tau}) = E[X] = |C_{12}|^2$. Also, the mean speckle intensity is C_{11}.

Finally, if we average N independent speckle measurements, the SNR is scaled by N. This is because basic statistics implies that when N independent samples $X_1, \ldots X_n$ of a random variable X are averaged,

$$E\left[\frac{1}{N}\sum_n X_n\right] = E[X] \tag{12}$$

$$Var\left[\frac{1}{N}\sum_n X_n\right] = \frac{1}{N}Var[X] \tag{13}$$

In Fig. 7 we test the effect of the window size on the estimated covariance. As predicted by theory, as the number of pixels N in the window increase, less noise is present. However, if the window size is too big the structure of the phase function is blurred out. As discussed above, to further eliminate noise, we capture multiple speckle images of the same homogenious material sample.

D COMPONENTS LIST

Our lab setup is composed of the following optical elements. For the laser light source we used a HeNe Laser, 632.8 nm, 21 mW (HNL210LB) from Thorlabs, coupled into an optical fiber. Two 4f relay systems are composed of four macro lenses (Nikon 105 mm f/2.8D). To capture speckle images we used the high resolution (6576 × 4384 pixels) Mono CCD Lumenera Camera (Lt29059HM),

with $5.5 \times 5.5\mu$m pixels. To translate the laser source and the scattering sample we used motorized translation stages from Thorlabs.