

Towards Practical Physical-Optics Rendering

SHLOMI STEINBERG, PRADEEP SEN, and LING-QI YAN, University of California, Santa Barbara, United States



Fig. 1. A scene rendered with our framework, viewed through a polarization filter (e.g., sunglasses), and lit by sunlight and the afternoon sky. Multiple diffraction optical effects are visible: (a) the glass window and (b) moulded plastic spoke guard admit stress birefringence, which results in iridescence depending on viewing direction (see Fig. 8b for a close-up rendering); and (c) the metal brake surfaces on the bicycle’s wheels act as imperfect diffraction gratings, dispersing scattered light (see Figs. 8 and 9 for photos and additional renderings of these effects). Unlike the state-of-the-art which still depends on classical materials for performance, all the materials in this scene are coherence-aware, physical optics materials (see Subsection 4.1), nevertheless, our rendering performance is close to classical radiometric renderers. The appearance of these materials depends on the radiometric, polarimetric, and coherence properties of light, see Fig. 10 for renderings under different lighting conditions.

Physical light transport (PLT) algorithms can represent the wave nature of light globally in a scene, and are consistent with Maxwell’s theory of electromagnetism. As such, they are able to reproduce the wave-interference and diffraction effects of real physical optics. However, the recent works that have proposed PLT are too expensive to apply to real-world scenes with complex geometry and materials. To address this problem, we propose a novel framework for physical light transport based on several key ideas that actually makes PLT practical for complex scenes. First, we restrict the spatial coherence shape of light to an anisotropic Gaussian and justify this

restriction with general arguments based on entropy. This restriction serves to simplify the rest of the derivations, without practical loss of generality. To describe partially-coherent light, we present new rendering primitives that generalize the radiometric radiance and irradiance, and are based on the well-known Stokes parameters. We are able to represent light of arbitrary spectral content and states of polarization, and with any coherence volume and anisotropy. We also present the wave BSDF to accurately render diffractions and wave-interference effects. Furthermore, we present an approach to importance sample this wave BSDF to facilitate bi-directional path tracing, which has been previously impossible. We show good agreement with state-of-the-art methods, but unlike them we are able to render complex scenes where all the materials are new, coherence-aware physical optics materials, and with performance approaching that of “classical” rendering methods.

Authors’ address: Shlomi Steinberg, p@shlomisteinberg.com; Pradeep Sen, psen@ece.ucsb.edu; Ling-Qi Yan, lingqi@cs.ucsb.edu, University of California, Santa Barbara, 2119 Harold Frank Hall, Santa Barbara, California, 93106, United States.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3528223.3530119>.

CCS Concepts: • **Computing methodologies** → **Rendering; Computer graphics**; • **Applied computing** → *Physics*.

Additional Key Words and Phrases: light transport, coherence, interference, diffraction, wave optics, scattering, path tracing, iridescence, polarization, Mueller, Stokes, spectral

ACM Reference Format:

Shlomi Steinberg, Pradeep Sen, and Ling-Qi Yan. 2022. Towards Practical Physical-Optics Rendering. *ACM Trans. Graph.* 41, 4, Article 1 (July 2022), 4 pages. <https://doi.org/10.1145/3528223.3530119>

1 INTRODUCTION

Modern rendering techniques strive to convincingly reproduce a wide range of optical phenomena. However, rendering effects that stem from the wave nature of light remains an open problem. Such effects include diffraction-grating dispersion by compact discs, holographs and LCD screens; iridescence effects in coated or painted materials; and, pleochroism and birefringent effects in moulded plastics and framed glass. Contemporary rendering methodologies have mostly either ignored the wave nature of light entirely, or have only considered wave effects *locally*. But, the process of wave interference, and the visual effects that arise, is driven by the light’s *optical coherence*: the statistical properties of light that dictate its ability to interfere and diffract. These properties change as light propagates through a scene and interacts with matter, and cannot be guessed or approximated locally with an *ad hoc* approach. Essentially, classical “physically-based” rendering approaches have basically ignored these effects by assuming that light in typical scenes is perfectly incoherent. However, no light is truly incoherent [Goodman 2015], and so all light we encounter in real-world scenes is at least *partially coherent*. In order to accurately reproduce the diffractive phenomena that emerge from such light, the coherence properties of light must be considered.

Physical light transport (PLT), introduced by Steinberg and Yan [2021a], is a formulation of light transport that is consistent with electromagnetism. By accounting for the coherence properties of light globally—throughout the scene—PLT facilitates accurate rendering of observable wave-interference effects. PLT borrows upon optical coherence theory [Wolf 2007], and models light as a statistical ensemble of waves and quantifies the correlations between the wavefronts that arise as an optical beam propagates and interacts with media. This approach gives rise to generalized, physical-optics rendering primitives, and Steinberg and Yan [2021a] showed that these new primitives superpose linearly under rather general conditions, and derived generalized linear rendering equations.

This linearity is crucial, as it implies that these primitives can indeed generalize and supersede the classical radiometric quantities that are used in traditional “physically-based” rendering algorithms, with a caveat: Helmholtz reciprocity no longer applies under physical optics (since diffraction is not reciprocal, see Steinberg and Yan [2021a, supplemental Appendix B]). Furthermore, the process of scattering of light by matter depends on the wave ensemble’s statistics—i.e., the light’s optical coherence—and, in general, this information is only available when emitting light from a light source and tracing its propagation. Therefore, path tracing under PLT has so far been restricted to source-to-eye (forward) tracing only, which is very slow to converge. Additional difficulties hinder the practical application of PLT: (i) The new rendering primitives

are now covariance functions, in contrast to simple numeric values; (ii) The modelling of scattering of light by matter abandons the classical radiometric treatment, and is now modelled as a more analytically-involved diffraction problem. These issues make the representation of the properties of both light and matter cumbersome, and frustrate practical, efficient formulation of the light-matter interaction process.

In this paper, we propose a way to overcome these difficulties and present a complete, practical framework for PLT. After a brief review of the theory of PLT in Subsections 3.1 and 3.2, we first introduce, in Subsection 3.3, a form of the *generalized Stokes parameters* as the generalized radiance that serves as our rendering primitive. This formalism extends the well-known Stokes-Mueller calculus [Pérez and Ossikovski 2016], which has been used in computer graphics for polarization-aware radiometric rendering [Jarabo and Arellano 2017]. The generalized radiance fully quantifies the radiometric, polarimetric, and wave (i.e. coherence) properties of light, and we show that all these properties work in unison to reproduce realistic appearance and optical phenomena.

To enable simple numeric representation in a renderer, we restrict the (spatial) coherence shape of light to that of an anisotropic Gaussian, which is a reasonable approximation because light from natural (spontaneous emission) light sources admits a coherence shape well approximated by an anisotropic Gaussian [Steinberg and Yan 2021b]. Furthermore, in Subsection 3.4, we present fundamental arguments based on thermodynamics that show that the passive optical elements we want to render actually scatter light such that the light’s coherence shape remains Gaussian. Then, in Subsection 3.5 we discuss scattering of light by matter. We will show that the familiar polarimetric BSDF (i.e. a classical Mueller matrix) as well as the matter’s power spectral density (PSD) are sufficient to accurately describe a matter’s scattering characteristics. We also discuss importance sampling, even when the light’s coherence properties are unknown.

In Section 4 we build upon these theoretical contributions and present an efficient rendering algorithm for physical light transport. Our renderer leverages bi-directional path transport to render complex scenes, achieving performance that far exceeds the state-of-the-art in PLT, and approaches the speed of classical “physically-based” rendering systems. Furthermore, unlike previous work [Steinberg and Yan 2021a,b], all the materials we render are new coherence-aware materials that we develop in Subsection 4.1.

2 RELATED WORK

Physical-optics light transport. The theoretical foundations of physical light transport (PLT) are presented by Steinberg and Yan [2021a]: Recognising the fact that single-point statistical descriptions of light, like the classical radiometric radiance (which is a time-averaged value), are insufficient to derive a formalism that is consistent with electromagnetism and physical optics [Wolf 2007], they introduce a rendering theory where the second-order statistics of a collection of waves that compose a beam of light—a *wave ensemble*—supplant the classical radiance as the core quantity of interest. Subsequent work [Steinberg and Yan 2021b] discusses a computationally-tractable representation of the functions that quantify these second-order

statistics as well as a practical, flexible representation of matter. Their discussion is limited to the scalar (non polarization-aware) case. Our work builds directly upon these theoretical foundations, and focuses on deriving a practical PLT framework that is able to render complex scenes, with modern coherence-aware materials, and without relying on classical tools.

Cuyper et al. [2012]; Oh et al. [2010] have also attempted to introduce physical optics to light transport. They model light using the Wigner-Ville Spectrum (WVD) in a bid to replace the classical radiance with a more physical variant. The WVD is essentially the Fourier transform pair of the cross-spectral density (CSD), and the disadvantages compared to PLT have been discussed by Steinberg and Yan [2021a]: The WVD was envisioned [Walther 1968] as a form of a generalized radiance that accounts for the coherence of light, but remains constant on free space propagation. However, it is not non-negative (for the difficulty of satisfying all of the physical requirements of classical radiance, see Friberg [1979]), meaning that it may not act as a power density. It is also a more analytically-complicated construct than the CSD, making the derivation of light-matter interaction formulae more cumbersome.

Asymptotic, physical optics solvers. Between the geometric optics (GO) approximation—prevalent in computer graphics—and full electromagnetism (EM), a large family of *asymptotic, physical optics* (PO) formalisms have been studied in applied optics literature. These methods aim to balance between the cost prohibitiveness of full-EM wave solvers and the inaccuracy of GO ray tracers. Of the most actively researched are the “*shooting-bouncing ray*” (SBR) family of methods. These operate by solving a PO light-matter interaction problem at a surface or medium, then transition to the far field and propagate light via simply GO ray tracing. The state-of-the-art SBR methods include the study of EM scatter of stealth aircraft [Bilal et al. 2019]; simulation of ground-penetrating radar [Warren et al. 2016]; simulation of assistive-driving radars [Castro et al. 2019]; analysis of WiFi antenna characteristics [Chen et al. 2019]; and, radar imaging [Feng and Guo 2021].

These hybrid GO/PO methods are fast (due to GO ray tracing), and apply best when the GO ray tracing is used to quantify multiple scattering effects (e.g., scattering cross-section analysis). However, they are wholly inadequate for the purposes of physical-optics light transport: The intensity of a light beam, that is quantified by the GO step, is time-averaged data (field fluctuations are neglected), which discards all waveform information and statistical correlations between spectral and transverse field components. This statistical information—the optical coherence of light—cannot be recovered from a GO treatment, needs to be propagated, and plays a central role in the appearance of materials. The framework presented in this paper acts as a *global* asymptotic-PO solver, but we do not relegate to a GO treatment, instead, we treat light as an ensemble of waves throughout and propagate coherence information globally.

Wave-optics rendering. The problem of reproducing the appearance of different wave-interference effects has given rise to a wide body of work. We start with work that considers the limited optical coherence of light. Levin et al. [2013] use a box kernel to approximate the spatial area over which light remains coherent, for

the purpose of synthesis of spatially-varying BRDFs. Gaussian kernels have also been used: to render diffractive scattering effects produced by surfaces with explicit microgeometry [Falster et al. 2020; Yan et al. 2018]; and, diffractive scratches [Werner et al. 2017]. Gaussian coherence kernels and the Jones calculus have been used to measure diffraction patterns of holographic surfaces [Toisoul et al. 2018]. Such work does not aim to transport coherence information throughout a scene, and only considers *local* effects.

Other work aims to solve a full electrodynamics problem while ignoring optical coherence. Precomputing solutions to Maxwell’s equations in very simple scenes using the finite-difference time-domain (FDTD) method is employed as an alternative rendering method [Musbach et al. 2013]. These precomputation are local, material specific, and do not address the problem of transporting the coherence properties of light through the scene. Auzinger et al. [2018] adapt FDTD for the nanofabrication of materials with structural (pigment-free) colours. Using such numeric solvers for rendering is extremely constrained: in practice, they are limited to data with up to tens of thousands of FDTD cells, and work on deterministic data. 3D FDTD has been performed, but only in exceedingly simple settings (e.g. Mann and Rastogi [2020]). Describing matter deterministically, at a sub-wavelength resolution, is neither feasible nor even desirable: It is the statistics of the a surface’s or medium’s scattering characteristics that induce its appearance and optical response. This difficulty is greatly exacerbated by the partial coherence of light, which requires performing FDTD and superposing the solutions for many independent radiators (typically thousands of samples).

Additional work tackles reproduction of material appearance via different simplified formalism. This includes the rendering of thin-film interference [Belcour and Barla 2017; Kneiphof et al. 2019]; birefringent dielectrics [Steinberg 2019; Weidlich and Wilkie 2008]; iridescent, and pearlescent materials [Guillén et al. 2020]; scratches [Velinov et al. 2018] and soap bubbles [Huang et al. 2020]. Toisoul and Ghosh [2017] formulate a method for the synthesis of diffraction-aware BSDFs. Moravec [1981]; Stam [1999] are among the first to consider diffractions in rendering.

Also of relevance is work that employs surface scatter theories in computer graphics [Holzschuch and Pacanowski 2017; Steinberg and Yan 2022; Yan et al. 2018]. Such theories make use of the power spectral density (PSD) of a surface in order to approximate the first-order diffraction that arises on scattering. We will show that the PSD plays a central role in accurately and efficiently formulating physical-optics scattering with partially-coherent light.

Vectorised, polarimetric light transport. The Jones calculus and the closely-related Mueller calculus have seen usage in computer graphics for different polarization-aware rendering methodologies. Jarabo and Arellano [2017]; Weidlich and Wilkie [2008]; Wilkie et al. [2001] discuss vectorised light transport and rendering, including bi-directional path tracing. A Jones calculus-based numeric method to solve for the vectorized transport through inhomogeneous birefringent media is presented by Steinberg [2020]. Different forms and models of the pBSDF (*polarimetric BSDF*) are proposed [Duncan et al. 2003; Lei et al. 2012; Priest and Gerner 2000]. The pBSDF, in the form of a Mueller matrix, is later introduced to

computer graphics by Baek et al. [2020] for polarimetric (polarization-aware) rendering and image-based acquisition of pBSDFs, and the importance sampling of such Mueller matrix pBSDFs is discussed by Mojzík et al. [2016]. Our light transport framework builds directly upon an extended form of the Stokes-Mueller calculus, and we also employ the classical pBSDF to formulate light-matter interactions.

3 THEORY OF PHYSICAL LIGHT TRANSPORT

In this section, we discuss the theory of physical light transport. Subsections 3.1 and 3.2 serve as a succinct overview of the background on physical light transport as necessary for this paper. For a more extensive discussion of the theoretical underpinning of PLT, please refer to Steinberg and Yan [2021a]. See Fig. 2 for a broad overview of how light is transported under PLT. Our contributions begin in Subsection 3.3.

3.1 Preliminaries

First, we briefly introduce our notation and the mathematical tools used throughout this paper, see Table 1 for a list of commonly used symbols. The field of real numbers is denoted as \mathbb{R} and the complex plane as \mathbb{C} . The Cartesian n -dimensional vector space is \mathbb{R}^n and the vector notation convention is as follows: Arbitrary vectors are accented with an arrow, viz. $\vec{r} \in \mathbb{R}^n$. A vector's magnitude is written as the scalar sharing the same letter, $r = |\vec{r}|$, and the circumflexed vector \hat{r} is the respective unit vector, i.e. $\hat{r} = \frac{1}{r}\vec{r}$. The Cartesian components of \vec{r} are written as r_0, r_1, \dots, r_{n-1} , or potentially as r_x, r_y, r_z when $\vec{r} \in \mathbb{R}^3$.

The vector spaces of $n \times m$ real and complex matrices are $\mathbb{R}^{n \times m}$ and $\mathbb{C}^{n \times m}$, respectively. A square matrix $A \in \mathbb{C}^{n \times n}$ is said to be *positive-definite*, written as $A > 0$, if it is Hermitian, $A = A^\dagger$ (where the dagger denotes the conjugate transpose), and $\vec{z}^\dagger A \vec{z} > 0$ for all non-zero $\vec{z} \in \mathbb{C}^n$. We denote the unique, *positive-definite square root* of a positive-definite matrix $A \in \mathbb{C}^{n \times n}$ as $A^{1/2}$, and A is real if and only if $A^{1/2}$ is real. I denotes the identity matrix.

Given a real, positive-definite matrix $\Sigma \in \mathbb{R}^{n \times n}$, the (unnormalized) *anisotropic Gaussian function* is defined as

$$g_\Sigma(\vec{r}) \triangleq e^{-\frac{1}{2}\vec{r}^\top \Sigma^{-1} \vec{r}}. \quad (1)$$

The n -dimensional multivariate *Dirac delta* arises naturally at the limit: $\delta(\vec{r}) \triangleq \lim_{\sigma \rightarrow 0^+} \frac{1}{\sigma^n} g(\sigma^2 I)(\vec{r})$.

Functional analysis. The *Fourier transform* operator, with unitary angular-frequency kernels, is denoted as \mathcal{F} and its inverse as \mathcal{F}^{-1} . Given $f : \mathbb{R}^3 \rightarrow \mathbb{C}$, an L^2 function, its Fourier transform is

$$\mathcal{F}\{f\}(\vec{\xi}) \triangleq \left(\frac{1}{2\pi}\right)^{\frac{3}{2}} \int_{\mathbb{R}^3} d^3\vec{r} f(\vec{r}) e^{-i\vec{r} \cdot \vec{\xi}}. \quad (2)$$

The Fourier transform operator defines the convolution operator via the well-known *convolution theorem*:

$$f * h \triangleq \mathcal{F}^{-1}\{\mathcal{F}\{f\}\mathcal{F}\{h\}\}. \quad (3)$$

Reference frames and the Mueller-Stokes calculus. We write a reference frame as $[\mu] = (\hat{x}, \hat{y}, \hat{z})$, which is an ordered tuple of orthonormal vectors that define the local coordinate frame. At times, we will need orthogonal *change-of-basis* matrices, denoted as $Q_{[\mu]}$ \in

$\mathbb{R}^{3 \times 3}$, that transform from the global coordinate system to a local reference frame $[\mu]$.

The well-known classical Mueller-Stokes calculus [Pérez and Osikovski 2016] is a formal method where the polarimetric state of light is represented via 4-dimensional, real *Stokes parameters vectors*, and the action of an optical element on light is quantified via a $\mathbb{R}^{4 \times 4}$ *Mueller matrix* acting upon these Stokes vector. This calculus is sensitive to the reference frame at which the Stokes vectors are defined. Therefore, we denote a Stokes vector as $\vec{S}^{[\mu]}$ —explicitly indicating the reference frame—and a Mueller matrix as $M^{[\mu_i] \rightarrow [\mu_o]}$. This Mueller matrix expects a Stokes parameter vector defined in frame $[\mu_i]$, and after acting upon the Stokes vector it transforms it to frame $[\mu_o]$. See Appendix D for more information on how Stokes vector are rotated between frames.

We also denote the common *unpolarized, linear-horizontal and linear-vertical* polarization states (in arbitrary reference frame), as

Table 1. List of symbols and notation (with location of definition)

NOTATION AND SYMBOLS	
NOTATION	
\vec{r}	Vectors: arrow accented boldface letters
$r = \vec{r} $	Vector's magnitude: scalars sharing the same letter as the vector
$\hat{r} = \vec{r}/r$	Unit vector: hat accented boldface, also sharing letter with a vector
A	Matrices: boldface, capital letters
\mathcal{M}	Operators: capital latin letters in script typeface
OPERATORS	
$\mathcal{F}\{\cdot\}$	Fourier transform operator [Eq. (2)]
$*$	Convolution operator [Eq. (3)]
z^*	Complex conjugate of $z \in \mathbb{C}$
A^\top, A^\dagger	Transpose and conjugate transpose of A
$\langle \cdot \rangle_\omega$	Ensemble-average over same-frequency constituents [Page 5]
SYMBOLS	
$\delta(\vec{r})$	Multivariate Dirac delta
g_Σ	Anisotropic Gaussian functions with covariance $\Sigma > 0$ [Eq. (1)]
$Q_{[\mu]}$	Orthogonal change-of-basis matrix from global coordinates to local reference frame $[\mu]$
PROPERTIES OF LIGHT	
\vec{E}	Electric field [Page 5]
c	The speed of light
\vec{k}	Wavevector [Page 5]
$k = \vec{k} $	Wavenumber [Page 5]
$\lambda = \frac{2\pi}{k}$	Wavelength [Page 5]
$\omega = ck$	Wave's angular frequency [Page 5]
$C_{\alpha\beta}$	Cross-transverse spectral density (CSD) function [Eq. (7)]
S, L	Classical, scalar irradiance and radiance
\vec{S}, \vec{L}	Classical irradiance- and radiance-carrying Stokes parameters vectors
$\vec{S}^{[\mu]}, \vec{L}^{[\mu]}$	Generalized irradiance and radiance [Eq. (8) and Definition 3.1]
Θ	Spatial coherence shape matrix [Definition 3.1]
PROPERTIES OF MATTER	
p	Stationary power spectral density (PSD) [Definition 3.3.(i)]
$M^{[\mu_i] \rightarrow [\mu_o]}$	Polarimetric BSDF (pBSDF) [Definition 3.3.(ii)]
\mathcal{W}	Wave BSDF (wBSDF) operator [Eq. (16)]

well as the shorthand \vec{S}_c describing the *cross-transverse* states:

$$\vec{S}_0 \triangleq \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \quad \vec{S}_{\text{LHP}} \triangleq \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \quad \vec{S}_{\text{LVP}} \triangleq \begin{pmatrix} 1 \\ -1 \\ 0 \\ 0 \end{pmatrix}, \quad \vec{S}_c(\chi, \zeta) \triangleq \begin{pmatrix} 0 \\ 0 \\ \chi \\ \zeta \end{pmatrix}, \quad (4)$$

where the arguments χ, ζ are real numbers that play the roles of the signed *degree-of-diagonal polarization* and the signed *degree-of-circular polarization*, respectively.

3.2 Background: Wave Ensembles and Beams

The primary goal of this subsection is to derive the fundamental rendering primitive—the wave packet—in a qualitative fashion directly from classical electromagnetism. We begin with a brief discussion about the electromagnetic nature of light.

Let $\vec{E}(\vec{p}, t)$ denote the electric field throughout spacetime (non-relativistic). Any physically-realizable field can be spectrally decomposed, by writing it as a superposition of monochromatic plane waves [Zangwill 2013]:

$$\vec{E}(\vec{p}, t) = \left(\frac{1}{2\pi}\right)^{\frac{3}{2}} \int_{\mathbb{R}^3} d^3\vec{k} \vec{a}_\perp(\vec{k}) e^{i(\vec{k}\cdot\vec{p}-ckt)}, \quad (5)$$

where $k = |\vec{k}|$ and c is the speed of light. $\vec{a}_\perp(\vec{k})$ is a vector-valued function that quantifies the vector amplitude (amplitude and direction of oscillation) of a plane wave, and must fulfil $\vec{a}_\perp(\vec{k}) \cdot \vec{k} = 0$. The above is simply a spatial inverse Fourier transform of the quantity $\vec{\psi}(\vec{k}, t) = \vec{a}_\perp(\vec{k}) e^{-i\omega t}$. The function $\vec{\psi}$ is a field of plane waves, where each plane wave propagates in the direction of the *wavevector* \vec{k} and with a peak (time-independent) amplitude $\vec{a}_\perp(\vec{k})$, *angular frequency* $\omega = ck$ and *wavelength* $\lambda = \frac{2\pi}{k}$.

Consider a light source illuminating a scene and giving rise to light that propagates to a position \vec{p} . The light source can be regarded as a collection of point sources, each radiating a plane wave $\vec{\psi}(\vec{k}, t)$ in direction of \vec{p} (the wavevector \vec{k} points from the respective point source towards \vec{p} and its magnitude defines the wavelength, i.e. $|\vec{k}| = k = \frac{2\pi}{\lambda}$). The electric field of the light that arrives at \vec{p} is then the superposition of all these plane waves, as described by Eq. (5). Non coherence-aware renderers aiming to reproduce wave effects might Monte Carlo integrate Eq. (5) directly—thereby sample the statistics of the stochastic process \vec{E} —by drawing many points on the source. However, the high-frequency complex exponent in Eq. (5) frustrates this process, leading to *wave-ensemble aliasing*, see Steinberg and Yan [2021b].

Wave packets. Let that light source be now enclosed in an imaginary unit sphere (much larger than the source). Consider a small continuous surface ΔS on that sphere, subtending a solid angle of Ω from the centre of the sphere. The *spectral flux* radiating into that solid angle is the electromagnetic power of wavelength λ that flows through ΔS (up to a constant):

$$\Phi \triangleq \oint_{\Delta S} d^2\hat{k} \left| \vec{\psi}(k\hat{k}) \right|^2. \quad (6)$$

Note that the wavenumber k above is kept constant, and we only integrate over the directions $\hat{k} \in \Delta S$. Hence, the above integrates the

power of all plane waves $\vec{\psi}(\vec{k})$, restricted to wavenumber k and directions \hat{k} that pass through ΔS . We omit time as $|e^{-i\omega t}| = 1$. The collection $\{\vec{\psi}(k\hat{k})\}_{\hat{k} \in \Delta S}$ of all these monochromatic plane waves is termed a *wave packet*. A wave packet is a statistical ensemble of plane waves that constitute a solution to Maxwell's equations and form a monochromatic packet of light that propagates into the solid angle Ω . For succinctness, we will use the term *beam* synonymously with wave packet. A beam should be understood as a carrier of optical energy that radiates into a small, but positive, solid angle Ω , and hence admits a positive cross-sectional area that grows quadratically with propagation.

Being Fourier transform-pairs, \vec{E} and $\vec{\psi}$ are different representations of the same data. Nevertheless, it is considerably simpler, analytically and conceptually, to understand light not as the field \vec{E} , where complex behaviour arises, but as that collection of plane waves quantified by $\vec{\psi}$. This spatio-temporal decomposition enables the formulation of wave packets with clear propagation direction and constant wavelength. This fixes the angular frequency ω , and thereby the time-dependent phase term $e^{-i\omega t}$ becomes identical for all plane wave constituents of the packet. As this propagator is the only time-dependent quantity, working piecewise with packets of monochromatic plane waves elicits a few key advantages: (i) in place of a time-dependent formalism, a wavelength-dependent formalism arises; and, (ii) as time can be ignored, temporal coherence does not need to be considered (monochromatic waves are perfectly temporally coherent). Tracing fixed-wavelength packets, instead of tracking the evolution of a time-dependent field, is the natural course of action in rendering. Temporal coherence effects are then reintroduced, at no additional effort, by sampling the entire temporal spectrum (i.e. wavelengths) of light. See Fig. 2 an illustrative intuitive explanation of how wave packets are formed and their use.

A wave packet, as formalised, forms an *invariant optical field* [Gutiérrez-Vega et al. 2000], meaning that the beam may diffract on free-space propagation only when partially blocked by an obstacle, otherwise it propagates solely into the solid angle Ω , as desired. When rendering with optical frequencies, we keep A and Ω small, therefore in this work we ignore partial occlusions and the free-space diffractions that would arise.

Optical coherence. The discussion thus far explains how to spectrally decompose an arbitrary electric field. This allows us to form wave packets—monochromatic beams of light—and we will now discuss and formalise the optical coherence of these beam.

When dealing with the partially-coherent light that we observe on a daily basis, the underlying electric field \vec{E} admits seemingly random fluctuations. A deterministic study of such fields and packets is neither feasible nor desirable: our sensors (the eye or a camera) do not observe the rapid individual oscillations of the electromagnetic field, but only time-averaged values, therefore observable optical phenomena are statistical in nature. Hence, we treat \vec{E} , as well as any derived wave packet, as a random function (stochastic process). PLT rendering primitives then aim to quantify the statistics of a wave packet. These statistics are devoid of high-frequency optical phase, thus are amenable to analytic tools, and are free of wave-ensemble aliasing. In contrast to radiometric quantities, we




Fig. 2. Physical light transport (PLT): (a) A light source emits electromagnetic radiation. The underlying fields of this radiation fluctuate rapidly, are disorganized and admit a considerable degree of random behaviour (not visualized). (b) We spectrally decompose (spatially and temporally) these fields, viz. Eq. (5), giving rise to a collection of plane wave constituents that compose the underlying electric field. (c) We select a subset of these plane waves, such that all plane waves in this subset are of the same frequency and radiate into a small solid angle Ω . This collection formally defines a *wave packet*—a *beam* of light radiated from projected source area A and with small divergence Ω . Such a construction of a wave packet is always possible for arbitrary physically-realizable electric fields. The spatial statistical correlation between the plane wave constituents of this wave packet is known as the *optical coherence* of light, and quantifies the ability of the beam to interfere with itself. (d) As A and Ω are kept small, the cross-sectional area of this beam is small as well (w.r.t. scene geometrical details). Therefore, to propagate the beam we employ standard ray tracing through the scene. The polarimetric and coherence properties of the beam are propagated as well. (e) Hence, once a matter is encountered, we have the necessary information to compute rigorous physical optics light-matter interaction, which is able to reproduce wave-interference and diffraction phenomena. (f) This interaction gives rise to a scattered field, and, as before, we may spectrally decompose this field and select a scattered beam that propagates into an arbitrary direction.

quantify *two-point* spatial statistics of the packet, thereby capturing sufficient information to remain consistent with the theory of electromagnetism.

We now formally restrict the stochastic process \vec{E} to be a wide-sense stationary process (implying steady-state rendering), and consider only its far-field statistics (i.e., far from the field's source). Since the statistics converge rapidly to the far-field statistics over the distance of several wavelengths [Agarwal et al. 2004; Charnotskii 2019], this is an excellent approximation. Let the wave packet's wavenumber be k . The plane wave constituents of this wave packet propagate with directions \hat{k} that have a small spread (i.e., all propagate into a small solid angle Ω , as before), and we denote the mean direction of propagation of the entire packet as the unit vector \hat{r} . We select a *transverse basis* $\{\hat{x}, \hat{y}\}$, such that it completes an orthonormal basis $\{\hat{x}, \hat{y}, \hat{r}\}$ of \mathbb{R}^3 , but otherwise may be arbitrarily chosen. This orthonormal basis is termed the *local frame* of the wave packet.

We may now define the scalar fields $E_x(\vec{p}, t) = \vec{E}(\vec{p}, t) \cdot \hat{x}$ and $E_y(\vec{p}, t) = \vec{E}(\vec{p}, t) \cdot \hat{y}$, which are the *transverse components* of the electric field. The core quantities of optical coherence theory are the four *cross-spectral density* (CSD) functions between these transverse components [Wolf 2007], viz.

$$C_{\alpha\beta}(\vec{p}, \vec{\xi}; \omega) \triangleq \left\langle E_{\alpha}\left(\vec{p} + \frac{1}{2}\vec{\xi}\right) E_{\beta}\left(\vec{p} - \frac{1}{2}\vec{\xi}\right)^{\star} \right\rangle_{\omega}, \quad (7)$$

with $\alpha, \beta \in \{x, y\}$. The operator $\langle \cdot \rangle_{\omega}$ denotes ensemble-averaging, where the subscript ω makes explicit the fact that we average over same-frequency plane-wave constituents of the wave packet. The CSD functions have units of spectral irradiance (spectral flux per

area) and are the second-order correlations that quantify the ability of the wave packet to superpose and interfere with a spatially-shifted version of itself. That is, they describe the statistical similarity between the waveforms of angular frequency ω that arrive at the points $\vec{p} \pm \frac{1}{2}\vec{\xi}$. The point \vec{p} is an arbitrary position in space, and $\vec{\xi}$ serves as the difference vector between the points. If these waveforms are similar, then constructive or destructive wave interference takes place and observable diffractive optical phenomena may arise. Otherwise, both constructive and destructive interference take place with equal probability and, on average, cancel out. It has been shown that a two-point formalism, as above, is necessary for consistency with electromagnetism [Wolf 2007]. We omit time t from the transverse components $E_{x,y}$ in Eq. (7), as ensemble averaging produces a time-independent quantity. It should be remembered that the definitions of the CSDs depend, implicitly, on the local transverse basis $\{\hat{x}, \hat{y}\}$.

The generalized Stokes parameters. The CSD functions fully describe the second-order statistics of the wave packet. We will now rewrite these quantities in the more convenient form of the *generalized Stokes parameters* (gSP), as introduced by [Korotkova and Wolf 2005]. A gSP vector serves as an extension of the classical, well-known Stokes parameters vectors to a two-point formalism, as follows:

$$\vec{S}^{|\mu|}(\vec{p}, \vec{\xi}; \omega) \triangleq \begin{pmatrix} C_{xx} + C_{yy} \\ C_{xx} - C_{yy} \\ C_{yx} + C_{xy} \\ i(C_{yx} - C_{xy}) \end{pmatrix}, \quad (8)$$

where $[\boldsymbol{\mu}] = (\hat{x}, \hat{y}, \hat{r})$ is the reference frame with respect to which the CSDs are defined and, for brevity, we omit the parameter lists to the CSDs $C_{\alpha\beta}$. The double arrow accent is used to denote quantities that generalize classical Stokes vectors (i.e., when evaluated $\vec{\mathcal{S}}^{[\boldsymbol{\mu}]}$ reduces to a classical Stokes vector, quantifying radiometric and polarimetric properties of light).

The gSP vector is a powerful representation of the wave packet's statistics, which serves to unify the theories of coherence and polarization into one formalism. This representation is not only convenient, but also both coherence and polarization dictate the observable properties of light, its propagation, and its interaction with matter. The first element of $\vec{\mathcal{S}}^{[\boldsymbol{\mu}]}$, i.e. $\mathcal{S}_0^{[\boldsymbol{\mu}]}$, quantifies the intensity, while the rest of the elements describe the polarization state [Korotkova 2017]. See Appendix A for explicit relations. Note that for the gSP to describe a physically-realizable wave packet, the following relation must be enforced:

$$\mathcal{S}_0^{[\boldsymbol{\mu}]}(\vec{p}, 0)^2 \geq \mathcal{S}_1^{[\boldsymbol{\mu}]}(\vec{p}, 0)^2 + \mathcal{S}_2^{[\boldsymbol{\mu}]}(\vec{p}, 0)^2 + \mathcal{S}_3^{[\boldsymbol{\mu}]}(\vec{p}, 0)^2, \quad (9)$$

where $\mathcal{S}_n^{[\boldsymbol{\mu}]}$ is the n -th element of the vector $\vec{\mathcal{S}}^{[\boldsymbol{\mu}]}$.

3.3 Generalized Radiance and Irradiance

Consider a beam, i.e., a collection of plane waves of wavenumber $k = 2\pi/\lambda$, radiating into a small solid angle Ω subtended by the surface ΔS , as before. Let the mean direction of propagation of the beam be \hat{r} and mean distance of propagation be r . The projected area (in direction \hat{r}) of the source producing the radiation is denoted A . That source may be primary—a radiating light source—or secondary—a scattering optical element. Then the cross-sectional area of the beam grows as $\sim (A + r^2\Omega)$ on propagation.

The wave packet's underlying electric field varies rapidly, spatially and temporally. On the other hand, the angular far-field statistics of that field vary slowly on ΔS , and thus can be assumed to be angularly wide-sense stationary. Formally, the gSP is a fast function of $\vec{\xi}$ but a slow function of \hat{p} . That is, the gSP quantifies the mutual coherence between the space points $\vec{p} \pm \frac{1}{2}\vec{\xi}$. If we were to shift these points by the same spatial offset (leaving $\vec{\xi}$ intact), then the gSP would change slowly, however if we were to change the difference vector $\vec{\xi}$ between these points, then the gSP would change far more rapidly. This is because the distribution of irradiance on the beam's cross section varies much slower than the correlations between the waveforms [Wolf 2007].

Hence, we may drop the directional dependence on \hat{p} and write $\vec{\mathcal{S}}^{[\boldsymbol{\mu}]}(\vec{p}, \vec{\xi}) \approx \vec{\mathcal{S}}^{[\boldsymbol{\mu}]}(\vec{r}, \vec{\xi})$. Note, the far-field condition mandates that $\hat{p} \approx \hat{r}$ and $|\vec{\xi}| \ll r$. Intuitively, we assume that the statistical properties of the beam are directionally invariant on ΔS . This is an excellent approximation for far-field statistics (small Ω) [Steinberg and Yan 2021b].

In line with the discussion above, we define the differential quantity as the *generalized radiance*

$$\begin{aligned} \vec{\mathcal{L}}^{[\boldsymbol{\mu}]}(\vec{r}, \vec{\xi}; \omega) &\triangleq -\frac{\partial^2}{\partial A \partial \Omega} \int_{r\Delta S} d^2\vec{p} \vec{\mathcal{S}}^{[\boldsymbol{\mu}]}(\vec{p}, \vec{\xi}; \omega) \\ &\approx r^2 \frac{\partial}{\partial A} \vec{\mathcal{S}}^{[\boldsymbol{\mu}]}(\vec{r}, \vec{\xi}; \omega). \end{aligned} \quad (10)$$

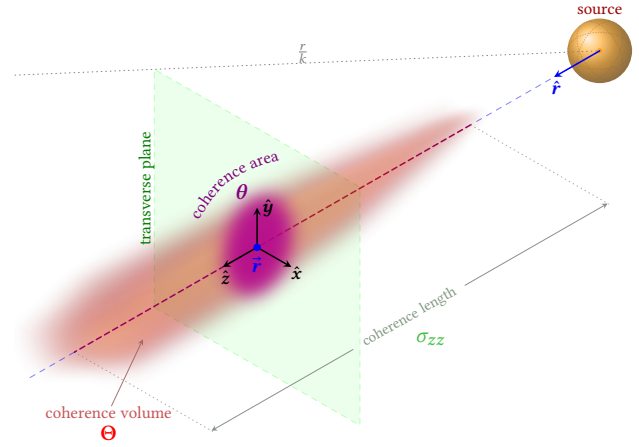


Fig. 3. Light with anisotropic Gaussian coherence is emitted by a source and propagates in direction \hat{r} to the point \vec{r} . The beam's local frame is spanned by \hat{x} , \hat{y} and $\hat{z} = \hat{r}$. The *coherence volume*—the spatial volume over which optical coherence is maintained—is depicted by the red volume and is quantified by the shape matrix Θ . The *coherence length*, in the direction of propagation, is dictated by the standard deviation σ_{zz} . The plane perpendicular to the direction of propagation is termed the *transverse plane*, and its intersection with the coherence volume is the *coherence area* (violet area). The shape of this area is dictated by the leading principal minor θ of the shape matrix and it grows linearly as $(\frac{r}{k})^2$ on free-space propagation. Figure adapted from Steinberg and Yan [2021b].

A brief dimensional analysis shows that $\vec{\mathcal{L}}^{[\boldsymbol{\mu}]}$ admits units of spectral radiance: the elements of the gSP vector are CSD functions with units of spectral irradiance. The integral over the surface $r\Delta S$ then has units of power (W) per frequency (Hz), i.e. spectral flux, and the derivative is per projected source area and solid angle. Steinberg and Yan [2021a] define a similar quantity, denoted the *radiance-carrying CSD function*. Clearly, when evaluated at $\vec{\xi} = 0$, $\vec{\mathcal{L}}^{[\boldsymbol{\mu}]}$ is simply the classical radiometric radiance. Therefore, just as the gSP generalizes the classical Stokes parameters to a two-point formalism, $\vec{\mathcal{L}}^{[\boldsymbol{\mu}]}$ generalizes the classical radiance to a two-point formalism. In similar manner, the (irradiance-carrying) gSP vector $\vec{\mathcal{S}}^{[\boldsymbol{\mu}]}$ is termed the *generalized irradiance*.

The properties that follow highlight that $\vec{\mathcal{L}}^{[\boldsymbol{\mu}]}$ can indeed be understood as a generalized radiance.

Property 3.1 (CONSERVATION OF RADIANCE ON FREE-SPACE PROPAGATION). *The spectral radiance quantified by $\vec{\mathcal{L}}^{[\boldsymbol{\mu}]}$ remains constant on far-field propagation.*

PROOF. The expression $r^2 \langle E_\alpha(\vec{r}) E_\beta(\vec{r})^* \rangle_\omega$, for $\alpha, \beta \in \{x, y\}$, remains constant in the far field on free-space propagation [Steinberg and Yan 2021a] (the the inverse square law for spherical waves), therefore $r^2 \vec{\mathcal{S}}^{[\boldsymbol{\mu}]}(\vec{r}, 0)$ also is a constant function of r . Then, the spectral radiance $L = \mathcal{L}_0^{[\boldsymbol{\mu}]}(\vec{r}, 0; \omega)$ (Eq. (37)) remains constant on free-space propagation. \square

Property 3.2 (LINEARITY UNDER SUPERPOSITION). *Given a pair of beams with generalized radiances $\vec{\mathcal{L}}_1^{[\boldsymbol{\mu}]}$ and $\vec{\mathcal{L}}_2^{[\boldsymbol{\mu}]}$. If these packets are*

mutually-incoherent, then the superposition of these wave packets is linear, viz. $\vec{\mathcal{L}}_1^{|\mu|} + \vec{\mathcal{L}}_2^{|\mu|}$.

PROOF. Let \vec{E}_1, \vec{E}_2 be the underlying electric fields of $\vec{\mathcal{L}}_1^{|\mu|}$ and $\vec{\mathcal{L}}_2^{|\mu|}$, respectively. Assume these fields are decomposed under the same transverse basis. Let $C_{1,\alpha\beta}, C_{2,\alpha\beta}$ be the CSDs of these fields and $C_{\alpha\beta}$ be the CSDs of the superposition $\vec{E} = \vec{E}_1 + \vec{E}_2$, for $\alpha, \beta \in \{x, y\}$. Let $\vec{r}_{1,2} = \vec{r} \pm \frac{1}{2} \vec{\xi}$ be a pair of points, then, mutual incoherence implies that $\langle E_{1,\alpha}(\vec{r}_1) E_{2,\beta}(\vec{r}_2) \rangle_\omega = 0$, therefore

$$\begin{aligned} C_{\alpha\beta}(\vec{r}, \vec{\xi}) &= \langle [E_{1,\alpha}(\vec{r}_1) + E_{2,\alpha}(\vec{r}_1)] [E_{1,\beta}(\vec{r}_2) + E_{2,\beta}(\vec{r}_2)] \rangle_\omega \\ &= C_{1,\alpha\beta}(\vec{r}, \vec{\xi}) + C_{2,\alpha\beta}(\vec{r}, \vec{\xi}). \end{aligned}$$

Hence, $\vec{\mathcal{L}}_1^{|\mu|} + \vec{\mathcal{L}}_2^{|\mu|}$ is the wave packet of the superposition. \square

Analytic representation of a beam. The quantity $\vec{\mathcal{L}}^{|\mu|}$, defined in Eq. (10), serves as a complete description of a beam of light. However, its elements are functions, in contrast to the numeric values admitted by classical radiometric quantities. In order to represent a wide family of CSD functions, Steinberg and Yan [2021b] propose expanding the CSD under the anisotropic Hermite-Gauss functional basis. This elicits multiple benefits: (i) representation reduces to tracking numeric coefficients; and (ii) the Hermite-Gauss functions admit useful analytic properties which greatly simplify the diffraction formulae. We proceed in a similar fashion, but restrict the expansion to the 0th-order coefficient only, i.e. a multivariate anisotropic Gaussian.

Definition 3.1 (GENERALIZED RADIANCE WITH GAUSSIAN COHERENCE). The radiance-carrying generalized Stokes parameters of a beam with anisotropic Gaussian spatial coherence is

$$\begin{aligned} \vec{\mathcal{L}}^{|\mu|}(r, \vec{\xi}; \omega) &\triangleq e^{ik\xi_z} \left[L_x g^{\Theta_x}(\vec{\xi}) \vec{S}_{\text{LHP}} + L_y g^{\Theta_y}(\vec{\xi}) \vec{S}_{\text{LVP}} \right. \\ &\quad \left. + \sqrt{L_x L_y} g^{\Theta_{1/2}}(\vec{\xi}) \vec{S}_c(\chi, \varsigma) \right], \end{aligned} \quad (11)$$

where r is the mean propagation distance, g is the (unnormalized) anisotropic Gaussian, defined in Eq. (1), and the common polarization states are defined in Eq. (4). The spectral radiances carried by each transverse component are $L_{x,y}$, the signed *degree-of-diagonal polarization* and the signed *degree-of-circular polarization* are χ and ς , respectively. See Appendix B for a brief derivation. The local frame is $[\mu] = (\hat{x}, \hat{y}, \hat{k})$, where \hat{x}, \hat{y} span the transverse plane and \hat{k} is the (mean) direction of propagation of the packet (the positive z direction in the local frame). The coherence properties of the wave packet are quantified by the *shape matrices*:

$$\Theta_{x,y,1/2} \triangleq \begin{pmatrix} \frac{r^2}{k^2} \theta_{x,y,1/2} & \\ & \sigma_{zz}^2 \end{pmatrix} \quad \text{with} \quad \theta_{1/2} \triangleq \frac{\theta_x + \theta_y}{2} \quad (12)$$

and where $\theta_{x,y}$ are positive-definite, real 2×2 matrices.

The total spectral radiance carried by the wave packet is $L = L_x + L_y$, while $L_x - L_y, \varsigma$ and χ partition the energy into different states of polarization, see Appendix A. These quantities might be wavelength dependent, but are otherwise constant—a restatement of the fact that radiance remains constant on propagation (Property 3.1).

These classical quantities are endowed with the transverse shape matrices $\Theta_{x,y}$ —tensors that serve to express the wave properties of the packet and quantify the first-order geometric properties of coherence [Steinberg and Yan 2021b]: the coherence area, shape and length. The leading principal minors $\theta_{x,y}$ are positive-definite 2×2 real matrices that fully describe the coherence area, while $\sigma_{zz} > 0$ is the standard variance of spatial coherence in the direction of propagation, i.e. the *coherence length* (see Fig. 3). $\theta_{x,y}$ and σ_{zz} also remain constant on propagation. The block-diagonal structure of the shape matrices reflects the quasi-homogenous beam assumption (the statistics are constants across the cross section of the beam, as quantified by Eq. (10)) that implies symmetry with respect to the transverse plane. Perfectly spatially-coherent and incoherent light are the limiting cases where the singular values of Θ tend to infinity and 0, respectively.

Discussion. Definition 3.1 defines an expression for an arbitrary partially-polarized, partially-coherent wave packet with anisotropic Gaussian coherence. In general, light does not admit anisotropic Gaussian spatial coherence. However, radiation from large (much greater than wavelength) light sources typically can be approximated well via a Gaussian coherence: This can be deduced by applying the Central Limit Theorem to a wave packet that is composed of very many statistically-independent contributions, each arising from a different radiator in the source; alternatively, the far-field spatial coherence function is analytically related to the source geometry via a Fourier transform, and for uniform isotropically-radiating spherical, cylindrical, or rectangular sources the coherence function takes the form of a sinc function or a Bessel function of the first kind [Steinberg and Yan 2021a], all of which admit good Gaussian fits (see Steinberg and Yan [2021b, Fig. 4]). Similar reasoning also applies to sources with a Gaussian distribution of radiating intensity (good approximation for thermal and gas-discharge sources). In the following subsections, we will discuss scattering and interaction with matter, and present additional fundamental arguments that suggest that spatial coherence retains its anisotropic Gaussian shape on scattering as well.

Hence, for general-purpose physical-optics rendering and computational optics applications, an assumption of anisotropic Gaussian spatial coherence is a decent approximation: we are able to represent light of arbitrary coherence area and length, and any states of polarization. This representation of a wave packet reduces to a few numeric constants: in addition to the common radiometric and polarimetric quantities $L_{x,y}, \chi, \varsigma$, we simply need to keep track of the constant shape matrices $\theta_{x,y}$ and σ_{zz} . Furthermore, as the anisotropic Gaussian is simply the 0th-order Hermite-Gauss function, we may take advantage of a simplified form of the light-matter interaction theory developed by Steinberg and Yan [2021b]. Note, only the spatial coherence is restricted to an anisotropic Gaussian. No assumptions are made regarding the temporal coherence (with the exception of wide-sense stationarity) and our rendering framework works with light of arbitrary spectral content.

For the rest of the paper, we assume that all generalized radiance and irradiance vectors (radiance- and irradiance-carrying gSP vectors) admit anisotropic Gaussian spatial coherence, as defined in Definition 3.1.




Fig. 4. Conservation of optical entropy on scattering: A beam is incident upon an optical element (a scattering material). The beam’s electric field (real part) on the xz -plane of the local frame (as defined in Fig. 3) is visualized: yellow and red colours for positive values, and cyan and blue for negative values. The beam scatters and gives rise to scattered beams. Any physically-realizable electric field can be written as a sum of its *principal modes*—coherent, but mutually-incoherent, fields—and the principal modes of the incident and scattered beams are visualised (amplitude and phase on the transverse plane) by the bottom left and bottom right, respectively, sets of insets. Due to reasons discussed in Subsection 3.4, an optical element may be understood as a deterministic finite-state automaton (DFA) that acts upon these modes. Hence, while the electric fields and modal decomposition of the incident and scattered beam in general differ, an important realization is that the *modal density* (and thus, entropy) remains invariant on scattering, and therefore the light’s coherence properties are conserved as well. Note that the shown modal decomposition is a simple example, real partially-coherent light will admit a far greater modal density.

3.4 Entropy and Scattering Conservation Laws

In this subsection we derive important conservation laws that set fundamental constraints on the how the coherence properties of light are transmitted (scattered) through an optical element. We present a general argument that is based on thermodynamic considerations. The motivation is two-fold: (i) show that if the incident light admits Gaussian spatial coherence, then, for a wide class of optical elements, the scattered light is also Gaussian in coherence; and (ii) derive constraints on the transformation of spatial coherence on scattering.

Consider a (monochromatic) partially-coherent beam of light that is incident upon an optical element (e.g., a surface or a medium) and is scattered by it. Due to the linearity of Maxwell’s equations, we may always write the underlying electric field E as a (finite) superposition of principal modes $E = \sum_j E_j$, such that each mode is a fully-coherent solution to Maxwell’s equations, but the modes are pair-wise mutually-incoherent. Such a (finite) modal decomposition depends on the optical element, but is always possible [Fabre and Treps 2020]. Assume that the incident beam consists of N such modes, which are then scattered by the optical element into M modes, that constitute the scattered beam.

We may understand these modes as information carriers: Each principal mode may contain, at most, constant bits of entropy (as it

is fully coherent over spacetime). The modes are pair-wise incoherent, hence they act as independent information carriers. The optical element is then a form of a computational device that generates M output modes based on the N input modes. Then, by the second law of thermodynamics, we immediately state that $N \leq M$: entropy may only increase, as this computational device has no means of concentrating N input channels into strictly less than N output channels (without a priori knowledge of the input’s information content).

When is entropy not conserved, i.e. $N < M$? Injection of entropy requires (i) information erasure, i.e. absorption; or (ii) active generation of entropy, via an external power source. Assume that the optical element does not have access to an external energy source. Hence, our computation device must be stateless and memoryless: it may not draw random data or access memory (as writing to memory requires information erasure, which by Landauer’s principle emits light or admit time-varying scattering characteristics. Then, such a passive optical element strongly-couples the input and *optical* output modes [Zhang et al. 2019]. Formally,

$$N \leq M = M_{\text{opt}} + \tilde{M} \leq N + \tilde{M}, \quad (13)$$

where \tilde{M} is count of *non-optical modes* (most often, such modes arise via absorption, which produces heat that is then radiated outside the visible spectrum). Strict $M_{\text{opt}} < N$ is only possible when the optical element perfectly absorbs one or more of the input modes. As $M_{\text{opt}} \leq N$, we conclude that *a passive optical element may not produce a less coherent scattered beam* (a consequence of the fact that real physical work is needed in order to inject entropy). The optical element may produce a more coherent scattered beam, but only at the price of absorbed modes.

Although such mode-specific absorption and scattering is theoretically possible to achieve, it is not easy to engineer: the computational device—which is stateless and memoryless, as we concluded, and hence acts as a deterministic finite-state automaton (DFA)—would need to recognize distinct input modes in order to perform different operations on these modes. In the realm of photonics, such careful engineering of mode-specific resonance is a very active area of research and is known as *material structuring* [Chao et al. 2021], and usually requires wavefront shaping of the incident beam [Yu et al. 2020], i.e. targeting a particular set of a priori chosen input modes (violating the assumption that the computational device encodes no a priori knowledge of the input).

As an aside, we note that conservation of (optical) entropy, i.e. $N = M_{\text{opt}}$, implies that the scattering process is (theoretically) reversible. For example, a simple diffuser does not increase entropy: with precise knowledge of the diffuser’s scattering characteristics, we may design a hologram that reverses the process by reconstructing the N input modes from the M_{opt} output modes. Some materials, like structured periodic gratings that disperse light, admit strong frequency-dependent, but not mode-dependent, response. Indeed, essentially all optical interactions of light with matter are frequency dependent, but frequency plays no role in the above analysis: recall that our framework deals with frequencies piece-wise (by spectrally decomposing light) and all the principal modes are monochromatic of a shared frequency. Other materials may couple the processes of absorption and optical emission (via fluorescence

or phosphorescence), however the emission is incoherent with the scattered modes (and usually of a different frequency) and hence has no effect on the scattered modes.

Under the assumption that the optical element absorbs in a *mode-agnostic* fashion, absorption behaves as a random process that is independent of the incident modes and its expected value is the total absorptivity of the optical element. In Appendix C we formalize this argument, and derive the following important scattering law:

THEOREM 3.2 (LOCAL CONSERVATION OF COHERENCE AREA). *Let a beam with cross-spectral density (CSD) function $C(\vec{p}, \vec{\xi})$ be incident upon a passive, linear, mode-agnostic optical element. This CSD is defined on the normalized transverse plane of the incident beam. Then, the CSD function on the normalized transverse plane of the scattered beam is $|\zeta|^2 C$, i.e. the CSD scaled by a coupling coefficient $\zeta \in \mathbb{C}$ (the ratio between the scattered and incident amplitudes).*

The normalized transverse plane is defined as the transverse plane transformed such that the beam’s cross section forms a unit disk (because the incident and scattered modes are defined in their own arbitrary spatial reference frames). In the far-field, such transformations are always linear, implying that the beam cross section always forms an ellipse. If the incident beam’s spatial coherence is of the form of an anisotropic Gaussian, then Theorem 3.2 implies that:

- (1) The scattered beam must also admit anisotropic Gaussian coherence.
- (2) The ratio between the coherence area and the beam’s cross-sectional area is a conserved quantity on propagation and scattering (by linear, passive, mode-agnostic optical elements).

If $\frac{r^2}{k^2} \theta$ is the leading 2×2 principal minor of the shape matrix of a beam with anisotropic Gaussian spatial coherence (as defined in Definition 3.1), then the *spatial coherence area* $|\theta|$ is a *conserved quantity* (because the beam cross section grows as r^2).

We have invoked arguments based on entropy to show that the spatial coherence of a beam relative to its cross-section area is a conserved quantity. The fact that the ratio between the coherence patch area and the beam cross section is a conserved quantity has also been noted before in very limited settings [Garcia-Sucerquia and Medina E 2003; Medina and Pozzi 1990]. However, to our knowledge, the general result we presented here has not been considered before and is a fundamental result in the theory of coherence transport. See Fig. 4 for an illustration.

Failure cases. As Gaussianity can be understood as a maximal-entropy model (indeed, this is a restatement of the Central Limit Theorem under Lindeberg’s condition [Linnik 1959]), our reasoning does not necessarily hold if the incident beam is not Gaussian in spatial coherence. This is because coherence may now transform both in structure and area, while still conserving entropy. However, as we discussed qualitatively and quantitatively in Subsection 3.3, light emitted by spontaneous emission sources is very well approximated by Gaussian spatial coherence, and, as stated by Theorem 3.2, that remains the case on scattering by essentially all optical elements that are of interest to use.

Injection of optical entropy is often trivially handled by partitioning the exitant energy into passively-scattered beams, emitted

beams (produced by a power source), and potentially re-emitted beams (e.g., fluorescence). As these classes of exitant beams are essentially always mutually-incoherent, this approach is accurate. Rarely, subtleties may arise: for example, under temporal entropy injection. Consider a mechanically-rotated etched glass plate (a diffuser). Once scattered by this diffuser, the observed, time-averaged light is in fact the superposition of contributions from distinct spatial regions on the diffuser. If the diffuser is rotated rapidly enough, then multiple statistically-independent regions contribute to the scattered beam, and decoherence takes place. That is, mechanical energy is transformed into optical entropy. In this work we ignore such materials. Any decoherence requires an injection of entropy, and future work could devise means to integrate over the source of entropy.

On the other hand, rare structured materials may produce a more coherent beam than the incident, as mentioned. In this work, we assume that the materials that we wish to render are entirely devoid of such structuring. We do note, however, that due to the hard physical limit set by Eq. (13) for passive (non-emitting, non-amplifying) optical elements, an increase in coherence would always come at a cost: to produce a more coherent beam, mode-dependant annihilation must take place. In other words, to double the coherence area, at least half the energy needs to be absorbed.

3.5 Diffraction by Matter

To quantify the matter’s response to incident light, we build upon the framework developed by Steinberg and Yan [2021b]. However, our restriction to a Gaussian spatial coherence as well as Theorem 3.2 serve to considerably simplify the formulae. As a matter of fact, we will show that classical (“non wave-aware”) quantities are sufficient to discuss and formalise scattering by matter.

Consider a spatial scattering region \mathbf{P} , centred around the origin. The matter’s scattering characteristics that emerge in that region are treated as a spatial stochastic process, and we restrict that process to the class of locally-stationary processes, as proposed by Steinberg and Yan [2021b]. A comprehensive introduction to the relevant aspects of scattering by locally-stationary matter as well as the related generalized Mueller calculus is presented in our supplemental material.

Let a beam, with generalized irradiance $\vec{\mathcal{S}}^{[\mu_i]}$ and anisotropic Gaussian spatial coherence, be incident upon a scatterer (a passive optical element) that occupies the region $\mathbf{P} \subset \mathbb{R}^3$. The shape matrices, irradiance carried by each transverse component and the degrees-of-diagonal and circular polarization of the incident beam are $\Theta_{x,y,z}^{(i)}$, $S_{x,y}^{(i)}$ and $\chi^{(i)}$, $\zeta^{(i)}$, respectively. Let \hat{s} and \hat{r} be the (mean) incident and scattering directions. We denote the incident and scattering local reference frames as $[\mu_i]$ and $[\mu_o]$, respectively.

Definition 3.3 (LOCALLY-STATIONARY MATTER). The quantities that follow describe the locally-stationary matter’s scattering characteristics:

- (i) The *power spectral density* (PSD) $p(\vec{\zeta}; \omega)$, where $\vec{\zeta}$ is spatial frequency, describes statistical perturbations in the scattering characteristics. While, in general, the PSD might be a (complex-valued) Mueller matrix, for simplicity we restrict it

to a scalar function. Generalizing to a Mueller matrix is not difficult, see Steinberg and Yan [2021b, Supplemental].

- (ii) The *polarimetric BSDF* (pBSDF) $\mathbf{M}^{|\mu_i| \rightarrow |\mu_o|}(\hat{s}, \hat{r}; \omega)$, is the real, average (over the entire matter) Mueller matrix. The pBSDF transforms from the incident to the scattering reference frames, and hence its definition depends on these frames.

The PSD is a well known and extensively studied quantity in optical scattering literature: It is convenient to formulate statistical models, like the K-correlation surface model [Stover 2012], via the PSD of the spatial features. Indeed, it is typically the PSD of a surface—in contrast to an explicit height field—that is measured from a physical sample [Siewert et al. 2008]. Likewise, the pBSDF has also been well studied and used in computer graphics, see Section 2.

The wave BSDF. Denote the shorthands

$$\Xi_{x,y,1/2}^{(i)} \triangleq \mathbf{Q}_{|\mu_i|} \Theta_{x,y,1/2}^{(i)} \mathbf{Q}_{|\mu_i|}^\top \quad \text{and} \quad \vec{h} \triangleq k(\hat{r} + \hat{s}), \quad (14)$$

i.e. Ξ are the incident shape matrices transformed to the global (matter's) reference frame, and define the *diffraction operator*:

$$\mathcal{D}\{\Sigma\} \triangleq \left(\frac{|\Sigma|}{8\pi^3} \right)^{1/2} \left(p * g \Sigma^{-1} \right) (\vec{h}), \quad (15)$$

which is a three-dimensional convolution between the PSD and an anisotropic Gaussian with some positive-definite matrix Σ , evaluated at the wave-optics analogue of the “half vector” \vec{h} . In Section S2 in our supplemental material we derive the following central result: The scattering *wave BSDF* (wBSDF) for radiation that undergoes scattering is defined as

$$\begin{aligned} \mathcal{W}_{\text{scat}} \left\{ \vec{\mathcal{S}}^{|\mu_i|} \right\} &\triangleq \cos \vartheta_o \mathbf{M}^{|\mu_i| \rightarrow |\mu_o|} \\ &\times \left[S_x^{(i)} \mathcal{D} \left\{ \Xi_x^{(i)} \right\} \vec{\mathcal{S}}_{\text{LHP}} + S_y^{(i)} \mathcal{D} \left\{ \Xi_y^{(i)} \right\} \vec{\mathcal{S}}_{\text{LVP}} \right. \\ &\left. + \sqrt{S_x^{(i)} S_y^{(i)}} \mathcal{D} \left\{ \Xi_{1/2}^{(i)} \right\} \vec{\mathcal{S}}_c \left(\chi^{(i)}, \varsigma^{(i)} \right) \right], \quad (16) \end{aligned}$$

where ϑ_o is the inclination angle and the shorthands $\vec{\mathcal{S}}_{\text{LHP}}$, $\vec{\mathcal{S}}_{\text{LVP}}$, $\vec{\mathcal{S}}_c$ were defined in Eq. (4). In similar manner to the classical BSDF, the wBSDF acts upon the incident generalized irradiance to produce a scattered radiance (see Fig. 5). For each polarization state, coherence-aware wave-interference effects arise via the diffractions process formalised by the diffraction operator \mathcal{D} , that acts upon the respective shape matrix. The produced radiance-carrying classical Stokes parameters vector $\vec{L}^{|\mu_o|} = \mathcal{W} \left\{ \vec{\mathcal{S}}^{|\mu_i|} \right\}$ fully quantifies the radiometric and polarimetric properties of the scattered radiation via the typical Stokes relations, viz.:

$$L_{x,y}^{(o)} = \frac{L_0^{|\mu_o|} \pm L_1^{|\mu_o|}}{2}, \quad \chi^{(o)} = \frac{L_2^{|\mu_o|}}{\sqrt{L_x^{(o)} L_y^{(o)}}}, \quad \varsigma^{(o)} = \frac{L_3^{|\mu_o|}}{\sqrt{L_x^{(o)} L_y^{(o)}}}.$$

Observe that the light's radiometric, polarimetric and wave properties all play a role in the process of scattering of light by matter. The interference process that is formalised by the scattering wBSDF gives rise to *coherence-induced spectral and polarimetric changes*. These effects have been previously observed and studied in the optical literature [Dogariu and Wolf 1998; Wang and Zhao 2010] and our rendering framework is able to capture such phenomena.

The shape matrices. The wBSDF produces a classical quantity, and we now turn our attention to the scattered shape matrices, specifically the 2×2 leading principal minors $\theta_{x,y}^{(o)}$. Theorem 3.2 fixes the most important—in terms of optical accuracy and visual response—aspect of coherence: the spatial coherence area on the transverse plane. To complete the picture, we now need to quantify the power transfer between the transverse components of the incident and scattered beams, as well as the geometric deformation that the beam undergoes on scattering. Our primary contribution here stems from the realization that this information can be approximated from the classical pBSDF $\mathbf{M}^{|\mu_i| \rightarrow |\mu_o|}$, and we do not require more complicated quantities to describe matter.

We start with a short intuitive discussion, see Fig. 6. Consider the scattered beam and the shape of its intensity distribution on the transverse plane a short distance away from the scattering matter. By the Van Cittert-Zernike (VCZ) theorem [Mandel and Wolf 1995] this intensity distribution is related via a Fourier transform to the spatial coherence of the beam once it propagates further away. That is, if the intensity shape is isotropic, then the spatial coherence is isotropic as well and the leading principal minors of the shape matrices must be proportional to identity. On the other hand, if the scattering process induces anisotropy in that intensity shape, then that Fourier relation means that the leading principal minors will encode an inverse anisotropy (a stretch becomes a contraction, and vice versa).

As Theorem 3.2 mandates that the coherence area is a conserved quantity, we are left with only two degrees-of-freedom for $\theta_{x,y}$: the eccentricity of the coherence ellipse and its orientation on the transverse plane. These are geometric deformations, hence the transformation of the intensity shape can be deduced from the curvature of the scattered intensity, i.e., via the second-order derivatives of the pBSDF. The inverse of that transformation, then, quantifies the transformation of the spatial coherence on the transverse plane. We will now formalise these arguments and present an approach to compute the scattered shape matrices. Our discussion remains general and applicable to any material that can be described using Definition 3.3, however simplified relations are noted for specific materials in Subsection 4.1.

Any physically-realizable Mueller matrix can be decomposed via the serial *Lu–Chipman decomposition* [Pérez and Ossikovski 2016, Ch. 8.2], as follows:

$$\mathbf{M}^{|\mu_i| \rightarrow |\mu_o|} = m_{00} \mathbf{M}_{\Delta p} \mathbf{M}_d \mathbf{M}_r, \quad (17)$$

where m_{00} is the *mean intensity coefficient* of the pBSDF, i.e. the top left element, $\mathbf{M}_{\Delta p}$ is the normalized (unit mean intensity) *depolarizer* with zero diattenuation, \mathbf{M}_d is a normalized pure (non-depolarizing) *diattenuator* and $\mathbf{M}_r = \text{diag} \{1, \mathbf{R}_M\}$ is a pure *retarder*. We denote \hat{r}^\perp as the plane that is perpendicular to \hat{r} , i.e., the transverse plane of the beam. The transformation of the scattered intensity of the transverse plane is then the Hessian of the diattenuator:

$$U_{x,y} \triangleq \frac{1}{\left| \tilde{\mathbf{U}}_{x,y}^{-1} \right|} \tilde{\mathbf{U}}_{x,y}^{-1}, \quad \text{with} \quad (18)$$

$$\tilde{\mathbf{U}}_{x,y} \triangleq \vec{\mathcal{S}}_0^\top \frac{d^2}{d\vec{r}'^2} \left[\cos \vartheta_o m_{00} \mathbf{M}_d \right] \Big|_{\hat{s}, \hat{r} + \vec{r}'}, \quad (19)$$




Fig. 5. The wave BSRDF (wBSDF): Plots of the radiance $L = L_x + L_y$ scattered by a scratched metal surface. Incident light with unit irradiance impinges upon the surface at normal incidence. Three different representative wavelengths are used as well as light of different spatial coherence area: (a) comparable to sunlight; (b) moderately coherent; (c) weakly coherent; and (d) highly diffused. Also plotted are Lambertian BRDF lobes, viz. $\frac{1}{\pi} \cos \vartheta_o$, where ϑ_o is the scattering angle, in gray, for comparison. The standard deviation σ of the characteristic length of spatial coherence on the transverse plane is listed. Incoherence is equivalent to an angular blurring of the wave BSRDF (convolution with the inverse coherence Gaussian), and we illustrate the angular magnitude of the blur using the purple lines at the top of the plots. An important distinction between the classical BSRDF and our wBSDF is that, due to constructive wave interference, the scattered power ratio $\partial\Phi^{(o)}/\partial\Phi^{(i)}$ may now surpass unity (assuming $\Omega^{(i)} = \Omega^{(o)}$, where Ω is the solid angle into which a beam propagate, and assuming no re-emission effects, like fluorescence). Nevertheless, energy must be conserved over the entire scattering codomain (the unit sphere or hemisphere).

with the derivative taken with respect to $\vec{r}' \in \hat{r}^\perp$, and \hat{s} (the incident direction) is held fixed. The vector on the right, \vec{S}_{LHP} or \vec{S}_{LVP} , is used to select a transverse component, while \vec{S}_0 masks out the total scattered radiance. These vectors are defined in Eq. (4). As discussed, the inversion of $\tilde{U}_{x,y}$ is a consequence of VCZ theorem's Fourier relation. The normalization to a unit determinant done by U is a restatement of the fact that U should not change the coherence area. To compute \tilde{U} , we may derive an analytic expression for the Hessian or simply employ a finite differences approximation (via Taylor's theorem). The second approach requires evaluating the pBSDF at an additional 5 closely-spaced values of \hat{r} , and, while crude, is an acceptable approximation, due to the normalization to unit determinant. The purpose of the decomposition (Eq. (17)) is to extract the pBSDF's diattenuating component, which dictates the geometric deformations discussed above, and induces rotation between the transverse components

Let φ_z be the Euler rotation angle around the z -axis encoded by the matrix M_r (as in the decomposition of the pBSDF, i.e. Eq. (17)). This angle describes the rotation of the transverse plane performed by the pBSDF and captures both the transformation between the incident and scattered reference frames, as well as any retardation. Let $R(\varphi_z)$ be a 2×2 rotation matrix through an angle φ_z . Putting it all together, the scattered shape matrices can be written in the following form:

$$\theta_{x,y}^{(o)} = U_{x,y} R(\varphi_z)^T \left(\cos^2 \varphi_z \theta_{x,y}^{(i)} + \sin^2 \varphi_z \theta_{y,x}^{(i)} \right) R(\varphi_z) U_{x,y}^T. \quad (20)$$

A quick check verifies that $\theta_{x,y}^{(o)}$ remain real positive-definite. The above equation first mixes the incident shape matrices based on the rotation between the incident and scattered transverse planes, then transforms each to the orientation of the transverse plane in the outgoing reference frame, and, finally, the discussed geometric deformations are applied to each matrix.

In Appendix D we provide a step-by-step overview of how to decompose a non-depolarising Mueller matrix as in Eq. (17) and present a simpler expression for \tilde{U} .

Relative roughness and specular lobes. Modern scatter theories (forms of the "T-matrix approach", see Sakurai and Napolitano [2021, Ch. 6] or Newton [1982, Ch. 7.2.2]) partition the contributions of the scatterer to the exitant energy into two: (i) radiation that interacts with scatterer; and (ii) "direct", or specular, contributions. The first type of contributions give rise to the coupled modes discussed in Subsection 3.4 and this active interaction with the scatterer is described by Eq. (16). The second contributions produces no interference effects. The radiance of the "direct" field is then

$$\mathcal{W}_{\text{direct}} \left\{ \vec{S}^{(i)} \right\} \triangleq M_{\text{direct}}^{[\mu_i] \rightarrow [\mu_o]} \vec{S}^{(i)}, \quad (21)$$

(i.e., the direct wBSDF is equivalent to the pBSDF) where $\vec{S}^{(i)}$ is the irradiance (in classical Stokes parameters form) of the incident radiation, i.e. $\vec{S}^{[\mu_i]}$ evaluated at $\vec{\xi} = 0$. The (non-depolarising) pBSDF M_{direct} quantifies the direct contributions, and typically is simply the Mueller matrix form of the Fresnel relations for specular reflection and refraction (in the directions dictated by the law of reflection and Snell's law, respectively), though any specular lobes are possible. The shape matrices of the direct contributions still need to be rotated, as in Eq. (20) (with respect to the direct pBSDF), but with a simplification: as the specular lobes are discrete, no anisotropies arise and we set $U_{x,y} \equiv I$.

The energy partitioning between these contributions depends on the matter's properties. Davies [1954] rigorously derived an exponential relation that depends on the *roughness* of a scattering surface. This remains the state-of-the-art, with only minor adjustments made by Harvey [2012] that replace that roughness with a band-limited *effective roughness*, and we use the same approach.

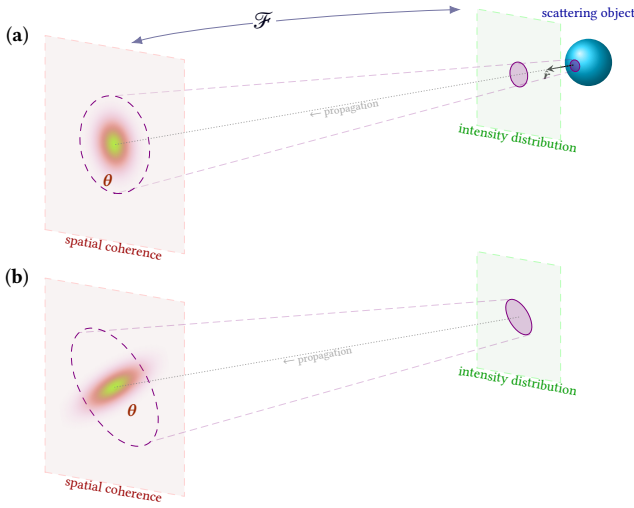


Fig. 6. Illustration of light that is scattered of the surface of an object. The intensity distribution on the transverse plane of the scattered beam is illustrated on the green plane on the right. This intensity distribution is related to the spatial coherence function of the light on the transverse plane after propagation (red plane on the left) via a Fourier transform. (a) If the material scatters isotropically and the intensity distribution is isotropic, then the spatial coherence must be isotropic as well. (b) On the other hand, anisotropy in the scattering characteristics might induce deformations in the beam. Then, the spatial coherence after propagation undergoes a transformation that is the inverse of the intensity distribution’s transformation. This transformation of the spatial coherence on the transverse plane is in fact the leading principal minor of the shape matrix θ .

While colloquially the term “roughness” is typically used when discussing surfaces, we define and use it for any matter, and in Appendix E we show how the relative roughness relates to the PSD.

We are now ready to formulate the total wBSDF:

$$\mathcal{W} \triangleq \alpha_{\text{direct}} \mathcal{W}_{\text{direct}} \left\{ \vec{\mathcal{S}}^{(i)} \right\} + (1 - \alpha_{\text{direct}}) \mathcal{W}_{\text{scat}} \left\{ \vec{\mathcal{S}}^{|\mu_i|} \right\}, \quad (22)$$

where $\alpha_{\text{direct}} = \exp[-(k \cos \vartheta_i q_{\text{rel}})^2]$ partitions between the direct and scattered contributions, with q_{rel} being the relative roughness (Eq. (63)) and ϑ_i the incident inclination angle.

Importance sampling the wBSDF. The total wBSDF consists of three quantities that need to be importance sampled: (i) the ratio partitioning into direct and scattered contributions, α_{direct} ; (ii) the direct wBSDF $\mathcal{W}_{\text{direct}}$; and (iii) the scattering wBSDF $\mathcal{W}_{\text{scat}}$. Importance sampling the first two is well understood: the direct wBSDF is simply a classical pBSDF that consists of, at most, a finite count of specular lobes (where no diffraction occurs and coherence plays no effect). More troublesome is the scattering wBSDF: the power contained in the scattered beam depends on the coherence properties of the incident light. Hence, when performing bi-directional path tracing, the scattering wBSDF cannot be evaluated when tracing from a sensor (as the shape matrices Θ are unknown), and exact importance sampling is not possible.

Nevertheless, the angular distribution of power quantified by the scattering wBSDF may be estimated. The diffraction operator \mathcal{D}

acts as a low-pass filter that performs an angular Gaussian blur of the PSD: lower spatial coherence results in greater blurring of the pBSDF, while perfectly-coherent radiation induces no filtering. Let σ_{min} be a global (scene wide) lower bound on the singular values of the coherence shape matrices Θ , implying that we do not render with light that admits a characteristic length of its coherence volume that is lower than σ_{min} (satisfying this constraint is discussed in Section 4). Then, we rewrite the scattering wBSDF as:

$$\mathcal{W}_{\text{scat}} \left\{ \vec{\mathcal{S}}^{|\mu_i|} \right\} = \mathcal{W}_{\text{MInc}} \left\{ \vec{\mathcal{S}}^{|\mu_i|} \right\} + \widetilde{\mathcal{W}} \left\{ \vec{\mathcal{S}}^{|\mu_i|} \right\}, \quad (23)$$

where the first term is the *maximally-incoherent* wBSDF, defined as

$$\mathcal{W}_{\text{MInc}} \triangleq \mathcal{W}_{\text{scat}} \Big|_{\mathbf{E}_{x,y,z}^{(i)} \rightarrow \sigma_{\text{min}} \mathbf{I}}, \quad (24)$$

i.e. the scattering wBSDF as in Eq. (16), but evaluated with respect to the shape matrices $\sigma_{\text{min}} \mathbf{I}$. The second term is then $\widetilde{\mathcal{W}} \triangleq \mathcal{W}_{\text{scat}} - \mathcal{W}_{\text{MInc}}$, which encodes the very high frequencies of the wBSDF and depends on the light’s coherence (see Fig. 7). Therefore, the maximally-incoherent wBSDF $\mathcal{W}_{\text{MInc}}$ serves as an envelope to the actual scattering wBSDF, in the sense that if the scattered radiance quantified by $\mathcal{W}_{\text{MInc}}$ is negligible then the actual scattered radiance (quantified by $\mathcal{W}_{\text{scat}}$) must be negligible as well, though the converse does not always hold. No coherence information is required to evaluate the maximally-incoherent wBSDF, hence it can be importance sampled when tracing paths from light or sensor.

4 METHODOLOGY OF RENDERING WITH PLT

Building upon the theory presented in Section 3, we outline the implementation of our rendering framework. Our implementation is a custom integrator plugin for Mitsuba 0.6 [Jakob 2010], based upon the bi-directional integrator `bdpt`.

Light emission. Light sources may be arbitrarily shaped, but must admit a positive spatial extent. We sample a position on the light source and treat that position as a small radiator with projected area A (in the direction of propagation). Effectively, we understand the light source as a very large collection of mutually-incoherent small radiators—an excellent model for spontaneous emission sources. That radiator emits a beam of light into a small solid angle Ω . Previous work [Steinberg and Yan 2021a] considered a beam emitted by the entire source, however for larger sources this leads to a pair of issues: (i) beams are traced from the centre of the source only, leading to paraxial errors; and (ii) beams might be highly-incoherent. The second point is an issue as our importance sampling method requires a global lower bound on the light’s spatial coherence. Our approach is more physically plausible and eliminates both problems with virtually no downside.

Spontaneous emission sources give rise to randomly polarized radiation, viz. $L_x = L_y$ and $\zeta = \chi = 0$. The shape matrices arise directly via the Hessian of the Fourier transform relation described by the VCZ theorem [Steinberg and Yan 2021b]:

$$\Theta_{x,y} = \begin{pmatrix} \frac{r^2}{k^2} \frac{2\pi\Omega}{A} & & \\ & \frac{r^2}{k^2} \frac{2\pi\Omega}{A} & \\ & & \sigma_{zz}^2 \end{pmatrix}, \quad (25)$$




Fig. 7. Illustration of wBSDF importance sampling procedure. The scattering wBSDF $\mathcal{W}_{\text{scat}}$, plotted in orange, depends on the coherence properties of light and can not be evaluated when these properties are unknown (e.g., when tracing paths from the sensor). However, as the spatial coherence area grows on propagation and passive optical elements may not induce decoherence (Eq. (13)), in practice, a reasonable lower bound on the spatial coherence can be set. This gives rise to the *maximally-incoherent wBSDF* $\mathcal{W}_{\text{Mlnc}}$, plotted in purple, which is coherence-independent—a classical pBSDF—and can be importance sampled. The difference between these wBSDFs captures the high-frequency details and is denoted $\overline{\mathcal{W}}$, with its absolute value is plotted in blue.

and we may verify that the coherence area grows linearly with the beam cross section (in line with Theorem 3.2), while radiation from what may be understood as an idealised point source, at the limit $A \rightarrow 0$, indeed gives rise to perfectly coherent radiation. Nevertheless, the underlying Fourier optics assumption mandates that $A \gg \lambda^2$, and our assumption that Ω dictates the cross-sectional area of the beam implies that $\Omega \gtrsim A$. The r^2/k^2 term above is the leading factor of the leading principal minors θ , as in Eq. (12).

Only the ratio Ω/A enters our equations: Unreasonably large values of Ω/A imply that Ω is too large (hence the beam’s cross-sectional area becomes too large) or A is too small (and we trace from very tiny radiators), either way coherence is overestimated leading to wave-ensemble aliasing (see Subsection 3.2). Nevertheless, the rendering will correctly converge (but slowly), as the effects of spatial coherence will be reproduced by sampling enough points on the source. On the other hand, unreasonably small values of Ω/A imply that A is too large, giving rise to paraxial errors and incorrectly underestimating the spatial coherence. Our framework is robust: noticeable changes in performance require order-of-magnitude changes in Ω/A .

For typical artificial light sources, we choose $A = 1 \text{ mm}^2$ and $\Omega = 10^{-6}$. For sunlight, we set $\Omega/A = d_{\text{sun}}^{-2} \sigma_{\text{sun}}$, where d_{sun} is the distance to the sun and $\sigma_{\text{sun}} \approx 50^2 \mu\text{m}^2$ is the sunlight’s spatial coherence variance on earth. Note that $r \approx d_{\text{sun}}$ (the light’s distance of propagation once it arrives on earth) and that Ω/A is multiplied by the r^2/k^2 term, leading to characteristic lengths of coherence after propagation of $\sim \sigma_{\text{sun}}$, as desired. Similarly for other distant sources. These are fairly reasonable values that should fit a wide-range of different scenes.

Path tracing. For light sourced from artificial sources, the above implies that after 1 m of propagation, the beam’s cross-sectional area is $\sim 1 \text{ mm}^2$. This is a tiny area, therefore we do not need to employ fully-fledged beam tracing, instead, for the purpose of propagating our beams throughout the scene, we treat these beams as rays and perform typical ray tracing. Classical light transport aims to solve the rendering equation, and PLT admits a modified transport equation (see Steinberg and Yan [2021a, Definition 4.1]) that governs the propagation of the generalized radiance at an interface between media:

$$\vec{\mathcal{L}}^{\mu_o} = \int_{\mathfrak{E}_+^2} d^2 \hat{s} \mathcal{W} \{ \vec{\mathcal{L}}^{\mu_i} \} \hat{s} \cdot \hat{n}, \quad (26)$$

where \mathfrak{E}_+^2 denotes the unit hemisphere and \hat{n} is the normal to the interface. Extension of the above to participating media has been done as well [Steinberg and Yan 2021a]. Due to the linearity of the diffraction operator and Eq. (10), we note that $d^2 \hat{s} \mathcal{W} \{ \vec{\mathcal{L}}^{\mu_i} \} \hat{s} \cdot \hat{n} = d^2 \mathcal{W} \{ \vec{\mathcal{S}}^{\mu_i} \}$. Finally, beams that take different paths through the scene are essentially always mutually incoherent: the characteristic length of spatial coherence of light from spontaneous emission sources is, at most, on the order of hundreds of micrometres. Therefore, Property 3.2 applies and distinct beams superpose linearly, enabling Monte-Carlo integration of the transport equation (Eq. (26)).

Our renderer performs spectral rendering with 16 spectral samples. Physical optics light-matter are heavily frequency dependent, furthermore to reproduce temporal coherence effects we need to accurately sample the spectrum of light (see Subsection 3.2). We stress again that temporal coherence effects are reproduced by our framework, at no additional cost, simply by sampling the spectrum of light. This is due to the spectral decomposition we employ (as discussed in Subsection 3.2). The parameter σ_{zz} only serves to limit the spatial coherence of light to a reasonable extent in the direction of propagation. We set $\sigma_{zz} = 1 \text{ mm}$ constant throughout. As expected, only order-of-magnitude changes in σ_{zz} have any impact on performance.

Importance sampling. The problem of importance sampling reduces to importance sampling the maximally-incoherent wBSDF $\overline{\mathcal{W}}$, which, as discussed in Subsection 3.5, acts as a classical pBSDF. For simplicity, our implementation ignores polarization during importance sampling and samples the scalar mean scattering intensity of $\overline{\mathcal{W}}$, viz. $\cos \vartheta_o m_{00} \mathcal{D} \{ \sigma_{\min} \mathbf{I} \}$, where σ_{\min} is the global lower bound on the spatial coherence variance. We set σ_{\min} to equal the spatial coherence variance of light after a short propagation of $r = 10 \text{ cm}$, i.e. $\sigma_{\min} = r^2 \lambda_{\min}^2 \frac{\Omega}{2\pi A}$, with λ_{\min} being the shortest wavelength that we render with (this follows from Definition 3.1 and Eq. (25)). Then, importance sampling is done by sampling the material-specific m_{00} and the convolved PSD, and we discuss it further in Subsection 4.1.

Note, in our implementation we sample with respect to the scalar maximally-incoherent wBSDF both backward and forward paths. This is done for simplicity, as we may ignore coherence and polarization entirely when tracing paths. Only once a complete path, light source to sensor, is found, we evaluate it by applying the full vectorized, coherence-aware formalism presented. This works for the vast majority of materials, because the maximally-incoherent

wBSDF provides a good approximation to the actual scattering wBSDF (see Fig. 7), however for some specific materials, like diffraction grating, which are highly coherence-sensitive, a better implementation would put the knowledge of the light's coherence properties to better use, when tracing forward paths.

4.1 Materials

We discuss the materials that were rendered in this paper.

4.1.1 Surfaces with varying-degrees of roughness. Let a wide-sense-stationary statistical surface be described by the Gaussian autocorrelation function $r(\vec{d}) = q^2 g^{\sigma_s^2}(\vec{d})$, with $\vec{d} \in \mathbb{R}^2$ a distance on the surface, q then becomes the *root-mean-square (rms) roughness* (as in Eq. (61)) and σ_s is the surface's *correlation length*. The PSD of the surface arises immediately via Eq. (62) (and Lemma S2.1 in the supplemental material):

$$p(\vec{\xi}) = \mathcal{F}\{r\}(\vec{\xi}) = q^2 \sigma_s^2 g^{\sigma_s^{-2}}(\vec{\xi}). \quad (27)$$

It is well-known that the probability density function of a sum of independent random variables is the convolution of their respective probability density functions—a fact that, in general, may facilitate finding analytic closed-form solutions or approximations to the diffraction operator in Eq. (15). For our surface, the diffraction operator is then simply

$$\mathcal{D}\{\Sigma\} = \frac{q^2}{2\pi} |\Sigma^{-1} + \sigma_s^{-2} \mathbf{I}|^{-1/2} g^{\Sigma^{-1} + \sigma_s^{-2} \mathbf{I}}(\vec{h}), \quad (28)$$

with Σ is restricted to the 2×2 -dimensional case (i.e. we extract the leading principal minor of Σ). We also require the (direction- and wavelength-dependent) relative roughness, defined in Eq. (63). For simplicity, and because for a Gaussian surface the integral over the entire surface approximates the band-limited integral in Eq. (63) very well for non-grazing angles, we set $q_{\text{rel}} = q$.

The scattering pBSDF that enters Eq. (16) for the surface is the Mueller form of the Fresnel relations, viz.:

$$\mathbf{M}^{[\mu_i] \rightarrow [\mu_{sp}]} \triangleq \frac{1}{\lambda^2} \begin{pmatrix} m_{00} & m_{01} & & \\ m_{01} & m_{00} & & \\ & & m_{22} & m_{23} \\ & & -m_{23} & m_{22} \end{pmatrix} \mathbf{T}^{[\mu_i] \rightarrow [\mu_{sp}]},$$

with $m_{00} = \frac{|f_{pp}|^2 + |f_{ss}|^2}{2}$, $m_{01} = \frac{|f_{pp}|^2 - |f_{ss}|^2}{2}$,

$$m_{22} = \text{Re}\{f_{ss} f_{pp}^*\} \quad \text{and} \quad m_{23} = \text{Im}\{f_{ss} f_{pp}^*\}, \quad (29)$$

where the exit reference frame $[\mu_{sp}]$ is defined as the standard s- and p-polarization frame: \hat{x} is perpendicular to both \hat{r} and \hat{s} and \hat{y} lies in the plane spanned by \hat{r} and \hat{s} ; T is the Mueller rotation matrix that rotates from the (arbitrary) incident frame to the sp frame (see Eq. (59) for an explicit definition). The quantities f_{ss}, f_{pp} are the well-known reflection or refraction Fresnel coefficients, which depend on the relative normal vector (the geometric half vector between \hat{r} and \hat{s} in the case of reflective scattering) and on the relative ratio of the refractive indices of the media, denoted η (which might be complex and is typically wavelength-dependent). The direct pBSDF that enters Eq. (21) is also as in Eq. (29), however it is restricted

to the typical specular lobes: perfect reflection and perfect refraction. The Hessians of the pBSDF are then (by applying Eq. (57))

$$\tilde{U}_x = \frac{d^2}{d\vec{r}'^2} |f_{pp}|^2 \Big|_{\hat{s}, \hat{r} + \hat{r}'} \quad \text{and} \quad \tilde{U}_y = \frac{d^2}{d\vec{r}'^2} |f_{ss}|^2 \Big|_{\hat{s}, \hat{r} + \hat{r}'}, \quad (30)$$

which we compute numerically, as discussed in Subsection 3.5.

Finally, we need to importance sample the envelope scattering wBSDF (Eq. (23)), which requires importance sampling the diffraction operator with $\Sigma = \sigma_{\text{min}} \mathbf{I}$. Because the variance matrix that enters the multivariate anisotropic Gaussian in Eq. (28) is now diagonal, this Equation becomes

$$\mathcal{D}\{\sigma_{\text{min}} \mathbf{I}\} = \frac{q^2}{2\pi w} g^w(h_x) g^w(h_y), \quad (31)$$

with $w = \sigma_{\text{min}}^{-1} + \sigma_s^{-2}$ the effective standard deviation. Therefore, we importance sample each univariate Gaussian above, which reduces to evaluating the inverse error function (see Winitzki [2003] for a numeric routine), yielding an outgoing direction \hat{r} . If the surface is transmissive, then we also importance sample a reflective or refractive scattering interaction, with respect to the surface's reflectivity. For simplicity, we do not importance sample the pBSDF $\mathbf{M}^{[\mu_i] \rightarrow [\mu_{sp}]}$, as the distribution of scattered power is dominated by the PSD. See our implementation.

The presented surface model is parametrized by the roughness q , the correlation length σ_s^2 (which, in term of optical response, inversely modulates the length of the “tail” of the highlight) as well as the wavelength-dependent refractive index η . This model is fully coherence-aware, and is easy to importance sample in a bi-directional path tracer. Nevertheless, real surfaces are rarely Gaussian, and typically their PSD admits a more “fractal” structure with a longer tail (more pronounced higher spatial frequencies). Deriving expressions for more realistic surface models is possible (as effectively the problem reduces to deriving closed-form expressions for the well-behaved diffraction operator Eq. (15) and importance sampling it), and is left for future work.

4.1.2 Perfectly-diffuse surfaces. A perfectly-diffuse surface is a special case of the rough surface discussed above, and can be understood at the infinite roughness and zero correlation length limit, viz. $q \rightarrow \infty$ and $\sigma_s \rightarrow 0$ simultaneously. Then, Eqs. (27) and (28) reduce to $p = \mathcal{D} = 1$. No direct contributions arise, therefore $\alpha_{\text{direct}} = 0$. The scattering pBSDF of a perfectly-diffuse surface is

$$\mathbf{M}^{[\mu_i] \rightarrow [\mu_{sp}]} \triangleq \frac{m_{00}}{\pi \lambda^2} \mathbf{T}^{[\mu_i] \rightarrow [\mu_{sp}]}, \quad (32)$$

where $0 \leq m_{00} \leq 1$ is the mean scattering intensity, T is the Mueller rotation matrix, as before, and we normalize by $1/\pi$.

4.1.3 Diffraction grating. We use a similar formulation to Steinberg and Yan [2021b]. Consider a (perfect) sinusoidal phase grating, with pitch (period) Λ and grating height of b , residing on an infinite plane. Assume that the reflectivity is constant over the entire surface, hence the pBSDF is the same Fresnel Mueller matrix as in Eq. (29). The complex transmission function for a plane wave incident at angle ϑ_i is well-known [Born and Wolf 1999]:

$$\phi(x) \triangleq \exp \left[ik \frac{b}{2 \cos \vartheta_i \cos \vartheta_o} \sin \left(2\pi \frac{x}{\Lambda} \right) \right], \quad (33)$$




Fig. 8. (left) Real photos of a road bicycle's rear wheel assembly. (a6) The translucent spoke guard is made of moulded plastic and admits a significant degree of birefringence. (a7) The brake surface on the Giant S-R3 wheel is made of an aluminium alloy with a grated surface that (a5) acts as a diffraction grating. (a1) When illuminated by sunlight and viewed through a polarization filter (sunglasses), visible birefringence-induced iridescence arises on the spoke guard, while the wheel disperses the light and diffraction gratings patterns are visible on ground. (a2) Same conditions but without a polarization filter: the iridescence on the spoke guard disappears almost entirely. (a3-4) In place of sunlight, a powerful artificial LED light source generates rather weakly-coherent incident light. With and without a polarization filter. The diffraction grating patterns disappear entirely, but, (a3) when viewed through a polarization filter, weak birefringent effects are still visible. (right) We render a similar close-up of a rear wheel assembly, (b2) with and (b1) without a polarization filter. The birefringence-induced iridescence on the spoke guard is high polarization sensitive, but the diffraction grating patterns are unaffected.

which quantifies the phase difference of the scattered wave due to the height differences across the grating, with x being the position on the grating. Applying the well-known identity of the Bessel function of the first kind, viz. $e^{iv \sin \varphi} = \sum_{n=-\infty}^{\infty} J_n(v) e^{in\varphi}$, serves to rewrite the transmission function ϕ as a Fourier series with coefficients $J_n(v)$. The PSD of a Fourier series is a series of Dirac deltas:

$$p(\zeta) = \sum_{n=-\infty}^{\infty} J_n \left(k \frac{b}{2 \cos \vartheta_i \cos \vartheta_o} \right)^2 \delta \left(\zeta - 2\pi \frac{n}{\Lambda} \right). \quad (34)$$

Evaluated at $\zeta = k(\sin \vartheta_o + \sin \vartheta_i)$, i.e. the projection of \vec{h} (as defined in Eq. (14)) onto the grating axis, the PSD of the grating above reproduces the known relation [Born and Wolf 1999]:

$$\sin \vartheta_o + \sin \vartheta_i = \frac{n}{\Lambda} \lambda, \quad (35)$$

which governs the scattering angle of lobe n at wavenumber k . The intensity of each lobe is given by the Bessel function in Eq. (34).

Clearly, the DC lobe $n = 0$ (the direct, non-diffracted lobe) scatters at $\vartheta_i = \vartheta_o$, regardless of wavelength. The diffraction lobes, $n \neq 0$, scatter at an angle that is wavelength dependent. The convolution of the PSD with the coherence Gaussian, viz. Eq. (15), is then a series of Gaussians, where the variance matrix is the shape matrix

of the spatial coherence of light, and with the mean of each Gaussian centred at the respective lobe. In practice, we restrict the rendered lobes such that $|n| \leq N$. Importance sampling then reduces to choosing a lobe, uniformly at random, from these $2N$ diffraction lobes as well as 1 DC lobe, and sampling the Gaussian of that lobe. It is important to note that the lobes generally overlap, which should be taken into account when calculating the probability density of the chosen importance sample. As all lobes are specular, $\vec{U}_{x,y} = I$.

With moderately-coherent light, like sunlight, these lobe Gaussians have very low variance. Indeed, a diffraction grating describes the most difficult kind of a material to importance sample: A multitude of Dirac lobes, with each restricted to a single wavelength and carrying only a small fraction of energy. See Fig. 8a for photos of a diffraction grating dispersing light.

4.1.4 Birefringent dielectrics. Birefringence is an optical phenomenon, where the permittivity of the matter admits spatial anisotropy, i.e., it depends on the direction of incidence. Birefringent dielectrics are quantified by a pair of refractive-indices, an *ordinary* η_o and an *extraordinary* η_e , as well as an *optic axis*, that describes the distinguished direction in the material around which anisotropy is induced. As light refracts into such a material, the beam splits into




Fig. 9. Real photos of a glass window pane, (a) with a polarization filter and (b) without. In a similar manner to the birefringent bicycle spoke guard, when viewed through a polarization filter, the glass window appears iridescent. We reproduce this interaction between the polarimetric and wave properties of light on the glass window in the rendered Figs. 1 and 10.



Fig. 10. The bicycle scene from Fig. 1 rendered with indoor lighting that consists of a pair of 3900 K halogen bulbs, as well as late evening sky. Also viewed through a polarization filter. These light give rise to radiation that far less coherent than sunlight, therefore the (red close-up) diffraction grating patterns disappear entirely. (blue close-up) The birefringence effects on the windows also diminish, but remain slightly visible.

a pair of beams, an *ordinary beam*, which refracts with respect to η_o , and an *extraordinary beam*, which encounters a refractive-index between η_o and η_e , depending on orientation. Once refracted out of the birefringent material, both beams encounter the same refractive-index of the external medium, and recombine. The difference in the optical path length travelled by these beams induces diffraction effects, which are known as *birefringence-induced iridescence*.

Birefringent dielectrics are implemented as thin dielectric slabs, with given thickness and with the surfaces being Gaussian surfaces of arbitrary roughness. The transport through the slab is then evaluated atomically, with no ray tracing performed inside the slab. Rendering is done similarly to a rough dielectric, Subsection 4.1.1,

and importance sampling remains unchanged. To compute birefringence, we use the framework developed by Steinberg [2019]. In place of the typical (isotropic) Fresnel coefficients in the pBSDF, we use the anisotropic Fresnel coefficients (see Steinberg [2019, supplemental] for a derivation). As the pair of beams refract out of the slab, we superpose them, with respect to an interference term that arises from the computed optical path difference. That interference term is modulated by the mutual coherence, i.e. the light's coherence Gaussian evaluated at the (mean) spatial distance between the ordinary and extraordinary beams. The Hessians of the pBSDF take the same form as with a rough surface, Eq. (30).

As the ordinary and extraordinary beams diverge only by a small amount, a birefringent dielectric is less sensitive to the spatial coherence of light, compared with a diffraction grating. On the other hand, the appearance of a birefringent dielectric depends heavily on the polarization properties of light. See Fig. 8a for a photo of such optical effect.

5 RESULTS

In contrast to previous PLT work, we are able to render complex scenes, where all the materials are the coherence-aware physical optics materials that were developed and discussed in Subsection 4.1. We are able to trace light beams of arbitrary spectrum, polarization and spatial coherence area, volume and anisotropy.

We render a bicycle that contains two surfaces with interesting diffractive properties: A birefringent plastic spoke guard and a aluminium wheel brake surface that acts as a diffraction grating. See Figs. 8a and 9 for photos of these materials. These effects are sensitive to the polarimetric and coherence properties of light: The glass panes and plastic bicycle spoke guard in Figs. 1 and 8b are birefringent and appear iridescent, with the effect most stark when viewed through cross polarizers, i.e., when polarized incident light (which arises on, e.g., reflection of a surface or the sky) transmits through the birefringent sample and then is passed through a polarization filter (like sunglasses). Our renderer reproduces these effects accurately, see Fig. 8b. As discussed in Subsection 4.1.4, subtle birefringence-induced iridescence may even when light is weakly-coherent. This effect is captured in the photos in Figs. 8a and 9 and reproduced by our framework in Fig. 10, where the bicycle scene is rendered at night with lighting dominated by weakly-coherent indoor halogen bulbs.

On the other hand, a diffraction grating is not sensitive to the polarization of light, but is highly coherence sensitive, as it scatters into many diffracted Dirac delta lobes. The metal brake surface on the Giant S-R3 bicycle wheel acts as a lateral diffraction grating (see Fig. 8a). This effect is readily visible when moderately-coherent sunlight is incident upon that surface, but disappears entirely with weakly-coherent light. We reproduce these optical phenomena in Figs. 8b and 10. See our supplemental material for full-resolution photographs, as well as high-resolution renderings. The diffraction grating surface gives rise to 6 diffraction lobes ($|n| \leq 3$, with n as in Subsection 4.1.3), with each lobes consisting of 16 dispersed lobes (one for each spectral sample). Together with the DC lobe, this amounts to 97 distinct lobes, which are highly coherence sensitive




Fig. 11. Comparison with the “twin scarab” scene from Steinberg and Yan [2021b]. The scene, illuminated by a cylindrical fluorescent light source of length L , was rendered using our framework with different light source sizes, leading to (left to right) strongly-coherent, moderately-coherent and weakly-coherent radiation incident upon the scarabs. The large scarab’s wings are coated with a $3\ \mu\text{m}$ perfectly smooth dielectric, resulting in visible interference effects. A mirror is placed on the right-hand side of the scene, reflecting light towards the scarabs. Light that arrives from the mirror has propagated farther, thus is more coherent, and more pronounced interference effects are visible on the right wing. (insets) Same scene rendered using Steinberg and Yan [2021b] framework. The material appearance between the renderings is not expected to match, as the material models are different (see Section 5), however the interference patterns match very well.




Fig. 12. Analysis of convergence performance. The modified “twin scarabs” scene (see Section 5 for details) is rendered by our framework as well as by the framework of Steinberg and Yan [2021b]. The small images are close-ups and are rendered with (left to right) an increasing count of samples-per-pixel (spp). Difference in material appearance is expected as the material models differ between the frameworks. The rendering times for the 1024spp images are 38.4 minutes and 9.8 minutes for ours and Steinberg and Yan [2021b], respectively. The greater time per sample for our framework is to be expected as we perform fully vectorized light transport that accounts for polarization, while theirs is a scalar renderer. The 10k spp image took 91 minutes of rendering time by their framework, however considerable noise is still present. Our framework approaches a converged rendering at 1024spp.

and hence are difficult to importance sample, making rendering the scenes in Figs. 1, 8b and 10 challenging.

We compare our framework against Steinberg and Yan [2021b]. To be able to render scenes with some complexity, most of the materials that they render are classical microfacet surfaces, rendered and sampled classically. For the materials that admit diffractive behaviour, they manually fit a microfacet “envelope” that is used for importance sampling. The transformation of the incident light’s coherence function in Hermite-Gauss space on interaction with matter is pre-computed for each Hermite-Gauss mode. This approach is (i) cumbersome, as it requires manual tuning and pre-computation for each distinct material; and (ii) inaccurate, as with

weakly-coherent light that microfacet envelope might not suffice (unless a global lower bound on the light’s coherence is ensured, as done by us). Our new materials are devoid of such issues. We render their “twin scarab” scene that includes coherence-dependent thin-film interference effects, and compare against their renderings, see Fig. 11. Slight differences in material appearance are to be expected, but we show that the interference patterns are precisely reproduced, suggesting that our coherence transport is accurate. The “twin scarab” scene contains a perfect mirror and the materials are mostly either diffuse or low-roughness conductors. To study convergence, we modify the “twin scarab” scene slightly in order to increase the complexity of the light transport: we make the mirror




Fig. 13. Analysis of convergence performance. We compare our framework against scalar, radiometric renderer (Mitsuba 0.6, using the bdpt integrator). Both renderers use 16 spectral samples. The small images are equal-time renderings. In these rendered scenes, our framework incurs a cost per sample that is about twice greater. The diffraction gratings on the wheels’ brake surfaces are slow to converge, as expected, but otherwise equal-quality rendering times of our framework are at about 3-4 times compared to the scalar, radiometric renderer. Differences in material appearance are expected.

rough, and tweak the materials such that there are surfaces with a variety of roughnesses. As expected, our framework converges significantly faster. Note, that even with 10 000 samples-per-pixel, the state-of-the-art is far from converged and admits significant noise. Our framework approaches convergence with 1024 samples. More complicated scenes with more complex materials, as in Figs. 1 and 8b, are impossible to render with the state-of-the-art, in any reasonable time frame.

Likewise, we also study convergence when compared to a scalar, radiometric renderer. Specifically, Mitsuba 0.6 with the bdpt integrator with 16 spectral samples. See Fig. 13. We use the night scene (Fig. 10) and the rear wheel assembly scene (Fig. 8b). These scenes are more complex than the “twin scarab” scene, hence integration cost is dominated by ray tracing, and the cost-per-sample of our renderer is about twice compared to the radiometric renderer. Overall, the convergence rate is roughly 3-4 times slower, see the RMSE plots in Fig. 14. It should be noted, however, that the slowest to converge are effects that arise due to the diffraction grating on the wheels. The radiometric renderer treats this material as a typical moderately rough conductor, while our renderer samples the 97 diffraction lobes, and the majority of the noise is induced by this material (this can be seen qualitatively in Fig. 13). Our renderer is also fully vectorized. The rendering times for all the scenes in this paper are listed in Table 2.

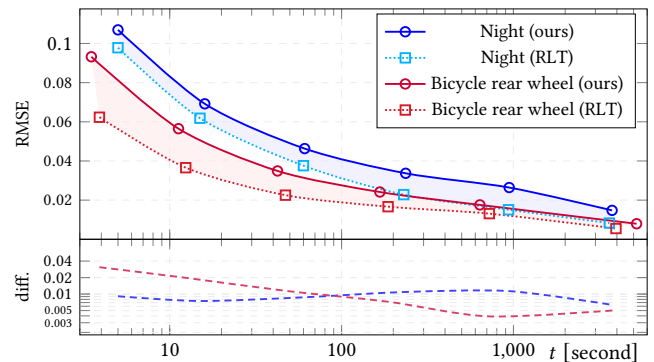


Fig. 14. Convergence performance of our framework compared against the scalar, radiometric renderer. (top) RMSE values as functions of rendering time for the (blue) “night” and (red) “bicycle rear wheel” scenes from Fig. 13, both when rendered under (dotted) RLT and our (solid) PLT method; and (bottom) the differences between the methods (log-log plot).

6 DISCUSSION AND CONCLUSION

The coupling relation between the incident and scattered principal modes, quantified by Eq. (13), is a hard physical limit for passive optical elements set by previous work [Zhang et al. 2019]. We

Table 2. Rendering runtimes of our framework for the scenes that appear in the paper. Rendering was done on an AMD[®] Ryzen[™] 5990X 64-core CPU (the “twin scarabs” scene was rendered on a AMD[®] Ryzen[™] 9 5950X CPU).

Scene		Samples	Rendering time
Bicycle day (3125 × 2000)	Fig. 1	150k	43 h 40 min
Bicycle night (2100 × 2100)	Fig. 10	90k	17 h 30 min
Rear wheel (2150 × 2150)	Fig. 8b	90k	18 h 50 min
Twin scarabs (1800 × 1200)	Fig. 11	8196	3 h 36 min

build upon that by noting that engineered materials that admit significant mode-specific absorption require very careful photonic engineering of *material structuring*. The assumption that the materials that we encounter and wish to render are not structured, i.e. are mode-agnostic, follows. Then, the conservation of spatial coherence, Theorem 3.2, can intuitively be understood as a form of a (functional) Central Limit Theorem: if the superposition of the incident optical modes tend to a Gaussian stochastic process, then a linear transform of these mode is expected to also tend towards a Gaussian process (because the coupling strengths, i.e. the singular values of this transform, are mode-agnostic).

By endowing the classical pBSDF with spectral information of the matter’s scattering characteristics—specifically its power spectral density—we introduce the wave BSDF and show that this information is sufficient for physical optics scattering. The derivations (available in our supplemental material) that lead to the general formulation of the wBSDF in Eq. (16) are essentially exact, up to the typical optical far field assumption (which, as mentioned, is very reasonable for our purposes and field statistics converge to far field statistics rapidly [Agarwal et al. 2004]). We have introduced a Gaussian surface model, however real surfaces are rarely Gaussian in their PSD. This means that we do not expect to achieve the visual fidelity of surface reproduction of, e.g., microfacet surfaces with a GGX distribution [Walter et al. 2007] (which was designed to match measured data), as used by the radiometric renderer. However, the appearance of statistical surfaces is dictated by the first-order diffractions that arise [Stover 2012]. The microfacet surface model lacks the ability to describe statistical correlations between the facets, and that information is needed to diffract light. Furthermore, microfacet surfaces simply do not arise in practice, and it serves to reason that this model should be superseded. Our wBSDF is able to account for these coherence- and frequency-sensitive diffractions, and we leave designing more expressive surfaces and material models for future work.

In this paper, we have presented the first practical algorithm for physical light transport that is consistent with Maxwell’s theory of electromagnetism and is able to reproduce physically-accurate wave-interference and diffraction effects at a global scale for complex scenes. To achieve this, we have proposed a new, wave-packet

rendering primitive based on Stokes parameters that can be easily integrated into a standard rendering system and which generalizes radiometric radiance and irradiance. To enable a simple analytic representation of a beam, we restrict the coherence shape of the wave packet to be a simple, anisotropic Gaussian, which we show using entropy conservation is a reasonable assumption. We have also presented a novel representation for the BSDFs of materials to make them coherence-aware so that they interact correctly with our wave packets, and showed how we can importance sample them in order to facilitate bi-directional path tracing, with only minor modifications to the modern offline path tracing pipeline.

We then show how all of these contributions can be integrated into a standard rendering system to simulate light transport in a complex scene in a physically accurate manner. In contrast to finite-difference time-domain (FDTD) approaches that are used in the physics community to accurately simulate Maxwell’s equations but require extremely simplified scenes (e.g., a single tiny sphere), or SBR methods (see Section 2) that do not transport coherence information, our method works for scenes of the complexity found in high-end production environments, and yet its performance approaches that of classical rendering methods. Our method is able to accurately represent light of arbitrary spectrum, polarization, and coherence area and volume.

Limitations. Our framework is designed for polychromatic (broad-spectrum), partially-coherent light of optical frequencies. In this domain, we are able to leverage, with excellent accuracy:

- (1) the short characteristic length of spatial coherence, implying that distinct beams are mutually-incoherent and superpose linearly, enabling numeric integration of the transport equation (Eq. (26));
- (2) furthermore, the low spatial coherence allows us to use ray tracing to propagate our beams, as their cross-sectional area is comparable to smallest scene geometrical features; and,
- (3) the very short temporal coherence of light, meaning that second-order statistical moments (the CSD function, Eq. (7)) are sufficient to quantify the statistics of a wave ensemble.

The last point arises from the fact that our sensors always time-average over periods long compared to the temporal coherence of light. In a different regime of the electromagnetic spectrum, these assumptions are not as accurate: for example, if we were to render with frequencies used for W-band radar radiation (over ×1000 longer wavelength than optical), then it is reasonable to assume that fully-fledged beam tracing would be required for accurate rendering. In addition, at these longer wavelengths, free-space diffractions become more significant, and can no longer be ignored.

The failure cases of Theorem 3.2 have been discussed in Subsection 3.4. In a typical scene, where material structuring or materials with time-varying characteristics are not present, Theorem 3.2 applies as is. An increase in coherence of a scattered beam can be simulated in a straightforward manner, though recall that Eq. (13) dictates that an increase in coherence must always be accompanied by an energy trade-off. Because our importance sampling strategy requires a global lower bound on the light’s spatial coherence, care should be taken in the rare circumstances where decoherence

might arise. As mentioned earlier, a possible approach would be to integrate over the source of entropy that induces decoherence.

Finally, while in this work we have kept A and Ω (the beam's source area and solid angle, see Eq. (25)) constant throughout, it is possible to trace beams with different values of Ω/A , as long as the constraints discussed in Section 4 are met. This would allow us to correctly trace beams through, for example, lens systems and pinholes, as long as the global lower bound on spatial coherence is not violated.

REFERENCES

- Garish S. Agarwal, Greg Gbur, and Emil Wolf. 2004. Coherence properties of sunlight. *Optics Letters* 29, 5 (Mar 2004), 459. <https://doi.org/10.1364/ol.29.000459>
- Thomas Auzinger, Wolfgang Heidrich, and Bernd Bickel. 2018. Computational design of nanostructural color for additive manufacturing. *ACM Transactions on Graphics* 37, 4 (Aug 2018), 1–16. <https://doi.org/10.1145/3197517.3201376>
- Seung-Hwan Baek, Tizian Zeltner, Hyun Jin Ku, Inseung Hwang, Xin Tong, Wenzel Jakob, and Min H. Kim. 2020. Image-based acquisition and modeling of polarimetric reflectance. *ACM Transactions on Graphics* 39, 4 (Jul 2020). <https://doi.org/10.1145/3386569.3392387>
- Laurent Belcour and Pascal Barla. 2017. A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence. *ACM Trans. Graph.* 36, 4, Article 65 (July 2017), 14 pages. <https://doi.org/10.1145/3072959.3073620>
- Ahmad Bilal, Syed Muhammad Hamza, Ziauddin Taj, and Shuaib Salamat. 2019. Comparison of SBR and MLFMM techniques for the computation of RCS of a fighter aircraft. *IET Radar, Sonar & Navigation* 13, 10 (2019), 1805–1810.
- Max Born and Emil Wolf. 1999. *Principles of optics : electromagnetic theory of propagation, interference and diffraction of light*. Cambridge University Press, Cambridge New York.
- Juan D. Castro, Sahitya Singh, Akshaj Arora, Sara Louie, and Damir Senic. 2019. Enabling Safe Autonomous Vehicles by Advanced mm-Wave Radar Simulations. In *2019 IEEE MTT-S International Microwave Symposium (IMS)*. IEEE. <https://doi.org/10.1109/mwsym.2019.8700843>
- Pengning Chao, Benjamin Strehka, Rodrick Kuate Defo, Sean Molesky, and Alejandro W. Rodriguez. 2021. Physical limits on electromagnetic response. arXiv:2109.05667 [physics.optics]
- Mikhail Charnotskii. 2019. Coherence of radiation from incoherent sources: I Sources on a sphere and far-field conditions. *Journal of the Optical Society of America A* 36, 8 (Jul 2019), 1433. <https://doi.org/10.1364/josaa.36.001433>
- Yen-Sheng Chen, Fei-Peng Lai, and Jing-Wei You. 2019. Analysis of antenna radiation characteristics using a hybrid ray tracing algorithm for indoor WiFi energy-harvesting rectennas. *IEEE Access* 7 (2019), 38833–38846.
- Tom Cuyppers, Tom Haber, Philippe Bekaert, Se Baek Oh, and Ramesh Raskar. 2012. Reflectance model for diffraction. *ACM Transactions on Graphics* 31, 5 (Aug 2012), 1–11. <https://doi.org/10.1145/2231816.2231820>
- H. Davies. 1954. The reflection of electromagnetic waves from a rough surface. *Proceedings of the IEE - Part IV: Institution Monographs* 101, 7 (Aug 1954), 209–214. <https://doi.org/10.1049/pi-4.1954.0025>
- Aristide Dogariu and Emil Wolf. 1998. Spectral changes produced by static scattering on a system of particles. *Optics Letters* 23, 17 (Sep 1998), 1340. <https://doi.org/10.1364/ol.23.001340>
- Donald D. Duncan, Daniel V. Hahn, and Michael E. Thomas. 2003. Physics-based polarimetric BRDF models. *Optical Diagnostic Methods for Inorganic Materials III* (Nov 2003). <https://doi.org/10.1117/1.2507317>
- C. Fabre and N. Treps. 2020. Modes and states in quantum optics. <https://doi.org/10.1103/revmodphys.92.035005>
- V. Falster, A. Jarabo, and J. R. Frisvad. 2020. Computing the Bidirectional Scattering of a Microstructure Using Scalar Diffraction Theory and Path Tracing. *Computer Graphics Forum* 39, 7 (Oct 2020), 231–242. <https://doi.org/10.1111/cgf.14140>
- Tiantian Feng and Lixin Guo. 2021. Multiview ISAR Imaging for Complex Targets Based on Improved SBR Scattering Model. (2021).
- Ari T. Friberg. 1979. On the existence of a radiance function for finite planar sources of arbitrary states of coherence. *Journal of the Optical Society of America* 69, 1 (Jan 1979), 192. <https://doi.org/10.1364/JOSA.69.000192>
- Jorge I. Garcia-Sucerquia and Francisco F. Medina E. 2003. M2 factor invariance through the geometric etendue for Gaussian Schell model sources. , 343–346 pages. <https://doi.org/10.1078/0030-4026-00279>
- Joseph Goodman. 2015. *Statistical optics*. John Wiley & Sons Inc, Hoboken, New Jersey.
- Ibón Guillén, Julio Marco, Diego Gutierrez, Wenzel Jakob, and Adrian Jarabo. 2020. A General Framework for Pearlescent Materials. *ACM Transactions on Graphics* 39, 6 (2020). <https://doi.org/10.1145/3414685.3417782>
- J. C. Gutiérrez-Vega, M. D. Iturbe-Castillo, and S. Chávez-Cerda. 2000. Alternative formulation for invariant optical fields: Mathieu beams. *Optics Letters* 25, 20 (Oct 2000), 1493. <https://doi.org/10.1364/ol.25.001493>
- James E. Harvey. 2012. Total integrated scatter from surfaces with arbitrary roughness, correlation widths, and incident angles. *Optical Engineering* 51, 1 (Feb 2012). <https://doi.org/10.1117/1.oe.51.1.013402>
- Nicolas Holzschuch and Romain Pacanowski. 2017. A Two-scale Microfacet Reflectance Model Combining Reflection and Diffraction. *ACM Trans. Graph.* 36, 4, Article 66 (July 2017), 12 pages. <https://doi.org/10.1145/3072959.3073621>
- Weizhen Huang, Julian Iseringhausen, Tom Kneiphof, Ziyin Qu, Chenfanfu Jiang, and Matthias B. Hullin. 2020. Chemomechanical simulation of soap film flow on spherical bubbles. *ACM Transactions on Graphics* 39, 4 (Jul 2020). <https://doi.org/10.1145/3386569.3392094>
- Wenzel Jakob. 2010. Mitsuba renderer. <http://www.mitsuba-renderer.org>
- Adrian Jarabo and Victor Arellano. 2017. Bidirectional Rendering of Vector Light Transport. *Computer Graphics Forum* 37, 6 (Dec 2017), 96–105. <https://doi.org/10.1111/cgf.13314>
- Tom Kneiphof, Tim Golla, and Reinhard Klein. 2019. Real-time Image-based Lighting of Microfacet BRDFs with Varying Iridescence. *Computer Graphics Forum* 38, 4 (2019), 77–85. <https://doi.org/10.1111/cgf.13772>
- Olga Korotkova. 2017. *Random Light Beams*. CRC Press, Boca Raton.
- Olga Korotkova and Emil Wolf. 2005. Generalized Stokes parameters of random electromagnetic beams. *Optics Letters* 30, 2 (Jan 2005), 198. <https://doi.org/10.1364/ol.30.000198>
- Gong Lei, Zhensen Wu, and Honglu Hou. 2012. Polarized bidirectional reflectance distribution function for optical substrate and different films. *6th International Symposium on Advanced Optical Manufacturing and Testing Technologies: Optical Test and Measurement Technology and Equipment* (Oct 2012). <https://doi.org/10.1117/12.974317>
- Anat Levin, Daniel Glasner, Ying Xiong, Frédo Durand, William Freeman, Wojciech Matusik, and Todd Zickler. 2013. Fabricating BRDFs at high spatial resolution using wave optics. *ACM Transactions on Graphics* 32, 4 (Jul 2013), 1–14. <https://doi.org/10.1145/2461912.2461981>
- Ju. V. Linnik. 1959. An Information-Theoretic Proof of the Central Limit Theorem with Lindeberg Conditions. *Theory of Probability & Its Applications* 4, 3 (Jan 1959), 288–299. <https://doi.org/10.1137/1104028>
- Leonard Mandel and Emil Wolf. 1995. *Optical coherence and quantum optics*. Cambridge University Press, Cambridge.
- Vidhi Mann and Vipul Rastogi. 2020. FDTD simulation studies on improvement of light absorption in organic solar cells by dielectric nanoparticles. *Optical and Quantum Electronics* 52, 5 (April 2020). <https://doi.org/10.1007/s11082-020-02328-2>
- Francisco F. Medina and Giulio Pozzi. 1990. Spatial coherence of anisotropic and astigmatic sources in interference electron microscopy and holography. , 1027 pages. <https://doi.org/10.1364/josaa.7.001027>
- David A. B. Miller. 2012. All linear optical devices are mode converters. , 23985 pages. <https://doi.org/10.1364/oe.20.023985>
- Michal Mojičik, Tomáš Skřivan, Alexander Wilkie, and Jaroslav Krivánek. 2016. Bidirectional Polarised Light Transport. In *Eurographics Symposium on Rendering - Experimental Ideas & Implementations*, Elmar Eisemann and Eugene Fiume (Eds.). <https://doi.org/10.2312/sre.20161215>
- Hans P. Moravec. 1981. 3D graphics and the wave theory. *Proceedings of the 8th annual conference on Computer graphics and interactive techniques - SIGGRAPH '81* (1981). <https://doi.org/10.1145/800224.806817>
- A. Musbach, G. W. Meyer, F. Reitich, and S. H. Oh. 2013. Full Wave Modelling of Light Propagation and Reflection. *Computer Graphics Forum* 32, 6 (Feb 2013), 24–37. <https://doi.org/10.1111/cgf.12012>
- Roger Newton. 1982. *Scattering theory of waves and particles*. Springer-Verlag, New York.
- Se Baek Oh, Sriram Kashyap, Rohit Garg, Sharat Chandran, and Ramesh Raskar. 2010. Rendering Wave Effects with Augmented Light Field. *Computer Graphics Forum* 29, 2 (May 2010), 507–516. <https://doi.org/10.1111/j.1467-8659.2009.01620.x>
- José Pérez and Razvigor Ossikovski. 2016. *Polarized light and the Mueller matrix approach*. CRC Press, Taylor & Francis Group, Boca Raton, FL. <https://doi.org/10.1201/b19711>
- Richard G Priest and Thomas A Gerner. 2000. *Polarimetric BRDF in the microfacet model: Theory and measurements*. Technical Report. NAVAL RESEARCH LAB WASHINGTON DC.
- J. J. Sakurai and Jim Napolitano. 2021. *Modern quantum mechanics*. Cambridge University Press, Cambridge New York.
- Frank Siewert, Heiner Lammert, and Thomas Zeschke. 2008. *The Nanometer Optical Component Measuring Machine*. Springer Berlin Heidelberg, 193–200. https://doi.org/10.1007/978-3-540-74561-7_11
- Jos Stam. 1999. Diffraction shaders. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques - SIGGRAPH '99*. ACM Press. <https://doi.org/10.1145/311535.311546>

- Shlomi Steinberg. 2019. Analytic Spectral Integration of Birefringence-Induced Iridescence. *Computer Graphics Forum* 38, 4 (Jul 2019), 97–110. <https://doi.org/10.1111/cgf.13774>
- Shlomi Steinberg. 2020. Accurate Rendering of Liquid-Crystals and Inhomogeneous Optically Anisotropic Media. *ACM Transactions on Graphics* 39, 3 (Jun 2020), 1–23. <https://doi.org/10.1145/3381748>
- Shlomi Steinberg and Ling-Qi Yan. 2021a. A Generic Framework for Physical Light Transport. *ACM Transactions on Graphics* 40, 4 (Aug 2021), 1–20. <https://doi.org/10.1145/3450626.3459791>
- Shlomi Steinberg and Ling-Qi Yan. 2021b. Physical Light-Matter Interaction in Hermite-Gauss Space. *ACM Trans. Graph.* 40, 6, Article 283 (dec 2021), 17 pages. <https://doi.org/10.1145/3478513.3480530>
- Shlomi Steinberg and Ling-Qi Yan. 2022. Rendering of Subjective Speckle Formed by Rough Statistical Surfaces. *ACM Trans. Graph.* 41, 1, Article 2 (feb 2022), 23 pages. <https://doi.org/10.1145/3472293>
- John Stover. 2012. *Optical scattering : measurement and analysis*. SPIE Press.
- Antoine Toisoul, Daljit Singh Dhillon, and Abhijeet Ghosh. 2018. Acquiring Spatially Varying Appearance of Printed Holographic Surfaces. *ACM Trans. Graph.* 37, 6, Article 272 (Dec. 2018), 16 pages. <https://doi.org/10.1145/3272127.3275077>
- Antoine Toisoul and Abhijeet Ghosh. 2017. Practical Acquisition and Rendering of Diffraction Effects in Surface Reflectance. *ACM Transactions on Graphics* 36, 5 (Jul 2017), 1–16. <https://doi.org/10.1145/3012001>
- Z. Velinov, S. Werner, and M. B. Hullin. 2018. Real-Time Rendering of Wave-Optical Effects on Scratched Surfaces. *Computer Graphics Forum* 37, 2 (2018), 123–134. <https://doi.org/10.1111/cgf.13347>
- Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction through Rough Surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques* (Grenoble, France) (EGSR '07). Eurographics Association, Goslar, DEU, 195–206.
- A Walther. 1968. Radiometry and coherence. *JOSA* 58, 9 (1968), 1256–1259. <https://doi.org/10.1364/JOSA.58.001256>
- Tao Wang and Daomu Zhao. 2010. Polarization-induced coherence changes of an electromagnetic light wave on scattering. *Optics Letters* 35, 18 (Sep 2010), 3108. <https://doi.org/10.1364/ol.35.003108>
- Craig Warren, Antonios Giannopoulos, and Iraklis Giannakis. 2016. gprMax: Open source software to simulate electromagnetic wave propagation for Ground Penetrating Radar. 209 (Dec 2016).
- Andrea Weidlich and Alexander Wilkie. 2008. Realistic rendering of birefringency in uniaxial crystals. *ACM Transactions on Graphics* 27, 1 (Mar 2008), 1–12. <https://doi.org/10.1145/1330511.1330517>
- Sebastian Werner, Zdravko Velinov, Wenzel Jakob, and Matthias Hullin. 2017. Scratch Iridescence: Wave-Optical Rendering of Diffractive Surface Structure. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 36, 6 (Nov. 2017). <https://doi.org/10.1145/3130800.3130840>
- Alexander Wilkie, Robert F. Tobler, and Werner Purgathofer. 2001. Combined Rendering of Polarization and Fluorescence Effects. *Rendering Techniques 2001* (2001), 197–204. https://doi.org/10.1007/978-3-7091-6242-2_18
- Serge Winitzki. 2003. *Uniform Approximations for Transcendental Functions*. Springer Science and Business Media LLC. https://doi.org/10.1007/3-540-44839-x_82
- Emil Wolf. 2007. *Introduction to the theory of coherence and polarization of light*. Cambridge University Press, Cambridge.
- Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering Specular Microgeometry with Wave Optics. *ACM Trans. Graph.* 37, 4, Article 75 (July 2018), 10 pages. <https://doi.org/10.1145/3197517.3201351>
- Sunkyu Yu, Cheng-Wei Qiu, Yidong Chong, Salvatore Torquato, and Namkyoo Park. 2020. Engineered disorder in photonics. , 226–243 pages. <https://doi.org/10.1038/s41578-020-00263-y>
- Andrew Zangwill. 2013. *Modern electrodynamics*. Cambridge University Press, Cambridge.
- Hanwen Zhang, Chia Wei Hsu, and Owen D. Miller. 2019. Scattering concentration bounds: brightness theorems for waves. , 1321 pages. <https://doi.org/10.1364/optica.6.001321>

A GENERALIZED STOKES PARAMETERS RELATIONS

Given generalized Stokes parameters vectors $\vec{S}^{|\mu|}$, $\vec{L}^{|\mu|}$, i.e. generalized irradiance and radiance, the following relations encode some useful radiometric and polarimetric properties:

- (1) the *observed* (non-instantaneous) spectral irradiance and radiance are

$$S(\vec{r}; \omega) \triangleq S_0^{|\mu|}(\vec{r}, 0; \omega), \quad (36)$$

$$L(\vec{r}; \omega) \triangleq L_0^{|\mu|}(\vec{r}, 0; \omega). \quad (37)$$

- (2) the complex *spectral degree-of-coherence* is

$$\gamma(\vec{r}, \vec{\xi}; \omega) \triangleq \frac{1}{[S(\vec{r}; \omega)S(\vec{r}; \omega)]^{1/2}} S_0^{|\mu|}(\vec{r}, \vec{\xi}; \omega), \quad (38)$$

with $|\gamma| \leq 1$; and,

- (3) the *degree of cross-polarization*, quantifying how polarized the radiation is, can be written as

$$\mu(\vec{r}, \vec{\xi}; \omega) \triangleq |S_0^{|\mu|}|^{-2} \left(|S_1^{|\mu|}|^2 + |S_2^{|\mu|}|^2 + |S_3^{|\mu|}|^2 \right)^{1/2}. \quad (39)$$

The relations in Eqs. (38) and (39) can also be written in an identical form using the generalized radiance $\vec{L}^{|\mu|}$. The last 3 elements of an (irradiance- or radiance-carrying) generalized Stokes parameters vector is a triplet that forms a coordinate on the Poincaré sphere and quantifies the partition of energy into different states of polarization. See Born and Wolf [1999] for additional ellipsometric relations.

B DERIVATION OF DEFINITION 3.1

We outline the derivation that leads to the generalized radiance with anisotropic Gaussian coherence as defined in Definition 3.1.

Let $\alpha, \beta \in \{x, y\}$ index the transverse components. Assume

$$\langle E_\alpha(\vec{r}_1) E_\alpha^*(\vec{r}_2) \rangle_\omega \propto g^{\Theta_\alpha}(\vec{\xi}), \quad (40)$$

i.e. stationary, anisotropic Gaussian spatial coherence with $\vec{\xi}$ being the difference vector between the points \vec{r}_1 and \vec{r}_2 in the local frame. This implies that the cross-transverse correlations are

$$\langle E_\beta(\vec{r}_1) E_\alpha^*(\vec{r}_2) \rangle_\omega \propto g^{\frac{\Theta_\beta + \Theta_\alpha}{2}}(\vec{\xi}). \quad (41)$$

The (ensemble-averaged) spectral radiances are

$$L_\alpha = \frac{d}{d\Omega} \langle E_\alpha(\vec{r}) E_\alpha^*(\vec{r}) \rangle_\omega, \quad (42)$$

with Ω a solid angle and \vec{r} a point on the cross-section of the beam. Then, the structure of $\vec{L}^{|\mu|}$ in Definition 3.1 follows from Eqs. (7) and (8).

The expression $\exp(ik\hat{r} \cdot \vec{\xi}) \equiv \exp(ik\xi_z)$ is the far-field propagator [Korotkova 2017]. The r^2/k^2 factor in the shape matrices is a consequence of the Van Cittert-Zernike theorem [Mandel and Wolf 1995]: the spatial coherence area increases linearly on free-space propagation and decreases with wavelength. See Steinberg and Yan [2021a, Supplemental Section 7.1] for an explicit derivation of the r^2/k^2 factor.

$L_{x,y}$ are real-valued, i.e. we ignore an irrelevant constant phase term. Physical realizability mandates that $\chi^2 + \varsigma^2 \leq 1$.

C PROOF OF THEOREM 3.2

Let the incident electric field be $E(\vec{r}, t) = \sum_j E_j(\vec{r}, t)$, where E_j are the N principal modes of the input. These principal modes are mutually-incoherent but perfectly self-coherent. Formally, this can be written as $\langle E_j(\vec{r}_1, t) E_l(\vec{r}_2, t)^* \rangle = \delta_{jl} a_j(\vec{r}_1, \vec{r}_2)$ with δ_{jl} being the Kronecker delta and a_j being the time-averaged mutual-intensity of mode j . We may write these modes in vector form as

$$\vec{E} \triangleq (E_1 \ E_2 \ \dots \ E_N)^\top. \quad (43)$$

A decomposition into principal mutually-incoherent but perfectly self-coherent modes is always possible, and is simply a restatement

of the fact that Maxwell's equations are linear. See Fig. 4 for a simple example of such decomposition.

Scattering by a linear, passive optical element is then formalised by an operator \mathcal{T} that acts upon $\vec{\mathcal{E}}$. This operator is necessary linear, finite and compact [Miller 2012]. Furthermore, as we neglect all non-optical output modes and, based on our discussion in Subsection 3.4, assume no perfectly absorbed modes, it must follow that $\mathcal{T} \in \mathbb{C}^{N \times N}$ and is full rank. Its singular value decomposition (SVD) is then $\mathcal{T} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\dagger$, with \mathbf{U}, \mathbf{V} unitary and a diagonal matrix of singular values $\mathbf{\Sigma} = \text{diag}\{\zeta_1, \zeta_2, \dots, \zeta_N\}$. Then, the $M_{\text{opt}} = N$ output (optical) modes are then

$$\vec{\mathcal{E}}' \triangleq \mathcal{T}\vec{\mathcal{E}}. \quad (44)$$

The total scattered electric field is then the sum of the scattered modes, viz.

$$E' = \vec{\mathbf{1}}^\top \vec{\mathcal{E}}', \quad (45)$$

where $\vec{\mathbf{1}}$ is the vector of all 1-s.

We understand an optical element that scatters in a mode-agonistic manner as an element where \mathcal{T} admits constant modal coupling coefficients, but the cross-modal coupling can be random. That is, all the singular values are constant $\zeta_j \equiv \zeta \in \mathbb{C}$ and the coupling power quantified by these singular values is simply the differential scattering power cross-section (i.e. the classical BSDF multiplied by a solid angle factor):

$$|\sigma|^2 \equiv \frac{d\Phi^{(o)}}{d\Phi^{(i)}}, \quad (46)$$

with Φ denoting spectral flux. The constant ζ can be understood as the expected value over the entire area of the beam's interaction with the matter. The cross-modal coupling is described by the unitary transformations \mathbf{U}, \mathbf{V} , which are random unitary matrices (circular unitary ensemble members) and are independent of the input modes. That is, $\langle \mathbf{U} \rangle = \mathbf{U}$, with the ensemble average taken over all realizations of the incident field, and similarly for \mathbf{V} .

Consider the cross-spectral density function (as defined in Eq. (7)) of the scattered beam. Directly from Eqs. (44) and (45), it is easy to see that

$$\langle E'(\vec{\mathbf{r}}_1, t) E'(\vec{\mathbf{r}}_2, t)^\star \rangle = \vec{\mathbf{1}}^\top \mathcal{T} \left\langle \vec{\mathcal{E}}(\vec{\mathbf{r}}_1, t) \vec{\mathcal{E}}(\vec{\mathbf{r}}_2, t)^\dagger \right\rangle \mathcal{T}^\dagger \vec{\mathbf{1}}. \quad (47)$$

Denoting the unitary transform $\mathbf{W} = \mathbf{U}\mathbf{V}^\dagger$ and taking the expected value (over all possible unitary couplings \mathbf{W}) of the CSD above, yields:

$$\begin{aligned} E[\langle E'(\vec{\mathbf{r}}_1, t) E'(\vec{\mathbf{r}}_2, t)^\star \rangle] \\ = \frac{1}{|\zeta|^2} \sum_{i,j,k} E[W_{ij} W_{kj}^\star] \langle E_j(\vec{\mathbf{r}}_1, t) E_j(\vec{\mathbf{r}}_2, t)^\star \rangle, \end{aligned} \quad (48)$$

with W_{ij} being the elements of \mathbf{W} . Holding j fixed, the sum over the indices i, k above is a product of the column-sum (of column j) with itself. By the unitarity of \mathbf{W} , the column are complex unit vectors. Using the well-known fact that a random unit vector can be written as a normalized N -variate normally distributed vector, viz. $W_{ij} = \frac{1}{(|X_1|^2 + |X_2|^2 + \dots + |X_N|^2)^{1/2}} X_i$ with $X_i \sim \mathcal{N}_C(0, 1)$ (complex

normal distribution), we write

$$E[W_{ij} W_{kj}^\star] = E\left[\frac{X_i X_k^\star}{|X_1|^2 + |X_2|^2 + \dots + |X_N|^2}\right] \rightarrow \begin{cases} \frac{1}{N} & \text{if } i = k \\ 0 & \text{otherwise} \end{cases}, \quad (49)$$

with rapid convergence as N increases (note that the X depend on the column, i.e. on j). Therefore, Eq. (48) reduces to

$$\begin{aligned} E[\langle E'(\vec{\mathbf{r}}_1, t) E'(\vec{\mathbf{r}}_2, t)^\star \rangle] &= \frac{1}{|\zeta|^2} \sum_{i,j} \frac{1}{N} \langle E_j(\vec{\mathbf{r}}_1, t) E_j(\vec{\mathbf{r}}_2, t)^\star \rangle \\ &= \frac{1}{|\zeta|^2} \left\langle \vec{\mathcal{E}}(\vec{\mathbf{r}}_1, t) \vec{\mathcal{E}}(\vec{\mathbf{r}}_2, t)^\dagger \right\rangle, \end{aligned} \quad (50)$$

where the second equality follows from the definition of $\vec{\mathcal{E}}$. The fact that $\langle E_l(\vec{\mathbf{r}}_1, t) E_n(\vec{\mathbf{r}}_2, t)^\star \rangle = \delta_{ln}$, by the mutual-incoherence of the incident modes, means that $\langle \vec{\mathcal{E}} \vec{\mathcal{E}}^\dagger \rangle = \langle \mathbf{E} \mathbf{E}^\star \rangle$ is the CSD of the incident beam. We conclude that the CSD of the input and the (expected) CSD of the output are related via a constant $|\zeta|^2$ factor. ■

D DECOMPOSITION OF THE PBSDF

In this appendix we discuss the decomposition of a non-singular Mueller matrix, i.e. $\mathbf{M}_{\Delta p} \equiv \mathbf{I}$. For general forward decomposition see Pérez and Ossikovski [2016, Ch. 8.2.2].

By definition, the diattenuator \mathbf{M}_d and retarder \mathbf{M}_r take the following forms:

$$\mathbf{M}_d \triangleq \begin{pmatrix} 1 & \vec{\mathbf{d}}^\top \\ \vec{\mathbf{d}} & D \end{pmatrix} \quad \mathbf{M}_r \triangleq \begin{pmatrix} 1 & \\ & \mathbf{R}_M \end{pmatrix}, \quad (51)$$

where $\vec{\mathbf{d}}$ is the *diattenuation vector* s.t. $|\vec{\mathbf{d}}| \leq 1$ and

$$D \triangleq \sqrt{1 - d^2} (\mathbf{I} - \vec{\mathbf{d}} \vec{\mathbf{d}}^\top) + \vec{\mathbf{d}} \vec{\mathbf{d}}^\top \quad (52)$$

if the *diattenuation matrix*, with $d = |\vec{\mathbf{d}}|$ as usual. Then, an arbitrary non-depolarising Mueller matrix \mathbf{M} can be written as

$$\mathbf{M} = m_{00} \begin{pmatrix} 1 & \vec{\mathbf{d}}^\top \mathbf{R}_M \\ \vec{\mathbf{d}} & \hat{\mathbf{M}}_3 \end{pmatrix}, \quad (53)$$

with $m_{00} \hat{\mathbf{M}}_3$ being the lower-right 3×3 minor of \mathbf{M} . We may immediately extract the mean intensity coefficient m_{00} and the diattenuation vector $\vec{\mathbf{d}}$ from Eq. (53). Then, Eq. (52) yields the matrix \mathbf{D} , and by inverting \mathbf{D} we get

$$\mathbf{R}_M = \mathbf{D}^{-1} \hat{\mathbf{M}}_3, \quad (54)$$

thereby completing the decomposition. The inversion of \mathbf{D} may be performed via consecutive applications of the Sherman-Morrison formula: Let

$$\mathbf{A} \triangleq \frac{1}{\sqrt{1 - d^2}} (\mathbf{I} - \vec{\mathbf{d}} \vec{\mathbf{d}}^\top)^{-1} = \frac{1}{\sqrt{1 - d^2}} \left(\mathbf{I} + \frac{1}{1 - d^2} \vec{\mathbf{d}} \vec{\mathbf{d}}^\top \right). \quad (55)$$

Then, noting that \mathbf{A} is symmetric, let $\vec{\mathbf{a}} = \mathbf{A} \vec{\mathbf{d}}$ and we get

$$\mathbf{D}^{-1} = \mathbf{A} - a^2 \left(\mathbf{1} + \vec{\mathbf{d}}^\top \vec{\mathbf{a}} \right)^{-1}. \quad (56)$$

Computing $\tilde{\mathbf{U}}$. Assuming \mathbf{M} is non-depolarizing, we may rewrite the formula for the Hessian matrix $\tilde{\mathbf{U}}$ (Eqs. (18) and (19)):

$$\begin{aligned} \tilde{U}_{x,y} &= \frac{d^2}{d^2\hat{\mathbf{r}}'} \tilde{\mathbf{S}}_0^\top \left[m_{00} \mathbf{M}_d \right] \Big|_{\hat{\mathbf{s}}, \hat{\mathbf{r}}+\hat{\mathbf{r}}'} \tilde{\mathbf{S}}_{\text{LHP,LVP}} \\ &= \frac{d^2}{d^2\hat{\mathbf{r}}'} \left[m_{00} \pm m_{10} \right] \Big|_{\hat{\mathbf{s}}, \hat{\mathbf{r}}+\hat{\mathbf{r}}'}, \end{aligned} \quad (57)$$

where $m_{10} = m_{00}d_x$ is the element in the second row, first column of \mathbf{M} .

Mueller rotation matrix. Finally, we define the Mueller rotation matrix between a pair of reference frames. Let $\{\hat{\mathbf{x}}, \hat{\mathbf{y}}\}$ and $\{\hat{\mathbf{x}}', \hat{\mathbf{y}}'\}$ be the incident and scattered, respectively, transverse bases, and let

$$\hat{\mathbf{x}}'_\perp \triangleq \frac{\hat{\mathbf{x}}' - \hat{\mathbf{s}}\hat{\mathbf{s}}^\top \hat{\mathbf{x}}'}{|\hat{\mathbf{x}}' - \hat{\mathbf{s}}\hat{\mathbf{s}}^\top \hat{\mathbf{x}}'|} \quad \text{and} \quad \hat{\mathbf{y}}'_\perp \triangleq \frac{\hat{\mathbf{y}}' - \hat{\mathbf{s}}\hat{\mathbf{s}}^\top \hat{\mathbf{y}}'}{|\hat{\mathbf{y}}' - \hat{\mathbf{s}}\hat{\mathbf{s}}^\top \hat{\mathbf{y}}'|} \quad (58)$$

be the normalized projections of the scattered transverse basis onto the incident transverse plane, with $-\hat{\mathbf{s}}$ being the (mean) direction of propagation of the incident beam (i.e., the vector normal to the transverse plane), as before. Let φ_z be the angle between $\hat{\mathbf{x}}'_\perp$ and $\hat{\mathbf{x}}$ (in the singular case where $\hat{\mathbf{x}}'$ is perpendicular to the incident transverse plane, we take the angle between $\hat{\mathbf{y}}'_\perp$ and $\hat{\mathbf{y}}$). That is, φ_z is the rotation around the z -axis encoded by \mathbf{R}_M . Then, the Mueller rotation matrix \mathbf{T} is defined via

$$\mathbf{R}_T \triangleq \begin{pmatrix} \cos 2\varphi_z & -\sin 2\varphi_z \\ \sin 2\varphi_z & \cos 2\varphi_z \end{pmatrix} \quad \text{and} \quad \mathbf{T} \triangleq \text{diag} \{1, \mathbf{R}_T, 1\}. \quad (59)$$

E THE RELATIVE ROUGHNESS AND THE PSD

Let the spatial function $f: \mathbb{R}^3 \rightarrow \mathbb{R}$ quantify the intrinsic scattering features of the matter, for example the height field of a surface or (optical) density of scatterers in a medium. In practice, f may be only known in a limited spatial extent (the scattering region), in which case we simply understand it as a periodic signal. The autocorrelation function of that signal is then

$$r(\vec{\mathbf{d}}) \triangleq \int_{\mathbb{R}^3} d^3\vec{\mathbf{r}} f(\vec{\mathbf{r}} + \vec{\mathbf{d}}) f(\vec{\mathbf{r}}) = (2\pi)^{\frac{3}{2}} [f(-\vec{\mathbf{r}}) * f(\vec{\mathbf{r}})](\vec{\mathbf{d}}), \quad (60)$$

where $\vec{\mathbf{d}}$ is a spatial distance. The above definition is with respect to a deterministic process, if f is in fact a stochastic process, then the typical definition of the autocorrelation function for stochastic processes is to be used. The *root-mean-square (rms) roughness* is then

$$q \triangleq \left(\int_{\mathbb{R}^3} d^3\vec{\mathbf{r}} f(\vec{\mathbf{r}})^2 \right)^{1/2} = \sqrt{r(0)}, \quad (61)$$

while the power spectral density of the signal is

$$p(\vec{\boldsymbol{\zeta}}) \triangleq \mathcal{F} \{r\}(\vec{\boldsymbol{\zeta}}), \quad (62)$$

and it is important to note that the PSD is well defined as a spectral decomposition of the autocorrelation function, even when the Fourier transform does not exist, by the Wiener-Khinchin theorem.

Eq. (61) can be rewritten using Eq. (62) as $q = (2\pi)^{-\frac{3}{4}} (\int p)^{\frac{1}{2}}$. However, note that when evaluating the diffraction operator (defined in Eq. (15)), the frequencies that enter the PSD are $\vec{\mathbf{h}}$ (Eq. (S2.18)) and clearly $|\vec{\mathbf{h}}| \leq 2k$. Spatial frequencies greater than $\sim 2k$ that

arise in the matter then never contribute to the scattered (far-field) radiation, but only produce evanescent waves. Furthermore, the *relative roughness* that is observed from the incident direction $\hat{\mathbf{s}}$ consists only of the frequencies that contribute to scattering, i.e. all the frequencies $\vec{\mathbf{h}} = k(\hat{\mathbf{r}} + \hat{\mathbf{s}})$ with fixed $\hat{\mathbf{s}}$. Hence, the relative roughness is the PSD integrated over the k -sphere centred at $k\hat{\mathbf{s}}$:

$$q_{\text{rel}} \triangleq \left(\frac{1}{2\pi} \right)^{\frac{3}{4}} \left(\int_{\mathcal{S}^2} d^2\hat{\mathbf{r}} p(k\hat{\mathbf{s}} + k\hat{\mathbf{r}}) \right)^{1/2}, \quad (63)$$

with \mathcal{S}^2 denoting the unit sphere.

In the case of a surface, the PSD is a two-dimensional quantity, i.e. $p(\vec{\boldsymbol{\zeta}}) \equiv p(\vec{\boldsymbol{\zeta}}^\perp)$ with $\vec{\boldsymbol{\zeta}}^\perp$ being the projection onto the surface. Therefore, the integral in Eq. (63) reduces in effect to an integral taken twice over the k -circle (up to a Jacobian term) in PSD space centred at $k\hat{\mathbf{s}}$: once for positive and once for negative $\hat{\mathbf{r}} \cdot \hat{\mathbf{n}}$ (with $\hat{\mathbf{n}}$ the normal vector), i.e. once for reflected and once for transmitted scattered contributions. If the surface is taken to be a perfect conductor, then we may assume that transmittance takes no place and integration is performed over the positive hemisphere only. For volumetric scatterers, the relative roughness remains rigorously defined as in Eq. (63) and constitutes of a rather sparse portion of the PSD: a k -sphere in PSD space, with the important conclusion being that the vast majority of frequencies do not contribute to scattering.

Eq. (63) is consistent with modern surface scatter literature, however it ignores the effects of coherence: the PSD is convolved by the spatial coherence of the incident radiation (Eq. (15)). This serves to perturb the incident direction $\hat{\mathbf{s}}$ by a small amount, hence it is more accurate to formulate the relative roughness as the contributions from a *thin spherical shell of the k -sphere in PSD space* with thickness proportional to the spatial coherence of the incident light. Such a coherence-aware formulation of relative roughness is left for future work. We only render surfaces in this work, and, for a surface, integrating over a thin spherical shell changes little, as the integral is projected onto the surface anyhow.