

Physical Light-Matter Interaction in Hermite-Gauss Space

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

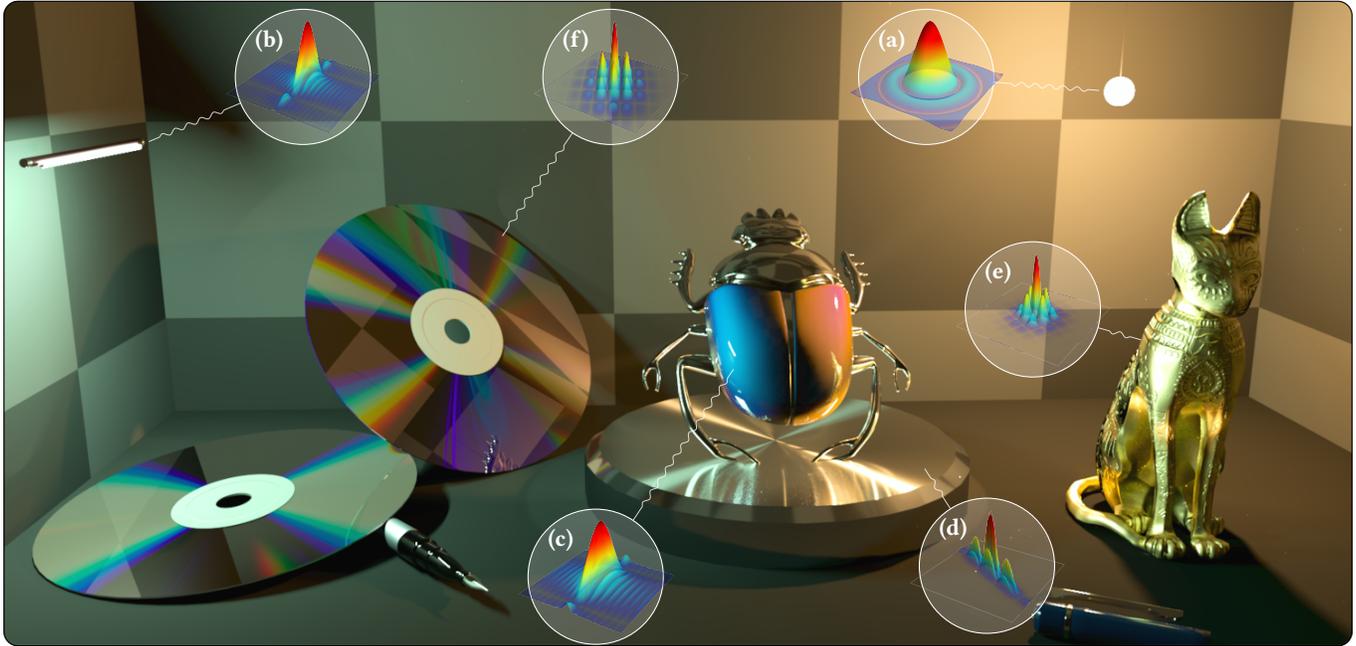


Fig. 1. Partially-coherent light transport in a scene with diffractive materials. The insets visualise the shape of the *coherence area* (on the plane transverse to propagation direction, see Fig. 5) of light that is sourced from a light source or scattered by matter. A spherical source gives rise to light with (a) highly isotropic spatial coherence, while a cylindrical source produces light that is (b) significantly more coherent in one transverse direction than in the other (a phenomenon we term *coherence anisotropy*). Light then propagates away from the source and interacts with matter. These physical processes—the coherence of light and light-matter interaction—are mutually-dependant processes: Coherence drives the optical response of the interaction of light with matter, and, conversely, the properties of matter alter the coherence properties of the scattered radiation. (c) Thin coating over the wings of a silver scarab induces interference. The distinct colours on the left and right wings arise solely due to the difference in the spectral composition and coherence of the incident light. The surface is smooth and the scattered light retains the coherence shape of the incident light. (d,e) On the other hand, scatter by rough surfaces induces coherence properties and anisotropies that are dictated by the surface parameters. (f) Diffraction grating by (unrecorded) DVD disks. Note that the secondary diffraction lobes diminish due to the limited spatial coherence of light. One of the primary theoretical conclusions of this paper is that it is the ensemble-averaged reflectivity of the matter that drives the coherence shape of the scattered light. This can be seen in (c) and (f), where the induced interference does not influence the scattered radiation’s coherence properties.

Our purpose in this paper is two-fold: introduce a computationally-tractable decomposition of the coherence properties of light; and, present a general-purpose light-matter interaction framework for partially-coherent light. In a recent publication, Steinberg and Yan [2021] introduced a framework that generalises the classical radiometry-based light transport to physical optics. This facilitates a qualitative increase in the scope of optical phenomena that can be rendered, however with the additional expressibility comes greater analytic difficulty: This coherence of light, which is the core quantity of physical light transport, depends initially on the characteristics of the light

source, and mutates on interaction with matter and propagation. Furthermore, current tools that aim to quantify the interaction of partially-coherent light with matter remain limited to specific materials and are computationally intensive. To practically represent a wide class of coherence functions, we decompose their modal content in Hermite-Gauss space and derive a set of light-matter interaction formulae, which quantify how matter scatters light and affects its coherence properties. Then, we model matter as a locally-stationary random process, generalizing the prevalent deterministic and stationary stochastic descriptions. This gives rise to a framework that is able to formulate the interaction of arbitrary partially-coherent light with a wide class of matter. Indeed, we will show that our presented formalism unifies a few of the state-of-the-art scatter and diffraction formulae into one cohesive theory. This formulae include the sourcing of partially-coherent light, scatter by rough surfaces and microgeometry, diffraction grating and interference by a layered structure.

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

© 2021 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/10.1145/3478513.3480530>.

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

Additional Key Words and Phrases: rendering, wave, optics, light transport, stochastic, partial, coherence, interference, diffraction, scattering, spectral, polarization, generalized stokes parameters

ACM Reference Format:

Shlomi Steinberg and Ling-Qi Yan. 2021. Physical Light-Matter Interaction in Hermite-Gauss Space. *ACM Trans. Graph.* 40, 6, Article 283 (December 2021), 8 pages. <https://doi.org/10.1145/3478513.3480530>

1 INTRODUCTION

The light that we observe daily consists of very many individual constituents—each an electromagnetic wave. The laws that govern the propagation of each such wave are well understood, however it is the orchestrated behaviour of an *ensemble of all these waves* that gives rise to observable visual phenomena that cannot be adequately explained classically (i.e., radiometrically) nor via a deterministic electromagnetic treatment. The statistical correlation between these constituents over time and space is known as the *optical coherence* of light, and is the primary quantity of interest in the study of the observable properties of light (see Fig. 1). Essentially all the light we employ and observe daily is composed of highly disorganised such constituents, due to the random, spontaneous emission nature of (non-lasing) light sources. Such light is termed *partially-coherent light*.

The limited spatial coherence of light is fundamental to the manner in which surface microgeometry reflects light, the appearance of coated materials and the metallic flakes embedded in some paints. See Fig. 2 for examples of how optical coherence affects material appearance. The importance of optical coherence goes beyond the reproduction of material appearance: radiation with limited coherence is used for optical coherence tomography, and these principles are applied for the decomposition of the light transport in a scene [Gkioulekas et al. 2015; Kotwal et al. 2020]. Partially-coherent optical speckle plays a role in a large variety of applications, e.g., non-line-of-sight imaging with weakly-coherent light [Katz et al. 2014] and diffraction-limited imaging. These applications go beyond optical frequencies: for example, the study of the propagation of radar in complex environment is a problem that can be formulated as rendering with partially-coherent light of non-optical frequencies.

In a recent work, Steinberg and Yan [2021] build upon established foundations in optical coherence theory to introduce a general light transport framework for partially-coherent light. The central quantity is a form of a *cross-spectral density* (CSD) function (to be introduced in Subsection 3.2), which quantifies the statistical similarity between spectral constituents of the wave ensemble at a pair of space points. This two-point information is necessary to transcend the limitation of classical radiometry and geometric optics, and establish a formalism that is consistent with electromagnetism [Wolf 2007]. Unlike traditional radiometry-based light transport, that framework provides the means and formulae needed to describe the *global* interaction of partially-coherent light throughout a scene, and render wave-optics phenomena hitherto not possible. However, as this two-point information is quantified by functions, propagation of light as well as light-matter interaction are formulated in terms of operators acting upon these functions. This requires implementations to either adopt a symbolic treatment or restrict the CSD

to analytic expressions chosen a-priori. But the coherence of light can take a wide variety of shapes (see Fig. 1): The characteristics of both the source and matter affect the coherence properties of light. Indeed, the coherence of light and the scattering of that light by matter are closely-related physical processes. Quantifying the coherence properties of a scattered beam with reasonable accuracy is important, as these properties affect subsequent interactions with matter and the observable properties of light.

Light transport in Hermite-Gauss space. A self-evident approach to represent a large family of functions, while avoiding an impractical symbolic treatment, is to work under a functional basis. The choice of basis is crucial for success. We elect to use a class of *Hermite-Gauss* (HG) functions, known as the *anisotropic HG functions* (formally introduced in Section 3), which endow the univariate HG functions with a linear operator. We term that operator the *shape matrix*, and we will show that this matrix has clear physical meaning: the coherence shape of light. Furthermore, of practical relevance is the fact that this shape matrix greatly expands the family of functions that we can represent with a limited number of basis functions. The HG functions admit highly attractive analytic properties, and in Section 4 we further discuss what makes them such a strong candidate for our functional basis.

Locally-stationary matter. When working with optical frequencies, a purely deterministic description of matter needs to be at a sub-micrometre resolution in order to quantify diffractive features. Often, such a detailed description of matter is infeasible. Instead, it is common to describe matter statistically. Surfaces are described via the spatial frequencies that arise, or a facet distribution, and the interaction of light with participating media may be parametrised by the random spatial distribution of some known types of particles. Such descriptions are convenient, however, a material conclusion emanates from the theory we develop: *Matter features that disappear on ensemble-averaging do not affect the coherence of the scattered beam.* But the ensemble-average of a stationary, at least in the wide-sense, stochastic process is constant by definition! Truly, it is the irregularities and imperfections—such as scratches and flakes—in the matter that produce interesting diffractions.

With that in mind, we present a novel model to describe matter and its interaction with light in a hybrid fashion: as a locally-stationary stochastic process. As will be shown, such a description facilitates the capacity to represent both the micro-scale statistics as well as larger, explicitly-defined features.

A theory of light-matter interaction. This pair of contributions, light transport in Hermite-Gauss space and a locally-stationary model of matter, go hand in hand: As mentioned, the scattering of light by matter and the evolution of a beam’s coherence characteristics are intrinsically connected processes. Therefore, the analytic means we use to model these processes influence both the precision of the optics we may reincarnate and the ease of the analysis.

Our models of choice complement each other well. The primary theoretical contribution (Theorem 5.1) of this paper states that one facet of the matter description—its ensemble-averaged behaviour—fully drives the diffraction process (and thus the coherence properties that emerge in the scattered beam), while another facet—the



Fig. 2. Photographs of (top) a compact disk and (bottom) soap bubbles, illuminated by (left) a very small white LED source that admits a moderate degree of spatial coherence and (right) a larger array of lower intensity LEDs, which, together, produce weakly-coherent radiation. Both sources have similar spectrum. The diffraction grating on the compact disk gives rise to primary diffraction lobes (outlined in orange). In addition, secondary diffraction lobes (outlined in cyan) appear as well under illumination by a more coherent source. The soap bubbles are made of an identical mixture and are akin in chemical composition. Nevertheless, the thin-film interference effects are clearly visible on the left and virtually not at all on the right, indicating the important role optical coherence plays. See our supplemental material where we include videos of these soap bubbles.

spatial correlation—dictates the interference process. This makes sense: an ensemble-average of scatter behaviour over small spatial regions (greater than wavelength) amounts to averaging the induced phase-shifts of the scattered radiation. Thus, as mentioned, if the former is constant, then the latter must annihilate. Consequently, it is the *correlation of spatial features* that gives rise to interference. This has strong practical consequences which we will discuss as we conclude, in Section 6. It will also be shown that our theory unifies some of the contemporary wave optics formulations of scatter into one (see Subsection 5.1), and has comparable computational complexity despite producing richer information: the coherence properties of the scattered light, and not just its intensity.

We start with a short overview of mathematical and optical background in Section 3. We then formulate light transport in Hermite-Gauss space and the theoretical underpinnings of our framework in Section 4. Then, we present our primary contributions in Section 5. Please also see our supplemental material, where we provide complete derivations of some of our theoretical proofs and discuss additional related topics.

2 RELATED WORK

Optical coherence in computer graphics. Steinberg and Yan [2021] have introduced a physical light transport framework that generalises classical light transport using a formalism that is able to account for the wave nature of light. Under that framework, the rendering equation is replaced with the spectral-density transport equation that formalises the transport of cross-spectral density (CSD) functions through the scene. We build upon these foundations to present a theory that is able to practically express a large variety of CSD functions, as well as the interaction of light (quantified by these CSD functions) with matter.

Cuyppers et al. [2012]; Oh et al. [2010] have introduced the Wigner-Ville Spectrum to computer graphics in a bid to replace the classical radiance with a more physical variant. The Wigner-Ville Spectrum—also known as the Wigner Distribution Function (WDF)—is a powerful bilinear space-(spatial) frequency representation of a signal, and is closely related to the CSD as its Fourier transform with respect to spatial separation. Using the fact that the CSD and the WDF are Fourier transform-pairs, it is easy to show that in Hermite-Gauss space the representation of the WDF is (up to the duality of the HG functions) the same as the representation of the CSD. We formally show so in Section S3 in our supplemental material. Noteworthy, this pleasant symmetry between the WDF and CSD in HG space highlights the analytic amenability of the HG representation under spectral decomposition (discussed in Section 4).

Other relevant work consider the optical coherence of light to reproduce a particular wave-optics effect. Different approaches are taken to roughly approximate the spatial area over which radiation remains coherent (the coherence area of light): A Gaussian kernel is used in order to render diffractive scattering of light by explicitly-defined surface microgeometry [Falster et al. 2020; Yan et al. 2018]; render diffractive scratches [Werner et al. 2017]; render biological diffraction grating surfaces [Dhillon et al. 2014]; and measure diffraction patterns of holographic surfaces using Jones calculus [Toisoul et al. 2018]. Alternatively, Levin et al. [2013] use simple square-pulse (box function) kernels for the synthesis of high-resolution, spatially-varying BRDFs.

Coherence modulation. Over the years, a substantial amount of optical work has studied the influence of different optical elements, surfaces and scattering media on the coherence properties of the scattered light. The following list is far from exhaustive.

The important work by Wolf et al. [1989] outlines the theory for scattering from random media within the accuracy of the first-order Born approximation, with closed-form expressions for the spectral density on scatter from a perfectly homogenous, stationary medium provided. A large body of follow-up work continue the study in different directions: e.g., predictable manipulation of coherence via scattering by specifically-designed statistical media [Pan et al. 2020], conditions on beam retention [Lahiri and Wolf 2009], and coherence changes of structured beams when scattered by deterministic media [Wang et al. 2016]. Korotkova and Yao [2020] explore the propagation of partially-coherent radiation through oceanic turbulence. Their approach quantifies the linear turbulent medium in terms of its power spectra and considers how these statistics affect the CSD. Gbur [2014] provide a review on the similar problem of studying

the effects of atmospheric turbulence on the statistics of propagating light. By rewriting the propagation integral through an ABCD optical system as a convolution, Ma et al. [2019] discuss a simple method for efficient numeric simulations of the CSD through that system. The influence of a rough surface spatial light modulator, which induces random phase modulations, on the coherence and polarization properties of light is investigated by Ma et al. [2015]. Their formalism employs the cross-spectral density matrix and ABCD optical elements, which we cover in our supplemental material. The effects of scattering by fluids with Brownian motion particles is explored [Popov and Vitkin 2016] by also employing the ABCD formalism. Betancur and Castañeda [2008] present a method for designing spatial modulator for preprescribed power distributions. All these works rely on the Huygens-Fresnel integral for free-space propagation, and (explicitly or implicitly) on the first-order Born approximation.

The work cited above shares some of the motivational basis for this paper in analysing the coherence of the scattered radiation. However, all these work consider specific types of matter, which either can be represented as simple ABCD systems or are constrained to purely statistical matter models. Furthermore, they are often restricted to Gaussian Schell-model sources or the methods rely on numeric integration of computationally non-trivial integrals. Our work focuses on computability and presents an expressive, novel locally-stationary matter model, that can describe both statistical and deterministic scattering features.

Appearance reproduction with wave optics. Rendering different diffraction phenomena has been a very active area of research in computer graphics. Some work aims to solve the full electromagnetic problem with coherent fields. The finite-difference time-domain (FDTD) method is used to formulate a solution to Maxwell's equations for the nanofabrication of materials with pigment-free structural colours [Auzinger et al. 2018], and numerically solving a light propagation problem in a simple scene [Musbach et al. 2013]. An approximate analytic solution to a scattering differential system is used to render liquid-crystal micrographs [Steinberg 2020].

Frequently, in addition to ignoring optical coherence, a simplified optical formalism ("wave optics") is employed in-place of full electromagnetism. Stam [1999] is of the first to present practical work that considers the wave nature of light. The large body of work that followed includes rendering thin-film interference [Belcour and Barla 2017; Kneiphof et al. 2019]; rendering soap bubbles [Huang et al. 2020]; rendering iridescent, pearlescent materials [Guilén et al. 2020]; formulating diffraction-aware BSDFs [Toisoul and Ghosh 2017; Velinov et al. 2018]; rendering of birefringent dielectrics [Steinberg 2019]; rendering effects that arise due to scattering by non-spherical particles, like rainbows [Sadeghi et al. 2012]; and, rendering optical speckle fields that arise on scattering of coherent illumination by random media [Bar et al. 2020]. Also of relevance is work that employs surface scatter theories in computer graphics [Holzschuch and Pacanowski 2017; Steinberg and Yan 2021a; Yan et al. 2018].

While ignoring optical coherence holds merit in many cases where the diffractive matter features are tiny, as we will show in

Section 6 caveats should be applied and coherence may play a role in virtually any diffraction effect.

Applications of the Hermite-Gauss functions. Takemura and Takeuchi [1988]; Ismail and Simeonov [2020] have studied some of the properties of the anisotropic Hermite-Gauss (Eq. (5)) functions, but otherwise these functions have seen virtually no application and been almost entirely ignored. The standard HG functions, on the other hand, have been employed by a great body of work, too numerous to exhaustively list here, for a variety of applications. The fact that the HG functions are the eigenfunctions of the fractional and standard Fourier transform has been used for image filtering, compression and encryption [Kang et al. 2015; Papari et al. 2012; van Dijk and Martens 1997].

It should be noted that the formalism that is presented in this paper bears no relation to these Hermite-Gauss beams that are common in optics: it is the modal content of the CSD (and not the beam!) that interests us.

3 THEORETICAL FOUNDATIONS

3.1 Preliminaries

The purpose of this subsection is to briefly introduce the reader to our notation and the mathematical tools central to this paper.

We denote the field of real numbers as \mathbb{R} , the complex plane as \mathbb{C} and the set of natural (non-negative) integers as \mathbb{N} . The Cartesian three-dimensional vector space then becomes \mathbb{R}^3 , and, similarly, the set of $n \times m$ real matrices is $\mathbb{R}^{n \times m}$. Our vector notation convention throughout the paper is as follows: Arbitrary spatial vectors are written with an arrow accent, $\vec{r} \in \mathbb{R}^3$. Vectors with a circumflex \hat{r} denote the unit vector in the direction of \vec{r} , while the magnitude of \vec{r} is denoted via the scalar sharing the same letter, viz. $r = |\vec{r}|$. The unit vectors that span the standard Cartesian basis of \mathbb{R}^3 are $\hat{x}, \hat{y}, \hat{z}$, and the Cartesian components of a vector \vec{r} are denoted as r_x, r_y, r_z .

A real, square matrix $A \in \mathbb{R}^{n \times n}$ is said to be *positive-definite* if $\vec{r}^\top A \vec{r} > 0$ for non-zero $\vec{r} \in \mathbb{R}^n$. Given a positive-definite matrix $A \in \mathbb{R}^{n \times n}$, we denote a *matrix square root* of A as a matrix $A^{1/2}$ that satisfies $A = A^{1/2} (A^{1/2})^\top$. Note that, as $A > 0$, such a square root always exists (but, in general, is not unique), and can be found via, e.g., the Cholesky factorization. In addition, our convention is that $A^{-1/2} = (A^{1/2})^{-1}$, i.e. the inverse of the square root. Given $A > 0$, a simple factorization argument can be used to show that $(A^{1/2})^{-\top} = (A^{-1})^{1/2}$.

Given a pair of L^2 functions, $f, g : \mathbb{R}^3 \rightarrow \mathbb{C}$, we write their inner product as

$$\langle f | g \rangle \triangleq \int_{\mathbb{R}^3} d^3 \vec{r}' f(\vec{r}') g^\star(\vec{r}'), \quad (1)$$

where the operator \star denotes complex conjugation and note that our notation reserves primed variables for local or integration variables.

Fourier analysis. We denote the spatial *Fourier transform* (FT) operator, with non-unitary angular-frequency kernels, as \mathcal{F} and its inverse as \mathcal{F}^{-1} . Let $f : \mathbb{R}^3 \rightarrow \mathbb{C}$ be an L^2 function, then the FT of f is

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

As a special case we define the *angular correlation transform* (ACT) operator, denoted \mathcal{A} , to be the Fourier transform of a function of two spatial variables, $F : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{C}$, with the second Fourier kernel conjugated, viz.

$$\mathcal{A}\{F\}(\vec{\zeta}_1, \vec{\zeta}_2) \triangleq \int_{\mathbb{R}^3} d^3\vec{r}'_1 \int_{\mathbb{R}^3} d^3\vec{r}'_2 F(\vec{r}'_1, \vec{r}'_2) e^{-i\vec{r}'_1 \cdot \vec{\zeta}_1} e^{i\vec{r}'_2 \cdot \vec{\zeta}_2}. \quad (3)$$

The definition of the ACT operator is motivated by the definition of the cross-spectral density function (see Eq. (10)).

The convolution operator arises from the FT via the well-known *convolution theorem*:

$$f * g \triangleq \mathcal{F}^{-1}\{\mathcal{F}\{f\}\mathcal{F}\{g\}\}. \quad (4)$$

Anisotropic Hermite-Gauss functions. We denote the n -th degree (univariate) *Hermite-Gauss* (HG) function as $\Psi_n(x)$, for a natural integer $n \in \mathbb{N}$. See Section S1 in our supplemental material for explicit formulae for the Ψ_n and a few of their relevant properties (further, see Szegő [1939] for a more comprehensive discussion). We would like to extend the nice properties admitted by the univariate HG functions to the three-dimensional space. To that end, we use a

Table 1. List of symbols and notation (location of definition on the right)

NOTATION AND SYMBOLS	
Notation	
\vec{r}	Vectors: arrow accented boldface, lower-case 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
\mathbf{A}	Matrices: boldface, capital letters
\mathcal{L}	Operators: typically capital latin letters in script typeface
Symbols	
c	The speed of light
λ	Wavelength (Page 6)
$k = \frac{2\pi}{\lambda}$	Wavenumber (Page 6)
$\omega = ck$	Wave's angular frequency (Page 6)
\mathcal{C}	Cross-spectral density (CSD) function of a wave ensemble (Eq. (10))
Ψ_{nml}^Θ	Anisotropic Hermite-Gauss function, of degree (n, m, l) and order $\mu = n + m + l$ (Eq. (5))
σ	Scattering amplitude function (Page 2)
Σ	Scattering mutual intensity function (Theorem 3.1)
$R_{\sigma\sigma}$	Scattering stationary autocorrelation function (Eq. (12))
\tilde{f}_σ	Scattering angular coherence function (Eq. (13))
Operators	
$\delta(x)$	Dirac delta
δ_{nm}	Kronecker delta
$\langle \cdot \cdot \rangle$	Inner product over L_2 -functions space (Eq. (1))
$\mathcal{F}\{\cdot\}$	Fourier transform (FT) operator (Eq. (2))
$\mathcal{A}\{\cdot\}$	Angular correlation transform (ACT) operator (Eq. (3))
$*$	Convolution operator (Eq. (4))
Re, Im	Real and imaginary value operators
z^*	Complex conjugate of $z \in \mathbb{C}$
$\mathbf{A}^\top, \mathbf{A}^\dagger$	Transpose and conjugate transpose of \mathbf{A}
$\langle \cdot \rangle$	Ensemble-averaging operator (Page 10)
$\langle \cdot \rangle_\omega$	Ensemble-average over same-frequency constituents (Page 6)

form of *anisotropic HG functions*, similar to the functions introduced by Takemura and Takeuchi [1988], which are defined with respect to some real positive-definite matrix Θ —the *shape matrix*—as follows:

$$\Psi_{nml}^\Theta(\vec{r}) \triangleq \frac{(-1)^\mu}{\pi^{\frac{3}{4}} 2^{\frac{\mu}{2}} \sqrt{m! n! l!}} \frac{e^{\frac{1}{2}\vec{r}^\top \Theta^{-1} \vec{r}}}{|\Theta|^{\frac{1}{4}}} \frac{\partial^\mu}{\partial r_x^n \partial r_y^m \partial r_z^l} e^{-\vec{r}^\top \Theta^{-1} \vec{r}}, \quad (5)$$

with *degree* $(n, m, l) \in \mathbb{N}^3$ (an ordered 3-tuple of natural integers), *order* $\mu = n + m + l$ and where r_x, r_y, r_z are the Cartesian components of \vec{r} . To complete the definition we also introduce the *dual* of the anisotropic HG function:

$$\tilde{\Psi}_{nml}^\Theta(\vec{r}) \triangleq |\Theta|^{-1/2} \Psi_{nml}^{\Theta^{-1}}(\Theta^{-1} \vec{r}). \quad (6)$$

Note that the partial derivatives that arise in the dual function are taken with respect to the components of the transformed vector.

The anisotropic HG functions are central to our discussion and their relevant properties are derived in Section S1. Some of these introduced properties are, to the best of our knowledge, novel. As the anisotropic HG functions form a complete basis of L^2 functions (see Section S1 for details), any L^2 function $f : \mathbb{R}^3 \rightarrow \mathbb{C}$ can be (uniquely) expanded under the *Hermite-Gauss basis* as

$$f(\vec{r}) = \sum_{n,m,l=0}^{\infty} \check{f}_{nml} \Psi_{nml}^\Theta(\vec{r}), \quad (7)$$

where the $\check{f}_{nml} \in \mathbb{C}$ are the *Hermite-Gauss coefficients* of degree (n, m, l) and order $n + m + l$, and our notation reserves inverted-hat accents for such HG coefficients. The HG coefficients can be recovered from f via the inner product with the dual function:

$$\check{f}_{nml} = \langle f | \tilde{\Psi}_{nml}^\Theta \rangle = \int_{\mathbb{R}^3} d^3\vec{r}' f(\vec{r}') \tilde{\Psi}_{nml}^\Theta(\vec{r}'). \quad (8)$$

In our supplemental material, we prove the important fact that the anisotropic HG functions and their dual serve as Fourier-transform pairs:

$$\mathcal{F}\{\Psi_{nml}^\Theta(\vec{r}')\}(\vec{\zeta}) = (2\pi)^{\frac{3}{2}} (-i)^{n+m+l} \tilde{\Psi}_{nml}^{\Theta^{-1}}(\vec{\zeta}). \quad (9)$$

As an aside, note that the anisotropic Hermite-Gauss are more expressive than the univariate ones: The partial derivatives in Eq. (5) are with respect to \vec{r} before the transformation by the shape matrix Θ , therefore Ψ_{nml}^Θ can be written as a simple product of univariate HG functions (Eq. (S1.1) in our supplemental material) when $\Theta = \mathbf{I}$, but not in general.

3.2 The Cross-Spectral Density

The study of optical coherence is the study of the statistical properties of light. Our sensors (a camera or the eye) do not observe the individual oscillations of the light's underlying electromagnetic fields, but only the time-averaged values. Therefore, for observable diffractive optical phenomena (like the interference patterns visible in Fig. 2) to arise, the light's waveform must remain statistically correlated over sufficiently large distances in spacetime. The optical coherence of light can then be understood as the correlation between the light's wave constituents over these spacetime regions. In order to formally describe that correlation, we use a function of a pair of points—the cross-spectral density—that quantifies the statistical similarity of waveforms that arrive at these points. This

subsection serves as a formal introduction of a wave ensemble and its cross-spectral density function. See Goodman [2015]; Wolf [2007] for a more in-depth introduction to optical coherence, as well as Steinberg and Yan [2021] for a comprehensive discussion of optical coherence under the context of physical light transport.

Let $u(\vec{r}, t)$ be a function describing the strength of the electromagnetic disturbances in spacetime. u is usually understood not as a deterministic function, but as a stochastic (random) process—a statistical *ensemble of waves*—reflecting the fact that light typically admits random fluctuations in its fields. Then, the central quantity of interest is the *cross-spectral density* (CSD) function of that stochastic process. The CSD is a spatial autocorrelation function, which quantifies the second-order statistics between spectral constituents of u (that is, between waves of an identical frequency in the ensemble), and is defined as the positional power spectrum:

$$C(\vec{r}; \vec{\xi}_1, \vec{\xi}_2; \omega) \triangleq \left\langle u(\vec{r} + \vec{\xi}_1) u^*(\vec{r} + \vec{\xi}_2) \right\rangle_{\omega}, \quad (10)$$

where the vectors $\vec{\xi}_{1,2}$ describe small offsets around \vec{r} (the point around which the CSD is defined). The operator $\langle \cdot \rangle_{\omega}$ is the *ensemble-average* operator (not to be confused with the inner product!) over the same-frequency constituents of the wave ensemble. $\omega = ck$ specifies the angular frequency of interest, with k being the *wavenumber*, which is related to the wavelength λ as $k = 2\pi/\lambda$. The (time-averaged) spectral intensity, that is observed at a point \vec{r} , is proportional to the complex magnitude squared of the underlying field:

$$I(\vec{r}; \omega) \triangleq \left\langle |u(\vec{r})|^2 \right\rangle_{\omega} = C(\vec{r}; 0, 0; \omega). \quad (11)$$

When dealing with partially-coherent light, C decays rapidly with increasing distance between the points at which we evaluate it (indicating loss of spatial coherence). It therefore makes practical sense to consider the CSD as the autocorrelation between points in some confined region in space. This is reflected in our notation: we arrange the arguments of C into groups. First, \vec{r} defines the position in space around which the CSD is defined, and, from the perspective of our analysis, we fix that position and treat it as a constant. The second group describes spatial perturbations around \vec{r} , with respect to which we evaluate the CSD. Finally, the last group only serves to make explicit the wavelength-dependence and the fact that we ensemble-average over same-frequency constituents. For brevity, we will at times neglect the first and third groups, however the reader should note the implied dependence on \vec{r} and ω .

Validity domain. Eq. (11) holds whenever the ensemble-average that appears in that equation is a sufficiently good approximation to a time-average over timescales that are large with respect to the temporal fluctuations of the wave ensemble. This is true when: (i) the wave ensemble is an ergodic stochastic process, and then ensemble-averaging is equivalent to infinite time-averaging; or, (ii) the radiation is polychromatic: in which case the temporal coherence is much shorter than the observation time, and second-order statistics are sufficient [Goodman 2015].

We clearly state the assumptions on the radiation's coherence when developing the formulae in Subsection 5.1, and due to the scale invariance of these formulae, the majority of our theory applies to light of any (spatial or temporal) coherence. Nevertheless, essentially

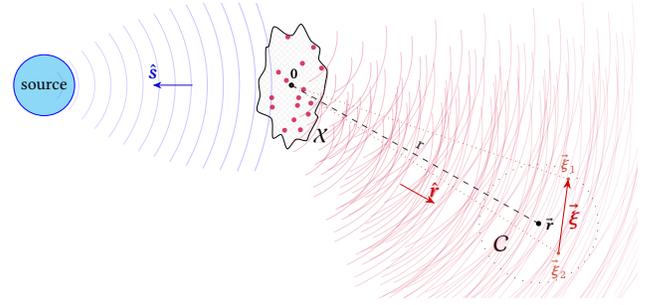


Fig. 3. Diffraction of natural light by a matter: A small spatial region X , centred at the origin 0 , scatters light incident from a source at point \vec{s} (wavefronts illustrated as blue lines). The incident radiation induces electromagnetic excitations in X (small purple balls), giving rise to scattered radiation that propagates away from the scattering region to some point \vec{r} (wavefronts illustrated as red lines). The scattering process induces changes in the statistical properties of the scattered radiation. A solution to the diffraction problem is the cross-spectral density function of the diffracted radiation, C , which is defined at a region around the point \vec{r} (black dotted circle) and quantifies the coherence properties and the statistical similarity of the wave ensemble at a pair of points $\vec{\xi}_1$ and $\vec{\xi}_2$ in that region.

all the emitters of interest (be it optical sources, like the Sun and artificial lights, or non-optical: for example, radar) produce partially-coherent radiation. Furthermore, the rendering equations of physical light transport, i.e. the *spectral-density transport equation* [Steinberg and Yan 2021], assume partially-coherent light, and hence our focus is limited to these kind of sources and such light.

3.3 Scattering and Diffraction

Our theory of light-matter interaction is formulated under the context of scalar diffraction theory. See Fig. 3 for an illustration of the geometry of a diffraction problem. Let the matter be confined to a scattering region $X \subset \mathbb{R}^3$. The core quantity that describes the matter response to incident radiation is the *scattering amplitude function*, denoted $\sigma(\hat{s}', \hat{r}'; \vec{p}')$, which describes the amplitude ratio of the scattered wave in direction \hat{r}' , produced due to a scattering event in position \vec{p}' in the matter or medium and under excitation by an incident wave from direction \hat{s}' . The scattering amplitude function is, in general, wavelength dependant and can take complex values, quantifying the phase-shifts induced by interaction with conducting particles.

We are now ready to derive a formula for the CSD function of the radiation diffracted by the scattering matter. To make analytic progress we first make the Born first-order approximation, as well as the usual Fraunhofer (optical far field) region assumption, implying that the characteristic length of that scattering region is small with respect to the distances to \vec{r} (the observation point) and \vec{s} (the radiation source that gives rise to u'). The analysis then proceeds in typical Fraunhofer region fashion, see Subsection S2.1 in our supplemental material for an extended discussion on scalar scattering, the assumptions involved and proof of Theorem 3.1 that follows.

Then, the formula for the diffracted and propagated CSD under our formalism of light-matter interaction becomes:

THEOREM 3.1 (SCATTERING OF THE CROSS-SPECTRAL DENSITY). *Let a scattering region (assumed to be centred around the origin) be described via a scattering amplitude function σ , and C' be the CSD of the radiation incident to the scattering region from direction \hat{s} . Under the Born first-order approximation and in the Fraunhofer region, the scattered CSD becomes*

$$C(\vec{r}; \vec{\xi}_1, \vec{\xi}_2; \omega) = \frac{e^{ik\hat{r} \cdot (\vec{\xi}_1 - \vec{\xi}_2)}}{\lambda^2 r^2} \mathcal{A} \left\{ \Sigma \cdot C' \right\} \left(k\hat{r} + \frac{k}{r} \vec{\xi}_1, k\hat{r} + \frac{k}{r} \vec{\xi}_2 \right),$$

where \mathcal{A} is the angular correlation transform operator (Eq. (3)) and we define $\Sigma(\vec{\xi}_1, \vec{\xi}_2) = \sigma(\vec{\xi}_1) \sigma^*(\vec{\xi}_2)$ as the scattering mutual intensity.

Note that \vec{r} , \vec{s} and ω are held constant and the integration of the ACT in Theorem 3.1 is done over the variables $\vec{\xi}_{1,2}$ of Σ and C' . It is assumed that the scattering amplitude function σ is a slow function of the incident and exitant directions, thus in the far-field those directions are held constant and are simply \hat{s} and \hat{r} , respectively. For brevity, we also drop the explicit dependence on these directions from the notation of σ and Σ .

4 LIGHT TRANSPORT IN HERMITE-GAUSS SPACE

Finding solutions to the relation in Theorem 3.1 is, in general, difficult: Numerically, the six-dimensional Fourier transform is expensive to compute and, when working with optical frequencies, accurate results require the computation to be done at a fine resolution. Analytic progress is frustrated by the fact that the shape of the incident CSD function, C' , depends on the properties of the emitter as well as the intermediate interactions of the radiation.

It has been shown [Steinberg and Yan 2021] that in the far-field region the CSD behaves radially (in the direction of the wave ensemble's propagation, \hat{r}) as a spherical wave, while the spatial coherence

of the wave ensemble is only a function of the directions from the source to the points $\vec{r} + \vec{\xi}_1$ and $\vec{r} + \vec{\xi}_2$. As a consequence, we also conclude that in the far-field the coherence properties of the light change slowly around the point of interest \vec{r} , and therefore the coherence properties are simply a function of the difference vector between the points $\vec{r} + \vec{\xi}_1$ and $\vec{r} + \vec{\xi}_2$ (i.e. a *quasi-homogeneous* source), denoted as $\vec{\xi} = \vec{\xi}_1 - \vec{\xi}_2$. See Fig. 3 for an illustration of diffraction by matter, and Fig. 5 for a visualization of the coherence properties of light.

Our choice of basis. As mentioned, we use a class of multivariate Hermite-Gauss functions as a functional basis. The motivation is as follows:

- (1) these functions form a complete, orthonormal basis;
- (2) CSDs will typically be oscillatory functions (as predicted by the Fourier optics relation between the incident and scattered CSDs), and so are the HG functions;
- (3) the shape matrix quantifies the spatial coherence of light;
- (4) the HG functions are a simple product of a Gaussian with a polynomial, making them easy to evaluate; and
- (5) the HG functions serve as eigenfunctions under the Fourier transform operator.

The last couple of points are crucial. The Fourier optics relation exhibited by Theorem 3.1 makes the Hermite-Gauss functions a compelling choice: It is this analytic amenability of these functions under the Fourier transform that allows us to make analytic progress when working in Hermite-Gauss space. In addition, a decomposition of the CSD under the HG basis is not only mathematically attractive; it also manifests the underlying physics via the shape matrix, which

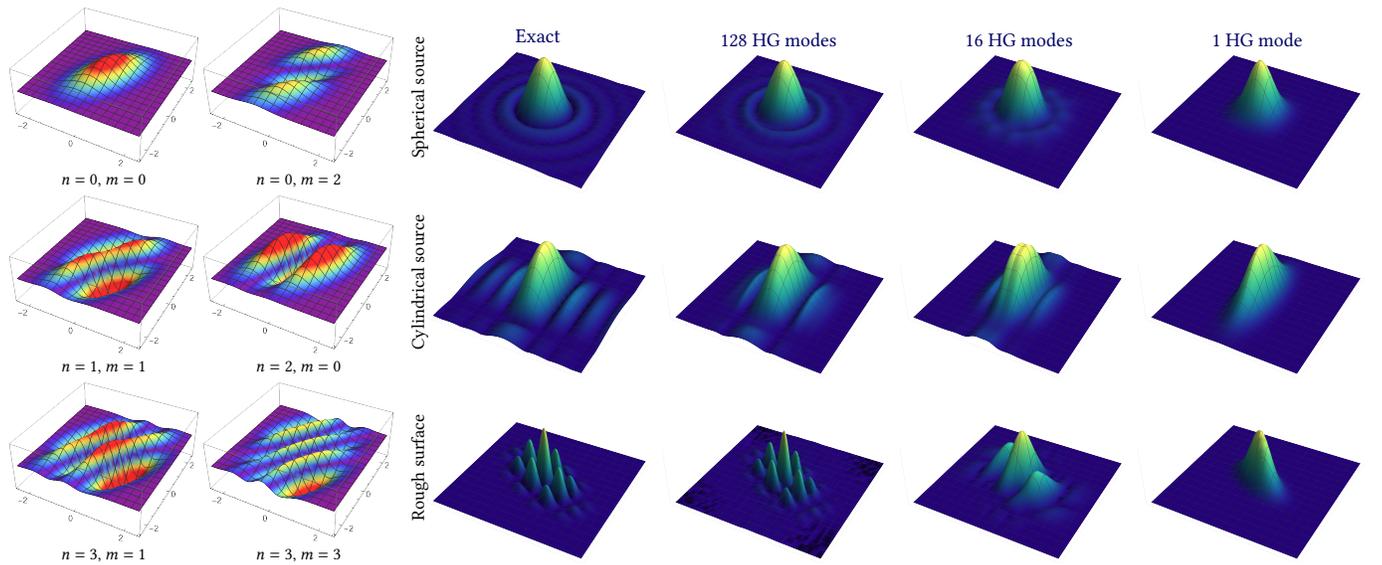


Fig. 4. (left) A few even-ordered anisotropic Hermite-Gauss transverse modes. The shape matrix Θ of the presented modes encodes a rotation and an anisotropic scale, i.e. $\Theta^{-1} = \mathbf{R}\mathbf{S}(\mathbf{R}\mathbf{S})^\top$, where \mathbf{R} is a rotation matrix of $\frac{\pi}{7}$ radians (around \hat{z}) and \mathbf{S} is a diagonal scale matrix with a scale of 2 in \hat{x} and 1 in \hat{y}, \hat{z} . (right) The HG modes form a functional basis that is able to express a wide variety of CSD functions using a limited number of modes. Observe that as the shape matrix encodes the coherence shape and area, even with a single mode we can capture some of the important characteristics of the CSD.

directly quantifies the first-order geometric properties of the spatial coherence (see Corollary 4.2).

Curiously, the anisotropic HG functions have received virtually no attention in literature. This can be explained by noting that when viewed as a functional basis, these functions give rise to an underdetermined system: given any choice of the shape matrix, the anisotropic HG functions still constitute a complete orthonormal basis. This mathematical awkwardness does not deter us, because in practice only a limited, finite count of HG functions may be used. See Fig. 4 for a visualization of the first few HG basis functions and a few CSD functions in HG space.

The CSD in Hermite-Gauss space. Inline with the discussion above, we can now write the general expression for the CSD in HG space, as follows:

Definition 4.1 (Far-field CSD with transverse Hermite-Gauss modes). Given quasi-homogeneous radiation that originates from the origin and has propagated to a point \vec{r} , its CSD takes the following analytic form:

$$C(\vec{r}; \vec{\xi}; \omega) = \frac{e^{ik\hat{r}\cdot\vec{\xi}}}{r^2} \sum_{n,m} \check{c}_{nm} \Psi_{nm}^{\Theta} \left(\frac{k}{r} Q \vec{\xi} \right),$$

with Q being a constant, orthogonal, basis change matrix that orients $\vec{\xi}$ to the *local frame* of the wave ensemble, i.e. where \hat{z} points in direction of propagation (which is \hat{r}) and \hat{x}, \hat{y} span the transverse plane (see Fig. 5).

Note that in-place of the pair of spatial parameters, $\vec{\xi}_{1,2}$, in Eq. (10), we now only have a single parameter $\vec{\xi} = \xi_1 - \xi_2$, i.e. the difference between the spatial points where we evaluate C . We slightly abuse notation in Definition 4.1 by writing $\Psi_{nm}^{\Theta} \equiv \Psi_{nm0}^{\Theta}$, that is, the order of the z component of the HG function (l in Eq. (5)) is always 0 and is dropped from the notation. Under the paraxial far-field approximation, the coherence properties remain constant in z , and extend to infinity. However, this fact is merely an artefact of the Fraunhofer region approximation. The z -dependant 0th-order HG function (a Gaussian) serves to confine the region over which C describes coherent radiation to a finite extent, but the exact shape in the z direction is of little importance. Higher-order HG functions are then reserved only for the transverse directions, and we refer to each $\check{c}_{nm} \neq 0$ as a *Hermite-Gauss transverse mode* of the CSD.

The shape matrix Θ serves two purposes: First, it scales the space of the local frame to match the extent over which the radiation remains spatially coherent. Second, it applies some (invertible) transformation to the transverse x, y components of $\vec{\xi}$. The purpose of that transform is to describe the shape of the transverse area over which the wave ensemble exhibits spatial coherence. It may encode space scaling and rotation, and, more interestingly, anisotropic scales, shears, etc. This greatly increases the expressiveness of the CSD, especially when, for practical reasons, we restrict ourselves to few HG transverse modes (see Fig. 4). The shape matrix together with the HG transverse modes fully describe the CSD, and the rest of this section is dedicated to studying these parameters.

The coherence properties. The coherence area, length and volume [Goodman 2015] quantify the geometric regions over which the

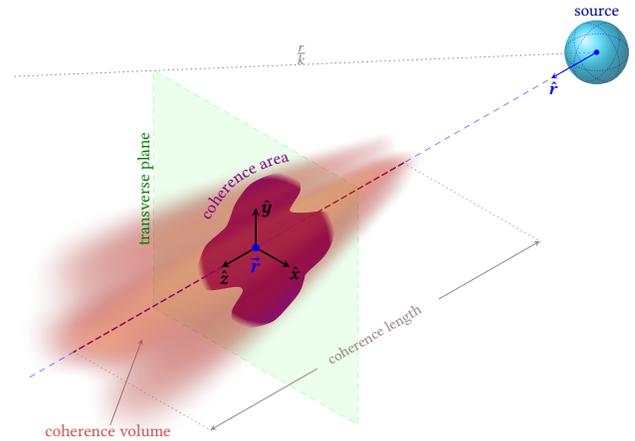


Fig. 5. The coherence properties of a beam of light are core to our discussion: The spatial volume over which light, that has propagated to a point \vec{r} , exhibits optical coherence is illustrated by the red volume and is the *coherence volume*. The distance on the propagation axis (in direction \hat{r}), over which the light remains coherent is known as the *coherence length*. This coherence length is dictated by the spectral content of the light: polychromatic light admits very short coherence length, while filtered (pseudo-monochromatic) light has longer coherence length. The plane that is perpendicular to the direction of propagation is the *transverse plane*, and the intersection of the coherence volume with that plane gives rise to the *coherence area* (violet area). The coherence area depends on the properties of the source or the scattering matter, and may exhibit distinct shapes and anisotropies. Observe that the coherence area grows linearly, as $\frac{r}{k}$, on free-space propagation. The beam's local frame, spanned by $\hat{x}, \hat{y}, \hat{z}$ is illustrated as well.

radiation exhibits coherence and are illustrated in Fig. 5. As previously alluded to, the shape matrix Θ plays a core role in quantifying these quantities: It can be understood as the first-order geometric properties of the beam's coherence, viz.

COROLLARY 4.2 (COHERENCE PROPERTIES). Given the CSD, formalised using transverse HG modes (Definition 4.1), of a beam of light, the geometric coherence properties of the light are quantified:

- (1) The coherence volume of the beam is $\sim \frac{4\pi}{3} \left| \frac{2r}{k} \Theta^{1/2} \right|$.
- (2) The coherence area is the area of the circle transformed by the (rank-2) matrix $\frac{2r}{k} \Theta^{1/2} \mathbf{P}_{\perp}$, where $\mathbf{P}_{\perp} = \text{diag}\{1, 1, 0\}$ projects onto the transverse plane. Explicitly, this can be written as $\sim \pi \left(\frac{2r}{k} \right)^2 \|\Theta^{1/2} \mathbf{P}_{\perp}\|_2 \|\Theta^{-1/2} \mathbf{P}_{\perp}\|_2^{-1}$.
- (3) Similarly, the coherence length is $\sim \left\| \frac{4r}{k} \Theta^{1/2} \mathbf{P}_{\parallel} \right\|_2$, with the projection $\mathbf{P}_{\parallel} = \text{diag}\{0, 0, 1\}$ onto the direction of propagation.
- (4) Coherence anisotropy (in the transverse plane) is characterised by the matrix condition number $\text{cond}_2(\Theta^{1/2} \mathbf{P}_{\perp})$, i.e. the ratio between the two non-zero singular values of $\Theta^{1/2} \mathbf{P}_{\perp}$.

PROOF. The Gaussian term of the HG functions (Eq. (5)) is the dominant term, hence the coherence volume, area and length are simply the regions where that Gaussian remains non-negligible, i.e. roughly while $\left(\frac{k}{r} \vec{\xi} \right)^T \Theta^{-1} \left(\frac{k}{r} \vec{\xi} \right) < 2$ (the cut-off value of 2 was

chosen ad hoc). The properties then follow via spectral analysis of the Jacobians of the transformations. \square

Therefore, given a CSD expanded in HG space, we can immediately extract these important light characteristic directly from the shape matrix. Note that the origin is set to be radiation source, as in Definition 4.1, thus r is the distance of (free-space) propagation.

The HG coefficients are real and of even order. Consider a CSD expressed using transverse HG modes (as in Definition 4.1). Observe that $C(\vec{\xi}) = C^*(-\vec{\xi})$, directly from the definition (Eq. (10)), that $\exp(ik\hat{r} \cdot \vec{\xi})$ is a Hermitian function and that an HG function of even order is an even function, otherwise it is an odd function (see Eq. (S1.8) in our supplemental material). These facts hold for any $\vec{\xi}$ and we deduce that (i) $\check{c}_{nm} = \check{c}_{nm}^*$, i.e. the HG modes are real; and (ii) odd-ordered HG modes must vanish, viz. $\check{c}_{nm} = 0$ when $n + m$ is odd.

For a physical insight, observe that the Fourier transform of an odd-ordered HG function is purely imaginary, and of an even-ordered HG function purely real (see Eq. (9)), with the interpretation being: Diffracted even-ordered modes fully propagate to the far field, while odd-ordered modes only give rise to evanescent waves, which do not propagate nor contribute to far-field radiation.

4.1 Light-Matter Interaction

We return to our discussion of light-matter interaction. As before, let a light source (primary or secondary) at \vec{s} give rise to radiation that is incident to a small scattering region, centred at the origin, with scattering amplitude function σ , and scattering mutual intensity $\Sigma(\vec{\xi}_1, \vec{\xi}_2) = \sigma(\vec{\xi}_1)\sigma^*(\vec{\xi}_2)$. The incident light interacts with matter and scatters and diffracts to a point \vec{r} .

Plug Definition 4.1 into Theorem 3.1 and in our supplemental material we outline the lengthy derivations that follow, which ultimately result in the following Lemma:

LEMMA 4.3 (LIGHT-MATTER INTERACTION: THE HERMITE-GAUSS TRANSVERSE MODES). *A scattered Hermite-Gauss transverse mode relates to the incident Hermite-Gauss transverse modes and shape matrix, \check{c}'_{nm} and Θ' , via the scattering mutual intensity, Σ , as follows:*

$$\check{c}_{nm} = \frac{(2\pi)^{\frac{3}{2}} (-i)^{n+m}}{s^2 \lambda^2} \sum_{n', m'} \check{c}'_{n'm'} \times \mathcal{A} \left\{ \Sigma \cdot \Psi_{nm}^{\Theta^{-1}} \left(\mathcal{Q} \frac{\vec{r}'_1 + \vec{r}'_2}{2} \right) \Psi_{n'm'}^{\Theta'} \left(k \mathcal{Q}' \frac{\vec{r}'_1 - \vec{r}'_2}{s} \right) \right\} \left(\vec{\phi}, \vec{\phi} \right),$$

with the shorthand $\vec{\phi} = k(\hat{r} + \hat{s})$. The variables $\vec{r}'_{1,2}$ are the ACT integration variables and the orthogonal matrices $\mathcal{Q}, \mathcal{Q}'$ transform to the local frame of the scattered and incident radiation, respectively.

PROOF. See Subsection S2.2 in our supplemental material. \square

The shape matrix. Lemma 4.3 holds for arbitrary shape matrices Θ . This is not a coincidence: recall that the anisotropic HG functions form a complete functional basis, regardless of the choice for Θ . In practice, only a finite count of HG modes will be used, hence, to complete the parameterization of the CSD of the scattered field, we turn our attention to devising a formula for the shape matrix Θ . Derived next is the following curious relation:

LEMMA 4.4 (LIGHT-MATTER INTERACTION: THE COHERENCE SHAPE MATRIX). *Given Θ' , the shape matrix of the incident radiation, and σ , the scattering amplitude function, the shape matrix of the scattered radiation is:*

$$\Theta = -\mathcal{Q}^T \cdot \left[\frac{\partial^2}{\partial \vec{w}^2} \ln \check{f}_{00} \right]_{\vec{w}=0}^{-1} \cdot \mathcal{Q},$$

with the 0th-order scattering mode:

$$\check{f}_{00}(\vec{w}) \triangleq \mathcal{A} \left\{ \Sigma \cdot \Psi_{00}^{\Theta'} \left(k \mathcal{Q}' \frac{\vec{r}'_1 - \vec{r}'_2}{s} \right) \right\} \left(\vec{\phi} + \frac{1}{2} \vec{w}, \vec{\phi} - \frac{1}{2} \vec{w} \right).$$

That is, the shape matrix is related via a similarity transform to the inverse of the Hessian matrix of the natural logarithm of the 0th-order scattered mode.

PROOF. See Subsection S2.3 in our supplemental material. \square

Lemmas 4.3 and 4.4 constitute the theoretical foundations of light transport and light-matter interaction in HG space.

Observe that the matrix produced by Lemma 4.4 is expected to be positive-definite, as required: The 0th-order scattering mode \check{f}_{00} is the autocorrelation of the scattered radiation with separation \vec{w} , which is expected to attain a local maximum at $\vec{w} = 0$. The logarithm of that concave function will also attain a local maximum at \vec{w} , thus the Hessian is negative-definite, its additive inverse is then positive-definite and definiteness is invariant under matrix inversion. Nevertheless, due to peculiarities in the matter behaviour and numerical issues, at times the generated matrix might become semi-definite or indefinite. In practice, it is then desired to ensure physical correctness by flipping the signs of negative singular values of Θ , if any, and, for numeric stability, clipping the singular values to a small $\epsilon > 0$.

5 SCATTERING WITH LOCALLY-STATIONARY MATTER

We now take a closer look at how we describe physical matter and its response to incident radiation.

As discussed, it is desirable to quantify the micro-scale scattering properties of the matter statistically. To that end, we now consider the scattering amplitude function σ as a (spatial) stochastic process, in which case the mutual scattering function is in fact the autocorrelation of the process, viz. $\Sigma(\vec{r}_1, \vec{r}_2) = \langle \sigma(\vec{r}_1) \sigma^*(\vec{r}_2) \rangle$. We also restrict σ to be a *locally-stationary stochastic process*. The class of locally-stationary processes was introduced by Silverman [1957] as a generalization of wide-sense-stationary processes, and the autocorrelation of such a process can be written as

$$\Sigma(\vec{r}_1, \vec{r}_2) = \left\langle \left| \sigma \left(\frac{\vec{r}_1 + \vec{r}_2}{2} \right) \right|^2 \right\rangle R_{\sigma\sigma}(\vec{r}_1 - \vec{r}_2), \quad (12)$$

where $R_{\sigma\sigma}$ is a normalised (i.e., $\int R_{\sigma\sigma} = 1$) spatial *stationary autocorrelation* function, which is a function of the spatial difference vector between the points $\vec{r}_{1,2}$. The ensemble-averaged scattering intensity above is the local, averaged (over the ensemble of all matter realizations) ratio between the scattered and incident intensity, and thus is directly related to the classical BSDF. Its (spatial) Fourier transform is the (averaged, spectral) scattering *angular coherence* (a

consequence of the Van Cittert–Zernike theorem [Born and Wolf 1999]), viz.

$$\tilde{f}_\sigma(\vec{\zeta}) \triangleq \mathcal{F} \{ \langle |\sigma|^2 \rangle \}. \quad (13)$$

Note the implicit dependence of all functions and quantities that describe the light-matter interaction process, σ , Σ , $R_{\sigma\sigma}$ and \tilde{f}_σ , on wavelength as well as the incident and scattered directions \hat{s} , \hat{r} remains. The ensemble-averaging operator $\langle \cdot \rangle$ that appears above averages over the ensemble of all surface realizations, and should not be confused with $\langle \cdot \rangle_\omega$, which averages over same-frequency constituents of the wave ensemble.

Before we discuss the practical and physical significance of locally-stationary matter, we first state that the stationary autocorrelation $R_{\sigma\sigma}$ and the angular coherence \tilde{f}_σ , together, fully describe the process of light interaction with the matter, and present the primary theoretical contribution we make in this paper:

THEOREM 5.1 (INTERACTION WITH LOCALLY-STATIONARY MATTER IN HERMITE-GAUSS SPACE). *Let σ be the matter's locally-stationary scattering amplitude, and $R_{\sigma\sigma}$ be its stationary autocorrelation. We set $\vec{\phi} = k(\hat{r} + \hat{s})$, as before. Then,*

(i) *the scattered HG transverse modes become:*

$$\check{c}_{nm} = \left\langle \tilde{f}_\sigma \left| \tilde{\Psi}_{nm}^\Theta(\mathbf{Q}\vec{r}') \right. \right\rangle \sum_{n',m'} \frac{\check{c}'_{n'm'}}{s^2 \lambda^2} \mathcal{F} \left\{ R_{\sigma\sigma} \Psi_{n'm'}^{\Theta'} \left(\frac{k\mathbf{Q}'\vec{r}'}{s} \right) \right\} (\vec{\phi}),$$

(ii) *and, the shape matrix is:*

$$\Theta = \mathbf{Q}^\top \cdot \left[\frac{\partial^2}{\partial \vec{\zeta}^2} \ln \tilde{f}_\sigma(\vec{\zeta}) \right]_{\vec{\zeta}=0}^{-1} \cdot \mathbf{Q}.$$

PROOF. Make the variable substitutions $\vec{r}_1'' = \vec{r}_1' - \vec{r}_2'$ and $\vec{r}_2'' = \frac{1}{2}(\vec{r}_1' + \vec{r}_2')$ in Lemmas 4.3 and 4.4, which decouple the ACT into independent FTs. Theorem 5.1.(ii) follows immediately by the fact that multiplicative constants in the logarithm annihilate under the derivative. Apply the convolution theorem as well as the HG function identities Eq. (9) and Eq. (S1.8) (in the supplemental material), the convolution then reduces to the inner product with the dual HG function, yielding Theorem 5.1.(i). \square

We summarise some notable features of the theorem above:

- Two independent terms arise: (i) The inner product of the angular coherence function \tilde{f}_σ with the dual HG functions. This term does not depend on the coherence of the incident light. (ii) The interference term, i.e. the FT of the stationary autocorrelation $R_{\sigma\sigma}$ with the incident HG transverse modes.
- The sum over the incident modes in Theorem 5.1.(i) does not depend on the indices n , m and can be computed once, thus the computational complexity is now linear in the count of HG modes, as opposed to quadratic.
- As the shape matrix (Theorem 5.1.(ii)) also depends only on \tilde{f}_σ , the ensemble-averaged scattering behaviour fully quantifies the coherence shape of the scattered radiation. Meanwhile, interference only depends on the autocorrelations of both the matter and the incident light's coherence.
- That interference term reduces to a simple, computationally-tractable Fourier optics relation. Furthermore, these formulae

are generally devoid of the highly-oscillatory phase terms that are typical with such diffraction integrals. This is because we only need to consider the statistical correlations of the matter's scatter behaviour and incident coherence.

The analytic form of the autocorrelation of locally-stationary matter, viz. Eq. (12), should be understood as matter where the ensemble-averaged scattering intensity, $\langle |\sigma(\vec{r})|^2 \rangle$, is a slow function compared to $R_{\sigma\sigma}$, i.e. it changes slowly with respect to the distances over which σ remains correlated. As previously discussed, a purely deterministic description of matter is infeasible due to the resolution needed to quantify interactions with light of optical frequencies. On the other hand, a purely statistical description is less interesting: *It is the irregularities in the matter that give rise to distinct diffractive phenomena.* Formally, this is a consequence of the fact that Θ only depends on an ensemble-averaged quantity (Theorem 5.1.(ii)), which averages away phase differences between scattered wavefronts, and thus only matter features that survive the ensemble-averaging contribute to Θ .

Thus, by endowing $\langle |\sigma|^2 \rangle$, a quantity closely related to the classical BSDF, with a stationary autocorrelation function—thereupon transforming it into a locally-stationary stochastic process—we introduce a matter model which allows us to practically superpose statistical perturbations of the matter over the broader explicit features that induce interesting scattering behaviour (e.g., imperfections and irregularities in the matter). Under that context, Theorem 5.1 provides computationally-tractable formulae for light-matter interaction in Hermite-Gauss space. While, so far, all of our derivations have been mathematically exact (up to the optical assumptions that led to Theorem 3.1) and very general, the simple analytic relations and well-behaved quantities in Theorem 5.1 suggest that for more specific types of matter, simpler relations could be found.

See Subsection S2.4 in our supplemental material for additional notes on computing the shape matrix via Theorem 5.1.(ii).

Perfectly specular reflectors. One type of surface that cannot be directly described by Theorem 5.1 is a perfect mirror. A perfect specular reflector is a notorious violator of the Fraunhofer region assumption, specifically the assumption that the scattering amplitude function σ —now a very sharp impulse, a Dirac delta—changes slowly with the directions \hat{r} , \hat{s} . This is an issue that afflicts any scatter theory formulated under Fresnel or Fraunhofer optics, and it is common to separate the interaction into a “specular lobe” and a “scattered lobe”: Harvey [2012] do so by introducing an exponential term that partitions the reflected energy between these lobes based on surface roughness. That very same approach was also used by Holzschuch and Pacanowski [2017].

A perfectly specular reflection (or refraction) does not change the coherence properties of light [Steinberg and Yan 2021]. Also, as the energy is concentrated in a sharp impulse regardless of the coherence of light, such surfaces constitute the simplest kind of light-matter interactions and are easy to path trace. For simplicity, in our implementation we classify matter into specular surfaces and everything else, however superposing specular and non-specular lobes could be done as well.

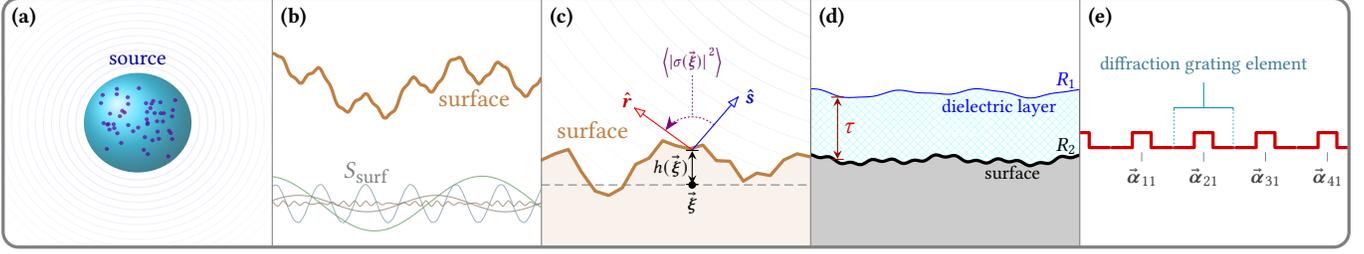


Fig. 6. Examples of light-matter interaction that we discuss in this paper. (a) Sourcing of partially-coherent light: Charges and current in a light source give rise to electromagnetic radiation. The amount of charge together with its spatial correlation fully describe the emitted radiation and its coherence properties. (b) Scatter by a statistical surface: A surface (brown) is decomposed into its spatial frequencies (illustrated in gray). The power of each frequency is given by the power spectral density (PSD) function S_{surf} , which plays an important role in describing the scattered radiation. (c) Scatter by explicitly-defined microgeometry: A surface patch is described by the (explicit) height perturbations $h(\vec{\xi})$, where $\vec{\xi}$ is a point on the surface plane (gray dashed line). As always: \hat{s} , \hat{r} are the incident and exitant directions and $\langle |\sigma(\vec{\xi})|^2 \rangle$ is the (ensemble-averaged) reflectivity (i.e. scattered intensity). (d) Layered structure: A dielectric layer over a substrate, where the surface of each is described statistically via the stationary autocorrelation functions R_1 and R_2 , respectively. The surfaces are assumed to be statistically independent. (e) Diffraction grating: A two-dimensional grid of small and identical, but otherwise arbitrary, optical elements produces a diffraction grating. The elements are regularly spaced and the position of the center of each element is denoted $\vec{\alpha}_{pq}$.

5.1 Examples

We now present a few examples and revisit some of the topical problems in computer graphics that involve rendering diffractive phenomena. The purpose of this subsection is two-fold: to present example applications as well as compare our formalism with some of the state-of-the-art. It should be remembered that the primary motivation behind our work is not the reproduction of material appearance, but the understanding of how light-matter interaction affects the coherence properties of light and the transport of these properties. Hence, there is a significant qualitative difference in the information produced by our methods compared with the state-of-the-art: We fully quantify the coherence properties of the scattered light and not merely the scattered intensity.

See Fig. 6 for an overview of the types of matter we discuss in this subsection.

Light sources. Emission of partially-coherent light by a light source is a special case of light-matter interaction: Electromagnetic radiation arises due to a time-varying distribution of charge and current within a source. The ensemble-averaged scattering intensity $\langle |\sigma(\vec{r})|^2 \rangle$ plays the role of the time-averaged power that flows through a differential volume element, and $R_{\sigma\sigma}$ is the autocorrelation function of that distribution of power within the source (see Fig. 6a). For spontaneous emission sources, that spatial correlation is negligible [Carminati and Greffet 1999] and we set $R_{\sigma\sigma} = \delta^3$, i.e. a spatial Dirac delta. The sourcing equation for partially coherent light then becomes the inner product

$$\check{c}_{nm} = \frac{1}{\lambda^2} \left\langle \tilde{f}_{\sigma} \left| \tilde{\Psi}_{nm}^{\Theta}(\mathcal{Q}\vec{r}') \right. \right\rangle, \quad (14)$$

which simply is the HG coefficient of degree $(n, m, 0)$ of $\frac{1}{\lambda^2} \tilde{f}_{\sigma}$. If the time-averaged power is mostly constant throughout the source (the entire source radiates with similar intensity), then $\tilde{f}_{\sigma} = \Lambda(\omega) \mathcal{F} \{ \mathbb{1}_{\mathcal{X}} \}$, where $\mathbb{1}_{\mathcal{X}}$ is the characteristic function of the source geometry \mathcal{X} (that function takes values of 1 inside the source and 0 outside), and Λ is the emitted power spectral density. For a spherical source, \tilde{f}_{σ} is proportional to a Bessel function of the first kind. This agrees with

Steinberg and Yan [2021], though note that our formalism is able to describe sources with spatially-varying emission characteristics as well. Once \tilde{f}_{σ} is computed, the shape matrix and HG coefficients follow immediately. See our sample implementation in our supplemental material for an example with spherical and cylindrical sources.

Scatter by statistical surfaces. Due to the previously discussed practical reasons, it is appealing to describe surfaces statistically. Over the last few years, the Harvey-Shack and related surface scatter theories have seen significant adoption in computer graphics (e.g., Holzschuch and Pacanowski [2017]; Steinberg and Yan [2021a]). These surface scatter theories operate on the *power spectral density* (PSD) of the surface, which is convenient, as it is often the PSD that is measured from a physical surface sample, and not explicit geometry [Siewert et al. 2008]. The PSD is simply the FT-pair of the stationary autocorrelation, hence any statistical description via a PSD can be carried out exactly under our formalism (a restatement of the fact that locally-stationary stochastic processes generalise wide-sense-stationary processes).

Assume that the surface spans the (local) xy -plane, and the surface's two-dimensional autocorrelation and PSD functions are denoted as R_{surf} and $S_{\text{surf}} = \mathcal{F} \{ R_{\text{surf}} \}$, respectively. The stationary autocorrelation of the scatter process is then $R_{\sigma\sigma} = R_{\text{surf}}(r_x, r_y) \delta(r_z)$ (assuming R_{surf} is normalised and constant surface reflectivity), with the Dirac delta encoding the fact that we deal with a surface. If we were to assume that the height perturbations of the surface and the spatial correlation length (the length over which the surface remains correlated) are small with respect to the coherence area of light, then the FT in Theorem 5.1.(i) becomes dominated by the $R_{\sigma\sigma}$ term, and a closed-form expression is derived

$$\begin{aligned} \check{c}_{nm} &\approx \left\langle \tilde{f}_{\sigma} \left| \tilde{\Psi}_{nm}^{\Theta}(\mathcal{Q}\vec{r}') \right. \right\rangle \sum_{n', m'} \frac{\check{c}'_{n' m'}}{s^2 \lambda^2} \Psi_{n' m'}^{\Theta'}(0) \mathcal{F} \{ R_{\sigma\sigma} \} (\vec{\phi}^{\perp}) \\ &= \frac{1}{\lambda^2} \left\langle \tilde{f}_{\sigma} \left| \tilde{\Psi}_{nm}^{\Theta}(\mathcal{Q}\vec{r}') \right. \right\rangle C'(0) S_{\text{surf}}(\vec{\phi}^{\perp}), \end{aligned} \quad (15)$$

where $C'(0)$ is the intensity of the incident radiation, $\vec{\phi} = k(\hat{r} + \hat{s})$ (with \hat{s} , \hat{r} being the direction to the source and scattering direction, respectively) as defined in Lemma 4.3 and $\vec{\phi}^\perp$ is the projection of $\vec{\phi}$ onto the surface (the xy -plane). This is indeed the relation predicted by the generalised Harvey-Shack and Rayleigh–Rice scatter theories [Krywonos 2006]: The scattered intensity is proportional to the surface PSD, evaluated at the spatial frequency of the diffraction lobe, which is precisely $\vec{\phi}^\perp$. As we require $R_{\sigma\sigma}$ to be normalised, the constants (describing the surface reflectivity) are captured by \tilde{f}_σ . Assuming the reflectivity remains roughly constant across the surface (which is the assumption taken by these theories), then the scattering surface patch acts as a homogeneous secondary light source, and indeed the shape matrix takes the same form as in the case of sourcing of light.

It is important to note that while matter that is purely stationary (at least in the wide-sense) is formally possible, such matter is aphysical, as stationarity implies that $\langle |\sigma|^2 \rangle$ is a positive constant and the scattering region extends to infinity. Moreover, the formula for Θ (Theorem 5.1.(ii)) diverges under stationarity, which indeed correctly reflects the physics: radiation sourced from a radiator with an infinite extent is perfectly-incoherent, and thus does not propagate [Wolf 2007] (in addition to requiring infinite energy). Therefore, \tilde{f}_σ should restrict the scattering region to a (small) finite region.

Diffraction microgeometry. Consider a surface with an explicit heightmap $h(x, y)$, and let a be the (potentially complex and spatially-varying) reflectivity of the surface. The scattering amplitude function can be formulated by encoding the phase-shifts induced by height variations, and written as $\sigma(\vec{\xi}) = a(\vec{\xi}) \exp[i(\vec{\xi} + h(\vec{\xi})\hat{z}) \cdot \vec{\phi}]$, where $\vec{\xi}$ is a position on the surface. The angular coherence becomes $\tilde{f}_\sigma = \mathcal{F}\{|a|^2\}$ (as the data is deterministic, the ensemble-averaging operator is dropped), the FT of BRDF of the surface. The shape matrix follows immediately (via Theorem 5.1.(ii)), and observe that in general coherence anisotropy (Corollary 4.2) will arise due to spatial variations in reflectivity or the Lambertian term. The stationary autocorrelation is then the convolution

$$R_{\sigma\sigma}(\vec{\xi}) = \delta(\xi_z) \frac{1}{|\mathcal{X}|} \int_{\mathcal{X}} d^2\vec{r}' \sigma(\vec{r}') \sigma^*(\vec{r}' - \vec{\xi}) \quad (16)$$

(up to a normalization of the convolution), where \mathcal{X} is the surface patch and $|\mathcal{X}|$ is its area. The convolution over the limited extent implies that we understand the surface as a periodically tiled—thus, stationary—signal. The tiling is merely an artefact, as \tilde{f}_σ restricts the light-matter interaction to the surface extent anyhow. Consequently, any matter described deterministically over a limited extent can be formulated in an exact manner as locally-stationary matter.

The above is an essentially identical expression to the one derived by Steinberg and Yan [2021], with same computational complexity. Unfortunately, the “nice” analytic properties of Theorem 5.1 are tainted by the fact that the highly-oscillatory exponent arises in the signal σ . This is unavoidable when deterministically describing the matter on sub-micro-scale resolution, and a statistical approach should be preferred. For this reason we do not render surfaces with explicitly-defined microgeometry.

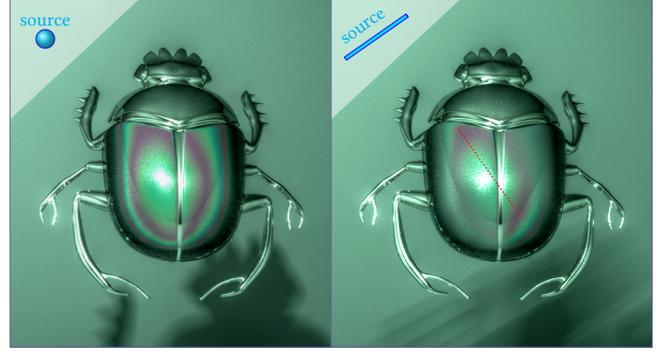


Fig. 7. Thin-film interference with partially-coherent light: A solid silver scarab, 14.60 cm in width (body), with a 1.23 μm coating applied to the wings is illuminated by a (left) spherical fluorescent light source of radius 0.60 cm (illustrated in blue), located about 60 cm away from the scarab, admitting a coherence area of about $15 \mu\text{m} \times 15 \mu\text{m}$. Interference patterns are clearly visible. (right) The same scarab is now illuminated by a cylindrical source of length 20 cm and radius 0.10 cm. Both sources radiate with identical spectrum and radiant power, however the coherence of the radiation from the cylindrical source is highly anisotropic: it remains coherent over $50 \mu\text{m}$ in one transverse direction but less than a micron in the other. The dotted red line indicates the scarab’s surface where the normal vector is perpendicular to the length of the light source. As the surface curves away from the indicated area, the interference pattern gradually disappears due to the coherence anisotropy, and as predicted by Eq. (18).

Layered structures. Just as any matter with a purely deterministic (periodic or of limited extent) or a purely stochastic (of a wide-sense-stationary nature) description can be captured by our formalism, we can also model matter with a hybrid description.

Consider a layer made of some homogeneous dielectric of thickness τ placed on top of a substrate (illustrated in Fig. 6d). Let the autocorrelation of the surface geometries of the top and bottom interfaces be R_1, R_2 , respectively, and for simplicity assume both surfaces have identical reflectivity. The autocorrelation of the surfaces then is $R_{\text{surf}} = R_1 + R_2$, assuming both are zero-mean, uncorrelated processes. Convolving in a like manner to Eq. (16), the autocorrelation in the z direction becomes

$$\bar{R}(z) = \delta(z) + \delta(z - \tau) \quad (17)$$

and the stationary autocorrelation is $R_{\sigma\sigma} = R_{\text{surf}}(r_x, r_y)\bar{R}(r_z)$. Let $S_{\text{surf}} = \mathcal{F}\{R_{\text{surf}}\}$ be the PSD, then

$$\mathcal{F}\left\{R_{\sigma\sigma} \Psi_{n'm'}^{\theta'}\left(\frac{kQ'\vec{r}'}{s}\right)\right\}(\vec{\phi}) = S_{\text{surf}}(\vec{\phi}^\perp) \times \left[\Psi_{n'm'}^{\theta'}(0) + \Psi_{n'm'}^{\theta'}\left(\frac{k}{s}\tau Q'\hat{z}\right)e^{-i\tau\phi_z}\right]. \quad (18)$$

The first term is the scatter of the statistical surfaces and the second is the interference between the layers, both expressed in closed-form. The real part of the exponent is $\cos(\tau\phi_z)$, which encodes the (average) phase-difference between the reflected wavefronts, in agreement with well-known optics. The HG function modulates the interference with respect to the spatial coherence at a distance of τ in the (local) z direction. Observe that this directional dependence implies that coherence anisotropies in the incident beam may induce

distinct visual responses, even if the incident spectral intensity and coherence area are held constant (see Fig. 7).

In the above simple example, the autocorrelation naturally decomposes into the surface autocorrelation and the autocorrelation of the deterministic layer structure. However, by using the fact that the autocorrelation of uncorrelated (zero-mean) processes is the sum of their autocorrelation functions, more interesting hybrid configurations can be constructed, e.g., the embedding of statistically distributed and oriented micro-flakes into a layer of resin.

Diffraction grating. Let a single scattering element be characterised by its scattering amplitude function. If we position such elements, adjacent and parallel to each other, in a two-dimensional grid extending to infinity on the xy -plane, we get an ideal diffraction grating. Naturally, this is a periodic structure and the autocorrelation $R_{\sigma\sigma}$ of the entire grating is computed simply as in Eq. (16). Consider now the FT in Theorem 5.1.(i). Rewrite that FT integral as a sum of integrals, each over the spatial extent of a single element in the grating. By making the assumption that the spatial extent of each element is much smaller than the spatial coherence of the incident light, the HG function remains effectively constant over an element. Then, by recalling that the autocorrelation of a periodic signal is also periodic with same period (the spatial extent over an element) and applying the translation property of the Fourier transform, we get

$$\mathcal{F}\left\{R_{\sigma\sigma}\Psi_{n'm'}^{\Theta'}\left(\frac{kQ'\vec{r}'}{s}\right)\right\}(\vec{\phi}) = \mathcal{F}\{R_{\sigma\sigma}\}(\vec{\phi}^\perp) \times \sum_{p,q} e^{-i\vec{\alpha}_{pq}\cdot\vec{\phi}} \Psi_{n'm'}^{\Theta'}\left(\frac{kQ'}{s}\vec{\alpha}_{pq}\right), \quad (19)$$

with $p, q \in \mathbb{Z}$ indexing the elements and $\vec{\alpha}_{pq}$ being the center position of an element (see Fig. 6e). The positions $\vec{\alpha}_{pq}$ are periodic, and we immediately recognise the sum above as a Fourier series. Applying our assumption that the spatial extent of the elements is small, and employing again the well-behaviour of the HG functions under the FT, we arrive at a closed-form expression:

$$\mathcal{F}\left\{R_{\sigma\sigma}\Psi_{n'm'}^{\Theta'}\left(\frac{kQ'\vec{r}'}{s}\right)\right\}(\vec{\phi}) = \left(\frac{s}{k}\right)^2 S_{\sigma\sigma}(\vec{\phi}^\perp) \widetilde{\Psi}_{n'm'}^{\Theta'}\left(\frac{sQ'}{k}\vec{\phi}^\perp\right), \quad (20)$$

with $S_{\sigma\sigma}$ being the PSD of the diffraction grating.

This result is very similar to the statistical surface scatter formula, Eq. (15), and the difference reflects the fact that a typical surface admits short correlation lengths, while diffraction grating remains correlated across the entire surface. Note that under partially-coherent illumination, diffraction lobes will attenuate with greater lobe orders (as the magnitude of the projection of $\vec{\phi}$ onto the surface becomes greater with each lobe), as expected (this effect is rendered in Fig. 1).

Optical speckle. Consider a volume $V \subset \mathbb{R}^3$ that contains N infinitesimal scattering particles, each scattering uniformly. The volume's scattering function can then be written as a sum of the isotropic scattering function multiplied by spatial dirac deltas, viz. $\sigma(\vec{r}) = \sum_j a_j \delta^3(\vec{r} - \vec{p}_j)$, where \vec{p}_j are the particles positions and a_j are the real, non-negative scattering coefficients. Then,

$$\left\langle \tilde{f}_\sigma \left| \widetilde{\Psi}_{nm}^\Theta(Q\vec{r}') \right. \right\rangle = \sum_{jl} e^{-i\vec{s}_l \cdot \vec{p}_j} a_j a_l \widetilde{\Psi}_{nm}^\Theta(Q\vec{p}_l), \quad (21)$$

which is a closed-form expression for the HG-coefficients of the expansion of the angular coherence function \tilde{f}_σ . More importantly,

$$\begin{aligned} & \mathcal{F}\left\{R_{\sigma\sigma}\Psi_{n'm'}^{\Theta'}\left(\frac{kQ'\vec{r}'}{s}\right)\right\}(\vec{\phi}) \\ &= \frac{1}{|V|} \int_{\mathbb{R}^3} \int_V d^3\vec{r}' d^3\vec{\xi}' \sigma(\vec{\xi}') \sigma^*(\vec{\xi}' - \vec{r}') \Psi_{n'm'}^{\Theta'}\left(\frac{kQ'\vec{r}'}{s}\right) e^{-i\vec{r}'\cdot\vec{\phi}} \\ &\approx \sum_{jl} \Psi_{n'm'}^{\Theta'}\left(kQ'\frac{\vec{p}_l - \vec{p}_j}{s}\right) a_j a_l e^{-i(\vec{p}_l - \vec{p}_j)\cdot\vec{\phi}}, \end{aligned} \quad (22)$$

where we first applied Eq. (16), as in the case of a deterministic surface (technically, the second equality is only exact in the $V \rightarrow \infty$ limit). The sum above is the superposition of partially-coherent speckle contributions, where the HG function simply quantifies how correlated are the waveforms that are scattered from a pair of particles. This equation defines partially-coherent speckle, and indeed it is the discretization (in HG space) of the expression for partially-coherent speckle intensity of Steinberg and Yan [2021a, p.8,eq.16]. Consequently, typical analysis of partially-developed speckle and sums of correlated speckle patterns apply [Goodman 2020] and the relevant speckle phenomena—like the well-known *memory effect*—can be reproduced without any special treatment. Monte-Carlo integrators developed by the computer graphics community to render (far-field) speckle [Bar et al. 2019; Steinberg and Yan 2021a] may be used to solve Eq. (22) when N is large.

The value of the analysis above is mostly academic: it serves to show that far-field, spatial speckle—like many far-field optical phenomena—is immediately describable and reproducible by our theory. However, the analysis above is confined to deterministically-described scattering particles. Speckle and scattering by media are phenomena that are mostly statistical in nature. Of more interest is an examination of partially-developed speckle produced by general locally-stationary media, which is left for future work.

5.2 Rendering and Implementation

Rendering was done using a proof-of-concept implementation in Mitsuba, building upon the foundations developed by Steinberg and Yan [2021]. We extend the bi-directional integrator bdpt to propagate coherence information in Hermite-Gauss space, and implement light-matter interaction using the core ideas formulated in this paper.

Light is sourced using Eq. (14) and each traced ray carries the shape matrix, transverse HG coefficients, the propagation direction and distance r . This is sufficient information to evaluate the CSD as defined by Definition 4.1. As discussed, some surfaces allow forward and backward path tracing: The rough statistical surfaces (materials (d) and (e) in Fig. 1) are rendered using Eq. (15), and as the coherence properties of the incident light play no role in that formula, standard bi-directional path tracing with importance sampling applies. To speed-up convergence, these surfaces were importance sampled via the classical microfacet BRDF with manually tweaked parameters. Importance sampling Eq. (15) is left for future work. The scattering region was assumed to be circular and the HG coefficients of the angular coherence function were precomputed. Also, the specular dielectric upper layer of the DVD disks is considered to be a perfect specular reflector, and scatters as a Dirac delta without affecting the coherence properties of light. When scattered by such a reflector, r is

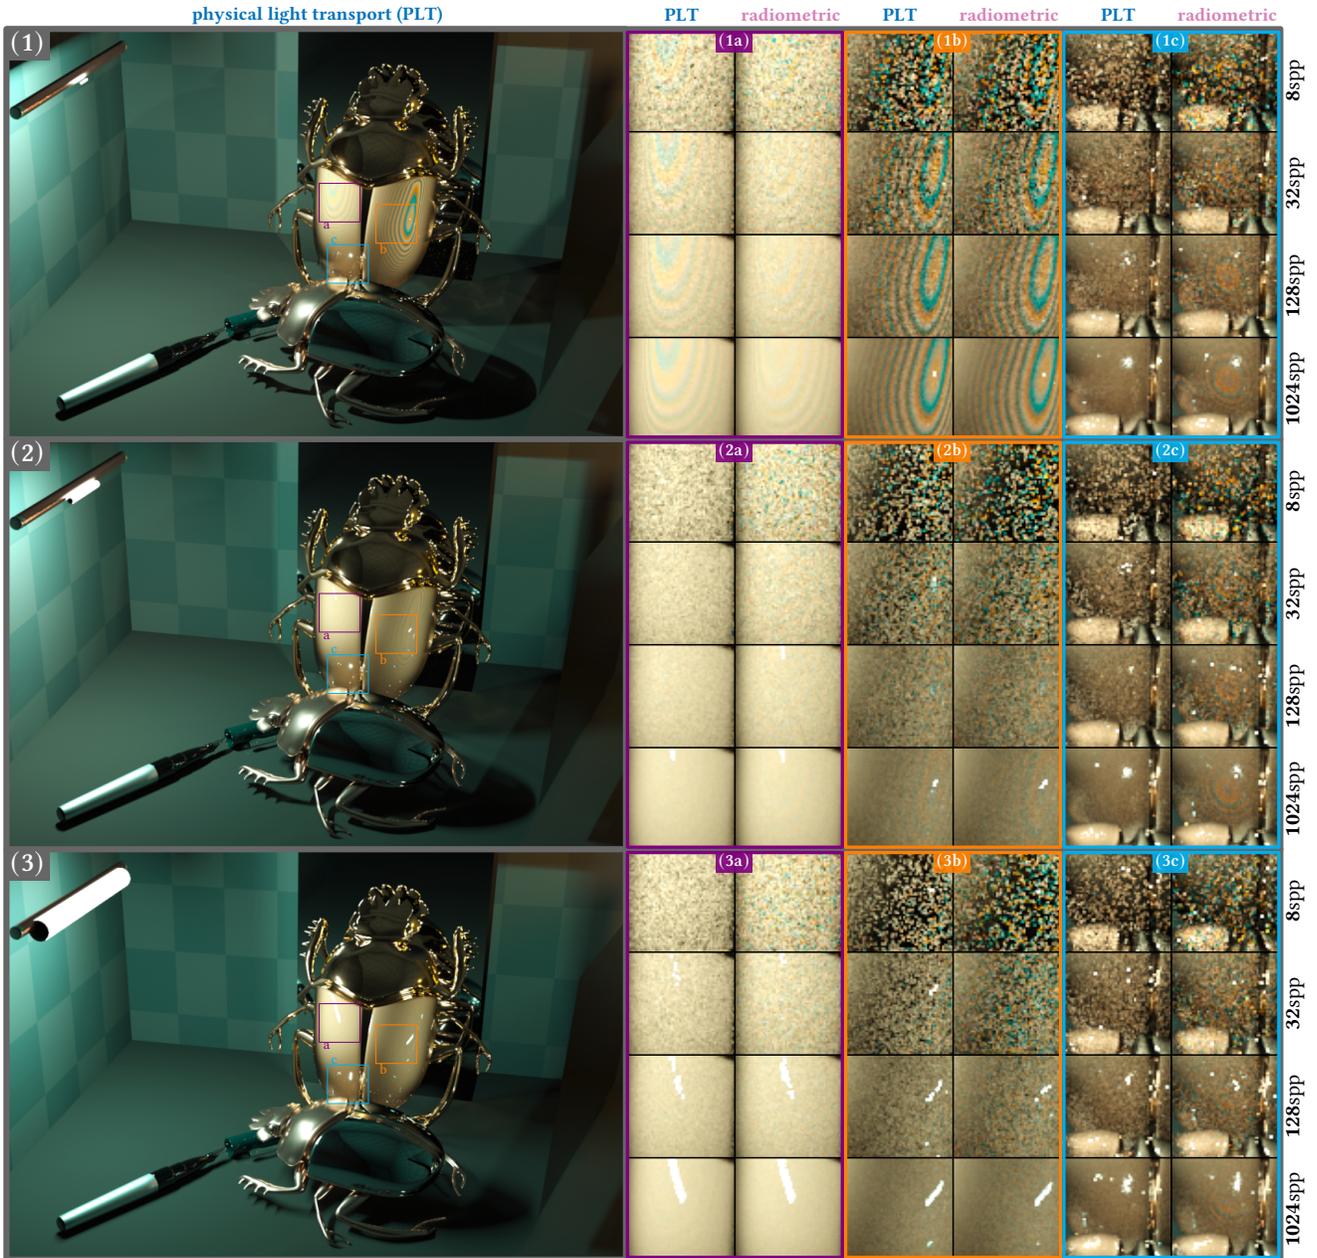


Fig. 8. (left) A large golden scarab with polished body and unpolished, rough wings, that are coated with a moderately thick dielectric, is illuminated by a short-spectrum light source of varying size. A smaller copper scarab with rough body and (uncoated) polished wings is placed at the bottom. A rectangular mirror, on the right, reflects light from the source back towards the golden scarab. The mirror and the polished small scarab's wings act as Dirac delta perfect reflectors (see Section 5). Hence, light incident upon the golden scarab after reflection off the mirror is more coherent than light that arrives directly from the source: the right wing showcases more pronounced interference patterns than the left. Rendered with 8192spp. (right) We compare our coherence-aware physical light transport (PLT) with radiometric light transport (RLT) where the golden scarab's wings are rendered the method by Belcour and Barla [2017]. Each of the insets is a close-up on an area outlined in the full image, rendered with increasing samples-per-pixel. RLT mimics spatial coherence by superpositioning many plane waves sourced from different points on the source and insets (a,b) show decent agreement between the produced patterns. However, to avoid aliasing many such samples are needed: observe the colourful noise artefacts in the RLT insets, which mostly disappear only with 1024spp and do not exist at all with PLT. The required sample count becomes particularly untenable with indirect weakly-coherent illumination, which is highlighted in insets (c): this area is illuminated indirectly by light diffused by the rough body of the small scarab. The close proximity of the small scarab and its body's rough surface mean that this light is very weakly-coherent (note the size of the diffused highlight on the rough body and invoke the VCZ theorem), an effect correctly captured by our method but not RLT: observe the erroneous diffraction patterns that appear. Both methods admit similar rendering times with identical sample count. High-resolution renderings are available in our supplemental material.

not reset to 0 and the effective propagation distance is accumulated with the propagation before scattering.

Other materials can only be rendered via forward path tracing: The coated scarab wings are rendered with Eq. (18) and the diffraction grating of the bottom layer of the DVDs using Eq. (20) (materials (c) and (f) in Fig. 1, respectively). The PSD of the diffraction grating lobes induced by the (unrecorded, for simplicity) disk pits is formulated in similar manner to Sun et al. [2000]. The coherence of the incident radiation gives rise to diminished visibility of the secondary diffraction lobes both on the right DVD disk in Fig. 1, and especially so on the left disk, as predicted by Eq. (20). Just as with the rough surfaces, for both the diffraction grating and the coated wings, the scattering region is assumed to be circular with constant ensemble-averaged reflectivity (this is most accurate at normal incidence, and only affects the scattered shape matrix and thus only the subsequent interactions).

For simplicity, the rendering of these figures was done using the 0th-order HG mode only.

Results. Fig. 1 was rendered using 90 000 samples and uniformly distributed 32 spectral samples (rendering time of ~ 3 days on a modern desktop computer). The renderings in Fig. 7 was rendered using 10 000 samples (about 4 hours of rendering time), though residual noise is still visible.

We also compare our physical light transport approach with a non coherence-aware renderer in Fig. 8. Radiometric transport may reproduce the effects of spatial coherence by sampling many points on the (primary or secondary) source—each acting as a coherent plane wave constituent of the wave ensemble. However, very many such samples are needed to accurately reflect the statistics of the wave ensemble. This becomes especially difficult when diffused secondary sources (e.g. reflection from a rough surface) illuminate diffractive materials. Both methods admit similar rendering time with an identical sample count, as the cost is dominated by the path-tracing, but our method doesn't suffer from wave-ensemble aliasing and produces more accurate results with lower sample counts.

The appearance of the materials in Figs. 1 and 7 is dominated by direct illumination from the sources, and for our spherical and cylindrical sources the sourced CSD is well approximated by a single HG mode (see Fig. 4). Therefore, to compare the difference between rendering with a single mode and 16 modes, we render a small scene that exhibits a greater degree of indirect illumination, see Fig. 9. The rendering time of the 16 HG modes image in Fig. 9 is about 7 hours with 16 384 samples, just over double the rendering time of the single HG mode image, indicating that evaluation of HG modes is not the primary cost. The visual differences between the two images are minor and are concentrated about the diffraction grating lobes and the surfaces that reflect this grating—to be expected with such a simple scene.

Unfortunately, present integrators are ill-suited for such light transport and rendering more complex scenes with strong indirect illumination is difficult with current tools: This is due to the fact that convergence is very slow when we are limited to forward-only path tracing. Nevertheless, the conclusions from this paper (see next section) pave way for the development of better path tracing approaches for physical light transport and interaction with matter.

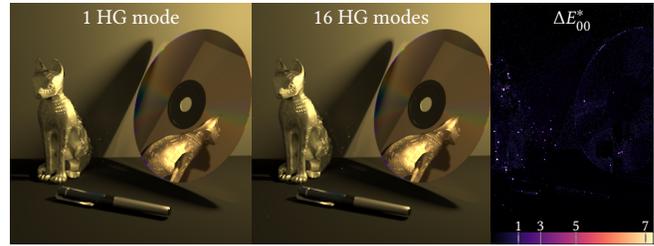


Fig. 9. A simple scene rendered using (left) the 0th-th order HG mode only and (middle) 16 HG modes. The lighting conditions are a mixture of direct and indirect illumination. (right) A difference image in ΔE_{00}^* (CIE DE2000) colour-difference space. The “just noticeable difference” is considered to be $\Delta E_{00}^* = 1$. Ignoring the path tracing noise, the differences between the rendered images is very minor. The small errors appear mostly on the DVD’s diffraction lobes and on the surfaces that face the DVD, as expected: the diffraction grating is the most coherence-sensitive material in this scene.

6 CONCLUSION

The primary result of this paper, Theorem 5.1, provides a complete formulation of light-matter interaction, in Hermite-Gauss space, with partially-coherent light of arbitrary spectral content. A compelling theoretical observation is that it is the stationary autocorrelation, together with the coherence of the incident light, that guides the interference process of the scattered radiation; while the ensemble-averaged scattering behaviour of the matter—a form of the classical BSDF—fully quantifies the diffraction process, and thus the shape of the scattered radiation. The fact that the diffraction process only depends on properties of the matter can be formulated as a stronger form of the Van Cittert–Zernike (VCZ) theorem for scattering with locally-stationary matter: *The Fourier transform of the ensemble-averaged scattering intensity distribution on the scattering medium is the coherence function of the scattered radiation.*

These conclusions have important practical consequences: Under physical optics path tracing is difficult. As noted, this is due to the fact that interference, and thus the angular distribution of scattered energy, depends on the coherence properties of the incident radiation. Nevertheless, the conclusions above imply that we are not limited to light tracing, but indeed we may *forward path trace from any point on a surface or medium*. Because each such scattering region acts as a secondary source. That is, if we start at a point on a surface, we have full information of the scattered beam (up to its spectral intensity), despite the fact that the coherence of an incident radiation can be arbitrary. Hereafter, all the information required to trace a path, to a sensor or other matter, is available. This is a significant step forward from Steinberg and Yan [2021] and lays the foundations for new practical bidirectional path tracing approaches for physical light transport.

We would also like to challenge the popular notion that for some scatter and diffraction formulations, like scattering by random surfaces and thin-film interference, the coherence properties of light can be ignored. It is usually assumed that the diffracting matter features of such materials are small with respect to the spatial coherence of light. However, as implied in Figs. 7 and 8, even with tiny features coherence may play a role. This should not be surprising:

there is no theoretical lower bound for the spatial coherence of light. Sub-wavelength spatial coherence can be produced by large sources or diffuse scatterers, and has been measured in practice [Morrill et al. 2016]. Hence, it should be remembered that the characteristic length of the coherence area of partially-coherent light spans orders-of-magnitude: from a fraction of a wavelength to hundreds of micrometres, and may play a more active role in the reproduction of material appearance than currently believed in computer graphics.

Polarization. We have neglected polarization in our discussion and treated all fields as scalar fields. Dealing with vector fields is straightforward: in our supplemental material (Section S4) we introduce the *generalized Stokes parameters* and show how to extend our theory to a fully-vectorized formalism.

Accuracy and limitations. The accuracy of our results is up to the same assumptions that led to Theorem 3.1: the Fraunhofer diffraction region and the Born first-order approximation. The rest of the derivations are mathematically exact. We made analytic progress chiefly by describing the physics under restricted, reasonable analytic models (HG basis and locally-stationary matter).

The main weakness of the presented work is the reliance on the Born first-order approximation. This implies that multiple interactions of light with the matter, like multiple inter-reflections in a layered material, are not accounted for. The Born first-order approximation is necessary to make progress and that weakness is inherent to all such formulations of scattering and asymptotic physical optics. Repeated interactions of the scattered radiation with the scattering matter render any Fresnel or Fraunhofer assumptions null and void, making this a difficult problem.

On the other hand, the Fraunhofer assumption is typically a very decent approximation for partially-coherent light [Charnotskii 2019]. Note that some optical effects, such as lens aberrations, are at times referred to as “near-field” effects. This terminology differs from theoretical disciplines, where “far field” is the region where the Fraunhofer approximation is valid. Therefore, despite the limitations of Born first-order approximations, repeated interactions at different locations in the matter—far with respect to the spatial coherence of light—can be carried out accurately as each interaction is independent from the other (e.g., path tracing in a participating medium that produces many individual scattering events).

It is interesting that despite the very different optical and mathematical approaches taken by us compared with the Harvey-Shack surface scatter theory [Krywons 2006], we arrive at similar results for the case of scatter by statistical surfaces, viz. Eq. (15), though our theory is able to quantify the impact of optical coherence.

Practical count of HG modes. In practice, the primary source of error would be the truncation of the expansion under the Hermite-Gauss basis. There are firm reasons to believe that no great number of modes are required: the precise oscillations of the coherence function on the transverse plane are less important than the total coherence area and its first-order geometric properties (rotation and scale), and the latter parameters are well captured by the shape matrix Θ . It is also suggested that some practical CSD functions may be represented with a reasonable degree of accuracy using rather few modes, see Fig. 4.

As a matter of fact, even with a single HG mode we are still able to represent a non-trivial family of coherence functions (see Fig. 4), propagate that coherence information throughout the scene and quantify interaction with matter in a computationally-tractable manner—the very problems we set out to solve in this paper. Furthermore, that 0th-order term is a simple Gaussian, thus the expressions in Theorem 5.1.(i) reduce to rather simple expressions. This might be advantageous to lower-accuracy applications, e.g., real-time rendering, and especially so if closed-form expressions can be found (or approximated) for \hat{f}_{σ} , $R_{\sigma\sigma}$ for some classes of matter.

Nevertheless, most coherence functions would not be approximated well by a single HG mode. The impact of this greatly depends on the application and used frequencies: For example, the accuracy of simulations performed with lower-frequency non-optical radiation (e.g., radar) depend more on accurate representation of the CSD than rendering with optical frequencies. The fact that small inaccuracies arise in a very simple scene, rendered with optical frequencies (Fig. 9), suggests that a single mode is not sufficient for optically-accurate simulations of physical light transport. We leave further investigation for future work.

REFERENCES

- 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>
- Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. 2019. A Monte Carlo framework for rendering speckle statistics in scattering media. *ACM Transactions on Graphics* 38, 4 (Jul 2019), 1–22. <https://doi.org/10.1145/3306346.3322950>
- Chen Bar, Ioannis Gkioulekas, and Anat Levin. 2020. Rendering near-field speckle statistics in scattering media. *ACM Transactions on Graphics* 39, 6 (Nov 2020), 1–18. <https://doi.org/10.1145/3414685.3417813>
- 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>
- Rafael Betancur and Roman Castañeda. 2008. Spatial coherence modulation. *Journal of the Optical Society of America A* 26, 1 (Dec 2008), 147. <https://doi.org/10.1364/josaa.26.000147>
- Max Born and Emil Wolf. 1999. *Principles of optics : electromagnetic theory of propagation, interference and diffraction of light*. Cambridge University Press, Cambridge New York.
- Rémi Carminati and Jean-Jacques Greffet. 1999. Near-Field Effects in Spatial Coherence of Thermal Sources. *Phys. Rev. Lett.* 82 (Feb 1999), 1660–1663. Issue 8. <https://doi.org/10.1103/PhysRevLett.82.1660>
- 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>
- Tom Cuypers, 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>
- D. S. Dhillon, J. Teyssier, M. Single, I. Gaponenko, M. C. Milinkovitch, and M. Zwicker. 2014. Interactive Diffraction from Biological Nanostructures. *Computer Graphics Forum* 33, 8 (2014), 177–188. <https://doi.org/10.1111/cgf.12425>
- 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>
- Greg Gbur. 2014. Partially coherent beam propagation in atmospheric turbulence [Invited]. *Journal of the Optical Society of America A* 31, 9 (Aug 2014), 2038. <https://doi.org/10.1364/josaa.31.002038>
- Ioannis Gkioulekas, Anat Levin, Frédo Durand, and Todd Zickler. 2015. Micron-scale light transport decomposition using interferometry. *ACM Transactions on Graphics* 34, 4 (Jul 2015), 1–14. <https://doi.org/10.1145/2766928>
- Joseph Goodman. 2015. *Statistical optics*. John Wiley & Sons Inc, Hoboken, New Jersey.
- Joseph Goodman. 2020. *Speckle phenomena in optics : theory and applications*. SPIE Press, Bellingham, Washington.
- 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>

- 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>
- Mourad E. H. Ismail and Plamen Simeonov. 2020. Multivariate holomorphic Hermite polynomials. *The Ramanujan Journal* 53, 2 (Sep 2020), 357–387. <https://doi.org/10.1007/s11139-020-00312-8>
- X.J. Kang, F. Zhang, L.Y. Xu, Y.S. Zhang, and Y.B. Li. 2015. Double image encryption based on the random fractional transform. In *IET International Radar Conference 2015*. Institution of Engineering and Technology. <https://doi.org/10.1049/cp.2015.1038>
- Ori Katz, Pierre Heidmann, Mathias Fink, and Sylvain Gigan. 2014. Non-invasive single-shot imaging through scattering layers and around corners via speckle correlations. *Nature Photonics* 8, 10 (Aug 2014), 784–790. <https://doi.org/10.1038/nphoton.2014.189>
- 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 and Jin-Ren Yao. 2020. Bi-static LIDAR systems operating in the presence of oceanic turbulence. *Optics Communications* 460 (Apr 2020), 125119. <https://doi.org/10.1016/j.optcom.2019.125119>
- Alankar Kotwal, Anat Levin, and Ioannis Gkioulekas. 2020. Interferometric transmission probing with coded mutual intensity. *ACM Transactions on Graphics* 39, 4 (Jul 2020). <https://doi.org/10.1145/3386569.3392384>
- Andrey Krywonos. 2006. *Predicting surface scatter using a linear systems formulation of non-paraxial scalar diffraction*. Ph.D. Dissertation. University of Central Florida.
- Mayukh Lahiri and Emil Wolf. 2009. Beam condition for scattering on random media. *Journal of the Optical Society of America A* 26, 9 (Aug 2009), 2043. <https://doi.org/10.1364/josaa.26.002043>
- 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>
- Ning Ma, Steen G. Hanson, Mitsuo Takeda, and Wei Wang. 2015. Coherence and polarization of polarization speckle generated by a rough-surfaced retardation plate depolarizer. *Journal of the Optical Society of America A* 32, 12 (Nov 2015), 2346. <https://doi.org/10.1364/josaa.32.002346>
- Pujuan Ma, Barbora Kacerovská, Raha Khosravi, Chunhao Liang, Jun Zeng, Xiaofeng Peng, Chenkun Mi, Yashar E. Monfared, Yuping Zhang, Fei Wang, and et al. 2019. Numerical Approach for Studying the Evolution of the Degrees of Coherence of Partially Coherent Beams Propagation through an ABCD Optical System. *Applied Sciences* 9, 10 (May 2019), 2084. <https://doi.org/10.3390/app9102084>
- Drew Morrill, Dongfang Li, and Domenico Pacifici. 2016. Measuring subwavelength spatial coherence with plasmonic interferometry. *Nature Photonics* 10, 10 (Sep 2016), 681–687. <https://doi.org/10.1038/nphoton.2016.162>
- A. Musbach, G. W. Meyer, F. Reitech, 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>
- 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>
- Xiaoning Pan, Ke Cheng, Xiaoling Ji, and Tao Wang. 2020. Manipulating far-zone scattered field by convolution of different types of weight function. *Optics Express* 28, 11 (May 2020), 16869. <https://doi.org/10.1364/oe.394967>
- Giuseppe Papari, Patrizio Campisi, and Nicolai Petkov. 2012. New Families of Fourier Eigenfunctions for Steerable Filtering. *IEEE Transactions on Image Processing* 21, 6 (Jun 2012), 2931–2943. <https://doi.org/10.1109/tip.2011.2179060>
- Ivan Popov and Alex Vitkin. 2016. Dynamic light scattering by flowing Brownian particles measured with optical coherence tomography: impact of the optical system. *Journal of Biomedical Optics* 21, 1 (Jan 2016), 017002. <https://doi.org/10.1117/1.jbo.21.1.017002>
- Iman Sadeghi, Adolfo Munoz, Philip Laven, Wojciech Jarosz, Francisco Seron, Diego Gutierrez, and Henrik Wann Jensen. 2012. Physically-based simulation of rainbows. *ACM Transactions on Graphics* 31, 1 (Jan 2012), 1–12. <https://doi.org/10.1145/2077341.2077344>
- 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
- R. Silverman. 1957. Locally stationary random processes. *IEEE Transactions on Information Theory* 3, 3 (Sep 1957), 182–187. <https://doi.org/10.1109/tit.1957.1057413>
- 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 Lingqi Yan. 2021a. Rendering of Subjective Speckle Formed by Rough Statistical Surfaces. *ACM Transactions on Graphics* (2021), To appear.
- Shlomi Steinberg and Ling-Qi Yan. 2021b. A Generic Framework for Physical Light Transport. *ACM Transactions on Graphics* 40, 4 (Aug 2021), 1–20. <https://doi.org/10.1145/3450626.3459791>
- Yinlong Sun, F. David Fracchia, Mark S. Drew, and Thomas W. Calvert. 2000. *Rendering Iridescent Colors of Optical Disks*. Springer Vienna, 341–352. https://doi.org/10.1007/978-3-7091-6303-0_31
- Gábor Szegő. 1939. *Orthogonal Polynomials*. American Mathematical Society. 432 pages.
- Akimichi Takemura and Kei Takeuchi. 1988. Some Results on Univariate and Multivariate Cornish-Fisher Expansion: Algebraic Properties and Validity. *Sankhyā: The Indian Journal of Statistics, Series A (1961-2002)* 50, 1 (1988), 111–136. <http://www.jstor.org/stable/25050684>
- 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>
- Antoon M. van Dijk and Jean-Bernard Martens. 1997. Image representation and compression with steered Hermite transforms. *Signal Processing* 56, 1 (Jan 1997), 1–16. [https://doi.org/10.1016/s0165-1684\(96\)00146-6](https://doi.org/10.1016/s0165-1684(96)00146-6)
- 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>
- Xun Wang, Zhirong Liu, Kelin Huang, and Deming Zhu. 2016. Spectral changes of cosine-Gaussian-correlated Schell-model beams with rectangular symmetry scattered on a deterministic medium. *Journal of the Optical Society of America A* 33, 10 (Sep 2016), 1955. <https://doi.org/10.1364/josaa.33.001955>
- 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>
- Emil Wolf. 2007. *Introduction to the theory of coherence and polarization of light*. Cambridge University Press, Cambridge.
- E. Wolf, J. T. Foley, and F. Gori. 1989. Frequency shifts of spectral lines produced by scattering from spatially random media. *Journal of the Optical Society of America A* 6, 8 (Aug 1989), 1142. <https://doi.org/10.1364/josaa.6.001142>
- Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering Spectral Microgeometry with Wave Optics. *ACM Trans. Graph.* 37, 4, Article 75 (July 2018), 10 pages. <https://doi.org/10.1145/3197517.3201351>