Rendering: From Ray Optics to Wave Optics

SHLOMI STEINBERG, NVIDIA, United States and University of Waterloo, Canada
RAVI RAMAMOORTHI, NVIDIA, United States and University of California San Diego, United States
BENEDIKT BITTERLI, NVIDIA, United States
EUGENE D'EON, NVIDIA, New Zealand
LING-QI YAN, University of California, Santa Barbara, United States
MATT PHARR, NVIDIA, United States

Fig. S1. **From ray optics to wave optics.** A scene rendered with our technique, showcasing various wave effects: (a) a Bornite ore with an interfering layer of copper oxide; (b) a Brazilian Rainbow Boa, whose scales are biological diffraction-grated surfaces; and (c) a Chrysomelidae beetle, whose colour arises due to multilayered interference reflectors in its elytron. The practical contribution of this paper is the ability to render such complex scenes, under rigorous wave-optical light transport, at a performance that surpasses the state-of-the-art by orders-of-magnitude. The objective of this paper is not the appearance reproduction of some material (a "diffractive BRDF"), but an accurate formulation of wave-optical light transport, where light is rigorously treated as waves globally throughout the entire scene. We propose the *generalized ray*: an extension of the classical ray to wave optics. The generalized ray retains the defining characteristics of the ray-optical ray: *locality* and *linearity*. These properties allow the generalized ray to serve as a "point query" of light's behaviour, the same purpose that the classical ray fulfils in rendering. Generalized rays enable the application of backward (sensor-to-source) light transport and sophisticated sampling techniques, which are impossible with the state-of-the-art. We indicate resolution and samples-per-pixel (spp) count in all figures rendered using our method. While these figures showcase converged (high spp) results, our implementation also allows interactive rendering of all these scenes at 1 spp. Frame times (at 1 spp) for interactive rendering are indicated. See our supplemental material for the implementation, and additional renderings and videos.

Authors' Contact Information: Shlomi Steinberg, p@shlomisteinberg.com, NVIDIA, San Francisco, United States and University of Waterloo, Ontario, Canada; Ravi Ramamoorthi, ravir@cs.ucsd.edu, NVIDIA, San Francisco, United States and University of California San Diego, California, United States; Benedikt Bitterli, benedikt.bitterli@gmail.com, NVIDIA, San Francisco, United States; Eugene d'Eon, ejdeon@gmail.com, NVIDIA, Wellington, New Zealand; Ling-Qi Yan, lingqi@cs.ucsb.edu, University of California, Santa Barbara, Santa Barbara, California, United States; Matt Pharr, matt@pharr.org, NVIDIA, San Francisco, United States.

S1 Introduction

Under the ray picture of light, light consists of "luminous" corpuscles. As such a particle evolves by propagating and interacting with optical systems, it traces a "light ray". A particle's position and *momentum*, i.e. the particle's direction of propagation, serve as a complete description of the light ray at a particular instant. Therefore, it is convenient to study the dynamics of ray optics in *phase space*: a $2n$-dimensional space that consists of n independent position coordinates, as well as n momenta coordinates (often referred to as the *canonically conjugate* variables). A ray, at a particular instant,

corresponds to a point in phase-space. This phase-space pictorial view of light is adopted, sometimes implicitly, by rendering theory: we perform point queries in phase space by tracing rays from a particular point, in a particular direction.

The concept of "*locality*" then becomes central to our discussion: Ray optics permits a precise, simultaneous knowledge of position and momentum. This perfect localization is what enables us to make use of spatial-subdivision acceleration structures for ray tracing, even achieving real-time performance.

On the other hand, under wave optics such locality is not possible. In wave optics, the basic descriptor of light is the wave function, which is the spatial function of the complex excitations of the underlying electric field, and the momentum space becomes the Fourier space. It is well known that a function and its Fourier conjugate (the Fourier transformed function) cannot both have finite support, leading to an *uncertainty relation*: position and momentum may not be both specified with perfect precision. Therefore, in sharp contrast to ray optics, where the descriptor of light—a ray or a collection of rays—is local, the wave function and its Fourier conjugate serve as a global description of light. This loss of locality in wave optics nullifies our ability to perform simple point queries in phase space, and indeed this inherent uncertainty is a major difficulty in devising a formalism of wave-optics rendering.

A rich history of research focuses on attempts to restore, to a degree, that "grainy" phase-space view of ray optics. Most notably, the *Wigner distribution function* [Wigner 1932] (also known as the Wigner-Ville distribution in mathematics) is a complete descriptor of light that simultaneously provides information about both the spatial and angular spectrum properties of the wave function. Thereby, the Wigner distribution function serves to define the dynamics of wave optics in a phase space. For a more comprehensive discussion on Wigner optics, as well as the role the Wigner distribution function plays in wave and quantum optics, the curious reader is referred to Testorf et al. [2010]; Torre [2005].

In this supplemental material, we will overview the Hamiltonian optics formalism of ray optics. We will then briefly "quantize" ray optics in-order to obtain wave optics. The Wigner distribution function will then be presented, and we will discuss its relevant properties. Then, we will identify the wave-optical analogue of the classical ray, and discuss the formal conditions under which pointwise sampling of the wave-optical phase space is possible. We will also show how optical coherence arises naturally when sampling a collection of such "wave-optical rays".

S2 Ray Optics

Hamiltonian optics are developed from Fermat's principle—the principle of extremal action, which in the optical context means the extremal optical path. Specifically, the path taken by a light ray from point $\vec{\boldsymbol{q}}_1$ to point $\vec{\boldsymbol{q}}_2$ fulfils

$$\delta \int_{\vec{\boldsymbol{q}}_1}^{\vec{\boldsymbol{q}}_2} \mathrm{d}s\, \eta(\vec{\boldsymbol{q}}') = 0\,, \tag{S2.1}$$

with η being the refractive index of the medium and s the arc length. That is, the path where the optical path length (path length times refractive index) is an extremum or is stationary, therefore the ray path must follow the refractive-index gradient:

$$\frac{\mathrm{d}}{\mathrm{d}s}\left[\eta(\vec{\boldsymbol{q}})\frac{\mathrm{d}\vec{\boldsymbol{q}}}{\mathrm{d}s}\right] = \nabla\eta(\vec{\boldsymbol{q}})\,. \tag{S2.2}$$

The above is reminiscent of Newton's second law, hence a ray behaves as a classical point particle, with the refractive-index of the medium serving as the mass of the particle. A force $\nabla\eta$ acts upon this particle, thereby light bends—traces an Eikonal—as it propagates through a refractive-index graded medium.

From Eq. (S2.2) we recognize the light particle's momentum as

$$\vec{\boldsymbol{p}} \triangleq \eta(\vec{\boldsymbol{q}})\frac{\mathrm{d}\vec{\boldsymbol{q}}}{\mathrm{d}s}\,. \tag{S2.3}$$

The momenta $\vec{\boldsymbol{p}}$ are the *canonically conjugate* variables to the position variables $\vec{\boldsymbol{q}}$, and are the optical direction cosines (ray direction scaled by the refractive index). We denote the vector

$$\vec{\boldsymbol{u}}(s) \triangleq \begin{pmatrix}\vec{\boldsymbol{q}}\\ \vec{\boldsymbol{p}}\end{pmatrix} \tag{S2.4}$$

as a *ray*. The ray $\vec{\boldsymbol{u}}$ lives in *phase-space*: a vector space defined as the cartesian product of the position and momentum space. The dynamics of that ray, as it evolves w.r.t. s, are quantified by the *Hamiltonian*

$$H(\vec{\boldsymbol{q}}, \vec{\boldsymbol{p}}\,; s) = -\sqrt{\eta^2(\vec{\boldsymbol{q}}) - p^2}\,, \tag{S2.5}$$

and Hamiltonian's equations

$$\frac{\mathrm{d}q_\beta}{\mathrm{d}s} = \frac{\partial H}{\partial p_\beta} \quad \text{and} \quad \frac{\mathrm{d}p_\beta}{\mathrm{d}s} = -\frac{\partial H}{\partial q_\beta}\,, \tag{S2.6}$$

with $\beta \in \{x, y, z\}$.

The above can be recast in operator notation into

$$\frac{\mathrm{d}}{\mathrm{d}s}\vec{\boldsymbol{u}} = \boldsymbol{\mathcal{H}}\vec{\boldsymbol{u}}\,, \tag{S2.7}$$

$$\text{with} \quad \boldsymbol{\mathcal{H}} \triangleq \sum_{\beta\in\{x,y,z\}} \left[\frac{\partial H}{\partial p_\beta}\frac{\partial}{\partial q_\beta} - \frac{\partial H}{\partial q_\beta}\frac{\partial}{\partial p_\beta}\right]. \tag{S2.8}$$

Eq. (S2.7) is the *ray equation*, an operator-valued differential first-order equation, and $\boldsymbol{\mathcal{H}}$ is the Lie operator associated with the optical Hamiltonian H. The ray equation yields the closely-related Eikonal equation, as well as the Snell's law of refraction and the law of reflection at an interface between two media.

The solution to the ray equation, representing the evolution of the ray from s_0 to some s, can be written via the *ray-transfer operator* $\boldsymbol{\mathcal{T}}$:

$$\vec{\boldsymbol{u}}(s) = \boldsymbol{\mathcal{T}}\vec{\boldsymbol{u}}(s_0)\,, \qquad \text{with} \quad \boldsymbol{\mathcal{T}} \triangleq \mathrm{e}^{(s-s_0)\boldsymbol{\mathcal{H}}}\,. \tag{S2.9}$$

The exponential map above maps the Lie algebra $\{\boldsymbol{\mathcal{H}}\}$ into the corresponding symplectic Lie group of ray-transfer operators.

Linear optical systems and quadratic Hamiltonian. When the light rays propagate roughly in the same direction, say the z-axis, we take a paraxial view: The ray evolution variable s is replaced with z, and $\vec{\boldsymbol{q}}, \vec{\boldsymbol{p}}$ become 2-dimensional vectors that live on the xy-plane at a particular instant $z = z'$ of a ray's evolution. Paraxiality implies

$p_x^2 + p_y^2 \ll \eta^2$, hence the optical Hamiltonian H (Eq. (S2.5)) can be written in the quadratic approximation:

$$H(\vec{\boldsymbol{q}}, \vec{\boldsymbol{p}}\,; z) = \frac{1}{2\eta(\vec{\boldsymbol{q}})} p^2 - \eta(\vec{\boldsymbol{q}})\,. \tag{S2.10}$$

An interesting special case of paraxial optical systems are simple systems, where $\boldsymbol{\mathcal{H}}$ does not dependant on z. Such systems are known as linear optical systems, or "ABCD" systems. The latter refers to the fact that $\boldsymbol{\mathcal{T}}$ can be written in the following block-structural form:

$$\begin{pmatrix} \vec{\boldsymbol{q}}(z) \\ \vec{\boldsymbol{p}}(z) \end{pmatrix} = \boldsymbol{\mathcal{T}} \begin{pmatrix} \vec{\boldsymbol{q}}(z_0) \\ \vec{\boldsymbol{p}}(z_0) \end{pmatrix} = \begin{pmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{C} & \boldsymbol{D} \end{pmatrix} \begin{pmatrix} \vec{\boldsymbol{q}}(z_0) \\ \vec{\boldsymbol{p}}(z_0) \end{pmatrix}, \tag{S2.11}$$

with $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{C}, \boldsymbol{D}$ being 2×2 real matrices, and for non-absorbing systems $|\boldsymbol{\mathcal{T}}| = 1$.

ABCD systems are of particular interest, as they include propagation, and reflection and refraction of light at simple interfaces, as well as curved interfaces (like lenses). For example, propagation through a medium with constant refractive-index, or focusing by a thin lens, admit the following ray-transfer matrices

$$\boldsymbol{\mathcal{T}}_{\text{propagation}} = \begin{pmatrix} 1 & \frac{d}{\eta} \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \boldsymbol{\mathcal{T}}_{\text{thinlens}} = \begin{pmatrix} 1 & 0 \\ \frac{1}{f} & 1 \end{pmatrix}, \tag{S2.12}$$

respectively. We use scalars for the ABCD elements of the matrices above to indicate that these systems are rotationally-invariant. d is the (scaled) distance of propagation, and f is the focal length of the lens.

S2.1 An Ensemble of Rays and Liouville's Theorem

The discussion above centred upon the dynamics of a single ray. We now extend the discussion to a statistical ensemble of rays. Let $\rho(\vec{\boldsymbol{q}}, \vec{\boldsymbol{p}}\,; z)$ be the *ray density* function, which is a probability density function quantifying the statistical distribution of rays over phase space. Given an arbitrary function of position and momentum $f(\vec{\boldsymbol{q}}, \vec{\boldsymbol{p}})$, the average value of f over the entire statistical ensemble of rays is

$$\langle f \rangle = \int \mathrm{d}^2\vec{\boldsymbol{q}}\, \mathrm{d}^2\vec{\boldsymbol{p}}\, f(\vec{\boldsymbol{q}}, \vec{\boldsymbol{p}}) \rho(\vec{\boldsymbol{q}}, \vec{\boldsymbol{p}}\,; z)\,, \tag{S2.13}$$

with the integration over the entire phase space of the system at instant z. The function f can be understood as an "observable", for example, the response of a camera sensor to light, or the reflectivity of a surface.

It can be shown that the dynamics of ρ are

$$\frac{\partial}{\partial z}\rho = -\boldsymbol{\mathcal{H}}\rho \tag{S2.14}$$

$$\frac{\mathrm{d}}{\mathrm{d}z}\rho = 0\,, \tag{S2.15}$$

which in Hamiltonian dynamics are referred to as *Liouville's equation* and *Liouville theorem*, respectively. The above illustrates important physics: Eq. (S2.14) means that the ray density evolves (up to a sign) just as a singular ray. As a mental model, the "optical flow" of light rays in phase-space can be thought of as the motion of an incompressible fluid. The total quantity of that fluid is the optical flux, while the phase-space volume occupied by that fluid is known as the *Étendue*. The Liouville theorem (Eq. (S2.15)) implies that the ray density in phase-space is a conserved quantity (ignoring absorption), both locally and globally:

- The "optical fluid" being incompressible means that Étendue is conserved, i.e. the optical fluid may move around in phase-space, but the volume it occupies is unchanged, leading to **global** conservation of density.
- Given any distinguished ray $(\vec{\boldsymbol{q}}, \vec{\boldsymbol{p}})$, the density $\rho(\vec{\boldsymbol{q}}, \vec{\boldsymbol{p}}\,; z)$ in an infinitesimal volume around that ray can be understood as a *property of that ray*, and propagates along the ray's trajectory, i.e. remains conserved **locally** around that ray as the system evolves.

If, at some particular instant of evolution z, we "scoop" some of the "optical fluid" out of the system, then the Étendue may decrease. Étendue may only increase if we add additional fluid into the system (i.e., inject optical flux).

We may relate the above to classical radiometry: the well-known radiometric radiance is defined as

$$L = \eta^2 \frac{\partial \Phi}{\partial G}\,, \tag{S2.16}$$

that is, the (differential) total quantity of fluid—the optical flux Φ—over the (differential) volume this fluid occupies—the Étendue G. The well-known conservation of basic radiance, viz. L/η^2, in simple, non-absorbing optical systems, is then an immediate consequence of Liouville theorem.

A more comprehensive formulation of Hamiltonian optics can be found in the textbooks: Buchdahl [1993] and Bass et al. [2009].

S3 Wave Optics

It is possible to recover aspects of ray optics as a limiting case of wave optics (specifically, the Helmholtz equation reduces to the Eikonal equation at the limit $\hbar \to 0$). However, wave optics cannot be formulated from ray optics purely via mathematical analysis. Instead, wave optics is typically brought forth from classical ray optics in a manner analogous to how quantum mechanics arises from classical Hamiltonian mechanics, and we briefly retrace these steps: A heuristic approach known as "quantization" (or "wavization"), first proposed separately by both Dirac and Heisenberg in different variations, is designed to obtain a quantum theory from a classical theory, hence quantization is a mapping between the theories. Quantization works by replacing the classical functions $f(q, p)$ on phase space with operators ("observables") $\hat{\boldsymbol{f}}$, acting upon wave functions, as well as replacing the classical dynamic laws (e.g., Eq. (S2.6)) with quantum dynamics.

Our starting point is a general classical non-relativistic particle with Hamiltonian

$$H(\vec{\boldsymbol{q}}, \vec{\boldsymbol{p}}) = \frac{p^2}{2m} + V(\vec{\boldsymbol{q}})\,, \tag{S3.1}$$

where m is mass and V is the potential function. The canonical position and momentum variables (Eq. (S2.3)) are mapped to their operator counterparts

$$q \mapsto \hat{\boldsymbol{q}} \triangleq q \qquad \text{and} \qquad p \mapsto \hat{\boldsymbol{p}} \triangleq -\mathrm{i}\hbar \frac{\partial}{\partial q}\,. \tag{S3.2}$$

The Hamiltonian operator then becomes

$$\hat{\boldsymbol{\mathcal{H}}}(\hat{\boldsymbol{q}}, \hat{\boldsymbol{p}}) = -\frac{\hbar^2}{2m}\nabla^2_{\vec{\boldsymbol{q}}} + V(\vec{\boldsymbol{q}}) \,, \tag{S3.3}$$

with $\nabla^2_{\vec{\boldsymbol{q}}}$ being the Laplacian w.r.t. the spatial variable $\vec{\boldsymbol{q}}$.

Switching to the optical context, we note that the classical quadratic optical Hamiltonian, Eq. (S2.10), is of the form of Eq. (S3.1), with $m = \eta$ and $V = -\eta$. Then, applying the mapping in Equation (S3.2),

$$\hat{\boldsymbol{\mathcal{H}}}(\hat{\boldsymbol{q}}, \hat{\boldsymbol{p}}) = -\frac{\hbar^2}{2\eta(\vec{\boldsymbol{q}})}\nabla^2_{\vec{\boldsymbol{q}}} - \eta(\vec{\boldsymbol{q}}) \tag{S3.4}$$

becomes the quadratic wave-optical Hamiltonian operator.

We denote $\psi(\vec{\boldsymbol{q}}\,;\, t)$ as the *wave function*, with $\vec{\boldsymbol{q}}$ being spatial position and t time. The wave function is a complex function that serves as a descriptor of light in the wave-optical context (it may be understood as the excitations of the underlying electric field). The evolution of the wave function is dictated by a time-evolution operator

$$\psi(\vec{\boldsymbol{q}}\,;\, t) = \hat{\boldsymbol{\mathcal{U}}}(t, t_0)\psi(\vec{\boldsymbol{q}}\,;\, t_0) \,. \tag{S3.5}$$

The time-evolution operator must fulfil the *Schrödinger wave equation*:

$$\mathrm{i}\hbar\frac{\partial}{\partial t}\hat{\boldsymbol{\mathcal{U}}} = \hat{\boldsymbol{\mathcal{H}}}\hat{\boldsymbol{\mathcal{U}}} \,, \tag{S3.6}$$

which takes a form reminiscent of its classical counterpart (Eq. (S2.7)), with time t now playing the role of the classical paraxial system evolution variable z. The Helmholtz equation of classical wave optics can be derived from the wave equation above.

We may identify the momentum space as the Fourier-conjugate of the position space: recognizing the eigenfunctions of $\hat{\boldsymbol{p}}$ as $\mathrm{e}^{\mathrm{i}\vec{\boldsymbol{q}}\cdot\vec{\boldsymbol{k}}}$, we may write

$$\psi(\vec{\boldsymbol{q}}) = \frac{1}{(2\pi)^{\frac{3}{2}}}\int \mathrm{d}^3\vec{\boldsymbol{k}}\,\tilde{\psi}\left(\vec{\boldsymbol{k}}\right)\mathrm{e}^{\mathrm{i}\vec{\boldsymbol{q}}\cdot\vec{\boldsymbol{k}}} \,, \tag{S3.7}$$

with the appropriate normalization constant added. The above is simply the inverse Fourier transform of $\tilde{\psi}(\vec{\boldsymbol{k}})$. Therefore, in wave optics it is often convenient to introduce the frequency operator

$$\hat{\boldsymbol{k}} \triangleq \frac{1}{\hbar}\hat{\boldsymbol{p}} \,, \tag{S3.8}$$

and we refer to $\tilde{\psi}(\vec{\boldsymbol{k}})$ as the Fourier-conjugate of the wave function, with $\vec{\boldsymbol{k}} = \frac{1}{\hbar}\vec{\boldsymbol{p}}$ being the *wavevector*, i.e. $|\vec{\boldsymbol{k}}| = \eta\frac{2\pi}{\lambda}$ is the *wavenumber*, where λ is the wavelength.

Uncertainty relation. Without loss of generality, assume that the signals $\psi, \tilde{\psi}$ are centred (zero mean). The variances of these signals, along a particular axis, say x, are

$$\sigma^2_{qx} \triangleq \int \mathrm{d}^3\vec{\boldsymbol{q}}\, q_x^2\left|\psi(\vec{\boldsymbol{q}})\right|^2 \,, \tag{S3.9}$$

$$\sigma^2_{kx} \triangleq \int \mathrm{d}^3\vec{\boldsymbol{k}}\, k_x^2\left|\tilde{\psi}\left(\vec{\boldsymbol{k}}\right)\right|^2 \,. \tag{S3.10}$$

Then, the Fourier relation between position and momentum, outlined by Eq. (S3.7), and a bit of analysis, gives rise to the important *uncertainty relation*:

$$\sigma_{qx}\sigma_{kx} \geq \frac{1}{2} \,. \tag{S3.11}$$

The uncertainty relation implies that the wave function and its conjugate cannot both be precisely localized in space.

ABCD optical systems and line-spread kernels. Let $\psi(\vec{\boldsymbol{q}}\,;\, z)$ be some wave function, under the paraxial approximation. Under the special case where the Hamiltonian $\hat{\boldsymbol{\mathcal{H}}}$ is not z-dependant (i.e., linear optical systems), the solution to the evolution of the system (Eq. (S3.5))

$$\hat{\boldsymbol{\mathcal{U}}}(z, z_0) = \exp\left(-\mathrm{i}\frac{z - z_0}{\hbar}\hat{\boldsymbol{\mathcal{H}}}\right) \,, \tag{S3.12}$$

can be rewritten (due to the linearity of the above) via a *line-spread function* acting on the wave function, viz.

$$\psi(\vec{\boldsymbol{q}}\,;\, z) = \int \mathrm{d}^2\vec{\boldsymbol{q}}'\, g(\vec{\boldsymbol{q}}, \vec{\boldsymbol{q}}')\psi(\vec{\boldsymbol{q}}'\,;\, z_0) \,. \tag{S3.13}$$

Note that time t is replaced by z as the system's evolution variable, under the paraxial setting. For Hamiltonians that admit only quadratic monomials in $\hat{\boldsymbol{q}}, \hat{\boldsymbol{p}}$, it can be shown that [Torre 2005]

$$g(q, q') = \sqrt{\frac{-\mathrm{i}}{2\pi\hbar B}}\exp\left[\mathrm{i}\frac{1}{2\hbar B}\left(Dq^2 + Aq'^2 - 2qq'\right)\right] \,, \tag{S3.14}$$

where separation into dimensions is implied. It can be shown that the conjugate $\tilde{\psi}$ transforms in similar manner to ψ, but the k-space (frequency space) *ABCD* parameters relate to the q-space (position space) via

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix}_k = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}\begin{pmatrix} A & B \\ C & D \end{pmatrix}_q\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \,. \tag{S3.15}$$

The above is the kernel of a linear canonical transform, which generalizes Fresnel transforms and fractional Fourier transforms. Therefore, any diffraction problem that can be solved via Fourier optics tools or the Huygens-Fresnel principle is, in fact, an ABCD system. The parametrization of the ABCD variables, from the ray-transfer matrix (Eq. (S2.11)) to the wave-optical line-spread function (Eq. (S3.14)), changes as $B \mapsto \hbar B$ and $C \mapsto \frac{1}{\hbar}C$, due to the remapping from p-space to k-space (Eq. (S3.8)).

Summary. As we transition from ray optics to wave optics, the position-momentum identification of a ray $\vec{\boldsymbol{u}} = \left(\vec{\boldsymbol{q}}, \vec{\boldsymbol{p}}\right)$ is replaced by the optical wave function and its conjugate $\left(\psi, \tilde{\psi}\right)$ as the descriptor of light. However, while under ray optics the precise and simultaneous specification of a ray's position and momentum is possible, the uncertainty relation implies that in wave optics, such a local specification is not possible. Therefore, while a ray is a *local* descriptor of light, the wave function and its conjugate serve as a *global* descriptor. Indeed, a (non-zero) wave function ψ or its conjugate $\tilde{\psi}$ will always admit infinite support.

The dynamics of the relevant systems are governed by the ray equation or wave equation. The special case of ABCD systems are of special interest for us, as these include the majority of the light transport (not accounting for interaction with materials) around a typical scene (perfect reflections and refractions, including curved surfaces, like lenses, and propagation in media with constant or slowly-varying refractive-index).

In attempt to regain a classical-like view of a wave optical system, where position-momentum pairs can be locally sampled, we will next introduce a wave-optical phase space. Crucially, we will show that point-queries in that wave-optical phase space evolve in identical fashion to their classical counterparts, under interaction by ABCD optical systems.

S3.1 The Wigner Distribution and the Wave-Optical Phase Space

The *Wigner distribution function* (WDF) [Wigner 1932] is defined as

$$\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right) \triangleq \frac{1}{(2\pi)^3}\int \mathrm{d}^3\vec{\boldsymbol{q}}'\;\psi^\star\left(\vec{\boldsymbol{q}}-\frac{\vec{\boldsymbol{q}}'}{2}\right)\psi\left(\vec{\boldsymbol{q}}+\frac{\vec{\boldsymbol{q}}'}{2}\right)\mathrm{e}^{-\mathrm{i}\vec{\boldsymbol{q}}'\cdot\vec{\boldsymbol{k}}}$$
$$\triangleq \frac{1}{(2\pi)^3}\int \mathrm{d}^3\vec{\boldsymbol{k}}'\;\tilde{\psi}^\star\left(\vec{\boldsymbol{k}}-\frac{\vec{\boldsymbol{k}}'}{2}\right)\tilde{\psi}\left(\vec{\boldsymbol{k}}+\frac{\vec{\boldsymbol{k}}'}{2}\right)\mathrm{e}^{\mathrm{i}\vec{\boldsymbol{q}}\cdot\vec{\boldsymbol{k}}'}\;, \quad \text{(S3.16)}$$

with both definitions equivalent. The WDF belongs to the wider *Cohen class* [Cohen 1994] of bilinear signal representations. Being a joint representation of the wave function both in q-space and k-space, the WDF gives rise to a wave-optical phase space. In this subsection, we will analyze the relevant properties of the WDF, and in-turn the wave-optical dynamics in this induced phase space. It is possible to recover the wave function, up to a phase term, from the WDF via an inverse transform.

When ψ is understood as a stochastic process—a statistical ensemble of waves—then a definition of the WDF in terms of the ensemble average is possible:

$$\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right) \triangleq \frac{1}{(2\pi)^3}\int \mathrm{d}^3\vec{\boldsymbol{q}}'\;\mathcal{C}\left(\vec{\boldsymbol{q}}-\frac{\vec{\boldsymbol{q}}'}{2},\vec{\boldsymbol{q}}+\frac{\vec{\boldsymbol{q}}'}{2}\right)\mathrm{e}^{-\mathrm{i}\vec{\boldsymbol{q}}'\cdot\vec{\boldsymbol{k}}}\;, \quad \text{(S3.17)}$$

where $\mathcal{C}$ is the *cross-spectral density* of light: the space-frequency formulation of optical coherence. Clearly, as the WDF and the cross-spectral density function are Fourier pairs, they contain the same information, and one can be recovered unequivocally from the other. For completeness, we explicitly note the inverse transform:

$$\mathcal{C}\left(\vec{\boldsymbol{q}}-\tfrac{1}{2}\vec{\boldsymbol{x}},\vec{\boldsymbol{q}}+\tfrac{1}{2}\vec{\boldsymbol{x}}\right) = \int \mathrm{d}^3\vec{\boldsymbol{k}}'\;\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}'\right)\mathrm{e}^{\mathrm{i}\vec{\boldsymbol{k}}'\vec{\boldsymbol{x}}}\;, \quad \text{(S3.18)}$$

or, equivalently, if we define $\vec{\boldsymbol{q}}_{1,2} = \vec{\boldsymbol{q}} \mp \frac{1}{2}\vec{\boldsymbol{x}}$:

$$\mathcal{C}\left(\vec{\boldsymbol{q}}_1,\vec{\boldsymbol{q}}_2\right) = \int \mathrm{d}^3\vec{\boldsymbol{k}}'\;\mathscr{W}\left(\frac{\vec{\boldsymbol{q}}_1+\vec{\boldsymbol{q}}_2}{2},\vec{\boldsymbol{k}}'\right)\mathrm{e}^{\mathrm{i}\vec{\boldsymbol{k}}'\left(\vec{\boldsymbol{q}}_2-\vec{\boldsymbol{q}}_1\right)}\;. \quad \text{(S3.19)}$$

It should be stressed that Eq. (S3.16) and Eq. (S3.17) are employed under different contexts: the former when we deal with a deterministic wave function, while the latter when the underlying field is modelled as a stochastic process. For more information about optical coherence theory, see Mandel and Wolf [1995]; Wolf [2007].

Given an arbitrary observable $\hat{\boldsymbol{f}}(\hat{\boldsymbol{q}},\hat{\boldsymbol{k}})$, its expectation value is

$$\left\langle\hat{\boldsymbol{f}}\right\rangle_\psi = \left\langle\psi\middle|\hat{\boldsymbol{f}}\middle|\psi\right\rangle = \int \mathrm{d}^3\vec{\boldsymbol{q}}\;\psi^\star(\vec{\boldsymbol{q}})\hat{\boldsymbol{f}}\left(\hat{\boldsymbol{q}},\hat{\boldsymbol{k}}\right)\psi(\vec{\boldsymbol{q}})\;. \quad \text{(S3.20)}$$

It is possible to map the observable $\hat{\boldsymbol{f}}$ to its corresponding "classical" phase-space function $f(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}})$ via the *Wigner-Weyl transform* [Cohen 1966]. Given such a pair, $\hat{\boldsymbol{f}}$ and f, the expectation value of the observable, i.e. Eq. (S3.20), can be recast as

$$\left\langle\hat{\boldsymbol{f}}\right\rangle_\psi = \int \mathrm{d}^3\vec{\boldsymbol{q}}\,\mathrm{d}^3\vec{\boldsymbol{k}}\;f\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right)\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right), \quad \text{(S3.21)}$$

which takes a similar form to the expectation of an observable w.r.t. the classical ray density ρ (Eq. (S2.13)). Note that Eq. (S3.20) is formulated in terms of operators, while Eq. (S3.21) is written in terms of c-functions, typically yielding a simpler expression that is more amenable to analytic tools. The WDF then serves a role similar to the classical ray density ρ: it allows us to "ask wave-optical questions", but in a manner resembling classical phase-space queries.

The properties of Wigner distribution function. The WDF fulfils most of the postulates expected from a phase-space density function.

(I) Realness — $\mathscr{W} \in \mathbb{R}$.

(II) Marginals — the position and momentum densities are the corresponding marginals of the WDF:

$$\left|\psi(\vec{\boldsymbol{q}})\right|^2 = \int \mathrm{d}^3\vec{\boldsymbol{k}}\;\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right) \quad \text{(S3.22)}$$

$$\left|\tilde{\psi}\left(\vec{\boldsymbol{k}}\right)\right|^2 = \int \mathrm{d}^3\vec{\boldsymbol{q}}\;\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right). \quad \text{(S3.23)}$$

(III) Unit measure — if the wave function is normalized, viz. $\int \mathrm{d}^3\vec{\boldsymbol{q}}\,|\psi(\vec{\boldsymbol{q}})|^2 = 1$ then the WDF integrates to one over the entire phase space:

$$\int \mathrm{d}^3\vec{\boldsymbol{q}}\,\mathrm{d}^3\vec{\boldsymbol{k}}\;\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right) = 1\;. \quad \text{(S3.24)}$$

The converse holds as well. In general, the WDF can be normalized as $\int \mathrm{d}\vec{\boldsymbol{q}}\,\mathrm{d}\vec{\boldsymbol{k}}\,\mathscr{W} = 0$ if and only if $\psi \equiv 0$.

(IV) Galilei invariance — the WDF is invariant under Galilean transformations:

$$\psi'(\vec{\boldsymbol{q}}) = \psi(\vec{\boldsymbol{q}}+\vec{\boldsymbol{q}}') \quad\Longrightarrow\quad \mathscr{W}'\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right) = \mathscr{W}\left(\vec{\boldsymbol{q}}+\vec{\boldsymbol{q}}',\vec{\boldsymbol{k}}\right) \quad \text{(S3.25)}$$

$$\tilde{\psi}'\left(\vec{\boldsymbol{k}}\right) = \tilde{\psi}\left(\vec{\boldsymbol{k}}+\vec{\boldsymbol{k}}'\right) \quad\Longrightarrow\quad \mathscr{W}'\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right) = \mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}+\vec{\boldsymbol{k}}'\right). \quad \text{(S3.26)}$$

(V) Support — Given convex $S_q, S_k \subseteq \mathbb{R}^3$ such that

$$\forall\vec{\boldsymbol{q}} \notin S_q,\;\psi(\vec{\boldsymbol{q}}) = 0 \qquad \text{and} \qquad \forall\vec{\boldsymbol{k}} \notin S_k,\;\tilde{\psi}\left(\vec{\boldsymbol{k}}\right) = 0\,,$$

the WDF vanishes outside these volumes as well:

$$\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right) \neq 0 \qquad \text{only if} \quad \left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right) \in S_q \bigtimes S_k\;. \quad \text{(S3.27)}$$

That is, the support of the WDF in q-space and k-space is the support of ψ and $\tilde{\psi}$, respectively.

(VI) Liouville transformation laws — under the paraxial approximation, given a quadratic Hamiltonian (with only quadratic monomials in $\hat{\boldsymbol{q}}$, $\hat{\boldsymbol{p}}$), the WDF obeys:

$$\frac{\partial}{\partial z}\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\;;z\right) = -\boldsymbol{\mathcal{H}}\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\;;z\right) \quad \text{(S3.28)}$$

$$\frac{\mathrm{d}}{\mathrm{d}z}\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\;;z\right) = 0\,, \quad \text{(S3.29)}$$

i.e. the Liouville's equation and Liouville theorem of Hamiltonian mechanics, viz. Eqs. (S2.14) and (S2.15), and note that $\boldsymbol{\mathcal{H}}$ above is the classical Hamiltonian of ray optics (Eq. (S2.8)). Also note that, as before, under the paraxial setting z replaces t as the system

evolution variable, and the q and k-spaces are now 2-dimensional, meaning the phase space becomes 4-dimensional.

(VII) Superposition — Given wave functions ψ_1 and ψ_2, the WDF of the superposition $\psi = \psi_1 + \psi_2$ is

$$\mathscr{W} = \mathscr{W}_1 + \mathscr{W}_2 + 2\,\mathrm{Re}\,\mathscr{W}_{12}\,, \tag{S3.30}$$

where $\mathscr{W}_{12}$ is the *cross-term*:

$$\mathscr{W}_{12}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right) \triangleq \frac{1}{(2\pi)^3}\int \mathrm{d}^3\vec{\boldsymbol{q}}'\;\psi_1^\star\left(\vec{\boldsymbol{q}}-\frac{\vec{\boldsymbol{q}}'}{2}\right)\psi_2\left(\vec{\boldsymbol{q}}+\frac{\vec{\boldsymbol{q}}'}{2}\right)\mathrm{e}^{-\mathrm{i}\vec{\boldsymbol{q}}'\cdot\vec{\boldsymbol{k}}}\,. \tag{S3.31}$$

The above highlights the bilinearity of the WDF.

Moments. Important information about the underlying wave functions, and the optical beams these wave functions encode, can be gleaned from the WDF moments. The total energy contained in the beam is

$$E \triangleq \int \mathrm{d}^3\vec{\boldsymbol{q}}\,\left|\psi(\vec{\boldsymbol{q}})\right|^2 = \int \mathrm{d}^3\vec{\boldsymbol{q}}\,\mathrm{d}^3\vec{\boldsymbol{k}}\,\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right). \tag{S3.32}$$

Clearly, when we understand the WDF strictly as a (quasi-)probability density function, then we only consider $E = 1$. First-order moments (mean) are

$$\begin{pmatrix}\bar{\boldsymbol{q}}\\ \bar{\boldsymbol{k}}\end{pmatrix} \triangleq \frac{1}{E}\int \mathrm{d}^3\vec{\boldsymbol{q}}\,\mathrm{d}^3\vec{\boldsymbol{k}}\begin{pmatrix}\vec{\boldsymbol{q}}\\ \vec{\boldsymbol{k}}\end{pmatrix}\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right). \tag{S3.33}$$

Second-order moments give information about the gyration of beam energy about the mean, in position and frequency spaces. The second-order moments are grouped into the real, symmetric *moments matrix* of the WDF:

$$\begin{aligned}\boldsymbol{M} &\triangleq \begin{pmatrix} m_{xx} & m_{xy} & m_{xz} & m_{x\tilde{x}} & m_{x\tilde{y}} & m_{x\tilde{z}}\\ m_{yx} & m_{yy} & m_{yz} & m_{y\tilde{x}} & m_{y\tilde{y}} & m_{y\tilde{z}}\\ & & \cdots & \cdots & & \\ m_{\tilde{z}x} & m_{\tilde{z}y} & m_{\tilde{z}z} & m_{\tilde{z}\tilde{x}} & m_{\tilde{z}\tilde{y}} & m_{\tilde{z}\tilde{z}}\end{pmatrix}\\ &\triangleq \frac{1}{E}\int \mathrm{d}^3\vec{\boldsymbol{q}}\,\mathrm{d}^3\vec{\boldsymbol{k}}\left[\begin{pmatrix}\vec{\boldsymbol{q}}\\ \vec{\boldsymbol{k}}\end{pmatrix}-\begin{pmatrix}\bar{\boldsymbol{q}}\\ \bar{\boldsymbol{k}}\end{pmatrix}\right]\left[\begin{pmatrix}\vec{\boldsymbol{q}}\\ \vec{\boldsymbol{k}}\end{pmatrix}-\begin{pmatrix}\bar{\boldsymbol{q}}\\ \bar{\boldsymbol{k}}\end{pmatrix}\right]^{\mathsf{T}}\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right), \end{aligned} \tag{S3.34}$$

where subscripts of the matrix elements $m_{\xi\zeta}$ are unaccented or accented with a tilde to indicate q-space or k-space axes, respectively. The second-order moments on the main diagonal of $\boldsymbol{M}$ provide information about the width of the beam in phase space, i.e. both in position and frequency spaces. Our interest lies primarily in these main diagonal moments. Mixed moments are used in the optical literature to characterize beam twist, curvature as well as beam quality. Furthermore, mixed moments quantify the longitudinal components of the orbital angular momentum.

Transformation of the WDF. The properties above suggest that the WDF can, to a degree, be understood as the classical phase-space density function ρ. A point-query of the wave-optical phase space, viz. $\mathscr{W}(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}})$, then can be understood as a "ray", and we write $\mathscr{W}(\vec{\boldsymbol{u}})$, with $\vec{\boldsymbol{u}} = \left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}\right)$ resembling its ray optical analogue (i.e. Eq. (S2.4)). Property (VI) then implies that the WDF transforms in a manner similar to a classical ray under interaction with an ABCD optical system:

$$\mathscr{W}(\vec{\boldsymbol{u}}\,;z) = \mathscr{W}\left(\boldsymbol{\mathcal{T}}^{-1}\vec{\boldsymbol{u}}\,;z_0\right), \tag{S3.35}$$

where the matrix $\boldsymbol{\mathcal{T}}$ is the appropriate ABCD ray-transfer matrix (Eq. (S2.11)), though note that it should be transformed to q-k representation of the phase space, from the q-p representation of the ray optical phase space, as discussed in Section S3.

Eq. (S3.35) means that a point-query $\vec{\boldsymbol{u}}$ in the wave-optical phase space (a "ray") transforms just as its ray-optical analogue, under quadratic Hamiltonian wave optics. Of particular interest is the fact that the WDF moments matrix (Eq. (S3.34)) also transforms via the ray-transfer matrix, as:

$$\boldsymbol{M}(z) = \boldsymbol{\mathcal{T}}\boldsymbol{M}(z_0)\boldsymbol{\mathcal{T}}^{\mathsf{T}}\,, \tag{S3.36}$$

on interaction with an ABCD optical systems.

The negativity of the WDF. One postulate of a probability density function not fulfilled by the WDF is non-negativity. The WDF may take negative values, a consequence of the uncertainty relation: The quantization process employed to promote the symplectic ray optics to metaplectic wave optics serves to quantize phase space into cells (the volume of which is dictated by the uncertainty relation, Eq. (S3.11)). These cells are not discrete cells with "sharp" boundaries, but overlap and interact with each other, therefore points within a phase-space cell do not constitute mutually-exclusive probability events (violating the σ-additivity of a probability measure), hence the WDF is only a quasi-probabilistic density function.

It can be shown that *anisotropic Gaussian Schell-model* (AGSM) beams are the only class of wave functions that admit non-negative Wigner distribution functions. Furthermore, AGSM beams have the most compact support in phase space (occupy the least phase space volume) relative to any other wave function. Thus, a Gaussian beam can be understood as an elementary construct that is the closest analogue of the classical ray: serving as a form of a "generalized ray" with non-singular extent in position and momentum spaces. It should be noted that a superposition of a pair of AGSM beams does not, in general, yield a non-negative WDF: the bilinearity of the WDF gives rise to cross-terms on superposition (Property (VII) above), and it is these interference terms that are the source of negative values.

See Bastiaans [1978]; Testorf et al. [2010]; Zhang and Levoy [2009] for additional discussion and applications of the WDF in optics.

S4 Wave-Optics Light Transport

The wave-optical phase space that arises via the Wigner distribution function admits attractive properties: it facilitates performing phase-space queries in a manner similar to classical ray optics, and these "rays" transform inline with Liouville's equations for ABCD optical systems. However, the WDF is not non-negative, frustrating its interpretation as an energy density. Furthermore, the WDF tends to be highly oscillatory: a consequence of the Fourier-like relation in the definition of the WDF (Eq. (S3.16)). As an example, consider a sample signal Φ, and its WDF $\mathscr{W}_\Phi$, and let a wave function be composed of two spatially- and frequency-shifted copies of this signal:

$$\psi(\vec{\boldsymbol{q}}) = \Phi(\vec{\boldsymbol{q}}-\vec{\boldsymbol{q}}_1)\mathrm{e}^{\mathrm{i}\vec{\boldsymbol{k}}_1\cdot\vec{\boldsymbol{q}}} + \Phi(\vec{\boldsymbol{q}}-\vec{\boldsymbol{q}}_2)\mathrm{e}^{\mathrm{i}\vec{\boldsymbol{k}}_2\cdot\vec{\boldsymbol{q}}}\,, \tag{S4.1}$$

where $\vec{\boldsymbol{q}}_{1,2}$ and $\vec{\boldsymbol{k}}_{1,2}$ are the spatial and frequency shifts, respectively. Using the shift properties of the WDF (Property (IV)), the WDF of

the wave function above is trivially:

$$\mathscr{W}\left(\vec{q},\vec{k}\right) = \mathscr{W}_{\Phi}\left(\vec{q}-\vec{q}_1,\vec{k}-\vec{k}_1\right) + \mathscr{W}_{\Phi}\left(\vec{q}-\vec{q}_2,\vec{k}-\vec{k}_2\right) + 2\,\mathrm{Re}\left[\mathrm{e}^{\mathrm{i}\left(\vec{k}_1-\vec{k}_2\right)\cdot\vec{q}-\mathrm{i}\left(\vec{q}_1-\vec{q}_2\right)\cdot\vec{k}'}\right]\mathscr{W}_{\Phi}\left(\vec{q}',\vec{k}'\right), \quad \text{(S4.2)}$$

with the shorthands $\vec{q}' = \vec{q} - \frac{1}{2}(\vec{q}_1+\vec{q}_2)$ and $\vec{k}' = \vec{k} - \frac{1}{2}(\vec{k}_1+\vec{k}_2)$. Note the complex exponent in the cross-term above: it is a heavily oscillatory term at optical frequencies ($k \gg 0$), with frequencies that grow greater as the separation in phase space between the two Φ signals increases. Hence, if light is composed of multiple partially-coherent components, as these propagate and their separation increases, the WDF becomes increasingly oscillatory.

The WDF as a "generalized radiance". As a brief aside, we note that the WDF was used to derive wave-optical radiometric quantities, in particular the radiance, first by Walther [1968]. Other definitions of such *generalized radiances* have been proposed, usually using other Cohen class joint space-frequency representations. A generalized radiance—in the form of the WDF—was also used in computer graphics to propagate partially-coherent fields.

However, it was shown that no such generalized radiance fulfils all the expected postulates: for example, it is not non-negative, or isn't conserved on non-paraxial propagation, or it is not a faithful representation of the signal (e.g., Property (II) does not hold). In general, such representations serve only as quasi-probability distributions. Furthermore, being "wasteful" representations (they are of double the dimensionality of the represented signal), in practical applications only a restricted parametrized class of functions are used. But under this constraint, there is no value to using the WDF as opposed to the cross-spectral density of light directly (which then is parameterized by the Fourier-conjugated class of functions). Whichever representation of light we chose to use, once we decide to quantify partially-coherent fields explicitly, we always suffer from the "sampling problem" (see Steinberg et al. [2022] and Section 1 in the paper), where backward path tracing is difficult, as importance sampling light-matter interactions require information about the coherence of light.

Instead of using the WDF as a descriptor of light, we are interested in the phase space that arises via the WDF. We would like to find a wave-optical analogue of the classical ray, formally discuss when the wave-optical phase space can adequately sampled via such "rays", and analyze the dynamics of these "rays". These rays facilitate a coherent-mode decomposition of light, allowing us to reason about the partially-coherent light that is of primary interest for us in rendering in a "classical" manner.

S4.1 Gaussian Beams as Rays

The Husimi Q representation. To combat the unattractive cross-terms that arise in the WDF on superposition of waves, the WDF can be convolved in phase space with a kernel function (a *Cohen kernel*), masking out the interference terms and producing a different representation. It can be shown that a convolution with a multivariate Gaussian, with position and frequency variances satisfying the uncertainty relation (Eq. (S3.11)), produces a representation that is strictly non-negative. The resulting distribution is known as the *Husimi Q distribution*:

$$\mathscr{Q}\left(\vec{q},\vec{k}\right) \triangleq \frac{1}{\pi^3}\int \mathrm{d}^3\vec{q}'\,\mathrm{d}^3\vec{k}'\,\mathscr{W}\left(\vec{q}',\vec{k}'\right)\mathrm{e}^{-\frac{1}{2}\vec{u}'^{\mathsf{T}}\boldsymbol{\Sigma}^{-1}\vec{u}'}, \quad \text{(S4.3)}$$

$$\text{with} \quad \vec{u}' \triangleq \begin{pmatrix}\vec{q}-\vec{q}'\\ \vec{k}-\vec{k}'\end{pmatrix} \quad \text{and} \quad \sqrt{|\boldsymbol{\Sigma}|} = \frac{1}{2^3},$$

where $\boldsymbol{\Sigma}$ is any positive-definite covariance matrix of the Gaussian low-pass filter that fulfils the above.

A wave-optical "ray". Consider the WDF that takes the form of a Dirac delta in phase space, viz. $\mathscr{W}_r = \delta^3(\vec{q}-\vec{q}_0)\,\delta^3(\vec{k}-\vec{k}_0)$, which is an aphysical construct that represents an idealised "ray" at position $\vec{q}_0$ with momentum $\vec{p}_0 = \hbar\vec{k}_0$. We stress that such a WDF is fictitious: it cannot arise from any physically-realizable wave function. However, its corresponding Husimi Q representation $\mathscr{Q}_r$, that arises from $\mathscr{W}_r$ via Eq. (S4.3), is physical. Let the covariance take the block-diagonal form $\boldsymbol{\Sigma} = \mathrm{diag}\{\boldsymbol{\Sigma}_q, \boldsymbol{\Sigma}_k\}$, then:

$$\mathscr{Q}_r\left(\vec{q},\vec{k}\,;\,t_0\right) = \frac{1}{\pi^3}\mathrm{e}^{-\frac{1}{2}\vec{q}'^{\mathsf{T}}\boldsymbol{\Sigma}_q^{-1}\vec{q}'-\frac{1}{2}\vec{k}'^{\mathsf{T}}\boldsymbol{\Sigma}_k^{-1}\vec{k}'}, \quad \text{(S4.4)}$$

which represent the phase space "picture" of the ray at an initial time t_0 of the system evolution. The shorthands $\vec{q}' = \vec{q}-\vec{q}_0$ and $\vec{k}' = \vec{k}-\vec{k}_0$ are the shifted coordinates. Clearly, this system is fully defined by its first 2 moments: the mean $\bar{\boldsymbol{u}}(t_0) = \left(\vec{q}_0,\vec{k}_0\right)$, and the moment matrix (Eq. (S3.34)) $\boldsymbol{M}(t_0) = \boldsymbol{\Sigma}$. The time evolution follows Eq. (S3.36):

$$\bar{\boldsymbol{u}}(t) = \boldsymbol{\mathcal{T}}(t,t')\bar{\boldsymbol{u}}(t'), \qquad \text{and} \quad \text{(S4.5)}$$

$$\boldsymbol{M}(t) = \boldsymbol{\mathcal{T}}(t,t')\boldsymbol{M}(t')\boldsymbol{\mathcal{T}}(t,t')^{\mathsf{T}}, \quad \text{(S4.6)}$$

for $t \geq t'$. For example, on propagation in a medium with a constant refractive-index η, viz.

$$\boldsymbol{\mathcal{T}}_{\text{propagation}}(t,t') = \begin{pmatrix}1 & \frac{t-t'}{\eta}\hbar c\\ 0 & 1\end{pmatrix}, \quad \text{(S4.7)}$$

(elements represent 3×3 matrices) which represents a phase space horizontal shear, with c being the speed of light. Hence, the evolution effectively constitutes propagating the centre-of-mass in phase space (mean) in direction $\vec{k}_0$ and spreading the spatial Gaussian footprint w.r.t. the variance in frequency.

Substitute the "smoothed WDF" $\mathscr{Q}_r$ of an idealised ray (Eq. (S4.4)) into the definition of the WDF (Eq. (S3.16)) and invert the transform:

$$\psi_r(\vec{q}\,;\,t_0) = \frac{1}{\psi_r^{\star}(\vec{q}_0\,;\,t_0)}\int \mathrm{d}^3\vec{k}'\,\mathscr{Q}_r\left(\frac{\vec{q}+\vec{q}_0}{2},\vec{k}'\,;\,t_0\right)\mathrm{e}^{\mathrm{i}\vec{k}'\cdot(\vec{q}-\vec{q}_0)} = \sqrt{\frac{2^3|\boldsymbol{\Sigma}_k|}{\pi^3}}\frac{\mathrm{e}^{-\mathrm{i}\vec{k}_0\cdot\vec{q}'}}{\psi_r^{\star}(\vec{q}_0\,;\,t_0)}\mathrm{e}^{-\frac{1}{8}\vec{q}'^{\mathsf{T}}\boldsymbol{\Sigma}_q^{-1}\vec{q}'-\frac{1}{2}\vec{k}'^{\mathsf{T}}\boldsymbol{\Sigma}_k\vec{k}'}. \quad \text{(S4.8)}$$

The value of the wave function at $\vec{q}_0$ is (up to a constant phase factor) computed via the respective marginal (Property (II) in Section S3.1):

$$\left|\psi_r(\vec{q}_0\,;\,t_0)\right|^2 = \int \mathrm{d}^3\vec{k}'\,\mathscr{Q}_r\left(\vec{q}_0,\vec{k}'\,;\,t_0\right) = \sqrt{\frac{2^3|\boldsymbol{\Sigma}_k|}{\pi^3}}. \quad \text{(S4.9)}$$

Plugging the above into Eq. (S4.8) yields the wave function that is the closest analogue to the Dirac delta in phase space, i.e. a *generalized ray*:

$$\psi_r\left(\vec{\boldsymbol{q}}\,;\bar{\boldsymbol{u}}\right)=\left(\frac{2^3|\boldsymbol{\Sigma}_k|}{\pi^3}\right)^{1/4}\mathrm{e}^{\mathrm{i}\varphi}\mathrm{e}^{-\mathrm{i}\vec{\boldsymbol{k}}_0\cdot\vec{\boldsymbol{q}}'}\,\mathrm{e}^{-\frac{1}{8}\vec{\boldsymbol{q}}'^{\mathsf{T}}\left(\boldsymbol{\Sigma}_q^{-1}+4\boldsymbol{\Sigma}_k\right)\vec{\boldsymbol{q}}'}\,, \tag{S4.10}$$

where we slightly abuse notation and make the mean $\bar{\boldsymbol{u}}=\left(\vec{\boldsymbol{q}}_0,\vec{\boldsymbol{k}}_0\right)$ at current time t explicit, with shifted position shorthand $\vec{\boldsymbol{q}}'=\vec{\boldsymbol{q}}-\vec{\boldsymbol{q}}_0$, as before, and $\varphi\in\mathbb{R}$ being an arbitrary initial phase. The evolution of that wave function to $t\geq t_0$ is dictated by Eqs. (S4.5) and (S4.6). The above should be understood as the *wave function that corresponds to the ray* $\vec{\boldsymbol{u}}=\bar{\boldsymbol{u}}$. Being a Gaussian beam, ψ_r is indeed a subclass of AGSM beams, and it has the most compact support possible in phase space, as discussed.

Coherent-modes phase-space decomposition. It is well-known that an arbitrary function in L^1 can be approximated arbitrary well by a finite sum of shifted Gaussians with identical variance (an immediate consequence of the Wiener's Tauberian theorem). In other words, multivariate Gaussians serve as an overcomplete functional basis. Therefore, the Husimi Q representation $\mathcal{Q}$ of an arbitrary WDF can be written as

$$\mathcal{Q}=\sum_{j=1}^{\infty}E_j\;\mathcal{Q}_r\Big|_{\bar{\boldsymbol{u}}_j,\boldsymbol{M}}\,, \tag{S4.11}$$

i.e. a superposition of the Husimi Q representations of generalized rays, all with the same moment matrix $\boldsymbol{M}$ but shifted via different means $\bar{\boldsymbol{u}}_j$. The moment matrix must fulfil the Husimi Q condition $|\boldsymbol{M}|=\frac{1}{2^3}$, but otherwise is chosen at will, we may set $\boldsymbol{M}(t_0)=\frac{1}{\sqrt{2}}\boldsymbol{I}$ initially, for simplicity. $E_j>0$ are the energies contained in each generalized ray. Positive energies are only possible because $\mathcal{Q}$ is always non-negative.

S4.2 Optical coherence

It is insightful to study how partially-coherent field effects arise under our formulation. Let $\mathcal{C}$ be the cross-spectral density of light, and $\mathscr{W}$ its corresponding WDF. The 3×3 spatial-coherence covariance matrix—termed the *shape matrix*—around a spatial point $\vec{\boldsymbol{q}}$ can be written as:

$$\begin{aligned}\boldsymbol{\Theta}(\vec{\boldsymbol{q}})=&\frac{1}{\mathcal{C}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{q}}\right)}\int\mathrm{d}^3\vec{\boldsymbol{q}}'\,\vec{\boldsymbol{q}}'\vec{\boldsymbol{q}}'^{\mathsf{T}}\,\mathcal{C}\left(\vec{\boldsymbol{q}}-\frac{1}{2}\vec{\boldsymbol{q}}',\vec{\boldsymbol{q}}+\frac{1}{2}\vec{\boldsymbol{q}}'\right)\\=&\frac{1}{\left|\psi(\vec{\boldsymbol{q}})\right|^2}\int\mathrm{d}^3\vec{\boldsymbol{q}}'\,\vec{\boldsymbol{q}}'\vec{\boldsymbol{q}}'^{\mathsf{T}}\int\mathrm{d}^3\vec{\boldsymbol{k}}'\,\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}'\right)\mathrm{e}^{\mathrm{i}\vec{\boldsymbol{q}}'\cdot\vec{\boldsymbol{k}}'}\,.\end{aligned} \tag{S4.12}$$

Formally-interchange the orders of integration, and note that

$$\int\mathrm{d}^3\vec{\boldsymbol{q}}'\,\vec{\boldsymbol{q}}'\vec{\boldsymbol{q}}'^{\mathsf{T}}\,\mathrm{e}^{\mathrm{i}\vec{\boldsymbol{q}}'\cdot\vec{\boldsymbol{k}}'}=-(2\pi)^3\frac{\partial^2}{\partial\vec{\boldsymbol{k}}'^2}\,\delta^3\left(\vec{\boldsymbol{k}}'\right), \tag{S4.13}$$

i.e., the Hessian matrix of the Dirac delta. Then, for "well-behaved" $\mathscr{W}$:

$$\begin{aligned}\boldsymbol{\Theta}(\vec{\boldsymbol{q}})=&-\frac{(2\pi)^3}{\left|\psi(\vec{\boldsymbol{q}})\right|^2}\int\mathrm{d}^3\vec{\boldsymbol{k}}'\,\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}'\right)\frac{\partial^2}{\partial\vec{\boldsymbol{k}}'^2}\,\delta^3\left(\vec{\boldsymbol{k}}'\right)\\=&-\frac{(2\pi)^3}{\left|\psi(\vec{\boldsymbol{q}})\right|^2}\left[\frac{\partial^2}{\partial\vec{\boldsymbol{k}}'^2}\mathscr{W}\left(\vec{\boldsymbol{q}},\vec{\boldsymbol{k}}'\right)\right]_{\vec{\boldsymbol{k}}'=0}\,,\end{aligned} \tag{S4.14}$$

that is, the Hessian (w.r.t. the frequency variable) of the WDF, evaluated at $\vec{\boldsymbol{q}}$ and $\vec{\boldsymbol{k}}=0$.

As mentioned, the cross-spectral density function and the WDF contain the same information (being Fourier-transform pairs), however we have shown that *spatial-coherence is dictated by the behaviour of the WDF in frequency-space only*, furthermore, for Gaussian signals the covariance of spatial coherence around a point $\vec{\boldsymbol{q}}$ is proportional to the inverse of covariance of angular spread of light at $\vec{\boldsymbol{q}}$.

S5 Analysis of Generalized Rays

In Section 4 we have derived our primary contributions: the theory of backward wave-optical light transport with generalized rays. We now analyse and validate these results.

Linearity. All formulations in the paper are linear: (i) Sourcing equations, Eqs. (16) and (17), describe a linear (incoherent) superposition of the measured intensities of each detector element; (ii) the rendering equation, Eq. (23), is linear; and, (iii) measurement of a generalized ray, i.e. integration over the sourcing distribution (Eq. (24)), as well as accumulation of these intensities are also linear operations. Generalized rays always carry positive intensities. We discuss the linearity of generalized ray further, from the perspective of Shannon sampling, in Section S5.1.

Locality. We need to show that all integrals in our formulae can be restricted to a well-defined finite integration region. A generalized ray is a sharply-peaked Gaussian, thus in practice we ignore its tails and assume that $\psi_{\beta,\rho}(\vec{\boldsymbol{r}})=0$ when $|\vec{\boldsymbol{r}}-\vec{\boldsymbol{r}}_0|>\varrho$, for some $\varrho>0$. Clearly, this can be done to arbitrary precision: the tail mass of $\psi_{\beta,\rho}$ decays rapidly as a function of ϱ. Therefore, all integrations are over the support of a generalized ray: in the sourcing (Eqs. (16) and (18)), measurement (Eq. (24)), and interaction operator acting upon a generalized ray (Eq. (11)), viz. $\mathcal{K}^{\mathrm{t}}\{g\}$, formulae are confined to the (finite) spatial extent of $\psi_{\beta,\rho}$.

We will now show that we may always restrict the integration that defines an arbitrary interaction kernel K^{t} to a finite region. Consider the action of that kernel (Eq. (11)) on a generalized ray g:

$$\begin{aligned}\mathscr{W}\triangleq&\,\mathcal{K}^{\mathrm{t}}\left\{g_{\beta,\rho}\left(\vec{\boldsymbol{r}}',\vec{\boldsymbol{k}}'\,;\vec{\boldsymbol{r}}_0,\vec{\boldsymbol{k}}_0\right)\right\}\\=&\int\mathrm{d}\vec{\boldsymbol{r}}'\,\mathrm{d}\vec{\boldsymbol{k}}'\,K^{\mathrm{t}}\left(\vec{\boldsymbol{r}},\vec{\boldsymbol{r}}',\vec{\boldsymbol{k}},\vec{\boldsymbol{k}}'\right)g_{\beta,\rho}\left(\vec{\boldsymbol{r}}',\vec{\boldsymbol{k}}'\,;\vec{\boldsymbol{r}}_0,\vec{\boldsymbol{k}}_0\right).\end{aligned} \tag{S5.1}$$

Substitute the definition of the kernel (Eq. (2)) and rewrite the generalized ray g using the definition of the WDF (Eq. (4)) via its wave function $\psi_{\beta,\rho}$ (Eq. (15)). With some basic algebra and proper variable changes the above becomes

$$\begin{aligned}\mathscr{W}=\int\mathrm{d}\vec{\boldsymbol{y}}'\,\mathrm{d}\vec{\boldsymbol{y}}''\,\mathrm{d}\vec{\boldsymbol{x}}_o\psi^{\star}_{\beta,\rho}\left(\vec{\boldsymbol{y}}'\right)h\left(\vec{\boldsymbol{y}}',\vec{\boldsymbol{r}}+\tfrac{1}{2}\vec{\boldsymbol{x}}_o\right)\\\times\psi_{\beta,\rho}\left(\vec{\boldsymbol{y}}''\right)h^{\star}\left(\vec{\boldsymbol{y}}'',\vec{\boldsymbol{r}}-\tfrac{1}{2}\vec{\boldsymbol{x}}_o\right)\mathrm{e}^{-\mathrm{i}\vec{\boldsymbol{k}}\cdot\vec{\boldsymbol{x}}_o}\,.\end{aligned} \tag{S5.2}$$

That is, the system's optical response function $h(\vec{\boldsymbol{r}}_o,\vec{\boldsymbol{r}}_i)$ only contributes to the result $\mathscr{W}$ when $\psi_{\beta,\rho}(\vec{\boldsymbol{r}}_o)$ is non-negligible. Apply the relation $h(\vec{\boldsymbol{r}}_o,\vec{\boldsymbol{r}}_i)=h^{\star}(\vec{\boldsymbol{r}}_i,\vec{\boldsymbol{r}}_o)$ to Eq. (S5.2), and we also deduce that $h(\vec{\boldsymbol{r}}_o,\vec{\boldsymbol{r}}_i)$ only contributes when $\psi_{\beta,\rho}(\vec{\boldsymbol{r}}_i)$ is non-negligible. Hence,

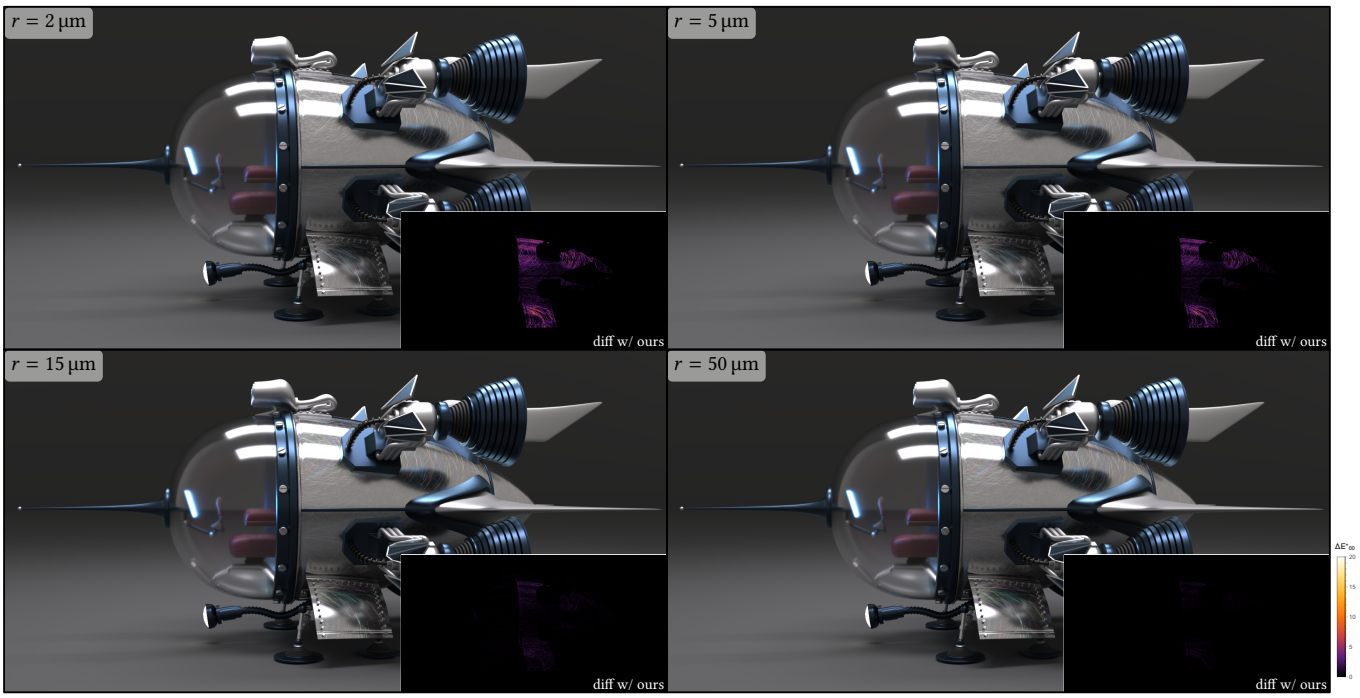

Fig. S2. Rendering of the spaceship (Fig. 1 in the paper) with Werner et al. [2017]. Difference insets are with respect to our method.

we may always replace the optical response function h with

$$\bar{h}(\vec{r}_o, \vec{r}_i) \triangleq \begin{cases} h(\vec{r}_o, \vec{r}_i) & \text{if} \quad |\vec{r}_o| < \varrho \text{ and } |\vec{r}_i| < \varrho \\ 0 & \text{otherwise} \end{cases} \tag{S5.3}$$

in the definition of a diffraction kernel (Eq. (2)), thereby limiting the integration to the spatial extent ϱ of the incident generalized ray, while remaining accurate to arbitrarily good precision. We have shown that our formalism achieves weak-locality.

Compare the above with related work: If we were to replace the generalized rays with plane waves, viz. $\psi_{\beta,\rho} \propto \exp(-\mathrm{i}\vec{k}\cdot\vec{r})$, locality would no longer be recoverable as the integration in Eq. (S5.2) must happen over the entire domain. Similarly, consider a Wigner-based perfectly-local formalism (for example Cuypers et al. [2012]) where we set $g = \delta(\vec{r} - \vec{r}_0)\,\delta(\vec{k} - \vec{k}_0)$. Then, Eq. (S5.1) becomes $\mathscr{W} = K^{-\mathrm{t}}$, and the integration region in the definition of K (Eq. (2)) can no longer be constrained (without sacrificing linearity).

Completeness. Given a photoelectric detector, whose detectable states (its WDF) can be written as Eq. (8), our derivations in Section 4 are exact. We stress that essentially all detectors of interest work via the process of photoelectric detection [Leonhardt 1997, Chapters 3 and 4]. Arbitrary detector geometry and detection properties, quantified by $\mathcal{D}$, are supported. Interactions of the WDF with the scene via a diffraction kernel K, as in Eqs. (1) and (2), is a general formalism [Testorf et al. 2010], and no restrictions are placed on the total interaction operator $\mathcal{K}^{-\mathrm{t}}$. Likewise, the sourcing distribution $\mathscr{W}_s$ can be arbitrary.

Because (multivariate) Gaussians serve as an overcomplete functional basis, an arbitrary Husimi Q distribution can be written as a finite superposition of Gaussians to arbitrarily good precision. We discuss this further in our supplemental material. Therefore, the recursive light transport process, formalised by Eq. (23), is well defined. The restriction of a generalized ray, as well as the definition of the interaction kernel (Eq. (2)), to a finite spatial region can also be done to an arbitrarily good precision. Therefore, our formalism is complete: able to reproduce any wave-optical effect observable by a photoelectric detector.

S5.1 Linearity of Generalized Rays

Consider Young's iconic double-slit interferometer (illustrated in Fig. S5a), and assume we use a coherent laser source. This experiment will be used to (i) lend insight into how generalized rays are always able to maintain linearity, even when the incident illumination is perfectly coherent; and, (ii) numerically validate our formalism (see Section S5.2). In addition, while our practical interactive rendering algorithm that we introduce in the paper (Section 5.1) neglects free-space diffractions, this experiment demonstrates that our formalism may reproduce such effects.

As the coherent light passes through the slits, it diffracts around the slits, resulting in a (coherent) superposition phasor φ at the screen. However, we do not observe φ at singular points, but only over regions, "pixels", on the screen (blue line on the screen in Fig. S5a). The spatial extent of these pixels must be positive, due to the uncertainty relation [Mandel and Wolf 1995]: we are never able to resolve light at infinite resolution. Because interference is averaged out over the spatial extent of a pixel, oscillations of φ that are more rapid than the pixel's extent do not contribute to observable interference effects.

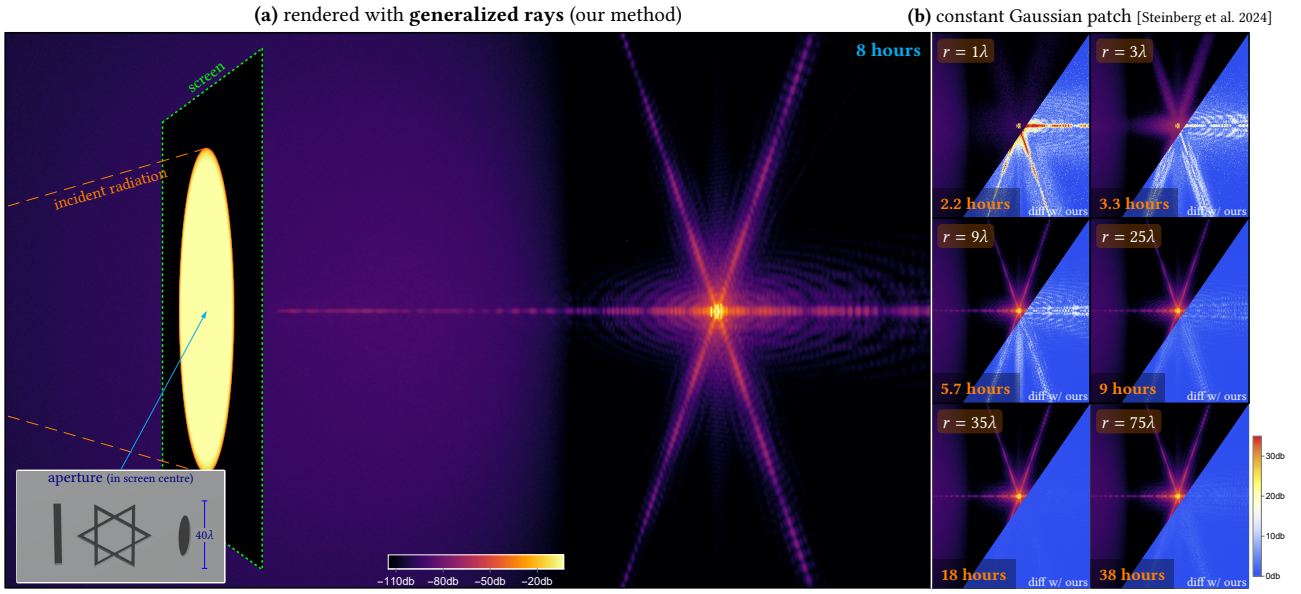

Fig. S3. **Locality in wave optics II.** Long-wavelength radiation impinges upon a conductive screen. The incident beam's extent is illustrated via red dashed lines, and the screen is marked with a dotted green outline. At the centre of the screen we cut out a complex aperture (visualized in the bottom left). Light diffracts through that aperture, giving rise to a diffraction pattern on the right wall. We visualize the colour-coded irradiance (in decibels) impinging upon scene surfaces. The screen is placed 1300λ from the right wall, and the right wall's edge length is 2000λ. The free-space diffraction BSDF is computed via the method by Steinberg et al. [2024]. We render the scene with (b) a constant BSDF integration path radius r, and compare with (a) our method. Accurately localizing light is imperative: when r is too small highly incorrect results are produced; on the other hand, when r is too big (i) BSDF integration becomes unnecessary expensive, and (ii) the sampling problem arises (see Fig. S4 for BSDF lobe visualizations). Generalized rays quantify the optimal integration patch, dynamically throughout the scene. All images are rendered at 1170×700 resolution with 64 samples per pixel.

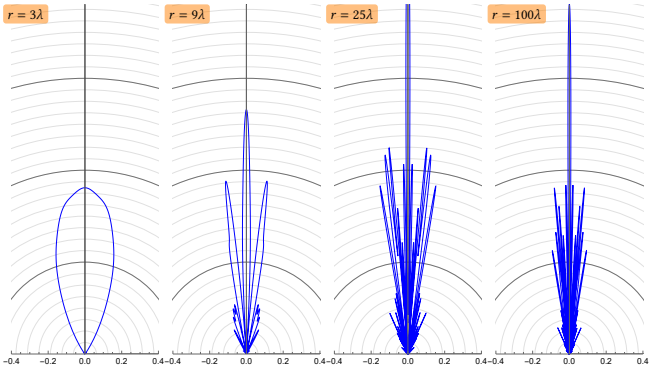

Fig. S4. Log-scale plots of free-space diffraction BSDF slices (see Fig. S3). Observe the *sampling problem*: there is a marked difference in the BSDF for $r = 25\lambda$ and $r = 100\lambda$, even though there is little difference in the observed optical response (diffraction patterns appear almost the same).

Let points $\vec{s}_{1,2}$ located on the slits act as point sources (i.e., the Huygens–Fresnel principle). Given identical peak amplitudes, their superposition phasor at a point $\vec{r}$ on the screen can be written as

$$\varphi \approx \frac{e^{-ikz}}{z}\left[\exp\left(-i\frac{k}{z}\vec{s}_1^{\perp}\cdot\vec{r}^{\perp}\right)+\exp\left(-i\frac{k}{z}\vec{s}_2^{\perp}\cdot\vec{r}^{\perp}\right)\right], \tag{S5.4}$$

where k is the wavenumber, as before, the superscript $\perp$ denotes projection upon the plate plane ($z = 0$), and we made the typical Fraunhofer (far-field) approximation [Born and Wolf 1999] (note, we only make this approximation for the illustrative analysis here, the results in Figs. S5b and S6 were obtained using exact formulae). The observed intensity of this phasor is then proportional to

$$|\varphi|^2 \propto \frac{2}{z^2}\left(1+\cos\left[\frac{k}{z}\left(\vec{s}_1^{\perp}-\vec{s}_2^{\perp}\right)\cdot\vec{r}^{\perp}\right]\right). \tag{S5.5}$$

The interference term above oscillates with an angular frequency of $k\frac{l_s}{z}$ (as a function of screen position $\vec{r}^{\perp}$), where $l_s = |\vec{s}_1^{\perp}-\vec{s}_2^{\perp}|$ is the distance between the points on the slits.

We now make contact with sampling theory: By the Shannon sampling theorem, we may resolve that interference term without aliasing, only if we observe the phasor over a spatial extent (i.e., a pixel) no greater than $\frac{1}{2k(l_s/z)}$. Under the setting of our small-angle approximation, l_s/z approximates the angle between $\vec{s}_1, \vec{s}_2$ subtended from $\vec{r}$, denoted θ_s. Denote the spatial extent of a pixel as l_r,and we establish the *sampling relation*: $l_r(k\theta_s) \leq 1/2$. This relation resembles the uncertainty relation [Mandel and Wolf 1995, Chapter 4] but with the inequality reversed (l_r is the spatial variance and $k\theta_s$ is the angular variance scaled by the wavenumber).

When the angular extent θ_s between the interference sources $\vec{s}_{1,2}$ is large enough to violate the sampling relation above, i.e. $l_r(k\theta_s) > 1/2$, the observer is no longer able to resolve the interference term. Generalized rays at the detector ($\rho = 0$) are minimum-uncertainty constructs, i.e. fulfil the equality in the sampling relation. Therefore, generalized rays quantify exactly the spatial and angular extent

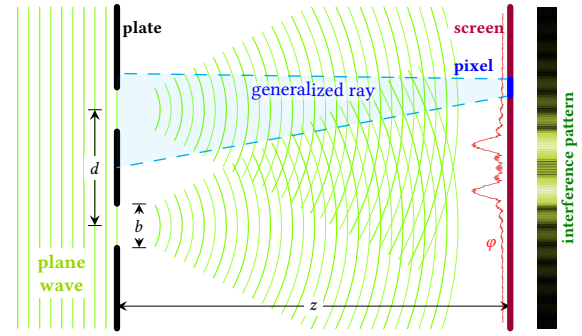

(a) experiment (top view)

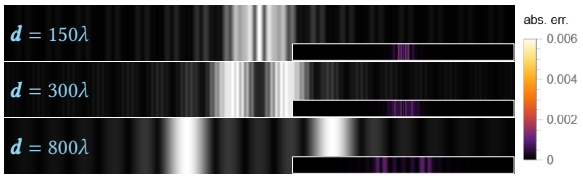

(b) interference patterns for various d

Fig. S5. **Diffraction through double slits.** (a) Schematic of Young's double slit experiment. A pair of slits, of width b and spaced a distance d apart, are cut in a thin, conductive plate. A coherent plane wave (illustrated in green) diffracts through the slits, and is observed upon a screen, placed at a distance z from the plate. The superposition of coherent light from both slits results in a rapidly-oscillating phasor φ (illustrated in red), producing an interference pattern. (b) The experiment is performed with increasing slit distances d, and we compare our method (sampling incident light with generalized rays) with a ground truth (explicitly diffracting the plane wave through the slits). Differences are plotted in the insets at the bottom right of each pattern; also see Fig. S6. The experiment was performed with wavelength $\lambda = 1$ (arbitrary units), $z = 10000\lambda$ and $b = 40\lambda$.

over which interference effects may be resolved. The arguments above are not limited to double-slit diffraction, and apply, in general, to any superposition from two or more point sources. To conclude: it is the *integration over the observer's spatial extent that induces decoherence*, allowing generalized rays to regain linearity under any illumination conditions.

We illustrate that decoherence in Fig. S5b, where we perform the experiment with increasing slit distances. The standard deviation of the spatial extent of a generalized ray (illustrated with dashed cyan lines in Fig. S5a) at the slits is about 100λ, i.e. a full width at half maximum of $\sim 235\lambda$. While the characteristic double-slit interference pattern is visible at first (when d is smaller than the spatial extent of a generalized ray), it slowly dissolves into a pair of (mutually-incoherent) single-slit diffraction patterns. Observe the increasing frequency of the secondary fringes in Fig. S5b with increasing d.

S5.2 Numeric Validation

We perform the described double-slit experiment using a pair of methods: (i) forward transport, where we diffract the incident plane wave directly via the exact Rayleigh-Sommerfeld (RS) explicit diffraction integral of the first kind [Mandel and Wolf 1995, Chapter 3.2],

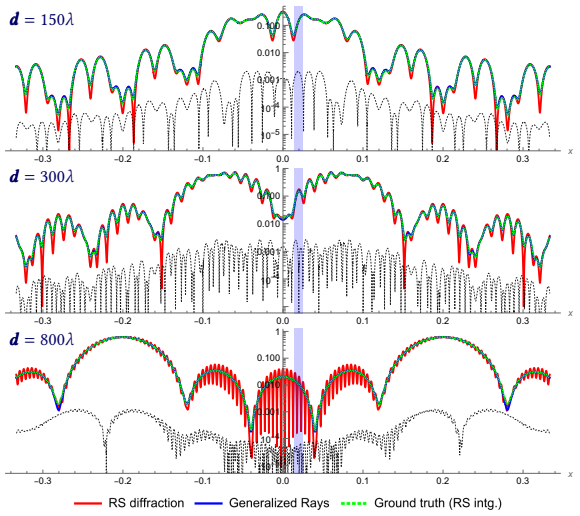

Fig. S6. Plots of light intensity as a function of x (position on screen) of the experiment in Fig. S5. In red we plot the exact Rayleigh-Sommerfeld (RS) diffraction, i.e. the *unobservable* interference that arises in singular points. The spatial extent of a detector on the screen (a pixel in each pattern in Fig. S5b) is illustrated as a cyan bar. Integrating the RS diffraction over that spatial extent of a pixel computes the numeric ground truth, plotted in dashed green. Results obtained with generalized rays are plotted in blue; absolute error with respect to the ground truth is plotted in dotted black.

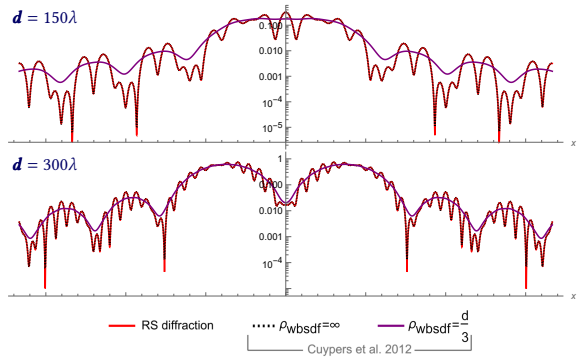

Fig. S7. **Loss of locality with Cuypers et al. [2012].** Red plot (exact RS diffraction) is as in Fig. S6. Let ρ_{wbsdf} be the radius of integration of a diffraction kernel (Eq. (2)) in Cuypers et al. [2012]. Limiting that integration radius produces erroneous results, hence *locality is entirely lost* with their method: correctness requires integration over the entire scene ($\rho_{\text{wbsdf}} = \infty$).

and integrate over the detector (each pixel); and (ii) backward transport using our formalism, i.e. sampling with generalized rays and diffracting them through the slits over their spatial extent. The absolute difference between the two methods is shown in the inset at the bottom right of each pattern in Fig. S5b. The differences (due to ignoring generalized rays' tail) are negligible.

Once the distance d becomes larger than the spatial extent of a generalized ray that reaches the slits, generalized rays no longer solve a double-slit diffraction problem. Indeed, they don't need to: interference between the two slits does arise at singular points, but is not observable over a pixel's extent. This can be seen in Fig. S6,

where for large d we can see that the rapidly-oscillating double-slit interference pattern arises everywhere, but is integrated out over the spatial extent of a pixel. In contrast to generalized rays, explicit diffraction integrals (first method above) have no means of quantifying the extent over which observable interference may arise. This means that the explicit method needs to (wastefully) integrate each sample over each pixel and over the entire plate. Therefore, our method (second method) is an order-of-magnitude faster than the explicit method (8.9 s compared to 246 s). This experiment demonstrates that generalized rays achieve well-defined weak locality.

Other methods suffer from similar deficiencies: they fail to derive machinery that is able to quantify the extent over which weak-locality can be maintained. For example, Cuypers et al. [2012] describe a WDF-based formalism with perfectly local primitives. However, as discussed in Section 2, such a formalism forgoes either locality or linearity: formulating a light-matter interaction requires integration over the entire scene, as defined by the interaction kernel (Eq. (2)). In Fig. S7 we show that if we were to limit the spatial extent of integration of an interaction kernel, while insisting on linearity (ignore the bilinear superposition term in Eq. (3)), incorrect results are produced. Indeed, no wave-optical formalism may be simultaneously perfectly local and linear.

S6 Interaction Kernels

Developing a light-matter interaction operator entails deriving an analytic method to compute the diffracted distribution $\mathcal{Q}_o$ which arises from the interaction operator acting on an arbitrary generalized ray, and express it as a finite sum of generalized rays, viz. Eq. (22). Formally, this can always be done with arbitrarily high accuracy (as discussed in Section 3). However, doing so in practice can be analytically involved, and at times application-specific assumptions are made in order to simplify the analysis. Often, an importance sampling strategy, that is used to sample a generalized ray out of the sum in Eq. (22), is desired. In this Section we present a few important example interactions.

Because time-reversal is equivalent to wavevector reversal and phase conjugation [Geru 2018, Chapter 2.3], we may rewrite the action of an arbitrary time-reversed interaction operator on a generalized ray (that enters the rendering equation Eq. (23)) as

$$\mathcal{K}^{-\mathrm{t}}\left\{g_{\beta,\rho}\left(\vec{\boldsymbol{r}}',\vec{\boldsymbol{k}}'\,;\vec{\boldsymbol{r}}_0,\vec{\boldsymbol{k}}_0\right)\right\}=\mathcal{K}\left\{g_{\beta,\rho}\left(\vec{\boldsymbol{r}}',\vec{\boldsymbol{k}}'\,;\vec{\boldsymbol{r}}_0,-\vec{\boldsymbol{k}}_0\right)\right\} \tag{S6.1}$$

(recall that the WDF is real). Instead of deriving time-reversed interactions, we only need to reverse the generalized ray's wavevector.

We classify light-matter interactions into two types:

(1) Simple — These arise with *linear optical systems* (also known as "ABCD systems"): propagation through a homogeneous medium with a slowly-varying refractive index (including free space), and reflection or refraction at a smooth interface. Under simple interactions the dynamics are identical to ray-optical dynamics (see our supplemental material).
(2) Diffractive — All other interactions, e.g., scattering by a rough surface: there are many ways to formalise a diffractive interaction, and we will discuss a few examples.

Generalized rays are unique in simultaneously (i) behaving like classical rays under simple interactions; (ii) and being *Wigner-representable* [Torre 2005], i.e. they admit well-defined wave functions. The first point only holds for constructs that are fully defined by their 2nd-order moments matrix (like a generalized ray, which is a Gaussian beam), while the second point only applies to phase-space constructs that fulfil the uncertainty relation. For example, the second point does not hold for perfectly-local Wigner-based formalisms (the WDF $\mathscr{W}\equiv\delta(\vec{\boldsymbol{r}})\,\delta(\vec{\boldsymbol{k}})$ admits no wave function).

Free-space propagation (simple interaction). Let a generalized ray be parameterized by its mean spatial position $\vec{\boldsymbol{r}}_0$, mean wavevector $\vec{\boldsymbol{k}}_0$, and β, ρ, as before. Using its phase-space representation (Eq. (5)), observe that we may express its spatial and wavevector variances as $\sigma_r^2=\beta^2/2$, and $\sigma_k^2=(1+\rho^2)/(2\beta^2)$, respectively. σ_r quantifies the spatial extent occupied by the generalized ray, and σ_k/k_0 the solid angle into which the generalized ray propagates.

Let $z>0$ be the distance of propagation, and $\bar{z}=z/|\vec{\boldsymbol{k}}_0|$ the distance normalized by the mean wavevector. Then, the corresponding kernel (Eq. (2)) is

$$K_{\substack{\text{FREE}\\\text{SPACE}}}\left(\vec{\boldsymbol{r}},\vec{\boldsymbol{r}}',\vec{\boldsymbol{k}},\vec{\boldsymbol{k}}'\right)\triangleq\delta\left(\vec{\boldsymbol{k}}'-\vec{\boldsymbol{k}}\right)\delta\left(\vec{\boldsymbol{r}}'+\bar{z}\vec{\boldsymbol{k}}-\vec{\boldsymbol{r}}\right), \tag{S6.2}$$

i.e. propagation in direction $\vec{\boldsymbol{k}}$. Applying Eq. (11), the diffracted (propagated) distribution after interaction becomes

$$\mathcal{Q}_o\left(\vec{\boldsymbol{r}},\vec{\boldsymbol{k}}\right)=\mathcal{K}_{\substack{\text{FREE}\\\text{SPACE}}}\left\{g_{\beta,\rho}\left(\vec{\boldsymbol{r}}',\vec{\boldsymbol{k}}'\,;\vec{\boldsymbol{r}}_0,\vec{\boldsymbol{k}}_0\right)\right\}=g_{\beta_o,\rho_o}\left(\vec{\boldsymbol{r}},\vec{\boldsymbol{k}}\,;\vec{\boldsymbol{r}}_0+\bar{z}\vec{\boldsymbol{k}}_0,\vec{\boldsymbol{k}}_0\right) \tag{S6.3}$$

$$\text{with}\qquad\beta_o^2=\beta^2+\bar{z}\left(2\rho+2\bar{z}\tfrac{1+\rho^2}{2\beta^2}\right)\quad\text{and}\quad\rho_o^2=\left(\rho+2\bar{z}\tfrac{1+\rho^2}{2\beta^2}\right)^2.$$

The reader may verify that after propagation $\vec{\boldsymbol{k}}_0,\sigma_k^2$ remain invariant, i.e. the direction and solid angle into which the generalized ray propagates do not change; and, the spatial variance transforms as $\sigma_r^2\to\sigma_r^2+O(\bar{z}^2\sigma_k^2)$, i.e. the space occupied by the generalized ray increases proportionally to the propagation distance times the solid angle. Finally, the mean position $\vec{\boldsymbol{r}}_0$ is shifted by z. That is, all the generalized ray's parameters transform under classical ray-optical dynamics, as would be expected with free-space propagation.

Reflection/refraction at an interface (simple interaction). Assume a smooth, flat interface. Let $\mathcal{K}$ be the interaction operator, quantifying the action of reflection or refraction at that interface. Let $\vec{\boldsymbol{k}}$ be the mean wavevector of the incident generalized ray, and we denote $\vec{\boldsymbol{k}}_0^{(\mathrm{r})}$ as that wavevector after the reflection (by the law of reflection) or refraction (by Snell's law) at the interface. Let a be the Fresnel coefficient for the interaction. Then,

$$\mathcal{Q}_o\left(\vec{\boldsymbol{r}},\vec{\boldsymbol{k}}\right)=\mathcal{K}_{\substack{\text{REFLECT}\\\text{REFRACT}}}\left\{g_{\beta,\rho}\left(\vec{\boldsymbol{r}}',\vec{\boldsymbol{k}}'\,;\vec{\boldsymbol{r}}_0,\vec{\boldsymbol{k}}_0\right)\right\}=a\,g_{\beta,\rho}\left(\vec{\boldsymbol{r}},\vec{\boldsymbol{k}}\,;\vec{\boldsymbol{r}}_0,\vec{\boldsymbol{k}}_0^{(\mathrm{r})}\right). \tag{S6.4}$$

The parameters $\vec{\boldsymbol{r}}_0,\beta,\rho$ are unchanged as no propagation takes place. That is, with generalized rays this interaction (and other simple interactions) mimics the classical dynamics. Reflection and refraction are polarization-dependent phenomena. Vectorization can be done trivially (generalized ray per each polarization component). We handle polarization differently, see Section 5.1.

Diffraction grating (diffractive interaction). A benefit of generalized rays is that they admit a well-defined wave function, $\psi_{\beta,\rho}$. Hence, to formulate a diffractive interaction one may work in phase

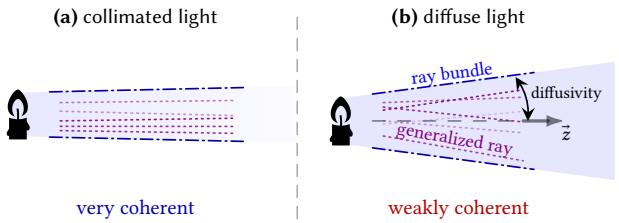

Fig. S8. **Relation to optical coherence.** A *ray bundle* is a collection of generalized rays. The optical coherence of that bundle is inversely proportional to the angular spread of generalized rays in the bundle. That is, (a) collimated light, where the rays propagated roughly in the same direction, is highly coherent; on the other hand, (b) a ray bundle, in which the rays have a large spread of propagation direction, is weakly coherent. The mean direction of propagation $\vec{z}$ of the bundle, as well as its *diffusivity*—variance of angular spread of generalized rays—are illustrated.

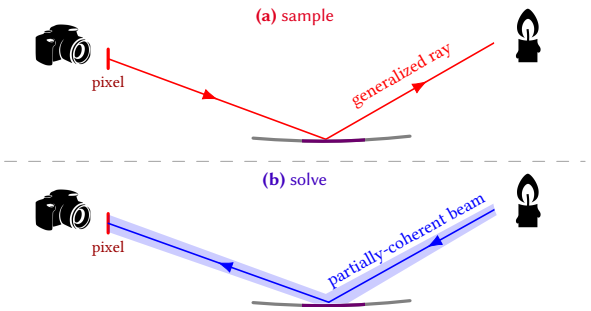

Fig. S9. **Sample-solve.** Our path tracing algorithm (a) uses generalized rays (dotted lines) to *sample* paths through the scene. Generalized rays are always linear, therefore classical sampling techniques apply essentially unchanged. Once a path is sampled (solid red path), we (b) *solve* for the partially-coherent light transport, by applying PLT [Steinberg et al. 2022] from the light source to the sensor across all intermediate interactions.

space with WDFs, or in position space with wave functions (via a diffraction integral, electromagnetic theory, etc.). Furthermore, the generalized ray's wave function is a simple (coherent) Gaussian beam, which has been extensively studied in optical literature.

Apply the Fraunhofer diffraction integral [Born and Wolf 1999] to a one-dimensional sinusoidal grating of period Λ and height h, yielding

$$\mathcal{Q}_o\left(\vec{\boldsymbol{r}},\vec{\boldsymbol{k}}\right) = \mathcal{K}_{\text{GRTN}}\left\{g_{\beta,\rho}\left(\vec{\boldsymbol{r}}',\vec{\boldsymbol{k}}'\,;\vec{\boldsymbol{r}}_0,\vec{\boldsymbol{k}}_0\right)\right\} \approx \sum\nolimits_j a_j g_{\beta,\rho}\left(\vec{\boldsymbol{r}},\vec{\boldsymbol{k}}\,;\vec{\boldsymbol{r}}_0,\vec{\boldsymbol{k}}_0^{(j)}\right), \tag{S6.5}$$

where $a_j = \mathrm{J}_j(hk/2)^2$ is the intensity of the j-order diffracted lobe, J_j is the Bessel function of 1$^{\text{st}}$ kind and $\vec{\boldsymbol{k}}_0^{(j)}$ is the diffracted wavevector in direction $\sin\theta_o = \sin\theta_i - j\frac{\lambda}{\Lambda}$, where $\theta_{i,o}$ are the incident and diffraction directions w.r.t. the grating direction and λ is the wavelength. The diffraction grating also very slightly enlarges the correlation constant ρ, but we ignore that effect for simplicity.

Scatter by moderately-rough surface (diffractive interaction). We make a simplifying assumption: ignore the beam curvature of a generalized ray, viz. $\psi_{\beta,\rho} \equiv \mathrm{e}^{\mathrm{i}\vec{\boldsymbol{k}}_0\cdot(\vec{\boldsymbol{r}}-\vec{\boldsymbol{r}}_0)}\mathrm{e}^{-\frac{1}{2\beta^2}|\vec{\boldsymbol{r}}-\vec{\boldsymbol{r}}_0|^2}$. This serves to transform the Gaussian beam into a spatially-modulated plane wave, easing the analysis. We then may apply the Harvey-Shack surface scatter theory [Krywonos 2006]:

$$\begin{aligned}\mathcal{Q}_o\left(\vec{\boldsymbol{r}},\vec{\boldsymbol{k}}\right) &= \mathcal{K}_{\text{SURFACE}}\left\{g_{\beta,\rho}\left(\vec{\boldsymbol{r}}',\vec{\boldsymbol{k}}'\,;\vec{\boldsymbol{r}}_0,\vec{\boldsymbol{k}}_0\right)\right\}\\ &\approx \int \mathrm{d}\vec{\boldsymbol{k}}_0^{(\mathrm{r})}\, f_{\text{HS}}\left(\vec{\boldsymbol{k}}_0 \to \vec{\boldsymbol{k}}_0^{(\mathrm{r})}\right) g_{\beta,\rho}\left(\vec{\boldsymbol{r}},\vec{\boldsymbol{k}}\,;\vec{\boldsymbol{r}}_0,\vec{\boldsymbol{k}}_0^{(\mathrm{r})}\right),\end{aligned} \tag{S6.6}$$

where f_{HS} is the Harvey-Shack BRDF, quantifying the scattering for incident and diffracted wavevectors $\vec{\boldsymbol{k}}_0$ and $\vec{\boldsymbol{k}}_0^{(\mathrm{r})}$, respectively. In practice, sampling generalized rays from the distribution $\mathcal{Q}_o$ is done via Monte-Carlo integrating the expression above, thereby rewriting it in the form of Eq. (22):

$$\mathcal{Q}_o\left(\vec{\boldsymbol{r}},\vec{\boldsymbol{k}}\right) \approx \frac{1}{J}\sum_j^J \frac{f_{\text{HS}}\left(\vec{\boldsymbol{k}}_0 \to \bar{\boldsymbol{k}}_{0,j}^{(\mathrm{r})}\right)}{p_j} g_{\beta,\rho}\left(\vec{\boldsymbol{r}},\vec{\boldsymbol{k}}\,;\vec{\boldsymbol{r}}_0,\bar{\boldsymbol{k}}_{0,j}^{(\mathrm{r})}\right) \tag{S6.7}$$

(up to a normalization constant), where $\bar{\boldsymbol{k}}_{0,j}^{(\mathrm{r})}$ are the set of J sampled mean scattering directions, and p_j are the sampling probabilities. In our implementation, the Harvey-Shack BRDF f_{HS} is importance sampled using the technique described by Holzschuch and Pacanowski [2017].

In Eqs. (S6.5) and (S6.6) we make optical approximations that are reasonable for our target application: rendering at optical frequencies. One consequence of using the (approximative) Harvey-Shack model for surfaces is that we only model the averaged scatter, where every generalized ray interacts with the entire distribution of surface frequencies. Therefore, surface imperfections, such as glints, or optical speckle, do not arise. We stress that the assumptions we made here are not a limitation of our light transport formalism. Also note that, in general, $\mathcal{K}$ may describe cross-wavelength scattering, e.g., due to fluorescence or phosphorescence, however we ignore such effects in our implementation.

Other materials that we use in our rendered scenes are also derived under the assumption that we ignore the curvature of a generalized ray. Their interaction formula then echoes Eq. (S6.6), but with the BRDF f_{HS} replaced with the BRDF of the relevant interaction that acts on plane waves.

S7 Sample-Solve

In this Subsection we draw a formal connection between our backward (sensor-to-source) formalism to *optical coherence*. The purpose of this is to connect our formalism to other forward-based models and computational optics tools, thereby enabling the application of bi-directional techniques. For our application of interest, we will show that we may use a forward PLT pass as a cheap variance reduction technique.

Relation to optical coherence. The central quantity in the study of optical coherence [Wolf 2007] is the *cross-spectral density* (CSD) of a statistical ensemble of light waves, viz.

$$\mathcal{C}\left(\vec{\boldsymbol{r}}_1,\vec{\boldsymbol{r}}_2\right) \triangleq \left\langle \psi(\vec{\boldsymbol{r}}_1)\psi^\star(\vec{\boldsymbol{r}}_2)\right\rangle, \tag{S7.1}$$

where $\langle\cdot\rangle$ denotes ensemble-averaging over the wave ensemble, and the wave function ψ is now understood as a realization from that

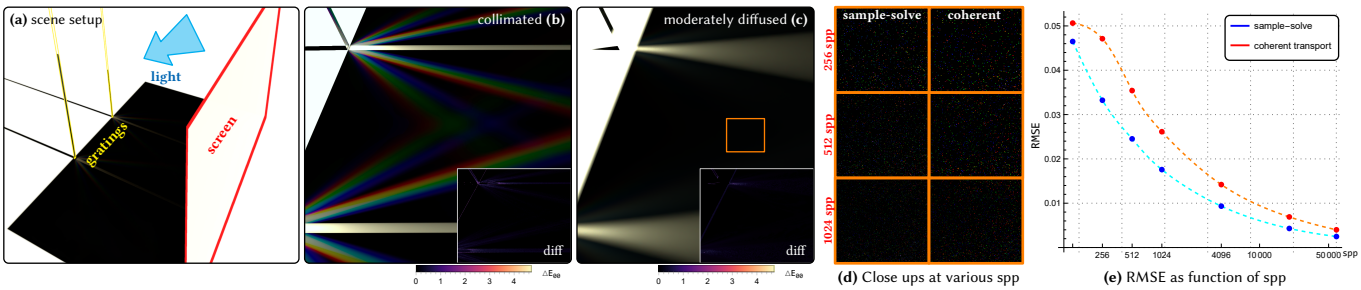

Fig. S10. **Sample-solve: PLT as a variance reduction technique.** During the solve step, we apply PLT to solve for the partially-coherent light transport over a sampled light path. (a) Light arrives from the right, illuminating a simple scene. A screen (on the right, red outline) shadows a rectangular area. On the left a pair of thin diffraction gratings (yellow outline) reflect and disperse light. (b) When the incident light is highly collimated (subtends a very small solid angle from the scene), hence is moderately coherent, the diffraction lobes are clearly visible. (c) As we increase the light's diffusivity (increase its solid angle), the reflection from the direct contribution lobe spreads out, while diffraction lobes mostly disappear, as expected. As the only light that arrives to the shadowed region is from the diffraction gratings, this is a good scene to study the benefits of the solve stage, and to do so we render the scene with only fully-coherent transport (no PLT applied). (b-c, insets) Difference images between the partially-coherent and fully-coherent sample-solve show that indeed both converge to an identical result. (d) However, close ups on the region outlined in orange in (c) show that, while the diffraction lobes are no longer visible, they still induce error, which the partially-coherent solve stage serves to reduce. (e) Plot of error in that area as function of sample count suggests that fully-coherent transport requires about 2-5 times the sample count to achieve similar-quality renderings.

ensemble. From the definitions of the WDF and the CSD, Eqs. (2) and (S7.1), it is easy to observe that the CSD is the Fourier transform of the ensemble-averaged WDF, therefore we may write

$$\mathcal{C}\left(\vec{\boldsymbol{r}}-\tfrac{1}{2}\vec{\boldsymbol{d}},\vec{\boldsymbol{r}}+\tfrac{1}{2}\vec{\boldsymbol{d}}\right)\propto\left\langle\mathscr{F}\left\{\mathscr{W}(\vec{\boldsymbol{r}},\vec{\boldsymbol{k}}')\right\}\left(\vec{\boldsymbol{d}}\right)\right\rangle \tag{S7.2}$$

(up to an irrelevant constant), where the Fourier transform $\mathscr{F}$ is w.r.t. the primed variable, $\vec{\boldsymbol{k}}'$. We may immediately conclude that optical coherence at a fixed point $\vec{\boldsymbol{r}}$ only depends on the (ensemble-averaged) light's wavevector distribution in phase space.

Under partial coherence, the observed values are often ensemble-averaged values [Wolf 2007] (and can be understood as time averaging over the period of detection), i.e. $\langle I\rangle$. Then, partial coherence is trivially accounted for by replacing the sourcing WDF $\mathscr{W}_s$ with its ensemble-averaged counterpart, $\langle\mathscr{W}_s\rangle$, in all our formulae. $\langle\mathscr{W}_s\rangle$ is the Fourier transform of the sourced CSD of light (Eq. (S7.2)).

We term the *diffusivity* $\boldsymbol{\Omega}$ of a WDF to be the angular variance of propagation from the mean direction of propagation. The diffusivity $\boldsymbol{\Omega}$ is a 2×2 positive-definite matrix $\boldsymbol{\Omega}$, allowing for anisotropy. For example, in the isotropic case, the diffusivity can be written as σ_k^2/k_0^2, where σ_k^2 is the distribution's wavevector variance. The variance in the solid angle into which the bundle propagates is then $\Omega=|\boldsymbol{\Omega}|$. Then, from Eq. (S7.2) we may formally derive the following result (see Section S4.2 in our supplemental material):

$$\boldsymbol{\Theta}=\lambda^2\boldsymbol{\Omega}^{-1}, \tag{S7.3}$$

where $\boldsymbol{\Theta}$ is the *shape matrix* from PLT theory [Steinberg et al. 2022], i.e. the (inverse) 2^{nd}-order moments of spatial variance. Deriving equality relations between higher-order moments is also possible, however we do not require higher-order moments.

A consequence of the above relation between optical coherence and the diffusivity of light is the well-known connection between the coherence area of light, i.e. $|\boldsymbol{\Theta}|$, and the solid angle subtended by a thermal source [Mandel and Wolf 1995]: $|\boldsymbol{\Theta}|=\frac{\lambda^2}{\Omega}$. Though note that the relation Eq. (S7.3) is more general, and establishes a direct connection between optical coherence and light's distribution in the wave-optical phase space, with no assumptions on the light sourcing process or its state-of-coherence at other regions in space.

Sample-solve. We present a simple two pass algorithm: first, we *sample* the wave-optical distribution of light using generalized rays, using the rendering algorithm we presented in the paper, Algorithm 1. The process continues until a generalized ray encounters a light source, thereby a path $\boldsymbol{\pi}$ that connects the detector to a light source is found. The effective diffusivity over that path is well defined. In-place of the measurement stage (line 9 in Algorithm 1), we now apply a forward pass: we use PLT machinery [Steinberg et al. 2022] to *solve* for the partially-coherent light transport over the path $\boldsymbol{\pi}$. More formally, instead of integrating over the ensemble-averaged sourcing WDF, $\langle\mathscr{W}_s\rangle$, we use PLT to source a partially-coherent beam that corresponds to $\langle\mathscr{W}_s\rangle$ from that light source, and retrace the steps forward (source-to-sensor) taken by the generalized ray over the path. As the PLT beam captures more information (it is "wider" than a generalized ray), this forward solve pass serves as a *variance-reduction technique*, by computing the partially-coherent optical response for the sampled path. See Fig. S10.

In our domain of interest—wave-optical rendering—applying PLT makes sense: The optical coherence of light is a primary factor in limiting our ability to resolve wave-interference effects. Other applications might find a different *solve* pass to be more appropriate: For example, integrating optical speckle statistics can be done via a sample-solve approach, where first we *sample* paths connecting the detector to a light source, and then integrate statistics in a *solve* pass. Such applications are beyond the scope of this paper, however they serve to highlight the generality of this simple sample-solve approach: it bridges a gap between classical path tracing tools and wave optics, via the generalized ray construct, and enables the application of powerful sampling techniques in a wider context.

References

Michael Bass, Casimer DeCusatis, Jay M Enoch, Vasudevan Lakshminarayanan, Guifang Li, Carolyn A MacDonald, Virendra N Mahajan, and Eric W Van Stryland. 2009. *Handbook of optics: Geometrical and physical optics, polarized light, components and instruments(set) v. 1* (3 ed.). McGraw-Hill Professional, New York, NY.

M J Bastiaans. 1978. The Wigner distribution function applied to optical signals and systems. *Opt. Commun.* 25, 1 (April 1978), 26–30.

Max Born and Emil Wolf. 1999. *Principles of optics : electromagnetic theory of propagation, interference and diffraction of light.* Cambridge University Press, Cambridge New York.

Hans Adolph Buchdahl. 1993. *An introduction to Hamiltonian optics.* Courier Corporation.

Leon Cohen. 1966. Generalized phase-space distribution functions. *J. Math. Phys.* 7, 5 (May 1966), 781–786.

Leon Cohen. 1994. *Time frequency analysis.* Prentice Hall, Philadelphia, PA.

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

Ion I. Geru. 2018. *Time-Reversal Symmetry.* Springer International Publishing. https://doi.org/10.1007/978-3-030-01210-6

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

Andrey Krywonos. 2006. *Predicting surface scatter using a linear systems formulation of non-paraxial scalar diffraction.* Ph. D. Dissertation. University of Central Florida.

Ulf Leonhardt. 1997. *Measuring the quantum state of light.* Cambridge University Press, Cambridge, England.

Leonard Mandel and Emil Wolf. 1995. *Optical coherence and quantum optics.* Cambridge University Press, Cambridge.

Shlomi Steinberg, Ravi Ramamoorthi, Benedikt Bitterli, Arshiya Mollazainali, Eugene d'Eon, Ling-Qi Yan, and Matt Pharr. 2024. A Free-Space Diffraction BSDF. (2024).

Shlomi Steinberg, Pradeep Sen, and Ling-Qi Yan. 2022. Towards Practical Physical-Optics Rendering. *ACM Transactions on Graphics* 41, 4 (Jul 2022), 1–13. https://doi.org/10.1145/3528223.3530119

Markus Testorf, Bryan Hennelly, and Jorge Ojeda-Castañeda. 2010. *Phase-space optics: fundamentals and applications.* McGraw-Hill Education.

Amalia Torre. 2005. *Linear ray and wave optics in phase space: bridging ray and wave optics via the Wigner phase-space picture.* Elsevier.

A Walther. 1968. Radiometry and coherence. *JOSA* 58, 9 (1968), 1256–1259. https://doi.org/10.1364/JOSA.58.001256

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

E. Wigner. 1932. On the Quantum Correction For Thermodynamic Equilibrium. *Phys. Rev.* 40 (Jun 1932), 749–759. Issue 5. https://doi.org/10.1103/PhysRev.40.749

Emil Wolf. 2007. *Introduction to the theory of coherence and polarization of light.* Cambridge University Press, Cambridge.

Zhengyun Zhang and Marc Levoy. 2009. Wigner distributions and how they relate to the light field. In *2009 IEEE International Conference on Computational Photography (ICCP).* 1–10. https://doi.org/10.1109/ICCPHOT.2009.5559007