

Rendering: From Ray Optics to Wave Optics

SHLOMI STEINBERG, RAVI RAMAMOORTHY, and 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

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 point-wise 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{q}_1 to point \vec{q}_2 fulfils

$$\delta \int_{\vec{q}_1}^{\vec{q}_2} ds \eta(\vec{q}') = 0, \quad (\text{S2.1})$$

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:

$$\frac{d}{ds} \left[\eta(\vec{q}) \frac{d\vec{q}}{ds} \right] = \nabla \eta(\vec{q}). \quad (\text{S2.2})$$

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

$$\vec{p} \triangleq \eta(\vec{q}) \frac{d\vec{q}}{ds}. \quad (\text{S2.3})$$

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

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

as a *ray*. The ray \vec{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*

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

Authors’ addresses: Shlomi Steinberg, p@shlomisteinberg.com; Ravi Ramamoorthy, ravir@cs.ucsd.edu; 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, 93106, United States; Matt Pharr, matt@pharr.org, NVIDIA, San Francisco, United States.

and Hamiltonian's equations

$$\frac{dq_\beta}{ds} = \frac{\partial H}{\partial p_\beta} \quad \text{and} \quad \frac{dp_\beta}{ds} = -\frac{\partial H}{\partial q_\beta}, \quad (\text{S2.6})$$

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

The above can be recast in operator notation into

$$\frac{d}{ds} \vec{u} = \mathcal{H} \vec{u}, \quad (\text{S2.7})$$

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

Eq. (S2.7) is the *ray equation*, an operator-valued differential first-order equation, and \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* \mathcal{T} :

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

The exponential map above maps the Lie algebra $\{\mathcal{H}\}$ 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{q}, \vec{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:

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

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

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

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

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

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

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 $\rho(\vec{q}, \vec{p}; z)$ 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 $f(\vec{q}, \vec{p})$, the average value of f over the entire statistical ensemble of rays is

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

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

$$\frac{\partial}{\partial z} \rho = -\mathcal{H} \rho \quad (\text{S2.14})$$

$$\frac{d}{dz} \rho = 0, \quad (\text{S2.15})$$

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*. Liouville theorem (Eq. (S2.15)) implies that the ray density in phase-space is a conserved quantity (ignoring absorption), both locally and globally:

- That "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{q}, \vec{p}) , the density $\rho(\vec{q}, \vec{p}; z)$ 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

$$L = \eta^2 \frac{\partial \Phi}{\partial G}, \quad (\text{S2.16})$$

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 $\hbar \rightarrow 0$). 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{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

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

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

$$q \mapsto \hat{q} \triangleq q \quad \text{and} \quad p \mapsto \hat{p} \triangleq -i\hbar \frac{\partial}{\partial q}. \quad (\text{S3.2})$$

The Hamiltonian operator then becomes

$$\hat{\mathcal{H}}(\hat{q}, \hat{p}) = -\frac{\hbar^2}{2m} \nabla_{\vec{q}}^2 + V(\vec{q}), \quad (\text{S3.3})$$

with $\nabla_{\vec{q}}^2$ being the Laplacian w.r.t. the spatial variable \vec{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),

$$\hat{\mathcal{H}}(\hat{q}, \hat{p}) = -\frac{\hbar^2}{2\eta(\hat{q})} \nabla_{\vec{q}}^2 - \eta(\vec{q}) \quad (\text{S3.4})$$

becomes the quadratic wave-optical Hamiltonian operator.

We denote $\psi(\vec{q}; t)$ as the *wave function*, with \vec{q} being spatial position and t time. The wave function is a complex function that serves as a descriptor of light under 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

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

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

$$i\hbar \frac{\partial}{\partial t} \hat{\mathcal{U}} = \hat{\mathcal{H}} \hat{\mathcal{U}}, \quad (\text{S3.6})$$

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{p} as $e^{i\vec{q} \cdot \vec{k}}$, we may write

$$\psi(\vec{q}) = \frac{1}{(2\pi)^{\frac{3}{2}}} \int d^3\vec{k} \tilde{\psi}(\vec{k}) e^{i\vec{q} \cdot \vec{k}}, \quad (\text{S3.7})$$

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

$$\hat{k} \triangleq \frac{1}{\hbar} \hat{p}, \quad (\text{S3.8})$$

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

$$\sigma_{qx}^2 \triangleq \int d^3\vec{q} q_x^2 |\psi(\vec{q})|^2, \quad (\text{S3.9})$$

$$\sigma_{kx}^2 \triangleq \int d^3\vec{k} k_x^2 |\tilde{\psi}(\vec{k})|^2. \quad (\text{S3.10})$$

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*:

$$\sigma_{qx} \sigma_{kx} \geq \frac{1}{2}. \quad (\text{S3.11})$$

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

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

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

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

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{q}, \hat{p} , it can be shown that [Torre 2005]

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

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

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix}_k = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} A & B \\ C & D \end{pmatrix}_q \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}. \quad (\text{S3.15})$$

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 $\vec{u} = (\vec{q}, \vec{p})$ is replaced by the optical wave function and its conjugate $(\psi, \tilde{\psi})$ 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

$$\begin{aligned} \mathcal{W}(\vec{q}, \vec{k}) &\triangleq \frac{1}{(2\pi)^3} \int d^3\vec{q}' \psi\left(\vec{q} - \frac{\vec{q}'}{2}\right) \psi^*\left(\vec{q} + \frac{\vec{q}'}{2}\right) e^{-i\vec{q}' \cdot \vec{k}} \\ &\triangleq \frac{1}{(2\pi)^3} \int d^3\vec{k}' \tilde{\psi}\left(\vec{k} - \frac{\vec{k}'}{2}\right) \tilde{\psi}^*\left(\vec{k} + \frac{\vec{k}'}{2}\right) e^{i\vec{q} \cdot \vec{k}'}, \end{aligned} \quad (\text{S3.16})$$

with both definitions equivalent. The WDF belongs to the wider *Cohen's 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:

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

where 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:

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

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

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

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 $\hat{f}(\hat{q}, \hat{k})$, its expectation value is

$$\langle \hat{f} \rangle_\psi = \langle \psi | \hat{f} | \psi \rangle = \int d^3\vec{q} \psi^*(\vec{q}) \hat{f}(\vec{q}, \hat{k}) \psi(\vec{q}). \quad (\text{S3.20})$$

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

$$\langle \hat{f} \rangle_\psi = \int d^3\vec{q} d^3\vec{k} f(\vec{q}, \vec{k}) \mathcal{W}(\vec{q}, \vec{k}), \quad (\text{S3.21})$$

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 — $\mathcal{W} \in \mathbb{R}$.

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

$$|\psi(\vec{q})|^2 = \int d^3\vec{k} \mathcal{W}(\vec{q}, \vec{k}) \quad (\text{S3.22})$$

$$|\tilde{\psi}(\vec{k})|^2 = \int d^3\vec{q} \mathcal{W}(\vec{q}, \vec{k}). \quad (\text{S3.23})$$

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

$$\int d^3\vec{q} d^3\vec{k} \mathcal{W}(\vec{q}, \vec{k}) = 1. \quad (\text{S3.24})$$

The converse holds as well. In general, the WDF can be normalized as $\int d^3\vec{q} d^3\vec{k} \mathcal{W} = 0$ if and only if $\psi \equiv 0$.

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

$$\psi'(\vec{q}) = \psi(\vec{q} + \vec{q}') \implies \mathcal{W}'(\vec{q}, \vec{k}) = \mathcal{W}(\vec{q} + \vec{q}', \vec{k}) \quad (\text{S3.25})$$

$$\tilde{\psi}'(\vec{k}) = \tilde{\psi}(\vec{k} + \vec{k}') \implies \mathcal{W}'(\vec{q}, \vec{k}) = \mathcal{W}(\vec{q}, \vec{k} + \vec{k}') . \quad (\text{S3.26})$$

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

$$\forall \vec{q} \notin S_q, \psi(\vec{q}) = 0 \quad \text{and} \quad \forall \vec{k} \notin S_k, \tilde{\psi}(\vec{k}) = 0 ,$$

the WDF vanishes outside these volumes as well:

$$\mathcal{W}(\vec{q}, \vec{k}) \neq 0 \quad \text{only if} \quad (\vec{q}, \vec{k}) \in S_q \times S_k . \quad (\text{S3.27})$$

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{q}, \hat{p}), the WDF obeys:

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

$$\frac{d}{dz} \mathcal{W}(\vec{q}, \vec{k}; z) = 0 , \quad (\text{S3.29})$$

i.e. the Liouville's equation and Liouville theorem of Hamiltonian mechanics, viz. Eqs. (S2.14) and (S2.15), and note that \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

$$\mathcal{W} = \mathcal{W}_1 + \mathcal{W}_2 + 2 \text{Re} \mathcal{W}_{12} , \quad (\text{S3.30})$$

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

$$\mathcal{W}_{12}(\vec{q}, \vec{k}) \triangleq \frac{1}{(2\pi)^3} \int d^3 \vec{q}' \psi_1\left(\vec{q} - \frac{\vec{q}'}{2}\right) \psi_2^*\left(\vec{q} + \frac{\vec{q}'}{2}\right) e^{-i\vec{q}' \cdot \vec{k}} . \quad (\text{S3.31})$$

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

$$E \triangleq \int d^3 \vec{q} |\psi(\vec{q})|^2 = \int d^3 \vec{q} d^3 \vec{k} \mathcal{W}(\vec{q}, \vec{k}) . \quad (\text{S3.32})$$

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

$$\left(\frac{\vec{q}}{\vec{k}}\right) \triangleq \frac{1}{E} \int d^3 \vec{q} d^3 \vec{k} \left(\frac{\vec{q}}{\vec{k}}\right) \mathcal{W}(\vec{q}, \vec{k}) . \quad (\text{S3.33})$$

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:

$$\mathbf{M} \triangleq \begin{pmatrix} m_{xx} & m_{xy} & m_{xz} & m_{x\tilde{x}} & m_{x\tilde{y}} & m_{x\tilde{z}} \\ m_{yx} & m_{yy} & m_{yz} & m_{y\tilde{x}} & m_{y\tilde{y}} & m_{y\tilde{z}} \\ & & \dots & \dots & & \\ m_{zx} & m_{zy} & m_{zz} & m_{z\tilde{x}} & m_{z\tilde{y}} & m_{z\tilde{z}} \end{pmatrix} \triangleq \frac{1}{E} \int d^3 \vec{q} d^3 \vec{k} \left[\begin{pmatrix} \vec{q} \\ \vec{k} \end{pmatrix} - \begin{pmatrix} \bar{q} \\ \bar{k} \end{pmatrix} \right] \left[\begin{pmatrix} \vec{q} \\ \vec{k} \end{pmatrix} - \begin{pmatrix} \bar{q} \\ \bar{k} \end{pmatrix} \right]^T \mathcal{W}(\vec{q}, \vec{k}) , \quad (\text{S3.34})$$

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 \mathbf{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. $\mathcal{W}(\vec{q}, \vec{k})$, then can be understood as a ‘‘ray’’, and we write $\mathcal{W}(\vec{u})$, with $\vec{u} = (\vec{q}, \vec{k})$ 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:

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

where the matrix \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{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:

$$\mathbf{M}(z) = \mathcal{T} \mathbf{M}(z_0) \mathcal{T}^T , \quad (\text{S3.36})$$

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] 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 \mathcal{W}_Φ , and let a wave function be composed of two spatially- and frequency-shifted copies of this signal:

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

where $\vec{q}_{1,2}$ and $\vec{k}_{1,2}$ 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:

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

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 (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*:

$$\begin{aligned} \mathcal{Q}(\vec{q}, \vec{k}) &\triangleq \frac{1}{(\sqrt{2\pi})^3} \int d^3\vec{q}' d^3\vec{k}' \mathcal{W}(\vec{q}', \vec{k}') e^{-\frac{1}{2}\vec{u}'^\top \Sigma^{-1} \vec{u}'}, \quad (\text{S4.3}) \\ \text{with } \vec{u}' &\triangleq \begin{pmatrix} \vec{q} - \vec{q}' \\ \vec{k} - \vec{k}' \end{pmatrix} \quad \text{and} \quad \sqrt{|\Sigma|} = \frac{1}{2^3}, \end{aligned}$$

where Σ 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. $\mathcal{W}_r = \delta^3(\vec{q} - \vec{q}_0) \delta^3(\vec{k} - \vec{k}_0)$, 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 \mathcal{Q}_r , that arises from \mathcal{W}_r via Eq. (S4.3), is physical. Let the covariance take the block-diagonal form $\Sigma = \operatorname{diag}\{\Sigma_q, \Sigma_k\}$, then:

$$\mathcal{Q}_r(\vec{q}, \vec{k}; t_0) = \frac{1}{(\sqrt{2\pi})^3} e^{-\frac{1}{2}\vec{q}'^\top \Sigma_q^{-1} \vec{q}' - \frac{1}{2}\vec{k}'^\top \Sigma_k^{-1} \vec{k}'}, \quad (\text{S4.4})$$

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 $\vec{u}(t_0) = (\vec{q}_0, \vec{k}_0)$, and the

moment matrix (Eq. (S3.34)) $\mathbf{M}(t_0) = \Sigma$. The time evolution follows Eq. (S3.36):

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

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

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

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

(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 $\bar{\mathbf{k}}_0$ and spreading the spatial Gaussian footprint w.r.t. the variance in frequency.

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

$$\begin{aligned} \psi_r(\bar{\mathbf{q}}; t_0) &= \frac{1}{\psi_r^*(\bar{\mathbf{q}}_0; t_0)} \int d^3\bar{\mathbf{k}}' \mathcal{Q}_r\left(\frac{\bar{\mathbf{q}} + \bar{\mathbf{q}}_0}{2}, \bar{\mathbf{k}}'; t_0\right) e^{i\bar{\mathbf{k}}' \cdot (\bar{\mathbf{q}} - \bar{\mathbf{q}}_0)} \\ &= \sqrt{\frac{|\Sigma_k|}{\pi^3}} \frac{e^{-i\bar{\mathbf{k}}_0 \cdot \bar{\mathbf{q}}'}}{\psi_r^*(\bar{\mathbf{q}}_0; t_0)} e^{-\frac{1}{8}\bar{\mathbf{q}}'^\top \Sigma_q^{-1} \bar{\mathbf{q}}' - \frac{1}{2}\bar{\mathbf{k}}'^\top \Sigma_k \bar{\mathbf{k}}'}. \end{aligned} \quad (\text{S4.8})$$

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

$$|\psi_r(\bar{\mathbf{q}}_0; t_0)|^2 = \int d^3\bar{\mathbf{k}}' \mathcal{Q}_r(\bar{\mathbf{q}}_0, \bar{\mathbf{k}}'; t_0) = \sqrt{\frac{|\Sigma_k|}{\pi^3}}. \quad (\text{S4.9})$$

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*:

$$\psi_r(\bar{\mathbf{q}}; \bar{\mathbf{u}}) = \left(\frac{|\Sigma_k|}{\pi^3}\right)^{\frac{1}{4}} e^{i\varphi} e^{-i\bar{\mathbf{k}}_0 \cdot \bar{\mathbf{q}}'} e^{-\frac{1}{8}\bar{\mathbf{q}}'^\top (\Sigma_q^{-1} + 4\Sigma_k) \bar{\mathbf{q}}'}, \quad (\text{S4.10})$$

where we slightly abuse notation and make the mean $\bar{\mathbf{u}} = (\bar{\mathbf{q}}_0, \bar{\mathbf{k}}_0)$ at current time t explicit, with shifted position shorthand $\bar{\mathbf{q}}' = \bar{\mathbf{q}} - \bar{\mathbf{q}}_0$, 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* $\bar{\mathbf{u}} = \bar{\mathbf{u}}$. 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

$$\mathcal{Q} = \sum_{j=1}^{\infty} E_j \mathcal{Q}_r \Big|_{\bar{\mathbf{u}}_j, \mathbf{M}}, \quad (\text{S4.11})$$

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

If we were to recover the corresponding cross-spectral density of the underlying field (via the inverse of Eq. (S3.17) applied to \mathcal{Q}), then by the linearity of the transform it becomes: $\mathcal{C} = \sum_{j=1}^{\infty} E_j \mathcal{C}_j$, where \mathcal{C}_j is simply the CSD of each of the shifted generalized rays. When a CSD can be written as a superposition of CSDs, the implication is that the light described by these CSDs is mutually incoherent. Therefore, we conclude that an arbitrary wave function ψ , with its phase space representation given by the husimi Q representation \mathcal{Q} can be written as

$$\psi(\bar{\mathbf{q}}) = e^{i\varphi} \sum_{j=1}^{\infty} E_j \psi_r(\bar{\mathbf{q}}; \bar{\mathbf{u}}_j), \quad (\text{S4.12})$$

$$\text{such that } \langle \psi_r(\bar{\mathbf{q}}_1; \bar{\mathbf{u}}_1) \psi_r(\bar{\mathbf{q}}_2; \bar{\mathbf{u}}_2) \rangle_{\omega} = \delta_{1j} \mathcal{C}_j(\bar{\mathbf{q}}_1, \bar{\mathbf{q}}_2),$$

for arbitrary spatial variables $\bar{\mathbf{q}}_{1,2} \in \mathbb{R}^3$, and where $\langle \cdot \rangle_{\omega}$ denotes the ensemble-averaging operator over same frequency realizations (as common in the space-frequency formalism of optical coherence theory), and δ_{1j} is the Kronecker delta. That is, the above is a superposition of mutually-incoherent generalized rays, up to an irrelevant global phase term φ .

As a single generalized ray is perfectly coherent w.r.t. itself (the field admits only a single possible realization, viz. Eq. (S4.10)), the above serves as a coherent-modes decomposition of an arbitrary wave function into a collection of Gaussian beams—generalized rays—with shifted origins (the mean at initial time t_0). We note again that this is only possible due to \mathcal{Q} being non-negative, resulting in positive energies E_j .

We would like to point out that the decomposition in Eq. (S4.12) is not the only possible way to understand Eq. (S4.11). We could have used the inverse of the transform that defines the WDF, viz. Eq. (S3.16), to recover a deterministic wave function, whose Husimi Q representation is also as in Eq. (S4.11). As a matter of fact, there are infinitely many such wave functions. There are also infinitely many stochastic processes, that would yield the desired Husimi Q representation. However, recall that in rendering we deal with time-averaged observables (we measure light over periods very long w.r.t. temporal frequency), therefore as all of the possible interpretations of Eq. (S4.11) admit the same cross-spectral density function, they are all indistinguishable in terms of their time-averaged optical phenomena. Therefore, we may interpret Eq. (S4.11) at will, and we chose the interpretation that is the most suitable for rendering: a coherent-mode decomposition into mutually-incoherent generalized rays.

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 \mathcal{W} its corresponding WDF. The 3×3 spatial-coherence covariance matrix—termed the *shape matrix*—around a spatial point $\bar{\mathbf{q}}$ can be

written as:

$$\begin{aligned}\Theta(\vec{q}) &= \frac{1}{c(\vec{q}, \vec{q})} \int d^3\vec{q}' \vec{q}' \vec{q}'^\top c\left(\vec{q} - \frac{1}{2}\vec{q}', \vec{q} + \frac{1}{2}\vec{q}'\right) \\ &= \frac{1}{|\psi(\vec{q})|^2} \int d^3\vec{q}' \vec{q}' \vec{q}'^\top \int d^3\vec{k}' \mathcal{W}(\vec{q}, \vec{k}') e^{i\vec{q}' \cdot \vec{k}'}.\end{aligned}\quad (\text{S4.13})$$

Formally-interchange the orders of integration, and note that

$$\int d^3\vec{q}' \vec{q}' \vec{q}'^\top e^{i\vec{q}' \cdot \vec{k}'} = -(2\pi)^3 \frac{\partial^2}{\partial \vec{k}'^2} \delta^3(\vec{k}'), \quad (\text{S4.14})$$

i.e., the Hessian matrix of the Dirac delta. Then, for “well-behaved” \mathcal{W} :

$$\begin{aligned}\Theta(\vec{q}) &= -\frac{(2\pi)^3}{|\psi(\vec{q})|^2} \int d^3\vec{k}' \mathcal{W}(\vec{q}, \vec{k}') \frac{\partial^2}{\partial \vec{k}'^2} \delta^3(\vec{k}') \\ &= -\frac{(2\pi)^3}{|\psi(\vec{q})|^2} \left[\frac{\partial^2}{\partial \vec{k}'^2} \mathcal{W}(\vec{q}, \vec{k}') \right]_{\vec{k}'=0},\end{aligned}\quad (\text{S4.15})$$

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

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{q} is proportional to the inverse of covariance of angular spread of light at \vec{q} .

S4.3 Generalized Rays as Coherent States

The Husimi Q representation (Eq. (S4.3)) spreads every point of the WDF in phase space into the minimum footprint that can be resolved. While this procedure yields a strictly non-negative distribution (unlike the corresponding WDF \mathcal{W}), \mathcal{Q} still remains only a quasi-probability density function, because it no longer reproduces the marginals of the underlying wave function (i.e., Property (II) no longer holds).

In quantum mechanics, the Gaussian construct defined in Eq. (S4.4), which we termed a generalized ray, is known as a *coherent state*. We have shown that these coherent states decompose the Husimi Q representation. However, as the WDF is not non-negative, we may not write the WDF as a sum of coherent states (with non-negative intensities), in a manner similar to the Husimi Q representation. This means that our observables, when working with the Husimi Q representation and coherent states, are no longer the position and momentum operators of the underlying wave function.

It might then appear that a generalized ray does not correspond to observable properties of light. However, it well known that coherent states are the eigenstates of the photon creation and annihilation operators [Mandel and Wolf 1995]. Essentially all our sensors operator via the process of photoelectric detection, i.e. by absorbing photons, hence it is the *artefacts of photoelectric detection that we observe*. Being eigenstates of the photon annihilation operator, coherent states then closely correspond to light properties measured via photoelectric detection [Mandel and Wolf 1995]. Generalized rays, under the coherent-modes decomposition outlined via Eq. (S4.12), then correspond to the observable (ensemble-averaged) response of

photoelectric detectors excited by a statistical ensemble of waves (a stream of photons).

S4.4 Propagation of Generalized Rays

Consider a generalized ray, with its phase space distribution quantified by Eq. (S4.4) and wave function as in Eq. (S4.10). It is easy to see that the expectations of its position and frequency operators (Eqs. (S3.2) and (S3.8)) are

$$\langle \hat{q} \rangle = \langle \psi_r | \hat{q} | \psi_r \rangle = \int d^3\vec{q} \vec{q} |\psi_r|^2 = \vec{q}_0, \quad (\text{S4.16})$$

$$\langle \hat{k} \rangle = \langle \psi_r | \hat{k} | \psi_r \rangle = -i \int d^3\vec{q} \psi_r^* \frac{\partial}{\partial \vec{q}} \psi_r = \vec{k}_0, \quad (\text{S4.17})$$

as expected. The variances are also trivially calculated,

$$\begin{aligned}\langle \hat{q}^2 \rangle - \langle \hat{q} \rangle^2 &= \langle \psi_r | \hat{q}^2 | \psi_r \rangle - \langle \hat{q} \rangle^2 \\ &= \int d^3\vec{q} \vec{q}^2 |\psi_r|^2 - \vec{q}_0 \vec{q}_0^\top = \Sigma_q,\end{aligned}\quad (\text{S4.18})$$

$$\begin{aligned}\langle \hat{k}^2 \rangle - \langle \hat{k} \rangle^2 &= \langle \psi_r | \hat{k}^2 | \psi_r \rangle - \langle \hat{k} \rangle^2 \\ &= - \int d^3\vec{q} \psi_r^* \frac{\partial^2}{\partial \vec{q}^2} \psi_r - \vec{k}_0 \vec{k}_0^\top = \Sigma_k.\end{aligned}\quad (\text{S4.19})$$

Finally, the covariance matrices are selected such that they fulfil the equality in the uncertainty relation, viz.

$$\left| \langle \hat{q}^2 \rangle - \langle \hat{q} \rangle^2 \right|^{\frac{1}{2}} \left| \langle \hat{k}^2 \rangle - \langle \hat{k} \rangle^2 \right|^{\frac{1}{2}} = \sqrt{|\Sigma_q| |\Sigma_k|} = \frac{1}{2^3}, \quad (\text{S4.20})$$

illustrating that a generalized ray initially occupies the least phase-space volume permissible.

Let that generalized ray propagate a distance of z in free space. In phase space, this constitutes a horizontal shear, as described (Eq. (S4.7)), hence the generalized ray’s phase-space distribution after propagation can be expressed in terms of its initial distribution, viz.

$$\mathcal{Q}_r(\vec{q}, \vec{k}; z) = \mathcal{Q}_r\left(\vec{q} - \frac{z}{k} \vec{k}, \vec{k}; z = 0\right). \quad (\text{S4.21})$$

The first and second moments become

$$\langle \hat{q} \rangle \Big|_z = \vec{q}_0 + \frac{\vec{k}_0}{k} z, \quad \left(\langle \hat{q}^2 \rangle - \langle \hat{q} \rangle^2 \right) \Big|_z = \Sigma_q + \frac{z^2}{k^2} \Sigma_k, \quad (\text{S4.22})$$

$$\langle \hat{k} \rangle \Big|_z = \vec{k}_0, \quad \left(\langle \hat{k}^2 \rangle - \langle \hat{k} \rangle^2 \right) \Big|_z = \Sigma_k. \quad (\text{S4.23})$$

The mean position has been translated by a distance of z in direction \vec{k}_0 , spatial variance is increased (i.e., the beam occupies greater spatial extent), while the frequency operator moments remain unchanged, as expected in free space. We may then define the *width* (i.e. spatial variance) of the generalized ray as

$$w(z) \triangleq \Sigma_q + \frac{z^2}{k^2} \Sigma_k. \quad (\text{S4.24})$$

Induced correlation. It is of particular insight to consider the correlation between the position and frequency operators for a generalized ray. Initially, a coherent state (generalized ray) is uncorrelated, meaning that the covariance matrix Σ of its phase space Gaussian is defined to be block diagonal, with no correlation arising

between position and momentum. We will show that propagation induces correlation.

To simplify the analysis, we consider only the one-dimensional case, and denote the initial variances $\sigma_{q,k}^2 = |\Sigma_{q,k}|$. These variances fulfil the equality in the uncertainty relation, which in the one-dimensional case reads $\sigma_q \sigma_k = \frac{1}{2}$. The correlation between position and frequency can then be written as

$$\omega_{qk} \triangleq \frac{\langle \hat{q}\hat{k} \rangle + \langle \hat{k}\hat{q} \rangle - 2\langle \hat{q} \rangle \langle \hat{k} \rangle}{2\left(\langle \hat{q}^2 \rangle \langle \hat{k}^2 \rangle\right)^{\frac{1}{2}}} = \frac{z}{\sqrt{4k^2\sigma_q^4 + z^2}}, \quad (\text{S4.25})$$

and we may now write the product of the variances in position and frequency space as

$$\begin{aligned} \left| \langle \hat{q}^2 \rangle - \langle \hat{q} \rangle^2 \right|^{\frac{1}{2}} \left| \langle \hat{k}^2 \rangle - \langle \hat{k} \rangle^2 \right|^{\frac{1}{2}} \Big|_z &= \sqrt{\left(\sigma_q^2 + \frac{z^2}{k^2} \sigma_k^2\right) \sigma_k^2} \\ &= \sqrt{\frac{1}{4} + \frac{z^2}{4^2 k^2 \sigma_q^4}} = \frac{1}{2} \frac{\sqrt{4k^2 \sigma_q^4 + z^2}}{2k \sigma_q^2} = \frac{1}{2} \frac{1}{\sqrt{1 - \omega_{qk}^2}}, \end{aligned} \quad (\text{S4.26})$$

where we use the fact $\sigma_k^2 = \frac{1}{4\sigma_q^2}$.

Observe, that when $z = 0$, no correlation arises, i.e. $\omega_{qk} = 0$, however after propagation, $z > 0$, the uncertainty between position and frequency is increased by the positive factor $(1 - \omega_{qk}^2)^{-1/2}$. In the limit $z \rightarrow \infty$, we may see that $\omega_{qk} = 1$ implying that uncertainty is infinite, and that the generalized ray occupies infinite phase-space volume. While a generalized ray in its initial state serves as a coherent state (as discussed in Subsection S4.3), after propagation it becomes a *correlated coherent state* (a form of a squeezed coherent state). As also noted by Man'ko and Wolf [2008], we may conclude that propagation induces correlation between position and momentum, and this correlation serves to enlarge the Heisenberg uncertainty constant \hbar by a factor of $(1 - \omega_{qk}^2)^{-1/2}$ (as evident from Eq. (S4.26)).

The induced correlation on propagation reduces our ability to resolve phase-space details via generalized rays after propagation. In other words, the phase space becomes increasingly “blurred” as we propagate away from a sensor. Intuitively, it may be understood as saying that the farther objects are from the sensor, the harder it is to resolve the object’s features. It also means that we may not “break” a generalized ray into smaller rays after propagation, but can only use generalized rays that fulfil the uncertainty relation scaled by the induced correlation factor $(1 - \omega_{qk}^2)^{-1/2}$.

S4.5 Summary

The purpose of the discussion above is to find the closest, physically-realizable analogue of the ray, i.e. an impulse response in phase space, in the wave-optical context. We formally derive an expression for such a generalized ray, and furthermore show that an arbitrary Husimi Q distribution can be written as a coherent-modes decomposition of such generalized rays.

Thus, we regain the ability to perform phase space point-queries in order to do light transport via the mapping

$$\vec{u} \mapsto \psi_r(\vec{q}; \vec{u}), \quad (\text{S4.27})$$

which maps between a classical ray and its corresponding wave function. The system evolution of these pair of constructs is similar, with respect to Liouville’s equation (Property (VI)). Observe that Liouville’s theorem (Eq. (S2.15)) implies that the classical probability $\rho(\vec{u}) d^6\vec{u}$ that a ray is found in the phase-space volume $d^6\vec{u}$ is in-fact a property of the ray \hat{u} , and remains invariant on system evolution. That is, the phase-space density around a ray transforms as the ray does. The generalized ray makes this fact explicit: the Gaussian spread of the generalized ray’s energy around its mean is a property of the ray, and this spread transforms in an identical fashion to how ρ transforms.

Light transport operations can be classified into two categories:

- (1) **Simple** — propagation in homogeneous media with slowly-varying refractive-index, and perfect reflections and refractions of a smooth (potentially curved) interface. These interactions are accurately handled by the ABCD formalism (Eqs. (S4.5) and (S4.6)), and we continue to understand light as a ray. The implication being that classical rendering and sampling techniques apply, under these interactions, unaltered.
- (2) **Diffractive** — any interaction with an optical system that may not be classified as an ABCD system. For example, scattering by a rough surface, interaction with particles in media, and free-space diffraction. When dealing with diffractive interactions, we model the perfectly-coherent interaction of the wave function of the generalized ray—the Gaussian beam—with the matter.

Note that generalized rays are mutually incoherent, hence they superpose linearly in terms of their intensities E , which can be equivalently rephrased in terms of radiometric quantities, just as classical rays do. Therefore, we do not need to propagate a wave’s phase, nor do we deal with partially-coherent interactions. All interactions are either incoherent, on superposition of rays, or perfectly-coherent, on diffractive interactions between a singular generalized ray and matter. It is important to remember that propagation induces correlation of position and momentum in phase space, meaning that our ability to resolve details diminishes, as discussed in Subsection S4.4.

We have formulated generalized rays to be coherent states ensemble-averaged over the stochastic process that quantifies the underlying electromagnetic field distribution. This stochastic process models randomness that arises due to rapid fluctuations of the field, i.e. randomness that is not observed over the course of a single measurement by a sensor. That is, generalized rays correspond to the observable response of photoelectric detectors. In fact, as coherent states are the only eigenstates of the photon creation and annihilation operators [Mandel and Wolf 1995], generalized rays are the only formal construct that corresponds to the observable (ensemble-averaged) response of photoelectric detectors.

In relation to optical coherence theory. A highly convenient consequence of the presented light transport formalism is that it is devoid of the cumbersome two-point formulation of optical coherence theory, viz. the cross-spectral density $C(\vec{q}_1, \vec{q}_2)$. In other words, we do not explicitly consider the spatial correlation between wavefronts. Instead, we show in Subsection S4.2 that the variance in spatial coherence is fully described by the spatial-frequency curvature of the

signal (2^{nd} -order derivative w.r.t. the k coordinate) in phase space. Given a beam of light that consists of generalized rays (recall that we may always decompose any wave function into a superposition of such rays), the angular spread of these rays—quantified by the k -moments of the moment matrix (lower right 3×3 minor of \mathbf{M} in Eq. (S3.34))—is known as the *diffusivity* of light.

That is, when the cross-spectral density is used to describe light in-order to do light transport, partially-coherent beams are propagated, and the cross-section of these beams should be at least similar to the extent over which spatial coherence is retained. Backward path tracing (sensor to light source), as well as importance sampling light-matter interactions, is then complicated by the fact that the coherence properties of light are unknown a priori, before connecting a path to a light source. On the other hand, if we use generalized rays to sample the wave-optical phase space (the natural setting in light transport), then, as discussed, all our interactions are either classical (incoherent) or perfectly-coherent, and no a priori knowledge of the coherence of light is required. While we technically still propagate beams (viz. Eq. (S4.10)), the extent of these beams in phase space is only dictated by the uncertainty relation, and is an order-of-magnitude less than the extent of partially-coherent beams.

Spatial-coherence effects are reproduced as we sample the diffusivity of light, i.e. the k space: for example, a perfectly collimated beam of zero diffusivity—where all the generalized rays have identical mean \vec{k}_0 and propagate in the same direction—is perfectly coherent, while a beam with a positive variance in its mean wavevectors \vec{k}_0 , and hence positive diffusivity, is partially coherent. As such a beam interacts with matter, the fully-coherent interaction is blurred as generalized rays with different mean wavevectors are sampled, inducing partial coherence.

As an aside we note that a similar relation between the angular spread (diffusivity) of light and optical coherence is described by the van Cittert–Zernike (VCZ) theorem. However, a typical setting of the VCZ theorem is relating an incoherent source distribution to the mutual coherence at the destination. Eq. (S4.15) is more general: the spatial variance in coherence of arbitrary light is related to the k moments in phase space. Furthermore, VCZ typically assumes a far-field setting where the implied decomposition is into plane waves. While the wavefront of a Gaussian beam can, in the far field (Fraunhofer region), be approximated by a plane wave, we do not make a far field assumption. Indeed, a mapping between a classical ray to a plane wave is highly unnatural: a ray is sharp impulse in phase space, while a plane wave has an equal probability of being measured throughout the entire space. Instead, by analysing the wave-optical phase space induced by the WDF, we map the classical ray to a Gaussian beam: a more general and a more pleasant analogue.

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.

- 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.
- Leonard Mandel and Emil Wolf. 1995. *Optical coherence and quantum optics*. Cambridge University Press, Cambridge.
- Vladimir I Man'ko and Kurt Bernardo Wolf. 2008. The influence of spherical aberration on gaussian beam propagation. In *Lie Methods in Optics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 207–225.
- 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>
- 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.