

A Generalized Ray Formulation For Wave-Optical Light Transport

SHLOMI STEINBERG, NVIDIA, United States and University of Waterloo, Canada

RAVI RAMAMOORTHY, 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. 1. **Locality in wave optics.** Spectral rendering of a spaceship model (10 cm in length) with a scratched aluminium fuselage. To render the scratches, we either use (a) our method, or for the bottom insets (b) we use a method similar to Werner et al. [2017]. Colourful wave effects arise when distinct wavefronts—scattered from the base surface and phase-shifted scatter from the scratches—superpose and interfere, hence reproduction of these effects requires integration over a positive region. Werner et al. [2017] use a Gaussian integration patch of a constant size, chosen ad hoc. We denote as r the radius of this integration region. Observe that when that region is too small, wave effects fail to materialize, and a purely geometric optics response arises, as expected. As we increase r , the correct effects emerge, up until further increases produce no observable response in the rendered appearance. This highlights two problems we address in this paper (discussed further in Section 1): (i) choosing the wavefront shape and size used for the integration of a wave-optical BSDF is not trivial; and (ii) the *sampling problem*: further increase of r has no observable effect but nevertheless induces a significant change in the integrated BRDF (note the log-scale BRDF lobe visualizations), meaning that *the BRDF used for importance sampling differs from the actual observed BRDF*. The sampling problem can make integrating wave effects a significant challenge (see Figs. 7 and 8). We present a formal framework that uses *generalized rays*—weakly-local constructs that exactly quantify the shape and size of wavefronts that BSDF need to act upon—to accurately sample and render any wave-optical distribution.

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.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3687902>.

Ray optics is the foundation of modern path tracing and sampling algorithms for computer graphics; crucially, it allows high-performance implementations based on ray tracing. However, many applications of interest in computer graphics and computational optics demand a more precise understanding of light: as waves. For example, accurately modelling scattering effects like diffraction or interference requires a model that provides the coherence of light waves arriving at surfaces. While recent work in Physical Light Transport [Steinberg et al. 2022; Steinberg and Yan 2021] has introduced such a model, it requires tracing light paths starting from the light sources, which is often less efficient than tracing them from the sensor, and does not allow the use of many effective importance sampling techniques.

We introduce a new model for wave optical light transport that is based on the fact that sensors aggregate the measurement of many light waves when capturing an image. This allows us to compactly represent the statistics of light waves in a *generalized ray*. Generalized rays allow sampling light paths starting from the sensor and applying sophisticated path tracing sampling techniques while still accurately modelling the wave nature of light. Our model is computationally efficient and straightforward to add to an existing path tracer; this offers the prospect of wave optics becoming the foundation of most renderers in the future. Using our model, we show that it is possible to render complex scenes under wave optics with high performance, which has not been possible with any existing method.

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

Additional Key Words and Phrases: wave optics, optical coherence, light transport, PLT

1 Introduction

Ever since the introduction of the rendering equation [Kajiya 1986], modern light transport algorithms have been built on an assumption of ray optics. Ray optics is able to accurately represent many forms of light scattering in the physical world, and it has two key properties that lead to efficient light transport simulation algorithms:

- (1) **Locality**: the light-transport primitive, i.e., the ray, exists at a singular position (a point) in space and propagates into a singular direction at any given time. This enables high-performance ray tracing algorithms that can be used to simulate light propagation. Today such algorithms can also be hardware-accelerated [Burgess 2020].
- (2) **Linearity**: the radiometric quantities associated with rays (e.g., intensity or radiance) sum linearly. Rays do not interact with other rays, which further enables high-performance and parallel implementations.

The ray optics model has the additional advantage that it is *complete*: any effect that can be described by the applicable physics may also be simulated.

Yet because ray optics does not simulate the wave nature of light, it is not able to accurately simulate many important optical phenomena. Such phenomena include the colourful glints that appear when light is scattered by scratches in metal; the colour of the wings and scales of some species of insects, snakes and fish; stress birefringence; and the appearance of an oil layer or metal oxide on a surface. While there is a significant amount of previous work in modelling such phenomena in computer graphics (discussed further in Section 2), this work has focused on modelling surface scattering and has required that significant assumptions be made about the wave properties of the incident light at points being shaded.

Fig. 1 shows how these assumptions can lead to errors in rendered images. In that scene, a BSDF that models interference is used for the scratched spaceship body. In order to accurately render its appearance, it is necessary to integrate over the correct area around each shaded point: this area of integration depends both on the wave coherence of the incident light as well as on the observer's ability to resolve optical details. If an arbitrary coherence is assumed, as has been done in previous work that is based on ray optics light transport, the image may be incorrect. As additional examples, Fig. S3 in our supplemental material demonstrates diffraction

of light around geometric obstacles, and Fig. 6 shows how some materials may exhibit different appearance depending on the light's coherence properties. In all these cases, choosing the correct integration region is essential, but doing so accurately is challenging.

The coherence of light arriving at a point on a surface depends not only on the emissive properties of the light sources, but it is also affected by scattering from other surfaces before it reaches the point. Thus, an accurate simulation of wave-optical light transport is required to provide this information. One might hope to formulate a wave-optical formalism that would provide locality, linearity, and completeness, but unfortunately that is impossible: a single-point formulation of light—one that is simultaneously both *local* and *linear*—is incompatible with electromagnetism [Wolf 2007].

Extensive research has been devoted to study alternative formulations of wave optics that regain some of the “grainy” picture of ray optics, where point queries of light's behaviour may be performed. While locality and linearity cannot be achieved simultaneously, many computational electromagnetic frameworks have been designed to be perfectly local: i.e., they employ a ray-like construct to conduct wave simulations. Unfortunately, any such formalism inevitably abandons linearity and frustrates the application of sampling and path tracing techniques (see Section 2).

Recent developments in Physical Light Transport (PLT) algorithms have provided a *weakly-local* wave optics light transport framework: rather than perfectly local rays, the transport primitive is a beam that occupies a small, well-defined spatial region and propagates into a solid angle [Steinberg et al. 2022; Steinberg and Yan 2021]. PLT sets the spatial region occupied by a beam to be proportional to the spatial coherence of light such that distinct beams become mutually incoherent and beams superpose linearly, enabling the application of some classical path tracing tools.

However, PLT is based on a *forward* model of light transport, where paths are sampled starting from light sources. Sampling paths from the light sources is inefficient in many scenes, for example if only a small part of a large scene is being rendered. Although it is possible to sample paths starting from the sensor with PLT if a global lower limit on the coherence of light is defined [Steinberg et al. 2022], such a limit must be conservative. Using such a limit leads to poor sampling for materials with strong wave-interference effects, which are precisely the materials of interest. In that case, the renderer uses a BSDF for importance sampling that significantly differs from the observable BSDF, as shown in the plots in Fig. 1. We call this incompatibility of weakly-local formalisms with backward transport the *sampling problem*.

In this paper, we develop a new formalism that efficiently and accurately models wave-optical light transport and allows path sampling from the sensor. Our formalism is simultaneously *weakly-local*, *linear*, and *complete*. This is a powerful result: completeness means that our formalism can simulate arbitrary wave-optical distributions of light, of any wavelength, any spectrum, any polarization, and any optical coherence.

Our approach takes advantage of the fact that sensors integrate incident light waves over a small area; it is not the individual waves

we are concerned with, but their distribution. This allows us to introduce the *generalized ray*, the wave-optical analogue of the classical ray. Light can be understood as a collection of such mutually-incoherent generalized rays, thereby regaining linearity and locality (as much as is physically allowed by the uncertainty relation). Conveniently, the weak locality of a generalized ray does not depend on the properties of light, and only depends on the properties of the sensor: its angular and spatial sensitivity, and sensed wavelength.

With generalized rays, we are able to propagate the detection states of a detector through the scene under time-reversed dynamics instead of propagating the emission distributions of light forward in time. However, in contrast to ray optics, where a ray of light has identical dynamics to its time-reversed counterpart, within the wave-optical context there is a fundamental change in physics between forward and backward light transport. In Section 4, we develop the theory of backward wave-optical light transport from optical principles, making essentially no approximating assumptions. This allows us to derive light transport equations for generalized rays that generalize their counterparts in classical ray optics.

This leads to our second major contribution: the first practical wave-optical light transport algorithm for rendering. Our theory also allows us to develop an algorithm for wave-optics light transport that allows the use of the effective importance sampling techniques that have been developed for path tracing over the past decades. This algorithm, *sample-solve*, has two steps: (i) Sample: generalized rays are perfectly coherent, so they can be used to sample paths in a coherence-agnostic manner, solving the sampling problem. (ii) Solve: once a path connecting a light source to the sensor has been sampled, we apply PLT tools in order to *solve* for the partially-coherent light transport over that path. This is possible due to the exact and formal relationship between generalized rays and optical coherence that we derive.

With our approach, several changes to a conventional path tracer are necessary to support wave-optics light transport:

- (1) A few additional quantities are associated with each ray.
- (2) The entire ray path from source to sensor is stored for the *solve* step.
- (3) Wave-optics BSDFs are implemented to benefit from the information our framework makes available to them.
- (4) The *solve* step is to be implemented: in this step, we retrace the sampled path and apply PLT in order to better resolve partially-coherent effects. The solve step acts as a variance-reduction technique.

We have implemented our algorithms in *Falcor*, a pre-existing GPU-accelerated path tracer; the implementation is available in the supplemental material. We have also implemented manifold sampling [Hanika et al. 2015; Zeltner et al. 2020] as an example to show the value of existing sampling techniques in the wave optics setting. We show in Section 5.1 that we are able to do wave-optical rendering with complex scenes, with high performance, and with a convergence rate that is multiple orders-of-magnitude faster compared to the state-of-the-art.

To date, a weakly-local, linear and complete wave-optical framework of light has not been formulated, and this is the primary theoretical contribution in this paper. While the primary focus of this paper is upon practical wave-optical rendering, our theoretical contributions are more general and serve as a link between path tracing techniques and wave-optical simulations. In the future, a wave-optical light transport framework, building upon our contributions, may be able to target many important unsolved problems that demand wave simulation in complex scenes, e.g., RADAR simulation in urban environments for automotive applications, and signal coverage simulation for efficient infrastructure deployment.

In the paper, the sections marked with an asterisk (*) are mathematically dense, and readers that are less interested in the theory might wish to (initially) skip these sections.

2 Related Work

Wave-optical light transport. In computer graphics one of the earliest formulations of wave-optical light transport is *diffractive shaders* [Stam 1999]. Light is propagated as rays, and a ray is converted into a plane wave (implicitly an asymptotic, far-field approximation) when formulating a diffractive shader, which quantifies the interaction of that ray with a surface. A plane wave exists throughout the entire space, therefore no degree of locality could be rigorously recovered.

Wigner distribution-based light transport methods have gained some attention both in computer graphics [Cuyppers et al. 2012] and optical literature [Jensen and Buot 1991; Mackay and Johnson 2021; Mout et al. 2018]. The *Wigner distribution function* (WDF) $\mathcal{W}(\vec{r}, \vec{k})$ (see formal definition in Eq. (4)) is a bilinear distribution defined w.r.t. some wave function $\psi(\vec{r})$, and is a function of spatial position $\vec{r}$ and wavevector $\vec{k}$. The WDF admits some useful ray-like properties, which has motivated the definition of the WDF as a form of a “generalized radiance” [Walther 1968]. However, locality is lost when the interaction with matter is considered, as we formally show next. The scattered WDF is [Torre 2005, Chapter 8]

$$\mathcal{W}_o(\vec{r}_o, \vec{k}_o) = \int d\vec{r}_i d\vec{k}_i K(\vec{r}_o, \vec{r}_i, \vec{k}_o, \vec{k}_i) \mathcal{W}_i(\vec{r}_i, \vec{k}_i), \quad (1)$$

where $\mathcal{W}_i, \mathcal{W}_o$ are the input and output WDFs, respectively, and K is a *diffraction kernel* that is completely defined by the scattering matter [Testorf et al. 2010, Chapter 1.6]:

$$K(\vec{r}_o, \vec{r}_i, \vec{k}_o, \vec{k}_i) \triangleq \int d\vec{x}_o d\vec{x}_i h\left(\vec{r}_o + \frac{1}{2}\vec{x}_o, \vec{r}_i + \frac{1}{2}\vec{x}_i\right) \times h^*\left(\vec{r}_o - \frac{1}{2}\vec{x}_o, \vec{r}_i - \frac{1}{2}\vec{x}_i\right) e^{-i(\vec{k}_o \cdot \vec{x}_o - \vec{k}_i \cdot \vec{x}_i)}, \quad (2)$$

where $\star$ denotes complex conjugation, and h is the *optical response function* of the system. Locality is lost because, by definition, K requires integration over the *entire scene*.

Cuyppers et al. [2012] claim that the negative values taken by the WDF allow “rays to interfere later for global illumination”. This is only true when a single WDF accounts for all light in the system, and the kernel K accounts for the entire scene. The WDF is a bilinear distribution, therefore the WDF of the sum $\psi_1 + \psi_2$ is

$$\mathcal{W} = \mathcal{W}_1 + \mathcal{W}_2 + \mathcal{W}_{12}, \quad (3)$$

where $\mathcal{W}_1, \mathcal{W}_2$ are the WDFs of ψ_1, ψ_2 , respectively, and $\mathcal{W}_{12}$ is a bilinear cross-term, that depends on both ψ_1 and ψ_2 (see Testorf et al. [2010, Chapter 1.8] for more details). If we were to limit the spatial integration in Eq. (2) to a finite region, then different kernels would account for different interactions, and to compute the total superposition WDF with contributions from rays scattered from distinct spatial regions, we must account for the bilinear term in Eq. (3). Cuyppers et al. [2012] implicitly neglect that term: by assuming that the WDFs from different rays add up *linearly*. The negative values that arise in the WDF of a light distribution depend only on that distribution, while the bilinear cross-term that integrates over both distributions is the term that accounts for their mutual interference. Discarding the cross-term is equivalent to neglecting wave interference. We numerically demonstrate in Fig. S7 in our supplemental material that limiting the integration region in Eq. (2) and not accounting for the bilinearity of the WDF produce incorrect results.

Both Stam [1999] and Cuyppers et al. [2012] recognize the need for linearity, however neither admits machinery that is able to limit the spatial integration region of light-matter interactions to a well-defined finite extent, hence neither admits locality of any form. Even if such machinery was to be developed, these frameworks would then suffer (like any other weakly-local forward formalisms) from the very same sampling problem that we address in this paper.

Physical light transport (PLT) [Steinberg et al. 2022; Steinberg and Yan 2021] does derive such machinery by quantifying the coherence of light. Under the assumption of weakly-coherent light, rigorous, well-defined weak locality is recovered. That is, PLT is able to quantify the spatial extent over which a light-matter interaction (like Eq. (2)) may produce observable interference. Treating light as partially-coherent beams, PLT gives rise to a weakly-local, linear formalism of wave-optical light transport. PLT is able to accurately and efficiently reproduce partially-coherent effects. However, in order to do that, PLT needs to propagate the coherence properties of light forward—properties that are needed in order to recover weak locality—hence the sampling problem inevitably arises, and practical backward light transport is not possible. Our method can be understood as generalizing optical coherence to the backward model of light transport, and we derive a formal connection to PLT and optical coherence in Section S7 in our supplemental.

Computational electrodynamics. A variety of wave solvers, many commercially available, aim to produce a solution to Maxwell’s equations. Common solvers are variants of the finite-difference time-domain (FDTD) [Yee 1966], finite element (FEM) [Jin 2015], or boundary element (BEM) methods. These methods are usually supra-linear in scene complexity, require deterministically described fields as well as geometry at sub-wavelength resolution. Furthermore, solving for the electromagnetic field directly is neither feasible nor desirable in most applications of interest: *observable properties* of light arise when integrated over the extent of a sensor as well as over the period of observation. These methods are practical only in exceedingly simple scenes.

Of more relevance are “physical optics” methods, most notably the “shooting-bouncing ray” (SBR) method. SBR is a perfectly-local

method, by design, and uses ray tracing to track field propagation to a target aperture or geometry (i.e., a ray captures the behaviour of the field as it propagates into some solid angle). Then, the electromagnetic field is reconstructed at the target in order to evaluate the scattered electromagnetic radiation. Applications of SBR methods are extensive, and include: simulation of radar for imaging [Feng and Guo 2021]; driving-assistive technology [Castro et al. 2019]; analysis of aircraft scattering cross-section [Bilal et al. 2019]; ground-penetrating radar [Warren et al. 2016]; and, indoor positioning using WiFi [Hossain et al. 2018]. Some methods point-sample an input field and propagate a ray tube in order to approximate a diffraction integral [Andreas et al. 2015].

Another popular class of perfectly-local methods are ray tracers that handle diffractions by employing the geometric/uniform theory of diffraction (UTD) [Bilibashi et al. 2020; Son and Myung 1999; Yi et al. 2022]. Some of these have become the state-of-the-art in the simulation of light of longer wavelengths: major applications include automotive-targeted simulation of RADAR [Boban et al. 2014; Guan et al. 2020], as well as simulation of WiFi/cellular radiation (e.g., for analysis of signal coverage in a city) [Choi et al. 2023; de Adana et al. 2005]. Only a few select works are cited. Due to their importance, these applications have garnered significant attention, and better solutions would be of real interest.

SBR and UTD-based ray tracing methods are perfectly local, and some may be complete, but as discussed, such methods can never be linear. That is, rays are always *mutually interfering*, and their interference on superposition must always be considered. As very many rays need to be traced, attempts to accelerate SBR methods have taken a research trajectory not dissimilar to early computer graphics work: Employing multi-resolution grids [Suk et al. 2001] and spatial-subdivision data structures to accelerate ray-facet intersections [Jin et al. 2006; Tao et al. 2008]; accelerating on GPUs [Gao et al. 2015; Tao et al. 2010]; and, sampling the initial ray directions using Halton sequences [Key et al. 2018]. Nevertheless, as linearity is unavoidably lost, this research greatly trails computer graphics in its ability to apply sophisticated sampling techniques, like importance sampling light-matter interactions or path guiding. Scene complexity remains a very limiting factor.

Material-level “diffractive BSDFs”. Also related is work that aims to reproduce the appearance of some diffractive materials. This includes the rendering of iridescent and pearlescent materials [Guilén et al. 2020]; diffractive scratches [Velinov et al. 2018; Werner et al. 2017]; statistical surface profiles [Holzschuch and Pacanowski 2017; Krywonos 2006; Steinberg and Yan 2022] or diffractive surfaces with explicit microgeometry [Falster et al. 2020; Yu et al. 2023]; free-space diffractions [Steinberg et al. 2024]; and, thin-film interference at a soap bubbles [Huang et al. 2020] or due to a dielectric layer over a conductor [Belcour and Barla 2017; Kneiphof et al. 2019]. Synthesis of BSDFs that account for wave interference was discussed by Toisoul and Ghosh [2017].

The cited work does not deal with light transport in a scene, but only considers the result of an interaction with a material. It is the light-transport framework’s responsibility to supply BSDFs that perform wave-optical interactions with the information about the waveforms these BSDFs need to diffract. The light-transport




Fig. 2. **The generalized ray in phase space.** It is often convenient to depict the distribution of light in a space known as *phase space*: an artificial space that can be loosely understood as the Cartesian product of position and direction of propagation. In this space, at a particular time instant, a ray-optical ray is a point (functionally a Dirac delta) as its position and direction are exactly known. Under wave optics, such a distribution of light is prohibited and the phase space is discretized into overlapping Gaussian cells. A *generalized ray*—the convolution of a Dirac with a *minimum uncertainty Gaussian* (a Gaussian with the least variance that fulfils the uncertainty relation, and occupies a single cell) is a physically-realizable *closest analogue to the ray* under wave optics.

framework needs to dynamically and accurately quantify the weak locality of light, and hence the spatial region BSDFs must integrate wave effects over.

3 Background: Theoretical Foundations

We briefly introduce the Wigner picture of wave optics, which is a useful setting for the discussion of the dynamics of light. See our supplemental material or Testorf et al. [2010]; Torre [2005] for additional information. We also formalize the process of photoelectric detection [Mandel and Wolf 1995; Ou and Kimble 1995]. Our interest in photoelectric detection is motivated by how our backward light transport formalism (introduced in Section 4) operates: In contrast to forward models of light transport, which propagate the emission distributions of light forward in time, we propagate the *detection states* of a detector under time-reversed dynamics. This is key in decoupling the light transport simulation from the properties of light, thereby solving the sampling problem.

3.1 Phase-Space Optics

The basic descriptor of light under wave optics is the *wave function*, denoted $\psi(\vec{r}; t)$, where $\vec{r}$ is position and t is time. The wave function quantifies the spatial excitations of a component of the associated electric field. Henceforth, we fix time and, for brevity, drop t from the argument lists. In his seminal work in 1932, Wigner formulated the *Wigner distribution function* (WDF) [Wigner 1932], which describes a *position-momentum* distribution of the signal ψ :

$$\mathcal{W}(\vec{r}, \vec{k}) \triangleq \frac{1}{(2\pi)^3} \int d\vec{r}' \psi^*(\vec{r} - \frac{1}{2}\vec{r}') \psi(\vec{r} + \frac{1}{2}\vec{r}') e^{-i\vec{r}' \cdot \vec{k}}, \quad (4)$$

where $\star$ denotes complex conjugation, $\vec{r}$ is a spatial position, and $\vec{k}$ is the *wavevector*, which quantifies light's temporal frequency and direction of propagation. $|\vec{k}| = \eta \frac{2\pi}{\lambda}$ is the *wavenumber*, with η being the refractive index of the medium in which light propagates

and λ is the light's wavelength. Up to a phase term, the wave function may always be recovered from the WDF, hence the WDF provides a complete description of light.

The domain of the WDF is a space of both positions and wavevectors, referred to as *phase space*. This phase-space treatment of wave optics admits stark similarities to the ray-optical depiction of light: $\mathcal{W}(\vec{r}, \vec{k})$ is the energy density of a “light particle” at position $\vec{r}$ with direction of propagation proportional to $\vec{k}$. $\mathcal{W}(\vec{r}, \vec{k})$ also propagates like the classical ray in free space and under some other interactions. Furthermore, it is well known that the WDF fulfils most of the postulates expected of a classical (ray optical) phase-space density function. The WDF does depart from the ray-optical picture in a couple of fundamental aspects:

- (1) **Bilinearity:** The WDF is a bilinear distribution, hence a linear superposition of WDFs is not the WDF of a superposition of waves (see Eq. (3)).
- (2) **Negativity:** The WDF takes somewhat counter-intuitive negative values, restricting its interpretation as an energy density. These negative values are not an exception, as a matter of fact, non-negative WDFs arise only with one specific class of wave functions [Torre 2005].

See Section S3.1 in our supplemental material for a more complete discussion of the WDF and its properties.

As discussed, the properties of the WDF have motivated other work to use the WDF to perform ray-like “point queries” of light behaviour for wave-optical light transport. However, either locality is entirely lost when interaction with matter considered (viz. Eq. (1)), because the kernel K must account for the entire scene; or, if we restrict K to a subset of the scene, then WDFs that arise from interactions with different such kernels must superpose bilinearly. No formalism can be simultaneously perfectly-local and linear. In addition, as a consequence of the WDF not being non-negative, it is a highly-oscillatory function, as shown in Section S4 in our supplemental material. This means that integration of the WDF with point samples heavily suffers from aliasing.

We are not interested in working with potentially-arbitrary emitted WDFs, as done by other Wigner-based formalisms. Instead, we work in phase space to study the dynamics of *detection states*, introduced next, and show that these states admit far nicer properties than a general WDF.

Gaussian distributions. A phase-space construct central to our discussion is the (symplectic) *Gaussian WDF* [Testorf et al. 2010]:

Phase-space Gaussian

$$g_{\beta, \rho}(\vec{r}, \vec{k}; \vec{r}_0, \vec{k}_0) \triangleq \frac{1}{\pi^3} \exp\left[-\frac{1+\rho^2}{\beta^2} |\vec{r} - \vec{r}_0|^2 - \beta^2 |\vec{k} - \vec{k}_0|^2\right] \times \exp[2\rho |\vec{r} - \vec{r}_0| |\vec{k} - \vec{k}_0|]. \quad (5)$$

As will be discussed in Section 3.2, the (non-negative) values measured by our detectors can always be written as, and only as, a convolution of such phase-space Gaussians with an arbitrary WDF. Therefore, we use these Gaussians as our weakly-local light transport primitives.

The parameterization of the Gaussian is chosen such that we may understand $\beta > 0$ as the initial spatial variance of the WDF g , and

the correlation parameter $\rho \geq 0$ is related to propagation distance. As will be seen later, propagation of light induces phase-space correlation between the spatial and wavevector variables, i.e. propagation increases ρ . In general, β, ρ can be positive-definite matrices, describing anisotropy (see our supplemental material). For simplicity, here we assume scalar parameters. g is normalized such that $\int d\vec{r} d\vec{k} g_{\beta, \rho} = 1$.

Phase-space Gaussians can also be understood as the most compact physically-realizable construct: note that the product of the spatial and wavevector variances of the Gaussian above is $\frac{1}{4}(1 + \beta^4 \rho^2)$, which fulfils the *uncertainty relation* [Mandel and Wolf 1995, Chapter 4], viz. $\sigma_r^2 \sigma_k^2 = 1/4$, if and only if $\rho = 0$.

3.2 Measurement of the WDF *

In sharp contrast to classical physics, we may not measure light without disturbing it: under wave optics, any measurement apparatus deduces information about light by interacting with it. The action of a detector on light can be quantified by the detector's WDF, and the intensity observed by the detector is [Dragoman 2005]

$$I = \int d\vec{r}' d\vec{k}' \mathcal{W}(\vec{r}', \vec{k}') \mathcal{W}_d(\vec{r}', \vec{k}'), \quad (6)$$

where $\mathcal{W}, \mathcal{W}_d$ are the WDFs of light and the detector, respectively. Even though the WDF takes negative values, the result of a measurement is always non-negative, i.e. $I \geq 0$, which formally follows directly from the Moyal formula [Torre 2005, Chapter 6.3.7].

We assume our detectors are *classical photoelectric* detectors. Photons impinging upon a photoelectric detector ionize free electrons, producing a current that is amplified and measured. All detectors of interest (e.g., a photoreceptor in the eye, a pixel in a CMOS array in a camera, or a RADAR antenna) operate in this fashion. Classicality implies that quantum effects are ignored: we assume that the photon flux is sufficiently high that electric pulses from single photons overlap and form a continuous electric current.

As photoelectric detection works by absorption of photons, it is the *photon annihilation* and *photon number* operators that are the observables [Mandel and Wolf 1995, Chapter 11]. The only eigenstates of these operators are the (squeezed) *coherent states* [Mandel and Wolf 1995, Chapters 11, 21], whose phase-space representation takes the form of uncorrelated (minimum-uncertainty) Gaussians, i.e. the Gaussians defined in Eq. (5) with $\rho = 0$, viz.

$$\hat{g}_{\beta}(\vec{r}, \vec{k}; \vec{r}_0, \vec{k}_0) \triangleq \frac{1}{\pi^3} e^{-\frac{1}{\beta^2} |\vec{r} - \vec{r}_0|^2} e^{-\beta^2 |\vec{k} - \vec{k}_0|^2}. \quad (7)$$

As with the general Gaussian g , the parameter $\beta > 0$ above quantifies the trade-off between the spatial and wavevector variances.

Let a detector be described by a distribution $\mathcal{D}$ of detection elements, each quantified by the parameters: (i) $\vec{k}_0$: mean wavevector of detection (quantifies the mean wavelength and direction of propagation of detected light); (ii) β : spatial extent; and, (iii) α : detection efficiency (quantum efficiency). The WDF $\mathcal{W}_d$ of the detector quantifies which light states are detectable [Dragoman 2005]. As the states detectable by a classical photoelectric detector are the coherent states, the detector's WDF may be written as

$$\mathcal{W}_d(\vec{r}, \vec{k}) = \int_{\mathcal{D}} d\mu_{\mathcal{D}} \alpha \hat{g}_{\beta}(\vec{r}, \vec{k}; \vec{r}_0, \vec{k}_0) \quad (8)$$

where $\mu_{\mathcal{D}}$ is the measure over $\mathcal{D}$. In the simplest case, $\vec{k}_0, \alpha, \beta$ are constant, and $\mathcal{D}$ quantifies the spatial region occupied by the detector. Then, the integration is over the spatial extent, with $\vec{r}_0$ being the integration variable. In general, $\mathcal{D}$ may quantify spatially-varying detection properties.

Denote the *Husimi Q distribution* of a WDF $\mathcal{W}$ as the convolution of the WDF with a phase-space Gaussian:

$$\mathcal{Q}(\vec{r}, \vec{k}, \beta) \triangleq \int d\vec{r}' d\vec{k}' \mathcal{W}(\vec{r}', \vec{k}') \hat{g}_{\beta}(\vec{r}', \vec{k}'; \vec{r}_0, \vec{k}_0). \quad (9)$$

Substitute Eqs. (8) and (9) into Eq. (6) and formally interchange the orders of integration:

$$\begin{aligned} I &= \int d\vec{r}' d\vec{k}' \mathcal{W}(\vec{r}', \vec{k}') \int_{\mathcal{D}} d\mu_{\mathcal{D}} \alpha \hat{g}_{\beta}(\vec{r}', \vec{k}'; \vec{r}_0, \vec{k}_0) \\ &= \int_{\mathcal{D}} d\mu_{\mathcal{D}} \alpha \mathcal{Q}(\vec{r}_0, \vec{k}_0, \beta). \end{aligned} \quad (10)$$

The above implies that a classical photoelectric detector directly measures the Husimi Q distribution, and not the WDF. The “smoothing” operation of the WDF over a minimum-uncertainty Gaussian (as in Eq. (9)) is a sufficient and necessary condition to produce a bandwidth-limited, non-negative distribution [Soto and Claverie 1983], and it is the smoothed distribution that classical photoelectric detectors observe [Dragoman 2004; Leonhardt 1997].

4 Theory of Backward Wave-Optical Light Transport

Note the symmetry in the measurement formula, Eq. (6): we may understand it as the detector acting upon the incident distribution of light; or, equivalently, as light acting upon detection states. A forward model works by sourcing a light distribution $\mathcal{W}$ from light sources, simulating its interaction with the scene, and finally integrating it over the detector's distribution. On the other hand, our backward model that we present in this section evolves the detector distribution $\mathcal{W}_d$ under time-reversed dynamics, and then integrates it over the sourcing distribution. In contrast to the sourced WDF, the WDF of the photoelectric detector $\mathcal{W}_d$ (Eq. (8)) admits much nicer analytic properties.

We term the detection states' wave function as the *generalized ray*, see Fig. 2 for an illustration. We formally define the generalized ray and discuss their sourcing from an arbitrary photoelectric detector $\mathcal{D}$ in Section 4.1. In Section 4.2, we formalize the workings of our backward light transport formalism with generalized rays. In Section 5 we build upon the formalism developed here, and present a practical, interactive wave-optical renderer. In addition, in our supplemental material, we analyze this formalism and show that it is weakly-local, linear and complete (Section S5), and interaction kernels will be discussed further (Section S6).

Our approach of reversing the dynamics in order to express wave-optical light transport under the backward model shares conceptual similarities with the *adjoint* and *importance* quantities used in classical rendering and light transport [Christensen 2003]. From this perspective, our generalized ray can be understood as carrying *wave-optical importance*. In similar manner to the classical importance, generalized rays serve to explore the regions in the scene that most contribute to the observed image. However, under the

wave-optical context, generalized rays additionally serve the crucial function of accurately quantifying the region over which interference effects need to be resolved—and thereby generalized rays allow us to regain weak locality.

Time-reversed dynamics. Consider the WDF $\mathcal{W}_s$ of light that is sourced from one or more light sources. Let the light quantified by $\mathcal{W}_s$ propagate around a scene, interact with that scene, and finally impinge upon a classical photoelectric detector, with detection WDF $\mathcal{W}_d$, as in Eq. (8). Let the kernel for the entire process be denoted K (as defined in Eq. (2)), and we denote the action of such kernel on a WDF (as in Eq. (1)) via the *interaction operator*:

$$\mathcal{K}\{\mathcal{W}\}(\vec{r}_o, \vec{k}_o) \triangleq \int d\vec{r}_i d\vec{k}_i K(\vec{r}_o, \vec{r}_i, \vec{k}_o, \vec{k}_i) \mathcal{W}(\vec{r}_i, \vec{k}_i). \quad (11)$$

Denote the WDF of the light that impinges upon the detector, after interaction with the scene, as $\mathcal{W}_o = \mathcal{K}\{\mathcal{W}_s\}$. Apply the measurement formula, Eq. (10), to $\mathcal{W}_o$, yielding an expression for the result of a measurement of light after interaction:

$$\begin{aligned} I &= \int_{\mathcal{D}} d\mu_{\mathcal{D}} \alpha \mathcal{Q}_o(\vec{r}_o, \vec{k}_o, \beta) \\ &= \int_{\mathcal{D}} d\mu_{\mathcal{D}} \alpha \int d\vec{r}' d\vec{k}' \hat{g}_{\beta}(\vec{r}', \vec{k}'; \vec{r}_o, \vec{k}_o) \mathcal{K}\{\mathcal{W}_s\}(\vec{r}', \vec{k}'), \end{aligned} \quad (12)$$

where $\mathcal{Q}_o$ is the Husimi Q distribution (Eq. (9)) of $\mathcal{W}_o$.

Time-reversal induces wavevector reversal, viz. $\vec{k} \rightarrow -\vec{k}$, and phase conjugation [Geru 2018, Chapter 2.3]. Hence, a system's optical response function h fulfils $h(\vec{r}_o, \vec{r}_i) = h^*(\vec{r}_i, \vec{r}_o)$. By applying this relation to Eq. (2), we define the *time-reversed diffraction kernel*:

$$K^t(\vec{r}_o, \vec{r}_i, \vec{k}_o, \vec{k}_i) \triangleq K^*(\vec{r}_i, \vec{r}_o, \vec{k}_i, \vec{k}_o). \quad (13)$$

The time-reversed interaction operator $\mathcal{K}^t$ is then defined to act upon a WDF via K^t , as in Eq. (11). Using Eqs. (2) and (11), we may observe that $\mathcal{K}^{-1} \equiv \mathcal{K}^t$, for lossless kernels. Formally interchange the integration order in Eq. (12) between the integral over the primed variables and the integral over the kernel K , and apply the time-reversed interaction operator, yielding

Backward light transport formalism

$$I = \int_{\mathcal{D}} d\mu_{\mathcal{D}} \alpha \int d\vec{r}' d\vec{k}' \mathcal{W}_s(\vec{r}', \vec{k}') \mathcal{K}^t\{\hat{g}_{\beta}\}(\vec{r}', \vec{k}'). \quad (14)$$

Subject to the physics of photoelectric detection (Section 3.2), all the derivations above are exact. Eq. (12) quantifies the *forward* model of light transport: a distribution of light sourced from light sources, $\mathcal{W}_o$, is evolved forward in time, and then measured by a detector. On the other hand, Eq. (14) quantifies the *backward* model: each detector state of a photoelectric detector is evolved backward, under the time-reversed wave-optical dynamics (as quantified by Eq. (13)), and then the overlap of this state with the sourcing distributions are computed. Both yield exactly the same answer. However, in contrast to the potentially arbitrary WDF $\mathcal{W}_o$, the artefacts of photoelectric detection—the states $\hat{g}$ —are much “better behaved” functions: they are simple Gaussians, bandwidth limited and always non-negative.

4.1 Generalized Rays *

By applying the inverse WDF transform to g [Torre 2005, Chapter 6.3.5], we derive the wave function of the phase-space Gaussian, which we term the *generalized ray* (illustrated in Fig. 3):

Generalized ray (wave function)

$$\psi_{\beta, \rho}(\vec{r}; \vec{r}_o, \vec{k}_o) \triangleq \left(\frac{1}{\pi\beta^2}\right)^{3/4} e^{i\vec{k}_o \cdot (\vec{r} - \vec{r}_o)} e^{-\frac{1}{2\beta^2}(1-i\rho)|\vec{r} - \vec{r}_o|^2}. \quad (15)$$

A general expression with anisotropic β, ρ is available in our supplemental material. The generalized ray is defined with respect to the mean spatial position $\vec{r}_o$, mean wavevector $\vec{k}_o$, and parameters β, ρ . The reader may verify via Eq. (4) that the phase-space representation of the generalized ray (its WDF) is $g_{\beta, \rho}(\vec{r}, \vec{k}; \vec{r}_o, \vec{k}_o)$.

Initially, i.e. given $\rho = 0$, a generalized ray is one of the detection states of the detector $\mathcal{D}$, and its phase-space representation is a coherent state $\hat{g}_{\beta}$ that minimizes the product of its space and wavevector variances. Thus, it is the *most local* construct permissible under wave optics [Torre 2005]. As its name suggests, the generalized ray can be understood as *the closest analogue to the classical ray under wave optics*. Once it is propagated away from the detector ($\rho > 0$), its WDF becomes the correlated Gaussian g , as in Eq. (5).

Sourcing. Let $\mathcal{W}_s$ be a sourcing WDF, as before. Under the backward model of light transport (Eq. (14)), the result of a measurement at a detector element $\{\vec{r}_o, \vec{k}_o, \beta, \alpha\} \in \mathcal{D}$ can be written as

$$I(\vec{r}_o, \vec{k}_o, \beta) \triangleq \int d\vec{r}' d\vec{k}' \mathcal{W}_s(\vec{r}', \vec{k}') \mathcal{K}^t\{\hat{g}_{\beta}\}(\vec{r}', \vec{k}'). \quad (16)$$

The total intensity integrated over the entire detector is then

$$I = \int_{\mathcal{D}} d\mu_{\mathcal{D}} \alpha I(\vec{r}_o, \vec{k}_o, \beta). \quad (17)$$

By Monte Carlo integrating the integral over the detector distribution $\mathcal{D}$ in Eq. (17), we *source* a generalized ray from the detector:

Sourcing

$$\begin{aligned} I &\approx \frac{1}{N} \sum_{n=1}^N \frac{\alpha^{(n)}}{p^{(n)}} I(\vec{r}_o^{(n)}, \vec{k}_o^{(n)}, \beta^{(n)}) \\ &= \frac{1}{N} \sum_{n=1}^N \frac{\alpha^{(n)}}{p^{(n)}} \int d\vec{r}' d\vec{k}' \mathcal{W}_s(\vec{r}', \vec{k}') \mathcal{K}^t\{g^{(n)}\}(\vec{r}', \vec{k}') \end{aligned} \quad (18)$$

(normalization constants accounted for by α), with $p^{(n)}$ being the sampling probabilities. Each sampled $\{\vec{r}_o^{(n)}, \vec{k}_o^{(n)}, \beta^{(n)}, \alpha^{(n)}\} \in \mathcal{D}$ is a detector state that quantifies a sourced generalized ray, whose WDF is denoted as

$$g^{(n)}(\vec{r}, \vec{k}) \triangleq \hat{g}_{\beta^{(n)}}(\vec{r}, \vec{k}; \vec{r}_o^{(n)}, \vec{k}_o^{(n)}) \quad (19)$$

(recall that detection state $\hat{g}_{\beta}$ is simply $g_{\beta, \rho}$ with $\rho = 0$). A sourced generalized ray is defined by its spatial position on the detector $\vec{r}_o^{(n)}$, the wavevector $\vec{k}_o^{(n)}$ that quantifies wavelength and direction of propagation of the generalized ray from the detector, and the variance $\beta^{(n)}$ that quantifies its spatial extent. At sourcing, we set $\rho^{(n)} = 0$. At the detector, α is the detection efficiency, however once a generalized ray is sourced $\alpha^{(n)}$ should be understood as the intensity carried by that generalized ray.




Fig. 3. **Light transport and dynamics in phase space.** A sensor's sensitivity function (gray) and a light source's emission function (orange) are visualized in the wave-optical phase space. In a sense, *generalized rays* (blue) generalize the concept of optical partial coherence to *backward* (sensor-to-source) sampling and transport of light: generalized rays quantify the phase-space extent over which interference may be observed. In Section S7 in our supplemental we derive a formal relation to optical coherence. Observe that free-space propagation constitutes a horizontal shear of the phase space (formalised by Eq. (S6.3)) just as in the ray-optical case.

4.2 Light Transport with Generalized Rays *

For each generalized ray sampled from a detector, viz. Eq. (18), we run a light transport simulation, as formalised by Eq. (14). That simulation is dictated by the time-reversed (total) interaction operator $\mathcal{K}^t$. To compute the contributions of the detection element to the total intensity observed by the detector, we finally integrate the generalized ray over the light sources' sourcing distributions $\mathcal{W}_s$.

Often it is more practical to formulate the interaction operator $\mathcal{K}$, that quantifies the total action on light of the entire scene, as a sequence of interactions, viz. $\mathcal{K} \equiv \mathcal{K}_1 \circ \mathcal{K}_2 \circ \dots \circ \mathcal{K}_M$. A single interaction $\mathcal{K}_j$ might now describe free-space propagation, or a light-matter interaction. The definition of the composed interaction operator $\mathcal{K}_1 \circ \mathcal{K}_m$ only depends on the respective interaction kernels K_l, K_m , and follows immediately from Eq. (1) (see Torre [2005, Chapter 8.3.2]). The time-reversed interaction operator trivially becomes $\mathcal{K}^t \equiv \mathcal{K}_M^t \circ \mathcal{K}_{M-1}^t \circ \dots \circ \mathcal{K}_1^t$.

We now recursively apply each kernel $\mathcal{K}_m^t$ to a generalized ray, in order to Monte Carlo integrate the light transport through the scene. This process is similar to classical light transport: A generalized ray undergoes interactions with the scene at discretized steps. Because a generalized ray is weakly-local, each interaction is confined to a small region in space. A sequence of interactions gives rise to a *path*, enumerated via the indices of the generalized rays chosen at each interaction, starting with n , the index of the sourced generalized ray from Eq. (18). Let a path of depth $m-1$ be denoted by a tuple $\pi = (n_1, \dots, n_m)$, where the indices n_i denote sampled generalized rays at each interaction. We define the addition operation of a path with an index as: $\pi + n_{m+1} = (n_1, \dots, n_m, n_{m+1})$, for some index n_{m+1} .

We denote a generalized ray's phase-space representation as the Gaussian WDF g^π , for an arbitrary path π , as

$$g^\pi(\vec{r}, \vec{k}) \triangleq g_{\beta^\pi, \rho^\pi}(\vec{r}, \vec{k}; \vec{r}_0^\pi, \vec{k}_0^\pi), \quad (20)$$

each defined w.r.t. the values $\beta^\pi, \rho^\pi, \vec{r}_0^\pi, \vec{k}_0^\pi$. Note that Eq. (20) generalizes Eq. (19).

Recursive light transport. Let π be some path of depth $m-1 \geq 0$, i.e. $|\pi| = m$, and g^π the generalized ray sampled using that path. For the initial case $m=1$, $g^\pi = g^{(n)}$ is a sampled sourced generalized ray from the detector, as in Eqs. (18) and (19). Denote the WDF that arises after interaction with $\mathcal{K}_m^t$:

$$\mathcal{Q}_m^\pi \triangleq \mathcal{K}_m^t \{g^\pi\} = \int d\vec{r}_i d\vec{k}_i K_m^t(\vec{r}, \vec{r}_i, \vec{k}, \vec{k}_i) g^\pi(\vec{r}_i, \vec{k}_i). \quad (21)$$

As interaction kernels fulfil all the properties of a WDF [Testorf et al. 2010, Chapter 1.6], the integration above is a phase-space convolution of a WDF with a Gaussian, hence $\mathcal{Q}_m^\pi$ is indeed a Husimi Q distribution (as defined in Eq. (9)). Note, the above holds only because our generalized ray is a phase-space Gaussian (a WDF is a non-negative Husimi Q distribution if, and only if, it can be written as a convolution of an arbitrary WDF with a Gaussian).

Gaussian functions form an overcomplete functional basis [Rudin 1990]. Therefore, a Husimi Q distribution, being a non-negative, bandwidth-limited distribution (a consequence of the convolution), may always be written as a finite sum of Gaussians to arbitrarily high accuracy, viz.

$$\mathcal{Q}_m^\pi(\vec{r}, \vec{k}) \approx \frac{1}{N_m} \sum_{n_m=1}^{N_m} \alpha^{\pi+(n_m)} g^{\pi+(n_m)}(\vec{r}, \vec{k}), \quad (22)$$

where $\alpha^{\pi+(n_m)} > 0$ are the intensities of each Gaussian above (positiveness is a consequence of the non-negativity of a Husimi Q distribution). Thereby, we have sampled N_m generalized rays with paths $\pi + (n_m)$ of depth m , arising from the action of the operator $\mathcal{K}_m$ on the generalized ray g^π .

Continue this process recursively, giving rise to sampled generalized rays g^{π_j} , with π_j being the sampled paths, and each generalized ray is defined by the tuple $\{\vec{r}_0^{\pi_j}, \vec{k}_0^{\pi_j}, \beta^{\pi_j}, \rho^{\pi_j}\}$ and α^{π_j} . The sampled path depth is $|\pi_j| - 1$: For $|\pi_j| = 1$ these generalized rays are sourced from the detector $\mathcal{D}$, while for $|\pi_j| > 1$ these generalized rays are sourced from a distribution that arises after interaction with $\mathcal{K}_{|\pi_j|-1}^t \circ \dots \circ \mathcal{K}_1^t$.

The rendering equation. The recursive *rendering equation for backward wave-optical light transport*, that drives the process above, takes the familiar form of a Fredholm integral equation:

Rendering equation

$$\mathcal{Q}_o \Big|_{\vec{r}_0} = \int d\vec{k}_0 d\beta d\rho \mathcal{K}_r^t \left\{ g_{\beta, \rho}(\vec{r}', \vec{k}'; \vec{r}_0, \vec{k}_0) \right\}, \quad (23)$$

where $\mathcal{K}_r^t$ is some interaction operator (acts w.r.t. the primed variables), g is an incident generalized ray and the integration is over all possible generalized rays' wavevectors $\vec{k}_0$ and parameters β, ρ . In contrast to the classical rendering equation that acts upon scalars, Eq. (23) acts upon functions, all generalized rays that are incident to mean position $\vec{r}_0$, and its result is a function as well, the distribution $\mathcal{Q}_o$. Outgoing generalized rays are then sampled from the resulting light distribution $\mathcal{Q}_o$, in an identical manner to Eq. (22). There is no emission term: interaction with light sources is handled by the measurement operator, introduced next. The interaction operator is discussed further in Section S6 in our supplemental material.

Measurement. The recursive light transport process above solves for the total light transport, viz. the term $\mathcal{K}^{-t}\{g^{(n)}\}$ in Eq. (18). The observed intensity is computed by integrating each generalized ray over the sourcing distribution $\mathcal{W}_s$, quantified by the following *measurement operator*:

Measurement

$$\begin{aligned} \mathfrak{I}\{g\} &\triangleq \int d\vec{r}' d\vec{k}' \mathcal{W}_s(\vec{r}', \vec{k}') g_{\beta, \rho}(\vec{r}', \vec{k}'; \vec{r}_0, \vec{k}_0) \\ &= \frac{1}{(2\pi)^3} \left| \int d\vec{r}' \psi_s(\vec{r}') \psi_{\beta, \rho}^*(\vec{r}'; \vec{r}_0, \vec{k}_0) \right|^2, \end{aligned} \quad (24)$$

where we apply the Moyal formula [Torre 2005, Chapter 6.3.7] in order to reformulate the phase-space integration using Wigner distributions into integration in position space using wave functions, and with ψ_s being the sourced wave function. In practice, the integration above happens once a generalized ray is incident upon a light source, and under the support of that light source's WDF.

4.3 Analysis

In our supplemental material, Section S5, we analyze the formulation presented above and formally show that it is indeed weakly-local, linear and wave-optically complete. We also connect the ability of generalized rays to always remain mutually-linear to sampling theory, and give an intuitive insight into the mutual decoherence induced by generalized rays. Furthermore, in Section S5.2 and Fig. S5 in our supplemental material we also numerically validate our formalism, and show that diffracting scatterers placed at arbitrary locations can always be accurately sampled with generalized rays. In Section S6 in our supplemental material we derive several light-matter interaction BSDFs that we use in this paper.

5 Wave-Optical Rendering

In this section we introduce our rendering algorithm. We start with the accurate algorithm that applies our light transport formalism, as derived in Section 4, exactly. Later, in Section 5.1 we show that with a few reasonable assumptions, our formalism gives rise to an interactive wave-optical rendering algorithm that is easy to implement in a modern path tracer.

Our wave-optical light transport formalism works by propagating backwards, from the detector, generalized rays and simulating their interactions with the scene under time-reversed dynamics, and finally measuring their contributions by integrating over the sourcing distributions of light sources. A generalized ray is a time-reversed Gaussian beam sourced from the detector: Eqs. (5) and (15) define its wave function and WDF, respectively.

Crucially, generalized rays are always **linear** (generalized rays never interfere), and **weakly local**: we limit the spatial extent of the generalized ray to some finite radius around its mean position $\vec{r}_0$. Being Gaussian beams, with rapidly-decaying tails, this can be easily done to arbitrarily high accuracy. For example, a radius of $\varrho = 2.5\beta$, where $\beta^2/2$ is the generalized ray's spatial variance, captures > 99.8% of its weight. Finally, this formalism is **complete**.

In Algorithm 1 we summarise our general-purpose, accurate light transport algorithm:

- **Sourcing.** A generalized ray is fully described by its mean position $\vec{r}_0$, mean wavevector $\vec{k}_0$ (quantifying direction of propagation and the wavelength), spatial variance $\beta^2/2$ and correlation parameter ρ . For each sample, these parameters are sourced via Eq. (18), thereby sourcing a generalized ray initiating a path.
- **Propagation.** We propagate the generalized ray through free space, until it (or a part of it) encounters matter. As we limit the spatial extent of a generalized ray to a region (ball of radius ϱ) around its mean $\vec{r}_0$, tracing it reduces to beam (or cone) tracing through the scene. We formalize the transformation of the generalized ray's parameters, $\{\vec{r}_0, \vec{k}_0, \beta, \rho\}$, under free-space propagation in Eq. (S6.3) in our supplemental material.
- **Measurement.** If the generalized ray (or part of it) encounters a light source, we integrate the generalized ray over the sourcing distribution, as formalised by the measurement operator (Eq. (24)). When no light source is present, $\mathfrak{I} \equiv 0$.
- **Light-matter interaction.** We compute the interaction operator $\mathcal{K}$ (Eq. (11)), that quantifies the action of matter that falls within the spatial extent of the beam (ball of radius ϱ centred at $\vec{r}_0$). This operator acts upon the traced generalized ray, giving rise to the scattered distribution $\mathcal{Q}_o$ (as in Eq. (21)). Then, a diffracted generalized ray is sampled out of this distribution and the light transport simulation continues.

The tuple π is a bookkeeping variable that keeps track of the indices of the generalized rays that were sampled at each step.

The algorithm presented above is a novel wave-optical rendering algorithm. As formulated, it is able to simulate rigorous wave-optical light transport accurately (to arbitrarily good precision) using weakly-local, linear generalized rays. This algorithm is employed

Algorithm 1: Our rendering algorithm.

```

1 foreach detector (pixel)  $\mathcal{D}$  do
2    $I \leftarrow 0$ ; /* will hold the computed detected intensity */
3   for  $n \leftarrow 1$  to  $N$  do /* sample the backward light transport */
4      $\{\vec{r}_0, \vec{k}_0, \beta, \rho\} \leftarrow \text{source generalized ray}$ ; /* Eq. (18) */
5      $\alpha \leftarrow \frac{\alpha^{(n)}}{N_p(n)}$ ; /* path weight; p is sampling probability */
6     while path depth < max depth do
7       propagate the generalized ray  $\{\vec{r}_0, \vec{k}_0, \beta, \rho\}$ ;
8       /* measurement (or solve pass) */
9        $I \leftarrow I + \alpha \cdot \mathfrak{I}\{g_{\beta, \rho}(\vec{r}', \vec{k}'; \vec{r}_0, \vec{k}_0)\}$ ; /* Eq. (24) */
10      /* light-matter interaction */
11       $\mathcal{K} \leftarrow \text{interaction operator around } \vec{r}_0$ ;
12       $\mathcal{Q}_o \leftarrow \mathcal{K}\{g_{\beta, \rho}(\vec{r}', \vec{k}'; \vec{r}_0, \vec{k}_0)\}$ ; /* rendering equation */
13       $\{\vec{r}_0, \vec{k}_0, \beta, \rho\} \leftarrow \text{sample } \mathcal{Q}_o$ ; /* Eq. (22) */
14       $\alpha \leftarrow \alpha \frac{\alpha^{\pi+(n_m)}}{N_m}$ ; /* update weight ( $\alpha^{\pi+(n_m)}$ ),  $N_m$  as in Eq. (22) */

```

for the simulation in Fig. S5 in our supplemental material. A pair of practical difficulties may arise:

- (1) In contrast to classical light transport, propagation cannot be done via ray tracing and requires significantly-slower beam tracing, as generalized rays are not perfectly local (indeed, no wave-optical formalism may be simultaneously perfectly local and linear).
- (2) Because generalized rays occupy a positive extent, the operator $\mathcal{K}$ in the **light-matter interaction** stage may need to account for multiple different materials simultaneously (for example, when a generalized ray is incident upon an edge or a slit, or upon a surface with spatially-varying scattering characteristics). This can be understood as the interaction being a composition of two or more operators, viz. $\mathcal{K} \equiv \mathcal{K}_1 \circ \mathcal{K}_2$. Formally, this can always be done, but, in practice, formalising the composed interaction and sampling generalized rays from it (thereby accounting for the mutual interference of the scattered generalized rays from all materials that are involved in the interaction) can be difficult.

Algorithm 1 describes the basic backward path tracing algorithm with generalized rays. As the rendering equation (Eq. (23)) is linear—by construction, as no interference arises between distinct generalized rays—many typical computer graphics rendering techniques (like Russian Roulette, next event estimation and importance sampling methods [Pharr et al. 2016]) can be applied to our rendering algorithm. We discuss this in Section 5.1.

An efficient implementation of beam tracing, designed for modern GPUs, as well as deriving some composite interactions of interest (e.g., free-space diffraction of a generalized ray around geometry) would be of great interest, and are left for future work. Such work would be especially important for wave-optical simulation at longer wavelengths (e.g., RADAR). In the rest of this paper our focus is on rendering at optical frequencies. We present next our interactive wave-optical rendering implementation, where we make a simplifying application-specific assumption.

5.1 Interactive Rendering

In similar manner to optical rendering with PLT [Steinberg et al. 2022], we make the assumption that the spatial extent ϱ of a generalized ray is smaller or comparable to the scene’s smallest geometric details (we refer to the explicit geometry against which we trace rays, and not analytically defined details like surface properties). Therefore, we may neglect the generalized ray’s spatial extent and use ray tracing to propagate it. Then, the **propagation** stage in our rendering algorithm, Algorithm 1, reduces to simple ray tracing, and in the **light-matter interaction** stage we do not consider composite interactions with multiple materials.

The simplifying assumption above is reasonable for our target application: high-performance rendering at optical frequencies. In terms of optical effects, this assumption means that we neglect free-space diffractions (light “bending” around geometry) and interference across distinct materials. We stress that this assumption is not a limitation of our wave-optical light transport formalism, as presented in Section 3.




Fig. 4. **Manifold sampling.** Example application of an advanced sampling technique. When performing next event estimation (NEE), manifold sampling (MS) [Hanika et al. 2015; Zeltner et al. 2020] enables finding a light path between a surface and a light source across one or more dielectric interfaces. In the rendered scene, such a sampled path, that refracts through the dispersive prism (outlined in blue), is visualized via the dotted blue line. We also employ MS for NEE on specular reflections: note the reflections off the paint brush’s metal handle and off the paint tubes, as well as the thin diffraction grating (outlined in yellow) dispersing light into multiple diffraction lobes. See full high-resolution rendering in our supplemental material. (inset) Colour-coded difference compared to a without-MS rendering. (bottom) The diffraction grating lobes at increasing samples-per-pixel (spp).

Under the assumption above, we may convert a classical unidirectional, backward path tracer into a path tracer compatible with wave optics: generalized rays replace the classical ray as the “point queries” of light’s behaviour. In practice, doing so requires the following changes: (i) transitioning to a spectral, vectorized (polarization-aware) path tracer; (ii) transforming generalized rays on free-space propagation; and (iii) reformulating all the BSDFs in terms of generalized rays.

In addition, in Section S7 in our supplemental material we make a connection between generalized rays and optical coherence, and detail a novel, simple variance reduction technique that we use in our implementation: *sample-solve*.

Implementation details. To **source** a generalized ray from the detector distribution $\mathcal{D}$ of a pixel: the variance $\beta^2/2$ is taken to be a fraction of the pixel spatial size; $\rho = 0$ at sourcing; for simplicity, we set the detection efficiency $\alpha = 1$; the initial mean position $\vec{r}_0$ is sampled uniformly at random on the pixel; and, $-\vec{k}_0$ is the (mean) wavevector. A wavelength λ_0 (such that $|\vec{k}_0| = \frac{2\pi}{\lambda_0}$) is drawn randomly at uniform from the visible spectrum. This wavelength drives the path sampling process, hence we refer to it as the “Hero wavelength” [Wilkie et al. 2014].

The **propagation** of a generalized ray remains unchanged, as described in Section 5. **Light-matter interactions** are done as described




Fig. 5. **Bike scene.** The bike scene from Steinberg et al. [2022] rendered using generalized rays. The scene contains a few materials with visible wave-interference effects, most notably the birefringent dielectrics, (a) like the plastic wheel spoke guard, (b-c) as well as the diffraction grating wheel brake surface. (b) when illuminated by direct sunlight, this grating disperses light into visible diffraction lobes; however, (c) when sunlight passes through a diffuser (like clouds), only the direct lobe is visible. (bottom) Low spp renderings, with (left side) and without (right side) a denoiser, showcasing the interactive rendering performance of our wave-optical renderer.

in Section S6 in our supplemental material, and we are able to importance sample these interactions.

We continue to transform the generalized ray’s parameters, viz. $\{\vec{r}_0, \vec{k}_0, \beta, \rho\}$, on propagation and on interaction with matter. Therefore, we are able to quantify exactly the spatial extent over which a generalized ray interacts with matter. Materials that describe deterministic scattering features, like a surface’s explicit microgeometry, or a multilayered stack, can be accurately rendered. This is not possible with any existing linear wave-optical formalisms: none of them are able to recover weak locality, and would require integration over the entire material.

In addition, when a generalized ray encounters matter, we attempt to connect the path to a light source via next event estimation (NEE). We also perform “Russian roulette” early path termination. We are also able to apply other advanced sampling techniques in our wave-optical rendering algorithm: To demonstrate that we have implemented manifold sampling [Hanika et al. 2015; Zeltner et al. 2020], see Fig. 4.

Under the context of our *sample-solve* algorithm (Section S7 in our supplemental material), the **measurement** stage in Algorithm 1 is replaced with a PLT solve pass. Every time a path has been connected to a light source, either organically or via NEE, the solve stage kicks in to compute the partially-coherent transport over the

sampled path. If the sampled path does not contain dispersive delta segments (e.g., refraction through a smooth dielectric interface), in addition to the Hero wavelength, we also importance sample three additional wavelengths from the light’s emission spectrum. We therefore transport one to four spectral samples per sampled path. The solve stage applies PLT in an essentially unchanged manner, see Steinberg et al. [2022] for more details.

Our implementation supports two types of light sources:

- (1) *Distant light sources* — These include analytic distant sources, with a predefined solid angle that they subtend from the scene, as well as environment maps, in which case light is sourced from a small cluster of a few pixels, defining the solid angle.
- (2) *Emissive geometry light sources* — where light is sourced from a small area on the emitting triangle.

As a design choice, our implementation ignores polarization during the sample pass (i.e., the backward transport with generalized rays). This is done for simplicity and performance: decomposing and inverting Mueller matrices, and other polarimetric interactions, can be cumbersome and expensive. No errors are introduced: the solve pass is fully vectorized, and hence correctly accounts for all polarization effects.




Fig. 6. **Diffraction materials under different illumination conditions.** The appearance of diffractive materials depends on light's spectral, polarimetric and coherence properties. (a) During daytime, the illumination reaching the snake is dominated by indirect sunlight. At some angles, where incident light has a narrower angular spread (and thus is more coherent), the diffractive scales show clear interference patterns. (b) At night, direct light from the fluorescent lamps is the primary source of illumination, however that light is too incoherent to produce visible diffraction patterns.

5.2 Results

Our results are comprised of three main scenes:

- (1) **Snake enclosure** (Fig. 6; Fig. S1 in our supplemental). This scene is illuminated by multiple light sources: the sun, the sky (diffused sunlight), as well as a pair of industrial 4100 K fluorescent lamps with a decent colour rendering index of 82 located at the back of the enclosure. Sunlight and skylight arrive from the opening at the top. This is a difficult scene to render: most of the light that arrives at the different diffractive materials is indirect.
- (2) **Manifold sampling** (Fig. 4). A highly-detailed scene that we use as our manifold-sampling playground.
- (3) **Bike** (Fig. 5). Adapted from Steinberg et al. [2022] (appearance is not expected to match, as our materials are different).

In addition to the above, the CD scene is used for analysis of partially-coherent sampling, Fig. 7, and comparison to the state-of-the-art, Fig. 8. The spaceship (Fig. 1) and Cornell box (Fig. S3 in our supplemental material) are simpler scenes used to demonstrate the crucial role of locality in the reproduction of wave-optical effects. Fig. 6 demonstrates that wave-optics is a global process: the appearance of a material, and the observable diffractive phenomena, depend on the wave properties of light. Accurate reproduction of such effects cannot be done at the material level, but instead *requires a wave-optical simulation throughout the entire scene*: it is the responsibility of the light transport framework to correctly track the propagation of waves, quantify their shapes and spatial extents, and feed that information to BSDFs for integration of light-matter interactions. The diffracting details in these scenes are quantified either statistically (e.g., a diffraction grating), or explicitly (e.g., scratches in a surface). Our supplemental material contains additional renderings, as well as animated videos of several of these scenes, showcasing the performance of our method.

Performance metrics, i.e. rendering resolution and samples-per-pixel (spp) count are given in each figure, and summarised in Table 1. For the figures in the paper we used high-quality, converged (i.e., very high samples-per-pixel) renders. Nevertheless, our method and implementation enable interactive wave-optical rendering at 1 spp, and the frame timings for interactive rendering are also given in the figures and in Table 1. Similar to other modern GPU-accelerated path tracers, we use a denoiser for interactive rendering. This enables generating acceptable images at 1 spp, and allows the user to interact with and edit the scene in real-time. The videos in our supplemental material were also rendered using a denoiser.

Low-spp images of the bike scene, with and without the denoiser, are shown in Fig. 5 (bottom). Low sample count is sufficient for the vast majority of the materials, including the diffractive birefringent dielectrics. An exception are the dispersive diffraction lobes, which are visible on the floor, and arise due to the diffraction-grated wheel brake surface. This is due to a couple of reasons: (i) A diffraction grating scatters into many wavelength-dependent lobes, hence requires a moderate amount of spectral samples. (ii) These are rendered via manifold sampling (MS), as discussed in Fig. 4. However, MS is only initiated during the sample stage when a path is scattered from the floor into the grating. The probability of finding such

Table 1. **Rendering performance.** Listed are the resolution and sample count used to generate the figures, as well as the interactive (1 spp) rendering frame times (at the indicated resolution). All rendering was done on a NVIDIA® GeForce RTX™ 3090 GPU.

Scene	Frame time (1 spp)	Figure	
Snake enclosure	116 ms	Figs. 1 and 6	1440p, 64k spp
MS playground	180 ms	Fig. 4	1440p, 16k spp
Bike	240 ms	Fig. 5	1440p, 4k spp
CD	42 ms	Fig. 7	1080p, 2M spp




Fig. 7. **Partially-coherent sampling.** (a) A scene contains a compact disk (CD) that rests next an open CD case, upon which another closed CD case is placed. A ceiling-mounted light source illuminates the scene. While simple, the scene admits interesting light transport. (b) We emulate partially-coherent (PC) sampling of BSDFs, in an identical manner to the state-of-the-art (PLT), which manifests the sampling problem: the diffraction lobes are very sharp lobes that are difficult to sample when doing backward path tracing (and the coherence properties of light are unknown) using existing tools. The close ups show the area marked in orange rendered at various sample counts. Note the sharp difference in noise between PC sampling and our proposed light transport formalism with generalized rays. (c) Plot of noise as function of sample count, quantifying the drastic improvement in sampling performance: sampling using generalized rays reduces sample count by a factor of about 4000 for similar-quality rendering. (a, inset) A difference image between the two sampling strategies. With the exception of the very high-frequency diffraction effects, the differences are minor, and are due to the PC rendering never converging.




Fig. 8. **Comparison to the state-of-the-art (PLT).** We render the CD scene using the PLT bidirectional path tracer (BDPT) [Steinberg et al. 2022]. Because the illumination of the diffractive CD surface is indirect, convergence is poor, as to be expected given the analysis in Fig. 7. The image was rendered with 56 000 samples over about 315 hours, on an Intel® Core™ i9-10980XE 18-core CPU. (a-c) Their renderer struggles with capturing the high-frequency details of the diffraction grating. Furthermore, because they propagate a fixed number (64) of spectral samples, sampled uniformly, clear banding artefacts are reproduced, suggesting that these high-frequency details will fail to converge due to spectral aliasing. (right) To analyze and compare the convergence performance to our approach, we plot error per sample and error per rendering-time graphs for the three different regions outlined in red: (d) The light transport that arrives to this area (red dotted line) is dominated by light that is not diffracted by the CD. Therefore, the per-sample convergence performance of their BDPT renderer (which performs much more work per sample) is significantly superior, as expected. (e-f) On the other hand, these regions do sample the CD and the convergence performance of their renderer is considerably inferior. A performance improvement is expected, as our renderer is GPU accelerated. The performance at region (d), plotted with open circle markers, should then be considered as baseline: up to 100 times faster, depending on sample count. The plots pertaining to the diffractive regions (filled markers) show a vastly greater improvement of about 1000 to 10 000 times faster convergence.

a connection organically is rather low (roughly about 1 in a few dozen), as the grated surface is quite small. The difficulty of rendering the diffracted lobes of a diffraction grating is also analysed in Fig. 4 (bottom). This problem is similar to the classical problem of rendering dispersive caustics, e.g., on the bottom of a pool with a disturbed water surface; a problem where classical unidirectional path tracers struggle as well.

Comparison with the state-of-the-art. We compare our interactive wave-optical renderer to PLT [Steinberg et al. 2022]. PLT employs a bi-directional path tracer, and samples partially-coherent BSDFs. As discussed, because the coherence of light cannot be known a priori when tracing paths backwards (i.e., the sampling problem), they set a global lower limit on the coherence of light. Partially-coherent BSDFs are then sampled with respect to that limit. We emulate such partially-coherent sampling in our renderer, and compare the performance in the simple CD scene, see Fig. 7. The difference is dramatic: we observe over a thousand-fold increase in convergence performance when sampling with generalized rays compared to partially-coherent sampling.

We also compare by rendering the same CD scene directly with PLT, see Fig. 8. An increase in performance is expected, as our implementation is GPU accelerated, nevertheless, we record up to a several thousand speed gain for a similar-quality rendering (for image areas that admit diffractive materials). Our equal-sample convergence performance is increased by a factor of 1 to 8. This is despite the fact that their path tracer is bi-directional (ours is unidirectional) and each sample propagates 64 uniformly sampled spectral samples (compared to the up to 4 per sample with our renderer), therefore they do a considerably greater amount of work per sample. Figs. 7 and 8 highlight the sampling problem and emphasize the need for better solutions: the state-of-the-art is able to achieve decent results when diffractive materials are directly illuminated, making good use of a bi-directional path tracer. However, it struggles greatly when it is not trivial to connect paths from a light source to these materials.

6 Conclusion

The first major contribution in this work is our backward (sensor-to-source) wave-optical light transport formalism. The derivations in Section 3 are exact, subject to a Wigner distribution treatment of optics and photoelectric detection, and our formalism is *complete*, subject to photoelectric detectors: any wave-optical phenomena that may be observed by a photoelectric detector can be simulated to arbitrarily high accuracy, with light of optical and non-optical frequencies, and of any state of coherence or polarization. At the core of our formalism is the generalized ray: a *weakly-local* and *linear* construct that samples the distribution of light starting from the detector. In Section S5.2 in our supplemental material we numerically validate this formalism, as well as delineate the highly-general process by which generalized rays are able to regain linearity in all conditions: only interfering phasors that can be resolved over the spatial extent of the detection element, without aliasing, contribute to observable interference effects.

Some practical considerations are left for future work: (i) efficient beam tracing; and, (ii) sampling the interaction of a generalized ray

with a composite interaction operator, like free-space diffraction through a slit. These considerations are especially important for simulations with longer, non-optical wavelengths, where the spatial extent of a generalized ray is significantly larger than at optical frequencies.

We show the application of our theory to wave-optical rendering in the visible spectrum, introducing an application-specific simplification: like PLT, we assume that the spatial extent of the generalized ray is small compared to the scene's geometric details, which allows the use of ray tracing to construct light paths. This leads to our second major contribution: making it possible for the first time to efficiently render complex scenes under an accurate wave-optical model. Our novel sample-solve approach makes it possible to apply decades of work in path tracing sampling algorithms to wave optics light transport, allowing the construction of paths from the camera and the use of advanced sampling techniques like manifold next event estimation. Our algorithms are efficient and amenable to GPU implementation; we show interactive one sample per pixel rendering of complex scenes on a modern GPU, at a performance that is orders-of-magnitude faster than the state-of-the-art. Conventional ray optics path tracers can be extended to wave-optical light transport using our algorithms via a small number of changes and our implementation is available with our supplemental material; we hope that this will allow wave-optical light transport to become widely used as the basis for rendering algorithms in place of ray optics.

References

- Birk Andreas, Giovanni Mana, and Carlo Palmisano. 2015. Vectorial ray-based diffraction integral. *Journal of the Optical Society of America A* 32, 8 (July 2015), 1403. <https://doi.org/10.1364/josaa.32.001403>
- Laurent Belcour and Pascal Barla. 2017. A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence. *ACM Trans. Graph.* 36, 4, Article 65 (July 2017), 14 pages. <https://doi.org/10.1145/3072959.3073620>
- Ahmad Bilal, Syed Muhammad Hamza, Ziauddin Taj, and Shuaib Salamat. 2019. Comparison of SBR and MLFMM techniques for the computation of RCS of a fighter aircraft. *IET Radar, Sonar & Navigation* 13, 10 (2019), 1805–1810.
- D Bilibashi, EM Vitucci, and V Degli-Esposti. 2020. Dynamic ray tracing: Introduction and concept. In *2020 14th European Conference on Antennas and Propagation (EuCAP)*. IEEE, 1–5.
- Mate Boban, Joao Barros, and Ozan K Tonguz. 2014. Geometry-based vehicle-to-vehicle channel modeling for large-scale simulation. *IEEE Transactions on Vehicular Technology* 63, 9 (2014), 4146–4164.
- John Burgess. 2020. RTX on—The NVIDIA Turing GPU. *IEEE Micro* 40, 2 (March 2020), 36–44.
- Juan D. Castro, Sahitya Singh, Akshaj Arora, Sara Louie, and Damir Senic. 2019. Enabling Safe Autonomous Vehicles by Advanced mm-Wave Radar Simulations. In *2019 IEEE MTT-S International Microwave Symposium (IMS)*. IEEE. <https://doi.org/10.1109/mwmsym.2019.8700843>
- Hyuckjin Choi, Jaeky Oh, Jaehoon Chung, George C Alexandropoulos, and Junil Choi. 2023. WiThRay: A Versatile Ray-Tracing Simulator for Smart Wireless Environments. *IEEE Access* (2023).
- Per H. Christensen. 2003. Adjoints and importance in rendering: an overview. *IEEE Transactions on Visualization and Computer Graphics* 9, 3 (2003), 329–340. <https://doi.org/10.1109/TVCG.2003.1207441>
- Tom Cuyppers, Tom Haber, Philippe Bekaert, Se Baek Oh, and Ramesh Raskar. 2012. Reflectance model for diffraction. *ACM Transactions on Graphics* 31, 5 (Aug 2012), 1–11. <https://doi.org/10.1145/2231816.2231820>
- FS de Adana, O Gutierrez, I Gonzalez, and MF Catedra. 2005. FASPRO: Fast computer tool for the analysis of propagation in mobile communications. In *INDIN'05. 2005 3rd IEEE International Conference on Industrial Informatics*, 2005. IEEE, 257–261.
- Daniela Dragoman. 2004. Phase space correspondence between classical optics and quantum mechanics. In *Progress In Optics*. Vol. 42. 424–486.
- Daniela Dragoman. 2005. Phase space formulation of quantum mechanics. Insight into the measurement problem. *Physica Scripta* 72, 4 (2005), 290.

- V. Falster, A. Jarabo, and J. R. Frisvad. 2020. Computing the Bidirectional Scattering of a Microstructure Using Scalar Diffraction Theory and Path Tracing. *Computer Graphics Forum* 39, 7 (Oct 2020), 231–242. <https://doi.org/10.1111/cgf.14140>
- Tiantian Feng and Lixin Guo. 2021. Multiview ISAR Imaging for Complex Targets Based on Improved SBR Scattering Model. (2021).
- Pengcheng Gao, Xiaobing Wang, Zichang Liang, and Wei Gao. 2015. Efficient GPU implementation of SBR for fast computation of composite scattering from electrically large target over a randomly rough surface. In *2015 IEEE International Symposium on Antennas and Propagation ; USNC/URSI National Radio Science Meeting*. IEEE. <https://doi.org/10.1109/aps.2015.7305222>
- Ion I. Geru. 2018. *Time-Reversal Symmetry*. Springer International Publishing. <https://doi.org/10.1007/978-3-030-01210-6>
- Junfeng Guan, Sohrab Madani, Suraj Jog, Saurabh Gupta, and Haitham Hassanieh. 2020. Through fog high-resolution imaging using millimeter wave radar. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11464–11473.
- Ibón Guillén, Julio Marco, Diego Gutierrez, Wenzel Jakob, and Adrian Jarabo. 2020. A General Framework for Pearlescent Materials. *ACM Transactions on Graphics* 39, 6 (2020). <https://doi.org/10.1145/3414685.3417782>
- Johannes Hanika, Marc Droske, and Luca Fascione. 2015. Manifold next event estimation. *Comput. Graph. Forum* 34, 4 (July 2015), 87–97.
- 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>
- Ferdous Hossain, Tan Kim Geok, Tharek Abd Rahman, Mhd Nour Hindia, Kaharudin Dimiyati, and Azlan Abdaziz. 2018. Indoor millimeter-wave propagation prediction by measurement and ray tracing simulation at 38 GHz. *Symmetry* 10, 10 (2018), 464.
- Weizhen Huang, Julian Iseringhausen, Tom Kneiphof, Ziyin Qu, Chenfanfu Jiang, and Matthias B. Hullin. 2020. Chemomechanical simulation of soap film flow on spherical bubbles. *ACM Transactions on Graphics* 39, 4 (Jul 2020). <https://doi.org/10.1145/3386569.3392094>
- KL Jensen and FA Buot. 1991. The methodology of simulating particle trajectories through tunneling structures using a Wigner distribution approach. *IEEE Transactions on Electron Devices* 38, 10 (1991), 2337–2347.
- Jian-Ming Jin. 2015. *The finite element method in electromagnetics*. John Wiley & Sons.
- K.-S. Jin, T.-I. Suh, S.-H. Suk, B.-C. Kim, and H.-T. Kim. 2006. Fast Ray Tracing Using A Space-Division Algorithm for RCS Prediction. *Journal of Electromagnetic Waves and Applications* 20, 1 (Jan. 2006), 119–126. <https://doi.org/10.1163/156939306775777341>
- James T. Kajiya. 1986. The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques - SIGGRAPH '86*. ACM Press. <https://doi.org/10.1145/15922.15902>
- Cam Key, Blake Troksa, Forest Kunkel, Slobodan V. Savic, Milan M. Ilic, and Branislav M. Notaros. 2018. Comparison of Three Sampling Methods for Shooting-Bouncing Ray Tracing Using a simple Waveguide Model. In *2018 IEEE International Symposium on Antennas and Propagation ; USNC/URSI National Radio Science Meeting*. IEEE. <https://doi.org/10.1109/apusncursinrm.2018.8609320>
- Tom Kneiphof, Tim Golla, and Reinhard Klein. 2019. Real-time Image-based Lighting of Microfacet BRDFs with Varying Iridescence. *Computer Graphics Forum* 38, 4 (2019), 77–85. <https://doi.org/10.1111/cgf.13772>
- 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.
- Jacob Mackay and David Johnson. 2021. Millimetre wave ray tracing simulator with phase and beam effects using the Wigner distribution function. In *Passive and Active Millimeter-Wave Imaging XXIV*, Vol. 11745. SPIE, 43–57.
- Leonard Mandel and Emil Wolf. 1995. *Optical coherence and quantum optics*. Cambridge University Press, Cambridge.
- Marco Mout, Michael Wick, Florian Bociort, Joerg Petschulat, and Paul Urbach. 2018. Ray tracing the Wigner distribution function for optical simulations. *Optical Engineering* 57, 1 (2018), 014106–014106.
- ZY Ou and HJ Kimble. 1995. Probability distribution of photoelectric currents in photodetection processes and its connection to the measurement of a quantum state. *Physical Review A* 52, 4 (1995), 3126.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically based rendering: From theory to implementation*. Morgan Kaufmann.
- Walter Rudin. 1990. *Functional Analysis* (2 ed.). McGraw Hill Higher Education, Maidenhead, England.
- Hae-Won Son and Noh-Hoon Myung. 1999. A deterministic ray tube method for microcellular wave propagation prediction model. *IEEE Transactions on Antennas and Propagation* 47, 8 (1999), 1344–1350.
- Francisco Soto and Pierre Claverie. 1983. When is the Wigner function of multidimensional systems nonnegative? *J. Math. Phys.* 24, 1 (1983), 97–100.
- Jos Stam. 1999. Diffraction shaders. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques - SIGGRAPH '99*. ACM Press. <https://doi.org/10.1145/311535.311546>
- Shlomi Steinberg, 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>
- Shlomi Steinberg and Ling-Qi Yan. 2021. A Generic Framework for Physical Light Transport. *ACM Transactions on Graphics* 40, 4 (Aug 2021), 1–20. <https://doi.org/10.1145/3450626.3459791>
- Shlomi Steinberg and Ling-Qi Yan. 2022. Rendering of Subjective Speckle Formed by Rough Statistical Surfaces. *ACM Trans. Graph.* 41, 1, Article 2 (feb 2022), 23 pages. <https://doi.org/10.1145/3472293>
- Sungha Suk, Tae-Il Seo, Hae-Sung Park, and Hyo-Tae Kim. 2001. Multiresolution grid algorithm in the SBR and its application to the RCS calculation. *Microwave and Optical Technology Letters* 29, 6 (2001), 394–397. <https://doi.org/10.1002/mop.1188>
- Yubo Tao, Hai Lin, and Hujun Bao. 2010. GPU-Based Shooting and Bouncing Ray Method for Fast RCS Prediction. *IEEE Transactions on Antennas and Propagation* 58, 2 (Feb. 2010), 494–502. <https://doi.org/10.1109/tap.2009.2037694>
- Yu-Bo Tao, Hai Lin, and Hu Jun Bao. 2008. Kd-tree based fast ray tracing for rcs prediction. *Electromagn. Waves (Camb.)* 81 (2008), 329–341. <https://doi.org/10.2528/pier08011305>
- Markus Testorf, Bryan Hennelly, and Jorge Ojeda-Castañeda. 2010. *Phase-space optics: fundamentals and applications*. McGraw-Hill Education.
- Antoine Toisoul and Abhijeet Ghosh. 2017. Practical Acquisition and Rendering of Diffraction Effects in Surface Reflectance. *ACM Transactions on Graphics* 36, 5 (Jul 2017), 1–16. <https://doi.org/10.1145/3012001>
- Amalia Torre. 2005. *Linear ray and wave optics in phase space: bridging ray and wave optics via the Wigner phase-space picture*. Elsevier.
- Z. Velinov, S. Werner, and M. B. Hullin. 2018. Real-Time Rendering of Wave-Optical Effects on Scratched Surfaces. *Computer Graphics Forum* 37, 2 (2018), 123–134. <https://doi.org/10.1111/cgf.13347>
- A Walther. 1968. Radiometry and coherence. *JOSA* 58, 9 (1968), 1256–1259. <https://doi.org/10.1364/JOSA.58.001256>
- Craig Warren, Antonios Giannopoulos, and Iraklis Giannakis. 2016. gprMax: Open source software to simulate electromagnetic wave propagation for Ground Penetrating Radar. 209 (Dec 2016).
- 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>
- A Wilkie, S Nawaz, M Droske, A Weidlich, and J Hanika. 2014. Hero wavelength spectral sampling. *Comput. Graph. Forum* 33, 4 (July 2014), 123–131.
- Emil Wolf. 2007. *Introduction to the theory of coherence and polarization of light*. Cambridge University Press, Cambridge.
- Kane Yee. 1966. Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media. *IEEE Transactions on antennas and propagation* 14, 3 (1966), 302–307.
- Haofan Yi, Danping He, P Takis Mathiopoulos, Bo Ai, Juan Moreno Garcia-Loygorri, Jianwu Dou, and Zhangdui Zhong. 2022. Ray tracing meets terahertz: Challenges and opportunities. *IEEE Communications Magazine* (2022).
- Yunchen Yu, Mengqi Xia, Bruce Walter, Eric Michielssen, and Steve Marschner. 2023. A Full-Wave Reference Simulator for Computing Surface Reflectance. , 17 pages. <https://doi.org/10.1145/3592414>
- Tizian Zeltner, Iliyan Georgiev, and Wenzel Jakob. 2020. Specular manifold sampling for rendering high-frequency caustics and glints. *ACM Trans. Graph.* 39, 4 (Aug. 2020).