



Contents lists available at ScienceDirect

Journal of Quantitative Spectroscopy and Radiative Transfer

journal homepage: www.elsevier.com/locate/jqsrt

Unifying radiative transfer models in computer graphics and remote sensing, Part I: A survey

Katherine Salesin^{a,*}, Kirk D. Knobelspiesse^b, Jacek Chowdhary^c, Peng-Wang Zhai^d, Wojciech Jarosz^a

^a Department of Computer Science, Dartmouth College, 15 Thayer Dr, Hanover, NH, 03755, USA

^b Ocean Ecology Laboratory, GSFC/NASA, Code 616, Greenbelt, MD, 20771, USA

^c Department of Applied Physics and Applied Mathematics, Columbia University, 2880 Broadway, New York, NY, 10025, USA

^d Department of Physics, University of Maryland Baltimore County, 1000 Hilltop Circle, Baltimore, MD, 21250, USA

ARTICLE INFO

Keywords:

Radiative transfer
Computer graphics
Remote sensing
Atmospheric and ocean optics
Polarization
Monte Carlo

ABSTRACT

The constellation of Earth-observing satellites continuously collects measurements of scattered radiance, which must be transformed into geophysical parameters in order to answer fundamental scientific questions about the Earth. Retrieval of these parameters requires highly flexible, accurate, and fast forward and inverse radiative transfer models. Existing forward models used by the remote sensing community are typically accurate and fast, but sacrifice flexibility by assuming the atmosphere or ocean is composed of plane-parallel layers. Monte Carlo forward models can handle more complex scenarios such as 3D spatial heterogeneity, but are relatively slower. We propose looking to the computer graphics community for inspiration to improve the statistical efficiency of Monte Carlo forward models and explore new approaches to inverse models for remote sensing. In Part 1 of this work, we examine the evolution of radiative transfer models in computer graphics and highlight recent advancements that have the potential to push forward models in remote sensing beyond their current periphery of realism.

1. Introduction

Light collected by sensors observing the Earth's atmosphere and oceans can provide enormous insight into the chemical, physical, and biological processes occurring within. In order to dissect the contributions of such processes to sensor measurements, accurate simulations of how light interacts with real-world phenomena must be developed. To illustrate the immensity of this task, consider that fully accurate simulations of atmosphere–ocean systems must account for scattering and absorption by gases, water droplets, and particulates (aerosols) in the atmosphere; scattering and absorption by gases, liquids, particulates, and organisms in the ocean; scattering at the ocean surface, which may be sprinkled with white caps or sea ice; as well as fluorescence and phosphorescence caused by biological processes. Simulations must also match the capabilities of the sensor to capture angular, spectral, and polarized properties of incoming light. Lastly, these simulations must be computationally efficient — although there is some leeway to sacrifice speed for complexity or vice versa.

In order to be useful to the scientific community, sensor measurements of light must be transformed via simulation into atmospheric

or oceanic parameters such as ocean surface roughness; chlorophyll-*a* concentration in the ocean; phase, height, optical thickness, and droplet radius of clouds; and shape, refractive index, optical depth, and vertical distribution of aerosols [1]. The retrieval of all parameters at once for a snapshot of place and time would be ideal, as it would capture their optically interrelated nature. However, this goal is not currently achievable due to both insufficient information per “snapshot” and the limitations of existing retrieval algorithms. More information can be gathered by developing more accurate or more advanced hardware, e.g. sampling more wavelengths in the visible, ultraviolet, and infrared range, sampling multiple viewing angles of the Earth, or detecting the polarization state and angle of incoming light. Many of these enhancements will be added to the instruments on NASA's upcoming Plankton, Aerosol, Cloud, Ocean Ecosystem (PACE) mission [2]: the Ocean Color Instrument (OCI) [3], a hyperspectral radiometer, and the Spectro-Polarimeter for Planetary Exploration (SPEXone) [4,5] and Hyper Angular Rainbow Polarimeter (HARP2) [5,6], multi-angle polarimeters of different angular and spectral resolutions. In order to harness the unprecedented volume and richness of incoming data to

* Corresponding author.

E-mail addresses: katherine.a.salesin.gr@dartmouth.edu (K. Salesin), kirk.d.knobelspiesse@nasa.gov (K.D. Knobelspiesse), jacek.chowdhary@nasa.gov (J. Chowdhary), pwzhai@umbc.edu (P. Zhai), wojciech.k.jarosz@dartmouth.edu (W. Jarosz).

<https://doi.org/10.1016/j.jqsrt.2023.108847>

Received 24 February 2023; Received in revised form 17 November 2023; Accepted 19 November 2023

Available online 29 November 2023

0022-4073/© 2023 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

its full potential, more advanced remote sensing algorithms must be developed.

Toward this goal, we look to a fresh source of inspiration: the field of *computer graphics*. Since the dawn of the computer era, the field of computer graphics has studied how to recreate what we see virtually, building upon the prior insights of artists and scientists over hundreds of years. In pursuit of “photorealism”, or computer-generated images indistinguishable from a photograph, computer graphics has evolved toward extremely realistic, physically based models. While applications in the past have typically been commercial (e.g. visual effects for movies and video games, artistic creation tools, virtual reality), the underlying Monte-Carlo-based radiative transfer models have finally reached the threshold of complexity and accuracy required for scientific applications. However, as of now, there is relatively little cross-pollination between radiative transfer literature in computer graphics and other scientific fields such as optics, astronomy, neutron transport, atmospheric science, oceanography, and remote sensing.

In Part 1 of this work, we create a Rosetta Stone for two scientific communities that have been working in parallel for decades. We review the radiative transfer models used in remote sensing and current capabilities of their existing frameworks in Section 2, and likewise for computer graphics in Section 3. While both communities use Monte Carlo models for radiative transfer, they are more heavily favored by the computer graphics community, who has spent decades of research improving their statistical efficiency via sophisticated sampling techniques — we highlight areas that are pertinent to remote sensing in Section 4.

In Part 2 of this work, we extend a forward and inverse modeling framework recently developed in the computer graphics community, Mitsuba 3 [7], to perform simulations of interest to remote sensing and test our framework on recent benchmark tests of simple atmosphere–ocean systems.

Through both the literature survey in Part 1 and the benchmark tests in Part 2, we demonstrate that Monte Carlo models in computer graphics can not only *match the accuracy* of those in remote sensing, but have yet-unexplored potential to model even *greater complexity* of real-world environments and *improve efficiency*.

2. Review of forward and inverse radiative transfer models

To prime our survey of relevant computer graphics literature, we first review the priorities of each community with respect to radiative transfer models and their current capabilities. A *radiative transfer model* is a mathematical and algorithmic framework for conducting realistic simulations of light. Both forward and inverse radiative transfer models are necessary in order to transform sensor measurements into other parameters of interest (see Fig. 1). A *forward* radiative transfer model maps *state space* to *measurement space*, mimicking the same transformation as a real-world sensor. State space encompasses all geophysical parameters such as scattering and absorption properties as well as scene composition parameters such as sensor angle, sun angle, geometry, etc., while measurement space encompasses all possible measurements of radiance that the sensor could produce, e.g. the space of all photos a camera can capture.

Naturally, an *inverse* radiative transfer model maps measurement space to state space, the desired transformation for remote sensing retrievals. Inverse models typically rely on embedded forward models to produce intermediate results for optimization.

Retrieval of atmosphere and ocean parameters requires forward and inverse radiative transfer models that are flexible, accurate, and fast. *Flexible* models are able to represent realistic light-based phenomena of arbitrary complexity, while *accurate* models are able to estimate radiance precisely under the assumptions of the underlying model. *Speed* is also important in order to deliver data of interest promptly to the scientific and applications community for further analysis. Typically, there exists a tradeoff between these desired qualities: favoring one requires sacrificing another to some extent.

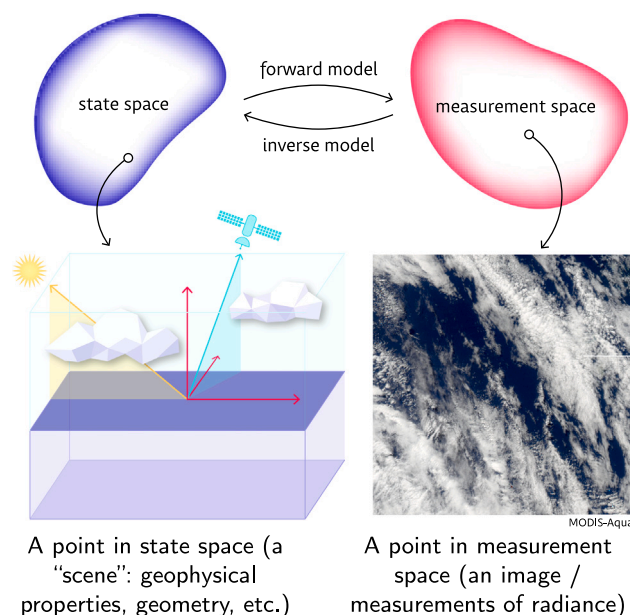


Fig. 1. A forward model maps state space, the space of all possible geophysical parameters and geometry, to measurement space, the space of all possible measurements of radiance, mimicking the functionality of a real-world sensor. An inverse model maps measurement space to state space, producing useful data for further analysis by the scientific community.

Source: Image taken by MODIS-Aqua of the South Pacific near Tonga (2006).

2.1. Review of forward radiative transfer models

The remote sensing community typically favors accurate and fast forward radiative transfer models that use *deterministic* solutions to a sub-problem of the classical radiative transfer equation (RTE), such as adding-doubling [8], successive orders of scattering [9,10], or the discrete-ordinate method [11]. We refer the reader to Chowdhary et al. [12] and references therein for a recent and comprehensive survey of such models available for remote sensing.

The accuracy of these models is controlled by the number of quadrature points used to numerically estimate the solutions of lower-dimensional integrals. The lower dimensionality of the new integral compared to the original RTE is critical because numerical integration suffers from the *curse of dimensionality*, in which the convergence rate of the estimate becomes exponentially slower as the dimension of the integral increases. While these methods may be able to handle complex effects such as polarization or fluorescence, they sacrifice flexibility by assuming the atmosphere or ocean is composed of *plane-parallel layers*, a stack of layers bounded by parallel planes whose properties are homogeneous within each layer. However, the assumption of plane-parallel layers breaks down both in the small scale (e.g. clouds are fluffy and organically shaped) and the large scale (e.g. the Earth is spherical).

An alternate solution to the RTE that is more general and does not suffer from the curse of dimensionality is a stochastic solution called *Monte Carlo integration*, which is heavily favored by the computer graphics community today. Monte Carlo models are able to handle more geometric complexity and spatial heterogeneity. One remaining assumption is geometric optics, in which light is modeled by rays traveling along straight paths that are diverted by scatter events. Geometric optics is compatible with some wave-based effects of light (e.g. polarization) but not others (e.g. diffraction) since it does not account for electromagnetic phase, although some phase-aware effects can be accounted for locally (e.g. Lorenz-Mie scattering [13,14]). Recent research in computer graphics has begun deconstructing this assumption to faithfully represent the wave nature of light in existing Monte Carlo frameworks [15–17].

Monte Carlo models are generally viewed as too computationally expensive for operational processing of satellite datasets, although some have been developed in remote sensing primarily for offline generation of data for look-up tables (see Chowdhary et al. [12] for a summary). Notably, some of the earliest atmosphere–ocean radiative transfer models developed in the late 1960s/early 1970s by Kattawar and Plass [18–21] were Monte-Carlo-based. However, Monte Carlo methods have progressed significantly in both accuracy and speed in the past two decades in computer graphics [22], partly due to the meteoric rise in computing power and partly due to research in physically based rendering (“to render” is to run a forward radiative transfer model). Some developments from computer graphics have just very recently trickled into Monte Carlo models in the remote sensing community, namely parallelization on the GPU and light/path tracing with next-event estimation [23–31] (notably, [27,30] use an earlier version of Mitsuba for forward modeling of clouds). However, they have not delved far into the huge library of variance reduction methods developed for computer graphics, like those we will highlight in Section 4. To prime further discussion, let us briefly review the basics of Monte Carlo. Monte Carlo integration approximates the integral of a function, $F := \int_D f(x) dx$, by averaging the function f evaluated at a number of randomly sampled locations x_i over the domain D :

$$\langle F \rangle := \frac{1}{N} \sum_{i=1}^N \frac{f(x_i)}{p(x_i)}. \quad (1)$$

As long as the probability density (PDF) of the samples $p(x) > 0$ for any location within the domain D where $f(x) \neq 0$, then the estimator $\langle F \rangle$ is *unbiased*, i.e. its expected value $\mathbb{E}[\langle F \rangle] = F$, but the estimator will still exhibit error due to *noise*, or random fluctuations. We can tame these fluctuations by averaging more samples, which should bring the estimate closer to the true solution. We can quantify error due to noise by the estimator’s *variance*, which expresses how far each sample is from the expected value, on average:

$$\mathbb{V}[\langle F \rangle] = \mathbb{E} \left[\left(\langle F \rangle - \mathbb{E}[\langle F \rangle] \right)^2 \right]. \quad (2)$$

Plugging Eq. (1) into Eq. (2) gives:

$$\mathbb{V}[\langle F \rangle] = \frac{1}{N^2} \sum_{i=1}^N \mathbb{V} \left[\frac{f(x_i)}{p(x_i)} \right] = \frac{1}{N} \mathbb{V} \left[\frac{f(x)}{p(x)} \right], \quad (3)$$

where the first step assumes the samples x_i are independent, and the second step assumes they are all drawn from the same PDF. This leads to the key observation that the variance (squared error) of a Monte Carlo estimator depends on the number of averaged samples N and the variance of the *quotient* $f(x)/p(x)$. We can reduce variance by increasing N or by choosing a $p(x)$ that more closely matches the shape of $f(x)$ across its domain. Ideally, when $p(x)$ is perfectly proportional to $f(x)$, $p(x) \sim f(x)$, the variance $\mathbb{V} \left[\frac{f(x_i)}{p(x_i)} \right]$ is zero, and so is that of the estimator.

Monte Carlo radiative transport codes approximate the integrals within the RTE using estimators like Eq. (1). In this setting, finding a neatly proportional $p(x)$ is quite hard since the RTE is the composition of many nested and recursive functions (see Part 2, Eqs. (3) and (5)). For example, the amount of light arriving at a point in space is the sum of many paths scattered throughout the scene and influenced by the materials and geometry of everything they have touched! Crafting a probability distribution well-suited to its integrand is called *importance sampling*, and is a key research concept in computer graphics. Many examples are given in Section 4.

2.2. Review of inverse radiative transfer models

For a comprehensive historical perspective on inverse radiative transfer models for remote sensing, we refer the reader to Chowdhary et al. [12], Frouin et al. [32] and Remer et al. [3,5] (for aerosol retrievals in particular). Available solutions fall roughly into four classes:

- *Heuristic models* use simple formulas to retrieve a single parameter based on its known properties, e.g. relating chlorophyll-*a* concentration to radiance measurements at its absorption peak of 443 nm [33].
- *Look-up tables* store the output of a forward model for every permutation of a densely sampled set of parameters.
- *Optimization-based models* optimize guesses of parameters of interest using a forward model to generate intermediate results. These typically use the gradients (sometimes referred to as Jacobians) of the forward model in order to update parameters at each iteration.

Optimization-based models are typically the most accurate, but also the most expensive. They use a loss function to quantify the difference between the forward model and some reference value, typically the radiance measured by a sensor. Gradients should always be used in the optimization when available: the convergence rate of optimizations with gradients are better than without (e.g. finite differences) on noisy objective functions [34]. However, gradients are not always available if the forward model is not differentiable, or computing them by automatic methods (see *differentiable rendering* in Section 4) is not possible due to computing resource constraints.

State-of-the-art inverse models for remote sensing target the retrieval of several atmosphere–ocean parameters simultaneously. For example, the MAPOL retrieval algorithm [35] uses a non-linear least squares optimization to retrieve aerosol (size, complex refractive index) and ocean (wind speed, chlorophyll-*a*, absorption and backscattering coefficients of plankton and detritus) parameters. Successor FastMAPOL [36] replaced the optimization backbone with a neural network. Retrieval algorithm MAPP [37] uses the optimal estimation method to retrieve aerosol (top height, optical depth, size, complex refractive index) and ocean (wind speed, chlorophyll concentration) parameters. Retrieval algorithm GRASP [38] uses a multi-term Least Square Method statistical model to retrieve aerosol (optical depth, single-scattering albedo, size, shape, composition, etc.), ocean (BRDF/BPDF), and land (BRDF/BPDF) parameters (BRDF: bidirectional reflectance distribution function/BPDF: bidirectional polarization distribution function). These algorithms have been applied successfully to real-world multi-angle polarimeter measurements from the recent 2016 ORACLES [39] and 2017 ACEPOL campaigns [36,40–43].

However, inverse models can only be as accurate as their underlying forward models, which are still dominated by those assuming plane-parallel layers. Recent surveys have pointed out the need for inverse models that can account for the three-dimensional geometry of clouds, which are especially important for aerosol retrievals near cloud edges [44]. Monte-Carlo-based inverse models, such as those that have been developed for computer graphics, could provide more realistic estimates for such cases. Some in the remote sensing domain have developed methods to estimate gradients for cloud tomography by ad-hoc approximation [27–29,31], and retrieval of spatially varying cloud properties via differential radiative transfer simulation [26]. Czerninski and Schechner [26] use the same underlying principles as the *radiative backpropagation* method introduced to computer graphics (see Section 4.6) to estimate derivatives for Monte Carlo simulations. Our framework in Part 2 also uses differential radiative transfer simulation to estimate derivatives; while Czerninski and Schechner [26] focus on retrieval of spatially varying cloud properties, our goal is to build an extensible framework general to many types of scenes and properties.

As in remote sensing, inverse methods for acquiring material properties have a long and bountiful history in computer graphics; we choose to limit our discussion to the subset of methods that use *differentiable rendering*, described in Section 4.6.

Table 1
Glossary of computer graphics terms.

BSDF	Bidirectional scattering distribution function.
Caustics	Bright patches caused by refraction through a glassy surface (e.g. the dancing light patterns at the bottom of a pool).
Delta tracking	A Monte Carlo transmittance estimator for heterogeneous media imported from neutron transport (Woodcock et al. [45]) that simulates collisions with a mixture of real and “virtual” particles (also called null scattering or Woodcock tracking).
Emitter	A light source. This could be directional, an emissive surface (“area light”), etc.
Global illumination	Light contributions from indirect sources, such as interreflections between surfaces.
Importance sampling	A sampling strategy for Monte Carlo estimation of an integral $\int f(x) dx$ that draws samples from a probability distribution approximating the shape of the integrand $f(x)$ in order to reduce the variance of the estimate.
Multiple importance sampling (MIS)	A method introduced by Veach [46] to combine different importance sampling strategies by weighting their respective function evaluations and PDFs.
Next-event estimation (NEE)	A variance reduction method often added to path tracing where at every scatter event, one ray direction is sampled proportional to the BSDF/phase function and another ray direction is sampled proportional to the direction of an emitter, and their radiance estimates are combined using MIS.
Participating medium/volume	A scattering medium; a bounded space defined by a phase function and scattering and absorption coefficients (or equivalently, single-scattering albedo and extinction coefficient).
Path tracing	A type of rendering algorithm in which paths are built incrementally starting at the sensor and scatter throughout a scene, collecting radiance.
(to) render	To run a forward radiative transfer algorithm. The input is a scene and the output is measurements of radiance (typically an image).
Rendering algorithm	A (typically Monte-Carlo-based) forward radiative transfer algorithm.
Scene	The input to a rendering algorithm, which describes the parameterized layout, geometry, and material properties of all surfaces, scattering media, sensors, and emitters. This could also be thought of as a point in “state space”.

3. The evolution of forward radiative transfer models in computer graphics

In the 1980s and 1990s, one of the main applications of computer graphics was to generate frames of synthetic images for video games, movies, and television. Like remote sensing, early production environments required image synthesis algorithms that were accurate and fast, but in this case favored *geometric* complexity over *lighting* complexity when allocating limited computational resources.

The dominant rendering algorithms at the time were various flavors of rasterization: a multi-step pipeline that loads each object in a scene individually, projects it onto the screen, and then “shades” (determines the final color of) each pixel. This system was advantageous for early computer systems that had little memory for storing geometry: each object could be processed one at a time and there was assumed to be no global interdependency of color. Dedicated graphics hardware was developed around this assumption and can execute the rasterization pipeline with extreme efficiency. However, this meant that effects such as shadows cast by light sources and interreflection between surfaces were missing and the final image was (by modern standards) quite crude in appearance. Over time, some of these global effects were wedged into the rasterization pipeline by either pre-processing algorithms (e.g. creating shadow maps for shadows [47]) or human effort (artists would place light sources strategically to create a false impression of global illumination). However, adding each effect piece by piece ultimately proved to be an inelegant solution to realistic lighting.

While rasterization dominated production, important research in the academic community in the 1980s began building the foundations of *physically based rendering*. This research progressed along two parallel tracks: “radiosity” and Monte Carlo. Finite-element-based “radiosity” methods, inspired by heat transfer literature, discretize the scene geometry into patches and solve for the exchange of light between them (notable early works include Goral et al. [48] and Rushmeier [49], while Cohen and Wallace [50] provides an early overview). These methods produced images of then-unprecedented realism, since the images were generated using a physical simulation of light. However, this method scales poorly with increasingly complex geometry, emitters, and materials, and hence failed to find widespread adoption within commercial computer graphics settings.

Meanwhile, others began to integrate Monte Carlo sampling into rendering algorithms to simulate glossy reflections, soft shadows, depth

of field, and motion blur [51,52]. Kajiya’s seminal work [53] (concurrently introduced by Immel et al. [54]) merged Whitted’s recursive ray tracing algorithm [55] with radiative transfer theory to introduce “the rendering equation” (a formulation of the RTE for surfaces in a vacuum) to computer graphics:

$$L_o(\mathbf{x}, \boldsymbol{\omega}) = L_e(\mathbf{x}, \boldsymbol{\omega}) + \underbrace{\int_{\Omega} L_i(\mathbf{x}, \boldsymbol{\omega}_i) f_s(\mathbf{x}, \boldsymbol{\omega}, \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i}_{L_r(\mathbf{x}, \boldsymbol{\omega}) : \text{reflected radiance}} \quad (4)$$

Here $L_o(\mathbf{x}, \boldsymbol{\omega})$ is the outgoing radiance leaving a surface from position \mathbf{x} in direction¹ $\boldsymbol{\omega}$. This is equal to the sum of emitted radiance L_e and reflected radiance L_r . The reflected radiance is itself an integral of incoming radiance $L_i(\mathbf{x}, \boldsymbol{\omega}_i)$ from all directions $\boldsymbol{\omega}_i$ over the sphere of directions Ω . This light is weighted by the surface’s cosine-weighted BSDF f_s . In a vacuum, the incoming and outgoing radiances are related as $L_i(\mathbf{x}, \boldsymbol{\omega}) = L_o(r(\mathbf{x}, \boldsymbol{\omega}), \boldsymbol{\omega})$, where $r(\mathbf{x}, \boldsymbol{\omega})$ is the ray tracing operator that returns the closest surface point $\mathbf{x}_r := \mathbf{x} - t\boldsymbol{\omega}$ along the ray starting at \mathbf{x} in direction $-\boldsymbol{\omega}$. Eq. (4) is therefore a recursive integral equation.

Kajiya [53] also introduced an algorithm to estimate the rendering equation by Monte Carlo integration that became known as *path tracing*. A basic form of this algorithm arises by simply using a 1-sample Monte Carlo estimator to approximate the integral in Eq. (4):

$$\langle L_o(\mathbf{x}, \boldsymbol{\omega}) \rangle = L_e(\mathbf{x}, \boldsymbol{\omega}) + \frac{\langle L_o(r(\mathbf{x}, \boldsymbol{\omega}_i), \boldsymbol{\omega}_i) \rangle f_s(\mathbf{x}, \boldsymbol{\omega}, \boldsymbol{\omega}_i)}{p(\boldsymbol{\omega}_i)} \quad (5)$$

Rather than tracing the paths of light or “photons” from light sources, path tracing is an *adjoint* method that instead shoots rays from the *sensor* into the scene to intersect the nearest objects. Then, a new ray direction $\boldsymbol{\omega}_i$ is chosen based on random sampling; recursively estimating L_o via Eq. (5) builds a random path by alternating between intersection and random directional sampling, until the path is absorbed, hits an emitter, or escapes the scene. The contributions of many paths from the sensor are averaged to produce a final Monte Carlo estimate of radiance. This method scales well with increasingly complex geometry, emitters, and materials; however, it requires enough memory to store the entire scene at once (so that paths can be intersected with the scene’s geometry) and time to take sufficient samples to get rid of

¹ We use the convention that unit vectors always point along the flow of light, so *into* a scattering event \mathbf{x} when measuring incident radiance L_i , and *out* when measuring outgoing radiance L_o or L_e .




Fig. 2. Common rendering algorithms in a scene with a sensor, an emitter, and several surfaces. (a) Path tracing begins paths at the sensor — only paths that hit an emitter gather radiance, unless next-event estimation is used (shown in red); (b) Light tracing begins paths at the emitter — likewise, only paths that hit the sensor contribute to the image/measurement of radiance, unless next-event estimation is used; (c) Bidirectional path tracing generalizes this concept by beginning paths from both the sensor and emitter and connecting these “subpaths” manually to form paths of varying lengths; (d) Photon mapping separates the bidirectional concept into two passes: a first pass where light paths are generated and their radiance stored in the scene, and a second pass where sensor paths are generated and nearby stored radiance is gathered at every scatter event.

objectionable noise (the variance in Eq. (3)). More early work drew connections between the rendering equation and particle transport theory in physics [56,57] and began to adapt well-established sampling strategies to the computer graphics context [58,59]. Others began to investigate ways to speed up the computation of “global illumination”, or indirect lighting: for example, Ward’s “irradiance caching” [60] performs the expensive indirect lighting calculations at sparse locations in a scene and stores them for later lookup and interpolation by nearby paths.

In the 2000s, as computer memory and processing power became more abundant and less expensive, Monte Carlo approaches such as path tracing became more favorable due to their elegance and realism (see [61–64] for more historical discussion). Subsequent research in the new millennium has focused on improving the *statistical efficiency* of Monte Carlo rendering algorithms (there are several examples of this in Section 4). Today, production renderers for animation and visual effects for movies (e.g. Pixar’s RenderMan [65], Disney’s Hyperion [66], Solid Angle’s Arnold [67]) are typically Monte-Carlo-based, often utilizing large parallel computing clusters called “render farms” to mitigate the increased computational cost. Video games and other “real-time” applications, meanwhile, have lagged behind the path tracing revolution since the faster rasterization pipeline is necessary in order to produce many frames per second. However, much like the movies in the 90s and 00s, real-time applications are starting to incorporate Monte Carlo methods to achieve specific isolated effects [68]. As graphics hardware continues to evolve rapidly in performance and new architectures are designed around ray tracing, a similar revolution may be on the horizon [69,70].

While production renderers have embraced physically based approaches to achieve photorealism, they are not good candidates for scientific research since they are typically proprietary/closed-source and are tailored to the artist-driven production environment. However, a number of open-source, extensible, physically based renderers have been developed for academic research, including PBRT [22] and Mitsuba 3 [7].

4. Highlights of relevant computer graphics research

From the above foundations, recent computer graphics research has developed Monte Carlo frameworks for polarization [71–76], fluorescence [77,78], Lorenz-Mie and T-matrix scattering [79,80], speckle

statistics [81,82], and other wave interference effects [15,16,83,84]. While this non-exhaustive list demonstrates that forward radiative transfer models in computer graphics are capable of the base level of realism required for remote sensing, we highlight here several threads of computer graphics research that are especially relevant and ripe for migration to remote sensing (see Table 1 for terminology).

4.1. Multiple importance sampling and the generalization of rendering algorithms

Unidirectional tracing. While the introduction of path tracing was an important milestone in computer graphics, in its most basic form it is not well-suited to certain scenarios, for example finding bright caustics or small/narrow light sources (Fig. 2(a)). One could imagine mitigating these issues by reversing the path direction and instead starting paths at a light source, a variant known as *light tracing* (Fig. 2(b)); however, this simply reverses the problem as paths often miss the sensor.

Next-event estimation. One way to significantly improve on this is to sample direct connections to emitters (in the case of path tracing) or to the sensor (in the case of light tracing) at every scattering event — a technique known as *next-event estimation* (NEE) in computer graphics, or *local estimate/directional estimate* in remote sensing [85,86]. Mathematically, this can be accomplished by expanding Eq. (4) into itself once and separating the integral of directly visible emissive surfaces (direct illumination) from reflective surfaces (indirect illumination):

$$L_r(\mathbf{x}, \boldsymbol{\omega}) = \int_{\Omega} L_e(r(\mathbf{x}, \boldsymbol{\omega}_i), \boldsymbol{\omega}_i) f_s(\mathbf{x}, \boldsymbol{\omega}, \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i \quad (6)$$

$$+ \int_{\Omega} L_r(r(\mathbf{x}, \boldsymbol{\omega}_i), \boldsymbol{\omega}_i) f_s(\mathbf{x}, \boldsymbol{\omega}, \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i. \quad (7)$$

Since we know a-priori where all emissive surfaces are in the scene, Monte Carlo estimation of the first integral can leverage improved “emitter importance sampling” techniques [87] that send sampled directions only toward emitters. This can provide a significant variance reduction in the presence of concentrated emitters or small sensors, which would otherwise have a very low probability of being sampled.

Multiple importance sampling. While NEE reduces variance caused by missing the light sources/sensor, a direct connection cannot always be made (e.g. connections through refractive surfaces) and it may not be the optimal importance sampling strategy in some cases — for example,

sampling Eq. (6) proportional to L_e may still result in high variance if the remaining parts of the integrand (the BSDF f_s and cosine term) vary significantly over the integration domain. Multiple importance sampling (MIS) [46] allows for multiple strategies to be combined by sampling according to several strategies and appropriately weighting the results. In its basic form, MIS generates samples from the *mixture density* of several sampling techniques. The key idea is that complex integrands like in the RTE are the product of many different functions (incident radiance, BSDF, cosine foreshortening term). MIS can mix several different strategies, where each could be tailored to different terms of the integrand. While MIS does not always reduce variance compared to the best strategy for a particular scenario, it allows the overall estimator to be more robust to a wide variety of scenarios than using any of the individual strategies alone.

MIS also allows for combining meta-sampling strategies such as rendering algorithms themselves. *Bidirectional path tracing* [53,88,89], for example, starts paths from both light sources and the sensor and connects them (Fig. 2(c)), then weights the full paths using MIS. Bidirectional path tracing is especially good at handling scenes with complex light transport paths where traditional next-event estimation may fail, such as bright caustics (especially relevant in scenes with an ocean surface).

Advanced techniques. In addition to the above path tracing variants, computer graphics has developed a cornucopia of Monte Carlo rendering algorithms unfamiliar to the outside scientific world (see [90] for a comprehensive overview). For example, *photon mapping* [91] breaks rendering into two passes: a preprocessing pass akin to light tracing, where the position, direction, and power of each light subpath vertex are stored as a “photon” in a “photon map” (typically a k-d tree). This is followed by a rendering pass akin to path tracing, but where recursion is short-circuited by performing density estimation of nearby photons stored in the first pass (Fig. 2(d)). This method can successfully sample complex light transport paths like bright caustics, at the cost of introducing *bias* to the radiance estimate (i.e. error that will not vanish as the sample count increases). Later variants like photon beams [92] and photon surfaces [93,94] store not just the path vertices but segments, resulting in lower-variance estimates, and, in some cases, no bias. Another example is *Metropolis light transport* [95], a Markov Chain Monte Carlo algorithm that begins with path tracing or some other variant, but then mutates existing paths to find those with high path throughput via the Metropolis–Rosenbluth–Hastings sampling routine [96,97].

Each rendering algorithm is tailored to reducing variance in certain scenarios; there is no universal rendering algorithm that performs better than all others under all circumstances (yet).

4.2. Lower-variance and multiple-scattering microfacet models

Microfacet models are commonly used to model the appearance of rough surfaces by describing the statistical behavior of the surface using a distribution of normals and a tunable “roughness” parameter. The Cox–Munk model [98], often used to represent the ocean surface in remote sensing, is a microfacet model in which the wind speed at the surface has been mapped to the roughness parameter by an observation-driven transformation. It is well-approximated by a Gaussian distribution of heights and slopes, which is called a Beckmann distribution [99] in the context of microfacet models in computer graphics [100,101]. Other microfacet distributions commonly used in computer graphics are Trowbridge-Reitz/GGX [102,103], which exhibit longer-than-Gaussian tails that have been shown to match measurements of some real-world surfaces better than a Beckmann distribution [102–104].

Microfacet models have received substantial attention in computer graphics. For example, the method of Heitz and d’Eon [105] samples only the distribution of normals *visible to the sensor/emitter*, reducing variance significantly. Microfacet models traditionally assume that

light scatters exactly once off a single microfacet and then exits the surface (or is “absorbed” by the masking-shadowing term). However, real-world rough surfaces exhibit significant multiple scattering, energy which is simply lost in a microfacet model. Heitz et al. [106] introduced a method to simulate multiple scattering in microfacet models by treating the surface as a special kind of “scattering medium” and performing a nested random walk at each scatter event. Later work introduced lower-variance [107,108] and analytic (for special cases) [109] estimators for multiple scattering in microfacet surfaces, as well as estimators for layered surfaces with different statistics [110]. Such extensions to microfacet models could provide more accurate estimates of water-leaving radiance, especially at grazing angles and in higher-wind conditions.

4.3. Lower-variance estimators of spectrally and spatially varying scattering media

In homogeneous scattering media, estimating transmittance (the fraction of light that successfully passes through a stretch of medium) can easily be done by sampling “free-flight distances” (distances to the next scatter event within the medium) from an exponential probability distribution perfectly proportional to the falloff of light. However, the heterogeneous case is much more challenging because no closed-form probability distribution is available when the medium’s optical properties vary spatially.

A popular algorithm for estimating the transmittance of heterogeneous media called delta tracking, Woodcock tracking, or null scattering was developed for electron and neutron transport [45,111–116] and imported to graphics [117–123]. This method fills the scattering medium with “virtual” particles until it is a homogeneous density — then, sampling free-flight distances proceeds as if the medium is homogeneous, except the sampled scatter event is accepted or ignored with a probability proportional to the fraction of real particles at that location.

However, until recently, there was no way to calculate the sampling probability distribution of this method, which rendered it incompatible with MIS. Miller et al. [124] derived this missing probability distribution and demonstrated how combining *spectral* and *spatial* importance sampling via MIS significantly reduces overall variance. This, and other recent MIS developments [125], could be especially relevant for simulating mixtures of gases, which have distinct, narrow absorption spectra that are challenging to simulate efficiently.

4.4. “Random” sample generation

The error convergence rate of a Monte Carlo estimator using uncorrelated random samples is $\mathcal{O}(\sqrt{N})$ [46] (from Eq. (3)), which crucially depends only on the number of samples (N) and *not* the dimensionality of the integral. This is the property that makes it a feasible option for high-dimensional integrals such as the radiative transfer equation. However, this convergence rate is relatively slow: it takes four times the number of samples to reduce error by half. Computer graphics has imported a number of sample generation techniques that have been shown to improve the variance and convergence rate of Monte Carlo estimates of radiance. These sample sets are typically stratified (e.g. jittered [52], multijittered [126], Latin hypercube [127]/N-rooks [128], orthogonal arrays [129]) or carefully crafted (e.g. Halton [130] and Sobol [131]) to avoid clumps and gaps in the sample domain. The latter category of “crafted” sample sets are called *quasi Monte Carlo* — they replace the pseudorandom numbers used in the sampling process with deterministic, low-discrepancy point sets (see Keller [132] for a good discussion).

In practice, rendering frameworks in computer graphics typically abstract the concept of a “sampler” into a base class that is called to generate a sample whenever one is required and can be extended to implement any of these variants. Any sample generation algorithm can then easily be swapped in for a default independent random sampler, which is simply a wrapper for a pseudorandom number generator.

4.5. Non-exponential transport through scattering media

Traditional radiative transfer models assume that scattering particles within a medium are statistically independent, which results in exponential light decay along paths in the medium. However, real-world materials can exhibit both positive and negative correlations between scattering particles; notably, clouds exhibit positive correlations between droplets and therefore slower-than-exponential light decay [133,134]. Work in computer graphics [135–138], inspired by neutron transport [139–141] and the aforementioned work from atmospheric sciences, developed Monte Carlo frameworks that remove this assumption and allow for arbitrary (non-exponential) light falloff through a scattering medium. Non-exponential transport has the potential to achieve even more realistic simulation of clouds than existing forward radiative transfer models in remote sensing.

4.6. Differentiable rendering

One of the most exciting threads of computer graphics research to emerge within the past few years is differentiable rendering: computing the derivatives of arbitrarily complex Monte Carlo radiative transfer simulations, which can then be used to build inverse models via gradient-based optimization.

Many optimization methods require the gradients of a loss function with respect to parameters of interest in order to guide them toward a more accurate solution. However, deriving those gradients in the context of radiative transfer is clearly challenging — a single Monte Carlo radiance measurement is the sum of many paths scattered throughout the scene and influenced by the materials and geometry of everything they have touched. Initial work in computer graphics tackled subproblems such as estimating gradients of irradiance on surfaces [142–144] and in scattering media [145–147], direct lighting shadows [148,149], or material parameters [150–152]. Deriving the gradients of radiance with respect to arbitrary scene parameters (e.g. material parameters, vertex positions, sensor pose) remained elusive due to non-differentiable visibility discontinuities between objects in a scene. Consider a ray traveling from one object toward another: if its direction were perturbed slightly, that ray may hit or miss an occluder edge or silhouette edge — this is the discontinuity, a binary hit or miss. Li et al. [153] finally overcame this limitation by proposing an *edge sampling* method that finds the relevant edges of geometry in a scene and samples points along them in order to estimate the visibility gradient. This method was later generalized beyond surfaces [154]. The edge sampling routine incurs a hefty computational expense, however, and follow-up works have re-parameterized parts of the integral to avoid such discontinuities [155,156]. See Zhao et al. [157] for further discussion of recent advances in differentiable rendering.

The differentiable renderer Mitsuba 2 [71] (now Mitsuba 3 [7]), developed concurrently to the above work, takes a different approach by calculating the derivatives of light transport using *reverse-mode automatic differentiation*, or *backpropagation* as it is called in the context of neural networks, which is essentially recursive application of the chain rule [158]. Loubet et al. [159] introduced a change of variables formulation that can be used to avoid integrating over discontinuities.

Subsequent work introduced a method called *radiative backpropagation* [160–162] that avoids a cumbersome computation graph entirely by performing a *differential radiative transfer* simulation analogous to an ordinary radiative transfer simulation. This method has significantly sped up computation time and memory overhead compared to the previous approach, and also opened up the potential to study *differential sampling strategies* for this adjoint transport problem akin to existing sampling strategies for normal light transport [163].

Differential radiative transfer methods have the advantage that effects, even complex ones such as Lorenz-Mie scattering, can be implemented normally in the forward sense and their derivatives can

be automatically calculated without extra effort, as opposed to hand-calculated. As paths are mutually independent, they are just as “embarrassingly parallel” as ordinary radiative transfer. However, as with all Monte Carlo estimators, differential radiative transfer will produce noisy gradients and the user must choose an appropriate sample count. Differential radiative transfer is also just as computationally expensive as ordinary radiative transfer: subsequent works have explored ways to reuse information from one iteration of optimization to the next [26, 164,165].

5. Conclusion

By tracing the parallels and contrasts between radiative transfer models in the remote sensing and computer graphics communities, we lay the groundwork for a bridge of collaboration. This moment in time is ripe for cross-pollination of ideas due to the maturity of physically based rendering techniques for computer graphics and the need for increasingly sophisticated retrievals for remote sensing. We have also drawn attention to a number of research threads that have the potential to improve the accuracy and efficiency of state-of-the-art radiative transfer models in remote sensing, both through more faithful modeling of real-world, light-based phenomena (e.g. multiple-scattering microfacet models, non-exponential light transport) and carefully crafted sampling schemes (e.g. multiple importance sampling, stratified sampling, lower-variance estimators for microfacet models and scattering media). We are especially interested in the potential application of *differentiable rendering* to inverse models in remote sensing, and that is the avenue we intend to explore next. In Part 2 of this work, we take first steps toward putting these concepts into practice by extending Mitsuba 3, a differentiable rendering framework from the computer graphics community, to perform forward simulations of interest to remote sensing and validate our framework on benchmark tests of simple atmosphere–ocean systems.

CRedit authorship contribution statement

Katherine Salesin: Conceptualization, Formal analysis, Funding acquisition, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Kirk D. Knobelspiesse:** Conceptualization, Funding acquisition, Supervision, Writing – review & editing, Methodology. **Jacek Chowdhary:** Supervision, Writing – review & editing. **Peng-Wang Zhai:** Supervision, Writing – review & editing. **Wojciech Jarosz:** Conceptualization, Funding acquisition, Methodology, Resources, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Code will be made publicly available at the author’s Github. Data will be made available upon request.

Acknowledgments

The authors thank Wenzel Jakob and the rest of the Mitsuba 3 development team for consultations on Mitsuba 3. We also thank Meng Gao for insightful discussions on radiative transfer codes.

Funding

This work was supported by the National Aeronautics and Space Administration (NASA) [grant number 80NSSC22K1154]; the National Science Foundation (NSF) [grant number 1844538]; and a Neukom Institute CompX faculty grant.

References

- [1] Werdell Jeremy. Data products table. 2022, URL: https://pace.oceansciences.org/data_table.htm.
- [2] PACE Science Definition Team. Pre-aerosol, clouds, and ocean ecosystem (PACE) mission science definition team report. Technical report NASA/TM-2018-219027/Vol. 2, NASA, Goddard Space Flight Center; 2018, URL: <https://ntrs.nasa.gov/citations/20190000977>.
- [3] Remer LA, Davis AB, Mattoo S, Levy RC, Kalashnikova OV, Coddington O, Chowdhary J, Knobelspiesse K, Xu X, Ahmad Z, Boss E, Cairns B, Dierssen HM, Diner DJ, Franz B, Frouin R, Gao B-C, Ibrahim A, Martins JV, Omar AH, Torres O, Xu F, Zhai P-W. Retrieving aerosol characteristics from the PACE mission, part 1: Ocean color instrument. *Front Earth Sci* 2019;7:152, doi:10/kn5r.
- [4] Hasekamp OP, Fu G, Rusli SP, Wu L, Di Noia A, aan de Brugh J, Landgraf J, Martijn Smit J, Rietjens J, van Amerongen A. Aerosol measurements by SPEXone on the NASA PACE mission: expected retrieval capabilities. *J Quant Spectrosc Radiat Transfer* 2019;227:170–84, doi:10/kn5g.
- [5] Remer LA, Knobelspiesse K, Zhai P-W, Xu F, Kalashnikova OV, Chowdhary J, Hasekamp O, Dubovik O, Wu L, Ahmad Z, Boss E, Cairns B, Coddington O, Davis AB, Dierssen HM, Diner DJ, Franz B, Frouin R, Gao B-C, Ibrahim A, Levy RC, Martins JV, Omar AH, Torres O. Retrieving aerosol characteristics from the PACE mission, part 2: multi-angle and polarimetry. *Front Environ Sci* 2019;7:94, doi:10/kn5s.
- [6] Martins J, Nielsen T, Fish C, Sparr L, Fernandez-Borda R, Schoeberl M, Remer L. HARP CubeSat – an innovative hyperangular imaging polarimeter for earth science applications. In: Proceedings of the small sat pre-conference workshop, Logan, Utah. 2014.
- [7] Jakob W, Speierer S, Roussel N, Nimier-David M, Vicini D, Zeltner T, Nicolet B, Crespo M, Leroy V, Zhang Z. Mitsuba 3 renderer. 2022, URL: <https://mitsuba-renderer.org>.
- [8] van de Hulst HC. A new look at multiple scattering. New York, NY: NASA Institute for Space Studies; 1963.
- [9] Raschke E. Multiple scattering calculation of the transfer of solar radiation in an atmosphere-ocean system. *Beitrage Phys Atmos* 1972;45:1–19.
- [10] Quenzel H, Kaestner M. Optical properties of the atmosphere: Calculated variability and application to satellite remote sensing of phytoplankton. *Appl Opt* 1980;19:1338–44, doi:10/fh45hs.
- [11] Stamnes K, Tsay S-C, Wiscombe W, Jayaweera K. Numerically stable algorithm for discrete-ordinate-method radiative transfer in multiple scattering and emitting layered media. *Appl Opt* 1988;27:2502–9, doi:10/ch7chf.
- [12] Chowdhary J, Zhai P-W, Boss E, Dierssen H, Frouin R, Ibrahim A, Lee Z, Remer LA, Twardowski M, Xu F, Zhang X, Ottaviani M, Espinosa WR, Ramon D. Modeling atmosphere-ocean radiative transfer: A PACE mission perspective. *Front Earth Sci* 2019;7. doi:10/gpq9v5.
- [13] Lorenz L. Det kongelige danske videnskabernes selskabs skrifter (trykt utg.): Naturvidenskabelig og matematisk afdeling. In: *Lysbevægelsen i og uden for en af plane Lysbølger belyst kugle*, Vol. 6. 1890, p. 2–62.
- [14] Mie G. Beiträge zur optik trüber medien, speziell kolloidaler metallösungen. *Ann Phys* 1908;330:377–445, doi:10/frvnpq.
- [15] Steinberg S, Yan L-Q. A generic framework for physical light transport. *ACM Trans Graph (Proc SIGGRAPH)* 2021;40:139:1–20, doi:10/gpzk56.
- [16] Steinberg S, Sen P, Yan L-Q. Towards practical physical-optics rendering. *ACM Trans Graph (Proc SIGGRAPH)* 2022;41:132:1–24, doi:10/gqjn7s.
- [17] Steinberg S, Ramamoorthi R, Bitterli B, d'Eon E, Yan L-Q, Pharr M. A generalized ray formulation for wave-optics rendering. 2023, doi:10/kn6s. arXiv:2303.15762.
- [18] Plass GN, Kattawar GW. Radiative transfer in an atmosphere–ocean system. *Appl Opt* 1969;8:455–66, doi:10/b9xdb3.
- [19] Plass GN, Kattawar GW. Monte Carlo calculations of radiative transfer in the earth's atmosphere-ocean system: I. Flux in the atmosphere and ocean. *J Phys Oceanogr* 1972;2:139–45, doi:10/cwx3ms.
- [20] Kattawar GW, Plass GN. Monte Carlo calculations of radiative transfer in the earth's atmosphere-ocean system: II. Radiance in the atmosphere and ocean. *J Phys Oceanogr* 1972;2:146–56, doi:10/bzr2wm.
- [21] Kattawar GW, Plass GN, Guinn John A. Monte Carlo calculations of the polarization of radiation in the earth's atmosphere-ocean system. *J Phys Oceanogr* 1973;3:353–72, doi:10/c6nnt2.
- [22] Pharr M, Jakob W, Humphreys G. Physically based rendering: From theory to implementation. 4th ed.. Cambridge, MA: MIT Press; 2023.
- [23] Ramon D, Steinmetz F, Jolivet D, Compiègne M, Frouin R. Modeling polarized radiative transfer in the ocean-atmosphere system with the GPU-accelerated SMART-G Monte Carlo code. *J Quant Spectrosc Radiat Transfer* 2019;222–223:89–107, doi:10/kn5k.
- [24] Goodenough AA, Brown SD. DIRSIG5: Next-generation remote sensing data and image simulation framework. *IEEE J Sel Top Appl Earth Obs Remote Sens* 2017;10:4818–33, doi:10/gcms6c.
- [25] Foster RJ. Karle fellowship research report. NRL memorandum report AD1129741, Washington, DC: U.S. Naval Research Laboratory; 2021.
- [26] Czerninski I, Schechner YY. PARS - path recycling and sorting for efficient cloud tomography. *Intell Comput* 2023;2. doi:10/gr2hq7.
- [27] Loeub T, Levis A, Holodovsky V, Schechner YY. Monotonicity prior for cloud tomography. In: *Computer vision – ECCV 2020*. Glasgow, United Kingdom: Springer-Verlag; 2020, p. 283–99, doi:10/kn6v.
- [28] Holodovsky V, Schechner YY, Levin A, Levis A, Aides A. In-situ multi-view multi-scattering stochastic tomography. In: *IEEE international conference on computational photography (ICCP)*. 2016, p. 1–12, doi:10/kn6z.
- [29] Levis A, Schechner YY, Davis AB, Loveridge J. Multi-view polarimetric scattering cloud tomography and retrieval of droplet size. *Remote Sens* 2020;12:2831, doi:10/gbcqf7.
- [30] Dandini P, Cornet C, Binet R, Fenouil L, Holodovsky V, Schechner YY, Ricard D, Rosenfeld D. 3D cloud envelope and cloud development velocity from simulated CLOUD (C3IEL) stereo images. *Atmos Meas Tech* 2022;15:6221–42, doi:10/kn62.
- [31] Tzabari M, Holodovsky V, Shubi O, Eytan E, Koren I, Schechner YY. Settings for spaceborne 3-D scattering tomography of liquid-phase clouds by the cloudnet mission. *IEEE Trans Geosci Remote Sens* 2022;60:1–16, doi:10/kn63.
- [32] Frouin RJ, Franz BA, Ibrahim A, Knobelspiesse K, Ahmad Z, Cairns B, Chowdhary J, Dierssen HM, Tan J, Dubovik O, Huang X, Davis AB, Kalashnikova O, Thompson DR, Remer LA, Boss E, Coddington O, Deschamps P-Y, Gao B-C, Gross L, Hasekamp O, Omar A, Pelletier B, Ramon D, Steinmetz F, Zhai P-W. Atmospheric correction of satellite ocean-color imagery during the PACE era. *Front Earth Sci* 2019;7. doi:10/ggcggs.
- [33] Blondeau-Patissier D, Gower JF, Dekker AG, Phinn SR, Brando VE. A review of ocean color remote sensing methods and statistical techniques for the detection, mapping and analysis of phytoplankton blooms in coastal and open oceans. *Prog Oceanogr* 2014;123:123–44, doi:10/f53ds6.
- [34] Jamieson KG, Nowak R, Recht B. Query complexity of derivative-free optimization. In: Pereira F, Burges CJC, Bottou L, Weinberger KQ, editors. *Advances in neural information processing systems*, Vol. 25. Curran Associates, Inc.; 2012.
- [35] Gao M, Zhai P-W, Franz B, Hu Y, Knobelspiesse K, Werdell PJ, Ibrahim A, Xu F, Cairns B. Retrieval of aerosol properties and water-leaving reflectance from multi-angular polarimetric measurements over coastal waters. *Opt Express* 2018;26:8968–89, doi:10/kn5n.
- [36] Gao M, Franz BA, Knobelspiesse K, Zhai P-W, Martins V, Burton S, Cairns B, Ferrare R, Gales J, Hasekamp O, Hu Y, Ibrahim A, McBride B, Puthukkudy A, Werdell PJ, Xu X. Efficient multi-angle polarimetric inversion of aerosols and ocean color powered by a deep neural network forward model. *Atmos Meas Tech* 2021;14:4083–110, doi:10/gkdpwc.
- [37] Stamnes S, Hostetler C, Ferrare R, Burton S, Liu X, Hair J, Hu Y, Wasilewski A, Martin W, van Diedenhoven B, Chowdhary J, Cetinic I, Berg LK, Stamnes K, Cairns B. Simultaneous polarimeter retrievals of microphysical aerosol and ocean color parameters from the MAPP algorithm with comparison to high-spectral-resolution lidar aerosol and ocean products. *Appl Opt* 2018;57:2394–413, doi:10/kn5w.
- [38] Dubovik O, Fuertes D, Litvinov P, Lopatin A, Lapyonok T, Dubovik I, Xu F, Ducos F, Chen C, Torres B, Derimian Y, Li L, Herreras-Giralda M, Herrera M, Karol Y, Matar C, Schuster GL, Espinosa R, Puthukkudy A, Li Z, Fischer J, Preusker R, Cuesta J, Kreuter A, Cede A, Aspetsberger M, Marth D, Bindreiter L, Hangler A, Lanzinger V, Holter C, Federspiel C. A comprehensive description of multi-term LSM for applying multiple a priori constraints in problems of atmospheric remote sensing: GRASP algorithm, concept, and applications. *Front Remote Sens* 2021;2. doi:10/kn6n.
- [39] Segal-Rozenhaimer M, Miller DJ, Knobelspiesse K, Redemann J, Cairns B, Alexandrov MD. Development of neural network retrievals of liquid cloud properties from multi-angle polarimetric observations. *J Quant Spectrosc Radiat Transfer* 2018;220:39–51, doi:10/gfm9kg.
- [40] Fan C, Fu G, Di Noia A, Smit M, Rietjens JHH, Ferrare RA, Burton S, Li Z, Hasekamp OP. Use of a neural network-based ocean body radiative transfer model for aerosol retrievals from multi-angle polarimetric measurements. *Remote Sens* 2019;11. doi:10/kn5m.
- [41] Gao M, Zhai P-W, Franz BA, Hu Y, Knobelspiesse K, Werdell PJ, Ibrahim A, Cairns B, Chase A. Inversion of multiangular polarimetric measurements over open and coastal ocean waters: A joint retrieval algorithm for aerosol and water-leaving radiance properties. *Atmos Meas Tech* 2019;12:3921–41, doi:10/kn5p.
- [42] Gao M, Zhai P-W, Franz BA, Knobelspiesse K, Ibrahim A, Cairns B, Craig SE, Fu G, Hasekamp O, Hu Y, Werdell PJ. Inversion of multiangular polarimetric measurements from the ACEPOL campaign: An application of improving aerosol property and hyperspectral ocean color retrievals. *Atmos Meas Tech* 2020;13:3939–56, doi:10/kn5q.

- [43] Puthukkudy A, Martins JV, Remer LA, Xu X, Dubovik O, Litvinov P, McBride B, Burton S, Barbosa HJM. Retrieval of aerosol properties from airborne hyper-angular rainbow polarimeter (AirHARP) observations during ACEPOL 2017. *Atmos Meas Tech* 2020;13:5207–36, doi:10/kn6m.
- [44] Marshak A, Ackerman A, Silva AD, Eck T, Holben B, Kahn R, Kleidman R, Knobelspiesse K, Levy R, Lyapustin A, Oreopoulos L, Remer L, Torres O, Varnai T, Wen G, Yorks J. Aerosol properties in cloudy environments from remote sensing observations: A review of the current state of knowledge. *Bull Am Meteorol Soc* 2021;1–57, doi:10/kn6f.
- [45] Woodcock ER, Murphy T, Hemmings PJ, Longworth TC. Techniques used in the GEM code for Monte Carlo neutronics calculations in reactors and other systems of complex geometry. In: Applications of computing methods to reactor problems. Argonne National Laboratory; 1965.
- [46] Veach E. Robust Monte Carlo methods for light transport simulation (Ph.D. thesis), Stanford University; 1997.
- [47] Williams L. Casting curved shadows on curved surfaces. *Comput Graph (Proc SIGGRAPH)* 1978;12:270–4, doi:10/ftz4zz.
- [48] Goral CM, Torrance KE, Greenberg DP, Battaile B. Modeling the interaction of light between diffuse surfaces. *Comput Graph (Proc SIGGRAPH)* 1984;18:213–22, doi:10/fsjssr.
- [49] Rushmeier HE. Realistic image synthesis for scenes with radiatively participating media (Ph.D. thesis), Ithaca, NY: Cornell University; 1988.
- [50] Cohen MF, Wallace JR. Radiosity and realistic image synthesis. NY: Academic Press; 1993.
- [51] Cook RL, Porter T, Carpenter L. Distributed ray tracing. *Comput Graph (Proc SIGGRAPH)* 1984;18:137–45.
- [52] Cook RL. Stochastic sampling in computer graphics. *ACM Trans Graph* 1986;5:51–72, doi:10/cqwhec.
- [53] Kajiyama JT. The rendering equation. *Comput Graph (Proc SIGGRAPH)* 1986;20:143–50, doi:10/cvf53j.
- [54] Immel DS, Cohen MF, Greenberg DP. A radiosity method for non-diffuse environments. *Comput Graph (Proc SIGGRAPH)* 1986;20:133–42, doi:10/dmjm9t.
- [55] Whitted T. An improved illumination model for shaded display. *Commun ACM* 1980;23:343–9.
- [56] Arvo JR. Transfer functions in global illumination. In: Global illumination, ACM SIGGRAPH courses. 1993, p. 1–28.
- [57] Arvo JR, Kirk D. Particle transport and image synthesis. *Comput Graph (Proc SIGGRAPH)* 1990;24:63–6, doi:10/dtp6gd.
- [58] Kirk D, Arvo J. Unbiased variance reduction for global illumination. In: Photorealistic rendering in computer graphics (proceedings of the eurographics workshop on rendering), Barcelona, 1991.
- [59] Shirley P. Physically based lighting calculations for computer graphics (Ph.D. thesis), Urbana–Champaign: University of Illinois; 1990.
- [60] Ward GJ, Rubinstein FM, Clear RD. A ray tracing solution for diffuse interreflection. *Comput Graph (Proc SIGGRAPH)* 1988;22:85–92, doi:10/dk6rt5.
- [61] Keller A, Fascione L, Fajardo M, Georgiev I, Christensen P, Hanika J, Eisner C, Nichols G. The path-tracing revolution in the movie industry. In: ACM SIGGRAPH courses. New York, NY, USA: ACM Press; 2015, p. 24:1–7, doi:10/gfzps5t.
- [62] Christensen PH, Jarosz W. The path to path-traced movies. *Found Trends® Comput Graph Vis* 2016;10:103–75, doi:10/gfjwjc.
- [63] Fascione L, Hanika J, Heckenberg D, Kulla C, Droske M, Schwarzhaupt J. Path tracing in production: Part 1: Modern path tracing. In: ACM SIGGRAPH courses. New York, NY: ACM Press; 2019, doi:10/ggd8nj.
- [64] Jakob W, Weidlich A, Beddini A, Pieké R, Tang H, Fascione L, Hanika J. Path tracing in production: Part 2: Making movies. In: ACM SIGGRAPH courses. New York, NY: ACM Press; 2019, doi:10/ggd8nh.
- [65] Christensen P, Fong J, Shade J, Wooten V, Schubert B, Kensler A, Friedman S, Kilpatrick C, Ramshaw C, Bannister M, Rayner B, Brouillat J, Liani M. RenderMan: An advanced path-tracing architecture for movie rendering. *ACM Trans Graph* 2018;37:30:1–21, doi:10/gfznsb.
- [66] Burley B, Adler D, Chiang MJ-Y, Driskill H, Habel R, Kelly P, Kutz P, Li YK, Teece D. The design and evolution of Disney's hyperion renderer. *ACM Trans Graph* 2018;37:33:1–22, doi:10/gfjm8w.
- [67] Georgiev I, Ize T, Farnsworth M, Montoya-Vozmediano R, King A, Lommel BV, Jimenez A, Anson O, Ogaki S, Johnston E, Herubel A, Russell D, Servant F, Fajardo M. Arnold: A brute-force production path tracer. *ACM Trans Graph* 2018;37:32:1–32:12, doi:10/gfznb2.
- [68] Wyman C, Kettunen M, Lin D, Bitterli B, Yuksel C, Jarosz W, Kozłowski P. A gentle introduction to ReSTIR path reuse in real-time. In: ACM SIGGRAPH courses. New York, NY: ACM Press; 2023, doi:10/knqq.
- [69] Haines E, Akenine-Möller T, editors. Ray tracing gems: High-quality and real-time rendering with DXR and other APIs. Berkeley, CA: A Press; 2019, doi:10/c54w.
- [70] Marrs A, Shirley P, Wald I, editors. Ray tracing gems II: Next generation real-time rendering with DXR, Vulkan, and OptiX. Berkeley, CA: A Press; 2021, doi:10/kn6q.
- [71] Nimier-David M, Vicini D, Zeltner T, Jakob W. Mitsuba 2: A retargetable forward and inverse renderer. *ACM Trans Graph (Proc SIGGRAPH Asia)* 2019;38:203:1–203:17, doi:10/ggd8m8.
- [72] Wilkie A, Weidlich A. A standardised polarisation visualisation for images. In: Proceedings of the spring conference on computer graphics (SCCG). New York, NY: Association for Computing Machinery; 2010, p. 43–50, doi:10/dx637c.
- [73] Wilkie A, Weidlich A. Polarised light in computer graphics. In: SIGGRAPH Asia 2012 courses. 2012, doi:10/gfz5nn.
- [74] Mojzík M, Skřivan T, Wilkie A, Křivánek J. Bi-directional polarised light transport. In: Eisemann E, Fiume E, editors. Proceedings of EGSR (experimental ideas & implementations). The Eurographics Association; 2016, doi:10/gfzsmx.
- [75] Jarabo A, Arellano V. Bidirectional rendering of vector light transport. *Comput Graph Forum* 2018;37:96–105, doi:10/gdvk6m.
- [76] Baek S-H, Zeltner T, Ku HJ, Hwang I, Tong X, Jakob W, Kim MH. Image-based acquisition and modeling of polarimetric reflectance. *ACM Trans Graph (Proc SIGGRAPH)* 2020;39, doi:10/gg8xcv.
- [77] Mojzík M, Fichet A, Wilkie A. Handling fluorescence in a uni-directional spectral path tracer. *Comput Graph Forum (Proc Eurograph Symp Render)* 2018;37:77–94, doi:10/gd5hcm.
- [78] Jung A, Wilkie A, Hanika J, Jakob W, Dachsbacher C. Wide gamut spectral upsampling with fluorescence. *Comput Graph Forum (Proc Eurograph Symp Render)* 2019;38:87–96, doi:10/kn5x.
- [79] Frisvad JR, Christensen NJ, Jensen HW. Computing the scattering properties of participating media using Lorenz-Mie theory. *ACM Trans Graph (Proc SIGGRAPH)* 2007;26:60, doi:10/bf3p26.
- [80] Guo Y, Jarabo A, Zhao S. Beyond mie theory: Systematic computation of bulk scattering parameters based on microphysical wave optics. *ACM Trans Graph (Proc SIGGRAPH Asia)* 2021;40, doi:10/kn59.
- [81] Bar C, Alterman M, Gkioulekas I, Levin A. A Monte Carlo framework for rendering speckle statistics in scattering media. *ACM Trans Graph (Proc SIGGRAPH)* 2019;38:39:1–22, doi:10/gf5jbk.
- [82] Bar C, Gkioulekas I, Levin A. Rendering near-field speckle statistics in scattering media. *ACM Trans Graph (Proc SIGGRAPH Asia)* 2020;39, doi:10/kn55.
- [83] Steinberg S. Accurate rendering of liquid-crystals and inhomogeneous optically anisotropic media. *ACM Trans Graph* 2020;39, doi:10/kn57.
- [84] Steinberg S. Analytic spectral integration of birefringence-induced iridescence. *Comput Graph Forum (Proc Eurograph Symp Render)* 2019;38:97–110, doi:10/kn58.
- [85] Marshak A, Davis A, editors. 3D radiative transfer in cloudy atmospheres, physics of earth and space environments. Berlin/Heidelberg: Springer-Verlag; 2005, doi:10/b859pq.
- [86] Emde C, Mayer B. Simulation of solar radiation during a total eclipse: A challenge for radiative transfer. *Atmos Chem Phys* 2007;7:2259–70, doi:10/btdtzw.
- [87] Shirley P, Wang C, Zimmerman K. Monte Carlo techniques for direct lighting calculations. *ACM Trans Graph* 1996;15:1–36, doi:10/ddgbbg.
- [88] Lafortune EP, Willems YD. Bi-directional path tracing. In: Proceedings of the international conference on computational graphics and visualization techniques (compugraphics), Volume 93, Alvor, Portugal. 1993, p. 145–53.
- [89] Veach E, Guibas LJ. Bidirectional estimators for light transport. In: Photorealistic rendering techniques (proceedings of the eurographics workshop on rendering). Springer-Verlag; 1995, p. 145–67, doi:10/gfznbh.
- [90] Dutré P, Jensen HW, Arvo J, Bala K, Bekaert P, Marschner S, Pharr M. State of the art in Monte Carlo global illumination. In: ACM SIGGRAPH courses. New York, NY: Association for Computing Machinery; 2004, doi:10/d6g74t.
- [91] Jensen HW. Realistic image synthesis using photon mapping. Natick, MA, USA: AK Peters, Ltd.; 2001.
- [92] Jarosz W. Efficient Monte Carlo methods for light transport in scattering media (Ph.D. thesis), San Diego: University of California; 2008.
- [93] Bitterli B, Jarosz W. Beyond points and beams: Higher-dimensional photon samples for volumetric light transport. *ACM Trans Graph (Proc SIGGRAPH)* 2017;36:112:1–112:12, doi:10/gfzbnr.
- [94] Deng X, Jiao S, Bitterli B, Jarosz W. Photon surfaces for robust, unbiased volumetric density estimation. *ACM Trans Graph (Proc SIGGRAPH)* 2019;38, doi:10/gf6rx9.
- [95] Veach E, Guibas LJ. Metropolis light transport. In: Annual conference series (proceedings of SIGGRAPH), Vol. 31. ACM Press; 1997, p. 65–76, doi:10/Bkjaj4.
- [96] Metropolis N, Rosenbluth A, Rosenbluth M, Teller A, Teller E. Equations of state calculations by fast computing machines. *J Chem Phys* 1953;21:1087–91, doi:10/ds736f.
- [97] Hastings WK. Monte Carlo sampling methods using markov chains and their applications. *Biometrika* 1970;vol. 57:97–109, doi:10/ds736f.
- [98] Cox C, Munk W. Measurement of the roughness of the sea surface from photographs of the sun's glitter. *J Opt Soc Amer* 1954;44:838–50, doi:10/ddbn9x.
- [99] Beckmann P, Spizzichino A. The scattering of electromagnetic waves from rough surfaces. NY: Pergamon Press; 1963.
- [100] Torrance KE, Sparrow EM. Theory for off-specular reflection from roughened surfaces. *J Opt Soc Amer* 1967;57:1105–14, doi:10/dw6f9n.
- [101] Blinn JF. Models of light reflection for computer synthesized pictures. *Comput Graph (Proc SIGGRAPH)* 1977;11:192–8, doi:10/dg64f5.
- [102] Trowbridge TS, Reitz KP. Average irregularity representation of a rough surface for ray reflection. *J Opt Soc Amer* 1975;65:531–6, doi:10/cx4c7.

- [103] Walter B, Marschner SR, Li H, Torrance KE. Microfacet models for refraction through rough surfaces. In: Rendering techniques (proceedings of the eurographics symposium on rendering). Eurographics Association; 2007, p. 195–206, doi:10/gfz4kg.
- [104] Dupuy J, Heitz E, Iehl J-C, Poulin P, Ostromoukhov V. Extracting microfacet-based BRDF parameters from arbitrary materials with power iterations. *Comput Graph Forum (Proc Eurograph Symp Render)* 2015;34:21–30, doi:10/f7mbrd.
- [105] Heitz E, d'Eon E. Importance sampling microfacet-based BSDFs using the distribution of visible normals. *Comput Graph Forum (Proc Eurograph Symp Render)* 2014;33:103–12, doi:10/f6fgph.
- [106] Heitz E, Hanika J, d'Eon E, Dachsbacher C. Multiple-scattering microfacet BSDFs with the smith model. *ACM Trans Graph (Proc SIGGRAPH)* 2016;35, doi:10/f89kkm.
- [107] Wang B, Jin W, Fan J, Yang J, Holzschuch N, Yan L-Q. Position-free multiple-bounce computations for smith microfacet BSDFs. *ACM Trans Graph (Proc SIGGRAPH)* 2022;41:134:1–134:14, doi:10/gqjn7v.
- [108] Bitterli B, d'Eon E. A position-free path integral for homogeneous slabs and multiple scattering on smith microfacets. *Comput Graph Forum (Proc Eurograph Symp Render)* 2022;41, doi:10/k646.
- [109] Lee JH, Jarabo A, Jeon DS, Gutierrez D, Kim MH. Practical multiple scattering for rough surfaces. *ACM Trans Graph (Proc SIGGRAPH Asia)* 2018;37:275:1–275:12, doi:10/db3d.
- [110] d'Eon E, Bitterli B, Weidlich A, Zeltner T. Microfacet theory for non-uniform heightfields. In: ACM SIGGRAPH conference papers. Los Angeles, CA: ACM Press; 2023, doi:10/kn65.
- [111] Butcher JC, Messel H. Electron number distribution in electron-photon showers. *Phys Rev* 1958;112:2096–106, doi:10/dmbvs8.
- [112] Butcher JC, Messel H. Electron number distribution in electron-photon showers in air and aluminium absorbers. *Nucl Phys* 1960;20:15–28, doi:10/bqxd9g.
- [113] Zerby CD, Curtis RB, Bertini HW. The relativistic doppler problem. Technical report ORNL-61-7-20, Oak Ridge, TN, USA: Oak Ridge National Laboratory; 1961, doi:c542.
- [114] Bertini HW. Monte Carlo simulations on intranuclear cascades. Technical report ORNL-3383, Oak Ridge, TN, USA: Oak Ridge National Laboratory; 1963, doi:c54v.
- [115] Miller LB. Monte Carlo analysis of reactivity coefficients in fast reactors; general theory and applications. Technical report ANL-7307, Argonne, IL, USA: Argonne National Laboratory; 1967, doi:c54x.
- [116] Coleman WA. Mathematical verification of a certain Monte Carlo sampling technique and applications of the technique to radiation transport problems. *Nucl Sci Eng* 1968;32:76–81, doi:10/gfzndg.
- [117] Raab M, Seibert D, Keller A. Unbiased global illumination with participating media. In: Keller A, Heinrich S, Niederreiter H, editors. Monte Carlo and Quasi-Monte Carlo methods. Springer-Verlag; 2008, p. 591–605, doi:10/bsxp2j.
- [118] Novák J, Selle A, Jarosz W. Residual ratio tracking for estimating attenuation in participating media. *ACM Trans Graph (Proc SIGGRAPH Asia)* 2014;33:179:1–179:11, doi:10/f6r2nq.
- [119] Kutz P, Habel R, Li YK, Novák J. Spectral and decomposition tracking for rendering heterogeneous volumes. *ACM Trans Graph (Proc SIGGRAPH)* 2017;36:111:1–111:16, doi:10/gbxjxg.
- [120] Szirmay-Kalos L, Tóth B, Magdics M. Free path sampling in high resolution inhomogeneous participating media. *Comput Graph Forum* 2011;30:85–97, doi:10/bpkdzm.
- [121] Szirmay-Kalos L, Georgiev I, Magdics M, Molnár B, Légrády D. Unbiased light transport estimators for inhomogeneous participating media. *Comput Graph Forum (Proc Eurograph)* 2017;36:9–19, doi:10/gbm5tr.
- [122] Misso Z, Li YK, Burley B, Teece D, Jarosz W. Progressive null-tracking for volumetric rendering. In: ACM SIGGRAPH conference papers. Los Angeles, CA: ACM Press; 2023, doi:10/kmdw.
- [123] Novák J, Georgiev I, Hanika J, Jarosz W. Monte Carlo methods for volumetric light transport simulation. *Comput Graph Forum (Proc Eurograph State Art Rep)* 2018;37:551–76, doi:10/gd2jqc.
- [124] Miller B, Georgiev I, Jarosz W. A null-scattering path integral formulation of light transport. *ACM Trans Graph (Proc SIGGRAPH)* 2019;38:44:1–44:13, doi:10/gf6rzb.
- [125] West R, Georgiev I, Gruson A, Hachisuka T. Continuous multiple importance sampling. *ACM Trans Graph (Proc SIGGRAPH)* 2020;39, doi:10/gg8xcr.
- [126] Chiu K, Shirley P, Wang C. Multi-jittered sampling. In: Heckbert PS, editor. Graphics gems IV. San Diego, CA, USA: Academic Press; 1994, p. 370–4.
- [127] McKay MD, Beckman RJ, Conover WJ. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* 1979;21:239–45, doi:10/gfzncd.
- [128] Shirley P. Discrepancy as a quality measure for sample distributions. In: Proceedings of eurographics. Eurographics Association; 1991, p. 183–94, doi:10/gfzncd.
- [129] Jarosz W, Enayet A, Kensler A, Kilpatrick C, Christensen P. Orthogonal array sampling for Monte Carlo rendering. *Comput Graph Forum (Proc Eurograph Symp Render)* 2019;38:135–47, doi:10/gf6rx5.
- [130] Halton JH. Algorithm 247: Radical-inverse quasi-random point sequence. *Commun ACM* 1964;7:701–2, doi:10/dd3674.
- [131] Sobol IM. The distribution of points in a cube and the approximate evaluation of integrals. *USSR Comput Math Math Phys* 1967;7:86–112, doi:10/crdj6j.
- [132] Keller A. Quasi-Monte Carlo image synthesis in a nutshell. In: Dick J, Ku FY, Peters GW, Sloan IH, editors. Monte Carlo and Quasi-Monte Carlo methods. Springer-Verlag; 2013, p. 213–49, doi:10/gfz9gw.
- [133] Davis AB, Marshak A, Gerber H, Wiscombe WJ. Horizontal structure of marine boundary layer clouds from centimeter to kilometer scales. *J Geophys Res: Atmos* 1999;104:6123–44, doi:10/fd2bcz.
- [134] Kostinski AB, Jameson AR. On the spatial distribution of cloud particles. *J Atmos Sci* 2000;57:901–15, doi:10/fhs4w5.
- [135] Bitterli B, Ravichandran S, Müller T, Wrenninge M, Novák J, Marschner S, Jarosz W. A radiative transfer framework for non-exponential media. *ACM Trans Graph (Proc SIGGRAPH Asia)* 2018;37:225:1–225:17, doi:10/gfz2cm.
- [136] d'Eon E. A reciprocal formulation of nonexponential radiative transfer. I: Sketch and motivation. *J Comput Theor Transp* 2018;47:84–115, doi:10/gg5kzj.
- [137] d'Eon E. A reciprocal formulation of nonexponential radiative transfer. 2: Monte Carlo estimation and diffusion approximation. *J Comput Theor Transp* 2019;48:201–62, doi:10/gg5kzf.
- [138] Jarabo A, Aliaga C, Gutierrez D. A radiative transfer framework for spatially-correlated materials. *ACM Trans Graph (Proc SIGGRAPH)* 2018;37:83:1–83:13, doi:10/gd52qq.
- [139] Larsen EW, Vasques R. A generalized linear Boltzmann equation for non-classical particle transport. *J Quant Spectrosc Radiat Transfer* 2011;112:619–31, doi:10/dfbvmm.
- [140] Vasques R, Larsen EW. Non-classical particle transport with angular-dependent path-length distributions. I: Theory. *Ann Nucl Energy* 2014;70:292–300, doi:10/gfz723.
- [141] Vasques R, Larsen EW. Non-classical particle transport with angular-dependent path-length distributions. II: Application to pebble bed reactor cores. *Ann Nucl Energy* 2014;70:301–11, doi:10/f55pxw.
- [142] Ward GJ, Heckbert PS. Irradiance gradients. In: Chalmers A, Paddon D, Sillion FX, editors. CE_EGWR93. Bristol, UK: Consolidation Express Bristol; 1992, p. 85–98.
- [143] Krivánek J, Gautron P, Bouatouch K, Pattanaik S. Improved radiance gradient computation. In: Proceedings of the spring conference on computer graphics (SCCG). New York, NY, USA: ACM Press; 2005, p. 155–9, doi:10/c8pb3t.
- [144] Krivánek J, Gautron P, Pattanaik S, Bouatouch K. Radiance caching for efficient global illumination computation. *IEEE Trans Vis Comput Graphics* 2005;11:550–61, doi:10/csf2sw.
- [145] Jarosz W, Donner C, Zwicker M, Jensen HW. Radiance caching for participating media. *ACM Trans Graph* 2008;27:7:1–7:11, doi:10/cwnw78.
- [146] Jarosz W, Zwicker M, Jensen HW. Irradiance gradients in the presence of participating media and occlusions. *Comput Graph Forum (Proc Eurograph Symp Render)* 2008;27:1087–96, doi:10/bg8nww.
- [147] Marco J, Jarabo A, Jarosz W, Gutierrez D. Second-order occlusion-aware volumetric radiance caching. *ACM Trans Graph* 2018;37:1–14, doi:10/gdv86k.
- [148] Arvo J. The irradiance Jacobian for partially occluded polyhedral sources. In: Annual conference series (proceedings of SIGGRAPH). ACM Press; 1994, p. 343–50, doi:10/bkjm2r.
- [149] Ramamoorthi R, Mahajan D, Belhumeur P. A first-order analysis of lighting, shading, and shadows. *ACM Trans Graph* 2007;26:2:1–21, doi:10/dhkp9c.
- [150] Gkioulekas I, Zhao S, Bala K, Zickler T, Levin A. Inverse volume rendering with material dictionaries. *ACM Trans Graph (Proc SIGGRAPH Asia)* 2013;32, doi:10/gbd53t.
- [151] Khungurn P, Schroeder D, Zhao S, Bala K, Marschner S. Matching real fabrics with micro-appearance models. *ACM Trans Graph* 2016;35, doi:10/fdsd.
- [152] Zhao S, Wu L, Durand F, Ramamoorthi R. Downsampling scattering parameters for rendering anisotropic media. *ACM Trans Graph (Proc SIGGRAPH Asia)* 2016;35:166:1–166:11, doi:10/f9cpqs.
- [153] Li T-M, Aittala M, Durand F, Lehtinen J. Differentiable Monte Carlo ray tracing through edge sampling. *ACM Trans Graph (Proc SIGGRAPH Asia)* 2018;37:222:1–222:11, doi:10/gfztzx.
- [154] Zhang C, Wu L, Zheng C, Gkioulekas I, Ramamoorthi R, Zhao S. A differential theory of radiative transfer. *ACM Trans Graph (Proc SIGGRAPH Asia)* 2019;38:227:1–227:16, doi:10/ggd8m9.
- [155] Zhang C, Miller B, Yan K, Gkioulekas I, Zhao S. Path-space differentiable rendering. *ACM Trans Graph (Proc SIGGRAPH)* 2020;39, doi:10/gg8xc2.
- [156] Zhang C, Yu Z, Zhao S. Path-space differentiable rendering of participating media. *ACM Trans Graph (Proc SIGGRAPH)* 2021;40:76:1–76:15, doi:10/gqjn75.
- [157] Zhao S, Jakob W, Li T-M. Physics-based differentiable rendering: From theory to implementation. In: ACM SIGGRAPH courses. Virtual Event USA: ACM Press; 2020, p. 1–30, doi:10/gjvrje.
- [158] Griewank A, Walther A. Evaluating derivatives: Principles and techniques of algorithmic differentiation. Society for Industrial and Applied Mathematics; 2008, doi:10/cthcq6.
- [159] Loubet G, Holzschuch N, Jakob W. Reparameterizing discontinuous integrands for differentiable rendering. *ACM Trans Graph (Proc SIGGRAPH Asia)* 2019;38:228:1–228:14, doi:10/ggfg2z.

- [160] Nimier-David M, Speierer S, Ruiz B, Jakob W. Radiative backpropagation: An adjoint method for lightning-fast differentiable rendering. *ACM Trans Graph (Proc SIGGRAPH)* 2020;39. doi:10/gg8xc5.
- [161] Vicini D, Speierer S, Jakob W. Path replay backpropagation: Differentiating light paths using constant memory and linear time. *ACM Trans Graph (Proc SIGGRAPH)* 2021;40:108:1–108:14, doi:10/gnngwp.
- [162] Stam J. Computing light transport gradients using the adjoint method. 2020, doi:10/kpk2. arXiv:2006.15059.
- [163] Zeltner T, Speierer S, Georgiev I, Jakob W. Monte Carlo estimators for differential light transport. *ACM Trans Graph (Proc SIGGRAPH)* 2021;40:78:1–78:16, doi:10/gqjn76.
- [164] Nicolet B, Rousselle F, Novák J, Keller A, Jakob W, Müller T. Recursive control variates for inverse rendering. *ACM Trans Graph (Proc SIGGRAPH)* 2023;42. doi:10/kn6t.
- [165] Chang W, Sivaram V, Nowrouzezahrai D, Hachisuka T, Ramamoorthi R, Li T-M. Parameter-space ReSTIR for differentiable and inverse rendering. In: *ACM SIGGRAPH conference papers*. Los Angeles, CA: ACM Press; 2023, doi:10/kn67.