

Ibón Guillén Serrano

Computational Light Transport for Forward and Inverse Problems.

Director/es

Jarabo Torrijos, Adrián
Gutiérrez Pérez, Diego

<http://zaguan.unizar.es/collection/Tesis>



Universidad
Zaragoza

Tesis Doctoral

COMPUTATIONAL LIGHT TRANSPORT FOR
FORWARD AND INVERSE PROBLEMS.

Autor

Ibón Guillén Serrano

Director/es

Jarabo Torrijos, Adrián
Gutiérrez Pérez, Diego

UNIVERSIDAD DE ZARAGOZA
Escuela de Doctorado

Programa de Doctorado en Ingeniería de Sistemas e Informática

2021

UNIVERSIDAD DE ZARAGOZA

DOCTORAL THESIS

**Computational Light Transport for
Forward and Inverse Problems**

Author:
Ibón Guillén

Supervisors:
Diego Gutiérrez
Adrián Jarabo

Departamento de Informática e Ingeniería de Sistemas
Escuela de Ingeniería y Arquitectura

May 7, 2021

Abstract

Computational light transport simulation comprises all techniques for calculating light flux in a virtual scene. It is ubiquitous in many applications, ranging from entertainment and advertising, to product design, engineering and architecture, to generating accurate simulation data for computational imaging techniques. However, accurately simulating light transport is an expensive process. This results into a trade-off between the physical fidelity of the simulation and computational resources. For example, it is common to assume geometric optics or infinite speed of light, or to simplify reflectance models ignoring specific phenomena. In this thesis, we introduce several contributions to light transport simulation, both in terms of improving the efficiency of the involved calculations, and expanding its range of practical applications. We pay special attention to uplift the infinite speed of light assumption, generalizing light transport to its transient state. With regard to efficiency, we present an improved method to calculate the light flux in surfaces arriving directly from luminaries in a Monte Carlo renderer, reducing significantly the variance of the resulting images with the same time budget. We also introduce a density estimation-based framework in the transient domain to better reuse temporal samples in participating media. In the application domain, we also introduce two uses of light transport: A model to simulate a special kind of goniochromatic pigments that exhibit pearlescent appearance, with the goal of provide intuitive editing for manufacture, and a framework for non-line-of-sight imaging using light time of flight information built over a wave-based light transport formulation.

Resumen

El transporte de luz computacional comprende todas las técnicas usadas para calcular el flujo de luz en una escena virtual. Su uso es ubicuo en distintas aplicaciones, desde entretenimiento y publicidad, hasta diseño de producto, ingeniería y arquitectura, incluyendo el generar datos validados para técnicas basadas en imagen por ordenador. Sin embargo, simular el transporte de luz de manera precisa es un proceso costoso. Como consecuencia, hay que establecer un balance entre la fidelidad de la simulación física y su coste computacional. Por ejemplo, es común asumir óptica geométrica o una velocidad de propagación de la luz infinita, o simplificar los modelos de reflectancia ignorando ciertos fenómenos. En esta tesis introducimos varias contribuciones a la simulación del transporte de luz, dirigidas tanto a mejorar la eficiencia del cálculo de la misma, como a expandir el rango de sus aplicaciones prácticas. Prestamos especial atención a remover la asunción de una velocidad de propagación infinita, generalizando el transporte de luz a su estado transitorio. Respecto a la mejora de eficiencia, presentamos un método para calcular el flujo de luz que incide directamente desde luminarias en un sistema de generación de imágenes por Monte Carlo, reduciendo significativamente la variancia de las imágenes resultantes usando el mismo tiempo de ejecución. Asimismo, introducimos una técnica basada en estimación de densidad en el estado transitorio, que permite reusar mejor las muestras temporales en un medio participativo. En el dominio de las aplicaciones, también introducimos dos nuevos usos del transporte de luz: Un modelo para simular un tipo especial de pigmentos gonicromáticos que exhiben apariencia perlescente, con el objetivo de proveer una forma de edición intuitiva para manufactura, y una técnica de imagen sin línea de visión directa usando información del tiempo de vuelo de la luz, construida sobre un modelo de propagación de la luz basado en ondas.

Measurable Contributions

This thesis has led to the following results, which can be found in detail in Section 1.4:

- 4 JCR-indexed journal publications [74; 75; 142; 149].
- A peer-reviewed conference publication [73].
- A research stay (three months) at École Polytechnique Fédérale de Lausanne (EPFL) in Switzerland.
- A research visit (one week) at University of Wisconsin-Madison in USA.
- A supervised final degree project (TFG).
- Participation in 3 research projects.
- A FPI (Formación de Personal Investigador) PhD grant.

Acknowledgements

During the years that I have been working on the present thesis, many people have helped and/or encouraged me. Their support have made it possible for me to reach the point where I am, and they are therefore in part responsible, for better or worse, for the contents of the present thesis. I would like to express my gratitude towards them, and for that purpose I include here a list as exhaustive as possible given my memory constraints. I apologize beforehand for any omission.

My supervisors, Diego and Adrian, for introducing me to the world of research, giving me the chance of working on it and putting up with me for so long. Working with me is not always easy, yet they have made a good job, of which this thesis is proof.

The members of the Graphics & Imaging Lab, the ones who stay behind now that I'm leaving and the ones who already left. It has been a good place to spend the last few years, and I have learned quite a lot during that time. Specially thanks to Manuel and Miguel for endless discussions and sharing of ideas, Víctor for his optimism and unorthodox programming, and Ana for always answering my questions.

My co-authors and collaborators, without whom none of the projects presented here could have been carried out. Thanks for guiding me, sharing your ideas and points of view, and working hand-in-hand with me.

The hosts during my internships: Andreas Velten at Winconsin and Wenzel Jakob at Lausanne, for giving me the chance to visit other workplaces and learn different approaches to research. And a special thanks the welcoming people I meet during my stay at EPFL: Merlin, Delio, Tizian and Guillaume.

My friends, for being there for me and giving me the chance to forget about work for a few hours, or even entire days in some cases.

Of course, my parents, who have always taken care of me and given me so much I don't even know where to start.

And finally, my sister for being always optimistic and reminding me constantly that when I feel totally lost, I'm just being too dramatic.

This work has been funded by the Ministerio de Economía, Industria y Competitividad through project LIGHTSPEED, the Defense Advanced Research Projects Agency through project REVEAL, and the European Research Council through Chameleon project.

Contents

Abstract	iii
Resumen	v
Contributions	vii
Acknowledgements	ix
I Introduction	1
1 Introduction	3
1.1 Physically-based rendering	3
1.1.1 Wave optics	5
1.1.2 Light transport in transient state	5
1.2 Beyond computer graphics	5
1.2.1 Predictive rendering	6
1.2.2 Inverse rendering	6
1.3 Goals and Overview	7
1.4 Contributions and measurable results	8
1.4.1 Publications	8
1.4.2 Research stays and visits	9
1.4.3 Projects	9
1.4.4 Other contributions	9
II Steady-state light transport	11
2 Area-preserving parameterizations for spherical ellipses	13
2.1 Introduction	13
2.2 Problem statement and previous work	14
2.3 Solid angle sampling of an oriented disk	16
2.3.1 Area-preserving mappings	17
2.3.2 Parallel Mapping	18
2.3.3 Radial mapping	20
2.4 Implementation	23
2.4.1 Tabulation	23
2.4.2 Efficiency	23
2.5 Results	24
2.6 Conclusions	26
2.A Derivation of Equation (2.9)	28
2.B Derivation of Equation (2.18)	28

2.C	Numerical computation of incomplete elliptic integrals using Carlson forms	29
2.C.1	Evaluating $\Omega_p^+(\phi_p)$ using Carlson functions	30
2.C.2	Evaluating $\Omega_r(\phi_r)$ using Carlson functions.	31
3	A general framework for pearlescent materials	33
3.1	Introduction	33
3.2	Related work	35
3.3	Pearlescent materials	36
3.4	Scattering by iridescent platelets	37
3.4.1	Scattering in substrate-based platelets	38
3.4.2	Analysis	39
3.5	Light transport in pearlescent materials	41
3.5.1	Modeling optical properties of pearlescent media	42
3.6	Implementation	44
3.6.1	Position-free Monte Carlo	44
3.6.2	Spectral rendering	44
3.6.3	Directional distribution.	44
3.6.4	Substrate thickness distribution	45
3.6.5	Platelet stratum extinction coefficient	46
3.7	Analysis and evaluation	46
3.7.1	Comparison with captured data	46
3.7.2	Exploration of the parameter space	47
3.7.3	Additional results	50
3.8	Conclusions	52
3.A	Tables	54
III	Transient state light transport	57
4	Progressive transient photon beams	59
4.1	Introduction	59
4.2	Related work	61
4.3	Transient radiative transfer	62
4.4	Transient photon beams	63
4.4.1	Our algorithm	64
4.5	Progressive transient photon beams	66
4.6	Results	70
4.7	Conclusions	73
4.A	Error in transient progressive photon beams	74
4.B	Variance and expected value of the error of the time-resolved beam radiance estimate	74
4.C	Convergence analysis of progressive transient photon beams	76
4.D	Minimizing Asymptotic Mean Squared Error	77
5	Non-line-of-sight imaging using Phasor Fields	79
5.1	Introduction	80
5.2	Related work	80
5.3	Phasor Fields in non-line-of-sight imaging	81
5.4	Results	83
5.5	Conclusions	85

5.A	Phasor Field propagation	87
5.A.1	Propagating broadband signals	88
5.A.2	Non-Lambertian surfaces	88
5.B	Line-of-sight template functions	89
5.B.1	Phase operator of an ideal computational thin lens	89
5.B.2	Projector and camera functions	91
5.C	Implementation of RSD solvers	92
5.C.1	Conventional photography camera using RSD	92
5.C.2	Confocal time-gated system using backprojection	93
5.C.3	Transient camera using backprojection	94
5.D	Details on data acquisition	94
5.D.1	Hardware configuration	94
5.D.2	NLOS measurement geometry	94
5.D.3	Exposure time	95
5.D.4	Collected data	95
5.D.5	Helmholtz reciprocity	99
5.E	Additional validation and discussion	100
5.E.1	Resolution limits	100
5.E.2	Effect of exposure time	100
5.E.3	Effect of strong interreflections	102
5.E.4	Non-Lambertian surfaces	103
5.E.5	Reconstruction comparison with other methods	103
IV	Conclusions	107
6	Conclusions	109
	Bibliography	113

List of Figures

1.1	Evolution of computer graphics	4
1.2	Transient imaging	6
1.3	Examples of predictive rendering	7
1.4	Transient imaging and simulation	8
2.1	A stratified unit-square sample pattern mapped onto the surface of a disk, using existing techniques and our proposed maps	15
2.2	Geometry of a disk, its solid angle projection, and associated reference frames	17
2.3	Surface area equivalence between the sphere and the cylinder, and different lateral projections of the spherical ellipse	18
2.4	Illustration of our parallel and radial maps	19
2.5	Illustration of our polar mapping	21
2.6	A scene illuminated by a double-sided disk light	25
2.7	A scene with participating media illuminated by a single-sided disk light	25
2.8	Colored disk lights rendered in Arnold using area sampling and our tabulated ld-radial map	25
2.9	Error and cost plots for the results in Figure 2.6 and Figure 2.7 w.r.t. sample count	26
2.10	View of the tangent ellipse	29
3.1	Examples of pearlescent cosmetic bottles	34
3.2	Electronic microscopic photographs of pearlescent materials and diagrams of our model	37
3.3	Scattering functions for different iridescent platelets	40
3.4	Changes in platelet response due to substrate thickness	41
3.5	Response of a mica-substrate platelet with fixed coating and varying substrate thickness distributions	45
3.6	Comparison of our model with ground-truth measurements	47
3.7	Chromaticity space of measured materials, our model and previous work	47
3.8	Effect of the different parameters of material C1 on its appearance	49
3.9	Schematic view of the material structure used for modeling materials C1 and C2	49
3.10	2D image slices representations of the reflectance of material C1	49
3.11	Variations of pearlescent materials due to changes on the thickness of the coating layer and the density of platelets	50
3.12	Comparison between the chromaticity of the diffuse and gloss components for the C1 and C2 materials	51
3.13	Variations of material C1 with increasing platelet rotation	51

3.14	Cars rendered with complex pearlescent paints, exhibiting platelet rotation and anisotropy	51
3.15	Equal-time renderings demonstrating the importance of both multiple importance sampling and spectral rendering using hero wavelength	52
4.1	The SOCCER scene	60
4.2	Photon emission and radiance estimation	63
4.3	Spatio-temporal kernel estimators	64
4.4	Steady-state renders for the scenes CORNELL SPHERES (Figure 4.5), MIRRORS (Figure 4.6), and PUMPKIN (Figure 4.7).	71
4.5	Comparison of CORNELL SPHERES scene using <i>camera-unwarping</i>	71
4.6	Comparison between Dirac delta and continuous (Heaviside) emission.	72
4.7	The PUMPKIN scene	72
4.8	The JUICE scene	73
5.1	NLOS as a virtual LOS imaging system	81
5.2	Reconstructions of a complex NLOS scene	83
5.3	Robustness of our technique	84
5.4	Additional novel NLOS imaging applications of our method	85
5.5	Effect of relay wall reflectance on Phasor Field reconstructions	89
5.6	Capture hardware used for the results shown in this chapter	95
5.7	Data Comparison for different exposure times	97
5.8	Visualization of the raw data for our long exposure office scene	97
5.9	Robustness to ambient light and noise	101
5.10	Short exposure reconstructions	101
5.11	Comparison of phasor fields to prior methods	101
5.12	Short exposure photon data	102
5.13	Robustness to multiple reflections	103
5.14	Robustness to scene reflectance	104
5.15	Reconstruction comparison on a public dataset 1	105

List of Tables

3.1	Error comparison of our model and Ergun et al.'s model	47
3.2	Total render time, samples per pixel, and resolution for the different scenes used throughout the chapter	52
A1	Characterization of the materials used throughout the chapter	54
A2	Optimized parameters for our model and Ergun's model	54
A3	Bottles: Materials description	55
A4	Cars: Materials description	55
5.1	Illumination wave functions for different light sources, used in our three example imaging systems.	90
5.2	Imaging operators to implement our three example imaging systems.	91
5.3	Photon statistics for the captured data	96

Part I

Introduction

Chapter 1

Introduction

Light transport simulation encompasses all techniques whose purpose is to calculate the light flux arriving to a virtual sensor on a given scene, taking into account the relevant interactions of light with the objects contained in the scene, up to the level of detail required for a particular application. Examples of the usage of light transport techniques include real-time video game graphics, special effect elements for movie production, or product previsualization in architecture and manufacturing. Each of these applications impose different constraints in terms of fidelity on the underlying physical simulation and computational complexity.

The goal of generating photorealistic graphics is to produce images as close as possible to real ones, generating computer images is a resource intensive task. For example, the computational budget available on real-time applications is limited to 16 ms of GPU time, while generating a single frame in a movie can take several hours on dedicated computer clusters. For this reason, the simulation systems need to balance between image fidelity, modelling as accurately as possible the underlying physical processes, and affordability, doing so in a reasonable time budget with the available computational resources.

Traditionally, the field has been focused on the generation of realistic images approximating the physical reality. However, as the field has matured, it has reached the point where the opposite has become possible, and starting from captured images from the real world, it is possible to estimate properties of the objects present on it such as its shape and material, or the illumination setting of the scene. This is generally called *inverse rendering* [166].

This dissertation has two complementary goals: First, to improve the efficiency of existing techniques for calculating light transport simulation; and second, to employ light transport to solve specific problems using light transport simulation. In particular, we predict the appearance of goniochromatic pigments by closely simulating light scattering inside them, and develop a new computational imaging method for capturing scenes outside the line-of-sight of the observer by posing it as a light transport problem.

1.1 Physically-based rendering

As computational power has become more readily available and less expensive, the focus of computer graphics has become not only generating synthetic images efficiently, but also to make them more credible. This pursuit of realism has resulted in incorporating accurate physical models in the rendering pipeline, slowly turning the generation of computer images into a full-fledged light transport simulation.




Figure 1.1: Evolution of computer graphics. Despite its relatively short history, the field has seen a constant evolution, driven both by the increase in available computing power and the demand for higher fidelity. Since the early ray traced images from the 80s (**left**), passing from the adoption of computer generated imagery in movie production in the 90s (**center**), to modern photorealistic characters (**right**). Left image from: [233]. Middle image from: [36] (© 1993 Industrial Light and Magic). Right image from: [58] (© 2017 Twentieth Century Fox Film).

Describing the exact nature of light is complex. It is known since old times that light does not propagate instantly but rather has a finite speed [184]. However in the spatial scale at which we humans live this propagation time is so small that can be easily dismissed. The same reasoning can be applied to the wave nature of light, which has also been known [247] and formalized [61] since the early XIX century. Yet in most situations a good estimate of radiance can be made without solving the complex equations of classical electromagnetic theory [152]. Even more evasive from day-to-day experiences is the quantum nature of light, describing light emission, absorption and scattering in discrete bundles of energy, or photons [51; 177]. These photons have individual properties, and interact in precise ways with every other particles, including other photons [1], although for most applications we are only interested in their bulk behavior.

The use of physically-based light transport approaches in rendering can be traced back to the 80s, with Whitted's ray tracing [233], which demonstrated how brute-force geometric optics, tracing the trajectory of thousands of light paths, could be used to achieve complex lighting effects (Figure 1.1, left). Ray tracing was later extended to account for multiple interreflections and non-local effects [38]. Kajiya [117] provided the link between image generation and classical particle transport, leading to the adoption of stochastic methods and more rigorous treatment of the rendering process [10].

In parallel with the adoption of physically-based light transport simulation, another effort towards realism was on the development of physically-based appearance models. Borrowing from the existing efforts in optics, the rendering community adopted the now-standard microfacet theory [13; 220] to model rough surfaces [22; 39], later generalized to anisotropic surfaces [93], multiple layered materials [14; 78; 229], or even discrete surfaces with glinty appearances [102; 244].

In computer graphics, in order to keep the balance between computability and accuracy, light is most often considered to be a continuous quantity that propagates instantly following straight paths, only changing directions when interacting with either solid objects or suspended particles. This is the so called *steady-state geometric optics*. This formulation has led to quite efficient calculations using Monte Carlo methods, and is the foundation of modern renderers [172]. However, this formulation has severe downsides, and some relevant effects cannot be simulated without

incorporating more complex light transport models on top of it. In the following, we briefly describe some of these generally ignored effects, which are addressed in this thesis.

1.1.1 Wave optics

The wave nature of light has been known since the XIX century [247], and its theoretical basis are well developed [61; 152]. The use of wave optics to calculate light transport in graphics was already proposed by Moravec [156], and He et al. [88] used it to model light scattering from rough surfaces, but it has not been widely adopted due to its high cost.

However, without taking into account wave optics some appearance effects simply cannot be reproduced. This has led to the development of simplified appearance models reproducing some of the most relevant wave-based effects, keeping mathematical consistency and being computationally affordable. Examples of these works include diffraction [44; 243], fluorescence [105; 114], interference [15; 204], diffraction [243], or small particle electromagnetic scattering [62].

Our own work on simulating pearlescent materials, presented in Chapter 3, fits into this line of work by incorporating wave optics based dispersion into a volumetric transport model to model goniochromatic effects on pigments.

1.1.2 Light transport in transient state

The fact the light does not propagate at an infinite speed have been observed since ancient time, and it is apparent at astronomical scales [184]. At terrestrial scale, and for most practical situations, the effects of non-instantaneous light propagation are imperceptible, and consequently computer graphics have traditionally overlooked the fact. However, simulating light in transient state is useful in areas outside of multimedia production, and an increasingly large number of computational imaging techniques use temporal information to capture extra information about the real world [55; 107]. An example of this kind of transient capture can be seen in Figure 1.2.

Removing the assumption of instant propagation requires keeping record of the propagation time between light scattering events. Unfortunately, there is little control on the propagation time: Therefore, naively simulating transient light transport can lead to unsuitable slow convergence [175; 200], and thus specialized sampling and reconstruction schemes are required. Jarabo et al. [106] introduced temporal-aware routines that uniformly distribute the sample on the temporal domain, leading to efficient reconstructions using density estimation.

In Chapter 4 we extend Jarabo et al's formulation of time resolved light transport, introducing a novel density estimation in both the spatial and temporal dimensions, allowing more efficient rendering of delta or almost delta volumetric light paths.

1.2 Beyond computer graphics

As rendering has developed more and more sophisticated models for image synthesis and appearance reproduction, it has reached a point where its utility goes far beyond its traditional multimedia applications. The achieved level of fidelity now allows its use for predicting the appearance of manufactured goods before fabricating




Figure 1.2: Transient imaging. Steady-stated image of a scene (in color) containing an alien toy and a mirror, and transient frames (in grayscale) of a light pulse traversing it, each one with exposure time of just 1.85 ps. Image from: [225].

them, allowing faster and cheaper workflows, or even designing computationally materials with user-defined appearances.

In addition, the accuracy and performance of current rendering techniques allow to inverse light transport processes. Starting from captured images from the real world, it is possible to invert the light transport and estimate the optical properties of the imaged objects. In the following, we briefly describe these two applications, that align to the contributions of this thesis.

1.2.1 Predictive rendering

Since the fidelity of the generated images with respect to the simulated reality is quite high, and the parameters used to generate them are linked to the physical properties of the objects, it is now possible to estimate accurately the appearance of an object before it gets manufactured. This new field is usually called *predictive rendering* [235], and has been successful applications in a wide variety of industries such as architecture [132], furniture catalogs [86], 3D printing [52; 174; 210], or even objects designed to project certain light patterns when illuminated [165; 189; 248]. Figure 1.3 shows examples of the last two mentioned applications.

Following this line of work, we approach the problem of modelling and predicting the appearance of diffraction-based gnochromatic pigments, widely used in car paint, cosmetics, and packaging in Chapter 3.

1.2.2 Inverse rendering

The traditional process of computer graphics is to create a synthetic image starting from a mathematical definition of a scene, using light transport algorithms to calculate the effect of all light interactions. Given how remarkably accurate modern rendering techniques and how close synthetic and real images are, the inverse process have become possible: starting with real photograph it is possible, and given that certain scene parameters such as illumination are known, to calculate the inverse light transport to recover the geometry and/or material properties of a real scene. In particular, in this thesis we focus on invert light transport in its transient state, assuming finite propagation speed. As with steady state images, transient-state simulations are considerably close to real captures, as can be appreciated in Figure 1.4.

Transient imaging While most imaging techniques focus on capturing a steady-state image, resulting of light propagating across the scene, recent advances in ultra-fast imaging have made it possible to capture the propagation of light itself, in what



Figure 1.3: Examples of predictive rendering. **Left:** A glass object designed to project a certain refraction pattern when illuminated from behind. The surface of the object was computationally optimized by simulating the light transport inside the object until its refraction converged to the desired pattern. **Right:** A 3D printed object designed to display a well-defined and high contrast texture pattern on its surface. The distribution of pigments in the inner layers of the object was optimized to reduce color bleeding and color washing, caused by light scattering by the highly transparent thermoplastic resins. Left images from: [189]. Right images from: [210].

is termed *transient imaging* [55; 107]. Each pixel has a temporal domain in which each bin represents the amount of light reaching the sensor at a given instant in time, which for the purpose of being able to capture the propagation of light needs to be in the picosecond scale.

Early attempts at capturing light in motion date back to the 70s, with Abramson’s holographic images of light pulses [2]. However, recent advances in sensing technology have enabled the widespread of imaging techniques built around the time-of-flight (ToF) of light. Different approaches have been proposed, based on varied technologies including streak cameras [224; 225], photonic mixed devices [90; 115], interferometry [67], modulated continuous waves [79] or gated sensors [135].

Applications of ToF imaging include applications in a wide range of industries, such as collision detection in autonomous cars [216], depth estimation for object detection in multimedia devices as the Kinect 2 [190], or material detection based on temporal responses [158; 208; 214].

From these applications, imaging of scenes outside the direct line of sight of the observer, called non-line-of-sight (NLOS) imaging, has received notable attention. The first practical demonstration of NLOS imaging was introduced by Velten et al. [224]. In the basic configuration, a light pulse is emitted by a controlled light source, bounces off a secondary visible wall, travels across the hidden scene, propagates back to the relay wall, and finally reaches the sensor. Knowing the geometry of the visible surfaces, it is possible to employ the time-of-flight of the light pulse to calculate the light transport across the hidden scene, recovering in the process information about the geometry and materials of the objects within it. This has opened a new imaging modality, allowing to see the invisible by exploiting hidden information in the temporal domain. In Chapter 5 we introduce a novel mathematical theory that allows to, for the first time, image hidden scenes at the meter scale, with arbitrary complexity, materials and significant multiple scattering.

1.3 Goals and Overview

This dissertation has two complementary goals: First, to improve the efficiency of existing techniques for calculating light transport simulation; and second, to solve specific problems using light transport simulation. In particular, we focus in two




Figure 1.4: Transient imaging and simulation. Each row shows a steady state image of a scene at the left, and three frames of a light pulse traversing the scene, with exposure times on the picosecond scale. The series at the **top** is a real capture, while the **bottom** one is a simulation of the same scene. While there are some visible differences between both series, mainly due to approximate materials and camera properties, this example illustrates how a simulation can accurately reproduce light transport on transient state, helping to understand the underlying phenomena and serving as a testbed for imaging applications. Top series from [225]. Bottom series from [106].

main problems: High-fidelity material modelling for predictive appearance modeling targeting industry, and non-line-of-sight imaging using time-resolved imaging. For clarity, we have divided our contribution along one main axis: The temporal domain. In particular, our contributions have been splitted in two parts: Part II (Chapters 2 and 3), containing the techniques assuming steady-state light transport, and Part III (Chapters 4 and 5) the ones exploring the transient state of light.

While I am not the lead author of many of the works included in this thesis, I have contributed on the development of all the projects. In order to give due credit, at the start of each chapter I contextualize the line of work that lead to the developments described within that particular section, and outline my personal contribution as well as that of my colleagues.

1.4 Contributions and measurable results

1.4.1 Publications

All the work presented in this dissertation have already been published on technical venues. Concretely, four articles published in JCR-indexed journals and a presentation on an international conference, which I describe in detail below:

- Area-preserving parameterizations for spherical ellipses [75]
Computer Graphics Forum, 36(4), 2017
This journal has an impact factor of 2.046, and its position in the JCR index 2017 is 22 out of 104 (Q1) in the category Computer Science, Software Engineering.
- Progressive transient photon beams [149]
Computer Graphics Forum, 38(6), 2019
This journal has an impact factor of 2.116, and its position in the JCR index 2019 is 38 out of 108 (Q2) in the category Computer Science, Software Engineering.
- Non-line-of-sight imaging using phasor-field virtual wave optics [142]
Nature, 572, 2019

This journal has an impact factor of 42.779, and its position in the JCR index 2017 is 1 out of 71 (Q1) in the category Multidisciplinary Sciences.

- On the effect of reflectance on phasor field non-line-of-sight imaging [73]
2020 IEEE International Conference on Acoustics, Speech and Signal Processing
This conference has a CORE 2018 class B rating.
- A general framework for pearlescent materials [74]
ACM Transactions on Graphics, 39(6), 2020
This journal has an impact factor of 5.084, and its position in the JCR index 2019 is 8 out of 108 (Q1) in the category Computer Science, Software Engineering.

1.4.2 Research stays and visits

This thesis has also involved the realization of the following research internships:

- A three-month research visit at *École Polytechnique Fédérale de Lausanne* in Lausanne, Switzerland, hosted by Prof. Wenzel Jakob. Our work on simulating pearlescent appearance, described in Chapter 3, resulted from this collaboration.
- A one-week research visit at *University of Wisconsin-Madison* in Madison, USA, hosted by Prof. Andreas Velten. This was part of an ongoing collaboration with Wisconsin's Computational Optics Group, within the research project REVEAL, that resulted in our work on NLOS imaging detailed in Chapter 5.

1.4.3 Projects

As part of my PhD thesis, I have collaborated in several research projects

- LIGHTSPEED: Computational transient imaging. *Ministerio de Economía, Industria y Competitividad (MINECO)*, project N° TIN2016-78753-P. PI: Diego Gutiérrez.
- REVEAL: Scene recovery using an extended plenoptic function. *Defense Advanced Research Projects Agency (DARPA)*, research Subcontract N° 678K904. PI: Diego Gutiérrez, Adrián Jarabo.
- CHAMELEON: Intuitive editing of visual appearance from real-world datasets. *European Research Council (ERC)*, grant agreement N° 682080. PI: Diego Gutiérrez.

1.4.4 Other contributions

Our improved light sampling method, described in Chapter 2, was developed in collaboration with Solid Angle and has been since then integrated into their Academy Award-winner rendering system Arnold, starting in the 5.0 version [5].

Part II

Steady-state light transport

Chapter 2

Area-preserving parameterizations for spherical ellipses

In this chapter, our goal is to improve the efficiency of physically based rendering, by developing new methods for calculating the illumination subtended by circular light sources. Specifically, we introduce new sampling schemes designed for their usage on Monte Carlo integration, using area-preserving mappings uniform in the solid angle subtended by the light source. Our method can be easily incorporated into a renderer, and we demonstrate how our method overcomes existing approaches, providing significantly lower variance with a small increase in runtime cost.

The project was done during my stay at Universidad Rey Juan Carlos in Madrid, and involved collaborators from Universidad de Granada and Solid Angle. My role as the first author involved developing the mapping described in Section 2.3.3, as well as generating results to compare the different techniques, while the alternative solid angle mapping described in Section 2.3.2 was proposed by Prof. Carlos Ureña. While most of the renders in this section were generated using my own implementation of the method, Figure 2.8 was created using the version currently integrated in the Arnold Renderer [5].

This work was originally published in Computer Graphics Forum and presented at the 28th Eurographics Symposium on Rendering (EGSR 2017).

Area-preserving parameterizations for spherical ellipses

Ibón Guillén, Carlos Ureña, Alan King, Marcos Fajardo, Iliyan Georgiev, Jorge López-Moreno & Adrian Jarabo

Computer Graphics Forum, Vol. 36(4), 2017

2.1 Introduction

Illumination from area light sources is among the most important lighting effects in realistic rendering, due to the ubiquity of such sources in real-world scenes. Monte Carlo integration is the standard method for computing the illumination from such luminaires [198]. This method is general and robust, supports arbitrary reflectance models and geometry, and predictively converges to the actual solution as the number of samples increases. Accurately sampling the illumination from area light sources is crucial for minimizing the amount of noise in rendered images.

Estimating the direct illumination at a point requires sampling the radiance contribution from directions inside the solid angle subtended by the given luminaire. A sensible strategy is to distribute those directions uniformly. This, however, is

hard to achieve for an arbitrary-shaped luminaire, as it involves first computing and then uniformly sampling its subtended solid angle. Specialized methods have been proposed for spherical [226], triangular [8; 222], rectangular [223], and polygonal lights [9]. These elaborate solid angle sampling techniques are more computationally expensive than naïve methods that uniformly sample the surface area of the luminaire. However, in most non-trivial scenes, where the sample contribution evaluation is orders of magnitude more costly than the sample generation, their lower variance improves overall efficiency.

Few works have focused on sampling oriented disk-shaped light sources. Disk lights are important in practice, both for their artistic expressiveness and their use in a number of real-world scenarios, generally including man-made light sources such as in architectural lighting, film and photography. Moreover, disk lights form the base for some approximate global illumination algorithms [87; 199]. So far, the only practical method for uniformly sampling the solid angle of disk lights is the work by [65], who proposed a rejection sampling approach that generates candidates using spherical quad sampling [223]. Unfortunately, achieving good sample stratification with this method requires special care.

In this work we present a set of methods for uniformly sampling the solid angle subtended by an oriented disk. We exploit the fact that a disk, as seen from a point, is bounded by an elliptical cone [50] and thus its solid angle defines a spherical ellipse whose properties have been analyzed in depth [24]. This allows us to define two different exact area-preserving mappings that can be used to transform stratified unit-square sample patterns to stratified directions on the subtended spherical ellipse. We describe how to efficiently implement these mappings in practice and demonstrate the lower variance they achieve compared to previous work.

2.2 Problem statement and previous work

Our goal is to compute the radiance L_s scattered at a point \mathbf{x} in direction ω_o due to irradiance from a disk-shaped luminaire D . This can be written as an integral over the solid angle Ω_D subtended by the luminaire:

$$L_s(\mathbf{x}, \omega_o) = \int_{\Omega_D} f(\mathbf{x}, \mathbf{x}_{\hat{\omega}}, \omega_o, \hat{\omega}) d\mu(\hat{\omega}), \quad (2.1)$$

where $\mathbf{x}_{\hat{\omega}}$ is the first visible point from \mathbf{x} in direction $\hat{\omega}$, μ is the solid angle measure, and the contribution function f is

$$f(\mathbf{x}, \mathbf{x}, \omega_o, \hat{\omega}) = \begin{cases} L_e(\mathbf{x}, -\hat{\omega}) f_s(\mathbf{x}, \omega_o, \hat{\omega}) |\hat{\omega} \cdot \hat{n}_x|, & \text{if } \mathbf{x} \text{ is on a surface,} \\ L_e(\mathbf{x}, -\hat{\omega}) \rho(\mathbf{x}, \omega_o, \hat{\omega}) T(\mathbf{x}, \mathbf{x}), & \text{if } \mathbf{x} \text{ is in a medium,} \end{cases}$$

with f_s , \hat{n}_x , and ρ being respectively the BSDF, surface normal, and medium phase function (times the scattering coefficient) at \mathbf{x} . $L_e(\mathbf{x}, -\hat{\omega})$ is the luminaire emission radiance at \mathbf{x} in direction $-\hat{\omega}$ and $T(\mathbf{x}, \mathbf{x})$ is the medium transmittance between \mathbf{x} and \mathbf{x} .

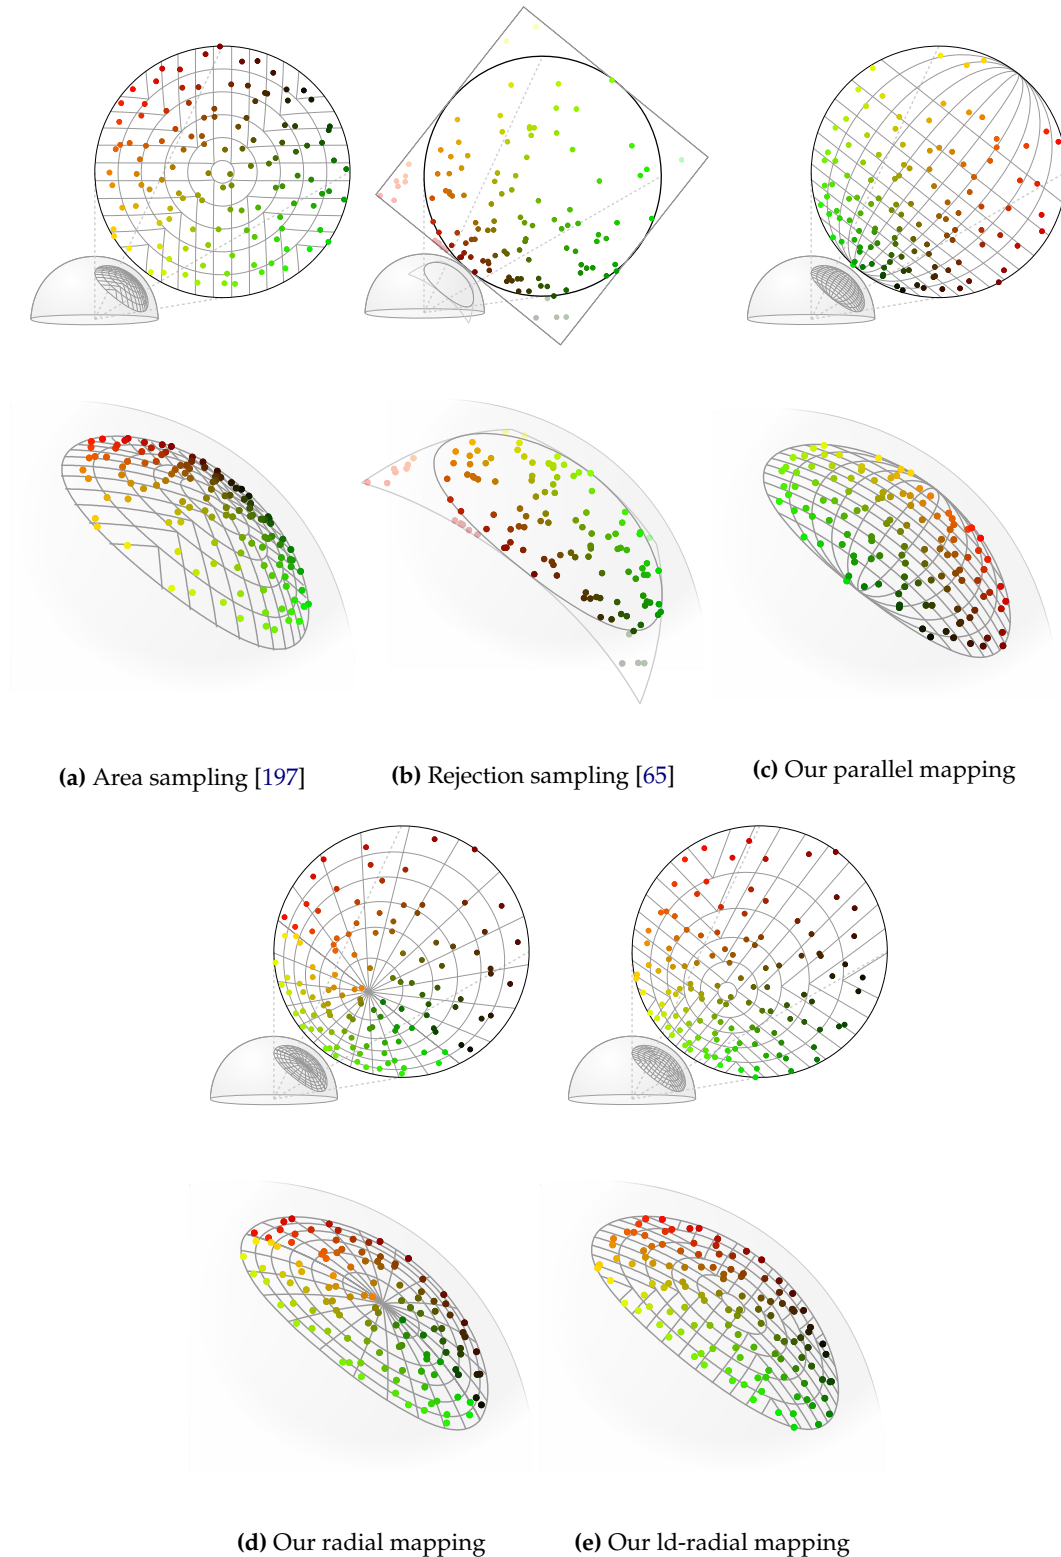


Figure 2.1: A stratified unit-square sample pattern mapped onto the surface of a disk, using existing techniques and our proposed maps (with solid angle projections on the bottom row). The points are colored according to their canonical $[0, 1]^2$ coordinates to illustrate the continuity of the maps. Gamito’s rejection sampling [65] does not allow for direct stratification, so we show the candidate low-discrepancy pattern for that case.

Solid angle sampling. Monte Carlo estimation of Equation (2.1) using N randomly sampled directions $\hat{\omega}_i$ has the following form:

$$L_s(\mathbf{x}, \omega_o) \approx \frac{1}{N} \sum_{i=1}^N \frac{f(\mathbf{x}, \mathbf{x}_{x\hat{\omega}_i}, \omega_o, \hat{\omega}_i)}{p(\hat{\omega}_i)}, \quad (2.2)$$

where $p(\hat{\omega})$ is the pdf for sampling $\hat{\omega}$. The choice of sampling density p is important, since a lower variation of f/p makes the estimator more efficient [198]. For disk lights the traditional choice is uniform density over the luminaire surface D . This *area sampling* technique is easy to implement and its resulting solid angle pdf is $p(\hat{\omega}) = \|\mathbf{x} - \mathbf{x}_{x\hat{\omega}}\| / (A(D) |\hat{\omega} \cdot \hat{n}_{x\hat{\omega}}|)$, where $A(D)$ is the area of D . This pdf can lead to very high variance in the radiance estimator (2.2), especially when the point \mathbf{x} is close to the luminaire. Our goal is to devise uniform *solid angle sampling* techniques that generate directions $\hat{\omega}$ with constant density $p(\hat{\omega}) = 1/|\Omega_D|$, yielding estimators with significantly lower variance than uniform area sampling.

Area-preserving mapping. Sample stratification can greatly improve the efficiency of Monte Carlo estimators [173; 196; 209]. Most existing stratification techniques generate samples in the canonical unit square $[0, 1]^2$, however our goal is to sample directions inside the solid angle Ω_D . Therefore, in order to take advantage of these techniques, we need to find a mapping M from $[0, 1]^2$ to Ω_D such that for any two regions $R_1, R_2 \subseteq [0, 1]^2$:

$$\frac{A(R_1)}{A(R_2)} = \frac{\mu(M(R_1))}{\mu(M(R_2))},$$

where A is the area measure, and μ is the solid angle measure as in Equation (2.1). We call such maps *area-preserving maps*. This key property makes it possible to generate stratified samples in Ω_D , because stratification is far more easily achieved in $[0, 1]^2$.

Area-preserving solid angle maps have been developed for triangles [8] and rectangles [223]. For sampling the solid angles of disks, Gamito [65] proposed to use a rectangle map [223] followed by rejection sampling. This technique cannot be used with fixed-size canonical point sets, and needs a low-discrepancy sequence capable of progressively generating stratified sample candidates. The rejection sampling also makes it very difficult to achieve good high-dimensional stratification in the presence of other distributed effects, e.g. volumetric scattering, where the coordination of the sample patterns of different effects is desired. In this work we focus on area-preserving maps for disks that do not require rejection sampling and work with any canonical sample pattern. Figure 2.1 compares our proposed maps against existing techniques.

For surface scattering points \mathbf{x} , an even better strategy is to importance sample the term $|\hat{\omega} \cdot \hat{n}_x|$ in the contribution f . Such uniform sampling of the *projected solid angle* has been described by Arvo [9] for polygonal lights. Extending our approach to projected solid angle sampling is an interesting avenue for future work.

2.3 Solid angle sampling of an oriented disk

We base our sampling techniques on the key observation that the projected area of any ellipse, including a disk, forms a spherical ellipse on the unit sphere around the shading point (Figure 2.2). Thus, in order to sample the solid angle subtended at




Figure 2.2: **Left:** The disk's local reference system $\mathcal{R}_d = (\hat{x}_d, \hat{y}_d, \hat{z}_d)$ and the local coordinates required to characterize its solid angle projection. **Center:** Projections of the relevant coordinates onto the unit sphere, defining the spherical ellipse and its local reference system $\mathcal{R}_e = (\hat{x}_e, \hat{y}_e, \hat{z}_e)$, where $\hat{x}_e \equiv \hat{x}_d$. **Right:** The spherical ellipse is defined by its semi-arcs α and β or, equivalently, by its semi-axes a and b in \mathcal{R}_e . The tangent ellipse (in red), which lies on a plane tangent to the sphere at \hat{z}_e (i.e. the spherical ellipse center), is defined by its semi-axes a_t and b_t .

point \mathbf{o} by an oriented disk with center \mathbf{c} , normal \hat{n} , and radius r , we will uniformly sample a point \mathbf{q} on the spherical ellipse and then backproject it to the disk.

Spherical ellipse. To compute the subtended spherical ellipse, we first define a local reference frame for the disk $\mathcal{R}_d = (\hat{x}_d, \hat{y}_d, \hat{z}_d)$:

$$\hat{z}_d = -\hat{n}, \quad \hat{y}_d = \hat{z}_d \times \frac{\mathbf{c} - \mathbf{o}}{\|\mathbf{c} - \mathbf{o}\|}, \quad \hat{x}_d = \hat{y}_d \times \hat{z}_d. \quad (2.3)$$

We then take the boundary disk coordinates y_0 and y_1 w.r.t. the \hat{y}_d axis and project them onto the sphere (Figure 2.2, left). From the coordinates y'_0, y'_1, z'_0, z'_1 of these projections (Figure 2.2, center) we can compute the spherical ellipse center: it is the result \hat{z}_e of normalizing the vector $(0, y'_h, z'_h)$, where $y'_h = (y'_0 + y'_1)/2$ and $z'_h = (z'_0 + z'_1)/2$. (Note that \hat{z}_e in general does not coincide with the spherical projection of the disk center \mathbf{c} .) Reprojecting \hat{z}_e onto the disk (Figure 2.2, left), the obtained y_h coordinate defines a chord $\overline{x_0x_1}$ parallel to \hat{x}_d . The chord endpoint projections onto the sphere, with \hat{x}_d -coordinates x'_0 and x'_1 , allow us to compute the lengths of the ellipse's semi-axes, a and b , and semi-arcs, α and β (Figure 2.2, right):

$$a = x'_1, \quad b = \frac{1}{2} \sqrt{(y'_1 - y'_0)^2 + (z'_1 - z'_0)^2}, \quad (2.4)$$

$$\alpha = \sin^{-1} a, \quad \beta = \sin^{-1} b. \quad (2.5)$$

Finally, from α and β we can compute the semi-axes $a_t = \tan \alpha$ and $b_t = \tan \beta$ of the ellipse tangent to the sphere at \hat{z}_e (Figure 2.2, right).

In the following, we use both the spherical and the tangent ellipses to derive two different mappings for uniformly sampling points \mathbf{q} on the spherical ellipse which we then map to the surface of the disk. These mappings operate in a coordinate system $\mathcal{R}_e = (\hat{x}_e, \hat{y}_e, \hat{z}_e)$, where $\hat{x}_e \equiv \hat{x}_d$ and $\hat{y}_e = \hat{z}_e \times \hat{x}_e$, shown in Figures 2.2 and 2.3.

2.3.1 Area-preserving mappings

Our new mappings are based on a generalization of the so-called *Archimedes Hat-Box theorem* [6]. It states that the area of a region on the sphere between two parallels is




Figure 2.3: Left: Any region (blue) on the unit sphere can be radially projected to another region (red) on a cylinder aligned with any axis (here \hat{x}). Any point \mathbf{q} on the sphere can be expressed in cylindrical coordinates as (ϕ, h, r) (azimuth angle, altitude, distance from center). This point can be mapped to a point \mathbf{p} on the unit cylinder with coordinates (ϕ, h) . The mapping preserves the areas of both finite and differential regions. Thus, to obtain a point on the spherical region, we can sample inside the cylindrical region and project back onto the sphere. **Center:** A spherical ellipse (blue), with center on the \hat{z}_e axis, projected onto a cylinder aligned with the \hat{y}_e axis. **Right:** The same spherical ellipse projected onto a \hat{z}_e -axis aligned cylinder. In this case, the projected region (red) has a ring-like shape.

equal to the area of that region’s projection onto a perpendicularly aligned bounding cylinder. This area-preserving property also holds for arbitrarily shaped regions on the sphere (Figure 2.3, left). The latter property was used by [218] to define compact metallic BRDFs (they provide a demonstration). It was also used (although not explicitly stated) by [8] and [223] to define area-preserving parameterizations for spherical triangles and rectangles, respectively.

We use this cylindrical projection property to derive our area-preserving mappings for a spherical ellipse centered on the \hat{z}_e axis. The ellipse can be radially projected onto a cylinder, obtaining a *cylindrical ellipse*. Two different unit-radius cylinders can be used. One is aligned with the \hat{y}_e axis (Figure 2.3, center), which we call a *parallel map*. The second one is aligned with the \hat{z}_e axis (Figure 2.3, right), which we call a *radial map*. We also propose a variant of the radial map that uses Shirley’s low-distortion map [197], which we call a *low-distortion radial map*, or *ld-radial map*.

Maps overview. The basic idea behind our maps is to first select a point \mathbf{p} on the cylindrical ellipse as a function of a canonical unit-square point $(\varepsilon_1, \varepsilon_2) \in [0, 1]^2$. We then project \mathbf{p} back onto the sphere perpendicularly to the cylinder axis (see Figure 2.3, left) to get the point \mathbf{q} . Let (ϕ, h) be the cylindrical coordinates of \mathbf{p} . We first obtain the azimuth angle ϕ by finding the lateral slice on the cylindrical ellipse whose solid angle is $\varepsilon_1 \Omega_D$ (Figure 2.4). With ϕ fixed, the altitude h is computed as a simple linear interpolation using ε_2 along the lateral line segment that is the intersection between the lateral plane at angle ϕ and the cylindrical ellipse (green line segment in Figure 2.4). The sampling of ϕ involves numerical inversion of incomplete elliptic integrals, as we show next.

2.3.2 Parallel Mapping

Our parallel mapping, whose cylinder axis is aligned with \hat{x}_e , operates by considering a portion (sector) of the cylindrical ellipse – the red-shaded region in Figure 2.4, left. This region is determined by the green line segment, whose endpoints have

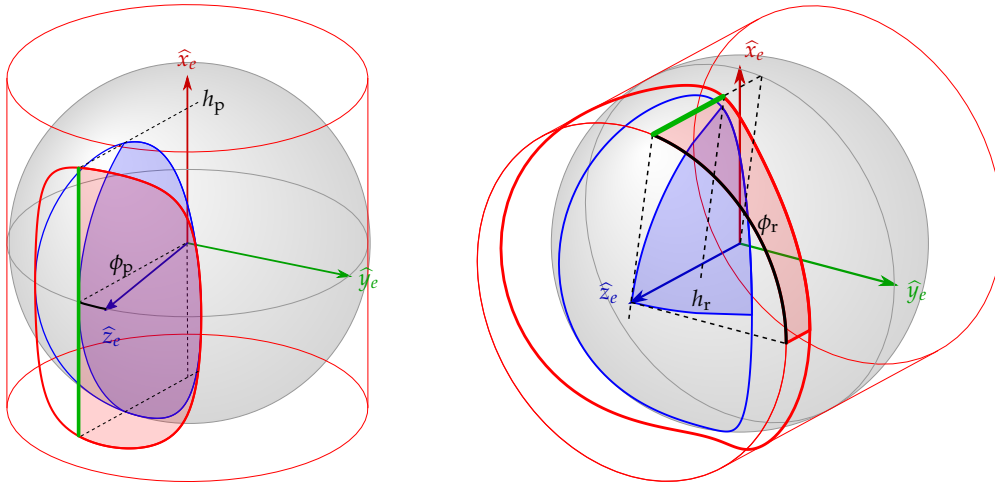


Figure 2.4: Illustration of our parallel (**left**) and radial (**right**) maps. Given a canonical sample $(\varepsilon_1, \varepsilon_2) \in [0; 1]^2$, we first find the azimuth angle ϕ_p , respectively ϕ_r , that cuts a region on the cylindrical ellipse with area $\varepsilon_1 \Omega_D$ (in red, determined by the green line segment). A sample point on the cylinder is then obtained by linearly interpolating the green segment endpoints using ε_2 . For the parallel map, the endpoint cylindrical coordinates are $(\phi_p, -h_p)$ and (ϕ_p, h_p) , with $\phi_p \in [-\beta, \beta]$. For the radial map, these coordinates are (ϕ_r, h_r) and $(\phi_r, 1)$, with $\phi_r \in [0, 2\pi]$ (we use $\phi_r \in [0, \pi/2]$ in each quadrant).

cylindrical coordinates (ϕ_p, h_p) and $(\phi_p, -h_p)$. The angle ϕ_p goes from $-\beta$ to β , since the spherical ellipse is centered on the \hat{z}_e axis.

Due to the Hat-Box theorem, the differential solid angle covered by the green segment is equal to its length, $2h_p$, which is in fact a function of ϕ_p . Thus, the solid angle subtended by the red region onto the spherical ellipse (the blue region in Figure 2.4, left) can be written as the integral of the segment length:

$$\Omega_p(\phi_p) = \int_{-\beta}^{\phi_p} 2h_p(\phi'_p) d\phi'_p, \quad (2.6)$$

where the full solid angle of the spherical ellipse is $\Omega_D = \Omega_p(\beta)$. Due to symmetry, for any angle $\phi_p \in [0, \beta]$ it holds $h_p(-\phi_p) = h_p(\phi_p)$. We use this to express $\Omega_p(\phi_p)$ as a sum of integrals $\Omega_p^+(\phi_p)$ over positive angles:

$$\Omega_p(\phi_p) = \begin{cases} \Omega_p^+(\beta) + \Omega_p^+(\phi_p) & : \phi_p \geq 0 \\ \Omega_p^+(\beta) - \Omega_p^+(-\phi_p) & : \phi_p < 0 \end{cases}, \quad (2.7)$$

where

$$\Omega_p^+(\phi_p) = \int_0^{\phi_p} 2h_p(\phi'_p) d\phi'_p. \quad (2.8)$$

In Appendix 2.A we derive an expression for $h_p(\phi_p)$:

$$h_p(\phi_p) = c_t \sqrt{\frac{1 - (p+1) \sin^2 \phi_p}{1 - (mp+1) \sin^2 \phi_p}}, \quad (2.9)$$

where

$$p = \frac{1}{b_t^2}, \quad m = \frac{a_t^2 - b_t^2}{a_t^2 + 1}, \quad c_t = \frac{a_t}{\sqrt{1 + a_t^2}}. \quad (2.10)$$

Substituting Equation (2.9) into Equation (2.8) and simplifying, we get:

$$\Omega_p^+(\phi_p) = \frac{2c_t}{b_t} \left[(1 - n) \Pi(n; \phi_p | m) - F(\phi_p | m) \right], \quad (2.11)$$

where $F(\phi | m)$ and $\Pi(n; \phi | m)$ are Legendre incomplete elliptic integrals of respectively the first and third kind, $m \in [0, 1)$, and

$$\phi_p = \sin^{-1} \left(\frac{\tan \phi_p}{b_t} \right), \quad n = -b_t^2. \quad (2.12)$$

Unfortunately, no closed-form expressions are known for $F(\phi | m)$ and $\Pi(n; \phi | m)$, so Equation (2.11) must be evaluated numerically.

Sampling. With the fractional spherical ellipse area Ω_p characterized, we can map a point on the unit square $(\varepsilon_1, \varepsilon_2) \in [0, 1]^2$ to a point on the spherical ellipse \mathbf{q} . We first need to find the angle ϕ_p that satisfies $\Omega_p(\phi_p) = \varepsilon_1 \Omega_D$, for which we need to evaluate the inverse function Ω_p^{-1} . This function has no analytical form, so we resort to numerically finding the roots of the equation

$$\Omega_p(\phi_p) - \varepsilon_1 \Omega_D = 0. \quad (2.13)$$

Having sampled ϕ_p , we get the point \mathbf{p} on the cylindrical ellipse by first computing $h_p(\phi_p)$ using Equation (2.9) and then linearly interpolating the altitude coordinate between $-h_p(\phi_p)$ and $h_p(\phi_p)$ using ε_2 :

$$\mathbf{p} = (\phi_p, (2\varepsilon_2 - 1) h_p(\phi_p)) = (\phi_p, h). \quad (2.14)$$

Finally, the corresponding point \mathbf{q} on the ellipse is obtained by radially projecting \mathbf{p} onto the sphere (see Figure 2.1a, left):

$$\mathbf{q} = \left(h, \sqrt{1 - h^2} \sin \phi_p, \sqrt{1 - h^2} \cos \phi_p \right). \quad (2.15)$$

Figure 2.1c shows the resulting map.

2.3.3 Radial mapping

The parallel mapping presented in Section 2.3.2 involves two elliptic integrals and introduces noticeable distortions (see the converging lines in Figure 2.1c), which can increase discrepancy and ruin any blue noise properties present in the input unit-square sample distribution. In this section we present an alternative radial mapping that uses a single elliptical integral and also exhibits less distortion. It is based on the analysis of the spherical ellipse topology by Booth [24].

We will exploit the fact that the four quadrants of the spherical ellipse are radially symmetric (see Figure 2.2, right), so its total area can be expressed as $\Omega_D = 4 \Omega_r$, with Ω_r being the area of each quadrant. Within a quadrant, the azimuth angle is $\phi_r \in [0, \pi/2]$.




Figure 2.5: Illustration of our polar mapping. Projected perpendicularly to its axis \hat{z}_e onto the $\hat{x}_e\hat{y}_e$ plane, the spherical ellipse forms a planar ellipse with semi-axes $a = \sin \alpha$ and $b = \sin \beta$. We use the Pythagorean theorem to express the altitude $h_r(\phi_r)$ of the cylindrical projection of the spherical ellipse's curve in terms of a , b , and ϕ_r .

We now consider a bounding cylinder aligned with the \hat{z}_e axis (Figure 2.4, right). Specifically, we are interested in the lateral region (in red in Figure 2.5) that is the radial projection of a fraction of the spherical quadrant. This region is determined by the position of the green line segment whose endpoints have cylindrical coordinates (ϕ_r, h_r) and $(\phi_r, 1)$. The segment length is $1 - h_r$, which is a function of ϕ_r . Similarly to Equation (2.6), we use the Hat-Box theorem to express the fractional quadrant area as the integral of this length:

$$\Omega_r(\phi_r) = \int_0^{\phi_r} [1 - h_r(\phi'_r)] d\phi'_r = \phi_r - \int_0^{\phi_r} h_r(\phi'_r) d\phi'_r. \quad (2.16)$$

Using the Pythagorean theorem, we express $h_r(\phi_r)$ as (see Figure 2.5)

$$h_r(\phi_r) = \sqrt{1 - r^2(\phi_r)}, \quad (2.17)$$

where $r(\phi_r)$ is the (planar) elliptical radius of the spherical ellipse with the following expression, which we derive in Appendix 2.B:

$$r(\phi_r) = \frac{ab}{\sqrt{a^2 \sin^2 \phi_r + b^2 \cos^2 \phi_r}}. \quad (2.18)$$

Plugging Equations (2.17) and (2.18) back into Equation (2.16), and using Booth's derivations [24], we can now express the fractional quadrant area $\Omega_r(\phi_r)$ using Legendre's incomplete elliptic integral of the third kind $\Pi(n; \varphi|m)$, so it becomes

$$\Omega_r(\phi_r) = \phi_r - \frac{b(1 - a^2)}{a\sqrt{1 - b^2}} \Pi(n; \varphi_r|m), \quad (2.19)$$

where

$$\varphi_r = \tan^{-1}\left(\frac{a_t}{b_t} \tan \phi_r\right), \quad n = \frac{a^2 - b^2}{a^2(1 - b^2)}, \quad m = \frac{a^2 - b^2}{1 - b^2}. \quad (2.20)$$

Above, φ_r is the parametric angle of the tangent ellipse, and n and m are the elliptic characteristic and module that characterize the elliptic integral [25].

Unfortunately, as with Equation (2.11), no closed-form expression is known for the general-case incomplete elliptic integral of the third kind, so we need to evaluate Equation (2.19) numerically.

Direct radial mapping. Having an expression for the fractional spherical ellipse Ω_r , we can map a unit-square point $(\varepsilon_1, \varepsilon_2) \in [0, 1]^2$ to a point \mathbf{q} on the spherical ellipse. Below we only consider sampling the first ellipse quadrant (shown in Figure 2.5); the entire ellipse can be covered by flipping the \hat{x}_e - and \hat{y}_e -coordinates of \mathbf{q} .

First, we need to find the azimuth angle $\phi_r \in [0, \pi/2]$ satisfying

$$\Omega_r(\phi_r) - \varepsilon_1 \Omega_r(\pi/2) = 0. \quad (2.21)$$

Since we do not have a method to analytically invert $\Omega_r(\phi_r)$, we compute ϕ_r by numerically finding the roots of the above equation.

Having sampled ϕ_r , we obtain point \mathbf{p} on the cylindrical ellipse by first computing $h_r(\phi_r)$ using Equation (2.17) and then linearly interpolating the altitude coordinate between $h_r(\phi_r)$ and 1 using ε_2 :

$$\mathbf{p} = (\phi_r, (1 - \varepsilon_2)h_r(\phi_r) + \varepsilon_2) = (\phi_r, h). \quad (2.22)$$

We find the corresponding point \mathbf{q} on the ellipse by projecting $\mathbf{p} = (\phi_r, h)$ using Equation (2.15) with swapped \hat{x}_e - and \hat{z}_e -coordinates. Figure 2.1d shows the resulting map.

Low-distortion radial mapping. As seen in Figure 2.1d, the direct mapping from above resembles the classical planar Cartesian-to-polar mapping. As such, it also exhibits the same distortion – the lines converging at the ellipse center, which does not preserve relative distances between samples and damages their stratification. In the planar case, the mapping of Shirley and Chiu [197] rectifies this distortion by warping concentric squares into concentric disks. To achieve the analogous mapping on the spherical ellipse, we first warp our input unit-square samples $(\varepsilon_1, \varepsilon_2)$ to the unit disk using Shirley and Chiu’s concentric mapping. We then move back to the unit square using the following *inverse polar mapping*:

$$u = \begin{cases} 2 \frac{\theta}{\pi} & : \theta \in [0, \frac{\pi}{2}) \\ 1 - 2 \frac{\theta - \pi/2}{\pi} & : \theta \in [\frac{\pi}{2}, \pi) \\ 2 \frac{\theta - \pi}{\pi} & : \theta \in [\pi, \frac{3\pi}{2}) \\ 1 - 2 \frac{\theta - 3/2\pi}{\pi} & : \theta \in [\frac{3\pi}{2}, 2\pi) \end{cases}, \quad v = r^2. \quad (2.23)$$

The result of this detour is a unit-square point set that, when warped using the classical (forward) planar polar mapping, gives Shirley and Chiu’s low-distortion concentric disk distribution. We, instead, feed this unit-square set to our direct radial mapping to get a concentric-like distribution on the spherical ellipse, which is shown in Figure 2.1e.

2.4 Implementation

We have implemented our maps as custom sampling procedures for disk lights in two different systems: the Mitsuba renderer [99] and the Arnold production renderer [56].

In order to sample from each mapping, we need to find the roots of Equations (2.13) and (2.21) respectively. Since the elliptic integrals they contain prevent analytical inversion, we resort to numerical root finding using an iterative Newton-Raphson method. However, this method can become very expensive, since for each iteration we need to numerically evaluate two and one incomplete elliptic integrals (for the parallel and radial mappings, respectively).

2.4.1 Tabulation

In order to reduce the significant cost of Newton-Raphson over area sampling (up to $10\times$ in simple scenes; see Figure 2.9) and avoid the expensive numerical inversion, we approximate our radial mapping by tabulating Equation (2.19). We choose to tabulate this mapping as it introduces less distortion in the output sample distribution than the parallel one, as shown in Figure 2.1.

We can write the fractional solid angle $\Omega'_r = \Omega_r(\phi_r)/\Omega_r(\pi/2)$ as a function of $\alpha \in [0, \pi/2]$, $\beta \in [0, \alpha]$ and $\phi_r \in [0, \pi/2]$. We can tabulate this function by discretizing each of the three parameters, producing a 3D array of values. A quick binary search based on ϕ_r (combined with interpolation) then allows us to get approximate values of Ω'_r with good accuracy. However, storing such a table would require a large amount of memory. To address this, we reparametrize Ω'_r in terms of α , the ratio $\beta' = \beta/\alpha$ (which is in $[0, 1]$), and ϕ_r . This version of Ω'_r has very low variation w.r.t. to α , so we can remove this parameter altogether, reducing the tabulation to a 2D array of Ω'_r values for a set of β' and ϕ_r values. Each entry in this array corresponds to a spherical triangle defined by the fraction of ϕ_r covered by the given entry and the value of $\theta(\phi_r)$ at the start of the entry's interval, which can be easily sampled [8]. This approximation causes some generated samples to lie outside the spherical ellipse, which we reject. Note that this rejection ensures unbiasedness; however, for practical reasons our production renderer implementation simply assigns zero weight to such invalid samples, resulting in a slight underestimation of the illumination. We have found the rejection ratio to be negligible, the storage requirement low, and the accuracy satisfactory for realistic rendering. In our implementation we use a 2D table with resolution 1024^2 , which we found to be accurate enough to provide an insignificant difference in variance compared to the analytic solution. Finally, note that in order to compute the samples' pdf $p(\hat{\omega}) = 1/|\Omega_D|$, we still need to compute Ω_D numerically. This computation is amortized among all samples for a given shading point.

2.4.2 Efficiency

Similarly to existing solid angle sampling techniques (e.g. for spherical triangles [8] and rectangles [223]), the cost of drawing a sample with our technique is higher than that of uniform area sampling (though some of it is amortized over multiple samples). This overhead pays off when the luminaire is close to the shading point (i.e. the subtended solid angle is relatively large).

As an optimization, our Arnold implementation (Figure 2.8) employs a simple heuristic to switch to uniform area sampling when the luminaire is far away. In order to provide a fair comparison against existing techniques (in Figures 2.6 and 2.7), our Mitsuba implementation does not take advantage of this optimization.

Even with the above tabulation scheme, we still need to compute the solid angle of the spherical ellipse Ω_D for the sampling pdf, using either Equation (2.11) or (2.19). The elliptical integrals involved can be computed using Carlson’s fast numerical algorithms [33].

2.5 Results

Figures 2.6 and 2.7 show a comparison between traditional area sampling [197], Gamito’s rejection-based solid angle sampling [65] and our techniques (Mitsuba implementation), without and with the presence of participating media. In both figures only direct illumination (single scattering) is computed, using 16 samples/pixel. Inside a medium, we first sample a point along the ray via equiangular sampling [129] w.r.t. the disk light center, and then use the corresponding disk sampling technique to generate a point on the light. In the case of uniform area sampling, a better strategy is to first sample the light surface and then perform equiangular sampling w.r.t. that chosen point. We therefore include this strategy in Figure 2.7 (called “Area sampling (first)”), which is incompatible with the solid angle mappings. The results show that our sampling methods outperform Gamito’s method on surfaces, and perform at least on par in participating media, where variance due to medium sampling dominates when using solid angle sampling. In all cases, area sampling yields much higher variance. In the supplemental document we provide global illumination comparisons between our tabulated radial sampling and Mitsuba’s built-in disk area sampling.

Figure 2.8 shows a scene rendered in Arnold, comparing our tabulated radial map implementation to the renderer’s built-in uniform area sampling. The scene features many production features, including high-resolution texture maps, fur, displacement, subsurface scattering, indirect surface and volume-to-surface light transport. In such cases the higher cost of our technique has a negligible impact on the total rendering performance. With 256 samples/pixel our tabulated radial map yields a noise-free image, while area sampling suffers from a substantial amount of noise.

Finally, Figure 2.9 shows a comparison between the convergence and the cost of the different techniques from Figures 2.6 and 2.7. For the same number of samples, our mappings produce lower error than Gamito on surfaces and perform virtually identically in media. In terms of cost, our tabulated version is almost as fast as area sampling, and the fully numerical implementation can be up to $10\times$ slower. Note that the performance of the iterative numerical inversion depends on the geometrical configuration: the starting point for the inversion affects the number of iterations required for convergence. The parallel and radial mappings take respectively 1-3 and 1-4 Newton-Raphson iterations in our tests. Also note that in scenes with higher geometric and shading complexity, the relative cost of the different methods has less impact on the overall rendering performance.

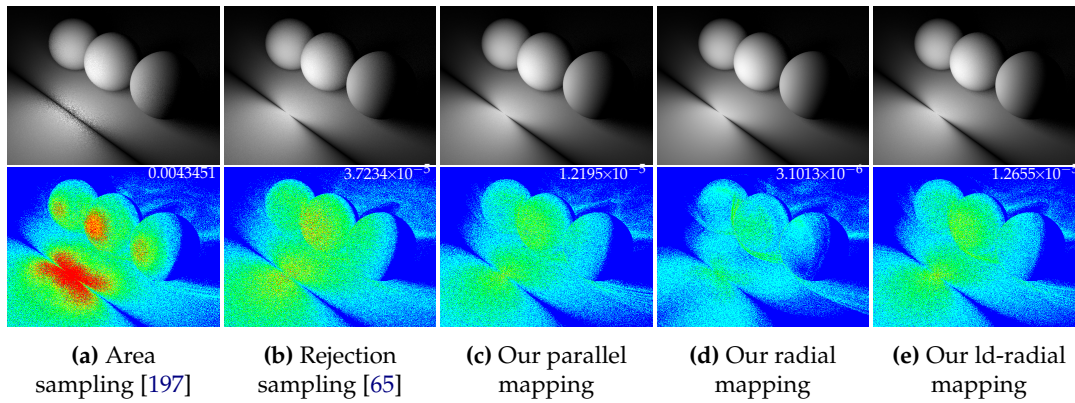


Figure 2.6: **Top:** A scene illuminated by a double-sided disk light, rendered with 16 samples/pixel. The light is perpendicular to the ground and invisible to camera rays. **Bottom:** False-color differences and MSE w.r.t. to a reference computed with 32K samples/pixel.

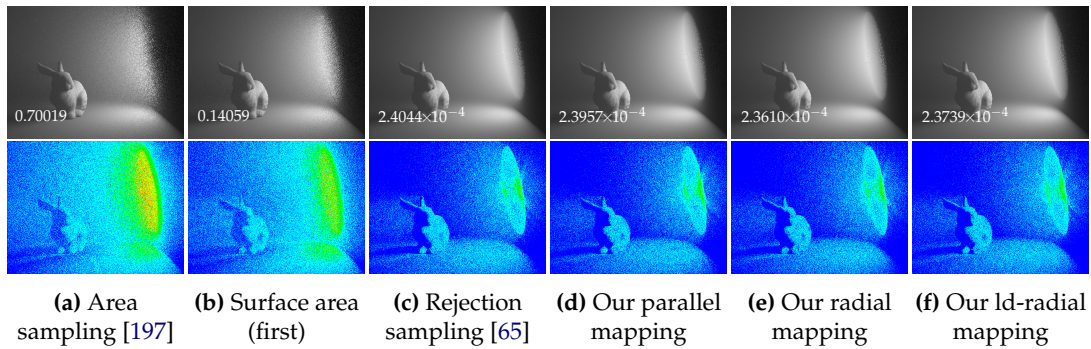


Figure 2.7: **Top:** A scene with participating media illuminated by a single-sided disk light (invisible), rendered with 16 samples/pixel. Please refer to Section 2.5 for details. **Bottom:** False-color differences and MSE w.r.t. to a reference computed with 64K samples/pixel.




Figure 2.8: Colored disk lights (invisible) rendered in Arnold, using area sampling (**left**) and our tabulated radial map (**right**) with 256 samples/pixel. Due to the use of complex surface and hair shaders, the higher cost of our technique has a negligible impact on performance.




Figure 2.9: Error (**top**) and cost (**bottom**) for the results shown in Figure 2.6 (**left**) and Figure 2.7 (**right**), w.r.t. sample count.

2.6 Conclusions

We have presented two new area-preserving mappings that enable the uniform solid angle sampling of oriented disk light sources. Following the key observation that this solid angle is a spherical ellipse, we make use of the Hat-Box theorem to transform canonical unit-square sample points onto the spherical ellipse in a way that preserves their stratification. To avoid costly numerical inversion, we develop a practical mapping tabulation that introduces little overhead over traditional uniform area sampling [197] while significantly reducing the variance of the illumination estimate. Our mappings are also competitive to existing disk solid angle sampling techniques [65], without imposing restrictions on the sample generator.

As a by-product of our work, we have proposed two new expressions for the subtended solid angle of a disk, which in addition to graphics is important in other fields such as particle transport. In this context, most previous analytic formulations [37; 167; 217] have included at least two incomplete elliptic integrals that need to be computed numerically. In contrast, our radial formulation, based on Booth’s spherical topology analysis [24], involves only one elliptic integral, making it more simple and practical than previous work.

While this work only considers circular disks, our approach could be extrapolated to other shapes whose subtended solid angle is also an ellipse, such as elliptical disks and ellipsoids [94]. Including these geometries would only require finding the spherical ellipses subtended by them. Moreover, our mappings could reduce variance of other shapes such as cylinders, following Gamito [65].

The main limitation of our method is the lack of analytical inversion of the proposed mappings, which requires using either costly numerical inversion or tabulation. Unfortunately, it seems impossible to find a spherical ellipse mapping that does not involve incomplete elliptic integrals, whose inversion is unknown. The presented mappings also only consider the solid angle, but not the other contribution terms in Equation (2.1), e.g. the BRDF or the foreshortening term. Developing methods for including at least some of these other terms is an interesting direction for future work. Furthermore, while our mappings are nearly optimal for uniformly emitting disk luminaires, it would be interesting to take into account spatially-varying emission profiles, in the spirit of the work of Bitterli et al. [20].

2.A Derivation of Equation (2.9)

Here we derive the expression for $h_p(\phi_p)$ in Equation (2.9), whose integral we then express as a combination of incomplete elliptic integral functions in Equation (2.11). We use the tangent ellipse, shown in Figure 2.10 and introduced in Section 2.3 and Figure 2.2, right. The ellipse semi-axes $a_t \geq b_t$ are aligned with \hat{x}_e and \hat{y}_e , respectively.

For any angle $\phi_p \in [-\beta, \beta]$, we first obtain a coordinate $y = \tan \phi_p$ along the \hat{y}_e axis. (We only consider $\phi_p > 0$, thus $y > 0$, and convert negative ϕ_p to positive using symmetry, as described in Section 2.3.2.) Using the ellipse equation $(x/a_t)^2 + (y/b_t)^2 = 1$, we can get the corresponding $x \geq 0$ coordinate along \hat{x}_e as a function of y :

$$x = a_t \sqrt{1 - (y/b_t)^2}. \quad (2.24)$$

We then consider the point $\mathbf{t} = (x, y, 1)$ on the tangent ellipse and its spherical projection $\mathbf{s} = \mathbf{t}/\|\mathbf{t}\|$. The \hat{x}_e -coordinate of \mathbf{s} , and also of its cylindrical projection (see Figure 2.4, left), is

$$h_p = \frac{x}{\sqrt{x^2 + y^2 + 1}}. \quad (2.25)$$

Substituting Equation (2.24) into (2.25):

$$h_p = a_t \frac{\sqrt{1 - py^2}}{\sqrt{y^2 + 1 + a_t^2(1 - py^2)}} = c_t \sqrt{\frac{1 - py^2}{1 - mpy^2}} \quad (2.26)$$

where p , m and c_t are as in Equation (2.10).

Using $y = \tan \phi_p$ and $0 \leq \phi_p \leq \beta \leq \pi/2$, in Equation (2.26) we can substitute y^2 by $(\sin^2 \phi_p)/(1 - \sin^2 \phi_p)$. With this we can finally write h_p explicitly as a function of ϕ_p :

$$h_p(\phi_p) = c_t \sqrt{\frac{1 - (p+1)\sin^2 \phi_p}{1 - (mp+1)\sin^2 \phi_p}} \quad (2.27)$$

which is exactly Equation (2.9).

2.B Derivation of Equation (2.18)

Here we derive the expression for $r(\phi_r)$ in Equation (2.18), which is used in the radial mapping (Section 2.3.3). We consider the planar ellipse resulting from the parallel projection of the spherical ellipse onto the $\hat{x}_e\hat{y}_e$ plane. This ellipse's semi-major and semi-minor axes are a and b , respectively (see Figure 2.5).

Consider a point \mathbf{r} whose polar coordinates $(\phi_r, r(\phi_r))$ and Cartesian coordinates (x, y) are related as

$$x = r(\phi_r) \cos \phi_r, \quad y = r(\phi_r) \sin \phi_r. \quad (2.28)$$

We want to define $r(\phi_r)$ in such a way that \mathbf{r} is on the planar ellipse curve. Thus, x and y must obey the ellipse equation, i.e.

$$\left(\frac{x}{a}\right)^2 + \left(\frac{y}{b}\right)^2 = 1. \quad (2.29)$$




Figure 2.10: View of the tangent ellipse (red). For a given angle ϕ_p , we first obtain the coordinate $y = \tan \phi_p$ along axis \hat{y}_e . Using the ellipse equation, we then find the corresponding coordinate x along \hat{x}_e . This gives point \mathbf{t} on the tangent ellipse, whose spherical projection \mathbf{s} has \hat{x}_e -coordinate h_p – the quantity we are interested in.

We can substitute x and y from Equation (2.28) into the ellipse equation, resulting in

$$r(\phi_r)^2 \left(\frac{b^2 \cos^2 \phi_r}{b^2 a^2} + \frac{a^2 \sin^2 \phi_r}{a^2 b^2} \right) = 1. \quad (2.30)$$

We can thus write

$$r(\phi_r) = \frac{ab}{\sqrt{b^2 \cos^2 \phi_r + a^2 \sin^2 \phi_r}}, \quad (2.31)$$

which is Equation (2.18).

2.C Numerical computation of incomplete elliptic integrals using Carlson forms

In this section we describe how to compute partial solid angles by directly using the functions introduced by Carlson [32]. With this method, expressions defining Ω_p^+ and Ω_r in terms of Incomplete Elliptic Integrals F and Π can be expressed as a combination of two Carlson functions R_F and R_J , which are defined as follows:

$$R_F(x, y, z) = \frac{1}{2} \int_0^\infty \frac{1}{\sqrt{(t+x)(t+y)(t+z)}} dt$$

$$R_J(x, y, z, q) = \frac{3}{2} \int_0^\infty \frac{1}{(t+q)\sqrt{(t+x)(t+y)(t+z)}} dt$$

values x, y and z must be non-negative, no more than one of them can be zero, and q must be non-zero. Obviously, parameters x, y and z can be exchanged, this is why these functions are called *symmetric forms*.

The first and third incomplete elliptic integrals (F and Π , respectively) can be written as combinations of R_F and R_J , as follows:

$$F(\varphi|m) = s R_F(x, y, 1) \quad (2.32)$$

$$\Pi(n; \varphi|m) = s R_F(x, y, 1) - \frac{1}{3} s^3 n R_J(x, y, 1, q) \quad (2.33)$$

where

$$s = \sin \varphi, \quad x = 1 - s^2, \quad y = 1 - ms^2, \quad q = 1 - ns^2. \quad (2.34)$$

These two equivalences are frequently used for computing F and Π , as there are quickly convergent series for approximating R_F and R_J [32]. In fact, this is the option used in the implementation of Boost library [23].

2.C.1 Evaluating $\Omega_p^+(\phi_p)$ using Carlson functions

The equation for $\Omega_p^+(\phi_p)$ in Section 2.3.2 (Equation (2.11)) can obviously be written as a combination of R_F and R_J , by using to (2.32). It holds

$$(1 - n) \Pi(n; \varphi_p|m) - F(\varphi_p|m) = w_F R_F(x, y, 1) + w_J R_J(x, y, 1, q) \quad (2.35)$$

with the weights w_F and w_J defined as:

$$w_F = -ns, \quad w_J = ns \frac{1}{3} (1 - n)s^2. \quad (2.36)$$

Note that, in this case

$$s = \sin \varphi_p = \frac{\tan \phi_p}{b_t} \leq 1, \quad n = -b_t^2. \quad (2.37)$$

These relations, together with $m < 1$, imply R_F and R_J parameters x, y, z and q obey these inequalities:

$$x = 1 - s^2 \geq 0 \quad (2.38)$$

$$y = 1 - ms^2 > 0 \quad (2.39)$$

$$z = 1 > 0 \quad (2.40)$$

$$q = 1 - ns^2 = 1 + \tan^2 \phi_p \geq 1, \quad (2.41)$$

thus $0 \leq x < y < z \leq q$, and both R_F and R_J can be safely evaluated.

The gain in efficiency by directly calling Carlson functions in Boost is due to the usage of two instead of three calls to Carlson functions implementations. Moreover, in this case $n < 0$, and Π implementation is fitted to the general case and thus needs to transform parameter n in a costly way, but this is not the case when directly using Carlson forms in our particular case.

We have used this to compute 10 millions values of expression (2.35), for a set of combinations of values of α, β and ϕ_p (with $0 \leq \phi_p \leq \beta \leq \alpha < \pi/2$). In each case we compare the computing time and resulting value obtained by using calls to Boost implementation of Π and F with the corresponding times and values by issuing direct calls to R_F and R_J . The values agree on 15 decimal digits on average (and a minimum of 13), with double precision, while computing time is halved by using direct calls

to Carlson functions (we have measured just time spent of Boost functions, not time devoted to precompute constant and parameters, which is negligible as compared to series approximation of Carlson functions).

2.C.2 Evaluating $\Omega_r(\phi_r)$ using Carlson functions.

Evaluating function Ω_r also involves computation of the incomplete elliptic integral of the first kind Π , as shown in Equation (2.19). In this case we also have two options: either to use a single call to an implementation of Π , or to use two calls to R_F and R_J implementation, by using (2.33).

We have compared the computing time and resulting values of both options, by using millions of evaluations uniformly distributed in the parameter space, with Boost library (in double precision), in a similar way as we have done for the parallel map. Resulting values agree in 12 decimal digits in the worst case. Regarding computing time, direct use of Carlson functions implementations takes around 89% of the time it takes to call Π implementation. In this case the time improvement is smaller than in the parallel map, but it is still noticeable.

We have also compared the average computing time for Ω_p^+ and Ω_r (the resulting values cannot be compared), in both cases by using Carlson functions implementations (the faster option). As a result, the average time it takes one evaluation of Ω_r is around 93% of the time it takes one evaluation of Ω_p^+ .

Chapter 3

A general framework for pearlescent materials

In this chapter we introduce a new technique for rendering materials manufactured with special effect goniochromatic pigments. Particularly, we simulate pearlescent pigments based on iridescent microscopic platelet-like flakes made up of several layers of dielectric materials, whose optical properties can be explained by phase shifts caused by the interreflection of light waves inside its structure. First, we present a thorough review of the properties of currently used pigments and manufacturing-related effects that influence pearlescence, then we propose a new model which expands the range of appearance that can be represented, and closely reproduces the behavior of measured materials, as we show in our comparisons.

This project started during an internship in the École Polytechnique Fédérale de Lausanne (EPFL) in Switzerland, hosted by Prof. Wenzel Jakob, and it was sparked by the inability of existing material models to fit real measurements of pearlescent and iridescent materials.

This work was originally published in ACM Transactions on Graphics and presented at SIGGRAPH Asia 2020.

A general framework for pearlescent materials

Ibón Guillén, Julio Marco, Diego Gutierrez, Wenzel Jakob & Adrian Jarabo
ACM Transactions on Graphics, Vol. 39(6), 2020

3.1 Introduction

Pearlescent materials have gained significant attention over the last decades. Despite being composed of relatively cheap materials, they are designed to exhibit a wide variety of attractive appearance, with a lustrous shade, metallic-like soft gloss, and vivid goniochromatic effects. These rich visual features have made these materials popular in the cosmetic industry since the 17th century [170], with a wide adoption in the 50's. More recently, the car and packaging industries have adopted them to create special-effect, luxurious looks [145]. In addition, they are currently starting to impact other fields such as clothing [144], or ink printing [176].

The intricate appearance of pearlescent materials is obtained by leveraging the optical properties of special pigments embedded into a hosting medium, called the container. These special pigments consist of oriented microscopic platelets with a layered structure of nanoscopic thickness. Given their scale and planar geometry, as



Figure 3.1: Synthetic objects rendered using our pearlescent material model, showing how changing a single parameter in each pair of bottles (as indicated by the brackets) leads to large, difficult-to-predict changes in appearance. From left to right, the changed parameters are: a) platelet density variation; b) substrate composition; c) deviation of the platelet normal distribution; d) deviation of the substrate thickness. The rightmost jar shows the values of our model parameters for its particular material. Please refer to Table A3 for a complete description of all the materials.

light scatters within the platelets it experiences interference effects. This results in highly directional and vivid colors.

Simulating the appearance of pearlescent materials is a complex task. A complete model needs to take into account the volumetric nature of such materials, including light transport inside the container, as well as volumetric scattering and complex anisotropic, interference-based light scattering involving the platelets. On top of that, the final appearance is highly dependent on the particular composition and structure of the platelets in the material, as well as their macroscopic orientation and distribution. This results in a highly non-linear appearance, which is difficult to predict and model. Figure 3.1 shows several examples of these strong, non-linear effects that arise by changing a single parameter in a pearlescent material.

Some existing works have proposed models of pearlescent materials, targeting the particular case of car paints [53; 54]. However, these models are rather limited in the range of appearance that can be simulated, since they are based on many simplifying assumptions that reduce their parameter space.

In this work, we present a general model for pearlescent materials. We review the existing literature from manufacturers (e.g., [144–146; 171]), which allows us to extract meaningful parameters that define the structure and optical properties of such materials. From this analysis, we develop a general radiative model for representing these materials, including a rigorous mathematical model for light scattering in platelets. Based on this model, we further analyze the gamut and behavior of the goniochromatic appearance of pearlescent materials, as their defining parameters change. We found that the angular effect of chromaticity lies on a narrow manifold, whose shape is directly related to the platelet reflectivity (gloss component) and transmittance (diffuse component). Furthermore, our analysis helps establish connections between low-level manufacturing parameters, and their effect on the material appearance.

We compare against real-world data, and show that our model reproduces ground-truth reflectance measurements more faithfully than prior work. Our model is fully spectral to avoid spectral aliasing, stochastic, and easy to integrate in any Monte Carlo renderer with multispectral support. Beyond the field of computer graphics, our work has potential applications in manufacturing: for instance, it might allow inverse methods to optimize appearance from actual fabrication parameters. This in turn has strong potential to enable predictive design of pearlescent materials in industries such as cosmetics, fine printing, effect paints, or packaging.

3.2 Related work

Volumetric materials. Light transport simulations involving volumetric materials typically build on the radiative transfer equation (RTE) [35], and its generalizations to anisotropic [100] and correlated [21; 104] media. Anisotropy in volumetric materials is generally modeled using *microflakes* [95; 250]. Wave effects such as polarization [105], or speckle [12] have also been incorporated into the radiative transfer framework. Our work models pearlescent polymers as a stack of dielectric layers that are each filled with an anisotropic medium that reproduces the behavior of the container and platelets.

Stratified materials. The complexity of real-world appearance is in part due to its stratified (layered) structure. Hanrahan and Krueger [85] proposed a general model for layered materials, based on expensive subsurface scattering computations, ignoring rough dielectric boundaries. Donner and Jensen [45] developed a significantly faster solution leveraging the diffusion approximation of light transport. Stam [205] generalized Hanrahan and Krueger’s work taking into account rough boundaries, in the context of skin rendering. A later model [101] proposed to handle arbitrary layer stacks using the adding-doubling method, handling all-frequency isotropic scattering. Zeltner and Jakob [249] then generalized this work to include anisotropic scattering. All of these methods require expensive precomputation. Other works [14; 77; 229; 242] proposed faster analytical solutions by stacking a set of BSDFs encoding the different effects of light transport. However, they rely on some approximations, and impose hard limitations on the type of materials that can be handled. Guo et al. [78] introduced an efficient Monte Carlo strategy for sampling the effective BRDF of a layered material that we also use in our work. Our work is based on precise modeling of light transport in stratified, pearlescent materials.

Wave-based scattering. Several works take into account relevant wave effects including diffraction-aware BSDFs [44; 57; 97; 204; 219; 232; 243], phase functions (based on Mie scattering) [62; 188], or goniochromatic patterns caused by birefringence [207]. Goniochromatic effects due to electromagnetic interference have been simulated for single-layer thin coatings [15; 69; 71; 128; 201; 213], and multiple-layer thin coatings [96; 212]. In our work we model the scattering in individual iridescent platelets by using multiple thin coatings; however, as opposed to the works by Sun et al. [212] and Hirayama et al. [96], we compute the exact reflectance and transmittance at run-time, without any precomputation.

Pearlescent Materials. The work of Gondek et al. [69] is the first to consider the problem of rendering pearlescent materials. This technique used a tabulated pre-computation of the appearance of such materials by brute-force simulation over explicit geometry. Rump et al. [186] and Kim et al. [122] used a data-driven representation of pearlescent paints based on bidirectional texture functions (BTF), limited to the reproduction of previously measured samples. Guo et al. [76] assumed plastic strata as a composite of coated discrete microfacets, omitting important effects like the global color shift caused by multiply scattered light. Volumetric models [53; 54] include single and multiply scattered light in a radiative framework, but require expensive precomputations, and impose assumptions on the distribution of the iridescent platelets and their composition, focusing on the particular pigments used in car paint. In contrast, we develop a general model that requires no precomputation, allows spectral rendering with minimal overhead, lifts most assumptions regarding the thickness of the platelets, and supports arbitrary platelets orientation including anisotropy and rotated platelets. Moreover, we show that our model provides a better match to real-world captured materials.

Several works in the optics community have characterized the directionally-varying reflectance gamut of pearlescent materials in terms of the incident and view angles, but they did not consider links to the underlying physical properties of the pigments. For instance, Medina [153] analyzed the spectral goniochromaticity of pearlescent materials by using principal component analysis. Ferrero et al. [60] measured a set of real-world pearlescent car paints; later, Ferrero et al. [59] extended this to diffraction-based pigments. We rely on their measured materials as a source of ground-truth reflectance data to conduct a comparative evaluation of several models in Section 3.7.1.

3.3 Pearlescent materials

In contrast to other diffractive effects [42; 204], or thin-film iridescence [15], pearlescence is the result of both volumetric absorption and electromagnetic interference effects in small layered, platelet-shaped structures with an average thickness of just a few nanometers. These *platelets* are uniformly distributed inside a base medium called the *container*, which creates the illusion of increased depth [145].

Figure 3.2 (a) shows a scanning electron microscope (SEM) capture of the cross-section of a manufactured pearlescent plastic, along with a magnified view of a single platelet (c).

Platelets. Platelets are made of transparent or semitransparent materials. They consist of a single optically homogeneous thin layer (*substrate-free* platelets), or form layered structures (*substrate-based* platelets). Substrate-free platelets have an optical thickness in the order of one fourth of the wavelength of visible light. They can be made up of single crystals, such as bismuth oxychloride (BiOCl), or polycrystalline platelets, made e.g. of titanium dioxide (TiO₂). However, these pigments are not commonly used in practice given the difficulty of crystallizing as thin platelets.

Substrate-based platelets, on the other hand, are much easier to fabricate. The substrate is a thin layer, typically between 100 and 1000 nm, made of a material with a low index of refraction. Approximately 98% of all fabricated substrates are transparent mica [146]. Other refractive materials with a higher index of refraction can

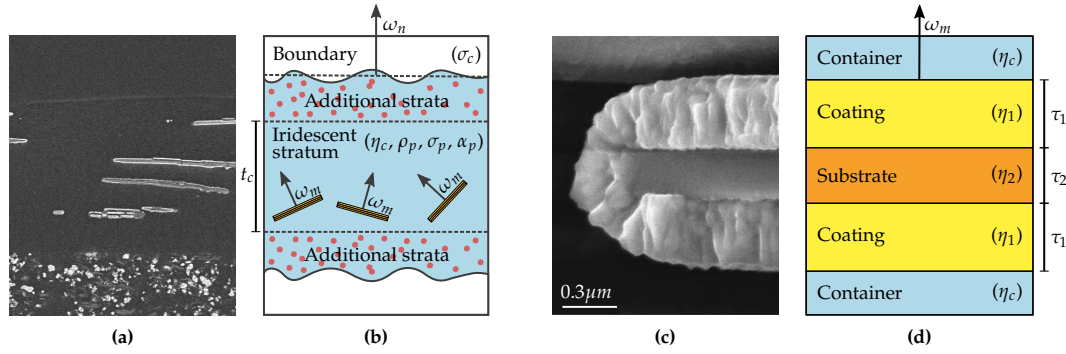


Figure 3.2: (a) Electronic microscopic capture of a manufactured plastic, where the embedded iridescent platelets can be observed. (b) Illustrative scheme of our model for pearlescent materials, where the layers and platelets are shown, along with their key parameters. Note that our model supports optional additional strata, which might also be pearlescent. (c) Electronic microscopic capture of a silica platelet. (d) Schematic view of our basic platelet model. Our model generalizes to include an arbitrary number of parallel coating layers, bounded at the top and bottom by the container media (see Figure 3.3 for more complex platelet structures). Image (a) from: [185]; used with permission. Image (c) from: [145]; used with permission.

be precipitated onto it, forming a coating layer around the substrate. The large difference in index of refraction between the substrate and the coating layer maximizes the desired interference effects. The most commonly used material for precipitation is titanium dioxide (TiO_2), which is almost transparent; this maximizes the color resulting from iridescence, while having minimum light loss. Some particular effects require the application of several levels of coating over the substrate [206]. Please refer to Maile's work [145] for more details about their chemical composition and fabrication processes.

Container. The container is usually manufactured with a cheap thermoplastic resin like polyethylene terephthalate (PET) in the case of plastics, or acrylic polyurethane for car paints. Among other mechanical reasons, these materials are chosen to maximize transparency with negligible absorption and scattering. The container may also be colored by adding wavelength-dependent absorbing dyes. This container is mixed with the platelets (which in general represent a small fraction of the total composite [227]). To achieve controllable and reproducible appearance, the thickness distribution of the platelets in the container must be narrow. During fabrication, the platelets get oriented roughly in parallel to the surface of the container as a result of stretching. However, this orientation can be altered during molding [41], yielding a broader normal distribution that alters the brightness of the surface [193].

In the following, we first describe our model of light scattering by individual iridescent platelets (Section 3.4). The scattering function derived from such model will then be incorporated into our global model for pearlescent materials (Section 3.5).

3.4 Scattering by iridescent platelets

Our platelet model consists of a stack of locally planar, parallel thin layers with normal ω_m (as seen in Figure 3.2d). Given that platelets are extremely thin, with thickness varying between 100 and 1000 nm [146], we assume that they are infinite in the horizontal domain. In addition, we build on the far-field approximation, assuming

an incoming planar field on the platelet. Each layer j in a platelet is characterized by its thickness τ_j and a wavelength-dependent complex index of refraction $\eta_j(\lambda)$.

Scattering function. The scattering function $\mathcal{F}_p(\omega_i, \omega_o, \lambda)$ of a platelet for wavelength λ is defined as the sum of two Dirac delta functions for reflection and transmittance

$$\mathcal{F}_p(\omega_i, \omega_o, \lambda) = \mathcal{R}(\omega_i, \lambda) \frac{\delta(\omega_m - \omega_h)}{4|\omega_h \cdot \omega_i|} + \mathcal{T}(\omega_i, \lambda) \delta(\omega_i - \omega_o), \quad (3.1)$$

where $\omega_h = (\omega_i + \omega_o)/|\omega_i + \omega_o|$ is the half vector, $|\cdot|$ represents the absolute dot product, and $\mathcal{R}(\omega_i, \lambda)$ and $\mathcal{T}(\omega_i, \lambda)$ quantify the proportions of the reflected and transmitted light in the platelet, respectively. The delta function $\delta(\omega_i - \omega_o)$ in the transmittance is consequence of the symmetric parallel structure of the platelet, which reverts the effect of the Snell's law until it cancels out at light leaving the platelet. Given the thinness of the platelet layers, we need to take into account wave optics for computing Equation (3.1). For substrate-free platelets, $\mathcal{R}(\omega_i, \lambda)$ and $\mathcal{T}(\omega_i, \lambda)$ can be computed using Airy summation, as done for thin-layer interference models [15; 201]. However, substrate-based platelets require a significantly more complex model, which we describe in the following subsection.

3.4.1 Scattering in substrate-based platelets

In the general case of substrate-based platelets, one or more coating layers surround the substrate, resulting in an N-layered structure (see Figure 3.2d). For a given direction ω_i and wavelength λ , the reflectivity $\mathcal{R}(\omega_i, \lambda)$ and transmissivity $\mathcal{T}(\omega_i, \lambda)$ of the N-layered platelet in Equation (3.1) are given by a plane-parallel solution of Maxwell's equations, tracking the influence of each layer on both amplitude and phase in order to express the effects of interreflection in terms of constructive and destructive interference. This solution can be efficiently evaluated using the *transfer matrix method* [26; 246]. The response of the layered structure for each polarization component of the electromagnetic field is given by the 2×2 *response matrix* M as

$$M = \prod_{j=1}^N M_j = \begin{pmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{pmatrix},$$

where M_j represents the individual response of each layer j . It is defined as (see [26, Ch.1.6.2] for details)

$$M_j = \begin{pmatrix} \cos \Delta\phi_j & -\frac{i}{q_j} \sin \Delta\phi_j \\ -i q_j \sin \Delta\phi_j & \cos \Delta\phi_j \end{pmatrix}, \quad (3.2)$$

where $\Delta\phi_j = 2\pi\lambda^{-1}\eta_j(\lambda)\tau_j \cos \theta_j$ is the phase shift of the incident light inside the layer, and θ_j is the transmission angle. The matrix M_j is different for each polarization component, with the tilted admittance q_j taking values $q_{\perp,j} = \eta_j(\lambda) \cos \theta_j$ and $q_{\parallel,j} = \cos \theta_j / \eta_j(\lambda)$ for the perpendicular and parallel polarization components, respectively. The value of $\cos \theta_j$ for each layer depends on the corresponding cosine of the previous layer, following Snell's law, with $\cos \theta_0 = \omega_m \cdot \omega_i$. To take into account the boundaries with the container, we consider two additional layers (top and bottom, see Figure 3.2) with $\eta_0 = \eta_{N+1} = \eta_c$. Given the symmetry of index of refraction changes in the layered structure, we have $\cos \theta_{N+1} = -\cos \theta_0$.

From the elements of the response matrix M , we obtain the complex reflection and transmission coefficients as

$$r = \frac{(m_{11} + m_{12} q_c) q_c + (m_{21} + m_{22} q_c)}{(m_{11} + m_{12} q_c) q_c - (m_{21} + m_{22} q_c)},$$

$$t = \frac{2 q_c}{(m_{11} + m_{12} q_c) q_c - (m_{21} + m_{22} q_c)},$$

where q_c refers to the tilted admittance of the container. Both M and q_c have to be obtained for the perpendicular (r_{\perp}, t_{\perp}) and parallel ($r_{\parallel}, t_{\parallel}$) components of the field. Finally, assuming unpolarized incident light, we compute the total reflectivity and transmissivity of the layered structure as

$$\mathcal{R}(\omega_i, \lambda) = \frac{1}{2} (|r_{\perp}|^2 + |r_{\parallel}|^2), \quad (3.3)$$

$$\mathcal{T}(\omega_i, \lambda) = \frac{1}{2} (|t_{\perp}|^2 + |t_{\parallel}|^2). \quad (3.4)$$

3.4.2 Analysis

Having presented the platelet, we showcase the complex nonlinear dependence of their reflectance properties on the layer structure before moving on to the full volumetric model in Section 3.5. Figure 3.3 shows the scattering behavior of different platelets using our model described in Section 3.4.1. We take titanium dioxide (TiO_2) as the main coating material since it is one of the most commonly used. All the structures shown (columns) are used in real applications such as cosmetics, plastic manufacturing, or car paint [171; 193; 195]. The only exception is the first column, which presents a simple substrate-free, TiO_2 platelet included for illustration purposes. In particular, each column shows the following:

1. Substrate-free, TiO_2 platelet
2. Mica substrate coated with TiO_2
3. Mica substrate coated with TiO_2 , in a PET container
4. Mica substrate coated with ferric oxide (Fe_2O_3) and TiO_2 , in a PET
5. Silica (SiO_2) substrate coated with TiO_2
6. Silica (SiO_2) substrate coated with TiO_2 , in a PET container
7. Aluminum substrate coated with TiO_2 , in a PET container
8. Aluminum substrate coated with TiO_2 , in a polyurethane container

The thickness of the TiO_2 layer increases from top to bottom in each column. Each row shows the directional reflectance with respect to the incident angle for ten uniformly distributed wavelengths (360-830 nm, top plots), as well as the resulting directional reflectance and transmittance in sRGB under a D65 illuminant. We do not plot the directional transmittance to avoid cluttering the plots; since mica and silica have a negligible absorption, transmittance curves are symmetric to the reflectance curves. Introducing aluminium and ferric oxide in the last two platelets adds some absorption, which breaks this symmetry.

This analysis reveals that multiple parameters have a large impact on the final appearance of the material. First, the effect of the coating thickness is very relevant.

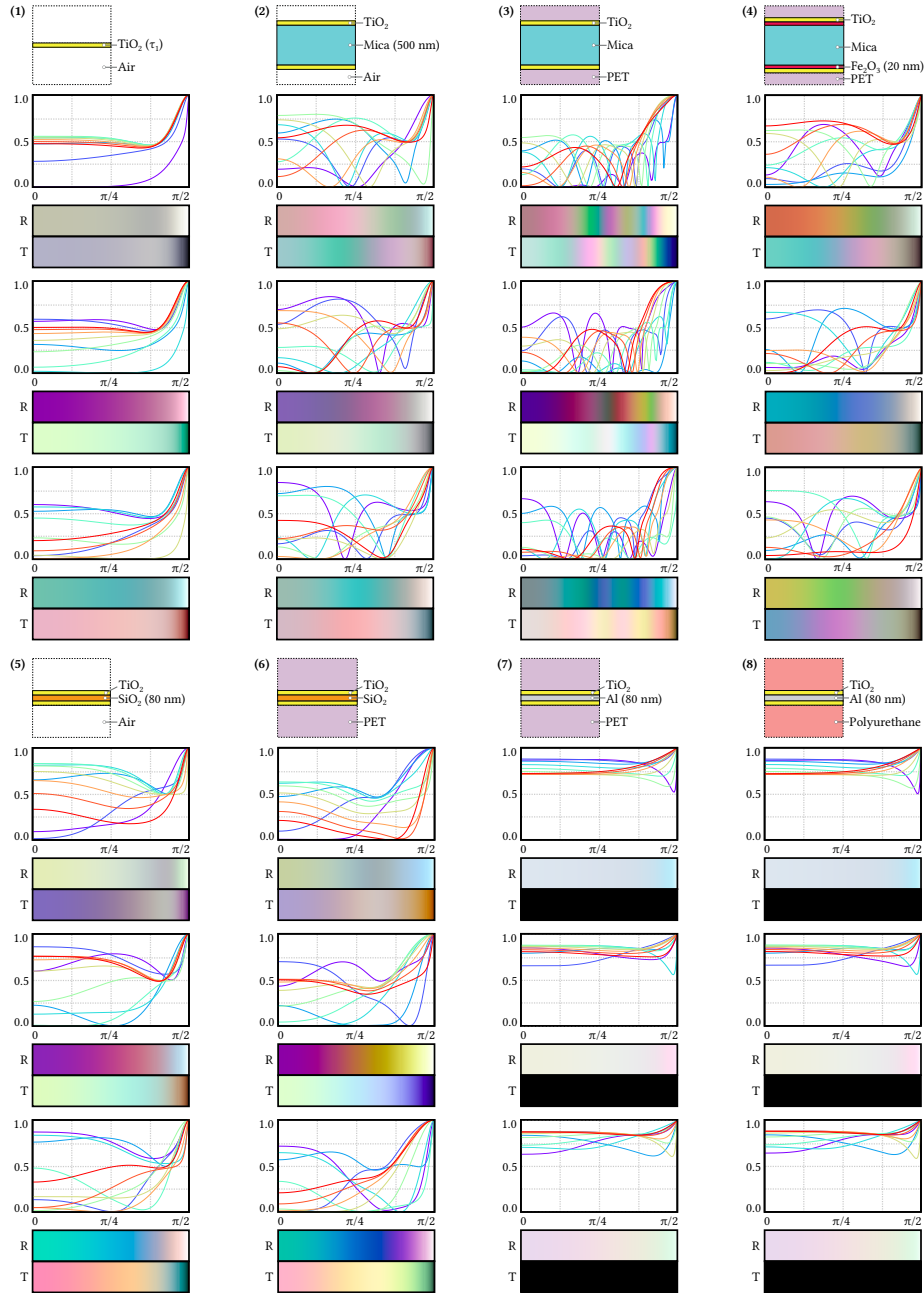


Figure 3.3: Scattering functions for different platelets. **(1)** Substrate-free platelet made of a single TiO_2 layer. **(2)** Mica substrate coated with TiO_2 . **(3)** Mica substrate coated with TiO_2 , in a PET container. **(4)** Mica substrate coated with ferric oxide (Fe_2O_3), and TiO_2 , in a PET container. **(5)** Silica (SiO_2) substrate coated with TiO_2 . **(6)** Silica (SiO_2) substrate coated with TiO_2 , in a PET container. **(7)** Metallic (aluminum) substrate coated with TiO_2 , in a PET container. **(8)** Metallic (aluminum) substrate coated with TiO_2 , in a polyurethane container. The mica substrates are 500 nm thick; silica and aluminum substrates are 80 nm; ferric oxide is 20 nm. The thickness of the TiO_2 coating layers varies in each column, from 60 nm (top), 100 nm (middle), and 140 nm (bottom). For each resulting platelet, the plots show the directional reflectance with respect to the incident angle for ten uniformly distributed wavelengths within the visible spectrum (360-830 nm, top plots), while the color gradients represent the resulting directional reflectance and transmittance in sRGB under a D65 illuminant.




Figure 3.4: Varying the substrate thickness leads to changes in platelet appearance. From left to right, mica substrate with thickness $\tau_s = 400$ nm, 500 nm and 600 nm, respectively (fixed TiO_2 coating of $\tau_1 = 60$ nm). Top: directional, wavelength-dependent reflectance curves. Bottom: reflectance and transmittance in sRGB under a D65 illuminant (similar to Figure 3.3).

Second, the substrate material also affects the platelet’s scattering, even for the same coating thickness (e.g., fifth and sixth columns): Transparent substrates such as mica or SiO_2 yield a strong transmittance while highly absorbing substrates such as aluminum are opaque. Third, the reflectance and transmittance gradients illustrate the rich variations in appearance of these materials, resulting from the combination of both. And fourth, the container also plays an important role in the final appearance (e.g., second and third columns; plastics generally use PET as container, while powder-based cosmetics use no container).

For artificial substrates such as SiO_2 , thickness can be carefully controlled within an error of ± 5 nm, and it is usually set to 80 nm. In contrast, for natural mica substrates τ_s is very difficult to control, resulting in a considerable thickness variation within the same material. Figure 3.4 illustrates the effect of these variations on the reflectance and transmittance.

3.5 Light transport in pearlescent materials

In this section we show how to incorporate the scattering function for a single platelet (Section 3.4) into our global model for pearlescent materials, and we conclude it with a discussion of important implementation-level details. As described in Section 3.3, pearlescent materials can be composed of a single or multiple strata (Figure 3.2b illustrates the simple case of a single-stratum material). The appearance of these materials is the result of the volumetric radiative transport in the container and the platelets of each stratum. The platelets can be assumed to be uncorrelated and randomly distributed throughout their medium [185], so non-exponential radiative transport [21; 104] does not need to be considered. Furthermore, we assume that wave-optical coherence effects play no role given the comparably large distance and random distribution of lengths between medium interactions.

In contrast to regular radiative transfer, which assumes isotropic spherical particles, pearlescent materials are composed of platelets that break the spherical symmetry assumption of traditional radiative transfer theory. As a consequence, we rely on the anisotropic formulation of the radiative transfer equation (RTE) introduced by Jakob et al. [100]. This formulation adds directional and positional dependency to all the relevant properties of the media, and couples them directly with the geometry of the scattering particles. In the absence of source terms, which can be removed in

pearlescent materials, this framework models the radiance L in direction

$$\omega_o \cdot \nabla L(\omega_o) + \Sigma_t(\omega_o)L(\omega_o) = \Sigma_s(\omega_o) \int_{\Omega} f_p(\omega_i \rightarrow \omega_o)L_i(\omega_i)d\omega_i, \quad (3.5)$$

where $\Sigma_t(\omega_o)$ and $\Sigma_s(\omega_o)$ are the directionally-dependent extinction and scattering parameters respectively, Ω is the sphere of directions, $f_p(\omega_i \rightarrow \omega_o)$ is the phase function, and ω_i is the incoming direction. Note that we have omitted the spatial and spectral dimensions for simplicity.

3.5.1 Modeling optical properties of pearlescent media

The appearance of a pearlescent medium is characterized by $\Sigma_t(\omega_o)$, $\Sigma_s(\omega_o)$, and $f_p(\omega_i \rightarrow \omega_o)$ in Equation (3.5). These three parameters depend on the optical parameters describing the container, including the presence of absorbing dyes, as well as the mixture of suspended platelets in it. In particular,

$$\Sigma_t = \Sigma_t^c + \sum_{j=1}^M \Sigma_t^{P_j}, \quad (3.6)$$

$$\Sigma_s = \Sigma_s^c + \sum_{j=1}^M \Sigma_s^{P_j}, \quad (3.7)$$

$$f_p(\omega_i \rightarrow \omega_o) = \frac{\Sigma_s^c}{\Sigma_s} f_p^c(\omega_i \rightarrow \omega_o) + \sum_{j=1}^M \frac{\Sigma_s^{P_j}}{\Sigma_s} f_p^{P_j}(\omega_i \rightarrow \omega_o), \quad (3.8)$$

where we use the superscripts c and p_j to refer to the container and platelet j respectively.

Container. To emphasize the iridescent properties of the material, the container is typically chosen to have minimum scattering, exhibiting a mostly transparent appearance. Therefore, we can assume negligible scattering, so that their optical properties are defined only by their complex index of refraction (IOR) η_c . The base container material is almost perfectly transparent, leading to a very small imaginary part in η_c . However, it is common to add some colouring dye. Its extinction, scattering, and phase function (Σ_t^d , Σ_s^d , and $f_p^d(\omega_i \rightarrow \omega_o)$, respectively) can be computed using Lorentz-Mie theory [62], parameterized by the concentration of pigment particles ρ_d , the particle size distribution $D(r_d)$, and the complex index of refraction $\eta_d(\lambda)$. In summary, the optical parameters of the container for a given wavelength λ are

$$\Sigma_t^c = 4\pi \text{Im}(\eta_d(\lambda)) \lambda^{-1} + \Sigma_t^d, \quad (3.9)$$

$$\Sigma_s^c = \Sigma_s^d, \quad (3.10)$$

$$f_p^c(\omega_i \rightarrow \omega_o) = f_p^d(\omega_i \rightarrow \omega_o), \quad (3.11)$$

where the first term in Equation (3.9) is direct consequence of the Beer-Lambert law. Note that these optical parameters do not have angular dependence.

Platelet stratum. The global optical properties of the platelet stratum are determined by the density of platelets ρ_p , their projected area $\sigma_p(\omega_i)$ in direction ω_i , and the directional distribution of their normals $D(\omega_m)$. These parameters define the

extinction and scattering coefficients as

$$\Sigma_t^p(\omega_i) = \rho_p \sigma_p(\omega_i), \quad (3.12)$$

$$\Sigma_s^p(\omega_i) = \alpha_p(\omega_i) \rho_p \sigma_p(\omega_i), \quad (3.13)$$

where $\sigma_p(\omega_i) = \int_{\Omega} \langle \omega_m, \omega_i \rangle D(\omega_m) d\omega_m$ [100], with $\langle \cdot, \cdot \rangle$ the clamped dot product, and $\alpha_p(\omega_i)$ the scattering albedo of the platelets. For a large number of real-world pearlescent platelets the absorption can be considered negligible, therefore $\alpha_p(\omega_i) = 1$. We parameterize the density of platelets as the fraction of platelets with respect to the total volume $C_p = \rho_p V_p$, with V_p the average volume of the platelets. We set $V_p = 400 \mu m^3$, based on SEM measurements [145]. In addition, the thickness of the platelets substrate $\tau_{p,s}$ might vary significantly (where s indicates the substrate), especially in natural substrates: We model this variability by using a substrate thickness distribution function $D(\tau_{p,s})$. On the other hand, the thickness of the coating layers can be carefully controlled given the chemical processes involved in coating [145]. This results into minimal variations on the platelet's coating thicknesses.

Since $D(\tau_{p,s})$ and $D(\omega_m)$ are uncorrelated, and using Equation (3.1) as the scattering function for an individual platelet, we can generalize the derivations by Heitz et al. [95, Eq. (5)] and compute the phase function of the platelets' substrate as

$$f_p^p = \int_0^\infty \int_{\Omega} \mathcal{F}_p(\omega_i, \omega_o | \omega_m, \tau_{p,s}) D(\tau_{p,s}) D_{\omega_i}(\omega_m) d\omega_m d\tau_{p,s} \quad (3.14)$$

$$\begin{aligned} &= \frac{D(\omega_m)}{4 \sigma_p(\omega_i)} \underbrace{\int_0^\infty \mathcal{R}(\omega_i | \tau_{p,s}) D(\tau_{p,s}) d\tau_{p,s}}_{\hat{\mathcal{R}}(\omega_i)} \\ &+ \delta(\omega_i - \omega_o) \underbrace{\int_0^\infty \int_{\Omega} \mathcal{T}(\omega_i | \omega_m, \tau_{p,s}) D(\tau_{p,s}) D_{\omega_i}(\omega_m) d\omega_m d\tau_{p,s}}_{\hat{\mathcal{T}}(\omega_i)} \end{aligned} \quad (3.15)$$

where $D_{\omega_i}(\omega_m) = \frac{D(\omega_m) \langle \omega_m, \omega_i \rangle}{\sigma_p(\omega_i)}$ is the distribution of visible normals; we parameterize \mathcal{F}_p given by (3.1) by the platelet's normal and the substrate thickness. The first term in Equation (3.15) is equivalent to a phase function based on specular microflakes [95; 100], though the $\hat{\mathcal{R}}(\omega_i)$ directionally varying albedo constitutes a key difference to prior work.

The second term $\hat{\mathcal{T}}(\omega_i)$ models an ideal forward-scattering peak based on a Dirac delta function. In this way, it behaves very similarly to null-scattering [155] and effectively reduces the optical density of the medium. However, in contrast to null-scattering, where this delta term is intentionally added to facilitate unbiased sampling of heterogeneous media, it is not fundamentally needed here. Furthermore, the additional sampling decision to choose between reflection and transmission components would add additional variance in a Monte Carlo framework. For this reason, we directly merge the effects of platelet transmission into Equations (3.12) and (3.13), which yields an equivalent description that is easier to sample:

$$\Sigma_t^p(\omega_i) = \rho_p \sigma_p(\omega_i) \left(1 - \widehat{\mathcal{T}}(\omega_i)\right), \quad (3.16)$$

$$\Sigma_s^p(\omega_i) = \alpha_p(\omega_i) \rho_p \sigma_p(\omega_i) \left(1 - \widehat{\mathcal{T}}(\omega_i)\right), \quad (3.17)$$

$$f_p^p(\omega_i \rightarrow \omega_o) = \frac{D(\omega_m)}{4\sigma_p(\omega_i)} \widehat{\mathcal{R}}(\omega_i). \quad (3.18)$$

3.6 Implementation

We implement our model as a BSDF in Mitsuba 2 [160], leveraging its native support for spectral rendering. In this section we provide details about the implementation of the different building blocks.

3.6.1 Position-free Monte Carlo

To implement our pearlescent material model, we leverage the generality of the recent position-free approach by Guo et al. [78] for rendering stratified materials. This method builds on the assumption that scattering takes place within a small surface region so that lateral effects can be neglected. This enables a simplified parameterization of the path integral formulation in terms of depth and orientation, where the invariance with respect to lateral displacement enables variance reduction through connection strategies such as next event estimation and multiple importance sampling (MIS). Light transport inside the material is computed stochastically in an unbiased manner, explicitly accounting for all interactions involving particles that permeate the interior of layers, as well as smooth or rough interfaces between layers. Finally, it supports arbitrary complexity and does not require the costly precomputation of prior work based on adding-doubling [53; 54].

3.6.2 Spectral rendering

Given the strong dependence on wavelength of the iridescent phase functions, using traditional RGB rendering for volumetric scattering might result in strong spectral aliasing. To avoid potential errors on the appearance reproduction, we implement our model in a fully spectral renderer, which fits very well into our stochastic model. We base our implementation on hero wavelength [234] for efficiently tracking four wavelengths at the same time. The sampling techniques for each wavelength are then combined via multiple importance sampling (MIS). This is crucial for obtaining low spectral variance (see [234] for details). In order to calculate the wavelength-dependent index of refraction for platelets and container $\eta(\lambda)$, we use a Cauchy polynomial fit from measured η and Abbe numbers V_d (lower values indicate higher chromatic dispersion) for transparent materials such as mica, SiO_2 , or TiO_2 . For absorbing materials such as Fe_2O_3 or aluminum, we resort to tabulated captured data (see Table A1).

3.6.3 Directional distribution.

The directional distribution $D(\omega_m)$ of the platelets in pearlescent materials has been previously analyzed by light microscope images or CT scans (e.g., [123]). It follows a Gaussian angular distribution [121; 124], with standard deviation between 7° and 30° . While by default the mean direction would be equal to the normal of

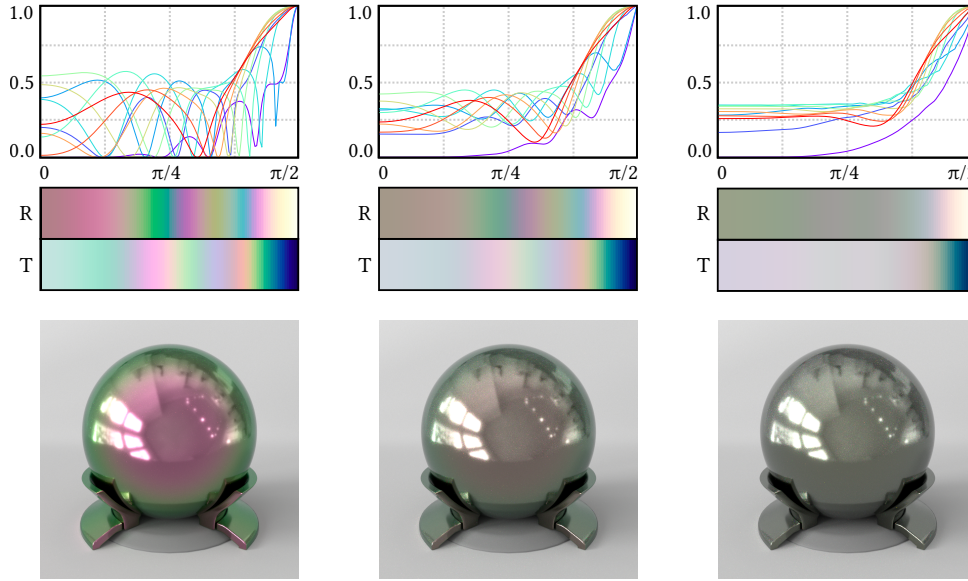


Figure 3.5: Response of a mica-substrate platelet with a thickness of $\tau_1 = 500$ nm, and a 60 nm TiO_2 coating, and varying substrate thickness distributions. The leftmost column shows how neglecting the effect of thickness variability [53] may lead to unrealistic appearance. The middle and right columns show results taking into account substrate thickness distributions $D(\tau_{p,s})$ with standard deviation $\gamma_{p,s}$ of 50 and 100 nm, respectively. In this example, interference effects are largely attenuated with increasing $\gamma_{p,s}$, due to loss of coherence of the light reflected within the mica substrate. Note that introducing a stochastic distribution of the thickness does not incur into significant additional variance during rendering (all renders have been done using 4K samples/pixel, 20 mins).

the pearlescent stratum ω_n , the orientation of the platelets can also be changed, as done in metallic paints to achieve complex appearance effects [193]. Gaussian-based microflake distributions exist [250]; in our work, we employ the roughly equivalent SGGX distribution [95], which models the distribution of normals as an ellipsoid, and allows simpler and more efficient evaluation and sampling of the distribution of visible normals $D_{\omega_i}(\omega_m)$. This model is parameterized by the effective cross section of the particles in each of the orthogonal directions (σ_p^x , σ_p^y , σ_p^z). We set σ_p^z to be the surface area of the platelet, orthogonal to the platelet mean normal $\hat{\omega}_m$ and allow σ_p^x and σ_p^y to be defined independently, thus supporting anisotropic deviation of the platelets. Unless stated otherwise (see Figure 3.13), the platelet mean normal $\hat{\omega}_m$ in results coincides with the stratum normal ω_n , and we assume an isotropic directional distribution of the platelets (i.e. $\sigma_p^x = \sigma_p^y = \sigma_p^{xy}$).

3.6.4 Substrate thickness distribution

Prior work assumes a constant substrate thickness [53; 54]. However, neglecting the effect of thickness variation leads to severe errors in the resulting appearance, as illustrated in Figure 3.5. Based on measurements of mica-based substrates using an Atomic Force Microscope [195], we approximate the distribution of substrate thickness $D(\tau_{p,s})$ as a Gaussian distribution with standard deviation $\gamma_{p,s}$. While synthetic substrates such as silica are characterized by narrow distributions [169], the thickness of mica substrates can range between 100 and 1000 nm, with a mean thickness of 560 nm and a standard deviation of 179 nm [195].

3.6.5 Platelet stratum extinction coefficient

As defined in Equation (3.16), the effective extinction coefficient $\Sigma_i^P(\omega_i)$ of an iridescent strata depends on the platelets transmission $\hat{T}(\omega_i)$, which in turn is defined not only by the structure of the layers, but also by their distribution of normals $D_{\omega_i}(\omega_m)$ and thicknesses $D(\tau_{p,s})$ (see Equation (3.15)). Given the complexity of the resulting expression, which involves the evaluation of the transmissivity of the layered structure over both distribution domains, we employ a stochastic evaluation procedure. In particular, when sampling the mean free path or calculating the attenuation, we sample a platelet direction from the distribution of visible normals $D_{\omega_i}(\omega_m)$ (see [95] for details), and a thickness from the Gaussian distribution of $D(\tau_{p,s})$. As mentioned earlier, the main source of variance is due to the strong wavelength dependence of extinction, whereas the variance introduced by the above procedure is insignificant in comparison.

3.7 Analysis and evaluation

We now turn to evaluation of our method beginning with a comparison to captured reflectance data of real-world iridescent materials. In the second part of this section, we study the behavior of the parameter space in greater detail.

3.7.1 Comparison with captured data

We compare the results of our model with data from Ferrero and colleagues [60], who measured a series of pearlescent material samples composed of TiO₂-coated platelets within a Silica substrate, providing discretized measurements of their reflectance field with an angular resolution of 10 degrees. We re-create three materials: *MCS1*, referring to Colorstream® T20-04 WNT Lapis Sunlight manufactured by Merck KGaA; *MCS2*, referring to Colorstream® T20-02 WNT Arctic Fire also manufactured by Merck KGaA; and *BASF1*, a burgundy and green material manufactured by BASF Coatings GmbH.

To populate the parameters of our model, we rely on all the information available from the manufacturers for each material. In particular, the container is made of Polysilazane ($\eta = 1.555$) with a thickness $t_c = 150 \mu\text{m}$ and roughness $\sigma_c = 0.01$, and the indices of refraction for TiO₂ and Silica are 2.6142 and 1.4585 respectively (see also Table A1). The remaining parameters are obtained by brute-force optimization within the limits defined by manufacturing parameters reported in the literature [145]: coating thickness $\tau_1 \in [60, 170]$ nm, substrate mean thickness $\tau_s \in [80, 500]$ nm, platelets density $C_p \in [3, 13]\%$, platelets normal deviation $\sigma_p^{xy} \in [0.01, 0.1]$, and substrate thickness standard deviation $\gamma_s \in [0, 70]$ nm. The fitted parameters for each measured material can be found in Table A2.

We also compare the results achieved with simpler models (Ergun et al.'s [53], which in turn is based of Ershov's model [54]), using their available parameters. As observed in Figure 3.6 and Table 3.1, our model provides better fits to the measured data, thus allowing us to represent real materials more accurately. Additionally, Figure 3.7 provides a visual comparison of the resulting distribution of chromaticities. Again, our model (green dots) produces reflectance values that are in closer agreement to ground-truth data than previous models, which tend to deviate significantly in both chromaticity and saturation.




Figure 3.6: Comparison of our model with ground-truth measurements [60] for three different materials: *MCS1*, *MCS2*, and *BASF1* (we keep the same naming as in the source article). The bottom row shows the results using a previous model [53]; our model allows to reproduce real measurements more faithfully (see also Figure 3.7).

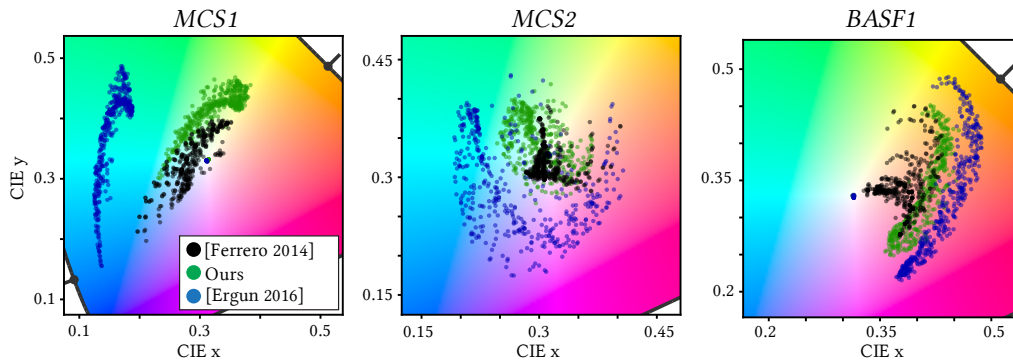


Figure 3.7: Chromaticity space of measured materials (black dots) [60], our model (green), and previous work (blue) [53]. Our model provides a closer match with measured data, with more accurate chromaticities and without excessive saturation.

Table 3.1: Error comparison of our model and Ergun et al.’s model [53] for the fits on Ferrero et al.’s [60] measurements. For each material we report error in CIE_{xy} chromaticity space, measured using mean squared error (MSE) and peak signal-to-noise ratio (PSNR).

Material	Model	MSE	PSNR (dB)
<i>MCS1</i>	Ergun et al. [53]	0.41400	11.53
	Ours	0.15100	14.74
<i>MCS2</i>	Ergun et al. [53]	0.27012	16.61
	Ours	0.08759	21.69
<i>BFS1</i>	Ergun et al. [53]	0.25252	17.05
	Ours	0.10579	20.74

3.7.2 Exploration of the parameter space

The set of optical phenomena that give rise to pearlescent appearance are complex and highly nonlinear: they involve the complex interplay of multiple anisotropic volume scattering, directionally varying interference within platelets, high-frequency spectral variation, interactions with the container, and refraction and internal reflection from layer boundaries. These aspects imply that predictive modeling of pearlescence remains a challenging and computationally intense task.

We performed a large set of simulations using our framework to study the space of pearlescent appearance, as parameterized by our model. We explore the resulting data and discuss our observations in the remainder of this section. We envision that such systematic exploration of material configurations could be a powerful ingredient in the computational design of pearlescence in the future. As explained in the previous sections, our framework supports an arbitrary number of strata. We restrict this analysis to a single-pearlescent-stratum material (such as manufactured plastic containers used by the cosmetic industry) only as a reasonable compromise to illustrate the capabilities of our model, while keeping the number of physical parameters to explore tractable.

We model a material consisting of a pearlescent stratum containing the iridescent platelets, on top of a base stratum (see Figure 3.9). This material (C1) contains platelets with an 80 nm SiO₂ substrate. The substrate is coated with a TiO₂ layer with a varying thickness τ_1 . We include a common PET thermoplastic container with thickness $t_c = 150 \mu\text{m}$, and interface roughness $\sigma_c = 0.01$. These pearlescent strata are applied over a base stratum, such as skin in cosmetic products, or the primer in automotive paints, which reflects light back to the pearlescent material. It thus plays an important role in its final appearance when transmittance is high, resulting in more vivid colors. We model this base stratum and set its diffuse reflectance albedo $\alpha_b = 0.7$.

For our analysis, we use a parameterization based on the half and difference angles (θ_h and θ_d , respectively) [187], which avoids redundant information and allows exploring isotropic BSDFs in a simpler 2D domain. Following Burley’s intuitive way to visualize materials by means of image slices [27], we observe two main goniochromatic effects along the same BRDF (see Figure 3.10): The first aligns with the difference angle θ_d along the material’s gloss component; the second appears on the diffuse reflection β , and is roughly radial with respect to $(\theta_h, \theta_d) = (0, 0)$. Based on these observations, we focus our analysis on the chromaticity and luminance, by exploring the effect of our model parameters on those two axes.

Chromaticity. Figure 3.8a illustrates the changes in diffuse chromaticity β , on a CIE xyY diagram as the thickness τ_1 of the platelets’ TiO₂ coating increases. We fix the platelet density $C_p = 3\%$ of the total volume, and the deviation to the normal distribution to $\sigma_p^{xy} = 0.1$. It can be seen how the diffuse chromaticity falls in a very narrow manifold, which is mostly outside the sRGB gamut (shown as a superimposed white triangle). Pairs of points on the different curves indicate equal diffuse coordinates, showing how the diffuse component travels along the manifold as τ_1 increases. Figure 3.8b shows a similar behavior for gloss chromaticity, except that the diffuse and gloss chromaticities change in opposite directions, presenting roughly complementary hues (see pairs of points in Figure 3.8c).

In addition, as the deviation of the normal distribution σ_p^{xy} increases, multiple scattering becomes dominant. As a consequence, colors become less saturated, both for the diffuse and gloss components (Figure 3.8d). On the other hand, the platelet density C_p only affects the diffuse component significantly (Figure 3.8e), especially for small σ_p^{xy} , where the saturation increases with density.

Luminance. Figure 3.8f shows how diffuse luminance decreases with the platelet density C_p , while the normal distribution deviation σ_p^{xy} has very little effect. Diffuse luminance is also higher on platelets with a thickness between 100 and 130 nm, since

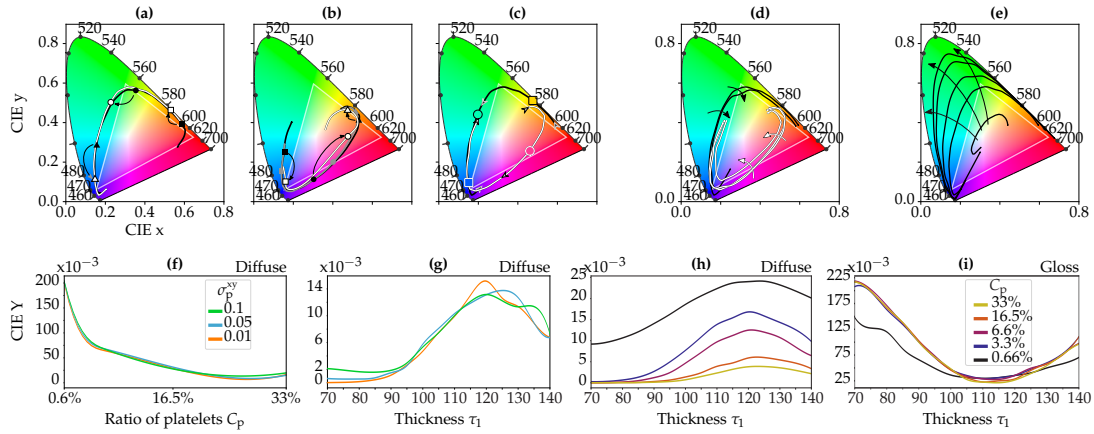


Figure 3.8: Effect of the different parameters of material C1 on its appearance (in all the diagrams, the triangle indicates the RGB gamut). **Top row (Chromaticity analysis):** (a) Influence of the TiO_2 coating thickness on the diffuse chromaticity, given a fixed $\sigma_p^{xy} = 0.01$ and $C_p = 3\%$. The two curves shown (black and white), correspond to two different thicknesses; pairs of points on the curves illustrate the changes in diffuse chromaticity as thickness increases. (b) Similar diagram for gloss chromaticity, for three curves of different thicknesses. These first two diagrams show how both chromaticities change, but remain within a narrow manifold. (c) Diffuse (black) and gloss (white) chromaticities as the thickness of the coating layer increases, for directions close to normal incidence. It can be seen how both components rotate in opposite directions, always showing roughly complementary colors. (d) Effect of the deviation of the normal distribution σ_p^{xy} for a fixed platelet density $C_p = 3\%$, showing both the diffuse (black) and gloss (white) components. As the deviation increases (as indicated by the arrows), both components become less saturated. (e) Effect of platelet density C_p for a fixed platelet normal deviation $\sigma_p^{xy} = 0.01\%$; as density increases (see arrows), so does saturation. **Bottom row (Luminance analysis):** Diffuse luminance, as a function of density (f), and coating thickness (g) for different platelet normal distribution deviations σ_p^{xy} . Diffuse (h) and gloss (i) luminance as a function of coating thickness τ_1 for different densities.




Figure 3.9: Schematic view of the material structure used for modeling materials C1 and C2. Both materials consist of an iridescent plastic stratum on top of a base diffuse layer. Note that we use single-stratum material for our analysis to keep the parameter space tractable; our model supports an arbitrary number of layers. We analyze the optical behavior of this pearlescent material, in particular the influence of our model parameters on the goniochromatic diffuse and gloss components, which are in turn affected by the platelet characteristics and the strata below the container.




Figure 3.10: 2D image slices representations [27] of the reflectance of material C1, for three different coating thickness τ_1 . Superimposed on the left, we can see the two main axes of goniochromaticity: One axis runs along the difference angle on the specular reflection (black); the second axis β (white) describes variations in diffuse reflection, and is radial from the normal reflection. The dots on top of each of both axes represent the angular location of the samples used in our analysis.




Figure 3.11: Variations of pearlescent materials due to changes on the thickness of the coating layer and the density of platelets. **C1 (top):** Silica substrate ($\tau_s = 80$ with zero deviation), TiO_2 coating. **C2 (bottom):** Mica substrate ($\tau_s = 560.44$, $\gamma_s = 179.32$). All samples have the same container roughness $\sigma_c = 0.01$, platelet normal deviation $\sigma_p^{xy} = 0.1$, and base stratum albedo $\alpha_b = 0.7$. Changes in the thickness of the coating layer produce large variations in chromaticity. On the other hand, as the density of the platelets increases the diffuse luminance decreases, and the gloss component of the chromaticity progressively dominates the final appearance. While the global behavior is similar with both types of substrate, the resulting appearance is vastly different (see the closeups of samples A and B in Figure 3.12, right). The material sample highlighted with the letter C (top) illustrates how the gloss hue first manifests around the highlights (please refer to the main text).

coatings in that range yield a higher transmittance (Figures 3.8g and 3.8h). The luminance of the gloss component behaves in a complementary manner (Figure 3.8i).

3.7.3 Additional results

Figure 3.11 (top) shows the significant variation in pearlescent appearance which emerges from the interactions of different parameters in our model. In particular, we show appearance changes in our material C1, due to changes in platelet density ρ_p and thickness τ_1 of the TiO_2 coating layer, for a base albedo $\alpha_b = 0.7$. Changes in thickness lead to strong changes in chromaticity, as shown in Figure 3.8a. On the other hand, as the platelet density C_p increases, the diffuse luminance decreases (as shown in Figure 3.8h). As a result, the diffuse hue becomes progressively less dominant, and the gloss hue emerges. This gloss hue first manifests subtly around the main highlights produced by the dielectric interface of the container: see for instance the reddish halo around the highlights in the object marked with a C in Figure 3.11, top; as density increases, this reddish hue progressively dominates the final appearance.

For comparison purposes, we model a second material C2, identical to material C1,





Figure 3.12: **Left:** Comparison between the chromaticity of the diffuse (black) and gloss (white) components for the C1 (solid) and C2 (dashed) materials. The different platelet structure results into a different manifold of appearance. **Right:** Side-by-side comparisons of the samples A and B in Figure 3.11. Despite the large deviation in appearance, only the substrate material varies in each pair.




Figure 3.13: Variations of material C1 with increasing platelet rotation (shown for the x -axis). Even slight rotations result in a visible angular offset in the highlights produced by the platelets (reddish for this material). **Top:** isotropic normal distributions of the platelets. **Bottom:** anisotropic normal distributions of the platelets. The last image shows anisotropy in the y -axis.



Figure 3.14: Cars rendered with complex pearlescent paints, exhibiting platelet rotation and anisotropy. Table A4 describes the parameters of each material.

except for the substrate; instead of silica, C2 has a mica substrate, with a mean thickness of 560 nm, and standard deviation 179 nm. It can be seen how changing the substrate has a profound impact on the underlying diffraction phenomena, resulting in a marked shift in overall chrominance (see Figure 3.11, bottom). This can be better observed in Figure 3.12, showing closeups of the objects marked with A and B in the previous figure. However, the global behavior of the diffuse and gloss components is essentially equivalent to our previous analysis (Figure 3.12, left).

In contrast to prior work, our method also admits configurations that lead to an anisotropic BSDF. Figure 3.13 illustrates changes in appearance due to rotations in the mean direction of the platelets' normals, combined with both isotropic (top) and anisotropic distributions (bottom). We observe that slight rotations in the x axis result in an angular offset between the chromatic gloss produced by platelets and the highlights produced by the roughness of the container.

Table A3 describes the parameters of each material used in Figure 3.1. Altering a




Figure 3.15: Equal-time renderings demonstrating the importance of both multiple importance sampling (MIS) and spectral rendering using hero wavelength (HW), for the scene in Figure 3.1 rendered at 4K samples per pixel.

single value per pair of objects, as shown in Figure 3.1 and highlighted in bold in the table, leads to noticeable appearance changes. Different from previous models, our support for thickness distributions allows us to represent appearance changes (see difference between 7th and 8th bottles in Figure 3.1) due to variations in substrate thickness, which naturally occur in common substrate compounds such as mica. Finally, Figure 3.14 demonstrates the generality of our model, showing three different cars rendered with complex pearlescent car paints. We use multi-layer platelets with metallic substrates (front and middle) and alumina substrate (back) [145], displaying different degrees of rotation and anisotropy. The exact parameters of each material are described in Table A4.

All renders have been computed on a dual Intel Xeon Gold 6140 using 64 threads. Execution times for all the results in this chapter can be found in Table 3.2. Equal-time comparisons demonstrating the importance of MIS and hero wavelength for efficiently rendering our model can be found in Figure 3.15.

Table 3.2: Total render time, samples per pixel, and resolution for the different scenes used throughout the chapter.

RENDER TIMES				
Scene	Time	Samples/pixel	Resolution	Figures
<i>Cosmetics</i>	64 minutes	4k	1600×720	3.1
<i>Knob</i>	17 minutes	4k	512×512	3.11, 3.12, 3.13
<i>Cars</i>	31 minutes	4k	1280×550	3.14

3.8 Conclusions

We have presented a general model that simulates the complex internal processes responsible for pearlescent appearance, including scattering from multi-layered platelets

subject to wave-optical interference, and internal reflection from smooth or rough layer boundaries. Our model is based on a thorough review of the structure and properties of real-world pearlescent pigments, and it accounts for stochastic variation in the local material properties that critically impacts accuracy. We demonstrate the practicality of our method and showcase a series of comparisons that showcase its superior performance compared to prior work. Our results also include an analysis of the influence of physical parameters on the resulting goniochromatic behavior. As this analysis shows, small perturbations of a single parameter often lead to disproportionately large changes in material appearance. We believe that detailed computational mapping of this highly nonlinear space will be a crucial component of future manufacturing application that seek to create pearlescent materials with desired optical properties. To foster future work on this topic, and to ensure the reproducibility of this article, we will release an open implementation of our full simulation pipeline.

Limitations & future work. Our model is implemented on top of a spectral renderer. To reduce variance arising from the stochastic nature of transmittance $\hat{T}(\omega_i)$ evaluation, the design of efficient sampling methods is an important avenue for future work. While our implementation builds on specialized sampling and MIS techniques for multilayered materials and spectral rendering, it would directly benefit from recent and future advances in both lines of research (e.g. [64; 240] and [130]). For applications where rendering efficiency is preferred over exact appearance, faster although less accurate solutions can be employed [53; 54], solving the volumetric transport inside the container strata using either precalculated [101] or approximated solutions [14; 230]. As for the parameter space, our model takes into account the main effects that can be controlled during manufacturing, such as the particular materials used, thickness distributions, or the deviation from the platelet normal distribution. Other manufacturing issues such as irregularities in the substrates, flaws in the layer precipitation process, small thermal and mechanical cracks of the pigments, or discontinuities and pores in the pigments may lead to additional scattering. This potentially replaces iridescence with an undesired hazy appearance, which could be characterized using Mie theory [146]. To make our model tractable we assume that lateral boundaries of platelets can be ignored. Removal of this approximation could be desirable to further improve accuracy, but this would entail replacing the layering computation by a significantly more costly wave-level simulation that would likely be impractical in the context of rendering. We also assume that all platelets suspended into the material are very small (on the order of a few microns), which is common in many pearlescent materials. Adding larger platelets would allow us to model glints (such as those in many car paints), and would require to replace the continuous distribution of platelets with a discrete counterpart [102]. Last, while intuitive models for editing BRDFs' appearance exist [120; 192], designing tools to enable intuitive editing of pearlescent materials remains a challenging open problem. This will involve translate the physical parameter space into perceptually-based appearance spaces [133; 168; 237]. Our model provides a useful foundation, but further work is required to address the high-dimensional and nonlinear nature of the underlying parameter space.

3.A Tables

In the following we describe the characteristics of container, substrate, and coating compounds used throughout the chapter.

Table A1: Characterization of the materials used throughout the chapter [178]. For transparent materials, we report the measured index of refraction η and Abbe number V_d . For absorbing materials (Fe_2O_3 , aluminum and copper) we report the average η from measured data. In the case of PET an polysilazane there are no accurate measurements and we resort to a constant η .

Material	η	V_d	Source
Air	1.0		
PET	1.5750		[202]
Polyurethane	1.5650	24.37	[98]
Polysilazane	1.5550		[157]
Mica	1.6137	54.56	[11]
SiO_2	1.4585	67.82	[147]
TiO_2	2.6142	9.87	[43]
Fe_2O_3	$3.3206 + 0.2192i$		[179]
Al_2O_3	1.7742	72.31	[148]
Al	$1.1978 + 7.0488i$		[180]
Cu	$0.7400 + 2.7071i$		[179]

Table A2: Optimized parameters for our model and Ergun's [53] (see Figure 3.6 and Figure 3.7 in the main text).

Material	Model	ρ_p	τ_s (nm)	γ_s	τ_1 (nm)	σ_p^{xy}
<i>MCS1</i>	Ergun et al. [53]	6.6%	160	0	120	0.1
	Ours	3.3%	280	40	60	0.1
<i>MCS2</i>	Ergun et al. [53]	13.2%	400	0	60	0.1
	Ours	13.2%	410	40	60	0.1
<i>BFS1</i>	Ergun et al. [53]	13.2%	160	0	60	0.1
	Ours	13.2%	160	20	60	0.1

Table A3: Description of the parameters used to render the objects in Figure 3.1 (bottles numbered in reading order), where C_p is the platelet density (fraction of platelets per total volume), τ_1 is the TiO_2 thickness, σ_p^{xy} is the deviation of the mean normal, σ_c is the dielectric roughness, γ_s is the deviation of the substrate thickness. We highlight in bold the varying parameter for each pair. All materials have base stratum albedo $\alpha_b = 0.6$, and a PET container.

BOTTLES: MATERIALS DESCRIPTION						
Bottle	C_p	τ_1 (nm)	σ_p^{xy}	σ_c	γ_s	Substrate
1	13.2%	140	0.05	0.2	179.32	Mica
2	1.3%	140	0.05	0.2	179.32	Mica
3	9.9%	85	0.02	0.1	0.0	SiO₂
4	9.9%	85	0.02	0.1	179.32	Mica
5	6.6%	90	0.01	0.1	179.32	Mica
6	6.6%	90	0.2	0.1	179.32	Mica
7	6.6%	70	0.07	0.07	22.41	Mica
8	6.6%	70	0.07	0.07	179.32	Mica
9	2.0%	90	0.05	0.02	179.32	Mica

Table A4: Description of the parameters used to render the cars in Figure 3.14. Front and middle cars have double-coating metallic platelets, and back car single-coating Alumina platelets [145].

CARS: MATERIALS DESCRIPTION			
Car	Front	Middle	Back
Coating 1	TiO ₂ (147 nm)	TiO ₂ (70 nm)	TiO ₂ (80 nm)
Coating 2	Fe ₂ O ₃ (80 nm)	Fe ₂ O ₃ (115 nm)	—
Substrate	Cu (80 nm)	Al (80 nm)	Al ₂ O ₃ (40 nm)
σ_p^{xy}	(0.05, 0.1)	(0.1, 0.01)	(0.02, 0.1)
Rotation (x,y,z)	(10°, 0°, 0°)	(10°, 0°, 0°)	(5°, -3°, 0°)
C_p	6.6%	6.6%	6.6%
σ_c	0.1	0.01	0.05
α_b	0.5	0.7	0.3

Part III

Transient state light transport

Chapter 4

Progressive transient photon beams

In this chapter we present a method for efficient transient rendering in the presence of participating media. Light transport in this kind of scenes is challenging to calculate due to the large amount of possible light paths, which results high variance even at high sample counts. Our method is based on density estimation, and let us reuse calculations using information from other light paths. We extend the concept of progressive photon beams [110] for density estimation in participating media to the transient state, allowing us to reuse light path close both in the temporal and spatial domains. Additionally, by introducing smoothing kernels we mollify delta or almost delta interactions across the light paths, allowing us to sample paths that are provably unsampleable using conventional Monte carlo techniques. Our extension results in consistent, robust method which provably converges to the correct solution using finite memory. We derive optimal convergence rates accounting for spatial and temporal kernels, and provide empirical results of our method improvement over existing approaches.

My role as the second author of this project involved improving the implementation of the proposed progressive algorithm, specially the 1D spatial variant. My implementation was used to generate many of the results found in Section 4.6, as well as the empirical convergence analysis.

This work was originally published in Computer Graphics Forum and presented at the 30th Eurographics Symposium on Rendering (EGSR 2019).

Progressive transient photon beams

Julio Marco, Ibón Guillén, Wojciech Jarosz, Diego Gutierrez & Adrian Jarabo
Computer Graphics Forum, Vol. 38(6), 2019

4.1 Introduction

The emergence of transient imaging has led to a vast number of applications in graphics and vision [107], where the ability of sensing the world at extreme high temporal resolution allows new applications such as imaging light in motion [225], appearance capture [158], geometry reconstruction [29; 150], or vision through media [28; 239] and around the corner [7; 224]. Sensing through media is one of the key applications: The ability of demultiplexing light interactions in the temporal domain is a very promising approach for important practical domains such as non-invasive medical imaging, underwater vision, or autonomous driving through fog.




Figure 4.1: The SOCCER scene (steady-state render on the right) features complex volumetric caustics due to multiple reflections and refractions off smooth dielectrics inside the medium. We are able to efficiently render the transient light transport (left sequence) by formulating a progressive, transient form of photon beam density estimation which provably eliminates error while working within a finite memory budget.

Accurately simulating light transport could help enormously in these applications, potentially serving as a benchmark, a forward model in optimization, or as a training set for machine learning.

Transient rendering in media is, however, still challenging: The increased dimensionality (time) increases variance dramatically in Monte Carlo algorithms, potentially leading to impractical rendering times. This variance is especially harmful in media, where the signal tends to be smooth due to the low-pass filtering behavior of scattering, in both the spatial and temporal domains. One of the major drawbacks of transient rendering is that it requires much higher sampling rates to fill up the extended temporal domain, specially when using $0D$ (photon) point samples, which are sparsely distributed across both time and space. We make the observation that $1D$ photon trajectories populate both space and time much more densely; hence, a technique based on photon beams [109] should significantly reduce the rendering time when computing a noise-free time-resolved render, and, given its density estimation nature, it could naturally combine with the temporal domain density estimation proposed by Jarabo et al. [106].

We present a new method for transient-state rendering of participating media, that leverages the good properties of density estimation for reconstructing smooth signals. Our work improves Jarabo et al. [106] by extending *progressive photon beams* (PPB) [110] to the transient domain, and combining it with temporal density estimation for improved reconstruction in both the spatial and temporal domains. Our technique is biased but consistent, converging to the ground truth using finite memory by taking advantage on the progressive [83; 127] nature of density estimation. We analyze the asymptotic convergence of our proposed space-time density estimation, computing the optimal kernel reduction ratios for both domains. Finally, we demonstrate our method on a variety of scenes with complex volumetric light transport, featuring high-frequency occlusions, caustics, or glossy reflections, and show its improved performance over naively extending PPB to the transient domain.

This technique is an extension of our previous work on rendering transient volumetric light transport [151], where we proposed a naive extension of photon beams to transient state. Here we increase the applicability of the method, by proposing a progressive version of the space-time density estimation, and rigorously analyze its convergence.

4.2 Related work

Rendering participating media is a long-standing problem in computer graphics, with a vast literature on the topic. Here we focus on works related directly with the scope of the project. For a wider overview on the field, we refer to the recent survey by Novák et al. [161].

Photon-based light transport. Photon mapping [112] is one of the most versatile and robust methods for rendering complex global illumination, with several extensions for making it compatible with motion blur [31], adapting the distribution of photons [72; 203], carefully selecting the radiance estimation kernel [103; 118; 203], combining it with unbiased techniques [66; 84], or making it progressive for ensuring consistency within a limited memory budget [83; 127]. Hachisuka et al.’s [82] recent SIGGRAPH course provides an in-depth overview.

Jensen and Christensen [113] were the first to extend photon mapping to media, and Jarosz and colleagues [111] significantly improved its efficiency with the beam radiance estimate, which replaces repeated point queries with one “beam” query finding all photons along the entire camera ray. Jarosz et al. [109] later applied this idea to the photon tracing process by storing full photon trajectories (photon beams), leading to a dramatic increase in photon density for the same photon tracing step. Their progressive and hybrid counterparts [110; 131] leveraged the benefits of photon beams while providing consistent solutions using finite memory. Recently, Bitterli and Jarosz [19] generalized 0D photon points and 1D photon beams to even higher dimensions, proposing the use of photon planes (2D), volumes (3D) and, in theory, higher-dimensional geometries, leading to unbiased density estimation. All these works are, however, restricted to steady-state renders; we instead focus on simulating light transport in transient state.

Transient rendering. Though the transport equations [35; 68] are time-resolved, most rendering algorithms focus on steady-state light transport. Still, several works have been proposed to deal with light transport in a time-resolved manner. In particular, most previous work on transient rendering has focused on simulating surfaces transport: Klein et al. [126] extended Smiths’ transient radiosity [200] for second bounce diffuse illumination, while other work has used more general methods based on transient extensions of Monte Carlo (bidirectional) path tracing [105; 106; 175] and photon mapping [154; 163]. Several works have also dealt with time-resolved transport on the field of neutron transport [16; 34; 49; 236]. Closer to our work, Ament and colleagues [4] rendered transient light transport in refractive media using volumetric photon mapping, but they do not provide an efficient approach that guarantees consistency. Jarabo et al. [106] proposed a transient extension of the path integral, and introduced an efficient technique for reconstructing the temporal signal based on density estimation. They also proposed a set of techniques for sampling media interactions uniformly in time. Their method is however limited to bidirectional path tracing and photon mapping, often failing to densely populate media in the temporal domain. Finally, Bitterli [18] and Marco et al. [151] proposed a transient extension of the photon beams algorithm, but these approaches are not progressive, therefore not converging to the correct solution in the limit. Our work extends the latter, proposing a progressive, consistent, and robust method for rendering transient light transport. We leverage beams continuity and spatio-temporal density

estimation to mitigate variance in the temporal domain, and derive the parameters for optimal convergence of the method.

4.3 Transient radiative transfer

The *radiative transfer equation* (RTE) [35] models the behavior of light traveling through a medium. While the original formulation is time-resolved, its integral form used in traditional rendering ignores this temporal dependence, and computes the radiance L reaching any point \mathbf{x} from direction ω as

$$L(\mathbf{x}, \omega) = T_r(\mathbf{x}, \mathbf{x}_s) L_s(\mathbf{x}_s, \omega) + \int_0^s \mu_s(\mathbf{x}_q) T_r(\mathbf{x}, \mathbf{x}_q) L_o(\mathbf{x}_q, \omega) dq, \quad (4.1)$$

where $\mathbf{x}_d = \mathbf{x} - d \cdot \omega$ is a point at distance d , μ_s is the scattering coefficient, and $T_r(\mathbf{x}, \mathbf{x}_d) = \exp(-\int_0^d \mu_t(\mathbf{x}_{d'}) dd')$ is the *transmittance* describing the fraction of photons that make it between \mathbf{x} and \mathbf{x}_d without undergoing extinction at any point $\mathbf{x}_{d'}$, determined by the *extinction coefficient* $\mu_t(\mathbf{x}_{d'})$. The outgoing radiance L_o in direction ω from a medium point \mathbf{x}_q at distance q is defined by the scattering integral:

$$L_o(\mathbf{x}_q, \omega) = L_e(\mathbf{x}_q, \omega) + \int_{\mathcal{S}} f_s(\mathbf{x}_q, \omega_i, \omega) L(\mathbf{x}_q, \omega_i) d\omega_i, \quad (4.2)$$

where \mathcal{S} is the spherical domain, and f_s is the phase function. L_s is defined analogously via the *rendering equation* [117], but integrated over the hemispherical domain, and using the cosine-weighted BSDF in place of the phase function.

Transient RTE Equations 4.1 and 4.2 assume that the speed of light is infinite. However, if we want to solve the RTE at time scales comparable to the speed of light we need to incorporate the different delays affecting light. In the following we review the main practical considerations for accounting time into the integral form of the RTE for its application in transient rendering. Light takes a certain amount of time to propagate through space, and therefore light transport from a point \mathbf{x}_0 towards a point \mathbf{x}_1 does not occur immediately. In the absence of scattering effects, transport between two points \mathbf{x}_0 and \mathbf{x}_1 occurs as

$$L(\mathbf{x}_1, \omega, t) = L(\mathbf{x}_0, -\omega, t - \Delta t), \quad (4.3)$$

where Δt is the time it takes the light to go from \mathbf{x}_0 to \mathbf{x}_1 . In turn, Δt is defined by

$$\Delta t(\mathbf{x}_0 \leftrightarrow \mathbf{x}_1) = \int_{\mathbf{x}_0}^{\mathbf{x}_1} \frac{\eta(\mathbf{x})}{c} d\mathbf{x}, \quad (4.4)$$

where $\eta(\mathbf{x})$ is the index of refraction at a medium point \mathbf{x} and c is the speed of light in vacuum. Note that in this case light does not travel in a straight line, but by following the Eikonal equation [4; 81]. In a medium with a constant index of refraction $\eta(\mathbf{x}) = \eta_m$, then $\Delta t(\mathbf{x}_0 \leftrightarrow \mathbf{x}_1)$ can be expressed as

$$\Delta t(\mathbf{x}_0 \leftrightarrow \mathbf{x}_1) = \frac{\eta_m}{c} \|\mathbf{x}_1 - \mathbf{x}_0\|. \quad (4.5)$$

The second form of delay occurs in the scattering events, and might occur from different sources, including electromagnetic phase shift, fluorescence and phosphorescence, or multiple scattering within the surface (or particle) microgeometry. To




Figure 4.2: Photon emission and radiance estimation. (a) A photon emitted from the light source will take a time $t_{b_0} = \frac{\eta m}{c}(s_1 + s_2 + s_3)$ to get to \mathbf{x}_b . (b) Radiance estimation in the medium is done by intersecting every ray against the photon beam map, and performing density estimations at the ray-beam intersections (red).

account for these sources of scattering delays, we introduce a temporal variable in the phase function as $f_s(\mathbf{x}, \omega_i, \omega, t)$, where t is the instant of light interacting with the particle before it is scattered. With those delays in place, we reformulate the RTE (Equations 4.1 and 4.2) introducing the temporal dependence as [68]

$$L(\mathbf{x}, \omega, t) = T_r(\mathbf{x}, \mathbf{x}_p) L_s(\mathbf{x}_p, \omega, t - \Delta t_p) + \int_0^p \mu_s(\mathbf{x}_q) T_r(\mathbf{x}, \mathbf{x}_q) L_o(\mathbf{x}_q, \omega, t - \Delta t_q) dq, \quad (4.6)$$

$$L_o(\mathbf{x}_q, \omega, t) = \int_{-\infty}^t L_e(\mathbf{x}_q, \omega, t') dt' + \int_{\mathcal{S}} \int_{-\infty}^t f_s(\mathbf{x}_q, \omega_i, \omega, t - t') L(\mathbf{x}_q, \omega_i, t) dt' d\omega_i, \quad (4.7)$$

with $\Delta t_p = \Delta t(\mathbf{x} \leftrightarrow \mathbf{x}_p)$ and $\Delta t_q = \Delta t(\mathbf{x} \leftrightarrow \mathbf{x}_q)$ (Equation (4.4)). L_s changes analogously. Note that we assume that the matter does not change at time-scales comparable to the speed of light, and therefore avoid any temporal dependence on μ_s and μ_t . Introducing temporal variation at such speeds would produce visible relativistic effects [108; 231].

4.4 Transient photon beams

Photon beams [109] provide a two-pass numerical solution for rendering participating media in steady state: In the first pass (Figure 4.2a), a series of random walk paths are traced from the light sources. These paths represent packages of light (photons) traveling through the medium. Every interaction of a photon within the medium is stored on a map as a *beam* with a direction ω_b , position \mathbf{x}_b and power Φ_b . In the second pass (Figure 4.2b), rays are traced from the camera against the scene, and Equation (4.1) is approximated by summing up the contribution of all near photon beams R_b of the eye ray defined by $r = (\mathbf{x}_r, -\omega_r)$

$$L(\mathbf{x}_r, \omega_r) \approx \sum_{b \in R_b} L_b(\mathbf{x}_r, \omega_r), \quad (4.8)$$

where $L_b(\mathbf{x}_r, \omega_r)$ is the contribution of photon beam b . Every photon beam b is considered to have certain radius R_b , and radiance seen by a camera ray is computed

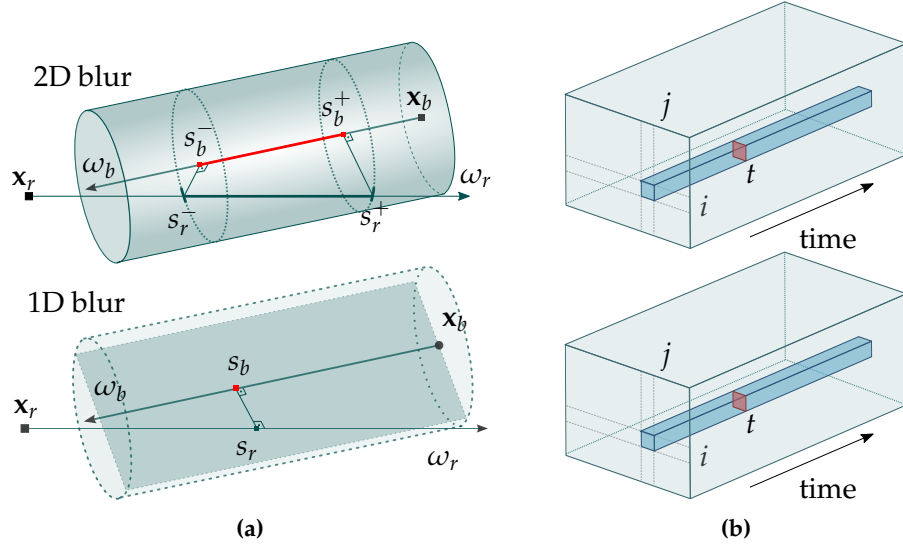


Figure 4.3: Spatio-temporal kernel estimators. (a) Ray-beam intersection for density estimation using a 2D kernel (top) and 1D kernel (bottom). Time delays t_b, t_r within these spatial density estimations will depend on the ray-beam orientation the blur region intersections s_b, s_r , the speed of light, and the index of refraction of the media. (b) Radiance estimate of a single beam at pixel ij using a 2D blur generates a temporal footprint over a time interval $[t^-, t^+]$ (top) while radiance estimate using a 1D blur occurs at a single time instant t (bottom).

by performing a density estimation on every ray-beam intersection. For 1D and 2D kernels, this radiance is computed as

$$L_b^{1D}(\mathbf{x}_r, \omega_c) = K_{1D}(R_b)\Phi_b f_s(\theta_b)\mu_s \frac{e^{-\mu_t s_b} e^{-\mu_t s_r}}{\sin \theta_b}, \quad (4.9)$$

$$L_b^{2D}(\mathbf{x}_r, \omega_r) = K_{2D}(R_b)\Phi_b f_s(\theta_b)\mu_s \frac{e^{-\mu_t(s_c^- - s_c^+)}(|\cos \theta_b| - 1) - 1}{e^{\mu_t(s_r^- + s_b^-)}\mu_t(|\cos \theta_b| - 1)}, \quad (4.10)$$

where the beam is defined by $\mathbf{x}_b + s_b\omega_b$ and the ray is defined by $\mathbf{x}_r + s_r\omega_r$ (see setups in Figure 4.3a).

4.4.1 Our algorithm

To generalize photon beams to the transient domain, we need to account for the duration of light paths. This requires considering propagation and scattering delays along the camera and light subpaths, but also the effect of time in the density estimation connecting these two subpaths.

Creating the photon map. We compute the photon propagation as a standard random walk through the scene, which can be modeled using the subpath formulation defined by Jarabo et al. [106]. Let us define a light subpath $\bar{\mathbf{x}}_l = \mathbf{x}_0 \dots \mathbf{x}_{k-1}$, with k vertices, where \mathbf{x}_0 is the light source. This light path defines $k - 1$ photon beams, in which a beam b_j is defined by its origin at $\mathbf{x}_{b_j} = \mathbf{x}_j$ and direction $\omega_{b_j} = \frac{\mathbf{x}_{j+1} - \mathbf{x}_j}{\|\mathbf{x}_{j+1} - \mathbf{x}_j\|}$. Using Jarabo's definition of the path integral (and therefore of the contribution of

the subpaths), we compute the flux of each photon as:

$$\Phi_{b_j} = \frac{f(\bar{\mathbf{x}}_j, \bar{\tau}_j)}{Mp(\bar{\mathbf{x}}_j, \bar{\tau}_j)} = \frac{L_e(\mathbf{x}_0 \rightarrow \mathbf{x}_1, \tau_0)T(\bar{\mathbf{x}}_j, \bar{\tau}_j)}{M \prod_{i=0}^j p(\mathbf{x}_i, \tau_i)}, \quad (4.11)$$

with $\bar{\mathbf{x}}_j$ the subpath of $\bar{\mathbf{x}}_l$ up the vertex j , f the subpath contribution function, $\bar{\tau}_j = \tau_0 \dots \tau_j$ the sequence of time delays up to vertex j , M the number of photon random walks sampled, $L_e(\mathbf{x}_0 \rightarrow \mathbf{x}_1, \tau_0)$ the emission function, $p(\mathbf{x}_i, \tau_i)$ the probability density of sampling vertex \mathbf{x}_i with time delay τ_i . The throughput, $T(\bar{\mathbf{x}}_j, \bar{\tau}_j)$, of subpath (\mathbf{x}_i, τ_j) is defined as:

$$T(\bar{\mathbf{x}}_j, \bar{\tau}_j) = \left[\prod_{i=1}^{j-1} f_s(\mathbf{x}_i, \tau_j) \right] \left[\prod_{i=0}^{j-1} G(\mathbf{x}_i, \mathbf{x}_{i+1}) V(\mathbf{x}_i, \mathbf{x}_{i+1}) \right], \quad (4.12)$$

with $f_s(\mathbf{x}_i, \tau_j)$ the scattering event at vertex \mathbf{x}_i with delay τ_j , and $G(\mathbf{x}_i, \mathbf{x}_{i+1})$ and $V(\mathbf{x}_i, \mathbf{x}_{i+1})$ the geometry and visibility terms between vertices \mathbf{x}_i and \mathbf{x}_{i+1} , respectively. Finally, for transient state we need to know the instant t_{b_j} at which the photon beam is created (through emission or scattering), defined as:

$$t_{b_j} = \sum_{i=0}^{j-1} \tau_j + \sum_{i=0}^{j-1} \Delta t(\mathbf{x}_i, \mathbf{x}_{i+1}). \quad (4.13)$$

Rendering. For rendering, we adapt Equation (4.8) to account for the temporal domain, as

$$L(\mathbf{x}_r, \omega_r, t) \approx \sum_{b \in R_b} L_b(\mathbf{x}_r, \omega_r, t), \quad (4.14)$$

with $L_b(\mathbf{x}_r, \omega_r, t)$ the radiance estimation for beam b to ray t at instant t . In essence, $L_b(\mathbf{x}_r, \omega_r, t)$ will return zero radiance if t is out of the temporal footprint of the density estimation kernel. Depending on the dimensionality of the density estimation, Jarosz and colleagues [110] proposed three different estimators based on 3D, 2D and 1D kernels. Since the 3D kernel results impractical due to costly 3D convolutions, we focus on 1D and 2D kernels (Equations (4.9) and (4.10)), and extend them to transient state, assuming homogeneous media.

Kernel 2D. We generalize Jarosz's et al.'s 2D estimate L_b^{2D} (Equation (4.10)) by introducing a temporal function $W(t)$ as

$$L_b^{2D}(\mathbf{x}_r, \omega_r, t) = K_{2D}(R_b) \Phi_b f_s(\theta_b, t) \mu_s \frac{e^{-\mu_t(s_r^- - s_r^+)} (|\cos \theta_b| - 1) - 1}{e^{\mu_t(s_r^- + s_b^-)} \mu_t (|\cos \theta_b| - 1)} W_{2D}(t), \quad (4.15)$$

where $[s_r^-, s_r^+]$ are the limits of the ray-beam intersection (Figure 4.3a), θ_b is the angle between ω_b and ω_r , and $K_{2D}(R_b)$ is a canonical 2D kernel with radius R_b . The temporal function $W_{2D}(t)$ models the temporal footprint of the 2D kernel as

$$W_{2D}(t) = \begin{cases} \frac{1}{t^+ - t^-} & \text{if } t \in (t^-, t^+) \\ 0 & \text{otherwise} \end{cases}, \quad (4.16)$$

where $t^- = t_b + t_r + \frac{\eta_m}{c}(s_r^- + s_b^-)$ and $t^+ = t_b + t_r + \frac{\eta_m}{c}(s_r^+ + s_b^+)$, and t_r and t_b are the initial times of the camera ray and beam, respectively. Note that due to transmittance, the photon energy varies as it travels across the blur region. Evenly distributing the integrated radiance L_b across this interval introduces temporal bias, in addition to the inherent spatial bias introduced by density estimation. However we observed this even distribution provides a good tradeoff between bias, variance, and computational overhead.

Kernel 1D. In the 1D kernel defined for density estimation by Jarosz et al. the spatial blur is performed over a line. Therefore, the energy of the beam is just spread on the ray on a single point at $r(s_r)$, from a single point of the beam $b(s_b)$ (see Figure 4.3a). In consequence, $s_r^\pm \rightarrow s_r$ and $s_b^\pm \rightarrow s_b$, which implies that $t^\pm \rightarrow t_{br}$, and the temporal function reduces to $W_{1D}(t - tb) = \delta(t)$, with $\delta(t)$ the Dirac delta function. With that in place, we transform Jarosz et al. 1D estimate to

$$L_b^{1D}(\mathbf{x}_r, \omega_r, t) = K_{1D}(R_b) \Phi_b f_s(\theta_b, t) \mu_s \frac{e^{-\mu_t s_b} e^{-\mu_t s_r}}{\sin \theta_b} \delta(t - t_b), \quad (4.17)$$

with $K_{1D}(R_b)$ a 1D kernel with radius R_b .

Implementation. Since photon beams correspond to full photon trajectories, they allow us to estimate radiance at any position $\mathbf{x}_b + s\omega_b$ of the beam, and therefore at any arbitrary time $t(\mathbf{x}_b + s\omega_b)$. As mentioned, one-dimensional radiance estimate corresponds to a single time across the beam. In a traditional rendering process where camera rays are traced through view-plane pixels against the beams map, the temporal definition *within* a pixel will be proportional to the amount of samples per pixel taken. Additionally, 2D blur requires distributing every radiance estimate along a time interval, which reduces variance in the time dimension of a pixel at the expense of introducing additional temporal bias.

Finally, note that the temporal footprint of the density estimation might be arbitrarily small, so the probability of finding a beam b at a specific time might be very low. We alleviate this issue using path reuse via density estimation [106]. In particular, for the non-progressive results we use histogram temporal density estimation. With this technique, the samples in the temporal domain are reused across all frames by evaluating their contribution functions, which correspond to the temporal window covered by each frame. In Section 4.5 we introduce temporal kernel-based density estimation, and combine it with the spatial density estimation of the beam.

4.5 Progressive transient photon beams

By means of Equations (4.30) and (4.32) we have introduced temporal dependence on the spatial density estimations that use 2D and 1D kernels, respectively. These density estimations reduce variance at the expense of introducing bias in the results, which means both Equations (4.8) and (4.29) will not converge to the correct solution, even with an infinite number of photons M . To avoid this, progressive density estimation aims to provide a biased, yet consistent technique, that in the limit converges to the expected value (in other words, the bias vanishes in the limit). The key idea is to average several render passes with a finite number of photon random walks M , progressively reducing the bias in each iteration while allowing variance to slightly increase.

Algorithm 1 Pseudo-code of our progressive spatio-temporal density estimation.

```

 $L_n \leftarrow 0$ 
 $R_b \leftarrow R_0$ 
 $\mathcal{T} \leftarrow \mathcal{T}_0$ 
for  $i \in [0..N)$  do
   $r \leftarrow \text{traceRay}()$ 
   $B \leftarrow \text{beamsMap}()$  (Eqs. (4.6), (4.7), (4.26)-(4.28))
   $R_b \leftarrow R_b \sqrt{\frac{i+2/3}{i+1}}$  (Eq. (4.20), left)
   $\mathcal{T} \leftarrow \mathcal{T} \sqrt{\frac{i+2/3}{i+1}}$  (Eq. (4.20), right)
   $L_b \leftarrow 0$ 
  for  $b \in B$  do
     $L_b \leftarrow L_b + \text{radiance}(r, b, R_b, \mathcal{T})$  (Eq. (4.18))
  end for
   $L_n \leftarrow L_n + L_b$ 
end for

```

In order to fully leverage a progressive approach, we propose to combine our time-resolved spatial density estimations (Section 4.4) with additional *temporal* density estimations. While our time-resolved 2D spatial kernel implicitly performs a temporal blur over the interval $[t^-, t^+]$, it is coupled with the spatial blur. This does not allow to choose its own initial kernel size for the temporal density estimation, which is a desirable degree of freedom since the temporal resolution may not be proportional to the spatial one. In contrast, our time-resolved 1D spatial kernel does not perform a temporal blur, since the footprint is a single instant in time. As we show in the remainder of this section, this allows us to perform additional progressive temporal density estimations with an independent initial kernel size, while keeping the same two-dimensionality (1D spatial and 1D temporal). In the following, we introduce our spatio-temporal beam density estimation based on our time-resolved 1D kernel, and then present our progressive approach.

Spatio-temporal beam estimation. Jarabo et al. [106] showed that progressive density estimations in the temporal domain can in fact improve the convergence rate for transient rendering, in particular when compared with the histogram method used in Section 4.4 for rendering the temporal domain. To combine such approach with the (progressive) spatial density estimation in photon beams [110], we reformulate the 1D kernel in Equation (4.32), by convolving it with a 1D temporal kernel $K_{\mathcal{T}}(t)$ so that

$$L_b^{\text{1D}}(\mathbf{x}_r, \omega_r, t) = K_{\text{1D}}(R_b) \Phi_b f_s(\theta_b, t) \mu_s \frac{e^{-\mu_t s_b} e^{-\mu_t s_r}}{\sin \theta_b} K_{\mathcal{T}}(t - t_b). \quad (4.18)$$

Progressive transient photon beams. We generalize the computation of $L(\mathbf{x}_r, \omega_r, t)$ (Equation (4.29)) using an iterative estimator, defined as

$$L(\mathbf{x}_r, \omega_r, t) \approx \widehat{L}_n(\mathbf{x}_r, \omega_r, t) = \frac{1}{n} \sum_{i=0}^n \sum_{b \in B_i} L_b(\mathbf{x}_r, \omega_r, t) \quad (4.19)$$

with \widehat{L}_n the estimate of L after n iterations, and B_i the set of photon beams in iteration i . Note that the previous equation assumes that the camera ray r is the same for all

iterations. That is not necessarily true (and in fact it is not) but for simplicity we express this way.

The error of the estimate \widehat{L}_n is defined by its bias and variance, which as shown in Appendix 4.B is dependent on the bandwidth of the spatial and temporal kernels. In particular, the variance of the error increases linearly with the bandwidth of the kernels, while bias is reduced at the same rate. Then, on each iteration we reduce the bias by allowing the variance to increase at a controlled rate of $(i + 1)/(i + \alpha)$, with $\alpha \in [0, 1]$ being a parameter that controls how much the variance is allowed to increase at each iteration. To achieve that reduction, on each iteration $i + 1$ we reduce the footprint of kernels K_{1D} and $K_{\mathcal{T}}$ ($R_{b|i}$ and \mathcal{T}_i) by

$$\frac{R_{b|i+1}}{R_{b|i}} = \left(\frac{i + \alpha}{i + 1}\right)^{\beta_R}, \quad \frac{\mathcal{T}_{i+1}}{\mathcal{T}_i} = \left(\frac{i + \alpha}{i + 1}\right)^{\beta_{\mathcal{T}}}, \quad (4.20)$$

where β_R and $\beta_{\mathcal{T}}$ control the individual reduction ratio of each kernel, with $\beta_{\mathcal{T}} = 1 - \beta_R$. A pseudo code of the main steps of our progressive approach can be found in Algorithm 1. In the following, we analyze the convergence rate of the method, and compute the optimal values for the parameters α , $\beta_{\mathcal{T}}$ and β_R .

Convergence analysis. We analyze the convergence of the algorithm as a function of the *asymptotic mean squared error* (AMSE) defined as

$$\text{AMSE}(\widehat{L}_n) = \text{Var}[\widehat{L}_n] + \text{E}[\epsilon_n]^2, \quad (4.21)$$

where $\text{Var}[\widehat{L}_n]$ is the variance of the estimate and $\text{E}[\epsilon_n]$ is the bias at iteration n . As shown in Appendix 4.C, the variance converges with rate

$$\text{Var}[\widehat{L}_n] \approx O(n^{-1}) + O(n^{-\alpha}) = O(n^{-\alpha}), \quad (4.22)$$

while the bias converges with rate

$$\text{E}[\epsilon_n] = O(n^{1-\alpha})^{-2\beta_{\mathcal{T}}} + O(n^{1-\alpha})^{2\beta_{\mathcal{T}}-2}. \quad (4.23)$$

Plugging Equation (4.22) and (4.23) into Equation (4.21), we can model the AMSE as

$$\text{AMSE}(\widehat{L}_n) = O(n^{-\alpha}) + \left(O(n^{1-\alpha})^{-2\beta_{\mathcal{T}}} + O(n^{1-\alpha})^{2\beta_{\mathcal{T}}-2}\right)^2. \quad (4.24)$$

Finally, by minimizing Equation (4.24) (see Appendix 4.D) we obtain the values for optimal asymptotic convergence $\beta_{\mathcal{T}} = 1/2$ and $\alpha = 2/3$, which by substitution gives us the final asymptotic convergence rate of our progressive transient photon beams

$$\text{AMSE}(\widehat{L}_n) = O(n^{-\frac{2}{3}}). \quad (4.25)$$

To generalize photon beams to the transient domain, we need to account for the duration of light paths. This requires considering propagation and scattering delays along the camera and light subpaths, but also the effect of time in the density estimation connecting these two subpaths.

Creating the photon map. We compute the photon propagation as a standard random walk through the scene, which can be modeled using the subpath formulation

defined by Jarabo et al. [106]. Let us define a light subpath $\bar{\mathbf{x}}_l = \mathbf{x}_0 \dots \mathbf{x}_{k-1}$, with k vertices, where \mathbf{x}_0 is the light source. This light path defines $k - 1$ photon beams, in which a beam b_j is defined by its origin at $\mathbf{x}_{b_j} = \mathbf{x}_j$ and direction $\omega_{b_j} = \frac{\mathbf{x}_{j+1} - \mathbf{x}_j}{\|\mathbf{x}_{j+1} - \mathbf{x}_j\|}$. Using Jarabo's definition of the path integral (and therefore of the contribution of the subpaths), we compute the flux of each photon as:

$$\Phi_{b_j} = \frac{f(\bar{\mathbf{x}}_j, \bar{\tau}_j)}{Mp(\bar{\mathbf{x}}_j, \bar{\tau}_j)} = \frac{L_e(\mathbf{x}_0 \rightarrow \mathbf{x}_1, \tau_0)T(\bar{\mathbf{x}}_j, \bar{\tau}_j)}{M \prod_{i=0}^j p(\mathbf{x}_i, \tau_i)}, \quad (4.26)$$

with $\bar{\mathbf{x}}_j$ the subpath of $\bar{\mathbf{x}}_l$ up the vertex j , f the subpath contribution function, $\bar{\tau}_j = \tau_0 \dots \tau_j$ the sequence of time delays up to vertex j , M the number of photon random walks sampled, $L_e(\mathbf{x}_0 \rightarrow \mathbf{x}_1, \tau_0)$ the emission function, $p(\mathbf{x}_i, \tau_i)$ the probability density of sampling vertex \mathbf{x}_i with time delay τ_i . The throughput, $T(\bar{\mathbf{x}}_j, \bar{\tau}_j)$, of subpath (\mathbf{x}_i, τ_j) is defined as:

$$T(\bar{\mathbf{x}}_j, \bar{\tau}_j) = \left[\prod_{i=1}^{j-1} f_s(\mathbf{x}_i, \tau_j) \right] \left[\prod_{i=0}^{j-1} G(\mathbf{x}_i, \mathbf{x}_{i+1}) V(\mathbf{x}_i, \mathbf{x}_{i+1}) \right], \quad (4.27)$$

with $f_s(\mathbf{x}_i, \tau_j)$ the scattering event at vertex \mathbf{x}_i with delay τ_j , and $G(\mathbf{x}_i, \mathbf{x}_{i+1})$ and $V(\mathbf{x}_i, \mathbf{x}_{i+1})$ the geometry and visibility terms between vertices \mathbf{x}_i and \mathbf{x}_{i+1} , respectively. Finally, for transient state we need to know the instant t_{b_j} at which the photon beam is created (through emission or scattering), defined as:

$$t_{b_j} = \sum_{i=0}^{j-1} \tau_j + \sum_{i=0}^{j-1} \Delta t(\mathbf{x}_i, \mathbf{x}_{i+1}). \quad (4.28)$$

Rendering. For rendering, we adapt Equation (4.8) to account for the temporal domain, as

$$L(\mathbf{x}_r, \omega_r, t) \approx \sum_{b \in R_b} L_b(\mathbf{x}_r, \omega_r, t), \quad (4.29)$$

with $L_b(\mathbf{x}_r, \omega_r, t)$ the radiance estimation for beam b to ray t at instant t . In essence, $L_b(\mathbf{x}_r, \omega_r, t)$ will return zero radiance if t is out of the temporal footprint of the density estimation kernel. Depending on the dimensionality of the density estimation, Jarosz and colleagues [109] proposed three different estimators based on 3D, 2D and 1D kernels. Since the 3D kernel results impractical due to costly 3D convolutions, we focus on 1D and 2D kernels (Equations (4.9) and (4.10)), and extend them to transient state, assuming homogeneous media.

Kernel 2D We generalize Jarosz's et al.'s 2D estimate L_b^{2D} (Equation (4.10)) by introducing a temporal function $W(t)$ as

$$L_b^{2D}(\mathbf{x}_r, \omega_r, t) = K_{2D}(R_b) \Phi_b f_s(\theta_b, t) \mu_s \frac{e^{-\mu_t(s_r^- - s_r^+)} (|\cos \theta_b| - 1) - 1}{e^{\mu_t(s_r^- + s_r^+)} \mu_t (|\cos \theta_b| - 1)} W_{2D}(t), \quad (4.30)$$

where $[s_r^-, s_r^+]$ are the limits of the ray-beam intersection (Figure 4.3a), θ_b is the angle between ω_b and ω_r , and $K_{2D}(R_b)$ is a canonical 2D kernel with radius R_b . The

temporal function $W_{2D}(t)$ models the temporal footprint of the 2D kernel as

$$W_{2D}(t) = \begin{cases} \frac{1}{t^+ - t^-} & \text{if } t \in (t^-, t^+) \\ 0 & \text{otherwise} \end{cases}, \quad (4.31)$$

where $t^- = t_b + t_r + \frac{\eta_m}{c}(s_r^- + s_b^-)$ and $t^+ = t_b + t_r + \frac{\eta_m}{c}(s_r^+ + s_b^+)$, and t_r and t_b are the initial times of the camera ray and beam, respectively. Note that due to transmittance, the photon energy varies as it travels across the blur region. Evenly distributing the integrated radiance L_b across this interval introduces temporal bias, in addition to the inherent spatial bias introduced by density estimation. However we observed this even distribution provides a good tradeoff between bias, variance, and computational overhead.

Kernel 1D In the 1D kernel defined for density estimation by Jarosz et al. the spatial blur is performed over a line. Therefore, the energy of the beam is just spread on the ray on a single point at $r(s_r)$, from a single point of the beam $b(s_b)$ (see Figure 4.3a). In consequence, $s_r^\pm \rightarrow s_r$ and $s_b^\pm \rightarrow s_b$, which implies that $t^\pm \rightarrow t_{br}$, and the temporal function reduces to $W_{1D}(t - t_b) = \delta(t)$, with $\delta(t)$ the Dirac delta function. With that in place, we transform Jarosz et al. 1D estimate to

$$L_b^{1D}(\mathbf{x}_r, \omega_r, t) = K_{1D}(R_b) \Phi_b f_s(\theta_b, t) \mu_s \frac{e^{-\mu_t s_b} e^{-\mu_t s_r}}{\sin \theta_b} \delta(t - t_b), \quad (4.32)$$

with $K_{1D}(R_b)$ a 1D kernel with radius R_b .

Implementation Since photon beams correspond to full photon trajectories, they allow us to estimate radiance at any position $\mathbf{x}_b + s\omega_b$ of the beam, and therefore at any arbitrary time $t(\mathbf{x}_b + s\omega_b)$. As mentioned, one-dimensional radiance estimate corresponds to a single time across the beam. In a traditional rendering process where camera rays are traced through view-plane pixels against the beams map, the temporal definition *within* a pixel will be proportional to the amount of samples per pixel taken. Additionally, 2D blur requires distributing every radiance estimate along a time interval, which reduces variance in the time dimension of a pixel at the expense of introducing additional temporal bias.

Finally, note that the temporal footprint of the density estimation might be arbitrarily small, so the probability of finding a beam b at an specific time might be very low. We alleviate this issue using path reuse via density estimation [106]. In particular, for the non-progressive results we use histogram temporal density estimation. In this technique, the samples in the temporal domain are reused across all frames by evaluating their contribution functions, which correspond to the temporal window covered by each frame. In Section 4.5 we introduce temporal kernel-based density estimation, and combine it with the spatial density estimation of the beam.

4.6 Results

In the following we illustrate the results of our proposed method in five scenes: CORNELL SPHERES, MIRRORS, PUMPKIN, SOCCER [211], PUMPKIN, and JUICE. See Figures 4.4, 4.1 (right), and 4.8 (left) for steady-state renders of the scenes. Results of Figures 4.5 and 4.6 were taken on a desktop PC with Intel i7 and 4GB RAM using a transient 2D kernel (Equation 4.30). Figures 4.1, 4.7, and 4.8 were rendered on an




Figure 4.4: Steady-state renders for the scenes CORNELL SPHERES (Figure 4.5), MIRRORS (Figure 4.6), and PUMPKIN (Figure 4.7).




Figure 4.5: Comparison of CORNELL SPHERES scene using *camera-unwarping* (top), where we do not take into account the camera time, and real propagation of light (bottom). In the bottom row the shape of the wavefront is altered by the camera time, as if we were scanning the scene from the viewpoint towards the furthest parts of the scene. Camera unwarping on the other hand illustrates more intuitively how light propagates locally.

Intel Xeon E5 with 256GB RAM, using our progressive spatio-temporal kernel density estimations (Section 4.5) derived from the transient spatial 1D kernel (Equation 4.32). In each iteration, we use a fixed radius for our spatio-temporal density estimators (instead of using a nearest neighbor approach). Please refer to the supplemental video for the full sequences of all the scenes.

Figure 4.5 shows a Cornell box filled with a scattering medium, and demonstrates the effect of *camera unwarping* [225] when rendering. Camera unwarping is an intuitive way of visualizing how light propagates *locally* on the scene without accounting for the time light takes to reach the camera. The scene consists of a diffuse Cornell box with a point light on the top, a glass refractive sphere (top, IOR = 1.5) and a mirror sphere (bottom). While Figure 4.5b shows the real propagation of light—including camera time—, Figure 4.5a depicts more intuitively how light comes out from the point light, travels through the refractive sphere, and the generated caustic bounces on the mirror sphere. Note how in the top sequence we can clearly see how light is slowed down through the glass sphere due to the higher index of refraction. We can also observe multiple scattered light (particularly noticeable in frames t=4ns and t=6ns) as a secondary wavefront.

Figure 4.6 compares visualizations of light propagation within the MIRRORS scene using Heaviside and Dirac delta light emission. The scene is composed by two colored mirrors and a glass sphere with IOR = 1.5, and was rendered using the previously mentioned camera unwarping. We can observe how delta emission generates wavefronts that go through the ball and bounce in the mirrors, creating wavefront holes where constant emission creates medium shadows. In the last frame of the top row Delta emission clearly depicts the slowed down caustic through the glass ball respect to the main wavefront.

Our progressive method combines time-resolved 1D spatial kernels of photon beams and temporal density estimations, reducing bias while providing consistent solutions in the limit with an optimal convergence rate of $O(n^{-\frac{2}{3}})$. In Figure 4.7 we analyze its convergence with respect to progressive transient path tracing with temporal KDE [106] (PTPT). In the middle graph we show the temporal profile on a single




Figure 4.6: Comparison between Dirac delta (top) and continuous (Heaviside) emission (bottom). Dirac delta emission lets us see how a pulse of light travels and scatters across the scene, depicting the light wavefronts bouncing on the mirrors and going through the glass ball. Continuous emission shows how light is emitted until it reaches every point in the scene, as if we were taking a picture with a camera at very slow-motion.




Figure 4.7: The PUMPKIN scene shows a jack o'lantern embedding a point light that creates hard shadows through the holes. The left frames show a sequence of the time-resolved renders after 4096 iterations of our algorithm (10k beams / iteration), and temporal KDE on a progressive transient path tracer (PTPT, 16spp / iteration) [106]. The middle plot compares the whole temporal footprint at the pink marker. Reference solution (dark grey) was obtained with a transient path tracer (no KDE) using 64M samples per pixel. Right plot shows MSE convergence with respect to the number of progressive iterations (in log-log scale), at 1 minute/iteration on each algorithm. As expected, the convergence of our method ($O(n^{-\frac{2}{3}})$) is slower than PTPT ($O(n^{-\frac{4}{5}})$); however, as shown in the equal-time comparison, our algorithm presents better temporal behavior with much less variance on later timings.

pixel for both our algorithm and PTPT after 4096 equal-time iterations, where both algorithms converge to the reference solution taken with transient path tracing (no temporal KDE) with 64 million samples. While PTPT presents faster convergence (see Figure 4.7, right graph), our algorithm presents a better behavior over time where variance increases due to the lack of samples (center graph). Additionally, it requires much fewer iterations than PTPT to achieve a similar MSE (see log-log right graph).

In Figure 4.1 we show a more complex scenario, with different caustics rendered, with our progressive algorithm. It contains a smooth dielectric figurine with different transmission albedos placed within a participating medium with an isotropic phase function. Our method is capable of handling complex caustics transmitted from light sources through the player, and then through the ball. Our algorithm progressively reduces bias and variance to provide a consistent solution.

Finally in Figure 4.8 we illustrate a setup combining different media properties, and specular refractive and reflective materials. The liquid has a very forward phase function, making the light first travel through the direction of the stream ($t = 4.6$ ns), and then going through the liquid inside the glass ($t = 5.1$ ns to $t = 6.3$ ns). The mirror surface makes the light to bounce back to the surrounding medium as a caustic through the water spills and ice cubes at $t = 5.1$ ns and $t = 6.6$ ns. Note that these are not fully observable in the steady-state render (left) due to the accumulated radiance from the surrounding medium and the adjusted exposure of the image.




Figure 4.8: We illustrate the potential of our method in the JUICE scene [17], which presents a scene very difficult to render for path tracing methods, but well-handled by photon-based methods. The scene is filled by a thin participating medium, while the glass contains ruby grapefruit juice as measured by Narasimhan et al. [159]. The highly forward phase function of the juice, as well as the delta interactions on the glass, ice cubes, and the mirror floor surface, generate complex caustic patterns which our method is able to simulate in transient state. Bottom row has increased exposure respect to top row to show the radiance at later timings.

4.7 Conclusions

In this chapter we have presented a robust progressive method for efficiently rendering transient light transport with consistent results. We derived our method based on progressive photon beams [110], extending its density estimators to account for light time-of-flight, and deriving a new progressive scheme. We then compute the convergence of the method, and derive the parameters for optimal asymptotic convergence. Our results demonstrate that combining continuous photon trajectories in transient state and our optimal spatio-temporal convergence rates allow to robustly compute a noise-free solutions to the time-resolved RTE for complex light paths. We believe that our work might be very useful for developing new techniques for transient imaging and reconstruction in media, as well as to obtain new insights on time-resolved light transport.

As future work it would be interesting to analyze more thoroughly the optimal performance and kernels for variance reduction and bias impact in transient state, under varying media characteristics. In addition, extending our method to leverage recent advances in media transport, such as transient-state adaptations of higher-dimensional photon estimators [19] as well as hybrid techniques [131], could improve performance of time-resolved rendering for a general set of geometries and media characteristics.

4.A Error in transient progressive photon beams

Here we analyze the consistency of the transient progressive photon beams algorithm described in Section 4.5. For our analysis on the error of the estimate, we use the *asymptotic mean squared error* (AMSE) defined as

$$\text{AMSE}(\widehat{L}_n) = \text{Var}[\widehat{L}_n] + \text{E}[\epsilon_n]^2, \quad (4.33)$$

where $\text{Var}[\widehat{L}_n]$ is the variance of the estimate and $\text{E}[\epsilon_n]$ is the bias at iteration n . We model $\text{Var}[\widehat{L}_n]$ as [127]

$$\text{Var}[\widehat{L}_n] = \frac{1}{n} \text{Var}[\Psi L] + \frac{1}{n^2} \sum_{j=1}^n \text{Var}[\Psi \epsilon_j], \quad (4.34)$$

where Ψ is the contribution of the eye ray, and ϵ_j is the bias for iteration j . The first term is the standard variance of the Monte Carlo estimate, which is unaffected by the kernel. The second term, on the other hand, is the variance of the error, and is dependent on density estimation. On the other hand, the estimated value of the error (bias) $\text{E}[\widehat{L}_n]$ is defined as

$$\text{E}[\widehat{L}_n] = L + \text{E}[\Psi] \text{E}[\epsilon_n], \quad (4.35)$$

where $\text{E}[\epsilon_n]$ is the bias of the estimator after n steps:

$$\text{E}[\epsilon_n] = \frac{1}{n} \sum_{j=1}^n \text{E}[\epsilon_j], \quad (4.36)$$

with $\text{E}[\epsilon_j]$ the expected error at iteration j . In the following, we first derive the variance and expected value of the error for a single iteration. Then, we analyze the asymptotic behavior of these terms, and compute the values for optimal convergence for $\beta_{\mathcal{T}}$, β_R and α .

4.B Variance and expected value of the error of the time-resolved beam radiance estimate

We first analyze the variance and expected value of the error (bias) introduced by the radiance estimate at each iteration. Let us first define the error in each iteration as:

$$\begin{aligned} \epsilon &= \widehat{L}_n(\mathbf{x}_r, \omega_r, t) - L(\mathbf{x}_r, \omega_r, t) \\ &= \sum_{i=1}^M K_{1D}(R_b) K_{\mathcal{T}}(t - t_i) \Phi_i - L(\mathbf{x}_r, \omega_r, t). \end{aligned} \quad (4.37)$$

Variance. We first define the variance of the error $\text{Var}[\epsilon]$ as (in the following, we omit dependences for clarity):

$$\begin{aligned} \text{Var}[\epsilon] &= \text{Var}\left[\sum_{i=1}^M K_{1D} K_{\mathcal{T}} \Phi - L\right] \\ &= (\text{Var}[K_{1D}] + \text{E}[K_{1D}]^2) (\text{Var}[K_{\mathcal{T}}] + \text{E}[K_{\mathcal{T}}]^2) \\ &\quad (\text{Var}[\Phi] + \text{E}[\Phi]^2) - \text{E}[K_{1D}]^2 \text{E}[K_{\mathcal{T}}]^2 \text{E}[\Phi]^2, \end{aligned} \quad (4.38)$$

In order to compute the variance of the error $\text{Var}[\epsilon]$ we need to make a set of assumptions: First, we assume that the beams' probability density is constant within the kernel K_{1D} in the spatial domain [110], and within $K_{\mathcal{T}}$ in the temporal domain [106]. We denote these probabilities as p_{R_b} and $p_{\mathcal{T}}$ respectively. We also assume that the distance between view ray and photon beam, time t_b and beams' energy Φ_i are independent samples of the random variables D , T and Φ , respectively, which are mutually independent. Finally, we assume that D and T have probability densities p_{R_b} and $p_{\mathcal{T}}$.

With these assumptions, and taking into account that $E[K_{1D}] = p_{R_b}$ and $E[K_{\mathcal{T}}] = p_{\mathcal{T}}$, we can model the the variance introduced by the temporal kernel $\text{Var}[K_{\mathcal{T}}]$ as [106]

$$\text{Var}[K_{\mathcal{T}}] = \frac{p_{\mathcal{T}}}{\mathcal{T}} \int_{\mathbb{R}} k_{\mathcal{T}}(\psi)^2 d\psi - p_{\mathcal{T}}^2, \quad (4.39)$$

where we express $K_{\mathcal{T}}$ as a canonical kernel $k_{\mathcal{T}}$ with unit integral such that $K_{\mathcal{T}}(\xi) = k_{\mathcal{T}}(\xi/\mathcal{T})\mathcal{T}^{-1}$. Analogously, $\text{Var}[K_{1D}]$ is, by [110],

$$\text{Var}[K_{1D}] = \frac{p_{R_b}}{R_b} \int_{\mathbb{R}} k_{1D}(\psi)^2 d\psi - p_{R_b}^2. \quad (4.40)$$

This allow us to express the variance of the error $\text{Var}[\epsilon]$ as:

$$\text{Var}[\epsilon] \approx (\text{Var}[\Phi] + E[\Phi]^2) \left(\frac{p_{R_b}}{R_b} \mathcal{C}_{1D} \right) \left(\frac{p_{\mathcal{T}}}{\mathcal{T}} \mathcal{C}_{\mathcal{T}} \right), \quad (4.41)$$

where \mathcal{C}_{1D} and $\mathcal{C}_{\mathcal{T}}$ are kernel-dependent constants. The last term can be neglected by assuming that the kernels cover small areas in their respective domains, which effectively means that $\mathcal{C}_{1D} \gg p_{R_b}$ and $\mathcal{C}_{\mathcal{T}} \gg p_{\mathcal{T}}$. Equation (4.41) shows that for transient density estimation, the variance $\text{Var}[\epsilon]$ is inversely proportional to $R_b \mathcal{T}$.

Bias. Bias at each iteration j is defined as the expected value of the error $E[\epsilon_j]$ as

$$\begin{aligned} E[\epsilon_j] &= E\left[\sum_{i=1}^M K_{1D} K_{\mathcal{T}} \Phi - L\right] \\ &= E[K_{1D}] E[K_{\mathcal{T}}] E[\Phi] - L. \end{aligned}$$

Using a second-order expansion of $p_{\mathcal{T}}$ and p_{R_b} , instead of the zeroth-order used when modeling variance, we can express the expected value of $K_{\mathcal{T}}$ as [106]

$$E[K_{\mathcal{T}}] \approx p_{\mathcal{T}} + \mathcal{T}^2 \int_{\mathbb{R}} k_{\mathcal{T}}(\psi) O(\|\psi\|^2) d\psi = p_{\mathcal{T}} + \mathcal{T}^2 \mathcal{C}_{\mathcal{T}}^{ii}, \quad (4.42)$$

while the expected value of K_{1D} is, by [110],

$$E[K_{1D}] \approx p_{R_b} + R_b \int_{\mathbb{R}^2} k_{1D}(\psi) O(\|\psi\|^2) d\psi = p_{R_b} + R_b \mathcal{C}_{1D}^{ii}, \quad (4.43)$$

where $\mathcal{C}_{\mathcal{T}}^{ii}$ and \mathcal{C}_{1D}^{ii} are constants dependent on the higher-order derivatives of the spatio-temporal light distribution. Using (4.42) and (4.43), and $L = p_{R_b} p_{\mathcal{T}} E[\Phi]$ we

finally compute $E[\epsilon_j]$ for iteration j as

$$\begin{aligned} E[\epsilon_j] &\approx (p_{R_b} + R_b^2 C_{1D}^{ii})(p_{\mathcal{T}} + \mathcal{T}^2 C_{\mathcal{T}}^{ii})E[\Phi] - p_{R_b} p_{\mathcal{T}} E[\Phi] \\ &= E[\Phi](p_{R_b} \mathcal{T}^2 C_{\mathcal{T}}^{ii} + p_{\mathcal{T}} R_b^2 C_{1D}^{ii} + \mathcal{T}^2 C_{\mathcal{T}}^{ii} R_b^2 C_{1D}^{ii}). \end{aligned} \quad (4.44)$$

4.C Convergence analysis of progressive transient photon beams

Based on the expressions for $\text{Var}[\epsilon]$ and $E[\epsilon_j]$ defined above (Equations (4.41) and (4.44)), we can now derive the asymptotic behaviour of Equation (4.21). For that, we will compute the variance $\text{Var}[\widehat{L}_n]$ and bias $E[\epsilon_n]$ after n iterations.

Variance. Assuming that the random variables Ψ and ϵ_j are independent, we model the variance of the estimator $\text{Var}[\widehat{L}_n]$ in Equation (4.34) as [127],

$$\begin{aligned} \text{Var}[\widehat{L}_n] &= \frac{1}{n} \text{Var}[\Psi L] + \frac{1}{n^2} \sum_{j=1}^n \text{Var}[\Psi \epsilon_j] \\ &= \frac{1}{n} \text{Var}[\Psi L] + \text{Var}[\Psi] \frac{1}{n^2} \sum_{j=1}^n \text{Var}[\epsilon_j] + \\ &\quad E[\Psi]^2 \frac{1}{n^2} \sum_{j=1}^n \text{Var}[\epsilon_j] + \text{Var}[\Psi] \frac{1}{n^2} \sum_{j=1}^n E[\epsilon_j]^2. \end{aligned} \quad (4.45)$$

Following [118], we can approximate $\text{Var}[\epsilon_n]$ as a function of the variance at the first iteration $\text{Var}[\epsilon_1]$ as:

$$\text{Var}[\epsilon_n] \approx \frac{\text{Var}[\epsilon_1]}{(2 - \alpha)n^\alpha} = O(n^{-\alpha}). \quad (4.46)$$

Finally, by applying $\text{Var}[\epsilon_n]$ and asymptotic simplifications, we can formulate $\text{Var}[\widehat{L}_n]$ (4.46) as

$$\begin{aligned} \text{Var}[\widehat{L}_n] &\approx \frac{1}{n} \text{Var}[\Psi L] + E[\Psi]^2 \text{Var}[\epsilon_n] \\ &\approx \frac{1}{n} \text{Var}[\Psi L] + \frac{\text{Var}[\epsilon_1]}{(2 - \alpha)n^\alpha} \\ &= O(n^{-1}) + O(n^{-\alpha}) = O(n^{-\alpha}). \end{aligned} \quad (4.47)$$

Bias. The expected value of the error $E[\epsilon_n]$ is modeled in Equation (4.35) as a function of the averaged bias introduced at each iteration $E[\epsilon_j]$ (4.44). Computing the kernels' bandwidth \mathcal{T}_j and R_{bj} at iteration j by expanding Equation (4.20) as a function of their initial value by we get

$$\mathcal{T}_j = \mathcal{T}_1 (j \propto B(\alpha, j))^{-\beta_{\mathcal{T}}}, \quad (4.48)$$

$$R_{bj} = R_{b1} (j \propto B(\alpha, j))^{-\beta_{R_b}}, \quad (4.49)$$

where $B(x, y)$ is the Beta function. Using (4.48) and (4.49) in Equation (4.44) we can express $E[\epsilon_j]$ as a function of the initial kernel bandwidths

$$\begin{aligned} E[\epsilon_j] &= E[\Phi] p_{R_b} C_{\mathcal{T}}^{ii} \mathcal{T}_1^2 \Theta(j^{1-\alpha})^{-2\beta_{\mathcal{T}}} \\ &\quad + E[\Phi] p_{\mathcal{T}} C_{1D}^{ii} R_{b_1}^2 \Theta(j^{1-\alpha})^{-2\beta_{R_b}} \\ &\quad + E[\Phi] C_{\mathcal{T}}^{ii} C_{1D}^{ii} \mathcal{T}_1^2 R_{b_1}^2 \Theta(j^{1-\alpha})^{-2(\beta_{\mathcal{T}} + \beta_{R_b})}. \end{aligned} \quad (4.50)$$

Finally, we use $\sum_{j=1}^n \Theta(j^x) = n O(n^x)$ to plug Equation (4.50) into Equation (4.36) to get the asymptotic behavior of $E[\epsilon_n]$ in transient progressive photon beams as

$$E[\epsilon_n] = O(n^{1-\alpha})^{-2\beta_{\mathcal{T}}} + O(n^{1-\alpha})^{-2\beta_{R_b}} + O(n^{1-\alpha})^{-2(\beta_{\mathcal{T}} + \beta_{R_b})},$$

which, by using the equality $\beta_{R_b} = 1 - \beta_{\mathcal{T}}$, becomes

$$\begin{aligned} E[\epsilon_n] &= O(n^{1-\alpha})^{-2\beta_{\mathcal{T}}} + O(n^{1-\alpha})^{2\beta_{\mathcal{T}}-2} + O(n^{1-\alpha})^{-2} \\ &= O(n^{1-\alpha})^{-2\beta_{\mathcal{T}}} + O(n^{1-\alpha})^{2\beta_{\mathcal{T}}-2}. \end{aligned} \quad (4.51)$$

4.D Minimizing Asymptotic Mean Squared Error

Using the asymptotic expression for variance and bias in Equations (4.47) and (4.51), we can express the AMSE (4.21) as

$$AMSE(\widehat{L}_n) = O(n^{-\alpha}) + \left(O(n^{1-\alpha})^{-2\beta_{\mathcal{T}}} + O(n^{1-\alpha})^{2\beta_{\mathcal{T}}-2} \right)^2. \quad (4.52)$$

which is a function of the parameters α and $\beta_{\mathcal{T}}$. Given that the variance is independent of $\beta_{\mathcal{T}}$, we first obtain the optimal value for this parameter that yields the highest convergence rate of the bias $E[\epsilon_n]$. We differentiate Equation (4.51), apply asymptotic simplifications and equating to zero, we obtain the optimal value $\beta_{\mathcal{T}} = 1/2$. By plugging this value in Equation (4.52), we obtain

$$AMSE(\widehat{L}_n) = O(n^{-\alpha}) + O(n^{-2(1-\alpha)}). \quad (4.53)$$

Finally, by finding the minimum again with respect to α we get the optimal parameter $\alpha = 2/3$, which results in the optimal convergence rate of the AMSE for our transient progressive photon beams as

$$AMSE(\widehat{L}_n) = O(n^{-\frac{2}{3}}) + O(n^{-2(1-\frac{2}{3})}) = O(n^{-\frac{2}{3}}). \quad (4.54)$$

Chapter 5

Non-line-of-sight imaging using Phasor Fields

This chapter introduces a new method for visualizing scenes which are outside the direct line of sight of the observer, which is performed using transient information captured using ultrafast imaging devices [40]. We introduce a new method based on virtual wave optics, which converts a delta laser pulse propagated across the hidden scene using secondary visible surfaces into a computationally modulated signal, which then can be computationally focused into a virtual image of the hidden scene. Compared with previous work, our method makes just the bare minimum assumptions about light transport in the hidden scene, recovers a considerable amount of detail, and it's robust to noise and other capture artifacts. We provide both the theoretical foundations of our technique, and empirical results generated from real transient captures and simulations.

The project was done in collaboration with our colleagues at the Computational Optics Group in the University of Wisconsin-Madison, Madison, which developed the initial idea and were in charge of capturing real scenes. My role as the second author of this project focused on improving the implementation of the computational camera systems, and generating the visualizations of the scene reconstructions that are present across the chapter. Our group was also in charge of generating transient simulations, as those present in Section 5.E.3 and Section 5.E.4, which were used as a testbed to validate both the method and the capture system.

This work was originally published in Nature, and presented at the 29th Congreso Español de Informática Gráfica (CEIG 2019). An extension of this work was presented at the 45th International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2020).

Non-line-of-sight imaging using phasor-field virtual wave optics

Xiaochun Liu, Ibón Guillén, Marco La Manna, Ji Hyun Nam, Syed Azer Reza, Toan Huu Le, Diego Gutierrez, Adrian Jarabo & Andreas Velten
Nature, Vol. 572, 2019

On the effect of reflectance on phasor field non-line-of-sight imaging

Ibón Guillén, Xiaochun Liu, Andreas Velten, Diego Gutierrez & Adrian Jarabo
Proceedings of the 2020 IEEE International Conference on Acoustics, Speech and Signal Processing

5.1 Introduction

Non-Line-of-Sight (NLOS) imaging allows to observe objects partially or fully occluded from direct view, by analyzing indirect diffuse reflections off a secondary, relay surface. Despite its many potential applications [7; 30; 80; 92; 119; 125; 135; 164; 224], existing methods lack practical usability due to several shared limitations, including the assumption of single scattering only, lack of occlusions, and diffuse reflectance. Line of sight imaging systems, on the other hand, can address these and other imaging challenges despite relying on the mathematically simple processes of linear diffractive wave propagation. In this work we show that the NLOS imaging problem can also be formulated as a diffractive wave propagation problem. This allows to image NLOS scenes from raw time of flight data by applying the mathematical operators that model wave propagation inside a conventional line of sight (LOS) imaging system. By doing this, we have developed a method that yields a new class of reconstruction algorithms mimicking the various capabilities of LOS cameras. To demonstrate our method, we derive three imaging algorithms, each with its own unique novel capabilities, modeled after three different LOS imaging systems. These algorithms rely on solving wave diffraction integrals, namely the Rayleigh-Sommerfeld Diffraction (RSD) integral. Fast solutions to RSD and its approximations are readily available, directly benefiting our method. We demonstrate for the first time NLOS reconstruction of complex scenes with strong multiple scattering and ambient light, arbitrary materials, large depth range, and occlusions. Our method handles these challenging cases without explicitly developing a light transport model. We believe that our approach will help unlock the potential of NLOS imaging, and the development of novel applications not restricted to lab conditions, as shown in our results.

5.2 Related work

We have recently witnessed large advances in transient imaging techniques [107], employing streak cameras [225], gated sensors [135], amplitude-modulated continuous waves [79], single-photon detectors (SPAD) [162], or interferometry [67]. Access to time-resolved image information has in turn led to advances in imaging of objects partially or fully hidden from direct view (NLOS imaging) [30; 80; 92; 125; 135; 139; 221; 224; 239; 241]. Other methods are able to use information encoded in the phase of continuous light and do not use time of flight [119]. In the basic configuration of an NLOS system, light bounces off a relay wall, travels to the hidden scene, then propagates back to the relay wall, and finally reaches the sensor.

Recent NLOS reconstruction methods are based on heuristic filtered backprojection [30; 80; 134; 135; 224], or attempt to compute inverse operators of simplified forward light transport models [91; 92; 164]. These simplified models do not take into account multiple scattering, surfaces with anisotropic reflectance or, with a few exceptions [91], occlusions and clutter. Moreover, the depth range that can be recovered is also limited, partially due to the difference in intensity between first- and higher-order reflections. Existing methods are thus limited to carefully controlled cases, imaging isolated objects of simple geometry with moderate or no occlusion. Moreover, while the goal of previous works is limited to the reconstruction of hidden geometry, we develop a novel theoretical framework for general NLOS imaging,

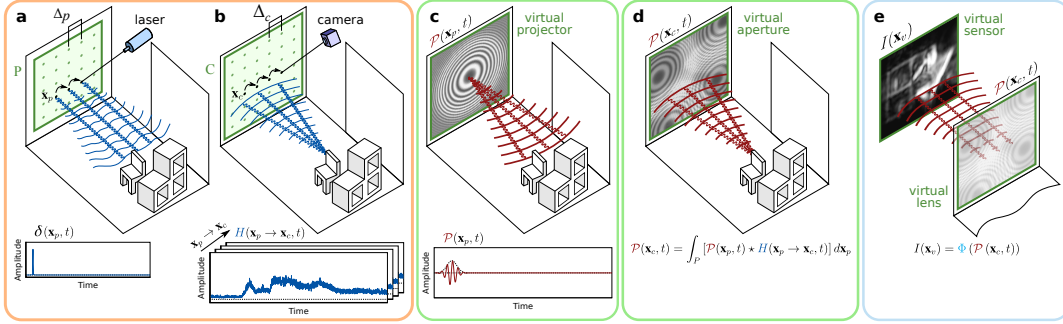


Figure 5.1: NLOS as a virtual LOS imaging system. Capturing scene data: **a**, A pulsed laser sequentially scans a relay wall; **b**, the light reflected back from the scene onto the wall is recorded at the sensor yielding an impulse response H of the scene. **c**, Virtual light source: The phasor field wave of a virtual light source $\mathcal{P}(\mathbf{x}_p, t)$ is modeled after the wavefront of the light source of the template LOS system. **d**, The scene response to this virtual illumination $\mathcal{P}(\mathbf{x}_c, t)$ is computed using H . **e**, The scene is reconstructed from the wavefront $\mathcal{P}(\mathbf{x}_c, t)$ using wave diffraction theory. The function $\Phi(\cdot)$ is also taken from the template LOS system.

reconstructing the irradiance at a virtual sensor; this enables applications beyond geometry reconstruction, as we demonstrate in this work. Our data and reconstruction code can be found in a figshare repository [141].

5.3 Phasor Fields in non-line-of-sight imaging

Time-of-flight LOS imaging has used a phasor formalism together with Fourier domain ranging [79] to describe the emitted modulated light signal. Kadambi et al. [116] extended this concept to reconstruct NLOS scenes using phasors to describe hardware intensity modulation. We show that a similar description can be used to model the physics of light transport through the scene. The key insight of our method is that propagation through a scene of intensity-modulated light can be modeled using a Rayleigh-Sommerfeld diffraction (RSD) operator acting on a quantity we term the *phasor field*. This allows us to formulate any NLOS imaging problem as a wave imaging problem (Figure 5.1), and to transfer well-established insights and techniques from classic optics into the NLOS domain. Given a captured time-resolved dataset of light transport through a NLOS scene, and a choice of a *template* LOS imaging system, our method provides a recipe that results in a NLOS imaging algorithm mimicking the capabilities of the corresponding LOS system. This template system can be any real or hypothetical wave imaging system that includes a set of light sources and detectors. The resulting algorithms can then be efficiently solved using diffraction integrals like the RSD, for which a variety of fast exact and approximate solvers exist [194].

We start by mathematically defining our phasor field $\mathcal{P}(\vec{x}, t)$. Let $\mathcal{E}(\mathbf{x}, t)$ [$\sqrt{\text{Wm}^{-2}}$] be a quasi-monochromatic scalar field at position $\mathbf{x} \in \mathcal{S}$ and time t , incident on (or reflected from) a Lambertian surface \mathcal{S} , with center frequency Ω_0 and bandwidth $\Delta\Omega \ll \Omega_0$. We can then define

$$\mathcal{P}(\mathbf{x}, t) \equiv \left\langle \frac{1}{\tau} \int_{t-\tau/2}^{t+\tau/2} |\mathcal{E}(\mathbf{x}, t')|^2 dt' \right\rangle - \left\langle \frac{1}{T} \int_{t-T/2}^{t+T/2} |\mathcal{E}(\mathbf{x}, t')|^2 dt' \right\rangle \quad (5.1)$$

as the mean subtracted irradiance [Wm^{-2}] at point \mathbf{x} and time t . The $\langle \cdot \rangle$ operator

denotes spatial speckle averaging (for the reflected case) accounting for laser illumination, and τ represents the averaging of the intensity at a fast detector, with $\tau \ll 1/\Delta\Omega \ll T$. The second integral in the equation above is a long-term average intensity over an interval $T \gg \tau$ of the signal as seen by a conventional non-transient photodetector. Now, let us define the Fourier component of $\mathcal{P}(\mathbf{x}, t)$ for frequency ω as

$$\mathcal{P}_{0,\omega}(\mathbf{x}) \equiv \int_{-\infty}^{+\infty} \mathcal{P}(\mathbf{x}, t) e^{-i\omega t} dt, \quad (5.2)$$

from which we can define a monochromatic component of the phasor field $\mathcal{P}_\omega(\mathbf{x}, t)$ as

$$\mathcal{P}_\omega(\mathbf{x}, t) \equiv \mathcal{P}_{0,\omega}(\mathbf{x}) e^{i\omega t}. \quad (5.3)$$

Using the above, our phasor field $\mathcal{P}(\mathbf{x}, t)$ can be expressed as a superposition of monochromatic plane waves as $\mathcal{P}(\mathbf{x}, t) = \int_{-\infty}^{+\infty} \mathcal{P}_\omega(\mathbf{x}, t) d\omega/2\pi$. Since $\mathcal{P}(\mathbf{x}, t)$ is a real quantity, the Fourier components $\mathcal{P}_{0,\omega}(\mathbf{x})$ are complex and symmetric about $\omega = 0$. Note that in many places of this manuscript we assign $\mathcal{P}(\mathbf{x}, t)$ an explicitly complex value; in these cases it is implied that the correct real representation is $\frac{1}{2}(\mathcal{P}(\mathbf{x}, t) + \mathcal{P}^*(\mathbf{x}, t))$. In practice the complex conjugate can be safely ignored in our calculations. As can be seen in Appendix 5.A, given an isotropic source plane S and a destination plane D , and assuming that the electric field at S is incoherent, the propagation of its monochromatic component $\mathcal{P}_\omega(\mathbf{x}, t)$ is defined by an RSD-like propagation integral

$$\mathcal{P}_\omega(\mathbf{x}_d, t) = \gamma \int_S \mathcal{P}_\omega(\mathbf{x}_s, t) \frac{e^{ik|\mathbf{x}_d - \mathbf{x}_s|}}{|\mathbf{x}_d - \mathbf{x}_s|} d\mathbf{x}_s, \quad (5.4)$$

where γ is an attenuation factor, and $k = 2\pi/\lambda$ is the wave number for wavelength $\lambda = 2\pi/\omega$, $\mathbf{x}_s \in S$ and $\mathbf{x}_d \in D$. Note that, as described in Appendix 5.A, we approximate γ as a constant over the plane S as $\gamma \approx 1/|\langle S \rangle - \mathbf{x}_d|$; this approximation has a minor effect on the signal amplitude at the sensor, but does not change the phase of our phasor field. While Equation (5.4) is defined for monochromatic signals, it can be used to propagate broadband signals by propagating each monochromatic component independently; this can be efficiently done by time-shifting the phasor field (more details provided in Appendix 5.A.1).

The key insight of Equation (5.4) is that, given the assumption of constant γ , the propagation of our phasor field is defined by the same RSD operator as any other physical wave. Therefore, in order to image a scene from a *virtual camera* with aperture at plane C , we can apply the image formation model of any wave-based LOS imaging system directly over the phasor field $\mathcal{P}(\mathbf{x}_c, t)$ at the aperture, with $\mathbf{x}_c \in C$. The challenge is how to compute $\mathcal{P}(\mathbf{x}_c, t)$ from an illuminating input phasor field $\mathcal{P}(\mathbf{x}_p, t)$, where \mathbf{x}_p is a point in the *virtual projector* aperture P , given a particular NLOS scene the one shown in Figure 5.1.

Since light transport is linear in space and time-invariant [163; 191], we can characterize light transport through the scene as an impulse response function $H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)$, where \mathbf{x}_p and \mathbf{x}_c are the positions of the emitter and detector, respectively. The phasor field at the virtual aperture $\mathcal{P}(\mathbf{x}_c, t)$ can thus be expressed as a function of the input phasor field $\mathcal{P}(\mathbf{x}_p, t)$ and $H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)$ as

$$\mathcal{P}(\mathbf{x}_c, t) = \int_P [\mathcal{P}(\mathbf{x}_p, t) \star H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)] d\mathbf{x}_p, \quad (5.5)$$

where \star denotes the convolution operator. Any imaging system can be characterized




Figure 5.2: Reconstructions of a complex NLOS scene. **a**, Photograph of the scene as seen from the relay wall. The scene contains occluding geometries, multiple anisotropic surface reflectances, large depth, and strong ambient and multiply scattered light. **b**, 3D visualization of the reconstruction using phasor fields ($\lambda = 6$ cm). We include the relay wall location and the coverage of the virtual aperture for illustration purposes. **c**, Frontal view of the scene, captured with an exposure time of 10 ms per laser position. **d**, Frontal view captured with just a 1 ms exposure time (24 seconds for the complete scan).

by its image formation function $\Phi(\cdot)$, which transduces the incoming field into an image

$$I(\mathbf{x}_v) = \Phi(\mathcal{P}(\mathbf{x}_c, t)), \quad (5.6)$$

where \mathbf{x}_v is the point being imaged (i.e., the point at the *virtual sensor*). This in turn can be formulated as an RSD propagator, requiring to solve a diffraction integral in order to generate the final image.

In an NLOS scenario, $H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)$ usually corresponds to 5D transients acquired via an ultrafast sensor focused on \mathbf{x}_c , and sequentially illuminating the relay wall with short pulses at different points \mathbf{x}_p as seen in Figure 5.1, and detailed in Appendix 5.D. Points \mathbf{x}_p and \mathbf{x}_c correspond to a virtual LOS imaging system projected onto the relay wall. Once $H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)$ has been captured, both the wavefront $\mathcal{P}(\mathbf{x}_p, t)$ and the imaging operator $\Phi(\cdot)$ can be implemented computationally, so they are not bounded by hardware limitations. We can leverage this to employ different $\mathcal{P}(\mathbf{x}_p, t)$ functions from any existing LOS imaging system [70] to emulate its characteristics in an NLOS setting.

5.4 Results

We illustrate the robustness and versatility of our method by implementing three virtual NLOS imaging systems based on common LOS techniques: a conventional photography camera capable of imaging NLOS scenes without knowledge of the timing or location of the illumination source; a transient photography system capable of capturing transient videos of the hidden scene revealing higher-order inter-reflections beyond 3rd bounce; and a confocal time-gated imaging system robust to




Figure 5.3: Robustness of our technique. **a**, Reconstruction in the presence of strong ambient illumination (all the lights on during capture). **b**, Hidden scene with a large depth range, leading to very weak signals from objects farther away.

interreflections. An in-depth description of these example imaging systems is provided in Appendix 5.B, including their corresponding $\mathcal{P}(\mathbf{x}_p, t)$ functions and imaging operators $\Phi(\cdot)$, while Appendix 5.C describes some practical solver examples.

The spatial resolution of our virtual camera is $\Delta_x = 0.61\lambda L/d$, where d is the virtual aperture diameter, and L is the imaging distance (more details in Appendix 5.E.1). The distance Δ_p between sample points \mathbf{x}_p in P (see Figure 5.1) has to be small enough to sample H at the phasor field wavelength. We fix $\Delta_p = 1$ cm, and unless stated otherwise $\lambda = 4$ cm. The minimum sampling rate is $\Delta_p < \lambda/2$; in practice we found $\Delta_p = \lambda/4$ to provide the best trade-off between reconstruction noise and resolution.

The computational cost of our algorithm is bounded by the RSD solver computing the image formation model $\Phi(\cdot)$. Fast diffraction integral solvers exist [194], with complexity $O(N^3 \log(N))$. For the particular case of our confocal system, we formulate the algorithm as a backprojection, as detailed in Appendix 5.C.2 for details, so we are bounded by the computational cost of the backprojection algorithm used.

One common application of NLOS imaging is the reconstruction of hidden geometry. Figure 5.2 shows the result on a complex scene with our virtual confocal camera. The scene contains multiple objects with occlusions distributed over a large depth, a wide range of surface reflectances and albedos, and strong interreflections. Despite this challenging scenario, our method is able to image many details of the scene, at the correct depths, even with an ultra-short, 1 ms exposure. More analysis on the robustness of our method to capture noise can be found in Appendix 5.E.2. For simpler scenes (no occlusions, limited depth, controlled reflectance, and no interreflections), our method yields results on par with current techniques, which already approach theoretical limits regarding reconstruction quality, as explained in Appendix 5.D.

In Figure 5.3 we demonstrate the robustness of our method when dealing with other challenging scenarios, including strong multiple scattering and ambient illumination (5.3.a), or a high dynamic range from objects spanning very large depths range (5.3.b). Last, our method also allows to implement novel NLOS imaging systems and applications, leveraging the wealth of tools and processing methods available in LOS imaging. Figure 5.4(top) demonstrates NLOS refocusing using our virtual photography camera, computed with both the exact RSD operator and using a faster Fresnel approximation, while the bottom row shows frames of NLOS femto-photography reconstructed using our virtual transient photography system, revealing 4th and 5th bounce components in the scene. The first, second and fourth frames, in green, show how light first illuminates the chair, then propagates to the shelf, and finally hits the




Figure 5.4: Additional novel NLOS imaging applications of our method. (Top) NLOS refocusing: The hidden letters (left) are progressively brought in and out of focus as seen from a virtual photography camera at the relay wall, using the exact lens integral (blue), and the faster Fresnel approximation (red). **(Bottom)** NLOS transient video: Example frames of light traveling through a hidden office scene when illuminated by a pulsed laser. Timestamps indicate the propagation time from the relay wall. Frames with a green border show third bounce objects, frames on orange show 4th and 5th bounce effects.

back wall three meters away. The frames in orange show higher order bounces. The third frame shows the chair being illuminated *again* by light bouncing back from the relay wall, while the last two frames show how the pulse of light travels from the wall back to the scene (see the supplementary video). A description of the Fresnel approximation to the RSD operator, as well as the LOS projector-camera functions used in these examples, appear in Appendices 5.C.1 and 5.B.2.

In the Appendix 5.E.3, we include comparisons against ground truth for two synthetic scenes, inside a corridor of 2 m x 2 m x 3 m to create interreflections, simulated using an open source transient renderer [106]; these scenes are included in a publicly available database [63]. We analyze the robustness of our method with and without such interreflections; the reconstruction mean square error does not increase, remaining below 5 mm. Last, in Appendix 5.E.4 we progressively vary the specularity of the hidden geometry, from purely Lambertian to highly specular; again, the quality of the reconstructions does not vary significantly (mean square error of about 2 mm).

The examples shown highlight the primary benefit of our approach: By turning NLOS into a virtual LOS system, the intrinsic limitations of previous approaches no longer apply, enabling a new class of NLOS imaging methods that leverage existing wave-based imaging methods. Formulating NLOS light propagation as a wave does not impose limitations on the types of problems that can be addressed, nor the datasets that can be used. Any signal can be represented as a superposition of phasor field waves; our formulation thus can be viewed as a choice of basis to represent any kind of NLOS data. Expressing the NLOS problem this way allows to create a direct analogy to LOS imaging, which can be exploited to derive suitable imaging algorithms, and to implement them efficiently.

5.5 Conclusions

We have shown three novel imaging algorithms derived from our method. Our results include significantly more complex scenes than any NLOS reconstruction

shown so far in the literature, as well as novel applications. In addition, our approach is flexible, fast, memory-efficient, and lacks computational complexity since it does not require inverting a light transport model. In the future, we anticipate its application to other LOS imaging systems to, for instance, separate light transport into direct and global components, or utilize the phase of \mathcal{P}_ω for enhanced depth resolution. Our virtual imaging system could also be used to create a *second* virtual imaging system to see around two corners, assuming the presence of a secondary relay Lambertian surface in the hidden scene, or to select and manipulate individual light paths to isolate specific aspects of the light transport in different NLOS scenes. In that context, combining our theory with light transport inversions, via, e.g., an iterative approach, could potentially lead to better results, and is an interesting avenue of future work.

5.A Phasor Field propagation

Here we derive the Rayleigh-Sommerfeld diffraction (RSD) integral for the phasor field. Consider a point light source at a location \mathbf{x}_s that emits light with a sinusoidal time-varying irradiance $L(\mathbf{x}, t) = \text{Re}[L_0 (e^{-i\omega t} + 1)]$ with amplitude L_0 and modulation frequency ω . More formally, $L(t)$ and L_0 relate with the electromagnetic field $\mathcal{E}(\mathbf{x}, t)$ by

$$L(\mathbf{x}, t) = \left\langle \frac{1}{\tau} \int_{t-\tau/2}^{t+\tau/2} |\mathcal{E}(\mathbf{x}, t')|^2 dt' \right\rangle \quad \text{and} \quad L_0(\mathbf{x}) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-T/2}^{+T/2} |\mathcal{E}(\mathbf{x}, t)|^2 dt,$$

with τ a sufficiently small value. The operator $\langle \cdot \rangle$ is the spatial averaging operator that takes into account multiple possible measurements for e.g. removing the effect of laser's speckle.

We define the real-valued phasor field $\mathcal{P}(\mathbf{x}, t)$ at a point in space as

$$\mathcal{P}(\mathbf{x}, t) = L(\mathbf{x}, t) - L_0(\mathbf{x}). \quad (5.7)$$

Since $L(\mathbf{x}, t)$ is modulated with a single frequency ω , this allows us to consider $\mathcal{P}(\mathbf{x}_s, t)$ as a monochromatic phasor field wave \mathcal{P}_ω emitted from a point light source at location \mathbf{x}_s , with amplitude $\mathcal{P}_{0,\omega}(\mathbf{x}_s)$ and oscillating at a frequency ω :

$$\mathcal{P}(\mathbf{x}_s, t) = \mathcal{P}_\omega(\mathbf{x}_s, t) = \mathcal{P}_{0,\omega}(\mathbf{x}_s) e^{i\omega t}. \quad (5.8)$$

In the following, wherever we write an explicitly complex expression for $\mathcal{P}(\mathbf{x}_s, t)$, it is implied that the actual real phasor field is $\frac{1}{2}(\mathcal{P}(\mathbf{x}_s, t) + \mathcal{P}^*(\mathbf{x}_s, t))$. In practice we can safely ignore the complex conjugate component in all our computations. Note that the constant $L_0(\mathbf{x})$ term is only necessary to link the phasor field wave to a measurable physical quantity, since real intensities cannot be negative. We can think of it as the monochromatic wave component at frequency $\omega = 0$. Since our propagator is linear and does not mix different frequency components it can safely be ignored as it can only create zero-frequency contributions.

To determine the light intensity and thereby the phasor field at any point in space and time (\mathbf{x}_d, t) we have to account for the travel time from \mathbf{x}_s to \mathbf{x}_d , defined as $t_p = |\mathbf{x}_d - \mathbf{x}_s|/c$, with c the propagation speed, and the radial drop-off in light intensity:

$$\mathcal{P}_\omega(\mathbf{x}_d, t) = \mathcal{P}_{0,\omega}(\mathbf{x}_s) \frac{e^{i\omega(t+t_p)}}{|\mathbf{x}_d - \mathbf{x}_s|^2} = \mathcal{P}_{0,\omega}(\mathbf{x}_s) \frac{e^{i\omega(t+|\mathbf{x}_d - \mathbf{x}_s|/c)}}{|\mathbf{x}_d - \mathbf{x}_s|^2} = \mathcal{P}_{0,\omega}(\mathbf{x}_s) \frac{e^{i\omega t + ik|\mathbf{x}_d - \mathbf{x}_s|}}{|\mathbf{x}_d - \mathbf{x}_s|^2},$$

where $k = 2\pi/\lambda$ is the wave number at the modulation wavelength, λ . If instead of a single light source we have a collection of incoherent sources comprising a surface S , we have

$$\mathcal{P}_\omega(\mathbf{x}_d, t) = \int_S \mathcal{P}_{0,\omega}(\mathbf{x}_s) \frac{e^{i\omega t + ik|\mathbf{x}_s - \mathbf{x}_d|}}{|\mathbf{x}_s - \mathbf{x}_d|^2} d\mathbf{x}_s. \quad (5.9)$$

This equation looks like the Rayleigh-Sommerfeld propagator, except for the squared denominator, and the missing $1/i\lambda$. The $1/i\lambda$ term is a global constant that does not qualitatively affect our propagator. We approximate $|\mathbf{x}_s - \mathbf{x}_d|^2 \approx |\mathbf{x}_s - \mathbf{x}_d| \langle S \rangle - \mathbf{x}_d|$ where $\langle S \rangle$ is the average position of all source points in S . Pulling this constant term

out of the integral, we obtain

$$\begin{aligned}\mathcal{P}_\omega(\mathbf{x}_d, t) &\approx \frac{1}{|\langle S \rangle - \mathbf{x}_d|} \int_S \mathcal{P}_{0,\omega}(\mathbf{x}_s) \frac{e^{i\omega t + ik|\mathbf{x}_s - \mathbf{x}_d|}}{|\mathbf{x}_s - \mathbf{x}_d|} d\mathbf{x}_s \\ &= \frac{1}{|\langle S \rangle - \mathbf{x}_d|} \int_S \mathcal{P}_\omega(\mathbf{x}_s, t) \frac{e^{ik|\mathbf{x}_s - \mathbf{x}_d|}}{|\mathbf{x}_s - \mathbf{x}_d|} d\mathbf{x}_s,\end{aligned}\quad (5.10)$$

which is the RSD in Equation (5.4) for scalar waves, with $\gamma = 1/|\langle S \rangle - \mathbf{x}_d|$. This approximation does not affect the phase term, causing only a slow-varying error in amplitude. Given a known source plane, this error can be precomputed. Since it does not alter the phase of the signal, it has no effect on the reconstructed geometry.

Furthermore, as we show in Section 5.B.1, most real imaging systems do not invert the $1/r$ term in the RSD propagator. Further research may also lead to alternative formulations of the phasor field that deal with this error in a more elegant way.

In the following, we use $\mathcal{R}_{\mathbf{x}_d}(\mathcal{P}(\mathbf{x}_s, t))$ as a shorthand for the RSD operator:

$$\mathcal{R}_{\mathbf{x}_d}(\mathcal{P}(\mathbf{x}_s, t)) = \frac{1}{|\langle S \rangle - \mathbf{x}_d|} \int_S \mathcal{P}_\omega(\mathbf{x}_s, t) \frac{e^{ik|\mathbf{x}_s - \mathbf{x}_d|}}{|\mathbf{x}_s - \mathbf{x}_d|} d\mathbf{x}_s. \quad (5.11)$$

5.A.1 Propagating broadband signals

The derived RSD operator propagates monochromatic waves $\mathcal{P}_\omega(\mathbf{x}_s, t) = \mathcal{P}_{0,\omega}(\mathbf{x}_s) e^{i\omega t}$. Any broadband signal can be propagated by first writing it as a superposition of monochromatic waves, then propagating each one individually. For a general $\mathcal{P}(\mathbf{x}, t)$ we therefore define the RSD operator as:

$$\mathcal{R}_{\mathbf{x}_d}(\mathcal{P}(\mathbf{x}_s, t)) = \mathcal{R}_{\mathbf{x}_d} \left(\int_{-\infty}^{+\infty} \mathcal{P}_\omega(\mathbf{x}_s, t) \frac{d\omega}{2\pi} \right) = \int_{-\infty}^{+\infty} \mathcal{R}_{\mathbf{x}_d}(\mathcal{P}_\omega(\mathbf{x}_s, t)) \frac{d\omega}{2\pi}. \quad (5.12)$$

Alternatively, a broadband RSD operator can be implemented in the time domain by shifting the components of \mathcal{P} in time as follows:

$$\begin{aligned}\mathcal{R}_{\mathbf{x}_d}(\mathcal{P}(\mathbf{x}_s, t)) &= \int_{-\infty}^{+\infty} \mathcal{R}_{\mathbf{x}_d}(\mathcal{P}_\omega(\mathbf{x}_s, t)) \frac{d\omega}{2\pi} \\ &= \int_{-\infty}^{+\infty} \int_S \mathcal{P}_{0,\omega}(\mathbf{x}_s) e^{i\omega t - ik|\mathbf{x}_s - \mathbf{x}_d|} d\mathbf{x}_s \frac{d\omega}{2\pi} \\ &= \int_S \int_{-\infty}^{+\infty} \mathcal{P}_{0,\omega}(\mathbf{x}_s) e^{i\omega(t - \frac{1}{c}|\mathbf{x}_s - \mathbf{x}_d|)} \frac{d\omega}{2\pi} d\mathbf{x}_s \\ &= \int_S \mathcal{P}(\mathbf{x}_s, t - \frac{1}{c}|\mathbf{x}_s - \mathbf{x}_d|) d\mathbf{x}_s\end{aligned}\quad (5.13)$$

We will make use of this property when deriving the camera functions for our example cameras (see Table 5.2).

5.A.2 Non-Lambertian surfaces

The RSD propagator we derive requires that the source plane S be Lambertian. Since all our cameras rely primarily on RSD propagators from the aperture of the relay




Figure 5.5: Light hitting the relay wall (in yellow) illuminates the target scene, and is scattered and bounced back by unknown objects in the scene (in blue). **a)** When the target object is perfectly Lambertian, it bounces light back to the entire captured surface of the relay (our virtual lens); our virtual imaging system then focuses the incoming irradiance. **b)** Increasingly specular surfaces may cause the returning light to be reflected towards specific directions; however, like a conventional camera, as long as such reflected light hits some area of the virtual lens, it will be imaged correctly, with a potential spatial resolution loss if few light sources are used. Increasing the number of light sources allows to obtain a progressively more complete irradiance reconstruction and improve resolution.

wall, the Lambertian constraint only applies to the relay wall. Rather than reconstructing the geometry and BRDF of the scene, our virtual cameras reconstruct phasor field irradiance from the scene towards our virtual aperture as a function of position and time, analogous to their LOS counterparts. The reconstructed signal thus corresponds to the averaged irradiance for the entire aperture. This is illustrated in Figure 5.5. Prior methods seek to reconstruct NLOS geometry, which requires correct modelling of albedo, BRDF, occlusions, and interreflections, resulting in a nonlinear inverse problem [91]. In the absence of such data from the hidden scene, these prior methods need to rely on simplifying assumptions, thus limiting the range of scenes that can be reconstructed. Since our method does not make any assumption about the surface properties of the hidden scene, the changes in material appearance do not significantly affect our irradiance reconstructions (see Results in Method section). Simulations with varying BRDFs can be found in the Method Section.

5.B Line-of-sight template functions

In this section we show how to model the outgoing illumination wavefront $\mathcal{P}(\mathbf{x}_p, t)$, as well as the image formation model $\Phi(\cdot)$ for our example *template* LOS imaging systems, using standard diffraction optics. We begin with some preliminary considerations regarding the phase transformation by an ideal lens, which is essential for deriving any arbitrary image formation model $\Phi(\cdot)$. We then derive $\mathcal{P}(\mathbf{x}_p, t)$ and $\Phi(\cdot)$ for three example systems: (1) Conventional photography camera, (2) Transient camera, and (3) Confocal time-gated camera.

5.B.1 Phase operator of an ideal computational thin lens

We define an ideal lens as an element that focuses a planar wavefront into a point at the focal distance f from the lens, on the optic axis. This is equivalent to converting light coming from a point \mathbf{x}_f and turning it into a planar wave, i.e., a wave with a phase that is independent of the position \mathbf{x}_l on the plane of the lens. Light leaving from a point at \mathbf{x}_f creates spherical wavefronts, i.e., the phase at a plane perpendicular to the z -direction at a distance f from \mathbf{x}_f is

System	$\mathcal{P}(\mathbf{x}_p, t)$
(1) Photo camera (ambient light)	$e^{i\omega t}$
(2) Transient camera (pulsed point light)	$e^{i\omega t} \delta(\mathbf{x}_p - \mathbf{x}_{l_s}) e^{-\frac{(t-t_0)^2}{2\sigma^2}}$
(3) Confocal time-gated camera (pulsed focused light)	$e^{i\omega(t - \frac{1}{c} \mathbf{x}_v - \mathbf{x}_p)} e^{-\frac{(t-t_0 - \frac{1}{c} \mathbf{x}_v - \mathbf{x}_p)^2}{2\sigma^2}}$

Table 5.1: Illumination wave functions for different light sources, used in our three example imaging systems.

$$\phi_\omega(\mathbf{x}_l, \mathbf{x}_f) = e^{i\omega \frac{|\mathbf{x}_f - \mathbf{x}_l|}{c}}. \quad (5.14)$$

The lens phase shift, $\phi_l(\mathbf{x}_l, \mathbf{x}_f)$, has to cancel this phase term, and thus the lens acts on the wavefront of a monochromatic wave $\mathcal{P}_\omega(\mathbf{x})$ as

$$\mathcal{P}'_\omega(\mathbf{x}_l, t) = \mathcal{P}_\omega(\mathbf{x}_l, t) \cdot \phi_l(\mathbf{x}_f, \mathbf{x}_l). \quad (5.15)$$

where $\mathcal{P}'_\omega(\mathbf{x}_l, t)$ is the wavefront *after* the lens and $\phi_l(\mathbf{x}_l, \mathbf{x}_f) = e^{-i\omega |\mathbf{x}_f - \mathbf{x}_l|/c}$.

To understand how this lens affects a general broadband signal, consider a wave $\mathcal{P}(\mathbf{x}, t)$ expressed as a superposition of monochromatic waves:

$$\mathcal{P}(\mathbf{x}, t) = \int_{-\infty}^{+\infty} \mathcal{P}_\omega(\mathbf{x}, t) \frac{d\omega}{2\pi} = \int_{-\infty}^{+\infty} \mathcal{P}_{0,\omega}(\mathbf{x}) e^{i\omega t} \frac{d\omega}{2\pi}; \quad (5.16)$$

Considering \mathbf{x}_f and $\mathbf{x}_l \in L$ and applying the phase shift of the lens to this wavefront, we find

$$\mathcal{P}(\mathbf{x}, t) \phi_l(\mathbf{x}_l, \mathbf{x}_f) = \int_{-\infty}^{+\infty} \mathcal{P}_{0,\omega}(\mathbf{x}) e^{i(\omega t - \omega \frac{|\mathbf{x}_f - \mathbf{x}_l|}{c})} \frac{d\omega}{2\pi} = \mathcal{P}(\mathbf{x}, t - \frac{|\mathbf{x}_f - \mathbf{x}_l|}{c}). \quad (5.17)$$

Like the RSD propagation, the phase shift of an ideal lens can thus also be described as a shift in time.

Imaging with a lens. A lens that images a point \mathbf{x}_v onto a sensor pixel \mathbf{x}_r can be described as a combination of two co-located lenses. One to collimate the light from \mathbf{x}_v , and one to focus it onto \mathbf{x}_r . This results in a phase shift $\phi(\mathbf{x}_l) = -\omega \frac{|\mathbf{x}_v - \mathbf{x}_l| - |\mathbf{x}_l - \mathbf{x}_r|}{c}$. In the camera, this is followed by a propagation from the lens to the sensor. If we use Equations 5.13 and 5.17 for the propagation and lens, we obtain:

$$\begin{aligned} \mathcal{P}(\mathbf{x}_r, t) &= \int_L \mathcal{P}(\mathbf{x}_l, t - |\mathbf{x}_v - \mathbf{x}_l| + |\mathbf{x}_l - \mathbf{x}_r| - |\mathbf{x}_l - \mathbf{x}_r|) d\mathbf{x}_l \\ &= \int_L \mathcal{P}(\mathbf{x}_l, t - |\mathbf{x}_v - \mathbf{x}_l|) d\mathbf{x}_l \\ &= \mathcal{R}_{\mathbf{x}_v}(\mathcal{P}(\mathbf{x}_l, t)), \end{aligned} \quad (5.18)$$

where the $1/r$ factor in the RSD propagator can be omitted. The imaging lens thus effectively propagates the field from the aperture \mathbf{x}_l back into the scene.

System	$\Phi(\mathcal{P}(\mathbf{x}_c, t))$
(1) Photo camera	$ \int_{\mathcal{C}} \mathcal{P}(\mathbf{x}_c, t - \frac{1}{c} \mathbf{x}_v - \mathbf{x}_c)d\mathbf{x}_c ^2 = \mathcal{R}_{\mathbf{x}_v}(\mathcal{P}(\mathbf{x}_c, t)) ^2$
(2) Transient camera	$ \int_{\mathcal{C}} \mathcal{P}(\mathbf{x}_c, t - \frac{1}{c} \mathbf{x}_v - \mathbf{x}_c)d\mathbf{x}_c ^2 = \mathcal{R}_{\mathbf{x}_v}(\mathcal{P}(\mathbf{x}_c, t)) ^2$
(3) Confocal time-gated camera	$ \int_{\mathcal{C}} \mathcal{P}(\mathbf{x}_c, -\frac{1}{c} \mathbf{x}_v - \mathbf{x}_c)d\mathbf{x}_c ^2 = \mathcal{R}_{\mathbf{x}_v}(\mathcal{P}(\mathbf{x}_c, -\frac{1}{c} \mathbf{x}_v - \mathbf{x}_c)) ^2$

Table 5.2: Imaging operators to implement our three example imaging systems. The evaluation functions essentially describe the imaging transform of a lens with the resulting image being read out at different times with respect to the illumination.

5.B.2 Projector and camera functions

Our theoretical model allows us to implement any arbitrary (virtual) camera system by defining the projector function $\mathcal{P}(\mathbf{x}_p, t)$ and imaging operator $\Phi(\cdot)$. Methods for modeling such function using Fourier optics are available in the literature [70]. In our work we implement three of them: (1) conventional photography camera, (2) Transient camera, and (3) Confocal time-gated camera. Each has capabilities never before demonstrated in NLOS imaging. The derived $\mathcal{P}(\mathbf{x}_p, t)$ and $\Phi(\cdot)$ functions are listed in Tables 5.1 and 5.2.

- (1) Our first example is a **conventional photography camera** system with a \mathcal{P} -field monochromatic illumination source at frequency ω . It reconstructs the hidden scene with low computational effort. Like a LOS photography camera it does not require knowledge of the position or timing of the light source.

In other words, the reconstruction is independent of the position of \mathbf{x}_p , which reduces the need of careful calibration of the laser positions in the relay wall. Since our illumination for this system is ambient light, the projector function $\mathcal{P}(\mathbf{x}_p, t)$ can be anything. However, like in conventional imaging, the resolution of the image is determined by the temporal bandwidth of $\mathcal{P}(\mathbf{x}_p, t)$, corresponding to the wavelength in the conventional camera. We thus choose a function with a short phasor field wavelength: $\mathcal{P}(\mathbf{x}_p, t) = e^{i\omega t}$. The camera operator is represented by a lens that creates an image on a set of detector pixels that record the absolute value squared of the field. Implementing the lens using the time shift property or the RSD propagator from Equation (5.18) yields $\Phi(\mathcal{P}(\mathbf{x}_c, t)) = |\int_{\mathcal{C}} \mathcal{P}(\mathbf{x}_c, t - \frac{1}{c}|\mathbf{x}_v - \mathbf{x}_c|)d\mathbf{x}_c|^2 = |\mathcal{R}_{\mathbf{x}_v}(\mathcal{P}(\mathbf{x}_c, t))|^2$. Note that this expression is constant with time, just like the intensity in the sensor of a LOS photography camera, so it can be evaluated at any time t .

- (2) The second example is a **NLOS transient camera** system. Like its LOS counterpart [225], this camera captures the propagation of light through the scene, revealing complex multibounce light transport phenomena. As a consequence this virtual camera may be used to identify direct and global components of such light transport. In this case we model a monochromatic point light source at a single point \mathbf{x}_{I_s} which illuminates the scene with a short gaussian shape flash of $\sigma = \frac{6\lambda}{2.36}$ at time t_0 . The illumination function is thus $\mathcal{P}(\mathbf{x}_p, t) = e^{i\omega t} \delta(\mathbf{x}_p - \mathbf{x}_{I_s}) e^{-\frac{(t-t_0)^2}{2\sigma^2}}$. The camera is the same as the conventional photography camera, except that the reconstructed intensity on the sensor now depends on time t , capturing frames at each t_f . We assume that the camera focus follows the light pulse.

- (3) Last, we implement a **confocal time-gated** imaging system, which images specific voxels \mathbf{x}_v of a volumetric space, illuminated with a focused ultrashort pulse of width σ . Note that our virtual imaging system is confocal, but the data for H is not necessarily captured with a confocal arrangement as in prior NLOS work [164].

In this case the illumination is a light pulse focused on a voxel \mathbf{x}_v , defined as $\mathcal{P}(\mathbf{x}_p, t) = e^{i\omega(t - \frac{1}{c}|\mathbf{x}_v - \mathbf{x}_p|)} e^{-\frac{(t - t_0 - \frac{1}{c}|\mathbf{x}_v - \mathbf{x}_p|)^2}{2\sigma^2}}$. In the design of this system we can choose the phasor field pulse width σ . As this width increases, the depth resolution of the virtual imaging system worsens, although the signal-to-noise ratio improves. In practice, we found that a pulse full width at half maximum of about six wavelengths $\sigma = \frac{6\lambda}{2.36}$ yields the best results. Longer pulses are effective for canceling more noise in the reconstruction. The camera is again implemented as an imaging lens, like the cases above. However, in this case the camera needs to focus on the same point \mathbf{x}_v as the light source. Since we are only concerned with the 3rd bounce return directly reflected by a scene surface at \mathbf{x}_v , we evaluate the signal only at a time $t = -\frac{1}{c}|\mathbf{x}_v - \mathbf{x}_c|$, when 3rd-bounce light from this location is seen. This results in a function $\Phi(\mathcal{P}(\mathbf{x}, t)) = |\int_C \mathcal{P}(\mathbf{x}_c, -\frac{1}{c}|\mathbf{x}_v - \mathbf{x}_c|) d\mathbf{x}_c|^2 = |\mathcal{R}_{\mathbf{x}_v}(\mathcal{P}(\mathbf{x}_c, -\frac{1}{c}|\mathbf{x}_v - \mathbf{x}_c|))|^2$.

5.C Implementation of RSD solvers

Here we describe an RSD diffraction integral solver to implement our conventional photographic camera system, and a backprojection solver for the transient confocal systems (refer to Section 5.B.2). Note that both solvers can be applied to any of the systems.

5.C.1 Conventional photography camera using RSD

Using Equation 5.19 we write

$$\begin{aligned} I(\mathbf{x}_v) &= \Phi \left(\int_P [\mathcal{P}(\mathbf{x}_p, t) \star H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)] d\mathbf{x}_p \right) \\ &= \Phi \left(\int_P \int_{-\infty}^{+\infty} \mathcal{P}(\mathbf{x}_p, t - \tau) H(\mathbf{x}_p \rightarrow \mathbf{x}_c, \tau) d\tau d\mathbf{x}_p \right). \end{aligned} \quad (5.19)$$

Plugging in the terms from Tables 5.1, and 5.2 for the conventional photographic camera, we find:

$$I(\mathbf{x}_v) = \left| \mathcal{R}_{\mathbf{x}_v} \left(\int_P [e^{i\omega t} \star H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)] d\mathbf{x}_p \right) \right|^2. \quad (5.20)$$

After this convolution, each time response can be represented entirely by a single complex number. The result is the phasor field (complex amplitude) at the virtual aperture, which is propagated back into the scene using an RSD propagator.

Solving the RSD propagator numerically for each voxel in the scene would be computationally expensive. For a voxel space of side-length N , and $N * N$ points \mathbf{x}_c , the complexity is N^5 . However, there are multiple algorithms that solve the RSD integral for a plane of voxels as a 2D convolution. For all planes making up the reconstruction space this results in a much lower complexity of $N^3 \log(N)$. While

there are efficient solvers for the exact RSD [194], we rely on the well-known Fresnel approximation [70], to implement an efficient solver.

The Fresnel diffraction from a source plane S to a parallel destination plane D at distance z can be approximated as

$$\mathcal{P}(u_d, v_d, z) \approx \gamma \frac{e^{ikz}}{z} \iint_S \mathcal{P}(u_s, v_s) e^{ik \frac{(u_d - u_s)^2 + (v_d - v_s)^2}{2z}} du_s dv_s, \quad (5.21)$$

where u and v are plane coordinates, and subscripts s and d refer to the coordinates in the source and destination planes, so that $\mathbf{x}_s = [u_s, v_s, 0] \in S$ and $\mathbf{x}_d = [u_d, v_d, z] \in D$. This can be interpreted as a 2D spatial convolution with a kernel $K(u, v) = \gamma \frac{e^{ikz}}{z} e^{ik \frac{u^2 + v^2}{2z}}$.

This approximation can be used for the RSD propagator in all our camera operators. The criteria for the validity of the Fresnel approximation is well known [70] and given by

$$\frac{d^4}{4L^3 \lambda} \ll 1, \quad (5.22)$$

where d is the effective aperture radius of the virtual camera, L is the distance between the aperture and the focal plane, and λ is the wavelength.

5.C.2 Confocal time-gated system using backprojection

Plugging in the corresponding terms from Tables 5.1, and 5.2 for the confocal time-gated imaging system in Equation 5.19 we obtain

$$I(\mathbf{x}_v) = \Phi \left(\int_P \int_{-\infty}^{+\infty} e^{i\omega(t - \tau - \frac{1}{c} |\mathbf{x}_v - \mathbf{x}_p|)} e^{-\frac{(t - \tau - t_0 - \frac{1}{c} |\mathbf{x}_v - \mathbf{x}_p|)^2}{2\sigma^2}} H(\mathbf{x}_p \rightarrow \mathbf{x}_c, \tau) d\tau d\mathbf{x}_p \right) = \left| \int_C \int_P \int_{-\infty}^{+\infty} e^{i\omega(-\frac{1}{c} |\mathbf{x}_v - \mathbf{x}_c| - \tau - \frac{1}{c} |\mathbf{x}_v - \mathbf{x}_p|)} e^{-\frac{(-\frac{1}{c} |\mathbf{x}_v - \mathbf{x}_c| - \tau - t_0 - \frac{1}{c} |\mathbf{x}_v - \mathbf{x}_p|)^2}{2\sigma^2}} H(\mathbf{x}_p \rightarrow \mathbf{x}_c, \tau) d\tau d\mathbf{x}_p d\mathbf{x}_c \right|^2.$$

There are multiple ways of solving this expression. We can simply numerically compute the integrals, or we can re-write the expression to include backprojection or diffraction operators. This is desirable since fast methods to execute these operators exist.

Let us first write the expression as a backprojection. We introduce a shifted time $t_s = -\frac{1}{c} |\mathbf{x}_v - \mathbf{x}_c| - \frac{1}{c} |\mathbf{x}_v - \mathbf{x}_p|$ to obtain

$$\Phi(\mathcal{P}(\mathbf{x}_c, t)) = \left| \int_C \int_P \int_{-\infty}^{+\infty} e^{i\omega(t_s - \tau)} e^{-\frac{(t_s - \tau - t_0)^2}{2\sigma^2}} H(\mathbf{x}_p \rightarrow \mathbf{x}_c, \tau) d\tau d\mathbf{x}_p d\mathbf{x}_c \right|^2. \quad (5.23)$$

We break up this expression into two steps. First we perform a convolution on all the collected time responses in H to obtain an intermediate result H' by

$$H'(\mathbf{x}_p \rightarrow \mathbf{x}_c, t) = (e^{i\omega t} e^{-\frac{(t - t_0)^2}{2\sigma^2}}) \star H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t'), \quad (5.24)$$

followed by shifting and summing the results:

$$\begin{aligned}
 I(\mathbf{x}_v) &= \left| \int_C \int_P H'(\mathbf{x}_p \rightarrow \mathbf{x}_c, -\frac{1}{c}|\mathbf{x}_v - \mathbf{x}_c| - \frac{1}{c}|\mathbf{x}_v - \mathbf{x}_p|) d\mathbf{x}_p d\mathbf{x}_c \right|^2 \\
 &\approx \left| \sum_{\mathbf{x}_c \in C} \sum_{\mathbf{x}_p \in P} H'(\mathbf{x}_p \rightarrow \mathbf{x}_c, -\frac{1}{c}|\mathbf{x}_v - \mathbf{x}_c| - \frac{1}{c}|\mathbf{x}_v - \mathbf{x}_p|) \right|^2. \quad (5.25)
 \end{aligned}$$

where the second term is a backprojection, for which efficient implementations exist [7].

5.C.3 Transient camera using backprojection

Last, we derive our NLOS transient system. Operating similarly to the confocal time-gated system, by plugging in the corresponding terms from Tables 5.1, and 5.2 in Equation 5.19 we obtain:

$$H'(\mathbf{x}_p \rightarrow \mathbf{x}_c, t) = (e^{i\omega t} \delta(\mathbf{x}_p - \mathbf{x}_{ls}) e^{-\frac{(t-t_0)^2}{2\sigma^2}}) \star H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t') \quad (5.26)$$

and

$$I(\mathbf{x}_v, t) = \left| \int_C H'(\mathbf{x}_p \rightarrow \mathbf{x}_c, t - \frac{1}{c}|\mathbf{x}_v - \mathbf{x}_c| - \frac{1}{c}|\mathbf{x}_v - \mathbf{x}_{ls}|) d\mathbf{x}_c \right|^2. \quad (5.27)$$

Besides the use of only one illumination point $\mathbf{x}_p = \mathbf{x}_{ls}$, this reconstruction differs from the confocal system in that it depends on time t . The reconstruction is 4-dimensional, resulting in a video of the light propagation in the 3D reconstruction space. To reduce computational cost, we can optionally locate empty voxels by first using our confocal imaging functions.

5.D Details on data acquisition

5.D.1 Hardware configuration

Our capture system, shown in Figure 5.6, consists of a Onefive Katana HP amplified diode laser (1 W at 532 nm, and a pulse width of about 35 ps used at a repetition rate of 10 MHz), and a gated SPAD detector processed by a Time-Correlated Single Photon Counter (TCSPC, PicoQuandt HydraHarp), with a time resolution of about 30 ps and a dead time of 100 ns. Two additional CCD cameras are used to calibrate the laser's position. The measured time resolution of our system is approximately 65 ps, a combination of the pulse width of the laser and the time jitter of the system.

5.D.2 NLOS measurement geometry

We obtain an impulse response function $H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)$ of the scene by sequentially illuminating points \mathbf{x}_p on the relay wall with a short pulse, and detecting the signal returning at points \mathbf{x}_c .

Our hardware device is located 2.5 m from the relay wall, with the NLOS scenes hidden from direct view. The FOV is 25 degrees. The walls are made of standard white styrofoam. The scanning area in the relay wall (virtual camera) is 1.8 m x 1.3 m, with laser points \mathbf{x}_p spaced by $\Delta_p = 1$ cm in each direction. The SPAD detector is focused at a position near the center of the grid. We avoid scanning a small square




Figure 5.6: Capture hardware used for the results shown in this chapter.

region around the SPAD focused position (the “confocal” position) since the signal becomes very noisy at this location. Figures 5.2 and 5.3 provide additional details for the specific scenes shown.

5.D.3 Exposure time

Our capture setup, shown in Figure 5.6, includes CCD cameras to confirm the 3D position of every laser during the measurement; these are a limiting factor in the speed of our experiments. Since the capture process runs in parallel, we use a very long 1 s exposure time per laser position for our main datasets. They are used for all results unless otherwise specified. In addition, we capture scenes without the additional CCD photographs that can be collected much faster and with much shorter exposure times. In Figure 5.2 we show datasets of an office scene captured with exposure times of 1 to 10 ms per laser position, which results into a total capture time as low as 24 s. Further reconstructions of a shelf dataset are shown later as additional results showing that we can reduce exposure times at least down to 50 ms per data point without a significant loss in quality even with ambient light. This results in less than 20 minutes of total capture time. Note also that in our current prototype, we capture data sequentially with a single SPAD. Prototype SPAD arrays are currently under development, and it seems likely that a 16x16 array will be available by the end of the year. We thus expect to be able to capture 256 data points in less than 0.1 s in the near future.

5.D.4 Collected data

In total, we use 10 experimental and two simulated datasets in the project. All experimental datasets use a single SPAD location and 130 by 180 laser positions. The datasets and exposure times are:

	Total Photons	Photons/bin	Max. bin	Avg./ laser pos.
Large depth scene	3215722952	9.7	552	13742
NLOS letters	6502986696	19.6	2889	27791
Shelf	6158590767	18.6	2074	26319
Office Scene	6201680972	18.7	1406	26503
Office Scene 10 ms	48017499	0.14	18	2716
Office Scene 5 ms	24012257	0.072	15	1026
Office Scene 1 ms	4801568	0.014	6	205

Table 5.3: Photon statistics for the captured data: used in the project. The first four scenes were captured with 1 s exposure time. The first column shows the total photons counted, the second the average photon count per time bin, the third is the maximum count over all time bins and the last contains the average number of photons collected in each laser position in the dataset.

- An Office Scene collected with 1 second exposure per laser position. This dataset is used to create the video shown in the supplementary video, frames of which are shown in Figure 5.4(bottom). A photograph and reconstruction of this scene is also shown in the supplementary video. The data is analyzed in Figure 5.7 and Table 5.3.
- An Office Scene collected with 10 ms, 5 ms, and 1 ms, used in Figures 5.2, 5.10, 5.12, 5.11, and Table 5.3.
- A scene of a bookshelf used in Figure 5.3(a) and Table 5.3.
- A scene of a bookshelf captured with various exposure times and ambient light conditions shown in Figures 5.7 and 5.9.
- A scene with letters distributed over a large depth used in Figure 5.3(b) and Table 5.3.
- A scene of the letters NLOS in a plane used in Figure 5.4(top) and Table 5.3.

To provide further insight into the noise and artifacts present in our data we go through an analysis of the raw data from our 1 second exposure office scene. We compare the maximum and average number of photons per second and laser position \vec{x}_p for our captured scenes in Table 5.3. The dark count rate of our detector is 10 photons per second. We do not explicitly subtract dark counts nor ambient light or backgrounds.

Note that the high total photon numbers in the transient responses (Table 5.3) are due to the long responses associated with the large depth and volume of the scenes, and not due to a particularly bright signal. Example data for a scene of a shelf is shown in Figure 5.7 (whose reconstruction can be found on Figure 5.9). In this scene, our longest 1 s exposure time peaks at about 150 photons/s (such peaks are probably due to the presence of specular surfaces), and the captured signal is extremely noisy. In comparison, the recent method by O’Toole et al. [164] acquires a brighter, cleaner signal in 0.1 s, peaking at about 600 photons/s, due to the use of retroreflective paint applied on the hidden objects (data from their data_resolution_chart_40cm dataset).

Let us further analyze the captured data. In Figure 5.8 we show a visualization of our data matrix for the 1s exposure office scene using the function `imagesc` in Matlab, where each row is the data collected for a different location of the laser illumination

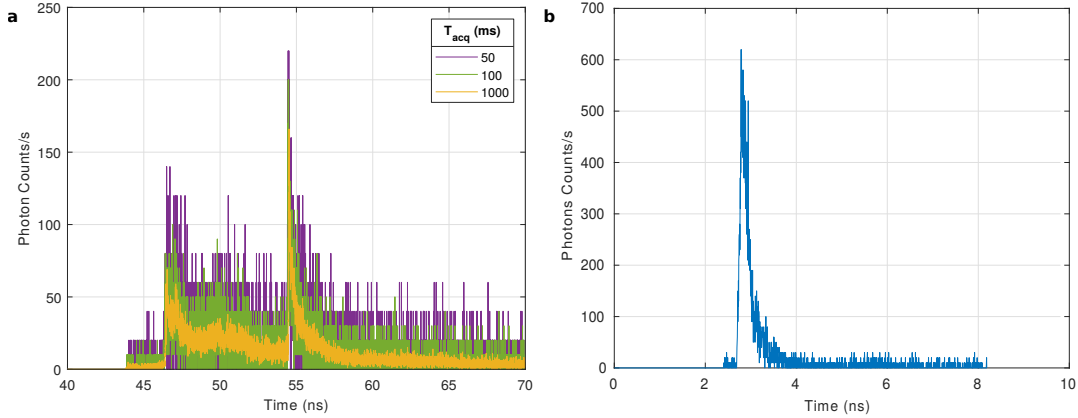


Figure 5.7: Data Comparison: **a**, Raw data for one of the laser positions x_p . Shown is the number of photons per second accumulated in each time bin (i.e. the collected histogram divided by the integration time in seconds). Time bins are 4 ps wide. As expected, all three curves appear to follow the same mean, but have a larger variance for lower exposure times. The raw data thus gets significantly noisier as exposure time decreases. The effects on the reconstruction, however, are rather minor as Figure 5.13 shows. **b**, Example dataset from O'Toole et al. [164] for comparison.

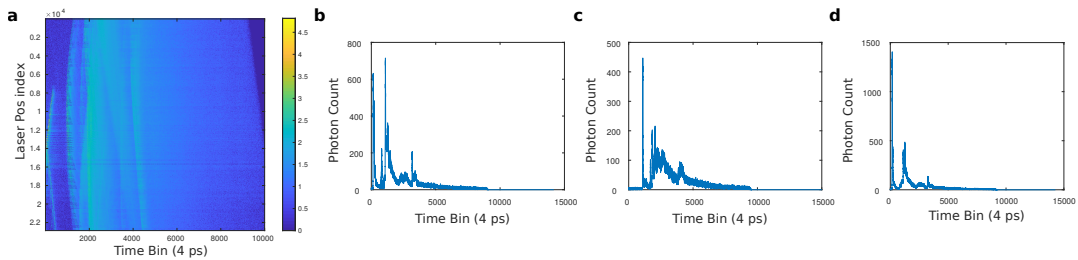


Figure 5.8: Visualization of the raw data for our long exposure office scene. From left to right: **a**, Base 10 logarithm of the photon counts in all time bins. **b**, After removing the first 833 time bins in each dataset the plots show the photon counts for the laser position that received the largest total number of photons in the dataset, **c**, the counts for the laser position that received the median number of photon counts, and **d**, the laser position that contains the time bin with the global maximum count in the entire set.

spot, and each column contains a different time bin. The first time bin corresponds to the time when the illumination laser pulse leaves the relay wall. In the images we do not show time bins 10001 to 15000 as they are mostly empty due to the closing of the gate. As can be seen there are some sparse, very large peaks in the dataset that saturate the counting registers of our TCSPC ($2^{16} - 1$ counts). As we will see, these artifacts in the data are likely due to imperfections in the gating or optical setup.

Let us focus on the first instants of the captured data shown in Figure 5.8(a), which reveal features that look like straight diagonal lines in the first few time bins. The fact that there are straight lines in this plot, indicates that they are likely related to a first-bounce signal, rather than the actual scene response. NLOS signals should show up as hyperbolas or sections of hyperbolas in this type of visualization and the curvature of the hyperbolas should be highest at the earliest time bins. The image contains many more features that look like straight lines that do not appear to have the correct hyperbolic curvature to be actual NLOS signals. Many of them also appear identically again in the other datasets which is another hint that they are probably not real NLOS data, but artifacts related to the measurement system.

Our algorithm is completely agnostic to the presence of these artifacts. The brightest peaks also appear too early in the data to be associated with a NLOS object. To see this, consider that the closest object in any of our scenes is the chair in the office scene and it is more than 50 cm away from the wall. Consequently the first time response from an actual object cannot arrive at the SPAD earlier than 3.3 ns after the laser illuminates the relay wall. Time bins are 4 ps wide. Any data before time bin 833 therefore can only be an artifact. We will speculate more about the origin of these artifacts later.

If we ignore those first 833 time bins that contain no useful data we obtain a dataset that can yield some meaningful statistics about the data. In this dataset the largest photon count in all our over 200 million time bins is smaller than 1400 photons. Note that as we show below, this 1400 maximum is likely still due to a gate artifact that just happened to occur slightly later than 3 ns into the dataset. Statistics for all datasets are shown in Table 5.3.

Maximum photon counts usually come from the objects in the scene closest to the wall. Considering the large depth and specularities of our scenes, most of the reconstructed scene volume is actually using signals much weaker than the maximum signal as voxels are further away from the wall. Signals from a given surface are expected to drop in magnitude with one over distance to the power of four as distance is increased. An object generating 100 photon peaks at 50 cm distance in the front of our scene would therefore only create $100/8$ photons if placed at at 1 m and $100/625 = 0.2$ photons at 2.5 m towards the end of our office scene. This ability to handle scenes with large dynamic range in the data is another stated advantage of our algorithm.

In extended data Figure 5.8(b) we show a plot of the photon counts over time bins for the laser position that received the most total photons. We again see the extreme peak of $2^{16} - 1$ counts in the beginning of the dataset. Again, this peak can not possibly be a real third bounce signal as it would require for the pulse to travel between laser position and SPAD position significantly faster than the speed of light. The actual NLOS data starting around time bin 1000 and peaking at just above 50 photons.

Finally, we show a plot of the laser position that received the total photon count closest to the median of all laser positions in Figure 5.8(c). We can see that this data generally stays below 150 photons with what are likely specular peaks reaching up to 200 photons and a large 450 photon peak at the beginning of the dataset that is either a specular peak or another gate artifact. At this point we want to also point out that since we illuminate only a grid of points at the wall, we do not capture all the specular peaks in our data. In order to see a specular reflection peak from a scene surface we have to get lucky and illuminate the exact spot on the wall that results in the specular reflection to overlap with the the SPAD position (see Figure 5.5 for an illustration). Therefore specular peaks in our measurements can vary greatly based on how close to the peak the laser actually sampled the wall. Again we want to point out that this type of uncontrolled artifact does not affect our algorithm.

As we stated above, the time bin with the highest photon count when ignoring obvious early artifacts contains about 1400 photons. Next we plot the laser position that contains this time bin as shown in Figure 5.8(d). Note that zero on the x-axis here corresponds to time bin 834. As we see the 1400 photon peak appears very close to the beginning of the transient and may be a gating artifact that occurs in the data just

after the opening of the gate. This type of data distortion is described further below. If not a gating artifact, the peak is likely a specular reflection as it is very short and could only be caused by a small isolated diffuse patch or a specular surface in the scene. Peaks from extended diffuse surfaces are necessarily longer in duration.

We thus conclude that while our data contains significant spurious artifacts, the actual photon counts useful for reconstructions are no higher nor cleaner than in previous methods. Note that the removal of early artifacts is only done here to generate Figure 5.8(b-d) to allow visualization. All reconstructions shown in the manuscript contain the full recorded data without the removal of any potential artifacts or time bins.

Even though an understanding of the origin of the artifacts in the data is not needed for our method, we can offer some additional speculation for sources of some of them:

- Many of the early peaks in our data are likely related to imperfections in our gating method. When the SPAD gate opens just after the laser pulse has passed, photoelectrons in the SPAD may cause a detection event that is not due to an actual photon, but rather might be due to the electrons excited by the first bounce light and trapped in long lived states in the SPAD. Even so these electrons are not amplified, they need to be transported off the SPAD junction or they can cause counts as soon as the gate opens.
- In certain cases it is also possible that the gate does not actually block the pulse for some laser positions. The gate has to be positioned such that it blocks the laser in all laser positions, while not blocking any actual signal. This is not always possible, and we do not re-adjust the gate for each position while scanning.
- In our past setups we have observed many effects inside the imaging system that can keep light trapped long enough to cause a peak at the time when the the NLOS data arrives. This can be due to multiple reflections between lenses, multiphoton fluorescence in the glass or coating of the lenses or stray light reflecting off a random surface at the right distance. We have confirmed some of these effects, but suspect there are many more.
- In particular we can see light that travels from the laser spot to the SPAD, reflects off the surface of the SPAD pixel, is imaged back to the relay wall and comes back to the SPAD. In confocal or near confocal configurations this can create a peak that is many times brighter than the actual data.
- Retro-reflective targets can be used to reduce many of these artifacts, most of which are created either by the laser or a first bounce reflection of the laser. If the hidden target is retroreflecting, the ratio between the brightness of the laser and its first bounce and the brightness of the 3rd bounce NLOS data is reduced by multiple orders of magnitude.

5.D.5 Helmholtz reciprocity

Ideally, we would capture $H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)$ sampling points on both the projector aperture $\mathbf{x}_p \in P$ and the camera aperture $\mathbf{x}_c \in C$. In our current setup with a single SPAD detector, we only sample a single point for \mathbf{x}_c . Using Helmholtz Reciprocity we can

interpret these datasets as having a single \mathbf{x}_p and an array of \mathbf{x}_c . The choice of capture arrangement is made for convenience since it is easier to calibrate the position of the laser spot on the wall. Improved results are anticipated once array sensors become available (currently under development).

5.E Additional validation and discussion

5.E.1 Resolution limits

The resolution limit for NLOS imaging systems with an aperture diameter d at imaging distance L is closely related to the Rayleigh diffraction limit [30]: $\Delta_x = 1.22c\sigma L/d$, with c the speed of light, for a pulse of full width at half maximum σ . O'Toole et al. [164] derive a criterion for a resolvable object based on the separability of the signal in the raw data, not in the reconstruction, resulting in a similar formula of $\Delta_x = 0.5c\sigma L/d \approx 0.5\lambda L/d$.

In our virtual LOS imaging system, we can formulate a resolution limit that ensures a minimum contrast in the reconstruction, based on the well-known resolution limits of wave based imaging systems. The resolution limit therefore depends on the particular choice of virtual imaging system. For an imaging system that uses focusing only on the detection or illumination side, this limit is approximated by the Rayleigh criterion. For an imaging system that provides focusing both on the light source and the detector side, the resolution doubles (as it does for example in a confocal or structured illumination microscope) and becomes $\Delta_x = 0.61\lambda L/d$.

5.E.2 Effect of exposure time

Ambient light: To analyze how well our technique works in ambient light and with much faster exposure times, we perform several additional measurements using progressively shorter exposure times, showing that we can reduce exposure times at least down to 50 ms per data point without a significant loss in quality, as can be seen in Figure 5.9. Raw data for one of the laser positions is shown in Figure 5.7. In particular, it shows the number of photons per second accumulated in each time bin (i.e. the collected histogram divided by the integration time in seconds). As expected, all three curves appear to follow the same mean, but have a larger variance for lower exposure times. The raw data thus gets significantly noisier as exposure time decreases. The effects on our reconstruction, however, are rather minor as Figure 5.9 shows.

Short exposure captured data: Figure 5.10 shows the reconstruction of the office scene (same as Figure 5.2) for short exposure times of 10 ms, 5 ms, and 1 ms for each of the roughly 24000 laser positions. This leads to total capture times of about 4 minutes, 2 minutes and 24 seconds respectively. Plots showing raw data from those datasets are shown in Figure 5.12.

We compare the results of our reconstructions on the 1 ms-data, against filtered backprojection with a Laplacian filter [224], as well as the LOG-filtered backprojection [135], which generally achieves better results. In fact, we are not aware of any reconstruction method that has been shown to consistently outperform a LOG-filtered backprojection. Figure 5.11 shows the result of this comparison.

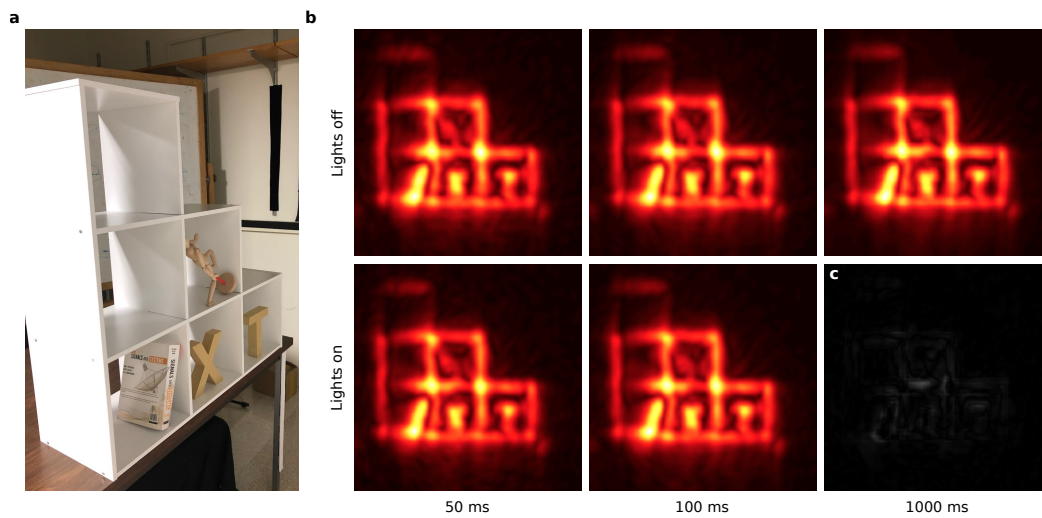


Figure 5.9: Robustness to ambient light and noise: **a**, Hidden bookshelf. **b**, Imaging results with increasingly higher exposure times; even at 50 ms, there is no significant loss in quality. **Top row:** Using only the pulsed laser as illumination source. **Bottom row:** adding a large amount of ambient light (same conditions as the photograph in (a)), the quality also remains constant. **c**, Difference between the 50 ms and 1000 ms exposure captures for the lights on case.




Figure 5.10: Short exposure reconstructions: Reconstruction of the office scene using very short capture times. **a**, Photograph of the captured scene. **b**, From left to right, reconstructions for data captured with 10, 5, and 1 ms exposure time per laser. The total capture time was about 4 minutes, 2 minutes, and 24 seconds, respectively.




Figure 5.11: Comparison to prior methods: Reconstruction of the office scene using very short capture times of 1 ms per laser (24 seconds in total). **a**, Filtered backprojection using the Laplacian filter. **b**, LOG-filtered backprojection. **c**, Our method.

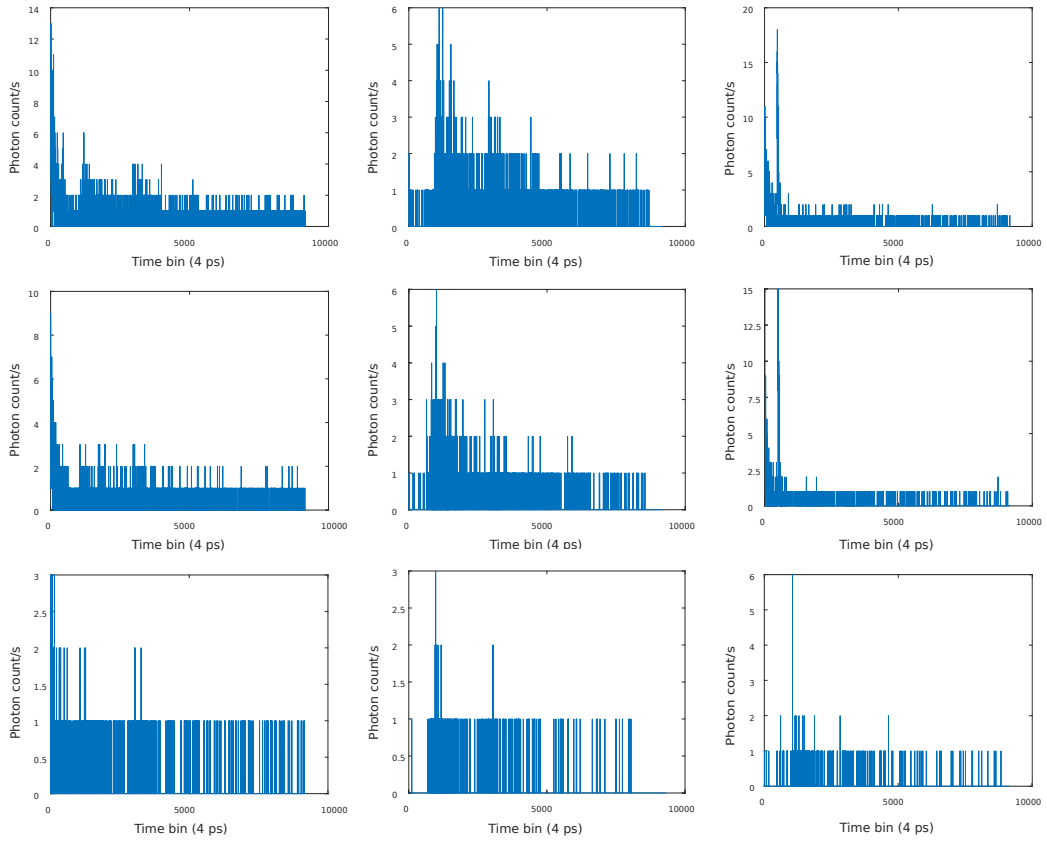


Figure 5.12: Short exposure data: Photon counts in the raw data for our office scene for 10 ms (top row), 5 ms (center row), and 1 ms (bottom row) exposure times per laser. After removing the first 833 time bins in each dataset the columns show (left) the photon counts for the laser position that received the largest total number of photons in the dataset, (center) the counts for the laser position that received the median number of photon counts, and (right) the laser position that contains the time bin with the global maximum count in the entire set.

5.E.3 Effect of strong interreflections

In order to confirm the presence and effect of strong interreflections in our captured data, we compare it qualitatively with primary data from a synthetic bookshelf scene, with and without interreflections. The bookshelf is placed in a corridor of 2 m x 2 m x 3 m, with only a single lateral aperture of 1 m x 2 m to allow imaging the hidden scene. The shelf has a size of 1.4 m x 0.5 m, placed at 1.7 m from the relay wall, and 0.3 m from the lateral walls. The virtual aperture has a size of 1.792 m x 1.7920 m and a granularity of 256 x 256 laser points; we use a $\lambda = 4$ cm, $\Delta_p = 2.8$ cm.

As can be seen in Figure 5.13, the synthetic data clearly shows how the presence of interreflections adds, as expected, low-frequency information resembling echoes of light. This same behavior can be seen in the real captured data, revealing the presence of strong interreflections.

Additionally, we evaluate the robustness of our method in the presence of such interreflections. Similar to recent work [164], we compare between a voxelization of the ground-truth geometry and a reconstructed voxel-grid obtained from our irradiance reconstructions, with and without including interreflections; the resulting mean square error is:




Figure 5.13: Robustness to multiple reflections: Result on the synthetic bookshelf scene. **a**, Without interreflections. **b**, Including high-order interreflections. The quality of the results is very similar. **c**, Primary data (streak images) from the same scene without (top), and with interreflections (middle). This synthetic data clearly shows how the presence of interreflections adds, as expected, low-frequency information resembling echoes of light. The bottom image shows primary data captured from the real office scene in Figure 5.2. It follows the same behavior as the middle image, revealing the presence of strong interreflections.

MSE without interreflections, Figure 5.13(a): 4.93 mm.

MSE with interreflections, Figure 5.13(b): 4.66 mm.

5.E.4 Non-Lambertian surfaces

To validate the robustness of our method in the presence of non-Lambertian materials in the hidden scene, we have created a synthetic scene made up of two letters, R and D, one partially occluding the other, placed in a corridor of 2m x 2m x 3m, with only a single lateral aperture of 1 m x 2 m to allow imaging the hidden scene. The letters have a size of 0.75 m x 0.8 m, placed at 1.25 m and 1.7 m from the relay wall, respectively, and 0.5m from the lateral walls (see extended data Fig.9.a). The virtual aperture has a size of 1.792 m x 1.792 m and a granularity of 128 x 128 laser points; we use a $\lambda = 4\Delta_p$ with $\Delta_p = 5.6$ cm. We start with purely Lambertian targets, and progressively increase their specularity; we use the Ward BRDF model [228], decreasing the surface roughness, using available transient rendering software [106]. The simulation includes up to the fifth indirect bounce.

Figure 5.14(b) shows the resulting irradiance reconstructions. Since our method does not make any assumption about the surface properties of the hidden scene, the changes in material appearance do not significantly affect our irradiance reconstructions. Similar to recent work [164] we compare between a voxelization of the ground-truth geometry, and the reconstructed voxel-grid; the resulting mean square error for each of the different reflectances is as follows:

MSE for a surface roughness of 1 (perfect Lambertian): 2.1 mm.

MSE for a surface roughness of 0.4: 2.2 mm.

MSE for a surface roughness of 0.2: 2.2 mm.

5.E.5 Reconstruction comparison with other methods

Our imaging system allows to reconstruct hidden geometry. For this application, we show here a comparison using the publicly available confocal dataset [164]. This set




Figure 5.14: Robustness to scene reflectance: **a**, Geometry of our experimental setup. **b**, From left to right, imaging results for the Lambertian targets (roughness 1), and increasingly specular surfaces (roughness 0.4 and roughness 0.2). The reconstructed irradiance is essentially the same for all cases.

can be reconstructed using different NLOS methods; we show results for the CNLOS deconvolution [164], filtered backprojection [30], and our proposed method. For these confocal measurements, backprojection can be expressed as a convolution with a pre-calculated kernel, and thus all three methods are using the same backprojection operator. Neither our method nor filtered backprojection are limited to confocal data, and can be acquired making use of simpler devices and capture configurations. They can thus be applied to a broader set of configurations and considerably more complex scenes. For the CNLOS deconvolution method [164], we leave the optimal parameters unchanged. For our proposed virtual wave method, we use the aperture size and its spatial sampling grid published in the supplementary materials to calculate the optimal phasor field wavelength. For the filtered backprojection it is important to choose a good discrete approximation of the Laplacian operator in the presence of noise. Previous works implicitly do the denoising step by adjusting the reconstruction grid size to approximately match the expected reconstruction quality [30; 80; 224], or by downsampling across the measurements [164]. All of them can be considered as proper regularizers. To provide a fair comparison without changing the reconstruction grid size, we convolve a Gaussian denoising kernel with the Laplacian kernel, resulting in a LOG filter, which we apply over the backprojected volume.

Note that a large improvement in reconstruction quality for the simple scenes included in the dataset (isolated objects with no interreflections) is not to be expected, since existing methods already deliver reconstructions approaching their resolution limits. We nevertheless achieve improved contrast and cleaner contours in our wave camera method, due to our better handling of multiply scattered light, which polutes the reconstructions in the other methods, as shown in Figure 5.15.

In noisy datasets, as the last two rows in Figure 5.15, filtered backprojection fails. CNLOS includes a Wiener filter that performs well removing uniform background noise, although a noise level must be *explicitly* estimated. Our phasor field virtual wave method, on the other hand, performs well automatically, without the need to explicitly estimate a noise level. This is particularly important in complex scenes with interreflections, where the background is not uniform across the scene, and the noise level cannot be reliably estimated.

Nevertheless, our main contribution is not improving the reconstruction for simple, 3rd-bounce scenes. Instead, our method allows to derive a new class of NLOS algorithms, which can successfully handle scenes of much larger complexity.




Figure 5.15: Reconstruction comparison on a public dataset. From left to right: CNLOS deconvolution, filtered (LOG) backprojection, and our proposed method. A large improvement in reconstruction quality for the simple scenes included in the dataset (isolated objects with no interreflections) is not to be expected, since existing methods already deliver reconstructions approaching their resolution limits. Nevertheless, our method achieves improved contrast and cleaner contours, due to better handling of multiply scattered light.

Part IV

Conclusions

Chapter 6

Conclusions

Through this work we have explored the possibilities of light transport, and presented several contributions that increase the capabilities of the existing models and their applications on other fields. This thesis is divided in parts, with Part II (Chapters 2 and 3) containing the contributions to steady-state light transport and Part III (Chapters 4 and 5) the ones exploring the transient domain.

Part II - Steady state We have explored two lines of work in this part, one related with improving efficiency and another focused on expanding the existing capabilities. First, in Chapter 2 we address lighting computation, developing new techniques to speed-up calculations when using stochastic methods. For that, we focus on the case of disk-shaped luminaries, and devise new solid-angle based techniques that reduces the overall variance of Monte Carlo computations. This results in a lower number of stochastic samples, and therefore execution time, required to obtain noise-free images. Our method can be incorporated into any Monte Carlo based rendering engine, and has been integrated into the commercial Arnold renderer starting on the 5.0 version [5].

Improving the efficiency of direct lighting computation is an active field of research, since even when computing global illumination a good next-event estimation strategy can reduce the overall variance in most practical scenarios.

In Chapter 3 we tackle the problem of reproducing the appearance of pearlescent pigments, a class of materials whose properties are impossible or too inefficient to replicate using existing models. We develop a new stochastic framework, with which it is possible to match captures of existing materials, and it is possible to represent a wide gamut of plausible appearance. In order to do so, we build over existing techniques for complex, multi-layered materials and incorporate the relevant parts from wave optics to create an accurate representation of the material inner structure. Our virtual model can be directly linked to physical properties of the simulated pigments, which in turn are determined by the manufacturing process, opening the possibility of computational design of pearlescence appearance. Being able to intuitively edit the aspect of these pigments, applied over a virtual representation of the object they are intended to be applied to, inside a computer simulation and then reproduce them accurately in manufacture, would considerable speed-up the design process, which currently relies heavily on trial-and-error with different pigment mixtures until a desirable look is achieved.

Part III - Transient state Similarly to the steady-state domain, in transient state we have worked in both efficient simulation of light transport and its applications. On

the simulation side, in Chapter 4 we introduce a density estimation method which reuses samples in the temporal and spatial domains, reducing the variance of the estimation at the cost of introducing a slight bias. This bias is progressively reduced as the number of samples used in the radiance calculation increases, resulting in a consistent method that converges asymptotically to the correct solution.

In Chapter 5 we introduce a new imaging technique for visualizing scenes that fall outside the line of sight of the observer. This task is performed by shooting short laser pulses towards secondary surfaces, which can reflect light towards the hidden scene, and capturing their time of arrival when the pulses bounce back from the relay surfaces after interacting with the objects in the hidden scene. The captured signal can be posed as an impulse response function, which serves to propagate computationally an arbitrary signal inside the hidden scene. In our method, this impulse function is used to propagate a virtual complex valued field with computationally controlled phase, the *phasor field*, which behaves almost identically to a modulated signal using the laser pulse as the carrier [183]. The phasor field is then focused computationally, using a virtual camera system, yielding a result equivalent to imaging the hidden scene using modulated light [183]. The resulting visualizations have higher or equivalent quality than the previous state-of-the-art methods [134; 164], while being robust in the presence of noise and handling large scenes with complex geometry and arbitrary materials.

While this line work is much more closely related to computational imaging than to light transport simulation, the technique itself builds over a wave optics model to represent the light propagation inside the scene, including the temporal domain. In order to validate the technique, several simulations of virtual scenes with similar characteristics of those captured were produced. The simulations served as a testbed to detect errors in capture and empirically narrow down the limits of the method and the kind of scenes that could be imaged.

This task was required because, despite the good performance of the method on non-line-of-sight (NLOS) imaging, the validity conditions and implicit assumptions of the phasor field model are still not fully understood. As a consequence, a considerable amount of work has been done in trying to have a complete theoretical description of the propagation and imaging process [46; 47; 143; 182]. Of special interest is the limitation that objects in the hidden scene need to fall within the *visible cone* [139] to be imaged, or otherwise are invisible to the computational camera. This limitation is shared by virtually all convolution-based NLOS imaging systems, a category which encompasses virtually all the state of the art [3], and it can be explained by missing frequencies in the captured signal. A better understanding of the underlying process could provide methods to infer those, leading to priors and iterative methods that allow the system to infer the missing frequencies. Also of interest is the effect that non-diffuse materials have in phasor field imaging. While we know that specular materials can be displayed using the framework, they have a considerable effect with the effective angular and spatial resolution of the virtual camera system, sometimes even improving the result [73]. Formalizing this interplay could lead to better imaging overall, and even open the possibility of classifying materials on the hidden scene based on their response to the phasor field. Interconnected with material properties is how speckle in the carrier signal, the laser pulse, after several bounces from rough surfaces could affect the phasor field. In our work we simply dismiss its influence since it doesn't seem to show up in our images, but a more in-depth analysis could lead to potentially exploit its effects [48; 215]

Closely related to our own work, other NLOS techniques based on wave-based propagation have been proposed. Lindell et al. [138] introduced an equivalent formulation for NLOS imaging using a virtual field, the main difference being the use of a fast approximate solver from seismology to bring the captured scene into focus [245]. Similarly, the base concepts from phasor fields can be translated to other kinds of wave-like carries, like modulated acoustic waves [137] and radio signals [89].

Finally, while our technique took the order of tens of minutes to provide high quality images, faster solvers have already been proposed that took mere seconds [140] and some NLOS techniques can produce results at rates near to real-time [138]. In addition to speed-ups the imaging process, the capture side of the system is also experiencing considerable improvements, moving away from the slow a-point-a-time scanning as employed in our project by incorporating SPAD arrays [181], able to capture a dense set of time-of-flight measurements concurrently. Not only that, but the scale of the scenes that can be imaged using SPAD sensors is continually increasing, and it has already been shown that NLOS reconstruction is viable at up to 1.43 km [238] and ToF imaging at over 200 km [136].

Bibliography

- [1] Georges Aad, Brad Abbott, Dale Charles Abbott, A. Abed Abud, K. Abelling, D. K. Abhayasinghe, S. H. Abidi, O. S. AbouZeid, N. L. Abraham, H. Abramowicz, et al. 2019. Observation of light-by-light scattering in ultraperipheral Pb+ Pb collisions with the ATLAS detector. *Physical Review Letters* 123, 5 (2019). <https://doi.org/10.1103/PhysRevLett.123.052001>
- [2] Nils Abramson. 1978. Light-in-flight recording by holography. *Optics Letters* 3, 4 (1978). <https://doi.org/10.1364/OL.3.000121>
- [3] Byeongjoo Ahn, Akshat Dave, Ashok Veeraraghavan, Ioannis Gkioulekas, and Aswin C. Sankaranarayanan. 2019. Convolutional approximations to the general non-line-of-sight imaging operator. In *Proceedings of the 2019 IEEE/CVF International Conference on Computer Vision (ICCV 2019)*. IEEE. <https://doi.org/10.1109/ICCV.2019.00798>
- [4] Marco Ament, Christoph Bergmann, and Daniel Weiskopf. 2014. Refractive Radiative Transfer Equation. *ACM Transactions on Graphics* 33, 2 (2014). <https://doi.org/10.1145/2557605>
- [5] Solid Angle. 2018. Arnold 5 - Rendering evolved. <https://www.arnoldrenderer.com/news/press-release-arnold-5/> Accessed: 2021-02-28.
- [6] Archimedes of Syracuse. -225. *On the Sphere and Cylinder*.
- [7] Victor Arellano, Diego Gutierrez, and Adrian Jarabo. 2017. Fast back-projection for non-line of sight reconstruction. *Optics Express* 25, 10 (2017). <https://doi.org/10.1364/OE.25.011574>
- [8] James Arvo. 1995. Stratified sampling of spherical triangles. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques (SIGGRAPH '95)*. ACM. <https://doi.org/10.1145/218380.218500>
- [9] James Arvo. 2001. Stratified sampling of 2-manifolds. In *State of the Art in Monte Carlo Ray Tracing for Realistic Image Synthesis*. ACM.
- [10] James Arvo and David Kirk. 1990. Particle transport and image synthesis. In *Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '90)*. ACM. <https://doi.org/10.1145/97879.97886>
- [11] Anita I. Bailey and Susan M. Kay. 1965. Measurement of refractive index and dispersion of mica, employing multiple beam interference techniques. *British Journal of Applied Physics* 16, 1 (1965). <https://doi.org/10.1088/0508-3443/16/1/307>

- [12] Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. 2019. A Monte Carlo framework for rendering speckle statistics in scattering media. *ACM Transactions on Graphics* 38, 4 (2019). <https://doi.org/10.1145/3306346.3322950>
- [13] Petr Beckmann and André Spizzichino. 1963. *The Scattering of Electromagnetic Waves from Rough Surfaces*. Artech House.
- [14] Laurent Belcour. 2018. Efficient rendering of layered materials using an atomic decomposition with statistical operators. *ACM Transactions on Graphics* 37, 4 (2018). <https://doi.org/10.1145/3197517.3201289>
- [15] Laurent Belcour and Pascal Barla. 2017. A practical extension to microfacet theory for the modeling of varying iridescence. *ACM Transactions on Graphics* 36, 4 (2017). <https://doi.org/10.1145/3072959.3073620>
- [16] George I. Bell and Samuel Glasstone. 1970. *Nuclear reactor theory*. Technical Report TID-25606. US Atomic Energy Commission, Washington, DC (United States).
- [17] Benedikt Bitterli. 2016. Rendering resources. <https://benedikt-bitterli.me/resources/>
- [18] Benedikt Bitterli. 2016. Virtual femto photography. <https://benedikt-bitterli.me/femto.html>
- [19] Benedikt Bitterli and Wojciech Jarosz. 2017. Beyond Points and Beams: Higher-Dimensional Photon Samples for Volumetric Light Transport. *ACM Transactions on Graphics* 36, 4 (2017). <https://doi.org/10.1145/3072959.3073698>
- [20] Benedikt Bitterli, Jan Novák, and Wojciech Jarosz. 2015. Portal-Masked Environment Map Sampling. *Computer Graphics Forum* 34, 4 (2015). <https://doi.org/10.1111/cgf.12674>
- [21] Benedikt Bitterli, Srinath Ravichandran, Thomas Müller, Magnus Wrenninge, Jan Novák, Steve Marschner, and Wojciech Jarosz. 2018. A radiative transfer framework for non-exponential media. *ACM Transactions on Graphics* 37, 6 (2018). <https://doi.org/10.1145/3272127.3275103>
- [22] James F. Blinn. 1977. Models of light reflection for computer synthesized pictures. In *Proceedings of the 4th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '77)*. ACM. <https://doi.org/10.1145/563858.563893>
- [23] Boost. 2021. Boost C++ Libraries. (2021). <https://www.boost.org/> Accessed: 2021-02-10.
- [24] James Booth. 1844. On the rectification and quadrature of the spherical ellipse. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 25, 163 (1844).
- [25] James Booth. 1852. Researches on the geometrical properties of elliptic integrals. *Philosophical Transactions of the Royal Society of London* 142 (1852).
- [26] Max Born and Emil Wolf. 1999. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light* (7 ed.). Cambridge University Press.

- [27] Brent Burley. 2012. Physically-based shading at Disney. In *Practical physically-based shading in film and game production (ACM SIGGRAPH 2012 Course Notes)*. ACM. <https://doi.org/10.1145/2343483.2343493>
- [28] Jens Busck. 2005. Underwater 3-D optical imaging with a gated viewing laser radar. *Optical Engineering* 44, 11 (2005). <https://doi.org/10.1117/1.2127895>
- [29] Jens Busck and Henning Heiselberg. 2004. Gated viewing and high-accuracy three-dimensional laser radar. *Applied Optics* 43, 24 (2004). <https://doi.org/10.1364/AO.43.004705>
- [30] Mauro Buttafava, Jessica Zeman, Alberto Tosi, Kevin Eliceiri, and Andreas Velten. 2015. Non-line-of-sight imaging using a time-gated single photon avalanche diode. *Optics express* 23, 16 (2015). <https://doi.org/10.1364/OE.23.020997>
- [31] Mike Cammarano and Henrik Wann Jensen. 2002. Time dependent photon mapping. In *Proceedings of the 13th Eurographics Workshop on Rendering (EGWR '02)*. Eurographics Association. <https://doi.org/10.2312/EGWR/EGWR02/135-144>
- [32] Bille C. Carlson. 1994. Numerical computation of real or complex elliptic integrals. *arXiv preprint math/9409227* (1994). <https://arxiv.org/abs/math/9409227>
- [33] Bille C. Carlson. 1995. Numerical computation of real or complex elliptic integrals. *Numerical Algorithms* 10, 1 (1995). <https://doi.org/10.1007/bf02198293>
- [34] Kenneth M. Case, George Placzek, and Frederic Hoffmann. 1953. *Introduction to the theory of neutron diffusion*. Vol. 1. Los Alamos Scientific Laboratory.
- [35] Subrahmanyam Chandrasekhar. 2013. *Radiative transfer*. Dover Publications.
- [36] Per Christensen, Julian Fong, Jonathan Shade, Wayne Wooten, Brenden Schubert, Andrew Kensler, Stephen Friedman, Charlie Kilpatrick, Cliff Ramshaw, Marc Bannister, Brenton Rayner, Jonathan Brouillat, and Max Liani. 2018. RenderMan: An advanced path-tracing architecture for movie rendering. *ACM Transactions on Graphics* 37, 3 (2018). <https://doi.org/10.1145/3182162>
- [37] John T. Conway. 2010. Analytical solution for the solid angle subtended at any point by an ellipse via a point source radiation vector potential. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 614, 1 (2010). <https://doi.org/10.1016/j.nima.2009.11.075>
- [38] Robert L. Cook, Thomas Porter, and Loren Carpenter. 1984. Distributed ray tracing. In *Proceedings of the 11th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '84)*. ACM. <https://doi.org/10.1145/800031.808590>
- [39] Robert L. Cook and Kenneth E. Torrance. 1982. A reflectance model for computer graphics. *ACM Transactions on Graphics* 1, 1 (1982). <https://doi.org/10.1145/357290.357293>

- [40] Sergio Cova, Massimo Ghioni, Andrea Lacaita, Carlo Samori, and Franco Zappa. 1996. Avalanche photodiodes and quenching circuits for single-photon detection. *Applied Optics* 35, 12 (1996). <https://doi.org/10.1364/AO.35.001956>
- [41] Guy Hubert Stephane Sylvain Culeron, Song Shuo, Ping Wang, Liang Yang, and Chun chuan Liang. 2016. Glossy container. US Patent Application 2016/0375624 A1.
- [42] Tom Cuypers, Tom Haber, Philippe Bekaert, Se Baek Oh, and Ramesh Raskar. 2012. Reflectance model for diffraction. *ACM Transactions on Graphics* 31, 5 (2012). <https://doi.org/10.1145/2231816.2231820>
- [43] James R. DeVore. 1951. Refractive indices of rutile and sphalerite. *Journal of the Optical Society of America* 41, 6 (1951).
- [44] Zhao Dong, Bruce Walter, Steve Marschner, and Donald P. Greenberg. 2015. Predicting appearance from measured microgeometry of metal surfaces. *ACM Transactions on Graphics* 35, 1 (2015). <https://doi.org/10.1145/2815618>
- [45] Craig Donner and Henrik Wann Jensen. 2005. Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics* 24, 3 (2005). <https://doi.org/10.1145/1073204.1073308>
- [46] Justin Dove and Jeffrey H. Shapiro. 2019. Paraxial theory of phasor-field imaging. *Optics Express* 27, 13 (2019). <https://doi.org/10.1364/OE.27.018016>
- [47] Justin Dove and Jeffrey H. Shapiro. 2020. Nonparaxial phasor-field propagation. *Optics Express* 28, 20 (2020). <https://doi.org/10.1364/OE.401203>
- [48] Justin Dove and Jeffrey H. Shapiro. 2020. Speckled speckled speckle. *Optics Express* 28, 15 (2020). <https://doi.org/10.1364/OE.398226>
- [49] James J. Duderstadt and William R. Martin. 1979. *Transport theory*. John Wiley & Sons.
- [50] David Eberly. 1999. Perspective projection of an ellipse.
- [51] Albert Einstein. 1905. Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. *Annalen der Physik* 17, 6 (1905). <https://doi.org/10.1002/andp.19053220607>
- [52] Oskar Elek, Denis Sumin, Ran Zhang, Tim Weyrich, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, and Jaroslav Křivánek. 2017. Scattering-aware texture reproduction for 3D printing. *ACM Transactions on Graphics* 36, 6 (2017). <https://doi.org/10.1145/3130800.3130890>
- [53] Serkan Ergun, Sermet Önel, and Aydin Ozturk. 2016. A general micro-flake model for predicting the appearance of car paint. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations (EGSR '16)*. Eurographics Association. <https://doi.org/10.2312/sre.20161211>
- [54] Sergey Ershov, Konstantin Kolchin, and Karol Myszkowski. 2001. Rendering pearlescent appearance based on paint-composition modelling. *Computer Graphics Forum* 20, 3 (2001). <https://doi.org/10.1111/1467-8659.00515>

- [55] Daniele Faccio and Andreas Velten. 2018. A trillion frames per second: the techniques and applications of light-in-flight photography. *Reports on Progress in Physics* 81, 10 (2018). <https://doi.org/10.1088/1361-6633/aacca1>
- [56] Marcos Fajardo. 2010. Ray tracing solution for film production rendering. *ACM SIGGRAPH Talks: Global Illumination Across Industries* (2010).
- [57] Viggo Falster, Adrian Jarabo, and Jeppe Revall Frisvad. 2020. Computing the bidirectional scattering of a microstructure using scalar diffraction theory and path tracing. *Computer Graphics Forum* 39, 7 (2020). <https://doi.org/10.1111/cgf.14140>
- [58] Luca Fascione, Johannes Hanika, Mark Leone, Marc Droske, Jorge Schwarzhaupt, Tomáš Davidovič, Andrea Weidlich, and Johannes Meng. 2018. Manuka: A batch-shading architecture for spectral path tracing in movie production. *ACM Transactions on Graphics* 37, 3 (2018). <https://doi.org/10.1145/3182161>
- [59] Alejandro Ferrero, Berta Bernad, Joaquin Campos Campos, Esther Perales, José Luis Velázquez, and Francisco M. Martínez-Verdú. 2016. Color characterization of coatings with diffraction pigments. *Journal of the Optical Society of America A* 33, 10 (2016). <https://doi.org/10.1364/JOSAA.33.001978>
- [60] Alejandro Ferrero, Esther Perales, Ana M. Rabal, Joaquin Campos Campos, Francisco Miguel Martínez-Verdú, Elizabet Chorro, and A. Pons. 2014. Color representation and interpretation of special effect coatings. *Journal of the Optical Society of America A* 31, 2 (2014). <https://doi.org/10.1364/JOSAA.31.000436>
- [61] Augustin Jean Fresnel. 1823. *Mémoire sur la loi des modifications que la réflexion imprime à la lumière polarisée*. De l’Imprimerie De Firmin Didot Frères.
- [62] Jeppe Revall Frisvad, Niels Jørgen Christensen, and Henrik Wann Jensen. 2007. Computing the scattering properties of participating media using Lorenz-Mie theory. *ACM Transactions on Graphics* 26, 3 (2007). <https://doi.org/10.1145/1276377.1276452>
- [63] Miguel Galindo, Julio Marco, Matthew O’Toole, Gordon Wetzstein, Diego Gutierrez, and Adrian Jarabo. 2019. A dataset for benchmarking time-resolved non-line-of-sight imaging. In *ACM SIGGRAPH 2019 Posters (SIGGRAPH ’19)*. ACM. <https://graphics.unizar.es/nlos>
- [64] Luis E. Gamboa, Adrien Gruson, and Derek Nowrouzezahrai. 2020. An Efficient Transport Estimator for Complex Layered Materials. *Computer Graphics Forum* 39, 2 (2020). <https://doi.org/10.1111/cgf.13936>
- [65] Manuel N. Gamito. 2016. Solid angle sampling of disk and cylinder lights. *Computer Graphics Forum* 35, 4 (2016). <https://doi.org/10.1111/cgf.12946>
- [66] Iliyan Georgiev, Jaroslav Krivánek, Tomáš Davidovič, and Philipp Slusallek. 2012. Light transport simulation with vertex connection and merging. *ACM Transactions on Graphics* 31, 6 (2012). <https://doi.org/10.1145/2366145.2366211>
- [67] Ioannis Gkioulekas, Anat Levin, Frédo Durand, and Todd Zickler. 2015. Micron-scale light transport decomposition using interferometry. *ACM Transactions on Graphics* 34, 4 (2015). <https://doi.org/10.1145/2766928>

- [68] Andrew S. Glassner. 1995. *Principles of digital image synthesis*. Morgan Kaufmann.
- [69] Jay S. Gondek, Gary W. Meyer, and Jonathan G. Newman. 1994. Wavelength dependent reflectance functions. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '94)*. ACM. <https://doi.org/10.1145/192161.192202>
- [70] Joseph W. Goodman. 2005. *Introduction to Fourier optics*. Roberts and Company Publishers.
- [71] Xavier Granier and Wolfgang Heidrich. 2003. A simple layered RGB BRDF model. *Graphical Models* 65, 4 (2003). [https://doi.org/10.1016/S1524-0703\(03\)00042-0](https://doi.org/10.1016/S1524-0703(03)00042-0)
- [72] Adrien Gruson, Mickaël Ribardière, Martin Šik, Jiří Vorba, Rémi Cozot, Kadi Bouatouch, and Jaroslav Krivánek. 2016. A spatial target function for metropolis photon tracing. *ACM Transactions on Graphics* 36, 4 (2016). <https://doi.org/10.1145/3072959.2963097>
- [73] Ibón Guillén, Xiaochun Liu, Andreas Velten, Diego Gutierrez, and Adrian Jarabo. 2020. On the effect of reflectance on phasor field non-line-of-sight imaging. In *Proceedings of the 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2020)*. IEEE.
- [74] 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>
- [75] Ibón Guillén, Carlos Ure na, Alan King, Marcos Fajardo, Iliyan Georgiev, Jorge López-Moreno, and Adrian Jarabo. 2017. Area-preserving parameterizations for spherical ellipses. *Computer Graphics Forum* 36, 4 (2017). <https://doi.org/10.1111/cgf.13234>
- [76] Jie Guo, Yanjun Chen, Yanwen Guo, and Jingui Pan. 2018. A physically-based appearance model for special effect pigments. *Computer Graphics Forum* 37, 4 (2018). <https://doi.org/10.1111/cgf.13476>
- [77] Jie Guo, Jinghui Qian, Yanwen Guo, and Jingui Pan. 2016. Rendering thin transparent layers with extended normal distribution functions. *IEEE Transactions on Visualization and Computer Graphics* 23, 9 (2016). <https://doi.org/10.1109/TVCG.2016.2617872>
- [78] Yu Guo, Miloš Hašan, and Shuang Zhao. 2018. Position-free monte carlo simulation for arbitrary layered BSDFs. *ACM Transactions on Graphics* 37, 6 (2018). <https://doi.org/10.1145/3272127.3275053>
- [79] Mohit Gupta, Shree K. Nayar, Matthias B. Hullin, and Jaime Martin. 2015. Phasor imaging: A generalization of correlation-based time-of-flight imaging. *ACM Transactions on Graphics* 34, 5 (2015). <https://doi.org/10.1145/2735702>
- [80] Otkrist Gupta, Thomas Willwacher, Andreas Velten, Ashok Veeraraghavan, and Ramesh Raskar. 2012. Reconstruction of hidden 3D shapes using diffuse reflections. *Optics express* 20, 17 (2012). <https://doi.org/10.1364/OE.20.019096>

- [81] Diego Gutierrez, Adolfo Muñoz, Oscar Anson, and Francisco Seron. 2005. Non-linear volume photon mapping. In *Proceedings of the 16th Eurographics Symposium on Rendering (EGSR '05)*. Eurographics Association.
- [82] Toshiya Hachisuka, Wojciech Jarosz, Iliyan Georgiev, Anton Kaplanyan, Derek Nowrouzezahrai, and Ben Spencer. 2013. State of the art in photon density estimation. In *SIGGRAPH Asia 2013 Courses (SIGGRAPH Asia '13)*. ACM. <https://doi.org/10.1145/2542266.2542281>
- [83] Toshiya Hachisuka, Shinji Ogaki, and Henrik Wann Jensen. 2008. Progressive photon mapping. *ACM Transactions on Graphics* 27, 5 (2008). <https://doi.org/10.1145/1409060.1409083>
- [84] Toshiya Hachisuka, Jacopo Pantaleoni, and Henrik Wann Jensen. 2012. A path space extension for robust light transport simulation. *ACM Transactions on Graphics* 31, 6 (2012). <https://doi.org/10.1145/2366145.2366210>
- [85] Pat Hanrahan and Wolfgang Krueger. 1993. Reflection from layered surfaces due to subsurface scattering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '93)*. ACM. <https://doi.org/10.1145/166117.166139>
- [86] Jens Hansegard. 2012. IKEA's new catalogs: Less pine, more pixels. *The Wall Street Journal* (23 Aug 2012). <https://www.wsj.com/articles/SB10000872396390444508504577595414031195148>
- [87] Miloš Hašan, Jaroslav Krivánek, Bruce Walter, and Kavita Bala. 2009. Virtual spherical lights for many-light rendering of glossy scenes. *ACM Transactions on Graphics* 28, 5 (2009).
- [88] Xiao D. He, Kenneth E. Torrance, François X. Sillion, and Donald P. Greenberg. 1991. A comprehensive physical model for light reflection. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '91)*. ACM. <https://doi.org/10.1145/122718.122738>
- [89] Ying He, Dongheng Zhang, Yang Hu, and Yan Chen. 2020. Non-line-of-sight Imaging with Radio Signals. In *Proceedings of the 12th Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC 2020)*. IEEE.
- [90] Felix Heide, Matthias B. Hullin, James Gregson, and Wolfgang Heidrich. 2013. Low-budget transient imaging using photonic mixer devices. *ACM Transactions on Graphics* 32, 4 (2013). <https://doi.org/10.1145/2461912.2461945>
- [91] Felix Heide, Matthew O'Toole, Kai Zang, David B. Lindell, Steven Diamond, and Gordon Wetzstein. 2019. Non-line-of-sight imaging with partial occluders and surface normals. *ACM Transactions on Graphics* 38, 3 (2019). <https://doi.org/10.1145/3269977>
- [92] Felix Heide, Lei Xiao, Wolfgang Heidrich, and Matthias B. Hullin. 2014. Diffuse mirrors: 3D reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors. In *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2014)*. IEEE. <https://doi.org/10.1109/CVPR.2014.418>
- [93] Eric Heitz. 2014. Understanding the masking-shadowing function in microfacet-based BRDFs. *Journal of Computer Graphics Techniques* 3, 2 (2014).

- [94] Eric Heitz. 2017. Analytical calculation of the solid angle subtended by an arbitrarily positioned ellipsoid to a point source. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 852 (2017). <https://doi.org/10.1016/j.nima.2017.02.004>
- [95] Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. 2015. The SGGX microflake distribution. *ACM Transactions on Graphics* 34, 4 (2015). <https://doi.org/10.1145/2766988>
- [96] Hideki Hirayama, Kazufumi Kaneda, Hideo Yamashita, and Yoshimi Monden. 2001. An accurate illumination model for objects coated with multilayer films. *Computers & Graphics* 25, 3 (2001).
- [97] Nicolas Holzschuch and Romain Pacanowski. 2017. A two-scale microfacet reflectance model combining reflection and diffraction. *ACM Transactions on Graphics* 36, 4 (2017). <https://doi.org/10.1145/3072959.3073621>
- [98] Nagaraj Goud Ireni, Ramanuj Narayan, Pratyay Basak, and K. V. S. N. Raju. 2016. Poly(thiourethane-urethane)-urea as anticorrosion coatings with impressive optical properties. *Polymer* 97 (2016).
- [99] Wenzel Jakob. 2010. Mitsuba renderer. <http://www.mitsuba-renderer.org>
- [100] Wenzel Jakob, Adam Arbree, Jonathan T Moon, Kavita Bala, and Steve Marschner. 2010. A radiative transfer framework for rendering materials with anisotropic structure. *ACM Transactions on Graphics* 29, 4 (2010). <https://doi.org/10.1145/1778765.1778790>
- [101] Wenzel Jakob, Eugene d'Eon, Otto Jakob, and Steve Marschner. 2014. A comprehensive framework for rendering layered materials. *ACM Transactions on Graphics* 33, 4 (2014). <https://doi.org/10.1145/2601097.2601139>
- [102] Wenzel Jakob, Miloš Hašan, Ling-Qi Yan, Jason Lawrence, Ravi Ramamoorthi, and Steve Marschner. 2014. Discrete stochastic microfacet models. *ACM Transactions on Graphics* 33, 4 (2014). <https://doi.org/10.1145/2601097.2601186>
- [103] Wenzel Jakob, Christian Regg, and Wojciech Jarosz. 2011. Progressive expectation-maximization for hierarchical volumetric photon mapping. *Computer Graphics Forum* 30, 4 (2011). <https://doi.org/10.1111/j.1467-8659.2011.01988.x>
- [104] Adrian Jarabo, Carlos Aliaga, and Diego Gutierrez. 2018. A radiative transfer framework for spatially-correlated materials. *ACM Transactions on Graphics* 37, 4 (2018). <https://doi.org/10.1145/3197517.3201282>
- [105] Adrian Jarabo and Victor Arellano. 2018. Bidirectional rendering of vector light transport. *Computer Graphics Forum* 37, 6 (2018). <https://doi.org/10.1111/cgf.13314>
- [106] Adrian Jarabo, Julio Marco, Adolfo Muñoz, Raul Buisan, Wojciech Jarosz, and Diego Gutierrez. 2014. A framework for transient rendering. *ACM Transactions on Graphics* 33, 6 (2014). <https://doi.org/10.1145/2661229.2661251>
- [107] Adrian Jarabo, Belen Masia, Julio Marco, and Diego Gutierrez. 2017. Recent advances in transient imaging: A computer graphics and vision perspective. *Visual Informatics* 1, 1 (2017). <https://doi.org/10.1016/j.visinf.2017.01.008>

- [108] Adrian Jarabo, Belen Masia, Andreas Velten, Christopher Barsi, Ramesh Raskar, and Diego Gutierrez. 2015. Relativistic effects for time-resolved light transport. *Computer Graphics Forum* 34, 8 (2015). <https://doi.org/10.1111/cgf.12604>
- [109] Wojciech Jarosz, Derek Nowrouzezahrai, Iman Sadeghi, and Henrik Wann Jensen. 2011. A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM Transactions on Graphics* 30, 1 (2011). <https://doi.org/10.1145/1899404.1899409>
- [110] Wojciech Jarosz, Derek Nowrouzezahrai, Robert Thomas, Peter-Pike Sloan, and Matthias Zwicker. 2011. Progressive photon beams. *ACM Transactions on Graphics* 30, 6 (2011). <https://doi.org/10.1145/2070781.2024215>
- [111] Wojciech Jarosz, Matthias Zwicker, and Henrik Wann Jensen. 2008. The beam radiance estimate for volumetric photon mapping. In *SIGGRAPH 2008 Classes (SIGGRAPH '08)*. ACM. <https://doi.org/10.1145/1401132.1401137>
- [112] Henrik Wann Jensen. 2001. *Realistic image synthesis using photon mapping*. AK Peters.
- [113] Henrik Wann Jensen and Per H. Christensen. 1998. Efficient Simulation of Light Transport in Scenes with Participating Media Using Photon Maps. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. ACM. <https://doi.org/10.1145/280814.280925>
- [114] Alisa Jung, Alexander Wilkie, Johannes Hanika, Wenzel Jakob, and Carsten Dachsbacher. 2019. Wide gamut spectral upsampling with fluorescence. *Computer Graphics Forum* 38, 4 (2019). <https://doi.org/10.1111/cgf.13773>
- [115] Achuta Kadambi, Refael Whyte, Ayush Bhandari, Lee Streeter, Christopher Barsi, Adrian Dorrington, and Ramesh Raskar. 2013. Coded time of flight cameras: sparse deconvolution to address multipath interference and recover time profiles. *ACM Transactions on Graphics* 32, 6 (2013). <https://doi.org/10.1145/2508363.2508428>
- [116] Achuta Kadambi, Hang Zhao, Boxin Shi, and Ramesh Raskar. 2016. Occluded imaging with time-of-flight sensors. *ACM Transactions on Graphics* 35, 2 (2016). <https://doi.org/10.1145/2836164>
- [117] James T. Kajiya. 1986. The rendering equation. In *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '86)*. ACM. <https://doi.org/10.1145/15922.15902>
- [118] Anton S. Kaplanyan and Carsten Dachsbacher. 2013. Adaptive progressive photon mapping. *ACM Transactions on Graphics* 32, 2 (2013). <https://doi.org/10.1145/2451236.2451242>
- [119] Ori Katz, Eran Small, and Yaron Silberberg. 2012. Looking around corners and through thin turbid layers in real time with scattered incoherent light. *Nature Photonics* 6, 8 (2012). <https://doi.org/10.1038/nphoton.2012.150>
- [120] William B. Kerr and Fabio Pellacini. 2010. Toward evaluating material design interface paradigms for novice users. *ACM Transactions on Graphics* 29, 4 (2010). <https://doi.org/10.1145/1778765.1778772>

- [121] Wilhelm H. Kettler and Gunter Richter. 1997. Investigation on topology of platelet-like effect-pigments in automotive surface-coatings. *Progress in Organic Coatings* 31, 4 (1997).
- [122] Duck Bong Kim, Myoung Kook Seo, Kang Yeon Kim, and Kwan H Lee. 2010. Acquisition and representation of pearlescent paints using an image-based goniospectrophotometer. *Optical engineering* 49, 4 (2010). <https://doi.org/10.1117/1.3407431>
- [123] Eric Kirchner. 2009. Film shrinkage and flake orientation. *Progress in Organic Coatings* 65, 3 (2009).
- [124] Eric Kirchner and Jacqueline Houweling. 2009. Measuring flake orientation for metallic coatings. *Progress in Organic Coatings* 64, 2-3 (2009).
- [125] Ahmed Kirmani, Tyler Hutchison, James Davis, and Ramesh Raskar. 2011. Looking around the corner using ultrafast transient imaging. *International journal of computer vision* 95, 1 (2011). <https://doi.org/10.1007/s11263-011-0470-y>
- [126] Jonathan Klein, Christoph Peters, Jaime Martín, Martin Laurenzis, and Matthias B Hullin. 2016. Tracking objects outside the line of sight using 2D intensity images. *Scientific Reports* 6 (2016). <https://doi.org/10.1038/srep32491>
- [127] Claude Knaus and Matthias Zwicker. 2011. Progressive photon mapping: A probabilistic approach. *ACM Transactions on Graphics* 30, 3 (2011). <https://doi.org/10.1145/1966394.1966404>
- [128] 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). <https://doi.org/10.1111/cgf.13772>
- [129] Christopher Kulla and Marcos Fajardo. 2012. Importance sampling techniques for path tracing in participating media. *Computer Graphics Forum* 31, 4 (2012). <https://doi.org/10.1111/j.1467-8659.2012.03148.x>
- [130] Peter Kutz, Ralf Habel, Yining Karl Li, and Jan Novák. 2017. Spectral and decomposition tracking for rendering heterogeneous volumes. *ACM Transactions on Graphics* 36, 4 (2017). <https://doi.org/10.1145/3072959.3073665>
- [131] Jaroslav Křivánek, Iliyan Georgiev, Toshiya Hachisuka, Petr Vévoda, Martin Šik, Derek Nowrouzezahrai, and Wojciech Jarosz. 2014. Unifying points, beams, and paths in volumetric light transport simulation. *ACM Transactions on Graphics* 33, 4 (2014). <https://doi.org/10.1145/2601097.2601219>
- [132] Raphaël Labayrade, Henrik Wann Jensen, and Claus Jensen. 2009. Validation of Velux daylight visualizer 2 against CIE 171:2006 test cases. In *Proceedings of the 11th International IBPSA Conference*. International Building Performance Simulation Association.
- [133] Manuel Lagunas, Sandra Malpica, Ana Serrano, Elena Garces, Diego Gutierrez, and Belen Masia. 2019. A similarity measure for material appearance. *ACM Transactions on Graphics* 38, 4 (2019). <https://doi.org/10.1145/3306346.3323036>

- [134] Martin Laurenzis and Andreas Velten. 2014. Feature selection and back-projection algorithms for nonline-of-sight laser-gated viewing. *Journal of Electronic Imaging* 23, 6 (2014). <https://doi.org/10.1117/1.JEI.23.6.063003>
- [135] Martin Laurenzis and Andreas Velten. 2014. Nonline-of-sight laser gated viewing of scattered photons. *Optical Engineering* 53, 2 (2014). <https://doi.org/10.1117/1.OE.53.2.023102>
- [136] Zheng-Ping Li, Jun-Tian Ye, Xin Huang, Peng-Yu Jiang, Yuan Cao, Yu Hong, Chao Yu, Jun Zhang, Qiang Zhang, Cheng-Zhi Peng, Feihu Xu, and Jian-Wei Pan. 2021. Single-photon imaging over 200 km. *Optica* 8, 3 (2021). <https://doi.org/10.1364/OPTICA.408657>
- [137] David B. Lindell, Gordon Wetzstein, and Vladlen Koltun. 2019. Acoustic non-line-of-sight imaging. In *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2019)*. IEEE. <https://doi.org/10.1109/CVPR.2019.00694>
- [138] David B. Lindell, Gordon Wetzstein, and Matthew O'Toole. 2019. Wave-based non-line-of-sight imaging using fast f-k migration. *ACM Transactions on Graphics* 38, 4 (2019). <https://doi.org/10.1145/3306346.3322937>
- [139] Xiaochun Liu, Sebastian Bauer, and Andreas Velten. 2019. Analysis of feature visibility in non-line-of-sight measurements. In *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2019)*. IEEE. <https://doi.org/10.1109/CVPR.2019.01038>
- [140] Xiaochun Liu, Sebastian Bauer, and Andreas Velten. 2020. Phasor field diffraction based reconstruction for fast non-line-of-sight imaging systems. *Nature Communications* 11, 1 (2020). <https://doi.org/10.1038/s41467-020-15157-4>
- [141] Xiaochun Liu, Ibón Guillén, Marco La Manna, Ji Hyun Nam, Syed Azer Reza, Toan Huu Le, Diego Gutierrez, Adrian Jarabo, and Andreas Velten. 2019. Datasets and reconstruction code for a virtual wave non-line-of-sight imaging approach. <https://doi.org/10.6084/m9.figshare.8084987.v1>
- [142] Xiaochun Liu, Ibón Guillén, Marco La Manna, Ji Hyun Nam, Syed Azer Reza, Toan Huu Le, Diego Gutierrez, Adrian Jarabo, and Andreas Velten. 2019. Non-line-of-sight imaging using phasor-field virtual wave optics. *Nature* 572 (2019). <https://doi.org/10.1038/s41586-019-1461-3>
- [143] Xiaochun Liu and Andreas Velten. 2020. The role of Wigner distribution function in non-line-of-sight imaging. In *Proceedings of the 2020 IEEE International Conference on Computational Photography (ICCP 2020)*. IEEE. <https://doi.org/10.1109/ICCP48838.2020.9105266>
- [144] Boris Mahltig, Jieyang Zhang, Linfei Wu, Daniel Darko, Miriam Wendt, Evelyn Lempa, Maike Rabe, and Hajo Haase. 2017. Effect pigments for textile coating: a review of the broad range of advantageous functionalization. *Journal of Coatings Technology and Research* 14, 1 (2017).
- [145] Frank J. Maile, Gerhard Pfaff, and Peter Reynders. 2005. Effect pigments—past, present and future. *Progress in organic coatings* 54, 3 (2005).
- [146] Frank J. Maile and Peter Reynders. 2003. Substrates for pearlescent pigments. *European coatings journal* 4 (2003).

- [147] Irving H. Malitson. 1965. Interspecimen comparison of the refractive index of fused silica. *Journal of the Optical Society of America* 55, 10 (1965).
- [148] Irving H. Malitson and Marilyn J. Dodge. 1972. Refractive-index and birefringence of synthetic sapphire. *Journal of the Optical Society of America* 62, 11 (1972).
- [149] Julio Marco, Ibón Guillén, Wojciech Jarosz, Diego Gutierrez, and Adrian Jarabo. 2019. Progressive transient photon beams. *Computer Graphics Forum* 38, 6 (2019). <https://doi.org/10.1111/cgf.13600>
- [150] Julio Marco, Quercus Hernandez, Adolfo Muñoz, Yue Dong, Adrian Jarabo, Min H. Kim, Xin Tong, and Diego Gutierrez. 2017. DeepToF: Off-the-shelf real-time correction of multipath interference in time-of-flight imaging. *ACM Transactions on Graphics* 36, 6 (2017). <https://doi.org/10.1145/3130800.3130884>
- [151] Julio Marco, Wojciech Jarosz, Diego Gutierrez, and Adrian Jarabo. 2017. Transient photon beams. In *Proceedings of the 27th Spanish Computer Graphics Conference (CEIG 2017)*. The Eurographics Association. <https://doi.org/10.2312/ceig.20171216>
- [152] James Clerk Maxwell. 1865. A dynamical theory of the electromagnetic field. *Philosophical transactions of the Royal Society of London* 155 (1865). <https://doi.org/10.1098/rstl.1865.0008>
- [153] José M. Medina. 2008. Linear basis for metallic and iridescent colors. *Applied optics* 47, 30 (2008). <https://doi.org/10.1364/AO.47.005644>
- [154] Stephan Meister, Rahul Nair, Bernd Jähne, and Daniel Kondermann. 2013. *Photon mapping based simulation of multi-path Reflection artifacts in time-of-flight sensors*. Technical Report. Heidelberg Collaboratory for Image Processing.
- [155] Bailey Miller, Iliyan Georgiev, and Wojciech Jarosz. 2019. A null-scattering path integral formulation of light transport. *ACM Transactions on Graphics* 38, 4 (2019). <https://doi.org/10.1145/3306346.3323025>
- [156] Hans P. Moravec. 1981. 3D graphics and the wave theory. In *Proceedings of the 8th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '81)*. <https://doi.org/10.1145/800224.806817>
- [157] Satoshi Naganawa and Yuta Suzuki. 2016. Modified polysilazane film and method for producing gas barrier film. US Patent 9,512,334.
- [158] Nikhil Naik, Shuang Zhao, Andreas Velten, Ramesh Raskar, and Kavita Bala. 2011. Single view reflectance capture using multiplexed scattering and time-of-flight imaging. *ACM Transactions on Graphics* 30, 6 (2011). <https://doi.org/10.1145/2070781.2024205>
- [159] Srinivasa G. Narasimhan, Mohit Gupta, Craig Donner, Ravi Ramamoorthi, Shree K. Nayar, and Henrik Wann Jensen. 2006. Acquiring scattering properties of participating media by dilution. *ACM Transactions on Graphics* 25, 3 (2006). <https://doi.org/10.1145/1141911.1141986>
- [160] Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. 2019. Mitsuba 2: A retargetable forward and inverse renderer. *ACM Transactions on Graphics* 38, 6 (2019). <https://doi.org/10.1145/3355089.3356498>

- [161] Jan Novák, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. 2018. Monte Carlo methods for volumetric light transport simulation. *Computer Graphics Forum* 37, 2 (2018). <https://doi.org/10.1111/cgf.13383>
- [162] Matthew O’Toole, Felix Heide, David B. Lindell, Kai Zang, Steven Diamond, and Gordon Wetzstein. 2017. Reconstructing transient images from single-photon sensors. In *Proceedings of the 2017 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2017)*. IEEE. <https://doi.org/10.1109/CVPR.2017.246>
- [163] Matthew O’Toole, Felix Heide, Lei Xiao, Matthias B Hullin, Wolfgang Heidrich, and Kiriakos N Kutulakos. 2014. Temporal frequency probing for 5D transient analysis of global light transport. *ACM Transactions on Graphics* 33, 4 (2014). <https://doi.org/10.1145/2601097.2601103>
- [164] Matthew O’Toole, David B. Lindell, and Gordon Wetzstein. 2018. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature* 555 (2018). <https://doi.org/10.1038/nature25489>
- [165] Marios Papas, Wojciech Jarosz, Wenzel Jakob, Szymon Rusinkiewicz, Wojciech Matusik, and Tim Weyrich. 2011. Goal-based caustics. *Computer Graphics Forum* 30, 2 (2011). <https://doi.org/10.1111/j.1467-8659.2011.01876.x>
- [166] Gustavo Patow and Xavier Pueyo. 2003. A survey of inverse rendering problems. *Computer Graphics Forum* 22, 4 (2003). <https://doi.org/10.1111/j.1467-8659.2003.00716.x>
- [167] Frank Paxton. 1959. Solid angle calculation for a circular disk. *Review of Scientific Instruments* 30, 4 (1959). <https://doi.org/10.1063/1.1716590>
- [168] Fabio Pellacini, James A. Ferwerda, and Donald P. Greenberg. 2000. Toward a psychophysically-based light reflection model for image synthesis. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH ’00)*. ACM. <https://doi.org/10.1145/344779.344812>
- [169] Gerhard Pfaff. 2003. Special effect pigments based on silica flakes. *Inorganic materials* 39, 2 (2003).
- [170] Gerhard Pfaff and Miriam Becker. 2012. Special effect pigments in cosmetics applications. *Household Personal Hold* 7, 1 (2012).
- [171] Gerhard Pfaff and Peter Reynders. 1999. Angle-dependent optical effects deriving from submicron structures of films and pigments. *Chemical reviews* 99, 7 (1999).
- [172] Matt Pharr, Brent Burley, Per Christensen, Marcos Fajardo, Luca Fascione, and Christopher Kulla. 2018. Design and implementation of modern production renderers. In *ACM SIGGRAPH 2018 Panels (SIGGRAPH ’18)*. ACM. <https://doi.org/10.1145/3209621.3214901>
- [173] Adrien Pilleboue, Gurprit Singh, David Coeurjolly, Michael Kazhdan, and Victor Ostromoukhov. 2015. Variance analysis for Monte Carlo integration. *ACM Transactions on Graphics* 34, 4 (2015). <https://doi.org/10.1145/2766930>

- [174] Michal Piovarči, Michael Foshey, Vahid Babaei, Szymon Rusinkiewicz, Wojciech Matusik, and Piotr Didyk. 2020. Towards spatially varying gloss reproduction for 3D printing. *ACM Transactions on Graphics* 39, 6 (2020). <https://doi.org/10.1145/3414685.3417850>
- [175] Phil Pitts, Arrigo Benedetti, Malcolm Slaney, and Phil Chou. 2014. *Time of flight tracer*. Technical Report MSR-TR-2014-142. Microsoft Research.
- [176] Petar Pjanic and Roger D Hersch. 2015. Color changing effects with anisotropic halftone prints on metal. *ACM Transactions on Graphics* 34, 6 (2015). <https://doi.org/10.1145/2816795.2818083>
- [177] Max Planck. 1901. Ueber das gesetz der energieverteilung im normalspectrum. *Annalen der Physik* 309, 3 (1901). <https://doi.org/10.1002/andp.19013090310>
- [178] Mikhail N. Polyanskiy. 2020. *Refractive index database*. Retrieved January 1, 2020 from <https://refractiveindex.info>
- [179] Marvin R. Querry. 1985. *Optical constants*. Technical Report CRDC-CR-85034. Missouri University, Kansas City, MO (United States).
- [180] Aleksandar D. Rakić. 1995. Algorithm for the determination of intrinsic optical constants of metal films: application to aluminum. *Applied Optics* 34, 22 (1995). <https://doi.org/10.1364/AO.34.004755>
- [181] Marco Renna, Ji Hyun Nam, Mauro Buttavava, Federica Villa, Andreas Velten, and Alberto Tosi. 2020. Fast-gated 16×1 SPAD array for non-line-of-sight imaging applications. *Instruments* 4, 2 (2020). <https://doi.org/10.3390/instruments4020014>
- [182] Syed Azer Reza, Marco La Manna, Sebastian Bauer, and Andreas Velten. 2019. Phasor field waves: A Huygens-like light transport model for non-line-of-sight imaging applications. *Optics Express* 27, 20 (Sept. 2019). <https://doi.org/10.1364/OE.27.029380>
- [183] Syed Azer Reza, Marco La Manna, Sebastian Bauer, and Andreas Velten. 2019. Phasor field waves: experimental demonstrations of wave-like properties. *Optics Express* 27, 22 (2019).
- [184] Ole Romer. 1676. Démonstration touchant le mouvement de la lumière trouvée par M. Römer de l'Académie Royale des Sciences. *Journal des Sçavans* 1676 (1676).
- [185] Michael Rösler, Frank J. Maile, and Adalbert Huber. 2008. The macroscopic appearance of effect coatings and its relationship to the local spatial and angular distribution of reflected light. In *Proceedings of American Coating Conference 2008*.
- [186] Martin Rump, Gero Mueller, Ralf Sarlette, Dirk Koch, and Reinhard Klein. 2008. Photo-realistic rendering of metallic car paint from image-based measurements. *Computer Graphics Forum* 27, 2 (2008). <https://doi.org/10.1111/j.1467-8659.2008.01150.x>
- [187] Szymon M. Rusinkiewicz. 1998. A new change of variables for efficient BRDF representation. In *Proceedings of the 9th Eurographics Workshop on Rendering*. Eurographics Association.

- [188] Iman Sadeghi, Adolfo Munoz, Philip Laven, Wojciech Jarosz, Francisco Seron, Diego Gutierrez, and Henrik Wann Jensen. 2012. Physically-based simulation of rainbows. *ACM Transactions on Graphics* 31, 1 (2012). <https://doi.org/10.1145/2077341.2077344>
- [189] Yuliy Schwartzburg, Romain Testuz, Andrea Tagliasacchi, and Mark Pauly. 2014. High-contrast computational caustic design. *ACM Transactions on Graphics* 33, 4 (2014). <https://doi.org/10.1145/2601097.2601200>
- [190] John Sell and Patrick O'Connor. 2014. The Xbox One system on a chip and Kinect sensor. *IEEE Micro* 34, 2 (2014). <https://doi.org/10.1109/MM.2014.9>
- [191] Pradeep Sen, Billy Chen, Gaurav Garg, Stephen R. Marschner, Mark Horowitz, Marc Levoy, and Hendrik Lensch. 2005. Dual photography. *ACM Transactions on Graphics* 24, 3 (2005). <https://doi.org/10.1145/1073204.1073257>
- [192] Ana Serrano, Diego Gutierrez, Karol Myszkowski, Hans-Peter Seidel, and Be-len Masia. 2016. An intuitive control space for material appearance. *ACM Transactions on Graphics* 35, 6 (2016). <https://doi.org/10.1145/2980179.2980242>
- [193] Christopher M. Seubert, Mark E. Nichols, J. Frey, Max Shtein, and Michael D. Thouless. 2016. The characterization and effects of microstructure on the appearance of platelet-polymer composite coatings. *Journal of Materials Science* 51, 5 (2016).
- [194] Fabin Shen and Anbo Wang. 2006. Fast-Fourier-transform based numerical integration method for the Rayleigh-Sommerfeld diffraction formula. *Applied Optics* 45, 6 (Feb. 2006). <https://doi.org/10.1364/AO.45.001102>
- [195] Hiroyuki Shiomi, Eiichirou Misaki, Maoya Adachi, and Fukuji Suzuki. 2008. High chroma pearlescent pigments designed by optical simulation. *Journal of Coatings Technology and Research* 5, 4 (2008).
- [196] Peter Shirley. 1991. Discrepancy as a quality measure for sample distributions. In *Proceedings of Eurographics '91*. Elsevier.
- [197] Peter Shirley and Kenneth Chiu. 1997. A low distortion map between disk and square. *Journal of graphics tools* 2, 3 (1997).
- [198] Peter Shirley, Changyaw Wang, and Kurt Zimmerman. 1996. Monte Carlo techniques for direct lighting calculations. *ACM Transactions on Graphics* 15, 1 (1996). <https://doi.org/10.1145/226150.226151>
- [199] Florian Simon, Johannes Hanika, and Carsten Dachsbacher. 2015. Rich-VPLs for improving the versatility of many-light methods. *Computer Graphics Forum* 34, 2 (2015).
- [200] Adam Smith, James Skorupski, and James Davis. 2008. *Transient rendering*. Technical Report UCSC-SOE-08-26. School of Engineering, University of California, Santa Cruz.
- [201] Brian E. Smits and Gary W. Meyer. 1992. Newton's colors: simulating interference phenomena in realistic image synthesis. In *Photorealism in Computer Graphics*. Springer.
- [202] James Speight. 2005. *Lange's handbook of chemistry* (16 ed.). McGraw-Hill Education.

- [203] Ben Spencer and Mark W. Jones. 2009. Into the blue: Better caustics through photon relaxation. *Computer Graphics Forum* 28, 2 (2009). <https://doi.org/10.1111/j.1467-8659.2009.01371.x>
- [204] Jos Stam. 1999. Diffraction shaders. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*. ACM. <https://doi.org/10.1145/311535.311546>
- [205] Jos Stam. 2001. An illumination model for a skin layer bounded by rough surfaces. In *Proceedings of the 12th Eurographics Workshop on Rendering (EGWR '01)*. Eurographics Association.
- [206] Katrin Steinbach and Ulrich Schmidt. 2010. Borosilicate pigments—Transparency meets brilliance and sparkle. *Cosmetic Science Technology* 2010 (2010).
- [207] Shlomi Steinberg. 2019. Analytic spectral integration of birefringence-induced iridescence. *Computer Graphics Forum* 38, 4 (2019). <https://doi.org/10.1111/cgf.13774>
- [208] Shuochen Su, Felix Heide, Robin Swanson, Jonathan Klein, Clara Callenberg, Matthias Hullin, and Wolfgang Heidrich. 2016. Material classification using raw time-of-flight measurements. In *Proceedings of the 2016 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2016)*. IEEE. <https://doi.org/10.1109/CVPR.2016.381>
- [209] Kartic Subr and Jan Kautz. 2013. Fourier analysis of stochastic sampling strategies for assessing bias and variance in integration. *ACM Transactions on Graphics* 32, 4 (2013). <https://doi.org/10.1145/2461912.2462013>
- [210] Denis Sumin, Tobias Rittig, Vahid Babaei, Thomas Nindel, Alexander Wilkie, Piotr Didyk, Bernd Bickel, Jaroslav Krivánek, Karol Myszkowski, and Tim Weyrich. 2019. Geometry-aware scattering compensation for 3D printing. *ACM Transactions on Graphics* 38, 4 (2019). <https://doi.org/10.1145/3306346.3322992>
- [211] Xin Sun, Kun Zhou, Stephen Lin, and Baining Guo. 2010. Line space gathering for single scattering in large scenes. *ACM Transactions on Graphics* 29, 4 (2010). <https://doi.org/10.1145/1778765.1778791>
- [212] Yinlong Sun. 2006. Rendering biological iridescences with RGB-based renderers. *ACM Transactions on Graphics* 25, 1 (2006). <https://doi.org/10.1145/1122501.1122506>
- [213] Yinlong Sun and Qiqi Wang. 2008. Interference shaders of thin films. *Computer Graphics Forum* 27, 6 (2008). <https://doi.org/10.1111/j.1467-8659.2007.01110.x>
- [214] Kenichiro Tanaka, Yasuhiro Mukaigawa, Takuya Funatomi, Hiroyuki Kubo, Yasuyuki Matsushita, and Yasushi Yagi. 2019. Material classification from time-of-flight distortions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 12 (2019). <https://doi.org/10.1109/TPAMI.2018.2869885>
- [215] Jeremy A. Teichman. 2019. Phasor field waves: a mathematical treatment. *Optics Express* 27, 20 (2019). <https://doi.org/10.1364/OE.27.027500>

- [216] Sebastian Thrun, Mike Montemerlo, Hendrik Dohlkamp, David Stavens, Andrei Aron, James Diebel, Philip Fong, John Gale, Morgan Halpenny, Gabriel Hoffmann, Kenny Lau, Celia Oakley, Mark Palatucci, Vaughan Pratt, Pascal Stang, Sven Strohband, Cedric Dupont, Lars-Erik Jendrossek, Christian Koenen, Charles Markey, Carlo Rummel, Joe van Niekerk, Eric Jensen, Philippe Alessandrini, Gary Bradski, Bob Davies, Scott Ettinger, Adrian Kaehler, Ara Nefian, and Pamela Mahoney. 2006. Stanley: The robot that won the DARPA Grand Challenge. *Journal of field Robotics* 23, 9 (2006). https://doi.org/10.1007/978-3-540-73429-1_1
- [217] Dan M. Timus, Maria Jose Prata, Shyam Lal Kalla, Mahmoud I. Abbas, Feda Öner, and Eduardo Galiano. 2007. Some further analytical results on the solid angle subtended at a point by a circular disk using elliptic integrals. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 580, 1 (2007). <https://doi.org/10.1016/j.nima.2007.05.055>
- [218] Robert F. Tobler, László Neumann, Mateu Sbert, and Werner Purgathofer. 1998. A new form factor analogy and its application to stochastic global illumination algorithms. In *Proceedings of the 9th Eurographics Workshop on Rendering (EGWR '98)*. Eurographics Association.
- [219] Antoine Toisoul and Abhijeet Ghosh. 2017. Practical acquisition and rendering of diffraction effects in surface reflectance. *ACM Transactions on Graphics* 36, 5 (2017). <https://doi.org/10.1145/3072959.3012001>
- [220] Kenneth E. Torrance and Ephraim M. Sparrow. 1967. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America* 57, 9 (1967). <https://doi.org/10.1364/JOSA.57.001105>
- [221] Chia-Yin Tsai, Aswin C. Sankaranarayanan, and Ioannis Gkioulekas. 2019. Beyond volumetric albedo — A surface optimization framework for non-line-of-sight imaging. In *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2019)*. IEEE. <https://doi.org/10.1109/CVPR.2019.00164>
- [222] Carlos Ureña. 2000. Computation of irradiance from triangles by adaptive sampling. *Computer Graphics Forum* 19, 2 (2000). <https://doi.org/10.1111/1467-8659.00452>
- [223] Carlos Ureña, Marcos Fajardo, and Alan King. 2013. An area-preserving parametrization for spherical rectangles. *Computer Graphics Forum* 32, 4 (2013). <https://doi.org/10.1111/cgf.12151>
- [224] Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Mounqi G Bawendi, and Ramesh Raskar. 2012. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communications* 3 (2012). <https://doi.org/10.1038/ncomms1747>
- [225] Andreas Velten, Di Wu, Adrian Jarabo, Belen Masia, Christopher Barsi, Chinmaya Joshi, Everett Lawson, Mounqi G. Bawendi, Diego Gutierrez, and Ramesh Raskar. 2013. Femto-photography: capturing and visualizing the propagation of light. *ACM Transactions on Graphics* 32, 4 (2013). <https://doi.org/10.1145/2461912.2461928>

- [226] Changyaw Wang. 1992. Physically correct direct lighting for distribution ray tracing. In *Graphics Gems III*. Academic Press Professional, Inc.
- [227] Ping Wang, Liang Yang, John Andrew McDaniel, Gian Armand Juliana DeBelder, and Gaoyang Wang. 2014. Pearlescent container. US Patent 8,859,067 B2.
- [228] Gregory J. Ward. 1992. Measuring and modeling anisotropic reflection. In *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '92)*. ACM. <https://doi.org/10.1145/133994.134078>
- [229] Andrea Weidlich and Alexander Wilkie. 2007. Arbitrarily layered micro-facet surfaces. In *Proceedings of the 5th international conference on Computer graphics and interactive techniques in Australia and Southeast Asia (GRAPHITE '07)*. ACM. <https://doi.org/10.1145/1321261.1321292>
- [230] Philippe Weier and Laurent Belcour. 2020. Rendering layered materials with anisotropic interfaces. *Journal of Computer Graphics Techniques* 9, 2 (2020).
- [231] Daniel Weiskopf, Ute Kraus, and Hanns Ruder. 1999. Searchlight and Doppler effects in the visualization of special relativity: A corrected derivation of the transformation of radiance. *ACM Transactions on Graphics* 18, 3 (1999). <https://doi.org/10.1145/336414.336459>
- [232] Sebastian Werner, Zdravko Velinov, Wenzel Jakob, and Matthias B. Hullin. 2017. Scratch iridescence: Wave-optical rendering of diffractive surface structure. *ACM Transactions on Graphics* 36, 6 (2017). <https://doi.org/10.1145/3130800.3130840>
- [233] Turner Whitted. 1980. An improved illumination model for shaded display. *Communications of the ACM* 23, 6 (1980). <https://doi.org/10.1145/358876.358882>
- [234] Alexander Wilkie, Sehara Nawaz, Marc Droske, Andrea Weidlich, and Johannes Hanika. 2014. Hero wavelength spectral sampling. *Computer Graphics Forum* 33, 4 (2014). <https://doi.org/10.1111/cgf.12419>
- [235] Alexander Wilkie, Andrea Weidlich, Marcus Magnor, and Alan Chalmers. 2009. Predictive rendering. In *SIGGRAPH Asia 2009 Courses (SIGGRAPH Asia '09)*. ACM. <https://doi.org/10.1145/1665817.1665829>
- [236] Michael Maurice Rudolph Williams. 1971. *Mathematical methods in particle transport theory*. Butterworths.
- [237] Josh Wills, Sameer Agarwal, David Kriegman, and Serge Belongie. 2009. Toward a perceptual space for gloss. *ACM Transactions on Graphics* 28, 4 (2009). <https://doi.org/10.1145/1559755.1559760>
- [238] Cheng Wu, Jianjiang Liu, Xin Huang, Zheng-Ping Li, Chao Yu, Jun-Tian Ye, Jun Zhang, Qiang Zhang, Xiankang Dou, Vivek K. Goyal, Feihu Xu, and Jian-Wei Pan. 2021. Non-line-of-sight imaging over 1.43 km. *Proceedings of the National Academy of Sciences* 118, 10 (2021). <https://doi.org/10.1073/pnas.2024468118>
- [239] Rihui Wu, Adrian Jarabo, Jinli Suo, Feng Dai, Yongdong Zhang, Qionghai Dai, and Diego Gutierrez. 2018. Adaptive polarization-difference transient imaging

- for depth estimation in scattering media. *Optics Letters* 43, 6 (2018). <https://doi.org/10.1364/OL.43.001299>
- [240] Mengqi Xia, Bruce Walter, Christophe Hery, and Steve Marschner. 2020. Gaussian Product Sampling for Rendering Layered Materials. *Computer Graphics Forum* 39, 1 (2020). <https://doi.org/10.1111/cgf.13883>
- [241] Shumian Xin, Sotiris Nousias, Kiriakos N Kutulakos, Aswin C Sankaranarayanan, Srinivasa G Narasimhan, and Ioannis Gkioulekas. 2019. A theory of fermat paths for non-line-of-sight shape reconstruction. In *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2019)*. IEEE. <https://doi.org/10.1109/CVPR.2019.00696>
- [242] Tomoya Yamaguchi, Tatsuya Yatagawa, Yusuke Tokuyoshi, and Shigeo Morishima. 2019. Real-time rendering of layered materials with anisotropic normal distributions. In *SIGGRAPH Asia 2019, Technical Briefs*. ACM.
- [243] Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering specular microgeometry with wave optics. *ACM Transactions on Graphics* 37, 4 (2018). <https://doi.org/10.1145/3197517.3201351>
- [244] Ling-Qi Yan, Miloš Hašan, Steve Marschner, and Ravi Ramamoorthi. 2016. Position-normal distributions for efficient rendering of specular microstructure. *ACM Transactions on Graphics* 35, 4 (July 2016). <https://doi.org/10.1145/2897824.2925915>
- [245] Dorian Chan Yao. 2021. On F-K migration and non-line-of-sight imaging. (2021).
- [246] Pochi Yeh. 1988. *Optical waves in layered media*. John Wiley & Sons, Ltd.
- [247] Thomas Young. 1807. *A course of lectures on natural philosophy and the mechanical arts: in two volumes*. Vol. 2. Johnson.
- [248] Yonghao Yue, Kei Iwasaki, Bing-Yu Chen, Yoshinori Dobashi, and Tomoyuki Nishita. 2014. Poisson-based continuous surface generation for goal-based caustics. *ACM Transactions on Graphics* 33, 3 (2014). <https://doi.org/10.1145/2580946>
- [249] Tizian Zeltner and Wenzel Jakob. 2018. The layer laboratory: a calculus for additive and subtractive composition of anisotropic surface reflectance. *ACM Transactions on Graphics* 37, 4 (2018). <https://doi.org/10.1145/3197517.3201321>
- [250] Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2011. Building volumetric appearance models of fabric using micro CT imaging. *ACM Transactions on Graphics* 30, 4 (2011). <https://doi.org/10.1145/2010324.1964939>