



**FACULTY
OF MATHEMATICS
AND PHYSICS**
Charles University

DOCTORAL THESIS

Mgr. Tomáš Iser

**Affordable optical measurement
methods for predictive rendering**

Department of Software and Computer Science Education

Supervisor of the doctoral thesis: doc. Dr. Alexander Wilkie

Study programme: Computer Science

Study branch: Visual Computing and Computer
Games

Prague 2023

I declare that I carried out this doctoral thesis independently, and only with the cited sources, literature and other professional sources. It has not been used to obtain another or the same degree.

I understand that my work relates to the rights and obligations under the Act No. 121/2000 Sb., the Copyright Act, as amended, in particular the fact that the Charles University has the right to conclude a license agreement on the use of this work as a school work pursuant to Section 60 subsection 1 of the Copyright Act.

In date

Author's signature

Αφιερωμένο στην κοπέλα μου, Δεσποινούλα.

Dedicated to my girlfriend, Despoinoula.

I would like to thank my former supervisor, Jaroslav Křivánek, who had an admirable passion and dedication to computer graphics. Thanks to him, his great knowledge and amazing advice, I was able to start the interesting research that eventually resulted in this thesis. My many thanks go to Alexander Wilkie, who, without hesitation, took over my supervision after the unfortunate passing of Jaroslav. I believe that without Alexander, the Computer Graphics Group at Charles University would not have survived, and he ensured that our research continued without interruption. I thank all my colleagues for their great support, namely Tobias Rittig, who spent many long days with me debugging our optical measurement setups. I thank Wenzel Jakob and his team in Lausanne for their endless support in rendering and for hosting my research visit to their material measurement laboratory. Finally, I thank my family and my girlfriend for their incredible love and support in my life and studies, and for ensuring that everything goes smoothly throughout the harsh deadlines.

The work contained in this thesis has received funding from the GA UK project 1165220 of Charles University, Faculty of Mathematics and Physics; from the Czech Science Foundation under grants numbers GAČR-19-07626S and GAČR-22-22875S; from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 956585; and from the Charles University grants SVV-260588 and SVV-260699.

In our renders, we use the “rabbit” 3D model from [Crealitty Cloud](#) licensed as [CC BY 4.0](#), and the [environment map](#) from Sergej Majboroda, licensed as [CC0](#).

Title: Affordable optical measurement methods for predictive rendering

Author: Mgr. Tomáš Iser

Department: Department of Software and Computer Science Education

Supervisor: doc. Dr. Alexander Wilkie, Department of Software and Computer Science Education

Abstract: Predictive rendering, a part of computer graphics, is based on the light transport equation and focuses on accurately predicting the appearance of objects and materials under various conditions. A variety of problems can be formulated as appearance prediction: from generating photorealistic images to enhancing color 3D printing. The accuracy relies on the materials' optical properties, which must either be estimated from first principles, or measured with expensive and sophisticated optical devices. Could we obtain these properties in an efficient and affordable way optimized for predictive rendering?

To answer the question, this thesis bridges the boundary between computer graphics and optics. We develop simpler and more affordable methods for measuring optical properties with a focus on color accuracy, thus making predictive rendering more accessible. We aim at two types of materials that are both ubiquitous but usually neglected because of their complex characteristics: *translucent materials* and *fluorescent materials*. For each, we present a separate measurement approach that only uses low-cost optical components, yet has a high spectral resolution for color-accurate applications.

Our first method is motivated by measuring translucent inks, which is required for accurate full-color 3D-printing algorithms. We develop an acquisition technique for the three unknown material parameters, namely, the absorption and scattering coefficients, and the phase function anisotropy factor. Only three point measurements with a spectrometer are required, as we found a three-dimensional appearance map, computed using Monte Carlo rendering, that allows the conversion between the three observables and the material parameters.

Our second method focuses on fluorescent materials. A Donaldson matrix is estimated for each material, which corresponds to a two-dimensional spectral characterization of the fluorescence and reflectance properties. Only a few measurements of the material's reflectance under a few illuminants are needed for the estimation with our algorithm. It is enabled by representing each Donaldson matrix with a multivariate Gaussian mixture model and bounded MESE (maximum entropy spectral estimate) and using gradient-descent optimization.

Keywords: predictive rendering, optical properties, measurement, translucent materials, fluorescent materials, affordable

Contents

| | |
|---|-----------|
| Preface | 4 |
| 1 Introduction | 5 |
| 1.1 High-level motivation | 5 |
| 1.2 Affordability | 6 |
| 1.3 Original contribution | 7 |
| 1.3.1 Translucent materials (Chapter 4) | 7 |
| 1.3.2 Fluorescent materials (Chapter 5) | 7 |
| 1.4 List of all publications | 8 |
| 1.4.1 First-authored journal publications | 8 |
| 1.4.2 Other publications | 8 |
| 1.5 Organization of the thesis | 12 |
| | |
| I Background | 13 |
| | |
| 2 Rendering materials | 14 |
| 2.1 Fundamental background | 14 |
| 2.1.1 Radiometric quantities | 14 |
| 2.1.2 Measurement equation | 16 |
| 2.1.3 Measuring colors | 17 |
| 2.1.4 Light transport equation | 18 |
| 2.1.5 Bi-directional scattering distribution function | 19 |
| 2.1.6 Solving the equations with Monte Carlo | 20 |
| 2.1.7 Implementation | 21 |
| 2.2 Translucent materials, participating media | 22 |
| 2.2.1 Light interactions | 22 |
| 2.2.2 Transmittance | 25 |
| 2.2.3 Phase function | 25 |
| 2.2.4 Volume rendering equation | 27 |
| 2.3 Fluorescent materials | 28 |
| 2.3.1 Light interactions | 28 |
| 2.3.2 Bi-spectral light transport equation | 29 |
| 2.3.3 Simplified BRRDF | 30 |
| 2.3.4 Donaldson matrix | 32 |
| | |
| 3 Measuring materials as an inverse problem | 33 |
| 3.1 Problem statement | 33 |
| 3.2 Ambiguities and similarity relations | 34 |
| 3.2.1 Spectral ambiguities, metamerism, and other shortcomings of RGB optical properties | 35 |
| 3.2.2 Donaldson matrix ambiguities | 37 |
| 3.2.3 Similarity relations in participating media | 38 |
| 3.3 Solving the inverse problem | 39 |
| 3.3.1 Analytical solutions | 39 |

| | | |
|-----------|--|-----------|
| 3.3.2 | Lookup tables | 40 |
| 3.3.3 | Learning-based approaches with deep networks | 42 |
| 3.3.4 | Gradient-based iterative approaches | 43 |
| 3.3.5 | Gradient-free iterative approaches | 44 |
| 3.4 | Differentiable rendering | 45 |
| 3.4.1 | Automatic differentiation | 46 |
| 3.4.2 | Challenges and solutions | 46 |
| II | Journal publications | 48 |
| 4 | Our measurement method for translucent materials | 49 |
| 4.1 | Introduction | 51 |
| 4.2 | Related work | 52 |
| 4.2.1 | Prediction models | 52 |
| 4.2.2 | Measurement methods | 53 |
| 4.2.3 | Translucency appearance and color 3D printing | 53 |
| 4.3 | Method | 53 |
| 4.3.1 | Material appearance against diffuse backgrounds | 54 |
| 4.3.2 | Phase function | 55 |
| 4.3.3 | Index of refraction | 56 |
| 4.4 | Measurement | 56 |
| 4.4.1 | Measurement setup overview | 56 |
| 4.4.2 | Measurement procedure | 57 |
| 4.4.3 | Fitting the optical properties | 58 |
| 4.5 | Results | 58 |
| 4.5.1 | Measuring the base inks | 58 |
| 4.5.2 | Validation | 60 |
| 4.6 | Conclusion | 62 |
| 5 | Our measurement method for fluorescent materials | 65 |
| 5.1 | Introduction | 67 |
| 5.2 | Problem statement and prior work | 68 |
| 5.2.1 | Modeling fluorescence with the Donaldson matrix | 69 |
| 5.2.2 | Sparse measurements | 69 |
| 5.2.3 | Previous approaches finding a solution of Eq. (4) | 70 |
| 5.3 | Estimating fluorescence with Gaussian distributions and bounded MESE | 71 |
| 5.3.1 | Overview | 71 |
| 5.3.2 | Fluorescence as a Gaussian mixture | 71 |
| 5.3.3 | Diagonal as a bounded MESE | 72 |
| 5.3.4 | Estimating the Donaldson matrix via gradient descent | 73 |
| 5.3.5 | Algorithm implementation | 74 |
| 5.4 | Evaluation | 74 |
| 5.4.1 | Optical setups | 74 |
| 5.4.2 | Synthetic dataset and simulated measurements | 76 |
| 5.4.3 | Real measurements | 77 |
| 5.4.4 | Quantitative evaluation | 77 |
| 5.4.5 | Qualitative evaluation | 79 |

| | |
|--------------------------|-----------|
| 5.5 Conclusion | 80 |
| 6 Conclusion | 84 |
| Bibliography | 86 |
| List of Figures | 97 |

Preface

When I started my Ph.D. in the Computer Graphics Group at Charles University in 2019, I faced the following problem. We wanted to accurately predict the appearance of 3D-printed objects without actually printing them, so we could improve the printing quality by algorithmically optimizing the ink placement in the 3D printer. Essentially, the goal was to follow up on the previous research of our group, mainly by [Elek et al. \[2017\]](#) and [Sumin et al. \[2019\]](#), but in a more general and accurate approach that would allow previously unseen applications.

It was obvious from the beginning that high-quality predictions will require measuring the optical properties of the printing materials with high accuracy. While [Elek et al. \[2017\]](#) had already measured the properties, their method had several limitations, such as the inability to acquire the properties spectrally. At that time, we also noticed that the printing materials were not only translucent, with significant subsurface scattering, but they were also showing signs of fluorescence, which was especially difficult to measure.

The goal was very clear: measure the material properties spectrally, including their translucency and fluorescence, and use these measurements in 3D printing optimizations. The obstacles were also very clear: the existing methods were limited or relied on expensive optical setups, and we had virtually no budget to afford sophisticated measurements. This gave birth to what eventually became the mission of my Ph.D.: figure out how the complex optical properties could be measured in an affordable way, optimized for predictive rendering.

Meanwhile, I collaborated on an important publication on 3D printing optimization based on differentiable rendering [[Nindel et al., 2021](#)]. As side projects, I also co-authored our sky radiance model [[Wilkie et al., 2021](#)], a deep network generator of cloud distributions [[Mirbauer et al., 2022](#)], and a method for reconstructing 3D wood textures from 2D photographs [[Nindel et al., 2023](#)].

Most importantly, we managed to achieve what had been my main mission: the affordable material measurements. In our two journal publications, we showed a new method for affordable measurements of translucent materials [[Iser et al., 2022](#)], published in ACM Transactions on Graphics and presented in South Korea, and of fluorescent materials [[Iser et al., 2023](#)], published in Optics Express.

This thesis is a so-called *cumulative thesis*, which is a concept at certain universities and institutions, including our faculty. Instead of being a monograph, the thesis is rather a collection of the two inserted publications, which are linked together with an extended introduction, background on material rendering and inverse problems, and conclusion. I hope that this thesis manages to communicate the motivation, effort, and achievements in our material measurements, but also helps anyone who finds themselves researching similar problems on the boundary between computer graphics and optics.

1 | Introduction

1.1 High-level motivation

Visual content such as movies, marketing materials and advertisements, images in catalogs and e-shops, are more and more often not real photographs or video recordings. Instead, they are based on computer-generated imagery tailored to look similar to a real photograph or video taken by an actual camera. This is called **photorealistic rendering**, and these days, artists can choose from various computer algorithms and software responsible for such results.

The term *photorealistic* denotes that the rendered image fooled the viewer into thinking it was an actual photograph. Historically, *photorealism* is a genre of art that responded to photography, and it includes paintings and drawings by artists such as Richard Estes, which have nothing to do with computers and algorithms. However, modern algorithms paired with powerful computers brought something that hand-painted images can hardly compete with: **physically-based rendering**. It is an approach that uses accurate physical equations and optical models to compute how light interacts with materials in the world around us, and based on that, pixel by pixel, creates an image corresponding to what a real camera would see, following the laws of physics.

Even in physically-based rendering, one often takes various shortcuts to simplify the simulation and make it more tangible and faster to compute at the cost of accuracy. To denote a precise light simulation whose results are desired to be indistinguishable from reality for the given intentions and purposes, we sometimes use the term **predictive rendering**. The term comes from the desire to *predict* the true appearance of a certain material or optical phenomenon.

Such accuracy then allows replacing a physical product with its virtual equivalent, often called the **digital twin**. This is often motivated financially: if we can reliably examine an object on the computer, we can save manufacturing costs and avoid material waste on physical prototypes. Examples range from presenting customers with accurate previews of bespoke furniture manufactured on demand, to simulating the exact sky radiance for architectural visualizations. Additionally, predictive rendering has applications in *computational optimization*. For example, by predicting the appearance of 3D printouts, their ink distribution can be numerically optimized to improve contrast and color fidelity.

In this thesis, we focus on an inherent problem of predictive rendering: to render an object, the **optical properties** of the object's material need to be known beforehand. Sometimes, the properties are well-known or can be derived from first principles. For example, the light absorption and scattering of molecules can be computed from the Mie solution to Maxwell's equations, given we know which molecules they are and what their statistical distribution is. Often, though, it is unfeasible to determine the precise composition, and the goal is to *measure* the optical properties of an exact sample at hand.

While properties like temperature or weight can be measured directly with simple tools like thermometers or scales, optical properties usually need to be measured indirectly. Indirect measurements rely on designing and solving an appropriate **inverse problem**. It is called *inverse* because it is reversed compared

to the rendering itself: in rendering, we know the optical properties and use them to compute the material’s appearance in an arbitrary scene; in the inverse problem, we observe the material in controlled conditions, and then we infer the previously unknown optical properties that the material most likely had.

The key is to design a **measurement method** optimized for predictive rendering. On its input, the method should only rely on a small set of simple, affordable, sparse observations of the material sample. Based on that input, the method should output the material’s optical properties with accuracy sufficient to predictively render the same material in any arbitrary virtual scene. This thesis first investigates the problem in general (Chapters 2 and 3) and then proposes solutions for two specific kinds of materials that exhibit non-trivial light-matter interactions and for which such solutions were not available: **translucent materials** (Chapter 4) and **fluorescent materials** (Chapter 5).

1.2 Affordability

The core idea of our thesis is not only to *somehow* measure the optical properties for predictive rendering but to design measurement methods that are *affordable*. That does not only concern the price of the optical components but also minimizing the number of measurements and the time and effort required for executing the method without compromising the accuracy. In this section, we interpret affordability in the broader sense and explain why it is a crucial research objective.

Low number of observations The method should require as few observations of the material sample as possible to achieve sufficient predictive rendering accuracy. This significantly impacts the usability of the method. If too many observations are required, the method is usually unsuitable for manual measurements, leading to the construction of specialized sensor arrays, robotic arms, or automated stages that are unnecessarily complicated. The ideal measurement method would require only one observation, yet that is usually impossible in practice. The key is to study how exactly the optical properties influence the appearance of the material and then reduce the complex problem to a low-dimensional one while minimizing ambiguities. That requires the understanding of both the forward and inverse problem, which we cover in Chapters 2 and 3, respectively.

Low-cost components As the problem’s dimensionality is reduced to a few observations, these observations must be acquirable mostly with low-cost and off-the-shelf optical components. This excludes expensive hardware such as hyperspectral cameras, tunable filters, monochromators, stabilized and calibrated light sources, and various automated rotational stages. As we will see in specific examples in Chapters 4 and 5, existing methods often rely on costly devices, which makes the whole method expensive to replicate.

Easy to build, calibrate, and use An affordable method should also be reasonably effortless and fast to build, calibrate, and use. We will see that existing methods often require sophisticated calibration procedures that may take hours or even days of work, or they require purchasing expensive reference materials.

We prefer a simpler physical setup at the cost of more complicated and longer computations, because computers are continuously becoming faster and more accessible, while manual labor is becoming more expensive.

Benefits and impact Measurement methods that are affordable, as defined above, have several benefits and a large potential impact. The most obvious benefit is the saving of time and budget. Our thesis enables accurate measurements optimized for predictive rendering without accessing a fully-equipped optical laboratory. Affordability also opens doors to miniaturizing and integrating the method into standalone measurement devices. Furthermore, a low-cost measurement setup can be left assembled and only used rarely, whereas expensive setups are often disassembled every time to reuse their expensive components for different measurements. Finally, if the calibration is not complicated, the measurement may be done by an inexperienced student without requiring a qualified physicist always to be around.

1.3 Original contribution

In Chapters 4 and 5, this thesis presents *two original methods* that we designed for affordable optical measurements targeted at spectrally precise predictive rendering. The first method aims at translucent and the second at fluorescent materials.

1.3.1 Translucent materials (Chapter 4)

First, we focus on translucent materials, specifically absorbing and light-scattering participating media, which includes common liquids, aerosols, tissue, and even many solids. Our main motivation is measuring inks used in full-color 3D printers, as their optical properties are crucial for enhancing printing quality. We cover the fundamental light transport theory in participating media in Section 2.2. Later, in Sections 3.2.1 and 3.2.3, we explain the problems that arise when measuring translucency as an inverse problem, including non-linearities in RGB approaches and the similarity relations. Especially, it becomes clear that for use in predictive rendering, our measurement method needs to resolve the unknown properties spectrally, which is a non-trivial problem.

To solve the problem, we designed what we called the *appearance map*, which is a relation between the unknown optical properties and the appearance of the material. With this relation, it becomes affordable to spectrally resolve the absorption and scattering coefficients and the phase function anisotropy of the translucent material. For that purpose, we built a simple and low-cost optical setup based on acquiring the spectral intensity of the material sample against a black background, white background, and against a collimated beam. We demonstrated the quality of our results on a specially designed 3D-printed color checker.

1.3.2 Fluorescent materials (Chapter 5)

Second, we focus on fluorescent materials, which includes natural and synthetic objects, specifically pigments that are intended for color enhancement and optical brightening. We cover the fundamentals of fluorescent materials in Section 2.3,

where it becomes clear that the Donaldson matrix, which characterizes fluorescence of a material, can be measured in a brute-force method by simply scanning through all incoming and outgoing wavelengths. Since that has significant practical problems, sparse measurement methods were eventually developed.

Our solution is also based on sparse acquisitions, specifically, we only illuminate the sample with a few spectra and measure the corresponding reflectances. However, our method significantly outperforms the quality of the state of the art thanks to our parametrization. While the previous methods typically reduce the dimensionality by building a set of basis functions based on a dataset of materials, we instead rely on Gaussian mixtures and bounded maximum entropy spectral estimates. That ensures that our estimates have the necessary degrees of freedom while still being constrained to avoid physically implausible results.

1.4 List of all publications

This section contains the list of all publications that Tomáš Iser contributed to during his doctoral studies in the period from October 2019 to the submission of this thesis in July 2023.

1.4.1 First-authored journal publications

This thesis is based on the following two journal publications. Tomáš Iser was their first author, and they were published in high-ranking journals in computer graphics and optics, namely ACM Transactions on Graphics and Optics Express.

- [Iser et al. \[2022\]](#)

Affordable Spectral Measurements of Translucent Materials.

Tomáš Iser, Tobias Rittig, Emilie Nogué, Thomas Klaus Nindel, and Alexander Wilkie.

ACM Transactions on Graphics, 41(6), December 2022.

The content of this paper is inserted as a part of Chapter 4.

- [Iser et al. \[2023\]](#)

Affordable method for measuring fluorescence using Gaussian distributions and bounded MESE.

Tomáš Iser, Loïc Lachiver, and Alexander Wilkie.

Optics Express, 31(15):24347–24362, July 2023.

The content of this paper is inserted as a part of Chapter 5.

1.4.2 Other publications

During his doctoral studies, Tomáš Iser also co-authored the following publications. While each publication is related to this thesis, which we explain for each of them separately, they together do not form any cohesive framework, and they are not the main contribution of this thesis.

A Gradient-Based Framework for 3D Print Appearance Optimization

[Nindel et al., 2021]

Authors Thomas Klaus Nindel, **Tomáš Iser**, Tobias Rittig, Alexander Wilkie, and Jaroslav Křivánek.

Journal ACM Transactions on Graphics, 40(4), 2021.

Abstract In full-color inkjet 3D printing, a key problem is determining the material configuration for the millions of voxels that a printed object is made of. The goal is a configuration that minimizes the difference between desired target appearance and the result of the printing process. So far, the techniques used to find such a configuration have relied on domain-specific methods or heuristic optimization, which allowed only a limited level of control over the resulting appearance. We propose to use differentiable volume rendering in a continuous material-mixture space, which leads to a framework that can be used as a general tool for optimizing inkjet 3D printouts. We demonstrate the technical feasibility of this approach and use it to attain fine control over the fabricated appearance, and high levels of faithfulness to the specified target.

Author’s contribution Tomáš Iser significantly contributed to this publication by researching and designing appropriate error metrics and a suitable halftoning algorithm for the material-mixture space. He also prepared and executed many experiments and contributed to the paper writing.

Relation to this thesis The publication shows applications of the concepts explained in this thesis. First, optimizing the 3D printouts in the publication is based on inverse rendering, just like our measurements are. However, the publication relies on differentiable rendering using radiative backpropagation, which we cover in Section 3.4, but do not directly use in our thesis. Second, the publication relies on the optical properties of the translucent 3D-printing inks, whose measurements we focus on in Chapter 4.

A fitted radiance and attenuation model for realistic atmospheres

[Wilkie et al., 2021]

Authors Alexander Wilkie, Petr Vévoda, Thomas Bashford-Rogers, Lukáš Hošek, **Tomáš Iser**, Monika Kolářová, Tobias Rittig, and Jaroslav Křivánek.

Journal ACM Transactions on Graphics, 40(4), 2021.

Abstract We present a fitted model of sky dome radiance and attenuation for realistic terrestrial atmospheres. Using scatterer distribution data from atmospheric measurement data, our model considerably improves on the visual realism of existing analytical clear sky models, as well as of interactive methods that are based on approximating atmospheric light transport. We also provide

features not found in fitted models so far: radiance patterns for post-sunset conditions, in-scattered radiance and attenuation values for finite viewing distances, an observer altitude resolved model that includes downward-looking viewing directions, as well as polarisation information. We introduce a fully spherical model for in-scattered radiance that replaces the family of hemispherical functions originally introduced by Perez e.a., and which was extended for several subsequent analytical models: our model relies on reference image compression via tensor decomposition instead.

Author’s contribution Tomáš Iser contributed to this publication by constructing and executing validation examples against the atmospheric simulation package `libRadtran` and preparing the supplemental document.

Relation to this thesis This publication is only weakly related to this thesis. The sky is a participating medium whose properties are related to the particle concentrations in the atmosphere. This work is a great example of predictive rendering, where the optical properties can be calculated from first principles from known distributions of the particular atoms and molecules.

SkyGAN: Towards Realistic Cloud Imagery for Image Based Lighting

[[Mirbauer et al., 2022](#)]

Authors Martin Mirbauer, Tobias Rittig, **Tomáš Iser**, Jaroslav Křivánek, and Elena Šikudová.

Conference Eurographics Symposium on Rendering, 2022.

Abstract Achieving photorealism when rendering virtual scenes in movies or architecture visualizations often depends on providing a realistic illumination and background. Typically, spherical environment maps serve both as a natural light source from the Sun and the sky, and as a background with clouds and a horizon. In practice, the input is either a static high-resolution HDR photograph manually captured on location in real conditions, or an analytical clear sky model that is dynamic, but cannot model clouds. Our approach bridges these two limited paradigms: a user can control the sun position and cloud coverage ratio, and generate a realistically looking environment map for these conditions. It is a hybrid data-driven analytical model based on a modified state-of-the-art GAN architecture, which is trained on matching pairs of physically-accurate clear sky radiance and HDR fisheye photographs of clouds. We demonstrate our results on renders of outdoor scenes under varying time, date, and cloud covers.

Author’s contribution Tomáš Iser contributed to this research topic by experimenting with the connection of volumetric differentiable rendering and deep networks, which was, in the end, not published yet and remains future work. For this particular publication, Tomáš Iser assisted with the methodology and prepared an automated processing and error correction pipeline that turns input

sky photographs into a cohesive and accurate dataset. He also contributed to the paper writing.

Relation to this thesis This publication is a great example of connecting accurate predictive rendering based on real optical properties (the sky radiance) with imprecise, artificially generated content from a deep network (clouds). In Section 3.3.3, we briefly talk about inverse transport networks, which combine differentiable rendering with deep learning. In this case, the principle is different, although one could, in principle, also use a differentiable renderer to constrain the physical plausibility of the generated cloud cover. This remains future work that was not yet published.

Automatic inference of an anatomically meaningful solid wood texture from a single photograph

[Nindel et al., 2023]

Authors Thomas Klaus Nindel, Mohcen Hafidi, **Tomáš Iser**, and Alexander Wilkie.

Venue This work is yet to be published, its current revision is available on arXiv.

Abstract Wood is a volumetric material with a very large appearance gamut that is further enlarged by numerous finishing techniques. Computer graphics has made considerable progress in creating sophisticated and flexible appearance models that allow convincing renderings of wooden materials. However, these do not yet allow fully automatic appearance matching to a concrete exemplar piece of wood, and have to be fine-tuned by hand. More general appearance matching strategies are incapable of reconstructing anatomically meaningful volumetric information. This is essential for applications where the internal structure of wood is significant, such as non-planar furniture parts machined from a solid block of wood, translucent appearance of thin wooden layers, or in the field of dendrochronology. In this paper, we provide the two key ingredients for automatic matching of a procedural wood appearance model to exemplar photographs: a good initialization, built on detecting and modelling the ring structure, and a phase-based loss function that allows to accurately recover growth ring deformations and gives anatomically meaningful results. Our ring-detection technique is based on curved Gabor filters, and robustly works for a considerable range of wood types.

Author’s contribution Tomáš Iser contributed to this publication by preparing the processes, data, and visualizations, assisting with the methodology, and contributing to the paper writing.

Relation to this thesis This publication demonstrates *wood* as a translucent material, whose internal composition needs to be studied and its optical properties understood for accurate predictive rendering of wooden objects, including

furniture. It is capable of reconstructing the internal structure based on a single photograph, utilizing the principles of differentiable rendering and automatic differentiation (Section 3.4).

1.5 Organization of the thesis

In Chapter 2, we briefly introduce the relevant background in rendering, from the elementary knowledge such as radiometric quantities and the light transport equation, to translucency and fluorescence. In Chapter 3, we review how material measurements can be understood as an inverse problem to rendering, and we discuss recent work on that topic, including differentiable rendering. The following two chapters contain the core of the thesis, the main contribution of the author. Chapter 4 describes our novel method on measuring translucent materials based on Iser et al. [2022], and Chapter 5 continues with our novel method on measuring fluorescent materials based on Iser et al. [2023]. The thesis is concluded in Chapter 6 together with a discussion on future research.

Part I

Background

2 | Rendering materials

In general, *rendering* is a process that generates an image based on objects and materials in a virtual scene. The material properties must be known in advance so that the rendering algorithm can simulate the image of the scene. This chapter is devoted to the relevant topics about rendering, such that later, in Chapter 3, we can introduce the inverse problem. Because rendering is a broad term with hundreds of books written on it, we only cover the topics that are critical for the context of the thesis. Throughout the text, we always refer to relevant publications that provide more detail.

From a broader perspective, rendering does not need to output an image. The rendering algorithms merely compute what a light sensor (e.g., a digital camera) would detect if it were placed in the scene. This involves calculating how light spreads within the scene, from one object to another, and how exactly the light interacts with the materials. The sensor can be purely hypothetical. For example, in 3D printing appearance optimization, we compute the light on the surface of a three-dimensional object, so the output is not an image, but a three-dimensional grid [Sumin et al., 2019, Nindel et al., 2021]. Or, in Chapter 4, in our translucency measurements, we simulated a monochromatic dataset of translucency appearance using a single-pixel monochrome sensor.

In all of these cases, the rendering is based on mathematical equations and radiometric quantities. The recursive equations contain multi-dimensional integrals, so they are often solved with Monte Carlo integration methods and path-tracing algorithms. Such equations and algorithms are generally suitable for simulating simple materials, including various kinds of diffuse and glossy surfaces. We covered all of these fundamentals in Section 2.1. In Sections 2.2 and 2.3, we focus on extending these equations to govern the light interactions in more complex materials: translucent materials (participating media) and fluorescent materials, which are of our interest in Chapters 4 and 5.

2.1 Fundamental background

Rendering materials is a complex research topic that may feel like a jigsaw puzzle: to truly understand rendering, one must assemble individual pieces of knowledge about different subtopics into a coherent mental picture. Instead of repeating the full theory and derivations found in excellent books, we attempt here to point out the most important pieces of the puzzle and show them in the context of our thesis and the material measurements. For the actual details, we refer the readers to two excellent publications: Pharr et al. [2023, primarily Chapters 2, 4, 5, and 13] for a more tutorial-like introduction, and Veach [1997, mainly Chapters 2 and 3] for a mathematically formal derivation.

2.1.1 Radiometric quantities

The mathematical equations behind rendering describe how light energy is distributed throughout the scene upon interacting with the scene and its materials.




Figure 2.1: Visualization of the fundamental radiometric quantities.

To objectively describe the light energy in rendering, we rely on *radiometric quantities* such as power, irradiance, and, most importantly, radiance (Figure 2.1).

Energy The light energy is denoted by Q and is measured in joules (J). Every single photon carries energy derived from its wavelength λ , the speed of light, and Planck’s constant. If we think of light as a stream of photons, then the total energy would refer to the cumulative energy of all the photons that reached a certain area, e.g., a sensor in a camera, over a certain time.

Power To express the amount of energy *in an instant*, with respect to time t , we use power (flux) $\Phi = Q/t$, measured in joules per second (J/s), a unit equivalent to watts (W). Power depends on the area: for example, sunlight detected by a large sensor will read a larger power than sunlight detected by a microsensor. In rendering, we often normalize the sensor such that the rendered image would have the same brightness no matter the sensor’s resolution or size.

Irradiance To express the density of the radiant power over a given surface area A , we often work with irradiance $E = \Phi/A$, measured in watts per meter squared (W/m^2). Irradiance at a surface position \mathbf{x} with a differential area dA is $E(\mathbf{x}) = d\Phi/dA$. Irradiance depends on the illumination angle, which is known as *Lambert’s law*: a sensor receives lower irradiance with a higher angle θ between the sensor and light source, and the factor is $\cos\theta$ (nicely illustrated in Pharr et al. [2023, Section 4.1.1]). It is also known as the *foreshortening factor*.

Radiance While irradiance measures the energy from all directions, *radiance* is a directionally-dependent quantity. Radiance is defined cleverly to eliminate the foreshortening factor $\cos\theta$, so it does not change with the mutual orientation of the light ray and the surface. For that purpose, we define radiance L as irradiance of a surface that is *perpendicular* to a ray in the given direction and with respect to the solid angle ω in the direction of the ray: $L(\mathbf{x}, \omega) = dE^\perp(\mathbf{x})/d\omega$. The unit is watts per meter squared per steradian ($\text{W}/\text{m}^2/\text{sr}$).

Significance of radiance Following the paragraphs above, radiance is the most fundamental of the quantities, and it is defined in a way that is convenient for at least three reasons. First, the quantity is independent of the $\cos\theta$ factor. Second, radiance is constant along a light ray traveling in a vacuum. Third, all the other quantities can be derived by integrating radiance (see below).

Relations Following the definitions above, flux and irradiance can be expressed by integrating radiance over the sphere \mathcal{S}^2 around a point \mathbf{x} :

$$E(\mathbf{x}) = \int_{\mathcal{S}^2} L(\mathbf{x}, \boldsymbol{\omega}) |\cos \theta| d\boldsymbol{\omega}, \quad (2.1)$$

$$\Phi = \int_A \int_{\mathcal{S}^2} L(\mathbf{x}, \boldsymbol{\omega}) |\cos \theta| d\boldsymbol{\omega} dA. \quad (2.2)$$

The expressions above are written without taking time t and wavelength λ into account. In reality, the variables are time and wavelength-dependent. In the context of our thesis, the time dependence is rather irrelevant, but the wavelength dependence is critical. For example, for a camera sensor of area A detecting light with an exposure time in a time range from t_0 to t_1 , sensitive to wavelengths in a certain range from λ_0 to λ_1 , the total detected energy would be:

$$Q = \int_{t_0}^{t_1} \int_{\lambda_0}^{\lambda_1} \Phi(t, \lambda) d\lambda dt = \int_{t_0}^{t_1} \int_{\lambda_0}^{\lambda_1} \int_A \int_{\mathcal{S}^2} L(\mathbf{x}, \boldsymbol{\omega}, t, \lambda) |\cos \theta| d\boldsymbol{\omega} dA d\lambda dt. \quad (2.3)$$

2.1.2 Measurement equation

The previous equation brings us to the *measurement equation* [Veach, 1997, Pharr et al., 2023], which describes what a sensor measures (Figure 2.2). The measurement equation has no exact definition, and its exact form differs between publications and implementations, but the overall concepts are equivalent.

In a sensor consisting of M pixels, each pixel will measure a certain value I_1, \dots, I_M representing the light intensity arriving at that pixel. Each pixel has a certain sensitivity to incoming light, and it may differ for photons that arrived with wavelength λ , at time t , at position \mathbf{x} , or under an angle $\boldsymbol{\omega}$. For example, a camera sensor typically consists of red, green, and blue pixels, which are each sensitive only to a small range of wavelengths. This is generally encompassed in the pixel *responsivity function* $W_e^{(j)}(\mathbf{x}, \boldsymbol{\omega}, t, \lambda)$, which may differ for each pixel j .

Using the responsivity function, we can then rewrite Eq. (2.3) to measure the j -th pixel's intensity as follows:

$$I_j = \int_t \int_{\lambda} \int_A \int_{\mathcal{S}^2} W_e^{(j)}(\mathbf{x}, \boldsymbol{\omega}, t, \lambda) L_i(\mathbf{x}, \boldsymbol{\omega}, t, \lambda) d\boldsymbol{\omega} dA d\lambda dt, \quad (2.4)$$




Figure 2.2: A sensor may consist of multiple pixels. Each pixel is sensitive to radiance incoming from a certain range of directions and wavelengths, and over a certain period of time. The pixels may have different sensitivities, which is common in color sensors sensitive to red, green, and blue.

where $L_i(\mathbf{x}, \boldsymbol{\omega}, t, \lambda)$ is radiance incoming at the sensor, specifically to point \mathbf{x} , from direction $\boldsymbol{\omega}$, at time t , and with wavelength λ .

For the purposes of rendering, a sensor is not a standalone element; rather, it is often paired with an aperture and lens, and together they form a *camera* (Figure 2.3). The measurement equation may be formally adjusted such that the pixel responsivity functions correspond to the whole camera geometry. We often distinguish between various models such as the *ideal pinhole camera* (an infinitesimally small aperture with no lens) or the *thin-lens camera* (based on the thin lens approximation, capable of simulating focal blur).




Figure 2.3: A simple model of light incoming to a sensor through the lens and aperture of a camera.

In our thesis, we only measure static material samples, so from this point onwards, we will consider radiance to be time-independent, and we will drop the time notation. Time is only important in specific cases, such as rendering dynamic scenes with motion blur or simulating phenomena such as phosphorescence.

2.1.3 Measuring colors

Spectral sensitivities The pixel’s responsivity W_e often depends on wavelength λ . A typical color camera has red, green, and blue pixels in a mosaic grid, and their spectral sensitivities are usually similar to those in Figure 2.4b. When measuring materials, it is often beneficial to use sensors with more spectral channels than just red, green, and blue, such that the material’s spectral behavior can be properly resolved (the shortcomings of RGB approaches are discussed later in Section 3.2.1). For example, the AMS AS7341 multispectral sensor, reviewed in Chapter 5, consists of eight pixels with narrowband sensitivities (Figure 2.4c). An even higher spectral resolution is possible with a spectrometer, whose pixels can have a sub-nanometer resolution, because the incoming light’s spectrum is split using a diffraction grating and projected onto the individual pixels. Nonetheless, we can use the same measurement equation if we treat the spectrometer as a black box, i.e., we are only interested in its inputs and outputs.

Converting a spectrum to a color Often, we want to visualize the measured intensities as a *color*, which is possible using *color spaces* and *color conversions*. Light intensity $L(\lambda)$ can be converted to color using color-matching functions. The CIE XYZ standard defines three color-matching functions $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, $\bar{z}(\lambda)$ (Figure 2.4a) [Carter et al., 2004] for the CIE 1931 standard colorimetric observer,

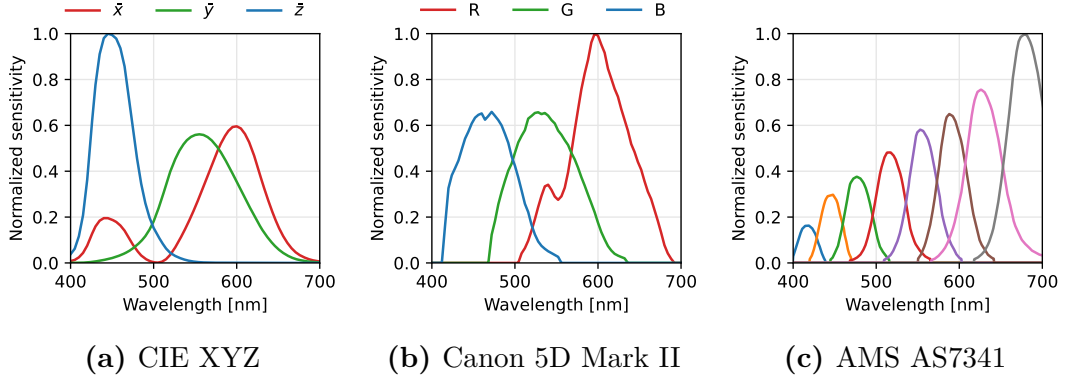


Figure 2.4: Spectral sensitivities of the CIE 1931 standard colorimetric observer based on real human vision (a), a color camera (b) [Kawakami et al., 2013], and a multispectral sensor (c).

based on real human vision. An (x,y,z) -triplet is then computed by integrating the actual incoming light spectrum $L(\lambda)$ with the color matching functions, e.g., $x = \int_{\lambda} L(\lambda)\bar{x}(\lambda)d\lambda$. The (x,y,z) -triplet can then be converted to a different color space, e.g., to the commonly used sRGB, using a 3×3 transformation matrix specific to each color space [Pharr et al., 2023].

Converting sensor measurements to a color The spectral sensitivities of a real sensor typically look different than the CIE XYZ curves (compare Figure 2.4b to 2.4a), or the sensor can have more than three channels (Figure 2.4c). Color conversions are then necessary. A transformation matrix that converts from the sensor’s space to CIE XYZ can be found, e.g., by running an optimization algorithm on a set of known reflectance spectra [Pharr et al., 2023, Section 5.4.2]. Afterward, converting the sensor’s image into sRGB is possible by first converting the sensor measurement into CIE XYZ and then from CIE XYZ to sRGB.

2.1.4 Light transport equation

The measurement equation (Eq. 2.4) formalizes what the sensor’s pixels measure given an incoming radiance L_i that reaches the camera. To compute the actual radiance L_i that comes from light interacting with the scene before it reaches the camera (Figure 2.5), we need to refer to the *light transport equation*, also known as LTE, which is introduced, for example, in Veach [1997, Section 3.7.2] and more extensively described in Pharr et al. [2023, Chapter 13].

Formally, we first need to define the *ray casting function* $r(\mathbf{x}, \boldsymbol{\omega})$ that returns the closest surface intersection for a ray that originated at \mathbf{x} and traveled in the direction $\boldsymbol{\omega}$. Then, because light in a vacuum travels in a straight line:

$$L_i(\mathbf{x}, \boldsymbol{\omega}) = L_o(r(\mathbf{x}, \boldsymbol{\omega}), -\boldsymbol{\omega}). \quad (2.5)$$

With that, we can now write the LTE as below. The LTE states that the radiance that leaves the point \mathbf{x} in the direction $\boldsymbol{\omega}_o$ is composed of two parts: first, the radiance that the object at point \mathbf{x} emits, in case it is a light source; second, the radiance that is reflected from or transmitted through the surface at point \mathbf{x} . The latter radiance is computed by integrating the incoming radiance from all




Figure 2.5: The radiance that reaches the camera sensor may come from anywhere in the scene, and it is common that the light reflects multiple times between the light source and the camera.

directions and weighting it by the bi-directional scattering distribution function (BSDF) $f(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$ (Section 2.1.5):

$$L_o(\mathbf{x}, \boldsymbol{\omega}_o) = \underbrace{L_e(\mathbf{x}, \boldsymbol{\omega}_o)}_{\text{emission}} + \underbrace{\int_{S^2} \underbrace{f(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)}_{\text{BSDF}} \underbrace{L_o(r(\mathbf{x}, \boldsymbol{\omega}_i), -\boldsymbol{\omega}_i) |\cos \theta_i|}_{\text{light incoming from } \boldsymbol{\omega}_i} d\boldsymbol{\omega}_i}_{\text{reflection and/or transmittance}}. \quad (2.6)$$

The LTE is recursive because the term L_o appears on both sides. This makes the equation infinitely-dimensional, because inside the two-dimensional integral, there is an infinite number of nested two-dimensional integrals. Fortunately, the equation can be solved numerically, which we show in Section 2.1.6.

2.1.5 Bi-directional scattering distribution function

The BSDF in Eq. (2.6) is a parameter of the material’s surface, and it is a statistical distribution function describing how light scatters (reflects and refracts) on the surface. Formally, it is the radiance leaving in the direction $\boldsymbol{\omega}_o$ per unit of irradiance incoming from $\boldsymbol{\omega}_i$:

$$f(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) = \frac{dL_o(\boldsymbol{\omega}_o)}{dE(\boldsymbol{\omega}_i)} = \frac{dL_o(\boldsymbol{\omega}_o)}{L_i(\boldsymbol{\omega}_i) |\cos \theta_i| d\boldsymbol{\omega}_i}. \quad (2.7)$$

If the object’s properties are not spatially varying, it is possible to drop the position \mathbf{x} from the notation and simply refer to the material’s BSDF as $f(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$. Further simplifications can be made for isotropic materials.

Different materials generally have different BSDF (Figure 2.6). The LTE and BSDF defined in these sections are enough to describe a large variety of materials, including conductors and dielectrics, or, more specifically, perfect mirrors, glossy, rough, and diffuse surfaces, transparent surfaces with refraction, and so on. The materials are either modeled with exact formulas for $f(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$ (and such formulas can be found, for example, in Pharr et al. [2023, Chapter 9]), or the BSDF can be even interpolated from measurements of real material samples, shown, for example, by Dupuy and Jakob [2018], whose database is publicly available¹.

However, more complex materials, such as participating media or fluorescent materials, cannot be described by these simple equations and require adjustments described later in Sections 2.2 and 2.3.

¹<https://rgl.epfl.ch/materials>





Figure 2.6: An identical shape rendered with four different materials: gold, plastic, diffuse, and glass. Each material scatters the incoming light differently, which is formally described by their BSDF. The diagrams above visualize the light interactions on each material’s surface.

2.1.6 Solving the equations with Monte Carlo

In predictive rendering, we typically rely on Monte Carlo methods to solve the equations above. Without going into formal mathematical definitions and proofs, which are available in [Veach \[1997, Chapter 2\]](#) and [Pharr et al. \[2023, Chapter 2\]](#), we can simply state that the idea behind Monte Carlo integration is to use random sampling to evaluate an integral of the following form:

$$I = \int_{\Omega} f(\mathbf{x})d\mu(\mathbf{x}), \quad (2.8)$$

where Ω is the integration domain and μ is a measure function on Ω , such as the surface area A , solid angle ω , or other measures mentioned earlier in this section.

The estimation then works by choosing an arbitrary probability density function p , with the only requirement being that for all \mathbf{x} for which the integrated function is non-zero, $f(\mathbf{x}) \neq 0$, the density must also be non-zero: $p(\mathbf{x}) \neq 0$. We then draw N independent samples X_1, \dots, X_N according to the density p . Finally, the value of the integral I is estimated as:

$$\hat{I} = \frac{1}{N} \sum_{i=1}^N \frac{f(X_i)}{p(X_i)}. \quad (2.9)$$

More formally, the estimator above is known as the Horvitz-Thompson estimator, and the estimate is a random variable whose properties have been extensively studied. The goal is to sample from such density functions that lead to the lowest variance, meaning only a few random samples are needed for the estimate to be close to the correct result. For that purpose, the commonly used techniques include *importance sampling* [[Pharr et al., 2023, Section 2.2.2](#)], *multiple importance sampling* (MIS) [[Veach, 1997, Kondapaneni et al., 2019](#)], *Russian roulette* [[Arvo and Kirk, 1990, Vorba and Křivánek, 2016, Pharr et al., 2023, Section 2.2.4](#)], and *path guiding* [[Müller et al., 2017, Vorba et al., 2019](#)].

2.1.7 Implementation

Finally, the aforementioned process must be implemented into the predictive renderer. The exact implementations differ as the mathematical equations may be adjusted and re-formulated to suit certain rendering scenarios more than others. The two fundamental approaches are *path tracing* (PT) [Kajiya, 1986] and *light tracing* (LT, also known as *particle tracing*) [Dutr e et al., 1993]. These two approaches together form *bi-directional path tracing* (BDPT) [Lafortune and Willems, 1993, Veach and Guibas, 1995, Veach, 1997]. Other variants include *photon mapping* (PM) [Jensen, 2001]; *vertex connection and merging* (VCM) [Georgiev et al., 2012]; *unified points, beams, and paths* (UPBP) [Křiv nek et al., 2014]; *Metropolis light transport* (MLT) [Veach and Guibas, 1997]; and many more, including modern implementations of the previous techniques.

Path tracing All rendered images in this thesis, including in the inserted publications in Chapters 4 and 5, were computed using path tracing, which is sufficient in most common scenarios. It is a straightforward technique implemented in common rendering software such as Mitsuba 3 [Jakob et al., 2022a], PBRT [Pharr et al., 2023], or ART [Wilkie and Tobler, 2022]. Path tracing follows naturally from the measurement and light transport equations (Eqs. 2.4 and 2.6). Specifically, the algorithm is constructed such that for each pixel of the sensor, Monte Carlo is employed to solve the measurement equation, which in turn depends on the light transport equation, which is solved by applying Monte Carlo recursively. In practice, the algorithm begins by randomly sampling directions originating at the sensor’s pixels and aiming in front of the camera. Following the light transport equation, we know that light travels in straight lines, so the incoming radiance can be traced by finding the closest intersection along the ray and recursively computing the reflected light by sampling the BSDF. A new ray is created at the new point, and this process continues. It is called *path tracing* because by *tracing* the new intersections along the random rays and directions, a *path* is constructed between the camera and a light source.

Light tracing Light tracing, also known as particle tracing, follows virtually the same principle as path tracing, with one critical difference: the sampling does not begin at the sensor’s pixels, but at the illuminants, so the light is traced in reverse. While light tracing resembles the actual physical process more, as photons indeed originate at the illuminants, this approach typically suffers from high variance (noise). However, it is useful in specialized scenarios such as simulating scenes with collimated light, which we relied on while computing the detected scattered light from a collimated beam in Chapter 4, Figure 6.

Sidenote: RGB rendering Instead of tracing the spectrally-varying radiance $L(\mathbf{x}, \boldsymbol{\omega}, \lambda)$, it is also possible to split radiance into three channels: red, green, and blue, essentially tracing the triplet $(L_R(\mathbf{x}, \boldsymbol{\omega}), L_G(\mathbf{x}, \boldsymbol{\omega}), L_B(\mathbf{x}, \boldsymbol{\omega}))$. This approach is significantly more efficient with less noise as it reduces the Monte Carlo dimensionality: the color information is pre-integrated into the scene and its materials, and the color is traced directly in the light paths instead of being reconstructed from individual Monte Carlo-sampled wavelengths. Working in RGB

color spaces is also natural to artists and designers, who usually prepare assets for rendering. Unfortunately, losing the original spectral information has significant shortcomings, introduced in more detail in Section 3.2.1, which makes RGB workflows generally unsuitable for predictive rendering. Furthermore, RGB rendering cannot support more advanced phenomena such as fluorescence. For these reasons, we always aim to measure the optical properties spectrally in this thesis, and we never refer to them in RGB.

2.2 Translucent materials, participating media

The introduction to predictive rendering in the previous section is sufficient for simulating many simple materials. In this section, we focus on more complex *participating media*, which are a special case of *translucent materials* that interact with light within their interior, under the surface. This includes many liquid materials like water, wine, or milk; aerosols such as fog, clouds, or smoke; biological matter like tissue, skin, or blood; and many solids such as plastic, wood, or porcelain. In Chapter 4, we focus on measuring such materials, which we demonstrate on translucent inks used in 3D printers.

The information in this section is based primarily on the following publications that contain more details on participating media: Pharr et al. [2023, Chapters 11 and 14], Novák et al. [2018], and Jarosz [2008, Chapter 4].

Sidenote: Translucent vs. participating We would like to point out that *translucent materials* and *participating media* are two terms that are sometimes used interchangeably. One could argue that *translucent material* is the more general term: for example, frosted glass is a translucent material, although it is not technically a participating medium: the light interacts only with its surface, so it is sufficient to describe the scattering with BSDF. On the other hand, not every participating medium is translucent; if the absorption coefficient is sufficiently high, the material will be fully opaque. The goal of this thesis is not to settle this debate, and since we focus on *materials*, we find the term *translucent materials* perhaps more understandable and appropriate than *participating media*.

2.2.1 Light interactions

We can no longer assume that light only interacts with surfaces and travels in straight lines without losing energy until it hits another surface. In participating media, light is partially absorbed and scattered along its whole travel throughout the material (Figure 2.7). The material is a collection of microscopic particles participating in the light transport, hence *participating media* (Figure 2.8). The individual material parameters significantly influence appearance (Figure 2.9).

Assumptions To build a reasonable mathematical model, we commonly make the following assumptions. We model a participating medium as a collection of microscopic, randomly positioned, and randomly oriented particles. For simplicity, we assume that the particle distribution is uncorrelated, although models for correlated media also exist [Bitterli et al., 2018].




Figure 2.7: A visualization of light being scattered through a participating medium enclosed in a glass container.

Boundary Many participating media, especially solids and liquids, have a well-defined boundary where the material starts and ends. For example, a glass of water has an exterior air-to-glass boundary and an interior glass-to-water boundary. At this boundary, the light is partially *reflected* and partially *transmitted* or *refracted*, which is described by the BSDF (Section 2.1.5), and it depends mainly on the *index of refraction* n . For our intentions and purposes, the BSDF is *independent* on the interior of the medium, meaning what the light does on the surface is one concept (covered by the previous section), and what happens on the inside is a separate concept (covered in this section).

Absorption As light travels through a medium, the photons have a certain probability of being absorbed by interacting with the medium particles. The probability density of the absorption event is given by the *absorption coefficient* $\sigma_a(\mathbf{x}, \boldsymbol{\omega}, \lambda)$, which may generally depend on the position \mathbf{x} inside the medium, the incoming light direction $\boldsymbol{\omega}$, and the wavelength λ . For simplicity, we often just write $\sigma_a(\mathbf{x})$. The unit is a reciprocal meter (m^{-1}), and statistically, we can think of the inverse value $1/\sigma_a$ as the mean free distance (in meters) that an average photon can travel until absorbed.

Scattering Similarly to absorption, the photons have a certain probability of being scattered (i.e., changing their direction) by interacting with the medium particles. The probability density of the scattering event is given by the *scattering coefficient* $\sigma_s(\mathbf{x}, \boldsymbol{\omega}, \lambda)$ with the same unit and principle as the absorption

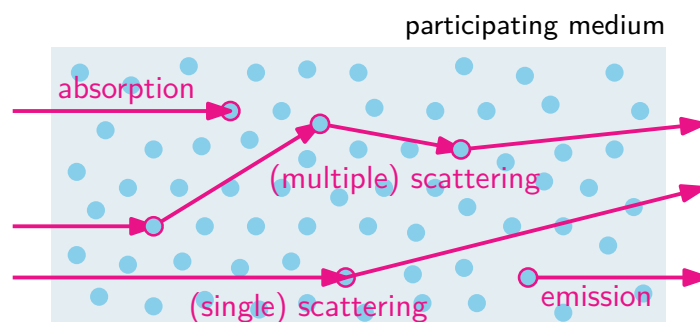




Figure 2.8: A diagram of the light interactions in a participating medium.

coefficient. In computer graphics, we almost exclusively assume that scattering is *elastic* for the visible wavelength range, meaning the photon is not losing any energy by being scattered, and its wavelength remains constant. In other fields, such as radiography working with X-rays, we also encounter *inelastic* scattering, such as Compton scattering, when the photon’s energy decreases and its wavelength changes [Mery, 2015, Choppin et al., 2013].


Emission Although not used in this thesis, one can also simulate an *emissive medium*, where light energy is created from within the medium itself. This allows for simulating media such as fire or explosions.



(a) The identical participating media with different boundaries with an increasing index of refraction n , from 1.00 (left) to 1.33, 1.50, and 2.40 (right).



(b) The effect of the single-scattering albedo $\alpha = \sigma_s / (\sigma_a + \sigma_s)$, ranging from 0.95 (left), to 0.80, 0.50, and 0.10 (right). The refractive index is $n = 1.00$.



(c) The effect of the extinction coefficient $\sigma_t = \sigma_a + \sigma_s$, which is the highest in the left image, and is $10\times$, $50\times$, and $250\times$ lower in the other images. The refractive index is $n = 1.50$.

Figure 2.9: Practical visualization of how the individual parameters (index of refraction, absorption coefficient, and scattering coefficient) influence the actual appearance of an object. The object’s shape and its surrounding are identical.

2.2.2 Transmittance

The energy of a light beam traveling through a medium decreases with the traveled distance due to the combined effect of absorption and scattering.

Extinction coefficient and albedo The absorption and scattering can be aggregated into a single quantity called the *extinction coefficient* $\sigma_t = \sigma_a + \sigma_s$. The ratio between the scattering events and all the events is called the *single-scattering albedo* $\alpha = \sigma_s/\sigma_t$.

Beer's law Following the extinction coefficient, one can derive the energy loss for a light beam traveling inside the medium between two points \mathbf{x} and $\mathbf{x}' = \mathbf{x} + t\boldsymbol{\omega}$ that are the distance t apart in the direction $\boldsymbol{\omega}$. According to Beer's law [Jarosz, 2008], the energy loss is exponential, and *transmittance* T defines the fraction of radiance that is transmitted along the beam between the two points (Figure 2.10):

$$T(\mathbf{x}, \mathbf{x}', \lambda) = \exp\left(-\int_0^t \sigma_t(\mathbf{x} + t\boldsymbol{\omega}, \boldsymbol{\omega}, \lambda) dt\right). \quad (2.10)$$

It is important to remember that transmittance only accounts for the *loss of energy* in a collimated beam. In a scattering medium, the beam will also *gain* certain energy from light that is scattered from outside *towards the beam*. To differentiate between light scattered out compared to towards the beam, we often use the terms *out-scattering* and *in-scattering*. The distinction will become important later in the volume rendering equation (Section 2.2.4).




Figure 2.10: The energy of a collimated pencil of light decreases as it travels through a medium.

2.2.3 Phase function

The final building block for deriving the modified light transport equation for participating media is the *phase function*, which is, very loosely speaking, an alternative to BSDF, except it does not describe scattering on a surface but inside a medium. When light scatters in the medium, the light is distributed in other directions, and the phase function describes the angular distribution of the scattering between the incoming $\boldsymbol{\omega}_i$ and outgoing $\boldsymbol{\omega}_o$ angles: $p(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$.


In general, we can denote the phase function $p(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o, \lambda)$, as it depends also on the position in the medium \mathbf{x} and the wavelength λ . The phase function always needs to be normalized, i.e.:

$$\forall \mathbf{x} \forall \boldsymbol{\omega}_i \forall \lambda : \int_{S^2} p(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o, \lambda) d\boldsymbol{\omega}_o = 1. \quad (2.11)$$


Asymmetry or anisotropy factor Statistically, for each phase function, we can compute the *average cosine* of the scattering angle, denoted $g \in [-1, 1]$:

$$g = \int_{S^2} p(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o) \cos \theta d\boldsymbol{\omega}_o, \quad (2.12)$$

where θ is the angle between the two directions $\boldsymbol{\omega}_i$ and $\boldsymbol{\omega}_o$. The quantity g has many names: apart from the scattering *average cosine*, it is also called the *asymmetry parameter* [Pharr et al., 2023, Section 11.3.1] or the *anisotropy factor* [Pickering et al., 1993]. We prefer the latter because an isotropic phase function has $g = 0$, whereas anisotropic phase functions have $g < 0$ or $g > 0$ for dominantly backward and forward scattering, respectively (Figure 2.11).



(a) Illumination facing the camera, $g = 0.6$ (left image), 0.3, 0.0, and -0.6 (right).



(b) Illumination behind the camera, $g = 0.6$ (left image), 0.3, 0.0, and -0.6 (right).

Figure 2.11: The effect of varying anisotropy factors g in different illumination conditions. When looking against the light source (a), forward scattering appears brighter than backward scattering. When looking from the light source position (b), backward scattering appears brighter as it reflects back toward the camera.

Isotropic phase function There exists exactly one isotropic phase function, i.e., one that scatters light in all directions uniformly, and its value is:

$$p(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o) = \frac{1}{4\pi}, \quad (2.13)$$

where the value $1/4\pi$ is a factor that ensures normalization. As already stated, its anisotropy factor is $g = 0$.

One-dimensional phase functions Whereas the isotropic phase function is technically zero-dimensional (it is a constant), the phase function's angular dependency can generally be four-dimensional (it has two two-dimensional directions). A common simplification is to reduce the dimensionality such that the angular

dependency is only one-dimensional. Instead of the phase function depending on both directions $p(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$, we write $p(\cos \theta)$, and the phase function then only depends on the angle θ between the two directions.

Henyey-Greenstein phase functions A special example of a one-dimensional phase function was introduced by [Henyey and Greenstein \[1941\]](#). It is not one phase function but rather a family of phase functions containing one function per each anisotropy factor g . They are defined by the following formula:

$$p(\cos \theta) = \frac{1}{4\pi} \frac{1 - g^2}{(1 + g^2 - 2g \cos \theta)^{3/2}}. \quad (2.14)$$

Because of the simplicity and universality, the Henyey-Greenstein (HG) phase functions are commonly used to approximate more complicated real phase functions. In [Chapter 4](#), our measurement method is based on the Henyey-Greenstein phase function, and it can estimate the anisotropy factor $g(\lambda)$ per wavelength λ .

2.2.4 Volume rendering equation

Using the building blocks from above, we now modify the light transport equation from [Section 2.1.4](#) to support the absorption, scattering, and emission of participating media. The resulting equation is also known as the *volume rendering equation* (VRE), and it is derived from the *radiative transfer equation* (RTE) [[Chandrasekhar, 1960](#)].

We begin with the RTE, which describes what happens to radiance $L(\mathbf{x}, \boldsymbol{\omega})$ in a light beam traveling through a participating medium ([Figure 2.12](#)). For a differential distance dz along the beam, the differential change in radiance can be written as:

$$\begin{aligned} \frac{dL(\mathbf{x}, \boldsymbol{\omega})}{dz} = & - \underbrace{\sigma_a(\mathbf{x}, \boldsymbol{\omega})L(\mathbf{x}, \boldsymbol{\omega})}_{\text{absorption}} - \underbrace{\sigma_s(\mathbf{x}, \boldsymbol{\omega})L(\mathbf{x}, \boldsymbol{\omega})}_{\text{out-scattering}} \\ & + \underbrace{\sigma_a(\mathbf{x}, \boldsymbol{\omega})L_e(\mathbf{x}, \boldsymbol{\omega})}_{\text{emission}} + \underbrace{\sigma_s(\mathbf{x}, \boldsymbol{\omega})L_s(\mathbf{x}, \boldsymbol{\omega})}_{\text{in-scattering}}, \end{aligned} \quad (2.15)$$

where L_s denotes light coming from all other directions towards the beam:

$$L_s(\mathbf{x}, \boldsymbol{\omega}) = \int_{S^2} p(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}) L_i(\mathbf{x}, \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i. \quad (2.16)$$

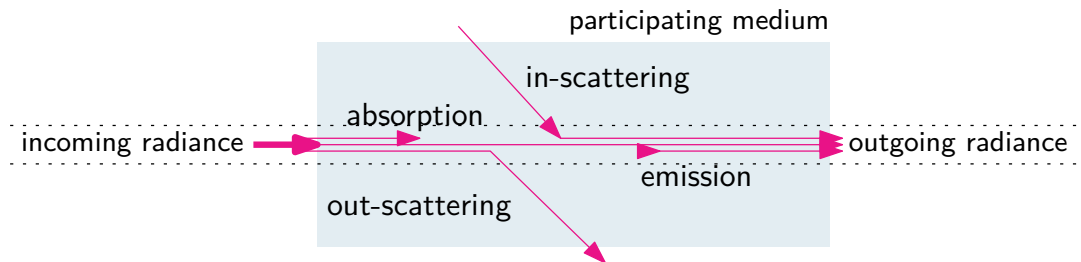


Figure 2.12: For a collimated pencil of light entering the medium, the outgoing radiance in the beam changes due to absorption, in-scattering, out-scattering, and emission.

The individual terms of the RTE, in the order of appearance, describe the amount of radiance that is: lost due to absorption; lost due to light being scattered from the beam towards outside; gained by emission of the medium; and gained by all the light that is in-scattered from all other directions towards the beam.

The final VRE, which can be used as a direct replacement of LTE (Eq. 2.6) for rendering participating media, is obtained by integrating the RTE (Eq. 2.15) along a line segment between an initial point \mathbf{x}_0 and another point \mathbf{x}_z , which is on a surface in the distance z (Figure 2.13):

$$L_o(\mathbf{x}_0, \boldsymbol{\omega}) = \tag{2.17}$$

$$\underbrace{\int_0^z T(\mathbf{x}_0, \mathbf{x}_y) [\sigma_a(\mathbf{x}_y, \boldsymbol{\omega})L_e(\mathbf{x}_y, \boldsymbol{\omega}) + \sigma_s(\mathbf{x}_y, \boldsymbol{\omega})L_s(\mathbf{x}_y, \boldsymbol{\omega})] dy}_{\text{total contribution from the medium}} \tag{2.18}$$

$$+ \underbrace{T(\mathbf{x}_0, \mathbf{x}_z)L_o(\mathbf{x}_z, \boldsymbol{\omega})}_{\text{contribution from the surface}} . \tag{2.19}$$

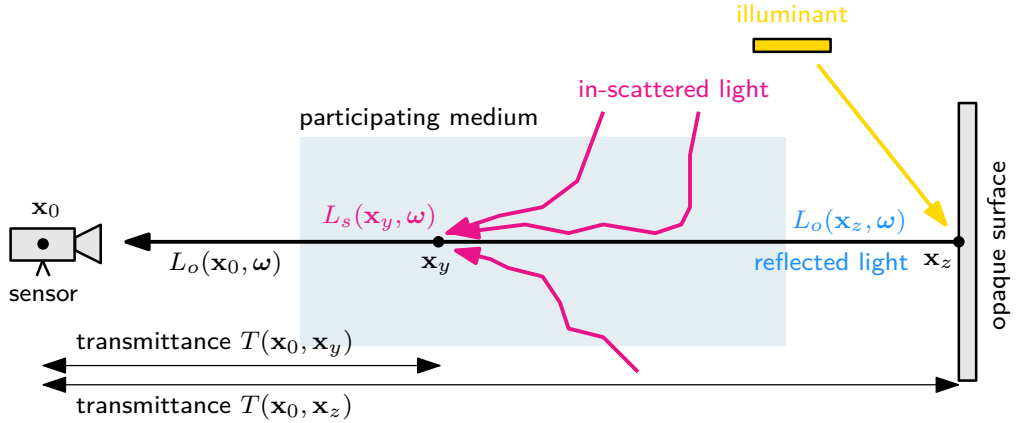


Figure 2.13: The volume rendering equation describes the light transport in a scene that contains not only surfaces but also participating media.

2.3 Fluorescent materials

This section extends the light transport theory to support *fluorescent materials*. We introduce their properties and a modified light transport equation, which will provide the baseline for our measurement method in Chapter 5.

The information in this section is based primarily on the following publications that contain more details on rendering fluorescent materials: Hullin et al. [2010] and the more recent extension by Mojzík et al. [2018].

2.3.1 Light interactions

When a *non-fluorescent* material is illuminated, a fraction of the incoming photons is reflected or refracted according to the BSDF (Section 2.1.5). While the scattering properties are wavelength-dependent, the photons never change their wavelength during the interactions, so the wavelength along the whole light path is always constant. In this section, the assumption is no longer valid. We focus

on materials that absorb and re-emit a fraction of the incoming photons at a *different wavelength* than the photons originally had. In other words, the photons change their wavelength during the interactions, throughout the light path, even multiple times. This is called *wavelength shifting*, and it can be caused by multiple phenomena, including *Compton scattering*, *Raman scattering*, *fluorescence*, and *phosphorescence*.

Fluorescence This thesis focuses specifically on *fluorescence* (Figure 2.14). It is a prominent phenomenon in the visible spectrum, common in natural (minerals, corals, tissues) and synthetic (optical brighteners, pigments in papers, textiles, plastics) objects. A fluorescent material absorbs an incoming photon, which excites the electrons for a certain period of time (in the order of 10^{-8} s), and upon relaxing back to the ground state, a new photon is emitted with lower energy (longer wavelength). Fluorescence is usually assumed to be mostly diffuse, i.e., almost independent on the angle [Hullin et al., 2010, Section 5], because the new photons are emitted in random directions without remembering under what angle the original photon arrived from.

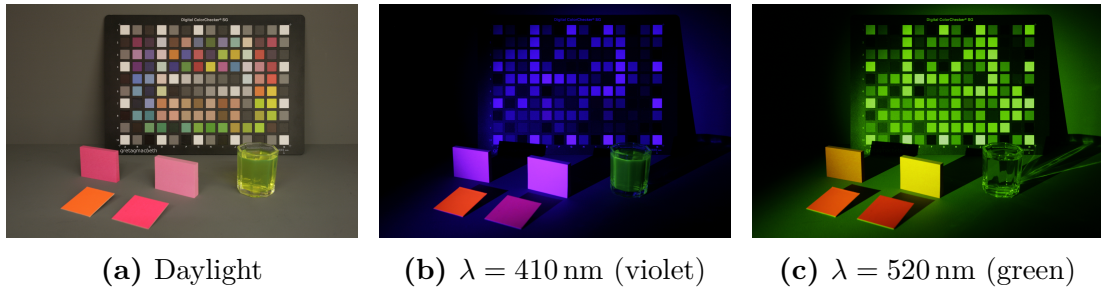


Figure 2.14: Photographs of fluorescent (sticky notes, liquid detergent) and non-fluorescent (color chart) materials under different illuminants. Under daylight, everything appears colorful. The monochrome illuminants reveal that the sticky notes and the liquid are fluorescent materials emitting different wavelengths than they received from the illuminants.

2.3.2 Bi-spectral light transport equation

To accommodate that photons can change their wavelength, we must adjust the original light transport equation (Eq. 2.6) accordingly. First, we must formally include the incoming and outgoing wavelengths λ_i and λ_o into the notation. Second, we must add an integral over all *incoming* wavelengths because the outgoing radiance may have a different wavelength than the incoming radiance. The equation then becomes [Hullin et al., 2010]:

$$L_o(\mathbf{x}, \boldsymbol{\omega}_o, \lambda_o) = L_e(\mathbf{x}, \boldsymbol{\omega}_o, \lambda_o) + \int_{S^2} \int_{\lambda} \underbrace{f(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o, \lambda_i, \lambda_o)}_{\text{bi-spectral BSDF}} \underbrace{L_o(r(\mathbf{x}, \boldsymbol{\omega}_i), -\boldsymbol{\omega}_i, \lambda_i)}_{\text{light incoming from } \boldsymbol{\omega}_i \text{ and } \lambda_i} |\cos \theta_i| d\lambda_i d\boldsymbol{\omega}_i. \quad (2.20)$$

The equation now contains a bi-spectral BSDF. If we limit ourselves to reflections only, it is also known as the BRRDF (bi-directional reflectance and

reradiation distribution function). Formally, we can adapt the standard BSDF formula (Section 2.1.5) and re-write it for a bi-spectral case:

$$f(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o, \lambda_i, \lambda_o) = \frac{d^2 L_o(\boldsymbol{\omega}_o, \lambda_o)}{L_i(\boldsymbol{\omega}_i, \lambda_i) |\cos \theta_i| d\lambda_i d\boldsymbol{\omega}_i}. \quad (2.21)$$

Implementing the equation into a renderer is a relatively straightforward modification, yet rendering fluorescence efficiently is still a challenging problem. Full implementation was described by Mojzík et al. [2018] and is available in the open-source ART renderer [Wilkie and Tobler, 2022].

Sidenote: Fluorescent participating media Nothing prevents combining fluorescence and participating media and merging the concept into a single *bi-spectral volume rendering equation*, which was introduced and implemented by Mojzík et al. [2018, Section 5] and is currently available in ART. In fact, translucent fluorescent materials exist in the real world, from rare specimens such as uranium glass to commonly used plastics with fluorescent pigments. The main problem seems to be measuring the optical properties of such complex materials, which has not been shown in computer graphics yet, at least to our knowledge. We believe that such measurements are demanding but achievable by combining our approaches in Chapters 4 and 5, but researching the exact methodology still remains a challenging future work.

2.3.3 Simplified BRRDF

The bi-spectral BSDF or BRRDF are high-dimensional, making them challenging to work with. First, representing fluorescence in memory is non-trivial: as the high-dimensional structure is demanding on storage, specialized compression techniques were invented for the wavelength shifting itself [Hua et al., 2022]. Second, while acquiring the full BRRDF is possible with bi-spectral goniophotometer setups such as the one from Hullin et al. [2010], that is far from affordable. Fortunately, the general solution can be simplified in a few ways.

Simple separability If we could split the BRRDF into independent wavelength shifting component $f^\lambda(\lambda_i, \lambda_o)$ and the angular component $f^\omega(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$, i.e.:

$$f(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o, \lambda_i, \lambda_o) = f^\lambda(\lambda_i, \lambda_o) f^\omega(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o), \quad (2.22)$$

we could not only save memory, but also acquire the material properties in an easier way: once for the spectral component, and once for the angular component. Real fluorescent materials are usually not separable in this way, because they contain both the diffuse fluorescence and specular highlights, which are not fluorescent. In other words, a real fluorescent material can be thought of as having at least two layers: the outer surface layer is like any other surface of a non-fluorescent material, it can be glossy, but underneath, the fluorescent layer is diffuse [Wilkie et al., 2006].

Series of separable terms Hullin et al. [2010, Equation (7)] proposes that if the material cannot be separated as in Eq. (2.22), then we can instead write it as a summed contribution of individual separable terms:

$$f(\omega_i, \omega_o, \lambda_i, \lambda_o) = \sum_n f_n^\lambda(\lambda_i, \lambda_o) f_n^\omega(\omega_i, \omega_o), \quad (2.23)$$

where some of the parts can be diffuse and some specular (Figure 2.15), which becomes applicable to a wide range of real materials.




Figure 2.15: The reflectance of a fluorescent surface can be modeled as a sum of individual contributions. The contributions can be, for example, a glossy reflection term with no fluorescence, and a separate diffuse fluorescent term.

Sidenote: Layered BSDFs In the previous paragraphs, we mentioned that fluorescent materials are more like layered materials, with glossy outer and diffuse inner layers. Layered models are a concept that is not limited to fluorescence: layered materials are commonly implemented in non-fluorescent renderers as well. There is a difference between simulating a layered BSDF model as described by Pharr et al. [2023, Section 14.3], and the blending model in Eq. (2.23), which is more in the realm of Wilkie et al. [2006] and Hullin et al. [2010]. In the former layered model, we simulate the light scattering between the two layers, which are assumed to be a certain distance apart (Figure 2.16, middle image). In the latter model, in Eq. (2.23), the layers are assumed to be at the same position, only blended together, which is a simplified approach (Figure 2.16, right image).




Figure 2.16: Regular reflection occurs on the surface of a material. In the layered model, fluorescence occurs under the surface in a certain depth in which the molecule or particle actually re-emitted the photon. In the simplified model, both the reflection and fluorescence are modeled to happen on the surface at the same position.

2.3.4 Donaldson matrix

The important problem that is unique to fluorescent materials is how to measure the spectral component $f^\lambda(\lambda_i, \lambda_o)$, which we call the *Donaldson matrix* and also denote $\Phi(\lambda_i, \lambda_o)$. Our method in Chapter 5 is devoted exactly to measuring this matrix, which is the critical ingredient to the measurements of fluorescent materials. We now assume that the position \mathbf{x} and the illumination and observation angles $\boldsymbol{\omega}_i, \boldsymbol{\omega}_o$ are fixed, such that we can focus purely on the spectral behavior. When a fluorescent material is illuminated by a light source with spectrum $l(\lambda_i) = L_i(\mathbf{x}, \boldsymbol{\omega}_i, \lambda_i)$, then the reflected intensity $r(\lambda_o) = L_o(\mathbf{x}, \boldsymbol{\omega}_o, \lambda_o)$ is:

$$r(\lambda_o) = \int \Phi(\lambda_i, \lambda_o) l(\lambda_i) d\lambda_i, \quad (2.24)$$

where $\Phi(\lambda_i, \lambda_o)$ jointly describes the material's reflectance and fluorescence, and λ_i and λ_o are incoming (excitation) and outgoing (reflected, emitted) wavelengths, respectively. For practical purposes, the whole equation is often discretized into N_i incoming and N_o outgoing wavelengths, such that $l \in \mathbb{R}^{N_i}$ and $r \in \mathbb{R}^{N_o}$ become vectors, $\Phi \in \mathbb{R}^{N_o \times N_i}$ becomes a matrix, and Eq. (2.24) becomes a matrix-vector multiplication visualized in Figure 2.17:

$$r = \Phi \cdot l. \quad (2.25)$$

The matrix Φ is called the *Donaldson matrix* after Donaldson [1954], who measured the matrix values directly using monochromatic light. The matrix implicitly exists also for non-fluorescent materials, where it is zero everywhere except for the diagonal ($\lambda_i = \lambda_o$), representing the pure reflectance. As Figure 2.17b shows, it is natural to separate the matrix $\Phi = \Phi_{\text{diag}} + \Phi_{\text{fluo}}$ into its diagonal Φ_{diag} (*pure reflectance*) and the off-diagonal elements Φ_{fluo} (*pure fluorescence*).

Sometimes, assumptions are made to simplify the shape of Φ [Donaldson, 1954, Blasinski et al., 2020, Hua et al., 2022]. First, we expect the emitted photons to have longer wavelengths than the excitation, hence $\Phi(\lambda_i, \lambda_o) = 0$ for $\lambda_i > \lambda_o$, or equivalently, the matrix is zero below the diagonal. Second, by following Kasha's rule, Φ_{fluo} would be a separable function $\Phi_{\text{fluo}}(\lambda_i, \lambda_o) = \phi_{\text{em}}(\lambda_i) \cdot \phi_{\text{ex}}(\lambda_o)$ with one-dimensional emission $\phi_{\text{em}}(\lambda_i)$ and excitation $\phi_{\text{ex}}(\lambda_o)$ spectra [Blasinski et al., 2020]. Unfortunately, it does not generally hold, especially for materials with multiple fluorophores that we also work with later in Chapter 5.




Figure 2.17: Illustration of Eqs. (2.24) and (2.25). The reflected spectrum (a) is computed from the Donaldson matrix (b) and illuminant spectrum (c). While the illuminant is green, the fluorescent reflection is orange and has two peaks.

3 | Measuring materials as an inverse problem

In Chapter 2, we covered how the appearance of various materials can be predicted or rendered, given we know the material’s optical parameters beforehand. This chapter is devoted to the *inverse problem*: how to measure unknown optical parameters of a material, given we know how the material looks in specific conditions. The problem is not hypothetical, in fact, it is exactly how various optical properties are measured, and it is how our methods in Chapters 4 and 5 work.

3.1 Problem statement

Forward problem The forward problem is the rendering itself: we know all the parameters of the scene, and we predict what the sensor or the camera in the scene would detect. We can denote the rendering function by $I(\boldsymbol{\pi})$, where $\boldsymbol{\pi}$ are the scene’s parameters. In this chapter, we will use the symbol $\boldsymbol{\pi}$ only for those parameters that are of our interest, for example the material parameters that we want to measure. The rendering function will implicitly assume all the other parameters that are our prior knowledge about the scene, including the object shapes, positions, spectra, and properties of the illuminants and the sensor.

Inverse problem The inverse problem is when we do not know the values of the parameters $\boldsymbol{\pi}$, and we want to infer them from a known reference image or sensor measurement I_{ref} . For example, we can use a real camera to take the picture I_{ref} of the material in specific conditions, and we now want to find what material parameters $\boldsymbol{\pi}$ correspond to that picture I_{ref} . In other words, we must find the parameters $\boldsymbol{\pi}$ for which the renderer generates a similar picture $I(\boldsymbol{\pi}) \approx I_{\text{ref}}$. Formally, this is an *optimization problem* of the following form:

$$\hat{\boldsymbol{\pi}} = \underset{\boldsymbol{\pi}}{\operatorname{argmin}} e(I_{\text{ref}}, I(\boldsymbol{\pi})), \quad (3.1)$$

where $\hat{\boldsymbol{\pi}}$ are the estimated parameters and $e(\cdot, \cdot)$ is the *objective function* (also known as *cost function* or *error metric*) that computes how much the two inputs differ from each other.

The objective function is typically chosen such that its lowest value is zero, and it only returns zero for identical inputs. Often, we use the ℓ^1 or ℓ^2 norm. If $\hat{\boldsymbol{\pi}}$ match the reference, then $I(\hat{\boldsymbol{\pi}}) = I_{\text{ref}}$, so $e(I_{\text{ref}}, I(\hat{\boldsymbol{\pi}})) = 0$. In real conditions, the predictions will not be completely identical to the real sensor, but it is still assumed that the error will be the lowest. Unfortunately, the reverse implication is not true, i.e., it does not hold that when $I(\hat{\boldsymbol{\pi}}) = I_{\text{ref}}$, then it must hold that $\hat{\boldsymbol{\pi}}$ are correctly estimated. As we will see in Section 3.2, this is caused by ambiguities in the parameter space.

Trivial example We demonstrate the inverse problem on a trivial example in Figure 3.1. Suppose we have a plastic rabbit model and a reference photograph I_{ref} captured in a controlled environment. We aim to measure how red the model is,

so our only parameter π is the rabbit’s *redness* $\pi \in [0, 1]$. We can render the exact scene with a varying redness π using a Monte Carlo renderer $I(\pi)$. As the objective function, we use the ℓ^2 difference between the pixels in the reference image and the render:

$$e(I_{\text{ref}}, I(\pi)) = \sqrt{\sum_j \left(I_{\text{ref}}^{(j)} - I(\pi)^{(j)} \right)^2}, \quad (3.2)$$

where $^{(j)}$ denotes the j -th pixel. In Figure 3.1, we can see the plot of the objective function with respect to the redness parameter π , and we can see that its global minimum corresponds to $\pi = 0.4$, which is indeed the correct parameter of the reference image I_{ref} .

Of course, in real scenarios, the situation is usually way more complicated: there are more parameters, so the global optimum cannot be found so trivially; the controlled environment is never perfect and introduces additional errors; and there might be ambiguities that make it impossible to find the real parameters. Our goal is to make the problem as simple yet unambiguous as possible: minimize the number of free parameters, make the controlled environment affordable and easy to replicate in a renderer, and remove any ambiguities.

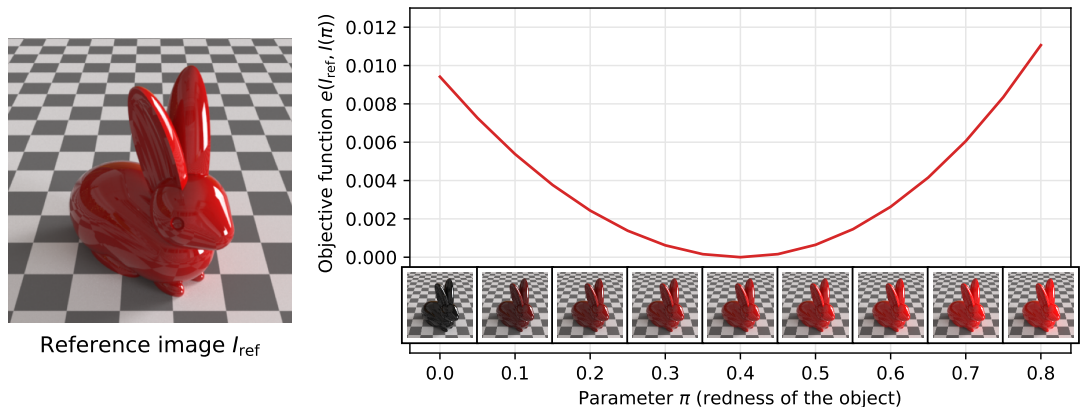


Figure 3.1: Demonstration of a trivial inverse problem. The parameter π of the reference image (left) is found by rendering the same scene with varying π and then picking the one whose image is the most similar to the reference.

3.2 Ambiguities and similarity relations

An ambiguity occurs when the optimization problem has at least two global minima $\pi_1 \neq \pi_2$, such that $I(\pi_1) = I(\pi_2) = I_{\text{ref}}$. In that situation, it becomes impossible to know which parameters correspond to the real material that was captured in the reference I_{ref} . In general, the existence of ambiguities means that the chosen methodology simply cannot distinguish between certain parameters. Apart from completely redesigning the whole scene, it can also be solved by introducing more information into the inverse problem, such as by choosing a fixed value for a subset of the parameters, or by optimizing more references $I_{\text{ref}}^{(1)}, I_{\text{ref}}^{(2)}, \dots$ taken under different conditions. In the rest of this section, we introduce examples of ambiguities that are relevant in the context of our thesis.

3.2.1 Spectral ambiguities, metamerism, and other shortcomings of RGB optical properties

A typical color camera does not capture the full spectrum of the scene. Instead, it captures a photograph whose individual pixels have their *colors* stored in the form of RGB triplets (Section 2.1.3). Working with RGB values became natural and has many benefits: a significantly smaller memory footprint, intuitive color editing for artists, or simpler rendering algorithms (Section 2.1.7). Furthermore, measuring material optical properties in RGB is common in computer graphics, as it allows acquiring the material using a standard color camera. Unfortunately, replacing full spectra with RGB values has critical shortcomings and serious consequences, which we introduce in the following paragraphs. Note that in our measurement methods in Chapters 4 and 5, we always reconstruct the full spectra and never rely only on RGB triplets.

Metamerism and illuminant dependence For any non-trivial RGB color value, there exists an infinite number of spectra $I(\lambda)$ that all correspond to the identical RGB value. This phenomenon is called *metamerism* [Wyszecki and Stiles, 1982]. One of the consequences is that two different materials with spectral reflectances $R_1 \neq R_2$, or transmittances $T_1 \neq T_2$, can both result in an identical *rgb*-triplet under a given illumination spectrum $L_1(\lambda)$, but they can have mismatched appearance under a different illuminant $L_2 \neq L_1$ (Figure 3.2). For example, in the red channel with a spectral sensitivity $\bar{r}(\lambda)$, it can hold:

$$\int \bar{r}(\lambda)R_1(\lambda)L_1(\lambda)d\lambda = \int \bar{r}(\lambda)R_2(\lambda)L_1(\lambda)d\lambda, \text{ although} \quad (3.3)$$

$$\int \bar{r}(\lambda)R_1(\lambda)L_2(\lambda)d\lambda \neq \int \bar{r}(\lambda)R_2(\lambda)L_2(\lambda)d\lambda. \quad (3.4)$$

Consequently, when we capture the material under one illuminant, the RGB optical properties are *fixed* to the illuminant and cannot be used to render the material’s appearance under a different illuminant. In Figure 3.2, we demonstrate this issue on two rabbit models observed under three illuminants. If we measured the rabbits with a color camera under D65, we would think that both rabbits are made of the same material. However, this would conceal the fact that the materials have, in fact, different reflectance spectra, and they do not appear identical under illuminants other than D65.

Spectral upsampling Attempting to reconstruct the actual spectra back from RGB values is called *spectral upsampling*, which is an important research topic with many recent publications [Jendersie, 2021, Jakob and Hanika, 2019, Jung et al., 2019]. As we showed above, the original spectrum cannot be found from only an RGB capture; however, estimates can be made by constraining the problem. For example, it is often assumed that a reflectance spectrum has to be smooth and cannot be spiky (however, this assumption is not true for illuminant spectra, which can be and often are spiky). One can also measure a large dataset of real reflectance spectra, decompose them into a set of basis functions, and then assume that the unknown reflectance spectrum is a linear combination of the basis functions. In Chapter 5, we constrain the reflectance spectra by assuming smooth curves represented by a small number of trigonometric moments, and then we recover the spectra using bounded MESE [Peters et al., 2019].

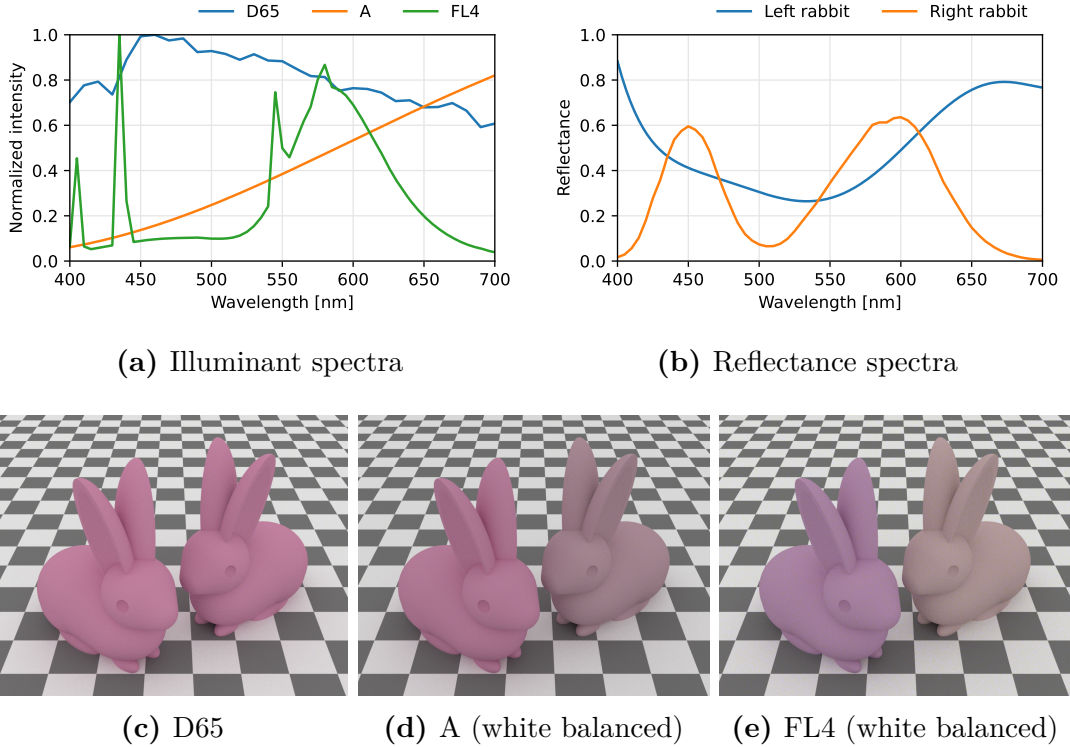


Figure 3.2: Two rabbit models observed under three different illuminants defined in the CIE standard: D65 (standard daylight), A (incandescent bulb), and FL4 (fluorescent lighting). While the materials look identical under D65, observing the rabbits under A and FL4 reveals that their reflectance spectra actually differ. The true illuminant and reflectance spectra are visible in (a), (b).

Color gamut and energy conservation Every RGB color space can represent a particular subset (gamut) of all possible colors. One of the issues with using common color spaces for measurements is that reflectances that are out of the gamut need to be represented with negative values or values above unity. Such values are not physically meaningful and break energy conservation by either creating or negating energy. Mapping these values back to the gamut, e.g., by clipping them to the permissible range, introduces inaccurate color shifts.

Transmission color shifts A problem that is specific to translucent materials and participating media stems from Beer’s law (Section 2.2.2, Eq. 2.10), which says that light transmittance through a translucent medium has an exponential dependency on the distance ℓ traveled through the material. Because this non-linear equation affects each wavelength independently, it causes color shifting as light propagates further through the medium with increasing ℓ – an effect that we call *spectral sharpening*. In Figure 3.3, we compare color predictions based on Eq. (2.10) applied to RGB extinction coefficients, in contrast to actual spectral $\sigma_t(\lambda)$. The results are visualized with their corresponding colored squares. The initial color is marked with a black outline, and the colors on the left and right correspond to shorter and longer distances ℓ , respectively. We can clearly see that the RGB and spectral predictions differ.

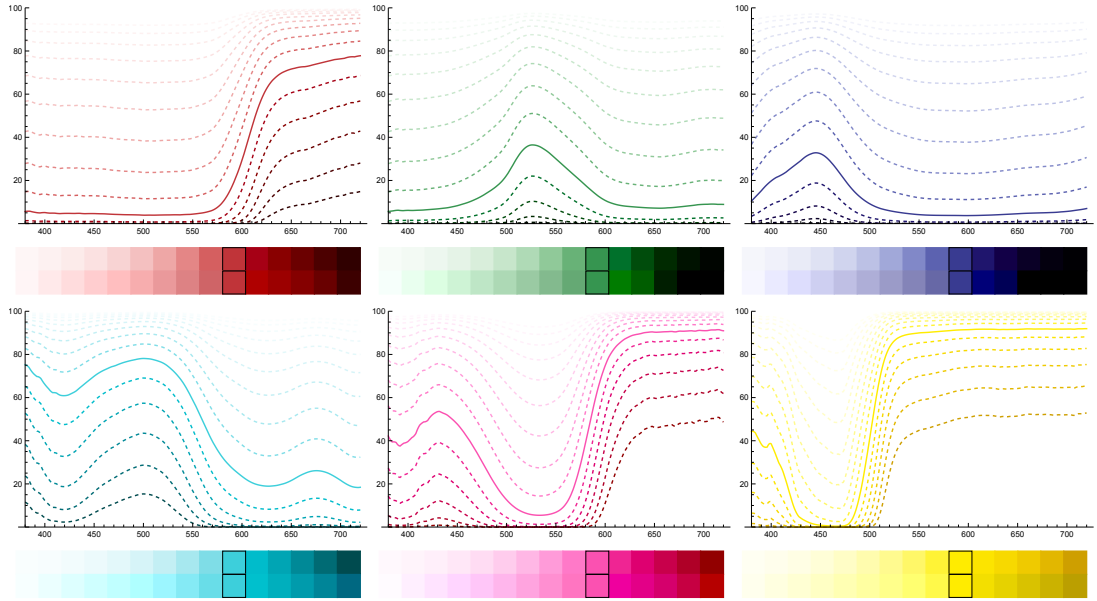


Figure 3.3: A comparison of applying Beer’s law, Eq. (2.10), to spectral extinction coefficients $\sigma_t(\lambda)$, versus applying it to RGB triplets $(\sigma_t^R, \sigma_t^G, \sigma_t^B)$. The solid lines in the plots show original spectral data for various colored materials. The dotted lines show the resulting spectral shapes for shorter (above the solid line) and longer (below) transmission distances. The colored squares show the resulting transmission colors based on spectral data (top) and on RGB data (bottom). While the outlined squares match as the RGB extinction coefficients were fitted to them, the colors of all other squares diverge from the correct spectral results.

3.2.2 Donaldson matrix ambiguities

A spectral ambiguity similar to metamerism occurs in the Donaldson matrix Φ in fluorescent materials (Section 2.3.4). By observing a fluorescent material under only a few illuminants and spectral channels, we cannot find the correct Donaldson matrix, as infinitely many matrices correspond to the observations. However, most of these matrices will not be physically plausible and will not correspond to any realistic materials and fluorophores.

In Chapter 5, we look into this problem in more detail. Some publications solve the ambiguity by splitting the matrix into two parts: the diagonal, which represents the reflectance, and the rest. For the reflectance, they use an approach like in the previous subsection, based on basis functions that were fitted to common reflectance spectra. For the fluorescence, they use basis functions that are fitted to a dataset of real fluorophores.

In our method, we also split the matrix into its two components, but we instead parametrize the diagonal using bounded MESE, and the fluorescence by Gaussian distributions. This ensures high smoothness and physically plausible shapes of the estimated Donaldson matrix and yields highly accurate estimates even from very sparse measurements.


3.2.3 Similarity relations in participating media

In participating media, we encounter ambiguities when it comes to measuring their absorption coefficient σ_a , scattering coefficient σ_s , and phase function p . These ambiguities are described by the so-called *similarity relations* [Wyman et al., 1989, Zhao et al., 2014], which show that there are equivalences in the appearance of translucent materials.


The theorem states that under certain illumination and geometry conditions that result in a linearly anisotropic radiance field, which may happen under a uniform diffuse light, there exist infinitely many combinations of material parameters (σ_a, σ_s, p) that yield identical appearance. For a single-parameter phase function like Henyey-Greenstein (HG, Eq. 2.14), it is enough to satisfy the first-order similarity with the anisotropy factor g :

$$\sigma_a^* = \sigma_a, \quad \sigma_s^*(1 - g^*) = \sigma_s(1 - g), \quad (3.5)$$

where an HG-scattering material with parameters (σ_a, σ_s, g) , and a different material with parameters $(\sigma_a^*, \sigma_s^*, g^*)$, will appear identical. We demonstrate this in Figure 3.4. However, notice in (b) that the appearance is not exactly identical: the difference can be observed in the rabbit’s ears, which are very thin geometric features, compared to the rabbit’s body which looks almost perfectly identical.



(a) Constant σ_a and σ_s , varying $g = 0.0$ (left), 0.3, 0.6, and 0.9 (right).



(b) Constant σ_a , varying σ_s following Eq. (3.5) for $g = 0.0$ (left), 0.3, 0.6, 0.9 (right).

Figure 3.4: Demonstrating the similarity relations using the Henyey-Greenstein phase function. Varying anisotropy factor g influences the appearance of the material (a). However, materials whose scattering coefficient σ_s and anisotropy g follow Eq. (3.5) look similar or even identical (b).

This has two interesting consequences. First, on the positive side, this means that there are circumstances under which one actually does not have to know the exact phase function shape to achieve correct rendering results: in such scenarios, assuming a simple HG phase function with any fixed parameter g is sufficient, with the other parameters being scaled accordingly. Zhao et al. [2014] apply this

principle to speed up Monte Carlo rendering of optically dense materials by recalculating their parameters to lower material densities with identical appearance.

On the flip side of the coin, the similarity theorem also means that *measuring* the phase function is a decidedly non-trivial problem: one has to identify the one correct parameter triplet (σ_a, σ_s, g) out of the potentially infinitely many possible solutions under the given measurement geometry.

Breaking the similarity relations The key to “breaking out of” the similarity relations to allow measuring the phase function seems to be using a combination of measurement geometries that together disambiguate between the forward and backward scattering in the material. In our method (Chapter 4), we solve this by illuminating the measured sample from the front using diffuse illuminants, and from the back using a collimated beam. Other methods rely on shining a collimated beam through the sample and observing it from various angles. Such angular (gonio-photometric) measurements were presented for example by Gkioulekas et al. [2013] and Leyre et al. [2014].

3.3 Solving the inverse problem

After introducing the ambiguities that influence whether the problem has a unique global optimum, we focus on methods for solving the inverse problem. In general, optimization problems have the following form [Antoniou and Lu, 2021]:

$$\text{minimize } f(\boldsymbol{\pi}) \text{ subject to } \boldsymbol{\pi} \in \boldsymbol{\Pi}, \quad (3.6)$$

where $\boldsymbol{\Pi}$ is the feasible parameter region that constrains the valid range of the parameters. In our context, following Eq. (3.1), we have $f(\boldsymbol{\pi}) = e(I_{\text{ref}}, I(\boldsymbol{\pi}))$, and the parameters $\boldsymbol{\pi}$ are inputs to the rendering algorithm $I(\boldsymbol{\pi})$. The feasible region $\boldsymbol{\Pi}$ is used to ensure that the parameters are physically meaningful. For example, in rendering participating media, we could constrain the input absorption and scattering coefficients to non-negative numbers.

3.3.1 Analytical solutions

When the whole forward problem can be described in mathematical terms and in analytical closed-form expressions (mainly without any integrals and limits), it might be feasible to find the inverse analytically.

Generally, there are two analytical approaches. First, we can try to rewrite the equations of the forward problem such that the unknown variables appear on the left side of the equations, which allows us to trivially compute the solution. Second, we can rely on computing the derivatives (gradient) of the minimization problem and finding the parameters that cause the derivatives (gradient) to be zero, or further also using the second derivatives (Hessian), hence obtaining parameters that correspond to a local minimum. The necessary and sufficient conditions for local minima in terms of gradients and Hessians are described for example by Antoniou and Lu [2021, Section 2.5].

We now show three examples of how an analytical solution can be used in computer graphics.

Measuring spectral reflectance A trivial example of an analytical inverse problem is measuring the reflectance spectrum of a diffuse object. The forward problem is simply defined as follows:

$$I(\lambda) = r(\lambda) l(\lambda), \quad (3.7)$$

where $I(\lambda)$ is the spectrum measured by a spectrometer that is pointed at a diffuse object with reflectance $r(\lambda)$ illuminated by an illuminant of spectrum $l(\lambda)$. The inverse problem is finding the reflectance $r(\lambda)$ given we know the illuminant spectrum $l(\lambda)$ and what the spectrometer measured $I_{\text{ref}}(\lambda)$. The solution is trivial:

$$r(\lambda) = I_{\text{ref}}(\lambda) l^{-1}(\lambda). \quad (3.8)$$

Beer’s law Another straightforward analytical solution is often used for measuring the extinction coefficient σ_t of homogeneous participating media. From the Beer’s law, Eq. (2.10), we can derive that the light intensity I of a collimated beam traveling through the medium for a distance ℓ is:

$$I = I_0 \cdot \exp(-\ell \cdot \sigma_t), \quad (3.9)$$

where I_0 is the beam intensity without any participating medium present. Inverting this equation allows us to measure the extinction coefficient σ_t of a homogeneous medium by simply capturing the beam’s intensity through the medium:

$$\sigma_t = \ell^{-1}(\ln I_0 - \ln I). \quad (3.10)$$

Kubelka and Munk While the previous two examples were rather trivial, there exist more sophisticated light transport models for which analytical solutions are available. An example is the model of Kubelka and Munk [1931]. It is a 1D radiative transfer model for translucent materials, which assumes a semi-infinite layer of a homogeneous medium illuminated by a diffuse hemispherical light. It is based on calculating the upward and downward flux, so it can approximate the total reflectance and transmission from that medium. Even though it is a very crude approximation, it used to be very popular in the textiles, paints, and printing industries for its simplicity, and it was later generalized to support refractive boundaries, non-diffuse illumination, and 3D radiative transfer [Nobbs, 1985, Yang and Hersch, 2008, Sandoval and Kim, 2015].

Most importantly, the KM formulas can be easily inverted, as shown for example by Vargas [2002]. The idea is to first measure the light that is reflected from a layer of the translucent material on a black absorbing background, and then on a white reflective background, which is rather similar to our method in Chapter 4, although we do not rely on Kubelka and Munk at all. Once these two measurements are available, it is possible to analytically compute the scattering parameters of the KM model for the given material.

3.3.2 Lookup tables

In a lot of cases, finding the solution analytically is not viable, so alternative approaches have to be used. A rather straightforward approach is based on

building a *lookup table*, also known as performing a *grid search* or a *parameter sweep*, as we simply search through the parameter space to find the optimal value. Formally, we could write that a lookup table \mathcal{T} consists of N pre-computed values of the rendering function I for varying parameters $\boldsymbol{\pi}$:

$$\mathcal{T} = \left\{ i \in \{1, \dots, N\} \mid T_i = I(\boldsymbol{\pi}_i) \right\}, \quad (3.11)$$

where $\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_N \in \boldsymbol{\Pi}$ are N arbitrarily sampled parameters anywhere from the feasible region. To find which parameters $\hat{\boldsymbol{\pi}}$ best correspond to the reference image I_{ref} using the error metric e , one simply iterates over all indices in the lookup table to find the minimum:

$$\hat{\boldsymbol{\pi}} = \boldsymbol{\pi}_k \quad \text{such that} \quad e(I_{\text{ref}}, T_k) \leq e(I_{\text{ref}}, T_i) \quad \forall i \in \{1, \dots, N\}. \quad (3.12)$$

Example Suppose we want to measure the absorption and scattering coefficients σ_a, σ_s of a diffusely illuminated cube. Assuming the parameters are in the range of $[0, 10] \text{ mm}^{-1}$, i.e., $\boldsymbol{\Pi} = [0, 10] \times [0, 10]$, we could decide to uniformly sample the parameter space into $11 \cdot 11 = 121$ parameters:

$$\begin{aligned} \left\{ \boldsymbol{\pi}_i \mid \boldsymbol{\pi}_i = (m, n) : m, n \in \{0, \dots, 10\} \right\} = \\ \{(0, 0), (1, 0), \dots, (10, 0), (0, 1), (1, 1), \dots, (10, 10)\}. \end{aligned} \quad (3.13)$$

The lookup table would then consist of 121 renders of the cube with the varying absorption and scattering coefficients. Given a reference image I_{ref} of a cube, we find the best match $\hat{\boldsymbol{\pi}}$ by comparing the reference to each image in the lookup table and picking the one with the lowest difference.

Applicability It is clear that this approach does not scale well: with each additional dimension of the parameter space, the memory demands to store the table increase exponentially, and so does the computation time required to pre-compute the whole table. However, this solution is still extremely useful because of its inherent simplicity and the ability to approximately find the global minimum even if multiple local minima are present. It is viable whenever the dimensionality is low and the rendering function only returns a small image with a tiny memory footprint. For example, in Chapter 4, we use three-dimensional lookup tables to store the absorption and scattering coefficients together with the scattering anisotropy factor. The rendering function only returns a single monochromatic intensity, so the values stored in the lookup table are just individual floating point numbers. Similarly, in Elek et al. [2021], a lookup table is built for one-dimensional scattering profiles, so each stored value is only a 1D curve.

Other uses Building a lookup table is useful even if the actual inverse problem is solved in a different way. The lookup table allows studying the behavior of the problem's objective function, specifically, whether it has a single minimum or there are any ambiguities (Section 3.2). Examples include Gkioulekas et al. [2013, Section 6.2, Figure 5] and Zhao et al. [2014, Figure 3].

3.3.3 Learning-based approaches with deep networks

Instead of building a lookup table, it is possible to build a *deep network* whose purpose is to estimate the matching parameters $\hat{\pi}$ on the output, given the reference image I_{ref} on the input. In the context of computer graphics, this approach was proposed by Che et al. [2020], who called it the *inverse transport network*.

Without going into too many details of deep networks, which are not the core topic of this thesis, we explain at least the basic principle (Figure 3.5). First, the network is given an image I as its input. The input is fed to an encoder, which is a sub-network that converts a large input (an image I , in this case) into a small encoding (the parameters π , in this case). If the network was trained well, these parameters are close to the actual parameters that we were looking for, so the network solved our inverse problem.

During the training of the network, the encoding π is fed through an actual renderer to give an output image I_o . The goal is that the input and output images match $I = I_o$, so the internal weights of the encoder sub-network are iteratively adjusted during the training to ensure this behavior.

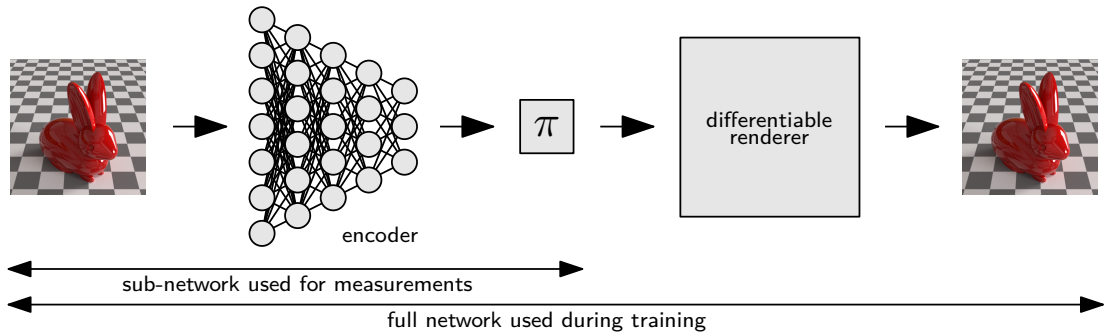


Figure 3.5: The inverse transport network consists of a sub-network that takes an image and returns its estimated parameters, and a differentiable renderer that converts the parameters back into an image. During training, the network is conditioned to ensure the input and output images are similar.

Applicability As with most deep networks, this approach is suitable mainly when it is plausible to generate a large training dataset on which the network can be trained. On low-dimensional inverse problems, it might be more efficient and more accurate to build a lookup table, as the performance of a lookup table is guaranteed and can be easily evaluated. On the other hand, a deep network typically has a much smaller memory footprint, as it is sufficient to store the architecture and its weights, and it is less rigid, meaning it might be able to recognize even inputs that it has never seen before, or that are slightly different from the assumptions. The publication of Che et al. [2020] is a prime example, since they trained the network on different object shapes than they used in the evaluation, proving that the network is not sensitive to specific details. This would be hardly achievable with a lookup table.

3.3.4 Gradient-based iterative approaches

The whole discipline that is encompassing the theory and practice of solving the general optimization problem, as in Eq. (3.6), is called *mathematical programming*. These days, the name is confusing as it has nothing to do with computer programming, but comes from the historical use of the word *programming* in the military. Based on the mathematical properties of the optimization problem, e.g., if it is a linear or a quadratic function, we have *linear programming*, *integer programming*, *quadratic programming*, *non-linear programming*, or *dynamic programming* [Antoniou and Lu, 2021].

In general, the algorithms for solving the problems are *iterative* and *descending* at the same time. The former refers to the fact that the algorithms obtain the solution iteratively, starting with an initial estimate $\boldsymbol{\pi}_1 \in \boldsymbol{\Pi}$, and then calculating a sequence of new estimates based on the previous estimates. The latter refers to each consecutive estimate resulting in a decrease in the objective function, so the algorithm descends towards the minimum. Formally, $f(\boldsymbol{\pi}_{k+1}) < f(\boldsymbol{\pi}_k)$.

General form Many iterative optimization algorithms compute the successive steps following a simple formula that can be generalized as follows:

$$\boldsymbol{\pi}_{k+1} := \boldsymbol{\pi}_k + \alpha_k S_k d_k, \quad (3.14)$$

where d_k is the *direction* along which the new parameter estimate is located, S_k is a transformation matrix that may further modify the direction, and α_k is the *step size* that modifies how far the new parameter is along the direction.

Gradient descent Many of the standard approaches are based on the *gradient* (first derivative) $\nabla f(\boldsymbol{\pi})$ of the objective function with respect to the parameters. Such an approach is called *gradient descent optimization*, and the direction in Eq. (3.14) is simply the opposite of the gradient: $d_k = -\nabla f(\boldsymbol{\pi}_k)$. The idea follows from calculus: if a function is continuous and has continuous first-order partial derivatives, then if $f(\boldsymbol{\pi})$ is a local minimum, then $\nabla f(\boldsymbol{\pi}) = 0$. Hence, following the path towards a lower gradient ensures that we will find a local minimum, although it is not guaranteed to be the global minimum, and it may also be a saddle point [Antoniou and Lu, 2021, Chapter 2]. In the simplest approach, called the *steepest descent*, the transformation matrix is not used and may be replaced by an identity matrix $S_k = 1$.

Newton direction In more advanced approaches based on the quadratic approximation of the Taylor series, such as in the *Newton*, *Gauss-Newton*, or *quasi-Newton methods* [Antoniou and Lu, 2021, Sections 5.3 and 5.4 and Chapter 7], the transformation matrix is equal to, or at least approximates, the inverse of the Hessian matrix (second-order partial derivatives) $S_k = H_k^{-1}$, and the final vector $S_k d_k = -H_k^{-1} \nabla f(\boldsymbol{\pi}_k)$ is called the *Newton direction*. Unfortunately, relying on the second-order derivatives is not always ideal: for high-dimensional problems, the Hessian may be a large matrix; numerical inaccuracies often occur due to a noisy input signal; and the inverse of the Hessian may not even exist. This is why the inverse of the Hessian is often only approximated.

Conjugate directions The basic approach based on $d_k = -\nabla f(\boldsymbol{\pi}_k)$ often results in the successive steps $\boldsymbol{\pi}_{k+1}$ essentially “zig-zagging” around the parameter space, which is not very efficient. Instead, to decrease the number of steps required for convergence, it is possible to use *conjugate directions* [Antoniou and Lu, 2021, Chapter 6], which are sequentially constructed using the previous direction and the gradient. Constructing the conjugate directions can be done in more ways, including those by Fletcher and Reeves [1964] and Powell [1964].

Heuristics for step sizes The step sizes α_k are an important ingredient of the optimization algorithms. Once we know the direction d_k , or rather $S_k d_k$, the problem collapses into a one-dimensional optimization along the direction, and the ideal step size α_k can be found numerically using a one-dimensional search method [Antoniou and Lu, 2021, Chapter 4]. Unfortunately, that requires additional evaluations of the objective function, which may take a very long time in the case of computer graphics, where the function may be a Monte Carlo renderer. A similar problem occurs in deep learning, which resulted in a growing interest in *finding the step sizes heuristically*. The most trivial heuristic is to use a *constant step size*, but that generally does not perform very well. In Chapter 5, we use a more sophisticated method: the adaptive moment estimation algorithm *Adam* [Kingma and Ba, 2017], which is very popular in training deep networks.

Applicability The gradient-based iterative optimization is applicable for differentiable objective functions whose gradient is reasonably fast to compute. In computer graphics, the use of gradient-based optimization has risen together with the new research in *differentiable rendering* (Section 3.4), which allows efficiently computing the gradient of a Monte Carlo renderer. Unfortunately, gradient-based optimization does not generally guarantee that the global optimum will be found, especially if the objective function has many ripples and saddle points. If that is the case, we can either initialize the optimization with an estimate $\boldsymbol{\pi}_1$ that is already close to the optimum, if we have such an estimate; or we can try to run the optimization algorithm progressively, starting from low-resolution and low-quality images, slowly converging to the actual optimum by increasing the resolution gradually, which may avoid being stuck in a local minimum.

3.3.5 Gradient-free iterative approaches

Finally, we have a look at special cases of iterative approaches that do not need information on the objective function’s gradient, and in fact, the function does not even need to be differentiable. That is beneficial if we want to treat the objective function as a black box without really knowing what happens inside, which is sometimes valuable in the case of complex Monte Carlo rendering. Compared to the lookup tables (Section 3.3.2), these iterative approaches are viable even for higher dimensions as they do not require excessive storage and evaluations.

Gradient-free conjugate directions In the previous subsection, we mentioned that certain conjugate directions methods do not require computing the Hessian matrix. In fact, some of them do not even require the gradient. An exam-

ple is the *Powell method* [Powell, 1964], which generates the conjugate directions through a series of line searches.

Hill climbing Hill climbing was described, for example, by Russell et al. [2010] for use in artificial intelligence, originally for maximization (*climbing*), which is equivalent to minimization with a negated objective function (*descent*). Following the description of Miao et al. [2016], the algorithm starts with initial parameters $\boldsymbol{\pi}_1$ and step sizes α_1 . The objective function is evaluated at $2 \cdot \|\boldsymbol{\pi}\|$ neighbors of $\boldsymbol{\pi}_k$ by changing a single parameter of $\boldsymbol{\pi}_k$ by $\pm\alpha_k$. The neighbor with a higher objective function value $f(\boldsymbol{\pi}_k \pm \alpha_k)$ is selected. If no such neighbor exists, α_k is reduced by half.

Nelder-Mead simplex The algorithm by Nelder and Mead [1965] [Miao et al., 2016, Johnson, 2019] uses a *simplex* with $\|\boldsymbol{\pi}\| + 1$ vertices in the $\|\boldsymbol{\pi}\|$ -dimensional parameter space. In one dimension, the simplex is a line segment; in two dimensions, it is a triangle; in three dimensions, it is a tetrahedron; and so on. We start with an initial simplex whose vertices correspond to the objective function values at that initial parameters. In each iteration, we replace the vertex with the highest objective function value by another point according to a set of rules. The simplex adapts itself and should contract to the final minimum. In the context of computer graphics, this method was used, for example, by Narasimhan et al. [2006] to measure participating media.

3.4 Differentiable rendering

To use the gradient-based iterative approaches in Section 3.3.4, the objective function’s first partial derivatives (the gradient) must be known, sometimes alongside the second-order derivatives (the Hessian). How exactly one can obtain the gradient of a Monte Carlo renderer is the core question of *differentiable rendering*, which is an important topic of computer graphics. Since our measurement methods in Chapters 4 and 5 do *not* rely on differentiable rendering, we only introduce the topic very briefly.

Derivation by hand In principle, it is possible to differentiate the light transport equation and its individual parts by hand. Unfortunately, this solution does not scale with the addition of new types of illuminants, material types with different BSDF, and so on. Deriving the gradient by hand was used, for example, in the volumetric representation of Vicini et al. [2021a]. A hybrid approach that combines manual derivatives with automatic differentiation was presented, for example, by Li et al. [2018], whose solution we will mention later again.

Finite differences Apart from deriving the gradient by hand, we can also rely on *numerical differentiation*. The most straightforward solution is based on *finite differences*. For example, for a one-dimensional function $f(x) : \mathbb{R} \rightarrow \mathbb{R}$, we can approximate its partial derivative $\partial f / \partial x$ using central differences [Schroeder, 2022, Section 3.2]:

$$\frac{\partial f}{\partial x} \approx \frac{f(x - \varepsilon) + f(x + \varepsilon)}{2\varepsilon}, \quad (3.15)$$

where $\varepsilon \in \mathbb{R}$ is a real number with a very low value. However, this approach is numerically unstable, choosing the value of ε is experimental, and evaluating the finite differences requires evaluating the objective function many times. In the case of rendering, such evaluations may be very costly, and the Monte Carlo noise may result in significant numerical issues. Hence, it makes sense to rely on more sophisticated automatic differentiation.

3.4.1 Automatic differentiation

A computer graphics renderer is a computer program, and a computer program is essentially a collection of elementary operations such as additions, multiplications, exponentiations, and similar. Derivatives of the elementary operations are trivially defined, for example, $\partial x^2 / \partial x = 2x$ or $\partial \sin x / \partial x = \cos x$. The *chain rule* then defines the partial derivative of the composition of two functions, which allows the nesting of the elementary operations into each other. Applying the chain rule repeatedly for the whole computer program execution then yields the partial derivatives of the whole program. Formally, a computer program takes N input parameters and outputs M parameters, so the program is a map $\mathbb{R}^N \rightarrow \mathbb{R}^M$. Its derivative is then the Jacobian matrix J of size $N \times M$, where the individual elements are partial derivatives: $J_{(n,m)} = \partial f_n / \partial \pi_m$.

Execution graph and its traversal The program’s execution can be thought of as a graph of elementary operations, with inputs on one side, and outputs on the other side. Implementing the automatic differentiation means applying the chain rule to the graph. Practically implementing this process was covered, for example, in the course of Schroeder [2022]. Generally, there are two major approaches: if the chain rule is applied from the inputs towards the outputs, it is called the *forward mode* differentiation; otherwise, it is the *reverse mode* differentiation [Bartholomew-Biggs et al., 2000]. Both approaches have benefits and downsides. For $N > M$, performing the reverse mode differentiation results in fewer operations than the forward mode. However, the reverse mode typically requires storing the whole computation graph, or constructing an adjoint pass.

Open-source implementations There are many open-source automatic differentiation frameworks, many of them motivated by deep learning. Examples include JAX [Bradbury et al., 2018], PyTorch [Paszke et al., 2019], TensorFlow [Martín Abadi et al., 2015], or approaches more targeted at computer graphics and physics include Dr. Jit [Jakob et al., 2022b] and Taichi [Hu et al., 2020]. In our method in Chapter 5, we use JAX to differentiate our objective function.

3.4.2 Challenges and solutions

Differentiable Monte Carlo rendering is, unfortunately, not as straightforward as simply applying automatic differentiation to the rendering algorithm. In the following paragraphs, in no particular order, we try to emphasize the key challenges that differentiable rendering has to overcome. We also refer the reader to two excellent courses on differentiable rendering that were prepared by Zhao et al. [2020] and Zhao et al. [2021].

Magnitude Suppose the algorithm renders an image in a relatively low resolution of 800×600 pixels, and each pixel has 128 samples, which is not many. Even then, the renderer had to simulate almost 62 million distinct light paths. Furthermore, suppose that the optimized parameters are the pixels of a spatially varying texture or the single-scattering albedo of a heterogeneous participating medium. If we relied on naive automatic differentiation, we would need to construct computation graphs and Jacobian matrices of more than 10^{12} elements, which is unacceptable. Specialized solutions have been developed to overcome this issue, including the *radiative backpropagation* [Nimier-David et al., 2020], and its later improvement *path replay backpropagation* [Vicini et al., 2021b], which has a constant memory and linear time complexity. In our 3D printing optimization publication [Nindel et al., 2021], we relied on radiative backpropagation to significantly reduce the time required for executing our method.

Discontinuities Another challenging problem occurs when the objective function is discontinuous with respect to the parameters π . For example, parametrizing the diffuse reflectance of an object is usually continuous. On the other hand, parametrizing the position or rotation of complex geometry is often discontinuous because the object suddenly occupies different pixels of the image. Solutions to this problem often rely on the *Reynolds transport theorem*, which allows decomposing an integral containing discontinuities into an interior and boundary integral, which may be solved independently. Probably the first method addressing the discontinuities in a general-purpose renderer was published by Li et al. [2018], and the discontinuities were handled via *edge sampling*. A later solution by Zhang et al. [2020] is based on a path-space formulation and detecting those paths that contain a boundary segment, which was also extended to participating media [Zhang et al., 2021]. Finally, *reparametrizing* the problem such that the discontinuities are eliminated or reduced was proposed by Loubet et al. [2019].

Sampling strategies When differentiating a Monte Carlo renderer, there are also more approaches to handling the random sampling. Zeltner et al. [2021] formally investigated the approaches and proposed terminology such as *detached* and *attached* strategies. In a more specialized publication, Nimier-David et al. [2022] proposed *differential ratio tracking* for sampling participating media with unbiased gradients.

Open-source implementations Currently, the state-of-the-art open-source implementation of a differentiable renderer seems to be Mitsuba3 [Jakob et al., 2022a], which is based on the aforementioned `Dr.Jit` automatic differentiation framework. An older differentiable renderer, which does not seem to be developed anymore, is Redner¹, which is based on the publication of Li et al. [2018].

¹<https://github.com/BachiLi/redner>

Part II
Journal publications

4 | Our measurement method for translucent materials

The first journal publication inserted in this thesis, Iser et al. [2022], focuses on measuring translucent materials, whose fundamental theory we covered in Section 2.2. Such materials are truly ubiquitous: examples include many liquid-state materials (water, wine, milk), aerosols (fog, clouds, smoke), biological matter (tissue, skin, blood), and many solids (plastic, wood, porcelain). Characterizing and measuring such materials is, therefore, an important topic not only for predictive rendering but also for various other fields, including biology and medicine [Tuchin, 1993, Jacques, 2013] or atmospheric sciences [Emde et al., 2016].

Problem To characterize the translucent materials, we use their bulk optical properties: the absorption and scattering coefficients, and phase function anisotropy (Section 2.2). These properties are spectrally varying, i.e., they depend on wavelength. The existing measurement methods are either very complicated, require expensive optical components, or the affordable methods rely on RGB cameras, so the results are not spectrally resolved. Relying on RGB properties leads to many inherent issues that we covered in Section 3.2.1, including transmission color shifts, metameric failures, or out-of-gamut colors. Hence, spectral measurements are a requirement for predictive rendering, but such a method that would also be affordable did not exist yet.

Our contribution and solution Our core question in the publication is whether it is possible to estimate the unknown material parameters using only three point measurements with a spectrometer. This is not a simple question, as we must account for ambiguities, including the similarity relations (Section 3.2.3). After studying the existing approaches and their behavior, we found and proved that there exists a simple and easy-to-visualize relation between material appearance and its optical properties, which we call the *appearance map*. This allows for building an affordable setup for fast spectral measurements of a given sample. The core of our approach is how light intensity changes in three different settings: first, a given material sample is diffusely illuminated and placed against a black background and a white background; then, it is observed with a collimated illumination. These three combinations are not only easy to simulate using a Monte Carlo renderer (Section 2.2), but also easy to capture in real life, which allows spectrally resolving the whole parameter triplet as an inverse problem using pre-computed lookup tables (Section 3.3.2). The results are demonstrated on materials used in full-color 3D printing. We measure the base CMYKW inks (cyan, magenta, yellow, black, white) and validate our method’s accuracy by predicting the spectral appearance of these inks mixed in various proportions within a halftoned color checker.

The full publication is inserted and begins on the next page.

The publication is published as open access and licensed under a Creative Commons Attribution International 4.0 License. The full citation follows.

Tomáš Iser, Tobias Rittig, Emilie Nogué, Thomas Klaus Nindel, and Alexander Wilkie. Affordable Spectral Measurements of Translucent Materials. *ACM Transactions on Graphics*, 41(6):1–13, December 2022. ISSN 0730-0301, 1557-7368. doi: 10.1145/3550454.3555499. URL <https://dl.acm.org/doi/10.1145/3550454.3555499>

Please note that the supplement of this publication is available in the ACM Digital Library under the link above.



Affordable Spectral Measurements of Translucent Materials

TOMÁŠ ISER, Charles University, Czech Republic

TOBIAS RITTIG, Charles University, Czech Republic

EMILIE NOGUÉ, Imperial College London, United Kingdom

THOMAS KLAUS NINDEL, Charles University, Czech Republic and Berufsakademie Sachsen, Germany

ALEXANDER WILKIE, Charles University, Czech Republic

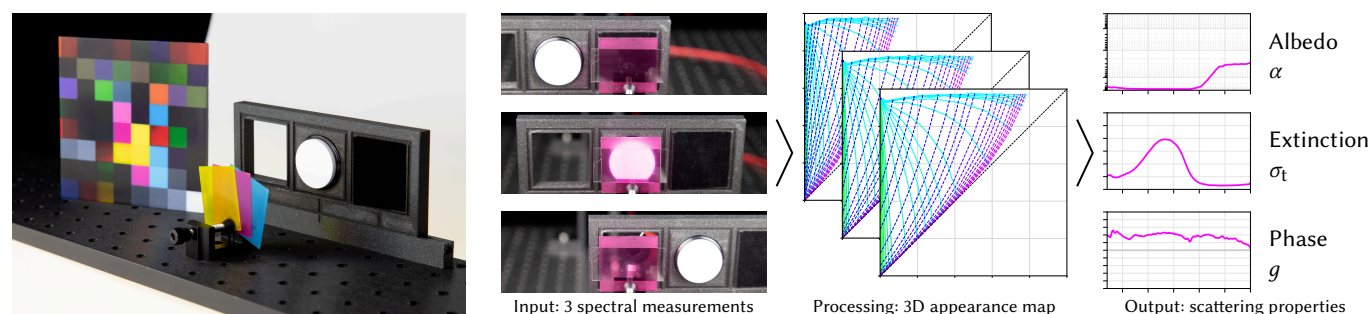


Fig. 1. We present a method for measuring bulk optical properties of translucent materials, such as the base inks of a full-color 3D printer. Starting from a thin sample, we measure three spectral intensities: on a black background, white background, and of a collimated light beam. After a fitting step through our precomputed 3D appearance map, we obtain the material’s spectrally-varying albedo, extinction coefficient, and phase function anisotropy.

We present a spectral measurement approach for the bulk optical properties of translucent materials using only low-cost components. We focus on the translucent inks used in full-color 3D printing, and develop a technique with a high spectral resolution, which is important for accurate color reproduction. We enable this by developing a new acquisition technique for the three unknown material parameters, namely, the absorption and scattering coefficients, and its phase function anisotropy factor, that only requires three point measurements with a spectrometer. In essence, our technique is based on us finding a three-dimensional *appearance map*, computed using Monte Carlo rendering, that allows the conversion between the three observables and the material parameters. Our measurement setup works without laboratory equipment or expensive optical components. We validate our results on a 3D printed color checker with various ink combinations. Our work paves a path for more accurate appearance modeling and fabrication even for low-budget environments or affordable embedding into other devices.

CCS Concepts: • **Computing methodologies** → **Rendering**; *Reflectance modeling*; • **Applied computing** → **Computer-aided manufacturing**; *Physics*.

Authors’ addresses: Tomáš Iser, tomas@cgg.mff.cuni.cz, Charles University, Faculty of Mathematics and Physics, Malostranské nám. 25, Prague, 118 00, Czech Republic; Tobias Rittig, tobias@cgg.mff.cuni.cz, Charles University, Faculty of Mathematics and Physics, Malostranské nám. 25, Prague, 118 00, Czech Republic; Emilie Nogué, e.nogue@imperial.ac.uk, Imperial College London, Exhibition Rd, South Kensington, London, SW7 2BX, United Kingdom; Thomas Klaus Nindel, thomas@cgg.mff.cuni.cz, Charles University, Faculty of Mathematics and Physics, Malostranské nám. 25, Prague, 118 00, Czech Republic and Berufsakademie Sachsen, Hans-Grundig-Strasse 25, Dresden, 01307, Germany; Alexander Wilkie, wilkie@cgg.mff.cuni.cz, Charles University, Faculty of Mathematics and Physics, Malostranské nám. 25, Prague, 118 00, Czech Republic.



This work is licensed under a Creative Commons Attribution International 4.0 License.
© 2022 Copyright held by the owner/author(s).
0730-0301/2022/12-ART199
<https://doi.org/10.1145/3550454.3555499>

Additional Key Words and Phrases: spectral measurements, optical properties, translucent materials, light scattering, appearance modeling, predictive rendering, 3D printing, computational fabrication

ACM Reference Format:

Tomáš Iser, Tobias Rittig, Emilie Nogué, Thomas Klaus Nindel, and Alexander Wilkie. 2022. Affordable Spectral Measurements of Translucent Materials. *ACM Trans. Graph.* 41, 6, Article 199 (December 2022), 13 pages. <https://doi.org/10.1145/3550454.3555499>

1 INTRODUCTION

The world around us is filled with translucent materials which partially absorb or scatter light that passes through them. Examples include many liquid-state materials (water, wine, milk), aerosols (fog, clouds, smoke), biological matter (tissue, skin, blood), and many solids (plastic, wood, porcelain). Characterizing and measuring the optical properties of such materials is therefore an important topic not only for appearance predictions in computer graphics, but also for various other fields, including biology and medicine [Jacques 2013; Tuchin 1993] or atmospheric sciences [Emde et al. 2016].

Its importance also shows in full-color 3D printing, where recent advances enabled fabricating customized objects with desired appearances. The appearance reproduction is achieved by depositing base printing inks layer-by-layer, forming a high-resolution optically heterogeneous grid. Each voxel of this grid is translucent, and finding their adequate arrangement is important to realize the desired color mixing. Hence, for accurate printouts without texture blurring or color inaccuracies, it is key to know the inks’ optical properties [Elek et al. 2017; Nindel et al. 2021; Sumin et al. 2019].

Measuring these optical properties is our main goal. More importantly, we focus on estimating them *spectrally*, allowing for accurate predictions of the colors when the base inks are *spatially mixed*

within the 3D printouts. As we show in our supplemental document, spectral simulations of translucent materials eliminate color inaccuracies caused by metameric ambiguities and attenuation non-linearity, as opposed to three-channel RGB simulations.

The optical properties are estimated by designing and solving an appropriate *inverse problem*, which consists of first measuring light interacting with the material particles in bulk, and then numerically finding which optical properties must have led to such results. For simplicity (unlike recently studied non-exponential media [Bitterli et al. 2018]), we model each material as *homogeneous*, with molecules distributed in an *uncorrelated, uniform* fashion. Under this assumption, the material is characterized by the wavelength (λ) dependent *extinction coefficient* $\sigma_t(\lambda)$ [mm^{-1}], the *single-scattering albedo* $\alpha(\lambda)$, and the *phase function* $p(\lambda, \theta)$. The extinction coefficient $\sigma_t(\lambda)$ characterizes the *exponential attenuation* (known as Beer’s law [Jarosz 2008]) of the light intensity $I(\ell)$ after traveling through the medium for a distance ℓ with the initial intensity I_0 :

$$I(\ell) = I_0 \cdot \exp(-\ell \cdot \sigma_t). \quad (1)$$

The attenuation is caused by absorption and out-scattering of the photons, parametrized by the *absorption and scattering coefficients* $\sigma_a(\lambda)$ and $\sigma_s(\lambda)$. It holds that $\sigma_t(\lambda) = \sigma_a(\lambda) + \sigma_s(\lambda)$, and $\alpha(\lambda) = \sigma_s(\lambda)/\sigma_t(\lambda)$. The phase function $p(\lambda, \theta)$ characterizes the angular θ distribution of the light scattering. We assume a simple single-parameter phase function of Henyey and Greenstein [1941], which approximates the more accurate Mie scattering model, but only requires a single parameter $g \in [-1, 1]$; $g < 0$, $g = 0$, and $g > 0$ for dominantly backward, isotropic, and dominantly forward scattering, respectively. We acquire the parameters $(\alpha(\lambda), \sigma_t(\lambda), g(\lambda))$ for the given materials with a high spectral resolution of approximately 1000 wavelengths within the visible spectrum.

Such measurements using the existing methods (Sec. 2.2) require complicated calibrations and/or expensive laboratory-grade equipment such as gonio-photometers, rotating platforms, integrating spheres, time-resolved detectors, or lasers. Although there are alternative approaches with emphasis on simplicity and affordability, they rely on acquiring one-dimensional spatial profiles, which is impossible with a spectrometer that can only measure a single point. Hence, these methods only use RGB cameras, and using them for spectral measurements would require expensive hyperspectral photography. Furthermore, many of these setups suffer from ambiguities that make them unable to estimate the phase function.

Our contribution. To enable high-resolution spectral measurements, our core question is how to resolve the three unknown material parameters $(\alpha(\lambda), \sigma_t(\lambda), g(\lambda))$ as quickly as possible per each wavelength, and with a minimum number of point measurements with a spectrometer. This is not a simple question, as we need to account for various factors, including the similarity relations [Wyman et al. 1989; Zhao et al. 2014]. After studying the existing approaches and their behavior, we found and proved that there exists a simpler and easy-to-visualize relation between material appearance and its optical properties (Sec. 3), which we call the *appearance map*. The core of our approach is how light intensity changes in three different settings: first, a given material sample is diffusely illuminated and placed against a black background, and a white background; then,

it is observed with a collimated illumination. These three combinations are not only easy to simulate using a Monte Carlo renderer, but also easy to capture in real life with a simple and affordable physical setup with single spectrometer (Sec. 4), which allows spectrally resolving the whole parameter triplet (α, σ_t, g) .

While we believe our method to be applicable in many fields, we remain concise and demonstrate the results on the inks for full-color 3D printing. We measure the base CMYKW (cyan, magenta, yellow, black, white) inks and validate our method’s accuracy by predicting the spectral appearance of their various mixtures within a halftoned color checker (Sec. 5). This implies the promised applicability of our method to the aforementioned printing optimization pipelines.

2 RELATED WORK

We begin by introducing relevant *prediction models* for simulating light behavior in translucent materials (Sec. 2.1). Using these models inversely is then the core of all *measurement methods* (Sec. 2.2), which show strategies for capturing materials and inversely fitting their matching optical properties. Lastly, we relate our work to the context of translucency appearance and 3D printing (Sec. 2.3).

2.1 Prediction models

We briefly look at prediction models that have been extensively used by methods for estimating the bulk optical properties in the past. We refer to Frisvad et al. [2020] for a complete survey.

The following models have been developed for a simplified situation of diffusely illuminated, infinitely wide, homogeneous layers of translucent materials. Kubelka and Munk [1931] show the computation of the total reflectance and transmission, a model popular for modeling textiles, paints, and 2D printing. It was later generalized to support refractive boundaries, non-diffuse illumination, and 3D radiative transfer [Nobbs 1985; Sandoval and Kim 2015; Yang and Hersch 2008]. More accurate is the *adding-doubling model* [Prah 1995; van de Hulst 1980], which generalizes the problem to multi-layer material stacks (stacked in one dimension), including correctly handling refractive boundaries. It is also useful for fast rendering of layered materials [Jakob et al. 2014; Zeltner and Jakob 2018]. While these models are fast, they lack flexibility and simplicity in defining the complete 3D geometry, including the light sources and sensors, and samples of finite dimensions, which are important in our method to accurately compute the appearance map.

The *diffusion approximation* [D’Eon and Irving 2011; Haskell et al. 1994; Jensen et al. 2001] is another often used model. It is based on the assumption that scattering events are more frequent than absorption. This makes it inaccurate for low-albedo materials, excluding the application for example to inks of dark colors.

In the end, we decided to rely on *Monte Carlo* methods that solve the radiative transfer by probabilistically sampling the space of possible light paths through the scene, and calculating their radiance contribution [Kajiya 1986]. This is the most flexible solution that supports various scenes including translucent materials [Elek et al. 2017; Novák et al. 2018], and offers accurately simulating a wide range of material properties within our method’s geometry. We used a modified fast and flexible implementation of Nimier-David et al. [2019] and Jakob et al. [2022].

2.2 Measurement methods

The simplest measurement can be performed by shining a collimated beam of intensity I_0 through a material sample of thickness ℓ , measuring its attenuated intensity I , and simply inverting Eq. (1):

$$\sigma_t = \ell^{-1} (\ln I_0 - \ln I), \quad (2)$$

but this approach gives no information on light scattering, necessitating a more sophisticated approach.

A common spectral measurement method is based on capturing the total *hemispherical reflectance and transmittance* of a material sample [Pickering et al. 1992, 1993; Prahl et al. 1993] using two integrating spheres, or a Coblentz hemisphere [Schröder et al. 2015]. Two material parameters (either α, σ_t ; or α, g ; assuming that the third parameter is known, e.g., from Eq. (2)) are then fitted using *inverse adding doubling* [Prahl 2011]. Our method does not require integrating spheres and instead uses significantly less expensive and easily switchable reflective and absorptive backgrounds.

Another group of methods relies on fitting the optical properties to *one-dimensional brightness curves*. They illuminate a sample in a specific way, and then capture how the brightness changes spatially over a certain region. These methods are affordable, because they use a simple RGB camera to take a photograph, from which the brightness curve is extracted. The main downside is that one cannot use a spectrometer as it can only measure a single point of interest, and hyperspectral cameras or various spectral filters would be an expensive alternative. Examples of such methods include the one of Narasimhan et al. [2006], which uses dilution to reach such low concentrations of the material that single scattering dominates over multiple scattering. Other methods combine fitting the one-dimensional profile to the diffusion-based prediction model with also measuring the total diffuse reflectance [Jensen et al. 2001; Weyrich et al. 2006]. Papas et al. [2013] used a similar approach with a custom-made container with five LEDs of different spectra, and used a combination of Monte Carlo and quantized diffusion [D'Eon and Irving 2011] as their prediction models.

Elek et al. [2021] presented an alternative method acquiring a *lateral scattering profile* on a step-edge black and white background. Similarly to us, they demonstrate their results on 3D printing, fit the properties to a Monte Carlo simulated dataset, and use contrasting backgrounds. However, they have exactly the same disadvantage as above, requiring a camera capture. A similar idea of using different backgrounds, including a mirror, was also opened in a short concurrent work by Pranovich et al. [2021]. However, they only used a simplistic light transport model, did not estimate phase functions, and did not analyze the appearance map like we did.

With the exception of the hemispherical reflectance and transmittance measurements, none of the methods above are capable of *measuring the phase function*. Elek et al. [2021] hinted that it might be possible with their profile, but our own experiments concluded that the discriminability is not high enough. In our supplemental document, we discuss similarity relations that are the general culprit: there exist equivalent optical properties that give the same material appearance under given conditions.

Methods that are primarily aimed at measuring the phase function are based on measurements of a collimated beam from various angles [Gkioulekas et al. 2013; Leyre et al. 2014], or assume that (α, σ_t)

is already known [Minetomo et al. 2018]. While Gkioulekas et al. [2013] accurately match various phase function shapes of reference materials, not limited to the simple Henyey-Greenstein model, they require rotation platforms, accurate calibrations, and bright collimated illuminants. Our method can only estimate a single-parameter phase function, but is affordable and less complicated.

2.3 Translucency appearance and color 3D printing

While our measurements aim to be objective and give physically meaningful parameters, it is important to note that the human visual system perceives translucency and translucent objects in a bigger context, combining parameters such as lighting direction or object's shape, especially its edges [Fleming and Bühlhoff 2005; Xiao et al. 2014, 2020]. This has led to the important question of how to define and measure translucency in a both physically and perceptually meaningful way [Urban et al. 2019]. Especially in full-color 3D printing, the complex heterogeneous light scattering inside the printouts poses challenges such as texture blurring or inaccurate color reproduction. Methods that aim to counteract these problems and control the printout's appearance can be roughly categorized into two classes: "top-down" phenomenological approaches [Brunton et al. 2018, 2015; Chen and Urban 2021; Urban et al. 2019], and "bottom-up" simulation-based methods [Elek et al. 2017; Nindel et al. 2021; Rittig et al. 2021; Sumin et al. 2019]. The latter works use Monte Carlo simulations to predict the appearance of a given heterogeneous 3D grid of solidified droplets of various base inks, which is then used inside an optimization loop. So far, these approaches relied on RGB measurements, which suffered from the issues discussed in our supplemental document. Our spectral results overcome these issues and are directly applicable and pluggable into the existing pipelines.

3 METHOD

Our method is based on constructing a mapping, or a relation, between optical properties of a given material sample, and light intensities in simple geometrical configurations that are easy to capture with a spectrometer. We have to ensure that the mapping is one-to-one between the optical properties and different measurements, at least for physically meaningful settings. We call the mapping the *appearance map* because it contains every possible material appearance in the given setting. In the following sections, we define and study a three-dimensional and an attendant two-dimensional appearance maps using Monte Carlo simulations, and we explain the actual measurement procedure later in Sec. 4.

Separation of surface and volume light transport. Solid and liquid translucent materials typically have a well-defined boundary where the material starts and ends. For example, a glass of water has an exterior air-to-glass boundary, and an interior glass-to-water boundary. Because light interacts not only with the medium itself, but also its boundary, it is beneficial to treat surface and internal optical properties independently. Our work focuses on volume properties of translucent materials. Hence, throughout this work, we eliminate the influence of the surface as much as possible by assuming that it is a simple, perfectly *smooth dielectric surface layer* governed by Fresnel equations. This assumption is valid both in theory and

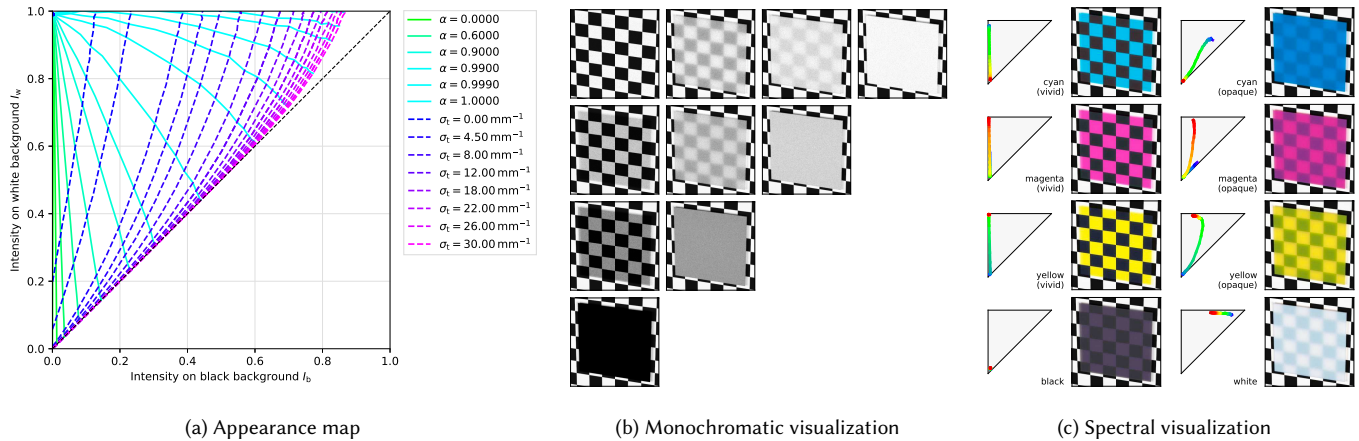


Fig. 2. Our two-dimensional appearance map relates the appearance of a material sample to its optical properties, here assuming the Henyey-Greenstein phase function $g = 0.4$, sample thickness $\ell = 0.5$ mm, and refractive index $n = 1.0$. The black and white backgrounds are assumed to be perfectly diffuse with reflectances 1% and 99%. The contours in (a) are of constant α and σ_t , respectively. Refer to Sec. 3.1 for a detailed analysis. In (b), we show rendered examples with their layout roughly corresponding to their position in the triangle in (a): the left edge contains purely absorbing, non-scattering materials; the top edge purely scattering, non-absorbing materials; and the diagonal perfectly opaque materials. In (c), we expand the visualization spectrally for 8 examples, with the colored curves in the small triangles representing per-wavelength coordinates in the appearance map, from shorter (blue) to longer (red) wavelengths.

in practice: liquids can be placed in a smooth glass container, and solid samples' rough surface can be sanded, polished, or optically smoothed by adding a thin layer of liquid and a microscope glass slide on top (Figs. 8 and 9, and Elek et al. [2021]; Pickering et al. [1993]; Prahl et al. [1993]).

3.1 Material appearance against diffuse backgrounds

A trivial property of translucent materials is that their appearance depends on the background behind them. For example, in Fig. 3, red wine appears red over a white background, but is black over a black background, while a strongly scattering milk remains white in both conditions. Our crucial observation is that there exists a *structured mapping* between the material's optical properties α, σ_t and the




Fig. 3. Three translucent materials in glass cuvettes placed against ideal black and white diffuse backgrounds. Water (left) is virtually non-absorbing and non-scattering, so one can see the backgrounds clearly without any color cast. Red wine (middle) is similar, but strongly absorbs light of wavelengths other than red, giving it a red color cast. Milk (right) is a strongly scattering material, so most of the light is scattered before reaching the backgrounds, which makes it appear white even against the black background.

observed light intensities I_b, I_w against wavelength-independent diffuse black and white backgrounds. We call this mapping the *two-dimensional appearance map* $A_2(I_b, I_w) \rightarrow (\alpha, \sigma_t)$, and we show its example diagram in Fig. 2. With this mapping, one can simply determine the material properties (α, σ_t) from the contours by measuring the intensities (I_b, I_w) on the horizontal and vertical axes, per each wavelength independently, since the diagram is wavelength-independent (assuming the black and white backgrounds are both wavelength-independent reflectors). Generating the mapping was done in the inverse way: we ran Monte Carlo simulations on pairs of (α, σ_t) and acquired the corresponding (I_b, I_w) , which led to the contours. The simulations used the same geometry as in Fig. 7a.

Basic properties. As shown in Fig. 2, this notion of appearance map is only valid in the upper left triangle, as anything below the diagonal would break the principle of energy conservation. The diagonal $I_b = I_w$ represents perfectly opaque materials, as they appear the same regardless of the background. The materials along the diagram's left edge $I_b = 0$ represent non-scattering absorptive materials with $\alpha = 0$, and the intensities along that edge trivially follow Eq. (1): $I_w \propto \exp(-\ell \cdot \sigma_t)$. Likewise, materials towards the top edge $I_w \rightarrow 1$ represent non-absorbing scattering materials with $\alpha \rightarrow 1$. The bottom-left corner $I_b = I_w = 0$ represents a perfectly absorbing material: $\sigma_t \rightarrow \infty, \alpha = 0$. The top-right corner $I_b = I_w = 1$ represents an ideal diffuse reflector: $\sigma_t \rightarrow \infty, \alpha = 1$. The top-left corner represents a perfectly transparent material: $\sigma_t = 0$.

Uniqueness. We conclude that there is a unique pair (α, σ_t) for each pair (I_b, I_w) , hence the appearance map is *one-to-one*, from the fact that the α and σ_t -contours in Fig. 2 always intersect at exactly one particular point. Very importantly, the individual contours are also not self-intersecting, with the following exceptions. The exceptions are the α -contours collapsing at $I_b = 0, I_w = 1$, and the

σ_t -contours collapsing along the diagonal, where $\sigma_t \rightarrow \infty$. This essentially means that there is a large σ_t gradient and uncertainty for highly-absorbing materials, because after a certain threshold, the material absorbs so much light that any further difference becomes negligible. The ambiguity likewise happens for a low α value, where the detected backscattered light intensity gets too low to discriminate the exact α . Also note that the α -contours show exponential spacing, with most of the map covered by $\alpha > 0.9$.

Triangle deformations. The appearance map contours' shapes also depend on other parameters (Fig. 4), e.g., the sample thickness ℓ , the phase function parameter g , or the material refractive index n ,

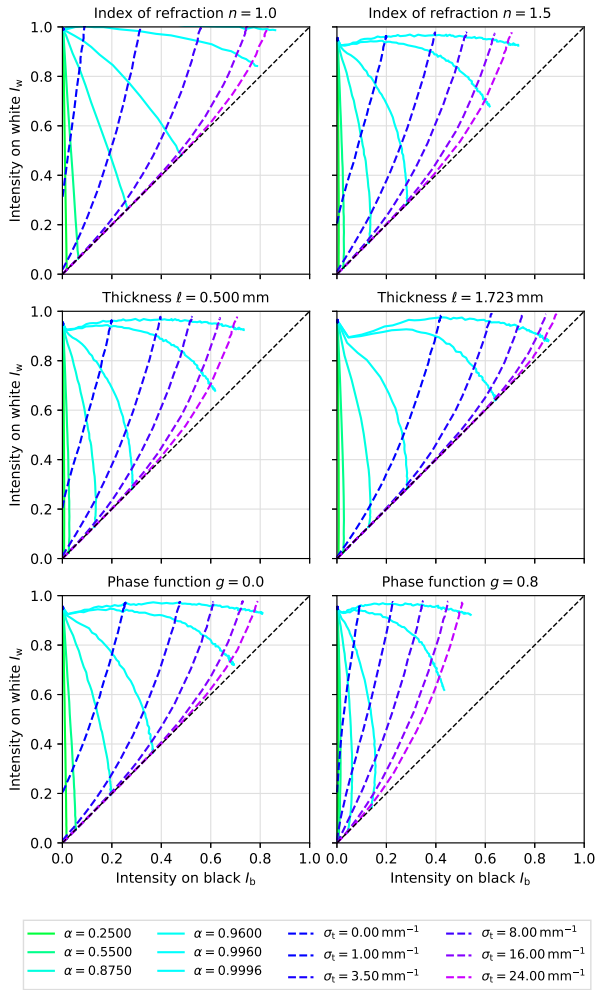


Fig. 4. Visualizing how the appearance map changes with different parameters n , ℓ , g . For example, increasing the refractive index n shrinks the map, as according to Fresnel equations, some energy is simply reflected away from the geometry (Sec. 3.3). Increasing the sample thickness ℓ shifts the σ_t -contours towards the diagonal, as more light is being absorbed in the thick sample. Changing the phase function anisotropy g shrinks the whole map, which we discuss in Sec. 3.2 and show in three dimensions in Fig. 5.




Fig. 5. The g -contours of the three-dimensional appearance map seen from two different angles. The three axes are $(I_b, I_w, \log_{10} I_c)$, the contours correspond to $g \in \{0.0, 0.2, 0.4, 0.6, 0.8\}$. Note how the contours are sufficiently spaced and non-intersecting, except for their collapse to a single line at $I_b = 0$, where no scattering occurs and the phase function plays no role.

which we discuss in detail later (Secs. 3.2 and 3.3). The contours also stretch if the assumed background materials are not ideally diffuse with reflectances 0% for the black, and 100% for the white background, which is impossible to achieve in real measurements. In the measurement procedure (Sec. 4), the true background reflectances need to be modeled in the simulations to prevent a bias.

3.2 Phase function

While the two-dimensional appearance map A_2 allows estimating (α, σ_t) , it assumes that we have chosen a known phase function with a fixed anisotropy g . While the map's contours change with various g (Fig. 4 bottom), there are infinitely many possibilities to which g was the correct one. To pin down this unknown g , we need to include a third observable to match the number of observables to the number of unknowns. We need to account for the similarity relations [Wyman et al. 1989; Zhao et al. 2014] (see our supplemental document), such that this third observable indeed allows for the disambiguation of g , while keeping the setup still affordable.

Three-dimensional appearance map. To disambiguate the phase function anisotropy factor g , we introduce a collimated beam placed at the backside of the material sample (Fig. 7b). The combination of front-illuminating diffuse lights and a back-illuminating collimated beam gives more specific information regarding backward and forward scattering. Together, this allows “breaking out of” the assumptions in the similarity relations. Formally, we extend the previous two-dimensional map A_2 into a *three-dimensional appearance map* $A_3(I_b, I_w, I_c) \rightarrow (\alpha, \sigma_t, g)$, i.e., A_2 is a partial map of A_3 , where I_c is an attenuated intensity of a collimated beam passing through the medium (Fig. 7b). We show an example diagram of A_3 in Fig. 5.

Unscattered transmission. We generate the contours similarly to the two-dimensional case. Even though we used a Monte Carlo simulation also for computing I_c , we discovered that computing it analytically using Eq. (1) has an accuracy indistinguishable from Monte Carlo. This analytical intensity, referred to as *unscattered transmittance* in literature [Prah 2011], seemingly differs from the actual measured intensity for highly-scattering materials ($\alpha \gg 0$) because of light in-scattering to the non-zero solid angle captured by a real sensor. However, as we compare in Fig. 6, the differences




Fig. 6. Transmittance observed with a 4 mm aperture detecting a 1 mm diameter collimated beam passing through a $\ell = 1.8$ mm thick sample of varying (α, σ_t) and a dominantly forward-scattering Henyey-Greenstein phase function $g = 0.8$. Notice how the transmittance of low-scattering materials with a small α follows the Beer's law, Eq. (1). For highly scattering materials, the in-scattering from the beam is superimposed on the detected transmittance, which causes the true measured intensity to be higher than Beer's law. However, this arguably only affects measurements whose dynamic range falls below what an ideal 16-bit sensor could capture. For lower g and thinner samples, this becomes even more negligible.

are not noticeable in our setup, even for a dominantly forward-scattering phase function and a relatively thick sample, though still thin enough compared to the inverse of the extinction coefficient.

Uniqueness. We show that the parameter g can be uniquely found in the A_3 -map. It follows from Fig. 5, as the g -contours are not self-intersecting and they only collapse into a single line at $I_b = 0$, corresponding to a non-scattering material with $\alpha = 0$; and at $I_w = 0$, a strongly absorbing material. It is reasonable, as one cannot measure a scattering phase function of a non-scattering material, or of a material that absorbs all the light it was illuminated with.

3.3 Index of refraction

The material's index of refraction n is the last factor that noticeably influences the appearance map. Our Monte Carlo simulations follow the Fresnel equations governing the reflection and transmission of light from the material boundary. As some incoming light gets completely reflected away, it is expected that the contours shrink towards lower values of I_w for increasing n (see Fig. 4, top). However, it may seem unintuitive that the peak of the top α -contours is not in the top-left corner, but rather the contours form a hill that raises with an increasing σ_t . This is a result of two factors: first, highly-scattering materials with a high α and σ_t back-scatter a lot of light before it even reaches the background boundary; second, our Monte Carlo simulations show that there are significant multiple reflection bounces occurring between the white background and the material boundary, which also explains why the top-left corner for $n = 1.5$ is higher than one would expect from the Fresnel equations alone.

4 MEASUREMENT

We design a simple, affordable physical setup and measurement procedure that follows from our proposed appearance maps A_2 and A_3 (Sec. 3). The end goal is to place a material sample into the

setup (Sec. 4.1), spectrally measure the intensities I_b , I_w , and I_c with a spectrometer (Sec. 4.2), and then interpolate within the appearance map to find the material's corresponding optical properties (α, σ_t, g) independently for each wavelength (Sec. 4.3).

For an affordable setup, we eliminated expensive and specialized equipment as much as possible: we used a simple pocket-size spectrometer, household illuminants, inexpensive electronics, a limited number of optical components, and a custom-made background holder. The calibrations that we require are simple and not time-consuming. In principle, our method is also usable for various samples, including liquids in glass cuvettes.

4.1 Measurement setup overview

We propose the following setup for capturing the intensities I_b , I_w , and I_c , illustrated in Fig. 7. The main components are the *material sample*, two *reflective diffuse backgrounds* (black and white), the *detector part*, and the *illuminants* (diffuse for I_w and I_b , and collimated for I_c). For detailed technical specifications of all the components mentioned in this section, please refer to the supplemental: this lists the exact supplier and part numbers for each component.

Our setup supports quickly switching between three different configurations: behind the sample, there can be a black background, a white background, or a collimated light beam (Fig. 7a, Fig. 7b). Using an inexpensive FDM 3D printer, we printed a custom sliding background holder to change between those three configurations in




Fig. 7. Diagram of our measurement setup, in two different configurations. The first configuration is used to measure I_b and I_w and is composed of the sample over a black or white diffuse background, illuminated by two diffuse illuminants. The reflected light is filtered through a pinhole, collected by a lens, and focused at an optical fiber connected to a spectrometer. The second configuration is used for the phase function estimation: the diffuse illuminants are covered, the background is removed, and the sample is backlit by a beam that was collimated from an LED.

a matter of seconds. The printable STL file for this holder is attached in the supplement. The holder, visible in Fig. 9, is split into three parts: a rectangular hole for collimated measurements, a Spectralon holder for the white background, and a rectangular space for gluing the black background.

In the end, measuring all three I_b , I_w , and I_c can be performed within a minute, with the only manual activity being sliding the background holder twice, and covering the diffuse illuminants for the collimated measurement. Since the intensities are measured spectrally with a spectrometer, the wavelength dependency is implicitly incorporated into each of the three measurements.

Material sample. The sample of the material needs to conform to certain requirements. First of all, samples of materials with a high σ_t have to be thin, otherwise we would detect zero transmitted light I_c . In our case, this was critical for the white ink that approaches $\sigma_t = 25 \text{ mm}^{-1}$ in the blue spectral region. Then, as mentioned in Sec. 3.1, we assume the sample to have a smooth surface with a known refractive index n . For liquids, we recommend using a rectangular glass cuvette. Its geometry has to be modeled in the appearance map simulation to compensate for the additional layers of material. For moderately rough solid samples, such as the 3D printed samples, we used a combination of manual polishing and then layering both sides with a small quantity of index-matched immersion oil and a thin borosilicate-glass microscope cover slide (Fig. 8). The glass slides adhere to the sample with the help of the oil. The sample and the slides can then be washed with soap and reused.

Reflective diffuse backgrounds. Ideally, the backgrounds would be perfectly diffuse materials with a uniform spectral response, with a 100% reflectance (for the white one), and a 0% reflectance (for the black one) over the whole visible range. As such materials do not physically exist, we use the closest commercially available alternatives. We use a 1-inch Spectralon disk with a 99% diffuse reflectance as the white material, but one could also use a much cheaper 92% high-reflectance PTFE sheet. For the absorbing black material, we use a black flocked paper with a reflectance around 2%, which unfortunately substantially increases for $\lambda > 650 \text{ nm}$. Alternatively, a black aluminum foil with a more uniform reflectance around 5% could be used, but it is noticeably more glossy than the flocked paper. Another option would be a larger cavity lined with black material, analogous to a beam dump in laser experiments.

Detector. The light that has transmitted or reflected from the sample first passes through a pinhole, which spatially filters only a small region of interest on the sample. This beam is then collected by a small camera lens mounted on an optical post and centered on the optical path. Finally, the light is focused on the entry of an optical fiber, which delivers it to the spectrometer.

Diffuse illumination. For the diffuse illumination for measuring I_b and I_w , we use two off-the-shelf $28.3 \times 14.1 \text{ cm}$ floodlight LED panels. They are placed at a distance of approximately 30 cm with an orientation of $\pm 45^\circ$ with respect to the sample surface.

Collimated illumination. The illumination for measuring I_c was created by collimating a simple, bright, warm-white LED with a small beam angle. Its light first passes through a diffuser immediately




Fig. 8. Solid materials, such as 3D printed samples, still have a somewhat rough surface even after moderate polishing. Such a rough refractive interface scatters light away from the measurement geometry, which is difficult to calibrate. To compensate, we “sandwich” solid samples between two borosilicate glass microscope cover slides (0.14 – 0.17 mm thick), and a thin layer of index-matched immersion oil. Since all these materials have a very similar index of refraction $n \approx 1.5$, the only significant reflections and refractions occur at the smooth outer layer.

followed by a pinhole, which approximates a point light source. The light is then gathered by a lens that collimates it. We further restrict the beam diameter with a bigger pinhole in order to limit the in-scattering intensity within the sample.

4.2 Measurement procedure

Measuring the material consists of a trivial illuminants calibration, and then measuring the sample against the three backgrounds.

Calibrations. We begin by turning all illuminants on and letting them stabilize and warm up for several minutes. We proceed by measuring the three backgrounds without any sample. This directly gives us the spectrum of the collimated LED, and indirectly gives us the spectrum of the diffuse light (by dividing the measured value by the white background’s uniform reflectance of 99%), and the reflectance of the black background (by dividing the measurement by the diffuse light spectrum acquired in the previous step). In case the illuminants are not perfectly stable, it is useful to perform these calibrations during the measurement of every sample again.

Sample measurement. Each sample is first prepared according to Sec. 4.1: in the case of the 3D printed inks, each sample is polished and then sandwiched with immersion oil and microscope




Fig. 9. Measuring a sample against a white background on our custom 3D printed background holder. Note how the appearance of the rough sample changes under the microscope slide, which optically smoothens the rough interface (see also Fig. 8).

glass slides. The sample is then inserted in a sample holder and centered within the optical path. The measurements can be done subsequently in any order by simply sliding the background holder. When measuring the collimated intensity I_c , we cover the diffuse illuminants by cardboards to make sure the spectrometer does not measure any parasitic diffuse backscattering from the sample.

Automation. While we only performed the described measurements manually, we believe that parts of the process could be easily automated, for example by using a motorized digitally controlled slider for the background holder. Such a setup could then be embedded into a self-contained device, to which a sample is inserted, and it then performs all three measurements fully automatically.

4.3 Fitting the optical properties

Once the relative intensities (I_b, I_w) , or (I_b, I_w, I_c) , are measured, the actual optical properties (α, σ_t) , or (α, σ_t, g) , respectively, have to be fitted via the appearance maps A_2 , or A_3 .

Datasets. We begin by constructing the appropriate appearance map for the given material refractive index n and sample thickness ℓ . As explained in Sec. 3, this is done using Monte Carlo simulations: first, we densely sample the space of possible parameters (α, σ_t, g) , and then we use Monte Carlo rendering to get the corresponding monochromatic intensities (I_b, I_w, I_c) on a single-pixel sensor, according to the geometry in Fig. 7. Note that following the argumentation in Sec. 3.2 and Fig. 6, we can also compute I_c analytically from Eq. (1) without Monte Carlo. The range of parameters that we simulated was inspired by Elek et al. [2021], with low values of σ_t and high values of α simulated with finer steps due to the exponential behavior of light attenuation and scattering. We used a modified Monte Carlo implementation of Nimier-David et al. [2019] and Jakob et al. [2022], which already contained a volumetric path tracer for the diffuse measurements, but we needed to implement a simple volumetric light tracer for the collimated measurements, as such a simulation is impossible with a path tracer.

Fitting. As the dataset points do not form a regular grid (it is impossible to know the spacing of the intensities beforehand), we use a multi-dimensional linear interpolation of an irregular grid. Specifically, we used a SciPy [Virtanen et al. 2020] implementation based on constructing the interpolants through a triangulation of the dataset, and then performing linear barycentric interpolation on each triangle. We experimented with non-linear interpolations, but we were unable to find a solution that would be more robust than the linear one. Note that the interpolation is wavelength-independent (Fig. 2b), which means that spectral measurements are fitted separately per each wavelength (Fig. 2c). For example, our spectrometer had a resolution of about 1000 wavelength bins, so we performed 1000 independent interpolations to fit the spectrum. This operation is very fast, only taking a few second for a full spectrum.

Ensuring robustness. The triangulation and linear interpolation are not stable in certain regions of the appearance map, especially towards the diagonal, where all extinction coefficients collapse into a single line. Hence, naïvely interpolating the optical parameters

from the exact measured values will not be robust, since any additive noise can cause instability of the σ_t estimate along the diagonal.

Our solution is based on *random sampling*: we randomly sample a small ε -neighborhood around the actual measured values (I_b, I_w, I_c) , we interpolate the optical properties for all of these samples, and then we choose the *median* result. We found that using 250 to 1000 random samples in an $\varepsilon \in [0.01, 0.04]$ neighborhood creates sufficiently robust results. Furthermore, we always combine data from at least two *different thicknesses* of the same material. For this, we simply combine the randomly sampled and interpolated values together, and choose the median of them. This also proves to be a useful consistency check, since both thicknesses should ideally result in identical optical parameter estimates. In Fig. 10, we show the results separately for each thickness, and for the combined estimate. It is also useful to visualize the σ_t estimates based on Eq. (2) to verify that the fitting was indeed robust.

5 RESULTS

We now demonstrate and validate our method on 3D printing inks from the Stratasys PolyJet material family. These inks have a wide range of properties, from almost non-scattering color primaries, to a strongly-scattering white material. This variety serves as a good stress test of our approach, and it also results in a wide gamut of appearances realizable by combining these inks together. For example, by combining a transparent yellow with a scattering white, we get an opaque yellow (Fig. 13).

In Sec. 5.1, we apply our measurement procedure to estimate the spectral optical properties of the cyan, magenta, yellow (VEROVIVID), black (BLACKPLUS), and white (PUREWHITE) inks. We also discuss the repeatability of our method, its performance, sources of error, and some practical remarks on processing these samples. In Sec. 5.2, we perform an indirect validation of our results. We 3D print a color checker that contains various mixtures of the base materials, and we compare their real appearance and spectral reflectances to Monte Carlo simulations of the same virtual mixtures, based on the properties estimated in Sec. 5.1.

5.1 Measuring the base inks

We 3D printed two samples of each of the five inks with the dimensions 40×40 mm, and thicknesses 0.4 mm and 0.8 mm. As the printed objects have a rough surface, we polished all samples with sandpapers of grit sizes 400 to 2000, which also reduced the sample thicknesses by about 0.02 mm on average. In theory, one could continue polishing the samples to an almost perfect finish, but that becomes laborious with diminishing returns, so we instead sandwiched the samples in immersion oil and glass slides as in Fig. 8.

After applying our method, we obtain the results in Fig. 10, which include the spectrally resolved single-scattering albedo α , extinction coefficient σ_t , and phase function anisotropy g . The plots feature estimates from both sample thicknesses individually, and a robust combined estimate based on the procedure described in Sec. 4.3. The σ_t plots also contain estimates obtained from Beer's law (Eq. 2), which offers another consistency check.

From these measurements, we can trivially conclude that the color primaries and black are virtually non-scattering materials,

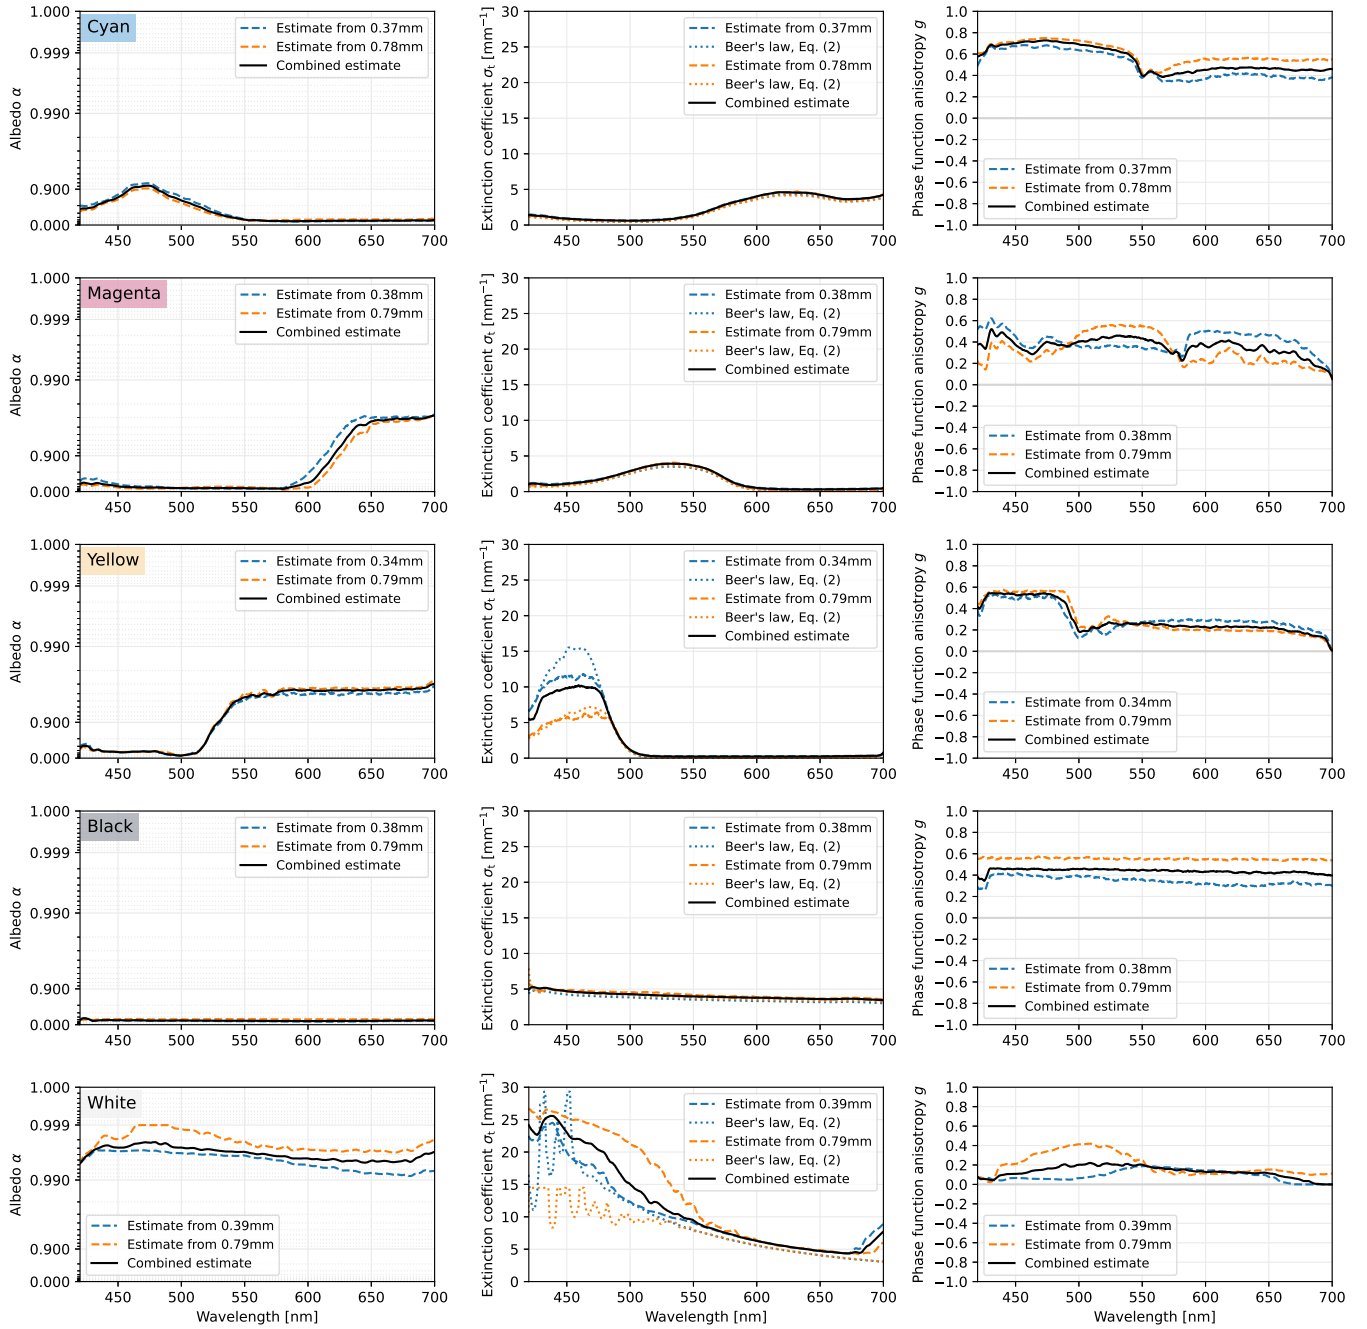


Fig. 10. Fitting the three optical properties (α , σ_t , g) of the base 3D printing materials, specifically, the Stratasys materials VEROVIVID (cyan, magenta, yellow), BLACKPLUS, and PUREWHITE. The thick black curves show a combined estimate from both thicknesses to ensure robustness (Sec. 4.3), while the dashed curves are based on only one thickness. In the σ_t plots, we also visualize the corresponding Beer's law estimates as a consistency check. Note that the α -axis is in an exponential scale, because changes in low α values have a very small effect on the material appearance compared to values above 0.9, as shown by Elek et al. [2017], who numerically fitted a mapping that linearizes the effect, which we use in the axis scaling.

while the white material is strongly scattering. This is in contrast to the older family of Stratasys materials, which had significantly

more scattering color primaries, as measured in RGB channels by Elek et al. [2021].

Sources of error. From Fig. 10, we can see that both the α and σ_t estimates are consistent throughout the materials, with a few exceptions. In the σ_t estimates, we can see inconsistencies in the blue wavelengths (around 400 to 500 nm). We suspect it may be caused by *fluorescence*, a wavelength-shifting effect of absorbing photons of lower wavelengths and re-emitting them at higher wavelengths. Our experiments with violet lasers indicated that wavelength shifting indeed occurs in these materials. This effect is not accounted for in the simulations, it is generally hard to measure, and none of the related methods could measure it either, so it remains as a challenging future work. This may also affect the white albedo, as the thicker sample would contain more fluorescent particles and appear more scattering than the thinner sample. We also observe some inconsistencies in the anisotropy g , however, one has to keep in mind that the estimates do not make physical sense for those wavelengths and materials, where scattering events seldomly occur, i.e., $\alpha(\lambda) \approx 0$ or $\sigma_t(\lambda) \approx 0$.

Repeatability. We assess the repeatability of our method by measuring the cyan ink several times, including different rotations of the sample. This experiment answers whether the samples are sufficiently homogeneous and rotationally invariant. From our results in Fig. 11, we conclude that the results are repeatable enough within the expectations from an affordable measurement method. For curious readers, Fig. 11 also includes a few failed measurements, where the sample was measured closer to its edge, where its thickness changes due to the manual polishing. It indicates the importance of accurately measuring the sample thickness at exactly the point that ends up being in the optical path of the measurement setup, for example using a micrometer screw gauge.

Time performance. The performance of our method can be split into two parts: precomputation of the appearance map, which is only done once per sample thickness, and the actual fitting process per material sample. The precomputation is using Monte Carlo, ideally with a very high sample count to avoid noise (for the final results, we used 640,000 samples per each datapoint). On a single machine, it takes less than a day, and it could be further improved by a significant factor by reusing the Monte Carlo path space samples for various α values. In contrast, the actual fitting process per each material sample is very fast, taking roughly two seconds for the whole spectrum of about 1000 wavelength bins.

5.2 Validation

We validate the measurements of the base inks by 3D printing a custom *color checker* and verifying its appearance. It is a 5 mm thick grid of $8 \times 8 = 64$ different mixtures of the base inks (Figs. 1 and 12). The mixtures correspond to 64 colors, each color being $15 \times 15 \times 5$ mm large and consisting of a different ratio of the base inks. Because 3D printers cannot physically mix the inks, the colors are printed in a 3D halftoned fashion: each color is made of voxels, and each voxel is randomly assigned a base ink such that the overall ink ratio is correct (similarly to mixture halftoning of Nindel et al. [2021]).

Using the optical properties measured in Sec. 5.1, we run a Monte Carlo simulation of this color checker to simulate each square's resulting reflectance. We then take the physically printed color




Fig. 11. Assessing the repeatability of our method by measuring the cyan ink several times, including various rotations. Poorly centered invalid measurements are also included. Refer to the text in Sec. 5.1 for discussion.

checker, and we measure each square with a real spectrometer. The goal is to compare the simulated reflectances to the real measurements, which we show for 20 of the squares in Fig. 13, visualized both with spectral curves, and with the corresponding sRGB colors.

Apart from visualizing the results obtained directly from the simulation, we also show these simulated curves multiplied by *correcting factors*, which were found individually for each square, minimizing the color error. The necessity for these factors reveals that there is a reasonable match in the wavelength dependency (in relative sense), but a mismatch in the overall intensity. The factors attempt to partially correct for geometry misalignments, e.g., because the color checker was manually polished and did not have a perfectly flat surface, and because it was manipulated by hand between every measurement. There are also additional sources of error such as the aforementioned fluorescence, which is unfortunately a wavelength-dependent effect, and the simplified phase function model, which only approximates the real light scattering. We also note that for




Fig. 12. Monte Carlo rendering of a demo scene containing our $120 \times 120 \times 5$ mm large color checker, and five 0.8 mm thick samples of the CMYKW base inks that the color checker is mixed from. The base inks are also located in the bottom row of the checker, leftmost four (CMYK) and the rightmost (W) square. We used a background that shows how some mixtures are significantly more translucent than others. A photograph of this color checker is in Fig. 1, and some of its spectral reflectances are shown in Fig. 13.

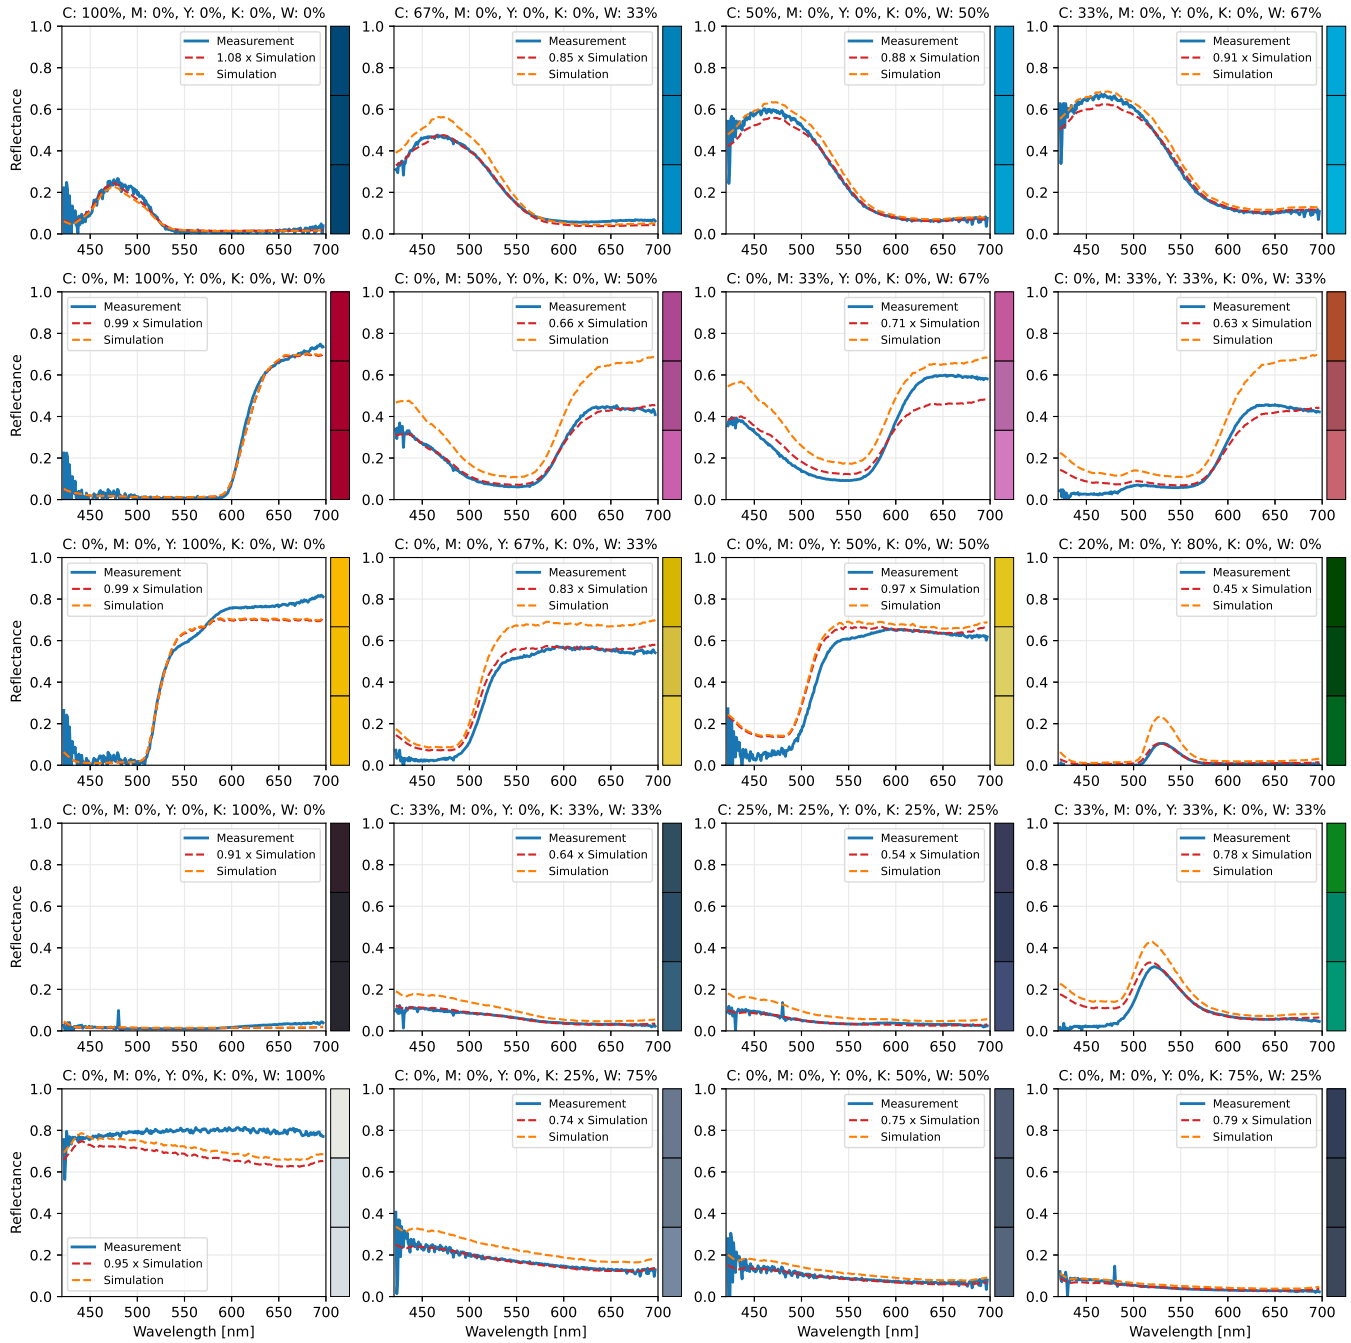


Fig. 13. A comparison between spectral reflectance measurements of a 3D printed color checker (Figs. 1 and 12) that mixes the five base inks (CMYKW) in various ratios, and a corresponding Monte Carlo simulation based on the estimated optical properties from Sec. 5.1, Fig. 10. Each plot also includes a curve with a correcting multiplicative factor (see Sec. 5.2 for discussion). The title of each plot describes the ink ratios corresponding to that measurement. The colored rectangles on the right visualize the (potentially clipped) sRGB colors corresponding to the reflectance spectra, assuming a D65 illuminant. Each top rectangle is the measured color, the middle one is the simulation with a correcting factor applied, and the bottom one is the pure simulation.

this simulation and rendering, our simulation software required

fixing a certain phase function, so we used $g = 0.4$, and the (α, σ_t) were refitted to that assumption.

Results with the base inks. The most critical result in Fig. 13 is that the predictions for the base inks themselves (left column, 100% ratios) are very accurate. Because we measured these inks on 0.4 and 0.8 mm thick samples, but the color checker is 5 mm thick, which is a significantly larger volume, the matching results are an important evidence of correctness. The most noticeable discrepancies are in the white and yellow inks, which we know to be fluorescent, which would explain why the real measurements always had a higher intensity in the longer wavelengths than the non-fluorescent simulations. The slight discrepancy in the black ink could be described by the fact that the absorbing black background used for the measurements in Sec. 5.1 is not perfect, and its reflectance increases around $\lambda > 650$ nm, causing a slight red shift.

Results with the mixtures. The highest accuracy was obtained in the mixtures of the cyan and white ink (top row). That is mainly because the cyan ink absorbs significantly above $\lambda > 550$ nm, so the error in the white ink, which is the highest in that range, is not being propagated to the final color. With all the other mixtures, we have achieved results corresponding to reasonably similar colors, although the exact color hues and saturation were mispredicted, mainly for magenta and green.

Rendering. Since the main use case of our method is Monte Carlo rendering and Monte Carlo based 3D printing optimization pipelines, we demonstrate such result in Fig. 12, a rendering of a demo scene containing the $120 \times 120 \times 5$ mm large color checker, and 0.8 mm thick samples of the base inks. The scene has a checkered background to visualize the translucency of some of the material mixtures.

6 CONCLUSION

We have presented a simple and affordable technique for estimating spectral bulk optical properties of translucent materials in a way that is sufficient for predictive rendering purposes. Our results are directly pluggable into existing rendering pipelines, and 3D printing optimization pipelines based on Monte Carlo simulations. We demonstrated our method on a set of 3D printing materials from the Stratasys PolyJet family, and showed rendering of their various mixtures in a 3D printed object.

The key discovery of our paper was that the bulk optical properties can be estimated from only three simple point measurements, which is simpler than the existing solutions. A byproduct of this is our two-dimensional appearance map, which we believe could also be useful for artists as an intuitive tool for visualizing or editing translucent appearance, because all relevant possible combinations are encoded in the triangle in a perceptually meaningful way.

For the future, we see two additional measurement problems that would benefit from a simple and more affordable solution. First, the acquisition of bi-spectral (fluorescence) optical properties. Second, measuring the index of refraction, which is currently a necessary input to our method. Both problems currently require expensive equipment such as a monochromator or an ellipsometer, which are hardly found outside of specialized laboratories.

ACKNOWLEDGMENTS

This work has received funding from the GA UK project 1165220 of Charles University, Faculty of Mathematics and Physics, and from the Czech Science Foundation under grants numbers GAČR-19-07626S and GAČR-22-22875S, and from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 956585. This work was further supported by the Charles University grant SVV-260588. We thank the reviewers and our "shepherd" for their valuable support while preparing the final revision of this publication. We also thank Wenzel Jakob and his team from École polytechnique fédérale de Lausanne (for hosting a research visit, and collaboration on their newest Mitsuba renderer), and Alexandr Lazarov from Czech Technical University (for 3D printing our samples).

REFERENCES

- 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), 225:1–225:17. <https://doi.org/10.1145/3272127.3275103>
- Alan Brunton, Can Ates Arikian, Tejas Madan Tanksale, and Philipp Urban. 2018. 3D Printing Spatially Varying Color and Translucency. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 37, 4 (July 2018), 157:1–157:13. <https://doi.org/10.1145/3197517.3201349>
- Alan Brunton, Can Ates Arikian, and Philipp Urban. 2015. Pushing the Limits of 3D Color Printing: Error Diffusion with Translucent Materials. *ACM Trans. Graph.* 35, 1 (Dec. 2015), 4:1–4:13. <https://doi.org/10.1145/2832905>
- Danwu Chen and Philipp Urban. 2021. Deep learning models for optically characterizing 3D printers. *Optics Express* 29, 2 (Jan. 2021), 615–631. <https://doi.org/10.1364/OE.410796> Publisher: Optica Publishing Group.
- Eugene D'Eon and Geoffrey Irving. 2011. A quantized-diffusion model for rendering translucent materials. *ACM Transactions on Graphics* 30, 4 (July 2011), 56:1–56:14. <https://doi.org/10.1145/2010324.1964951>
- 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 (Nov. 2017), 1–15. <https://doi.org/10.1145/3130800.3130890>
- Oskar Elek, Ran Zhang, Denis Sumin, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, Jaroslav Křivánek, and Tim Weyrich. 2021. Robust and practical measurement of volume transport parameters in solid photo-polymer materials for 3D printing. *Optics Express* 29, 5 (March 2021), 7568–7588. <https://doi.org/10.1364/OE.406095> Publisher: Optical Society of America.
- Claudia Emde, Robert Buras-Schnell, Arve Kylling, Bernhard Mayer, Josef Gasteiger, Ulrich Hamann, Jonas Kylling, Bettina Richter, Christian Pause, Timothy Dowling, and Luca Bugliaro. 2016. The libRadtran software package for radiative transfer calculations (version 2.0.1). *Geoscientific Model Development* 9, 5 (May 2016), 1647–1672. <https://doi.org/10.5194/gmd-9-1647-2016> Publisher: Copernicus GmbH.
- Roland W Fleming and Heinrich H Bühlhoff. 2005. Low-level image cues in the perception of translucent materials. *ACM Transactions on Applied Perception (TAP)* 2, 3 (2005), 346–382. <https://doi.org/10.1145/1077399.1077409> Publisher: ACM New York, NY, USA.
- J. R. Frisvad, S. A. Jensen, J. S. Madsen, A. Correia, L. Yang, S. K. S. Gregersen, Y. Meuret, and P.-E. Hansen. 2020. Survey of Models for Acquiring the Optical Properties of Translucent Materials. *Computer Graphics Forum* 39, 2 (2020), 729–755. <https://doi.org/10.1111/cgf.14023>
- Ioannis Gkioulekas, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin. 2013. Inverse Volume Rendering with Material Dictionaries. *ACM Trans. Graph.* 32, 6 (Nov. 2013), 162:1–162:13. <https://doi.org/10.1145/2508363.2508377>
- Richard C. Haskell, Lars O. Svaasand, Tsong-Tseh Tsay, Ti-Chen Feng, Bruce J. Tromberg, and Matthew S. McAdams. 1994. Boundary conditions for the diffusion equation in radiative transfer. *Journal of the Optical Society of America A* 11, 10 (Oct. 1994), 2727. <https://doi.org/10.1364/JOSAA.11.002727>
- L. G. Henyey and J. L. Greenstein. 1941. Diffuse radiation in the Galaxy. *Astrophysical Journal* 93 (Jan. 1941), 70–83. <https://doi.org/10.1086/144246>
- Steven L. Jacques. 2013. Optical properties of biological tissues: a review. *Physics in Medicine & Biology* 58, 11 (May 2013), R37. <https://doi.org/10.1088/0031-9155/58/11/R37> Publisher: IOP Publishing.
- Wenzel Jakob, Eugene d'Eon, Otto Jakob, and Steve Marschner. 2014. A comprehensive framework for rendering layered materials. *ACM Transactions on Graphics* 33, 4 (July 2014), 1–14. <https://doi.org/10.1145/2601097.2601139>

- Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, and Delio Vicini. 2022. DR.JIT: a just-in-time compiler for differentiable rendering. *ACM Transactions on Graphics* 41, 4 (July 2022), 1–19. <https://doi.org/10.1145/3528223.3530099>
- Wojciech Jarosz. 2008. *Efficient Monte Carlo Methods for Light Transport in Scattering Media*. PhD Thesis. UC San Diego.
- Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. 2001. A practical model for subsurface light transport. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques - SIGGRAPH '01*. ACM Press, 511–518. <https://doi.org/10.1145/383259.383319>
- James T. Kajiya. 1986. The rendering equation. *ACM SIGGRAPH Computer Graphics* 20, 4 (Aug. 1986), 143–150. <https://doi.org/10.1145/15886.15902>
- Paul Kubelka and Franz Munk. 1931. Ein Beitrag zur Optik der Farbanstriche. *Zeitschrift für Technische Physik* 12 (1931), 593–601.
- Sven Leyre, Youri Meuret, Guy Durinck, Johan Hofkens, Geert Deconinck, and Peter Hanselaer. 2014. Estimation of the effective phase function of bulk diffusing materials with the inverse adding-doubling method. *Applied Optics* 53, 10 (April 2014), 2117. <https://doi.org/10.1364/AO.53.002117>
- Yuki Minetomo, Hiroyuki Kubo, Takuya Funatomi, Mikio Shinya, and Yasuhiro Mukaigawa. 2018. Acquiring non-parametric scattering phase function from a single image. *Computational Visual Media* 4, 4 (Dec. 2018), 323–331. <https://doi.org/10.1007/s41095-018-0122-z>
- 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 (July 2006), 1003–1012. <https://doi.org/10.1145/1141911.1141986>
- Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. 2019. Mitsuba 2: a retargetable forward and inverse renderer. *ACM Trans. Graph* 38, 6 (2019), 203:1–203:17. <https://doi.org/10.1145/3355089.3356498>
- Thomas Klaus Nindel, Tomáš Iser, Tobias Rittig, Alexander Wilkie, and Jaroslav Krivánek. 2021. A Gradient-Based Framework for 3D Print Appearance Optimization. *ACM Trans. Graph.* 40, 4 (July 2021). <https://doi.org/10.1145/3450626.3459844> Place: New York, NY, USA Publisher: Association for Computing Machinery.
- James H Nobbs. 1985. Kubelka–Munk Theory and the Prediction of Reflectance. *Review of Progress in Coloration and Related Topics* 15, 1 (1985), 66–75. <https://doi.org/10.1111/j.1478-4408.1985.tb03737.x>
- Jan Novák, Iliyan Georgiev, Johannes Hanika, Jaroslav Krivánek, and Wojciech Jarosz. 2018. Monte Carlo Methods for Physically Based Volume Rendering. In *ACM SIGGRAPH 2018 Courses (SIGGRAPH '18)*. ACM, New York, NY, USA, 14:1–14:1. <https://doi.org/10.1145/3214834.3214880>
- Marios Papas, Christian Regg, Wojciech Jarosz, Bernd Bickel, Philip Jackson, Wojciech Matusik, Steve Marschner, and Markus Gross. 2013. Fabricating translucent materials using continuous pigment mixtures. *ACM Transactions on Graphics* 32, 4 (July 2013), 1–12. <https://doi.org/10.1145/2461912.2461974>
- John W. Pickering, Christian J. M. Moes, H. J. C. M. Sterenborg, Scott A. Prah, and Martin J. C. van Gemert. 1992. Two integrating spheres with an intervening scattering sample. *JOSA A* 9, 4 (April 1992), 621–631. <https://doi.org/10.1364/JOSA.9.000621> Publisher: Optical Society of America.
- John W. Pickering, Scott A. Prah, Niek van Wieringen, Johan F. Beek, Henricus J. C. M. Sterenborg, and Martin J. C. van Gemert. 1993. Double-integrating-sphere system for measuring the optical properties of tissue. *Applied Optics* 32, 4 (Feb. 1993), 399–410. <https://doi.org/10.1364/AO.32.000399> Publisher: Optical Society of America.
- Scott Prah. 2011. Everything I think you should know about Inverse Adding-Doubling. (March 2011), 74. <https://omlc.org/software/iad/manual.pdf>
- Scott A. Prah. 1995. The Adding-Doubling Method. In *Optical-Thermal Response of Laser-Irradiated Tissue*, Ashley J. Welch and Martin J. C. Van Gemert (Eds.). Springer US, Boston, MA, 101–129. https://doi.org/10.1007/978-1-4757-6092-7_5
- Scott A. Prah, Martin J. C. van Gemert, and Ashley J. Welch. 1993. Determining the optical properties of turbid media by using the adding–doubling method. *Applied Optics* 32, 4 (Feb. 1993), 559. <https://doi.org/10.1364/AO.32.000559>
- Alina Pranovich, Sasan Gooran, Jeppe Revall Frisvad, and Daniel Nyström. 2021. Optical properties and appearance of fused deposition modelling filaments. *Advances in Printing and Media Technology - Printing in the Digital Era* (2021), 134–140. Publisher: International Association of Research Organizations for the Information, Media and Graphic Arts Industries.
- Tobias Rittig, Denis Sumin, Vahid Babaei, Piotr Didyk, Alexey Voloboy, Alexander Wilkie, Bernd Bickel, Karol Myszkowski, Tim Weyrich, and Jaroslav Krivánek. 2021. Neural Acceleration of Scattering-Aware Color 3D Printing. *Computer Graphics Forum* 40, 2 (2021), 205–219. <https://doi.org/10.1111/cgf.14262>
- Christopher Sandoval and Arnold D. Kim. 2015. Extending generalized Kubelka–Munk to three-dimensional radiative transfer. *Applied Optics* 54, 23 (Aug. 2015), 7045. <https://doi.org/10.1364/AO.54.007045>
- Sven Schröder, Alexander von Finck, and Angela Duparré. 2015. Standardization of light scattering measurements. *Advanced Optical Technologies* 4, 5–6 (Oct. 2015), 361–375. <https://doi.org/10.1515/aot-2015-0041> Publisher: De Gruyter.
- 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), 111:1–111:14. <https://doi.org/10.1145/3306346.3322992>
- Valery V. Tuchin. 1993. Laser Light Scattering in Biomedical Diagnostics and Therapy. *Journal of Laser Applications* 5, 2 (Oct. 1993), 43–60. <https://doi.org/10.2351/1.4745330> Publisher: Laser Institute of America.
- Philipp Urban, Tejas Madan Tanksale, Alan Brunton, Bui Minh Vu, and Shigeki Nakauchi. 2019. Redefining A in RGBA: Towards a Standard for Graphical 3D Printing. *ACM Trans. Graph.* 38, 3 (June 2019), 21:1–21:14. <https://doi.org/10.1145/3319910>
- H. C. van de Hulst. 1980. *Multiple light scattering: tables, formulas, and applications*. Vol. 1. Academic Press, New York.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, Ilhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, António H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Tim Weyrich, Wojciech Matusik, Hanspeter Pfister, Bernd Bickel, Craig Donner, Chien Tu, Janet McAndless, Jinho Lee, Addy Ngan, Henrik Wann Jensen, and Markus Gross. 2006. Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics* 25, 3 (July 2006), 1013–1024. <https://doi.org/10.1145/1141911.1141987>
- Douglas R Wyman, Michael S Patterson, and Brian C Wilson. 1989. Similarity relations for anisotropic scattering in Monte Carlo simulations of deeply penetrating neutral particles. *J. Comput. Phys.* 81, 1 (March 1989), 137–150. [https://doi.org/10.1016/0021-9991\(89\)90067-3](https://doi.org/10.1016/0021-9991(89)90067-3)
- Bei Xiao, Bruce Walter, Ioannis Gkioulekas, Todd Zickler, Edward Adelson, and Kavita Bala. 2014. Looking against the light: how perception of translucency depends on lighting direction. *Journal of Vision* 14, 3 (March 2014), 17. <https://doi.org/10.1167/14.3.17>
- Bei Xiao, Shuang Zhao, Ioannis Gkioulekas, Wenyan Bi, and Kavita Bala. 2020. Effect of geometric sharpness on translucent material perception. *Journal of vision* 20, 7 (2020), 10–10. <https://doi.org/10.1167/jov.20.7.10> Publisher: The Association for Research in Vision and Ophthalmology.
- Li Yang and Roger David Hersch. 2008. Kubelka-Munk Model for Imperfectly Diffuse Light Distribution in Paper. *Journal of Imaging Science and Technology* 52 (May 2008). [https://doi.org/10.2352/J.ImagingSci.Technol.\(2008\)52:3\(030201\)](https://doi.org/10.2352/J.ImagingSci.Technol.(2008)52:3(030201))
- 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 (Aug. 2018), 1–14. <https://doi.org/10.1145/3197517.3201321>
- Shuang Zhao, Ravi Ramamoorthi, and Kavita Bala. 2014. High-order similarity relations in radiative transfer. *ACM Transactions on Graphics* 33, 4 (July 2014), 104:1–104:12. <https://doi.org/10.1145/2601097.2601104>

The inserted publication ends here.

5 | Our measurement method for fluorescent materials

The second journal publication inserted in this thesis, Iser et al. [2023], focuses on fluorescent materials, whose fundamental theory we covered in Section 2.3. Fluorescence is common in both natural (minerals, corals, tissues) and synthetic (pigments in papers, textiles, plastics) objects, it is often used in optical brighteners, and it plays a crucial role in material appearance. Characterizing and measuring such materials is hence important for predictive rendering [Johnson and Fairchild, 1999, Hullin et al., 2010, Mojzík et al., 2018, Jung et al., 2019, 2020, König et al., 2020, Hua et al., 2022]. Other uses include remote sensing of vegetation [Mohammed et al., 2019] and fluorescence imaging in biology and medicine, including diagnostics in animals and humans [Choyke and Kobayashi, 2012], plants [Delalieux et al., 2009], and coral reefs [Roth and Deheyn, 2013].

Problem A seemingly trivial solution exists for measuring fluorescence: a sample is illuminated by monochromatic light, the whole reflected spectrum is measured, and this process is repeated for each illumination wavelength [Donaldson, 1954]. However, there are several downsides: it needs repeating for each wavelength, it requires an expensive monochromator with a powerful broadband light source, and the monochromatic light is weak, leading to even weaker fluorescent emission and substantial measurement noise. Alternative solutions are based on sparse measurements, but they often rely on expensive equipment, can only support one fluorophore, or they otherwise suffer from accuracy issues.

Our contribution and solution Our core question is whether one could use only a few sparse measurements with low-cost equipment to robustly estimate the material’s fluorescence with high accuracy. This is a high-dimensional inverse problem with infinitely many ambiguous solutions (Section 3.2.2), so finding the correct Donaldson matrix is challenging. Several previous approaches use a dataset of reflectances and fluorophores to reduce the dimensionality [Blasinski et al., 2020, Zheng et al., 2014, Lam and Sato, 2013, Zheng et al., 2015, Fu et al., 2016]. Instead, we represent the fluorescence with a *multivariate Gaussian mixture model* and the reflectance with a *bounded MESE* (maximum entropy spectral estimate), which, to our knowledge, is their first application in such measurements. It parametrizes and constrains the estimate in a robust and simple way, allowing the use of gradient-descent optimization (Section 3.3.4). We demonstrate how our method adapts to different optical setups that anyone can easily build from low-cost illuminants and various possible detectors, including a standard spectrometer or a low-cost multispectral sensor. We evaluate all on a combination of real and synthetic data and compare to state of the art, showing that we reach consistently lower errors and avoid the oscillations and ripples in the estimates.

The full publication is inserted and begins on the next page.

The publication is published as open access and licensed under the Optica Open Access Publishing Agreement. The full citation follows.

Tomáš Iser, Loïc Lachiver, and Alexander Wilkie. Affordable method for measuring fluorescence using Gaussian distributions and bounded MESE. *Optics Express*, 31(15):24347–24362, July 2023. doi: 10.1364/OE.495459. URL <https://opg.optica.org/oe/abstract.cfm?URI=oe-31-15-24347>. Publisher: Optica Publishing Group



Affordable method for measuring fluorescence using Gaussian distributions and bounded MESE

TOMÁŠ ISER, *  LOÏC LACHIVER, AND ALEXANDER WILKIE

Charles University, Faculty of Mathematics and Physics, Malostranské nám. 25, Prague, 11800, Czech Republic

*tomas@cgg.mff.cuni.cz

Abstract: We present an accurate and low-cost method for measuring fluorescence in materials. Our method outputs an estimate of the material's *Donaldson matrix*, which is a commonly used two-dimensional spectral characterization of its fluorescence and reflectance properties. To find the estimate, only a few measurements of the material's reflectance under a few illuminants are needed, which we demonstrate using low-cost optical components. Internally, our algorithm is based on representing each Donaldson matrix with a *multivariate Gaussian mixture model* and its diagonal with a *bounded MESE* (maximum entropy spectral estimate). It parametrizes and constrains the estimate in a robust and simple way, allowing the use of gradient-descent optimization. We evaluate our algorithm on a combination of real and synthetic data, and four examples of distinct optical components. We reach significantly lower errors than the current state of the art on the exact same inputs, our estimates do not suffer from artifacts such as oscillations of the spectra, and they are stable and robust.

© 2023 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

1. Introduction

When a material is illuminated, a fraction of the incoming photons is reflected. The ratio between the reflected and incoming photons is called the material's *surface reflectance*. In general, the reflectance varies per wavelength, which gives materials their various colors. Materials may also absorb and re-emit a fraction of the incoming photons at a different wavelength than they originally had, which is called *wavelength shifting*. For instance, minerals such as fluorite or calcite will glow in color even when illuminated only with invisible ultraviolet light. This effect is called *fluorescence*, and it is immediate and observable with bare eyes.

Fluorescence is common in both natural (minerals, corals, tissues) and synthetic (pigments in papers, textiles, plastics) objects. Distinct *fluorophores* (chemicals) can also be combined in a single object. Due to the re-emission of incoming photons into longer wavelengths, fluorescent materials can appear brighter and more saturated than their non-fluorescent counterparts, and their color or reflectance may significantly change under various illuminants (Fig. 1). This has been studied in computer graphics to allow accurate visualizations and color rendering of such objects [1–7], but also in remote sensing of vegetation [8], and in fluorescence imaging in biology and medicine, including diagnostics in animals and humans [9], plants [10], and coral reefs [11].

All of these methods rely on measuring the fluorescence in the materials. A simple solution is shining monochromatic light on the sample, measuring the whole reflected spectrum, and repeating this for each illumination wavelength, e.g., in 10 nm steps (Sec. 2.1; [12]; partially illustrated in Fig. 1(bc)). It has several downsides: it needs repeating for each wavelength, it requires an expensive monochromator with a powerful broadband light source, and the monochromatic light is weak leading to even weaker fluorescent emission and substantial measurement noise.

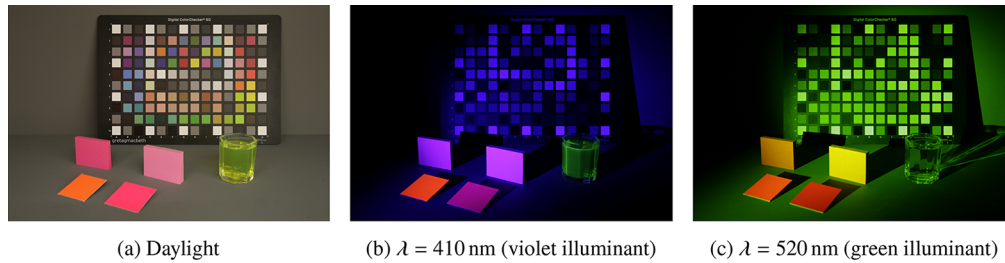


Fig. 1. Photographs of fluorescent (sticky notes, liquid detergent) and non-fluorescent (color chart) materials under different illuminants. The monochrome illuminants reveal which objects are fluorescent: they emit different wavelengths than the illuminants had.

Alternative solutions are based on *sparse measurements* (Sec. 2.2): instead of scanning all combinations, only a few illuminants are used and the reflectances may be measured by a sensor with a coarse resolution. The fluorescence is then estimated by solving an *inverse problem*, i.e., by numerically finding which fluorescence could have resulted in the observed reflectances. Some of these approaches still require expensive equipment such as programmable filters [13] or high-frequency spectral illuminants [14], only assume RGB data [15], or can inherently only support one fluorophore in each material [16–20]. The most important work in our context is by Blasinski et al. [21], who not only summarize the previous approaches, but also unify and generalize several of them into a new cohesive framework. Their method can estimate fluorescence, even from multiple fluorophores simultaneously, using any arbitrary illuminant spectra and sensor spectral sensitivities, which makes it compatible with low-cost optical setups.

However, their estimates can be significantly improved, which we show both quantitatively and qualitatively in our evaluations, where our new method outperforms their results on the exact same inputs. For example, their estimated spectra often suffer from oscillations and ripples: a flat spectrum is reconstructed as a wavy spectrum oscillating around the ground truth value, or a single peak is reconstructed as multiple peaks instead. Similar problems are common in methods that represent the spectra with basis functions that were built from fluorophore and reflectance datasets [17–21]. In our method, we represent fluorescence in a different way, using Gaussian distributions and maximum entropy spectral estimates, which we picked to ensure that the estimated fluorescence is accurate, including fluorescent peaks and potential flat spectra.

Our contribution We develop a robust algorithm for accurately estimating fluorescence and reflectance from sparse measurements acquirable with low-cost optical setups. Our algorithm is based on solving an inverse problem. While several previous approaches use a dataset of reflectances and fluorophores to reduce the dimensionality of the problem [17–21], we instead represent the fluorescence with a *multivariate Gaussian mixture model* and the reflectance with a *bounded MESE* (maximum entropy spectral estimate), which, to our knowledge, is their first application in such measurements. We show that our approach is simple, yet accurate and robust, and even though the inverse problem is non-convex, it behaves well with our chosen gradient-descent optimization strategy. We demonstrate how our method adapts to optical setups of different equipment costs. We evaluate all on a combination of real and synthetic data and compare to state of the art, showing consistently lower errors and higher-quality spectra.

2. Problem statement and prior work

The reflectance from fluorescent materials can be mathematically modeled (Sec. 2.1) and the measurement problem can then be written and solved as an inverse problem (Secs. 2.2 and 2.3).

2.1. Modeling fluorescence with the Donaldson matrix

Assuming fixed illumination and observation angles, when a fluorescent material is illuminated by a light source with spectrum $I(\lambda_i)$, then the reflected intensity $r(\lambda_o)$ is (Fig. 2):

$$r(\lambda_o) = \int \Phi(\lambda_i, \lambda_o) I(\lambda_i) d\lambda_i, \quad (1)$$

where $\Phi(\lambda_i, \lambda_o)$ is a two-dimensional function jointly describing the material's reflectance and fluorescence, and λ_i and λ_o are incoming (excitation) and outgoing (reflected, emitted) wavelengths, respectively. For practical purposes, we discretize the spectral dimension into N_i incoming and N_o outgoing wavelengths, and then $I \in \mathbf{R}^{N_i}$ and $r \in \mathbf{R}^{N_o}$ become vectors, $\Phi \in \mathbf{R}^{N_o \times N_i}$ becomes a matrix, and Eq. (1) becomes a matrix-vector multiplication (Fig. 2):

$$r = \Phi \cdot I. \quad (2)$$

Finding the exact values of the Φ matrix for the given material is our goal. The matrix is often called the *Donaldson matrix* after Donaldson [12], who measured the matrix values directly using monochromatic light (as introduced in Sec. 1). Note that the matrix exists also for non-fluorescent materials: it is simply zero everywhere except for the diagonal ($\lambda_i = \lambda_o$), which represents pure reflectance. As Fig. 2 shows, it is natural to separate the diagonal Φ_{diag} (*pure reflectance*) and the off-diagonal Φ_{fluo} (*pure fluorescence*), and write $\Phi = \Phi_{\text{diag}} + \Phi_{\text{fluo}}$.




Fig. 2. Illustration of Eqs. (1) and (2). The reflected spectrum (a) is computed from the Donaldson matrix (b) and illuminant spectrum (c). While the illuminant is green, the reflection is orange with two peaks, which results from fluorescence.

Sometimes, assumptions are made to simplify the shape of Φ [7,12,21]. First, we expect the emitted photons to have longer wavelengths than the excitation, hence $\Phi(\lambda_i, \lambda_o) = 0$ for $\lambda_i > \lambda_o$, or equivalently, the matrix is zero below the diagonal. Second, by following Kasha's rule, Φ_{fluo} would be a separable function $\Phi_{\text{fluo}}(\lambda_i, \lambda_o) = \phi_{\text{em}}(\lambda_i) \cdot \phi_{\text{ex}}(\lambda_o)$ with one-dimensional emission $\phi_{\text{em}}(\lambda_i)$ and excitation $\phi_{\text{ex}}(\lambda_o)$ spectra [21]. But this does not hold in general, especially for materials with multiple fluorophores, so we do not assume such separability in this publication.

2.2. Sparse measurements

As measuring the Donaldson matrix Φ directly is expensive, slow, and suffers from noise, we focus on the idea of estimating Φ from only *sparse measurements* (Fig. 3) that are typically faster to acquire and do not always rely on expensive components [13–21]. In general, such a measurement setup consists of:

- a set of n_i illuminants with spectra $I_1(\lambda), \dots, I_{n_i}(\lambda)$, and
- a detector of the reflected light, which can detect n_s spectral channels with spectral sensitivities $s_1(\lambda), \dots, s_{n_s}(\lambda)$.

In our method, we impose no implicit restrictions. The detector can be an RGB camera with only 3 channels, as well as a spectrometer with a sub-nanometer resolution and hundreds of channels. The illuminants can be individual LEDs, but also a single broadband lamp with switchable filters.




Fig. 3. Illustration of Eq. (3) (sparse measurements). In this example, the Donaldson matrix (a) is unknown, and the sparse measurements are based on 8 illuminants (b) and an 8-channel sensor (c), resulting in 8×8 detected intensity values $I_{j,k}$ (d).

By discretizing the problem and using the notation above, we can write that when the fluorescent sample is illuminated by the j -th illuminant, then the k -th channel of the detector should detect the reflected intensity of (Fig. 3):

$$(\Phi \cdot l_j) \odot s_k, \quad (3)$$

where \odot is an element-wise multiplication (Hadamard product). In order to estimate an unknown Φ , we can hence illuminate the sample once with each illuminant and detect the reflected intensities per channel, which gives $n_l \times n_s$ measurements denoted as $I_{j,k}$. Estimating the Donaldson matrix $\hat{\Phi}$ then becomes an inverse optimization problem, for which we chose the Euclidean distance (ℓ^2 norm) between the actual measurements and their estimation:

$$\hat{\Phi} = \underset{\Phi}{\operatorname{argmin}} \sqrt{\sum_j^{n_l} \sum_k^{n_s} (I_{j,k} - (\Phi \cdot l_j) \odot s_k)^2}. \quad (4)$$

In general, this problem is significantly underdetermined and has infinitely many solutions, most of which are not even physically plausible. Therefore, a good method needs not only to minimize the error, but mainly converge to a realistic and plausible solution.

2.3. Previous approaches finding a solution of Eq. (4)

To simplify the space of possible solutions of Eq. (4), many approaches use basis functions derived from a database of a priori known reflectances and fluorophores and essentially interpolating between them [17–21]. As we explained in Sec. 1 and show in Sec. 4.5, fitting the spectra onto the small linear bases results in artifacts such as a flat spectrum being reconstructed with oscillations.

From the broader perspective, Blasinski et al. [21] categorized the existing sparse approaches into two groups: *bispectral separation* [14,16–19] and *computational separation* [13,15,20,21]. Using our notation from Sec. 2.2, we could say that bispectral separation methods rely on high-resolution spectral detectors, meaning $n_s > 30$. Such tailored methods can operate under a very small number of illuminants n_l , e.g., two broadband illuminants [16], two high-frequency illuminants [14], or even just one spiky illuminant [18]. On the other hand, the computational separation methods employ more complex algorithms to allow cheaper detectors, e.g., just a simple RGB camera with two broadband illuminants [15] or a set of narrowband illuminants [20]. Using this terminology, our method would certainly fall into computational separation.

Out of all of these methods, only these by Suo et al. [13], Blasinski et al. [21], and ours are general enough to resolve materials with more than one fluorophore, mainly because they do not

build on the assumption of Kasha's rule and separability (Sec. 2.1). The algorithm of Blasinski et al. reaches higher accuracy than Suo et al., it employs and unifies the concepts from the other methods, and it is the most recent method. For the purpose of comparisons, we refer to it as the *state of the art*.

3. Estimating fluorescence with Gaussian distributions and bounded MESE

We design a sparse measurement algorithm that finds a solution to Eq. (4). It ensures that the Donaldson matrix is smooth and plausible by parametrizing it via Gaussian mixtures and bounded MESE, and it finds the final estimate using gradient-descent optimization.

3.1. Overview

Input The inputs of our algorithm are sparse measurements of the material following Fig. 3 and Sec. 2.2. The inputs include the spectra $l_1(\lambda), \dots, l_{n_l}(\lambda)$ of the n_l illuminants, the spectral sensitivities $s_1(\lambda), \dots, s_{n_s}(\lambda)$ of the detector's n_s spectral channels, and the $n_l \times n_s$ detected values of the material's reflectances corresponding to $I_{j,k}$ in Eq. (4) and Fig. 3(d).

Output The output of our algorithm is an estimate of the material's Donaldson matrix Φ . Because it is represented parametrically in our algorithm, it has a small memory footprint, and a fully continuous signal Φ can be easily reconstructed with any spectral resolution.

Processing We split the estimate $\Phi = \Phi_{\text{diag}} + \Phi_{\text{fluo}}$ into the off-diagonal elements Φ_{fluo} (Fig. 4(ab)) and the diagonal Φ_{diag} (Fig. 4(c)). The fluorescence Φ_{fluo} is represented by a multivariate Gaussian mixture model (GMM), which consists of individual two-dimensional Gaussian distributions parametrized by their intensity, mean, and covariance (Sec. 3.2). The diagonal Φ_{diag} is represented by a small set of trigonometric moments that define a continuous signal given by the bounded maximum entropy spectral estimate (bounded MESE) (Sec. 3.3). Using Eq. (4) then allows estimating the whole matrix via gradient-descent optimization (Sec. 3.4).




Fig. 4. Illustration of our parametric model. The Gaussian mixture model (GMM) consists of independent distributions. Their contours and mean values (\times symbols) are visualized in (a), and their superposition Φ_{fluo} in (b). The pure reflectance Φ_{diag} (c) is visualized with its discrete samples (blue) together with the final bounded MESE (red).

3.2. Fluorescence as a Gaussian mixture

A two-dimensional Gaussian mixture model (GMM) is a linear superposition of Gaussian distributions (Fig. 4(ab)). Following Hua et al. [7], we define it as a weighted sum of n_g

individual Gaussians:

$$f(\mathbf{x}) = \sum_{m=1}^{n_g} w_m \mathcal{N}(\mathbf{x} \mid \mu_m, \Sigma_m), \quad (5)$$

where $\mathcal{N}(\mathbf{x} \mid \mu, \Sigma): \mathbf{R}^2 \rightarrow \mathbf{R}$ is a two-dimensional Gaussian distribution:

$$\mathcal{N}(\mathbf{x} \mid \mu, \Sigma) = \left(2\pi\sqrt{|\Sigma|}\right)^{-1} \exp\left(-\frac{1}{2}(\mathbf{x} - \mu)^T \Sigma^{-1}(\mathbf{x} - \mu)\right), \quad (6)$$

and $w_m > 0$, $\mu_m \in \mathbf{R}^2$, and $\Sigma_m \in \mathbf{R}^{2 \times 2}$ are the weight, mean, and covariance matrix of the m -th distribution in the mixture. In our context, the GMM and Eq. (6) are interpreted to be in the wavelength domain, meaning $\mathbf{x} = (\lambda_i, \lambda_o)$ is a vector of the incoming and outgoing wavelengths, and $\mu = (\mu_i, \mu_o)$ represents the mean incoming and outgoing wavelengths. The pure fluorescence of the Donaldson matrix is then defined as $\Phi_{\text{fluo}}(\lambda_i, \lambda_o) = f(\mathbf{x})$.

This concept follows the recent publication of Hua et al. [7], who worked on compressing fluorescent textures in photorealistic rendering. They observed that the fluorescence distribution's shape resembles Gaussian distributions and it can be modeled as a sum of a few Gaussians. We use this observation in the new context of measurements as it allows using only a few parameters to control the fluorescence while also ensuring that the result will be physically plausible to a certain degree. Furthermore, the estimate naturally supports multiple fluorophores, it is continuous and can be evaluated at any wavelength resolution, and it is directly compatible with efficient importance sampling in Monte Carlo rendering [7]. We observed that the GMM is a smooth, differentiable function that behaves well in the gradient-descent optimization setting.

3.3. Diagonal as a bounded MESE

We represent the diagonal Φ_{diag} as a one-dimensional bounded maximum entropy spectral estimate (bounded MESE, Fig. 4(c)). Following Peters et al. [22], we know that a bounded 2π -periodic signal $g(\varphi) \in [0, 1]$ can be represented by $m + 1$ complex Fourier coefficients $\mathbf{c} = (c_0, \dots, c_m)$, which can also be viewed as the signal's trigonometric moments:

$$\mathbf{c} = \int_{-\pi}^{\pi} g(\varphi) \mathcal{F}(\varphi) d\varphi \in \mathbf{C}^{m+1}, \quad (7)$$

where $\mathcal{F}(\varphi)$ is the Fourier basis:

$$\mathcal{F}(\varphi) = \frac{1}{2\pi} (\exp(-ij\varphi))_{j=0}^m \in \mathbf{C}^{m+1}. \quad (8)$$

With the whole signal $g(\varphi)$ reduced to only $m + 1$ complex coefficients, we also need a way to solve the inverse problem: finding a signal that corresponds to the given coefficients. Broadly speaking, MESE [22,23] is the result of such an inverse process: it is an estimated signal that best corresponds to the given coefficients \mathbf{c} , and by "best" we mean that it maximizes Burg entropy, so the estimate is a relatively smooth signal. This is a memory compression mechanism, as a whole curve is represented by only a few numbers, but it also acts as a smoothing operator, because after computing the first few moments of a noisy signal, its MESE will be a smoother signal (Fig. 4(c), compare the blue samples to the red curve).

Spectral reflectances, in our case $\Phi_{\text{diag}}(\lambda) \in [0, 1]$, are generally not periodic signals and they are additionally bounded due to energy conservation. Hence, in our algorithm, we follow Peters et al.'s method [22] for computing bounded MESE and transforming them to reflectance spectra by mirroring and wavelength-mapping the signals. Such transformed moments also lose their imaginary part, so they become real instead of complex numbers. The meaningful ranges of the moment values are $c_0 \in [0, 1]$ and $c_j \in [-\frac{1}{\pi}, \frac{1}{\pi}]$ for $j \in \{1, \dots, m\}$. We confirm Peters et al.'s observation that 8 moments ($m = 7$) are sufficient for reflectance spectra.

3.4. Estimating the Donaldson matrix via gradient descent

The GMM and bounded MESE models from Secs. 3.2 and 3.3 together form a *parametric model* that describes a continuous Donaldson matrix Φ with only a handful of parameters (Fig. 4). Plugging our model into Eq. (4) leads to an inverse problem, in which these model parameters are the unknown variables that we need to solve for.

This problem has significantly fewer parameters and is more constrained than just naïvely discretizing the whole matrix, which in our early experiments did not lead to any meaningful results. For example, a Donaldson matrix in the wavelength range of 400 to 700 nm discretized with a 10 nm resolution has 496 free parameters (the whole matrix has $31 \cdot 31 = 961$ elements, out of which 31 are the diagonal, 465 are the upper left triangle, and 465 are zero), whereas our parametric model has almost ten times fewer parameters.

Our parametric model is non-linear and non-convex, since a sum of multiple Gaussian distributions can have more local extrema. Fortunately, it has well-defined derivatives, which allows using numerical optimization algorithms based on *gradient descent* [24]. We take inspiration from the field of deep learning, where non-convex optimizations are often solved by adaptive step sizes. Specifically, we noticed that the adaptive moment estimation algorithm *Adam* [25] converges well to estimates with high accuracy.

Optimized parameters In order to minimize Eq. (4), the optimization algorithm estimates the following parameters:

- for each Gaussian distribution, its weight $w > 0$, mean wavelength $\mu \in [400, 700]^2$, and covariance matrix $\Sigma \in \mathbf{R}^{2 \times 2}$ (we used only two non-zero elements to prevent rotations of the distribution), in total 5 real numbers per distribution,
- for the diagonal, the trigonometric moments $c_0 \in [0, 1]$ and $c_j \in [-\frac{1}{\pi}, \frac{1}{\pi}]$ for $j \in \{1, \dots, m\}$ defining the bounded MESE curve, in our case using 8 moments ($m = 7$).

While estimating the trigonometric moments directly via the optimization algorithm is possible and leads to valid results, in practice, it is computationally expensive and also hinders the accuracy. Hence, during the optimization, we instead discretize the diagonal with a 10 nm resolution, leading to 31 parameters (visualized by the blue crosses in Fig. 4(c)), and we let the optimizer estimate them individually with a simple penalization to keep them in the bounded $[0, 1]$ range. Only then, we compute the 8 trigonometric moments corresponding to the discretized diagonal, which finally gives a continuous smooth bounded MESE curve (red curve in Fig. 4(c)).

Initial estimates For the optimization to converge, each parameter needs to be initialized within a valid and meaningful range. Experimentally, we verified that the specific initialization does not play such an important role. We initialized the diagonal as a constant function of $0.5 = 50\%$ reflectance, and each Gaussian distribution with a covariance matrix $\Sigma = \{\{4000, 0\}, \{0, 4000\}\}$ to make their support large enough to reach from the edges to the center, weight $w = 100$, and their mean wavelengths covering the whole triangle, making sure every corner is covered, i.e., for 6 Gaussian distributions: (400, 400), (550, 550), (700, 700), (550, 400), (700, 400), and (700, 550) nm.

Iterations and step sizes To ensure fast convergence, we ran each optimization in three loops with decreasing step sizes. Hence, the Adam algorithm is executed 3 times per Donaldson matrix: first for 100 iterations with a step size of 0.10; second for 200 iterations and 0.05; and third for 800 iterations and 0.02. The second and third loops are initialized with the parameters that yielded the lowest error in the previous loop. Furthermore, the step sizes need to be scaled for each parameter separately, because different parameters have different ranges, e.g., the reflectance is only valid between 0 and 1, whereas the Gaussian covariance matrix can reach large values. The exact values are not critical, because Adam is an adaptive algorithm. In our case, we multiplied the reflectance's step size by 0.05, and each Gaussian parameter's step size by 200, e.g., the actual step size of the Gaussian mean wavelength in the first loop would be $0.10 \cdot 200 = 20$.

3.5. Algorithm implementation

The algorithm itself was implemented in Python and the source codes together with accompanying data are attached in [Code 1](#), Ref. [26]. Most of the implementation is straightforward application of the equations in the paper, with one exception, which is the optimization itself. The gradient-descent algorithm requires computing the gradients (partial derivatives) of Eq. (4) w.r.t. the parameters of the Gaussian distributions and the diagonal. We achieved that using the `jax` auto-differentiation framework [27]. With such a naive implementation, the run-time per material is a couple of minutes, but since it can be trivially executed in parallel, estimating a batch of 32 materials on a 32-core CPU takes the same time as a single material. For comparison, the implementation of Blasinski et al. [21,28] is faster, taking a couple of seconds per material, but that is for the price of lower accuracy.

4. Evaluation

We evaluate our method on a combination of synthetic data and real measurements. We first introduce examples of optical setups compatible with our method (Sec. 4.1), then our synthetic dataset (Sec. 4.2), and then our real measurements (Sec. 4.3). The evaluation itself is both quantitative, based on commonly used error metrics (Sec. 4.4), and qualitative, based on visual examinations of the matrices, spectra, and predictive image rendering (Sec. 4.5).

4.1. Optical setups

We introduce two sets of illuminants (LEDs and color filters) and two kinds of detectors (a spectrometer and a multispectral sensor chip) to demonstrate the variability and compatibility of our method. These illuminants and detectors can be used interchangeably, which gives $2 \cdot 2 = 4$ different examples of compatible measurement setups. We illustrate some of them in Fig. 5 and their spectra in Fig. 6.




Fig. 5. Examples of optical setups with different optical components compatible with our algorithm. The abbreviations L1, L2, S1, and S2 are from Sec. 4.1.

Measurement geometry Figure 5 shows what is often called a $45^\circ/0^\circ$ measurement: a diffuse material sample is illuminated under approximately 45° , while the detector is placed orthogonally to the sample to avoid specular reflections. To measure how the material performs under varying angles, our method could in principle be extended to a fully goniophotometric setup, e.g., with the detector on a rotating arm. Our algorithm would then reconstruct one Donaldson matrix per each angle separately, and these results could be interpolated if needed.

Calibration The Donaldson matrix is a unitless distribution that describes ratios of the reflected and re-emitted energy, so the exact units of $l_1(\lambda), \dots, l_{n_l}(\lambda)$ and $s_1(\lambda), \dots, s_{n_s}(\lambda)$ are arbitrary. It is critical to normalize the values with respect to a single value to ensure that the estimated ratios in the Donaldson matrix are meaningful. For example, Fig. 6(bd) are normalized to the peak of the strongest spectrum. To calibrate the illuminants, we used Spectralon, which has a

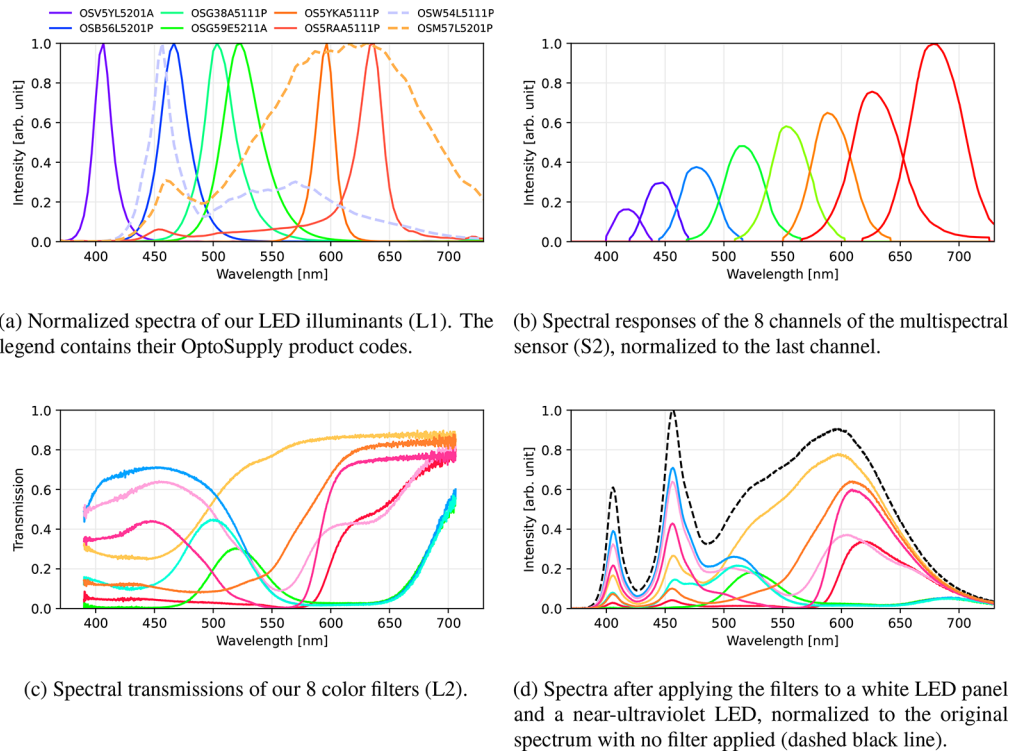


Fig. 6. Measured spectra of our exemplar optical setups from Sec. 4.1 and Fig. 5.

guaranteed 99% diffuse reflectance in the whole visible spectrum. By illuminating the Spectralon with each of the illuminants, the illuminant spectra are obtained from the known reflectance of 99%. It is also important that the detector's spectral sensitivity is calibrated beforehand, such that the detected peaks have correct intensities. Many spectrometers and detectors come factory pre-calibrated, or their calibration is possible with a stabilized light source of a known spectrum.

Examples of illuminants (L1, L2) The goal is to have a set of illuminants distinctly covering the intended spectral range, in our case the visible range from around 400 to 700 nanometers. We demonstrate two affordable options. The first option (denoted **L1**) is *light-emitting diodes* (LEDs), which are readily available with various spectra. In our example (measured in Fig. 6), we use 6 narrowband and 2 broadband (warm white and cold white) LEDs from OptoSupply with the cost of about 0.50 U.S. dollars per piece. The second option (denoted **L2**) is placing different *color filters* in front of a white illuminant. In our example (measured in Fig. 6), we purchased colored foils from a photography equipment store (about 25 U.S. dollars) and we placed them in front of a 200 W white LED panel (about 140 U.S. dollars). Note that it was necessary to add an additional near-ultraviolet LED (about 0.50 U.S. dollars), otherwise the spectrum would not cover the wavelengths around 400 nm. The main difference between L1 and L2 is that L1 relies on multiple LEDs where each is electronically controlled, whereas in L2, the filters can be exchanged quickly by hand.

Examples of detectors (S1, S2) An optimal spectral detector in our case can differentiate between reflected light throughout the whole spectral range. The most straightforward option (denoted **S1**) is a *spectrometer*, which can easily have a sub-nanometer resolution and hundreds of spectral channels. An entry-level USB spectrometer (Ocean Insight USB-650) can be purchased for around 2000 U.S. dollars. A significantly more low-cost solution (denoted **S2**) relies on a

multi-spectral sensor, which is a small chip containing a few pixels, each with a different spectral sensitivity. In our example (measured in Fig. 6), we used the 8 narrowband channels of the AMS AS7341 sensor (about 25 U.S. dollars).

Cost comparison In Table 1, we summarize the aforementioned rough estimated costs of the individual optical setups L1-S1, L1-S2, L2-S1, and L2-S2. We also compare these costs to the reference setup that is used for ground truth measurements, which includes a monochromator, a broadband Xenon light source including its power supply, and a spectrometer. Because the manufacturer of our reference setup does not exist anymore, we instead decided to use the publicly available price lists of Newport Corporation (newport.com, sections “CS130B 1/8m Configured Monochromator” and “Low Power Xenon (Xe) Research Light Sources”) to estimate the prices for the monochromator and light source.

Table 1. Rough estimated costs for the individual optical setups.

| | Reference setup | L1-S1 | L1-S2 | L2-S1 | L2-S2 |
|-------------|--------------------------|-----------|--------|-----------|---------|
| Illuminants | Monochromator: 8,500 USD | 4 USD | 4 USD | 165 USD | 165 USD |
| | Xe light: 10,000 USD | | | | |
| Detector | 2,000 USD | 2,000 USD | 25 USD | 2,000 USD | 25 USD |
| Total | 20,500 USD | 2,004 USD | 29 USD | 2,165 USD | 190 USD |

4.2. Synthetic dataset and simulated measurements

For evaluating fluorescence estimation methods, it is common to simulate measurements on synthetic datasets of Donaldson matrices [14,16,21], i.e., instead of performing an actual measurement, the optical setup is only simulated to evaluate the method’s performance. Such an approach has many benefits: the ground truth matrices are precisely known, the simulations are efficient, and the synthetic datasets can contain many materials, so the evaluation is statistically meaningful compared to real measurements on only a few isolated samples.

Base dataset We use a dataset of real materials that were measured by Gonzales [29], who directly acquired the Donaldson matrices using a specialized bi-spectral spectrophotometer (Labsphere BFC-450), which has two monochromators measuring the ground truth as shown by Donaldson [12] (Sec. 2.1). These ground truth matrices are available in a 10 nm resolution and they cover our range of interest of 400 to 700 nm. As the matrices were acquired with a real setup, they already contain significant noise, hence no synthetic noise was added. We picked 32 samples with strong fluorescence, including color pigments, polymer clays, golf balls, textiles, papers, and index cards, all showing various reflectances and levels of fluorescence. Out of these samples, 28 contained a single fluorophore, and we judged 4 to have two fluorophores.

Enlarging the dataset To enlarge the dataset, we further created 56 synthetic Donaldson matrices containing two fluorophores. First, we sorted the original 28 matrices based on their emission peaks (to prevent entirely overlapping spectra), and then we summed the Donaldson matrices while preserving the original diagonals (reflectances). Physically, this would be roughly equivalent to overlaying two transparent fluorescent slides with different fluorophores, which was done by Blasinski et al. [21].

Simulating the measurements Because we know the parameters of the optical setups from Sec. 4.1, mainly the illuminant and sensor spectra l_j, s_k (Fig. 6), and we know the ground truth matrix Φ from the synthetic dataset, we can use Eq. (3) to compute $I_{j,k}$ (the reflectances that the sensor detects). These values are then used as the sparse inputs to the actual fluorescence estimation method. The method then outputs the estimated matrix $\hat{\Phi}$. The similarity between $\hat{\Phi}$ and Φ is then evaluated (Secs. 4.4 and 4.5). A perfect estimation method would output $\hat{\Phi} = \Phi$. The data and the algorithm for these simulations are part of the attachment Code 1, Ref. [26].

4.3. Real measurements

Apart from the synthetic measurements, we also performed real measurements. Specifically, we measured fluorescent papers of different colors and fluorophores. To obtain ground truth measurements of the same materials, we relied on a tunable monochromator by AMKO GmbH, with a spectral range from 350 nm to 780 nm. We illuminated the sample in 10 nm steps, while the reflected light was being measured by a spectrometer. From these data, the ground truth Donaldson matrix was reconstructed [12]. When obtaining the ground truth matrices, the material needs to be illuminated under the same measurement geometry as when obtaining the sparse measurements. We solved this by using optical fiber to direct the light from the monochromator and from the LEDs into the exact same spot on the material samples. Each measurement was then performed once with the monochromator (for the ground truth Donaldson matrix), and once with the LED spectra (for the Donaldson matrix estimated using our method). Our experiments confirmed that our method works in real conditions, and we discuss the accuracy of the estimated matrices of real and synthetic data together in Sec. 4.5.

4.4. Quantitative evaluation

We evaluate the accuracy of our method on the various low-cost optical setups from Sec. 4.1. We focus on the synthetic dataset as it allows for drawing statistically meaningful conclusions. The exact same inputs were also evaluated using `fiToolbox` [28], an open-source implementation of the state-of-the-art method by Blasinski et al. [21]. This allows us to directly compare the estimated matrices between our method, the state of the art, and ground truth.

Root-mean-square error A commonly used quantitative metric is the *root-mean-square error* (RMSE) of the estimated matrices compared to the ground truths. Since the diagonal Φ_{diag} and fluorescence Φ_{fluo} can be separated, we evaluate their errors independently, which gives more insight into the algorithms' behaviors. We compare the RMSE on all combinations of the exemplar optical setups from Sec. 4.1, namely L1-S1, L1-S2, L2-S1, and L2-S2, which vary by their illuminants (L) and sensors (S). The resulting errors are visualized in Fig. 7 using standard boxplots, which show the medians, quartiles, minimum and maximum errors, and outliers, all separately for the four optical setups, and for one and two fluorophores.

Following Fig. 7, we can conclude that our algorithm reaches *lower errors* than the state of the art, and it is also more stable with *fewer outliers*. The biggest difference is visible in the L1-S1 setup on the single fluorophore dataset, where our fluorescence median error is *three times lower* than Blasinski et al. Notice that our algorithm's error significantly decreases with better optical setups (L1-S1 vs. L2-S2, the worst setup), whereas Blasinski et al. reaches similar errors in most setups. In the original publication by Blasinski et al. [21], Fig. 4, we can see that for a small number of illuminants, increasing the sensor resolution (camera filters) above some point does not improve the accuracy anymore, whereas our algorithm still takes advantage of the extra resolution to lower the estimation error. The only test in which our algorithm underperformed was the most difficult L2-S2 setup on the two-fluorophores dataset, where Blasinski et al. benefitted from their pre-learned basis functions.

Color accuracy Another important aspect is the *color accuracy under narrowband illumination* provided by our estimates (Fig. 8) to ensure that our method is suitable for predictive rendering of such scenarios in computer graphics. Given the estimated Donaldson matrices, we evaluated the error in predicting the color that is reflected from the fluorescent materials when illuminated by varying monochromatic illuminants. We calculated the color difference using the standardized perceptual metric CIEDE2000 (denoted ΔE_{00} , [30]). Following Fig. 8, our color predictions are consistently and significantly more accurate than in the estimates of Blasinski et al. The highest differences can be seen for low wavelengths, because most fluorophores in our dataset

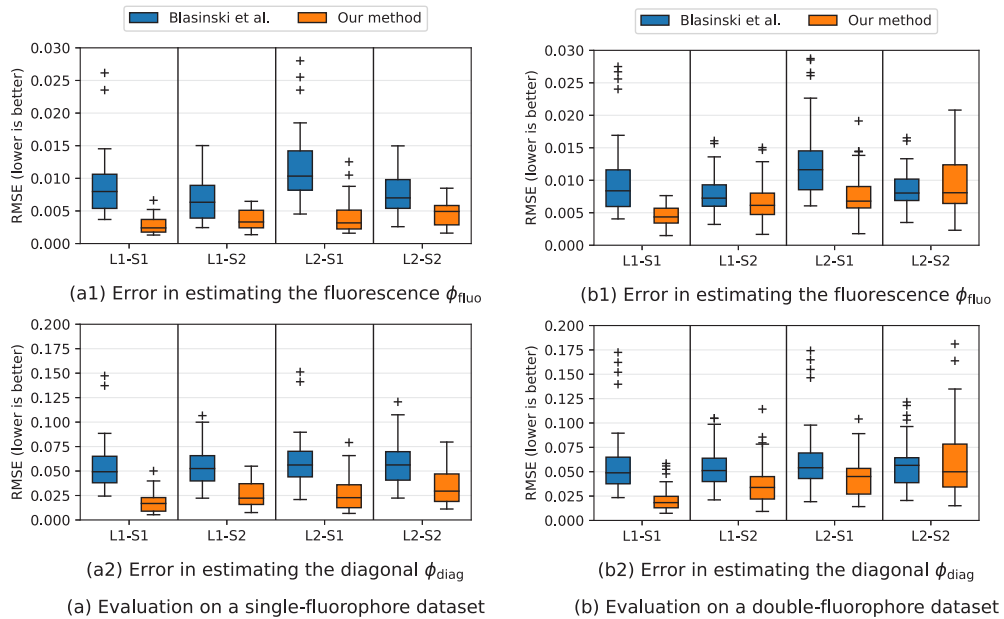


Fig. 7. Quantitative evaluation comparing the root mean square errors (RMSE) of Donaldson matrices estimated with our method (orange) and the state of the art (blue, Blasinski et al. [21]). The boxes show quartiles, horizontal lines are median errors, whiskers minimum and maximum errors, and the plus symbols are outliers. The L1-S1, L1-S2, L2-S1, and L2-S2 correspond to different optical setups per Sec. 4.1.

are triggered with a blue illuminant. Towards red wavelengths, the reflected color becomes independent of the fluorescence, it is given mostly by the pure reflectance (diagonal), and hence the overall color difference is lower.

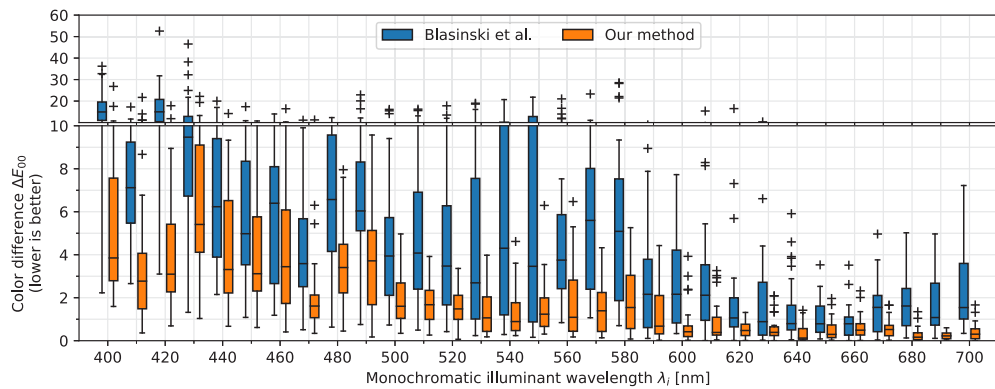


Fig. 8. Quantitative evaluation measuring the color difference (CIEDE2000, ΔE_{00} [30], lower is better) between the ground truth materials and our estimates (orange boxplots), and the state-of-the-art estimates (blue boxplots, Blasinski et al. [21]). The vertical axis shows how much the predicted reflected color is wrong, given the material was monochromatically illuminated with the wavelength on the horizontal axis. The dataset in this plot is the single-fluorophore, L1-S1 optical setup.

4.5. Qualitative evaluation

While the previous section focused on objective quantitative metrics using the synthetic dataset, we now perform a qualitative examination of all our data to understand the behavior of the estimation methods in various contexts.

Visualized matrices In Figs. 9 and 10, we show a small subset of ground truth Donaldson matrices and their estimates using our algorithm and state of the art. We chose 2 examples from the synthetic one-fluorophore dataset, 1 example from the synthetic two-fluorophores dataset, and 2 examples from real measurements. While one can see that the RMSE is typically lower for our estimates compared to Blasinski et al., the major observation is that the estimates of Blasinski et al. suffer from the aforementioned oscillations.

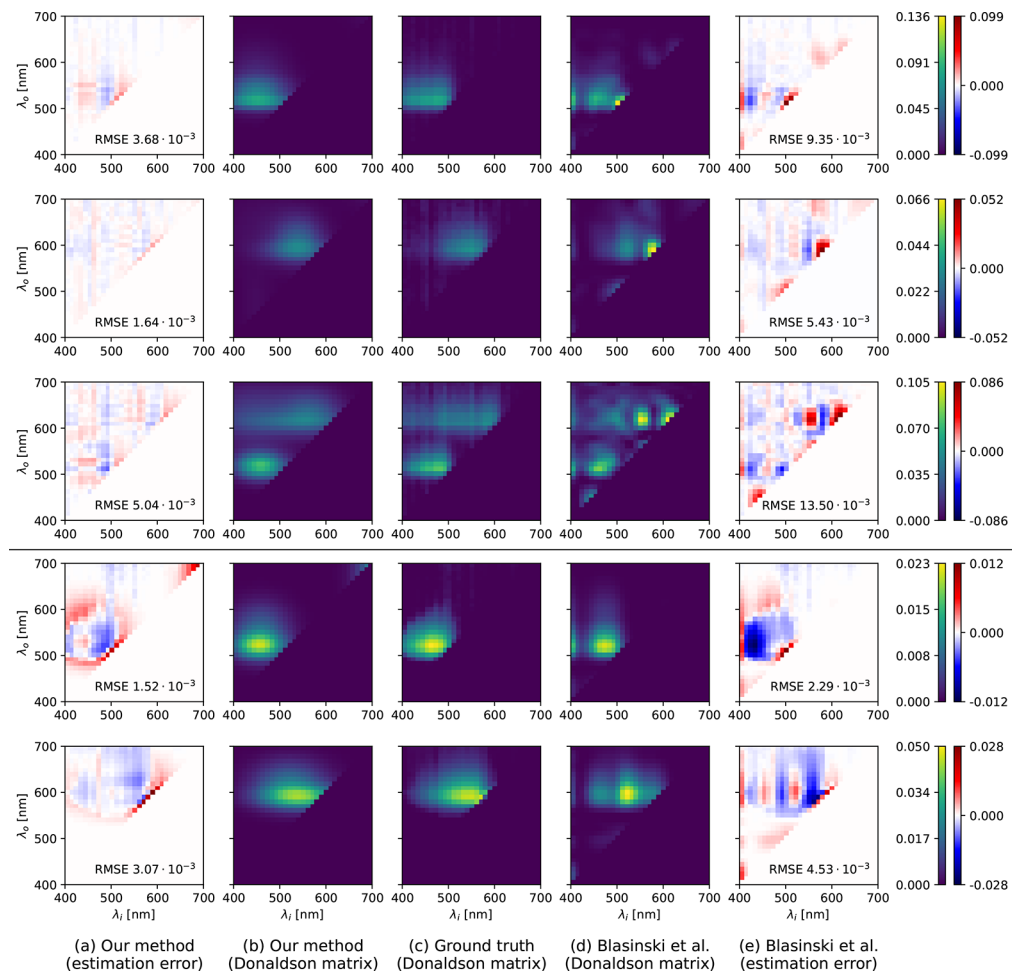


Fig. 9. Examples of Donaldson matrices of five different materials and their estimates using the L1-S1 optical setup, comparing our method (a,b) to ground truth (c) and the state of the art (d,e). The top three measurements are synthetic, the bottom two are real.

For example, in Fig. 9, all the ground truth matrices have either a single fluorescent peak or two peaks. While our estimates manage to reconstruct the number of peaks with a high accuracy, Blasinski et al. often overestimate the number of peaks, and the Donaldson matrices then resemble ripples on the water surface. In Fig. 10, we can see a similar effect in the reflectance

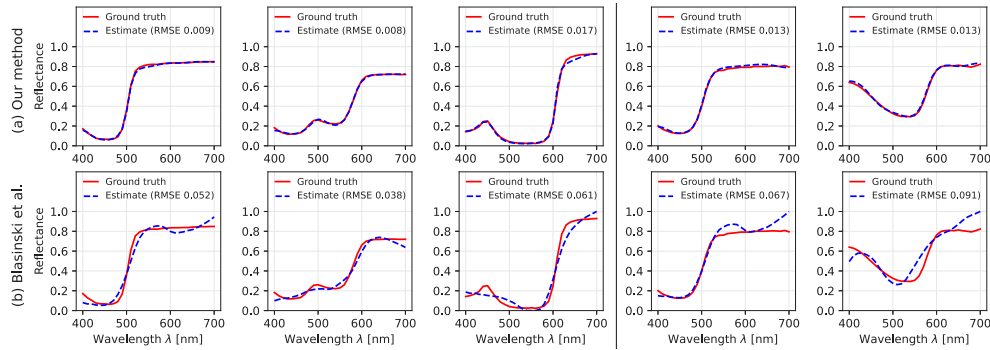


Fig. 10. Examples of ground truth and estimated pure reflectances (diagonals) of the five materials from Fig. 9 under identical conditions, comparing our method (a) to the state of the art (b). The left three measurements are synthetic, the right two are real.

spectra: where the ground truth and our estimates have a relatively flat spectrum, the estimates of Blasinski et al. often oscillate and create waves, which are not present in the ground truth spectra.

Using the estimates in predictive rendering Furthermore, we also present an example of using our estimated Donaldson matrices in the Monte Carlo predictive renderer ART [31], which is open-source and supports fluorescent materials. In Fig. 11, we show a comparison of six images rendered with ground truth matrices, and matrices estimated with our method and by Blasinski et al. [21] for two fluorescent materials lit by a monochromatic illuminant at 560 nm. This is an illustration of Fig. 8: the importance of having a consistently low Δ_E color error over the illuminants to ensure high color accuracy. The renders have been tonemapped with the exact same parameters to ensure that the exposure and contrast do not vary from result to result.




Fig. 11. A fluorescent ball is monochromatically illuminated (560 nm). The ball in each rendered image has a different Donaldson matrix. The balls (a) and (d) use the ground truth matrices from the material dataset, and the balls (b)-(c) and (e)-(f) use matrices estimated with our method, and with Blasinski et al. [21].

5. Conclusion

We have presented a simple and affordable technique for estimating fluorescence in materials. The inputs for our algorithm are easily acquirable with low-cost optical setups, and the output is in the form of a parametrized Donaldson matrix with a small memory footprint, usable in any standard application including predictive and photorealistic rendering. We showed that our estimates have objectively lower errors than the state-of-the-art algorithm on the exact same input data, our algorithm is also more stable and robust with a lower number of outliers, and our estimates are qualitatively better and more suitable for color-accurate predictive rendering.

Our paper's core principle was the parametrization of the Donaldson matrix by a Gaussian mixture model (GMM) and bounded maximum entropy spectral estimate (bounded MESE), which can be done in a differentiable way compatible with gradient-descent optimization algorithms.

This makes it compatible with sparse measurements and makes it possible to measure fluorescence accurately without specialized and expensive bi-spectral spectrophotometers.

Future work We can see several minor directions for future work. One improvement would consist of a faster implementation that does not rely on auto-differentiation, and instead computes derivatives using equations derived by hand, possibly accelerated on the GPU. Furthermore, the optimization itself could consist of more adaptive steps, such that only one Gaussian distribution is estimated at first, and more Gaussians are added subsequently to allow faster convergence for materials with only one fluorophore. Finally, our method allows anyone to quickly build a dataset of Donaldson matrices of common materials, which would be useful in computer graphics for predictive and photorealistic rendering.

Funding. Grantová Agentura České Republiky (GAČR-22-22875S); Horizon 2020 Framework Programme (956585); Univerzita Karlova v Praze (SVV-260699).

Disclosures. The authors declare no conflicts of interest.

Data availability. Data underlying the results presented in this paper are available in [Code 1](#), Ref. [26], along with Python source codes used to generate the reported values.

References

1. G. Johnson and M. Fairchild, "Full-spectral color calculations in realistic image synthesis," *IEEE Comput. Graph. Appl.* **19**(4), 47–53 (1999).
2. M. B. Hullin, J. Hanika, B. Ajdin, H.-P. Seidel, J. Kautz, and H. P. A. Lensch, "Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions," *ACM Trans. Graph.* **29**(4), 1–7 (2010).
3. M. Mojzík, A. Fichet, and A. Wilkie, "Handling Fluorescence in a Uni-directional Spectral Path Tracer," *Comput. Graph. Forum* **37**(4), 77–94 (2018).
4. A. Jung, A. Wilkie, J. Hanika, W. Jakob, and C. Dachsbacher, "Wide Gamut Spectral Upsampling with Fluorescence," *Comput. Graph. Forum* **38**, 1–10 (2019).
5. A. Jung, J. Hanika, and C. Dachsbacher, "Spectral Mollification for Bidirectional Fluorescence," *Comput. Graph. Forum* **39**(2), 373–384 (2020).
6. L. König, A. Jung, and C. Dachsbacher, "Improving Spectral Upsampling with Fluorescence," (2020).
7. Q. Hua, V. Tázlar, A. Fichet, and A. Wilkie, "Efficient Storage and Importance Sampling for Fluorescent Reflectance," *Computer Graphics Forum* (2022).
8. G. H. Mohammed, R. Colombo, and E. M. Middleton, *et al.*, "Remote sensing of solar-induced chlorophyll fluorescence (SIF) in vegetation: 50 years of progress," *Remote. Sens. Environ.* **231**, 111177 (2019).
9. P. L. Choyke and H. Kobayashi, "Medical Uses of Fluorescence Imaging: Bringing Disease to Light," *IEEE J. Sel. Top. Quantum Electron.* **18**(3), 1140–1146 (2012).
10. S. Delalieux, A. Auwerkerken, W. W. Verstraeten, B. Somers, R. Valcke, S. Lhermitte, J. Keulemans, and P. Coppin, "Hyperspectral Reflectance and Fluorescence Imaging to Detect Scab Induced Stress in Apple Leaves," *Remote Sens.* **1**(4), 858–874 (2009).
11. M. S. Roth and D. D. Deheyn, "Effects of cold stress and heat stress on coral fluorescence in reef-building corals," *Sci. Rep.* **3**(1), 1421 (2013).
12. R. Donaldson, "Spectrophotometry of fluorescent pigments," *Br. J. Appl. Phys.* **5**(6), 210–214 (1954).
13. J. Suo, L. Bian, F. Chen, and Q. Dai, "Bispectral coding: compressive and high-quality acquisition of fluorescence and reflectance," *Opt. Express* **22**(2), 1697–1712 (2014).
14. Y. Fu, A. Lam, I. Sato, T. Okabe, and Y. Sato, "Separating Reflective and Fluorescent Components Using High Frequency Illumination in the Spectral Domain," in *2013 IEEE International Conference on Computer Vision*, (2013), pp. 457–464. ISSN: 2380–7504.
15. C. Zhang and I. Sato, "Separating reflective and fluorescent components of an image," in *CVPR 2011*, (2011), pp. 185–192. ISSN: 1063–6919.
16. S. Tominaga, K. Hirai, and T. Horiuchi, "Estimation of bispectral Donaldson matrices of fluorescent objects by using two illuminant projections," *J. Opt. Soc. Am. A* **32**(6), 1068–1078 (2015).
17. Y. Zheng, I. Sato, and Y. Sato, "Spectra Estimation of Fluorescent and Reflective Scenes by Using Ordinary Illuminants," in *Computer Vision – ECCV 2014*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, eds. (Springer International Publishing, Cham, 2014), Lecture Notes in Computer Science, pp. 188–202.
18. Y. Zheng, Y. Fu, A. Lam, I. Sato, and Y. Sato, "Separating Fluorescent and Reflective Components by Using a Single Hyperspectral Image," in *2015 IEEE International Conference on Computer Vision (ICCV)*, (2015), pp. 3523–3531. ISSN: 2380–7504.
19. A. Lam and I. Sato, "Spectral Modeling and Relighting of Reflective-Fluorescent Scenes," in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, (2013), pp. 1452–1459. ISSN: 1063–6919.
20. Y. Fu, A. Lam, I. Sato, T. Okabe, and Y. Sato, "Reflectance and Fluorescence Spectral Recovery via Actively Lit RGB Images," *IEEE Trans. Pattern Anal. Mach. Intell.* **38**(7), 1313–1326 (2016).

21. H. Blasinski, J. Farrell, and B. Wandell, "Simultaneous Surface Reflectance and Fluorescence Spectra Estimation," *IEEE Trans. on Image Process.* **29**, 8791–8804 (2020).
22. C. Peters, S. Merzbach, J. Hanika, and C. Dachsbacher, "Using moments to represent bounded signals for spectral rendering," *ACM Trans. Graph.* **38**(4), 1–14 (2019).
23. J. P. Burg, "Maximum Entropy Spectral Analysis," PhD Thesis, Stanford University (1975).
24. A. Antoniou and W.-S. Lu, *Practical Optimization: Algorithms and Engineering Applications*, Texts in Computer Science (Springer US, New York, NY, 2021).
25. D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *ArXiv*, ArXiv:1412.6980 (2017).
26. T. Iser, L. Lachiver, and A. Wilkie, "Source code and data for the fluorescence measurement method using Gaussian distributions and bounded MESE," figshare (2023), <https://doi.org/10.6084/m9.figshare.23567073>.
27. J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, "JAX: composable transformations of Python+NumPy programs," (2018).
28. H. Blasinski, "Fluorescence Imaging Toolbox for Image Systems Engineering Toolbox (fiToolbox for ISET)," (2020). GitHub repository: <https://github.com/hblasins/fiToolbox>.
29. S. Gonzalez, "Evaluation of bispectral spectrophotometry for accurate colorimetry of printing materials," Ph.D. thesis, Rochester Institute of Technology (2000).
30. G. Sharma, W. Wu, and E. N. Dalal, "The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations," *Color Res. Appl.* **30**(1), 21–30 (2005).
31. A. Wilkie and R. F. Tobler, "ART: The Advanced Rendering Toolkit," (2022). URL: <https://cgg.mff.cuni.cz/ART/>.

The inserted publication ends here.

6 | Conclusion

Predictive rendering plays an important role in the modern digital era. It allows the creation of digital twins, such that the appearance of real objects can be precisely represented in the virtual world. That has numerous applications for designers, prototyping, visualizations, or e-commerce, but also industrial applications that rely on computational optimization, such as in 3D printing.

Rendering, whose fundamentals were covered in Chapter 2, requires optical properties of the materials in the scene. It is evident that rendering can only be as accurate as the properties themselves. Measuring the optical properties is often straightforward if we have unlimited time and finances and we can afford to build specialized measurement setups in a fully equipped optical laboratory with expensive equipment and calibrated samples. Unfortunately, measurement methods that would be affordable yet accurate enough for predictive rendering were not available for more complex materials exhibiting translucency or fluorescence.

In this thesis, we interpreted material measurements as an inverse problem to rendering (Chapter 3), which allowed us to carefully design two affordable measurement methods. In Chapter 4, we included our publication specialized in translucent materials [Iser et al., 2022]. Using only three point measurements with a spectrometer, we recover the spectral absorption and scattering coefficients and the phase function anisotropy factor of a material sample. The estimates are based on Monte Carlo predictions and pre-computed lookup tables, and the accuracy was demonstrated on full-color 3D printing. In Chapter 5, we included our publication on affordable fluorescence measurements [Iser et al., 2023]. The unknown Donaldson matrix of a material sample is estimated from sparse measurements under a small set of illuminants, which is allowed by parametrizing the matrix using Gaussian distributions and bounded MESE, and using gradient-descent optimization. The accuracy significantly surpassed previous work.

We believe that both publications together have demonstrated that it is possible to design specialized measurement methods that are affordable, yet do not compromise the accuracy that is required in predictive rendering. Working with low-cost optical components is challenging and requires carefully examining the problem domain, but the results have a significant potential impact by improving the accessibility of measurements, showing which relations between optical properties and appearance are important, and allowing the potential miniaturization and integration of the methods into affordable standalone optical devices. And it is possible that one day, every smartphone will be capable of simply scanning the materials around us and immediately integrating them into virtual reality.

Future challenges

In the grand scheme of things, we believe that the major future challenge is to find a method that can estimate the optical properties of arbitrary materials from as few observations as possible. Eventually, it may become possible with modern inverse rendering approaches. Currently, our methods are specialized in specific materials and do not scale to generic materials that may combine more of these effects. For example, we are not aware of any method that would be capable

of measuring fluorescent participating media that exhibit absorption, scattering, and fluorescence together. While it may sound like a hypothetical problem at first, such materials are actually common because many fabrics and plastics in everyday items contain optical brighteners.

Fluorescence is still a phenomenon that is not integrated into many frameworks, simply because its effect was usually considered negligible for photorealistic rendering. However, in predictive rendering, the discrepancies caused by not taking fluorescence into account are, unfortunately, significant, especially when we are expecting an exact match between a virtual material and a real one.

Even in 3D printing, there are still unsolved problems in the predictions. Even though our method in Chapter 4 is capable of measuring the optical properties of the base inks, when the inks are halftoned and printed together, various inaccuracies that stem from the printing process itself propagate to the appearance of the final printout. In the future, we would like to measure and characterize the printing process including its inaccuracies, for example, the surface roughness and anisotropy due to the printhead movements.

There are also other optical properties that are not trivial to measure without specialized equipment. Examples include spectral measurements of spatially varying translucent materials, such as fabrics, especially those with complex patterns, or anisotropic materials that must be scanned from many angles. We look forward to future research on material measurements optimized for predictive rendering, and we hope that affordability remains an important factor.

Bibliography

Please note that this bibliography only contains publications referenced in the main text of the thesis. Otherwise, each of the two inserted journal publications has its own bibliography at the end of the insert. We decided to not merge the bibliographies as it could lead to confusion.

Andreas Antoniou and Wu-Sheng Lu. *Practical Optimization: Algorithms and Engineering Applications*. Texts in Computer Science. Springer US, New York, NY, 2021. ISBN 978-1-07-160841-8 978-1-07-160843-2. URL <https://link.springer.com/10.1007/978-1-0716-0843-2>.

James Arvo and David Kirk. Particle transport and image synthesis. *ACM SIGGRAPH Computer Graphics*, 24(4):63–66, September 1990. ISSN 0097-8930. doi: 10.1145/97880.97886. URL <https://dl.acm.org/doi/10.1145/97880.97886>.

Michael Bartholomew-Biggs, Steven Brown, Bruce Christianson, and Laurence Dixon. Automatic differentiation of algorithms. *Journal of Computational and Applied Mathematics*, 124(1):171–190, December 2000. ISSN 0377-0427. doi: 10.1016/S0377-0427(00)00422-2. URL <https://www.sciencedirect.com/science/article/pii/S0377042700004222>.

Benedikt Bitterli, Srinath Ravichandran, Thomas Müller, Magnus Wrenninge, Jan Novák, Steve Marschner, and Wojciech Jarosz. A radiative transfer framework for non-exponential media. *ACM Transactions on Graphics*, 37(6):225:1–225:17, 2018. ISSN 0730-0301. doi: 10.1145/3272127.3275103. URL <https://doi.org/10.1145/3272127.3275103>.

Henryk Blasinski, Joyce Farrell, and Brian Wandell. Simultaneous Surface Reflectance and Fluorescence Spectra Estimation. *IEEE Transactions on Image Processing*, 29:8791–8804, 2020. ISSN 1941-0042. doi: 10.1109/TIP.2020.2973810.

James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.

E. C. Carter, Y. Ohno, M. R. Pointer, A. R. Robertson, R. Seve, J. D. Schanda, and K. Witt, editors. *Colorimetry*. Number 15 in Publication / CIE. Comm. Internat. de l’éclairage, Wien, 3rd ed edition, 2004. ISBN 978-3-901906-33-6.

S. Chandrasekhar. *Radiative Transfer*. Dover Publications, New York, January 1960. ISBN 978-0-486-60590-6.

Chengqian Che, Fujun Luan, Shuang Zhao, Kavita Bala, and Ioannis Gkioulekas. Towards Learning-based Inverse Subsurface Scattering. In *2020 IEEE International Conference on Computational Photography (ICCP)*, pages 1–12, Saint Louis, MO, USA, April 2020. IEEE. ISBN 978-1-72815-230-1. doi: 10.1109/

ICCP48838.2020.9105209. URL <https://ieeexplore.ieee.org/document/9105209/>.

Gregory R. Choppin, Jan-Olov Liljenzin, Jan Rydberg, and Christian Ekberg. *Radiochemistry and nuclear chemistry*. Elsevier/AP, Academic Press is an imprint of Elsevier, Amsterdam ; Boston, 4th edition edition, 2013. ISBN 978-0-12-405897-2. OCLC: ocn852806072.

Peter L. Choyke and Hisataka Kobayashi. Medical Uses of Fluorescence Imaging: Bringing Disease to Light. *IEEE Journal of Selected Topics in Quantum Electronics*, 18(3):1140–1146, May 2012. ISSN 1558-4542. doi: 10.1109/JSTQE.2011.2164900.

Stephanie Delalieux, Annemarie Auwerkerken, Willem W. Verstraeten, Ben Somers, Roland Valcke, Stefaan Lhermitte, Johan Keulemans, and Pol Coppin. Hyperspectral Reflectance and Fluorescence Imaging to Detect Scab Induced Stress in Apple Leaves. *Remote Sensing*, 1(4):858–874, December 2009. ISSN 2072-4292. doi: 10.3390/rs1040858. URL <https://www.mdpi.com/2072-4292/1/4/858>.

R. Donaldson. Spectrophotometry of fluorescent pigments. *British Journal of Applied Physics*, 5(6):210–214, June 1954. ISSN 0508-3443. doi: 10.1088/0508-3443/5/6/303. URL <https://doi.org/10.1088/0508-3443/5/6/303>.

Jonathan Dupuy and Wenzel Jakob. An adaptive parameterization for efficient material acquisition and rendering. *ACM Transactions on Graphics*, 37(6):274:1–274:14, December 2018. ISSN 0730-0301. doi: 10.1145/3272127.3275059. URL <https://dl.acm.org/doi/10.1145/3272127.3275059>.

Philip Dutré, Eric P Lafortune, and Yves D Willems. Monte Carlo light tracing with direct computation of pixel intensities. 1993.

Oskar Elek, Denis Sumin, Ran Zhang, Tim Weyrich, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, and Jaroslav Křivánek. Scattering-aware Texture Reproduction for 3D Printing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 36(6):241:1–241:15, November 2017. ISSN 0730-0301. doi: 10.1145/3130800.3130890.

Oskar Elek, Ran Zhang, Denis Sumin, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, Jaroslav Křivánek, and Tim Weyrich. Robust and practical measurement of volume transport parameters in solid photo-polymer materials for 3D printing. *Optics Express*, 29(5):7568–7588, March 2021. ISSN 1094-4087. doi: 10.1364/OE.406095. URL <https://www.osapublishing.org/oe/abstract.cfm?uri=oe-29-5-7568>. Publisher: Optical Society of America.

Claudia Emde, Robert Buras-Schnell, Arve Kylling, Bernhard Mayer, Josef Gasteiger, Ulrich Hamann, Jonas Kylling, Bettina Richter, Christian Pause, Timothy Dowling, and Luca Bugliaro. The libRadtran software package for radiative transfer calculations (version 2.0.1). *Geoscientific Model Development*, 9(5):1647–1672, May 2016. ISSN 1991-959X. doi: 10.5194/gmd-9-1647-2016.

- URL <https://gmd.copernicus.org/articles/9/1647/2016/>. Publisher: Copernicus GmbH.
- R. Fletcher and C. M. Reeves. Function minimization by conjugate gradients. *The Computer Journal*, 7(2):149–154, January 1964. ISSN 0010-4620. doi: 10.1093/comjnl/7.2.149. URL <https://doi.org/10.1093/comjnl/7.2.149>. eprint: <https://academic.oup.com/comjnl/article-pdf/7/2/149/959725/070149.pdf>.
- Ying Fu, Antony Lam, Imari Sato, Takahiro Okabe, and Yoichi Sato. Reflectance and Fluorescence Spectral Recovery via Actively Lit RGB Images. *IEEE transactions on pattern analysis and machine intelligence*, 38(7):1313–1326, July 2016. ISSN 1939-3539. doi: 10.1109/TPAMI.2015.2439270.
- Iliyan Georgiev, Jaroslav Krivánek, Tomáš Davidovič, and Philipp Slusallek. Light transport simulation with vertex connection and merging. *ACM Transactions on Graphics*, 31(6):192:1–192:10, November 2012. ISSN 0730-0301. doi: 10.1145/2366145.2366211. URL <https://dl.acm.org/doi/10.1145/2366145.2366211>.
- Ioannis Gkioulekas, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin. Inverse volume rendering with material dictionaries. *ACM Transactions on Graphics (TOG)*, 32(6):162:1–162:13, 2013. ISSN 0730-0301. doi: 10.1145/2508363.2508377. Publisher: ACM New York, NY, USA.
- L. G. Henyey and J. L. Greenstein. Diffuse radiation in the Galaxy. *Astrophysical Journal*, 93:70–83, January 1941. doi: 10.1086/144246.
- Yuanming Hu, Luke Anderson, Tzu-Mao Li, Qi Sun, Nathan Carr, Jonathan Ragan-Kelley, and Frédo Durand. DiffTaichi: Differentiable Programming for Physical Simulation. *ICLR*, 2020. URL: <https://github.com/yuanming-hu/taichi>.
- Q. Hua, V. Tázlar, A. Fichet, and A. Wilkie. Efficient Storage and Importance Sampling for Fluorescent Reflectance. *Computer Graphics Forum*, 2022. ISSN 1467-8659. doi: 10.1111/cgf.14716. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14716>.
- Matthias B. Hullin, Johannes Hanika, Boris Ajdin, Hans-Peter Seidel, Jan Kautz, and Hendrik P. A. Lensch. Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions. *ACM Transactions on Graphics*, 29(4):97:1–97:7, July 2010. ISSN 0730-0301. doi: 10.1145/1778765.1778834. URL <https://doi.org/10.1145/1778765.1778834>.
- Tomáš Iser, Tobias Rittig, Emilie Nogué, Thomas Klaus Nindel, and Alexander Wilkie. Affordable Spectral Measurements of Translucent Materials. *ACM Transactions on Graphics*, 41(6):1–13, December 2022. ISSN 0730-0301, 1557-7368. doi: 10.1145/3550454.3555499. URL <https://dl.acm.org/doi/10.1145/3550454.3555499>.
- Tomáš Iser, Loïc Lachiver, and Alexander Wilkie. Affordable method for measuring fluorescence using Gaussian distributions and bounded MESE. *Optics Express*, 31(15):24347–24362, July 2023. doi: 10.1364/OE.495459. URL <https://doi.org/10.1364/OE.495459>.

- [//opg.optica.org/oe/abstract.cfm?URI=oe-31-15-24347](http://opg.optica.org/oe/abstract.cfm?URI=oe-31-15-24347). Publisher: Optica Publishing Group.
- Steven L. Jacques. Optical properties of biological tissues: a review. *Physics in Medicine & Biology*, 58(11):R37, May 2013. ISSN 0031-9155. doi: 10.1088/0031-9155/58/11/R37. URL <https://iopscience.iop.org/article/10.1088/0031-9155/58/11/R37/meta>. Publisher: IOP Publishing.
- Wenzel Jakob and Johannes Hanika. A Low-Dimensional Function Space for Efficient Spectral Upsampling. *Computer Graphics Forum (Proceedings of Eurographics)*, 38(2), March 2019.
- Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. Mitsuba 3 renderer, 2022a. <https://mitsuba-renderer.org>.
- Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, and Delio Vicini. Dr.Jit: A Just-In-Time Compiler for Differentiable Rendering. *Transactions on Graphics (Proceedings of SIGGRAPH)*, 41(4), July 2022b. doi: 10.1145/3528223.3530099.
- Wojciech Jarosz. *Efficient Monte Carlo Methods for Light Transport in Scattering Media*. PhD Thesis, UC San Diego, September 2008.
- Johannes Jendersie. Fast Spectral Upsampling of Volume Attenuation Coefficients. In Adam Marrs, Peter Shirley, and Ingo Wald, editors, *Ray Tracing Gems II: Next Generation Real-Time Rendering with DXR, Vulkan, and OptiX*, pages 153–159. Apress, Berkeley, CA, 2021. ISBN 978-1-4842-7185-8. doi: 10.1007/978-1-4842-7185-8_13. URL https://doi.org/10.1007/978-1-4842-7185-8_13.
- Henrik Wann Jensen. *Realistic Image Synthesis Using Photon Mapping*. A. K. Peters, Ltd., USA, 2001. ISBN 1-56881-147-0.
- G.M. Johnson and M.D. Fairchild. Full-spectral color calculations in realistic image synthesis. *IEEE Computer Graphics and Applications*, 19(4):47–53, July 1999. ISSN 1558-1756. doi: 10.1109/38.773963.
- Steven G. Johnson. The NLOpt nonlinear-optimization package, 2019. URL <https://github.com/stevengj/nlopt>.
- A. Jung, J. Hanika, and C. Dachsbacher. Spectral Mollification for Bidirectional Fluorescence. *Computer Graphics Forum*, 39(2):373–384, 2020. ISSN 1467-8659. doi: 10.1111/cgf.13937. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13937>.
- Alisa Jung, Alexander Wilkie, Johannes Hanika, Wenzel Jakob, and Carsten Dachsbacher. Wide Gamut Spectral Upsampling with Fluorescence. *Computer Graphics Forum (Proceedings of Eurographics Symposium on Rendering)*, 38(4), October 2019. doi: 10.1111/cgf.13773.

- James T. Kajiya. The rendering equation. *ACM SIGGRAPH Computer Graphics*, 20(4):143–150, August 1986. ISSN 0097-8930. doi: 10.1145/15886.15902. URL <https://dl.acm.org/doi/10.1145/15886.15902>.
- Rei Kawakami, Hongxun Zhao, Robby T. Tan, and Katsushi Ikeuchi. Camera Spectral Sensitivity and White Balance Estimation from Sky Images. *International Journal of Computer Vision*, 105(3):187–204, December 2013. ISSN 1573-1405. doi: 10.1007/s11263-013-0632-1. URL <https://doi.org/10.1007/s11263-013-0632-1>.
- Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization, January 2017. URL <http://arxiv.org/abs/1412.6980>. arXiv:1412.6980 [cs].
- Ivo Kondapaneni, Petr Vevoda, Pascal Grittmann, Tomáš Skřivan, Philipp Slusallek, and Jaroslav Křivánek. Optimal multiple importance sampling. *ACM Transactions on Graphics*, 38(4):37:1–37:14, July 2019. ISSN 0730-0301. doi: 10.1145/3306346.3323009. URL <https://dl.acm.org/doi/10.1145/3306346.3323009>.
- Paul Kubelka and Franz Munk. Ein Beitrag zur Optik der Farbanstriche. *Zeitschrift für Technische Physik*, 12:593–601, 1931.
- Lars König, Alisa Jung, and Carsten Dachsbacher. Improving Spectral Upsampling with Fluorescence. 2020. ISBN 978-3-03868-108-3. doi: 10.2312/mam.20201139. URL <https://diglib.eg.org:443/xmlui/handle/10.2312/mam20201139>.
- Jaroslav Křivánek, Iliyan Georgiev, Toshiya Hachisuka, Petr Vévoda, Martin Šik, Derek Nowrouzezahrai, and Wojciech Jarosz. Unifying Points, Beams, and Paths in Volumetric Light Transport Simulation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 33(4), July 2014. doi: 10/f6cz72.
- Eric P Lafortune and Yves D Willems. Bi-directional path tracing. 1993.
- Antony Lam and Imari Sato. Spectral Modeling and Relighting of Reflective-Fluorescent Scenes. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1452–1459, June 2013. doi: 10.1109/CVPR.2013.191. ISSN: 1063-6919.
- Sven Leyre, Youri Meuret, Guy Durinck, Johan Hofkens, Geert Deconinck, and Peter Hanselaer. Estimation of the effective phase function of bulk diffusing materials with the inverse adding-doubling method. *Applied Optics*, 53(10):2117, April 2014. ISSN 1559-128X, 2155-3165. doi: 10.1364/AO.53.002117. URL <https://www.osapublishing.org/abstract.cfm?URI=ao-53-10-2117>.
- Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable Monte Carlo Ray Tracing through Edge Sampling. *ACM Transactions on Graphics (TOG)*, 37(6):125:1–125:12, July 2018. ISSN 0730-0301. doi: 10.1145/3306346.3322954. Publisher: ACM New York, NY, USA.

- Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. Reparameterizing discontinuous integrands for differentiable rendering. *ACM Transactions on Graphics*, 38(6):1–14, November 2019. ISSN 0730-0301, 1557-7368. doi: 10.1145/3355089.3356510. URL <https://dl.acm.org/doi/10.1145/3355089.3356510>. <https://rgl.epfl.ch/publications/Loubet2019Reparameterizing>.
- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, 2015. URL <http://tensorflow.org/>.
- Domingo Mery. *Computer Vision for X-Ray Testing: Imaging, Systems, Image Databases, and Algorithms*. Springer International Publishing, 2015. ISBN 978-3-319-20746-9. URL <https://www.springer.com/gp/book/9783319207469>.
- S. Miao, A. Tuysuzoglu, Z. J. Wang, and R. Liao. Real-time 6DoF pose recovery from X-ray images using library-based DRR and hybrid optimization. *International Journal of Computer Assisted Radiology and Surgery*, 11(6):1211–1220, June 2016. ISSN 1861-6410, 1861-6429. doi: 10.1007/s11548-016-1387-2. URL <http://link.springer.com/10.1007/s11548-016-1387-2>.
- Martin Mirbauer, Tobias Rittig, Tomáš Iser, Jaroslav Křivánek, and Elena Šikudová. SkyGAN: Towards Realistic Cloud Imagery for Image Based Lighting. In *Eurographics Symposium on Rendering*, page 10 pages. The Eurographics Association, 2022. doi: 10.2312/SR.20221151. URL <https://diglib.eg.org/handle/10.2312/sr20221151>.
- Gina H. Mohammed, Roberto Colombo, Elizabeth M. Middleton, Uwe Rascher, Christiaan van der Tol, Ladislav Nedbal, Yves Goulas, Oscar Pérez-Priego, Alexander Damm, Michele Meroni, Joanna Joiner, Sergio Cogliati, Wouter Verhoef, Zbyněk Malenovský, Jean-Philippe Gastellu-Etchegorry, John R. Miller, Luis Guanter, Jose Moreno, Ismael Moya, Joseph A. Berry, Christian Frankenberg, and Pablo J. Zarco-Tejada. Remote sensing of solar-induced chlorophyll fluorescence (SIF) in vegetation: 50 years of progress. *Remote Sensing of Environment*, 231:111177, September 2019. ISSN 0034-4257. doi: 10.1016/j.rse.2019.04.030. URL <https://www.sciencedirect.com/science/article/pii/S0034425719301816>.
- M. Mojzík, A. Fichet, and A. Wilkie. Handling Fluorescence in a Uni-directional Spectral Path Tracer. *Computer Graphics Forum*, 37(4):77–94, July 2018. ISSN 01677055. doi: 10.1111/cgf.13477. URL <http://doi.wiley.com/10.1111/cgf.13477>.

- Thomas Müller, Markus Gross, and Jan Novák. Practical path guiding for efficient light-transport simulation. In *Computer Graphics Forum*, volume 36, pages 91–100. Wiley Online Library, 2017. Issue: 4.
- Srinivasa G. Narasimhan, Mohit Gupta, Craig Donner, Ravi Ramamoorthi, Shree K. Nayar, and Henrik Wann Jensen. Acquiring scattering properties of participating media by dilution. *ACM Transactions on Graphics*, 25(3): 1003–1012, July 2006. ISSN 0730-0301. doi: 10.1145/1141911.1141986. URL <https://doi.org/10.1145/1141911.1141986>.
- J. A. Nelder and R. Mead. A Simplex Method for Function Minimization. *The Computer Journal*, 7(4):308–313, January 1965. ISSN 0010-4620, 1460-2067. doi: 10.1093/comjnl/7.4.308. URL <https://academic.oup.com/comjnl/article-lookup/doi/10.1093/comjnl/7.4.308>.
- Merlin Nimier-David, Sébastien Speierer, Benoît Ruiz, and Wenzel Jakob. Radiative Backpropagation: An Adjoint Method for Lightning-Fast Differentiable Rendering. *Transactions on Graphics (Proceedings of SIGGRAPH)*, 39(4): 146:1–146:15, July 2020. ISSN 0730-0301. doi: 10.1145/3386569.3392406.
- Merlin Nimier-David, Thomas Müller, Alexander Keller, and Wenzel Jakob. Unbiased inverse volume rendering with differential trackers. *ACM Transactions on Graphics*, 41(4):44:1–44:20, July 2022. ISSN 0730-0301. doi: 10.1145/3528223.3530073. URL <https://dl.acm.org/doi/10.1145/3528223.3530073>.
- Thomas K. Nindel, Mohcen Hafidi, Tomáš Iser, and Alexander Wilkie. Automatic inference of an anatomically meaningful solid wood texture from a single photograph, February 2023. URL <http://arxiv.org/abs/2302.01820>. arXiv:2302.01820 [cs].
- Thomas Klaus Nindel, Tomáš Iser, Tobias Rittig, Alexander Wilkie, and Jaroslav Křivánek. A Gradient-Based Framework for 3D Print Appearance Optimization. *ACM Transactions on Graphics*, 40(4), July 2021. ISSN 0730-0301. doi: 10.1145/3450626.3459844. URL <https://doi.org/10.1145/3450626.3459844>. Place: New York, NY, USA Publisher: Association for Computing Machinery.
- James H Nobbs. Kubelka—Munk Theory and the Prediction of Reflectance. *Review of Progress in Coloration and Related Topics*, 15(1): 66–75, 1985. ISSN 1478-4408. doi: 10.1111/j.1478-4408.1985.tb03737.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1478-4408.1985.tb03737.x>.
- Jan Novák, Iliyan Georgiev, Johannes Hanika, Jaroslav Křivánek, and Wojciech Jarosz. Monte Carlo Methods for Physically Based Volume Rendering. In *ACM SIGGRAPH 2018 Courses*, SIGGRAPH '18, pages 14:1–14:1, New York, NY, USA, 2018. ACM. ISBN 978-1-4503-5809-5. doi: 10.1145/3214834.3214880. URL <http://doi.acm.org/10.1145/3214834.3214880>.

- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library, December 2019. URL <http://arxiv.org/abs/1912.01703>. arXiv:1912.01703 [cs, stat].
- Christoph Peters, Sebastian Merzbach, Johannes Hanika, and Carsten Dachsbacher. Using moments to represent bounded signals for spectral rendering. *ACM Transactions on Graphics*, 38(4):136:1–136:14, July 2019. ISSN 0730-0301. doi: 10.1145/3306346.3322964. URL <https://dl.acm.org/doi/10.1145/3306346.3322964>.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically Based Rendering, fourth edition: From Theory to Implementation*. The MIT Press, Cambridge, 4th edition edition, March 2023. ISBN 978-0-262-04802-6.
- John W. Pickering, Scott A. Prahl, Niek van Wieringen, Johan F. Beek, Henricus J. C. M. Sterenborg, and Martin J. C. van Gemert. Double-integrating-sphere system for measuring the optical properties of tissue. *Applied Optics*, 32(4):399–410, February 1993. ISSN 2155-3165. doi: 10.1364/AO.32.000399. URL <https://www.osapublishing.org/ao/abstract.cfm?uri=ao-32-4-399>. Publisher: Optical Society of America.
- M. J. D. Powell. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *The Computer Journal*, 7(2):155–162, January 1964. ISSN 0010-4620. doi: 10.1093/comjnl/7.2.155. URL <https://doi.org/10.1093/comjnl/7.2.155>. eprint: <https://academic.oup.com/comjnl/article-pdf/7/2/155/959784/070155.pdf>.
- Melissa S. Roth and Dimitri D. Deheyn. Effects of cold stress and heat stress on coral fluorescence in reef-building corals. *Scientific Reports*, 3(1):1421, March 2013. ISSN 2045-2322. doi: 10.1038/srep01421. URL <https://www.nature.com/articles/srep01421>.
- Stuart J. Russell, Peter Norvig, and Ernest Davis. *Artificial intelligence: a modern approach*. Prentice Hall series in artificial intelligence. Prentice Hall, Upper Saddle River, 3rd ed edition, 2010. ISBN 978-0-13-604259-4.
- Christopher Sandoval and Arnold D. Kim. Extending generalized Kubelka–Munk to three-dimensional radiative transfer. *Applied Optics*, 54(23):7045, August 2015. ISSN 0003-6935, 1539-4522. doi: 10.1364/AO.54.007045. URL <https://www.osapublishing.org/abstract.cfm?URI=ao-54-23-7045>.
- Craig Schroeder. Practical course on computing derivatives in code. In *ACM SIGGRAPH 2022 Courses*, SIGGRAPH ’22, pages 1–28, New York, NY, USA, August 2022. Association for Computing Machinery. ISBN 978-1-4503-9362-1. doi: 10.1145/3532720.3535643. URL <https://dl.acm.org/doi/10.1145/3532720.3535643>.

- Denis Sumin, Tobias Rittig, Vahid Babaei, Thomas Nindel, Alexander Wilkie, Piotr Didyk, Bernd Bickel, Jaroslav Křivánek, Karol Myszkowski, and Tim Weyrich. Geometry-aware scattering compensation for 3D printing. *ACM Transactions on Graphics*, 38(4):111:1–111:14, 2019. ISSN 0730-0301. doi: 10.1145/3306346.3322992. URL <https://doi.org/10.1145/3306346.3322992>.
- Valery V. Tuchin. Laser Light Scattering in Biomedical Diagnostics and Therapy. *Journal of Laser Applications*, 5(2):43–60, October 1993. ISSN 1042-346X. doi: 10.2351/1.4745330. URL <https://lia.scitation.org/doi/abs/10.2351/1.4745330>. Publisher: Laser Institute of America.
- William E. Vargas. Inversion methods from Kubelka–Munk analysis. *Journal of Optics A: Pure and Applied Optics*, 4(4):452, June 2002. ISSN 1464-4258. doi: 10.1088/1464-4258/4/4/314. URL <https://dx.doi.org/10.1088/1464-4258/4/4/314>.
- Eric Veach. *Robust Monte Carlo methods for light transport simulation*. PhD Thesis, Stanford University, Stanford, CA, USA, December 1997. ISBN: 0-591-90780-1.
- Eric Veach and Leonidas Guibas. Bidirectional Estimators for Light Transport. In Georgios Sakas, Stefan Müller, and Peter Shirley, editors, *Photorealistic Rendering Techniques*, pages 145–167. Springer Berlin Heidelberg, Berlin, Heidelberg, 1995. ISBN 978-3-642-87827-5 978-3-642-87825-1. doi: 10.1007/978-3-642-87825-1_11. URL http://link.springer.com/10.1007/978-3-642-87825-1_11.
- Eric Veach and Leonidas J Guibas. Metropolis light transport. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 65–76, 1997.
- Delio Vicini, Wenzel Jakob, and Anton Kaplanyan. A non-exponential transmittance model for volumetric scene representations. *ACM Transactions on Graphics*, 40(4):136:1–136:16, July 2021a. ISSN 0730-0301. doi: 10.1145/3450626.3459815. URL <https://dl.acm.org/doi/10.1145/3450626.3459815>.
- Delio Vicini, Sébastien Speierer, and Wenzel Jakob. Path Replay Backpropagation: Differentiating Light Paths using Constant Memory and Linear Time. *Transactions on Graphics (Proceedings of SIGGRAPH)*, 40(4):108:1–108:14, August 2021b. doi: 10.1145/3450626.3459804.
- Jiří Vorba and Jaroslav Křivánek. Adjoint-driven Russian roulette and splitting in light transport simulation. *ACM Transactions on Graphics*, 35(4):42:1–42:11, July 2016. ISSN 0730-0301. doi: 10.1145/2897824.2925912. URL <https://dl.acm.org/doi/10.1145/2897824.2925912>.
- Jiří Vorba, Johannes Hanika, Sebastian Herholz, Thomas Müller, Jaroslav Křivánek, and Alexander Keller. Path guiding in production. In *ACM SIGGRAPH 2019 Courses*, SIGGRAPH ’19, pages 1–77, New York, NY, USA, July 2019. Association for Computing Machinery. ISBN 978-1-4503-6307-5. doi: 10.1145/3305366.3328091. URL <https://dl.acm.org/doi/10.1145/3305366.3328091>.

- A. Wilkie and R. F. Tobler. ART: The Advanced Rendering Toolkit, 2022. URL <https://cgg.mff.cuni.cz/ART/>. URL: <https://cgg.mff.cuni.cz/ART/>.
- Alexander Wilkie, Andrea Weidlich, Caroline Larboulette, and Werner Purgathofer. A reflectance model for diffuse fluorescent surfaces. In *Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia*, GRAPHITE '06, pages 321–331, New York, NY, USA, November 2006. Association for Computing Machinery. ISBN 978-1-59593-564-9. doi: 10.1145/1174429.1174484. URL <https://dl.acm.org/doi/10.1145/1174429.1174484>.
- Alexander Wilkie, Petr Vévoda, Thomas Bashford-Rogers, Lukáš Hošek, Tomáš Iser, Monika Kolářová, Tobias Rittig, and Jaroslav Křivánek. A fitted radiance and attenuation model for realistic atmospheres. *ACM Transactions on Graphics*, 40(4):1–14, August 2021. ISSN 0730-0301, 1557-7368. doi: 10.1145/3450626.3459758. URL <https://dl.acm.org/doi/10.1145/3450626.3459758>.
- Douglas R Wyman, Michael S Patterson, and Brian C Wilson. Similarity relations for anisotropic scattering in Monte Carlo simulations of deeply penetrating neutral particles. *Journal of Computational Physics*, 81(1):137–150, March 1989. ISSN 0021-9991. doi: 10.1016/0021-9991(89)90067-3. URL <https://www.sciencedirect.com/science/article/pii/0021999189900673>.
- G. Wyszecki and W. S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulae*. Wiley, 1982.
- Li Yang and Roger David Hersch. Kubelka-Munk Model for Imperfectly Diffuse Light Distribution in Paper. *Journal of Imaging Science and Technology*, 52, May 2008. doi: 10.2352/J.ImagingSci.Technol.(2008)52:3(030201).
- Tizian Zeltner, Sébastien Speierer, Iliyan Georgiev, and Wenzel Jakob. Monte Carlo Estimators for Differential Light Transport. *Transactions on Graphics (Proceedings of SIGGRAPH)*, 40(4), August 2021. doi: 10.1145/3450626.3459807.
- Cheng Zhang, Bailey Miller, Kai Yan, Ioannis Gkioulekas, and Shuang Zhao. Path-space differentiable rendering. *ACM Transactions on Graphics*, 39(4):143:143:1–143:143:19, August 2020. ISSN 0730-0301. doi: 10.1145/3386569.3392383. URL <https://dl.acm.org/doi/10.1145/3386569.3392383>.
- Cheng Zhang, Zihan Yu, and Shuang Zhao. Path-space differentiable rendering of participating media. *ACM Transactions on Graphics*, 40(4):76:1–76:15, July 2021. ISSN 0730-0301. doi: 10.1145/3450626.3459782. URL <https://dl.acm.org/doi/10.1145/3450626.3459782>.
- Shuang Zhao, Ravi Ramamoorthi, and Kavita Bala. High-order similarity relations in radiative transfer. *ACM Transactions on Graphics (TOG)*, 33(4):1–12, 2014. Publisher: ACM New York, NY, USA.

- Shuang Zhao, Wenzel Jakob, and Tzu-Mao Li. Physics-based differentiable rendering: from theory to implementation. In *ACM SIGGRAPH 2020 Courses*, SIGGRAPH '20, pages 1–30, New York, NY, USA, August 2020. Association for Computing Machinery. ISBN 978-1-4503-7972-4. doi: 10.1145/3388769.3407454. URL <https://dl.acm.org/doi/10.1145/3388769.3407454>.
- Shuang Zhao, Ioannis Gkioulekas, and Sai Bangaru. CVPR 2021 Tutorial on Physics-Based Differentiable Rendering, June 2021. URL <https://diff-render.org/tutorials/cvpr2021/>.
- Yinqiang Zheng, Imari Sato, and Yoichi Sato. Spectra Estimation of Fluorescent and Reflective Scenes by Using Ordinary Illuminants. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, Lecture Notes in Computer Science, pages 188–202, Cham, 2014. Springer International Publishing. ISBN 978-3-319-10602-1. doi: 10.1007/978-3-319-10602-1_13.
- Yinqiang Zheng, Ying Fu, Antony Lam, Imari Sato, and Yoichi Sato. Separating Fluorescent and Reflective Components by Using a Single Hyperspectral Image. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 3523–3531, December 2015. doi: 10.1109/ICCV.2015.402. ISSN: 2380-7504.

List of Figures

| | | |
|------|--|----|
| 2.1 | Visualization of the fundamental radiometric quantities. | 15 |
| 2.2 | A sensor may consist of multiple pixels. Each pixel is sensitive to radiance incoming from a certain range of directions and wavelengths, and over a certain period of time. The pixels may have different sensitivities, which is common in color sensors sensitive to red, green, and blue. | 16 |
| 2.3 | A simple model of light incoming to a sensor through the lens and aperture of a camera. | 17 |
| 2.4 | Spectral sensitivities of the CIE 1931 standard colorimetric observer based on real human vision (a), a color camera (b) [Kawakami et al., 2013], and a multispectral sensor (c). | 18 |
| 2.5 | The radiance that reaches the camera sensor may come from anywhere in the scene, and it is common that the light reflects multiple times between the light source and the camera. | 19 |
| 2.6 | An identical shape rendered with four different materials: gold, plastic, diffuse, and glass. Each material scatters the incoming light differently, which is formally described by their BSDF. The diagrams above visualize the light interactions on each material’s surface. | 20 |
| 2.7 | A visualization of light being scattered through a participating medium enclosed in a glass container. | 23 |
| 2.8 | A diagram of the light interactions in a participating medium. | 23 |
| 2.9 | Practical visualization of how the individual parameters (index of refraction, absorption coefficient, and scattering coefficient) influence the actual appearance of an object. The object’s shape and its surrounding are identical. | 24 |
| 2.10 | The energy of a collimated pencil of light decreases as it travels through a medium. | 25 |
| 2.11 | The effect of varying anisotropy factors g in different illumination conditions. When looking against the light source (a), forward scattering appears brighter than backward scattering. When looking from the light source position (b), backward scattering appears brighter as it reflects back toward the camera. | 26 |
| 2.12 | For a collimated pencil of light entering the medium, the outgoing radiance in the beam changes due to absorption, in-scattering, out-scattering, and emission. | 27 |
| 2.13 | The volume rendering equation describes the light transport in a scene that contains not only surfaces but also participating media. | 28 |
| 2.14 | Photographs of fluorescent (sticky notes, liquid detergent) and non-fluorescent (color chart) materials under different illuminants. Under daylight, everything appears colorful. The monochrome illuminants reveal that the sticky notes and the liquid are fluorescent materials emitting different wavelengths than they received from the illuminants. | 29 |

| | | |
|------|---|----|
| 2.15 | The reflectance of a fluorescent surface can be modeled as a sum of individual contributions. The contributions can be, for example, a glossy reflection term with no fluorescence, and a separate diffuse fluorescent term. | 31 |
| 2.16 | Regular reflection occurs on the surface of a material. In the layered model, fluorescence occurs under the surface in a certain depth in which the molecule or particle actually re-emitted the photon. In the simplified model, both the reflection and fluorescence are modeled to happen on the surface at the same position. | 31 |
| 2.17 | Illustration of Eqs. (2.24) and (2.25). The reflected spectrum (a) is computed from the Donaldson matrix (b) and illuminant spectrum (c). While the illuminant is green, the fluorescent reflection is orange and has two peaks. | 32 |
| 3.1 | Demonstration of a trivial inverse problem. The parameter π of the reference image (left) is found by rendering the same scene with varying π and then picking the one whose image is the most similar to the reference. | 34 |
| 3.2 | Two rabbit models observed under three different illuminants defined in the CIE standard: D65 (standard daylight), A (incandescent bulb), and FL4 (fluorescent lighting). While the materials look identical under D65, observing the rabbits under A and FL4 reveals that their reflectance spectra actually differ. The true illuminant and reflectance spectra are visible in (a), (b). | 36 |
| 3.3 | A comparison of applying Beer's law, Eq. (2.10), to spectral extinction coefficients $\sigma_t(\lambda)$, versus applying it to RGB triplets ($\sigma_t^R, \sigma_t^G, \sigma_t^B$). The solid lines in the plots show original spectral data for various colored materials. The dotted lines show the resulting spectral shapes for shorter (above the solid line) and longer (below) transmission distances. The colored squares show the resulting transmission colors based on spectral data (top) and on RGB data (bottom). While the outlined squares match as the RGB extinction coefficients were fitted to them, the colors of all other squares diverge from the correct spectral results. | 37 |
| 3.4 | Demonstrating the similarity relations using the Henyey-Greenstein phase function. Varying anisotropy factor g influences the appearance of the material (a). However, materials whose scattering coefficient σ_s and anisotropy g follow Eq. (3.5) look similar or even identical (b). | 38 |
| 3.5 | The inverse transport network consists of a sub-network that takes an image and returns its estimated parameters, and a differentiable renderer that converts the parameters back into an image. During training, the network is conditioned to ensure the input and output images are similar. | 42 |