Reconsidering Light Transport

Acquisition and Display of Real-World Reflectance and Geometry

Dissertation zur Erlangung des Grades des Doktors der Ingenieurwissenschaften der Naturwissenschaftlich-Technischen Fakultäten der Universität des Saarlandes

Vorgelegt durch

Matthias B. Hullin Max-Planck-Institut Informatik Campus E1 4 66123 Saarbrücken Deutschland

am 7. Oktober 2010 in Saarbrücken

Betreuender Hochschullehrer -- Supervisor

Prof. Dr. Hans-Peter Seidel, MPI Informatik, Saarbrücken, Deutschland

Gutachter -- Reviewer

Prof. Dr. Hans-Peter Seidel, MPI Informatik, Saarbrücken, Deutschland Prof. Dr. Hendrik P. A. Lensch, Universität Ulm, Deutschland Prof. Dr. Wolfgang Heidrich, University of British Columbia, Vancouver, Kanada

Dekan -- Dean

Prof. Dr. Holger Hermanns, Universität des Saarlandes, Saarbrücken, Deutschland

Kolloquium -- Examination

Datum -- Date: 15. Dezember 2010

Vorsitzender -- Chair: Prof. Dr. Philipp Slusallek, Universität des Saarlandes, Saarbrücken, Deutschland

Prüfer -- Examiners: Prof. Dr. Hans-Peter Seidel, MPI Informatik, Saarbrücken, Deutschland Prof. Dr. Hendrik P. A. Lensch, Universität Ulm, Deutschland

Protokoll -- Reporter: Dr. Ivo Ihrke, Universität des Saarlandes, Saarbrücken, Deutschland

The flick book animations at the bottom of this thesis show a dataset of a glass mouse acquired by the method of Fluorescent Immersion Range Scanning presented in Chapter 3. Odd pages depict slices through the object volume as seen by the camera as the laser illuminates different planes. When the book is closed, this stack of slices forms a volumetric model of the object. On even pages, a rendering of the reconstructed geometry can be seen.



Abstract

In this thesis, we cover three scenarios that violate common simplifying assumptions about the nature of light transport.

We begin with the first ingredient to any 3D rendering: a geometry model. Most 3D scanners require the object-of-interest to show diffuse reflectance. The further a material deviates from the Lambertian model, the more likely these setups are to produce corrupted results. By placing a traditional laser scanning setup in a participating (in particular, fluorescent) medium, we have built a light sheet scanner that delivers robust results for a wide range of materials, including glass.

Further investigating the phenomenon of fluorescence, we notice that, despite its ubiquity, it has received moderate attention in computer graphics. In particular, to date no datadriven reflectance models of fluorescent materials have been available. To describe the wavelength-shifting reflectance of fluorescent materials, we define the bispectral bidirectional reflectance and reradiation distribution function (BRRDF), for which we introduce an image-based measurement setup as well as an efficient acquisition scheme.

Finally, we envision a computer display that shows *materials* instead of *colours*, and present a prototypical device that can exhibit anisotropic reflectance distributions similar to common models in computer graphics.

3

Kurzzusammenfassung

In der Computergraphik und Computervision ist es unerlässlich, vereinfachende Annahmen über die Ausbreitung von Licht zu machen. In dieser Dissertation stellen wir drei Fälle vor, in denen diese nicht zutreffen.

So wird die dreidimensionale Geometrie von Gegenständen oft mit Hilfe von Laserscannern vermessen und dabei davon ausgegangen, dass ihre Oberfläche diffus reflektiert. Dies ist bei den meisten Materialien jedoch nicht gegeben, so dass die Ergebnisse oft fehlerhaft sind. Indem wir das Objekt in einem fluoreszierenden Medium einbetten, kann ein klassischer 3D-Scanner-Aufbau so modifiziert werden, dass er verlässliche Geometriedaten für Objekte aus verschiedensten Materialien liefert, einschließlich Glas.

Auch die akkurate Nachbildung des Aussehens von Materialien ist wichtig für die photorealistische Bildsynthese. Wieder interessieren wir uns für Fluoreszenz, diesmal allerdings für ihr charakteristisches Erscheinungsbild, das in der Computergraphik bislang kaum Beachtung gefunden hat. Wir stellen einen bildbasierten Aufbau vor, mit dem die winkelund wellenlängenabhängige Reflektanz fluoreszierender Oberflächen ausgemessen werden kann, und eine Strategie, um solche Messungen effizient abzuwickeln.

Schließlich befassen wir uns mit der Idee, nicht nur *Farben* dynamisch anzuzeigen, sondern auch *Materialien* und ihr je nach Lichteinfall und Blickwinkel unterschiedliches Aussehen. Einer generellen Beschreibung des Problems folgt die konkrete Umsetzung in Form zweier Prototypen, die verschiedene Reflektanzverteilungen auf einer Oberfläche darstellen können.



Summary

In computer vision and graphics, many common techniques are based on simplifying assumptions about the nature of light and its propagation. In this thesis, we treat a selection of cases where these assumptions are violated, causing standard methods to break down. For three example scenarios, we provide solutions that do not only counteract the undesired effects but also leverage some of the new properties in order to obtain novel results.

In the first part of the thesis, we deal with the acquisition of surface geometry. The image formation model behind most 3D scanning techniques assumes diffuse (Lambertian) reflectance. These techniques fail when confronted with glossy materials, translucent or transparent media, or surfaces of very low albedo, causing the resulting geometry to be incomplete and corrupted by noise. Building on a traditional laser scanning setup, our *fluorescent immersion range scanner* uses a participating medium to make light rays visible before they hit the object surface. The resulting light-sheet scanner allows to robustly acquire a wide range of non-diffuse materials, including glass objects. For homogeneous transparent materials, we can directly obtain high-quality volume sections and assemble volumetric models akin to tomographic reconstructions.

The second part concerns the acquisition and rendering of fluorescent materials themselves. Fluorescence is ubiquitous in every-day life, yet in computer graphics its colourshifting behaviour constitutes a major inconvenience in any rendering pipeline. A few researchers have dealt with the problem of fluorescence rendering, but to date no data-driven reflectance models of these materials have been available. Representing the wavelength alteration on fluorescent surfaces, and its dependence on the angles of incidence and exitance, calls for a six-dimensional bispectral and bidirectional reflectance and reradiation distribution function (*bispectral BRRDF*). We propose an image-based setup for the acquisition of such distributions for isotropic fluorescent surfaces. The dataset is of high dimensionality (three angular and two spectral dimensions), but can be approximated efficiently by a low-rank decomposition. We acquire a sparse subset of the full data and use it to steer the remaining measurement, allowing for streamlined acquisition of such materials.

After investigating the acquisition of geometry and appearance, we conclude with an outlook into the future of computer display. So far, only *colours* have been displayed, but along with an increasing demand for realism and immersion, at some point there may be a need to display *materials* and the characteristic way they reflect light under different angles. We define the problem of displaying reflectance, relate it to existing techniques, and propose a few approaches to its solution. Finally, we demonstrate a prototypical device that can be programmed to exhibit various degrees of surface roughness. We show that our device can exhibit anisotropic bell-shaped microfacet statistics, through which it is related to many common reflectance models in the field of computer graphics and beyond.

Zusammenfassung

In der Computergraphik besteht ein zentrales Problem in der Berechnung synthetischer Bilder, die die Realität, entsprechend etwa einem Foto, möglichst akkurat abbilden. Technisch gesehen entspricht das einer numerischen Simulation der Ausbreitung von Licht in der realen Welt. Je näher die verwendeten Modelle an der physikalischen Wirklichkeit sind, umso verlässlicher ist das Ergebnis dieser Simulation. Die Schwesterdisziplin der Graphik, die Computervision, befasst sich im Gegenzug mit der Interpretation von visuellen Messdaten, die die reale Welt beschreiben. Auch hier spielt die Lichtausbreitung eine wichtige Rolle, denn nur wenn das der Analyse zu Grunde liegende Modell die Wirklichkeit hinreichend beschreibt, lassen sich die Messdaten korrekt interpretieren.

Nun ist die Realität jedoch so komplex, dass sie unmöglich vollständig beschrieben werden kann. Keine der beiden Disziplinen kommt daher umhin, explizit oder implizit stark vereinfachende Annahmen über die Natur des Lichts zu machen.

Wir befassen uns in dieser Dissertation mit solchen Situationen, in denen etablierte Techniken scheitern, weil die Wirklichkeit eben nicht immer den Annahmen entspricht. Anhand von drei Beispielfällen aus verschiedenen Gebieten zeigen wir Lösungen auf, mit denen wir nicht nur bestehende Verfahren um je einen neuen optischen Effekt bereichern, sondern bestimmte Phänomene und Eigenschaften ganz gezielt ausnutzen, um neuartige Ergebnisse zu erreichen.

Die wichtigste Zutat für die dreidimensionale Darstellung jedes Gegenstands ist das Wissen über seine geometrische Form. Für ein gegebenes Objekt kann die Geometrie beispielsweise unter Verwendung eines 3D-Scanners erfasst werden. Die meisten Techniken gehen allerdings davon aus, dass das Objekt eine diffuse Oberfläche besitzt, und liefern entsprechend fehlerhafte Daten zurück, wenn dies so nicht gegeben ist. Insbesondere stark glänzende Materialien, durchscheinende oder transparente Objekte – und nicht zuletzt sehr dunkle Oberflächen – stellen eine große Herausforderung für traditionelle 3D-Scanner dar. Wir stellen daher eine Erweiterung für laserbasierte Systeme vor, bei der das Objekt in einer fluoreszierenden Flüssigkeit eingebettet wird. Dadurch wird jeder Lichtstrahl, mit dem das Objekt angeleuchtet wird, auch vor seinem Auftreffen bereits sichtbar. Anstatt uns auf Reflexionen von der Oberfläche zu verlassen, definieren wir die Oberfläche des Objekts als das sichtbare Ende der Lichtstrahlen im Volumen. Damit können nun Objekte aus vielen verschiedenen Materialien zuverlässig vermessen werden, insbesondere solche aus Glas.

Geometriedaten allein reichen aber nicht aus, um photorealistische Bilder zu erzeugen; ebenso wichtig ist ein Modell für das Erscheinungsbild, also die Reflektanz des Materials, aus dem ein Gegenstand besteht. Nachdem wir Fluoreszenz schon als Hilfsmittel verwendet haben, um dreidimensionale Geometrie zu vermessen, wenden wir uns nun dem charakteristischen Aussehen von Objekten zu, die selbst fluoreszieren. Obwohl Fluoreszenz im Alltag eine enorm wichtige Rolle spielt, ist sie in der Computergraphik bislang eher stiefmütterlich behandelt worden. Ihre Simulation erfordert es nämlich, die Wellen-



längenverschiebung nachzubilden, die fluoreszent reflektiertes Licht erfährt. Die meisten Verfahren zur Bildsynthese gehen jedoch davon aus, dass die Wellenlänge des Lichts stets konstant bleibt. In der wissenschaftlichen Literatur findet man eine Handvoll Techniken, die das Erscheinungsbild fluoreszierender Substanzen simulieren können, jedoch keinerlei Reflektanzmodelle für reale Materialien, die auf gemessenen Daten basierten. In dieser Arbeit stellen wir einen bildbasierten gonioreflektometrischen Aufbau vor, mit dem die Reflektanz fluoreszierender Materialien abhängig von der Einfalls- und der Ausfallsrichtung des Lichts ausgemessen werden kann. Mathematisch beschreiben wir dies in Form der bispektralen bidirektionalen Reflektanz- und Reradiationsverteilungsfunktion (bispektrale BRRDF). Mit ihren vier Winkel- und zwei spektralen Argumenten ist die vollständige Ausmessung einer bispektralen BRRDF sehr zeitaufwendig. Wir zeigen daher, wie eine Hauptkomponentenanalyse auf einer Untermenge der Daten genutzt werden kann, um einen kompletten Datensatz effizient zu erfassen. Mit dem vorgestellten Verfahren ist es erstmals möglich, das visuelle Erscheinungsbild realer fluoreszierender Oberflächen in synthetischen Renderings nachzubilden.

Nach diesen Problemen, die sich um die Vermessung von Geometrie und Aussehen ranken, begeben wir uns nun auf die Ausgabeseite. Hier hat man Anzeigegeräte bislang so ausgelegt, dass sie unter einer Vielzahl von Beleuchtungs- und Betrachtungsbedingungen möglichst immer die gleichen Farben anzeigen sollten. Die wirkliche Welt, deren realistische Abbildung ja oft das Ziel ist, kann aber je nach Lichteinfall und Blickwinkel sehr unterschiedlich aussehen. Im letzten unserer Beispiele entwickeln wir auf dieser Grundlage die Idee, nicht nur verschiedene Farben, sondern auch verschiedene Materialien und ihre charakteristische Reflektanz auf einer Oberfläche dynamisch anzuzeigen. Dazu definieren wir zunächst das Problem im Allgemeinen und skizzieren verschiedene Ansätze zu seiner Lösung, um schließlich zwei Prototypen solcher Reflektanz-Displays vorzustellen, die einfallendes Licht in genau kontrollierbaren, anisotropen Winkelverteilungen reflektieren können. Mit ähnlichen Verfahren könnte es in Zukunft möglich werden, Fabrikationsprozesse vorab dynamisch zu visualisieren.

7

Acknowledgements

My sincere thanks go to my supervisor Prof. Dr. Hans-Peter Seidel for his support and most valuable advice during my time at the MPI, and for providing an excellent research environment for this work.

My co-advisor Prof. Dr. Hendrik P. A. Lensch, who lit my interest in computer graphics and supported me all the way through this work, deserves my deepest gratitude. Not only did he contribute a great many of ideas but also a lot of hands-on help in critical situations.

I also thank Prof. Dr. Wolfgang Heidrich, who agreed to serve as external reviewer.

My thanks also go to Prof. Paul Debevec (University of Southern California) and Prof. Ramesh Raskar (Massachusetts Institute of Technology), who welcomed me as a visitor in their groups.

In times of transition, it is invaluable to have supportive colleagues. In particular, I thank Dr. Christel Weins, Dr. Christian Fuchs and Dr. Martin Fuchs who helped me to quickly find a place in my new environment. In the last year, Dr. Ivo Ihrke kindly adopted me as a member of his newly founded group, which helped me in many ways. All in all, the work constituting this thesis would not have been possible if it weren't for the cooperation and discussion with, and the motivation and critical feedback from my colleagues. I would like to thank the following people (in alphabetical order): Boris Ajdin, Dr. Tongbo Chen, Miguel Granados, Peter Grosche, Johannes Hanika, Dr. Karol Myszkowski, Dr. Thomas Schultz, Dr. Carsten Stoll, Dr. Robert Strzodka, Martin Sunkel, Art Tevs, and Dr. Michael Wand.

The Max Planck Institute is run by an excellent team of technicians and supporting staff who make working here a great experience. The service delivered by these people is truly exceptional, even when compared to other top-notch institutions. I therefore thank our technicians Axel Koeppel, Michael Laise and Uwe Meiser, the IT department (especially Stefan Bender, Uwe Brahm, Simon Hirth and Thomas Hirtz), Rüdiger Gronwald, Annette Hartz and Heike Löx, and our secretaries Sabine Budde, Ellen Fries and Conny Liegl.

Financial support for parts of this dissertation has been granted by the Max-Planck-Gesellschaft zur Förderung der Wissenschaften e.V., by the DFG Emmy Noether Scholarship Le 1341/1-1 and the Max Planck Center for Visual Computing and Communication (BMBF-FKZ011MC01) through the support of Dr. Lensch's research group.

Finally, I thank Katharina and my family for their tireless support in all questions that *really* matter.



CONTENTS

9

Contents

1	Introduction							
	1.1	Contributions						
	1.2	Fluore	scent Immersion Range Scanning	14				
	1.3	Acquis	sition and Analysis of Bispectral BRRDFs	14				
	1.4	amic Display for BRDFs	15					
	1.5	Outlin	e	15				
2	Background 17							
	2.1	Physic	al Description of Light	17				
		2.1.1	Geometric Optics	17				
		2.1.2	Wave Optics	18				
		2.1.3	Ouantum Optics	22				
		2.1.4	So Which One Is It?	23				
	2.2	Measu	ring Light	25				
		2.2.1	Radiometric Quantities	25				
		2.2.2	Spectrometry	26				
		2.2.3	Photometry	26				
	2.3	 23 Interaction between Light and Matter 						
	2.0	2.3.1	Kev Mechanisms	 27				
		2.3.2	Abstraction: Bidirectional Distribution Functions	28				
	2.4	utational Light Transport	30					
	2.1	2.4.1	Solving Maxwell's Equations	31				
		2.4.2	Rendering Equations	31				
		2.1.2		51				
3	Fluo	rescent	Immersion Range Scanning	35				
	3.1	3D Sca	Scanning of Optically Challenging Materials					
	3.2 Related Work		d Work	37				
		3.2.1	Range Scanning: Acquisition of Surface Geometry	39				
		3.2.2	The Struggle with Non-Lambertian Materials	39				
		3.2.3	Tomography	41				
	3.3	Makin	g Light Rays Visible	41				
		3.3.1	Elastic Scattering in Dilute Milk	42				
		3.3.2	Inelastic Fluorescence Scattering	43				

		3.3.3	Comparison
	3.4	Measur	rement Setup
		3.4.1	General Setup
		3.4.2	Calibration 48
		3.4.3	Triangulation
		3.4.4	Data Acquisition
	3.5	3D Sur	face Scanning
		3.5.1	Influence of the Material 50
		3.5.2	Detection Algorithm
	3.6	scan Direct Volume Capturing	
		3.6.1	Matching the Refractive Index
		3.6.2	Volume Processing
	3.7	Results	56
		3.7.1	Comparison to Existing Methods 59
		3.7.2	Applicability and Limitations
	3.8	Conclu	sion
4	Bisp	ectral B	RRDFs 63
	4.1	Related	l Work
		4.1.1	Analytical and Data-Driven BRDF Models
		4.1.2	Colour Characterisation and Spectral Imaging 65
		4.1.3	Bispectral Measurement: Fluorometry 65
		4.1.4	Rendering of Fluorescence65
	4.2 Reflection and Reradiation		ion and Reradiation
			ral Reflectance and Reradiation
		4.3.1	Bispectral Rendering Equation
		4.3.2	Bispectral BRRDF 68
	4.4	Measur	rement and Reconstruction
		4.4.1	Setup
		4.4.2	Geometric and Spectral Calibration
		4.4.3	Measurement and Data Processing 71
	4.5	PCA-ba	ased Acquisition
		4.5.1	Insights
	4.6	Results	77
		4.6.1	Acquired Datasets
		4.6.2	Spectral Detail Level
	4.7	Conclu	sion
_	_		
5 Dynamic BRDF Display		amic BF	RDF Display 83
	5.1	Related	l Work
		5.1.1	Representation and acquisition of reflectance
		5.1.2	View- and light-dependent rendering and display 85
		5.1.3	Fabrication 86
		5.1.4	Our Contribution 86



CONTENTS

5.2	Displaying Reflectance			
	5.2.1	Approaching the Problem	88	
5.3	Characterisation of Reflectance by Surface Waves			
	5.3.1	Single Sine Wave in One Dimension	91	
	5.3.2	Multiple Sine Waves in One or Two Dimensions	92	
	5.3.3	Connection to Analytical BRDF Models	93	
5.4	Constr	ruction of Devices	94	
5.5 Results		S	95	
	5.5.1	Characterisation	95	
	5.5.2	Reflectance	98	
	5.5.3	Diffuse + Specular	98	
	5.5.4	Liquid Metal	99	
5.6	Discus	ssion	99	
Con	clusion 10			
Refle	ectance	of Sinusoidal Height Fields	105	
A.1	Reflect	tion from a Sine Wave	105	
A.2 Linearisation, Mean and Variance of the Directional Distribution				
A.3Superposition of Sine WavesA.4Fresnel Reflection				
				A.5
16	Two I)imensional Height Fields	111	
A.0	TWO-L		111	
	 5.2 5.3 5.4 5.5 5.6 Cone Reflet A.1 A.2 A.3 A.4 A.5 A.6 	 5.2 Displa 5.2 Displa 5.3.1 5.3.2 5.3.3 5.4 Constribution 5.5 Result 5.5.1 5.5.2 5.5.3 5.5.4 5.6 Discuss Conclusion Reflectance A.1 Reflect A.2 Linear A.3 Superp A.4 Fresne A.5 Relation A 6 Two-F 	 5.2 Displaying Reflectance	

11

CONTENTS

Chapter 1 Introduction

In computer vision and graphics, many common methods are based on simplifying assumptions about the nature of light, such as the following:

- The object we are about to 3D scan has a diffuse surface.
- The wavelength of light always remains constant.
- The output of display devices should be as invariant as possible to changes in light or observer position.

In this work, we deliberately violate the above, which causes standard techniques to break down and, for instance, deliver corrupted 3D models or incorrect renderings. We will show that deviations from these common assumptions can not only be cured, but also exploited in order to obtain unprecedented results. The choice of scenarios covers a range of graphics and vision research fields, including geometry and appearance acquisition, rendering, and display.

1.1 Contributions

Some of the techniques covered in this work have already been published and presented at international research conferences [Hullin08a, Hullin08b, Hullin10], or are currently in submission[Hullin11]. In this thesis, we present them in revised and extended form. Our key contributions are:

• A simple extension to a conventional 3D scanning setup to allow for robust geometry measurements on transparent, translucent and very dark materials. We can directly sample geometry slices of homogeneous clear objects with results better than tomographic techniques.

- A bispectral extension of the BRDF formalism to describe the appearance of fluorescent materials, as well as a setup and a practical method for its acquisition based on a low-rank approximation.
- The concept of BRDF display, and a prototypical device that "displays" different degrees of surface roughness, exhibiting a range of anisotropic reflectance distributions.

1.2 Fluorescent Immersion Range Scanning

The key ingredient in the reproduction of real-world objects is a geometry model. In Chapter 3 we investigate the problem of obtaining robust 3D scans of uncooperative materials. Traditional triangulation-based laser scanning systems rely on the diffuse reflectance of the acquired surface. In the presence of strong surface gloss, sub-surface light transport or very low albedo, the detection of surface points may not work reliably, causing the resulting geometry to be incomplete and corrupted by noise. Instead of attempting to detect surface points directly, our approach traces the propagation of light before it hits the object. We have modified a traditional laser scanning setup by placing the object in a participating (in particular, fluorescent) medium, which renders the illuminating light sheet visible in the volume. The surface of the object is defined by the cut-off line of this laser sheet, which can be detected using a few simple criteria. This enables a reliable surface acquisition for a wide variety of materials that are otherwise difficult to handle, such as translucent, transparent and very dark materials and mixtures thereof. For homogeneous clear materials, a simple extension of this approach allows for the direct sampling of the object volume, yielding results that outperform tomographic approaches in quality, but do not require a numerical reconstruction step.

1.3 Acquisition and Analysis of Bispectral BRRDFs

After using fluorescence as a tool for the acquisition of geometry, Chapter 4 deals with the appearance of fluorescent materials themselves and its representation in a computer graphics context. While these materials play an important role in everyday life, they have received moderate attention in computer graphics due to their peculiar optical properties. In particular, they may absorb energy in the invisible ultraviolet range and reradiate it as visible light, resulting in their typical brighter-than-white appearance. We define this wavelength-shifting behaviour and its directional dependence in terms of a bispectral bidirectional reflectance and reradiation distribution function (bispectral BRRDF). Using a bispectral and bidirectional measurement setup, we have acquired such distributions for a set of fluorescent materials. The measurement setup is steered by a low-rank decomposition in order to speed up the acquisition of these high-dimensional datasets.



1.4 A Dynamic Display for BRDFs

We now leave the acquisition of geometry and reflectance aside, and turn to output, i.e., display, instead. Traditionally, most display devices have been limited to showing a *colour* per pixel that should be as invariant to the surrounding light and the viewing angle as possible. However, the real world does not behave like this: the visual impression of a material is defined by the way its appearance changes under different lighting and viewing conditions. In order to achieve a true immersive experience, at some point it may therefore become desirable to display *materials* instead of *colours*. In Chapter 5, we present two prototypes of devices that can be programmed so as to exhibit different reflectance characteristics, mimicking the different ways in which real-world materials reflect light. We provide the theoretical foundations required for understanding the connections between the height fields generated by our displays, and the resulting reflectance distributions, and show that our design can be used to achieve a Gauss-like highlight shape. We characterise our prototypes with regard to the theoretical predictions, discuss their limitations and finally propose possible directions to making such display technology more practical and increasing the range of displayable BRDFs.

1.5 Outline

This thesis is structured as follows: After this introduction, we proceed with the physical foundations that are of importance throughout the thesis (Chapter 2). From Chapter 3 to Chapter 5, our own contributions are presented in detail. Each of these chapters comes with its own detailed review of related work, as well as a discussion of the limitations and future possibilities of the respective technique. The conclusions are drawn in Chapter 6. As an appendix, we provide an in-depth derivation of the Fresnel reflectance of sinusoidal height fields as used by our BRDF displays.

CHAPTER 1. INTRODUCTION

Chapter 2 Background

In this chapter, we will review the theory most commonly used to describe light and related phenomena. We will focus on the aspects that are most important for the understanding of this thesis, and place them in the context of computer graphics. Readers looking for in an in-depth treatment of physical background are kindly referred to a text book, for instance the optics classics by [Born99] and [Hecht01]. As an additional interesting read, we recommend the article by [Banks07] which closes the gap to the very foundations of quantum electrodynamics and is specifically targeted to a computer graphics audience.

2.1 Physical Description of Light

Light is a phenomenon that can be described by a variety of different models, the most common of which will be introduced in the following. Although they might seem contradictory, each of these models has its domain of applicability. Together they form a practical system to describe the vast majority of optical phenomena.

2.1.1 Geometric Optics

The most basic way to think about light is in terms of rays. In vacuum or a homogeneous optical medium, light is assumed to propagate along straight lines. Any ray can be characterised by a point \vec{r} and a direction \vec{d} in space. On interfaces between different media, rays are mirrored and/or refracted, the latter being governed by Snell's Law:

$$\frac{n_1}{n_2} = \frac{\sin(\alpha_2)}{\sin(\alpha_1)} = \frac{c_2}{c_1}$$
(2.1)

where n_i are the refractive indices, α_i the angles of the rays with respect to the normal, and c_i the speeds of light in the respective media. We will find in Chapter 3.4.2 that

CHAPTER 2. BACKGROUND

this nonlinear relation invalidates common assumptions about camera projections when looking into refractive media.

The geometric model for light allows for very simple and effective computation of optical elements, image formation, and, in fact, most macroscopic phenomena (see, for instance, Figure 2.1). Therefore, this is by far the most commonly used light model in computer graphics and vision.



Figure 2.1: A parabolic mirror reflects a parallel bundle of light rays into one point. Settings like this can be constructed on a purely geometrical basis by mirroring ray direction vectors with respect to the surface normal.

When it comes to quantifying amounts of light (Section 2.2.1), discrete rays can be replaced by a differential formalism.

The geometric model fails when light interacts with structures on a scale in the order of the wavelength, or for coherent light, where the electromagnetic wave carries a welldefined phase in space and/or time. These cases give rise to phenomena such as diffraction and interference that can only be described with waves. Note that even Snell's law (Eq. 2.1), although applicable in a purely geometric framework, can only be derived using a notion of light waves.

2.1.2 Wave Optics

Like many other physical phenomena, light can be described in terms of waves: disturbances that propagate through time and space. While some types of waves are bound to a carrier matter (sound waves, water waves), the electromagnetic field that is modulated by light waves can also exist in vacuum. What constitutes a wave, as opposed to any arbitrary spatio-temporal function, is the coupling between the spatial and temporal dimensions, and the constant exchange between two forms of energy. For electromagnetic waves such as light, these are the energies of the electric and magnetic fields.

The evolution of the electromagnetic field over time and space is governed by Maxwell's



equations:

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\varepsilon} \qquad \nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$$
$$\nabla \cdot \mathbf{B} = 0 \qquad \nabla \times \mathbf{B} = \mu \mathbf{J} + \mu \varepsilon \frac{\partial \mathbf{E}}{\partial t}$$

where **E** and **B** are the electric and magnetic fields, ρ is the density of the electric charge and **J** the electric current density (**J** = 0 in charge-free space). $\varepsilon = \varepsilon_r \varepsilon_0$ is the electric permittivity of the medium and composed of the relative permittivity ε_r , a material constant, and the free-space permittivity ε_0 (formerly also called the "dielectric constant of vacuum"). Similarly, the magnetic permeability $\mu = \mu_r \mu_0$ is the product of the material constant μ_r (for non-magnetic materials \approx 1) and the magnetic permeability of free space, μ_0 .

Let us assume that there is no charge in space: $\rho = 0$, $\mathbf{J} = 0$. By taking the curl of the curl equations and using the vector identity $\nabla \times (\nabla \times \mathbf{V}) = \nabla (\nabla \mathbf{V}) - \nabla^2 \mathbf{V}$, we obtain a pair of wave equations for the electric and magnetic fields:

$$\left(\nabla^2 - \mu \varepsilon \frac{\partial^2}{\partial t^2}\right) \mathbf{E} = 0 \tag{2.2}$$

$$\left(\nabla^2 - \mu \varepsilon \frac{\partial^2}{\partial t^2}\right) \mathbf{B} = 0 \tag{2.3}$$

The phase velocity, or speed of light, is given as $c = (\mu \varepsilon)^{-1/2}$. Assuming a nondispersive medium ($\mu = \text{const}$, $\varepsilon = \text{const}$), the general solution to the wave equation for the electric field (Eq. 2.2) is a superposition of planar waves of the form

$$\mathbf{E}(\mathbf{r},t) = g(\varphi(\mathbf{r},t)) = g(\omega t - \mathbf{k} \cdot \mathbf{r})$$
(2.4)

where ω is the angular frequency (radians per second) and **k** is the so-called *wave vector*. The speed of light *c* relates the spatial frequency *k* with the temporal frequency ω as $\omega = ck$. Such planar waves are the equivalent of the rays in the geometric model, with the direction of the ray being defined by the direction of **k** and the light spectrum given by the function $g(\varphi)$. The wave fronts (surfaces of constant phase $\varphi = \omega t - \mathbf{k} \cdot \mathbf{r}$) are planes of infinite extent that are perpendicular to the ray direction. In the real world, this assumption will not hold, which leads to diffraction, as explained later in this section.

Using a one-dimensional Fourier decomposition, each planar wave $g(\varphi)$ can be written as a superposition of sinusoidal terms, such that

$$\mathbf{E}(\mathbf{r},t) = \sum_{j} \hat{\mathbf{E}}_{j} e^{i(v_{j}t - \mathbf{k}_{j} \cdot \mathbf{r} + \varphi_{j}^{(0)})}$$
(2.5)

with polarisation (and amplitude) and wave vectors $\hat{\mathbf{E}}_j \perp \mathbf{k}_j$, angular frequency v_j , and phase offset $\varphi_i^{(0)}$ per term. *i* is the imaginary unit. Each individual term describes a



monochromatic, linearly polarised light wave where the **E**-field oscillates in the plane that is spanned by $\hat{\mathbf{E}}$ and \mathbf{k} . For such waves, ν equals to the corresponding angular frequency, and a wavelength λ can be defined such that $k = 2\pi/\lambda = \nu/c$. Note that the energy content of each term (cf. Section 2.2.1) is proportional to $\hat{\mathbf{E}}^2$.

The linearly polarised planar waves form a complete basis in which all light propagation in homogeneous media can be expressed. In particular, any polarisation state can be described as superposition of linearly polarised components of different direction and phase. The possible states include linear polarisation in any plane, circular polarisation as sum of two terms with perpendicular polarisation vectors $\hat{\mathbf{E}}_1 \perp \hat{\mathbf{E}}_2$ and a relative phase shift $\Delta \varphi = \varphi_1^{(0)} - \varphi_2^{(0)} = \pm \pi/2$, and the more general case of elliptic polarisation. There exist elegant formalisms to deal with light of arbitrary polarisation, such as the Jones calculus [Jones41] that employs two-dimensional complex vectors that contain all amplitude and phase information.

A dual basis according to Huygens' principle are spherical waves: every point on a propagating wavefront is itself a source of an outgoing spherical wave, and all such elementary waves superimposed form the shape of the actual wavefront. We refer the reader to literature about Fourier optics [Goodman05, Hecht01] to learn more about this dualism.



 n_1 n_2 α_2

Figure 2.2: A planar wave is diffracted by passing through a slit. According to Huygens' principle, each point on the slit cross-section emits a spherical wave. Their superposition results in a wave that is non-planar and extends into the region behind the occluder that is geometrically shadowed.

Figure 2.3: A beam of monochromatic light with wavelength λ_1 impinges on a boundary between media of different refractive indices $n_{1,2}$ under an angle α_1 with respect to the normal. The spherical waves emitted according to Huygens' principle (shown in red) form a new wavefront. Since the temporal frequency ν remains constant throughout, the wavelengths on either side are $\lambda_{1,2} = 2\pi c_{1,2}/\nu$, and relate as $\lambda_2/\lambda_1 = n_1/n_2$. Simple trigonometric relations yield $\sin(\alpha_1)/\sin(\alpha_2) = \lambda_1/\lambda_2$, from which immediately follows Eq. 2.1.

Diffraction. The wave-optical counterpart to a light ray is a planar wave of infinite lateral extent. This assumption, however, is intrinsically violated in any optical system. A real, finite, optical beam can only be expressed as a superposition of planar waves ("angular



spectrum") as given by a Fourier decomposition. In other words, confining the extent of a light beam to a finite measure will produce a superposition of wave components of varying directionality.

This effect can be observed in the form of *diffraction*, whenever a light beam traverses a narrow slit or a different structure that is not orders of magnitude larger than the wavelength (Figure 2.2). The light wave will then protrude into regions that are geometrically shadowed.

Refraction. Huygens' principle holds, in fact, through boundaries between media of different phase velocity c_1 and c_2 , and provides an intuitive explanation as to why light is refracted at these interfaces. With the refractive index n of a medium defined as the ratio between the speeds of light in vacuum and the medium, $n = c_{\text{medium}}/c_{\text{vacuum}}$, Snell's Law (Eq. 2.1) can be derived as in Figure 2.3. The direction of the refracted ray is perpendicular to the refracted wavefront which, in turn, is formed by elementary waves emanating from the interface between the media.

Coherence/Interference. In practice, it can often be assumed that the superimposing waves are *incoherent*, i.e. their relative phase varies randomly over time and space. This is fulfilled for most spectrally broad light sources. In this case, for the superposition of two beams of light, the resulting intensity will be the sum of the individual intensities.



Figure 2.4: Two waves of identical frequency and amplitude superimpose to form a standing wave. The wave travelling to the right (blue) and the one travelling to the left (red) add up to the stationary wave drawn in black. The vertical lines mark nodes (short dashes) and antinodes (long dashes), the places where the resulting amplitude is extremal.

Narrowband light sources exhibit a strong *coherence* in that the relative phase of the electromagnetic wave is strongly correlated over large regions in time and space. These waves can interfere in a way similar to other wave-like phenomena. In particular, waves can either add up constructively, or cancel each other out. Superimposing light beams of opposite direction can generate so-called standing waves with a fixed spatial amplitude pattern (Figure 2.4). Ultraprecise measurement techniques often make use of interference phenomena in order to achieve subwavelength resolution. Interference of coherent light is also the cause for the typical speckle effect known from many laser light sources, which,



depending on the application, can be difficult to deal with. In our setup used in Chapter 3, we circumvent the problem of speckle formation by the use of a fluorescent medium which emits incoherent broadband light.

Dispersion. In all of the above, we assumed that either the light was monochromatic, or the material constants ε_r and μ_r do not depend on the frequency of the light. In practice, however, most optical media exhibit *dispersion*, i.e. the speed of light will vary with the wavelength. Dealing with dispersion plays an important role in the design of optical systems. We encountered it when trying to match the refractive index of a clear solid by immersing it in a liquid (Figure 3.20).

2.1.3 Quantum Optics

Whenever light interacts with charged matter, for instance in semiconductor-based light sources and photodetectors of almost any kind, both the geometric and wave models fail to describe its behaviour. The reason is that light is not a continuous flow of energy but quantised as per Einstein's findings, whose explanation of the photoelectric effect [Einstein05] was later awarded a Nobel prize. He described light as a flux of particles that each carry a specific amount of energy, Q. While the radiant flux (to be introduced in Section 2.2.1) controls the amount of these particles called *photons*, the per-photon energy solely depends on the frequency v, i.e. the wavelength λ :



$$Q = \hbar v = \frac{hc}{\lambda} \tag{2.6}$$

where $\hbar = h/(2\pi)$ is the reduced Planck constant.



Figure 2.5: Left: Photons can be absorbed by lifting a bound electron to a higher energetic level. Likewise, when an electron drops to a lower level, a photon can be emitted. Middle: Depending on the electronic structure of an atom or molecule and how it is embedded in its environment, the relaxation from a higher to a lower energy state can occur without emission of a photon, for instance by generating heat. Right: Fluorescence is a combination of both effects where part of the energy is dissipated and the rest emitted in a photon of longer wavelength.

The implications are manifold: Under low-light conditions light is a statistical process that follows Poisson's law of small numbers [Poisson37]. Whenever N photons are counted by a detector, the uncertainty of the measurement (*photon shot noise*) is \sqrt{N} .

Furthermore, the energy carried by a photon can interact with the electron hull of an atom. Electrons can absorb incoming photons and be lifted to a higher energetic level. As an electron falls to a state of lower energy, it can emit a photon that carries exactly the difference amount in energy. This makes light a powerful tool for the manipulation and understanding of electronic systems. The phenomenon of fluorescence that plays a major role in this text is a sequence of a quantum-optical absorption and emission event (Figure 2.5).

2.1.4 So Which One Is It?

All of them. Depending on the situation, each of the three models can be useful to describe and utilise light:

- In the macroscopic limit, the most convenient ray model often covers all relevant effects.
- On small scales, or in the coherent case (e.g. when using a narrowband laser source), the wave-like nature of light comes into play.
- When light interacts with bound electrons, or for very low intensities, it needs to be treated as a flux of particles in the quantum mechanical framework.

We will conclude this section by showing how all three paradigms play together in one of the most important discoveries of the 20th century, the laser [Schawlow58]. In this work, we use lasers primarily as source of well-defined rays (Chapters 3 and 5). The key components of all laser systems are shown in Figure 2.6.



Figure 2.6: Laser principle: an amplifying medium is placed in an optical resonator consisting of two mirrors. Inside the resonator, the light forms a standing wave.

The core of any laser is a medium where an effect called *stimulated emission* leads to coherent amplification of light passing through. In many laser media, the underlying quantum-optical mechanism is closely related to fluorescence.

The amplifying medium is placed in an optical resonator that reflects the light wave into itself, generating a standing wave (see also Figure 2.4). Of all the light in the medium, only those waves will resonate that are aligned around the optical axis and whose frequency matches a longitudinal mode of the resonator, i.e. the round-trip distance must be a multiple of the wavelength.

One of the laser mirrors is not perfectly reflecting and acts as an output coupler for part of the light. The higher the resonator quality (i.e. the lower the losses on a round-trip between the mirrors), the better the definition of the outgoing wave will be geometrically as well as spectrally (coherence). Depending on the design, the output beam can be of such low divergence that we can treat it like an almost perfect geometric ray.

The example of the laser demonstrates that each of the light models may play a role in different situations, and that sometimes it is not possible, or useful, to make a clear distinction. In the example scenarios presented in this thesis, we are lucky to get away with the most convenient ray model. Specifics such as fluorescence will be treated mostly phenomenologically, without the need for a full quantum optical formalism.



2.2 Measuring Light

No matter which of the above models may be most useful in a given application, being able to quantify amounts of light is always crucial. Light propagation involves a transport of energy, which can be set in relation to various geometric quantities. In the following, we will review the most important terms that will be used throughout this work.

2.2.1 Radiometric Quantities

Radiant Energy *Q*: The total energy of light is denoted by *Q* and measured in *Joule* [J].

Radiant Flux or Power Φ : The amount of energy per unit time or radiant power is measured in *Watt* (1 W = 1 J/s).

 $\Phi \coloneqq \mathrm{d}Q/\mathrm{d}t$

Radiant Intensity *I*: The radiant intensity relates the radiant flux to a solid angle ω and is measured in W/sr.

$$I \coloneqq \mathrm{d}\Phi/\mathrm{d}\omega$$

Radiant Flux Area Density *u*: This quantity expresses the differential flux per differential area *A* and is measured in W/m^2 .

$$u \coloneqq \mathrm{d}\Phi/\mathrm{d}A$$

The *Irradiance E*(incident flux per area) and *Radiant Emittance M*, often also called *Ra-diosity B*(exitant flux per area), are both instances of the area density of radiant flux.

In quantum optics, it is often convenient to express the much smaller energies in electronvolts (1 eV \approx 1.602 × 10⁻¹⁹ J). Typical photon energies are between 1.8 eV (700 nm) and 3.1 eV (400 nm).

A monochromatic light source ($\lambda = 500 \text{ nm}; Q_{\text{photon}} \approx 2 \text{ eV}$) with a power of 1 W emits about 3×10^{18} photons per second.

The point light source is a popular model in computer graphics. Isotropic point lights are of constant intensity: their radiant flux Φ is distributed evenly over the full spherical solid angle. $I = \Phi/(4\pi sr)$

For a point light source, the area density of radiant flux, unlike the radiant intensity, is not constant but decreases with the distance d from the source. $u = \Phi/(4\pi d^2)$.

Directional lights, another light source model frequently used in graphics, have a constant area density of radiant flux.





Figure 2.7: Geometry around the definition of radiance.

Radiance *L*: Radiance is the radiant flux per unit projected area per unit solid angle. If θ is the angle of incidence or exitance with respect to the surface normal, *L* is defined as

$$L := \frac{\mathrm{d}^2 \Phi}{\cos \theta \cdot \mathrm{d} A \cdot \mathrm{d} \omega}$$

The corresponding unit is $1W/(m^2 \cdot sr)$. In empty space, the radiance is constant along a ray.

The radiance expresses the amount of light per time that passes through a small area from a certain direction. This quantity is of utmost importance in graphics and sensing, since it corresponds to what is seen by an observer, whether in the form of the human eye or other type of sensor. The constancy of radiance is the reason why objects (in vacuum) appear equally bright from all distances.

2.2.2 Spectrometry

All of the above quantities relate to the total amount of energy in the entire electromagnetic spectrum. The discipline of spectrometry investigates how much energy, radiance, etc. is present in a certain spectral part. Each radiometric quantity can be resolved spectrally. For instance, the relation between the integral radiant emittance M and the spectral radiant emittance $M_{\text{spectral}}(\lambda)$ is

$$M = \int_{\Lambda} M_{\rm spectral}(\lambda) d\lambda \tag{2.7}$$

Since the spectral versions of the radiometric quantities are derivatives with respect to the wavelength, their units comprise an additional reciprocal length, e.g. $[M_{\text{spectral}}] = 1 \text{ W}/(\text{m}^2 \cdot \text{nm}).$

2.2.3 Photometry

Photometry is the field of research that investigates how bright something appears *to the eye*, i.e., light is measured taking into account the response of the human visual system. Every radiometric quantity has its photometric counterpart, for instance *radiant emittance* M



 \leftrightarrow *luminant emittance* M_{ν} . A photometric unit is obtained by weighting the corresponding spectral radiometric quantity with a standardised [CIE32] dimensionless spectral sensitivity curve or luminosity function $V(\lambda)$, then integrating over the full spectral range Λ , e.g.:

$$M_{\nu} = 683.002 \frac{\text{lm}}{\text{W}} \int_{\Lambda} M_{\text{spectral}}(\lambda) V(\lambda) d\lambda$$
(2.8)

Note that this definition assumes a "standard observer" under bright conditions (photopic vision). Under low light, the human eye has a different (scotopic) response.

2.3 Interaction between Light and Matter

Even the most sophisticated simulation techniques for light transport can only produce physically correct results when accurate material models are used. In the following, we will review the most important mechanisms of light-matter interaction, and how they can be abstracted as a bidirectional reflectance distribution function (BRDF).

2.3.1 Key Mechanisms

Metals and Dielectrics. Maxwell's equations feature a pair of material constants, the electric permittivity (dielectric constant) ε_r and the magnetic permeability μ_r . While μ_r is approximately 1 for almost any non-magnetic material, ε_r can vary over a wide range, and influences the way a surface reflects and transmits light.

Material	Relative <i>static</i> permittivity (< 1 kHz)		
	$\operatorname{Re}(\varepsilon_r)$	$\operatorname{Im}(\varepsilon_r)$	
Air	1.0006	0	
Glass	4–10	0	
Aluminium	-1300	$1.3 imes 10^{14}$	

On a planar interface between dielectrics, part of an incoming wave is reflected and the other part transmitted. While Snell's Law describes the angular relation, it provides no insight about the relative amounts of light. By imposing continuity constraints for the magnetic and electric fields on Maxwell's equations, Fresnel's equations (Eq. 2.9) are obtained. Here we present them in their squared form which relates to the relative *energy flux* that are reflected for light of s-type and p-type polarisation¹. The reflection coefficients are given as

$$R_{s} = \left(\frac{n_{1}\cos\theta_{i} - n_{2}\cos\theta_{t}}{n_{1}\cos\theta_{i} + n_{2}\cos\theta_{t}}\right)^{2} \quad \text{and} \quad R_{p} = \left(\frac{n_{1}\cos\theta_{t} - n_{2}\cos\theta_{i}}{n_{1}\cos\theta_{t} + n_{2}\cos\theta_{i}}\right)^{2} \tag{2.9}$$

¹With regard to a surface of normal **n**, light is called s-polarised when $\hat{\mathbf{E}} \parallel \mathbf{k} \times \mathbf{n}$, and p-polarised when $\hat{\mathbf{E}} \perp \mathbf{k}$ and $\hat{\mathbf{E}} \perp \mathbf{k} \times \mathbf{n}$. $\hat{\mathbf{E}}$ is the direction of oscillation of the electric field and **k** the wave vector.



Figure 2.8: Fresnel reflection coefficients for s- and p-polarised light and transition from refractive index n_1 to n_2 .

The Fresnel equations are the reason behind virtually all things shiny and specular. Metals have a negative permittivity and therefore form a special group of materials. They are not penetrated by light but reflect it regardless of the angle. Absorption in a material (including coloured metals) can be expressed as an imaginary component of a complex-valued permittivity, causing an exponential decay as per Eq. 2.5.

Scattering. Scattering is an interaction of light and matter that is governed by small structures. Consider the two cases illustrated in Figure 2.9, one example being a volume that contains small reflecting particles, the other one a rough reflecting surface. The "close-up" drawings to the left in this figure show that the direction into which an impinging light ray is redirected is very sensitive with regard to the exact position where the ray interacts with the surface or medium; any minor change in the ray geometry can result in a completely different reflection direction. From a macroscopic point of view, the path of each photon is to some extent random. The ensemble of all reflected light can be characterised quantitatively in terms of a probability distribution, the medium's so-called *phase function*.

Absorption and Emission. In classical scattering media, the underlying mechanism is dielectric reflection and refraction on a microscopic scale. In addition, light may also interact with the charge carriers in a medium, be absorbed and emitted. Often both processes take place in combination, for instance in the case of fluorescence which is an absorption event followed by emission. While the mechanisms behind absorption and emission can only be understood and described in terms of quantum physics (see Section 2.1.3), the resulting optical behaviour can often simply be considered a part of the overall macroscopic reflectance. Wavelength-shifting processes such as fluorescence require a bispectrally resolved treatment (Chapter 4).

2.3.2 Abstraction: Bidirectional Distribution Functions

The monochromatic reflectance of a point \vec{x} on an opaque surface, i.e., the scattering at a reflective interface with microgeometry, can be completely described by the *bidirectional*





Figure 2.9: Scattering on a rough surface (top row) and particles in a volume (bottom row). The presence of microscopic features (left) leads to macroscopic reflectance and transmittance distributions or phase functions (middle). The photos on the right show example cases visualised in a participating medium [Hullin08a]



Figure 2.10: From left to right: reflection lobes of mirroring, diffuse, and glossy reflectance.

reflectance distribution function (BRDF):

$$f_{r}^{(\vec{x})}(\hat{\omega}_{i},\hat{\omega}_{o}) = \frac{dL_{o}^{(\vec{x})}(\hat{\omega}_{o})}{dE^{(\vec{x})}(\hat{\omega}_{i})} = \frac{dL_{o}^{(\vec{x})}(\hat{\omega}_{o})}{L^{(\vec{x})}(\hat{\omega}_{i})\cos\theta d\hat{\omega}_{i}}$$
(2.10)

For homogeneous materials, the BRDF is a four-dimensional function, $f_r(\hat{\omega}_i, \hat{\omega}_o)$. Intuitively, slices of the BRDF can be visualised as a polar diagram for one incoming ray direction (Figure 2.10), but in fact the BRDF describes the reflectance for all incoming and outgoing ray directions. If the reflectance varies over a 2-dimensional surface, two additional spatial parameters can be added to obtain a 6D function: $f_r(\vec{x}, \hat{\omega}_i, \hat{\omega}_o)$.

The ray direction $\hat{\omega}$ can be parametrised using spherical coordinates θ (elevation) and φ (azimuth). A large class of materials reflects light *isotropically*, i.e. the BRDF does not depend on the absolute azimuthal angles of either ray, only the relative angle $\varphi_{rel} = \varphi_o - \varphi_i$. The BRDF for these materials can therefore be reduced to a 3D (homogeneous) or 5D (spatially varying) function $f_r^{(\vec{x})}(\theta_i, \theta_o, \varphi_{rel})$.

The field of research that deals with the measurement and characterisation of reflectance



CHAPTER 2. BACKGROUND

distributions is called *gonioreflectometry*.

For a thin scattering material as illustrated in Figure 2.9, the *bidirectional transmittance* distribution function (BTDF) $f_t^{(\vec{x})}(\hat{\omega}_i, \hat{\omega}_o)$ is defined analogously. While the BRDF covers the hemisphere above the illuminated surface, Ω^+ , the BTDF takes care of the lower hemisphere Ω^- . Combined they form the *bidirectional scattering distribution function* (BSDF).

These definitions assume that all scattering occurs in a localised manner, i.e. all light leaves the surface at the exact point of incidence. In practice, this is often not the case, and light is scattered underneath the surface to a different location. The eight-dimensional function covering general non-local light transport is called the *bidirectional subsurface scattering (reflectance) distribution function* (BSSRDF).

2.4 Computational Light Transport



Figure 2.11: Illumination of a diffuse scene. From left to right: direct light, one indirect bounce added, second indirect bounce added. The steady-state limit of this progression is called "global illumination".

We use the term *light transport* to describe the propagation of radiant energy through the world and its interaction with obstacles along its path. As radiant energy hits an object, part of it is reflected, scattered, or transmitted in a specific way.

Figure 2.11 shows a 2-dimensional scene consisting of a round shape surrounded by a box. Through a hole, light enters the box and hits the round object. Assuming that all surfaces are of diffuse reflectance (see also Section 2.3.2), the round object will reflect the incoming light into all directions and therefore indirectly illuminate the inner walls of the box. As the light bounces on, the entire inner of the box will gradually be illuminated. The resulting state of radiant equilibrium is called *global illumination*. In most real-world scenarios, only the steady state will be observed since the propagation of light is much faster (speed of light) than any changes in illumination.

Simulating global illumination is a crucial step in obtaining photorealistic computer graphics.



2.4.1 Solving Maxwell's Equations

In principle, Maxwell's equations describe everything that is needed to predict global illumination on a physical level. On small scales, where wave effects are predominant, an exact solution of Maxwell's equations may be required. In most (macroscopic) cases, however, this level of description is much too detailed to be of practical use, and a treatment based on geometrical optics is more convenient.

2.4.2 Rendering Equations

[Kajiya86] introduced the rendering equation that expresses all light transport in a scene. It draws its inspiration from the theory of radiative heat transfer. The radiance *L* leaving a point \vec{x} in the direction $\hat{\omega}_o$ is composed of the radiance emitted by the point itself, L_e , and the radiance reflected from other sources at the same point into the same direction, L_r :



$$L(\vec{x}, \hat{\omega}_{o}) = L_{e}(\vec{x}, \hat{\omega}_{o}) + L_{r}(\vec{x}, \hat{\omega}_{o})$$
(2.11)

For opaque surfaces, L_r is obtained by weighting the incident radiance L_i with the BRDF of the surface and integrating over the hemisphere of incident angles, Ω^+ :

$$L_r = \int_{\Omega^+} L_i(\vec{x}, \hat{\omega}_i) f_r(\vec{x}, \hat{\omega}_i, \hat{\omega}_o) \cos\theta d\hat{\omega}_i$$
(2.12)

Hence, for points on opaque surfaces Eq. 2.11 becomes

$$L(\vec{x}, \hat{\omega}_o) = L_e(\vec{x}, \hat{\omega}_o) + \int_{\Omega^+} L_i(\vec{x}, \hat{\omega}_i) f_r(\vec{x}, \hat{\omega}_i, \hat{\omega}_o) \cos\theta d\hat{\omega}_i$$
(2.13)

The surfaces to be rendered are often represented in terms of differential patches dA that are centered at position \vec{P} , \vec{Q} , \vec{R} , ... with the normal pointing in the direction \hat{n}_P , \hat{n}_Q , \hat{n}_R , Light sources are represented as surfaces with non-zero emission term L_e . Any incident radiance at point \vec{P} then corresponds to radiance exiting at a different point \vec{Q} in the direction of \vec{P} , and the integral $\int_{\Omega^+} \dots d\hat{\omega}_i$ turns into an integral over all surfaces:

$$L(\vec{P} \to \vec{R}) = L_e(\vec{P} \to \vec{R}) + \int_S L(\vec{Q} \to \vec{P}) f_r(\vec{Q} \to \vec{P} \to \vec{R}, \hat{n}_P) V(\vec{P} \leftrightarrow \vec{Q}) dA$$
(2.14)

where $V(\vec{P} \leftrightarrow \vec{Q})$ is a geometry term that comprises the mutually subtended projected differential solid angles and the mutual visibility:

$$V(\vec{P}\leftrightarrow\vec{Q}) = \frac{\langle \hat{n}_P, \widehat{PQ} \rangle \langle \hat{n}_Q, \widehat{QP} \rangle}{\|\widehat{PQ}\|^2} \cdot \begin{cases} 1 & \text{if } \vec{P} \text{ visible to } \vec{Q} \\ 0 & \text{if } \vec{P} \text{ not visible to } \vec{Q} \end{cases}$$
(2.15)

For rendering tasks such as participating volumes [Kajiya84], specialised forms of the rendering equation have been developed.

Solution

The major challenge in solving Eq. 2.13 and Eq. 2.14 is that they are recursive and therefore cannot be evaluated without already knowing the solution. However, the equations can be rewritten such that an iterative solution becomes possible. In operator notation, the rendering equation reads as

$$L = L_e + TL \tag{2.16}$$

where L contains the outgoing radiance for all surface points and all directions, and L_e is the emissive radiance for all surface points. T is a light transport operator that, in the case of opaque surfaces, comprises all surface reflectance and free-space light propagation, including all shadowing terms and view factors for any pair of surface points.



We move *TL* to the left hand side, and expand the inverse of the operator (1 - T), 1 being the identity operator, into a Neumann series.

$$L - TL = (1 - T)L = L_e$$

$$L = (1 - T)^{-1}L_e$$

$$L = (1 + T + T^2 + T^3 + \cdots)L_e$$

$$= L_e + TL_e + T^2L_e + T^3L_e + \cdots$$

Starting with the emissive radiance, global illumination can now be obtained by applying the light transport operator once, twice, and so forth. In this notation, T^i stands for the i^{th} bounce (see also Figure 2.11). Convergence for the series is guaranteed if the spectral radius of T is less than 1, i.e., if the reflection operator satisfies the conservation of energy.

In this form, the rendering equation can be solved using many techniques. Most of them belong to either the finite element or random walk (Monte Carlo) families [Dutre02]. In practice, satisfactory results are often obtained by computing only the first few terms of the series.

Spectral and Bispectral Rendering

For all of the above equations, spectral formulations can be obtained by using spectral instead of monochromatic radiances. For bispectral rendering, which is required for the reproduction of fluorescence, we provide the rendering equation in Chapter 4.3.



CHAPTER 2. BACKGROUND



Chapter 3

Fluorescent Immersion Range Scanning



Figure 3.1: From left to right: Photo of a glass goblet, light sheet in fluorescent immersion intersecting the object, reconstruction from 6 scans under different orientations.

The first ingredient to the rendering of any three-dimensional scene is knowledge about its geometry. There exists a large number of techniques to digitise the objects around us, most prominently laser line 3D scanners based on the triangulation principle. Traditionally, these systems have assumed the measured objects to be of homogeneous opaque and diffuse (Lambertian) reflectance. Most real materials, however, exhibit more complex, possibly non-local, light transport and therefore violate this assumption, resulting in corrupted measurements. By placing the object in a participating medium, we turn a traditional laser line scanner into a light sheet scanner. This allows us to robustly measure the free space in front of the surface and acquire robust 3D scans, regardless of the material reflectance. For homogeneous clear materials, entire slices of the geometry can be acquired at once by matching the refractive index of the participating medium to the one of the object.



3.1 3D Scanning of Optically Challenging Materials

The acquisition of 3-dimensional surfaces is a problem which occurs frequently in mechanical and optical engineering, computer vision and many other application fields. Today's most precise methods involve active illumination by means of laser beams, lines or more sophisticated projection patterns. Usually, the light reflected or scattered by the surface is captured by an off-axis camera, so that the depth information can be recovered by triangulation.

The usability of most of these methods is limited by the object material, which defines the reflectance properties. An ideal surface for this class of 3D scanners scatters an incoming ray of light diffusely into all directions so that each impinging light ray results in a well-defined hit point visible from any viewing direction. However, many objects and materials exhibit a reflectance that is highly uncooperative with regard to range scanning. Complex effects such as transparency, subsurface light transport, specular reflectivity, and interreflections between surface points pose a challenge for the design of a general and robust system for shape acquisition.

Some of these effects have been addressed by modified scanning systems, more robust analysis of the scanning pattern [Curless95, Nayar06], or e.g. by exploiting other physical properties of light such as polarisation [Chen07].

In this chapter, we take a different approach: Instead of analysing the intensity of the reflection, we observe light rays as they propagate through space *before* they hit the surface. The intersection point with the surface is located where the ray is interrupted, see Figure 3.2. This holds for all kinds of surfaces except perfect mirrors.

In order to directly observe light rays, we have to embed the object into a participating medium. In fact, sparse sets of light rays have already been used as probes for optical density measurements, e.g. of smoke [Hawkins05, Fuchs07]. However, media such as smoke or dilute milk exhibit multiple scattering, resulting in a visible halo around the actual ray and a significant decrease in perceived contrast. The reason is that scattering events in these media are *elastic*, i.e. upon scattering of a photon its wavelength is conserved. Therefore, photons can be scattered any number of times before they leave the volume or are absorbed.

Fluorescent media also scatter incoming light, but they interact *inelastically* with photons, absorbing them in a particular wavelength range and emitting photons of larger wavelength. Compared to elastic scattering, fluorescence has a number of advantages:

- The phase function is almost perfectly isotropic.
- Multiple scattering is negligible due to a small overlap of the absorption and emission spectra.
- The spectral separation of fluorescence and directly reflected light is simple using an optical long-pass filter.




Figure 3.2: (a) Traditional laser range scanning: A sheet of light is projected onto the object surface and its reflection is observed by a camera (not shown), (b) Fluorescent immersion range scanning: A sheet of light is projected onto an object immersed in a participating medium. The light sheet becomes visible and its cut-off is observed by a camera (not shown), enabling triangulation of surface points.

Based on this principle of observing the attenuation along a ray in a fluorescent fluid rather than detecting the direct reflection of light from the surface, we present a 3D range scanning system that can detect the surface of objects composed of a wide range of materials. Our system enables the acquisition of dark, glossy, subsurface-scattering or even fully transparent objects.

Exploring the same principle, we propose a second setup that can perform direct volume slicing for transparent objects with a constant index of refraction (Section 3.6). After matching the refractive index of the fluid and the object, the light rays in the volume form straight lines. This matching has previously been used in optical tomography approaches [Sharpe02, Trifonov06]. In contrast to tomographic reconstruction, we directly capture individual silhouette slices through the object as the laser sheet is only visible outside the object, in regions where it interacts with the fluorescent fluid. In Figure 3.19, we show the captured volumes of two intricate glass objects. The input data captured by our setup can be seen as a flick-book style animation on odd pages, the reconstructed 3D rendering on even pages.

3.2 Related Work

In this chapter we investigate 3D range scanning of objects with non-cooperative materials, such as refractive surfaces, objects made of materials exhibiting a significant specular component, light absorbing materials such as diffusely reflecting, dark surfaces and mixtures thereof. In the following, we will give an overview of current state of the art, especially with respect to those notoriously challenging cases.





Figure 3.3: By immersing objects into a fluorescent liquid we have acquired high quality range scans for a set of surfaces that are difficult to acquire with traditional methods.



3.2.1 Range Scanning: Acquisition of Surface Geometry

3D range scanning has been investigated for several decades, the majority of approaches assuming a diffuse reflectance of the object's surface. The literature on the subject is most commonly divided into active and passive techniques. Active light techniques include laser range scanning, coded structured light systems and time-of-flight scanners. Overviews of the state of the art in active light range scanners can be found in [Chen00], or more recently in [Blais04]. The other broad class of object acquisition techniques uses passive sensing. A recent survey of image based 3D acquisition, covering techniques such as structure from motion, stereo, multi-view stereo and shape from shading, is given in [Remondin06].

3.2.2 The Struggle with Non-Lambertian Materials

The further a surface deviates from the Lambertian reflectance assumption, the less accurate standard 3D range scanning techniques become [Curless95, Beraldin04]. [Curless95] show that varying reflectance of the scanned surface, among other factors, results in systematic depth errors. They propose to analyse the raw laser range data in tilted space-time images of single scanlines to stabilise the detection process. Techniques that exploit both spatial and temporal analysis have also been proposed for stereo applications [Davis05a].

While coating objects with paint or removable powder is an option, clearly there are situations where this is not desirable. Several researchers have adressed the shortcomings of traditional range scanning techniques, extending the class of objects that can be successfully acquired by active light techniques. However, most proposed techniques only covered one out of the challenging material classes. They will be discussed in the following.

Specular Materials

Particular effort has been spent on materials with specular gloss and shiny highlights. In the field of laser-based setups, attempts have been made to explain and exploit multiple brightness peaks caused by specular highlights [Park04, Park08], or to mask them out by optical modifications to the setup[Baba01]. Others have used polarisation to remove reflections [Trucco94], or employed multi-camera setups in order to disambiguate range measurements of specular surfaces [Clark97].

Other techniques involving active, or at least, well-known illumination, have also been presented. Some use additional assumptions about the shape of the surface [Halstead96]. [Bonfort03, Bonfort06] placed specular objects in known scenes and reconstructed voxel models or point clouds with normal information. Similarly, [Tarini05] surrounded mirroring objects in a setup involving stripe pattern displays, and observed the distortion of the patterns in order to generate depth and normal maps.

Passive range sensing techniques have also been extended to account for non-lambertian materials. For instance, [Graves07, Jin05, Davis05b, Zickler02, Bhat95, Schultz94] pro-



posed stereo techniques tailored to specular objects. However, the accuracy of the acquired object geometry varies widely and is in general lower than that achievable with active light techniques.

Sub-Surface Scattering and Transparent Materials

Many materials exhibit light transport that is non-local: light impinging on one point may be scattered underneath the surface and leave the object in a different point. Those materials are a challenge to most 3D scanning systems and have therefore received some attention as well. Again, results can be improved when the influence of the sub-surface scattering effect on the data is reduced. [Chen07] took into account the polarisation of the reflected light and were able to reliably remove specular reflections and subsurface scattering effects, respectively.

Scattering and transparency play a central role in this chapter in that we place the object in a medium that is both scattering (fluorescent) and transparent for different respective parts of the visible spectrum.

The challenges imposed by photogrammetry in clear media are mostly of geometric nature. Solutions to these problems date back as far as 1971, when Höhle et al. [Höhle71] posed the reconstruction of under-water objects as a bundle adjustment problem. They assumed that the object of interest would be surrounded by a medium of well-known refractivity that is bounded by a single refracting surface. The requirement of a known refractive index can be relaxed when the interface is planar [Maas95]. [Treibitz08] provided an indepth treatment of the geometric relations when looking through flat refractive interfaces, which leads to projections that do not have a single centre. Sometimes the interest is in the shape of the transparent surface itself, for which an approach was exposed by [Murase92]. The author analyses a distorted pattern at the bottom of a water tank to estimate the shape of a water surface. [Morris05] use the same principle to deal with dynamic water surfaces exhibiting an unknown index of refraction. The reconstruction of single refracting surfaces is a special case of multi-interface reconstruction which has been theoretically analysed by [Kutulakos08]. The method is based on ray measurements using an LCD screen mounted onto a precision translation stage.

A different technique for reconstructing refractive objects or fluids is based on inverse ray-tracing. Assuming an object's surface at a certain position, an image is generated using a forward model of light transport. The surface is then optimised such that the synthetically generated image matches the photograph as closely as possible. One example of this technique is introduced by [Ihrke05] who reconstruct free-flowing water surfaces using measurements of the optical path length inside the water column. Similarly, [Finckh10] optimise refractive geometry to generate a desired caustic shape. Here, the objective is not in the exact reconstruction of geometry but in faithful reconstruction of the caustics. Another related approach uses polarised light in conjunction with polarisation ray-tracing [Miyazaki05]. A different polarisation-based technique is described by [Saito99].



The authors measure the polarisation state of surface highlights and compute surface shape from this information.

[Ben-Ezra03] use a model-based approach to reconstruct simple shapes such as spheres, cylinders and lens-shapes based on the apparent movement of refracted feature points. [Morris07] describe an active light technique called scatter trace photography that can be used to acquire the geometry of objects with difficult material properties. The method is based on analysis of the materials' reflectance properties using a moving light source.

The acquisition of object geometry in the presence of participating media, as presented by [Narasimhan05], is another field of research that shares certain aspects with our work. This problem is most prominent in the fields of aerial and under-water imaging, where the scattering environment leads to a degradation in image quality. However, rather than causing objectionable effects, scattering media can also be used as a tool for vision and measurement purposes. For instance, methods for acquiring 3D density distributions within heterogeneous participating or fluorescent media have been presented before [Hawkins05, Fuchs07, Deusch01]. In our case, we employ a homogeneous medium in order to ease the detection of the scanned object's surface.

3.2.3 Tomography

Finally, 3D shapes can be acquired by means of computerised tomography, e.g. using X-rays that do not underlie refraction [Kanitsar02]. [Trifonov06] transferred the CT methodology to optical wavelengths by immersing transparent objects in a liquid of matching refractive index. Our volume scanning approach is inspired by this work, however, we acquire the desired volume data directly instead of sampling projections of it. Therefore, our technique does not require tomographic reconstruction, or an inverse formulation of any kind.

3.3 Making Light Rays Visible

Our light sheet range scanning approach is based on the visualisation of light rays propagating through space in front of the object surface. We achieve this visualisation by employing a participating (scattering) medium that, at every point in the volume, diverts a small portion of the incident illumination towards the camera. In Chapter 2.3, we have seen that scattering can often be described in terms of interaction with very small geometry: smoke particles, fog droplets or fat globules in dilute milk. For our task of 3D scanning, it turns out that fluorescent media are particularly well suited. In this section we compare the characteristics of dilute milk and fluorescent liquids using a test scene composed of different materials (Figures 3.6 and 3.7).







Figure 3.4: Light sheet ($\lambda = 532 \text{ nm}$) projected into a glass cylinder filled with dilute milk (left) and fluorescent Eosin Y solution (right). Notice the colour shift, and the contrast improvement on the lower edge of the light sheet due to the absence of multiple scattering.

3.3.1 Elastic Scattering in Dilute Milk

Scattering in milk mainly occurs due to the presence of particles such as fat globules. They absorb, reflect and refract incoming photons without altering the wavelength. Light transport in such media is usually described by two parameters, an average phase function per scattering event and the mean free path of a photon in the medium. The single-scattering phase function is intrinsically given by the kind of milk being used (size and shape of particles). The mean free path l_{free} depends on the scattering cross-section σ and on the concentration of scatterers in the water n:

$$l_{\text{free}} = \frac{1}{\sigma n} \tag{3.1}$$

Knowing the mean free path length allows for basic quantitative considerations. In particular, it is directly related with the probability p of a scattering event along a given distance l:

$$p = 1 - e^{-l/l_{\rm free}}$$
(3.2)

For example, assume that a photon's probability of being scattered on its way through the volume to an object is $p_1 \ll 1$. Then, for a photon that has been scattered at a location close to the object, the probability of a second scattering event on its way back is again $p_2 = p$. Hence, for low concentrations, a given photon is much more likely to undergo a single scattering event than multiple scattering events $(p_1 p_2 \dots p_n = p^n \ll p)$. On the other hand, the probability of a photon to reach the object surface without being scattered is $1 - p \approx 1$. This results in surface reflections that are typically much brighter than the volumetric contribution, which undoes the potential advantages of a light sheet based approach. If the milk concentration is chosen high enough for the light sheet to reach a significant brightness $(p_1 \not\ll 1)$, multiple scattering will occur inevitably (please refer to Figure 3.7 for experimental evidence).





3.3.2 Inelastic Fluorescence Scattering

Figure 3.5: Normalised absorption and emission spectra of Eosin Y and transmission curve of the longpass filter. Data obtained from [TU Graz; Thorlabs]

In the ideal case for our application, multiple scattering and contributions from the surface would be non-existing and single scattering would determine the image formation. We can approximate this case by using a solution of a fluorescent dye.

In Chapter 2.1.3, we explained that fluorescence is a process where a photon is absorbed by an electronic system (atom, molecule, ion), which then emits a photon of different wavelength. The change of wavelength always occurs towards long waves (*Stokes shift*), as required by the conservation of energy. This results in a set of very desirable properties: Multiple scattering can almost completely be suppressed by using an excitation wavelength that is close to the long wavelength tail of the absorption spectrum (see Figure 3.5). Due to the loss of energy, a photon that takes part in one scattering event is unlikely to be scattered again, rendering the fluorescent fluid close to transparent for the emitted wavelength. Using an optical longpass filter, we can conveniently separate the fluorescence from surface reflections. Finally, as a side effect, the phase function (see Chapter 2.3) is almost perfectly isotropic, since the excited state of the dye molecule typically has a lifetime in the order of a few nanoseconds, during which the molecule can rotate numerous times before it emits a photon into an uncorrelated direction.

3.3.3 Comparison

In order to compare elastic and inelastic scattering, we set up a test scene composed of an opaque checkerboard pattern, a translucent object and a mirror in a 251 water tank filled with scattering media of varying concentration (Figure 3.6). An interactive version of this experiment (*Virtual Scattering Lab*) is available online [Hullin08c].





Figure 3.6: Test scene for comparison between dilute milk and Eosin solution. The intersection lines between light sheet and surfaces are marked in red.



Figure 3.7: (top) Test scene in dilute milk. For very low concentrations, multiple scattering is rather unlikely. In this case, however, the light sheet is very dim and hardly visible in the presence of the much brighter surface features. (bottom) Scene in Eosin Y solution, captured through longpass filter and normalised for print.



- *Choice of concentration.* The scattering behaviour of the additive is crucial for the image quality. While it is difficult to find a suitable concentration of milk, we can vary the concentration of Eosin Y over a wide range without affecting the image contrast (Figure 3.7, bottom series).
- *Intensity*. For milk, the intensity that reaches the camera decreases with the distance to the scattering location. The fluorescent liquid is virtually transparent for the fluorescence emission.
- Contrast. The same holds for the contrast (see also Figure 3.4).
- *Surface reflections*. Reflections from object surfaces (including subsurface light transport) can become much brighter than the light sheet, see e.g. the dolphin in Figure 3.7 (top). Using a fluorescent medium, a longpass filter removes most of these contributions because of their shorter wavelength.
- *Directionality.* Dilute milk scatters light anisotropically with strong preference on forward and backward scattering. The phase function of fluorescence events is close to isotropic.

3.4 Measurement Setup

In the following we discuss our physical measurement setup, the choice of chemicals involved as well as the calibration of our range scanning system.

3.4.1 General Setup

The setup consists of a computer-controlled laser projector, a CCD camera and a glass tank, see Figure 3.8. The glass tank is filled with a liquid mixed with fluorescent dye in which we immerse the object to be measured. The fluorescent solution renders sheets of light, projected by the laser system, visible. For our experiments, we used different liquids mixed with the chemical Eosin Y, which exhibits fluorescent characteristics. The type of liquid used for surface scanning, Section 3.5, differs from the ones employed in volume scanning, Section 3.6. However, there are some common considerations.

The intensity of emitted light depends on the concentration of the fluorescent agent. A higher concentration increases the scattering coefficient σ_s of the solution, resulting in a higher fluorescent emission intensity. On the other hand, the absorption of light along the laser ray also increases. In order to minimise CCD noise, we find a trade-off that minimises intensity fall-off along the illuminating rays while keeping exposure times at reasonable levels.

As a light source, we use a frequency doubled Nd:YAG laser (P = 60 mW, $\lambda = 532 \text{ nm}$) and a computer controlled XY pair of galvanometer scanners. Our imaging sensor is a





Figure 3.8: Our measurement setup consists of a computer controlled laser projector, a CCD camera and a glass tank containing a fluorescent liquid and the object.

Canon EOS 5D digital SLR camera equipped with a 100 mm f/2.8 macro lens. In order to capture only the fluorescence and filter out any reflections of laser light, we equip the lens with an optical longpass filter (Thorlabs FEL0550) with a cut-off wavelength of $\lambda_c = 550$ nm.

The scanned object is fixed onto a precisely manufactured, manual turn-table positioned in the glass tank. By rotating and scanning the object, we acquire surface scans from different orientations, similar to traditional range scanning approaches. The single scans have to be merged in order to produce a full mesh.

Our two setups mainly differ in the arrangement of the laser projector with respect to the imaging sensor and in the choice of liquid that fills the glass tank.

After capture, the raw data from the camera is processed through the freeware *dcraw* [Coffin10]. Proper interpolation of the interleaved RGB data from the image sensor is a complex issue on its own [Ajdin08]; since we found a quarter of the original image resolution to be sufficient for all purposes, we use dcraw's demosaicking-by-downsampling option (dcraw -h) in order to save time and storage space.

Surface Scan Specifics

For fluorescent immersion surface scanning, Section 3.5, we employ plain water as a basis for the fluorescent solution. The concentration of Eosin Y is approximately 0.5 mg/l. We position the camera such that its optical axis is approximately aligned with the normal of the bounding plane of the glass tank. This way, we minimise image aberrations and achieve a relatively homogenous sampling rate in the image plane. We position the laser at an angle





Figure 3.9: Top view on three glass cylinders intersecting a laser plane within a fluorescent liquid (illustration). Three different cases occur: (1) $n_{\text{glass}} > n_{\text{liquid}}$, (2) $n_{\text{glass}} < n_{\text{liquid}}$, (3) good match ($n_{\text{glass}} \approx n_{\text{liquid}}$). The arrow indicates the incoming light direction.

of approximately $30^{\circ}-45^{\circ}$ to the camera's optical axis. As in traditional laser scanning there is a trade-off between triangulation accuracy and occlusion problems. Additionally, in our case, we want to avoid grazing angles of the laser plane with respect to the front face of the glass tank. Our setup for surface acquisition is shown in Figure 3.8.

Volume Scan Specifics

In the case of fluorescent immersion volume scanning for refractive objects with constant index of refraction, Section 3.6, we match the refractive index of the immersing liquid to that of the object before scanning it. This measure straightens the ray paths and enables us to directly observe volumetric slices of the object.

Refractive index matching liquids are available for a wide range of indices. However, they tend to get more and more poisonous as the refractive index increases. [Trifonov06] suggest the use of potassium thiocyanate to obtain a refractive index of around 1.5, which is sufficient for most light glasses. From a saturated solution (around 80 %) the refractive indices of borosilicate glass (n = 1.48) and quartz glass (n = 1.46) can easily be matched by gradually adding water until the refractive index of liquid and glass is approximately equal (see Figure 3.9).

Note, however, that the refractive index can only be matched perfectly for the single wavelength of the laser. Because of the broadband spectrum of the fluorescence light, dispersion may cause a slight disturbance of light rays on their way from the illuminated plane to the camera through glass structures (see Figure 3.20).

A drawback of potassium thiocyanate is that it weakens the fluorescence of Eosin significantly. We also experimented with glycerol, which, unlike potassium thiocyanate, is non-ionic and therefore induces less chemical stress on the dye molecules, resulting in a higher efficiency (more emission per absorption). However, due to the high viscosity of glycerol, heavy stirring is necessary to blend it with water, which in turn produces numerous small bubbles.

In addition to matching the refractive index of the immersing liquid to the object, we change the laser-camera setup to an approximate 90° constellation. The laser projects



sheets of light through one side of the glass tank, while the camera observes the object slices through the other. This is possible because occlusion is not going to occur for transparent objects immersed in a refractive index matched liquid. The objects are close to invisible to the naked eye due to negligible refraction.

3.4.2 Calibration



Figure 3.10: Two views of an olympic swimming pool. Note how the nonlinearity in Snell's law introduces severe geometric distortions. On the left photo, the refraction causes apparent intersections of lines that are in fact parallel.



Figure 3.11: An example of a laser sheet being projected onto the calibration target. The sheet of light intersects the planar target in a curve. In the background, part of the calibration pattern is visible. The dots have a distance of 10 mm centre to centre.

Since both camera and laser source are looking through a refractive interface, their projections are severely distorted [Treibitz08]. We calibrate our range scanning system using a two-step procedure. First, the viewing rays are calibrated using an image-based approach similar to [Trifonov06].

A precisely manufactured, planar calibration target with known world coordinates is afixed to the turntable and recorded from different orientations. This step results in a vector of 3D world coordinates (x, y, z) per pixel (u, v) of our imaging sensor, defining mappings

$$T_{\alpha}: \mathbb{R}^2 \to \mathbb{R}^3$$



from image space to world space. α denotes the orientation of the calibration target and T_0 and T_{180} define the front and the back of our reconstruction volume, respectively. Transforming the pixel coordinates (u, v) to world coordinates using the transformations T_{α} , we obtain a set of 3D points (x, y, z) for each camera pixel. We compute the camera rays within the glass tank by fitting a three-dimensional line, Eq. 3.5, to these points. To obtain the ray equations inside the glass tank for every image pixel, we fit a line to these 3D point measurements.

A similar image-based technique can be used in the second step to calibrate the laser sheets inside the recording volume [Trucco94]. However, since we project sheets of light into a refractive medium, single rays within the sheet form different incident angles with the medium. This results in a deformation of the sheets, which no longer form planes in space, see Figure 3.11.

In order to intersect camera rays with the curved light sheets and thus triangulate 3D points, we need an algebraic description of the curved laser illumination. To compute this description, we record a sparsely sampled set of laser sheets projected onto the planar calibration target, again in different orientations. The intersection points of the laser sheet with the calibration target in the image plane are computed as described in Section 3.5. We denote the set of sampled intersection points by $(u_{\alpha}^{i}, v_{\alpha}^{i}, t_{\alpha})$ i = 1...N, where N is the number of detected 2D points for laser sheet t_{α} with the calibration target rotated into position α . We then compute the positions of these sample points in world coordinates by applying the appropriate mapping T_{α} :

$$(x, y, z, t) = [T_{\alpha}(u_{\alpha}, v_{\alpha}); t_{\alpha}].$$
(3.3)

For better readability we drop the sample index *i* from the equation. Applying Eq. (3.3) we obtain a number of world points (x, y, z) that were hit by a particular laser sheet *t*. We use these sampled points to interpolate this information in the whole reconstruction volume. We model the behaviour of the light sheets inside the glass tank as a quadratic polynomial in *y* and *t*, while keeping the dependency on *x* linear:

$$z = f(x, y, t) = a_0 y^2 x + a_1 y^2 + a_2 y x + a_3 y + a_4 x + a_5,$$

$$a_i = b_{i0} t^2 + b_{i1} t + b_{i2}.$$
(3.4)

Eq. (3.4) is linear in its 18 coefficients b_{00} to b_{52} . Thus, we obtain one linear equation per sample point (x, y, z, t) and can solve for the polynomial's coefficients in a least squares sense. The choice to keep f linear in its argument x is motivated by the fact that in our setup the laser sheet forms a constant angle with the boundary of the glass tank for this variable. Therefore, the curvature of the laser sheet will not be affected by changes in the x coordinate.

3.4.3 Triangulation

The task of triangulation is to determine a 3D point from pixel positions (u, v) on the image sensor and a laser sheet *t*. The first step of the calibration procedure yields a ray for each



image pixel:

$$\mathbf{r}(u,v) = \mathbf{p}(u,v) + s\mathbf{d}(u,v). \tag{3.5}$$

Rearranging Eq. (3.4) into an implicit form

$$f(x, y, t) - z = 0, (3.6)$$

and combining it with Eq. (3.5) we obtain a cubic polynomial in *s*. Solving it results in three solutions, which we insert into the ray equation, Eq. (3.5). We check the resulting points against our reconstruction volume and discard the ones that fall outside of it.

3.4.4 Data Acquisition

The scanning process involves the sweeping of a vertical laser sheet through the acquisition volume. We start the scanning process with the laser sheet t = 0 closest to the camera and proceed towards the back of the glass tank, i.e. with increasing t the distance between the laser sheet and the camera increases as well. For each projected laser sheet we take an image with the imaging sensor. The output of this system is a stack of images $I_t(u, v)$, which we analyse to determine the intersection points between the object and the laser sheet.

In the following, we describe a robust surface detection algorithm, Section 3.5, and a volume slicing method for refractive objects, Section 3.6.

3.5 3D Surface Scanning

Using the setup described in Section 3.4.1 we first present a system for performing a surface range scan which operates on surfaces that are otherwise challenging for traditional 3D scanners, namely transparent, translucent or dark materials.

Due to the use of fluorescent Eosin Y the laser sheet is visible to the camera as long as it travels through un-occluded space. As soon as the incident light interacts with the surface, most of its energy is absorbed by or transmitted through the object. Since the object itself is supposed not to be fluorescent, we observe a significant fall-off in the measured intensity along the laser sheet exactly at the intersection point with the object's surface as seen in Figure 3.12. Our surface detection will therefore concentrate on detecting a negative change in intensity rather than detecting a peak as in traditional range scanning [Curless95].

3.5.1 Influence of the Material

This drop in intensity can be observed regardless of the surface properties, as long as the material is not perfectly mirroring. It is, however, superposed by a number of additional effects caused by global light transport. Examples of such effects are specular or diffuse



3.5. 3D SURFACE SCANNING



Figure 3.12: Laser sheet hitting various surfaces. The intensity along the laser line drops significantly when interacting with the surface. The intersection line is clearly visible for transparent (Figure 3.1), dark (a), translucent (b), and partially diffuse (c) surfaces. Note that for the first three cases the intensity does not increase right at the intersection line. The die in (c) is transparent but slightly fluorescent itself, resulting in a bright profile.

reflections of the light sheet on the object surface, as well as volume caustics of refracted or reflected light rays (see e.g. Figure 3.1 (middle)). In Section 3.5.2, we propose a method to deal with these effects.

Furthermore, diffuse surfaces of high albedo exhibit a bright line caused by volumetric scattering in front of the surface. This causes a small systematic error, as discussed in Section 3.7.2.

3.5.2 Detection Algorithm

In order to detect the surface, we analyse the captured images $I_t(u, v)$ in a space-time stack, and process them initially per camera scanline. For each scanline v, we obtain a slice $S_v(u, t)$ through the space-time volume as depicted in Figure 3.13.

It is instructive to look at the gradient g with components (g_u, g_v, g_t) of the recorded intensity of one pixel in the space time volume. $g_t(u, t)$ and $g_u(u, t)$ are shown in Figure 3.13(c) and (d). We compute the gradients on a median filtered space-time volume (kernel size $3 \times 3 \times 3$). For any surface we can state the following conditions:

- The component g_u has a negative peak. (We are looking for an intensity drop.)
- The component gt has a negative peak.
 (It has to move if the laser sheet is moved.)
- The 2D gradient (g_u, g_v) must not be perpendicular to the laser direction l
 in the image.
 (*Mistrust features that are aligned with laser rays.*)
- 4. (optional) The detected surface line in the slice must be strictly monotonic.





Figure 3.13: (a) Space-time slice for one scan line through the crystal goblet. (b) Shows the detected surface in red which is nicely aligned with the local minimum of the gradient component g_t (c). At the surface, the gradient component g_u (d) also has a minimum, but at other places it is more corrupted by stationary features.

As can be seen in Figure 3.13(d), the component g_u can be corrupted by caustics and reflections that intersect with the actual edge in image space, resulting in the vertical lines in the gradient image. This is not the case in the g_t component, where the wanted surface appears much cleaner. We therefore determine all local minima of g_t that are below a threshold θ_t with sub-pixel precision, reporting a fractional t value for each camera pixel. Blurring the gradient images before the thresholding allows to trade resolution against smoother results.

In the next step, we discard all candidate points who do not meet the first and third criteria with the corresponding thresholds $g_u \leq \theta_u$ and $\leq (\vec{l}, (g_u, g_v)) \leq \theta_\alpha$. In order to close some occasional gaps, we soften this selection process by also keeping points which have an 8-neighbor that meets the criteria.

Finally, we traverse the slice from bottom to top, and collect all first surface points encountered. We require the surface to be monotonic, which is necessary for transparent objects. Otherwise it might happen that distorted features behind the first surface are detected as well. For nontransparent materials, however, the monotonicity constraint can be dropped, resulting in a higher yield of valid surface points.

After performing this routine for each space-time slice, we obtain a map containing a t value per camera pixel (Figure 3.14). It might still contain a number of patches that correspond to volume caustics due to reflection or refraction. Since those unwanted surfaces are typically small in size, we can eliminate most of them by segmenting the depth map estimate into continuous patches and discarding or filling up segments that contain less than n_{\min} pixels.





Figure 3.14: Scan maps generated from space-time stack. From left to right: depth map t(u, v) before and after segment size filtering, world coordinate map (x, y, z). Colours are normalised for print.

The absolute threshold, denoising and segment size filter parameters have to be found to meet demands of the respective measurement setup (camera resolution, measurement geometry, light sheet brightness, etc.). After downscaling all camera images by 50% (corresponding to approx. 3 megapixels), the following parameters performed well for our setup: $\theta_{\alpha} = 85^{\circ}$, $n_{\min} = 1000$. The optimal choice of the thresholds θ_t and θ_u depends on the observed contrast in the space-time images. We determined them experimentally for our setup and used the same set of parameters for all surface scans in this chapter.

The filtered depth map t(u, v) is converted into a 3D surface mesh using the camera calibration as described in Section 3.4.3.

3.6 Single-scan Direct Volume Capturing

In this section we describe a method for capturing *transparent* objects volumetrically, without the need for tomographic reconstruction. We present a technique that directly records one slice through the volume for every laser sheet.

3.6.1 Matching the Refractive Index

So far, our scanning was limited to the first surface, where the light is being refracted and no longer propagates in a straight line (see Figure 3.9). We change our setup to the one described in Section 3.4.1. By matching the refractive index of the liquid to that of the object the ray direction is preserved even at interfaces between liquid and solid, so that all intersections between a light ray and the object occur along a straight line (cf. [Trifonov06]).





Figure 3.15: Capturing the mouse model: Photo of the object and two slices of the volume stack after matching the refractive index. The contour of the mouse is clearly visible in the surrounding fluorescent liquid.

Using a fluorescent liquid, we can now observe entire slices of clear objects at once when illuminating the tank with a laser sheet. Since the index of refraction is matched, the viewing rays will not bend at the surface. The object overall, and in particular in front or behind the laser plane, therefore is not visible at all. Since the object itself is not fluorescent, only the surrounding liquid lights up while the object remains black (see Figure 3.15). The object's surface is located at the boundary between light and dark image regions. In order to capture the entire volume we sweep the laser plane through the entire volume once.

3.6.2 Volume Processing

As with the surface acquisition, we obtain our raw data with a stereo setup of a camera and a laser projector. This time, however, since occlusion does not play a role for clear objects in a matched environment, we can position camera and projector perpendicularly to each other. Thus, we achieve a maximum stereo baseline and hence the highest possible triangulation accuracy.

Our volume acquisition pipeline is organised as follows:

- 1. Recording of a space-time stack by sweeping a light plane through the volume and taking a sequence of images through the longpass filter
- 2. Binary segmentation
- 3. Filtering in spatial and space-time domain (see Figure 3.16) to reduce noise and artifacts
- 4. Resampling of the filtered stack into world coordinates (x, y, z)
- 5. Extraction of the isosurface using Marching Cubes [Lorensen87]

Given an input stack of images, (for examplary slices see Figure 3.15), we perform a binary segmentation to separate the object from the background. A stable separation is



obtained by using a median filtered reference image in which the laser sheet does not hit the object.

The segmented slices might still contain a set of misclassified pixels which are mostly filtered out in the pipeline illustrated in Figure 3.16. The following effects have to be accounted for:

Noise. Particles in the liquid, bubbles, and other effects lead to the presence of a "salt and pepper"-type noise which forms connected regions in the order of a few pixels. They can be removed by inverting all regions that are smaller than 200 pixels (Figure 3.16 (b)). Alternatively, a median filter could be applied.



Figure 3.16: The volume filtering pipeline: A binary segmentation is performed using a simple thresholding at 50% of a median filtered reference light sheet (a). As a first denoising step, all four-connected image regions smaller than 200 pixels are discarded (b). Shadows cast by inclusions are detected as they penetrate the space to the right of the object (green box) which should be unoccluded (c). After tracing them back towards the light source, the traversed space-time volume segments are filled by 3D diffusion. Remaining horizontal features of insignificant height are removed in the last filtering step (d).

Needle-shaped shadows. Inclusions in the glass, such as bubbles of air or particles, or regions with different refraction index will cast shadows. Similarly, a slight mismatch in the refractive index can cause such shadows, as depicted in Figure 3.9. We propose two different approaches to detect and to remove these kinds of artifacts.

Most of these artifacts extend to the far right side of each volume slice. Any feature that is detected in the region outside the object's bounding box is treated as a shadow region.



55

We project all of these regions onto a shadow map. Then we trace each shadow region back towards the laser, until its entire cross-section is lit. All pixels on this trace are marked as invalid.

After marking all possible shadow pixels they are filled using a simple volumetric diffusion in the space-time stack of images, i.e. even across different laser sheets. The process is illustrated in Figure 3.16 (c) and (d).

Some small horizontal lines that do not extend to the outside are missed by this approach. We simply remove all horizontal features that are less than five pixels in height which is very small compared to any part of the real object (see Figure 3.16 (d)).

The filtering so far was performed on a space-time stack of images, which, as the final step, is resampled into world coordinates using the calibration data.

For isosurface searching methods such as Marching Cubes, it can be useful to downsample the binary volume further so as to obtain smoother normals from the grey values.

3.7 Results

In order to examine the versatility of our approach, we have acquired the surface shape of several challenging objects. Figure 3.17 illustrates the performance of our prototype fluorescent immersion range scanner compared to traditional scans acquired with a commercial laser range scanner (Minolta Vi910). The figure shows a scan from a single direction for both acquisition systems. Please note that the data produced by our system is essentially unfiltered, except for the simple image processing operations explained in Sections 3.5 and 3.6.

The horse model (top row) consists of two different materials. The first material is a black, rough plastic, while the second one has a higher albedo and a somewhat smoother surface. Both materials are glossy, showing significant surface highlights. The Minolta scanner captures the higher albedo material quite well but has problems with the dark material. In comparison, our approach captures the details of the surface very well and even captures the surface roughness of the dark plastic. The dolphin (middle row) consists of a translucent stone and is an example for a sub-surface scattering material with visible crystal structure beneath the surface. Again, the commercial scanner fails to provide a dense depth map, whereas our result reveals significant detail. In this example, however, we observe small holes in the fins of the dolphin which cannot be attributed to occlusion alone. The reason is an overly aggressive thresholding during the construction of the depth map where pixels with a low gradient component g_{μ} are removed. The lines on the fins are thus classified as stationary features. Our third example (bottom row) is a transparent die with diffuse, white dots. Actually, the object's transparent part is fluorescent as well, but the material's extinction coefficient happens to be high enough that all fluorescent emission is limited to the surface. The Minolta scanner can only capture the diffuse, white part of the surface, while our method produces a dense depth map for both material types. However,





Figure 3.17: By immersing objects into a fluorescent liquid, we have acquired high quality range scans for a set of surfaces that are difficult to acquire with traditional methods. From left to right: photo, fluorescent immersion range scan, traditional range scan (Minolta Vi910).

due to the fluorescence of the material, a conventional laser scanner equipped with a green laser would probably be able to capture the full surface of this object. The crystal goblet, Figure 3.18, is completely transparent and exhibits significant refraction due to its faceted surface. Nevertheless, the depth map recovered by our method (Figure 3.14) has only a few minor holes and is virtually free of noise. Using only 6 scans from different directions, almost the entire outer surface of the goblet can be reconstructed with high accuracy. By simple superposition of the individual calibrated scans and without the need for advanced mesh alignment methods, we achieve a low noise level in the order of 0.05% of the overall object size (Figure 3.18).

The volume slicing technique presented in Section 3.6 is geared towards capturing transparent objects. We demonstrate its effectiveness on two objects. The first is a bowl of quartz glass (n = 1.46) with a set of quartz glass cylinders of 10 mm diameter. Our other glass object is a mouse figurine, made from borosilicate glass (n = 1.48). Figure 3.19 shows the results of the two volume scans. Both objects were scanned with a potassium thiocyanate solution as refractive index matching liquid.

The bowl is a challenging object because of the numerous overlapping, lens-like cylinders that focus the light, requiring very precise refractive index matching. The size of the acquired volume model is $1100 \times 720 \times 1400$ voxels, the real world size of the object being





Figure 3.18: For a quantitative analysis of our first-surface scanning method, we took a cross-section of the goblet model (red line in (a)) and fitted an ellipse against it. Given 3000 sample points (b), we obtained an RMS error of 0.046 mm at an object diameter of 43.2 mm. The slight misalignment of the individual scans could probably be leveraged by geometry post-processing.

92 mm. This results in a voxel resolution of about 0.08 mm in each dimension. However, the overall resolution achieved in practice is a little lower. As can be seen in Figure 3.19 (c), in the back of the volume (from the camera's viewpoint) some artifacts appear. They result from the differing wavelength of the laser illumination and the light emitted by Eosin Y. As discussed in Section 3.4.1, the refractive index matching is only accurate for one particular wavelength. We have to match the refractive indices for the wavelength of the illuminating laser light to ensure planar light sheets while scanning the volume, and to suppress edge shadows (see Section 3.6). This, however, results in a slight mismatch of the refractive index for the light emitted via fluorescence, giving rise to minor artifacts in the back of the scanned volume. A comparison of slice images from different positions in the scanning volume is shown in Figure 3.20, illustrating the effect. The gap that is visible in the bowl in Figure 3.19(d) is caused by the fact that this part of the bowl was outside the calibrated volume. We missed it because the bowl is not visible in the refractive index matched liquid without the illumination of the laser sheet.

The second object, a hand-made mouse, has intricate detail and a complicated surface topology. Its length is approximately 40 mm and the acquired volume has a size of $500 \times 300 \times 500$ voxels, resulting in a similar effective resolution as in the previous example.

The filtered volumes can be visualised with traditional volume rendering techniques producing the x-ray images in Figure 3.19(b). After downsampling the data by a factor of 0.5 in all dimensions, we extract the isosurface to generate surface models (c) which we render using a glass shader in a synthetic environment in (d). For quantitative analysis, we extracted two of the glass cylinders, one close to the camera and the other further away, and fitted analytical models against them. Given 11500 samples on each cylinder, we obtained a RMS error of 0.037 mm for the near and 0.285 mm for the far cylinder.



3.7. RESULTS



Figure 3.19: The volume acquisition method enables the acquisition of an entire volume for transparent objects with a constant refractive index. The scan can be performed in one sweep with negligible post-processing. From left to right: photo (a), volume rendering of the acquired data set (b), isosurface mesh (c), glass shader rendering (d).



Figure 3.20: Slices of the glass cylinders dataset. Notice the difference in edge quality depending on the distance of the scanning laser sheet to the camera: near (left) and far (right).

3.7.1 Comparison to Existing Methods

While there have been many previous methods specialising on individual types of materials, e.g. for glossy [Chen06, Park08] or transparent or translucent scenes [Chen07, Morris07], only [Morris07] attempted the acquisition of objects with mixtures of different materials. The resolution obtained by our surface scanning method is on par with existing standard techniques for diffuse objects.

Comparing our volume scanning procedure with existing methods such as CT or MRI is problematic, since our algorithm is tailored to recovering a binary density volume. However, given the comparatively small acquisition and reconstruction effort of our method and the high resolution it provides, we believe that for objects made of clear, homogeneous materials our method is an attractive alternative to existing techniques.



3.7.2 Applicability and Limitations

Immersing Objects in Liquids

For most objects composed of plastic, glass, ceramics, stone, or metals the immersion in a liquid is uncritical. Care has to be taken with porous materials and substances that are soluble in water or prone to corrosion. On the other hand, the same holds for covering the surface with diffuse spray paint which is often alcohol-based and, as a side effect, modifies the surface geometry. Depending on the requirements of the object, the solvent for the Eosin molecules can be chosen from a wide range of fluids with different chemical properties, e.g. (deaerated) water, ethanol, glycerol, other alcohols or nonpolar oils, some of which are used by archaeologists to preserve organic objects.

Mirroring Materials

Strongly mirroring materials are still difficult to acquire. In the case of curved mirroring surfaces, strong caustics may appear in front of the surface, masking out any intersection line behind them. Given a perfect planar mirror, the method described above might fail because the change of intensity at the surface vanishes. These problems can probably be resolved using more sophisticated light patterns and/or multiple cameras. In fact, binary coded patterns might reveal more information about the surface, for instance normal directions (Figure 3.21.



Figure 3.21: Using a series of stipple patterns, it might become possible to identify metallic reflections and recover not only surface points but also projections of the normal vector.



High-Albedo Materials

For diffuse materials with high albedo (such as the calibration target in Figure 3.11), two effects become apparent. First, the surface is brightened up by the laser sheet in front of it (compare e.g. the brightness of the white surface to the one of the holes). Since we are dealing with diffuse reflection, this is a stationary, low-frequency effect which does not affect the detection routine. Even more prominent is the bright line at the intersection between laser sheet and surface, caused by volumetric scattering of the diffusely reflected light sheet. Under the assumption of isotropic scattering it can be shown that this effect can never become brighter than the light sheet itself. Nevertheless, its superposition with the incident light sheet leads to a small shift of the detected edge towards larger distances (by max. half the ray diameter). So far, we have not found a way to circumvent this small systematic error while still keeping the detection method as general as it currently is. For the calibration, however, using a diffuse red calibration target instead of the white one would be a perfect workaround.

3.8 Conclusion

The technique of fluorescent immersion 3D scanning offers novel means for capturing the geometry of objects that due to their material are difficult to capture with traditional range scanning techniques. Examples of such surfaces are dark, translucent or even transparent surfaces. Two techniques were proposed in order to scan surfaces and volumes, respectively. The key idea of both is to place the object into a tank filled with a fluorescent fluid. Rather than detecting peak intensities as in structured light scanning, the object's surface will appear as an intensity cut-off. In order to detect the edge, we have proposed a robust detection scheme based on the space-time gradient. We have found it to work reliably for a wide range of surface reflection properties of the object, producing high quality surfaces and volumes with moderate acquisition effort.

We see possible applications of our range scanning technique in reverse engineering, preservation of cultural heritage, and quality control, whenever traditional techniques fail and a modification of the object itself, for instance covering the surface with opaque powder, is not an option.

A limitation of the technique is that the object of interest has to be immersed in a liquid, which may not always be desirable. For instance, valuable historic artifacts or substances that corrode easily or are soluble in water may require careful handling. However, for such objects other liquids, such as de-aerated water (free of dissolved oxygen) or glycerol (which is actually commonly used to preserve organic tissues), might be an option. Alternatively, one could think of using a fluorescent gas, which to our knowledge does not exist at room temperature and with fluorescence in the visible spectrum, but might be possible in the infrared.

Using the fluorescent fluid, light rays can be made visible without the disturbing effects



of multiple scattering. Visualising light rays this way might inspire other novel acquisition methods, for example for visualising and capturing reflection properties of materials (see also Figure 2.9), or the joint acquisition of surface positions and normals of strongly mirroring objects.

Knowledge about the geometry of an object alone is not enough to reproduce its appearance in a synthetic rendering. Equally important is a reflectance model that describes the visual appearance of materials. After using fluorescence as a tool for geometry measurements, we will now turn to the phenomenon of fluorescence itself and try to reproduce the appearance of fluorescent objects.



Chapter 4 Bispectral BRRDFs



Figure 4.1: Fluorescence is a common property among many human-made and natural materials. Photos taken by author, except scorpion (Kenton Elliot), flowers (Bjørn Rørslett), and fish (NOAA/ Edie Widder).

We have seen that the effect of fluorescence can be used as a tool to robustly measure object geometry. However, the appearance of fluorescent materials can be visually appealing on its own right. We have learnt in Chapter 2.1.3 that fluorescence is essentially the sequence of an absorption and an emission event, and that fluorescent materials reradiate part of the incident light at a different wavelength due to their physical and chemical nature. In this chapter, we deal with the gonioreflectometric description of such materials, i.e., investigate how their reflectance depends on the illuminating and outgoing angle in order to faithfully reproduce their appearance in computer renderings.

Fluorescent objects play an important role in everyday life. Still, rendering them in



CHAPTER 4. BISPECTRAL BRRDFS

a physically correct way is costly and far from trivial. Given that their appearance under white light can be emulated by extreme (and often unphysical) colour values, the phenomenon of fluorescence has received limited attention within the field of computer graphics. In fact, only a handful of CG publications have dealt with this problem so far, mostly describing the phenomenon itself through manually tuned analytical models.



Figure 4.2: Fluorescent materials absorb part of the incoming light at wavelength λ_i and reradiate it at the longer wavelength λ_o . We have measured bispectral BRRDFs to capture such materials and render them in the spectral environment map of a winter sunset. The bottom row depicts slices of the bispectral BRRDF, showing one rendered sphere for each pair of incident and reflected or reradiated wavelengths (λ_o , λ_i) \in [400 nm; 720 nm] × [380 nm; 720 nm]. Fluorescence is represented by the off-diagonal entries.

Inspired by data-driven BRDF representations, we use a bispectral and bidirectional measurement setup to acquire reflectance and reradiation distributions of real fluorescent materials. Since the acquisition of these high-dimensional datasets (2 spectral plus 3 or 4 angular dimensions) is rather costly, we propose a low-rank decomposition that allows for more efficient data capture.

4.1 Related Work

4.1.1 Analytical and Data-Driven BRDF Models

The history of data-based BRDF models in the context of computer graphics goes back to the early nineties, when [Ward92] measured and modelled the BRDF of anisotropic materials. The first larger material database of 61 different, albeit sparsely sampled BRDFs emanated from the CUReT project [CUReT96]. Later, using a setup similar to [Marschner00], [Matusik03] measured more than 100 different materials, from which they derived a generative BRDF model. [Ngan05] compiled an overview of different models and how well they approximate BRDF measurements. Many of these BRDF models allow for spectrally



varying reflectance distributions, such as the Cook-Torrance BRDF [Cook81], but do not model bispectral distributions as needed for the reproduction of fluorescence.

4.1.2 Colour Characterisation and Spectral Imaging

It is widely acknowledged that the reproduction of colours benefits from spectral resolution above the usual red, green, blue (RGB) channels. A very comprehensive treatment of many questions within the field of colour science was given by [Nassau83]. More within the field of reflectance capture, [Proctor96] constructed a spectral BRDF measuring gantry featuring a tunable monochromatised light source and broadband receiver, which makes it suitable for spectral, but not for bispectral measurements.

The acquisition of spectrally resolved image data has been of interest in many fields such as life sciences, astronomy, aerial imaging, and many more, and there exists a large body of work [Gat00, Sonka07, Hearnshaw90]. The introduction of computation into the imaging process gave rise to novel, more efficient acquisition techniques of spectral imagery [Imai98, Hardeberg99, Rump10]. In this chapter, we apply the principle of low-rank capture in order to acquire high-dimensional bispectral bidirectional distribution functions in a practical way.

4.1.3 Bispectral Measurement: Fluorometry

Within the field of fluorometry, bispectral measurements are a long-established technique [Leland97]. In fact, the concept of a reradiation matrix dates back over half a century [Donaldson54]. Due to the high dimensionality of the reflectance and reradiation function, researchers usually put more focus on the spectral dimension and constrained themselves to very sparse angular sampling of BRDFs, typically at $0^{\circ}/45^{\circ}$ or $0^{\circ}/10^{\circ}$ when performing spectral or bispectral measurements [Angelopoulou01, Gundlach94, Hersch07]. In order to vary between these angular settings [Holopainen08] proposed a carefully calibrated bispectral goniometer setup, but the limited angular range and resolution prevents sampling a full BRRDF.

4.1.4 Rendering of Fluorescence

[Glassner94] was the first to adapt the rendering equation [Kajiya86] to include fluorescence (as well as phosphorescence). Within the field of computer graphics, the next significant step was taken by [Wilkie01], who introduced a rendering framework that is capable of treating polarisation and fluorescence effects in a joint fashion. More recently, Wilkie et al. went on to derive a physically motivated multilayer reflectance model that allows for the inclusion of fluorophores in the diffuse layer [Wilkie06]. In this work, they also performed a qualitative analysis of a fluorescent material, optimising their model parameters



to visually match the reflection of a green laser on an orange fluorescent cardboard. The authors conclude that while fluorescence is mostly a diffuse effect, it may still depend on the view/light directions, e.g., due to additional Fresnel effects.

In this chapter, our goal is not the derivation of parameters for an analytical model. Rather, we demonstrate how it is possible in a data-driven setting to acquire real-world reflectance data efficiently by taking into account the complexity of the given material.

4.2 Reflection and Reradiation

Fluorescent materials change the wavelength of light upon reflection. This affects many everyday materials, for instance human teeth, utility vehicle paints, detergents (fabric whiteners), or even ordinary photocopying paper. This shift of wavelength causes compelling visual effects if it occurs within the visible spectrum or turns UV radiation into visible light. In particular, many fluorescent surfaces appear brighter than perfectly white surfaces (Figure 4.3).

The underlying physical mechanism is well understood. A fluorescent medium consists of atoms or molecules that absorb incident photons at a given wavelength, and re-emit them after a short time (in the order of 10^{-8} s). During this time interval, the electrons of the fluorescent molecule remain in an excited state above the ground energy level. The re-emission of a photon occurs as the fluorophore relaxes to its ground state. Due to mechanical interaction with the surrounding molecules, some of the excitation energy is lost during this process, leading to a change of wavelength, or Stokes shift. For an illustration of the process, see Chapter 2.1.3. As required for conservation of energy, except in the case of multi-photon interactions, this shift always occurs towards longer wavelengths, corresponding to a loss in per-photon energy.

The wavelength-shifting behaviour of a fluorescent material can be intuitively described *bispectrally* using a so-called *reradiation matrix* [Donaldson54], specifying for each combination of incoming and outgoing wavelengths the amount of reradiated light. At the same time, however, the reflectance of every real-world material also depends on the directions of the incident and reflected light rays with regard to the surface, as usually described by the *bidirectional reflectance distribution function* (BRDF) [Nicodemus77] which we introduced earlier in Chapter 2.3.2. In this chapter, we integrate these well-known concepts into the *bispectral BRRDF*¹ that can describe general fluorescent (and non-fluorescent) materials and the bidirectional dependency of their wavelength-preserving reflectance and their wavelength-shifting reradiation.



¹In optics, the *bispectral luminescent radiance factor* is commonly used to describe fluorescent materials. This is inconvenient for our purposes, as it defines fluorescence relative to a perfect, non-fluorescent diffuser.

Note that reflection and reradiation are different physical mechanisms and are treated as such throughout the scientific literature. In a computer graphics context, however, it makes sense to abstract reradiation as an instantaneous phenomenon and treat it in the same way as reflectance.



Figure 4.3: Spheres covered in three different multilayer paint varieties for utility vehicles ("pearl white" primer, fluorescent red/yellow/orange paint, clear coat). Note the outstanding chroma of the paints, and their luminance, which even in the non-specular part is higher than the one of the 90% white ColorChecker patch.

While there have been attempts to approximate fluorescent materials with analytical models, to our knowledge the full bidirectional and bispectral reflectance of real materials has not been measured so far. By equipping a traditional BRDF measurement setup with spectral filters for light source and camera, we acquire bispectral BRRDF datasets for a range of isotropic fluorescent materials.

Given the measured data, we project them into a sparse spectral basis using principal component analysis (PCA). Such a decomposition, when performed on a partially acquired dataset, allows for efficient planning of the remaining acquisition task.

4.3 **Bispectral Reflectance and Reradiation**

For computer graphics purposes the phenomena of reradiation and reflection can be treated in an unified manner, albeit being physically different. We will henceforth refer to both phenomena as "bispectral reflectance" while keeping our terminology as compatible as possible with the metrology and physics literature.



Radiance	$L(\omega)$	$\left[\frac{W}{\mathrm{sr}\cdot\mathrm{m}^2}\right]$
Spectral Radiance	$L(\omega,\lambda)$	$\left[\frac{W}{\text{sr}\cdot\text{m}^2\cdot\text{nm}}\right]$
Spectral Irradiance	$E(\lambda) = \int_{\Omega} L(\omega, \lambda) \mathrm{d}\omega$	$\left[\frac{W}{m^2 \cdot nm}\right]$
Irradiance	$E = \int_{\Lambda} \int_{\Omega} L(\omega, \lambda) d\omega d\lambda$	$\left[\frac{W}{m^2}\right]$

Table 4.1: Definitions of spectral quantities; ω refers to directions and λ to wavelengths.

4.3.1 Bispectral Rendering Equation

We have learnt that the light transport in a scene can be expressed by the rendering equation (Eqs. 2.13 and 2.14). Light transport considering energy transfer from one wavelength to another, in order to account for fluorescence, can be expressed by the bispectral rendering equation:

$$L(\omega_{o},\lambda_{o}) = L_{e}(\omega_{o},\lambda_{o}) +$$

$$\int_{\Omega} \int_{\Lambda} L(\omega_{i},\lambda_{i}) f_{r}(\omega_{o},\omega_{i},\lambda_{o},\lambda_{i}) cos\theta d\lambda_{i} d\omega_{i},$$
(4.1)

which requires an additional integration over all incident wavelengths λ_i . (In optics and other engineering fields, the letters λ and μ are often used to denote the incident and outgoing wavelengths. In a global illumination setting, however, the μ reflected on one surface will become the next reflection's λ , and so we choose λ_o ("out") and λ_i ("in") in order to avoid confusion.)

Solving the bispectral rendering equation is relatively straightforward when using spectral path tracing, e.g. [Bendig08], but of course other rendering techniques can be adopted as well.

4.3.2 Bispectral BRRDF

In the monochromatic rendering equation, the reflectance of surfaces is described by the BRDF. Here, it is replaced by the *bispectral BRRDF* $f_r(\omega_0, \omega_i, \lambda_0, \lambda_i)$ that describes the angularly dependent reflectance for any pair of wavelengths. A mere *spectral BRDF* as the spectrometric extension (Chapter 2.2.2) cannot represent fluorescent materials. Likewise, reradiation matrices as proposed by [Donaldson54] fail to describe the directional dependence of the observed reradiation.

Before we provide the general bispectral BRRDF, let us briefly recall the definition of a spectral BRDF. The required spectral quantities are defined in Table 4.1. We have seen in Chapter 2.2.2 that the spectral quantities feature a different unit compared to their nonspectral counterparts since they are essentially derivatives with respect to the wavelength. Following [Nicodemus77], the differential reflected spectral radiance $dL_o(\omega_o, \lambda_o)$ due to



4.4. MEASUREMENT AND RECONSTRUCTION

the incident differential spectral irradiance $dE(\lambda)$ from direction ω_i is given as:

$$dL_{o}(\omega_{o},\lambda) = dE(\lambda)f_{r}(\omega_{o},\omega_{i},\lambda) \left[\frac{W}{\mathrm{sr}\cdot\mathrm{m}^{2}\cdot\mathrm{nm}}\right], \qquad (4.2)$$

with ω_i and ω_o being the incident and outgoing directions. The spectral BRDF $f_r(\omega_o, \omega_i, \lambda)$ for a single wavelength is therefore defined as the ratio of differential reflected spectral radiance to differential incident spectral irradiance:

$$f_{\rm r}(\omega_{\rm o},\omega_{\rm i},\lambda) = \frac{dL_{\rm o}(\omega_{\rm o},\lambda)}{dE(\lambda)} = \frac{dL_{\rm o}(\omega_{\rm o},\lambda)}{L_{\rm i}(\omega_{\rm i},\lambda)\cos(\theta_{\rm i})d\omega_{\rm i}} \left[\frac{1}{\rm sr}\right].$$
(4.3)

It follows that the unit of the spectral BRDF is $\left[\frac{1}{sr}\right]$, which is the same as for non-wavelength dependent BRDFs $f_r(\omega_0, \omega_i)$, although the units for *L* and *E* differ in the spectral vs. non-spectral case.

We now generalise Nicodemus' derivation of the BRDF to account for cross-wavelength energy transfer by the bispectral BRRDF and show that its unit differs from the spectral BRDF. Referring to the bispectral rendering equation (Eq. 4.1) the differential reflected and reradiated spectral radiance (differential with regard to the incident direction ω_i and the incident wavelength λ_i) is due to incident double differential (*non-spectral*) irradiance for ω_i and λ_i :

$$d^{2}L_{o}(\omega_{o},\lambda_{o}) = d^{2}E \cdot f_{r}(\omega_{o},\omega_{i},\lambda_{o},\lambda_{i}) \left[\frac{W}{\mathrm{sr}\cdot\mathrm{m}^{2}\cdot\mathrm{nm}}\right], \qquad (4.4)$$

and hence the bispectral BRRDF may be defined as

$$f_{\rm r}(\omega_{\rm o},\omega_{\rm i},\lambda_{\rm o},\lambda_{\rm i}) = \frac{{\rm d}^2 L_{\rm o}(\omega_{\rm o},\lambda_{\rm o})}{L_{\rm i}(\omega_{\rm i},\lambda_{\rm i})\cos(\theta_{\rm i}){\rm d}\omega_{\rm i}\,{\rm d}\lambda_{\rm i}} \left[\frac{1}{{\rm sr}\cdot{\rm nm}}\right]. \tag{4.5}$$

The bispectral BRRDF is a general way to represent fluorescent materials as it does not make any assumptions about the material.

In the discretised case, an individual sample of the bispectral BRRDF for the directions (ω_i , ω_o) expresses the energy transfer from the incoming spectrum to the reflected spectrum as a matrix over λ_o and λ_i , see Figure 4.2. While the diagonal entries refer to reflection at the same wavelength, the fluorescent effect is represented by the off-diagonal part. As there is typically no transfer from longer to shorter wavelengths (towards higher energy), the upper triangle will remain black.

4.4 Measurement and Reconstruction

4.4.1 Setup

In order to acquire isotropic bispectral BRRDFs, we have built a fully automated imagebased measurement device. It follows the design of [Matusik03] for isotropic BRDFs but with the added capability to emit and sense at specific wavelength bands (Figure 4.4).





Figure 4.4: A depiction of our setup. A sample sphere (1) is mounted on a turntable (2), to which a digital monochrome still camera (3) is attached. The camera is equipped with a visible-spectrum tunable filter (4). The sphere is illuminated by a light guide coupled xenon light source (5) with another tunable filter (6) mounted in front; near-UV light is generated with LEDs that can be selected using a motorised wheel (7). On the exit aperture of (6) and the entry aperture of (4), we attach optical depolarisers.

The spectral filters used are LCD-based Lyot filters (CRi VariSpec VIS10/35 mm) whose transmission bands are about 10 nm–20 nm wide and range from 400 nm to 720 nm. We apply additional polarisation scrambling optics to undo the linear polarisation from the LCD filters so as not to bake any unwanted side effects into the measured BRRDF. Figure 4.5 illustrates the strong influence of polarisation both on the specular and the diffuse reflection for a sample with clear coat. Our experiments show that even non-coated, apparently diffuse surfaces do not necessarily completely decorrelate the polarisation state.

As light source we employ a xenon arc lamp coupled into a light fiber (XION medical Xenon R180), whose light has a flat and stable spectrum (measured using a spectroradiometer) but rather weak blue and UV output, especially after passing the spectral filter. We therefore add LEDs for better coverage of this range (370 nm-420 nm in 10 nm steps). The camera is a monochrome, digital still camera (Jenoptik ProgRes MFcool), with which we acquire high-dynamic-range images using exposure series from 1 ms to 16 s.

4.4.2 Geometric and Spectral Calibration

Calibrating such a measurement setup geometrically and spectrally is not a trivial task. On the geometrical side, [Havran05] propose Helmholtz reciprocity as a validity criterion for the measurement setup. However, luminescent materials (that do not preserve the wavelength of light) do not fulfill Helmholtz reciprocity by nature [Springsteen99, Clarke85].





Figure 4.5: Influence of polarisation on reflection from a sphere: Average image and difference images (red: negative) depending on the polarisation state of light source and observer, where "H/V" stands for "horizontal in, vertical out". Near the Brewster angle, the specular highlight is contained almost exclusively in the V/V component. Also, note the variation in the diffuse regions.

Therefore, we chose to treat the spectral and geometric calibration independently of each other. Geometric validation of our BRRDF measurement setup is provided by a reciprocity check using a non-fluorescent sample sphere.

The spectral response of the system is determined by taking into account the spectral transmission or sensitivity curves for all of the components [DeRose07, Holopainen08], either as obtained from the manufacturers or calibrated using a photospectrometer. We obtain the absolute scaling from a bispectral measurement on a flat Spectralon target, which preserves the wavelength and has a well-defined reflectance (99% quasi-Lambertian reflectivity across the spectrum). Due to a slight crosstalk between neighboring filter bands, we observe a small contribution of the wavelength-preserving reflection to non-diagonal elements in the bispectral matrix. For normalisation, we require the rows of the captured bispectral Spectralon BRRDF sample to add up to 0.99 sr⁻¹, which, along with the transmission curves for the illumination and detection branch, gives us a spectral scaling factor $s(\lambda_0, \lambda_i)$ for each wavelength pair. Any captured camera value is scaled according to $s(\lambda_0, \lambda_i)$ before generating one entry in the matrix of a bispectral BRRDF sample.

4.4.3 Measurement and Data Processing

The straightforward way of acquiring a bispectral BRRDF is to capture images at all turntable rotations β for every pair of wavelengths (λ_0 , λ_1). For practical reasons we constrain ourselves to 20 nm steps in the range from 380 nm to 720 nm for λ_1 and 400 nm to 720 nm for λ_0 , amounting to 170 images per β as the upper triangle of the bispectral matrix can be ignored. We vary β in the range of 5° to 170°. For highly specular materials a stepping of 5° is chosen to sufficiently sample the sharp highlight while we take a coarser sampling of 10° for materials of lesser angular bandwidth. The acquired data then undergo geometric processing and resampling.

Adaptive measurement. In the case of a PCA-steered measurement (Section 4.5), we first acquire full bispectral data sets for a small number of turntable angles. After performing the PCA decomposition, we acquire dense angular data only for the sparse bispectral





Figure 4.6: Average and the first 10 principal components B_i.

basis that is required for a good reconstruction. Only a small set of wavelength pairs needs to be measured.

Geometric processing. Depending on the material, we use two different sample geometries: coated spheres for the paints and a custom-made piecewise cylindrical object (Figure 4.14) wrapped in stripes of paper or fabric.

Due to the varying normals of the shapes each surface point will be illuminated and viewed from a slightly different direction. From simple geometric considerations we can determine (ω_0, ω_i) for every pixel captured under a specific turntable rotation β .

For storage and further processing, we discretise the data for each wavelength pair in 32^3 bins using the (θ_0 , θ_i , ϕ_{diff}) parameterisation chosen in [Matusik03]. For the strongly specular materials, 64^3 bins are used. Bins that are not populated due to the coarse sampling of the turntable position are filled in by diffusion.

In the rendering step, a specific reflectance sample $f_r(\theta_0, \theta_i, \phi_{\text{diff}}, \lambda_0, \lambda_i)$ is obtained by multilinear interpolation from our bispectral BRRDF representation.

4.5 PCA-based Acquisition

In this work, we opt for a data-driven representation of the bispectral BRRDF because we do not want to make strong *a priori* assumptions about the spectral and angular behaviour of the material. Particularly, we know that due to the different nature of fluorescent reradiation and specular reflection, the overall spectral and angular variation of most bispectral BRRDFs (say, a diffuse red material with a white highlight) is not strictly separable in the


form

$$f_{\rm r}(\omega_{\rm o},\omega_{\rm i},\lambda_{\rm o},\lambda_{\rm i}) = f^{\lambda}(\lambda_{\rm o},\lambda_{\rm i})f^{\omega}(\omega_{\rm o},\omega_{\rm i})$$
(4.6)

This is unfortunate, since it would have allowed us to measure angular and spectral dependencies f^{ω} and f^{λ} separately and simply compute the bispectral BRRDF as the outer product of both functions. On the other hand, it is always possible to expand f_r into a series of separable terms:

$$f_{\rm r}(\omega_{\rm o},\omega_{\rm i},\lambda_{\rm o},\lambda_{\rm i}) = \sum_{n} f_{n}^{\lambda}(\lambda_{\rm o},\lambda_{\rm i}) f_{n}^{\omega}(\omega_{\rm o},\omega_{\rm i})$$
(4.7)

Since fluorescence itself is weakly directional due to its physical nature [Zastrow81], a PCA will likely yield such a decomposition of low rank. We exploit this insight and perform a dense bispectral measurement under only a sparse set of turntable angles $(0^{\circ}, 70^{\circ} \text{ and } 150^{\circ},$ each of which corresponds to a 2D slice of the BRDF). The measurements from this sparse set of turntable positions contain samples from a lot of different angles of incidence and exitance. The bispectral correlations found in these measurements can then be transferred to other angles. Figure 4.12 illustrates this for measurements taken at two turntable positions.

We assemble a matrix F which contains all bispectrally-valued BRRDF samples with the average value \overline{f} subtracted, and compute its SVD. Selecting the *n* eigenvectors with the greatest eigenvalues the basis *B* is assembled. Examples of such eigenvectors can be seen in Figure 4.6.

Measurement basis. Since our narrowband filter assembly only allows for the sampling of wavelength pairs $\underline{\lambda}_k = (\lambda_0, \lambda_i)_k$, we cannot measure in the PCA basis directly. This calls for a basis transformation, and a selection of wavelength pairs that best represents the PCA basis.

Assume we have a set of bispectral vectors **B** as returned by the PCA. Given a set of mixture coefficients \mathbf{x} , we can compute a resulting reradiation matrix \mathbf{b} where each entry corresponds to a wavelength pair:

$$\mathbf{B}\mathbf{x} = \mathbf{b} \tag{4.8}$$

$$\left(\left(\begin{array}{c} \\ \end{array}\right)\left(\begin{array}{c} \\ \end{array}\right)\cdots\right)\left(\begin{array}{c} x_1 \\ x_2 \\ \vdots \end{array}\right) = \left(\begin{array}{c} b_{\underline{\lambda}_1} \\ b_{\underline{\lambda}_2} \\ \vdots \end{array}\right)$$
(4.9)

In this notation, we think of reradiation matrices in a vectorised way, such that



although they represent bispectral data, are treated as column vectors.

By picking a sparse set of wavelength pairs p_i , we set most of the entries in **b** to 0. For illustration, let us assume we pick three wavelength pairs, $p_1 \dots p_3$, and measure the



corresponding samples $m_1 \dots m_3$. Doing so corresponds to selecting a subset of the rows of **B**, from which we can assemble a new reduced matrix **R**:

$$\begin{pmatrix} \begin{pmatrix} \bullet & P_1 \\ P_2 & \bullet \\$$

Our wavelength pair samples p_i now have to be chosen in a way that the invertibility of **R** is maximised. Starting with the wavelength pair corresponding to the greatest entry in **B**, we follow a greedy strategy selecting the set of samples which minimises cond(**R**).

The mixture coefficients x_i in the PCA basis are obtained through the pseudoinverse of **R**:

$$\begin{pmatrix} x_1 \\ x_2 \\ \vdots \end{pmatrix} = \mathbf{R}^+ \begin{pmatrix} m_1 \\ m_2 \\ m_3 \end{pmatrix}$$

We can now use a few wavelength pair samples as a basis to reconstruct the full bispectral information from the PCA basis vectors.

4.5.1 Insights

Same-angle reconstruction using standard PCA. In Figure 4.7, we show a reconstruction of a single-angle bispectral measurement from the PCA basis for the same angle. The error measure provided is given by the residual energy as determined by the SVD and relates to the full bispectral dataset with 170 wavelength pairs, not just the resulting sRGB colour vectors as shown in the figure. Although the fidelity increases with the cardinality of the measurement basis, the visual appearance even at a numerical error as low as 1% or 0.1% is not always fully satisfactory. It is due to the least-squares nature of the SVD that materials with particularly strong specular highlights (Orange, Speckled) attract the attention of the first few eigenvectors at the expense of a slower convergence in the nonspecular regions. Remarkably, a decent reconstruction is often achieved even before the inclusion of the first measurement of an off-diagonal wavelength pair. In all five examples, the first 9 components were based purely on wavelength-preserving measurements.

Speckled dataset. The "Speckled" material (Figure 4.8) is a fluorescent yellow sphere onto which we manually applied a non-uniform layer of red speckles. While this sample does not have a homogeneous BRRDF, we included it as it represents a material with





Figure 4.7: Image-based same-angle reconstruction of full bispectral data (20 nm resolution) from a small number of acquired bispectral samples. As a reference, we provide ground truth in the top half of each image, and the number of basis vectors required for the residual energy to drop below 1% and 0.1%, respectively. Note that in the case of strong specular highlights (Orange, Speckled), the numerical error does not reflect the visual difference well.





Figure 4.8: The *Speckled* sample, while not a homogeneous BRRDF, represents the class of mixed materials with angularly dependent visibility of the individual components. Note the red shift towards grazing angles in the defocused shot (right). Photos taken under 420nm light.



Figure 4.9: First 5 eigenvectors for the Speckled dataset with the specular highlight masked out. Intuitively speaking, the first component constitutes a diffuse BRDF, while with the second the different material speckles are resolved. The number below each image is its relative energy as per PCA.

complex microstructure resulting in directionally dependent fluorescence, as can be seen in the defocused image (Figure 4.8 right). Again, the sharp highlight dominates the first few eigenvectors of a classical PCA. However, when the highlight region is masked out, we obtain a meaningful decomposition that nicely reflects the "eigenpaint" character of this material, see Figure 4.9.

Taming the specular highlight. It is to be expected that a signal decomposition based on a L^1 measure would no longer overemphasise the importance of the specular highlight. However, the few L^1 -PCA approaches in existence are computationally rather expensive and, due to their nonlinear optimisation scheme, offer no guarantee of global convergence [Baccini96, Ke03]. In order to emulate a similar behaviour using a standard PCA, we apply a saturation function on the values in the covariance matrix (and its inverse after the reconstruction), which dampens the high pixel values in the highlight region (Figure 4.10). As a consequence, the visual convergence is sped up considerably (Figure 4.11).





Figure 4.10: As saturation function $\tilde{x} = \operatorname{sat}(x)$, we chose the inverse of $x = \operatorname{sign}(\tilde{x})(\tilde{x}^2 + |\tilde{x}|)$, which is linear for small values of |x|, but of $\mathcal{O}(\sqrt{x})$ asymptotically.

Angular dependency. For most materials (except Speckled), the resulting principal components show that the fluorescent entries in the reradiation matrix $\lambda_i \neq \lambda_o$, taken for themselves, carry a rather weak angular dependency, i.e., indicating reasonable separability. However, as soon as the full matrix, the full bispectral BRRDF including the non-fluorescent elements, is considered, at least two, most often even more separable functions are required for faithful reproduction.

Angle transfer. Earlier on, we assumed that the bispectral decomposition of an angular BRRDF slice can be transferred to the spectrally sparse measurements taken under different angles. In Figure 4.12, we provide experimental evidence for this assumption. Using a PCA basis that was computed using fully-bispectral measurements at turntable angles of 0° and 70° , we reconstruct a fully bispectral intermediate slice at 35° out of only five measured wavelength pairs.

Our PCA-guided measurement routine allows us to drastically reduce the acquisition cost. As an example, let us assume a sampling of 5° steps from 0° to 150° for the turntable position. If all 170 bispectral entries are captured for each angle, the total time of optical exposure amounts to approx. 45 hours. By performing the full bispectral measurement under three turntable orientations only, and by measuring only 5 out of 170 wavelength pairs for the remaining angles, the acquisition time drops to 5.5 hours.

4.6 Results

4.6.1 Acquired Datasets

Using our measurement device and acquisition scheme, we have captured bispectral BR-RDFs of a number of fluorescent materials, including fluorescent paints with or without clear coating as well as paper and white cloth. As shown in Figure 4.2, they all have different wavelength ranges for excitation and emission.





Figure 4.11: By applying a saturation function on the covariance matrix, the convergence in the non-specular parts is significantly improved.

The strength of the reradiation therefore depends heavily on the illuminating spectrum as demonstrated in Figure 4.13 and Figure 4.14. In Figure 4.14, RGB photographs are compared against renderings under a 5600K illuminant. As our illumination system contains near-UV LEDs, we can even capture materials such as the paper sample which exhibit significant reradiation in the blue to UV range. The effect is clearly visible in Figure 4.13. Our captured bispectral BRRDFs faithfully reproduce the fluorescence in both images.

4.6.2 Spectral Detail Level

In Figure 4.16, we demonstrate that fluorescence does in fact require bispectral modelling of sufficient resolution. After reducing the measured bispectral BRRDF samples to a 3×3 (RGB×RGB) matrix by integrating over the RGB spectral curves, the renderings show clear differences to the full bispectral BRRDFs. Especially for the Yellow sample, the coarse RGB×RGB representation is unable to reproduce reradiation, which is sharply centered around 540 nm. We also integrate the bispectral BRRDF into a spectral BRDF by assuming a uniform illuminant spectrum. Again, the differences can clearly be seen. While they are less pronounced because the illumination in this scene is similar to the spectrum used for the conversion, slight deviations in colour and intensity can still be made out. Reducing the bispectral BRRDF to a simple RGB BRDF (again assuming a uniform incident spectrum) shows obvious differences. These differences are most pronounced for non-white spectra as demonstrated in Figure 4.15.





Figure 4.12: Reconstruction of an intermediate angle, which was not sampled for the PCA, from 5 wavelength pairs using the standard and saturated PCA approaches. Again, the top half of the reconstructed images is ground truth. The PSNR has been computed separately for the highlight area and the rest.

4.7 Conclusion

The common optical phenomenon of fluorescence calls for a bispectral description in order to obtain faithful computer renderings. We have provided the definition of the bispectral BRRDF, which models direction-dependent fluorescent behaviour, and a setup to measure such distributions for isotropic fluorescent materials.

Even though real-world BRRDFs are not directly separable into spectral and angular functions, they are intrinsically of low rank. We can therefore apply a PCA-steered acquisition scheme that only measures relevant bispectral samples of the BRRDF, resulting in a significant speedup (approx. 9:1), rendering such acquisition practical.

One of the shortcomings of a standard PCA in this context is its optimisation of an L^2 error function. As a consequence, the importance of specular highlights is often overemphasised, which leads to slower convergence in the non-specular regions. While a L^1 -based decomposition could potentially resolve this issue, we reached a simpler solution that is just as effective. By reweighting the covariance matrix before performing the SVD, we are able to improve the reconstruction fidelity for very small numbers of measured components.

For most of our materials, we can thus reach a visually indistinguishable reconstruction using only 5 or even fewer measurements per angle (as compared to 3 for a standard RGB measurement). This should allow for efficient acquisition of even higher-dimensional functions (anisotropic bispectral BRRDFs, spatially varying distributions or reflectance fields) in the future.

As a possible extension to our setup, we consider replacing the narrowband tunable spectral filters with hyperspectral devices. Measurements could then be directly performed



CHAPTER 4. BISPECTRAL BRRDFS



Figure 4.13: Renderings (top row) and photos (bottom row) of different materials under UV light (400 nm).



Figure 4.14: Renderings (top row) and photos (bottom row) of different materials under 5600K illumination.

using bi-hyperspectral basis vectors rather than individual wavelength pairs. The resulting optical path would have a higher overall transmission, which would lead to faster exposure times and/or better noise figures. Furthermore, bi-hyperspectral measurements would allow for advanced computational sensing approaches such as multiplexed illumination or compressed sensing.

With that, we conclude the part of this work that was aimed at the acquisition of realworld material and geometry data with the goal of producing computer-generated renderings. For the final scenario, we will do the exact opposite and bend real-world geometry so as to mimic the reflectance of materials in a computer-controllable way.





Figure 4.15: Measured fluorescent red bispectral BRRDF (a) compared to a simple RGB vector valued BRDF (b) under blue illumination. Note the cross-colour reflectance from blue to red in the case of the full bispectral BRRDF. The RGB BRDF cannot represent these complex colour shifts and fails to reproduce the fluorescent effect.



Figure 4.16: Comparison renderings using 3 different measured fluorescent materials. Full bispectral BRRDF measurements (right half of each sphere) are compared to: spectral measurements, RGB×RGB reradiation matrices, and standard RGB BRDFs.



CHAPTER 4. BISPECTRAL BRRDFS

Chapter 5 Dynamic BRDF Display



Figure 5.1: Left: Idea of a "reflectance display" that reacts to its environment like real-world materials do. Right: Photo of our device that can exhibit different degrees of anisotropic surface roughness.

After learning about the acquisition of reflectance and geometry in the previous chapters, we will now turn to the output or display side. In a time where technology is constantly evolving, the emerging technical possibilities bring along new paradigms of interaction and immersion, gradually blurring out the boundaries between real and virtual worlds. Somewhere in the future, content displayed on computer displays may become visually indistinguishable from the real world, and there are in fact substantial efforts in academic and industrial research to work towards such hyper-realism.

In this chapter, we focus on one major clue that helps distinguish real from virtual objects: reflectance. If we move around an object, and its appearance, the highlights, etc. change consistently with what we are used to and in accordance with the surrounding world, this object is more likely to be perceived as "real".

In order to achieve full realism for computer generated content, the display of the fu-



ture will need to behave like a showcase window through which the real and virtual worlds can interact with each other. While there have been efforts to reproduce the observer dependency, or at least achieve an impression of depth through stereoscopic means or other techniques such as head tracking, we want to focus on a different aspect. So far, all types of computer displays have shown pixels of different *colours*. Our goal is to physically mimic the characteristic way different surfaces reflect light, often described in terms of the Bidirectional Reflectance Distribution Function (BRDF), and display *materials* instead of colours (Figure 5.1).

In the recent years, many methods have been developed to fabricate materials with custom reflectance and sub-surface scattering properties. Yet, to our knowledge these properties have not been displayed dynamically. In Section 5.1, we discuss the prior work in related fields. The contribution of this chapter is a very first step towards the dynamic display of materials by means of a physical device that can be programmed to exhibit a range of reflectance distributions. Before we elaborate on our own approach, we lay the foundations in Section 5.2 with a general definition of the problem, along with a few sketches to its solution.

Of major importance to the appearance of real-world materials is their microstructure, e.g. in the form of surface roughness. By shaping surfaces, specific reflectance distribution functions can be achieved, as in the case of polished or brushed steel surfaces, or sandblasted glass. Our approach to dynamically displaying reflectance is also based on roughness modulation: we start with a liquid surface and excite surface waves on it. The space of possible appearances is defined by the reflectance of the base material, and the achievable surface structures that are in turn governed by the physics of wave propagation on liquids. In Section 5.3, we provide the underlying theory, and make a few basic predictions that will later be checked in experiment.

In Section 5.4, we present two single-texel prototypes that modulate the angular variation of reflected light in an optically passive way. We show in Section 5.5 that our devices can produce a range of anisotropic BRDFs that match our theoretical expectations. Using implementations of the same principle on different scales, we show that miniaturisation is not just possible but also desirable.

We conclude by discussing the limitations of our devices in Section 5.6, and providing directions towards more practical implementation and a wider range of displayable BRDFs.

An in-depth treatment of the relations between dielectric surfaces and the resulting reflectance is provided in Appendix A.



Display approach	Optics	Light	Geometry /Material	Viewpoint
Virtual reality	active	virtual	virtual dynamic	simulated
[Jones07, Ito10]	active	virtual	virtual dynamic	real
Augmented reality [Raskar01]	active	virtual	real static	real
[Debevec98]	active	real	virtual dynamic	simulated
[Cossairt08, Koike08, Hirsch09]	active	real	virtual dynamic	real
[Fuchs08] and fabrication	passive	real	real static	real
BRDF display (ours)	passive	real	real dynamic	real

Table 5.1: Light and observer-dependent rendering and display techniques at a glance.

5.1 Related Work

5.1.1 Representation and acquisition of reflectance

From the very early age of computer graphics research, it has been recognised that reflectance models are a crucial ingredient for realistic rendering. Torrance and Sparrow [Torrance67] had been the first to provide a shading model based on microfacet geometry, Phong [Phong75] and Blinn [Blinn77] showed its first applications in computer graphics. As the technical possibilities grew and the demand for physical accuracy increased, researchers began fitting model parameters against sparse gonioreflectometric data [Ward92] and provided databases of measured BRDFs [CUReT96, Matusik03]. Ngan et al. [Ngan05] related many of the previously introduced analytical BRDF models to dense reflectance measurements taken from real materials.

We use the established Ward BRDF [Ward92] as a reference model for the reflectance distributions exhibited by our device.

5.1.2 View- and light-dependent rendering and display

Rendering with measured real-world environment lighting can yield a great degree of realism with moderate technical and artistic effort, as first demonstrated by [Debevec98].



Raskar et al. [Raskar01] used computer-controlled lighting to make real objects look in a desired way. Their "shader lamps" require total darkness for best performance. On the observer side, the virtual viewpoint can be controlled by the user through different means, e.g. through game input devices or head tracking as in many early Virtual Reality installations. In the recent years, autostereoscopic displays have been constructed (optionally combined with head tracking) to achieve freely viewable 3D imagery [Jones07, Ito10] and even more based on lenticular or parallax barrier principle [Hirsch10, Kim10]. Various active devices have been demonstrated that combine light field sensing and/or display with intermediate processing [Cossairt08, Koike08, Hirsch09].

5.1.3 Fabrication

The last three years have spawned a large amount of work dedicated to fabricating materials and objects with custom properties. The efforts include milling of height fields to reproduce reflectance distributions [Weyrich09], printing of spatially varying BRDFs on paper [Matusik09], fabrication of subsurface scattering materials [Hašan10, Dong10], surface reliefs that show lighting-dependent images [Alexa10], and even custom deformability [Bickel10]. The reflectance field assembly by Fuchs et al. [Fuchs08] stacked purely passive "pixels", each encapsulating a full 4D transmittance field, albeit at very limited resolution.

5.1.4 Our Contribution

As the key contribution of this work, we see the introduction of the BRDF display as a dynamic alternative to fabrication. We outline the main characteristics that any BRDF display should possess, and demonstrate a design that meets most of these requirements. To our knowledge, our device is the first of its kind in that it is both optically passive and programmable (see Table 5.1).

5.2 Displaying Reflectance



Figure 5.2: Left: optically active setup with a camera, a processing stage (here: negation) and a projector. Right: optically passive setup consisting of an imaging lens and a diffusor sheet to achieve a blurred image.



We have learned in Chapter 2.3 that the visual appearance of materials is defined by their interaction with light on a microscopic scale. Even homogeneous materials can reflect incoming light in different ways, if only their microscopic surface roughness follows a different statistic.

The most common mathematical description of surface reflectance is the BRDF (Chapter 2.3.2). Consequently, a BRDF display is a device that can be programmed to exhibit varying reflectance distributions. We propose the following criteria that any BRDF display should meet:

- 1. **View and light dependence.** The reflectance of a BRDF display must vary with the viewing and lighting direction in a physically plausible manner.
- 2. **Bell-shaped highlights.** For the vast majority of glossy real-world materials (as well as most analytical BRDF models, for that matter), the highlights follow Gaussian or power-of-cosine distributions, or sums thereof.
- 3. **Spatial extent.** In order for an observer to appreciate the reflectance function, a display needs to have a certain minimum size to display the shape of highlights.
- 4. Light efficiency and contrast. Any setup will fail to convince if it has insufficient light throughput, especially when competing against undesired effects such as first-surface reflections.
- 5. **Optically passive construction** (Figure 5.2). Setups with an optical path involving electronic imaging devices and computation allow for the most general ways of modulation. Purely optical setups, on the other hand, offer immediate response (all "processing" being done at the speed of light) and virtually unlimited dynamic range. Since real-world reflectance is optically passive, it must be possible to mimic it by passive means.
- 6. Virtuality/programmability. Lastly, the display has to be controlled by a computer in a nonpermanent way.

The following features would not be amiss, but we consider them optional:

- 7. Multi-lobe BRDFs, for instance the popular case of diffuse+glossy+specular.
- 8. Anisotropy, where the reflectance varies with the tangent orientation of the sample.
- 9. **Spatial variation,** where the display is formed by an array of individually controllable "texels".
- 10. **Normal variation,** where the perception of reflectance is further supported by displaying non-flat surfaces.



5.2.1 Approaching the Problem

While our task is rather well-defined—modulate the angular distribution of light reflected on a surface—it is hard to come up with "the" one ideal path to its solution. In fact, a host of different approaches can be imagined, each one with its own set of advantages and drawbacks. However, for optically passive setups it appears as if there are only very few families of operation principles.

Holography. In principle, holographic techniques could be applied to display different information, i.e. reflectances, for different *viewing* directions. These techniques are referred to as angle-multiplexed holograms, e.g. [Mok93]. Volume holographic storage devices [Orlov00] are reacting to selective directional *illumination*. Holographic *wavelength* multiplexing [Rakuljic92] could even store wavelength-dependent BRDFs [Hullin10]. However, these techniques do not work outside of high-end laser laboratories. To our best knowledge, the combination of these different holographic techniques to form a reflectance field display has also, so far, not been demonstrated. A computer controlled holographic display [Nwodoh00] could be combined with light field sensing devices to build an active BRDF display system. We aim, however, for a passive realisation. Thus, holographic techniques can be considered hypothetical to date.



Figure 5.3: Concept of a multiplexing setup consisting of a dual-pitch lens array, an elliptic mirror and a modulation mask (stipple pattern). Within the resolution limits, the transmitted radiance for all combinations of incident and outgoing angles can be modulated multiplicatively in the focal manifold. Since the the lenslet assembly transmits in both ways, the transmittance of a light path (red) is averaged with its reciprocal counterpart (blue). The modulation can be static (printed transparencies as in [Fuchs08]) or dynamic, e.g., using an LCD panel.

Integral Photography inspired approach. Using an optical multiplexer, the 4D space spanned by the incident and outgoing hemispheres is mapped to a plane, where the corresponding radiance values can be modulated by a 2D mask. A similar idea was implemented in transmission by [Fuchs08]; a reflective counterpart could for instance be imagined using



lenslet and mirrorlet arrays. While this approach allows for almost arbitrary modulation, its resolution is inherently limited. An elaborate system of diffusors may be necessary in order to achieve an impression of continuity, since a lenslet array that shoots out rays from different locations will not *look* like reflectance, unless a massive array of identical pixels is being viewed from a large distance. More importantly, though, any design based on this principle will exhibit very poor light throughput when it comes to displaying sharp highlights. Since for every possible combination of incident and outgoing rays a light path is provided internally, the only way to display a sharp highlight (delta distribution) is to attenuate all light paths that do not fulfill the reflection condition, i.e., pretty much all of them. Considering the above, any implementation of the Integral Photography idea will have massive difficulties meeting Criteria 3, 4 and 9. Also, a dynamic but optically passive device has yet to be demonstrated.



Figure 5.4: Two instances of the redistribution principle. Left: a switchable diffusor (blue) turns into a binary BRDF display when mounted on a mirror (grey). Right: crumpling a piece of aluminium foil modifies its BRDF.



Figure 5.5: Deforming the shape of a droplet through electro-wetting: the contact angle θ_c decreases as a voltage is applied to the electrodes (grey). The lower electrode and the droplet form a capacitor that strives to maximise its surface in order to minimise potential energy. In competition with the surface tension of the droplet, an equilibrium is reached.

Redistribution. Since real BRDFs span only a small subspace of all possible 4D distributions, general 4D modulation may not be necessary. A different, more natural approach to our problem is to redistribute the available light dynamically. Inspired by how real materials function on a microscopic level, a desired reflectance distribution can be obtained through a controlled scattering process [Weyrich09]. There are countless (static and dynamic) ways of achieving this, and the principle by itself does not collide with any of our criteria. For





Figure 5.6: Three lab samples of bistable cholesteric liquid crystal panels in reflection (top) and transmission (bottom). The green panel has one row of pixels set to transparent. Similar technologies could be used for additional colour modulation in a multi-layer BRDF display.

instance, instead of milling microgeometry, a dynamic version of the same idea could be implemented using a gimbal-mounted analogue mirror [Texas Instruments].

Mechanisms such as electro-wetting (Figure 5.5, [de Gennes02]) could be used to modulate the shape of droplets scattered on a substrate, changing the way they refract and reflect light.

What appears particularly interesting in this context is the idea of stacking functional layers to extend the space of achievable BRDFs. In Figure 5.4, we show how an optical element with a switchable transmittance profile can be turned into a reflectance display by adding a mirroring layer. Similarly, colour modulation could be added through layers of liquid crystal displays (Figure 5.6) or electronic paper.

Our own design, which will be covered in the rest of this chapter, belongs to the *redis-tribution* family of BRDF displays. We use a liquid surface as the reflecting base geometry, and reshape it over space and time by inducing surface waves. The prototypes meet Criteria 1–6 and 8, and further miniaturisation may allow for 7 and 9 as well. In Section 5.5.3, we present our device in a multi-layer operation mode allowing diffuse colour modulation.

5.3 Characterisation of Reflectance by Surface Waves

The operation principle of our device is to excite surface waves in a medium that supports relatively free travel of these waves. For our experiments, we use an interface between air and water.

We are not aiming at producing standing waves that would generate oscillating, yet



stationary, microgeometry. Instead we rely on *time-averaging* of travelling waves. If the generated height field varies fast enough, this results in an impression of a static microfacet distribution at every surface point. Mathematically, this is akin to averaging over a static height field of infinite extent.

Traditionally, microfacet-based models have been expressing reflectance as the product of a *probability density function* (PDF) of reflection directions, and additional material, geometry and normalisation terms (see e.g. Eq. 11 in [Torrance67]). As a physical surface and Fresnel reflector, our device naturally takes care of all of these, so we can focus on controlling the PDF.

In the following, we discuss the PDF generated by a single sine wave of small amplitude. For small angles ϕ , e.g.

$$\phi < 5^{\circ} \approx 0.0873 \,\mathrm{rad},\tag{5.1}$$

the trigonometric functions can be approximated to an error of less than 0.4% by the first term of their Taylor series:

$$\sin(\phi) \approx \tan(\phi) \approx \phi_{(\text{rad})} \quad \text{and} \quad \cos(\phi) \approx 1.$$
 (5.2)

Also, we can safely assume that interreflections are absent.

5.3.1 Single Sine Wave in One Dimension



Figure 5.7: On a half-wave height field $h(x) = a \sin(kx)$ (red), an incident ray (blue) with an angle β to the vertical (green) is reflected.

Assume a sine wave in *x* direction as depicted in Figure 5.7:

$$h(x) = a\sin(kx)$$
, where $k = \frac{2\pi}{\lambda}$, (5.3)

a being the amplitude, *k* the wave number and λ the wavelength of the excited wave. The angle $\alpha(x)$ of the surface normal at position *x*, measured in a mathematically positive sense with respect to the vertical axis, is related to the slope of the function as follows:

$$\tan \alpha(x) = h'(x) = \frac{\mathrm{d}h}{\mathrm{d}x} = ak\cos(kx). \tag{5.4}$$

Eq. 5.1 is met if the roughness $ak =: \rho$ is small. Then, a light ray incident at (x, h(x)) under a fixed angle β to the vertical will be mirrored into the reflection angle δ :

$$\delta(x) = 2\alpha(x) - \beta \overset{(2)(4)}{\approx} 2\rho \cos(kx) - \beta$$
(5.5)



The PDF $f_{\hat{\Delta}_{\beta}}(\delta)$, i.e., the likelihood of a ray in an *ensemble* of rays incident under the angle β to be reflected into the angle δ , can now be imagined as the limit case of a value histogram of $\delta(x)$: the probability for δ to lie within an infinitesimal interval $d\delta$ is the combined measure of all intervals dx which are mapped to $d\delta$.

For symmetry reasons, it is sufficient to look at the first half-wave of our height field $(x \in [0, \pi/k])$, where $\delta(x)$ is monotonically decreasing and therefore bijective, so that $f_{\hat{\Delta}_{\beta}}(\delta)$ can be obtained as the derivative of the inverse of $\delta(x)$, normalised by the measure of the interval:

$$f_{\hat{\Delta}\beta}(\delta) = \frac{k}{\pi} \left| \frac{\mathrm{d}x(\delta)}{\mathrm{d}\delta} \right| \approx \frac{1}{\pi} \left| \frac{\mathrm{d}}{\mathrm{d}\delta} \cos^{-1} \left(\frac{\delta + \beta}{2\rho} \right) \right|$$
$$= \frac{1}{\pi\sqrt{4\rho^2 - (\delta + \beta)^2}} \tag{5.6}$$

A plot of this function can be seen in Figure 5.8 (blue curve). Without the assumption of Eqs. 5.2, the PDF turns out slightly bulkier. Please refer to Appendix A for a full derivation.

5.3.2 Multiple Sine Waves in One or Two Dimensions



Figure 5.8: PDFs for sum-of-sinusoid functions in the linear limit (small amplitude, normal incidence, $\rho = 0.5$). As more sinusoidal terms are added up, the distribution converges to a Gaussian profile.

When we shoot an ensemble of rays at our height field, the exact shape of the surface is not of importance. In fact, we can treat the orientation of microfacets as a random variable that follows a probability distribution of Eq. 5.6. As *n* sinusoidal terms are superimposed in



our device, this corresponds to an addition of random variables $\hat{\Delta}^1_{\beta} + \dots + \hat{\Delta}^n_{\beta}$. If we can ensure that the $\hat{\Delta}^i_{\beta}$ are *independent and identically distributed* (iid), the central limit theorem states that the PDF of the their sum is the *n*-fold convolution of the individual PDFs, and that it approaches a Gaussian distribution for a large number of terms [Grimmett01]. In practice, we observe a satisfactory bell shape already for 5 superimposed sinusoidal waves, see Fig. 5.8.

Note that the distribution $f_{\hat{\Delta}_{\beta}}$ only depends on ρ , i.e. the *product* of amplitude and wave number. Since the wavenumber is directly related to the excitation frequency through the dispersion relation of the medium, we can generate identical distributions (same ρ) using different combinations of wave number and amplitude.

Independence, in our setting, translates to a vertical motion of every point on the dynamic height field surface that is as non-repetitive as possible. We choose excitation frequencies that relate like large prime numbers to approximate this.

The variance of a superposition of *n* identical distributions is related to the roughness as $\sigma_n^2 = n\sigma_0^2 = n \cdot 2\rho^2$, where $\sigma_0^2 = 2\rho^2$ is the variance of the single-sine distribution (Eq. 5.6). Within the linear approximation, we can thus generate Gaussian reflectance profiles of a desired variance by scaling all amplitudes of the sinusoidal terms uniformly.

In our device we are using orthogonally travelling planar waves in x- and y-direction, respectively. At sufficiently small ρ values, the waves decouple and the two-dimensional PDF of the reflection directions simply becomes the product of two distributions $f_{\hat{\Delta}_{\beta}}^{x}$ and $f_{\hat{\Delta}_{\gamma}}^{y}$, both of the form in Eq. 5.6.

5.3.3 Connection to Analytical BRDF Models

Since the variances on the *x*- and *y*-axes can be chosen independently, our display can reproduce BRDFs with elliptical highlights as popularised by Ward [Ward92]. Note that our variance σ_n^2 describes the reflection angle whereas Ward's $\alpha_{x,y}^2$ is related to the half-angle. Hence, the distributions are comparable when $\sigma_{x,y}^2 \approx 4\alpha_{x,y}^2$.

While the surface waves modulate the specular or glossy part of these models, the diffuse term can be realised by using a ground plane made of Labsphere Spectralon, an almost Lambertian reflector. The colour of the diffuse reflection can be influenced by a transmissive filter; we demonstrate this by dyeing the water prior to modulating the water surface.

In conclusion, our device is capable of displaying micro-facet BRDF models with anisotropic Gaussian microfacet distribution. The parameters of the model can be directly related to the parameters that control our device.





Figure 5.9: Various views of Setup 2 built from a pair of discarded 2.5-inch hard disk drives. Clockwise from top left: Components, water surface and checkerboard target, close-up on actuator 2, top view.

5.4 Construction of Devices

We built two incarnations of the same principle on different scales. Both setups consist of a flat water surface on which a pair of actuators excites crossing planar waves. We use voice coils that are fed with an amplified audio signal from the computer. Our signal source is the free software Puredata [Puckette] running a patch that synthesises a stereo signal from sine wave terms of different frequency and amplitude.

Setup 1 (large, slow) consists of the ripple tank system WA-9897 by Pasco, Inc., a device designed for demonstration experiments in physics classes. We use a pair of ripple generators (each with a bar-shaped lever and modified to accept audio input) on a flat water tank with wedge-shaped soft foam beaches at the borders to suppress reflections, and a surface of approximately 23×23 cm².

Setup 2 (small, fast) is a downscaled version built from two 2.5-inch hard disk drives (Figure 5.9). The platters and the controller boards as well as part of the aluminium frames were removed, leaving only the arm assemblies in place. To each arm we attached a small bar-shaped piece of plastic to dip into the water, and mounted both frames crosswise. A small water receptacle (approx. $2 \times 2 \text{ cm}^2$ in size) is placed underneath the actuators.



5.5 Results



Figure 5.10: Resonance profile of one of the wave generators in Setup 1, as observed through deflection of a laser beam. Trajectories were recorded in 1 Hz steps.

5.5.1 Characterisation

By deflecting a laser beam, we characterised the achieved surface normal variation for the "slow" Setup 1. The response of the actuator and its coupling to the water surface varies with the frequency of the signal. In Figure 5.10, we see that the efficiency peak is located around 20 Hz when the system is in contact with water. The resonance frequency of the uncoupled actuators is approximately 40 Hz. In Figure 5.11, we investigate the 1-dimensional case as described in Section 5.3.1. The driving signal is a sum of sinusoids with different amplitudes and frequencies. The amplitudes of each term were individually adjusted to yield a constant roughness parameter ρ on the water surface by adjusting the peak deflection angle of a laser beam. The resulting brightness profiles agree surprisingly well with the theoretical predictions.

For two-dimensional excitation, we observed the momentary shape of the reflection pattern, as well as its temporal average, which approximates an elliptical Gaussian (Figure 5.12). The corresponding synthesis patch is shown in Figure 5.13.

Due to its smaller size, Setup 2 responds considerably faster than Setup 1. We managed to excite water waves at frequencies as high as 800 Hz, although viscous damping limits the reach of such high-frequent waves to a few millimetres or centimetres. We found the setup to be most efficient for frequencies around 200 Hz, with deflection angles of up to 30° for single sine waves, or $\rho \approx 0.13$. If distortions of the trajectory can be tolerated, deflection angles of 50° ($\rho \approx 0.20$) are achievable. Above that, the actuators become unstable but droplet formation is not observed even for much higher amplitudes. As we relate our deflection angle to Ward's BRDF model through the variance (Section 5.3.2), we obtain a range of $0 \le \alpha_{x,y} < 0.14$, which overlaps with the values measured by Ward ($0.04 \le \alpha_{x,y} \le 0.26$).



Figure 5.11: We verified the insights from Section 5.3.1 by superimposing sinusoidal waves on the water surface of Setup 1 and observing the deflection of a laser beam. Top left: idealised temporal profile of excited wave; top right: photos of deflected beam; bottom: intensity profiles. Note the similarity to Figure 5.8.





Figure 5.12: Top: Momentary snapshots of a laser beam deflected from a surface excited in x- and y-direction, each by 5 sine terms. Averaging the trajectory over time results in a near-Gaussian elliptical profile (bottom). The synthesis patch used to generate the signal is shown in Figure 5.13.



Figure 5.13: Puredata patch for additive synthesis using 5 sinusoidal terms per audio channel. Note the choice of prime frequency ratios (red). The amplitudes (blue) were adjusted in order to achieve a constant peak deflection angle per term.



Figure 5.14: Photos of Setup 1 (left, checkerboard pitch 10 mm) and raytraced simulations (right) of various degrees of anisotropic blur.

5.5.2 Reflectance

In Figure 5.14, we compare the reflectance of our Setup 1 reflecting a checkerboard pattern with a raytraced simulation of a comparable setting. The photos are long-exposure shots, which, given the slow response of the device, were required in order to achieve satisfactory temporal averaging. The anisotropic blur of the reflection is very similar in nature to the simulated result.

The smaller Setup 2 allows for comparable results at a much faster speed, enabling the observer (or a video camera) to directly perceive the angular blur. Figure 5.15 shows four representative frames from a live video recording where the amplitudes in X and Y direction are manually adjusted in real time.

Setup 2 shows that miniaturisation brings a lot of benefits, since the achievable frequencies are approximately reciprocal to the geometric scale. Scaling down the setup by another order of magnitude may for instance enable temporal multiplexing of different lobes.

5.5.3 Diffuse + Specular

Also using Setup 2, we placed a diffuse reflector underneath the water surface and dyed the liquid. As can be seen in Figure 5.16, this adds greatly to the range of displayable BRDFs, however in our case the shallow water layer leads to an increase of viscous damping.





Figure 5.15: Real-time blur as displayed by Setup 2. The checkerboard scale is 1 mm.

5.5.4 Liquid Metal

In order to obtain a higher reflectivity especially for near-normal directions, we replaced the water in Setup 2 with an eutectic alloy of gallium, indium and tin. The nontoxic substance has a melting point of -18° C and is therefore liquid at room temperature. The increased viscosity, mass density, strong surface tension and incessant formation of an oxide layer make it difficult to control the surface shape. In particular, planar waves are rather hard to obtain even at very small amplitudes. However, we can still control a slight variation in the reflectance (Figure 5.17).

5.6 Discussion

The limitations of our device can be broadly classified into two categories: practical issues of our prototypical implementation, and principal limitations of the general design. In the following we discuss the major practical limitations of our prototype and suggest alternative ways of implementation.

1. Our display exhibits a limited range of surface roughness. With higher amplitudes, the nonlinearities in our physical system become hard to predict. By investigating the dominating effects and modelling them in the predictive model, it could be possible to enhance the scattering capabilities of the device beyond the 30° we achieved





Figure 5.16: Macrophotos of Setup 2 (checker size 1 mm). By adding a diffuse white substrate and injecting coloured inks (here done manually), we obtain a combination of an anisotropic specular and a diffuse lobe.



5.6. DISCUSSION



Figure 5.17: Reflection of a human eye in Setup 2b filled with liquid metal. Left: resting. Right: in motion. The distortions are caused by surface tension.

in our experiments.

- 2. We are currently limited to BRDFs consisting of a single white Gaussian lobe in conjunction with a fixed coloured diffuse component. Further miniaturisation, e.g. using piezos as actuators, might enable temporal multiplexing of multiple lobes. The ink-based colouring layer currently used to model the diffuse component can be replaced by a passive transflective display panel.
- 3. The Fresnel reflection factor is currently the one of water. By adding refractive index altering agents to the medium in our device, the reflectivity could approach solid materials like plastics, resins or coatings.
- 4. Currently, we are only using two orthogonally crossing planar waves to generate BRDFs separable in the two dimensions. A more general implementation of the principle could make use of multiple point-like actuators, generating spherical waves, implementing Huygens principle. This way, more general reflectance distributions could be generated, for instance a rotation of the tangent frame, which is currently missing in our design.
- 5. The use of liquids constrains us to horizontal mounting. Exchanging the water for solid jelly-like substances or elastic films might allow for arbitrary mounting angles. This may require the theoretical model to account for effects such as internal stresses and strains.

The only principal limitation we are aware of is the lack of the missing possibility of varying the normal of the simulated macro-surface, i.e. Criterion 10 is most likely not achievable with the proposed system.

In conclusion, this chapter has introduced the concept of a reflectance display. We have discussed theoretical requirements, advantages and drawbacks of different potential implementations. We have characterised a promising design both theoretically and practically by

CHAPTER 5. DYNAMIC BRDF DISPLAY

building two prototypes at different scales, and established an analytic link between Wardlike anisotropic BRDFs and the class of BRDFs displayable by our device. In experiment, we have verified that dynamic time-averaged micro-facet distributions can give the impression of dynamically changing, programmable reflectance in real time. Our prototype meets most of the requirements for a BRDF display and offers room for many extensions. We are confident that this presents a first step towards future hyper-realistic displays. However, much work remains to be done.



Chapter 6 Conclusion

In this thesis, we have questioned common assumptions about the nature of light transport and shown up their failure cases. By making creative use of unusual optical phenomena, we were able to make existing techniques work with uncollaborative materials, yielding results of unprecedented quality. This comprises solutions for a choice of problems in the fields of geometry and reflectance acquisition and reproduction, as well as the conceptual introduction of a BRDF display.

We consider the following to be the three main contributions of this work. First, we used the phenomenon of fluorescence as an optical tool for the acquisition of geometry. We modified a 3D scanning system to deliver robust measurements of objects made from almost any material by placing the object in a fluorescent volume. Instead of relying on surface reflections, our new 3D scanning technique is based on the observation of light transport in front of the surface. Since accurate appearance models are just as important for the generation of photorealistic imagery as is the geometry of an object, we then investigated the appearance of surfaces that are themselves fluorescent. We introduced a setup and measurement scheme for the practical acquisition of reflectance and reradiation distribution functions of fluorescent surfaces, allowing the data-driven rendering of measured fluorescent materials. Finally, we introduced the idea of displaying materials instead of colours, defined the problem of BRDF display, and presented two prototypes that can display a range of anisotropic reflectance distributions.

Common to all of the above problems and our solutions is the combination and interplay of elements from disciplines as diverse as physics, chemistry, computer science, mathematics and engineering. The choice of phenomena and techniques may appear unusual from a computer science point of view, but all of them were long-established and well-understood in their respective disciplines. Now, add to this all the technologies that are still emerging: organic semiconductors, ultrafast optics and electronics, photonic crystals, nano- and biotechnology, micromechanics and microfluidics, cognitive science and artificial intelligence, and robotics, just to name a few. In optics alone, there are countless unusual phenomena that have received little attention so far. While some might be interesting to simulate in a rendering framework, others could be exploited in novel ways to solve

CHAPTER 6. CONCLUSION

seemingly unrelated problems. We are convinced that through "out-of-the-box" thinking and interdisciplinary exchange there is much more to be gained in computer graphics and beyond. The techniques presented in this work, taken by themselves, may be of limited reach and applicability, but we hope for the ideas to serve as a source of inspiration to other researchers.



Appendix A Reflectance of Sinusoidal Height Fields

In Chapter 5, we derived the reflectance distribution for the sine-shaped height fields generated by our BRDF display prototypes. We assumed a small amplitude because it would allow us to approximate the trigonometric functions by their first Taylor expansion terms. For completeness, the full derivation without approximations is provided in the following.



Figure A.1: On a half-wave height field $h(x) = a \sin(kx)$ (red), an incident ray (blue) with an angle β to the vertical (green) is reflected.

A.1 Reflection from a Sine Wave

Assume the following setup: light is illuminating a one-dimensional, sinusoidal height field vertically. We want to derive the distribution of the reflection directions assuming a perfectly mirroring surface.

Let the height field be given by

$$h(x) = a\sin(kx)$$
, where $k = \frac{2\pi}{\lambda}$.

For simplicity, we restrict the discussion to one half-wave, i.e. $x \in [0, \lambda/2]$, as depicted in Fig. A.1. The other half-wave can be treated similarly and yields an identical distribution. For vertically incident lighting, the reflection direction of the incident rays equals 2α . The angle α can be obtained by the relation $\tan \alpha = h'(x)$, i.e. the reflection directions are given by

$$\delta(x) = 2\tan^{-1}h'(x). \tag{A.1}$$

Slightly more general, let the illumination be incident at an angle β , where α and β are angles with the vertical taken in the mathematically positive sense, then the reflection angle is given by

$$\delta_{\beta}(x) = \beta + 2(\alpha - \beta) = 2\alpha - \beta = 2\tan^{-1}h'(x) - \beta.$$
(A.2)

Now, consider the wave to generate a micro-facet distribution. In our experiments we achieve this situation by time-averaging over the waves travelling by. As seen from Eq. A.1, the reflection directions directly depend on the derivative of the height field h'(x). Thus, knowledge of the distribution of the height field derivatives yields the desired distribution of reflection directions, i.e. the BRDF caused by the height field.

This can be achieved by considering a random process which describes photons as they hit the height field at random positions. (Equally well the process can be seen as describing photons at a fixed position hitting a height field moving at a constant speed randomly in time.) The problem then becomes to compute the probability density of a function of a random variable. The initial photon distribution is assumed to be the unit distribution with density function

$$f_X(X) = \begin{cases} 2/\lambda & X \in [0, \lambda/2] \\ 0 & \text{else} \end{cases},$$
(A.3)

to which Eq. A.1 is applied to obtain the directional distribution. The choice of a unit distribution is valid if the amplitude of the wave is small such that the micro-facets can be considered as lying in the plane of the macro-surface, only exhibiting the orientations of a sine wave. From probability theory [Grimmett01] the probability density of the random variable Y generated by applying a function g to another random variable X is

$$f_Y(y) = f_X(g^{-1}) \cdot \left| \frac{dg^{-1}}{dy}(y) \right|.$$
 (A.4)

In our case this yields

$$f_{\Delta}(\delta) = \frac{2}{\lambda} \left| \frac{\mathrm{d}\{\delta(x)\}^{-1}(\delta)}{\mathrm{d}\delta} \right|,\tag{A.5}$$

where $\{\delta(x)\}^{-1}(\delta)$ denotes the inverse of the direction function, Eq. A.1 or A.2, and f_{Δ} is the probability density of the reflection directions,

$$\{\delta(x)\}^{-1}(\delta) = \frac{\cos^{-1}\left(\frac{\tan(\delta/2)}{ak}\right)}{k},\tag{A.6}$$



and

$$\frac{\mathrm{d}\{\delta(x)\}^{-1}}{\mathrm{d}\delta}(\delta) = -\frac{1}{2} \frac{1 + \tan(\delta/2)^2}{ak^2 \left(1 - \frac{\tan(\delta/2)^2}{a^2k^2}\right)^{\frac{1}{2}}}.$$
 (A.7)

Thus, the probability density function for the reflection direction due to a sine wave is given by:

$$f_{\Delta}(\delta) = \frac{1}{2\pi} \frac{1 + \tan(\delta/2)^2}{ak \left(1 - \frac{\tan(\delta/2)^2}{(ak)^2}\right)^{\frac{1}{2}}},$$
(A.8)

where the absolute value $|\cdot|$ has been dropped because $\frac{d\{\delta(x)\}^{-1}}{d\delta}(\delta)$ is strictly negative and $\lambda = 2\pi/k$ has been used. As we can see, amplitude and frequency enter the equation as a common factor $\rho = ak = 2\pi a/\lambda$. Thus, amplitude and wavelength are directly proportional: halving or doubling both leaves the same directional PDF, i.e. the BRDF is invariant to those changes.

The corresponding PDF for Eq. A.2 is

$$f_{\Delta_{\beta}}(\delta) = \frac{1}{2\pi} \frac{1 + \tan((\delta + \beta)/2)^2}{\sqrt{\rho^2 - \tan((\delta + \beta)/2)^2}}.$$
 (A.9)

We see that changing the incident angle β only results in a shift of the distribution. Its shape remains unaltered.

A.2 Linearisation, Mean and Variance of the Directional Distribution

If the slope of the waves is reasonably small, we can assume that $\tan^{-1}h'(x) \approx h'(x)$. For the sine wave the maximum slope can be controlled by the amplitude *a* and the wavenumber *k* since

$$h'(x) = ka\cos(kx) \in [-ka, ka].$$

The requirements for linearisation of the inverse tangent function can thus be met experimentally. There are additional sources of non-linearities if the wave amplitude becomes large. First, the source distribution of rays, Eq. A.3, ceases to be a unit distribution. This assumption is only valid if the height of intersection of the ray with the height field surface is negligible. Second, the mean and variance of $f_{\Delta_{\beta}}$, Eq. A.9, are not analytically computable. Thus, a direct link between σ_0 (ρ) and σ_n can only be tabulated (Another calibration step, linking applied voltage to resulting wave amplitude introduces another scale factor). Third, results from probability theory as used in Sect. A.3 are only valid in the case of linearised angles. For these reasons we choose to work with the angle-linearised version of the PDF. Essentially, this implies that we are working with slopes instead of angles. The angle-linearised version of Eq. A.9 is

$$f_{\hat{\Delta}_{\beta}} = \frac{1}{\pi \sqrt{4\rho^2 - (\delta + \beta)^2}},$$
 (A.10)

which is analytically integrable. In particular, the mean and the variance of $f_{\hat{\Delta}_{\beta}}$ can be computed. As expected, the mean is found at $-\beta$ ($\mu_0 = 0$ at $\beta = 0$ and substitution of variables), i.e. the main reflection direction agrees with the law of specular reflection. For the variance we consider the case $\beta = 0$. This case is sufficient since the PDF, Eq. A.10 is shift-invariant. The variance can be calculated by

$$\sigma_{0}^{2} = \int_{-2\rho}^{2\rho} \delta^{2} \cdot f_{\hat{\Delta}_{0}} d\delta$$

$$= \int_{-2\rho}^{2\rho} \frac{\delta^{2}}{\pi \sqrt{4\rho^{2} - \delta^{2}}} d\delta$$

$$= \left[\frac{1}{2\pi} \cdot \left(4\rho^{2} \tan^{-1} \left(\frac{\delta}{\sqrt{4\rho^{2} - \delta^{2}}} \right) - \delta \sqrt{4\rho^{2} - \delta^{2}} \right) \right]_{-2\rho}^{2\rho}$$

$$= \left[\frac{4\rho^{2}}{2\pi} \cdot 4\rho^{2} \tan^{-1} \left(\frac{\delta}{\sqrt{4\rho^{2} - \delta^{2}}} \right) \right]_{-2\rho}^{2\rho}$$
(A.11)

where the integration bounds are due to the domain of the inverse tangent. To compute the integral we have to take limits from the left at the right boundary and from the right at the left boundary of the interval $[-2\rho \dots 2\rho]$.

$$\lim_{\delta \to \pm 2\rho} \frac{4\rho^2}{2\pi} \tan^{-1} \left(\frac{\delta}{\sqrt{4\rho^2 - \delta^2}} \right) = \pm \rho^2$$
(A.12)

The variance of the angle-linearised directional distribution is thus given by $\sigma_0^2 = 2 \cdot \rho^2$.

A.3 Superposition of Sine Waves

In our reflectance display we use several super-positioned, statistically independent, sinewaves. Since we are working in the angle-linearised regime, by super-positioning of the




Figure A.2: Exemplary height field generated from the 5-term Puredata patch in Figure 5.13.

waves we achieve an addition of random variables. As is well known from probability theory, the probability density of the sum of two i.i.d. random variables is simply the convolution of their respective probability densities [Grimmett01]:

$$f_{X+Y} = f_X \otimes f_Y. \tag{A.13}$$

Now, for the addition of *n* different, i.i.d. random variables with the same distribution f_{Δ} , the PDF becomes, in the limit as $n \to \infty$, a Gaussian distribution (central limit theorem).

For our purposes the super-position of 5 to 7 sine waves is sufficient to achieve a sufficiently Gaussian distribution. The variance σ_n^2 of the distribution after *n* convolutions is linked to the variance σ_0^2 of the initial distribution f_{Δ} by

$$\sigma_n^2 = n \cdot \sigma_0^2. \tag{A.14}$$

Therefore, we can establish a link between the original distribution f_{Δ} , Eq. A.10, dependent on $\rho = ak$ and the variance σ_n^2 of a Gaussian-like distribution. By choosing a fixed number *n* of super-positioned, statistically independent, sinusoids, this enables us to compute proper amplitude-wavelength combinations $\rho = a_i k_i$ $i = 1 \dots n$ to achieve a Gaussian distribution with variance σ_n as used in many standard BRDF models. Note that ρ is constant for any particular choice of targeted Gaussian variance σ_n . The combinations $\rho = a_i k_i$ are chosen such that the resulting *n* sinusoids are statistically independent. An example height field is shown in Figure A.2.

A.4 Fresnel Reflection

Up to now we have dealt with perfectly mirroring materials. We have derived the probability density function for the distribution of the reflection directions. If the reflecting material is a dielectric, both reflection and transmission occur when photons interact with the material surface. The ratios of transitted and reflected light are described by the Fresnel equations



Figure A.3: Fresnel reflection for unpolarised light in air incident on a water surface.

$$r_{s} = \frac{(n_{i}\cos\theta_{i} - n_{t}\cos\theta_{t})}{(n_{i}\cos\theta_{i} + n_{t}\cos\theta_{t})}$$
(A.15)

$$r_p = \frac{(n_t \cos \theta_i - n_i \cos \theta_t)}{(n_t \cos \theta_i + n_i \cos \theta_t)}$$
(A.16)

$$F_r = \frac{1}{2} \left(r_s^2 + r_p^2 \right)$$
 (A.17)

$$F_t = 1 - F_r, \tag{A.18}$$

where F_r is the Fresnel reflectance for a particular incident light direction $\cos \theta_i$ with respect to the surface normal, which in our setting equals to $\alpha - \beta$. $\cos \theta_t$ is the corresponding angle for the transmitted light direction and n_i and n_t are the refractive index of the incident (air) and transmissive (water) material, respectively. A plot of the Fresnel reflectivity for an air-to-water interface can be seen in Figure A.3.

A.5 Relation to Synthetic BRDF Models

Since we model the reflection directions as random variables, the Fresnel reflection factor simply multiplies the reflected intensity *physically*. This is consistent with existing synthetic



BRDF models for non-diffuse dielectrics.

Additionally, these models usually contain a geometry term modelling self-occlusion of the surface, typically by assuming some very coarse geometry model like V-shaped grooves. Our surfaces have a more complex structure and thus the agreement between computational rendering by a computer and physical rendering by our device might not be perfect. However, since we are generating semi-repetitive structures, see Fig. A.2 for an example, the coarse form of our surfaces resembles the general structure of V-shaped grooves and thus the results are comparable.

A.6 Two-Dimensional Height Fields

In the realisation of our device we are using surface waves that are travelling orthogonally. Our height field is thus of the form

$$h(x, y) = a\sin(k_x x) + b\sin(k_y y). \tag{A.19}$$

Again, we have to derive the reflection direction function. The normal to a height field is given by

$$\vec{n} = \begin{pmatrix} 1\\ 0\\ \frac{dh}{dx} \end{pmatrix} \times \begin{pmatrix} 0\\ 1\\ \frac{dh}{dy} \end{pmatrix} = \begin{pmatrix} -\frac{dh}{dx}\\ -\frac{dh}{dy}\\ 1 \end{pmatrix}.$$
 (A.20)

Assume that the incident light direction is given by \vec{l} , $|\vec{l}| = 1$, then the reflection direction is given by

$$\vec{\delta}(\vec{l}) = \vec{l} + 2\vec{n} \cdot \vec{l}(\vec{n} - \vec{l}).$$
 (A.21)

We see that the *x*- and *y*-components of the reflection direction are influenced by only dh/dx and dh/dy, respectively. The equations have the same structure as Eq. A.2 and thus the two-dimensional case is reduced to the one-dimensional one discussed before.





Bibliography

[Ajdin08] BORIS AJDIN, MATTHIAS B. HULLIN, CHRISTIAN FUCHS, HANS-PETER SEIDEL, AND HENDRIK P. A. LENSCH. Demosaicing by smoothing along 1D features. In Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2008), 2008. [Alexa10] MARC ALEXA AND WOJCIECH MATUSIK. Reliefs as images. ACM Trans. Graph. (Proc. ACM SIGGRAPH), 2010. [Angelopoulou01] Elli Angelopoulou and Rana Molana. Multispectral Skin Color Modeling. In IEEE Conference on Computer Vision and Pattern Recognition, pages 635-642, 2001. [Baba01] Mitsuru Baba, Kozo Ohtani, Makoto Imai, and Tadataka KONISHI. New Laser Rangefinder for Three-Dimensional Shape Measurement of Specular Objects. Optical Engineering, 40(1):53-60, January 2001. [Baccini96] A. BACCINI, PH. BESSE, AND A. DE FALGUEROLLES. An L1-norm PCA and a heuristic approach. In Ordinal and Symbolic Data Analysis, pages 359-368. Springer, 1996. [Banks07] DAVID C. BANKS AND LAITH ABU-RADDAD. The foundations of photo-realistic rendering: from quantum electrodynamics to Maxwell's equations. In GVE '07: Proceedings of the IASTED International Conference on Graphics and Visualization in Engineering, pages 137-144, Anaheim, CA, USA, 2007. ACTA Press. [Ben-Ezra03] M. BEN-EZRA AND S.K. NAYAR. What Does Motion Reveal About Transparency? In *Proc. of ICCV'03*, volume 2, pages 1025–1032, 2003. [Bendig08] MARION BENDIG, JOHANNES HANIKA, HOLGER DAMMERTZ, JAN CHRISTOPH GOLDSCHMIDT, MARIUS PETERS, AND MICHAEL Simulation of Fluorescent Concentrators. WEBER. In IEEE Symposium on Interactive Ray Tracing, pages 93–98, August 2008.

- [Beraldin04] J.-A. BERALDIN. Integration of Laser Scanning and Close-Range Photogrammetry - the Last Decade and Beyond. In Proceedings of the XXth ISPRS Congress, pages 972-983, 2004. [Bhat95] D. BHAT AND S NAYAR. Stereo in the Presence of Specular Reflection. In International Conference on Computer Vision (ICCV), pages 1086-1092, 1995. [Bickel10] Bernd Bickel, MORITZ BÄCHER, Miguel A. Otaduy, HYUNHO RICHARD LEE, HANSPETER PFISTER, MARKUS GROSS, AND WOJCIECH MATUSIK. Design and fabrication of materials with desired deformation behavior. ACM Trans. Graph. (Proc. ACM SIGGRAPH), 2010. [Blais04] F. BLAIS. Review of 20 Years of Range Sensor Development. Journal of Electronic Imaging, 13(1):231-243, 2004. [Blinn77] JAMES F. BLINN. Models of light reflection for computer synthesized pictures. Proc. ACM SIGGRAPH, 1977. [Bonfort03] THOMAS BONFORT AND PETER STURM. Voxel Carving for Specular Surfaces. In Proc. of ICCV'03, pages 591–596, 2003. [Bonfort06] THOMAS BONFORT, PETER STURM, AND PAU GARGALLO. General Specular Surface Triangulation. In Proceedings of the Asian Conference on Computer Vision, volume 2, pages 872-881, jan 2006. [Born99] MAX BORN AND EMIL WOLF. Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light (7th Edition). Cambridge University Press, 7th edition, October 1999. [Chen00] FRANK CHEN, GORDON M. BROWN, AND MUMIN SONG. Overview of Three-Dimensional Shape Measurement using Optical Methods. Optical Engineering, 39(1):10-22, January 2000. [Chen06] TONGBO CHEN, MICHAEL GOESELE, AND HANS-PETER SEIDEL. Mesostructure from Specularity. In Proc. of CVPR '06, pages 17-22, 2006. [Chen07] TONGBO CHEN, HENDRIK P. A. LENSCH, CHRISTIAN FUCHS, AND HANS-PETER SEIDEL. Polarization and Phase-Shifting for 3D Scanning of Translucent Objects. In Proc. of CVPR '07, pages 1-8, 2007. [CIE32] COMMISION INTERNATIONALE DE L'ECLAIRAGE – CIE. Proceedings of
- [Clark97] J. CLARK, E. TRUCCO, AND L. B. WOLFF. Using Light Polarization in Laser Scanning. *Image and Vision Computing*, 15(1):107–117, 1997.

the 1931 Conference. Cambridge University Press, Cambridge, 1932.



[Clarke85]	F.J.J. CLARKE AND D.J. PARRY. Helmholtz Reciprocity: its validity and application to reflectometry. <i>Lighting Research and Technology</i> , 17(1):1–11, 1985.
[Coffin10]	D. Coffin. dcraw. http://www.cybercom.net/~dcoffin/dcraw/, 2010.
[Cook81]	ROBERT COOK AND KENNETH TORRANCE. A Reflectance Model for Computer Graphics. In <i>SIGGRAPH</i> , pages 307–316, 1981.
[Cossairt08]	O. COSSAIRT, S. K. NAYAR, AND R. RAMAMOORTHI. Light Field Trans- fer: Global Illumination Between Real and Synthetic Objects. <i>ACM</i> <i>Trans. Graph. (Proc. ACM SIGGRAPH)</i> , 2008.
[CUReT96]	CURET. Columbia Utrecht Texture Database. Web-Page. http://www1.cs.columbia.edu/CAVE/software/curet/index.php, 1996.
[Curless95]	BRIAN CURLESS AND MARC LEVOY. Better Optical Triangulation Through Spacetime Analysis. In <i>Proc. of ICCV'95</i> , pages 987–994, 1995.
[Davis05a]	J. DAVIS, D. NEHAB, R. RAMAMOORTHI, AND S. RUSINKIEWICZ. Spacetime stereo: a unifying framework for depth from triangula- tion. <i>IEEE Trans. PAMI</i> , 27(2):296–302, 2005.
[Davis05b]	J. DAVIS, R. YANG, AND L. WANG. BRDF Invariant Stereo using Light Transport Constancy. In <i>Proc. of ICCV'05</i> , pages 436–443, 2005.
[de Gennes02]	PIERRE-GILLES DE GENNES, FRANÇOISE BROCHARD-WYART, AND DAVID QUÉRÉ. <i>Gouttes, bulles, perles et ondes</i> . Belin, 2002. Also available in English: <i>Capillarity and wetting phenomena: Drops, bub- bles, pearls, waves</i> . Springer, 2004.
[Debevec98]	PAUL DEBEVEC. Rendering synthetic objects into real scenes. <i>Proc. ACM SIGGRAPH</i> , 1998.
[DeRose07]	PAUL C. DEROSE, EDWARD A. EARLY, AND GARY W. KRAMER. Qual- ification of a fluorescence spectrometer for measuring true fluores- cence spectra. <i>Review of Scientific Instruments</i> , 78, 2007.
[Deusch01]	S. DEUSCH AND T. DRACOS. Time resolved 3D passive scalar concentration-field imaging by laser induced fluorescence (LIF) in moving liquids. <i>Meas. Sci. Technol.</i> , (12):188–200, 2001.
[Donaldson54]	R. DONALDSON. Spectrophotometry of fluorescent pigments. <i>British Journal of Applied Physics</i> , 5(6):210–214, 1954.



[Dong10]	YUE DONG, JIAPING WANG, FABIO PELLACINI, XIN TONG, AND BAIN- ING GUO. Fabricating spatially-varying subsurface scattering. <i>ACM</i> <i>Trans. Graph. (Proc. ACM SIGGRAPH)</i> , 2010.
[Dutre02]	PHILIP DUTRE, KAVITA BALA, AND PHILIPPE BEKAERT. Advanced Global Illumination. A. K. Peters, Ltd., Natick, MA, USA, 2002.
[Einstein05]	ALBERT EINSTEIN. Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. <i>Annalen der Physik</i> , 322(6):132–148, 1905.
[Finckh10]	M. FINCKH, H. DAMMERTZ, AND H. P. A. LENSCH. Geometry Re- construction from Caustic Images. In <i>Proc. of ECCV 2010</i> , page to appear, 2010.
[Fuchs07]	CHRISTIAN FUCHS, TONGBO CHEN, MICHAEL GOESELE, HOLGER THEISEL, AND HANS-PETER SEIDEL. Density Estimation for Dynamic Volumes. <i>Computers & Graphics</i> , 31(2):205–211, April 2007.
[Fuchs08]	MARTIN FUCHS, RAMESH RASKAR, HANS-PETER SEIDEL, AND HEN- DRIK P. A. LENSCH. Towards passive 6D reflectance field displays. ACM Trans. Graph. (Proc. ACM SIGGRAPH), 2008.
[Gat00]	NAHUM GAT. Imaging Spectroscopy using Tunable Filters: A Review. In H. H. Szu, M. Vetterli, W. J. Campbell, and J. R. Buss, editors, <i>SPIE Conference Series</i> , volume 4056, pages 50–64, April 2000.
[Glassner94]	ANDREW GLASSNER. A Model for Fluorescence and Phosphores- cence. In <i>Proc. of the Eurographics Workshop on Rendering</i> , pages 57–68, June 1994.
[Goodman05]	JOSEPH W. GOODMAN. <i>Introduction to Fourier Optics</i> . Roberts & Company Publishers, 3 edition, December 2005.
[Graves07]	KYNAN GRAVES, ROMESH NAGARAJAH, AND PAUL R. STODDART. Analysis of Structured Highlight Stereo Imaging for Shape Measure- ment of Specular Objects. <i>Optical Engineering</i> , 46(8):published on- line, 30. August 2007, August 2007.
[Grimmett01]	GEOFFREY GRIMMETT AND DAVID STIRZAKER. "Probability and Ran- dom Processes". 2001.
[Gundlach94]	DIETRICH GUNDLACH AND HEINZ TERSTIEGE. Problems in mea- surement of fluorescent materials. <i>Color Research & Application</i> , 19(6):427–436, 1994.



[Halstead96]	MARK A. HALSTEAD, BRIAN A. BARSKY, STANLEY A. KLEIN, AND ROBERT B. MANDELL. Reconstructing Curved Surfaces From Spec- ular Reflection Patterns Using Spline Surface Fitting of Normals. In <i>Proc. of ACM SIGGRAPH 96</i> , pages 335–342, August 1996.
[Hardeberg99]	J. Y. HARDEBERG, F. J. SCHMITT, AND H. BRETTEL. Multispectral im- age capture using a tunable filter. In R. Eschbach & G. G. Marcu, edi- tor, <i>Society of Photo-Optical Instrumentation Engineers (SPIE) Confer-</i> <i>ence Series</i> , volume 3963 of <i>Presented at the Society of Photo-Optical</i> <i>Instrumentation Engineers (SPIE) Conference</i> , pages 77–88, Decem- ber 1999.
[Havran05]	VLASTIMIL HAVRAN, ATTILA NEUMANN, GEORG ZOTTI, WERNER PURGATHOFER, AND HANS-PETER SEIDEL. On cross-validation and resampling of BRDF data measurements. In SCCG '05: Proceedings of the 21st Spring Conference on Computer Graphics, pages 161–168, 2005.
[Hašan10]	MILOŠ HAŠAN, MARTIN FUCHS, WOJCIECH MATUSIK, HANSPETER PFISTER, AND SZYMON RUSINKIEWICZ. Physical reproduction of ma- terials with specified subsurface scattering. <i>ACM Trans. Graph. (Proc.</i> <i>ACM SIGGRAPH)</i> , 2010.
[Hawkins05]	TIM HAWKINS, PER EINARSSON, AND PAUL DEBEVEC. Acquisition of Time-Varying Participating Media. In <i>Proc. of ACM SIGGRAPH 2005</i> , pages 812–815. ACM, 2005.
[Hearnshaw90]	J. B. HEARNSHAW. <i>The Analysis of Starlight : One Hundred and Fifty Years of Astronomical Spectroscopy</i> . Cambridge University Press, Cambridge ; New York :, 1st Pbk. Ed. (with corrections) edition, 1990.
[Hecht01]	EUGENE HECHT. <i>Optics (4th Edition)</i> . Addison Wesley, 4 edition, August 2001.
[Hersch07]	ROGER D. HERSCH, PHILIPP DONZÉ, AND SYLVAIN CHOSSON. Color Images Visible under UV Light. ACM Trans. Graph. (Proc. SIG- GRAPH), 26(3):75, 2007.
[Hirsch09]	MATTHEW HIRSCH, DOUGLAS LANMAN, HENRY HOLTZMAN, AND RAMESH RASKAR. BiDi Screen: A Thin, Depth-Sensing LCD for 3D Interaction using Lights Fields. <i>ACM Trans. Graph. (Proc. ACM SIG-GRAPH Asia)</i> , 2009.
[Hirsch10]	MATTHEW HIRSCH AND DOUGLAS LANMAN. Build your own 3D dis- play. In ACM SIGGRAPH 2010 Courses, 2010.



[Höhle71]	J. HÖHLE. Reconstruction of the Underwater Object. <i>Photogrammet-</i> <i>ric Engineering</i> , 37:948–954, 1971.
[Holopainen08]	SILJA HOLOPAINEN, FARSHID MANOOCHERI, AND ERKII IKONEN. Go- niofluorometer for characterization of fluorescent materials. <i>Applied</i> <i>Optics</i> , 47(6):835–842, 2008.
[Hullin08a]	MATTHIAS B. HULLIN, MARTIN FUCHS, IVO IHRKE, BORIS AJDIN, HANS-PETER SEIDEL, AND HENDRIK P. A. LENSCH. Direct visual- ization of real-world light transport. In <i>Proceedings of VMV</i> , pages 363–371, 2008.
[Hullin08b]	MATTHIAS B. HULLIN, MARTIN FUCHS, IVO IHRKE, HANS-PETER SEI- DEL, AND HENDRIK P. A. LENSCH. Fluorescent immersion range scanning. <i>ACM Transactions on Graphics (Proceedings of SIGGRAPH</i> 2008), 27(3):87:1 – 87:10, 2008.
[Hullin08c]	MATTHIAS B. HULLIN, MARTIN FUCHS, IVO IHRKE, HANS-PETER SEIDEL, AND HENDRIK P. A. LENSCH. Virtual Scattering Lab. http://www.mpi-inf.mpg.de/resources/FIRS/virtuallab.html, 2008.
[Hullin10]	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. <i>ACM Transactions on Graphics (Proceedings</i> <i>of SIGGRAPH 2010)</i> , 29(3), 2010.
[Hullin11]	MATTHIAS B. HULLIN, HENDRIK P. A. LENSCH, RAMESH RASKAR, HANS-PETER SEIDEL, AND IVO IHRKE. Dynamic display of BRDFs. <i>Proceedings of Eurographics</i> , 2011. To appear.
[Ihrke05]	I. IHRKE, B. GOLDLUECKE, AND M. MAGNOR. Reconstructing the Geometry of Flowing Water. In <i>Proc. of ICCV'05</i> , pages 1055–1060, 2005.
[Imai98]	FRANCISCO H. IMAI, ROY S. BERNS, AND CHESTER F. CARLSON. High-Resolution Multi-Spectral Image Archives - A Hybrid Approach. pages 224–227, 1998.
[Ito10]	KATSUHISA ITO, HIROKI KIKUCHI, HISAO SAKURAI, IZUSHI KOBAYASHI, HIROAKI YASUNAGA, HIDENORI MORI, KAZUTATSU TOKUYAMA, HIROTAKA ISHIKAWA, KENGO HAYASAKA, AND HI- ROYUKI YANAGISAWA. Sony RayModeler: 360° Autostereoscopic Display. In <i>SIGGRAPH Emerging Technologies</i> , 2010.

[Jin05] H. JIN, S. SOATTO, AND A. J. YEZZI. Multi-View Stereo Reconstruction of Dense Shape and Complex Appearance. *International Journal of Computer Vision*, 63(3):175–189, Jul 2005.



[Jones41]	R. CLARK JONES. A New Calculus for the Treatment of Optical Systems. <i>J. Opt. Soc. Am.</i> , 31(7):488–493, 1941.
[Jones07]	ANDREW JONES, IAN MCDOWALL, HIDESHI YAMADA, MARK BOLAS, AND PAUL DEBEVEC. An interactive 360° light field display. In <i>SIG-GRAPH Emerging Technologies</i> , 2007.
[Kajiya84]	JAMES T. KAJIYA AND BRIAN P VON HERZEN. Ray tracing volume densities. <i>SIGGRAPH Comput. Graph.</i> , 18(3):165–174, 1984.
[Kajiya86]	J. KAJIYA. The Rendering Equation. In <i>Computer Graphics (Proc. of SIGGRAPH '86)</i> , pages 143–150, August 1986.
[Kanitsar02]	A. KANITSAR, T. THEUSSL, L. MROZ, M. ŠRÁMEK, A. VILANOVA BARTROLÍ, B. CSÉBFALVI, J. HLÅDUVKA, S. GUTHE, M. KNAPP, R. WE- GENKITTL, P. FELKEL, S. RÖTTGER, D. FLEISCHMANN, W. PURGATH- OFER, AND M. E. GRÖLLER. Christmas Tree Case Study: Computed Tomography as a Tool for Mastering Complex Real World Objects with Applications in Computer Graphics. In <i>Proc. of Visualization'02</i> , pages 489–492, 2002.
[Ke03]	QIFA KE AND TAKEO KANADE. Robust Subspace Computation Using L1 Norm. <i>Carnegie Mellon Tech Report CMU-CS-03-172</i> , 2003.
[Kim10]	YOUNGMIN KIM, KEEHOON HONG, AND BYOUNGHO LEE. Recent re- searches based on integral imaging display method. <i>3D Research</i> , 1:17–27, 2010. 10.1007/3DRes.01(2010)2.
[Koike08]	TAKAFUMI KOIKE AND TAKESHI NAEMURA. BRDF display: interac- tive view dependent texture display using integral photography. In <i>IPT/EDT '08: Proceedings of the 2008 Workshop on Immersive Pro-</i> <i>jection Technologies/Emerging Display Technologies</i> , pages 1–4, New York, NY, USA, 2008. ACM.
[Kutulakos08]	KIRIAKOS N. KUTULAKOS AND E. STEGER. A Theory of Refractive and Specular 3D Shape by Light-Path Triangulation. <i>International</i> <i>Journal of Computer Vision (IJCV)</i> , 76(1):13–29, 2008.
[Leland97]	JAMES E. LELAND, NORBERT L. JOHNSON, AND ANGELO V. ARECCHI. Principles of bispectral fluorescence colorimetry. <i>Photometric Engi-</i> <i>neering of Sources and Systems</i> , 3140(1):76–87, 1997.
[Lorensen87]	WILLIAM E. LORENSEN AND HARVEY E. CLINE. Marching cubes: A high resolution 3D surface construction algorithm. In <i>Proc. of ACM SIGGRAPH 87</i> , pages 163–169. ACM, 1987.



[Maas95]	HANS-GERD MAAS. New Developments in Multimedia Photogram- metry. In A. Grün and H. Kahmen, editors, <i>Optical 3D Measurement</i> <i>Techniques III</i> . Wichmann Verlag, 1995.
[Marschner00]	S. MARSCHNER, S. WESTIN, E LAFORTUNE, AND K. TORRANCE. Image-Based Bidirectional Reflectance Distribution Function Mea- surement. <i>Applied Optics</i> , 39(16):2592–2600, 2000.
[Matusik03]	WOJCIECH MATUSIK, HANSPETER PFISTER, MATT BRAND, AND LEONARD MCMILLAN. A Data-Driven Reflectance Model. ACM Trans. Graph. (Proc. SIGGRAPH), 22(3):759–769, 2003.
[Matusik09]	WOJCIECH MATUSIK, BORIS AJDIN, JINWEI GU, JASON LAWRENCE, HENDRIK P. A. LENSCH, FABIO PELLACINI, AND SZYMON RUSINKIEWICZ. Printing spatially-varying reflectance. ACM Trans. Graph. (Proc. ACM SIGGRAPH Asia), 2009.
[Miyazaki05]	D. MIYAZAKI AND K. IKEUCHI. Inverse Polarization Raytracing: Esti- mating Surface Shapes of Transparent Objects. In <i>Proc. of CVPR '05</i> , volume 2, pages 910–917, 2005.
[Mok93]	FAI H. Mok. Angle-Multiplexed Storage of 5000 Holograms in Lithium Niobate. <i>Optics Letters</i> , 18(11):915–917, 1993.
[Morris05]	N. J. W. MORRIS AND K. N. KUTULAKOS. Dynamic Refraction Stereo. In <i>Proc. of ICCV'05</i> , pages 1573–1580, 2005.
[Morris07]	NIGEL J. W. MORRIS AND KYROS N. KUTULAKOS. Reconstructing the Surface of Inhomogeneous Transparent Scenes by Scatter-Trace Photography . In <i>Proc. of ICCV'07</i> , pages 1–8, 2007.
[Murase92]	H. MURASE. Surface Shape Reconstruction of a Nonrigid Transparent Object Using Refraction and Motion. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 14(10):1045–1052, October 1992.
[Narasimhan05]	S. G. NARASIMHAN, S. K. NAYAR, B. SUN, AND S. J. KOPPAL. Struc- tured Light in Scattering Media. <i>Proc. of ICCV'05</i> , I:420–427, 2005.
[Nassau83]	KURT NASSAU. <i>The Physics and Chemistry of Color: The Fifteen Causes of Color</i> . John Wiley and Sons, New York, 1983.
[Nayar06]	S. K. NAYAR, G. KRISHNAN, M. D. GROSSBERG, AND R. RASKAR. Fast Separation of Direct and Global Components of a Scene Using High Frequency Illumination. In <i>Proc. of ACM SIGGRAPH 2006</i> , pages 935–944, 2006.



[Ngan05]	ADDY NGAN, FRÉDO DURAND, AND WOJCIECH MATUSIK. Experi- mental Analysis of BRDF Models. In <i>Proc. of the Eurographics Sym-</i> <i>posium on Rendering</i> , pages 117–226, 2005.
[Nicodemus77]	F. NICODEMUS, J. RICHMOND, J. HSIA, I. GINSBERG, AND T. LIMPERIS. Geometrical considerations and nomenclature for reflectance. <i>Na-</i> <i>tional Bureau of Standards, Washington, DC. Inst. for Basic Standards</i> , 1977.
[Nwodoh00]	THOMAS A. NWODOH AND STEPHEN A. BENTON. Chidi Holographic Video System. In <i>Practical Holography XIV and Holographic Materials VI</i> , pages 167–176, 2000.
[Orlov00]	PETER ORLOV. Volume Holographic Data Storage. <i>Communications of the ACM</i> , 43(11):46–54, 2000.
[Park04]	JOHNNY PARK AND AVINASH C. KAK. Specularity Elimination in Range Sensing for Accurate 3D Modeling of Specular Objects. In <i>Proceedings of 3DPVT'04</i> , pages 707–714, 2004.
[Park08]	J. PARK AND C. KAK. 3D Modeling of Optically Challenging Objects. <i>IEEE Trans. on Visualization and Computer Graphics</i> , 14(2):246–262, March/April 2008.
[Phong75]	BUI TUONG PHONG. Illumination for computer generated pictures. <i>Commun. ACM</i> , 18(6):311–317, 1975.
[Poisson37]	S. D. POISSON. Recherches sur la probabilité des jugements en matière criminelle et en matière civile, précédées des règles générales du calcul des probabilités. Bachelier, Paris, 1837.
[Proctor96]	JAMES PROCTOR AND YVONNE BARNES. NIST High Accuracy Reference Reflectometer-Spectrophotometer. <i>Journal of Research of the Nat. Institute of Standards and Technology</i> , 101(5):619–626, 1996.
[Puckette]	MILLER SMITH PUCKETTE. PureData (Pd). http://puredata.info.
[Rakuljic92]	GEORGE A. RAKULJIC, VICTOR LEYVA, AND AMNON YARIV. Optical Data Storage by using Orthogonal Wavelength-Multiplexed Volume Holograms. <i>Optics Letters</i> , 17(20):1471–1473, 1992.
[Raskar01]	RAMESH RASKAR, GREG WELCH, KOK-LIM LOW, AND DEEPAK BANDYOPADHYAY. Shader Lamps: Animating Real Objects With Image-Based Illumination. In <i>Proceedings of the 12th Eurographics</i> <i>Workshop on Rendering Techniques</i> , pages 89–102, London, UK, 2001. Springer-Verlag.



BIBLIOGRAPHY

[Remondino06]	F. REMONDINO AND S. EL-HAKIM. Image Based 3D Modeling: A Review. <i>The Photogrammetric Record</i> , 21(115):269–291, 2006.
[Rump10]	MARTIN RUMP AND REINHARD KLEIN. Spectralization: Reconstruct- ing spectra from sparse data. In Jason Lawrence and Marc Stam- minger, editors, <i>Rendering Techniques (Proc. EGSR 2010)</i> , pages 1347– 1354, Saarbruecken, Germany, June 2010. Eurographics Association.
[Saito99]	M. SAITO, YOICHI SATO, KATSUSHI IKEUCHI, AND H. KASHIWAGI. Measurement of Surface Orientations of Transparent Objects using Polarization in Highlight. In <i>Proc. of CVPR '99</i> , volume 1, pages 381– 386, 1999.
[Schawlow58]	A. L. SCHAWLOW AND C. H. TOWNES. Infrared and Optical Masers. <i>Phys. Rev.</i> , 112(6):1940–1949, Dec 1958.
[Schultz94]	H. SCHULTZ. Retrieving Shape Information from Multiple Images of a Specular Surface. <i>IEEE Transactions on Pattern Analysis and Ma-</i> <i>chine Intelligence</i> , 16(2):195–201, February 1994.
[Sharpe02]	J. SHARPE, U. AHLGREN, P. PERRY, B. HILL, A. ROSS, J. HECKSHER- SORENSEN, R. BALDOCK, AND D. DAVIDSON. Optical Projection To- mography as a Tool for 3D Microscopy and Gene Expression Studies. <i>Science</i> , 296(19):541–545, 2002.
[Sonka07]	MILAN SONKA, VACLAV HLAVAC, AND ROGER BOYLE. <i>Image Process-</i> <i>ing, Analysis, and Machine Vision</i> . Thomson-Engineering, 2007.
[Springsteen99]	ART SPRINGSTEEN. Introduction to measurement of color of fluores- cent materials. <i>Analyt. Chimica Acta</i> , 380(2–3):183–192, 1999.
[Tarini05]	MARCO TARINI, HENDRIK P. A. LENSCH, MICHAEL GOESELE, AND HANS-PETER SEIDEL. 3D Acquisition of Mirroring Objects. <i>Graphical Models</i> , 67(4):233–259, 2005.
[Texas Instruments]	TEXAS INSTRUMENTS. TALP1000B analog mirror datasheet. http://focus.ti.com/lit/ds/symlink/talp1000b.pdf.
[Torrance67]	K. E. TORRANCE AND E. M. SPARROW. Theory for Off-Specular Re- flection From Roughened Surfaces. <i>J. Opt. Soc. Am.</i> , 57(9):1105–1112, 1967.
[Treibitz08]	T. TREIBITZ, Y.Y. SCHECHNER, AND H. SINGH. Flat refractive geometry. pages 1–8, jun. 2008.
[Trifonov06]	BORISLAV TRIFONOV, DEREK BRADLEY, AND WOLFGANG HEIDRICH. Tomographic Reconstruction of Transparent Objects. In <i>Proc. of EGSR'06</i> , pages 51–60, 2006.



BIBLIOGRAPHY

[Trucco94]	EMANUELE TRUCCO AND ROBERT B. FISHER. Acquisition of Con- sistent Range Data Using Local Calibration. In <i>IEEE International</i> <i>Conference on Robotics and Automation</i> , pages 3410–3415, 1994.
[Ward92]	GREGORY J. WARD. Measuring and Modeling Anisotropic Reflection. <i>Computer Graphics (Proc. SIGGRAPH)</i> , 26(2):265–272, 1992.
[Weyrich09]	TIM WEYRICH, PIETER PEERS, WOJCIECH MATUSIK, AND SZYMON RUSINKIEWICZ. Fabricating Microgeometry for Custom Surface Re- flectance. <i>ACM Trans. Graph. (Proc. ACM SIGGRAPH)</i> , 2009.
[Wilkie01]	ALEXANDER WILKIE, ROBERT TOBLER, AND WERNER PURGATHOFER. Combined Rendering of Polarization and Fluorescence Effects. In <i>Proc. of the Eurographics Workshop on Rendering</i> , pages 197–204, 2001.
[Wilkie06]	ALEXANDER WILKIE, ANDREA WEIDLICH, CAROLINE LARBOULETTE, AND WERNER PURGATHOFER. A Reflectance Model for Diffuse Flu- orescent Surfaces. In <i>Proceedings of Graphite 2006</i> , pages 321–328, 11 2006.
[Zastrow81]	ARMIN ZASTROW. Physikalische Analyse der Energieverlustmech- anismen im Fluoreszenzkollektor. <i>PhD thesis, Albert-Ludwigs-</i> <i>Universität Freiburg, Germany, 1981</i> , 1981.
[Zickler02]	TODD ZICKLER, PETER N. BELHUMEUR, AND DAVID J. KRIEGMAN. Helmholtz Stereopsis: Exploiting Reciprocity for Surface Recon- struction. <i>International Journal of Computer Vision (IJCV)</i> , 49(2- 3):215–227, 2002.

123

