
New Acquisition Techniques for Real Objects and Light Sources in Computer Graphics

Michael Goesele

**Max-Planck-Institut für Informatik
Saarbrücken, Germany**

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

Eingereicht am 8. Juni 2004 in Saarbrücken.

Betreuender Hochschullehrer — Supervisor

Prof. Dr. Hans-Peter Seidel,
Max-Planck-Institut für Informatik, Saarbrücken, Germany

Gutachter — Reviewers

Prof. Dr. Hans-Peter Seidel,
Max-Planck-Institut für Informatik, Saarbrücken, Germany

Prof. Dr. Wolfgang Heidrich,
The University of British Columbia, Vancouver, Canada

Prof. Dr. Roberto Scopigno,
Istituto Scienza e Tecnologie dell'Informazione,
National Research Council, Pisa, Italy

Dekan — Dean

Prof. Dr. Jörg Eschmeier,
Universität des Saarlandes, Saarbrücken, Germany

**Promovierter akademischer Mitarbeiter —
Academic Member of the Faculty having a Doctorate**

Dr. Marcus Magnor,
Max-Planck-Institut für Informatik, Saarbrücken, Germany

Datum des Kolloquiums — Date of Defense

14. Juli 2004 — July 14th, 2004

Michael Gösele
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken, Germany
goesele@mpi-sb.mpg.de

Abstract

Accurate representations of objects and light sources in a scene model are a crucial prerequisite for realistic image synthesis using computer graphics techniques. This thesis presents techniques for the efficient acquisition of real world objects and real world light sources, as well as an assessment of the quality of the acquired models.

Making use of color management techniques, we setup an appearance reproduction pipeline that ensures best-possible reproduction of local light reflection with the available input and output devices. We introduce a hierarchical model for the subsurface light transport in translucent objects, derive an acquisition methodology, and acquire models of several translucent objects that can be rendered interactively. Since geometry models of real world objects are often acquired using 3D range scanners, we also present a method based on the concept of modulation transfer functions to evaluate their accuracy.

In order to illuminate a scene with realistic light sources, we propose a method to acquire a model of the near-field emission pattern of a light source with optical prefiltering. We apply this method to several light sources with different emission characteristics and demonstrate the integration of the acquired models into both, global illumination as well as hardware-accelerated rendering systems.

Kurzfassung

Exakte Repräsentationen der Objekte und Lichtquellen in einem Modell einer Szene sind eine unerlässliche Voraussetzung für die realistische Bilderzeugung mit Techniken der Computergraphik. Diese Dissertation beschäftigt sich mit der effizienten Digitalisierung von realen Objekten und realen Lichtquellen. Dabei werden sowohl neue Digitalisierungstechniken als auch Methoden zur Bestimmung der Qualität der erzeugten Modelle vorgestellt.

Wir schlagen eine Verarbeitungskette zur Digitalisierung und Wiedergabe der Farbe und Spekularität von Objekten vor, die durch Ausnutzung von Farbmanagement-Techniken eine bestmögliche Wiedergabe des Objekts unter Verwendung der gegebenen Ein- und Ausgabegeräte ermöglicht. Wir führen weiterhin ein hierarchisches Modell für den Lichttransport im Inneren von Objekten aus durchscheinenden Materialien sowie eine zugehörige Akquisitionsmethode ein und digitalisieren mehrere reale Objekte. Die dabei erzeugten Modelle können in Echtzeit angezeigt werden. Die Geometrie realer Objekte spielt eine entscheidende Rolle in vielen Anwendungen und wird oftmals unter Verwendung von 3D Scannern digitalisiert. Wir entwickeln daher eine Methode zur Bestimmung der Genauigkeit eines 3D Scanners, die auf dem Konzept der Modulationstrans-

ferfunktion basiert.

Um eine Szene mit realen Lichtquellen beleuchten zu können, schlagen wir ferner eine Methode zur Erfassung der Nahfeldabstrahlung einer Lichtquelle vor, bei der vor der Digitalisierung ein optischer Filterungsschritt durchgeführt wird. Wir wenden diese Methode zur Digitalisierung mehrerer Lichtquellen mit unterschiedlichen Abstrahlcharakteristika an und zeigen auf, wie die dabei erzeugten Modelle in globalen Beleuchtungsberechnungen sowie bei der Bildsynthese mittels moderner Graphikkarten verwendet werden können.

Summary

Accurate representations of objects and light sources in a scene model are a crucial prerequisite for realistic image synthesis using computer graphics techniques. Such models can be generated using a variety of methods (e.g., in a CAD based construction process or by physical simulation techniques). But given a *real world* object or a *real world* light source, the most reliable method to create an accurate representation of the object or light source is to acquire its properties using a suitable acquisition technique. This thesis therefore presents techniques for the efficient acquisition of real world objects and real world light sources. It also includes an assessment of the quality of the acquired models.

The appearance of an object showing only local light reflection can be modeled by a bidirectional reflectance distribution function (BRDF). We incorporate color management techniques into the appearance acquisition in order to reference the acquired models to a well-defined standard color space. On the output side of the appearance reproduction pipeline, we again use color management techniques to adapt the generated image data to the output device. Under the assumption that the color management system correctly describes the properties of the involved input and output devices we can achieve best-possible reproduction of local light reflection with the available input and output devices. We validate both, the acquired models as well as final renderings of the models, by comparing them to ground truth.

Translucent objects are characterized by diffuse light transport *inside* the object. To describe their light interaction behavior, instead of a BRDF the more complex bidirectional scattering-surface reflectance distribution function (BSSRDF) is required. We introduce a hierarchical model for diffuse subsurface light transport in translucent objects that approximates the BSSRDF by a diffuse reflectance function. We derive an acquisition methodology and setup an acquisition system that illuminates individual surface points of a translucent object and observes the object's response. The system is used to acquire models of several translucent objects and renderings of the acquired models are compared to photographs of the real objects under similar illumination conditions. The acquired models can be rendered interactively using modern graphics hardware but can also be incorporated into many other rendering systems.

Geometry models of real world objects are often acquired using 3D range scanning systems. Their accuracy is not only important for realistic rendering but also for many other acquisition techniques that require a geometry model as part of the input data. We present a method based on the concept of modulation transfer functions (MTF) to evaluate the accuracy of a 3D range scanner. We acquire a scan of a sharp edge and compute its Fourier transform. By comparing the resulting frequency spectrum with the frequency spectrum of an ideal sharp

edge we determine the MTF of the 3D range scanner. The computed MTF is a particularly relevant accuracy measure as it determines how well sharp features, that are contained in many real world objects, can be acquired.

Apart from realistic models of real world objects, realistic illumination is very important for high quality scene rendering. In order to be able to illuminate a scene with realistic light sources, we propose a method to acquire the near-field emission characteristics of real world light sources. Such a near-field model describes the emitted radiance for any position in space and any emission direction. In order to ensure correct sampling and to avoid aliasing artifacts, we perform an optical prefiltering step that projects the emitted radiance into a finite set of basis functions. The acquired models can be efficiently integrated into both global illumination as well as hardware-accelerated rendering systems. We acquire models of several light sources with different emission characteristics and compare the reconstructed emission pattern projected onto flat surfaces to photographs of the real world light sources illuminating similar diffuse surfaces.

In summary, this thesis contributes to both, acquisition techniques for real world objects and real world light sources and their validation. It covers a broad range of properties of real world objects including their geometry, their interaction with light, and their light emission characteristic. It helps to improve the quality of the models used as foundation for realistic image synthesis and allows for using accurate models of real world objects and real world light sources in rendering applications.

Zusammenfassung

Exakte Repräsentationen der Objekte und Lichtquellen in einem Modell einer Szene sind eine unerläßliche Voraussetzung für die realistische Bilderzeugung mit Techniken der Computergraphik. Solche Modelle von Objekten und Lichtquellen können durch einer Reihe von Methoden erzeugt werden (zum Beispiel in einem Konstruktionsprozess mittels CAD-Systemen oder mit physikalischen Simulationstechniken). Die zuverlässigste Methode, um ein exaktes Modell eines *echten* Objektes oder einer *echten* Lichtquelle zu erzeugen, ist jedoch die Digitalisierung ihrer Eigenschaften unter Verwendung geeigneter Akquisitionstechniken. Diese Arbeit stellt daher Techniken zur effizienten Digitalisierung von echten Objekten und echten Lichtquellen vor. Sie enthält außerdem Methoden, die die Beurteilung der Qualität der digitalisierten Modelle erlauben.

Das Aussehen eines Objekts, das Licht an der Oberfläche nur lokal reflektiert, kann mittels einer bidirektionalen Reflektanzverteilungsfunktion (Bidirectional Reflectance Distribution Function, BRDF) beschrieben werden. Wir integrieren Farbmanagementtechniken in ein Verfahren zur Digitalisierung von Farbe und Spekularität eines Objektes mit dem Ziel, daß die digitalisierten Datensätze in einem wohldefinierten Farbraum vorliegen. Wir wenden weiterhin zur Ausgabe der erzeugten Bilddaten Techniken des Farbmanagements an, um die Bilder an die Eigenschaften der verwendeten Ausgabegeräte anzupassen. Unter der Voraussetzung, daß das Farbmanagementsystem die Eigenschaften der verwendeten Ein- und Ausgabegeräte korrekt beschreibt, können wir so eine bestmögliche Reproduktion der lokalen Reflektionseigenschaften mit den gegebenen Geräten erzielen. Wir überprüfen die Genauigkeit der erzeugten Modelle und der aus den Modellen erzeugten Bilder, indem wir sie mit unabhängig gemessenen Werten beziehungsweise dem wahren Aussehen unter gegebenen Beleuchtungsbedingungen vergleichen.

Objekte aus durchscheinenden Materialien sind durch diffusen Lichttransport *im Inneren* der Objekte charakterisiert. Um dieses Verhalten zu beschreiben, wird anstatt einer BRDF die mächtigere bidirektionale Reflektanzverteilungsfunktion für streuende Materialien (Bidirectional Scattering-Surface Reflectance Distribution Function, BSSRDF) benötigt. Wir führen ein hierarchisches Modell für den diffusen Lichttransport in durchscheinenden Objekten ein, das die BSSRDF mittels einer diffusen Reflektanzfunktion annähert. Wir leiten daraus eine Digitalisierungsmethode ab und konstruieren einen Akquisitionsaufbau, der einzelne Oberflächenpunkte eines durchscheinenden Objektes beleuchtet und die Impulsantwort des Objektes aufzeichnet. Wir digitalisieren mehrere durchscheinende Objekte mit diesem Aufbau und vergleichen aus den Modellen erzeugte Bilder mit Photographien der echten Objekte unter ähnlichen Beleuchtungsbedingungen. Die erzeugten Modelle können unter Verwendung moderner Graphikkarten inter-

aktiv dargestellt werden. Sie können aber auch direkt in viele andere Systeme zur Bildsynthese integriert werden.

Geometriemodelle echter Objekte werden oft unter Verwendung von 3D Scannern erzeugt. Die Genauigkeit dieser Modelle ist nicht nur wichtig für realistische Bildsyntheseverfahren sondern auch für viele Digitalisierungsverfahren, die ein Geometriemodell als Teil der Eingabedaten voraussetzen. Wir stellen daher eine Methode zur Messung der Genauigkeit eines 3D Scanners vor, die auf dem Konzept der Modulationstransferfunktion (Modulation Transfer Function, MTF) beruht. Wir führen dazu einen einzelnen 3D Scan einer scharfen Kante durch und berechnen daraus die Fouriertransformation. Wir vergleichen das resultierende Frequenzspektrum mit dem Frequenzspektrum einer idealen scharfen Kante und berechnen daraus die MTF des 3D Scanners. Die Modulationstransferfunktion ist ein besonders wichtiges Genauigkeitsmaß, da sie angibt, wie gut scharfe Details digitalisiert werden können, die in vielen echten Objekten vorkommen.

Abgesehen von exakten Modellen echter Objekte, ist auch eine realistische Beleuchtung für die qualitativ hochwertige Bildsynthese erforderlich. Wir schlagen eine Methode zur Digitalisierung des Nahfeldes einer Lichtquelle vor, um eine Szene mit einer echten Lichtquelle beleuchten zu können. Das Nahfeld beschreibt die Lichtabstrahlung für alle Punkte und alle Abstrahlungsrichtungen. Wir führen einen optischen Filterungsschritt durch, der die Lichtabstrahlung in einen endlichdimensionalen Raum von Basisfunktionen projiziert, um eine korrekte Abtastung zu erhalten und Aliasingartefakte zu vermeiden. Die digitalisierten Modelle eignen sich sowohl zur effizienten Verwendung im Rahmen von globalen Beleuchtungsberechnungen als auch für die Bildsynthese mittels moderner Graphikkarten. Wir digitalisieren verschiedene Lichtquellen mit unterschiedlichen Abstrahleigenschaften und vergleichen das rekonstruierte Abstrahlverhalten mit dem Abstrahlverhalten der echten Lichtquellen, indem wir die Projektionen auf eine flache Oberfläche betrachten.

Zusammenfassend enthält diese Dissertation Beiträge zu Digitalisierungstechniken für echte Objekte und echte Lichtquellen sowie Methoden zur Validierung dieser Techniken. Die Arbeit deckt ein breites Spektrum von Eigenschaften realer Objekte einschließlich Geometrie, ihrer Interaktion mit Licht sowie ihrer Lichtabstrahlcharakteristik ab. Sie trägt dazu bei, die Qualität von Modellen, die als Grundlage für die realistische Bilderzeugung verwendet werden, zu verbessern und ermöglicht es, Modelle echter Objekte und echter Lichtquellen in der Bilderzeugung zu verwenden.

Acknowledgements

First of all, I would like to thank my supervisor Prof. Dr. Hans-Peter Seidel for his interest in this work, his valuable comments, his continuous support, and the freedom to pursue my own ideas.

I would also like to thank the two external reviewers Prof. Dr. Wolfgang Heidrich and Prof. Dr. Roberto Scopigno who agreed to review this thesis. Prof. Dr. Wolfgang Heidrich was my academic adviser during the first part of my thesis and hosted me twice at his group in Vancouver. An important part of my work on light sources is based on valuable discussions and close cooperation with him. Prof. Dr. Roberto Scopigno played also a vital part in the ViHAP3D project on acquisition and presentation of cultural heritage.

I would especially like to thank all my present and former colleagues at the Computer Graphics Group at MPI who help to make it such a great place. In particular, I owe special thanks to Philippe Bekaert, Christian Fuchs, Xavier Granier, Jan Kautz, Jochen Lang, Hendrik Lensch, Marcus Magnor and Marco Tarini who co-authored some of my previous publications. I am very thankful for the opportunity to collaborate with all these splendid and creative researchers. Without them, many of the projects described in this dissertation would not have been feasible.

Special thanks to Axel Köppel and Michael Laise from the MPI technical staff and the MPI Rechnerbetriebsgruppe for building some of the measurement apparatus and patiently answering to numerous requests. I am very grateful to John Lloyd from the University of British Columbia for his help in setting up and programming ACME for light source measurements. I would also like to thank Oliver Ashoff, Samir Hammann, Stoyan Mutafchiev, Thomas Neumann, Michael Schultz, Jan Uschok and Heiko Wanning who worked as research assistants at MPI, contributed to our hard- and software systems, and performed part of the measurements.

Many thanks also to Kristina Scherbaum who created most of the illustrations used in this thesis and to Xavier Granier, Jochen Lang, Karol Myszkowski, Holger Theisel and Ingo Wald for proof-reading drafts of this dissertation.

From the people not connected with my work, I would first like to thank my parents, Erika and Helmut Gösele, and my whole family who encouraged and supported me during my whole life. They also rose my interest in the wonders of nature already at an early age and started my career as a scientist. Many thanks to Kirsten Lemke who became a close friend during the last months of the work on this dissertation.

This work was funded by the Max-Planck-Gesellschaft zur Förderung der Wissenschaften e. V., by the DFG Schwerpunktprogramm V3D2 “Verteilte Verarbeitung und Vermittlung Digitaler Dokumente” and by the European Union within the scope of the ViHAP3D Project IST-2001-32641 “Virtual Heritage: High-

Quality 3D Acquisition and Presentation.” Thanks also to Konica Minolta Europe (Langenhagen, Germany) for providing 3D scanners and to Villeroy & Boch (Mettlach, Germany), the Musée de la Faïence (Sarreguemines, France) and the Archaeological Museum (Florence, Italy) who gave us access to some of the test objects used in this dissertation.

Contents

1	Introduction	1
1.1	Problem Statement	2
1.2	Main Contributions	2
1.3	Chapter Overview	4
2	Background	7
2.1	Models of Light	7
2.2	Radiometry and Photometry	9
2.2.1	Radiometric Quantities	9
2.2.2	Photometric Quantities	10
2.3	Light Sources	11
2.3.1	Light Generation and Illuminants	11
2.3.2	Spectral Properties	12
2.3.3	Radiation Patterns of Luminaires	13
2.3.4	Light Sources and Light Source Representations	14
2.3.5	Near-Field and Far-Field of a Luminaire	16
2.4	Interaction of Light and Matter	17
2.4.1	Physics of Reflection and Refraction at a Material Boundary	18
2.4.2	Bidirectional Scattering-Surface Reflection Distribution Function (BSSRDF)	19
2.4.3	Bidirectional Reflectance-Distribution Function (BRDF) .	21
2.4.4	Reflection Models	22
2.4.5	Reflection Models for Subsurface-Scattering Material . . .	23
2.5	Rendering Equation	25
2.6	Inverse Rendering	26
2.6.1	Photometric Stereo	27
2.6.2	Inverse Lighting	27
2.6.3	Inverse Reflectometry	27

3	Acquisition Facilities	29
3.1	Geometry Acquisition	29
3.1.1	Performance and Accuracy	30
3.1.2	3D Geometry Processing Pipeline	31
3.2	Digital Camera Systems	32
3.2.1	Camera Setup and Components	32
3.2.2	Geometric Calibration	35
3.2.3	Photometric Calibration	37
3.2.4	Camera Systems	39
3.3	Illumination and Environment	43
3.3.1	Point and Area Light Sources	43
3.3.2	Laser System	43
3.3.3	Environment	44
3.4	ACME Robotic Measurement Facility	45
4	MTF Estimation for 3D Range Scanners	47
4.1	Previous Work	48
4.1.1	Range Scanner Accuracy	48
4.1.2	OTF Measurement	49
4.2	Measurement Principle	49
4.2.1	Slanted Edge Analysis for 3D Range Scanners	50
4.2.2	Data Processing	52
4.2.3	Data Analysis	52
4.3	The “Ideal” Modulation Transfer Function	53
4.3.1	Three-Dimensional Systems	54
4.4	Empirical Analysis	56
4.4.1	Analysis of a Single Scan	56
4.4.2	Repeatability and Robustness	56
4.5	Further Examples	58
4.5.1	Horizontal and Vertical MTF	58
4.5.2	Scan Processing	59
4.5.3	Scanner Comparison	60
4.6	Discussion	60
4.7	Conclusion	63
5	Color Calibrated BRDF Acquisition and Reproduction	65
5.1	Previous Work	66
5.1.1	High-Dynamic Range Imaging	66
5.1.2	BRDF Measurement	67
5.2	Acquisition and BRDF Generation	68
5.2.1	BRDF Acquisition and Modeling Pipeline	68

5.2.2	BRDFs in CIEXYZ Color Space	68
5.3	Color Managed HDR Imaging	69
5.4	Experimental Validation of the BRDF Models	70
5.4.1	Spectrophotometric Measurements	71
5.4.2	Comparison to a Photographic Image	71
5.5	Calibrated Rendering and Output	73
5.5.1	Rendering of Calibrated BRDF Models	73
5.5.2	Validation Including Output Devices	74
5.6	Discussion	76
5.6.1	Accuracy of Color Management	76
5.6.2	CIEXYZ versus sRGB Color Space	76
5.7	Conclusion	76
6	Light Source Acquisition	79
6.1	Previous Work	80
6.2	Overview	81
6.3	Theoretical Framework	83
6.3.1	Measured Irradiance	84
6.3.2	Exact Reconstruction	85
6.3.3	Shift-Invariant Approximation	85
6.4	Filter Design	87
6.4.1	Piecewise Quadratic Basis	88
6.4.2	B-Spline Basis	89
6.4.3	Filter Fabrication and Calibration	90
6.5	Acquisition Setup	91
6.5.1	Direct Acquisition (Setup A)	91
6.5.2	Reflector-Based Acquisition (Setup B)	91
6.6	Rendering	94
6.6.1	Importance Sampling	95
6.6.2	Estimation of Direct Illumination	96
6.7	Results	98
6.7.1	Acquisition	98
6.7.2	Rendering	104
6.8	Conclusions	106
7	Translucent Objects	107
7.1	Previous Work	108
7.1.1	Models for Translucent Objects	109
7.1.2	Acquisition Techniques	109
7.1.3	Rendering	110
7.2	Hierarchical BSSRDF Model	111

7.2.1	Discretization of Light Transport	113
7.2.2	Modulation of Smooth Hierarchy Levels	114
7.2.3	Implemented Model	114
7.3	Generation of Synthetic Datasets	118
7.3.1	Local Response	119
7.3.2	Global Response	119
7.3.3	Modulation Texture	120
7.3.4	Examples	120
7.4	Acquisition of Translucent Objects	120
7.4.1	Acquisition Setup	121
7.4.2	Geometric and Photometric Calibration	122
7.5	Efficient Data Access	123
7.6	Global Term Post-Processing	124
7.6.1	Data Resampling	125
7.6.2	Interpolation	125
7.6.3	Modulation Texture	127
7.7	Local Term Post-Processing	128
7.7.1	Data Resampling	128
7.7.2	Interpolation	129
7.8	Rendering	130
7.8.1	Computing the Illumination	132
7.8.2	Radiosity Reconstruction	132
7.9	Acquisition Results	133
7.10	Conclusion	138
8	Summary, Conclusion and Future Work	141
8.1	Summary	141
8.2	Conclusion	142
8.3	Future Work	143
	Bibliography	145
	Curriculum Vitae – Lebenslauf	161

Chapter 1

Introduction

Photorealistic rendering of complex scenes can nowadays be performed even on commodity PC hardware at interactive rates. Offline rendering algorithms can generate images that are virtually indistinguishable from reality. As improvements in hardware capability follow or even surpass Moore's law [Moore65] and more efficient rendering algorithms are constantly introduced, there is also a rising demand for more and more complex and realistic input data modeling all relevant properties of a scene. This generally includes the scene geometry as well as light emission and its interaction with objects consisting of various materials.

Accurate scene models can originate from a variety of sources: If highly detailed CAD data is generated in an engineering process such as the construction of a new car, aircraft or building, it can also be used for rendering purposes provided that the optical properties are modeled as well. Artists are able to build compelling models of individual objects and whole environments that are tailored to a specific application – an approach often used by the movie industry. Furthermore, physical simulations of natural phenomena or mathematical models such as fractal geometry can create realistically looking input data.

But given a *real world* object, none of the above methods is well suited to generate a realistic model of its geometry and interaction with light or its light emission characteristics for the case of a light source. In particular, small imperfections characteristic to an individual specimen – even if it is generated by industrial mass production methods – pose a huge challenge. This thesis focuses therefore on measurement based modeling approaches that capture properties of real world objects and real world light sources relevant for realistic image synthesis.

1.1 Problem Statement

The complexity of the properties of real world objects make acquiring a *general* and *complete* representation extremely hard or even impossible. The first challenge in acquiring the properties of a real world object is therefore normally to define a suitable model that allows for an accurate representation of the property or properties under consideration. For computer graphics applications, it is often possible to use a simplified, approximative model as long as it is still sufficiently accurate for realistic image synthesis.

The next task is then to devise an acquisition method that estimates the model parameters so that it best approximates the properties of the real world object. This normally requires the use of various physical devices that need to be carefully calibrated and operated under controlled conditions. The acquired data is then used in a modeling or optimization step to determine the correct model parameters.

The acquired models must then be converted into a format that allows for efficient access to the relevant information by the target applications. The requirements of different applications, e.g., a global illumination system or a hardware-accelerated rendering system, may differ quite strongly so that the models need to be adapted to each application scenario. It is therefore advisable to be already aware of the final requirements when designing a model and an acquisition approach.

Finally, the quality of the acquired models should be validated to ensure that they accurately represent the given object properties. Validation is usually performed by comparing the models to ground truth or by assessing the quality in a user study.

This dissertation covers four topics related to real world acquisition techniques that are also represented in Figure 1.1. We first present a method to estimate the accuracy of 3D range scanners commonly used to acquire the geometry of an object. We then extend an appearance acquisition and reproduction pipeline by a color management system in order to ensure correct color and appearance acquisition. Finally, we introduce methods to acquire the near-field of a light source and the reflection properties of translucent materials. All methods result in models that can be efficiently integrated into both global illumination systems and hardware-accelerated rendering approaches.

1.2 Main Contributions

Throughout the course of this dissertation, parts have already been published at different conferences and in various journals [Goesele00, Goesele01, Heidrich01, Goesele03b, Granier03, Lensch03b, Goesele03a, Goesele04a,




Figure 1.1: *Topics addressed in this dissertation. Top left: 3D range scan of a sharp edge used for 3D scanner accuracy estimation. Top right: Color-calibrated (left) and uncalibrated appearance dataset of a carafe. Middle: Rendering using an acquired light source model of a car headlight and bike light in acquisition setup. Bottom: Partially illuminated model of a translucent horse with local and global light propagation.*

Lensch04, Goesele04b]. These publications are the foundation of this thesis which incorporates them under the topic of real world acquisition techniques and presents improvements and updated results.

The main contributions of this dissertation can be summarized as follows:

- The setup of an acquisition lab tailored to the acquisition of real world objects and real world light sources and a description of the used hardware.
- An extension of the concept of modulation transfer functions (MTF) to 3D scanning systems and an efficient measurement methodology. Given a single 3D scan of a slanted edge, the accuracy of the scanning system can be determined for the practically relevant case of sharp feature edges in an object.
- The integration of an acquisition algorithm for spatially varying bidirectional reflectance distribution functions (BRDF) into a completely color-managed processing pipeline. The final models can thus be validated through ground-truth color measurements using a spectrophotometer and the error between rendered model and the real object under identical illumination conditions can be quantified.
- An acquisition algorithm for real light sources that takes specifically includes correct prefiltering before the sampling stage. The acquired light source models can be used to illuminate arbitrary scenes in a global illumination framework but can also be rendered using programmable shader hardware of current graphics cards.
- A hierarchical model for the bidirectional scattering-surface reflectance distribution function (BSSRDF) for translucent objects and a corresponding acquisition algorithm to acquire such a model for translucent real-world objects. The model can be integrated into many current rendering algorithms for translucent materials.

1.3 Chapter Overview

The remainder of this dissertation is structured as follows: Chapter 2 summarizes the fundamentals of physics and computer graphics relevant to the covered topics. We then discuss the acquisition hardware and software used as foundation for our work (Chapter 3). Our method to determine the modulation transfer function of a 3D scanning system is introduced in Chapter 4 and the color-managed BRDF pipeline is presented in Chapter 5. Next, we discuss our image-based lightsource model and the corresponding acquisition and rendering system (Chapter 6). We

derive then in Chapter 7 a hierarchical reflection model for translucent objects, show how its parameters can be measured for real objects and introduce an efficient rendering algorithm. The dissertation ends with the conclusion and an outlook on future work (Chapter 8).

Chapter 2

Background

The image we perceive if we look at an object depends on a whole series of effects (Figure 2.1): Light is emitted by a light source and illuminates an object. A portion of the energy is absorbed by the object whereas the remainder is reflected – either directly at the surface or after passing through parts of the volume occupied by the object. The light finally reaches an observer’s eye. Alternatively, it is digitized by some acquisition device such as a digital camera.

This thesis focuses on acquisition techniques for computer graphics applications. This includes capturing the light emission properties of real-world light sources and acquiring the interaction behavior of light with translucent objects. But it draws also on all other parts of the “visual pipeline”. We therefore review first the underlying physics of light and give then an overview over light sources and their properties. We next describe ways and models of how light interacts with objects. We briefly review the rendering equation and show how it is used in inverse rendering to gain information about the properties of light sources and materials.

2.1 Models of Light

One of the fundamental aspects of computer graphics are the properties of light, the interaction of light with various materials, and the reception of light by either the human visual system or by an imaging system. We will in the following review the basic physical terms and concepts of light relevant to computer graphics. More in depth descriptions can be found in standard physics textbooks such as Born and Wolf [Born99], specific computer graphics applications are for example described in Glassner [Glassner94] or Cohen and Wallace [Cohen93].

The physics of light can be described by several different models. *Quantum Optics*, the most comprehensive model, is derived from the principles of quantum




Figure 2.1: The “visual pipeline”. Light is emitted by a light source and illuminates an object. The object interacts with the light, both, on its surface and within its whole volume and reflects a portion of the light. A human observer looks at the scene and interprets the physical signal that arrives at his eyes. Alternatively, it is acquired by a digital camera system.

mechanics. Its relevance to computer graphics and vision is limited to effects such as the interaction of individual photons (light particles) and matter in an imaging sensor (photoeffect).

Light can also be described as an electromagnetic wave since it is a form of electromagnetic radiation. *Wave Optics* can model a broad range of optical phenomena such as interference or diffraction. For example, star-shaped patterns in captured images of bright light sources are often caused by diffraction at the edges of the aperture blades (see Figure 2.2).

Wave optics can be approximated by *Geometrical Optics* or *Ray Optics* if the wavelength λ of light is much shorter than the size of relevant features.¹ Light is modeled as independent rays transporting energy through space; contributions of multiple rays can be summed based on the principle of superposition. The interactions of light with matter can be described in geometrical terms. Most computer graphics and vision applications use this model. It is, e.g., the foundation for ray

¹Born and Wolf [Born99] define geometrical optics as an approximation to wave optics neglecting the wavelength of light, i.e., observing the limiting case $\lambda \rightarrow 0$.




Figure 2.2: 10-sided star around a bright light source caused by diffraction at the edges of the camera lens' pentagonal aperture. The image also shows other effects such as lens flare caused by reflections inside the optical system.

tracing systems [Appel68, Rubin80] but also for the OpenGL rendering system [Segal98].

Particles are a discretization of geometrical optics often used in computer graphics with no physical equivalent. Light is modeled as a flow of discrete particles (similar to photons) carrying energy. In contrast to particles in quantum optics, there is no interaction between the particles. Important examples are Monte Carlo simulations of light transport [Dutr 03] or photon mapping [Jensen01b].

2.2 Radiometry and Photometry

Radiometry is the measurement of electromagnetic energy and covers a wide range of wavelengths including infrared, visible, and ultraviolet light. In contrast to that, photometry covers only the visible spectrum and takes the sensitivity of the human visual system into account.

2.2.1 Radiometric Quantities

Radiant Energy Q : Basic quantity describing the energy transported by light of all wavelengths measured in *Joule* [$J = W \cdot s = kg \cdot m^2/s^2$].

Radiant Flux or Radiant Power Φ : Radiant energy per unit time or power of the radiation measured in *Watt* [$W = J/s = kg \cdot m^2/s^3$].

$$\Phi := \frac{dQ}{dt}$$

Radiant Intensity I : Radiant Flux per solid angle ω measured in $[W/sr]$.

$$I := \frac{d\Phi}{d\omega}$$

The radiant intensity is often used to describe a point light source. The radiant intensity of an isotropic point light source is for example $I = \Phi/(4\pi \cdot sr)$ since a full sphere has a solid angle of $4\pi \cdot sr$.

Radiant Flux Area Density u : Differential flux per differential area A measured in $[W/m^2]$.

$$u := \frac{d\Phi}{dA}$$

Two important special cases are the (*incident*) *Irradiance E* (differential flux per differential area incident at a surface location) and *Radiant Exitance* or *Radiosity B* (differential flux per differential area exiting at a surface location).

Radiance L : Radiant flux per unit projected area per unit solid angle (arriving or leaving from a surface location at an angle θ to the surface normal) measured in $[W/(m^2 \cdot sr)]$.

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

Radiance can be interpreted as the number of photons per time passing through a small area from a certain direction. It is exactly what is perceived by an observer – be it a human eye or an electronic device – and is therefore one of the most important quantity for image-based modeling. Radiance has the important property of being constant along a ray in empty space. This is the inherent reason for the fact that the signal seen by an imaging system is independent of the distance to the object. On the rendering side, radiance is (at least implicitly) used by most rendering systems such as ray tracing.

All of the above radiometric quantities can also be evaluated for a specific wavelength λ instead of the whole spectrum. The unit of a spectral term is the unit of the corresponding radiometric term divided by the unit of length $[m]$. Spectral radiosity B_λ or $B(\lambda)$ has therefore the unit $[W/m^3]$.

2.2.2 Photometric Quantities

Photometry is the psychophysical measurement of the visual sensation produced by the electromagnetic spectrum. It therefore only covers the visible spectrum

Radiometric Quantity	Radiometric Unit	Photometric Quantity	Photometric Unit
Radiant Energy	[J]	Luminous Energy	[$talbot$]
Radiant Flux	[$W = J/s$]	Luminous Flux	[$lm = talbot/s$]
Radiant Intensity	[W/sr]	Luminous Intensity	[$cd = lm/sr$]
Irradiance Radiosity	[W/m^2]	Illuminance Luminosity	[$lx = lm/m^2$]
Radiance	[$W/(m^2 \cdot sr)$]	Luminance	[$nit = cd/m^2$]

Table 2.1: Basic radiometric quantities and their photometric counterparts. The new units talbot, lumen (lm), lux (lx) and nit can be derived from the SI unit candela (cd).

which is generally considered to be wavelengths from about 380 to 780 nm . Different wavelengths within the visible spectrum correspond to different colors ranging from violet to red. Each photometric quantity can be computed from the corresponding spectral radiometric quantity by convolving it with the spectral efficiency function $V(\lambda)$ [CIE86]. Luminosity B_v – the photometric quantity corresponding to radiosity – is therefore defined as

$$B_v = \int B(\lambda) \cdot V(\lambda) d\lambda$$

and is measured in lux [lx]. Table 2.1 summarizes and compares radiometric and photometric quantities and their units.

2.3 Light Sources

A light source or *luminaire* contains generally one or several *illuminants* – the actual devices that emit visible light. The remaining parts of a luminaire such as focusing optics or diffusing elements shape the emitted radiation pattern, provide the mechanical structure for the luminaire, or might even serve purely artistic purposes.

2.3.1 Light Generation and Illuminants

A variety of mechanisms can be used to generate visible light. We will in this section discuss the mechanisms most relevant to modern illuminants.

Heat generated by nuclear fission processes is responsible for sun and star light. The heat caused by a combustion process is also the oldest way used by humans to generate light – burning wood or the flames of oil lamps were already




Figure 2.3: Scene with plastic patches illuminated by an HMI arc lamp (near daylight spectrum, left) and green LEDs (right). Note the “black and green” effect in the right image due to the narrow spectrum of the LEDs. The patches on the right show fluorescence: the green LED light is absorbed and emitted at a different wavelength.

used in ancient times to illuminate human dwellings. Modern lamps use the same principle: Inside a tungsten lamp a piece of wire is heated up by an electrical current so that it emits light; in an arc lamp, the electrical current leads to hot plasma that emits light.

In neon tubes and other *gas discharge* lamps, gas inside a glass tube is excited and emits light of distinct wavelength. This principle is also the base of fluorescent tubes and modern low-energy bulbs. Hereby mercury vapor is excited and emits invisible, ultraviolet light. A layer of fluorescent material on the inside of the glass tube absorbs this radiation and immediately re-emits it in the visible spectrum using *fluorescence*.

Due to their rapidly increasing power, high efficiency and durability, light emitting diodes (LED) are currently becoming an important illuminant. The *re-combination* of holes and electrons in a semiconductor diode leads to excessive energy which is emitted from the diode as light within a narrow frequency spectrum.

A huge variety of other physical and chemical processes – lasers, organical LEDs, chemoluminescence to name a few – can also lead to the emission of light. They are however of limited practical relevance for illumination purposes and therefore not discussed.

2.3.2 Spectral Properties

The spectrum of light directly influences the color of a light source that is observed by a human observer or an imaging system. It is however also important



Figure 2.4: *Common elements to modify the radiation pattern of a luminaire. Left: A lamp screen with diffuse transmission. Middle: Office lamp with reflective components. Right: Spotlight with polygonal aperture and projective lens.*

for the observed color of surfaces that are illuminated by a light source. In the absence of fluorescence or phosphorescence effects, this process can be modeled as a convolution of the illuminant spectrum with the reflective spectrum of the material. If the illuminant spectrum contains sharp peaks as is the case for many LEDs or gas discharge lamps, a colorful scene might be observed as being “black and white” (or rather red, green, blue, . . . – see Figure 2.3 for an example). But even smooth spectra such as the reddish spectrum of a tungsten bulb have a significant influence on the observed color.

The international commission on illumination (Commission Internationale de l’Éclairage, CIE) therefore defined multiple standard illuminants such as illuminant A (corresponding to a tungsten filament lamp with a correlated color temperature of 2856 Kelvin) or illuminant D_{65} (daylight with a correlated color temperature of approximately 6500 Kelvin) [CIE86]. These illuminants can be used in order to analyze the general interaction behavior between light and materials.

2.3.3 Radiation Patterns of Luminaires

Luminaires can fulfill different purposes: A car headlight should illuminate the relevant part of the road without blinding other drivers. The emission characteristics should therefore be very directionally dependent. In contrast to that the main purpose of a brake light is to be a clearly visible warning sign for other drivers; its radiation pattern should be in general diffuse although individual peaks can improve visibility.

Luminaires contain various elements modifying the emission characteristics of the illuminants in order to shape its radiation pattern. *Diffusely reflecting* or *diffusely transmitting elements* provide a large emissive area leading to directionally as well as spatially uniform illumination (see Figure 2.4 (left)). The element




Figure 2.5: *Light Source Configurations.* Left: A uniform point light source at \vec{x}_e illuminating a point \vec{x}_i . \hat{n}_i is the surface normal. $\hat{\omega}_e = -\hat{\omega}_i$ is the normalized direction vector from \vec{x}_e to \vec{x}_i . Middle: A spotlight with opening angle α and center direction $\hat{\omega}_s$. Right: Area light source. The plane A is used to parameterize the emission position and direction.

itself is easily visible from all directions. The emission direction can be controlled within some limits, e.g., by a diffuse reflector with concave shape.

Specular reflectors (mirrors) allow for a much better control with hardly any loss of intensity and are therefore used in some office lamps where a uniform illumination of the working area is required (see Figure 2.4 (middle)). Parabolic reflectors are also often employed in flashlights or car headlights to focus the emitted light.

Refractive elements such as lenses can also focus light or change the light direction. Strong spotlights often use a lens system and a changeable aperture to achieve maximum flexibility in lighting effects (Figure 2.4 (right)). Lenses are also used in slide projectors to project an image of an object.

2.3.4 Light Sources and Light Source Representations

Light sources can be characterized according to their emission characteristics into point and area light sources.

Point Light Sources

A *point light source* is characterized by the fact that all light is emitted from a single point in space \vec{x}_e . There are several types of point light sources defined by their emission characteristics.

If light is emitted equally in all directions, the light source is a *uniform point light source*. The irradiance caused by such a light source at a receiver at \vec{x}_i is given as:

$$E(\vec{x}_i) = I_e \cdot V(\vec{x}_e, \vec{x}_i) \cdot \frac{\langle \hat{\omega}_i \cdot \hat{n}_i \rangle}{|\vec{x}_e - \vec{x}_i|^2}$$

I_e is the radiant intensity and $|\vec{x}_e - \vec{x}_i|^2$ is the squared distance between the light source position \vec{x}_e and \vec{x}_i . The unit direction vector from \vec{x}_i to \vec{x}_e is denoted by $\hat{\omega}_i$.² The visibility term $V(\vec{x}_e, \vec{x}_i)$ encodes both the visibility between the two points (i.e., whether the line of sight is blocked) and the orientation of the surface. It is defined as:

$$V(\vec{x}_e, \vec{x}_i) := \begin{cases} 1 & \vec{x}_e \text{ visible from } \vec{x}_i \text{ and } \langle \hat{\omega}_i \cdot \hat{n}_i \rangle > 0 \\ 0 & \text{else} \end{cases}.$$

We assume that the HMI light source described in Section 3.3.1 acts as an uniform point light source within the working area.

A *spot light* is a uniform point light source that emits light only within a cone of directions around a center direction $\hat{\omega}_s$, i.e.,

$$I_e(\hat{\omega}_e) = \begin{cases} I_e & \langle \hat{\omega}_s \cdot \hat{\omega}_e \rangle > \cos \alpha \\ 0 & \text{else} \end{cases}.$$

The opening angle of the spotlight is α . Note that $\hat{\omega}_e = -\hat{\omega}_i$ due to the geometry of the setup.

For a *textured point light source* [Segal92], the intensity $I_e(\hat{\omega}_e)$ can vary freely with the emission direction $\hat{\omega}_e$. *Goniometric diagrams* [Verbeck84] are textured point light sources often used to characterize the emission characteristics of a (real or artificial) luminaire. They are frequently provided by light source manufacturers for their luminaires.

Area Light Source

While point light sources are convenient for modeling and rendering purposes, most real world light sources are *area light sources* with finite spatial extent. The irradiance caused at position \vec{x}_i by an area light source can be computed as follows:

$$E(\vec{x}_i) = \int_A V(\vec{x}_e, \vec{x}_i) \cdot L_e(\vec{x}_e, \hat{\omega}_e) \cdot \frac{\langle \hat{\omega}_e \cdot \hat{n}_e \rangle \langle \hat{\omega}_i \cdot \hat{n}_i \rangle}{|\vec{x}_e - \vec{x}_i|^2} d\vec{x}_e.$$

²Vectors \vec{x} with arrow $\vec{}$ refer to positions in space whereas vectors \hat{n} or $\hat{\omega}$ with hat $\hat{}$ refer to normalized directions.




Figure 2.6: 2D plot of a goniometric diagram. The yellow circle marks the position of the point light source, the arrows denote the emission intensity for a given direction.

We use a plane A to parameterize the emitted radiance $L_e(\vec{x}_e, \hat{\omega}_e)$ at position \vec{x}_e within A . $\hat{\omega}_e$ is the emission direction (see Figure 2.5).

Note that this approach can also be used to parameterize a luminaire with volumetric emission characteristics.

2.3.5 Near-Field and Far-Field of a Luminaire

The emission characteristics of an area light source shows in general both angular and spatial variation and can therefore be described by a four-dimensional function $L_e(\vec{x}_e, \hat{\omega}_e)$ where \vec{x}_e determines a surface position and $\hat{\omega}_e$ gives the emission direction. $L_e(\vec{x}_e, \hat{\omega}_e)$ encodes the *near-field* of a luminaire.

In contrast to that, the *far-field* of a luminaire is only parameterized by the emission direction assuming that all light is emitted from a single point in space. Consequently, the far-field $L_e(\hat{\omega}_e)$ varies only with $\hat{\omega}_e$. Goniometric diagrams therefore only represent the far-field of a luminaire.

As a rule of thumb, the emission characteristics of a luminaire can be approximated by its far field, if it is observed from a distance of at least five times its active diameter [Ashdown95]. The active diameter is defined as the largest diameter of the light emitting area of the luminaire. This *five times rule* ensures that the error in estimated irradiance is below 1 % provided that the emission does not vary over the luminaire's surface [Murdoch81]. If the distance between emitter and receiver is too small to apply this approximation, the luminaire can be logically split into

several smaller units so that each of them can be approximated by a point light source.

Near-Field Effects

For a typical lighting design application such as computing the irradiance in an office environment, an area light source can be approximated with sufficient accuracy by its far-field. But there are several effects that can only be reproduced using a luminaire's near-field representation.

For any point \vec{x} in space, a point light source is either visible or the line-of-sight between the point light source's position and \vec{x} is blocked by an obstacle. \vec{x} is therefore either illuminated or in shadow leading to *hard shadows*. If the light is however emitted by an area light source, the view of the light source from \vec{x} can also be partially blocked. The shadow becomes a *soft shadow* consisting of a *penumbra region* where the light source is partially blocked and an *umbra region* where it is blocked completely. Soft shadows can therefore only occur if the near-field of a luminaire is modeled independent of the distance between a luminaire and \vec{x} .

Caustics are brighter areas where light is focused (most often due to reflection or refraction). A common example is light shining through a glass of water and producing a bright pattern on an object behind the glass. Caustics can also be caused by elements of a luminaire such as a reflector that focuses the emitted light at some points in 3D space. A faithful model of a luminaire including caustics must therefore include light emitted from various positions on the light source hitting the same point in space. This is impossible using the far-field assumption.

A slide projector can be seen as a special type of luminaire. If it is modeled as a textured point light source, which seems to be at first glance the corresponding model, the projected image will always be *in focus*. In order to model the fact that the projection is only focused within some volume (similar to caustics), again a near-field representation is required.

2.4 Interaction of Light and Matter

The interaction of light and matter is a complex physical process. Various aspects can be studied using the models of light presented in Section 2.1. We will focus here on the parts relevant to computer graphics and the topics covered in this dissertation and refer to the literature (e.g., the books by Glassner [Glassner94] or Born and Wolf [Born99]) for a more general description.




Figure 2.7: Light ray hitting a material boundary. The ray is split into a reflected and a refracted ray.

2.4.1 Physics of Reflection and Refraction at a Material Boundary

When light crosses the boundary between different materials its behavior is mainly determined by the indices of refraction of the two materials. The index of refraction η_λ is defined by the ratio of the speed of light in a material v_λ to the speed of light in vacuum c

$$\eta_\lambda = c/v_\lambda$$

and depends on the wavelength λ .

A single-wavelength light ray hitting a material boundary will in general be split into a reflected and a refracted ray (see Figure 2.7). The three rays and the surface normal are all in the same plane, the *plane of incidence*. *Snell's law* determines the direction of the two new rays while magnitude of the reflected and refracted flux can be computed using the *Fresnel equations*.

Snell's Law

The angle θ_t between the surface normal \hat{n} and the transmitted ray \vec{t} can be computed from the angle θ_i between the incoming ray \vec{i} and the surface normal:

$$\eta_1 \sin \theta_i = \eta_2 \sin \theta_t.$$

The angle θ_r between the reflected ray \vec{r} and the surface normal is equal to θ_i :

$$\theta_i = \theta_r.$$

Fresnel Equations

In addition to the indices of refraction η , the magnitude of reflected and refracted flux depend also on polarization:

$$\rho_{\parallel} = \frac{\eta_2 \cos \theta_1 - \eta_1 \cos \theta_2}{\eta_2 \cos \theta_1 + \eta_1 \cos \theta_2}$$

$$\rho_{\perp} = \frac{\eta_1 \cos \theta_1 - \eta_2 \cos \theta_2}{\eta_2 \cos \theta_1 + \eta_1 \cos \theta_2}.$$

ρ_{\parallel} is the reflection coefficient for light with the electric field parallel to the plane of incidence whereas ρ_{\perp} is the reflection coefficient for light with the electric field perpendicular to the plane of incidence.

The specular reflectance (also called *Fresnel reflection coefficient* F_r) for unpolarized light becomes then:

$$F_r(\theta) = \frac{1}{2}(\rho_{\parallel}^2 + \rho_{\perp}^2) = \frac{d\Phi_r}{d\Phi_i}.$$

For the refracted ray, the *Fresnel transmission coefficient* can be computed as:

$$F_t(\theta) = 1 - F_r(\theta).$$

Schlick [Schlick94] proposed a good approximation to the Fresnel reflection coefficient for unpolarized light based on the Fresnel reflection coefficient at normal incidence F_0 :

$$F_r(\theta) \approx F_0 + (1 - F_0)(1 - \cos \theta)^5.$$

It is possible to describe the interaction of light and matter entirely using their physical interaction principles such as Snell's law or the Fresnel equations. Alternatively, we can observe the actual interaction using suitable techniques and derive models that reflect this interaction behavior. The remainder of this section describes two functions (the BSSRDF and the BRDF) that can describe light interaction with different materials.

2.4.2 Bidirectional Scattering-Surface Reflection Distribution Function (BSSRDF)

The *bidirectional scattering-surface reflectance distribution function* (BSSRDF) [Nicodemus77] provides a general model of light transport inside an object and is defined as follows:

$$S(\vec{x}_i, \hat{\omega}_i; \vec{x}_o, \hat{\omega}_o) := \frac{dL^{\rightarrow}(\vec{x}_o, \hat{\omega}_o)}{d\Phi^{\leftarrow}(\vec{x}_i, \hat{\omega}_i)} \quad (2.1)$$




Figure 2.8: *Geometric configuration of a BSSRDF.*

The BSSRDF S is the ratio of reflected radiance $L^\rightarrow(\vec{x}_o, \hat{\omega}_o)$ leaving the surface at a point \vec{x}_o in direction $\hat{\omega}_o$ to the incident flux $\Phi^\leftarrow(\vec{x}_i, \hat{\omega}_i)$ arriving at a point \vec{x}_i from a direction $\hat{\omega}_i$. The unit of S is $[m^{-2} \cdot sr^{-1}]$.

Given the incoming radiance $L^\leftarrow(\vec{x}_i, \hat{\omega}_i)$ and the BSSRDF S for an object, the outgoing radiance $L^\rightarrow(\vec{x}_o, \hat{\omega}_o)$ at a point \vec{x}_o in direction $\hat{\omega}_o$ can be computed by integrating over the whole surface A and all incoming directions Ω

$$L^\rightarrow(\vec{x}_o, \hat{\omega}_o) = \int_A \int_\Omega L^\leftarrow(\vec{x}_i, \hat{\omega}_i) S(\vec{x}_i, \hat{\omega}_i; \vec{x}_o, \hat{\omega}_o) \langle \hat{n}_i \cdot \hat{\omega}_i \rangle d\hat{\omega}_i d\vec{x}_i. \quad (2.2)$$

The BSSRDF varies also with the wavelength λ of the incoming light. Apart from multi-spectral applications, this dependence is in computer graphics often approximated by defining and evaluating the BSSRDFs separately per color channel.

To include *fluorescence* and *phosphorescence* where light within a range of wavelengths is absorbed and re-emitted at some later time in a different range of wavelengths would require additional parameters.³

Approximations for Scattering Materials

A BSSRDF is often used to model the light transport in scattering materials such as skin or marble. The photons interact in this case with the object's material. After multiple scattering events, their propagation direction does not depend any more on $\hat{\omega}_i$ [Ishimaru78]. It is therefore sometimes convenient to split S into

³Fluorescent materials re-emit the light no later than 10 ns after absorption whereas phosphorescent re-emit the light after more than 10 ns – sometimes even after hours.

a sum of a single scattering contribution $S^{(1)}$ and a multiple scattering component S_d

$$S(\vec{x}_i, \hat{\omega}_i; \vec{x}_o, \hat{\omega}_o) = S^{(1)}(\vec{x}_i, \hat{\omega}_i; \vec{x}_o, \hat{\omega}_o) + S_d(\vec{x}_i, \hat{\omega}_i; \vec{x}_o, \hat{\omega}_o)$$

in order to treat both of them separately [Jensen01c].

2.4.3 Bidirectional Reflectance-Distribution Function (BRDF)

Without subsurface light transport, all light arriving at an object's surface is either reflected or absorbed at the point of incidence. This behavior can be described by the *bidirectional reflectance-distribution function* (BRDF) [Nicodemus77] which simplifies the BSSRDF by setting $\vec{x}_i = \vec{x}_j$. The BRDF is defined as

$$f_r(\vec{x}; \hat{\omega}_i, \hat{\omega}_o) := \frac{dL^\rightarrow(\vec{x}, \hat{\omega}_o)}{dE^\leftarrow(\vec{x}, \hat{\omega}_i)}. \quad (2.3)$$

The BRDF f_r is the ratio of reflected radiance $L^\rightarrow(\vec{x}, \hat{\omega}_o)$ leaving the surface at a point \vec{x} in direction $\hat{\omega}_o$ to the irradiance $E^\leftarrow(\vec{x}, \hat{\omega}_i)$ arriving at the same point \vec{x} from a direction $\hat{\omega}_i$. The unit of f_r is $[sr^{-1}]$.

Given the incident radiance $L^\leftarrow(\vec{x}, \hat{\omega}_i)$ and the BRDF f_r for an object, the outgoing radiance $L^\rightarrow(\vec{x}_o, \hat{\omega}_o)$ at \vec{x}_o in direction $\hat{\omega}_o$ can be computed by integrating over all incoming directions Ω

$$L^\rightarrow(\vec{x}, \hat{\omega}_o) = \int_{\Omega} L^\leftarrow(\vec{x}, \hat{\omega}_i) f_r(\vec{x}; \hat{\omega}_i, \hat{\omega}_o) \langle \hat{n}_i \cdot \hat{\omega}_i \rangle d\hat{\omega}_i. \quad (2.4)$$

Simplifications for Isotropic and Homogeneous Materials

The BRDF in its most general form is a six-dimensional function: Two parameters fix the position \vec{x} on the surface. The directions of the incident and exitant ray require another two parameters each. If an object has uniform reflection behavior at all surface positions (*homogeneous material*), the BRDF can be represented by a four-dimensional function $f_r(\hat{\omega}_i, \hat{\omega}_o)$.

For *isotropic materials*, incoming and outgoing directions $\hat{\omega}_i = (\theta_i, \phi_i)$ and $\hat{\omega}_o = (\theta_o, \phi_o)$ can be rotated around the surface normal without change:

$$f_r(\vec{x}; \theta_i, \phi_i, \theta_o, \phi_o) = f_r(\vec{x}; \theta_i, \phi_i + \delta, \theta_o, \phi_o + \delta).$$

An isotropic BRDF $f_r(\vec{x}; \theta_i, \theta_o, \phi_i - \phi_o)$ is therefore a five-dimensional function, and a homogeneous and isotropic BRDF $f_r(\theta_i, \theta_o, \phi_i - \phi_o)$ is a three-dimensional function.

Physical Properties of a BRDF

In order to be physically valid, a BRDF must fulfill the *Helmholtz reciprocity rule* which states that the roles of incident and reflected energy can be exchanged:

$$f_r(\vec{x}; \hat{\omega}_i, \hat{\omega}_o) = f_r(\vec{x}; \hat{\omega}_o, \hat{\omega}_i).$$

Furthermore, a BRDF must not reflect more energy than it receives, i.e., it must *conserve energy*:

$$\int_{\Omega} f_r(\vec{x}; \hat{\omega}_i, \hat{\omega}_o) \langle \hat{n}_i \cdot \hat{\omega}_i \rangle d\hat{\omega}_o \leq 1 \quad \forall \hat{\omega}_i \in \Omega.$$

Note that a material may well absorb part of the energy and reflect less energy than it receives.

Reflectance

While the BRDF as defined in Equation 2.3 is a derivative quantity that can take on arbitrary large values, the *reflectance* $\rho(\vec{x})$ is the ratio of reflected to incident flux [Nicodemus77]:

$$\rho(\vec{x}) := \frac{d\Phi^{\rightarrow}(\vec{x})}{d\Phi^{\leftarrow}(\vec{x})} = \frac{\int_{\Omega} \int_{\Omega} L^{\leftarrow}(\vec{x}, \hat{\omega}_i) f_r(\vec{x}; \hat{\omega}_i, \hat{\omega}_o) d\hat{\omega}_i d\hat{\omega}_o}{\int_{\Omega} L^{\leftarrow}(\vec{x}, \hat{\omega}_i) d\hat{\omega}_i}.$$

$\rho(\vec{x})$ is the fraction of light that is reflected by a surface. The remaining part is either transmitted or absorbed. For non-emitting surfaces, $\rho(\vec{x})$ must be in the range $[0..1]$ to ensure energy conservation.

2.4.4 Reflection Models

Real objects can have a complex reflection behavior that can in the most general case only be modeled by a tabulated BRDF that explicitly gives the reflection coefficients for a large number of configurations. For many materials, the reflection behavior can however be approximated by various reflection models.

Diffuse Reflection

A surface with ideal diffuse reflection reflects incident light equally in all directions, i.e., the BRDF is constant:

$$f_r(\vec{x}_i, \hat{\omega}_i, \hat{\omega}_o) = k_d = \frac{\rho_d}{\pi}.$$

For this so called *Lambertian reflection*, outgoing radiance depends only on k_d and the area foreshortening given by the incident polar angle θ_i :

$$L^{\rightarrow}(\vec{x}, \hat{\omega}_o) = k_d \int_{\Omega} L^{\leftarrow}(\vec{x}, \hat{\omega}_i) \cos \theta_i d\hat{\omega}_i.$$

Glossy Reflection

The Phong reflection model [Phong75] was historically one of the most popular shading models. It consists of a diffuse term and a specular lobe. Unfortunately, it is neither energy conserving nor reciprocal and cannot model the reflection behavior of most real materials. Several improvements to the Phong model have been proposed such as the Blinn-Phong model [Blinn77]:

$$f_r(\vec{x}, \hat{\omega}_i, \hat{\omega}_o) = k_d + k_s \frac{\langle \hat{n} \cdot \hat{h} \rangle^n}{\langle \hat{n} \cdot \hat{\omega}_i \rangle}.$$

The halfway vector $\hat{h} = (\hat{\omega}_i + \hat{\omega}_o)/2$ lies inbetween the incoming and outgoing light direction. The specular exponent n controls the width of the specular lobe and therefore the extent of the highlights.

In Chapter 5, we use the Lafortune model [Lafortune97]

$$f_r(\vec{x}, \hat{\omega}_i, \hat{\omega}_o) = k_d + \sum_j [C_{x,j} \cdot \hat{\omega}_{i,x} \hat{\omega}_{o,x} + C_{y,j} \cdot \hat{\omega}_{i,y} \hat{\omega}_{o,y} + C_{z,j} \cdot \hat{\omega}_{i,z} \hat{\omega}_{o,z}]^{n_j}.$$

to model the reflection properties of real world objects. Similar to the Blinn-Phong model, the Lafortune model consists of a diffuse term and a set of specular lobes. The lobe is described by a weighted inner product of the directions $\hat{\omega}_i$ and $\hat{\omega}_o$ that is risen to the n th power. Depending on the coefficients C , a lobe can be forward-reflective ($C_{x,j} < 0$), retro-reflective ($C_{x,j} > 0$), or anisotropic ($C_{x,j} \neq C_{y,j}$). The ratio of $C_{x,j}$ or $C_{y,j}$ to $C_{z,j}$ determines the off-specularity of the lobe.

2.4.5 Reflection Models for Subsurface-Scattering Material

Unlike for BRDFs, general BSSRDF models that can be used to describe the reflection properties of subsurface-scattering material are rare. Jensen et al. [Jensen01c] introduced a model for the diffuse light propagation S_d in translucent materials to computer graphics that is now widely used.

The model assumes an infinite half-space of homogeneous, optically thick material with given material properties (reduced scattering coefficient σ'_s , absorption coefficient σ_a , relative index of refraction η – see Ishimaru [Ishimaru78] and




Figure 2.9: *The geometry of the setup for the BSSRDF model of Jensen et al. [Jensen01c].*

$$\begin{aligned}
 z_r &= 1/\sigma'_t \\
 z_v &= z_r + 4AD \\
 d_r &= \|\vec{x}_r - \vec{x}_o\|, \text{ with } \vec{x}_r = \vec{x}_i - z_r \hat{n}_i \\
 d_v &= \|\vec{x}_v - \vec{x}_o\|, \text{ with } \vec{x}_v = \vec{x}_i + z_v \hat{n}_i \\
 A &= \frac{1 + F_{dr}}{1 - F_{dr}} \\
 F_{dr} &= -\frac{1.440}{\eta^2} + \frac{0.710}{\eta} + 0.668 + 0.0636\eta \\
 D &= 1/3\sigma'_t \\
 \sigma_{tr} &= \sqrt{3\sigma_a\sigma'_t} \\
 \sigma'_t &= \sigma_a + \sigma'_s \\
 \alpha' &= \sigma'_s/\sigma'_t
 \end{aligned}$$

Figure 2.10: *Additional definitions used in the BSSRDF model of Jensen et al. [Jensen01c]. See text for more details.*




Figure 2.11: The BSSRDF model of Jensen et al. [Jensen01c] evaluated for marble. The graphs show the diffuse reflectance due to subsurface scattering R_d for a measured sample of marble with $\sigma'_s=(2.19,2.62,3.00)$, $\sigma_a=(0.0021,0.0041,0.0071)$ and $\eta=1.5$ (parameters extracted from [Jensen01c]). $R_d(r)$ indicates the radiosity at a distance r in a plane, due to unit incident power at the origin.

Jensen et al. [Jensen01c] for a description of these parameters and a method to recover them for real materials). Light propagation is modeled by a dipole approximation as follows:

$$S(\vec{x}_i, \hat{\omega}_i; \vec{x}_o, \hat{\omega}_o) = \frac{1}{\pi} F_t(\eta, \hat{\omega}_i) R_d(\vec{x}_i, \vec{x}_o) F_t(\eta, \hat{\omega}_o)$$

$$R_d(\vec{x}_i, \vec{x}_o) = \frac{\alpha'}{4\pi} \left[z_r(1 + \sigma_{tr}d_r) \frac{e^{-\sigma_{tr}d_r}}{d_r^3} + z_v(1 + \sigma_{tr}d_v) \frac{e^{-\sigma_{tr}d_v}}{d_v^3} \right]$$

Figure 2.9 shows the geometrical configuration of the setup and Figure 2.10 defines the remaining parameters. An example evaluation of the model using measured parameters for marble taken from [Jensen01c] is given in Figure 2.11.

2.5 Rendering Equation

The rendering equation in it's most basic form

$$L^\rightarrow(\vec{x}, \hat{\omega}) = L_e(\vec{x}, \hat{\omega}) + L_r(\vec{x}, \hat{\omega}) \quad (2.5)$$

states that the outgoing radiance $L^\rightarrow(\vec{x}, \hat{\omega})$ is equal to the sum of emitted radiance $L_e(\vec{x}, \hat{\omega})$ and reflected radiance $L_r(\vec{x}, \hat{\omega})$. It is the mathematical foundation for all rendering algorithms.

If the reflection of all objects in a scene can be described by a BRDF, we can insert Equation 2.4 in the above definition to get [Kajiya86]:

$$L^\rightarrow(\vec{x}, \hat{\omega}) = L_e(\vec{x}, \hat{\omega}) + \int_{\Omega} L^\leftarrow(\vec{x}, \hat{\omega}_i) f_r(\vec{x}; \hat{\omega}_i, \hat{\omega}) \langle \hat{n} \cdot \hat{\omega}_i \rangle d\hat{\omega}_i.$$

If the rendered scene contains objects with non-local reflection properties, we need to revert to using a BSSRDF to describe the reflection properties. Inserting Equation 2.2 results in:

$$L^\rightarrow(\vec{x}, \hat{\omega}) = L_e(\vec{x}, \hat{\omega}) + \int_A \int_{\Omega} L^\leftarrow(\vec{x}_i, \hat{\omega}_i) S(\vec{x}_i, \hat{\omega}_i; \vec{x}, \hat{\omega}) \langle \hat{n}_i \cdot \hat{\omega}_i \rangle d\hat{\omega}_i d\vec{x}_i.$$

The double integral increases the effort considerably but allows for taking effects such as subsurface scattering into account.

2.6 Inverse Rendering

Given a complete description of a scene including the illumination and reflection properties, the rendering equation (Equation 2.5) can be used to synthesize an image of a scene by computing the radiance per pixel. The goal of *inverse rendering* is to reverse this process and to derive a (partial) scene description from a set of input images.

Given an image (e.g., captured by a digital camera) with pixel value p_{uv} for a pixel at position (u, v) , we can compute the corresponding radiance $L(\vec{x}_{uv}, \hat{\omega}_{uv})$ using a set of geometric and photometric calibration procedures (see Section 3.2 for more details). $(\vec{x}_{uv}, \hat{\omega}_{uv})$ define a ray leaving from a scene object's surface and hitting the sensor element corresponding to pixel p_{uv} . The scene description can then in principle be recovered by minimizing the following error function:

$$E_{uv} = [L(\vec{x}_{uv}, \hat{\omega}_{uv}) - [L_e(\vec{x}_{uv}, \hat{\omega}_{uv}) + L_r(\vec{x}_{uv}, \hat{\omega}_{uv})]]^2. \quad (2.6)$$

If we assume no distant light transport inside objects, i.e, if the reflection behavior can be described by a BRDF, this results in:

$$E_{uv} = [L(\vec{x}_{uv}, \hat{\omega}_{uv}) - [L_e(\vec{x}_{uv}, \hat{\omega}_{uv}) + \int_{\Omega} L^\leftarrow(\vec{x}_{uv}, \hat{\omega}_i) f_r(\vec{x}_{uv}; \hat{\omega}_i, \hat{\omega}_{uv}) \langle \hat{n}_{uv} \cdot \hat{\omega}_i \rangle d\hat{\omega}_i]]^2 \quad (2.7)$$

Inverse rendering algorithms can be classified using three categories which we describe in the following sections.

2.6.1 Photometric Stereo

Photometric stereo refers to the problem of reconstructing the geometry of a scene given a set of images with known illumination and has been extensively studied in the area of computer vision [Woodham81]. Given known lighting and reflection properties f_r , Equation 2.7 can be used to estimate a normal \hat{n}_{uv} for each pixel. The resulting normal map contains the spatial derivative of the geometry and can be integrated to obtain a smooth geometry.

Photometric stereo algorithms frequently assume that the target object is illuminated by a single parallel or point light source and has Lambertian reflection behavior. Surfaces with more complex reflection behavior can also be acquired using a corresponding reflection model.

2.6.2 Inverse Lighting

Inverse lighting tries to recover information about the incident illumination from acquired images. Schoeneman et al. [Schoeneman93] determine the intensity of discrete light sources to achieve a given lighting scenario while Kawai et al. [Kawai93] additionally take different light source types into account. Several authors addressed the problem to recover distant illumination from observed lighting effects [Marschner97, Marschner98, Ramamoorthi01].

In Chapter 6, we present a method to acquire the near-field of luminaires which employs a prefiltering step that selects different portions of the emitted light field and captures the projection onto a diffuse surface. The prefiltering step helps avoiding many of the classical problems of inverse lighting.

2.6.3 Inverse Reflectometry

The goal of *inverse reflectometry* is to recover the reflection properties of objects in a scene. Both, the BRDF acquisition algorithm employed in Chapter 5 and the acquisition method for translucent objects presented in Chapter 7 belong into this category. More details and related work are described in the respective previous work sections (Sections 5.1 and 7.1).

Note that these three problem categories can also be addressed together if several parameters of the inverse problem such as incident lighting and the reflection properties are unknown. An extensive survey of inverse lighting and inverse reflectometry techniques can be found in [Patow03].

Chapter 3

Acquisition Facilities

In this chapter, we describe some basic hardware and software systems that we used as tools to perform the measurements our models are based on. This includes a 3D scanning system and the processing software used to acquire the geometry models presented in this thesis. We furthermore summarize basic properties of digital camera systems and discuss the strengths and weaknesses of the actual cameras used. We discuss the lighting systems including the measurement environment and give an overview over the robotic measurement facility employed for some of the measurements.

Although these tools are assumed to be given for the projects described in this dissertation, their performance and accuracy is crucial for the quality of the results. We therefore include also measurements and benchmarks where appropriate.

3.1 Geometry Acquisition

All geometry models used in this dissertation except for the bird model were acquired using a Minolta VI-900 or VI-910 laser range scanner. These active optical range scanning systems sweep a laser line over the target object and observe the reflection of the line with a video camera. The position of the reflected laser line in the recorded image depends then on the geometry of the setup and the distance of the illuminated points on the target object from the 3D scanner (see Figure 3.1). The laser line defines a plane P in space whereas each pixel in an image captured by the video camera defines a ray \vec{r} in space. If the reflection of the laser line on the target object is visible in a pixel, the intersection of P and \vec{r} yields a single 3D point on the surface of the target object. A single complete scan returns a depth map where each valid pixel contains a distance value from the 3D scanner to the object.

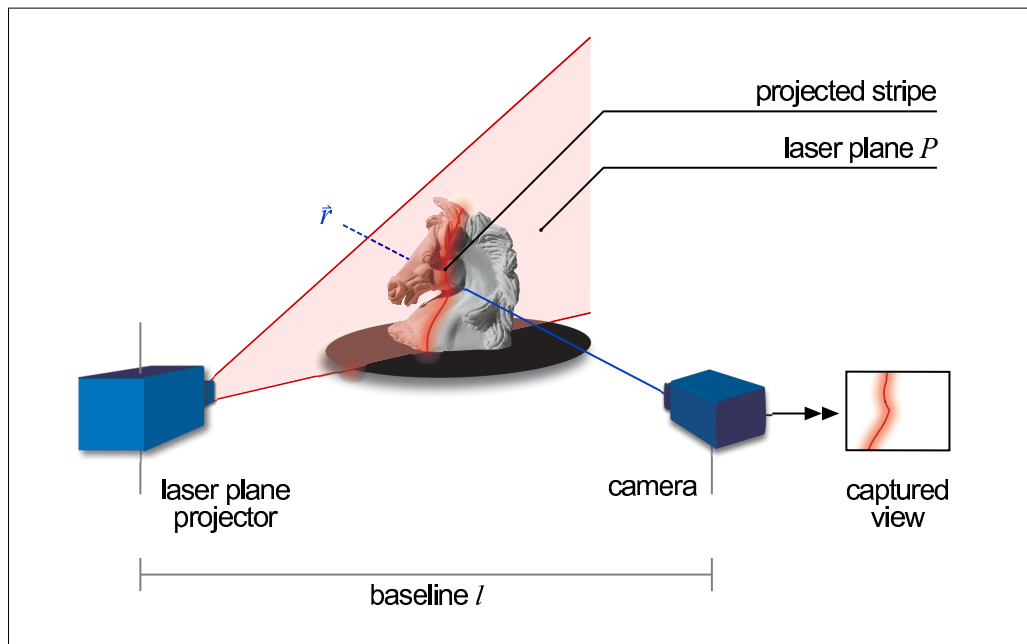


Figure 3.1: *The principle of operation of the 3D laser range scanner used to acquire geometry models. Valid surface points are defined by the intersection of the laser plane P with a ray \vec{r} (corresponding to positions in the captured image).*

The bird model was acquired using a Steinbichler Tricolite structured light 3D scanner. The scanner is based on the same measurement principle but – instead of sweeping a laser line – it projects a series of stripe patterns onto the object to define the intersection plane P .

3.1.1 Performance and Accuracy

The underlying principle of active triangulation-based optical range scanners is the stereo vision principle [Klette98]. (The laser projector could be replaced by a second camera leading, however, to the in general difficult stereo correspondence problem.) A surface point can therefore only be scanned successfully if it is visible from both, the position of the laser projector and the position of the video camera. Concave objects may lead to occlusions which often result in holes in the depth maps and are therefore hard to scan. The larger the baseline (the distance between laser projector and video camera) of a triangulation-based 3D scanning system, the more difficult concave objects are to scan. There is however a trade-off as a larger baseline leads to higher accuracy in the resulting models.

A laser range scanner relies on the constant, near-diffuse local light reflection on the surface of an object. Unfortunately, many of our test objects do not fulfill



Figure 3.2: *The alabaster horse. Left and middle: Photographs of the sculpture with and without subsurface scattering under identical illumination conditions. The sculpture was covered with fine dust to “turn off” subsurface scattering. Much surface detail is masked by the translucency and active optical range scanning systems have grave problems to scan the translucent object. Right: Rendering of the acquired geometry model (100 k triangles) with flat shading.*

this requirement leading to noisy or incomplete depth maps: Colorful objects can absorb a large portion of the red laser light so that the reflected signal becomes too weak. Likewise, specular surfaces reflect hardly any light in off-specular directions so that the laser line may not be recorded by the video camera depending on surface orientation. Even worse, translucent objects exhibit non-local light reflection, i.e., light enters and leaves the surface at possibly distinct locations. In order to acquire good quality models even for difficult to scan objects, some of the objects were sprayed with white dust (Helling Met-L-Chek Developer D70) in order to achieve a surface with Lambertian reflection. Figure 3.2 shows an example of a translucent alabaster horse sculpture. The translucent behavior masks many geometric details that can be easily seen when the object is covered with dust.

We will present in Chapter 4 a method based on the concept of a modulation transfer function (MTF) that allows us to estimate the accuracy of a 3D scanning system in frequency space.

3.1.2 3D Geometry Processing Pipeline

A single 3D scan of an object is usually not sufficient to cover its whole surface. To achieve complete coverage, the object is scanned several times from different directions (possibly using a turntable to rotate the object in front of the 3D scanner). The individual scans are then aligned and combined into a single data set. Based on this point cloud, a surface triangulation is constructed and remaining holes are filled manually or semi-automatically. Various postprocessing opera-

tions can be applied to the final triangle mesh, e.g., to remove noise or to reduce the complexity of the mesh.

The geometry models used in this thesis were assembled and processed using existing software at the Max-Planck-Institut für Informatik as well as a variety of commercial software tools (e.g., Minolta Polygon Editing Tool, Steinbichler COMETplus, INUS Rapidform, Innovmetric Polyworks, TowerGraphics Easy3DScan, and ViHAP3D's mesh processing tools). A more detailed overview over tools and methods for geometry processing can be found in [Kobbelt00].

3.2 Digital Camera Systems

A digital camera can be seen as a measurement device consisting of millions of input sensors that convert luminance arriving at the pixels into digital values. The general assumption in image-based modeling applications is that after some calibration steps the acquired images are a faithful representation of reality.

We will in this section first discuss the general setup of a digital camera concentrating on the lens system and the sensor chip. We summarize then common calibration techniques for geometric and photometric camera calibration.¹ We also mention techniques to judge the accuracy of acquired images. We finally present the cameras used for this work and discuss their specific properties.

3.2.1 Camera Setup and Components

The two most important components of a digital camera are the lens system and the imaging sensor that ideally should be tailored to each other. Several mechanical components such as the aperture and the shutter as well as the camera electronics and host software influence the system performance, too, but will not be discussed in detail.

Imaging Sensor

An imaging sensor converts irradiance incident on a light sensitive area into an electrical signal whose exact nature depends on the sensor technology. Imaging sensors in digital cameras consist generally of a one or two dimensional array of sensor elements (sensels) or picture elements (pixels).

¹There is some confusion about the meaning of the word calibration: In some areas such as color science *calibration* means to set a device into a defined state. In an additional *profiling* step, this state is characterized and the characterization is later taken into account in order to achieve a defined behavior of a device. In computer graphics, *calibration* is tuning a general model of the physical device in order to best describe the specific instance on the device. Unless otherwise noted, we use this second semantics in this thesis.



Figure 3.3: The Bayer pattern [Bayer76] for color image capture. Left: Arrangement of the individual color sensors. Middle: Color interpolation artifacts in an image acquired with a Bayer pattern sensor (Jenoptik ProgRes C14, single shot, fast reconstruction). Right: Ground truth image acquired using 4-shot mode. Note also the visually higher resolution of the ground truth image.

In the case of a *charge coupled device (CCD)* sensor, light is absorbed and creates free electrons that are stored inside the sensor element that serves as capacitor (e.g., [Janesick01]). At the end of an exposure, the collected charge is shifted towards one end of the sensor using an analog shift register and digitized using an analog-digital converter. Due to the physical principles involved, CCD sensors are highly linear, i.e., the amount of charge collected is proportional to the irradiance. The main performance limitations are leakage currents caused by chip defects and thermal noise. Cooling the sensor will greatly improve its performance as the noise doubles with every increase in temperature of about 6 Kelvin. Both, the Kodak DCS 560 and the Jenoptik ProgRes C14 digital cameras (see Section 3.2.4) employ CCD imaging sensors.

The dynamic range, i.e., the ratio of brightest to darkest non-zero recordable intensity values of a CCD sensor is not only limited by noise but also by the well capacity. A sensel of the Sony ICX085AK chip used in the Jenoptik ProgRes C14 camera can typically store about 25,000 electrons so that the dynamic range even of a noise free sensor is limited to about 1:25,000 requiring less than 15 bit for digitization. Very bright spots in an image can also lead to blooming artifacts if the irradiance creates a strong current and charge is overflowing from one sensel to neighboring sensels.

In contrast to CCD sensors, *complementary metal oxide semiconductor (CMOS)* sensors are based on the technology used for modern random access memory chips, processors, or analog circuitry. They are more versatile and flexible but often suffer from stronger image noise. In particular, modern CMOS cameras can achieve a dynamic range of more than 140 dB which is often required for computer graphics applications (see, e.g., Kang et al. [Kang03] for an

overview of camera systems with extended dynamic range). The Silicon Vision Lars III camera is such a camera with a CMOS imaging sensor.

Color image acquisition is difficult as a single sensor element can capture only a single value (defined by the convolution of its spectral sensitivity with the spectrum of the incident light) but not a full set of RGB tristimulus values. The most common solution is to cover individual sensels with color filters arranged in a Bayer pattern [Bayer76] (see Figure 3.3) and to capture individual color components at different spatial locations. Spatial resolution is in this case traded for color information. Tristimulus color values (RGB) are reconstructed for each pixel using information from neighboring sensels. This reconstruction can introduce color artifacts so that the pixel values do not necessarily correspond to the irradiance incident at the pixel location.

Alternatively, the image can be simultaneously captured by 3 sensors with corresponding color filters (3 chip cameras) or at different times using, e.g., a color filter wheel or tunable filters. Only recently, a new technology was introduced that makes use of the different penetration depths of light into the chip and stacks three sensels on top of each other [Lyon02]. It allows to capture tristimulus color values without spatial or temporal aliasing artifacts.

Lens System

The lens system projects an image of a scene onto the imaging sensor (see Figure 3.4). A lens with focal length f focuses a point at distance g onto a point at distance b according to

$$\frac{1}{f} = \frac{1}{g} + \frac{1}{b}$$

leading to a sharp image. At a point b' farther away, the image is blurred within a circle of confusion whose diameter d is proportional to the radius of the aperture.

While achieving a sharp image is desirable for an analog, film-based camera system, it can be harmful for the sampled imaging system of a digital camera: If the light sensitive portion of a sensor element covers not its full area, a small feature could be focused on a light-insensitive part and would not be captured (see Figure 3.5). It has also the potential of severe aliasing if the fill factor, the percentage of light sensitive area of a pixel, is much smaller than 100 %. The fill factor can be increased by adding small lenses on top of the sensor elements as shown in Figure 3.5. This is however not sufficient for color sensors with Bayer pattern [Bayer76] where even a small feature should be registered by red, green, and blue color sensors in order to minimize interpolation artifacts. An additional antialiasing filter can be used to ensure this by slightly blurring the projected image.

To achieve good image quality, it is therefore important that the lens system matches the imaging sensor. A way to check this is to determine the modulation transfer function (MTF) [Williams99, Burns01] of an imaging system that describes its properties in frequency space.

The lens system can furthermore lead to a variety of artifacts such as chromatic aberration (the focal length f varies for different wave lengths of light leading to color artifacts at image edges), geometric distortions (the acquired image is distorted compared to an image generated using a pin-hole projection), or diffraction at the aperture blades (visible as star-shaped artifacts around bright features). It is therefore mandatory to employ high quality lenses and apply where possible correction techniques in order to acquire high quality data.

3.2.2 Geometric Calibration

The lens system of a digital camera projects an image of a scene onto the imaging sensor. The mapping between points in space and 2D image coordinates depends on the position and orientation of the camera (so called *extrinsic parameters*) and the *intrinsic parameters* such as focal length, distortion parameters, and the position of the optical axis on the image plane (principal point). Various authors such as [Tsai87, Heikkilä97, Zhang99] proposed camera models and calibration techniques to estimate extrinsic and intrinsic camera parameters from a set of images of a calibration target.

For this thesis, we used a calibration toolbox [Bouguet03] (which is based on the estimation technique by Zhang [Zhang99] and uses a camera model similar to the model of Heikkilä and Silvén [Heikkilä97]) to determine the actual focal length $fc = (fc(1), fc(2))^T$ and the distortion coefficients $kc = (kc(1), kc(2))^T$ of a lens. The camera model projects a 3D point $P = (P_x, P_y, P_z)^T$ onto a 2D image coordinate $p_p = (x_p, y_p)^T$ as follows:

$$\begin{pmatrix} x_p \\ y_p \end{pmatrix} = d(r^2) \begin{pmatrix} fc(1) \cdot P_x/P_z \\ fc(2) \cdot P_y/P_z \end{pmatrix} + \begin{pmatrix} cc(1) \\ cc(2) \end{pmatrix}$$

We assume that the lens is centered with respect to the image sensor and set the center of projection $cc = (cc(1), cc(2))^T$ to the image center. The radial distortion $d(r^2)$ is modeled by

$$d(r^2) = 1 + kc(1)r^2 + kc(2)r^4$$

and depends on $r^2 = (P_x/P_z)^2 + (P_y/P_z)^2$. We assume that the tangential distortion is not significant.

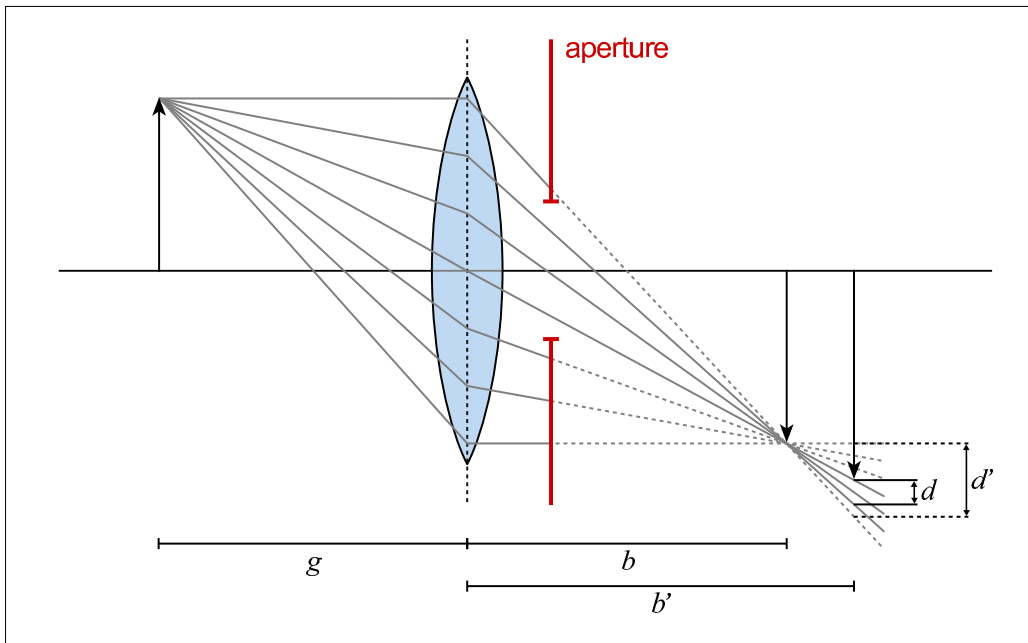


Figure 3.4: Projection by a lens system with adjustable aperture. See text for a more detailed description.




Figure 3.5: Projection onto a sampled imaging sensor with light insensitive areas (fill factor < 100 %). Left: A small feature is either focused on a light sensitive or insensitive area leading potentially to aliasing. Middle: A slightly out-of-focus projection reduces the aliasing potential but will still result in a focused digital image if the circle of confusion is smaller than the sensel spacing. Right: Some cameras include an antialiasing filter to prefilter in the optical domain.

3.2.3 Photometric Calibration

The pixel values stored in a digital image are normally not linearly related to intensity values in the captured scene. The sRGB standard [IEC99] defines for example a gamma value of 2.2 in order to avoid quantization artifacts, i.e., a 2.2-power function needs to be applied to an sRGB color value in order to get a value linearly related to intensity.² If linear values are needed it is therefore mandatory to recover the response curve (also called transfer function or opto-electronic conversion function (OECF) [ISO99]) that performs the mapping from digital image values to linear intensity values.

A simple method to recover the response function is to acquire an image of a uniformly illuminated test target with patches of known reflectance and to compute the response curve from the digital image values of the patches. For the purpose of this dissertation, we apply however high-dynamic range imaging methods that also recover the response curve.

High-Dynamic Range Imaging

Standard digital cameras can only capture intensity values within a quite limited dynamic range: For a given exposure setting, too dark areas of a scene are usually mapped to zero (or vanish in the device's noise floor) while too bright areas are mapped to some maximum value. For both cases, the digital image values cannot be mapped back to linear intensity values and therefore are only of limited use for image-based measurement methods.

The principle of high-dynamic range (HDR) imaging [Debevec97, Robertson99, Robertson03] is therefore to acquire multiple images of the same scene with different exposure settings ensuring that at least one meaningful digital value is available for each pixel. The camera's response curve I is then computed from this image series. Using [Robertson99], a HDR image with linear floating point intensity values of pixel x_j at image position j is generated as weighted average of the linearized input images scaled according to their exposure times t_i :

$$x_j = \frac{\sum_i w_{ij} t_i I_{y_{ij}}}{\sum_i w_{ij} t_i^2}.$$

$I_{y_{ij}}$ is the linearized pixel value at position j of input image i . The weighting function w_{ij} is a Gaussian-like function that emphasizes pixel values with high confidence.

²Due to numerical considerations, the sRGB standard [IEC99] specifies more exactly that small values up to a threshold of 0.0031308 (on a scale of [0..1]) are kept linear while a gamma of 2.4 is applied to larger values.




Figure 3.6: *The ICC color management pipeline. Each device is linked to the common profile connection space by an ICC profile. Color data can thus be converted between arbitrary devices minimizing the introduced error.*

Color Calibration

The tristimulus color value (e.g., the RGB triplet) recorded for a color patch in a scene depends not only on the spectral reflectance of the patch but also on the spectrum of the illumination and on the spectral response of the imaging sensors. The most basic method to ensure correct color acquisition is therefore to perform a white balancing step. To this end, a neutral white or gray color patch is acquired under the given illumination and multiplicative scaling factors for the linearized color values are determined that lead to equal intensity values.

Color Management Systems

In graphical arts and the printing industry, color management systems based on the ICC profile system [ICC03] are routinely used to achieve consistent colors throughout the processing work flow.

The principal problem to be addressed is that all input and output devices usually have their own, uncalibrated color space, typically some form of RGB or CMYK. An ICC-based color management system provides a profile for each device that allows to convert between its color space and a common profile connection space (PCS) – either the CIE XYZ or the CIE Lab color space [CIE86]. This conversion takes also other properties of the devices such as their gamut (the range of possible colors) or the whitepoint into account. Figure 3.6 shows an overview over such a color management system. Many current operating systems have built-in color management systems or are at least able to use color management for individual applications.

A crucial part of any ICC-based color management system is the generation



Figure 3.7: *The GretagMacbeth Colorchecker DC used for ICC profile generation imaged under diffuse illumination.*

of appropriate profiles. For an input device such as a digital camera or a scanner, a test target with a set of known color patches such as the GretagMacbeth Colorchecker DC (see Figure 3.7) is imaged. The profile generation software analyzes this data and encodes the transformation between the color space of the input device and the PCS in an input profile. The core element of an input profile is in most cases a large lookup table which allows to encode a wide range of transformations.

There are different ways to obtain an output profile that characterizes a display or printing device. A high-end calibration requires displaying or printing a well-defined test pattern and measuring the emitted spectrum of each color patch using a spectrophotometer. By matching the intended color to the measured color, the ICC output profile is generated. This calibration routine is both time consuming and expensive since a dedicated hardware is required. There are also several low-end software solutions such as the Little CMS profiler [Saguer04] where the spectrophotometer is replaced by the user who has to compare colors and patterns yielding an output profile of less quality.

3.2.4 Camera Systems

We use three different digital cameras to complete the projects described in this dissertation – a professional photographic camera, a high-end color video camera for scientific applications, and a prototype implementation of a high-dynamic

range video camera.

Kodak DCS 560

The Kodak DCS 560 is a professional photographic SLR camera based on a Canon EOS-1N body. It contains a 3040×2008 pixel CCD sensor (Bayer pattern) that produces images with 12 bit color depth in a Kodak proprietary TIFF format and can be used with most Canon EF autofocus lenses. The camera can be remotely controlled via a IEEE 1394 FireWire connection using our own software based on an SDK provided by the manufacturer. A detailed technical specification can be found in the camera's manual [Kodak98].

When it was acquired in 1999, the camera provided superior image quality due to the high resolution of the imaging sensor and the high-quality lenses originally developed for film-based SLR cameras. The uncooled CCD sensor however suffers from severe noise artifacts for longer exposure times (starting at about 1 s exposure time). As the camera is tailored to professional photographers, it contains undisclosed image optimization software and provides no access to raw sensor data.

Jenoptik ProgRes C14

The Jenoptik ProgRes C14 video camera was designed for scientific photography and features a C-Mount lens connector. Its 1300×1030 pixel CCD sensor is cooled by a Peltier element and an additional fan to keep image noise at a minimum enabling exposure times of up to 10 min without visible noise. The Bayer pattern color sensor is mounted onto a piezo micro-scanning actuator and can be shifted by a whole pixel to sequentially capture per-pixel RGB information. Additional shifting by sub-pixel steps allows to increase image resolution up to 3900×3090 pixels. The camera records 14 bit intensity values with a typical dynamic range of 66 dB. It can be controlled remotely via IEEE 1394 FireWire and custom software using the provided SDK. The camera is usually used either with the high-resolution Jenoptik Lametar 2.8/25 mm lens or with a Schneider Cinegon 1.4/4.8 mm wide-angle lens.

The excellent noise characteristics make the C14 an ideal camera for low-light level applications. Static scenes can be captured with high resolution in multi-shot mode requiring no color interpolation. Dynamic scenes can however only be captured in single shot mode with up to 10 frames per second (fps). Scenes with a strong infrared signal might require an additional infrared cut-off filter to achieve good color fidelity.




Figure 3.8: Scene with a contrast of 1:3,700,000 captured with the Lars III camera. Left: HDR image with medium exposure. Right: Color-coded version (logarithmic) encoding the full dynamic range.

Silicon Vision Lars III

The Silicon Vision Lars III is a prototype implementation of a black-and-white high-dynamic range video camera (up to 30 fps) with C-Mount lens connector. Its 768×496 pixel CMOS sensor achieves a fill factor of 100% by placing the light sensitive photo diode on top of the per-pixel circuitry. Each sensel decides at up to 12 time steps (usually powers of 2) whether or not the exposure should be continued. The collected charge is digitized using 12 bit A/D converters leading to a linear dynamic range of 120 dB. Figure 3.8 shows a color-coded version of a HDR image captured with the Lars III camera.

Due to its prototype status, the camera requires extensive tuning using the supplied SDK and tuning software. The camera is integrated into the acquisition setup using custom software. We generally use the Jenoptik Lametar 2.8/25 mm lens with fully opened aperture to minimize artifacts introduced by the interplay of lens system and HDR scene.

Lars III Photometric Calibration

Performing a photometric calibration to evaluate the linearity of a HDR camera is difficult as it requires both, a HDR calibration target as well as suitable software. Furthermore, the exposure time of the Lars III video camera cannot be changed in a controlled way.

To compute its response curve, we therefore captured 3 video streams (483 images total) of a static scene. While one of the streams was captured without modifications, we added a neutral density filter (gray filter) with $1.41 \times$ and $10 \times$ extension factor, respectively, in the optical system for the other two streams. The extension factors are not powers of 2 in order to avoid interference with the camera's switching behavior. As shown in Figure 3.10, the scene contains an




Figure 3.9: Recovered response curve for the Lars III camera (horizontal axis: intensity, vertical axis: pixel values). Three of the 483 input image for the response curve recovery are shown in Figure 3.10.




Figure 3.10: Three of the input images used to recover the response curve of the Lars III camera (Figure 3.9). All images show the same scene consisting of a bright light source as well as specular and diffuse reflecting objects in order to achieve a good luminance distribution. For the middle and right image, the lens was covered with a $1.41\times$ extension factor gray filter (middle) and a $10\times$ extension factor gray filter (right), respectively, to emulate varying exposure times.

HMI light with a metallic reflector (Section 3.3.1) and other specular and diffuse reflecting objects to achieve a good luminance distribution.

The images were encoded in a 4+12 bit exponent-mantissa format and fed into a 16 bit version of an HDR recovery algorithm [Robertson99] to compute the response curve (Figure 3.9). It confirms the linear response but also shows artifacts correlated to the sensor's switching points. The non-linearities at high intensities are probably caused by intensity noise in the input images and by the small number of samples per intensity. The bend at low intensities is at least partially caused by a problem of the measurement methodology: The neutral density filters influence only the irradiance at the image sensor but do not change the exposure time. The dark current noise (charge collected in some pixels independent of irradiance) is consequently not influenced by the filters. As the camera contains a significant number of pixels with dark current noise of some 10s or 100s, the recovered response will be unreliable in this range.

3.3 Illumination and Environment

In this section, we first present the illumination systems used for the projects in this dissertation and give some details about our acquisition lab.

3.3.1 Point and Area Light Sources

Many measurement applications require control over or at least precise knowledge of illumination conditions. A uniform point light source which can be assumed to equally emit light in all directions from a single point in space is well suited for that purpose: All direct illumination hitting an object arrives from a single direction. This strongly simplifies the inversion of Equation 2.6 required for inverse rendering (see Section 2.6).

We included therefore in our lab a pair of K5600 Joker-Bug 800 light sources. Their 800 W HMI burners serve as a good approximation to a point light source with daylight spectrum. Additional diffusors transform them in area light sources in order to diffusely illuminate a scene.

3.3.2 Laser System

For the acquisition of translucent objects (see Chapter 7) we used a custom-built laser system. It consists of three individual lasers – a red diode laser and green and blue solid-state lasers with wavelengths of 635 nm, 532 nm, and 476 nm and 10 mW optical power per laser. All laser beams are fed into a single optical fiber to ensure exact alignment of the resulting beams. An adjustable collimator at the



Figure 3.11: *The acquisition lab at the Max-Planck-Institut für Informatik. A test object on the table is illuminated by two HMI light sources with different reflectors.*

other end of the fiber focuses the beam. A 2D galvanometer scanner deflects the beam with high precision. Care is taken in the whole laser system to reduce laser speckle to a minimum in order to avoid measurement artifacts.

The intensity of the lasers varies due to noise and shows a drift due to thermal effects leading to noise and inaccuracies in the measured data. These issues can be resolved in the acquisition stage by either locking the lasers' intensity or by monitoring and calibrating for the variations. The noise can also be suppressed by smoothing of the acquired data in later stages of a processing pipeline.

3.3.3 Environment

Unless noted otherwise, all acquisitions have been performed in the acquisition lab at the Max-Planck-Institut für Informatik. The lab's acquisition area consists of an $4.8\text{ m} \times 4.65\text{ m}$ sized room (see Figure 3.11). Its walls and ceiling are covered with black felt that diffusely reflects about 1.4 % of the incoming light. The floor is covered with black needle fleece carpet. This ensures that the influence of the environment on the measurements is reduced to negligible levels.




Figure 3.12: *The ACME measurement facility at the University of British Columbia configured for light source measurements.*

3.4 ACME Robotic Measurement Facility

Some of the light source measurements presented in Chapter 6 of this dissertation have been performed using the UBC Active Measurement Facility (ACME) [Pai99] at the University of British Columbia (UBC) in Vancouver, Canada. This robotic measurement platform consists of several linear and angular positioning stages that can be arranged into various configurations on a Newport optical table. ACME also includes a variety of sensors and actuators that however have not been used for our measurements. Application software can control the whole system via a Java API using a client-server architecture.

The light source measurements utilized only the linear translation stages which were used to precisely move a light source and a small filter relative to a projection screen. The positioning accuracy and repeatability of the translation stages is on the order of $\pm 0.00635 \text{ mm}$ which is clearly sufficient given the minimum step size of several millimeters employed in the measurements.

The whole ACME setup is located in a large measurement lab that is not tailored to light sensitive measurements. Furthermore, most of the system consists of shiny metal parts which could influence the measurements. The whole platform was therefore wrapped in black cloth to shut off ambient light. The actual measurement area with the light source, the projection screen and the recording camera was additionally enclosed in a second layer of black cloth that effectively

removed all stray light and reduced internal reflections to a minimum. Figure 3.12 shows the whole setup with partially removed cloth to make also the internal area accessible.

Chapter 4

Modulation Transfer Function Estimation for 3D Range Scanners

Numerous 3D range scanning devices utilizing many different acquisition technologies have been built in the last decades by industrial as well as academic developers [Bernardini02, Blais03, Chen00, Scopigno02]. Comparing their performance in terms of accuracy is however difficult as each system has its specific strengths and weaknesses so that accuracy depends heavily on the test targets and the test conditions [Chen00]. A system that can faithfully digitize a smooth surface of an object might not be able to acquire a sharp edge between two surfaces with sufficient accuracy. Beraldin et al. [Beraldin95] used therefore multiple test targets to determine the accuracy – the deviation from the exact shape – of range scanning systems for various conditions.

In contrast to previous work, we determine the accuracy of a range scanning system in terms of its spatial frequency response or its optical transfer function (OTF). The OTF determines the ability of the scanning system to reproduce different spatial frequencies present in an object. Sharp edges of an object to be scanned contain arbitrarily high spatial frequency components. A range scanning system which samples a finite number of surface points will never be able to reproduce a sharp edge faithfully unless it makes some assumptions about features present in the original scene. Whether the reconstructed edge looks blurry or shows aliasing artifacts depends on the OTF of the scanning system – especially on the modulation transfer function (MTF), the amplitude of the OTF.

Our work is inspired by the elegant and accurate slanted edge OTF measurement technique [Reichenbach91] for 2D imaging devices such as digital cameras or scanners and extends it into the third dimension. In principle, a single scan of a




Figure 4.1: *The setup used for test scans consisting of a Minolta VI-910 laser range scanner and the 90° calibration target.*

test target containing a straight sharp edge is sufficient to compute an estimate of the OTF. Once such a test target is available, a very simple acquisition setup such as depicted in Figure 4.1 is sufficient to perform this analysis.

4.1 Previous Work

We first review previous work on 3D range scanner accuracy and focus then on methods to determine the OTF of continuous and sampled systems.

4.1.1 Range Scanner Accuracy

Accuracy is commonly defined as the closeness of agreement between the result of a measurement and the value of the measurand [Taylor94]. Standard accuracy tests for 3D range scanners are often based on length measurement tests developed for coordinate measurement systems or for photogrammetry [Bernardini02, Keferstein98]. Hebert and Krotkov [Hebert92] for example use multiple measurements of targets with known ground truth distance and compute the RMS error. Most commonly, accuracy is specified as range accuracy or as a multiple of the standard deviation σ of a Gaussian fitted to the measurement results [Dalton98].

Chen et al. [Chen00] point out that these accuracy values can only be com-

pared if a standard measurement procedure is used including an exact specification of the test targets. The test targets should furthermore contain the same features such as curved surfaces or sharp edges that are encountered during the practical use of the scanner [Beraldin95].

4.1.2 OTF Measurement

Characterizing imaging devices by their OTF is an established technique since many decades – an overview can be found for example in [Williams99]. There are many techniques based on the observation of edge gradients, sine waves, square waves, or other features to determine the OTF of a continuous system. Using the Rayleigh criterion [Born99], the resolution of the system can be easily determined from the OTF [Williams01]. Direct application of these techniques to sampled systems is however difficult due to noise, undersampling, and the anisotropic sampling performed by most sampled imaging systems [Reichenbach91].

Reichenbach et al. [Reichenbach91] introduced a knife-edge technique for OTF measurement of sampled imaging systems where a superresolution edge profile is computed from an image of a slanted edge. Multiple scan lines are registered to each other and combined into a single edge profile to overcome the difficulties due to undersampling and anisotropic sampling problems. Noise can be suppressed by smoothing or by more advanced filtering techniques [Fischer94]. The method is now part of several ISO standards for resolution measurements such as [ISO00] but can also reveal many other characteristics of an imaging system [Burns01, Williams01, Okano98].

Miles et al. [Miles02] determined the OTF of an intensity image acquired by a ladar range imaging system. They did however not analyze the OTF of the associated range data.

In contrast to these approaches, we measure the modulation transfer function of a 3D scanning device using a single acquired range image. We use the MTF to derive an accuracy measure for range scanning systems that specifically takes the performance at object edges into account. Some work on this has already been reported by Paakkari and Moring [Paakkari93] who included a set of 3D bar gratings into their performance tests of range imaging devices and carried out some limited experiments to determine the frequency response of a range imaging sensor.

4.2 Measurement Principle

The optical transfer function (OTF) describes the spatial response of a linear system. It is commonly used to describe the performance of an optical sys-

tem [Williams99]. The OTF consists of the modulation transfer function (MTF) encoding the change in amplitude and the phase transfer function (PTF) encoding the phase difference:

$$\begin{aligned}\text{MTF}(s) &= |O(s)/I(s)| \\ \text{PTF}(s) &= \phi(O(s)) - \phi(I(s))\end{aligned}$$

where I is an input sine wave with frequency s and phase ϕ , O is the corresponding output sine wave with the same frequency but possibly different amplitude and phase. Using the principle of superposition for linear systems, an arbitrary input signal can be decomposed in its frequency components via Fourier transform, multiplied with the OTF, and transformed back using inverse Fourier transform to compute the system's response to the input signal.

4.2.1 Slanted Edge Analysis for 3D Range Scanners

The OTF of a sampled imaging system such as a digital camera can be computed using the knife-edge technique. While serving well for traditional analog systems, this technique has certain problems when dealing with sampled systems such as digital cameras. As one of these problems is undersampling, it is necessary to improve the spatial resolution of the device. The OTF can be measured beyond the Nyquist frequency of the device by placing an edge target slightly off-perpendicular to the scanning axes and averaging multiple scan lines (see Figure 4.2). Hereby whole dimension is sacrificed to overcome the Nyquist limit.

We are using a slightly slanted edge of a cube instead of a flat knife-edge target in order to measure the OTF of a 3D scanner. A step edge function corresponding directly to a knife-edge target would be mathematically better to analyze the frequency response. It introduces however the problem of mixed pixels [Hebert92] which is of interest for long range scanners but not necessarily for object acquisition.

The target edge should be as sharp and straight as possible and the two surfaces adjacent to the edge should be diffuse reflecting and flat. The angle of the edge can also influence the measurements as shown in Section 4.4.2.

The larger the distance to the target, the further away are the sample points from each other. Note that the Nyquist frequency can only be determined if the sample points are evenly spaced, i.e., the scan is not an arbitrary unstructured point cloud and if the scan shows a flat surface perpendicular to the scanning direction. If the scanning device returns points on a grid, it is however possible to compute the average distance between neighboring points in a certain region from which an estimate of the Nyquist frequency can be computed.

The focus of the device also affects the frequency response at the target edge. The edge should therefore be exactly in focus. If the scanner is equipped with an




Figure 4.2: Registration and superposition of scanned points. After sampling the slanted edge with a regular sampling grid (top), the points are registered along the edge location (middle), and projected onto a plane to form a superresolution edge profile (bottom).

autofocus system, the focus setting should be verified because focusing exactly on a sharp edge can be difficult.

4.2.2 Data Processing

Our data processing starts once a single scan of a slanted edge has been acquired. The scan is saved as a range image retaining the individual scan lines. Next, it is automatically segmented by a histogram based technique using local normals. Two least-squares planes are fitted to the reliable data points on both sides of the edge which excludes the central region near the edge and noisy or incompletely sampled areas. The intersection of the planes is used to register the individual scan lines to one of the coordinate axes of the scanner's imaging sensor by a simple shearing operation. All data points are then projected onto a plane perpendicular to the registered line of intersection to form the superresolution edge profile (see Figure 4.2). Likewise, a "perfect" edge profile is computed from the sheared version of the fitted planes.

4.2.3 Data Analysis

The two generated profiles are then resampled perpendicular to the bisecting line of the perfect profile into 2^n bins centered around the peak of the profile using twofold oversampling compared to the estimated Nyquist frequency. Hereby some samples are discarded at both ends of the profiles. The ends of the binned profiles are moved to the horizontal axis and the profiles are componentwise multiplied with a modified Welch window [Press94]. Then both profiles are rotated around one end point by 180° in order to continue them periodically. The goal of these operations is that the only high frequency features of the profiles are the two peaks. Continuation errors around the zero line are suppressed. Figure 4.3 shows the two profiles.

Now the Fourier transform can be applied to each profile. If the whole process contains no numerical errors, the even Fourier coefficients should all be zero and we therefore neglect them when calculating the MTF estimate. The MTF estimate is then calculated by dividing the absolute value of the odd Fourier coefficients of the scanned profile by the corresponding Fourier coefficients of the perfect profile. Figure 4.5 shows plots of the Fourier coefficients and the calculated MTF estimate.




Figure 4.3: *The superresolution and the “perfect” edge profile for a single scan. The insert shows a magnified view of the positive peak.*

4.3 The “Ideal” Modulation Transfer Function

Before we continue by analyzing different MTF measurements, we first raise the question how an MTF plot can be interpreted and what the “ideal” MTF looks like.

For the case of a 2D imaging system such as a digital camera, Burns and Williams [Burns01] separate the area of an MTF plot into several regions as given in Figure 4.4. The ideal MTF should fall completely inside a sweet spot starting at a unity modulation transfer for a constant signal. The modulus at the Nyquist frequency should be around 0.1 – the resolution limit given by the Rayleigh criterion [Born99, Williams01]. If the MTF falls off too fast, high spatial frequencies are suppressed and the image looks blurry. An MTF that falls partially into the sharpening region above the sweet spot indicates that there is a danger of over sharpening, edge ringing, and increased noise. A significant amount of signal beyond the Nyquist frequency indicates undersampling of the scene leading to aliasing artifacts in the acquired images.




Figure 4.4: *The ideal MTF for a 2D imaging system (adapted from Burns and Williams [Burns01]). An MTF should ideally fall inside the “sweet spot”. Different regions around it can be correlated with visible artifacts in the acquired images.*

4.3.1 Three-Dimensional Systems

Whereas most images are viewed as captured or with minor post-processing applied, 3D range scans are often only the input data for a long chain of processing steps. Multiple scans are for example aligned and merged into a single point cloud from which eventually a three-dimensional model of the scanned object is generated. Each operation influences the overall MTF.

A generally desirable property is however that signals up to a certain frequency should be preserved, i.e., the MTF should be unity to ensure a faithful representation. Likewise, the MTF should be zero for high frequencies to suppress noise. Whether a steep falloff with the potential of ringing artifacts or a slow falloff that smoothly suppresses higher frequencies is better remains to be determined.

Merging multiple range scans can potentially increase the overall sampling rate. The merged scan can in that way sample a signal correctly even if it contains substantial frequency components beyond the Nyquist frequency of each individual scan. This approach is related to the way the superresolution edge is constructed for our MTF measurement technique.




Figure 4.5: Data computed from the profiles in Figure 4.3. Odd (top) and even (middle) coefficients of the FFT on a logarithmic scale. Note the different orders of magnitude. Bottom: MTF plot for the odd coefficients. The scale of the horizontal axis is derived from the estimated Nyquist frequency and shows the twofold supersampling.

4.4 Empirical Analysis

We tested our approach with a Minolta VI-910 laser triangulation scanner that acquires up to 307,000 points in a single scan by sweeping a laser plane through space and observing the reflected light. A 25 mm lens was mounted on the scanner and no filtering was applied to the scans unless noted otherwise. More details about the scanner specification are available at the Minolta web site [Minolta]. The scanner was controlled by Minolta's Polygon Editing Tool software.

Finding suitable test targets is difficult as they have to contain a sharp edge with a resolution larger than the measured MTFs. For our tests we used dark gray blocks of rigid PVC with precision milled edges of 80°, 90°, and 100° edge angle. The edge lengths varied between 75 and 108 mm. The results below and in Section 4.5 suggest that the edges are of sufficient quality, however an independent test should be performed in the future.

The complete test setup for one configuration is depicted in Figure 4.1.

4.4.1 Analysis of a Single Scan

The profiles in Figure 4.3 were computed from a single range scan of the 90° target consisting of 159,780 surface points to which the *noise filter* of the Polygon Editing Tool was applied (see Section 4.5.2). After binning the data into 2048 bins and applying the Fourier transform, even and odd Fourier coefficients are plotted in separate graphs (see Figure 4.5). The horizontal axis is labeled in multiples of the sampling resolution. Note that the even coefficients of the perfect edge are several orders of magnitude smaller than the odd coefficients. Furthermore, the even coefficients of the scanned edge stay on a constant noise level which is also reached by the odd coefficients around 0.9 times the sampling resolution.

The MTF plot is quite noisy which could be improved by smoothing it, by averaging multiple plots, or by filtering the input data such as in [Fischer94]. Nevertheless, the general shape of the curve is clearly visible and shows for example a quite substantial contribution beyond the estimated Nyquist limit which is a potential source of aliasing. Furthermore, the plot has a clear minimum between 0.8 and 1.0 caused by the smoothing effect of the noise removal (see Section 4.5.2 for a detailed comparison). The signal beyond the sampling frequency seems to be caused by remaining noise in the input data which is confirmed by the results of the repeatability test in Section 4.4.2.

4.4.2 Repeatability and Robustness

As a repeatability test, we performed 20 scans of the 90° target with constant scanner settings where we slightly varied the position and orientation of the test




Figure 4.6: Average MTF for 20 scans and standard deviation per frequency.

target. Figure 4.6 shows a plot of the average MTF computed from these scans and the standard deviation per coefficient. While the technique gets unreliable for high frequencies beyond twice the Nyquist limit, it shows that the MTF can be recovered quite robustly for frequencies up to this point.

To evaluate how much the results depend on the edge angle of the test target, we performed a test series with the 80° , 90° , and 100° test targets and the VI-910 with no filtering applied. Each target was scanned 3 times and analyzed (see Figure 4.7). The averaged MTF plots show that the general shape is largely independent of the edge angle. Higher frequencies are however suppressed if the target angle gets smaller. In principle, this could be a quality problem of the target. We presume however that the main problem lies in the fact that sharper angled edges are harder to digitize. It is therefore likely that smoothing errors occur at higher frequencies. Although this behavior violates the linear system assumption (there is no unique $\text{MTF}(s)$ for a given frequency s), we still believe that our analysis is applicable and can help to understand the behavior of a scanning device.




Figure 4.7: Average MTF for 3 scans and different test targets with edge angles of 80°, 90°, and 100°.

4.5 Further Examples

We show in this section several effects that can be observed by analyzing the MTF of a 3D scanner. We first compare the MTFs computed for horizontal and vertical edges, show then the effects of different filtering operations, and finally compare the MTFs of two different scanner models.

4.5.1 Horizontal and Vertical MTF

The scanning principle of sweeping a laser plane and observing the reflections on the objects is a highly anisotropic process. This can lead to asymmetries in the MTFs for horizontal and vertical edge directions. This phenomenon can for example be observed for flatbed image scanners where a linear sensor is moved across the scan [Williams01]. It is therefore advisable to at least compute MTFs for both directions using horizontal and vertical edges [Reichenbach91].

The asymmetry of the MTF is confirmed by the plots in Figure 4.8; the falloff is steeper for the horizontal edge MTF. This can be caused by the fact that the vertical edge MTF is mainly determined by the resolution of the imaging sensor




Figure 4.8: Comparison between the MTFs for a horizontal and a vertical edge.

while the horizontal edge MTF depends as well on several other factors such as the number of positions of the laser plane or the width of the laser line. Furthermore, laser position detection can be more difficult when the projected laser line is almost parallel to the edge.

4.5.2 Scan Processing

The Minolta Polygon Editing Tool provide several filters that can be applied to a captured 3D scan after the acquisition but before exporting the data: *noise filtering*, *high quality*, and a combination of the two. The effects of this filtering on the scan and on the computed MTFs are shown in Figure 4.9. It can be clearly seen that the noise filter suppresses a frequency band around and beyond the Nyquist frequency whereas the high quality filter has a rather small influence on the MTF. The combination of both filters has almost no additional effect on the MTF compared to the noise filter alone.

After exporting, the Polygon Editing Tool also offers the possibility to smooth the scanned object. We applied this filter 0, 1, 5, and 10 times with a weight of 0.5 to the same scan used above. The resulting MTFs are shown in Figure 4.10. It can be seen that the MTF drops much earlier if more smoothing is applied. The

	VI-910	VI-700
number of points used for superresolution	159,778	19,336
number of bins used for superresolution	512	128
average distance between sampled points in mm	0.167972	0.440337

Table 4.1: Values used to compute the MTF estimates shown in Figure 4.11

images of the smoothed edges show less aliasing than the unfiltered version.

4.5.3 Scanner Comparison

When comparing MTFs of different scanners with different resolutions and consequently also different Nyquist limits, either a relative or an absolute comparison can be performed. For a *relative* comparison, the MTFs have to be normalized to the Nyquist frequency. Then the plots allow for a comparison of how well the scanners behave within their respective theoretical limits. For an *absolute* comparison, the MTFs have to be normalized to the same resolution so that the behavior for same scale details can be compared.

We included a Minolta VI-700 – an earlier model than the Minolta VI-910 with lower resolution – in the test. Table 4.1 gives the details of two analyzed test scans whose MTFs are presented in Figure 4.11. The plot shows that the response of the VI-700 is quite flat up to the Nyquist limit and then drops quickly which means that the scanner should not suffer much from aliasing. Some frequencies are however exaggerated where the MTF is above unity. Due to its higher Nyquist limit, the VI-910 can sample higher frequencies.

4.6 Discussion

We have presented a method to determine the modulation transfer function of a 3D range scanner using a slanted edge technique. Although dealing with three-dimensional data poses some additional problems compared to the MTF measurements for 2D imaging devices, the general simplicity of the slanted edge technique prevails. The experimental results show that a single scan of a target provides sufficient data for the MTF calculation.

As expected, analysis of the acquired MTFs shows analogies to the behavior of 2D MTFs for sampled imaging systems. Nevertheless, it is not clear what the “ideal” MTF of a 3D range scanning system should look like and how such an MTF should therefore be interpreted. As already mentioned in Section 4.3, this will strongly depend on further processing steps applied to the scan. Although




Figure 4.9: Influence of export filters on the MTF. No filter, the high quality filter, the noise filter, and the noise + high quality filter were applied to the same scan. The bottom row shows shaded scans of parts of the edge where some smoothing effect of the noise filter is visible at the edge (from left to right: no filter, high quality, noise, high quality + noise)




Figure 4.10: Influence of smoothing. The smoothing operation of the Polygon Editing Tool is 0, 1, 5, and 10 times applied to the unfiltered edge scan from Figure 4.9 with parameter setting 0.5. The corresponding renderings of parts of the edge. The effect of the smoothing operation is clearly visible.




Figure 4.11: An MTF computed for the VI-700 scanner. Note the differences compared to the VI-910. The horizontal axes have been rescaled to allow for an absolute comparison of the two scanners (see Section 4.5.3).

a theoretical analysis can be performed, we expect that an extensive experimental study with various combinations of 3D scanning device and postprocessing software is required in order to arrive at a well-founded answer.

But even without such a study, comparing the MTFs acquired for a single device under varying conditions (e.g., for different target materials or with different filtering options) will at least give a hint about the relative quality of the results.

4.7 Conclusion

Accurate 3D models of real world objects are a fundamental requirement for many algorithms: Both, the color managed acquisition of reflection properties presented in Chapter 5 as well as the acquisition algorithm for translucent objects shown in Chapter 7 require a 3D model of the acquired object. Apart from purely image-based methods, most rendering use the geometry of the objects in a rendered scene.

The presented method directly improve the accuracy of 3D models acquired

with a 3D range scanner but it allows to the scanners accuracy for the important case of sharp features in a real world object. It can be used as a tool to compare different 3D range scanning systems but can also give guidance during the development of such scanning systems.

Chapter 5

Color Calibrated BRDF Acquisition and Reproduction

A visually realistic 3D model of a physical object represents both its geometry as well as its surface properties. Traditionally, a diffuse texture is pasted onto the geometry model but the realism of such a model is limited because a diffuse texture cannot represent an object's true reflection properties. Additionally, specular reflections make capturing a diffuse texture challenging and highlights caused by the illumination are often “baked” into the surface leading to visible artifacts. In contrast to that, a bidirectional reflectance distribution function (BRDF) [Nicodemus77] (see Section 2.4.3) is capable of representing all local lighting effects on an object's surface leading to a realistic appearance without artifacts. Several image-based acquisition techniques [Marschner99, Lensch03a, Lensch03c] have been introduced in recent years to acquire a spatially varying BRDF for real objects.

We build on the technique by Lensch et al. [Lensch03a, Lensch03c] but focus on the quality and color fidelity of the acquired models. In order to achieve consistent color reproduction, we use in the acquisition phase a color management system based on the International Color Consortium's ICC profiles [ICC03] (see Section 3.2.3) which allows us to compute high-dynamic range images in CIEXYZ color space [CIE86] from the acquired input images. The BRDF model is then also computed in the CIEXYZ color space. We validate our approach by comparing the model at selected surface points with spectral measurements. We furthermore compute difference images (ΔE) between renderings of the model and real views thereby checking the integrity of both the BRDF modeling as well as the color management workflow. Using color managed output devices, we are able to generate synthetic renderings of real objects that match the original closely both in color as well as in specularity.



Figure 5.1: *Calibrated (left) and uncalibrated (right) carafe dataset. Both datasets are generated from the same input images – once using the color managed workflow and once using standard techniques. Variations in the highlights are caused by the optimization process used to generate the models.*

Figure 5.1 illustrates the problem of correct color reproduction: The two images show renderings of BRDF models generated from the same input data that differ mainly in color. Given any of the two images (or even both), it is impossible to determine the color of the original object. The objective of our approach is not only to acquire a BRDF model with correct colors but also to ensure correct reproduction of the rendered BRDF model on arbitrary output devices.

5.1 Previous Work

The approach described in this chapter draws on previous work in the areas of high-dynamic range imaging and BRDF measurement and reproduction.

5.1.1 High-Dynamic Range Imaging

High-dynamic range imaging was first introduced by Wyckoff [Wyckoff63] in the 1960s who invented an analog film with several emulsions of different sensitivity levels. His false color film had an extended dynamic range of about $1:10^8$.

More recently, several authors proposed methods to extend the dynamic range of digital images by combining multiple images of the same scene that differ only in exposure time. Madden [Madden93] assumes linear response of a CCD imager

and selects for each pixel an intensity value from the brightest non-saturated image. Mann and Picard [Mann95] construct the response curve of the camera by observing the intensity values of single pixels in an exposure series. A more robust method to determine the response curve, which selects a small number of pixels from the images and performs an optimization with a smoothness constraint, was introduced by Debevec and Malik [Debevec97]. They also introduced a white balance step in order to determine the relationship between the channels of an RGB color image. Robertson et al. [Robertson99, Robertson03] improved on the response curve recovery by optimizing over all pixels in all images and using a different weighting function.

The above methods for constructing an HDR image from an exposure series are either assuming a linear response of the imaging system or determine the response curve in a separate step, treating multiple color channels independently. If the relationship between the color channels is considered at all, it is only done in a white balancing step for the final HDR image.

5.1.2 BRDF Measurement

Several methods for the acquisition of surface properties have been described in the literature. Some of these methods have focused on homogeneous materials [Lafortune97, Ward92, Lu98, Marschner99] neglecting spatially varying properties of objects. Other methods assume that only the diffuse part of the BRDF (i.e., the color of the object) varies over the surface while the specular part (reflections around the mirror direction) remains fixed [Yu99, Sato97, Tominaga01] – an assumption that does not hold for many practical cases.

More general approaches which can capture unrestricted spatially varying reflection properties even of complex surfaces are reflectance fields [Debevec00] or BTFs (bidirectional texture functions) [Dana99]. However, these methods require several hundred input images for a single point of view leading to a huge effort for the acquisition and storage of the resulting models.

Far less images are required if spatially varying BRDFs are expressed using analytic BRDF models. An image-based technique for measuring spatially varying reflection properties of 3D objects requiring only 20 to 40 views has been presented by Lensch et al. [Lensch03a, Lensch03c]. In this chapter, we build on this technique and combine it with a color management system in order to faithfully reproduce the appearance of 3D objects.

5.2 Acquisition and BRDF Generation

We focus here on the color calibration aspects of the BRDF acquisition and rendering. Compared to the original approach published by Lensch et al. [Lensch03a, Lensch03c], this requires mainly adding color-related calibration steps and color space transformations. The overall acquisition, modeling, and rendering workflow remains however unchanged. We will therefore review the applied BRDF measurement technique only briefly. Please refer to the original publications for an in-depth description of the basic approach.

5.2.1 BRDF Acquisition and Modeling Pipeline

During the BRDF acquisition, the target object is illuminated by a point light source in a dark and low reflecting room (see Section 3.3.3). A digital camera is used to acquire HDR images of the object from different viewpoints and for varying light source positions. Several calibration steps ensure that the position of both camera and light source relative to the object are known. Furthermore, a geometry model of the target object is generated using 3D scanning technology.

Given this input data, the surface is clustered in different basis material clusters and the parameters for the analytical Lafortune BRDF model [Lafortune97]

$$f_r(\vec{x}, \hat{\omega}_i, \hat{\omega}_o) = k_d + \sum_j [C_{x,j} \cdot \hat{\omega}_{i,x} \hat{\omega}_{o,x} + C_{y,j} \cdot \hat{\omega}_{i,y} \hat{\omega}_{o,y} + C_{z,j} \cdot \hat{\omega}_{i,z} \hat{\omega}_{o,z}]^{n_j}.$$

are determined for each cluster (see Section 2.4.4 for a description of the parameters of the Lafortune model). The actual measurements per surface point are then projected into a basis of these per-cluster BRDFs in order to recover a truly spatially varying and highly detailed model of the reflection properties.

5.2.2 BRDFs in CIEXYZ Color Space

The Lafortune BRDF model is originally defined in linear RGB color space which is closely related to the CIEXYZ color space used in the profile connection space (PCS) of an ICC-based color management system. The similarity of the color spaces makes BRDF generation in CIEXYZ straightforward. We only need to generate HDR input images in CIEXYZ color space as described in the following section and use them for the BRDF computation. Final rendered images must also be converted back to a suitable RGB color space as described in Section 5.5.

The different topologies of the color spaces (e.g., different definitions of the luminance channel) require however some parameter changes in the non-linear fitting of the BRDF model to the input data which among other things considers




Figure 5.2: *The modified Gaussian weighting function $w(Y)$.*

at some stage chromaticity and luminance separately. This could also be a reason for the differences in the highlight regions in Figure 5.1.

5.3 Color Managed HDR Imaging

The key idea for color managed BRDF acquisition is to combine the color calibration techniques employed in a color management system with techniques for HDR image generation from a series of input images. Similar to the response curve that is used to linearize the input images in classical HDR image generation systems, the ICC profile of a color management system converts the input images also into a linear color space – the CIEXYZ color space used in the PCS. These linearized images can then be combined to a HDR image using standard techniques.

Given a set of images $\{I_j\}$ with different exposure times T_j of the same scene (i.e., of a single view used in the BRDF acquisition), we first mark all pixels in an input image as invalid for which at least one channel is overexposed or for which all channels are underexposed (or below a certain threshold). All valid pixels are then converted to CIEXYZ color space [CIE86] using the Little CMS color management engine [Saguer04] with high precision setting and relative colorimetric intent. The images are then transformed to xyY color space in order to separate

chromaticity and luminance channels:

$$\begin{aligned} x &= \frac{X}{X + Y + Z} \\ y &= \frac{Y}{X + Y + Z} \\ Y &= Y \end{aligned} \quad (5.1)$$

We define a modified Gaussian weighting function

$$w(Y) = \frac{5}{\sqrt{8\pi}} \cdot e^{-\frac{0.5 \cdot (Y-55)^2}{400}} \cdot \begin{cases} 0 & \text{invalid pixel} \\ \frac{Y^2}{900} & 0 \leq Y \leq 30 \\ 1 & 30 < Y \leq 80 \\ \frac{(100-Y)^2}{400} & 80 < Y \leq 100 \end{cases}$$

using the luminance channel Y as a confidence measure (see Figure 5.2).¹ The weighting function $w(Y)$ is centered around $Y = 55$ to emphasize reliable pixels and suppress less reliable pixels. It is zero for invalid pixels.

The xyY value of a pixel i in the final HDR image is then computed as

$$\begin{aligned} x_i &= \frac{\sum_j x_{i,j} w(Y_{i,j})}{\sum_j w(Y_{i,j})} \\ y_i &= \frac{\sum_j y_{i,j} w(Y_{i,j})}{\sum_j w(Y_{i,j})} \\ Y_i &= \frac{\sum_j Y_{i,j} T_j^{-1} w(Y_{i,j})}{\sum_j w(Y_{i,j})} \end{aligned} \quad (5.2)$$

Note that due to the division by the exposure time T_j , the range of Y_i is no longer limited to $[0..100]$ but can take on arbitrary large values. The generated HDR image is converted back to CIEXYZ color space using the inverted version of Equation 5.1 (again without limiting the values to some fixed range) and processed for BRDF generation.

5.4 Experimental Validation of the BRDF Models

Color calibrated BRDF models were acquired for several objects – a clay model of a bird, a corroded bronze bust, and a porcelain carafe (see also Figures 5.3

¹The range of valid values for the luminance channel Y for a standard image in xyY color space is $[0..100]$.

and 5.5) – using the presented technique. The acquisition was performed with the Kodak DCS 560 digital camera and a K5600 Joker Bug 800 W HMI lamp as point light source. The bird and the carafe were captured in the MPI acquisition lab whereas the bust was captured in a large hall in Florence where it is currently under restoration. An ICC input profile [ICC03] was generated using Gretag-Macbeth’s Colorchecker DC and ProfileMaker 3.1.5. All parts of the acquisition setup are described in more detail in Section 3.

We first compare the acquired appearance model with spectrophotometric measurements of the real objects in order to check color fidelity. In a second step, we render the model illuminated by a point light source and compare it to a photograph acquired under identical lighting conditions. These comparisons are based on the model in CIEXYZ color space and do not take the properties of possible output devices into account.

5.4.1 Spectrophotometric Measurements

We used a GretagMacbeth Eye-One spectrophotometer with a 45/0 measurement geometry (see [Hunt98]) to perform spectral measurements of the target objects at selected surface points. This proved to be difficult due to the curved surface geometry of the test objects and due to the bulky measurement head. In order to minimize errors, we selected near-flat regions and averaged multiple measurements. The BRDF data was generated by sampling at several locations and averaging the results. The BRDFs are computed up to a linear scale factor. We therefore determine an optimal scale factor per object based on the luminance channel and scale the BRDF data accordingly. Table 5.1 shows the resulting CIEXYZ values. The color difference between the measured and modeled values is given in ΔE (computed from the rescaled CIEXYZ values). All measurements were performed for illuminant D50 and 2 degree standard observer.

The results show that a ΔE of about 4 can be achieved for some of the materials. The orange and white surfaces of the bird show a much larger error. This is at least partly caused by the inhomogeneous color of the real model (often the yellow color of the base material shines through). Furthermore, the complex geometry of the orange and white parts of the bird and the white parts of the carafe make sampling with the spectrophotometer difficult and less reliable.

5.4.2 Comparison to a Photographic Image

To evaluate the quality of the complete BRDF model (i.e. including the modeling of highlights), we acquired additional images of the objects illuminated by a point light source. The camera pose and the position of the light source were com-

	BRDF (CIEXYZ)	Spectrophotometer (CIEXYZ)	ΔE
Bird yellow	59.41, 61.23, 14.58	61.40, 61.51, 13.42	4.158
Bird orange	42.28, 33.17, 8.02	43.73, 32.73, 5.95	7.923
Bird blue	34.99, 40.09, 46.65	32.74, 37.48, 48.16	4.283
Bird white	68.79, 73.35, 61.04	76.67, 78.75, 58.77	7.168
Carafe blue	17.63, 18.67, 33.26	17.72, 19.34, 36.325	3.398
Carafe white	57.16, 58.33, 52.11	54.00, 56.23, 43.27	7.003

Table 5.1: Comparison between the BRDF model in CIEXYZ color space and spectrophotometric measurements. CIEXYZ values are computed for illuminant D50 and 2 degree standard observer. The ΔE values are computed from the CIEXYZ values.

puted using the same calibration steps employed in the original BRDF acquisition process.

Figure 5.3 shows a side-by-side comparison of a photograph of the corroded bust and the acquired model (both images were converted to sRGB space for display using the color management workflow). The false-color image shows the difference between them in ΔE .

The BRDF model is clearly able to capture not only the general appearance of the object but also many details: The overall color reproduction is quite good, highlights are preserved and have the correct shape. The face and neck region appear however to be desaturated (especially the dark green area on the bust’s neck) while the helmet is reproduced correctly. The highlights are generally not crisp enough. These problems lead to an average ΔE of 9.8995 for which all non-black pixels with $\Delta E < 30$ were counted (we assume that larger differences are mainly caused by other problems – see below).

Some effects not related to the appearance model can be observed as well: By far the largest error occurs in regions where the shadow-boundaries are misaligned. This can be caused by a misregistration of light source and object as seen at the shadow of the nose. Another error source is inaccurate geometry which leads for example to the loss of detail in the hair region. The 3D geometry model used in this example was smoothed to remove noise and contains about 120k triangles which was not sufficient to represent small-scale details. These problems unrelated to the actual color reproduction prevented also the use of the Visual Difference Predictor [Daly93] as a psychometric quality metric for this validation.



Figure 5.3: *From left to right: Photograph of the corroded bust, rendered BRDF model, and difference image (ΔE). The false-color difference image covers the range of $[0, 50]$; areas with larger difference (mostly misaligned shadow boundaries) were cropped.*

5.5 Calibrated Rendering and Output

Up to this point, we only validated renderings generated from the actual BRDF model that did not consider the properties of the output devices. We compare therefore in this section renderings displayed on a CRT screen and color printouts to the real objects under identical illumination conditions.

5.5.1 Rendering of Calibrated BRDF Models

The high-dynamic range images rendered by a BRDF rendering system require clamping and scaling in order to generate a low-dynamic range version. The resulting images are then transformed by the ICC profile for low-dynamic display into a device dependent RGB or CMYK color space. Again, a suitable ICC profile needs to be generated for each output device (see Section 3.2.3). If no such profile is available, the sRGB color space [IEC99] can be used as default.

An alternative approach to clamping and scaling would be to apply a tone-mapping operator [Devlin02] which compresses the dynamic range such that structures in dark and in bright regions still remain visible. To our knowledge, nobody so far has reported the application of a tone-mapping operator in the context of color management.




Figure 5.4: Gamut plot of the camera system (black), the monitor (green), and the color laser printer (red) used in Figure 5.5.

5.5.2 Validation Including Output Devices

Figure 5.5 shows two photographs of a comparison setup with a CRT screen on the left, the real object in the middle, and a printout on the right. (The exposure times of both halves differ in order to compensate for the different luminance levels of the CRT and the rest of the scene). The models were rendered under identical conditions (i.e., same camera view point and illumination) as the real object.

The rendered image is displayed on a calibrated CRT monitor. The printout was performed by a calibrated HP Color LaserJet 8550 color laser printer. The system utilizes the complete color management pipeline as outlined in Figure 3.6 and no manual color tuning was performed. Under this conditions we can expect best possible color reproduction within the limits of the devices involved. Figure 5.5 shows that our BRDF measurement pipeline based on ICC color management succeeded to reproduce the appearance on different output device with only small color deviations. Note that part of these are also caused by the digital camera used to capture the validation images.

The comparison between CRT display and printout shows also some of the limitations of the output devices: The dynamic range of the color laser printer is only 1:28 so that highlights appear very dim. In addition, Figure 5.4 shows a plot of the gamut (the range of reproducible colors) of the camera system (black), the monitor (green), and the color laser printer (red). The printer gamut is smallest so that colors are most likely wrong on the printout.



Figure 5.5: *Photographic comparison of the appearance model with the original object. The images show a rendered image on a CRT screen (left), the original object under the same pose and illumination conditions (middle), and a color laser printer printout of the rendered model (right). The luminance levels required a longer exposure times for the screen compared to the other parts. The lower image is composed from 3 individual images taken under identical conditions as the model was broken during the acquisition.*

5.6 Discussion

The overhead to add color management to the presented appearance acquisition system is limited as it essentially replaces only the standard HDR computation routine. Compared to the benefit to reference the recovered appearance models to a standardized color space and to achieve best possible color reproduction, it should be a worthwhile effort. This is particularly evident in the comparison to spectrophotometric measurements shown in Section 5.4.1.

5.6.1 Accuracy of Color Management

The accuracy of the color management system both in color reproduction as well as in linearity is crucial for the approach. In the current system, we therefore use a rather aggressive weighting function $w(Y)$ to achieve reliable colors. Strongly overexposed or underexposed pixel values that are quite common in an HDR image series are problematic for the chosen color management setup and can bias the result. We expect however, that this problem could be addressed by optimizations in the profile generation phase.

5.6.2 CIEXYZ versus sRGB Color Space

The fact that CIEXYZ is one of the color spaces used in the PCS of ICC based color management system makes it a convenient and obvious choice of working color space. Nevertheless, the data could also be transferred to (linear) sRGB color space at various stages of the processing pipeline.

We opted to perform all computations in CIEXYZ mainly due to conceptual simplicity but also due to its large gamut. Our goal is to create an accurate appearance model of an object for which we utilize the full gamut of our acquisition system that is larger than sRGB even if sRGB covers the gamut of most current display and output systems.

5.7 Conclusion

Image-based techniques are an important possibility to model the complex behavior of real-world objects. Referencing the acquired models to ground-truth standards and validating their quality is therefore imperative to achieve realistic rendering and output.

We presented a color management workflow for a long and complex image-based BRDF acquisition pipeline starting from imaging the target object and converting from the physical (spectral) representation to tristimulus values and end-

ing with the output on calibrated devices. We showed both, quantitatively and visually, that the output matches the original quite closely proving that color management techniques can be successfully applied for image-based modeling. The benefit of correct color reproduction without extensive manual tuning makes the effort worthwhile and we believe that most image-based acquisition techniques can benefit from this development.

Chapter 6

Light Source Acquisition

Convincing and realistic light source models are an important prerequisite for photorealistic image synthesis. Given the exact design of a luminaire (including geometry, materials, and emission characteristics of the illuminant), one can apply standard global illumination computations to incorporate complex luminaires into synthetic scenes. This approach is however impractical in most real-world scenarios since the precise design data represents an intellectual property of significant value, and is unlikely to be shared for rendering purposes. But even if such design data is made available, it does not account for manufacturing tolerances. However, relatively minor variations of parameters like the glass thickness in halogen light bulbs can have a significant impact on the illumination patterns generated. In addition, rendering based on the geometric model can be computationally very demanding.

As an alternative, one can incorporate measured light data from complex luminaires directly into the rendering system. Such data is provided by many companies in the form of goniometric diagrams [Verbeck84] (see Section 2.3.5), which represent measurements of the far-field of a luminaire. Unfortunately, the far-field is only a faithful approximation of the emitted light when the light source is sufficiently far away from the object to be illuminated – as a rule of thumb at least five times the maximum luminaire dimension [Ashdown95]. Furthermore, some near-field effects such as soft shadows cannot be generated by a far-field approximation independent of the distance between emitter and receiver as discussed in Section 2.3.5. Even improved models using spatial intensity maps or complex emitting geometry [Rykowski97] are still too weak to represent the near-field of a complex light source.

On the other hand, a light field [Gortler96, Levoy96] completely represents near- and far-field illumination of a light source, and can thus be used to represent the light source without knowing its geometry and internal structure.



Figure 6.1: *The cloister scene illuminated by two car headlights.*

We propose therefore to model the emission characteristics of complex luminaires with a light field. We describe a complete workflow including acquisition of real world light sources and rendering scenes illuminated by captured light sources. Special attention is paid to signal processing issues and suitable sampling and reconstruction filters are incorporated into the system.

6.1 Previous Work

In his work on *near-field photometry* [Ashdown93, Ashdown95], Ashdown has presented methods for measuring the light field of luminaires in both its near-field and its far-field components. The basic principle of near-field photometry is that a number of cameras are pointed at the luminaire to be measured (or, more practically, a single camera is moved around). The camera positions correspond to a sampling of some virtual sampling surface \mathcal{S} . At each position, the irradiance incident at the imaging sensor is recorded yielding information about the angular distribution of irradiance at that particular point in space.

In Ashdown's acquisition setup, a digital camera mounted on a robot arm observes a light source from many positions on a surrounding (hemi-)sphere. Both Rykowski and Wooley [Rykowski97], and Jenkins and Mönch [Jenkins00] employ a similar setup called source imaging goniometer to acquire the light field

of a luminaire. In the system of Siegel and Stock [Siegel96], the camera lens is replaced with a pinhole.

A problem of all these approaches is that a real world light source may produce arbitrarily high spatial frequencies on the sampling surface \mathcal{S} . Sampling at discrete locations will therefore introduce aliasing unless an appropriate low-pass filter is applied in the spatial domain before the sampling step. This issue is relevant to all light field methods and has already been raised by Halle [Halle94]. It is also noted by Levoy and Hanrahan [Levoy96] in their original paper on light fields. They both propose to use the finite aperture of the camera lens as a low-pass filter by choosing the aperture size equal to the size of a sample on the camera plane. In terms of sampling theory, the shape of the spatial low-pass filter is thus a side effect of the lens system rather than a deliberate choice.

In this chapter, we describe a measurement system that significantly improves on near-field photometry by projecting the light field emitted by the light source into a finite basis *before* sampling using a simple optical system. The shape and support of the basis functions acting as low-pass filters are specifically designed for a particular sampling scheme. Based on the resulting measurements we can *exactly* reconstruct the least-squares approximation of the true light field in our basis.

In comparison to previous acquisition approaches, our optical setup lets us move in very close to the surface of the light source. In particular, the distance between the light and the aperture of our system is smaller than the distance between the aperture and the imaging plane. This allows us to capture a much wider field of view *in a single photograph*.

Our method is also closely related to the *canned light sources* proposed by Heidrich et al. [Heidrich98] as a way to integrate complex light sources into computer graphics rendering systems. A canned light source is a light field representation of a luminaire's emission characteristics and can act as a "black box" for any type of light source (similar to the representations used in near-field photometry).

Using projective textures [Segal92], a scene illuminated by a canned light source can be efficiently rendered using hardware-accelerated rendering techniques [Heidrich98]. We introduce therefore a series of approximations necessary to render a scene containing our acquired light sources in a similar way. We furthermore derive an efficient importance sampling algorithm that allows to use the measured light sources in many global illumination algorithms.

6.2 Overview

Our conceptual setup to filtered light source measurement is depicted in Figure 6.2. Light rays are emitted from the light source and hit a filter in the sampling




Figure 6.2: Cross section through the conceptual acquisition setup.

plane \mathcal{S} . This plane is opaque except for the finite support area of the filter which can be moved to a number of discrete positions on a regular grid. The filter is a semi-transparent film, similar to a slide, containing the 2D image of a basis function $\Phi'_{ij}(u, v)$. Light rays are passing straight through the filter without scattering but are attenuated according to $\Phi'_{ij}(u, v)$. They then hit the measurement plane \mathcal{M} , on which we are able to measure the irradiance with a high spatial resolution. We assume that sampling plane \mathcal{S} and measurement plane \mathcal{M} are parallel so that $d = R \cdot \cos \theta$.

When comparing this setup with the work on near field photometry [Ashdown93, Siegel96], \mathcal{S} roughly corresponds to the plane on which the pinholes or finite apertures of the cameras are located, while \mathcal{M} corresponds to the imaging sensor of the camera. The fundamental difference in our approach is that we can select an arbitrary filter $\Phi'_{ij}(u, v)$ which projects the light field emitted by the luminaire into an arbitrarily chosen function space. This results in an optical low-pass filtering of the light field *before* sampling, and thus helps avoiding aliasing artifacts.

In principle, one also has to perform low-pass filtering on the measurement plane \mathcal{M} . This, however, is optically very difficult or even impossible to implement if the filter kernels overlap. Fortunately, it is technically easy to achieve a very high measurement resolution on \mathcal{M} , so that we can consider the measured data as point samples, and implement an arbitrary filter kernel while downsampling to the final resolution.

Note that the filter kernels for \mathcal{S} and \mathcal{M} can be chosen independently as required by the reconstruction process. We are using in this thesis a box filter for

Symbol	Meaning
$\Psi_{ijkl}(u, v, s, t)$	basis function for approximating the light field
Φ_i	1D basis used for reconstruction
$\Phi_{ijkl}(u, v, s, t)$	4D tensor product basis for reconstruction
Φ'_i	biorthogonal 1D basis used for measurement
$\Phi'_{ijkl}(u, v, s, t)$	4D tensor product basis for measurement
\mathcal{M}	surface on which the irradiance is measured (<i>measurement surface</i> , or (s, t) -plane)
\mathcal{S}	surface on which the optical filters are placed (<i>sampling surface</i> , or (u, v) -plane)
$L(u, v, s, t)$	radiance passing through (u, v) on \mathcal{S} and (s, t) on \mathcal{M}
$\tilde{L}(u, v, s, t)$	projection of $L(u, v, s, t)$ into basis $\{\Psi_{ijkl}(u, v, s, t)\}$
$L_{mn}(u, v, s, t)$	radiance $\Phi'_{mn}(u, v) \cdot L(u, v, s, t)$ projected through <i>one</i> filter $\Phi'_{mn}(u, v)$
$E_{mn}(s, t)$	irradiance caused by $L_{mn}(u, v, s, t)$ on the measurement surface \mathcal{M}

Table 6.1: The notation used in this chapter (see also Figure 6.2 for an overview over the conceptual acquisition setup).

$\Phi'_{kl}(s, t)$ assuming sufficiently high measurement resolution on \mathcal{M} . To construct $\Phi'_{ij}(u, v)$, we use a more advanced basis such as piecewise quadratic basis. Section 6.4 describes the design of the filters in more detail.

6.3 Theoretical Framework

Before we discuss the theory behind the proposed method in detail we first introduce the mathematical notation used throughout this chapter. This notation is summarized in Table 6.1, an overview over the measurement setup is given in Figure 6.2.

In our approach, we assume that the light field emitted by the light source is well represented by a projection into a basis $\{\Psi_{ijkl}(u, v, s, t)\}_{ijkl \in \mathbb{Z}}$:

$$L(u, v, s, t) \approx \tilde{L}(u, v, s, t) := \sum_{i,j,k,l} \Psi_{ijkl}(u, v, s, t) \cdot L_{ijkl}. \quad (6.1)$$

We assume that Ψ_{ijkl} has local support, and i, j, k , and l roughly correspond to translations in u, v, s , and t , respectively. \mathbb{Z} is the set of intergers. Note, however, that the translated basis functions will not in all cases have the same shape, i.e., $\Psi_{i'j'k'l'}(u, v, s, t)$ may not be an exact copy of $\Psi_{ijkl}(u, v, s, t)$ in general.

We also define two additional sets of basis functions, one for *measuring* and one for *reconstruction*. For reconstruction we use a 1D basis $\{\Phi_i\}_{i \in \mathbb{Z}}$ with the property $\Phi_i(x) = \Phi(x + i)$. The 4D reconstruction basis is then given as the tensor product basis

$$\Phi_{ijkl}(u, v, s, t) := \Phi_{ij}(u, v) \cdot \Phi_{kl}(s, t) = \Phi_i(u) \cdot \Phi_j(v) \cdot \Phi_k(s) \cdot \Phi_l(t).$$

For measurement, we use the *biorthogonal* (or *dual*) $\{\Phi'_i(x)\}_{i \in \mathbb{Z}}$ of the reconstruction basis with

$$\int_{-\infty}^{\infty} \Phi'_i(x) \cdot \Phi_j(x) dx = \delta_{ij}. \quad (6.2)$$

and again we use a tensor-product construction for the 4D dual basis.¹

6.3.1 Measured Irradiance

Our approach is based on measuring the irradiance $E_{mn}(s, t)$ in the (s, t) -plane that is caused by the incident radiance

$$L_{mn}(u, v, s, t) = \Phi'_{mn}(u, v) \cdot L(u, v, s, t).$$

The result of such a measurement is

$$\begin{aligned} E_{mn}(s, t) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{\cos^2 \theta}{R^2} \cdot \Phi'_{mn}(u, v) \cdot L(u, v, s, t) du dv \\ &\approx \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{\cos^2 \theta}{R^2} \cdot \Phi'_{mn}(u, v) \cdot \tilde{L}(u, v, s, t) du dv \\ &= \sum_{i,j,k,l} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{\cos^2 \theta}{R^2} \cdot \Phi'_{mn}(u, v) \cdot \Psi_{ijkl}(u, v, s, t) \cdot L_{ijkl} du dv. \end{aligned} \quad (6.3)$$

The geometric term $\cos^2 \theta / R^2$ is composed of the distance R of the point on the (u, v) -plane from the point on the (s, t) -plane, as well as the cosine of the angle θ between the plane normals and the vector connecting the two points (see Figure 6.2 for a sketch of the setup). Note that this term also accounts for any differences in the parameterizations on the two planes (i.e., different grid spacings).

In Section 6.5, we will discuss two physical setups for performing this kind of measurement.

¹The Kronecker delta function δ_{uv} is defined as $\delta_{uv} = 1$ if $u = v$ and $\delta_{uv} = 0$ otherwise.

6.3.2 Exact Reconstruction

We now describe an exact reconstruction algorithm given the measurements E_{mn} . To this end, we first define what the relationship between the basis functions Ψ_{ijkl} and the reconstruction and measurement bases should be. We define

$$\Psi_{ijkl}(u, v, s, t) := \frac{R^2}{\cos^2 \theta} \cdot \Phi_{ij}(u, v) \cdot \Phi_{kl}(s, t). \quad (6.4)$$

Inserting this definition into Equation 6.3, and using the biorthogonality relationship (Equation 6.2) yields

$$E_{mn}(s, t) = \sum_k \sum_l \Phi_{kl}(s, t) \cdot L_{mnkl}. \quad (6.5)$$

To determine which reconstruction filter to use, we now rewrite Equation 6.1 using Equations 6.4 and 6.5:

$$\begin{aligned} \tilde{L}(u, v, s, t) &= \sum_{m,n,k,l} \Psi_{mnkl}(u, v, s, t) \cdot L_{mnkl} \\ &= \sum_{m,n,k,l} \frac{R^2}{\cos^2 \theta} \cdot \Phi_{mn}(u, v) \cdot \Phi_{kl}(s, t) \cdot L_{mnkl} \\ &= \sum_{m,n} \frac{R^2}{\cos^2 \theta} \cdot \Phi_{mn}(u, v) \cdot E_{mn}(s, t). \end{aligned} \quad (6.6)$$

This indicates that we can *exactly* reconstruct \tilde{L} , the projection of L into the basis $\{\Psi_{ijkl}\}$, by using the reconstruction filter

$$\frac{R^2}{\cos^2 \theta} \cdot \Phi_{mn}(u, v).$$

Note that this reconstruction filter does contain a shift-variant component in form of the geometric term. Interestingly, this geometric term will cancel out for the Photon Map particle emission algorithm we discuss in Section 6.6. However, for the hardware accelerated rendering algorithm shown in Section 6.6.2 the reconstruction step should ideally only involve a convolution with a shift-invariant filter kernel. We will therefore introduce in the next section further approximations that remove this shift-variant term.

6.3.3 Shift-Invariant Approximation

In order to introduce a shift-invariant reconstruction based on the same measured irradiance as defined in the previous sections, we define the relationship between Ψ_{ijkl} , the measurement and the reconstruction bases differently. Starting from

$$\Psi_{ijkl}(u, v, s, t) := \Phi_{ij}(u, v) \cdot \Phi_{kl}(s, t),$$

we get

$$E_{mn}(s, t) = \sum_{ijkl} \int_S \frac{\cos^2 \theta}{R^2} \cdot \Phi'_{mn}(u, v) \cdot \Phi_{ij}(u, v) \cdot \Phi_{kl}(s, t) \cdot L_{ijkl} du dv.$$

Note that both, θ and R , depend not only on (s, t) but also on (u, v) .

To provide a real shift-invariant approximation, we use the following assumption: If the distance d between the (u, v) -plane and the (s, t) -plane is large compared to the support of $\Phi_{ij}(u, v)$, and if θ is small, then the geometric term $\cos^2 \theta / R^2$ is well approximated by one constant for each point on the (s, t) -plane:

$$g(u, v, s, t) := \frac{\cos^2 \theta(u, v, s, t)}{R(u, v, s, t)^2} \approx \frac{\cos^2 \theta_{kl}(s, t)}{R_{kl}(s, t)^2}. \quad (6.7)$$

This yields an approximation of the measured irradiance:

$$E_{mn}(s, t) \approx \sum_{kl} \frac{\cos^2 \theta_{kl}(s, t)}{R_{kl}(s, t)^2} \cdot \Phi_{kl}(s, t) \cdot L_{mnkl}.$$

Since both the geometric term and $\Phi_{kl}(s, t)$ are known, it is in principle possible to compute the approximate incoming radiance

$$E'_{mn}(s, t) \approx \sum_{kl} \Phi_{kl}(s, t) \cdot L_{mnkl}$$

by de-convolution. In practice, this is only feasible for basis functions $\Phi_{kl}(s, t)$ with a small support. This is not a major problem, however, since the practical measurement setups presented in the previous section have a very high resolution on the (s, t) -plane, so that a bilinear or even a box filter can be used. Like in the previous section, we apply the definition of Ψ_{ijkl} to determine the appropriate reconstruction filter:

$$\begin{aligned} \tilde{L}(u, v, s, t) &= \sum_{mnkl} \Psi_{mnkl}(u, v, s, t) \cdot L_{mnkl} \\ &= \sum_{mnkl} \Phi_{mn}(u, v) \cdot \Phi_{kl}(s, t) \cdot L_{mnkl} \\ &\approx \sum_{mn} \Phi_{mn}(u, v) \cdot E'_{mn}(s, t). \end{aligned}$$

Error Analysis for the Shift-Invariant Approximation

The quality of this approximation depends on the error introduced by assuming the geometric term constant over the support of the basis function $\Phi_{ij}(u, v)$ in Equation 6.7. To evaluate the validity of this approximation, we define the following

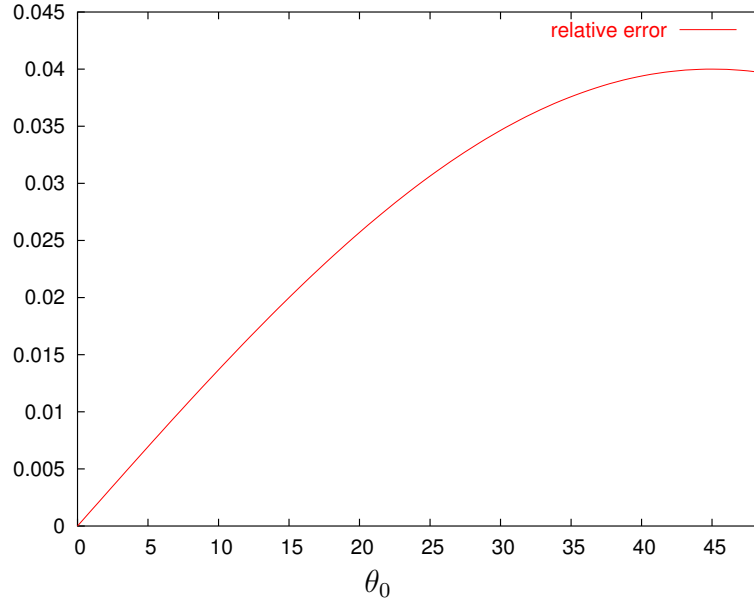


Figure 6.3: Error analysis for the geometric term. The plots show the relative error $E_r(\theta_0)$ for angles θ_0 (in degrees). The setup corresponds to Setup B in Table 6.2 (distance $\mathcal{M} - \mathcal{S}$: 750 mm, size of measurement plane: $1.2 \text{ m} \times 1.2 \text{ m}$, filter width of reconstruction filter: 15 mm).

relative error

$$E_r(\theta_0) := \frac{\max_{\theta \in \mathcal{F}} \cos^4(\theta) - \min_{\theta \in \mathcal{F}} \cos^4(\theta)}{\cos^4(\theta_0)}$$

where \mathcal{F} is the support of the basis Φ_{ij} , θ_0 is the angle at the center of this support, and $\cos^4(\theta)$ corresponds to the geometric term g of Equation 6.7 for a distance of 1 between the planes. Its evaluation shows that the error is below 4% if the ratio s between filter width and the distance of the two planes \mathcal{S} and \mathcal{M} equals 0.02 (see Figure 6.3). This corresponds to Setup B described in Section 6.5.2 with large quadratic filter (see also Table 6.2 in Section 6.7.1).

6.4 Filter Design

The theoretical framework and the implementations detailed above allow for the use of arbitrary filters Φ_i and Φ'_i with finite support. It is therefore possible to design filters adapted to the specific properties of the light source at hand, for example taking the structure and smoothness of the light source's emission pattern into account. The intended use of the data – especially the subsequent reconstruction and rendering steps – should also be considered here.




Figure 6.4: A quadratic reconstruction basis and one dual. Left: Function Φ_{bellq} (red) and a dual Φ'_{bellq} (green) used for most of the measurements. Middle and right: 2D filter images used for the acquisition (split into positive and negative components and normalized).

In this section, we will introduce two alternative bases we used for our measurements – a piecewise quadratic basis and a B-Spline basis. We furthermore give some details about the fabrication of the physical filters used in the acquisition setups.

6.4.1 Piecewise Quadratic Basis

The first series of the measurements shown in this thesis were performed using manually operated translation stages. To reduce acquisition time we therefore aimed at minimizing the number of measurements so that a small support of the measurement filter kernel (relative to the filter spacing) was desired. This could for example be achieved with the classical bilinear basis. We chose however to use a piecewise quadratic basis function Φ_{bellq}

$$\Phi_{bellq} = \begin{cases} 1 - 2x^2 & |x| \leq \frac{1}{2} \\ 2(|x| - 1)^2 & \frac{1}{2} < |x| \leq 1 \\ 0 & \text{else} \end{cases} \quad (6.8)$$

due to its C^1 continuity that helps preventing Mach bands. Figure 6.4 shows a plot of this function and one dual. Note that the optimal dual that performs an orthogonal projection on the basis has infinite support and can therefore not be manufactured as a filter. We therefore use an approximate dual with finite support (shown in Figure 6.4 middle and right) that still fulfills Equation 6.2.

As shown in Figure 6.4, the dual Φ'_{bellq} of Φ_{bellq} which we need to use as measurement filter contains negative values which can not be implemented optically as




Figure 6.5: A B-Spline reconstruction basis $\Phi_{b\text{-spline}}$ (green) and one dual $\Phi'_{b\text{-spline}}$ (red).

an opacity. We therefore split the filters into a positive and a negative component

$$\Phi'_{ij(+)}(u, v) := \begin{cases} \Phi'_{ij}(u, v) & \Phi'_{ij}(u, v) \geq 0 \\ 0 & \text{else} \end{cases}$$

$$\Phi'_{ij(-)}(u, v) := \begin{cases} 0 & \Phi'_{ij}(u, v) \geq 0 \\ -\Phi'_{ij}(u, v) & \text{else} \end{cases}$$

and perform separate measurements for both parts. The results are then subtracted in software. Similarly, we can rescale basis functions with values > 1 and correct for it in software.

6.4.2 B-Spline Basis

In a later stage of the measurements, we changed the reconstruction basis to a B-Spline basis $\Phi_{b\text{-spline}}$. Consider the C^1 continuous quadratic B-Spline functions over the uniform knot vector $(\dots, -\frac{5}{3}, -\frac{3}{3}, -\frac{1}{3}, \frac{1}{3}, \frac{3}{3}, \frac{5}{3}, \dots)$. This way, one representative of them can be written as

$$\Phi_{b\text{-spline}} = \begin{cases} \frac{3}{4} - \frac{9}{4}x^2 & |x| \leq \frac{1}{3} \\ \frac{9}{8}(|x| - 1)^2 & \frac{1}{3} < |x| \leq 1 \\ 0 & \text{else} \end{cases}.$$

Figure 6.5 shows a plot of this function and one dual. Unfortunately, the support of the optimal dual basis $\Phi'_{b\text{-spline}}$ required for capture is infinite. Even using an approximative dual with finite support as acquisition filter is quite impractical




Figure 6.6: Left: *Normalized 2D filter image of the B-Spline reconstruction basis $\Phi_{b\text{-spline}}$ used for the acquisition.* Middle and right: *Testpattern for the film writer and one of the acquired images.*

since it would not only require manufacturing a large filter but also lead to a much higher number of sample locations consequently slowing down the acquisition process.

The dual functions $\Phi'_{i,b\text{-spline}}$ are however linear combinations of the basis functions $\Phi_{j,b\text{-spline}}$ since they belong to the same function space:

$$\Phi'_{i,b\text{-spline}} = \sum_j a_{ij} \Phi_{j,b\text{-spline}}.$$

If we make a scalar product of this equation with $\Phi_{k,b\text{-spline}}$ and apply Equation 6.2, we get

$$\delta_{ij} = \sum_j a_{ij} \langle \Phi_{j,b\text{-spline}} \cdot \Phi_{k,b\text{-spline}} \rangle.$$

The coefficients a_{ij} can thus be computed by inverting the matrix

$$M = (\langle \Phi_{j,b\text{-spline}} \cdot \Phi_{k,b\text{-spline}} \rangle).$$

This allows us to use the basis functions $\Phi_{b\text{-spline}}$ as filter in the acquisition process and to compute correctly filtered input images as linear combinations of the acquired images using the weighting factors a_{ij} .

6.4.3 Filter Fabrication and Calibration

The filters were printed with 1016 dpi resolution on high density film using a LightJet 2080 Digital Film Recorder. To achieve linear film response, a test pattern containing patches with all 256 possible gray values was printed. The density of the patches was determined by capturing multiple high dynamic range images of the backilluminated pattern. Brighter areas were successively covered to reduce the effects of lens flare. The resulting filters achieve a dynamic range of more than 1 : 50,000.

6.5 Acquisition Setup

In this section, we explain how the theoretical framework can be transformed into a practical acquisition method. We propose two different hardware configurations and discuss their advantages and disadvantages.

6.5.1 Direct Acquisition (Setup A)

The first setup we propose – Setup A – is based on a standard photographic camera where the lens is removed and replaced by the filter Φ'_{ij} (see Figure 6.7). As with conventional light fields, the camera has to be positioned at the vertices of a regular grid within the sampling plane \mathcal{S} and images can be captured.

Figure 6.8 shows our implementation of Setup A. A custom-built filter holder is attached to the camera instead of a lens, so that the filters can easily be exchanged without affecting the geometry of the setup. The light source to be measured is mounted on manually operated translation stage so that it can be shifted in front of the camera with high precision and repeatability (see Figure 6.14).

Although this setup is compact and easy to handle, it does have some drawbacks. The fixed and often quite small size of the imaging sensor imposes a strong limitation on the size of the filter kernel and the incident angles of illumination covered by the measurement. In addition, the distance between the filter plane and the imaging plane determines the focal length of the camera (similar to a pin-hole camera). A large distance leads to a tele-lens effect and further limits the useful incident angles.

6.5.2 Reflector-Based Acquisition (Setup B)

These problems can be avoided with Setup B. Here, the filter Φ'_{ij} is used to project an image of the light source onto a diffusely reflecting screen. A standard camera system records the projected image. Depending on the actual setup, a substantial calibration effort is required to rectify, undistort and register them to the setup before they can be used.

For the implementation of Setup B shown in Figure 6.10, we used a $1.2\text{ m} \times 1.2\text{ m}$ flat board that was spray-painted with ultra matte paint as a projection screen. A black board with an integrated filter holder was used to mount the filter and shield the light source from the projection screen. To avoid occlusion problems we mounted the camera above the whole setup and equipped it with a wide angle lens allowing us to capture the whole projection screen. The camera parameters were recovered using standard camera calibration techniques and the images were rectified and registered to the filter setup.




Figure 6.7: Setup A for light source acquisition: The camera lens is replaced by a filter and the light field is directly projected onto the imaging sensor inside the camera.




Figure 6.8: Implementation of Setup A (Figure 6.7). A filter holder is mounted to the camera instead of the lens. The camera with the filter holder is moved relative to the light source.




Figure 6.9: Setup B for light source acquisition: The filter $\Phi'_{ij}(u, v)$ projects an image on a diffuse reflector which is captured by a camera.




Figure 6.10: Implementation of Setup B (Figure 6.9). The projection screen is observed by the camera located above the setup. The black board contains the filter holder. The light sources can be moved on the table parallel to the filter and the projection screen.




Figure 6.11: *Alternative implementation of Setup B using the UBC Active Measurement Facility (ACME, see Section 3.4). The projection screen is observed by the camera located above the setup. The black board contains the filter holder. The light sources can be moved on the table parallel to the filter and the projection screen.*

In order to be able to perform the acquisition in a fully automatic way, we built a second version of Setup B utilizing the ACME system (see Section 3.4 for a description of ACME and Figure 6.11 for an image of the setup). The system allows for computer controlled movement of both, the filter and the light source.

If a high quality diffuse (Lambertian) reflector material is available, the filter size and viewing angles are (almost) only limited by the size of the reflector. On the other hand, the larger setup makes it potentially harder to use, and errors due to indirect light reflecting from parts of the equipment such as the reflection screen or the environment tend to increase. Furthermore, the calibration effort increases since the intrinsic and extrinsic parameters of the camera system relative to the projected image need to be recovered.

6.6 Rendering

In this section, we discuss the integration of our measured light sources into both, a particle-based global illumination system and a hardware-accelerated rendering

system.

6.6.1 Importance Sampling

The key to include our measured light source datasets into a particle based global illumination system is efficient importance sampling of the light source. We included the importance sampling approach described in the next section into an implementation of the *photon map* algorithm [Jensen96, Jensen01a, Jensen01b] which we chose because it is one of the most common and simple to use methods available. Note that our importance sampling approach does not only apply to photon maps, but to any algorithm that requires shooting from the light source such as particle tracing [Walter97] or bidirectional path-tracing [Lafortune93, Veach94].

For the scene description, we attach the light field to a virtual geometric object, representing an impostor for the area from which light is emitted similar to the parametrization introduced for area light sources in Section 2.3.4. This can be, for example, the glass surface of a car head light, or simply a planar polygon representing one of the light field planes. This geometry helps in positioning the light source in the 3D scene and can also be used to aid the sampling in our rendering algorithm.

To emit a particle from our light source representation, we could then use the standard method of selecting a random point and random direction on the impostor geometry and perform rejection sampling. While this approach might be sufficient for a relatively uniform light source, it can be inefficient when high spatial and directional frequencies are present. We therefore develop an importance sampling strategy for our light field representation.

The particles that we want to emit have to estimate the total energy E coming from the light field:

$$E = \int_{u,v,s,t} \tilde{L}(u, v, s, t) \frac{\cos^2 \theta}{R^2} \quad (6.9)$$

In order to achieve a constant energy for each photon, the probability density function (PDF) of the photon emission has to be:

$$p(u, v, s, t) = \frac{1}{E} \tilde{L}(u, v, s, t) \frac{\cos^2 \theta}{R^2} \quad (6.10)$$

To generate random photons distributed according to this PDF, we use the standard approach of generating uniformly distributed random points in 4D, and passing them through the inverted cumulative density function (CDF, see [Pitman92]). If we have pre-inverted and tabulated the CDF, then generating a new random photon according to this importance is possible in constant time, i.e., independent of the light field resolution.




Figure 6.12: A light field cell for PDF computation formed by a single pixel on the measurement plane \mathcal{M} and by $1/4$ of a filter on the filter plane \mathcal{S} for the piecewise quadratic basis. This corresponds to the region on \mathcal{S} where $\Phi_{i,bellq}(u)$, $\Phi_{i+1,bellq}(u)$, $\Phi_{j,bellq}(v)$, and $\Phi_{j+1,bellq}(v)$ overlap.

One complication in our case is, however, that we have the PDF represented in terms of some complicated and overlapping basis functions instead of the usual point sampled (i.e., tabulated) representation. We therefore do not directly generate and invert the CDF from the basis representation, but rather choose a two-steps approach.

We first compute a 4D table of constant probabilities for selecting a 4D cell from the light field (see Figure 6.12). A single 4D cell corresponds to all the directions coming from one pixel on the measurement plane to $1/4$ of a filter on the filter plane for the piecewise quadratic basis and $1/9$ of a filter on the filter plane for the B-Spline basis, respectively. Inside these areas, the same 4 reconstruction filters $\Phi_{ij,bellq}$ (or 9 reconstruction filters $\Phi_{ij,b-spline}$) overlap. The probability on this cell can be computed by some simple summations over the light field coefficients. Then, as we now have a constant probability for each cell, we can invert this table more easily than the continuous representation of the PDF.

Note that all these steps are performed only once in a preprocessing stage. The resulting table is part of the light source description that we save on disk.

6.6.2 Estimation of Direct Illumination

In order to efficiently compute the reflected radiance due to direct illumination from the shift-invariant approximation (Section 6.3.3), we have to evaluate the




Figure 6.13: Configuration of a direct reflection.

following equation for a visible point \vec{x} and viewing direction $\hat{\omega}_e$:

$$\begin{aligned}
 L(\vec{x}, \hat{\omega}_e) &= \sum_{mn} \int_S f_r(\vec{x}, \hat{\omega}_e, u, v) V(\vec{x}, u, v) \frac{\cos(\theta) \cos(\theta')}{R^2(\vec{x}, u, v)} \\
 &\quad \tilde{L}(u, v, s(\vec{x}, u), t(\vec{x}, v)) dudv \\
 &= \sum_{mn} \int_S f_r(\vec{x}, \hat{\omega}_e, u, v) V(\vec{x}, u, v) \frac{\cos(\theta) \cos(\theta')}{R^2(\vec{x}, u, v)} \\
 &\quad E'_{mn}(s(\vec{x}, u), t(\vec{x}, v)) \Phi_{mn}(u, v) dudv,
 \end{aligned} \tag{6.11}$$

Figure 6.13 shows the geometry of the setup and the notation. The BRDF $f_r(\vec{x}, \hat{\omega}_e, u, v)$ describes the object's reflection properties. $V(\vec{x}, u, v)$ and $R(\vec{x}, u, v)$ encode the visibility and the distance between the light sample (u, v) and the position \vec{x} , respectively.

To simplify this expression, we make the common assumption that the reflection properties, visibility, and the geometric term are constant on the support of Φ_{mn} and are estimated at its center:

$$\begin{aligned}
 f_{r,mn}(\vec{x}, \hat{\omega}_e) &\approx f_r(\vec{x}, \hat{\omega}_e, u, v) \\
 V_{mn}(\vec{x}) &\approx V(\vec{x}, u, v) \\
 g_{mn}(\vec{x}) &\approx \frac{\cos(\theta) \cos(\theta')}{R^2(\vec{x}, u, v)}.
 \end{aligned}$$

Given these assumptions, Equation 6.11 can be approximated by:

$$L(\vec{x}, \hat{\omega}_e) \approx \sum_{mn} V_{mn}(\vec{x}) f_{r,mn}(\vec{x}, \hat{\omega}_e) g_{mn}(\vec{x}) \cdot \int_{\mathcal{S}} E'_{mn}(s(\vec{x}, u), t(\vec{x}, v)) \Phi_{mn}(u, v) dudv \quad (6.12)$$

We then compute $\overline{E'_{mn}}(\vec{x})$, the average value of $E'_{mn}(s(\vec{x}, u), t(\vec{x}, v))$ on the support of Φ_{mn} for a position \vec{x} in order to make our final approximation of Equation 6.11:

$$L(\vec{x}, \hat{\omega}_e) \approx A \sum_{mn} V_{mn}(\vec{x}) f_{r,mn}(\vec{x}, \hat{\omega}_e) g_{mn}(\vec{x}) \overline{E'_{mn}}(\vec{x}), \quad (6.13)$$

where $A = \int_{\mathcal{S}} \Phi_{mn}(u, v) dudv$.

Each term of this sum corresponds now to the reflection of a textured spot-light [Segal92], located at the center of a filter support, and pointing toward \vec{x} (similar to the canned light source approach [Heidrich98]).

6.7 Results

We first give details about the acquired datasets and show then rendering results using both, the photon map based global illumination system and the hardware-accelerated rendering system.

6.7.1 Acquisition

We acquired several datasets with the presented setups. Table 6.2 summarizes the acquisition geometry, Table 6.3 gives an overview over the number of views captured and the filters used. The datasets were acquired using high-dynamic range imaging techniques. The flashlight was acquired using the Kodak DCS 560 digital camera and the bike- and car headlight were acquired with the Jenoptik ProgRes C14. (See Section 3.2.4 and 3.2.3 for a description of these cameras and of high-dynamic range imaging). A 5×5 median filter was used to suppress the noise in the images captured by the Kodak DCS 560 camera resulting from long exposure times. Remaining noise appears as background illumination in the light field. All acquired images were downsampled to achieve a resolution of 300×300 pixels on the measurement plane \mathcal{M} .

Setup	Distance $\mathcal{M} - \mathcal{S}$	Size of \mathcal{M}	Filter Width (Acquisition Filter)	Filter Type
Setup A	55 mm	18 mm \times 27 mm	10 mm	Quadratic
Setup B	750 mm	1.2 m \times 1.2 m	20 mm (medium) 30 mm (large)	Quadratic Quadratic
Setup B (ACME)	750 mm	1.22 m \times 1.07 m	49 mm	B-Spline

Table 6.2: Geometry of the acquisition setups. The filter spacing on the filter plane \mathcal{S} is $1/4$ of the width of the acquisition filter (the dual filter) for the piecewise quadratic basis Φ'_{bellq} and $1/3$ of the width of the acquisition filter for the B-Spline basis $\Phi'_{\text{b-spline}}$.

Data Set	Setup	# Views	Filter Width (Basis Φ_i)	Filter Width (Dual Basis Φ'_i)
Flashlight	A	9 \times 9	5 mm	10 mm
Flashlight Medium	B	7 \times 7	10 mm	20 mm
Flashlight Large	B	5 \times 5	15 mm	30 mm
LED Flashlight	B	5 \times 5	15 mm	30 mm
Bike Light	B	9 \times 7	15 mm	30 mm
Car Headlight	B (ACME)	13 \times 9	49 mm	—

Table 6.3: Parameters for the acquired data sets.




Figure 6.14: Left and middle: Image of the flashlight and the bike light. Right: The negative part of the filter Φ'_{bellq} illuminated by the bike light.




Figure 6.15: Flashlight acquisition with Setup A. Left: A projection on a distant plane (3 m). Right: A global illumination solution. The acquired data set was converted to grayscale.




Figure 6.16: Left: Photograph of the flashlight illuminating a plane at 75 cm distance. Middle: Reconstruction of the same setup from data measured with a filter of 30 mm width. Right: Reconstruction from data measured with a filter of 20 mm width.

Setup A

Only the flashlight dataset was captured with Setup A. Due to geometry constraints imposed by the dimensions of the camera (see Table 6.2), we were only able to capture a small part of the light field. Figure 6.15 shows on the left the acquired light source pointed at a projection screen and on the right onto the crypt. The captured part corresponds roughly to the central dark region in the projection of the real light source (see Figure 6.16, left). The low resolution is due to the comparably large filter size. Using a smaller filter would increase both the resolution and the number of views required. In this configuration, Setup A serves as a proof-of-concept but can probably not be used to acquire a complete model of a




Figure 6.17: *The bike light. Left: Photograph of the bike light illuminating a plane at 75 cm distance. Right: Reconstruction of the same setup from data measured with a filter of 30 mm width.*

real light source.

Setup B

With Setup B, we can generally capture a much larger portion of the light field of the light source due to the larger viewing angle of the configuration. Additionally, a much larger filter kernel can be used for the filter plane while still improving the resolution of the data set.

In Figure 6.16 we show the light pattern that a flashlight causes on a plane at 75 cm distance for a specific focus setting of the flashlight. The left image is a photograph of the flashlight pointed at a diffuse surface. The remaining images show two reconstructions from measurements. The middle image was generated from a dataset where a dual filter width of 30 mm has been used, while a dual filter of 20 mm width was used for the right image. In comparing the reconstructed images with the photograph, we can see that there are some tone-mapping issues that result in slightly different colors. However, the main features of the light cone, such as the black ring in the center, are represented very well in both reconstructions.

In Figure 6.17 we show similar images for the light field emitted by a bike light. The left image again shows a photograph of the light source illuminating a plane at 75 cm distance. The image on the right is a reconstruction for the same distance. The narrow vertical stripes are an interesting feature of the bike light that should be preserved in a model. They are mainly emitted from the upper and lower edges of the luminaire and are reproduced very well when rendered. The




Figure 6.18: *Light field acquired for the bike light dataset (rotated counter-clockwise by 90°).*




Figure 6.19: Left and middle: *Photograph and reconstruction of the LED flashlight acquired using the 30 mm acquisition basis Φ'_{bellq} .* Right: *Reconstruction of the car headlight acquired with the 49 mm B-Spline acquisition basis $\Phi'_{b-spline}$.*

image shows furthermore some artifacts due to misalignment caused by the manual acquisition process and due to remaining non-linearities in the manufactured filters.

Due to the effective size of the luminaire (5 cm×2.5 cm) and its strongly varying emission characteristics, the near field is quite extended and inhomogeneous. This is particularly evident in the overview over the captured light field in Figure 6.18. (The individual sub-images of the light field were normalized for display. Note that they also contained negative values due to the acquisition filter so that the black or gray background in each image corresponds to 0.) Different parts of the narrow vertical stripes are emitted from different areas of the reflector. Their behavior could probably be even modeled using a single textured spotlight. The central bright spot is in contrast to that emitted throughout the reflector.

The LED flashlight contains 3 LEDs instead of a light bulb (see left of Figure 6.19 - the fourth, smaller LED is a low battery light, and is not used in normal operation). Although the actual flashlight is identical to the one from Figure 6.16 except for the light bulb and the reflector, the cone of light produced by this setup is quite different from the one before. In particular, the illumination is much more uniform, although roughly triangular due to the LED layout. This illustrates the importance of using actual measurements from real world light sources in image synthesis, since even small changes in the reflectors can cause major differences in appearance.

Figure 6.19 shows furthermore on the right a reconstruction of the car headlight pointing at a plane. This dataset is the only one acquired using the B-Spline acquisition basis $\Phi'_{b-spline}$. It shows a strong peak pointing straight ahead and a much dimmer spread-out pattern to illuminate the close area ensuring that drivers are not blinded when the headlight is used in the car.




Figure 6.20: *Exploring the cloister with the bike light.*

6.7.2 Rendering

Figures 6.1 and 6.20 show a global illumination solution based on the acquired datasets. The global illumination system is based on the photon mapping approach [Jensen01b] and uses the importance sampling approach described in Section 6.6.1. Note that the distinctive patterns generated by the real world light sources are reproduced faithfully.

The hardware-accelerated rendering approach described in Section 6.6.2 was implemented on a Linux workstation with an Intel Xeon 1.7 GHz processor, 512 MB memory and an NVIDIA GeForce FX 5800 Ultra graphics board. The acquired images are downsampled to a size of 300×300 pixels. The current implementation allows for rendering the flashlight and bike light datasets interactively with a frame rate from 3.2 fps down to 0.55 fps at a window size of 400×400 pixels (see Figures 6.21 and 6.22). Rendering speed depends on both, the size of the datasets and the polygon count of the rendered geometry model.




Figure 6.21: *Rendering of the 5×5 flashlight dataset in the crypt environment (500 polygons) with different light source positions. Rendering speed: 2.7 to 3.2 frames per second at a window size of 400×400 pixels.*




Figure 6.22: *Left: Rendering of the 5×5 flashlight dataset zoomed on a detail with small soft shadow. Right: An earlier version of the bike light dataset captured with the Kodak DCS 560 and rendered at 0.55 frames per second in the cloister model (8000 polygons). The color difference compared to Figure 6.20 is caused by noise in the input images.*

6.8 Conclusions

We presented a novel approach for accurate real world light source acquisition. A model is measured using a simple optical system to project the real light source into a finite basis. This projection allows a correct optical pre-filtering of the acquired light field. We introduced a theoretical framework for this approach and two possible setups for real acquisitions. The tests we performed to validate this approach have shown that our setup can faithfully represent rather complex near- and far-field information with a relatively small number of measurements. The fact that we perform optical low-pass filtering avoids aliasing artifacts such as missing directional peaks. In our evaluations we found that Setup B, although larger and harder to calibrate, is more practical for real-world light sources, since it covers a larger field-of-view.

We also presented an efficient algorithm for integrating our new light representation into particle-based global illumination systems. Using a pre-inverted cumulative density function, this method allows a constant time importance sampling to emit those particles. Direct lighting can be computed interactively using graphics hardware based on a shift-invariant approximation of the acquired data.

Given the combination of acquisition and rendering, we introduced a complete pipeline to incorporate real world light sources into virtual worlds.

Chapter 7

Translucent Objects

Light interacts with ideal opaque objects purely locally – it is scattered at the point of incidence according to the bidirectional reflectance distribution function (BRDF, see Section 2.4.3) [Nicodemus77]. In contrast, light passes straight through transparent objects (e.g., glass) and is only reflected or refracted at material boundaries. Many daily life objects (e.g., milk, skin or marble) are translucent and belong to neither of these categories. They are characterized by multiple light scattering inside the object. This *subsurface scattering* behavior leads to a very distinct appearance which is visually important: light shines through objects, they appear smooth and surface details are hidden.

Translucent objects can be rendered using a variety of physical simulation techniques. Most recent rendering systems are based on the dipole approximation to a diffusion model [Jensen01c]. This approach enables interactive evaluation while providing good visual quality. Furthermore, physically correct parameters can be determined for homogeneous materials by a point-based measurement setup [Jensen01c]. To our knowledge, there exists however no technique that can derive the necessary input data for a real object with spatially varying properties such as the alabaster block shown in Figure 7.1. Several image-based acquisition systems are able to capture objects that exhibit subsurface scattering. Unfortunately, they do not take the specific properties of translucent objects into account and are unable to provide a comprehensive model of such an object. They are particularly not able to render translucent objects from arbitrary viewpoints and under arbitrary lighting conditions where the incident illumination varies per surface point (e.g., illumination with textured point lights, or projected shadow borders).

We therefore propose a method to acquire a comprehensive model of a translucent object with diffuse light propagation inside the object and local illumination variation due to shadows or directionally dependent light sources. To this end we sequentially illuminate a dense set of locations on an object’s surface and observe




Figure 7.1: *Side view of an alabaster block illuminated from the top by a laser. Light is non-uniformly scattered depending on the material properties.*

its impulse response function with a high-dynamic range video camera. The acquired data is resampled into a hierarchical model taking advantage of the specific properties of translucent material such as the exponential fall-off near the point of incidence and the smooth global response. We introduce methods to consistently interpolate holes (mainly caused by occlusion) in the acquired data and to reduce the variance due to noise in the acquisition process. We implemented a rendering method that allows us to display the acquired objects interactively. Finally, we discuss how the acquired data could be incorporated into many current rendering algorithms as a model for real objects exhibiting subsurface scattering.

7.1 Previous Work

The theory of radiation in scattering media is a well studied problem [Ishimaru78] and is of interest in many applications (e.g., wave propagation for communication, medicine, remote sensing). The goal in medicine and remote sensing is to draw conclusions about the interior of a region by analyzing the scattered response to incoming radiation which requires solving a difficult inverse problem. The predominant goal in computer graphics is to acquire enough data to achieve a realistic visual model of an object's response to incident illumination. While knowledge of the interior structure of an object is required to perform exact simulations of the physical behavior, it is neither mandatory nor necessarily efficient. Our approach is entirely based on external measurements and does not depend on solving an

inversion problem.

7.1.1 Models for Translucent Objects

Translucent objects can be modeled by defining their basic physical properties (e.g., the absorption and scattering cross sections σ_a and σ_s [Ishimaru78]) for each point inside their volume. Rendering such an object could then make use of physical simulations or appropriate approximations.

Alternatively, it is sufficient to record the visible effects of subsurface scattering. Translucent objects can be modeled using a bidirectional scattering-surface reflectance distribution function $S(\vec{x}_i, \hat{\omega}_i; \vec{x}_o, \hat{\omega}_o)$ (BSSRDF, see Section 2.4.2) [Nicodemus77] that relates irradiance at a surface position \vec{x}_i to the reflected radiance at \vec{x}_o for arbitrary incoming and outgoing directions $\hat{\omega}_i$ and $\hat{\omega}_o$. The complexity of this 8D function makes handling it quite cumbersome. The BSSRDF model of Jensen et al. [Jensen01c] consists of a directionally dependent single scattering term $S^{(1)}$ and a directionally independent multiple scattering term S_d (see also Section 2.4.5). However, for the case of optically dense material, the directional dependency is negligible since the response is dominated by multiple scattered particles. Omitting the single scattering term, the BSSRDF of a translucent object can be collapsed to the 4D diffuse subsurface reflectance function $R_d(\vec{x}_i, \vec{x}_o)$ that depends only on the incoming and outgoing surface positions.

7.1.2 Acquisition Techniques

Many image-based acquisition and rendering techniques are able to record and reproduce some aspects of translucent objects: Techniques based on light fields [Levoy96] or lumigraphs [Gortler96] achieve this for a given set of viewing and lighting conditions. Surface light fields [Miller98, Wood00] record objects under a single fixed illumination for arbitrary viewpoints. In contrast to surface light fields, reflectance fields [Debevec00] as well as polynomial texture maps [Malzbender01] capture an object illuminated by a set of distant point light sources seen from a fixed viewpoint.

Strongly translucent and transparent objects with reflection and refraction effects can be recorded for a single viewpoint using environment matting [Zongker99, Chuang00]. Matusik et al. [Matusik02] combine environment matting and reflectance fields to acquire an image-based representation of transparent and refractive objects including their 3D shape. Even though some of the above techniques record objects under varying illumination, they record only cases where the whole object is illuminated by distant light sources. Hence, as noted by Debevec et al. [Debevec00], they are unable to faithfully reproduce the

effects of local illumination variation such as shadow boundaries projected onto an object.

Masselus et al. [Masselus03] capture the reflectance field of a scene for a fixed viewpoint and arbitrary illumination parameterized as a 4D incident light field. This allows for lighting the scene with arbitrary light sources including local illumination within the relatively coarse resolution limit of the incident light field.

Jensen et al. [Jensen01c] measure the scattering parameters for common materials such as marble or skin. They illuminate a single point on the surface, capture the reflected radiance with a digital camera, and compute the absorption cross section σ_a and the reduced scattering cross section σ'_s . Their results allow for modeling homogeneous objects made from these materials. Arbitrary inhomogeneous objects are beyond the scope of their measurement approach.

In contrast, we take advantage of the specific properties of translucent materials and acquire the diffuse subsurface reflectance function $R_d(\vec{x}_i, \vec{x}_o)$ that has no angular variance at incident and exitant surface locations. This allows us to densely sample the incoming and outgoing surface locations. We furthermore go beyond a pure image-based representation and transform the acquired data into a hierarchical data structure which reduces the storage cost considerably and allows us to interpolate missing data points consistently.

7.1.3 Rendering

A variety of rendering techniques such as finite element methods [Rushmeier90, Sillion95, Blasi93], bidirectional path tracing [Hanrahan93, Lafortune96], photon mapping [Jensen98, Dorsey99], Monte Carlo simulations [Pharr00, Jensen99], or diffusion [Stam95, Stam01] are able to simulate the effects of subsurface scattering based on physical principles. Subsurface scattering can also be integrated in the framework of precomputed radiance transfer [Sloan03].

Jensen et al. [Jensen01c] introduced a practical BSSRDF model consisting of a dipole approximation of the diffuse, multiple scattering term and an extension of the Hanrahan and Krueger BRDF model [Hanrahan93] for the local, directional dependent effects caused by single scattering. The simplicity of the model implies drastically reduced rendering times compared to a full Monte Carlo Simulation. Although the dipole approximation is only correct for the surface of a homogeneous, infinite half-space, it yields visually convincing results. The degree of realism can be further improved by adding a modulation texture to simulate spatially varying materials.

Rendering with BSSRDFs is expensive since contributions from the whole surface must be taken into account. Jensen and Buhler [Jensen02] sample therefore the irradiance first and hierarchically evaluate the diffusion approximation in a second pass. Hao et al. [Hao03] use a local illumination model and precompute

the transfer factors for all vertex neighborhoods. Mertens et al. [Mertens03a] use a clustering hierarchy which is re-built in real time to render deformable translucent objects interactively. They also propose a local subsurface scattering algorithm that integrates the samples in image space at interactive rates using hardware acceleration [Mertens03b]. The translucent shadow map [Dachsbacher03] augments a shadow map by the surface normal and the irradiance entering the object. This information is used to evaluate the dipole approximation by hierarchically filtering the translucent shadow map in real time on graphics hardware. Carr et al. [Carr03] built an adaptive link structure in a preprocessing phase encoding the overall response to incoming light and evaluate it at rendering time using 3 passes on the graphics board.

Apart from Jensen et al. [Jensen01c], none of the above rendering methods are based on extensive sets of measured input data. The main focus of our approach is therefore to capture the exact behavior of real translucent objects including effects that cannot be captured by simplified models such as the dipole approximation (e.g., cracks in an object, volumetrically varying properties, hollow objects). The acquired model is comprehensive and versatile. We will also discuss how it can be integrated into a variety of the above rendering algorithms and present a interactive rendering technique for the acquired models.

7.2 Hierarchical BSSRDF Model

The bidirectional scattering-surface reflectance distribution function (BSSRDF, see Section 2.4.2) S provides a general model for basic light transport. Omitting spectral effects, polarization, and fluorescence, the outgoing radiance $L^\rightarrow(\vec{x}_o, \hat{\omega}_o)$ at a surface location \vec{x}_o and a direction $\hat{\omega}_o$ is related to the incident radiance $L^\leftarrow(\vec{x}_i, \hat{\omega}_i)$ at all surface locations \vec{x}_i and for all directions $\hat{\omega}_i$ via

$$L^\rightarrow(\vec{x}_o, \hat{\omega}_o) = \int_A \int_{\Omega_+(\vec{x}_i)} L^\leftarrow(\vec{x}_i, \hat{\omega}_i) S(\vec{x}_i, \hat{\omega}_i; \vec{x}_o, \hat{\omega}_o) \langle \hat{n}_i \cdot \hat{\omega}_i \rangle d\hat{\omega}_i d\vec{x}_i.$$

A denotes the surface of the object and $\Omega_+(\vec{x}_i)$ is the hemisphere of directions on the outside of the surface at \vec{x}_i .

Under the assumption of multiple light scattering inside the object which removes the angular dependencies at incoming and exiting surface locations we can replace S with a 4D diffuse subsurface reflectance function $R_d(\vec{x}_i, \vec{x}_o)$:

$$L^\rightarrow(\vec{x}_o, \hat{\omega}_o) = \frac{1}{\pi} F_{t,o}(\eta, \hat{\omega}_o) \int_A R_d(\vec{x}_i, \vec{x}_o) \cdot \int_{\Omega_+(\vec{x}_i)} L^\leftarrow(\vec{x}_i, \hat{\omega}_i) F_{t,i}(\eta, \hat{\omega}_i) \langle \hat{n}_i \cdot \hat{\omega}_i \rangle d\hat{\omega}_i d\vec{x}_i. \quad (7.1)$$




Figure 7.2: Light propagation in translucent objects. All photons incident at \vec{x}_i and exiting at \vec{x}_o must pass through two small volumes dV around \vec{x}_i and \vec{x}_o while their paths are spread out inbetween.

The Fresnel transmittance factors $F_t(\eta, \hat{\omega})$ model what fraction of the flux or radiosity is transmitted through the surface boundary for a direction $\hat{\omega}$ and a given relative index of refraction η . The factor $1/\pi$ converts radiosity into exitant radiance. The area foreshortening is represented by $\langle \hat{n}_i \cdot \hat{\omega}_i \rangle$.

The goal of our work is to measure $R_d(\vec{x}_i, \vec{x}_o)$ per color channel for all incoming and outgoing surface locations \vec{x}_i and \vec{x}_o . If we illuminate a target object at a single surface point \vec{x}'_i with known incident radiance $L^\leftarrow(\vec{x}'_i, \hat{\omega}_i)$, we can observe $L^\rightarrow(\vec{x}_o, \hat{\omega}_o)$ on the object's surface. Inverting Equation 7.1 becomes trivial and we can record the impulse response $R_d(\vec{x}'_i, \vec{x}_o)$ for all \vec{x}_o .

Storing and rendering the full 4D function R_d with high resolution is impractical. The nature of optically dense translucent objects (see Figure 7.2) however makes R_d well suitable for a hierarchical modeling approach: R_d will typically vary strongly in the proximity of \vec{x}_i due to the exponential fall-off of radiosity inside the scattering volume. In distant areas, R_d varies quite smoothly. Sharp features are usually caused by inhomogeneous material properties within a small volume dV on and right below the object's surface at \vec{x}_i and \vec{x}_o . All photons pass through these volumes and their influence is significant (see Feng et al. [Feng93] for a study of photon path distributions within scattering material).

7.2.1 Discretization of Light Transport

The hierarchical BSSRDF model is based on a discrete version of the integral expression in Equation 7.1 at which we arrive with a Galerkin type approach [Heckbert91, Zatz93]. For each hierarchy level, we introduce a set of basis functions for the discretization. Multiple discrete representations of the same data are of course redundant but allow us to adapt the sampling density to the smoothness of R_d . Appropriate blending will be necessary for correct results when combining them into a single model.

On the finest level of the hierarchy we generally employ a set of piecewise constant basis functions assuming a sufficiently high sampling density to model the objects response in the vicinity of an incoming light impulse. At coarser levels, which will be used to represent the smoother global response, we use more elaborate basis functions such as hat functions.

For the discretization, we first reformulate Equation 7.1 in order to be able to address the irradiance $E(\vec{x}_i)$ collected at a surface point \vec{x}_i and the radiosity $B(\vec{x}_o)$ at the outgoing surface position \vec{x}_o individually:

$$L^-(\vec{x}_o, \hat{\omega}_o) = \frac{1}{\pi} F_{t,o}(\eta, \hat{\omega}_o) B(\vec{x}_o) \quad (7.2)$$

$$B(\vec{x}_o) = \int_A R_d(\vec{x}_i, \vec{x}_o) E(\vec{x}_i) d\vec{x}_i \quad (7.3)$$

$$E(\vec{x}_i) = \int_{\Omega_+(\vec{x}_i)} L^-(\vec{x}_i, \hat{\omega}_i) F_{t,i}(\eta, \omega_i) \langle \hat{n}_i \cdot \hat{\omega}_i \rangle d\hat{\omega}_i. \quad (7.4)$$

For each of the discretizations we first choose a set of spatial basis functions $\psi_u(\vec{x})$ and project the irradiance $E(\vec{x}_i)$ onto the chosen basis:

$$\tilde{E}_u(\vec{x}_i) = \sum_u E_u \psi_u(\vec{x}_i) \approx E(\vec{x}_i)$$

The coefficients E_u are found by calculating scalar products of $E(\vec{x}_i)$ with dual basis functions $\tilde{\psi}_u(\vec{x})$:¹

$$E_u = \int_A E(\vec{x}_i) \tilde{\psi}_u(\vec{x}_i) d\vec{x}_i \quad (7.5)$$

Equation 7.3 is then transformed into a matrix-vector multiplication:

$$B_v = \sum_u E_u \mathbf{F}_{u,v} \quad (7.6)$$

¹The dual basis functions are the unique set of linear combinations of the primary basis functions $\psi(\vec{x})$ that fulfill the orthonormality relations $\int_A \psi_u(\vec{x}) \tilde{\psi}_v(\vec{x}) d\vec{x} = \delta_{uv}$. The Kronecker delta function δ_{uv} is defined as $\delta_{uv} = 1$ if $u = v$ and $\delta_{uv} = 0$ otherwise.

with throughput factors

$$\mathbf{F}_{u,v} = \int_A \int_A \psi_u(\vec{x}_i) R_d(\vec{x}_i, \vec{x}_o) \tilde{\psi}_v(\vec{x}_o) d\vec{x}_o d\vec{x}_i. \quad (7.7)$$

The radiosity $B(\vec{x}_o)$ at surface position \vec{x}_o is reconstructed as

$$B(\vec{x}_o) \approx \sum_v B_v \psi_v(\vec{x}_o). \quad (7.8)$$

The conversion of radiosity into outgoing radiance $L^\rightarrow(\vec{x}_o, \hat{\omega}_o)$ for a particular direction $\hat{\omega}_o$ can then be performed according to Equation 7.2.

7.2.2 Modulation of Smooth Hierarchy Levels

Using a coarse hierarchy level to model the scattering behavior of an object will lead to a smooth radiosity distribution on its surface (depending on the sample spacing and the chosen set of basis functions $\psi_u(\vec{x})$). For an object with inhomogeneous scattering properties, that exhibits high frequency variations on the surface as shown in Figure 7.2, this will lead to a too smooth reconstruction hiding important details.

If we assume that the high frequency variations depend on the properties of a small volume $dV(\vec{x}_o)$ surrounding an exitant surface location \vec{x}_o , the effect of the details is independent of \vec{x}_i and applies to all discretized radiosity contributions B_v . We can therefore introduce a modulation factor $T_\rho(\vec{x}_o)$ that modulates the computed radiosity for each surface location. Using Equation 7.3, $T_\rho(\vec{x}_o)$ should be chosen such that the objective function

$$O = \int_A R_d(\vec{x}_i, \vec{x}_o) E(\vec{x}_i) d\vec{x}_i - T_\rho(\vec{x}_o) \cdot \sum_v \left(\psi_v(\vec{x}_o) \cdot \sum_u E_u \mathbf{F}_{u,v} \right)$$

is minimal for all incoming surface positions \vec{x}_i .

Overall, the outgoing radiance $L^\rightarrow(\vec{x}_o, \hat{\omega}_o)$ of one hierarchy level is now approximated by

$$L^\rightarrow(\vec{x}_o, \hat{\omega}_o) \approx \frac{1}{\pi} F_{t,o}(\eta, \hat{\omega}_o) \cdot T_\rho(\vec{x}_o) \cdot \sum_v \left(\psi_v(\vec{x}_o) \cdot \sum_u E_u \mathbf{F}_{u,v} \right). \quad (7.9)$$

7.2.3 Implemented Model

For the remainder of this chapter we use an implementation of the above model with two hierarchy levels – a model with high sampling density in the vicinity of

the incoming impulse (*local term*) and a much coarser sampling in distant, smooth areas (*global term*) possibly modulated by T_ρ . The radiosity due to the local and global terms is denoted B^l and B^g , respectively.

The parameterization and the sets of basis functions used in the model are chosen with the requirements of an hardware-accelerated rendering approach in mind: We use the vertices of a triangle mesh representing the object's geometry as sample locations for the global term. The local term is parameterized as mesh texture stored in a texture atlas. We assume furthermore, that during rendering the irradiance is available for each non-empty texel in the texture atlas, i.e., that an irradiance map corresponding to the texture atlas is available.

Global Term

The smooth global term of the subsurface light transport is represented by vertex-to-vertex form factors $F_{i,j}$ using linear interpolation of vertex radiosities. The form factors are stored in a form factor matrix F . Linear interpolation of vertex colors is well-known in graphics under the name of Gouraud interpolation. On a triangle mesh, it corresponds to representing a color function by its coefficients with respect to the following basis functions:

$$\begin{aligned}\psi_1^g(\vec{x}) &= \beta_1(\vec{x}) \\ \psi_2^g(\vec{x}) &= \beta_2(\vec{x}) \\ \psi_3^g(\vec{x}) &= \beta_3(\vec{x})\end{aligned}$$

where

$$\vec{x} = \beta_1(\vec{x})\vec{v}_1 + \beta_2(\vec{x})\vec{v}_2 + \beta_3(\vec{x})\vec{v}_3.$$

v_1, v_2, v_3 are the three vertices of the triangle containing \vec{x} . The β 's are the barycentric coordinates of \vec{x} in the triangle $\vec{v}_1, \vec{v}_2, \vec{v}_3$. Note that

$$\beta_3(\vec{x}) = 1 - \beta_1(x) - \beta_2(x).$$

We associate a basis function ψ_i^g with every vertex i in the triangle mesh. $\psi_i^g(\vec{x})$ is zero except for the three vertices in the triangle containing \vec{x} . When plotted, these basis functions look like “hats” centered at each mesh vertex.

In order to achieve a good sampling density, an input triangle mesh with typically several 100 k triangles needs to be simplified to an appropriate vertex count (on the order of 10 k vertices for interactive rendering). During simplification we try to obtain equilateral triangles of similar size in order to improve the uniformity of the basis functions.

Local Term

To capture the strong variations of the local term in the immediate neighborhood of an incident light impulse, we use a surface texture stored in a texture atlas to parameterize the object. We assign constant basis functions per texel and model the light transport by means of 7×7 texel-to-texel throughput filter kernels $K_{(u,v)}$ centered at each non-empty texel (u, v) in the texture atlas.

Mathematically, this corresponds to projecting on a piecewise constant basis functions $\psi_{(u,v)}^l$. The basis functions are 1 on the part of the model surface $S(u, v)$ projected in a single texture atlas texel (u, v) (the footprint of the texel). Everywhere else the basis functions are 0. The dual basis functions $\tilde{\psi}_{(u,v)}^l$ are piecewise constant in the same way, except that they take a value $1/A(u, v)$ instead of 1 on $S(u, v)$. $A(u, v)$ is the area of $S(u, v)$.

The texture atlas is generated by first splitting the triangular mesh of the model into different partial meshes and orthographically projecting each partial mesh onto a suitable plane. The angle between the normals of the projected triangles and the plane normal are kept small to limit distortion and ensure best sampling rate. Starting with a random triangle, we add an adjacent triangle to a partial mesh if the deviation of the triangle's normal compared to the average normal of the partial mesh is below some threshold (commonly $\frac{\pi}{4}$ leading to a maximum distortion factor of $1/\cos \frac{\pi}{4} = \sqrt{2}$). We also add to each partial mesh a border formed by adjacent triangles. The width of the border is required to be at least 3 texels to provide sufficient support for applying the 7×7 filter kernels to the core of the partial mesh. The border triangles may be distorted in order to fulfill this constraint.

All projected partial textures are rotated to ensure that the area of their axis-aligned bounding box is minimal [Toussaint83, Pirzadeh99]. A generic packing algorithm generates a dense packing of the bounding boxes into a square texture of predefined size. The algorithm is able to scale the size of the bounding boxes using a global scaling factor in order to ensure dense packing. Figure 7.9 shows an example of a texture atlas generated for the alabaster horse model.

Blending Global and Local Term

The global and the local response cannot be simply added to obtain the correct result. In the regions of direct illumination, both contributions will add up to approximately twice the correct result. However, the radiosity B^g will have the largest interpolation error near the point of light incidence while B^l returns the more accurate response for points close to direct illumination (see Figure 7.3 (b) and (c)). B^l is actually only available for those points.

Our choice is to keep the texel accurate filter kernel for the local radiosity

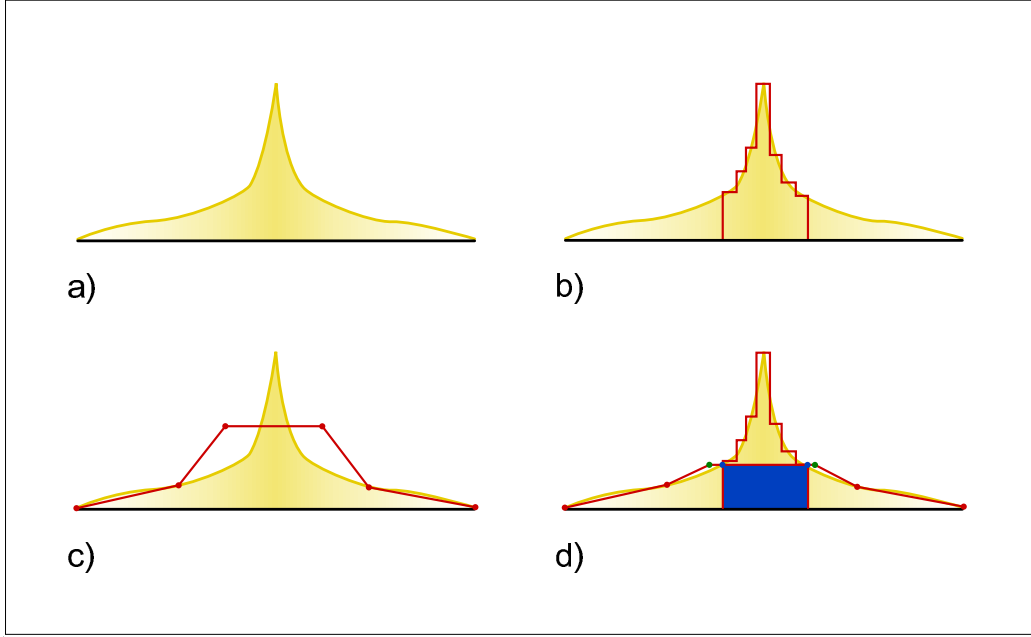


Figure 7.3: a) Ideal impulse response. b) Local response B^l modeled by the filter kernel in texture space with constant basis functions (red) c) Linear interpolation of the global response B^g resulting from distributing the irradiance and evaluating the form factor matrix \mathbf{F} . d) Optimized global and local response: The diagonal of \mathbf{F}^0 is set to zero, the weights for B_i^d (green dots) are optimized to interpolate the boundary of the filter kernel (blue dots), the blue area is subtracted from the filter kernel.

since it represents the best local response of our model. Thus, we somehow have to reduce the influence of the global part at small scattering distances and must ensure smooth blending between the local and global response where the influence of the local response ends.

The global radiosity B^g due to direct illumination corresponds to the diagonal elements $\mathbf{F}_{i,i}$ of the form factor matrix \mathbf{F} . Since the radiosity due to direct illumination will be modeled by the local term B^l , the diagonal entries of \mathbf{F} are set to zero yielding \mathbf{F}^0 and a new global radiosity B^{g0} without direct illumination

$$B_\nu^{g0} = \sum_i E_i^g \mathbf{F}_{i,\nu}^0. \quad (7.10)$$

To obtain smooth blending we introduce a new radiosity vector B_i^d which is directly derived from the irradiance at the kernel centers (u, v) :

$$B_i^d = \sum_{(u,v)} E(u, v) \psi_i^g(\vec{x}(u, v)) w_i(u, v).$$

The weighting function $w_i(u, v)$ specifies the contribution of the irradiance incident at (u, v) to B_i^d for the three vertices of the triangle containing (u, v) . Using this new radiosity vector, the combined radiosity response will be obtained as

$$B(\vec{x}) = B^l(\vec{x}) + B^d(\vec{x}) + B^{g0}(\vec{x}). \quad (7.11)$$

Our approach to determine the weighting functions $w_\nu(u, v)$ is to minimize the difference between the interpolated global radiosity and the *correct* radiosity for each texel (s, t) on the boundary Γ of its filter kernel $K_{(u,v)}(s, t)$. The *correct* radiosity at Γ is found by calculating or measuring a larger filter kernel $K_{(u,v)}^{9 \times 9}(s, t)$. Stated mathematically, the problem is to find $w_\nu(u, v)$ such that

$$\begin{aligned} O &= \sum_{(s,t) \in \Gamma} \left[K_{(u,v)}^{9 \times 9}(s, t) E(u, v) - B^{g0}(\vec{x}_c(s, t)) - B^d(\vec{x}_c(s, t)) \right]^2 \\ &= \sum_{(s,t) \in \Gamma} \left[K_{(u,v)}^{9 \times 9}(s, t) E(u, v) - B^{g0}(\vec{x}_c(s, t)) - \sum_{\nu=1}^3 E(u, v) \psi_\nu^g(\vec{x}_c(s, t)) w_\nu(u, v) \right]^2 \end{aligned}$$

is minimal. $\vec{x}_c(s, t)$ is the surface point corresponding to the center of texel (s, t) . The new global radiosity $B^{g0}(\vec{x})$ at a point \vec{x} is the sum of B_ν^{g0} (see Equation 7.10) over the vertices ν of the triangle containing \vec{x} :

$$B^{g0}(\vec{x}) = \sum_{\nu} \psi_\nu^g(\vec{x}) B_\nu^{g0}.$$

After correcting the global response, we also have to change the filter kernels. The interpolated values of the global response have to be subtracted from each kernel, corresponding to the blue area in Figure 7.3 (d).

The irradiance at each texel is now distributed to two different vectors: to E_i^g using the dual basis functions $\tilde{\psi}^g$ and to B_i^d using the weights w_i described in this section.

Modulation Texture

For the modulation texture, we make use of the parameterization given by the texture atlas. Modulation factors $T_\rho(u, v)$ can thus be defined using the basis functions $\psi_{(u,v)}^l$ defined for the local term.

7.3 Generation of Synthetic Datasets

Given a method to evaluate the diffuse subsurface scattering reflectance function $R_d(\vec{x}_i, \vec{x}_o)$ for arbitrary incoming and outgoing surface locations, we can compute

the necessary parameters of the proposed model, i.e., spatially varying filter kernels $K_{(u,v)}$, the vertex-to-vertex throughput factor matrix \mathbf{F} , and the modulation texture T_ρ .

7.3.1 Local Response

According to Equation 7.5, the irradiance coefficients $E^l(u, v)$ correspond to the average irradiance on $S(u, v)$. We will approximate them by the value at the center point in the texel. The texel-to-texel throughput factor (Equation 7.7) between texel (u, v) and (s, t) is approximated as $K_{(u,v)}(s, t) = A(u, v)R_d(\vec{x}_c(u, v), \vec{x}_c(s, t))$ with R_d being evaluated at the surface points \vec{x}_c corresponding to the center of the texels. These texel-to-texel throughput factors can be viewed as non-constant irradiance texture filter kernels. Equation 7.6 then corresponds to a convolution of the irradiance texture. The convolved (blurred) irradiance map shows the locally scattered radiosity $B^l(\vec{y})$.

7.3.2 Global Response

We will need to project the irradiance $E(x)$ into the basis of hat functions:

$$E(\vec{x}) \approx \sum_i E_i^g \psi_i^g(\vec{x}).$$

The coefficients E_i^g are given by scalar products with the dual basis functions $\tilde{\psi}_i^g$ of ψ_i^g (see Equation 7.5). These dual functions are zero except for the three vertices of the triangle containing \vec{x} . The three non-zero values are

$$\tilde{\psi}_\nu^g(\vec{x}) = \frac{3}{A_\nu}(4\beta_\nu(\vec{x}) - 1)$$

where A_ν is the sum of the areas of the triangles sharing vertex ν , ν being a vertex of the triangle containing \vec{x} .

The throughput factors (7.7) are approximated in the following way, which requires to evaluate the diffuse subsurface scattering reflectance R_d only once for each pair of mesh vertices (v_i, v_j) :

$$\mathbf{F}_{i,j} \approx R_d(v_i, v_j) \cdot \int_S \psi_i^g(\vec{x}) d\vec{x} \cdot \int_S \tilde{\psi}_j^g(\vec{y}) d\vec{y} = \frac{A_i}{3} R_d(v_i, v_j).$$

The matrix-vector product (Equation 7.6) then results in the scattered radiosity B_j^g at the mesh vertices v_j . The global radiosity $B^g(\vec{y})$ for intermediate surface points \vec{y} is found by linear interpolation:

$$B^g(\vec{y}) = \sum_j B_j^g \psi_j^g(\vec{y}) = \sum_\nu B_\nu^g \beta_\nu(\vec{y})$$

where the latter sum is over the three vertices of the triangle containing \vec{y} .




Figure 7.4: *Synthetic datasets. The bust on the right is rendered as skim milk. Differences in the scattering behavior between the horse model with uniform marble and with added veins are clearly visible.*

7.3.3 Modulation Texture

Since the form factor matrix \mathbf{F} already computes the radiosity at the vertices correctly, we have to ensure that those values are not changed by the modulation texture T_ρ . Therefore we divide the vertex radiosity by its corresponding texture value prior to multiplication with the texture:

$$B_i^T = B_i^g / T_\rho(v_i) \quad (7.12)$$

The complete low-frequency response is then given by

$$B^g(\vec{x}) = T_\rho(\vec{x}) \sum_{\nu=1}^3 \psi_\nu^g(\vec{x}) B_j^T. \quad (7.13)$$

7.3.4 Examples

The synthetic datasets shown in Figure 7.4 were computed using the dipole approximation [Jensen01c] (see Section 2.4.5) for $R_d(\vec{x}_i, \vec{x}_o)$ and the parameters from [Jensen01c]. The bust is made out of skim milk, the uniform horse is computed using the parameters of marble. For the right horse, the marble was mixed with a much darker material using a Perlin noise function [Perlin85]. The rendering was performed using the technique presented in Section 7.8.

7.4 Acquisition of Translucent Objects

We described in the previous section how synthetic data for the hierarchical BSSRDF model can be generated. In the next sections we will introduce a method

to acquire a model for real world translucent objects.

The system employs a laser projection system to sequentially illuminate individual surface locations \vec{x}'_i with a narrow beam of light. Under the assumption of diffuse light propagation inside the translucent object, the objects impulse response corresponds to the diffuse subsurface scattering reflectance function R_d . The outgoing radiance can be captured by a high-dynamic range video camera. The properties of this setup allow us to simplify Equation 7.1 to

$$\begin{aligned} L^\rightarrow(\vec{x}_o, \hat{\omega}_o) &= \frac{1}{\pi} F_{t,o}(\eta, \hat{\omega}_o) \int_A R_d(\vec{x}_i, \vec{x}_o) \delta(\vec{x}_i - \vec{x}'_i) \\ &\quad \cdot \int_{\Omega_+(\vec{x}_i)} L^\leftarrow(\vec{x}_i, \hat{\omega}_i) F_{t,i}(\eta, \hat{\omega}_i) \langle \hat{n}_i \cdot \hat{\omega}_i \rangle \delta(\hat{\omega}_i - \hat{\omega}'_i) d\hat{\omega}_i d\vec{x}_i \\ &= \frac{1}{\pi} F_{t,o}(\eta, \hat{\omega}_o) R_d(\vec{x}'_i, \vec{x}_o) F_{t,i}(\eta, \hat{\omega}'_i) \Phi_{\text{laser}}. \end{aligned}$$

We assume hereby that an infinitesimally small laser beam is incident at surface location \vec{x}'_i from $\hat{\omega}'_i$. Φ_{laser} is the radiant flux emitted by the laser which we assume to be constant. $\delta(\vec{x}_i - \vec{x}'_i)$ denotes the Dirac delta function.² In order to compute the Fresnel factors $F_t(\eta, \hat{\omega})$ that represent an approximation to the real behavior of a surface [Schlick94] we need to know the relative index of refraction η . We follow here Jensen et al. [Jensen01c] and set $\eta = 1.3$ but acknowledge that this is not exact.

Given the observations of the HDR video camera, we can now determine $R_d(\vec{x}'_i, \vec{x}_o)$ for all visible surface locations.

7.4.1 Acquisition Setup

Figure 7.5 shows an overview over the acquisition setup in the MPI acquisition lab: The target object is illuminated by a laser projection system (see Section 3.3.2) that sweeps a laser beam over the object's surface. The object is placed on a turntable in order to illuminate as many surface regions as possible. The object's response to the illumination is recorded by the Lars III high-dynamic range video camera (Section 3.2.4). The camera is manually placed at different locations relative to the projection system in order to minimize occlusions and to record the object from all sides. We rely on the repeatability of the laser projector and the turntable to ensure that identical surface positions are lit for all camera positions.

²The Dirac delta function is the derivative of the Heavyside step function $H(x)$ with $H(x) = 0$ for $x < 0$, $H(x) = 1/2$ for $x = 0$, $H(x) = 1$ for $x > 0$. We use here the following property of the delta function: $\int_{-\infty}^{\infty} f(x)\delta(x-a)dx = \int_{a-\epsilon}^{a+\epsilon} f(x)\delta(x-a)dx = f(a)$. A delta function with non-scalar arguments can be expanded as a product of scalar delta functions, e.g., $\delta(\vec{x} - \vec{x}') = \delta(x-x')\delta(y-y')\delta(z-z')$ for $\vec{x} = (x, y, z)^T$.




Figure 7.5: *Acquisition setup. The laser projector illuminates the target object (resting on the turntable). The scene is observed by the HDR video camera from several positions. The two spotlights visible on both sides of the camera in the right image are used to illuminate the object during silhouette detection.*

We record one image per view (a combination of turntable and camera position) where the object is fully illuminated by additional spotlights for registration of the object with respect to the camera. The whole acquisition runs at 15 fps (half the camera speed) as we discard all frames during which the laser position was changed. Additional breaks occur when the laser color is changed (due to the warmup curve of the lasers) or when the turntable is rotated.

7.4.2 Geometric and Photometric Calibration

We performed a geometric camera calibration [Bouguet03] to recover the intrinsic parameters of the HDR camera as described in Section 3.2.2. Due to the lack of features on translucent objects and the ease of silhouette detection caused by the global light transport, we use a silhouette based registration algorithm [Lensch01] to recover the pose of the target object relative to the camera. Given a non-degenerated set of laser hit points on the object’s surface and the corresponding deflection settings of the laser projector, we are able to recover the position of the laser projector relative to the setup.

For the photometric calibration, we rely on the overall linearity of the camera output which is confirmed by the response curve given in Section 3.2.4. We furthermore assume that the radiant flux emitted by the laser is constant over time. We then need to perform a white balancing step for the individual lasers taking the spectral response of the camera into account. To this end, we sequentially illuminate a white, scattering calibration target with the three lasers and sum up the contribution of all image pixels.

7.5 Efficient Data Access

A typical acquisition yields an uncompressed data volume of several hundreds of gigabytes. (Data compression can reduce the size of the raw data to a few tens of gigabytes while our final models are typically only a few hundreds of megabytes.) It is therefore mandatory to use efficient algorithms to access the input data when estimating our hierarchical model. In this section, we describe the measures we took to speed up this access before we describe the post-processing of the global and local term in Sections 7.6 and 7.7, respectively.

Each complete acquisition consists of a small number of views (combinations of camera position and turntable rotation, typically 20–30 views). Most of the essential information for the further processing steps is constant per view and needs to be computed only once as the camera observes the object always from the same perspective. This information includes the position of the object and the laser projector relative to the camera. We precompute the Fresnel term for all visible surface locations and assign all pixels in the input image a confidence value based on the viewing and lighting directions that is used as weighting factor for the input data. We also reject at this stage all pixels that are close to a specular highlight or seen under grazing angles and generate the mapping from our texture atlas into the input images. The mapping from vertices to image coordinates and the visibility are precomputed as well.

In estimating our hierarchical model, we evaluate all images for a given turntable and laser projector position at the same time. These images show the object under identical illumination conditions and contain all information that is available for a specific illumination condition. We first decide whether the image tuple is valid, i.e., whether the laser spot is visible, and compute the laser spot position in the image with sub-pixel accuracy. From this information, we can derive the 3D position of the laser spot on the object and the corresponding texture atlas coordinate. Note that Godin et al. [Godin01] have shown that 3D laser range scanning of translucent objects suffers from a systematic bias since the peak location is moved for varying viewing and lighting directions. The computed 3D laser spot position will therefore not necessarily correspond to the true incident laser position but will be the best estimate derived from the input images.

The data can now be efficiently resampled to derive the global and local term as given in the following sections using the precomputed information. Our current implementation processes the input data streams with up to 50 images per second on a PC with a 3 GHz Xeon CPU which is more than three times the speed of the acquisition.




Figure 7.6: *The global term for the starfruit model. Left: A single column of the global throughput factor matrix. The green vertex in the center marks the point of incidence, red areas denote missing values. Right: The same vertex after interpolation. Note that values on the ridge are interpolated consistently.*

7.6 Global Term Post-Processing

As discussed in Section 7.2.3, the object’s diffuse subsurface reflectance function $R_d(\vec{x}_i, \vec{x}_o)$ away from the point of incidence is represented with a global term as a matrix of throughput factors \mathbf{F} . During resampling, \mathbf{F} is filled with subsurface reflectance functions – one function per vertex stored in a column of \mathbf{F} . In general, the columns of \mathbf{F} are only partially filled since $R_d(\vec{x}_i, \vec{x}_o)$ is only available for a finite set of incoming surface locations \vec{x}_i . Furthermore, only a subset of exiting surface locations \vec{x}_o is recorded due to incomplete coverage and occlusions.

The task in post-processing is therefore to estimate the complete throughput factor matrix \mathbf{F} based on the available measurements. Post-processing combines individual observations for a specific point of incidence \vec{x}_i , it interpolates within a single subsurface reflectance function (a column of \mathbf{F}) and between subsurface reflectance functions of neighboring vertices. During interpolation the distribution of energy within the material needs to be taken into account. In homogeneous material the energy falls off exponentially with distance from the point of incident light according to the diffusion approximation. However, our method is aimed at inhomogeneous objects with behavior which deviates from the smooth diffusion approximation. Figure 7.6 shows an individual subsurface reflectance function for the starfruit model before and after interpolation of missing values. Figure 7.7 shows renderings with the full matrix \mathbf{F} before and after interpolation.




Figure 7.7: *The global term for the starfruit model. Left: Irradiance. Middle: The global term before interpolation. Missing data leads to artifacts. Right: The global term after interpolation. The artifacts are reduced and missing data is filled in.*

7.6.1 Data Resampling

A high-dynamic range image of the object is already a scaled subsurface reflectance function for light entering at a point \vec{x}_i . The data consists however of samples at discrete locations of the image plane and covers only the part of the surface which is in a view. Knowing the 3D geometry of the measurement setup allows us to resample the data on the object's surface. The surface is represented as a triangular surface mesh where each triangle represents a similar surface area. Resampling then is the look-up of the bi-linearly interpolated image intensity at each vertex location. The position of the incident light needs also to be resampled and is assumed to contribute to the three vertices of the enclosing triangle. We weight the light according to the barycentric coordinates of the point inside the triangle. We combine measurements of the same subsurface reflectance function observed from different viewpoints and with varying laser colors. The result of the resampling are columns of the throughput factor matrix \mathbf{F} in RGB color space with some missing entries due to unobserved surface area and completely missing columns for vertices on the surface which were never lit. The interpolation of the matrix of throughput factors which addresses these cases is described next.

7.6.2 Interpolation

Interpolation of the throughput factors within a column of the matrix is the task of function interpolation on an irregularly meshed surface. In the diffusion approximation for homogeneous media the function is a sum of exponential decays

which non-trivially depend on distance. Jensen et al. [Jensen01c] report difficulties in fitting the diffusion approximation with dipole lighting to the measured responses of homogeneous material samples. The function is more complicated for inhomogeneous objects and, in our case the interpolation has to potentially fill in large unseen areas. We conclude that function fitting seems inappropriate. Instead, we use filtering of throughput factors on the mesh (similar to mesh fairing [Taubin95]) with different averaging filter kernels. The edge length in the mesh is taken into account as by Desbrun et al. [Desbrun99]. The filter operates on the logarithm of the transfer coefficients because of the exponential decay of the subsurface reflectance function. The filtering operates on color information which originates from separate images with sequential laser illumination. Therefore, the noise in the color information will affect luminance and chrominance. We choose the *CIE YU*V** model [CIE86] since it yields linear luminance values and allows interpolation of chrominance values in a near-linear space with respect to the human visual system.

On the symmetric matrix \mathbf{F} we fill all the missing entries $\mathbf{F}_{r,k}$ in columns k where we observe the incident light point $\mathbf{F}_{k,k}$. Filling is performed by weighted averaging fixing the observation values. In particular, we solve the following iteration (similar to Perona and Malik [Perona90] but on a mesh domain)

$$\mathbf{F}_{r,k}^{t+1} = \mathbf{F}_{r,k}^t + \alpha(1 - c_r) \cdot \sum_{n \in \mathcal{N}} \psi \left(\frac{\mathbf{F}_{n,k} - \mathbf{F}_{r,k}}{e_{n,r}}, \sigma \right)$$

The neighbors n of a vertex r are its one-ring neighborhood \mathcal{N} connected with edges of length $e_{n,r}$. The weighting function ψ can be understood as the derivative of an error norm ρ with the scale parameter σ [Sapiro01]. We employ least squares with $\psi(x, \sigma) = 2 * x / \sigma^2$ and median filtering with $\psi(x) = \text{sign}(x)$. The choice of *confidence* c_r in the current transfer coefficient $\mathbf{F}_{r,k}$ controls the update rate and $c_r = 1$ keeps the existing measurement fixed. We observed satisfactory results with this approach. For highly non-smooth transfer matrices image inpainting techniques (e.g., [Bertalmio00]) may be able to interpolate structure as well.

Subsurface reflectance functions for vertices for which the laser did not reach the triangle surface fan are interpolated iteratively from neighboring vertices. This is justified since a throughput factor $\mathbf{F}_{k,c}$ far from the point of illumination is typically similar to the throughput factor $\mathbf{F}_{k,n}$ connecting the same vertex k with incident light positions n close to the position of illumination c . The neighborhood \mathcal{N} is the one-ring of the vertex c . This approach will break down for the diagonal entry $\mathbf{F}_{k,k}$ and close-by points. The neighborhood \mathcal{N} of diagonal elements $\mathbf{F}_{k,k}$ are the neighboring diagonal $\mathbf{F}_{n,n}$ of the one-ring. We define the vertices of the one-ring neighborhood as close points and blend diagonal interpolation and far interpolation. We define diagonal *neighbors* of a close point as the

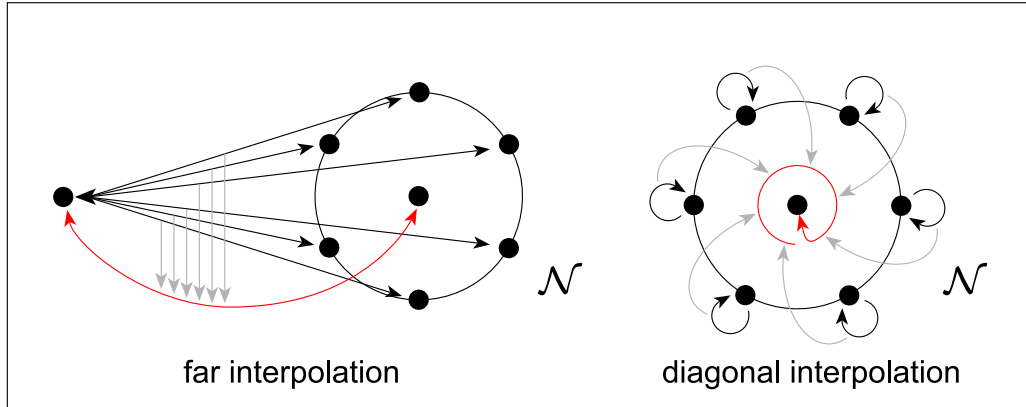


Figure 7.8: Far and diagonal interpolation of throughput factor matrix \mathbf{F} . The throughput factor shown in red is interpolated based on the neighboring factors shown in black.

distance-weighted average response of the vertices of the one-ring neighborhood of diagonal neighbors. (This is similar to the interpolation of Green’s functions of an elastic solid described by Pai et al. [Pai01]). Far and diagonal interpolation are illustrated in Figure 7.8.

We ensure symmetry of the throughput factor matrix by setting $\mathbf{F}^{t+1} = \frac{1}{2}(\mathbf{F}^t + (\mathbf{F}^t)^T)$ between each kind of interpolation, as well as at the beginning and the end of the post-processing. The interpolation for missing entries in a column of \mathbf{F} and for missing complete columns is achieved with the same filtering framework. The differences between the two tasks are limited to different neighborhood definitions. Our filtering framework could also be easily extended to diffusion solved by forward Euler iterations and diffusion solved with implicit schemes [Desbrun99].

7.6.3 Modulation Texture

The texture values of T_ρ are derived during resampling from object regions distant from the point of incidence of the illumination. For a texel in T_ρ that is visible and smoothly lit in an input image we record its RGB value $t_{u,v}$. We select the appropriate mesh triangle and interpolate t_{int} based on the throughput factors at the corresponding vertex locations. The weighted average of the ratio $t_{u,v}/t_{\text{int}}$ over all appropriate input images is stored in T_ρ . At rendering time, we multiply the global response with T_ρ . The construction ensures that the radiosity remains unchanged at mesh vertices and is consistently modulated at all other locations. Figure 7.9 shows as an example the recovered detail texture for the horse model.




Figure 7.9: Left: *Texture atlas with borders for the horse model.* Right: *Recovered modulation texture (color coded version of the green channel: blue areas become darker, red areas become brighter, green areas remain unchanged).*

7.7 Local Term Post-Processing

The filter kernels are estimated in the texture atlas domain based on images where the laser spot is visible. In general, the peak of the laser spot will fall between the discrete pixel locations in the texture atlas. In the following, we describe the resampling necessary to ensure accurate filter kernels.

7.7.1 Data Resampling

Figure 7.10 (left) shows an example of the peak of a filter kernel mapped to the nearest neighbor texel. All texels in an 8×8 pixel neighborhood around the peak location are plotted according to their 3D distance to the peak texel. A clearly structured fall-off is not visible. If we however plot them according to their distance to the sub-pixel location of the peak (see Figure 7.10, right), the general shape of the fall-off of subsurface scattering material is recognizable (see Figure 2.11 for a comparison to the dipole approximation with arbitrary parameters). The measurement data also reveals strong variations due to surface detail at a given texel location. Surface detail is associated with a texel location and should not be resampled to preserve the sharpness of features.

In order to recover plausible filter kernels, we therefore shift the peak location to all four neighboring pixel locations while keeping the surface details spatially fixed. To separate the illumination from the surface detail, we fit the approxima-




Figure 7.10: *Building the filter kernels. Left: Location of the laser peak mapped to its nearest neighbor texel. The plot shows monochrome luminance values from the 8×8 texel neighborhood around the highlight. The horizontal axis is the distance in mm to a neighboring texel in the texture atlas. The values show a fall-off with distance but no clear structure. Right: If we plot the same data over the distance to the sub-pixel peak location, a fall-off resembling the sum of two exponential functions appears (see Figure 2.11 for comparison to the dipole approximation [Jensen01c]).*

tion

$$m(d) = c_1 \cdot e^{\alpha_1 d} + c_2 \cdot e^{\alpha_2 d}$$

to the data. d is the 3D distance from the peak location, c_1 and c_2 are fit per color channel, α_1 , and α_2 are identical for all color channels. In order to achieve a stable fit we first set c_2 to 0 and fit α_1 and c_1 in log space to all data points with $d > \xi_1$. We then fix α_1 and c_1 and fit the remaining parameters to all data points with $d < \xi_2$ ($\xi_1 = 3 \text{ mm}$ and $\xi_2 = 1 \text{ mm}$ for the example in Figure 7.10). This typically results in a stable fit of the exponents. In the next step, we fit only the color values c_1 and c_2 to all data points. The difference between the measured data points and $m(d)$ is then encoded in a 8×8 pixel multiplicative correction texture C .

In order to compute the four filter kernels for the neighboring vertices, we evaluate $m(d)$ for their respective center location and divide the result values by C . If some pixels in the neighborhood were not visible in the input images, we can interpolate them using $m(d)$.

7.7.2 Interpolation

Interpolation of filter kernels is performed by vector filtering over the surface of the object. Each 7×7 filter kernel is represented as a 49-vector. We use the same

filtering framework as for the transfer coefficients \mathbf{F} but on the texture atlas domain instead of the triangular mesh. The neighborhood \mathcal{N} in the texture atlas contains a texel's four-neighbors except for texture atlas boundaries. The boundaries need to be treated specially in order to ensure the correct neighborhood information. Each boundary texel links therefore to a corresponding non-boundary texel in a different map of the atlas that is used instead of the boundary texel during filtering. Filter kernels also differ in size since different areas in the atlas are not isometric. The functional approximation allows for an easy interpolation independent of size at texture map boundaries.

Beside the vector of function parameters, we also propagate the multiplicative correction textures C . Typically, most of the multiplicative textures overlap in our examples due to their size and the laser spot sample spacing so that little interpolation is necessary.

7.8 Rendering

We carefully chose the implementation of the hierarchical BSSRDF model to facilitate interactive rendering using a modern PC supported by graphics hardware (see Section 7.2.3). Interactive rendering of both, synthetic and acquired datasets, is therefore straightforward: We first compute an illumination map and split the computation into two branches. One branch derives the irradiance at each vertex from the illumination map and computes the global response. The second branch evaluates the local response by filtering the illumination map. Both branches can be executed in parallel on a dual processor machine. Finally, the global and the local response are combined. Figure 7.11 shows an overview over the rendering system.

Note however that the acquisition method for translucent objects is not limited to this particular rendering technique. Although it does not contain the basic parameters needed for a physical simulation, it can still be treated as a “black box” by a Monte Carlo raytracer or a photon mapping system. Both require only a probabilistic evaluation of the local and global throughput factors which can be sped up using an inverted cumulative density function [Pitman92]. The model can also be directly integrated in the preprocessing phase of precomputed radiance transfer [Sloan03], the local illumination framework of Hao et al. [Hao03] or the adaptive link technique by Carr et al. [Carr03]. The model can substitute an evaluation of the dipole approximation [Jensen01c] although the size of the data structure makes an evaluation on current graphics hardware difficult.




Figure 7.11: Workflow during rendering: first, incident illumination is computed and projected into the texture atlas. The resulting illumination map is processed in two ways. The local response (left branch) is computed by filtering the incident illumination map with spatially varying 7×7 texel-to-texel filter kernels. The global response (right branch) is computed by projecting the illumination to the mesh vertices and multiplying the vertex irradiance vector with a vertex-to-vertex throughput factor matrix. Finally the global and the local response are combined.

7.8.1 Computing the Illumination

Since we require an illumination map based on the texture atlas, we need to convert the illumination from object space into texture space. Furthermore, we have to integrate over the illumination map in order to compute the irradiance at the vertices.

The illumination map can easily be created by rendering the object not using its 3D vertex positions but its 2D texture coordinates from the texture atlas. The result is a 2D representation of the incident illumination in the form of an illumination map texture. Some care has to be taken that the lighting is correctly computed even though the geometry is projected into 2D. We do this by computing the lighting in a vertex shader [Lindholm01] using the original 3D vertex position and normal. Furthermore we include a Fresnel term in the lighting calculations for which we use Schlick's approximation [Schlick94] (see Section 2.4.1) which can also be computed in the vertex shader. The rendered illumination map is then read back from the frame buffer and stored. In Figure 7.11, the illumination on the object and the corresponding illumination map derived using the texture atlas are shown.

Once the irradiance at each texel (u, v) is computed, we can integrate it to obtain the irradiance for each vertex. In order to distribute the texel irradiance correctly to vertex irradiance, we follow Equation 7.5. The vertex irradiance E_i^g is given as

$$E_i^g = \sum_{(u,v)} \tilde{\psi}_i^g(u, v) E(u, v) A(u, v), \quad (7.14)$$

the sum over all texels in the illumination map times the value at the current texel of the dual basis function corresponding to the vertex, times the area $A(u, v)$ of the model surface $S(u, v)$ covered in the texel. As a result, the illumination at each texel is distributed to exactly three different vertices. The weights $\tilde{\psi}_i^g(u, v) A(u, v)$ are precomputed into an image of the same resolution as the texture atlas.

The same distribution mechanism is also applied to obtain the second radiosity vector B_i^d . This time, the weights $w_i(u, v)$ are used instead of the dual basis function. Distributing the illumination map to two vectors instead of just one does not significantly influence rendering performance.

7.8.2 Radiosity Reconstruction

Given the irradiance at the texels and vertices, the radiosity at the object's surface can be computed according to Equation 7.11 as the sum of local term B^l , global term B^{g0} and the radiosity vector B^d .

The global response is calculated by multiplying the throughput factor matrix \mathbf{F} with the irradiance at the vertices:

$$B_v^g = \sum_i E_i^g F_{iv}.$$

As previously discussed, the radiosity at a particular point \vec{x} on a triangle is interpolated using the barycentric basis $\psi_v^g(\vec{x})$ with respect to the vertices of the triangle.

Depending on the number of vertices n_v (given by the geometry model) and its scattering behavior, the entries in the matrix \mathbf{F} may drop to very small values. In these cases a full $n_v \times n_v$ matrix-vector multiplication may be too costly for interactive rendering. A higher rendering speed can be achieved by ignoring form factors below a threshold and using a sparse matrix-vector multiplication. In our experiments the overhead of representing a sparse matrix paid off if 40 percent of the form factors could be ignored.

According to the factorization described in Section 7.2.3, the local response to a light impulse impinging on the surface at point \vec{x} is represented by a 7×7 filter kernel $K_{(u,v)}(s, t)$ in texture space centered at the corresponding texel (u, v) . The resulting radiosity due to local response is thus computed by a kind of convolution for each texel:

$$B^l(\vec{x}) = K_{(u,v)}(s, t) \otimes E(s, t) = \sum_{(s,t) \in 7 \times 7} K_{(u,v)}(s, t) E(s, t).$$

Since the filter kernels for each texel are different, we currently implement this step in software. After the convolution, the filtered illumination map is reloaded as texture and applied during the final composition.

Local and global response are combined in one hardware-accelerated rendering step using multi-texturing. Register combiners are set up in such a way that the vertex radiosity is multiplied with the surface texture T_ρ and at the same time the filtered illumination map corresponding to the local response is added.

7.9 Acquisition Results

We acquired the following objects to demonstrate our approach: An alabaster horse sculpture, a rubber duck and a starfruit (carambola). Figure 7.12 shows all objects under indoor illumination, as well as illuminated by the three lasers (i.e., the object's impulse response).

The alabaster horse has strong variation in its subsurface scattering properties and complex geometry. It consists of regions with varied translucency and contains cracks inside the material. The base material of the rubber duck is uniform,



Figure 7.12: Top: *The test objects under indoor illumination.* Bottom: *Test objects illuminated by a white laser beam (all three lasers combined). The laser speckle is strongly visible due to the low fill factor of the camera.*

	Horse	Duck	Starfruit
# input views	24	25	20
# input images	1,065,744	541,125	401,220
input size (compressed)	31G	14G	12G
acquisition time	20.5h	11.25h	8h
# vertices	8924	5002	5001
# filter kernels	82,390	115,151	112,538
processing time (resampling)	7.8h	3.6h	3.4h

Table 7.1: *Some statistics about the acquired models.*



Figure 7.13: *A model of an alabaster horse sculpture acquired by our method. Differences in material are clearly visible when the model is lit from behind.*

strongly scattering rubber. The beak and the eyes are painted on the outside with glossy, opaque paint. Incoming light is so strongly diffused at the rubber surface that the assumption of diffuse, multiple scattering remains valid although the duck is empty inside. The starfruit is an example of a translucent biological object with concave geometry and a relatively specular surface.

Some relevant details about the acquired models are given in Table 7.1.

Alabaster Horse Sculpture

Our acquisition system captures models with local variation and significant detail. The model of the alabaster horse sculpture in Figure 7.13 nicely shows areas of differing translucency. There is a more opaque region at the head around the eyes while the muzzle is quite translucent. The left side of the head (when viewed from the muzzle) has various opaque areas towards the support and in the mane. There is a crack running top to bottom in the center at the neck as well. The model captures these volume effects.

Figure 7.14 demonstrates how the local light transport adds surface detail and gives an impression of the varying optical densities at the object's surface. The




Figure 7.14: *The horse model. Top: Irradiance, local light transport by filter kernels, and global light transport due to the throughput factor matrix. Bottom left: Combined rendering of local and global term. The local light transport adds surface detail and gives an impression of the material properties at the object's surface whose basic structure is also visible in the global term. Bottom right: Photograph of the real object under similar conditions. A slide projector was used to produce the sharp shadow boundary.*




Figure 7.15: *The duck model illuminated with a spot light. Left: Local filter kernels. Center: Global throughput factors. Right: Combined rendering.*




Figure 7.16: *Details of the duck model. Left: Head of the duck illuminated with a small spotlight from the back. The beak and the eyes are clearly visible. The spot near the tail is caused by diffuse reflection from the back of the head during the acquisition. Middle: Body of the duck model illuminated with a small spotlight from the back. Right: Photograph of the duck taken under similar conditions. The images show the specific property of the real duck that light illuminating the head is scattered mainly inside the head and light illuminating the body is scattered mainly inside the body.*

structure is also visible in the global throughput factor matrix albeit smoother. The side-by-side comparison of a rendering of the model and a photograph shows that our method recovers fine structure detail (see lower images in Figure 7.14). The highly translucent veins as well as the strongly diffuse patches which are visible in the photograph are present in our model. The slightly brownish region in the center is also captured well by our model in both color and shape.



Figure 7.17: Starfruit model with global transfer function and local term (7×7 filter kernels). Left: The hierarchical model of the starfruit is illuminated from the front right. Right: The same model illuminated from the back. The global transfer function F alone is rendered in Figures 7.6 (illumination at a single vertex) and 7.7.

Rubber Duck

The duck model in Figures 7.15 and 7.16 shows how our method can deal with completely heterogeneous objects without representing the material distribution within the object explicitly. The rubber duck is made of regular material overall but head and body appear nearly separate when lit from behind. The wings block more light probably because of the extra thickness. The beak and eyes are marked with an opaque paint layer. The model captures these major deviations from homogeneous behavior while being smooth overall.

Starfruit

The starfruit is visually quite interesting because of its shape. The model in Figure 7.17 is of good fidelity despite missing capture data due to the geometrically complex shape. Additional input images would fill gaps better than our data interpolation technique.

7.10 Conclusion

We presented a complete framework for the acquisition and rendering of translucent objects. The framework is based on a hierarchical model of the diffuse sub-

surface scattering reflectance function R_d that models light transport in optically dense, scattering material. We built an acquisition system to acquire R_d for real world objects. In a post-processing step, missing information is consistently interpolated and noise artifacts of the acquisition are reduced. The resulting models are suitable for the proposed hardware-accelerated rendering approach but can also be used by a variety of rendering systems. We validated the approach by acquiring three translucent objects with strongly differing behavior.

With the notable exception of Jensen et al. [Jensen01c], most current rendering systems for translucent materials suffer from a lack of realistic input data. We are now able to supply real world data which can not only be used to drive the rendering systems but also to validate common assumptions such as the validity of the dipole approximation. Artificial models are definitely able to mimic the overall behavior of translucent objects but they are often too regular and smooth compared to real objects. We believe that the gain in realism makes the effort to acquire real translucent objects worthwhile.

Chapter 8

Summary, Conclusion and Future Work

In the following we first briefly summarize the contributions of this dissertation. We then draw some conclusions and end with an outlook on future work.

8.1 Summary

We first introduce a method to determine the accuracy of a 3D range scanner based on the concept of modulation transfer functions. In contrast to approaches that analyze the quality by acquiring a smoothly varying surface, the method determines the system behavior when acquiring a sharp edge (or other sharp features in an object). Since this is a common case that nevertheless poses difficulties for many 3D range scanning systems, the method is of high practical relevance.

Next, we combine an appearance acquisition and reproduction system for objects showing only local reflection with color management technology in order to reference the acquired model to a well-defined color space and to ensure best possible reproduction with the given input and output devices. One of the key ideas of this approach is to introduce color management into the high-dynamic range imaging computation. This idea is naturally not limited to the presented application but can also be applied in many other acquisition methods.

We propose an acquisition system capable of acquiring the near-field of a light source that performs optical prefiltering before sampling. Due to the geometry of the setup, we can acquire a much wider field of view than previous approaches thereby speeding up the acquisition. Using optical prefiltering also solves the aliasing problem for general light field acquisition by decoupling the filtering function from the optical system used for capture. On the rendering side,

the acquired light source model can be either integrated into a global illumination system or used in a hardware-accelerated renderer.

We finally present a first acquisition approach for translucent objects with spatially varying properties that takes the specific behavior of diffuse light transport inside translucent materials into account. The underlying hierarchical model can be adapted to the structure of the object, the resolution of the acquisition system and to the requirements of the final rendering application. We implement a hardware-accelerated rendering system for the acquired objects and show how models can be integrated into many other current rendering systems for translucent objects.

8.2 Conclusion

The main motivation of the work presented in this thesis was to acquire realistic models of real world objects and real world light sources for use in computer graphics applications. Since the capabilities of general purpose computers as well as specialized graphics hardware are increasing rapidly, more and more complex models of real world objects can be included in a scene without slowing down the rendering process. This leads to a rising demand for such models, e.g., in the area of translucent objects. While a large number of specialized rendering algorithms for translucent objects have been developed in recent years, there were only very few methods available to determine the correct physical parameters.

In addition, acquisition methods themselves benefit from better hardware capabilities for acquiring, storing, and processing input data. The properties and capabilities of the camera systems used in this dissertation are a prominent example for the improvements made within a couple of years. Specialized high-dynamic range camera systems are only now becoming available. The goal of an acquisition system will therefore always be to fulfill the demands of the computer graphics community in terms of complexity and accuracy of the acquired datasets while at the same time making good use of the available technology.

The question whether the achievable results justify the effort is frequently posed in the context of real world acquisition techniques. But current experience clearly demonstrates that renderings of purely synthetic scenes look either artificial or too perfect (at least when compared to a similar real scene or a photograph thereof) or require an excessively large effort for modeling which should be seen in proportion to the effort for acquisition. This is all the more true if an accurate model of one specific real object is required – a case frequently occurring in the cultural heritage field. We expect that in the long run both, synthetic and acquired models, will be used depending on the requirements of the applications.

8.3 Future Work

Future work for any real world acquisition system will always be to make it faster, cheaper, more robust, easier to use, more automatic, applicable to a larger class of objects, and achieving better quality results. These issues are continuous challenges for researchers and application engineers. But more specifically, there are several directions of future work that could be followed – some of them were already briefly mentioned in the conclusions of the respective chapters.

Conclusions about the interpretation of modulation transfer functions of 3D range scanners have until now only been drawn by analogy to the case of a digital camera. In order to be able to draw well-founded conclusions, an extensive study involving a large number of 3D range scanning systems would be required. Furthermore, tests and theoretical analyses have to be performed in order to determine the requirements of the scanning post-processing software in terms of MTF.

The examples of color calibrated BRDF acquisition shown in Chapter 5 suggest that there is still room for improvement in the processing pipeline. In particular, the color managed high-dynamic range imaging system used to acquire color calibrated BRDFs depends directly on the properties of the color management system that could probably be better tuned to this specific application. Ultimately, one would however like to extend the capabilities of color management systems also to high-dynamic range images in order to link directly with high-dynamic range camera systems.

The central part of the light source acquisition system is the projection into an appropriate basis. The influence of the basis on the quality of the acquisition could be further studied and other sampling strategies could be explored.

The acquisition algorithm for translucent objects assumes purely diffuse light transport to reduce the dimensionality of the problem. Unfortunately, this assumption does not hold for a large number of real world objects leading to artifacts in the acquisition and rendering stages. Including a BRDF to model local light reflection could considerably improve on this. Furthermore, the acquired datasets of real world objects can be used to validate the dipole approximation for translucent materials that is currently the foundation for many rendering systems for translucent objects.

Bibliography

- [Appel68] ARTHUR APPEL. Some Techniques for Shading Machine Renderings of Solids. *SJCC*, pages 27–45, 1968.
- [Ashdown93] IAN ASHDOWN. Near-Field Photometry: A New Approach. *Journal of the Illuminating Engineering Society*, 22(1):163–180, Winter 1993.
- [Ashdown95] IAN ASHDOWN. Near-Field Photometry: Measuring and Modeling Complex 3-D Light Sources. In *ACM SIGGRAPH '95 Course Notes - Realistic Input for Realistic Images*, pages 1–15. ACM, 1995.
- [Bayer76] BRYCE E. BAYER. Color imaging array. US Patent 3,971,065, 1976.
- [Beraldin95] JEAN-ANGELO BERARDIN, SABRY F. EL-HAKIM, AND FRANÇOIS BLAIS. Performance Evaluation of Three Active Vision Systems Built at the National Research Council of Canada. In *Proceedings of the Conference on Optical 3-D Measurements Techniques*, pages 352–361, 1995.
- [Bernardini02] FAUSTO BERNARDINI AND HOLLY RUSHMEIER. The 3D Model Acquisition Pipeline. *Computer Graphics Forum*, 21(2):149–172, 2002.
- [Bertalmio00] MARCELO BERTALMIO, GUILLERMO SAPIRO, VINCENT CASELLES, AND COLOMA BALLESTER. Image Inpainting. In *SIGGRAPH 2000*, pages 417–424, 2000.
- [Blais03] FRANÇOIS BLAIS. A Review of 20 Years of Range Sensor Development. In *Videometrics VII, Proceedings of SPIE-IS&T Electronic Imaging*, volume 5013, pages 62–76, 2003.

- [Blasi93] PHILIPPE BLASI, BERTRAND LE SAËC, AND CHRISTOPHE SCHLICK. A Rendering Algorithm for Discrete Volume Density Objects. *Computer Graphics Forum*, 13(3):201–210, 1993.
- [Blinn77] JAMES F. BLINN. Models of Light Reflection for Computer Synthesized Pictures. In *Proceedings of the 4th annual conference on Computer graphics and interactive techniques*, pages 192–198, 1977.
- [Born99] MAX BORN AND EMIL WOLF. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press, 7. (expanded) edition, 1999.
- [Bouguet03] JEAN-YVES BOUGUET. Camera Calibration Toolbox for Matlab. See <http://www.vision.caltech.edu/bouguetj>, 2003.
- [Burns01] PETER D. BURNS AND DON WILLIAMS. Image Resolution and MTF for Digital Cameras and Scanners. In *Tutorial Notes, IS&T PICS Conference 2001*, 2001.
- [Carr03] NATHAN A. CARR, JESSE D. HALL, AND JOHN C. HART. GPU Algorithms for Radiosity and Subsurface Scattering. In *Graphics Hardware 2003*, pages 51–59, 2003.
- [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.
- [Chuang00] YUNG-YU CHUANG, DOUGLAS E. ZONGKER, JOEL HINDORFF, BRIAN CURLESS, DAVID H. SALESIN, AND RICHARD SZELISKI. Environment Matting Extensions: Towards Higher Accuracy and Real-time Capture. In *SIGGRAPH 2000*, pages 121–130, 2000.
- [CIE86] Colorimetry. Publication CIE No. 15.2, 1986.
- [Cohen93] MICHAEL F. COHEN AND JOHN R. WALLACE. *Radiosity and Realistic Image Synthesis*. Academic Press Professional, 1993.
- [Dachsbacher03] CARSTEN DACHSBACHER AND MARC STAMMINGER. Translucent Shadow Maps. In *Rendering Workshop 2003*, pages 197–201, 2003.

- [Dalton98] GRAHAM DALTON. Reverse Engineering Using Laser Metrology. *Sensor Review*, 18(2):92–96, 1998.
- [Daly93] S. DALY. The Visible Differences Predictor: An Algorithm for the Assessment of Image Fidelity. In A.B. Watson, editor, *Digital Image and Human Vision*, pages 179–206. Cambridge, MA: MIT Press, 1993.
- [Dana99] K. DANA, B. VAN GINNEKEN, S. NAYAR, AND J. KOENDERINK. Reflectance and Texture of Real-World Surfaces. *ACM Transactions on Graphics*, 18(1):1–34, January 1999.
- [Debevec97] PAUL E. DEBEVEC AND JITENDRA MALIK. Recovering High Dynamic Range Radiance Maps from Photographs. *Proceedings of SIGGRAPH 97*, pages 369–378, August 1997.
- [Debevec00] PAUL DEBEVEC, TIM HAWKINS, CHRIS TCHOU, HAARM-PIETER DUIKER, WESTLEY SAROKIN, AND MARK SAGAR. Acquiring the Reflectance Field of a Human Face. *Proceedings of SIGGRAPH 2000*, pages 145–156, July 2000.
- [Desbrun99] MATHIEU DESBRUN, MARK MEYER, PETER SCHRÖDER, AND ALAN H. BARR. Implicit Fairing of Irregular Meshes using Diffusion and Curvature Flow. In *SIGGRAPH 1999*, pages 317–324, 1999.
- [Devlin02] K. DEVLIN, A. CHALMERS, A. WILKIE, AND W. PURGATHOFER. Tone Reproduction and Physically Based Spectral Rendering. In *Eurographics 2002: State of the Art Reports*, pages 101–123, 2002.
- [Dorsey99] JULIE DORSEY, ALAN EDELMAN, HENRIK WANN JENSEN, JUSTIN LEGAKIS, AND HANS KOHLING PEDERSEN. Modeling and Rendering of Weathered Stone. In *SIGGRAPH 1999*, pages 225–234, 1999.
- [Dutr 03] PHILIP DUTR , PHILIPPE BEKAERT, AND KAVITA BALA. *Advanced Global Illumination*. A.K. Peters, 2003.
- [Feng93] SHESHAO FENG, FANAN ZENG, AND BRITTON CHANCE. Monte Carlo Simulations of Photon Migration Path Distributions in Multiple Scattering Media. In *Photon Migration and Imaging in Random Media and Tissues, Proc. of SPIE Vol. 1888*, pages 78–89, 1993.

- [Fischer94] TIMOTHY A. FISCHER AND JACK HOLM. Electronic Still Picture Camera Spatial Frequency Response Measurement. In *Proceedings of IS&T's 47th Annual Conference/ICPS*, pages 626–630, 1994.
- [Glassner94] ANDREW S. GLASSNER. *Principles of Digital Image Synthesis*. Morgan Kaufmann Publishers Inc., 1994.
- [Godin01] GUY GODIN, MARC RIOUX, J.-ANGELO BERARDIN, MARC LEVOY, LUC COURNOYER, AND FRANCOIS BLAIS. An Assessment of Laser Range Measurement on Marble Surfaces. In *5th Conf. on Optical 3D Measurement Techniques*, 2001.
- [Goesele00] MICHAEL GOESELE, WOLFGANG HEIDRICH, HENDRIK P.A. LENSCH, AND HANS-PETER SEIDEL. Building a Photo Studio for Measurement Purposes. In *Proceedings of the 5th Conference on Vision, Modeling, and Visualization (VMV-00)*, November 2000.
- [Goesele01] MICHAEL GOESELE, WOLFGANG HEIDRICH, AND HANS-PETER SEIDEL. Color Calibrated High Dynamic Range Imaging with ICC Profiles. In *Proceedings of the 9th Color Imaging Conference Color Science and Engineering: Systems, Technologies, Applications*, pages 286–290, Scottsdale, USA, 2001. Society for Imaging Science and Technology.
- [Goesele03a] MICHAEL GOESELE, CHRISTIAN FUCHS, AND HANS-PETER SEIDEL. Accuracy of 3D Range Scanners by Measurement of the Slanted Edge Modulation Transfer Function. In *4th International Conference on 3-D Digital Imaging and Modeling*, 2003.
- [Goesele03b] MICHAEL GOESELE, XAVIER GRANIER, WOLFGANG HEIDRICH, AND HANS-PETER SEIDEL. Accurate Light Source Acquisition and Rendering. *ACM Trans. Graph.*, 22(3):621–630, 2003.
- [Goesele04a] MICHAEL GOESELE, HENDRIK P. A. LENSCH, JOCHEN LANG, CHRISTIAN FUCHS, AND HANS-PETER SEIDEL. DISCO – Acquisition of Translucent Objects. *ACM Trans. Graph.*, 23(3), 2004. To be published.

- [Goesele04b] MICHAEL GOESELE, HENDRIK P. A. LENSCH, AND HANS-PETER SEIDEL. Validation of Color Managed 3D Appearance Acquisition. In *Proceedings of the 12th Color Imaging Conference on Color Science and Engineering Systems, Technologies and Applications (CIC-04)*. IS&T, 2004. To be published.
- [Gortler96] STEVEN J. GORTLER, RADEK GRZESZCZUK, RICHARD SZELISKI, AND MICHAEL F. COHEN. The Lumigraph. In *Proceedings of the SIGGRAPH 96 annual conference*, pages 43–54. ACM Press, 1996.
- [Granier03] XAVIER GRANIER, MICHAEL GOESELE, WOLFGANG HEIDRICH, AND HANS-PETER SEIDEL. Interactive Visualization of Complex Real-World Light Sources. In *Proceedings of Pacific Graphics 2003*, pages 59–66, 2003.
- [Halle94] MICHAEL W. HALLE. Holographic Stereograms as Discrete Imaging Systems. In *Practical Holography VIII*, volume 2176, pages 73–84. SPIE, February 1994.
- [Hanrahan93] PAT HANRAHAN AND WOLFGANG KRUEGER. Reflection from Layered Surfaces due to Subsurface Scattering. In *SIGGRAPH 1993*, pages 165–174, 1993.
- [Hao03] XUEJUN HAO, THOMAS BABY, AND AMITABH VARSHNEY. Interactive Subsurface Scattering for Translucent Meshes. In *Proc. of the 2003 Symposium on Interactive 3D Graphics*, pages 75–82, 2003.
- [Hebert92] MARTIAL HEBERT AND ERIC KROTKOV. 3D Measurements from Imaging Laser Radars: How Good are They? *Intl. Journal of Image and Vision Computing*, 10(3):170–178, April 1992.
- [Heckbert91] P. S. HECKBERT AND J. WINGET. Finite Element Methods for Global Illumination. Technical Report UCB/CSD 91/643, Computer Science Division (EECS), University of California, Berkeley, California, USA, July 1991.
- [Heidrich98] WOLFGANG HEIDRICH, JAN KAUTZ, PHILIPP SLUSALLEK, AND HANS-PETER SEIDEL. Canned Lightsources. In *Rendering Techniques '98*, pages 293–300. Eurographics, June 1998.

- [Heidrich01] W. HEIDRICH AND M. GOESELE. Image-Based Measurement of Light Sources With Correct Filtering. Technical Report TR-2001-08, Department of Computer Science, The University of British Columbia, 2001.
- [Heikkilä97] J. HEIKKILÄ AND O. SILVÉN. A Four-Step Camera Calibration Procedure With Implicit Image Correction. In *CVPR97*, 1997.
- [Hunt98] R. W. G. HUNT. *Measuring Colour*. Fountain Press, 3rd edition, 1998.
- [ICC03] Specification ICC.1:2003-09, File Format for Color Profiles (Version 4.1.0). International Color Consortium, available from <http://www.color.org>, 2003.
- [IEC99] Multimedia Systems and Equipment – Colour Measurement and Management – Part 2-1: Colour Management – Default RGB Colour Space – sRGB. Publication IEC 61966-2-1, 1999.
- [Ishimaru78] A. ISHIMARU. *Wave Propagation and Scattering in Random Media*. Academic Press, 1978.
- [ISO99] ISO 14524:1999: Photography – Electronic Still-Picture Cameras – Methods for Measuring Opto-Electronic Conversion Functions (OECFs), 1999.
- [ISO00] ISO 12233:2000: Photography – Electronic Still-Picture Cameras – Resolution Measurements, 2000.
- [Janesick01] JAMES R. JANESICK. *Scientific Charge-Coupled Devices*. SPIE, 2001.
- [Jenkins00] DAVID R. JENKINS AND HOLGER MÖNCH. Source Imaging Goniometer Method of Light Source Characterization for Accurate Projection System Design. In *Proc. of SID Conference*, pages 862–865, 2000.
- [Jensen96] HENRIK WANN JENSEN. Global Illumination Using Photon Maps. In *Rendering Techniques '96 (Proceedings of the Seventh Eurographics Workshop on Rendering)*, pages 21–30, New York, NY, 1996. Eurographics, Springer-Verlag/Wien.

- [Jensen98] HENRIK WANN JENSEN AND PER H. CHRISTENSEN. Efficient Simulation of Light Transport in Scenes with Participating Media using Photon Maps. In *SIGGRAPH 1998*, pages 311–320, 1998.
- [Jensen99] HENRIK WANN JENSEN, JUSTIN LEGAKIS, AND JULIE DORSEY. Rendering of Wet Materials. In *Rendering Workshop 1999*, pages 281–290, 1999.
- [Jensen01a] HENRIK WANN JENSEN. A Practical Guide to Global Illumination Using Photon Mapping. In *SIGGRAPH 2001 Course Notes CD-ROM*. ACM, ACM SIGGRAPH, August 2001. Course 38.
- [Jensen01b] HENRIK WANN JENSEN. *Realistic Image Synthesis using Photon Mapping*. A.K. Peters, 2001.
- [Jensen01c] HENRIK WANN JENSEN, STEPHEN R. MARSCHNER, MARC LEVOY, AND PAT HANRAHAN. A Practical Model for Sub-surface Light Transport. In *SIGGRAPH 2001*, pages 511–518, 2001.
- [Jensen02] HENRIK WANN JENSEN AND JUAN BUHLER. A Rapid Hierarchical Rendering Technique for Translucent Materials. In *SIGGRAPH 2002*, pages 576–581, 2002.
- [Kajiya86] JAMES T. KAJIYA. The Rendering Equation. In *Proceedings of SIGGRAPH 1986*, pages 143–150. ACM Press, 1986.
- [Kang03] SING BING KANG, MATTHEW UYTTENDAELE, SIMON WINDER, AND RICHARD SZELISKI. High Dynamic Range Video. *ACM Transactions on Graphics*, 22(3):319–325, 2003.
- [Kawai93] JOHN K. KAWAI, JAMES S. PAINTER, AND MICHAEL F. COHEN. Radioptimization: Goal Based Rendering. In *Proceedings of SIGGRAPH 1993*, pages 147–154, 1993.
- [Keferstein98] CLAUS P. KEFERSTEIN AND MICHAEL MARXER. Testing Bench for Laser Triangulation Sensors. *Sensor Review*, 18(3):183–187, 1998.
- [Klette98] REINHARD KLETTE, KARSTEN SCHLÜNS, AND ANDREAS KOSCHAN. *Computer Vision: Three-Dimensional Data from Images*. Springer, 1998.

- [Kobbelt00] LEIF P. KOBBELT, STEPHAN BISCHOFF, MARIO BOTSCH, KOLJA KÄHLER, CHRISTIAN RÖSSL, ROBERT SCHNEIDER, AND JENS VORSATZ. Geometric Modeling Based on Polygonal Meshes. Technical Report MPI-I-2000-4-002, Max-Planck-Institut für Informatik, July 2000.
- [Kodak98] EASTMAN KODAK. *Kodak Professional DCS 500 Series Digital Cameras – User’s Guide*, 1998.
- [Lafortune93] ERIC P. LAFORTUNE AND YVES D. WILLEMS. Bidirectional Path Tracing. In *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics ’93)*, pages 145–153, Alvor, Portugal, December 1993.
- [Lafortune96] ERIC P. F. LAFORTUNE AND YVES D. WILLEMS. Rendering Participating Media with Bidirectional Path Tracing. In *Rendering Workshop 1996*, pages 91–100, 1996.
- [Lafortune97] E. LAFORTUNE, S.-C. FOO, K. TORRANCE, AND D. GREENBERG. Non-Linear Approximation of Reflectance Functions. In *Proceedings of SIGGRAPH 1997*, pages 117–126, August 1997.
- [Lensch01] HENDRIK LENSCH, WOLFGANG HEIDRICH, AND HANS-PETER SEIDEL. Silhouette-based Algorithm for Texture Registration and Stitching. *Graphical Models*, 63(4):245–262, 2001.
- [Lensch03a] HENDRIK P. A. LENSCH, JAN KAUTZ, MICHAEL GOESELE, WOLFGANG HEIDRICH, AND HANS-PETER SEIDEL. Image-based Reconstruction of Spatial Appearance and Geometric Detail. *ACM Trans. Graph.*, 22(2):234–257, 2003.
- [Lensch03b] HENDRIK P.A. LENSCH, MICHAEL GOESELE, PHILIPPE BEKAERT, JAN KAUTZ, MARCUS A. MAGNOR, JOCHEN LANG, AND HANS-PETER SEIDEL. Interactive Rendering of Translucent Objects. *Computer Graphics Forum*, 22(2):195–206, 2003.
- [Lensch03c] HENDRIK PETER ASMUS LENSCH. *Efficient, Image-Based Appearance Acquisition of Real-World Objects*. PhD thesis, 2003.

- [Lensch04] HENDRIK P. A. LENSCH, MICHAEL GOESELE, AND HANS-PETER SEIDEL. Color-Corrected 3D Appearance Reproduction. 2004. Submitted.
- [Levoy96] MARC LEVOY AND PAT HANRAHAN. Light field rendering. In *Proceedings of the SIGGRAPH 96 annual conference*, pages 31–42. ACM Press, 1996.
- [Lindholm01] ERIK LINDHOLM, MARK J. KILGARD, AND HENRY MORETON. A User-Programmable Vertex Engine. In *Proceedings of ACM SIGGRAPH 2001*, pages 149–158, August 2001.
- [Lu98] R. LU, J. KOENDERINK, AND A. KAPPERS. Optical Properties (Bidirectional Reflectance Distribution Functions) of Velvet. *Applied Optics*, 37(25):5974–5984, September 1998.
- [Lyon02] RICHARD F. LYON AND PAUL M. HUBEL. Eyeing the Camera: Into the Next Century. In *Proceedings of the 10th Color Imaging Conference on Color Science and Engineering Systems, Technologies and Applications (CIC-02)*, pages 349–355. IS&T, 2002.
- [Madden93] BRIAN C. MADDEN. Extended Intensity Range Imaging. Technical report, University of Pennsylvania, GRASP Laboratory, 1993.
- [Malzbender01] TOM MALZBENDER, DAN GELB, AND HANS WOLTERS. Polynomial Texture Maps. In *SIGGRAPH 2001*, pages 519–528, 2001.
- [Mann95] STEVE MANN AND ROSALIND W. PICARD. On Being ‘Undigital’ with Digital Cameras: Extending Dynamic Range by Combining Differently Exposed Pictures. In *IS&T’s 48th Annual Conference*, pages 422–428, 1995.
- [Marschner97] STEPHEN R. MARSCHNER AND DONALD P. GREENBERG. Inverse Lighting for Photography. In *Fifth Color Imaging Conference: Color Science, Systems and Applications*, pages 262–265, 1997.
- [Marschner98] STEPHEN R. MARSCHNER. *Inverse Rendering for Computer Graphics*. PhD thesis, 1998.

- [Marschner99] STEPHEN R. MARSCHNER, STEPHEN H. WESTIN, ERIC P. F. LAFORTUNE, KENNETH E. TORRANCE, AND DONALD P. GREENBERG. Image-based BRDF Measurement Including Human Skin. In *Rendering Workshop 1999*, pages 139–152, 1999.
- [Masselus03] VINCENT MASSELUS, PIETER PEERS, PHILIP DUTRÉ, AND YVES D. WILLEMS. Relighting with 4D Incident Light Fields. *ACM Trans. Graph.*, 22(3):613–620, 2003.
- [Matusik02] WOJCIECH MATUSIK, HANSPETER PFISTER, REMO ZIEGLER, ADDY NGAN, AND LEONARD MCMILLAN. Acquisition and Rendering of Transparent and Refractive Objects. In *Rendering Workshop 2002*, pages 267–278, 2002.
- [Mertens03a] TOM MERTENS, JAN KAUTZ, PHILIPPE BEKAERT, HANSPETER SEIDEL, AND FRANK VAN REETH. Interactive Rendering of Translucent Deformable Objects. In *Rendering Workshop 2003*, pages 130–140, 2003.
- [Mertens03b] TOM MERTENS, JAN KAUTZ, PHILIPPE BEKAERT, FRANK VAN REETH, AND HANSPETER SEIDEL. Efficient Rendering of Local Subsurface Scattering. In *Pacific Graphics 2003*, pages 51–58, 2003.
- [Miles02] BRIAN MILES, JAY LAND, ANDREW HOFFMAN, WILLIAM HUMBERT, BRIAN SMITH, ANDREW HOWARD, JOE COX, MIKE FOSTER, DAVE ONUFFER, SAMMIE THOMPSON, TOM RAMRATH, CLARKE HARRIS, AND PAUL FREEDMAN. Field Testing Protocols for Evaluation of 3D Imaging Focal Plane Array Ladar Systems. In *Laser Radar Technology and Applications VII, Proceedings of SPIE*, volume 4723, pages 43–56, 2002.
- [Miller98] G. S. P. MILLER, S. RUBIN, AND D. PONCELEON. Lazy Decompression of Surface Light Fields for Precomputed Global Illumination. In *Rendering Workshop 1998*, pages 281–292, 1998.
- [Minolta] MINOLTA. Minolta Web Site, www.minolta-3d.com.
- [Moore65] GORDON E. MOORE. Cramming More Components onto Integrated Circuits. *Electronics*, 38(8), April 1965.

- [Murdoch81] JOSEPH B. MURDOCH. Inverse Square Law Approximation of Illuminance. *Journal of the Illuminating Engineering Society*, 10(2):96–106, 1981.
- [Nicodemus77] F. E. NICODEMUS, J. C. RICHMOND, J. J. HSIA, I. W. GINSBERG, AND T. LIMPERIS. Geometrical Considerations and Nomenclature for Reflectance. National Bureau of Standards, 1977.
- [Okano98] YUKIO OKANO. Influence of Image Enhancement Processing on SFR of Digital Cameras. In *Proceedings of IS&T's PICS Conference 1998*, pages 74–78, 1998.
- [Paakkari93] JUSSI PAAKKARI AND ILKKA MORING. Method for evaluating the performance of range imaging devices. In *Industrial Applications of Optical Inspection, Metrology and Sensing, Proceedings of SPIE*, volume 1821, pages 350–356, 1993.
- [Pai99] DINESH K. PAI, JOCHEN LANG, JOHN .E. LLOYD, AND ROBERT J. WOODHAM. ACME, A Telerobotic Active Measurement Facility. In P. Corke and J. Trevelyan, editors, *6th International Symposium on Experimental Robotics*, 250, pages 391–400. Springer, 1999.
- [Pai01] DINESH K. PAI, KEES VAN DEN DOEL, DOUG L. JAMES, JOCHEN LANG, JOHN E. LLOYD, JOSHUA L. RICHMOND, AND SOM H. YAU. Scanning Physical Interaction Behavior of 3D Objects. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 87–96. ACM Press, 2001.
- [Patow03] GUSTAVO PATOW AND XAVIER PUEYO. A Survey of Inverse Rendering Problems. *Computer Graphics Forum*, 22(4):663–688, 2003.
- [Perlin85] KEN PERLIN. An Image Synthesizer. In *Proceedings of SIGGRAPH 1985*, pages 287–296. ACM Press, 1985.
- [Perona90] P. PERONA AND J. MALIK. Scale-Space and Edge Detection using Anisotropic Diffusion. *IEEE PAMI*, 12:629–639, 1990.
- [Pharr00] MATT PHARR AND PAT HANRAHAN. Monte Carlo Evaluation of Non-linear Scattering Equations for Subsurface Reflection. In *SIGGRAPH 2000*, pages 75–84, 2000.

- [Phong75] BUI TUONG PHONG. Illumination for Computer Generated Pictures. *Commun. ACM*, 18(6):311–317, 1975.
- [Pirzadeh99] HORMOZ PIRZADEH. Computational Geometry with the Rotating Calipers. Master’s thesis, School of Computer Science, McGill University, Montreal, Quebec, Canada, November 1999.
- [Pitman92] J. PITMAN. *Probability*. Springer, 1992.
- [Press94] WILLIAM H. PRESS, SAUL A. TEUKOLSKY, WILLIAM T. VETTERLING, AND BRIAN P. FLANNERY. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge Univ. Press, 2nd edition, 1994.
- [Ramamoorthi01] RAVI RAMAMOORTHI AND PAT HANRAHAN. A Signal-Processing Framework for Inverse Rendering. In *Proceedings of SIGGRAPH 2001*, pages 117–128, 2001.
- [Reichenbach91] STEPHEN E. REICHENBACH, STEPHEN K. PARK, AND RAMKUMAR NARAYANSWAMY. Characterizing Digital Image Acquisition Devices. *Optical Engineering*, 30(2):170–177, February 1991.
- [Robertson99] MARK A. ROBERTSON, SEAN BORMAN, AND ROBERT L. STEVENSON. Dynamic Range Improvement Through Multiple Exposures. In *Proceedings of ICIP 1999*, 1999.
- [Robertson03] MARK A. ROBERTSON, SEAN BORMAN, AND ROBERT L. STEVENSON. Estimation-Theoretic Approach to Dynamic Range Enhancement using Multiple Exposures. *Journal of Electronic Imaging*, 12(2):219–285, April 2003.
- [Rubin80] STEVEN M. RUBIN AND TURNER WHITTED. A 3-Dimensional Representation for Fast Rendering of Complex Scenes. In *Proceedings of SIGGRAPH 1980*, pages 110–116. ACM Press, 1980.
- [Rushmeier90] HOLLY E. RUSHMEIER AND KENNETH E. TORRANCE. Extending the Radiosity Method to Include Specularly Reflecting and Translucent Materials. *ACM Trans. Graph.*, 9(1):1–27, 1990.

- [Rykowski97] RONALD F. RYKOWSKI AND C. BENJAMIN WOOLEY. Source Modeling for Illumination Design. In *Lens Design, Illumination, and Optomechanical Modeling*, volume 3130, pages 204–208. SPIE, 1997.
- [Saguer04] MARTÍ MARIA SAGUER. *Little CMS Engine – How to use the Engine in Your Application*, 2004. Available from <http://www.littlecms.com>.
- [Sapiro01] GUILLERMO SAPIRO. *Geometric Partial Differential Equations and Image Analysis*. Cambridge University Press, 2001.
- [Sato97] Y. SATO, M. WHEELER, AND K. IKEUCHI. Object Shape and Reflectance Modeling from Observation. In *Proc. SIGGRAPH*, pages 379–388, August 1997.
- [Schlick94] CHRISTOPHE SCHLICK. An Inexpensive BRDF Model for Physically-based Rendering. *Computer Graphics Forum*, 13(3):233–246, 1994.
- [Schoeneman93] CHRIS SCHOENEMAN, JULIE DORSEY, BRIAN SMITS, JAMES ARVO, AND DONALD GREENBURG. Painting with Light. In *Proceedings of SIGGRAPH 1993*, pages 143–146, 1993.
- [Scopigno02] ROBERTO SCOPIGNO, CARLOS ANDUJAR, MICHAEL GOESELE, AND HENDRIK P. A. LENSCH. 3D Data Acquisition. In *Tutorial Notes, Eurographics Conference 2002*, 2002.
- [Segal92] MARK SEGAL, CARL KOROBKIN, ROLF VAN WIDENFELT, JIM FORAN, AND PAUL HAEBERLI. Fast shadows and lighting effects using texture mapping. In *Proceedings of the SIGGRAPH 92 annual conference*, pages 249–252. ACM Press, 1992.
- [Segal98] MARK SEGAL AND KURT AKELEY. *The OpenGL Graphics System: A Specification (Version 1.2)*, 1998.
- [Siegel96] M. W. SIEGEL AND R. D. STOCK. A General Near-Zone Light Source Model and its Application to Computer Automated Reflector Design. *SPIE Optical Engineering*, 35(9):2661–2679, September 1996.

- [Sillion95] FRAN COIS X. SILLION. A Unified Hierarchical Algorithm for Global Illumination with Scattering Volumes and Object Clusters. *IEEE Trans. Visualization and Computer Graphics*, 1(3):240–254, 1995.
- [Sloan03] PETER-PIKE SLOAN, JESSE HALL, JOHN HART, AND JOHN SNYDER. Clustered Principal Components for Precomputed Radiance Transfer. *ACM Trans. Graph.*, 22(3):382–391, 2003.
- [Stam95] JOS STAM. Multiple Scattering as a Diffusion Process. In *Rendering Workshop 1995*, pages 51–58, 1995.
- [Stam01] JOS STAM. An Illumination Model for a Skin Layer Bounded by Rough Surfaces. In *Rendering Workshop 2001*, pages 39–52, 2001.
- [Taubin95] G. TAUBIN. A Signal Processing Approach to Fair Surface Design. *Proceedings of SIGGRAPH 1995*, pages 351–358, 1995.
- [Taylor94] BARRY N. TAYLOR AND CHRIS E. KUYATT. Guidelines for Evaluation and Expressing the Uncertainty of NIST Measurement Results. Technical Report NIST 1297, National Institute of Standards and Technology (NIST), 1994.
- [Tominaga01] S. TOMINAGA, T. MATSUMOTO, AND N. TANAKA. 3D Recording and Rendering for Art Paintings. In *Proceedings of the Ninth Color Imaging Conference: Color Science and Engineering: Systems, Technologies, Applications*, pages 337–341. IS&T, 2001.
- [Toussaint83] GODFRIED T. TOUSSAINT. Solving Geometric Problems with the Rotating Calipers. In *Proceedings of IEEE MELECON'83, Athens, Greece, May 1983*.
- [Tsai87] R.Y. TSAI. A Versatile Camera Calibration Technique for High-Accuracy 3D Machine Vision Metrology Using Off-the-Shelf TV Cameras and Lenses. *RA*, 3(4):323–344, 1987.
- [Veach94] ERIC VEACH AND LEONIDAS J. GUIBAS. Bidirectional Estimators for Light Transport. In *Eurographics Rendering Workshop 1999*, pages 147–162, June 1994.

- [Verbeck84] C. P. VERBECK AND D. P. GREENBERG. A Comprehensive Light Source Description for Computer Graphics. *IEEE Computer Graphics & Applications*, 4(7):66–75, July 1984.
- [Walter97] BRUCE WALTER, PHILIP M. HUBBARD, PETER SHIRLEY, AND DONALD P. GREENBERG. Global Illumination Using Local Linear Density Estimation. *ACM Transactions on Graphics*, 16(3):217–259, July 1997.
- [Ward92] G. WARD. Measuring and modeling anisotropic reflection. In *Proceedings of SIGGRAPH 1992*, pages 265–272, July 1992.
- [Williams99] TOM L. WILLIAMS. *The Optical Transfer Function of Imaging Systems*. Institute of Physics Publishing, London, 1999.
- [Williams01] DON WILLIAMS AND PETER D. BURNS. Diagnostics for Digital Capture using MTF. In *Proceedings of IS&T's PICS Conference 2001*, pages 227–232, 2001.
- [Wood00] DANIEL N. WOOD, DANIEL I. AZUMA, KEN ALDINGER, BRIAN CURLESS, TOM DUCHAMP, DAVID H. SALESIN, AND WERNER STUETZLE. Surface Light Fields for 3D Photography. In *SIGGRAPH 2000*, pages 287–296, 2000.
- [Woodham81] ROBERT J. WOODHAM. Analysing Images of Curved Surfaces. *Artificial Intelligence*, 17:117–140, August 1981.
- [Wyckoff63] CHARLES W. WYCKOFF AND STAN A. FEIGENBAUM. An Experimental Extended Exposure Response Film. *SPIE*, 1:117–125, 1963.
- [Yu99] YIZHOU YU, PAUL DEBEVEC, JITENDRA MALI, AND TIM HAWKINS. Inverse Global Illumination: Recovering Reflectance Models of Real Scenes From Photographs. *Proceedings of SIGGRAPH 99*, pages 215–224, August 1999.
- [Zatz93] HAROLD R. ZATZ. Galerkin Radiosity: A Higher Order Solution Method for Global Illumination. In *Proceedings of SIGGRAPH 93*, pages 213–220, 1993.
- [Zhang99] ZHENGYOU ZHANG. A Flexible New Technique for Camera Calibration. Technical Report MSR-TR-98-71, Microsoft Research, 1999. Updated version of March 25, 1999.

- [Zongker99] DOUGLAS E. ZONGKER, DAWN M. WERNER, BRIAN CURLESS, AND DAVID H. SALESIN. Environment Matting and Compositing. In *SIGGRAPH 1999*, pages 205–214, 1999.

Curriculum Vitae – Lebenslauf

Curriculum Vitae

1972 born in Heidenheim an der Brenz, Germany
1979 – 1983 Hillerschule (Primary School), Steinheim am Albuch
1983 – 1992 Hellenstein-Gymnasium (High School), Heidenheim
1992 – 1993 Civil Service, Lebenshilfe Heidenheim
1993 – 1999 Study of Computer Science, University of Ulm, Germany
1997 – 1998 Exchange Student, University of North Carolina
 at Chapel Hill, Chapel Hill, USA
1999 Diploma (Diplom-Informatiker, Dipl.-Inf.)
since 1999 Ph.D. Student at the Max-Planck-Institut für Informatik
 (MPI Informatik), Saarbrücken, Germany

Lebenslauf

1972 geboren in Heidenheim an der Brenz
1979 – 1983 Hillerschule Steinheim am Albuch
1983 – 1992 Hellenstein-Gymnasium, Heidenheim
1992 – 1993 Zivildienst, Lebenshilfe Heidenheim
1993 – 1999 Informatikstudium, Universität Ulm
1997 – 1998 Studium als Austauschstudent, University of North Carolina
 at Chapel Hill, Chapel Hill, USA
1999 Abschluß als Diplom-Informatiker (Dipl.-Inf.)
seit 1999 Promotion am Max-Planck-Institut für Informatik, Saarbrücken