
Advanced Methods for
Relightable Scene Representations
in Image Space

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

vorgelegt von

Martin Fuchs

MPI Informatik
Campus E1 4
66123 Saarbrücken

am 04. November 2008 in Saarbrücken.

Betreuende Wissenschaftler

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

Prof. Dr. Volker Blanz
Universität Siegen

Prof. Dr. Ramesh Raskar
MIT Media Lab, Cambridge, MA, USA

Dr.-Ing. Hendrik P. A. Lensch
MPI Informatik, Saarbrücken

Gutachter

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

Dr. Hendrik P. A. Lensch
MPI Informatik, Saarbrücken

Prof. Dr. Szymon Rusinkiewicz
Princeton University, Princeton, NJ, USA

Dekan

Prof. Dr. Joachim Weickert
Universität des Saarlandes, Saarbrücken

Mitglieder des Prüfungsausschusses

Prof. Dr. Philipp Slusallek
Universität des Saarlandes, Saarbrücken
Vorsitzender des Prüfungsausschusses

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

Dr. Hendrik P. A. Lensch
MPI Informatik, Saarbrücken

Dr. Elmar Eisemann
Universität des Saarlandes, Saarbrücken

Tag des Kolloquiums

15. Dezember 2008

Abstract

The realistic reproduction of visual appearance of real-world objects requires accurate computer graphics models that describe the optical interaction of a scene with its surroundings. Data-driven approaches that model the scene globally as a reflectance field function in eight parameters deliver high quality and work for most material combinations, but are costly to acquire and store. Image-space relighting, which constrains the application to create photos with a virtual, fixed camera in freely chosen illumination, requires only a 4D data structure to provide full fidelity.

This thesis contributes to image-space relighting on four accounts: (1) We investigate the **acquisition** of 4D reflectance fields in the context of sampling and propose a practical setup for pre-filtering of reflectance data during recording, and apply it in an adaptive sampling scheme. (2) We introduce a feature-driven image synthesis algorithm for the **interpolation** of coarsely sampled reflectance data in software to achieve highly realistic images. (3) We propose an **implicit** reflectance data representation, which uses a Bayesian approach to relight complex scenes from the example of much simpler reference objects. (4) Finally, we construct novel, passive devices out of optical components that **render** reflectance field data in real-time, shaping the incident illumination into the desired image.

Kurzzusammenfassung

Die realistische Wiedergabe der visuellen Erscheinung einer realen Szene setzt genaue Modelle aus der Computergraphik für die Interaktion der Szene mit ihrer Umgebung voraus. Globale Ansätze, die das Verhalten der Szene insgesamt als Reflektanzfeldfunktion in acht Parametern modellieren, liefern hohe Qualität für viele Materialtypen, sind aber teuer aufzuzeichnen und zu speichern. Verfahren zur Neubeleuchtung im Bildraum schränken die Anwendbarkeit auf fest gewählte Kameras ein, ermöglichen aber die freie Wahl der Beleuchtung, und erfordern dadurch lediglich eine 4D - Datenstruktur für volle Wiedergabetreue.

Diese Arbeit enthält vier Beiträge zu diesem Thema: (1) wir untersuchen die **Aufzeichnung** von 4D Reflektanzfeldern im Kontext der Abtasttheorie und schlagen einen praktischen Aufbau vor, der Reflektanzdaten bereits während der Messung vorfiltert. Wir verwenden ihn in einem adaptiven Abtastschema. (2) Wir führen einen merkmalgesteuerten Bildsynthesearchiv für die **Interpolation** von grob abgetasteten Reflektanzdaten ein. (3) Wir schlagen eine **implizite Beschreibung** von Reflektanzdaten vor, die mit einem Bayesschen Ansatz komplexe Szenen anhand des Beispiels eines viel einfacheren Referenzobjektes neu beleuchtet. (4) Unter der Verwendung optischer Komponenten schaffen wir passive Aufbauten zur **Darstellung** von Reflektanzfeldern in Echtzeit, indem wir einfallende Beleuchtung direkt in das gewünschte Bild umwandeln.

Acknowledgement

Just like any larger work, a dissertation is never the product of a single person in isolation; it is always shaped by the environment in which it comes into being. Accordingly, I would like to thank my advisors, colleagues, friends and family alike; and while I am grateful to many more supporters than I can list here, some of them stand out so much that I would like to thank them by name.

I thank my advisors and coauthors, Hans-Peter Seidel, Hendrik Lensch, Volker Blanz and Ramesh Raskar for their ongoing support throughout the different phases of my research; I thank the people shouldering the administrative chores at both the MPI Informatik and Mitsubishi Electric Research Labs, especially our administrative assistant, Sabine Budde, whose effort was a strong shield against the hindrances bureaucracy so often entails. Likewise, I thank the students and employees of our Information Services and Technology department, who constantly fight to get our computer infrastructure to fulfill our permanently growing needs.

Credit is due to the technicians who supported us in constructing our setups and building our prototypes, namely Uwe Meiser, Axel Köppel and Michael Laise at the MPI, and John Barnwell at Mitsubishi Electric Research Labs; further I thank Douglas Lanman, who, during the Christmas Holidays 2007, kept feeding the 3D printer out of our reach so that it could produce the components we needed for the passive display project just in time, and Boris Ajdin, inventor of the 5 o'clock AM tea, who so patiently assembled the display prototypes, tested them, took them apart and assembled them again.

I thank Matthias Hullin and Thomas Annen for feedback and proof-reading in the production of this thesis, and Kristina Scherbaum for some of the illustrations and generally good advice on the typesetting process.

Further, I owe thanks to the patient models we recorded for the implicit relighting project, Bettina Stiller-Weishaupt, Katrin Dedden, and Conny Liegl, who also lent us her voice talent for video footage, helped perform some experiments and so often provided support in finding the single available place to stay at in cities booked out long ago.

While conducting the research for this thesis, I also served as an administrator of the computer infrastructure in the Computer Graphics Department at the MPI. My special thanks therefore go to Johannes Günther, Martin Sunkel and Carsten Stoll, who shared and helped balance this existing, but at times also burdensome duty. Carsten also deserves special credit for sharing his expertise on geometrical modeling, and for, instead of heading into a well-deserved vacation after submitting work of his own, volunteering to create an enthralling video from our results of the interpolating image synthesis project, hours before the deadline.

Just as I wish to thank the anonymous reviewers of our publications, who pushed the quality and provided valuable insight, I also wish to thank the many authors of the many software programs, free not only in cost but also in spirit, that I used in my research and in the writing of this thesis, especially Rafał Mantiuk and Grzegorz Krawczyk for the

PFS toolkit that processed the HDR images we used, and Ken Clarkson, who provided the program `hull` which we used for the triangulation computation that is part of our software rendering.

Finally, I would like to thank my family and close friends – for bearing with me and supporting me during the ups and downs a project so long as this one entails, especially my parents Karin and Josef and my brother Andreas, Annemarie and Fritz Rupertus, Dieter Klersy and my aunt Inge Klersy, whose refreshing criticism I fondly remember.

The work presented in this thesis was partially funded by the Max Planck Center for Visual Computing and Communication (BMBF-FKZ01IMC01) and by the DFG Emmy Noether fellowship (Le 1341 / 1-1). The work presented in chapter 8 received additional funding by Mitsubishi Electric Research Laboratories, Cambridge, MA, USA.

Contents

1	Introduction	1
2	The Principles of Imaging	7
2.1	Naming Conventions	7
2.2	Radiometry	7
2.2.1	Physical Quantities	7
2.2.2	The Plenoptic Function	9
2.3	Imaging	9
2.3.1	Light Fields	9
2.3.2	Digital Images	10
2.4	Modeling Appearance	11
2.4.1	Local Appearance Descriptors	12
	Properties of Scattering Materials	12
	The Rendering Equation	13
2.4.2	Global Appearance Modeling	13
	Reflectance Fields	13
	Image-Space Relighting	15
3	Related Work	17
3.1	Overview	17
3.2	Fixed-Illumination Models	17
3.2.1	Light Field Techniques	17
3.2.2	The Lumigraph	18
3.2.3	Surface Light Fields	19
3.2.4	Environment Maps	19
3.2.5	Image-Based Rendering	19
3.3	Local Light Transport Descriptors for Relightable Scenes	20
3.3.1	Evaluating the Rendering Equation	20
	Rendering by Drawing Primitives	20
	Ray Tracing	21
	Global Illumination Methods	22
	Pre-computed Radiance Transfer	22
3.3.2	BRDF Methods	23
	Modeling the BRDF	23
	Editing the BRDF	24
	Measuring the BRDF	24

3.3.3	SVBRDF Extensions	25
	Measuring the SVBRDF	25
	Editing the SVBRDF	26
3.3.4	Full BSSRDF techniques	26
3.4	Global Light Transport Descriptors with Relighting Capability	27
3.4.1	Models in World-Space	27
	4 → 4 Reflectance Fields	27
	2 → 4 Reflectance Fields	27
	4 → 2 Reflectance Fields	27
3.4.2	Image-Space Relighting with 2 → 2 Reflectance Fields	28
	Concepts	28
	Measuring	29
	Rendering	30
	Applications	31
	Image-space Reflectance Fields as Input Data	31
3.5	Hybrid Approaches	32
3.5.1	Surface Reflectance Fields	32
3.5.2	BTF Methods	32
	Measuring the BTF	33
	Compressing the BTF	33
	Synthesizing and Editing the BTF	33
3.6	Discussion	33
4	A Measurement Setup with Programmable Light Source Extent	35
4.1	Introduction	35
4.2	Construction	36
4.3	Calibration	37
4.4	Rendering Scheme	39
5	Sampling Reflectance Fields	43
5.1	Introduction	43
	5.1.1 Related Work	45
5.2	Fixed Pattern Sampling	46
	5.2.1 Results of Fixed Pattern Sampling	46
5.3	Adaptive Sampling	47
	5.3.1 Algorithm	48
	5.3.2 Results of the Adaptive Sampling Algorithm	49
	Quantitative Evaluation	50
	Qualitative Evaluation	53
	Timings / Performance	58
5.4	Two-Level Reflectance Fields	58
5.5	Discussion	59
	5.5.1 Evaluation using sampling theory	59
	5.5.2 Limitations of this technique	64

5.6	Conclusion	65
6	Feature-Guided Image Synthesis for Reflectance Field Interpolation	67
6.1	Introduction	67
6.1.1	Related Work	68
6.1.2	Overview	69
6.2	Extracting Image Semantics	71
6.2.1	Specular Highlights	71
6.2.2	Shadow Boundaries	72
6.3	Image Feature Upsampling	73
6.3.1	Upsampling Specular Highlights	74
6.3.2	Upsampling shadow data	74
6.4	Image Synthesis	74
6.4.1	Reconstruction Based on Shadow Maps	75
6.4.2	Regularization Using Image-Based Priors	77
6.5	Results	78
6.5.1	Limitations	82
6.6	Conclusion	82
7	Implicit Relighting with a Bayesian Approach	85
7.1	Introduction	85
7.2	Implicit Relighting	87
7.3	Bayesian Relighting	88
7.4	Experiments	91
7.5	Results	94
7.6	Conclusion	96
8	Towards Passive 6D Reflectance Field Displays	99
8.1	Introduction	99
8.1.1	Related Concepts	101
8.1.2	Overview	102
8.2	Observer-Invariant displays	104
8.2.1	Design I	104
8.2.2	Results for Design I	105
8.2.3	Design II with Correcting Lenses	107
8.2.4	Results for Design II	109
8.3	Observer-Variant Displays	110
8.3.1	Design III	110
8.3.2	Results for Design III	113
8.3.3	Design IV	113
8.3.4	Results for Design IV	113
8.4	Implementation details	115
8.4.1	Pattern/Lens Registration	116
8.5	Discussion	116

8.5.1	Limitations	116
8.5.2	Future Challenges	117
8.5.3	Applications	117
8.6	Conclusion	118
9	Concluding Remarks	119
9.1	Applications	119
9.2	Future Challenges	120
9.3	Summary	121
	Bibliography	123
	Index	147

List of Figures

1.1	Image-space relighting example	1
1.2	Reflectance field example	3
2.1	Appearance models taxonomy	12
4.1	Indirect illumination measurement setup	37
4.2	Achievable light distributions	38
5.1	Point light relighting results	44
5.2	Regular sampling results	45
5.3	Sequence of generated sample positions.	50
5.4	Spatial error plots	51
5.5	Global error drop	52
5.6	Adaptive sampling quality comparison	54
5.7	Adaptive sampling quality comparisons (detail)	55
5.8	Adaptive sampling failure cases	56
5.9	Adaptive sampling in high frequency illuminations	57
5.10	Adaptive sampling: combination of scales	60
5.11	Adaptive sampling: real world comparison	61
5.12	Light source power spectra	62
5.13	Distribution of angular distances	63
6.1	Improvements through interpolation	67
6.2	Recorded and interpolated light source positions	70
6.3	Interpolation processing pipeline	71
6.4	Highlight separation	72
6.5	Shadow separation	73
6.6	Illustration of the reconstruction process	75
6.7	Polar plot of blurred shadow maps	76
6.8	Upsampling results	79
6.9	Complex scene upsampling results	80
7.1	Dinner scene after implicit relighting	86
7.2	Concept of implicit relighting	88
7.3	Implicit relighting: influence of regularization	90
7.4	Implicit relighting setup	92
7.5	Implicit relighting: comparison of HDR and LDR computations	92

7.6	Light probe generalization	93
7.7	Relighting a human face	95
7.8	Implicit relighting results	95
7.9	Implicit relighting: comparison of extended and point light sources	97
8.1	4D reflectance field display demonstration	100
8.2	Display dimensionalities	101
8.3	Flat reflectance field embedding	103
8.4	Lens arrays in light and reflectance field displays	103
8.5	Design I ray diagram	105
8.6	Design I construction	106
8.7	Transmissivity and contrast experiment	106
8.8	Design II ray diagram	108
8.9	Design II construction	108
8.10	Design II results	109
8.11	Design III ray diagram	111
8.12	Design III construction	112
8.13	Design III result	112
8.14	Design IV results	114
8.15	Design IV pattern design	115

1 Introduction



Figure 1.1: Image-space relighting virtually sets the illumination in a picture. It delivers renderings for extended light sources and directional lights as well as complex illuminations from real world environments (bottom right).

Mathematical models for the visual appearance of the real world are an important topic in the field of **computer graphics**. These models enable us to create – or **render** – life-like imagery from abstract representations. For achieving this, a large body of research has come into existence in an astonishing development which began in the second half of the twentieth century.

Today, synthetic imagery has become an integral part of most, if not all, commercial movie productions, and creates increasingly realistic virtual worlds in computer entertainment applications. As the bar on what may be considered photo-realism continues to rise with every production which is released, seamless integration of the real world into the realm of imagination becomes ever more important. Increasing demands on realism, though, go hand-in-hand with increasing model complexity, which makes manual model creation harder and harder, and calls for automated solutions.

One of the most exciting challenges in computer graphics that provides such solutions is the **acquisition of real-world appearance**. Given a real-world scene, appearance

acquisition addresses the problem of recording its optical interactions with its surroundings, expressing these in a method-specific model of light transport. This model can then be applied to predict the scene's appearance in novel conditions, such as for novel viewpoints or illuminations, which usually are not accessible at acquisition time. In this context, one is not only facing the intricacy of the accurate measurement of the scene, but also suitable representations and efficient methods for subsequent rendering.

This thesis concentrates on change of appearance with illumination (see Figure 1.1); specifically, given a real-world scene, to derive a mathematical model of it that permits to **relight** it, that is, to render the scene in novel illumination.

A multitude of possible paths towards solutions have been explored in the past. In general, these methods can be categorized with respect to their generality, expressiveness and robustness. Here, "*generality*" refers to how many assumptions they make, i. e. for which kinds of scenes they work well and how strictly the measurement setup needs to be defined. "*Expressiveness*" states a method's range in predicting interactions with the recorded scene – such as re-illumination or novel views, especially in so far as they were not part of the initial acquisition conditions. Finally, the measure of "*robustness*" looks at how well a method performs if the assumptions no longer apply and its sensitivity to imprecise acquisition conditions. In addition, the compactness in representation and computational expense has a profound impact on the applicability of any given approach.

Of course, these objectives are in conflict and usually cannot be satisfied simultaneously. Chapter 3 will discuss a selection of the approaches known today, along with their trade-offs. However, we can already state at this point that **image-space relighting** methods stand out in that they are most radically geared towards generality, often requiring the input scene only to be static over the time needed to measure it, and making only few assumptions on the color formation process. In turn, they sacrifice the ability to generate novel views – usually a camera with the exact same optical properties as the optical system used for measuring is used for rendering as well – but maintaining a high expressiveness for novel light conditions, commonly restricted only to the assumption of novel illumination to be distant with respect to the scene. In this case, the scene can be expressed by a four-dimensional reflectance field, as shown in Figure 1.2.

The first natural extensions of such general methods towards local illumination (Masselus et al., 2003) and even, at the same time, novel viewpoints (Garg et al., 2006), have meanwhile become available, but usually they are still most expensive when applied to achieve the high sampling density and range of views that are prerequisites for full expressiveness and accuracy in rendering. Therefore, the work presented within this thesis is still bound by the limits of the customary assumptions, i. e. a fixed observer and distant illumination, while pushing the envelope on ease of use, accuracy of rendering, generality of scenes and quality of the generated results, culminating in the introduction of a novel display device that renders reflectance fields with optical means only.

In the subsequent paragraphs, we will concisely list the contributions of this thesis, which will follow this introduction after recalling the basic concepts of imaging (Chapter 2) and an in-depth discussion of related work (Chapter 3).

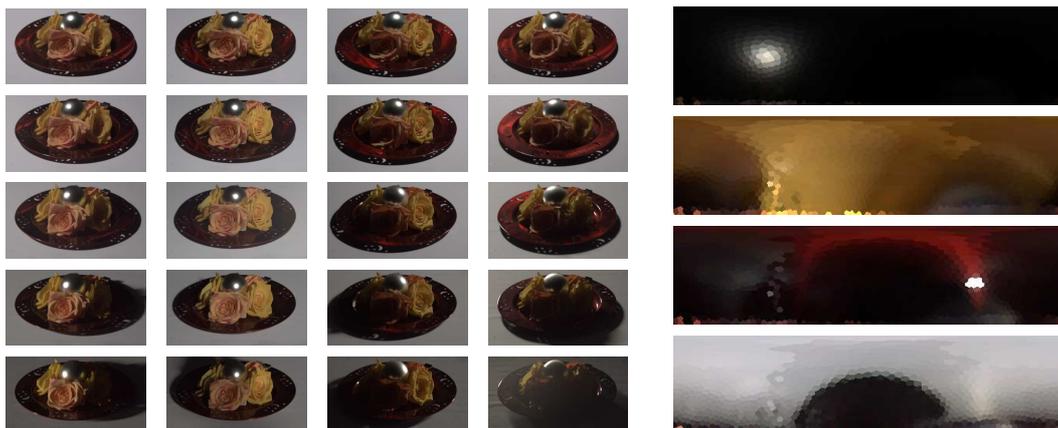


Figure 1.2: 4D reflectance fields can be interpreted as a collection of images for basis illuminations (left), or fixing pixel positions, as illumination-dependent appearance at these pixels (right); from top to bottom: pixel on the specular ball, on a flower, the metal plate, and a pixel on the table, partially in shadow.

Acquisition

The acquisition procedure, which establishes the scene-dependent relationship between the scene appearance and the light that illuminates it, is crucially important for the applicability and quality of any image-space relighting method, as it defines the information available for analysis and rendering. The previous work on this topic can be understood as adhering to one of three distinct approaches, with typical references for each:

- *Passive recordings in variant illuminations* (Matusik et al., 2004): illumination during the acquisition step varies naturally, outside of the influence of the method. From the combined observations of scene and illuminations, a relightable model of the scene can be deduced.

This is highly general with respect to acquisition conditions, permitting even the relighting of an entire city panorama, but it is limited in rendering quality due to the strong assumptions that need to be taken in the data representation.

- *Recordings in controlled, complex basis illumination conditions* (Peers and Dutré, 2003, 2005): typical methods for this approach create extended illumination in precisely controlled patterns, thus generating a compressed model of the light transport already during the acquisition step.

While providing a high data storage efficiency already during the measurement process, it has been only demonstrated for small subareas of the sphere of possible incident illumination directions, and the high sensitivity to inevitable noise strongly suggest it to be impractical to scale, precluding many relevant applications.

- *Recording in controlled, directional illumination* (Debevec et al., 2000): this approach, which employs a moving directional light source, typically implemented as a point light in a reasonable distance, is longest in use and provides the most direct observation of the **reflectance field**, which is the underlying data structure in most image-space relighting techniques.

It immediately delivers high contrast and maximal signal to noise ratios among all three approaches, however at the cost of either restricting the input scene to be low-frequency in its response to illumination changes, or requiring exceptionally large input data sets, with the number of observations in the hundreds or even thousands.

Our work aims at the improvement of the latter method: in analyzing it in the context of the sampling theorem (Chapter 5), we will see that we can modify the angular extent of the incident illumination in a way which effectively pre-filters the reflectance field before it is acquired, avoiding the artefacts introduced by undersampling without increasing the measurement expense. We will in detail describe a practical setup to achieve this (Chapter 4), including calibration and other implementation issues, evaluate it with respect to its sampling properties and demonstrate its usefulness.

With the modification of the light source shape we have achieved an improvement on the measurement expense vs. quality tradeoff from the quality side, avoiding aliasing artefacts by sampling and reconstructing only up the limit that we could actually observe.

However, the high number of input images required for faithful reconstructions are only correct for a worst case scenario, where dense sampling of illumination directions is required for the entire domain of incident light directions. As we will see, there are scenes which require this high a density only in very small areas; we will therefore introduce an adaptive algorithm, which during acquisition automatically determines where to take additional samples. Therefore, we can again improve on the expense vs. quality tradeoff, but this time maintaining a high quality while reducing the required measurement effort. Using ground truth measurements with 10 000 input images as reference, we can in detail demonstrate the improvement.

The work in chapters 4 and 5 has been published in the ACM Transactions on Graphics (Fuchs et al., 2007b).

Interpolation by Guided Image Synthesis

There are classes of scenes where the use of adaptive sampling is limited, because all light directions are almost equally important for a faithful reconstruction – for instance a mirror sphere. There are also application cases, such as recording of motion sequences or fast changing scenes (Wenger et al., 2005), where the measurement setup will not permit to move the light source in time for an adaptive acquisition to take place.

Still, it would be desirable to stay clear of undersampling artefacts and blur. In Chapter 6, we will therefore introduce a software solution to that problem. By analyzing the captured image data, we can extract separate information on shadows and highlights from a scene, and interpolate and upsample them separately. Then, we can use the new data for a guided image synthesis, inserting in-between images into the original data set until a plausible reconstruction becomes possible.

We have published this method in the Eurographics 2007 special issue of Computer Graphics Forum (Fuchs et al., 2007c).

Implicit Relighting

Looking at the entire image space relighting problem, it reduces to creating the impression of a scene in a novel illumination condition; but instead of explicitly modeling the light distributions and the scattering behavior of the input scene, we can learn its behavior relative to a simple reference object. Then, we can take a picture of the reference object in a novel light condition in which we have never observed the whole scene, and faithfully reconstruct the scene in that light. By expressing the light transport with the implicit scheme that we present in Chapter 7, we can thus solve the relighting problem without recording a reflectance field, and greatly relax the requirements on the measurement efforts while adding only a slight computational burden. A Bayesian maximum a-posteriori approach provides control over noise in the measurement and the reconstruction.

We have published this method at the Eurographics Symposium on Rendering (Fuchs et al., 2005b).

Rendering with Optical Means

Up to now, we have investigated the front and the middle of the image pipeline, the acquisition and modeling steps. We will now turn our attention to its end, concerning the output steps creating the result images. Traditionally, the algorithm rendering the images has been a direct consequence of the internal data representation, resulting in the evaluation of an integral over illumination and reflectance data in software. Specifically, Nayar et al. (2004) have introduced an integrated display system, which measures the real-world incident illumination on their display, and then renders a scene in that light, thus creating the experience of a virtualized object.

As we will see in Chapter 8, we can take this idea further and actually transform the incident illumination itself into the desired image, without in-between steps of measurement, analysis and software rendering. Instead, we can arrange a set of optical elements, so that the incident illumination is modulated by a pattern encoding of a reflectance field

and the individual images of respective illumination directions are superimposed to form the end result.

We will further show a proof of concept optical setup that generates different images for different observer positions, which in the limit forms a light dependent light field display, providing a 6D display experience.

This work has been published in the SIGGRAPH 2008 special issue of the ACM Transactions on Graphics (Fuchs et al., 2008).

2 The Principles of Imaging

Before we explain our contributions in more detail, we let us address the principles of imaging and introduce the necessary terminology. We will discuss the fundamentals of computer graphics, especially rendering, and establish the necessary relationship between physical descriptions of the real world and computational models for their implementation.

2.1 Naming Conventions

In the following chapters, we use the following naming conventions for variables in order to describe geometric relationships:

(θ, φ)	a direction in space, expressed in a polar coordinate system so that the pole identified as “up” is at $\theta = 0$. In the context of surfaces, “up” is aligned with the surface normal.
ω	a solid angle patch
(x_w, y_w, z_w)	a location in space, in a Cartesian world coordinate system
(x, y)	a position on the image plane, in image space
(u, v)	a position on a reference plane in space, or on a 2D parameterizable surface
A	an area patch
\square_i	a variable associated with an incident illumination
\square_o	a variable associated with outgoing / exitant / radiant illumination

2.2 Radiometry

2.2.1 Physical Quantities

Light transport can be measured with physical quantities; this process is referred to as **radiometry**. An alternative description of light transport is provided by **photometry** which models it according to the perception by the human eye. As we will not intensively discuss perceptual issues, we will for the sake of simple computation remain with radiometric terminology. A short, but concise summary of radiometric and photometric

quantities with additional references is provided by Palmer (2008); an in-depth discussion of light transport is available from Hecht (2002).

Radiometry describes the propagation of **energy**

$$Q, \quad [Q] = 1 \text{ J} \quad (2.1)$$

through space by the electromagnetic **radiation** process. This energy is carried in discrete form by **photons**; the amount of energy each of these elementary particles carries is characterized by an intrinsic property called the **wavelength**.

To describe the time-dependence of light transport, we can observe the **radiant power**

$$P = \frac{dQ}{dt}, \quad [P] = 1 \frac{\text{J}}{\text{s}} = 1 \text{ W} \quad (2.2)$$

of the light which is emitted or received. The radiant power is also called **radiant flux** with symbol Φ , which we will use to avoid confusion with the definition of the plenoptic function below in Equation 2.7.

In order to describe the spatial or angular distributions of radiation, we need to define differential quantities that express the radiation at an infinitesimal point in space or an infinitesimal angle direction, rather than a finite volume or a finite solid angle. For instance, a light source may emit radiant power that changes with an emission direction (θ, φ) . The radiation emitted into an infinitesimal solid angle $d\omega$ is described with the **radiant intensity**

$$I = \frac{d\Phi}{d\omega}, \quad [I] = 1 \frac{\text{W}}{\text{steradians}} = 1 \frac{\text{W}}{\text{sr}} \quad (2.3)$$

Considering the flux through a surface element of infinitesimal area A from the entire hemisphere of incident directions, we can measure **irradiance**

$$E = \frac{d\Phi}{dA}, \quad [E] = \frac{\text{W}}{\text{m}^2} \quad (2.4)$$

as the flux arriving at the surface and also define the corresponding **radiant exitance**, or **radiosity**

$$B = \frac{d\Phi}{dA}, \quad [B] = \frac{\text{W}}{\text{m}^2} \quad (2.5)$$

as the flux emitted from the surface. Either may be referred to as **flux density**.

Distinguishing light transported to or from a surface from a specific direction under an angle θ against the surface normal, we can finally define **radiance**

$$L = \frac{d^2\Phi}{d\omega dA} \cos \theta, \quad [L] = \frac{\text{W}}{\text{sr m}^2}. \quad (2.6)$$

It has the important property of being constant along rays in space in which light propagates without obstruction, and therefore is an important basis measure in Computer Graphics. **Spectral radiance** refers to the differential radiance at a specific wavelength.

2.2.2 The Plenoptic Function

There is a range of optical effects, such as variation in **phase** and **polarization** that human observers cannot directly perceive. If we omit these from our consideration, we will not exhaustively model the real world, but still be able to treat the most relevant subset of light-matter interaction. Then, the spectral radiance in any given scene can be exhaustively expressed by means of the seven-dimensional **plenoptic function**

$$P(\theta_w, \varphi_w, x_w, y_w, z_w, \lambda, t) \quad (2.7)$$

defined by Adelson and Bergen (1991). It states the spectral radiance of wavelength λ along the ray that passes through the point in space (x_w, y_w, z_w) in direction (θ_w, φ_w) at time t .

2.3 Imaging

2.3.1 Light Fields

Human color perception emerges as the result of simultaneous stimulation of three distinct receptor types in the human eye. These receptor types have different spectral sensitivity, which we in practice model as three **primary colors** red, green and blue. Perception of composed colors emerges as the result of simultaneous stimulation of the separate receptors for the primary colors.

In this work, we will assume that we do not encounter materials which, when taking part in light transport, scatter light in a wavelength different than the input wavelength. This explicitly excludes fluorescent material, but it allows us to model the result of light transport separated into three discrete wavelength bands, or **color channels** which we associate again with the observable primaries red, green and blue. We further assume that the illumination we encounter also to be separatable into these channels with broad support in the spectrum; while this does not permit to model the behavior of materials in narrow-band illuminations (such as from low-pressure sodium vapor lamps), it is sufficient both for most natural light sources (such as sunlight or fire) and for many artificial indoor illuminations that are chosen for their comfort. For notational simplicity, we will in general not introduce respective indices, but describe light transport neutrally for all wavelengths and imply a separate, but otherwise identical treatment of each channel.

As we have already stated in Section 2.2.1, radiance along an unobstructed ray in space remains constant. As a ray in space can be defined by four scalars, such as positions (u, v) in a plane and a direction (θ, φ) , the plenoptic function, when restricting observers to the outside of a given scene, is actually redundant. For a static scene – this drops the variation with t – and the conventional independent description of color channels – dropping λ – this yields a four-dimensional description of the radiance distribution in the

scene

$$L(u, v, \theta, \varphi) \quad (2.8)$$

which is called the **light field** (Levoy and Hanrahan, 1996).

For observer points restricted to one side of the defining plane, the light field is often alternatively parameterized by the intersection of rays with a second plane as $L(s, t, u, v)$ with intersection points (s, t) and (u, v) in the respective planes.

2.3.2 Digital Images

This permits us now to formalize an **image** with width w and height h as a mapping

$$I : [0, w) \times [0, h) \rightarrow \mathbb{R} \quad (2.9)$$

that usually represents a visual signal on a 2D sensor or on the human retina. It can be recorded using a **digital imaging sensor**, a device that integrates over slices of the plenoptic function. It can generally be modeled as a function

$$\begin{aligned} (\mathbb{R}^7 \rightarrow \mathbb{R}) &\rightarrow (\mathbb{R}^2 \rightarrow \mathbb{R}) \\ P(\theta_w, \varphi_w, x_w, y_w, z_w, \lambda, t) &\mapsto I(x, y) \end{aligned} \quad (2.10)$$

The most common type is a **digital camera**, which performs this mapping as

$$\begin{aligned} I(x, y) = \int_{\theta} \int_{\varphi} \int_{\lambda} \int_t &w_a(x, y, \theta, \varphi) \cdot w_s(x, y, \lambda) \cdot w_e(x, y, t) \\ &\cdot P(G(x, y, \theta, \varphi), \lambda, t) dt d\lambda d\varphi \cos \theta d\theta \end{aligned} \quad (2.11)$$

for pixel locations (x, y) , where $w_a(x, y, \theta, \varphi)$ encodes the relative transparency of the camera's aperture from angles (θ, φ) in camera-local coordinates, $w_s(x, y, \lambda)$ models the spectral response, and $w_e(x, y, t)$ expresses the generalized exposure time. G expresses the camera geometry here as a function which maps the integration angles (θ, φ) and the pixel positions (x, y) to the world coordinates (x_w, y_w, z_w) and the world directions (θ_w, φ_w) .

In practice, digital cameras discretize the image by **sampling** only at integer positions $(x, y) \in [0, w) \times [0, h) \subset \mathbb{Z}^2$. These positions are called **pixel positions**, and the image is thus composed of discrete **picture elements**, or, in short, **pixels**.

Mimicking the human observer, digital color cameras are also sensitive to light in separate wavelength bands; usually, also in the red, green and blue part of the light spectrum. Some cameras provide this separation for each pixel (x, y) , but most employ a **Bayer pattern**, i. e. they are only sensitive to a single one of the color bands in each of the pixels they record, and require a software to reconstruct all color channels in a single pixel by incorporating several neighboring pixels. In this thesis, we will consider this problem solved and refer the reader to Ajdin et al. (2008) who have recently contributed to the constantly evolving subject.

The relationship between the radiant energy that the image sensor for a single pixel accumulates and the resulting value returned by the camera is expressed by the camera's **response function**. In this paper, we employ cameras that we can either assume to be directly linear in response, or for which the response can be linearized using a known response curve. Digital cameras discretize the pixel energy for the representation in finite data types, in a process called **quantization**. This causes noise, but is usually masked by more noise inherent to the acquisition process. For ease of use, we will nevertheless process digital camera data using floating point numbers. Finally, cameras have a limited **dynamic range**, which is the ratio of the maximum energy to the minimum energy their output can possibly represent. In static scenes, however, the radiant flux in each pixel is constant over time; by combining the observations of several images of the same scene with different camera exposure durations (Robertson et al., 2003), we can therefore obtain an accurate estimate of the flux. Thus, we can simulate any exposure pattern and achieve **high dynamic range** or, in short, **HDR** imaging.

The space of possible camera geometries is discussed by Adams and Levoy (2007). For the purpose of this thesis, though, we only need to keep the information that digital cameras may be treated as linear mappings from real-world light to the space of images.

For a static scene, a fixed color channel and an external observer, knowing the light field is sufficient to synthesize any image of the scene. Conversely, a light field can be understood as a collection of images taken in positions (u, v) so that the angularly parameterized directions (θ, φ) map to positions (x, y) in image space. Within the range of the rays for which it is known, it permits to simulate any digital imaging post-hoc by collecting and integrating over the rays of the light field in a virtual digital camera.

2.4 Modeling Appearance

For acquiring the appearance of a scene, we need to generate a model of the scene's interaction with light, so that we can, given a specified illumination L_i incident on the scene, compute the generated illumination L_o in a process referred to as **rendering**. Two principles for scene description exist: local scene descriptions express the scene as a collection of surfaces of known properties, while global descriptions treat the light transport in the scene as a whole.

Local descriptions permit to easily assemble scenes in a bottom-up approach, and are well suited for predictive rendering in Computer Aided Design, where abstract scene descriptions exist long before a real-world object is created. In contrast to that, global descriptions are particularly apt for acquiring already existing scenes, discovering scene structure in a top-down fashion.

Both approaches provide simplifications to reduce the dimensionality of the required data structures (see Figure 2.1). While local approaches maintain expressiveness, and provide information sufficient to render images for arbitrary view points, camera parameters or

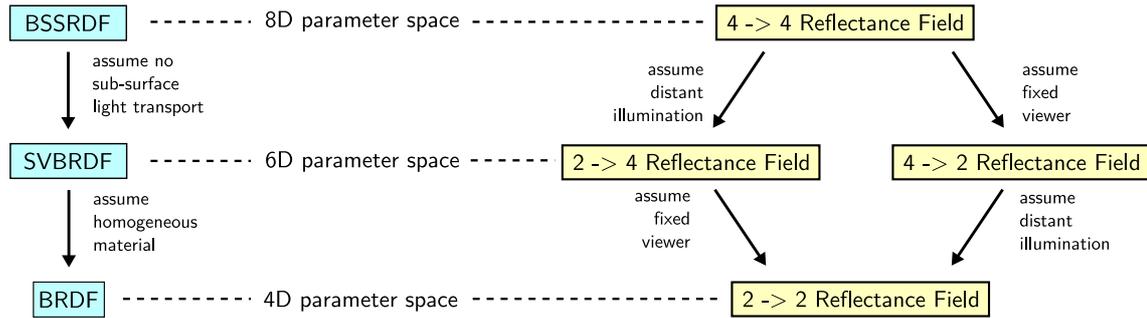


Figure 2.1: Mind map for relightable scene representations, as a subset of the taxonomies presented in the work of Rusinkiewicz and Marschner (2000) and Lensch (2003). Local modeling approaches (cyan, to the left) simplify by restricting the generality of the model, while maintaining its expressiveness, while global approaches (yellow, to the right) restrict the expressiveness, but maintain generality.

illumination choices, they restrict the generality of the possible material classes. Global approaches, on the other hand, work with any material even in the 4D case, but constrain the possible viewer parameters or the freedom of choice of illumination.

2.4.1 Local Appearance Descriptors

Properties of Scattering Materials

We will first look at local appearance descriptors, and imagine our scene to consist of surfaces, which scatter light according to their material properties. For the materials we are concerned with in this thesis, the **superposition principle** holds, and light transport is **linear** in the space of possible incident illuminations. Consider a surface where each surface point (u_i, v_i) is illuminated by a light field $L_i(u_i, v_i, \theta_i, \varphi_i)$ with incident angles (θ_i, φ_i) . The outgoing radiance at a point (u_o, v_o) in directions (θ_o, φ_o) , can then be described using the **reflection equation**

$$L_o(u_o, v_o, \theta_o, \varphi_o) = \int_{u_i} \int_{v_i} \int_{\theta_i} \int_{\varphi_i} L_i(u_i, v_i, \theta_i, \varphi_i) \cdot S(u_i, v_i, \theta_i, \varphi_i; u_o, v_o, \theta_o, \varphi_o) d\varphi_i \cos \theta_i d\theta_i dv_i du_i \quad (2.12)$$

where S is called the **Bi-Directional Scattering-Surface Reflectance Distribution Function**, or **BSSRDF** (Nicodemus et al., 1977) with unit $[S] = \frac{1}{\text{sr m}^2}$

For opaque surfaces lacking sub-surface light transport, S collapses for $(u_i, v_i) \neq (u_o, v_o)$;

in this case, for any point (u, v) on the surface,

$$L_o(u, v, \theta_o, \varphi_o) = \int_{\theta_i} \int_{\varphi_i} L_i(u, v, \theta_i, \varphi_i) \cdot f_r(\theta_i, \varphi_i; u, v; \theta_o, \varphi_o) d\varphi_i \cos \theta_i d\theta_i \quad (2.13)$$

for the six-dimensional **Shift-Variant Bi-Directional Reflectance Distribution Function** or **SVBRDF** $f_r, [f_r] = \frac{1}{\text{sr}}$. It is also called the **Spatially Varying Bi-Directional Reflectance Distribution Function**.

The SVBRDF is defined also for transparent surfaces; the transmissive component of light transport can then be modeled by the analogous **Shift-Variant Bi-Directional Transmittance Distribution Function** or **SVBTDF**.

For materials of uniform surface appearance, the location (u, v) is often dropped, yielding the four-dimensional **Bi-Directional Reflectance Distribution Function** or **BRDF** and the **Bi-Directional Transmittance Distribution Function** or **BTDF**, respectively. BRDFs which are simply constant are called **Lambertian**, in honor of Johann Heinrich Lambert who first investigated it in 1760. The value of such a BRDF is referred to as the **surface albedo**.

The Rendering Equation

Augmenting 2.12 by possible local light emission $L_e(u_o, v_o, \theta_o, \varphi_o)$ delivers

$$L_o(u_o, v_o, \theta_o, \varphi_o) = L_e(u_o, v_o, \theta_o, \varphi_o) + \int_{x_i} \int_{y_i} \int_{\theta_i} \int_{\varphi_i} L_i(u_i, v_i, \theta_i, \varphi_i) \cdot S(u_i, v_i, \theta_i, \varphi_i; u_o, v_o, \theta_o, \varphi_o) d\varphi_i \cos \theta_i d\theta_i du dv \quad (2.14)$$

This equation is known as the **rendering equation**. It was introduced by Kajiya (1986) in a slightly different formulation integrating over the discrete surfaces of a scene. Its recursive nature – the incident light field L_i stems from radiant light fields L_o exitant from different scene surface points – makes it in general challenging to evaluate. Note that in the notation of the rendering equation, all external illumination is expressed as a scene surface with appropriately chosen emission component L_e .

2.4.2 Global Appearance Modeling

Reflectance Fields

Treating light transport for a scene as a whole removes the requirement of modeling the scene as a collection of surfaces with known BSSRDFs; instead, we consider the direct relationship between any (external) incident light field L_i and the exitant light

field L_o (observed outside the scene) that results from in-scene scattering. In search for a function

$$\begin{aligned} (\mathbb{R}^4 \rightarrow \mathbb{R}) &\rightarrow (\mathbb{R}^4 \rightarrow \mathbb{R}) \\ L_i(u_i, v_i, \theta_i, \varphi_i) &\mapsto L_o(u_o, v_o, \theta_o, \varphi_o) \end{aligned} \quad (2.15)$$

that expresses this relationship, we can again make use of the linearity of light transport. We may consider the contribution of each ray of incident radiance to the outgoing light field separately, and observe in the **reflectance equation** that

$$\begin{aligned} L_o(u_o, v_o, \theta_o, \varphi_o) \\ = \int_{u_i} \int_{v_i} \int_{\theta_i} \int_{\varphi_i} L_i(u_i, v_i, \theta_i, \varphi_i) \cdot R(u_i, v_i, \theta_i, \varphi_i ; u_o, v_o, \theta_o, \varphi_o) d\varphi_i \cos \theta_i d\theta_i dv_i du_i \end{aligned} \quad (2.16)$$

for a scene-dependent function $R : \mathbb{R}^8 \rightarrow \mathbb{R}$ that we call **reflectance field** (Debevec et al., 2000). R states how much of the incident radiance along the ray specified by $(u_i, v_i, \theta_i, \varphi_i)$ contributes to the exitant radiance along the ray specified by $(u_o, v_o, \theta_o, \varphi_o)$. Note the similarity between equations 2.16 and 2.12 – the reflectance field is equivalent to the BSSRDF in expressiveness, and only differs with respect to the parameterization; the reflectance field is parameterized in arbitrary world coordinates and does not require the definition of a surface. In contrast to the rendering equation 2.14, the reflectance equation is not recursive, as L_o and L_i are separate entities defined outside the modeled scene. This makes the evaluation much simpler, at the cost of a scene representation which requires more memory.

Just as in the relationship between BSSRDF and BRDF, it is possible to reduce the dimensionality of reflectance fields by imposing constraints. In contrast, though, these assumptions do not restrict the generality – the range of modeled material types is maintained – but the expressiveness, that is, either the space of possible incident illuminations shrinks or the space of possibly simulated cameras is reduced. To distinguish between them, we label them according to the dimensionality of the incident and exitant light fields; the general case then is represented by a $4 \rightarrow 4$ **reflectance field**.

One customary assumption is that the incident light field be distant, i. e.

$$L_i(u_i, v_i, \theta_i, \varphi_i) = L_i(u'_i, v'_i, \theta_i, \varphi_i) \quad \forall u_i, v_i, u'_i, v'_i \in \mathbb{R} \quad (2.17)$$

Then, the dependence of L_i from u_i and v_i may be dropped and the **reflectance equation for distant illumination** reads

$$L_o(u_o, v_o, \theta_o, \varphi_o) = \int_{\theta_i} \int_{\varphi_i} L_i(\theta_i, \varphi_i) \cdot R(\theta_i, \varphi_i ; u_o, v_o, \theta_o, \varphi_o) d\varphi_i \cos \theta_i d\theta_i \quad (2.18)$$

for a six-dimensional $2 \rightarrow 4$ **reflectance field** R . As the distant light fields form a linear subspace of the general light fields, this distant-light reflectance field is still a light

transport model that is linear in its input.

An alternative assumption restricts the dimensionality of the modeled exitant light field L_o . Here, we will not assume that it is identical in every position (u_o, v_o) , as this would impose unrealistic restrictions on the possible scenes. Instead, we will rather restrict the possible evaluation of the output field to a fixed, implicitly given (u_o, v_o) and obtain

$$L_o(\theta_o, \varphi_o) = \int_{u_i} \int_{v_i} \int_{\theta_i} \int_{\varphi_i} L_i(u_i, v_i, \theta_i, \varphi_i) \cdot R(u_i, v_i, \theta_i, \varphi_i; \theta_o, \varphi_o) d\varphi_i \cos \theta_i d\theta_i dv_i du_i \quad (2.19)$$

again reducing R to six dimensions, but this time defining the $4 \rightarrow 2$ **reflectance field**. It maintains the ability to simulate the illumination of the scene with arbitrary light rays, but cannot simulate arbitrary digital cameras anymore; it is only suitable for camera configurations that measure rays passing through (u_o, v_o) .

Image-Space Relighting

As we may recall from Equation 2.11 and the closing remark from 2.3.2, digital imaging itself is a linear process. That means, that instead of fixing (u_o, v_o) , we can also fix an arbitrarily defined digital camera, along with exposure characteristics w_e , aperture function w_a and geometrical definition G , and bake it into the definition of R , describing light transport including the imaging step as

$$I(x, y) = \int_{u_i} \int_{v_i} \int_{\theta_i} \int_{\varphi_i} L_i(u_i, v_i, \theta_i, \varphi_i) \cdot R(u_i, v_i, \theta_i, \varphi_i; x, y) d\varphi_i \cos \theta_i d\theta_i dv_i du_i \quad (2.20)$$

In this case, R is a six-dimensional $4 \rightarrow 2$ reflectance field in image space, again being a linear descriptor of light transport.

Taking both the assumptions of distant illumination and a fixed camera simultaneously reduces R to a four-dimensional $2 \rightarrow 2$ **reflectance field** and, dropping no longer needed indices, yields the **reflectance equation for image-space relighting**

$$I(x, y) = \int_{\theta} \int_{\varphi} L(\theta, \varphi) \cdot R(\theta, \varphi, x, y) d\varphi \cos \theta d\theta \quad (2.21)$$

When implementing it in an application, it is approximated by a finite sum. This can be achieved by expressing the incident light $L(\theta, \varphi)$ as a linear combination

$$L(\theta, \varphi) \approx \sum_{0 \leq i < n} l_i L_i(\theta, \varphi) \quad (2.22)$$

of n basis illuminations $L_i(\theta, \varphi)$ with appropriately chosen linear coefficients l_i . Inserting

this into equation 2.21 and exchanging summation and integration yields

$$\begin{aligned}
 I(x, y) &= \int_{\theta} \int_{\varphi} \sum_{0 \leq i < n} l_i L_i(\theta, \varphi) \cdot R(\theta, \varphi, x, y) d\varphi \cos \theta d\theta \\
 &= \sum_{0 \leq i < n} l_i \underbrace{\int_{\theta} \int_{\varphi} L_i(\theta, \varphi) x \cdot R(\theta, \varphi, x, y) d\varphi \cos \theta d\theta}_{=: R_i(x, y)}
 \end{aligned} \tag{2.23}$$

This reveals two important properties of this approximation: for one, no matter how small n is chosen, the approximate rendered image I corresponds to the accurate image for some physically realizable illumination, which approximates the desired target illumination. Second, the pre-integrated R_i , as defined in Equation 2.23, are computed from L_i and R analogously to Equation 2.21, implying that the R_i are actual images of the scene as illuminated by the basis illuminations L_i .

Accordingly, we can easily define a practical pipeline for acquiring and rendering with $2 \rightarrow 2$ reflectance fields in image space:

acquisition step: Expose a scene to a set of different distant illumination conditions $L_i(\theta, \varphi)$ and record corresponding images $R_i(x, y)$ in these illuminations with a digital camera.

modeling step: Infer an approximative description of R . Due to the relationship in Equation 2.23, this step is optional and may be skipped.

rendering step: For any desired target illumination L , generate an approximative image I of the scene in that light.

This pipeline, as a whole, defines the method of **image-space relighting**.

3 Related Work

3.1 Overview

In this chapter, we will review different alternatives to the acquisition of relightable scene representations. While we will discuss image-space relighting techniques in depth, we will coarsely discuss the other methods by representative work due to the broadness of the subject. In the discussion of related methods, we will group them according to the prevalent data structure or basis problem they solve. As the development of rendering already spans several decades, some remarkably similar concepts have turned up independently from one another in the context of different applications; we will discuss them in a single place and reference them accordingly.

Lehtinen (2007) discusses a unified framework for measured and simulated data. Nevertheless, as in Section 2.4, we will discern local strategies (Section 3.3), that model scenes as compositions of surface primitives with attached reflectance descriptors, and global strategies (Section 3.4), that describe light transport within scenes as a whole. In addition, we will look into hybrid approaches (Section 3.5) which combine massively acquired data with coarse scene geometry.

Before we can turn our attention to relightable scene descriptions, though, we review the work on light fields (Section 3.2) and image-based rendering, as they form basic components in many global strategies.

3.2 Fixed-Illumination Models

3.2.1 Light Field Techniques

McMillan and Bishop (1995) recorded images in real-world scenes and treated them as observations of the plenoptic function; novel views could then be synthesized by re-sampling it. While this can be seen as a precursor of all techniques discussed in this subsection, the terminology of later articles changed; as most scenes are only modeled to be observed from the outside, modeling by a light field (see Section 2.3.1) is appropriate. The light field was first used for rendering by Levoy and Hanrahan (1996), and their article already defined two typical application cases: light fields can be used for recording real scene appearance, but also as an intermediate rendering data structure, where the

light field is computed as an evaluation of the Rendering Equation and stored to facilitate fast rendering at a later time.

As a data structure, light fields always form a discrete, sampled representation of the abstract concept, for instance, by storing a collection of images taken from different view points. Therefore, special care needs to be taken during rendering to avoid artefacts such as **ghosting**, i. e. non-corresponding overlays of neighboring images. Isaksen et al. (2000) describe a dynamical re-parameterization enabling fast change of focal parameters at rendering time. Kubota et al. (2004) combine different views into an all-in-focus representation, and Liu et al. (2006) employ a stereo reconstruction for improved quality. Peter and Straßer (2001) compressed light fields in wavelet streams, permitting fast rendering.

An analysis of optimal reconstruction kernels for light field rendering for diffuse scenes based on depth complexity is provided by Chai et al. (2000).

Next to camera arrays, light fields have been recorded by hand-held, modified digital SLR cameras (Georgiev et al., 2006; Ng, 2006; Ng et al., 2005; Veeraraghavan et al., 2007) and used in microscopy (Levoy et al., 2006). Time-variant light fields can be recorded by distributed synchronized (Yang et al., 2002) or unsynchronized camera arrays (Wang et al., 2007a), and transmitted to a real-time display (Arai et al., 1998; Matusik and Pfister, 2004).

In addition to measuring object appearance, light fields have also been used as illuminants for other rendering methods. Heidrich et al. (1998) computed the light field of a lamp from its geometric description and stored it for later use, Goesele (2004); Goesele et al. (2003) measured the emitted light field of real light sources. Unger et al. (2003) recorded the light field of a complex illumination environment for use in later rendering.

Enabling new applications, Heidrich et al. (1999) analyzed a light field to model reflective and refractive objects. Wang et al. (2005b) developed a light field morphing algorithm. Ziegler et al. (2007) demonstrated how to convert light field data into a representation suitable for the use in holograms.

3.2.2 The Lumigraph

In parallel to the light field article (Levoy and Hanrahan, 1996), Gortler et al. published "The Lumigraph" (1996). The **lumigraph** is also a 4D function describing radiance, and equivalent to the light field. In addition, the lumigraph paper describes a parameterization of the function close to the actual surface geometry, which creates a higher rendering quality.

Schirmacher et al. (1999b) presented an adaptive sampling scheme for synthetic scenes, investigated high-quality rendering with interactive performance (Schirmacher et al., 2000), and later added on-line acquisition to their pipeline (Schirmacher et al., 2001). While many algorithms re-sample acquired data into a regular structure before display, Buehler et al. (2001) directly synthesize output images from input camera data.

3.2.3 Surface Light Fields

Surface light fields (Miller et al., 1998) define the parameterization of the light field directly on scene surface primitives, which avoids most ghosting artefacts on scene contours, but once more requires the surface geometry to be known. Related techniques on surface light fields encompasses recording and compressing (Wood et al., 2000), fast rendering (Chen et al., 2002) and morphing (Zhang et al., 2002), and incremental acquisition (Coombe et al., 2005). Vlasic et al. (2003) added varying opacity to the model.

3.2.4 Environment Maps

When used to describe illumination on scenes, distant light fields (as in Equation 2.17) take an important role as a 2D radiance data structure. Distant light fields are also known as **environment maps** (Greene, 1986) and were introduced by Blinn and Newell (1976) for displaying surface effects for BRDFs with mirroring components. While this work did not approximate the full reflection equation for arbitrary BRDFs, this can be achieved quickly by sampling environment maps with directional light sources (Agarwal et al., 2003). Recently, Annen et al. (2008) introduced a real-time rendering algorithm for environment map illumination that also permits soft shadows to appear.

Environment maps were first acquired by taking a photograph of a mirror sphere by Miller and Hoffman (1984). Debevec (1998) called such a sphere a **light probe**, acquiring image sequences for an HDR reconstruction method (Debevec and Malik, 1997) to get a full-range radiance map. Nishino and Nayar (2004) extracted environment maps from photographs of human eyes. Sato et al. (1999b, 2003) obtained them by an analysis of cast shadows.

Assuming an isotropically emitting environment, Sato et al. (1999a) extended the idea to reconstruct a 3D distribution of incident illuminators using stereo.

In order to enhance the immersion of real-time interactive applications, Ghosh et al. (2005) created an approximative immersive display of a virtually generated environment map by modulating indirect room illumination accordingly.

3.2.5 Image-Based Rendering

Without the ability to relight scenes, **image-based rendering** incorporates image-space encoded data with surface descriptions¹. These descriptions take the form of **billboards**, i. e. flat scene components carrying images, or **impostors**, general geometry proxies of

¹We note that, throughout the literature, the term “image-based rendering” is often employed for various types of scene representation and rendering approaches that involve image-space data, including BRDF measurement with cameras and image-space relighting. For this thesis, we employ a more strict terminology.

coarse resolution that express how the input images need to be warped to avoid blending artifacts (Schauffler, 1998). Scene geometry itself can also be expressed in image space, for instance assign a depth to each pixel (Chang et al., 1999; Oh et al., 2001; Oliveira and Bishop, 1999)

Image-based rendering has been applied to the modeling of architecture (Debevec et al., 1996, 1998) as well as to the display of synthetic scenes (Lischinski and Rappoport, 1998).

Khan et al. (2006) have investigated approximations for material edits using a single image with annotations as user input, creating perceptually plausible results in the fixed input view.

3.3 Local Light Transport Descriptors for Relightable Scenes

Light fields only describe visual appearance of a scene when illuminated with an illumination fixed at acquisition time. We now review approaches that provide the capability to relight scenes, and render them in novel input illuminations. First, we discuss techniques that describe scenes as collections of surfaces with locally defined illumination response.

3.3.1 Evaluating the Rendering Equation

In contrast to the global models, where the definition of the data structure itself defines the achievable quality of the rendering and the expressiveness – for instance, for which viewpoints the scene can be rendered – local schemes model the scene largely independent of the rendering, and require an evaluation of the rendering equation (2.14). As this equation is recursive, this evaluation may incur high costs. Therefore, a range of approximations have been found which differ with respect to the precision provided, the manageable scene complexity and other properties. We will now discuss the most prominent examples.

Rendering by Drawing Primitives

First, we will focus on approaches that render 2D, planar images. Quite a few of these approximations share not only the expression of the scene surfaces as collection of surface primitives – triangles being a popular choice – but also the idea of mapping the surface primitives from scene space to the image plane, and drawing them there in a 2D data structure referred to as the **frame buffer**. The algorithm determining the pixel values at each coordinate is called **shader**.

In scenes with non-trivial surface configurations, rendering algorithms must solve the **visibility problem**: in each pixel, the opaque surface closest to the camera needs to

occlude those behind it, which may not overdraw it. The **painter's algorithm** solves this by sorting the surfaces according to their distance to the camera along the viewing direction, which is referred to as "depth", and drawing the distant surfaces first. The closer surfaces, being drawn later, occlude the distant surfaces in a natural fashion.

The painter's algorithm is easy to understand, but it cannot handle arbitrarily intersecting surfaces or configurations that cannot be depth-ordered. Further, the pre-ordering before rendering can be an expensive operation. The **Z buffer** algorithm (Catmull, 1974) therefore solves the visibility problem in a different fashion: the so-called Z buffer – with the same dimensions as the frame buffer – stores the depth of drawn surface primitives in each pixel. Subsequent primitives are only drawn if they are closer to the camera. This method is supported in hardware by most graphics processing units deployed today, and part of the OpenGL® rendering standard. For several years now, this method has provided affordable real-time renderings in consumer hardware for increasingly complex scenes.

The first available implementations of these algorithms were restricted to local evaluations of the reflection equation 2.13, coping with the expense of the integral by restricting the evaluation to a discrete set of non-zero illumination, corresponding to **directional** or **point light sources**. The effect of shadows can be simulated then by testing the visibility of the light source from the surface point currently being rendered. Popular algorithms for this employ **shadow volumes** (Crow, 1977), which render a stencil of the shadow into the scene, or **shadow maps** (Williams, 1978) which represent a variation of the Z buffer technique for the scene as seen from the light source.

Since the early days of computer graphics, solutions for **texturing** of surfaces (Catmull, 1974, 1975) have been available; for a long time though, only as shift-variant Lambertian BRDFs – basically 2D images fixed to the surface, looking the same from each observation direction.

In the 1990, with the advent of graphics cards with more programmability, real-time solutions for more complex illumination models became feasible, for instance, through the work of Heidrich (1999).

Ray Tracing

Unlike the per-primitive approaches, **ray tracing** (Whitted, 1979, 1980) renders scenes one pixel at a time. It employs ray optics to trace light rays that arrive in the camera pixels backwards through the scene until they hit scene surfaces. For the found locations, one again computes appearance locally with a **shader**, testing the occlusion of light by different scene surfaces with the same method that is used for finding the first surface. This ray-surface intersection, along with sorting for the first surface, is an expensive operation if performed in a brute-force manner. Luckily, efficient acceleration structures such as the k -D tree (Bentley, 1975; Bentley and Friedman, 1979) exist that, after some pre-computation on the scene geometry, accelerate the process greatly.

Ray tracing is not only used to solve visibility; shooting additional, secondary rays from the first hit surface can be used to approximate the reflection equation 2.13 better. For mirroring BRDFs and refracting BTDFs, only one further ray each needs to be generated and traced to accurately render reflections. Shooting several rays for the approximation of the reflection equation (Cook et al., 1984) can be used for a better rendition of effects such as soft shadows from extended light sources and glossy materials, and for the approximation of the integration in the simulated digital camera.

Due to the sub-linear performance of the acceleration structures, ray tracing scales excellently with geometric scene complexity. As the values for different pixels can be computed independently from each other, ray tracing is highly parallelizable. Being concerned with the traversal of individual rays, ray tracing works for very general camera geometries, and is also used as building block in more complicated rendering approaches.

Global Illumination Methods

The approaches that we have discussed so far compute the interaction of light with surfaces only locally for the surfaces directly observed in each camera pixel. Given the recursive nature of the rendering equation, this is not sufficient for photorealism. **Global illumination** methods therefore simulate light transport also between scene surfaces – at the expense of higher computation time and more complex data structures.

The **radiosity method** (Goral et al., 1984) establishes **form factors** for describing the amount of radiative exchange between pairs of Lambertian surface primitives and solves a linear system to describe the equilibrium state of light transport in a closed scene. It was later extended (Rushmeier and Torrance, 1990) for specular materials.

The article on the rendering equation (Kajiya, 1986) discusses Monte Carlo solutions to the problem; later solutions are, for instance, **bi-directional path tracing** (Lafortune and Willems, 1993) or **photon mapping** (Jensen and Christensen, 1995).

Fast approximations are also possible with the help of wavelets (Christensen et al., 1996; Kontkanen et al., 2006; Lewis and Fournier, 2000).

Pre-computed Radiance Transfer

Bridging the gap between accurate, but slow, and approximate, but real-time techniques, **pre-computed radiance transfer** or **PRT** methods (Sloan et al., 2002) compute an approximate solution for the reflection integral in a pre-processing step, and perform a change of basis in order to approximate the evaluation of the integral in a lower-dimensional subspace, where it can be quickly evaluated at run-time. Possible choices for basis functions include spherical harmonics (Sloan et al., 2002), wavelets (Ng et al., 2003, 2004), piece-wise constant functions (Xu et al., 2008) or spectral mesh bases (Wang et al., 2007b).

Later extensions added initially lacking support for more BRDF types (Liu et al., 2004b; Wang et al., 2004), incremental rendering (Overbeck et al., 2006) and real-time geometry animation (Feng et al., 2007; Iwasaki et al., 2007).

While PRT techniques usually pre-compute all scene interactions up to the point where the real-time evaluation takes place at the visible surface only, Mei et al. (2004) added support for PRT with inter-object light transport for shadows and reflection, which was extended by Pan et al. (2007).

For diffuse surfaces, steerable bases (Ashikhmin and Shirley, 2002) provide extra compression for directional light sources. In what can be considered an early predecessor of PRT, Cabral et al. (1999), employed a set of BRDF pre-convolved environment maps, however without fully modeling occlusion.

3.3.2 BRDF Methods

After discussing the options for evaluating the rendering equation, we can now focus on the illumination response models inside. We will start with the BRDF, and then follow alternative descriptors with increasing complexity.

Modeling the BRDF

Tailored to specific use cases, there are different kinds of BRDF representations. The most compact representations take the form of instances of parametrized families of functions; only few parameters must be stored so that the function can be evaluated algorithmically. The earliest example of this model type is the constant **Lambertian BRDF** (Lambert, 1760) which just needs to store the surface albedo. Some existing BRDF models have been constructed phenomenologically, i. e. with the purpose of achieving a visual effect with few computations or direct parameter choice (Blinn, 1977; Lafortune et al., 1997; Phong, 1975). Others are based on physical models (Cook and Torrance, 1981; He et al., 1991) and their approximations (He et al., 1992), or have been created by simulating micro-geometric effects (Oren and Nayar, 1994).

In more general function spaces, BRDFs have been represented using spherical harmonics (Westin et al., 1992), wavelets (Lalonde and Fournier, 1997), or other functions with local support (Noe and Peroche, 2000).

BRDFs have also been constructed in a **data-driven** manner as linear combinations of basis BRDFs that resulted from real-world measurements (Matusik, 2003; Matusik et al., 2003a,b). There, the parameterization of the 4D input space plays an important role for the data compression efficiency (Rusinkiewicz, 1998), on which physical plausibility constraints may be imposed (Edwards et al., 2006).

Anisotropic BRDFs, which are variant under consistent rotation of the incident and outgoing directions by the surface normal, have been explicitly addressed by the models

of Kajiya (1985) and Poulin and Fournier (1990). While BRDF models are usually separated into the three RGB color channels, Gondek et al. (1994) and Xu and Sun (2006) have given spectral effects explicit treatment.

For rendering efficiency in real-time applications, BRDFs can be approximated by separable models (Kautz and McCool, 1999; Latta and Kolb, 2002; McCool et al., 2001).

Editing the BRDF

The choice of parameters for BRDF models can quickly become a difficult task, as in many cases, the parameters expose their meaning only in complex test renderings. This is addressed by visual parameter space navigation techniques (Ben-Artzi et al., 2006; Ngan et al., 2006). In 2008, Ben-Artzi et al. introduced a polynomial BRDF representation that even permits previewing with global illumination effects.

Colbert et al. (2006) suggested a free-form painting system for BRDF appearance which maintains physical plausibility constraints.

An early example of automatic parameter fitting of a BRDF model is given by Cabral et al. (1987), who reconstructed parameters for the Cook and Torrance model from a synthetic surface micro-geometry; later approaches often fit BRDF parameters to object measurements, which we will in detail discuss in the following paragraphs.

Measuring the BRDF

As a component of appearance acquisition, local surface descriptors such as the BRDF of a homogeneous object must be acquired as well. Gonioreflectometers can be used for that task; however, they are slow; Ward (1992) introduced a special device for parallelization of the acquisition, and fitted parameters for an anisotropic BRDF model that he published in the same paper. Image-based approaches with digital cameras are easier to set up, but only deliver sparse samples (Karner et al., 1996) which can however be used to fit model parameters. On homogeneous objects with curved object geometries, digital cameras can be used more effectively, yielding more samples faster, see the work of Lu et al. (1998, for cylinders) and Marschner et al. (1999, for (approximate) spheres).

If model parameters are sought, a single input image may be sufficient (Tominaga and Tanaka, 2000); a comparison of different models and their performance is provided by (Schirmacher et al., 1999a) and (Ngan et al., 2005).

The task of inferring reflectance properties from a possibly more complex scene is referred to as **inverse rendering**, which is intensively discussed by Ramamoorthi and Hanrahan (2001, 2004).

Extensions to the standard BRDF include time variance (Sun et al., 2007), and, for the special case of car paint, a synthetic sparkling layer introducing shift variance (Günther et al., 2005).

3.3.3 SVBRDF Extensions

Only for a small category of objects and scenes, which are locally composed of the same, smooth material, defining a BRDF is sufficient. In the general case, at the least a shift-variant BRDF is required. It can be expressed as point-wise defined BRDFs on object surfaces. For the parametrized models we have discussed above, this is usually a sufficiently compact representation. Being 6D instead of 4D, it is however considerably more challenging to measure. Therefore, the works we will discuss now mostly concentrate on the acquisition and measurement of a SVBRDF.

Measuring the SVBRDF

In contrast to the BRDF measurement methods we discussed before, SVBRDF approaches rarely focus on the acquisition of a material to be combined with an independent geometry model – this role is usually taken by the BTF, which we will discuss in Section 3.5.2. Rather, SVBRDF methods are typically geared towards modeling the appearance of a specific real-world object. Its geometry is sometimes estimated as an integral part of the approach, but often separately acquired, for instance using a structured light scanner.

As Lambertian components can be estimated more robustly from few input data, many approaches combine an estimation of non-Lambertian BRDF components over extended surface patches with a locally varying diffuse texture.

Sato et al. (1997) fitted a parametric BRDF model where specular BRDF components change slowly over the surface, recording their object with the help of a robotic arm. Boivin and Gagalowicz (2001) used a single image with geometry annotations as input, Okabe et al. (2006) used a single image with user-supplied normal information. For flat surfaces, Gardner et al. (2003) introduced an approach with a linear, moving light source; Paterson et al. (2005) describe a simpler setup with a camera, flash, and markers placed around the surface.

Inverse rendering has been applied to entire scenes; both outdoors (Yu and Malik, 1998), where light scattered back from the scene into the illuminating environment can be neglected, and indoors (Loscos et al., 2000, 1999; Yu et al., 1999), where global illumination needs to be inverted. For increased precision, this step can also be performed for outdoor scenes (Debevec et al., 2004; Tchou et al., 2004). For objects which can be rotated in their illuminating environment, few input images suffice (Nishino et al., 2001)

More general approaches aim for variation also of non-diffuse, difficult to estimate BRDF parameters. This can be achieved by estimating linear combinations of basis materials: (Lensch, 2003; Lensch et al., 2001, 2003a,b) automatically cluster surface materials, for which Lafortune et al. model parameters are fitted, and then describe the scene surfaces as linear combinations of the cluster's BRDFs. Weistroffer et al. (2007) decompose sparsely observed reflectance data into linear combinations of BRDF bases composed of radial basis functions or measured materials. Zickler et al. (2005, 2006) share reflectance

data on a recorded object directly between neighboring surface elements, but smooth out specular contributions over larger surface areas.

If the object to be modeled can be a-priori assumed to belong to a specific class of objects, it is possible to acquire information on the object class in a preparation step, and thus ease the burden of measurement for the actual acquisition. Blanz and Vetter (1999) have performed this for human faces, thus being able to obtain a geometry model plus diffusely varying surface data for the Phong BRDF from a single color image. With this model, facial appearance can be transferred between images (Blanz et al., 2004). In a previous publication (Fuchs et al., 2005a), we have used this model to also reconstruct non-diffuse facial BRDFs in a-priori defined surface zones on the face, using several images of the same face as input. Lee et al. (2005) constructed a morphable face model that also includes more densely measured reflectance data. Theobalt et al. (2007) use multi-view input video to estimate Phong BRDF parameters and diffuse texture of a human actor, enabling relightable free-viewpoint video.

Yu et al. (2006) expressed the SVBRDF as a linear combination of diffuse and specular basis materials, and acquired their scene in a-priori unknown incident illumination. Paris et al. (2003) describe a simplified model geared towards real-time rendering performance on human faces. Gu et al. (2006) measured time-variant SVBRDF data.

Editing the SVBRDF

For efficient editing, Lawrence et al. (2006) decomposed measured SVBRDFs into shade trees (Cook, 1984). Tan et al. (2005, 2008) discussed filtering parameters of micro-facet models in parameter space.

Wang et al. (2006) synthesized a time-variant SVBRDF from labeled examples.

3.3.4 Full BSSRDF techniques

With two more dimensions, the BSSRDF is again more difficult to measure than the SVBRDF. Hanrahan and Krueger (1993) discuss scattering behavior of layered surfaces. Jensen et al. (2001) approximate it in a 4D apparent BRDF (see 3.4.1 below) assuming uniform, distant illuminations which allows to pre-integrate the full BSSRDF. For human skin, Donner and Jensen (2006) analytically derived a parametric model, Weyrich et al. (2006) augmented a surface light transport model for human faces with sub-surface scattering components.

Treating sub-surface light transport as a light transport phenomenon between two surface points independent from the incident and exitant angles, Goesele et al. (2004) found a simplification which made measurement feasible.

Simplifying the BSSRDF to a function of volume density, non-surface light transport of smoke can be measured with time variance (Fuchs et al., 2007a; Hawkins et al., 2005a) or even a single image as input (Fuchs et al., 2006).

3.4 Global Light Transport Descriptors with Relighting Capability

The BSSRDF is a very general representation of surface light interaction, but is still locally defined. We will now return our attention to global approaches, that express the optical response of a scene to illumination in a global fashion.

3.4.1 Models in World-Space

4 → 4 Reflectance Fields

Reflectance fields are the classical global light transport descriptors (see Section 2.4.2). While defined for the general $4 \rightarrow 4$ case by the seminal work of Debevec et al. (2000), they have not been exhaustively measured as of today. Garg (2006) and colleagues (Garg et al., 2006) describe a feasible approach that includes an adaptive measurement system performing compression during the acquisition process, but have only demonstrated it for a narrow directional coverage of incident illumination.

The concept of a linearly separated light transport descriptor can be traced further back; if the parameterization of the dimensions is defined on object surfaces, the reflectance field is equivalent to the **GRDF**, or **Global Reflection Distribution Function** (Lafortune and Willems, 1994)

2 → 4 Reflectance Fields

Restricting the domains of incident light fields to distant light, $2 \rightarrow 4$ reflectance fields as a six-dimensional data structure are easier to acquire; adding time as an additional dimension, Einarsson et al. (2006) recorded the periodic motion of human actors.

Again, equivalent concepts of the $2 \rightarrow 4$ reflectance field can be found in earlier work; already in 1995, Fournier wrote: *“In fact, one can easily conceive of a bidirectional room reflectance distribution function (**BRRDF**) which for any directional light sources gives the reflected light towards an eye at infinity”*. Assuming distant illumination, the BRRDF coincides with a $2 \rightarrow 4$ reflectance field parameterized on scene surfaces. It is also sometimes called an **apparent BRDF** (Wong et al., 1997). Using a scene-global parameterization, Koudelka et al. (2001) employ an apparent BRDF and a depth map of the scene to simulate 4D illumination.

4 → 2 Reflectance Fields

With full 4D variability of incident illumination, but only a 2D observer variation, such as provided by a fixed camera, $4 \rightarrow 2$ reflectance fields again are six-dimensional. They have

been first measured by Masselus et al. (2003); for a narrow angular sampling, Sen et al. (2005) have demonstrated that faster acquisition is possible with a parallel recording of several cameras for a single projector light source. While their setup naturally records a $2 \rightarrow 4$ field, they exchange the roles of projector and camera virtually, exploiting Helmholtz reciprocity.

For $4 \rightarrow 2$ data, Chen and Lensch (2005) give an interpolation scheme for changing light source positions.

Being in the core interest of this thesis, $2 \rightarrow 2$ reflectance fields are discussed in more depth in a separate section 3.4.2 below.

3.4.2 Image-Space Relighting with $2 \rightarrow 2$ Reflectance Fields

Concepts

After the discussion of global and local light transport descriptors, we can now turn to the work most related to this thesis: the global modeling of light interaction for a fixed camera, or **image-space relighting**. Let us recall Equation 2.21:

$$I(x, y) = \int_{\theta} \int_{\varphi} L(\theta, \varphi) \cdot R(\theta, \varphi, x, y) d\varphi \cos \theta d\theta$$

which expresses the image brightness in pixel (x, y) as integral over a distant illumination distribution $L(\theta, \varphi)$ multiplied by a $2 \rightarrow 2$ reflectance field R . Image-space relighting is concerned with the efficient recovery, editing and rendering of R . In addition to the reflectance field terminology, mathematically equivalent descriptions are apparent BRDFs parameterized over a 2D image space (Wong et al., 2001) and environment mattes. **Environment mattes** (Zongker et al., 1999) solve the same problem as $2 \rightarrow 2$ reflectance fields, but have a different application in mind: extending classical matting techniques which aim at separating an object from its background for latter compositing with a new background, environment matting also takes the interaction of the object with its background into account, modeling refraction and reflections. The early environment matting publications (Chuang et al., 2000; Zongker et al., 1999) handle also illumination from the side, but have a simple reflection model, while later (Peers and Dutré, 2003), precision of the model gained importance, and the angular extent of the expressible illumination was reduced to cover the scene background only.

Even before these publications, Belhumeur and Kriegman (1996) investigated the space of images spanned by the response to varying illumination, and concluded that for convex Lambertian scenes, the entire space is spanned by only three basis images.

Measuring

As we have concluded from Equation 2.23, the reflectance field R can be directly acquired in slices R_i by recording images for a set of illuminations L_i . We can now distinguish existing work according to the choice of the L_i .

Directional light methods: Using a directional light source illuminating from the direction (θ, φ) , usually approximated by a small, isotropic point light source in sufficient distance from the scene, $R(\theta, \varphi, \cdot, \cdot)$ can be directly observed. The first measurements of this type (Debevec et al., 2000) were performed with a rig the authors called the **light stage**, where a point light was manually moved on a spiral path around a human head. Masselus et al. (2002) described a somewhat relaxed measurement setup; instead of a-priori defining the path the light should take and the direction they wanted to record R for, they moved a hand-held light source, and observed their position in calibration objects. Winnemöeller et al. (2005) even removed the calibration objects and inferred a illumination distribution by interpreting the recorded images directly.

Schechner et al. (2003) discussed multiplexed illumination as a means of increasing the possible signal-to-noise ratio. Gardner et al. (2004) added time-dependence to the measurement process; extracting local surface descriptors such as normals and albedo, Wenger et al. (2005) achieved editability and an efficient relighting of the performance of a human actor.

The work (Fuchs et al., 2007b) which we will describe in detail in chapters 4 and 5 also fits into the directional light category; modeling the reflectance field acquisition as a sampling problem, we show how choosing the illuminations in a way that enables pre-filtering reduces aliasing artefacts, and how to distribute the sampling directions non-uniformly for a more efficient measurement process. Reflectance field acquisition as a sampling problem was also discussed by Peers (2006). For objects consisting of glossy surfaces, Dumont et al. (2005) investigated the required sampling density of directional illuminations depending on object glossiness.

Explicit basis methods express the space of possible incident illuminations in relighting directly as a linear combination of basis illuminants. Wavelet environment matting (Peers, 2006; Peers and Dutré, 2003) is a prominent example: there, a wavelet basis is constructed adaptively during the measurement, performing scene-adaptive compression during recording. The high compression rate is an advantage of that approach; however, the illumination must be controlled most precisely. The wavelet environment matting paper uses a CRT display to achieve this, which has only a small angular coverage. The reflective light stage approach (Peers et al., 2006) extends this explicit control to a larger domain.

Inferred reflectance fields: The previously discussed approaches record reflectance fields as direct observations in given illumination conditions. Alternatively, one may postulate a strong assumption on the possible shape of R . While this restricts the generality of the method in the sense of constraining the class of describable objects, it facilitates the acquisition process. The original environment matting (Zongker et al.,

1999) technique assumed a non-zero response in a rectangle, its first extension (Chuang et al., 2000) employed Gaussian functions. Image-based environment matting (Wexler et al., 2002) estimated the properties of refractive objects from observations in front of moving photo backgrounds, assuming R to have small support (i. e. , small non-zero contributions) for which a diffuse object would present a failure case. Matusik et al. (2004) model R as a composition of piece-wise constant functions with rectangular support. While this imposes a strong restriction on the shape of R , the measurement actually works with observed, but uncontrolled illumination, and therefore can relight an entire cityscape. Peers and Dutré (2005) controlled the illumination explicitly, but used wavelet noise as illuminants. This permits to express $R(\theta, \varphi, x, y)$ in a wavelet base which is chosen independently for each coordinate pair (x, y) and achieves efficient compression.

Bayesian relighting (Fuchs et al., 2005b), which we will discuss in detail in Chapter 7, represents R only implicitly, describing the reflectance field of a scene relative to the behavior of a simple reference object. Hawkins et al. (2005b) exploit Helmholtz reciprocity for the measurement of R . Instead of sampling densely in image space (x, y) , and coarsely sampling the illumination space (θ, φ) , they virtually exchange the role of illumination (created by a laser) and camera (observing a diffuse sphere surrounding the scene). Thus, they obtain an instant observation of a densely sampled (θ, φ) slice of the reflectance field – at the cost of recording a separate image for each image pixel position (x, y) .

Wang et al. (2008) observed a reflectance field of a human subject in an infra-red wavelength band – invisible, and therefore unobtrusive – and transfer the gained knowledge to video data taken in room illumination, thus creating a more balanced, virtual illumination for a video-conferencing application.

Mann (1995) took a non-digital approach to acquiring and displaying illumination-variance of scenes, recording illumination-variant holograms.

Rendering

While simple rendering of reflectance field data is rather straight-forward and its implementation choices dictated by the underlying data structure, it pays off to invest extra effort. Masselus et al. (2004) investigated several general interpolation and representation methods in the context of 4D reflectance fields. We have introduced a specific approach (Fuchs et al., 2007c) that separately upsamples highlight and shadow information, thus improving the quality of an undersampled reflectance field by a software-only approach. An in-detail description of this method is found in Chapter 6.

By warping several basis reflectance fields in image space, Hawkins et al. (2004) created facial animations. Inferring additional data using 3D geometry, Jones et al. (2006) simulated 4D illumination on a human face. Peers et al. (2007) transferred reflectance field data on a video of a human face for relighting purposes.

The rendering of reflectance fields can be used to virtualize scenes, i. e. re-creating object appearance in the current, real-world illumination. Nayar et al. (2004) defined a setup consisting of a digital camera recording the scene illumination, and an electronic display showing a rendering of a reflectance field in this illumination. As we have recently demonstrated (Fuchs et al., 2008), it is possible to perform such a rendering with optical components only, shaping incident illumination from the back side of a display into the rendered image of a reflectance field (see Chapter 8).

Applications

Image-based relighting methods, can be evaluated quickly, and have therefore found their applications as in-between data structures for rendering even before the formalization as a reflectance field. By performing costly renderings in basis illuminations as a preprocessing step, fast rendering at run-time becomes possible. Nimeroff et al. (1994) pre-computed nine images in basis illuminations which approximated natural sky light, while being chosen in a way that made them steerable (i. e. , linear combinations of the basis images created a perceived motion of the sun over the sky). Teo et al. (1997) performed investigations into choices of basis illuminants for further light source types.

Such technology spurred applications in interactive lighting design (Dorsey et al., 1995), both for indoor (Dobashi et al., 1995) and outdoor scenes (Dobashi et al., 1996), the latter article adding handling of specularities and incorporating compression by vector quantization.

Schoeneman et al. (1993) investigated virtual illumination design in a synthetic scene: given a painted target specification, they chose illumination so that the rendered image in that illumination approximates the target. Anrys et al. (2004a,b) used this approach for virtual illumination of real-world scenes, in this way obtaining a specification how real world light sources would have to be dimmed or colored to generate the image in the real world. Mohan et al. (2007, 2005) replaced fixed light sources by indirect illumination with a moving head projector in a white box; while it is difficult to replicate the optimal virtual illumination that was the result of computation in reality, the box approach provides richer possibilities for lighting than the discrete set of light sources could.

As a means for stage lighting, the virtual lighting concept can be turned around, illuminating a human actor in the real world with intensity-controlled light sources (Debevec et al., 2002) or the bounce light from an umbrella, which is in turn lit by a projector (Okabe et al., 2007), so as to give the impression of the actor standing in a different place.

Image-space Reflectance Fields as Input Data

Image collections that vary with illumination have served as useful input data for **non-photorealistic rendering (NPR)** (Fattal et al., 2007; Raskar et al., 2004). Estimating

surface normals, Malzbender et al. (2006) exaggerated specularities to emphasize surface detail. Akers et al. (2003) blended the images for different illuminations into a single image, varying the blending factors with the pixel locations for an NPR rendering effect. Sunkavalli et al. (2007) decomposed time-lapse video into different illumination effects, providing artistic freedom in the recombination.

Marschner and Greenberg (1997) demonstrated that, knowing the appearance of a scene for a set of light basis illuminants, one can compute the incident illumination in an image of the scene within the subspace spanned by the basis.

3.5 Hybrid Approaches

Hybrid approaches take a role between global and local light transport descriptors. They describe scene appearance in a data-driven fashion and apply intensive pre-computation or globally acquired models. At the same time, they require some geometric information of the scene, in some cases to the level of defining individual surfaces. Thus, they provide a compromise between generality of the expressible light transport, and the achievable rendering quality.

3.5.1 Surface Reflectance Fields

Just like light fields, reflectance fields can be parameterized along scene surfaces in order to improve coherence during rendering. This structure is called **surface reflectance field**. Matusik et al. (2002a) proposed a parameterization on the **visual hull**, i. e. the volumetric intersection of object silhouettes as seen from several camera positions, of a scene and recorded objects with varying opacity, extending it (Matusik et al., 2002b) by environment matting (see Section 3.4.2 below). Weyrich et al. (2005) propose a rendering approach for deformable surface reflectance fields.

3.5.2 BTF Methods

The **Bi-Directional Texture Function** or **BTF** is in many respects equivalent to a $2 \rightarrow 4$ surface reflectance field. It generalizes the concept of surface texture to include illumination response, describing the 2D exitant light field from 2D surface points in response to a 2D, distant illumination, and subsuming the effect of local geometry (occlusions, BSSRDF etc.). It is different to the surface reflectance field in that it usually does not model an entire object, but instead is understood as a surface property. Therefore, its acquisition is often performed from a small, planar patch of a material sample, and can be applied to synthetic surface geometries for predictive rendering with realistic materials.

It was defined by Dana et al. (1996, 1999), who recorded a database of real-world BTFs which is publicly available².

Measuring the BTF

Han and Perlin (2003) proposed employing a kaleidoscope for simultaneously recording multiple views. Wang et al. (2005a) augmented it with an extra representation for geometry, improving the quality at object silhouettes. Rump et al. (2008) expressed it as a delta to a fitted BRDF model for the special case of car paint.

Compressing the BTF

With the large amount of data comprising a BTF, its compression is an important topic. This was addressed by Furukawa et al. (2002), also providing a fast measurement system, and by Müller et al. (2006), introducing an efficient compression system aligning coordinate systems. Suykens et al. (2003) found a factorized representation suitable for fast rendering.

Daubert and Seidel (2002) and colleagues 2001, discussed simplifications with parametric BRDF models for the special case of cloth rendering.

Synthesizing and Editing the BTF

As BTFs provide a data structure that can be evaluated locally, but encodes also global illumination effects, it can be used as an acceleration structure for rendering of synthetic scenes (Dischler, 1998). It can also be synthesized using small example patches, taking care of the local surface curvature (Leung et al., 2007; Liu et al., 2004a; Tong et al., 2002), incorporating user-painted input (Zhou et al., 2005), or from sparse measurements (Ngan and Durand, 2006).

Procedural (Mueller et al., 2007) and interactive (Kautz et al., 2007) editing methods have also recently become available. Sloan et al. (2003) presented a combined rendering method including pre-computed radiance transfer for non-local effects.

3.6 Discussion

Global and local scene descriptors target separate applications. While local descriptors are compact, by design usually easy to edit and transfer between surfaces, they can only work within specific, constraining assumptions. For instance, SVBRDFs only work for

²At the time of this write-up, it can be found at
<http://www.cs.columbia.edu/CAVE/software/curet/>

scenes composed of opaque surfaces. For scenes which cannot be modeled as collections of surfaces, even BSSRDFs are inapplicable – or are identical to $4 \rightarrow 4$ reflectance fields. Further, local approaches require possibly costly evaluation for full quality.

Global descriptors, on the other hand, are rarely constrained to specific material types. Even the simple approach we will encounter in Chapter 7 works for glass, metal, cloth, human skin, and cross-sections of oranges alike. These materials belong to very different classes which are hard to treat uniformly in a direct evaluation of the rendering equation. Another strength of global approaches is the quality of the rendering: for illumination conditions that were part of the initial measurement, they are able to deliver the image corresponding to the measurement, and therefore provide true photo-realism.

As a notable disadvantage, they are hard to edit, consume large amounts of memory, and, in the general case of a $4 \rightarrow 4$ reflectance field, require a measurement effort which, at the time of writing of this thesis, are impractical for full angular coverage of observer and illumination variation.

Here, image-space relighting methods cover an interesting middle ground. Even the reference data from Figure 6.9, comprising of 14116 high dynamic range images, consumes, when stored with the original 1360×1024 pixels per image, less than 222 Gigabytes – it can be stored on an inexpensive consumer-class storage device. While image-space relighting sacrifices the ability to generate novel views, we have seen that it has relevant applications in computer-assisted lighting design and in the context of free-form illumination photography.

As it shares many of the challenges in acquisition and rendering with full 8D reflectance fields, it permits to research them with a manageable investment in data size and acquisition time. Therefore, we will investigate image-space relighting exclusively in the remainder of the thesis, advancing the state of the art with respect to acquisition, interpolation, and rendering.

4 A Measurement Setup with Programmable Light Source Extent

4.1 Introduction

After reviewing the basic terminology and the related work on relightable object representations, we can now start with our contributions to the subject of image-space relighting. In this chapter, we will describe a practical measurement setup for the recording of $4D \rightarrow 2$ reflectance fields in directional illumination conditions.

It needs to have the following properties:

1. **Programmable distribution of directions:** We require a setup which is able to create incident illumination for an arbitrary specification of input illumination directions with high precision.
2. **Programmable light source extent:** The perfect directional illumination corresponds to a Dirac distribution of non-zero illumination directions; a real-world realization will always have a certain extent of non-zero directions. To investigate reflectance field recording in a sampling theory context, we also need to control this extent.
3. **Overlapping illumination conditions:** In order to experiment with different sampling kernels, we need to be able to create neighboring illumination conditions with sufficient overlap.
4. **Unsupervised recording:** A recording of a pre-defined sequence should not involve any human supervision after the initial scene setup, as we will need to run long sequences with different sample distributions.
5. **Reproducible results:** Within the limits of physical object persistence, we would like to be able to repeat measurements, for instance with new illumination directions, and maintain comparability to previous results.

We can afford the following limitations:

Upper hemisphere only: We will put the objects we measure on opaque surfaces, so illumination from below is not required. Nevertheless, our observations carry forth to the entire sphere of possible incident directions in a straight-forward way.

Long acquisition times: We will record only static scenes, and will not collect a large database of hundreds of scenes. As the setup runs automatically, several minutes recording time for coarsely sampled fields (for instance, with $n = 256$ input HDR images) or even hours for reference data sets ($n \geq 10000$ input HDR images) can be dealt with.

These requirements motivate a robotic setup. However, physically moving a light source in a distance to a scene requires a large robotic installation which is highly expensive. We therefore chose to create our illumination indirectly, by illuminating an opaque tent of cloth from the inside with off-the-shelf projection systems with programmable direction and aperture (see also Mohan et al. (2005)). Thus, the possibility to overlap neighboring illumination distributions comes for free.

A faster construction may be possible by using statically mounted video projectors that illuminate the tent from within with varying patterns, however, such an approach would either require more projectors to cover the entirety of the incident illumination domain, or be much less efficient with respect of achievable light intensity for a given light source power.

4.2 Construction

Figure 4.1 displays a photograph of our measurement setup. The illumination comes from so-called “moving head” controllable spotlights. By using two spot light emitters our setup covers a large part of the hemisphere. We put the object that we want to measure in the center of the room, and illuminate it indirectly by light which bounces off the enclosing tent, creating distant light conditions. As light sources, we use two JB Lighting Varyscan P3 moving head spotlights containing 250 Watts HTI light sources, at a distance of 65cm from the object.

Our setup is constructed inside a tent of dark, uniform cloth, which covers an area of $4\text{ m} \times 3.8\text{ m}$ and has a height of 1.9m. The floor is covered by dark carpet. We used dark surfaces, as opposed to the white walls used in previous methods (Mohan et al., 2005), in order to reduce the power of higher order bounce light relative to the power of the first order bounce light used for measurement.

The spotlights can be controlled using the DMX protocol. They can be rotated over 430 degrees pan and 270 degrees tilt, have a programmable iris, shutter, focus and frost effect, which allows us to control the light position in the room, its extent, and the smoothness of the fall-off of the spotlight. The diameter of the emitted light cone can be varied from 4.3° at maximally closed aperture up to 18° at fully open aperture. Due to the fact that the spotlight is located next to the object, the size of the incoming light at the object varies between 3.1° and 4.3° for closed aperture, and between 13° and 18° for open aperture. Figure 4.2 shows some of the light conditions that can be achieved. In our experiments concerning sampling, we used two settings for the light sources: the



Figure 4.1: Photograph of our measurement setup. Spot light projectors on the floor shine on a tent made from black cloth, and the bounce light illuminates the scene, which is arranged on top of the boxes. A tripod-mounted camera records the scene, and is protected against lens flares by occluders. In order to show the inside of the tent, we have added additional illumination (lamps outside the picture) which is not present during the acquisition.

wide setting, with the iris fully open and the frost effect filter applied, and the setting where the iris was 30% opened and the frost effect switched off.

For image capture, we run HDR sequences (Robertson et al., 2003) on Jenoptik CFcool or C14plus cameras. We choose the length of the HDR sequences depending on the aperture of the spotlights. Typically, we use four exposures with durations of 20, 100, 500 and 1000 ms, respectively, if the aperture is fully open and the frost filter is enabled. Choosing the exposure times as multiples of 20 ms cancels out 50 Hz flickering in the light sources of the spotlights.

4.3 Calibration

In order to apply the setup to reflectance field measurements, we need to determine the relationship between projector pan/tilt settings and light directions, as seen from the

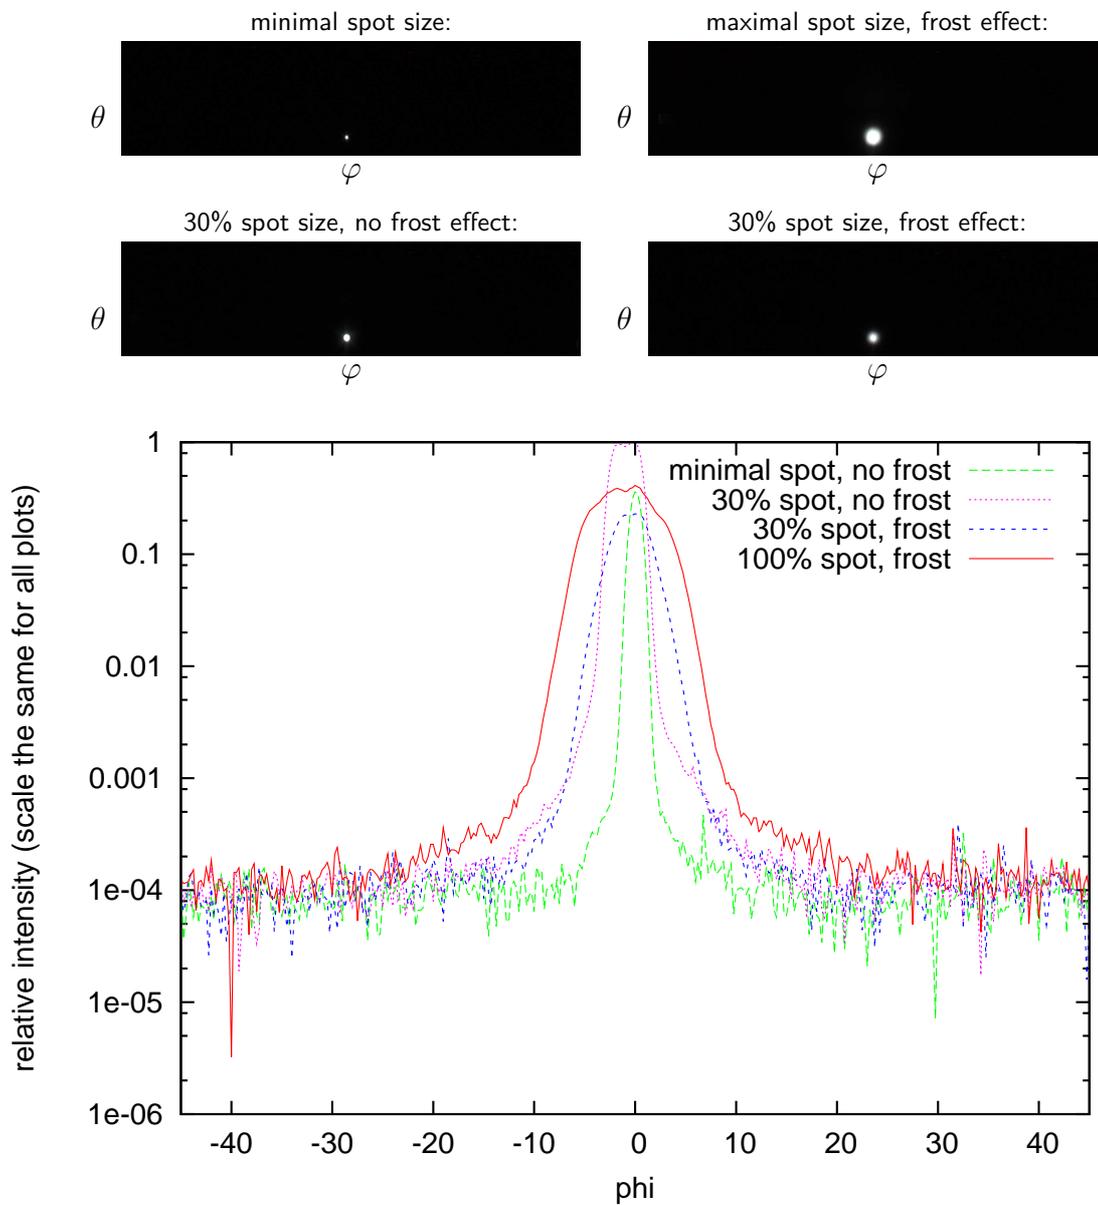


Figure 4.2: Illustration of achievable light distributions. On the top, environment maps in polar parameterization of the lighting situations are shown. Below, the intensity distribution along the centrally lit θ angle is plotted for varying ϕ values. Note the high contrast and the smooth fall-off.

object. Also, as our tent is non-spherical, the total brightness of the bounce light arriving at the object depends on the light direction (it is darker in the corners of the tent), and needs to be taken into account. The influence of pan/tilt on the exact size and shape of the light distribution could also be measured, but in our setup we may assume that it has no significant effect.

For the calibration, we place a camera with a fish eye lens at the center of the measurement tent facing the upper hemisphere. Methods for its precise geometric (Kannala and Brandt, 2004) and vignetting (Goldman and Chen, 2005) calibration are available; for our experiments, we only performed a manual calibration for the most important parameters for the mapping of environment map directions to camera pixels. As long as the same lens is used for the acquisition for the target environment maps, most geometric and vignetting effects will cancel out.

The mapping between pan and tilt positions of the projectors to (θ, φ) is obtained by rotating the projector in steps of 10 degrees and recording images of the bounce light from the black cloth. In the captured images, we can compute the average direction of all pixels which have a brightness above a noise threshold, and thus obtain an estimate for the incoming light direction which corresponds to the selected pan/tilt settings. Linear interpolation between these values makes it possible to estimate for any given (θ, φ) the corresponding pan/tilt configuration for each of the spotlights.

The non-spherical structure of our measurement environment makes it necessary to compensate for different brightnesses of the bounce light. Photometric calibration is performed by establishing the camera's response curve and measuring the white point. The observed brightness in the calibration images (sum of pixel intensities that are above the noise threshold) gives a relative brightness scaling factor for each direction.

During calibration, we also infer which of the two spotlights is to be preferred for each (θ, φ) , based on the spacing of (θ, φ) samples in the calibration images. Due to the field of view of the fish eye lens, we can only calibrate the upper hemisphere, so we only measure there. However, for most scenes this is the most significant region.

4.4 Rendering Scheme

Rendering the reflectance field R for an arbitrary environment map may require evaluating $R(\theta, \varphi)$ for light angles (θ, φ) which have never been observed. For this interpolation problem, Masselus et al. (2004) have investigated several methods. We will now discuss those properties that are required in our setup.

Let $(R_i)_{1 \leq i \leq n}$ be the images of the scene observed for distant light at angles (θ_i, φ_i) . Let (θ, φ) be the angle we want to interpolate for. While we cannot guarantee to predict correct values of the reflectance field, we can aim for physically plausible interpolation for which we know that a light configuration exists which creates it. This is the case

for linear combinations of the input images with non-negative coefficients $l_i(\theta, \varphi)$ chosen depending on that angle; so, R will be approximated as

$$R(\theta, \varphi, x, y) = \sum_i l_i(\theta, \varphi) R_i(x, y). \quad (4.1)$$

Equation 2.21 can now be rewritten as

$$\begin{aligned} I(x, y) &= \int_{\theta} \int_{\varphi} L(\theta, \varphi) \cdot R(\theta, \varphi, x, y) d\varphi \cos \theta d\theta \\ &= \int_{\theta} \int_{\varphi} L(\theta, \varphi) \sum_i l_i(\theta, \varphi) R_i(x, y) d\varphi \sin \theta d\theta \\ &= \sum_i \underbrace{\int_0^{\pi} \int_0^{2\pi} L(\theta, \varphi) \cdot l_i(\theta, \varphi) d\varphi \sin \theta d\theta}_{=:k_{i,L}} R_i(x, y) \end{aligned} \quad (4.2)$$

As the $l_i(\theta, \varphi)$ only depend on the structure of the R_i , and not on the pixel position (x, y) , we can pre-compute them. At rendering time, we can then compute the $k_{i,L}$ for the given environment map L , and efficiently compute all pixels $I(x, y)$ in the target image as linear combination of the input images (with pixel coordinate-independent linear coefficients k_i).

In addition, we require the following properties of the interpolation weights l_i :

- **local support:** $l_i(\theta, \varphi)$ should be non-zero only if (θ, φ) is close to the measurement light direction (θ_i, φ_i) of the corresponding image.
- **partition of unity:** for each direction (θ, φ) , the weights should add up to one:

$$\sum_i l_i(\theta, \varphi) = 1 \quad (4.3)$$

in order to conserve energy.

- **interpolation:** if the position of the novel direction (θ, φ) coincides with a observed direction in (θ_i, φ_i) , the corresponding weight $l_i(\theta, \varphi)$ should be one.
- **continuity:** the weights should change C^0 -smooth over changes in directions (θ, φ) .
- **generality:** the approximation scheme should work with arbitrary, possibly non-regular, distributions of sample directions.

We have chosen a linear interpolation on barycentric coordinates (Masselus et al., 2004) defined on solid angle relations on the sphere throughout the following discussion. While

smoother reconstructions are possible, this type of reconstruction can quickly and efficiently be evaluated, and is sufficient to achieve the quality we show in our results. A detailed description of the computation of the interpolation weights is given below.

The light directions (θ_i, φ_i) used to measure the reflectance field can be interpreted as normalized vectors $d_i \in \mathbb{R}^3$ on the unit sphere, forming a set D . Let $d \in \mathbb{R}^3$ be the unit vector in the direction (θ, φ) we want to interpolate the reflectance field R for.

In the following, we will explain how to select indices j, k, l and interpolation weights $l_j(\theta, \varphi), l_k(\theta, \varphi), l_l(\theta, \varphi)$, such that we can approximate

$$R(x, y, \theta, \varphi) \approx l_j(\theta, \varphi) \cdot R_j(x, y) + l_k(\theta, \varphi) \cdot R_k(x, y) + l_l(\theta, \varphi) \cdot R_l(x, y). \quad (4.4)$$

In a preprocessing step, we compute a triangulation T of the hemispherical hull of D . We do that by computing a Delaunay tetrahedralization of $D \cup \{g\}$, where g is a guardian element below the hemisphere, at the $(\theta, \varphi) = (\pi, 0)$ pole. Each tetrahedron which contains a point at infinity is on the outside of D , and contains a triangle t which consists only of points on $D \cup \{g\}$. If and only if t does not contain g , t lies entirely on the upper hemisphere, and thus defined to be part of T . All vertices d which are connected to g by an edge in the tetrahedralization are marked as boundary vertices for further use.

Now, we can intersect the ray defined by d and the origin with all triangles in T . Two cases are possible:

1. the ray intersects with the triangle created by (d_j, d_k, d_l) . Let

$$s : \mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}_0^+ \quad (4.5)$$

now be the function which maps three points on the unit sphere to the non-negative solid angle covered by the inside spherical triangle defined by the points.

We can then define

$$l_j(\theta, \varphi) := \frac{s(d, d_k, d_l)}{s(d_j, d_k, d_l)}, \quad l_k(\theta, \varphi) := \frac{s(d_j, d, d_l)}{s(d_j, d_k, d_l)}, \quad l_l(\theta, \varphi) := \frac{s(d_j, d_k, d)}{s(d_j, d_k, d_l)} \quad (4.6)$$

yielding a interpolation with barycentric coordinates on spherical triangle solid angles.

2. The ray does not intersect with any triangle. This is possible if d is close to the border of the hemisphere. In that case, we select j and k so that d_j and d_k are boundary vertices so that

$$\varphi_j \leq \varphi < \varphi_k \wedge \exists m : \varphi_j < \varphi_m < \varphi_k \quad (4.7)$$

By duplicating the vertex d_i , where $\varphi_i \leq \varphi_m \forall m$, and assigning a value of $\varphi := \varphi_i + 2\pi$ to it, we can handle the wrap-around condition correctly.

Now, we can define

$$l_j(\theta, \varphi) := \frac{\varphi_k - \varphi}{\varphi_k - \varphi_j}, \quad l_k(\theta, \varphi) := \frac{\varphi - \varphi_j}{\varphi_k - \varphi_j}, \quad l_l(\theta, \varphi) := 0 \quad (4.8)$$

as arc interpolation weights.

5 Sampling Reflectance Fields

After defining the physical components of our reflectance field acquisition method in the previous chapter, we can now focus on the software side. In this chapter, we discuss how to define control patterns for the measurement setup, and compare results for fixed-pattern sampling with extended light sources (Section 5.2) and an adaptively sampling scheme (Section 5.3). Then, we combine narrow and extended illuminations (Section 5.4), and discuss the sampling properties of the found patterns in (Section 5.5).

5.1 Introduction

While directional sampling-based reflectance models discussed in Section 3.4.2 (p. 29) are robust and achieve a high quality, they require large numbers of input images even for low-frequency target illuminations. This is a result of two problems: gaps in the coverage of the incident light domain, and aliasing artefacts due to undersampling.

In this chapter, we will address both by the way our input light sources are constructed. The coverage problem can be mitigated by using extended light sources, the aliasing problem can be reduced by using light sources with a smooth fall-off. This is actually a consequence of the results of sampling theory, which is centered around the sampling theorem by Shannon (1949) which states that a function, which does not contain frequencies over a frequency limit, can be completely reconstructed if sampled at twice the frequency of the limit. If less samples are available, the reconstruction of the function induces aliasing artefacts, which can be mitigated by prefiltering.

One way to achieve such a prefiltering can be deduced from Equation 2.21. It resembles a convolution, which corresponds to a multiplication in Fourier space. Accordingly, high frequency content of the reflectance field can physically be suppressed using extended light conditions L_i with smooth fall-offs for its measurement. Thus, while we cannot influence the frequency content of the reflectance field itself, prefiltering can be established through the choice of the illuminations L_i . By this, we limit the maximum frequency in each illumination pattern, and thus observe an approximation of R which is low-pass filtered in the (θ, φ) domain (see Ramamoorthi and Hanrahan (2001)).

Aliasing artefacts predominantly occur at shadow boundaries (Figure 5.1) or on shiny objects where the original light sources can be seen as mirror reflections (Figure 5.2). We show how they can be avoided by using extended light sources with smooth fall-offs as input illuminants (Section 5.2) at the cost of fidelity in high frequency illumination.

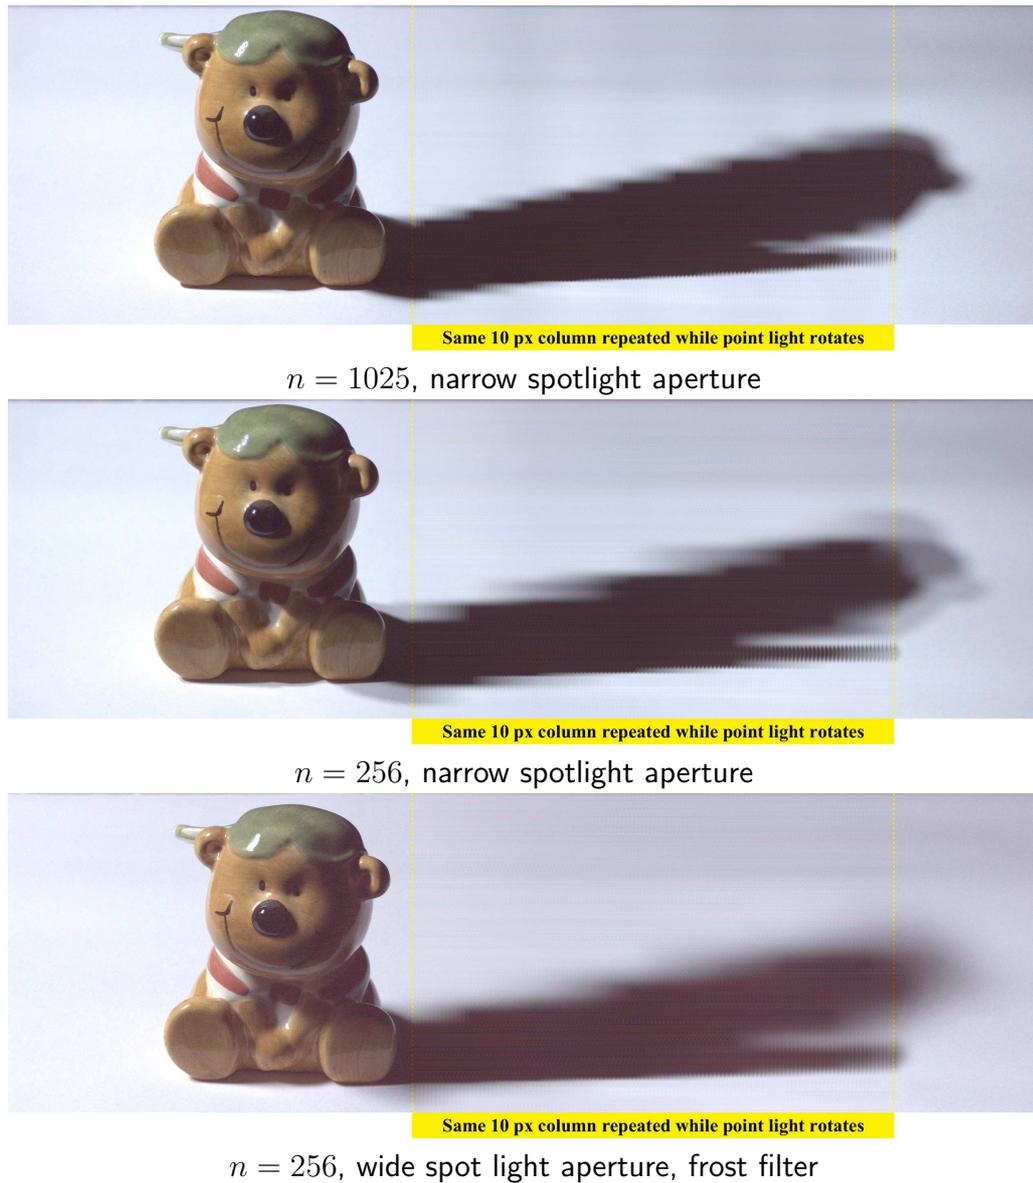


Figure 5.1: Results of relighting in point light conditions. Between the two yellow lines, the same vertical strip of 10 pixels width is rendered repeatedly, while the point light is rotated around the up axis over 60 degrees with 0.5 degrees separation. Note the differences between the images at the shadow boundary. For the narrow aperture, $n = 1025$ images deliver an almost smooth boundary, for $n = 256$ images, the boundary is strongly jagged, and several overlapping shadow regions can be seen to the right of the image. For the same number of input images, and for the same incoming light directions, but for full aperture and activated frost filter, the jaggedness is almost not perceivable, and although the shadow to the right of the image is noticeably blurred, it shows a continuous penumbra zone without ghosting artefacts.

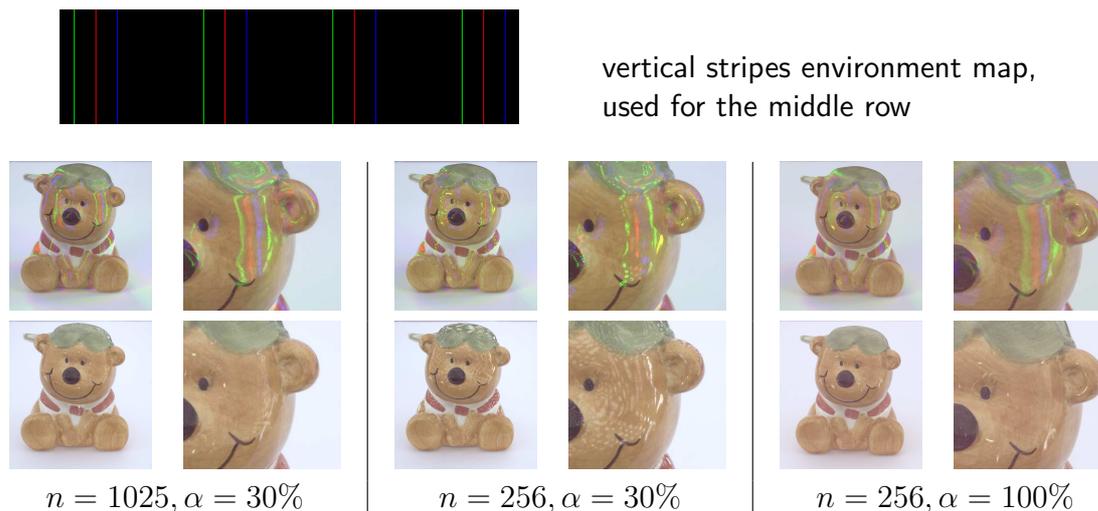


Figure 5.2: Regular sampling results for a ceramic teddy bear, for varying numbers of light sampling positions n and varying relative aperture size α . For full aperture, the frost filter was switched on, off otherwise. The middle row of images has been rendered with the stripe illumination on display in the first row, the bottom row has been rendered in ambient white light. The brightness of the result images has been rescaled to match each other.

An alternative way to reduce the effort of measurements, while retaining high frequency reflections, is to distribute the sample directions in an irregular pattern that is adapted to the sensitivity of the reflectance function to changes in incident light direction. We present an adaptive algorithm for this in Section 5.3, and perform a quantitative and qualitative analysis of its benefits in Section 5.3.2.

Our algorithm starts with a coarse set S_0 of regularly distributed illumination directions (θ_i, φ_i) on the upper hemisphere. For each direction, an image R_i is recorded. Then, in each iteration j , we analyze the observed slices of the reflectance field and augment S_j by several new sample directions in those directions where the reflectance field is not smooth. The iteration terminates at a user-specified fixed number of samples.

The adaptive sampling technique may still produce the artefacts that we mentioned above if the observed object has small, but shiny surfaces that contribute little to the overall error. Therefore, we propose to combine adaptive sampling at narrow-aperture lighting with sparsely sampled extended light source conditions (Section 5.4). This technique creates a reflectance field with two levels of physical prefiltering. Sampling theory motivates a discussion of the point distributions generated by the adaptive sampling algorithm that is provided in Section 5.5.

5.1.1 Related Work

Chai et al. (2000) have discussed optimal reconstruction kernels for light fields of diffuse

scenes, based on the depth complexity. Reflectance fields, however, can contain arbitrary high frequencies. In order to suppress aliasing artefacts, prefiltering at acquisition time is required, as we demonstrate in Figure 5.2.

Adaptive sampling of the incident light direction in the context of measuring individual homogeneous BRDF has been proposed by Matusik et al. (2003b) based on a principle component analysis of a large collection of captured BRDFs.

Lensch et al. (2003b) computed the next best viewing and lighting directions for capturing spatially varying BRDFs of 3D objects by analyzing the covariance matrices of the parameters obtained by fitting an analytic BRDF model to the current set of measurements of each surface point.

Peers and Dutré (2003) have used adaptively refined wavelets. Instead of capturing a single basis image, Sen et al. (2005) have emitted several illumination patterns simultaneously, parallelizing the acquisition.

Our adaptive scheme is inspired by the work by Schirmacher et al. (1999b) who proposed an adaptive light field acquisition method, that adaptively subdivides a mesh of light view positions according to an image-space error metric.

5.2 Fixed Pattern Sampling

In our setup, we can measure reflectance fields at variable densities and illumination shapes, which we will refer to as sampling patterns in the following. Unlike previous approaches with light sources mounted on fixed rigs, we can now choose the density both in elevation angle θ and in azimuth φ freely. Unlike hand-held approaches, we can precisely distribute and repeat the sampling.

In order to create patterns at different densities of sampling directions (θ, φ) and at an approximately uniform spacing, our algorithm places virtual charges on pseudo-random positions on a sphere, and iteratively moves them until an equilibrium is reached. The directions on the upper hemisphere are used as sampling positions. This randomized approach generates an approximately regular lattice with a hexagonal pattern (see Figure 5.3; the black points from iteration 0 have been generated with this algorithm).

5.2.1 Results of Fixed Pattern Sampling

Figure 5.2 shows some results of fixed sampling with different numbers of input images and spotlight apertures. For the narrow aperture – which corresponds to a light cone of about 8 degrees – we need about 1025 images until the highlights on the ceramic figure become continuous for the vertical stripes case. While few isolated highlights are still visible in the ambient rendering, the general appearance is smooth. If we reduce the number of captured images to 256, strong artefacts occur on the glossy surface. The highlights are not continuous anymore, and in some places yellow highlights show up

(as the available resolution is not sufficient to separate the contributions from a red and green colored input image), some other highlights are missing altogether.

If the aperture is fully opened and the frost filter activated, the highlights are blurred out and high-frequency detail is lost. However, for only $n = 256$ input images, the aforementioned artefacts vanish.

For animations where the environment map changes continuously, the large aperture and smooth falloff induced by the frost filter are highly advantageous, as it reduces temporal aliasing. We demonstrate this by rendering an image sequence with a rotating point light source (Figure 5.1). In the top image, for $n = 1025$ light samples, we perceive a slight jaggedness in the shadow contour. The reason for this becomes visible when we decrease the number of samples to $n = 256$, as in the center image: temporal aliasing artefacts can be seen, the shadow jumps, or its contours fade in and out. This is due to the low number of input light conditions, combined with the narrow angular distribution, causing aliasing problems in the incident light domain.

The bottom image shows again a rendering for only $n = 256$ input images, with wide spotlight aperture and activated frost filter. Although the sampling density is the same as in the center image, the renderings are, though blurry, almost free from aliasing artefacts.

On a side note, Figure 5.1 also demonstrates that we can resolve the high contrast between lit areas and the cast shadow, which is achieved as a result of using dark cloth as reflective material in our tent which reduces the intensity of scattered, ambient light.

5.3 Adaptive Sampling

The previous results show that for environments with a low-frequency variation in the angular domain, it is sufficient to work with few input images and to use extended light sources. If the incident illumination contains high frequencies, however, these will be blurred, and detail cannot be preserved. This problem can be alleviated by using a more narrow light distribution, but if the hemisphere of input illumination is still to be covered fully, the sample density on the hemisphere needs to be increased by the square of the factor by which the spot radius is reduced. Using a regular sampling pattern, a significant number of the samples might be wasted, though, since they might excessively sample smooth (diffuse) regions of the reflectance field.

We therefore propose an adaptive sampling scheme, which creates a subset chain of sets of incident light directions $S_o \subset \dots \subset S_j \subset \dots \subset S_n$ with elements from $[0, \frac{\pi}{2}] \times (0, 2\pi)$. After the corresponding images for one set S_j are taken, it is analyzed and new directions are inserted to create S_{j+1} . For now, we work with a fixed, narrow spot size. In Section 5.4, we will consider combinations of different spot sizes.

5.3.1 Algorithm

In this section, we explain in more detail how the sets $S_0 \subset \dots \subset S_j \subset \dots \subset S_n$ of sampling directions are constructed. The first set S_0 contains samples which are regularly distributed on the upper hemisphere as before.

Based on the images R_i for light directions $(\theta_i, \varphi_i) \in S_j$, the set S_{j+1} is constructed by including additional sampling directions, for which we will acquire the respective images in the following iteration. The detailed algorithm is as follows:

1. For each direction (θ_i, φ_i) , compute the one-ring N_i

$$N_i := \{k \in \mathbb{N} : (\theta_i, \varphi_i) \text{ and } (\theta_k, \varphi_k) \text{ share an edge in the triangulation} \} \quad (5.1)$$

as the neighborhood of directions around (θ_i, φ_i) in the triangulation used for the interpolation (see Section 4.4).

2. For each direction (θ_i, φ_i) , evaluate an error metric describing the smoothness of the reflectance field in the neighborhood of (θ_i, φ_i) .

A comparison of R_i to the average of the images $R_k : k \in N_i$ would be a straightforward choice. However, while developing our algorithm, we have observed such a metric causing endless subdivisions at one-sided discontinuities, e.g. where a shadow or caustic suddenly appears or disappears due to occlusions in the scene, while neglecting high-frequency effects which are changing rapidly along all directions, for instance moving highlights. This is in contrast to our perceptual preferences, as human observers can easily spot jumping highlights in animations or missing highlights in mirroring surfaces, but are not as sensitive to the exact location of one-sided discontinuities.

We have therefore chosen a metric which yields low error values if R_i is similar to *some* of the $R_k : k \in N_i$, but results in a high error if R_i is different to *each* of the R_k .

In order to achieve this, we interpret the images as vectors and compute the projection error of R_i into $\langle R_k \rangle_{k \in N_i}$, that is, we chose linear coefficients μ_k such that

$$\left\| R_i - \sum_{k \in N_i} \mu_k R_k \right\|^2 \quad (5.2)$$

is minimized, and interpret the value of the entire term as local error. While this heuristic is not perfectly adapted for discontinuities, it produces good results in our experiments, and its value can be very quickly computed using a Singular Value Decomposition on the matrix composed of the R_k .

3. For each triangle in the triangulation, assign the sum of the corners' error values as error value.

4. For each of the m triangles T_i with the largest error, insert a new direction (θ_i, φ_i) at T_i 's circumcenter unless it is close to a direction already inserted in the current iteration, in which case the next most important direction is inserted.

This results in a locally even distribution (Lloyd, 1982): the newly inserted direction is equally distant to points already present in the local neighborhood, as the circumcenter of T_i is on the intersection of the perpendicular bisectors of T_i .

In areas where the local point density suddenly changes, the (θ_i, φ_i) may even be outside T_i ; however, inserting a sampling direction there improves the distribution of the sampling directions, and further refinements inside T_i will take place in future iterations if required.

S_{j+1} is then defined as the union of the new sampling directions and S_j . This process is iterated until S_n has been constructed, and the corresponding images have been taken. In each iteration, we add $m = 16$ directions, which is more efficient than adding one direction $m = 1$ at a time. We skip all triangles where the radius of the circumcircle would be below a threshold of 2 degrees. This imposes an upper limit on sampling density, and makes sure we do not sample denser than the reference recording we will record later.

5.3.2 Results of the Adaptive Sampling Algorithm

In order to evaluate our sampling scheme, we have conducted a set of experiments. We have measured reflectance fields of three objects, which we will discuss further:

1. a bottle containing a colored liquid, which, due to its cylindrical shape and glossy surface material, is expected to have an anisotropic reflectance field, i.e. it should demonstrate the benefits of adaptive sampling,
2. a graphics card, which has complex self-shadowing and is covered by a diverse set of different materials, and
3. a ceramic figure, which serves as a difficult example: on its surface, we find practically every surface normal direction, as opposed to the cylindrical bottle. Moreover, its surface is shiny, so each direction in the illuminating environment is almost equally important.

For all these objects, we measured reflectance fields once using fixed sampling configurations, and once applying our adaptive sampling scheme. In the adaptive case, we initialized the algorithm with 256 regularly distributed samples, and added 16 light situations in each of 48 iterations, yielding 1024 total images. For the bottle data set, we added sample points until the data set was densely sampled and only less than 16 images could have been inserted, which happened after iteration 198. Figure 5.3 illustrates the insertion sequence for the different data sets, and shows the objects rendered in ambient

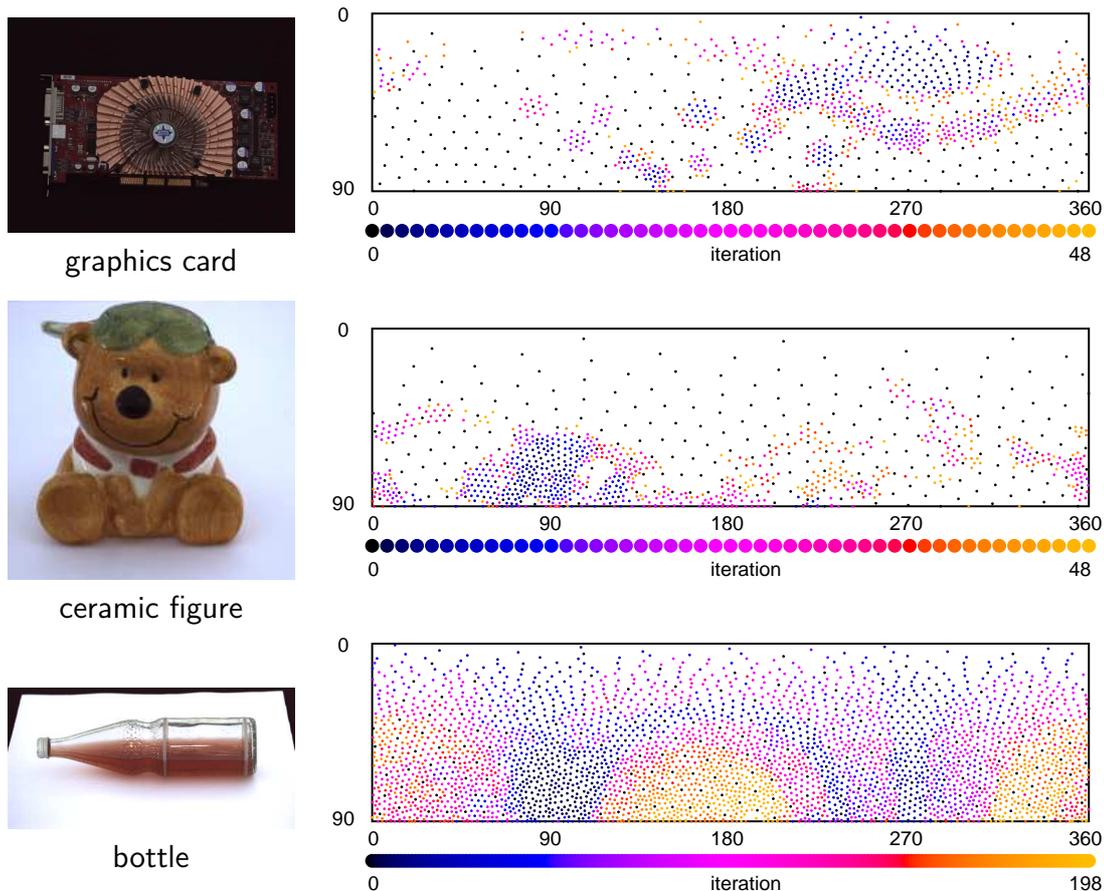


Figure 5.3: The plots illustrate the sequence in which light directions have been generated. For initialization, iteration 0 consisted of 256 uniformly distributed samples (drawn in black), later samples grow in regions of interest. The renderings in the left column are for $n = 1024$ total images (after 48 iterations), rendered in ambient light.

light after 48 iterations. As can be seen, clusters of sampling directions occur, indicating regions of directions where the appearance changes drastically. These correspond to light directions which trigger extended highlights to appear or disappear, and to shadow areas that move over large regions in image space as a result of small variation of the incident light direction.

Quantitative Evaluation

For each of the objects mentioned above, we have recorded a ground truth measurement, consisting of images in 10 000 light directions densely distributed over the upper hemisphere on a regular grid, which is visualized as the positions of dots in Figure 5.4.

In order to obtain a numerical measure for the quality of a given reflectance field, we

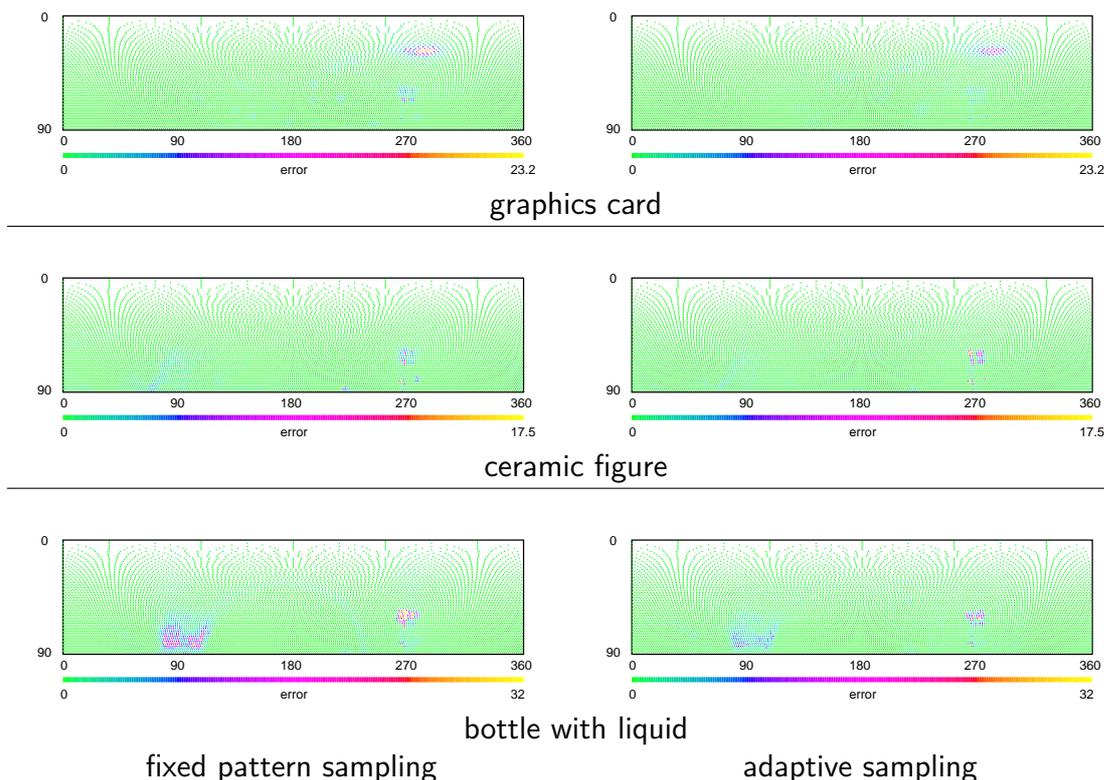


Figure 5.4: Spatially plotted error values. Each dot represents the light direction for which a ground truth image was taken, its color gives the sum of squared pixel differences between ground truth and the reconstructed image for this direction. The adaptive sampling column displays values for a reflectance field with 256 initial, regularly distributed samples, into which groups of 16 samples each were inserted over 48 iterations, yielding a total of 1024 images; the fixed pattern sampling column shows a regularly sampled reflectance field of 1025 input images.

render an image for each of the captured ground truth directions, and sum up the squared differences between the ground truth pixel values and the rendered pixel values. The color-coded differences are plotted in Figure 5.4 at the corresponding (θ, φ) locations. As the aspect ratio of the objects does not match the aspect ratio of the camera, we have cropped the input images to the same region of interest which is used for the renderings in Figure 5.3. In order to produce the plots efficiently, we have downsampled the images using either a 3×3 or a 4×4 box filter to smaller resolutions: 453×224 pixels for the bottle data set, 453×341 pixels for the graphics card, and 200×196 for the ceramic figure. This restricts the precision of the least squares error metric, but it also increases the robustness against pixel noise generated by the camera. As a result, the absolute numbers of SSD error are only comparable between the same data set.

In the error plots of Figure 5.4, we can observe that for the ceramic figure, the error distribution and the amount of error is comparable both for the adaptive sampling scheme and

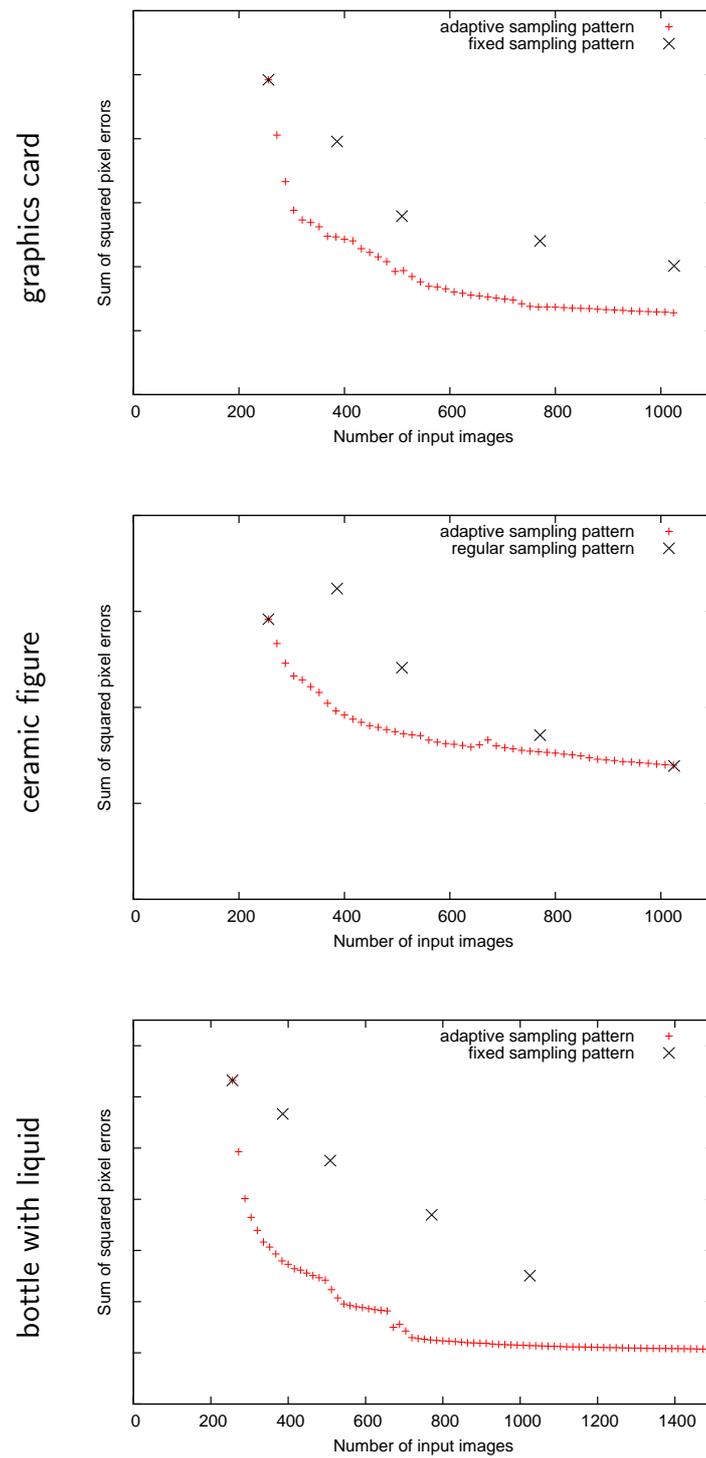


Figure 5.5: Plots of the global error value as sum of squared pixel differences between ground truth and adaptive / fixed pattern sampling.

a fixed sampling scheme with similar amounts of input images. For the other data sets, however, there is considerably less residual error in the adaptively sampled reflectance field than in the fixed pattern case. Compared to Figure 5.3, we can see that the remaining error zones are to be found in the (θ, φ) regions where samples have been placed densely up to the density limit, and that these samples have also been placed in early iterations.

In order to evaluate the performance of the adaptive sampling scheme as the iterations proceed, we have summed up the errors depicted in Figure 5.4 for each iteration, which gives a total numerical error value. As the sampling distribution of the ground truth plots is arranged so as to ensure an equal density over the hemisphere, this is a meaningful error value. Figure 5.5 shows these global error values as the iterative, adaptive algorithm proceeds, and as reference points the total error values for control measurements with a fixed sampling pattern.

Even for the ceramic figure as worst-case scenario, the adaptive sampling scheme stays competitive to fixed pattern sampling, and for low numbers of input images, it is consistently better. For the other two reflectance fields, which are more anisotropic, the advantages of adaptive sampling are even more clear. Over a wide range, the same amount of error can be achieved with about half or even less than half of the input images as compared to fixed pattern sampling.

Especially in the bottle data set, a plateau is reached at some point close to 700 images, or about 28 iterations. As can be seen in Figure 5.3, the light directions which have been sampled up to this point (marked in black and dark blue) coincide with the regions which expose the highest remaining error in Figure 5.4. Also, the final sampling density has already been reached there. Because of that, the error can not be reduced much more, as samples can only be inserted in areas where the error already was low.

Qualitative Evaluation

So far, we have only shown results on a quantitative evaluation of the different sampling schemes. Equally important, however, are qualitative evaluations, which demonstrate improvements in visible image quality. Therefore, we have created a synthetic high frequency environment map, which is displayed in the top right of Figure 5.6. It consists of cells that are filled by one solid color each, using the colors red, green and blue. Both the frequency of color changes and the purity of the resulting colors can be used as quality indicators in highlight regions of rendered images. Outside highlight regions, the integration over the product of the object BRDF and the environment map will create a diffuse object color since the illumination from a region of the environment will average to white.

Figure 5.6 shows renderings of the bottle scene in this environment map for different sample numbers both for fixed pattern and adaptive sampling. In the electronic version of this thesis, you can zoom in until the full resolution is visible; in order to make the

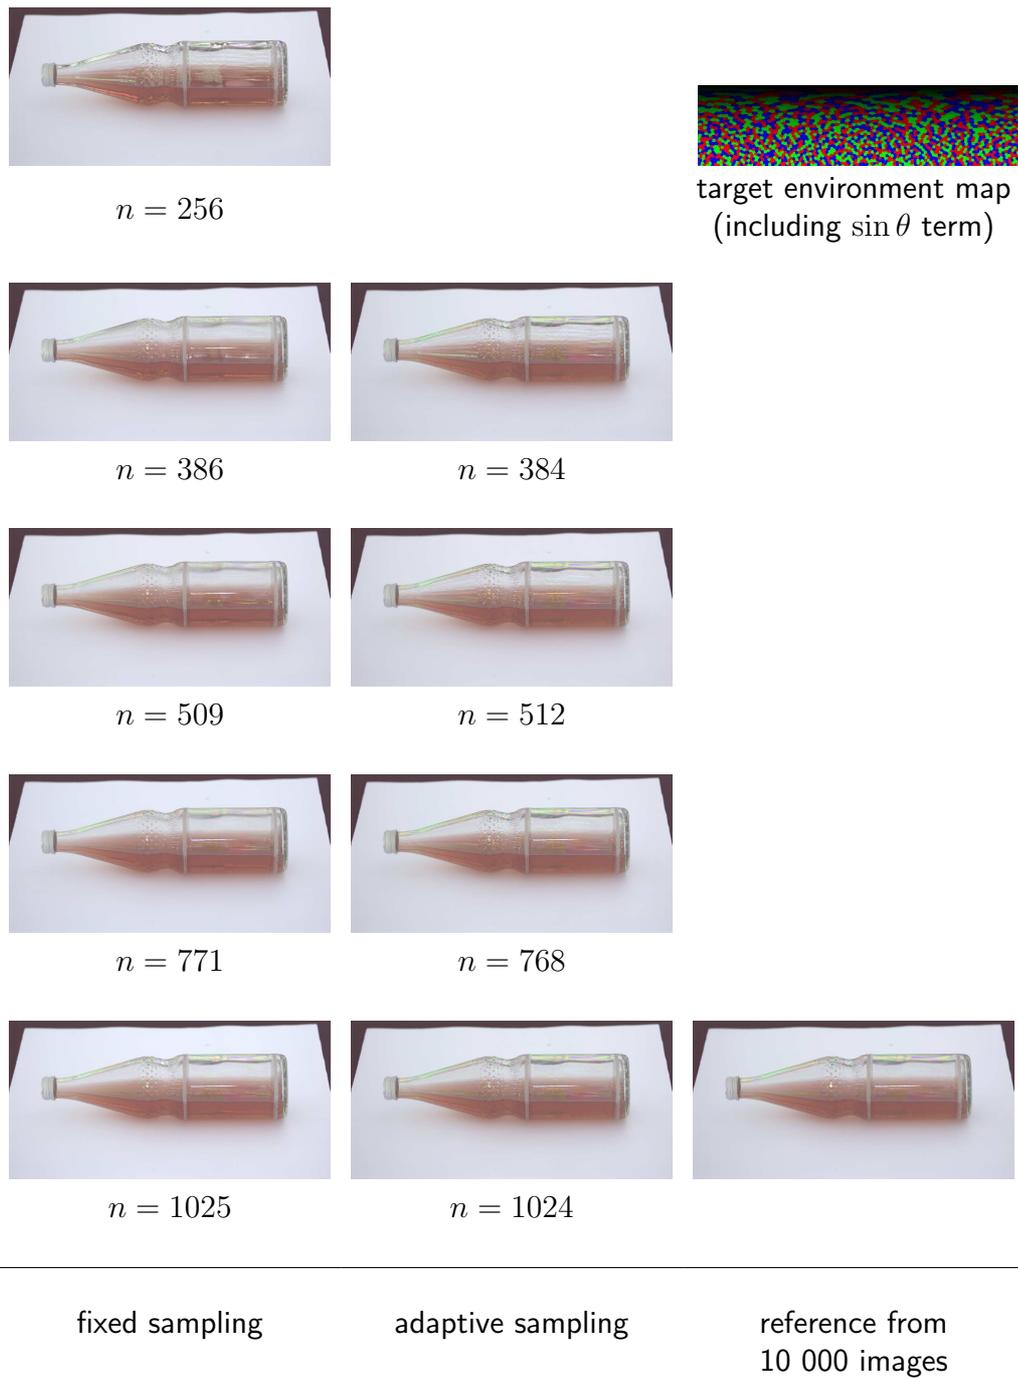


Figure 5.6: Renderings of the bottle scene in a high frequency environment map (top right), for different regular sampling distribution (left column) and different adaptive sampling distribution. The adaptive sampling was initialized by acquiring a regularly sampled reflectance field with $n = 256$ images. To the bottom right, a rendering created from the measurement with 10 000 input images is shown as reference. See Figures 5.7 and 5.8 for detail enlarge details.

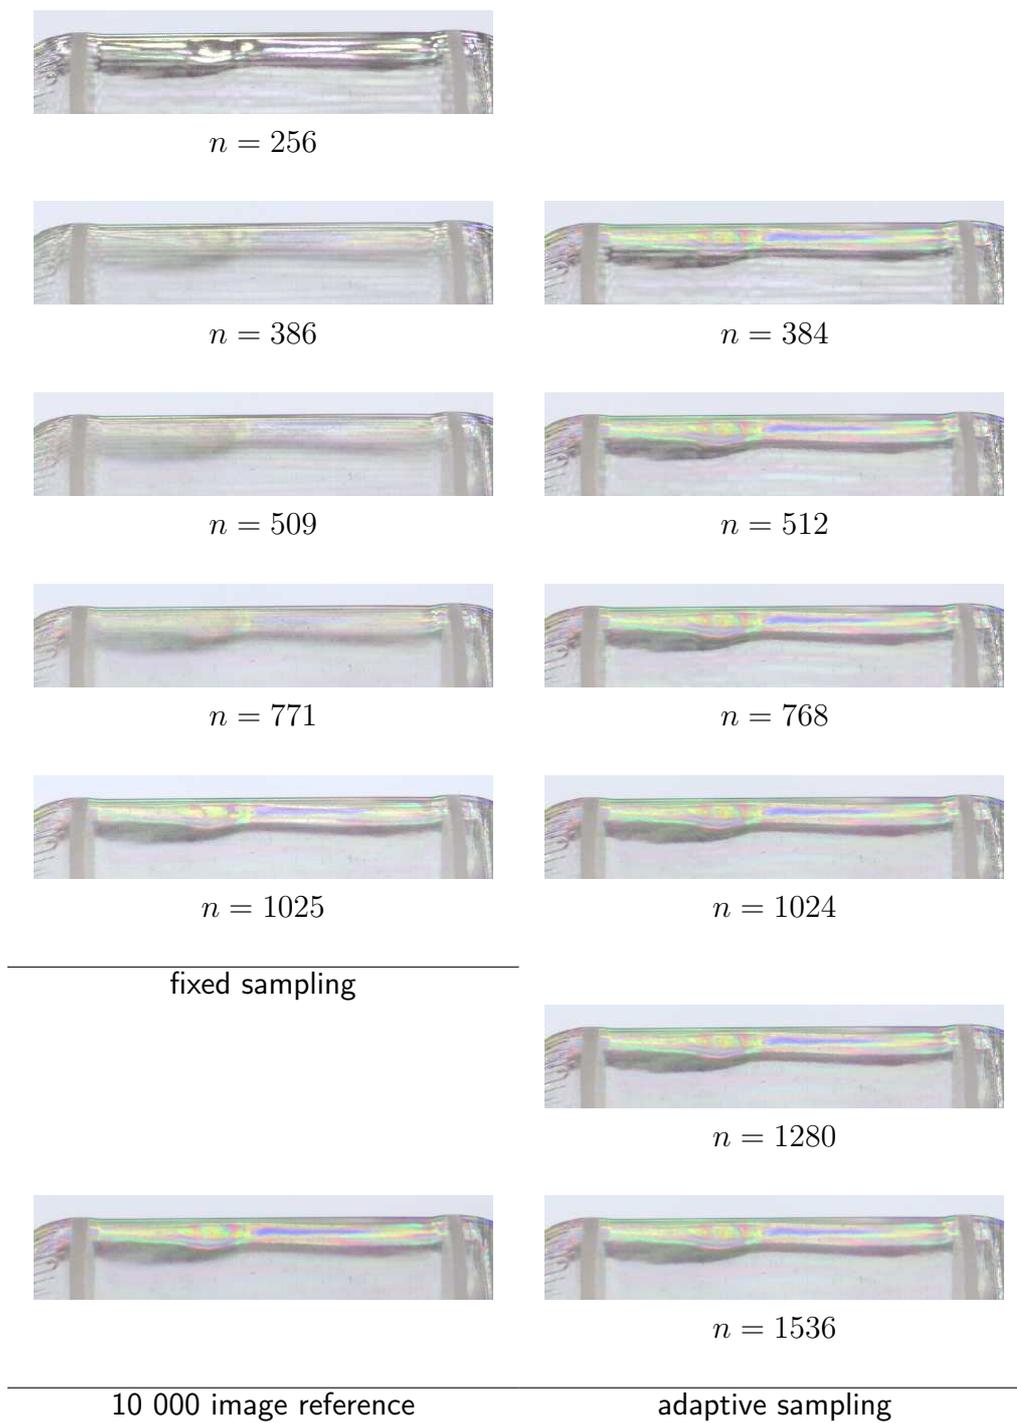


Figure 5.7: Detail pictures for Figure 5.6. Note how the detailed reconstruction of the highlights on the bottle surface is achieved much earlier by adaptive sampling.

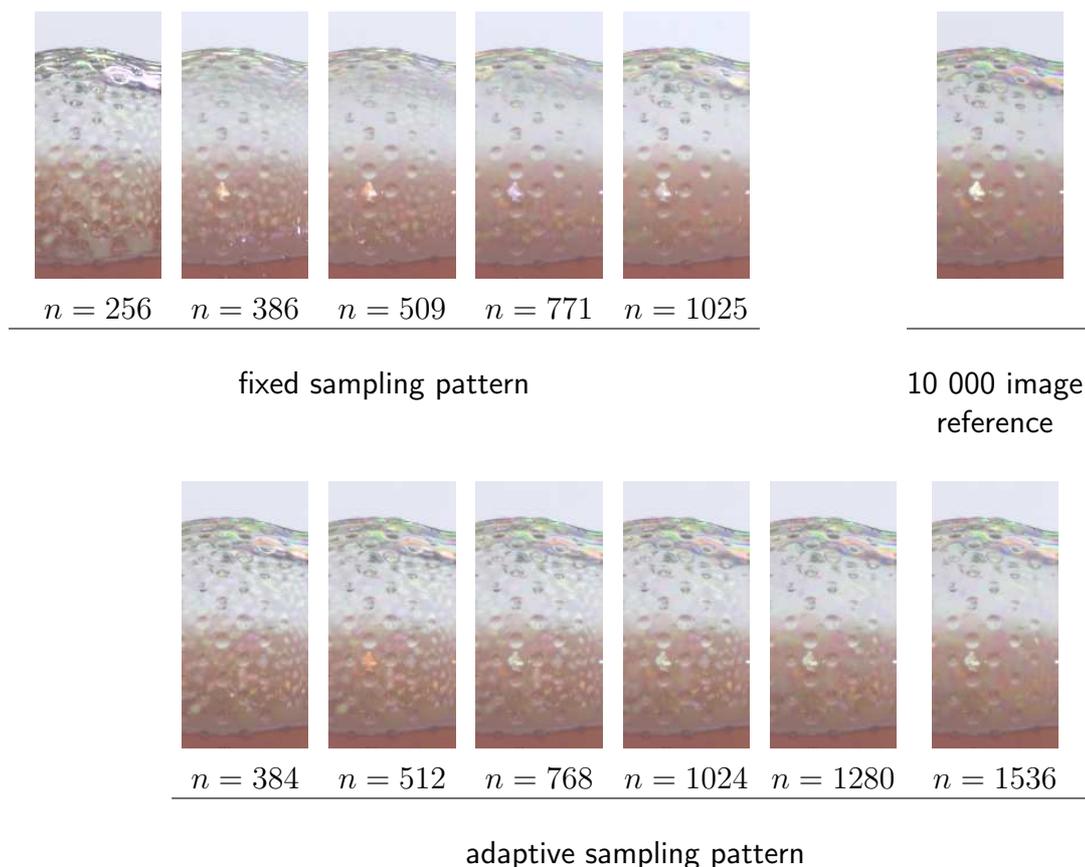


Figure 5.8: Detail pictures from Figure 5.6. Unlike the reflections on the top part of the images, which are captured nicely by adaptive sampling, other artefacts are more persistent: These are small colored spots distributed over a large area in the center of each image, best seen in front of the dark liquid (lower half of each image). In adaptive sampling, they disappear later than in fixed pattern sampling, because the algorithm tends to focus on more intense highlights first.

specific differences more visible, we provide Figure 5.7, which enlarges the highlight region on top of the bottle.

The reflectance field composed of $n = 256$ regularly distributed samples, which also serves as input data for the adaptive sampling algorithm, clearly lacks quality: the highlights are visibly separated, and where highlights are present, they show up as composed colors (yellow, magenta, cyan), which means that the same input image was used for two different input light conditions by the interpolation algorithm; this shows that the sampling resolution is much too low.

As the number of samples increases, the fixed pattern scheme closes the gaps between the highlights, but it takes up to $n = 1025$ input images until a clear separation into the colors red, green and blue takes place. The adaptive sampling scheme, however, reaches full color fidelity much earlier, even after only 128 samples are inserted ($n = 384$).



10 000 image
reference

$n = 1025$,
fixed sampling

$n = 1024$,
adaptive sampling

Figure 5.9: Renderings of the ceramic figure in the high frequency environment map shown in Figure 5.6. As can be seen, the renderings of the colored reflections on the cap is much closer to the reference than in the regular setting for the adaptive sampling case. Still, the narrow light source causes spotty artefacts in the remaining parts, as highlights are disconnected. This can be avoided by using a wide aperture for these samples (see Figure 5.2), which motivates two-level reflectance fields.

As can be seen in Figure 5.8, which shows a detail left to the bottle's center, there are also regions where adaptive sampling reaches full quality slightly later than fixed pattern sampling. This occurs if the intensity and the image size of highlights is small, so these regions are refined later in the algorithm. While the fixed pattern sampling gradually reduces the distance between light source positions as the sample number increases, the adaptive algorithm may focus on other regions first, and refine small highlights only relatively late.

In Figure 5.9, renderings for the ceramic figure data set are displayed. While, due to its isotropic reflectance field, the adaptive scheme does not achieve a sufficient density to obtain continuous highlights for $n = 1024$ total images, the highlights on the cap of the figure are reconstructed more precisely than for $n = 1025$ uniformly distributed samples.

Timings / Performance

The acquisition time required to record a data set in our setup not only varies with the total number of recorded images, but also with the sequence type. We will now list timings at the example of the bottle data set. For the HDR reconstruction, two standard PCs were used, one of which also ran the refinement computations and control loop, the projectors and camera were controlled by a laptop.

Acquiring $n = 1025$ sample images in the fixed pattern distribution took 131 minutes – including the multiple exposures required for the HDR reconstruction – plus 25 minutes of cropping and scaling for the area of interest. Adaptive sampling of $n = 1024$ images, of which 256 are distributed regularly in the initial set, and 16 are added in each following iteration, can be done in 245 minutes including all processing. The acquisition of the 10 000 image ground truth data took 1023 minutes plus 248 minutes of processing.

The additional time for adaptive sampling is dominantly spent on the following items: the re-initialization of the measurement devices in each iteration, the unordered rotation of the moving heads to the unevenly distributed target positions, and the time for evaluating the distance measure in Equation 5.2. A more efficient implementation with more advanced pipelining and reduced I/O overhead should cover the first two points; the analysis of the reflectance field, point insertions and triangulation operations in total take about 1.9s / image.

5.4 Two-Level Reflectance Fields

The previous results show already an improvement over a fixed pattern sampling, as sampling concentrates at the most important light directions. However, as we use a global measurement process, undersampling deficiencies can still occur in areas where small image regions did not contribute sufficiently to the global error measure so as to warrant further subdivision, but are sensitive to only sparsely sampled light directions. Such artefacts may, for example, show up as discontinuous highlights in high curvature regions.

Following the discussions in Sections 5.1 and 5.2, we can mask those effects by replacing those samples by samples with a stronger prefiltering. We therefore propose a two-level representation for the reflectance field which contains samples from two different filter kernels, which we will discuss in more detail in the next paragraphs.

We partition the final set S_n of sampling light directions S_i into two subsets,

$$S_n = S_{\text{narrow}} \cup S_{\text{wide}} \quad (5.3)$$

so that S_{narrow} contains densely sampled regions of important directions, and S_{wide} contains sparse samples of directions in smooth regions, filling up the hemisphere. Then, we can create a combined reflectance field of narrow and wide spot light apertures: for the

first field, we retain the images already taken during the adaptive sampling step. For the second field, we acquire new images with fully open spotlight aperture and active frost filter. For each of the directions in S_{wide} , we scale the brightness of the full-aperture image so that it matches the average of the respective previous images in narrow lighting. Then, we use it as a replacement in the reflectance field. For rendering, we interpolate the field with the same algorithm as before (see Section 4.4).

The detailed process is illustrated by an example: we start with one of the adaptively sampled reflectance fields of the bottle scene from above, and again consider neighborhoods N_i for the directions $(\theta_i, \varphi_i) \in S_n$ as in Section 5.3.1. For each of the directions which were in the initial sample set S_0 , and for which the difference angle to any of the directions in N_i exceeds a threshold (9 degrees in our experiment), we record an additional image lit from the same light direction, but with full spot aperture and active frost filter.

By restricting ourselves to $S_{\text{wide}} \subseteq S_0$, we can sample S_0 once in the wide aperture setting, and quickly create renderings for different values of n . This restriction, however, would not be used in an application case.

Figure 5.10 shows renderings for such fields, in the same high frequency environment map as before, and lists the number of total images needed (both from the initial measurement and from the additional measurement in wider light conditions). For $n = 256$, all light directions fulfill the criterion for re-measurement, so they are all replaced, and the rendering shows dull highlights, as expected.

The next row shows a large improvement already. For $n = 384$, the rendering of the resulting reflectance field shows both the high quality highlights at the top of the bottle, and a continuous illumination in the right-column detail enlargement. While 244 images need to be re-recorded, driving the measurement costs up to 628 images, the result is visually almost on the same quality level as uniformly lit, adaptively or regularly sampled reflectance fields with higher image counts. Since some of the directions have been captured both with narrow and wide spotlights only 384 images need to be finally stored and accumulated in the rendering step. Figure 5.11 shows a result for a real-world environment map.

5.5 Discussion

5.5.1 Evaluation using sampling theory

In this section, we examine the relationship between sampling density and the shapes of the light sources. More specifically, we have computed the power spectra of the light intensity plots on display in Figure 4.2, as shown in Figure 5.12 as log-scale plots.

Here, the advantage of using extended light sources with smooth fall-offs becomes apparent. While, for each frequency, some power is present in the spectrum, there is a quick

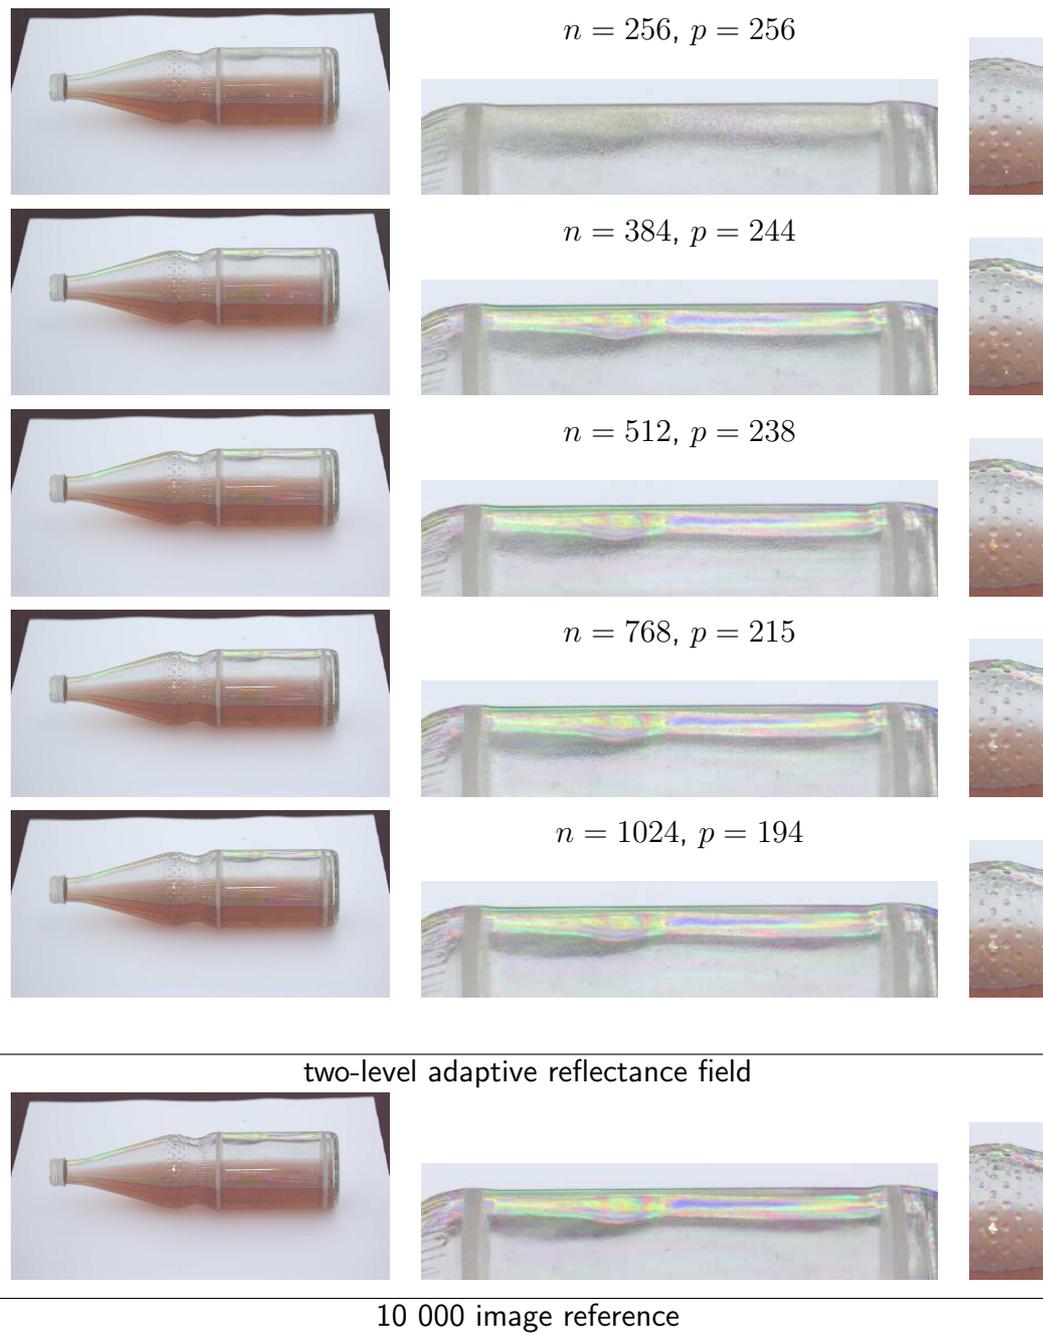


Figure 5.10: Reflectance fields of different illuminant spot sizes can be combined into one in the adaptive scheme. Here, n is again the number of images used in the data set, p is the number of pictures which were re-recorded for wide aperture. For instance, in the top row, being the initial data set, all pictures were replaced.

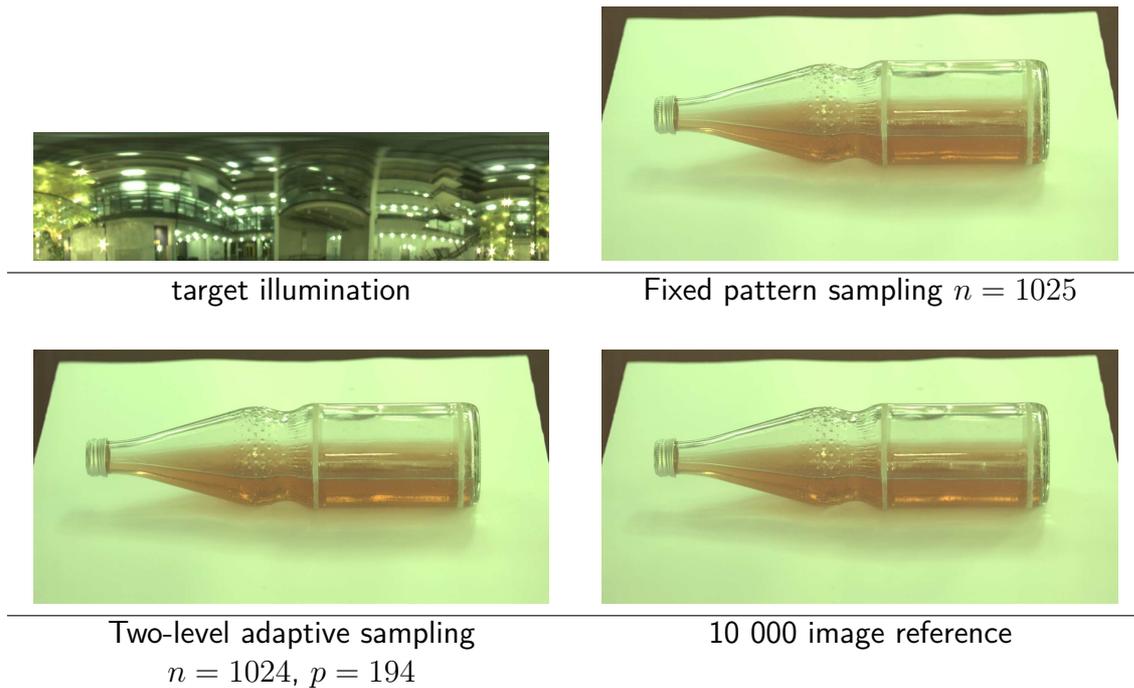


Figure 5.11: Rendering of the bottle scene in a real-world environment map. While the differences are subtle, the fused sampling field is closer to the 10 000 reference than a fixed-pattern sampling, especially in the bright red reflection at the lower border of the bottle.

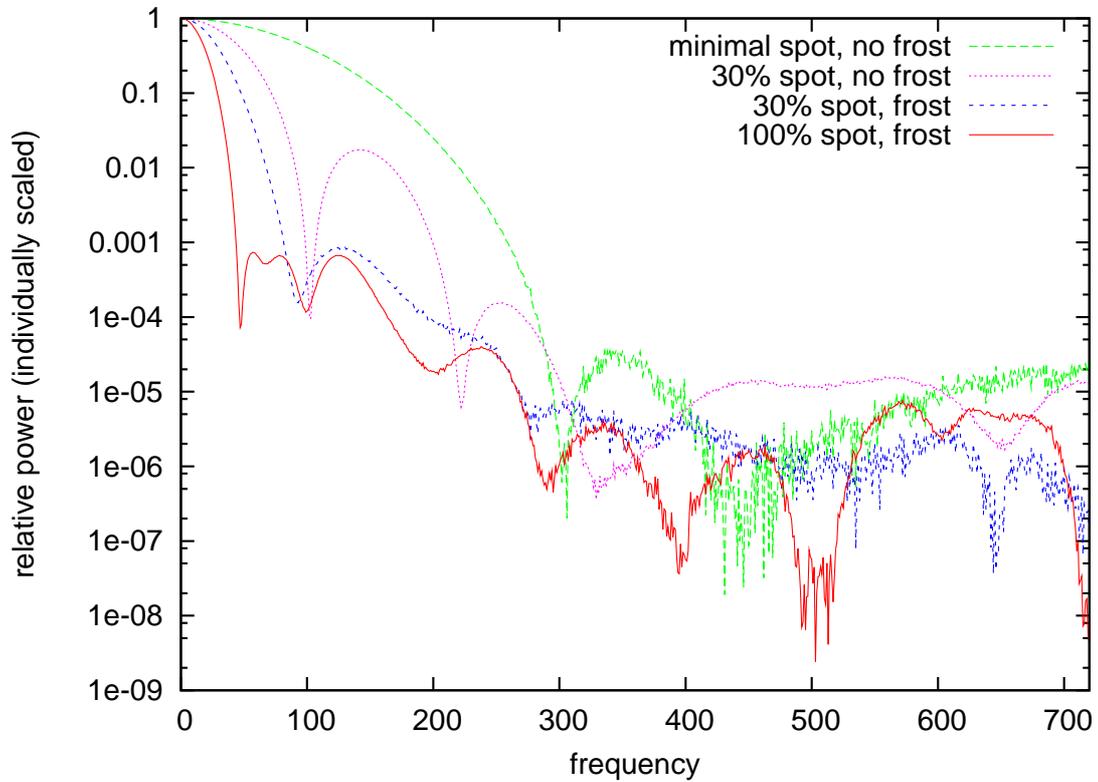
decay as the frequency increases. Thus, we can use the sampling theorem to motivate a maximum sampling frequency.

If, for instance, we choose to ignore frequencies whose contribution is less than 10% and we use the minimal spot size, we need to deal with frequencies of up to 160 cycles per 360° ¹, requiring to place the samples at a distance of at most $d_{10\%} = \frac{360^\circ}{160 \cdot 2} = 1.125^\circ$. The table in Figure 5.12 lists the respective values for the four listed light source configurations.

We can now evaluate the direction distributions that we used by computing their Delaunay triangulations and looking at the angles the edges span. They represent distances between sampling directions which are adjacent in the Voronoi diagram. Figure 5.13 shows percentile plots of these distributions. As can be seen, the reference data is sufficiently dense so that, when recorded with the narrow light source, the reflectance field data will expose aliasing only in frequencies which contribute less than 10% to the transmitted power. The distances are densely concentrated, corresponding to an approximately even distribution.

For the fixed pattern, the samples are also evenly distributed; for the narrow light source,

¹In order to observe the light source fully, we have positioned it few degrees above the horizon of the recording fish eye lens; as a result, the measured intensity distribution covers less than 360° arc length. Still, the reported minimal required densities are misestimated by less than 4%.



type	$\nu_{10\%}$	$d_{10\%}$	$\nu_{50\%}$	$d_{50\%}$
minimal spot	160	1.125°	88	2.045°
30% spot ("narrow")	72	2.5°	42	4.286°
30% spot, frost	48	3.75°	26	6.923°
100% spot, frost ("wide")	29	6.207°	17	10.588°

Figure 5.12: Power spectrum for a selection of light source configurations computed with a Discrete Fourier Transform on the data sampled at 1440 locations along φ for a fixed θ . Each spectrum is rescaled so that maximally transportable power equals 1. The table to the right lists, for each configuration, the first frequency where the contained power drops below 10% / 50%, respectively, and the maximum allowable distance between two sampling directions so that this frequency can be represented according to the sampling theorem.

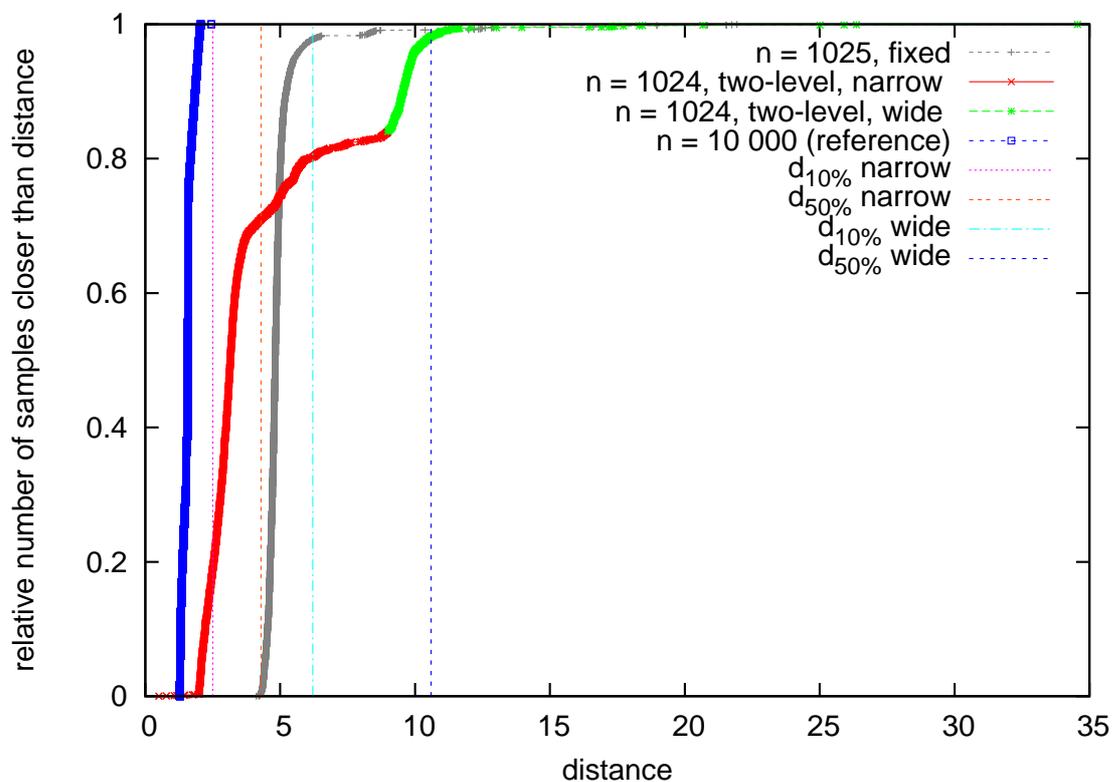


Figure 5.13: Distribution of angles between neighboring sample directions for the bottle data set in quantile plots. The dashed vertical lines illustrate required sampling densities for the narrow and the wide light configurations, if frequencies contributing less than 10% or 50%, respectively, may be ignored. The blue line shows the distribution of the 10 000 reference points set, the grey line shows the distribution of 1025 sampling directions in a fixed pattern, the red/green line shows the distribution of the adaptive bottle data set directions after 48 iterations with 1024 recorded images. The red part is recorded with narrow light sources, the green part is recorded with wide light sources in the two-level field.

they are merely dense enough to measure a signal which may still expose aliasing in frequencies which are present with 50% or more of relative power.

The adaptively sampled two-level data set has an uneven distribution. In the densely sampled area it is close to the $d_{10\%}$ limit for narrow light sources. At 9° degrees and above, we switch to the wide light-source recordings which require less dense sampling. This threshold separates between the adaptively inserted directions and the start set of directions S_0 (see knee in Figure 5.13) and is appropriate if we may assume that the adaptive sampling algorithm has reliably sampled the most important regions. A more conservative threshold would be e.g. $d_{50\%} = 4.286^\circ$ for the narrow light source, as it would impose a limit on the aliasing that can possibly occur.

5.5.2 Limitations of this technique

The above discussion of sampling densities in the context of sampling theory includes several simplifications. Especially, it analyzes sampling densities as if the reflectance field was a 1D signal; a 2D analysis would be more appropriate, but this significantly more complicated analysis would be beyond the scope of this chapter. Also, as the tent that we used for illumination is not spherical, and the projectors are not precisely in its center, an analysis of sampling effects might also include variations in the light distributions with θ and φ .

Unlike authors such as Peers and Dutré (2005), we do not estimate independent reconstruction functions for each image pixel. This avoids fitting errors, but restricts us to a global data representation and a global metric for the refinement of the reflectance field sampling. Accordingly, small image regions may expose interesting reflectance features which are not precisely sampled by the adaptive sampling algorithm. Still, we hide aliasing artefacts by sampling sparsely observed regions with large kernels in the two-level approach. As the large kernels overlap, we can be sure that we do not miss features altogether.

The adaptive sampling algorithm is inherently online, i.e. it can only take decisions on where to refine the reflectance field based on data it has already observed; thus, important features like isolated highlights may be missed. Still, they will be captured by the wide kernel measurement in the two-level algorithm. Also, we have not exhaustively investigated the refinement criterion. Yet, our evaluations show a good performance of the adaptive sampling.

The evaluation is limited in the sense that we do not take possible target illuminations into account, but use an evenly distributed, dense reference reflectance field. The least squares metric for the comparison of reference and rendering is not perceptually motivated.

We have shown that Reflectance Fields can consist of samples in different resolutions / taken with different filter kernels. In order to evaluate the improvements with increasing sampling counts – see Figure 5.10 – effectively, we have only combined samples of two different kernels. A more principled way is to use all available light configurations

and analyze the corresponding Nyquist Frequencies. After the adaptive algorithm has determined the sampling positions, each sample can then be recorded with an optimal light source.

Finally, the measurement setup construction involves a trade-off between optimal (close to spherical) geometry and easy construction, which led to the tent design. This non-uniformity is partially mitigated by the light source intensity calibration.

5.6 Conclusion

In image-based relighting, the quality of the renderings depends critically on the characteristics of the illumination basis used for measuring the reflectance field. In this chapter, we present a way of numerically evaluating the quality of sampling patterns of different light directions. We have introduced an adaptive sampling scheme and shown, both quantitatively and qualitatively, how it improves on the trade-off between measurement effort, i. e. the required number of input images, and measurement quality when compared to fixed pattern sampling.

The automated measurement setup described in Chapter 4 makes the required data available efficiently, scaling up to thousands of images. The acquisition system allows for measuring with an extended light source covering any possible direction of the entire hemisphere by indirect illumination reflected of dark cloth. The dark cloth reduces the power of ambient light effectively. We have shown that the use of extended light sources can significantly improve the rendering quality especially at low sampling rates.

A new quality metric is required to correctly compare the quality of a two- or multi-level reflectance field with varying light source diameter to one acquired with a spotlight of fixed width. In this context, the conceptually simple least-squares metric used in the proposed technique could also be replaced with something perceptually more meaningful. Still, the simple scheme presented here already produces useful results for very low image counts.

Our findings demonstrate that adaptive measurement schemes, which take the specific properties of objects into account, have a significant potential to improve the efficiency of reflectance field methods.

6 Feature-Guided Image Synthesis for Reflectance Field Interpolation

6.1 Introduction

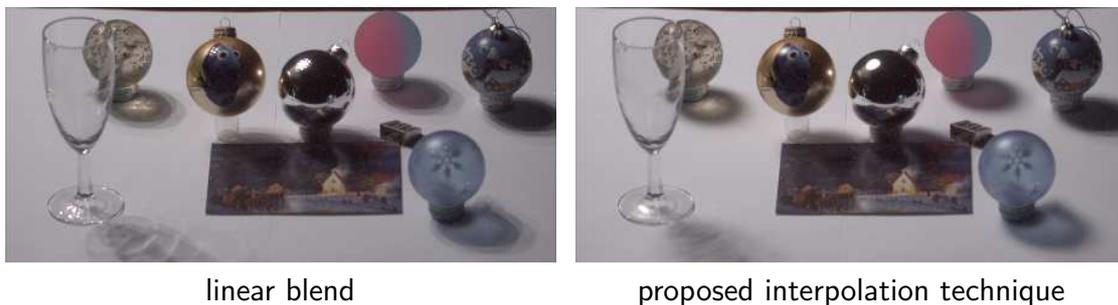


Figure 6.1: Relighting a sparsely sampled reflectance field (230 input images) with a small area light source. Top: Linear interpolation of the input samples results in banded shadows and disconnected highlights. Bottom: Our novel interpolation technique upsamples the same input reflectance field, effectively increasing its resolution to 3547 samples, and is capable of plausibly reproducing smooth shadows, shadows of semi-transparent objects, highlights and caustics.

As we have previously seen, undersampling of reflectance fields can cause ugly artefacts, and one should choose the sampling frequency and the sampling kernel accordingly. However, there are situations where neither is possible, for instance, if acquisition time constraints only allow few input images to be taken, and prefiltering would unduly blur out illumination features, especially highlights and shadows. While adaptive sampling may contribute to mitigating this problem, it cannot help in scenes where the illumination directions are roughly equally important (see Figure 6.1 for an example).

We therefore propose a novel interpolation and upsampling scheme that takes as input a sparsely sampled reflectance field captured with a narrow light source and produces a plausible reflectance field with much higher resolution, for instance, supersampling it from 230 captured to more than 3300 synthesized illumination directions. With the constructed superresolution reflectance field, the motion and appearance of high frequency illumination effects can be approximated without requiring to capture thousands of samples.

The intermediate reflectance images are created by first separating the input images into regions where simple linear interpolation is sufficient, and regions where it is not. As

undersampling artefacts are most visible along shadow boundaries and close to highlights, we treat those separately. We use a texture inpainting method to approximate object appearance beneath highlights, extract a highlight layer and perform an optical flow computation to warp the highlights smoothly to light directions that were not recorded during the acquisition step. For shadows, we use a heuristic to determine shadow pixels and run a morphological blending algorithm to interpolate the “shadowed” semantics to new directions. By separating the prediction whether a pixel should be shadowed or not from the estimation of its appearance when actually shadowed, we maintain local texture detail while reconstructing a sharp shadow boundary. We regularize the shadow upsampling using texture likelihood priors extracted from the input images.

The benefit of our approach is that from low sampling rates with high-spatial frequency illumination, we can simulate intermediate illuminations both at high spatial frequencies, such as point lights, and simulate extended light sources, such as those found in many light fields. From a theoretical perspective, the algorithm taps new sources of information about reflectance from sparse samples by modeling and separating the tight connection between the effects in the image plane and in the angular domain.

6.1.1 Related Work

Reflectance fields can be acquired using directional light, explicit basis illuminations, or can be inferred (refer to Section 3.4.2, page 29 for a detailed discussion). Inferring methods construct a reflectance model in a method-dependent process. While it may be possible to achieve super-resolution in such a context, an interpolation after the model construction seems difficult. Explicit basis illuminations create images in conditions that possibly vary greatly in their extent, which makes it difficult to interpolate in a unified approach.

This leaves directional light methods as candidates for upsampling. Masselus et al. (2004) discussed several local interpolation methods for $2 \rightarrow 2$ reflectance fields; an example for this is provided in Figure 6.1 for linear interpolation. While other interpolation methods may be smoother, local interpolators generally fail to move features in image space to pixels that did not observe them; therefore, highlights on curved, shiny surfaces stay disconnected under area light sources. In this chapter, we instead propose an advanced non-linear interpolation technique that can correctly reproduce the object’s appearance for in-between light source positions from a sparse sampling.

One way to estimate intermediate samples is to approximate the reflectance field locally by an explicit model. Various models have been proposed to describe the apparent BRDF for the surface point visible at a given camera pixel, such as analytic BRDF models (Debevec et al., 2000), which, however, do not account for global effects. Reflectance sharing (Zickler et al., 2005, 2006) does not require an analytic model, but has so far been demonstrated for opaque surfaces of known geometry only. In this chapter, we present interpolation techniques for reflectance fields of non-opaque surfaces and unknown geometry.

The problem of interpolation between light source directions is related to the problem of view interpolation, for which various techniques have been proposed, e.g. optical flow (Black and Anandan, 1996; Brox et al., 2004) or level set blending (Whitaker, 2000). Special solutions to handle illumination effects include flow for specular surfaces (Roth and Black, 2006) or interpolation for specularly refracting materials (Matusik et al., 2002b).

However, applying view interpolation techniques to the input images of a reflectance field directly has the inherent problem that the scene itself and all its texture is static while the illumination creates apparent motion on top of the static structures. Barrow and Tenenbaum (1978) therefore proposed the separation into intrinsic and illumination images. For reflectance fields of mostly diffuse scenes, Matsushita et al. (2004) used the illumination image to detect shadows which are then moved according to the position of the light source with the help of an explicit 3D model of the scene. Applying view-interpolation techniques on thresholded binary images of detected shadow regions, Chen and Lensch (2005) generated smoothly moving shadow regions. They worked on 6D reflectance fields where the illumination is controlled by a projector. This allowed them to turn off the direct illumination to the shadow regions while still considering the indirect illumination from other scene parts.

Our upsampled reflectance fields are regularized by applying image-based priors introduced by Fitzgibbon et al. (2003). They performed view interpolation and implicitly reconstructed a depth map subject to the constraint that the interpolated view is locally consistent with the recorded image data. The constraint enforces that any pixel neighborhood in the interpolated view has to occur somewhere in the input views. Wexler et al. (2004) further extended this idea to fill in holes in a space-time video cube. Without any additional information, relying on image-based priors only, the holes are filled in by a multi-resolution framework using 3D neighborhoods. At each level the solution of the previous level is regularized to match the image-based priors at that resolution.

Upsampling reflectance fields can be viewed as filling holes in a 4D structure. In contrast to the work by Wexler et al., our input data unfortunately provides a sparse set of densely sampled 2D slices only, so we only have priors for 2D neighborhoods. We therefore include the reconstruction step based on highlight and shadow maps in order to recover 4D consistency.

6.1.2 Overview

Reflectance fields which were captured with directional illumination consist of a number of n slices $R_i(x, y)$, corresponding to images capturing the appearance of the scene for a single illumination direction (θ_i, φ_i) . In our approach, we address the interpolation problem by synthesizing plausible intermediate images R_k at a much higher density in the domain of light directions (see Figure 6.2) and inserting them into the measured field. From there, we can proceed with any interpolation scheme; like before, we show results

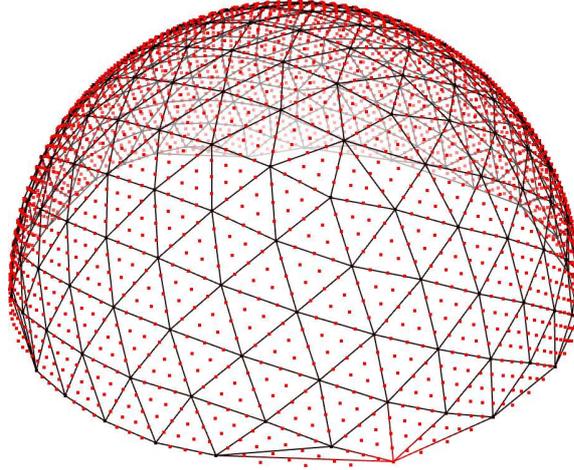


Figure 6.2: The incident light directions of a sparsely sampled ($n = 230$) reflectance field visualized as black dots, and a connecting Delaunay triangulation which defines barycentric coordinates on which linear interpolation can be performed. We increase the resolution by synthesizing samples for in-between light directions (red dots).

based on linear interpolation in barycentric coordinates on a Delaunay triangulation of the input images.

We generate the intermediate images R_k by processing the input images R_i according to the pipeline depicted in Figure 6.3. The pipeline consists of three main stages: segmentation and labeling of the input data, separate upsampling of the highlight and shadowmaps, and reconstruction at the target resolution subject to image-based priors.

The segmentation stage (Section 6.2) separates out highlights and shadows which are treated separately in our pipeline. From the input images R_i we extract the images H_i , that just contain the highlights, and reconstruct images D_i where the information beneath the highlights is plausibly reconstructed using image-based priors. In a second step, a binary shadow map S_i is computed, which encodes whether a pixel is likely to be in shadow ($S_i(x, y) = 0$), or not ($S_i(x, y) = 1$). The shadow labeling does not need to be very precise since it is used mainly to obtain the correct movement of shadow boundaries, but not the interior intensity.

In the second stage of the pipeline (Section 6.2), we upsample the highlight and shadow data to higher resolutions in the θ, φ domain: we apply optical flow to warp the highlight images H_i and perform level-set image blending (Whitaker, 2000) on the shadow maps S_i .

In the third stage, a tentative intermediate image R_k^{reco} is synthesized. From the warped shadow images we determine for each pixel how to interpolate its reflectance from shadowed or lit samples. We then apply image-based priors in order to regularize our solution, removing intensity or high frequency artefacts introduced by incorrect segmentation. Finally, we add the warped highlight layer and arrive at the interpolated image R_k .

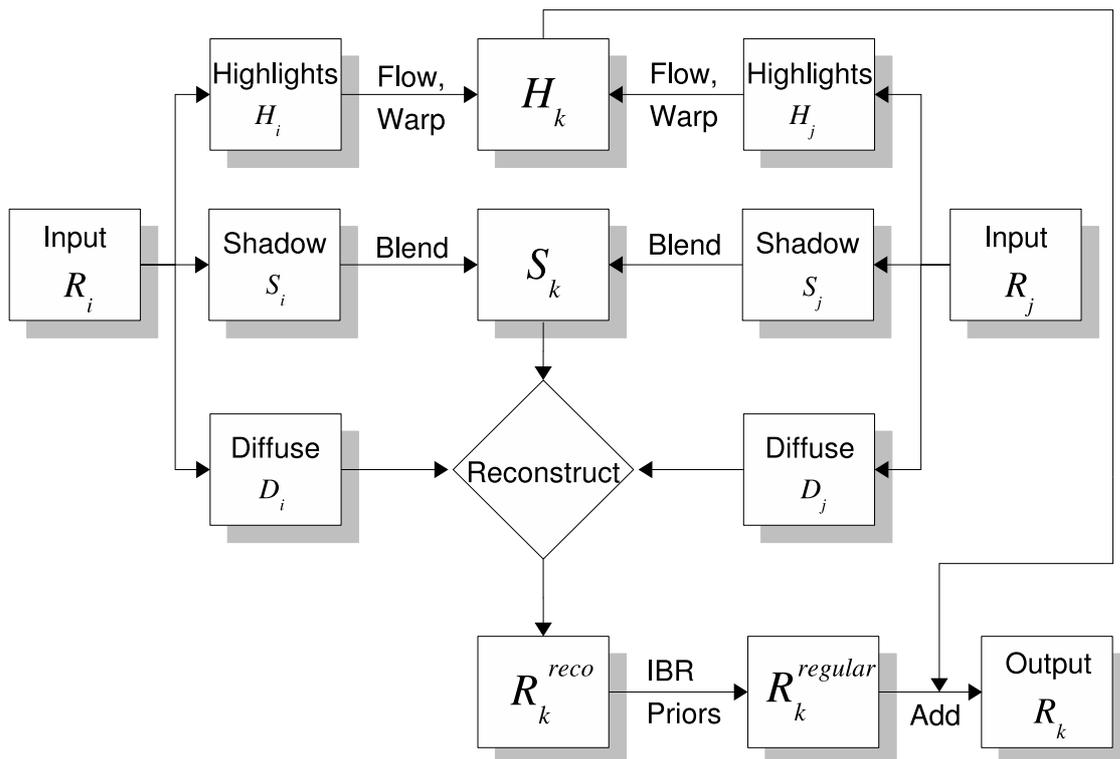


Figure 6.3: Input slices R_i , R_j of the reflectance field for incident light angles θ_i, φ_i and θ_j, φ_j , respectively, are separated into highlights, shadows and a diffuse component. We interpolate the components separately and combine them to a new slice R_k for an intermediate light direction.

6.2 Extracting Image Semantics

Our method begins with the extraction of semantic maps of highlights and shadows for our input images R_i . In order to obtain a criterion to determine the status of each pixel, we sort the measured reflectance values $R_i(x, y)$ for each pixel, resulting in images P_i , where P_0 contains the dimmest reflectance for each pixel and P_{n-1} the brightest. From these pictures we select a picture P where no pixel is a highlight nor a shadow. For an input reflectance field consisting of $n = 256$ images, we found $P = P_{230}$ to be a robust choice. Based on P , shadow and highlight maps are generated in the following steps.

6.2.1 Specular Highlights

The separation of highlights is done separately for each input image R_i in several steps. At first, a binary map $M_i^{high}(x, y)$ encodes whether a pixel $R_i(x, y)$ contains highlight

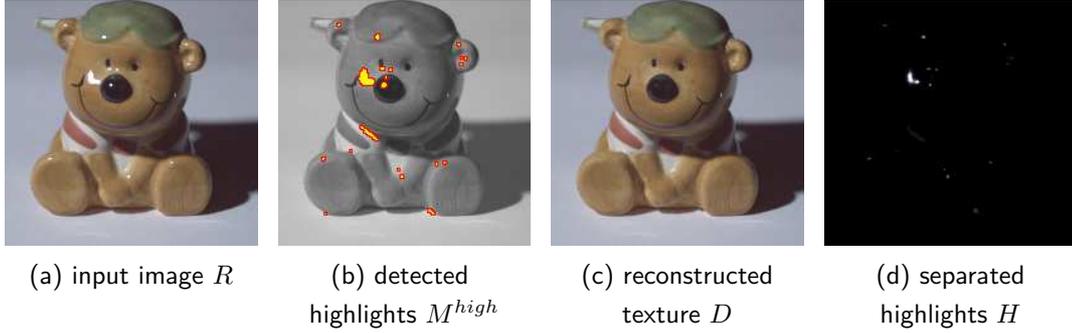


Figure 6.4: Separation of highlights. By intensity analysis in the input image R (a), a region of definite highlights M_{high} is detected (b, colored in yellow) which is dilated (b, colored in red) for robustness. Texture reconstruction allows to estimate the appearance D below the highlights (c). By subtracting D from R , the highlight layer H for this image is estimated (d). Note that the reconstructed highlights have smooth boundaries.

information or not. Initially, $M_i^{high}(x, y)$ is 1 iff $R_i(x, y) \geq f_{highlight} \cdot P(x, y)$, where $f_{highlight}$ is a constant factor for all images; in our experience, $f_{highlight} = 2$ or 4 is a good choice. Then, we dilate $M_i^{high}(x, y)$ with a 3×3 or 5×5 box filter, growing the highlight regions to include possible haze around them.

Based on M_i^{high} , we generate the highlight-free images D_i by removing all possible highlight pixels from R_i and applying the hole filling algorithm by Wexler et al. (2004) (see Section 6.4.2) to reconstruct plausible texture information replacing the highlights. As texture references, we use R_i and the images for neighboring light configurations, but we remove potential highlight pixels. In modification to Wexler et al.'s original algorithm, we do not perform an outlier detection, and add a term to the region lookup which favors close-by regions, as this improves the performance of the lookup structure.

Finally, the highlight map $H_i(x, y) = R_i(x, y) - D_i(x, y)$ is computed as the per-pixel difference. Figure 6.4 illustrates the highlight detection and removal process.

6.2.2 Shadow Boundaries

In contrast to the highlight maps, the shadow maps are generated not independently for each input image, but independently for each pixel at position (x, y) . We perform a region growing on lit regions and treat the remainder as shadow.

More precisely, a shadow map $S_i(x, y)$ is generated for all directions (θ_i, φ_i) , with initially $S_i(x, y) = 1$ for “lit” iff $R_i(x, y) \geq P(x, y)$, and $S_i(x, y) = 0$, otherwise. In order to mitigate camera noise one can perform this segmentation based on a prefiltered version of R_i . We then perform the following iterative update: Let (θ_i, φ_i) be the direction for which pixel (x, y) is currently labeled as shadowed ($S_i(x, y) = 0$), and (θ_j, φ_j) a neighboring direction for which the pixel has been labeled as lit ($S_j(x, y) = 1$). For image R_i , the pixel is re-labeled as lit ($S_i(x, y) := 1$) if the following criterion is met: $R_i(x, y) > f_{lit} \cdot R_j(x, y)$. We usually chose $0.66 \leq f_{lit} \leq 0.9$. The region growing is

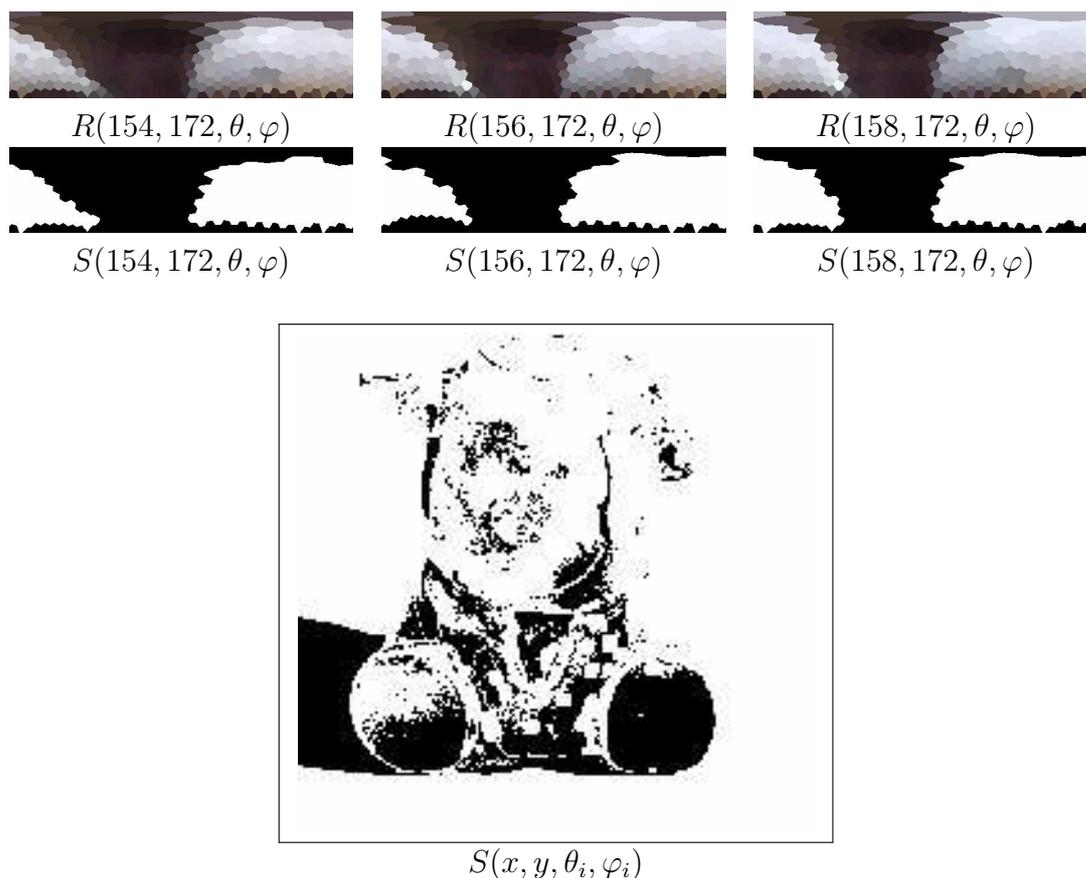


Figure 6.5: Computation of shadow maps. At the top, the top row shows a plot of acquired reflectance in polar parameterization for three pixels (to the right of the figure’s left foot). Below, the segmentation in shadow (black) and lit (white) regions is displayed. The large image shows a slice of the shadow map in image space for a fixed incident light angle (θ, φ) . While noisy in image space, the shadow computation is stable along the (θ, φ) directions.

iterated until no more updates take place. Thus, the lit regions grow until the difference to neighboring directions becomes so large that linear blending is likely to fail, indicating a shadow boundary at this location.

Figure 6.5 illustrates the detection of shadow maps. Note that while the results may be noisy in the (x, y) domain, the scheme is continuous in the (θ, φ) domain, where the reconstruction later takes place.

6.3 Image Feature Upsampling

Before we can synthesize reflectance images I_k at the full resolution, we need to procure estimations of the shadow and highlight distributions for this resolution. Let S_i and S_j be shadow maps, and H_i and H_j be highlight maps for two input images R_i, R_j that have

adjacent light directions (θ_i, φ_i) and (θ_j, φ_j) in the mesh provided in Figure 6.2. In this section, we will explain how to generate intermediate maps S_k and H_k for the in-between direction. The full target resolution is then later obtained by iteratively subdividing the mesh edges and recomputing the triangulation until the desired resolution is achieved, i.e. information is available at all red dots in Figure 6.2.

6.3.1 Upsampling Specular Highlights

The generation of H_k is rather straight-forward. We apply the optical flow algorithm by Brox et al. (2004) both to obtain a flow field from H_i towards H_j , and from H_j to H_i . We clamp the HDR data of the highlights to obtain a smaller dynamic range, which makes matching easier. Using the flow fields, we warp H_i and H_j with full dynamic range to the halfway position, and blend the result linearly to obtain H_k .

The flow algorithm is capable of generating very smooth fields, which is useful for our application, as it allows to drag highlight data along even if only neighboring parts matched. As we segment our highlights carefully and the separated highlight image is mostly black, this does not introduce artefacts.

6.3.2 Upsampling shadow data

For the interpolation of the shadow maps S_i , we employ the level-set blending approach by Whitaker (2000), which generates intermediate shapes of level sets within images. In contrast to optical flow, this algorithm is more tolerant to noise in the image domain, estimating the correct movement of the shadow boundaries while keeping fixed structures constant. In contrast to the original article, we replace the up-wind scheme for the computation of partial derivatives with simple central differences, as this yields more stable results for our application. Since the level-set blending algorithm produces continuous images even for sharply segmented input data 0 or 1, we re-quantize the output again. This produces plausible shadow boundaries for in-between images (see the central row of Figure 6.6).

6.4 Image Synthesis

Having generated the S_k and H_k for the target resolution, we can now synthesize the output images R_k . Let (θ_k, φ_k) be the light direction for which the synthesis is required, and (θ_a, φ_a) , (θ_b, φ_b) , and (θ_c, φ_c) be the triangle of observed input directions from Figure 6.2 which contains (θ_k, φ_k) . The reconstruction takes three steps: first, reconstruct the diffuse appearance R_k^{reco} using the estimated shadow maps, second, regularize the estimation using IBR priors, yielding $R_k^{regular}$, and third, add the highlight layer H_k resulting in the reconstruction R_k . Figure 6.6 illustrates the different reconstruction

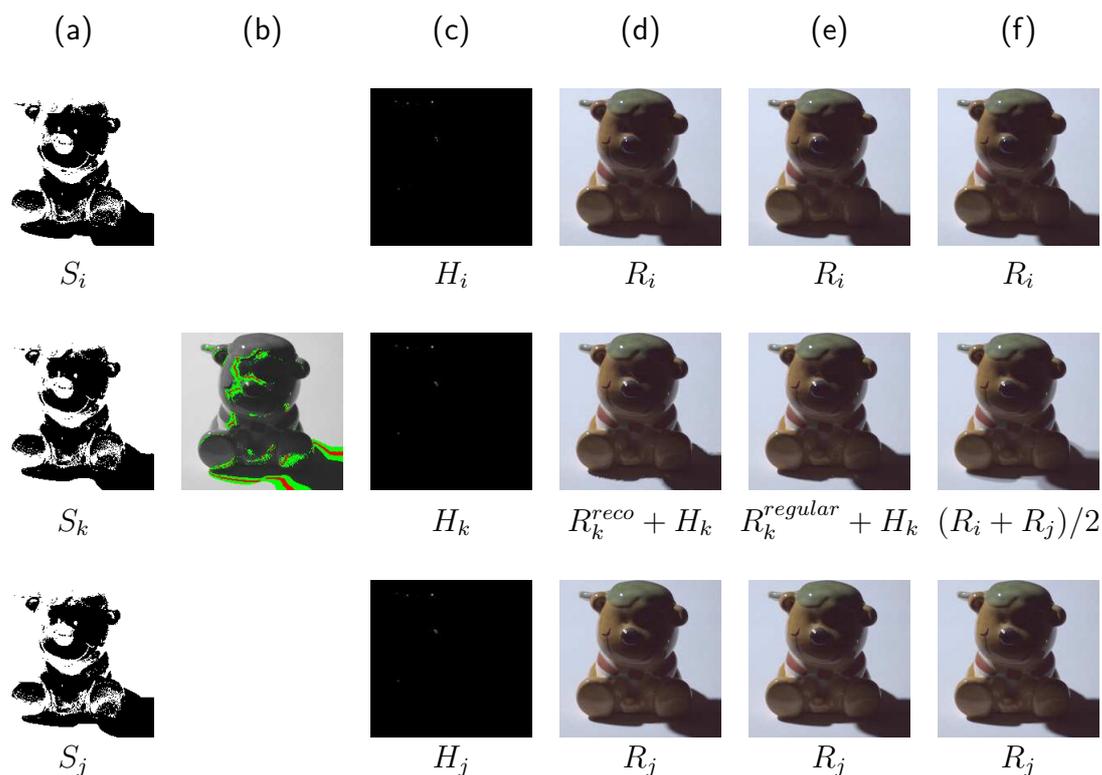


Figure 6.6: Illustration of the reconstruction process. The bottom and top row contain information from the initial, sparse input field. The in-between row is synthesized. By analyzing the shadow maps (a) we can obtain areas (b) where we either can locally interpolate (gray) or need to locally extrapolate (green and red). The green area is predicted to be fully lit or fully shadowed, the red area is the estimated location of the shadow boundary. After adding the highlight map (c), reconstruction yields a rough picture (d) which is improved using image-based priors (e). Especially the shadow boundary is much more pleasing than in a direct linear interpolation (f).

steps. In the figure, the highlight map H_k has already been added to the reconstruction results based on the shadow map.

6.4.1 Reconstruction Based on Shadow Maps

We first reconstruct the initial diffuse appearance R_k^{reco} , given the upsampled shadow maps S . For the construction of R_k^{reco} , we need to obtain estimates for every pixel's appearance both for the hypothesis R_k^{lit} that the pixel may be lit, and for the hypothesis $R_k^{shadowed}$ that it is in shadow. In this step, the shadow maps serve two purposes: a prediction whether a given pixel (x, y) should be shadowed or lit, and a statement whether the input images $R_{a,b,c}(x, y)$ describe the pixel in lit or shadowed state. However, as the shadow maps are binary and possibly imprecise, they cannot be used directly. Instead, we compute a blurred shadow map \bar{S}_k by averaging over the 10 nearest light

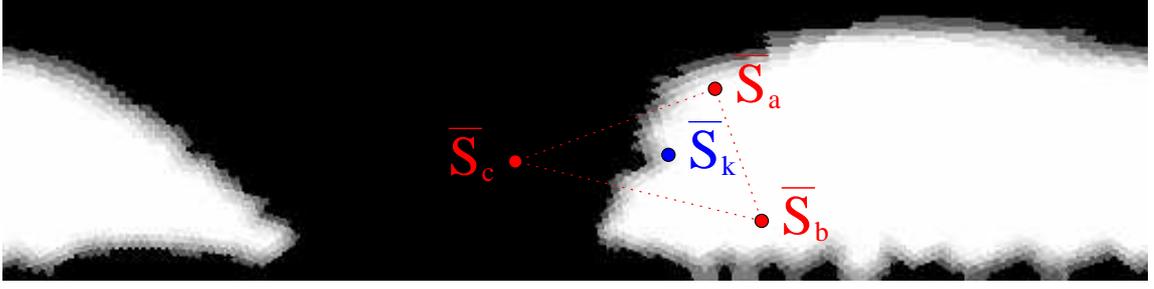


Figure 6.7: Polar plot of all blurred shadow maps at a fixed pixel position (white is lit, black shadowed, gray uncertain). The reflectance at light direction k is reconstructed according to the locally estimated shadow state \bar{S}_k by extrapolation from nearby fully lit (white) and fully shadowed (black) measured samples. If the shadow state of all observed conditions a, b, c agrees, direct linear interpolation takes place.

directions, which we use as smooth approximation of the relative shadowiness in pixel $R_k(x, y)$. Analogously, we also compute smoothed shadow maps $\bar{S}_{a,b,c}$, which are used for estimating the shadow state of $R_{a,b,c}(x, y)$; optionally, one can average over more neighbors in order to be more conservative (see Figure 6.7).

We can now compute the values $R_k^{shadowed}(x, y)$ and R_k^{lit} for the appearance of $R_k(x, y)$ both in shadowed state and in lit state. A set of cases is possible:

1. $\bar{S}_a(x, y) = \bar{S}_b(x, y) = \bar{S}_c(x, y) = 1$, i.e. all neighboring input images have reliably observed the pixel while it was lit: We compute $R_k^{lit}(x, y)$ using a direct interpolation using barycentric coordinates α, β, γ as $R_k^{lit}(x, y) = \alpha R_a(x, y) + \beta R_b(x, y) + \gamma R_c(x, y)$.
2. $\bar{S}_a(x, y) = \bar{S}_b(x, y) = \bar{S}_c(x, y) = 0$: Similarly, we compute $R_k^{shadowed}(x, y) = \alpha R_a(x, y) + \beta R_b(x, y) + \gamma R_c(x, y)$.
3. For all other cases, we extrapolate a value for $R_k^{lit}(x, y)$ (and $R_k^{shadowed}(x, y)$) by robustly fitting a linear spherical harmonic reconstruction using two SH bands to the observations of nearby light directions, for which a lit state (shadowed state) have been observed.

The output pixel is finally estimated as $R_k^{reco}(x, y) = \bar{S}_k(x, y) \cdot R_k^{lit}(x, y) + (1 - \bar{S}_k(x, y)) \cdot R_k^{shadowed}(x, y)$

Note that in cases where we reconstruct either a fully lit pixel from fully lit neighbors, or a fully shadowed pixel from fully shadowed neighbors, we effectively perform a direct linear interpolation of the pixel value. This implies that outlier pixels, which are wrongly estimated to be shadowed or lit, but are consistently misestimated for a set of close neighbors, tend to be robustly linearly interpolated as well.

6.4.2 Regularization Using Image-Based Priors

The reconstruction explained above provides an estimate for the output image, but this estimate may be unreliable, e.g. due to misregistrations of shadow or interpolation defects. The affected areas consist of pixels for which a pixel value was computed from extrapolated reflectance data, or in the vicinity of shadow boundaries in $\overline{S}_k(x, y)$. For these, we run a regularization algorithm, which we will describe in more detail below.

We have adapted the approach by Wexler et al. (2004) to perform hole filling and regularization on 2D slices of a reflectance field: The algorithm iteratively updates all pixels within the hole region to better fulfill the given priors: Let E_0 be the initial estimate image a series of iterative updates E_i is constructed. For each window $\mathcal{W}^{E_i}(p, q)$ centered around pixel (p, q) we search for the best matching window $\mathcal{W}^{I_k}(s, t)$ centered around some other pixel in the entire set of input images I_k based on the modified distance measure

$$d(\mathcal{W}^{E_i}(p, q), \mathcal{W}^{I_k}(s, t)) = \arg \min_{\lambda} \sum_{u=-L}^L \sum_{v=-L}^L \|E_i(p+u, q+v) - \lambda I_k(s+u, t+v)\|^2, \quad (6.1)$$

where L is the radius of the window. In contrast to the original approach we search for the best matching window up to some intensity scale λ . Since our input data is relatively sparse the extension is necessary to find appropriate matches for regions which change their intensity slowly due to the cosine fall-off under diffuse illumination. An efficient lookup can still be achieved by normalizing the input window to unit intensity and then searching for the nearest neighbor using a k -D tree which is built out of normalized windows of the input images. The distance measure is furthermore translated into a similarity measure

$$s(W^{E_i}(p, q), W^{I_k}(s, t)) = e^{-\frac{d(W^{E_i}(p, q), W^{I_k}(s, t))}{2\sigma^2}} \quad (6.2)$$

which is used in the following to update all pixels (x, y) within the window $W^{E_i}(p, q)$. Let W_j^E be the set of all windows that contain (p, q) . Each corresponding matched window λW_j^I yields an estimate c_j on what the final color for (p, q) should be. Using the corresponding similarity measures s_j the update is computed as the weighted average:

$$E_{i+1}(p, q) = \frac{\sum_j s_j c_j}{\sum_j s_j} \quad \text{-- compare Wexler et al. (2004)} \quad (6.3)$$

Usually, three iterations of this algorithm are sufficient, and using only input images for neighboring light conditions is sufficient. The quality improvement is visible in Figures 6.8 and 6.9. Note that in most cases the image-based priors remove high frequency noise but do not corrupt apparent motions which is important to achieve temporal consistency. In some cases however, the shape of some shadow edges cannot be matched to any

action	avg. time	
initial sort for threshold image	53 sec.	t
generation of the initial shadow map S_i	62 sec.	t
separation into highlight H_i and D_i	129 sec.	i
level-set blending between two shadow maps	23 sec.	o
optical flow and warp for highlight map	9 sec.	o
blurring of all full resolution shadow maps	7 min.	t
reconstruction of one image without priors	50 sec.	o
reconstruction of one image with priors	< 15 min.	o

Table 6.1: Timings for the data set in Figure 6.9. Processing time either for (t)otal field, per (i)nput or (o)utput image.

input image due to a missing sample with similar orientation. This can lead to slightly deformed shadow edges, as in the slightly too flat shadow of the shiny sphere in the rightmost image in the second row of Figure 6.9. One approach to circumvent this problem would be to include rotated neighborhoods in the nearest neighbor search. Once the reconstruction is complete, the warped highlight layer H_k is simply added to the result, thus obtaining the final reconstruction.

6.5 Results

In order to demonstrate the performance of the proposed method we constructed super-resolution reflectance fields for two different scenes: a ceramic bear with highly specular coating (256 input images) and a second, more complicated scene composed of a set of spheres with drastically different reflection and transmission properties, as well as a champagne glass (230 input images). In both cases, the illumination directions were spread equidistantly over the upper hemisphere. Iterating the subdivision process two times, we synthesize about 4000 intermediate positions. The bear data set has a resolution of 200×196 and the sphere set of 453×211 pixels. Timings for a PC from the year 2006 are available in Table 6.1. The most expensive part of the algorithm is the application of the texture inpainting for the highlight segmentation and the regularization of the output. The table lists the runtimes split into separable units. While I/O operations are considerable if the steps are split this way, it is easy to exploit parallelization and run independent operations for several images concurrently, e.g. on a cluster of PCs. Due to the smaller resolution, the runtime required for the bear data set is considerable lower. For the renderings in Figure 6.8, we have only generated samples needed for the light situation, and could therefore afford to create a supersampling of three subdivision steps.

In Figure 6.8 we synthesized the superresolution reflectance field and compare it to the result one would obtain for linear blending and to a reflectance field captured at the full resolution. In the right column, we illuminated the scene from a single intermediate



Figure 6.8: Reconstruction results using linear interpolation from 256 images (top row), our upsampling scheme applied three times (center row) compared to a reference rendering with 10 000 input images (bottom row). Left column: simulated extended light source with 15 degrees radius, right column: renderings with a directional source. The linear interpolation produces shadow banding and significant gaps in the highlights. Our upsampling scheme improves significantly on the visual quality.

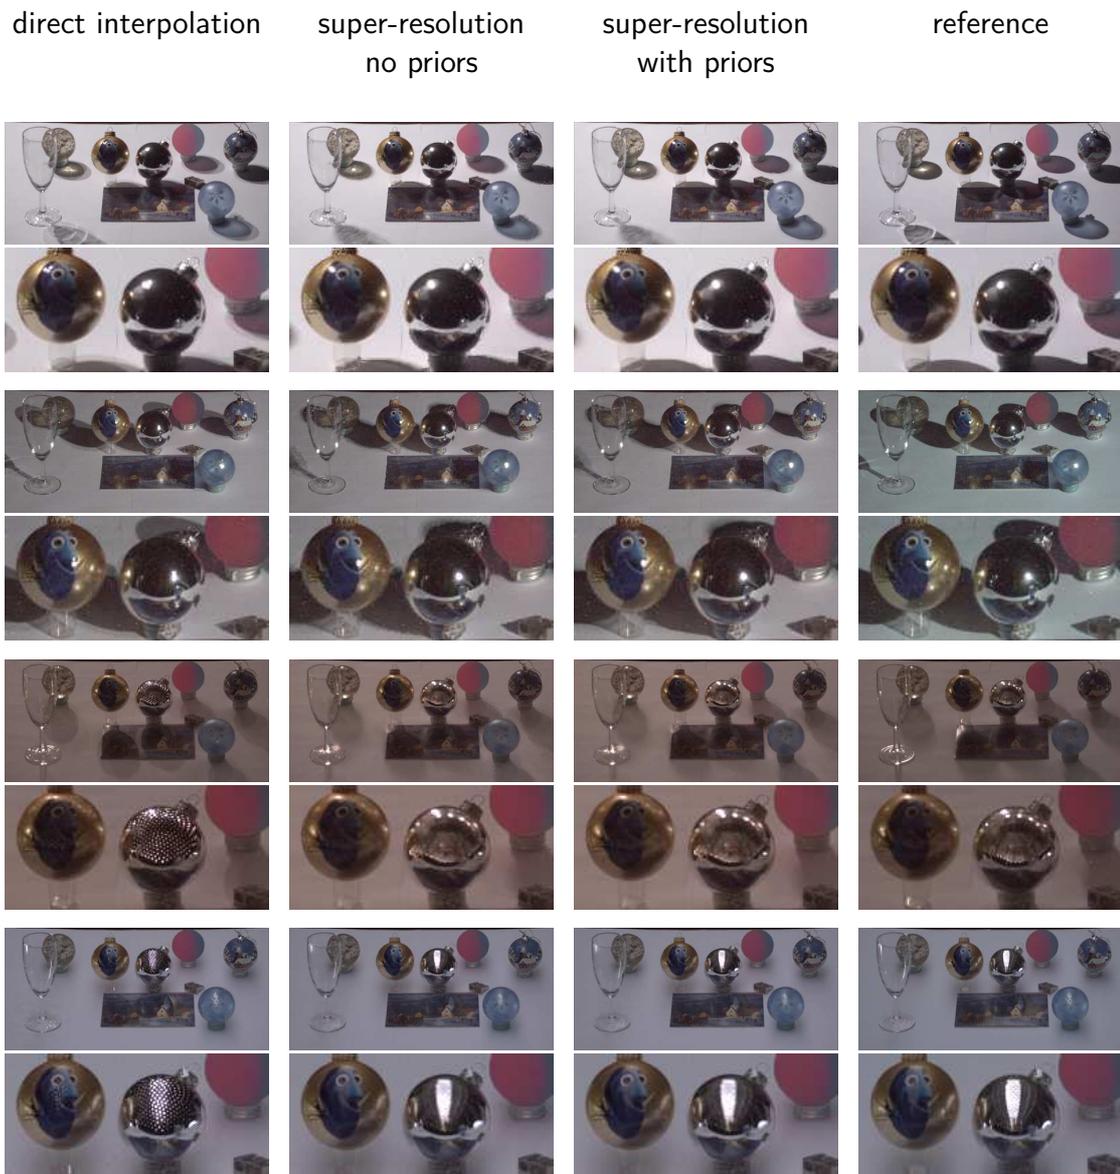


Figure 6.9: Results of a complex scene with detail zoom-ins for renderings in two point light conditions (upper four rows) and in the Galileo's Tomb and Uffizi environment maps (Debevec, 1998). The super-resolution reflectance field is upsampled from 230 to 3547 images, the reference has 14116 images. Our upsampling approach produces superior results compare to linear interpolation. Applying the regularization (second column from the right) further improves shadow boundaries; for real-world illuminations, the non-regularized version is already close in quality. Note that the reference exposes aliasing which the super-resolution images do not show due to sub-pixel blurring during the upsampling steps.

light direction. In the locally, linearly interpolated image double/triple contours from the three neighboring input images are visible which are removed by our upsampling scheme. While the reconstructed shadow boundary does not perfectly match the reference image, it is still plausible and provides a sharp contour. The left column shows the result after integrating over an extended light source. The linear blending results in a set of clearly separated highlights while our reconstruction shows smooth and connected reflections of the light source. There, the difference to the ground truth is only marginal. Not only the highlights, but also the cast shadows in the scene are reconstructed with superior quality.

The second scene (Figure 6.9) contains examples for much more complicated transport paths. The scene is rendered in a set of environment maps. The figure demonstrates that we synthesize intermediate reflectance samples for all illumination directions in the hemisphere. The first column shows the result obtained from linearly blending between the input images, the images in the center column are synthesized based on the initial reconstruction R^{reco} with highlights added in. In the third column we present results rendered with our final reconstruction after regularization with image-based priors. The image-base priors remove noise and other artefacts that occur at shadow boundaries in the initial reconstruction due to imprecise shadow segmentation and labeling. The improved quality due to the priors is best visible under illumination with a single point light (first two rows). When illuminated with a complex environment map the noise in the initial reconstruction already averages out, and the expensive post-processing using image-based priors might be skipped, though the shadow boundaries are still slightly sharper in the third column.

We now analyze the differences between linear blending and our interpolation algorithm. The mirror sphere in the front reflects the environment and shows strong interreflections with the neighboring textured, slightly more glossy sphere. Extracting and warping the highlight maps, a sharp reflection of the environment is rendered with our approach (second and third column) while linear blending clearly produces artefacts revealing the original input samples even on the glossy sphere. With our method for shadow extraction, blending and reconstruction, we are able to correctly move and interpolate shadows avoiding the triple contours and banded shadows that are visible in all renderings using the linear blending. The shadow cast by the two spheres onto the postcard moves with the light source direction while the texture itself stays fixed.

The selection of different interpolation schemes derived from the warped shadow maps (Section 6.4.2) is powerful enough to even cope with non-trivial shadows as cast by the semi-transparent refracting spheres in the front right and in the back. They produce shadows tinted by the color of the spheres and furthermore exhibit an interesting intensity variation within the shadowed regions caused by caustics. Both the relatively diffuse caustic of the blue sphere in the front and the sharp caustic of the sphere behind the glass are faithfully reproduced with our method, moving smoothly. In contrast, they fade in and out in the linear blend.

6.5.1 Limitations

We observed some artefacts in the reconstruction of the shadow and caustics of the champagne glass. While the caustics have high spatial frequency, they are relatively dim. Our algorithm fails to detect them as highlights which need to be flowed, and therefore produces an incorrect interpolation. Improving on the segmentation algorithm, which is currently based on two manually selected thresholds f_{lit} and $f_{highlight}$, could help in this respect.

Another problem that is mostly noticeable at grazing angles is that illumination features, e.g. caustics or shadow edges, might move too far in image space between two neighboring measured light samples to be handled correctly by our interpolation scheme. If there is no sufficient overlap in the shadow regions the applied level-set blending cannot produce the expected smooth motion. Capturing more samples in these regions would help.

The combined effect is visible for the champagne glass under point light source illumination where the complicated structures within in the shadow are not as expected, but they do not show the banding effects visible in the linear blend either. When illuminated by a complex environment map, the artefacts are barely visible.

6.6 Conclusion

In this chapter, we have presented a novel and powerful interpolation framework for 4D reflectance fields. Our algorithm augments sparsely sampled reflectance fields by synthetic intermediate slices, simulating the displacement of highlights and shadows in the image plane as the direction of incident light changes. We have demonstrated that the augmented reflectance field can be used for creating realistic images both for extended light sources and for directional light.

The method interpolates highlights, cast and attached shadows in a plausible way, and is flexible enough to interpolate even non-trivial illumination effects such as shadows cast by semi-transparent or translucent objects. As we only use local pixel neighborhoods to determine whether a pixel is shadowed or not, but otherwise decouple its local appearance from the appearance of nearby pixels, we can precisely maintain texture information along shadow boundaries (such as on the postcard) without dragging it along as the shadow moves. Our method does not require any previous knowledge on the scene geometry and operates only on the reflectance field data. At the same time, the reduced number of samples makes the measurement significantly faster. We hope that these developments help to make reflectance fields applicable to a wide range of new applications in the future.

The output of our algorithm is currently a reflectance field of full resolution which consumes significant memory space. As the output is deterministically generated from the sparse input data, it is largely redundant. While it is not economical to re-run the algorithm for each rendered illumination situation, a first step towards successful compression

could be to exploit that most pixels are computed as linear combinations of the neighboring input images, and store their location in a compressed mask. This could also allow for faster rendering times, as the integration of linearly changing reflectance values can be sped up.

Due to the demanding computation for a full reflectance field with up to 4000 intermediate images, we currently apply the image-based priors in a post processing step only. In the future we plan to integrate image-based priors in every step of the pipeline to avoid errors in the labeling and in the estimated flow.

7 Implicit Relighting with a Bayesian Approach

7.1 Introduction

Explicit acquisition of reflectance fields for image-space relighting ensures high rendering quality, but it requires a carefully calibrated measurement setup under tightly controlled conditions. In this chapter, we propose an implicit method instead that learns the response of a complex scene to distant illumination relative to the response of a simple reference object such as a snooker ball. Then, we can render the scene in any illumination for which we have an observation or synthetic rendering of the reference object (see Figure 7.1). This technique as a whole is inexpensive – a feasible setup merely requires a hand-held light source, a mostly white, diffusely coated room, a snooker ball and a digital camera.

In a Bayesian approach, our algorithm takes into account the statistical properties of the sample data in order to find the most plausible image of the target object (maximum-a-posteriori estimate), given the image of the probe object in a novel illumination. This statistical criterion avoids artefacts that would occur in a direct reconstruction, for example due to noise. Specifically, we form a tradeoff between reproducing the probe image as faithfully as possible, and maximizing the prior probability of the illumination that might have given rise to this image, based on the distribution learned from examples.

We do not only express the $2 \rightarrow 2$ reflectance field in an implicit fashion, but also the distant light field illuminating the scene. Therefore, while an object such as the snooker ball or the classical light probe in form of a mirroring sphere serves well as a reference object, almost any object can perform this function, as long as pixel-to-pixel correspondence between acquisition and rendering can be maintained (see Figure 7.6 below).

In the nomenclature of Section 2.4.2, page 16, we change the image-spaced relighting pipeline as follows:

acquisition step: we simultaneously record the reflectance field slices R_i of the scene we are interested in, and reflectance field slices S_i for the reference object, illuminated by the same incident distant light field $L_i(\theta, \varphi)$ that illuminates the scene. In the learning theory terminology, this step might also be called a **training step**.



Figure 7.1: Four relighting examples (upper rows) as linear combination of 272 images, the coefficients being defined by novel images of a probe object (lower rows, left image of each pair) which are reconstructed with the sampled probe images (right).

rendering step: instead of providing an explicit, novel target illumination L , we use an image S of the reference object in that target illumination as input, and generate an approximative image I of the scene in that light. In the learning theory terminology, this step is a **prediction step**.

The **modeling step** can be omitted.

7.2 Implicit Relighting

In order to motivate the concept of implicit relighting, we recall Equation 2.23

$$\begin{aligned} I(x, y) &= \int_{\theta} \int_{\varphi} \sum_{0 \leq i < n} l_i L_i(\theta, \varphi) \cdot R(\theta, \varphi, x, y) d\varphi \cos \theta d\theta \\ &= \sum_{0 \leq i < n} l_i \underbrace{\int_{\theta} \int_{\varphi} L_i(\theta, \varphi) x \cdot R(\theta, \varphi, x, y) d\varphi \cos \theta d\theta}_{=: R_i(x, y)} \end{aligned}$$

and make two key observations:

1. Rendering a $2 \rightarrow 2$ reflectance field can be performed by evaluating a linear combination of input images R_i ,

$$I(x, y) = \sum_{0 \leq i < n} l_i R_i(x, y)$$

While the linear factors l_i depend on the relationship between the desired target illumination $L(\theta, \varphi)$ and the illuminations $L_i(\theta, \varphi)$ present at acquisition time via

$$L(\theta, \varphi) \approx \sum_{0 \leq i < n} l_i L_i(\theta, \varphi)$$

it is not necessary to explicitly know the angular distribution of either either, provided that the l_i can be computed in another way.

2. As the linear factors l_i *only* depend on the relationship between $L(\theta, \varphi)$ and the $L_i(\theta, \varphi)$, they will be the same for any reflectance field we want to render in $L(\theta, \varphi)$, as long as it was recorded in the same incident illuminations $L_i(\theta, \varphi)$.

We therefore propose to record $2 \rightarrow 2$ reflectance field slices $R_i(\theta, \varphi, x, y)$ of a static scene while placing a simple reference object close by, so that it is illuminated in the same fashion as the scene. We call its reflectance field slices $S_i(\theta, \varphi, x, y)$; they can be easily recorded within the same images that the camera takes for the scene. Then, for a picture S of the reference object in a novel illumination condition, approximate S as a linear combination of the S_i with factors l_i , and use the same linear combination to render the entire scene, as illustrated in Figure 7.2.

As an exact choice of the l_i is only possible if $L(\theta, \varphi)$ is in the linear span of the training illuminations $L_i(\theta, \varphi)$, we propose to approximate it in the least squares-optimal sense by minimizing

$$E(l_i, S) = \left\| \sum_{0 \leq i < n} l_i \cdot S_i - S \right\|^2 = \sum_{x, y} \left(\sum_{0 \leq i < n} l_i \cdot S_i(x, y) - S(x, y) \right)^2 \quad (7.1)$$

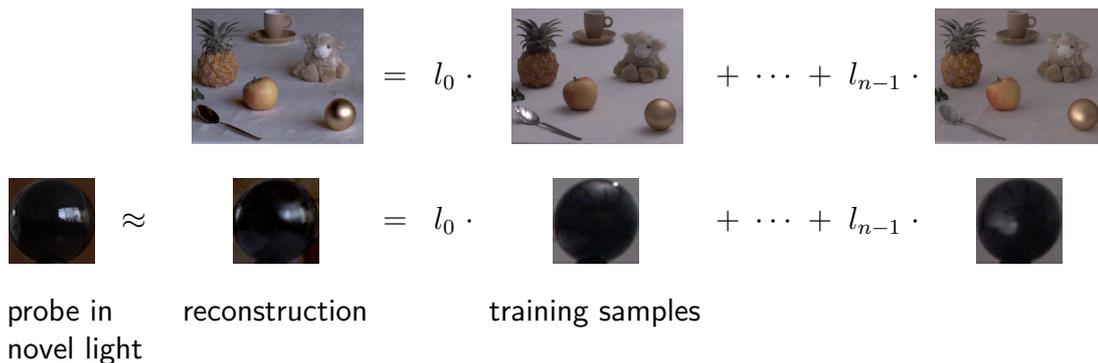


Figure 7.2: Relit scene in novel illumination as linear combination of $n = 272$ previously recorded images (top row). The coefficients l_j are found by reproducing an image of a probe object (bottom row), using a maximum-a-posteriori estimate.

The minimum of $E(l_i, S)$ can be found by solving a simple linear system using standard methods such as the pseudo-inverse matrix. In the following section, we describe a more appropriate technique that employs regularization.

7.3 Bayesian Relighting

Solving directly for the parameters l_i by a pseudo-inverse produces overfitting artefacts, as shown in Figure 7.3 (where $\eta = 0$): first, the images of the probe object are noisy, so the system attempts to reproduce this noise. Second, the samples do not span the full space of possible illuminations, so a least-squares reconstruction of the novel illumination would involve extreme coefficients l_i far from the convex hull of examples. This implies that noise in the sample images of the probe would be scaled with large factors l_i of opposite sign, causing actual noise amplification in the result.

Therefore, we take a maximum-a-posteriori approach (MAP, see the textbook by Duda et al. (2001)) to relighting: given an image S of the reference object in a novel illumination, we find the image I of the target object that maximizes the conditional probability $p(I|S)$ (posterior probability), based on an estimate of the prior probability p of lighting conditions from the sample set. We do not need to know the incident light field $L(\theta, \varphi)$ explicitly, but only in terms of a linear combination of sample illuminations. Based on the prior probability, a regularization parameter controls how conservative our estimate will be.

To estimate the prior, we perform a principal component analysis (PCA) on the set of reference object samples $(S_i)_{i=0, \dots, n-1}$: let $\bar{S} = \frac{1}{n} \sum_i S_i$, and \mathbf{A} be the matrix formed by the columns $(S_i - \bar{S})$. PCA is based on a diagonalization of the covariance matrix: $\mathbf{C} = \frac{1}{n} \mathbf{A} \mathbf{A}^T = \mathbf{U} \text{diag}(\sigma_i^2) \mathbf{U}^T$, where σ_i are the standard deviations of the data along the orthogonal principal component vectors \mathbf{u}_i given by the columns of \mathbf{U} . This

diagonalization is achieved by a Singular Value Decomposition (Press et al., 1992)

$$\mathbf{A} = \mathbf{U}\mathbf{W}\mathbf{V}^T \quad (7.2)$$

with a diagonal matrix $\mathbf{W} = \text{diag}(w_i)$, $\sigma_i = \frac{1}{\sqrt{n}}w_i$, and an orthogonal matrix \mathbf{V} . The reference image S can be written as a linear combination of the principal components

$$S = \sum_i c_i \mathbf{u}_i + \bar{S} = \mathbf{U}\mathbf{c} + \bar{S}, \quad (7.3)$$

where c_i are the linear coefficients. The estimated normal distribution of samples is, with a normalization factor ν_p ,

$$p(\mathbf{c}) = \nu_p e^{-\frac{1}{2} \sum_i \frac{c_i^2}{\sigma_i^2}} \quad (7.4)$$

within the linear span of examples. Since the coefficients of a linear combination of reference images S_i also describe the combination of light maps $L_i(\omega)$, p also captures the estimated probability density of light distributions within the span of L_i . With additive Gaussian pixel noise in the reference images S , the likelihood of an incident light map $L(\omega)$ producing S is

$$p(S|\mathbf{c}) = \nu_l \cdot \prod_{x,y} e^{-\frac{1}{2\sigma_N^2} (\sum_i c_i u_{i,x,y} + \bar{S}_{x,y} - S_{x,y})^2} = \nu_l \cdot e^{-\frac{1}{2\sigma_N^2} \|\mathbf{U}\mathbf{c} + \bar{S} - S\|^2}, \quad (7.5)$$

with a standard deviation σ_N and a normalization factor ν_l . The norm $\|\cdot\|^2$ denotes the sum of squared pixel differences.

According to Bayes' theorem, the posterior probability is

$$p(\mathbf{c}|S) \sim p(S|\mathbf{c}) \cdot p(\mathbf{c}), \quad (7.6)$$

which is maximized if a cost function given by the negative, rescaled logarithm is minimized:

$$E(\mathbf{c}, S) = \|\mathbf{U}\mathbf{c} + \bar{S} - S\|^2 + \eta \sum_i \frac{c_i^2}{\sigma_i^2}, \quad (7.7)$$

where $\eta = \sigma_N^2$ is a regularization parameter that can be used to control how conservative the estimate is supposed to be, which depends on the anticipated measurement noise and the properties of the sampled illuminations, such as their angular distribution and angular overlap. The more complete and smooth the basis of samples, the smaller an η we may choose without producing artefacts. Figure 7.3 illustrates the effect of different values of η in our system.

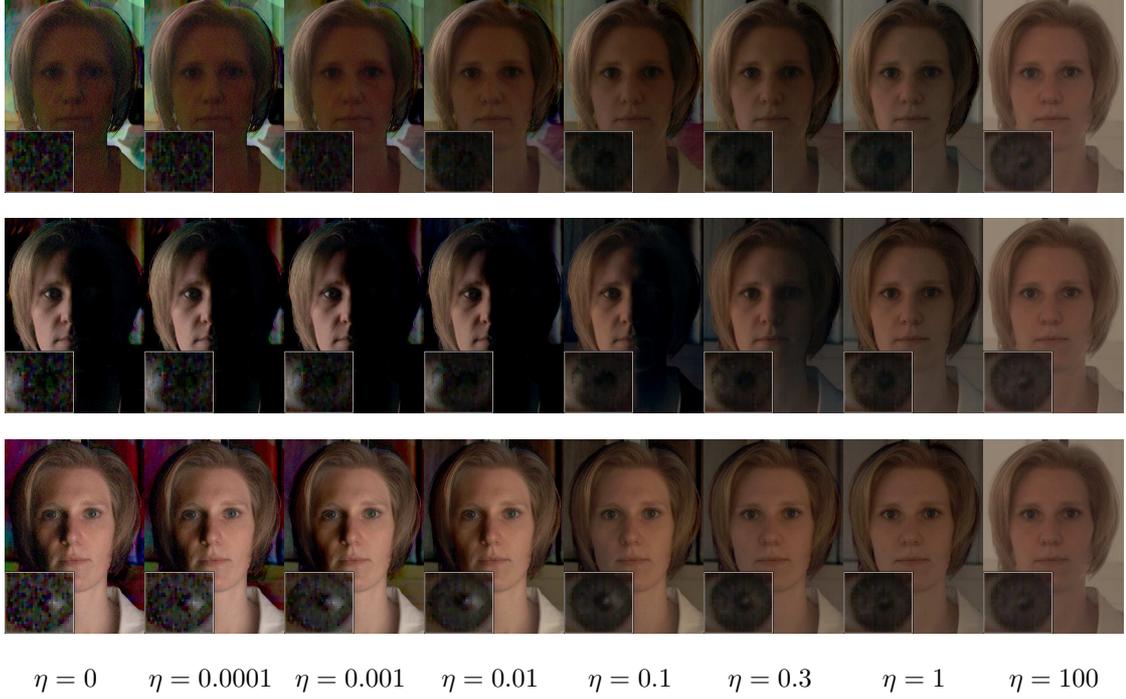


Figure 7.3: Each row shows a predicted image with close-up on the test subject's right eye. The parameter η is used for regularization in the renderings. The highlight in the eye and the shadow distribution demonstrate the tradeoff between a detailed lighting (low values of η) and low noise (high values η) in the rendering. For these renderings, $n = 625$ input images were used.

E is minimal if $\frac{\partial E}{\partial c_i} = 0$ for all i :

$$\frac{\partial E}{\partial c_i} = 2 \left\langle \mathbf{u}_i, \sum_k c_k \mathbf{u}_k + \bar{S} - S \right\rangle + 2\eta \frac{c_i}{\sigma_i^2} = 0, \quad (7.8)$$

which is achieved for

$$\mathbf{c} = \text{diag} \left(\frac{\sigma_i^2}{\sigma_i^2 + \eta} \right) \mathbf{U}^T (S - \bar{S}). \quad (7.9)$$

The conservative best fit can be rewritten in terms of the original basis, using $\mathbf{V}^T \mathbf{V} = \text{id}$:

$$S_{\text{MAP}} = \mathbf{U} \mathbf{c} + \bar{S} = \mathbf{A} \mathbf{V} \mathbf{W}^{-1} \mathbf{c} + \bar{S} = \mathbf{A} \tilde{\mathbf{c}} + \bar{S} \quad (7.10)$$

$$\text{where } \tilde{\mathbf{c}} = \mathbf{V} \text{diag} \left(\frac{1}{\sqrt{n}} \cdot \frac{\sigma_i}{\sigma_i^2 + \eta} \right) \mathbf{U}^T (S - \bar{S}). \quad (7.11)$$

Using the definition of \bar{S} , we obtain

$$S_{\text{MAP}} = \sum_i l_i S_i, \quad \lambda_i = \tilde{c}_i + \frac{1}{n} (1 - \sum_k \tilde{c}_k). \quad (7.12)$$

These coefficients l_i also provide the maximum-a-posteriori prediction for the target object image at a novel illumination:

$$I_{\text{MAP}} = \sum_i l_i I_i . \quad (7.13)$$

7.4 Experiments

For collecting images of the probe and target objects at different illuminations, we use an inexpensive setup with widely available equipment. Probe and target objects should be relatively close together to make sure that they are illuminated in the same way in the sense of Equation 2.21. Probe and target can be captured either in the same picture, as we did, or in separate pictures taken with two cameras. We used an Olympus C5050Z digital camera for still images at a resolution of 2576×1925 pixels, an Imperx MDC 1004 video camera at 1004×1004 pixels for the data set shown in Figure 7.3, and an HDRC VGAX high dynamic range 640×480 video camera, courtesy of IMS-CHIPS¹, for the face data set shown in the video that we presented at the Eurographics Symposium on Rendering 2005. For all cameras, images were captured in raw format, and a linearization and Bayer reconstruction were performed. The renderings in this chapter, as elsewhere in this thesis, are subject to an sRGB non-linear transform, approximating a gamma value of 2.2.

The illumination in our measurement was indirect light from the white walls and ceiling of a seminar room in our lab (Figure 7.4). Walking around the room, we illuminate different parts of the room with a hand-held HMI light source (Joker-Bug 800 by K5600). The method should work with any bright light source, and as the illumination may change during exposure, long exposures do not deteriorate the measurements for static objects. We avoid to hit the objects or the camera directly by using a reflector and pointing the light away from the measurement setup. In the seminar room, ceiling and walls were far enough to approximately satisfy the assumption of distant light.

While our approach does not require calibrated illumination with a known distribution $L(\theta, \varphi)$, and neither ambient light nor smaller objects or darker regions in the room affect the measurements, there are two issues to take care of: first, the incident light should cover as much of the sphere around the objects as possible across different measurements. Regions that were left out cannot contribute incident light during relighting. Second, the illumination patches, which essentially define the basis $L_i(\theta, \varphi)$ of light distributions, should be overlapping and smooth: if the scene is illuminated by point lights or by small patches of indirect light from the walls, novel probe images with specular reflections between those that were measured cannot be reconstructed, and the new light directions will be missed altogether. Therefore, we started off by illuminating large portions of the room from a larger distance in overlapping patches, and then lit overlapping sequences of smaller and smaller patches.

¹<http://www.ims-chips.de/>



Figure 7.4: Measurement setup for the still camera: an Olympus C5050Z digital camera records objects on a table which are indirectly lit by a hand-held spotlight pointed at the white walls, ceiling and floor. The reference object, a black snooker ball, is mounted on the small tripod next to the table.

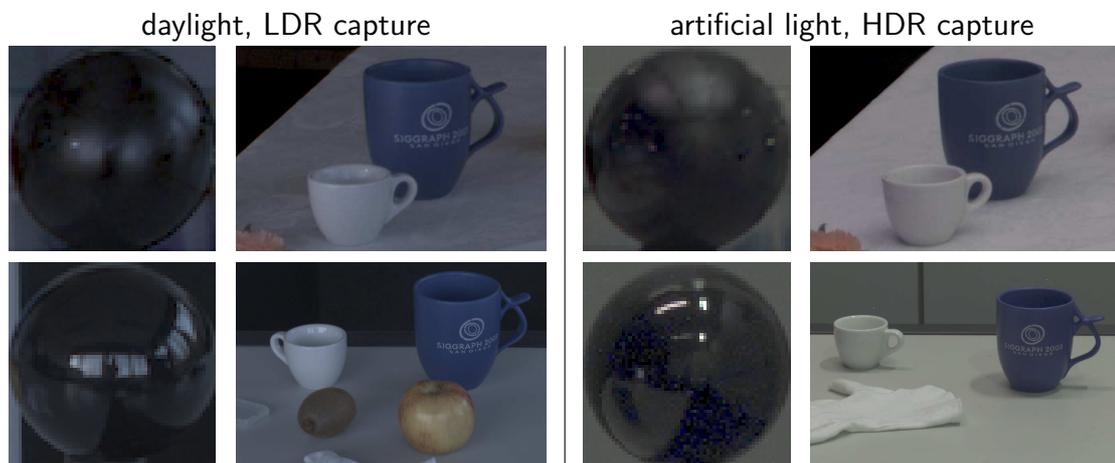


Figure 7.5: Reconstruction of two light situations (top row) and ground truth images (bottom row) of a scene in similar arrangement. The left situation is captured in low dynamic range in daylight, the right situation is taken as multi-exposure image in artificial illumination. While the predicted images approximate the lighting condition only with a sparse set ($n = 272$) of input images, they match the overall brightness and distribution of highlights and shadows.



Figure 7.6: Arbitrary reference object (left, 128×128 pixels), sparse sampling: an arbitrary part of the scene can be used as a probe object in the case of a fixed camera setup. The top row shows the reconstruction for sparse sampling ($n = 250$), the bottom row ground truth. While the sparseness causes blurred highlights and shadows, the reconstruction does not produce multiple blended shadow boundaries or comparable artefacts as point-based lighting sometimes does.

The reference object can be any object that is sensitive to illumination changes, as Figure 7.6 demonstrates. For most measurements, we chose to use a sphere, since the pixel-to-pixel mapping between sampled images and images of the object at novel illumination can be established easily due to rotational symmetry, without fixing the object to the camera. For this mapping, which is needed to find the linear combination of samples that reproduces the novel probe image best, we select the sphere by a bounding box in the images, and apply a scale and translation operation, assuming orthographic projection of the sphere. We achieve good results scaling the sphere to 64×64 pixels, and masking the non-sphere parts of the images.

In capturing incident light distributions, the dynamic range of the camera is an important issue. Most authors record images of a metallic sphere with high-dynamic range imaging to avoid saturated – and therefore underestimated – highlights on the sphere. In order to reduce the radiance at highlights, we prefer to use a black snooker ball, which reflects only a small portion of the incident light to the camera (Tchou et al., 2004): according to the Fresnel formulas, the specular reflectance of the non-metallic snooker ball is 1.0 at tangent directions, and falls off rapidly to a value of 0.04 in the center (at an index of refraction of 1.5). As a result, our probe object produces relatively dim specular reflections that are likely to be within the dynamic range of a digital camera in images that, at the same time, capture the target scene appropriately.

So, more precisely, our implicit relighting pipeline comprises the following two steps:

Training Step: Record a set of n images at different illumination with fixed cameras and static objects, define a bounding box around the probe sphere in the first image, crop

and scale the probe in all images, perform a PCA on the probe images and store the result.

Prediction Step: Given a photograph of the probe in a novel illumination, we crop and scale the probe again from the image, compute c , \tilde{c} and l_i (Equations 7.9,7.10,7.12) and form the weighted sum of sampled images I_i (Equation 7.12).

The variation in overall brightness in our sample sets turned out to be sufficient to cover the variations in novel illuminations without rescaling.

Unlike most previous methods, our linear combinations may involve negative coefficients l_i : in the classical light stage approach (Debevec et al., 2000) and most subsequent methods, linear coefficients are non-negative, since they are weights proportional to the incident radiance in the neighborhoods of discrete light directions. Matusik et al. (2004) enforce constraints on the model coefficients in the reconstruction of R , Mohan et al. (2005) run a constraint optimization directly on the l_i .

In our setting, negative coefficients arise from the overlapping base functions. Still, they do not imply physically invalid results: consider an image with two lights A and B , and one with A only. The difference image reproduces the situation with B only, and all resulting pixel values represent valid positive radiances. However, negative color values may still occur in our least-squares framework within the range of approximation errors. In a second iteration, we therefore alleviate this problem by fitting against an image consisting of the inverse of negative result pixels on the probe (with a smaller value for η), and adding the resulting coefficients to the previous results.

7.5 Results

Figure 7.1, 7.2 and 7.8 show images of objects that were relighted with our system, demonstrating the high spatial detail that can be achieved with still camera measurements.

The training step for these images contained 272 sample images, taken in about 30 minutes. The computation time for the PCA on the probe pictures took about 20 seconds per color channel on a PC with 3 GHz Intel Xeon Processor. Determining the coefficients l_i for some target image S takes less than 2 seconds, and reconstructing an image takes between 1 (for 64×64 pixels) and 16 seconds (for 708×560 pixels). These numbers are for the data set from Figure 7.2, but are comparable to the others. The timings are performed after transferring the pictures from the camera's Compact Flash card, and reconstructing the Bayer pattern in each of the input images.

Figure 7.8 shows the wide range of material appearances which are captured by our approach: cloth (napkin), polished metal (cutlery), glossy objects (orange, plate), transparent objects (wine inside glass), reproducing near-field caustics (jelly) and simulating sub-surface light transport (orange). All of these are plausibly relighted. For a ground truth comparison, we reproduced lighting of two scenes, as shown in Figure 7.5.



Figure 7.7: Results of relighting a test subject from $n = 75$ still images. Some artefacts arise because of movement (see edges), but lighting remains realistic.

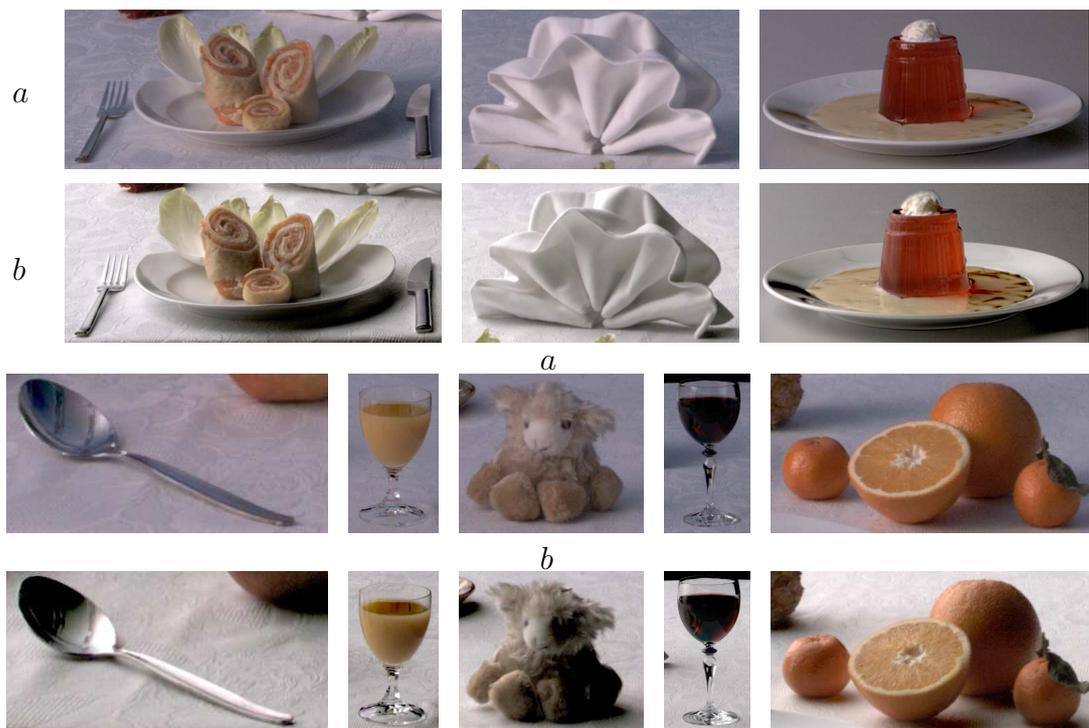


Figure 7.8: Various materials relighted, $n = 272$. The rows *a* show reproduced natural lighting (as in Figure 1, second column), the rows *b* show synthetic lighting of a directional dominant light source in an otherwise totally dark room. The objects are from four different data acquisitions: (i) dish, napkin and wine, (ii) jelly, (iii) spoon, sheep and juice, (iv) oranges.

For live objects, such as human faces, the acquisition time is an issue. Figure 7.7 demonstrates that from $n = 75$ sample illuminations, which were captured in about 8 minutes, interesting effects of human skin structure can already be captured. However, especially for the perceptually important visual properties of human eyes, a setup where the eyes are open is more appropriate; therefore, we also performed experiments with two faster cameras. In one experiment, the test subject (Figure 7.3) was recorded for 25 seconds with a video camera, yielding 625 input images which allow us to recreate even specular highlights in the eyes. In another experiment we used an HDR video camera yielding 1000 frames in a comparable time.

Even though the cameras gave us abundant image data in a short period of time, we had to record for about 25 seconds to sweep the light source's cone over the wall manually, covering a sufficient set of light conditions. Residual movements of the test subject who was recorded with the still camera (Figure 7.7) are less prominent in the video acquisition setup (Figure 7.3) due to shorter recording intervals.

In the background of the synthetic images, behind the target objects, our technique tends to produce ghost images that show the experimenter and the spotlight, as the experimenter becomes part of the distant incoming light environment.

Although it is designed for relighting with natural illuminations, our approach can also be used for synthetic relighting, based on renderings of a snooker ball. For Figures 7.8 (row *b*), and 7.3, we rendered a sphere with Phong BRDF and an additional Fresnel term for a refractive index $n_{\text{refract}} = 1.5$. The Phong exponent gives us an easy control of the distribution of incoming light; by choosing a low exponent, an extended light source is simulated. Synthetic images of the ball created with a ray-tracer or global illumination techniques could be used as well.

For illumination design, as shown in Figure 7.1 (right), the user draws patterns of incoming light with standard imaging software into an image of the probe object, which is then reconstructed by our Bayesian method for transferring the lighting on the target. This is unlike previous methods (Anrys et al., 2004a,b), where the lighting was designed in the target image directly. For practical applications, both approaches are useful, but they address different design purposes.

7.6 Conclusion

The contributions of the method we presented in this chapter are a new theoretical approach for relighting, and a low-cost system that requires no light stage or other sophisticated setup or equipment. From a maximum-a-posteriori approach, we have derived a simple mathematical formula which makes the relighting algorithm easy to implement. We hope that our technique helps to make relighting more available to a broad range of users.

While our implicit approach does not exercise the tight control on the prefiltering caused by the incident illumination that our setup from Chapter 4 provides, its extended light

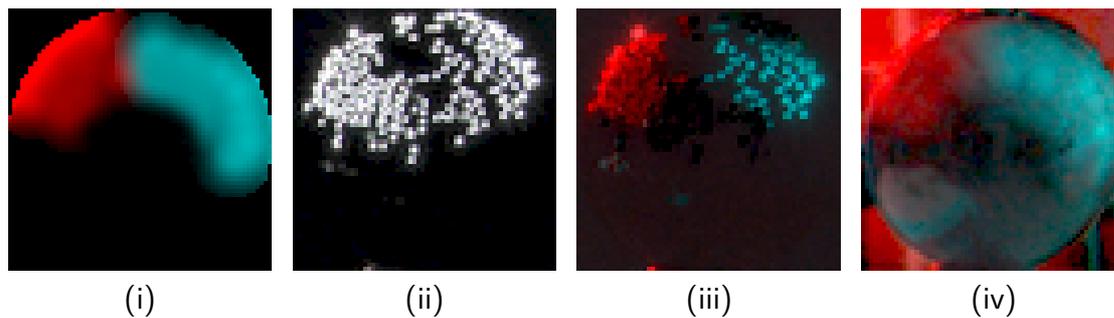


Figure 7.9: Comparison to impulse-response sampling: in order to reconstruct a target (i), sampling $n = 271$ non-uniformly distributed points (ii) inside the target image’s active area creates holes and non-smooth artefacts (iii), even though we sampled light directions favorably. In contrast, our method with extended incoming light sources (iv) gives a continuous reconstruction.

sources nevertheless tend to mitigate undersampling problems, as Figure 7.9 illustrates. As a result of these extended light sources, however, our method tends to blur highlights and shadow edges, as can be seen in Figure 7.6.

In animations with moving directional light sources, our method produces smoothly moving, but slightly broadened highlights in comparison to appropriate sampling.

As a conceptual advantage to most techniques, our implicit approach learns the mapping between light probe and output images directly, rather than investing in the estimate of intermediate information, such as incident light (Debevec, 1998) or reflectance (Matusik et al., 2004). We have proposed a new, more general notion of a light probe object, which makes the method interesting for new applications in fixed camera setups. We presented a result employing toy figures for that purpose.

Our method fits seamlessly into existing acquisition pipelines that measure incident light distributions explicitly, as the mapping of the light distribution to the snooker ball is straight-forward. However, this is not our primary goal, since we propose a different measurement process for determining illumination that is equally simple as the conventional method of capturing a mirror sphere. It is easy to improve the speed of the illumination sampling by technical means, as the measurement setup is uncalibrated.

From a given set of sampled illumination conditions, our statistical approach enables us to predict a relighted image in an optimal sense without explicit knowledge of the object or lighting properties, making it a consequent implementation of learning-based computer graphics.

8 Towards Passive 6D Reflectance Field Displays

8.1 Introduction

In the previous chapters, the acquisition and computational improvement of image-space reflectance fields were our concern. Now, we will turn to rendering. The reflectance equation for image-space relighting (2.21) prescribes a direct rendering scheme as an evaluation of the integral within in software. In this chapter, we will instead describe techniques for a rendering in hardware.

Specifically, we will measure incident illumination and render an image all in one passive display device akin to stained glass. Using optical components for the computation, we can actually convert the light falling on the back plane of our display into the rendering displayed on the front (see Figure 8.1).

Our designs have been inspired by autostereoscopic displays that produce view dependent images by placing a set of lenses on top of a pattern encoding the outgoing light field. In contrast to autostereoscopic displays, the output of our display depends on the incident illumination (see Figure 8.2). We describe an optical setup which changes the appearance depending on the illumination angle, and define a mechanism that superimposes the contributions from different angles, effectively computing an expensive integration with optical means.

Using a simple configuration of optical components – mainly lenses and lenslet arrays – we discretize the 4D space of incident light rays, and embed it into a 2D plane in which a printed pattern performs a direction-dependent modulation. Another set of lenses then shapes the outgoing light field both in the angular or the spatial domain. The pattern can be quickly modified, for the efficient re-use of optical components. Our lens-based designs are considerably more light efficient than designs using slit masks.

In particular we propose the following designs and prototypes:

- *Design I*: a 4D display for $2 \rightarrow 2$ reflectance fields in image space. It renders a 2D image depending on the 2D distant incident light field for a fixed observer (Section 8.2.1),
- *Design II*: an improvement of Design I with improved precision (Section 8.2.3)

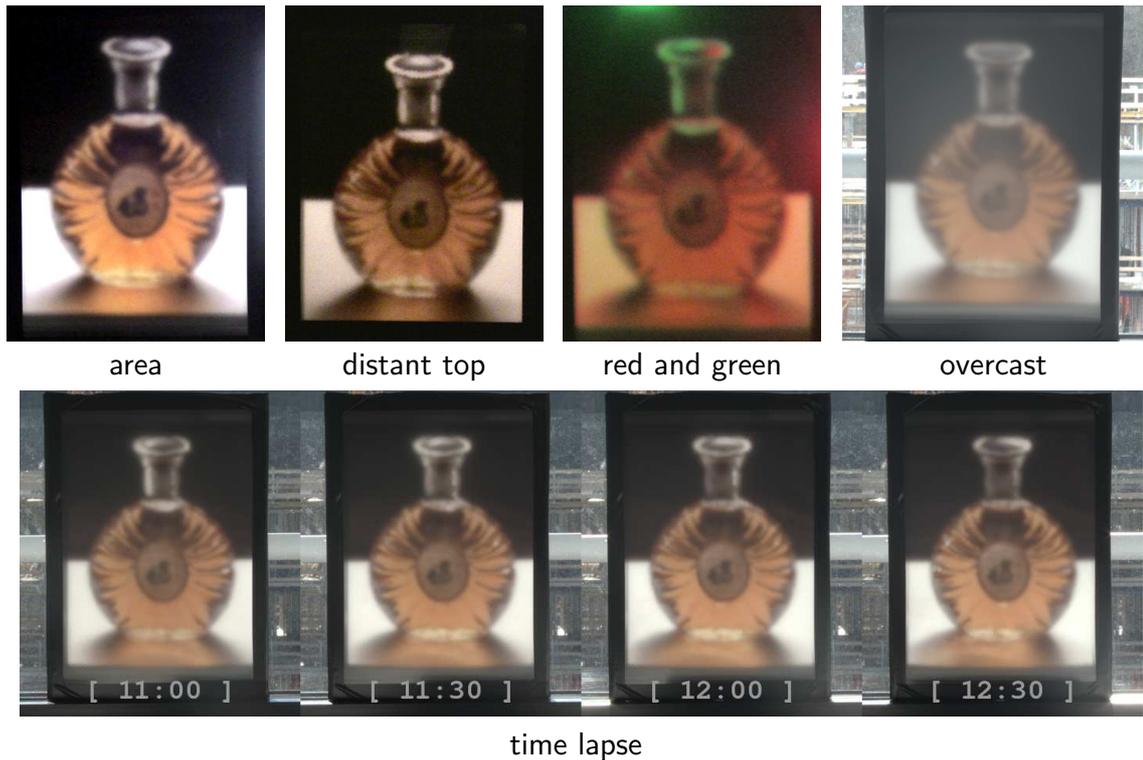


Figure 8.1: A 4D reflectance field display (Design I) illuminated from behind with different incident directions and light source types. For the “time lapse” and “overcast” results, we taped the display to a window and recorded pictures with direct sun illumination and in overcast sky. Note how the highlights on the bottle cap and the shadows move as the illumination direction changes, and how the shadow boundary smoothes for the area light source.

- *Design III*: a 4D display that at one location modulates the 2D incident distribution, projecting out a different 2D angular distribution for each incident angle (Section 8.3.1), and
- *Design IV*: a 6D proof of concept for a $2 \rightarrow 4$ reflectance field display, which combines multiple 4D displays to modulate the incident 2D illumination into a 4D emitted light field, varying both in the angular and the spatial domain (Section 8.3.3)

Both Design I and Design II have deficiencies with respect to contrast and transmissivity, which we will discuss in Section 8.2.2. Still, they demonstrate the concept and are fully applicable in sunlight.

Some of the possible applications of this technology are time-variant transmission in window sheets forming fully passive mood lights using the sun as light source. The displays can show time-variant pictures for natural illuminations. They can act as physical representations of captured 4D reflectance fields and environment mattes, passively rendering

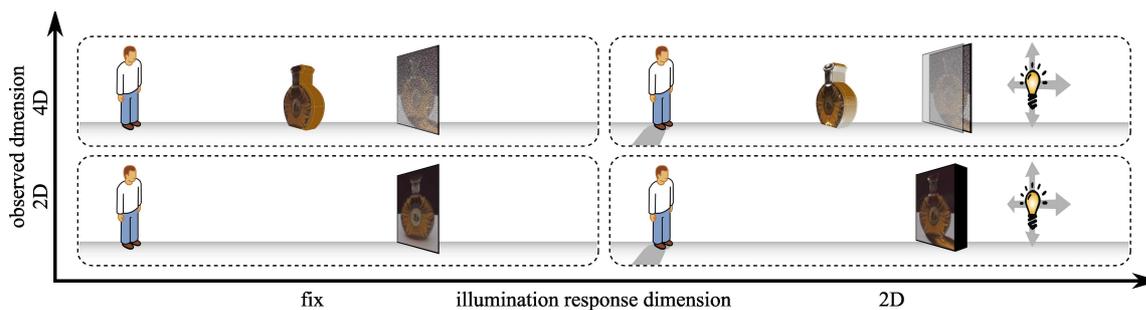


Figure 8.2: Display technologies: 2D displays (bottom left) are independent from viewer and light. A 4D light field display (top left) creates a parallax illusion of an object floating in space, while a 4D reflectance field display (bottom right) changes its flat appearance as its illumination changes. Combined, they become a 6D display (top right): a display which creates an illumination-dependent light field rendering.

the object of interest in the ambient illumination using only optical components for the rendering process.

8.1.1 Related Concepts

Flattening of High Dimensional Visual Data

High dimensional data structures play an important role in our everyday life: in images, volumes, light fields or reflectance field data sets. However, many optical visualizations and recording techniques are limited to a 2D structure. Therefore, methods have been presented in the past which address this problem by flattening the high dimensional data, embedding it in a planar, 2D representation.

Integral photography (Lippmann, 1908) is an early approach which records a 4D light field on a photographic plate. The main concept is adding an array of lenses to the plate, both discretizing the spatial coordinates (x, y) and embedding the angular domain (θ, φ) within. In recent years, this approach has been successfully applied to digital recording of light fields (see Section 3.2.1).

Light Field Displays

Planar encodings of light fields are since the days of integral photography closely coupled to the development of displays which create a 3D impression by projecting a light field into space. Nakajima et al. (2001) described a lens array on top of a computer display for a 3D viewing experience. In 2004, Matusik and Pfister presented an end-to-end system which records a 3D light field, streams it over the network and then displays it

on a lenticular array screen. Their article also gives a good overview of current multi-dimensional display techniques. Javidi and Okano, F. (Ed.) (2001) discussed a range of related techniques.

Illumination-Variant Displays

The display by Nayar et al. (2004) measures the distant room illumination, approximated as environment map, and interactively renders an image in this illumination. Koike and Naemura (2007) propose an extension towards emitting a light field in a similar fashion. Both displays are electronic and rely on software and hardware evaluating the illumination and rendering the reflectance field. Scharstein et al. (1996) obtained a patent on a device which is passive: it employs optics in order to create a numeral display of the current time. This is achieved by encoding a pattern in a slit mask so that natural sunlight direction produces different symbols. However, this construction inherently blocks the majority of incident light rays. Mann (1995) describes a setup for holographic acquisition of the reflectance field of a real scene.

In this work, we also follow a passive approach to illumination-variant displays, but, in contrast to Scharstein et al. (1996), we use lenses and colored patterns, thus using a larger portion of the available light for a higher contrast display of more expressive patterns, in contrast to Mann (1995), we maintain a direct, local reflectance encoding that works for measured and synthetic reflectance fields alike.

8.1.2 Overview

The display types we propose in this chapter modulate the incident light field both spatially and angularly. We assume the incident light to be distant, and are either concerned with rendering a 4D $2 \rightarrow 2$ reflectance field (see page 15) or a 6D $2 \rightarrow 4$ reflectance field (see page 14). Our displays implement the reflectance equations for distant illumination (2.18, 2.21) physically, using only passive optical elements. They operate on real-world illumination incident from behind the display and transmit a controlled, modified light field.

We present a low-resolution 6D reflectance field display in Section 8.3.3, and reduce the dimensionality of the reflectance field for the other prototypes by creating the same image for all observer positions (Section 8.2.1), or causing only angular but no spatial variation (Section 8.3.1).

All of the different designs make use of common building blocks, which are inspired by previous integral photography approaches. The high dimensional reflectance field function is first discretized and then flattened into a 2D plane, trading the spatial resolution in the plane for encoding both the angular and the spatial variation in various dimensions (see Figure 8.3).



Figure 8.3: After discretizing, a 4D pattern is embedded into a 2D image (left): macro-pixels capture the variation along the spatial coordinates (x, y) , within, micro-pixels vary according to (θ, φ) . To the right, the cap of the bottle shown in Figure 8.1 is displayed. (x, y) are distributed according to the hexagonal lattice arrangement of the lenslet array.

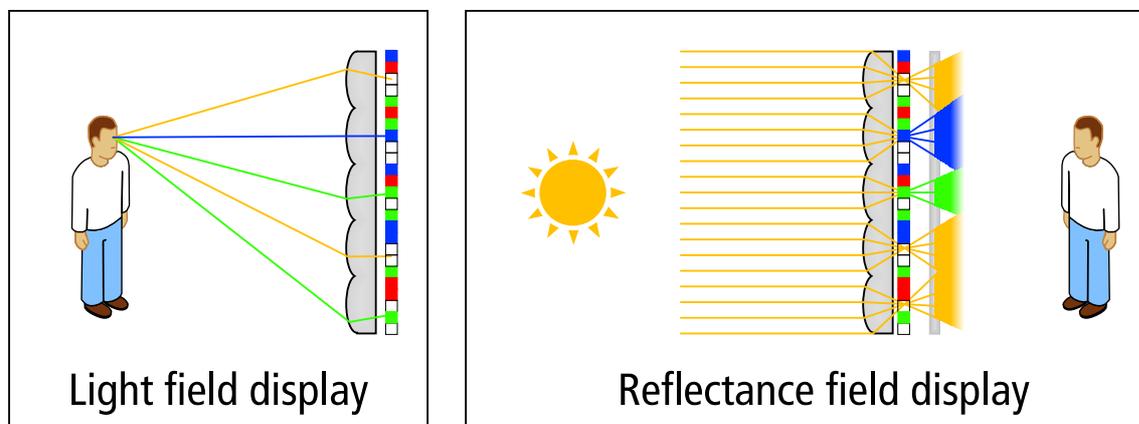


Figure 8.4: Lens array technology has been used by the integral photography approach to modulate the light field emitted by the display. We propose a modulation according to the illumination angle.

Each location p within this plane encodes exactly one sample of the reflectance field. Each point is responsible for the modulation of the light transport for one particular incident direction (θ_o, φ_o) and one outgoing direction (θ_i, φ_i) in one particular location (x, y) in image space. The modulation is carried out physically by letting the incident light field shine through a printed transparency. Printing a different pattern will realize a different reflectance field.

The second common component of our designs is a set of lenses or lenslet arrays which have the task of discretizing and mapping the incident and the reflected light field to the corresponding locations in the plane.

8.2 Observer-Invariant displays

On the way to a full 6D display, we introduce observer-invariant displays (Design I and Design II), that output pictures varying with the incident illumination, but display the same 2D picture for any observer position, providing a 4D viewing experience. Extensions to observer-variant displays will be discussed in Section 8.3.

8.2.1 Design I

In order to flatten the 4D reflectance field into a plane, we discretize the spatial dimensions (x, y) into macro-pixels, and, within each macro-pixel, encode the angular variation along (θ, φ) in micro-pixels, as illustrated in Figure 8.3.

For guiding the incident light directions to be modulated by the correct micro- and macro-pixel the incident light field is discretized and mapped to this planar representation as well. We apply the same concept as used in integral photography, and use a lens array for this task (see Figure 8.4). While in traditional light field displays the lenslet array dispatches the different 4D *viewing* rays into a 2D plane we use them to guide the *incident light* rays to specific locations depending on the light direction. In our case, each lens corresponds to a final macro-pixel, i. e. it will correspond to one pixel as seen by the observer. Assuming a distant, parallel light beam, the lens focuses the light onto a single spot in its focal plane behind the lens. The location of the spot, however, will move depending on the angle of incidence, which is exactly the behavior required to produce the mapping.

The printed transparency is placed in the focus plane. It modulates the color and attenuates the intensity only of that light beam which hits the lens from the corresponding direction. After the modulation the beam again diverges, illuminating a pixel-sized spot on the diffuser surface.

Figure 8.5 illustrates this design.¹ A parallel light beam hits an array of lenses (A), which focuses the light on a plane with an embedded pattern (B), flattening the 4D

¹The ray diagrams are to be read from left-to-right from the light source to the observer. For clarity,

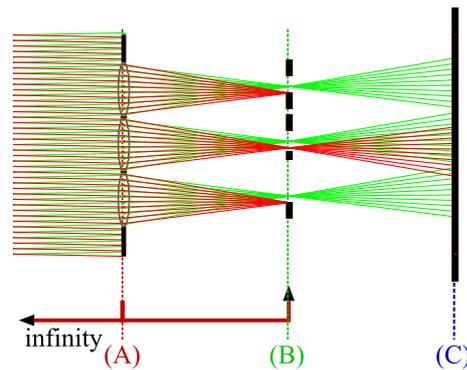


Figure 8.5: Design I : As the incident light angle changes, the transmissivity of pixels changes. For sufficiently small angles, the individual contributions on the diffuser plane overlap.

field. Depending on the angular direction and the spatial position, the transparency then modulates the incident light beam with a $2 \rightarrow 2$ reflectance field $R(\theta, \varphi, x, y)$.

In order to provide a view independent experience, a diffuser surface (C) is added. It ensures that the modulated beam can be observed from a wide range of viewing angles. Furthermore, it physically integrates over all incoming directions of one macro-pixel as illustrated in Figure 8.5. Thus, the setup physically evaluates the reflectance equation for image-space relighting (Eq. 2.21) in real-time.

8.2.2 Results for Design I

We have implemented a prototype display, which consists of five components (Figure 8.6): an inexpensive hexagonal lens array (Fresnel Technologies Inc., 20.77 lenses per cm^2 , focal length 3 mm), a printed transparency (using an Epson Stylus Pro 4800 printer set to 2880×1440 dpi resolution), a plastic spacer, a diffuser sheet (Intelliccoat Technologies DMBF5UV), and a selectively transmissive sheet (see below). In order to improve contrast, we printed a mirrored version on a second transparency and mounted them back-to-back, which is a technique occasionally found in commercial printing. Transparent plastic sheets maintain equal distance between the transparency and the diffuser sheet. The distance is chosen to level sharpness and anti-aliasing. We encoded the reflectance field of a bottle into the pattern, which is shown in Figure 8.3. Details of this process will be found in Section 8.4.

we always draw a 2D cut through the device and align the spatial location y with the vertical axis, hiding the x direction. In our drawings, we indicate conjugateness for relevant planes by drawing a colored bracket, so that a dotted line indicates the plane where the lens lies, and two arrows indicate the conjugate planes. In cases where one plane is the focal plane of the lens, we label the opposite arrow with “infinity”. While in our experiments, we modulate the light field with different colors and varying attenuation, we illustrate the patterns in the ray diagrams either fully transparent or fully opaque in order to prevent additional visual clutter.

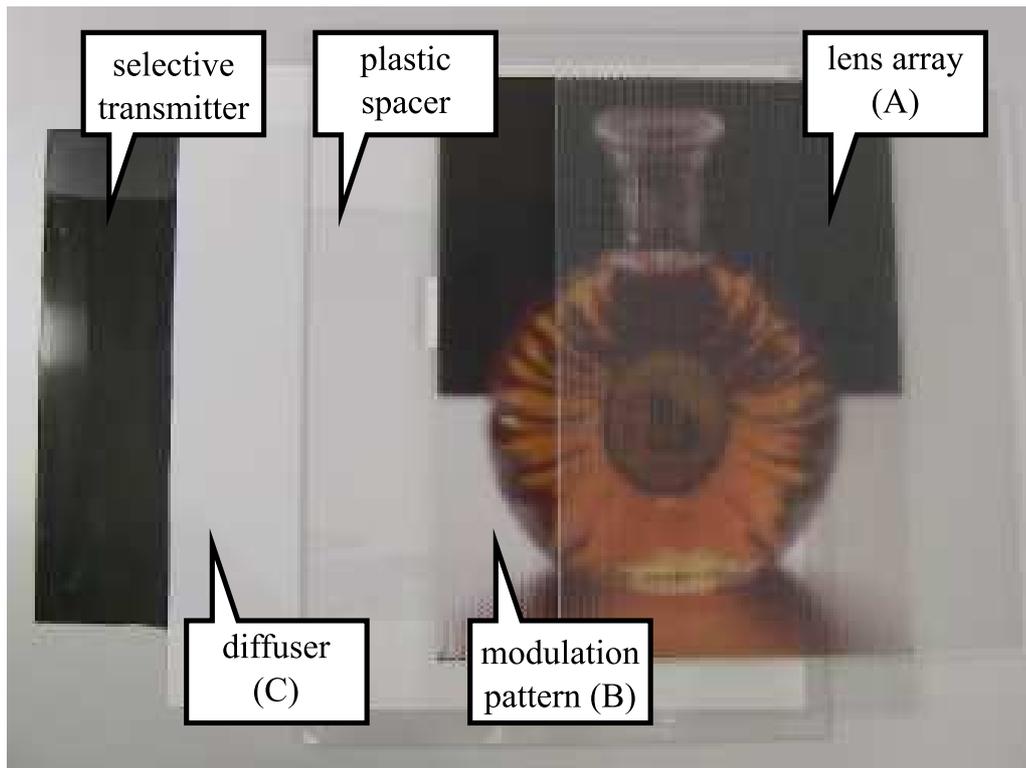


Figure 8.6: 4D construction for the display from Figure 8.1, with separated layers. (Design I)

(a)	(b)	(c)	(d)	(e)	(f)
transmissivity	60 %	28 %	18 %	61 %	37 %
contrast in dark room	223 : 1	80 : 1	84 : 1	222 : 1	211 : 1
contrast with front ill.	88 : 1	2.9 : 1	5.5 : 1	5.7 : 1	12 : 1

Figure 8.7: Design I: Transmissivity and contrast experiment: background illumination (a), display with high-contrast pattern without any diffusers (b), with Intellicoat diffuser (c), Intellicoat and 3M (d), holographic diffuser (e), holographic diffuser and 3M (f). Top row: light box in dark room, bottom row: front illumination added.

Unfortunately, the diffuser screen does not only scatter the illumination along the intended light path, i. e. from the back, but integrates all light from the viewer side as well. This drastically reduces the achievable contrast. This effect can be partially reduced by adding a “Notebook Privacy Computer Filter” by 3M to the observer side. It is a thin, transparent sheet, which largely reduces transmission for non-normal viewing angles. Still, the best result is achieved with either direct sunlight as back-illumination, or with an artificial light source in an otherwise darkened room. Preliminary experiments with a 30 degrees light shaping holographic diffuser screen from BFi OPTiLAS suggest possible efficiency and contrast gains (see Figure 8.7). The best material combination will heavily depend on the specific application case.

Figure 8.1 summarizes the results obtained in various settings, e.g. illuminated with point and area light sources in a darkened lab. The figure further shows the display mounted to an office window for different sky light conditions. As the light direction changes, highlights and shadows of our rendering move in real-time as a result of the optical computation. The point light sources produce crisp and correctly moving highlights. Illuminating with a red and a green light bulb, one sees how the display integrates simultaneously over different light source directions, resulting in separate red or green highlights but rendering the table top in yellow.

The contrast in the window setting is slightly reduced. The overall image resolution is limited to the resolution of the lens array. One can further make out some small blur in the overcast illumination and in the area light source setting. As explained in the next section, without a corrective lens layer, the rendered images tend to shift slightly on the diffuse plane as the incident light direction changes.

8.2.3 Design II with Correcting Lenses

Figure 8.5 also demonstrates one shortcoming of this simple design. On the diffuser, the modulated beams of one macro-pixel are not perfectly aligned, but only overlap partially. In the worst case, even contributions from neighboring macro-pixels are combined. This can be corrected by an improved design (Figure 8.8), where a second array of lenses is introduced in plane (B). Their task is to project the image of the respective main lens on the diffuser surface. This implies that planes (C) and (A) are conjugate with respect to the corrective lens arrays, and their focal length must therefore be half the focal length of the main lens. If the main lenses are plano-convex, we can re-use parts with the same specification in (B), putting them back-to-back and embedding the pattern in-between.

This additional lens array guarantees a one-to-one mapping between the macro-pixel observed at (C) and its corresponding lenslet in (A). Overlap to a neighboring pixel can only occur if the diffuser plane is moved, or if the incident angle is so steep that the beam through one lens in (A) will hit the neighboring macro-pixel in (B) (see Figure 8.8).

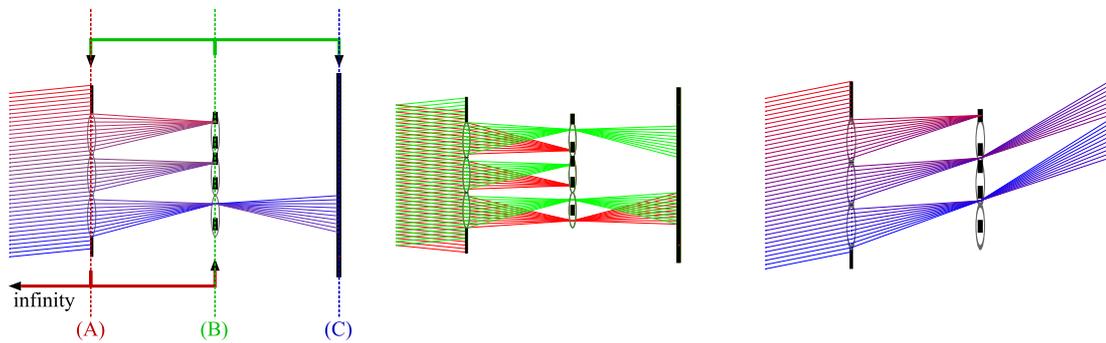


Figure 8.8: Design II: Adding a lens array with half the focal length of the main lenses in the pattern plane (B) keeps the modulated light cones from moving (left diagram), generating a precise overlap on the diffuser (center). For steeper angles, a wrap-over into the next lenslet can still occur (right).

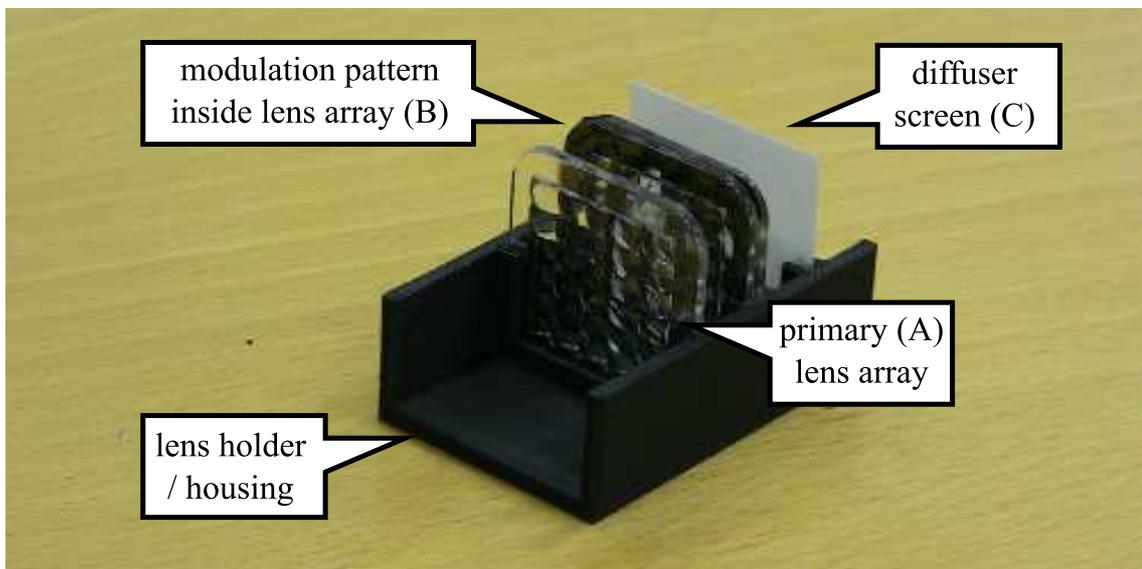


Figure 8.9: Design II for a small display, with top half of housing removed.

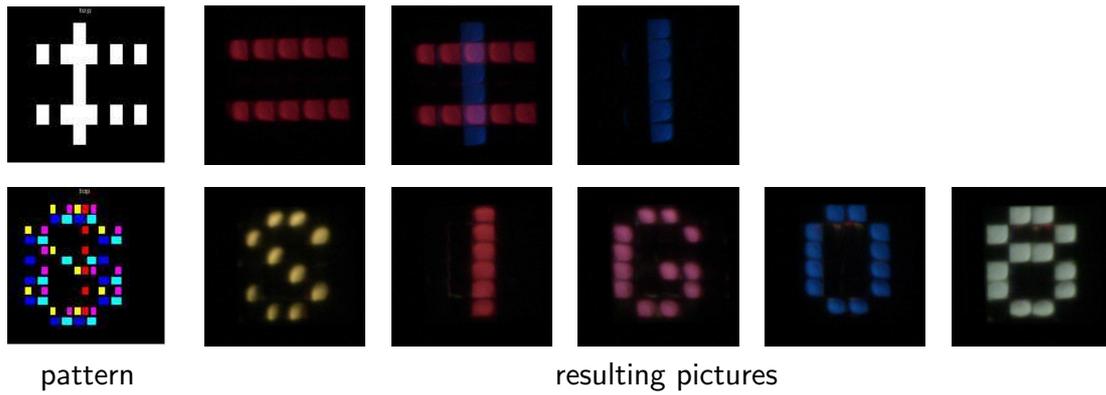


Figure 8.10: Demonstration of Design II. Top: a simple pattern, demonstrating the exact alignment of the partial pictures when illuminating with a blue light source from the right and a red light source from the left. Bottom: a more complicated pattern, illuminated by white light from five different directions.

8.2.4 Results for Design II

The second prototype display includes this corrective lens array. Since it involves more parts of higher optical quality and precision, we have manufactured a prototype of smaller resolution.

The setup is depicted in Figure 8.9. This time, we use a plano-convex lens array with $24\text{ mm} \times 24\text{ mm}$ edge length and 6×5 lenses of a focal length of 10.5 mm as primary lenses in plane (A), and two lens arrays of that type as correction layer in plane (B). They are made of glass by Moulded Optics GmbH. We use the same diffuser as in the first prototype. In order to hold the lenses, the pattern and the diffuser precisely in place, a custom-built lens holder has been printed in plastic with a 3D printer. This lens holder also blocks unwanted stray light. As the spatial resolution is much lower than in Design I, we have encoded simpler patterns.

Again, we performed experiments with hand-held light sources in a darkened room. The results are shown in Figure 8.10. The first row demonstrates the precise alignment that can be achieved with the corrective lenses: we have created a black and white pattern which shows two horizontal bars when illuminated from the left, and a vertical bar when illuminated from the right. Using red and blue light for the respective directions demonstrates that the individual contributions add up, and the blue and right pixels are precisely aligned. Figure 8.10 also shows results for a colored pattern which encodes letters for five different illumination directions. When illuminated by white light from different directions, they create the picture of the letters “SIG08”.

Some vignetting is noticeable at the borders of the macro-pixels but otherwise the light efficiency is rather good since the light path for every incident direction is controlled by lenses rather than blocking apertures as in the light dependent display of Scharstein et al. (1996).

8.3 Observer-Variant Displays

In the previous section, we have introduced illumination-variant displays which generate planar pictures only. Now, we will show how they can be extended to display different pictures for different viewing angles.

8.3.1 Design III

First, we limit the spatial variation to a single point, describing a design for a display which implements a 4D reflectance field parameterized by $R(\theta_{in}, \varphi_{in}, \theta_{out}, \varphi_{out})$. For different incident light directions, it projects out different angular light distributions. Similarly to the previous design, this effect is achieved by mapping each incident light direction into a different spatial location. For each location a different output distribution will be generated by modulation. The proposed design is depicted in Figure 8.11.

Since the design describes just a single macro-pixel, we put a single big lens in plane (A) where we had the lenslet array in the previous design. The function of this main lens is to map each incident light direction to one particular location on the plane (C). In the example of Figure 8.11, the incident directions are discretized into three bins which are mapped to three meso-pixels in plane (C).

Instead of modulating the beam in this plane, as we did before, the beam is widened by an array of small lenses in (C). They project the image of the main lens (A) onto the modulation plane (D). For all incident light directions which hit the same meso-pixel this area will be the same, producing the same output distribution. The outgoing distribution is controlled by modulation in plane (D). Each micro-pixel modulates the light for one output direction. As different patterns can be placed beneath each meso-pixel lens, the modulation is different for each incident direction.

The described design does everything necessary to implement a $R(\theta_i, \varphi_i, \theta_o, \varphi_o)$ display. As before, we add additional lenses to limit the shifting and the rotation of the outgoing distribution which currently slightly changes with the incident angle. We add a field lens (B) in plane (C) with the same focal length as the lens in (A), in order to suppress vignetting. It also rotates the incident light cone at (C) in such a way that it aligns the optical axes of all directions within one meso-pixel. Further, we add a lens array in the pattern plane (D) to project the single spot in the meso-pixel (B) onto a single spot in plane (E). This ensures that each incident light beam of the same bin will undergo the same modulation. Finally, the lens array in (E) widens the beam again for projecting out the angular variation.

This design guarantees that all output distributions of one meso-pixel will be perfectly aligned in the angular domain. Spatially they will shift slightly. However, for a large observation distance and a small macro-pixel size, this is hardly noticeable. Under these conditions, the contributions from neighboring meso-pixels, i. e. different incident angles, can hardly be resolved spatially. This effectively accumulates the different output distributions. Given this design of a single pixel which modulates the light in the angular

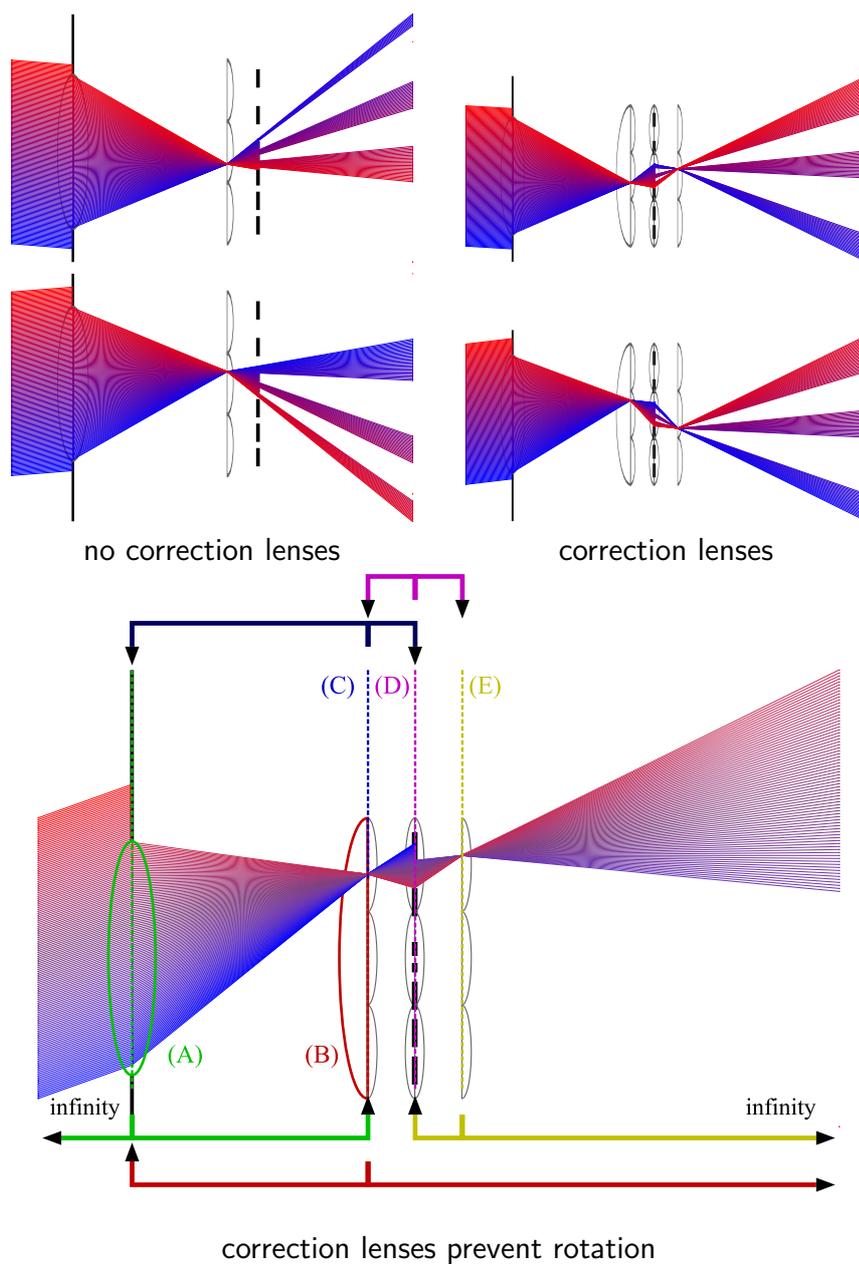


Figure 8.11: Design III ray diagrams. The simplest design (top left) rotates the outgoing light beams, which is fixed by adding additional lenses (bottom), as depicted on the top right.

domain, a 6D display can be designed by simply stacking multiple of these macro-pixels in two dimension, thus allowing for spatial and angular light modulation.

8.3.2 Results for Design III

Figure 8.12 shows a prototype for an observer-variant 4D display, i. e. a single macro-pixel of the 6D display, which we have again built using a 3D printed lens holder. In its design, we assume that the distance between planes (A) and (B) is large in comparison to the distance between (B) and (D), and choose the distances (B)-(D) and (D)-(E) to be the same. This allows for the use of the same lenslet arrays with focal length 10.5 mm as used in the prototype of Design 2. In plane (D), the arrays are again arranged in a back-to-back configuration to double its focal length. The reuse of the same lenslet arrays in all four places lowers manufacturing costs. Choosing the distance between (A) and (B) to be 50 mm, we can use off-the-shelf coated plano-convex lenses from Edmund Optics Inc. with 50 mm focal length, and 25 mm diameter in plane (A) and 30 mm diameter in plane (B), respectively.

In order to demonstrate the 4D design with angular variation, we put a pattern consisting of digits and letters in plane (D) (see Figure 8.13), and illuminated the prototype with a distant spot light. The output distribution is made visible by projection onto a diffuser screen which is not directly illuminated by the spot light. We translate the light source, that is, we change the angle of incidence, following the sequence indicated by the small arrows in Figure 8.13. The result is a time sequence of well readable letters projected onto the screen. In the corners ("3" and "7") some vignetting as well as some distortion is noticeable. This could be removed if the field lens would be enlarged to equally cover all meso-pixels in plane (B).

The results in Figure 8.13 share some similarity with the patterns produced by the Design 2 in Figure 8.10. Design 2 implements an angular-to-spatial reflectance field $R(\theta, \varphi, x, y)$ where every output pixel individually depends on the incident light direction. Design III, on the other hand implements an angular-to-angular reflectance field $R(\theta_i, \varphi_i, \theta_o, \varphi_o)$. Each letter is projected out by a different meso-pixel. Therefore, the letters are continuous and not composed of sub-pixels.

8.3.3 Design IV

Combining 7×7 macro-pixels of Design III we built a prototype for a 6D display implementing a rendering for a $2 \rightarrow 4$ reflectance field $R(\theta_i, \varphi_i; u_o, v_o, \theta_o, \varphi_o)$ (Figure 8.14, Design IV). Every macro-pixel independently produces some output distribution depending on the incident direction. Combined, the display project out different 4D light fields depending on the incident illumination.

8.3.4 Results for Design IV

We created a 6D pattern (see Figure 8.15) of simple illumination-dependent shapes which grow when the illumination is moved downwards, and appear increasingly hollow when

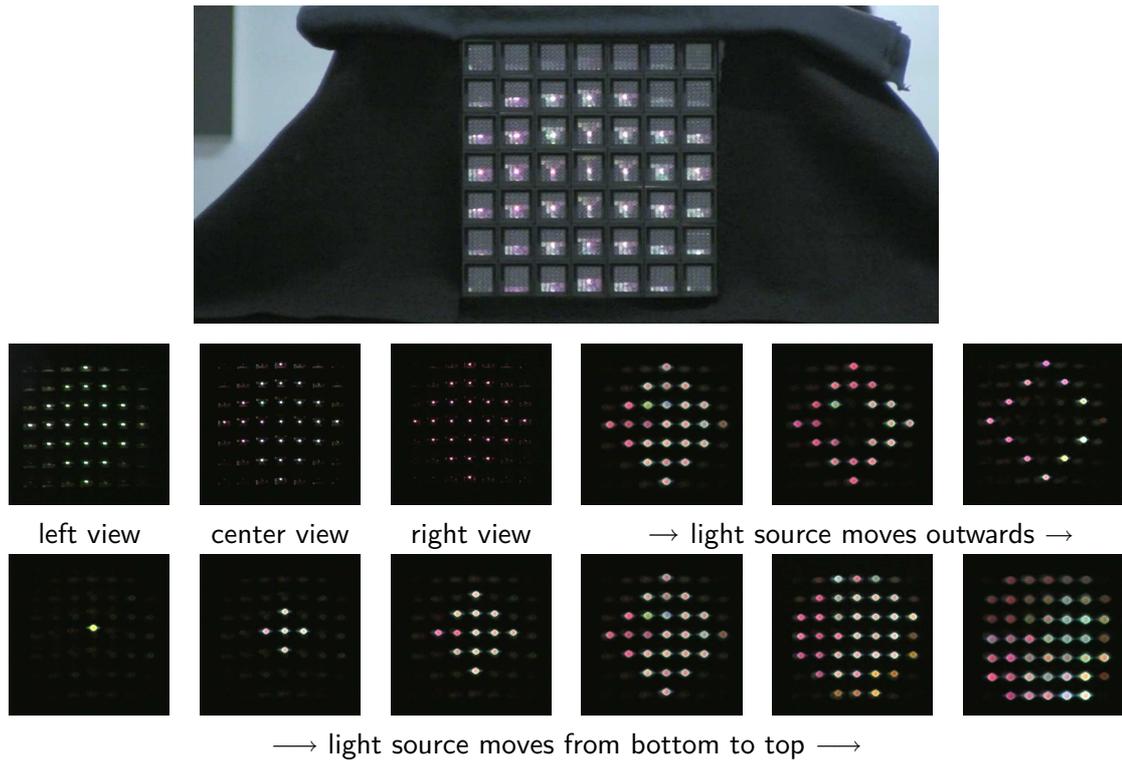


Figure 8.14: 6D experiment and result (Design IV). Top: fully assembled prototype in mostly diffuse room illumination. The image resembles the pattern from Figure 8.15. Center row: changing observer direction induces a green to magenta color shift; horizontal light movement empties the patterns inside. Bottom row: vertical light movement induces different patterns. For light movements, we employed focal blur to visualize the effect better.

illuminated from the sides. In each output pixel, and for each incident direction we further encode some color variation that depends on the viewing angle. For the sake of simplicity we chose the variation in color to be the same for every macro-pixel. The design does allow for a free configuration though.

In order to produce the results in Figure 8.14 we surrounded the device by dark cloth and then illuminated and observed the display from distance. In the top image of Figure 8.14 one actually sees the integration of the patterns produced by the diffuse illumination in the room. At this resolution the structure of the meso-pixels is still visible. In the two bottom lines, the device is again illuminated with a hand-held light bulb. The camera is slightly out of focus which makes the light transmitted by the individual meso-pixels more visible. As the camera is moved from left to right the color shift in each pixel can be observed. Moving the light source produces the intended shapes.

Since the observer is not sufficiently distant to the display, we observe not only the white, but also red and green colored light rays, in the pictures for the center view.

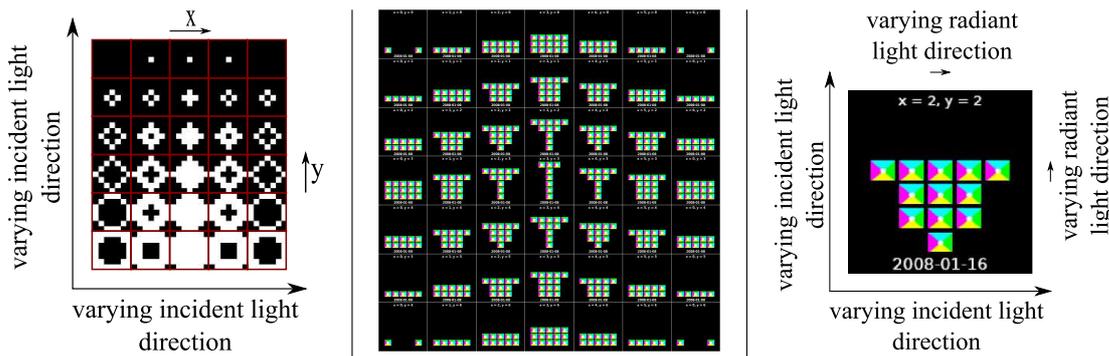


Figure 8.15: 6D pattern for Design IV. The display shows simple diamond shapes for varying incident directions (left), and adds an angular color modulation, so that they appear magenta when seen from the right, white when seen from the center, and green when seen from the left. This yields the 6D pattern in the center, with a single pixel zoom-in on the right.

8.4 Implementation details

In the previous sections, we have introduced the optical design for our prototypes. Now, we will discuss how to generate the patterns from a reflectance field structure for the 4D case. For our analysis, as well as for the ray diagrams that we have previously shown, we will apply the assumption of paraxial optics, i. e. we assume that the angles of all light rays in our optical system to the optical axes of the lenses within are sufficiently small. In paraxial optics the change of direction and position of each ray due to a lens can be easily modeled with matrix methods (Gerrard and Burch, 1975).

The encoding procedure for the pattern is as follows: for each pixel position (p_x, p_y) in the pattern

- determine the index (l_x, l_y) of the lenslet above the pixel
- determine the pixel position (s_x, s_y) relative to the optical axis of the lenslet
- simulate refraction at the lenslet in this pixel for a ray perpendicular to the pattern plane through (s_x, s_y) , obtaining a ray in direction (θ, φ)
- color the pixel at (p_x, p_y) proportional to $\frac{R(\theta, \varphi, l_x, l_y)}{\cos \kappa}$, where κ is the angle of the refracted ray to the glass surface normal. This factor compensates for less light falling into the lenslet under flat angles.

Finally, we tonemap the pattern by scaling its intensity linearly to a comfortable brightness and performing a gamma and blacklevel correction for the output device. We have tested printers based on laser, ink jet and solid ink technology, and achieved best results with an ink jet printer, putting two sheets back to back for increased contrast.

The pattern used in the demonstrator of Design I in Figure 8.1 has been computed from a captured reflectance field. It has been recorded with indirect, extended illumination from

the entire hemisphere around the object using a setup as described in chapter 4. The display however can only render a smaller range of incident angles. In order to produce a more interesting experience of moving highlights at the top, we have virtually rotated the incident light by 45 degrees during the calculation of the patterns. In principle, one can also re-parameterize the reflectance field along the angular dimension in order to emphasize or suppress the angular dependence.

8.4.1 Pattern/Lens Registration

The pattern and the lenses in our prototype have to be precisely aligned in order to modulate the correct light paths. For Design I, we have first measured the exact lens distance in the lenslet arrays to a higher precision than given by the manufacturer: we have printed patterns for a set of hypotheses and then chosen the distance which minimized the Moiré effects that occur when rotating the lens array on top of the pattern. For precise alignment of the final pattern, we have drilled holes into four lenses at the corners of the lenslet arrays under a microscope, and augmented the pattern by markers which show up below the holes. The patterns for the 6×5 lenslet arrays used in Design III and Design IV are manually cut and aligned.

8.5 Discussion

8.5.1 Limitations

Our display designs come with some inherent limitations:

For one, we only modulate the light that comes in from the back and project light out to the front side. Light integration is only performed and controlled from one hemisphere. Any light incident from the front hemisphere will limit the contrast to some extent. This is especially true for Designs I and II, which make use of a diffuser which inherently integrates the light from both sides.

Another limitation is given by field of view of the front lenslet array or main lens. The displays currently use much less than the potential 180 degrees of illumination and observation. The same holds for the projection lenses in plane (E) of Design III which produce a rather narrow projection cone.

As we use lenslet arrays without further angular blocking technology, wrapping can occur for rather shallow incident angles. In this case the light beam might cross over to the light path of a neighboring macro-pixel resulting in incorrect and distorted patterns.

Since our display is passive, we can not emit more light from a single macro-pixel than the flux that hits the primary lens. Thus, we cannot display reflectance fields of complicated optical components that focus light from afar on a single pixel with full intensity.

8.5.2 Future Challenges

There is a set of interesting extensions to the presented display prototypes. Most fundamental would likely be an extension to full 8D rendering, with arbitrary incident, non-distant 4D light fields. The design however is not obvious. As long as the light can be assumed distant and 2D parameterizable, we know that each macro lens is hit by the same incident light distribution, and therefore each pattern below the macro lens encodes arbitrarily distant global illumination effects (such as caustics or shadows). For a full 8D display that responds to a 4D incident light field, though, one would need to implement a coupling in the spatial domain. Some means of exchanging energy between different macro-pixels is required in order to faithfully reproduce global illumination effects such as caustics or subsurface scattering. This intrinsically requires a different design.

Another interesting direction of future research might implement the controlled integration of the front hemisphere. The incident light could be reflected back, by a mirror for example. Separate modulation of viewing and illumination rays remains a challenge.

In all the designs we presented, we manually cut out the patterns and aligned them with the optical components. This restricts the precision which we can build the prototypes with and keeps us from shrinking the single pixels, which would be required for higher resolutions. For industrial production, this should be less of an issue.

A principle limit on the resolution of Design III and IV is given by the diffraction limit. Too small features in the modulation plane will widen the transmitted rays resulting in a blurred angular distribution. The diffraction limit is however less a problem for Design I as the resulting image is observed right after the modulation plane.

The presented designs are all based around using lens systems in order to embed high dimensional information in planar surfaces. Holograms exploit variations in the phase of wave fronts; it should be interesting to research possible extensions which incorporate holographic techniques for higher data densities and resolutions. Another technology that might influence this area of multi-dimensional displays is nanotechnology which can already produce surfaces with particularly designed reflection functions (Min et al., 2007; Sun and Kim, 2004).

8.5.3 Applications

There is a large set of applications which the presented display designs might enable:

Using Design I we can produce a relightable flattened copy of any 3D object which will automatically adapt to the incident illumination, providing a relightable object virtualization. This could be useful to present novel products, valuable artefacts, or for instruction manuals in a rather novel way.

Manufacturing Design I at a larger scale, one might cover a full window which automatically block or dims the direct sun light according to the time of day, as expressed by the

sun position as it follows its arc on the sky. Since the modulation can be both in color and intensity one could also create mood lights.

Inserting a high resolution LCD panel instead of the printed transparencies one can implement even a dynamic reflectance field display. In particular, one could control the transport of individual light rays and dispatch them into different directions.

8.6 Conclusion

In this chapter, we have presented design and prototypes for a set of different reflectance field displays. They are built using passive devices only, combining lenses and transparencies so as to modulate the outgoing light field depending on the incident light direction. Once built, the displays do not require sophisticated electronics, expensive computations or even electric current.

The proposed designs make use of lenses to guide incident light beams to a particular location in the modulation plane. This design makes use of most of the light that hits the main lens or lens array. Compared to a possible design based on blockers it is inherently more light efficient.

While the optical quality of some of the prototypes is currently limited, it can be expected to improve with industrial manufacture. With the presented designs we hope to inspire further exploration and developments of higher dimensional display technology.

9 Concluding Remarks

9.1 Applications

Image-space relighting has a rich set of applications, as it enables us to abstract visual appearance of a scene from the optical environment we observe it in. This can be used for the purpose of digital archiving of valuable cultural artifacts – a future development of our passive display technology may make it possible one day to share the experience of unique cultural artifacts simultaneously by many people all over the world.

With globally networked computer systems, digital information can be transmitted quickly and has become truly ubiquitously available – this implies that we can now also communicate visual appearance in a new fashion. At the time of writing of this thesis, the websites of car manufacturers provide interactive, free-viewpoint visualizations of car interiors to attract potential customers as a standard feature. It is easily conceivable to augment existing product advertisement also by demonstrating the product appearance for different light environments; especially for items such as jewelry which show their true beauty only in the interaction with illumination.

However, the value of image-space relighting is not only in its ability to store visual appearance digitally. Defining the illumination of a scene *after* its appearance was initially recorded also provides new ways of artistic expression, and can complement the emerging technology to define camera parameters such as focus and depth-of-field (Ng, 2006) on the way towards achieving true post-hoc photography.

Relighting also makes it possible to bring elements of real scenes into virtual entertainment scenarios with unprecedented integration. Sony Corporation markets the EyeToy®, a simple digital camera, with great success, as it can be used to transport the likeness of a player into a digital game in real-time; augmenting this technology with relighting techniques could make the gaming experience much more persuasive.

Facing the task of implementing an image-space relighting pipeline for a particular application, one will have to carefully weigh the trade-offs of the design choices. For instance, for the acquisition of a human person, or even a live performance, high speed recording (Einarsson et al., 2006) is paramount to success; here the off-line processing of our interpolating image synthesis may increase the quality significantly (Chapter 6). For maximal accuracy, an adaptive sampling algorithm as in Section 5.3 can improve the quality for a constant number of input images.

If live streaming is required, both post-processing and adaptive sampling may be too expensive; in this case, aliasing artefacts can be easily avoided with the help of extended

light sources (Section 5.2) without any computational overhead. For the simulation of those techniques and the comparison to high-resolution ground truth data, our measurement setup (Chapter 4) provides a good starting point.

A light-weight, but portable recording setup may be constructed using the implicit relighting approach (Chapter 7), reducing expenses on hardware to a minimum. Passive reflectance field displays (Chapter 8) can contribute to virtualize the appearance of objects completely, creating a printable appearance representation that reacts to illumination in a natural way, without ad-hoc computations, even without the consumption of electrical energy in the rendering process.

9.2 Future Challenges

The discoveries presented in this thesis have pushed the envelope for the recording, modeling and rendering of 4D $2 \rightarrow 2$ reflectance fields, which enable relighting for a fixed choice of view point and camera parameters. While it is known today how to record in principle a general 8D $4 \rightarrow 4$ reflectance field which is free from these restrictions, the problem of finding a practical, inexpensive implementation is not yet solved. Also, for this most general case, efficiently compressing measurement schemes are needed – if the resolution of a simple 1000×1000 pixel 2D image is sought for across all eight free parameters of an 8D reflectance field, information in the order of 10^{24} point-to-point light transport observations must be stored, which in the raw form will remain impractical for years to come. For the development of these techniques, the study of image-space relighting may prove most useful.

With the data once available, new challenges for their editing arise. For instance, there are techniques available today that enable to morph the shapes of 2D images into each other, employing blending and warping. For higher dimensional data, though, these are not directly applicable. It might be possible to estimate geometrical descriptions for reflectance field data, and then incorporate approaches that are already available for meshed animations with surface reflectance fields. However, this requires intensive investigations, for instance, in order to detect and align features.

Another important topic is the efficient integration of global scene measurement into surface-level descriptions. While Lehtinen (2007) has already proposed a unifying framework for measured and synthetic scene descriptors, a method for the inference of locally editable light transport properties that maintains the fidelity of a global descriptor still remains to be discovered. Such research might trigger the development of a new approach towards rendering, where reflectance fields form the primitives of scene description.

Finally, most measurement setups available today work well in a laboratory environment, and require expensive devices and careful, manual calibration. In order to make them applicable by the general public and spread their use, new, simple, consumer devices will be required to record rich scene representations with the same casual attitude that we have today towards photography, thanks to the advent of digital cameras.

9.3 Summary

Image-space relighting, as a tool for globally modeled scene representation, shares the same traits as other global approaches. It is general, coping with a wide range of materials; with sufficient sampling density, it can model highly specular objects made from polished steel, refractive objects such as glass, correctly handle subsurface scattering in objects such as human skin and fruit, handle cast shadows from directional and area light sources, and any intra-surface light transport within the recorded scene.

It lacks somewhat with respect to expressiveness – it is an image-space method, and therefore does not provide the ability to generate novel views. However, in the way that we have presented it in this thesis, it is very robust. While in a technique that involves reconstructing a BRDF, SVBRDF or even BSSRDF from pictures, even a slight mis-calibration of geometric position or light source parameters has a strong non-linear effect and can break the method, this is not the case for image-space relighting: small errors in the acquisition process only cause small errors in the rendering. With the techniques that we have employed in this thesis, the illumination expressed in a rendered image corresponds to an actually realizable illumination in the real-world. Therefore, a mis-calibrated light source direction merely causes a warp between the desired incident illumination and the illumination in the image; likewise, mis-estimated light source brightness only locally influences the brightness of the illumination in the output picture.

With the work presented in this thesis, we have advanced the state of the art in the following respects:

- we have shown how to construct a **practical measurement setup** that performs a pre-filtering on reflectance data with directional illumination of controllable extent, thus **countering aliasing artefacts** created by low numbers of input pictures,
- we have proposed an **adaptive directional-light sampling algorithm** for reflectance fields that places illumination samples in the most important directions, reducing the measurement time,
- with our interpolating **image synthesis** algorithm, the rendering quality of reflectance fields is greatly improved in post-processing,
- we have proposed a new, **implicit representation** of reflectance data that further reduces the required effort in recording measurements,
- our **passive devices for rendering** reflectance fields have proven a novel paradigm for the cost-effective rendering of relightable images.

While restricted to the 4D case, our discoveries should carry forward in many respects to higher dimensional methods. Thus, we hope that our contribution motivates further research on the computer graphics pipeline – in its continuing improvement end-to-end.

Bibliography

- Adams, A. and Levoy, M. (2007). General Linear Cameras with Finite Aperture. In *Rendering Techniques 2007 (Proc. Eurographics Symposium on Rendering)*, pages 121–126. 11
- Adelson, E. H. and Bergen, J. R. (1991). The plenoptic function and the elements of early vision. In Landy, M. S. and Movshon, J. A., editors, *Computational models of visual processing*, chapter 1. MIT Press. 9
- Agarwal, S., Ramamoorthi, R., Belongie, S., and Jensen, H. W. (2003). Structured importance sampling of environment maps. *ACM Transactions on Graphics (Proc. SIGGRAPH 2003)*, 22(3):605–612. 19
- Ajdin, B., Hullin, M. B., Fuchs, C., Seidel, H.-P., and Lensch, H. P. A. (2008). Demosaicing by smoothing along 1d features. *Computer Vision and Pattern Recognition, 2008. IEEE Computer Society Conference on.*, 1:1–8 Vol. 1. 10
- Akers, D., Losasso, F., Klingner, J., Agrawala, M., Rick, J., and Hanrahan, P. (2003). Conveying shape and features with image-based relighting. In *VIS '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*, page 46, Washington, DC, USA. IEEE Computer Society. 32
- Annen, T., Dong, Z., Mertens, T., Bekaert, P., Seidel, H.-P., and Kautz, J. (2008). Real-time, all-frequency shadows in dynamic scenes. *ACM Transactions on Graphics (Proc. SIGGRAPH 2008)*, 27(3). 19
- Anrys, F., Dutré, P., and Willems, Y. D. (2004a). Image based lighting design. In *The 4th IASTED Int. Conf. on Visualization, Imaging, and Image Proc.* 31, 96
- Anrys, F., Dutré, P., and Willems, Y. D. (2004b). Lighting design by simulated annealing. Technical report CW 393, Departement Computerwetenschappen, KU.Leuven. 31, 96
- Arai, J., Okano, F., Hoshino, H., and Yuyama, I. (1998). Gradient-index lens-array method based on real-time integral photography for three-dimensional images. *Applied Optics*, 37:2034–2045. 18
- Ashikhmin, M. and Shirley, P. (2002). Steerable illumination textures. *ACM Transactions on Graphics*, 21(1):1–19. 23

- Barrow, H. and Tenenbaum, J. (1978). Recovering intrinsic scene characteristics from images. In Hanson, A. and Riseman, E., editors, *Computer Vision Systems*. Academic Press. 69
- Belhumeur, P. and Kriegman, D. (1996). What is the set of images of an object under all possible lighting conditions. 28
- Ben-Artzi, A., Egan, K., Ramamoorthi, R., and Durand, F. (2008). A precomputed polynomial representation for interactive BRDF editing with global illumination. *ACM Transactions on Graphics*, 27(2):1–13. 24
- Ben-Artzi, A., Overbeck, R., and Ramamoorthi, R. (2006). Real-time BRDF editing in complex lighting. *ACM Transactions on Graphics (Proc. SIGGRAPH 2006)*, 25(3):945–954. 24
- Bentley, J. (1975). Multidimensional binary search trees used for associative searching. *CACM*, 18(9):509–517. 21
- Bentley, J. L. and Friedman, J. H. (1979). Data structures for range searching. *ACM Computing Surveys*, 11(4):397–409. 21
- Black, M. J. and Anandan, P. (1996). The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer Vision and Image Understanding, CVIU*, 63(1):75–104. 69
- Blanz, V., Scherbaum, K., Vetter, T., and Seidel, H.-P. (2004). Exchanging faces in images. *Computer Graphics Forum (Proc. Eurographics 2004)*, 23(3):669–676. 26
- Blanz, V. and Vetter, T. (1999). A morphable model for the synthesis of 3D faces. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194. ACM Press/Addison-Wesley Publishing Co. 26
- Blinn, J. F. (1977). Models of light reflection for computer synthesized pictures. In *SIGGRAPH '77: Proceedings of the 4th annual conference on Computer graphics and interactive techniques*, pages 192–198, New York, NY, USA. ACM. 23
- Blinn, J. F. and Newell, M. E. (1976). Texture and reflection in computer generated images. *Commun. ACM*, 19(10):542–547. 19
- Boivin, S. and Gagalowicz, A. (2001). Image-based rendering of diffuse, specular and glossy surfaces from a single image. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 107–116, New York, NY, USA. ACM. 25
- Brox, T., Bruhn, A., Papenber, N., and Weickert, J. (2004). High accuracy optical flow estimation based on a theory for warping. In Pajdla, T. and Matas, J., editors, *European Conference on Computer Vision (ECCV)*, volume 3024 of *LNCS*, pages 25–36, Prague, Czech Republic. Springer. 69, 74

- Buehler, C., Bosse, M., McMillan, L., Gortler, S., and Cohen, M. (2001). Unstructured lumigraph rendering. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 425–432, New York, NY, USA. ACM. 18
- Cabral, B., Max, N., and Springmeyer, R. (1987). Bidirectional reflection functions from surface bump maps. In *SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 273–281, New York, NY, USA. ACM. 24
- Cabral, B., Olano, M., and Nemec, P. (1999). Reflection space image based rendering. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 165–170, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 23
- Catmull, E. E. (1974). *A Subdivision Algorithm for Computer Display of Curved Surfaces*. PhD thesis, Dept. of CS, U. of Utah. 21
- Catmull, E. E. (1975). Computer display of curved surfaces. In *Proceedings of the IEEE Conference on Computer Graphics, Pattern Recognition, and Data Structure*, pages 11–17. 21
- Chai, J.-X., Chan, S.-C., Shum, H.-Y., and Tong, X. (2000). Plenoptic sampling. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 307–318, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 18, 45
- Chang, C.-F., Bishop, G., and Lastra, A. (1999). LDI tree: a hierarchical representation for image-based rendering. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 291–298, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 20
- Chen, B. and Lensch, H. P. A. (2005). Light Source Interpolation for Sparsely Sampled Reflectance Fields. In *Proc. Vision, Modeling and Visualization*, pages 461–469. 28, 69
- Chen, W.-C., Bouguet, J.-Y., Chu, M. H., and Grzeszczuk, R. (2002). Light field mapping: efficient representation and hardware rendering of surface light fields. *ACM Transactions on Graphics (Proc. SIGGRAPH 2002)*, 21(3):447–456. 19
- Christensen, P. H., Stollnitz, E. J., Salesin, D. H., and DeRose, T. D. (1996). Global illumination of glossy environments using wavelets and importance. *ACM Transactions on Graphics*, 15(1):37–71. 22
- Chuang, Y.-Y., Zongker, D. E., Hindorff, J., Curless, B., Salesin, D. H., and Szeliski, R. (2000). Environment matting extensions: towards higher accuracy and real-time

- capture. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 121–130, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 28, 30
- Colbert, M., Pattanaik, S., and Krivánek, J. (2006). BRDF-shop: Creating physically correct bidirectional reflectance distribution functions. *IEEE Computer Graphics and Applications*, 26(01):30–36. 24
- Cook, R. L. (1984). Shade trees. *SIGGRAPH Comput. Graph.*, 18(3):223–231. 26
- Cook, R. L., Porter, T., and Carpenter, L. (1984). Distributed ray tracing. *ACM SIGGRAPH Computer Graphics (Proc. SIGGRAPH 1984)*, 18(3):137–145. 22
- Cook, R. L. and Torrance, K. E. (1981). A reflectance model for computer graphics. In *SIGGRAPH '81: Proceedings of the 8th annual conference on Computer graphics and interactive techniques*, pages 307–316, New York, NY, USA. ACM. 23, 24
- Coombe, G., Hantak, C., Lastra, A., and Grzeszczuk, R. (2005). Online construction of surface light fields. *Rendering Techniques 2005 (Proc. Eurographics Symposium on Rendering)*, pages 83–90. 19
- Crow, F. C. (1977). Shadow algorithms for computer graphics. *Computer Graphics (SIGGRAPH '77 Proceedings)*, 11(2):242–248. 21
- Dana, K., van Ginneken, B., Nayar, S., and Koenderink, J. (1996). Reflectance and texture of real world surfaces. Technical report, Department of Computer Science, Columbia University. 33
- Dana, K. J., van Ginneken, B., Nayar, S. K., and Koenderink, J. J. (1999). Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34. 33
- Daubert, K., Lensch, H. P. A., Heidrich, W., and Seidel, H.-P. (2001). Efficient cloth modeling and rendering. *Rendering Techniques 2001 (Proc. Eurographics Rendering Workshop)*, pages 63–70. 33
- Daubert, K. and Seidel, H.-P. (2002). Hardware-based volumetric knit-wear. *Computer Graphics Forum (Proc. Eurographics 2002)*, 21(3). 33
- Debevec, P. (1998). Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 189–198, New York, NY, USA. ACM. 19, 80, 97
- Debevec, P., Hawkins, T., Tchou, C., Duiker, H.-P., Sarokin, W., and Sagar, M. (2000). Acquiring the reflectance field of a human face. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 145–156, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 4, 14, 27, 29, 68, 94

- Debevec, P., Tchou, C., Gardner, A., Hawkins, T., Poullis, C., Stumpfel, J., Jones, A., Yun, N., Einarsson, P., Lundgren, T., Fajardo, M., and Martinez, P. (2004). Estimating surface reflectance properties of a complex scene under captured natural illumination. *ICT Technical Report ICT-TR-06.2004*. 25
- Debevec, P., Wenger, A., Tchou, C., Gardner, A., Waese, J., and Hawkins, T. (2002). A lighting reproduction approach to live-action compositing. *ACM Transactions on Graphics (Proc. SIGGRAPH 2002)*, 21(3):547–556. 31
- Debevec, P. E. and Malik, J. (1997). Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 369–378, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 19
- Debevec, P. E., Taylor, C. J., and Malik, J. (1996). Modeling and rendering architecture from photographs: a hybrid geometry- and image-based approach. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 11–20, New York, NY, USA. ACM. 20
- Debevec, P. E., Yu, Y., and Boshukov, G. D. (1998). Efficient view-dependent image-based rendering with projective texture-mapping. *Rendering Techniques '98 (Proc. Eurographics Rendering Workshop)*, pages 105–116. 20
- Dischler, J.-M. (1998). Efficiently rendering macro geometric surface structures with bi-directional texture functions. *Rendering Techniques '98 (Proc. Eurographics Rendering Workshop)*, pages 169–180. 33
- Dobashi, Y., Kaneda, K., Nakashima, T., Yamashita, H., and Nishita, T. (1995). A quick rendering method using basis functions for interactive lighting design. *Computer Graphics Forum (Proc. Eurographics 95)*, 14(3):229–240. 31
- Dobashi, Y., Kaneda, K., Yamashita, H., and Nishita, T. (1996). Method for calculation of sky light luminance aiming at interactive architectural design. *Computer Graphics Forum (Proc. Eurographics 1996)*, 15(3):109–118. 31
- Donner, C. and Jensen, H. W. (2006). A spectral BSSRDF for shading human skin. *Rendering Techniques 2006 (Proc. Eurographics Symposium on Rendering)*, pages 409–417. 26
- Dorsey, J., Arvo, J., and Greenberg, D. (1995). Interactive design of complex time-dependent lighting. *IEEE Computer Graphics and Applications*, 15(2):26–36. 31
- Duda, R., Hart, P., and Stork, D. (2001). *Pattern Classification*. John Wiley & Sons, New York, 2nd edition. 88
- Dumont, O., Masselus, V., Zaenen, P., Wagemans, J., and Dutré, P. (2005). A perceptual quality scale for image-based relighting of glossy surfaces. *Katholieke Universiteit Leuven Tech Report*. 29

- Edwards, D., Boulos, S., Johnson, J., Shirley, P., Ashikhmin, M., Stark, M., and Wyman, C. (2006). The halfway vector disk for BRDF modeling. *ACM Transactions on Graphics*, 25(1):1–18. 23
- Einarsson, P., Chabert, C.-F., Jones, A., Ma, W.-C., Lamond, B., Hawkins, T., Bolas, M., Sylwan, S., and Debevec, P. (2006). Relighting human locomotion with flowed reflectance fields. In *Rendering Techniques 2006 (Proc. Eurographics Symposium on Rendering)*, pages 183–194. 27, 119
- Fattal, R., Agrawala, M., and Rusinkiewicz, S. (2007). Multiscale shape and detail enhancement from multi-light image collections. *ACM Transactions on Graphics (Proc. SIGGRAPH 2007)*, 26(3):51. 31
- Feng, W.-W., Peng, L., Jia, Y., and Yu, Y. (2007). Large-scale data management for PRT-based real-time rendering of dynamically skinned models. *Rendering Techniques 2007 (Proc. Eurographics Symposium on Rendering)*, pages 23–34. 23
- Fitzgibbon, A. W., Wexler, Y., and Zisserman, A. (2003). Image-based rendering using image-based priors. In *Proc. International Conference on Computer Vision (ICCV)*, pages 1176–1183. 69
- Fournier, A. (1995). From local to global illumination and back. *Rendering Techniques '95 (Proc. Eurographics Rendering Workshop)*, pages 127–136. 27
- Fuchs, C., Chen, T., Goesele, M., Theisel, H., and Seidel, H.-P. (2006). Volumetric density capture from a single image. In Möller, T., Machiraju, R., Chen, M.-S., and Ertl, T., editors, *Volume Graphics 2006 : Eurographics / IEEE VGTC Workshop Proceedings ; Fifth International Workshop on Volume Graphic*, pages 17–22, Boston, USA. Eurographics / IEEE Computer Society VGTC, Eurographics. 26
- Fuchs, C., Chen, T., and Seidel, H.-P. (2007a). Density estimation for dynamic volumes. *Computers and Graphics*, 31(2):205–211. 26
- Fuchs, M., Blanz, V., Lensch, H., and Seidel, H.-P. (2005a). Reflectance from images: A model-based approach for human faces. *IEEE Transactions on Visualization and Computer Graphics*, 11(03):296–305. 26
- Fuchs, M., Blanz, V., Lensch, H. P., and Seidel, H.-P. (2007b). Adaptive sampling of reflectance fields. *ACM Transactions on Graphics*, 26(2):10. 4, 29
- Fuchs, M., Blanz, V., and Seidel, H.-P. (2005b). Bayesian relighting. In *Rendering Techniques 2005 (Proc. Eurographics Symposium on Rendering)*, pages 157–164. 5, 30
- Fuchs, M., Lensch, H. P. A., Blanz, V., and Seidel, H.-P. (2007c). Superresolution reflectance fields: Synthesizing images for intermediate light directions. *Computer Graphics Forum (Proc. Eurographics 2007)*, 26(3):447–456. 5, 30

- Fuchs, M., Raskar, R., Seidel, H.-P., and Lensch, H. P. A. (2008). Towards passive 6d reflectance field displays. *ACM Transactions on Graphics (Proc. SIGGRAPH 2008)*, 27(3). 6, 31
- Furukawa, R., Kawasaki, H., Ikeuchi, K., and Sakauchi, M. (2002). Appearance based object modeling using texture database: Acquisition, compression and rendering. *Rendering Techniques 2002 (Proc. Eurographics Rendering Workshop)*, pages 257–266. 33
- Gardner, A., Tchou, C., Hawkins, T., and Debevec, P. (2003). Linear light source reflectometry. *ACM Transactions on Graphics (Proc. SIGGRAPH 2003)*, 22(3):749–758. 25
- Gardner, A., Tchou, C., Wenger, A., Hawkins, T., and Debevec, P. (2004). Postproduction re-illumination of live action using interleaved lighting. In *SIGGRAPH 2004 (Poster)*. 29
- Garg, G. (2006). *Efficiently Acquiring Reflectance Fields Using Patterned Illumination*. PhD thesis, Stanford University. 27
- Garg, G., Talvala, E.-V., Levoy, M., and Lensch, H. P. A. (2006). Symmetric photography: Exploiting data-sparseness in reflectance fields. In *Rendering Techniques 2006 (Proc. Eurographics Symposium on Rendering)*, pages 251–262. 2, 27
- Georgiev, T., Zheng, K. C., Curless, B., Salesin, D., Nayar, S., and Intwala, C. (2006). Spatio-angular resolution trade-offs in integral photography. In *Rendering Techniques 2006 (Proc. Eurographics Symposium on Rendering)*, pages 263–272. 18
- Gerrard, A. and Burch, J. M. (1975). *Introduction to Matrix Methods in Optics*. John Wiley & Sons, New York - Chichester - Brisbane. 115
- Ghosh, A., Trentacoste, M., Seetzen, H., and Heidrich, W. (2005). Real Illumination from Virtual Environments. In *Rendering Techniques 2005 (Proc. Eurographics Symposium on Rendering)*, pages 243–252. 19
- Goesele, M. (2004). *New Acquisition Techniques for Real Objects and Light Sources*. PhD thesis, Universität des Saarlandes. 18
- Goesele, M., Granier, X., Heidrich, W., and Seidel, H.-P. (2003). Accurate light source acquisition and rendering. *ACM Transactions on Graphics (Proc. SIGGRAPH 2003)*, 22(3):621–630. 18
- Goesele, M., Lensch, H. P. A., J.Lang, Fuchs, C., and Seidel, H.-P. (2004). DISCO: Acquisition of translucent objects. *ACM Transactions on Graphics (Proc. SIGGRAPH 2004)*, 23(3):835–844. 26

- Goldman, D. B. and Chen, J.-H. (2005). Vignette and exposure calibration and compensation. In *The 10th IEEE International Conference on Computer Vision*, pages 899–906. 39
- Gondek, J. S., Meyer, G. W., and Newman, J. G. (1994). Wavelength dependent reflectance functions. In *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 213–220, New York, NY, USA. ACM. 24
- Goral, C. M., Torrance, K. E., Greenberg, D. P., and Battaile, B. (1984). Modeling the interaction of light between diffuse surfaces. *SIGGRAPH Computer Graphics (Proc. SIGGRAPH 1984)*, 18(3):213–222. 22
- Gortler, S. J., Grzeszczuk, R., Szeliski, R., and Cohen, M. F. (1996). The lumigraph. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 43–54, New York, NY, USA. ACM. 18
- Greene, N. (1986). Environment mapping and other applications of world projections. *IEEE Computer Graphics and Applications*, 6(11):21–29. 19
- Gu, J., Tu, C.-I., Ramamoorthi, R., Belhumeur, P., Matusik, W., and Nayar, S. (2006). Time-varying surface appearance: acquisition, modeling and rendering. *ACM Transactions on Graphics (Proc. SIGGRAPH 2006)*, 25(3):762–771. 26
- Günther, J., Chen, T., Goesele, M., Wald, I., and Seidel, H.-P. (2005). Efficient acquisition and realistic rendering of car paint. In Greiner, G., Hornegger, J., Niemann, H., and Stamminger, M., editors, *Vision, Modelling, and Visualization 2005 (VMV), Proceedings, November 16-18, Erlangen, Germany*, pages 487–494, Berlin. Akademische Verlagsgesellschaft Aka GmbH. 24
- Han, J. Y. and Perlin, K. (2003). Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Transactions on Graphics (Proc. SIGGRAPH 2003)*, 22(3):741–748. 33
- Hanrahan, P. and Krueger, W. (1993). Reflection from layered surfaces due to subsurface scattering. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 165–174, New York, NY, USA. ACM. 26
- Hawkins, T., Einarsson, P., and Debevec, P. (2005a). Acquisition of time-varying participating media. *ACM Transactions on Graphics (Proc. SIGGRAPH 2005)*, 24(3):812–815. 26
- Hawkins, T., Einarsson, P., and Debevec, P. (2005b). A dual light stage. In *Rendering Techniques 2005 (Proc. Eurographics Symposium on Rendering)*, pages 91–98. 30
- Hawkins, T., Wenger, A., Tchou, C., Gardner, A., GÅransson, F., and Debevec, P. (2004). Animatable facial reflectance fields. In *Rendering Techniques 2004 (Proc. Eurographics Symposium on Rendering)*, pages 309–319. Eurographics Assoc. 30

- He, X. D., Heynen, P. O., Phillips, R. L., Torrance, K. E., Salesin, D. H., and Greenberg, D. P. (1992). A fast and accurate light reflection model. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 253–254, New York, NY, USA. ACM. 23
- He, X. D., Torrance, K. E., Sillion, F. X., and Greenberg, D. P. (1991). A comprehensive physical model for light reflection. In *SIGGRAPH '91: Proceedings of the 18th annual conference on Computer graphics and interactive techniques*, pages 175–186, New York, NY, USA. ACM. 23
- Hecht, E. (2002). *Optics*. Addison Wesley. 8
- Heidrich, W. (1999). *High-quality Shading and Lighting for Hardware-accelerated Rendering*. PhD thesis, Univeristy of Erlangen, Erlangen, Germany. 21
- Heidrich, W., Kautz, J., Slusallek, P., and Seidel, H.-P. (1998). Canned lightsources. *Rendering Techniques '98 (Proc. Eurographics Rendering Workshop)*, pages 293–300. 18
- Heidrich, W., Lensch, H., Cohen, M. F., and Seidel, H.-P. (1999). Light field techniques for reflections and refractions. *Rendering Techniques '99 (Proc. Eurographics Rendering Workshop)*, pages 187–196. 18
- Isaksen, A., McMillan, L., and Gortler, S. J. (2000). Dynamically reparameterized light fields. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 297–306, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 18
- Iwasaki, K., Dobashi, Y., Yoshimoto, F., and Nishita, T. (2007). Precomputed radiance transfer for dynamic scenes taking into account light interreflection. *Rendering Techniques 2007 (Proc. Eurographics Symposium on Rendering)*, pages 35–44. 23
- Javidi, B. and Okano, F. (Ed.) (2001). *Three-dimensional video and display: devices and systems*. SPIE, Washington, USA. 102
- Jensen, H. W. and Christensen, N. J. (1995). Photon maps in bidirectional monte carlo ray tracing of complex objects. *Computers & Graphics*, 19(2):215–224. 22
- Jensen, H. W., Marschner, S. R., Levoy, M., and Hanrahan, P. (2001). A practical model for subsurface light transport. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 511–518, New York, NY, USA. ACM. 26
- Jones, A., Gardner, A., Bolas, M., McDowall, I., and Debevec, P. (2006). Simulating spatially varying lighting on a live performance. In *3rd European Conference on Visual Media Production (CVMP 2006)*. 30

- Kajiya, J. T. (1985). Anisotropic reflection models. In *SIGGRAPH '85: Proceedings of the 12th annual conference on Computer graphics and interactive techniques*, pages 15–21, New York, NY, USA. ACM. 24
- Kajiya, J. T. (1986). The rendering equation. *SIGGRAPH Comput. Graph.*, 20(4):143–150. 13, 22
- Kannala, J. and Brandt, S. (2004). A generic camera calibration method for fish-eye lenses. In *Proceedings of the 2004 Virtual Reality*. IEEE. 39
- Karner, K. F., Mayer, H., and Gervautz, M. (1996). An image based measurement system for anisotropic reflection. *Computer Graphics Forum (Proc. Eurographics 1996)*, 15(3):119–128. 24
- Kautz, J., Boulos, S., and Durand, F. (2007). Interactive editing and modeling of bidirectional texture functions. *ACM Transactions on Graphics (Proc. SIGGRAPH 2007)*, 26(3):53. 33
- Kautz, J. and McCool, M. D. (1999). Interactive rendering with arbitrary BRDFs using separable approximations. *Rendering Techniques '99 (Proc. Eurographics Rendering Workshop)*, pages 247–260. 24
- Khan, E. A., Reinhard, E., Fleming, R. W., and Bühlhoff, H. H. (2006). Image-based material editing. *ACM Transactions on Graphics (Proc. SIGGRAPH 2006)*, 25(3):654–663. 20
- Koike, T. and Naemura, T. (2007). BRDF display. *SIGGRAPH 2007 poster presentation*. 102
- Kontkanen, J., Turquin, E., Holzschuch, N., and Sillion, F. X. (2006). Wavelet radiance transport for interactive indirect lighting. *Rendering Techniques 2006 (Proc. Eurographics Symposium on Rendering)*, pages 161–171. 22
- Koudelka, M., Belhumeur, P. N., Magda, S., and Kriegman, D. J. (2001). Image-based modeling and rendering of surfaces with arbitrary BRDFs. In *Proceedings of IEEE CVPR 2001*, pages 568–575. 27
- Kubota, A., Takahashi, K., Aizawa, K., and Chen, T. (2004). All-focused light field rendering. *Rendering Techniques 2004 (Proc. Eurographics Symposium on Rendering)*, pages 235–242. 18
- Lafortune, E. P. and Willems, Y. D. (1993). Bi-directional path tracing. In *Proceedings of Computergraphics '93*, pages 145–153, Alvor, Portugal. 22
- Lafortune, E. P. and Willems, Y. D. (1994). A theoretical framework for physically based rendering. *Computer Graphics Forum*, 13(2):97–107. 27

- Lafortune, E. P. F., Foo, S.-C., Torrance, K. E., and Greenberg, D. P. (1997). Non-linear approximation of reflectance functions. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 117–126, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 23, 25
- Lalonde, P. and Fournier, A. (1997). A wavelet representation of reflectance functions. *IEEE Transactions on Visualization and Computer Graphics*, 3(04):329–336. 23
- Lambert, J. H. (1760). *Photometria Sive de Mensura et Gradibus Luminus, Colorum et Umbrae*. Eberhard Klett. 13, 23
- Latta, L. and Kolb, A. (2002). Homomorphic factorization of BRDF-based lighting computation. *ACM Transactions on Graphics (Proc. SIGGRAPH 2002)*, 21(3):509–516. 24
- Lawrence, J., Ben-Artzi, A., DeCoro, C., Matusik, W., Pfister, H., Ramamoorthi, R., and Rusinkiewicz, S. (2006). Inverse shade trees for non-parametric material representation and editing. *ACM Transactions on Graphics (Proc. SIGGRAPH 2006)*, 25(3):735–745. 26
- Lee, J., Pfister, H., Moghaddam, B., and Machiraju, R. (2005). Estimation of 3D faces and illumination from single photographs using a bilinear illumination model. *Rendering Techniques 2005 (Proc. Eurographics Symposium on Rendering)*, pages 73–82. 26
- Lehtinen, J. (2007). A framework for precomputed and captured light transport. *ACM Transactions on Graphics*, 26(4):13. 17, 120
- Lensch, H. P. A. (2003). *Efficient, Image-Based Appearance Acquisition of Real-World Objects*. PhD thesis, Universität des Saarlandes. 12, 25
- Lensch, H. P. A., Kautz, J., Goesele, M., Heidrich, W., and Seidel, H.-P. (2001). Image-based reconstruction of spatially varying materials. *Rendering Techniques 2001 (Proc. Eurographics Rendering Workshop)*, pages 103–114. 25
- Lensch, H. P. A., Kautz, J., Goesele, M., Heidrich, W., and Seidel, H.-P. (2003a). Image-based reconstruction of spatial appearance and geometric detail. *ACM Transactions on Graphics*, 22(2):234–257. 25
- Lensch, H. P. A., Lang, J., Sá, A. M., and Seidel, H.-P. (2003b). Planned sampling of spatially varying BRDFs. *Computer Graphics Forum (Proc. Eurographics)*, 22(3):473–482. 25, 46
- Leung, M.-K., Pang, W.-M., Fu, C.-W., Wong, T.-T., and Heng, P.-A. (2007). Tileable BTF. *IEEE Transactions on Visualization and Computer Graphics*, 13(5):953–965. 33
- Levoy, M. and Hanrahan, P. (1996). Light field rendering. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 31–42, New York, NY, USA. ACM Press. 10, 17, 18

- Levoy, M., Ng, R., Adams, A., Footer, M., and Horowitz, M. (2006). Light field microscopy. *ACM Transactions on Graphics (Proc. SIGGRAPH 2006)*, 25(3):924–934. 18
- Lewis, R. R. and Fournier, A. (2000). Wavelet radiative transfer and surface interaction. *Computer Graphics Forum*, 19(2). 22
- Lippmann, G. (1908). Epreuves reversibles donnant la sensation du relief. *Journal of Physics*, 7:821–825. 101
- Lischinski, D. and Rappoport, A. (1998). Image-based rendering for non-diffuse synthetic scenes. *Rendering Techniques '98 (Proc. Eurographics Rendering Workshop)*, pages 301–314. 20
- Liu, X., Hu, Y., Zhang, J., Guo, X. T. B., and Shum, H.-Y. (2004a). Synthesis and rendering of bidirectional texture functions on arbitrary surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 10(03):278–289. 33
- Liu, X., Sloan, P.-P., Shum, H.-Y., and Snyder, J. (2004b). All-frequency precomputed radiance transfer for glossy objects. *Rendering Techniques 2004 (Proc. Eurographics Symposium on Rendering)*, pages 337–344. 23
- Liu, Y., Chen, G., Max, N., Hofsetz, C., and McGuinness, P. (2006). Undersampled light field rendering by a plane sweep. *Computer Graphics Forum*, 25(2):225–236. 18
- Lloyd, S. P. (1982). Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28:128–137. 49
- Loscos, C., Drettakis, G., and Robert, L. (2000). Interactive virtual relighting of real scenes. *IEEE Transactions on Visualization and Computer Graphics*, 6(04):289–305. 25
- Loscos, C., Frasson, M.-C., Drettakis, G., Walter, B., Granier, X., and Poulin, P. (1999). Interactive virtual relighting and remodeling of real scenes. *Rendering Techniques '99 (Proc. Eurographics Rendering Workshop)*, pages 329–340. 25
- Lu, R., Koenderink, J. J., and Kappers, A. M. L. (1998). Optical properties (bidirectional reflection distribution functions) of velvet. *Applied Optics*, 37(25):5974–5984. 24
- Malzbender, T., Wilburn, B., Gelb, D., and Ambrisco, B. (2006). Surface enhancement using real-time photometric stereo and reflectance transformation. *Rendering Techniques 2006 (Proc. Eurographics Symposium on Rendering)*, pages 245–250. 32
- Mann, S. (1995). Recording lightspace so shadows and highlights vary with varying viewing illumination. *Optics Letters*, 20(24). 30, 102
- Marschner, S. R. and Greenberg, D. P. (1997). Inverse lighting for photography. In *Proceedings of IS&T/SID Fifth Color Imaging Conference*, pages 262–265, Scottsdale, AZ, USA. 32

- Marschner, S. R., Westin, S. H., Lafortune, E. P. F., Torrance, K. E., and Greenberg, D. P. (1999). Image-based BRDF measurement including human skin. *Rendering Techniques '99 (Proc. Eurographics Rendering Workshop)*, pages 131–144. 24
- Masselus, V., Dutré, P., and Anrys, F. (2002). The Free-form Light Stage. In Debevec, P. Gibson, S., editor, *Proc. Eurographics Rendering Workshop*, pages 257–265, Pisa, Italy. 29
- Masselus, V., Peers, P., Dutré, P., and Willems, Y. D. (2003). Relighting with 4d incident light fields. *ACM Transactions on Graphics (Proc. SIGGRAPH 2003)*, 22(3):613–620. 2, 28
- Masselus, V., Peers, P., Dutré, P., and Willems, Y. D. (2004). Smooth reconstruction and compact representation of reflectance functions for image-based relighting. In *Rendering Techniques 2004 (Proc. Eurographics Symposium on Rendering)*, pages 287–298. 30, 39, 40, 68
- Matsushita, Y., Kang, S.-B., Lin, S., Shum, H.-Y., and Tong, X. (2004). Lighting and shadow interpolation using intrinsic lumigraphs. *Journal of Image and Graphics (IJIG)*, 4(4):585–604. 69
- Matusik, W. (2003). *A Data-Driven Reflectance Model*. PhD thesis, Massachusetts Institute of Technology. 23
- Matusik, W., Loper, M., and Pfister, H. (2004). Progressively-refined reflectance functions from natural illumination. In *Rendering Techniques 2004 (Proc. Eurographics Symposium on Rendering)*, pages 299–308. 3, 30, 94, 97
- Matusik, W. and Pfister, H. (2004). 3D TV: a scalable system for real-time acquisition, transmission, and autostereoscopic display of dynamic scenes. *ACM Transactions on Graphics (Proc. SIGGRAPH 2004)*, 23(3):814–824. 18, 101
- Matusik, W., Pfister, H., Brand, M., and McMillan, L. (2003a). A data-driven reflectance model. *ACM Transactions on Graphics (Proc. SIGGRAPH 2003)*, 22(3):759–769. 23
- Matusik, W., Pfister, H., Brand, M., and McMillan, L. (2003b). Efficient isotropic BRDF measurement. In *Rendering Techniques 2003 (Proc. Eurographics Symposium on Rendering)*, pages 241–248. 23, 46
- Matusik, W., Pfister, H., Ngan, A., Beardsley, P., Ziegler, R., and McMillan, L. (2002a). Image-based 3D photography using opacity hulls. *ACM Transactions on Graphics (Proc. SIGGRAPH 2002)*, 21(3):427–437. 32
- Matusik, W., Pfister, H., Ziegler, R., Ngan, A., and McMillan, L. (2002b). Acquisition and Rendering of Transparent and Refractive Objects. In *Thirteenth Eurographics Rendering Workshop*, pages 267–278. 32, 69

- McCool, M. D., Ang, J., and Ahmad, A. (2001). Homomorphic factorization of BRDFs for high-performance rendering. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 171–178, New York, NY, USA. ACM. 24
- McMillan, L. and Bishop, G. (1995). Plenoptic modeling: an image-based rendering system. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 39–46, New York, NY, USA. ACM. 17
- Mei, C., Shi, J., and Wu, F. (2004). Rendering with spherical radiance transport maps. *Computer Graphics Forum (Proc. Eurographics 2004)*, 23(3):281–290. 23
- Miller, G. and Hoffman, C. (1984). Illumination and reflection maps: Simulated objects in simulated and real environments. *SIGGRAPH 84 Advanced Computer Graphics Animation seminar notes*. 19
- Miller, G., Rubin, S., and Ponceleon, D. (1998). Lazy decompression of surface light fields for precomputed global illumination. *Rendering Techniques '98 (Proc. Eurographics Rendering Workshop)*, pages 281–292. 19
- Min, C., Wang, P., Jiao, X., Deng, Y., and Ming, H. (2007). Beam manipulating by metallic nano-optic lens containing nonlinear media. *Opt. Express*, 15(15):9541–9546. 117
- Mohan, A., Bailey, R., Waite, J., Tumblin, J., Grimm, C., and Bodenheimer, B. (2007). Tabletop computed lighting for practical digital photography. *IEEE Transactions on Visualization and Computer Graphics*, 13(04):652–662. 31
- Mohan, A., Tumblin, J., Bodenheimer, B., Grimm, C., and Bailey, R. (2005). Table-top computed lighting for practical digital photography. In *Rendering Techniques 2005 (Proc. Eurographics Symposium on Rendering)*, pages 165–172. 31, 36, 94
- Mueller, G., Sarlette, R., and Klein, R. (2007). Procedural editing of bidirectional texture functions. *Rendering Techniques 2007 (Proc. Eurographics Symposium on Rendering)*, pages 219–230. 33
- Müller, G., Sarlette, R., and Klein, R. (2006). Data-driven local coordinate systems for image-based rendering. *Computer Graphics Forum (Proc. Eurographics 2006)*, 25(3):369 – 378. 33
- Nakajima, S., Nakamura, K., Masamune, K., Sakuma, I., and T., T. D. (2001). Three-dimensional medical imaging display with computer-generated integral photography. *Computerized Medical Imaging and Graphics*, 25:235–241. 101
- Nayar, S. K., Belhumeur, P. N., and Boulton, T. E. (2004). Lighting sensitive display. *ACM Transactions on Graphics*, 23(4):963–979. 5, 31, 102
- Ng, R. (2006). *Digital light field photography*. PhD thesis, Stanford University. 18, 119

- Ng, R., Levoy, M., Brédif, M., Duval, G., Horowitz, M., and Hanrahan, P. (2005). Light field photography with a hand-held plenoptic camera. *Stanford University Computer Science Tech Report*. 18
- Ng, R., Ramamoorthi, R., and Hanrahan, P. (2003). All-frequency shadows using non-linear wavelet lighting approximation. *ACM Transactions on Graphics (Proc. SIGGRAPH 2003)*, 22(3):376–381. 22
- Ng, R., Ramamoorthi, R., and Hanrahan, P. (2004). Triple product wavelet integrals for all-frequency relighting. *ACM Transactions on Graphics (Proc. SIGGRAPH 2004)*, 23(3):477–487. 22
- Ngan, A. and Durand, F. (2006). Statistical acquisition of texture appearance. *Rendering Techniques 2006 (Proc. Eurographics Symposium on Rendering)*, pages 31–40. 33
- Ngan, A., Durand, F., and Matusik, W. (2005). Experimental analysis of BRDF models. *Rendering Techniques 2005 (Proc. Eurographics Symposium on Rendering)*, pages 117–126. 24
- Ngan, A., Durand, F., and Matusik, W. (2006). Image-driven navigation of analytical BRDF models. *Rendering Techniques 2006 (Proc. Eurographics Symposium on Rendering)*, pages 399–407. 24
- Nicodemus, F. E., Richmond, J. C., Hsia, J. H., Ginsberg, I. W., and Limperis, T. (1977). Geometrical considerations and nomenclature for reflectance. Monograph 161. Published by the U.S. Department of Commerce; National Bureau of Standards. 12
- Nimeroff, J., Simoncelli, E., and Dorsey, J. (1994). Efficient re-rendering of naturally illuminated environments. *Photorealistic Rendering Techniques (Proc. Eurographics Rendering Workshop)*, pages 373–388. 31
- Nishino, K. and Nayar, S. K. (2004). Eyes for relighting. *ACM Transactions on Graphics (Proc. SIGGRAPH 2004)*, 23(3):704–711. 19
- Nishino, K., Zhang, Z., and Ikeuchi, K. (2001). Determining reflectance parameters and illumination distribution from a sparse set of images for view-dependent image synthesis. *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, 1:599–606 vol.1. 25
- Noe, N. and Peroche, B. (2000). Hierarchical reconstruction of BRDFs using locally supported functions. *Computer Graphics Forum*, 19(2). 23
- Oh, B. M., Chen, M., Dorsey, J., and Durand, F. (2001). Image-based modeling and photo editing. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 433–442, New York, NY, USA. ACM. 20

- Okabe, M., Takayama, K., Ijiri, T., and Igarashi, T. (2007). Light shower: A poor man's light stage built with an off-the-shelf umbrella and projector. *SIGGRAPH Sketch*. 31
- Okabe, M., Zeng, G., Matsushita, Y., Igarashi, T., Quan, L., and Shum, H.-Y. (2006). Single-view relighting with normal map painting. *Pacific Graphics*. 25
- Oliveira, M. M. and Bishop, G. (1999). Image-based objects. In *I3D '99: Proceedings of the 1999 Symposium on Interactive 3D graphics*, pages 191–198, New York, NY, USA. ACM. 20
- Oren, M. and Nayar, S. K. (1994). Generalization of Lambert's reflectance model. In *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 239–246, New York, NY, USA. ACM. 23
- Overbeck, R., Ben-Artzi, A., Ramamoorthi, R., and Grinspun, E. (2006). Exploiting temporal coherence for incremental all-frequency relighting. *Rendering Techniques 2006 (Proc. Eurographics Symposium on Rendering)*, pages 151–160. 23
- Palmer, J. M. (2008). Radiometry and photometry FAQ. Retrieved on October 6th from: <http://www.optics.arizona.edu/Palmer/rpfaq/rpfaq.htm>. 8
- Pan, M., Wang, R., Liu, X., Peng, Q., and Bao, H. (2007). Precomputed radiance transfer field for rendering interreflections in dynamic scenes. *Computer Graphics Forum (Proc. Eurographics 2007)*, 26(3):485–493. 23
- Paris, S., Sillion, F., and Quan, L. (2003). Lightweight face relighting. In *Proceedings of Pacific Graphics*. 26
- Paterson, J. A., Claus, D., and Fitzgibbon, A. W. (2005). BRDF and geometry capture from extended inhomogeneous samples using flash photography. *Computer Graphics Forum (Proc. Eurographics 2005)*, 24(3):383–391. 25
- Peers, P. (2006). *Sampling Reflectance Functions for Image-based Relighting*. PhD thesis, Department of Computer Science, Katholieke Universiteit Leuven. 29
- Peers, P. and Dutré, P. (2003). Wavelet environment matting. In *Rendering Techniques 2003 (Proc. Eurographics Symposium on Rendering)*, pages 157–166. 3, 28, 29, 46
- Peers, P. and Dutré, P. (2005). Inferring reflectance functions from wavelet noise. In *Rendering Techniques 2005 (Proc. Eurographics Symposium on Rendering)*, pages 173–182. 3, 30, 64
- Peers, P., Hawkins, T., and Debevec, P. (2006). A reflective light stage. *ICT Technical Report ICT-TR-04.2006*. 29
- Peers, P., Tamura, N., Matusik, W., and Debevec, P. (2007). Post-production facial performance relighting using reflectance transfer. *ACM Transactions on Graphics (Proc. SIGGRAPH 2007)*, 26(3):52. 30

- Peter, I. and Straßer, W. (2001). Thewavelet stream: Interactive multi resolution light field rendering. *Rendering Techniques 2001 (Proc. Eurographics Rendering Workshop)*, pages 127–138. 18
- Phong, B. T. (1975). Illumination for computer generated pictures. *Commun. ACM*, 18(6):311–317. 23, 26
- Poulin, P. and Fournier, A. (1990). A model for anisotropic reflection. In *SIGGRAPH '90: Proceedings of the 17th annual conference on Computer graphics and interactive techniques*, pages 273–282, New York, NY, USA. ACM. 24
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (1992). *Numerical Recipes in C*. Cambridge University Press, Cambridge. 89
- Ramamoorthi, R. and Hanrahan, P. (2001). A signal-processing framework for inverse rendering. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 117–128, New York, NY, USA. ACM Press. 24, 43
- Ramamoorthi, R. and Hanrahan, P. (2004). A signal-processing framework for reflection. *ACM Transactions on Graphics*, 23(4):1004–1042. 24
- Raskar, R., Tan, K.-H., Feris, R., Yu, J., and Turk, M. (2004). Non-photorealistic camera: depth edge detection and stylized rendering using multi-flash imaging. *ACM Transactions on Graphics (Proc. SIGGRAPH 2004)*, 23(3):679–688. 31
- Robertson, M. A., Borman, S., and Stevenson, R. L. (2003). Estimation-theoretic approach to dynamic range improvement using multiple exposures. *Journal of Electronic Imaging*, 12(2):219–228. 11, 37
- Roth, S. and Black, M. J. (2006). Specular flow and the recovery of surface structure. In *Proc. the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 1869–1876. 69
- Rump, M., Müller, G., Sarlette, R., Koch, D., and Klein, R. (2008). Photo-realistic rendering of metallic car paint from image-based measurements. *Computer Graphics Forum (Proc. Eurographics 2008)*, 27:527 – 536. 33
- Rushmeier, H. E. and Torrance, K. E. (1990). Extending the radiosity method to include specularly reflecting and translucent materials. *ACM Transactions on Graphics*, 9(1):1–27. 22
- Rusinkiewicz, S. and Marschner, S. (2000). Measurement i - BRDFs. 12
- Rusinkiewicz, S. M. (1998). A new change of variables for efficient BRDF representation. *Rendering Techniques '98 (Proc. Eurographics Rendering Workshop)*, pages 11–22. 23

- Sato, I., Sato, Y., and Ikeuchi, K. (1999a). Acquiring a radiance distribution to superimpose virtual objects onto a real scene. *IEEE Transactions on Visualization and Computer Graphics*, 5(01):1–12. 19
- Sato, I., Sato, Y., and Ikeuchi, K. (1999b). Illumination distribution from shadows. *Computer Vision and Pattern Recognition, 1999. IEEE Computer Society Conference on.*, 1:306–312 Vol. 1. 19
- Sato, I., Sato, Y., and Ikeuchi, K. (2003). Illumination from shadows. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 25(3):290–300. 19
- Sato, Y., Wheeler, M. D., and Ikeuchi, K. (1997). Object shape and reflectance modeling from observation. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 379–387, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 25
- Scharstein, H., Krotz-Vogel, W., and Scharstein, D. (1996). Digital sundial. In *US patent 5 590 093, German patent 4 431 817*. 102, 109
- Schaufler, G. (1998). Per-object image-warping with layered impostors. *Rendering Techniques '98 (Proc. Eurographics Rendering Workshop)*, pages 145–156. 20
- Schechner, Y. Y., Nayar, S. K., and Belhumeur, P. N. (2003). A theory of multiplexed illumination. In *ICCV '03: Proceedings of the Ninth IEEE International Conference on Computer Vision*, page 808, Washington, DC, USA. IEEE Computer Society. 29
- Schirmacher, H., Heidrich, W., Rubick, M., Schiron, D., and Seidel, H.-P. (1999a). Image-based BRDF reconstruction. In *Proc. of the 4th VMV Conference*, pages 285–292. 24
- Schirmacher, H., Heidrich, W., and Seidel, H.-P. (1999b). Adaptive acquisition of lumigraphs from synthetic scenes. *Computer Graphics Forum (Proc. Eurographics)*, 18(3):151–160. 18, 46
- Schirmacher, H., Heidrich, W., and Seidel, H.-P. (2000). High-quality interactive lumigraph rendering through warping. *Proceedings of Graphics Interface 2000*, pages 87–94. 18
- Schirmacher, H., Ming, L., and Seidel, H.-P. (2001). On-the-fly processing of generalized lumigraphs. *Computer Graphics Forum (Proc. Eurographics 2001)*, 20(3). 18
- Schoeneman, C., Dorsey, J., Smits, B., Arvo, J., and Greenberg, D. (1993). Painting with light. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 143–146, New York, NY, USA. ACM. 31
- Sen, P., Chen, B., Garg, G., Marschner, S. R., Horowitz, M., Levoy, M., and Lensch, H. P. A. (2005). Dual photography. *ACM Transactions on Graphics (Proc. SIGGRAPH 2005)*, 24(3):745–755. 28, 46

- Shannon, C. E. (1949). Communication in the presence of noise. In *Proc. IRE*, volume 37, pages 10–21. 43
- Sloan, P.-P., Kautz, J., and Snyder, J. (2002). Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics (Proc. SIGGRAPH 2002)*, 21(3):527–536. 22
- Sloan, P.-P., Liu, X., Shum, H.-Y., and Snyder, J. (2003). Bi-scale radiance transfer. *ACM Transactions on Graphics (Proc. SIGGRAPH 2003)*, 22(3):370–375. 33
- Sun, B., Sunkavalli, K., Ramamoorthi, R., Belhumeur, P. N., and Nayar, S. K. (2007). Time-varying BRDFs. *IEEE Transactions on Visualization and Computer Graphics*, 13(03):595–609. 24
- Sun, Z. and Kim, H. K. (2004). Refractive transmission of light and beam shaping with metallic nano-optic lenses. *Appl. Phys. Lett.*, 85:652. 117
- Sunkavalli, K., Matusik, W., Pfister, H., and Rusinkiewicz, S. (2007). Factored time-lapse video. *ACM Transactions on Graphics (Proc. SIGGRAPH 2007)*, 26(3):101. 32
- Suykens, F., Berge, K., Lagae, A., and Dutré, P. (2003). Interactive rendering with bidirectional texture functions. *Computer Graphics Forum (Proc. Eurographics 2003)*, 22(3). 33
- Tan, P., Lin, S., Quan, L., Guo, B., and Shum, H.-Y. (2005). Multiresolution reflectance filtering. *Rendering Techniques 2005 (Proc. Eurographics Symposium on Rendering)*, pages 111–116. 26
- Tan, P., Lin, S., Quan, L., Guo, B., and Shum, H.-Y. (2008). Filtering and rendering of resolution-dependent reflectance models. *IEEE Transactions on Visualization and Computer Graphics*, 14(2):412–425. 26
- Tchou, C., Stumpfel, J., Einarsson, P., Fajardo, M., and Debevec, P. (2004). Unlighting the parthenon. In *SIGGRAPH 2003, Sketches and Applications*. 25, 93
- Teo, P. C., Simoncelli, E. P., and Heeger, D. J. (1997). Efficient linear re-rendering for interactive lighting design. Technical Report STAN-CS-TN-97-60, Stanford University, Stanford, CA. 31
- Theobalt, C., Ahmed, N., Lensch, H., Magnor, M., and Seidel, H.-P. (2007). Seeing people in different light – joint shape, motion, and reflectance capture. *IEEE Transactions on Visualization and Computer Graphics*, 13(04):663 – 674. 26
- Tominaga, S. and Tanaka, N. (2000). Estimating reflection parameters from a single color image. *IEEE Computer Graphics and Applications*, 20(05):58–66. 24
- Tong, X., Zhang, J., Liu, L., Wang, X., Guo, B., and Shum, H.-Y. (2002). Synthesis of bidirectional texture functions on arbitrary surfaces. *ACM Transactions on Graphics (Proc. SIGGRAPH 2002)*, 21(3):665–672. 33

- Unger, J., Wenger, A., Hawkins, T., Gardner, A., and Debevec, P. (2003). Capturing and rendering with incident light fields. In *Rendering Techniques 2003 (Proc. Eurographics Symposium on Rendering)*, pages 141–149. 18
- Veeraraghavan, A., Raskar, R., Agrawal, A., Mohan, A., and Tumblin, J. (2007). Dappled photography: mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *ACM Transactions on Graphics (Proc. SIGGRAPH 2007)*, 26(3):69. 18
- Vlasic, D., Pfister, H., Molinov, S., Grzeszczuk, R., and Matusik, W. (2003). Opacity light fields: interactive rendering of surface light fields with view-dependent opacity. In *I3D '03: Proceedings of the 2003 symposium on interactive 3D graphics*, pages 65–74, New York, NY, USA. ACM. 19
- Wang, H., Sun, M., and Yang, R. (2007a). Space-time light field rendering. *IEEE Transactions on Visualization and Computer Graphics*, 13(04):697–710. 18
- Wang, J., Tong, X., Lin, S., Pan, M., Wang, C., Bao, H., Guo, B., and Shum, H.-Y. (2006). Appearance manifolds for modeling time-variant appearance of materials. *ACM Transactions on Graphics (Proc. SIGGRAPH 2006)*, 25(3):754–761. 26
- Wang, J., Tong, X., Snyder, J., Chen, Y., Guo, B., and Shum, H.-Y. (2005a). Capturing and rendering geometry details for BTF-mapped surfaces. *The Visual Computer*, 21(8–10):559–568. 33
- Wang, L., Lin, S., Lee, S., Guo, B., and Shum, H.-Y. (2005b). Light field morphing using 2D features. *IEEE Transactions on Visualization and Computer Graphics*, 11(01):25–34. 18
- Wang, O., Davis, J., Chuang, E., Rickard, I., de Mesa, K., and Dave, C. (2008). Video relighting using infrared illumination. *Computer Graphics Forum (Proc. Eurographics 2008)*, 27:271 – 279. 30
- Wang, R., Tran, J., and Luebke, D. (2004). All-frequency relighting of non-diffuse objects using separable BRDF approximation. *Rendering Techniques 2004 (Proc. Eurographics Symposium on Rendering)*, pages 345–354. 23
- Wang, R., Zhu, J., and Humphreys, G. (2007b). Precomputed radiance transfer for real-time indirect lighting using a spectral mesh basis. *Rendering Techniques 2007 (Proc. Eurographics Symposium on Rendering)*, pages 13–21. 22
- Ward, G. J. (1992). Measuring and modeling anisotropic reflection. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 265–272, New York, NY, USA. ACM. 24
- Weistroffer, R. P., Walcott, K. R., Humphreys, G., and Lawrence, J. (2007). Efficient basis decomposition for scattered reflectance data. *Rendering Techniques 2007 (Proc. Eurographics Symposium on Rendering)*, pages 207–218. 25

- Wenger, A., Gardner, A., Tchou, C., Unger, J., Hawkins, T., and Debevec, P. (2005). Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactions on Graphics (Proc. SIGGRAPH 2005)*, 24(3):756–764. 4, 29
- Westin, S. H., Arvo, J. R., and Torrance, K. E. (1992). Predicting reflectance functions from complex surfaces. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 255–264, New York, NY, USA. ACM. 23
- Wexler, Y., Fitzgibbon, A. W., and Zisserman, A. (2002). Image-based environment matting. In *Rendering Techniques 2002 (Proc. Eurographics Rendering Workshop)*, pages 279–290. 30
- Wexler, Y., Shechtman, E., and Irani, M. (2004). Space-Time Video Completion. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'04)*, volume 1, pages 120–127. 69, 72, 77
- Weyrich, T., Matusik, W., Pfister, H., Bickel, B., Donner, C., Tu, C., McAndless, J., Lee, J., Ngan, A., Jensen, H. W., and Gross, M. (2006). Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics (Proc. SIGGRAPH 2006)*, 25(3):1013–1024. 26
- Weyrich, T., Pfister, H., and Gross, M. (2005). Rendering deformable surface reflectance fields. *IEEE Transactions on Visualization and Computer Graphics*, 11(1):48–58. 32
- Whitaker, R. T. (2000). A level-set approach to image blending. *IEEE Transactions on Image Processing*, 9:1849–1861. 69, 70, 74
- Whitted, T. (1979). An improved illumination model for shaded display. *ACM SIGGRAPH Computer Graphics (SIGGRAPH 1979)*, 13(2):14. 21
- Whitted, T. (1980). An improved illumination model for shaded display. *Communications of the ACM*, 23(6):343–349. 21
- Williams, L. (1978). Casting curved shadows on curved surfaces. *Computer Graphics (SIGGRAPH '78 Proceedings)*, 12(3):270–274. 21
- Winnemöeller, H., Mohan, A., Tumblin, J., and Gooch, B. (2005). Light waving: Estimating light positions from photographs alone. *Computer Graphics Forum (Proc. Eurographics)*, 24(3):433–438. 29
- Wong, T.-T., Heng, P.-A., and Fu, C.-W. (2001). Interactive relighting of panoramas. *IEEE Computer Graphics and Applications*, 21(02):32–41. 28
- Wong, T.-T., Heng, P.-A., Or, S.-H., and Ng, W.-Y. (1997). Image-based rendering with controllable illumination. *Rendering Techniques '97 (Proc. Eurographics Rendering Workshop)*, pages 13–. 27

- Wood, D. N., Azuma, D. I., Aldinger, K., Curless, B., Duchamp, T., Salesin, D. H., and Stuetzle, W. (2000). Surface light fields for 3D photography. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 287–296, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 19
- Xu, H. and Sun, Y. (2006). Compact representation of spectral BRDFs using fourier transform and spherical harmonic expansion. *Computer Graphics Forum*, 25(4):759–775. 24
- Xu, K., Jia, Y.-T., Fu, H., Hu, S.-M., and Tai, C.-L. (2008). Spherical piecewise constant basis functions for all-frequency precomputed radiance transfer. *IEEE Transactions on Visualization and Computer Graphics*, 14(2):454–467. 22
- Yang, J. C., Everett, M., Buehler, C., and McMillan, L. (2002). A real-time distributed light field camera. *Rendering Techniques 2002 (Proc. Eurographics Rendering Workshop)*, pages 77–86. 18
- Yu, T., Wang, H., Ahuja, N., and Chen, W.-C. (2006). Sparse lumigraph relighting by illumination and reflectance estimation from multi-view images. *Rendering Techniques 2006 (Proc. Eurographics Symposium on Rendering)*, pages 41–50. 26
- Yu, Y., Debevec, P., Malik, J., and Hawkins, T. (1999). Inverse global illumination: recovering reflectance models of real scenes from photographs. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 215–224, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 25
- Yu, Y. and Malik, J. (1998). Recovering photometric properties of architectural scenes from photographs. In *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 207–217, New York, NY, USA. ACM. 25
- Zhang, Z., Wang, L., Guo, B., and Shum, H.-Y. (2002). Feature-based light field morphing. *ACM Transactions on Graphics (Proc. SIGGRAPH 2002)*, 21(3):457–464. 19
- Zhou, K., Du, P., Wang, L., Matsushita, Y., Shi, J., Guo, B., and Shum, H.-Y. (2005). Decorating surfaces with bidirectional texture functions. *IEEE Transactions on Visualization and Computer Graphics*, 11(05):519–528. 33
- Zickler, T., Enrique, S., Ramamoorthi, R., and Belhumeur, P. (2005). Reflectance sharing: Image-based rendering from a sparse set of images. In *Rendering Techniques 2005 (Proc. Eurographics Symposium on Rendering)*, pages 253–264. 25, 68
- Zickler, T., Ramamoorthi, R., Enrique, S., and Belhumeur, P. N. (2006). Reflectance sharing: Predicting appearance from a sparse set of images of a known shape. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(8):1287–1302. 25, 68

- Ziegler, R., Bucheli, S., Ahrenberg, L., Magnor, M., and Gross, M. (2007). A bidirectional light field - hologram transform. *Computer Graphics Forum (Proc. Eurographics 2007)*, 26(3):435–446. 18
- Zongker, D. E., Werner, D. M., Curless, B., and Salesin, D. H. (1999). Environment matting and compositing. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 205–214, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 28, 29

Index

- 4 → 2 reflectance field, 15
- 4 → 4 reflectance field, 14
- 2 → 2 reflectance field, 15
- 2 → 4 reflectance field, 14

- acquisition step, 16, 85
- apparent BRDF, 27
- appearance acquisition, 1

- Bayer pattern, 10
- bi-directional path tracing, 22
- Bi-Directional Reflectance Distribution Function, 13
- Bi-Directional Scattering-Surface Reflectance Distribution Function, 12
- Bi-Directional Texture Function, 31
- Bi-Directional Transmittance Distribution Function, 13
- billboards, 19
- BRDF, 13
- BRRDF, 26
- BSSRDF, 12
- BTDF, 13
- BTF, 31

- color channel, 9
- computer graphics, 1

- data-driven, 23
- digital camera, 10
- digital imaging sensor, 10
- directional light source, 21
- dynamic range, 11

- energy, 8
- environment map, 19
- environment matte, 27

- flux density, 8
- form factor, 22
- frame buffer, 20

- ghosting artefact, 18
- global illumination, 22
- Global Reflection Distribution Function, 26
- GRDF, 26

- HDR imaging, 11
- high dynamic range imaging, 11

- image, 10
- image-based rendering, 19
- image-space relighting, 2, 16, 27
- impostor, 19
- inverse rendering, 24
- irradiance, 8

- Lambertian BRDF, 13, 23
- light field, 10
- light probe, 19
- light stage, 28
- light transport, 7
- linearity of light transport, 12
- lumigraph, 18

- modeling step, 16, 86

- non-photorealistic rendering, 30
- NPR, 30

- painter's algorithm, 20
- phase, 9
- photometry, 7
- photon, 8
- photon mapping, 22

- picture elements, 10
- pixel, 10
- pixel positions, 10
- plenoptic function, 9
- point light source, 21
- polarization, 9
- pre-computed radiance transfer, 22
- prediction step, 86
- primary colors, 9
- PRT, 22

- quantization, 11

- radiance, 8
- radiant exitance, 8
- radiant flux, 8
- radiant intensity, 8
- radiant power, 8
- radiation, 8
- radiometry, 7
- radiosity, 8
- radiosity method, 22
- ray tracing, 21
- reflectance equation, 14
- reflectance equation for distant illumination, 14
- reflectance equation for image-space relighting, 15
- reflectance field, 4, 14, 26
- reflection equation, 12
- relighting, 2
- rendering, 1, 11
- rendering equation, 13
- rendering step, 16, 86
- response function, 11

- sampling, 10
- shader, 20, 21
- shadow map, 21
- shadow volume, 21
- Shift-Variant Bi-Directional Reflectance Distribution Function, 12
- Shift-Variant Bi-Directional Transmittance Distribution Function, 13

- Spatially Varying Bi-Directional Reflectance Distribution Function, 13
- spectral radiance, 8
- superposition principle, 12
- surface albedo, 13
- surface light field, 18
- surface reflectance field, 31
- SVBRDF, 12
- SVBTDF, 13

- texturing, 21
- training step, 85

- visibility problem, 20
- visual hull, 31

- wavelength, 8

- Z buffer, 20