

Realistic Materials and Illumination Environments

by

Abhijeet Ghosh

B.E., Gujarat University, India, 2000
M.S., Stony Brook University, N.Y., 2002

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

The Faculty of Graduate Studies

(Computer Science)

The University Of British Columbia

June, 2007

© Abhijeet Ghosh 2007

Abstract

Throughout its history, the field of computer graphics has been striving towards increased realism. This goal has traditionally been described by the notion of *photo-realism*, and more recently and in many cases the more ambitious goal of *perceptual realism*. Photo-realistic image synthesis involves many algorithms describing the phenomena of light transport in a scene as well as its interaction with various materials. On the other hand, research in perceptual realism typically involves various tone mapping algorithms for display devices as well as algorithms that mimic the natural response of the human visual system in order to recreate the visual experience of a real scene.

An important aspect of realistic rendering is the accurate modeling of the scene elements such as light sources and material reflectance properties. This dissertation proposes a set of new techniques for efficient acquisition of material properties as well as new algorithms for high quality rendering with acquired data. Here, we are mostly concerned with the acquisition and rendering of local illumination effects. In particular, we propose a new optical setup for efficient acquisition of the bidirectional reflectance distribution function (BRDF) with basis illumination and various Monte Carlo strategies for efficient sampling of direct illumination.

The dissertation also looks into the display end of the image synthesis pipeline and proposes algorithms for displaying scenes on high dynamic range (HDR) displays for visual realism, and for tying the room illumination with the viewing environment for a sense of presence and immersion in a virtual environment. Here, we develop real-time rendering algorithms for driving the HDR displays as well as for active control of room illumination based on dynamic scene content. Thus, we propose contributions to the acquisition, rendering, and display end of the image synthesis pipeline while targeting real-time rendering applications, as well as high quality off-line rendering with realistic

materials and illumination environments.

Contents

Abstract	ii
Contents	iv
List of Tables	vii
List of Figures	viii
Acknowledgements	xii
1 Introduction	1
1.1 Realistic Image Synthesis Pipeline	2
1.2 Chapter Overview	5
2 Background and Related Work	7
2.1 The Bidirectional Reflectance Distribution Function	7
2.1.1 Analytic BRDF Models	9
2.1.2 BRDF Acquisition	13
2.1.3 Basis Representations and Illumination	15
2.2 Monte Carlo Integration and Sampling	17
2.2.1 Monte Carlo Integration	17
2.2.2 Sampling of Direct Illumination	20
2.3 Display Algorithms and Viewing Conditions	25
2.3.1 Dynamic Range of Display Devices	26
2.3.2 Tone Mapping and Visual Adaptation	26
2.3.3 Perceptual Studies on Room Lighting	27

2.3.4	Bridging Real and Virtual Illumination	28
3	BRDF Acquisition with Basis Illumination	30
3.1	Overview	31
3.2	Measurement with Basis Functions	33
3.2.1	Orthonormal Zonal Basis	34
3.2.2	Basis Conversion to Spherical Harmonics	36
3.2.3	Fitting Analytical Reflection Models	37
3.3	Measurement Setup and Calibration	38
3.3.1	Calibration	41
3.4	Results	43
3.4.1	Rendering with Acquired BRDFs	46
3.5	Discussion	50
4	Sampling Techniques for Direct Illumination	53
4.1	Importance Sampling of Direct Illumination	54
4.2	Bidirectional Importance Sampling	56
4.2.1	Sample Generation through Rejection	57
4.2.2	Sample Generation through SIR	59
4.2.3	Bidirectional Sampling Results	60
4.3	Correlated Visibility Sampling	63
4.3.1	Bias	68
4.3.2	Correlated Sampling Results	69
4.4	Discussion	72
5	Sequential Sampling of Environment Maps	75
5.1	Sequential Monte Carlo Sampling	77
5.1.1	Direct Illumination Estimate	81
5.2	Variance Reduction with Intermediate Distributions	83
5.3	Unoccluded Illumination Estimate with Path Sampling	86
5.4	Implementation	87
5.5	Results	88

5.6	Discussion	94
6	Real-time Rendering for HDR Displays	95
6.1	Human Contrast Perception	97
6.1.1	Local Contrast Perception	97
6.2	Driving the Projector Display	98
6.3	Driving the LED Display	101
6.4	Applications	103
6.5	Discussion	104
7	Real Illumination from Virtual Environments	107
7.1	Method	108
7.1.1	Physical Setup	108
7.1.2	Calibration	111
7.1.3	Rendering	113
7.2	Content Creation	114
7.3	Experiments	116
7.4	Discussion	121
8	Conclusions	124
	Bibliography	129
A	Radiometric and Photometric Terms	143
B	Zonal Basis Function Plots	145
C	Normalization Constant	146

List of Tables

5.1	Summary of notation used in Chapter 5.	78
-----	------------------------------------------------	----

List of Figures

1.1	Overview of realistic image synthesis pipeline.	2
2.1	The BRDF geometry.	8
2.2	The Ward imaging gonio-reflectometer.	14
2.3	Sampling via CDF inversion.	20
3.1	Physical setup and 2-D mockup of BRDF acquisition device.	32
3.2	The measurement zone \mathbf{Z}	33
3.3	Suppression of ringing artifacts through fitting analytical BRDF models.	38
3.4	Photograph of the proposed BRDF acquisition setup.	39
3.5	Iterative design process for the profile of the reflective dome.	40
3.6	Plot of measured coefficients of gray card vs. that of pure Lambertian diffuse reflectance.	42
3.7	Representative set of BRDFs acquired with lower order zonal basis functions.	43
3.8	Specular chocolate wrapping papers acquired using higher order zonal basis functions, and then fit to an analytical model for rendering.	44
3.9	Visual comparison of two kinds of acquired satin samples wrapped around a cylinder as lit by a point source against real photographs. Left column: Photographs of the red and blue satin samples. Right column: Rendering of the D-BRDF fit to the acquired data.	45
3.10	Real-time rendering of acquired BRDFs.	47
3.11	Comparison of importance sampling strategies for acquired BRDF in high frequency HDR lighting of Grace Cathedral EM.	48

3.12	Various BRDFs acquired with our prototype setup.	51
3.13	Audi-TT model rendered with acquired paint BRDFs.	52
4.1	Angular plots of the importance function of EM, BRDF and product. .	56
4.2	Sample generation by rejection sampling	58
4.3	Sampling-importance resampling (SIR)	59
4.4	Rejection and SIR quality comparison	61
4.5	Convergence plot of RMS error for bidirectional sampling.	62
4.6	Quality comparison of SIR with previous techniques	63
4.7	Quality comparison of bidirectional sampling and correlated visibility sampling.	64
4.8	Visibility masks for correlated sampling.	65
4.9	Lens perturbation within a 5×5 transition tile.	67
4.10	Visual comparison of biased and unbiased correlated visibility sampling.	69
4.11	Quality comparison of correlated sampling with previous techniques. .	70
4.12	Dilated vs. undilated mask.	72
4.13	Dragon model rendered with and without dilation of the visibility mask.	74
5.1	Quality comparison of our SMC sampling algorithm with bidirectional importance sampling for a sky probe sequence.	76
5.2	Quality comparison of SMC sampling with single distribution vs se- quence of intermediate distributions at every time step.	89
5.3	Convergence plots of RMS errors for single distribution with multi- ple MCMC moves and sequence of intermediate distributions with one MCMC move each.	90
5.4	Quality comparison between bidirectional sampling and our SMC sam- pling algorithm for a specular BRDF in the sky probe gallery sequence.	91
5.5	Quality comparison between bidirectional sampling and our SMC sam- pling algorithm for a BRDF with a diffuse component in the sky probe gallery sequence.	92

5.6	SMC sampling for a camera animation sequence with the David model in the Grace Cathedral EM.	93
6.1	Two alternative designs for HDR displays.	96
6.2	The PSF of the human eye.	98
6.3	Response function of both the LCD panel and the DLP projector in the projector-based display.	99
6.4	Rendering algorithm for the projector-based display.	100
6.5	Point spread function of the projector.	100
6.6	Factoring an HDR image for projector-based display.	101
6.7	Rendering algorithm for the LED-based display.	102
6.8	Factoring an HDR photograph for LED-based display.	103
6.9	Comparison of image displayed on projector-based HDR display vs. a regular display	104
6.10	Screen photographs of HDR display applications	106
7.1	Room layout for active illumination control.	109
7.2	Light pattern generated by an iColor Cove light	110
7.3	Opened iColor Cove light source.	110
7.4	Light probe images acquired for each of the 24 light sources at the intended viewing position.	111
7.5	Room lighting corresponding to NFS underground 2 driving game. . .	114
7.6	Room lighting corresponding to HDR driving video.	115
7.7	Room lighting corresponding to HDR environment map.	116
7.8	User preferences regarding constant or uniform dynamic illumination for HDR video.	118
7.9	User preferences for directional vs. uniform illumination in an HDR panorama viewer.	119
7.10	User preferences for directional vs. uniform illumination when watching low-dynamic-range video game footage.	120
7.11	Photograph of room housing the active illumination system.	123

- B.1 The plots of zonal basis functions Z_l^m defined over the measurement space $[\pi/20, \pi/2] \times [0, 2\pi]$, for $l \leq 2$ 145

Acknowledgements

First and foremost, I thank my supervisor Assoc. Prof. Wolfgang Heidrich for his mentoring and keen involvement and guidance on all aspects of this dissertation. I also thank the members of my supervisory and examining committees, Prof. Bob Woodham, Assist. Prof. Robert Bridson, Prof. David Lowe, Prof. Peter Lawrence and Prof. Philip Dutré for taking interest in my work and allocating the time to review this thesis and for their suggestions. I would like to particularly thank Matthew Trentacoste, Helge Seetzen, David Burke, Assoc. Prof. Arnaud Doucet, and Shruthi Achutha with whom I have collaborated on various projects as part of this dissertation. I gained valuable experience while supervising the undergraduate honours thesis of Matthew O'Toole. Ciaran Llachlan Leavitt deserves special mention for proof-reading all my papers. Thanks also to Dr. Paul Debevec for providing the HDR environment maps. This dissertation work was supported by a Precarn fellowship (2003-04), an ATI Technologies fellowship (2004-05), and a UBC university graduate fellowship (2005-06).

I would also like to thank all my past and present colleagues and faculty members in the Imager Lab (large and small), the PSM Lab and the LCI lab for providing a stimulating environment for conducting research at UBC. Vladislav Kraevoy, David Burke, Kristian Hildebrand, Ciaran Llachlan Leavitt, Fred Kimberley, James Slack, Chen Yang, Tyson Brochu, Matthew Trentacoste, Derek Bradley, Brad Atcherson, Cheryl Lau and Allan Rempel made Imager a really fun place to be. Thanks to my fellow residents of *Die Kommune* - Eric Brochu, Hendrik Kueck, David Pritchard, Florence Bertails, Tamy Boubekour, and frequent flyers Tyson Brochu and Gillian Mahoney for making life in Vancouver that much more interesting! Finally, I take this opportunity to thank my parents, Anushree and Anup Kumar Ghosh, for all their support and encouragement over the years. This dissertation is dedicated to them.

Chapter 1

Introduction

The pursuit of *realism* has been a central theme throughout the history of computer graphics. The goal of realism in computer generated imagery has traditionally been described by the notion of *photo-realism* - creating or rendering images that are indistinguishable from real photographs. Applications of realistic rendering include entertainment such as movies and games, simulation and virtual reality applications, architectural design, and scientific visualization. More recently, the scope of realism in image synthesis has been extended to the notion of *perceptual realism* - synthesizing scenes that are perceived as real by a human observer. The goal of perceptual realism is, in many cases, even more challenging than photo-realism as its aim is to afford the viewer the same visual experience as if they were actually immersed in the scene.

Photo-realistic rendering typically involves the development of algorithms for the simulation of physically accurate light transport in a scene. However, the quality of rendering produced by these algorithms is limited by the quality of the input scene descriptions such as materials and illumination models. Therefore, photo-realistic rendering also requires accurate modeling of scene attributes such as geometry, lighting, and material properties, such as surface reflectance. Realistic illumination and reflectance can be very complex and hard to describe analytically. Therefore, the best way to obtain high quality representation of complex illumination and materials is by acquiring real world data through measurements.

At the other end of the image synthesis pipeline, images rendered using realistic material and illumination representations need to be shown on display devices. Since displayed images are meant to be viewed by a human observer, this stage typically needs to account for the characteristics of the human visual system along with that of

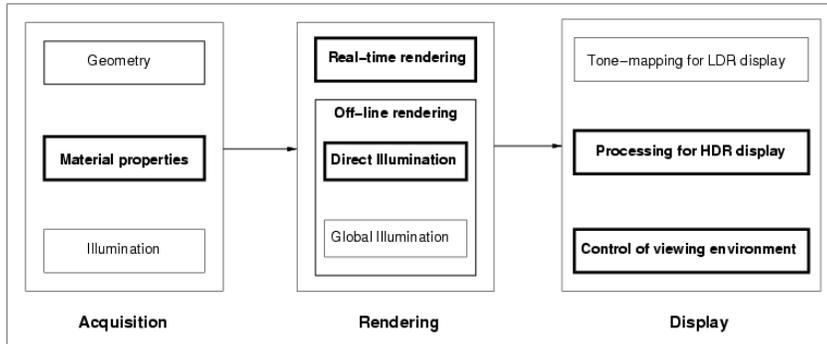


Figure 1.1: Overview of stages of the realistic image synthesis pipeline including acquisition, rendering and display. Individual components of the pipeline that are discussed as part of this dissertation are highlighted.

the display device. Hence, the notion of perceptual realism assumes importance for image display.

1.1 Realistic Image Synthesis Pipeline

Figure 1.1 presents a general overview of the realistic image-synthesis pipeline as discussed in this dissertation. As illustrated, the three major stages of the pipeline are: acquisition, rendering and display. The figure also lists the main components of the pipeline while highlighting those components toward which we make a contribution in this dissertation. As discussed above, photo-realistic rendering typically involves the acquisition and rendering stages of the pipeline, while perceptually realistic rendering involves the display stage.

Acquisition and Rendering for Photo-Realism: With the advances within the field of digital photography over the last decade, there has been significant interest in acquiring material and illumination models from photographs. This acquisition method has led to the development of image-based modeling and rendering techniques for realistic rendering. Image-based acquisition techniques have become very popular with the development of high dynamic range (HDR) imaging techniques [112]. This is par-

ticularly true for acquisition of complex real world illumination in the form of HDR environment maps [29]. These environment maps, also called radiance maps, represent the absolute intensities of the surrounding illumination, and can be used to render scenes with natural illumination, as well as for realistic relighting of virtual objects with major applications in movies and games. However, the availability of such a representation for complex natural illumination requires its integration into existing rendering systems, as well as the development of new algorithms for efficient rendering with such a representation.

In recent years, image-based techniques have also become popular for acquisition of material properties, such as reflectance functions and reflectance fields. However, most previous work has been conducted under highly controlled lighting conditions with multiple light sources carefully positioned to cover a hemisphere of incident directions [84]. The material samples sometimes also need to be translated with respect to the light sources using some mechanical gantry for dense sampling of incident illumination [85, 93]. Hence, acquiring reflectance properties of materials has typically been on the scale of a few hours of acquisition time in a dedicated environment. Also, the acquired data are huge and typically suffer from noise in the measurements resulting in the requirement of post-processing of the data and resampling into a more compact form for usage in rendering.

Image-based techniques are widely used for real-time photorealistic rendering applications such as games. The challenge here is to integrate the huge amounts of acquired illumination and material data into the real-time rendering framework. Many researchers have used frequency-domain analysis for encoding the acquired data into compact basis function representations for real-time rendering. Examples of popularly used basis functions are *spherical harmonics* [110, 142] and *wavelets* [72, 92]. Such basis function approaches are now becoming popular in commercially available games, and 3D programming APIs such as DirectX [11].

This dissertation proposes a set of new techniques for efficient acquisition of material properties as well as new algorithms for high quality rendering with acquired data. Here, we are mostly concerned with the acquisition and rendering of local illumination effects. In particular, we propose a new optical setup for efficient image-based

acquisition of the bidirectional reflectance distribution function (BRDF) using curved reflective surfaces, such that there are no moving parts, and the entire setup fits in a small, enclosed area. In addition, in order to speed up acquisition time and to avoid any post-processing of the data, we propose directly *optically* sampling the BRDF data into a suitable basis function during acquisition. The aim here is to acquire the data in terms of coefficients of a suitable basis function that can then be used to reconstruct the reflectance function during rendering. The advantage of such an acquisition is that the data are directly low-pass filtered during acquisition and encoded into a compact form that can be used in rendering systems within *a few minutes*. The dissertation also proposes various Monte Carlo strategies for sampling the acquired reflectance data as well as natural HDR lighting for efficient rendering of scenes with complex direct illumination.

Display for Perceptual Realism: Perceptually realistic rendering is mainly concerned with algorithms that aim to recreate the visceral experience of a real scene for a human observer. This includes the development of appropriate tone-mapping algorithms for display devices (e.g., [35, 73, 111, 116, 135]), as well as development of algorithms that mimic the natural response of the human visual system to displayed scenes [39, 60, 99, 114]. Tone-mapping algorithms map the computed scene radiance values to the dynamic range of a display device while preserving the contrast and perceived scene intensities. Perceptually realistic rendering has to also account for the response of the visual system to these intensities, such as light adaptation and glare simulations.

One of the most significant advances in perceptual realism is the development of high dynamic range display technology [119]. The human visual system can simultaneously capture approximately five orders of magnitude of dynamic range which cannot be represented by standard display devices. Tone mapping operators for these display devices alleviate this problem to an extent but they cannot recreate the natural adaptive responses in the visual system to such dynamic range. HDR displays are thus the key to perceptually realistic representation of real world HDR scenes. As part of this dissertation, we also develop algorithms for processing HDR images in real-time for driving

the HDR display devices, and demonstrate its use in many potential applications such as photo-realistic rendering and medical imaging.

It is our hypothesis that the HDR display acts as a window into a virtual world. However, this sense of realism and immersion in a virtual environment is lost outside this window as the HDR display cannot account for the viewing conditions. A true sense of immersion can only be achieved if the illumination levels in the real and virtual world are compatible. As part of this dissertation, we aim to connect the room illumination to the viewing environment to achieve a sense of presence and immersion in a virtual environment.

Hence, this dissertation aims to take both the aspects of realism, photo-realism and perceptual realism, into account while targeting real-time rendering applications as well as high-quality offline renderings with realistic materials and illumination environments.

1.2 Chapter Overview

The remainder of this thesis is organized as follows. Chapter 2 discusses the relevant background for the topics covered in this dissertation. In Chapter 3, we describe our novel image-based BRDF acquisition technique for acquiring basis reflectance of material samples. We develop a set of appropriate basis functions for this purpose and discuss how to project the acquired data into either spherical harmonics or an analytical model as a way of extrapolation for rendering. We also present a working prototype and some acquisition results.

We then discuss a series of Monte Carlo techniques for efficient sampling of direct illumination targeting high quality offline rendering in Chapters 4 and 5. In Section 4.2, we discuss how sampling from the *product* distribution of illumination and the reflectance function, which we call *bidirectional* sampling, is more efficient than sampling from the individual distributions. We discuss two Monte Carlo strategies for sampling according to the bidirectional importance. In Section 4.3, we extend our bidirectional sampling strategy to also account for the binary visibility function. In this case, we employ Metropolis sampling to establish a correlation in the energy estimate

of neighboring pixels on the image plane, thereby reducing noise in partially occluded regions. In Chapter 5, we extend product distribution sampling in the temporal domain to efficiently generate coherent samples for a video sequence in the presence of dynamic illumination.

Real-time perceptually realistic rendering is the theme of Chapters 6 and 7. We present our algorithm for processing HDR images in real-time for driving the HDR displays in Chapter 6 and demonstrate its use in many potential applications. In Chapter 7, we propose to tie the room illumination into the display and viewing environment. Here, we actively control the room illumination according to illumination in a virtual environment, thereby triggering the natural visual adaptation processes and providing an increased sense of immersion in the virtual environment.

Finally, we conclude with a discussion on the scope and contribution of the dissertation in Chapter 8.

Chapter 2

Background and Related Work

We first present an overview of the relevant background and previous work in this chapter. We start by describing the Bidirectional Reflectance Distribution Function (BRDF) and how it is used to represent the surface reflectance property of materials. We review some analytical BRDF models used in computer graphics followed by a discussion on existing BRDF acquisition techniques and various basis representations for BRDFs. We then review some concepts of Monte Carlo integration and importance sampling in the context of computing the direct illumination integral. Finally we review various tone mapping algorithms for display devices and some studies on viewing conditions that are relevant to our work on enhancing perceptual realism.

2.1 The Bidirectional Reflectance Distribution Function

Accurate descriptions of how light reflects off a surface are a fundamental prerequisite for realistic rendering. Real world materials exhibit characteristic surface reflectance, such as glossy or specular highlights, anisotropy, or retro-reflection, which need to be modeled for visual realism. The surface reflectance of a material is formalized by the notion of the Bidirectional Reflectance Distribution Function as described by Nicodemus et. al. [94], which is a 4 dimensional function describing the response of a surface in a certain exitant direction to illumination from a certain incident direction over a hemisphere of directions.

Figure 2.1 depicts the geometry of light interaction at a surface that describes the BRDF. Light arriving at a differential surface dA from an incident direction (θ_i, ϕ_i)

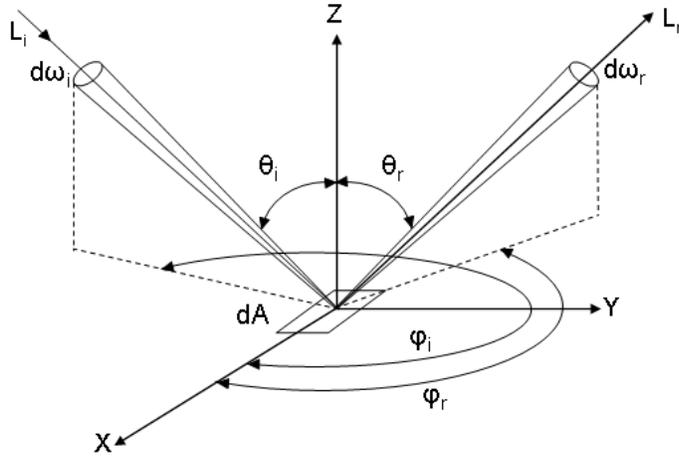


Figure 2.1: The BRDF geometry.

through a solid angle $d\omega_i$ is reflected in the direction (θ_r, ϕ_r) centered within a cone of solid angle $d\omega_r$. The BRDF is mathematically defined as the ratio of the directionally reflected *radiance* to the directionally incident *irradiance*. Please refer to Appendix A for a brief definition of some relevant radiometric terms.

$$f_r(\lambda, \omega_i, \omega_r) = \frac{dL_{\lambda,r}(\lambda, \omega_r)}{dE_{\lambda,i}(\lambda, \omega_i)} \quad (2.1)$$

Here $dE_{\lambda,i}$ is the incident spectral irradiance (i.e., the incident flux of a given wavelength per unit area of the surface) and $dL_{\lambda,r}$ is the reflected spectral radiance (i.e., the reflected flux of a given wavelength per unit area per unit solid angle). Since the definition above includes a division by solid angle, the units of a BRDF are inverse steradian [1/sr]. We can drop the wavelength dependence for notational simplicity. The BRDF then becomes,

$$f_r(\omega_i, \omega_r) = \frac{dL_r(\omega_r)}{dE_i(\omega_i)} \quad (2.2)$$

The BRDF is thus a function of four variables: two variables describing the incoming light direction, and two specifying the reflected light direction. Note that this model of reflectance assumes a homogenous material. For materials with spatially

varying reflectance property, the BRDF is also a function of the position x on the surface and hence a higher dimensional function. BRDFs maintain the laws of Helmholtz reciprocity, i.e.,

$$f_r(\omega_i, \omega_r) = f_r(\omega_r, \omega_i) \quad (2.3)$$

and energy conservation,

$$\int_{\Omega} f_r(\omega_i, \omega_r) \cos\theta_r d\omega_r \leq 1, \forall \omega_i \in \Omega. \quad (2.4)$$

2.1.1 Analytic BRDF Models

Analytical reflection models attempt to describe certain classes of BRDFs using a mathematical representation involving a small number of parameters. These parameters can be either adjusted manually or obtained from fitting to measured BRDF data. These analytical models mostly fall under two categories. (1) Empirical models that are not based on the underlying physics, but provide a class of functions that can be used to approximate reflectance. (2) Physics based models that take into account the physical properties of the material while modeling a specific phenomenon or class of materials.

Empirical Models

One of the earliest models to express specular reflections which is still widely used today in computer graphics is the Phong model [105]. The model is a sum of a diffuse component and a cosine weighted specular lobe. It can be expressed as

$$f_r(\hat{l}, \hat{v}) = \rho_d + \rho_s \frac{(\hat{r} \cdot \hat{v})^s}{(\hat{n} \cdot \hat{l})}, \quad (2.5)$$

where \hat{l} is the normalized vector towards the light source, \hat{v} is the view vector, \hat{r} is the vector obtained by reflecting the light vector about the surface normal \hat{n} and lies in the same plane as \hat{l} and \hat{n} , s is the specular exponent, and ρ_d and ρ_s are the coefficients of diffuse and specular reflectance respectively. The model is based on ad hoc observation of the behavior of reflectance and is neither reciprocal nor energy preserving.

Blinn [10] adopted the model for a more physically accurate reflection by computing the specular component based on the halfway vector \hat{h} :

$$f_r(\hat{l}, \hat{v}) = \frac{\rho_d}{\pi} + \rho_s \frac{(\hat{n} \cdot \hat{h})^s}{(\hat{n} \cdot \hat{l})}, \quad (2.6)$$

Lafortune et al. [71] presented a more elaborate generalized cosine lobe model that can account for off-specular peaks, retro-reflection and anisotropy:

$$f_r(\hat{l}, \hat{v}) = \frac{\rho_d}{\pi} + \sum_i [C_{x,i}(\hat{l}_x \cdot \hat{v}_x) + C_{y,i}(\hat{l}_y \cdot \hat{v}_y) + C_{z,i}(\hat{l}_z \cdot \hat{v}_z)]^{s_i}. \quad (2.7)$$

An advantage of this model is that it is well suited for fitting to measured BRDF data.

Ward [140] proposed a model for anisotropic reflection based on an elliptical Gaussian distribution of normals. It is both energy conserving and reciprocal. It can be expressed as

$$f_r(\hat{l}, \hat{v}) = \frac{\rho_d}{\pi} + \rho_s \frac{1}{\sqrt{\cos \theta_i \cos \theta_r}} \frac{\exp[-\tan^2 \delta (\frac{\cos^2 \phi}{\alpha_x^2} + \frac{\sin^2 \phi}{\alpha_y^2})]}{4\pi \alpha_x \alpha_y}, \quad (2.8)$$

where δ is the angle between the half vector and the normal; ϕ is the azimuth angle of the half vector projected into the surface plane; and α_x, α_y are the standard deviations of the surface slope in the principal \hat{x} and \hat{y} directions, respectively. This does not model Fresnel effects or retroreflection.

Ashikhmin and Shirley [6] have proposed an anisotropic Phong model that observes the laws of reciprocity and energy conservation, includes a Fresnel reflectance term and a non constant diffuse term. The specular part of the model is given by

$$\rho_s(\hat{k}_1, \hat{k}_2) = \frac{\sqrt{(n_u + 1)(n_v + 1)}}{8\pi} \frac{(\hat{n} \cdot \hat{h})^{n_u \cos^2 \phi + n_v \sin^2 \phi}}{(\hat{h} \cdot \hat{k}) \max((\hat{n} \cdot \hat{k}_1), (\hat{n} \cdot \hat{k}_2))} F((\hat{k} \cdot \hat{h})), \quad (2.9)$$

where \hat{k}_1 is the light vector, \hat{k}_2 is the view vector, \hat{h} the halfway vector, F is the Fresnel term, and n_u, n_v are two Phong like exponents that control the shape of the specular lobe in orthogonal directions u and v .

Physically-based Models

There has been a whole body of work, initially in the field applied optics and later in computer graphics, in developing physically accurate reflection models. The illumination model by Torrance and Sparrow [133] is one of the most important physically-based models for the inter-reflection of light at rough surfaces. It derives the specular component by assuming the reflecting surface to be composed of microfacets based on a Gaussian distribution. It includes a Fresnel term for off-specularity and also accounts for shadowing and masking with respect to the microfacet distribution. It is given as

$$f_r(\hat{l}, \hat{v}) = \frac{F \cdot G \cdot D}{\pi \cdot (\hat{n} \cdot \hat{l}) \cdot (\hat{n} \cdot \hat{v})}, \quad (2.10)$$

where F , the Fresnel reflectance term as derived from Snell's laws, is given in the form

$$F = \frac{(g - c)^2}{2(g + c)^2} \left[1 + \frac{(c(g + c) - 1)^2}{(c(g - c) + 1)^2} \right], \quad (2.11)$$

with $c = (\hat{h} \cdot \hat{v})$, and $g^2 = n^2 + c^2 - 1$ where n is the index of refraction. The term D is the distribution term for the microfacets and assumes a Gaussian distribution of the angle between normal and the halfway vector. This can also be interpreted as a distribution of the halfway vector itself. Because the microfacets are perfectly specular, only those with a normal equal to the half vector \hat{h} cause perfect specular reflection from \hat{l} to \hat{v} . Finally, the term G describes the geometrical attenuation caused by self-shadowing and masking of the microfacets. Under the assumption of symmetric, v-shaped grooves, G is given as

$$G = \min \left\{ 1, \frac{2(\hat{n} \cdot \hat{h})(\hat{n} \cdot \hat{v})}{(\hat{h} \cdot \hat{v})}, \frac{2(\hat{n} \cdot \hat{h})(\hat{n} \cdot \hat{l})}{(\hat{h} \cdot \hat{v})} \right\}. \quad (2.12)$$

Several variations of this model have been proposed. Cook and Torrance [20] have extended the model for spectral rendering. Poulin and Fournier [106] presented an anisotropic reflection model assuming a microgeometry of oriented cylindrical grooves. Schlick [115] has proposed a model in which all terms are approximated by simple rational polynomial formulae to improve performance. Oren and Nayar [97] have pro-

posed a microfacet distribution model for diffuse reflection. Ashikhmin et al. [7] proposed an expression for the shadowing and masking for any arbitrary distribution of microfacet normals. He et al. [57, 58] have developed an even more comprehensive model to account for arbitrary polarization of incident light to describe effects like interference.

Distribution-based BRDF Model

Recently, Ashikhmin [5] has proposed a BRDF model that is a generalization of the Ashikhmin-Shirley Phong model in that it allows for arbitrary microfacet distributions in the spirit of [7], while preserving a much simpler mathematical formulation:

$$\rho(\hat{k}_1, \hat{k}_2) = \frac{cp(\hat{h})F((\hat{k} \cdot \hat{h}))}{(\hat{k}_1 \cdot \hat{n}) + (\hat{k}_2 \cdot \hat{n}) - (\hat{k}_1 \cdot \hat{n})(\hat{k}_2 \cdot \hat{n})} \quad (2.13)$$

Here, $p(\hat{h})$ is a particular distribution of half-vectors that describes the material, c is an RGB scaling constant, and $\hat{k}_1, \hat{k}_2, \hat{h}, F$ are as defined in the AS Phong model. The Fresnel term $F(\hat{k} \cdot \hat{h})$ is given by Schlick's polynomial approximation [115]:

$$F((\hat{k} \cdot \hat{h})) = r_0 + (1 - r_0)(1 - (\hat{k} \cdot \hat{h}))^5 \quad (2.14)$$

where r_0 is the reflectance at normal incidence. An advantage of the model is that the distribution $p(\hat{h})$ can be extracted from measured data using a very simple procedure without requiring any numerical fitting. $p(\hat{h})$ can be extracted just from the backscattering measurements as $\hat{k}_1 = \hat{k}_2 = \hat{k} = \hat{h}$ for backscattering geometry. The D-BRDF takes the form

$$\rho(\hat{k}, \hat{k}) = \frac{cr_0p(\hat{h})}{2(\hat{k} \cdot \hat{n}) - (\hat{k} \cdot \hat{n})^2}, \quad (2.15)$$

providing a function that is proportional to the distribution $p(\hat{h})$. We employ this D-BRDF framework to fit measured data in Chapter 3.

2.1.2 BRDF Acquisition

Despite the richness of the various analytical models, they still do not capture the reflectance property of all kinds of day-to-day materials. Also, the parameters required by these models to approximate a certain kind of material are not easily obtainable. Measurement is the most straightforward approach of obtaining BRDF data for a broad class of materials and we review some measurement techniques in this section.

Dense Measurements

As an alternative to analytical BRDF models, one can use dense measurements of BRDFs directly in a rendering system. Such data is available from many sources typically using some version of a gonio-reflectometer (e.g., the Cornell [21] and STARR [96] databases), or by doing dense measurements with a camera (e.g., CURET [24] database). The gonio-reflectometer measurement setup usually involves a photometer and a light source that can both be moved with respect to a surface sample to cover a hemisphere of measurement directions. A dense set of measurement with this setup can take a large amount of time. Also, when using data from such a setup there is generally a need to account for any missing information in directions for which there are no available measurements. For example, the BRDF data in the CURET database represents 205 reflectance measurements uniformly distributed over the hemisphere of 60 different materials. This amounts to a relatively sparsely sampled BRDF due to which there is a need to fit the data to analytical models [26].

Image-based Acquisition

The wide availability and decreasing cost of digital cameras has recently led researchers to explore various image based BRDF acquisition approaches. One way of reducing the number of images that need to be taken is by using curved surfaces for acquiring BRDFs. Marschner et al. [84] constructed an image-based acquisition device for isotropic BRDFs that photographed a curved object with uniform BRDF from many illumination directions to efficiently collect a dense sample set. Lensch et al. [77] presented a clustering procedure to model spatially varying BRDFs and fit each cluster

to a Lafortune model [71]. They then used *principal component analysis* to compute basis BRDFs for material clusters.

Matusik et al. [85] presented a radically different data driven approach to modeling BRDFs. They developed an acquisition device for isotropic BRDFs that sampled the illumination directions much more densely than Marschner and acquired BRDF data for over a 100 representative materials. They then employed linear and non-linear dimensionality reduction techniques to obtain a low-dimensional manifold that characterizes the BRDFs and developed ways for users to navigate over this manifold to generate new BRDFs over the space spanned by their acquired data.

Generally, these methods require knowledge of the geometric shape, and are not well-suited for capturing fabric or sheet materials. Such materials can be measured by wrapping them around a cylinder at various orientations as recently done by Ngan et al. [93]. In this case, they acquire measurements over the hemisphere populating approximately 25% of the sampling bins with data and also fit the acquired data to the anisotropic Phong model [6].

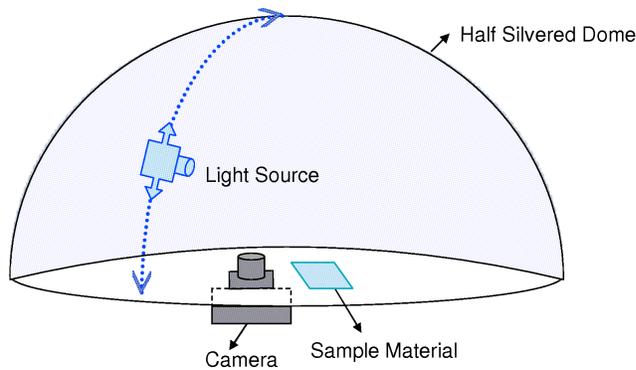


Figure 2.2: The Ward imaging gonioreflectometer.

In many cases, planar samples are, however, more convenient. Other researchers have therefore focused on special optics to cover a large range of incident or exitant light directions for a planar sample in a single photograph. Ward's *imaging gonioreflectometer* [140] as shown in Figure 2.2 was the first image-based measurement

device for planar samples. He used a semi-silvered hemispherical mirror, a CCD camera with a fisheye lens and a movable collimated light source to capture the reflectance data. This setup enables every outgoing direction to be measured with a single image, greatly reducing the acquisition time. The drawback of the design is the difficulty in measuring BRDF values near the grazing angles making it unsuitable for measuring very specular materials. Malzbender et al. [80] use a dome with attached, individually controlled light sources to photograph a surface under varying lighting conditions and recover the per pixel reflectance function which they encoded as coefficients of a polynomial for reconstruction.

Han and Perlin [53] developed a device to capture bidirectional texture functions (BTFs) based on a kaleidoscope. Their setup is similar to ours in the sense of not involving any movement of the camera as well as light source. However, their device measures a sparse set of orientations over the hemisphere and is not suitable for BRDF acquisition. Dana [25] designed an acquisition device using a parabolic mirror that densely covers a relatively small solid angle. The system also involves planer translations of the light source to cover various incident directions and translations of the sample in order to scan the surface for spatial variations in reflectance. Recently, Kuthirummal and Nayar [70] have developed a class of radial imaging systems for image-based acquisition of geometry, texture, and BRDFs. Their BRDF measurement setup can image 4 radial lines of reflectance of a given material for a fixed light source direction. Our design, as discussed in Chapter 3 [45], is most closely related to the last two papers. It can however measure a much larger zone of directions and can be extended to acquire spatially varying BRDFs like Dana.

2.1.3 Basis Representations and Illumination

Independent of the acquisition process, the acquired data is generally not used directly due to problems such as missing measurements and noise in the measurement process. Furthermore, the inherent dimensionality of the BRDF data, and the need to sample it at a high resolution leads to unwieldy storage problems. Researchers have therefore either resorted to fitting the data to analytical models [41, 71, 93, 140] or sampled

the data into some suitable basis function representation. We will review some such representations in this section.

Spherical harmonics have been a popular choice for representing BRDFs in computer graphics [110, 142]. They are the analogue of the Fourier series as defined on a sphere. The BRDF function is approximated by a finite number of terms of the spherical harmonic series as

$$f(\theta, \phi) = \sum_0^{\infty} c_l^m Y_l^m(\theta, \phi) \quad (2.16)$$

where Y_l^m is the spherical harmonic basis function of order l and degree m and c_l^m is the corresponding coefficient. An advantage of this representation is that arbitrary rotations over a sphere are well defined and in the context of BRDFs this is useful for local lighting computations. The major disadvantage of this representation is that for specular BRDFs, the spherical harmonics basis functions suffer from oscillations around the true function value, also known as Gibbs phenomenon. These oscillations are visible in the reconstruction as undesirable ringing artifacts (Figure 3.3).

Wavelets are suitable for representing functions containing high frequency content because they localize in the spatial and the frequency domain. Schröder and Sweldens [117] extended wavelets to spherical domain to efficiently represent spherical functions. They used these spherical wavelets to represent a 2D slice of BRDF by keeping the viewing direction constant. Lalonde and Fournier [72] and more recently Ng et al. [92] have proposed solutions to represent the four-dimensional BRDF with wavelets. The problem with the wavelet representation is that arbitrary rotations are not efficient.

Another approach is to map points on a hemisphere onto a disk. Keondrink et al. [68] take this approach of representing BRDFs using orthonormal basis functions on the unit disk. They project the Zernike polynomial functions onto a hemisphere and use tensor-products of these functions to represent BRDFs. Rotation matrices are not defined for the Zernike polynomial representation.

Gautron et al. [42] recently developed a basis over the hemisphere that is a more natural choice for representing BRDFs than spherical harmonics (SH). However, there

is a need to convert the data to SH for arbitrary rotations. The development of our zonal basis function for BRDF measurement in Chapter 3 can be seen as a further generalization of the hemispherical basis. Our proposed basis transformation method can also be seen as an improvement over their approach of explicitly fitting the SH basis to be zero below the hemisphere.

Somewhat related to our approach, Wenger et al. [141] project basis illumination from a discretized geodesic sphere to acquire the time-varying reflectance fields of an actor’s performance for post-production relighting. Basis illumination in the form of a wavelet noise pattern has also been used for inferring the per-pixel reflectance function of a scene for relighting and environment matting applications [101].

2.2 Monte Carlo Integration and Sampling

In this section, we present an overview of Monte Carlo (MC) integration and sampling as applied to realistic image synthesis. Many treatises exist in the literature that introduce the broad concepts of probability theory and Monte Carlo methods and we point the reader here to the texts by Kalos and Whitlock [62] and Hammersley and Handscomb [52]. For introduction to Monte Carlo methods in ray tracing and image synthesis, we refer the reader to [36, 37, 103, 121].

2.2.1 Monte Carlo Integration

The goal in image synthesis is to compute the reflected radiance for every surface point visible to the viewer. In this thesis, we will restrict ourselves to evaluating the direct illumination integral. In the presence of arbitrary HDR illumination and BRDFs, the direct illumination integral is too computationally intensive for a brute-force solution. Instead, we employ Monte Carlo integration, a stochastic technique that allows us to compute approximations for integrals where an exact solution is intractable.

Suppose we wish to evaluate the integral

$$I(f) = \int_S f(x)p(x)dx, \quad (2.17)$$

where $f(x)$ is a function defined over the domain S and p is the underlying *probability density function* (PDF) of f . One can approximate such an integral with a tractable sum $I_N(f)$. Given a set of identically independently distributed samples $X = \{x_1, x_2, \dots, x_N\}$ drawn from some density $p(x)$ defined over S , the tractable sum

$$I_N(f) = \frac{1}{N} \sum_{i=1}^N f(x_i) \quad (2.18)$$

converges asymptotically as $N \rightarrow \infty$.

The integral $I(f)$ is also referred to as the average or the *expected value* $E(x)$ of a random variable x over the domain S . The expected value of a random variable satisfies the following properties:

$$E(\alpha X) = \alpha E(X); \quad (2.19)$$

$$E\left(\sum_i X_i\right) = \sum_i E(X_i), \quad (2.20)$$

where X is a random variable and α is any constant.

Now, let us examine the *variance* of a random variable. The variance $var(X)$ of a random variable X is the expected value of the squared difference of the realization of a variable and its expected value

$$\begin{aligned} var(X) &= E((X - E(X))^2) \\ &= E(X^2) - E(X)^2. \end{aligned}$$

For independent random variables X_i , variance observes the following properties:

$$var\left(\sum_i X_i\right) = \sum_i var(X_i); \quad (2.21)$$

$$var(\alpha X) = \alpha^2 var(X). \quad (2.22)$$

With these properties in mind, we see that

$$\begin{aligned}
\text{var}(I_N(f)) &= \text{var}\left(\frac{1}{N} \sum_{i=1}^N f(x_i)\right) \\
&= \frac{1}{N^2} \sum_{i=1}^N \text{var}(f(x_i)) \\
&= \frac{1}{N} \text{var}(f(x_i)).
\end{aligned}$$

In Monte Carlo rendering, variance manifests itself as per-pixel noise in the output image. The above equation shows that the variance in a Monte Carlo estimator of an integral $I(f)$ is inversely proportional to the sample size N . Moreover, the error in the estimate behaves similarly to the *standard deviation* of the function, i.e., the square-root of the variance. This highlights the fundamental problem with Monte Carlo integration: the notion of diminishing returns. Because image quality depends on N^2 , one must quadruple N to halve the error.

This brings us to the concept of *importance sampling* which is an important variance reduction technique employed in Monte Carlo integration. Consider an integral $I(f)$ that needs to be evaluated for a function f over the domain $x \in S$:

$$I(f) = \int_S f(x) dx. \quad (2.23)$$

An unbiased Monte Carlo estimate of $I(f)$ can be obtained from the following sum:

$$I(f) \approx \frac{1}{N} \sum_{i=1}^N \frac{f(x_i)}{p(x_i)}. \quad (2.24)$$

where p is a PDF, and $x_i \sim p$. p is referred to as the proposal distribution for the N random variables x_i . In order to have a good estimate of $I(f)$, we need $\frac{f}{p}$ to have low variance. Choosing p appropriately such that f and p have similar shape, and that p is large when f is large is called importance sampling. In the context of computing the direct illumination integral, p has traditionally been chosen according to the distribution of the incident illumination or the distribution of the BRDF.

2.2.2 Sampling of Direct Illumination

All rendering systems, both global and local, must at some point compute the direct illumination in the scene. Unfortunately, this task remains expensive, especially for complex light sources such as environment maps and other image-based representations. Much effort has focused on the development of more efficient techniques for completing this task.

Sampling of Environment Maps

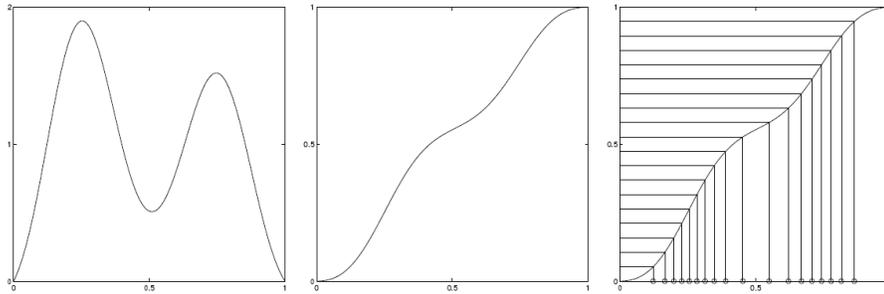


Figure 2.3: Example of sampling via CDF inversion. Left: a 1-D PDF, $p(x)$. Center: the corresponding CDF, $C(x)$. Right: uniform samples along the vertical axis transformed by $C^{-1}(x)$ onto the horizontal axis and re-distributed according to $p(x)$. Image courtesy of [118].

Illumination from environment maps has been a topic of much recent research. Most of this work focuses on interactive applications and therefore uses expensive precomputation [51, 59, 64, 66]. In some recent work, the illumination and/or BRDF are projected into finite bases such as spherical harmonics (e.g., [109, 110, 122]) and wavelets [92].

Other researchers have used importance sampling techniques to distribute samples according to the energy distribution in the environment map, either by using point relaxation schemes [19, 69] based on Lloyd’s clustering algorithm [79] or by using an efficient hierarchical Penrose tiling scheme [98].

Agarwal et al. [2] introduced a sampling method for environment maps taking

into account both the energy distribution in the environment map and the solid angle separating the samples. In this way, close clustering of environment map samples is avoided, which reduces redundant shadow tests.

A common technique for sampling from an arbitrary discrete distribution $p(x)$ is to build the corresponding *cumulative density function* (CDF). For every real number x , the CDF is given by

$$C(x) = p(X \leq x). \quad (2.25)$$

In other words, $C(x)$ is the probability that the random variable X takes on a value less than or equal to x . For $X \in [0, 1]$, the CDF is computed as

$$C(x) = \int_0^x p(x)dx. \quad (2.26)$$

In order to sample from $p(x)$, we can choose a uniform variate $u_i \in [0, 1]$, and then transform it according to the inverse of the CDF to generate $x_i = C^{-1}(u_i)$. This is illustrated in Figure 2.3. Note that since $p(x) \geq 0$, $C(x)$ is a monotonically increasing function and $C^{-1}(x)$ always exists.

Secord et al. [118] described an algorithm for computing and inverting a 2-D CDF based on image intensities in the context of stippling. This is a simple and efficient method, a variant of which we use in our work for drawing samples from environment maps. Inversion of CDFs is also used by Lawrence et al. [75] to sample from environment maps. For details on the CDF inversion method for sampling from environment maps, we refer the reader to Burke's Master's thesis [12].

Sampling of BRDFs

Importance sampling from the BRDF is a common operation particularly for scenes with uniform illumination, such as outdoor scenes. Simple analytical models such as diffuse, Phong and generalized cosine lobe models can be easily sampled analytically. For example, given a Phong BRDF with exponent n , the corresponding PDF of the

specular lobe is given by

$$p(\theta, \phi) = \frac{n+1}{2\pi} \cos^n \theta. \quad (2.27)$$

We can sample from the above PDF by transforming a pair of uniform random variables $(u_1, u_2) \in [0, 1]$ according to

$$(\theta, \phi) = (\arccos((1 - u_1)^{\frac{1}{n+1}}), 2\pi u_2). \quad (2.28)$$

Note that the sampled direction (θ, ϕ) is distributed about the polar axis $+Z$. However, the specular term in the BRDF is parameterized about the reflection direction. Hence, the final step in sampling involves transforming the generated direction from the global polar axis to the local frame of the BRDF. Similarly, more sophisticated analytical models such as the Ward model [140] and the Ashikhmin-Shirley-Phong model [6] also provide means for analytical inversion of the specular lobe.

Microfacet distribution models based on Gaussian distributions [7, 133] provide an analytical form of the PDF $p(\hat{h})$ that can be inverted, while an arbitrary distribution such as that of the D-BRDF model needs to be inverted using a 2-D CDF inversion technique similar to that employed for environment maps. The viewing vector \hat{k}_2 then needs to be reflected about the generated half vector \hat{h} to obtain the sampling direction \hat{k}_1 :

$$\hat{k}_1 = 2(\hat{h} \cdot \hat{k}_2)\hat{h} - \hat{k}_2 \quad (2.29)$$

The final step involves converting the PDF over the half-vector $p(\hat{h})$ into a PDF over the viewing direction $p(\hat{k}_2)$ for importance sampling:

$$p(\hat{k}_2) = \frac{p(\hat{h})}{4(\hat{h} \cdot \hat{k}_1)} \quad (2.30)$$

For tabulated BRDFs, kd-tree representations [86] and more recently, factored representations [74] have been used for efficient importance sampling, while cosine lobe approximations have been used for importance sampling of procedural shaders [124].

In our work as discussed in Chapter 4, we use only Phong and diffuse reflection models. However, our method could easily be extended to incorporate more sophisticated materials using any of the above methods.

Multiple Sampling Approaches

Veach & Guibas [136] proposed multiple importance sampling (MIS) as an effective variance reduction technique. They combined different sampling distributions such as illumination and BRDF distributions in an optimal manner using their proposed *balance* heuristics. However, their method results in a mixing of the variance of the individual distributions while methods that directly sample from the product distribution of the illumination and the BRDF generally reduce variance further than MIS [13, 16].

Szecsni et al. [128] sample the unoccluded illumination using correlated sampling and the difference due to visibility using importance sampling. This method generally performs well in fully visible regions, but rather poorly in occluded or partially occluded regions, since the sampling of visibility does not follow a special sampling pattern. Our work, by contrast, focuses visibility tests according to the direct illumination integral.

Sampling from Product Distributions

We introduce the notion of *bidirectional importance sampling* [13] in Section 4.2 that takes into account the energy of incident illumination as well as the BRDF in the sampling process. We present two Monte Carlo algorithms for sampling from the product distribution - one based on rejection sampling and the other based on sampling-importance resampling (SIR). The SIR algorithm is also used by Talbot et al. [129] for variance reduction of direct illumination estimate.

Clarberg et al. [16] present a technique for efficiently sampling the product of the illumination and the BRDF using a hierarchical wavelet representation. Their method is very efficient for tabulated BRDFs but requires significant precomputation for environment maps. Lawrence et al. [75] present an approach for compressing cumulative distribution functions for efficient inversion and they apply it to sampling from many

precomputed environment map PDFs for different surface orientations, which is a step towards approximating the product distribution. Recently, Cline et al. [17] presented an efficient technique for generating samples according to the product distribution based on recursive splitting of the environment map at the BRDF peaks. They employ a warping step to distribute samples into the split regions similar to Clarberg et al., but without requiring the expensive precomputation step for a wavelet decomposition.

We extend product distribution sampling in the temporal domain in Chapter 5 with our proposed *sequential Monte Carlo* (SMC) [46] mechanism for sampling dynamic illumination. Our solution is very general and can be applied in combination with any sampling scheme discussed above for proposing samples in the first time step. In fact, sequential sampling would also help overcome the precomputation requirements of some of these techniques for a dynamic sequence.

Metropolis Sampling for Global Illumination

Veach & Guibas [137] first applied Metropolis sampling to the problem of image synthesis and developed a general, robust and unbiased algorithm called Metropolis Light Transport (MLT) that was well suited for hard cases for sampling because of its localized exploration and path re-usage properties. Fan et al. [38] recently applied the Metropolis algorithm for efficiently sampling coherent light paths for photon mapping.

Cline et al. [18] presented an efficient unbiased method to solve correlated integral problems with a hybrid algorithm that uses Metropolis sampling-like mutation strategies in a standard Monte Carlo integration setting, overcoming the startup bias problem of MLT. They apply energy redistribution over the image plane to reduce variance of path tracing for global illumination. Our work as discussed in Section 4.3 is similar in spirit in the sense of initial Monte Carlo sampling followed by Metropolis sampling, except that we apply this to direct illumination with a specific focus on efficient exploration of visibility in partially occluded regions [47]. Similarly, our proposed SMC sampling mechanism for dynamic illumination also includes a Metropolis sampling step for exploration of the new direct illumination target distribution.

Sampling from Dynamic Illumination

Several researchers have considered the problem of sampling from dynamic illumination. For example, Sbert et al. [113] presented a method for reusing light paths computed in one frame of a light animation sequence in all other frames using multiple importance sampling. The indirect illumination in each frame of the sequence is approximated as weighted contributions from these precomputed virtual point lights. This method works on moving point lights, not complex environments.

Some researchers have looked at path re-usage in the context of global illumination for a sequence of camera animation via reprojection of the primary ray hits on to the image plane. Here, techniques that are possibly biased [54] as well as unbiased techniques [87] using multiple importance sampling have been proposed for the path re-usage. Our proposed SMC sampling mechanism can also be applied to a camera animation sequence for sample propagation in an unbiased manner, while being more efficient than multiple importance sampling.

Recent work on dynamic environment maps, including Wan et al. [139] and Havran et al. [55], approximates environments with a set of point lights that are drawn according to the energy distribution in the environment map, and evolve smoothly over time. However, this procedure introduces a systematic error for specular materials if none of the chosen point lights resides inside the specular lobe. In this case, one would expect to see a specular reflection of a dimmer part of the environment, but these methods cannot produce this result.

2.3 Display Algorithms and Viewing Conditions

Finally, in this section we review some work on tone-mapping algorithms for display devices and studies on viewing conditions and visual adaptation that are relevant to the enhancement of perceptual realism of a virtual scene. Creation of a sense of presence and immersion in a virtual environment is the major focus of perceptually realistic rendering. Research targeting perceptual realism has so far had to deal with two major factors. First is the dynamic range of display technology and its capacity for represent-

ing the full range of dark and light intensities found in the real world. And secondly, accounting for the viewing conditions for displayed images.

2.3.1 Dynamic Range of Display Devices

The limited dynamic range of both imaging devices and displays has received a lot of attention in recent years. Algorithms have been developed for capturing both photographs [29, 81, 89, 112] and videos [63] with extended dynamic range. The challenge has been to faithfully reproduce the range of intensities in the scene on the display end. Tone mapping operators alleviate the problem of limited dynamic range of conventional display devices to an extent, but are unable to compensate fully for these shortcomings [76].

For the most part, the dynamic range problem has been addressed recently by new high dynamic range (HDR) display technology [119]. We discuss how to drive the projector-based HDR displays in real-time in Chapter 6 while also briefly describing the algorithm for the LED-based design.

2.3.2 Tone Mapping and Visual Adaptation

The class of image processing techniques for coping with the discrepancy between real world luminances and those that fit within the limited gamut of a conventional output device is collectively called *tone-mapping*. A significant body of research in realistic rendering has focused on tone mapping operators for displaying a wide range of intensities on conventional displays. Much of this work is intended primarily for still images (e.g., [35, 73, 111, 116, 135] and others).

Ferwerda et al. [39] pioneered work on tone mapping operators that *explicitly* take visual adaptation into account. Their work uses threshold vs. intensity functions to map threshold contrasts from the original intensity range to that of the display. Patanaik et al. [99] developed a model that also incorporates supra-threshold brightness, color, and visual acuity. Both models could, in principle, be used to track changes in visual adaptation over time, although they are computationally rather expensive. Other researchers [4, 34, 60, 100, 114] have developed methods to model the time-dependent

state of adaptation based on the recent history of viewing conditions in the virtual world.

One limitation of these methods is that they have no information of the user's *actual* state of adaptation, since it depends largely on the real-world illumination in the room, rather than the virtual world. The viewing conditions are largely unknown, meaning that parameters such as the viewer's light and color adaptation cannot be considered in the image generation process. An image generated for a dark-adapted viewer will not be perceived as realistic in a bright room. The use of HDR displays largely removes the need for tone mapping operators, since they can reproduce intensities from mesopic to medium photopic vision levels. However, this does not solve the problem of unknown viewing conditions.

2.3.3 Perceptual Studies on Room Lighting

A vast body of literature in the perceptual psychology community deals with the impact of room lighting conditions. Many of these studies were performed to analyze the impact of room lighting on ergonomic factors such as screen visibility, eye strain, and so forth. There has also been work on using light sensors to adjust the display brightness and contrast (e.g. [3, 8]). These studies try to either minimize the influence of room lighting on displays or compensate at the display end for the illumination conditions. We, on the other hand, deliberately focus on tying the room lighting into the viewing environment to enhance a viewer's sense of immersion [49].

Other studies analyze the perceived brightness vs. luminance levels of images viewed under different room illumination (e.g. [9, 32]). Most existing studies only cover static illumination. An exception is [108], which discusses the need for adjusting for dynamic lighting changes in critical applications such as reading controls for cockpits. However, all studies we are aware of implicitly assume that room lighting is the primary factor for adaptation, and that it has a significant impact on the display surface itself. Both of these assumptions do not hold for HDR displays, and as a consequence the findings from these studies do not apply to our work.

2.3.4 Bridging Real and Virtual Illumination

One way of integrating real and virtual objects is to change the illumination in the virtual world to make it consistent with the real world. Nayar et al. [91] recently developed a lighting sensitive display that tracks changes of illumination in the room, including directional changes, and uses this information to shade virtual objects. Our work, as discussed in Chapter 7, takes the opposite approach and adapts the room illumination to be consistent with a virtual world. We believe there is room for both approaches: while Nayar et al.'s method creates opportunities for new user interaction metaphors, ours is useful whenever a specific virtual world needs to be displayed.

There has been other work that uses separate light sources to augment computer displays. Philips produces a series of high-end flatpanel TVs which have light sources on the back to illuminate the wall behind the TV [104]. The lights are driven uniformly based on the average intensity of the screen content, thereby essentially reducing contrast between the wall and the TV screen. In contrast, our system creates directional illumination based on actual information from the virtual environment. Light sources have also been used to augment displays in theme park rides. There, lights are often used together with other physical props to show a fixed scene. We, on the other hand can deal with dynamically generated content, but aim only at creating low resolution information for the peripheral field of view.

Other related work includes Debevec et al.'s Light Stage 3 [30]. That work uses a number of computer-controlled lights to illuminate actors or objects such that they appear on camera as if they were in a certain real-world environment. An actor in Light Stage 3 sees only a number of point lights, while we aim at producing a smooth environmental illumination that can convincingly represent the real environment in the user's peripheral view. This goal requires a different physical setup, as well as different calibration and rendering algorithms.

Also related to our work are fully immersive, CAVE-like environments [23]. Unfortunately, CAVEs, like other VR displays, have very limited contrast which makes them unsuitable for representing the kind of adaptation processes we are interested in. Due to engineering constraints such as power consumption and heat production [119],

its seems unlikely that immersive HDR environments will be feasible in the near or even medium term future. Space and cost present further obstacles. Such a system would also require high resolution, omnidirectional illumination information, which is hard to generate, for example, in live action film. Therefore, in Chapter 7, we focus on conventional, limited field-of-view displays with high contrast which we augment with low-resolution directional illumination.

Chapter 3

BRDF Acquisition with Basis

Illumination

Realistic image synthesis requires both complex and realistic representation of surface reflectance as described by the bidirectional reflectance distribution function (BRDF). Real world materials exhibit characteristic surface reflectance, such as glossy or specular highlights, anisotropy, or retro-reflection, which need to be modeled for visual realism. Numerous analytical models of BRDFs have been used in computer graphics that observe the laws of energy conservation and reciprocity, and hence are physically plausible. [6, 7, 20, 57, 58]. However, these models generally do not capture the reflectance properties of all kinds of materials. Furthermore, selecting appropriate model parameters for representing different kinds of real-world materials can be a non-intuitive and time-consuming process. Therefore, acquisition of various real world BRDF data has been a very active area of research over the last few years. This task has typically involved measuring the response of various samples using some version of a gonio-reflectometer [21, 24, 96, 140]. More recently, several researchers have employed image based techniques in order to make acquisition of BRDFs more efficient [25, 53, 77, 84, 85].

Independent of the acquisition process, the acquired data is generally not used directly due to its large size, the noise present in the measurement process, and missing data for certain incident and exitant directions. The data is either fitted to an analytical model for rendering [71, 93, 140] or sampled into some suitable basis function [92, 110, 142]. This fitting process results in the loss of some of the captured high frequency details in the original data, possibly making the high sampling density of

acquisition an overkill. At the same time, reducing the sampling density during acquisition would result in aliasing artifacts for sharp features that would then fall below the Nyquist limit.

In this chapter, we propose an alternative approach to the acquisition of reflectance data where we *optically* project the data into a suitable basis function directly *during* the capture process [45]. This approach results in optical low-pass filtering of the data at capture time, and thus addresses aliasing issues and minimizes high frequency noise. An added benefit is that this prevents any redundancy in data capture as we can use all of the data we acquire. For this purpose, we design a compact optical setup that can project illumination as continuous basis functions over a spherical zone of directions. Likewise, the light reflected off the sample is measured over a continuous zone of directions. BRDF values outside the zonal directions are extrapolated by re-projecting the zonal measurements into a spherical harmonics basis, or by fitting analytical reflection models to the data. Our approach speeds up acquisition time to one or two *minutes* compared to a few hours required by previous acquisition approaches.

Here, we outline the main contributions of our approach:

- The theory behind, and a practical implementation of the concept of measuring the response of a surface to a basis function as a way of optically filtering and encoding the BRDF data.
- Development of a set of orthogonal basis functions defined over the measurement space, as well as basis transformation as a way of data extrapolation.
- A novel design for a curved reflector catadioptric imaging setup resulting in an efficient image based BRDF acquisition without involving any moving parts.

3.1 Overview

The distinguishing characteristic of our BRDF measurement system is that it captures the response of the surface to illumination in the form of smooth basis functions, while existing methods measure impulse response using thin pencils of light that approximate

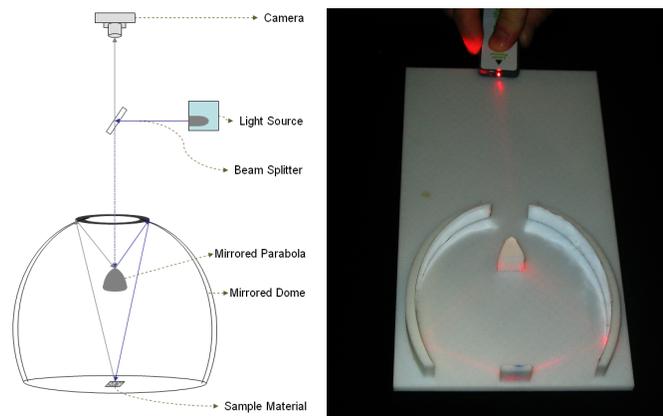


Figure 3.1: Left: Physical setup of our reflectance acquisition device. A camera focused on the mirrored components views a zone of reflected directions. A projector illuminates the corresponding zone of incident directions using a beam splitter. Right: A prototype demonstrating the concept in 2-D. Here, we focus illumination on the mirrored components using a laser pointer and observe that the beam bounces back to its origin.

Dirac peaks. For this concept to be practical, we require an optical setup that allows us to simultaneously project light onto the sample from a large range of directions, and likewise to measure the reflected light distribution over a similarly large range of directions. Developing such optics also has the advantage that no moving parts are required, which is one reason for the speed of our acquisition.

In this work, we choose a spherical zone of directions as the acquisition region for both incident and exitant light directions. Spherical zones have several advantages over regions of other shape. First, they allow us to develop basis functions that align nicely with the symmetries present in many BRDFs, thus minimizing the number of basis functions required to represent a given BRDF. Alignment also simplifies extrapolation of data into missing regions. Second, a zonal setup allows us to design optics that could, in principle, cover over 98% of the hemisphere, with only a small hole near the zenith, where BRDF values are usually smoother compared to more tangential directions. Our optical design for covering this zone, in principle, enables high sampling density close to the tangential directions which are hard to measure with conventional setups. The

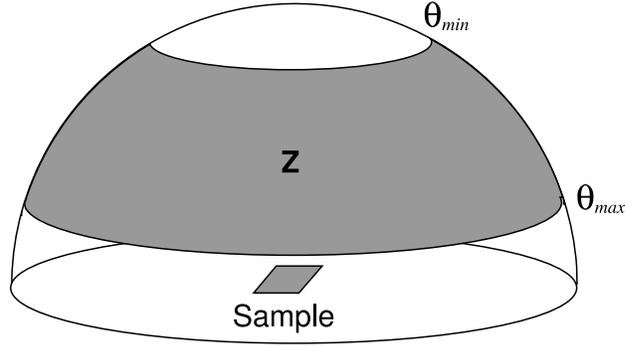


Figure 3.2: The measurement zone Z .

manufacturing process that we used for our optical components allowed us to produce a section of that range corresponding to 51% of the hemisphere.

Figure 3.1 shows a diagram and a 2D mockup of such an optical setup. A camera focused on the mirrored components can capture the full zone of reflected directions in our setup. Simultaneously, a projector focused on the the mirrored components can cover the corresponding zone of incident directions.

In the following, we will first discuss the theoretical underpinnings for basis function BRDF acquisition (Section 3.2), and then describe the physical setup (Section 3.3). Finally, we present some results in Section 3.4 and conclude with a discussion in Section 3.5.

3.2 Measurement with Basis Functions

In this section, we discuss the mathematical concepts behind a basis function approach for BRDF measurement, and derive the specific basis that we use in our work. Section 3.3 then deals with the physical realization of these concepts.

Assume that we want to measure a BRDF $f_r(\omega_i, \omega_r)$ for combinations of incident light direction ω_i and exitant light direction ω_r restricted to a spherical zone Z centered around the surface normal. Z corresponds to longitudinal angles $\phi \in [0 \dots 2\pi]$ and latitudinal angles $\theta \in [\theta_{min} \dots \theta_{max}]$, as shown in Figure 3.2.

We would like to approximate the BRDF over this zone with a linear combination of basis functions $\{Z_k(\omega_i)\}$ over the incident light directions. We will include the $\cos \theta_i$ term in this basis representation for convenience and numerical stability, i.e.,

$$\hat{f}_r(\omega_i, \omega_r) = f_r(\omega_i, \omega_r) \cos \theta_i \approx \sum_k Z_k(\omega_i) z_k(\omega_r), \quad (3.1)$$

so that we can write the reflected radiance for any outgoing direction ω_r as

$$L_r(\omega_r) = \int_{\mathbf{Z}} f_r(\omega_i, \omega_r) L_i(\omega_i) \cos \theta_i d\omega_i \quad (3.2)$$

$$= \int_{\mathbf{Z}} \left(\sum_k Z_k(\omega_i) z_k(\omega_r) \right) L_i(\omega_i) d\omega_i \quad (3.3)$$

$$= \sum_k z_k(\omega_r) \int_{\mathbf{Z}} Z_k(\omega_i) L_i(\omega_i) d\omega_i. \quad (3.4)$$

In this framework, BRDF measurement can be seen as the process of determining the coefficients $z_k(\omega_r)$ for each basis Z_k and each exitant light direction ω_r . If we have chosen the Z_k such that they form an orthonormal basis *over the zone* \mathbf{Z} , then the coefficients are given as

$$z_k(\omega_r) = \int_{\mathbf{Z}} Z_k(\omega_i) f_r(\omega_i, \omega_r) \cos \theta_i d\omega_i. \quad (3.5)$$

In other words, we can measure $z_k(\omega_r)$ by recording the reflected light along each direction $\omega_r \in \mathbf{Z}$ for different incident illumination patterns $Z_k(\omega_i)$.

For practical applications, we of course need to extrapolate from the data measured over the zone to incident and exitant directions that have not been measured. In general, we would also like to transform the data into a different representation that is more convenient for rendering purposes, such as a tensor-product Spherical Harmonics (SH) basis, or coefficients of an analytical reflection model. Interestingly, format conversion and extrapolation can be achieved in a single, inexpensive step, as described below. First, however, we introduce the orthonormal zonal basis that we have developed.

3.2.1 Orthonormal Zonal Basis

Like the Spherical Harmonic basis, our Zonal Basis (ZB) is derived from the Associated Legendre Polynomials (ALP) $P_l^m(x)$, $m \in \{0, \dots, l\}$, which are orthogonal over

$x \in [-1, 1]$ with

$$\int_{-1}^1 P_l^m(x) P_l^m(x) dx = \frac{2(l+|m|)!}{(2l+1)(l-|m|)!} \delta_{ll'}. \quad (3.6)$$

For defining spherical harmonics Y_l^m , the P_l^m are scaled so that they are orthogonal over $[0, \pi]$, with

$$Y_l^m(\theta, \phi) = \begin{cases} \sqrt{2} K_l^m \cos(m\phi) P_l^m(\cos \theta) & \text{if } m > 0 \\ \sqrt{2} K_l^m \sin(-m\phi) P_l^{-m}(\cos \theta) & \text{if } m < 0 \\ K_l^0 P_l^0(\cos \theta) & \text{if } m = 0 \end{cases}, \quad (3.7)$$

where K_l^m is the SH normalization constant:

$$K_l^m = \sqrt{\frac{(2l+1)(l-|m|)!}{4\pi(l+|m|)!}}. \quad (3.8)$$

For our zonal basis, we follow the same principle, and rescale the ALP to the range $[\theta_{min} \dots \theta_{max}]$.

$$\hat{P}_l^m(x) = P_l^m(n_1 \cdot x - n_2), \quad (3.9)$$

with

$$n_1 = \frac{2}{\cos \theta_{min} - \cos \theta_{max}}$$

$$n_2 = \frac{2 \cos \theta_{min}}{\cos \theta_{min} - \cos \theta_{max}} - 1.$$

The ZB functions $Z_l^m(\phi, \theta) \in [0, 2\pi] \times [\theta_{min}, \theta_{max}]$ then become

$$Z_l^m(\theta, \phi) = \begin{cases} \sqrt{2} \hat{K}_l^m \cos(m\phi) \hat{P}_l^m(\cos \theta) & \text{if } m > 0 \\ \sqrt{2} \hat{K}_l^m \sin(-m\phi) \hat{P}_l^{-m}(\cos \theta) & \text{if } m < 0 \\ \hat{K}_l^0 \hat{P}_l^0(\cos \theta) & \text{if } m = 0 \end{cases}, \quad (3.10)$$

where the zonal normalization constant \hat{K}_l^m can be shown to be

$$\hat{K}_l^m = \sqrt{\frac{(2l+1)(l-|m|)!}{2\pi \cdot (\cos \theta_{min} - \cos \theta_{max}) \cdot (l+|m|)!}}. \quad (3.11)$$

3.2.2 Basis Conversion to Spherical Harmonics

While the zonal basis is a convenient basis for measurement due to its orthonormality over the zone, it is not convenient for representing the full BRDF, since it does not cover the full hemisphere of incident and exitant directions. We therefore need to extrapolate the missing data, and transform the measurements into a basis that is convenient for rendering. The spherical harmonics are one suitable basis, which has been used extensively in the past for representing BRDF data.

Converting ZB coefficients to SH coefficients requires us to project the continuous function \hat{f}_r over the zone \mathbf{Z} into the basis Y_l^m , i.e., $\hat{f}_r(\omega_i, \omega_r) \approx \sum_{l,m} y_l^m(\omega_r) Y_l^m(\omega_i)$. Unlike the ZB basis functions, however, the restrictions of the SH basis functions to \mathbf{Z} are *not* orthonormal, and therefore, the equivalent of Equation 3.5 does *not* hold for spherical harmonics. Instead, we have

$$y_l^m(\omega_r) = \int_{\mathbf{Z}} \hat{Y}_l^m(\omega_i) \hat{f}_r(\omega_i, \omega_r) d\omega_i, \quad (3.12)$$

where $\{\hat{Y}_l^m(\omega_i)\}$ is the *dual basis* to the spherical harmonics *over the zone \mathbf{Z}* , i.e. the basis that fulfills the conditions

$$\int_{\mathbf{Z}} Y_l^m(\omega) \hat{Y}_p^q(\omega) d\omega = \begin{cases} 1 & \text{if } l = p \text{ and } m = q \\ 0 & \text{otherwise} \end{cases}. \quad (3.13)$$

Since $\{\hat{Y}_l^m(\omega_i)\}$ is a basis for the same function space as the SH basis, we also have

$$\hat{Y}_p^q = \sum_{l,m} c_{l,p}^{m,q} Y_l^m. \quad (3.14)$$

Equations 3.13 and 3.14 together describe a sparse linear system $Ax = b$ where x is the vector of $l, m \times p, q$ unknown linear weights $c_{l,p}^{m,q}$ that define the duals \hat{Y}_p^q and b is the vector of $lm \times pq$ constraints defined by the RHS in Equation 3.13. The matrix A is large with $N^2 \times N^2$ elements where N is the number of SH bases being used to represent a function. However it is very sparse with only $O(N^2)$ non-zero values occurring for combinations of ZB and SH functions with matching degree.

Once we have the duals \hat{Y}_p^q in terms of the linear weights $c_{l,p}^{m,q}$, we can project any function defined over Z into the SH basis. In particular, we have

$$\hat{f}_r = \sum_{l,m} y_l^m Y_l^m, \quad \text{with } y_l^m = \int_{\mathbf{Z}} \hat{f}_r \hat{Y}_l^m d\omega_i. \quad (3.15)$$

Given this SH representation of the BRDF, we can now use the coefficients y_l^m both inside the zone \mathbf{Z} , and in the region outside the measurement area, which extrapolates the measurements to the full hemisphere. The exitant directions ω_r can be projected into a spherical harmonic basis in a similar fashion, which results in an extrapolation of missing data on the outgoing hemisphere.

Moreover, we can also project the individual functions of the ZB into the SH basis in a precomputation step. Conversion from ZB to SH is then a simple linear transformation of the zonal coefficients z_p^q of a function \hat{f}_r by a basis change matrix C into corresponding SH coefficients y_l^m . Each element of this matrix is defined by

$$C_{l,p}^{m,q} = \int_{\mathbf{Z}} z_p^q \hat{Y}_l^m d\omega. \quad (3.16)$$

The basis conversion matrix C ($N \times N$) is quite sparse, since it only has non-zero coefficients when the degrees of the ZB and the SH function match.

3.2.3 Fitting Analytical Reflection Models

For relatively low frequency BRDFs, the spherical harmonic representation produces very good results. For specular materials, it is well known that basis functions such as spherical harmonics suffer from Gibbs oscillations that are visible in the reconstruction as undesirable ringing artifacts. The zonal basis functions developed here also exhibit the same oscillations for high frequency data. Hence, for specular materials, we cannot directly use the acquired coefficients or transform them into SH for rendering.

Instead, we propose to fit the higher order zonal representation of specular BRDFs to an analytical model, thereby overcoming the problem of oscillations in the reconstruction. In our case, we chose the distribution-based BRDF model proposed by Ashikhmin [5], due to the simplicity of the fitting procedure. The D-BRDF model is also a generalization of the Ashikhmin-Shirley-Phong model [6], which was recently

found to be particularly well-suited for fitting to measured data [93]. However, the measured zonal data can be fitted to any other suitable analytic model using a numerical fitting procedure such as Levenberg-Marquardt [107].

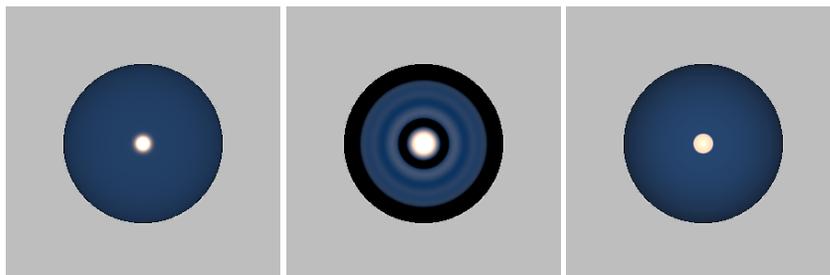


Figure 3.3: An illustration of the suppression of ringing through fitting analytical BRDF models. Left: original acrylic blue paint BRDF. Center: 10^{th} order zonal reconstruction, rendered after transformation into SH, exhibiting severe ringing artifacts. Right: Corresponding D-BRDF fit to the zonal reconstruction.

Figure 3.3 shows a synthetic example of the D-BRDF fitting. The ground truth data as shown on the left is the acrylic blue isotropic BRDF acquired by Matusik et al. [85]. The figure in the center is a reconstruction of this BRDF after projection into a higher order zonal basis function, followed by a transformation into the spherical harmonics basis for rendering. The black rings represent negative values due to ringing. The figure on the right corresponds to fitting the zonal reconstruction directly into the D-BRDF analytical model. Since Gibbs phenomenon results in oscillations around the true value of the function, the least-squares fit of a BRDF model that does not exhibit this kind of oscillation is very effective at approximating the original shape.

3.3 Measurement Setup and Calibration

The primary components of our image-based acquisition setup are a convex parabolic mirror suspended inside a mirrored dome. This optical setup can cover a zone of incident as well as exitant directions of measurement. In addition to the mirrored components, the acquisition system consists of a FireWire machine vision camera (Prosil-

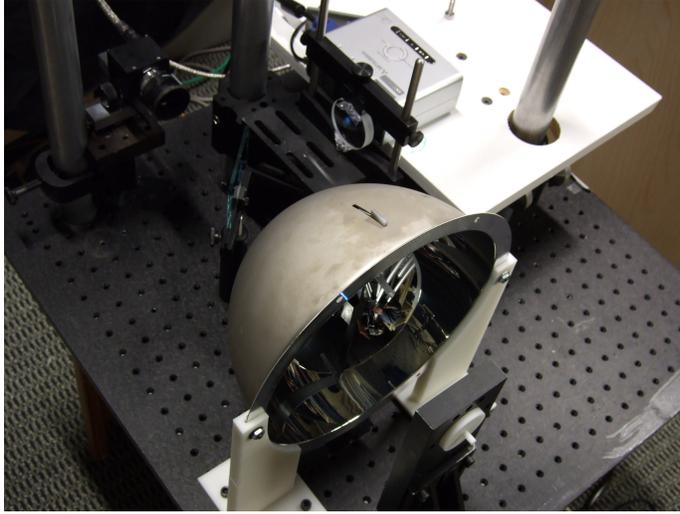


Figure 3.4: Photograph of the proposed BRDF acquisition setup including a camera, a projector, a beam-splitter, and two curved reflectors mounted on a $40\text{ cm} \times 40\text{ cm}$ optical bench.

ica EC 1350C), an LED RGB PocketProjector (Mitsubishi PK1), and a beam splitter (Figure 3.4). The camera has a resolution of 1360×1024 and an acquisition rate of 15 frames per second at 12-bits per color channel. The projector has a resolution of 800×600 with peak illumination intensity specified at 200 Lux. An external 350 mm lens was used to focus the projector at the required focal distance. All reflectance measurements are performed with multiple exposures [29] for high dynamic range (HDR) acquisition.

Our optical setup consists of two mirrored components, a convex parabola and a concave reflective dome as shown in Figures 3.1 and 3.5. The dome has a rotationally symmetric shape with a freeform profile, as detailed in the following.

Dome Shape: For a fixed configuration of parabola, sample, camera, and projector, the freeform profile of the dome is determined as follows. First, the location of the dome’s rim D_1 is found by intersecting a camera ray reflecting off the bottom edge P_1 of the paraboloid with the tangent plane of the sample (Figure 3.5, left). The surface normal at the rim defines a tangent plane in D_1 . For the next camera ray reflecting

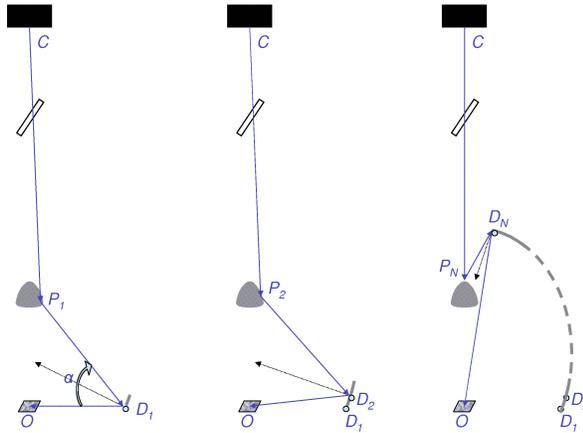


Figure 3.5: Iterative process for designing the profile of the reflective dome for a fixed convex parabolic reflector.

of P_2 , we compute the intersection D_2 of the reflected ray with the tangent plane of D_1 (Figure 3.5, center). The normal in D_2 defines a new tangent plane that we can use in the same way to obtain the next point on the dome. Proceeding iteratively with this approach, we can determine the full shape of the dome (Figure 3.5, right) in what amounts to an Euler integration procedure. Note that these simulations are run at orders of magnitude higher resolution than actual camera or projector pixel resolution.

Design Simulations: The design parameters, i.e., the spatial location of parabola, sample, camera, and projector, were optimized using detailed simulations with a ray-tracer. We modeled the camera and projector as thin lens devices. Our simulations took into account various parameters such as focal distances, finite apertures and pixel resolutions of cameras and projectors, and stability under minor misalignments of the various optical components to the optical axis.

Final design: After extensive simulations, we decided on a design that lets us project over 100 pixels between the vertex and the tangent of the parabolic mirror in order to provide at least 1 measurement per degree along the latitudinal directions. For this setup, the distance between the center of projection of the camera and the vertex of the parabolic mirror is 27 cm, and the distance between the parabola vertex and the sample at the bottom is 13.5 cm.

The dimensions of the full dome are $11'' \times 11'' \times 10''$ for this setup. Our design provides us > 1 pixel/degree measurements over the full measurement zone. The full dome as simulated in Figure 3.5 would cover the zone from 9° to 90° off the normal to the sample. This range corresponds to over 98% of the full hemisphere. Zonal basis function plots defined over this *maximal* zone are given in Appendix B.

Physical implementation: For the manufacturing of the dome and parabola, we chose electroforming process, in which a mandrel of the dome is first machined and polished, and then the actual dome is deposited on this mandrel in an electrolyte bath. This process allows the production of optical quality free-form surfaces at moderate cost. Electroforming also helps with the *reproducibility* of the design from the machined mandrel. However, a downside of this approach is that it only allows for convex holes, since the mandrel has to be removed after the electroforming process. For this reason, we were only able to build a dome covering the zone from 9° to 57° off normal, corresponding to about 51% of the hemisphere (Figure 3.4).

3.3.1 Calibration

Geometric calibration is necessary in order to align the camera and the projector to the optical axis of the acquisition setup. We also need to perform photometric calibration in order to recover the absolute scaling factors for our measurements with respect to some known reflectance standard.

Optical Axis Calibration: The optical axis of the camera and projector need to be aligned with that of the parabolic mirror and dome. We mount the dome on an optical table, and mark its optical axis with crosses that are attached to the dome with precision mounts. The camera is moved with a manual translation stage until all crosses line up. Likewise, the projector is moved manually until the shadows of all crosses line up. For details of the design simulation and optical axis calibration process, we refer the reader to Achutha's Master's thesis [1].

Sample Mounting: Due to the large aperture of our optical system, the depth-of-field is very shallow, about 2 *mm*. As a result, the material samples have to be mounted with fairly high precision, which is easily achieved with a mechanical stop.

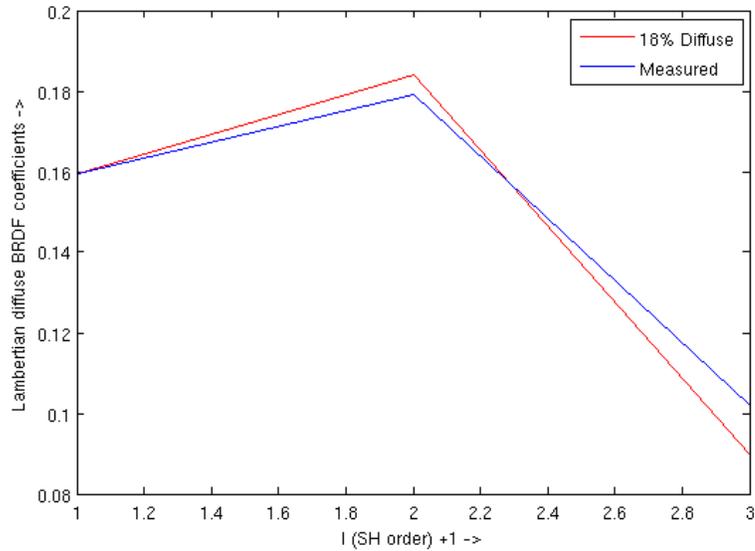


Figure 3.6: Plot of measured coefficients of gray card vs. that of pure Lambertian diffuse reflectance. Note that the measured coefficients have been scaled so that the DC term ($l = 0, m = 0$) matches the DC term of 18% Lambertian diffuse reflectance.

Projector Flat-Fielding: We account for any spatial variation of the projector illumination by acquiring an HDR photograph of a full screen image set to medium gray, projected on to a diffuse white screen at the required focal distance of 28cm. All the basis images are then modulated by this image.

Reflectance Calibration: An important aspect of the calibration is to recover the relative scaling factors for our measurements with respect to some known reflectance standard. For this, we take advantage of an 18% diffuse gray card commonly used in photography. We measure the diffuse reflectance of the gray card with our setup using the first 9 zonal basis functions [109]. The relative scaling factors for each color channel are obtained by white-balancing the results of the gray card measurements. As an additional step, we plot the first 3 ($l \leq 2, m = 0$) measured coefficients of the gray card against coefficients of 18% Lambertian diffuse reflectance for validation of the accuracy of the measurement system. As seen in Figure 3.6, once we scale the measured DC term to match that of 18% gray, the other two terms are off by around

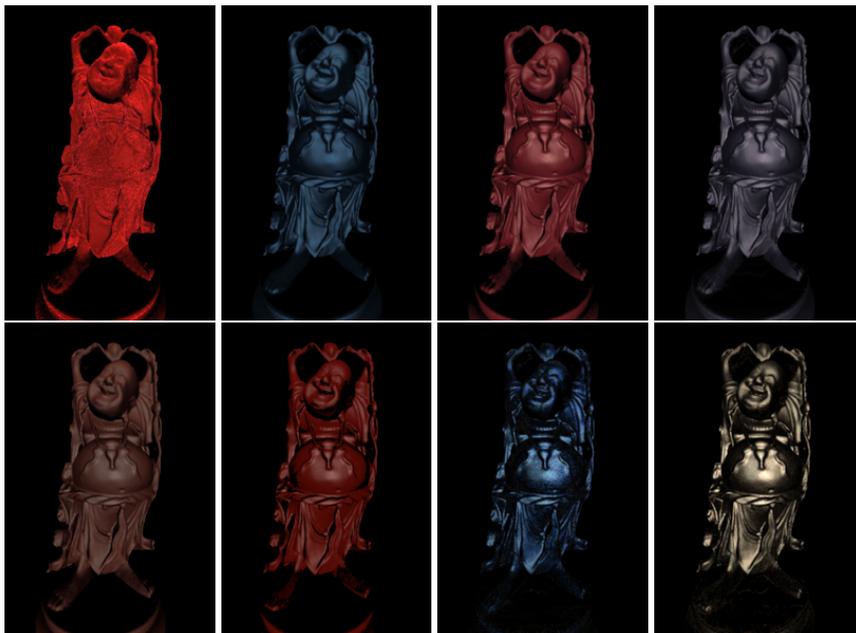


Figure 3.7: Representative set of BRDFs acquired with lower order zonal basis functions rendered under directional lighting. Top row: from left to right - red velvet, dark dark blue synthetic fabric, red printer toner, magenta plastic sheet. Bottom row: from left to right - brown leather, glossy red paper, *KrylonTM* blue paint, chrome gold dust automotive paint.

2% and 10% respectively.

3.4 Results

Using our prototype setup, we have acquired the BRDFs of various types of materials, including velvet, anisotropic synthetic, silk and satin fabrics, leather, various kinds of glossy and shiny papers, paint and plastic samples, printer toners, wax, highly specular metal foil wrapping papers, and anisotropic samples such as a guitar pick and a copper coin. Figure 3.12 presents a selection of BRDFs as rendered on a sphere under a directional light source. Most of the materials were acquired using lower order ($l \leq 6$) zonal basis functions. The silk and satin fabrics, and the guitar pick were acquired

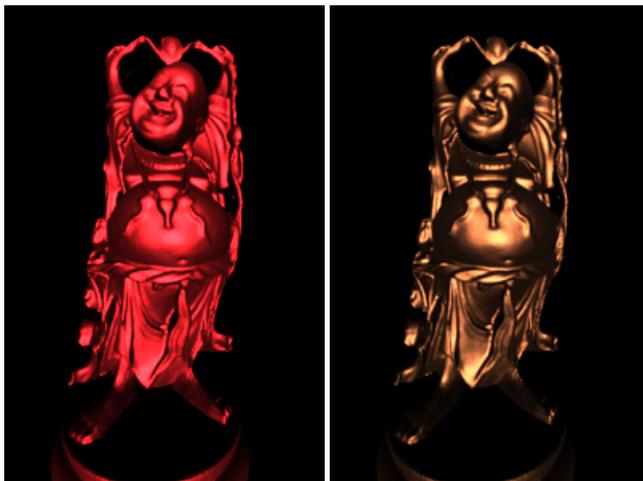


Figure 3.8: Specular chocolate wrapping papers acquired using higher order zonal basis functions, and then fit to an analytical model for rendering. Left: Red *KitKat*TM wrapping paper. Right: Copper colored *Lindt*TM chocolate wrapping paper.

with order $l = 8$ zonal basis function, while the shiny wrapping papers and anisotropic copper coin required acquisition with order $l = 10$ zonal basis function.

The entire process of acquisition consists of separately projecting the positive and negative lobes of the basis functions and recording the response to these lobes in HDR. We used 3 exposures each at 2 f-stops apart for the HDR sequence. Thereafter, in a post-process we construct HDR images from the 3 different exposures and then subtract the responses to the negative lobes from the responses to the corresponding positive lobes in software. Hence, for a 4th order zonal basis acquisition this results in the processing of $25 \times 2 \times 3 = 150$ response images, which takes about one minute. Higher order basis acquisitions for very specular materials take 5 – 7 minutes.

Figure 3.13 presents the BRDFs of 2 different paint samples that we acquired using 4th order zonal basis functions, rendered on the Audi-TT car model, and illuminated by an HDR environment map using the Physically Based Ray Tracing (PBRT) system [103].

A representative set of the BRDFs acquired using lower order ($l \leq 6$) zonal basis functions is shown in Figure 3.7. For this class of materials, the entire process of acqui-

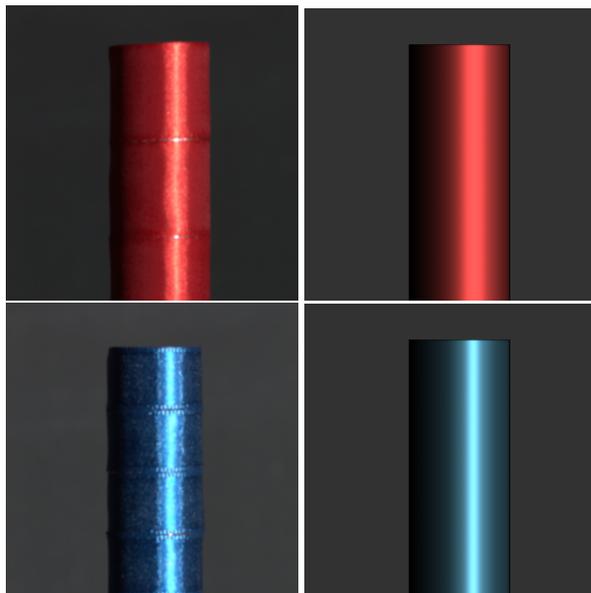


Figure 3.9: Visual comparison of two kinds of acquired satin samples wrapped around a cylinder as lit by a point source against real photographs. Left column: Photographs of the red and blue satin samples. Right column: Rendering of the D-BRDF fit to the acquired data.

sition followed by a basis transformation into the SH basis took under three minutes.

Figure 3.8 demonstrates the specular materials, in this case shiny metal foil chocolate wrapping papers, which we then fit to the D-BRDF analytical model. The D-BRDF fitting procedure consists of constructing the distribution of the half-vector ω_h between the incident light direction ω_i and exitant viewing direction ω_r as a function of the back-scattering direction measurements, i.e., the directions where $\omega_i = \omega_r$. In our case, we extract the zonal half-vector distribution from the measured data, and then extrapolate that to cover the full hemisphere of half-vector directions. The entire acquisition and fitting procedure took only a few minutes to complete in all examples. Similarly, we also fit the anisotropic guitar pick, the copper coin and the red satin sample to the D-BRDF analytical model (Figure 3.12, bottom row).

Finally, as a way of validating our measurement and fitting approach, we photographed two satin samples wrapped around a cylinder in a dark room and lit by a col-

limited point light source. Figure 3.9 presents the comparisons of these photographs with the corresponding renderings of the D-BRDF fits to these samples. The highlights in the rendered images are a close match to the real photographs validating our approach.

3.4.1 Rendering with Acquired BRDFs

Once we have acquired the BRDF data using our basis acquisition approach, we need to integrate the data into existing rendering systems. In this section we discuss how to use such data, either encoded in the SH basis representation or as a D-BRDF fit, in real-time rendering systems and ray-tracing systems for high quality off-line renderings.

Real-time Rendering

Most of the BRDF data acquired using our system is encoded in the SH basis representation after undergoing a basis transformation from the acquired zonal basis representation. Similar to the approach of Kautz et al. [65], we discretize the exitant viewing direction (θ_r, ϕ_r) and encode the incident directions corresponding to each exitant direction as SH coefficients $y_l^m(\theta_r, \phi_r)$ for run-time efficiency of evaluation of the reflectance function. During rendering, the reflectance function is reconstructed for every discretised exitant direction from the corresponding SH coefficients. The reflected radiance in the viewing direction L_r for a local lighting model is then given by:

$$\begin{aligned} L_r(\theta_r, \phi_r) &= \int_{\Omega} \hat{f}_r(\omega_i, \omega_r) L_i(\omega_i) d\omega \\ &= \int_{\Omega} \sum_{l,m} y_l^m(\theta_r, \phi_r) Y_l^m(\omega_i) L_i(\omega_i) d\omega \end{aligned} \tag{3.17}$$

For real-time rendering on current generation graphics hardware, we need to discretize both the exitant and the incident directions and pack the coefficients and the basis functions into an RGBA 3-D float texture. Here, the discretised directions are indexed as 2-D coordinates within a slice of the 3-D texture and the l, m slices each correspond a separate SH basis function. Values for arbitrary directions need to be bilinearly interpolated explicitly within each slice as this is not currently supported in

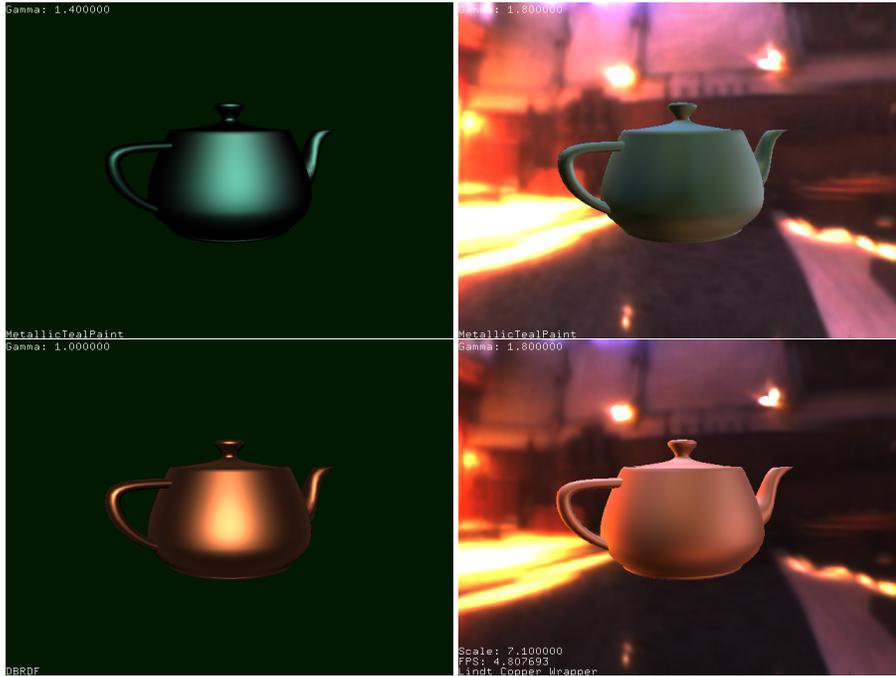


Figure 3.10: Real-time rendering of acquired BRDFs. Top row: Metallic teal automotive paint rendered using the SH basis representation. Bottom row: D-BRDF fit to copper colored *Lindt*TM chocolate wrapping paper. Left column: single directional light source. Right column: Grace cathedral HDR environment map lighting approximated with 128 directional lights generated by the Median Cut algorithm.

hardware. For efficiency of packing as well as memory access pattern, we pack the coefficients $y_l^m(\theta_r, \phi_r)$ into the RGB channel and the corresponding SH basis function value $Y_l^m(\theta_i, \phi_i)$ in the alpha channel of the 3-D texture (Figure 3.10, top row). In our implementation, we have used a resolution of 90×90 for the spherical directions (θ, ϕ) .

For integrating environment map illumination in real-time rendering, one approach would be to transform the environment map into the SH basis and evaluate the reflected radiance on the fly as a dot product of the BRDF coefficients and the environment map coefficients [42, 110]. However, such an approach results in ringing artifacts for high frequency lighting. For such situations, it is preferable to approximate the environment map with a number of point lights either using point relaxation [2, 69] or importance

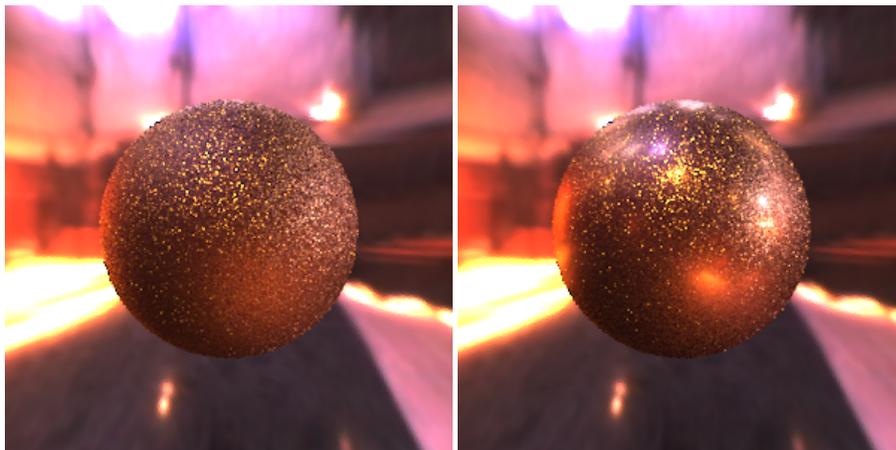


Figure 3.11: Comparison of importance sampling strategies for acquired BRDF of copper colored *Lindt™* chocolate wrapping paper in high frequency HDR lighting of Grace Cathedral EM. Left : Importance sampling from a cosine lobe (default option in PBRT). Right : Importance sampling from the extracted half-vector distribution $p(\hat{h})$ of the D-BRDF fit to acquired data. Both images were generated with 100 samples per pixel in 20 seconds. Note the significant reduction in overall image noise, particularly around the specular highlights, with sampling according to the extracted distribution.

sampling [98] schemes. In our implementation, we use the simple procedure of the Median Cut algorithm proposed by Debevec [28] to generate a point light approximation of the environment map (Figure 3.10, right column).

We use the D-BRDF representation to fit the acquired zonal basis data for highly specular BRDFs instead of transforming to the SH basis representation. Real-time rendering with the D-BRDF fit (Figure 3.10, bottom row) involves storing the extracted 2-D half-vector distribution $p(\hat{h})$ in graphics hardware as a 2-D float texture indexed by (θ_h, ϕ_h) . In addition, we need to precompute and store the Fresnel term $F(\hat{k} \cdot \hat{h})$ as a 1-D float texture since the exponentiation involved in evaluating the Fresnel term (Equation 2.14) is currently not supported in hardware.

Monte Carlo Ray-tracing

Monte Carlo ray-tracing is often the choice of algorithm for high-quality offline rendering. In this discussion, we will restrict ourselves to rendering with complex direct illumination in the form of HDR environment maps. Integrating the acquired BRDF data into a ray-tracing system is fairly straightforward. We have extended the PBRT raytracing system (in terms of its *BxDF* class) to incorporate our acquired BRDFs in the SH representation or as a D-BRDF fit. For the SH representation, we employ the discretization scheme of Section 3.4.1 over the exitant direction ω_r while evaluating the SH function $Y_l^m(\omega_i)$ for the incident direction ω_i at run-time. Similarly, for the D-BRDF representation we load the half-vector distribution $p(\hat{h})$ as a discretised table over (θ_h, ϕ_h) while evaluating all other terms including the Fresnel term at run-time.

When using acquired BRDFs to describe materials in a scene, a ray-tracing system needs to be able to sample from the specific BRDF representation. An efficient importance sampling scheme becomes particularly important in the presence of complex environment map illumination. One way to sample from BRDF data encoded in the SH basis representation would be to do a convex combination of samples drawn from individual SH lobes. Such a scheme would involve first choosing an SH lobe to sample from based on the distribution of the coefficients y_l^m for a given exitant direction ω_r , and then generating a sample from that SH lobe using a 2-D CDF inversion method. This method is further complicated for the SH representation due to the presence of negative lobes of the SH functions in addition to the positive lobes. We implemented such a sampling scheme in PBRT and found it to be not very efficient, particularly for low frequency BRDFs and small sample counts.

Instead, we make the observation that the proposal distribution for importance sampling needs to simply be similar to the BRDF function, and as long as one can efficiently draw samples from the proposal distribution we do not need to use the exact same BRDF representation for rendering and importance sampling. Hence, we advocate to draw samples from our acquired BRDFs by constructing the underlying half-vector distribution $p(\hat{h})$ from the measured reflectance data in the D-BRDF framework as discussed in Section 2.1.1. At run-time, we sample a half-vector direction (θ_h, ϕ_h)

using a 2-D CDF inversion scheme, reflect the view vector ω_r about this sampled half-vector \hat{h} to obtain the sampling direction ω_i . The final step is to appropriately normalize the importance sampling weight p according to Equation 2.30. Figure 3.11 demonstrates the benefit of importance sampling from the corresponding D-BRDF distribution of the acquired data over sampling from a cosine lobe, which is the default option for BRDFs in PBRT.

3.5 Discussion

In this chapter, we have presented a novel basis function approach to BRDF measurement. Our contributions include a novel theory for basis function BRDF acquisition, the development of an orthonormal basis for spherical zones, and the design of an optical setup that allows for basis function illumination of BRDF samples. We also discuss how to integrate the acquired BRDF data into real-time and off-line rendering systems.

The dome we use in our prototype setup covers a sufficient percentage of the hemisphere to obtain high quality BRDF measurements with our basis function approach. To further increase quality by reducing the amount of extrapolation, a dome with a larger coverage could be used. It would be interesting to look into manufacturing techniques that are able to produce such domes.

Due to the basis function approach and the dispensing of all moving parts, BRDF measurement with our setup is very fast, reducing the acquisition time to a few minutes even for high-frequency materials. Moreover, the physical dimensions of the setup are quite compact, so that the whole apparatus could be enclosed in a box the size of a small suitcase. Such a variant of the system would be quite mobile, and could be used on movie sets and in similar settings.

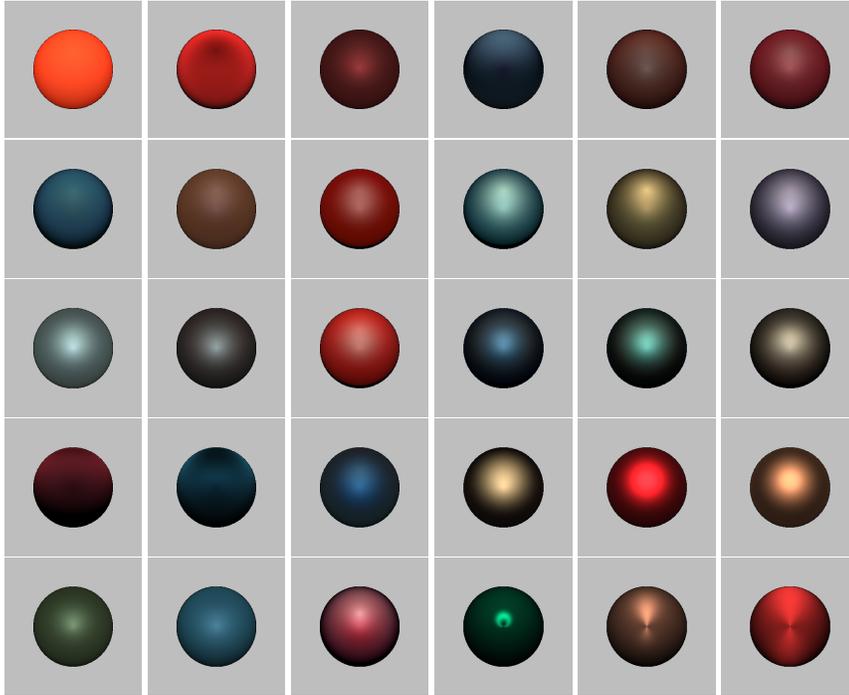


Figure 3.12: Various BRDFs acquired with our prototype setup using zonal basis functions. Top row: from left to right - bright orange paper, red velvet, maroon synthetic fabric, anisotropic dark blue synthetic fabric, brown leather, red printer toner. Second row: from left to right - blue printer toner, coated brown paper, glossy red paper, glossy blue-gray paper, *LindtTM* chocolate box paper, magenta plastic with grain finish. Third row: from left to right - retro-reflective plastic, dark brown plastic coffee lid, *KrylonTM* banner red paint, *KrylonTM* true blue paint, metallic teal automotive paint, chrome gold dust automotive paint. Fourth row: from left to right - purple anisotropic silk fabric, blue anisotropic silk fabric, shiny blue paper, shiny golden paper, red *KitKatTM* wrapping paper, copper colored *LindtTM* chocolate wrapping paper. Bottom row: from left to right - glossy succulent plant leaf, blue rubber band, red wax, anisotropic plastic guitar pick, anisotropic copper coin, anisotropic red satin.



Figure 3.13: The Audi-TT model rendered with acquired BRDFs of 2 different paint samples. The BRDFs were acquired using 25 4^{th} order basis functions as defined in this chapter, and then rendered with a basis transformation into spherical harmonics. Top: Metallic teal automotive paint. Bottom: *KrylonTM* true blue paint. In each case, the time taken for the entire BRDF measurement process including data capture and re-projection into the spherical harmonic basis was about one minute.

Chapter 4

Sampling Techniques for Direct Illumination

In Chapter 3 we described a novel method for efficient image-based acquisition of BRDFs, and a way to integrate the acquired reflectance data into both real-time as well as offline rendering systems. Besides requiring such realistic representations of day-to-day materials, a rendering system also requires realistic representation of complex real-world illumination. In recent years, image-based representations for illumination such as high dynamic range (HDR) environment maps, textured area lights, and light fields have received considerable attention because images can capture complex real-world illumination that is difficult to represent in other forms. In this chapter, we focus on high quality offline rendering with such complex direct illumination using Monte Carlo ray-tracing.

When integrating image-based lighting such as environment maps into a rendering system, the use of a good sampling strategy for illumination is paramount. While several researchers have recently worked on this problem, the approach taken in most of that work is an importance sampling strategy based on the energy distribution in the image. Unfortunately, such an approach performs poorly for highly specular surfaces, since samples chosen this way have a low probability of residing within the specular lobe. Similarly, if importance sampling is based solely on the BRDF of the surface, then the sampling will not perform well for high frequency illumination. In either case, costly visibility tests are required for directions that contribute little to the surface illumination for a particular viewpoint.

We introduce bidirectional importance sampling in Section 4.2, a method that sam-

ples visibility according to an importance derived from the product of BRDF and environment map illumination [13]. The advantage of this approach is that costly visibility tests are restricted to directions that can contribute significantly to the illumination. The number of visibility tests required for high quality rendering can be reduced drastically as a result.

The reduced visibility tests from bidirectional sampling, however, result in noise in partially occluded regions as these samples do not take visibility into account. Hence, we introduce correlated visibility sampling in Section 4.3, a method that additionally takes visibility into account in the sampling process for partially occluded regions, thereby improving the quality of the shadowed regions [47]. We first introduce some preliminaries of importance sampling of direct illumination before we discuss our techniques.

4.1 Importance Sampling of Direct Illumination

The direct illumination at a point for a given observer direction ω_r is given by Kajiya's *rendering equation* [61] as the following integral:

$$L_r(\omega_r) = \int_{\Omega} f_r(\omega_i \rightarrow \omega_r) \cos \theta_i L_i(\omega_i) V(\omega_i) d\omega_i, \quad (4.1)$$

with L_i denoting the incident illumination from an environment map, f_r representing the BRDF, and V being the binary visibility term. Note that we treat L_i and f_r as scalar-valued functions here and throughout the text. In practice, L_i and f_r are color-valued functions, from which we derive the scalar-valued ones by averaging the color channels.

In order to estimate the reflected radiance L_r , conventional approaches perform importance sampling either solely from the intensity in the lighting or solely from the BRDF. In the former case, we get the importance function

$$q_L(\omega_i) := \frac{L_i(\omega_i)}{\int_{\Omega} L_i(\omega_i) d\omega_i} \quad (4.2)$$

with the corresponding Monte Carlo estimator:

$$\begin{aligned} L_{N,L}(\omega_r) &= \frac{1}{N} \sum_{j=1}^N \frac{f_r(\omega_{i,j} \rightarrow \omega_r) \cos \theta_{i,j} L_i(\omega_{i,j}) V(\omega_{i,j})}{q_L(\omega_{i,j})} \\ &= \frac{\int_{\Omega} L_i(\omega_i) d\omega_i}{N} \sum_{j=1}^N f_r(\omega_{i,j} \rightarrow \omega_r) \cos \theta_{i,j} V(\omega_{i,j}). \end{aligned}$$

The variance of this estimator $\omega_{i,j} \sim q_L(\omega_i)$ can then be computed using standard results for importance sampling (e.g. [121]) as:

$$\text{var}(L_{N,L}) = \frac{\int_{\Omega} L_i^2}{N} \text{var}(f_r(\omega_i \rightarrow \omega_r) \cos \theta_i V(\omega_i)).$$

In other words, when proposing samples from the environment only, the resulting variance is proportional to the variance in the BRDF. Similarly, when proposing solely from the BRDF, variance is proportional to the lights. Either approach will produce significant noise if both the BRDF *and* the illumination contain any high frequency information. The solution of Veach and Guibas [136] was to combine samples drawn exclusively from either the lights or the BRDF. However, a mix of samples still suffers from dependence on the variances of the individual techniques. Ideally, we would like to directly sample from the product of the BRDF and the illumination which, given that the visibility function is generally unknown, is the target distribution for the direct illumination integral.

Figure 4.1 shows angular plots of the probability densities corresponding to the various proposal distributions. The top image depicts samples drawn from a Phong BRDF overlaid onto the energy distribution of an environment map. It is obvious that sampling from the BRDF alone misses the bright lights in the environment. The center image shows samples drawn from an environment map, rendered into the importance function for the Phong BRDF at a specific viewing direction. It can be seen that most of these samples are placed outside the specular lobe of the BRDF. Finally, the bottom image represents samples drawn from the product distribution, as well as the product distribution itself. With this method, all samples reside on bright spots of the environment map but also inside the specular lobe.

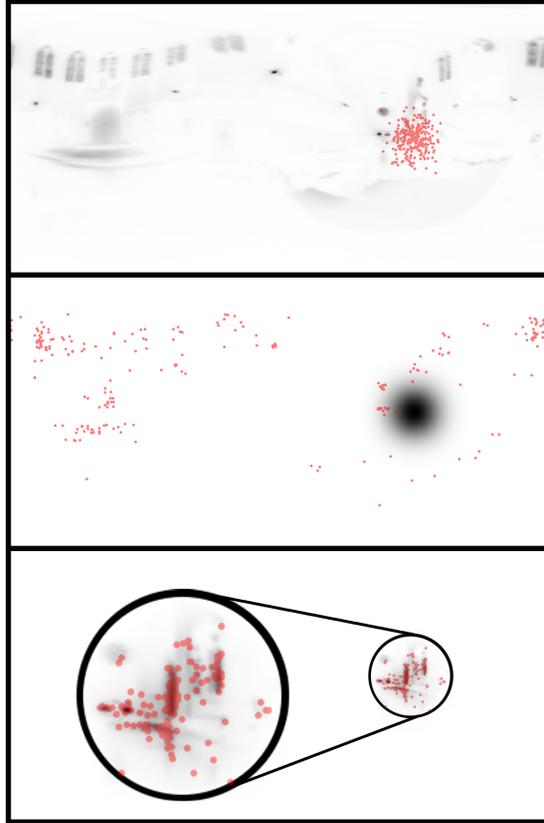


Figure 4.1: From top to bottom: angular plots of the importance function of the Grace Cathedral EM, a specular Phong BRDF of exponent 50, and their product. Samples (red discs) drawn solely from the BRDF or the environment vastly undersample the product distribution. The sample set in the bottom image was generated with our SIR technique (described in Section 4.2.2).

4.2 Bidirectional Importance Sampling

We propose a bidirectional importance sampling approach for direct illumination in which both the energy distribution in the environment map and the reflectance of the BRDF are taken into account. This is a two-step approach: we initially create samples according to only either the BRDF or the environment map, and then adjust these samples to be proportional to the product distribution. The adjusted samples are then

used for visibility testing.

Our approach is to perform importance sampling using the product of the incident light distribution and the BRDF as the importance function:

$$p(\omega_i) := \frac{f_r(\omega_i \rightarrow \omega_r) \cos \theta_i L_i(\omega_i)}{\int_{\Omega} f_r(\omega_i \rightarrow \omega_r) \cos \theta_i L_i(\omega_i) d\omega_i}. \quad (4.3)$$

Observe that the normalization term in the denominator is the direct illumination integral with the visibility term $V(\omega_i)$ omitted. In other words, this term is the exitant radiance in the absence of shadows. We refer to it as L_{ns} (“radiance, no shadows”):

$$L_{ns} := \int_{\Omega} f_r(\omega_i \rightarrow \omega_r) \cos \theta_i L_i(\omega_i) d\omega_i. \quad (4.4)$$

If we draw sample directions $\omega_{i,j} \sim p(\omega_i)$ according to the product distribution in Equation 4.3, we can estimate Equation 4.1 with $L_{N,p}$, where

$$\begin{aligned} L_{N,p}(\omega_r) &= \frac{1}{N} \sum_{j=1}^N \frac{f_r(\omega_{i,j} \rightarrow \omega_r) \cos \theta_{i,j} L_i(\omega_{i,j}) V(\omega_{i,j})}{p(\omega_{i,j})}, \\ &= \frac{L_{ns}}{N} \sum_{j=1}^N V(\omega_{i,j}). \end{aligned} \quad (4.5)$$

We refer to $L_{N,p}$ as the *bidirectional estimator* for the direct illumination integral. The evaluation of Equation 4.5 can be interpreted as taking the unoccluded reflected radiance L_{ns} and scaling it by the average result of N visibility tests performed along directions that contribute most significantly to the radiance. Note that the variance of the bidirectional estimator is now independent of the variance in the illumination and the BRDF.

We have developed two solutions for generation of samples according to the bidirectional importance, one based on rejection sampling and the other on the sampling-importance resampling (SIR) algorithm. We now describe our two realizations of bidirectional importance sampling in the following two sections.

4.2.1 Sample Generation through Rejection

Our first approach for sampling from the product distribution is through *rejection sampling*. To create samples $\omega_{i,j} \sim p(\omega_i)$, we can approximate $p(\omega_i)$ with a PDF $q(\omega_i)$,

such that $p(\omega_i) < c \cdot q(\omega_i)$ for some constant c and all directions ω_i . We then generate random samples $\omega_{i,j} \sim q(\omega_i)$ and accept them with a probability of $p(\omega_{i,j}) / (c \cdot q(\omega_{i,j}))$.

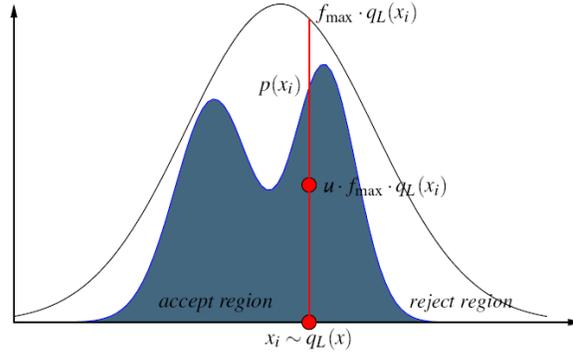


Figure 4.2: Sample generation by rejection sampling. A sample $x_i \sim q_L(x)$ is accepted as being a valid sample of the target distribution $p(x)$ if a uniform sample in $[0, f_{\max} \cdot q_L(x)]$ falls under the product distribution $p(x_i)$.

In our particular case, a simple way of bounding $p(\omega_i)$ from Equation 4.3 is to use q_L , the energy distribution of the light sources (Equation 4.2), as the approximation. The bounding constant in this case is $f_{\max} := \max_{\omega_i} q_f(\omega_i)$, the largest value of the BRDF distribution over all incident light directions but for a given fixed exitant direction. Clearly, $p(\omega_i) < f_{\max} \cdot q_L(\omega_i)$. Figure 4.2 illustrates rejection sampling using this approach.

In order to accept N visibility samples, on average we have to create $M \approx f_{\max} \cdot N$ environment map samples $\omega_{i,j}$ through importance sampling, and then accept each sample individually with probability

$$\frac{p(\omega_{i,j})}{f_{\max} \cdot q_L(\omega_{i,j})} = \frac{f_r(\omega_{i,j}) \cos \theta_{i,j} \cdot \int_{\Omega} L_i(\omega_i) d\omega_i}{f_{\max} \cdot L_{ns}}$$

Both this formula and the final radiance estimate from Equation 4.5 require the normalization term L_{ns} from Equation 4.4. We can estimate this term with the M samples drawn from the environment map during rejection sampling.

In general, we should draw the initial samples in such a way that the bounding constant is minimized. That is, if $f_{\max} < L_{\max}$ we should importance sample from

the environment map; otherwise, we importance sample from the BRDF. One inherent downside of using rejection sampling is that one cannot guarantee bounds on the execution time for creating a new sample. If the area between $c \cdot q(\omega_{i,j})$ and $p(\omega_{i,j})$ is large, the probability of sample acceptance will be low.

One way of dealing with this is to choose a maximum number of sample attempts in the rejection sampling. If no samples are accepted, a possible strategy could be to test visibility for a random subset. A less expensive but biased possibility is to use the unoccluded illumination wherever visibility has not been tested at all. The rationale behind this approach is that the rejection process will fail mostly in very dark areas, where the product of illumination and BRDF is very small. In these areas, the visibility term will not have significant impact anyway. In practice, we have not found it necessary to resort to these biased methods, since the rejection sampling acceptance probability has been sufficiently high even in the presence of highly specular BRDFs and complex environments.

4.2.2 Sample Generation through SIR

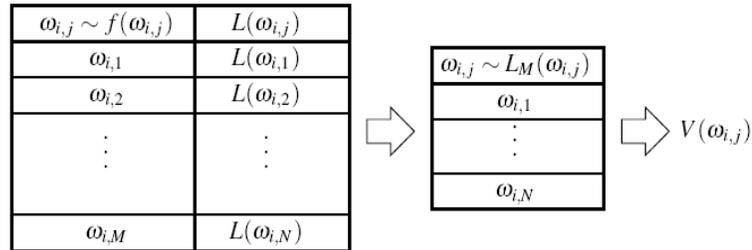


Figure 4.3: Sampling-importance resampling (SIR). First, M samples are proposed from q_f , the PDF of the BRDF. The candidate directions are then resampled based on the incoming light along those directions, producing N samples for visibility testing. N is generally much less than M .

Our second method from sampling the product distribution does not suffer from the unbounded execution time of the rejection sampling. This method uses the so-called *sampling-importance resampling* (SIR) algorithm [43, 125, 131].

SIR first draws a set of M samples $X = \{x_1, \dots, x_M\}$ from a simple distribution $q(x)$. The actual target distribution $p(x)$ is evaluated at these M samples, and the resulting values are used to approximate p . In a second step, a smaller set of N samples $Y = \{y_1, \dots, y_N\}$ is drawn from X with sample probabilities $w(x_i)$ proportional to their importance ratio $p(x_i)/q(x_i)$. As the number of first-round samples M approaches infinity, the sample set Y can be shown to have been drawn directly from p . The closer q approximates p , the faster the method converges.

We can apply SIR to the problem of drawing samples from the bidirectional distribution. We can use either q_L (i.e., sampling from the light sources) or q_f (i.e., sampling from the BRDF) for the first stage. As in the rejection sampling approach, starting with q_L is advantageous if the illumination contains higher frequencies than the BRDF and vice versa, since the higher frequency factor better approximates the shape of the product distribution.

Figure 4.3 summarizes the approach. The total number of samples generated for each pixel is exactly $M + N$. This is an improvement over rejection sampling for two reasons. First, execution time is tightly bounded. We no longer have to wait an indeterminate time for the rejection criterion to accept a sample. Using the SIR algorithm, samples can be drawn directly from the product distribution in constant time.

The second improvement over rejection sampling is that the sample sizes M and N can be chosen freely, yielding fine control over the tradeoff between quality and time. The sample size M dictates the quality of the estimate of L_{ns} , and hence the quality of unoccluded regions. Also, it is possible to directly select N — the target number of visibility rays traced per pixel — based on, for example, scene complexity. Typical values of M are one to two orders of magnitude larger than N . Note that conventional importance sampling from either the BRDF or the illumination alone are just special cases of the SIR technique where $M = N = 1$.

4.2.3 Bidirectional Sampling Results

In the following, we compare the results of our bidirectional sampling techniques with previous sampling strategies for rendering from environment maps. Our comparisons

examine the output quality of the various rendering algorithms for a fixed amount of computing time. We performed these tests on a 3.0 GHz P4 running Linux.

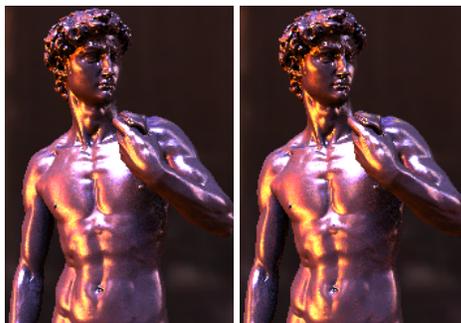


Figure 4.4: Quality comparison between our two proposed bidirectional sampling methods. Left: Rejection sampling. Right: Sampling-importance resampling (SIR). 176×248 images computed in 13.0 seconds using 15/800 rejection and SIR samples.

Figures 4.4 and 4.6 contain images of Michelangelo’s David in the Grace Cathedral environment. We use the version of David with 700k-triangles acquired from the Digital Michaelangelo Project [126]. In our implementation, intersecting a ray with the David model takes, on average, $6.1 \mu s$ on our test machine. The Grace Cathedral environment is a 1024×512 HDR map with a contrast ratio of $10^7 : 1$. In all tests, each algorithm was given 13.0 seconds to render a 176×248 image.

In a first test, we compared rejection sampling and SIR (Figure 4.4). Both the algorithms produced images of indistinguishable quality at the same computing time for a variety of combinations of materials and environment maps.

Figure 4.6 compares bidirectional sampling to earlier methods: sampling only from either the lights or BRDFs, and Veach&Guibas’ multiple importance sampling [136]. In the latter case, the weights for choosing between lights and BRDF were optimized manually through trial and error. For bidirectional importance sampling, we used SIR with $M = 800$ primary samples and $N = 15$ final samples for which visibility was tested.

The first row of the figure uses a glossy Phong BRDF with an exponent of 10. In this case, sampling from the environment map only (left column) is still preferable to sampling from the BRDF (center left), since the environment map contains higher fre-

quencies than BRDF. Even so, sampling from the environment map only results in visible noise. Multiple importance sampling produces a result comparable to environment map sampling, while bidirectional sampling clearly outperforms all other methods.

The second row of Figure 4.6 shows the same scene with a shinier BRDF (Phong exponent of 50). Now, sampling from the BRDF produces better results than sampling from the environment map. Multiple importance sampling further improves on this result. However, bidirectional sampling again outperforms all other methods. In the last row of the figure, we added a diffuse component. This significantly lowers the quality of BRDF sampling. Again, bidirectional sampling is superior to the other strategies without having to adjust weights as in the case of multiple importance sampling.

Finally in Figure 4.5, we present a comparison of the convergence in terms of RMS errors for importance sampling and bidirectional sampling. The plot here was computed for the David model (Phong exponent 50, $k_s = 0.5, k_d = 0.5$) in the Grace Cathedral environment, with first round sampling from the illumination and resampling based on the BRDF. It is clear from the figure that the RMS error converges faster for bidirectional sampling. We found similar behavior for other materials and environment maps.

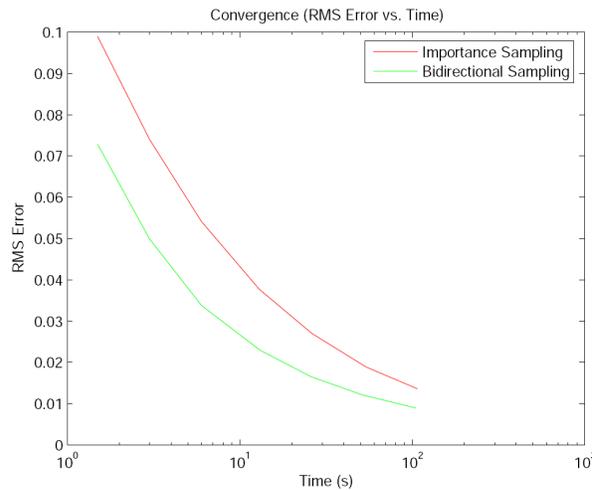


Figure 4.5: Convergence plots of RMS errors for importance sampling and bidirectional sampling. Note how the RMS error reduces faster for bidirectional sampling.

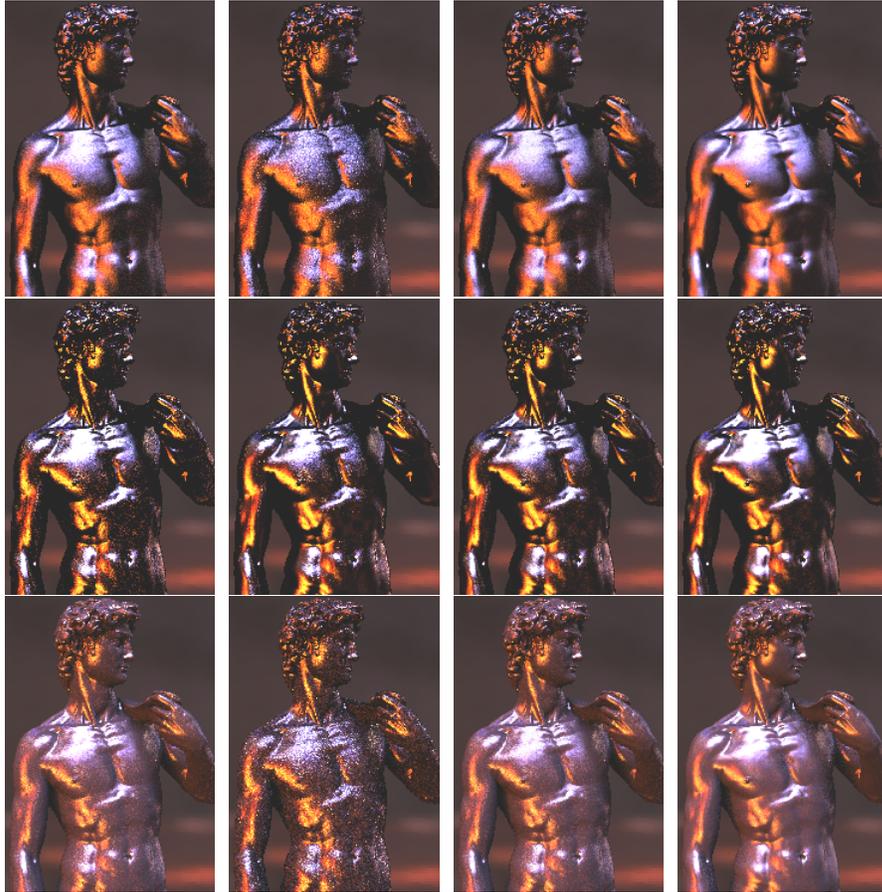


Figure 4.6: David in Grace Cathedral – 176×248 images rendered in 13.0 seconds. Left column: Importance sampling purely from the illumination (100 samples). Center left: Importance sampling purely from the BRDF (75 samples). Center right: Combined sampling (Veach&Guibas) with manually fine-tuned weights. Right: Bidirectional importance sampling with SIR (15/800 samples). Top row: Phong exponent 10, $k_s = 1.0, k_d = 0.0$. Center: Phong exponent 50, $k_s = 1.0, k_d = 0.0$. Bottom row: Phong exponent 50, $k_s = 0.5, k_d = 0.5$.

4.3 Correlated Visibility Sampling

The variance of the bidirectional estimator for the reflected radiance is proportional to the variance in the visibility function (Equation 4.5). This is an improvement over sampling techniques that only consider either the illumination or the BRDF in the sam-

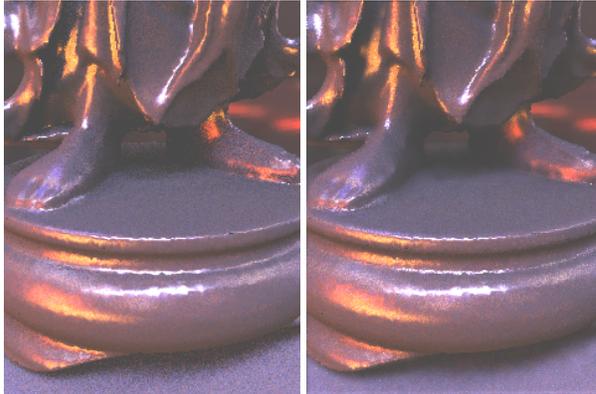


Figure 4.7: Buddha model (Phong BRDF $s = 50$, $k_s = k_d = 0.5$) in Grace Cathedral EM. Left: Bidirectional importance sampling, 20 samples/pixel. Right: Correlated visibility sampling, 16 bidirectional samples (1st pass) and 16 Metropolis samples per unoccluded sample (2nd pass). Rendering times are identical (24 seconds).

pling process. This is because the variance of these techniques depends in addition on the variance in the function that they do not sample from, BRDF or illumination respectively. However in regions with complex visibility, estimates with bidirectional sampling will still suffer from considerable variance (Figure 4.7, left).

In this section we introduce *correlated visibility sampling*, a method that additionally takes visibility into account in the sampling process for partially occluded regions, thereby improving the quality of the shadowed regions (Figure 4.7, right). The aim of this technique is to develop an efficient means of drawing samples from the triple product of the incident illumination, BRDF and visibility, which we achieve by employing the Metropolis-Hastings (MH) algorithm [88]. We describe two variants of the method, one which is unbiased, and a more efficient one that is consistent, but may exhibit a small amount of bias. Note that in our work on sampling the direct illumination, we make the assumption that the visibility function is binary. In general, for scenes with semi-transparent objects, visibility can be continuous. However, it should be noted that binary visibility is the worst case scenario for a sampling algorithm as it results in a very high frequency function at the shadow boundaries and our visibility sampling approach is equally applicable to scenes with continuous visibility.

Our solution is a two-step approach. In the first step, energy estimates for each pixel

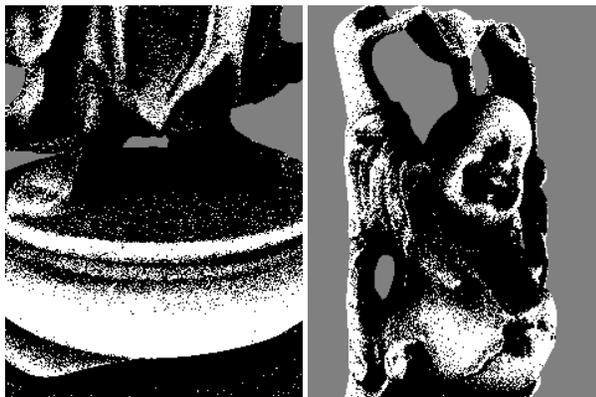


Figure 4.8: Visibility masks for the images presented in Figures 4.7 and 4.11. The white pixels correspond to unoccluded pixels at the end of first round of bidirectional sampling while the black pixels correspond to the partially occluded pixels.

are created using samples drawn from the bidirectional importance (product distribution) of the incident illumination and the surface BRDF. This estimate is built using the sampling-importance resampling (SIR) algorithm, as discussed in Section 4.2.2. We create a visibility mask and mark pixels for which one or more of the visibility tests failed, i.e., pixels which are partially occluded (Figure 4.8). If desired, any image-space operation such as dilation can be applied to the visibility mask. In the second step, Metropolis sampling is started for the partially occluded pixels in order to locally explore the shadowed regions more extensively. Metropolis sampling is only started for the partially occluded pixels, thereby avoiding unnecessary visibility tests in unoccluded regions.

Given a non-negative function f , the MH algorithm generates a series of samples $X = \{x_1, x_2, \dots, x_n\}$ from a distribution proportional to f , which is also referred to as the target distribution, without requiring to normalize f and invert the resulting PDF. It is thus applicable to a wide variety of sampling problems and was first applied in computer graphics by Veach & Guibas [137] to the problem of image synthesis. Given a current sample x , the next sample x' in the sequence is generated by randomly mutating x and then accepting or rejecting the mutation. The mutations are accepted or rejected in such a way that the samples converge to the target distribution. For a description of

the MH algorithm, we refer the reader to the chapter on Metropolis Sampling by Pharr in the SIGGRAPH 2004 course notes [37].

The MH algorithm generally suffers from a *startup bias* as the initial samples in the sequence are not drawn according to the target distribution and thus need to be discarded. Despite the startup bias, integral estimates according to the MH algorithm are asymptotically unbiased as long as *detailed balance* is maintained [137]. Detailed balance defines an *acceptance probability* of a mutation strategy:

$$a(x \rightarrow x') = \min\left\{1, \frac{f(x') \cdot T(x' \rightarrow x)}{f(x) \cdot T(x \rightarrow x')}\right\}, \quad (4.6)$$

where x is the current sample and x' is the mutated sample, $f(x)$ is the function being integrated and $T(x \rightarrow x')$ is the cumulative *transition probability* of mutating from x to x' . Note that the acceptance probability accounts for changes in surface orientation and surface BRDF from one pixel to another, for example between the diffuse ground plane and the specular Buddha model in Figure 4.7.

Since we begin our Metropolis sampling from an unbiased Monte Carlo estimate, our method does not suffer from startup bias. However, in general, a small amount of bias may originate from using samples in unoccluded regions for both estimating the direct illumination, and for making a decision about entering the correlated sampling stage. This issue is discussed in more detail in Section 4.3.1

We employ *lens perturbation* as the mutation strategy for our algorithm. Since there is correlation in the visibility of points in neighboring pixels, using this strategy to transfer energy of samples $\omega_{i,x}$ to neighboring pixels x' can be an effective means of reducing variance. We partition the image plane into 5×5 tiles (Figure 4.9) for lens perturbation and carry out mutations only between the partially occluded pixels within each tile. First a mutation of a valid unoccluded sample (obtained from first round of bidirectional sampling) is proposed. Visibility is then sampled in the same direction (for environment map illumination) for the pixel that the sample gets mutated to. If the visibility test passes, the mutation is accepted with a probability a , else it is rejected. If the mutation is accepted, energy is transferred from pixel coordinate x to x' .

In our case, the cumulative transition probability $T(x \rightarrow x')$ needs to account for

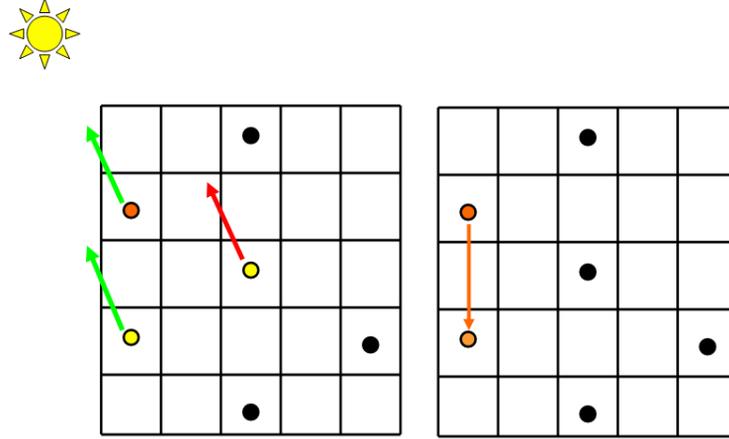


Figure 4.9: Lens perturbation within a 5×5 transition tile. Left: The source (orange) pixel selects two other (yellow) pixels within the transition tile for energy transfer. Right: Only one pixel is selected for energy transfer based on visibility test in the same direction. Green arrows refer to unoccluded light directions, red arrows to occluded ones.

both the probability of choosing a neighboring pixel x' for transition from pixel x , which we will call $t(x \rightarrow x')$, and for the probability of choosing the same sample direction ω_i to sample illumination at the two pixels according to bidirectional importance p :

$$T(x \rightarrow x') = t(x \rightarrow x') \cdot p(\omega_{i,x'})$$

By restricting mutations to happen only within 5×5 tiles, we ensure that every partially occluded pixel has the same number of neighbors for energy transfer. This ensures that $t(x \rightarrow x') = t(x' \rightarrow x)$.

However, imposing fixed transition tiles could potentially lead to block artifacts at the tile boundaries. Hence, in practice we employ a *moving* tile mechanism for transition centered around the current pixel x . For example, instead of performing $C = 16$ path mutations on a single tile, we chose 16 different partitions of the image plane into 5×5 tiles with different offsets, and perform one mutation each. Following the argument from above, this yields an estimate for every pixel and each of the C tile offsets. The total estimate for one pixel is then computed as an average of $N \cdot C$ transitions from the individual tile offsets, which does not introduce bias.

The acceptance probability of the above mutation strategy then reduces to:

$$a(x \rightarrow x') = \min\left\{1, \frac{f(x') \cdot p(\omega_{i,x})}{f(x) \cdot p(\omega_{i,x'})}\right\}, \quad (4.7)$$

where

$$f(x) = f_r(\omega_{i,x} \rightarrow \omega_{r,x}) \cos \theta_{i,x} L_i(\omega_{i,x}),$$

since the visibility term $V(\omega_{i,x}) = 1$, and

$$p(\omega_{i,x}) = f_r(\omega_{i,x} \rightarrow \omega_{r,x}) \cos \theta_{i,x} L_i(\omega_{i,x}) / L_{ns,x},$$

where $L_{ns,x}$ is the unoccluded radiance in the viewing direction given in Equation 4.4. The numerator of $p(\omega_{i,x})$ cancels out with $f(x)$ in Equation 4.7, further reducing the equation to

$$a(x \rightarrow x') = \min\left\{1, \frac{L_{ns,x'}}{L_{ns,x}}\right\}. \quad (4.8)$$

$L_{ns,x}$ can be estimated from the first phase of bidirectional sampling for each pixel and hence does not need to be recomputed during the correlated sampling phase.

The reflected radiance at each partially occluded pixel is then estimated as

$$L_{vis}(\omega_r) = \frac{1}{N \cdot C} \sum_{j=1}^N \sum_{k=1}^C V(\omega_{i,x}^{(j)}) L_{ns,x'} a(\omega_{i,x'}^{(j)} \rightarrow \omega_{i,x}^{(j)}), \quad (4.9)$$

where L_{vis} is the *visibility estimator* in the viewing direction ω_r , $L_{ns,x'} a(\omega_{i,x'}^{(j)} \rightarrow \omega_{i,x}^{(j)})$ is the fraction of energy received at pixel x from a neighboring pixel x' during each transition, N is the number of bidirectional samples chosen from first round sampling, and C is the number of energy transitions (Markov chains of length 1) employed in the second round to spread the energy of unoccluded samples, i.e., the valid samples of the target distribution.

4.3.1 Bias

As we have introduced the method this far, it produces results consistent with the true illumination, but it may exhibit a small amount of bias for finite sample sizes. We use samples for both estimating the direct illumination, and for deciding whether to start Metropolis mutations in partially occluded regions. Such dual use results in a bias, as



Figure 4.10: Visual comparison of the biased (left) and the unbiased version (right) of our method. Note that the highlights are crisper in the unbiased solution.

pointed out by Kirk and Arvo [67], although the bias is typically small – it is less than the standard deviation of the Monte Carlo estimate in the first stage.

In practice, we find this bias to be small enough to be accepted, but if deemed necessary, we can derive an unbiased variant of our method by splitting the MC sample set from the first phase into two partitions: one for deciding whether to apply the Metropolis algorithm, and one used for estimating the illumination in case Metropolis is not necessary. Since we are now using a smaller set of samples to estimate visibility in partially occluded regions, we have to use slightly larger sample chains to achieve the same quality of results. Since the number of visibility tests remains the same, this can be done at low additional cost. Figure 4.10 shows a comparison of the biased and the unbiased version of the algorithm.

4.3.2 Correlated Sampling Results

In this section we compare the results of our correlated visibility sampling with bidirectional importance sampling for rendering from HDR environment maps. Images were generated with a reasonably well-optimized ray tracer using a voxel grid as the acceleration data structure for intersection queries. Our comparisons examine the output quality of the two discussed rendering algorithms for a fixed amount of computing



Figure 4.11: Buddha model (Phong BRDF $s = 50$, $k_s = k_d = 0.5$) in an indoor HDR EM. Left: Importance sampling from EM, 200 samples/pixel. Center Left: Multiple importance sampling from EM (140 samples/pixel) and BRDF (40 samples/pixel). Center Right: Bidirectional importance sampling, 20 samples/pixel. Right: Correlated visibility sampling, 16 bidirectional samples (1^{st} pass) and 16 Metropolis samples per unoccluded sample (2^{nd} pass). Rendering times are identical (16 seconds).

time. We performed these tests on a 3.0 GHz Xeon running Linux SuSE 9.0.

Figure 4.8 presents the visibility masks obtained from first round bidirectional sampling for the images in Figures 4.7 and 4.11. The white pixels represent unoccluded pixels which would not be processed in our second round of sampling. The gray pixels correspond to the background environment map. Finally, the black pixels correspond to those where one or more visibility samples were occluded during first round of bidirectional sampling. These pixels are deemed partially occluded and are processed during our second round of correlated visibility sampling.

Figure 4.7 presents a complex visibility scenario with the Buddha model in the Grace Cathedral Environment. The Buddha model has 300K triangles, while the Grace Cathedral environment is a 1024×512 HDR map with a contrast ratio of $10^7 : 1$. In this test, both the bidirectional sampling and the correlated sampling algorithms were given 24 seconds to render one 176×248 image each. The time budget was chosen so as to allow good quality in unoccluded regions. For bidirectional sampling, this time budget allowed for visibility to be tested with $N = 20$ samples, and these N samples were chosen after resampling from a larger set of $M = 800$ samples. Note how the shadows between Buddha's feet as well as on the ground-plane are noisy with bidirectional sampling. For the same compute time, the partially occluded regions are very

smooth with our correlated sampling approach. Here, we used $N = 16$ first round bidirectional samples for the unoccluded regions, and then $C = 16$ Metropolis samples to spread the energy of the unoccluded samples in our second round of sampling. The un-occluded reflected radiance L_{ns} was estimated using fewer samples ($M=725$) with our correlated sampling approach. Hence, our algorithm produces slightly noisier results in these un-occluded regions. However, the overall tradeoff is better with this approach. The RMS error compared to a converged image reduced from 0.066 when using bidirectional sampling to 0.058 when using the correlated sampling approach for Figure 4.7. And visually, the image rendered with correlated sampling are much more pleasing due to lower noise in the shadowed regions.

Figure 4.11 presents a scene with visibility not as complex as that of Figure 4.7 and with lower frequency illumination. Here we compare the performance of our correlated sampling approach with standard importance sampling from EM, multiple importance sampling from EM and BRDF, as well as bidirectional sampling. Due to high frequencies in both the EM and the BRDF, multiple importance sampling has better performance than sampling only according to the EM. Bidirectional sampling does better than both these approaches in reducing image noise as it samples according to the product distribution. Even then, our correlated sampling approach is more effective than bidirectional sampling in reducing noise in partially occluded regions such as the inside of Buddha’s arms and regions around the face and chest.

Figure 4.13 presents another visibility situation with the Dragon model (870K triangles) in the Grace Cathedral environment. Here, the regions on the Dragon’s neck underneath the head as well as on the body are partially occluded by other parts of the Dragon’s body. Again, our correlated sampling approach nicely cleans up the shadowed regions that remain noisy with the bidirectional sampling approach. Figure 4.12 presents the visibility masks for the images in Figure 4.13. The mask on the left shows many pixels in generally unoccluded areas on the Dragon’s body as well as its head are marked as partially occluded after first round of sampling. The mask on the right is obtained after applying a simple dilation operation with a 3×3 kernel to the original mask and is better representative of the visibility situation. The dilation operation reduces the variance in the penumbra region, and reproduces sharper shadow bound-

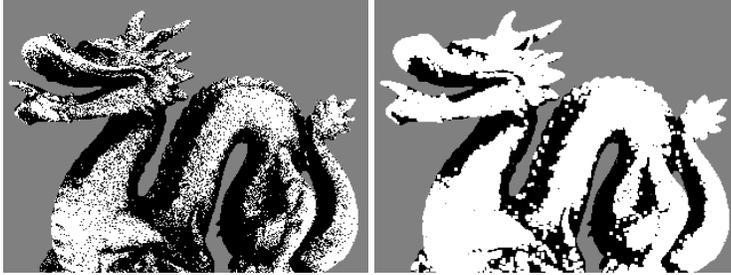


Figure 4.12: Visibility mask for the images in Figure 4.13. Left: Undilated mask. Right: Dilated mask.

aries in regions that are generally unoccluded such as the Dragon’s head (Figure 4.13, bottom).

The implementation of our correlated sampling approach involves the usual time vs. memory tradeoff. Compared to the bidirectional sampling approach that processes each pixel independently, the correlated sampling approach needs to store information about neighboring pixels and the visibility mask at the end of the first stage of sampling. For efficiency, we store the N bidirectional samples for each pixel obtained from first stage sampling as well as the estimate of L_{ns} for each pixel. In addition, in order to prevent having to trace primary rays for every transition during correlated sampling, we also store the information corresponding to primary rays such as vertex position, vertex normal and view vector for every pixel. Thus our implementation incurs an additional memory overhead of $\sim W \times H \times N \times Sample$, where $W \times H$ is the resolution of the image plane and $Sample$ is a triple of floats used for storing a sample/position/normal.

With these memory overheads, our correlated sampling stage only required an additional 5 – 10% computation time after the first stage bidirectional sampling. This additional time was mostly spent in areas with high occlusion such as ground planes occluded by geometry in Figure 4.7.

4.4 Discussion

In this chapter, we discussed the problem of sampling the direct illumination, especially in the presence of HDR environment maps. We demonstrate how sampling only

from the distribution of illumination or only the BRDF can be sub-optimal in this scenario. We proposed a bidirectional importance sampling approach for sampling from the product distribution of the illumination and BRDF in Section 4.2, and described two Monte Carlo techniques for realizing bidirectional sampling. While this approach makes the sample selection process more expensive, we drastically reduce the number of visibility tests required to obtain good image quality. As a consequence, we achieve significant quality improvements over previous sampling strategies for the same compute time.

In Section 4.3, we extended bidirectional sampling to also account for the visibility function. In this case, we employed the Metropolis-Hastings algorithm for establishing a correlation in the energy estimates of neighboring pixels, thereby reducing noise in partially occluded regions. To our knowledge, this is the first proposed solution for sampling the direct illumination integral according to the triple product distribution that does not require any precomputation of visibility.

The correlated sampling method incurs additional memory overheads for storing samples generated from first round of bidirectional sampling. However, this overhead is linear in the number of pixels in the image, and hence is not that significant. Given this small memory overhead, the correlated sampling stage reduces noise in shadowed regions with complex visibility with a small (5 – 10%) additional computation time. For this extra computation time, our method of employing Metropolis sampling after an initial phase of Monte Carlo sampling provides much greater benefit in terms of image quality in shadowed regions compared to pure Monte Carlo sampling.

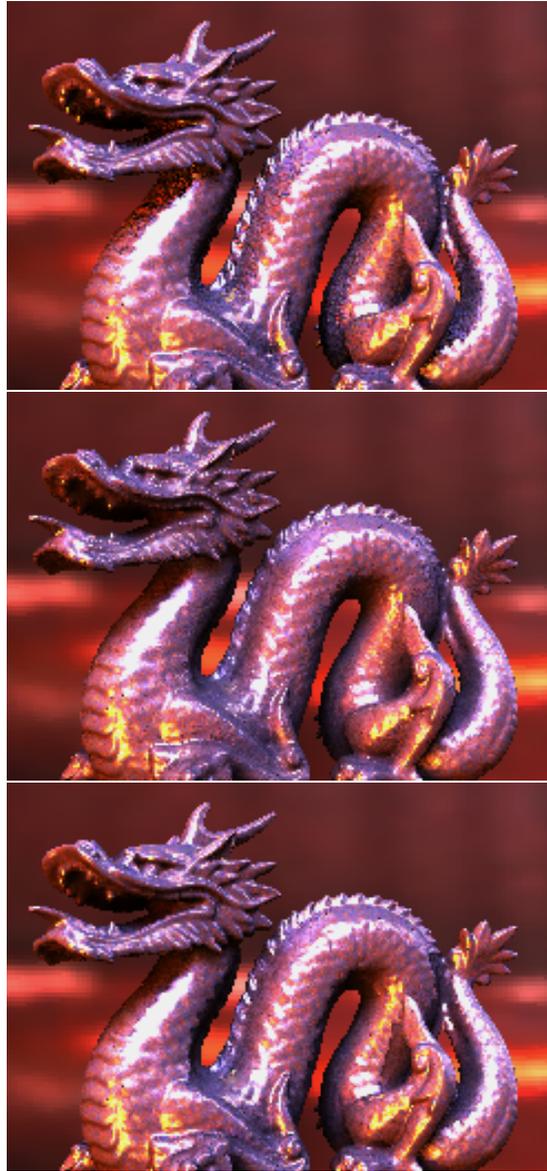


Figure 4.13: Dragon model (Phong BRDF $s = 50$, $k_s = k_d = 0.5$) in the Grace Cathedral HDR EM. Top: Bidirectional importance sampling, 20 samples/pixel. Center and Bottom: Correlated visibility sampling, 16 bidirectional samples (1st pass) and 16 Metropolis samples per unoccluded sample (2nd pass). Center: Undilated visibility mask. Bottom: Dilated visibility mask. Rendering times are identical (16 seconds).

Chapter 5

Sequential Sampling of Environment Maps

As discussed in Chapter 4, realistic rendering with complex direct illumination has received significant attention in recent years with major applications in realistic re-lighting. The best known techniques for direct illumination sample from the product of the incident illumination and the surface reflectance [13, 16, 129]. With advancements in HDR acquisition technologies, HDR video environments are becoming increasingly available. This availability has spawned recent work on sampling dynamic illumination from such HDR video sequences [56, 139]. However, these techniques only take the dynamic importance of the illumination into account while proposing samples for the video sequence. Such techniques are problematic in the presence of high frequencies in both the illumination as well as the surface BRDFs. While one can produce noise-free images by using one set of samples for all surface locations, doing so eliminates the impact of dimmer light sources, which should dominate the reflection for certain surface orientations.

In this chapter, we aim to efficiently sample from the product distribution of the illumination and the BRDF in a video sequence with dynamic illumination using a *Sequential Monte Carlo (SMC)* sampling strategy [46]. The basic idea is to generate samples according to the product distribution in the first frame of the sequence, and thereafter to filter these samples (particles) in the subsequent frames according to the dynamic product distribution. This sequential sampling mechanism is more efficient than independently re-regenerating the samples for every time step (Figure 5.1), especially for scenes with high frequencies in both the dynamic illumination and the

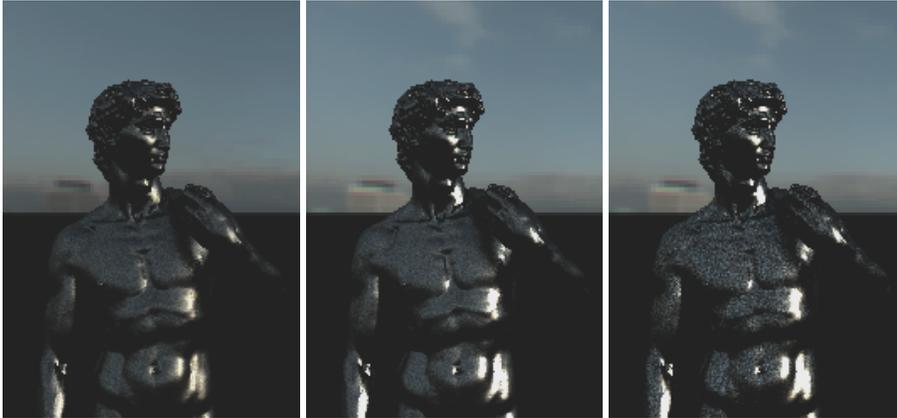


Figure 5.1: Quality comparison of our SMC sampling algorithm with bidirectional importance sampling for a sky probe sequence. Left: 1st frame rendered using bidirectional importance sampling in 8 seconds. Center: 5th frame of sequence rendered using SMC sampling in 4 seconds. Right: Comparison image for 5th frame generated using bidirectional importance sampling in 4 seconds.

BRDF. At the same time, our method avoids systematic under-estimation of reflections at certain angles, which is common to dynamic importance methods generating point light approximations of the environment. Another very important benefit of the SMC approach for rendering a video sequence is the temporal coherence of the sampling pattern that results in significant reduction of sparkling noise in an animation compared to independent sampling.

Our solution to sampling from the dynamic product distribution is a two-step approach. We assume that we have already obtained a sample set according to the product distribution of the previous frame. For the first frame of the animation, we generate this sample set with bidirectional importance sampling (Section 4.2). In the first step of our algorithm, samples distributed according to the product distribution of the previous time step are propagated in time using sequential importance sampling. The product distribution at the new time step is incrementally estimated using the weights of the sequential importance. The second step extends the path of each of these samples using a Markov Chain Monte Carlo (MCMC) kernel whose invariant distribution is the target distribution at the current time step. The MCMC kernel is implemented using

the Metropolis-Hastings (MH) algorithm [88]. No visibility tests are performed during either of the two steps. Visibility is finally tested at the end of the second step in order to obtain a Monte Carlo estimate of the direct illumination. This approach has the following benefits:

- We need to propagate only a small number of samples in time to estimate the direct illumination as these samples are distributed according to the target distribution at each time step. This makes sample propagation very efficient.
- The normalization constant for the product distribution at each time step can be incrementally computed using the sequential importance weights. Thus, the normalization constant, i.e., the un-occluded reflected radiance at each time step, can be estimated very efficiently without drawing a large number of samples.
- Sample generation cost at each time step is independent of the cost of sampling from the BRDF representation since the algorithm only requires to evaluate the BRDF but does not require to sample from it. Thus, any complex BRDF representation can be used without impacting sample generation cost.
- The method creates samples on the fly and does not require any expensive pre-computation.

5.1 Sequential Monte Carlo Sampling

As already mentioned, we propose a sequential Monte Carlo (SMC) sampling algorithm for sampling according to the dynamic product distribution of the illumination and the surface reflectance during an animation sequence. Traditional SMC algorithms in the literature deal with the case where the target distribution of interest at time n , defined on Ω_n , is of a higher dimension than the target distribution at time $n - 1$ [33]. A classical example of this is an SMC algorithm applied to sequential Bayesian inference. In our case, the target distribution at every time step, i.e., the direct illumination integral, is defined on a common space of the hemisphere of directions Ω . Hence, we employ a class of SMC samplers recently developed for a common domain [31] to the

problem of sampling from the product distribution of incident illumination and BRDF in the presence of dynamic illumination.

Symbol	Description
$\omega_i^{(j)}$	j^{th} sample (incident direction)
$W_n^{(j)}$	Weight for j^{th} sample in frame n
$\tilde{w}_n^{(j)}$	Unnormalized weight for j^{th} sample in frame n
N	Number of samples per pixel
p_n	Target PDF for frame n
\tilde{p}_n	Unnormalized importance (product of illumination and BRDF)
q_n	Proposal distribution at time n
$P_{n-1:n}^{k/P}$	The k^{th} of P intermediate distributions between frame $n-1$ and frame n

Table 5.1: Summary of notation used in this chapter.

Our SMC sampling algorithm is a two-step approach: we start with samples created according to the product distribution in the previous time step, and propagate them in time using sequential importance sampling followed by a possible resampling step. We then employ an appropriate MCMC transition kernel to redistribute the samples according to the product distribution at the new time step. Thereafter these samples, now distributed according to the new target distribution, are used for visibility testing.

Once again, consider the direct illumination at a point for a given observer direction ω_r :

$$L_r(\omega_r) = \int_{\Omega} f_r(\omega_i \rightarrow \omega_r) \cos \theta_i L_i(\omega_i) V(\omega_i) d\omega_i, \quad (5.1)$$

with L_i denoting the incident illumination from an environment map, f_r representing the BRDF, and V being the binary visibility term.

The target distribution of interest for direct illumination is the product distribution p of incident illumination and the BRDF:

$$p(\omega_i) := \frac{f_r(\omega_i \rightarrow \omega_r) \cos \theta_i L_i(\omega_i)}{\int_{\Omega} f_r(\omega_i \rightarrow \omega_r) \cos \theta_i L_i(\omega_i) d\omega_i} = \frac{\tilde{p}(\omega_i)}{L_{ns}}. \quad (5.2)$$

Here, $\tilde{p}(\omega_i) = f_r(\omega_i \rightarrow \omega_r) \cos \theta_i L_i(\omega_i)$ is the un-normalized importance function, and $L_{ns} = \int_{\Omega} f_r(\omega_i \rightarrow \omega_r) \cos \theta_i L_i(\omega_i) d\omega_i$ is the un-occluded reflected radiance in the

viewing direction and is the normalization constant of the target distribution. Following the convention of Section 4.2, we refer to L_{ns} as "radiance, no-shadows".

Our SMC algorithm works as follows: we start with samples $\omega_{i,n-1}^{(j)}$ and weights $W_{n-1}^{(j)}, j = 1, \dots, N$, such that the weighted samples are proportional to the product distribution of BRDF and illumination in frame $n - 1$. These samples represent the incident light directions for one surface point, that is, the point visible at a specific pixel. For the first frame, these samples are obtained by bidirectional importance sampling, and all weights are $1/N$. The two steps of our method are then *sample propagation* followed by *MCMC transition* to adjust the samples to the product distribution in the next frame.

Step 1: Sample Propagation: We propagate the samples $\omega_{i,n-1}^{(j)}$ in time using sequential importance sampling. The un-normalized *incremental* weight \tilde{w}_n of every sample for sequential importance at time n is given by the following ratio:

$$\tilde{w}_n^{(j)} = \frac{\tilde{p}_n(\omega_{i,n-1}^{(j)})}{\tilde{p}_{n-1}(\omega_{i,n-1}^{(j)})}, \quad (5.3)$$

This weight is just the ratio of the target function evaluated at the sample point at time n to that evaluated at time $n - 1$ for the same point, and represents the change in weighting of a sample due to changes in the target distribution. The normalized weights for the N samples at time n are then given by:

$$W_n^{(j)} = \frac{W_{n-1}^{(j)} \cdot \tilde{w}_n^{(j)}}{\sum_{k=1}^N W_{n-1}^{(k)} \cdot \tilde{w}_n^{(k)}}. \quad (5.4)$$

This tracking of weights to represent evolving target distributions is called *Sequential Monte Carlo Sampling*, or SMC for short [31]. Note that we can use the SMC mechanism for sampling from a dynamic sequence even if we use a different proposal distribution q instead of p for the first frame. The only difference here would be that we would have to appropriately weight the samples of the first frame for sequential importance. The un-normalized weights $\tilde{w}_1^{(j)}$ would then need to be computed as:

$$\tilde{w}_1^{(j)} = \frac{\tilde{p}_1(\omega_{i,1}^{(j)})}{\tilde{q}_1(\omega_{i,1}^{(j)})},$$

and then normalized to obtained $W_1^{(j)}$.

Step 1a: Resampling. As the variance between the proposal distribution q_n and the target distribution p_n tends to increase with n , the variance of the un-normalized importance weights tends to increase resulting in a potential degeneracy of the particle approximation. This degeneracy can be measured using the criterion of *effective sample size* (ESS) $(\sum_{j=1}^N (W_n^{(j)})^2)^{-1}$ [78]. The ESS takes values between 1 and N . If the ESS is below a pre-specified threshold, say $N/2$, we resample the N samples according to the weights $W_n^{(j)}$ and set the weights of the resampled samples equal to $1/N$. This resampling step discards samples with low weights while copying the samples with high weights multiple times, thus keeping the samples according to the target distribution. Note that the mutations in the next step make sure we do not keep multiple identical samples.

Step 2: MCMC Transitions. After sample propagation and potential resampling, we apply an MCMC kernel $K_n(\omega_i, \omega'_i)$ of invariant distribution p_n to every sample $\omega_{i,n-1}^{(j)}$ in order to obtain new samples $\omega_{i,n}^{(j)}$. The new samples $\omega_{i,n}$ are marginally distributed according to

$$p_n(\omega'_i) = \int_{\Omega} p_{n-1}(\omega_i) K_n(\omega_i, \omega'_i) d\omega_i. \quad (5.5)$$

We employ the Metropolis-Hastings algorithm (MH) for these transitions, with a mix of local random walk moves and independent proposal moves. When using the MH algorithm, the MCMC kernel K_n of invariant distribution p_n is described in terms of an acceptance probability of a proposed transition:

$$a(\omega_i^{(j)} \rightarrow \omega'_i{}^{(j)}) = \min\left\{1, \frac{\tilde{p}_n(\omega'_i{}^{(j)})q(\omega_i^{(j)} \rightarrow \omega'_i{}^{(j)})}{\tilde{p}_n(\omega_i^{(j)})q(\omega'_i{}^{(j)} \rightarrow \omega_i^{(j)})}\right\}, \quad (5.6)$$

where $\omega_i^{(j)}$ is the current sample, $\omega'_i{}^{(j)}$ is the proposed sample using the transition kernel q , and a is the acceptance probability of the proposed transition.

We mix local walk moves with independent proposal moves as these independent moves are required to prevent the SMC samples from getting stuck in local possibly narrow modes of the target distribution. For high frequency dynamic lighting, we choose local walk moves with uniform random directional perturbations of up to a few degrees of the samples, while choosing samples from the environment map (EM) for

the independent proposal moves. When using local random walk moves, there is an equal probability of transition between $\omega_i^{(j)}$ and $\omega_i'^{(j)}$, i.e., $q(\omega_i^{(j)} \rightarrow \omega_i'^{(j)}) = q(\omega_i'^{(j)} \rightarrow \omega_i^{(j)})$. Thus, the acceptance probability a_{local} of the local walk moves is given by:

$$a_{local}(\omega_i^{(j)} \rightarrow \omega_i'^{(j)}) = \min\left\{1, \frac{\tilde{p}_n(\omega_i'^{(j)})}{\tilde{p}_n(\omega_i^{(j)})}\right\}. \quad (5.7)$$

When using independent proposal moves from the EM, the transition probability of the proposed sample is given by $q(\omega_i^{(j)} \rightarrow \omega_i'^{(j)}) = L_{i,n}(\omega_i'^{(j)}) / \int_{\Omega} L_{i,n} d\omega_i$. Hence the acceptance probability of the independent move is given by:

$$\begin{aligned} a_{EM}(\omega_i^{(j)} \rightarrow \omega_i'^{(j)}) &= \min\left\{1, \frac{\tilde{p}_n(\omega_i'^{(j)})L_{i,n}(\omega_i'^{(j)}) / \int_{\Omega} L_{i,n} d\omega_i}{\tilde{p}_n(\omega_i^{(j)})L_{i,n}(\omega_i^{(j)}) / \int_{\Omega} L_{i,n} d\omega_i}\right\} \\ &= \min\left\{1, \frac{f_r(\omega_i'^{(j)} \rightarrow \omega_r) \cos \theta_i'^{(j)}}{f_r(\omega_i^{(j)} \rightarrow \omega_r) \cos \theta_i^{(j)}}\right\}. \end{aligned} \quad (5.8)$$

In practice, we propose several MCMC moves per sample for good exploration of the target distribution.

Note the SMC algorithm as described above is unbiased: Step 1 corresponds to an importance sampling step, which produces the correct distribution at time n using the distribution at time $n - 1$ as importance. The variance of this distribution is reduced by applying the MCMC algorithm in Step 2. Since MCMC algorithm works on an unbiased distribution, startup-bias is avoided. A formal argument can be found in [31].

5.1.1 Direct Illumination Estimate

With the sample sets and weights derived above, we can now estimate the reflected radiance at a surface location as

$$L_{N,n,smc}(\omega_r) = L_{ns,n} \sum_{j=1}^N W_n^{(j)} \cdot V(\omega_{i,n}^{(j)}). \quad (5.9)$$

Equation 5.9 can be interpreted as the scaling of the un-occluded reflected radiance $L_{ns,n}$ by the weighted average of N visibility tests performed along the directions contributing most significantly to the radiance at time n . Here, the normalization constant

$L_{ns,n}$ at time n can be incrementally estimated as

$$\frac{L_{ns,n}}{L_{ns,n-1}} = \frac{\int_{\Omega} \tilde{p}_n(\omega_{i,n-1}) d\omega_{i,n-1}}{\int_{\Omega} \tilde{p}_{n-1}(\omega_{i,n-1}) d\omega_{i,n-1}} \approx \sum_{j=1}^N W_{n-1}^{(j)} \cdot \tilde{w}_n^{(j)}. \quad (5.10)$$

The derivation of this result is given in Appendix C. It is worth pointing out that this incremental estimate of the normalization constant $L_{ns,n}$ at time n according to Equation 5.10 provides the crucial advantage for the SMC algorithm in terms of computational expense over a pure MC approach such as bidirectional sampling where the proper estimate of L_{ns} requires drawing large number of samples. To estimate $L_{ns,n}/L_{ns,1}$, one can use the product of estimates given by Equation 5.10 from time $t = 2$ to n . $L_{ns,1}$ is estimated in our case during bidirectional sampling for the first time step.

A summary of the method can be found in Algorithm 1. The complexity of this algorithm is in $O(N)$.

Algorithm 1 SMC Sampling for Dynamic Illumination

1: INITIALIZATION

- Set $n = 1$.
- For $j = 1, \dots, N$ draw $\omega_{i,1}^{(j)} \sim p_1$ using bidirectional sampling and set $W_1^{(j)} = 1/N$.

Iterate steps 2 and 3:

2: WEIGHTING AND RESAMPLING

- Set $n = n + 1$.
- Compute new weights $W_n^{(j)}$ according to Equations 5.3 and 5.4.
- If $ESS < Threshold$, resample and set $W_n^{(j)} = 1/N$.

3: SAMPLING

- For $j = 1, \dots, N$ draw $\omega_{i,n}^{(j)} \sim K_n(\omega_{i,n-1}, \omega_{i,n})$.
- Estimate $L_{ns,n}$ according to Equation 5.10.
- Estimate reflected radiance according to Equation 5.9.

5.2 Variance Reduction with Intermediate Distributions

The aim of the SMC sampling algorithm as discussed in Section 5.1 is to “smoothly” move samples from the target distribution at time $n - 1$ to the target distribution at time n . More formally, we have samples

$$\omega_{i,n-1}^{(j)} \sim p_{n-1}(\omega_i) = \frac{\tilde{p}_{n-1}(\omega_i)}{L_{ns,n-1}},$$

and we want to move towards samples

$$\omega_{i,n}^{(j)} \sim p_n(\omega_i) = \frac{\tilde{p}_n(\omega_i)}{L_{ns,n}}.$$

This transition is smooth under the assumption that $p_{n-1} \approx p_n$. However, this may not be true in practice, especially in the case where the dynamic illumination is in the form of a high frequency HDR video environment. If the discrepancy between the two successive distributions is too high, this will result in high variance in the un-normalized incremental weights $\tilde{w}_n^{(j)}$, and thus indirectly result in high variance in the normalized weights $W_n^{(j)}$. The variance in $W_n^{(j)}$ can be countered with the resampling step after sequential importance sampling, resulting in good estimates of the posterior \tilde{p}_n . However, the resampling step does not affect the variance in the un-normalized weights $\tilde{w}_n^{(j)}$, which can lead to high variance in the incremental estimate of the normalization constant $L_{ns,n}$ according to Equation 5.10.

In this scenario, we introduce a sequence of intermediate distributions [44] between the original distribution p_{n-1} and the new one p_n in order to select a smooth transition that the sample can follow. These intermediate distributions are blends of the original distributions:

$$p_{n-1:n}^\gamma(\omega_i) \propto \tilde{p}_{n-1:n}^\gamma(\omega_i) = [\tilde{p}_{n-1}(\omega_i)]^{1-\gamma}[\tilde{p}_n(\omega_i)]^\gamma, \quad (5.11)$$

such that

$$p_{n-1:n}^0(\omega_i) = p_{n-1}(\omega_i), \quad p_{n-1:n}^1(\omega_i) = p_n(\omega_i).$$

In practice, we introduce P discrete intermediate distributions: $p_{n-1:n}^{k/P}(\omega_i)$, where $k = 1, \dots, P$. We can use these new distributions to reduce variance with little additional cost. The idea is to reduce the variance in the incremental weights \tilde{w}_n by computing them as a product of P incremental weights \tilde{w}^k of these intermediate distributions. The consecutive intermediate distributions $p_{n-1:n}^{k/P}(\omega_i)$ and $p_{n-1:n}^{(k-1)/P}(\omega_i)$ are closer to each other by construction than $p_{n-1}(\omega_i)$ is to $p_n(\omega_i)$, resulting in flatter weights \tilde{w}^k , as compared to \tilde{w}_n .

The SMC algorithm with P intermediate distributions requires slight modifications to the algorithm discussed in Section 5.1. Instead of first computing the un-normalized

weights $\tilde{w}_n^{(j)}$ and the normalized weights $W^{(j)}$ and then doing the MCMC transitions, the algorithm for P intermediate distributions computes the un-normalized weights $\tilde{w}_n^{(j)}$ as a product of P intermediate un-normalized weights $\tilde{w}^{k,(j)}$ that each involve an MCMC kernel of invariant distribution $p_{n-1:n}^{k/P}$. Here, the intermediate un-normalized weights $\tilde{w}^{k,(j)}$ are computed as:

$$\tilde{w}^{k,(j)} = \frac{\tilde{p}_{n-1:n}^{k/P}(\omega_i^{k-1,(j)})}{\tilde{p}_{n-1:n}^{(k-1)/P}(\omega_i^{k-1,(j)})}. \quad (5.12)$$

Assuming we have samples $\{W_{n-1}^{(j)}, \omega_{i,n-1}^{(j)}\}$ approximating p_{n-1} . The algorithm then proceeds as follows:

Algorithm 2 SMC with Intermediate Distributions

1: INITIALIZATION:

- We write $\omega_i^{0,(j)} = \omega_{i,n-1}^{(j)}$ and set $\tilde{w}_n^{(j)} = 1$.

2: ITERATION: for $k = 1, \dots, P$

- Compute $\tilde{w}^{k,(j)}$ according to Equation 5.12.
 - Set $\tilde{w}_n^{(j)} = \tilde{w}_n^{(j)} \cdot \tilde{w}^{k,(j)}$.
 - Sample $\omega_i^{k,(j)} \sim K_k(\omega_i^{k-1,(j)}, \omega_i^{k-1,(j)})$ of invariant distribution $p_{n-1:n}^{k/P}(\omega_i)$.
-

At the end of the P iterations of intermediate distributions, the normalized weights $W_n^{(j)}$ are still computed according to Equation 5.4, and resampled if the ESS is below the pre-specified threshold. Finally, the normalization constant $L_{ns,n}$ and the direct illumination estimate are computed according to Equations 5.10 and 5.9 respectively.

In general, there is greater benefit in terms of variance reduction in the estimate of the target distribution with the introduction of a sequence of P intermediate distributions involving one MCMC move each than with a single distribution involving P MCMC moves, while being only a slightly more expensive in terms of computation time (Figure 4.5 in the results section). This benefit is because a sequence of interme-

diating distributions simultaneously reduces the variance in the un-normalized weights $\tilde{w}_n^{(j)}$ while exploring the target distribution at time n .

5.3 Unoccluded Illumination Estimate with Path Sampling

In this section, we explore an alternative estimate for the unoccluded radiance L_{ns} , which can be used in place of the method described in Section 5.1.1. We obtain this alternate solution by path sampling [44]. In the statistics literature, the path of a sample is defined as the continuous trajectory of a sample over time. It should not be confused with light paths in classical global illumination literature. With this definition of path, *path sampling* refers to smoothly moving samples from one distribution to the next.

Considering a continuous path of distributions

$$p_{n-1:n}^\gamma(\omega_i) = \frac{\tilde{p}_{n-1:n}^\gamma(\omega_i)}{L_{ns,n-1:n}^\gamma},$$

the following path sampling identity holds [44]:

$$\log\left(\frac{L_{ns,n}}{L_{ns,n-1}}\right) = \int_0^1 \int \frac{d \log(\tilde{p}_{n-1:n}^\gamma(\omega_i))}{d\gamma} p_{n-1:n}^\gamma(\omega_i) d\omega_i d\gamma. \quad (5.13)$$

In our case, the logarithm of the target function \tilde{p} is given by

$$\log(\tilde{p}_{n-1:n}^\gamma(\omega_i)) = (1 - \gamma) \log(\tilde{p}_{n-1}(\omega_i)) + \gamma \log(\tilde{p}_n(\omega_i))$$

according to Equation 5.11.

Thus, the derivative of the logarithm of the function is

$$\frac{d \log(\tilde{p}_{n-1:n}^\gamma(\omega_i))}{d\gamma} = \log\left(\frac{\tilde{p}_n(\omega_i)}{\tilde{p}_{n-1}(\omega_i)}\right). \quad (5.14)$$

When considering a discrete path of P intermediate distribution, we can approxi-

mate Equation 5.13 with

$$\begin{aligned} \log\left(\frac{L_{ns,n}}{L_{ns,n-1}}\right) &= \frac{1}{P} \sum_{k=1}^P \int \log\left(\frac{\tilde{p}_n(\omega_i)}{\tilde{p}_{n-1}(\omega_i)}\right) p_{n-1;n}^{k/P}(\omega_i) d\omega_i \\ &\approx \frac{1}{P} \sum_{k=1}^P \sum_{j=1}^N W^{k,(j)} \log\left(\frac{\tilde{p}_n(\omega_i^{k,(j)})}{\tilde{p}_{n-1}(\omega_i^{k,(j)})}\right). \end{aligned} \quad (5.15)$$

Note that Equation 5.15 involves computing the normalized weights $W^{k,(j)}$ for every intermediate distribution $p_{n-1;n}^{k/P}$, which is not required when using the standard form of the intermediate distributions. The un-normalized intermediate weights $\tilde{w}^{k,(j)}$ are still computed according to Equation 5.12.

When using path sampling, we can also obtain an estimate of $\log(L_{ns}/L_{ns,1})$ as follows

$$\log\left(\frac{L_{ns,n}}{L_{ns,1}}\right) = \sum_{t=2}^n \log\left(\frac{L_{ns,t}}{L_{ns,t-1}}\right). \quad (5.16)$$

Computing the normalization constant using path sampling is a bit more expensive than the standard introduction of the P intermediate distributions as discussed in Section 5.2. However, the estimate of L_{ns} according to Equation 5.15 generally has lower variance than using P intermediate distributions with the standard ratio according to Equation 5.10.

5.4 Implementation

We have implemented the algorithm described above in a system that offers two rendering modes: a relighting mode, and a mode that allows for free camera movement.

Relighting. In relighting mode, the camera and object are in fixed locations, and only the environment can change. The initial frame is rendered using bidirectional importance sampling. For all subsequent frames, the samples for each pixel are propagated and mutated as described above, and form the sample set for the same pixel in the next frame (Figures 5.4 and 5.5).

Free camera movement. Once the camera is allowed to move freely, a surface point will project to different pixels in different frames. We take this into account by tracking

the motion of the surface points at which the samples were generated. When we ray-trace a sample for one frame, we also store the corresponding information into the next frame. To this end, we compute which pixel the surface point will project to at the next time step, and store all samples from the current frame into that pixel. The memory requirement for this procedure is about 300 Bytes/pixel. When we want to compute the illumination for the next frame, most pixels will therefore have a sample set from the previous frame associated with them. We propagate and mutate those as discussed. Other pixels might not have a sample set due to dis-occlusion, or differences in sampling rate. We start bidirectional importance sampling for these pixels only (Figure 5.6).

General object movements can be dealt with the same way as camera movements: by knowing where object points will be located in the next time step, we can store sample information at the appropriate pixel locations. Currently, our implementation does not support this kind of object motion.

5.5 Results

In this section we compare the results of our unbiased SMC sampling algorithm with bidirectional importance sampling for rendering from HDR video environments. Images were generated with a reasonably well-optimized ray tracer using a voxel grid as the acceleration data structure for intersection queries. Our comparisons examine the output quality of the two rendering algorithms for a fixed amount of computing time. We performed these tests on a 3.6 GHz Xeon running Linux SuSE 9.0.

Figure 5.1 presents a comparison of our SMC algorithm with bidirectional importance sampling for a sequence of the sky probe gallery [127]. The image on the left is the first frame of the sequence rendered at a high quality using bidirectional sampling ($N = 16$, $M = 800$) in 8 seconds. The image in the center is the 5th frame of the sequence rendered using our SMC algorithm ($N = 16$, $P = 5$) with path sampling in 4 seconds. The BRDF of the David model in these images has a high specular exponent (Phong $s = 50$) and no diffuse component. Under these conditions of high frequency

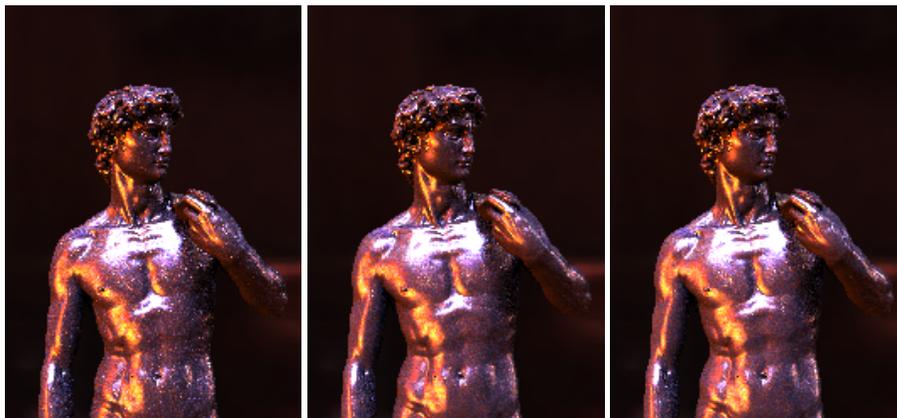


Figure 5.2: Quality comparison of single distribution vs sequence of intermediate distributions with David model (Phong BRDF $s = 50$, $k_s = k_d = 0.5$) in Grace Cathedral HDR EM. Left: Single distribution with 10 MCMC moves. Center: 10 intermediate distributions with L_{ns} computed with standard ratio. Right: 10 intermediate distributions with L_{ns} computed with path sampling. Rendering times are identical (5 seconds).

lighting and highly specular BRDF, our SMC algorithm does much better than bidirectional importance sampling for the same computation time of 4 seconds. In this case, bidirectional sampling could only use a smaller number of samples ($M = 200$) to estimate L_{ns} for the same compute time (right).

Figure 5.2 presents the quality comparison of renderings produced with our SMC algorithm when using a single distribution (left) versus when using a sequence of intermediate distributions with standard ratio (center) and path sampling (right). These images correspond to the first frame rendered by our SMC algorithm after rotating the EM by 1.5° along the radial direction simulating a small change in the HDR illumination of the Grace Cathedral. The sequence of intermediate distributions greatly help in reducing the variance in the incremental computation of L_{ns} compared to a single distribution, while path sampling improves the quality of the estimate a bit more. Here, the BRDF of the David model has a significant diffuse component. Hence, the incremental estimate of the L_{ns} has higher variance compared to Figure 5.1 as the SMC algorithm uses only a very small number of samples to approximate the L_{ns} .

In Figure 5.3, we present a comparison of the convergence in terms of RMS errors

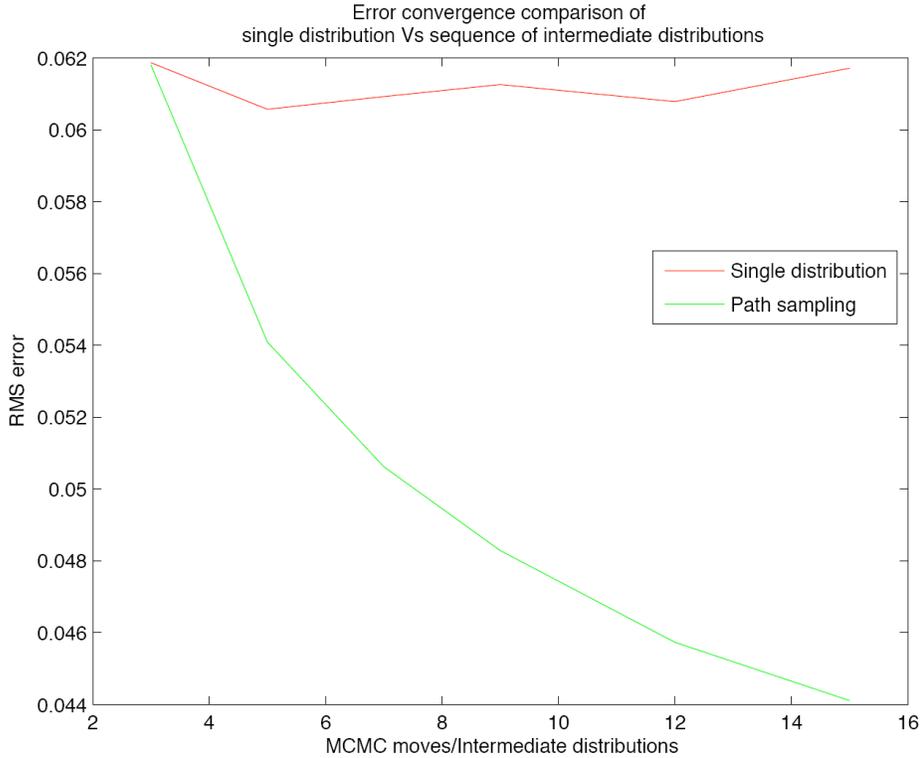


Figure 5.3: Convergence plots of RMS errors for single distribution with multiple MCMC moves and sequence of intermediate distributions with one MCMC move each. Note how the RMS error reduces fast when using a sequence of intermediate distributions while the error does not really reduce much with just one distribution.

for a single distribution with multiple MCMC moves and a sequence of intermediate distributions with path sampling. The RMS error plot was computed for same frame rendered in Figure 5.2 with our SMC algorithm. It is clear from the plot that multiple MCMC moves for a single distribution do not help much in reducing the variance in the incremental computation of L_{ns} , though they may help in exploring the target function \tilde{p} . A sequence of intermediate distributions with 1 MCMC move per distribution, on the other hand, is effective in reducing the variance in the computation of L_{ns} .

The introduction of a sequence of intermediate distributions also helps in reducing the degeneracy of the samples. We tracked the sample degeneracy, in terms of ESS,

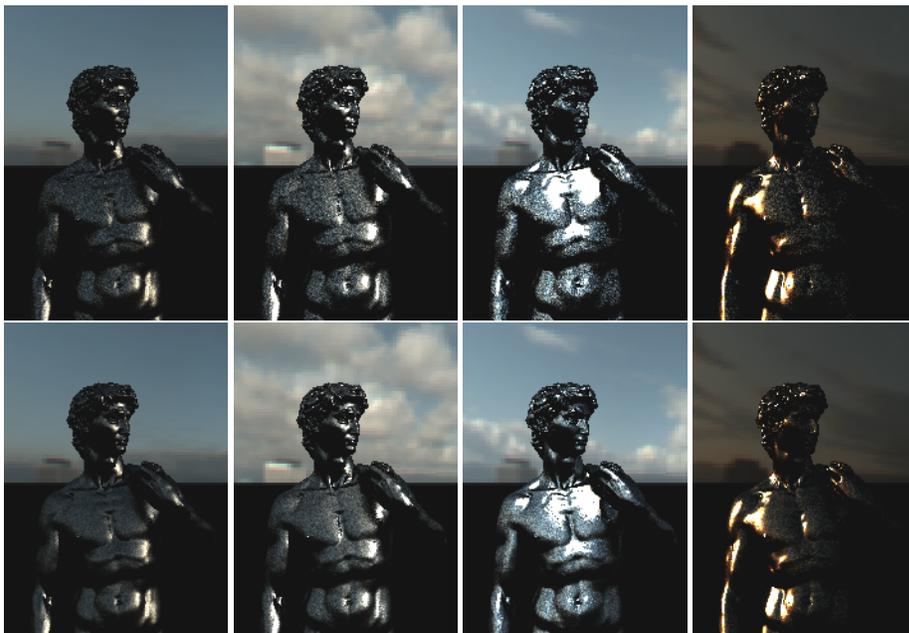


Figure 5.4: Quality comparison between bidirectional sampling and our SMC sampling algorithm for a specular BRDF in the sky probe gallery sequence. Top row: Bidirectional sampling ($N = 16$, $M = 200$). Bottom row: SMC sampling ($N = 16$, $P = 5$, path sampling). All images took the same compute time of 4 seconds.

for sequence of 100 frames while rendering the David model (Phong BRDF $s = 50$, $k_s = k_d = 0.5$) in the Grace Cathedral HDR EM with rotations to the EM by 1.5° per frame. We observed that, on an average, the samples corresponding to 36% of the pixels required resampling when using a single distribution with 5 MCMC moves per sample. This fraction reduced to 18% when using a sequence of 5 intermediate distributions and 1 MCMC move per sample with path sampling. Hence, the additional cost of computing a sequence of intermediate distributions is offset to an extent by having to resample fewer samples. In practice, we used only up to 5 – 7 intermediate distributions as this was enough to reduce the variance of most pixels. However, in order to maintain the quality of the renderings over time, we tracked the incremental estimate of L_{ns} and explicitly computed $L_{ns,n}$ using a large number of samples whenever the ratio $L_{ns,n}/L_{ns,1}$ altered significantly.

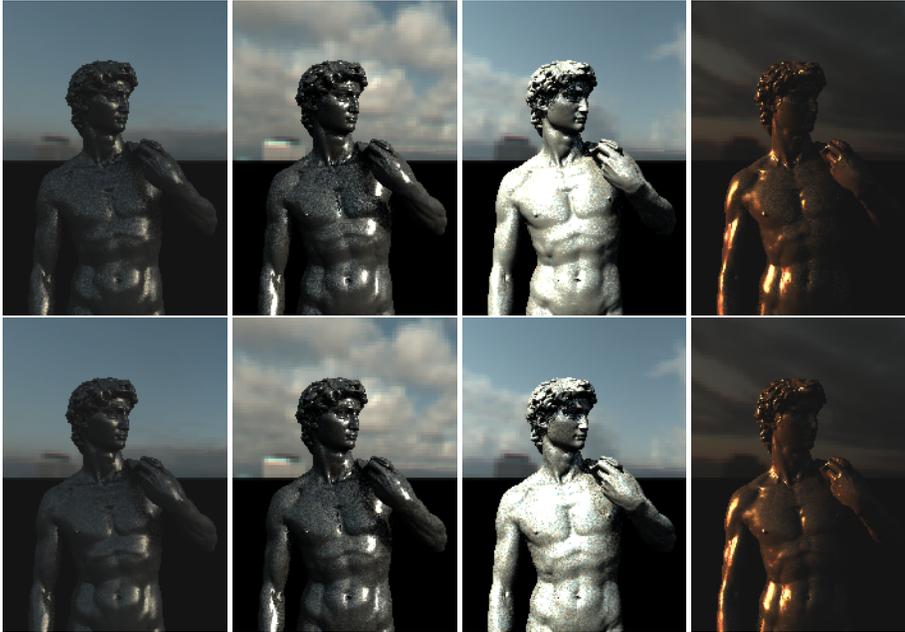


Figure 5.5: Quality comparison between bidirectional sampling and our SMC sampling algorithm for a BRDF with a diffuse component in the sky probe gallery sequence. Top row: Bidirectional sampling ($N = 16$, $M = 200$). Bottom row: SMC sampling ($N = 16$, $P = 5$, path sampling). All images took the same compute time of 4 seconds.

Figure 5.4 presents the quality comparison between bidirectional sampling and our SMC sampling algorithm for a dynamic environment sequence from dawn to dusk from the sky probe gallery [127]. We used the sample HDR sky probe images from the gallery that have been captured at 10 minute intervals as key frames of our sequence and interpolated to create 3 additional frames between each key frame. The David model in these images has the same highly specular BRDF as in Figure 5.1, and in this situation, our SMC algorithm performs much better than bidirectional importance sampling for the same computation time of 4 seconds. The SMC samples also have a lot more temporal coherence, greatly reducing flickering in the resulting animation.

Figure 5.5 presents the quality comparison between bidirectional sampling and our SMC sampling algorithm for the same sky probe sequence from sunrise to sunset,

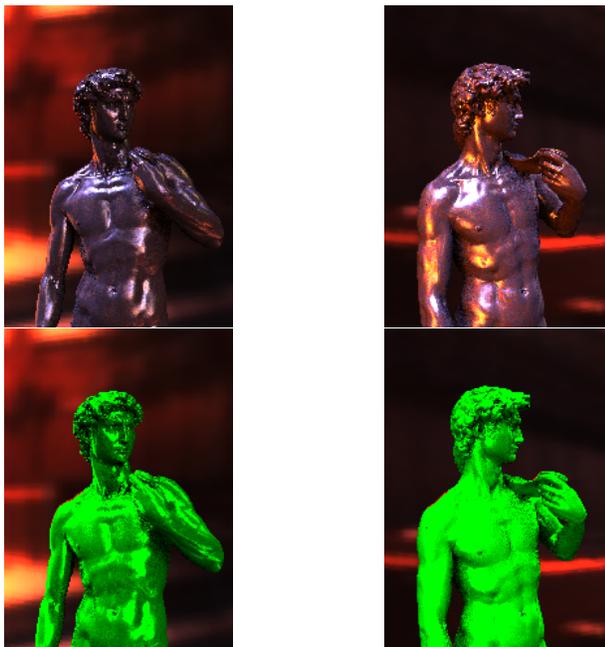


Figure 5.6: SMC sampling for a camera animation sequence with the David model in the Grace Cathedral EM. Top row: Images rendered with SMC sampling for dynamic viewpoints. Bottom row: False color visualization of bidirectional samples (red) and SMC samples (green).

except that the BRDF of the David model now has a significant diffuse component (Phong BRDF $s = 50$, $k_s = k_d = 0.5$). In this case, the difference in the quality of renderings produced by the two algorithms is not much as the SMC algorithm needs to estimate the reflected radiance due to a wider lobe with a small number of samples. However, the images corresponding to the SMC algorithm still have lower variance than bidirectional sampling wherever the specular contribution is high.

In Figure 5.6, we present example renderings with our SMC sampling algorithm adapted for changing viewpoint as discussed in Section 5.4. We rendered a sequence with the camera moving from right to left of the David model with 2° rotation of the camera between every frame. The images in the bottom row are false color visualizations of the pixels corresponding to re-projected points that used SMC sampling (green) and new exposed pixels that used bidirectional sampling (red) for the two viewpoints.

As shown here, in a sequence involving a slowly moving camera, most pixels can effectively use samples propagated from neighboring pixels in subsequent frames.

The speedup that we get in the case of moving cameras is lower than the speedup for relighting, since we cannot use SMC for all pixels due to occlusion, and since the reprojection of samples into the next frame consumes time. While we find a speedup of about a factor of 2 for relighting, the speedup for moving cameras is only 1.6. Note that a moving camera is in some sense the worst case scenario, since all points in the scene move. In general scenes with only few objects moving, one would expect a speedup somewhere between the extremes of relighting and the camera movement.

5.6 Discussion

In this chapter, we have introduced the use of sequential Monte Carlo methods for efficiently computing direct illumination in the presence of both high frequency Illumination and BRDF. By propagating samples over time, the method makes efficient use of coherence across frames. We demonstrate that this approach results in significantly reduced variance for the same compute time compared to other state-of-the-art methods.

Sequential Monte Carlo samplers have been the focus of recent research activities in statistics and machine learning. The sampling strategies used in this work are at the leading edge of methods developed in those areas. We believe that these methods are promising for solving other sampling problems in computer graphics, for example for global illumination with path tracing or photon mapping.

Chapter 6

Real-time Rendering for HDR

Displays

Up until now, we have discussed novel approaches towards photo-realistic rendering with acquisition of reflectance properties of real materials (Chapter 3), and rendering scenes with complex direct illumination in the form of environment maps (Chapters 4 and 5). We now shift our focus to the display end of the image-synthesis pipeline where the goal is to translate the photo-realistic reproduction of a scene by the acquisition and rendering stages of the pipeline into a perceptually realistic, possibly non-linear, reproduction on a display device. Perceptually realistic rendering is mainly concerned with algorithms that aim to recreate the *visceral experience* of a real scene for a human observer. A major factor affecting this visceral experience is the dynamic range of display technology and its capacity for representing the full range of dark and light intensities found in the real world.

The dynamic range of many real-world environments exceeds the capabilities of current display technology by several orders of magnitude. Tone mapping operators alleviate the problem of limited dynamic range of conventional display devices to an extent by simulating the non-linearity of the human visual system, but are unable to compensate fully for these shortcomings [76]. Recently, Seetzen et al. [119] described two alternative designs for high dynamic range (HDR) display systems that are capable of displaying images with a dynamic range much more similar to that encountered in the real world.

Both display designs are based on the fundamental idea of using an LCD panel as an optical filter of programmable transparency to modulate a high intensity but low

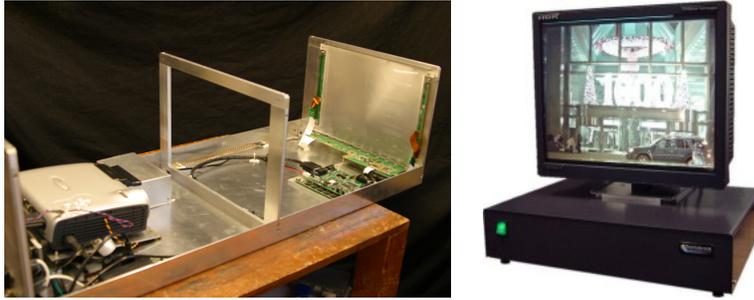


Figure 6.1: Two alternative designs for HDR displays. Left: Projector-based display. Right: LED-based display.

resolution image from a second display. If the first display has a contrast range of $c_1 : 1$ between the darkest and the brightest producible color, and if we now put an LCD panel with a contrast ratio of $c_2 : 1$ in front of the first one, then in principle, the contrast of the combined system is $(c_1 \cdot c_2) : 1$. Based on this principle, Seetzen et al. derived two alternative designs for the HDR displays, one based on a digital light projection (DLP) technology for the backlighting and the second design replacing the projector with a matrix of ultra-bright LEDs. While the DLP-based design is easier to drive, the LED-based design is more suitable for commercial applications due to factors such as significantly smaller form factor and power consumption of the LED-based displays. Prototypes of the two displays are shown in Figure 6.1. The two displays have dynamic ranges well beyond $50,000 : 1$, and a maximum luminous intensity of $2700cd/m^2$ and $8500cd/m^2$, respectively. This compares to a dynamic range of about $300 : 1$, and a maximum intensity of about $300cd/m^2$ for a typical desktop display.

In this chapter, we describe the real-time rendering algorithm that we developed for driving the projector-based HDR display [119], while briefly describing the algorithm for the LED-based design for completeness. We first review some relevant issues on human contrast perception in Section 6.1 before describing the rendering algorithms for the HDR displays. We also present a few possible applications for the projector-based display in Section 6.4.

6.1 Human Contrast Perception

The human visual system has tremendous capabilities but also some limitations. Some of these limitations are an integral part of the theory underlying the HDR displays. In general, the human eye has evolved to deal with the vast dynamic range available in the daily environment, ranging from starlight to sunlight over at least an eight order of magnitude luminance range. To cope with this range, the eye uses a complex adaptation system that mainly works at two time scales: mechanisms working at time scale of the order of minutes and those with a shorter time scale. The adaptation mechanisms at shorter time scales are very interesting from the point of view of HDR display development as they are the primary reason why current displays cannot provide realistic representations of real world HDR scenes. The eye can capture approximately 5 orders of magnitude of dynamic range effectively *simultaneously*. No conventional display technology comes close to this. Yet, there are limitations to this capability as described below.

6.1.1 Local Contrast Perception

While we can see a vast dynamic range across a scene, we are unable to see more than a small portion of it locally in small regions (corresponding to small angles). Different researchers report different values for the threshold past which we cannot make out high contrast boundaries, but most agree that the maximum perceivable contrast is somewhere around 150 : 1 [138]. Scene contrast boundaries above this threshold appear blurry and indistinct, and the eye is unable to judge the relative magnitudes of the adjacent regions.

This inherent limitation can be explained by the scattering properties of the eye. From Moon&Spencer's original work on glare [90], we know that any high contrast boundary will scatter at least 4% of its energy on the retina to the darker side of the boundary, obscuring the visibility of the edge and details within a few degrees of it (Figure 6.2). If the contrast of an edge is 25 : 1, then details on the darker side will be competing with an equal amount of light scattered from the brighter side, reducing visible contrast by a factor of 2 in the darker region. When the edge contrast reaches

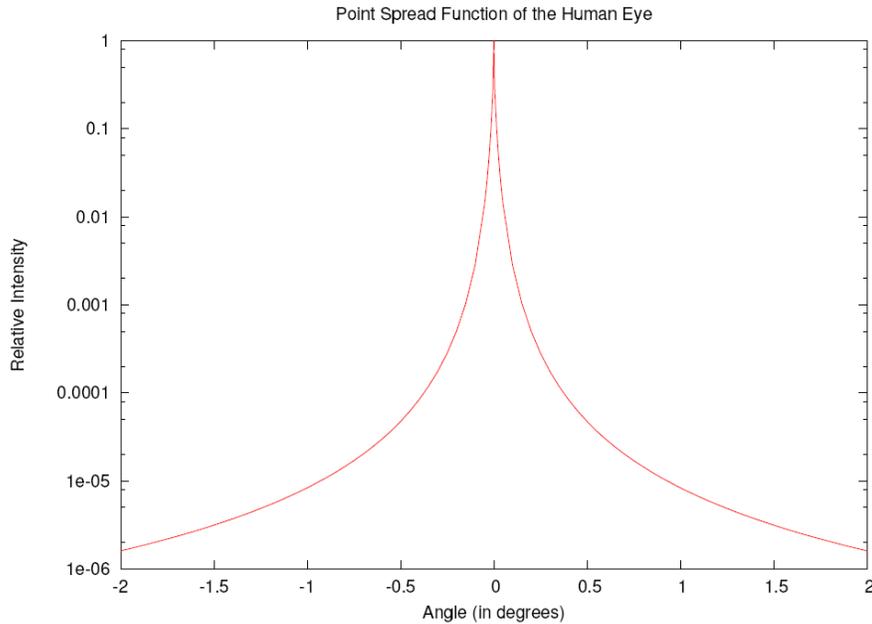


Figure 6.2: The point spread function of the human eye according to Moon&Spencer [90].

a value of 150 : 1, the visible contrast on the dark side is reduced by a factor of 12, rendering details indistinct or invisible.

The HDR display technology exploits this inability of humans to see detail in the immediate vicinity of a high-contrast boundary. It maintains relative (and even absolute) brightnesses, and reproduces edges exactly when they are below the maximum contrast of the front display - about 400 : 1 in the prototypes described in [119]. Only when this range is exceeded is some fidelity lost near high contrast boundaries. But this effect is well below the detectable threshold of the eye, as has been verified in user experiments [120].

6.2 Driving the Projector Display

To correctly render HDR images on the projector-based display, we need to analyze the image formation process of the system. Let us assume for the moment that both the

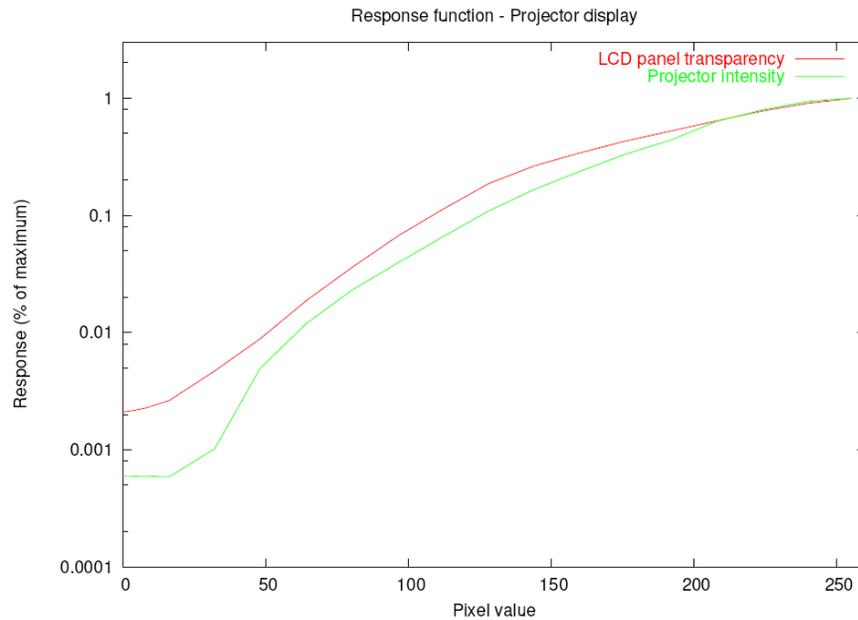


Figure 6.3: Response function of both the LCD panel and the DLP projector in the projector-based display.

projector and the LCD panel are perfectly linear, and that both have the same dynamic range. Under these assumptions, we can achieve the target intensity by scaling the HDR image to the range of $0 \dots 1$, and using the square root of this normalized intensity to drive both the projector and the LCD panel. This even split between pixel values on the projector and the LCD panel is preferable to a scenario where one value is very large and the other is very small, since quantization artifacts are relatively larger for small values.

In reality, neither the projector nor the LCD have a linear response (Figure 6.3). Also, although the projector-based HDR display provides mechanisms for alignment of the DLP (backlight) and LCD (front-panel) pixels, a sub-pixel match is hard to achieve and almost impossible to maintain. To avoid moire patterns and alignment artifacts associated with even a minor misalignment, we have to deliberately blur the projector image and compensate for the blur in the LCD image. We do this in the following way: we choose a simple estimate of what the projector intensity should be, and then

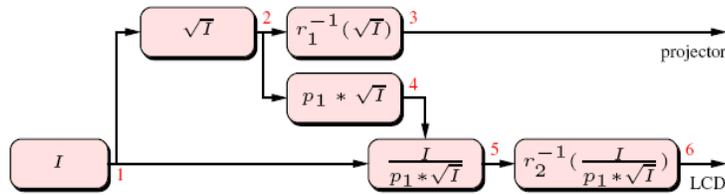


Figure 6.4: Rendering algorithm for the projector-based display.

simulate the effect of response function and blurring. Finally, the pixel values of the LCD panel are chosen such that they compensate for these effects.

The complete rendering algorithm then works as follows (also see Figure 6.4): we take square root of the original HDR image with intensity I (1). The resulting image (2) represents the target intensity \sqrt{I} for the projector. We map these intensities into projector pixel values by applying the inverse of the projector's response function r_1 (3). The projector now produces an image of intensity $r_1(r_1^{-1}(\sqrt{I})) = \sqrt{I}$, except that the image is actually blurred according to p_1 , the projector's point spread function (PSF). To simulate this blurring, we convolve the projector intensities with the PSF (4) and divide the result out from the original HDR image to get the target LCD transparency (5). For the final pixel values of the LCD, we apply the inverse of the panel's response function r_2 (6).

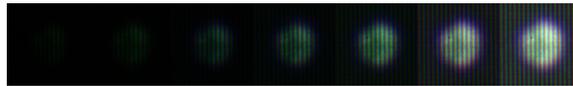


Figure 6.5: Point spread function of the projector.

An exposure sequence of the PSF p_1 is depicted in Figure 6.5. Note the vertical lines visible in images with larger exposure times. These are the RGB subpixels of the LCD panel. In order to speed up the computation, we do not use the measured PSF directly, but fit a tensor-product Gaussian of standard deviation of 2 – 3.5 pixels to p_1 ,

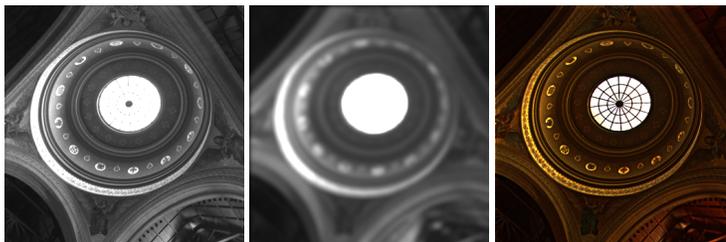


Figure 6.6: Factoring an HDR photograph for projector-based display. Left: square root of the intensity. Center: blurred image which is predicted to be the image generated by the de-focused projector. Right: edge enhanced LCD panel image that corrects for the blurriness of the projector image.

so that we can use a 2D separable filter of width 13 for the convolution.

These image processing steps are possible both completely in software and using recent programmable graphics hardware. The convolution, with the 2D separable approximation, can be readily implemented as a pixel shader on both the latest ATI and NVIDIA chips. We achieve real-time frame rates on graphics hardware with this approach.

Figure 6.6 shows the results of this image factorization on a portion of Debevec's Stanford Memorial Church HDR photograph. On the left is shown a grayscale image that corresponds to the square root of the original intensity values. Convoluting that image with the PSF of the projector yields the center image. This is the predicted image that will be produced by the projector. Finally, the right image is the color LCD panel image that corrects for the blurriness of the projector. It is interesting to note that the LCD panel image is essentially an edge-enhanced image with low frequency components attenuated or removed. This is particularly noticeable for the widths of the window frames. Thus, the algorithm that we apply here is very similar in principle to a local tone mapping operator for the LCD panel image.

6.3 Driving the LED Display

In this section, we briefly describe the rendering algorithm for the LED-based HDR displays. For a detailed analysis of the rendering algorithms for the LED-based dis-

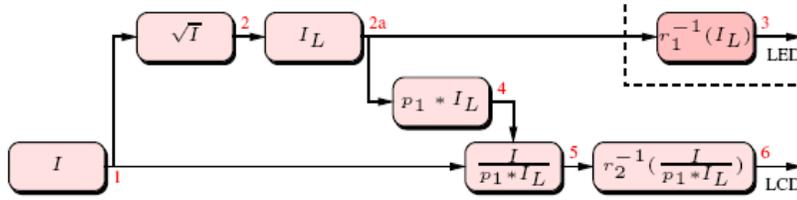


Figure 6.7: Rendering algorithm for the LED-based display.

plays, we refer the reader to Trentacoste’s Master’s thesis [134]. The principal rendering algorithm for the LED-based system is quite similar to that for the projector-based display, as shown in Figure 6.7. The primary difference between the two display systems from a rendering perspective is that the PSF of an LED has a much wider support than the one for a pixel of the projector. Also of importance is the fact that the LEDs are arranged on a hexagonal grid rather than a rectangular grid. These differences have two consequences. Firstly, because of the wider support of the PSF, it is advisable to come up with a better way to choose the LED values. Since the supports of the PSFs for neighboring LEDs overlap, determining the optimal LED value is essentially a de-convolution problem, as explained below. Secondly, because of both the hexagonal geometry and wider support of the PSF, the convolution (4) has to be implemented differently.

The first issue is addressed by adding an additional stage (2a) to derive the target intensities I_L for every individual LED. To this end, the image is first down-sampled to the resolution of the LED array, and then we need to solve for the values, taking the overlap of the PSFs into account. This is essentially a de-convolution problem, the full solution of which would require solving a sparse linear equation system with as many unknowns as there are LEDs. This de-convolution is not an option for interactive applications. We therefore simply choose the intensity of every LED as a simple weighted sum of the neighboring pixels in the down-sampled image.

As before, we rely on the LCD panel to compensate for any differences between the LED values and the target image. To this end, we need to forward-simulate the low-



Figure 6.8: Factorization of an HDR photograph for the LED-based display. Left: LED contribution as a result of convolving LED values with the LED point spread function. Right: edge enhanced LCD panel image that corrects for the blurriness of the LED image.

frequency image (4) generated by the LED panel in order to derive the LCD pixel values. We use a splatting approach to do this on graphics hardware. We simply draw screen aligned quadrilaterals with textures of the PSF into the framebuffer. Alpha blending is then used to accumulate the results.

Figure 6.8 shows the factorization of an HDR image into LED component and LCD panel component. Due to the wider support of the PSF of the LEDs compared to the PSF of a projector pixel in the first setup, the LED image is even more low-pass filtered than before. As a result, the compensation performed in the LCD panel is more significant, and is visible in the right image.

6.4 Applications

We have developed a few simple applications to test our projector-based display algorithm, and to demonstrate its potential in a number of application domains. The first one is a simple HDR image viewer (Figure 6.10, top). It allows the user to load an HDR image and view it while interactively adjusting the exposure settings (i.e., the absolute scale of intensity). This application was also used to generate the images in Figure 6.9. The three images to the right of that figure show a color-coded comparison of an original radiance map, an HDR photograph taken off the projector-based display, and finally a photograph taken off a standard LCD screen. The intensities of the photograph from the HDR display are similar but not identical to the values in the original radiance map.

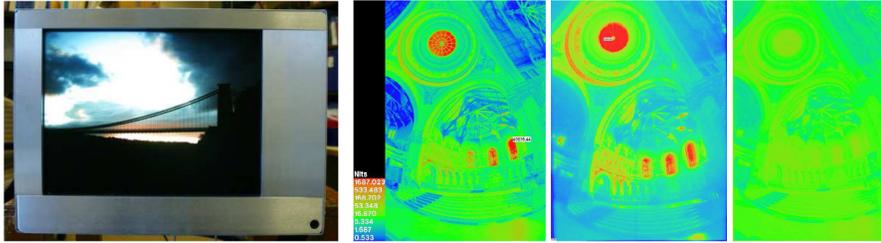


Figure 6.9: From left to right: our projector-based display showing an HDR image; a color-coded original HDR image; HDR photograph taken off the screen of our projector-based system; HDR photograph taken off a conventional monitor displaying the tone-mapped image.

The differences are mostly due to imperfections in both the calibration of the display to absolute intensities, and in the image acquisition process. Clearly both intensity and dynamic range of the HDR display are vastly superior to the standard monitor.

The second application we developed is related to interactive photorealistic rendering. We modified a DirectX application for displaying BRDFs measured with linear light source reflectometry [41], and replaced its tone mapping step with the rendering algorithm for the HDR display (Figures 6.10, center). Other interactive applications that can render into floating point buffers can easily be modified in a similar fashion.

Based on a similar principle, we built a simple volume ray-caster that runs on a GPU (Figure 6.10, bottom). It allows for operations such as rotation, slicing, and adjustments to the transfer function. Both the actual volume rendering algorithm and the processing for the projector-based display is implemented in a single Cg shader [83]. This has been extended to automatic transfer function generation for volume rendering with HDR displays [48].

6.5 Discussion

The dynamic range issue of display devices for perceptually realistic display has recently been addressed by the HDR display devices. In this chapter, we discuss the basis principle of image formation for the HDR displays and describe in detail the

real-time rendering algorithm for the projector-based design. We also develop a few simple applications to test and validate our rendering algorithm. We believe these applications demonstrate the potential application areas for the HDR displays including photo-realistic rendering and scientific visualization.

It should be noted that the rendering algorithm for the LED-based display, as described in Section 6.3, is computationally more demanding than the algorithm for the projector-based display due to the larger support of the PSF. However, the latest generation GPUs have support for floating point blending and bilinear interpolation. Hence, a fullscreen image factorization into LCD and LED components can be carried out in real-time on these chips. The commercial versions of the LED-based HDR displays from Brightside Technologies [132] now have custom firmware on the displays for real-time processing of HDR images, thereby off-loading this computationally intensive task from the graphics cards.

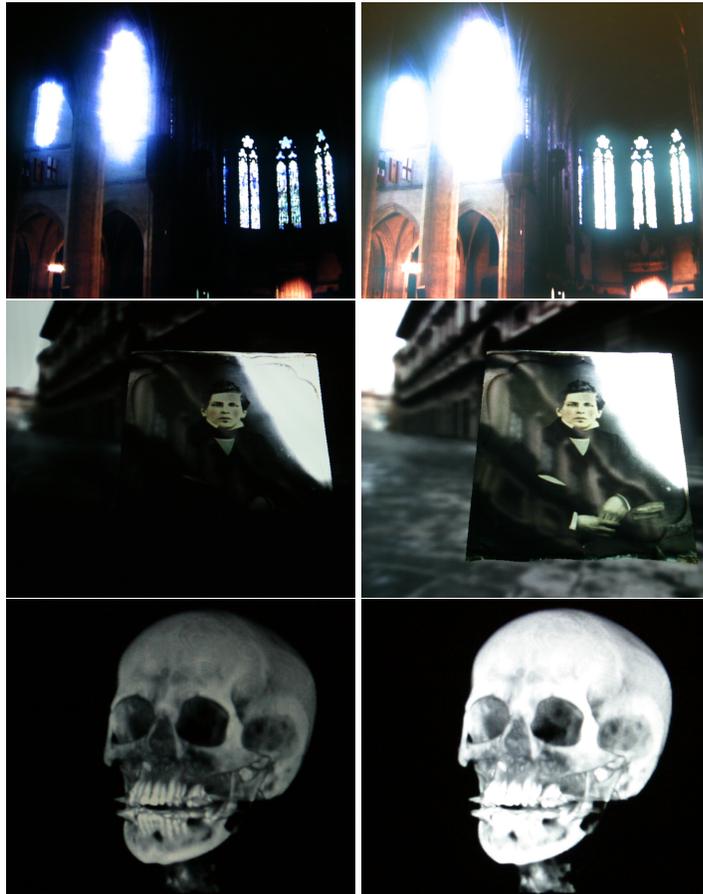


Figure 6.10: Screen photographs of the different applications we implemented. The exposure times of the two images in each pair differ by 4 stops. Top: HDR image viewer. Center: interactive rendering of measured BRDFs. Bottom: volume rendering.

Chapter 7

Real Illumination from Virtual Environments

Creation of a sense of presence and immersion in a virtual environment has been a major theme for graphics research targeting perceptual realism. As discussed in Chapter 6, the high dynamic range (HDR) display technology goes a long way to achieve this goal by matching the intensity and contrast range of many real world scenes. Apart from the dynamic range issue of display devices, another major factor affecting perceptual realism is that the viewing conditions for a virtual scene are largely unknown. This means that parameters such as the viewer's light and color adaptation cannot be considered in the image generation process.

It is our hypothesis that the HDR display behaves like a window into a virtual world. However, this sense of realism and immersion in a virtual environment is lost outside this window as the HDR display cannot account for the viewing conditions. A true sense of immersion can only be achieved if the illumination levels in the real and virtual world are compatible. A night driving simulation, for example, should happen in a darkened room, while the same application in a daylight setting should take place in a bright room. Ideally, it would be possible to adjust the room illumination over time to simulate, for example, the car entering or leaving a tunnel.

In this chapter, we propose to actively control the illumination in a room so that it is consistent with a virtual world (see Figure 7.11) [49]. In combination with an HDR display, the system produces both uniform and directional illumination at intensity levels covering a wide range of real-world environments. It thereby allows natural adaptation processes of the human visual system to take place, for example when moving between

bright and dark environments. In addition to enhancing the sense of immersion in the virtual environment, the directional illumination provides additional information in the user's peripheral field of view. This is the main concept behind our approach.

To evaluate the potential for the proposed method in an entertainment setting, we conducted a survey of user preference. All participants of the survey preferred the system with dynamic, directional illumination over a room of constant brightness. The participants also believed that the additional cues provided by directional illumination helped them keep track of orientation in the virtual world.

In the following sections, we first describe the system, including hardware, calibration, and rendering algorithms (Section 7.1). We then describe several ways to acquire the relevant lighting information for both virtual worlds and film sequences (Section 7.2). Finally, in Section 7.3, we discuss the user survey that we performed to test the concept.

7.1 Method

In our prototype implementation, we use computer-controlled LED lights that are distributed throughout the room. All lights are individually programmable to a 24 bit RGB color. The computer-controlled lights are programmed such that, for a specific real-world viewing position, the room illumination resembles a blurred environment map for the virtual world at the virtual viewing position.

Our method has three major components: the physical setup of the light sources in a room, calibration methods, and rendering algorithms. These aspects are discussed in the following sections.

7.1.1 Physical Setup

We assembled our prototype system in a separate room, approximately 15.5' long, 9' wide, and 9.5' high. The room remained as-is: the walls were kept in the original pastel color, and specular objects such as a whiteboard and the reflectors of the houselights were left in place. A window cover was used to block out daylight (Figure 7.11, left).

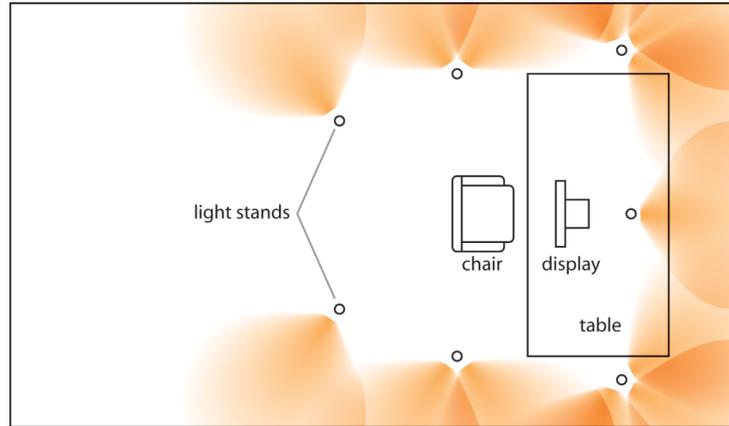


Figure 7.1: Room layout: additional lights are mounted below the ceiling, pointing upwards.

The room contained several pieces of furniture, including two tables and several chairs. One of the tables was located at one end of the room and held the computer console. We used both an 18” LED HDR display [119] and a standard 18” flatpanel (NEC Multisync LCD 1850e) for our experiments.

The lighting system consists of 24 RGB LED lights (ColorKinetics iColor Cove), each of which can be individually programmed to a 24-bit RGB color value. Instead of pointing the light sources directly at the viewer, which would create high intensity illumination from very specific directions, we aimed the lights at the walls in order to diffuse the light output over a large range of directions. This corresponds to our goal of creating a lighting system that can produce a low-frequency version of the illumination in the virtual world.

We used seven poles with stands to mount the 24 light sources. The lights were positioned and oriented such that they predominantly illuminated the ceiling, as well as the walls to the left, right and in front of the viewer (see Figure 7.1). Experiments quickly confirmed our intuition that illumination from behind has only a comparatively small impact, and hence we used only a few light sources for those directions. No illuminators were aimed towards the floor for similar reasons. Our arrangement roughly mimics the change in photoreceptor density in the retina from foveal to peripheral view.

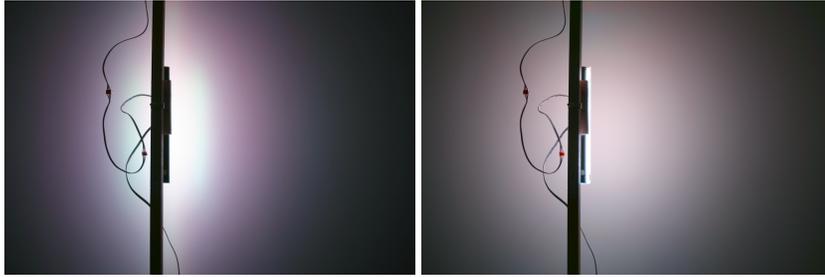


Figure 7.2: Left: light pattern generated by a single iColor Cove light. Note the narrow light spot and color banding. Right: pattern generated using a diffuser.

However, it might be interesting future work to design the physical setup by formally taking into account the resolution of the human eye [22].

To create a smoothly varying illumination pattern we used strong diffusers at the light sources, which also reduce color separation of the RGB elements (Figure 7.2). The diffuser for each light consist of 2" diameter transparent acrylic tubing that was cut in half along its axis, and spray-painted lightly on the outside with white plastic paint (*KrylonTM Fusion* for plastic). To avoid internal reflection losses we used reflective film to coat the inside of the light source.

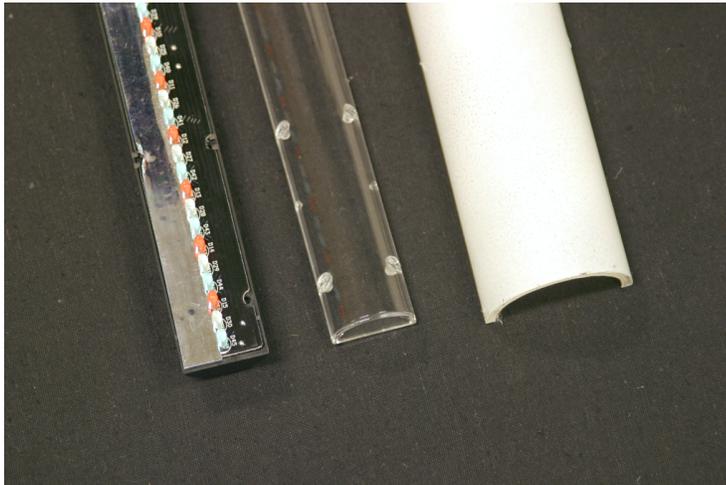


Figure 7.3: Left: opened iColor Cove light source. The left side of the circuit board is already covered with reflective foil. Center: clear cover for the light source. Right: diffuser built from acrylic tubing and white spray-paint.

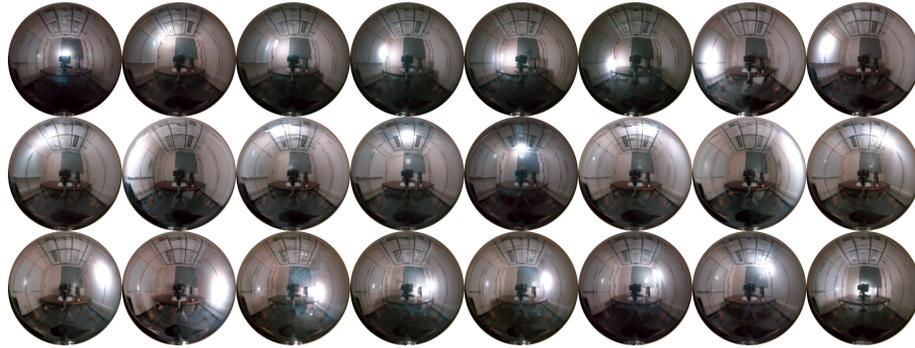


Figure 7.4: Light probe images acquired for each of the 24 light sources at the intended viewing position.

The typical light output of each light is specified as 52.4 lumens for full white. This is not bright enough to match the top intensity of the HDR display, but since every light is set to an average intensity over a moderately large cone of directions, this limited top intensity has not been an issue¹. It should also be noted that brighter LEDs than the ones we use are readily available. We chose the iColor Cove system primarily because it includes off-the-shelf electronics for computer control.

7.1.2 Calibration

Some calibration steps are necessary in order to control the system in a way consistent with the virtual world and the image shown on the display. Geometric calibration is necessary to determine the positions of the light sources relative to the viewer, and their spread, which is modeled as a Gaussian. Photometric calibration is subsequently performed to match white points and illumination levels between the light sources and the display.

Light Position. The rendering algorithm, as described below, requires information about the impact of every individual light on the illumination as seen from the location of the viewer. To obtain this information, we place a reflective ball at the intended

¹Typical HDR scenes have only a few small bright regions (corresponding to windows or skylights) that are to full white. Experiments we conducted with the HDR display indicate that overall light output is typically less than 10% of peak intensity.

viewer position to act as a light probe. We take photographs with a web camera (Creative NX Ultra) while switching on one light at a time. The resulting environment maps appear in Figure 7.4.

We then model the impact of every light source by fitting a Gaussian to the environment map. It is centered around the direction corresponding to the brightest point in the environment map, and its standard deviation is chosen such as to minimize the RMS error. Other directional bases such as cosine lobes could be used instead of Gaussians. However, we found Gaussians to be convenient, since they can capture more distant contributions caused by indirect illumination, while smoothing over high frequency details such as object boundaries.

White Point and Intensity Calibration. An important part of the calibration step is to match the white points of the display and the lighting system. At the same time, we need to establish the relative intensities of light sources and display. For both tasks, we use an 18% gray card commonly used in photography. Under the assumption of uniform hemispherical illumination ($L_i(\omega_i) = \text{const}, \omega_i \in \Omega$), the reflected radiance of a (diffuse!) 18% gray card is

$$L_o(\omega_o) = \int_{\Omega} L_i \cdot \frac{0.18}{\pi} \cos \theta_i d\omega_i = 0.18 \cdot L_i.$$

Since uniform hemispherical illumination can be approximated by setting all our lights to the same intensity, the calibration task is implemented as a uniform adaptation of the intensity of the lights until the gray card reflection matches an 18% monitor gray. Note that the response function of the display needs to be taken into account, i.e., the monitor color is set to 18% of the top *intensity*, not 18% of the top pixel value. The color matching can be automated by using the same web cam as above. We cover one half of the screen with the gray card, and show 18% red on the other half of the display. We then adjust the red intensity of the light sources using binary search until the camera observes the same intensities for both the monitor image and the gray card. The same steps are repeated for the green and blue channels.

From this procedure we recover the relative scaling factors for the light sources

that correspond to the full intensity of the individual color channels on the display. During rendering, these can either be used directly to adjust the light intensities, or their reciprocals can be applied to the image shown on the display. If an absolute white point calibration is desired, the former method should be used, and the monitor should be calibrated with standard tools.

Since the lighting setup is indirect, the color of the walls and other large objects does influence the color temperature of the illumination. If the walls or other large objects in the room show a great variation in color temperature, then the effective contributions of the individual lights have a different color. In that case, the color difference for the individual lights needs to be calibrated first by sequentially switching them on, and comparing the color of the resulting illumination on the grey card. Only then can the intensity be calibrated as described above. Note that the white point and intensity calibration should ideally be repeated every time a light source moves, or even when large, colored objects get moved around in the room. Fortunately, all calibration steps are automated and can be completed within a few minutes.

7.1.3 Rendering

There are two algorithms for driving the calibrated lighting system during rendering. The first option is to uniformly adjust the light source intensity to the average intensity of the scene. The second option is to drive every light source individually by sampling the scene's illumination in the region affected by the light. In both cases, we use an environment map for the location of the viewer as a representation of the scene illumination. This is a convenient choice, since many realtime graphics applications already create those maps for shading objects near the viewer. It also does not require any scene geometry, and could therefore easily be painted by an artist to augment old footage (see Section 7.3). Ideally, the environment map should be in high dynamic range format, but low dynamic range information can also be used, especially if the environment map is split into different layers. This is discussed further in Section 7.2.

In the case of uniform illumination, we simply integrate the intensities in the environment map, and use the resulting value to drive all light sources.

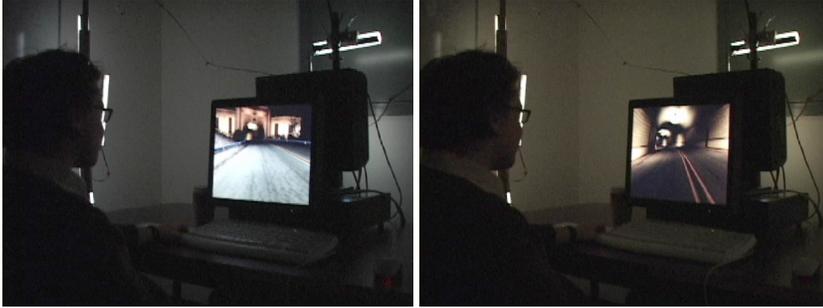


Figure 7.5: Left: room lighting corresponding to a passing street lamp on the left side in the driving game NFS Underground 2. Right: lighting corresponding to overhead lights and lights on the walls inside a tunnel in the game.

In the case of directional illumination, we precompute an importance sampling pattern for the Gaussians that we fit to every light in the calibration process (Section 7.1.2). For every frame, we then sample the illumination from these patterns, and use the resulting integral to drive the light sources. We can also blend uniform and directional illumination to control the degree of directional dependence.

7.2 Content Creation

We can create the environment maps required to control the light sources in a number of ways depending upon the application.

Synthetic Environments: The system is easy to integrate into fully synthetic scenes, such as in computer games or animated films. Many games already generate these environment maps for shading objects in the scene [59]. Future games will likely generate these environment maps in an HDR format. Current games typically use low-dynamic-range representations due to the lack of support for floating point textures and framebuffers in older graphics hardware. However, the environments are often split into multiple layers, corresponding to different parts of the scene. This layered information can be used to reconstruct HDR lighting information.

For our experiments we used footage from Electronic Arts' racing game "Need for Speed Underground 2" (Figure 7.5), which features a particularly wide range of

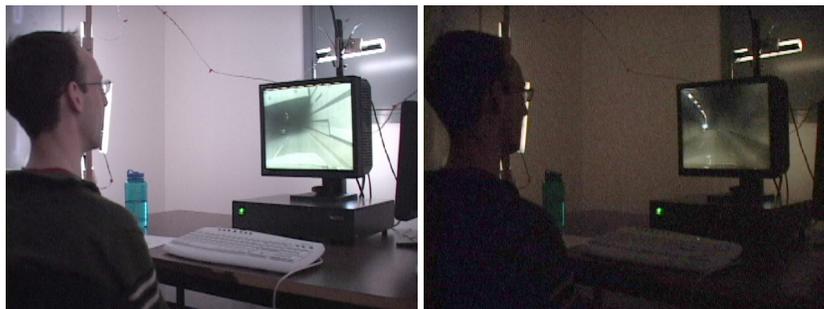


Figure 7.6: Left: uniform lighting corresponding to direct daylight outside a tunnel in the HDR driving video. Right: uniform room lighting corresponding to lighting inside the tunnel.

differently illuminated environments. We used captured environment maps generated by the existing shading system. The layers of the environment map, which correspond to lights, sky, and objects in the scene, were scaled by different factors and added up. The resulting HDR information was used to program the lighting system.

Legacy Film and Video Footage: For some applications it is interesting to retrofit conventionally shot film and video material with environment lighting information. Sometimes it might be possible to automatically extract the required information from the image sequences themselves. For example, Nishino and Nayar’s approach for extracting environment maps from reflections in eyes [95] might be used for sequences in which human or animal faces are visible at a high enough resolution.

In other cases, the required information can be generated manually, since only approximate low-resolution lighting information is necessary. A uniform brightness can be estimated for a set of key frames and interpolated across the sequence. The generation of directional information is more tedious, but as shown by Sloan et al. [123], artists can control lighting in a scene by painting spherical environment maps. This could be done for a set of key frames, and the resulting environments could then be interpolated using a method similar to the one proposed by Cabral et al. [14]. We leave this for future work.

For our experiments we manually created environment map information for HDR video footage of a car ride, and played it back on an 18” LED HDR display [119].



Figure 7.7: Left: directional lighting corresponding to the bright windows in the Grace Cathedral HDR environment. Right: directional lighting corresponding to the altar in the Grace Cathedral.

The video was shot with an HDR camera (HDRC CamCube from IMS Vision), and compressed using the method described by Mantiuk et al. [82]. It features both direct daylight and a dark tunnel (Figure 7.6). We augmented this video with environment lighting information by uniformly adjusting the brightness level for every frame. We mostly used two brightness values, one for the inside of the tunnel and one for the daylight part. At the entrance and exit of the tunnel, we interpolated the uniform intensity linearly over a few seconds.

Filming New Footage: When shooting new films, a light probe can be used to capture the surrounding environment. Unlike in relighting applications [27, 30], the environment maps for our system (Figure 7.7) should ideally be centered at the viewer position, that is, the main camera used for filming the scene. This should make it feasible to record a light-probe video with all additional components outside the field-of-view of that of the main camera. In some cases it might be necessary or desirable to post-process the sequences to account for specific lighting effects.

7.3 Experiments

To test the concept, we conducted a user survey with a set of three experiments. As the evaluation criterion, we chose user preference rather than other possible criteria such as perceived realism. This choice was made due to our primary interest in entertainment

applications for the current system; ultimately the proposed method can be successful for these applications if and only if potential users like the results independent of how realistic they believe those are. This means that other, more formal studies with a different aim are required before the system could be used, for example, for design purposes.

The first experiment focused on the impact of uniform changes in room illumination, while the second one emphasized the directional aspects of our approach. The third experiment was designed to test whether the lighting system could also be useful in combination with conventional low-dynamic-range displays. The participants were 12 graduate and undergraduate students, none of whom work in computer graphics or related areas. All participants had normal or corrected-to-normal vision. The participants entered the room at least 5 minutes before the start of the actual experiments in order to allow them to adapt to a slightly dimmed environment. Questions were asked after individual experiments. After all experiments had been completed, the participants had the opportunity to provide additional comments.

Uniform Illumination. We designed the first experiment to test whether the participants would prefer a dynamic, but directionally uniform illumination level over constant room illumination. To this end, the participants were shown the HDR driving video (see Section 7.2) on the HDR display once in a dark room, and once with a uniform brightness change generated by the lighting system. The participants were then asked to indicate their preference for either the constant or the dynamic illumination on a 5-level scale (strong preference for the dark room, weak preference for the dark room, undecided, weak preference for dynamic lighting, and strong preference for dynamic lighting). The same experiment was repeated for a room illuminated at normal brightness levels. Again, both variants were shown back-to-back, and the participants were asked to state their preference.

The answers given by the participants are summarized in Figure 7.8. All participants preferred or strongly preferred dynamic illumination over a constant brightness level. In general, the preference was even stronger in the comparison with a dark room, although one participant was undecided. This stronger preference can be explained by

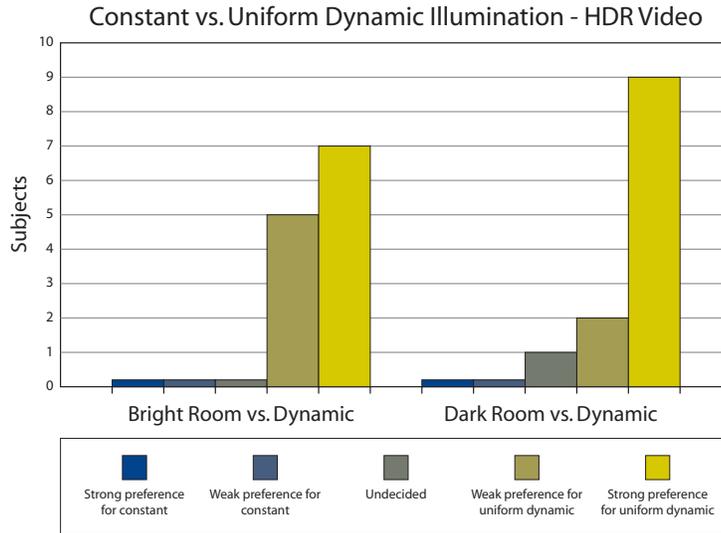


Figure 7.8: User preferences regarding constant or uniform dynamic illumination for HDR video.

the ability of the HDR display to produce light levels that start to be uncomfortable in very dark environments. From this results we can conclude a significant preference of the participants for dynamic illumination compared to a constant light level.

Directional Illumination. The second set of experiments tested whether the users would prefer directionally localized illumination changes over uniform brightness changes. This experiment was based on using a simple viewer similar to Quicktime VR [15]: the program loads an HDR panorama, shows it on the HDR display, and lets the user look around with a simple mouse interface for rotation (see Figure 7.7). Note that a ‘dynamic’ illumination approach with uniform intensities for all lights would result in a constant illumination pattern for this application, since total brightness would not change under rotations of the viewing direction.

The participant was then asked to use the rotation interface to locate the brightest point in the panorama. The application was run 20 times, and each time a different panorama was selected at random. The application also randomly decided whether or not to use the lighting system to create directional information (if not, the lights

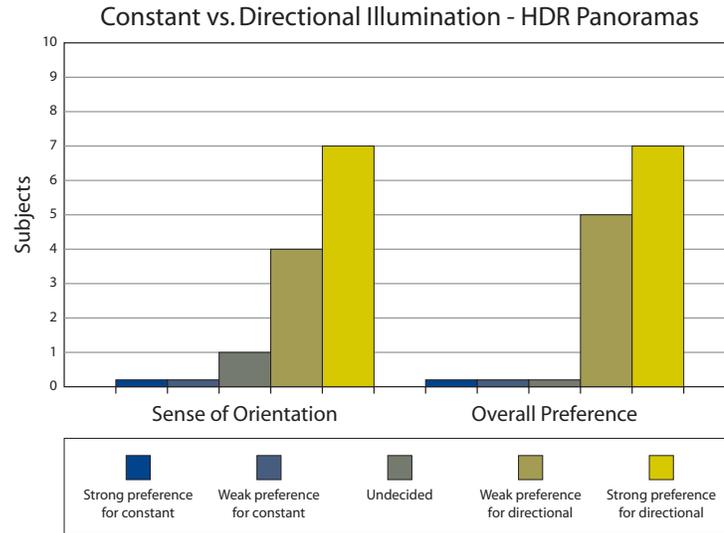


Figure 7.9: User preferences for directional vs. uniform illumination in an HDR panorama viewer.

were switched to a medium intensity). After 20 runs, the participants were asked two questions: first, whether they felt that the directional illumination helped them with orientation (on a scale from strongly disagree to strongly agree), and second whether they preferred the directional or the uniform illumination.

Figure 7.9 shows the answers for both questions. All participants preferred or strongly preferred the directional illumination over the uniform one. With one exception, all participants answered the question regarding an improved sense of orientation identically to the way they answered the question for overall preference. One participant was uncertain whether the directional lighting had improved his sense of orientation, but nonetheless preferred directional over uniform lighting.

From the results of this experiment, it is clear that dynamic directional illumination is preferable over both constant and uniform adaptive lighting (note again, that the latter would have produced constant illumination as well in this application).

Low-Dynamic-Range Footage. Finally, we also wanted to determine whether the lighting system is useful in combination with conventional low-dynamic-range dis-

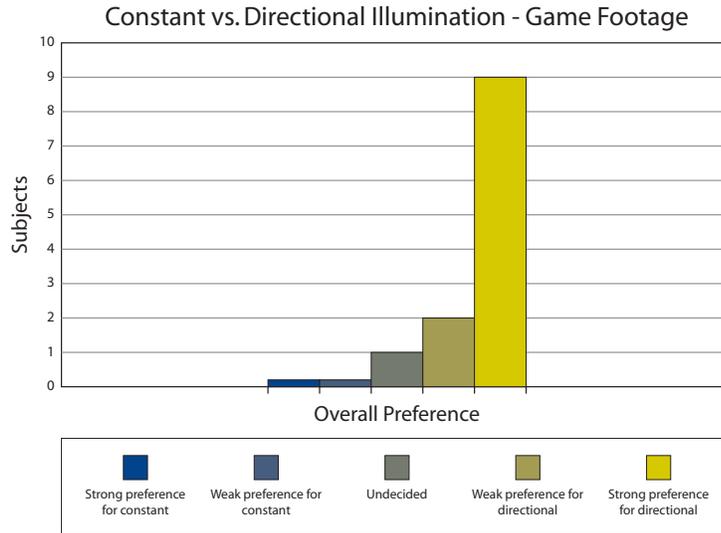


Figure 7.10: User preferences for directional vs. uniform illumination when watching low-dynamic-range video game footage.

plays. To this end, the participants were shown the footage from “Need for Speed Underground 2” (see Section 7.2) on a conventional display. The segment contained a tunnel sequence, in which widely spaced street lights caused the scene to get brighter and darker at regular intervals. First, we showed the segment under constant illumination, and then with the dynamic, directional illumination, and the participants were asked for their preference.

In this experiment, the preference for directional dynamic illumination was very strong, as indicated in Figure 7.9. One participant was undecided, as also revealed in his written comment: “[I] did not like the flickering lights [when passing by the street lights in Need for Speed] – very realistic, but very annoying (it’s distracting enough when you’re driving in real life). [It is] more annoying in real life. [I] really enjoyed the feeling of motion”. This subject did not express similar concerns in the directional experiment using the HDR display, where he strongly preferred the directional illumination. We believe that this ambivalence is in some sense caused by the lighting system overpowering the conventional display, which cannot produce the same intensities as the HDR display.

Based on the overwhelmingly positive response of the other participants, we do believe that the system has potential even in a low-dynamic-range setting. However, more studies are required to determine how the lights should be controlled in this scenario to avoid irritations for some users.

Other comments from participants not attributed to a specific experiment included the following: “The dynamic lighting immerses you in the experience”, “[Dynamic lighting] makes you really feel like you’re there, especially in light-dark transitions”, and “One day, this will be in the movies!”

7.4 Discussion

In this chapter, we have introduced an approach for actively controlling the lighting in a room to match illumination in a virtual world. In doing so, we are able to reproduce illumination levels similar to the ones we experience every day in the real world. This triggers natural adaptation processes in the human visual system, for example, when moving between bright and dark environments. In addition, we can generate directional illumination patterns, such as light that appears to come from the sides or from behind. We believe that when used in conjunction with an HDR display, such a system greatly enhances the perceptual realism of a virtual scene.

Our user survey shows overwhelming support for this concept in combination with an HDR display: *all* of our participants preferred the lighting system over constant room illumination. We believe that this combination of HDR display and lighting system comprises the best setup, since it makes it possible to create similar brightnesses both on the display and in the surrounding room.

Even in combination with a conventional low-dynamic-range display, the participants were predominantly positive about the lighting system. With the display dimmer than the light sources, one subject was irritated by the dynamic illumination, although he strongly preferred the system in combination with an HDR display. This indicates that, while the lighting system is promising even in combination with conventional displays, the algorithms for driving the system in such a setting require more research.

One solution could be the use of a tone mapping operator for the light sources.

On the hardware side, several variants of the current system are possible. At the moment, we use 24 individually packaged LED light sources. One could easily imagine repackaging those lights into a single housing with multiple independently controllable spotlights, to be hung from the ceiling. Such a package could also contain light sensors for the calibration, eliminating the need for an external camera. Since LED lights are becoming more and more popular as standard room illumination, one could also imagine directly plugging into an existing home automation system to control those light sources. We envision that such systems could be used in home theaters and gaming environments, while screening rooms and higher end systems would use dedicated light sources such as in our system.



Figure 7.11: Top: photograph of the room housing our system, with all lights switched on. Center: illumination programmed to resemble the Grace Cathedral environment. Bottom: A user viewing the Grace Cathedral environment on an HDR display in a room illuminated by our system. Note how the room illumination is consistent with the virtual environment shown on the screen.

Chapter 8

Conclusions

In this dissertation, we focus on the domain of realistic image synthesis in computer graphics while targeting both photo-realistic as well as perceptually realistic rendering. We identify the principal stages of the image synthesis pipeline - acquisition, rendering and display, and have developed novel techniques and algorithms addressing various problems at every stage of the pipeline.

An interesting observation is that all techniques discussed have one thing in common: they are specially designed for dealing with high dynamic range (HDR) data. In Chapter 3, the acquisition of BRDF response to basis functions relies on HDR imaging. HDR environment maps are used as representations of realistic direct illumination in Chapters 4 and 5. The Monte Carlo sampling techniques developed in these chapters are specially designed for high frequencies in lighting and BRDFs as captured by HDR imaging. Chapter 6 discusses how to process HDR images, either produced by rendering algorithms or directly acquired from the real world, in real-time for driving the HDR displays while Chapter 7 describes a technique for extending the displayed dynamic range on to the real viewing environment for a sense of immersion in a virtual environment.

Also related is the dissertation's focus on realistic rendering with complex direct and directional illumination effects from static as well as dynamic HDR environment maps. This is in accordance with the increasing popularity of image based rendering techniques for realistic image synthesis, particularly for representing complex real-world illumination that is difficult to model in other forms. Here, we focus on efficiently computing the local lighting in a virtual (or real) environment for both real-time rendering applications (Chapters 3 and 7) and high quality off-line rendering with

ray-tracing (Chapters 3, 4 and 5).

Chapter 3 describes a novel technique for the acquisition of surface reflectance properties of every-day materials. Unlike existing work on BRDF acquisition that measure impulse responses to point lights approximating Dirac peaks, the distinguishing characteristic of our measurement system is that it captures the response of the surface to illumination in the form of smooth basis functions. This speeds up acquisition time by an order of magnitude from a few hours to just a *few minutes*. The required optical setup for our approach has the side benefit of being compact and not having any moving parts, hence making it portable.

Our optical setup, *in principle*, enables sampling over 98% of the hemisphere with high sampling density even close to grazing angles. However, our choice of manufacturing process with electroforming for the required free-form surface of the dome limited our measurements to a smaller subset of directions. It would be interesting to investigate other moderate cost manufacturing processes for such concave holes. Another design option would be to optimize over the hemispherical coverage for a concave hemispherical dome that can be electroformed easily.

It would also be interesting to look at alternative basis functions for acquisition of specular materials. A hierarchical wavelet-based acquisition would be suitable for this purpose in principle, as wavelets localize well in both spatial and frequency domain. However, the wavelet bases would not be very suitable for extrapolation of the data in the zone of missing measurements due to their localized support. While we can take advantage of reciprocity to fill in a subset of the missing values, the data extrapolation problem is particularly tricky when both the incident light direction as well as the viewing direction are outside the measurement zone. Other possible approaches can include experimenting with non-linear non-orthogonal basis functions such as a mixture of Phong lobes [40, 71] or Gaussians [130], although extrapolation of the reflectance function in the zone of missing measurements with such bases would also be non-trivial. Another interesting approach would be to optically fit the BRDF to an analytical model as a basis function. An advantage of such an approach over using a linear zonal basis function is that the choice of basis is data-driven, thus making the encoding efficient. Most BRDF models, however, are not easily separated into a structured illu-

mination basis and a reflected light basis, which presents a challenge for deriving the appropriate basis illumination. However, we believe that this may be an excellent area for future research.

High quality off-line rendering with realistic materials and complex direct illumination has major applications in realistic relighting for motion pictures and architectural design. Hence, as part of this dissertation, we introduce several state-of-the-art Monte Carlo sampling techniques from statistics and machine learning literature to efficiently compute the direct illumination integral. Note that, in general, the proper simulation of global illumination is very important for realistic rendering. However, we focus on direct illumination in this thesis because it is a required component of all global illumination renderers, and direct illumination computation can be computationally more expensive than global illumination in the presence of complex illumination models. In Chapter 4, we demonstrate how sampling for Monte Carlo ray-tracing only from the distribution of illumination or only the BRDF can be sub-optimal in scenarios where both are high frequency functions. We propose a bidirectional importance sampling approach for sampling from the product distribution of the illumination and BRDF, which is the target distribution for sampling in the absence of visibility information. We describe two alternative Monte Carlo strategies for bidirectional sampling, one based on rejection sampling and the another based on sampling-importance resampling.

One downside of both these strategies is that sample selection process is more expensive compared to sampling from individual distributions, as well as compared to another simultaneously developed approach to product distribution sampling based on wavelets [16]. However, sampling from the product distribution drastically reduces the number of visibility tests required to obtain good image quality compared to importance sampling from a single distribution. And unlike wavelet importance sampling, our bidirectional sampling strategy requires no pre-computation of the environment map making it suitable for rendering a video sequence with high resolution environment maps, and scenes with spatially-varying BRDFs.

In Section 4.3, we extend bidirectional sampling to also efficiently account for the visibility function. In this case, we employ the Metropolis-Hastings algorithm for establishing a correlation in the energy estimates of neighboring pixels on the image-

plane, thereby reducing noise in partially occluded regions. To our knowledge, this is the first proposed solution for sampling the direct illumination integral according to the triple product distribution that does not require any precomputation of visibility. Similar to the work by Cline et al. [18], we demonstrate how Metropolis sampling after an initial phase of Monte Carlo sampling is an effective technique for variance reduction. Such a hybrid algorithm also overcomes the startup-bias of Metropolis sampling [137], and hence demonstrates fast convergence.

We extend product distribution sampling in the temporal domain in Chapter 5 with our proposed sequential Monte Carlo mechanism for sampling dynamic illumination. By propagating samples over time, the method makes efficient use of coherence across frames. The method is very general, requires no precomputation, and the sample propagation process is much more efficient than selecting samples from scratch every frame using Monte Carlo sampling. We believe that such SMC methods are also promising for solving other sampling problems in computer graphics, for example for global illumination with path tracing or photon mapping, or for rendering motion blur effects.

The dissertation also contributes to the display end of the image-synthesis pipeline where the goal is to translate the photo-realistic reproduction of a scene by the acquisition and rendering stages of the pipeline into a perceptually realistic reproduction on a display device. Perceptually realistic rendering is mainly concerned with recreating the *visceral experience* of a real scene for a human observer and a major factor affecting this visceral experience is the dynamic range of display technology. Here, we develop an algorithm for processing HDR images in real-time for driving the projector-based HDR displays [119], and demonstrate its application in several domains including photo-realistic rendering and scientific visualization. The real-time rendering algorithm for driving the commercially developed LED-based HDR displays [132, 134] is very similar in principle, albeit more computationally intensive.

Apart from the dynamic range issue of display devices, another major factor affecting perceptual realism is the unknown viewing conditions for a virtual scene. Hence, in Chapter 7, we propose to actively control the illumination in a room so that it is consistent with a virtual world. This triggers natural adaptation responses in the human visual system, thereby enhancing the sense of realism and immersion in a virtual

environment. The method provides directional control over the room illumination and hence goes further than related commercially available systems [104]. We have conducted an informal user preference survey in order to verify this concept. Our participants overwhelmingly supported this concept, especially in combination with an HDR display.

We believe that the work presented here also creates a variety of promising directions for future research. One important area is artistic tools for content creation, in particular for augmenting existing film material with information about directional illumination for driving the room lights. We have described some ideas on this topic in Section 7.2, but more sophisticated approaches such as optical flow and particle filtering should be possible. At the moment, we focus on entertainment-style applications, where user preference is arguably all that matters. An interesting topic for future work is to analyze whether the system can also be helpful in task-oriented applications, for example ones that require navigation in space. Our user survey indicates that this might be possible since the dynamic illumination can help with orientation. However, more studies are required to fully assess the potential of the proposed method in such applications. This dissertation also does not explicitly look into the role played by color in the perceptually realistic representation of a scene. Thus, there is scope for more formal investigation into the relative importance of luminance and color for visual and perceptual realism, especially in high dynamic range settings.

To conclude, this dissertation takes both the aspects of realism, photo-realism and perceptual realism, into account while targeting real-time rendering applications as well as high-quality offline renderings with realistic materials and illumination environments. Most of the techniques developed here are very general and should be applicable to other problems in computer graphics. For example, the Monte Carlo sampling techniques developed in this dissertation can be readily applied to global illumination algorithms. Similarly, the concept of acquisition of material reflectance with basis functions can potentially be extended to acquisition of light fields and reflectance fields. Researchers are already experimenting with optical setups for such applications [102].

Bibliography

- [1] S. Achutha. BRDF acquisition with basis illumination. Master's thesis, The University of British Columbia, 2006.
- [2] S. Agarwal, R. Ramamoorthi, S. Belongie, and H. W. Jensen. Structured importance sampling of environment maps. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 22(3):605–612, July 2003.
- [3] J.V.C. Antwerp. Automatic brightness control apparatus. *United States Patent 4,514,727*, April 1985.
- [4] A. Artusi, J. Bittner, M. Wimmer, and A. Wilkie. Delivering interactivity to complex tone mapping operators. In *Eurographics Symposium on Rendering*, pages 38–44, 2003.
- [5] M. Ashikhmin. Distribution-based BRDFs, 2006. <http://jesper.kalliope.org/blog/library/dbrdfs.pdf>.
- [6] M. Ashikhmin and P. Shirley. An anisotropic phong BRDF model. *J. Graph. Tools*, 5(2):25–32, 2000.
- [7] M. Ashikhmin, S. Premoše, and P. Shirley. A microfacet-based BRDF generator. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 65–74, 2000.
- [8] B. Barbier, S. Ediar, and J. Brun. System for the display of luminous data with improved readability. *United States Patent 5,057,744*, October 1991.
- [9] C. J. Bartelson and E. J. Breneman. Brightness perception in complex fields. *Journal of the Optical Society of America*, 57(7):953–957, March 1967.

- [10] J. F. Blinn. Models of light reflection for computer synthesized pictures. In *Computer Graphics (Proc. of ACM SIGGRAPH '77)*, pages 192–198, 1977.
- [11] D. Blythe. The Direct3D 10 system. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 25(3):724–734, 2006.
- [12] D. Burke. Bidirectional importance sampling for illumination from environment maps. Master's thesis, The University of British Columbia, 2004.
- [13] D. Burke, A. Ghosh, and W. Heidrich. Bidirectional importance sampling for direct illumination. In *Eurographics Symposium on Rendering*, pages 147–156, June 2005.
- [14] B. Cabral, M. Olano, and P. Nemeč. Reflection space image based rendering. In *Proc. of ACM SIGGRAPH '99*, pages 165–170, 1999.
- [15] S. E. Chen. Quicktime VR – an image-based approach to virtual environment navigation. In *Proc. of ACM SIGGRAPH '95*, pages 29–38, 1995.
- [16] P. Clarberg, W. Jarosz, T. Akenine-Möller, and H. Wann Jensen. Wavelet importance sampling: Efficiently evaluating products of complex functions. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 24(3):1166–1175, August 2005.
- [17] D. Cline, P.K. Egbert, J.F. Talbot, and D.L. Cardon. Two stage importance sampling of direct lighting. In *Proc. of Eurographics Symposium on Rendering*, pages 103–113, June 2006.
- [18] D. Cline, J. Talbot, and P.K. Egbert. Energy redistribution path tracing. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 24(3):1186–1195, August 2005.
- [19] J. Cohen and P. Debevec. Light-gen. HDRshop plugin, 2001. <http://www.ict.usc.edu/~jcohen/lightgen/lightgen.html>.
- [20] R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. *ACM Transactions on Graphics*, 1(1):7–24, 1982.

-
- [21] Cornell. CORNELL light measurement laboratory, 2005. <http://www.graphics.cornell.edu/research/measure/>.
- [22] C. A. Crucio, K. R. Sloan, R. E. Kalina, and A. E. Hendrickson. Human photoreceptor topography. *J. of Comparative Neurology*, 292:497–523, 1990.
- [23] C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *Proc. of ACM SIGGRAPH '93*, pages 135–142, 1993.
- [24] CURET. CURET: Columbia-Utrecht Reflectance and Texture, 1999. <http://www.cs.columbia.edu/CAVE/curet/>.
- [25] K. Dana. BRDF/BTF measurement device. In *ICCV*, pages 460–466, 2001.
- [26] K.J. Dana, B. van Ginneken, S.K. Nayar, and J.J. Koenderink. Reflectance and texture of real world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, 1999.
- [27] P. Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proc. of ACM SIGGRAPH '98*, pages 189–198, July 1998.
- [28] P. Debevec. A median cut algorithm for light probe sampling, 2005. SIGGRAPH 2005 Poster.
- [29] P. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Proc. of ACM SIGGRAPH '97*, pages 369–378, 1997.
- [30] P. Debevec, A. Wenger, C. Tchou, A. Gardner, J. Waese, and T. Hawkins. A lighting reproduction approach to live-action compositing. *ACM Transactions on Graphics (Proc. of SIGGRAPH '02)*, 21(3):547–556, 2002.
- [31] P. del Moral, A. Doucet, and A. Jasra. Sequential monte carlo samplers. *J. Royal Statist. Soc. B*, 68(3):1–26, 2006.

- [32] L. E. DeMarsh. Optimum telecine transfer characteristics. *J. of the SMPTE*, 81(10):784–787, October 1972.
- [33] A. Doucet, N. de Freitas, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, 2001.
- [34] F. Durand and J. Dorsey. Interactive tone mapping. In *Eurographics Workshop on Rendering*, pages 219–230, 2000.
- [35] F. Durand and J. Dorsey. Fast bilateral filtering for the display of high-dynamic-range images. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2002)*, 21(3):257–266, 2002.
- [36] P. Dutré, P. Bekaert, and K. Bala. *Advanced Global Illumination*. AK Peters, August 2003.
- [37] P. Dutré, H. W. Jensen, J. Arvo, K. Bala, P. Bekaert, S. Marschner, and M. Pharr. *SIGGRAPH Course Notes: State of the Art in Monte Carlo Global Illumination*. ACM SIGGRAPH, 2004.
- [38] S. Fan, S. Chenney, and Y. Lai. Metropolis photon sampling with optional user guidance. In *Eurographics Symposium on Rendering*, pages 127–138, 2005.
- [39] J. A. Ferwerda, S. N. Patanaik, P. Shirley, and D. P. Greenberg. A model of visual adaptation for realistic image synthesis. In *Proc. of ACM SIGGRAPH '96*, pages 249–258, 1996.
- [40] A. Fournier. Filtering normal maps and creating multiple surfaces. Technical Report TR-92-41, The University of British Columbia, 1992.
- [41] A. Gardner, C. Tchou, T. Hawkins, and P. Debevec. Linear light source reflectometry. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2003)*, 22(3):749–758, 2003.
- [42] P. Gautron, J. Křivánek, S.N. Pattanaik, and K. Bouatouch. A novel hemispherical basis for accurate and efficient rendering. In *Eurographics Symposium on Rendering*, pages 321–330, June 2004.

-
- [43] A. Gelman, J. Carlin, H. Stern, and D. Rubin. *Bayesian data analysis*. Chapman and Hall, London, 1995.
- [44] A. Gelman and X. L. Meng. Simulating normalization constants: From importance sampling to bridge sampling to path sampling. *Statist. Sci.*, 13:163–185, 1998.
- [45] A. Ghosh, S. Achutha, W. Heidrich, and M. O’Toole. BRDF acquisition with basis illumination. Technical Report TR-2007-10, Department of Computer Science, The University of British Columbia, April 2007.
- [46] A. Ghosh, A. Doucet, and W. Heidrich. Sequential sampling of dynamic environment map illumination. In *Proc. of Eurographics Symposium on Rendering*, pages 115–126, June 2006.
- [47] A. Ghosh and W. Heidrich. Correlated visibility sampling for direct illumination. *The Visual Computer (Proc. Pacific Graphics)*, 22(9–11):693–701, September 2006.
- [48] A. Ghosh, M. Trentacoste, and W. Heidrich. Volume rendering for high dynamic range displays. In *Volume Graphics 2005*, pages 91–98, 2005.
- [49] A. Ghosh, M. Trentacoste, H. Seetzen, and W. Heidrich. Real illumination from virtual environments. In *Proc. of Eurographics Symposium on Rendering*, pages 243–252, June 2005.
- [50] A. S. Glassner. *Principles of Digital Image Synthesis*. Morgan Kaufmann Publishers Inc., 1994.
- [51] N. Greene. Environment mapping and other applications of world projections. *IEEE CG&A*, 6(11):21–29, 1986.
- [52] J. Hammersley and D. Handscomb. *Monte Carlo Methods*. Methuen, London, 1964.
- [53] J. Y. Han and K. Perlin. Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Transactions on Graphics*, 22(3):741–748, 2003.

- [54] V. Havran, C. Domez, K. Myszkowski, and H. Seidel. An efficient spatio-temporal architecture for animation rendering. In *Proc. of Eurographics Symposium on Rendering*, pages 106–117, June 2003.
- [55] V. Havran, K. Dmitriev, and H. Seidel. Goenometric diagram mapping of hemisphere. In *Short Presentations (Eurographics 2003)*, 2003.
- [56] V. Havran, M. Smyk, G. Krawczyk, K. Myszkowski, and H. Seidel. Interactive system for dynamic scene lighting using captured video environment maps. In *Proc. of Eurographics Symposium on Rendering*, pages 31–42, June 2005.
- [57] X. D. He, P. O. Heynen, R. L. Phillips, K. E. Torrance, D. H. Salesin, and D. P. Greenberg. 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, 1992.
- [58] X. D. He, K. E. Torrance, F. X. Sillion, and . P. Greenberg. 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, 1991.
- [59] W. Heidrich and H.-P. Seidel. Realistic, hardware-accelerated shading and lighting. In *Proc. of ACM SIGGRAPH '99*, pages 171–178, August 1999.
- [60] P. Irawan, J. A. Ferwerda, and S. R. Marschner. Perceptually based tone mapping of high dynamic range image streams. In *Proc. of Eurographics Symposium on Rendering*, pages 231–242, 2005.
- [61] J. T. Kajiya. The rendering equation. In *SIGGRAPH '86: Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, pages 143–150, 1986.
- [62] M. Kalos and P. Whitlock. *Monte Carlo Methods*. John Wiley & Sons, New York, 1986.

-
- [63] S. B. Kang, M. Uyttendaele, S. Winder, and R. Szeliski. High dynamic range video. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2003)*, 22(3):319–325, 2003.
- [64] J. Kautz and M. McCool. Approximation of glossy reflection with prefiltered environment maps. In *Proc. of Graphics Interface*, pages 119–126, May 2000.
- [65] J. Kautz, P.-P. Sloan, and J. Snyder. Fast arbitrary BRDF shading for low-frequency lighting using spherical harmonics. In *Eurographics Workshop on Rendering*, pages 291–296, 2002.
- [66] J. Kautz, P.-P. Vázquez, W. Heidrich, and H.-P. Seidel. Unified approach to prefiltered environment maps. In *Eurographics Workshop on Rendering*, pages 185–196, 2000.
- [67] J. Kirk and J. Arvo. Unbiased sampling techniques for image synthesis. In *Proc. of ACM SIGGRAPH '91*, pages 153–156, July 1991.
- [68] J.J. Koenderink, A.J. van Doorn, and M. Stavridi. Bidirectional reflection distribution function expressed in terms of surface scattering modes. In *ECCV '96. 4th European Conference on Computer Vision*, volume 2, pages 28–39, 1996.
- [69] T. Kollig and A. Keller. Efficient illumination by high dynamic range images. In *Eurographics Symposium on Rendering*, pages 45–51, June 2003.
- [70] S. Kuthirummal and S. K. Nayar. Multiview Radial Catadioptric Imaging for Scene Capture. *ACM Trans. on Graphics (also Proc. of ACM SIGGRAPH)*, 25(3):916–923, Jul 2006.
- [71] E. Lafortune, S.-C. Foo, K. Torrance, and D. Greenberg. Non-linear approximation of reflectance functions. In *Proc. of ACM SIGGRAPH '97*, pages 117–126, August 1997.
- [72] P. Lalonde and A. Fournier. A wavelet representation of reflectance functions. *IEEE Trans. on Visualization and Computer Graphics*, 3(4):329–336, 1997.

- [73] G. Ward Larson, H. Rushmeier, and C. Piatko. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Trans. on Visualization and Computer Graphics*, 3(4):291–306, 1997.
- [74] J. Lawrence, S. Rusinkiewicz, and R. Ramamoorthi. Efficient BRDF importance sampling using a factored representation. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 23(3):496–505, August 2004.
- [75] J. Lawrence, S. Rusinkiewicz, and R. Ramamoorthi. Adaptive numerical cumulative distribution functions for efficient importance sampling. In *Eurographics Symposium on Rendering*, pages 11–20, June 2005.
- [76] P. Ledda, A. Chalmers, and H. Seetzen. A psychophysical validation of tone mapping operators using a high dynamic range display. In *Symposium on Applied Perception in Graphics and Visualization*, 2004.
- [77] H. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel. Image-based reconstruction of spatially varying materials. In *Eurographics Workshop on Rendering*, pages 104–115, 2001.
- [78] J. Liu and R. Chen. Sequential monte carlo for dynamic systems. *J. Amer. Statist. Assoc.*, 93:1032–1044, 1998.
- [79] S. Lloyd. An optimization approach to relaxation labelling algorithms. *Image and Vision Computing*, 1(2):85–91, 1983.
- [80] T. Malzbender, D. Gelb, and H. Wolters. Polynomial texture maps. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 519–528, 2001.
- [81] S. Mann and R.W. Picard. Being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures. Technical Report 323, M.I.T. Media Lab Perceptual Computing Section, 1994. Also appears, IS&T's 48th annual conference, Cambridge, MA, May 1995.

-
- [82] R. Mantiuk, G. Krawczyk, K. Myszkowski, and H.-P. Seidel. Perception-motivated high dynamic range video encoding. *ACM Transactions on Graphics (Proc. SIGGRAPH '04)*, pages 733–741, 2004.
- [83] W. R. Mark, R. S. Glanville, K. Akeley, and M. J. Kilgard. Cg: A system for programming graphics hardware in a C-like language. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 22(3):896–907, 2003.
- [84] S. Marschner, S. Westin, E. Lafortune, and K. Torrance. Image-based measurement of the bidirectional reflection distribution function. *Applied Optics*, 39(16):2592–2600, 2000.
- [85] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 22(3):759–769, 2003.
- [86] M. D. McCool and P. K. Harwood. Probability trees. In *Proc. Graphics Interface*, pages 37–46, 1997.
- [87] A. Méndez-Feliu, M. Sbert, and L. Szirmay-Kalos. Reusing frames in camera animation. *Journal of WSCG*, 14, 2006.
- [88] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [89] T. Mitsunaga and S. K. Nayar. Radiometric self calibration. In *Proc. of IEEE CVPR*, pages 472–479, 1999.
- [90] P. Moon and D. Spencer. The visual effect of non-uniform surrounds. *Journal of the Optical Society of America*, 35(3):233–248, 1945.
- [91] S. K. Nayar, P. N. Belhumeur, and T. E. Boult. Lighting sensitive display. *ACM Transactions on Graphics*, 23(4):963–979, 2004.

- [92] R. Ng, R. Ramamoorthi, and P. Hanrahan. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 22(3):376–381, July 2003.
- [93] A. Ngan, F. Durand, and W. Matusik. Experimental analysis of BRDF models. In *Proceedings of the Eurographics Symposium on Rendering*, pages 117–226, 2005.
- [94] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Geometric considerations and nomenclature for reflectance. *NBS Monograph*, 160, 1977.
- [95] K. Nishino and S. K. Nayar. Eyes fo relighting. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2004)*, 23(3):704–711, 2004.
- [96] NIST. NIST reference reflectometer: STARR facility, 2003. <http://www.physics.nist.gov/Divisions/Div844/facilities/brdf/starr.html>.
- [97] M. Oren and S. K. Nayar. Generalization of lambert’s reflectance model. In *Computer Graphics (Proc. ACM SIGGRAPH '94)*, pages 239–246, 1994.
- [98] V. Ostromoukhov, C. Donohue, and P.-M. Jodoin. Fast hierarchical importance sampling with blue noise properties. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 23(3):488–495, August 2004.
- [99] S. N. Pattanaik, J. A. Ferwerda, M. D. Fairchild, and D. P. Greenberg. A multi-scale model of adaptation and spatial vision for realistic image display. In *Proc. of ACM SIGGRAPH '98*, pages 287–298, 1998.
- [100] S. N. Pattanaik, J. Tumblin, H. Yee, and D. P. Greenberg. Time-dependent visual adaptation for fast realistic display. In *Proc. of ACM SIGGRAPH 2000*, pages 47–54, 2000.
- [101] P. Peers and P. Dutré. Inferring reflectance functions from wavelet noise. In *Proc. Eurographics Symposium on Rendering*, pages 173–182, 2005.

-
- [102] P. Peers, T. Hawkins, and P. Debevec. A reflective light stage. Technical Report ICT-TR-04.2006, USC Institute for Creative Technologies, 2006.
- [103] M. Pharr and G. Humphreys. *Physically Based Rendering*. Morgan Kaufmann, New York, 2004.
- [104] Philips. Ambient light technology. <http://www.flattv.philips.com/>.
- [105] B. Phong. Illumination for computer generated images. *Communications of the ACM*, 18(6):311–317, June 1975.
- [106] P. Poulin and A. Fournier. A model for anisotropic reflection. In *Computer Graphics (Proc. of ACM SIGGRAPH '90)*, volume 24, pages 273–282, 1990.
- [107] W. Press, B. Flannery, S. Teukolsky, and W. Vetterling. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, 1992.
- [108] J.S. Prince, H.L. Herz, and E.P. Coleman. Brightness control system for crt video display. *United States Patent 4,589,022*, May 1986.
- [109] R. Ramamoorthi and P. Hanrahan. An efficient representation for irradiance environment maps. In *Proc. of ACM SIGGRAPH '01*, pages 497–500, 2001.
- [110] R. Ramamoorthi and P. Hanrahan. Frequency space environment map rendering. In *Proc. of ACM SIGGRAPH '02*, pages 517–526, 2002.
- [111] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic tone reproduction for digital images. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2002)*, 21(3):267–276, 2002.
- [112] M. Robertson, S. Borman, and R. Stevenson. Dynamic range improvements through multiple exposures. In *Proc. of International Conference on Image Processing (ICIP) '99*, pages 159–163, 1999.
- [113] M. Sbert, L. Szecsi, and L. Szirmay-Kalos. Real-time light animation. *Computer Graphics Forum (Eurographics 04 Proceedings)*, 23(3):291–299, September 2004.

- [114] A. Scheel, M. Stamminger, and H.-P. Seidel. Tone reproduction for interactive walkthroughs. In *Computer Graphics Forum (Proc. Eurographics)*, pages 301–312, 2000.
- [115] C. Schlick. An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum*, 13(3):233–246, 1994.
- [116] C. Schlick. Quantization techniques for visualization of high dynamic range pictures. In *Proc. of Eurographics Workshop on Rendering '94*, pages 7–20, 1994.
- [117] P. Schröder and W. Sweldens. Spherical wavelets: Efficiently representing functions on the sphere. In *Computer Graphics 29, Annual Conference Series*, pages 161–172, 1995.
- [118] A. Secord, W. Heidrich, and L. Streit. Fast primitive distribution for illustration. In *Eurographics Workshop on Rendering*, pages 215–226, June 2002.
- [119] H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, and A. Vorozcovs. High dynamic range display systems. In *ACM Transactions on Graphics (Proc. of SIGGRAPH '04)*, pages 760–768, August 2004.
- [120] H. Seetzen, L. Whitehead, and G. Ward. A high dynamic range display using low and high resolution modulators. In *Society for Information Display International Symposium Digest of Technical Papers*, pages 1450–1453, 2003.
- [121] P. Shirley. *Realistic Ray Tracing*. A K Peters, Natick, 2000.
- [122] P. Sloan, J. Kautz, and J. Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency environments. In *Proc. of ACM SIGGRAPH '02*, pages 527–536, 2002.
- [123] P.-P. Sloan, W. Martin, A. Gooch, and B. Gooch. The lit sphere: a model for capturing NPR shading from art. In *Proc. Graphics Interface 2001*, pages 143–150, 2001.

-
- [124] P. Slusallek, T. Pflaum, and H.-P. Seidel. Using procedural renderman shaders for global illumination. In *Proc. of Eurographics*, volume 14, pages 311–324, 1995.
- [125] A. Smith and A. Gelfand. Bayesian statistics without tears: a sampling-resampling perspective. volume 46, pages 84–88, 1992.
- [126] Stanford Graphics Group. Digital Michaelangelo project, 2001. <http://graphics.stanford.edu/projects/mich/>.
- [127] J. Stumpfel, A. Jones, A. Wenger, C. Tchou, T. Hawkins, and P. Debevec. Direct hdr capture of the sun and sky, 2004. SIGGRAPH 2004 Poster.
- [128] L. Szecsi, M. Sbert, and L. Szirmay-Kalos. Combined correlated and importance sampling in direct light source computation and environment mapping. In *Proc. of Eurographics*, volume 23, pages 585–593, 2004.
- [129] J. Talbot, D. Cline, and P. Egbert. Importance resampling for global illumination. In *Eurographics Symposium on Rendering*, pages 139–146, June 2005.
- [130] P. Tan, S. Lin, B. Guo L. Quan, and H-Y. Shum. Multiresolution reflectance filtering. In *Proc. Eurographics Symposium on Rendering*, pages 111–116, 2005.
- [131] M. Tanner. *Tools for statistical inference*. Springer Verlag, New York, 3rd edition, 1996.
- [132] Brightside Technologies. DR37-P High Dynamic Range Display, 2005. <http://www.brightsidetech.com/>.
- [133] K.E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. 57(9):1105–1114, September 1967.
- [134] M. Trentacoste. Photometric image processing for high dynamic range displays. Master’s thesis, The University of British Columbia, 2005.
- [135] J. Tumblin and G. Turk. LCIS: A boundary hierarchy for detail-preserving contrast reduction. In *Proc. of ACM SIGGRAPH ’99*, pages 83–90, 1999.

- [136] E. Veach and L. Guibas. Optimally combining sampling techniques for monte carlo rendering. In *Proc. of ACM SIGGRAPH '95*, pages 419–428, August 1995.
- [137] E. Veach and L. Guibas. Metropolis light transport. In *Proc. of ACM SIGGRAPH '97*, pages 65–76, 1997.
- [138] J. Vos. Disability glare – a state of the art report. *CIE Journal*, 3(2):39–53, 1984.
- [139] L. Wan, T. Wong, and C. Leung. Spherical q2-tree for sampling dynamic environment sequences. In *Eurographics Symposium on Rendering*, June 2005.
- [140] G. J. Ward. Measuring and modeling anisotropic reflection. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 265–272, 1992.
- [141] A. Wenger, A. Gardner, C. Tchou, J. Unger, T. Hawkins, and P. Debevec. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 24(3):756–764, July 2005.
- [142] S. Westin, J Arvo, and K. Torrance. Predicting reflectance functions from complex surfaces. In *Computer Graphics 26, Annual Conference Series*, pages 255–264, 1992.

Appendix A

Radiometric and Photometric Terms

In computer graphics, the interaction of light with matter is often modeled geometrically using ray-optics. In this section we provide a brief description of some of the radiometric terms used in this thesis.

Radiant Energy, denoted by Q , is the most basic radiometric unit, measured in Joules [J]. For a photon of wavelength λ , the particle model of light gives the energy Q in terms of Planck's constant h and speed of light in vacuum c , as

$$Q = \frac{hc}{\lambda} \quad (\text{A.1})$$

Radiant Flux (or Radiant Power), denoted by ϕ , is the energy flowing through a surface per unit time and is measure in Watts [W].

$$\phi = \frac{dQ}{dt} \quad (\text{A.2})$$

Radiant Flux Area Density, denoted by u , is a measure of energy flow given by radiant flux per unit area, measured in [W/m^2].

$$u = \frac{d\phi}{dA} \quad (\text{A.3})$$

If the energy flow is toward the surface, it is referred to as *irradiance* (denoted by E) and if the energy flow is away from the surface, it is referred to as *radiosity or radiant exitance* (denoted by B).

Intensity, I , is the measure of flux with respect to solid angle instead of area and is measured in [W/sr]

$$I = \frac{d\phi}{d\vec{\omega}} \quad (\text{A.4})$$

This is useful in describing point light sources, since the area goes to zero.

Radiance, denoted by L , is a measure of radiant flux per unit projected area per unit solid angle. Its unit is $[W/m^2sr]$.

$$L = \frac{d^2\phi}{dA d\omega \cos\theta} \quad (\text{A.5})$$

where θ is the angle between the normal N of the surface area element dA and the direction of the flux ϕ . The cosine term represents the foreshortening with respect to the flux direction. *Spectral radiance* is the radiance per unit wavelength interval and is measured in $[W/m^3sr]$ units.

For each radiometric quantity, there exists a photometric quantity where the different light frequencies are weighted by the luminous efficiency function, i.e., the relative perceived brightness of light with the same power at different frequencies. Some of these quantities and their units are listed below:

Luminous Flux is the flux weighted by the luminous efficiency function and is measured in Lumens $[lm]$.

Illuminance and **Luminosity** are luminous flux densities, and correspond to irradiance and radiosity, respectively. They are measured in Lux $[lx = lm/m^2]$.

Luminance is luminous flux density per solid angle, and therefore corresponds to radiance. It is measured in Candela $[cd = lm/m^2sr = lx/sr]$.

Luminous Intensity is the photometric quantity corresponding to intensity and it is measured in cd/m^2 .

For a more details, please refer to Glassner [50].

Appendix B

Zonal Basis Function Plots

Our proposed reflectance measurement setup, *in principle*, allows measurement of reflectance in the interval $\theta \in [\pi/20, \pi/2]$, i.e. $9^\circ, \dots, 90^\circ$ from the surface normal. Angular plots of the first few ZB functions defined over this measurement space are provided in the Figure B.1.

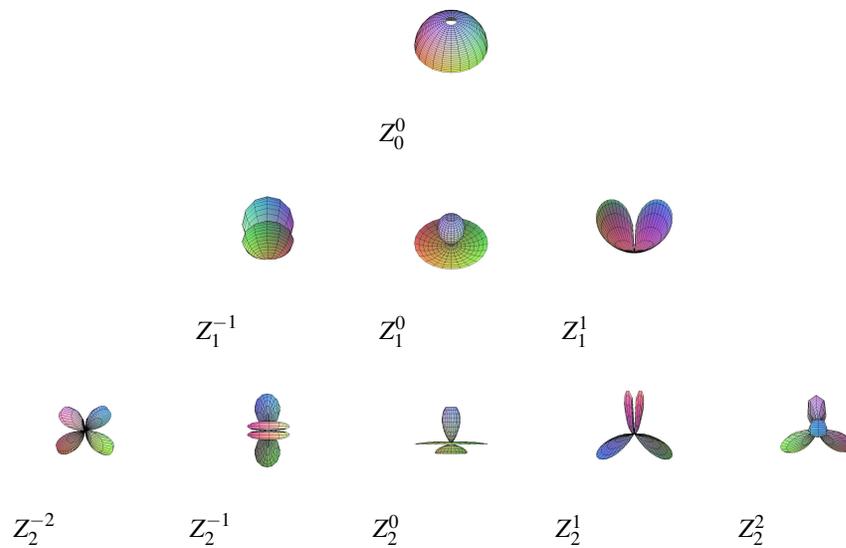


Figure B.1: The plots of zonal basis functions Z_l^m defined over the measurement space $[\pi/20, \pi/2] \times [0, 2\pi]$, for $l \leq 2$.

Appendix C

Normalization Constant

The sequential normalization constant $L_{ns,n}$ at time n of the direct illumination integral can be incrementally estimated according to Equation 5.10 in Section 5.1.1. This result can be explained as follows: the weighted empirical distribution $\{W_{n-1}^{(j)}, \omega_{i,n-1:n}^{(j)}\}$ obtained after the MCMC sampling step is an approximation of $p_{n-1}(\omega_{i,n-1})K_n(\omega_{i,n-1}, \omega_{i,n})$ according to Equation 5.5. The expectation $E(\tilde{w}_n)$ of the incremental weights \tilde{w}_n with respect to this joint distribution $p_{n-1} \cdot K_n$ is

$$\begin{aligned}
 E(\tilde{w}_n) &= \int_{\Omega} \int_{\Omega} \tilde{w}_n p_{n-1}(\omega_{i,n-1}) K_n(\omega_{i,n-1}, \omega_{i,n}) d\omega_{i,n} d\omega_{i,n-1} \\
 &= \int_{\Omega} \int_{\Omega} \frac{\tilde{p}_n(\omega_{i,n-1})}{\tilde{p}_{n-1}(\omega_{i,n-1})} p_{n-1}(\omega_{i,n-1}) K_n(\omega_{i,n-1}, \omega_{i,n}) d\omega_{i,n} d\omega_{i,n-1} \\
 &= \frac{1}{L_{ns,n-1}} \int_{\Omega} \int_{\Omega} \tilde{p}_n(\omega_{i,n-1}) K_n(\omega_{i,n-1}, \omega_{i,n}) d\omega_{i,n} d\omega_{i,n-1} \\
 &= \frac{1}{L_{ns,n-1}} \int_{\Omega} \tilde{p}_n(\omega_{i,n-1}) d\omega_{i,n-1} \\
 &= \frac{L_{ns,n}}{L_{ns,n-1}}.
 \end{aligned}$$

The Monte Carlo estimate of this expectation is given by

$$E(\tilde{w}_n) \approx \sum_{j=1}^N W_{n-1}^{(j)} \cdot \tilde{w}_n^{(j)}. \quad (\text{C.1})$$