A SIGNAL-PROCESSING FRAMEWORK FOR FORWARD AND INVERSE RENDERING

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF COMPUTER SCIENCE
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Ravi Ramamoorthi
August 2002

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Pat Hanrahan
(Principal Adviser)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Marc Levoy

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Jitendra Malik
(UC Berkeley)

Approved for the University Committee on Graduate Studies:

Abstract

The study of the computational aspects of reflection, and especially the interaction between reflection and illumination, is of fundamental importance in both computer graphics and vision. In computer graphics, the interaction between the incident illumination and the reflective properties of a surface is a basic building block in most *rendering* algorithms, i.e. methods that create synthetic computer-generated images. In computer vision, we often want to undo the effects of the reflection operator, i.e. to invert the interaction between the surface reflective properties and lighting. In other words, we often want to perform *inverse rendering*—the estimation of material and lighting properties from real photographs. Inverse rendering is also of increasing importance in graphics, where we wish to obtain accurate input illumination and reflectance models for (forward) rendering.

This dissertation describes a new way of looking at reflection on a curved surface, as a special type of convolution of the incident illumination and the reflective properties of the surface (technically, the bi-directional reflectance distribution function or BRDF). The first part of the dissertation is devoted to a theoretical analysis of the reflection operator, leading for the first time to a formalization of these ideas, with the derivation of a convolution theorem in terms of the spherical harmonic coefficients of the lighting and BRDF. This allows us to introduce a signal-processing framework for reflection, wherein the incident lighting is the signal, the BRDF is the filter, and the reflected light is obtained by filtering the input illumination (signal) using the frequency response of the BRDF filter.

The remainder of the dissertation describes applications of the signal-processing framework to forward and inverse rendering problems in computer graphics. First, we address the forward rendering problem, showing how our framework can be used for computing and displaying synthetic images in real-time with natural illumination and physically-based

BRDFs. Next, we extend and apply our framework to inverse rendering. We demonstrate estimation of realistic lighting and reflective properties from photographs, and show how this approach can be used to synthesize very realistic images under novel lighting and viewing conditions.

Acknowledgements

Firstly, I would like to thank my advisor, Pat Hanrahan, for convincing me to come to Stanford, for insightful discussions and technical advice during all these years on the whole spectrum of graphics topics, and for the perhaps more important discussions and advice on following appropriate scientific methodology, and development as a scientist and researcher. It has indeed been an honor and a privilege to work with him through these four years.

I would also like to thank my other committee members, Marc Levoy, Jitendra Malik, Ron Fedkiw, and Bernd Girod, for the advice, support, encouragement and inspiration they have provided over the years. I especially want to thank Marc for the many discussions we had in the early stages of this project.

Over the course of my time here, it has been exciting and energizing to work with and be around such an amazing group of people in the Stanford graphics lab. I want to thank Maneesh Agrawala, Sean Anderson, James Davis, Ziyad Hakura, Olaf Hall-Holt, Greg Humphreys, David Koller, Bill Mark, Steve Marschner, Matt Pharr, Kekoa Proudfoot, Szymon Rusinkiewicz, Li-Yi Wei and many others for being such wonderful friends, colleagues and co-workers. In particular, I want to thank Steve and Szymon for the many hours spent discussing ideas, their patience in reviewing drafts of my papers, and their help with data acquisition for the inverse rendering part of this dissertation.

Over the last four years, the Stanford graphics lab has really been the epicenter for graphics research. I can only consider it my privilege and luck to have been a part of this golden age.

Contents

Abstract **iv**

Acknowledgements **vi**

1 Introduction **1**

- 1.1 Theoretical analysis of Reflection: Signal Processing 3
- 1.2 Forward Rendering 5
- 1.3 Inverse Rendering 7

2 Reflection as Convolution **10**

- 2.1 Previous Work 12
- 2.2 Reflection Equation 15
  - 2.2.1 Assumptions 16
  - 2.2.2 Flatland 2D case 18
  - 2.2.3 Generalization to 3D 21
- 2.3 Frequency-Space Analysis 25
  - 2.3.1 Fourier Analysis in 2D 25
  - 2.3.2 Spherical Harmonic Analysis in 3D 28
  - 2.3.3 Group-theoretic Unified Analysis 35
  - 2.3.4 Alternative Forms 36
- 2.4 Implications 43
  - 2.4.1 Forward Rendering with Environment Maps 43
  - 2.4.2 Well-posedness and conditioning of Inverse Lighting and BRDF . . 44
  - 2.4.3 Light Field Factorization 47

2.5 Conclusions and Future Work 50

3 Formulae for Common Lighting and BRDF Models 52

3.1 Background 53

3.1.1 Reflection Equation and Convolution Formula 53

3.1.2 Analysis of Inverse Problems 55

3.2 Derivation of Analytic Formulae 56

3.2.1 Directional Source 57

3.2.2 Axially Symmetric Distribution 59

3.2.3 Uniform Lighting 60

3.2.4 Mirror BRDF 62

3.2.5 Lambertian BRDF 65

3.2.6 Phong BRDF 70

3.2.7 Microfacet BRDF 73

3.3 Conclusions and Future Work 78

4 Irradiance Environment Maps 80

4.1 Introduction and Previous Work 81

4.2 Background 83

4.3 Algorithms and Results 87

4.3.1 Prefiltering 87

4.3.2 Rendering 89

4.3.3 Representation 91

4.4 Conclusions and Future Work 93

5 Frequency Space Environment Map Rendering 94

5.1 Introduction 95

5.2 Related Work 97

5.3 Preliminaries 97

5.3.1 Reparameterization by central BRDF direction 99

5.3.2 4D function representations 101

5.4 Spherical Harmonic Reflection Maps 103

5.5 Analysis of sampling rates/resolutions 107
    5.5.1 Order of spherical harmonic expansions 108
    5.5.2 Justification for SHRM representation 111
5.6 Prefiltering 112
    5.6.1 Main steps and insights 113
    5.6.2 Prefiltering Algorithm 114
    5.6.3 Computational complexity 116
    5.6.4 Validation with Phong BRDF 118
5.7 Results 120
    5.7.1 Number of coefficients for analytic BRDFs 120
    5.7.2 Number of coefficients for measured BRDFs 122
    5.7.3 SHRM accuracy 123
    5.7.4 Speed of prefiltering 123
    5.7.5 Real-time rendering 125
5.8 Conclusions and Future Work 126

6 Inverse Rendering Under Complex Illumination 133
6.1 Taxonomy of Inverse problems and Previous Work 135
    6.1.1 Previous Work on Inverse Rendering 137
    6.1.2 Open Problems 141
6.2 Preliminaries 143
    6.2.1 Practical implications of theory 144
6.3 Dual angular and frequency-space representation 146
    6.3.1 Model for reflected light field 147
    6.3.2 Textures and shadowing 148
6.4 Algorithms 150
    6.4.1 Data Acquisition 151
    6.4.2 Inverse BRDF with known lighting 152
    6.4.3 Inverse Lighting with Known BRDF 157
    6.4.4 Factorization—Unknown Lighting and BRDF 163
6.5 Results on Complex Geometric Objects 166

6.6 Conclusions and Future Work 169

7 Conclusions and Future Work **174**

7.1 Computational Fundamentals of Reflection 175

7.2 High Quality Interactive Rendering 178

7.3 Inverse Rendering 180

7.4 Summary 182

A Properties of the Representation Matrices **183**

Bibliography **185**

List of Tables

1.1 Common notation used throughout the dissertation. 9

4.1 Scaled RGB values of lighting coefficients for a few environments. 89

5.1 Notation used in chapter 5. 98
5.2 Comparison of different 4D representations. 102
5.3 Computational complexity of prefiltering. 117
5.4 Comparison of timings of angular and frequency-space prefiltering for different values of the Phong exponent s. 124
5.5 Times for angular-space and our frequency-space prefiltering 124

List of Figures

1.1 A scene rendered in real time using our approach, described in chapter 5. . 5
1.2 Inverse Rendering. 7

2.1 Schematic of reflection in 2D. 18
2.2 Diagram showing how the rotation corresponding to (α, β, γ) transforms between local (primed) and global (unprimed) coordinates. 22
2.3 The first 3 orders of real spherical harmonics ($l = 0, 1, 2$) corresponding to a total of 9 basis functions. 30
2.4 Reparameterization involves recentering about the reflection vector. 41

3.1 The *clamped cosine* filter corresponding to the Lambertian BRDF and successive approximations obtained by adding more spherical harmonic terms. 65
3.2 The solid line is a plot of $\hat{\rho}_l$ versus l, as per equation 3.32. 67
3.3 Numerical plots of the Phong coefficients $\Lambda_l \hat{\rho}_l$, as defined by equation 3.37. 72

4.1 The diffuse shading on all the objects is computed procedurally in real-time using our method. 82
4.2 A comparison of irradiance maps from our method to standard prefiltering. 90
4.3 Illustration of our representation, and applications to control appearance. . 92

5.1 These images, showing many different lighting conditions and BRDFs, were each rendered at approximately 30 frames per second using our Spherical Harmonic Reflection Map (SHRM) representation. 95
5.2 An overview of our entire pipeline. 103
5.3 The idea behind SHRMs. 105

5.4 Renderings with different lighting and BRDF conditions. 108
5.5 Accuracy $(1-\epsilon)$ versus frequency F for an order F approximation of the reflected light field B 111
5.6 Comparing images obtained with different values for P for a simplified microfacet BRDF model with surface roughness $\sigma = 0.2$. 121
5.7 Accuracy of a spherical harmonic BRDF approximation for all 61 BRDFs in the CURET database. 130
5.8 Comparing the correct image on the left to those created using SHRMs (middle) and the 2D BRDF approximation of Kautz and McCool (right). . 131
5.9 Comparing the correct image to those created using SHRMs and icosahedral interpolation (Cabral's method). 132

6.1 **Left:** Schematic of experimental setup **Right:** Photograph 151
6.2 Direct recovery of BRDF coefficients. 153
6.3 Comparison of BRDF parameters recovered by our algorithm under complex lighting to those fit to measurements made by the method of Marschner et al. [55]. 154
6.4 Estimation of dual lighting representation. 159
6.5 Comparison of inverse lighting methods. 163
6.6 Determining surface roughness parameter σ. 165
6.7 BRDFs of various spheres, recovered under known (section 6.4.2) and unknown (section 6.4.4) lighting. 166
6.8 Spheres rendered using BRDFs estimated under known (section 6.4.2) and unknown (section 6.4.4) lighting. 167
6.9 Comparison of photograph and rendered image for cat sculpture. 168
6.10 Comparison of photographs (middle column) to images rendered using BRDFs recovered under known lighting (left column), and using BRDFs (and lighting) estimated under unknown lighting (right column). 169
6.11 BRDF and lighting parameters for the cat sculpture. 170
6.12 Recovering textured BRDFs under complex lighting. 171