Chapter 1

Introduction

The central problem in computer graphics is creating, or *rendering*, realistic computer-generated images that are indistinguishable from real photographs, a goal referred to as *photorealism*. Applications of *photorealistic rendering* extend from entertainment such as movies and games, to simulation, training and virtual reality applications, to visualization and computer-aided design and modeling. Progress towards photorealism in rendering involves two main aspects. First, one must develop an algorithm for physically accurate light transport simulation. However, the output from the algorithm is only as good as the inputs. Therefore, photorealistic rendering also requires accurate input models for object geometry, lighting and the reflective properties of surfaces.

Over the past two decades, there has been a signficant body of work in computer graphics on accurate light transport algorithms, with increasingly impressive results. One consequence has been the increasing realism of computer-generated special effects in movies, where it is often difficult or impossible to tell real from simulated.

However, in spite of this impressive progress, it is still far from routine to create photorealistic computer-generated images. Firstly, most light transport simulations are slow. Within the context of interactive imagery, such as with hardware rendering, it is very rare to find images rendered with natural illumination or physically accurate reflection functions. This gap between interactivity and realism impedes applications such as virtual lighting or material design for visualization and entertainment applications.

A second difficulty, which is often the limiting problem in realism today, is that of obtaining accurate input models for geometry, illumination and reflective properties. Entertainment applications such as movies often require very laborious fine-tuning of these parameters. One of the best ways of obtaining high-quality input illumination and material data is through measurements of scene attributes from real photographs. The idea of using real data in graphics is one that is beginning to permeate all aspects of the field. Within the context of animation, real motions are often measured using *motion capture* technology, and then retargetted to new characters or situations. For geometry, it is becoming more common to use *range scanning* to measure the 3D shapes of real objects for use in computer graphics. Similarly, measuring rendering attributes—lighting, textures and reflective properties—from real photographs is increasingly important. Since our goal is to invert the traditional rendering process, and estimate illumination and reflectance from real images, we refer to the approach as *inverse rendering*. Whether traditional or image-based rendering algorithms are used, rendered images now use measurements from real objects, and therefore appear very similar to real scenes.

In recent years, there has been significant interest in acquiring material models using inverse rendering. However, most previous work has been conducted in highly controlled lighting conditions, usually by careful active positioning of a single point source. Indeed, complex realistic lighting environments are rarely used in either forward or inverse rendering. While our primary motivation derives from computer graphics, many of the same ideas and observations apply to computer vision and perception. Within these fields, we often seek to perform *inverse rendering* in order to use intrinsic reflection and illumination parameters for modeling and recognition. Although it can be shown that the perception of shape and material properties may be qualitatively different under natural lighting conditions than artificial laboratory settings, most vision algorithms work only under very simple lighting assumptions—usually a single point light source without any shadowing.

It is our thesis that a deeper understanding of the computational nature of reflection and illumination helps to address these difficulties and restrictions in a number of areas in computer graphics and vision. This dissertation is about 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. Although this idea has long been known qualitatively, this is the first time this notion of reflection as convolution has been formalized with an analytic convolution formula in the spherical domain. The dissertation includes both a theoretical analysis of reflection in terms of signal-processing, and practical applications of this frequency domain analysis to problems in forward and inverse rendering. In the rest of this chapter, we briefly discuss the main areas represented in the dissertation, summarizing our contributions and giving an outline of the rest of the dissertation. At the end of the chapter, table 1.1 summarizes the notation used in the rest of the dissertation.

1.1 Theoretical analysis of Reflection: Signal Processing

The computational study of the interaction of light with matter is a basic building block in rendering algorithms in computer graphics, as well as of interest in both computer vision and perception. In computer vision, previous theoretical work has mainly focussed on the problem of estimating shape from images, with relatively little work on estimating material properties or lighting. In computer graphics, the theory for forward global illumination calculations, involving all forms of indirect lighting, has been fairly well developed. The foundation for this analysis is the rendering equation [35]. However, there has been relatively little theoretical work on inverse problems or on the simpler reflection equation, which deals with the direct illumination incident on a surface.

It should be noted that there is a significant amount of qualitative knowledge regarding the properties of the reflection operator. For instance, we usually represent reflections from a diffuse surface at low resolutions [59] since the reflection from a matte, or technically *Lambertian*, surface *blurs* the illumination. Similarly, we realize that it is essentially impossible to accurately estimate the lighting from an image of a Lambertian surface; instead, we use mirror surfaces, i.e. gazing spheres.

In this dissertation, we formalize these observations and present a mathematical theory of reflection for general complex lighting environments and arbitrary reflection functions in terms of signal-processing. Specifically, we describe a signal-processing framework for analyzing the reflected light field from a homogeneous convex curved surface under distant illumination. Under these assumptions, we are able to derive an analytic formula for the reflected light field in terms of the spherical harmonic coefficients of the BRDF and the lighting. Our formulation leads to the following theoretical results:

Signal-Processing Framework for Reflection as Convolution: Chapter 2 derives our signal-processing framework for reflection. The reflected light field can therefore be thought of in a precise quantitative way as obtained by convolving the lighting and reflective properties of the surface (technically, the bi-directional reflectance distribution function or BRDF), i.e. by filtering the incident illumination using the BRDF. Mathematically, we are able to express the frequency-space coefficients of the reflected light field as a product of the spherical harmonic coefficients of the illumination and the BRDF. We believe this is a useful way of analyzing many forward and inverse problems. In particular, forward rendering can be viewed as *convolution* and inverse rendering as *deconvolution*.

Analytic frequency-space formulae for common lighting conditions and BRDFs Chapter 3 derives analytic formulae for the spherical harmonic coefficients of many common lighting and BRDF models. Besides being of practical interest, these formulae allow us to reason precisely about forward and inverse rendering in the frequency domain.

Well-posedness and Conditioning of Forward and Inverse Problems: Inverse problems can be ill-posed—there may be several solutions. They are also often numerically ill-conditioned, which may make devising practical algorithms infeasible. From our signal-processing framework, and the analytic formulae derived by us for common lighting and BRDF models, we are able to analyze the well-posedness and conditioning of a number of inverse problems. This analysis is presented in chapter 3, and explains many previous empirical observations, as well as serving as a guideline for future research in inverse rendering. We expect fruitful areas of research to be those problems that are well-conditioned. Additional assumptions will likely be required to address ill-conditioned or ill-posed inverse problems. This analysis is also of interest for forward rendering. An ill-conditioned inverse problem corresponds to a forward problem where the final results are not sensitive to certain components of the initial conditions. This often allows us to approximate the initial conditions—such as the incident illumination—in a principled way, giving rise to more efficient forward rendering algorithms.

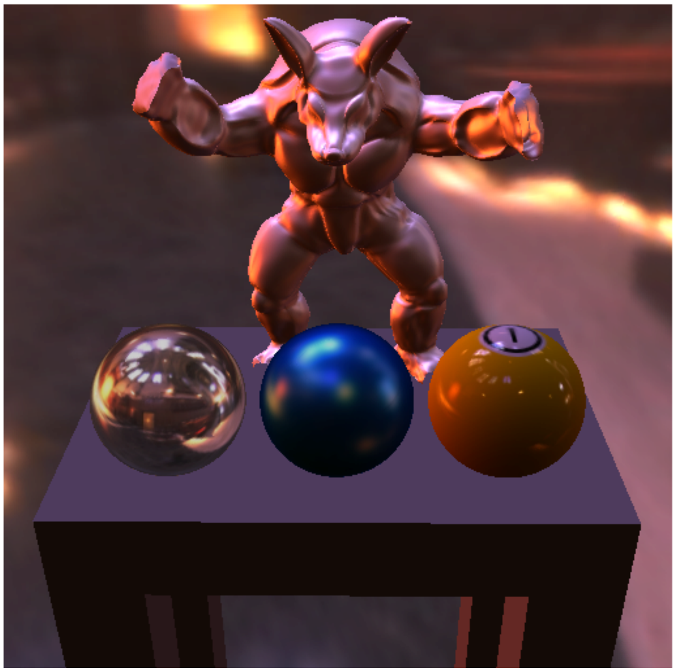

Figure 1.1: *A scene rendered in real time using our approach, described in chapter 5. The illumination is measured light in the Grace Cathedral in San Francisco, obtained by photographing a mirror sphere, courtesy of Paul Debevec. The surface reflective properties include a number of measured BRDFs.*

1.2 Forward Rendering

Lighting in most real scenes is complex, coming from a variety of sources including area lights and large continuous lighting distributions like skylight. But current graphics hardware only supports point or directional light sources, and very simple unphysical surface reflection models. One reason is the lack of simple procedural formulas for efficiently computing the reflected light from general lighting distributions. Instead, an integration over the visible hemisphere must be done for each pixel in the final image.

Chapters 4 and 5 apply the signal-processing framework to interactive rendering with natural illumination and physically-based reflection functions. An example image is shown in figure 1.1. That image includes natural illumination and a number of physically-based or measured surface reflective properties.

As is common with interactive rendering, we neglect the effects of global effects like cast shadows (self-shadowing) and interreflections, and restrict ourselves to distant illumination. Since the illumination corresponds to captured light in an environment, such rendering methods are frequently referred to as *environment mapping*.

Chapter 4 demonstrates how our signal-processing framework can be applied to *irradiance environment maps*, corresponding to the reflection from perfectly diffuse or Lambertian surfaces. The key to our approach is the rapid computation of an analytic approximation to the irradiance environment map. For rendering, we demonstrate a simple procedural algorithm that runs at interactive frame rates, and is amenable to hardware implementation. Our novel representation also suggests new approaches to lighting design and image-based rendering.

Chapter 5 introduces a new paradigm for prefiltering and rendering environment mapped images with general isotropic BRDFs. Our approach uses spherical frequency domain methods, based on the earlier theoretical derivations. Our method has many advantages over the angular (spatial) domain approaches:

Theoretical analysis of sampling rates and resolutions: Most previous work has determined reflection map resolutions, or the number of reflection maps required, in an ad-hoc manner. By using our signal-processing framework, we are able to perform error analysis, that allows us to set sampling rates and resolutions accurately.

Efficient representation and rendering with Spherical Harmonic Reflection Maps: We introduce *spherical harmonic reflection maps (SHRMs)* as a compact representation. Instead of a single color, each pixel stores coefficients of a spherical harmonic expansion encoding view-dependence of the reflection map. An important observation that emerges from the theoretical analysis is that for almost all BRDFs, a very low order spherical harmonic expansion suffices. Thus, SHRMs can be evaluated in real-time for rendering. Further, they are significantly more compact and accurate than previous methods.

Fast precomputation (prefiltering): One of the drawbacks of current environment mapping techniques is the significant computational time required for prefiltering, or precomputing reflection maps, which can run into hours, and preclude the use of these approaches

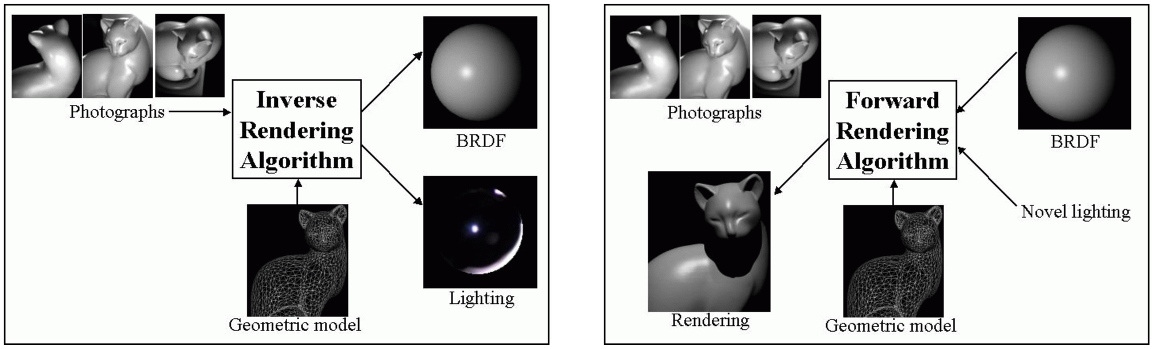

Figure 1.2: *Inverse Rendering. On the left, we show a conceptual description, while the right shows how illumination may be manipulated to create realistic new synthetic images.*

in applications involving lighting and material design, or dynamic lighting. We introduce new prefiltering methods based on spherical harmonic transforms, and show both empirically, and analytically by computational complexity analysis, that our algorithms are orders of magnitude faster than previous work.

1.3 Inverse Rendering

Finally, chapter 6 discusses practical inverse rendering methods for estimating illumination and reflective properties from sequences of images. The general idea in inverse rendering is illustrated in figure 1.2. On the left, we show a conceptual description. The inputs are photographs of the object of interest, and a geometric model. The outputs are the reflective properties, visualized by rendering a sphere with the same BRDF, and the illumination, visualized by showing an image of a chrome-steel or mirror sphere. On the right, we show how new renderings can be created under novel lighting conditions using the measured BRDF from the real object. We simply need to use a standard physically based rendering algorithm, with the input BRDF model that obtained by inverse rendering.

There are a number of applications of this approach. Firstly, it allows us to take input photographs and relight the scene, or include real objects in synthetic scenes with novel illumination and viewing conditions. This has applications in developing new methods for image-processing, and in entertainment and virtual reality applications, besides the usefulness in creating realistic computer-generated images.

The dissertation uses the theoretical framework described earlier, and the formal analysis of the conditioning of inverse problems, to derive new robust practical algorithms for inverse rendering. Our specific practical contributions are:

Complex Illumination: As stated in the introduction, one of our primary motivations is to perform inverse rendering under complex illumination, allowing these methods to be used in arbitrary indoor and outdoor settings. We present a number of improved algorithms for inverse rendering under complex lighting conditions.

New Practical Representations And Algorithms: The theoretical analysis is used to develop a new low-parameter practical representation that simultaneously uses the spatial or angular domain and the frequency domain. Using this representation, we develop a number of new inverse rendering algorithms that use both the spatial and the frequency domain.

Simultaneous Determination Of Lighting And BRDF: In many practical applications, it might be useful to determine reflective properties without knowing the lighting, or to determine both simultanously. We present the first method to determine BRDF parameters of complex geometric models under unknown illumination, simultaneously finding the lighting.

$\vec{L}, \vec{N}, \vec{V}, \vec{R}$	Global incident, normal, viewing, reflected directions
B	Reflected radiance
B_{lp}	Coefficients of Fourier expansion of B in 2D
B_{lmpq}, B_{lmnpq}	Coefficients of isotropic, anisotropic basis-function expansion of B in 3D
L	Incoming radiance
L_l	Coefficients of Fourier expansion of L in 2D
L_{lm}	Coefficients of spherical-harmonic expansion of L in 3D
ρ	Surface BRDF
$\hat{\rho}$	BRDF multiplied by cosine of incident angle
$\hat{\rho}_{lp}$	Coefficients of Fourier expansion of $\hat{\rho}$ in 2D
$\hat{\rho}_{lpq}, \hat{\rho}_{ln;pq}$	Coefficients of isotropic, anisotropic spherical-harmonic expansion of $\hat{\rho}$
θ_i', θ_i	Incident elevation angle in *local*, *global* coordinates
ϕ_i', ϕ_i	Incident azimuthal angle in *local*, *global* coordinates
θ_o', θ_o	Outgoing elevation angle in *local*, *global* coordinates
ϕ_o', ϕ_o	Outgoing azimuthal angle in *local*, *global* coordinates
$\vec{X}$	Surface position
α	Surface normal parameterization—elevation angle
β	Surface normal parameterization—azimuthal angle
γ	Orientation of tangent frame for anisotropic surfaces
R_α	Rotation operator for surface orientation α in 2D
$R_{\alpha,\beta}, R_{\alpha,\beta,\gamma}$	Rotation operator for surface normal (α, β) or tangent frame (α, β, γ) in 3D
$F_k(\theta)$	Fourier basis function (complex exponential)
$F_k^*(\theta)$	Complex Conjugate of Fourier basis function
$Y_{lm}(\theta, \phi)$	Spherical Harmonic
$Y_{lm}^*(\theta, \phi)$	Complex Conjugate of Spherical Harmonic
$f_{lm}(\theta)$	Normalized θ dependence of Y_{lm}
$D_{mn}^l(\alpha, \beta, \gamma)$	Representation matrix of dimension $2l + 1$ for rotation group $SO(3)$
$d_{mn}^l(\alpha)$	$D_{mm'}^L$ for y-axis rotations
Λ_l	Normalization constant, $\sqrt{4\pi/(2l+1)}$
I	$\sqrt{-1}$

Table 1.1: *Common notation used throughout the dissertation.*