

SPECTRAL ANALYSIS OF INTERREFLECTION

by

Ian A. Harder

B.Sc. Simon Fraser University 1989

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF SCIENCE
in the School
of
Computing Science

© Ian A. Harder 1992

SIMON FRASER UNIVERSITY

August 1992

All rights reserved. This work may not be
reproduced in whole or in part, by photocopy
or other means, without the permission of the author.

APPROVAL

Name: Ian A. Harder
Degree: Master of Science
Title of Thesis: Spectral Analysis of Interreflection

Examining Committee:

Dr. Fred Popowich
Chair

Dr. Brian V. Funt
Senior Supervisor
Professor, Computing Science

Dr. Ze-Nian Li
Supervisor
Assistant Professor, Computing Science

Dr. Jacques Vaisey
Assistant Professor, Engineering Science

Date Approved:

August 10, 1992

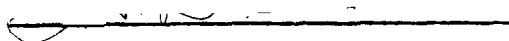
PARTIAL COPYRIGHT LICENSE

I hereby grant to Simon Fraser University the right to lend my thesis, project or extended essay (the title of which is shown below) to users of the Simon Fraser University Library, and to make partial or single copies only for such users or in response to a request from the library of any other university, or other educational institution, on its own behalf or for one of its users. I further agree that permission for multiple copying of this work for scholarly purposes may be granted by me or the Dean of Graduate Studies. It is understood that copying or publication of this work for financial gain shall not be allowed without my written permission.

Title of Thesis/Project/Extended Essay

Spectral Analysis of Interreflection.

Author:


(signature)

Ian August Harder

(name)

August 13, 1992

(date)

Abstract

Interreflection (or *mutual illumination*) occurs when two or more object surfaces are illuminated both by a light source and the light reflected from other surfaces. As the distance or angle between two interreflecting surfaces decreases, the intensity of interreflected light increases, with a corresponding shift in colour known as *colour bleeding*. For computer vision algorithms that assume spatially invariant surface reflectances, this plays a confounding role. As an example, in the presence of interreflection, “shape-from-shading” methods will incorrectly reconstruct surfaces such that the orientation of their surface normals will appear to be closer to the direction of the illuminant than they actually are.

Rather than treating interreflection as noise, surface colours can be analysed to provide additional information such as the illuminant spectra and surface shading. In this thesis, a finite dimensional model is employed to recover the surface spectral reflectances of two interreflecting Lambertian surfaces under a known illuminant. The resulting reflectances are used to construct *colour basis vectors* for linear decomposition of colour channel intensities for each surface, from which the coefficients of the *no-bounce* colour components (*shading fields*) are extracted. The robustness of this simple, straightforward algorithm is tested on both synthetic and real interreflecting planar surfaces, and the improvement the recovered shading field provides over image intensity is demonstrated using a simplified shape-from-shading scheme.

Acknowledgements

I would like to express my gratitude to Brian Funt for his patience and thoughtfulness, and perceptiveness in the field of computer vision he is always willing to share. I am also indebted to Mark Drew and Graham Finlayson for their mathematical insight and enthusiastic input.

I extend my love and appreciation to my parents, my wife Louise, and our daughter Phaedra who have provided me with the incentive and environment for pursuing my education.

Finally, I would like to acknowledge the Natural Sciences and Engineering Research Council and Simon Fraser University for their generous financial support.

Contents

Abstract	iii
Acknowledgements	iv
1 Introduction	1
1.1 Motivation	1
1.2 Scope of Thesis	3
1.3 Thesis Outline	4
2 Preliminaries	6
2.1 Symbol Definitions	6
2.2 Colour Signal Decomposition	7
2.3 Lambertian Surfaces	8
2.4 Error Measurement	8
2.4.1 Just-Noticeable Difference	9
2.4.2 Angle Between Vectors	10
2.4.3 Root Mean Square	10
3 Previous Work	11
3.1 Radiosity	11
3.2 Spectral Methods	12
3.2.1 A Qualitative Analysis	12
3.2.2 Colour Constant Analysis	13
3.3 Reduced Component Methods	14

3.3.1	Intensity and Interreflection	14
3.3.2	Intersection of Planes	15
4	Finite Dimensional Models	17
4.1	Background	17
4.1.1	Modeling Surface and Illuminant Spectra	17
4.1.2	Basis Functions	18
4.2	Basis Function Calculation	20
4.2.1	Implementation	20
4.2.2	Results	21
5	Models of Interreflection	23
5.1	One-bounce model	24
5.2	Two-bounce model	27
6	Multibounce Colour Space	30
6.1	Implementation	30
6.2	Results	33
7	Reconstructing Multibounce Colours	35
7.1	Assumptions	36
7.2	Implementation	37
7.3	Results	38
7.3.1	One-bounce Colour Reconstruction	38
7.3.2	Two-bounce Colour Reconstruction	39
8	Surface Colour Decomposition	42
8.1	Assumptions	42
8.2	Algorithm	43
8.3	Selected Cases	44
8.3.1	Test Case 1	45
8.3.2	Test Case 2	45
8.3.3	Test Case 3	46

8.4	Reconstructing Surface Shape	48
8.5	Synthetic Images	50
8.5.1	Constructing the Images	50
8.5.2	Implementation	51
8.5.3	Results	52
8.6	Real Images	58
8.6.1	Obtaining Linear Responses	58
8.6.2	Implementation	60
8.6.3	Results	60
9	Discussion and Conclusions	67

List of Tables

6.1	Colour vector angles θ and ϕ under 3200K illuminant.	33
6.2	Colour vector angles θ and ϕ under D_{65} illuminant.	33
7.1	Error statistics for reconstructed $C^{(1)}(\lambda)$ under 3200K illuminant. . .	40
7.2	Error statistics for reconstructed $C^{(1)}(\lambda)$ under D_{65} illuminant.	40
7.3	Error statistics for reconstructed $\rho^{(1)}$ under 3200K illuminant.	40
7.4	Error statistics for reconstructed $\rho^{(1)}$ under D_{65} illuminant.	40
7.5	Error statistics for reconstructed $C^{(2)}(\lambda)$ under 3200K illuminant. . .	41
7.6	Error statistics for reconstructed $C^{(2)}(\lambda)$ under D_{65} illuminant.	41
7.7	Error statistics for reconstructed $\rho^{(2)}$ under 3200K illuminant.	41
7.8	Error statistics for reconstructed $\rho^{(2)}$ under D_{65} illuminant.	41
8.1	Colour signal reconstruction errors for selected cases.	44
8.2	Colour response reconstruction errors for selected cases.	45
8.3	Comparison of angles (degrees) for ρ of actual model and for reconstruction of test cases.	45

List of Figures

4.1	Cohen's first three basis vectors obtained through linear component analysis of 150 randomly selected Munsell colour chips.	22
4.2	The first three basis vectors obtained through singular value decomposition of 462 Munsell colour chips.	22
5.1	Interreflection between two coloured surfaces with spectral reflectances $S^A(\lambda)$ and $S^B(\lambda)$, showing no-bounce colour signals $C^{(A)}$ and $C^{(B)}$, and one-bounce colour signal $C^{(AB)}$	24
5.2	Sample multibounce colour components of signal $C^A(x, \lambda)$ for increasing number of bounces i , resulting from interreflection between Munsell colours 10BG6/4 (blue-green) and 5P5/6 (purple) under standard illuminant D_{65}	25
5.3	Angle θ between $\rho^{(0)}$ and $\rho^{(1)}$, and angle ϕ between $\rho^{(2)}$ and the plane described by $\rho^{(0)}$ and $\rho^{(1)}$ (shaded area) in three-dimensional RGB space. 29	29
6.1	Relative radiant power distributions of the D_{65} standard illuminant and a 3200K Planckian (blackbody) radiator.	31
6.2	Relative response curves for CIE 1931 standard observer filters.	31
6.3	Relative response curves for Sony DXC-151 filters.	32
6.4	Relative response curves for Kodak #47b (blue), #58 (green), and #25 (red) filters.	32
8.1	Actual and modeled spectral reflectances $S(\lambda)$ for case 1 surfaces A (blue-green) and B (purple).	46

8.2	Actual and modeled spectral reflectances $S(\lambda)$ for case 2 surfaces A (red) and B (blue).	47
8.3	Actual and modeled spectral reflectances $S(\lambda)$ for case 2 surfaces A (grey) and B (red).	48
8.4	Ray geometry for luminous intensity I of a Lambertian surface resulting from illumination by luminous flux E . ϵ is the angle between surface normal \mathbf{n} and the direction of E	49
8.5	Geometry of a semi-infinite edge used to simulate interreflection between two planar surfaces A and B with opening angle β	51
8.6	Test case 1 shading fields using one-bounce model. Vertical line represents edge between surface A (blue-green) on left and surface B (purple) on right.	53
8.7	Test case 1 shading fields using two-bounce model. Vertical line represents edge between surface A (blue-green) on left and surface B (purple) on right.	53
8.8	Shape from intensity and from one-bounce shading field for test case 1, with surface A on left and surface B on right.	54
8.9	Shape from one-bounce shading field and from two-bounce shading field.	54
8.10	Test case 2 shading fields using one-bounce model. Vertical line represents edge between surface A (red) on left and surface B (blue) on right.	56
8.11	Shape from intensity and from one-bounce shading field for test case 2, with surface A (red) on left and surface B (blue) on right.	56
8.12	Test case 3 shading fields using the one-bounce model. Vertical line represents edge between surface A (grey) on left and surface B (red) on right. Decomposition is not possible for surface B using this model.	57
8.13	Shape from intensity and from one-bounce shading field for test case 3, with surface A (grey) on left and surface B (red) on right. One-bounce reconstruction cannot be performed for surface B	57
8.14	Transmittance curve of Kodak 301-A infrared filter.	59
8.15	Shading fields of test case 1 using the one-bounce model on a real image.	61

8.16	Shape from intensity and from one-bounce shading fields for real test case 1.	61
8.17	Shading fields of test case 2 using the one-bounce model on a real image.	63
8.18	Shape from intensity and from one-bounce shading fields for real test case 2.	63
8.19	Shading fields of test case 3 using the one-bounce model on a real image. Decomposition is not possible for surface <i>B</i>	64
8.20	Shape from intensity and from one-bounce shading fields for real test case 3. Shape can be only be recovered using shading fields for surface <i>A</i>	64

Chapter 1

Introduction

1.1 Motivation

When two or more objects are illuminated by both a light source and light reflected from other objects, a *colour bleeding* effect is observed, with an increase in colour intensity and a shift in hue as distances or angles between surfaces decrease. This is known as *interreflection* (or *mutual illumination*).

There are two general problems associated with interreflection. The graphics problem is that of rendering realistic images in closed environments. First introduced by Goral *et al* [15], *radiosity methods* require knowledge of three-dimensional geometry, surface reflectance properties, and signal spectra. The task is to project this information onto a two-dimensional image. The inverse vision problem is the extraction of three-dimensional information from a two-dimensional image in the presence of interreflection.

For computer vision algorithms that assume spatially invariant surface reflectances, interreflection plays a confounding role. As an example, “shape-from-shading” methods that use image intensity as indicators of shape incorrectly reconstruct surfaces such that the orientation of their surface normals appears to be closer to the direction of the illuminant than they actually are [12, 22]. A measure of shading that is independent of the effects of interreflection would thus be most useful to these methods for improved accuracy of reconstruction.

For most natural surfaces, interreflection intensity is significantly attenuated each time a colour signal is reflected from a surface, therefore it is reasonable to use a *finite-bounce* interreflection model. A *one-bounce* model approximates the actual signal as a linear combination of a *no-bounce colour* and a *one-bounce colour*. The no-bounce colour is that which results from the reflection of the source illuminant off a single surface. The one-bounce colour results from the reflection of a no-bounce colour from a second surface. Decomposition of the response (eg. *RGB* from a camera) to an interreflected colour into no-bounce and one-bounce components will result in two coefficients. The no-bounce coefficient can be described as a *shading field* as it represents the shading for each point on a surface independent of interreflection. This field can be used to provide improved shape reconstruction over image intensity for shape-from-shading methods.

If both the no-bounce and one-bounce colour signals or their filter responses are known, then recovery of their coefficients from an interreflected colour can be performed directly by simple linear decomposition, such a matrix elimination. However, the one-bounce colour exists only in the presence of the no-bounce colour (assuming surfaces are not shaded from the illuminant). Furthermore, it is the product of surface reflectance and a colour signal over all wavelengths. There is no simple transformation between the responses of no-bounce and one-bounce responses.

A straightforward approach to obtaining the one-bounce response is to read two responses from the interreflecting surfaces at points where interreflection is negligible, reconstruct surface reflectances and colour signals, then calculate the one-bounce colour and its resulting filter responses. The task then becomes one of reconstructing reflectances and colour signals. If the illuminant and filter spectra are known, a finite dimensional model can be used to recover surface reflectance.

Although we expect the one-bounce model to closely approximate the actual colour signal obtained from an interreflecting surface, as the concavity of the surfaces increases or the distance between them decreases, other multibounce colours contribute more to the final surface colour and the model begins to break down [6, 13]. We are thus motivated to progress to a two-bounce model. This is simply an extension of the one-bounce model with the addition of a third colour component: that consisting

of light reflected from one surface, then from a second surface, and again from the first surface. For decomposition, the two-bounce response can be calculated using the same simple technique as the one-bounce response, with an additional reflectance used to create the colour signal. However, if a finite dimensional model is used to reconstruct reflectances, additional errors in the recovery of the two-bounce response can be expected, and could possibly outweigh improved shading fields a two-bounce model might provide.

1.2 Scope of Thesis

In this thesis a finite dimensional model of surface reflectance is presented as a tool for colour signal decomposition, where the colour signals result from interreflection of two planar surfaces by a single known light source. The decomposed signal is used to demonstrate an improvement to surface shape recovery methods that use image intensity as indicators of surface orientation.

The colour response bases used to decompose interreflected colours are not orthogonal, and are affected by many factors, such as illuminant spectra, filter sensitivities, and surface reflectances. The effects of changing these properties is explored by analysing the angles between basis vectors encompassed by interreflection between pairs of Munsell surfaces under two different illuminants and detected through three sets of filters.

As an integral part of the creation of these colour basis vectors, the finite dimensional model must reconstruct surface reflectances and their resulting colour signals. The error in decomposition is affected directly both by the model and by these reconstructions. The effects of the three sets of filters and two illuminants are statistically analysed for reconstruction of both one-bounce and two-bounce colour signals and their filter responses for pairs of Munsell surfaces.

Using test cases selected for their comparative performances with other pairs of interreflecting surfaces, colour reconstruction is performed in both synthetic and real images. The improvement from using recovered shading fields rather than image colour intensity is demonstrated using a simple shape reconstruction algorithm. The

feasibility of extending the one-bounce model to two bounces is also explored.

1.3 Thesis Outline

Chapter 2 provides the reader with a description of symbols and their descriptions used throughout this thesis. In addition, an introduction to colour signal composition, Lambertian surfaces, and methods used for measuring errors are presented.

Chapter 3 summarizes some of the more important research related to the study of interreflection, and that which forms much of the groundwork for this thesis.

Chapter 4 introduces the reader to finite dimensional models as applied to surface reflectances and illumination. The basis functions used to reconstruct surface spectral reflectances within the thesis are generated from 462 Munsell colours and compared with existing Munsell basis functions constructed by Cohen [4].

Chapter 5 describes the role of colour in interreflection. The one-bounce model is presented and the incentive for investigating a two-bounce model is discussed.

Chapter 6 contains a statistical analysis of multibounce colour space. For all pairwise combinations of the set of Munsell colours, the angles between the tristimulus vectors resulting from no-bounce and one-bounce colours are examined. The angles between the two-bounce vectors and the plane described by the no-bounce and one-bounce vectors are also investigated. Statistics of both sets of data are tabulated and observations made with respect to the one-bounce and two-bounce interreflection models.

Chapter 7 investigates the reconstruction of multibounce colours using a finite dimensional model of surface reflectance. A statistical analysis of both one-bounce and two-bounce colour reconstruction is performed.

Chapter 8 presents an algorithm for the analysis of interreflecting surfaces using a finite dimensional model of surface reflectance. From the results in previous chapters, three sets of test cases are chosen, and detailed analysis of the colour signal and tristimulus value reconstructions are made for each. The decomposition algorithm is then applied to synthetic images and real images for the illuminant, filters, and surfaces selected. The improvement the shading field can provide for surface shape reconstruction is also demonstrated.

Chapter 9 contains a disussion of the results obtained from previous chapters. The feasibility of both one-bounce and two-bounce models is examined, and some conclusions drawn. Possible areas for future research are suggested.

Chapter 2

Preliminaries

2.1 Symbol Definitions

The following table of symbol definitions provides a summary of notations commonly used throughout the thesis.

SYMBOL DEFINITION

λ Wavelength (*nm*). In this thesis λ ranges from 400 - 700 *nm*, in 10 *nm* increments.

$E(\lambda)$ Spectral power distribution of ambient light.

$S(\lambda)$ Spectral reflectance of a surface.

$S^A(\lambda)$ Spectral reflectance of surface *A*.

$S_i(\lambda)$ Spectral reflectance basis function *i*.

σ_i Coefficient of reflectance basis function *i*.

σ Vector of reflectance basis coefficients.

$C(\lambda)$ Colour signal, $C(\lambda) = E(\lambda)S(\lambda)$.

$C^A(x, \lambda)$ Colour signal emanating from surface *A* at location *x*.

$C^{(i)}(\lambda)$ i^{th} -bounce component of a colour signal resulting from interreflection.

$C^{(A)}(\lambda)$ $C^{(0)}(\lambda)$ specific to surface A .

$C^{(AB)}(\lambda)$ $C^{(1)}(\lambda)$ for either surface A or surface B , as it is symmetric.

$C^{(ABA)}(\lambda)$ $C^{(2)}(\lambda)$ specific to surface B .

$R_k(\lambda)$ Spectral sensitivity of k^{th} sensor (filter) class.

$R(\lambda)$ Matrix of spectral sensitivities.

ρ_k Response of k^{th} sensor, $\rho_k = \int C(\lambda)R_k(\lambda)d\lambda$.

ρ Vector of sensor responses.

$\rho^{(i)}$ Vector of sensor responses to colour signal $C^{(i)}(\lambda)$. Notation for no-bounce, one-bounce, and two-bounce responses for specific surfaces is identical to $C^{(i)}(\lambda)$. For example, the specific response $\rho^{(0)}$ for surface A is $\rho^{(A)}$.

Λ Matrix whose ki^{th} entry is $\int E(\lambda)S_i(\lambda)R_k(\lambda)d\lambda$.

$a_i^A(x)$ Coefficient of the i^{th} -bounce colour signal component emanating from surface A at position x .

β Opening angle of the two interreflecting planar surfaces A and B .

θ Angle between $\rho^{(0)}$ and $\rho^{(1)}$ in three-dimensional sensor response space.

ϕ Angle between $\rho^{(2)}$ and the plane described by $\rho^{(0)}$ and $\rho^{(1)}$.

2.2 Colour Signal Decomposition

Light reflecting from a surface results in a colour signal $C(\lambda)$ as the spectral product

$$C(\lambda) = E(\lambda)S(\lambda) \quad (2.1)$$

where $E(\lambda)$ is the spectral power distribution of the source illuminant, and $S(\lambda)$ is the surface spectral reflectance function at wavelength λ . A set of photoreceptors with sensitivity functions $R_k(\lambda)$, $k=1..x$ sensing the colour signal will produce responses

$$\rho_k = \int E(\lambda) S(\lambda) R_k(\lambda) d\lambda, \quad k=1..x \quad (2.2)$$

where the integral is taken over the visible spectrum. In humans, this is approximately 400 – 700 nm. In both humans and camera systems, the number of photoreceptors, x , is typically three. Using CIE standard observer colour-matching functions \bar{x} , \bar{y} , and \bar{z} , the resulting tristimulus values are identified as XYZ . However, these filter functions are not generally available to camera systems which employ red, green, and blue channel filters. These are identified as RGB .

2.3 Lambertian Surfaces

A common simplifying assumption for colour vision work is that surfaces appear equally bright regardless of viewing position. Such surfaces are uniformly diffusing, with the luminous intensity I_ϵ at each wavelength varying with the cosine of the angle ϵ from the surface normal ($\epsilon = 0$) [31]:

$$I_\epsilon = I_{\epsilon=0} \cos \epsilon. \quad (2.3)$$

Equation 2.3 is commonly referred to as *Lambert's Law*, and materials that satisfy this equation are said to be *Lambertian*.

An assumption used throughout this thesis is that surface reflectance functions are independent of viewing position, that is, they are Lambertian. In reality this is only approachable by some surfaces. A more realistic model of surface reflectance is described in [27].

2.4 Error Measurement

There are various methods with which one can obtain quantitative measurements of a mathematical model's accuracy and of experimental error. Here we will mainly be

concerned with the accuracy of reconstructed colour signals and tristimulus values. Methods have been selected that produce both quantized values of human perception, and those that are purely mathematical.

2.4.1 Just-Noticeable Difference

In much of this work we will want to compare ideal colour signals with those reconstructed by measurements of specific response filters. In order to obtain meaningful results, a measurement that corresponds to our own perception of differences between colours should be used. However, most colour coordinate systems have not been devised with this requirement in mind. As a typical example, consider the x, y chromaticity diagram based on the *1931 CIE Standard Observer* colour-matching functions, $x(\lambda), y(\lambda), z(\lambda)$ [1, 10, 31]. Humans perceive differences in colours on the x, y plane toward the blue and red corners of the diagram more easily than with greens and yellows [31]. Consequently, if we were to use some measure of *just-noticeable difference*, then one unit of this measure would correspond to a much larger Euclidian distance along the x, y plane in the green area than the blue area of the diagram.

In an attempt to reconcile Euclidian distances with human perceptual differences, a transformation from CIE standard observer XYZ tristimulus coordinates to the more uniform CIE $L^*u^*v^*$ colour coordinate system was developed as follows [31]:

$$\begin{aligned} L^* &= 116\left(\frac{Y}{Y_n}\right)^{1/3} - 16 \\ u^* &= 13L^*(u' - u'_n) \\ v^* &= 13L^*(v' - v'_n) \end{aligned} \tag{2.4}$$

where

$$\begin{aligned} u' &= \frac{4X}{X + 15Y + 3Z} \\ v' &= \frac{9Y}{X + 15Y + 3Z} \\ u'_n &= \frac{4X_n}{X_n + 15Y_n + 3Z_n} \\ v'_n &= \frac{9Y_n}{X_n + 15Y_n + 3Z_n} \end{aligned}$$

and X_n, Y_n, Z_n are the tristimulus values produced by one of the CIE standard illuminants, with Y_n scaled to 100. The standard illuminant used in this thesis is D_{65} .

The perceptible difference between two colours can thus be represented by the following *colour-difference formula*:

$$\Delta E = \sqrt{(\Delta L^*)^2 + (\Delta u^*)^2 + (\Delta v^*)^2}. \quad (2.5)$$

To evaluate differences independent of colour intensity (ie. luminous value L), the XYZ are scaled such that their Y components are equivalent, and thus $\Delta L = 0$.

2.4.2 Angle Between Vectors

Tristimulus values can be viewed as vectors in three-dimensional colour space, with coordinate systems corresponding to the filter system used (eg. $RGB, XYZ, L^*u^*v^*$). With two vectors rooted at the origin of their coordinate system, their difference can be measured as the angle between them. With one vector considered the ideal and the second as a reconstruction, the angle between them can be used as a measurement of error.

If we consider two vectors \mathbf{a} and \mathbf{b} , then the angle θ they subtend is described in the equation

$$\cos \theta = \frac{\mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|}. \quad (2.6)$$

Here we will be representing angles in degrees.

2.4.3 Root Mean Square

Root Mean Square (RMS) is a measure of the fit of an approximated curve to an actual curve, and can be used as an alternative measure of error for both colour signals and tristimulus values. Given a theoretical vector \mathbf{v}^0 and its approximation \mathbf{v} , we define the RMS error as

$$RMS \equiv \sqrt{\frac{(\mathbf{v}^0 - \mathbf{v}) \cdot (\mathbf{v}^0 - \mathbf{v})}{\mathbf{v}^0 \cdot \mathbf{v}^0}} \times 100\%. \quad (2.7)$$

Chapter 3

Previous Work

3.1 Radiosity

While we are ultimately interested in decomposing colour signals of interreflecting surfaces, it is helpful to gain some insight into the forward graphics problem: modeling interreflection between object surfaces for image rendering. There are two basic types of techniques used to approach this problem. *Ray tracing* typically simulates interreflection between specular surfaces, tracing light rays between light sources and the eye, while *radiosity* methods simulate view-independent interreflections in diffuse environments. To provide realistic renderings, variations of techniques that model both specular and diffuse interreflections have also been used. Alternatively, the results of ray-tracing and radiosity methods can be combined. Since the analysis described in this thesis is restricted to interreflection between Lambertian surfaces, we are interested primarily in radiosity in diffuse environments.

First introduced by Goral *et al* [15], radiosity methods assume the conservation of light energy in a closed environment. All energy emanating from a surface must then be either reflected from or absorbed by another surface. The *radiosity* of a surface is the rate at which energy leaves that surface, and is the sum of the rates at which it emits energy and reflects it from other surfaces. For each wavelength this is described

by the following *radiosity equation* [10]:

$$B_i = E_i + \rho_i \sum_{j=0}^{n-1} B_j F_{ji} \frac{A_j}{A_i}. \quad (3.1)$$

B_i and B_j are the radiosities (typically measured in W/m^2) of surface patches i and j having areas A_i and A_j . E_i is the rate at which light is emitted from patch i , in the same units as B_i and B_j . ρ_i is the dimensionless reflectivity of patch i , and F_{ji} is the dimensionless *configuration factor*, specifying the fraction of energy leaving patch j and arriving at patch i , and is a function of image geometry. In chapter 8, a modified form of equation 3.1 will be used to construct a synthetic, one-dimensional image of two interreflecting surfaces.

For the graphics problem of image rendering, three-dimensional geometry, surface reflectance properties, and signal spectra are known. The task is to project this information onto a two-dimensional image. Conversely, computer vision algorithms commonly attempt to extract three-dimensional information from two-dimensional images. This is a notoriously underconstrained problem, such that researchers must search for constraints where they exist, and make assumptions where they do not.

3.2 Spectral Methods

3.2.1 A Qualitative Analysis

A useful qualitative analysis of interreflection between two surfaces in the presence of specular (or *surface*) reflections was presented by Novak [25]. She builds upon Shafer's *dichromatic reflection model* [27], separating the colour signal C emanating from a surface into factors due to incident illumination E and reflection characteristics S_b and S_s (*body* and *surface* reflection):

$$C(\lambda, i, c, g) = E(\lambda, i, c, g)[S_b(\lambda, i, c, g) + S_s(\lambda, i, c, g)] \quad (3.2)$$

where λ is wavelength and i , c , and g are geometric factors. In the *primary* (no-bounce) case, E consists only of the illumination from the light source. In the *secondary* (one-bounce) case, with both objects obeying the dichromatic reflection model,

the secondary reflection is broken into four components: $S_b^A S_b^B$, $S_b^A S_s^B$, $S_s^A S_b^B$, and $S_s^A S_s^B$. Together with the two primary components, there are a total of six reflection components. In general, for n -ary reflection there are $2^{n+1} - 2$ of these full-spectrum reflection components.

Although this model is useful for understanding the various components of interreflecting surfaces, obtaining a solution to unknown surfaces and illuminants is difficult, since the number of unknowns is greater than the number of equations a three-sensor imaging system can provide. As a result, Novak concentrated on obtaining histograms of synthetic and actual images where interreflection and specularities occur, and identifying the various components according to the model.

3.2.2 Colour Constant Analysis

Researchers Funt, Drew, and Ho [13] presented an algorithm for obtaining colour constancy between two semi-infinite planar Lambertian surfaces in the presence of interreflection, solving two problems simultaneously. The first, and the object of the paper, is to recover the spectral power distribution of the illuminant from which surface spectral reflectances can be recovered independently of the light source. An integral part of this algorithm is the presence of interreflection. As a byproduct of the process, the geometric (or *shading*) component of interreflection is extracted.

This work is based on previous colour constancy research by Maloney and Wandell [19, 20, 30], using finite dimensional models of illumination and surface reflectance. A weakness of the algorithm presented by Maloney and Wandell is the requirement that the number of sensor classes be greater than the number of reflectance basis functions. However, since human and other vision systems are generally constrained by three classes of sensors, at least three basis functions are required to adequately model most surfaces [4]. Using interreflection, Funt *et al* were able to require that the number of sensor classes be $p = (2n + m)/3$, where m and n are the number of basis functions used to model illumination and reflectance respectively. This can be satisfied with $p = n = m = 3$.

The initial relationship between the illuminant, surface reflectances, and inter-reflection is quite complex, and involves three nonlinear equations. To obtain a more easily solved form, a one-bounce model of interreflection is used (where contributions of colour signals reflected more than once from each surface are assumed to be negligible). Furthermore, by using an iterative scheme, the nonlinear equations can be reduced to a sequential set of linear equations. Ultimately, a solution for the basis weights of the illuminant ϵ_i , both surfaces $\sigma_j^{(1)}$ and $\sigma_j^{(2)}$, and configuration factor α converge. To reduce the number of variables, and to stabilize convergence, the constraint $\epsilon_1 = 1$ is imposed. Thus, while the illuminant spectrum is obtained, its overall intensity remains unknown.

Although it was shown that convergence occurs, when the underlying finite dimensional model poorly describes both surfaces the algorithm is slow to converge, and does a poor job of recovering surface reflectances. Likewise, if the illuminant is poorly modeled, poor results can also be expected.

3.3 Reduced Component Methods

3.3.1 Intensity and Interreflection

Interreflection of surfaces has a confounding effect on surface shape recovery using shape-from-shading methods. These methods typically assume Lambertian surfaces, so that surface orientation is directly related to relative intensity. As such, shape-from-shading methods generally use grey-scale images, with pixel intensity representing an object's shading. The surface appears brighter where interreflection occurs, incorrectly indicating an orientation closer to that perpendicular to the illuminant. As a consequence, concave surfaces tend to appear shallower than they actually are. A solution to this problem was proposed by Nayar *et al* [22]. Using grey-scale images, initial estimates of shading and surface shape were obtained by photometric stereo. By modeling intensity as radiosity at each facet, as in equation 3.1, the interreflection component was removed and new estimates of shading and resulting shape were obtained. These estimates were iteratively refined to converge to the correct shape and

shading.

Unfortunately, Nayar’s original approach disregards surface spectral reflectance [12]. For example, consider a concave edge formed by two surfaces with narrow-band spectral reflectances at opposite ends of the visible spectrum, say red and blue. The product of their reflectance curves will be zero at every wavelength, thus any light reflected from one surface is completely absorbed by the other. However, the colour intensity of each surface is nonzero. Unless either narrow-band sensors or illuminant are employed, Nayar’s approach will assume interreflection occurs and incorrectly “deepen” the actual shape. Thus we can see that colour plays an important role in interreflection, and cannot be ignored.

In a subsequent approach to the problem of interreflecting multispectral coloured signals, Nayar and Gong [21] presented a modified shape recovery algorithm, where the incident light is assumed to be monochromatic. Recovery is thus performed for three separate narrow-band channels, again using an iterative approach.

3.3.2 Intersection of Planes

More recent work by Funt and Drew [12] uses three-component *RGB* analysis of two interreflecting surfaces. They employ a one-bounce model to decompose each surface facet into its no-bounce and one-bounce colours, as a linear combination of the two. For Lambertian surfaces under a spectrally invariant illuminant, these colours will lie on a plane in three-dimensional *RGB* colour space (see figure 5.3).

Integral to their method is the observation that the one-bounce colour is identical to both surfaces, and the two planes must intersect along the vector defined by its *RGB* values. Using singular value decomposition (SVD), the planes are detected for each surface and their intersection determined. The no-bounce colour is assumed to exist at some point on each surface. With both colours known, decomposition of each facet of the surfaces can quickly be performed.

While this intersection of planes works well in synthetic images, using the same method presented some problems with real images. For instance, a SVD of the data showed that colours did not lie as close to a plane as expected. The authors speculate

that this may be due to specularities from the surfaces not accounted for by the model, and by other factors such as signal noise and stray light. Since most colours lie clustered together much closer to the no-bounce component than the one-bounce component, extrapolation becomes prone to significant error resulting from incorrect identification of the planes.

Chapter 4

Finite Dimensional Models

Finite-dimensional models are frequently used to describe illumination and surface spectral reflectances through weighted sums of basis functions (generally represented in the discrete domain as vectors) [6, 13, 14, 20, 30]. Such a model can represent discretely sampled signals with reasonable accuracy using significantly fewer coefficients than the number of signal samples. For example, a signal consisting of 31 samples (400–700 *nm* in 10 *nm* increments) might be well modeled using three basis vectors.

4.1 Background

4.1.1 Modeling Surface and Illuminant Spectra

Most natural surface spectral reflectances can be well-modeled by a finite dimensional linear model of low dimension [20, 30]. Using a set of linearly independent basis functions S_i , $i=1..m$, we can model surface spectral reflectance $S(\lambda)$ as

$$S(\lambda) \simeq \sum_{i=1}^m S_i(\lambda) \sigma_i \quad (4.1)$$

and the illuminant spectral power distribution as

$$E(\lambda) \simeq \sum_{j=1}^n E_j(\lambda) \sigma_j. \quad (4.2)$$

Substituting equation 4.1 into equation 2.2 we get

$$\rho = \Lambda \sigma \quad (4.3)$$

where ρ is a column vector formed by the quantum catches of the x sensor classes, σ is a column vector containing the m reflectance basis coefficients, and Λ is an $x \times m$ matrix whose ki^{th} entry is $\int E(\lambda)S_i(\lambda)R_k(\lambda)d\lambda$. We know $R_k(\lambda)$ and have chosen $S_i(\lambda)$. Assuming we also have $E(\lambda)$, Λ will contain only known values, so it can be precomputed. We thus need only to solve a set of m equations in m unknowns to obtain σ_i . If $E(\lambda)$ is unknown then equation 4.3 is unconstrained and we must search for alternate methods for obtaining unique solutions to σ_i .

A technique proposed by Maloney and Wandell [19, 20, 30] assumes surface reflectance $S(\lambda)$ can be well-modeled with one less degree of freedom than the number of receptor classes. That is, $m = x - 1$. If we also represent the illuminant by a linear model as in equation 4.2 where $n = x$, then sampling at s unique spatial locations provides $s(m + 1)$ data values and sm unknowns. When $s > n$ samples are obtained, there are more data values than unknowns, which may then be solved for.

Although this method provides solutions to both unknown illuminant and reflectance spectra, the restriction on surface reflectance may be unreasonable. Consider the case where the number of sensor classes x is restricted to three, as in human photopic vision and most camera systems. $S(\lambda)$ would then have to be modeled in two dimensions, such that only surfaces with extremely smooth variations over visible wavelengths are modeled well.

4.1.2 Basis Functions

For any finite dimensional treatment of spectra we require a set of basis functions (or vectors in the discrete domain) from which a surface or illuminant may be modeled. While there are various ways in which one can select this set, we would like to obtain those that provide the best approximation to the signals they will be modeling. An obvious choice may seem to be the first three Fourier functions (a constant function, a sine wave, and a cosine wave). However, can find an alternative set that provides a better representation of the data we expect to see.

Dimensionality of Surfaces and Illuminants

If there are no constraints on signals we represent using finite dimensional models then the number of basis functions could be arbitrarily large. Fortunately, most naturally occurring colour signals and reflectances are band-limited, varying smoothly over the visual range. The number of basis functions required to obtain reasonable approximations of natural surfaces is thus quite small. Buchsbaum and Gottschalk [3] suggested that as few as three functions can provide metamers to every Munsell chip spectral reflectance, and cover a large central area of the CIE chromaticity diagram. Cohen [4] showed that three properly selected basis reflectances encapsulated 99.2% of the overall variance of 150 randomly selected Munsell chips. These same three basis reflectances were shown to provide good approximations to 337 of the naturally occurring objects measured by Krinov [18, 20].

Singular Value Decomposition

Given a set of data (eg. colour signals or spectral reflectances), we would like to be able to condense and summarize this information as a linear combination of a set of basis vectors. *Singular Value Decomposition* will provide us with such a set. It is based on the following theory [26]:

Any $M \times N$ matrix A , whose number of rows M is greater than or equal to its number of columns N , can be written as the product of an $M \times N$ column-orthogonal matrix U , an $N \times N$ diagonal matrix W with positive or zero coefficients, and the transpose of an $N \times N$ orthogonal matrix V .

$$A = U W V^T . \quad (4.4)$$

Although the SVD can provide solution vectors to singular sets of equations (fewer equations than unknowns) and linear least squares solutions to overconstrained sets (fewer unknowns than equations), for our purposes the rows of the resulting matrix V^T contain an ordered set of basis vectors. Furthermore, we can obtain the cumulative variance $v(k)$, of the fit of the first k basis vectors from the elements w_{ij} of diagonal

matrix W as

$$v(k) = \frac{\sum_{i=1}^k w_{ii}^2}{\sum_{j=1}^N w_{jj}^2} \quad (4.5)$$

While the SVD minimizes the total variance with respect to the entire input data set, there is no guarantee that the fit will be good for a particular sample. To make this method work best for arbitrary inputs, it is important to provide it with as many representative signals as possible so that these will more likely be similar to those within the sample set, and consequently more likely to be well-modeled.

4.2 Basis Function Calculation

A frequently chosen set of reflectance basis functions is that provided by Cohen [4]. Of the spectral reflectance curves of 433 colour chips from the Munsell Book of Colour made available by Nickerson [24], Cohen randomly chose 150 of these (the capacity of his computer) to perform a linear component analysis.

Without the capacity constraints of Cohen's computer, it is reasonable to perform an analysis on the entire set of data provided by Nickerson rather than restrict ourselves to 150. This was done, and the resulting basis vectors used for multibounce colour reconstruction in the process of decomposing the colours of interreflecting surfaces.

4.2.1 Implementation

The entire set of Nickerson data as provided in Brainard and Wandell's Colour Analysis Package [2] consists of 462 sets of spectra, sampled from 380 – 770 nm in 10 nm increments. This was truncated to the range 400 nm – 700 nm (31 samples per spectrum). Singular value decomposition was performed on the resulting 31×462 matrix. The first three rows of the 31×31 orthogonal matrix V^T were then extracted as the basis vectors.

4.2.2 Results

The analysis performed by Cohen resulted in the first three basis vectors shown in figure 4.1. The singular value decomposition performed here provided those shown in figure 4.2. There is an expected similarity between the sets, with minor exceptions. One might observe that the basis vectors of figure 4.2 approach those of a Fourier basis set consisting of a constant signal, a sine wave, and a cosine wave.

The SVD analysis performed on the 462 samples captured a greater proportion of the cumulative variance than Cohen's 150 random samples. The Cohen vectors obtained 92.7%, 97.25%, and 99.18% for one, two, and three vectors respectively. SVD, on the other hand, produced values of 92.52%, 97.45%, and 99.63%. We must be careful when making comparisons between the two sets of results since Cohen performed his analysis with different endpoints (380 - 770 *nm*). Because we will be using a finite dimensional model with three degrees of freedom within the smaller range of wavelengths, the latter set of vectors will be selected.




Figure 4.1: Cohen's first three basis vectors obtained through linear component analysis of 150 randomly selected Munsell colour chips.




Figure 4.2: The first three basis vectors obtained through singular value decomposition of 462 Munsell colour chips.

Chapter 5

Models of Interreflection

For simplicity, consider an edge formed by two Lambertian surfaces having distinct spectral reflectances. Let the angle β subtended by the two surfaces be acute, so that interreflection will occur. Figure 5.1 depicts this situation.

Let surface A have spectral reflectance $S^A(\lambda)$, and surface B have spectral reflectance $S^B(\lambda)$. If we have three classes of receptors with spectral response functions $R_k(\lambda)$, $k=1..3$, and an illuminant with spectral power distribution $E(\lambda)$, then we will call the response of the k^{th} receptor class ρ_k . Restricting ourselves to a single location x on surface A , this response is the result of direct and interreflected colour signals reflected from that point:

$$\rho_k^A(x) = \int C^A(x, \lambda) R_k(\lambda) d\lambda \quad (5.1)$$

where the colour signal $C^A(x, \lambda)$ is:

$$C^A(x, \lambda) = \sum_{i=0}^{\infty} a_i^A(x) E(\lambda) S^A(\lambda)^{\lceil \frac{i+1}{2} \rceil} S^B(\lambda)^{\lfloor \frac{i}{2} \rfloor} . \quad (5.2)$$

Letting $a_i^A(x) = 1$ we can observe the relative spectral distributions of the multibounce colour signal components as i increases in figure 5.2. For $i = 0$, the colour signal contains no contributions from interreflection. The coefficient $a_0^A(x)$ thus represents the intensity variation (*shading field*) at position x of side A . If we maintain the assumption that the surfaces are Lambertian then this shading field varies with the cosine of surface orientation. The interpretation of the remaining coefficients is not as




Figure 5.1: Interreflection between two coloured surfaces with spectral reflectances $S^A(\lambda)$ and $S^B(\lambda)$, showing no-bounce colour signals $C^{(A)}$ and $C^{(B)}$, and one-bounce colour signal $C^{(AB)}$.

straightforward. There are a number of factors that can affect the magnitude of the interreflected colour signals, such as the surface orientation at position x with respect to the illuminant, the shading of both surfaces, and geometric factors such as possible self-reflection from concave surfaces and occlusion of points between the surfaces.

It is obvious that simplifying assumptions are required to calculate the colour signal, since we cannot solve for an infinite number of coefficients. A common assumption is the rapid attenuation of the signal as it is reflected between surfaces. Reference to figure 5.2 confirms the feasibility of such an assumption. This leads to a *finite-bounce* interreflection model.

5.1 One-bounce model

The one-bounce model of interreflection (for which $i = 1$ in equation 5.2) is based on the assumption that there are no significant colour signal contributions beyond that obtained from the weighted sums of the illuminant directly reflected from a surface




Figure 5.2: Sample multibounce colour components of signal $C^A(x, \lambda)$ for increasing number of bounces i , resulting from interreflection between Munsell colours 10BG6/4 (blue-green) and 5P5/6 (purple) under standard illuminant D_{65} .

and from the resulting colour signal reflected from a second surface [6, 8, 13, 12]. The combined signal measured from surface A would thus be defined as

$$C^A(x, \lambda) \simeq E(\lambda) \left(a_0^A(x) S^A(\lambda) + a_1^A(x) S^A(\lambda) S^B(\lambda) \right). \quad (5.3)$$

Substituting equation 5.3 into equation 5.1 gives us

$$\rho_k^A(x) \simeq a_0^A(x) \rho_k^{(A)} + a_1^A(x) \rho_k^{(AB)} \quad (5.4)$$

where

$$\rho_k^{(A)} = \int E(\lambda) S^A(\lambda) R_k(\lambda) d\lambda$$

and

$$\rho_k^{(AB)} = \int E(\lambda) S^A(\lambda) S^B(\lambda) R_k(\lambda) d\lambda.$$

The responses ρ_k^B for side B are formed in similar manner. In the general case, $\rho^{(0)}$ and $\rho^{(1)}$ will refer to the no-bounce and one-bounce receptor responses, and $C^{(0)}(\lambda)$ and $C^{(1)}(\lambda)$ will refer to their respective colour signals.

According to this model, the colour responses $\rho(x)$ obtained from the two surfaces are approximately linear combinations of their respective $\rho(0)$ and $\rho(1)$. Thus, if the $\rho(x)$ are plotted in three-dimensional vector space, they should fall on a plane defined by $\rho^{(0)}$ and $\rho^{(1)}$ [13, 27]. Furthermore, since we are only looking at positive signals, these responses will lie between these two vectors. An important assumption here is that $\rho^{(0)}$ and $\rho^{(1)}$ are linearly independent. In general, this problem can be avoided by ensuring that neither surface has a constant (ie. grey or white) spectral reflectance.

Note that $\rho^{(1)}$ is identical for both surfaces, representing the response for a colour signal reflected once from each side before striking the receptors. In fact, referring to equation 5.2, we see that for all odd i :

$$\left\lceil \frac{i+1}{2} \right\rceil = \left\lfloor \frac{i}{2} \right\rfloor$$

and thus

$$S^A(\lambda)^{\lceil \frac{i+1}{2} \rceil} S^B(\lambda)^{\lfloor \frac{i}{2} \rfloor} = S^A(\lambda)^{\lfloor \frac{i}{2} \rfloor} S^B(\lambda)^{\lceil \frac{i+1}{2} \rceil}. \quad (5.5)$$

Funt and Drew [13] observed the symmetry of the one-bounce colour, noting that the colour response plane of surfaces A and B will intersect along $\rho^{(1)}$.

A reasonable assumption is that we know the no-bounce colours, or at least know the intensity of response for each class of receptor. In other words, there exists an identifiable point on each surface having the same spectral reflectance and light source where interreflection is nonexistent or negligible.

Obtaining the one-bounce colour is not as simple. Since it only exists in the company of the no-bounce colour it cannot be measured directly. Even colours very close to an interreflecting edge consist mainly of the no-bounce component. One method [12] uses the knowledge that colours of a surfaces fall close to the plane described by $\rho^{(0)}$ and $\rho^{(1)}$. The intersection of the planes for two surfaces would thus be the common vector $\rho^{(1)}$. A weakness of this method is that colour measurements are grouped much closer to $\rho^{(0)}$ so that noise, specularities, and other factors contributing to response measurement may translate to large imprecisions and errors as planes are projected toward their intersection. Contributions of other multibounce colour components (especially near the edge) can also result in measurements further from the plane than those expected with a simple one-bounce model.

If both $E(\lambda)$ and $R(\lambda)$ are known, then projection of these planes is unnecessary. Spectral reflectances of each surface can be reconstructed using equation 4.3, so that any multibounce component can be calculated directly. Details of this method are investigated in chapter 7.

Once $\rho^{(1)}$ has been estimated, we can use equation 5.4 with $k = 3$ to obtain solutions for our shading and interreflection coefficients a_0 and a_1 . Since there are three equations with two unknowns, it can be solved using a least squares fit. We can view this as projecting $\rho(x)$ orthogonal to, and onto, the plane described by $\rho^{(0)}$ and $\rho^{(1)}$, and obtaining the coefficients of this vector.

Although the use of a one-bounce model of interreflection can provide a superior estimate of surface shading than colour intensity, it begins to break down close to the surface edges. As the edge between surfaces is approached, contribution from other multibounce components increases. Since the one-bounce model attributes all image intensity to the two colour components of the model, an increased coefficient error results. This may take the form of either an increased or decreased shading field, depending on the direction of the other multibounce components relative to the first two.

5.2 Two-bounce model

A possible solution to the limitations of the one-bounce model is to extend it to two bounces, including the effect of light rays that bounce off one surface, then the second surface, and back to the original surface again before contacting the receptors. The colour signal obtained from surface A would thus be defined as:

$$C^A(x, \lambda) \simeq E(\lambda) \left(a_0^A(x) S^A(\lambda) + a_1^A(x) S^A(\lambda) S^B(\lambda) + a_2^A(x) S^A(\lambda)^2 S^B(\lambda) \right) . \quad (5.6)$$

Substituting equation 5.6 into equation 5.1 thus gives us

$$\rho_k^A(x) \simeq a_0^A(x) \rho_k^{(A)} + a_1^A(x) \rho_k^{(AB)} + a_2^A(x) \rho_k^{(ABA)} \quad (5.7)$$

where $\rho_k^{(A)}$ and $\rho_k^{(AB)}$ are as described in equation 5.4 and

$$\rho_k^{(ABA)} = \int E(\lambda) S^A(\lambda)^2 S^B(\lambda) R_k(\lambda) d\lambda .$$

The colour response $\rho_k^B(x)$ would be solved similarly. As with $\rho^{(0)}$ and $\rho^{(1)}$, for the general case we will refer to $\rho^{(ABA)}$ and $\rho^{(BAB)}$ as $\rho^{(2)}$. Note that, unlike $\rho^{(1)}$, $\rho^{(2)}$ is not symmetrical and must be obtained independently for each surface.

Since equation 5.7 with three receptor classes forms a linear set of three equations in three unknowns, we may be able to find a_1 , a_2 , and a_3 using Gauss-Jordan elimination or any other method for solving linear equations.

One drawback to this model is that we cannot be certain the two-bounce colour associated with a surface is linearly independent of its no-bounce and one-bounce colours. That is, in three-dimensional space $\rho^{(2)}$ may lie on the plane described by $\rho^{(0)}$ and $\rho^{(1)}$. While the assumption that neither surface has a constant (ie. ideal white) spectral reflectance is made to avoid linear dependencies between the first two colours, there are no apparent assumptions that will avoid it between these and $\rho^{(2)}$.

Although we cannot enforce linear independence, we can determine whether it holds. In addition to observing the angle θ between $\rho^{(0)}$ and $\rho^{(1)}$, we must also look at the angle ϕ between $\rho^{(2)}$ and the vector resulting from its projection onto the plane described by $\rho^{(0)}$ and $\rho^{(1)}$ (see figure 5.3). For each test, the angle should be greater than some threshold approaching 0° . If the former condition is violated, then interreflection cannot be measured. If the latter condition is violated, then a one-bounce model can still be used. These tests can be performed independently for each surface, and the appropriate model used for each.

Since a two-bounce model expects to see colours in three-dimensional space rather than on a plane, we can no longer use the intersection of two planes to find the one-bounce colour. Furthermore, there seems to be no extension to this method that would allow us to find the two-bounce colour. The known-illuminant assumption, on the other hand, allows us to separate a colour response into its illuminant, filter, and reflectance constituents, so any number of multibounce colours can be obtained. However, the accuracy of these colours diminishes for each bounce, since the spectral reflectance reconstruction of each surface is unlikely to be precise, and multiplication of signals will increase this error.

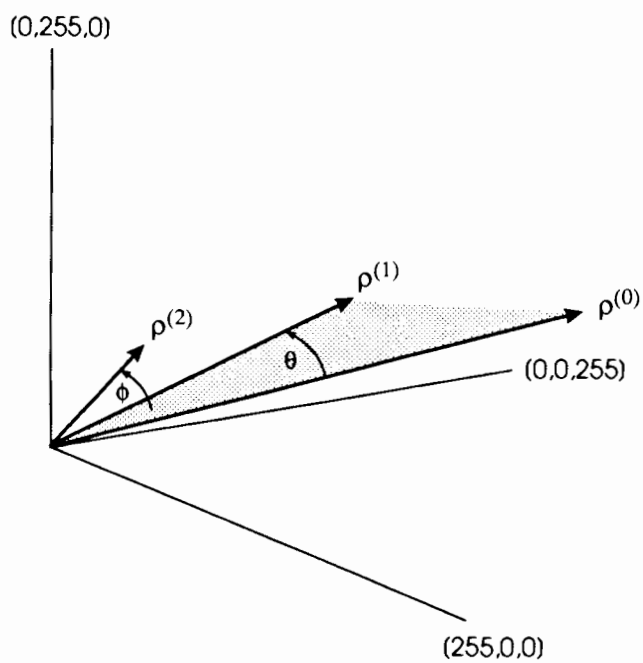


Figure 5.3: Angle θ between $\rho^{(0)}$ and $\rho^{(1)}$, and angle ϕ between $\rho^{(2)}$ and the plane described by $\rho^{(0)}$ and $\rho^{(1)}$ (shaded area) in three-dimensional *RGB* space.

Chapter 6

Multibounce Colour Space

Since the $\rho^{(0)}$, $\rho^{(1)}$, $\rho^{(2)}$ vectors have not been selected as orthogonal bases, but rather are determined by the two interreflecting surfaces, it is useful to analyse the space they describe. Statements can then be made regarding the relative performances of filters and illuminants, and on the suitability of the one-bounce and two-bounce models of interreflection. Here we will examine the angle θ between $\rho^{(0)}$ and $\rho^{(1)}$, and the angle ϕ between $\rho^{(2)}$ and the plane described by $\rho^{(0)}$ and $\rho^{(1)}$.

6.1 Implementation

Using the spectral power distributions of a 3200K Planckian radiator (fig 6.1), all pairwise combinations of 462 Munsell surface spectral reflectances [2], and each of the three sets of filters with response curves shown in figures 6.2-6.4 the theoretical ρ vectors were calculated. For surface A , $\rho^{(0)}$ and $\rho^{(1)}$ were obtained as in equation 5.4, and $\rho^{(2)}$ as in equation 5.7. Values for surface B were obtained similarly.

The angle θ was calculated between vectors $\rho^{(0)}$ and $\rho^{(1)}$ for each surface. The angle between $\rho^{(2)}$ and the plane was obtained by first projecting it onto the plane to obtain ρ' . ϕ was calculated as the angle between these two vectors. Statistics were then gathered for all θ and ϕ .

To observe the effect of $E(\lambda)$ on the results, the experiment was repeated using the D_{65} standard illuminant. The spectral power distribution of this light source is

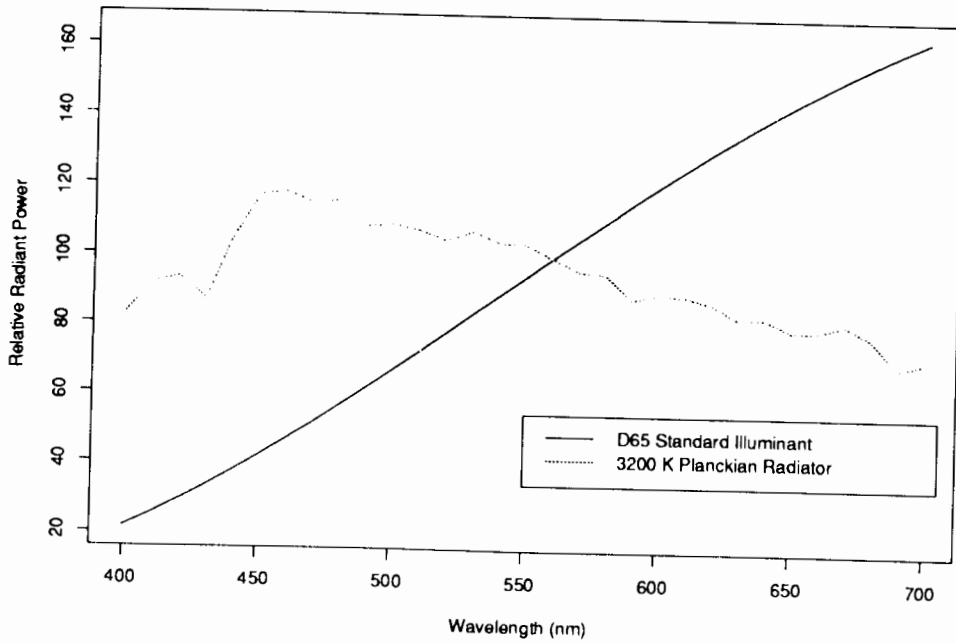


Figure 6.1: Relative radiant power distributions of the D_{65} standard illuminant and a 3200K Planckian (blackbody) radiator.




Figure 6.2: Relative response curves for CIE 1931 standard observer filters.




Figure 6.3: Relative response curves for Sony DXC-151 filters.




Figure 6.4: Relative response curves for Kodak #47b (blue), #58 (green), and #25 (red) filters.

<i>Response Filters</i>	<i>min.</i>		<i>max.</i>		<i>mean</i>		<i>std. dev.</i>	
	θ	ϕ	θ	ϕ	θ	ϕ	θ	ϕ
Std. Observer	0.0	0.0	44.7	22.9	8.9	3.6	6.2	3.5
Sony DXC-151	0.0	0.0	44.5	25.5	9.6	3.3	6.6	3.5
Kodak Filters	0.0	0.0	57.8	30.9	12.3	4.2	8.6	4.3

Table 6.1: Colour vector angles θ and ϕ under 3200K illuminant.

<i>Response Filters</i>	<i>min.</i>		<i>max.</i>		<i>mean</i>		<i>std. dev.</i>	
	θ	ϕ	θ	ϕ	θ	ϕ	θ	ϕ
Std. Observer	0.0	0.0	52.1	20.6	11.5	4.1	8.0	3.1
Sony DXC-151	0.0	0.0	48.7	15.7	13.2	4.6	8.4	3.3
Kodak Filters	0.0	0.0	60.2	22.3	16.6	6.2	10.4	4.3

Table 6.2: Colour vector angles θ and ϕ under D_{65} illuminant.

much flatter, yet less smooth, than that of a 3200K blackbody.

6.2 Results

The colour vector angles resulting from the application of a 3200K illuminant to all pairwise combinations of interreflected surfaces are recorded in table 6.1. Those resulting from illumination by the D_{65} standard illuminant are shown in table 6.2.

Under both illuminants, the Kodak filters result in significantly greater mean angles than both the Standard Observer and Sony. Referring to the filter response curves of figures 6.2, 6.3, and 6.4, we observe that the Kodak filters have less overlap between responses than the others. It would be interesting to apply *spectral sharpening* [9] to the filters to investigate the effect the resulting “sharpened” filters have on θ and ϕ .

Comparison of the mean angles under the two illuminants shows much improved results from using the D_{65} standard illuminant. As Forsyth observed [11], the gamut of photoreceptor responses is restricted under a given illuminant. For instance, a light with its power distributed close to the red end of the visible spectrum would not produce a strong blue response for any surface. Such is the case with the 3200K illuminant, whereas the D_{65} illuminant has a much flatter distribution. Since the

angles θ and ϕ are measured between responses in three dimensional colour space a restricted gamut would reduce the range of possible angles under that illuminant.

Regardless of the filters and illuminants used, the minimum angles for θ and ϕ are 0^0 , confirming our expectations that our models of interreflection will fail for some surface combinations. In particular, when at least one surface has a constant spectral reflectance curve we will observe these angles. However, under all configurations the mean and maximum angles provide motivation to investigate both the one-bounce and two-bounce models.

As we see from these results, the vector space for decomposition is far from orthogonal. A surface whose θ is close to 0^0 will be unable to be analysed with a one-bounce model of interreflection. Likewise, a surface whose ϕ is close to 0^0 will be unable to be analysed by a two-bounce model. As the mean for ϕ is much smaller than θ , we can expect there to be surfaces that cannot be analysed with a two-bounce model, yet can still be analysed with a one-bounce model. The thresholds for both θ and ϕ must rely on the accuracy of the finite dimensional model of surface reflectance, and (for the case of real images), experimental precision.

Chapter 7

Reconstructing Multibounce Colours

Let us now revisit equation 4.3. As stated earlier, if we know $E(\lambda)$ then the matrix Λ can be precomputed. With 3 sensor classes and 3 basis vectors this will be a 3×3 matrix. To obtain a solution to σ for any set of tristimulus values ρ , we can simply augment Λ with ρ and solve the resulting augmented matrix. Knowing σ allows us to reconstruct a finite-dimensional approximation to $S(\lambda)$. Doing so for surfaces A and B would provide us with $S^A(\lambda)$ and $S^B(\lambda)$. Knowing $E(\lambda)$ and these two surface reflectances allows us to construct any number of multi-bounce colour components, limited only by the precision of our reflectance estimates.

If both $E(\lambda)$ and $R_k(\lambda)$ contribute no errors (ie. are known precisely), then errors in $S(\lambda)$ will be due to the finite dimensional model of surface reflectance through limitations on the number of basis vectors $S_i(\lambda)$ and through our choice of these vectors. Of course, outside of a synthetic environment there are other contributing factors such as noise and measurement errors.

To measure the effect of the reconstruction of $S^{(A)}$ and $S^{(B)}$ on the recovery of receptor responses to each i -bounce colour signal we construct each i^{th} signal component $C^{(i)}(\lambda)$ of 5.2 separately, where $C^{(i)}(\lambda)$ is:

$$C^{(i)}(\lambda) = E(\lambda) S^A(\lambda) \left[\frac{i+1}{2} \right] S^B(\lambda) \left[\frac{i}{2} \right]. \quad (7.1)$$

The response of the k^{th} receptor class p_k to $C^{(i)}(\lambda)$ is:

$$\rho_k^{(i)} = \int C^{(i)}(\lambda) R_k(\lambda) d\lambda \quad (7.2)$$

We construct the theoretical $\rho^{(i)}$ as a three component vector for which $E(\lambda)$, $R(\lambda)$, $S^A(\lambda)$, and $S^B(\lambda)$ are known. The estimated $\rho^{(i)}$ is constructed similarly, with the exception that $S^A(\lambda)$ and $S^B(\lambda)$ are obtained through a finite dimensional model of surface reflectances.

The accuracy of the estimated $C^{(i)}(\lambda)$ from which $\rho^{(i)}$ is measured varies according to that of the the two reconstructed reflectances $S^A(\lambda)$ and $S^B(\lambda)$. While the two reconstructed $C^{(0)}(\lambda)$ are metamers of the actual no-bounce colour signals, reconstructed $C^{(i)}(\lambda)$ for $i \geq 1$ will not, in general, be metamers of product signals. To examine the effect of the difference between the theoretical onebounce colour and that produced through reconstruction of surface reflectances, all possible combinations of the set of 462 Munsell surface reflectances obtained from Brainard and Wandell [2] were analysed.

7.1 Assumptions

While interreflection can provide additional information for approaching the problem of colour constancy [6, 8, 13, 12], simultaneous solutions invariably lead to poorer estimates of shading and interreflection fields than could be attained by approaching interreflection as a problem in its own right. With the moderate success and continuing research being performed on colour constancy [14, 16, 20, 29, 30], it is not unreasonable to assume the relative spectral power distribution of the illuminant is known. However, we need not know the actual intensity of this source. To further simplify the problem, we also assume there is a single illuminant.

Simplifying assumptions may also be made regarding the properties of the inter-reflecting surfaces themselves. The first is that both surfaces are Lambertian, having spectral reflectance functions invariant of viewing position.

Secondly, since an integral part of decomposing interreflected colour signals is the reconstruction of a unique one-bounce signal, the spectral reflectances of both surfaces

must not be constant (ie. white or grey). The effect of relaxing this for one surface will be investigated in chapter 8.

We must also make some assumptions regarding the imaging geometry. With respect to reconstructing the multibounce colours, interreflection is limited to that between two surfaces. Furthermore, there must be at least one facet on each surface where interreflection is either absent or negligible, and we have obtained filter responses for those facets. Both facets must be illuminated by the single known illuminant.

Finally, we assume the filter sensitivity functions $R_k(\lambda)$, $k=1..3$ are known. This is a common assumption, often not explicitly stated and rarely relaxed.

7.2 Implementation

Experiments were performed using the spectral power distribution of a 3200K Planckian radiator and the first three basis vectors obtained through singular value decomposition of 462 Munsell surface spectral reflectances [2]. Results were obtained for three sets of filters (CIE 1931 standard observer, Kodak #25 (red), #58 (green) and #47b (blue), and the Sony DXC-151 CCD). As in chapter 6, tests were repeated using the D_{65} standard illuminant, having a flatter, but less smooth spectral power distribution $E(\lambda)$.

The matrix Λ was precomputed, where the ki^{th} entry of Λ is $\sum E(\lambda)S_i(\lambda)R_k(\lambda)$, and for each pair of surfaces the channel responses ρ_k^A and ρ_k^B were calculated according to equation 2.2. The coefficients σ_i of the basis vectors $S_i(\lambda)$ were then obtained for equation 4.3 by Gauss-Jordan elimination. A finite dimensional model with three degrees of freedom was then used to reconstruct the surface reflectance according to equation 4.1. The one-bounce colour

$$C^{(AB)}(\lambda) = E(\lambda)S^A(\lambda)S^B(\lambda) \quad (7.3)$$

was then approximated as

$$C^{(AB)}(\lambda) \simeq E(\lambda) \sum_{i=1}^3 S_i(\lambda)\sigma_i^A\sigma_i^B \quad (7.4)$$

and the two-bounce colour for surface A

$$C^{(ABA)}(\lambda) = E(\lambda) S^A(\lambda) S^B(\lambda) S^A(\lambda) \quad (7.5)$$

was approximated as

$$C^{(ABA)}(\lambda) \simeq E(\lambda) \sum_{i=1}^3 S_i(\lambda) \sigma_i^A \sigma_i^B \sigma_i^A. \quad (7.6)$$

The two-bounce colour for surface B was reconstructed in a similar manner. The *RMS* and ΔE errors between the signal of equation 7.3 and the reconstructed metamer of equation 7.4 were then measured, and statistical results gathered. Statistics were also obtained of errors for the two-bounce signals as described by equations 7.5 and 7.6, and the surface B equivalent equations.

In addition to those of colour signals, errors between theoretical and reconstructed tristimulus values for the three sets of filters were measured. As ΔE describes differences of colour signals measured through CIE Standard Observer filters this was not used for the tristimulus values. Instead, an angular error was employed, describing the angle (in degrees) between the theoretical and reconstructed vectors if rooted at the origin of three dimensional colour space.

From the exhaustive set of surfaces, test cases were selected to display specific types of errors.

7.3 Results

7.3.1 One-bounce Colour Reconstruction

Examination of the mean *RMS* error in tables 7.1 and 7.2 shows similar errors between theoretical and reconstructed $C^{(1)}(\lambda)$ for the three sets of filters under both illuminants. A slightly lower error under 3200K might be attributed to its smooth power distribution as shown in figure 6.1.

A first glance at the mean ΔE errors of the same tables seems to indicate a better performance for the Standard Observer than the other filter systems. However, we must realize that ΔE is obtained by first taking the responses of the colour signal

$C^{(1)}(\lambda)$ as seen through the Standard Observer and translating to $L^*u^*v^*$ coordinates by the set of equations 2.4. We should thus use this measure of error for comparison of different illuminants, and for different $C^{(i)}(\lambda)$ rather than as a comparison between sets of filters. It is also useful for providing us with a perception of the error in colour signal reconstruction, as $\Delta E = 1.0$ indicates a just-noticeable difference, as defined by the $L^*u^*v^*$ coordinate system.

Minimum and maximum errors show considerable spread for all filter systems under both illuminants. While some colour signals are nearly perfectly modeled by a finite dimensional model with three degrees, others are quite poorly modeled. Very high errors can be expected with a product of two poorly modeled surfaces.

Tables 7.3 and 7.4 indicate a slightly better mean *RMS* error for reconstruction of $\rho^{(1)}$ under D_{65} than $3200K$. However, since our goal is the decomposition of surface ρ as linear combinations of $\rho^{(i)}$, the most significant indicator of one-bounce colour reconstruction is the angular error. Using this measure, the $3200K$ illuminant gives slightly better mean errors under all illuminants, with much smaller maximum errors. For both mean and maximum errors, the Kodak filters provide better results.

7.3.2 Two-bounce Colour Reconstruction

For application of a two-bounce model of interreflection, a third colour signal $C^{(2)}(\lambda)$ and resulting response vector $\rho^{(2)}$ must be reconstructed.

Tables 7.5 and 7.6 demonstrate one problem with extending the model of interreflection to two bounces: as i increases, the error in reconstructed signal $C^{(i)}(\lambda)$ increases considerably. The relative performances of the three filters under the two illuminants remains much the same.

The reconstructed $\rho^{(2)}$ of tables 7.7 and 7.8 also show similar increases. Under the $3200K$ illuminant, the angular error almost precisely doubles, while D_{65} provides a moderate increase.

<i>Response Filters</i>	<i>minimum</i>		<i>maximum</i>		<i>mean</i>		<i>std. dev.</i>	
	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>
Std. Observer	0.0	0.9	138.6	124.6	4.6	13.9	6.2	10.9
Sony DXC-151	0.0	0.8	118.5	67.9	8.2	12.1	9.4	8.2
Kodak Filters	0.0	0.8	127.8	64.2	12.1	12.4	11.2	8.5

Table 7.1: Error statistics for reconstructed $C^{(1)}(\lambda)$ under 3200K illuminant.

<i>Response Filters</i>	<i>minimum</i>		<i>maximum</i>		<i>mean</i>		<i>std. dev.</i>	
	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>
Std. Observer	0.0	0.7	133.1	145.6	5.0	14.7	6.8	10.8
Sony DXC-151	0.0	0.7	154.4	77.9	9.6	13.1	9.4	8.4
Kodak Filters	0.0	0.6	190.6	75.7	14.3	13.4	13.2	8.8

Table 7.2: Error statistics for reconstructed $C^{(1)}(\lambda)$ under D_{65} illuminant

<i>Response Filters</i>	<i>minimum</i>		<i>maximum</i>		<i>mean</i>		<i>std. dev.</i>	
	<i>deg.</i>	<i>RMS</i>	<i>deg.</i>	<i>RMS</i>	<i>deg.</i>	<i>RMS</i>	<i>deg.</i>	<i>RMS</i>
Std. Observer	0.0	0.0	9.6	16.8	0.5	0.9	0.7	1.1
Sony DXC-151	0.0	0.0	8.9	15.5	0.6	1.1	0.7	1.2
Kodak Filters	0.0	0.0	6.0	10.4	0.3	0.6	0.4	0.7

Table 7.3: Error statistics for reconstructed $\rho^{(1)}$ under 3200K illuminant.

<i>Response Filters</i>	<i>minimum</i>		<i>maximum</i>		<i>mean</i>		<i>std. dev.</i>	
	<i>deg.</i>	<i>RMS</i>	<i>deg.</i>	<i>RMS</i>	<i>deg.</i>	<i>RMS</i>	<i>deg.</i>	<i>RMS</i>
Std. Observer	0.0	0.0	20.8	11.9	1.0	0.6	1.2	0.7
Sony DXC-151	0.0	0.0	14.9	7.4	0.8	0.6	1.1	0.7
Kodak Filters	0.0	0.0	12.0	77.0	0.7	0.4	0.8	0.6

Table 7.4: Error statistics for reconstructed $\rho^{(1)}$ under D_{65} illuminant.

<i>Response Filters</i>	<i>minimum</i>		<i>maximum</i>		<i>mean</i>		<i>std. dev.</i>	
	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>
Std. Observer	0.0	0.7	393.0	143.9	11.4	18.2	14.3	13.9
Sony DXC-151	0.0	0.8	286.8	98.4	15.3	16.7	18.6	12.0
Kodak Filters	0.0	0.9	548.8	101.8	18.7	17.2	18.7	12.7

Table 7.5: Error statistics for reconstructed $C^{(2)}(\lambda)$ under 3200K illuminant.

<i>Response Filters</i>	<i>minimum</i>		<i>maximum</i>		<i>mean</i>		<i>std. dev.</i>	
	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>
Std. Observer	0.0	0.9	304.7	178.9	12.6	18.9	16.1	13.6
Sony DXC-151	0.0	0.9	281.5	121.0	16.3	17.5	18.3	11.7
Kodak Filters	0.0	0.9	512.4	149.6	19.7	17.9	19.0	12.3

Table 7.6: Error statistics for reconstructed $C^{(2)}(\lambda)$ under D_{65} illuminant.

<i>Response Filters</i>	<i>minimum</i>		<i>maximum</i>		<i>mean</i>		<i>std. dev.</i>	
	<i>deg.</i>	<i>RMS</i>	<i>deg.</i>	<i>RMS</i>	<i>deg.</i>	<i>RMS</i>	<i>deg.</i>	<i>RMS</i>
Std. Observer	0.0	0.0	20.6	35.8	1.0	1.8	1.2	2.1
Sony DXC-151	0.0	0.0	16.2	28.1	1.2	2.2	1.3	2.3
Kodak Filters	0.0	0.0	178.8	200.0	0.7	1.1	1.8	2.4

Table 7.7: Error statistics for reconstructed $\rho^{(2)}$ under 3200K illuminant.

<i>Response Filters</i>	<i>minimum</i>		<i>maximum</i>		<i>mean</i>		<i>std. dev.</i>	
	<i>deg.</i>	<i>RMS</i>	<i>deg.</i>	<i>RMS</i>	<i>deg.</i>	<i>RMS</i>	<i>deg.</i>	<i>RMS</i>
Std. Observer	0.0	0.0	31.9	54.9	1.2	2.0	1.4	2.4
Sony DXC-151	0.0	0.0	16.8	29.3	1.1	2.0	1.2	2.2
Kodak Filters	0.0	0.0	177.3	199.9	0.8	1.4	2.0	2.6

Table 7.8: Error statistics for reconstructed $\rho^{(2)}$ under D_{65} illuminant.

Chapter 8

Surface Colour Decomposition

With statistical information gathered for selection of test cases and the underlying method now in place for reconstruction of multibounce colours, we proceed with decomposition of colours from interreflecting surfaces into shading and interreflection fields. Synthetic images will first be used to test the algorithm with controlled assumptions, and in the absence of experimental noise. Real image counterparts of the synthetic images are then decomposed. For all decompositions the recovered shading field is examined, and the improvement this field provides over “shape-from-shading” (i.e. shape reconstruction using image intensity) demonstrated.

8.1 Assumptions

The majority of assumptions necessary for surface colour decomposition are necessary for reconstruction of multibounce colour responses $\rho^{(i)}$, for $i > 0$, are discussed in section 7.1. Here we will itemize these assumptions, including others more specific to the task of decomposition.

1. We assume this problem is part of a larger one in which patches of interreflecting surfaces are segmented and identified.
2. Surfaces are Lambertian, such that surface appears to be equally bright, with an equal colour spectrum from all viewing positions. The Lambertian assumption

is stronger than simply requiring the relative reflectance spectrum $S(\lambda)$ to be invariant. The latter assumption can be used if the recovered shading fields are not to be used for surface shape reconstruction.

3. The sensitivity functions $R_k(\lambda)$, $k=1..3$ are known.
4. The relative spectral power distribution $E(\lambda)$ is known. We need not know the absolute intensity.
5. Illuminant intensity remains constant over each surface patch.
6. Only two surfaces are interreflecting. This assumption can be relaxed to three surfaces using a one-bounce model, but this will not be investigated here.
7. Surfaces do not have a flat spectral reflectance curve (ie. are not white or grey). The affect of relaxing this assumption for one surface will be demonstrated.
8. There exists a point on each surface where interreflection is insignificant or nonexistent. In real images this will be taken as a point furthest from the adjacent surface, and the results will be affected to the extent the assumption holds.

8.2 Algorithm

1. Precompute the 3×3 matrix Λ , where the ki^{th} entry is $\int E(\lambda)S_i(\lambda)R_k(\lambda)d\lambda$. If $E(\lambda)$ is unknown, it will first need to be reconstructed.
2. Obtain $\rho^{(A)}$ and $\rho^{(B)}$ from points on the two surfaces A and B where interreflection is nonexistent or negligible.
3. Using equation 4.3, obtain σ^A and σ^B . Reconstruct $S^A(\lambda)$ and $S^B(\lambda)$ using equation 4.1.
4. Reconstruct $\rho^{(AB)}$, $\rho^{(ABA)}$, and $\rho^{(BAB)}$ as described in equations 5.4 and 5.7.

<i>Colour</i>	<i>Case 1</i>		<i>Case 2</i>		<i>Case 3</i>	
<i>Signal</i>	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>	ΔE	<i>RMS</i>
$C^{(AB)}(\lambda)$	2.3	4.8	116.5	21.8	19.8	11.0
$C^{(ABA)}(\lambda)$	7.7	7.6	54.8	9.6	16.0	11.7
$C^{(BAB)}(\lambda)$	2.8	5.9	264.7	63.4	53.3	7.9

Table 8.1: Colour signal reconstruction errors for selected cases.

5. For both surfaces, calculate the vector angles θ and ϕ using the method described in chapter 6 and compare with predetermined thresholds. If θ is below its threshold, then a shape-from-shading model will be used for that surface. Otherwise, if ϕ is below its threshold, then a one-bounce model will be used. Otherwise, a two-bounce model will be used.
6. For each pixel at location x on each surface, calculate the field coefficients depending on the selected model for that surface. For a two-bounce model use three equations in three unknowns of the form described in equation 5.7, and solve for the shading field and one-bounce and two-bounce interreflection fields using any linear method. For a one-bounce model use three equations in two unknowns of the form described in equation 5.4, and solve for the shading field and one-bounce interreflection field using a linear least-squares fit. A no-bounce model simply uses colour intensity as the shading field.

8.3 Selected Cases

Of the three filters and two light sources under which the coloured surfaces were statistically evaluated, the Sony DXC-151 filters and 1000W tungsten-halogen (3200K blackbody) illuminant were selected for case studies, simply because this equipment was available in the lab. Surfaces for the three test cases were chosen, however, for their results relative to the overall performances of the entire set, and their Munsell colour coordinates identified [23, 17]. Tables 8.1-8.3 describe the selection criteria.

Colour Response	Case 1		Case 2		Case 3	
	deg.	RMS	deg.	RMS	deg.	RMS
$\rho^{(AB)}(\lambda)$	0.4	4.5	4.5	7.9	0.6	1.0
$\rho^{(ABA)}(\lambda)$	0.7	1.7	2.9	1.2	2.1	1.1
$\rho^{(BAB)}(\lambda)$	0.6	1.1	10.0	17.5	0.8	1.4

Table 8.2: Colour response reconstruction errors for selected cases.

Angle	Case 1			Case 2			Case 3		
	model	recon.	err.	model	recon.	err.	model	recon.	err.
θ^A	9.5	9.7	0.2	10.0	12.1	2.1	24.3	23.7	-0.6
θ^B	12.8	13.0	0.2	44.4	42.8	-1.6	0.1	0.5	0.4
ϕ^A	13.0	13.5	0.5	1.4	1.6	0.2	0.0	0.0	0.0
ϕ^B	6.5	6.6	0.1	6.2	14.0	7.8	0.2	0.2	0.0

Table 8.3: Comparison of angles (degrees) for ρ of actual model and for reconstruction of test cases.

8.3.1 Test Case 1

Munsell surfaces *10BG 6/4* (blue-green) and *5P 5/6* were chosen as representative of those surfaces considered “well-behaved” with respect to the selection criteria.

As observed in table 8.1, the errors in reconstruction of the colour signals compare favourably with the mean errors for the Sony DXC-151 found in tables 7.1 and 7.5. This can be attributed to their smoothly varying spectral reflectances, as shown in figure 8.1. These results directly affect the colour responses in table 8.2, which also compare favourably to mean errors found in 7.3 and 7.7.

In table 8.3, the error between θ for an ideal model (where ρ are known precisely), and those of reconstructed ρ are relatively small. Furthermore, if we refer back to table 6.1, we find mean values of 9.6° and 3.3° for θ and ϕ respectively. θ and ϕ for both surfaces are thus large enough to evaluate the performance of a both the one-bounce and two-bounce models.

8.3.2 Test Case 2

The surfaces of the second test case, Munsell colours *5R 4/14* (red), and *5PB 3/12* (blue), were selected for their poor colour signal reconstruction.

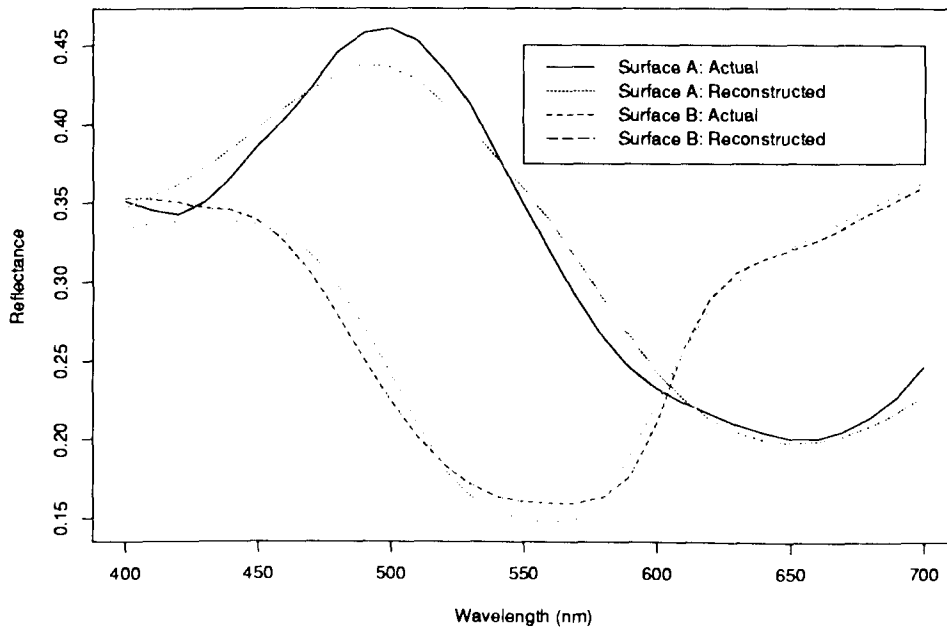


Figure 8.1: Actual and modeled spectral reflectances $S(\lambda)$ for case 1 surfaces A (blue-green) and B (purple).

In contrast to the surfaces of Case 1, the spectral reflectances for Case 2 are poorly modeled by a finite dimensional model of degree three (see figure 8.2). As a result, the one-bounce and two-bounce colour signal errors in table 8.1, and the tristimulus values in table 8.2 are large.

Interestingly, both θ of the reconstructed ρ in table 8.3 are quite close to those of the model, providing the possibility of obtaining reasonable results from a one-bounce model in spite of the poor colour signal reconstruction. Values of ϕ^A and ϕ^B , on the other hand, indicate that a two-bounce model should not be attempted with these surfaces. The former angle is much too small, while the latter has an unacceptable error.

8.3.3 Test Case 3

The third test case was selected to observe the interreflection between a surface having constant spectral reflectance (ie. white or grey) and another non-white surface. To provide such a case, Munsell colours $N6/$ (grey) and $5R\ 4/14$ (red) were used. As in




Figure 8.2: Actual and modeled spectral reflectances $S(\lambda)$ for case 2 surfaces A (red) and B (blue).

the previous case, the red surface also provides the additional property of having a poorly reconstructed colour signal.

As figure 8.3 shows, the grey surface is modeled very well by the finite dimensional model, but red is poorly modeled. This results in relatively large signal reconstruction errors in table 8.1. This is not reflected in the one-bounce colour response in table 8.2 because the product of the two signals is close to being a scaled response of that for the original red surface, which we reconstructed from tristimulus values obtained directly from the image.

In table 8.3, we find θ^B to be close to 0^0 . This is expected, since the no-bounce red signal reflecting from the grey surface is just an attenuated red signal with the same spectra. We thus only hope to use a no-bounce model for surface B . We can, however, still use a one-bounce model for Surface A . Obviously a two-bounce model is not possible given the values of ϕ .

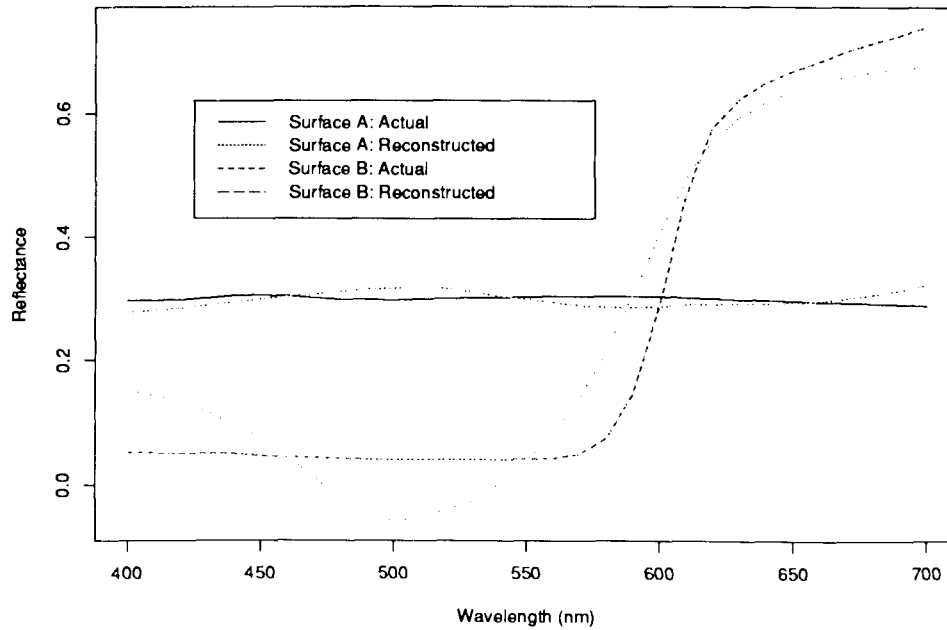


Figure 8.3: Actual and modeled spectral reflectances $S(\lambda)$ for case 2 surfaces A (grey) and B (red).

8.4 Reconstructing Surface Shape

To illustrate the improvement the recovered shading field can provide over image intensity for “shape-from-shading” methods in the presence of interreflection, a simple surface shape recovery algorithm can be implemented along the one-dimensional surfaces. That is, given known x and constant y , the third dimension z can be found.

For Lambertian surfaces illuminated by a parallel beam of luminous flux E , with angle an ϵ between E and the surface normal \mathbf{n} (see figure 8.4), the resulting luminous intensity I of that surface is given by [31]:

$$I = \frac{E}{\pi} \cos \epsilon . \quad (8.1)$$

Normally, equation 8.1 would not be valid for a one-dimensional synthetic edge generated by the radiosity method, since an isotropic rather than parallel beam illuminant is used. In this case, however, both generated surfaces are flat, and the effects associated with isotropic and parallel beam illumination are identical. For the purposes of illustration the assumption can be made that they were illuminated by a




Figure 8.4: Ray geometry for luminous intensity I of a Lambertian surface resulting from illumination by luminous flux E . ϵ is the angle between surface normal \mathbf{n} and the direction of E .

parallel light source placed directly overhead. This assumption would not have been possible had the individual surface shapes been curved.

Since the illuminant for real images was positioned directly above the interreflecting edge and is assumed to be directly above the edge for synthetic images, the orientation ϵ_0 of facets x_0 furthest from the edges on each surface is known to be 45° . We assume that shape has been reconstructed correctly up to these points since interreflection is negligible.

Equation 8.1 can be reformulated as

$$k = \frac{I_0}{\cos \epsilon_0} \quad (8.2)$$

where $\epsilon_0 = 45^\circ$ and $k = E/\pi$. If reconstruction is to be performed using the recovered shading field, then I_0 is the shading field at x_0 , otherwise it is the colour intensity.

The orientation ϵ_i of each successive facet x_i , for $i > 0$ toward the edge can then be calculated as

$$\epsilon_i = \arccos \left(\frac{I_i}{k} \right) \quad (8.3)$$

where I_i is either the shading field or colour intensity of facet x_i .

To complete the reconstruction, surfaces must also be assumed to vary smoothly (which is consistent with the assumption that image segmentation has already been performed). Each successive facet is thus connected to the previous facet, with reconstruction beginning furthest from the edge.

8.5 Synthetic Images

8.5.1 Constructing the Images

To evaluate the algorithm in the absence of experimental noise and measurement errors, synthetic images of the selected test cases were generated. These are multispectral images of theoretical edges composed of two semi-infinite planar surfaces illuminated by diffuse light. The radiosity method described in [6] is used to provide a good simulation of interreflection between two surfaces. This is an iterative method in which each surface is divided into a number of facets and the colour spectra of each facet results from that emitted from the surface and those reflected from colour signals being emitted from all others (as described in the general radiosity equation 3.1). Using the image geometry of figure 8.5, the number of interreflecting surfaces is reduced to two, and a one-dimensional interreflecting edge (ie. a single slice through the two planes) was constructed. Colour signal $C^A(\omega_A, \lambda)$ emanating from surface A is iteratively generated using isotropic illuminant $E(\lambda)$ as

$$\begin{aligned} C^{A^t}(\omega_A, \lambda) &= S^A(\lambda)E(\lambda) + S^A(\lambda) \int_0^{\Omega_B} C^{B^{t-1}}(\omega_B, \lambda) \gamma(\omega_A, \omega_B) d\omega_B, \text{ for } t > 0 \\ C^{A^0}(\omega_A, \lambda) &= S^A(\lambda)E(\lambda) \end{aligned} \quad (8.4)$$

where $\gamma(\omega_A, \omega_B)$ is the wavelength-independent *configuration factor* representing the contribution from infinitesimal area ω_B to infinitesimal area ω_A . This is given by [6, 13, 28]:

$$\gamma(\omega_A, \omega_B) = \frac{1}{2} \left[1 + \frac{\Omega_B \cos \beta - \omega_A}{(\omega_A^2 + \Omega_B^2 - 2\omega_A \Omega_B \cos \beta)^{\frac{1}{2}}} \right]. \quad (8.5)$$

$C^B(\omega_B, \lambda)$ is constructed simultaneously in the same manner.




Figure 8.5: Geometry of a semi-infinite edge used to simulate interreflection between two planar surfaces A and B with opening angle β .

8.5.2 Implementation

For each selected test case a synthetic image was generated with surfaces having spectral reflectances $S^A(\lambda)$ and $S^B(\lambda)$. The spectral power distribution of the D_{65} standard illuminant was used as isotropic light source $E(\lambda)$. Edges were generated for the three sets of filters (CIE 1931 standard observer, Kodak #25 (red), #58 (green) and #47b (blue), and the Sony DXC-151 CCD). An opening angle β of 90° was used in all simulations. This angle was chosen as being representative of a large percentage of man-made objects. For example, the walls of a room typically meet at 90° . The effect of varying this angle has been investigated [6, 21], with interreflection increasing as β decreases.

The decomposition algorithm described in section 8.2 was applied to the synthetic images. The resulting shading fields were then plotted against those obtained through calculations using the known values of $S^A(\lambda)$ and $S^B(\lambda)$ and the same model of interreflection.

8.5.3 Results

Test Case 1

Applying a one-bounce interreflection model to the synthetic image of test case 1 (blue-green and purple surfaces) resulted in shading fields very close to the actual constant value of 1.0, as shown in figure 8.6. The theoretical model (with known $S^A(\lambda)$ and $S^B(\lambda)$), represents the best we can hope for using the model. As we see, the recovered fields of both surfaces follow this very closely. While surface A has a smaller error with respect to the model, both the model and recovered shading fields of surface B are closer to the actual model.

Application of a two-bounce model to the test image is shown in figure 8.7. While the recovered fields are both slightly improved over those of a one-bounce model, the theoretical model results are mixed. We see that toward the edge, surface A shows an improvement in absolute distance from the actual shading field, yet surface B is poorer. Furthermore, the error between actual and theoretical models has approximately doubled. This is likely a more significant factor in a real image environment which is prone to additional sources of error.

To put the shading field results into perspective, observe figure 8.8. Reconstruction of surface shape using colour intensity results in the surface being “shallower” than it actually is. The one-bounce reconstruction is shown for comparison. Figure 8.9 is a magnification of the surfaces close to the edge, displaying reconstructed surface shape from the recovered one-bounce and two-bounce shading fields. Although the two-bounce reconstruction is slightly better than the one-bounce, one should note the lack of significance of such an improvement, since both models reconstruct the surface very close to the actual shape.

Test Case 2

The second test case (red and blue surfaces) was chosen because of poor colour signal reconstruction for both surfaces, as shown in figure 8.2. However, table 8.3 describes large θ with small errors for both surfaces, leading us to expect good results. Referring to figure 8.10 we see this is so: the absolute shading field errors are close to those of




Figure 8.6: Test case 1 shading fields using one-bounce model. Vertical line represents edge between surface *A* (blue-green) on left and surface *B* (purple) on right.

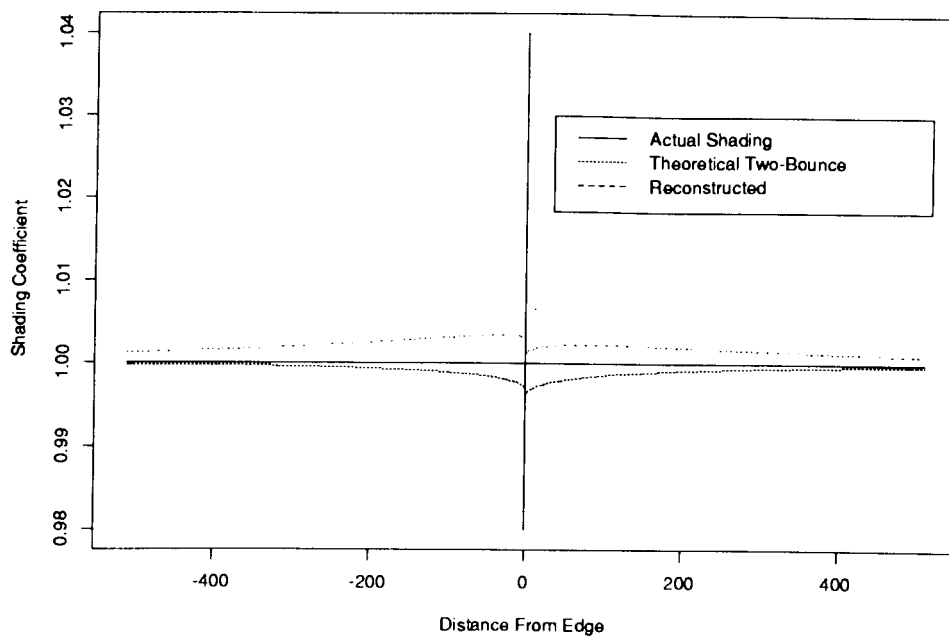


Figure 8.7: Test case 1 shading fields using two-bounce model. Vertical line represents edge between surface *A* (blue-green) on left and surface *B* (purple) on right.

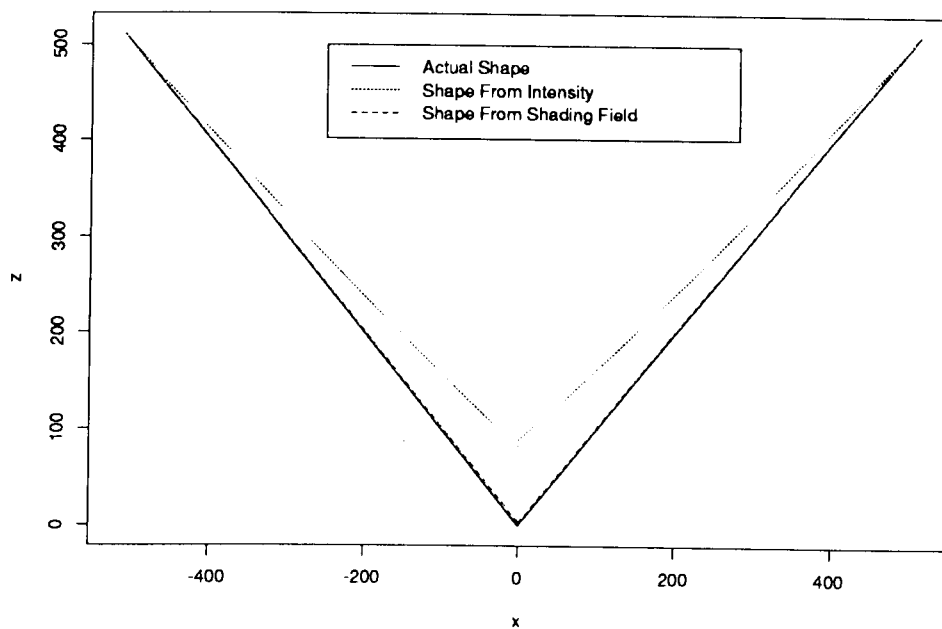


Figure 8.8: Shape from intensity and from one-bounce shading field for test case 1, with surface *A* on left and surface *B* on right.

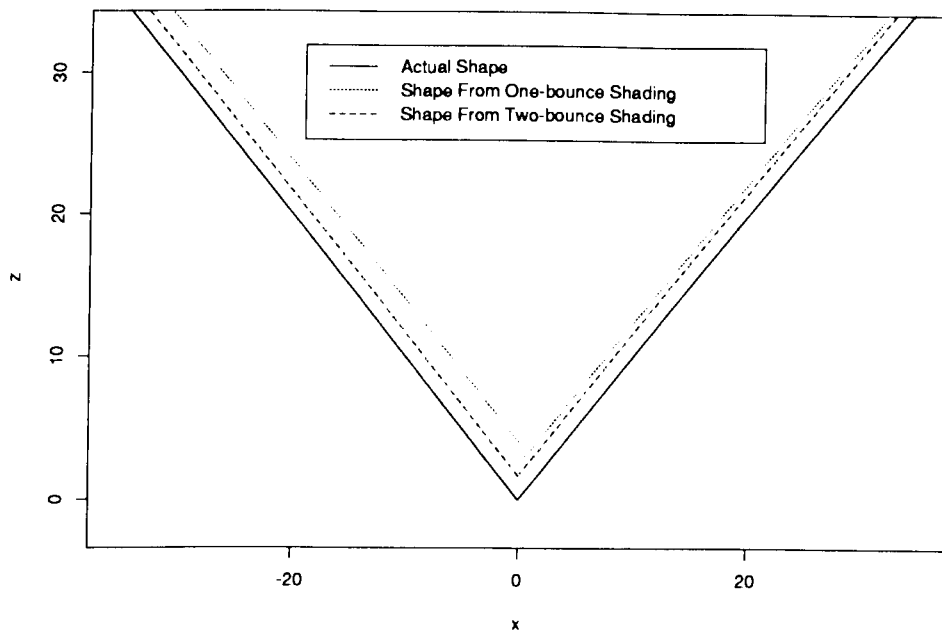


Figure 8.9: Shape from one-bounce shading field and from two-bounce shading field.

case 1 in figure 8.6.

In figure 8.11 we observe a smaller shape from intensity error than the previous case, indicating a lesser amount of interreflection between the red and blue surfaces. This is expected since the product of two spectra with distributions close to opposing ends is one with a small magnitude throughout the spectrum and reduced filter responses. The effect is most apparent on the red surface, as it has a greater basic colour intensity with respect to its one-bounce component. Even with a lesser amount of interreflection, the shading field of the one-bounce model is significantly better as an estimator of actual shape than simple colour intensity.

Test Case 3

Test case 3 is the degenerate case where surface A has a flat spectral reflectance curve (see figure 8.3). Since $C^{(A)}(\lambda)$ has the same spectral distribution as $E(\lambda)$, $C^{(B)}(\lambda)$ and $C^{(AB)}(\lambda)$ are indistinguishable except for their intensities. We see this in table 8.3, with small θ^B for both the model and reconstruction. This small angle indicates that $\rho^{(B)}$ and $\rho^{(AB)}$ are close to being linearly dependent, so ρ^B cannot be decomposed into its no-bounce and one-bounce components.

Although surface B cannot be analyzed with a one-bounce model, surface A can. Table 8.3 shows a reconstructed θ^A of 23.7° with a very small error with respect to the model. We cannot use a two-bounce model however, since ϕ^A is 0.0. The resulting shading fields are displayed in figure 8.12.

Figure 8.13 demonstrates the higher degree of error in shape from intensity due to greater intensities of common wavelengths with respect the surfaces in case 2, in particular between 600 and 700 nm . Reconstruction of surface A using the shading field of the one-bounce model provides a much closer estimate than that of colour intensity. Unfortunately, we cannot use the one-bounce model to improve shape reconstruction for surface B .




Figure 8.10: Test case 2 shading fields using one-bounce model. Vertical line represents edge between surface *A* (red) on left and surface *B* (blue) on right.




Figure 8.11: Shape from intensity and from one-bounce shading field for test case 2, with surface *A* (red) on left and surface *B* (blue) on right.




Figure 8.12: Test case 3 shading fields using the one-bounce model. Vertical line represents edge between surface *A* (grey) on left and surface *B* (red) on right. Decomposition is not possible for surface *B* using this model.




Figure 8.13: Shape from intensity and from one-bounce shading field for test case 3, with surface *A* (grey) on left and surface *B* (red) on right. One-bounce reconstruction cannot be performed for surface *B*.

8.6 Real Images

Surface colour decomposition has proved to be fairly robust within a synthetic environment, where control is maintained over the environment and assumptions can be certain to hold. Real images, on the other hand, provide us with conditions that are somewhat less than ideal. Random noise, imperfect Lambertian surfaces, varying illumination, and imprecise knowledge of camera sensitivity functions all contribute to errors in the decomposition process. Although we can only hope to approach the results obtained with synthetic images, we still expect the one-bounce model of mutual illumination to be approximated.

Wherever possible, the equipment properties and imaging geometry of the synthetic environment were duplicated for obtaining real images. Five 3×5 inch matte Munsell papers corresponding to the chosen surface reflectance curves were obtained, and a 1000W 3200K tungsten-halogen lamp and Sony DXC-151 CCD camera were used. To reduce the possibility of responses to signals beyond 700 nm, a Kodak 301-A infrared filter was placed in front of the camera lens. Its transmittance curve is shown in figure 8.14. In addition, the camera required calibration to produce linear responses.

8.6.1 Obtaining Linear Responses

There are two properties of a camera that must be considered when collecting images. The first, and most obvious, is the three spectral sensitivity curves. These were provided by the Sony Corporation upon request, and will be assumed to be reasonably accurate. The second property is the linearity of the three channel responses. The response of a channel is linear if its response under lighting condition A is ρ_A , under lighting condition B is ρ_B , under combined lighting conditions $A + B$ is ρ_C , and the following equation holds:

$$\rho_C = \rho_A + \rho_B \quad (8.6)$$

When tests of this nature were performed on the DXC-151, it was found to have




Figure 8.14: Transmittance curve of Kodak 301-A infrared filter.

nonlinear responses, with each channel exhibiting its own nonlinearity. Since consistent colour responses under varying illuminant intensity are integral to the success of this algorithm, correction functions are required.

A straightforward approach to linear correction is the construction of three lookup tables with entries between 0 and 255, corresponding to the camera's initial responses. Given the example in the previous paragraph, an equation of the following form can be constructed:

$$x_{\rho_A} + x_{\rho_B} - x_{\rho_C} = 0 \quad (8.7)$$

where x_{ρ_i} is the linearly corrected value for response ρ_i .

Correction functions were obtained by taking images under all 16 combinations of four separate illuminants. The object taken in this case was Munsell paper *N6/* on a black cloth background. Random samples were obtained for various linear combinations of illuminants, and a $k \times 256$ homogeneous matrix obtained for each channel, where $k \geq 256$. Since all 256 intensities were not represented, the empty columns corresponding to missing ones were removed, and the remaining matrix solved by least-squares approximation. Missing values were then interpolated or extrapolated,

and the resulting monotonically increasing function smoothed and scaled such that $x_{255} = 255$.

References to *linearized responses* in subsequent sections refers to those obtained through the three correction tables calculated by this process.

8.6.2 Implementation

Before obtaining images of interreflecting surfaces, an image of the illuminant was taken. A scalar was obtained for each linearly corrected response such that the product of the response and this scalar produced the theoretical response for a 3200K illuminant using normalized spectral sensitivity functions. Thus, a perfect white light would produce equivalent responses for each channel.

For each test case, the two corresponding Munsell papers were mounted in a supporting stand, such that the surfaces were flat, with an opening angle β of 90° , each at 45° to the image plane. The camera was mounted directly above, with the illuminant beside it about one metre from the surfaces.

Images were taken of each test case, and cross-sections perpendicular to the edge were obtained. The two segments corresponding to the surfaces were extracted and smoothed. Finally, each pixel was linearized then corrected with the corresponding channel scalar. The tristimulus values furthest from the edge taken as the no-bounce colours. The shading fields along each surface were then extracted using the one-bounce decomposition algorithm of section 8.2.

8.6.3 Results

Test Case 1

The first test case was selected because of favourable synthetic properties. $S^A(\lambda)$ (blue-green) and $S^B(\lambda)$ (purple) were both reconstructed well using a finite dimensional model, with colour vectors $\rho^{(A)}$, $\rho^{(B)}$, and $\rho^{(AB)}$ oriented in *RGB* space such that decomposition of ρ^A and ρ^B was feasible.

Results for the real image are somewhat poorer. Figure 8.15 indicates a maximum




Figure 8.15: Shading fields of test case 1 using the one-bounce model on a real image.




Figure 8.16: Shape from intensity and from one-bounce shading fields for real test case 1.

absolute error for shading coefficients a_0^A and a_0^B of 0.1 and 0.2 respectively, with a_0^A hovering close to the ideal value of 1.0, while a_0^B consistently increases. The effect of applying these coefficients to the recovery of surface shape is shown in figure 8.16. We see that a_0^A still shows a significant improvement over intensity for surface A , whereas a_0^B shows only minimal improvement for surface B .

Test Case 2

As in the synthetic environment, this case demonstrates the limited interreflection between red and blue surfaces. Figure 8.17 indicates maximum absolute errors of 0.07 and 0.02 for a_0^A and a_0^B . Furthermore, as observed in figure 8.18, with the limited interreflection we cannot hope to make significant improvements using shading fields over intensity for shape recovery.

Test Case 3

The third test case demonstrates the interreflection between white and red surfaces. While recovery of a_0^B is not possible, a_0^A of the white surface was obtained with a maximum absolute error of 0.03 (figure 8.19). Improvements in shape recovery are shown for surface A , but since we have not been able to recover a_0^B , we must be content with using intensity for shape recovery of surface B .

Sources of Error

The preceding results were affected by errors resulting from various factors violating the stated assumptions in a real image environment. Unlike synthetic images, the sources of these errors cannot always be stated definitively, and invariably have a collective effect.

Random Noise: For some machine vision techniques such as edge detection, excessive smoothing can reduce sharp gradients necessary for proper detection of features. Since we assume here that previous segmentation has been performed, we are left with surfaces containing gradients we expect to result only from interreflection or gradual

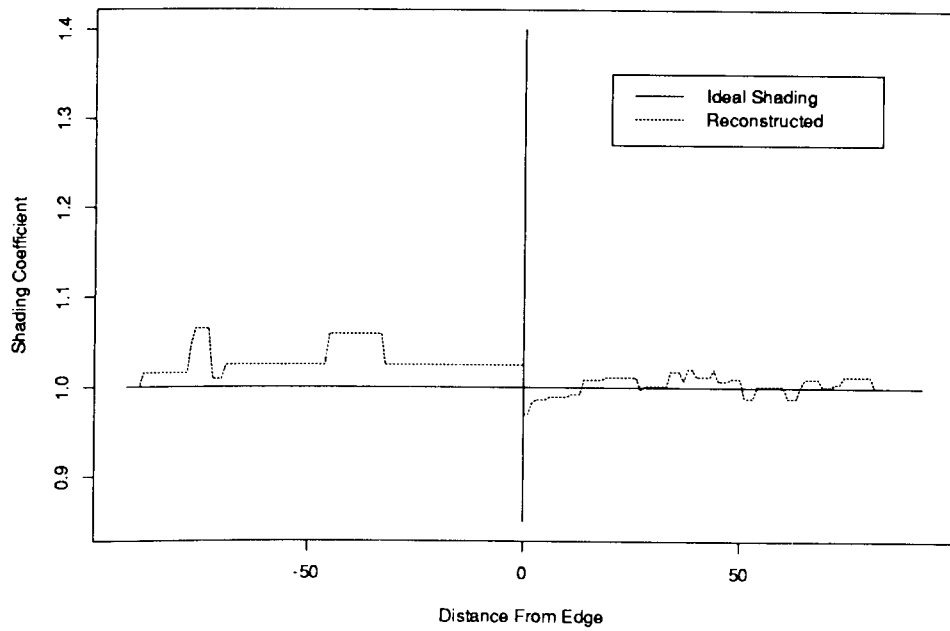


Figure 8.17: Shading fields of test case 2 using the one-bounce model on a real image.

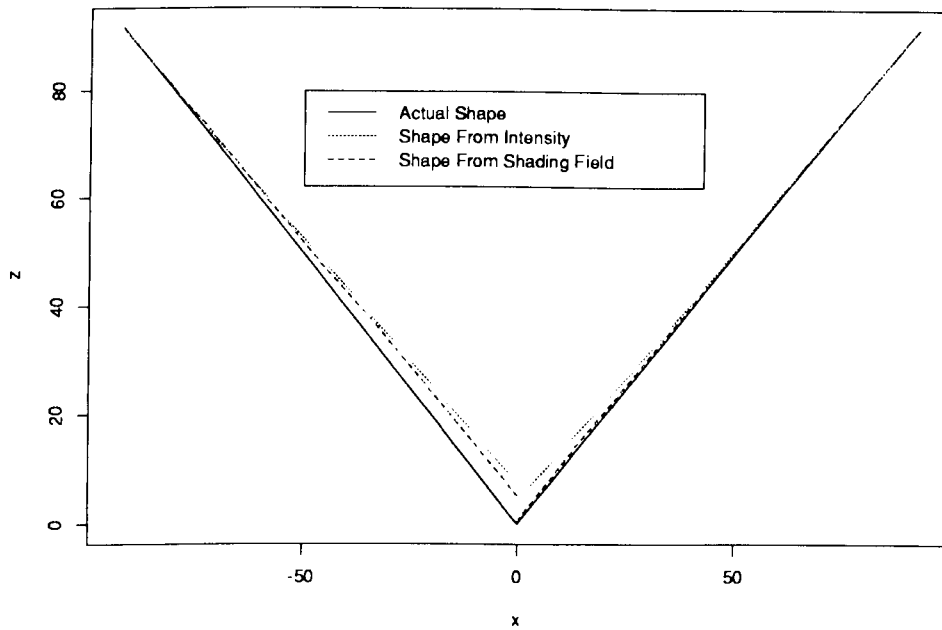


Figure 8.18: Shape from intensity and from one-bounce shading fields for real test case 2.




Figure 8.19: Shading fields of test case 3 using the one-bounce model on a real image. Decomposition is not possible for surface *B*.




Figure 8.20: Shape from intensity and from one-bounce shading fields for real test case 3. Shape can be only be recovered using shading fields for surface *A*.

changes in surface orientation. Random noise is thus not expected to contribute significantly to erroneous results, as minor perturbations are easily removed.

Surface Properties: The Munsell papers used for this experiment are described as *matte*. This type of surface texture approaches Lambertian qualities by minimizing specular reflections. However, a perfect Lambertian surface appears equally bright from all directions, according to Lambert's Law (equation 2.3). The extent to which the Munsell papers follow this law is unknown, and may vary from paper to paper. Imperfections and marks from handling will also affect their reflective properties. The one-bounce model also assumes the two interreflecting surfaces are semi-infinite. That is, the distances Ω_A and Ω_B in figure 8.5 are finite, but the distance along the common edge is infinite. The Munsell papers are five inches long. We can expect the interreflection toward the two supposedly-infinite ends to be reduced.

Illuminant Properties: During construction of the synthetic edge by the radiosity method, an isotropic illuminant was used. Consequently, every facet of the edge was illuminated with equal intensity. The one-bounce reconstruction algorithm does not require this property to hold over all surfaces. However, if we wish to use the shading field as an indicator of surface orientation, we must assume colour intensity gradients are the result of changes in surface orientation with respect to the illuminant or from interreflection. The proximity of the light source used in this experiment violates this assumption, producing a gradient in the direction of the source. Ambient light was eliminated as a possible source of error by obtaining images with the $3200K$ illuminant turned off. The resulting pixel intensities were insignificant. However, contributions of light reflecting from nearby immovable objects with the illuminant turned on are unknown. Finally, the spectral power distribution of a $1000W$ tungsten filament project lamp is listed as having a correlated colour temperature in the range $3100 - 3300K$ [31]. The spectral power distribution of colour temperatures is defined by Planck's Formula [31]. This was used to produce $E(\lambda)$ for the $3200K$ illuminant. As no calibration of the illuminant has been performed, the conformity of the actual distribution to this theoretical value is unknown.

Spectral Sensitivities and Transmittances: The camera's channel responses are affected not only by the photoreceptor sensitivity functions, but by all materials through which light rays must pass before being detected. The former were provided by the Sony Corporation, and taken to be correct. The transmittance of the Kodak 301-A infrared filter may be considered accurate but not precise, as it was manually extracted from a small-scale logarithmic graph [5], and for which the original data are no longer published. The camera lens is expected to have 100% transmittance, yet this is undoubtedly a false assumption. The divergence of the combined sensitivities and transmittances likely contributes enormously to incorrect results.

Selection of No-Bounce Colours: In the synthetic environment, the assumption that a point exists where no interreflection occurs was not violated. With real images the no-bounce colours are extracted from points furthest from the edge. If interreflection still occurs, this will affect reconstruction of the one-bounce colour, and subsequent extraction of the shading field.

Chapter 9

Discussion and Conclusions

This thesis has investigated interreflection between coloured surfaces through the use of a finite dimensional model of surface spectral reflectance. The use of this model to recover reflectances permits decomposition of interreflected colour signals that is simple, straightforward, and robust.

An analysis of the multibounce colour space resulting from interreflection between pairs of Munsell spectral reflectances under both D_{65} and $3200K$ illuminants illustrated the limitations of colour signal decomposition. As expected, one-bounce signals exist that cannot be decomposed as linear combinations of their no-bounce components. Such a condition results when one or both of the no-bounce colours has a flat (ie. grey or white) spectral reflectance curve. Comparison of the angles θ and ϕ subtended by colour responses as vectors in RGB space suggest a greater colour space is spanned using a “whiter” illuminant (D_{65}) and filters having less “overlap” (Kodak).

The problem of attaining multibounce colour responses $\rho^{(i)}$ for $i > 0$ is integral to the linear decomposition of interreflected colours, yet they cannot be attained simply through algebraic manipulation of the three components of $\rho^{(0)}$. Using finite dimensional models of $S^A(\lambda)$ and $S^B(\lambda)$, and with a known illuminant and filters, multibounce signals $C^{(i)}(\lambda)$ and resulting $\rho^{(i)}$ were obtained, with errors compounding as i increases. A quantitative analysis of the recovered response vectors indicates a mean error less than 1.0^0 for $\rho^{(1)}$ and 1.2^0 for $\rho^{(2)}$ over all illuminants and filters. As

with the colour space spanned by multibounce colours, Kodak filters generally provide superior results. However, a “smoother” (3200K) rather than “whiter” illuminant is more important for recovery of reflectances.

Based on the analysis of multibounce colour space and colour reconstruction, three test cases were chosen to demonstrate the decomposition algorithm, with each having separate characteristics and sources of error. A “best” case, a case where the finite dimensional model does a poor job of reconstructing $S^A(\lambda)$ and $S^B(\lambda)$, and a case where $\rho^{(0)}$ and $\rho^{(1)}$ are linearly dependent (ie. one surface is white) were all presented. In the synthetic environment, shading fields were recovered with good precision except where recovery was not possible. Real images resulted in poorer recovery, yet for all cases, the shading field was shown to provide a superior measure of surface shading than colour intensity as demonstrated by a simplified shape-from-shading method. Future work in shape recovery might incorporate the use of shading fields to enhance reconstruction in scenes where interreflection between two or three surfaces is known to exist.

It is questionable whether a two-bounce model can provide better shading field recovery than a one-bounce model since $\rho^{(2)}$ is more sensitive to errors in reconstruction of $\rho^{(A)}$ and $\rho^{(B)}$. Furthermore, even for perfect values of $\rho^{(2)}$ the angle ϕ is typically quite small, discouraging the use of a third component vector for linear decomposition. The synthetic test case for which both models were compared showed little benefit to advancing beyond one-bounce, with an improved shading field only slightly overcoming the error in reconstruction of $\rho^{(2)}$.

Investigation of a two-bounce model was made possible by the availability of three components of ρ for every pixel. The one-bounce model applied to interreflection between two surfaces needs to recover only two coefficients: a_0 and a_1 corresponding to the shading and interreflection fields. Future research in this area could involve interreflection between three surfaces (such as found in the corner of a room), where recovery of a shading field and two interreflection fields may be possible.

Although the decomposition algorithm is simple, direct, and robust as claimed, there are two major weaknesses that must be acknowledged. The first concerns the known illuminant assumption. While it is true that research in colour constancy

continually provides promising illuminant recovery methods, one might claim that the simplicity of decomposition presented here is realized by knowing the illuminant. However, the flexibility exists to integrate independent illuminant recovery methods, a topic of possible future research. The second weakness is the assumption that surfaces are Lambertian. In reality, surfaces have varying degrees of specularities, and the extent to which they stray from a simple Lambertian model will affect decomposition, likely resulting in shading coefficients of disproportionately high magnitude. This is perhaps the most critical weakness, precluding the analysis of arbitrary surfaces having unknown reflectance characteristics.

Bibliography

- [1] F.W Billmeyer and M. Saltzman. *Principles of Color Technology*. John Wiley and Sons, New York, 2nd edition, 1981.
- [2] D. Brainard and B. Wandell. *The Color Analysis Package*. Stanford University, 1990.
- [3] G. Buchsbaum and A. Gottschalk. Chromaticity coordinates of frequency-limited functions. *J. Opt. Soc. Am. A*, 1:885, 1984.
- [4] J. Cohen. Dependency of the spectral reflectance curves of the Munsell color chips. *Psychonomic Science*, 1:369, 1964.
- [5] Eastman Kodak company. *Kodak filters: for scientific and technical uses*, 2nd edition, 1981.
- [6] M.S. Drew and B.V. Funt. Calculating surface reflectance using a single-bounce model of mutual illumination. In *Proceedings: International Conference on Computer Vision*, pages 393–399, Osaka, December 1990. IEEE.
- [7] M.S. Drew and B.V. Funt. Natural metamers. *Computer Vision, Graphics, and Image Processing*, in press.
- [8] M.S Drew and B.V Funt. Variational approach to interreflection in color images. *J. Opt. Soc. Am. A*, August 1992.
- [9] G.D. Finlayson, M.S. Drew, and B.V. Funt. Spectral sharpening: Sensor transformations for improved color constancy. Technical Report CSS/LCCR TR91-14, Simon Fraser University, 1991.

- [10] J.D. Foley, A. van Dam, S.K. Feiner, and J.F. Hughes. *Computer Graphics, Principles and Practice*. Addison-Wesley, 2nd edition, 1990.
- [11] D.A. Forsyth. A novel algorithm for color constancy. *International Journal of Computer Vision*, 5(1):5–36, 1990.
- [12] B.V. Funt and M.S. Drew. Color space analysis of mutual illumination. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (To appear 1992).
- [13] B.V. Funt, M.S. Drew, and J. Ho. Color constancy from mutual reflection. *International Journal of Computer Vision*, 6(1):5–24, 1991.
- [14] R. Gershon, A.D. Jepson, and J.K. Tsotsos. From [R,G,B] to surface reflectance: Computing color constant descriptors in images. In *Proceedings of IJCAI Conference*, pages 755–758, 1987.
- [15] C.M. Goral, K.E. Torrance, D.P. Greenberg, and B. Battaile. Modeling the interaction of light between diffuse surfaces. *ACM Computer Graphics*, 18(3), July 1984.
- [16] J. Ho, B.V. Funt, and M.S. Drew. Separating a colour signal into illumination and surface reflectance components: Theory and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:966–977, 1990.
- [17] K.L. Kelly, K.S. Gibson, and D. Nickerson. Tristimulus specification of the Munsell Book of Color from spectrophotometric measurements. *J. Opt. Soc. Am.*, 33(7):355–376, 1943.
- [18] E.L. Krinov. Spectral reflectance properties of natural formations. *Technical Translation TT-439, National Research Council of Canada*, 1947.
- [19] L.T. Maloney. *Computational Approaches to Color Constancy*. PhD thesis, Stanford University, 1985.
- [20] L.T. Maloney and B.A. Wandell. Color constancy: a method for recovering surface spectral reflectance. *J. Opt. Soc. Am. A*, 3:29–33, 1986.

- [21] S.K. Nayar and Y. Gong. Colored interreflections and shape recovery. In *Image Understanding Workshop*, San Diego, January 1992. DARPA.
- [22] S.K. Nayar, K. Ikeuchi, and T. Kanade. Shape from interreflections. In *Proceedings: International Conference on Computer Vision*, pages 2–11, Osaka, December 1990. IEEE.
- [23] S.M. Newhall, D. Nickerson, and D.B. Judd. Final report of the O.S.A. Subcommittee on the spacing of the Munsell colors. *J. Opt. Soc. Am.*, 33(7):385–418, 1943.
- [24] D. Nickerson. *Spectrophotometric Data for a Collection of Munsell Samples*. U.S. Department of Agriculture, Washington, D.C., 1957.
- [25] C.L. Novak. *Understanding Interreflection in Machine Vision*. PhD thesis, Carnegie Mellon University, 1990.
- [26] W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, Cambridge, 1988.
- [27] S.A. Shafer. Using color to separate reflection components. *Color Res. Appl.*, 10(4):210–218, 1985.
- [28] R. Siegel and J. Howell. *Thermal Radiation Heat Transfer*. Hemisphere, Washington, DC, 1981.
- [29] S. Tominaga and B.A. Wandell. Component estimation of surface spectral reflectance. *J. Opt. Soc. Am. A*, 7:312–317, 1990.
- [30] B.A. Wandell. The synthesis and analysis of color images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9:2–13, 1987.
- [31] G. Wyszecki and W.S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulac*. Wiley, New York, 2nd edition, 1982.