

**PHENOMENOLOGICAL MODELING OF
IMAGE IRRADIANCE FOR
NON-LAMBERTIAN SURFACES UNDER
NATURAL ILLUMINATION**

By

Shireen Y. Elhabian

B.Sc. 2002, Faculty of Computers and Information, Cairo University

M.Sc. 2005, Faculty of Computers and Information, Cairo University

A Dissertation

Submitted to J.B. Speed

School of Engineering of University of Louisville

in Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

Department of Electrical and Computer Engineering

University of Louisville

Louisville, Kentucky

December 2012

**PHENOMENOLOGICAL MODELING OF
IMAGE IRRADIANCE FOR
NON-LAMBERTIAN SURFACES UNDER
NATURAL ILLUMINATION**

By

Shireen Y. Elhabian

B.Sc. 2002, Faculty of Computers and Information, Cairo University

M.Sc. 2005, Faculty of Computers and Information, Cairo University

A Dissertation

Submitted to J.B. Speed

School of Engineering of University of Louisville

in Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

Department of Electrical and Computer Engineering

University of Louisville

Louisville, Kentucky

December 2012



© Copyright by Shireen Y. Elhabian
All Rights Reserved

PHENOMENOLOGICAL MODELING OF IMAGE
IRRADIANCE FOR NON-LAMBERTIAN SURFACES UNDER
NATURAL ILLUMINATION

By

Shireen Y. Elhabian

B.Sc. 2002, Faculty of Computers and Information, Cairo University

M.Sc. 2005, Faculty of Computers and Information, Cairo University

A Dissertation Approved on

November 26, 2012


by the Following Reading and Examination Committee:




Aly A. Farag, Ph.D., Dissertation
Director




Thomas L. Starr, Ph.D.



James H. Graham, Ph.D.



Robert W. Cohn, Ph.D.



Hossam Eldin H. Abd El Munim,
Ph.D.



DEDICATION

In the loving memory of my dear mother (may ALLAH shower His mercy on her), to my great father and to my beloved husband.

ACKNOWLEDGEMENTS

This research was developed by the grace of *Almighty God*, the merciful, the compassionate for the uncountable gifts given to me, who gave me the knowledge and wit to finish and establish this research.

My sincere gratitude and thanks go to my advisor, **Prof. Aly A. Farag**, for the support and encouragement he gave me during my study at the University of Louisville. Thank you for introducing me to Computer Vision and giving me the opportunity to explore this captivating field. Your ongoing confidence in me and my work has been an inspiration. It has been an honor to be one of your Ph.D. students. I appreciate your support to make my Ph.D. experience productive and stimulating. You always aim high, "*sky is the limit*", and I hope, with this dissertation, I met a portion of your expectations. *Thank you Doc.*

I would like to acknowledge the other members of my Ph.D. committee for spending time and effort in reading and reviewing my work. The insights and differing points of view of **Prof. Thomas L. Starr** were a great help in developing new ideas, while his simultaneous enthusiasm and support were also much appreciated. I would like to acknowledge **Prof. Robert Cohn** for giving me resources which helped me learn about surface measures and surface topographic relief.

During my course of study I have been privileged by the interaction with **Dr. Petter Nillius**, a researcher in the Physics of Medical Imaging lab at the Department of Physics at KTH, who helped me understand bits and pieces from his Ph.D. thesis. He never hesitated replying my emails especially during Christmas time (2011). Communicating with him have been profoundly an educative experience.

During my study at UofL, I was surrounded with people and technology which made my research possible. I owe a great debt to **William Michael Miller**, who was always there with stimulating discussions and new suggestions. It has been a great privilege to work with someone like him. *Thanks, Mike.* Thanks to **Chuck Sites** for his technical support and guidance.

I would also like to thank **Dr. Asem Ali** and **Dr. Moumen Elmelegy** with whom I enjoyed very fruitful and stimulating research collaborations. I have enjoyed working with **Ham Rara**; our discussions were always a great pleasure, and led to significant improvements in the content and exposition of this work. He gave important and fruitful guidance during my first steps into harmonic theory. His concepts have had a remarkable influence on my entire dissertation research. *Thank you, Ham.*

All members of the Computer Vision and Image Processing Lab have helped in making these past few years a thoroughly enjoyable experience for which I am truly grateful. It has been an absolute pleasure and privilege to work with all of them. The group has been a source of friendships as well as good advice and collaboration. I am especially grateful to **Mostafa Abdelrahman**, **Eslam Mostafa**, **Aly Abdelrehim**, **Cambron Carter**, **Ali Helmi** and **Ahmed Shalaby** for their great help and patience

in data acquisition and fruitful discussions. I really appreciate the help of **Ahmed ElBarkouky** who tolerated my discussions related to mathematical derivations. I would like to acknowledge **Travis Gault** with whom I shared fruitful discussions related to dissertation preparation and writing along with enormous technical support.

Research life is full of ups and downs and no one can move on just being alone. One of the great things that happened to me is meeting someone like **Marwa Ismail** whom I consider a sister and a friend. Both of us share moments of frustrations as well as happiness. I wish her the best for her life and family. Meanwhile, I have enjoyed and enriched during working with her in the virtual colonoscopy project.

I would also like to acknowledge my new friend and little sister, **Heba Farag** who spreads a spirit of joy and pleasure to her surroundings. I really admire her pure spirit as well as her baby acts. I am going to miss you so much.

Being far from my own home country and family, **Mrs. Salwa Elshazly** and **Amal Farag** have not hesitated to provide me guidance and help to adapt living in the States during my first two years in Louisville. I still remember moments when I became frustrated being alone and find Mrs. Salwa cheering me up. Thank you both for all great memories you have been a part of.

I am forever indebted to my parents, **Youssef Elhabian** and **Eman Mazyouna**, for their love, support, guidance and unfailing faith in me. Words cannot do justice to their impact in forging my personal and academic outlook. I am sure my mother would have been very happy seeing me graduating, may ALLAH shower His mercy on her. I would also like to acknowledge my parents-in-law for their love and encouragement.

I have been blessed with a marriage to one of the most decent and lovable persons God ever created. My beloved husband, **Samir Abdelrahman**, supported me in times of doubt, and has enriched my life beyond measure. He has cheerfully accommodated all manner of long hours and inconvenient absences. His unconditional love and constant encouragement are what keep me going. Basically he was the person I counted on for support and encouragement. No words can express my love and appreciation for him.

ABSTRACT

PHENOMENOLOGICAL MODELING OF IMAGE IRRADIANCE FOR NON-LAMBERTIAN SURFACES UNDER NATURAL ILLUMINATION

Shireen Y. Elhabian

November 26th, 2012

Various vision tasks are usually confronted by appearance variations due to changes of illumination. For instance, in a recognition system, it has been shown that the variability in human face appearance is owed to changes to lighting conditions rather than person's identity. Theoretically, due to the arbitrariness of the lighting function, the space of all possible images of a fixed-pose object under all possible illumination conditions is infinite dimensional. Nonetheless, it has been proven that the set of images of a convex Lambertian surface under distant illumination lies near a low dimensional linear subspace. This result was also extended to include non-Lambertian objects with non-convex geometry. As such, vision applications, concerned with the recovery of illumination, reflectance or surface geometry from images, would benefit from a low-dimensional generative model which captures appearance variations w.r.t. illumination conditions and surface reflectance properties. This enables the formulation of such inverse problems as parameter estimation.

Typically, subspace construction boils to performing a dimensionality reduction scheme, e.g. Principal Component Analysis (PCA), on a large set of (real/synthesized) images of object(s) of interest with fixed pose but different illumination conditions. However, this approach has two major problems. First, the acquired/rendered image ensemble should be statistically significant vis-à-vis capturing the full behavior of the sources of variations that is of interest, in particular illumination and reflectance. Second, the curse of dimensionality hinders numerical methods such as Singular Value Decomposition (SVD) which becomes intractable especially with large number of large-sized realizations in the image ensemble.

One way to bypass the need of large image ensemble is to construct appearance subspaces using phenomenological models which capture appearance variations through mathematical abstraction of the reflection process. In particular, the harmonic expansion of the image irradiance equation can be used to derive an analytic subspace to represent images under fixed pose but different illumination conditions where the image irradiance equation has been formulated in a convolution framework. Due to their low-frequency nature, irradiance signals can be represented using low-order basis functions, where Spherical Harmonics (SH) has been extensively adopted.

Typically, an ideal solution to the image irradiance (appearance) modeling problem should be able to incorporate complex illumination, cast shadows as well as realistic surface reflectance properties, while moving away from the simplifying assumptions of Lambertian reflectance and single-source distant illumination. By handling arbitrary complex illumination and non-Lambertian reflectance, the appearance model proposed in this dissertation moves the state of the art closer to the ideal solution.

This work primarily addresses the geometrical compliance of the hemispherical basis for representing surface reflectance while presenting a compact, yet accurate representation for arbitrary materials. To maintain the plausibility of the resulting appearance, the proposed basis is constructed in a manner that satisfies the Helmholtz reciprocity property while avoiding high computational complexity. It is believed that having the illumination and surface reflectance represented in the spherical and hemispherical domains respectively, while complying with the physical properties of the surface reflectance would provide better approximation accuracy of image irradiance when compared to the representation in the spherical domain.

Discounting subsurface scattering and surface emittance, this work proposes a surface reflectance basis, based on hemispherical harmonics (HSH), defined on the cartesian product of the incoming and outgoing local hemispheres (*i.e. w.r.t.* surface points). This basis obeys physical properties of surface reflectance involving reciprocity and energy conservation. The basis functions are validated using analytical reflectance models as well as scattered reflectance measurements which might violate the Helmholtz reciprocity property (this can be filtered out through the process of projecting them on the subspace spanned by the proposed basis, where the reciprocity property is preserved in the least-squares sense).

The image formation process of isotropic surfaces under arbitrary distant illumination is also formulated in the frequency space where the orthogonality relation between illumination and reflectance bases is encoded in what is termed as *irradiance harmonics*. Such harmonics decouple the effect of illumination and reflectance from the underlying pose and geometry. Further, a bilinear approach to analytically construct irradiance subspace is proposed in order to tackle the inherent problem of small-sample-size and curse of dimensionality. The process of finding the analytic subspace is posed as establishing a relation between its principal components and that of the irradiance harmonics basis functions. It is also shown how to incorporate prior information about natural illumination and real-world surface reflectance characteristics in order to capture the full behavior of complex illumination and non-Lambertian reflectance.

The use of the presented theoretical framework to develop practical algorithms for shape recovery is further presented where the hitherto assumed Lambertian assumption is relaxed. With a single image of unknown general illumination, the underlying geometrical structure can be recovered while accounting explicitly for object reflectance characteristics (*e.g.* human skin types for facial images and teeth reflectance for human jaw reconstruction) as well as complex illumination conditions. Experiments on synthetic and real images illustrate the robustness of the proposed appearance model vis-à-vis illumination variation.

Keywords: computer vision, computer graphics, shading, illumination modeling, reflectance representation, image irradiance, frequency space representations, (hemi)spherical harmonics, analytic bilinear PCA, model-based bilinear PCA, 3D shape reconstruction, statistical shape from shading.

TABLE OF CONTENTS

Dedication	iii
Acknowledgements	iv
Abstract	vi
Table Of Contents	viii
List of Figures	xiii
List of Tables	xvii
List of Algorithms	xviii
1 Prologue	1
1.1 Shading-based Inference	3
1.2 Modeling Image Formation	3
1.2.1 Illumination Modeling	3
1.2.2 Surface Reflectance Modeling	4
1.2.3 Image Irradiance Modeling	7
1.3 Problem Statement	12
1.4 Related Work	13
1.5 Research Questions	16
1.6 Contributions	18
1.7 Dissertation Organization	20
I Theoretical Background	23
2 Photometric Aspect of Image Formation	24
2.1 Light in Space	25
2.1.1 Hemisphere of Directions	25
2.1.2 Foreshortening	26
2.1.3 Solid Angle	27
2.1.4 Radiance	28
2.2 Light At Surfaces	30
2.2.1 Surface Irradiance	30

2.2.2	Surface Scattering	31
2.2.3	Surface Radiance	33
2.3	Image Irradiance	33
2.4	Assumptions	36
2.4.1	Homogeneous Convex Surfaces	36
2.4.2	Distant Illumination	36
2.4.3	Orthographic Projection	37
2.5	Image Irradiance Equation	38
2.6	Summary	38
3	Illumination Representation	39
3.1	Light Source Models	39
3.1.1	Point Sources	39
3.1.2	Distant Sources	41
3.1.3	Area Sources	43
3.1.4	Infinite Area Sources	43
3.2	Illumination in Frequency Space	47
3.2.1	Spherical Harmonics	47
3.2.1.1	Definition	47
3.2.1.2	Rotation of Spherical Functions	49
3.2.1.3	Rotational Invariance	52
3.2.2	Illumination Harmonic Expansion	54
3.3	Illumination Signal Energy Content	55
3.4	Frequency Properties of Light Sources	57
3.4.1	Distant Sources	58
3.4.2	Point Sources	59
3.4.3	Infinite Area Sources	60
3.5	Summary	60
II	Image Irradiance Representation	63
4	Phenomenological Representation of Arbitrary Reflectance	64
4.1	Introduction	65
4.2	Related Work	69
4.3	Hemispherical Harmonics (HSH)	71
4.3.1	Shifted Associated Legendre Polynomials	71
4.3.2	HSH Definition	72
4.4	Proposed Helmholtz HSH-based Reflectance Basis	75
4.4.1	Basis for Arbitrary Reflectance	76
4.4.2	Reflectance Basis For Isotropic Surfaces	80
4.5	On Basis Complexity	82
4.6	Representation of Analytic Reflectance Models	84
4.6.1	Ideal Diffuse Reflectance	86
4.6.2	Ideal Specular Reflectance	88

4.6.3	Example of Non-Ideal Physical Reflectance Models	94
4.7	Modeling Scattered Reflectance Data	98
4.7.1	Experimentation on Isotropic-Merl Reflectance Data	99
4.7.2	Experimentation on CURET Reflectance Data	103
4.8	Directional Hemispherical Reflectance	104
4.8.1	Directional Hemispherical Basis	104
4.8.2	Experimentation on Compact Description of BRDF Physical models	107
4.9	Summary	111
5	Image Irradiance Harmonics	114
5.1	Introduction	114
5.2	Related Work	117
5.3	Irradiance Harmonics Definition	120
5.4	Irradiance Harmonics for HSH-based Reflectance Basis	122
5.5	Irradiance Signal Energy Content	123
5.6	Experimental Results	125
5.7	Summary	128
6	Analytic Appearance Subspace Construction	134
6.1	Introduction	135
6.2	Analytic Bilinear Subspace Construction	139
6.2.1	Problem Formulation	140
6.2.2	Analytic Bilinear PCA Derivation	141
6.2.3	Model-based Bilinear PCA	144
6.3	Connection with Analytic PCA	145
6.4	Time and Space Complexities	147
6.4.1	Compression Ratio	147
6.4.2	Time Complexity	148
6.5	Experimental Results	149
6.5.1	Effect of Initialization	150
6.5.2	Effect of Irradiance Harmonics	150
6.5.3	Effect of Q%	152
6.5.4	Effect of Noise	152
6.6	Summary	154
7	Harmonics Projection from Incomplete Irradiance	161
7.1	Introduction	161
7.2	Harmonic Projection Image Definition	164
7.3	HP Images With Missing Information	166
7.3.1	Irradiance Corruption as Statistical Outlier	166
7.3.2	Irradiance Corruption as Erroneous Continuous Region	168
7.4	Experimental results	171
7.4.1	HP Images Under Occlusion	173
7.4.2	Shape Recovery Under Occlusion	174

7.5	Summary	178
III	Applications	180
8	Model-based Facial Shape Recovery from Single Image	181
8.1	Introduction	181
8.2	Image Irradiance Basis for Torrance-Sparrow Reflection	185
8.3	Harmonic Projection Images	188
8.4	Model-based Shape Recovery	190
8.5	Experimental Results	191
8.5.1	Model Construction	191
8.5.2	Synthetic Images	192
8.5.3	Real Images	193
8.6	Summary	194
9	Clinical Crowns Shape Reconstruction - An Image-based Approach	197
9.1	Introduction	198
9.2	Related Work	200
9.3	Contributions	202
9.4	Anatomical Jaw Landmarks - Proposed Definition	203
9.4.1	Dental Nomenclature	203
9.4.2	Landmarks Definition	205
9.4.2.1	Teeth Landmarks	206
9.4.2.2	Teeth-Gum Landmarks	206
9.4.3	Landmark Localization in Optical Images	207
9.5	Beyond Lambertian Reflection	207
9.5.1	Microscopic Surface Model	207
9.5.2	Towards a Realistic Model for Human Teeth Reflectance	210
9.6	Jaw Statistical Model Construction	214
9.6.1	Shape Model	215
9.6.2	Albedo Model	217
9.7	Model-based Appearance Subspace Construction	219
9.8	Model-based Shape and Albedo Recovery	221
9.8.1	SFS Brightness Constraint	221
9.8.2	Harmonic Projection (HP) Irradiance Model	224
9.8.3	Shape and Albedo Recovery	225
9.9	Application to Dental Restoration	227
9.10	Experimental Results	229
9.10.1	Visible Crowns Reconstruction	230
9.10.2	Results on Tooth Restoration	232
9.11	Summary	232

10 Epilogue	235
10.1 Summary of Contributions	236
10.2 Limitations and Suggested Future Directions	238
10.3 Summary	240
References	241
Appendix A: The Polar Integral of Shifted Associated Legendre Polynomials	258
Appendix B: Torrance Sparrow Spectrum Coefficients	260
Appendix C: Approximation Accuracy of Directional Hemispherical Basis	264
Appendix D: Analytic Derivation of $C_{npr}^q(\alpha)$	267
Appendix E: Irradiance Signal Energy Content	272
Appendix F: Proof of Analytic Bilinear PCA Theorem	279
Appendix G: Proof of Analytic PCA Theorem	282
Appendix H: (Hemi)Spherical Functions	283
Appendix I: Projection Onto Basis - Monte Carlo Integration	288
Nomenclature	309
Curriculum Vitae	310

LIST OF FIGURES

1.1	Image formation trilogy	2
1.2	A simplified image formation model	5
1.3	Illumination cone illustration	7
1.4	Illustration of the first approach for linear subspace construction	8
1.5	Illustration of the second approach for linear subspace construction	10
1.6	Illustration of the third approach for linear subspace construction	12
1.7	Dissertation Layout Illustration	21
2.1	How a surface point sees its surrounding world	26
2.2	Foreshortening illustration	27
2.3	How a light source sees a surface patch	28
2.4	Solid angle illustration	29
2.5	Local hemisphere	31
2.6	Shadows illustration	34
2.7	Integral domains illustrations	35
3.1	Rendering with point light source	40
3.2	Point light source illustration	41
3.3	Rendering with distant light source	42
3.4	Area light source illustration	44
3.5	Rendering with area light sources	45
3.6	Rendering with infinite area light sources	46
3.7	Environment maps parameterizations	47
3.8	Spherical harmonics illustration	50
3.9	Visualization of spherical harmonics basis functions	51
3.10	Illustration of 3D rotation using Euler angles.	52
3.11	Frequency properties of Debevec illumination maps	61
3.12	Frequency properties of sIBL illumination maps	62
4.1	Hemispherical harmonics illustration	73
4.2	Visualization of hemispherical harmonics basis functions	74
4.3	Visualization of proposed reflectance basis functions	79
4.4	Visualization of proposed isotropic reflectance basis functions	81
4.5	Number of basis functions	83
4.6	Surface types illustration	85

4.7	Reconstruction of the ideal specular reflectance	92
4.8	Reconstruction of the ideal specular reflectance with different incident angles	93
4.9	Torrance-Sparrow spectrum	98
4.10	Average approximation accuracy of Torrance Sparrow reflection model	99
4.11	Approximation accuracy of Merl database materials	100
4.12	Required minimum reflectance order for the materials in the Merl database	101
4.13	Approximation accuracy of CURET database materials	102
4.14	Required minimum reflectance order for the materials in the CURET database	103
4.15	Oren-Nayar BRDF approximation accuracy	108
4.16	Cook-Torrance BRDF sample reconstruction	110
4.17	Cook-Torrance BRDF approximation accuracy	110
4.18	Cook-Torrance approximation accuracy versus order	111
4.19	Cook-Torrance approximation accuracy with different specular reflectivity	112
5.1	Illustration of an object's surface is illuminated under distant lighting function	115
5.2	Lambertian attenuation factors	125
5.3	Attenuation factors of single reflectance modes	126
5.4	Average approximation accuracy of image irradiance as a function of the illumination order	127
5.5	Rendered sphere images of sample BRDFs	128
5.6	Rendered Camel images of beige-fabric BRDF	129
5.7	Rendered Hippo images of blue-metallic paint BRDF	130
5.8	The irradiance approximation accuracy of CURET database materials	131
5.9	The irradiance approximation accuracy of Merl database materials	132
5.10	Required minimum illumination order for the materials in the CURET database	133
6.1	Numeric subspace construction	135
6.2	Subspace construction to model illumination and reflectance	137
6.3	The average reconstruction error for each material in the Merl BRDF database	151
6.4	The average reconstruction error as a function of Q%	153
6.5	Sample of image reconstruction with different Q% using Ramamoorthi and Hanrahan basis	154
6.6	Sample of image reconstruction with different Q% using Nillius and Eklundh basis	155
6.7	Sample of image reconstruction with different Q% using proposed Helmholtz HSH-based basis	156
6.8	The average reconstruction error as a function of the signal-to-noise (SNR) ratio in dB	157

6.9 Sample of image reconstruction with different signal-to-noise (SNR) ratios using Ramamoorthi and Hanrahan basis	158
6.10 Sample of image reconstruction with different signal-to-noise (SNR) ratios using Nillius and Eklundh basis	159
6.11 Sample of image reconstruction with different signal-to-noise (SNR) ratios using proposed Helmholtz HSH-based basis	160
7.1 Illustration for harmonic-based image reconstruction	162
7.2 Harmonics Projection (HP) Images	165
7.3 Samples of occluded images	166
7.4 Geman-McClure function	167
7.5 Lorentzian function	167
7.6 \mathbf{f} defines the spatial support of the error ϵ	168
7.7 Results of harmonic projection with incomplete irradiance	174
7.8 Sample results of harmonic projection with different levels of occlusion .	175
7.9 Results of shape recovery with occlusion (height error)	176
7.10 Results of shape recovery with occlusion (surface orientation error) . . .	176
7.11 Shape and albedo recovery results	177
7.12 Shape and albedo recovery results on Yale database	178
7.13 Shape and albedo recovery results on Yale database with random occluder	178
8.1 Block diagram of the proposed model-based facial shape recovery	182
8.2 Error histograms of the recovered shapes	193
8.3 Recovered shapes for different out-of-training USF samples.	194
8.4 Recovered shapes for different YaleB Database subjects	195
9.1 Dental nomenclature	204
9.2 Teeth anatomy	204
9.3 Tooth surfaces	205
9.4 Illustration of the proposed human jaw anatomical landmarks	206
9.5 Surface height variations of anterior	208
9.6 Surface height variations of posterior	209
9.7 Teeth roughness parameter	210
9.8 Angle definitions of micro-facet models	211
9.9 The average approximation accuracy of Wolff-Oren-Nayar reflection model	214
9.10 A sample of extracted iso-surface and a re-meshed (coarsened) version .	215
9.11 A sample of estimated reflectance of an input occlusal image	218
9.12 A sample of albedo being warped to a jaw surface	218
9.13 Computations of the proposed appearance subspace construction algorithm	223
9.14 Block diagram of the proposed model-based human jaw shape recovery .	227
9.15 Illustrative example of inlay versus onlay restorations	228
9.16 Sample reconstruction result of an upper (post-repair) jaw (bottom row shows the top-view of the occlusal surface).	231
9.17 Sample reconstruction result of a lower (post-repair) jaw (bottom row shows the top-view of the occlusal surface).	231

9.18 Sample inlay reconstruction result of an upper (post-repair) jaw (bottom row shows a zoom-in view of the damaged teeth). 233

9.19 Sample onlay reconstruction result of a lower (post-repair) jaw (bottom row shows a zoom-in view of the damaged teeth). 234

10.1 The spherical polar coordinate system: it is defined by the polar angle from the z -axis θ and the azimuthal angle ϕ in the xy -plane. The distance from the origin of the coordinate system is defined by the radius r 284

10.2 An example of spherical function, $f^s(\theta, \phi) = \frac{1}{4} [\cos^3(6\theta) + \sin^4(\phi) + 1]$ visualized as (a) a unit sphere being deformed such that the distance between the sphere's center and the surface points encodes the function absolute values at different spherical coordinates, this is accomplished by the relation between the spherical and cartesian coordinates, (b) a textured unit sphere where the color of its surface points encodes the function values. Note that $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi]$ 285

10.3 An example of hemispherical function, $f^h(\theta, \phi) = \frac{1}{4} [\cos^3(6\theta) + \sin^4(\phi) + 1]$ visualized as (a) a unit hemisphere being deformed such that the distance between the hemisphere's center and the surface points encodes the function absolute values at different spherical coordinates, this is accomplished by the relation between the spherical and cartesian coordinates, (b) a textured unit hemisphere where the color of its surface points encodes the function values. Note that $\theta \in [0, \pi/2]$ and $\phi \in [0, 2\pi]$ 285

10.4 The projection of $f(x) = 1.75x^3 + 0.6x^2 + 0.21x - 0.06$ in red onto normalized Legendre polynomials up to the third order, where the expansion/projection coefficients are evaluated using Monte Carlo integration. 293

10.5 The reconstruction of $f(x) = 1.75x^3 + 0.6x^2 + 0.21x - 0.06$ (dashed red) using the normalized Legendre polynomials up to the third order. . . . 293

10.6 Left: the amount of energy captured by each individual expansion coefficient, evaluated as $\frac{a_n^2}{\int_{-1}^1 |f(x)|^2 dx}$. Right: function approximation accuracy e^N of N -order function approximation using normalized Legendre polynomials, where $e^0 = 0.0301$, $e^1 = 0.8433$, $e^2 = 0.8925$ and $e^3 = 1.00$ 294

LIST OF TABLES

1.1	Linear Subspace Construction Approaches for Appearance Modeling . . .	11
1.2	Different Approaches for Illumination Cone Approximation	17
5.1	Assumptions and Harmonic Basis of Different Formulations of Image Irradiance Equation (θ'_i is the angle between surface normal and incident light direction)	119
8.1	Mean and standard deviation of the mean height error for 80 out-of-training USF samples under Yale light conditions	192
8.2	Mean and standard deviation of the mean surface orientation error for 80 out-of-training USF samples under Yale light conditions	192
9.1	Average <i>whole jaw</i> surface reconstruction accuracy (RMS) in <i>mm</i> . . .	232
9.2	Average <i>inlay</i> surface reconstruction accuracy (RMS) in <i>mm</i>	233
9.3	Average <i>onlay</i> surface reconstruction accuracy (RMS) in <i>mm</i>	233
10.1	Legendre Polynomials up to 8-th order	292

LIST OF ALGORITHMS

1	d^n Matrices Computation	55
2	Model-based Analytic Bilinear PCA	146
3	Compute Harmonic Projection (HP) Image	165
4	Compute Robust Harmonic Projection (HP) Image	168
5	Compute Harmonic Projection (HP) Image with MGRF and Sparsity	172
6	Principal Component Regression Framework for 3D Shape Recovery	190
7	Model-based Appearance Subspace Construction of Human Visible Crowns	222
8	Model-based Shape and Albedo Recovery	226

PROLOGUE

The human vision system is one of the gates to the external world, allowing one to navigate through and interact with the surrounding environment, detecting and recognizing objects and people. In a glimpse of human's eyes, information about the surrounding objects is available including their shapes, materials and relative poses.

Being inspired from human and other biological visual systems, computer vision aims at making the computer *see*. This involves the extraction of information about the three dimensional world given two dimensional images. This is analogical to the human visual system which uses several cues to infer the same information from the 2-D images projected onto the retina. Different cues can be used to accomplish such a task. Looking at the world from slightly different positions, providing different views, gives an important depth cue known as the *disparity cue*. Another cue is the *motion parallax*, where images of nearby objects appear to move faster than these of distant objects.

In some situations, a single still image is available, e.g. a passport image, leaving the image brightness to process. While edges and corners have been widely used in computer vision, such features cannot explain local shape perception provided by the *shading cue*. Shading can be defined as the gradual variation of image brightness as a function of the surface normal relative to the light source and the imaging sensor. Such brightness variation mainly depends on the surface reflectance properties and the illumination conditions. The extraction of information from shading primarily depends on understanding the image formation process; defined as the response of an imaging sensor to radiation. The core motivation behind the presented work is formulating an imaging model which describes the relationship between object's attributes (shape and reflectance properties) and the formed image brightness.

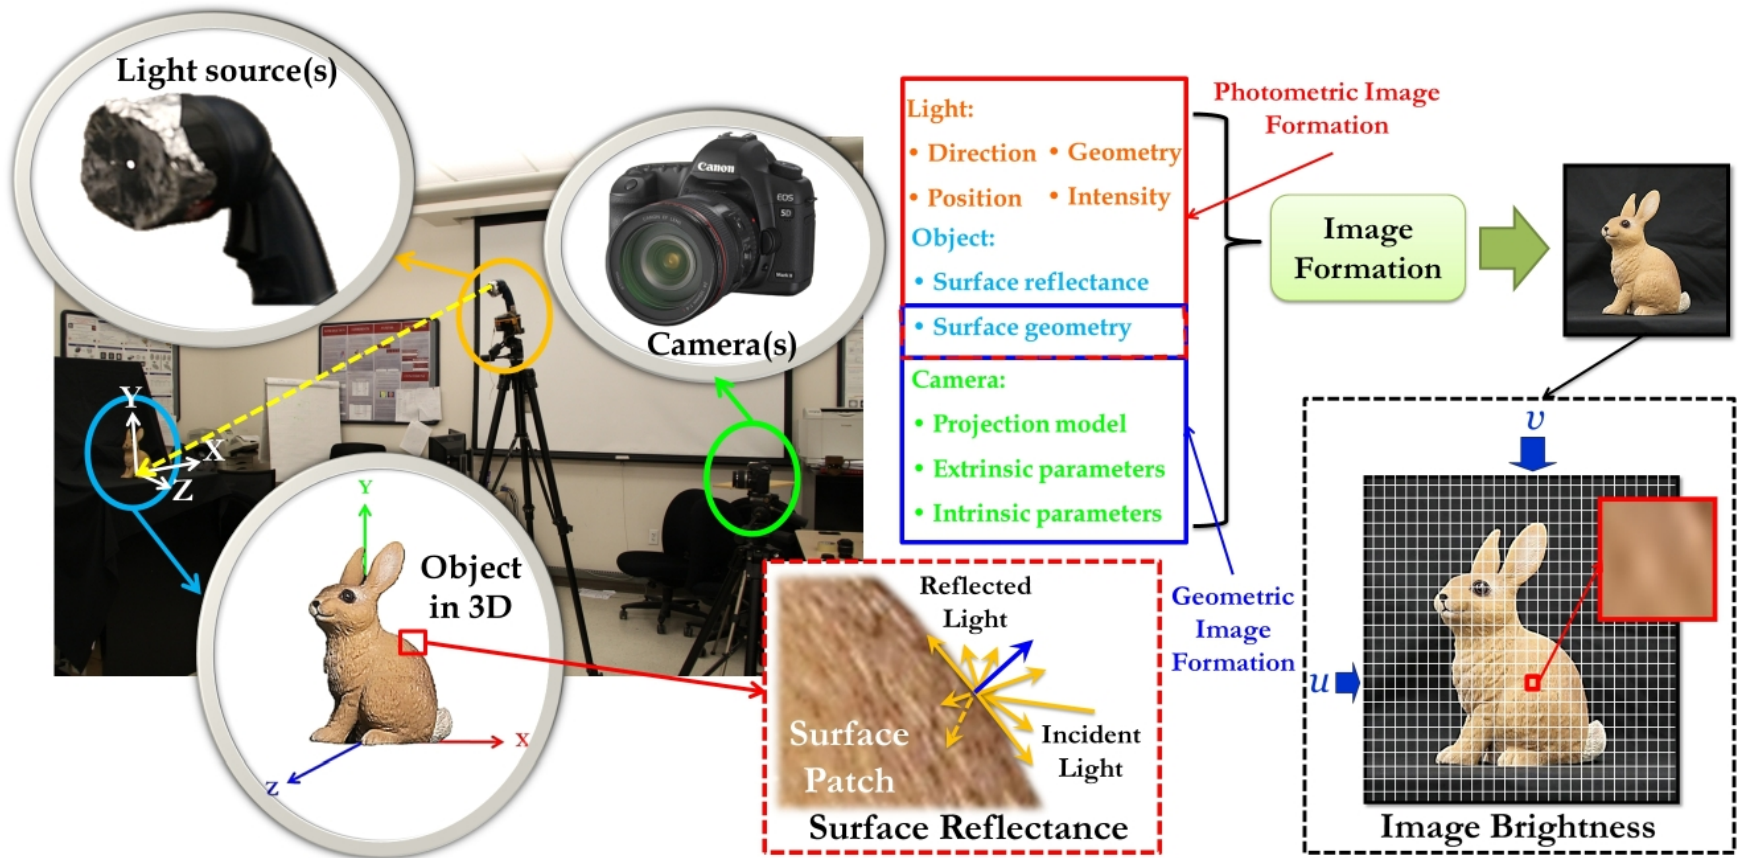


Figure 1.1: Image formation triology, light source(s) are used to illuminate a 3D object in space which will be captured by an imaging system, i.e. camera, to form the captured/observed image. The formed image depends on three main components; (1) The camera (viewer) which is characterized by its projection model (orthographic versus perspective projection), intrinsic and extrinsic parameters, (2) Illumination (light source(s)) used to illuminate the object to be captured, characterized by direction, position, geometry and intensity, and (3) The object itself defined by its 3D structure, i.e. shape/surface geometry, and surface reflectance properties. The formed image at the end is a two-dimensional pattern of brightness.

1.1 Shading-based Inference

Shading is the visible output of the physical process of image formation, depending on the the shape of the surface, the surface' reflectance properties, the lighting conditions and the parameters of the imaging sensor. Physics-based vision [1] is concerned about extracting information about an image's contents based on an understanding of the underlying physics which governs how the image was formed. However, such a problem is under-constrained because the formation process involves numerous unknowns in contrast to the available information at hand, *i.e.* image brightness value per pixel. Thus information inference based on shading commonly involves adding assumptions to the image formation process such as distant illumination and homogenous reflectance (where surface patches are assumed to have the same reflectance properties).

Under the assumption of fixed pose and fixed camera parameters, shading-based inference involves the three components of the photometric image formation (shape, reflectance properties and illumination). Usually it is assumed that one or more components are known beforehand, e.g. Lambertian reflectance and known point light source were assumed in the pioneering work of Horn [2] to solve for the surface shape. The main contribution sought in this dissertation work is to present appearance models and representations which can be used for non-Lambertian surface reflectance and complex illumination.

1.2 Modeling Image Formation

Image formation can be defined as the response of an imaging sensor to radiation. It consists of two main aspects [3]; (1) *Geometric* image formation aspect which determines the location of a 3D point measured in the world frame (coordinate system) when projected onto the image plane, and (2) *Photometric* image formation aspect which determines the brightness of a point in the image plane relying on the interplay between the light scattering and the surface reflectance properties. See Figure 1.1.

1.2.1 Illumination Modeling

The simplest illumination model is a point light source which provides a good approximation to direct sunlight. Although inadequate for most real situations, it has been

widely adopted in most work on shading. Formally, illumination is a spherical function, defining the intensity of incident light from all possible incident directions. Such a function can never be recovered from the shading alone. Nonetheless, it has been found that surfaces acts as a low-pass filter on the incident illumination [4, 5], yielding a band-limited reflected light. This introduced the frequency-space illumination representation when D’Zmura [5] expressed the shading equations using spherical harmonics (SH) as an analogy to Fourier series on the real line. Further, independent works of (Basri and Jacobs) in [6] and (Ramamoorthi and Hanrahan) in [7] used SH to analyze the reflected light for convex Lambertian surfaces where they showed that it can be approximated as a convex combination of up to 2nd order SH (*i.e.* nine basis functions), while the appearance of non-Lambertian surfaces depends on higher frequencies of the illumination.

1.2.2 Surface Reflectance Modeling

Modeling how light is reflected from surfaces is a central theme in both computer graphics and computer vision. When a light ray (measured by radiance) hits a surface, three effects might occur; (1) it may be absorbed inside the surface, (2) it may be reflected from the surface, and (3) it may be scattered at various depths/levels inside the surface. See Figure 1.2 for illustration. Usually a combination of these effects occurs.

The situation can even be more complicated through the tendency of some surfaces to absorb light at one wavelength and then radiate it at a different wavelength. This effect is known as fluorescence. In addition, a warm enough surface can emit light in the visible range.

It is common in vision to assume that all effects are local with neither emission nor fluorescence. A reasonable model for such kind of surfaces assumes the following [8]; (1) the radiance leaving a point on a surface is due to the radiance arriving at this point, (2) light leaving a surface at a given wavelength is due to light arriving at that wavelength, (3) light sources are treated separately, and (4) surfaces do not generate light internally.

How the surface reflects the incident light is described by notion of the bidirectional reflectance distribution function (BRDF) [9]. It is considered the most general and common mathematical model of local reflection which entails how much light is reflected in an outgoing direction due to light being incident from another direction. This function is specified by two directions (outgoing and incoming), hence the name bidirectional.




Figure 1.2: A simplified image formation model. When a light ray hits the surface of an object in space, a portion of light energy is absorbed by the object’s surface and the remaining is reflected from the surface in different directions. The reflected light rays in the direction of the camera affect the brightness of the projected point in the image plane.

It ranges from 0, where there is no outgoing/reflected radiance at the exit direction, to infinity, where arbitrary small irradiance is received from the incoming direction. The BRDF is symmetric in the incoming and outgoing directions; this is referred to as *Helmholtz reciprocity* principle. An important subclass is the *isotropic* BRDF where the surface appearance remains the same under rotations of the tangential plane about the surface normal. In such a case, BRDFs become functions of only three angles (the inclination angles of the incident and reflected illuminations and the absolute difference of their azimuthal angles). Although isotropy is not a universal property for all surfaces, many real surfaces exhibit such a property, while the reflectance of anisotropic surfaces can be approximated as being isotropic [10].

BRDF is also a function of position; however, sometimes this positional variance is not included in the description of a BRDF, where position/shift-invariant BRDFs are assumed. When the spatial position is not included as a parameter to the BRDF function, an assumption is implicitly made that the reflectance properties of a surface do not vary with the spatial position, this is only valid for homogeneous surfaces having homogeneous materials. Non-homogeneous surfaces with spatially varying reflectance properties are often approximated in vision applications by using a spatially varying texture which modulates the surface radiance [11].

In literature, analytic reflectance models (physically-based analytic [12, 13, 14] or empirical [15, 16]) are used to provide BRDFs to approximate the reflectance of real materials. Yet such models only describe the phenomena for which they are designed. The simplest model is the Lambertian one (perfect diffuser) where the surface is assumed to distribute the reflected light equally over all outgoing directions regardless of the incoming direction. At the other end of the spectrum, mirror-like surfaces are modeled by specular spike where the reflection only occurs along the mirrored incident direction. The surface roughness can also be captured through assuming that the surface exhibits a micro-facet geometry which accounts for self-shadowing and inter-reflections at micro-structure. Glossy rough surfaces can be modeled by Torrance-Sparrow [12] and Cook-Torrance [13] models, while diffuse rough surfaces can be modeled by Oren-Nayar model [14]. Such analytical models are based on material parameters which could be measured in principle, but difficult to acquire in practice [17].

An alternative path is to measure the surface reflectance for different combinations of incoming and outgoing directions then fit the measured data to specific analytic model through various optimization approaches [16, 18]. Nonetheless, the measure-and-fit approach is lacking in some aspects; the underlying assumption that there is an inherent noise in the measuring process which is filtered out by the fitting process ignores significant modeling errors due to the approximation imposed by the analytical model itself. Further, the choice of the error function in the optimization process is not intuitive¹ while the initial solution of the model parameters has a significant impact on the final solution of the fitting process.

Another approach to modeling reflectance is acquiring dense measurements of reflectance which are tabulated in a lookup table to be used as a BRDF. While this approach maintains BRDF detailed structure which might be lost in the fitting process, it becomes expensive in terms of lookup time and storage with increasing number of materials to be used in a system.

As a compromise between storage and representation, an arbitrary BRDF can be represented as a linear combination of a complete set of orthonormal basis functions, analogous to Fourier basis representing functions over the real line *e.g.* [19, 20, 21, 22].

¹Metrics depending on Euclidean distance tend to overemphasize the weight of the specular peaks in the optimization process [17].

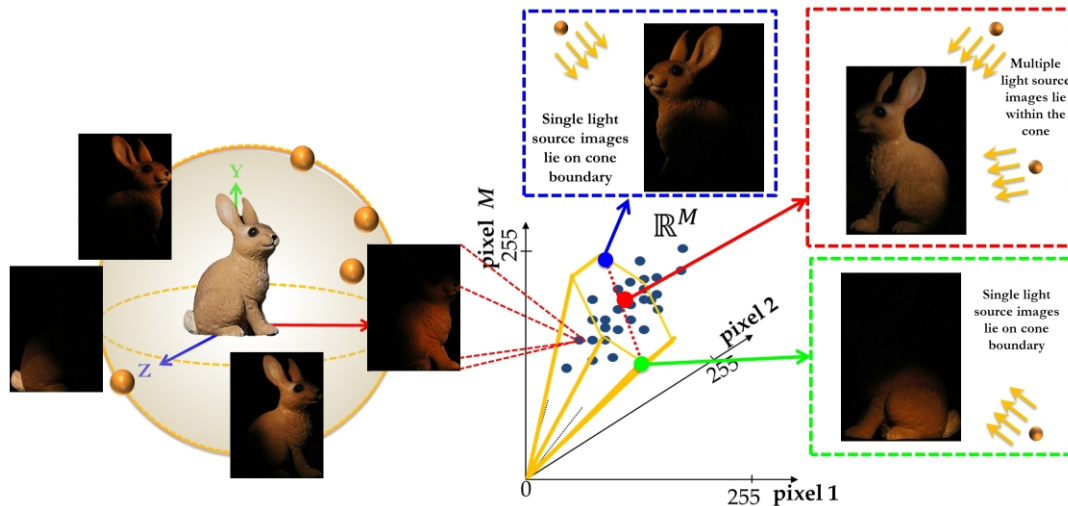


Figure 1.3: Images of the same object under fixed pose appear differently due to the change in the lighting conditions. The set of M -pixel images of any convex Lambertian object seen under all possible lighting conditions, while kept at fixed pose with respect to the imaging sensor, spans a convex cone known-as the *illumination cone* [23] in \mathbb{R}^M . In the imaging setup at the CVIP-UofL Lab, a coated torch (flash light) is used with a very tiny hole to narrow down light rays in order to simulate a single directional light source, modeled as a delta function. The light source and the camera are kept at a large distance compared to the object’s size, where viewer-centered coordinate system is used, to simulate distant illumination and far away viewer where orthographic projection can be assumed.

1.2.3 Image Irradiance Modeling

The appearance of an object under fixed pose depends primarily on the trilogy of photometric image formation process; the object’s geometrical structure (shape), the surface reflectance properties (material) and illumination, where the shading is given by the image irradiance equation [24]. Belhumer and Kriegman in [23] posed the question of finding some underlying generative structure to the set of images of a fixed-pose object under variable illumination conditions. While theoretically, the space of all such images is infinite dimensional due to the arbitrariness of the lighting functions [25], they proved that such set of images lie in a convex cone, termed as *illumination cone*, in the space of images.

Due to its convexity, the illumination cone for a given fixed-pose object is characterized by a finite set of *extreme rays* defining images of that object under appropriately chosen single distant point light sources. Whereas all other images inside the cone are formed by convex combinations of these extreme rays, see Figure 1.3 for illustration. This convex polyhedral cone can be used to generate and recognize images with novel

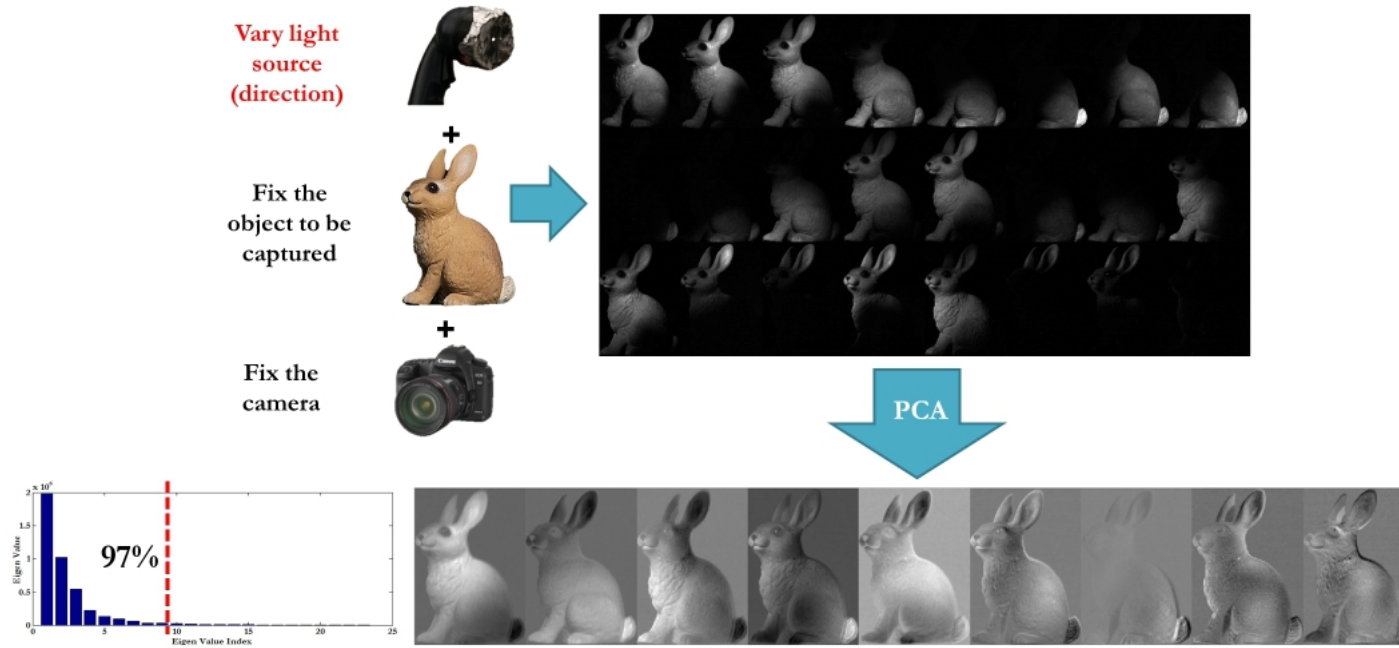


Figure 1.4: The linear subspace which represent all possible images of a specific object under all illumination conditions, while pose kept fixed, can be constructed by performing principal component analysis (PCA) on a large set of images of an object under different illumination conditions. The CVIP Lab bunny was imaged with fixed pose and fixed camera, while the light direction was varied. The light source (coated torch with a tiny hole in the middle to narrow down the light rays in order to simulate a delta function, *i.e.* single directional light source) was kept at a large distance compared to the object size to simulate distant/directional light source. PCA was then performed using 24 images with different light source directions. The eigenvalue spectrum is shown on the bottom left where nine-basis are found to maintain 97% of image variability, emphasizing low-dimensional representation of images of known object with fixed pose and camera but under arbitrary lighting conditions. This complies with the results of [6] and [7].

illumination conditions [26] while the relative pose to the imaging sensor is kept fixed.

Whereas Belhumer and Kriegman in [23] provided analytical justification for using low-dimensional linear subspaces to model image variation of convex surfaces under different illumination (this also has been validated by others, *e.g.* [6, 7, 25, 26]). This was also empirically justified by Yuille *et al.* [27] where they used a human face and other objects to perform Principal Component Analysis (PCA) on images acquired under a distant point light source moving on the locus of a sphere surrounding the object. They empirically illustrated that image variations due to illumination of Lambertian objects can be explained by the first five principal components. Thus linear subspaces can be used to represent a certain class of images, capturing variations due to different imaging conditions [6, 25, 26]. Further, Epstein *et al.* [28] showed that non-convex surfaces, with cast as well as attached shadows, are still well approximated by a relatively low-dimensional subspace, yet with a bit higher dimension compared to convex surfaces [29].

Typically, the illumination cone, as an object-specific representation, provides a means to predict object appearance under unseen imaging conditions. As such, its bases serve as a generative appearance model which can be used to render/synthesize images of a fixed-pose object under new illumination conditions, *i.e.* extrapolate appearance. While objects of the same class (*e.g.* human faces) have illumination cones with similar shapes [29], different objects would have different extreme rays associated to their individual illumination cones, hence such cones can be used to encode the identity of an object within a given class under varying illumination, *i.e.* object recognition can be performed by assigning the identity of the closest illumination cone based on Euclidean distance.

As outlined by Lee *et al.* [29], basis images spanning such a low-dimensional linear subspace can be obtained using three different ways; (1) Performing PCA on a large set of images of object(s) of interest with fixed pose but different illumination conditions, see Figure 1.4. This approach necessitates a fixed calibrated lighting rig while keeping the camera and the object fixed. Besides operating on a discrete set of images, this may produce biased subspace towards the particular illumination directions used. (2) Using 3D models (shape and texture), image formation process can be simulated to render synthetic images under assumed imaging conditions, while PCA is again used to compute the required subspace, see Figure 1.5. Such an approach needs photo-realistic rendering

algorithms, *e.g.* ray-tracing [30], to accurately simulate the effect of different illumination conditions with all possible configurations (3D rotations) on the formed images. Since the continuous space of illumination directions should be sampled, the constructed subspace might also be biased towards such sampling. (3) Assuming certain surface reflectance distribution function, the harmonic expansion of the image irradiance equation [2] can be used to derive an analytic subspace to approximate images under fixed pose but different illumination conditions, see Figure 1.6. While this approach does not involve complicated acquisition setup other than the availability of groundtruth shape (which can be obtained by Cyberware scanners or shape-from-X approaches) and reflectance (which can be obtained from publicly available BRDF databases, *e.g.* [10, 17]), the main challenge is to incorporate realistic assumptions in the image formation model. This involves non-Lambertian reflectance and complex illumination.

Table 1.1 summarizes different approaches for linear subspace construction used to approximate/represent the illumination cone using finite set of orthonormal basis. This work is stemmed from the third approach where an analytical approach is defined for constructing the illumination cone of non-Lambertian objects under fixed pose relative to the imaging sensor.

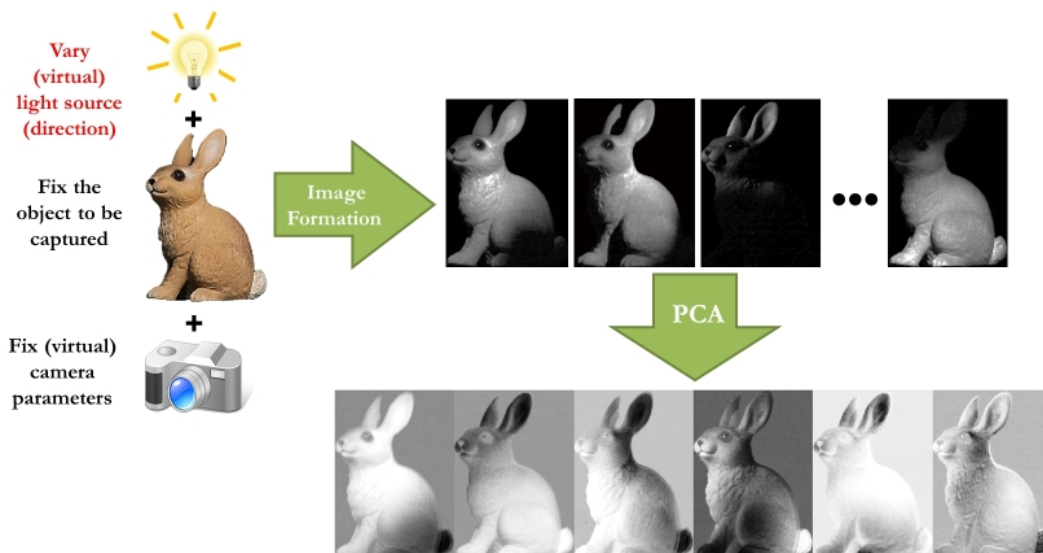


Figure 1.5: The linear subspace which represent all possible images of a specific object under all illumination conditions, while pose is kept fixed, can be constructed by simulating image formation process using 3D models to render synthetic images under various illumination conditions, while PCA is used to compute the required subspace. The Cyberware scanner was used to obtain the shape and albedo of the CVIP Lab bunny.

Table 1.1: Linear Subspace Construction Approaches for Appearance Modeling

	Object	Light	Camera	Image Formation	Basis Construction	Limitations and Challenges
Approach (1) [27]	Real/physical Object	Moving directional (physical) light source	Physical camera (orthographic projection assumed)	Physical image capturing	PCA on captured images	Special acquisition setup is needed to control illumination directions. Biased subspace construction towards the used illumination directions.
Approach (2) [26]	3D Model (acquired or estimated)	Moving directional (virtual) light source	Virtual camera (orthographic projection assumed)	Image (forward) rendering	PCA on synthesized images	Needs photo-realistic image formation algorithms. Biased subspace construction towards illumination directions sampling.
Approach (3) [6, 7, 11, 25, 31, 32]	3D Model	Analytically modeled	Camera model embedded in the image irradiance equation (orthographic projection assumed)	Analytically formulated	Harmonic expansion of image irradiance equation	The need of shape and albedo of the object of interest. Incorporating realistic assumptions in the image formation model

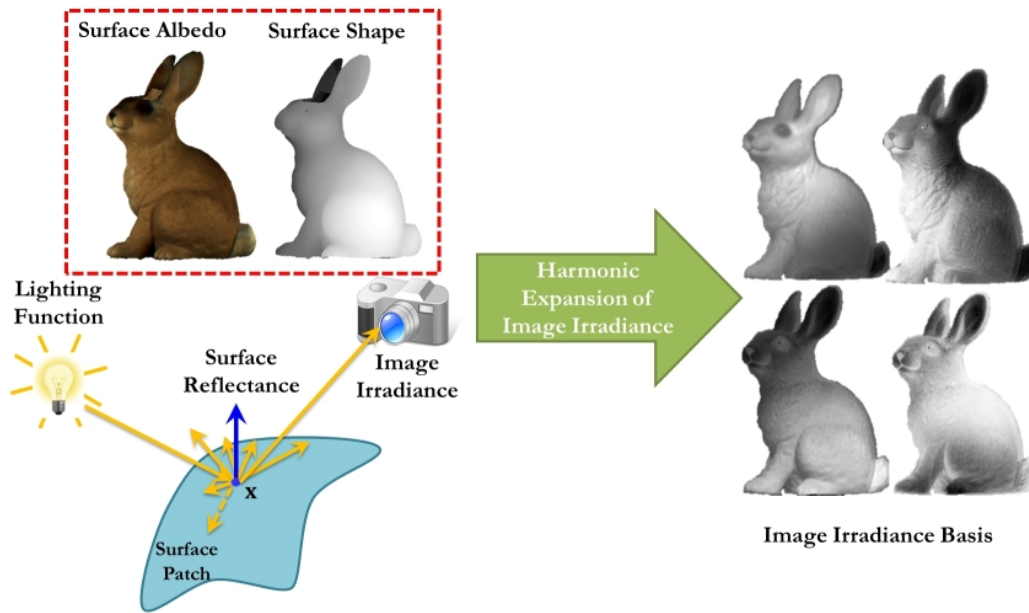


Figure 1.6: The linear subspace which represent all possible images of a specific object under all illumination conditions, while pose is kept fixed, can be constructed by assuming certain surface reflectance distribution function, the harmonic expansion of the lighting function can be used to derive an analytic subspace to approximate images under fixed pose but different illumination conditions. The Cyperware scanner was used to obtain the shape and albedo of the CVIP Lab bunny.

1.3 Problem Statement

A typical framework which estimates a low-dimensional subspace for object appearance involves sampling a finite set of the images in the illumination cone then use PCA. Nonetheless, this approach suffers from the need “large enough” image ensemble to decrease the bias of constructed subspace towards illumination sampling, yet this renders numerical methods intractable. Further, with a small-sample-size, a biased set of samples would produce an ineffective subspace in terms of encoding the object identity. Hence, recognition applications, for example, would benefit from an economical solution for subspace construction which resolves the trade-off between the small-sample-size problem and numerical estimation of appearance subspace basis. Lee *et al.* [29] addressed the issue of acquiring such subspaces by specifying a universal illumination configuration where appropriately chosen 9 point light sources were used to acquire images that are considered as a good approximation to the subspace basis images, *i.e.* extreme rays. While as few as 5 images can be acquired, this necessitates a calibrated lighting rig which might be possible in an industrial setting.

Typically, an ideal solution to the appearance modeling problem should be able to incorporate complex illumination, cast shadows as well as realistic surface reflectance properties, while moving away from the simplifying assumptions of Lambertian reflectance and single-source distant illumination. Current literature is comprised of models which relax one assumption while restricting to the others. Primarily, this work focuses on bringing out the underlying assumptions of the current image irradiance models with respect to the problem of modeling the appearance of objects under unknown general illumination while keeping the object at fixed pose relative to the imaging sensor. As such, the problem statement of this dissertation can thus be stated as follows.

Derive an analytic form of appearance subspaces which define the illumination cone, given 3D object model (shape and texture), under realistic imaging conditions, e.g. complex illumination and non-Lambertian reflectance.

In particular, the frequency space based representation of image irradiance accounts for complex illumination, thus a point light source assumption can be relaxed, leading to a finite-dimensional linear model to represent the image irradiance under arbitrary illumination.

This low-dimensional representation of image irradiance under unknown arbitrary lighting can be considered as a fundamental process for many computer vision tasks such as illumination modeling [6], surface reflectivity estimation/analysis [11, 32], statistical shape and albedo recovery [33, 34, 35], shape from shading [36], photometric stereo [37, 38], object detection and recognition [39], to name a few.

1.4 Related Work

By definition, the illumination cone of a given fixed-pose objects contains all image variations under all possible lighting conditions. As such, an accurate representation of such cone would benefit vision tasks which need appearance extrapolation under previously-unseen imaging conditions. While the basic building block of an illumination cone is the estimation of its extreme rays, their number is proportional to the number of independent (distinct) surface normal which can be very large (order of thousands). Hence, in order to fully define the exact illumination cone of an object, a large number of extreme rays that make up its cone is needed to be computed. For example, a convex-

Lambertian surface has $O(n^2)$ extreme rays where n is the number of distinct surface normals [29].

Nonetheless it was conjectured in [23] that the illumination cone for typical objects is flat, therefore it lies near a low-dimensional linear subspace of the image space. As such, the computation of the full cone may not be necessary where the illumination cone of a fixed-pose object can be approximated by a low-dimensional subspace whose basis can be estimated using a generative appearance model.

Georghiadis *et al.* [26] defined a generative appearance model for object recognition under variation in lighting. They used a small number of training images per subject taken under different *known* lighting directions with their variant of photometric stereo to reconstruct the corresponding shape and albedo. They were, in turn, used to render images under new illumination conditions which are used to estimate the basis images spanning the illumination cone using Singular Value Decomposition (SVD). Aside from assuming point light sources and Lambertian reflectance, this approach has inherent curse of dimensionality where a “large enough” rendered ensemble is needed to decrease the effect of illumination sampling on the resulting subspace.

Lee *et al.* [29] addressed the issue of arranging physical lighting so that the acquired basis images for each object can be used directly as basis for low-dimensional appearance subspace. In the context of face recognition, they showed that such physical arrangements are qualitatively very similar for different individuals. They demonstrated both theoretically and empirically that few real images taken under well-chosen lighting configurations can be directly used to model illumination effects. Yet the optimization process which solved for such lighting configurations (locations and directions) needs to uniformly sample light directions in addition to rendering images under single directional light sources to construct an initial set of extreme rays (≥ 1000 samples). With Lambertian assumption, such rendering is as simple as the dot product between surface normals and light directions. Generalizing their approach to non-Lambertian surfaces would require computationally expensive rendering module to initiate the optimization procedure.

Independent works of Basri and Jacobs in [6] and Ramamoorthi and Hanrahan in [7] provided analytical justification that the illumination cone of a convex-Lambertian surface can be accurately approximated by a 9D linear subspace derived based on spherical harmonics (SH). They formulated the image irradiance equation [2] in a convolution

framework where the lighting function acts as a signal filtered by the Lambertian kernel, opening the whole field of signal processing and filtering to the image irradiance modeling. They provided an analytical expression of an image of a convex-Lambertian object illuminated by distant lighting. They proved that the Lambertian reflectance kernel acts as a low pass filter, presenting the image irradiance as a band-limited signal being represented by a finite number of basis functions. Both [6] and [7] used SH to represent the surface reflectance function in the frequency domain. While Basri and Jacobs [6] formulated this process in global coordinates with respect to global reference frame, Ramamoorthi in [7, 31] made the distinction between such a formalization in global and local coordinates (with respect to surface points).

While [6] and [7] provided a single expression for the harmonic basis, Ramamoorthi in [25], under the assumption of distant light, linearly combined higher order SH basis using analytical PCA, to build an optimal orthonormal eigenfunctions (basis) for image spaces. On the other hand, QR-decomposition is used in [6] to numerically orthogonalize up-to 2nd order SH basis.

Recently Elhabian *et al.* [40, 41] avoided the use of higher order SH basis and numerical orthogonalization to model the reflectance of convex-Lambertian surfaces under single viewpoint. They used hemispherical harmonics (HSH) which forms an orthonormal set of basis over a unit hemisphere, thus PCA or basis orthogonalization is no longer needed. Nonetheless, the incident illumination is restricted to the upper global hemisphere facing the camera while Lambertian reflectance is assumed.

For non-Lambertian surfaces, Ramamoorthi and Hanrahan in [31] used SH basis functions to represent illumination and surface reflectance for real-time rendering of objects with complex isotropic non-Lambertian reflectance under distant natural illumination. However the surface reflectance of non-emitting surfaces (discounting subsurface scattering) is defined on the cartesian product of two hemispheres corresponding to the incident and outgoing directions, hence the natural way to represent such a hemispherical function is to use hemispherical basis. On the other hand, Nillius [32] used the hemispherical basis proposed by Koenderink and van Doorn [21]. Such bases are based on Zernike radial polynomials, where bases for the unit disk are mapped onto the upper hemisphere. Yet, such polynomials have high computational cost [20] when compared to associated Legendre polynomials used for (hemi)spherical harmonics. In particular, Zernike radial polynomials require an amount of CPU time proportional to $\mathcal{O}(N^2)$ [42]

in contrast to associated Legendre polynomials which needs $\mathcal{O}(N \log^2 N)$ [43] for cutoff frequency $N = n^2$ with order $n \geq 0$.

Table 1.2 summarizes main different approaches in literature which solves for the basis images spanning the illumination cone of a given object as compared to the proposed formulation.

1.5 Research Questions

By handling arbitrary complex illumination and non-Lambertian reflectance, the appearance model proposed in the presented dissertation will move the state of the art closer to the ideal solution. Discounting subsurface scattering and surface emittance, this work seeks to address the following research questions:

RQ1: How can image irradiance be efficiently and accurately represented, given that a surface point sees its surrounding world through a unit hemisphere? It has been shown that hemispherical functions present discontinuities at the boundary of the hemisphere when represented in the spherical domain. Consequently, the functional representation requires more coefficients for accurate representation.

RQ2: How can surface reflectance be represented in a manner that guarantees the Helmholtz reciprocity property, which should be maintained by physically plausible surface bidirectional reflectance distribution functions?

RQ3: How can illumination and reflectance be decoupled from the underlying geometry and pose?

RQ4: How can prior information about natural illumination and real world surface materials be incorporated to construct appearance models in an analytical manner?

RQ5: How can the curse of dimensionality for subspace construction be tackled?

RQ6: How can incomppliance to the imaging model assumptions such as non-convexity be handled?

RQ7: How can computer vision tasks benefit from the proposed imaging model?

Table 1.2: Different Approaches for Illumination Cone Approximation

	Input	Assumed Illumination	Assumed Reflectance	Appearance Subspace	Notes
Georghiades et al. [26]	7 images of known light directions	Point light sources	Lambertian	Numerical orthogonalization via SVD on rendered images	Needs large enough ensemble for unbiased subspace estimation
Basri and Jacobs [6]	3D object (shape and texture)	Complex illumination	Lambertian	Numerical orthogonalization via QR-decomposition of texture-modulated SH-based harmonic basis	Using spherical basis to represent hemispherical domain. Can't be extended to non-Lambertian surfaces unless local hemisphere is considered.
Ramamoorthi and Hanrahan [7]	3D object (shape and texture)	Complex illumination	Lambertian	No orthogonalization performed	Using spherical basis to represent hemispherical domain.
Ramamoorthi [25]	3D object (shape and texture)	Complex illumination	Lambertian	Analytical Linear PCA	Using spherical basis to represent hemispherical domain.
Ramamoorthi and Hanrahan [31]	3D object (shape and texture)	Complex illumination	non-Lambertian	No orthogonalization performed	Using spherical basis to represent hemispherical domain.
Nillius and Eklundh [32, 44]	3D object (shape and texture)	Complex illumination	non-Lambertian	Analytical Linear PCA	Zernike polynomials are known to have high computational complex in addition of being undefined for all combination of harmonic orders and degrees. Rotation matrices are not defined for Zernike polynomials [20], hence they can not be used in the global hemispherical domain.
Elhajian et al. [40, 41]	3D object (shape and texture)	Complex illumination (Investigated Near lighting)	Lambertian	No orthogonalization performed	Incident light is restricted to the upper hemisphere facing the camera while Lambertian reflectance is assumed. Can't be extended to non-Lambertian surfaces
Proposed Appearance Model	3D object (shape and texture)	Complex illumination	non-Lambertian	Analytical <i>Bilinear</i> PCA	Hemispherical reflectance basis is proposed which maintain physically plausible reflectance properties. Shape and albedo recovery is also proposed to construct appearance models given a single image under unknown lighting

1.6 Contributions

To address the aforementioned research questions, an analytic formulation of appearance subspace construction is sought to capture the full behavior of complex illumination and real-world reflectance, while in the process accounting for the hemispherical nature of surface reflectance functions to model physically plausible surface materials.

This work started with examining a very restrictive case of convex Lambertian surfaces when image irradiance is captured from a single viewpoint with distant [40] and near [41] illumination assumptions. In such a case, the visible part of the object’s surface constructs the upper hemisphere of the surface normals where spherical harmonics (SH) is no longer orthonormal. Assuming no inter-reflections and considering light source(s) distributed uniformly over the upper hemisphere, the use of HemiSpherical Harmonics (HSH) was proposed to model such a case while it has been proved analytically, and validated experimentally, that the Lambertian reflectance kernel has a more compact harmonic expansion in the hemispherical domain when compared to its spherical counterpart.

Moving away from a Lambertian assumption, it was realized the need for formulating the image irradiance equation with regard to the local-coordinate frame where illumination is perceived with regard to surface points. The contributions of this dissertation can be outlined as follows:

- In **Chapter 4** (based on [45]), a novel surface reflectance basis is proposed based on hemispherical harmonics which is defined on the cartesian product of the incoming and outgoing local hemispheres (*i.e. w.r.t. surface points*). This basis obeys physical properties of surface reflectance involving reciprocity and energy conservation. The proposed basis functions are validated using analytic reflectance models as well as scattered reflectance measurements. The form of the proposed basis is derived in case of directional hemispherical reflectance. *This addresses RQ1 and RQ2.*
- In **Chapter 5** (based on [46]), the image formation process of isotropic surfaces under arbitrary distant illumination is further formulated in the frequency space. The term *irradiance harmonics* is defined which enables decoupling illumination and reflectance from the underlying geometry and pose. A closed form of the

energy content being maintained by different reflectance modes of the proposed irradiance harmonics is defined. The approximation accuracy of the irradiance harmonics is evaluated based on the proposed reflectance basis compared to similar basis in literature where BRDF measurements is used which are directly measured from real surfaces. *This addresses RQ3.*

- In **Chapter 6**, an analytical formulation for low-dimensional subspace construction is proposed in which shading cues lie while preserving the natural structure of an image sample. Thanks to the frequency-space representation of the image irradiance equation, the process of finding such subspace can be cast as establishing a relation between its principal components and that of a deterministic set of basis functions. This resolves the issue of dimensionality since the source of randomness in the imaging process becomes the irradiance harmonics coefficients rather than the whole image realization. Representing images in their natural dimension, *i.e.* matrices, further lessen the number of parameters to be estimated to define a bilinear projection which maps the image sample to a lower-dimensional bilinear subspace. Since irradiance harmonics enables decoupling illumination and reflectance from the underlying geometry and pose; this enables the incorporation of prior information about natural illumination and real world surface materials. *This addresses RQ4 and RQ5.*

- In **Chapter 7** (based on [47]), the problem of reconstructing the image irradiance signal from incomplete irradiance signal is considered where sources of incompleteness might include, for example, occlusions due to wearing apparel and makeup for facial images, or even incomppliance to the imaging model assumptions such as non-convexity which introduce cast shadows. The errors introduced due to irradiance incompleteness are cast as: (1) statistical outliers which are determined and rejected using robust statistics and (2) local spatial erroneous continuous regions where Markov Gibbs random field with the homogenous isotropic Potts model is adopted to model the incompleteness's spatial interaction. *This addresses RQ6.*

- In **Chapter 8** (based on [48]), the main objective is aiming at recovering 3D facial shape from a single image of unknown general illumination while relaxing the non-realistic assumption of Lambertian reflectance. Prior shape, albedo and reflectance models from real data, which are metric in nature, are incorporated into

the shape recovery framework. Adopting a frequency-space based representation of the image irradiance equation, an appearance model is proposed, termed as Harmonic Projection Images, which accounts explicitly for different human skin types as well as complex illumination conditions. Assuming skin reflectance obeys Torrance-Sparrow model, it is proved analytically that it can be represented by at most 5th order harmonic basis whose closed form is provided. The recovery framework is a non-iterative approach which incorporates regression-like algorithm in the minimization process. *This addresses RQ7.*

- In **Chapter 9** (preliminary work appears in [49]), a model-based shape-from-shading approach is proposed which allows for the construction of plausible human jaw models in vivo, without ionizing radiation, using fewer sample points in order to reduce the cost and intrusiveness of acquiring models of patients teeth/jaws over time. Human teeth reflectance is assumed to obey Wolff-Oren-Nayar model where it is experimentally proved that teeth surface obeys the microfacet theory. While most shape-from-shading (SFS) approaches assume known parameters of surface reflectance and point light source with known direction, this work relaxes such assumptions using the harmonic expansion of the image irradiance equation where it is feasible to incorporate prior information about natural illumination and teeth reflectance characteristics. The results demonstrate the effect of adding statistical prior as well as appearance (illumination and reflectance) modeling on the accuracy of the recovered shape. The applicability of the proposed approach to a dental application encompassing tooth restoration is further investigated. *This addresses RQ7.*

1.7 Dissertation Organization

This document is divided into three main parts. The first part (Chapters 2 and 3) covers the fundamentals of the photometric aspect of image formation, introducing the necessary terminology to understand the domain of the research presented and discussing how illumination can be represented using spherical harmonics. The second part (Chapters 4 through 7) presents the proposed solution for efficient image irradiance modeling while relaxing the simplifying Lambertian assumption which is widely assumed. The third part (Chapters 8 and 9) shows how the proposed imaging model would be beneficial to

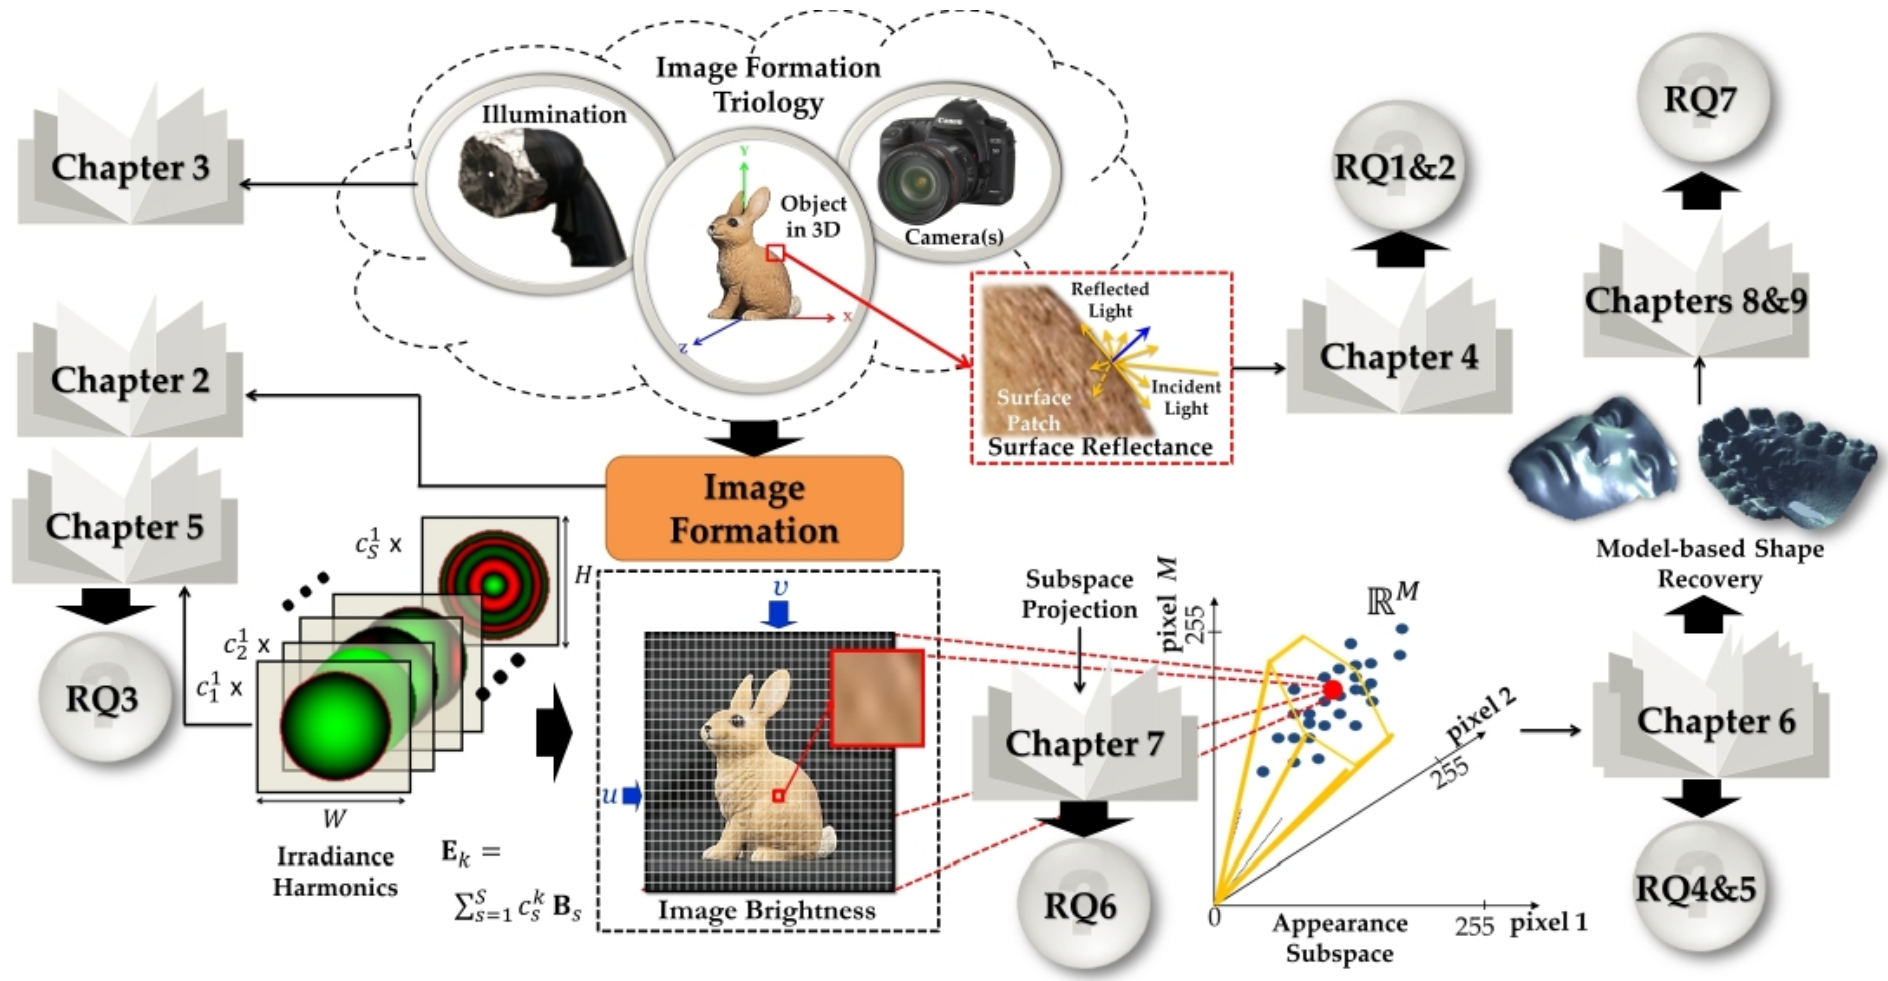


Figure 1.7: Dissertation Layout Illustration

vision tasks such as shape recovery. The conclusion and future directions are outlined in Chapter 10. Relevant proofs for theorems and formulae formulated in this dissertation along with some background materials are included in the appendices. Dissertation layout is illustrated in Figure 1.7.

PART I

THEORETICAL BACKGROUND

PHOTOMETRIC ASPECT OF IMAGE FORMATION

In order to use images to infer information about the 3D world, one needs to study how pixel brightness in the image is related to the physical world. To accomplish such a task, two questions needed to be addressed; (1) where some point in 3D will appear in the 2D image, and (2) how bright this image point will be. The former question is related to the camera/viewer properties, *i.e.* geometric image formation, while the latter one is governed by surface properties and illumination conditions, *i.e.* photometric image formation. This chapter presents the fundamentals of photometric image formation necessary for the analysis performed in the presented work, where it is common in computer vision to consider a viewer-centered coordinate system with the viewer located at a distance relatively large compared to the object size [24]. Assuming, without loss of generality, that the viewing direction coincides with the z -axis of the global reference frame, where the orthographic projection can be used to approximate the geometric imaging process.

The interplay between light source(s) illuminating an object and the object reflectance properties construct the fundamentals of the photometric aspect of image formation [3]. The image brightness primarily depends on the physics related to the behavior of light in 3D space, defined by how energy is transferred from a light source to a surface patch.

Consider an omnidirectional light source emitting light rays in all directions, where one light ray is directed towards a surface point \mathbf{x} . When a light ray hits a surface point, a portion of its energy is absorbed by the object, while the remaining is reflected from the surface. Such portion is a function of surface reflectance properties. The brightness

of the projected point in the image plane is affected by the reflected light rays in the direction of the camera.

Usually, the physics of light is explained using different models [50], such as: (1) *Geometric or ray optics* which models light as independent rays traveling in an optical media according to geometric rules, describing the processes of reflection and refraction. (2) *Wave optics* which models light as electromagnetic waves, describing interference and diffraction, in addition to reflection and refraction. It is necessary to understand the interaction of light with objects having size comparable to the light wavelength. (3) *Electromagnetic Optics* which includes wave optics while adding polarization and dispersion. (4) *Photon or quantum optics* which provides the foundation of the interaction of light with atoms and molecules of the matter, where the light is assumed to consist of particle, or photons. In computer vision, ray (or geometric) optics has been exclusively used to model the nature of light and how it scatters in the environment, and this work is no exception, ignoring effects such as diffraction, interference and polarization.

In the following sections, the vocabulary used to describe the photometric aspect of image formation is reviewed. Starting off with the behavior of light rays in space, followed by the process of light reflection at surfaces, ending by formulating the image irradiance equation. Due to the ill-posedness of the shading-based inference problems, this chapter is ended by enumerating the assumptions of the image formation process employed by the presented work, and how such assumptions affect the image irradiance equation.

2.1 Light in Space

Image brightness primarily depends on physics related to the behavior of light in 3D space, which is defined by how energy is transferred from a light source to a surface patch. In order to study such a behavior, it is required to establish the terminology of light in space.

2.1.1 Hemisphere of Directions

A point \mathbf{x} on a 3D surface patch sees its surrounding world through a hemisphere of directions Ω' centered at that point, where light rays can arrive or leave. The local frame

(coordinate system)¹ of this point is defined by the surface normal $\vec{n}(\mathbf{x})$ along with two orthogonal tangential vectors $\vec{t}(\mathbf{x})$ (corresponding to y' -axis) and $\vec{b}(\mathbf{x}) = \vec{n}(\mathbf{x}) \times \vec{t}(\mathbf{x})$ (which is the surface binomial, corresponding to the x' -axis) at this point, where the surface normal is used to define the orientation of the hemisphere (*i.e.* the z' -axis). A light ray $\vec{\omega}'$ arriving to or leaving from the surface point is defined by its inclination/slant angle θ' measured from the surface normal and its azimuth/tilt angle ϕ' of its orthogonal projection on the surface tangent plane at this point measured from a fixed reference direction on that plane. See Figure 2.1 for illustration.




Figure 2.1: A point \mathbf{x} on a surface patch sees its surrounding world through a unit hemisphere Ω' centered at that point. The surface normal $\vec{n}(\mathbf{x})$ at the surface point \mathbf{x} defines the orientation of the hemisphere. A light ray $\vec{\omega}'$ incident to the point \mathbf{x} is defined by its direction (θ', ϕ') in spherical/angular coordinates.

2.1.2 Foreshortening

Foreshortening refers to the visual effect that a surface patch is smaller than it actually is, because it is angled toward the viewer. This means the surface does not appear directly in front of the viewer, *i.e.* tilted, hence it appears smaller than it is. Consider a surface patch tilted with respect to the direction in which the illumination is traveling, such a patch looks smaller to the light source due to foreshortening. Thus the source sees the foreshortened/projected area instead of the actual area of the surface patch. This concept is important since two different surface patches having the same foreshortened area will receive the same amount of radiation from the light source seeing them. Similarly, two different light sources will have the same effect on a surface patch if they

¹Primed coordinates are used to denote local reference frame with respect to a surface point, in contrast to unprimed coordinates which denote the global reference frame. These two coordinate systems are related by a rotation.

look exactly the same to the surface patch. See Figure 2.2 for illustration.




Figure 2.2: A light source sees the foreshortened area of a surface patch instead of its actual area. The foreshortened area depends on the actual area and the angle between the light ray and the surface normal.

2.1.3 Solid Angle

The differential solid angle, $d\vec{\omega}'$, represents the angular size of a light ray as well as its direction, it can be viewed as an infinitesimal area on the unit sphere, measured in steradian (abbreviated as *sr*). It can be expressed in the spherical coordinates using the surface point local frame with z' -axis aligned to the surface normal \vec{n} and the two orthogonal vectors in the tangential plane, namely \vec{t} and \vec{b} , where the size of a differential solid angle in spherical coordinates is given by¹:

$$d\vec{\omega}' = \sin \theta' d\theta' d\phi' \quad (2.1)$$

i.e. the infinitesimal area on the unit sphere is the product of the length of the longitudinal arc $d\theta'$ and the length of the latitudinal arc $\sin \theta' d\phi'$

Given the spherical coordinates, the direction $\vec{\omega}'$ of the solid angle can be computed as,

$$\vec{\omega}' = \sin \theta' \cos \phi' \vec{t} + \sin \theta' \sin \phi' \vec{b} + \cos \theta' \vec{n} \quad (2.2)$$

An arbitrary small planar surface patch of area dA at a distance r from the origin of an illuminating unit sphere, where the light source is centered at the origin of such sphere,

¹Using spherical polar coordinates causes infinitesimal patches on the surface of the sphere to be bigger as the equator is approached and vanishes towards the poles, such an effect is encoded by $\sin \theta'$.

subtends a solid angle given by Equation 2.3 where $\cos \theta'$ accounts for the foreshortening effect¹.

$$d\vec{\omega}' = \frac{dA \cos \theta'}{r^2} \quad (2.3)$$

The light source here is considered as the observer/viewer of the surface patch; hence the solid angle measures how the light source sees such a patch. See Figure 2.3 for illustration. It is important to note that two surface patches subtends the same solid angle receives the same amount of radiation from the light source. The roles of the light source and the surface patch can be interchanged to indicate how a surface patch sees a light source. See Figure 2.4.




Figure 2.3: The solid angle is a two dimensional angle in the three dimensional space which an object (*e.g.* a small planar surface patch) subtends at a point. It measures how large an object appears to a viewer centered at that point. The solid angle subtended by a small planar surface patch of area dA at a light source is a function of the distance between this light source and the surface patch, taking into consideration the foreshortening effect since the patch is tilted with respect to surface area subtended by the solid angle.

2.1.4 Radiance

In radiometry, the field of measuring light, the **radiant flux/power** Φ is defined as the power of electromagnetic radiation emitted/radiated from a source. It is measured in *Watt (W)* having the dimension of energy per unit time or *joules/second*. Hence the **radiant intensity** I_o is defined by the radiant power per unit solid angle, *i.e.* power coming out from a source through a unit solid angle, it is given by;

$$I_o = \frac{d\Phi}{d\vec{\omega}'} \quad W.sr^{-1} \quad (2.4)$$

¹To understand why an angle is related to area/distance, consider the 2D case with a circle of radius r . An infinitesimal arc length dl is related to its subtended angle by $d\theta = dl/r$.




Figure 2.4: An infinitesimal surface patch centered at surface point \mathbf{x} with area dA sees an incident light ray $d\vec{\omega}'$ through the solid angle subtended by the light source on the unit hemisphere centered at \mathbf{x} .

The distribution of light in space is defined by *radiance*, $L(\mathbf{x}, \vec{\omega}')$, measured in $Wm^{-2}sr^{-1}$. According to [8], radiance is the amount of energy traveling at some point \mathbf{x} in a specific direction $\vec{\omega}'^1$, *i.e.* $d^3E_{\vec{\omega}'}$, per unit time, per unit area perpendicular to the direction of travel (*i.e.* foreshortened area), per unit solid angle.

$$L(\mathbf{x}, \vec{\omega}') = \frac{d^3E_{\vec{\omega}'}}{(\cos \theta' dA) d\vec{\omega}' dt} \quad W.m^{-2}.sr^{-1} \quad (2.5)$$

where $d^3E_{\vec{\omega}'}$ is the energy transmitted by a patch/source into an infinitesimal region of solid angle $d\vec{\omega}'$ around the direction $\vec{\omega}' = (\theta', \phi')$ in time dt through a foreshortened area $\cos \theta' dA$.

This definition implies that the amount of energy incident on some patch depends on (1) how large the source is viewed from the patch and (2) how large the patch is viewed from the source. Assuming the light doesn't interact with the traveling medium, *i.e.* vacuum, radiance is characterized by being constant along a straight line, *i.e.* the radiance leaving a point \mathbf{x}_1 in the direction of another point \mathbf{x}_2 is the same as the radiance arriving at \mathbf{x}_2 from \mathbf{x}_1 [8]. This assumption is employed by the majority of computer vision problems.

¹A direction in 3D can be specified in one of two ways; using spherical coordinates $\vec{\omega} = (\theta, \phi)$ or the direction from one point to another, *i.e.* $\vec{\omega} = \mathbf{x}_1 \rightarrow \mathbf{x}_2$.

2.2 Light At Surfaces

Modeling the light scattering locally at surface points, *i.e.* what happens when a light ray strikes a given surface, gives the fundamentals of what so called *local illumination*, which is formulated by the *reflectance equation* under the assumption of no surface emittance. On the other hand, *global illumination* takes into consideration the light being reflected from the whole scene and how this affects each surface point, this is modeled by the *rendering equation*, from which the reflectance function is a special case. The advantage of local illumination models is that they are easy to manipulate. Moreover, they support quite simple theories of how shape information can be extracted from shading. On the contrary, global models are more accurate. However, it is extremely difficult to infer shape information from shading values.

It is common in vision to assume that all effects are local with neither emission nor fluorescence. A reasonable model for such kind of surfaces assumes the following [8]; (1) the radiance leaving a point on a surface is due to the radiance arriving at this point, (2) light leaving a surface at a given wavelength is due to light arriving at that wavelength, (3) light sources are treated separately, and (4) surfaces do not generate light internally.

Since the illumination is assumed to affect the surface locally. Consider a point \mathbf{x} on a surface, seeing the surrounding world through a local unit hemisphere Ω' oriented by the surface normal $\vec{n}(\mathbf{x})$ at this point. The incoming and outgoing/reflected illumination (light rays) at this point are defined in the point's local frame as $\vec{\omega}'_i = (\theta'_i, \phi'_i)$ and $\vec{\omega}'_o = (\theta'_o, \phi'_o)$ respectively, where $\theta'_i, \theta'_o \in [0, \pi/2]$ and $\phi'_i, \phi'_o \in [0, 2\pi]$, see Figure 2.5. In order to study/investigate the illumination effect on the surface at this point, the relationship between the incoming and outgoing light at such a point is described as a function of the direction in which light arrives $\vec{\omega}'_i$ and the direction in which light leaves $\vec{\omega}'_o$.

2.2.1 Surface Irradiance

The amount of light falling on a surface is called *irradiance*, *i.e.* incident radiance, which is the unit used to represent the incoming power. It is defined as the incident power (energy per unit time) per unit area not foreshortened. It is measured in $W.m^{-2}$. Hence a surface at point \mathbf{x} illuminated by radiance $L_i(\mathbf{x}, \vec{\omega}'_i)$ ¹ coming in from a directional

¹The subscript i is used here to denote incident radiance to the surface, in contrast to radiance being emitted from a light source.




Figure 2.5: A point \mathbf{x} on a surface patch sees its surrounding world through a unit hemisphere Ω' centered at the point. The surface normal \vec{n} at the point defines the orientation of the hemisphere. Using spherical coordinates, a light ray incident to the point \mathbf{x} is defined by its unit direction $\vec{\omega}'_i = (\theta'_i, \phi'_i)$, subtending an incident solid angle $d\vec{\omega}'_i$. Similarly, an outgoing/reflected light ray is defined by its unit direction $\vec{\omega}'_o = (\theta'_o, \phi'_o)$, subtending an outgoing solid angle $d\vec{\omega}'_o$.

region of solid angle $d\vec{\omega}'_i$ at direction $\vec{\omega}'_i = (\theta'_i, \phi'_i)$ receives irradiance,

$$dE_i(\mathbf{x}, \vec{\omega}'_i) = L_i(\mathbf{x}, \vec{\omega}'_i) \cos \theta'_i d\vec{\omega}'_i \quad (2.6)$$

2.2.2 Surface Scattering

When a light ray strikes a surface, it usually enters the surface and then scatters internally before leaving the surface at a different spatial position, *e.g.* translucent materials such as marble. This scattering process is described by the *Bidirectional Scattering Surface Reflectance Distribution Function* or BSSRDF [9]. The BSSRDF, S , relates the differential reflected radiance, *i.e.* surface radiance, dL_r , at surface point \mathbf{x}_o in direction $\vec{\omega}'_o$, to the differential incident flux, $d\Phi_i$ at \mathbf{x}_i from direction $\vec{\omega}'_i$,

$$S(\mathbf{x}_i, \vec{\omega}'_i, \mathbf{x}_o, \vec{\omega}'_o) = \frac{dL_r(\mathbf{x}_o, \vec{\omega}'_o)}{d\Phi_i(\mathbf{x}_i, \vec{\omega}'_i)} \quad (2.7)$$

The BSSRDF is the most general mathematical abstraction of surface interaction with incident light field, it is a function of the incoming position and direction as well as the outgoing position and direction. However, an eight-dimensional function is costly to

evaluate. Under the assumption that light hitting a surface location is reflected at that same location, Nicodemus *et al.* [9] defined the *Bidirectional Reflectance Distribution Function*, BRDF, to describe light reflection at a surface as an approximation of the BSSRDF, thus reducing the domain to a six-dimensional one.

The BRDF, f_r , defines the relationship between the reflected radiance and irradiance, it is the ratio of the surface radiance in the outgoing direction $\vec{\omega}'_o$ to the irradiance at the same surface point¹.

$$f_r(\mathbf{x}, \vec{\omega}'_i, \vec{\omega}'_o) = \frac{dL_r(\mathbf{x}, \vec{\omega}'_o)}{dE_i(\mathbf{x}, \vec{\omega}'_i)} = \frac{dL_r(\mathbf{x}, \vec{\omega}'_o)}{L_i(\mathbf{x}, \vec{\omega}'_i) \cos \theta'_i d\vec{\omega}'_i} \quad (2.8)$$

This function is specified by two directions (outgoing and incoming), hence the name bidirectional. It ranges from 0, where there is no outgoing/reflected radiance at the exit direction, to infinity, where arbitrary small irradiance received from the incoming direction. BRDF is also a function of position; however, sometimes this positional variance is not included in the description of a BRDF, where position/shift-invariant BRDFs are assumed. When the spatial position is not included as a parameter to the BRDF function, an assumption is implicitly made that the reflectance properties of a surface do not vary with the spatial position. Thus assuming homogeneous surfaces having position/shift-invariant BRDFs, the spatial parameter can be excluded from the BRDF definition, reducing the function to a four-dimensional domain.

Based on its geometrical and physical nature, a BRDF should satisfy the *Helmholtz's* law of reciprocity [51, 52], which guarantees the invariance of the BRDF to the permutation of incident and outgoing directions, *i.e.* ,

$$f_r(\mathbf{x}, \vec{\omega}'_i, \vec{\omega}'_o) = f_r(\mathbf{x}, \vec{\omega}'_o, \vec{\omega}'_i) \quad (2.9)$$

Due to energy conservation, a surface cannot reflect more light that it receives from the incident upper hemisphere Ω'_i , thus a BRDF must satisfy the following,

$$\int_{\Omega'_i} f_r(\mathbf{x}, \vec{\omega}'_i, \vec{\omega}'_o) \cos \theta'_i d\vec{\omega}'_i < 1, \quad \forall \vec{\omega}'_o \quad (2.10)$$

The BRDF of the surfaces, which are isotropic with respect to light scattering, depends only on the polar angles of the incident and outgoing directions, whereas the

¹ r in f_r stands for reflection.

azimuthal dependence is restricted to the absolute difference between the azimuths of the incident and outgoing light rays [21, 53],

$$f_r(\mathbf{x}, \vec{\omega}'_i, \vec{\omega}'_o) = f_r(\mathbf{x}, \theta'_i, \phi'_i, \theta'_o, \phi'_o) = f_r(\mathbf{x}, \theta'_i, \theta'_o, |\phi'_i - \phi'_o|) = f_r(\mathbf{x}, \theta'_i, \theta'_o, \Delta\phi) \quad (2.11)$$

Although isotropy is not a universal property for all surfaces, many real surfaces exhibit such a property, while the reflectance of anisotropic surfaces can be approximated as being isotropic [10].

2.2.3 Surface Radiance

The surface radiance results from the contribution of all light rays being reflected in the direction towards the camera. The incident radiance at a surface point \mathbf{x} in the direction $\vec{\omega}'_i$, $L_i(\mathbf{x}, \vec{\omega}'_i)$, defines the surface irradiance $dE_i(\mathbf{x}, \vec{\omega}'_i)$ through a differential solid angle $d\vec{\omega}'_i$ in the direction of incidence, it depends on the incident polar angle due to the foreshortening effect, see Equation 2.6. According to the BRDF definition in Equation 2.8, this gives rise to differential surface radiance, $dL_r(\mathbf{x}, \vec{\omega}'_o)$, defined according to the differential incident solid angle. The total surface radiance can thus be computed by integrating this differential surface radiance over the hemisphere of possible incident directions Ω'_i ,

$$L_r(\mathbf{x}, \vec{\omega}'_o) = \int_{\Omega'_i} L_i(\mathbf{x}, \vec{\omega}'_i) f_r(\mathbf{x}, \vec{\omega}'_i, \vec{\omega}'_o) \cos \theta'_i d\vec{\omega}'_i \quad (2.12)$$

2.3 Image Irradiance

The brightness of an image pixel is determined by the *image irradiance* received by the corresponding element of the image sensor. Using Horn model of a single lens camera [24], the image irradiance is (approximately) directly proportional to the surface radiance with the factor of proportionality depending on the angle φ between the ray from the surface to the image and the optical axis, *i.e.* ,

$$E(\mathbf{x}, \vec{\omega}'_o) \propto \cos^4 \varphi L_r(\mathbf{x}, \vec{\omega}'_o) \quad (2.13)$$

When the distance between the camera and the surface is large compared to the size of the surface, such factor is negligible. Thus the image irradiance equation, under the

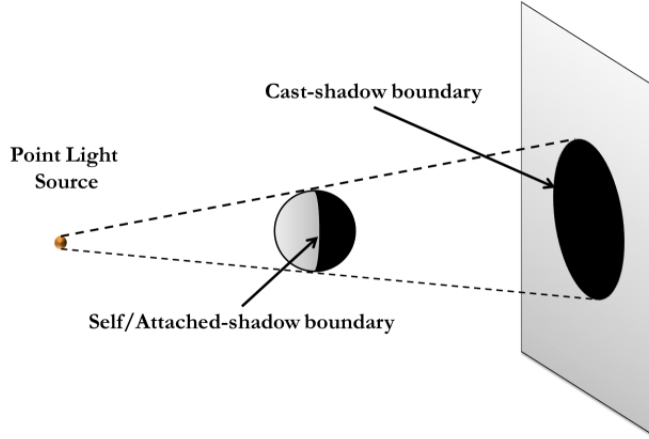


Figure 2.6: Self-shadow boundaries occur when the surface is turning away from the source while cast-shadow boundaries occur when a surface patch does not see the source due to occluding object. This figure is re-produced from [8] for illustration purposes.

assumption of no surface emission, can be defined as,

$$E(\mathbf{x}, \vec{\omega}'_o) = \int_{\Omega'_i} L_i(\mathbf{x}, \vec{\omega}'_i) f_r(\mathbf{x}, \vec{\omega}'_i, \vec{\omega}'_o) \cos \theta'_i d\vec{\omega}'_i \quad (2.14)$$

In a local illumination model, shadows occur at surface points which cannot see one or more light sources. Any surface patch is said to be in shadow, if the line of sight connecting such a patch and the light source passes through one or more objects. Shadows generated from point sources are characterized by crisp boundaries, where there are two kinds of shadow boundary; (1) *attached-shadow*¹ boundaries which occur when the surface is turning away from the light source and the ray connecting the patch to the source becomes tangent to the surface, and (2) *cast-shadow* boundaries where the source suddenly disappears behind an occluding object causing shadowing effect of that object cast on the surface patch. See Figure 2.6 for illustration. In case of more than one source illuminating the scene, shadows will be less dark with very dark areas occur if no source is visible. Ignoring cast-shadows, the effect of attached shadows on the image irradiance is a negative intensity value, where the angle between the surface normal and the incident light direction exceeds $\pi/2$. While attached shadows can be accounted for by controlling the integration domain of Equation 2.14 and the foreshortening term, a visibility term should be added to the integrand in order to handle cast shadows which indicates whether a surface point sees the source. This work primarily focuses on attached shadows whereas cast shadows can be considered as a potential future direction

¹a.k.a. self-shadow since it occurs due to the surface geometry.




Figure 2.7: The image irradiance equation can be formulated such that the integration domain is over the global sphere. This is commonly used in case of Lambertian surfaces, where such formulation is switched to local spherical domain with respect to surface points in case of non-Lambertian surfaces. To ensure non-negative intensity values, the image irradiance equation Equation 2.14 is usually modified to include $\max(0, \cos \theta'_i)$. While under the assumption of non-emitting objects, a surface point sees the surrounding world through the local hemisphere oriented by the surface normal at that point, thus integration can be defined over the incident local hemisphere. This formulation accounts implicitly for attached shadows without explicitly modifying the foreshortening (cosine) term.

of research.

The image irradiance equation in Equation 2.14 can be formulated using three different integration domains according to the assumptions imposed on the image formation process. In case of convex-Lambertian surfaces, it is common to express the image irradiance in the global coordinate frame [6, 11, 25], *i.e.* global sphere, see Figure 2.7. For non-Lambertian surfaces, it is useful to change this formulation to the local coordinate frame oriented by the surface normal at each surface point [11]. To account for attached shadows, the image irradiance equation Equation 2.14 is usually modified to include $\max(0, \cos \theta'_i)$ to ensure non-negative intensity values, while the integration domain takes into account the whole global/local sphere [6, 7, 11]. On the other hand, under the assumption of non-emitting surfaces, a surface point only sees the surrounding environment through the local upper hemisphere oriented by the surface normal at this point, see Figure 2.7. Thus the integration in Equation 2.14 can be defined over the incident local hemisphere where $\theta'_i \in [0, \pi/2]$. This formulation accounts implicitly for attached shadows without explicitly modifying the foreshortening (cosine) term.

2.4 Assumptions

Due to the illposedness of shading-based inference problems, assumptions are commonly used to present a simplified version of the image formation process. This work assumes homogeneous, convex surfaces under distant illumination conditions. While these assumptions are common in computer vision, non-homogeneous surfaces with spatially varying BRDFs are often approximated in vision applications by using a spatially varying texture which modulates the surface radiance [11]. These assumptions are considered a reasonable approximation to various computer vision tasks while being flexible enough for analytical formulation. In the following subsections, the effect of such assumptions on the image irradiance equation in Equation 2.14 is elaborated.

2.4.1 Homogeneous Convex Surfaces

Assuming homogeneous surfaces allows modeling the surface at hand with the same BRDF everywhere. Thus the dependency of the BRDF on the spatial position can be dropped, leading to a four-dimensional function, $f_r(\vec{\omega}'_i, \vec{\omega}'_o)$.

The convex assumption implies that there is no cast-shadows or inter-reflection. Thus the incident illumination at a surface point is only due to a distant light source. This assumption allows parameterizing the surface by orientation. For isotropic surfaces, the surface orientation can be fully defined by its normal direction at each surface point. While for anisotropic surfaces, the direction of anisotropy is specified by the orientation of the local tangent frame.

2.4.2 Distant Illumination

Local surface patches have sizes relatively smaller than the distance to the light source, thus the incident light field can be assumed to be constant on these local patches. Hence distant illumination is a common assumption in many computer vision algorithms.

The image irradiance in Equation 2.14 is defined in a local coordinate frame oriented by the surface normal. The surface orientation with angular coordinates (α, β, γ) defines the rotational relation between the global reference frame and the local frame at the surface point \mathbf{x} where $\alpha = \alpha(\mathbf{x})$, $\beta = \beta(\mathbf{x})$ and $\gamma = \gamma(\mathbf{x})$. Noting that γ is only important for anisotropic surfaces, which controls the rotation of the local tangential plane about the surface normal. Thus the image irradiance can be re-parameterized as

$$E(\mathbf{x}, \vec{\omega}'_o) \rightarrow E(\vec{n}(\mathbf{x}), \vec{\omega}'_o) = E(\alpha, \beta, \gamma, \theta'_o, \phi'_o).$$

Under the assumption of distant light source, all surface points receive the same amount of light, hence the positional variance of the lighting function can be omitted. Thus the lighting function can be re-parameterized as $L_i(\mathbf{x}, \vec{\omega}'_i) \rightarrow L_i(\vec{n}(\mathbf{x}), \vec{\omega}'_i) = L_i(\vec{\omega}_i)$ ¹. Now, the lighting function represents the source radiance from incident directions defined in the global reference frame, however the surface irradiance from the surrounding environment at any surface point is defined locally on the upper hemisphere, yielding a mixed-frames image irradiance equation,

$$E(\vec{n}(\mathbf{x}), \vec{\omega}'_o) = \int_{\Omega'_i} L_i(\vec{\omega}_i) f_r(\vec{\omega}'_i, \vec{\omega}'_o) \cos \theta'_i d\vec{\omega}'_i \quad (2.15)$$

Using Euler angle representation of 3D rotations with ZYZ-convention, corresponding to three consecutive counterclockwise rotations about the z -, y - and z - axis respectively [54], the image irradiance equation in local coordinates can be written as,

$$E(\vec{n}(\mathbf{x}), \vec{\omega}'_o) = \int_{\Omega'_i} L_i(R_{\alpha, \beta, \gamma} \vec{\omega}'_i) f_r(\vec{\omega}'_i, \vec{\omega}'_o) \cos \theta'_i d\vec{\omega}'_i \quad (2.16)$$

where $R_{\alpha, \beta, \gamma} = R_Z(\beta)R_Y(\alpha)R_Z(\gamma)$. Note that (α, β) are the spherical coordinates of the surface normal $\vec{n}(\mathbf{x})$ in the global reference frame. For isotropic surfaces, γ has no physical meaning; it can be set arbitrarily to zero. Thus the rotation operator is redefined as $R_{\alpha, \beta, 0} = R_{\alpha, \beta} = R_Z(\beta)R_Y(\alpha)$.

2.4.3 Orthographic Projection

Considering a viewer-centered coordinate system with the viewer/camera located at a distance relatively large compared to the object size. Assuming, without loss of generality, that the viewing direction (*i.e.* the optical axis) coincides with the z -axis of the global reference frame, where orthographic projection can be used to approximate the imaging process. This viewing geometry implies that the reflected beam aligns with the optical axis, meaning that $\theta'_o = \alpha$ [32]. Moreover, setting $\gamma = 0$ for isotropic surface leads to $\phi'_o = \pi$. Thus the image irradiance equation for isotropic surfaces can be completely determined by the surface normal as follows,

¹The angular coordinates of the surface normal is used to rotate the local incident direction to be represented in the global coordinate frame.

$$E(\vec{n}(\mathbf{x})) = E(\alpha, \beta) = \int_{\Omega'_i} L_i(R_{\alpha,\beta,0}(\theta'_i, \phi'_i)) f_r(\theta'_i, \phi'_i, \alpha, \pi) \cos \theta'_i d\vec{\omega}'_i \quad (2.17)$$

where $R_{\alpha,\beta,0}(\theta'_i, \phi'_i) \triangleq R_{\alpha,\beta,0}\vec{\omega}'_i$.

2.5 Image Irradiance Equation

The image irradiance of a convex surface with spatially varying albedo is then formed by allowing each surface point to inherit its intensity from that point on a unit sphere having the same normal, such intensity is scaled by the point's albedo. Formally, let $\rho(\mathbf{x})$ be the albedo/texture which modulates the surface radiance at point \mathbf{x} . Thus the image irradiance can be written as¹,

$$E(\vec{n}(\mathbf{x})) = \rho(\mathbf{x}) \int_{\Omega'_i} L_i(R_{\alpha,\beta,0}(\theta'_i, \phi'_i)) f_r(\theta'_i, \phi'_i, \alpha, \pi) \cos \theta'_i d\vec{\omega}'_i \quad (2.18)$$

In this dissertation work, an accurate and compact representation of this function in the frequency domain will be investigated for the purposes of illumination modeling and image analysis.

2.6 Summary

In this chapter, the fundamentals of the photometric aspect of image formation are reviewed, introducing the necessary terminology to understand the domain of the research presented. Starting with the light being radiated from a light source, then how light rays being received by a surface patch and then reflected. The chapter ends with the assumptions employed in this dissertation research and how they affected the image irradiance equation. In the following chapters, the representation of illumination and reflectance in the frequency domain will be investigated while deriving the image irradiance basis which spans the image subspace.

¹Sometimes $E(\vec{n}(\mathbf{x}))$, $E(\mathbf{x})$ and $E(\alpha, \beta)$ are used interchangeably to refer to the image irradiance from surface point \mathbf{x} .

ILLUMINATION REPRESENTATION

Light sources are used to emit light rays which are eventually reflected back to the camera to form the respective image of visible objects in the scene. There are variety of light source models which are primarily characterized by their geometry. Assuming simple source geometries is beneficial for two reasons [8]; (1) this will facilitate the generative process of image formation, i.e. image/view synthesis, while producing fairly effective images and (2) sources with simple geometries can still yield complex visual effects. In this chapter, a brief description of some source models is given which are used in vision and graphics. The ultimate goal is how such models can be represented as a part of the irradiance integral given in Chapter 2.

3.1 Light Source Models

Among various source models which primary defined in graphics literature, four main types of geometrical source models are of interest; point, area, distant and infinite-area sources. Different geometrical sources lead to different expressions for surface radiance obtained from the same surface patch. In what follows, the appearance of the light source from the surface patch is considered.

3.1.1 Point Sources

Since a light source is physically small compared to the scene it is illuminating, it is safe to assume a light source as a point with respect to the scene. In order to model the effect of a point source on a surface patch, the source can be represented as a very small sphere S^2 which emits light at each point on the sphere with a constant exitance¹

¹Exitance $E_\chi(\mathbf{x})$ at point \mathbf{x} by definition is the total internal power (energy per time) leaving a light source positioned at a point \mathbf{x} per unit area on the radiating surface. It has units Wm^{-2} . It is essential

i.e. $E_\chi(\mathbf{x}) = E_o \quad \forall \mathbf{x}$ over the sphere. This is also known as isotropic point light source, emitting the same amount of light in all directions, in contrast to anisotropic point sources. See Figure 3.1 for a sample rendering of this type of source.




Figure 3.1: Path tracing [30] is used to render Killeroo model being illuminated by a point light source. Notice the hard cast shadow boundaries (as shown in the zoom-in views) produced for this source model. Killeroo model courtesy of headus/Rezard (<http://www.headus.com/au/>).

Consider a surface patch centered at surface point \mathbf{x} viewing a sphere centered at \mathbf{x}_s of radius ϵ at a distance $r(\mathbf{x}) = \|\mathbf{x}_s - \mathbf{x}\|$ such that $r(\mathbf{x}) \gg \epsilon$. See Figure 3.2 for illustration. The solid angle subtended by the source $d\bar{\omega}'_i$ is approximately proportional to $\pi\epsilon^2/r(\mathbf{x})^2$, recall that solid angle is related to foreshortened area/distance. As the sphere moves away from the surface patch, light rays leaving/reflected from the surface patch and arriving to the sphere move closer to each other. Recall that surface radiance is obtained by integrating the incident light rays multiplied by the foreshortening factor over the incoming hemisphere. However in this situation, it is important to note that the integral is over a small domain on the incoming hemisphere of the surface patch and the radiance does not change much over such a domain (due to constant exitance from the source), hence it is roughly equal to one value times the area of the patch. Further, for a point light source, the incident radiance comes from only a single direction since the source is only visible from its given location. Thus point sources are usually defined by a delta distribution where surface radiance due to point source can be defined as;

to notice the difference between exitance and radiosity. Exitance measures the internally generated power radiated by a light source per unit area on the radiating surface, while radiosity measures the total power leaving/reflected from a surface point per unit area. Sources can have both radiosity and exitance, because energy may be reflected off the source as well as generated within it [8].

$$L_r(\mathbf{x}, \vec{\omega}'_o) = \int_{\Omega'_i} \delta(\vec{\omega}'_i - \vec{\omega}'_s) f_r(\vec{\omega}'_i, \vec{\omega}'_o) \left(\frac{\epsilon^2}{r^2(\mathbf{x})} \right) E_o \cos \theta'_i d\vec{\omega}'_i \quad (3.1)$$

where $\vec{\omega}'_s$ is unit vector defining the point light direction as perceived locally from the surface patch centered at \mathbf{x} .

Let $\vec{n}(\mathbf{x})$ be the unit normal to the surface at point \mathbf{x} and $\vec{s}(\mathbf{x})$ be a vector from the surface point \mathbf{x} and the source with length $|\vec{s}(\mathbf{x})| = \epsilon^2 E_o$. Thus $\vec{n}(\mathbf{x}) \cdot \vec{s}(\mathbf{x}) = \cos \theta'_i$. This will yield the standard *nearby point source model*;

$$L_r(\mathbf{x}, \vec{\omega}'_o) = \int_{\Omega'_i} \delta(\vec{\omega}'_i - \vec{\omega}'_s) f_r(\vec{\omega}'_i, \vec{\omega}'_o) \frac{\cos \theta'_i}{r^2(\mathbf{x})} d\vec{\omega}'_i = f_r(\vec{\omega}'_s, \vec{\omega}'_o) \frac{\cos \theta'_s}{\|\mathbf{x}_s - \mathbf{x}\|^2} \quad (3.2)$$

where $\vec{\omega}'_s = \frac{\vec{s}(\mathbf{x})}{|\vec{s}(\mathbf{x})|} = \text{cart}(1, \theta'_s, \phi'_s)$.

According to this model, the unknown exitance term is folded in the source vector and an explicit relationship between surface radiance and surface shape (the normal term) is established. Although each surface point \mathbf{x} has a different distance to the nearby source, it is common, however incorrect, to omit the dependency on such distance.




Figure 3.2: A surface patch sees a distant sphere with radius ϵ at distance r . The sphere produces a small illuminated patch on the incoming hemisphere of the surface patch.

3.1.2 Distant Sources

Distant, *a.k.a.* directional, source emits illumination from the same direction at every surface point. As a point source becomes progressively farther away, *i.e.* located at

infinity, *e.g.* the sun, or its distance to the object is large compared to the object's size, the terms $\vec{s}(\mathbf{x})$ and $r(\mathbf{x})$ become constant with respect to surface points, *i.e.* $\vec{s}(\mathbf{x}) = \vec{s}_o + \Delta\vec{s}(\mathbf{x})$ with $|\vec{s}_o| \gg |\Delta\vec{s}(\mathbf{x})|$ and $r(\mathbf{x}) = r_o + \Delta r(\mathbf{x})$ with $|r_o| \gg |\Delta r(\mathbf{x})|$. Thus,

$$\frac{\vec{n}(\mathbf{x}) \cdot \vec{s}(\mathbf{x})}{r^2(\mathbf{x})} = \frac{\vec{n}(\mathbf{x}) \cdot (\vec{s}_o + \Delta\vec{s}(\mathbf{x}))}{(r_o + \Delta r(\mathbf{x}))^2} \approx \frac{\vec{n}(\mathbf{x}) \cdot \vec{s}_o}{r_o^2} \quad (3.3)$$

With a single distant source of direction $\vec{\omega}_s$, the lighting function is described as a delta function [11]. A surface point sees the light source through its upper hemisphere from a single direction $\vec{\omega}'_s = R_{\alpha,\beta,\gamma}\vec{\omega}_s$, as such only one term in the radiance integral (Equation 3.2) will survive. Thus the surface radiance due to a *distant light source* becomes (r_o can be discarded since it is constant for all surface points while light source can be assumed to be of unit exitance);

$$L_r(\mathbf{x}, \vec{\omega}'_o) = \int_{\Omega'_i} \delta(\vec{\omega}'_i - \vec{\omega}'_s) f_r(\vec{\omega}'_i, \vec{\omega}'_o) \cos \theta'_i d\vec{\omega}'_i = f_r(\vec{\omega}'_s, \vec{\omega}'_o) \cos \theta'_s \quad (3.4)$$

where $\cos \theta'_s = \vec{n}(\mathbf{x}) \cdot \vec{s}_o$.




Figure 3.3: Path tracing [30] is used to render Killeroo model being illuminated by a distant light source. Along with the apparent hard cast shadow boundaries, one can observe the shape details of Killeroo model from the shadow cast on the ground. This reveals the difference between appearance resulting from distant and point light source (see Figure 3.1 for comparison). Notice the shadow of Killeroo is different according to the surrounding environment. Killeroo model courtesy of headus/Rezard (<http://www.headus.com/au/>).

Distant source (a.k.a. point source at infinity) is a valid model when the distance between objects in the scene is smaller in magnitude than the distance to the source, otherwise the nearby point source model should be assumed. The difficulty with the

nearby model is that surface radiance changes very sharply over space [8].

3.1.3 Area Sources

An overcast sky in natural scenes and fluorescent light box in synthetic environments are good examples of area light sources. Area sources are commonly modeled as surface patches whose emitted radiation is independent of position and direction [8], generating uniform illumination distribution, thus they can be described by their exitance.

Computing surface radiance due to area lights involves computing integrals over the surface of the light source which often can not be computed in closed form [30]. At the expense of computational complexity, soft shadows and more realistic lighting effects, rather than hard shadows, can be achieved by area light sources. See Figure 3.5 for sample renderings of Killeroo model being illuminated by a disk shaped area light sources with different radii.

Consider a surface patch illuminated by an area source of exitance $E_\chi(\mathbf{u})^1$, where \mathbf{u} is a point on the source measured with respect to the source local frame, in contrast to the surface point \mathbf{x} on the patch. Thus surface radiance due to area source can be defined as;

$$L_r(\mathbf{x}, \vec{\omega}'_o) = \int_{\Omega'_i} L_i(\mathbf{x}, \mathbf{u} \rightarrow \mathbf{x}) f_r(\mathbf{u} \rightarrow \mathbf{x}, \vec{\omega}'_o) \cos \theta'_i d\vec{\omega}'_i \quad (3.5)$$

where $\mathbf{u} \rightarrow \mathbf{x}$ defines the direction from the source to the surface patch. Going from the integration over the incoming hemisphere at the surface patch to the outgoing hemisphere at the source gives [8], (note that L_i is now emitted from the source and becomes source exitance L_e)

$$\begin{aligned} L_r(\mathbf{x}, \vec{\omega}'_o) &= \int_{\Omega'_s} L_e(\mathbf{u}, \mathbf{u} \rightarrow \mathbf{x}) f_r(\mathbf{u} \rightarrow \mathbf{x}, \vec{\omega}'_o) \cos \theta'_i d\vec{\omega}'_s \\ &= \int_{source} \frac{E_\chi(\mathbf{u})}{\pi} f_r(\mathbf{u} \rightarrow \mathbf{x}, \vec{\omega}'_o) \cos \theta'_i \left(\frac{dA_u \cos \theta'_s}{r^2} \right) \end{aligned} \quad (3.6)$$

Note that the integration is now over the sources area/region. See Figure 3.4 for illustration.

3.1.4 Infinite Area Sources

Usually referred to as an *environment map*, it describes an infinitely far away area light source casting light into the scene from all directions surrounding the entire scene.

¹Since area sources are approximated as rectangular region which emits light, \mathbf{u} is used to denote any point on such region similar to texture-mapping notations, however meanings are implicitly different.

Typically, an omnidirectional image (referred to as a *light probe*) is used to record the incident illumination from all incident directions at a particular point in space. These images can be captured from real world scenes [55] and several freely available light probe images can be found in <http://www.debevec.org/Probes/>. Further, a larger collection of environment maps are released from sIBL (smart Image Based Lighting) archive¹ including a wide variety of illumination conditions such as indoor, outdoor, summer and winter conditions.

See Figure 3.6 for sample renderings for Killeroo model at different times of the day. One can notice the change of appearance due to illumination variation while keeping the geometry and reflectance (*e.g.* Lambertian) fixed.




Figure 3.4: An area light source illuminating a diffuse surface patch. In order to compute the patch radiance due to such a source, the integral of incoming radiance at the surface is transformed into an integral over the source area dA_u .

Light probes, also sometimes called “Light Maps” as well, are 360×180 -degree spherical (or latlong) high dynamic range (HDR) files. They have been captured by photographing a mirror sphere from several directions and at several exposures. The images were then calibrated and merged to form a high dynamic range image of the full view sphere. These probes are considered nowadays the de-facto-standard for encoding natural illumination in vision and graphics literature.

¹<http://www.hdrlabs.com/sibl/archive.html>




Figure 3.5: Path tracing [30] is used to render Killeroo model being illuminated by a disk-shaped area light source of different radii. Starting with a disk radius which is relatively small compared to the model size; the shadow has soft penumbrae but otherwise the image looks very similar to the one with point source. The effect of using larger disk appears in larger penumbrae and noticeably different appearances when wider range of directions is used for illumination. Killeroo model courtesy of headus/Rezard (<http://www.headus.com/au/>).




Figure 3.6: Path tracing [30] is used to render Killeroo model illuminated by different environment maps. (a) Blue skylight, (b) day skylight and (c) sunset skylight. Maps (courtesy of [30]) are shown in the first row. Using a realistic illumination distribution gives appearance which is more realistically compelling compared to the ones of area and directional light sources. In addition, changing just the environment map gives different renderings revealing the illumination conditions surrounding the scene object. Killeroo model courtesy of headus/Rezard (<http://www.headus.com/au/>).

While several parameterizations are available, see Figure 3.7¹, in this work it is preferred to stick with the angular representation whose representation complies with the subsequent analysis. It is worth noting that sIBL maps are available as latlong files where HDRShop² was used to convert them to the angular representation.




Figure 3.7: Sample light probe image (Galileo’s Tomb, Santa Croce, Florence) with different parameterizations courtesy of Paul Debevec [55].

3.2 Illumination in Frequency Space

Typically, the lighting function $L_i : \mathcal{S}^2 \rightarrow \mathbb{R}$ is a spherical function defined on the surface of the *global* sphere. Thus spherical harmonics (SH) has been an attractive fit for illumination representation in the frequency domain since they form an orthonormal basis for spherical functions.

3.2.1 Spherical Harmonics

3.2.1.1 Definition

An infinite series of associated Legendre polynomials $P_n^m(x)$ can be used to express any piecewise continuous function over the interval $[-1, 1]$. They are defined in terms of the m th derivative of the Legendre polynomials as³,

$$P_n^m(x) = \frac{1}{2^n n!} (1 - x^2)^{m/2} \frac{d^m}{dx^m} P_n(x) \tag{3.7}$$

where $n \geq 0$ and $m \in [0, n]$. n here is used as the band index (a.k.a. order), dividing the polynomial set into bands of functions with $(n + 1)n$ polynomials for n -th band series.

¹Light probes are courtesy of Paul Debevec, <http://www.pauldebevec.com/Probes/>.

²<http://ict.debevec.org/debevec/HDRShop/download/>

³In [56], the associated Legendre polynomials were defined with an additional $(-1)^m$, which seems a complication that can be ignored at this point. It will be included in the definition of the spherical harmonics later.

Whereas m denotes the polynomial degree.

In case of either band-limited functions or band-limited approximation, a finite number of polynomials can be used. Thus any circular-symmetric function with no azimuthal dependence can be expressed in terms of associated Legendre polynomials by mapping the polar angle θ to the interval $[-1, 1]$, this can be accomplished by using $x = \cos \theta$ as the independent variable of the polynomial $P_n^m(x)$. In order to guarantee orthogonality in case of non-circular symmetric functions, associated Legendre polynomials are combined with sinusoidal functions for the azimuthal part dependency [57]. As such, the real SH functions can be written as¹,

$$Y_n^m(\theta, \phi) = \begin{cases} \sqrt{2}N_n^m P_n^m(\cos \theta) \cos(m\phi) & m > 0 \\ N_n^0 P_n^0(\cos \theta) & m = 0 \\ (-1)^m \sqrt{2}N_n^{|m|} P_n^{|m|}(\cos \theta) \sin(-m\phi) & m < 0 \end{cases} \quad (3.8)$$

with N_n^m being defined as the normalization factor (to guarantee orthonormality),

$$N_n^m = \sqrt{\frac{2n+1}{4\pi} \frac{(n-m)!}{(n+m)!}} \quad (3.9)$$

Eqn. Equation 3.8 can be short-handed as,

$$Y_n^m(\theta, \phi) = \Theta_n^m(\theta) \Phi_m(\phi) \quad (3.10)$$

where $n \geq 0$ denoting the harmonic order, $m \in [-n, n]$ denoting the harmonic degree, $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi]$. The *polar part* is given by,

$$\Theta_n^m(\theta) = N_n^{|m|} P_n^{|m|}(\cos \theta) \quad (3.11)$$

and the *azimuthal part* is defined as,

$$\Phi_m(\phi) = \begin{cases} \sqrt{2} \cos(m\phi) & m > 0 \\ 1 & m = 0 \\ (-1)^m \sqrt{2} \sin(-m\phi) & m < 0 \end{cases} \quad (3.12)$$

¹The complex-valued SH functions y_n^m are defined in terms of complex sinusoids. This work mainly deals with real-valued functions, hence according to Chisholm [58], the definition of the real-form of the SH from their complex analogy is used.

By construction, the azimuthal part obeys the orthogonality condition such that,

$$\int_0^{2\pi} \Phi_m(\phi)\Phi_{m'}(\phi)d\phi = 2\pi\delta_{mm'} \quad (3.13)$$

Figure 3.8 demonstrates the functional dependencies of both the polar and azimuthal parts by considering the spherical basis $Y_4^2(\theta, \phi)$, where the polar part depends on $P_4^2(\cos \theta) = \frac{15}{2}(7 \cos^2 \theta - 1) \sin^2 \theta$ and the azimuthal part depends on $\cos(2\phi)$.

Figure 3.9 shows the visualization of the real-form of the spherical harmonic basis up to the fourth order, where the distribution of the basis function values are demonstrated via the textured unit spheres displayed on the top right corner of each basis function (refer to [Appendix H: \(Hemi\)Spherical Functions](#) from more details about spherical functions visualization). It can be observed that there are three main classes of SH basis; (1) The *zonal harmonics* which are these of degree $m = 0$, *i.e.* have no azimuthal dependence, thus they are circular symmetric. (2) The *sectoral harmonics* which have the form of $Y_{|m|}^m$. (3) The *tesseral harmonics* including all other harmonics, where the distribution of the function values over the unit sphere appears to be divided into several blocks in longitudinal and latitudinal directions.

3.2.1.2 Rotation of Spherical Functions

Rotations of the unit sphere S^2 are equivalent to rotations in \mathbb{R}^3 , they form a special orthogonal group $SO(3)$ whose elements are 3×3 real orthogonal matrices of unit determinant parameterized by the Euler angles (α, β, γ) with $\alpha \in [0, \pi]$, $\beta \in [0, 2\pi]$ and $\gamma \in [0, 2\pi]$. Any rotation $R_{\alpha, \beta, \gamma} \in SO(3)$ can be written as a product of three matrices $R_{\alpha, \beta, \gamma} = R_Z(\beta)R_Y(\alpha)R_Z(\gamma)$ where R_Z and R_Y represent the rotation matrices about the z -axis and y -axis, respectively, defined as,

$$R_Z(\beta) = \begin{pmatrix} \cos \beta & -\sin \beta & 0 \\ \sin \beta & \cos \beta & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (3.14)$$

and

$$R_Y(\alpha) = \begin{pmatrix} \cos \alpha & 0 & \sin \alpha \\ 0 & 1 & 0 \\ -\sin \alpha & 0 & \cos \alpha \end{pmatrix} \quad (3.15)$$




Figure 3.8: An illustration of a real spherical harmonic basis function with order/band-index $n = 4$ and degree $m = 2$. (a) The azimuthal part, $\cos(2\phi)$. (b) The polar part $P_4^2(\cos \theta)$, which is an associated Legendre polynomial having θ mapped to the interval $[-1, 1]$ through $\cos \theta$. (c) The spherical harmonic basis resulted from combining the azimuthal and the polar parts with a normalization factor $\sqrt{2} N_4^2$. Notice that textured unit sphere is used to visualize the spherical functions where the greenish and reddish regions correspond to positive and negative values, respectively.




Figure 3.9: Visualization of spherical harmonics (SH) basis up to the fourth order, where greenish and reddish regions correspond to positive and negative values, respectively. Small spheres displayed on the top right show unit spheres textured with the basis function values.

The rotation $R_{\alpha,\beta,\gamma}$ takes the north pole of the unit sphere with Cartesian coordinates $(0,0,1)$ to the point on a unit sphere having a unit direction $\vec{\omega} = \text{cart}(1,\alpha,\beta)$ with Cartesian coordinates $(\cos\beta\sin\alpha, \sin\beta\sin\alpha, \cos\alpha)$, while the equator plane (which is perpendicular to the north pole) is rotated by γ . See Figure 3.10 for illustration.




Figure 3.10: Illustration of 3D rotation using Euler angles.

Consider the spherical domain, any two unit directions $\vec{\omega}_1, \vec{\omega}_2 \in \mathcal{S}^2$ are related by some rotation $R_{\alpha,\beta,\gamma} \in SO(3)$ such that $R_{\alpha,\beta,\gamma}\vec{\omega}_1 = \vec{\omega}_2$. Thus $SO(3)$ forms a transformation group of \mathbb{R}^3 which causes rotation of vectors in \mathbb{R}^3 . For each rotation $R_{\alpha,\beta,\gamma} \in SO(3)$, square integrable functions f defined on \mathcal{S}^2 , *i.e.* the Hilbert space $L^2(\mathcal{S}^2)$, can be rotated by the rotation operator $f(R_{\alpha,\beta,\gamma}\vec{\omega})$ which acts point-wise $\forall \vec{\omega} \in \mathcal{S}^2$.

3.2.1.3 Rotational Invariance

Due to basis orthonormality, SH is known to be rotationally invariant [30]; rotation of spherical function coefficients is the same as rotating the function then projecting it onto the SH basis. This enables rotating spherical functions given their harmonics spectrum. The rotation process is thus a linear operation where coefficients between distinct bands don't interact. Thus, after rotation, a SH basis $Y_n^m(\cdot)$ of band n can be expressed as a linear combination of other un-rotated SH bases with weights depending on the rotation Euler angles such that

$$Y_n^m(R_{\alpha,\beta,\gamma}\vec{\omega}) = \sum_{m'=-n}^n \mathcal{D}_{mm'}^n(\alpha, \beta, \gamma) Y_n^{m'}(\vec{\omega}) \quad (3.16)$$

where $\vec{\omega} = \text{cart}(1, \theta, \phi)$ and \mathcal{D}^n are Wigner's rotation matrices [59] satisfying,

$$\mathcal{D}_{mm'}^n(\alpha, \beta, \gamma) = \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi} Y_n^{m'}(R_{\alpha,\beta,\gamma}\vec{\omega}) Y_n^m(\vec{\omega}) \sin\theta d\theta d\phi \quad (3.17)$$

For a given order n , \mathcal{D}^n is a $(2n+1) \times (2n+1)$ matrix which encodes how to express a rotated SH basis function in terms of all other SH bases of the same order. The orthogonality relation of \mathcal{D}^n matrices is given as [11],

$$\int_{\gamma=0}^{2\pi} \int_{\beta=0}^{2\pi} \int_{\alpha=0}^{\pi} \mathcal{D}_{mq}^n(\alpha, \beta, \gamma) \mathcal{D}_{m'q'}^{n'}(\alpha, \beta, \gamma) \sin \alpha d\alpha d\beta = \frac{8\pi^2}{2n+1} \delta_{nn'} \delta_{mm'} \delta_{qq'} \quad (3.18)$$

Since the z -axis rotation is only affecting the azimuthal part of a spherical function, contributing to the sin/cos part of the real SH, the rotation matrix for the n th band for real SH is defined as [60],

$$\begin{aligned} \mathcal{D}_{mm'}^n(\alpha, \beta, \gamma) &= \text{sgn}(m') \Phi_m(\beta) \Phi_{m'}(\gamma) \left(\frac{d_{|m'|, |m|}^n(\alpha) + (-1)^m d_{|m|, -|m'|}^n(\alpha)}{2} \right) \\ &- \frac{1}{2} \text{sgn}(m) \Phi_m(\beta) \Phi_{-m'}(\gamma) \left(\frac{d_{|m'|, |m|}^n(\alpha) - (-1)^m d_{|m|, -|m'|}^n(\alpha)}{2} \right) \end{aligned} \quad (3.19)$$

where,

$$\text{sgn}(x) = \begin{cases} 1 & x \geq 0 \\ -1 & x < 0 \end{cases} \quad (3.20)$$

and d^n is a matrix defining the y -axis rotation, it is defined explicitly as [61],

$$\begin{aligned} d_{mm'}^n(\alpha) &= \sqrt{(n+m)!(n-m)!(n+m')!(n-m')!} \\ &\sum_s \frac{(-1)^{m-m'+s} \left(\cos \frac{\alpha}{2}\right)^{2n+m'-m-2s} \left(\sin \frac{\alpha}{2}\right)^{m-m'+2s}}{(n+m'-s)!s!(m-m'+s)!(n-m-s)!} \end{aligned} \quad (3.21)$$

with the range of s determined such that all factorials are non-negative.

The computation of d^n matrices dominates the evaluation of SH rotations where numerical instability arises when computed explicitly as in Equation 3.21. Consequently, existing approaches tend to rely on recurrence relations such as [60, 61, 62]. According to Lessig *et al.* [63], Ivanic's [62] error grows proportionally with the spherical harmonic order compared to that of Blanco's [60] which in turn shows stable numerical performance for higher orders. Meanwhile Ivanic's approach consumes more time compared to that of Blanco's, such time difference increases with the harmonic order. In particular, at lower orders, rotation matrices evaluated by Ivanic's relations exhibit lower execution time, yet such performance is not maintained for $n \geq 3$.

As such, this work makes use of recurrence relations defined by Blanco *et al.* in [60] which are enumerated in Algorithm 1. Since $\lim_{\alpha \rightarrow \pi} \tan \alpha/2 = \infty$, steps (d) and (e)

in Algorithm 1 becomes invalid recurrence relations. However all elements of the last two columns are zeros due to the following relations [60] except for $d_{-n,-n}^n(\alpha) = 1$ and $d_{n-1,1-n}^n(\alpha) = n - 1$.

$$d_{n,m'}^n(\alpha) = \sqrt{\frac{(2n)!}{(n+m')!(n-m')!}} \left(\cos \frac{\alpha}{2}\right)^{n+m'} \left(-\sin \frac{\alpha}{2}\right)^{n-m'} \quad (3.22)$$

$$d_{n-1,m'}^n(\alpha) = (n \cos \alpha - m') \sqrt{\frac{(2n-1)!}{(n+m')!(n-m')!}} \left(\cos \frac{\alpha}{2}\right)^{n-1+m'} \left(-\sin \frac{\alpha}{2}\right)^{n-1-m'} \quad (3.23)$$

3.2.2 Illumination Harmonic Expansion

The spherical harmonic basis functions form a complete set of functions which are orthonormal over the surface of the unit sphere \mathbb{S}^2 , where the term *complete* implies that any well-behaved/smooth function can be expanded as an infinite series of these basis functions. The orthonormality property connotes that,

$$\langle Y_n^m, Y_n^{m'} \rangle = \int_0^{2\pi} \int_0^\pi Y_n^m(\theta, \phi) Y_n^{m'*}(\theta, \phi) \sin \theta d\theta d\phi = \delta_{nn'} \delta_{mm'} \quad (3.24)$$

holds for real SH functions.

A lighting function $L_i(\vec{\omega}_i)$ can be represented as a convex combination of the spherical harmonic (SH) basis functions such that,

$$L_i(\vec{\omega}_i) = \sum_{n=0}^{\infty} \sum_{m=-n}^n l_n^m Y_n^m(\vec{\omega}_i), \quad \vec{\omega}_i = R_{\alpha,\beta,\gamma} \vec{\omega}'_i \quad (3.25)$$

where $\{l_n^m\}$ are the illumination spectrum coefficients when projected onto the subspace spanned by the real SH basis. Thus the lighting function in Equation 5.2 can be rewritten as,

$$L_i(\vec{\omega}_i) = L_i(R_{\alpha,\beta,\gamma} \vec{\omega}'_i) = \sum_{n=0}^{\infty} \sum_{m=-n}^n \sum_{m'=-n}^n l_n^m \mathcal{D}_{mm'}^n(\alpha, \beta, \gamma) Y_n^{m'}(\vec{\omega}'_i) \quad (3.26)$$

For isotropic surfaces, where rotation of the tangential plane about surface normal doesn't affect surface appearance, γ can be arbitrary set to 0. Hence the SH rotation matrices are re-written as [11],

$$\mathcal{D}_{mm'}^n(\alpha, \beta) = \mathcal{D}_{mm'}^n(\alpha, \beta, 0) \quad (3.27)$$

where the orthogonality relation is modified such that,

$$\int_{\beta=0}^{2\pi} \int_{\alpha=0}^{\pi} \mathcal{D}_{mq}^n(\alpha, \beta) \mathcal{D}_{m'q}^{n'}(\alpha, \beta) \sin \alpha d\alpha d\beta = \frac{4\pi}{2n+1} \delta_{nn'} \delta_{mm'} \quad (3.28)$$

where the orthogonality relation is no longer held for index q due to the absence of γ integration [11].

Algorithm 1 d^n Matrices Computation

INPUT: (a) Euler rotation angle about y -axis, α , (b) Maximum SH order, N .

OUTPUT: Rotation matrices $d^n \forall n \in [0, N]$.

- 1: **Compute zero order matrix:** $d_{0,0}^0(\alpha) = 1$
- 2: **Compute the matrix of the first order:**

$$\begin{aligned} d_{0,0}^1(\alpha) &= \cos \alpha & d_{1,-1}^1(\alpha) &= \sin^2 \frac{\alpha}{2} \\ d_{1,0}^1(\alpha) &= -\frac{1}{\sqrt{2}} \sin \alpha & d_{1,1}^1(\alpha) &= \cos^2 \frac{\alpha}{2} \end{aligned}$$

To obtain other matrix entries, use symmetric relations defined as,

$$d_{m,m'}^n(\alpha) = d_{-m,-m'}^n(\alpha) = (-1)^{m+m'} d_{-m,-m'}^n(\alpha) = (-1)^{m+m'} d_{m,m'}^n(\alpha)$$

- 3: **Compute higher order matrices:** For $n = 2, 3, \dots, N$

- (a) Compute $d_{m,m'}^n(\alpha)$ for $m = 0, \dots, n-2$ and $m' = -m, \dots, m$ such that

$$\begin{aligned} d_{m,m'}^n(\alpha) &= \frac{n(2n-1)}{\sqrt{(n^2-m^2)(n^2-m'^2)}} \left\{ \left(d_{0,0}^1(\alpha) - \frac{mm'}{n(n-1)} \right) d_{m,m'}^{n-1}(\alpha) \right. \\ &\quad \left. - \frac{\sqrt{[(n-1)^2-m^2][(n-1)^2-m'^2]}}{(n-1)(2n-1)} d_{m,m'}^{n-2}(\alpha) \right\} \end{aligned}$$

- (b) Compute $d_{n,n}^n(\alpha)$

$$d_{n,n}^n(\alpha) = d_{1,1}^1(\alpha) d_{n-1,n-1}^{n-1}(\alpha)$$

- (c) Compute $d_{n-1,n-1}^n(\alpha)$

$$d_{n-1,n-1}^n(\alpha) = (n d_{0,0}^1(\alpha) - n + 1) d_{n-1,n-1}^{n-1}(\alpha)$$

- (d) Compute $d_{n,m'}^n(\alpha)$ for $m' = n-1, \dots, -n$

$$d_{n,m'}^n(\alpha) = -\sqrt{\frac{(n+m'+1)d_{1,-1}^1(\alpha)}{(n-m')d_{1,1}^1(\alpha)}} d_{n,m'+1}^n(\alpha)$$

- (e) Compute $d_{n-1,m'}^n(\alpha)$ for $m' = n-2, \dots, 1-n$

$$d_{n-1,m'}^n(\alpha) = -\frac{n d_{0,0}^1(\alpha) - m'}{n d_{0,0}^1(\alpha) - m' - 1} \sqrt{\frac{(n+m'+1)d_{1,-1}^1(\alpha)}{(n-m')d_{1,1}^1(\alpha)}} d_{n-1,m'+1}^n(\alpha)$$

- (f) Use symmetric relations in step 4 to obtain the rest of the elements with $m > 0$.
-

3.3 Illumination Signal Energy Content

The concept of the signal's energy content normally arises in constructing linear subspaces [64] where the signal's frequency spectrum portrays the range of frequencies for

which the signal has significant energy content. In case of illumination representation, it is of interest to maintain the minimum number of bases which captures a significant energy content of the lighting function. The illumination energy content can be defined as,

$$e_L = \int_{S^2} [L_i(\vec{\omega}_i)]^2 d\vec{\omega}_i \quad (3.29)$$

Using SH-based representation and basis orthonormality criterion, the energy content can be expressed as,

$$\begin{aligned} e_L &= \int_{S^2} \left(\sum_{n=0}^{\infty} \sum_{m=-n}^n l_n^m Y_n^m(\vec{\omega}_i) \right) \left(\sum_{n'=0}^{\infty} \sum_{m'=-n'}^{n'} l_{n'}^{m'} Y_{n'}^{m'}(\vec{\omega}_i) \right) d\vec{\omega}_i \\ &= \sum_{n=0}^{\infty} \sum_{n'=0}^{\infty} \sum_{m=-n}^n \sum_{m'=-n'}^{n'} l_n^m l_{n'}^{m'} \underbrace{\int_{S^2} Y_n^m(\vec{\omega}_i) Y_{n'}^{m'}(\vec{\omega}_i) d\vec{\omega}_i}_{\delta_{nn'} \delta_{mm'}} \\ &= \sum_{n=0}^{\infty} \sum_{m=-n}^n (l_n^m)^2 \end{aligned} \quad (3.30)$$

In order to eliminate the effect of the illumination configuration, the lighting function undergoing all possible 3D rotations is taken into account. Consider a light source with spectrum coefficients \tilde{l}_n^m such that,

$$L_i(\vec{\omega}_i) = \sum_{n=0}^{\infty} \sum_{m=-n}^n \tilde{l}_n^m Y_n^m(\vec{\omega}_i) \quad (3.31)$$

After a 3D rotation $R_{\alpha,\beta,\gamma} \in SO(3)$, this light source becomes,

$$\begin{aligned} L_i(R_{\alpha,\beta,\gamma}\vec{\omega}_i) &= \sum_{n=0}^{\infty} \sum_{m=-n}^n \sum_{o=-n}^n \tilde{l}_n^m \mathcal{D}_{mo}^n(\alpha, \beta, \gamma) Y_n^o(\vec{\omega}_i) \\ &= \sum_{n=0}^{\infty} \sum_{m=-n}^n \underbrace{\sum_{o=-n}^n \tilde{l}_n^o \mathcal{D}_{mo}^n(\alpha, \beta, \gamma)}_{l_n^m} Y_n^m(\vec{\omega}_i) \end{aligned} \quad (3.32)$$

Thus the average signal energy of the lighting function over all possible 3D rotations can be written as,

$$E\{e_L\} = \int_{SO(3)} e_L dR_{\alpha,\beta,\gamma} = \sum_{n=0}^{\infty} \sum_{m=-n}^n \int_{SO(3)} (l_n^m)^2 dR_{\alpha,\beta,\gamma} \quad (3.33)$$

where,

$$\int_{SO(3)} (l_n^m)^2 dR_{\alpha,\beta,\gamma}$$

$$\begin{aligned}
 &= \int_{SO(3)} l_n^m l_{n'}^{m'} dR_{\alpha,\beta,\gamma} \\
 &= \int_{SO(3)} \left(\sum_{o=-n}^n \tilde{l}_n^o \mathcal{D}_{mo}^n(R_{\alpha,\beta,\gamma}) \right) \left(\sum_{o'=-n'}^{n'} \tilde{l}_{n'}^{o'} \mathcal{D}_{m'o'}^{n'}(R_{\alpha,\beta,\gamma}) \right) dR_{\alpha,\beta,\gamma} \\
 &= \sum_{o=-n}^n \sum_{o'=-n'}^{n'} \tilde{l}_n^o \tilde{l}_{n'}^{o'} \int_{SO(3)} \mathcal{D}_{mo}^n(R_{\alpha,\beta,\gamma}) \mathcal{D}_{m'o'}^{n'}(R_{\alpha,\beta,\gamma}) dR_{\alpha,\beta,\gamma} \quad (3.34)
 \end{aligned}$$

From the orthogonality relation of \mathcal{D}^n matrices Equation 3.18, it can be shown that,

$$\int_{SO(3)} \mathcal{D}_{mo}^n(R_{\alpha,\beta,\gamma}) \mathcal{D}_{m'o'}^{n'}(R_{\alpha,\beta,\gamma}) dR_{\alpha,\beta,\gamma} = \frac{1}{2n+1} \delta_{nn'} \delta_{mm'} \delta_{oo'} \quad (3.35)$$

where,

$$dR_{\alpha,\beta,\gamma} = \frac{1}{8\pi^2} \sin \alpha d\alpha d\beta d\gamma \quad (3.36)$$

Thus the average simply becomes,

$$E\{e_L\} = \sum_{n=0}^{\infty} \frac{1}{2n+1} \underbrace{\sum_{o=-n}^n (\tilde{l}_n^o)^2}_{e_L^\mu(n)} \quad (3.37)$$

where $e_L^\mu(n)$ is the average energy content maintained by the n -th illumination order.

Thus let's define the representation approximation accuracy to be the ratio of the signal energy content of the truncated series to that of the infinite one.

$$Acc_L(N) = \frac{\sum_{n=0}^N e_L^\mu(n)}{\sum_{n=0}^{\infty} e_L^\mu(n)} \quad (3.38)$$

In the computations below, the infinite series in the denominator of Equation 3.38 is computed up to an asymptotic illumination order where infinite precision rational numbers are used for such high orders. In the sequel, the average energy profile of different lighting functions is investigated including point light sources and natural illuminations.

3.4 Frequency Properties of Light Sources

A light source defines a directionally varying incident radiance function $L_i(\vec{\omega}_i)$. For point light sources (near or distant), the incident radiance comes from a single direction while for area or infinite-area sources, different directions carry different incident illumination. This section investigates the frequency properties of different light source models when

being represented in the space spanned by the SH basis functions.

3.4.1 Distant Sources

A single directional source can be defined as a delta function. Let $\vec{\omega}_s$ be a unit vector denoting the source direction in the global coordinates. Consider a surface patch centered at the global origin and oriented such that the global and local frames coincide¹. Hence the lighting function can be defined as,

$$L_i(\vec{\omega}_i) = \delta(\vec{\omega}_i - \vec{\omega}_s) \quad (3.39)$$

As such the lighting spectrum can be obtained by integrating the lighting function in Equation 3.39 with each harmonic basis function. Since a directional source is defined as a delta function, the integral boils to a single term where the light coefficients are the basis functions evaluated at the light source direction.

$$\tilde{l}_n^m = \int_{\mathbb{S}^2} L_i(\vec{\omega}_i) Y_n^m(\vec{\omega}_i) d\vec{\omega}_i = Y_n^m(\vec{\omega}_s) \quad (3.40)$$

Since all possible 3D rotations are taken into account when computing the average signal energy content, it is safe to assume that the light source is positioned at the north pole with $\vec{\omega}_s = \text{cart}(1, 0, 0)$, where,

$$Y_n^m(\vec{\omega}_s) = Y_n^m(0, 0) = \begin{cases} \sqrt{\frac{2n+1}{4\pi}} & m = 0 \\ 0 & m \neq 0 \end{cases} \quad (3.41)$$

Consequently, distant sources have a non-decayable energy profile² where each illumination order maintains the same average energy content such that,

$$e_L^\mu(n) = \frac{1}{2n+1} \sum_{o=-n}^n \left(\tilde{l}_n^o\right)^2 = \frac{1}{4\pi} \quad (3.42)$$

While this implies that an infinite number of basis is needed in order to fully recover a directional light source, the goal is to represent the image irradiance integral where the surface reflectance acts as a low pass filter on the lighting spectrum. As such, a

¹Having a surface patch deviating from these assumptions would involve 3D rotation of the illumination direction *w.r.t.* the surface patch. This is accounted for when computing the average signal energy content where all 3D rotations are taken into consideration.

²This corresponds to the Fourier transform of a delta function in 1D.

finite-dimensional model will be derived later in the dissertation.

3.4.2 Point Sources

Harmonic basis, in general, are known to provide optimal basis for images of convex objects under distant light [65]. Nonetheless, light coming from close by objects, indoor environments and reflections from surrounding surfaces are examples of common near light scenarios, where the lighting function is dominated by higher frequencies. Thus it is crucial to relax the distant light assumption to determine what distance to the light source is sufficient to be considered infinite.

Consider a point light source located on the locus of a global sphere with radius r surrounding a surface patch centered at the global origin and oriented such that the global and local frames coincide. According to the standard nearby point source model, the light intensity arrives at the surface patch is attenuated by a factor of $\frac{1}{r^2}$ where r^2 denotes the square the distance between the surface patch and the light source.

In case of a single point light source, the incident radiance comes from only a single direction; it can be defined as a delta function. Let $\vec{\omega}_s$ be a unit vector denoting the source direction in the global coordinates. Hence the lighting function can be defined as,

$$L_i(\vec{\omega}_i) = \frac{1}{r^2} \delta(\vec{\omega}_i - \vec{\omega}_s) \quad (3.43)$$

Therefore the lighting spectrum can be obtained by,

$$\tilde{l}_n^m = \int_{S^2} L_i(\vec{\omega}_i) Y_n^m(\vec{\omega}_i) d\vec{\omega}_i = \frac{1}{r^2} Y_n^m(\vec{\omega}_s) \quad (3.44)$$

Assuming a light source located at the north pole, the average energy content maintained by the n -th order expansion of a point light source can be obtained as¹,

$$e_L^\mu(n; r) = \frac{1}{4\pi r^4} \quad (3.45)$$

While the energy content depends on how near the light source is, point light sources still have a non-decayable energy profile similar to distant light sources.

¹We use a semicolon notation in the function argument to denote given constant quantities.

3.4.3 Infinite Area Sources

Recall that infinite area sources are represented by two-dimensional texture maps defining the directionally-varying emitted radiance as a function of incident directions in spherical coordinates. As such, Monte Carlo integration¹ is considered to be a reasonable approach to solve for the lighting coefficients [30]. Figure 3.11 shows the average energy content and the approximation accuracy of Debevec’s environment maps [55]. It can be noticed that while the energy content of point light sources remains constant over frequency, *i.e.* illumination order, the energy content of infinite area sources (*i.e.* natural illumination) rapidly decrease at high frequencies. The rate of decay primarily depends on the frequency content of the light source. For example Funston Beach illumination map can be considered as a low-frequency light source compared to other maps since its approximation accuracy saturates at 100% at lower frequencies, meanwhile the approximation accuracy of Galileo Tomb and St. Peter Basilica maps asymptotically approaches 100% due to their higher frequency content. Such results are emphasized in Figure 3.12 which shows a sample of a wider variety of environment maps taken from sIBL archives. Hence it can be concluded that most of natural (real world) illuminations tend to have low-frequency content where most of the lighting signal content is captured using lower illumination orders. This motivates the concept of deriving lower-dimensional subspace to model natural lighting conditions.

3.5 Summary

This chapter reviewed the definition of different light source models deployed in vision and graphics applications. Spherical harmonics (SH) was also reviewed as the key ingredient which is widely used to represent illumination in the frequency domain. According to the frequency properties of different light source models, it has been shown that directional/distant sources have non-decayable spectrum whereas natural illumination has a spectrum which rapidly decrease with frequency. This discussion benefits subsequent chapters when explaining the image formation process and how the image irradiance equation can be represented in the frequency domain.

¹Refer to [Appendix I: Projection Onto Basis - Monte Carlo Integration](#) for more details about Monte Carlo integration and projection onto basis functions.

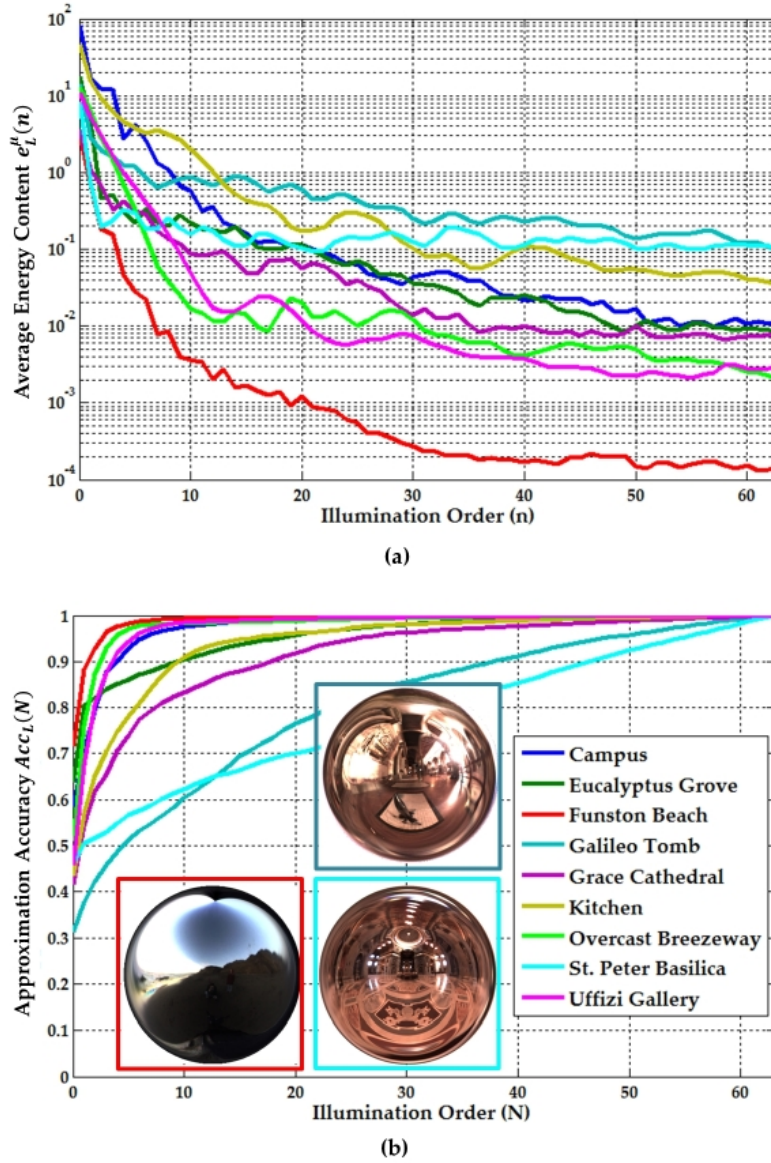


Figure 3.11: Frequency properties of sample infinite area light sources [55]: (a) Average energy content $e_L^\mu(n)$ (log scale) as a function of illumination order n . (b) Approximation accuracy $Acc_L(N)$ as a function of the truncating illumination order N . Notice that Funston Beach illumination map (shown in red box) is considered a low-frequency environment map due to the rapidly decrease of its average energy content at high frequencies, *i.e.* illumination order, compared to other sources, hence its approximation accuracy saturates at 100% at lower orders. Meanwhile the approximation accuracy of Galileo Tomb and St. Peter Basilica maps asymptotically approaches 100% due to their higher frequency content. Note: Lighting coefficients are solved for using Monte Carlo integration of number of sample cells $J = 250$.

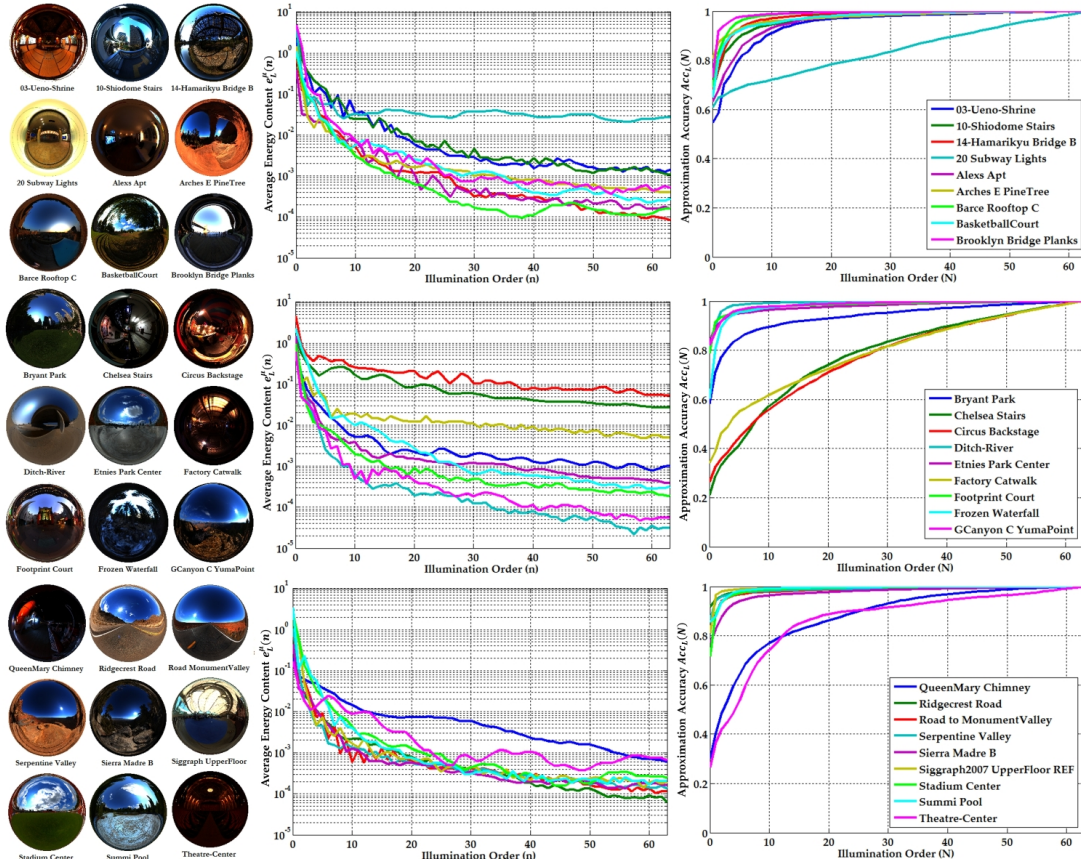


Figure 3.12: Frequency properties of sample environment maps from sIBL whose angular maps are shown in the left. Middle: Average energy content $e_L^\mu(n)$ (log scale) as a function of illumination order n . Right: Approximation accuracy $Acc_L(N)$ as a function of the truncating illumination order N . Notice that the approximation accuracy of low-frequency environment maps saturate at 100% at lower illumination orders, *e.g.* Brooklyn Bridge Planks, due to the rapidly decrease of its average energy content at high frequencies. On the other hand, high-frequency illuminations such as Subway Lights and Chelsea Stairs maps asymptotically approaches 100% due to their higher frequency content. Note: Lighting coefficients are solved for using Monte Carlo integration of number of sample cells $J = 250$.

PART II

IMAGE IRRADIANCE REPRESENTATION

PHENOMENOLOGICAL REPRESENTATION OF ARBITRARY REFLECTANCE

The visual appearance of real-world surfaces is the net result of surface reflectance characteristics when exposed to illumination. Appearance models, under local illumination assumption, can be constructed using phenomenological models which capture surface appearance through mathematical modeling of the reflection process. This yields an integral equation, known as the *reflectance equation*, describing the surface radiance, which depends on the interaction between the incident light field and the surface bidirectional reflectance distribution function (BRDF). The BRDF embodies all the information about the reflectance behavior of a surface. It is a function defined on the cartesian product of two hemispheres corresponding to the incident and outgoing directions; the nature way to represent such a hemispherical function is to use hemispherical basis. However, due to their compactness in the frequency space, spherical harmonics (SH) has been extensively used for this purpose.

This chapter addresses the geometrical compliance of hemispherical basis for representing plausible surface BRDFs. A Cartesian product of the hemispherical harmonics (HSH) is proposed to provide a compact, yet accurate representation for arbitrary BRDFs, while satisfying the Helmholtz reciprocity property. An analytical analysis and experimental justification is provided that for a given truncating reflectance order, the proposed hemispherical basis provide better approximation accuracy of the BRDF when compared to similar bases in literature. In case of isotropic surfaces, where the surface appearance remains the same under rotations of the tangential plane about the surface

normal, it is proved that the number of basis is reduced drastically, however the proposed Helmholtz HSH-based basis contains more basis functions than the Zernike-based ones at a given approximating order. The presented validation is extended to model scattered reflectance data which might violate the Helmholtz reciprocity property; this is filtered out through the process of projecting them on the subspace spanned by the HSH-based basis, where the reciprocity property is preserved in the least-squares sense. This chapter ends by investigating the closed form of the proposed basis in case of directional hemispherical reflectance where the BRDF domain is reduced to be two dimensional; *i.e.* at fixed incoming or outgoing direction. In such a case, an implicit assumption holds where there is a smooth BRDF variation *w.r.t.* the fixed direction.

4.1 Introduction

Various types of surfaces can be identifiable by their appearance, which is the net result of the surface reflectance characteristics when exposed to illumination. There are three fundamental ways to construct appearance models; (1) understanding the underlying physical process of light-surface interaction; leading to *Fresnel's equations*, (2) understanding the structure of the surface's material at a micro-scale level, where *layered models* can be investigated, and (3) using phenomenological models to capture surface appearance through mathematical modeling of the reflection process, leading to the *reflectance equation*. The theoretical formulation and modeling of the reflectance equation is of a great interest to numerous computer vision tasks concerned with inverse rendering problems¹, such as illumination modeling [6], surface reflectivity estimation/analysis [11, 32], shape and albedo recovery [33, 34], shape from shading [36], photometric stereo [37, 38], object detection and recognition [39], to name a few.

Using geometric optics, the *reflectance equation*, under the assumption of no surface emittance, is an integral equation describing the surface reflected radiance which results from the interaction between the incident light field and the surface bidirectional reflectance distribution function (BRDF). Being considered as the most general model of local reflection, the surface BRDF is a mathematical and physical abstraction which captures the positional and directional dependencies of reflection, while in the meantime, hides the complexity of the interaction of light with matter.

¹Inverse rendering problems appear in applications in which scene properties, such as the illumination, scene 3D geometry and/or reflectance properties, are inferred from images.

A BRDF model can be obtained by (1) deriving an analytic formula based on either physical principles or simple formulation designed to mimic some kind of reflection, (2) simulating an assumed surface micro-geometry model, or (3) measuring BRDF values based on empirical observation [66]. In all cases, it is required to devise a compact, yet accurate, BRDF representation of real-world materials which maintains the physical characteristics of surface reflectance. As such, BRDF representation has received much research attention from computer graphics and vision communities [67], where low-dimensional parametric models enable formulating inverse rendering problems as parameter estimation [68].

A tabular representation, as simple as is, can be used to represent an arbitrary BRDF. However due to its four-dimensionality, full BRDF data are very scarce where the measurements can be restricted to the plane of incidence and is usually limited by some angular resolution. Further, the process of BRDF measurements acquisition could take hours to mechanically vary light source and sensor positions [66]. Interpolation, extrapolation and smoothing are usually used to benefit from measured BRDF data in a convenient manner.

At the other end of the spectrum, analytic formula can present a very compact model for BRDF representation. Such formula can be either based on empirical observations, such as Lambertian model with a constant BRDF [69] and Blinn/Phong model [15, 70] (based on cosine-lobes), or physically based modeling of the microscopic surface geometry (*i.e.* distribution of micro-facet orientations) such as Torrance-Sparrow [12] for rough specular surfaces and Oren-Nayar [14] for rough diffuse surfaces. Despite the representation compactness, the lack of generality is the main drawback for such models, where there is no guarantee that these analytical models can represent arbitrary measured reflectance data [67], thus they only represent limited classes of surfaces.

On the middle ground, phenomenological models represent an arbitrary/real-world BRDF as a linear combination of a complete set of orthonormal basis functions, analogous to Fourier basis representing functions over the real line. Such low-dimensional representation has been used as an alternative to both measured/tabulated BRDFs and analytic models for three main reasons: First, most BRDFs are smooth functions [19], *i.e.* they depend slowly on the directions and hence reflectance information is encoded via low frequency components. As such, they are considered as good candidates for representation using smooth orthonormal basis functions. Second, the human visual system

is, in general, insensitive to appearance fine details [67], *i.e.* high frequency components of surface reflectance, as long as features such as specularity and color are maintained. Matusik *et al.* [17] supported this proposition by representing a BRDF as a linear combination of BRDF basis derived from densely sampled reflectance measurements of 100 real-world materials. Third, the lack of acquired data in areas such as near grazing angles, near retro-reflection directions (*i.e.* scattering in the backward direction) and along specular directions have justified the use of basis functions for BRDF representation to extrapolate missing BRDF measurements while filtering out violations of physical reflectance requirements. Whereas analytical models are compact in principle, they only approximate the reflectance phenomena they were designed for. Moreover, there is no single model that is expected to represent all real-world materials. In that regard, phenomenological models are considered a practical approach for reflectance representation [21].

Surface BRDFs are not arbitrary functions; in theory they are governed by basic principles of physics [71]. (1) *Non-negativity* - A BRDF should attain non-negative values since both radiance and irradiance are non-negative. (2) *Energy conservation* - Under the assumption of no surface emittance, a surface cannot reflect more light that it receives from the incident upper hemisphere. Maintaining such property is important to avoid color bleeding while producing plausible rendered images with accurate reflection behavior especially near grazing incident angles [72]. (3) *Reciprocity* - Helmholtz's law of reciprocity [52] implies that the BRDF is invariant to permutation of incident and outgoing light directions. BRDF models which obey these principles are said to be *physically plausible BRDFs* [73] since they can reproduce natural reflection [66], in contrast to non-physical ones which do not exist in nature. Another advantage of adopting a representation which maintains reciprocity is that, in forward rendering applications, light beams being reflected from a fixed outgoing/viewing direction can be traced while accumulating the incident radiance at different incident directions [4].

A BRDF can be classified as an *isotropic* one if the surface appearance remains the same under rotations of the tangential plane about the surface normal, or an *anisotropic* one when this does not hold, *i.e.* appearance is affected by the azimuthal orientation of the surface. It is advantageous to split the BRDF spectrum into an isotropic part and an anisotropic part, nonetheless it has been noted in [21] that the contribution of the anisotropic part is much smaller compared to that of the isotropic one. Further, while

anisotropy prevails that there are preferred reflectance angles, isotropy holds when the medium of interest is composed of a large number of particles [74] since there is no longer preferred orientations. Hence the reflectance of anisotropic surfaces can be approximated as being isotropic.

Isotropy and reciprocity have been adopted in vision to reduce the BRDF domain for surface reconstruction purposes [75, 76]. They have also been used in graphics in image-based rendering [77, 78]. For example, Helmholtz reciprocity property, implying symmetry of surface reflectance when light and camera positions are swapped, is exploited to recover accurate surface normals [76]. This is termed as *Helmholtz Stereopsis* which is primarily used for image-based reflectance measurement and rendering [79]. Tan *et al.* [80] exploited isotropy and reciprocity in a broader sense to resolve the generalized base-relief (GBR) ambiguity which is inherent to uncalibrated photometric stereo. As such, devising a low-dimensional phenomenological model which yields physically plausible BRDFs goes to the heart of various vision and graphics applications.

The BRDF is a function defined on the cartesian product of two hemispheres corresponding to the incident and outgoing directions; the nature way to represent such a hemispherical function is to use hemispherical basis. However, due to their compactness in the frequency space, spherical harmonics (SH) has been extensively used for this purpose [4, 19, 81]. Whereas SH is a complete set of orthonormal basis on the full unit sphere, hemispherical functions present discontinuities at the boundary of the hemisphere when represented in the spherical domain [20], demanding more coefficients for accurate representation.

The goal of this chapter is thus to derive a low-dimensional, *i.e.* compact, phenomenological BRDF model which can (1) address the physical compliance of the model to yield physically plausible BRDFs, (2) adhere to the geometrical nature of surface BRDF as being defined on a hemispherical domain rather than a spherical one, (3) address the trade-off between approximation accuracy and model compactness (memory consumption) and (4) achieve accuracy comparable to the state-of-art phenomenological models. This suggests a linear model in terms of hemispherical orthonormal basis functions which obey symmetries such as isotropy and reciprocity. An inherent advantage of linear models is avoiding non-linear optimization processes used when employing polynomial functions [66]. In particular, a cartesian product of the hemispherical harmonics (HSH) is proposed to provide a compact, yet accurate, representation for arbitrary

BRDFs, while satisfying the physical characteristics of surface reflectance. The proposed BRDF model can be used in place of simple Lambertian models in algorithms such as shape-from-shading and photometric stereo.

This work is similar in spirit to that of Koenderink and van Doorn [21] where they used hemispherical Zernike polynomials. Whereas the polar part of the proposed basis is defined in terms of shifted associated Legendre polynomials. The set of such polynomials is distinguished by the property that it contains a polynomial for every combination of order and degree [82], compared to Zernike polynomials which are restricted to even differences between polynomial order and degree. Thus, an analytical analysis and experimental justification is provided that for a given approximation order, the proposed HSH-based basis provide better approximation accuracy when compared to the Zernike-based ones [21], while, avoiding the high computational complexity inherent from Zernike polynomials [20].

The accuracy of the proposed basis functions is assessed analytically as well as numerically, starting with a theoretical analysis of the representation accuracy of the proposed model on different analytic physical BRDFs ranging from ideal diffuse and specular reflection to micro-facet based reflection models. The accuracy of such approximations is then evaluated using measured reflectance where scattered BRDF data might violate the Helmholtz reciprocity property.

4.2 Related Work

Whereas real-world materials can be used as basis for general BRDF representation [17], using orthonormal set of basis functions robustify the process of coefficient estimation where BRDF coefficients, *i.e.* spectrum, can be estimated individually using Monte Carlo integration as opposed to estimating them simultaneously using least squares. Another intrinsic property of orthonormal basis is that it amounts for efficient truncated decomposition where the omission of high-frequency components is likely to yield good reflectance representation.

While several attempts adopted SH for BRDF representation, such models do not comply with the geometrical characteristics of non-emitting surface reflectance. For example, Kajiyama and Herzen [74] used SH to derive an analytical scattering function describing the radiation scattering in volume densities such as the case of clouds and

fog. However this method used the wave theory of light as apposed to geometric optics. Cabral *et al.* [4] and Sillion *et al.* [83] used such representation as a numerical approximation of the BRDF. In particular, assuming constant outgoing (*a.k.a.* viewing) direction, Cabral *et al.* [4] derived SH coefficients for isotropic clamped BRDFs¹ based on a height field geometry without relying on neither analytic BRDF model nor measured reflectance data. They used a tabulated (binned) version of BRDFs while dividing the outgoing hemisphere into specific number of bins. This amounted for using smaller lookup tables while smoothing and interpolation took place. The hemispherical topology of reflectance was handled by replacing $\cos \theta'_i$ with the nonlinear term $\max(0, \cos \theta'_i)$. While Cabral *et al.* [4] used this representation for the purpose of simulating diffuse and glossy reflections of the environment. Sillion *et al.* [83] tried to use such representation to solve the global illumination problem for arbitrarily complex reflectance models. By expanding the BRDF at fixed incoming direction, they represented the SH coefficients as functions of the incident angle which were stored as one dimensional cubic splines.

Westin *et al.* [19] simulated optical scattering to introduce a physically-based Monte Carlo algorithm to approximate arbitrary BRDFs. They used SH to define the basis over the cartesian product of two spheres, taking the advantage of symmetry and reciprocity to reduce the non-zero coefficients representing the BRDF (isotropic as well as anisotropic ones). This representation does not require discretizing scattering directions as in the work of Cabral *et al.* [4]. The reciprocity property was enforced by using a combined basis function $Y_p^q(\vec{\omega}'_i)Y_r^s(\vec{\omega}'_o)$. The dependence on the incident and outgoing directions was encoded in a large matrix which stores the SH coefficients. On the other hand, to avoid on-the-fly evaluation of high-order basis functions, Kautz *et al.* [81] parameterized the clamped BRDF by the outgoing direction to represent the 4D space of arbitrary BRDF (isotropic and anisotropic) in a 2D offline-table of SH coefficients.

On the other hand, several hemispherical basis have been proposed in literature to represent hemispherical functions. Sloan *et al.* [84] used SH as in [19] to represent an even-reflected (about xy -plane) version of a hemispherical function. Coefficients were found using least squares SH in contrast to Monte Carlo integration in [19], however this leads to non-zero values in the lower hemisphere. Koenderink *et al.* [21, 53] used Zernike polynomials [85], which are basis functions defined on a disk, to build hemispherical basis. Yet, such polynomials have high computational cost. Makhotkin [57] and Gautron *et*

¹Clamped BRDF is the product of the BRDF with the foreshortening factor $\cos \theta'_i$.

al. [20] proposed hemispherically orthonormal basis through mapping the negative pole of the sphere to the border of the hemisphere. Such contraction was achieved through shifting the adjoint Jacobi polynomials [57] and the associated Legendre polynomials [20] without affecting the orthogonality relationship. Recently, Habel and Wimmer [22] used the SH as an intermediate basis to define polynomial-based hemispherical basis. They used the SH basis functions which are symmetric to the $z = 0$ plane since they are orthogonal over the upper hemisphere. While other basis functions are shifted the same way that was proposed by [20]. Although such basis definition leads to polynomial basis, this inhibits deriving an analytical expression of the harmonic expansion of the surface reflectance function and in turn of the irradiance integral.

This work adopts the hemispherical basis defined by Gautron *et al.* [20] to propose the surface reflectance basis. While [20] introduced the hemispherical harmonics (HSH) in the context of forward rendering, the primary focus of this work is on the reverse problem. Along the same line, Gautron *et al.* [20] used two HSH transformations to define hemispherical basis. Nonetheless, such definitions do not guarantee Helmholtz reciprocity. On the other hand, Koenderink *et al.* [21, 53] used Zernike polynomials [85] to define hemispherical basis that satisfy such a property, however such polynomials are known to have high computational complexity [20] and in the meantime, they are not defined for all combinations of polynomial orders and degrees.

4.3 Hemispherical Harmonics (HSH)

4.3.1 Shifted Associated Legendre Polynomials

The idea of using linear transformation to shift orthogonal polynomials has been commonly used to change the domain of well-known basis. For example, associated Legendre polynomials and adjoint Jacobi polynomials are shifted to define hemispherical basis functions as an adapted version of the spherical ones [20, 57].

Consider an orthogonal set of functions $\{B_k(x)\}$ defined over the interval $x \in [a, b]$ with weighting function $w(x)$, a linear transformation of the independent variable x to $\kappa_1 x + \kappa_2$ with $\kappa_1 \neq 0$ causes the set of functions $\{B_k(\kappa_1 x + \kappa_2)\}$ to be orthogonal over the interval $[\frac{a-\kappa_2}{\kappa_1}, \frac{b-\kappa_2}{\kappa_1}]$ with weighting function $w(\kappa_1 x + \kappa_2)$ [86]. Moreover, if the set $\{B_k(x)\}$ is orthonormal then $\{(\text{sign}(\kappa_1))^k \sqrt{|\kappa_1|} B_k(\kappa_1 x + \kappa_2)\}$ maintains the orthonormality property.

With $\kappa_1 = 2$ and $\kappa_2 = -1$, shifted associated Legendre polynomials $\tilde{P}_n^m(x)$ can be defined over the interval $x \in [0, 1]$ ¹,

$$\tilde{P}_n^m(x) = P_n^m(2x - 1) \quad (4.1)$$

The orthogonality relationship of the shifted associated Legendre polynomials is defined with weighting function $\tilde{w}(x) = w(2x - 1) = 1$ as [20],

$$\int_0^1 \tilde{P}_n^m(x) \tilde{P}_{n'}^m(x) dx = \frac{(n+m)!}{(2n+1)(n-m)!} \delta_{nn'} \quad (4.2)$$

4.3.2 HSH Definition

An infinite series of shifted associated Legendre polynomials can be used to express any piecewise continuous function over the interval $[0, 1]$, where band-limited functions can be exactly reconstructed using finite number of polynomials. Shifted Legendre polynomials can be combined with sinusoidal functions to define the hemispherical harmonics representation for non-circular symmetric functions where the polar angle $\theta \in [0, \pi/2]$ is mapped to the interval $[0, 1]$ by using $x = \cos \theta$ as the independent variable of the polynomial $\tilde{P}_n^m(x)$. In case of circular-symmetric functions, zonal harmonics are used which reduce to the shifted associated Legendre polynomials with zero degree.

Analogical to SH definition, the real hemispherical harmonics (HSH) can be written as [20]²,

$$H_n^m(\theta, \phi) = \begin{cases} \sqrt{2} \tilde{N}_n^m \tilde{P}_n^m(\cos \theta) \cos(m\phi) & m > 0 \\ N_n^0 \tilde{P}_n^0(\cos \theta) & m = 0 \\ (-1)^m \sqrt{2} \tilde{N}_n^{|m|} \tilde{P}_n^{|m|}(\cos \theta) \sin(-m\phi) & m < 0 \end{cases} \quad (4.3)$$

with \tilde{N}_n^m being defined as the normalization factor,

$$\tilde{N}_n^m = \sqrt{\frac{2n+1}{2\pi} \frac{(n-m)!}{(n+m)!}} \quad (4.4)$$

¹The main idea behind this shifting is limiting the domain of the inclination angle θ to $[0, \pi/2]$ when substituting x by $\cos \theta$.

²Note the difference between this work and [20] in the factor of $(-1)^m$, since it is preferred to include this factor in the harmonics definition rather than including it in the associated Legendre polynomial definition.

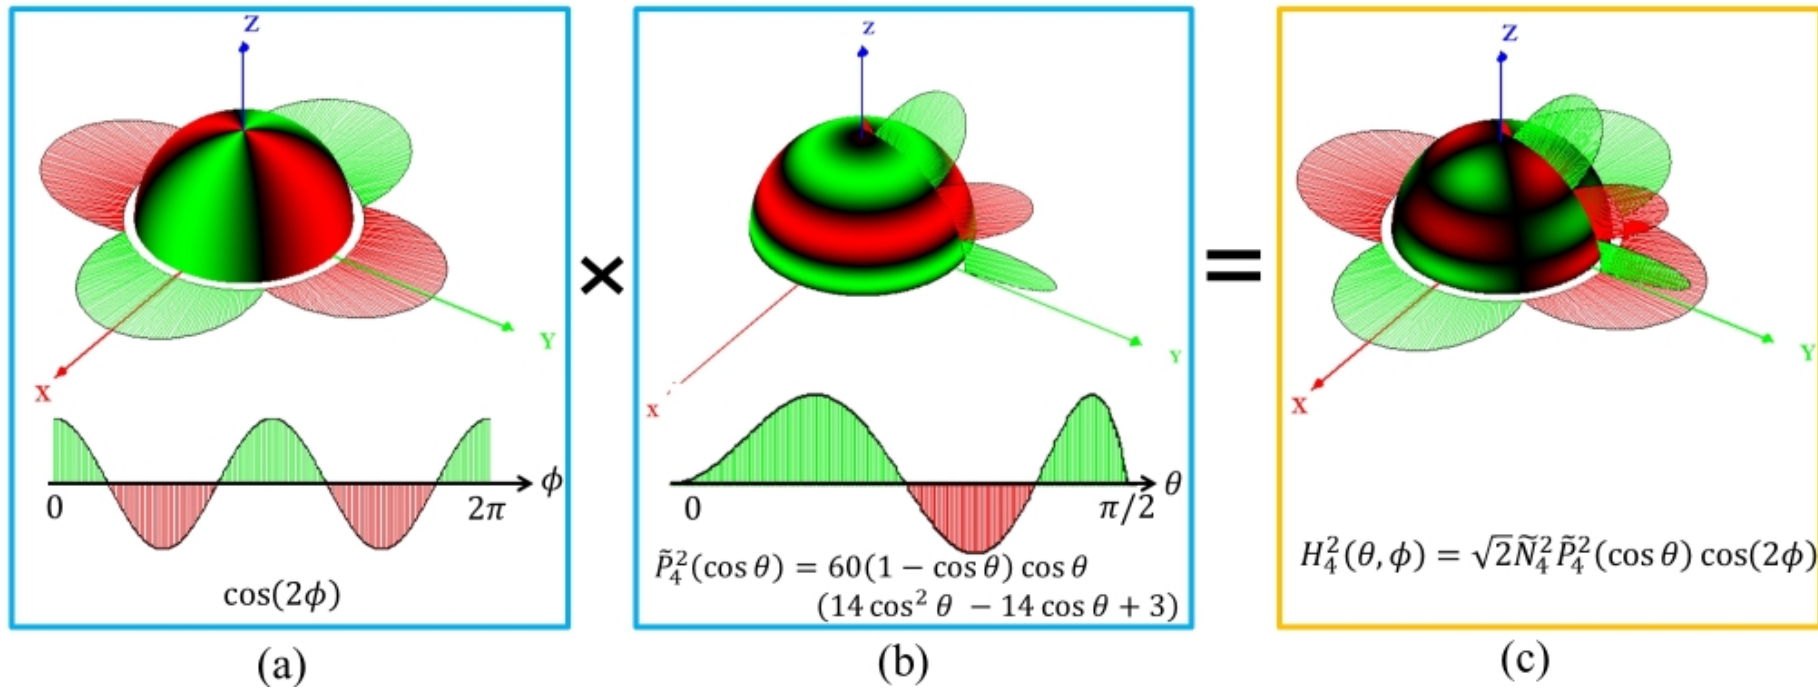


Figure 4.1: An illustration of a real hemispherical harmonic basis function with order/band-index $n = 4$ and degree $m = 2$. (a) The azimuthal part, $\cos(2\phi)$. (b) The polar part $\tilde{P}_4^2(\cos \theta)$, which is a shifted associated Legendre polynomial having θ mapped to the interval $[0, 1]$ through $2 \cos \theta - 1$. (c) The hemispherical harmonic basis resulted from combining the azimuthal and the polar parts with a normalization factor $\sqrt{2} \tilde{N}_4^2$. Notice that textured unit hemisphere is used to visualize the hemispherical functions where the greenish and reddish regions correspond to positive and negative values, respectively.




Figure 4.2: Visualization of hemispherical harmonics basis up to the fourth order, where greenish and reddish regions correspond to positive and negative values, respectively. Small hemispheres displayed on the top right show unit spheres textured with the basis function values.

Eqn. Equation 4.3 can be short-handed as,

$$H_n^m(\theta, \phi) = \tilde{\Theta}_n^m(\theta)\Phi_m(\phi) \quad (4.5)$$

where $n \geq 0$ denoting the harmonic order, $m \in [-n, n]$ denoting the harmonic degree, $\theta \in [0, \pi/2]$ and $\phi \in [0, 2\pi]$. The *polar part* is given by $\tilde{\Theta}_n^m(\theta) = \tilde{N}_n^{|m|}\tilde{P}_n^{|m|}(\cos \theta)$ and the *azimuthal part* is defined as Equation 3.12.

Figure 4.1 demonstrates the functional dependencies of both the polar and azimuthal parts by considering the hemispherical basis $H_4^2(\theta, \phi)$, where the polar part depends on $\tilde{P}_4^2(\cos \theta) = P_4^2(2 \cos \theta - 1) = 60(1 - \cos \theta) \cos \theta(14 \cos^2 \theta - 14 \cos \theta + 3)$ and the azimuthal part depends on $\cos(2\phi)$.

Figure 4.2 shows the visualization of the real-form of the hemispherical harmonic basis up to the fourth order, where the distribution of the basis function values are demonstrated via the textured unit hemispheres displayed on the top right corner of each basis function.

The HSH basis functions form a complete set of functions which are orthonormal over the surface of the unit hemisphere Ω , implying that any hemispherical smooth function can be expanded as an infinite series of these basis functions. The orthonormality property connotes that,

$$\langle H_n^m, H_{n'}^{m'} \rangle = \int_0^{2\pi} \int_0^{\pi/2} H_n^m(\theta, \phi) H_{n'}^{m'*}(\theta, \phi) \sin \theta d\theta d\phi = \delta_{nn'} \delta_{mm'} \quad (4.6)$$

holds for real HSH functions.

4.4 Proposed Helmholtz HSH-based Reflectance Basis

As noted earlier, the surface bidirectional reflectance distribution function (BRDF) is defined as the ratio of the reflected radiance exiting along the outgoing direction $\vec{\omega}'_o$ to the surface irradiance incident from the incoming direction $\vec{\omega}'_i$, where both directions are defined with respect to the surface normal¹. While it is often preferable to work with the clamped BRDF [11], *i.e.* $f_r(\vec{\omega}'_i, \vec{\omega}'_o) \cos \theta'_i$, to enforce C^1 continuity at the equator and hence reduce the ringing in the approximation (*a.k.a.* Gibbs phenomenon) [83], especially at grazing incident angles where singular performance is expected [21]. Yet, it is

¹Note that the primed coordinates are used to indicate measurements with respect to the local frame of a surface point.

not required to deal with the clamped BRDF since the proposed basis is hemispherical in nature, *i.e.* it adheres to the geometrical structure of surface reflectance. In addition, clamped BRDF leads to a reflectance kernel which does not obey the reciprocity property. One way to handle this is to multiply both sides of the image irradiance equation by $\cos \theta'_o$ [11], yet this inhibits modeling the image irradiance signal itself. As such, this work opt for dealing with the BRDF itself without incorporating the foreshortening factor, $\cos \theta'_i$, which will be included in later analysis when dealing with the irradiance equation.

4.4.1 Basis for Arbitrary Reflectance

Consider the product of two HSH basis functions to give a mapping $\Omega'_i \times \Omega'_o \rightarrow \mathbb{R}$ from the Cartesian product of two hemispheres to the real line, one can define a combined basis function $H_p^q(\vec{\omega}'_i)H_r^s(\vec{\omega}'_o)$. Although these functions construct a complete orthogonal basis for the cartesian product of two hemispheres, this work seeks a set of bases which span the subspace of functions characterized by maintaining the Helmholtz reciprocity property. Thus the term *Helmholtz HSH-based basis* is defined by symmetrizing the combined basis *w.r.t.* the incident and outgoing directions, *i.e.*

$$\mathcal{H}_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o) = H_p^q(\vec{\omega}'_i)H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i)H_p^q(\vec{\omega}'_o) \quad (4.7)$$

The orthogonality property of this basis is inherited from the orthonormality of the HSH basis in Equation 4.6, where,

$$\begin{aligned} & \langle \mathcal{H}_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o), \mathcal{H}_{p'r'}^{q's'}(\vec{\omega}'_i, \vec{\omega}'_o) \rangle \\ &= \int_{\Omega'_o} \int_{\Omega'_i} [H_p^q(\vec{\omega}'_i)H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i)H_p^q(\vec{\omega}'_o)] \\ & \quad [H_{p'}^{q'}(\vec{\omega}'_i)H_{r'}^{s'}(\vec{\omega}'_o) + H_{r'}^{s'}(\vec{\omega}'_i)H_{p'}^{q'}(\vec{\omega}'_o)] d\vec{\omega}'_i d\vec{\omega}'_o \\ &= \left(\int_{\Omega'_i} H_p^q(\vec{\omega}'_i)H_{p'}^{q'}(\vec{\omega}'_i) d\vec{\omega}'_i \right) \left(\int_{\Omega'_o} H_r^s(\vec{\omega}'_o)H_{r'}^{s'}(\vec{\omega}'_o) d\vec{\omega}'_o \right) \\ & \quad + \left(\int_{\Omega'_i} H_p^q(\vec{\omega}'_i)H_{r'}^{s'}(\vec{\omega}'_i) d\vec{\omega}'_i \right) \left(\int_{\Omega'_o} H_r^s(\vec{\omega}'_o)H_{p'}^{q'}(\vec{\omega}'_o) d\vec{\omega}'_o \right) \\ & \quad + \left(\int_{\Omega'_i} H_r^s(\vec{\omega}'_i)H_{p'}^{q'}(\vec{\omega}'_i) d\vec{\omega}'_i \right) \left(\int_{\Omega'_o} H_p^q(\vec{\omega}'_o)H_{r'}^{s'}(\vec{\omega}'_o) d\vec{\omega}'_o \right) \\ & \quad + \left(\int_{\Omega'_i} H_r^s(\vec{\omega}'_i)H_{r'}^{s'}(\vec{\omega}'_i) d\vec{\omega}'_i \right) \left(\int_{\Omega'_o} H_p^q(\vec{\omega}'_o)H_{p'}^{q'}(\vec{\omega}'_o) d\vec{\omega}'_o \right) \end{aligned}$$

$$= 2\delta_{pp'}\delta_{qq'}\delta_{rr'}\delta_{ss'} + 2\delta_{pr}\delta_{qs} \quad (4.8)$$

Hence, the *normalized* Helmholtz HSH-based basis can be defined as,

$$\mathcal{H}_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o) = N_{pr}^{qs} [H_p^q(\vec{\omega}'_i)H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i)H_p^q(\vec{\omega}'_o)] \quad (4.9)$$

where N_{pr}^{qs} is a normalization factor which guarantee the basis orthonormality. Using the orthogonality condition in Equation 4.8, the normalization factor can be defined as,

$$N_{pr}^{qs} = \frac{1}{\sqrt{2 + 2\delta_{pr}\delta_{qs}}} \quad (4.10)$$

Thus, an arbitrary surface BRDF can be represented in terms of the Helmholtz HSH-based basis as follows,

$$f_r(\vec{\omega}'_i, \vec{\omega}'_o) = \sum_{p=0}^{\infty} \sum_{r=0}^{\infty} \sum_{q=-p}^p \sum_{s=-r}^r a_{pr}^{qs} \mathcal{H}_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o) \quad (4.11)$$

In case of measured BRDFs, where scattered data are available, the Helmholtz reciprocity property might be violated. The process of projecting such scattered data on the subspace spanned by the proposed basis will filter out noisy components yielding the closest function which fits the data yet maintains the reciprocity property in the least-squares sense. The expansion coefficients of the series in Equation 4.11 can be obtained through projecting the surface BRDF on the Helmholtz HSH-based basis such that¹,

$$a_{pr}^{qs} = \langle f_r(\vec{\omega}'_i, \vec{\omega}'_o), \mathcal{H}_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o) \rangle = \int_{\Omega'_o} \int_{\Omega'_i} f_r(\vec{\omega}'_i, \vec{\omega}'_o) \mathcal{H}_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o) d\vec{\omega}'_i d\vec{\omega}'_o \quad (4.12)$$

By construction one has,

$$\int_{\Omega'_o} \int_{\Omega'_i} \mathcal{H}_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o) \mathcal{H}_{p'r'}^{q's'}(\vec{\omega}'_i, \vec{\omega}'_o) d\vec{\omega}'_i d\vec{\omega}'_o = \delta_{pp'}\delta_{rr'}\delta_{qq'}\delta_{ss'} \quad (4.13)$$

Hence, the BRDF energy content, which is defined as the integral $[f_r(\vec{\omega}'_i, \vec{\omega}'_o)]^2$ over its

¹Complex conjugate is dropped since this work deals with the real form of the basis.

entire domain¹, can be written as,

$$\begin{aligned}
 e_B &= \int_{\Omega'_o} \int_{\Omega'_i} [f_r(\vec{\omega}'_i, \vec{\omega}'_o)]^2 d\vec{\omega}'_i d\vec{\omega}'_o \\
 &= \sum_{prqs} \sum_{p'r'q's'} a_{pr}^{qs} a_{p'r'}^{q's'} \int_{\Omega'_o} \int_{\Omega'_i} \mathcal{H}_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o) \mathcal{H}_{p'r'}^{q's'}(\vec{\omega}'_i, \vec{\omega}'_o) d\vec{\omega}'_i d\vec{\omega}'_o \\
 &= \sum_{prqs} (a_{pr}^{qs})^2
 \end{aligned} \tag{4.14}$$

Thus the approximation accuracy (or the energy captured by the P th order approximation) can be defined as²,

$$Acc_B(P) = \frac{\sum_{p=0}^P \sum_{r=0}^P \sum_{q=-p}^p \sum_{s=-r}^r (a_{pr}^{qs})^2}{\int_{\Omega'_o} \int_{\Omega'_i} [f_r(\vec{\omega}'_i, \vec{\omega}'_o)]^2 d\vec{\omega}'_i d\vec{\omega}'_o} \tag{4.15}$$

One can observe that the approximation accuracy of surface BRDF looks little bit different from that of illumination. This is because in the latter case it is required to take into consideration all possible 3D rotations for the lighting function, hence the average energy content is of the main concern. This is in contrast to the former case where all function parameters (incoming and outgoing directions) are already included in the BRDF definition which in turn are defined *w.r.t.* the local coordinates, hence the approximation accuracy of the BRDF becomes independent of the hemispheres orientation.

The list of the coefficients $\{a_{pr}^{qs}\}$ represents the bidirectional surface reflectance spectrum (BSRS) which characterize the BRDF in a manner similar to the Fourier spectrum, where low-order spectrum maintains the overall properties of the underlying kernel, whereas increasing the order adds more details.

For scattered reflectance data, high-order spectrum components would add noise rather than details [21], thus truncation/smoothing is mostly used to improve scattered data description. This resulting harmonic expansion defines a smooth function which can be considered as a least squares approximation of the underlying BRDF. Figure 4.3 visualizes up-to 3rd order of the proposed basis at a fixed incident direction (more details about basis at fixed direction will be given in later sections).

¹Using the orthonormality property of the HSH-based basis, all terms in this integral vanishes except for $p = p', q = q', r = r'$ and $s = s'$.

²Note that the integral in the denominator can be evaluated using Monte Carlo integration.

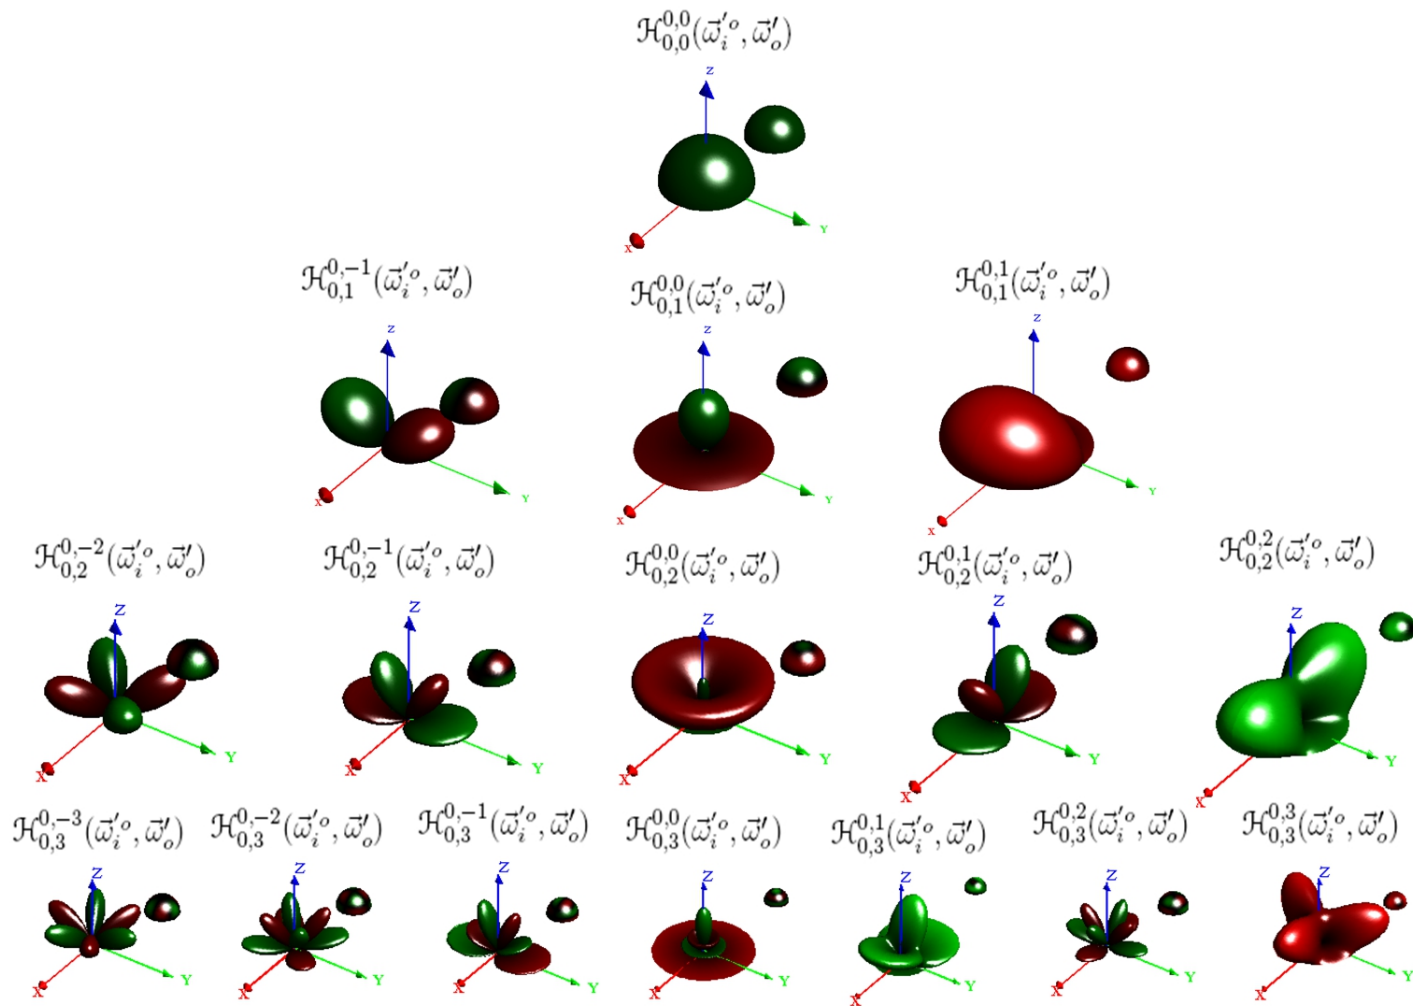


Figure 4.3: Visualization of up-to 3rd order of the proposed Helmholtz surface reflectance basis at an incident direction $\vec{\omega}'_o = (\pi/3, 0)$.

4.4.2 Reflectance Basis For Isotropic Surfaces

Isotropic BRDFs can be assumed when rotating the local tangent plane doesn't affect the surface reflectance properties, *i.e.* the BRDF becomes a function of only three variables $f_r(\vec{\omega}'_i, \vec{\omega}'_o) = f_r(\theta'_i, \phi'_i, \theta'_o, \phi'_o) = f_r(\theta'_i, \theta'_o, |\phi'_o - \phi'_i|)$. Thus depending only on $|\phi'_o - \phi'_i|$ allows negating both the incident and outgoing azimuthal angles without affecting the surface BRDF.

Furthermore, all terms in the harmonic expansion Equation 4.11 vanish except for those terms which satisfy the isotropy assumption, *i.e.* rotational invariant *w.r.t.* the azimuthal part (adding $\Delta\phi'$ to both ϕ'_i and ϕ'_o makes no difference), this requires $q = s$. Hence the expansion coefficients of the BRDF are now defined, using three indices, by $a_{pr}^{qs} = a_{pr}^{qq} = a_{pr}^{(-q)(-q)} = a_{pr}^q$. From isotropy it can be shown that,

$$\begin{aligned} \mathcal{H}_{pr}^{qs}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) &= \mathcal{H}_{pr}^{qq}(\theta'_i, \Delta\phi', \theta'_o, 0) \\ &= \mathcal{H}_{pr}^{qq}(\theta'_i, 0, \theta'_o, \Delta\phi') \\ &= \mathcal{H}_{pr}^q(\theta'_i, \theta'_o, \Delta\phi') \end{aligned} \quad (4.16)$$

where $\Delta\phi' = |\phi'_o - \phi'_i|$.

Therefore, the harmonic expansion of isotropic BRDFs can be defined as,

$$\begin{aligned} f_r(\vec{\omega}'_i, \vec{\omega}'_o) &= f_r(\theta'_i, \theta'_o, \Delta\phi') \\ &= \sum_{p=0}^{\infty} \sum_{r=0}^{\infty} \sum_{q=-\min(p,r)}^{\min(p,r)} a_{pr}^q \mathcal{H}_{pr}^q(\theta'_i, \theta'_o, \Delta\phi') \end{aligned} \quad (4.17)$$

where the basis for isotropic surfaces can be defined in terms of the polar and azimuthal parts of the HSH as,

$$\mathcal{H}_{pr}^q(\theta'_i, \theta'_o, \Delta\phi') = N_{pr}^q \left[\tilde{\Theta}_p^q(\theta'_i) \tilde{\Theta}_r^q(\theta'_o) + \tilde{\Theta}_r^q(\theta'_i) \tilde{\Theta}_p^q(\theta'_o) \right] \Phi_q(\Delta\phi') \quad (4.18)$$

where

$$N_{pr}^q = N_{pr}^{qq} \Phi_q(0) = \frac{1}{\sqrt{2 + 2\delta_{pr}}} \begin{cases} \sqrt{2} & q > 0 \\ 1 & q = 0 \\ 0 & q < 0 \end{cases} \quad (4.19)$$

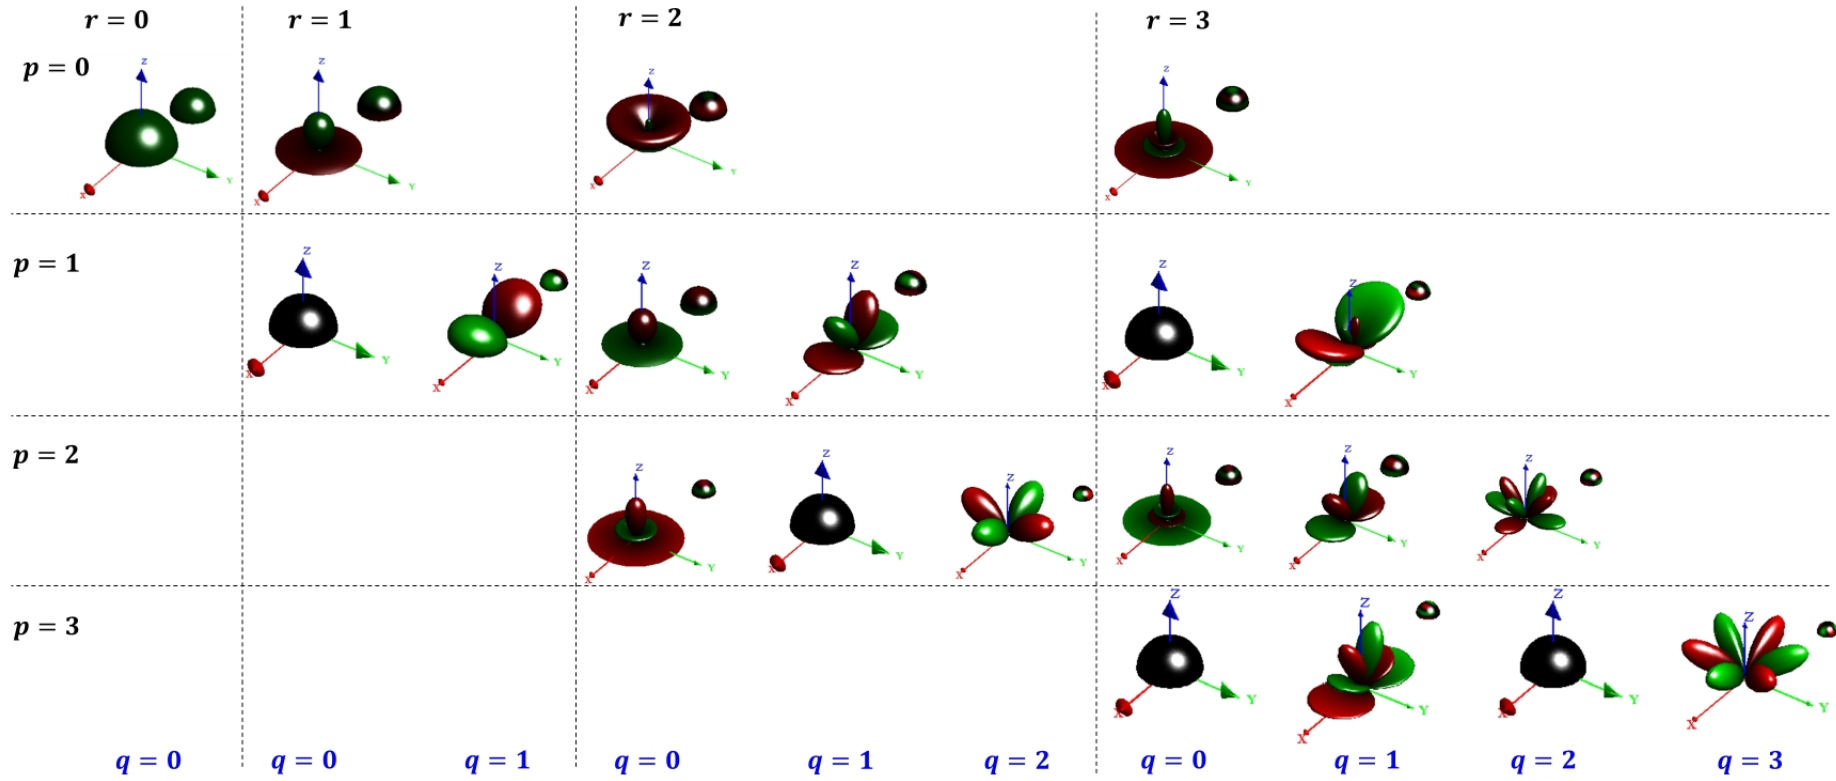


Figure 4.4: Visualization of up-to 3rd order of the proposed *isotropic* Helmholtz surface reflectance basis at an incident direction $\vec{\omega}_i^o : (\pi/3, 0)$. The order p runs from top to bottom while the associated order r runs from left to right. The azimuthal order q , or known as degree, also runs from left to right for each order r .

This reduced domain causes the number of independent spectrum components to be reduced drastically [21]. The normalization factor for negative degree basis is zero, i.e. $N_{pr}^q = 0 \quad \forall \quad q \in [-\min(p, r), -1]$. Moreover, due to the Helmholtz reciprocity, one has $\mathcal{H}_{pr}^q(\cdot) = \mathcal{H}_{rp}^q(\cdot)$. Thus the expansion in Equation 4.17 can be rewritten as,

$$\begin{aligned} f_r(\vec{\omega}'_i, \vec{\omega}'_o) &= f_r(\theta'_i, \theta'_o, \Delta\phi') \\ &= \sum_{p=0}^{\infty} \sum_{r=0}^p \sum_{q=0}^r a_{pr}^q \mathcal{H}_{pr}^q(\theta'_i, \theta'_o, \Delta\phi') \end{aligned} \quad (4.20)$$

Whereas the isotropic bidirectional surface reflectance spectrum (IBSRS) can be obtained by,

$$\begin{aligned} a_{pr}^q &= \langle f_r(\theta'_i, \theta'_o, \Delta\phi'), \mathcal{H}_{pr}^q(\theta'_i, \theta'_o, \Delta\phi') \rangle \\ &= \int_{\Omega'_o} \int_{\Omega'_i} f_r(\theta'_i, \theta'_o, \Delta\phi') \mathcal{H}_{pr}^q(\theta'_i, \theta'_o, \Delta\phi') d\vec{\omega}'_i d\vec{\omega}'_o \end{aligned} \quad (4.21)$$

Figure 4.4 shows a visualization for up-to 3rd order of the proposed Helmholtz basis in case of isotropic surface BRDF. It can be observed that higher orders are entertained with larger number of basis adding more details to the BRDF representation. This in contrast to the isotropic basis proposed by Koenderink *et al.* [21] which maintains only 8 basis at the 3rd order representation due to the definition of the Zernike polynomials.

Figure 4.5 shows the number of available orthonormal basis as a function of the reflectance order p (where $r = p$) for isotropic and anisotropic cases. One can find that the proposed basis definition provides more orthonormal basis functions when compared to that of Koenderink *et al.* [21]. This give insights that the proposed reflectance basis would capture similar energy content to that of the Zernike-based basis at lower orders of reflectance, leading to a more compact representation of surface BRDF. It is worth noting that the proposed basis share the same number of orthonormal basis with [19, 20, 22].

4.5 On Basis Complexity

The numerical computation of expansion coefficients in general and reflectance coefficients (*a.k.a.* spectrum) in particular is one of the important computational issues in many applied research areas, *e.g.* computer vision, medical imaging and statistical analysis. As such, this motivated much research to be conducted in developing fast

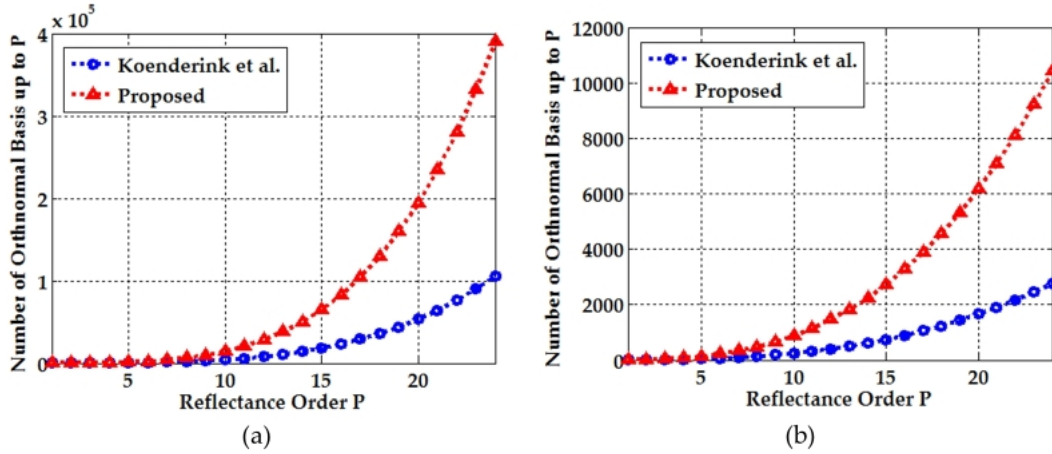


Figure 4.5: Number of orthonormal reflectance basis available up to reflectance order $p = r = P$ for (a) anisotropic and (b) isotropic cases. It can be noted that the proposed basis has more orthonormal bases at lower orders compared to that of Koenderink *et al.* [21] leading to a more compact representation for surface BRDF.

computational algorithms for computing expansions, *e.g.* [87, 88]. Since spherical and hemispherical reflectance bases share the same azimuthal functions, the core computational component of different reflectance basis in comparison is the polar part. Whereas the spherical-based bases depend on associated Legendre polynomials, *e.g.* [19, 81], hemispherical-based bases depend on either (shifted) associated Legendre polynomials, *e.g.* [20, 22] or Zernike radial polynomials [21]. Thus the problem of fast computation of reflectance spectrum reduces to efficient computation of the polar functions of the respective basis definition. As noted earlier, the core justification of the proposed basis is that the associated Legendre polynomials are known to have lower computational complexity when compared to Zernike ones. This section presents such a justification in more details.

Let $n \geq 0$ be the band index (*a.k.a.* order) dividing the associated Legendre polynomial set into bands of functions with $(n + 1)n$ polynomials for n -th band series. Using the standard arithmetic complexity model¹, Healy *et al.* in [87] and [43] provided algorithmic tools which improved the asymptotic complexity of the complete set of Legendre polynomials from $\mathcal{O}(N^{3/2})$ to $\mathcal{O}(N \log^2 N)$, for cutoff frequency $N = n^2$. Their algorithms were driven from the factorization of the Legendre functions obtained from their recurrence relations where a high-degree Legendre function was decomposed as a linear combination of lower-degree Legendre functions. Suda and Takami in [89] proposed a fast approximate algorithm for the associated Legendre transform by means

¹Single operation is defined as a complex multiplication followed by a complex addition.

of polynomial interpolation accelerated by the Fast Multipole Method (FMM). Their divide-and-conquer approach gives the same computational complexity of $\mathcal{O}(N \log^2 N)$.

On the other hand, the calculation of Zernike polynomials using the series definition of the radial polynomials, *i.e.* Direct method, requires $\mathcal{O}(N^3)$ for N -th order polynomial. Different approaches have been proposed to expedite the computation of Zernike polynomials including Belkasim's [90], Kintner's [91] and Prata's [92]. Belkasim's algorithm is based on factorizing some of the redundant terms in the radial and angular expansion of Zernike polynomials, however the total number of multiplications involved remains $\mathcal{O}(N^3)$. Kintner's method uses a recurrence relation which uses polynomials with varying lower orders and with fixed degree reducing the number of operations required to $\mathcal{O}(N^2)$. Prata, on the other hand, used a different recurrence relation which uses radial polynomials of lower orders, yet the complexity remains at $\mathcal{O}(N^2)$.

Therefore, one finds that bases which are based on (shifted) associated Legendre polynomials require less amount of CPU time when compared to the ones which are based on Zernike radial polynomials, especially for high orders.

4.6 Representation of Analytic Reflectance Models

Surfaces can be classified according to their reflectance properties. Sillion *et al.* [83] introduced the classification of reflectance types into ideal diffuse, ideal specular and directional diffuse components. Diffuse, a.k.a Lambertian/body, reflection occurs when the surface reflects the incident light rays equally in all directions regardless of the incident direction, *i.e.* it has a constant BRDF and thus can be represented by a single scalar value (*a.k.a.* albedo). see Figure 4.6(a). On the other hand, the specular reflection, also called interface reflectance, is caused by mirror-like reflection at the surface air interface, *i.e.* it has a Dirac delta function as its BRDF and thus can be represented by the ratio of the outgoing radiation in the specular direction to the incoming one (*a.k.a.* specular reflectance). Figure 4.6(b) demonstrates the ideal specular reflection. In practice, the specular reflection can be described by two overlapping components: the specular spike and specular lobe. The specular spike occurs in a very narrow range of solid angles around the perfect specular direction (mirror-like) reflection and it is more dominant in the smooth surfaces. The specular lobe arises from the scattering reflection caused by the surface roughness, see Figure 4.6(c). Most surface materials exhibit a




Figure 4.6: Different reflectance components: (a) Diffuse/Lambertian reflectance, BRDF is independent of outgoing direction. Hence the radiance leaving the surface is angle-independent. Thus a Lambertian surface will look equally bright from any direction. (b) and (c) Specular reflectance, ideally radiance arriving along a particular incoming direction can only leave the surface along the specular direction. However in practice radiation arriving in a particular direction leaves the surface in a small lobe of directions surrounding the specular direction. (d) Directional diffuse reflectance excludes the specular component. (e) The BRDF at a surface point can be approximated by a combination of a diffuse and a specular component. (f) A zoom-in view on the surface where the specular lobe arises from the scattering reflection caused by the surface roughness.

more elaborate directionality than these limiting cases [83]. As such directional diffuse reflectance is widely used to describe general reflectance while excluding the specular component, see Figure 4.6(d). Such type of reflectance depend primarily on surface roughness, wavelength and the material electrical properties [83]. One way to add the specular effect is superimposing the specular component and the (directional) diffuse one, see Figure 4.6(e).

The approximating orders of the proposed reflectance basis expansion of an arbitrary BRDF primarily depend on the analytic form of the BRDF at hand, where smooth BRDFs require fewer coefficients for accurate representation compared to non-smooth/complex ones. This section investigates the representation power of the proposed basis to approximate perfect diffuse and specular reflectance models. This analysis is also extended to investigate the spectrum of physical analytic models of surface reflectance such as Torrance-Sparrow specular reflection model [12] for rough surfaces. Take note that one is mainly interested in physical (*a.k.a.* theoretical) models which try to accurately simulate reflectance using physics laws, as opposed to empirical models which provide a simple analytic formulation designed to mimic some reflection phenomena, *e.g.* Phong [15] and Blinn [70] models.

4.6.1 Ideal Diffuse Reflectance

Ideal diffuse, i.e. Lambertian, surfaces have the property that their BRDF is independent of the incoming and outgoing directions. Formally, this model is given by a surface whose BRDF is independent of outgoing direction and, by the reciprocity principle, of incoming direction as well, i.e. $f_r(\vec{\omega}'_i, \vec{\omega}'_o) = f_r$. Hence the radiance leaving the surface is angle-independent. This implies that a Lambertian surface will look equally bright from any direction, taking into account that brightness is proportional to surface radiance.

For Lambertian surfaces, *radiosity*¹ is used to measure the radiance of the surface. The BRDF being independent of the incoming and outgoing directions; is commonly

¹In general, if the radiance leaving a surface is independent of the exit angle, *radiosity* is used to measure this leaving power, it is defined as the total power (energy per time) leaving a point \mathbf{x} on a surface per unit area on the surface. It has units Wm^{-2} . It is obtained by integrating the surface radiance leaving the surface at that point over the whole exit hemisphere Ω'_o , where the radiance is multiplied by the foreshortening factor to turn the foreshortened area into an actual area.

referred to as *albedo* ρ_d which is defined as follows; where d refer to diffuse.

$$\begin{aligned}\rho_d &= \int_{\Omega'_o} f_r(\vec{\omega}'_i, \vec{\omega}'_o) \cos \theta'_o d\vec{\omega}'_o \\ &= f_r \int_0^{2\pi} \int_0^{\pi/2} \cos \theta'_o \sin \theta'_o d\theta'_o d\phi'_o = \pi f_r\end{aligned}\quad (4.22)$$

Thus the BRDF of a Lambertian surface is related to surface albedo by a constant, *i.e.* ,

$$f_r^L(\vec{\omega}'_i, \vec{\omega}'_o) = \frac{\rho_d}{\pi} \quad (4.23)$$

Since the Lambertian BRDF has no azimuthal dependence, it can be represented in terms of zonal-basis where $q = 0$. The isotropic surface reflectance spectrum, *i.e.* the expansion coefficients, can be obtained as follows.

$$\begin{aligned}a_{pr}^L &= a_{pr}^0 = \int_{\Omega'_o} \int_{\Omega'_i} f_r^L(\theta'_i) \mathcal{H}_{pr}^0(\theta'_i, \theta'_o, \Delta\phi') d\vec{\omega}'_i d\vec{\omega}'_o \\ &= N_{pr}^0 \frac{\rho_d}{\pi} \int_{\Omega'_i} \int_{\Omega'_o} \left[\tilde{\Theta}_p^0(\theta'_i) \tilde{\Theta}_r^0(\theta'_o) + \tilde{\Theta}_r^0(\theta'_i) \tilde{\Theta}_p^0(\theta'_o) \right] \Phi_0(\Delta\phi') d\vec{\omega}'_i d\vec{\omega}'_o \\ &= N_{pr}^0 \frac{\rho_d}{\pi} \left[\int_{\Omega'_o} \tilde{\Theta}_r^0(\theta'_o) d\vec{\omega}'_o \int_{\Omega'_i} \tilde{\Theta}_p^0(\theta'_i) d\vec{\omega}'_i \right. \\ &\quad \left. + \int_{\Omega'_o} \tilde{\Theta}_p^0(\theta'_o) d\vec{\omega}'_o \int_{\Omega'_i} \tilde{\Theta}_r^0(\theta'_i) d\vec{\omega}'_i \right] \\ &= N_{pr}^0 \frac{\rho_d}{\pi} [\varsigma_r \varsigma_p + \varsigma_p \varsigma_r] = N_{pr}^0 \frac{2}{\pi} \rho_d \varsigma_r \varsigma_p\end{aligned}\quad (4.24)$$

where (for proof refer to [Appendix A: The Polar Integral of Shifted Associated Legendre Polynomials](#)),

$$\begin{aligned}\varsigma_p &= \int_{\Omega'_o} \tilde{\Theta}_p^0(\theta'_o) d\vec{\omega}'_o \triangleq \int_{\Omega'_i} \tilde{\Theta}_p^0(\theta'_i) d\vec{\omega}'_i \\ &= (-1)^p \sqrt{2\pi(2p+1)} \sum_{k=0}^p \frac{(-1)^k (p+k)!}{(k+1)(k!)^2(p-k)!} \\ &= \begin{cases} 2.5066 & p = 0 \\ 0 & p > 0 \end{cases}\end{aligned}\quad (4.25)$$

Similarly it can be shown that,

$$\varsigma_r = \int_{\Omega'_o} \tilde{\Theta}_r^0(\theta'_o) d\vec{\omega}'_o \triangleq \int_{\Omega'_i} \tilde{\Theta}_r^0(\theta'_i) d\vec{\omega}'_i$$

$$= \begin{cases} 2.5066 & r = 0 \\ 0 & r > 0 \end{cases} \quad (4.26)$$

Therefore, the only non-vanishing spectral component is at the zeroth-order, thus one zero-order basis is only needed to construct the Lambertian reflectance kernel, scaled by¹

$$a_{pr}^L = \begin{cases} \frac{\rho_d}{\pi} \varsigma_0^2 & p = 0, r = 0 \\ 0 & p > 1, 1 < r \leq p \end{cases} \quad (4.27)$$

Thus the DC component of the BRDF spectrum is simply the Lambertian reflectance function, hence the Lambertian model is the lowest-order approximation of any reflectance function. This complies with the conclusion presented in [21].

4.6.2 Ideal Specular Reflectance

An ideal specular surface behaves like an ideal mirror, where radiance arriving along a particular incoming direction can only leave the surface along the specular direction $\vec{\omega}'_s$. This direction is obtained by reflecting the incoming direction about the surface normal at the point of interest. This implies that the outgoing polar angle θ'_o equals the incident polar angle θ'_i and the incident and outgoing rays lie in a plane containing the surface normal. Therefore the surface radiance in the outgoing direction is the source radiance in the corresponding outgoing direction [93], *i.e.* the surface forms a virtual image of the light source².

$$L_r^S(\theta'_o, \phi'_o) = L_i(\theta'_o, \phi'_o + \pi) \quad (4.28)$$

Comparing to the definition of surface radiance in Equation 2.12, the specular BRDF can be written as [93],

$$f_r^S(\theta'_i, \phi'_i, \theta'_o, \phi'_o) = \frac{\delta(\theta'_i - \theta'_o)\delta(\phi'_i - \phi'_o + \pi)}{\sin \theta'_i \cos \theta'_i} \quad (4.29)$$

According to Koenderink *et al.* [21], a specular reflectance kernel defined in Equation 4.30 is used instead to define perfect mirror BRDF.

$$f_r^S(\theta'_i, \phi'_i, \theta'_o, \phi'_o) \approx \frac{\delta(\theta'_i - \theta'_o)\delta(\phi'_i - \phi'_o + \pi)}{\sin \theta'_i} \quad (4.30)$$

¹ $N_{00}^0 = 1/2$

²Normally, a portion of the incoming radiance is absorbed by the surface. An ideal specular surface would absorb the same fraction of the incoming radiance for every direction.

The spectral components (expansion coefficients) of the mirror-like reflectance kernel can be obtained as follows.

$$\begin{aligned}
 a_{pr}^{S,qs} &= \int_{\Omega'_o} \int_{\Omega'_i} f_r^S(\vec{\omega}'_i, \vec{\omega}'_o) \mathcal{H}_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o) d\vec{\omega}'_i d\vec{\omega}'_o \\
 &= \int_{\phi'_o=0}^{2\pi} \int_{\theta'_o=0}^{\pi/2} \int_{\phi'_i=0}^{2\pi} \int_{\theta'_i=0}^{\pi/2} \delta(\theta'_i - \theta'_o) \delta(\phi'_i - \phi'_o + \pi) \\
 &\quad \mathcal{H}_{pr}^{qs}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) \sin \theta'_o d\theta'_i d\phi'_i d\theta'_o d\phi'_o \quad (4.31)
 \end{aligned}$$

Since integration is performed over delta functions, the integral vanishes except for $\theta' = \theta'_i = \theta'_o$ and $\phi' = \phi'_i = \phi'_o - \pi$.

$$a_{pr}^{S,qs} = \int_{\phi'=0}^{2\pi} \int_{\theta'=0}^{\pi/2} \mathcal{H}_{pr}^{qs}(\theta', \phi', \theta', \phi' + \pi) \sin \theta' d\theta' d\phi' \quad (4.32)$$

Due to the separability of HSH basis functions in Equation 4.5, using Equation 4.9, the integral can be factorized into polar and azimuthal parts as follows.

$$\begin{aligned}
 a_{pr}^{S,qs} &= N_{pr}^{qs} \int_{\phi'=0}^{2\pi} \Phi_q(\phi') \Phi_s(\phi' + \pi) d\phi' \int_{\theta'=0}^{\pi/2} \tilde{\Theta}_p^q(\theta') \tilde{\Theta}_r^s(\theta') \sin \theta' d\theta' \\
 &\quad + N_{pr}^{qs} \int_{\phi'=0}^{2\pi} \Phi_s(\phi') \Phi_q(\phi' + \pi) d\phi' \int_{\theta'=0}^{\pi/2} \tilde{\Theta}_r^s(\theta') \tilde{\Theta}_p^q(\theta') \sin \theta' d\theta' \\
 &= N_{pr}^{qs} \int_{\theta'=0}^{\pi/2} \tilde{\Theta}_p^q(\theta') \tilde{\Theta}_r^s(\theta') \sin \theta' d\theta' \\
 &\quad \left[\int_{\phi'=0}^{2\pi} \Phi_q(\phi') \Phi_s(\phi' + \pi) d\phi' + \int_{\phi'=0}^{2\pi} \Phi_s(\phi') \Phi_q(\phi' + \pi) d\phi' \right] \quad (4.33)
 \end{aligned}$$

According to the definition of the azimuthal part of the HSH in Equation 3.12, the following relation can be inferred,

$$\Phi_q(\phi + \pi) = \begin{cases} \Phi_q(\phi) & q \text{ is even} \\ -\Phi_q(\phi) & q \text{ is odd} \end{cases} \quad (4.34)$$

Therefore, the azimuthal part of the integral in Equation 4.33 vanishes if q and s do not have the same parity, *i.e.* both are even or both are odd. Using the orthogonality of Φ_q functions in Equation 3.13 leads to,

$$\int_{\phi'=0}^{2\pi} \Phi_q(\phi') \Phi_s(\phi' + \pi) d\phi' + \int_{\phi'=0}^{2\pi} \Phi_s(\phi') \Phi_q(\phi' + \pi) d\phi'$$

$$\begin{aligned}
 &= \begin{cases} (-1)^s \int_0^{2\pi} \Phi_q(\phi') \Phi_s(\phi') d\phi' & q + s \text{ even} \\ 0 & q + s \text{ odd} \end{cases} \\
 &= \begin{cases} (-1)^s 2\pi \delta_{qs} & q + s \text{ even} \\ 0 & q + s \text{ odd} \end{cases} \quad (4.35)
 \end{aligned}$$

Thus the integral in Equation 4.33 can be written as,

$$a_{pr}^{S,qs} = \begin{cases} (-1)^s 2\pi \delta_{qs} N_{pr}^{qq} \int_0^{\pi/2} \tilde{\Theta}_p^q(\theta') \tilde{\Theta}_r^q(\theta') \sin \theta' d\theta' & q + s \text{ even} \\ 0 & q + s \text{ odd} \end{cases} \quad (4.36)$$

where

$$\begin{aligned}
 &\int_0^{\pi/2} \tilde{\Theta}_p^q(\theta') \tilde{\Theta}_r^q(\theta') \sin \theta' d\theta' \\
 &= \tilde{N}_p^{|q|} \tilde{N}_r^{|q|} \int_0^{\pi/2} \tilde{P}_p^{|q|}(\cos \theta') \tilde{P}_r^{|q|}(\cos \theta') \sin \theta' d\theta' \quad (4.37)
 \end{aligned}$$

Switching to the cartesian representation, one has $z = \cos \theta'$, thus $dz = -\sin \theta' d\theta'$, while the integration domain changes from $[0, \pi/2]$ to $[1, 0]$. Using the orthogonality relation of the shifted associated Legendre polynomials in Equation 4.2 and the definition of the normalization factor $\tilde{N}_r^{|q|}$ in Equation 4.4 yield,

$$\begin{aligned}
 &\int_0^{\pi/2} \tilde{\Theta}_p^q(\theta') \tilde{\Theta}_r^q(\theta') \sin \theta' d\theta' \\
 &= \left(\tilde{N}_r^{|q|} \right)^2 \frac{(r + |q|)!}{(2r + 1)(r - |q|)!} \delta_{pr} = \frac{1}{2\pi} \delta_{pr} \quad (4.38)
 \end{aligned}$$

Therefore, the spectral components of the mirror-like reflectance kernel can be written as,

$$a_{pr}^{S,qs} = \begin{cases} \frac{(-1)^q}{2} \delta_{pr} \delta_{qs} & q + s \text{ even} \\ 0 & q + s \text{ odd} \end{cases} \quad (4.39)$$

We can note that all non-vanishing spectrum components have the same absolute value, *i.e.* the specular reflectance kernel has a flat spectrum as an analogy to the Fourier spectrum of an impulse. The simplicity of the result originates from the assumption of ideal specular surface, however, such an expression becomes more complicated when taking into consideration the off-specular reflection for rough surfaces [12].

Figure 4.7(a) plots the approximations of the ideal specular reflectance kernel obtained by truncating the infinite series in Equation 4.11 at orders $P = 0, 1, 2, 4, 6$ and 8. For comparison purposes, Figure 4.7(b) shows the approximation of the same kernel using the Zernike-based basis proposed by [21]¹. In both case, the lowest order approximation leads to the Lambertian kernel, *i.e.* perfectly diffuse, where the reflected radiance is isotropically distributed and no specularity is observed. With higher orders, the radiance tends to reflect in the general direction of specularity, with narrower lobe.

It can be observed from Figure 4.7 that with the HSH-based basis, the reflected radiance becomes more concentrated towards the specularity direction at lower orders when compared to the Zernike-based basis of [21], while being less sensitive to the abrupt truncation of the series causing less ringing effect. Further the reflected radiance have smaller spurious lobes in directions other than that of the perfect specular direction which vanish faster than that of the Zernike-based basis at higher orders.

Figure 4.8 shows the reconstruction of the ideal specular lobe at different incident directions while truncating the approximating series at the 8th order. It is evident that the HSH-based basis provide a more compact approximation with lower ringing effect when compared to the Zernike-based basis.

It is important to note that, as in Fourier synthesis, the width of the reconstructed specular lobe primarily depend on the order at which the series is truncated where an infinite number of terms causes the reconstruction of the perfect specular spike. This leads to a relation between width and order.

Ideal specular surfaces rarely exist in the real world. Usually, radiation arriving in a particular direction leaves the surface in a small lobe of directions surrounding the specular direction, where the incoming radiance is shared over all outgoing directions within this lobe, see Figure 4.6(c). This appears as a bright blob, also called specularity, along the specular direction. Phong model [15] is commonly used to model the shape of the specular lobe, where the radiance leaving a specular surface is proportional to $\cos^l(\theta'_o - \theta'_i)$, where θ'_o is the exit polar angle, θ'_i is the specular polar angle which is the same as the incident polar angle. l is a parameter defining the width of the specular lobe, where larger values of l defines narrower lobes and sharper specularities.

¹Refer to [21] for the expansion coefficients of the reflectance function.

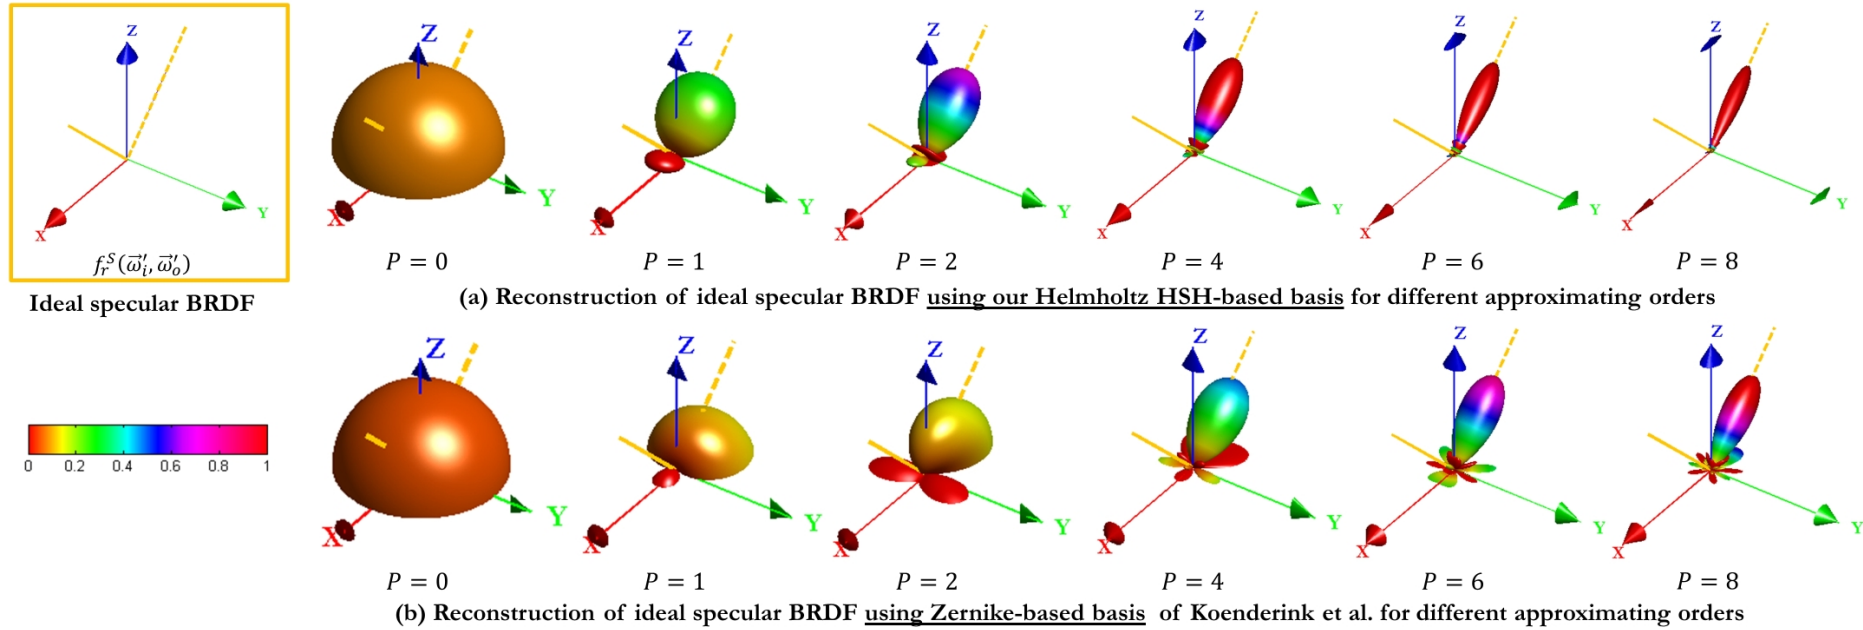


Figure 4.7: The reconstruction of the ideal specular reflectance kernel, where the light ray (in solid orange) is incident from direction $\theta_i' = \pi/4$ and $\phi_i' = 0$ and the perfect specular direction is plotted in dashed-orange. The approximating series has been truncated at different orders P , using (a) the proposed HSH-based basis, versus using (b) Zernike-based basis of [21]. In both cases, the lowest order entails the perfect diffuse/Lambertian reflectance while the specular lobe begins to be apparent for orders 1 and higher. Notice the ringing effect due to the abrupt truncation of the approximating series, however the HSH-based basis has lower ringing effect compared to the Zernike-based basis having smaller spurious lobes in directions other than that of the perfect specular direction, vanishing faster than that of the Zernike-based basis at higher orders.

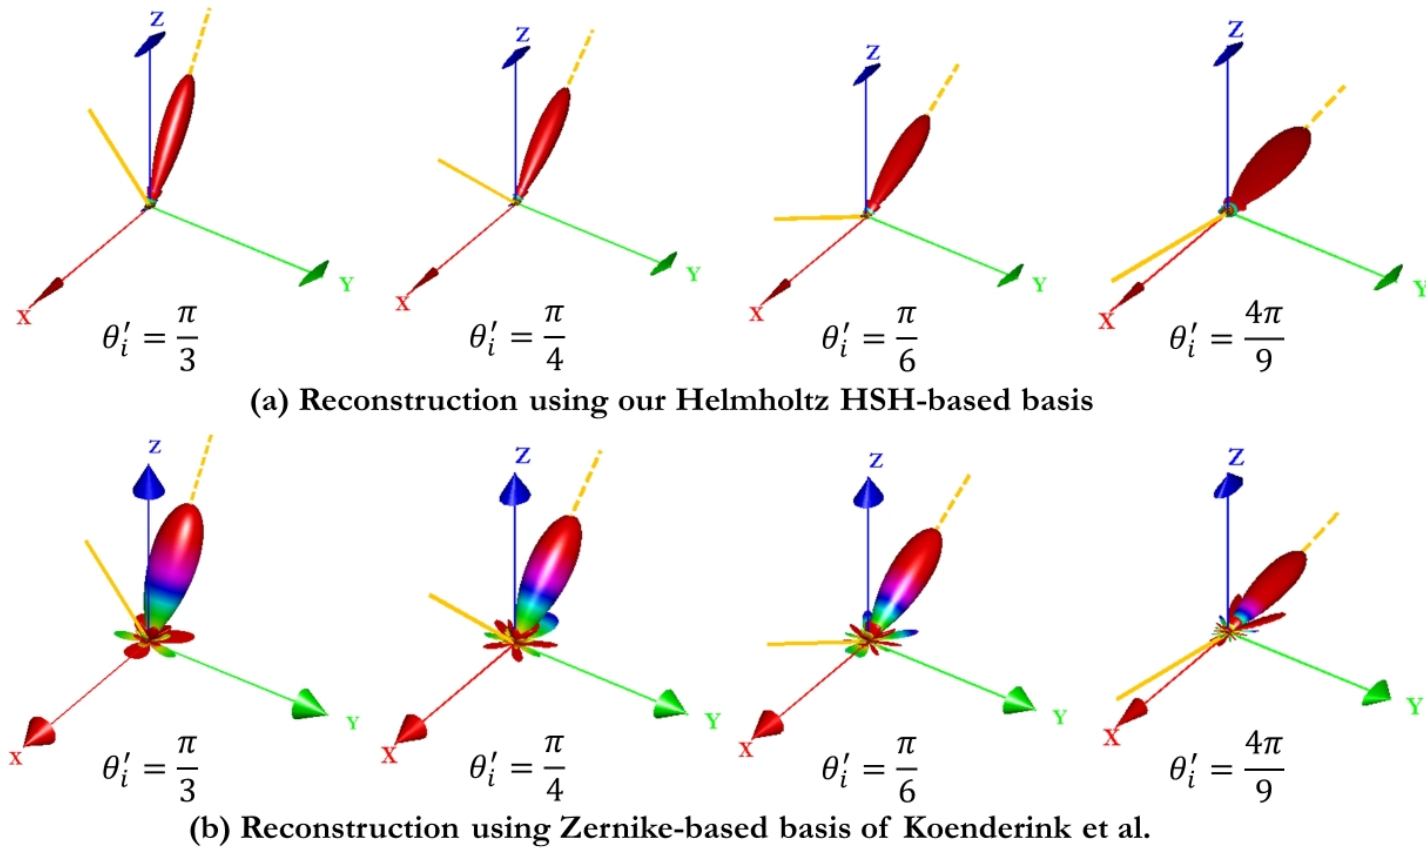


Figure 4.8: The reconstruction of the ideal specular reflectance kernel, where the light ray (in solid orange) is incident from different directions and the perfect specular direction is plotted in dashed-orange. The approximating series has been truncated at order $P = 8$, using (a) the proposed HSH-based basis, versus using (b) Zernike-based basis of [21]. It can be noted that the HSH-based basis provides a compact approximation for the specular lobe with lower ringing effect when compared to the Zernike-based basis.

4.6.3 Example of Non-Ideal Physical Reflectance Models

Pure diffuse or specular surfaces are non-existent. Many surfaces can be approximated by having a BRDF consisting of a diffuse/Lambertian component and a specular component, see Figure 4.6(e). Diffuse albedo and specular albedo are used to weight these components respectively in this combination.

Despite the simplicity of the Lambertian model, Oren and Nayar [14] had proved that such a model provides an inaccurate approximation to the diffuse component of the surface reflectance. This was accomplished through a set of experiments carried out on real samples, such as plaster, clay, sand and cloth, demonstrating significant deviation from Lambertian behavior.

The surface roughness model by Torrance and Sparrow [12] is built on the assumption that the surface is composed of a collection of long symmetric V-cavities, see Figure 4.6(f). Each V-cavity has two opposing facets with the width of each facet is much smaller than its length. The roughness of the surface is specified using probability function for the orientations of the facets. In order to use geometric optics, the area of each facet is assumed to be much larger than the wavelength of the incident light beam, and at the same time much smaller than the area of the surface patch being projected onto one pixel, thus the facets covered by one pixel can be described by statistical distributions.

The physically-based micro-facet model proposed by Torrance and Sparrow [12] is focused on glossy reflectance, where the V-cavity geometry implies that only facets facing in direction of the halfway vector $\vec{\omega}'_h = (\theta'_h, \phi'_h)$ affect the BRDF, where the reflection from each facet is described by the Fresnel equation $F(\cos \theta'_h; \eta)$ where η is the classical refractive index from geometric optics. The fraction of the facets oriented in the direction of $\vec{\omega}'_h$ is described by the facet slope distribution function $D(\theta'_h; \sigma)$, where σ is the root mean square slope of the micro-facets describing the surface's roughness. Whereas masking and shadowing of micro-facets are included in the geometric attenuation factor $G(\vec{\omega}'_i, \vec{\omega}'_o)$. Thus the Torrance-Sparrow (TS) BRDF is given by [12],

$$f_r^{TS}(\vec{\omega}'_i, \vec{\omega}'_o) = \frac{\rho_d}{\pi} + \frac{\rho_s}{\pi} \frac{F(\cos \theta'_h; \eta) D(\theta'_h; \sigma) G(\vec{\omega}'_i, \vec{\omega}'_o)}{\cos \theta'_i \cos \theta'_o} \quad (4.40)$$

where ρ_d is the diffuse albedo, ρ_s is the specular albedo and the half-way vector is given

by,

$$\vec{\omega}'_h = \frac{\vec{\omega}'_i + \vec{\omega}'_o}{\|\vec{\omega}'_i + \vec{\omega}'_o\|} \quad (4.41)$$

An approximation for Fresnel formula for dielectrics used by Cook and Torrance [13] can be written as,

$$F(c; n) = \frac{1}{2} \frac{(g - c)^2}{(g + c)^2} \left(1 + \frac{(c(g + c) - 1)^2}{(c(g - c) + 1)^2} \right) \quad (4.42)$$

with $g^2 = n^2 + c^2 - 1$.

The distribution function of TS-BRDF is used to model the facet slope distribution which is based on the physical theory of scattering electromagnetic waves, the distribution function is given by,

$$D(\theta'_h; \sigma) = \frac{1}{\sigma^2} \exp \left[- \left(\frac{\theta'_h}{\sigma} \right)^2 \right] \quad (4.43)$$

The geometric attenuation factor G ranges from 0 (total shadowing) to 1 (no shadowing), defining how the incoming light beam interact with the surface locally; the entire beam can reflect, a portion of the outgoing beam can be blocked (masking), or a portion of the incoming beam can be blocked (shadowing). Thus the geometric attenuation factor is defined as,

$$G(\vec{\omega}'_i, \vec{\omega}'_o) = \min \left(1, \frac{2 \cos \theta'_h \cos \theta'_o}{\vec{\omega}'_o \cdot \vec{\omega}'_h}, \frac{2 \cos \theta'_h \cos \theta'_i}{\vec{\omega}'_o \cdot \vec{\omega}'_h} \right) \quad (4.44)$$

For analytical simplicity, a simplified 4-parameter TS model is widely used in computer graphics, where the aforementioned model is modified in the following aspects while maintaining the physical properties of surface reflectance [7]. The Fresnel for a refractive index η is normalized to be 1 at normal exitance. While F depends on the angle *w.r.t.* the half-way vector; in practice, this angle is very close to θ'_o , thus the Fresnel term becomes $F(\cos \theta'_o; \eta) / F(1; \eta)$. Ashikhmin [94] believed that the distribution function has much greater impact on surface appearance than the geometric attenuation term, as such the geometric term can be omitted for simplicity [7]. This also comply with the distribution-based BRDF [67] where the shape of the reflection is dominated by the distribution function. Assuming, without loss of generality, that the viewer/camera is located at a distance relatively large compared to the object size such that the viewing direction coincides with the z -axis of the global reference frame; normal-exitance (*i.e.* $\theta'_o = 0$) can

be used to approximate the imaging process. Thus the half-way angle can be written as $\theta'_h = \theta'_i/2$ [7] since there is no azimuthal dependance in Equation 4.40. Hence the modified TS-BRDF can be written as,

$$f_r^{TS}(\vec{\omega}'_i, \vec{\omega}'_o) = \frac{\rho_d}{\pi} + \frac{\rho_s}{\pi\sigma^2 \cos \theta'_i \cos \theta'_o} \times \frac{F(\cos \theta'_o; \eta)}{F(1; \eta)} \exp \left[- \left(\frac{\theta'_i}{2\sigma} \right)^2 \right] \quad (4.45)$$

Since there is no azimuthal dependance in Equation 4.45, *i.e.* TS is radially symmetric, it can be expanded using zonal basis of the isotropic Helmholtz HSH-based basis of \mathcal{H}_{pr}^q , centered at surface normal $\vec{n}(\mathbf{x})$, where q is set to zero. The zonal basis can be written as,

$$\mathcal{H}_{pr}^0(\vec{\omega}'_i, \vec{\omega}'_o) = \frac{1}{2\pi} \sqrt{\frac{(2p+1)(2r+1)}{2+2\delta_{pr}}} \times \left[\tilde{P}_p(\cos \theta'_i) \tilde{P}_r(\cos \theta'_o) + \tilde{P}_r(\cos \theta'_i) \tilde{P}_p(\cos \theta'_o) \right] \quad (4.46)$$

where \tilde{P}_p is shifted associated Legendre polynomials [20] with order $p \geq 0$. Thus TS-BRDF can be expanded in the subspace spanned by the isotropic Helmholtz HSH-based basis as,

$$f_r^{TS}(\vec{\omega}'_i, \vec{\omega}'_o) = \sum_{p=0}^{\infty} \sum_{r=0}^{\infty} a_{pr} \mathcal{H}_{pr}^0(\vec{\omega}'_i, \vec{\omega}'_o) \quad (4.47)$$

The spectrum coefficients of the diffuse component of TS model can be derived using the orthogonality relation of the shifted associated Legendre polynomials, while their series representation can be used for the spectrum of the specular component. The spectrum coefficients can be obtained as (for proof refer to [Appendix B: Torrance Sparrow Spectrum Coefficients](#)),

$$\begin{aligned} a_{pr}(\theta'_o) &= \int_{\Omega'_o} \int_{\Omega'_i} f_r^{TS}(\vec{\omega}'_i, \vec{\omega}'_o) \mathcal{H}_{pr}^0(\vec{\omega}'_i, \vec{\omega}'_o) d\vec{\omega}'_i d\vec{\omega}'_o \\ &= 2\rho_d \delta_{p0} \delta_{r0} + \frac{\rho_s}{\pi\sigma^2 \cos \theta'_o} \frac{F(\cos \theta'_o; \eta)}{F(1; \eta)} B_{pr}(\theta'_o, \sigma) \end{aligned} \quad (4.48)$$

where ,

$$B_{pr}(\theta'_o, \sigma) = 2\pi \sqrt{\frac{(2p+1)(2r+1)}{2+2\delta_{pr}}} \left[\frac{\delta_{r0}}{2r+1} A_p(\theta'_o, \sigma) + \frac{\delta_{p0}}{2p+1} A_r(\theta'_o, \sigma) \right] \quad (4.49)$$

with

$$A_p(\theta'_o, \sigma) = \int_0^{\pi/2} \frac{\sin \theta'_i}{\cos \theta'_i} \exp \left[-\frac{\theta'^2_i}{4\sigma^2} \right] \tilde{P}_p(\cos \theta'_i) d\theta'_i \quad (4.50)$$

This means that $a_{p0} = a_{0p}$, thus a single index for the spectrum coefficients can be used where,

$$B_p(\theta'_o, \sigma) = \begin{cases} 2\pi A_0(\theta'_o, \sigma) & p = 0 \\ 2\pi \sqrt{\frac{2p+1}{2}} A_p(\theta'_o, \sigma) & p > 0 \end{cases} \quad (4.51)$$

Hence TS-BRDF can be represented as,

$$f_r^{TS}(\vec{\omega}'_i, \vec{\omega}'_o) = a_0(\theta'_o) + 2 \sum_{p=1}^{\infty} a_p(\theta'_o) \mathcal{H}_{p0}^0(\vec{\omega}'_i, \vec{\omega}'_o) \quad (4.52)$$

Note that the spectrum coefficients depends on the 4-parameters of TS model. Using Merl skin reflectance BRDF [95] which provide TS parameters of six different skin types (from very white to black), the harmonic spectrum is investigated in case of normal-exitance of each individual sample of skin BRDFs where Figure 4.9 shows the graphical representation of the first thirteen coefficients (up to the 12th order) averaged over skin types. The approximation accuracy, *i.e.* cumulative energy, of the p th order expansion as a function of the harmonic order is shown. It can be observed that the spectrum of TS model decays along the harmonic order. While the approximation accuracy exceeds 99% using the 5th order approximation. It is worthwhile to note that more spectrum coefficients are needed due to the specular component of TS model, while the diffuse component is fully encoded in the zeroth order approximation.

Using the skin BRDF parameters of Merl Skin BRDF database[95], Figure 4.10 compares the average approximation accuracy of Torrance Sparrow reflection model under distant illumination using the proposed isotropic basis in comparison to bases of Westin *et al.* [19], Gautron *et al.* [20], Koenderink *et al.* [21], Habel and Wimmer [22] and the Helmholtz basis of Koenderink *et al.* [21] where their isotropic version is used. The average is taken over 100 samples for each skin type at different facial regions while the BRDF spectrum is obtained using Monte Carlo integration. While the proposed basis provide a comparable accuracy to the Helmholtz basis of Koenderink *et al.* [21], one can observe that the former basis shows higher approximation accuracy especially at lower truncating reflectance orders compared to other bases. It is important to note

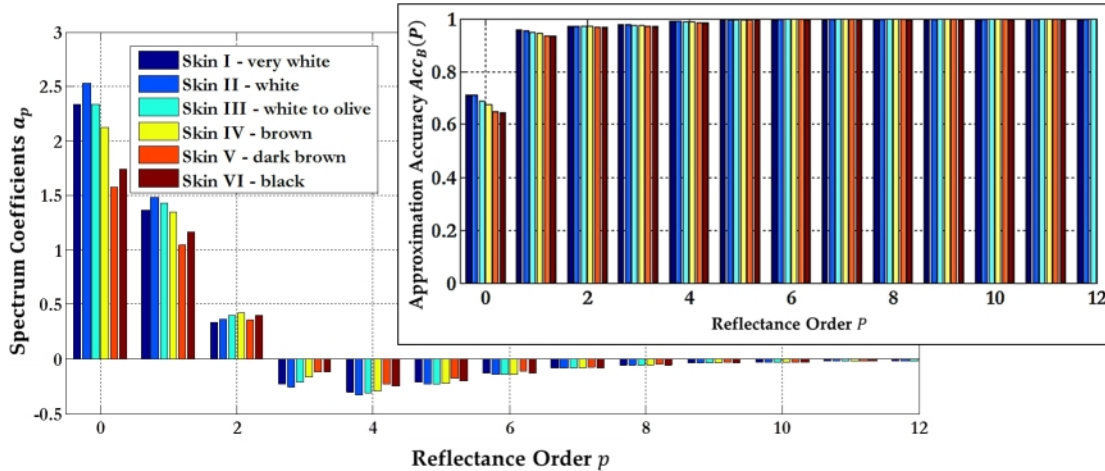


Figure 4.9: A graphical representation of the first thirteen harmonic expansion coefficients of Torrance Sparrow reflection model under distant illumination using skin BRDF parameters of Merl Skin BRDF database[95]. Note that TS acts as a low pass filter, where the 5th order expansion is sufficient to encode more than 99% of the BRDF energy content regardless of the skin type.

that the spherical basis Westin *et al.* [19] shows the lowest accuracy compared to the hemispherical ones, this justifies the need of a representation which complies with the geometric structure of reflectance functions. As such, the proposed basis provides a compact BRDF representation requiring fewer coefficients to accurately model surface BRDFs.

4.7 Modeling Scattered Reflectance Data

Scattered BRDF data might violate the Helmholtz reciprocity property; this can be filtered out through the process of projecting them onto the subspace spanned by the HSH-based basis, where the reciprocity property is preserved in the least-squares sense. Furthermore, in many practical cases, reflectance data are only available for plane-of-incidence geometries, where the incident and outgoing directions are coplanar with the surface normal. The reflectance spectrum components provide a phenomenological extrapolation from the available data in a unique manner [21]. This is used implicitly when assuming a surface to be Lambertian with diffuse albedo obtained from few or even single measurements; this is equivalent to taking the zero-order approximation of the reflectance kernel. Thus using the reflectance spectrum components can be considered as a way of refining the reflectance kernel representation beyond the zero-order.

This section further evaluates the accuracy and compactness of the HSH-based basis

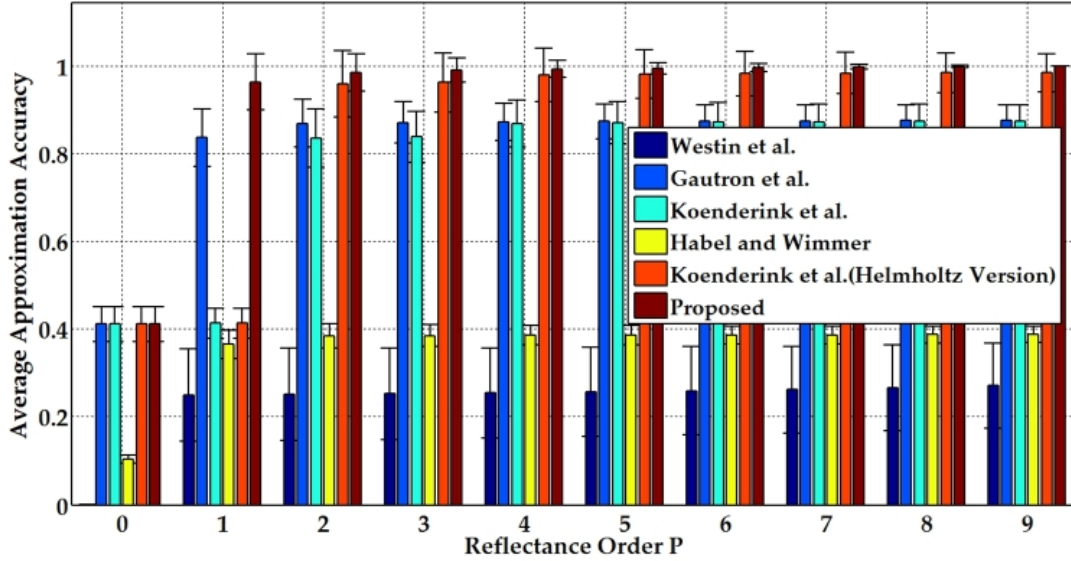


Figure 4.10: The average approximation accuracy of Torrance Sparrow reflection model under distant illumination as a function of the truncating reflectance order P . The skin BRDF parameters of Merl Skin BRDF database [95] are used. The average is taken over 100 samples for each skin type at different facial regions. The BRDF spectrum is obtained by projecting the TS function using Monte Carlo integration onto the subspace spanned by the proposed isotropic basis in comparison to bases of Westin *et al.* [19], Gautron *et al.* [20], Koenderink *et al.* [21], Habel and Wimmer [22] and the Helmholtz basis of Koenderink *et al.* [21] where their isotropic version is used. Note that the proposed basis provide higher approximation accuracy at lower reflectance orders compared to other bases, hence the proposed representation is capable of providing a compact representation of reflectance functions.

using BRDF measurements which are directly measured from real surfaces. This work makes use of BRDF databases which are available, free of charge, for academic purpose. In this dissertation, two BRDF databases are used for reflectance modeling. The first one is the database provided by Mitsubishi Electric Research Laboratories (Merl) [17] containing isotropic materials. The second one is offered by Columbia University and Utrecht University named as CURET database [10] containing a mix of isotropic and anisotropic materials.

4.7.1 Experimentation on Isotropic-Merl Reflectance Data

Merl reflectance data [17] is based on 100 isotropic materials to represent a wide variety of surface materials with different diffuse and specular reflection properties. Based on uniform spacing, 1,458,000 BRDF measurements are provided in 3D angular space using half-angle parameterization of Rusinkiewicz [96]. These BRDF samples are used as a lookup table since they are dense enough.

For each Merl material, Figure 4.11 compares between the approximation accuracies

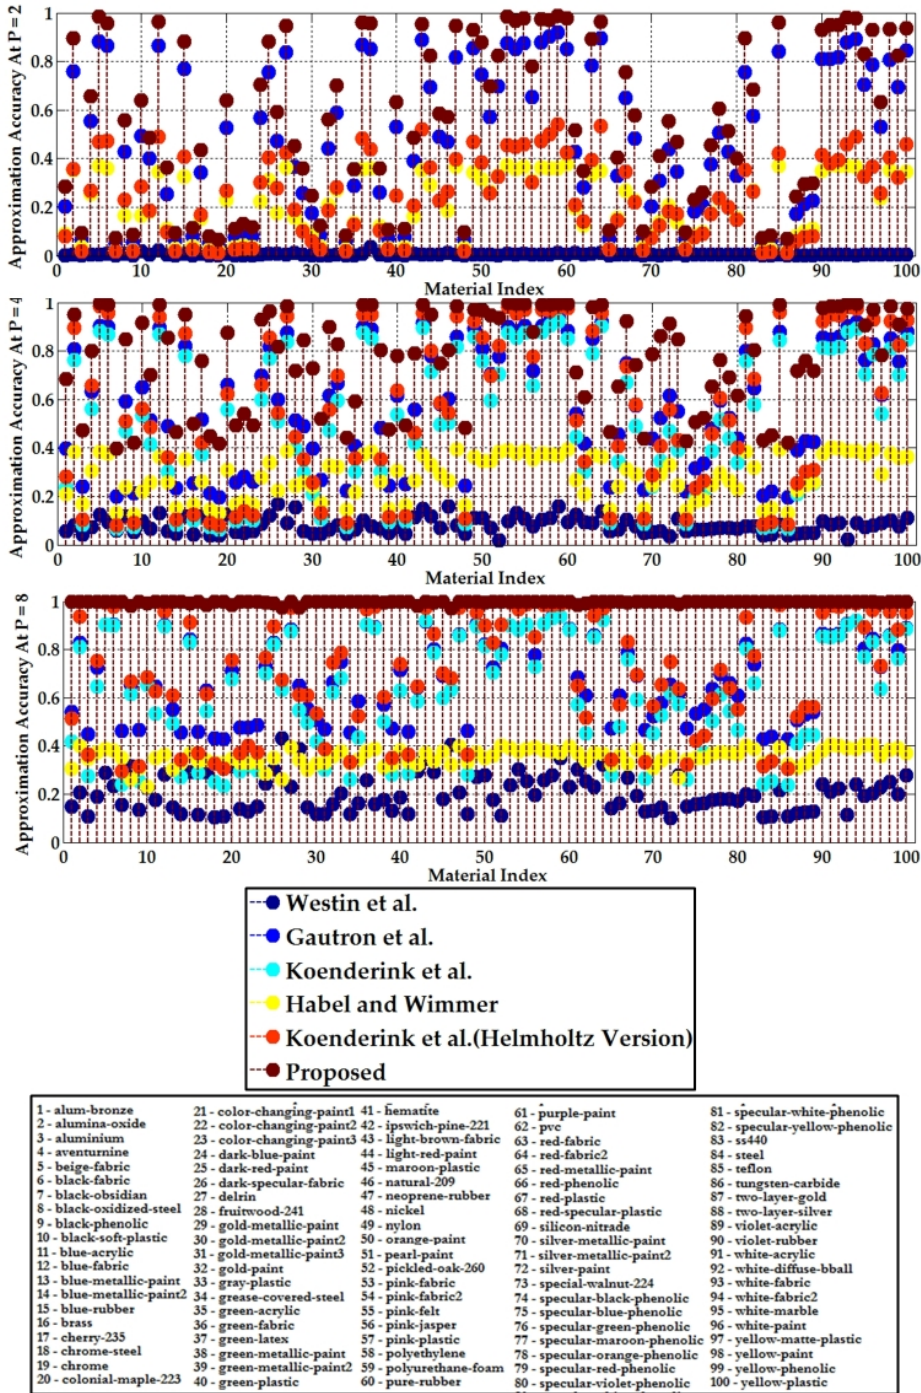


Figure 4.11: The approximation accuracy of Merl database materials [17] for different truncating reflectance orders P . The reflectance spectrum for each material is obtained by projecting randomly drawn Merl BRDF measurements using Monte Carlo integration onto the subspace spanned by the proposed *isotropic* basis in comparison to bases of Westin *et al.* [19], Gautron *et al.* [20], Koenderink *et al.* [21], Habel and Wimmer [22] and the Helmholtz basis of Koenderink *et al.* [21] where their isotropic version is used. The sampling points are drawn from the cartesian product of the incoming and outgoing (hemi)spheres (according to the basis definition). It can be observed that the proposed basis provides the highest approximation accuracy for most of the Merl materials especially at reflectance order $P = 8$.

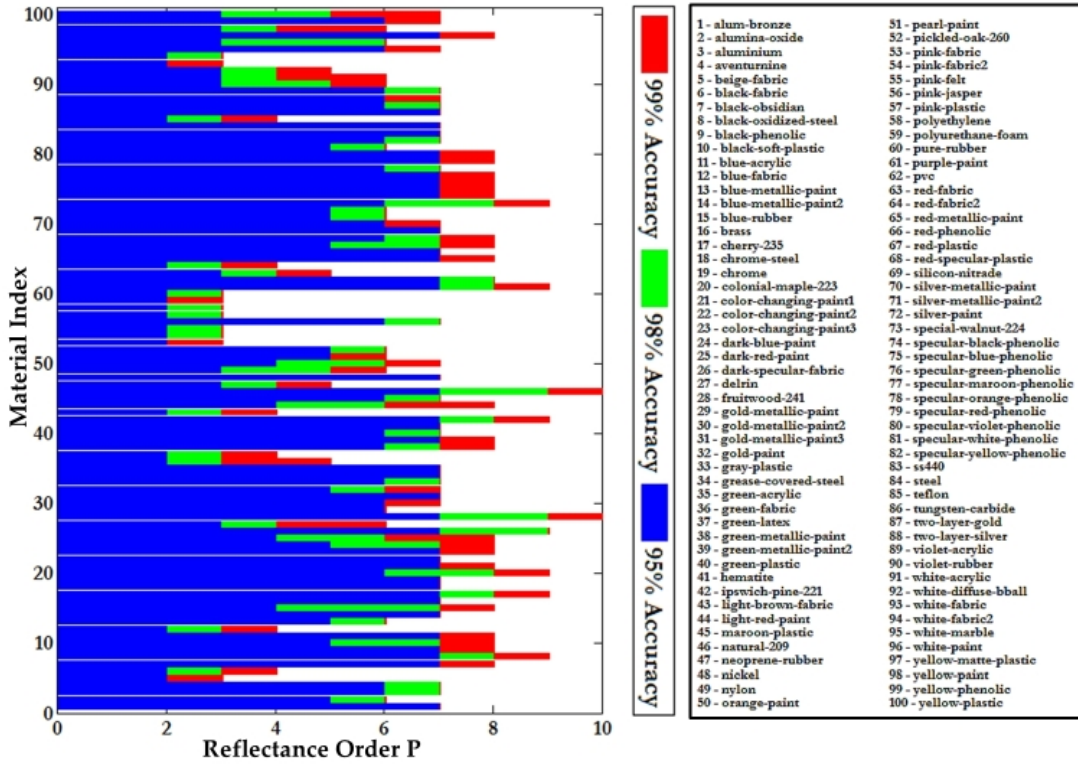


Figure 4.12: Required minimum reflectance order for the materials in the Merl database [17] when using the proposed *isotropic* basis. The blue, green and red bars show the required reflectance order P to achieve approximation accuracy of 95%, 98% and 99%, respectively.

provided by different *isotropic* reflectance bases in contrast to that of the proposed *isotropic* basis with different truncating reflectance orders P . Some materials show extreme specular properties such as chrome-steel where it displays low approximation accuracies in Figure 4.11 at lower reflectance orders regardless of the reflectance basis. In contrast, diffuse materials such as beige-fabric shows high approximation accuracies at lower reflectance orders. Nonetheless, in general, one can observe how the proposed basis provide the highest approximation accuracy for most of Merl materials especially at $P = 8$, while the spherical basis shows the worst accuracy for BRDF representation. These observations hold for both diffuse and specular materials.

Figure 4.12 shows the required minimum truncating reflectance order to achieve specific level of accuracy when using the proposed *isotropic* basis. It can be noted that all Merl materials can be represented with accuracy of 95% with reflectance order less than or equal $P = 7$, while diffuse materials such as fabric materials can be represented with accuracy of 99% with reflectance order less than or equal $P = 4$.

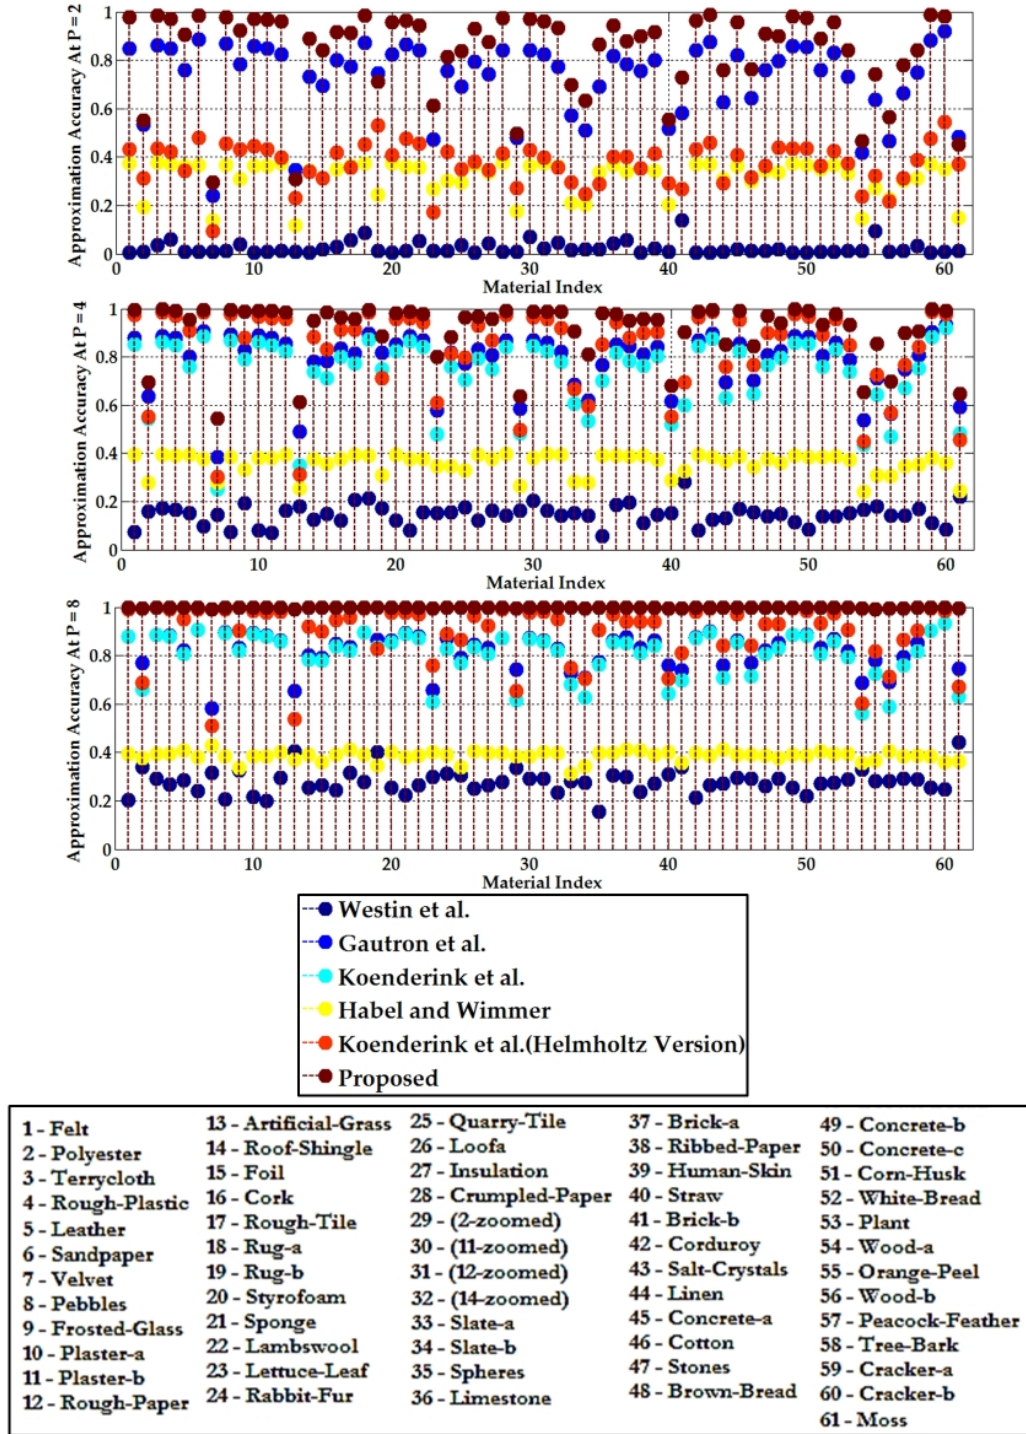


Figure 4.13: The approximation accuracy of CURET database materials [10] for different truncating reflectance orders P . The reflectance spectrum for each material is obtained by projecting randomly drawn CURET BRDF measurements using Monte Carlo integration onto the subspace spanned by the proposed basis in comparison to bases of Westin *et al.* [19], Gautron *et al.* [20], Koenderink *et al.* [21], Habel and Wimmer [22] and the Helmholtz basis of Koenderink *et al.* [21]. Note that BRDF sparse measurements are interpolated at the drawn samples using the provided fitted measurements. The sampling points are drawn from the cartesian product of the incoming and outgoing (hemi)spheres (according to the basis definition). It can be observed that the proposed basis provides the highest approximation accuracy for all CURET, while the spherical basis [19] provides the lowest accuracy compared to others.

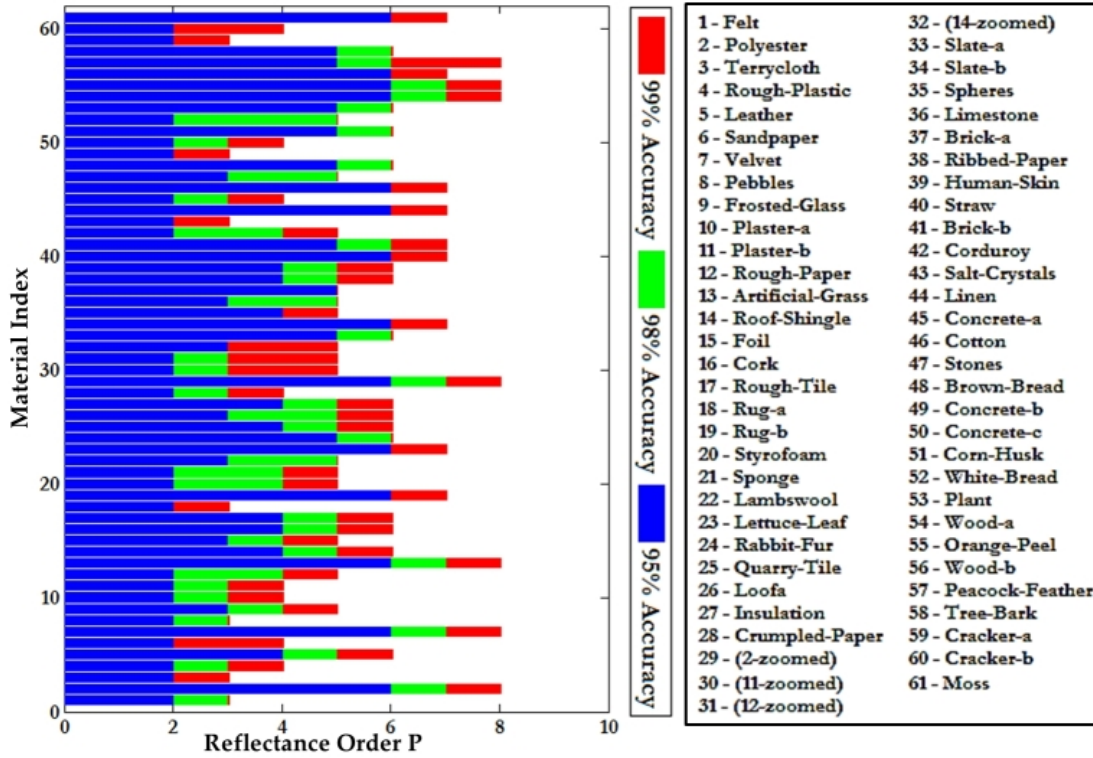


Figure 4.14: Required minimum reflectance order for the materials in the CURET database [10] when using the proposed basis. The blue, green and red bars show the required reflectance order P to achieve approximation accuracy of 95%, 98% and 99%, respectively.

4.7.2 Experimentation on CURET Reflectance Data

The CURET database [10] consists of 61 BRDFs with sparse measurements of 205 measurements each over varying incident and outgoing directions. It also offers a BRDF parameter database which fits the sparse measurements to the Koenderink reflectance model [21]. In the presented experimentation, the parameter database is used since the samples are not dense enough to be used directly as a BRDF lookup table. One opts for using the anisotropic parameters since this database contains anisotropic materials.

Figure 4.13 shows the approximation accuracy for each CURET material. Material spectrums are computed based on different anisotropic reflectance bases compared to that of the proposed anisotropic basis. It can be observed that such basis provide the best approximation accuracy levels for all CURET materials especially with higher orders. It is again the spherical basis which provide the worst accuracy levels compared to the hemispherical ones. Notice that an anisotropic material such as vevlet is represented with higher approximation accuracy using the proposed basis when compared to others.

Figure 4.14 presents the required minimum truncating reflectance order to represent CUReT materials using the anisotropic basis where all the materials can be represented with accuracy of 99% with $P \leq 8$.

4.8 Directional Hemispherical Reflectance

The light leaving many surfaces is almost independent of the exit angle [8]. With the reciprocity property maintained, this can be rephrased as being independent of the incident angle. This gives an answer to what fraction of incident light is reflected. The fraction reflected primarily depends on the directional distribution of the incident light. Hence, surface reflectivity can be measured by the directional hemispherical reflectance function, where one direction is involved, instead of using the bidirectional reflectance distribution function depending on two. *Directional Hemispherical Reflectance* $f_r(\vec{\omega}'_i, \vec{\omega}'_o)$ of a surface is defined as the fraction of the incident irradiance in a given direction that is reflected by the surface, irrespective of the direction of reflection. It is dimensionless with values ranging from 0 to 1. It is obtained by integrating the radiance leaving the surface over all directions and dividing by the irradiance in the direction of illumination.

Representing BRDF spectrum as a function of either incident or outgoing/reflected direction is usually used to accelerate the evaluation of the reflectance integral in the image irradiance equation which forms the bottleneck in real time rendering. It is worth mentioning that assuming dependence on either incoming or outgoing directions implicitly assumes low frequency content on the respective domain and hence requires fewer coefficients.

4.8.1 Directional Hemispherical Basis

In order to translate the four-dimensional problem into a two-dimensional one, for a fixed incident direction, the BRDF is a hemispherical function of the outgoing direction¹, thus to obtain a better insight on the representation power of the proposed basis and how the approximation order affects the function reconstruction accuracy, it is helpful to derive the basis in such a frame. For a fixed incident direction $\vec{\omega}'_i$, the BRDF expansion coefficients can be written as,

$$a_{pr}^{qs}(\vec{\omega}'_i) = \int_{\Omega'_o} f_r(\vec{\omega}'_i, \vec{\omega}'_o) \left(\int_{\Omega'_i} \mathcal{H}_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o) d\vec{\omega}'_i \right) d\vec{\omega}'_o \quad (4.53)$$

¹This is also valid if the outgoing direction is fixed due to the Helmholtz reciprocity property.

where,

$$\begin{aligned}
 & \int_{\Omega'_i} \mathcal{H}_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o) d\vec{\omega}'_i \\
 &= N_{pr}^{qs} \int_{\Omega'_i} [H_p^q(\vec{\omega}'_i) H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i) H_p^q(\vec{\omega}'_o)] d\vec{\omega}'_i \\
 &= N_{pr}^{qs} H_r^s(\vec{\omega}'_o) \int_{\Omega'_i} H_p^q(\vec{\omega}'_i) d\vec{\omega}'_i \\
 & \quad + N_{pr}^{qs} H_p^q(\vec{\omega}'_o) \int_{\Omega'_i} H_r^s(\vec{\omega}'_i) d\vec{\omega}'_i
 \end{aligned} \tag{4.54}$$

Using the orthonormality of HSH basis Equation 4.6 and the fact that $H_0^0(\vec{\omega}'_i) = \frac{1}{\sqrt{2\pi}}$, one has,

$$\begin{aligned}
 \int_{\Omega'_i} H_p^q(\vec{\omega}'_i) d\vec{\omega}'_i &= \sqrt{2\pi} \int_{\Omega'_i} H_p^q(\vec{\omega}'_i) H_0^0(\vec{\omega}'_i) d\vec{\omega}'_i \\
 &= \sqrt{2\pi} \delta_{n0} \delta_{m0}
 \end{aligned} \tag{4.55}$$

Thus Equation 4.54 can be simplified as,

$$\int_{\Omega'_i} \mathcal{H}_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o) d\vec{\omega}'_i = \sqrt{2\pi} N_{pr}^{qs} [H_r^s(\vec{\omega}'_o) \delta_{p0} \delta_{q0} + H_p^q(\vec{\omega}'_o) \delta_{r0} \delta_{s0}] \tag{4.56}$$

Hence, the BRDF expansion coefficients under fixed incident light direction can be written as,

$$a_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o) = \sqrt{2\pi} N_{pr}^{qs} \int_{\Omega'_o} f_r(\vec{\omega}'_i, \vec{\omega}'_o) [H_r^s(\vec{\omega}'_o) \delta_{p0} \delta_{q0} + H_p^q(\vec{\omega}'_o) \delta_{r0} \delta_{s0}] d\vec{\omega}'_o \tag{4.57}$$

If $p = 0$ and $q = 0$ ¹, using Equation 4.10,

$$a_{0r}^{0s}(\vec{\omega}'_i, \vec{\omega}'_o) = \begin{cases} \sqrt{\pi} \int_{\Omega'_o} f_r(\vec{\omega}'_i, \vec{\omega}'_o) H_r^s(\vec{\omega}'_o) d\vec{\omega}'_o & r \neq 0, s \neq 0 \\ \int_{\Omega'_o} f_r(\vec{\omega}'_i, \vec{\omega}'_o) d\vec{\omega}'_o & r = 0, s = 0 \end{cases} \tag{4.58}$$

and if $p \neq 0$ and $q \neq 0$,

$$a_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o) = \begin{cases} \sqrt{\pi} \int_{\Omega'_o} f_r(\vec{\omega}'_i, \vec{\omega}'_o) H_p^q(\vec{\omega}'_o) d\vec{\omega}'_o & r = 0, s = 0 \\ 0 & r \neq 0, s \neq 0 \end{cases} \tag{4.59}$$

We can observe that the reciprocity property is also maintained in the expansion coefficients, where $a_{p0}^{q0} = a_{0p}^{0q}$. Therefore, as an analogy to the Fourier series expansion,

¹ $H_0^0(\vec{\omega}'_o) = \frac{1}{\sqrt{2\pi}}$

the BRDF at a given incident direction can be written as,

$$\begin{aligned}
 f_r(\vec{\omega}'_i, \vec{\omega}'_o) &= a_{00}^{00}(\vec{\omega}'_i, \vec{\omega}'_o) \mathcal{H}_{00}^{00}(\vec{\omega}'_i, \vec{\omega}'_o) \\
 &\quad + \sum_{p=1}^{\infty} \sum_{q=-p}^p a_{p0}^{q0}(\vec{\omega}'_i, \vec{\omega}'_o) \mathcal{H}_{p0}^{q0}(\vec{\omega}'_i, \vec{\omega}'_o) \\
 &\quad + \sum_{r=1}^{\infty} \sum_{s=-r}^r a_{0r}^{0s}(\vec{\omega}'_i, \vec{\omega}'_o) \mathcal{H}_{0r}^{0s}(\vec{\omega}'_i, \vec{\omega}'_o) \quad (4.60)
 \end{aligned}$$

Using the reciprocity property maintained by the expansion coefficients and the basis functions, the two series in Equation 4.60 are equal,

$$\begin{aligned}
 f_r(\vec{\omega}'_i, \vec{\omega}'_o) &= a_{00}^{00}(\vec{\omega}'_i, \vec{\omega}'_o) \mathcal{H}_{00}^{00}(\vec{\omega}'_i, \vec{\omega}'_o) \\
 &\quad + 2 \sum_{r=1}^{\infty} \sum_{s=-r}^r a_{0r}^{0s}(\vec{\omega}'_i, \vec{\omega}'_o) \mathcal{H}_{0r}^{0s}(\vec{\omega}'_i, \vec{\omega}'_o) \quad (4.61)
 \end{aligned}$$

where

$$\begin{aligned}
 \mathcal{H}_{00}^{00}(\vec{\omega}'_i, \vec{\omega}'_o) &= \frac{1}{\sqrt{2 + 2\delta_{00}\delta_{00}}} \left[H_0^0(\vec{\omega}'_i) H_0^0(\vec{\omega}'_o) + H_0^0(\vec{\omega}'_o) H_0^0(\vec{\omega}'_i) \right] \\
 &= \frac{1}{2} \left[\frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{2\pi}} + \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{2\pi}} \right] \\
 &= \frac{1}{2\pi} \quad (4.62)
 \end{aligned}$$

and for $r \geq 1, s \in [-r, r]$ one has,

$$\begin{aligned}
 \mathcal{H}_{0r}^{0s}(\vec{\omega}'_i, \vec{\omega}'_o) &= N_{0r}^{0s} \left[H_0^0(\vec{\omega}'_i) H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i) H_0^0(\vec{\omega}'_o) \right] \\
 &= \frac{1}{\sqrt{2 + 2\delta_{0r}\delta_{0s}}} \left[\frac{1}{\sqrt{2\pi}} H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i) \frac{1}{\sqrt{2\pi}} \right] \\
 &= \frac{1}{2\sqrt{\pi}} \left[H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i) \right] \quad (4.63)
 \end{aligned}$$

Hence, the HSH-based basis under fixed incident direction can be defined as,

$$\mathcal{H}_{0r}^{0s}(\vec{\omega}'_i, \vec{\omega}'_o) = \begin{cases} \frac{1}{2\pi} & r = 0, s = 0 \\ \frac{1}{2\sqrt{\pi}} \left[H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i) \right] & r \geq 1, s \in [-r, r] \end{cases} \quad (4.64)$$

Therefore the BRDF can be represented in terms of the expansion coefficients a_{0r}^{0s} in Equation 4.58 as follows,

$$f_r(\vec{\omega}'_i, \vec{\omega}'_o) = \frac{1}{2\pi} a_{00}^{00}(\vec{\omega}'_i, \vec{\omega}'_o) + \frac{1}{2\sqrt{\pi}} \sum_{r=1}^{\infty} \sum_{s=-r}^r a_{0r}^{0s}(\vec{\omega}'_i, \vec{\omega}'_o) \left[H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i) \right] \quad (4.65)$$

Considering the R th order function approximation, the approximation accuracy (or the energy captured by the R th order approximation) of the BRDF representation under fixed incident direction using HSH-based basis can be defined as (refer to [Appendix C: Approximation Accuracy of Directional Hemispherical Basis](#) for proof),

$$Acc_B(R; \vec{\omega}'_i) = \begin{cases} \frac{(a_{00}^{00}(\vec{\omega}'_i))^2}{\int_{\Omega'_o} \int_{\Omega'_i} [f_r(\vec{\omega}'_i, \vec{\omega}'_o)]^2 d\vec{\omega}'_i d\vec{\omega}'_o} & R = 0 \\ \frac{(a_{00}^{00}(\vec{\omega}'_i))^2 + \sum_{r=1}^R \sum_{s=-r}^r \left\{ 2\sqrt{\pi} a_{00}^{00}(\vec{\omega}'_i) a_{0r}^{0s}(\vec{\omega}'_i) H_r^s(\vec{\omega}'_i) + (a_{0r}^{0s}(\vec{\omega}'_i))^2 \left(\frac{1}{2} + \pi [H_r^s(\vec{\omega}'_i)]^2 \right) \right\}}{\int_{\Omega'_o} \int_{\Omega'_i} [f_r(\vec{\omega}'_i, \vec{\omega}'_o)]^2 d\vec{\omega}'_i d\vec{\omega}'_o} & R > 0 \end{cases} \quad (4.66)$$

4.8.2 Experimentation on Compact Description of BRDF Physical models

This subsection investigates the representation power of the proposed directional hemispherical basis to approximate physical analytical models of BRDF such as Oren-Nayer BRDF [14] and Cook-Torrance BRDF [13]. They are physically-based micro-facet models where the former models facets as perfectly diffuse while the latter is focused on glossy reflectance.

Consider a BRDF model with a parameter set $\mathcal{P} = \{p_1, p_2, \dots, p_d\}$ whose incoming hemisphere being sampled, where $\mathcal{P}_{CT} = \{\rho_d, \rho_s, m, n\}$ ¹ and $\mathcal{P}_{ON} = \{\rho_d, \sigma\}$ ² for Cook-Torrance and Oren-Nayer BRDFs respectively. The average amount of energy captured by the R th order approximation, *i.e.* average approximation accuracy, can be expressed as,

$$E\{Acc_B(R; \mathcal{P})\} = \int_{\Omega_i} Acc_B(R, \vec{\omega}'_i; \mathcal{P}) d\vec{\omega}'_i \quad (4.67)$$

Thus in order to study the effect of a specific parameter p_k , one can marginalize the average approximation accuracy to become a function of the parameter p_k . In order to assess the representation power of the proposed basis, the parameter-set domain of

¹ n is the classical refractive index from geometric optics, m is the root mean square slope of the micro-facets describing the surface's roughness, ρ_s is the surface specular reflectivity and ρ_d is the surface diffuse reflectivity.

² σ is the standard deviation of the Gaussian distribution used to represent surface roughness and ρ_d is the surface diffuse reflectivity.

Cook-Torrance is sampled based on the statistics for Torrance-Sparrow BRDF model, provided by the MERL/ETH skin reflectance database developed by Weyrich *et al.* [95], which was derived from measuring the skin reflectance of 156 subjects with varying age, gender and race. In the presented experimentation, the distribution of surface roughness m and specular reflectivity ρ_s are sampled, where $\mu(m) = 0.3032$, $\text{std}(m) = 0.0891$, $\mu(\rho_s) = 0.3753$ and $\text{std}(\rho_s) = 0.1655$, while ρ_d was taken as $1 - \rho_s$ and skin refractive index $n = 1.4433$ according to [97]. On the other hand, the domain of surface roughness parameter for the Oren-Nayar BRDF was uniformly sampled where $\sigma \in [0, 1]$ where unit diffuse albedo is assumed.¹

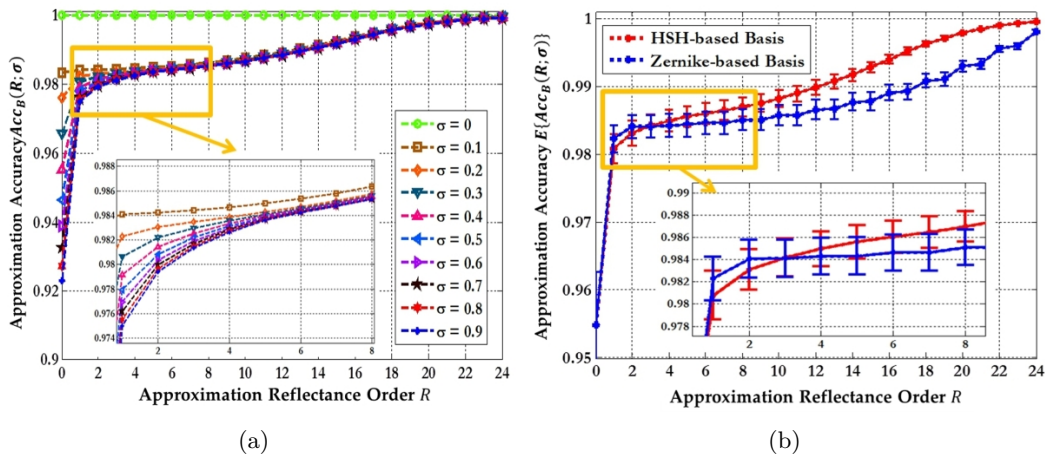


Figure 4.15: The Oren-Nayar BRDF approximation accuracy of the Helmholtz HSH-based basis as a function of the approximating order, where in (a) each curve corresponds to the average approximation accuracy over all sampled incident directions for several values of surface roughness, and in (b) it is compared to Zernike-based basis proposed by Koenderink *et al.* [21], where the mean and standard deviation were taken over $e_N(\sigma)$.

Figure 4.15(a) shows the Oren-Nayar BRDF approximation accuracy of the Helmholtz HSH-based basis as a function of the approximating order R for several values of surface roughness σ . Each curve corresponds to the average approximation accuracy over sampled incident directions. It can be observed that; (1) for zero surface roughness, the Oren-Nayar BRDF boils to a Lambertian model whose total energy is captured by the zero-order approximation, (2) as the surface roughness increase, zero-order approximation captures less energy, thus with rougher surfaces, higher orders are needed to maintain a specific level of accuracy, (3) beyond the sixth order, excluding smooth surfaces, the surface roughness parameter has insignificant impact on the Oren-Nayar BRDF approximation accuracy. Figure 4.15(b) shows mean and standard deviation of

¹We used Monte Carlo integration to evaluate the expansion coefficients of the BRDFs.

$Acc_B(R; \sigma)$ of the Oren-Nayar BRDF for the Helmholtz HSH-based basis in contrast to the basis based on Zernike polynomials, as functions of the approximation order R . It can be observed that (1) both basis have the same zero-order accuracy, since the zero-order basis corresponds to the BRDF DC component, (2) Zernike-based basis outperforms HSH-based basis up to the second order, after which HSH-based basis performance grows faster than that of the Zernike-based basis. This emphasizes the effect of reflectance basis at all combination of orders and degrees as in the case of HSH when compared to Zernike polynomials.

In case of Cook-Torrance BRDF, the grazing incident angles represents a problem in the BRDF construction, where the shape of the BRDF varies from Lambertian-like shape in case of very low specular reflectivity to specular-lobe in case of high specular reflectivity. This introduces a source of discontinuity/non-smoothness in the BRDF. Since harmonic-based basis are compact in the frequency domain, they are capable of representing smooth BRDFs with fewer non-zero coefficients when compared to complex BRDFs. It has been known that both SH and Zernike polynomials require large number of basis to represent specular BRDFs [96]. Since HSH is an adapted version of SH, it is anticipated to behave similarly. Rusinkiewicz [96] proposed a change of variables to re-parameterize the BRDF in order to yield efficient decompositions/representations of specular BRDFs, this can be viewed as one of the potential future efforts. This work investigates the proposed reflectance basis in the original/intuitive parameterizations of incident and outgoing angles. For purposes of comparison to Zernike-based basis, one has two situations; (1) non-grazing and (2) grazing incident angles, where throughout the experimentations, an incident angle of $\theta'_i > \frac{7\pi}{18}$ was considered to be a grazing incident angle. Figure 4.16 shows an example of fitting HSH-based basis versus the Zernike-based ones on a Cook-Torrance BRDF with high surface specular reflectivity.

Figure 4.17 shows the approximation accuracy of the Cook-Torrance BRDF in case of non-grazing and grazing incident angles, where each curve shows the mean and standard deviation of $Acc_B(R; \rho_s, m)$. One can observed that the HSH-based basis functions provide better representation when compared to the Zernike-based ones in case of non-grazing incident angles, however, this is not the case for grazing angles where specularly perturbs the smoothness of the BRDF, one way to tackle such a problem is to re-parameterize the BRDF as proposed by [96].

In the following, discussion is restricted to the case of non-grazing incident angles.

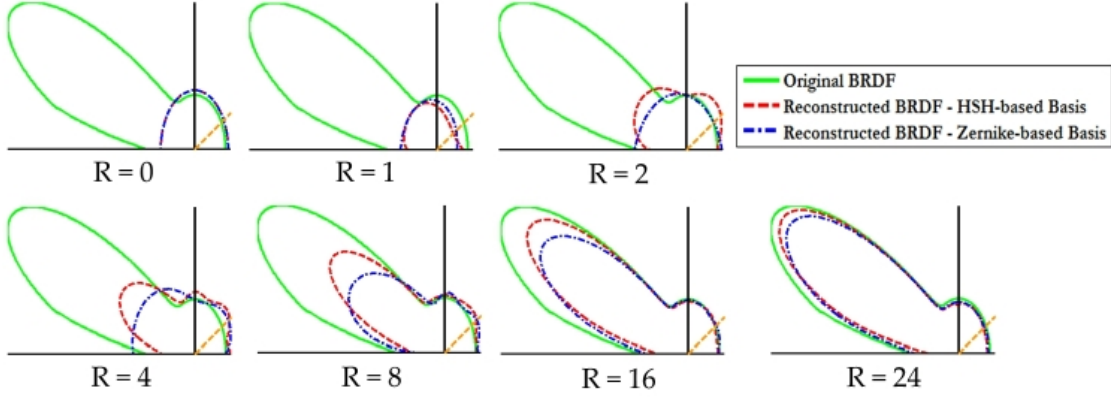


Figure 4.16: Example of fitting HSH-based basis (dashed red) versus the Zernike-based ones (dashed blue) on a Cook-Torrance BRDF (solid green) with high surface specular reflectivity, where $\rho_s = 0.458$, $m = 0.125$, $\theta_i = \pi/3$, $\phi_i = 0$. Notice how HSH-based basis provide better fitting at lower orders when compared to the Zernike-based basis. For visualization purposes, the incident plane is plotted and dashed orange line represents the light incident direction.

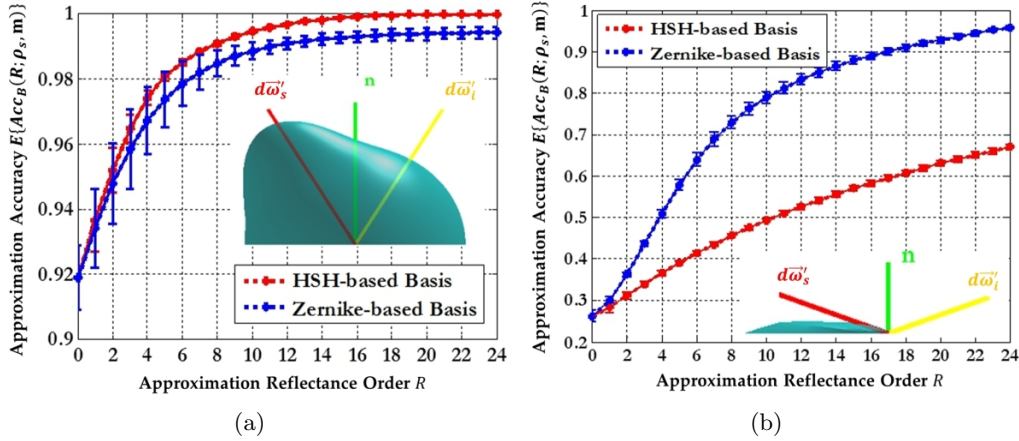


Figure 4.17: The Cook-Torrance BRDF approximation accuracy of the Helmholtz HSH-based basis compared to Zernike-based basis proposed by Koenderink *et al.* [21] as a function of the approximating order, where the mean and standard deviation were taken over $e_N(\rho_s, m)$: (a) non-grazing incident angles, where $\theta_i \in [0, \frac{7\pi}{18}]$, (b) grazing incident angles, where $\theta_i > \frac{7\pi}{18}$. An example of the BRDF is visualized for the two cases where $m = 0.2$, $n = 1.44$, $\rho_s = 0.6$ and $\rho_d = 0.4$.

Figure 4.18 illustrates the effect of surface roughness of Cook-Torrance model on the approximation accuracy of HSH-based basis, Figure 4.18(a), versus Zernike-based ones, Figure 4.18(b), where each curve represents the average approximation accuracy taken over several values of surface specular reflectivity. In general, the rougher the surface, the less specular it looks, this results in having more energy captured at lower approximation orders. In the close-up views, it can be observed that the effect of surface roughness on the approximation accuracy decreases at lower orders in case of HSH-based basis

when compared to the Zernike-based ones. In other words, Zernike-based basis are more sensitive to surface roughness parameter of the Cook-Torrance model.

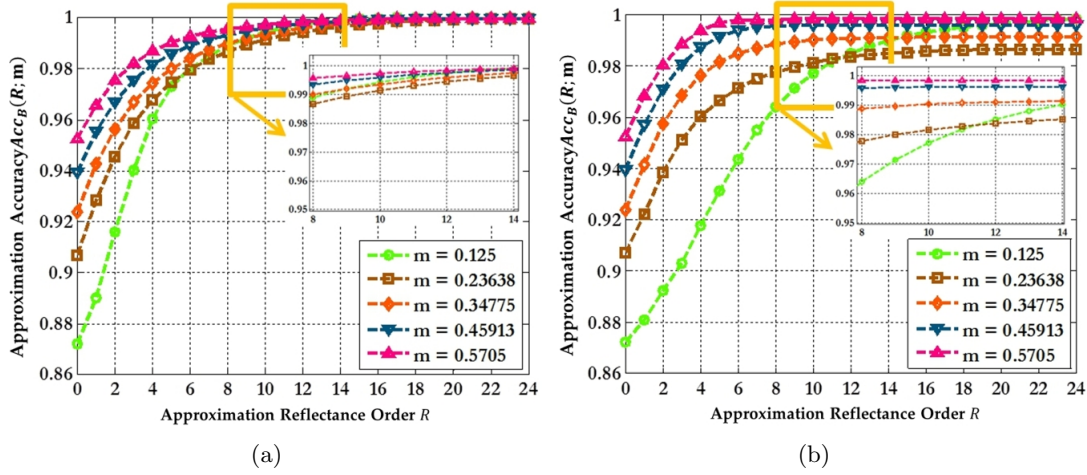


Figure 4.18: The Cook-Torrance BRDF approximation accuracy of (a) the Helmholtz surface reflectance basis, (b) Zernike-based basis, as a function of the approximating order, where each curve corresponds to the average approximation accuracy over all sampled incident directions for several values of surface roughness.

Figure 4.19 shows the effect of surface specular reflectivity on the approximation accuracy of HSH-based basis Figure 4.19(a) versus that of Zernike-based basis Figure 4.19(b), where each curve represents the average approximation accuracy taken over several values of surface roughness. With lower specular reflectivity, the performance of the approximation accuracy tends toward the Lambertian case, where all the function's energy is captured with the zero-order coefficient. This performance degrades with more specular reflectivity. It can be observed that the Helmholtz basis captures more function energy at lower orders when compared to Zernike-based ones. In the close-up views, the sensitivity of the Zernike-based basis to the specular reflectivity can be noticed, while on the other hand the effect of the specular reflectivity on the approximation accuracy decreases at lower orders in case of HSH-based basis.

4.9 Summary

This chapter proposed a complete, orthonormal basis to provide a compact and efficient representation for surface bidirectional reflectance distribution function (BRDF), which is defined on the cartesian product of two hemispheres. The proposed basis, which are defined in terms of hemispherical harmonics (HSH), preserve the Helmholtz reciprocity

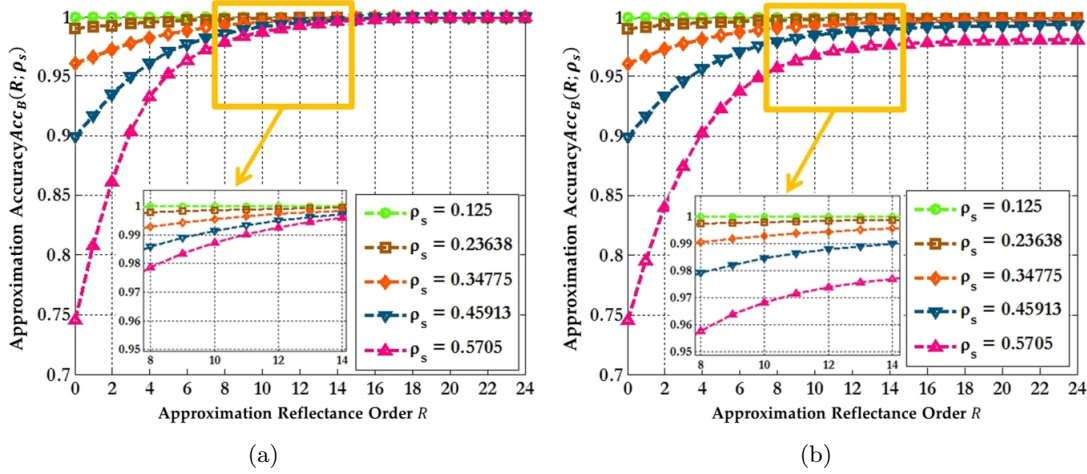


Figure 4.19: The Cook-Torrance BRDF approximation accuracy of (a) the Helmholtz surface reflectance basis, (b) Zernike-based basis, as a function of the approximating order, where each curve corresponds to the average approximation accuracy over all sampled incident directions for several values of specularity reflectivity.

property of BRDFs while avoid the computational complexity inherited from Zernike polynomials that are usually used to construct hemispherical basis. An analytical as well as experimental justification was presented such that for a given truncating reflectance order, the proposed hemispherical basis provide better approximation accuracy of the BRDF when compared to similar bases in literature. While hemispherical basis provide higher approximation accuracies when compared to spherical ones, basis maintaining helmholtz property was observed to provide higher accuracy levels compared to others. The closed form of the proposed basis was presented in case of isotropic and directional hemispherical reflectance. The proposed basis was further validated using scattered reflectance data which might violate the Helmholtz reciprocity property; where such property is maintained in the least-squares sense in the process of fitting the BRDF measurements to the HSH-based basis. The basis was also validated using micro-facet physical BRDF models: Oren-Nayar for diffuse rough surfaces and Torrance-Sparrow and Cook-Torrance for specular rough surfaces. Based on the fact that associated Legendre polynomials are defined for all combinations of polynomial order and degree in contrast to Zernike polynomials, HSH-based basis showed high BRDF approximation accuracy at lower orders. It was observed that the significance of BRDF parameters decayed with the approximation order versus the case of Zernike-based basis. The grazing incident angles represented a problem in case of Cook-Torrance BRDF which introduces a source of non-smoothness in the BRDF, where the HSH-based basis functions provide

better representation when compared to the Zernike-based ones in case of non-grazing incident angles, however, this is not the case for grazing angles. Modeling measured reflectance was further investigated where the Helmholtz reciprocity might be violated. It was shown that the proposed basis captured almost all diffuse and specular materials with reflectance order less than $P = 10$. Ongoing efforts are directed towards handling grazing incident angles using BRDF re-parameterization.

IMAGE IRRADIANCE HARMONICS

Phenomenological appearance models capture surface appearance through mathematical modeling of the reflection process. Theoretically, the space of all possible images of a fixed-pose object under all possible illumination conditions is infinite dimensional. Nonetheless, due to their low-frequency nature, irradiance signals can be represented using low-order basis functions. Discounting subsurface scattering and surface emittance, this work seeks to address the question; how to compactly and accurately represent image irradiance under unknown general illumination, given that a surface point sees its surrounding world through the local upper hemisphere oriented by the surface normal at this point. This chapter formulates the image formation process of isotropic surfaces under arbitrary distant illumination in the frequency space while addressing the physical compliance of hemispherical basis for representing surface reflectance, e.g. Helmholtz reciprocity and isotropy. The term *irradiance harmonics* is also defined which enables decoupling illumination and reflectance from the underlying geometry and pose. This work provides a closed form of the energy content being maintained by different reflectance modes of the proposed irradiance harmonics. Since specular materials tend to require more basis functions when compared to diffuse ones, the presented harmonics captures same cumulative energy content, by providing larger number of orthogonal irradiance basis, at lower illumination orders when compared to similar basis in literature.

5.1 Introduction

The crux of *Physics-based vision* [1] is inferring information about an image's contents based on modeling the underlying physical process of image formation. However, such a problem is under-constrained since the formation process involves numerous unknowns

in contrast to the available information at hand, *i.e.* image brightness value per pixel. Thus information inference based on shading commonly involves adding assumptions to the image formation process such as distant illumination and homogenous reflectance.

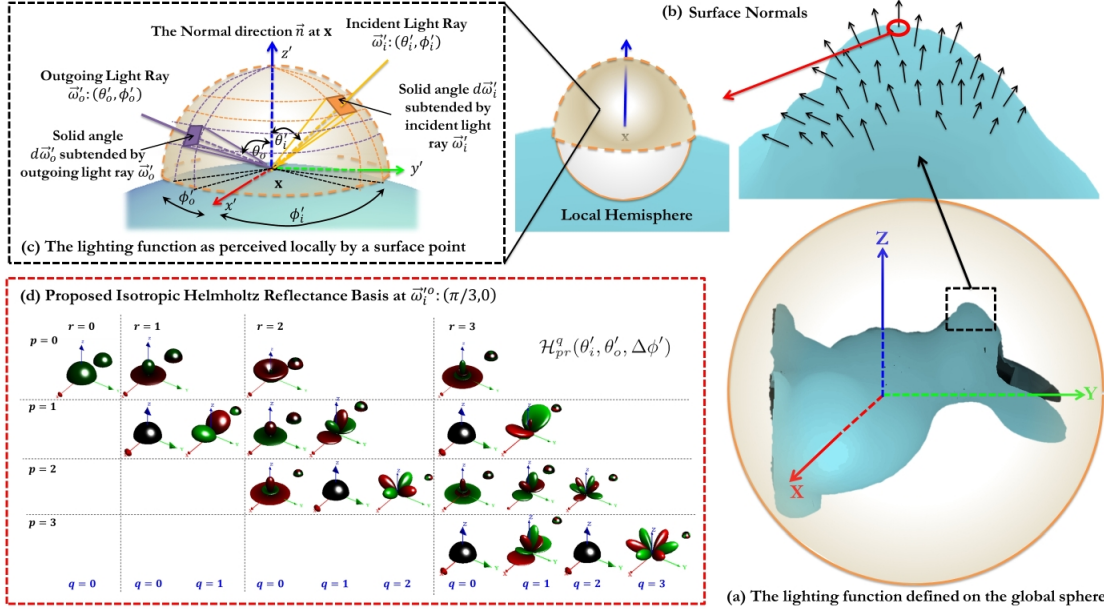


Figure 5.1: (a) An object’s surface is illuminated under distant lighting function defined on the global sphere. (b) An in-depth view of a surface patch showing surface normals at each surface point. Under the assumption of non-emitting surfaces, a surface point only sees the surrounding environment through the local upper hemisphere oriented by the surface normal at this point, thus the reflection integral can be defined over the incident local hemisphere where $\theta_i' \in [0, \pi/2]$. (c) A zoom-in view at a surface point \mathbf{x} seeing its surrounding world through a unit hemisphere Ω' centered at the point and oriented by the surface normal \vec{n} at that point. Using spherical coordinates, a light ray incident to the point \mathbf{x} is defined by its unit direction $\vec{\omega}_i' : (\theta_i', \phi_i')$, subtending an incident solid angle $d\vec{\omega}_i'$. Similarly, an outgoing/reflected light ray is defined by its unit direction $\vec{\omega}_o' : (\theta_o', \phi_o')$, subtending an outgoing solid angle $d\vec{\omega}_o'$. (d) Visualization of up-to 3rd order of the proposed *isotropic* Helmholtz surface reflectance basis at an incident direction $\vec{\omega}_i' : (\pi/3, 0)$. The order p runs from top to bottom while the associated order r runs from left to right. The azimuthal order q , or known as degree, also runs from left to right for each order r .

Appearance is the net result of the surface reflectance characteristics when exposed to illumination. Theoretically, due to the arbitrariness of the lighting function, the space of all possible images of a fixed-pose object under all possible illumination conditions is infinite dimensional [25]. Yet, statistical modeling is widely devised to provide a low-dimensional linear subspace of image irradiance by performing Principal Component Analysis (PCA) on a large set of images under various imaging conditions [27]. However, this approach suffers from the need of special acquisition setup to control and measure the lighting function while keeping the camera and the object fixed. In addition, the

constructed subspace is biased towards the sampled illumination directions. A problem which can be alleviated by using a “large-enough” image ensemble rendering numerical methods intractable due to the curse of dimensionality. Moreover, either surface reflectance properties are assumed or known a-priori.

On the other hand, appearance models can be constructed using phenomenological models which capture variations of surface appearance with respect to illumination conditions and surface reflectance properties through mathematical abstraction of the reflection process. In particular, the harmonic expansion of the image irradiance equation can be used to derive an analytic subspace to approximate images under fixed pose but different illumination conditions [6, 11, 25, 44]; it has been shown that surfaces act as a low-pass filter on the incident illumination [4, 5], yielding a band-limited reflected light, introducing the frequency-space representation of image irradiance.

Discounting subsurface scattering and surface emittance, this work seeks to address the question; how to compactly and accurately represent image irradiance under unknown general illumination, given that a surface point sees its surrounding world through the local upper hemisphere oriented by the surface normal at this point, see Figure 5.1. This work is after deriving an analytic form for the image irradiance basis which spans the subspace representing the set of images under all possible illumination conditions of non-Lambertian surfaces; taking into account that hemispherical functions present discontinuities at the boundary of the hemisphere when represented in the spherical domain [20], demanding more spectrum coefficients for accurate representation. In contrast to the case where spherical harmonics (SH) is used to represent hemispherical reflectance kernels [6, 11], the physical compliance of the hemispherical basis is addressed for representing surface arbitrary bidirectional reflectance distribution functions.

The previous chapter addressed the physical compliance of reflectance hemispherical basis by defining a Cartesian product of hemispherical harmonics (HSH) to provide a compact and accurate representation for arbitrary reflectance, while satisfying the Helmholtz’s law of reciprocity [52]. This chapter defines the image irradiance basis (or what is termed as *irradiance harmonics*) based on the proposed reflectance basis. While maintaining isotropy and reciprocity properties reduces drastically the number of basis, the proposed isotropic Helmholtz reflectance basis contains more basis functions than the widely used Zernike-based ones [32] at a given approximating order. This is based on the fact that the set of associated Legendre polynomials is distinguished by

the property that it contains a polynomial for every combination of order and degree [82], compared to Zernike polynomials which are restricted to even differences between polynomial order and degree. This chapter proposes a compact representation of image irradiance of isotropic surfaces by adopting frequency domain based representation for the purposes of illumination modeling and image analysis.

Due to the illposedness of shading-based inference problems, assumptions are commonly used to present a simplified version of the image formation process. For instance, since local surface patches have sizes relatively smaller than the distance to the light source, the incident light field can be assumed to be constant on these local patches. Hence distant illumination is a common assumption in many vision algorithms. This chapter assumes homogeneous, convex surfaces under distant illumination conditions. Non-homogeneous surfaces with spatially varying reflectance properties are often approximated in vision applications by using a spatially varying texture which modulates the surface radiance [11]. These assumptions are considered a reasonable approximation to various computer vision tasks while being flexible enough for analytical formulation.

5.2 Related Work

An image formed by a convex surface is an albedo-modulated version of the *surface reflectance function*, which is an analytic expression representing the process of turning the surface irradiance into outgoing radiance (reflection). It has been shown that surfaces acts as a low-pass filter on the incident illumination [4, 5], yielding a band-limited reflected light. This introduces the frequency-space representation of the surface reflectance function and in turn the image irradiance equation, where spherical harmonics (SH) based representation is widely used as a powerful tool to analyze shading theoretically [32], accounting for arbitrary illumination.

D’Zmura [5] expressed the image irradiance equation using SH as an analogy to Fourier series on the real line, to explore ambiguities in lighting. Since then, research has been conducted to analytically derive a linear subspace for the image irradiance equation to aid computer vision tasks to handle illumination variations. Basri and Jacobs in [6] and Ramamoorthi and Hanrahan in [7, 11] formulated the image irradiance equation [2] in a convolution framework where the lighting function acts as a signal filtered by the Lambertian kernel, opening the whole field of signal processing and fil-

tering to the image irradiance modeling. They provided an analytical expression of an image of a convex-Lambertian object illuminated by distant lighting using SH. They proved that the Lambertian reflectance kernel acts as a low pass filter, presenting the image irradiance as a band-limited signal being represented by a finite number of basis functions. While Basri and Jacobs [6] formulated this process in global coordinates with respect to global reference frame, Ramamoorthi and Hanrahan [7, 11] made the distinction between such a formulation in global and local coordinates (*w.r.t.* surface points).

Whereas Basri and Jacobs in [6] and Ramamoorthi in [11] provides a single expression for the harmonic basis, Ramamoorthi [25], under the assumption of distant light, linearly combined the SH basis to build an optimal basis for image spaces illuminated by uniformly distributed point light sources to create a new set of orthonormal eigenfunctions over this restricted domain. In particular, Ramamoorthi in [25] used analytical principal component analysis (PCA) in order to provide orthogonal set of basis in such a domain, this resulted in a convex combination of higher order spherical harmonic basis functions. On the other hand, QR-decomposition is used in [6] to numerically orthogonalize up-to 2nd order SH basis.

Under the assumption of no surface emittance, surface reflectance is defined on the cartesian product of two hemispheres corresponding to the incident and outgoing directions; the nature way to represent such a hemispherical function is to use hemispherical basis. However, due to their compactness in the frequency space, SH has been extensively used for this purpose. Recently Elhabian *et al.* [40, 41] avoided the use of higher order SH basis and numerical orthogonalization to model the reflectance of convex-Lambertian surfaces under single viewpoint. They used hemispherical harmonics (HSH) which forms an orthonormal set of basis over a unit hemisphere, thus PCA or basis orthogonalization is no longer needed. Nonetheless, the incident illumination is restricted to the upper global hemisphere facing the camera while Lambertian reflectance is assumed.

For non-Lambertian surfaces, in contrast to [11], Nillius [32] used the hemispherical basis proposed by Koenderink and van Doorn [21]. Such bases are based on Zernike radial polynomials, where bases for the unit disk are mapped onto the upper hemisphere. Yet, such polynomials have high computational cost [20] when compared to associated Legendre polynomials used for (hemi)spherical harmonics. In particular, Zernike radial polynomials require an amount of CPU time proportional to $\mathcal{O}(N^2)$ [42] in contrast

Table 5.1: Assumptions and Harmonic Basis of Different Formulations of Image Irradiance Equation (θ'_i is the angle between surface normal and incident light direction)

	Light	Camera	Object	Integration Domain	Harmonic Basis	Attached Shadows	Notes
Basri and Jacobs [6]	Distant (directional)	Distant (Orthographic projection assumed)	Convex, Lambertian	Global sphere	Spherical harmonics	$\max(0, \cos \theta'_i)$	Using spherical basis to represent hemispherical domain. Cant be extended to non-Lambertian surfaces unless local hemisphere is considered.
Ramamoorthi and Hanrahan [7, 11, 31]	Distant (directional) versus uniform lighting	Distant (Orthographic projection assumed)	Convex, non-Lambertian	Global vs. local sphere	Spherical harmonics	$\max(0, \cos \theta'_i)$	Using spherical basis to represent hemispherical domain.
Ramamoorthi [25]	Distant (directional)	Distant (Orthographic projection assumed)	Convex, Lambertian	Global sphere vs. global hemisphere	Spherical harmonics	$\max(0, \cos \theta'_i)$	Using spherical basis to represent hemispherical domain.
Nillius and Eklundh [32, 44]	Distant (directional)	Distant (Orthographic projection assumed)	Convex, non-Lambertian	Local hemisphere	Zernike-based hemispherical basis	Integrating over hemisphere	Zernike polynomials are known to have high computational complex in addition of being undefined for all combination of harmonic orders and degrees. Rotation matrices are not defined for Zernike polynomials [20], hence they can not be used in the global hemispherical domain.
Elhabian et al. [40, 41]	Distant (directional)/Investigate Near lighting	Distant (Orthographic projection assumed)	Convex, Lambertian	Global hemisphere	Hemispherical Harmonics basis	Integrating over hemisphere	Incident light is restricted to the upper hemisphere facing the camera while Lambertian reflectance is assumed.
Our Model	Distant (directional)	Distant (Orthographic projection assumed)	Convex, non-Lambertian	Local hemisphere	Helmholtz Hemispherical Harmonics-based basis	Integrating over hemisphere	Work can be extended to handle near illumination.

to associated Legendre polynomials which needs $\mathcal{O}(N \log^2 N)$ [43] for cutoff frequency $N = n^2$ with order $n \geq 0$. Table 5.1 summarizes the assumptions and harmonic basis of different formulations of the image irradiance equation compared to the proposed formulation.

5.3 Irradiance Harmonics Definition

A parametric model for the image irradiance involves a frequency-space representation of the lighting function and the surface BRDF, where their spectrum coefficients in the frequency domain aid inferring one of them given the other, *i.e.* inverse rendering [11]. The main challenge will be deriving a frequency-space basis which efficiently and accurately models the image irradiance of any surface under arbitrary illumination.

Consider a distant viewer where the orthographic projection can be used to approximate the geometric imaging process. According to Horn model of a single lens camera [24], the image irradiance equation of an object with isotropic reflectance can be written as¹

$$E(\alpha, \beta) = \int_{\mathbb{D}'} L_i(R_{\alpha, \beta}(\theta'_i, \phi'_i)) f_r(\theta'_i, \alpha, |\phi'_i - \pi|) t_f(\theta'_i) d\vec{\omega}'_i \quad (5.1)$$

where (α, β) are the spherical coordinates of the surface normal $\vec{n}(\mathbf{x})$ in the global reference frame, $L_i(R_{\alpha, \beta}(\vec{\omega}'_i))$ denotes the incident radiance at a surface point \mathbf{x} in the direction $\vec{\omega}'_i : (\theta'_i, \phi'_i) \in \Omega'_i$ with Ω'_i denoting the local hemisphere oriented by $\vec{n}(\mathbf{x})$, $R_{\alpha, \beta} = R_Z(\beta)R_Y(\alpha)$ is the Euler angle representation of 3D rotations with ZYZ-convention and f_r denotes the Bidirectional Reflectance Distribution Function (BRDF). Refer to Figure 8.1 for illustration.

To account for attached shadows, the image irradiance equation in Equation 5.1 can be formulated using two different local integration domains \mathbb{D}' *w.r.t.* the surface point. A spherical domain \mathcal{S}^2 requires the foreshortening term $t_f(\theta'_i)$ to be defined as $\max(0, \cos \theta'_i)$, *e.g.* [6, 7, 11]. Whereas under the assumption of non-emitting surfaces, a surface point only sees the surrounding environment through the local upper hemisphere Ω'_i oriented by the surface normal at this point. This formulation accounts implicitly for attached shadows without imposing the nonlinear max operator, *i.e.* $t_f(\theta'_i) = \cos \theta'_i$, *e.g.* [44].

¹Primed coordinates are used to denote local reference frame with respect to a surface point.

The lighting function can be written in terms of spherical harmonics (SH) basis as,

$$L_i(R_{\alpha,\beta}(\theta'_i, \phi'_i)) = \sum_{n=0}^{\infty} \sum_{m=-n}^n \sum_{m'=-n}^n l_n^m \mathcal{D}_{mm'}^n(\alpha, \beta, \gamma) Y_n^{m'}(\theta'_i, \phi'_i) \quad (5.2)$$

where $\{l_n^m\}$ are the illumination spectrum coefficients when projected onto the subspace spanned by the real SH bases $\{Y_n^m\}$ and \mathcal{D}^n are Wigner's rotation matrices [59] which encodes how to express a rotated SH basis function in terms of all other SH bases of the same order.

Let the surface BRDF f_r be represented by (hemi)spherical basis functions denoted as \mathcal{F}_{pr}^q up to reflectance order P s.t. $p \in [0, P]$ and $\{r, q\}$ indices are given according to the type of reflectance basis, e.g. [7, 44, 45]. Hence its harmonic expansion can be written as,

$$f_r(\theta'_i, \alpha, |\phi'_i - \pi|) = \sum_{prq} a_{pr}^q \mathcal{F}_{pr}^q(\theta'_i, \alpha, |\phi'_i - \pi|) \quad (5.3)$$

where $\{a_{pr}^q\}$ are the BRDF spectrum coefficients when projected onto the subspace spanned by the reflectance bases $\{\mathcal{F}_{pr}^q\}$. Thus the image irradiance can be written as a linear combination of precomputed basis which is termed as **image irradiance harmonics**. Such harmonics encodes the orthogonality relation between SH and the deployed BRDF representation. It can be written as follows where due to the orthogonality relation of the azimuthal functions, all terms in the integral vanish except for $m' = q$.

$$\mathcal{B}_{npr}^{mq}(\alpha, \beta) = \mathcal{D}_{mq}^n(\alpha, \beta) \int_{\mathbb{D}'} Y_n^q(\vec{\omega}'_i) \mathcal{F}_{pr}^q(\theta'_i, \alpha, |\phi'_i - \pi|) t_f(\theta'_i) d\vec{\omega}'_i \quad (5.4)$$

Hence, the frequency-space representation of the image irradiance in Equation 5.1 for isotropic surfaces can be written as follows,

$$E(\alpha, \beta) = \sum_{nm,prq} l_n^m a_{pr}^q \mathcal{B}_{npr}^{mq}(\alpha, \beta) = \sum_s c_s \mathcal{B}_s(\alpha, \beta) \quad (5.5)$$

where $c_s = l_n^m a_{pr}^q$ with s and its corresponding indices n, m, p, r and q are given by an ordering function based on the average power content (i.e. second moment) $P\{\mathcal{B}_{npr}^{mq}\}$ of the irradiance harmonics. This is defined as the integral of the respective harmonic over the global hemisphere of visible surface normals weighted by the surface normal

distribution $p(\alpha, \beta)$,

$$P\{\mathcal{B}_{npr}^{mq}\} = \int_{\Omega} [\mathcal{B}_{npr}^{mq}(\alpha, \beta)]^2 p(\alpha, \beta) d\vec{\omega} \quad (5.6)$$

5.4 Irradiance Harmonics for HSH-based Reflectance Basis

Representing the lighting function using SH and the surface reflectance using the Helmholtz HSH-based basis, the frequency-space representation of the image irradiance in Equation 5.1 for isotropic surfaces can be written as follows,

$$E(\alpha, \beta) = \int_{\Omega'_i} \left(\sum_{n=0}^{\infty} \sum_{m=-n}^n \sum_{m'=-n}^n l_n^m \mathcal{D}_{mm'}^n(\alpha, \beta) Y_n^{m'}(\vec{\omega}'_i) \right) \left(\sum_{p=0}^{\infty} \sum_{r=0}^p \sum_{q=0}^r a_{pr}^q \mathcal{H}_{pr}^q(\theta'_i, \alpha, |\phi'_i - \pi|) \right) \cos \theta'_i d\vec{\omega}'_i \quad (5.7)$$

Moving terms which are independent of the local incident direction out of the integral yields,

$$E(\alpha, \beta) = \sum_{\substack{nm \\ prq}} l_n^m a_{pr}^q \mathcal{D}_{mm'}^n(\alpha, \beta) \times \int_{\Omega'_i} Y_n^{m'}(\vec{\omega}'_i) H_{pr}^q(\theta'_i, \alpha, |\phi'_i - \pi|) \cos \theta'_i d\vec{\omega}'_i \quad (5.8)$$

Based on the orthogonality relation of the azimuthal functions, all terms in the integral vanish except for $m' = q$. The polar part of the integral can be derived as follows. Using variable substitution $z = \cos \theta'_i$, inserting the explicit forms of (shifted) associated Legendre polynomials [98]¹, using the binomial theorem which constrains $|q|$ to be even and using the Euler Beta integral [99] lead to (proof is included in [Appendix](#)

¹Mulder in [98] defined the series representation of associated Legendre polynomials, by replacing $\cos \theta'_i$ with $2 \cos \theta'_i - 1$, the explicit form of shifted associated Legendre can be derived using the identity $\sin^{|q|} \theta'_i = (1 - \cos^2 \theta'_i)^{|q|/2}$ for $\theta'_i, |q| \geq 0$

D: Analytic Derivation of $C_{npr}^q(\alpha)$,

$$\begin{aligned}
 \mathcal{C}_{npr}^q(\alpha) &= \int_{\Omega'_i} Y_n^q(\vec{\omega}'_i) \mathcal{H}_{pr}^q(\theta'_i, \alpha, |\phi'_i - \pi|) \cos \theta'_i d\vec{\omega}'_i \\
 &= \pi N_{pr}^q (-1)^q \sqrt{2(1 + \delta_{q0})} \\
 &\quad \left\{ \tilde{\Theta}_r^q(\theta'_o) N_n^{|q|} \tilde{N}_p^{|q|} \mathcal{J}_{np}^q + \tilde{\Theta}_p^q(\theta'_o) N_n^{|q|} \tilde{N}_r^{|q|} \mathcal{J}_{nr}^q \right\}
 \end{aligned} \tag{5.9}$$

where¹,

$$\mathcal{J}_{np}^q = \int_0^{\pi/2} P_n^{|q|}(\cos \theta'_i) \tilde{P}_p^{|q|}(\cos \theta'_i) \cos \theta'_i \sin \theta'_i d\theta'_i \tag{5.10}$$

Hence, the image irradiance of isotropic reflectance under arbitrary distant illumination, where surface points can be parameterized by their normals, can be represented as a sum of the irradiance harmonics \mathcal{B}_{npr}^{mq} as follows,

$$E(\alpha, \beta) = \sum_{nm} \sum_{prq, q \text{ is even}} l_n^m a_{pr}^q \underbrace{[\mathcal{D}_{mq}^n(\alpha, \beta) \mathcal{C}_{npr}^q(\alpha)]}_{\mathcal{B}_{npr}^{mq}(\alpha, \beta)} \tag{5.11}$$

Note the effect of the light and BRDF spectrum coefficients on the image irradiance representation. While theoretically, an infinite number of basis is needed to perfectly represent the irradiance signal, most natural BRDFs have a band-limited spectrum, thus acting as a low-pass filter on the illumination signal, leading to finite number of basis capturing a large portion of the irradiance signal energy content.

5.5 Irradiance Signal Energy Content

The concept of the signal's energy content normally arises in constructing linear subspaces [64]. In case of image irradiance representation, one is interested in maintaining the minimum number of basis which portrays the range of frequencies for which the signal has a significant energy content. Since the illumination function can be arbitrary while the surface reflectance is known to act as a low-pass filter, one assumes a directional light source, *i.e.* delta function, with a non-decayable spectrum along illumination bands. Thus the decay of the irradiance spectrum will be due to the BRDF spectrum. To eliminate the effect of the light direction, a light source undergoing all possible 3D

¹Proof and closed form is included in [Appendix D: Analytic Derivation of \$C_{npr}^q\(\alpha\)\$](#) .

rotations $SO(3)$ is considered.

Considering the surface normals perceived by the camera, the average energy content of the image irradiance signal over all possible directional light sources can be defined as follows, where the representation in Equation 5.11 and the orthogonality relation of the \mathcal{D}^n matrices¹ are used. Note that the integral is over the surface normals constructing the visible hemisphere Ω .

$$\varepsilon = \int_{\Omega} \int_{SO(3)} [E(\alpha, \beta)]^2 dR d\vec{\omega} = \sum_{n=0}^{\infty} \left(\frac{\sum_{o=-n}^n \tilde{l}_n^o}{2n+1} \right) \mathcal{A}_B(n) \quad (5.12)$$

where \tilde{l}_n^o are the lighting coefficients before source rotation and $\mathcal{A}_B(n)$ are the attenuation factors imposed by the surface BRDF on the illumination order n , defined as:

$$\mathcal{A}_B(n) = \sum_{prq} \sum_{p'r'q'} a_{pr}^q a_{p'r'}^{q'} \sum_{m=-n}^n \int_{\Omega} \mathcal{B}_{npr}^{mq}(\alpha, \beta) \mathcal{B}_{n p'r'}^{m q'}(\alpha, \beta) d\vec{\omega} \quad (5.13)$$

Using the isotropic Helmholtz HSH-based reflectance bases, the attenuation factors can be written as²,

$$\begin{aligned} \mathcal{A}_B(n) = & 2\pi^2 \sum_{\substack{prq \\ q \text{ is even}}} (1 + (-1)^q \delta_{q0}) (a_{pr}^q)^2 (N_{pr}^q)^2 (N_n^{|q|})^2 \\ & \left\{ \left(\tilde{N}_p^{|q|} \right)^2 (\mathcal{J}_{np}^q)^2 + \left(\tilde{N}_r^{|q|} \right)^2 (\mathcal{J}_{nr}^q)^2 \right. \\ & \left. + 2\tilde{N}_p^{|q|} \tilde{N}_r^{|q|} \mathcal{J}_{np}^q \mathcal{J}_{nr}^q \delta_{pr} \right\} \end{aligned} \quad (5.14)$$

In case of directional light source, the average signal energy maintained by the illumination order n is constant, *i.e.* $\frac{\sum_{o=-n}^n \tilde{l}_n^o}{2n+1} = \frac{1}{4\pi}$ [32]. Thus the approximation accuracy of the N th order expansion of the image irradiance can be written as a function of the BRDF attenuation factors where the energy due to light can be factored out.

$$Acc_E(N) = \frac{\sum_{n=0}^N \mathcal{A}_B(n)}{\sum_{n=0}^{\infty} \mathcal{A}_B(n)} \quad (5.15)$$

Figure 5.2 compares the attenuation factor of the Lambertian reflectance of the proposed irradiance basis in contrast to [32] as a function of the illumination order n where all degrees $m \in [n, n]$ are summed for visualization purposes. Although both bases shows

¹Proof is included in [Appendix E: Irradiance Signal Energy Content](#).

²See [Appendix E: Irradiance Signal Energy Content](#).

that 9 bases (2nd order illumination) suffice to capture most of the irradiance energy content. Nonetheless, the proposed basis provides a more compact representation due to smaller attenuation factors at higher order illumination *i.e.* $n > 2$ suppressing higher frequencies of illumination.

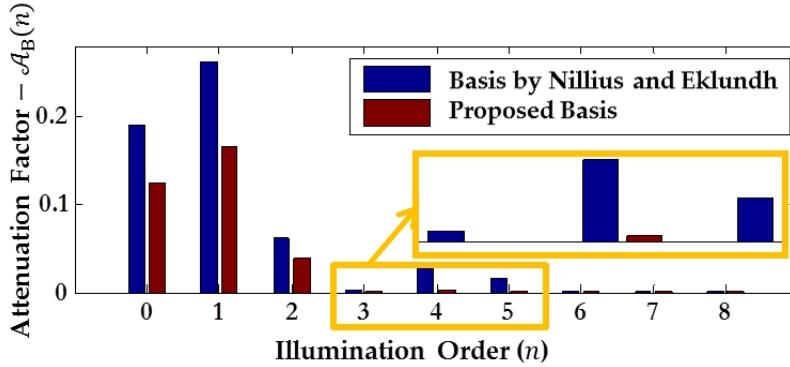


Figure 5.2: Attenuation factors $\mathcal{A}_B(n)$ in case of Lambertian reflectance as a function of the illumination order n . Blue shows the factors based on using Nillius and Eklundh [44]. reflectance basis while red shows the factors based on the proposed basis.

While the Lambertian case only need 9 basis for accurate representation, surfaces with non-Lambertian reflectance need more reflectance modes/basis through which more illumination frequencies can survive. Consider the light being reflected from a single reflectance mode \mathcal{H}_{pr}^q , Figure 5.3 shows the attenuation factors $\mathcal{A}_B(n)$ of reflectance modes (p, r, q) up to $p = 7$. For comparison purposes, the factors of the reflectance modes in [32] are also included. While the proposed basis agrees with [32] that the spectrum tend to attain a maxima at around the same illumination order n , since this basis is based on the Legendre polynomial family, more basis are available for non-Lambertian surfaces when compared to the Zernike-based basis of [32].

5.6 Experimental Results

This section evaluates the approximation accuracy of the irradiance harmonics based on the HSH-based reflectance basis. BRDF measurements are used which are directly measured from real surfaces [10, 17]. For accuracy computation, $\mathcal{A}_B(n)$ ¹ is computed as in Equation 5.13 using Monte Carlo integration. In case of the proposed reflectance basis, the analytical definition in Equation 5.14 is used instead for faster evaluation. In order to compensate for the infinite series in the denominator of Equation 5.15, the

¹Assuming single directional light source which has a non-decayable spectrum.

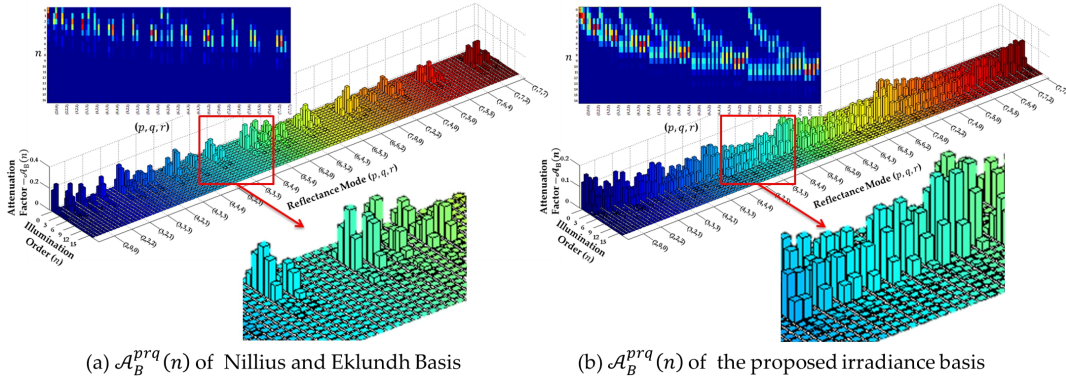


Figure 5.3: Attenuation factors $\mathcal{A}_B(n)$ of single reflectance modes (p, r, q) . The upper left corner of (a) and (b) shows a 2-dimensional view of the attenuation factors where the color encodes the level of attenuation. Note that the proposed basis provide larger number of basis with decayable spectrum when compared to the Zernike-based basis of [32].

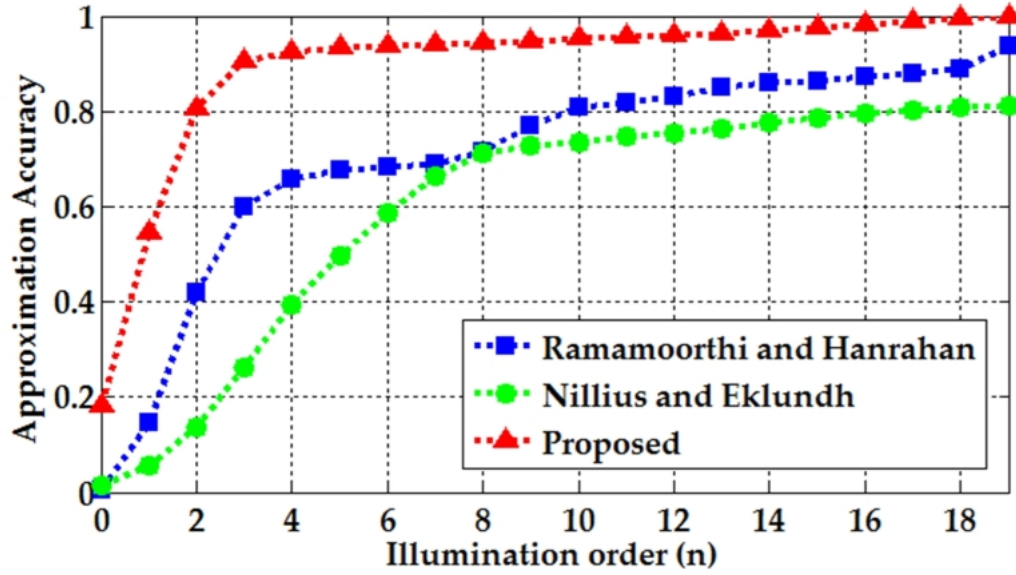
tail of the series is estimated by fitting an exponential curve to the series following an asymptotic illumination order at $n = 30$.

Figure 5.5 shows rendered images of lettuce-leaf and felt BRDFs [10] using single directional light source which was performed similarly to [31]. By visual inspection, the rendered image approaches the groundtruth rendering as the illumination order increases.

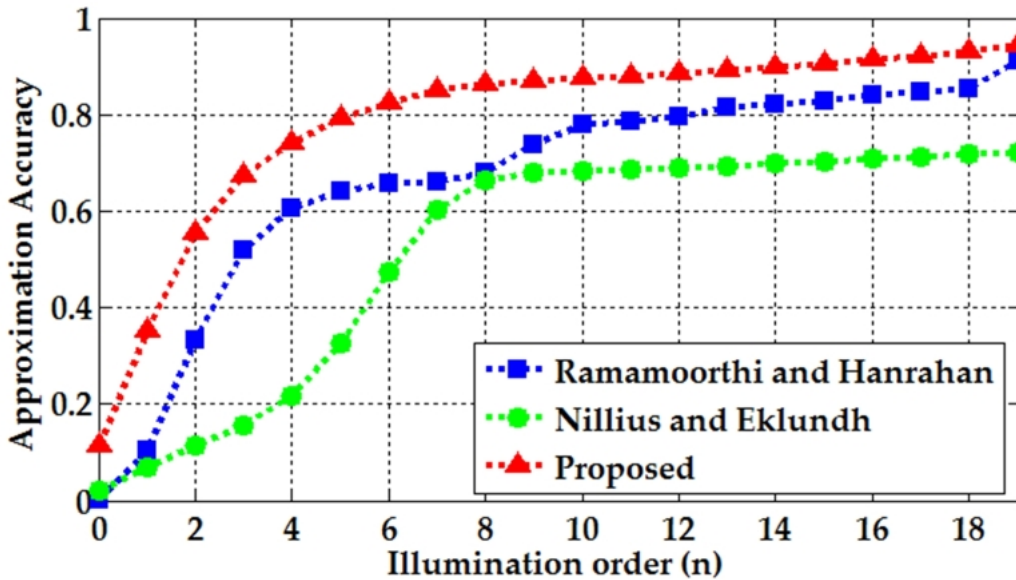
Figure 5.4 portrays the average approximation accuracy of the irradiance harmonics based on the proposed reflectance basis compared to those of Ramamoorthi and Hanrahan in [11, 31] and Nillius ad Eklundh in [32, 44]. The average is computed over the materials in the given BRDF database as a function of the illumination order n where all degrees $m \in [-n, n]$ are summed for visualization purposes. Note that the irradiance harmonics based on the proposed reflectance basis provide higher accuracy at lower illumination orders when compared to the others.

Figs 5.8 and 5.9 compare between the approximation accuracies of irradiance harmonics based on the proposed basis versus Ramamoorthi and Hanrahan in [11, 31] and Nillius ad Eklundh in [32, 44] with different truncating illumination orders N . While, for most of CURET and Merl materials, the spherical basis of [31] and Zernike-based basis of [44] have comparable accuracies, yet the proposed basis shows higher accuracy levels for most of the materials. These observations hold for both diffuse and specular materials.

Figure 5.10 shows the minimum required truncating illumination order to achieve



(a)



(b)

Figure 5.4: Average approximation accuracy of image irradiance as a function of the illumination order n based on (a) CURET BRDF database [10] and (b) Merl BRDF Database [17]. Average is computed over the database materials under the assumption of single distant light source having a non-decayable constant spectrum. Note that the proposed reflectance basis maintains higher irradiance approximation accuracy at lower illumination orders.

specific level of accuracy when using the proposed isotropic basis compared to those of Ramamoorthi and Hanrahan in [11, 31] and Nillius and Eklundh in [32, 44]. It can be noted that irradiance signals of most of CURET materials can be represented with accuracy of 95% with illumination order less than or equal $N = 6$, while anisotropic materials such as ribbed-paper maintain this level of accuracy at $N = 12$, yet other




Figure 5.5: Rendered sphere images of (a) lettuce-leaf and (b) felt BRDFs [10] using single directional light source with increasing illumination order $N = 0, 2, 4, 6, 8$ and increasing reflectance order $P = 0, 2, 4, 6$. Note as the illumination order increase, the rendered image approaches the groundtruth rendering. Note the zero illumination order acts as the DC component of the irradiance signal.

reflectance bases need higher orders.

Figs 5.6 and 5.7 show rendered images of the camel and hippo toys, respectively, of the “Weizmann Photometric Stereo Database” [100]. In Figure 5.6, the image is rendered under a high-frequency illumination map (Galileo Tomp [55]) with a non-specular BRDF (beige-fabric [17]) using irradiance harmonics of illumination order up-to $N = 12$. While in Figure 5.7, the image is rendered under a low-frequency illumination map (Funston Beach [55]) with a specular BRDF (blue-metallic paint [17]) using irradiance harmonics of illumination order up-to $N = 6$. It can be observed that spherical reflectance basis performs poorly at capturing the reflectance appearance while the proposed basis achieved lowest mean absolute error compared to the groundtruth rendering (achieved by path-tracing [30]).

5.7 Summary

This chapter formulated the image irradiance basis of isotropic surfaces under arbitrary distant illumination. Such bases are termed as image irradiance harmonics. The physical compliance of hemispherical basis for representing BRDFs of isotropic surfaces was

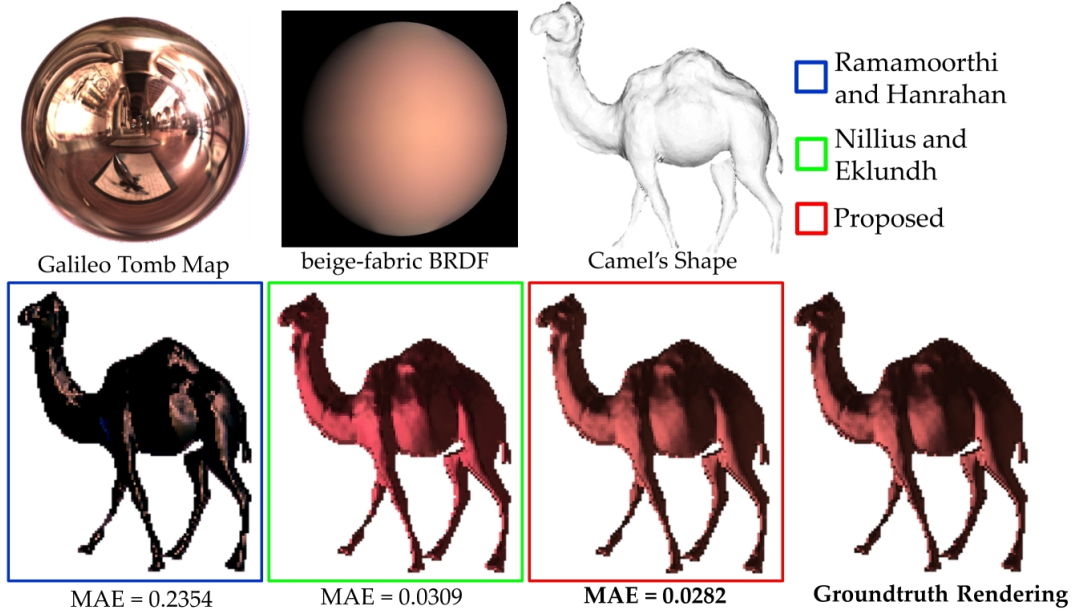


Figure 5.6: Sample of frequency-space rendering: Rendered images for the camel toy from “Weizmann Photometric Stereo Database” [100] with beige-fabric BRDF [17] under Galileo Tomb (high-frequency) illumination map [55] using the proposed reflectance basis (red box) versus that of Ramamoorthi and Hanrahan [31] (blue box) and Nillius and Eklundh [44] (green box) where $N = 12$. Path tracing [30] is used to render a unit sphere with the BRDF at the top-middle as well as the groundtruth rendering at the bottom-left. Mean absolute error (MAE) is shown below each rendered image where all intensities are normalized in the range $[0, 1]$. Note that the proposed basis captures the appearance of the surface reflectance under high-frequency illumination compared to the others.

addressed while satisfying the Helmholtz reciprocity property. The number of basis was proven to be reduced drastically due to isotropy, nonetheless the proposed isotropic Helmholtz reflectance basis contains more basis functions when compared to the Zernike-based ones at a given illumination order. By investigating the energy content maintained by the proposed image irradiance basis for natural materials, one found that specular materials tend to require more basis functions when compared to diffuse ones. Yet the presented basis captures similar cumulative energy content at lower illumination orders when compared to similar representations in literature.

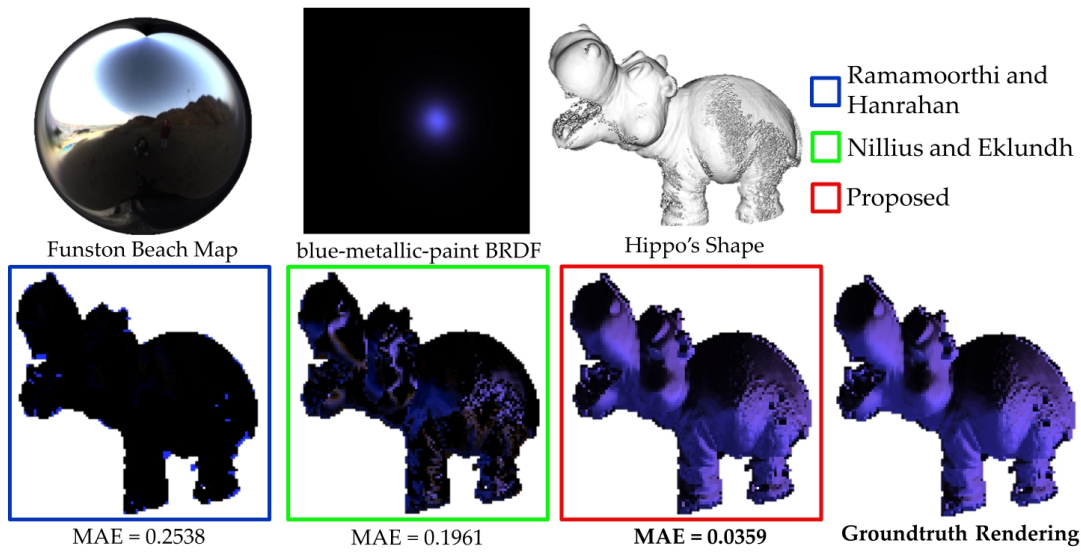


Figure 5.7: Sample of frequency-space rendering: Rendered images for the hippo toy from “Weizmann Photometric Stereo Database” [100] with blue-metallic paint BRDF [17] under Funston Beach (low-frequency) illumination map [55] using the proposed reflectance basis (red box) versus that of Ramamoorthi and Hanrahan [31] (blue box) and Nillius and Eklundh [44] (green box) where $N = 6$. Path tracing [30] is used to render a unit sphere with the BRDF at the top-middle as well as the groundtruth rendering at the bottom-left. Mean absolute error (MAE) is shown below each rendered image where all intensities are normalized in the range $[0, 1]$. Note that the proposed basis captures the appearance of the surface reflectance under low-frequency illumination compared to the others.

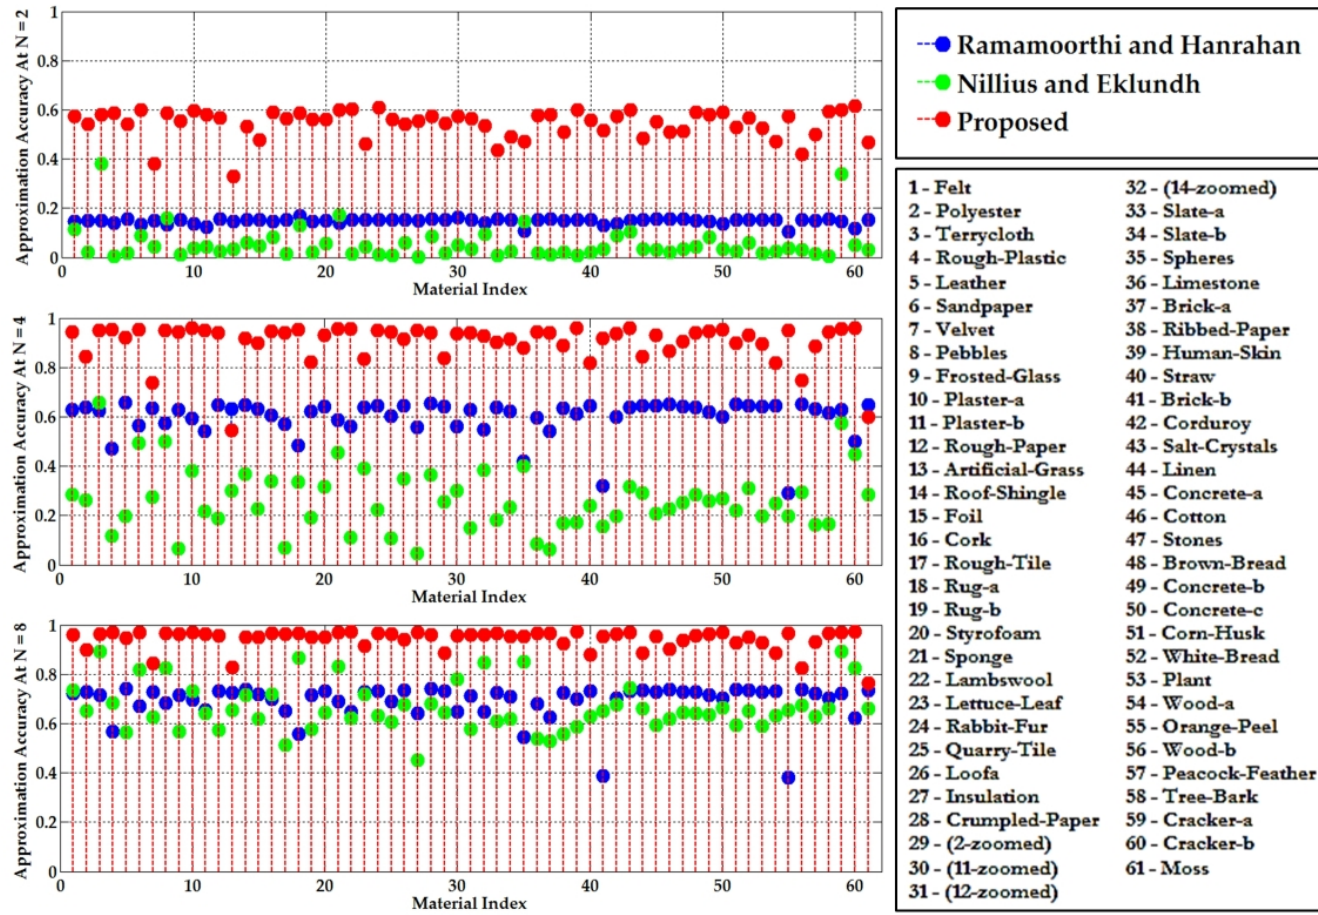


Figure 5.8: The irradiance approximation accuracy of CURET database materials [10] for different truncating illumination orders N . It can be observed that the proposed reflectance basis captures most of the irradiance energy content at lower truncating illumination orders for almost all CURET materials

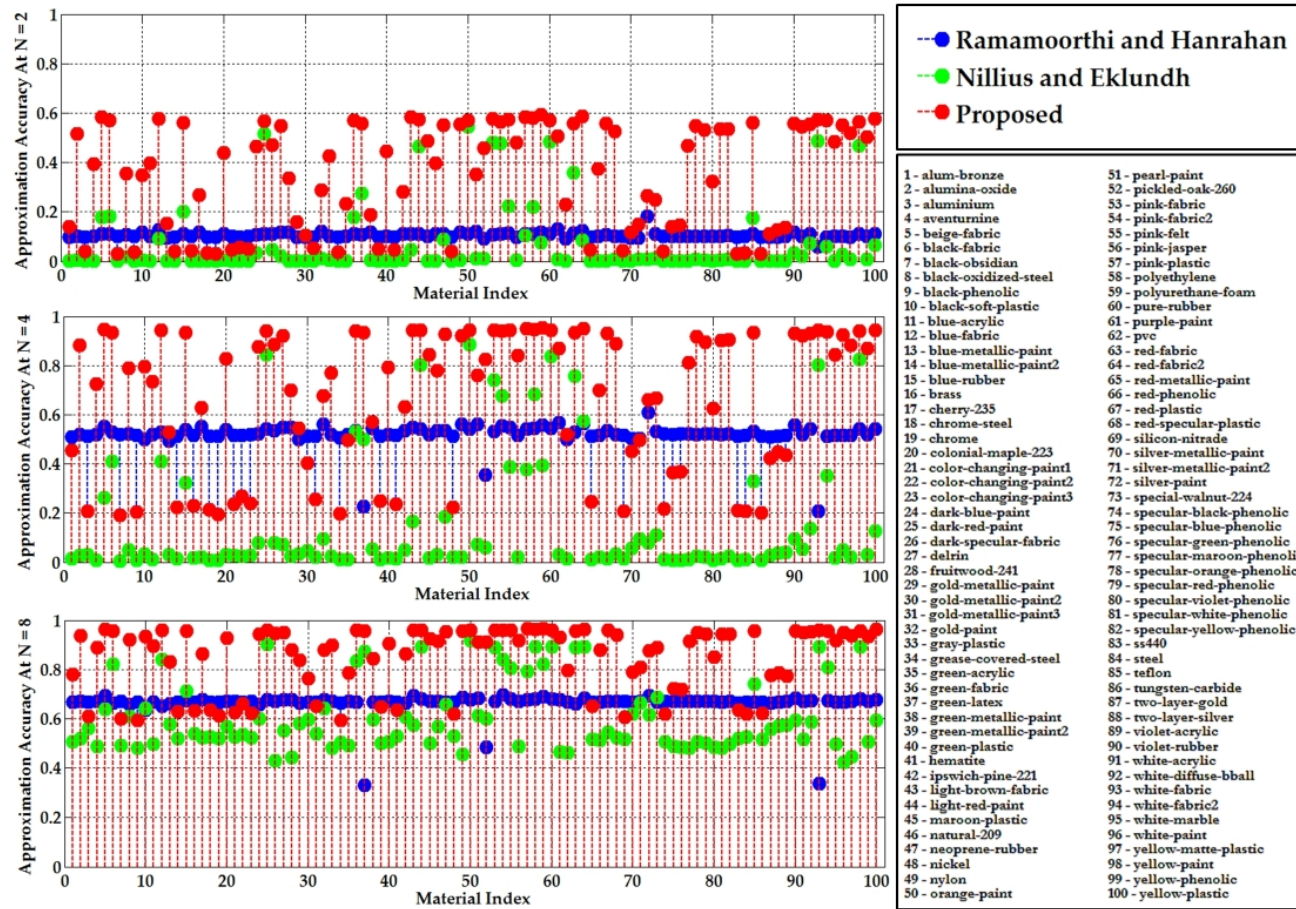


Figure 5.9: The irradiance approximation accuracy of Merl database materials [17] for different truncating illumination orders N . It can be observed that the proposed reflectance basis captures most of the irradiance energy content at lower truncating illumination orders for almost all Merl materials

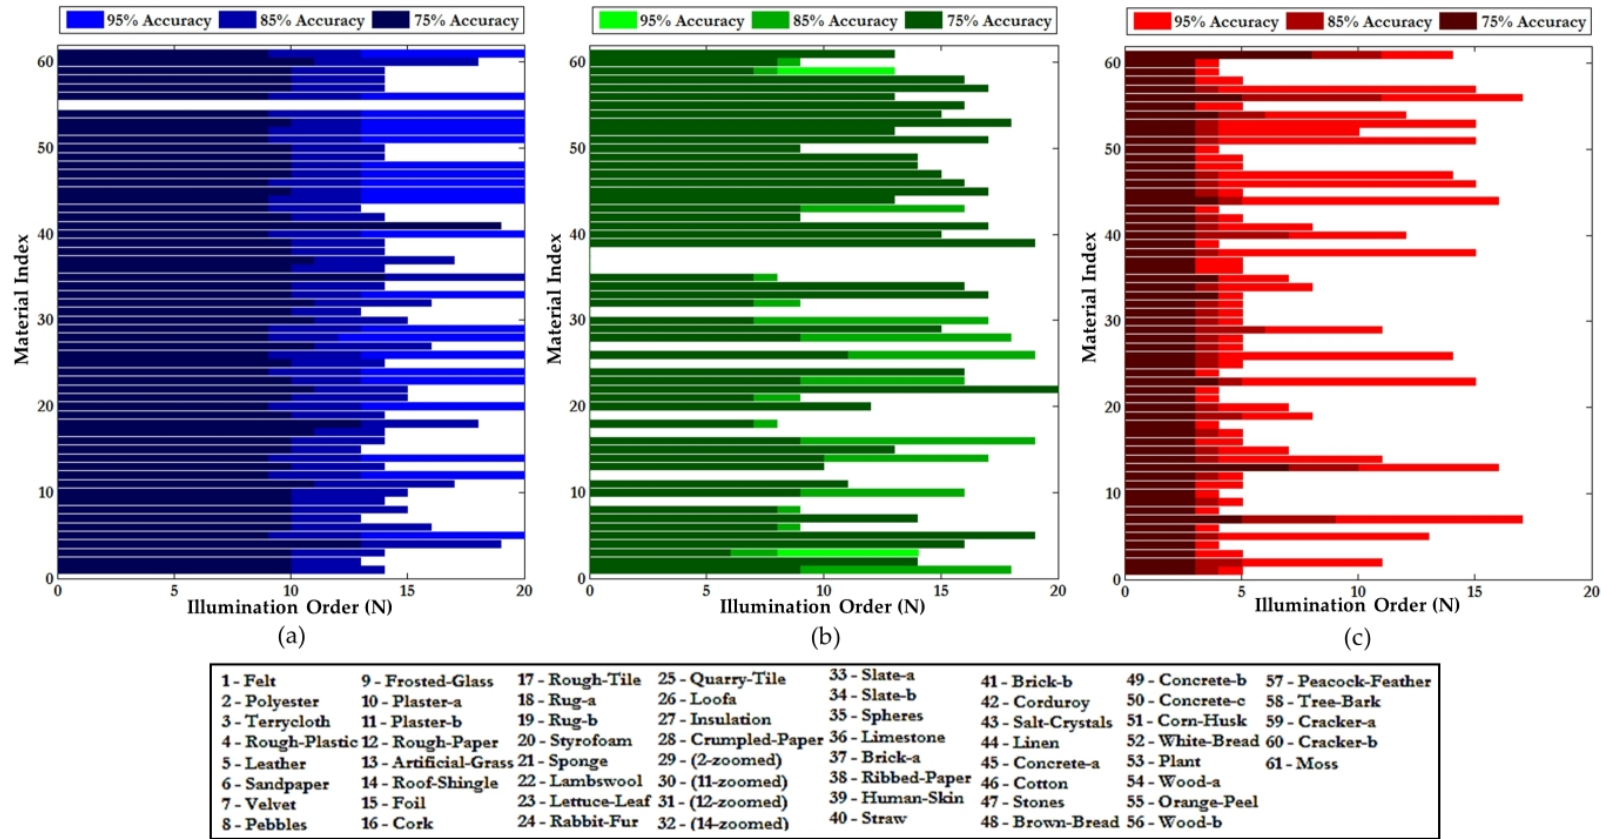


Figure 5.10: Required minimum illumination order for the materials in the CURET database [10] when using the irradiance harmonics based on (a) the proposed *isotropic* basis, (b) spherical basis deployed by Ramamoorthi and Hanrahan in [11, 31] and (c) Zernike-based hemispherical basis deployed by Nillius ad Eklundh in [32, 44]. The dark, medium and light shaded bars show the required illumination order N to achieve irradiance approximation accuracy of 75%, 85% and 95%, respectively. Note that missing bars implies the inability of the respective basis to capture the required accuracy levels at illumination orders less than 20.

ANALYTIC APPEARANCE SUBSPACE CONSTRUCTION

Appearance variation due to illumination changes is an inherent challenge in many vision tasks such as recognition. Applications concerned with inverse rendering would benefit from low-dimensional generative model which captures appearance variations *w.r.t.* illumination conditions and surface reflectance properties. However, conventional subspace construction approaches suffer from the need of “large-enough” image ensemble rendering numerical methods intractable. This chapter proposes an analytic formulation for low-dimensional subspace construction in which shading cues lie while preserving the natural structure of an image sample. Thanks to the frequency-space representation of the image irradiance equation where the process of finding such subspace can be cast as establishing a relation between its principal components and that of a deterministic set of basis functions. This resolves the issue of dimensionality since the source of randomness in the imaging process becomes the irradiance harmonics coefficients rather than the whole image realization. Representing images in their natural dimension, *i.e.* matrices, further lessen the number of parameters to be estimated to define a bilinear projection which maps the image sample to a lower-dimensional bilinear subspace. Since *irradiance harmonics* enables decoupling illumination and reflectance from the underlying geometry and pose; this enables the incorporation of prior information about natural illumination and real world surface materials. Results show significant impact on dimensionality reduction with minimal loss of information as well as robustness against noise.

6.1 Introduction

Various vision tasks are usually confronted by appearance variations due to changes of illumination. For instance, in a recognition system, Moses *et al.* [101] showed that the variability in human face appearance is owed to changes to lighting conditions rather than person’s identity. Theoretically, due to the arbitrariness of the lighting function, the space of all possible images of a fixed-pose object under all possible illumination conditions is infinite dimensional [25]. Nonetheless, Belhumer and Kriegman [23] proved that the set of images of a convex Lambertian surface under distant illumination lies near a low dimensional linear subspace. This result was also extended to include non-Lambertian objects with non-convex geometry. As such, vision applications, concerned with the recovery of illumination, reflectance or surface geometry from images, would benefit from low-dimensional generative model which captures appearance variations *w.r.t.* illumination conditions and surface reflectance properties, enabling the formulation of such inverse problems as parameter estimation.

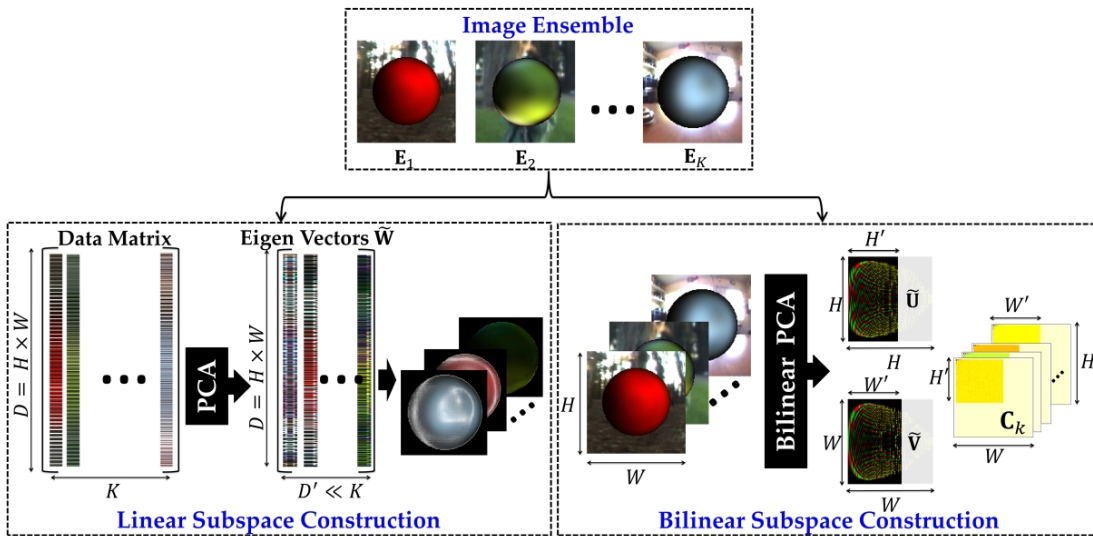


Figure 6.1: Numeric subspace construction typically requires a large set of captured or synthesized images under all possible illuminations and materials.

Typically, subspace construction entails performing a dimensionality reduction scheme, *e.g.* Principal Component Analysis (PCA), on a large set of real or synthesized images of object(s) of interest with fixed pose but different illumination conditions *e.g.* [27], see Figure 6.1. Despite the great success of such numerical construction, this approach has two major problems. First, the acquired or rendered image ensemble should be

statistically significant vis-à-vis capturing the full behavior of the sources of variations that is of interest, in particular illumination and reflectance¹. When the sample data dimension is much higher than the number of samples in the ensemble, PCA is prone to be overfitted to the training ensemble [102]. This is known as the *small sample size* (SSS) problem. Second, the curse of dimensionality hinders numerical methods such as Singular Value Decomposition (SVD) which becomes intractable especially with a large number of large-sized realizations in the image ensemble.

One way to bypass the need for large image ensembles is to construct appearance subspaces using phenomenological models which capture appearance variations through mathematical abstraction of the reflection process, see Figure 6.2. In particular, the harmonic expansion of the image irradiance equation [24] can be used to derive an analytic subspace to represent images under fixed pose but different illumination conditions [6, 25, 44]; it has been shown that surface reflectance acts as a low-pass filter on the incident illumination [4, 5] where the image irradiance equation has been formulated in a convolution framework [6, 7]. This yielded a band-limited reflected light while accounting for complex illumination, thus a point light source assumption can be relaxed.

For specific object geometry under fixed pose, the convolution theory implies a multiplicative framework in the frequency domain where an image is represented as a linear combination of pre-computed basis functions, which are termed as *irradiance harmonics*, that are pose and geometry dependent. This leads to a finite-dimensional linear model to represent the image irradiance under arbitrary illumination. Usually a large number of basis functions is required in order to capture high frequency lighting, sharp shadows and glossy reflectance. Hence irradiance harmonics-based representation, although finite dimensional, would result in a huge number of parameters than can be estimated given shading cues from captured images. This highlights the need for deriving a mechanism to find a lower-dimensional subspace in which shading cues lie.

This chapter takes advantage of the two-fold benefit of the frequency-space representation of the image irradiance equation. First, it decouples the image formation process such that the illumination conditions and surface reflectance characteristics are encoded into the coefficients of the irradiance harmonics. This allows the incorporation of prior information about natural illumination and real world surface materials in the

¹Relying on an image ensemble to be enough to generate statistics has an implicit assumption that the image formation process is an ergodic one.

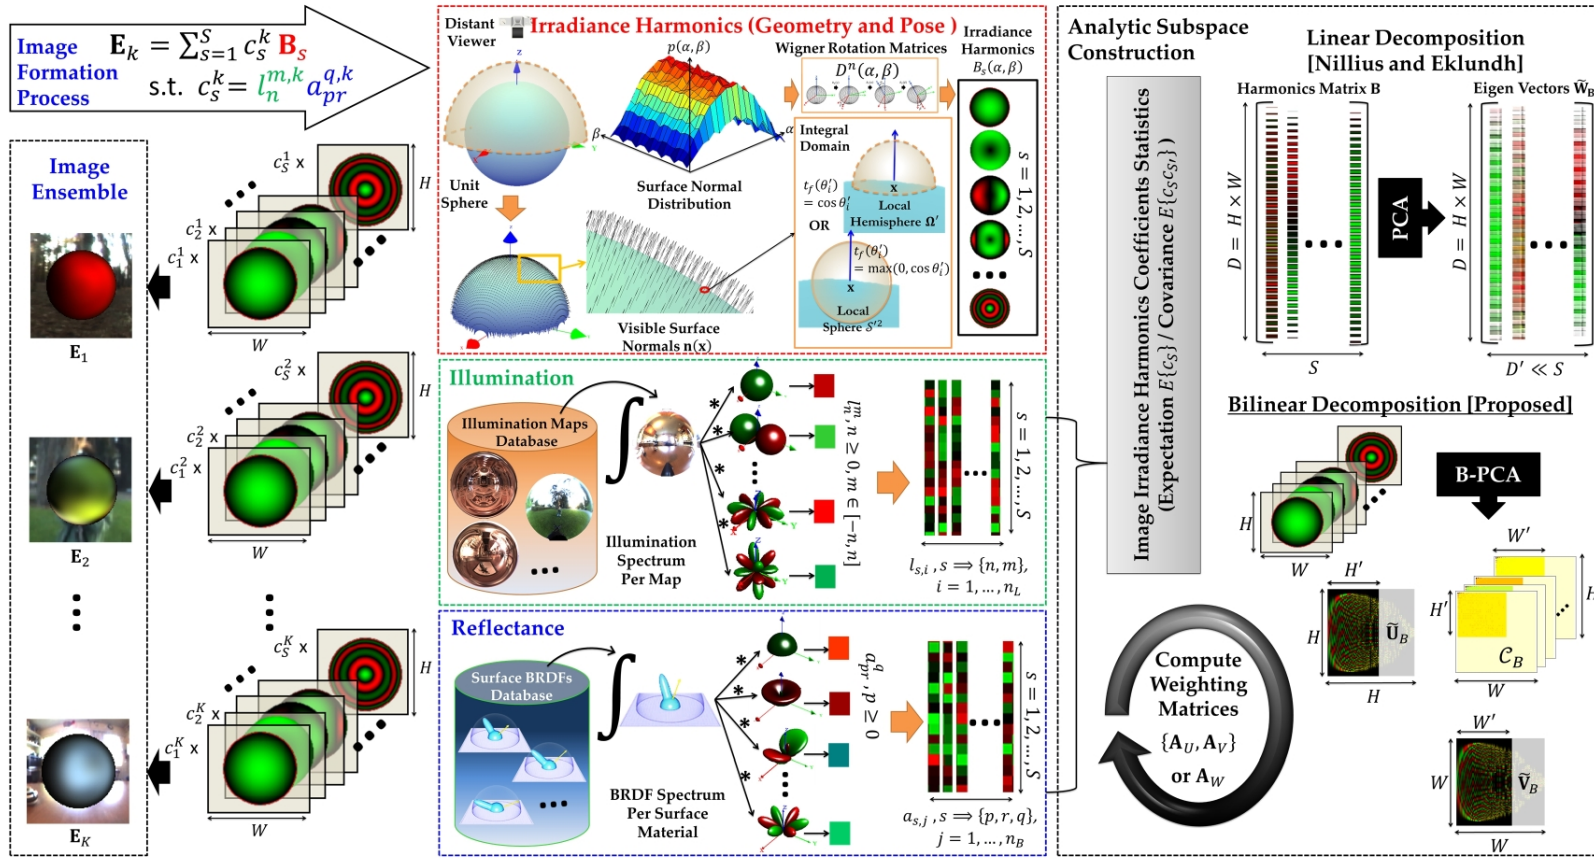


Figure 6.2: Subspace construction to model illumination and reflectance: assuming fixed pose and surface geometry, the frequency space representation of the image irradiance decouples the image formation process such that the illumination and reflectance are encoded into the coefficients c_s^k of the irradiance basis functions \mathbf{B}_s which are geometry and pose dependent. In case of analytic subspace construction (b), the image coefficients c_s^k become the random variable instead of the image realization. Thus the inherent curse of dimensionality numerical methods such as Singular Value Decomposition can be handled.

subspace construction process. Second and more importantly, the process of finding such subspace can be cast as establishing a relation between its principal components and that of the irradiance harmonics. This resolves the issue of dimensionality since the source of randomness in the imaging process becomes the irradiance harmonics coefficients rather than the whole image realization.

PCA, a linear subspace learning algorithm, seeks an optimal linear mapping to a lower-dimensional subspace by representing images as samples in a vector space. This leads to the estimation of a large number of parameters while breaking the natural structure of the two-dimensional image samples [103]. Recently, multilinear subspace learning has attracted the community's attention where images are represented in their natural dimension, second/third order tensorial data, preserving their intrinsic dimensionality. This further lessens the number of parameters to be estimated to define a bilinear projection which maps the image sample to a lower-dimensional bilinear subspace. The presented work is after deriving the subspace of image irradiance in the image space taking into account the two-dimensional structure of samples in the image space rather than representing an image in the vector space to obtain more compact representation with minimal loss of information.

It is worth mentioning that this work shares the work of Vasilescu and Terzopoulos (TensorTextures) [104] in representing appearance variation using multilinear PCA (bilinear as a special case). Nonetheless, in contrast to their work, an analytic construction of the appearance principal components is proposed where all possible illumination configurations are incorporated into such construction by analytically rotating light sources; an inherent benefit of analytical approaches compared to numerical ones where the synthesized or acquired images should exhibit wide variations of illumination conditions. Further, Vasilescu and Terzopoulos dealt with the image sample in the vector space, ignoring the intrinsic structure of the image space. This work studies the spatial relations between image pixels with the effect of illumination and reflectance. As such, one focuses on demonstrating the benefit of considering the image space (bilinear subspace) versus vector space (linear subspace).

This chapter proposes an approach for analytic bilinear subspace construction to capture the full behavior of appearance variation resulting from non-Lambertian surface reflectance when exposed to complex illumination. The spatial correlation of image pixels is taken into account while connecting the spatial constraints to the irradiance

constraints. This implies applying two linear transforms (hence bilinear) to both the left and right sides of the input image irradiance. Whereas the process of finding such subspace is cast as establishing a relation between its principal components and that of the irradiance harmonics. This resolves the issue of dimensionality since the source of randomness in the imaging process becomes the irradiance harmonics coefficients rather than the whole image realization. Natural illumination maps from [105] are analytically rotated, while surface reflectance is modeled using a database of real world surface materials [17]. One focus on isotropic reflectance yielding a parametric representation of image irradiance with smaller number of parameters whilst being a good enough approximation for the general case. Results show the superiority of the proposed approach compared to the analytic linear one, *e.g.* [25, 44], in terms of providing significant decrease in subspace dimensionality while maintaining higher approximation accuracy.

Potential applications for the proposed model involve constructing generative appearance models which can be used to render/synthesize images of a fixed-pose object under new illumination conditions, *i.e.* extrapolate appearance. Further, it can be considered as an economical solution for subspace construction for recognition applications which resolves the trade-off between the small-sample-size problem and numerical estimation of appearance subspace basis. Whereas the proposed model can be used to construct the basis for the illumination cone [23] which provides a mean to predict object appearance under unseen imaging conditions. As such, object recognition can be performed by assigning the identity of the closest illumination cone based on Euclidean distance.

This low-dimensional representation of image irradiance under unknown arbitrary lighting can also be considered as a fundamental process for many computer vision tasks such as illumination modeling [6], surface reflectivity estimation/analysis [11, 32], statistical shape and albedo recovery [33, 34, 35], shape from shading [36], photometric stereo [37, 38], object detection and recognition [39], to name a few.

6.2 Analytic Bilinear Subspace Construction

The main idea of analytic subspace construction is to use the representation in Equation 5.8 directly in an unsupervised subspace learning scheme seeking orthonormal principal components spanning the subspace where shading cues lie. This work derives such

a subspace taking into account the spatial correlation of the image pixels within a localized neighborhood. Matrix-based representation is used in contrast to a vector formed by stacking the image's columns. It is worth mentioning that vector space model admits to a simple and closed form solution for subspace learning using SVD of the matrix whose columns are realizations from the image space. Nonetheless, such computation restricts its applicability to small data size.

In the sequel, notations follow the conventions in the multilinear algebra where vectors are denoted by lowercase boldface letters, *e.g.* \mathbf{x} and matrices by uppercase boldface, *e.g.* \mathbf{X} . Indices are denoted by lowercase letters spanning the range from 1 to the uppercase letter of the index, *e.g.* $i = 1, 2, \dots, I$.

6.2.1 Problem Formulation

Let $\mathbf{E} \in \mathbb{R}^{H \times W}$ be a matrix representation of the image irradiance of the visible surface normals to the viewer such that H denotes height and W denotes width. The objective is to solve for two linear transformations, hence the name *bilinear*, $\tilde{\mathbf{U}} \in \mathbb{R}^{H \times H'}$ and $\tilde{\mathbf{V}} \in \mathbb{R}^{W \times W'}$ which map the image space $\mathbb{R}^{H \times W}$ into a lower-dimensional subspace $\mathbb{R}^{H' \times W'}$ with $H' \leq H$ and $W' \leq W$ where¹,

$$\mathbf{Y} \triangleq \tilde{\mathbf{U}}^T \mathbf{E} \tilde{\mathbf{V}} \quad (6.1)$$

such that this low-dimensional subspace captures most of the variations observed in the image space due to illumination and reflectance.

Let the image space variation be measured by the *total matrix scatter* defined as the variance of the image projection $\mathbf{Y} \in \mathbb{R}^{H' \times W'}$ in the lower-dimensional subspace. Thus the objective is to determine the orthonormal projection matrices which maximize the total scatter matrix,

$$\{\tilde{\mathbf{U}}, \tilde{\mathbf{V}}\} = \underset{\tilde{\mathbf{U}}, \tilde{\mathbf{V}}}{\operatorname{argmax}} \Psi_{\mathbf{Y}} \quad s.t. \quad \Psi_{\mathbf{Y}} = E\{\|\mathbf{Y} - \bar{\mathbf{Y}}\|_F^2\} \quad (6.2)$$

where F denotes the matrix Frobenius norm [106] and $\bar{\mathbf{Y}} = E\{\mathbf{Y}\}$ is the origin of the desired subspace. One needs to point out that this centering is different from [44] where the mean of the image (pixel-centering) rather than the mean image of the data (image-

¹ denotes lower-dimensional.

centering) is subtracted. With image-centering, the joint presence of two images with different appearance has a larger impact in the subspace construction when compared to the impact of two similar images. With pixel-centering, on the other hand, the impact of similar images in a small image ensemble is much pronounced. As such, one opted for image-centering to capture appearance variations which can be used for recognition applications where the contribution of each image is proportional to its deviation from the subspace origin.

6.2.2 Analytic Bilinear PCA Derivation

The optimization problem in Equation 6.2 does not admit to a closed form solution allowing solving for the projection matrices simultaneously. Lu *et al.* [106] used alternating projection method to derive a numerical method for multilinear PCA for M -order tensors. This work derives the analytical counterpart for $M = 2$ of such approach where the irradiance harmonics-based representation makes it possible to avoid acquiring a large ensemble of images under different illumination conditions. Note that, in principle, all possible lighting configurations are required to be taken into account in order to model illumination variation. Hence analytical construction of image subspace allows for analytically rotating light sources over the domain of 3D rotations $SO(3)$.

Consider a specific object defined by its shape (*i.e.* surface normals) under fixed pose relative to the imaging sensor, the irradiance harmonics $\{\mathbf{B}_s\}$ can be viewed as the backbone of its appearance, *i.e.* different appearance is only owed to different illumination conditions and/or reflectance characteristics.

Let $\mathbf{B}_s \in \mathbb{R}^{H \times W}$ be the s -th irradiance harmonics of the visible surface normals be represented as a matrix. Hence it can be decomposed into a core matrix \mathbf{C}_s^B and two orthonormal projection matrices $\tilde{\mathbf{U}}_B \in \mathbb{R}^{H \times H'}$ and $\tilde{\mathbf{V}}_B \in \mathbb{R}^{W \times W'}$ such that,

$$\mathbf{B}_s \triangleq \tilde{\mathbf{U}}_B \mathbf{C}_s^B \tilde{\mathbf{V}}_B^T \quad (6.3)$$

The harmonics projection matrices $\{\tilde{\mathbf{U}}_B, \tilde{\mathbf{V}}_B\}$ (column and row projectors, respectively) and the core matrices $\mathbf{C}_s^B \forall s$ are solved for using Bidirectional PCA (BD-PCA) [102] in an offline stage using all irradiance harmonics up to illumination and reflectance order (N and P). This process can be summarized as follows. Let $\{\mathbf{B}_1, \dots, \mathbf{B}_S\}$ be the set of S -irradiance harmonics. Let the s -th harmonics \mathbf{B}_s be represented as an H -set

of $1 \times W$ row vectors; the row total scatter matrix $\Psi_B^H \in \mathbb{R}^{W \times W}$ is defined as,

$$\Psi_B^H = \frac{1}{S \times H} \sum_{s=1}^S (\mathbf{B}_s - \bar{\mathbf{B}})^T (\mathbf{B}_s - \bar{\mathbf{B}}) \quad (6.4)$$

where $\bar{\mathbf{B}}$ is the mean irradiance harmonics. Therefore $\tilde{\mathbf{V}}_B$ is the matrix containing the row eigenvectors corresponding to the first W' largest eigenvalues of Ψ_B^H . By representing the same harmonics as a W -set of $H \times 1$ column vectors, the column total scatter matrix $\Psi_B^W \in \mathbb{R}^{H \times H}$ is defined as,

$$\Psi_B^W = \frac{1}{S \times W} \sum_{s=1}^S (\mathbf{B}_s - \bar{\mathbf{B}}) (\mathbf{B}_s - \bar{\mathbf{B}})^T \quad (6.5)$$

where $\tilde{\mathbf{U}}_B$ is the first H' column eigenvectors corresponding to the first largest H' eigenvalues of Ψ_B^W .

BD-PCA reduces the dimensionality in both column and row directions while numerically overcomes the overfitting problem in contrast to PCA. The immediate benefit of this step is the determination of the lower dimensions H' and W' based on the variations inherited from a deterministic ensemble of matrices, *i.e.* the irradiance harmonics themselves. It is important to emphasize the computational advantage of BD-PCA over PCA where BD-PCA requires $H \times W \times H' + H \times H' \times W'$ multiplications while PCA requires $H \times W \times D'$ multiplications for $D' \leq H \times W$, $H' \leq H$ and $W' \leq W$ [102]. In addition, such decomposition is performed once independent of any color-channel (wavelength). As such the image irradiance of the visible surface normals can be rewritten as,

$$\mathbf{E} = \tilde{\mathbf{U}}_B \left(\sum_s c_s \mathbf{C}_s^B \right) \tilde{\mathbf{V}}_B^T \quad (6.6)$$

According to the linearity property of the expectation operator, the origin of the image irradiance subspace can be given by,

$$\bar{\mathbf{Y}} = E\{\mathbf{Y}\} = E\left\{ \tilde{\mathbf{U}}^T \mathbf{E} \tilde{\mathbf{V}} \right\} = \tilde{\mathbf{U}}^T \tilde{\mathbf{U}}_B \left(\sum_s E\{c_s\} \mathbf{C}_s^B \right) \tilde{\mathbf{V}}_B^T \tilde{\mathbf{V}} \quad (6.7)$$

where the harmonics coefficients c_s , encoding light and surface BRDF, becomes the random variable rather than the image itself.

In order to simplify the process of finding the projection matrices $\{\tilde{\mathbf{U}}, \tilde{\mathbf{V}}\}$ which span the image space of the object of interest, one establishes a relation between the

orthonormal columns of $\tilde{\mathbf{U}}$ and $\tilde{\mathbf{V}}$ and those of $\tilde{\mathbf{U}}_B$ and $\tilde{\mathbf{V}}_B$, respectively, such that,

$$\tilde{\mathbf{U}}^T = \mathbf{A}_U \tilde{\mathbf{U}}_B^T \quad , \quad \tilde{\mathbf{V}}^T = \mathbf{A}_V \tilde{\mathbf{V}}_B^T \quad (6.8)$$

where $\mathbf{A}_U \in \mathbb{R}^{H' \times H'}$ and $\mathbf{A}_V \in \mathbb{R}^{W' \times W'}$ are the weighting matrices encoding the relation of principal components spanning the row and column subspaces of the image space with those spanning the row and column subspaces of the irradiance harmonics. This renders another benefit of the analytic construction where the optimization problem in Equation 6.2 will be solved for less number of parameters, *i.e.* $H' \times H' + W' \times W'$ rather than $H \times H' + W \times W'$ with $H' \leq H$ and $W' \leq W$. It is worth noting that the decomposition of the irradiance harmonics accounts for dimensionality reduction whereas the weighting matrices relate the principal components of the row and column subspaces to the image formation model while encoding prior information about illumination and reflectance through the expectation operator.

From Equation 6.8 and Equation 6.6, the image projection can be simplified to,

$$\mathbf{Y} = \mathbf{A}_U \left(\sum_s c_s \mathbf{C}_s^B \right) \mathbf{A}_V^T \quad (6.9)$$

where the solution for $\{\mathbf{A}_U, \mathbf{A}_V\}$ matrices is given by the following theorem.

Theorem 6.1 (Analytic Bilinear PCA). *Let $\{\mathbf{A}_U, \mathbf{A}_V\}$ be the solution to Equation 6.2. Then, based on Alternating Least Squares [107], given \mathbf{A}_V , the weighting matrix \mathbf{A}_U consists of the H' -eigenvectors of Υ_V^T such that,*

$$\Upsilon_V = \sum_s \sum_{s'} E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} \mathbf{C}_s^B \mathbf{A}_V^T \mathbf{A}_V (\mathbf{C}_{s'}^B)^T \quad (6.10)$$

and given \mathbf{A}_U , the weighting matrix \mathbf{A}_V consists of the W' -eigenvectors of Υ_U^T such that,

$$\Upsilon_U = \sum_s \sum_{s'} E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} (\mathbf{C}_s^B)^T \mathbf{A}_U^T \mathbf{A}_U \mathbf{C}_{s'}^B \quad (6.11)$$

where $E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} = E\{c_s c_{s'}\} - E\{c_s\}E\{c_{s'}\}$.

Proof. The proof is given in [Appendix F: Proof of Analytic Bilinear PCA Theorem](#). \square

Since the optimization of both matrices are inter-dependent, this does not admit to a closed form solution. Rather an iterative procedure can be devised to solve for the

weighting matrices. Truncated full projection [106] can be used to initialize this iterative scheme, nonetheless in the presented framework allows for analytic full projection.

In case of full projection, one has $H' = H$ and $W' = W$. In order not to abuse notation, let $(\cdot)^*$ denote matrices with the full dimensions H and W as opposed to $\tilde{(\cdot)}$. Thus \mathbf{U}^* and \mathbf{V}^* become square matrices where $\mathbf{U}^{*-1} = \mathbf{U}^{*T}$ and $\mathbf{V}^{*-1} = \mathbf{V}^{*T}$. As a result one has,

$$\mathbf{A}_V^{*T} \mathbf{A}_V^* = (\mathbf{V}^{*T} \mathbf{V}_B^*)^T \mathbf{V}^{*T} \mathbf{V}_B^* = \mathbf{V}_B^{*T} \mathbf{V}^* \mathbf{V}^{*T} \mathbf{V}_B^* = \mathbf{I}_W \quad (6.12)$$

where \mathbf{I}_W is an identity matrix of size $W \times W$. Similar argument for $\mathbf{A}_U^{*T} \mathbf{A}_U^* = \mathbf{I}_H$. Thus Υ_V^* and Υ_U^* becomes independent of the weighting matrices such that,

$$\Upsilon_V^* = \sum_s \sum_{s'} E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} \mathbf{C}_s^{*B} (\mathbf{C}_{s'}^{*B})^T \quad (6.13)$$

and,

$$\Upsilon_U^* = \sum_s \sum_{s'} E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} (\mathbf{C}_s^{*B})^T \mathbf{C}_{s'}^{*B} \quad (6.14)$$

The optimal \mathbf{A}_U^* and \mathbf{A}_V^* are then obtained, directly with no iterations, as the matrices containing the eigenvectors of Υ_V^{*T} and Υ_U^{*T} , respectively. While the total scatter in the image space would be fully captured, there is no dimensionality reduction gained using full projection. Hence, the first H' and W' columns of the full projection matrices corresponding to the largest eigenvalues of Υ_V^{*T} and Υ_U^{*T} are kept to provide initial weighting matrices \mathbf{A}_U and \mathbf{A}_V , respectively.

6.2.3 Model-based Bilinear PCA

The major advantage of analytic subspace construction is the explicit relation between the principal components spanning the image space and the illumination and reflectance coefficients, allowing for a model-based framework for generic subspace generation. Further, the rotation of the lighting function can be done analytically, reducing the need for acquired data drastically.

Assuming that the lighting function and the surface material/reflectance are independent, one would have [44],

$$E\{c_s\} = E\{l_n^m\} E\{a_{pr}^q\} \quad , \quad E\{c_s c_{s'}\} = E\{l_n^m l_{n'}^{m'}\} E\{a_{pr}^q a_{p'r'}^{q'}\} \quad (6.15)$$

where the respective indices are given by the ordering function of the irradiance harmonics functions. Nillius and Eklundh [44] derived a closed form for the expectations and covariances in Equation 6.15. As such, databases of illumination maps, *e.g.* [55, 105] and real world materials, *e.g.* [10, 17] can be devised to incorporate prior information in a model-based framework for analytic subspace construction. The proposed approach is summarized in Algorithm 2.

6.3 Connection with Analytic PCA

Nillius and Eklundh in [44] proposed an analytic PCA framework which depends on vector spaces. It can be seen that the proposed analytical solution generalizes the case of vector spaces in [44] where the image is considered as a vector $\mathbf{e} \in \mathbb{R}^D$ with $D = H \times W$. The objective is then to find orthonormal projection matrix $\tilde{\mathbf{W}} \in \mathbb{R}^{D \times D'}$ which maps the original vector space \mathbb{R}^D into a vector subspace $\mathbb{R}^{D'}$ with $D' \leq D$ where $\mathbf{y} = \tilde{\mathbf{W}}^T \mathbf{e}$. Thus the projection matrix $\tilde{\mathbf{W}}$ maximizes the total vector scatter such that,

$$\tilde{\mathbf{W}} = \underset{\tilde{\mathbf{W}}}{\operatorname{argmax}} \Psi_{\mathbf{y}} \quad \text{s. t.} \quad \Psi_{\mathbf{y}} = E\{\|\mathbf{y} - \bar{\mathbf{y}}\|_F^2\} \quad (6.16)$$

Let the s -th irradiance harmonics be represented in the vector space as $\mathbf{b}_s \in \mathbb{R}^D$ which can be written as $\mathbf{b}_s = \tilde{\mathbf{W}}_B \mathbf{c}_s^B$ with $\tilde{\mathbf{W}}_B \in \mathbb{R}^{D \times D'}$ and $\mathbf{c}_s^B \in \mathbb{R}^{D'}$. The optimal $\tilde{\mathbf{W}}_B$ is determined from the D' -eigenvectors of $\mathbf{B} = [\mathbf{b}_1 \dots \mathbf{b}_S]$ corresponding to its largest D' -eigenvalues. Relating the principal components of the image space to that of the irradiance harmonics in the vector space results in $\tilde{\mathbf{W}}^T = \mathbf{A}_W \tilde{\mathbf{W}}_B^T$ where $\mathbf{A}_W \in \mathbb{R}^{D' \times D'}$. The solution of this matrix is given by the following theorem.

Theorem 6.2 (Analytic PCA). *Let \mathbf{A}_W be the solution to Equation 6.16. Then, the weighting matrix \mathbf{A}_W consists of the D' -eigenvectors of Υ^T such that,*

$$\Upsilon = \sum_s \sum_{s'} E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} \mathbf{c}_s^B (\mathbf{c}_{s'}^B)^T \quad (6.17)$$

where $E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} = E\{c_s c_{s'}\} - E\{c_s\}E\{c_{s'}\}$.

Proof. The proof is given in [Appendix G: Proof of Analytic PCA Theorem](#). □

Comparing the residual of approximating the image matrix \mathbf{E} in the vectorized form by the low dimensional subspace $\tilde{\mathbf{W}}$ with the matrix form by a low dimensional row

Algorithm 2 Model-based Analytic Bilinear PCA

INPUT:

- (a) Visible surface normals $\vec{n}(\mathbf{x}) = (\alpha, \beta)$ where $\mathbf{x} = \{(x, y) : 1 \leq x \leq W, 1 \leq y \leq H\}$ and $\alpha \in [0, \pi]$ and $\beta \in [0, 2\pi]$,
- (b) Illumination map dataset $L_i^k(\vec{\omega}_i)$ with $k = 1, 2, \dots, n_L$, (c) Maximum illumination order N ,
- (d) Surface materials dataset $f_r^j(\vec{\omega}'_i, \vec{\omega}'_o)$ with $j = 1, 2, \dots, n_B$, (e) Maximum reflectance order P ,
- (f) Q : percentage of eigenvalues to be maintained in subspace construction,
- (g) Type of reflectance basis functions, *e.g.* Helmholtz HSH-based basis.

OUTPUT: Subspace projection matrices $\{\tilde{\mathbf{U}}, \tilde{\mathbf{V}}\}$.

1: Build illumination prior:

- (a) Compute illumination spectrum $l_{n,k}^m$ for $n \in [0, N], m \in [-n, n]$ and $k = 1, 2, \dots, n_L$.
- (b) Compute illumination expectation $E\{l_n^m\}$ and covariances $E\{l_n^m l_{n'}^{m'}\}$ as in [44].

2: Build reflectance prior:

- (a) Compute BRDF spectrum $a_{pr,j}^q$ for $p \in [0, P]$ for $j = 1, 2, \dots, n_B$ while $\{r, q\}$ are determined according to the specified reflectance basis functions.
- (b) Compute BRDF expectation $E\{a_{pr}^q\}$ and covariances $E\{a_{pr}^q a_{p'r'}^{q'}\}$ as in [44].

3: Compute irradiance harmonics:

- (a) Compute surface normal distribution $p(\alpha, \beta)$.
- (b) Compute Wigner rotation matrices $\mathcal{D}^n(\alpha, \beta)$ where $n \in [0, N]$ as in [61].
- (c) Compute irradiance harmonics $\mathcal{B}_{npr}^{mq}(\alpha, \beta)$ for $n \in [0, N]$ and $p \in [0, P]$ and sort them according to their average power content to obtain $\mathcal{B}_s(\alpha, \beta)$ for $s = 1, 2, \dots, S$.
- (d) Compute harmonics coefficients expectations $E\{c_s\}$ and covariances $E\{c_s c_{s'}\}$ where indices are given by $s \leftrightarrow \{n, m, p, r, q\}$.
- (e) Decompose the irradiance harmonics based on BD-PCA [102] with $Q\%$ eigenvalues to be maintained to obtain $\tilde{\mathbf{U}}_B \in \mathbb{R}^{H \times H'}$, $\tilde{\mathbf{V}}_B \in \mathbb{R}^{W \times W'}$ and $\mathbf{C}_s^B \in \mathbb{R}^{H' \times W'}$ for $s = 1, 2, \dots, S$.

4: Initialization:

- (a) Decompose the irradiance harmonics based on BD-PCA [102] with $Q = 100\%$ to obtain full projection core matrices $\mathbf{C}_s^{*B} \in \mathbb{R}^{H \times W}$ for $s = 1, 2, \dots, S$.
- (b) Compute initial weighting matrix \mathbf{A}_U as the H' -eigenvectors of Υ_V^{*T} in Equation 6.13.
- (c) Compute initial weighting matrix \mathbf{A}_V as the W' -eigenvectors of Υ_U^{*T} in Equation 6.14.

5: repeat

- 6: Compute the weighting matrix \mathbf{A}_U as the H' -eigenvectors of Υ_V^T in Equation 6.10.
- 7: Compute the weighting matrix \mathbf{A}_V as the W' -eigenvectors of Υ_U^T in Equation 6.11.
- 8: Compute $\Psi_{\mathbf{Y}}^i = \min\{\text{trace}(\mathbf{A}_U \Upsilon_V \mathbf{A}_U^T), \text{trace}(\mathbf{A}_V \Upsilon_U \mathbf{A}_V^T)\}$ where i is the iteration index.
- 9: **until** $(\Psi_{\mathbf{Y}}^i - \Psi_{\mathbf{Y}}^{i-1}) / \Psi_{\mathbf{Y}}^{i-1} < \eta$

10: Compute $\tilde{\mathbf{U}}^T = \mathbf{A}_U \tilde{\mathbf{U}}_B^T$ and $\tilde{\mathbf{V}}^T = \mathbf{A}_V \tilde{\mathbf{V}}_B^T$.

and column subspaces $\{\tilde{\mathbf{U}}, \tilde{\mathbf{V}}\}$, the former case is considered a special case where the column-mode projection matrix is identity such that,

$$\mathbf{e} = \tilde{\mathbf{W}}\mathbf{y} \iff (\tilde{\mathbf{I}} \otimes \tilde{\mathbf{W}})\mathbf{y} = \mathbf{e} \quad \text{v. s.} \quad \mathbf{E} = \tilde{\mathbf{U}}\mathbf{Y}\tilde{\mathbf{V}}^T \iff (\tilde{\mathbf{V}} \otimes \tilde{\mathbf{U}})\mathbf{y} = \mathbf{e} \quad (6.18)$$

where \otimes is the Kronecker product [108] of two matrices, $\tilde{\mathbf{I}}$ is the identity matrix and the equivalence is due to a property of the Kronecker product [109]. One can see the difference between standard PCA and the proposed bilinear approach where the latter encapsulates the structure of the image irradiance in the row and column dimensions while jointly models different types of spatial structures presented in the irradiance signal. In the former case, the image data is projected in the row-mode only while the projection in the column-model is ignored. This results in poor expressiveness since the spatial interaction between image pixels are ignored.

6.4 Time and Space Complexities

In this work, the main objective is to alleviate the expensive computations of linear subspaces in case of vector spaces where SVD is considered the main ingredient, whereas bilinear transformation leads to lower computational cost in comparison to SVD. Further, with matrix-based representation, such bilinear subspace is able to capture intrinsic local properties in the image irradiance signal which leads to good classification performance for recognition applications [110]. Nonetheless, it is important to analyze the time and space complexity of the proposed bilinear subspace construction in comparison to the linear one.

6.4.1 Compression Ratio

For large and high dimensional datasets, the lack of available space becomes a critical issue, as such compression ratio becomes an important factor in applications such as recognition and image retrieval.

Consider a set of K -images in the vector space $\{\mathbf{e}_k\}_{k=1}^K$. Since each $\mathbf{e}_k \in \mathbb{R}^D$ is approximated by $\tilde{\mathbf{W}}\mathbf{y}_k$ where $\tilde{\mathbf{W}} \in \mathbb{R}^{D \times D'}$ is common for all images, one needs to keep $\tilde{\mathbf{W}}$ and $\{\mathbf{y}_k\}_{k=1}^K$ for all approximations where $\mathbf{y}_k \in \mathbb{R}^{D'}$. Hence this requires $D \times D' + K \times D' = (D + K) \times D'$ scalars to store the reduced representation. As such

the compression ratio (storage saved) using linear representation is $\frac{K \times D}{(D+K) \times D'}$.

In case of bilinear representation, the matrices $\tilde{\mathbf{U}} \in \mathbb{R}^{H \times H'}$ and $\tilde{\mathbf{V}} \in \mathbb{R}^{W \times W'}$ and $\{\mathbf{Y}_k\}_{k=1}^K$ can be used to recover the original image set $\{\mathbf{E}_k\}_{k=1}^K$ where each $\mathbf{E}_k \in \mathbb{R}^{H \times W}$ is approximated by $\tilde{\mathbf{U}}\mathbf{Y}_k\tilde{\mathbf{V}}^T$. Hence to store the bilinear representation, one needs $H \times H' + W \times W' + K \times (H' \times W')$ scalars with compression ratio $\frac{K \times H \times W}{H \times H' + W \times W' + K \times (H' \times W')}$, where $H' \ll D'$, $W' \ll D'$ and $H \times W = D$.

For instance, for a 100×100 gray-level image and 20 principal components, a linear representation would require $10,000 \times 20 + 20 = 200,020$ scalars, while the bilinear one would require $100 \times 20 + 100 \times 20 + 20 \times 20 = 4400$ scalars. As such, the latter presents a significant impact on dimensionality reduction while capturing the intrinsic properties of an image sample. Note that as H' and W' decrease, the compression ratio increases, yet small values of H' and W' may lead to low of the intrinsic information in the original image sample.

6.4.2 Time Complexity

The most computationally expensive steps in the bilinear subspace construction is the formation of the matrices Υ_V and Υ_U in steps 6 and 7 in Algorithm 2 where it takes $\mathcal{O}(S^2 \times H' \times W'(H' + W'))$ time for their computations for S -irradiance harmonics. Nonetheless the key to the low space complexity of the proposed bilinear representation is that the formation of the matrices Υ_V and Υ_U is based on the lower-rank decomposition of the irradiance harmonics.

The computation time for the bilinear representation \mathbf{Y} is $\mathcal{O}(H' \times W(H + W'))$ in comparison to $\mathcal{O}(D \times D')$ in case of linear representation where $H' \ll D'$, $W' \ll D'$ and $H \times W = D$.

Note that the proposed algorithm involves two eigen problems of sizes $H' \times H'$ and $W' \times W'$ while the linear counterpart involves an eigen problem of size $D' \times D'$ where $H' \ll D'$ and $W' \ll D'$. Further, the decomposition of irradiance harmonics in the linear case needs to solve an eigen problem of size $D \times D$ where $D = H \times W$ while the bilinear case needs two eigen problems with sizes $H \times H$ and $W \times W$. As such, the proposed bilinear representation is computationally more efficient when compared to the linear one.

6.5 Experimental Results

In the sequel, surface reflectance is modeled using the database provided by Mitsubishi Electric Research Laboratories (Merl) [17] which represent a wide variety of surface materials with different diffuse and specular reflection properties ($n_B = 100$). The BRDF measurements are fit up to $P = 8$ to (1) *spherical* harmonics basis [7], (2) *hemispherical* Zernike-based basis [44] and (3) the proposed isotropic *hemispherical* harmonics (HSH)-based Helmholtz reflectance basis. The main difference between the three types of basis is modeling the dependency of the surface BRDF *w.r.t.* the polar coordinates $\{\theta'_i, \theta'_o\}$ where associated Legendre polynomials is used in (1) while Zernike polynomials and shifted associated Legendre polynomials in (2) and (3), respectively.

We compute the irradiance harmonics as in Equation 5.4 for the visible part of a unit sphere, nonetheless, this analysis is applicable to any other geometrical structure. One uses illumination order up to $N = 10$ and reflectance order up to $P = 8$ from which S -harmonics of the highest average power content are selected where S is chosen such that at least 70% of the cumulative power content is maintained. The three types of harmonics¹ are decomposed using BD-PCA [102] to obtain their orthonormal projection matrices spanning the row and column subspaces of the respective harmonics, *i.e.* $\tilde{\mathbf{U}}_B, \tilde{\mathbf{V}}_B$ and their corresponding core matrices \mathbf{C}_s^B . In addition to their corresponding vector subspaces $\tilde{\mathbf{W}}_B^T$ along with \mathbf{c}_s^B .

In the following set of experiments, testing images are rendered using an out-of-training scenario where the illumination maps from sIBL [105] (including a wide variety of illumination conditions such as indoor, outdoor, summer and winter conditions) are used in training the illumination prior ($n_L = 54$) whereas ten illumination maps from [55] are used to render testing images with randomly drawn views. Further, the BRDF that is used to render a testing image is excluded from training the reflectance prior.

In order to assess the representation accuracy of the proposed appearance subspace, testing images are rendered, in a similar manner as in [31], of a sphere for the 100 materials in the Merl BRDF database under 10 natural illumination maps from [55] where for each map 10 randomly drawn views were used resulting in 100 testing images per material.

Given an image with the same geometry but under unknown natural illumination

¹According to the devised reflectance basis

and reflectance. The trained appearance subspace is fit to such an image using number of principal components which maintain $Q\%$ of harmonics variation. One assesses the reconstruction accuracy in accordance to an error estimator which is based on the Mean Absolute Error (MAE) between the testing image and the reconstructed one where color values are normalized in the range $[0, 1]$.

6.5.1 Effect of Initialization

The iterative scheme implied by Theorem 6.1 raises the concern of whether its convergence is dependent on the initialization step. Both random matrices and pseudo (truncated) identity matrices are tested along with truncated full projection. Despite different initializations, one found that the objective function $\Psi_{\mathbf{Y}}$ converged to the same point within no more than three iterations. Nonetheless, using truncated full projection affects the speed of convergence where only one iteration was needed. This result is consistent with that of Lu *et al.* in [106]. In the following, pseudo (truncated) identity matrices are used for initialization while the appearance bilinear subspace is trained using three iterations.

6.5.2 Effect of Irradiance Harmonics

Figure 6.3 shows the average reconstruction errors for each material in the Merl BRDF database based on linear and bilinear subspace construction. It can be noted that bilinear subspace, generally, attain lower error levels for all surface materials when compared to the linear one. This highlights the ability of bilinear representation to encode the intrinsic spatial properties of an image sample compared to the linear (vector) one. Further, it can be observed that HSH-based irradiance harmonics (diamond) provides minimal reconstruction error compared to SH-based (circle) and Zernike-based ones (square). This emphasizes the importance of accounting for the physical properties of non-emitting surfaces where a surface point receives incident illumination from the incoming hemisphere oriented by the surface normal at that point. In addition, the spectrum based on HSH-based basis captures more BRDF energy content compared to that of Zernike-based ones. This is based on the fact that at a specific illumination order, the set of associated Legendre polynomials is distinguished by the property that it contains a polynomial for every combination of order and degree [82], compared to Zernike polynomials which are restricted to even differences between polynomial order

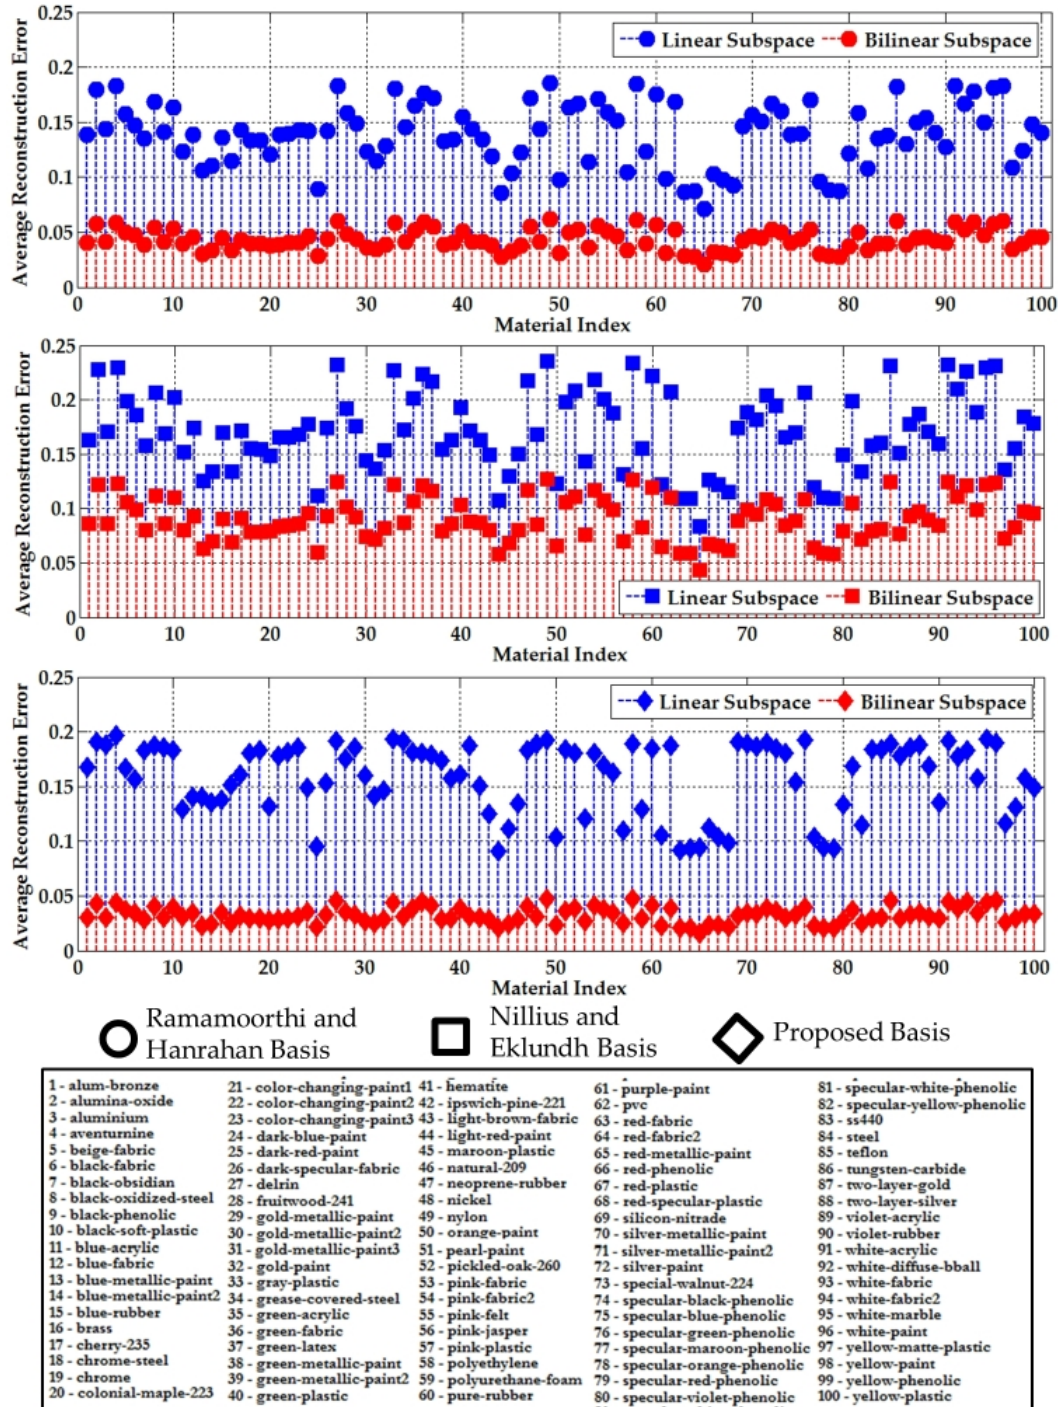


Figure 6.3: The average reconstruction error for each material in the Merl database based on the 100 rendered testing images per material where subspaces were constructed based spherical and hemispherical irradiance harmonics (refer to the legend) where $Q = 98\%$. Note that bilinear subspaces surpass the linear ones while the proposed HSH-based irradiance harmonics attains minimal reconstruction errors for all the surface materials.

and degree.

6.5.3 Effect of $Q\%$

An important factor which affect the subspace dimensionality is the percentage of harmonic variation Q being maintained when the irradiance harmonics is decomposed (using PCA or BD-PCA in case of linear or bilinear subspace, respectively). While lower values of Q would inhibit capturing appearance variation due to illumination and reflectance, higher values would result in model over-fitting. Typical values lie in the range $95\% \pm 4\%$.

Figure 6.4 portrays the average reconstruction error as a function of maintained harmonic variation percentage for linear and bilinear subspaces, respectively, with range $Q \in [85\%, 99\%]$. It can be noted that a bilinear subspace can capture the appearance accurately with lower harmonic variation percentage (*i.e.* lower-dimensional subspace) compared to the linear one regardless of the basis used to model surface reflectance.

Figures 6.5, 6.6 and 6.7 show sample reconstructions for different toys from “Weizmann Photometric Stereo Database” [37] using different BRDFs and illumination maps [55]. It is worth noting that the proposed bilinear appearance representation is capable of capturing the true appearance when compared to the linear one regardless of the percentage of harmonic variations being maintained. Also Figure 6.5 shows an example of over-fitting when the reconstruction error increases at $Q = 98\%$. Note that neither the testing BRDF nor the testing illumination map is included in the subspace construction process.

6.5.4 Effect of Noise

To test the robustness of the proposed subspace *w.r.t.* to noise, white Gaussian noise is added with different signal-to-noise ratio (SNR) levels to each testing image. To conduct a fair assessment, one has fixed the harmonic variation that each subspace captures such that $Q = 98\%$. Figure 6.8 shows the average reconstruction error as a function of SNR levels. One can observe the superiority of bilinear construction in capturing surface appearance even at low SNR levels compared to the linear one. This highlights the benefit of the proposed bilinear appearance model in terms of robustness against noise. Figures 6.9, 6.10 and 6.11 show sample reconstructions for different toys from “Weizmann Photometric Stereo Database” [37] using different BRDFs and illumination maps [55] where the proposed bilinear representation attains minimal reconstruction errors regardless of surface geometry and the deployed reflectance basis.

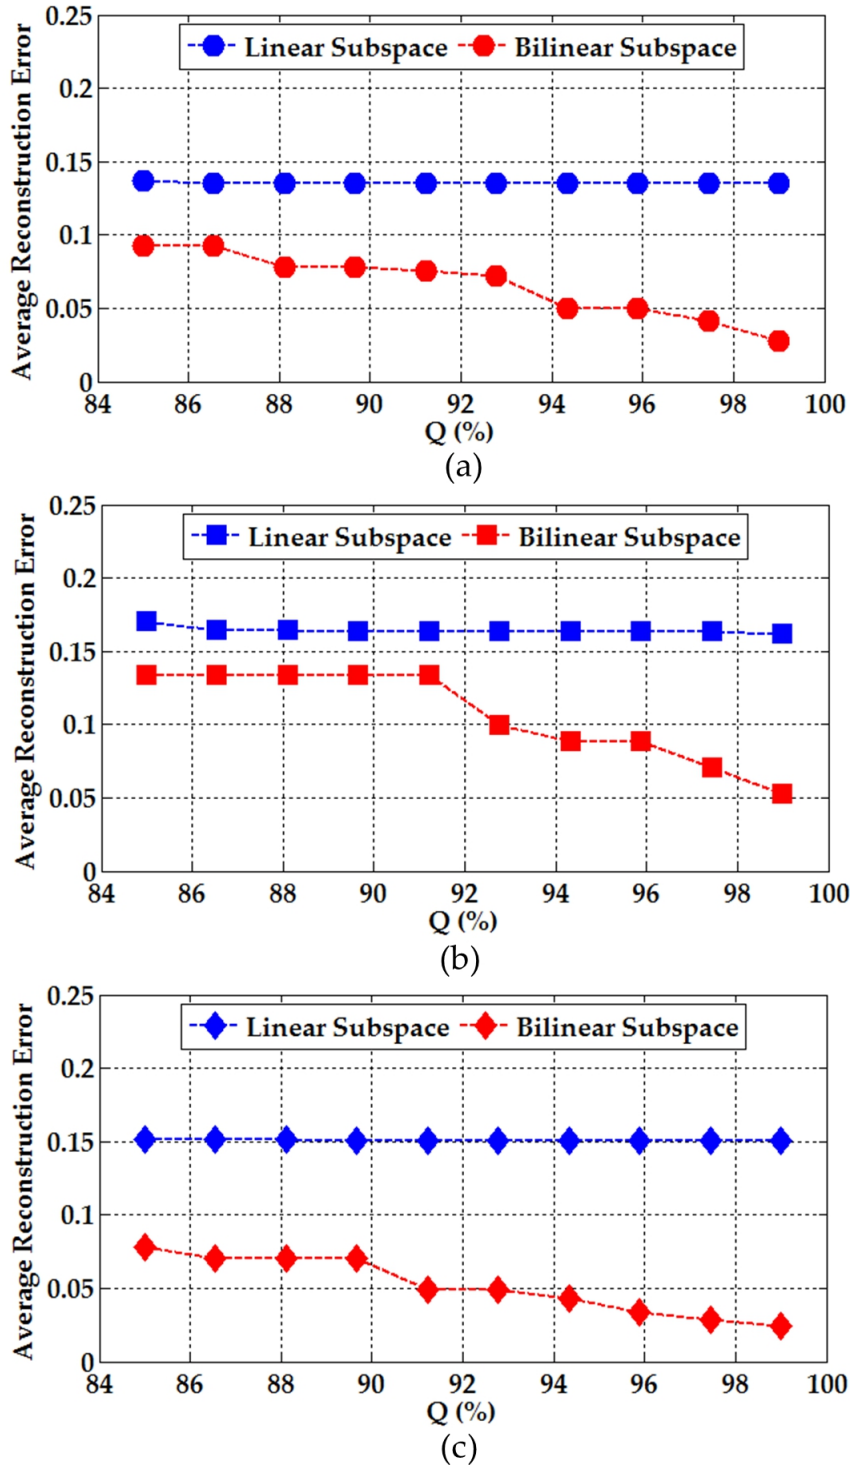


Figure 6.4: The average reconstruction error as a function of the harmonic variation percentage ($Q\%$) maintained by the subspace. Averages were computed based on the 100 rendered testing images for 50 randomly drawn materials from the Merl BRDF database. Surface reflectance is represented using (a) *spherical* harmonics basis [7], (b) *hemispherical* Zernike-based basis [44] and (c) the proposed isotropic *hemispherical* harmonics (HSH)-based Helmholtz reflectance basis.

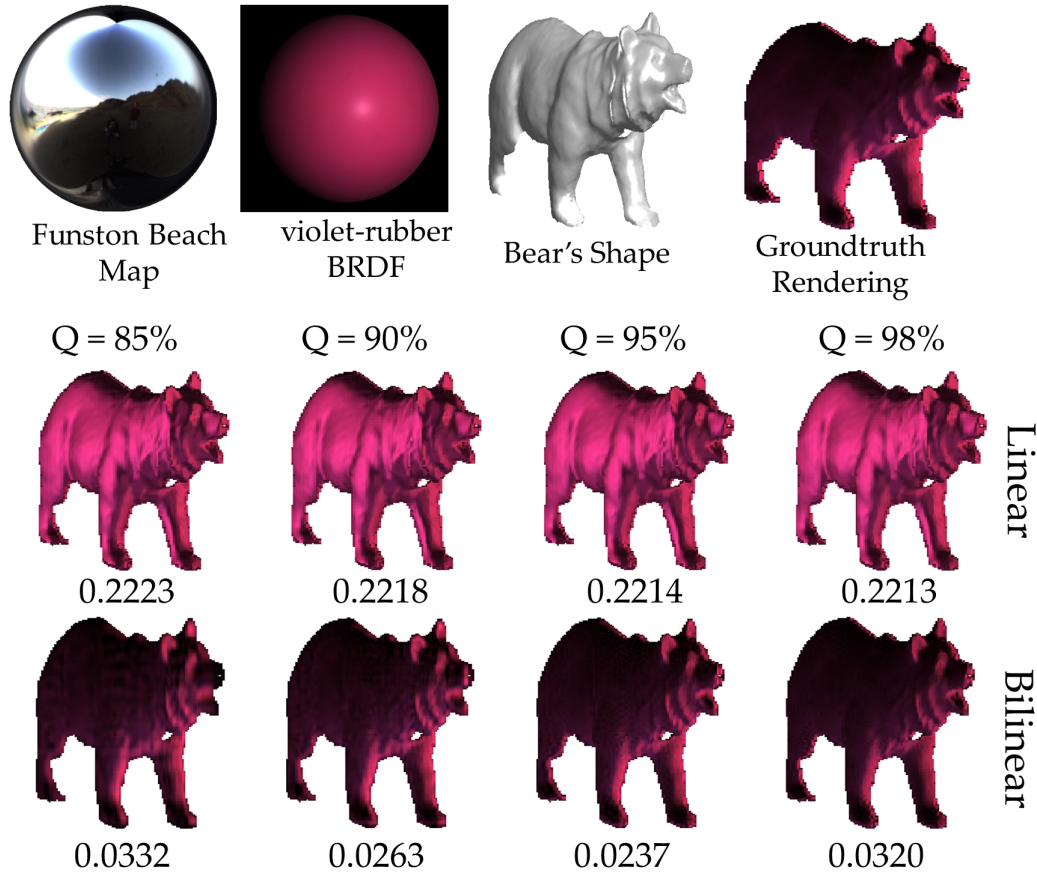


Figure 6.5: Sample reconstructions of the bear toy from “Weizmann Photometric Stereo Database” [37] using the violet-rubber BRDF [17] under the Funston Beach illumination map [55] with different $Q(\%)$ using irradiance harmonics that are based on spherical basis deployed by Ramamoorthi and Hanrahan in [11, 31]. Mean absolute error (MAE) is shown below each rendered image where all intensities are normalized in the range $[0, 1]$.

6.6 Summary

This chapter proposed an analytic formulation for subspace reconstruction to capture the full behavior of complex illumination and non-Lambertian reflectance. Thanks to the frequency-space representation of the image irradiance equation, one was able to incorporate prior information about natural illumination and real world surface materials. The process of finding the analytic subspace was cast as establishing a relation between its principal components and that of the irradiance harmonics basis functions to resolve the issue of dimensionality. By representing images as matrices rather than vectors, one was able to lessen the number of parameters to be estimated to define a bilinear projection which maps the image sample to a lower-dimensional bilinear subspace. Despite admitting to an iterative scheme, the proposed approach showed robustness

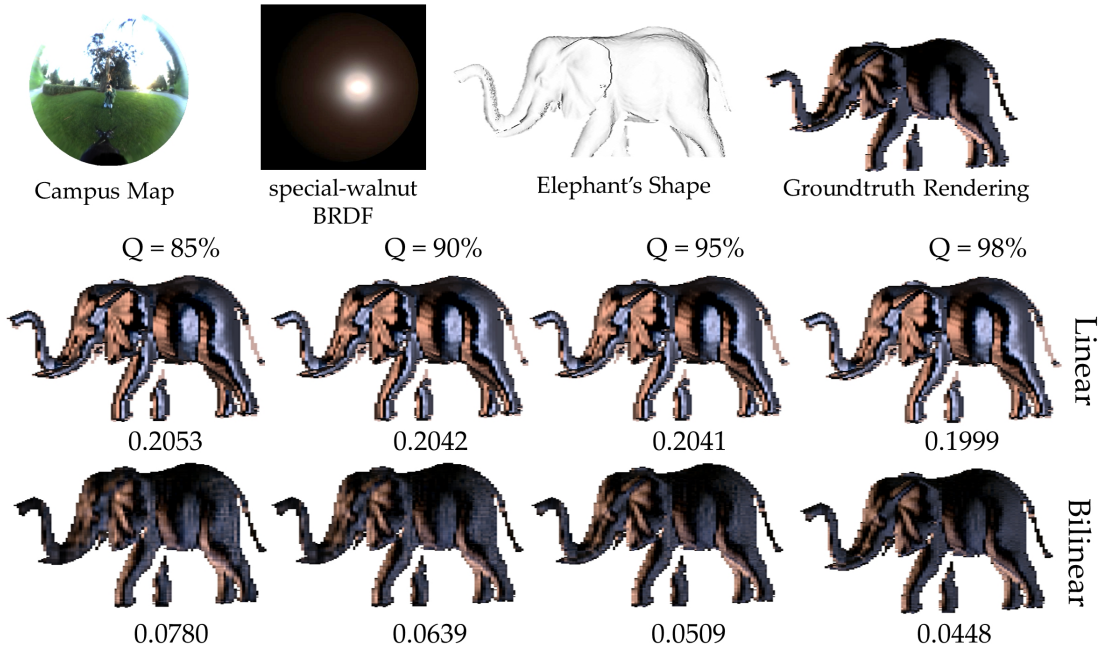


Figure 6.6: Sample reconstructions of the elephant toy from “Weizmann Photometric Stereo Database” [37] using the special-walnut BRDF [17] under the Campus illumination map [55] with different $Q(\%)$ using irradiance harmonics that are based on Zernike-based hemispherical basis deployed by Nilius ad Eklundh in [32, 44]. Mean absolute error (MAE) is shown below each rendered image where all intensities are normalized in the range $[0, 1]$.

w.r.t. initialization while being able to converge in just one iteration when using truncated full projection. The proposed analytic bilinear PCA showed significant decrease in dimensionality when compared to the linear counterpart while attaining the lowest reconstruction errors. It was further tested against noisy input showing robust image representation even at low SNR levels.

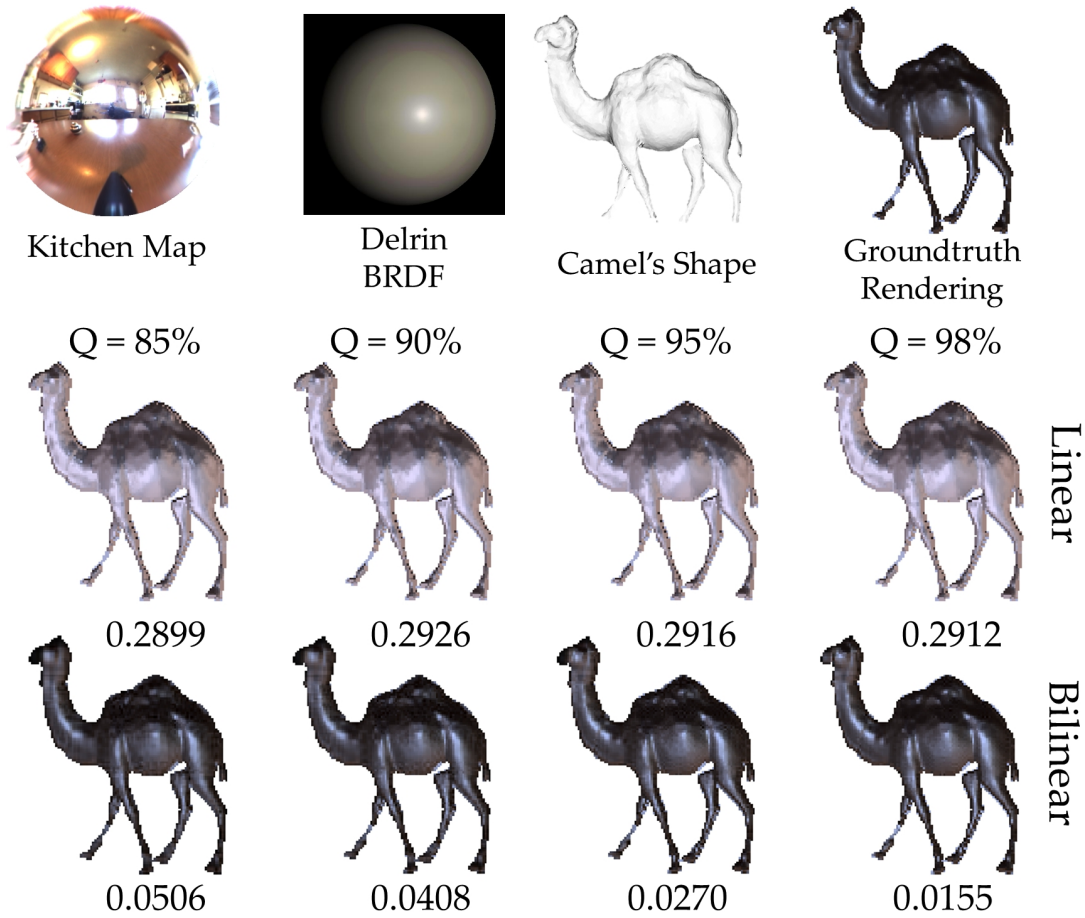


Figure 6.7: Sample reconstructions of the camel toy from “Weizmann Photometric Stereo Database” [37] using the derlin BRDF [17] under the Kitchen illumination map [55] with different Q(%) using irradiance harmonics that are based on the proposed Helmholtz HSH-based basis. Mean absolute error (MAE) is shown below each rendered image where all intensities are normalized in the range [0, 1].

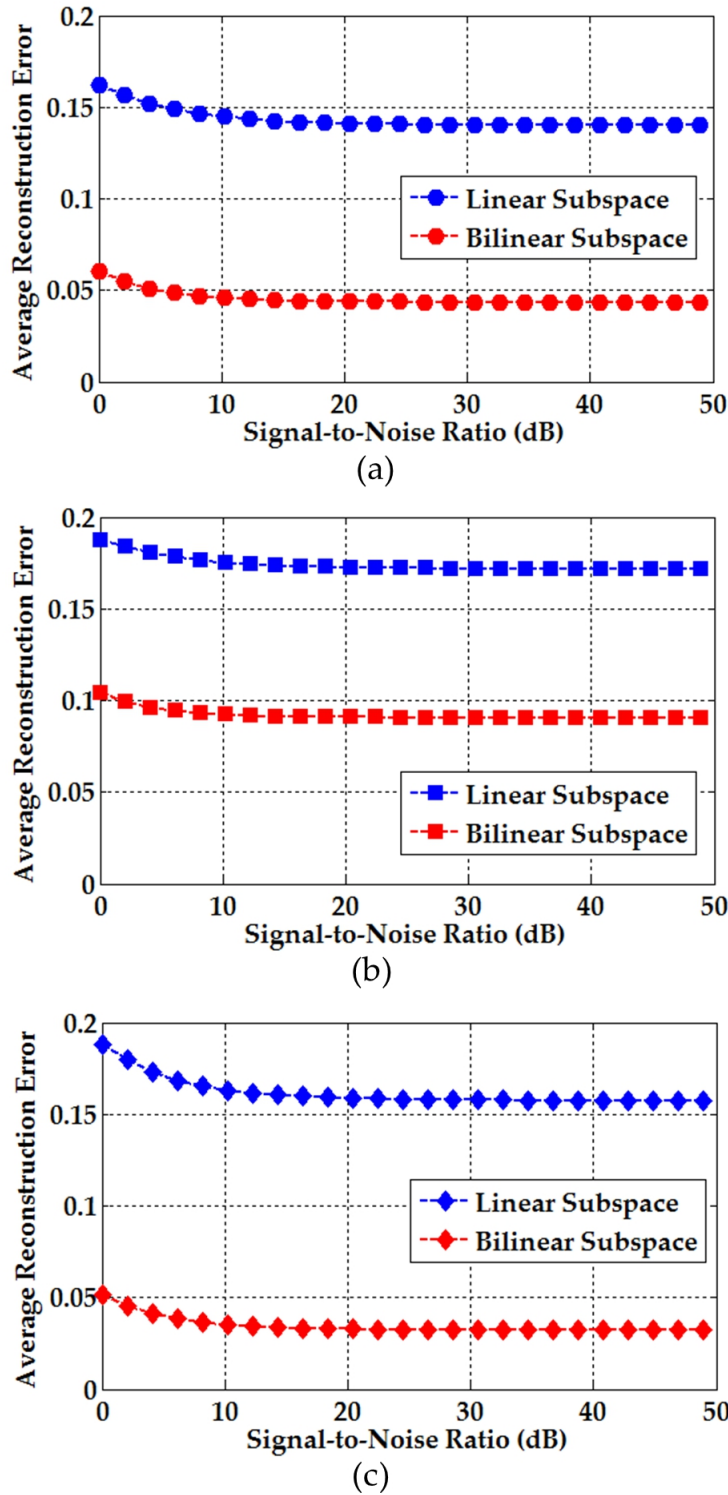


Figure 6.8: The average reconstruction error as a function of the signal-to-noise (SNR) ratio in dB. Averages were computed based on the 100 rendered testing images for 50 randomly drawn materials from the Merl BRDF database. Surface reflectance is represented using (a) *spherical* harmonics basis [7], (b) *hemispherical* Zernike-based basis [44] and (c) the proposed isotropic *hemispherical* harmonics (HSH)-based Helmholtz reflectance basis. White Gaussian noise was added with different levels of SNR. Linear and bilinear subspaces were trained using $Q = 98\%$ of harmonic variations. One can observe that bilinear subspace still capture appearance even with low SNR levels.




Figure 6.9: Sample reconstructions of the camel toy from “Weizmann Photometric Stereo Database” [37] using the yellow-phenolic BRDF [17] under the Campus illumination map [55] with different SNR levels using irradiance harmonics that are based on spherical basis deployed by Ramamoorthi and Hanrahan in [11, 31]. Mean absolute error (MAE) is shown below each rendered image where all intensities are normalized in the range $[0, 1]$.




Figure 6.10: Sample reconstructions of the hippo toy from “Weizmann Photometric Stereo Database” [37] using the pink-fabric BRDF [17] under the Eucalyptus Grove illumination map [55] with different SNR levels using irradiance harmonics that are based on Zernike-based hemispherical basis deployed by Nillius ad Eklundh in [32, 44]. Mean absolute error (MAE) is shown below each rendered image where all intensities are normalized in the range [0, 1].




Figure 6.11: Sample reconstructions of the bear toy from “Weizmann Photometric Stereo Database” [37] using the green-latex BRDF [17] under the Funston Beach illumination map [55] with different SNR levels using irradiance harmonics that are based on the proposed Helmholtz HSH-based basis. Mean absolute error (MAE) is shown below each rendered image where all intensities are normalized in the range [0, 1].

HARMONICS PROJECTION FROM INCOMPLETE IRRADIANCE

This chapter considers the problem of reconstructing the image irradiance signal from incomplete irradiance signal. Sources of incompleteness might include, for example, occlusions due to wearing apparel and makeup for facial images, or even incomppliance to the imaging model assumptions such as non-convexity which introduce cast shadows. This introduces errors in the reconstructed *Harmonics Projection* (HP) images which in turn have a direct impact on subsequent analysis steps. In particular, cast shadow pixels and those which belong to occluded areas will be incorporated into the projection process unless detected and excluded.

We propose to cast errors introduced due to irradiance incompleteness as: (1) statistical outliers which are determined and rejected using robust statistics and (2) local spatial erroneous continuous regions where Markov Gibbs random field with the homogeneous isotropic Potts model is adopted to model the incompleteness's spatial interaction. The results show the effectiveness of the proposed algorithms in handling high levels of contiguous missing information compared to the conventional image projection.

7.1 Introduction

After the conclusion drawn by Basri and Jacobs [6], harmonic expansion of the image irradiance equation has been an attractive fit for illumination modeling. Nonetheless, the main challenge is the computation of what is defined the *Harmonics Projection* (HP) images to be robust against imaging conditions other than illumination. An HP image denotes the image resulted from; (1) projecting a given image onto an analytically

constructed lower-dimensional subspace, then (2) reconstructing the image from the computed projection and the subspace basis functions. This process would benefit many applications of practical interests such as *illumination transfer* for augmented reality applications (*a.k.a.* inverse lighting [111]), *image relighting* for shape recovery and recognition applications [34, 112] and illumination recovery from images [113].

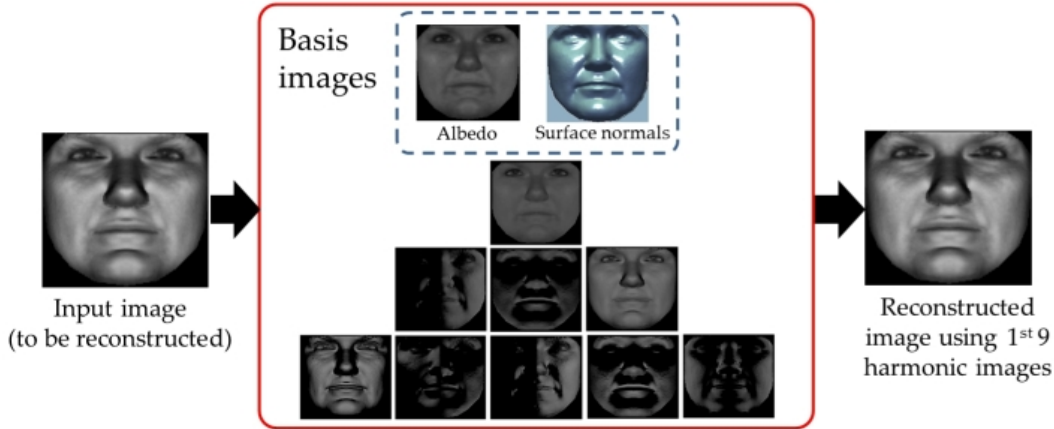


Figure 7.1: Illustration for harmonic-based image reconstruction, the input image is reconstructed by projecting it first to the 1st nine spherical harmonics basis images and then taking the reverse process by summing the scaled basis images. The reconstructed image is visually similar to the input image.

This chapter considers the problem of reconstructing the image irradiance signal from incomplete irradiance signal. The image formation process itself (e.g. perspective projection, cast shadows due to surface non-convexity) generates intensity discontinuity in the observed image [114]. Other sources of incompleteness include wearing apparel such as sunglasses, hats and eyeglasses in case of facial images. Even in the absence of an occluding object, the violation of the imaging model assumptions can also be considered as a source of incompleteness. This introduces errors in the reconstructed HP images which in turn have a direct impact on subsequent analysis steps. In particular, cast shadow pixels and those which belong to occluded areas will be incorporated into the projection process unless detected and excluded. Hence, robustness against missing information is a crucial issue for HP images reconstruction.

Yuen et al. [114] proposed a heuristic rule within the framework of fast marching to handle occlusion due to perspective projection for shape recovery application. With face recognition as an application, Zhou et al. [115] used Markov random fields in order to detect the spatial support of the occlusion which was discarded in subsequent processing. They adopted the classical Ising model for the probability mass function of occlusion

spatial support whose parameters were manually set.

This chapter poses the problem of missing irradiance information as a sparse portion of the image is being corrupted. There is no assumption made for the prior knowledge of the color, shape, size or number of missing pixels. One is mainly concerned with the computation of the HP images to handle reconstruction from partial irradiance information. Missing information is viewed as either (1) statistical outliers to an assumed model where the field of robust statistics [116] can be employed, or (2) local spatial erroneous continuous region, assuming that the corrupted pixels are likely to be adjacent to each other in the observed image. Their spatial interaction is modeled using Markov random fields. Inspired by Zhou et al. [115] work, the spatial support of the corrupted region is identified using Markov Gibbs Random Field (MGRF) with the homogenous isotropic Potts model using an asymmetric Gibbs potential function. Asymmetric Potts model is adopted to guarantee that the Gibbs energy function is submodular, hence it can be minimized using a standard graph cuts approach [117] in polynomial time. In contrast to Zhou [115], the MGRF parameter is automatically estimated which controls the strength of spatial interaction between adjacent pixels; in addition the joint probability density function of the occlusion is modeled as a mixture of Gaussians whose parameters are estimated using the classical Expectation-Maximization (EM) algorithm. Various levels of contiguous missing information are simulated from 1% to 96% by replacing a randomly located block of a facial image with the image of (1) a baboon used in [115] and (2) random noise.

We consider the statistical shape-from-shading framework proposed by Rara *et al.* [34] as one of the potential applications to the proposed approach. This framework makes use of the fact that general lighting can be expressed using low-order spherical harmonics (SH) for convex Lambertian objects. As such incomppliance to assumptions such as diffuse (Lambertian) reflection and non-convexity would introduce discontinuities in the given image, rising the issue of incomplete irradiance information. Figure 7.1 illustrates the process of SH image reconstruction based on the first nine harmonic basis images. Using the USF HumanID 3D face database [118] (for simulations) and the extended Yale face database B, the results illustrate the sensitivity of Rara et al. [34] approach with respect to low levels of missing information. On the contrary, incorporating the proposed approach in their framework can handle high levels of missing information with small error percentage in the recovered shape.

7.2 Harmonic Projection Image Definition

According to the central hypothesis of this dissertation which comply with previous works such as [6, 31, 44, 119], the irradiance signal can be represented by a low-dimensional subspace driven from the harmonic expansion of the image irradiance equation. Thus the image irradiance E produced by a convex surface of arbitrary reflectance with spatially varying albedo can be expressed as a linear combination of pre-computed basis functions $\{\mathcal{B}_s\}$ as,

$$E(\mathbf{x}) = \rho(\mathbf{x}) \sum_{s=1}^S c_s \mathcal{B}_s(\mathbf{x}) \quad (7.1)$$

where $\rho(\mathbf{x})$ denotes albedo at surface point \mathbf{x} .

In matrix notation, Let $\mathbf{E} \in \mathbb{R}^{H \times W}$ be a matrix representation of the image irradiance of the visible surface normals to the viewer such that H denotes height and W denotes width. Let the image \mathbf{E} be considered as a vector $\mathbf{e} \in \mathbb{R}^D$ with $D = H \times W$. Let the s -th albedo-modulated subspace basis function $\mathbf{B}_s = [\rho(\mathbf{x})\mathcal{B}_s(\mathbf{x})] \forall \mathbf{x}$ be represented in the vector space as $\mathbf{b}_s \in \mathbb{R}^D$. Consider $\mathbf{B} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_S] \in \mathbb{R}^{D \times S}$ be the matrix of subspace basis functions as its columns, denoted as basis matrix hereafter, where S is the number of basis functions, and $\mathbf{c} \in \mathbb{R}^{S \times 1}$ vector of image coefficients when projected onto the subspace spanned by $\{\mathcal{B}_s\}$.

This results in an over-determined linear system of equations $\mathbf{e} = \mathbf{B}\mathbf{c}$ which can be solved using Singular Value Decomposition (SVD) (see Algorithm 3). If the input image and the basis matrix used to compute the coefficients \mathbf{c} belong to the same object, one can reconstruct the input image from these coefficients, *i.e.* $\mathbf{h} = \mathbf{B}\mathbf{c} = \mathbf{e}$. However in applications such as illumination transfer and shape recovery, the basis matrix \mathbf{B} would belong to an object which is different from the one in the input image \mathbf{e} , thus the reconstructed image \mathbf{h} provide a mean of encoding the illumination and/or the reflection¹ of the input image while maintaining the geometrical structure of the object used to derive the basis matrix, See Figure 7.2 for illustration.

As such, Harmonic Projection (HP) image can be defined as the image resulted from; (1) computing the subspace coefficients \mathbf{c} given the input image and the basis matrix of some object, then (2) reconstructing the image from the computed coefficients and the basis matrix. Thus, for example in the framework of statistical shape-from-shaping [34],

¹In case of assuming Lambertian reflectance, the coefficients \mathbf{c} encode only the illumination conditions of the input image.

the process of finding the HP image for each human subject in a given database allows subjects in this database to share the illumination conditions of the input image. The SVD approach is outlined in Algorithm 3.

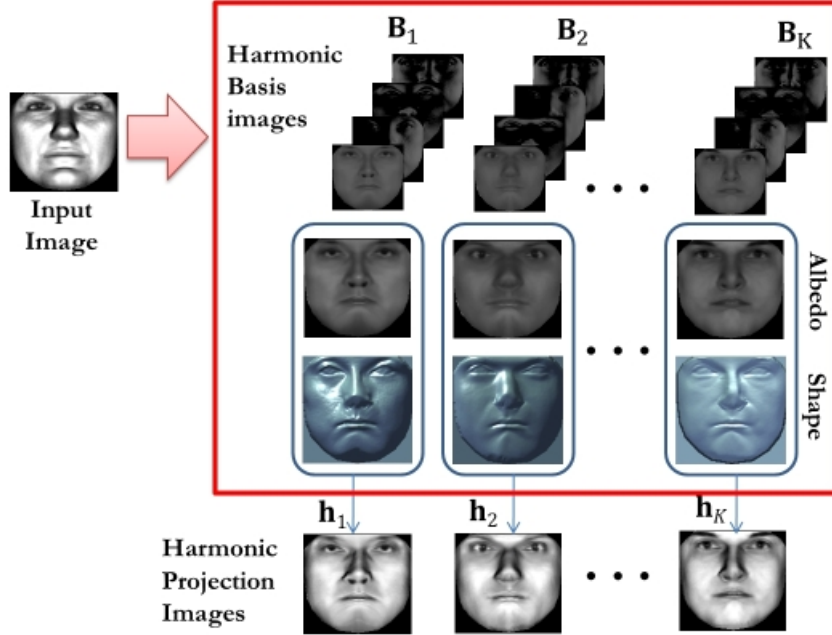


Figure 7.2: Harmonics Projection (HP) Images: when the basis images of a different object are used to reconstruct the input image, the HP images encode the illumination and/or reflectance conditions of the input but retain the identity of the object involved. The pipeline in Algorithm 3 is extended to distinct subjects of the USF database [118]. Note that Active Appearance Modeling (AAM) [120] is used to automatically localize facial landmarks which guide the process of aligning the input image to the aligned albedo samples of the USF database [118] using thin-plate-splines warping.

Algorithm 3 Compute Harmonic Projection (HP) Image

- 1: **Input:** (a) Input image, $\mathbf{e} \in \mathbb{R}^D$ with $D = H \times W$, (b) Basis matrix, \mathbf{B}
 - 2: **Output:** HP image, \mathbf{h}
 - 3: Compute the SVD decomposition of the basis matrix \mathbf{B} : $[\mathbf{U}, \mathbf{S}, \mathbf{V}] = \text{svd}(\mathbf{B})$
 - 4: Retain the first S -columns of \mathbf{U} : $\mathbf{U} = \mathbf{U}(:, 1 : S)$, where S is the number of basis functions in the underlying subspace
 - 5: Retain the first S -singular values in \mathbf{S} : $\mathbf{S} = \mathbf{S}(1 : S, 1 : S)$
 - 6: Solve for the coefficient vector $\hat{\mathbf{c}}$ using the least-squares estimate of $\mathbf{e} = \mathbf{B}\mathbf{c}$ where
$$\hat{\mathbf{c}} = \mathbf{V}\mathbf{S}^{-1}\mathbf{U}^T\mathbf{e}$$
 - 7: Compute the HP image; $\mathbf{h} = \mathbf{B}\hat{\mathbf{c}}$
-

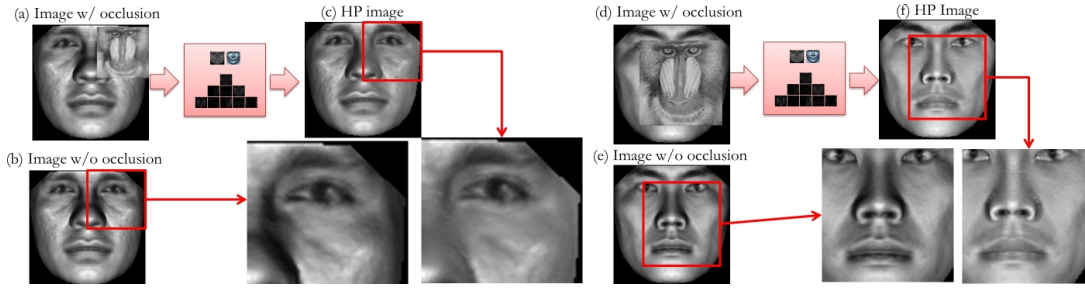


Figure 7.3: Two samples of input images (b) and (e) illuminated using combined light source at $(0,0,1)$ and $(0,0.5,0.9)$ are corrupted by 25% (a) and 50% (d) contiguous occlusion, Algorithm 3 was used to reconstruct the input image from the basis images of the same subject (see Figure 7.1) (c) and (f), it can be observed that the input image was not correctly recovered due to the presence of occluding region which was taken into consideration when constructing the linear system of equation $\mathbf{e} = \mathbf{Bc}$.

7.3 HP Images With Missing Information

Consider the case where the input image is partially occluded or is due to a non-Lambertian object, HP algorithm based on SVD (Algorithm 3) would fail especially having the basis matrix being driven under the diffuse reflectance assumption. This is because such an algorithm is based on the construction of a linear system of equations which depends on the holistic view of the input image, see Figure 7.3. However, if the spatial support of the occlusion/missing information can be reliably determined, the corrupted pixels can be discarded from this system of equations and estimating the coefficients can proceed using the non-corrupted pixels.

7.3.1 Irradiance Corruption as Statistical Outlier

In [47], the field of robust statistics [116] is deployed to solve the problem of estimating the projection coefficients when outliers exist in the input image. The system of equations $\mathbf{e} = \mathbf{Bc}$ can be translated into the robust estimation framework with the following minimization;

$$\min_{\mathbf{c}} E_R(\mathbf{c}) = \min_{\mathbf{c}} \sum_{\mathbf{x}} \rho(\mathbf{e} - \mathbf{Bc}; \sigma) \quad (7.2)$$

where σ is a scale parameter and ρ is the robust estimate of the error norm.

Different ρ -functions result into various robust estimators in the sense of being insensitive to outliers. Two examples are used in this work which are the Geman-McClure¹

¹ $\rho(x) = \frac{x^2}{\sigma + x^2}$




Figure 7.4: Geman-McClure ρ and ψ functions. Left: $\rho(x) = \frac{x^2}{\sigma+x^2}$. Right: $\psi(x) = \rho'(x) = \frac{2x\sigma}{(\sigma+x^2)^2}$




Figure 7.5: Lorentzian ρ and ψ functions. Left: $\rho(x) = \log\left(1 + \frac{1}{2}\left(\frac{x}{\sigma}\right)^2\right)$. Right: $\psi(x) = \rho'(x) = \frac{2x}{2\sigma^2+x^2}$

(Figure 7.4) and Lorentzian¹ ρ -functions (Figure 7.5) because both are twice differentiable and gradient descent-based approaches such as *simultaneous over-relaxation* can be used to find the local minima. These ρ -functions also have the scale parameter ρ which makes it attractive to use *continuation methods*, which can be used to find a globally optimal solution of non-convex objective function. Blake and Zisserman [121] developed the Graduated Non-Convexity algorithm, a type of continuation method which constructs a parameterized piecewise polynomial approximation to the truncated quadratic. Instead of using truncated quadratic, this work proposes to use both Geman-McClure and Lorentzian ρ -functions and create convex approximations out of them.

Using simultaneous over-relaxation, the iterative update equations for this minimization with respect to the s -th coefficient c_s can be written as;

$$c_s^{t+1} = c_s^t - \omega \frac{1}{T(c_s)} \frac{\partial E_R}{\partial c_s} \quad (7.3)$$

¹ $\rho(x) = \log\left(1 + \frac{1}{2}\left(\frac{x}{\sigma}\right)^2\right)$

where $0 < \omega < 2$ is an over-relation parameter that is used to overcorrect the estimate c_s^{t+1} . $T(c_s)$ is an upper bound on the second partial derivative of E_R . The complete algorithm to compute the robust HP images is shown in Algorithm 4.

Algorithm 4 Compute Robust Harmonic Projection (HP) Image

- 1: **Input:** (a) Input image, \mathbf{e} , (b) Basis matrix, \mathbf{B}
 - 2: **Output:** HP image, \mathbf{h}
 - 3: **Initial coefficients:** Solve for \mathbf{c}^0 using the least-squares estimate of $\mathbf{e} = \mathbf{B}\mathbf{c}^0$ (use Algorithm 3)
 - 4: **Compute partial derivatives:** Compute $\partial E_R / \partial c_s$ for $s \in [1, S]$.
 - 5: **Coefficient vector update:** Update using $c_s^{t+1} = c_s^t - \omega \frac{1}{T(c_s)} \frac{\partial E_R}{\partial c_s}$ where ω is an over-relation parameter. $T(c_s) \geq \partial^2 E_R / \partial c_s^2$.
 - 6: **Robust function parameter update (Graduated Non-convexity):** $\sigma^{t+1} = \kappa \sigma^t$ with $\kappa \in (0, 1)$.
 - 7: **Repeat steps 2-4 until maximum number of iterations or convergence.**
 - 8: **Test for outliers:** Determine outlying measurements using $|\mathbf{e} - \mathbf{B}\mathbf{c}| \geq \tau$ where τ is determined by the error norm and the control parameter σ .
 - 9: **Compute coefficients from non-outlier pixels:** solve for $\hat{\mathbf{c}}$ using the least-squares estimate of $\mathbf{e} = \mathbf{B}\hat{\mathbf{c}}$ (use Algorithm 3) after discarding outlier pixels from both the input image and the basis images.
 - 10: **Compute the HP image:** $\mathbf{h} = \mathbf{B}\hat{\mathbf{c}}$
-

7.3.2 Irradiance Corruption as Erroneous Continuous Region

In the presence of discontinuous intensity regions due to irradiance corruption, the HP image can be modeled as $\mathbf{e} = \mathbf{B}\mathbf{c} + \boldsymbol{\epsilon} = \mathbf{h} + \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon} \in \mathbb{R}^D$ is a sparse error vector occurred due to some contiguous irradiance corruption. Its nonzero entries are both sparse and spatially continuous, thus it can be recovered by solving the ℓ^1 -norm minimization problem [115];

$$\hat{\boldsymbol{\epsilon}} = \operatorname{argmin} \|\boldsymbol{\epsilon}\|_1 \quad \text{s. t.} \quad \mathbf{e} = \mathbf{h} + \boldsymbol{\epsilon} \quad (7.4)$$




Figure 7.6: \mathbf{f} defines the spatial support of the error $\boldsymbol{\epsilon}$

Let $\mathbf{f} \in \{-1, 1\}^D$ defines the error spatial support vector where $\mathbf{f}(\mathbf{x}) = -1$ for non-corrupted pixels and $\mathbf{f}(\mathbf{x}) = 1$ for corrupted pixels, see Figure 7.6. Finding this spatial support can be posed as an image labeling problem where the objective is to assign corrupted or non-corrupted label to each pixel in the error vector. Let $\mathcal{P} = 1, 2, \dots, D$ be the set of D -pixels representing the error vector. The structure of the input image helps to define a geometric neighborhood system \mathfrak{N} consisting of a set of all neighboring pairs $\{\mathbf{p}, \mathbf{q}\}$ where $\mathbf{p}, \mathbf{q} \in \mathcal{P}$. The second order neighborhood system is used where the eight nearest neighbors are sharing a side with a given pixel are considered.

Using the maximum-a-posteriori estimation framework [122], one can describe the error vector $\boldsymbol{\epsilon}$ and its spatial support \mathbf{f} by a joint Markov Gibbs random field (MGRF) model of independent image signals and interdependent region labels. A two-level probability model of the error vector and its spatial support is given by a joint distribution $P(\boldsymbol{\epsilon}, \mathbf{f}) = P(\mathbf{f})P(\boldsymbol{\epsilon}|\mathbf{f})$ where $P(\mathbf{f})$ defines the unconditional probability distribution of the spatial support, while $P(\boldsymbol{\epsilon}|\mathbf{f})$ denotes the conditional distribution of the error vector given its spatial support. The Bayesian maximum-a-posteriori estimate of the spatial support \mathbf{f} can thus be expressed as;

$$\mathbf{f}^* = \underset{\mathbf{f} \in \mathcal{F}}{\operatorname{argmin}} P(\boldsymbol{\epsilon}|\mathbf{f})P(\mathbf{f}) \quad (7.5)$$

where \mathcal{F} is the set of all possible labeling. In order to assure that the posterior distribution $P(\boldsymbol{\epsilon}|\mathbf{f})$ is a Markov random field, one assumes independent noise at each pixel, therefore $P(\boldsymbol{\epsilon}|\mathbf{f}) = \prod_{\mathbf{p} \in \mathcal{P}} P(\epsilon_{\mathbf{p}}|f_{\mathbf{p}})$. Gibbs random field is used to provide a global model for the error vector by specifying the probability mass function of its spatial support as $P(\mathbf{f}) = \frac{1}{Z} \exp\left(-\sum_{\{\mathbf{p}, \mathbf{q}\} \in \mathfrak{N}} V_{\mathbf{p}\mathbf{q}}(f_{\mathbf{p}}, f_{\mathbf{q}})\right)$, where $V_{\mathbf{p}\mathbf{q}}(\cdot, \cdot)$ is a two-pixel potential controlling the spatial interaction between neighboring pixels. Following Ali et al. [123], one adopts the asymmetric pairwise homogenous isotropic MGRF model with Potts prior where the asymmetric Potts model is chosen to guarantee that the Gibbs energy function is submodular, hence it can be minimized using a standard graph cuts approach [117] in polynomial time. The Gibbs potential for asymmetric pairwise co-occurrences of labels can be described as $V(f_{\mathbf{p}}, f_{\mathbf{q}}) = \zeta \delta(f_{\mathbf{p}} \neq f_{\mathbf{q}})$ where $\delta(\cdot)$ is the indicator function and ζ is a parameter which influences the interaction between neighboring pairs. Note that the $\mathbf{p}\mathbf{q}$ -subscript in $V_{\mathbf{p}\mathbf{q}}$ is dropped since one is dealing with homogenous isotropic

models. Hence the Gibbs probability distribution can be written as [123];

$$\begin{aligned} P(\mathbf{f}) &= \frac{1}{Z} \exp \left(- \sum_{\{\mathbf{p}, \mathbf{q}\} \in \mathfrak{N}} V_{\mathbf{p}\mathbf{q}}(f_{\mathbf{p}}, f_{\mathbf{q}}) \right) \\ &= \frac{1}{Z} \exp(-\zeta |\mathcal{T}| \mathfrak{F}_{\text{neq}}(\mathbf{f})) \end{aligned} \quad (7.6)$$

where $\mathcal{T} = \{\{\mathbf{p}, \mathbf{q}\} : \mathbf{p}, \mathbf{q} \in \mathcal{P}, \{\mathbf{p}, \mathbf{q}\} \in \mathfrak{N}\}$ is the family of the neighboring pixel pairs supporting the Gibbs potentials, with $|\mathcal{T}|$ as its cardinality, $\mathfrak{F}_{\text{neq}}(\mathbf{f})$ denotes the relative frequency of the non-equal labels in pixel pair of \mathcal{T} . It is defined as; $\mathfrak{F}_{\text{neq}}(\mathbf{f}) = \frac{1}{|\mathcal{T}|} \sum_{\{\mathbf{p}, \mathbf{q}\} \in \mathcal{T}} \delta(f_{\mathbf{p}} \neq f_{\mathbf{q}})$. In contrast to the work done by Zhou et al [115], the model parameter ζ is estimated based on the analytical solution proposed by Ali et al. [123], where the approximate maximum likelihood estimate (MLE) of ζ is computed as $\zeta^* = K(1 - 2\mathfrak{F}_{\text{neq}}(\mathbf{f}))$, where $K = 2$ denotes the number of labels.

Using the log-likelihood of Equation 7.5, the labeling problem can be formulated as the following maximization,

$$\mathbf{f}^* = \underset{\mathbf{f} \in \mathcal{F}}{\text{argmin}} \left(\sum_{\mathbf{p} \in \mathcal{P}} \log(P(\epsilon_{\mathbf{p}}|f_{\mathbf{p}})) - \sum_{\{\mathbf{p}, \mathbf{q}\} \in \mathfrak{N}} V(f_{\mathbf{p}}, f_{\mathbf{q}}) \right) \quad (7.7)$$

This is equivalent to minimizing the following energy function;

$$E_M(\mathbf{f}) = \sum_{\{\mathbf{p}, \mathbf{q}\} \in \mathfrak{N}} V(f_{\mathbf{p}}, f_{\mathbf{q}}) + \sum_{\mathbf{p} \in \mathcal{P}} \overbrace{-\log(P(\epsilon_{\mathbf{p}}|f_{\mathbf{p}}))}^{\text{Data Penalty Term}} \quad (7.8)$$

Zhou et al. [115] used a piecewise-constant likelihood function to approximate the conditional probability of the error $P(\epsilon|\mathbf{f})$. While the precise form of the approximation was not essential for their application, through the conducted experimentation, one found that such rough approximation affects the data penalty term leading to inaccurate HP images. This work approximates the marginal probability density functions of each class using a mixture of Gaussians whose parameters are estimated using the classic Expectation-Maximization (EM) algorithm [124].

In the processing of minimizing $E_M(\mathbf{f})$, one needs to set a threshold τ which indicates the accepted level of error to consider a nonzero entry in the error vector as resulting from occlusion. The conditional probabilities can be used in order to derive such threshold. However, in the general case one will be comparing an input image \mathbf{e} with an HP image

obtained from a subspace constructed from different geometrical structure, hence an error results due object identity rather than pixel corruption. However these errors are not separable, thus one cannot rely on the error probability density function to give such a threshold using conditional probabilities. Following the strategy of Zhou et al. [115], the error value is normalized to be in the range of $[0, 1]$. Therefore τ is chosen adaptively within that range, starting with a relatively large value and reducing it by a constant step size at each iteration.

We use the project-out approach of [120] by iterating between the estimation of the error spatial support and the projection coefficients. The complete algorithm to compute the MGRF-based HP images is shown in Algorithm 5.

7.4 Experimental results

To compare the proposed algorithm to that of Rara et al [34], various levels of contiguous occlusion are simulated ranging from 1% to 96% by replacing a randomly located block of a given facial image with (1) the image of a baboon used in [115] and (2) random noise (in Yale experiments). For Algorithm 4, the parameters are set as follows; $\omega = 1.995$, $\kappa = 0.95$, $\sigma = 1$ and $\tau = \sigma/\sqrt{3}$ in case of Geman-McClure norm and $\tau = \sigma\sqrt{2}$ in case of Lorentzian norm. For Algorithm 5, one uses $\tau_o = 0.2$ and $\Delta\tau = 0.03$ (values recommended by Zhou in [115]). HP images are reconstructed in 30 iterations in case of robust estimation (Algorithm 4) while MGRF takes 12 iterations to estimate a stable occlusion spatial support. One uses $\epsilon = 0.7$ to compute the overall occluding region given the occlusion spatial support of individual subjects in the database. Ground-truth images \mathbf{E}^G are formed using the Lambertian model with combined light source directions at $(0, 0, 1)$ and $(0, 0.5, 0.9)$. The input image \mathbf{E} is then formed by replacing a randomly located block with an occluder image in order to occlude a certain percentage of the facial region.

We use the following codes; (1) SVD: Algorithm 3, (2) RSL: Algorithm 4 with Lorentzian function, (3) RSG: Algorithm 4 with Geman-McClure function, (4) MGRF0: Algorithm 5 with $\zeta = 2$ (to compare to the case [115]) with piecewise-constant likelihood function, (5) MGRF1: same as MGRF0 but with automatic estimation of the Gibbs potential parameter ζ [123], (6) MGRF2: same as MGRF1 with mixture-of-Gaussian likelihood function (7) MGRF3: same as MGRF2 with error threshold τ adaptively

Algorithm 5 Compute Harmonic Projection (HP) Image with MGRF and Sparsity

- 1: **Input:** (a) Input image, \mathbf{e} , (b) Basis matrix, \mathbf{B}
- 2: **Output:** HP image, \mathbf{h}
- 3: **Initial coefficients:** Solve for \mathbf{c}^0 using the least-squares estimate of $\mathbf{e} = \mathbf{B}\mathbf{c}^0$ (use Algorithm 3)
- 4: **Initialize the error threshold:** $\tau^0 = \tau_o$
- 5: **Compute initial error vector:** $\hat{\boldsymbol{\epsilon}} = |\mathbf{e} - \mathbf{B}\mathbf{c}^0|$
- 6: **Compute initial error spatial support:**

$$\mathbf{f}^0 = \begin{cases} -1 & \hat{\epsilon} < \tau^0 \\ 1 & \hat{\epsilon} \geq \tau^0 \end{cases}$$

- 7: **repeat**
- 8: **Exclude occluding pixels:** $\mathbf{B}^* = \mathbf{B}(\mathbf{f}^{t-1} = -1, :)$, $\mathbf{e}^* = \mathbf{e}(\mathbf{f}^{t-1} = -1)$.
- 9: **Solve the ℓ^1 -norm minimization problem:**
 $(\hat{\mathbf{c}}, \hat{\boldsymbol{\epsilon}}) = \operatorname{argmin} \|\boldsymbol{\epsilon}\|_1 \quad \text{s. t.} \quad \mathbf{e}^* = \mathbf{B}^* \mathbf{c} + \boldsymbol{\epsilon}^*$.
- 10: **Compute the complete error vector:** $\hat{\boldsymbol{\epsilon}} = |\mathbf{e} - \mathbf{B}\hat{\mathbf{c}}|$
- 11: **Find an initial labeling** for the error spatial support

$$\mathbf{f} = \begin{cases} -1 & \hat{\epsilon} < \tau^t \\ 1 & \hat{\epsilon} \geq \tau^t \end{cases}$$

- 12: **Compute the initial MGRF Potts model parameter:**

$$\zeta^* = K(1 - 2\mathfrak{F}_{\text{neq}}(\mathbf{f}))$$

- 13: **Fit a mixture of two Gaussians** using EM-algorithm to obtain conditional probabilities $P(\boldsymbol{\epsilon}|\mathbf{f} = -1)$ and $P(\boldsymbol{\epsilon}|\mathbf{f} = 1)$.
- 14: **Update the error spatial support using graph cuts:**

$$\mathbf{f}^* = \operatorname{argmax}_{\mathbf{f} \in \mathcal{F}} \exp \left(\sum_{\mathbf{p} \in \mathcal{P}} \log(P(\boldsymbol{\epsilon}_{\mathbf{p}}|f_{\mathbf{p}})) - \sum_{\{\mathbf{p}, \mathbf{q}\} \in \mathfrak{N}} V(f_{\mathbf{p}}, f_{\mathbf{q}}) \right)$$

- 15: **Update the error threshold:** $\tau^{t+1} = \tau^t + \Delta\tau$
 - 16: **until** maximum number of iterations or convergence.
 - 17: **Compute coefficients from non-outlier pixels:** solve for $\hat{\mathbf{c}}$ using the least-squares estimate of $\mathbf{e} = \mathbf{B}\hat{\mathbf{c}}$ (use Algorithm 3) after discarding outlier pixels from both the input image and the basis matrix.
 - 18: **Compute the HP image:** $\mathbf{h} = \mathbf{B}\hat{\mathbf{c}}$
-

lowered at each iteration with a constant step size.

7.4.1 HP Images Under Occlusion

Assuming that HP images are reconstructed using the basis matrix of the same subject. The mean absolute error between the ground-truth image \mathbf{E}^G and the reconstructed HP image \mathbf{h} . Figure 7.7 shows the mean and standard deviation of the mean absolute error measured between the ground-truth images and the reconstructed ones using occlusion percentages ranging from 1% to 96%. One can conclude the following: (1) SVD is unstable even with low levels of occlusion. (2) MGRF based on automatic estimation of Potts potential parameter and fitting mixture-of-Gaussians to the error probability distribution (MGRF3) provides significant improvement compared to other approaches especially for high levels of occlusion ($> 40\%$) (3) The use of robust statistics (RSL and RSG) acts as a midway between SVD and MGRF3. (4) Lorentzian robust estimator (RSL) is less sensitive to outliers when compared to Geman-McClure (RSG). (5) The use of piecewise-constant approximation of the likelihood functions provides similar results in case of automatically estimation the Potts potential parameter (MGRF1) or manually setting it (MGRF0). Hence the estimation of the potential parameter does not have a significant impact when using such approximation for the likelihood functions. (6) Piecewise-constant likelihood functions seem to be more robust to low levels occlusion ($< 40\%$). However with higher levels of occlusion, the choice of such approximation imposes a significant influence on the reconstruction accuracy.

Figure 7.8 shows sample results of the HP image reconstruction under different occlusion percentages. The following can be observed: (1) As the occlusion percentage increases, the performance of the classic SVD becomes inferior when compared to the other proposed approaches. (2) Modeling the spatial interaction of the occluding region using MGRF provides the best results in terms of the reconstructed HP image and the estimated occlusion spatial support. (3) The estimation of the Potts potential parameter provides better results (MGRF1, MGRF2 and MGRF3) since different spatial supports will have different values for potential parameter. (4) The estimation of the error threshold from the conditional probabilities (MGRF2) leads to inability to accurately estimate the occlusion spatial support. This is evident especially in case of very high occlusion level where the identity of the subject is mistaken with that of the baboon. (5) Piecewise-constant likelihood functions (MGRF0 and MGRF1) provide rough

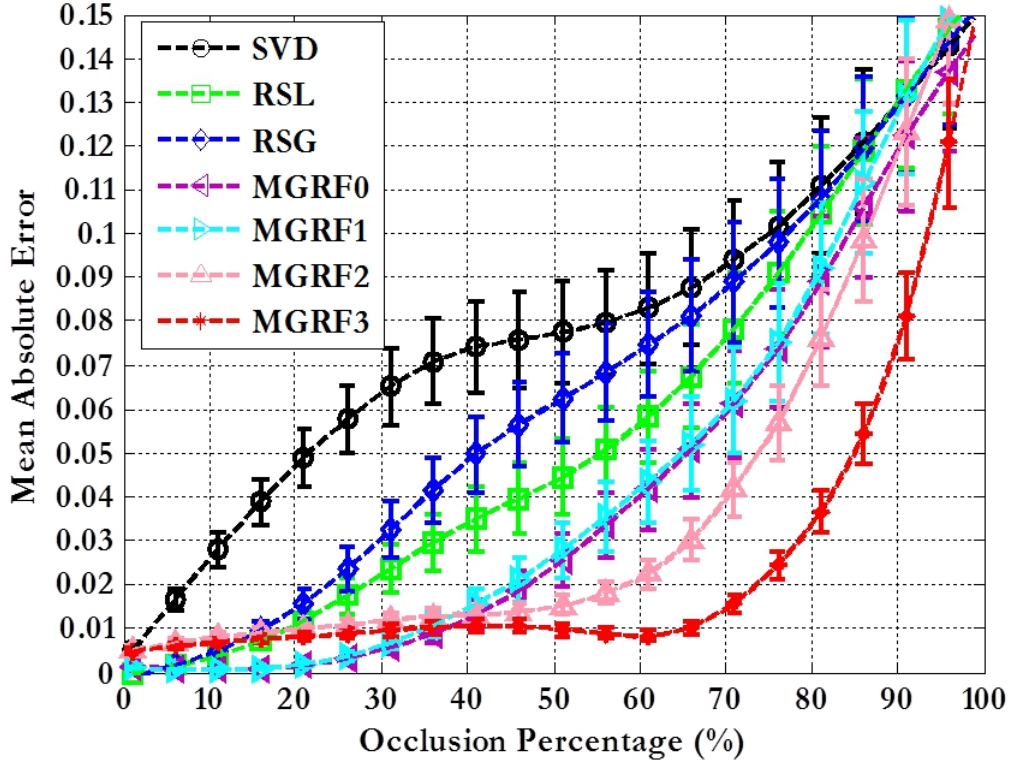


Figure 7.7: The mean and standard deviation of the mean absolute error measured between the ground-truth images and the reconstructed ones using occlusion percentages ranging from 1% to 96%.

estimation for the data penalty term which makes the estimation of the spatial support unreliable when compared to the mixture-of-Gaussian cases (MGRF2 and MGRF3).

7.4.2 Shape Recovery Under Occlusion

We evaluate the performance of the proposed method in recovering the 3D shape given a partially occluded facial image. The face models are built using the USF 3D Face Database [118]. The database contains 100 subjects of diverse gender and ethnicity. The 3D shape for 100 out-of-training is recovered. In each experiment, The following error metrics are computed: (1) Height Error (reported as percentages) where the recovered height map is compared with the ground-truth height map, and (2) Surface Orientation Error where the directions of the recovered surface normal vectors are compared to the ground-truth ones. The results from SVD, RSL, RSG and MGRF3 are reported, where shapes recovered from other variants of MGRF are not included due to their performance in previous subsection compared to MGRF3.

From Figures 7.9 and 7.10, it can be observed that: (1) Algorithms tend to be divided into two main groups, (a) SVD and RSL being sensitive to occlusion and (b)

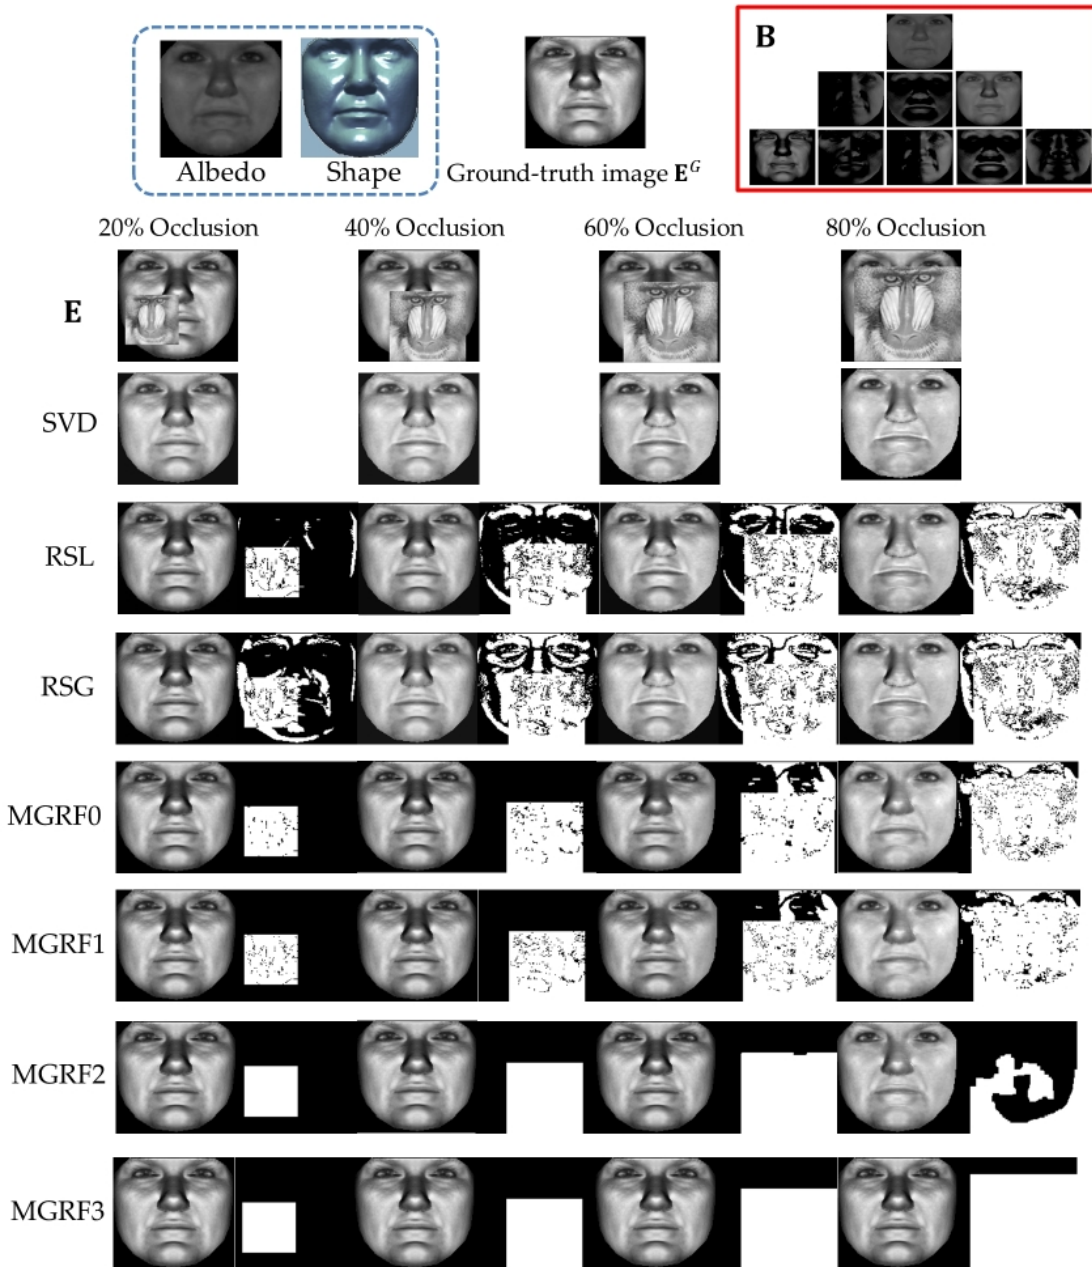


Figure 7.8: Sample results of HP images reconstruction with different levels of occlusions. The estimated occlusion spatial support of each algorithm is shown in the even-ordered columns. MGRF3 provides the best results in terms of the reconstructed HP image and the estimated occlusion spatial support. Note that such spatial support does not exist for the SVD algorithm.

RSG and MGRF3 being robust to contiguous occlusion even in the presence of high levels occlusion. (2) RSG provides better performance when compared to RSL. This shows the robustness of Geman-McClure function over Lorentzian function in handling the errors due to different identities (where basis images used to reconstruct the input image does not have the same identity as the input image).

Figure 7.11 shows sample results in terms of recovered shape and albedo of one out-

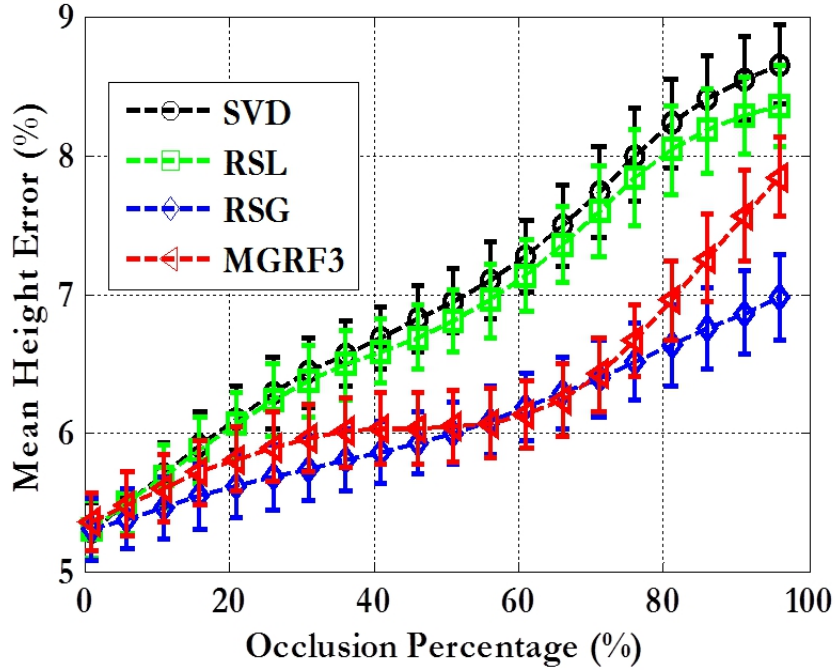


Figure 7.9: Mean and standard deviation of the mean absolute height error (reported as percentages) where the recovered height map (shape) is compared with the ground-truth height map.




Figure 7.10: Mean and standard deviation of the surface orientation error (reported in radians) where the directions of the recovered surface normal vectors are compared to the ground-truth ones.

of-training subject using different levels of occlusions. The profile comparison of the recovered versus the ground-truth shape is also shown. One can observe the following; (1) SVD behaves poorly especially with increasing occlusion level. (2) Geman-McClure reconstructions are closer to the ground-truth when compared to that of Lorentzian. (3)

The power of MGRF is manifested with high levels of occlusion.

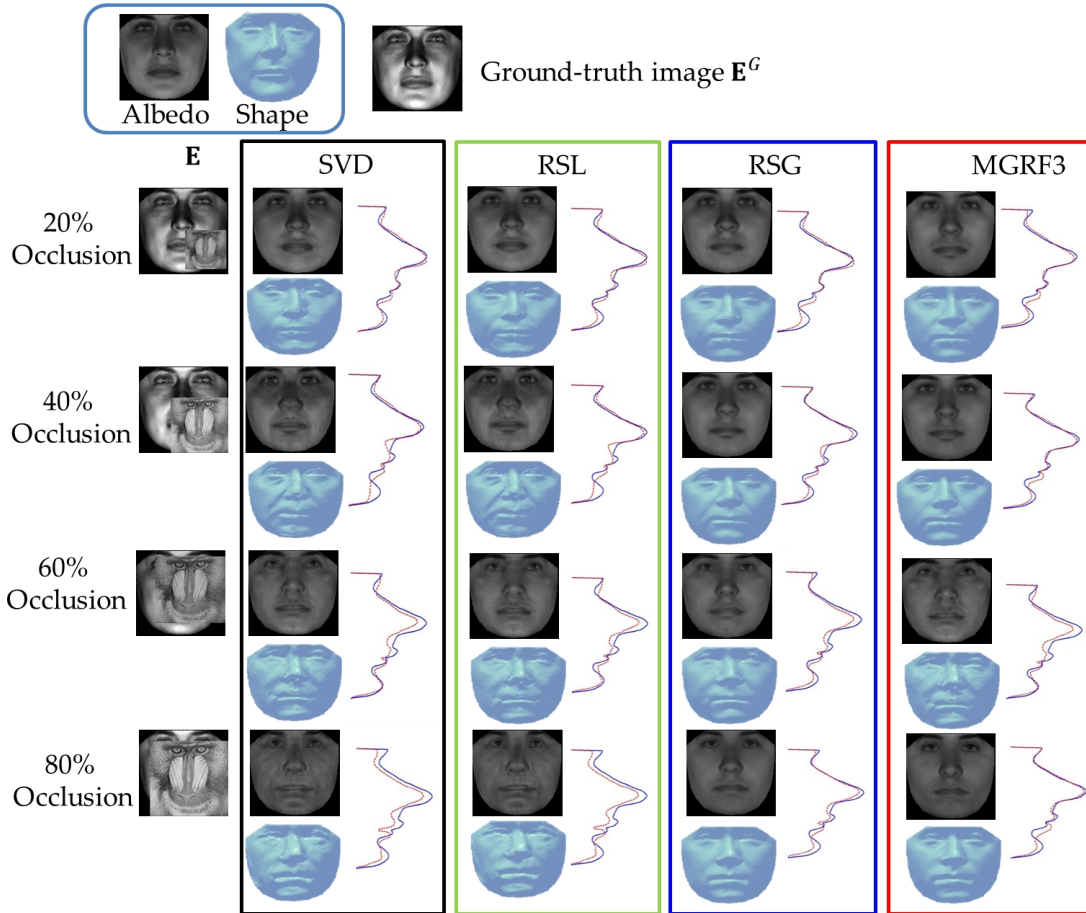


Figure 7.11: Shape and albedo recovery results for one out-of-training sample with combined light sources at $(0,0,1)$ and $(0,0.5,0.9)$. Shape and albedo are recovered using the SHP images computed from Algorithm 3 (SVD), Algorithm 4 with Lorentzian ρ -function (RSL) and with Geman-McClure ρ -function (RSG) and Algorithm 5 (MGRF3). The profile comparison is also shown where solid blue is used for ground-truth shape and dashed red is used for the recovered shape.

The next step is to apply the proposed method to real images of the Extended Yale Database B [29]. Figures 7.12 and 7.13 show recovered shapes and albedos of out-of-training samples using SVD versus the proposed algorithms for SHP computation, where baboon and random images are used as occluders. Visual inspection illustrates that (1) the albedo images belong to the same subject of the input, and (2) the recovered shapes based on SVD is significantly affected by higher levels of occlusion, which emphasizes the need of occlusion handling.

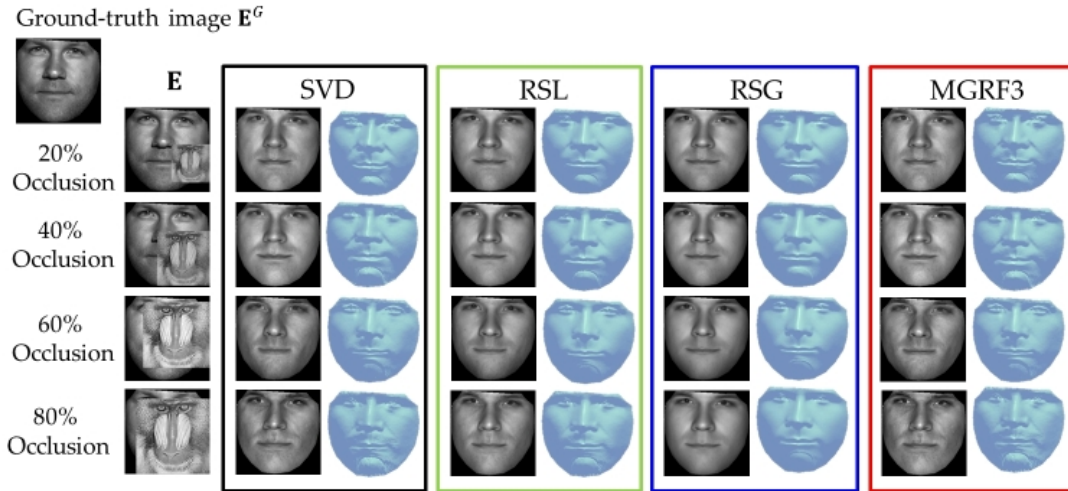


Figure 7.12: Shape and albedo recovery results for one out-of-training sample of Yale Database, where baboon image is used as the occluder.

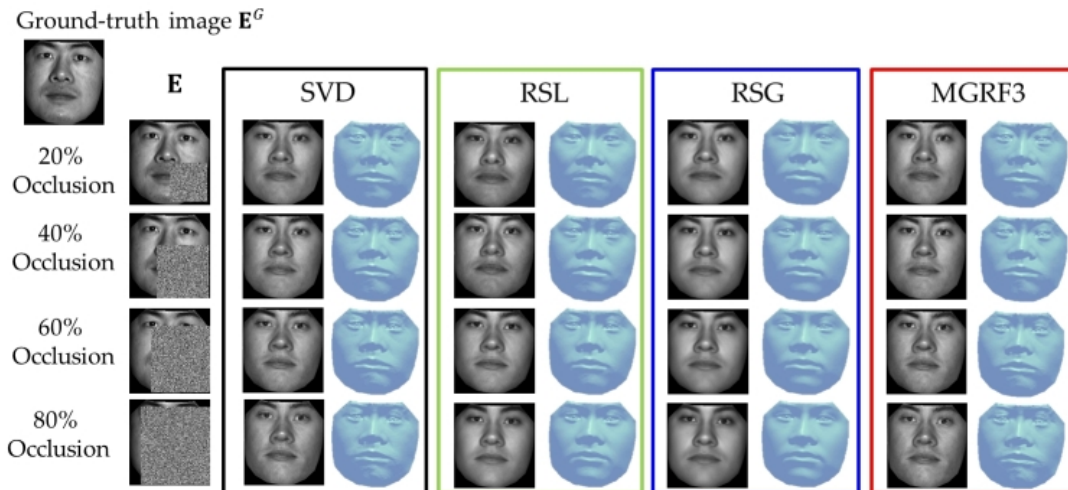


Figure 7.13: Shape and albedo recovery results for one out-of-training sample of Yale Database, where random occluder is used.

7.5 Summary

This chapter considers the problem of reconstructing the image irradiance signal from incomplete irradiance signal where sources of incompleteness might include occlusion and violation to the assumed image formation assumptions. This introduces errors in the reconstructed HP images which in turn has a direct impact on subsequent analysis steps. In particular, cast shadow pixels and those which belong to occluded areas will be incorporated into the projection process unless detected and excluded. This chapter proposed to use robust estimators based on Lorentzian and Geman-McClure

functions within the Graduated Non-Convexity algorithm to estimate the SH coefficients with the presence of statistical outliers (corrupted pixels). On the other hand, assuming corrupted pixels to be likely adjacent to each other, their spatial interaction was modeled using Markov Gibbs random field with Potts model. The results showed the effectiveness of the proposed algorithms in handling high levels of contiguous missing information compared to the conventional image projection. In particular, MGRF and robust estimation using Geman-McClure function outperformed the Singular Value Decomposition (SVD) performance approach which is very sensitive to the presence of missing information even at low levels. In the meantime, the performance of Lorentzian function approaches SVD due to the presence of errors caused by fitting an input image to subspace bases constructed using different geometrical object.

To evaluate the performance of the method, it is tested on both relighted (and real) images of the USF (and Yale) database. Visual inspection on the Yale database reconstructions revealed realistic recovered shape and albedo under high levels of occlusion. The following conclusions are drawn: (1) SVD performance is very sensitive to the presence of occlusion even at low levels. (2) Geman-McClure function compared to Lorentzian function is more robust to errors due to using basis functions of different subjects to recover the shape and albedo of an input image. (3) MGRF and robust estimation using Geman-McClure function outperform the SVD approach. (4) The performance of Lorentzian function approaches SVD due the presence of errors due to using basis of different identity than the shape to be reconstructed.

PART III

APPLICATIONS

MODEL-BASED FACIAL SHAPE RECOVERY FROM SINGLE IMAGE

Through depth perception, humans have the ability to determine distances based on a single 2D image projected on their retina, where shape-from-shading (SFS) provides a mean to mimic such a phenomenon. The goal of this chapter is to recover 3D facial shape from a single image of unknown general illumination, while relaxing the non-realistic assumption of Lambertian reflectance. Prior shape, albedo and reflectance models from real data, which are metric in nature, are incorporated into the shape recovery framework. Adopting a frequency-space based representation of the image irradiance equation, an appearance model is proposed, termed as *Harmonic Projection Images*, which accounts explicitly for different human skin types as well as complex illumination conditions. Assuming skin reflectance obeys Torrance-Sparrow model, it is proved analytically that it can be represented by at most 5th order harmonic basis whose closed form is provided. The recovery framework is a non-iterative approach which incorporates regression-like algorithm in the minimization process. The experiments on synthetic and real images illustrate the robustness of the proposed appearance model vis-à-vis illumination variation.

8.1 Introduction

The shape-from-shading (SFS) problem, formally introduced by the seminal work of Horn [24], involves recovering the 3D shape of an object using the cues of lighting and shading available in a single image. Notwithstanding, such a problem is under-constrained because the image formation process involves numerous unknowns in con-

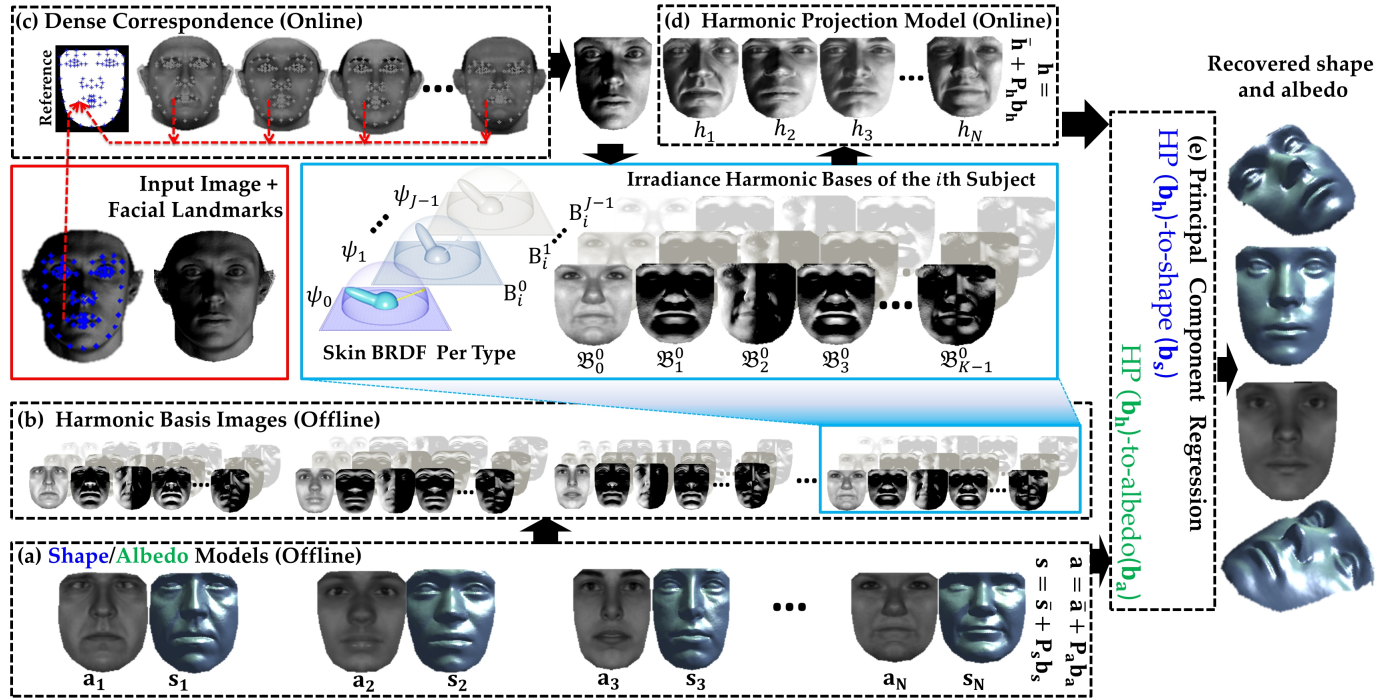


Figure 8.1: Block diagram of the proposed model-based facial shape recovery: (a) An aligned ensemble of facial shapes and albedos is used to build shape/albedo models. (b) Given the albedo and surface normals of a certain subject in the ensemble, harmonic basis images are constructed to encode reflectance properties of different skin types. Given an input image under general unknown illumination and a set of facial landmark points including both anatomical and pseudo-landmarks: (c) Dense correspondence is established between the input image and each subject in the ensemble, where each pixel position within the convex hull of a reference shape corresponds to a certain point on a sample facial scan (shape and albedo) and in the same time to a certain point on the input image. (d) The input image, in the reference frame, is projected onto the subspace spanned by the harmonic basis of each subject in the ensemble which are scaled (using the projection coefficients) and summed-up to construct the harmonic projection (HP) images. A HP image maintains the subject identity of the corresponding harmonic basis while encodes the illumination conditions and the reflectance properties of the input image. The HP images are then used to construct a harmonic projection model of the input image. (e) The inherit relation between the HP images and the corresponding shape/albedo is cast as a regression framework where the two regression models at work here, namely, HP-to-shape and HP-to-albedo models. Principal component regression is used to solve for shape and albedo coefficients to recover the shape and albedo of the input image.

trast to the available information at hand. Thus information inference based on shading commonly involves adding assumptions to the image formation process such as directional light sources and uniform albedo.

Previous works in shape recovery can be broadly categorized in two classes; image-centered and model-based approaches. Starting from the pioneering work of Horn [24], *image-centered* algorithms rely entirely on the available shading cues in the given image. Zhao and Chellapa [125] proposed a symmetric SFS algorithm using the known bilateral symmetry of frontal faces as a geometric constraint. Nonetheless, Prados and Faugeras [126] showed that constraining the SFS algorithm to specific class of objects can improve the accuracy of the recovered shape, while relaxing the uniform albedo assumption. In particular, there has been a significant interest in the recovery of facial shapes given single image. Following Zhao and Chellapa, Ahmed and Farag [127] formulated the symmetric SFS problem as a partial differential equation solved by Lax-Friedrichs sweeping method [128].

On the other hand, *model-based* approaches use statistical prior models which is rooted in the cue theory of depth perception; humans learn the connection between cue and depth through previous experiences. Atick et al. [129] proposed the first statistical SFS method by parameterizing the set of all possible facial surfaces using principal component analysis (PCA). Dovgard and Basri [130] added the geometric constraint of facial symmetry. Smith and Hancock [131] embedded a statistical model of surface normals within a SFS framework. The morphable model framework of [132] estimates the shape and texture coefficients from an input 2D image, together with other scene parameters, using an optimization method based on stochastic gradient descent. It is a 3D extension of the seminal work of Cootes et al. [133] on Active Appearance Models (AAM), where a coupled statistical model is generated to describe the 2D shape and appearance (albedo) information of faces. Castelan et al. [134] developed a coupled statistical model to recover 3D shapes from intensity images with a frontal pose where the 2D shape model in AAM is replaced with a height map model. The main advantage of the Castelan approach over the 3D morphable model framework is in the straightforward recovery of the 3D face shape, which does not go through a costly iterative optimization process, i.e., shape recovery can be performed using a series of matrix operations.

One of the main challenges that confront SFS algorithms is dealing with arbitrary illumination. Basri and Jacobs in [6] and Ramamoorthi and Hanrahan in [7] proved that

images of convex Lambertian object taken under arbitrary illumination conditions can be approximated accurately using low-dimensional linear subspace based on spherical harmonics (SH). Since then, SH has been incorporated in SFS framework to alleviate the problem of illumination assuming Lambertian reflectance [33, 34, 37, 135, 136, 137]. Zhang and Samaras [135] primary goal was to recover the SH basis images from a single test image. Basri and Jacobs [37] relies on multiple images, i.e. photometric stereo, to infer the 3D shape. Ahmed and Farag [136] extended Castelan’s coupled statistical model [134] by combining shape, appearance/albedo and SH in order to parameterize facial surfaces under arbitrary illumination. Rara et al. [137] further extended the work of [136] to include 2D shape information in the model. In subsequent work, Rara et al. [33] decoupled the coupled model of [134] and [136] to obtain a separate model for shape and albedo where the classic brightness constraint in SFS is approximated using SH basis images. Castelan and Horebeek [138] and Rara *et al.* [34] cast their models in a regression framework using the Partial Least Squares (PLS) method which uses a few matrix operations for shape reconstruction to provide a computationally efficient alternative to the iterative methods used in [33, 134].

Despite the simplicity of the Lambertian model, accurately modeling the appearance of objects in general and human faces in particular goes far beyond the Lambertian reflectance model. The human face is neither Lambertian nor convex, introducing cast shadows in concave regions. As such, this inspired oneself to move closer to the realistic situation with a non-Lambertian reflectance model, where non-homogeneous surfaces with spatially varying reflectance properties are often approximated in vision applications by using a spatially varying texture/albedo which modulates the surface radiance [7]. Weyrich *et al.* [95] provided an experimental proof that skin reflectance deviates from the Lambertian assumption by measuring the skin reflectance of 156 subjects of varying age, gendre and race. Their reflectance model consists of spatially-varying analytic reflectance model based on Torrance-Sparrow model¹ [12].

This chapter proposes a framework for 3D face shape and albedo recovery from a single image under unknown general illumination. The frequency-space based representation of the image irradiance equation enables deriving a low-dimensional subspace for human faces with different skin types. Assuming skin reflectance obeys Torrance-

¹The reflection of many dielectrics materials including human skin can be modeled using Torrance-Sparrow [12] reflectance model [139], since they do not exhibit any significant back-scattering properties [140].

Sparrow model [95], it is proved analytically that it can be compactly represented by up to a 5th order harmonic basis which are derived by representing the lighting and the surface reflectance functions using spherical harmonics (SH) and the Helmholtz hemispherical harmonics (HSH)-based reflectance basis [45], respectively. The core component of the proposed framework resides in the appearance model, termed as *Harmonic Projection (HP) Image*, which encodes the illumination conditions and the reflectance properties of an input image. The classical shape-from-shading iterative equation is cast as a regression framework, which can be solved efficiently and non-iteratively using Principal Component Regression (PCR). The experiments show the robustness of the proposed non-Lambertian appearance model vis-à-vis illumination variation compared to the earlier work in [47] which assumed Lambertian reflectance.

8.2 Image Irradiance Basis for Torrance-Sparrow Reflection

Recall that the image irradiance equation for a non-emitting surface with a spatially-varying albedo $\rho(\mathbf{x})$ is defined as,

$$E(\mathbf{x}, \vec{\omega}'_o) = \rho(\mathbf{x}) \int_{\Omega'_i} L_i(R_{\alpha,\beta,\gamma}(\vec{\omega}'_i)) f_r(\vec{\omega}'_i, \vec{\omega}'_o) \cos \theta'_i d\vec{\omega}'_i \quad (8.1)$$

Eqn. Equation 8.1 can be represented in the frequency space by means of expressing the lighting function L_i and the surface reflection f_r in terms of spherical harmonics [11] and the proposed Helmholtz HSH-based basis, respectively. Due to SH basis orthonormality, the rotation process defined in Equation 8.1 is thus a linear operation where coefficients between distinct illumination bands don't interact. Thus the frequency-space representation of the image irradiance in Equation 8.1 can be written as follows¹,

$$E(\mathbf{x}) = \rho(\mathbf{x}) \int_{\Omega'_i} \left(\sum_{n=0}^{\infty} \sum_{m=-n}^n \sum_{m'=-n}^n l_n^m \mathcal{D}_{mm'}^n(\alpha, \beta) Y_n^{m'}(\vec{\omega}'_i) \right) \left(a_0 + 2 \sum_{p=1}^{\infty} a_p \mathcal{H}_{p0}^{00}(\vec{\omega}'_i, \vec{\omega}'_o) \right) \cos \theta'_i d\vec{\omega}'_i \quad (8.2)$$

where l_n^m are the illumination spectrum coefficients when projected onto the subspace spanned by the real SH basis, \mathcal{D}^n are Wigner's rotation matrices [59] and $\vec{\omega}'_o = (0, \pi)$.

¹ γ is arbitrarily set to zero since TS model is isotropic while normal-exitance is assumed.

Note that Torrance-Sparrow (TS) spectrum coefficients has two components; the zeroth order a_o which is due to the diffuse component of the BRDF, while the summation term is due to the specular part. Thus it is anticipated to have such division being reflected on the image irradiance basis.

The diffuse component is similar to the Lambertian case in [6] and [11], moving terms which are independent of the local incident direction out of the integral results in¹,

$$\begin{aligned} \sum_{n=0}^{\infty} \sum_{m=-n}^n \sum_{m'=-n}^n l_n^m a_0 \mathcal{D}_{mm'}^n(\alpha, \beta) \int_{\Omega'_i} Y_n^{m'}(\vec{\omega}'_i) \cos \theta'_i d\vec{\omega}'_i \\ = \sum_{n=0}^{\infty} \sum_{m=-n}^n l_n^m a_0 \mathcal{D}_{m0}^n(\alpha, \beta) c_n \end{aligned} \quad (8.3)$$

where,

$$c_n = \begin{cases} \frac{\sqrt{\pi}}{2} & n = 0 \\ \sqrt{\frac{\pi}{3}} & n = 1 \\ \frac{(-1)^{\frac{n}{2}+1} \sqrt{\pi(2n+1)}}{2^n(n-1)(n+2)} \binom{n}{n/2} & n \geq 2, \text{ even} \\ 0 & n \geq 2, \text{ odd} \end{cases} \quad (8.4)$$

Using series representation of the associated Legendre polynomial, the specular component can be written as,

$$\begin{aligned} \sum_{n=0}^{\infty} \sum_{m=-n}^n \sum_{m'=-n}^n \sum_{p=1}^{\infty} 2l_n^m a_p \mathcal{D}_{mm'}^n(\alpha, \beta) \\ \int_{\Omega'_i} Y_n^{m'}(\vec{\omega}'_i) \mathcal{H}_{p0}^{00}(\vec{\omega}'_i, \vec{\omega}'_o) \cos \theta'_i d\vec{\omega}'_i \\ = \sum_{n=0}^{\infty} \sum_{m=-n}^n \sum_{p=1}^{\infty} 2l_n^m a_p \mathcal{D}_{m0}^n(\alpha, \beta) C_{np} \end{aligned} \quad (8.5)$$

¹All terms vanish except for $m' = 0$ due to the orthogonality relation of the azimuthal functions *w.r.t.* polynomial degree.

where,

$$\begin{aligned}
 C_{np} = & \sqrt{\frac{2p+1}{2}} \left[\left(\frac{(-1)^p}{2^n} \sqrt{\frac{2n+1}{4\pi}} \right. \right. \\
 & \left. \left. \sum_{k=0}^{\lfloor n/2 \rfloor} \sum_{k'=0}^p \frac{(-1)^{k+k'} \binom{n}{k} \binom{2n-2k}{n}}{n-2k+k'+2} \right) \right. \\
 & \left. + \begin{cases} \frac{\sqrt{\pi}}{2} & n=0 \\ \sqrt{\frac{\pi}{3}} & n=1 \\ \frac{(-1)^{\frac{n}{2}+1} \sqrt{\pi(2n+1)}}{2^n(n-1)(n+2)} \binom{n}{n/2} & n \geq 2, \text{ even} \\ 0 & n \geq 2, \text{ odd} \end{cases} \right] \quad (8.6)
 \end{aligned}$$

This suggests that the image irradiance equation of a convex surface with albedo $\rho(\cdot)$ and surface normals $\vec{n}(\cdot)$ and whose reflectance obeys TS reflection model can be represented as a linear combination of harmonic basis functions of the form:

$$\mathcal{B}_n^m(\rho(\mathbf{x}), \vec{n}(\mathbf{x})) = \rho(\mathbf{x}) \mathcal{D}_{m0}^n(\alpha, \beta) \left(a_0 c_n + 2 \sum_{p=1}^{\infty} a_p C_{np} \right) \quad (8.7)$$

In the image space, assuming the light source and the viewer are far from the object where orthographic projection is a good model for the imaging process, the image intensity I can inherit the image irradiance E of the corresponding surface point.

Formally, let a 3D object with D surface points be defined by its shape $\mathbf{s} = [\vec{n}(\mathbf{x}_0), \vec{n}(\mathbf{x}_1), \dots, \vec{n}(\mathbf{x}_{D-1})]^T$, albedo $\mathbf{a} = [\rho(\mathbf{x}_0), \rho(\mathbf{x}_1), \dots, \rho(\mathbf{x}_{D-1})]^T$ and TS-BRDF parameters $\psi = \{\rho_d, \rho_s, \sigma\}^1$. One computes the harmonic basis functions in Equation 8.7 for each surface point which is orthographically projected on the retinal plane through the imaging process. Hence the harmonic basis is referred to as images in the image space.

Using the N th order harmonic approximation of the image irradiance equation, the image of this object under fixed-pose with arbitrary lighting function can be efficiently approximated as a linear combination of the harmonic basis images such that,

$$I(\mathbf{x}) \approx \sum_{n=0}^N \sum_{m=-n}^n l_n^m \mathcal{B}_n^m(\rho(\mathbf{x}), \vec{n}(\mathbf{x})) \quad (8.8)$$

In matrix notation; Let \mathbf{B} denotes $D \times K$ matrix whose columns are the harmonic

¹Note that the BRDF parameters affect the reflectance spectrum $a_p \forall p \geq 0$.

basis images $\mathcal{B}_k \forall k = n(n+1) + m$ with $K = (N+1)^2$, typically $K = 36$ in case of human skin (see Figure 4.9). Let $\ell = (l_0, l_1, \dots, l_{K-1})^T$ denote the lighting coefficients. Thus Equation 8.8 can be rewritten as follows,

$$I \approx \mathbf{B} \ell \tag{8.9}$$

The next section presents how to devise this representation to transfer the illumination conditions of an input image to a given sample of the statistical prior to aid the shape recovery process.

8.3 Harmonic Projection Images

The classical brightness constraint in SFS measures the deviation of the input image from the image irradiance under the assumed image formation process, it can be defined as;

$$\epsilon = \int \int [I(\mathbf{x}) - E(\mathbf{x})]^2 d\mathbf{x} \tag{8.10}$$

Representing the image irradiance in terms of harmonic basis Equation 8.7 allows inferring the illumination conditions of a single image as follows. Given an input image I of D -pixels and the harmonic basis matrix \mathbf{B} of a corresponding 3D object (a human face in particular), the lighting coefficients $\hat{\ell} = [l_0, \dots, l_{K-1}]^T$ can be deduced as the minimal solution of the over-determined linear system of equations $I = \mathbf{B}\ell$ that best matches the input image. The least squares solution can be obtained using Singular Value Decomposition such that $\hat{\ell} = \mathbf{V}\mathbf{S}^{-1}\mathbf{U}^T I$ where $\mathbf{B} = \mathbf{U}\mathbf{S}\mathbf{V}^T$. The reconstructed image $h = \mathbf{B}\hat{\ell}$ encodes the illumination conditions of the input image I .

The harmonic basis matrix works best when a single variable in the image formation process is only considered, *i.e.* illumination. However different human skin types (*i.e.* surface reflectance properties) affect the appearance of the facial surface even if illumination is held fixed. It has been shown in [17, 141] that the BRDF of many objects can be expressed in terms of fundamental materials in a data-driven context. Thus the image irradiance corresponding to the input image can be written as,

$$E(\mathbf{x}) \approx \sum_{j=0}^{J-1} w_j \sum_{k=0}^{K-1} l_k \mathcal{B}_k^j(\rho(\mathbf{x}), \vec{n}(\mathbf{x})) \tag{8.11}$$

where $\varpi = [w_0, \dots, w_{J-1}]^T$ are the weighting factors for the BRDF of different skin types.

This inspired oneself to exploit the dependency of the harmonic basis images on the BRDF parameters, where one can infer surface reflectance properties from the input image by incorporating the reflectance properties of different skin types in the prior model. This is achieved by assigning multiple BRDFs $\Psi = [\psi_0, \dots, \psi_{J-1}]$ to a pair of shape and albedo, where ψ_j corresponds to TS parameters of the j th skin type¹. Thus \mathbf{B}_i^j denotes the basis matrix of the i th shape/albedo-pair with the j th skin type.

As such, the i th *harmonic projection (HP) image* \mathbf{h}_i of an input image is defined as the reconstructed image from the process of projecting the input image to the subspace spanned by the harmonic basis matrices $\mathbf{B}_i^j \forall j$ then taking the reverse process by summing up the scaled basis using the computed harmonic coefficients $\widehat{\ell}^j$ and the weighting factors $\widehat{\varpi}^i$, *i.e.* with $\mathbf{B}_i^j = \mathbf{U}_i^j \mathbf{S}_i^j (\mathbf{V}_i^j)^T$ one has,

$$\mathbf{h}_i = \sum_{j=0}^{J-1} \widehat{\varpi}_j^i \mathbf{B}_i^j \widehat{\ell}^j \quad ; \quad \widehat{\ell}^j = \mathbf{V}_i^j (\mathbf{S}_i^j)^{-1} (\mathbf{U}_i^j)^T I \quad (8.12)$$

where the BRDF weights are the least squares solution of the over-determined system $I = \mathbf{H}_i \varpi^i$ with $\mathbf{H}_i = [\mathbf{h}_i^0, \dots, \mathbf{h}_i^{J-1}]$ be the matrix whose columns $\mathbf{h}_i^j = \mathbf{B}_i^j \widehat{\ell}^j$ are the reconstructed images from the subspace spanning the j th skin type for the i th prior sample.

It is worth noting that \mathbf{h}_i^j encodes the illumination conditions of the input image being transferred to the 3D object used to generate basis \mathbf{B}_i^j while maintaining the i th object identity and reflectance. Meanwhile \mathbf{h}_i encodes both the illumination and reflectance properties of the input image while maintaining the i th object identity. Thus the shape recovery problem can then be formulated as finding the harmonic projection image \mathbf{h} which maintains the illumination, reflectance and identity of the input image. As such, the discrete version of the brightness constraint becomes,

$$\epsilon = \sum_{\mathbf{x}} [I(\mathbf{x}) - h(\mathbf{x})]^2 \quad (8.13)$$

The existence of a statistical prior of shape and albedo aids resolving the bas-relief ambiguity [142] which normally occurs when both lighting and shape are unknown. Meanwhile, it enables finding the shape and albedo which minimizes the objective func-

¹In the conducted experiments, average parameters for the cheek region of the skin types defined in [95] were used, thus $J = 6$.

tion in Equation 8.13 in a regression-like framework which is detailed in the next section.

8.4 Model-based Shape Recovery

The statistical SFS framework in this paper involves two *offline* models, namely, the shape and albedo models and one *online* model, termed HP-model. By offline one means that such models are constructed during the training stage compared to the online one which is formed once the input image is given. The process of constructing the shape and albedo models involves two major steps: (a) establishing a dense correspondence among an ensemble of human faces and (b) statistical modeling. The model construction in this paper is similar to the morphable model construction in [132, 143].

Algorithm 6 Principal Component Regression Framework for 3D Shape Recovery

INPUT: (a) Input image, I_{inp} (b) Shape and albedo samples: $(\mathbf{s}_1, \mathbf{a}_1)$ to $(\mathbf{s}_n, \mathbf{a}_n)$

OUTPUT: (a) Recovered shape, $\tilde{\mathbf{s}}$ (b) Recovered albedo, $\tilde{\mathbf{a}}$

- 1: **Build the shape and albedo models from the samples using PCA:** construct $\mathbf{s} = \bar{\mathbf{s}} + \mathbf{P}_s \mathbf{b}_s$ and $\mathbf{a} = \bar{\mathbf{a}} + \mathbf{P}_a \mathbf{b}_a$.
 - 2: **Construct the basis images for each pair i and for each assigned BRDF ψ_j :** $\mathbf{B}_i^j = [\mathcal{B}_{0,i}^j, \dots, \mathcal{B}_{K-1,i}^j] \in \mathbb{R}^{D \times K}$ for each pair $(\mathbf{s}_i, \mathbf{a}_i)$
 - 3: **Compute the lighting coefficients for each pair i and for each assigned BRDF ψ_j :** $\hat{\ell}_i^j = \mathbf{V}_i^j (\mathbf{S}_i^j)^{-1} (\mathbf{U}_i^j)^T I_{inp}$ where $\mathbf{B}_i^j = \mathbf{U}_i^j \mathbf{S}_i^j (\mathbf{V}_i^j)^T$.
 - 4: **Compute the reconstructed images from the j th skin type subspace for each pair i :** $\mathbf{H}_i = [\mathbf{h}_i^0, \dots, \mathbf{h}_i^{J-1}]$ where $\mathbf{h}_i^j = \mathbf{B}_i^j \hat{\ell}_i^j$.
 - 5: **Compute the BRDF weights for each pair i :** $\hat{\omega}_i = \mathbf{V}_i^h (\mathbf{S}_i^h)^{-1} (\mathbf{U}_i^h)^T I_{inp}$ where $\mathbf{H}_i = \mathbf{U}_i^h \mathbf{S}_i^h (\mathbf{V}_i^h)^T$.
 - 6: **Compute the harmonic projection (HP) image for each pair i :** $\mathbf{h}_i = \sum_{j=0}^{J-1} \hat{\omega}_j^i \mathbf{B}_i^j \hat{\ell}_i^j$.
 - 7: **Build the HP model:** construct $\mathbf{h} = \bar{\mathbf{h}} + \mathbf{P}_h \mathbf{b}_h$.
 - 8: **Replace the shape samples with its coefficients:** solve for $\mathbf{b}_{si} = P_s^T (\mathbf{s}_i - \bar{\mathbf{s}})$
 - 9: **Replace the albedo samples with its coefficients:** solve for $\mathbf{b}_{ai} = P_a^T (\mathbf{a}_i - \bar{\mathbf{a}})$
 - 10: **Replace the HP images with its coefficients:** solve for $\mathbf{b}_{hi} = P_h^T (\mathbf{h}_i - \bar{\mathbf{h}})$
 - 11: **Setup matrices for PCR:** $\mathbf{T} = [\mathbf{b}_{h_1}^T, \dots, \mathbf{b}_{h_n}^T]$, $\mathbf{U}_{sh} = [\mathbf{b}_{s_1}^T, \dots, \mathbf{b}_{s_n}^T]$ and $\mathbf{U}_{al} = [\mathbf{b}_{a_1}^T, \dots, \mathbf{b}_{a_n}^T]$
 - 12: **Build two PCR models:** construct $\tilde{\mathbf{C}}_{sh} = (\mathbf{T}^T \mathbf{T})^{-1} \mathbf{T}^T \mathbf{U}_{sh}$ and $\tilde{\mathbf{C}}_{al} = (\mathbf{T}^T \mathbf{T})^{-1} \mathbf{T}^T \mathbf{U}_{al}$
 - 13: **Solve for the SHP coefficients of the input image:** get $\mathbf{b}_{h,inp} = P_h^T (I_{inp} - \bar{\mathbf{h}})$
 - 14: **Solve for the shape and albedo coefficients:** get $\tilde{\mathbf{b}}_s^T = \tilde{\mathbf{b}}_h^T \tilde{\mathbf{C}}_{sh}$ and $\tilde{\mathbf{b}}_a^T = \tilde{\mathbf{b}}_h^T \tilde{\mathbf{C}}_{al}$
 - 15: **Solve for the recovered shape and albedo:** $\tilde{\mathbf{s}} = \bar{\mathbf{s}} + \mathbf{P}_s \mathbf{b}_s$ and $\tilde{\mathbf{a}} = \bar{\mathbf{a}} + \mathbf{P}_a \mathbf{b}_a$
-

The shape and albedo models can be constructed offline by performing principal component analysis (PCA) on the aligned set of database samples, the resulting shape model is $\mathbf{s} = \bar{\mathbf{s}} + \mathbf{P}_s \mathbf{b}_s$ where $\bar{\mathbf{s}}$ is the mean shape, \mathbf{P}_s are the shape eigenvectors and \mathbf{b}_s is the set of shape coefficients. Similarly $\mathbf{a} = \bar{\mathbf{a}} + \mathbf{P}_a \mathbf{b}_a$ is the albedo model where $\bar{\mathbf{a}}$ is

the mean albedo, \mathbf{P}_a are the albedo eigenvectors and \mathbf{b}_a is the set of albedo coefficients. These two models can be trained offline.

The HP model, on the other hand, is trained online to incorporate the illumination and reflectance conditions of the input image; given an image I and the basis images \mathbf{B}_i^j of object instance i , the HP image \mathbf{h}_i is obtained as in Equation 8.12. After reconstructing the HP images of all the instances in the database, the HP images can be modeled using PCA as $\mathbf{h} = \bar{\mathbf{h}} + \mathbf{P}_h \mathbf{b}_h$ where $\bar{\mathbf{h}}$ is the mean HP image, \mathbf{P}_h are the HP images eigenvectors and \mathbf{b}_h is the set of HP coefficients.

There is an inherited relation between the HP images \mathbf{h} and shape \mathbf{s} and albedo \mathbf{a} simultaneously, hence shape and albedo recovery can be cast in a regression framework having two regression models (a) HP-to-shape and (b) HP-to-albedo. This chapter uses Principal Component Regression (PCR) to obtain the relationship between the independent data (HP images) and the dependent data (shape/albedo) through decomposing them into a low-dimensional subspace, i.e., replacing the high-dimensional vectors \mathbf{s}_i , \mathbf{a}_i , and \mathbf{h}_i by their respective PCA coefficients (\mathbf{b}_{s_i} , \mathbf{b}_{a_i} , and \mathbf{b}_{h_i}). Then, standard multiple linear regression is performed between the low-dimensional representations. This method is summarized in Algorithm 6.

8.5 Experimental Results

8.5.1 Model Construction

The shape and albedo models in this paper are derived from 3D scans of the USF database [118], where 80 subjects of its 100 samples were deemed to be acceptable to build the shape and albedo models. The original 3D scans are converted into a Monge patch format [8] which represents the surface as $(x, y, f(x, y))$ ¹. The front (for albedo information) and depth buffers (for the corresponding height information) within the frame buffer in a computer graphics system are simulated [144]. The approach of Patel et al. [143] is used to establish the dense correspondence between database samples, where a set of sparse landmark points (anatomical and pseudo-landmark points [145]) are manually annotated on the resulting albedo data for all database samples. Generalized

¹This representation is convenient because a unique point on the surface can be determined by specifying only the image coordinates. Castelan et al. [134] uses the term, height map, for a similar concept.

Procrustes Analysis (GPA)[146] is first performed to align the set of shapes to a common reference frame. The average of the aligned shapes define the reference shape which is crucial in establishing dense correspondence between USF samples, see Figure 8.1. Each pixel within the convex hull of the reference shape corresponds to a certain point on each USF sample scan through a physically motivated thin-plate splines [146] warping function.

8.5.2 Synthetic Images

To quantify the reconstruction accuracy of the proposed framework, one recovers the 3D shape for 80 out-of-training USF samples illuminated by directional light sources defined in the Extended YaleB database [29], while reflectance properties are randomly drawn from the skin reflectance database [95]. For each experiment, the following measures are computed: (a) Height Error - the recovered height map is compared with the ground truth height and the mean absolute error \bar{s}_{err} is then computed, and (b) Surface Orientation Error - the directions of the recovered normal vectors are compared with the ground truth data. The average of the difference angle $\bar{\theta}_{err}$ is then computed.

The histogram of the mean height absolute error and mean surface orientation error are displayed in Figure 8.2. The maximum height error is less than 6%, with most samples in the (1.5-5%) range, while the mean surface orientation error does not exceed 0.08 rad. Note the increasing error levels throughout different illumination subsets where subset 4 is considered the most challenging compared to other subsets. It is worth mentioning that based on one’s earlier work in [47], the Lambertian case reported average height error of 5.4% with 0.15 rad as average surface orientation error.

Table 8.1: Mean and standard deviation of the mean height error for 80 out-of-training USF samples under Yale light conditions

	Subset 1	Subset 2	Subset 3	Subset 4
$\mu_{\bar{s}_{err}}$ (%)	3.03	3.09	3.24	3.37
$\sigma_{\bar{s}_{err}}$ (%)	1.03	1.02	1.04	1.01

Table 8.2: Mean and standard deviation of the mean surface orientation error for 80 out-of-training USF samples under Yale light conditions

	Subset 1	Subset 2	Subset 3	Subset 4
$\mu_{\bar{\theta}_{err}}$ (rad)	0.05	0.05	0.051	0.052
$\sigma_{\bar{\theta}_{err}}$ (rad)	0.0065	0.0069	0.0077	0.0081

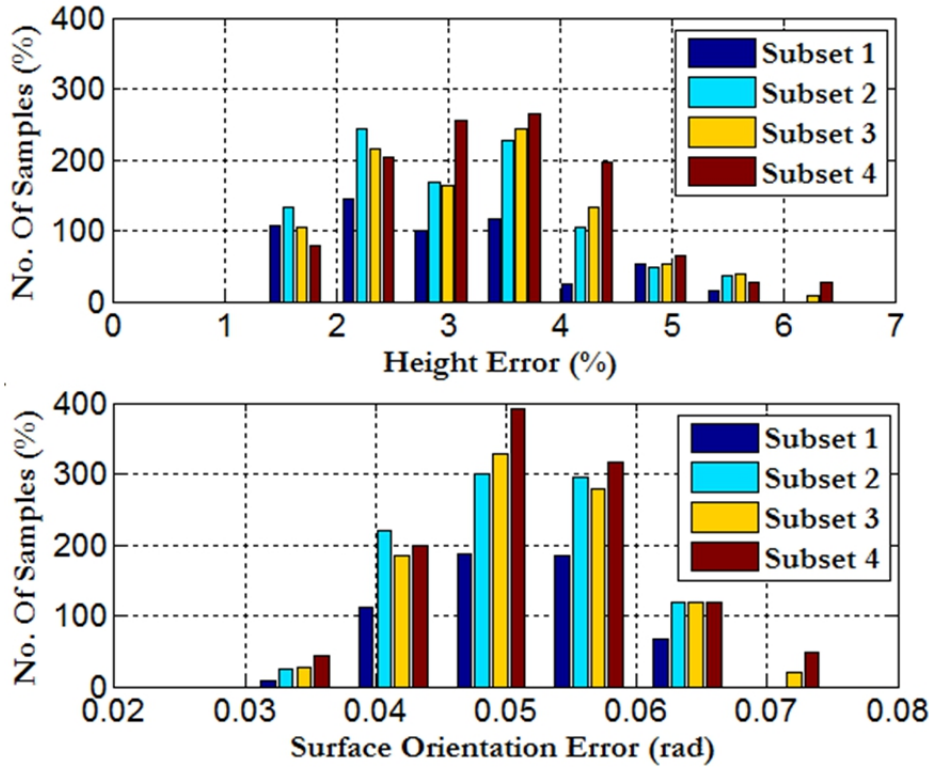


Figure 8.2: Error histograms of the recovered shape for 80 out-of-training USF-samples under illumination conditions of the extended YaleB database: Top: height error histogram. Bottom: surface orientation error histogram.

Figure 8.3 shows the recovered shapes for different out-of-training USF samples together with the input image and the ground-truth shape. Notice that the results are visually acceptable. The shape error maps indicate that the recovered shapes are close to ground-truth values.

8.5.3 Real Images

The next step is to apply the proposed method to real images of the Extended Yale Database B [29] which is a standard database for illumination studies in vision, esp. face recognition. Note that the database subjects are not included in the prior model. Figure 8.4 shows the input images under the same illumination conditions as in Figure 8.3, together with the recovered shape and albedo. For visual comparison, the shape of each sample is reconstructed using photometric stereo [8]. It is worthwhile to note that photometric stereo needs multiple images of a fixed-pose subject under different known illumination conditions. This explains the detailed shape reconstruction in Figure 8.4(b) when compared to Figure 8.4(d). However such luxury might not be available in case of having a single image under unknown illumination, hence the proposed

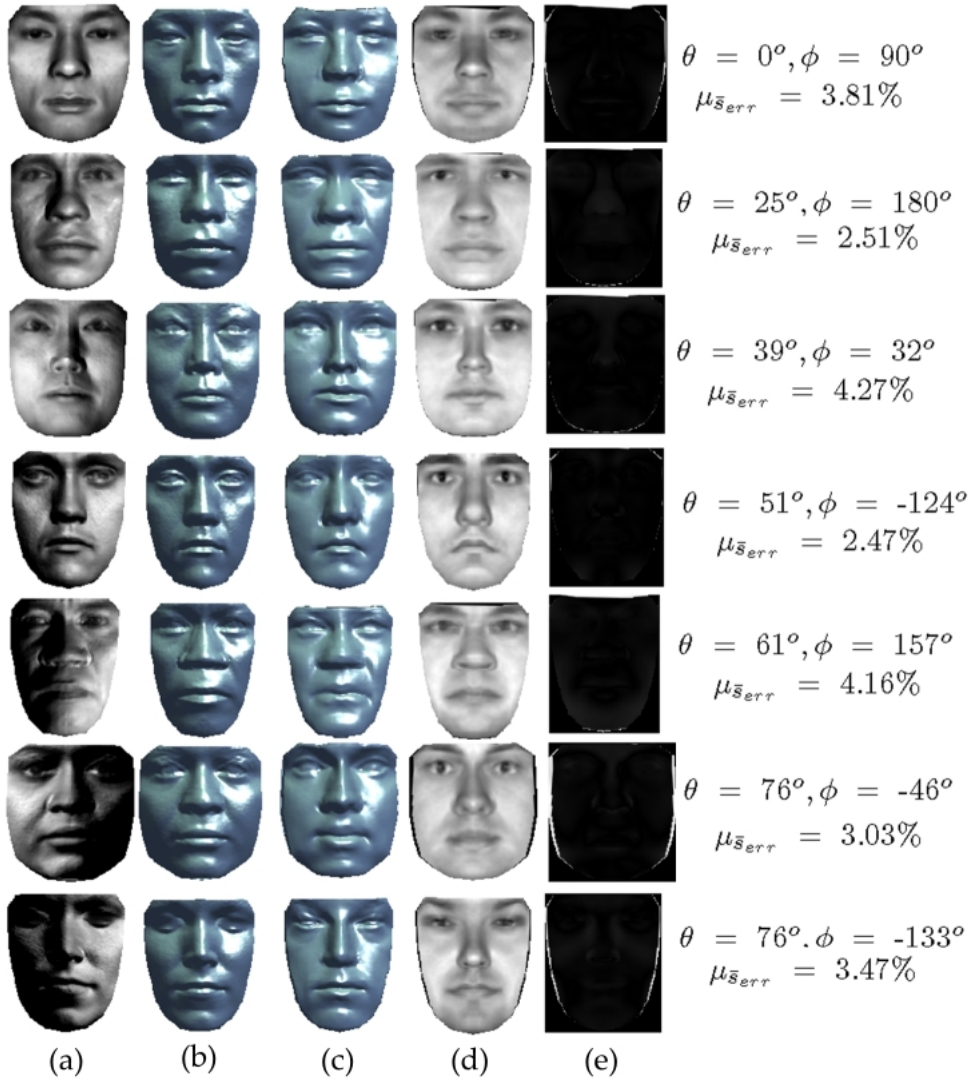


Figure 8.3: Recovered shapes for different out-of-training USF samples, (a) Input image under directional light source of inclination angle θ and azimuthal angle ϕ , (b) ground-truth shape, (c) recovered shapes, (d) recovered albedos and (e) error maps with average height error shown in percentages.

framework provides a solution given the image brightness value per pixel.

8.6 Summary

This chapter presented a model-based facial shape recovery framework for general and unknown illumination which is made possible through the frequency space representation of the image irradiance equation. Based on recent studies of human skin reflectance, image irradiance harmonic basis is derived based on Torrance-Sparrow reflection model. The classical shape-from-shading equation was cast as a regression problem which results in a sequence of matrix operations, enabling shape recovery at higher rate compared to



Figure 8.4: Recovered shapes for different YaleB Database subjects under same illumination conditions as in Figure 8.3, (a) Input image, (b) recovered shapes from photometric stereo using eight images from subsets 1 and 2 with known light directions, (c) recovered shapes and (d) recovered albedos.

the iterative counterparts. To evaluate the performance of the proposed harmonic basis for the task of shape recovery, the proposed framework was tested on both relighted (and real) images of the USF (and Yale) database. Experiments on the out-of-training samples of the USF database provide convincing results. Visual inspection on the Yale database

reconstructions reveal realistic recovered shape and albedo compared to the situation of photometric stereo which requires multiple images of known light sources. Ongoing efforts are directed towards investigating the concave/convex ambiguity resulting from the existence of singular points (*i.e.* points with maximum brightness).

CLINICAL CROWNS SHAPE RECONSTRUCTION - AN IMAGE-BASED APPROACH

Several existing 3D systems for dental applications rely on obtaining an intermediate solid model of the jaw (cast or teeth imprints) from which the 3D information can be captured. This chapter proposes a model-based shape-from-shading (SFS) approach which allows for the construction of plausible human jaw models *in vivo*, without ionizing radiation, using fewer sample points in order to reduce the cost and intrusiveness of acquiring models of patients teeth/jaws over time. The inherent relation between the photometric information and the underlying 3D shape is formulated as a statistical model where the coupled effect of illumination and reflectance is modeled using the proposed Helmholtz Hemispherical Harmonics (HSH)-based irradiance harmonics whereas the Principle Component Regression (PCR) approach is deployed to carry out the estimation of dense 3D shapes. Moreover, shape and texture alignment is accomplished using a proposed definition of anatomical jaw landmarks which can be automatically detected. Vis-à-vis dental applications, the results demonstrate a significant increase in accuracy in favor of the proposed approach. In particular, the proposed approach is able to recover geometrical details of tooth occlusal surface as well as mouth floor and ceiling as compared to SFS-based approaches. Further, the robustness of the proposed approach is investigated in an automated framework for tooth restoration from a single optical image. The system is evaluated on database of 16 jaws for inlay and onlay restorations. Results shows a promising performance for using the proposed approach

in clinical application¹.

9.1 Introduction

Object modeling from a single image, augmented with prior information, facilitates various studies and applications in art, design, reverse engineering, rapid prototyping and basic analysis of deformations and uncertainties. Without the use of ionizing radiation (e.g. X-ray and Computer Tomography - CT), object modeling involves constructing a 3D representation for the information conveyed in the given 2D images. This problem has been studied in the past four decades resulting in many solutions bundled under the name *shape-from-X*. In particular, techniques, such as shape-from-shading (SFS) provide promise of image-based 3D reconstruction when the imaging environment is somewhat precise.

To motivate the contribution of this work, one considers a dental application; 3D reconstruction of the visible part of the human jaw, *a.k.a. clinical crown*. Dentistry usually require accurate 3D representation of the teeth and jaw for diagnostic and treatment purposes. For instance, the precise knowledge of the 3D shape of the tooth is becoming crucial for the treatment of malocclusion problems and several endodontic procedures where accurate location and orientation of the tooth 3D shape must be determined for the sake of the treatment plan. Further, this 3D representation could play a vital role in constructing tooth implants where crowns and bridges of high quality are needed [147]. Meanwhile, since teeth are the most durable parts in the body, they are frequently used in forensic medicine [148, 149]. As such a good Computer Aided-Design (CAD) system should integrate the advantages of constructing surfaces which fit to the existing articulation of the tooth along with retaining its morphological features [148].

Several existing 3D systems for dental applications found in literature rely on obtaining an intermediate solid model of the jaw (cast or teeth imprints) and then capturing the 3D information from that model, e.g. [150]. There may therefore be a demand for intraoral measurement that could be fulfilled by photogrammetry which seems to offers a reduced cost technique while avoiding the need for castings.

Our argument of image-based approach for 3D reconstruction as an alternative to CT-scanning is based on the following. During the exposure to diagnostic imaging using

¹This work has been funded by NSF funding Grant NSF IIS-0513974 and University of Louisville, "3D Modeling of the Human Jaw".

X-ray (ionizing/electromagnetic radiation), the patient's body is penetrated by millions of X-ray photons whose ionization can damage the body's molecules especially DNA in chromosomes. Most DNA damage is repaired immediately, but rarely a portion of a chromosome may be permanently altered (a mutation) leading ultimately to the formation of a tumor [151]. While doses and risks for dental radiology are small, a number of epidemiological studies have provided evidence of an increased risk of brain [152], salivary gland [153] and thyroid tumors [154] for dental radiography. Also, pregnant mothers undergoing diagnostic or therapeutic procedures involving ionizing radiation, or who may be exposed to environmental radiation, there is a great potential for damage to the early embryo [155]. These effects are believed to have no threshold radiation dose below which they will not occur [156]. On the other hand, CT-scanning is considered expensive and not paid by insurance companies unless disease oriented. Meanwhile, dental offices in rural areas do not have such a luxury. Thus the intent is to develop a purely image-based reconstruction mechanism as a cost-effective information tool for the dentist.

This chapter aims at making it easy and feasible for doctors, dentists, and researchers to obtain models of a person's jaw *in vivo*, without ionizing radiation, using fewer sample points in order to reduce the cost and intrusiveness of acquiring models of patients teeth/jaws over time. This is a challenging problem due to the "unfriendly" environment of taking measurements inside a person's mouth [157]. Further assumptions of the presence of distinct features or texture regions on the object in stereo images and the photo consistency in space carving are rarely valid in practice.

Due to the lack of surface texture, shape-from-shading (SFS) algorithms have been used to reconstruct the 3D shape of human teeth/jaw due to the significant shading cue presented in an intra-oral image, e.g. [158]. Nonetheless, in principle, SFS is an ill-posed problem, Prados and Faugeras [126] showed that constraining the SFS problem to a specific class of objects can improve the accuracy of the recovered shape. Thus the main objective of the presented work is to develop and validate a holistic approach for image-based 3D reconstruction of the human jaw based on statistical shape-from-shading (SSFS), covering regions which the classical SFS approach does not handle, using scanned molds and images of the oral cavity to estimate the shape of a human jaw in order to create a more accurate jaw 3D model. In specific, the structure of the human jaw reveals what can be acquired in terms of prior information to enhance the

SFS process where the upper and lower jaws are symmetric and lined up according to specific anatomical features and landmarks. This approach is believed to have the potential to greatly improve plausibility of the resulting shape from shading models.

9.2 Related Work

There has been a substantial amount of work regarding statistical shape recovery for human face modeling and biomedical structures with distinct shapes - e.g., modeling the corpus callosum, the kidney and spinal cord; it is an active research area under shape and appearance modeling (e.g., [34, 131]). Atick et al. [129] proposed the first statistical shape-from-shading (SSFS) method where Principal Component Analysis (PCA) was used to parameterize the set of all possible facial surfaces. Seminal work of Cootes et al. [133] used a coupled statistical model to describe the 2D shape and appearance/albedo of faces. In the 3D venue, Blanz and Vetter [132] refined Attick's work by modeling both facial shape and texture separately. Then scene parameters such as pose and illumination were estimated in the process of a morphable model fitting using a stochastic gradient descent-based optimization. By considering the statistical constraint of [129] and the geometric constraint of symmetry in [125], Dovgard and Basri [130] introduced a statistical symmetric SFS method. Smith and Hancock [131] modeled surface normals within the framework of SSFS. Based on Active Appearance Models (AAM) concept of Cootes et al. [133], Castelan et al. [134] developed a coupled statistical model to recover the 3D shape from intensity images with frontal light source, where the 2D shape model in [133] is replaced with a 3D shape model composed of height maps. The main advantage of Castelan's approach over the 3D morphable model framework [132] is the straightforward recovery of the 3D face shape, without undergoing a costly optimization process.

One of the main challenges that confront SFS algorithms is dealing with arbitrary illumination. Basri and Jacobs [6] proved that images of convex Lambertian object taken under arbitrary distant illumination conditions can be approximated accurately using low-dimensional linear subspace based on Spherical Harmonics (SH). This has also been validated for near illumination conditions [65]. Since then, SH was incorporated in SFS framework to tackle the problem of illumination [33, 34, 135, 136, 138]. Zhang and Samaras [135] primary goal was to recover the SH basis images from a single test

image. Basri and Jacobs [37] relied on multiple images, i.e. photometric stereo, to infer the 3D shape. Ahmed and Farag [136] extended Castelan's coupled statistical model [134] by combining shape, albedo and SH in order to parameterize facial surfaces under arbitrary illumination. Rara et al. [137] further extended the work of Ahmed and Farag [136] to include 2D shape information in the model. In subsequent work, Rara et al. [33] decoupled the coupled model of [159] and [136] to obtain a separate model for shape and albedo where the classic brightness constraint in SFS is approximated using SH basis images. Castelan and Horebeek [138] and Rara *et al.* [34] cast their models in a regression framework using the Partial Least Squares (PLS) method which uses a few matrix operations for shape reconstruction to provide a computationally efficient alternative to the iterative methods used in [33, 134].

On the other hand, several works have been conducted in the reconstruction of tooth occlusal surface based on 3D surface measurements and a training teeth dataset [148, 149, 160, 161, 162]. For instance, Zheng *et al.* [148, 160] considered teeth anatomical features where the snake model proposed in [163, 164] was improved to automatically capture the features on the tooth surface such as marginal ridges, cusps and groove lines. Constraints were then established based on the corresponding features between the standard tooth and a prepared dataset where Radial Basis Function (RBF) was applied to define a one-to-one point mapping while interpolating the intermediate points.

Buchaillard *et al.* [149] relied on manual alignment between each training tooth and a generic tooth, in order to have a point-to-point correspondence between different specimens. Nonrigid registration using free form deformations and splines was then applied to deform the generic tooth to each training tooth in the set where each resulting example is represented by a vector. A Point Distribution Model (PDM) was then computed using PCA to describe the shape variation of a set of aligned sample models where major deformation modes were defined by the eigenvectors associated with the maximum eigenvalues. To reconstruct the 3D tooth shape, initial rigid registration was conducted using Iterative Closest Point (ICP) [165] while nonrigid registration was then conducted for finding optimum parameter values based on an energy function that minimized the distance between the crowns of the two volumes.

Sporrying and Jensen in [161] proposed a statistical model to reconstruct missing tooth surface using a selection of tooth shapes and including information the position and anatomy of other teeth. They studied shape variations of teeth landmarks using

Bayes PCA where the estimation of the Bayes covariance matrix was conducted with a Gaussian likelihood while an inverted Wishart distribution was used as prior. Their system depended on hand picked landmarks. Moreover, it was evaluated on cast models scanned by a laser scanner.

Alternatively, Blanz and Vetter in [162] proposed a statistical model of a selection tooth shapes by warping each training shape landmark to a template shape, where the mesh of this template is projected onto the shape space before warping landmarks and mesh vertices back. Based on this, a 3D morphable model was introduced in [162] that used 3D meshes as training data which are morphed to fit the input optical images. Bayes estimation was then used to construct the missing tooth, by aligning the patients teeth with the statistical model for the non-missing teeth. Yet they rely on manual annotation to achieve onlay and inlay restoration while they did not handle the recovery of missing crown as their model is a tooth-based model.

9.3 Contributions

This work proposes an image-based approach for reconstruction of the tooth crowns which can serve as a fully automated framework for tooth restoration from a single optical image. The presented approach is based on building statistical models for 3D shape and 2D texture of human's visible jaw (clinical crowns). The process starts with annotating the jaw at known anatomical landmarks, in order to co-register the shapes and textures needed to construct the corresponding models. The key challenge is then the estimation of non-rigid transformation between an input image and the mean jaw (the origin of the object space). Such transformation is estimated using the physically motivated thin-plate splines [146] warping function using a proposed definition of anatomical human jaw landmarks that are automatically detected. The need for hand picked jaw features is then eliminated using an extended version of Active Shape Model (ASM) [166].

Our approach can recover the 3D shape of clinical crowns from a single occlusal image of the human jaw under arbitrary illumination while relaxing the non-realistic assumption of Lambertian reflectance which is commonly assumed, *e.g.* [34, 49]. Prior shape, albedo and appearance (net result of illumination and reflectance) models from real data, which are metric in nature, are incorporated into the shape recovery framework

in order to resolve the concave/convex ambiguity of conventional SFS approaches.

With the frequency-space based representation of the image irradiance equation, an appearance model is proposed which accounts explicitly for teeth reflectance as well as complex illumination conditions. Albedo-modulated irradiance harmonics is used to provide the optimal basis for illumination and reflectance representation¹. Assuming enamel reflectance obeys Wolff-Oren-Nayar reflectance model [167], where surface roughness is physically measured using optical surface profiler, it is proved that it can be represented by at most 3rd order harmonic basis. The recovery framework is a non-iterative approach which incorporates regression-like algorithm in the minimization process.

It is worth mentioning that human face and jaw modeling bear similarities as well as differences. Facial images can be easily obtained and databases of various imaging conditions are already in place, along with a significant body of algorithmic development. Human faces are easy to annotate and automate the process of face cropping and feature extraction. On the other hand, the human jaw is not a friendly environment to image, as indicated before, while no databases exist to carry out a SSFS methodology.

While in one's earlier work [49], a model-based shape reconstruction for the human jaw was proposed, the simple Lambertian model for tooth reflection was assumed whereas the original 3D scans were converted into a Monge patch (2.5D) format which represents the surface as $(x, y, z(x, y))$. This inhibits the reconstruction of labial, buccal and lingual surfaces as well. This chapter shows how to build the shape model based on full 3D information of the jaw molds.

9.4 Anatomical Jaw Landmarks - Proposed Definition

9.4.1 Dental Nomenclature

Human teeth can be classified in two groups [168]; *anterior teeth* which rip food apart and posterior/molar teeth which help chewing food, see Figure 9.1 for illustration. A tooth consists of several layers [169] (see Figure 9.2), however it has been believed that the enamel and dentine layers have a major impact on the visual appearance of teeth [170]; contributions from layers such as pulp may be neglected [171].

¹By modulating irradiance harmonics with surface albedo one manipulates the low-frequency information that can be predicted by the proposed imaging model, especially illumination effects, while preserving the high-frequency texture information of the object at hand




Figure 9.1: Dental nomenclature, common language used in dentistry. Teeth are categorized into two main classes, anterior teeth which rip food apart and posterior/molar teeth which chew the food. All anterior teeth exhibit four types of surfaces and one edge on their crowns (mesial, distal, lingual and buccal), while posterior teeth have five surfaces on their crowns (mesial, distal, lingual, buccal and occlusal). Proximal surface refer to any surface between two teeth, thus it can be either mesial or distal according whether it is towards or away from the midline.




Figure 9.2: Teeth anatomy: different layers of a human model teeth are shown where the enamel and dentine layers are believed to play an important role in characterizing teeth appearance.




Figure 9.3: Tooth surfaces: anterior teeth have four types of surfaces , towards and away from the midline (mesial and distal resp.), towards the lips (labial) and towards the tongue (lingual). While posteriors exhibit an additional surface for chewing (occlusal), besides the inside surface is named buccal.

All anterior teeth have four types of surfaces [172] named: (1) *Mesial* - the surface toward the midline, (2) *Distal* - the surface away from the midline, (3) *Labial* - the outside surface towards the lips and (4) *Lingual* - the inside surface toward the tongue. Posteriors exhibit mesial, distal and lingual surfaces in addition to *Buccal* surface which is outside toward the cheek and *Occlusal* surface which is the chewing surface. See Figure 9.3 for illustration.

9.4.2 Landmarks Definition

This work mainly focuses on the reconstruction of the *clinical crowns* which are defined to be the portion of the teeth that is visible in the human mouth. As such, the jaw's anatomical landmarks are limited to such a space according to their location, i.e. on the tooth surface or on the interface between the tooth and the gum. Typically a landmark represents a distinguishable point which is present in most of the images under consideration, for example, the location of central grooves of each tooth. Figure 9.4 illustrates the location of 62 landmark points for a jaw containing twelve-teeth.




Figure 9.4: Illustration of the proposed human jaw anatomical landmarks. The upper left corner shows a sample a detected landmarks on a real image using Active Shape Model [166].

9.4.2.1 Teeth Landmarks

In case of posterior teeth (i.e. cuspids, premolars and molars) which are responsible for chewing food, the coalescence of the crown lobes are of interest. In particular, a *central pit or groove* can be considered as a landmark which is the deepest portion of a tooth fossa. While anterior teeth (i.e. incisors) whose job is to rip food apart is identifiable by a convex elevation of the crown surface which forms the biting edge. Hence the midpoint of the *incisal edge or ridge* is considered as a landmark for an anterior tooth.

9.4.2.2 Teeth-Gum Landmarks

The fibrous tissue covering the alveolar bone and surrounds the necks of the teeth, i.e. the gum, forms what is termed as *gingival line*. This line marks the level of termination of the non-attached soft tissue surrounding the tooth. It separates the clinical crown and the root. The *gingival line midpoint* is defined to be the minima (for lower jaws) or maxima (for upper jaws) point on the gingival line formed by a single tooth. On the other hand, *gingival embrasure* is the respective point in the open space between the proximal surfaces of two adjacent teeth in the same dental arch.

9.4.3 Landmark Localization in Optical Images

In the online stage of the proposed approach, a single image of the visible crowns is given from which the defined landmarks should be identified. This guides the alignment of the input image to the prior model, e.g. [162]. Hence, it is essential to automate the detection of such landmarks. In the training set, one manually annotates an ensemble of human jaws surfaces (based on CT-scanning of molds) in order to construct a sparse version of the jaw shape. These landmarks serve as a correspondence operator between individual training samples where the generalized Procrustes analysis [173] is used to filter out translation, scale and rotation. An extended version of ASM [166] is deployed to search for the landmarks in the given image. The ASM repeats the following two steps until convergence: (i) suggest a tentative shape by adjusting the locations of shape points by template matching of the image texture around each point (ii) conform the tentative shape to a global shape model. The individual template matches are unreliable and the shape model pools the results of the weak template matchers to form a stronger overall classifier. The entire search is repeated at each level in an image pyramid, from coarse to fine resolution. The initialization of the mean shape onto the given image is accomplished by segmenting the teeth region based on fitting a Gaussian mixture to the image intensity with two dominant classes; jaw and background.

9.5 Beyond Lambertian Reflection

Accurately modeling the appearance of objects in general and human teeth in particular goes far beyond the Lambertian reflectance model. Human teeth are neither Lambertian nor convex, introducing cast shadows in concave regions. As such, this inspired oneself to move closer to the realistic situation with a spatially varying non-Lambertian reflectance model.

9.5.1 Microscopic Surface Model

Surface reflection depends on, among other factors, the microscopic surface characteristics [174]. Rough surfaces tend to scatter incident light in various directions, favoring some directions than others, while smooth surfaces may reflect incident light in a single direction. Modeling the microscopic surface irregularities aids the construction of realistic/physical reflectance models. A surface can be modeled as a random process, defined

by a statistical distribution of some robust feature derived from the surface profile. The following discusses the physical measurement of human teeth surface roughness, then one of the popular analytic reflectance models is reviewed which take surface roughness into consideration.

Roughness is a measure of the topographic relief of a surface [175]. It can be viewed as a texture measure of a given surface. It can be quantified as the vertical deviation (*i.e.* height variation) of that surface from the mean surface level along the surface spatial wavelength. A relative smooth surface has typically peak-to-valley roughness of $\approx 0.1\mu\text{m}$ [175].

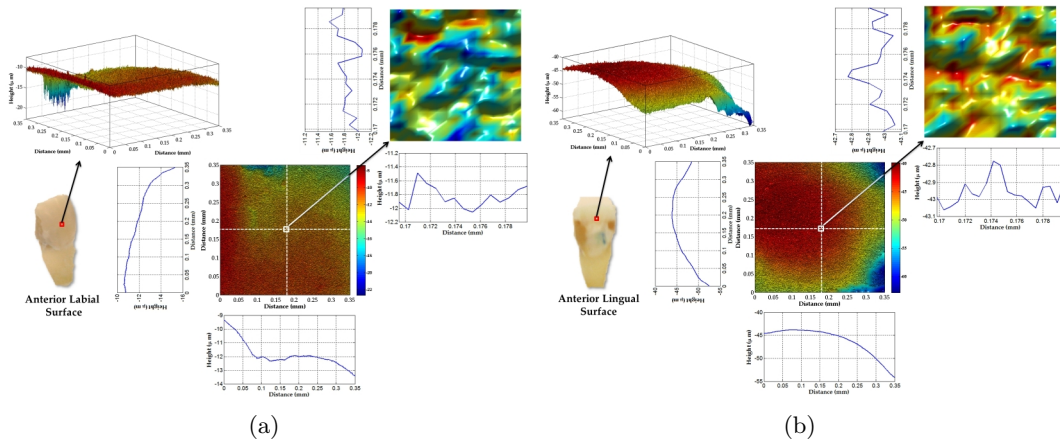


Figure 9.5: Surface height variations of anterior (a) labial and (b) lingual surfaces. Average horizontal and vertical surface profiles are shown on the measured area (0.35mm^2). A zoom-in view on an area 0.01mm^2 is also shown along with its surface profiles.

Surface profile (*i.e.* height variations) measurements, on the *microscopic* scale, can be used to measure the roughness using either one of two probes; light beam for optical/non-contact profilers or mechanical stylus (using diamond stylus) for contact profilers. In the latter case, height calibration is required and deconvolution is also needed (yet not successful) since the measured profile is a convolution of the true surface profile and the stylus radius [175]. On the other hand, neither calibration nor deconvolution are needed in case of optical profilers, whereas the requirement of flat surfaces with non-specular reflection become the major issue of measuring height variations.

The University of Louisville has a 3D optical surface profiler, NewView 700s from Zygo company which is based on Scanning White-Light Interferometry technology. It offers fast, non-contact, high-accuracy, 3D metrology of surface features and include proprietary MetroPro software for the ultimate in advanced data analysis. The field of

view ranges from 0.35mm to 3.5mm .

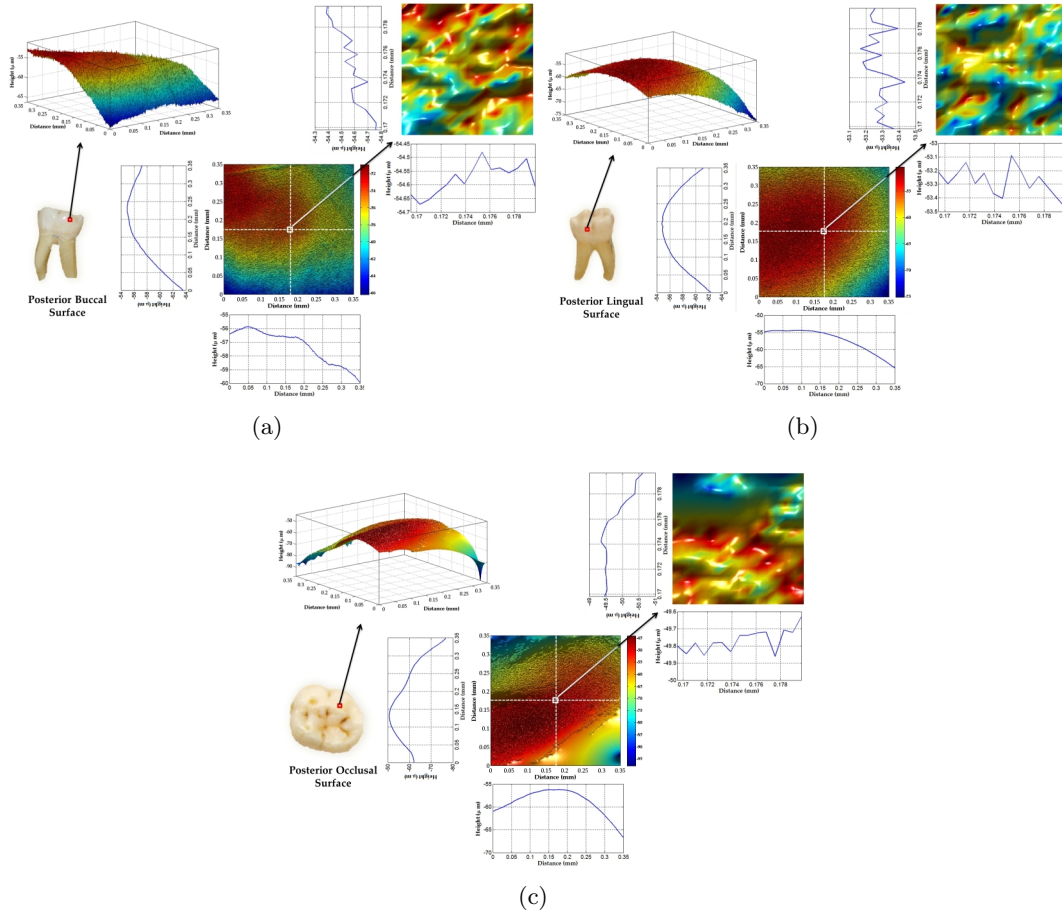


Figure 9.6: Surface height variations of posterior (a) buccal (b) lingual and (c) occlusal surfaces. Average horizontal and vertical surface profiles are shown on the measured area (0.35mm^2). A zoom-in view on an area 0.01mm^2 is also shown along with its surface profiles.

We measured height variations of different visual tooth surfaces (*i.e.* mesial and distal surfaces are ignored) from an area of 0.35mm^2 using 10X optical zoom. Figure 9.5 and Figure 9.6 show the measured height variations of a sample central incisor and 3rd molar. It can be noticed that on a tiny area (0.35mm^2), the tooth surface is not planar causing the deviation from a macroscopically flat surface profile. A zoom-in view on an area of 0.01mm^2 shows local variations of surface height with respect to surface point having close to normal incidence. The average surface profiles provide a physical validation that the appearance of a tooth surface can be modeled using the microfacet-theory [13, 70] which assumes that the surface consists of a large number of small flat facets. Microfacet reflectance models tend to be intuitive with tractable analytical expressions [94].

Vis-à-vis surface reflection, the slope distribution model is the most popular rough-

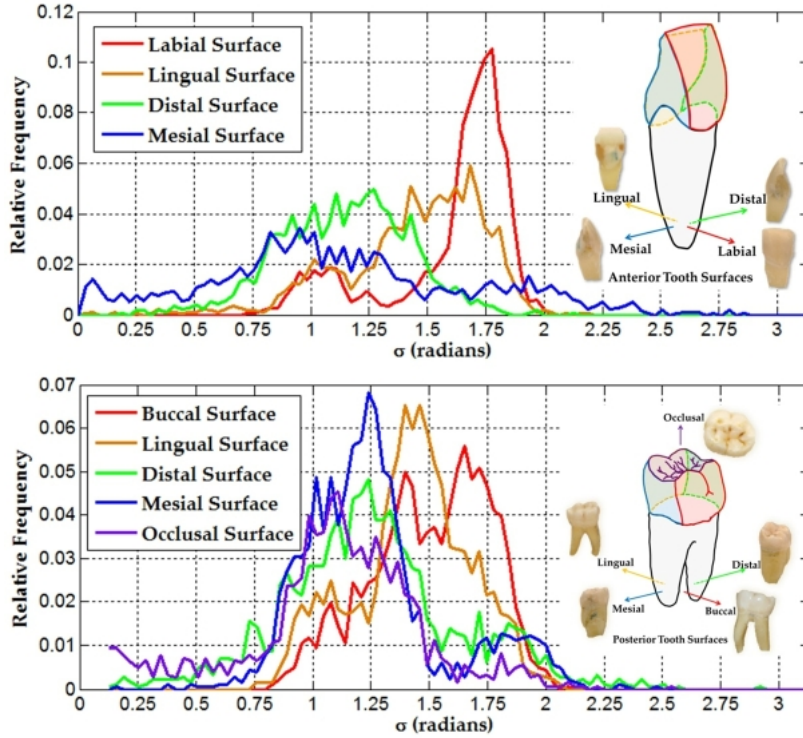


Figure 9.7: The roughness parameter is estimated based on the measurement of microscopic height variation of a $0.35mm^2$ surface patches of different surface types for incisor and molar teeth. The intra-oral camera pixel size was estimated to cover approximately $0.0075mm^2$. The measured patch is divided into smaller patches with pixel size where the roughness parameter σ is computed. According to the distribution, the parameter tends to lie between 0.7 to 2 radians regardless the tooth surface type.

ness model [14], where the surface is considered as a collection of planar facets whose normal vectors deviate from the surface geometric normal (*i.e.* mean surface orientation). The distribution of the deviation angle is assumed to follow a Normal distribution with zero mean and standard deviation governed by the surface roughness. Figure 9.7 shows the distribution of the computed roughness parameter for different teeth surface types using the microscopic surface height variations, where the deviation angle of the surface microfacets compared to the surface geometric normal tends to lie between 0.7 to 2 radians regardless the tooth surface type.

9.5.2 Towards a Realistic Model for Human Teeth Reflectance

Light interaction with tooth surface is a complex process which involves Fresnel theory¹ and subsurface scattering controlled by biological parameters of the enamel and dentin layers. While computer graphics community has developed comprehensive models to

¹Teeth are usually wet, giving rise to Fresnel reflection due to different refractive indices of the saliva and the tooth material.

account for translucent materials such as human teeth for photo-realistic image synthesis, such models require computationally expensive rendering techniques. Thus computer vision work has opted for discount subsurface scattering where the notion of surface reflectance can be analytically modeled by surface BRDF.

Despite the simplicity of Lambert’s model, it has been proven to be an inaccurate approximation to the diffuse component of the surface reflectance [176]. Through a set of experiments carried out on real samples, such as plaster, clay, sand and cloth, Oren and Nayar [176, 177, 178] showed that all these surfaces demonstrate significant deviation from Lambertian behavior. These results motivated them to develop a comprehensive reflectance model for rough diffuse surfaces. They used the roughness model proposed by Torrance and Sparrow [12] which assumes that the surface is composed of a collection of long symmetric V-cavities. Each V-cavity has two opposing facets. The roughness of the surface is specified using probability function for the orientations of the facets. In order to use geometric optics, the area of each facet is assumed to be much larger than the wavelength of the incident light beam, and at the same time much smaller than the area of the surface patch being projected onto one pixel, thus the facets covered by one pixel can be described by statistical distributions.

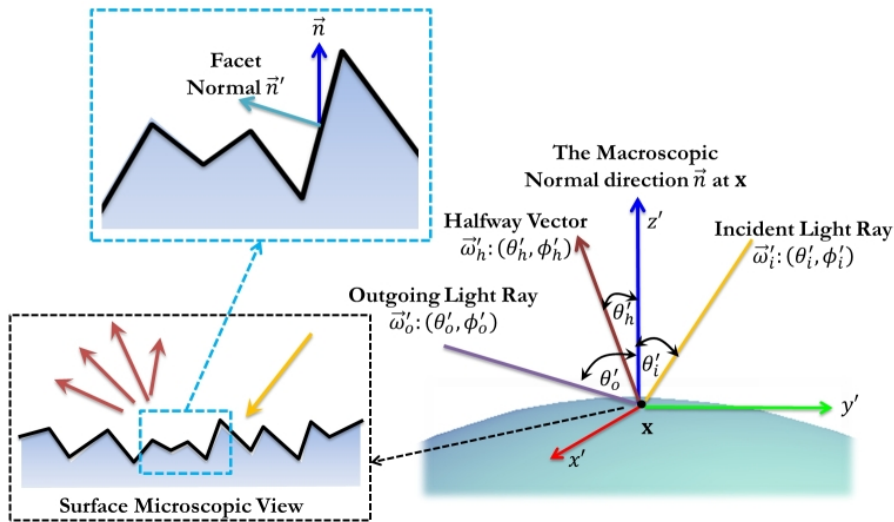


Figure 9.8: Definitions of angles and vectors used in microfacet models. Vector $\vec{\omega}'_h$ is the bisector between the light direction and the viewer direction. Note that rough surfaces tend to scatter incident light as compared to smooth surfaces.

The V-cavity geometry implies that only facets facing in direction of the halfway vector $\vec{\omega}'_h = (\theta'_h, \phi'_h)$ affect the BRDF (see Figure 9.8) where masking and shadowing of micro-facets are included in the geometric attenuation factor $G(\vec{\omega}'_i, \vec{\omega}'_o, \vec{\omega}'_h)$. Oren and

Nayar [14] conducted analysis of roughness for diffuse surfaces, where the facets covered by one pixel described by the statistical distribution of the slope area. The distribution function $D_\sigma(\vec{\omega}'_h)$ describes the fraction of the facets oriented in the direction of $\vec{\omega}'_h$. It is described by a zero mean Gaussian with variance σ^2 representing surface roughness, determining the fraction of surface area which the facets of the same orientation covers. Thus the Oren-Nayar BRDF is given by,

$$f_r^{ON}(\vec{\omega}'_i, \vec{\omega}'_o) = \frac{D_\sigma(\vec{\omega}'_h)G(\vec{\omega}'_i, \vec{\omega}'_o, \vec{\omega}'_h)}{4 \cos \theta'_i} \quad (9.1)$$

The geometric attenuation factor G ranges from 0 (total shadowing) to 1 (no shadowing), defining how the incoming light beam interact with the surface locally; the entire beam can reflect, a portion of the outgoing beam can be blocked (masking), or a portion of the incoming beam can be blocked (shadowing). Thus the geometric attenuation factor can be defined as,

$$G(\vec{\omega}'_i, \vec{\omega}'_o, \vec{\omega}'_h) = \min \left(1, \frac{2 \cos \theta'_h \cos \theta'_o}{\vec{\omega}'_o \cdot \vec{\omega}'_h}, \frac{2 \cos \theta'_h \cos \theta'_i}{\vec{\omega}'_o \cdot \vec{\omega}'_h} \right) \quad (9.2)$$

A simplified version presented by Oren and Nayer [14] is usually used which is defined as,

$$f_r^{ON}(\vec{\omega}'_i, \vec{\omega}'_o) = \frac{\rho_d}{\pi} \left(A + B \max[0, \cos(\phi'_o - \phi'_i)] \sin \check{\alpha} \tan \check{\beta} \right) \quad (9.3)$$

where,

$$A = 1 - 0.5 \frac{\sigma^2}{\sigma^2 + 0.33}, \quad B = 0.45 \frac{\sigma^2}{\sigma^2 + 0.09} \quad (9.4)$$

and $\check{\alpha} = \max[\theta'_o, \theta'_i]$, $\check{\beta} = \min[\theta'_o, \theta'_i]$ and ρ_d is the diffuse albedo.

While Oren-Nayar model [14] modulates the Lambert's cosine law by a term which depends on the squared sine of the incidence angle, resulting in apparent brightening at surface patches which move away from the light source; assumed to be in self shadow. Wolff [167], on the other hand, has developed a reflection model for smooth surfaces which are modeled as a collection of scatterers contained in a uniform medium with index of refraction different from that of air. Wolff's model uses an angle dependent Fresnel term to account for the refractive attenuation of incident light at the surface-air boundary [179]. This Fresnel term modifies in a multiplicative way the Lambertian cosine model. The effect is to depress the surface radiance for near-normal incidence.

Wolff's model assumes the subsurface inhomogeneities isotropically scatter light where the Wolff's BRDF for a dielectric medium with index of refraction η is given by,

$$f_r^W(\vec{\omega}'_i, \vec{\omega}'_o) = \rho_d (1 - F(\theta'_i; \eta)) \times \left\{ 1 - F \left(\sin^{-1} \left(\frac{\sin \theta'_o}{\eta} \right); 1/\eta \right) \right\} \quad (9.5)$$

An approximation for Fresnel formula for dielectrics used by Cook and Torrance [13] can be written as,

$$F(c; n) = \frac{1}{2} \frac{(g - c)^2}{(g + c)^2} \left(1 + \frac{(c(g + c) - 1)^2}{(c(g - c) + 1)^2} \right) \quad (9.6)$$

with $g^2 = n^2 + c^2 - 1$.

Wolff *et al.* in [180] has a physically deeper model for diffuse reflectance from shiny but slightly rough surfaces where they combined the Oren-Nayar model for rough surfaces [14] and the Wolff model for smooth surfaces [167]. This combined model (Wolff-Oren-Nayar) makes the assumption that each V-groove microfacet reflects light according to the Wolff model. As such, it is defined as,

$$f_r^{WON}(\vec{\omega}'_i, \vec{\omega}'_o) = \frac{\rho_d}{\pi} \left(\tilde{A} + B \max[0, \cos(\phi'_o - \phi'_i)] \sin \check{\alpha} \tan \check{\beta} \right) \quad (9.7)$$

where,

$$\tilde{A} = A (1 - F(\theta'_i; \eta)) \times \left\{ 1 - F \left(\sin^{-1} \left(\frac{\sin \theta'_o}{\eta} \right); 1/\eta \right) \right\} \quad (9.8)$$

Since the azimuthal part in Equation 9.7 depends on the azimuthal angles difference, Wolff-Oren-Nayar BRDF can be expanded in the subspace spanned by the proposed *isotropic* Helmholtz HSH-based basis as,

$$f_r^{WON}(\vec{\omega}'_i, \vec{\omega}'_o) = \sum_{p=0}^{\infty} \sum_{r=0}^{\infty} \sum_{q=0}^r a_{pr}^q \mathcal{H}_{pr}^q(\vec{\omega}'_i, \vec{\omega}'_o) \quad (9.9)$$

Figure 9.9 compares the average approximation accuracy of Wolff-Oren-Nayar reflection model under distant illumination using the proposed isotropic basis in comparison to bases of Westin *et al.* [19], Gautron *et al.* [20], Koenderink *et al.* [21], Habel and Wimmer [22] and the Helmholtz basis of Koenderink *et al.* [21] where their isotropic version is used. The average is taken over 3750 BRDF samples taken by uniformly sampling the roughness (according to Figure 9.7) and the enamel's refractive index domain¹, where

¹As noted by Wang *et al.* [181], the refractive index of the tooth enamel layer lies in the range

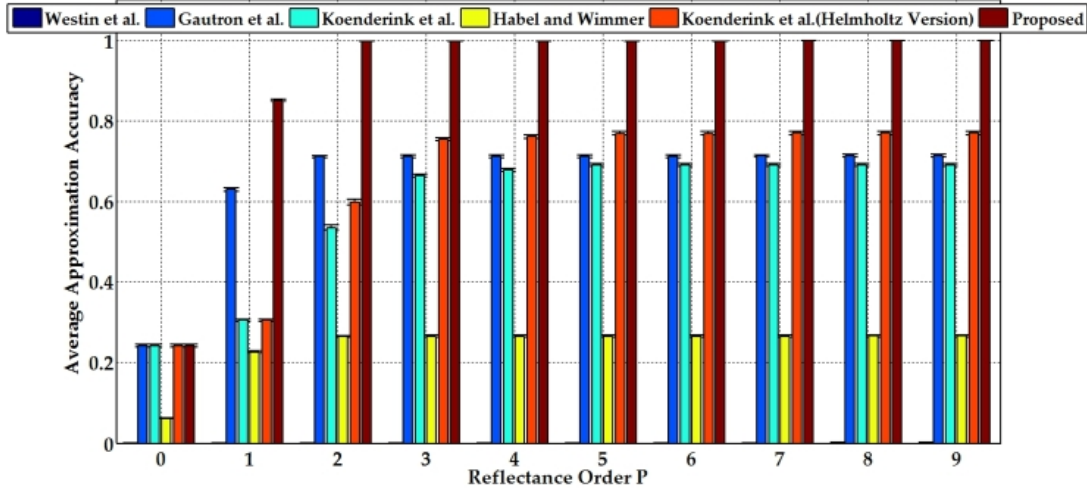


Figure 9.9: The average approximation accuracy of Wolff-Oren-Nayar reflection model under distant illumination as a function of the truncating reflectance order P . The roughness (according to Figure 9.7) and the enamel’s refractive index (1.62 ± 0.02 [181]) domains are uniformly sampled. The average is taken over 3750 BRDF samples where spectrum is obtained by projecting f_r^{WON} using Monte Carlo integration onto the subspace spanned by the proposed isotropic basis in comparison to bases of Westin *et al.* [19], Gautron *et al.* [20], Koenderink *et al.* [21], Habel and Wimmer [22] and the Helmholtz basis of Koenderink *et al.* [21] where their isotropic version is used. Note that the proposed basis provide higher approximation accuracy at lower reflectance orders compared to other bases, hence the proposed representation is capable of providing a compact representation of reflectance functions.

the BRDF spectrum coefficients can be obtained using Monte Carlo integration. One can observe that the proposed basis shows higher approximation accuracy especially at lower truncating reflectance orders compared to other bases. As such, the presented basis provide a compact BRDF representation requiring fewer coefficients to accurately model surface BRDFs.

9.6 Jaw Statistical Model Construction

The proposed model-based shape recovery involves the construction of three models; namely the shape, albedo (also referred to as texture) and appearance (net result of illumination and reflectance) models. While the first two models are constructed in an offline stage, the appearance model is constructed at runtime when an input image is presented to the shape recovery framework. This section presents how the offline jaw models can be constructed. The process of constructing these models primarily depends on two major steps; (1) establishing dense correspondence between different samples in

1.62 ± 0.02.

the training ensemble, and (2) statistical modeling of shape/albedo variations in the given ensemble.

9.6.1 Shape Model

The jaw's shape model is constructed from a training data ensemble of 3D triangular meshes where each mesh is obtained from a high resolution computer tomography (CT) scan of human jaw molds where the Expectation-Maximization (EM) algorithm [124] is used for segmentation. In order to obtain a compact shape representation, Valette *et al.* approach in [182] is used to re-mesh the extracted iso-surfaces from the segmented CT volumetric data where more mesh triangles are maintained where prevailing features exist on the teeth surface, see Figure 9.10.

Hereafter, each 3D jaw surface is represented as a 2-manifold triangular mesh $\mathcal{G} = (\mathcal{V}, \mathcal{F})$, where $\mathcal{V} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_V\}$ is a set of V -vertices with $\mathbf{x}_v \in \mathbb{R}^3$ and $\mathcal{F} = \{t_1, t_2, \dots, t_F\}$ is a set of F -triangular faces, with the f -th face $t_f = \{\mathbf{x}_{f0}, \mathbf{x}_{f1}, \mathbf{x}_{f2}\}$ is constructed from three vertices with indices $f0, f1$ and $f2$.

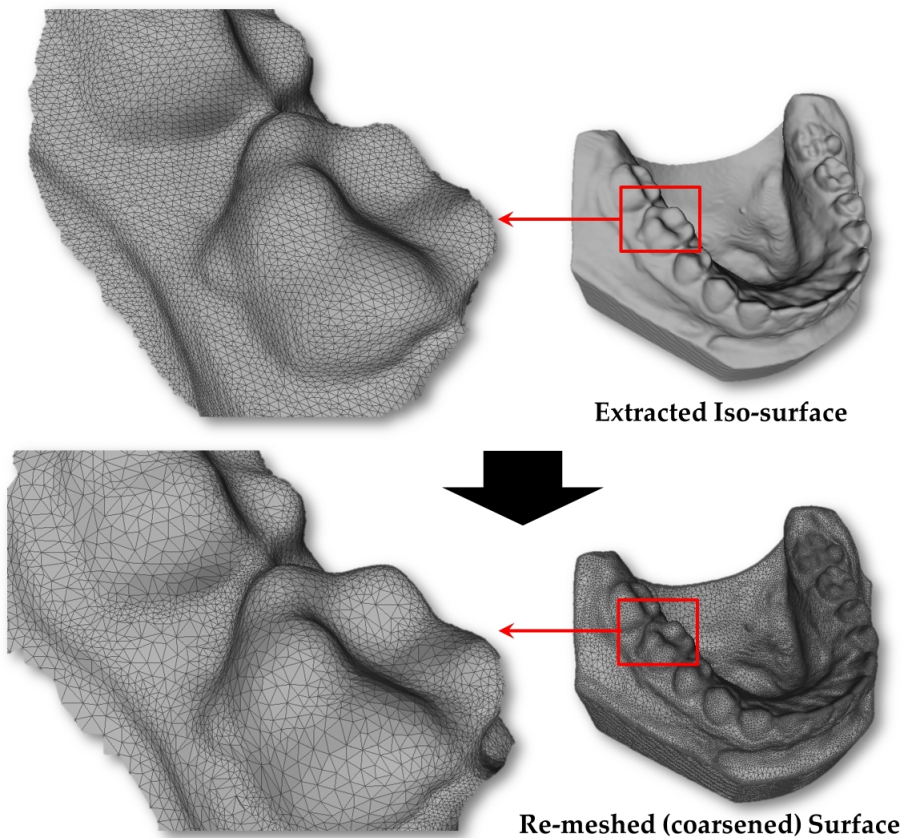


Figure 9.10: A sample of extracted iso-surface and a re-meshed (coarsened) version using [182] where 25% of mesh vertices are retained.

We follow the work by Patel et al. [143] in obtaining dense correspondence between different jaw surfaces where a finite set of sparse landmark points, (x'_v, y'_v, z'_v) , are manually annotated for all the database samples. Note that this is done once in an offline stage¹. There is a total of 72 manually annotated landmark points on a 14 teeth jaw. Generalized Procrustes Analysis (GPA)[146] is then performed to provide an initial rigid alignment of the dense shapes to a common reference frame where the alignment procedure is guided by the sparse set of anatomical jaw landmarks (Figure 9.4).

To obtain point-to-point dense correspondence between two rigidly aligned jaw shapes \mathcal{V}_1 and \mathcal{V}_2 , a warping function, $\mathbf{f}: \mathbb{R}^3 \rightarrow \mathbb{R}^3$, based on physically motivated thin-plate splines [146, 183] is constructed using the landmark points (e.g., $(x'_v{}^1, y'_v{}^1, z'_v{}^1) \in \mathcal{V}_1$ and $(x'_v{}^2, y'_v{}^2, z'_v{}^2) \in \mathcal{V}_2$) as control points, i.e., $\mathbf{f}(\mathbf{x}'_v{}^1) = \mathbf{x}'_v{}^2$, where $\mathbf{x}'_v{}^1 = [x'_v{}^1, y'_v{}^1, z'_v{}^1]^T$ and $\mathbf{x}'_v{}^2 = [x'_v{}^2, y'_v{}^2, z'_v{}^2]^T$. Once the warping function is solved, this warp is applied to all vertices in \mathcal{V}_1 from the first shape, determining their corresponding locations on the second shape \mathcal{V}_2 . Note that the warping operation is performed in 3D space.

In order to obtain a dense correspondence between all shapes in the database, the 3D thin-plate spline is applied in an iterative manner as follows, where according to the experimentation, this algorithm converges in a few iterations.

-
- 1 Choose one sample as an initial estimate of the mean shape - one may use the first shape in the ensemble.
 - 2 Solve for the warping function \mathbf{f} between the current mean shape and all other samples in the database.
 - 3 Re-calculate the estimate of the mean from the aligned shapes.
 - 4 If there is a significant change in the mean, return to step 2.
-

The shape information for each aligned shape \mathcal{V}_k can be stacked into a vector, \mathbf{S}_k where for a V -vertices shape one has,

$$\mathbf{S}_k = [x_1^k, y_1^k, z_1^k, \dots, x_V^k, y_V^k, z_V^k]^T \quad (9.10)$$

Principal Component Analysis (PCA) [184, 185] is then performed on the set of shape vectors, \mathbf{S}_k . The average ($\bar{\mathbf{s}} = \frac{1}{K} \sum_{k=1}^K \mathbf{S}_k$) is first subtracted from each shape vector, $\mathbf{d}_k = \mathbf{S}_k - \bar{\mathbf{s}}$, and form the data matrix $\mathbf{D} = (\mathbf{d}_1, \dots, \mathbf{d}_K)$. PCA is performed

¹A possible direction to pursuit is to automatically localize these landmarks in a given triangular mesh.

by computing the the eigenvectors $(\mathbf{s}_1, \mathbf{s}_2, \dots)$ of the covariance matrix $(\mathbf{C} = \frac{1}{K}\mathbf{D}\mathbf{D}^T)$, which can be achieved using Singular Value Decomposition (SVD) [186] of \mathbf{D} . The eigenvalues $(\sigma_{s,k}^2)$ of \mathbf{C} are related to the variance of the data across each eigenvector direction. The resulting shape model can be written as,

$$\mathbf{s} = \bar{\mathbf{s}} + \sum_{k=1}^{K-1} \lambda_k^s \cdot \mathbf{s}_k = \bar{\mathbf{s}} + \mathbf{P}_s \boldsymbol{\lambda}_s \quad (9.11)$$

where $\mathbf{P}_s = [\mathbf{s}_1, \mathbf{s}_2, \dots]$ are the shape eigenvectors and $\boldsymbol{\lambda}_s$ is the set of shape coefficients. Note that the subscript s refer to shape.

9.6.2 Albedo Model

Given human jaw molds obtained from a high resolution CT scans and intra-oral images of the visible crown (occlusal) surface, it is important to address the issue of incorporating texture information into the prior information. This involves mapping such texture from the image domain to the triangular mesh domain in order to associate each mesh vertex \mathbf{x}_v with albedo in the red-green-blue channels, $\rho_r(\mathbf{x}_v)$, $\rho_g(\mathbf{x}_v)$ and $\rho_b(\mathbf{x}_v)$.

Due to the unfriendly image acquisition scenario for jaw occlusal images, the resulting images can not be considered as pure albedo where there are shadows due to non-convex jaw regions in addition to non-uniform distribution of illumination. As such, it is needed to factor out, as much as possible, the effect of shape and illumination from the given image. This can be accomplished by extracting the reflectance information (albedo) from the given texture where the recently proposed intrinsic image decomposition proposed by Barron and Malik¹ in [187] is deployed. Their approach decomposes a single image into its constituent images: shape, reflectance and illumination where one is only interested in the resulting reflectance image. Figure 9.11 shows a sample of an input occlusal image and the estimated reflectance according to [187]. Note that the albedo is only solved for using the pixels which lie within the convex hull of the jaw landmarks defined on the occlusal image.

In order to provide a dense correspondence between an occlusal image and its corresponding jaw surface, the sparse jaw landmarks which are located in the image plane are assumed to be on the xy-plane in the 3D space where 3D thin-plate spline is used to provide a warping function between image pixels and surface points using image landmarks

¹We would like to acknowledge the authors for providing their software and help to apply their approach on jaw images.




Figure 9.11: A sample of estimated reflectance of an input occlusal image using [187]. Note the effect of shape and illumination on the resulting image (right) is minimal compared to the input image on the left.

and surface landmarks as control points. Due to the projective nature of the occlusal image, the same image pixel can be mapped to an occlusal surface point as well as the mold base surface. As such, according to cervical landmark points, a least-squares plane is fit to remove such ambiguity, see Figure 9.12.

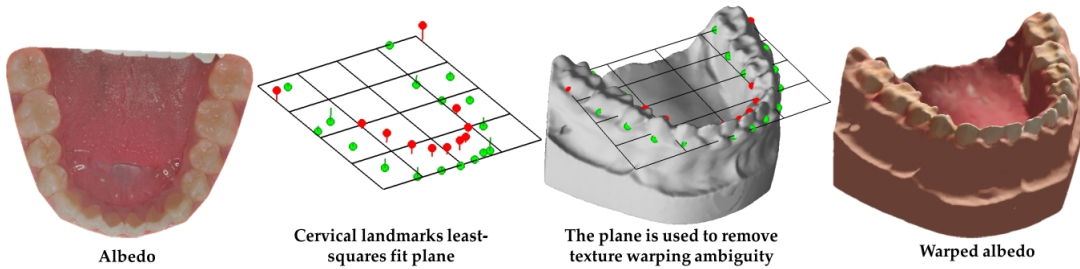


Figure 9.12: A sample of albedo being warped to a jaw surface while the cervical landmarks least-squares fit plane is used to resolve the ambiguity of having a pixel being mapped to two surface points (one above and the other below the plane) where mapping in such cases is restricted to the points above the plane. Red cervical landmarks lie above the cervical plane while the green ones lie below the plane. Note that vertices having no texture mapped to are set of average albedo value, this is apparent in vertices belonging to the mold base. Whereas for vertices belonging to the mouth floor (or ceiling), the vertices which are closest to the cervical plane are chosen.

The albedo information for each data sample can be stacked into a vector¹, \mathbf{A}_k , i.e.,

$$\mathbf{A}_k = \left[\rho_r(\mathbf{x}_1^k), \rho_g(\mathbf{x}_1^k), \rho_b(\mathbf{x}_1^k), \dots, \rho_r(\mathbf{x}_V^k), \rho_g(\mathbf{x}_V^k), \rho_b(\mathbf{x}_V^k) \right]^T \quad (9.12)$$

Exactly the same procedure is done to obtain the texture eigenvectors (\mathbf{a}_k) and variances

¹We used the Lab color space instead of the RGB one since the latter suffers from strong correlation among its color channels as well as non-linearity, *i.e.* a small change in one channel does not necessarily mean another shade of the same perceived color.

($\sigma_{a,k}^2$). The resulting albedo model can be written as,

$$\mathbf{a} = \bar{\mathbf{a}} + \sum_{k=1}^{K-1} \lambda_k^a \cdot \mathbf{a}_k = \bar{\mathbf{a}} + \mathbf{P}_a \boldsymbol{\lambda}_a \quad (9.13)$$

where $\mathbf{P}_a = [\mathbf{a}_1, \mathbf{a}_2, \dots]$ are the albedo eigenvectors and $\boldsymbol{\lambda}_a$ is the set of albedo coefficients. Note that the subscript a refer to albedo.

9.7 Model-based Appearance Subspace Construction

Appearance subspaces can be constructed using phenomenological models which capture appearance variations through mathematical abstraction of the reflection process. In particular, the harmonic expansion of the image irradiance equation [24] can be used to derive an analytic subspace to represent images under fixed pose but different illumination conditions [6, 25, 44]; it has been shown that surface reflectance acts as a low-pass filter on the incident illumination [4, 5] where the image irradiance equation has been formulated in a convolution framework [6, 7]. This yielded a band-limited reflected light while accounting for complex illumination, thus a point light source assumption can be relaxed.

For a specific object geometry under fixed pose, the convolution theory implies a multiplicative framework in the frequency domain where an image is represented as a linear combination of pre-computed basis functions, which is termed as *irradiance harmonics*, $\{\mathcal{B}_s\}$, that are pose and geometry dependent. This leads to a finite-dimensional linear model to represent the image irradiance under arbitrary illumination. Representing the illumination by its spherical harmonics (SH) coefficients l_n^m as in [7, 44] and the surface reflectance by its coefficients a_{pr}^q in the proposed Helmholtz HSH-based basis, the image irradiance can be defined as,

$$E(\alpha, \beta) = \sum_s c_s \mathcal{B}_s(\alpha, \beta) \quad (9.14)$$

where (α, β) are the spherical coordinates of the surface normal $\vec{n}(\mathbf{x})$ in the global reference frame, $c_s = l_n^m a_{pr}^q$ with s and its corresponding indices n, m, p, r and q are given by an ordering function of the basis functions based on their average power content. The irradiance harmonics $\{\mathcal{B}_s\}$ encode the orthogonality relation between SH and the

deployed reflectance representation. They are defined as,

$$\mathcal{B}_{npr}^{mq}(\alpha, \beta) = \mathcal{D}_{mq}^n(\alpha, \beta) \int_{\Omega'_i} Y_n^q(\vec{\omega}'_i) \mathcal{H}_{pr}^q(\theta'_i, \alpha, |\phi'_i - \pi|) \cos \theta'_i d\vec{\omega}'_i \quad (9.15)$$

where $\{Y_n^q\}$ are the real SH bases, $\{\mathcal{H}_{pr}^q\}$ are the Helmholtz HSH-based reflectance bases and \mathcal{D}^n are Wigner's rotation matrices [59].

This work takes advantage of the two-fold benefit of the frequency-space representation of the image irradiance equation. First, it decouples the image formation process such that the illumination conditions and surface reflectance characteristics are encoded into the coefficients of the irradiance harmonics. This allows the incorporation of prior information about natural illumination and real world surface materials in the subspace construction process. Second and more importantly, the process of finding such subspace can be cast as establishing a relation between its principal components and that of the irradiance harmonics. This resolves the issue of dimensionality since the source of randomness in the imaging process becomes the irradiance harmonics coefficients rather than the whole image realization.

The main idea behind constructing a model-based analytic appearance subspace is to use the frequency representation of image irradiance directly in an unsupervised subspace learning scheme seeking orthonormal principal components spanning the subspace where shading cues lie. Let a D -pixel image irradiance be represented in the vector space as $\mathbf{e} \in \mathbb{R}^D$. The objective is to define an orthonormal projection matrix $\tilde{\mathbf{W}} \in \mathbb{R}^{D \times D'}$ which maps the image space to a lower-dimensional subspace $\mathbb{R}^{D'}$, with $D' \leq D$, which captures most of the variations due to illumination and reflectance¹. Thus the projection matrix $\tilde{\mathbf{W}}$ should maximize the total vector scatter such that²,

$$\tilde{\mathbf{W}} = \underset{\tilde{\mathbf{W}}}{\operatorname{argmax}} \Psi_{\mathbf{y}} \quad \text{s. t.} \quad \Psi_{\mathbf{y}} = E\{\|\mathbf{y} - \bar{\mathbf{y}}\|_F^2\} \quad \text{where} \quad \mathbf{y} = \tilde{\mathbf{W}}^T \mathbf{e} \quad (9.16)$$

Let the s -th irradiance harmonics be represented in the vector space as $\mathbf{b}_s \in \mathbb{R}^D$ which can be written as $\mathbf{b}_s = \tilde{\mathbf{W}}_B \mathbf{c}_s^B$ with $\tilde{\mathbf{W}}_B \in \mathbb{R}^{D \times D'}$ and $\mathbf{c}_s^B \in \mathbb{R}^{D'}$. The optimal $\tilde{\mathbf{W}}_B$ is determined from the D' -eigenvectors of $\mathbf{B} = [\mathbf{b}_1 \dots \mathbf{b}_S]$ corresponding to its largest D' -eigenvalues. Relating the principal components of the image space to that of the irradiance harmonics in the vector space results in $\tilde{\mathbf{W}}^T = \mathbf{A}_W \tilde{\mathbf{W}}_B^T$ where $\mathbf{A}_W \in$

¹ denotes lower-dimensional.

²Note that D here connotes the number of visible surface points to the viewer, *i.e.* $D \leq V$ where V is the total number of vertices for a given 3D object.

$\mathbb{R}^{D' \times D'}$. The solution of this matrix is given by the D' -eigenvector of Υ^T such that,

$$\Upsilon = \sum_s \sum_{s'} E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} \mathbf{c}_s^B (\mathbf{c}_{s'}^B)^T \quad (9.17)$$

where $E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} = E\{c_s c_{s'}\} - E\{c_s\}E\{c_{s'}\}$. As such, given the shape and albedo of a data sample, one can construct its appearance subspace while incorporating natural illumination (*e.g.* [55, 105]) and teeth reflectance properties. The proposed approach is summarized in Algorithm 7. Steps are also illustrated in Figure 9.13.

9.8 Model-based Shape and Albedo Recovery

Given a jaw occlusal image, its sparse landmarks are used to infer a warping function between image pixels and the vertices of the mean jaw shape $\bar{\mathbf{s}}$ using 3D thin-plate spline where the landmarks are assumed to be located on the xy-plane in the 3D space. The cervical least-squares plane is then used to remove mapping ambiguity in a similar manner as in constructing the albedo model.

9.8.1 SFS Brightness Constraint

When the light source and the viewer are far from the object, the image irradiance E from surface point \mathbf{x} can be defined as the surface radiance being modulated by the surface albedo $\rho(\mathbf{x})$, i.e. $E(\mathbf{x}) = \rho(\mathbf{x})\mathcal{R}(\vec{n}(\mathbf{x}))$. The classical brightness constraint in SFS measures the total brightness of the reconstructed image irradiance compared to the input irradiance, it can be defined as;

$$\epsilon = \int \int (E(\mathbf{x}) - \rho(\mathbf{x})\mathcal{R}(\vec{n}(\mathbf{x})))^2 d\mathbf{x} \quad (9.18)$$

where $\rho(\cdot)$ is the surface albedo at point \mathbf{x} while $\mathcal{R}(\cdot)$ is the radiance of the surface patch with unit normal $\vec{n}(\mathbf{x}) = \text{cart}(\alpha, \beta)$, also known as *surface reflectance function* [6].

The brightness constraint in Equation 9.18 can be rewritten in the discrete domain as a linear combination of pre-computed basis resulted from the harmonic expansion of the reflectance function. Thus the image intensity E can be expressed as;

$$E(\mathbf{x}) = \sum_{s=1}^S y_s \mathcal{W}_s(\vec{n}(\mathbf{x})) \quad (9.19)$$

Algorithm 7 Model-based Appearance Subspace Construction of Human Visible Crowns

INPUT:

- (a) Surface normals of the k -th jaw sample: $\{\vec{n}(\mathbf{x}_v^k) = (\alpha_v^k, \beta_v^k)\}$ where $\mathbf{x}_v^k = \{(x_v^k, y_v^k, z_v^k) : 1 \leq v \leq V\} \in \mathcal{V}^k$ and $\alpha_v^k \in [0, \pi]$ and $\beta_v^k \in [0, 2\pi]$,
- (b) The spectrum coefficients $\{l_n^m(i)\}$ of n_L natural illumination maps (e.g. [55, 105]) with $i = 1, 2, \dots, n_L$, $n \in [0, N]$ and $m \in [-n, n]$ such that N is the maximum illumination order,
- (c) The spectrum coefficients of teeth reflectance $\{a_{pr}^q(j)\}$ with $p, r \in [0, P]$ and $q \in [0, \min(p, r)]$ where P is the maximum reflectance order. The j -th BRDF is defined as Wolff-Oren-Nayar model with surface roughness $\sigma \in [0.7, 2]$ and enamel refractive index $\eta \in 1.62 \pm 0.02$,
- (d) Q : percentage of eigenvalues to be maintained in subspace construction.

OUTPUT: Subspace projection matrix $\tilde{\mathbf{W}} = [\mathbf{w}_1 | \mathbf{w}_2 | \dots]$ such that a jaw occlusal image can be written as $\mathbf{e} = \tilde{\mathbf{W}}\mathbf{y}$ in the vector form.

- 1: **Build illumination prior:** Compute illumination expectation $E\{l_n^m\}$ and covariances $E\{l_n^m l_{n'}^{m'}\}$ as in [44].
- 2: **Build reflectance prior:** Compute BRDF expectation $E\{a_{pr}^q\}$ and covariances $E\{a_{pr}^q a_{p'r'}^{q'}\}$ as in [44].
- 3: **Compute irradiance harmonics:**
 - (a) Compute surface normal distribution $p(\alpha_v^k, \beta_v^k)$.
 - (b) Compute Wigner rotation matrices $\mathcal{D}^n(\alpha_v^k, \beta_v^k)$ where $n \in [0, N]$ as in [61].
 - (c) Compute *albedo*-modulated irradiance harmonics $\mathcal{B}_{npr}^{mq}(\alpha_v^k, \beta_v^k)$ for $n \in [0, N]$ and $p \in [0, P]$ and sort them according to their average power content to obtain $\mathcal{B}_s(\alpha_v^k, \beta_v^k)$ for $s = 1, 2, \dots, S$.
- 4: **Decompose irradiance harmonics:** Perform PCA on $\mathbf{B} = [\mathbf{b}_1 \dots \mathbf{b}_S]$ to obtain the optimal $\tilde{\mathbf{W}}_B$ from the D' -eigenvectors corresponding to its largest D' -eigenvalues where D' is set to maintain $Q\%$ of harmonics variation such that $D' \leq S$.
- 5: **Compute harmonic coefficients statistics:** Compute expectation $E\{c_s\} = E\{l_n^m\}E\{a_{pr}^q\}$ and covariances $E\{c_s c_{s'}\} = E\{l_n^m l_{n'}^{m'}\}E\{a_{pr}^q a_{p'r'}^{q'}\}$ where the respective indices are given by the ordering function of the irradiance harmonics functions $s \rightarrow \{n, m, p, r, q\}$.
- 6: **Compute Υ matrix:**

$$\Upsilon = \sum_s \sum_{s'} E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} \mathbf{c}_s^B (\mathbf{c}_{s'}^B)^T$$

where $E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} = E\{c_s c_{s'}\} - E\{c_s\}E\{c_{s'}\}$.

- 7: **Compute weighting matrix:** Compute the weighting matrix \mathbf{A}_W as the D' -eigenvectors of Υ^T .
 - 8: **Compute subspace projection matrix:** Compute $\tilde{\mathbf{W}}^T = \mathbf{A}_W \tilde{\mathbf{W}}_B^T$.
-

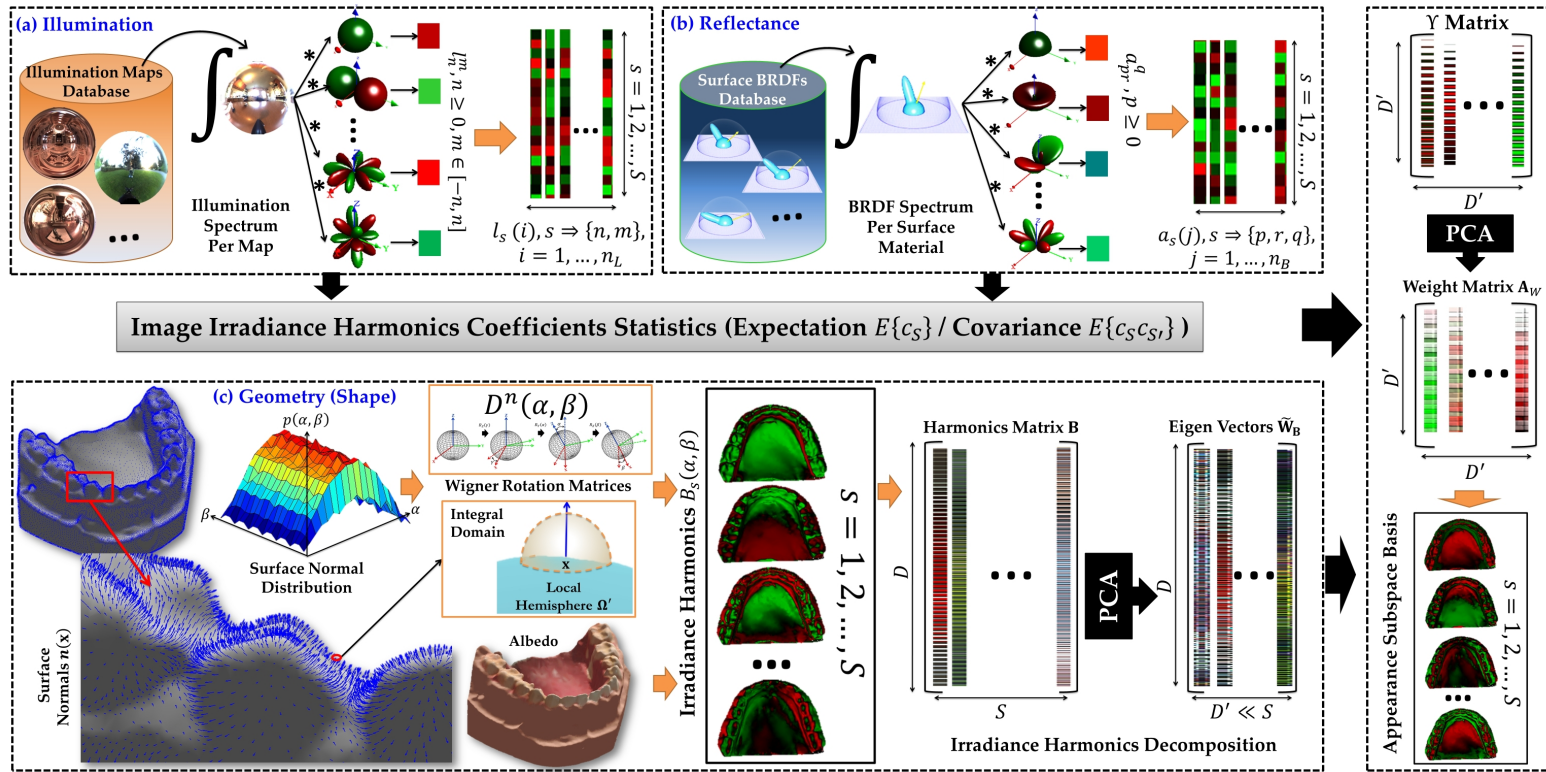


Figure 9.13: Computations of the proposed appearance subspace construction algorithm. Spherical harmonics is used to compute the illumination spectrum of a database of environment maps [55, 105] while the Helmholtz HSH-based basis is used to compute the reflectance spectrum of a database of teeth reflectance. Image irradiance harmonics are computed based on the visible surface normals of the object of interest and then sorted according to their average power content. The illumination and reflectance spectrum components are sorted accordingly to provide a single index-based notation. The irradiance harmonics are then decomposed to obtain \tilde{W}_B while the Υ matrix is computed according to the harmonics coefficients statistics and the projection of the irradiance harmonics onto the subspace spanned by the columns of \tilde{W}_B . The weighting matrix is then computed based on the eigenvectors of the Υ matrix. Note greenish and reddish parts in the harmonics, basis functions and spectrum components connotes positive and negative values, respectively.

where $\vec{n}(\mathbf{x}) = (\alpha(\mathbf{x}), \beta(\mathbf{x}))$, \mathcal{W}_s is the s -th column of $\tilde{\mathbf{W}}$ and $\{y_s\}$ is the weighting vector which results from projecting E onto the subspace spanned by $\{\mathcal{W}_s\}$.

In matrix notation, let $\mathbf{e} \in \mathbb{R}^D$ be an image vector with D -pixels, $\tilde{\mathbf{W}} = [\mathcal{W}_1(\vec{n}(\mathbf{x})), \dots, \mathcal{W}_S(\vec{n}(\mathbf{x}))] \in \mathbb{R}^{D \times S}$ be the subspace projection matrix having appearance basis as its columns, where S is the number of basis, and $\mathbf{y} \in \mathbb{R}^S$ be a vector of irradiance coefficients. Hence the discrete version of the brightness constraint becomes,

$$\epsilon = \sum_{\mathbf{x}} \left[E(\mathbf{x}) - \tilde{\mathbf{W}}(\vec{n}(\mathbf{x}))\mathbf{y} \right]^2 = \|\mathbf{e} - \tilde{\mathbf{W}}\mathbf{y}\| \quad (9.20)$$

9.8.2 Harmonic Projection (HP) Irradiance Model

Representing the surface reflectance function in terms of the proposed appearance basis allows inferring the coupled effect of illumination and reflectance of the input irradiance signal as follows. Given an input image irradiance E and the appearance basis matrix $\tilde{\mathbf{W}}$ of a 3D object (a human jaw in particular) derived from its shape and albedo, the irradiance coefficients $\hat{\mathbf{y}}$ are deduced to best match the input irradiance. This results in an over-determined linear system of equations $\mathbf{e} = \tilde{\mathbf{W}}\mathbf{y}$ which can be solved for $\hat{\mathbf{y}}$ using singular value decomposition (SVD) (See Algorithm 3).

If the input irradiance and the appearance basis used to compute the coefficients $\hat{\mathbf{y}}$ belong to the same object, one can reconstruct the input irradiance from these coefficients, i.e. $\mathbf{h} = \tilde{\mathbf{W}}\hat{\mathbf{y}} = \mathbf{e}$, where \mathbf{h} denotes what is called *harmonics projection* (HP) irradiance. However in the general case, the basis images $\tilde{\mathbf{W}}$ would belong to an object which is different from the one that was imaged, nonetheless they belong to the same object class e.g. different realizations of a human jaw. Thus the reconstructed irradiance \mathbf{h} provides a mean of encoding the illumination and reflectance of the input irradiance while maintaining the identity of the object whose basis are used in the reconstruction process.

While the shape model $\mathbf{s} = \bar{\mathbf{s}} + \mathbf{P}_s\boldsymbol{\lambda}_s$ and the albedo model $\mathbf{a} = \bar{\mathbf{a}} + \mathbf{P}_a\boldsymbol{\lambda}_a$ are constructed in a pre-processing (offline) step, the HP model is constructed when the input irradiance is given to the framework in order to incorporate the illumination and reflectance conditions of the given irradiance into the prior information. In particular, given the irradiance vector \mathbf{e} and the appearance basis matrix $\tilde{\mathbf{W}}_k$ of object instance k ,

the HP irradiance for each data sample can be stacked into a vector, \mathbf{H}_k , i.e.,

$$\mathbf{H}_k = \tilde{\mathbf{W}}_k (\mathbf{V}\mathbf{S}^{-1}\mathbf{U}^T) \mathbf{e} \quad (9.21)$$

where $[\mathbf{U}, \mathbf{S}, \mathbf{V}] = \text{svd}(\tilde{\mathbf{W}}_k)$. The HP eigenvectors (\mathbf{h}_k) and variances ($\sigma_{h,k}^2$) are obtained using the same procedure discussed in case of shape model construction. The resulting HP model can be written as,

$$\mathbf{h} = \bar{\mathbf{h}} + \sum_{k=1}^{K-1} \lambda_k^h \cdot \mathbf{h}_k = \bar{\mathbf{h}} + \mathbf{P}_h \boldsymbol{\lambda}_h \quad (9.22)$$

where $\mathbf{P}_h = [\mathbf{h}_1, \mathbf{h}_2, \dots]$ are the HP eigenvectors and $\boldsymbol{\lambda}_h$ is the set of HP coefficients. Note that the subscript h refer to harmonic.

9.8.3 Shape and Albedo Recovery

While Equation 9.18 can be solved in an iterative manner to infer the underlying shape as in [33], the inherit relation between the HP irradiance \mathbf{h} and the corresponding shape \mathbf{s} and albedo \mathbf{a} can be cast into a regression framework resulting into the HP-to-shape and HP-to-albedo models. In this case, the shape and albedo is solved for using a series of matrix operations guaranteeing faster shape recovery when compared to its iterative counterpart. This was proven to yield comparable results in terms of reconstruction accuracy [34].

Dimensionality reduction is performed using PCA to construct 3D shape, albedo and HP models where the shape/albedo/HP coefficients are used to build the regression models rather than the original shape, albedo and HP instances. Thus, instead of using the high dimensional vectors \mathbf{s}_k , \mathbf{a}_k and \mathbf{h}_k into the regression, they are replaced by their respective coefficients $\boldsymbol{\lambda}_s^k$, $\boldsymbol{\lambda}_a^k$ and $\boldsymbol{\lambda}_h^k$, where the HP coefficients are considered the independent variable while the shape and albedo coefficients are the dependent variables. Principal Component Regression (PCR) [138] is used to avoid random noise which might exist in the dependent and independent variables. It also deals with the small-sample-size (SSS) problem where the ratio between observations and variables is usually low. Steps are enumerated in Algorithm 8 and Figure 9.8.2 shows a block diagram of the offline/online processes for the proposed shape/albedo recovery approach.

Algorithm 8 Model-based Shape and Albedo Recovery

INPUT:

- (a) Input image irradiance, E ,
- (b) Anatomical jaw landmarks being detected/annotated on E ,
- (c) Aligned shape and albedo K -samples: $(\mathbf{S}_1, \mathbf{A}_1)$ to $(\mathbf{S}_K, \mathbf{A}_K)$,
- (d) The corresponding appearance basis matrix $\tilde{\mathbf{W}}_k$ for each pair $(\mathbf{S}_k, \mathbf{A}_k)$.

OUTPUT:

- (a) Recovered shape, $\tilde{\mathbf{s}}$,
- (b) Recovered albedo, $\tilde{\mathbf{a}}$

- 1: **Input irradiance warping:** Use the given sparse landmarks to infer a warping function \mathbf{f} between image pixels and the vertices of the mean jaw shape $\bar{\mathbf{s}}$ using 3D thin-plate spline such that $\mathbf{f}(\mathbf{x}^*) = \mathbf{x}$, where $\mathbf{x}^* = [x^*, y^*, 1]^T$ is an image pixel on the xy -plane and $\mathbf{x} = [x, y, z]^T$ is the mean shape vertex.
 - 2: **Shape and albedo prior:** Build the shape and albedo models from the given aligned samples using PCA where $\mathbf{s} = \bar{\mathbf{s}} + \mathbf{P}_s \boldsymbol{\lambda}_s$ and $\mathbf{a} = \bar{\mathbf{a}} + \mathbf{P}_a \boldsymbol{\lambda}_a$.
 - 3: **HP prior:** Given the warped irradiance \mathbf{e} in the vector form, solve for the HP irradiance for each data sample such that $\mathbf{H}_k = \tilde{\mathbf{W}}_k (\mathbf{V}\mathbf{S}^{-1}\mathbf{U}^T) \mathbf{e}$, where $[\mathbf{U}, \mathbf{S}, \mathbf{V}] = \text{svd}(\tilde{\mathbf{W}}_k)$. Then use PCA to construct the HP model such that, $\mathbf{h} = \bar{\mathbf{h}} + \mathbf{P}_h \boldsymbol{\lambda}_h$.
 - 4: **Samples shape coefficients:** Solve for $\boldsymbol{\lambda}_s^k = P_s^T (\mathbf{S}_k - \bar{\mathbf{s}})$
 - 5: **Samples albedo coefficients:** Solve for $\boldsymbol{\lambda}_a^k = P_a^T (\mathbf{A}_k - \bar{\mathbf{a}})$
 - 6: **Samples HP irradiance coefficients:** Solve for $\boldsymbol{\lambda}_h^k = P_h^T (\mathbf{H}_k - \bar{\mathbf{h}})$
 - 7: **Setup matrices for PCR:**
 - (a) $\mathbf{T} = [(\boldsymbol{\lambda}_h^1)^T, \dots, (\boldsymbol{\lambda}_h^K)^T]$
 - (b) $\mathbf{U}_s = [(\boldsymbol{\lambda}_s^1)^T, \dots, (\boldsymbol{\lambda}_s^K)^T]$
 - (c) $\mathbf{U}_a = [(\boldsymbol{\lambda}_a^1)^T, \dots, (\boldsymbol{\lambda}_a^K)^T]$
 - 8: **Build two PCR models:**
 - (a) Construct $\tilde{\mathbf{C}}_s = (\mathbf{T}^T \mathbf{T})^{-1} \mathbf{T}^T \mathbf{U}_s$
 - (b) Construct $\tilde{\mathbf{C}}_a = (\mathbf{T}^T \mathbf{T})^{-1} \mathbf{T}^T \mathbf{U}_a$
 - 9: **Solve for the HP coefficients of the input irradiance:** $\tilde{\boldsymbol{\lambda}}_h = P_h^T (\mathbf{e} - \bar{\mathbf{h}})$
 - 10: **Solve for the shape coefficients of the input irradiance:** $\tilde{\boldsymbol{\lambda}}_s^T = \tilde{\boldsymbol{\lambda}}_h^T \tilde{\mathbf{C}}_s$
 - 11: **Solve for the albedo coefficients of the input irradiance:** $\tilde{\boldsymbol{\lambda}}_a^T = \tilde{\boldsymbol{\lambda}}_h^T \tilde{\mathbf{C}}_a$
 - 12: **Solve for the recovered shape and albedo:**
 - (a) Compute $\tilde{\mathbf{s}} = \bar{\mathbf{s}} + \mathbf{P}_s \tilde{\boldsymbol{\lambda}}_s$
 - (b) Compute $\tilde{\mathbf{a}} = \bar{\mathbf{a}} + \mathbf{P}_a \tilde{\boldsymbol{\lambda}}_a$
-

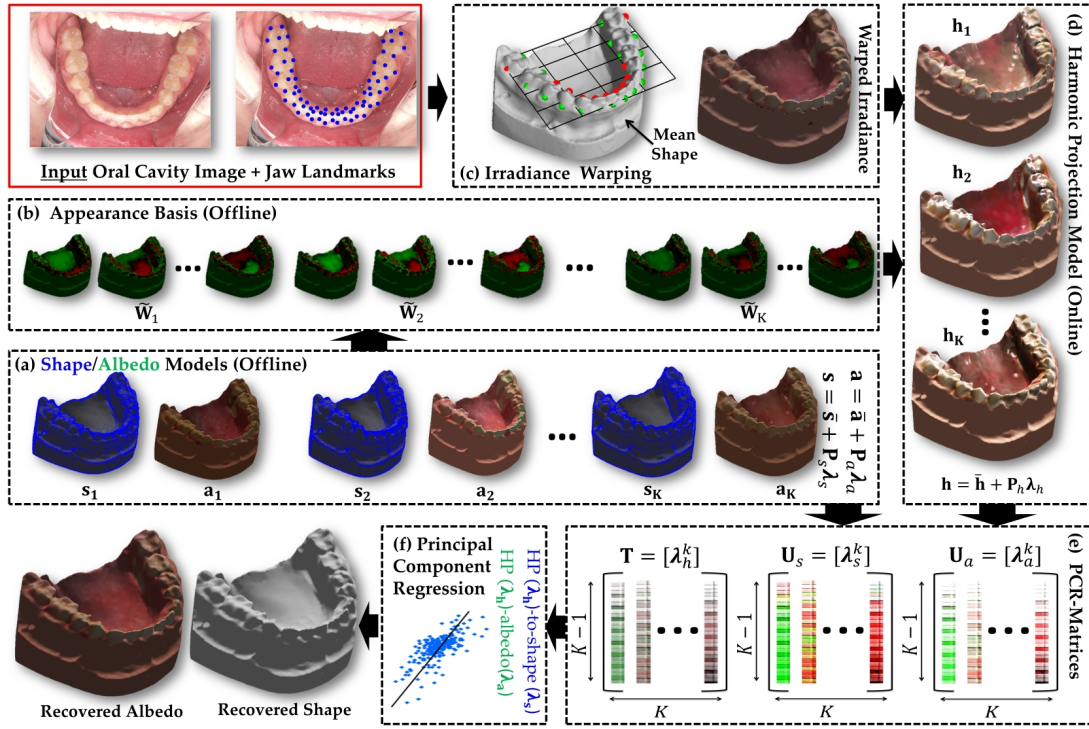


Figure 9.14: Block diagram of the proposed model-based human jaw shape recovery: (a) An aligned ensemble of the shapes and albedos of human jaws is used to build the 3D shape and albedo models. (b) Given the albedo and surface normals (defining the shape) of a certain jaw in the ensemble, appearance bases are constructed using Algorithm 7. Given an input oral cavity image under general unknown illumination and a set of human jaw anatomical landmark points: (c) Dense correspondence is established between the input irradiance and the mean jaw shape using 3D thin-plate splines. (d) The input image, in the reference frame, is projected onto the subspace spanned by the appearance basis of each sample in the ensemble which are scaled (using the projection coefficients) and summed-up to construct the harmonic projection (HP) irradiance which encodes the illumination and reflectance conditions of the input image. Such images are then used to construct an HP model of the input image. (e,f) The inherit relation between the HP irradiance and the corresponding shape and albedo is cast as in a regression framework where principal component regression is used to solve for shape and albedo coefficients to recover the shape and albedo of the input image.

9.9 Application to Dental Restoration

Dental restoration, in particular inlays and onlays (indirect fillings), is a conservative alternative to full coverage of dental crowns where it offers a long lasting reparative solution to tooth decay or similar structural damage. In contrast to traditional (direct) dental fillings which are molded during a dental visit, inlays and onlays are fabricated in a dental lab where an appropriate standard tooth model is picked from a tooth database. Such process requires dental technicians who are highly trained experts in tooth anatomy. *Inlay* refers to a restoration process where the material is fitted and

bonded within the tooth groove, on the other hand *onlay* restoration is recommended when the structural damage requires the inclusion of one or more tooth cusps or even full coverage of the biting surface. See Figure 9.9.

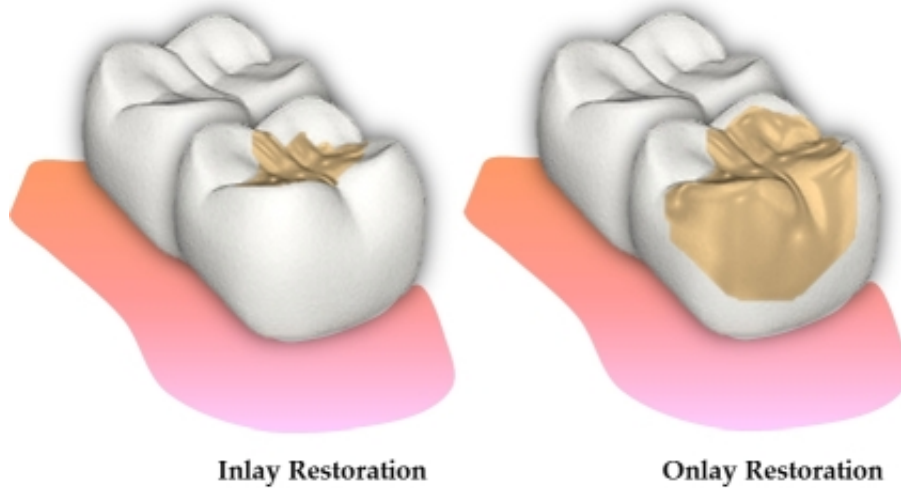


Figure 9.15: Illustrative example of inlay versus onlay restorations. Whereas inlay involves fitting and bonding a restorative material within the center of the tooth, onlay restoration is recommended when the tooth damage extends to include one or more cusps.

Over the past decade, Computer-Aided Design (CAD) and Computer-Aided Manufacturing (CAM) have become ubiquitous in many major dental groups and laboratories [188, 189] where CAD/CAM technologies are used to produce different types of dental restorations including inlays and onlays. A typical CAD system consists of three main components acquisition, modeling, and milling modules. The success of a restoration process primarily rely on the acquisition and modeling of an accurate 3D shape of the occlusal surface of interest for manufacturing purposes.

In literature, various methods have been studied for tooth surface design. The first dental CAD/CAM system of Duret, developed in France, is very complex. It can only be used for simple restoration design [190]. In the Minnesota system, affine transformation of a 3D tooth model was used to adapt to the scanned tooth. Instead of covering all kinds of restoration, they mainly considered the production of crowns. In the GN-I system [191], the contour line of the tooth surface was determined by selecting the size of the reference line and the design area and the occlusal surface was designed through simulation. It is difficult to establish a relation due to lack of different forms of tooth surface design. In the CEREC system [192], three shape design steps were proposed:

extrapolation, correlation, and function, but due to the interactive construction process, the surface of missing part had to be modified manually after inserted into tooth cavity. This system is based on multi sensor scanner. Sporrying *et al.* [161] proposed a statistical model of a selection of tooth shapes and a reconstruction of missing data by including information the position and anatomy of other teeth. Their system depends on hand picked landmark detection. Moreover, it is evaluated on cast models scanned by 3 shape laser scanner. Mehl *et al.* [193] proposed a statistical model of a selected tooth shapes which relies on optical images instead of 3D. The manual annotation is still a requirement in their system which is on onlay and inlay restoration. It does not include the missing crown since it is a tooth based model. Zheng *et al.* [148] proposed a 3D morphing approach which eliminates the manual annotation. Their system based on accurate 3D acquisition system. The approach proposed in this chapter can serve as a fully automated framework for tooth restoration from a single optical image where the need for hand picked jaw features is eliminated using active shape model [166].

9.10 Experimental Results

This section shows experiments to evaluate the performance of the proposed framework in recovery 3D models for human jaws. Upper and lower jaw models are constructed from eight subjects (5 males and 3 females with ages range from 16 to 46 years old) using their oral cavity images and the CT-scan of their respective molds (lower and upper jaws). There are two samples per subject, one *pre-repair* jaw and another *post-repair* jaw, referring to the jaw status *before* and *after* applying an orthodontic teeth alignment process, respectively. The statistical priors (shape, albedo and appearance models) are trained using out-of-training samples with pre- and post-repair instances.

The core hypothesis of this work is that using prior information would enable obtaining a plausible and more accurate jaw model for the human jaw while covering regions that the classical SFS approach does not handle. As such, one compares the proposed approach with a recently evaluated SFS approach [194] for tooth surface reconstruction based on the work of Ahmed and Farag in [195, 196]. In their work, the image irradiance equation was formulated as a Partial Differential Equation (PDE) to solve for surface gradients, where the theory of viscosity solutions for Hamilton-Jacobi type equations provide a good framework of SFS algorithms [197]. Carter *et al.* in [194] concluded

that, based on quantitative error analysis, the Oren-Nayar reflectance model has been proven to be the most ideal SFS formulation for reconstructing tooth crown surface from a single image. Nonetheless, their work did not incorporate the object physical characteristics such as surface roughness and Fresnel reflection due to different refractive indices of the saliva and the tooth material. Hence the proposed approach is compared with the one proposed by Ahmed and Farag [196] where Wolff-Oren-Nayar reflectance model was assumed. Moreover, the presented work is compared with one’s earlier work in [49] which assumes Lambertian reflectance.

In order to share the same metric coordinate frame, the average jaw shape \bar{s} (along with its anatomical landmarks) is used as a reference to establish a dense correspondence between the groundtruth CT scan of the jaw mold corresponding to each testing image¹ and the reconstructed shape. The alignment proceeds as follows. Procrustes-based rigid registration [173] is used to filter out translation, scale and rotation followed by 3D thin-plate splines [146, 183] for non-rigid registration.

We assess the reconstruction accuracy in accordance to an error estimator which is based on the Root Mean Square (RMS) error between the 3D points from the CT scan and the corresponding reconstructed surface points. Note that errors are computed based on the surface points of the visible crowns to exclude errors which might rise from the reconstruction of the mold base.

9.10.1 Visible Crowns Reconstruction

To evaluate the proposed approach, out-of-training jaw samples are reconstructed and compared against the ground truth CT-scan. Four types of samples are considered: (a) pre-repair and (b) post-repair lower jaw, (c) pre-repair and (d) post-repair upper jaw. Along with the groundtruth shapes, Figures 9.16 and 9.17 show a sample of shape and albedo reconstruction of upper and lower jaws, respectively. It important to note that SFS only recovers a height map (2.5D) of the input image where there is no metric information reserved. With the metric prior used to train the offline shape model, the proposed approach reconstructs the triangular mesh (3D) corresponding to the input image. Further, most of SFS approaches assume known parameters of surface reflectance and point light source with known direction. In contrast, one’s approach explicitly accounts for complex illumination and models surface reflectance without the need of the

¹Note that out-of-training testing scenario is followed.

exact reflectance parameters. In Figures 9.16 and 9.17, one can observe the closeness of the SSFS-based reconstruction to the groundtruth shape when compared to the recovery shape from SFS. This emphasizes the role of incorporating prior-information for shape recovery as well as appearance modeling.

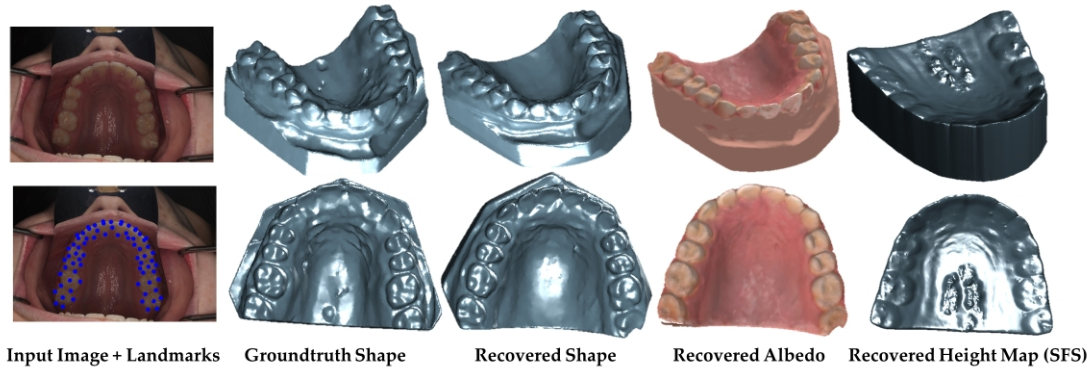


Figure 9.16: Sample reconstruction result of an upper (post-repair) jaw (bottom row shows the top-view of the occlusal surface).

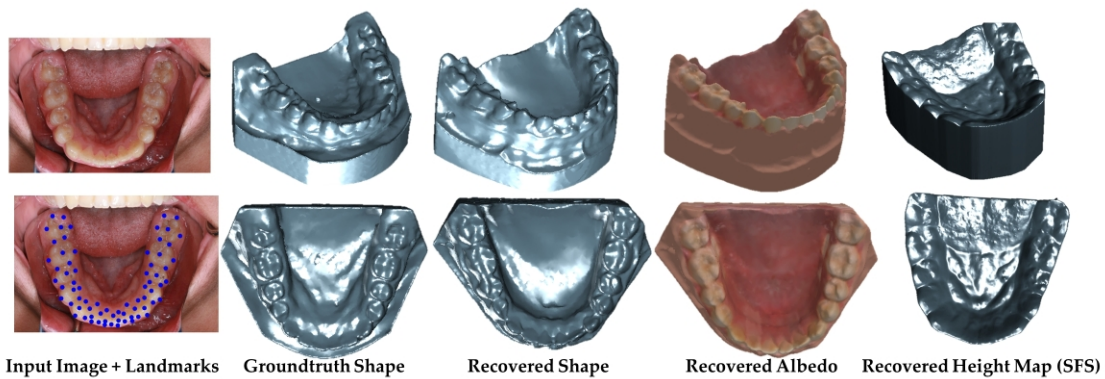


Figure 9.17: Sample reconstruction result of a lower (post-repair) jaw (bottom row shows the top-view of the occlusal surface).

Table 9.1 reports the RMS error in *mm* between the 3D points from the CT scan and the corresponding reconstructed surface points. For the sake of comparison, one also includes one’s earlier results in case of assuming Lambertian reflectance [49]. Notice that the error values of the SSFS-based reconstructions are minimal when compared to SFS-based reconstruction. Pre-repair error values are also smaller than post-repair values in most of the samples, indicating that the statistical prior capability of capturing irregular tooth shapes and locations.

A natural question to be asked is how to make use of SFS results and SSFS? Of course, SFS is based on the visible surface of the jaw; at best the crown would be

Table 9.1: Average *whole jaw* surface reconstruction accuracy (RMS) in *mm*

Jaw Type	Proposed non-Lambertian SSFS	Lambertian SSFS [49]	SFS [196]
Upper, Pre-repair	0.6289	2.08999	15.2995
Upper, Post-repair	0.6689	2.02334	16.3098
Lower, Pre-repair	0.6714	3.11911	12.1241
Lower, Post-repair	0.8073	2.57112	13.4959

possibly constructed, while SSFS constructs the entire jaw. On the other hand, SFS provides the object-specific constructions. A logical thing would be to enhance the SSFS with SFS, by morphing the upper part of the model with the crown portion generated from SFS. Another question arises where even smaller reconstruction errors are needed, to that end it is worth pointing out that these results are based on a model that is being trained on a small ensemble of 16 jaws (8 pre-repair and 8 post-repair). With a large enough ensemble of objects, credible shape, albedo and appearance models would be possible, which when morphed to the crown reconstructions would produce a more realistic jaw.

9.10.2 Results on Tooth Restoration

In this section, the performance of the proposed framework in restoration of missing and broken human teeth is assessed. The reconstruction experiments are conducted on real images of jaw that are manually distorted to mimic the missing and broken tooth as shown in Figures 9.18 and 9.19. This allows validation with respect to groundtruth 3D information. The test is done on 8 subjects with different distortion, inlays and onlays, on different premolars and molars. One-of-out training is still used such that the testing subject is not included in the model training.

Tables 9.2 and 9.3 show the average reconstruction error in *mm* for inlay and onlay restorations, respectively. The results show that, on average, the small tooth damage (*i.e.* inlay) leads to smaller errors when compared to bigger ones where one or more cusps are missing, *i.e.* onlay.

9.11 Summary

This chapter presented an affordable, flexible, automatic dental tool for the reconstruction of the clinically visible part of the human jaw. It was based on a single captured

Table 9.2: Average *inlay* surface reconstruction accuracy (RMS) in *mm*

Jaw Type	Proposed SSFS	SFS [196]
Upper, Pre-repair	0.6808	13.1934
Upper, Post-repair	0.5859	12.6267
Lower, Pre-repair	0.7559	13.7182
Lower, Post-repair	0.6901	12.8085

Table 9.3: Average *onlay* surface reconstruction accuracy (RMS) in *mm*

Jaw Type	Proposed SSFS	SFS [196]
Upper, Pre-repair	0.6337	16.0437
Upper, Post-repair	0.6399	16.8294
Lower, Pre-repair	0.8143	15.0697
Lower, Post-repair	0.9675	15.4122

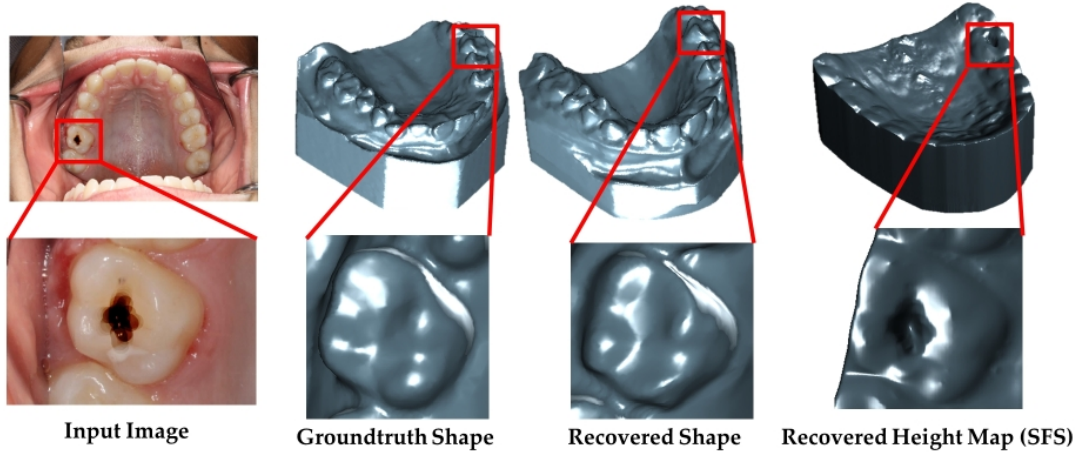


Figure 9.18: Sample inlay reconstruction result of an upper (post-repair) jaw (bottom row shows a zoom-in view of the damaged teeth).

optical image and a statistical shape recovery approach which makes use of a small number of measured points to construct a plausible 3D model through a learned correspondence based on a measured human jaw dataset. While most shape-from-shading (SFS) approaches assume known parameters of surface reflectance and point light source with known direction, the presented work has relaxed such assumptions using the harmonic expansion of the image irradiance equation where one was able to incorporate prior information about natural illumination and real world surface materials. To provide optimal appearance basis of occlusal surfaces, an analytic formulation for appearance subspace reconstruction is proposed to capture the full behavior of complex illumination and teeth reflectance characteristics. The brightness constraint was then cast as a




Figure 9.19: Sample onlay reconstruction result of a lower (post-repair) jaw (bottom row shows a zoom-in view of the damaged teeth).

regression problem, which allows for the rapid computation of the solution. Principal Component Regression (PCR) algorithm was deployed which is composed of a sequence of matrix operations; the approach in this work can recover 3D shapes much faster than its iterative counterpart, without compromising the integrity of the results. The results demonstrated the effect of adding statistical prior as well as appearance (illumination and reflectance) modeling on the accuracy of the recovered shape. The applicability of the proposed approach to a dental application encompassing tooth restoration was investigated. The system showed ability to handle different missing and broken parts, inlay and onlay, with comparable accuracy to the existing methods without manual intervention from a single image. The restoration surfaces match smoothly well with adjacent tooth, the morphological character of the tooth are retained. The next step is to investigate the fusion of SFS and SSFS where SFS provides the object-specific constructions while SSFS is perform shape recovery based on partial information.

EPILOGUE

The shape and reflectance of a three-dimensional object provide invariant properties to changes in the imaging formation process, *e.g.* illumination and viewpoint changes. Extraction of such information given a single image, as an under-constrained problem, is an inherent challenge in many vision tasks such as recognition. This involves adding assumptions to the image formation process such as point light source, known reflectance model with homogenous characteristics.

The appearance of an object, *i.e.* shading, is given by the *image irradiance equation* [24], where its low-dimensional representation under *unknown* arbitrary lighting and reflectance is a fundamental process for many computer vision tasks including illumination modeling [6], surface reflectivity estimation/analysis [11, 32], shape and albedo recovery [33, 34, 35, 37], shape from shading [36, 196, 198], photometric stereo [38], object detection and recognition [39], to name a few.

Theoretically, due to the arbitrariness of the lighting function, the space of all possible images of a fixed-pose object under all possible illumination conditions is infinite dimensional [25]. Yet, statistical modeling is widely devised to provide a low-dimensional linear subspace of image irradiance by performing Principal Component Analysis (PCA) on a large set of images under various imaging conditions [27]. However, this approach suffers from the need of special acquisition setup to control and measure the lighting function while keeping the camera and the object fixed. In addition, the constructed subspace is biased towards the sampled illumination directions. Moreover, either surface reflectance properties are assumed or known a-priori.

On the other hand, appearance models can be constructed using phenomenological models which capture variations of surface appearance with respect to illumination conditions and surface reflectance properties through mathematical abstraction of the

reflection process. In particular, the harmonic expansion of the image irradiance equation can be used to derive an analytic subspace to approximate images under fixed pose but different illumination conditions; it has been shown that surfaces acts as a low-pass filter on the incident illumination, yielding a band-limited reflected light, introducing the frequency-space representation of image irradiance.

This dissertation has presented a theoretical framework allowing for the derivation of low-dimensional appearance subspace under natural illumination and arbitrary surface reflectance while taking into account the intrinsic dimensionality of the imaging domain. With a low-dimensional generative model for image irradiance, one was able to incorporate prior information about complex illumination and real-world materials. One also demonstrated how the presented work can be applied in the context of vision algorithms that determine surface geometry.

With the theoretical and practical components of this dissertation, the sequel discusses the dissertation's main contributions, the bigger picture, highlight some limitations and end by discussing directions for future research.

10.1 Summary of Contributions

This dissertation has proposed a number of advances in modeling image irradiance along with potential applications which go some way to addressing the research questions formulated in Chapter 1. The findings from this dissertation make several contributions to the current literature which can be summarized as follows:

- **Surface reflectance basis:** In Chapter 4 a complete, orthonormal basis was proposed to provide a compact and efficient representation for surface bidirectional reflectance distribution function (BRDF), which is defined on the cartesian product of two hemispheres. The proposed basis, which are defined in terms of hemispherical harmonics (HSH), preserve the Helmholtz reciprocity property of BRDFs while avoid the computational complexity inherited from Zernike polynomials that are usually used to construct hemispherical basis. By analytical and experimental comparison, the proposed basis captured surface reflectance characteristics with lower reflectance order compared to similar basis in literature. These results were validated based on analytical reflectance models as well as scattered BRDF measurements.

- **Image irradiance harmonics:** Through the mathematical modeling of the reflection process, in Chapter 5 the definition of image irradiance harmonics was formulated which enabled decoupling illumination and reflectance from the underlying geometry and pose. These harmonics encode the orthogonality relation between spherical harmonics basis used for illumination representation and the deployed reflectance basis. The image irradiance representation was further validated based on the proposed reflectance basis compared to similar representations in literature using different surface materials. By investigating the energy content maintained by the proposed image irradiance basis for natural materials, one found that specular materials tend to require more basis functions when compared to diffuse ones. Yet the presented basis captures similar cumulative energy content at lower illumination orders when compared to similar representations in literature.
- **Analytic subspace construction:** In Chapter 6, a bilinear approach was proposed to analytically construct irradiance subspace in order to tackle the inherent problem of small-sample-size and curse of dimensionality. The process of finding the analytic subspace was posed as establishing a relation between its principal components and that of the irradiance harmonics basis functions. One further showed how to incorporate prior information about natural illumination and real-world surface reflectance characteristics in order to capture the full behavior of complex illumination and non-Lambertian reflectance. The proposed analytic bilinear PCA showed significant decrease in dimensionality when compared to the linear counterpart while attaining the lowest reconstruction errors. It was further tested against noisy input showing robust image representation even at low SNR levels.
- **Missing irradiance information:** In Chapter 7, one proposed to cast errors introduced due to irradiance incompleteness as: (1) statistical outliers which were determined and rejected using robust statistics and (2) local spatial erroneous continuous regions where Markov Gibbs random field with the homogenous isotropic Potts model was adopted to model the incompleteness's spatial interaction. Results demonstrated the effectiveness of the proposed algorithms in handling high levels of contiguous missing information compared to the conventional image projection. In particular, MGRF and robust estimation using Geman-McClure func-

tion outperformed the Singular Value Decomposition (SVD) performance approach which is very sensitive to the presence of missing information even at low levels. In the meantime, the performance of Lorentzian function approaches SVD due to the presence of errors caused by fitting an input image to subspace bases constructed using different geometrical object. The approach was further incorporated in a model-based shape from shading framework which assumed Lambertian reflectance which gave rise to incomppliance to the assumed model. One showed that conventional SVD was very sensitive to the presence of missing information even at low levels while the proposed harmonic projection handled corrupted irradiances.

- **Shape recovery based on appearance modeling:** In Chapters 8 and 9, the application of the proposed imaging model was presented to construct appearance models for human faces and clinical crowns of human jaws. While most shape-from-shading (SFS) approaches assume known parameters of surface reflectance and point light source with known direction, this work has relaxed such assumptions using the harmonic expansion of the image irradiance equation where one was able to incorporate prior information about natural illumination and real world surface materials. The results showed the effectiveness of the proposed appearance model in reconstructing the geometrical structure of the object of interest.

One of the more significant findings to emerge from this dissertation is that frequency-based image irradiance representation has made constructing appearance models affordable where the need for capturing a “large enough” ensemble of images under different illumination conditions is eliminated. The relevance of representing hemispherical surface reflectance using hemispherical basis is clearly supported by the current dissertation findings.

10.2 Limitations and Suggested Future Directions

Our analysis also lead to a number of interesting observations, some of which may be considered as lines of future research. Below one outlines some limitations as well as several important directions of further research.

- **Global illumination:** This dissertation has addressed complex illumination as well as arbitrary reflectance. Further steps may be taken towards encompassing

effects such as inter-reflections and cast shadows to lower the complexity of global illumination simulations while achieving interactive rendering.

- **Pose and geometry:** While the proposed appearance model captures variations due to illumination and reflectance, in order to complete the view, one might need to investigate adding variable surface geometry and camera poses. As such completing the two aspects of image formation, photometric (this dissertation) and geometric aspects.
- **High frequency illumination:** Whereas one's focus was to accurately model surface reflectance, using spherical harmonics to represent light assumes low-frequency illumination. One suggests investigating appearance models based on different illumination representation to capture all-frequency light and non-convex geometrical objects.
- **Human perception:** This work has focused on the physical computational aspect of the image irradiance. Nonetheless, human perception response is logarithmically related to the physical intensity. As such, an important future direction of research is considering the perceptual aspect of the image formation process to bridge the gap between photorealism and interactivity.
- **More generalized reflectance model:** One suggests employing a more generalized reflectance model which may be obtained if the Lambertian component in Cook-Torrance model is replaced by a more realistic diffuse model such as Oren-Nayar model or Wolff-Oren-Nayar model. The new model will have three parameters that control the percentages of the diffuse and the specular reflections, in addition to the degree of the surface roughness. However, various types of validation tests and experiments on real data samples for object reflectance are needed to be performed before this model can be used to model visual appearance.
- **Image-based reflectance modeling:** There are different approaches exist for BRDF estimation from one or more images. Most of them rely on the availability of prior information such as object's shape and known/controlled illumination. The relaxation of the required apriori lighting information is still a challenge which needs to be addressed. One believes that modeling surface roughness will aid the

illumination recovery process. Furthermore, cast shadows inherited from the non-convexity of the object's surface can be exploited to recover lighting directions.

- **Subsurface scattering:** This work ignored subsurface transport to reduce the domain of surface reflectance. This is a reasonable approximation for most of computer vision tasks. Yet, one wonders if similar analysis can be done to the general rendering equation which incorporate subsurface scattering where physically-based models for subsurface scattering of light provide attractive tools for accurate image synthesis.
- **Statistical modeling:** there are many open questions related to statistical modeling of prior information. This involves: (1) the amount of training data needed to build a statistical model that accurately generalize to any class realization, (2) the effect of dimensionality of the object space and (3) the effect of nonlinear analysis in contrast to linear subspace learning.
- **Simultaneous estimation of scene parameters:** One interesting question is whether one can incorporate the proposed appearance model in a unified framework which allow the estimation of illumination and reflectance with geometry estimation.

10.3 Summary

In brief, this dissertation has presented an analytical framework for appearance subspace construction which lead to a low-dimensional generative model for image irradiance while taking into account the intrinsic dimensionality of the imaging domain. One showed how to incorporate prior information about illumination and reflectance in the process of constructing appearance bases. The presented approach can be used in graphics as well as vision where the presented theory has lead to a new model-based shape recovery where the hitherto assumed Lambertian assumption is relaxed. One believes that there is still much progress to be made in the future.

REFERENCES

- [1] L. Wolff, S. Shafer, and G. Healy. *Physics-based Vision: Principles and Practice: Shape Recovery*. Jones and Bartlett Publishers, 1992. [3](#), [114](#)
- [2] B. Horn. *Shape from Shading, A Method for Obtaining the Shape of a Smooth Opaque Object from One View*. PhD thesis, Massachusetts Inst. of Technology, Cambridge, Massachusetts, 1970. [3](#), [10](#), [14](#), [117](#)
- [3] C. Liang, Y. Shih, and H. Chen. Light field analysis for modeling image formation. *IEEE Transactions on Image Processing*, 20(2):446 – 460, Feb. 2011. [3](#), [24](#)
- [4] B. Cabral, N. Max, and R. Springmeyer. Bidirectional reflection functions from surface bump maps. *SIGGRAPH Computer Graphics*, 21(4):273–281, August 1987. [4](#), [67](#), [68](#), [70](#), [116](#), [117](#), [136](#), [219](#)
- [5] M. D’Zmura. Shading ambiguity: reflectance and illumination. *In Computational Models of Visual Processing*, pages 187–207, 1991. [4](#), [116](#), [117](#), [136](#), [219](#)
- [6] R. Basri and D. W. Jacobs. Lambertian reflectance and linear subspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(2):218–233, 2003. [4](#), [8](#), [9](#), [11](#), [13](#), [14](#), [15](#), [17](#), [35](#), [65](#), [116](#), [117](#), [118](#), [119](#), [120](#), [136](#), [139](#), [161](#), [164](#), [183](#), [186](#), [200](#), [219](#), [221](#), [235](#)
- [7] R. Ramamoorthi and P. Hanrahan. On the relationship between radiance and irradiance: determining the illumination from images of a convex lambertian object. *Journal of the Optical Society of America A*, 18(10):2448–2459, 2001. [4](#), [8](#), [9](#), [11](#), [14](#), [15](#), [17](#), [35](#), [95](#), [96](#), [117](#), [118](#), [119](#), [120](#), [121](#), [136](#), [149](#), [153](#), [157](#), [183](#), [184](#), [219](#)
- [8] D. Forsyth and J. Ponce. *Computer Vision: A Modern Approach*. Prentice Hall, 2nd edition, 2011. [4](#), [29](#), [30](#), [34](#), [39](#), [40](#), [43](#), [104](#), [191](#), [193](#)
- [9] F. Nicodemus, J. Richmond, J. Hsia, I. Ginsberg, and T. Limperis. *Geometrical Considerations and Nomenclature for Reflectance*. US Dept. of Commerce, National Bureau of Standards: for sale by the Supt. of Docs., US Govt. Print. Off., 1977. [4](#), [31](#), [32](#)
- [10] K. Dana, B. Ginneken, S. Nayar, and J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, January. [5](#), [10](#), [33](#), [99](#), [102](#), [103](#), [125](#), [126](#), [127](#), [128](#), [131](#), [133](#), [145](#)

-
- [11] R. Ramamoorthi. *A Signal-Processing Framework for Forward and Inverse Rendering*. Ph.D. Dissertation. Stanford University, Stanford, CA, USA, Advisor(s) P. Hanrahan, 2002. [5](#), [11](#), [13](#), [35](#), [36](#), [42](#), [53](#), [54](#), [55](#), [65](#), [75](#), [76](#), [116](#), [117](#), [118](#), [119](#), [120](#), [126](#), [127](#), [133](#), [139](#), [154](#), [158](#), [185](#), [186](#), [235](#), [272](#)
- [12] K. Torrance and E. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of American Optical Society*, 57(9):1105–1114, 1967. [6](#), [66](#), [86](#), [90](#), [94](#), [184](#), [211](#)
- [13] R. Cook and K. Torrance. A reflectance model for computer graphics. *ACM Transaction of Graphics*, 1(1):7–24, 1982. [6](#), [95](#), [107](#), [209](#), [213](#)
- [14] M. Oren and S. Nayar. Generalization of lambert’s reflectance model. In *Computer Graphics, 28(Annual Conference Series)*, pages 239–246, 1994. [6](#), [66](#), [94](#), [107](#), [210](#), [212](#), [213](#)
- [15] B. Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, June 1975. [6](#), [66](#), [86](#), [91](#)
- [16] G. Ward. Measuring and modeling anisotropic reflection. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques, SIGGRAPH’92*, pages 265–272, 1992. [6](#)
- [17] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. *ACM Transactions on Graphics*, 22(3):759–769, July 2003. [6](#), [10](#), [67](#), [69](#), [99](#), [100](#), [101](#), [125](#), [127](#), [128](#), [129](#), [130](#), [132](#), [139](#), [145](#), [149](#), [154](#), [155](#), [156](#), [158](#), [159](#), [160](#), [188](#)
- [18] Y. Sato, M. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques, SIGGRAPH ’97*, pages 379–387, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co. [6](#)
- [19] S. Westin, J. Arvo, and K. Torrance. Predicting reflectance functions from complex surfaces. In *Proc. of SIGGRAPH, ACM Press*, pages 255–264, 1992. [6](#), [66](#), [68](#), [70](#), [82](#), [83](#), [97](#), [98](#), [99](#), [100](#), [102](#), [213](#), [214](#)
- [20] P. Gautron, J. Krivanek, S.N. Pattanaik, and K. Bouatouch. A novel hemispherical basis for accurate and efficient rendering. In *Proceedings of the Fifteenth Eurographics conference on Rendering Techniques*, pages 321–330. Eurographics Association, 2004. [6](#), [15](#), [17](#), [68](#), [69](#), [71](#), [72](#), [82](#), [83](#), [96](#), [97](#), [99](#), [100](#), [102](#), [116](#), [118](#), [119](#), [213](#), [214](#), [260](#), [268](#)
- [21] J. Koenderink and A. van Doorn. Phenomenological description of bidirectional surface reflection. *Journal of the Optical Society of America*, 15(11):2903–2912, 1998. [6](#), [15](#), [33](#), [67](#), [69](#), [70](#), [71](#), [75](#), [78](#), [82](#), [83](#), [88](#), [91](#), [92](#), [93](#), [97](#), [98](#), [99](#), [100](#), [102](#), [103](#), [108](#), [110](#), [118](#), [213](#), [214](#)

- [22] R. Habel and M. Wimmer. Efficient irradiance normal mapping. In *Proceedings of the 2010 ACM SIGGRAPH symposium on Interactive 3D Graphics and Games*, pages 189–195, 2010. [6](#), [71](#), [82](#), [83](#), [97](#), [99](#), [100](#), [102](#), [213](#), [214](#)
- [23] P. Belhumeur and D. Kriegman. What is the set of images of an object under all possible illumination conditions? *International Journal of Computer Vision*, 28(3):245–260, July 1998. [7](#), [9](#), [14](#), [135](#), [139](#)
- [24] B. Horn. *Robot Vision*. McGraw-Hill, 1986. [7](#), [24](#), [33](#), [120](#), [136](#), [181](#), [183](#), [219](#), [235](#)
- [25] R. Ramamoorthi. Analytic pca construction for theoretical analysis of lighting variability in images of a lambertian object. *IEEE Transaction of Pattern Analysis and Machine Intelligence*, 24(10):1322–1333, October 2002. [7](#), [9](#), [11](#), [15](#), [17](#), [35](#), [115](#), [116](#), [118](#), [119](#), [135](#), [136](#), [139](#), [219](#), [235](#)
- [26] A. Georghiadis, P. Belhumeur, and D. Kriegman. From few to many: illumination cone models for face recognition under variable lighting and pose. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(6):643–660, jun 2001. [9](#), [11](#), [14](#), [17](#)
- [27] A. Yuille, D. Snow, R. Epstein, and P. Belhumeur. Determining generative models of objects under varying illumination: Shape and albedo from multiple images using svd and integrability. *International Journal of Computer Vision*, 35(3):203–222, December 1999. [9](#), [11](#), [115](#), [135](#), [235](#)
- [28] R. Epstein, P.W. Hallinan, and A.L. Yuille. 5 plus or minus 2 eigenimages suffice: An empirical investigation of low-dimensional lighting models. In *IEEE Workshop on Physics-Based Modeling in Computer Vision*, pages 108–116, June 1995. [9](#)
- [29] K-C. Lee, J. Ho, and D. Kriegman. Acquiring linear subspaces for face recognition under variable lighting. *IEEE Transaction of Pattern Analysis and Machine Intelligence*, 27(5):684–698, May 2005. [9](#), [12](#), [14](#), [177](#), [192](#), [193](#)
- [30] M. Pharr and G. Humphreys. *Physically Based Rendering: From Theory to Implementation, 2nd edition*. 2nd edition, 2010. [10](#), [40](#), [42](#), [43](#), [45](#), [46](#), [52](#), [60](#), [128](#), [129](#), [130](#), [289](#)
- [31] R. Ramamoorthi and P. Hanrahan. Frequency space environment map rendering. *ACM Transactions on Graphics*, 21(3):517–526, July 2002. [11](#), [15](#), [17](#), [119](#), [126](#), [127](#), [129](#), [130](#), [133](#), [149](#), [154](#), [158](#), [164](#)
- [32] P. Nillius. *Image Analysis using the Physics of Light Scattering*. Doctoral Thesis, KTH Numerical Analysis and Computer Science, Stockholm, Sweden, 2004. [11](#), [13](#), [15](#), [17](#), [37](#), [65](#), [116](#), [117](#), [118](#), [119](#), [124](#), [125](#), [126](#), [127](#), [133](#), [139](#), [155](#), [159](#), [235](#), [269](#), [274](#), [276](#)
- [33] H. Rara, S. Elhabian, T. Starr, and A. Farag. Model-based shape recovery from single images of general and unknown lighting. In *16th IEEE International Conference on Image Processing (ICIP2009)*, pages 517–520, 7-10 Nov. 2009. [13](#), [65](#), [139](#), [184](#), [200](#), [201](#), [225](#), [235](#)

- [34] H. Rara, S. Elhabian, T. Starr, and A. Farag. 3d face recovery from intensities of general and unknown lighting using partial least squares. In *17th IEEE International Conference on Image Processing (ICIP2010)*, Sept. 2010. [13](#), [65](#), [139](#), [162](#), [163](#), [164](#), [171](#), [184](#), [200](#), [201](#), [202](#), [225](#), [235](#)
- [35] M. Lee and C. Choi. Fast facial shape recovery from a single image with general, unknown lighting by using tensor representation. *Pattern Recognition.*, 44(7):1487–1496, July 2011. [13](#), [139](#), [235](#)
- [36] I. Kemelmacher-Shlizerman and R. Basri. 3d face reconstruction from a single image using a single reference face shape. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 33(2):394–405, 2010. [13](#), [65](#), [139](#), [235](#)
- [37] R. Basri, D. Jacobs, and I. Kemelmacher. Photometric stereo with general, unknown lighting. *International Journal of Computer Vision*, 72(3):239–257, 2007. [13](#), [65](#), [139](#), [152](#), [154](#), [155](#), [156](#), [158](#), [159](#), [160](#), [184](#), [201](#), [235](#)
- [38] N. Aldrin, T. Zickler, and D. Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8, june 2008. [13](#), [65](#), [139](#), [235](#)
- [39] P. Nillius, J. Sullivan, and A. Argyros. Shading models for illumination and reflectance invariant shape detectors. In *In Proc. IEEE Computer Vision and Pattern Recognition (CVPR08)*, Anchorage, June 2008. [13](#), [65](#), [139](#), [235](#)
- [40] S. Elhabian, H. Rara, and A. Farag. Modeling lambertian surfaces under unknown distant illumination using hemispherical harmonics. In *Eighth Canadian Conference on Computer and Robot Vision (CRV)*, St. John's, Newfoundland, Canada, pages 293–300, 2011. [15](#), [17](#), [18](#), [118](#), [119](#)
- [41] S. Elhabian, H. Rara, and A. Farag. Towards accurate and efficient representation of image irradiance of convex-lambertian objects under unknown near lighting. In *International Conference of Computer Vision (ICCV)*, Barcelona, Spain, November 2011. [15](#), [17](#), [18](#), [118](#), [119](#)
- [42] C. Chong, P. Raveendran, and R. Mukundan. A comparative analysis of algorithms for fast computation of zernike moments. *Pattern Recognition*, 36(3):731–742, 2003. [15](#), [118](#)
- [43] Jr. D. Healy, P. Kostelec, and D. Rockmore. Towards safe and effective high-order legendre transforms with applications to ffts for the 2-sphere. *Advances in Computational Mathematics*, 21(2):59–105, 2004. [16](#), [83](#), [120](#)
- [44] P. Nillius and J. Eklundh. Phenomenological eigenfunctions for image irradiance. In *Proceedings of the Ninth IEEE International Conference on Computer Vision - Volume 2, ICCV '03*, pages 568–575, Washington, DC, USA, 2003. IEEE Computer Society. [17](#), [116](#), [119](#), [120](#), [121](#), [125](#), [126](#), [127](#), [129](#), [130](#), [133](#), [136](#), [139](#), [140](#), [144](#), [145](#), [146](#), [149](#), [153](#), [155](#), [157](#), [159](#), [164](#), [219](#), [222](#)

- [45] S. Elhabian, H. Rara, and A. Farag. Towards efficient and compact phenomenological representation of arbitrary bidirectional surface reflectance. In *Proceedings of the British Machine Vision Conference (BMVC)*, pages 89.1–89.11. BMVA Press, 2011. [18](#), [121](#), [185](#)
- [46] S. Elhabian and A. Farag. Modeling image irradiance under natural illumination and isotropic surface reflectance. In *International Conference of Image Processing (ICIP)*, 2012. [18](#)
- [47] S. Elhabian, H. Rara, A. Ali, and A. Farag. Illumination-invariant statistical shape recovery with contiguous occlusion. In *Eighth Canadian Conference on Computer and Robot Vision (CRV)*, pages 301–308, 2011. [19](#), [166](#), [185](#), [192](#)
- [48] S. Elhabian, E. Mostafa, H. Rara, A. Ali, and A. Farag. Non-lambertian model-based facial shape recovery from single image under unknown general illumination. In *Ninth Canadian Conference on Computer and Robot Vision (CRV)*, pages 252–259, 2012. [19](#)
- [49] A. Farag, S. Elhabian, A. Abdelrahim, W. Aboelmaaty, A. Farman, and D. Tazman. Model-based human teeth shape recovery from a single optical image with unknown illumination. In *MICCAI Medical Computer Vision Workshop (MCV)*, 2012. [20](#), [202](#), [203](#), [230](#), [231](#), [232](#)
- [50] B. Saleh and M. Teich. *Fundamentals of Photonics*. New York: John Wiley & Sons, 1991. [25](#)
- [51] M. Minnaert. The reciprocity principle in lunar photometry. *Journal of Astrophys*, 93:403410, 1941. [32](#)
- [52] H. von Helmholtz. Treatise on physiological optics. *Dover, New York*, I:231, 1962. [32](#), [67](#), [116](#)
- [53] J. Koenderink, A. Van Doorn, and M. Stavridi. Bidirectional reflection distribution function expressed in terms of surface scattering modes. In *European Conference of Computer Vision (ECCV)*, pages 28–39, 1996. [33](#), [70](#), [71](#)
- [54] J. Craig. *Introduction to Robotics*. Addison-Wesley, 1989. [37](#)
- [55] P. Debevec. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques, SIGGRAPH '98*, pages 189–198, New York, NY, USA, 1998. ACM. [44](#), [47](#), [60](#), [61](#), [128](#), [129](#), [130](#), [145](#), [149](#), [152](#), [154](#), [155](#), [156](#), [158](#), [159](#), [160](#), [221](#), [222](#), [223](#)
- [56] M. Abramowitz and I. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. New York: Dover, 1965. Chapter 8, pp. 332. [47](#), [267](#)

- [57] O. Makhotkin. Analysis of radiative transfer between surfaces by hemispherical harmonics. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 56(6):869–879, 1996. [48](#), [70](#), [71](#)
- [58] C. Chisholm. *Group theoretical techniques in quantum chemistry*. Academic Press, London, New York, 1976. [48](#)
- [59] E. Wigner (Author) and J. Griffin (Translator). *Group Theory and Its Application to the Quantum Mechanics of Atomic Spectra*. Academic Press; Expanded & improved edition, 1959. [52](#), [121](#), [185](#), [220](#), [272](#)
- [60] M. Blanco, M. Flrez, and M. Bermejo. Evaluation of the rotation matrices in the basis of real spherical harmonics. *Journal of Molecular Structure: THEOCHEM*, 419(1-3):19 – 27, 1997. [53](#), [54](#), [262](#)
- [61] Z. Gimbutas and L. Greengard. A fast and stable method for rotating spherical harmonic expansions. *Journal of Computational Physics*, 228(16):5621 – 5627, 2009. [53](#), [146](#), [222](#)
- [62] J. Ivanic and K. Ruedenberg. Rotation matrices for real spherical harmonics. direct determination by recursion. *Journal of Physical Chemistry*, 100(15):6342–6347, 1996. [53](#)
- [63] C. Lessig, T. de Witt, and E. Fiume. Efficient and accurate rotation of finite spherical harmonics expansions. *Journal of Computational Physics*, 231(2):243 – 250, 2012. [53](#)
- [64] L. Franks. *Signal Theory*. Prentice-Hall, 1969. [55](#), [123](#)
- [65] D. Frolova, D. Simakov, and R. Basri. Accuracy of spherical harmonic approximations for images of lambertian objects under far and near lighting. In *European Conference on Computer Vision, Zofin Palace, Slovansky ostrov, Prague 1, Czech Republic*, pages 574–587, 2004. [59](#), [200](#)
- [66] R. Montes and C. Urena. An overview of brdf models. Technical report, Dept.Lenguajes y Sistemas Informaticos, University of Granada, Granada, Spain, 2012. [66](#), [67](#), [68](#)
- [67] M. Ashikhmin and S. Premoz. Distribution-based brdfs. *Energy*, 2006. [66](#), [67](#), [95](#)
- [68] K. Nishino and S. Lombardi. Directional statistics-based reflectance model for isotropic bidirectional reflectance distribution functions. *J. Opt. Soc. Am. A*, 28(1):8–18, Jan 2011. [66](#)
- [69] J. Lambert. *Photometria sive de mensura de gratibus luminis colorum et umbrae*. Eberhard Klett, 1760. [66](#)
- [70] J. Blinn. Models of light reflection for computer synthesized pictures. *SIGGRAPH Computer Graphics*, 11(2):192–198, July 1977. [66](#), [86](#), [209](#)

- [71] E. Veach. *Robust monte carlo methods for light transport simulation*. PhD thesis, 1997. [67](#)
- [72] D. Edwards, S. Boulos, J. Johnson, P. Shirley, M. Ashikhmin, M. Stark, and C. Wyman. The halfway vector disk for brdf modeling. *ACM Transactions on Graphics*, 25(1):1–18, January 2006. [67](#)
- [73] R. Lewis. Making shaders more physically plausible. *Computer Graphics Forum*, 13(2):109–120, 1994. [67](#)
- [74] J. Kajiya and B. Von Herzen. Ray tracing volume densities. *SIGGRAPH Computer Graphics*, 18(3):165–174, January 1984. [68](#), [69](#)
- [75] J. Lu and J. Little. Reflectance and shape from images using a collinear light source. *International Journal Computer Vision*, 32(3):213–240, August 1999. [68](#)
- [76] T. Zickler, P. Belhumeur, and D. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *International Journal of Computer Vision (2002)*, 49(2):215–227, 2002. [68](#)
- [77] T. Hawkins, P. Einarsson, and P. Debevec. A dual light stage. In *Proceedings of the Eurographics Symposium on Rendering Techniques, Konstanz, Germany*, pages 91–98. Eurographics Association, June 29 - July 1, 2005. [68](#)
- [78] P. Sen, B. Chen, G. Garg, S. Marschner, M. Horowitz, M. Levoy, and H. Lensch. Dual photography. In *ACM SIGGRAPH 2005 Papers*, SIGGRAPH '05, pages 745–755, New York, NY, USA, 2005. ACM. [68](#)
- [79] W. Barros, R. Carceroni, L. Coelho, and J. Queiroz-Neto. An opto-mechanical apparatus for binocular helmholtz stereopsis in static and dynamic scenes. In *Brazilian Symposium on Computer Graphics and Image Processing (SIBGRAPI'07)*, pages 213–220, 2007. [68](#)
- [80] P. Tan, S. Mallick, D. Kriegman, L. Quan, and T. Zickler. Isotropy, reciprocity and the gbr ambiguity. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2007. [68](#)
- [81] J. Kautz, P-P. Sloan, and J. Snyder. Fast, arbitrary brdf shading for low-frequency lighting using spherical harmonics. In *Proceedings of the 13th Eurographics workshop on Rendering, EGRW '02*, pages 291–296, Aire-la-Ville, Switzerland, Switzerland, 2002. Eurographics Association. [68](#), [70](#), [83](#)
- [82] A. Bhatia and E. Wolf. On the circle polynomials of zernike and related orthogonal sets. In *Mathematical Proceedings of the Cambridge Philosophical Society*, number 50. [69](#), [117](#), [150](#)
- [83] F. Sillion, J. Arvo, S. Westin, and D. Greenberg. A global illumination solution for general reflectance distributions. *SIGGRAPH Computer Graphics*, 25(4):187–196, July 1991. [70](#), [75](#), [84](#), [86](#)

-
- [84] P. Sloan, J. Hall, J. Hart, and J. Snyder. Clustered principal components for precomputed radiance transfer. *ACM Transactions of Graphics*, 22(3):382–391, 2003. [70](#)
- [85] J. Wyant and K. Creath. Basic wavefront aberration theory for optical metrology. In *Applied optics and Optical Engineering*, Academic Press, Inc. [70](#), [71](#)
- [86] G. Szego. *Orthogonal polynomials*. American Mathematical Society, Providence, Rhode Island, Fourth Edition, 1975. pp. 29. [71](#)
- [87] Jr. D. Healy, P. Kostelec, D. Rockmore, and S. Moore. Ffts for the 2-sphere-improvements and variations. *Journal of Fourier Analysis and Applications*, 9(4):341–385, 2003. [83](#)
- [88] A. Mohammed and Y. Jie. Practical fast computation of zernike moments. *Journal of Computer Science and Technology*, 17(2):181–188, March 2002. [83](#)
- [89] R. Suda and M. Takami. A fast spherical harmonics transform algorithm. *Journal Mathematics of Computation*, 71(238):703–715, April 2002. [83](#)
- [90] S. Belkasim, M. Ahmadi, and M. Shridhar. Efficient algorithm for fast computation of zernike moments. *Journal of the Franklin Institute*, 333(4):577 – 581, 1996. [84](#)
- [91] E. Kintner. On the mathematical properties of the zernike polynomials. *Optica Acta: International Journal of Optics*, 23(8):679–680, 1976. [84](#)
- [92] Jr. A. Prata and W. Rusch. Algorithm for computation of zernike polynomials expansion coefficients. *Applied Optics*, 28(4):749–754, Feb 1989. [84](#)
- [93] B. Horn and R. Sjoberg. Calculating the reflectance map. *Applied Optics*, 18:1770 – 1779, 1979. [88](#)
- [94] M. Ashikmin, S. Premoze, and P. Shirley. A microfacet-based brdf generator. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '00, pages 65–74, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co. [95](#), [209](#)
- [95] T. Weyrich, W. Matusik, H. Pfister, B. Bickel, C. Donner, C. Tu, J. McAndless, J. Lee, A. Ngan, H. Jensen, and M. Gross. Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics*, 25(3):1013–1024, July 2006. [97](#), [98](#), [99](#), [108](#), [184](#), [185](#), [189](#), [192](#)
- [96] S. Rusinkiewicz. A new change of variables for efficient BRDF representation. In *Rendering Techniques (Proc. Eurographics Workshop on Rendering)*, June 1998. [99](#), [109](#)
- [97] E. Eibenberger and E. Angelopoulou. Beyond the neutral interface reflection assumption in illuminant color estimation. In *17th IEEE International Conference on Image Processing (ICIP)*, number 50. [108](#)

- [98] J. Mulder. Closed-form spherical harmonics: Explicit polynomial expression for the associated legendre functions. *Journal of Chemical Education*, 7(2):244, Feb 2000. [122](#), [269](#)
- [99] F. Olver, D. Lozier, R. Boisvert, and C. Clark. *The NIST Handbook of Mathematical Functions*. Cambridge University Press, New York, NY, May, 2010. Ch5, pp. 142. [122](#), [262](#), [271](#)
- [100] [128](#), [129](#), [130](#)
- [101] Y. Moses, Y. Adini, and S. Ullman. Face recognition: The problem of compensating for changes in illumination direction. In Jan-Olof Eklundh, editor, *Computer Vision ECCV '94*, volume 800 of *Lecture Notes in Computer Science*, pages 286–296. Springer Berlin / Heidelberg, 1994. [135](#)
- [102] W. Zuo, D. Zhang, and K. Wang. Bidirectional pca with assembled matrix distance metric for image recognition. *Transactions on Systems, Man and Cybernetics - Part B*, 36(4):863–872, August 2006. [136](#), [141](#), [142](#), [146](#), [149](#)
- [103] H. Lu, K. Plataniotis, and A. Venetsanopoulos. A survey of multilinear subspace learning for tensor data. *Pattern Recognition*, 44(7):1540–1551, July 2011. [138](#)
- [104] M. Vasilescu and D. Terzopoulos. Tensortextures: multilinear image-based rendering. In *ACM SIGGRAPH 2004 Papers*, SIGGRAPH '04, pages 336–342, New York, NY, USA, 2004. ACM. [138](#)
- [105] Smart Image based Lighting. <http://www.hdrlabs.com/sibl/archive.html>. [139](#), [145](#), [149](#), [221](#), [222](#), [223](#)
- [106] H. Lu, K. Plataniotis, and A. Venetsanopoulos. MPCA: Multilinear principal component analysis of tensor objects. *IEEE Transactions on Neural Networks*, 19(1):18–39, jan. 2008. [140](#), [141](#), [144](#), [150](#)
- [107] P. Kroonenberg and J. de Leeuw. Principal component analysis of three-mode data by means of alternating least squares algorithms. *Psychometrika*, 45(1):69–97, 1980. [143](#), [279](#)
- [108] G. Kolda and B. Bader. Tensor decompositions and applications. *SIAM Rev.*, 51(3):455–500, August 2009. [147](#)
- [109] B. J. Broxson. The kronecker product. Master of science in mathematics, University of North Florida, 2006. [147](#)
- [110] J. Ye. Generalized low rank approximations of matrices. In *Proceedings of the twenty-first international conference on Machine learning*, ICML '04, pages 112–137, New York, NY, USA, 2004. ACM. [147](#)
- [111] T. Okabe, I. Sato, and Y. Sato. Spherical harmonics vs. haar wavelets: Basis for recovering illumination from cast shadows. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004. Washington, DC, USA*, volume 1, pages 50–57. IEEE, 2004. [162](#)

-
- [112] Z. Wen, Z. Liu, and T. Huang. Face relighting with radiance environment maps. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, pages 58–65, june 2003. [162](#)
- [113] X. Mei, H. Ling, and D. Jacobs. Illumination recovery from image with cast shadows via sparse representation. *Transactions on Image Processing*, 20(8):2366–2377, August 2011. [162](#)
- [114] S. Yuen, Y. Tsui, and C. Chow. A fast marching formulation of perspective shape from shading under frontal illumination. *Pattern Recognition Letters*, 28(7):806 – 824, 2007. [162](#)
- [115] Z. Zhou, A. Wagner, H. Mobahi, J. Wright, and M. Yi. Face recognition with contiguous occlusion using markov random fields. In *IEEE 12th International Conference on Computer Vision.*, pages 1050 –1057, 29 2009-oct. 2 2009. [162](#), [163](#), [168](#), [170](#), [171](#)
- [116] F. Hampel, E. Ronchetti, P. Rousseeuw, and W. Stahel. *Robust Statistics: The Approach Based on Influence Functions*. Wiley-Interscience, New York, April 2005. [163](#), [166](#)
- [117] Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(9):1124 –1137, sept. 2004. [163](#), [169](#)
- [118] S. Sarkar. Usf humanid 3d face database. University of South Florida, Tampa, FL, USA. [163](#), [165](#), [174](#), [191](#)
- [119] P. Belhumeur and D. Kriegman. What is the set of images of an object under all possible lighting conditions? In *IEEE Conference of Computer Vision and Pattern Recognition*, pages 270–277, 1996. [164](#)
- [120] I. Matthews and S. Baker. Active appearance models revisited. *International Journal of Computer Vision*, 60(2):135–164, November 2004. [165](#), [171](#)
- [121] M. Black and P. Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer Vision and Image Understanding*, 63(1):75 – 104, 1996. [167](#)
- [122] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6):721–741, November 1984. [169](#)
- [123] A. Ali, A. Farag, and G. Gimel’Farb. Optimizing binary mrfs with higher order cliques. In *Proceedings of the 10th European Conference on Computer Vision: Part III, ECCV ’08*, pages 98–111, Berlin, Heidelberg, 2008. Springer-Verlag. [169](#), [170](#), [171](#)

-
- [124] A. Dempster, N. Laird, and D. Rdin. Maximum likelihood from incomplete data via the em-algorithm. In *Journal Of The Royal Statistical Society, Series B*, volume 39, pages 1–38, 1977. [170](#), [215](#)
- [125] W. Zhao and R. Chellappa. Symmetric shape-from-shading using self-ratio image. *International Journal on Computer Vision*, 45(1):55–75, 2001. [183](#), [200](#)
- [126] E. Prados and O. Faugeras. Unifying approaches and removing unrealistic assumptions in shape from shading: Mathematics can help. In *European Conference on Computer Vision ECCV04*, Prague, Czech Republic, May 2004. [183](#), [199](#)
- [127] A. Ahmed, A. Farag, and T. Starr. A new symmetric shape from shading algorithm with an application to 3-d face reconstruction. In *Proc. of IEEE International Conference on Image Processing (ICIP08)*, San Diego, California, pages 201–204, October 12-15, 2008. [183](#)
- [128] C. Kao, S. Osher, and J. Qian. Lax-friedrichs sweeping scheme for static hamilton-jacobi equations. *Journal of Computational Physics*, 196(1):367–391, May 2004. [183](#)
- [129] J. Atick, P. Griffin, and A. Redlich. Statistical approach to shape from shading: Reconstruction of three-dimensional face surfaces from single two-dimensional images. *Neural Computing*, 8(6):1321–1340, 1996. [183](#), [200](#)
- [130] R. Dovgard and R. Basri. Statistical Symmetric Shape from Shading for 3D Structure Recovery of Faces. In *Proceedings of the 8th European Conference on Computer Vision (volume II)*, volume 3022 of *Lecture Notes in Computer Science*, pages 99–113, Prague, Czech Republic, May 2004. [183](#), [200](#)
- [131] W. Smith and E. Hancock. Recovering facial shape using a statistical model of surface normal direction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(12):1914–1930, 2006. [183](#), [200](#)
- [132] V. Blanz and T. Vetter. Face recognition based on fitting a 3d morphable model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(9):1063–1074, 2003. [183](#), [190](#), [200](#)
- [133] T. Cootes, G. Edwards, and C. Taylor. Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23:681–685, 2001. [183](#), [200](#)
- [134] M. Castelan, W. Smith, and E. Hancock. A coupled statistical model for face shape recovery from brightness images. *IEEE Transaction on Image Processing*, 16(1):1139–1151, 2007. [183](#), [184](#), [191](#), [200](#), [201](#)
- [135] L. Zhang and D. Samaras. Face recognition from a single training image under arbitrary unknown lighting using spherical harmonics. *IEEE Transaction on Pattern Analysis Machine Intelligence*, 28(3):351–363, 2006. [184](#), [200](#)

- [136] A. Ahmed and A. Farag. A new statistical model combining shape and spherical harmonics illumination for face reconstruction. In *Proc. ISVC*, pages 531–541, 2007. [184](#), [200](#), [201](#)
- [137] H. Rara, S. Elhabian, T. Starr, and A. Farag. A statistical model combining shape and spherical harmonics for face reconstruction and recognition. In *Biomedical Engineering Conference, CIBEC2008. Cairo International*, pages 1–4, 18-20 Dec. 2008. [184](#), [201](#)
- [138] M. Castelan and J. Van Horebeek. Relating intensities with three-dimensional facial shape using partial least squares. *Computer Vision, IET*, 3(2):60–73, June 2009. [184](#), [200](#), [201](#), [225](#)
- [139] A. Georghiadis. Recovering 3-d shape and reflectance from a small number of photographs. In *Proceedings of the 14th Eurographics workshop on Rendering*, EGRW '03, pages 230–240, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association. [184](#)
- [140] S. Marschner, S. Westin, E. Lafortune, K. Torrance, and D. Greenberg. Image-based brdf measurement including human skin. In Dani Lischinski and Gregory Ward Larson, editors, *Rendering Techniques 99, Proceedings of the Eurographics Workshop in Granada, Spain, June 21-23, 1999*, pages 131–144. Springer, 1999. [184](#)
- [141] D. Goldman, B. Curless, A. Hertzmann, and S. Seitz. Shape and spatially-varying brdfs from photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(6):1060–1071, 2010. [188](#)
- [142] P. Belhumeur, D. Kriegman, and A. Yuille. The bas-relief ambiguity. *International Journal of Computer Vision*, 35(1):33–44, November 1999. [189](#)
- [143] A. Patel and W. Smith. 3d morphable face models revisited. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 0:1327–1334, 2009. [190](#), [191](#), [216](#)
- [144] E. Angel. *Interactive Computer Graphics: A Top-Down Approach with OpenGL*. Addison Wesley Longman, Inc., 2000. [191](#)
- [145] M. Stegmann. Active appearance models: Theory, extensions and cases. Master's thesis, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, August 2000. [191](#)
- [146] T. Cootes and C. Taylor. Statistical models of appearance for computer vision. Technical report, Imaging Science and Biomedical Engineering, University of Manchester, Manchester M13 9PT, U.K., March 2004. [192](#), [202](#), [216](#), [230](#)
- [147] D. Ning, C. Sheng, L. He, Y. Qing, and W. Yong. Deformation design technology of dental restoration model. In *International Conference on BioMedical Engineering and Informatics(BMEI)*., volume 2, pages 793 –797, may 2008. [198](#)

- [148] S-X. Zheng, J. Li, and Q-F. Sun. A novel 3d morphing approach for tooth occlusal surface reconstruction. *Computer Aided Design*, 43(3):293–302, March 2011. [198](#), [201](#), [229](#)
- [149] S. Buchaillard, S. Ong, Y. Payan, and K. Foong. 3d statistical models for tooth surface reconstruction. *Computers in Biology and Medicine*, 37(10):1461–1471, October 2007. [198](#), [201](#)
- [150] A. Goshtasby, S. Nambala, W. deRijk, and S. Campbell. A system for digital reconstruction of gypsum dental casts,. In *IEEE Transactions on Medical Imaging*, volume 16, pages 664–674, Oct 1997. [198](#)
- [151] European Commission. European guidelines on radiation protection in dental radiology. Radiation Protection Issue number 136, 2004. [199](#)
- [152] H. Maillie and J. Gilda. Radiation-induced cancer risk in radiographic cephalometry. *Oral Surg Oral Med Oral Pathol* 75: 631-7, 1993. [199](#)
- [153] P. Horn-Ross, B. Ljung, and M. Morrow. Environmental factors and the risk of salivary gland cancer. *Epidemiology* 8:414-9, 1997. [199](#)
- [154] A. Hallquis, L. Hardell, A. Degerman, G. Wingren, and L. Boquist. Medical diagnostic and therapeutic ionizing radiation and the risk for thyroid cancer: a case-control study. *Eur J Cancer Prevention* 3:259-67, 1994. [199](#)
- [155] K. Wilson, N. Sun, M. Huang, W. Zhang, A. Lee, A. Li, S. Wang, and J. Wu. Effects of ionizing radiation on self renewal and pluripotency of human embryonic stem cells. *Cancer Res.* 70(13): 55395548, 2010. [199](#)
- [156] European Commission. Low dose ionizing radiation and cancer risk. Radiation Protection Issue number 125, 2001. [199](#)
- [157] S. M. Yamany, A. Farag, D. Tasman, and A. Farman. A 3-d reconstruction system for the human jaw using a sequence of optical images. *IEEE Transactions on Medical Imaging*, 19:538–547, May 2000. [199](#)
- [158] A. Abdelrahim, M. Abdelrahman, H. Abdelmunim, A. Farag, and M. Miller. Novel image-based 3d reconstruction of the human jaw using shape from shading and feature descriptors. In *British Machine Vision Conference (BMVC)*, pages 41.1–41.11, 2011. [199](#)
- [159] M. Castelan and J. Horebeek. 3d face shape approximation from intensities using partial least squares. *Computer Vision and Pattern Recognition Workshop*, 0:1–8, 2008. [201](#)
- [160] S. Zheng, L. Jia, and S. Qingfeng. A space mapping deformation technology for dental surface restoration. In *3rd International Conference on Bioinformatics and Biomedical Engineering (ICBBE)*, pages 1–4, june 2009. [201](#)

-
- [161] J. Sporring and K. Jensen. Bayes reconstruction of missing teeth. *Journal of Mathematical Imaging and Vision*, 31(2-3):245–254, July 2008. [201](#), [229](#)
- [162] V. Blanz, A. Mehl, T. Vetter, and H. Seidel. A statistical method for robust 3d surface reconstruction from sparse data. In *In International Symposium on 3D Data Processing, Visualization and Transmission*, pages 293–300, 2004. [201](#), [202](#), [207](#)
- [163] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988. [201](#)
- [164] D. Williams and M. Shah. A fast algorithm for active contours and curvature estimation. *CVGIP: Image Understanding.*, 55(1):14–26, January 1992. [201](#)
- [165] P. Besl and H. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, feb 1992. [201](#)
- [166] S. Milborrow and F. Nicolls. Locating facial features with an extended active shape model. In *Proceedings of the 10th European Conference on Computer Vision: Part IV, ECCV '08*, pages 504–513, Berlin, Heidelberg, 2008. Springer-Verlag. [202](#), [206](#), [207](#), [229](#)
- [167] L. Wolff. Diffuse-reflectance model for smooth dielectric surfaces. *Journal of Optical Society fo America A*, 11(11):2956–2968, Nov 1994. [203](#), [212](#), [213](#)
- [168] R. Marino. Dental anatomy: Quick study academic, 2004. Barcharts Inc, illustrated by: Vincent Perez. [203](#)
- [169] J. Zijp and J. Rudolf. *Optical properties of dental hard tissues*. Dissertations, University of Groningen, 2001. [203](#)
- [170] J. Ragain and W. Johnston. Accuracy of kubelka-munk reflectance theory applied to human dentin and enamel. *Journal of Dental Research*, 80(2):449–452, 2001. [203](#)
- [171] C. Thode Larsen. *Modelling and Rendering of Teeth - A Study in Sub Surface Scattering*. Academic dissertation (Master thesis), Technical University of Denmark (DTU), Kgs. Lyngby, Denmark, 2011. [203](#)
- [172] J. Fuller, G. Denehy, and T. Schulein. *Concise Dental Anatomy and Morphology, 4th edition*. University of Iowa Pubns Dept, 2001. [205](#)
- [173] J. Gower. Generalized procrustes analysis. *Psychometrika*, 40(1):33–51, 1975. [207](#), [230](#)
- [174] S. Nayar., K. Ikeuchi, and T. Kanade. Surface reflection: physical and geometrical perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):611–634, jul 1991. [207](#)
- [175] J. Bennett and L. Mattsson. *Introduction to surface roughness and scattering*. Optical Society of America in Washington, D.C ., 1999. [208](#)

- [176] M. Oren and S. Nayar. Generalization of the lambertian model and implications for machine vision. *International Journal of Computer Vision*, 14(3):227 – 251, 1995. [211](#)
- [177] M. Oren and S. Nayar. Diffuse reflectance from rough surfaces. In *International Conference on Computer Vision and Pattern Recognition CVPR93*, New York,USA, 1993. [211](#)
- [178] M. Oren and S. Nayar. seeing beyond lambert’s law. In *European Conference on Computer Vision ECCV’1994*, Stockholm,Sweden, 1994. [211](#)
- [179] H. Ragheb and E. Hancock. Surface radiance correction for shape from shading. *Pattern Recognition*, 38:1574–1595, 2005. [212](#)
- [180] L. Wolff, S. Nayar, and M. Oren. Improved diffuse reflection models for computer vision. *International Journal of Computer Vision*, 30(1):55–71, October 1998. [213](#)
- [181] X-J. Wang, T. Milner, J. de Boer, Y. Zhang, D. Pashley, and J. Nelson. Characterization of dentin and enamel by use of optical coherence tomography. *Applied Optics*, 38(10):2092–2096, Apr 1999. [213](#), [214](#)
- [182] S. Valette, J. Chassery, and R. Prost. Generic remeshing of 3d triangular meshes with metric-dependent discrete voronoi diagrams. *IEEE Transactions on Visualization and Computer Graphics*, 14(2):369 –381, march-april 2008. [215](#)
- [183] F. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11:567–585, 1989. [216](#), [230](#)
- [184] R. Duda, P. Hart, and D. Stork. *Pattern Classification*. Wiley, New York, 2. edition, 2001. [216](#)
- [185] I. Jolliffe. *Principal Component Analysis*. Springer, second edition, October 2002. [216](#)
- [186] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical recipes in C (2nd ed.): the art of scientific computing*. Cambridge University Press, New York, NY, USA, 1992. [217](#)
- [187] J. Barron and J. Malik. Color constancy, intrinsic images, and shape estimation. In *12th European Conference on Computer Vision, Florence, Italy*, pages 57–70, October 7-13, 2012. [217](#), [218](#)
- [188] T. Miyazaki, Y. Hotta, J. Kunii, S. Kuriyama, and Y. Tamaki. A review of dental cad/cam: current status and future perspectives from 20 years of experience. *Dental Materials Journal*, 28(1):44–56, 2009. [228](#)
- [189] A. Tao, L. Wenhe, and L. Peijun. Computer aided design and manufacturing of the framework of pfm fixed bridge. In *7th International Conference on Computer-Aided Industrial Design and Conceptual Design (CAIDCD)*, pages 1 –5, nov. 2006. [228](#)

- [190] E. Rekow, D. R. Riley, A. G. Erdman, and B.E. B. E. Klamecki. Cad/cam for dental restorations-challenges and possibilities. In *Engineering in Medicine and Biology Society, 1989. Images of the Twenty-First Century., Proceedings of the Annual International Conference of the IEEE Engineering in*, pages 1642 –1643 vol.5, nov 1989. [228](#)
- [191] K. Hikita, Y. Uchiyama, K. Iiyama, and F. Duret. Function and clinical application of dental cad/cam "gn-1". *Institute of Medical Science, Health Sciences University of Hokkaido, Ainosato 2-5, Kita-ku, Sapporo, Japan*, 2002. [228](#)
- [192] T. Nakamura, N. Dei, T. Kojima, and K. Wakabayashi K. Marginal and internal fit of cerec 3 cad/cam all-ceramic crowns. *Division of Oromaxillofacial Regeneration, Osaka University Graduate School of Dentistry, Suita, Japan*, 2003. [228](#)
- [193] A. Mehl and V. Blanz and Hickel. Biogeneric tooth:a new mathematical representation for tooth morphology in lower first molars. pages 333–340, 2005. [229](#)
- [194] C. Carter, R. Pusateri, D. Chen, A. Ahmed, and A. Farag. Shape from shading for hybrid surfaces as applied to tooth reconstruction. In *International Conference on Image Processing (ICIP10)*, pages 4049 – 4052, Hong Kong, 2010. [229](#)
- [195] A. Ahmed and A. Farag. A new formulation for shape from shading for non-lambertian surfaces. In *International Conference on Computer Vision and Pattern Recognition CVPR06*, pages 1817–1824, NY, USA, 2006. [229](#)
- [196] A. Ahmed and A. Farag. Shape from shading under various imaging conditions. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR'07), Minneapolis, MN*, pages X1–X8, June 18-23, 2007. [229](#), [230](#), [232](#), [233](#), [235](#)
- [197] J. Durou, M. Falcone, and M. Sagona. Numerical methods for shape-from-shading: A new survey with benchmarks. *Computer Vision and Image Understanding*, 109(1):22 – 43, 2008. [229](#)
- [198] A. Ahmed and A. Farag. Shape from shading for hybrid surfaces. In *Proc. of IEEE International Conference on Image Processing (ICIP07), St. Antonio, Texas*, pages II–525–II–528, September 16-19, 2007. [235](#)
- [199] P. Kythe and M. Schaferkötter. *Handbook of Computational Methods for Integration*. Chapman & Hall/CRC, 2004. [259](#), [262](#)
- [200] G. Chirikjian and A. Kyatkin. *Engineering Applications of Non-commutative Harmonic Analysis*. CRC Press, 2001. pp. 296. [274](#)
- [201] I. Dokmanić and D. Petrinović. Convolution on the n-sphere with application to pdf modeling. *Transactions of Signal Processing*, 58(3):1157–1170, March 2010. [286](#)
- [202] J. Driscoll and D. Healy. Computing fourier transforms and convolutions on the 2-sphere. *Advances in Applied Mathematics*, 15(2):202 – 250, 1994. [286](#), [287](#)

- [203] W. Magnus, F. Oberhettinger, and R. Soni. *Formulas and Theorems for the Special Functions of Mathematical Physics*. Springer-Verlag; Enlarged 3rd edition, 1966. Chapter V, pp. 204. [288](#)
- [204] G. Arfken and H. Weber. *Mathematical Methods for Physicists: A Comprehensive Guide*. Elsevier Academic Press - Mathematics, Sixth Edition, 2005. Chapter 12. [291](#), [292](#)
- [205] W. Jarosz, N. Carr, and H. Jensen. Importance sampling spherical harmonics. *Computer Graphics Forum (Proceedings of Eurographics 2009)*, 28(2):577–586, April 2009. [291](#)
- [206] Robin Green. Spherical harmonic lighting: The gritty details. In *Archives of the Game Developers Conference*, volume 2, pages 2–3, 2003. [295](#)

APPENDIX A: THE POLAR INTEGRAL OF SHIFTED ASSOCIATED LEGENDRE POLYNOMIALS

In this appendix, the closed form of the following integral is derived,

$$\varsigma_n = \int_{\Omega} \tilde{\Theta}_n^0(\theta) d\omega \quad (\text{A-1})$$

where,

$$\tilde{\Theta}_n^m(\theta) = \tilde{N}_n^{|m|} \tilde{P}_n^{|m|}(\cos \theta) \quad (\text{A-2})$$

with a normalization factor defined as,

$$\tilde{N}_n^m = \sqrt{\frac{2n+1}{2\pi} \frac{(n-m)!}{(n+m)!}} \quad (\text{A-3})$$

and the shifted associated Legendre polynomials are given by,

$$\tilde{P}_n^m(x) = P_n^m(2x-1) \quad (\text{A-4})$$

Hence it can be shown that,

$$\begin{aligned} \varsigma_n &= \tilde{N}_n^0 \int_{\Omega} \tilde{P}_n^0(\cos \theta) d\omega \\ &= \tilde{N}_n^0 \int_0^{2\pi} d\phi \int_0^{\pi/2} \tilde{P}_n^0(\cos \theta) \sin \theta d\theta \\ &= \sqrt{2\pi(2n+1)} \int_0^{\pi/2} \tilde{P}_n^0(\cos \theta) \sin \theta d\theta \end{aligned} \quad (\text{A-5})$$

Switching to the cartesian representation, such that $z = \cos \theta$, thus $dz = -\sin \theta d\theta$,

while the integration domain changes from $[0, \pi/2]$ to $[1, 0]$,

$$\begin{aligned}\varsigma_n &= \sqrt{2\pi(2n+1)} \int_1^0 \tilde{P}_n^0(z)(-dz) \\ &= \sqrt{2\pi(2n+1)} \int_0^1 \tilde{P}_n(z)dz\end{aligned}\tag{A-6}$$

Using the series representation of the shifted Legendre polynomials [199],

$$\tilde{P}_n(z) = (-1)^n \sum_{k=0}^n (-1)^k \binom{n}{k} \binom{n+k}{k} z^k\tag{A-7}$$

Thus, the integration in Equation A-6 can be simplified as,

$$\begin{aligned}\int_0^1 \tilde{P}_n(z)dz &= \int_0^1 \left((-1)^n \sum_{k=0}^n (-1)^k \binom{n}{k} \binom{n+k}{k} z^k \right) dz \\ &= (-1)^n \sum_{k=0}^n (-1)^k \binom{n}{k} \binom{n+k}{k} \left(\int_0^1 z^k dz \right)\end{aligned}\tag{A-8}$$

where,

$$\begin{aligned}\int_0^1 z^k dz &= \left. \frac{z^{k+1}}{k+1} \right|_0^1 \\ &= \frac{1}{k+1}\end{aligned}\tag{A-9}$$

Therefore, it can be shown that,

$$\begin{aligned}\varsigma_n &= (-1)^n \sqrt{2\pi(2n+1)} \sum_{k=0}^n (-1)^k \binom{n}{k} \binom{n+k}{k} \frac{1}{k+1} \\ &= (-1)^n \sqrt{2\pi(2n+1)} \sum_{k=0}^n \frac{(-1)^k (n+k)!}{(k+1)(k!)^2(n-k)!}\end{aligned}\tag{A-10}$$

where the first few terms are¹,

$$\begin{aligned}\varsigma_0 &\approx 2.5066 \\ \varsigma_n &\approx 0 \quad \forall n \geq 1\end{aligned}\tag{A-11}$$

□

¹These values were validated using Monte Carlo integration.

APPENDIX B: TORRANCE SPARROW

SPECTRUM COEFFICIENTS

Since there is no azimuthal dependence in Equation 4.45, it can be expanded using zonal basis of the isotropic Helmholtz HSH-based basis \mathcal{H}_{pr}^q , centered at surface normal $\vec{n}(\mathbf{x})$, where q is set to zero. The zonal basis can be written as,

$$\begin{aligned} \mathcal{H}_{pr}^0(\vec{\omega}'_i, \vec{\omega}'_o) &= \frac{1}{2\pi} \sqrt{\frac{(2p+1)(2r+1)}{2+2\delta_{pr}}} \\ &\times \left[\tilde{P}_p(\cos \theta'_i) \tilde{P}_r(\cos \theta'_o) \right. \\ &\quad \left. + \tilde{P}_r(\cos \theta'_i) \tilde{P}_p(\cos \theta'_o) \right] \end{aligned} \quad (\text{B-1})$$

where \tilde{P}_p is shifted associated Legendre polynomials [20] with order $p \geq 0$. Thus TS-BRDF can be expanded in the subspace spanned by the Helmholtz HSH-based basis as,

$$f_r^{TS}(\vec{\omega}'_i, \vec{\omega}'_o) = \sum_{p=0}^{\infty} \sum_{r=0}^{\infty} a_{pr} \mathcal{H}_{pr}^0(\vec{\omega}'_i, \vec{\omega}'_o) \quad (\text{B-2})$$

where the spectrum coefficients can be obtained as¹,

$$\begin{aligned} a_{pr} &= \int_{\Omega'_o} \int_{\Omega'_i} f_r(\vec{\omega}'_i, \vec{\omega}'_o) \mathcal{H}_{pr}^0(\vec{\omega}'_i, \vec{\omega}'_o) d\vec{\omega}'_i d\vec{\omega}'_o \\ &= \frac{\rho_d}{\pi} d_{pr} + \frac{\rho_s}{\pi \sigma^2 \cos \theta'_o} \frac{F(\cos \theta'_o; \eta)}{F(1; \eta)} B_{pr}(\theta'_o, \sigma) \end{aligned} \quad (\text{B-3})$$

where d_{pr} and B_{pr} denote the diffuse and specular integrals of the TS spectrum coefficients, respectively.

¹Note that the complex-conjugate is discarded since the focus here is real-form basis.

The diffuse integral can be derived as follows;

$$\begin{aligned}
 d_{pr} &= \int_{\Omega'_o} \int_{\Omega'_i} \mathcal{H}_{pr}^0(\vec{\omega}'_i, \vec{\omega}'_o) d\vec{\omega}'_i d\vec{\omega}'_o \\
 &= \frac{1}{2\pi} \sqrt{\frac{(2p+1)(2r+1)}{2+2\delta_{pr}}} \left\{ \int_{\Omega'_i} \tilde{P}_p(\cos \theta'_i) d\vec{\omega}'_i \int_{\Omega'_o} \tilde{P}_r(\cos \theta'_o) d\vec{\omega}'_o \right. \\
 &\quad \left. + \int_{\Omega'_i} \tilde{P}_r(\cos \theta'_i) d\vec{\omega}'_i \int_{\Omega'_o} \tilde{P}_p(\cos \theta'_o) d\vec{\omega}'_o \right\} \quad (\text{B-4})
 \end{aligned}$$

Using variable substitution $z = \cos \theta \rightarrow dz = -\sin \theta d\theta$ and the orthogonality relation of the shifted-associated Legendre polynomial (with zero-degree),

$$\int_{\Omega} \tilde{P}_n(\cos \theta) d\vec{\omega} = 2\pi \int_0^1 \tilde{P}_n(z) dz = \frac{2\pi}{2n+1} \delta_{n0} \quad (\text{B-5})$$

Substituting in d_{pr} will lead to,

$$\begin{aligned}
 d_{pr} &= \frac{4\pi}{\sqrt{(2+2\delta_{pr})(2p+1)(2r+1)}} \delta_{p0} \delta_{r0} \\
 &= \begin{cases} 2\pi & p=0, r=0 \\ 0 & \text{otherwise} \end{cases} \quad (\text{B-6})
 \end{aligned}$$

Thus the diffuse/Lambertian component of the TS-BRDF is encoded in the first component. On the other hand, the specular integral B_{pr} can be written as,

$$\begin{aligned}
 B_{pr}(\theta'_o, \sigma) &= \int_{\Omega'_o} \int_{\Omega'_i} \overbrace{\frac{1}{\cos \theta'_i} \exp \left[-\left(\frac{\theta'_i}{2\sigma} \right)^2 \right]}^{g(\theta'_i; \sigma)} \mathcal{H}_{pr}^0(\vec{\omega}'_i, \vec{\omega}'_o) d\vec{\omega}'_i d\vec{\omega}'_o \\
 &= \frac{1}{2\pi} \sqrt{\frac{(2p+1)(2r+1)}{2+2\delta_{pr}}} \times [C_{pr}(\theta'_o, \sigma) + C_{rp}(\theta'_o, \sigma)] \quad (\text{B-7})
 \end{aligned}$$

with

$$\begin{aligned}
 C_{pr}(\theta'_o, \sigma) &= \int_{\Omega'_o} \int_{\Omega'_i} g(\theta'_i; \sigma) \tilde{P}_p(\cos \theta'_i) \tilde{P}_r(\cos \theta'_o) d\vec{\omega}'_i d\vec{\omega}'_o \\
 &= \int_{\Omega'_o} \tilde{P}_r(\cos \theta'_o) d\vec{\omega}'_o \int_{\phi'_i=0}^{2\pi} d\phi'_i \\
 &\quad \times \int_0^{\pi/2} g(\theta'_i; \sigma) \tilde{P}_p(\cos \theta'_i) \sin \theta'_i d\theta'_i \quad (\text{B-8})
 \end{aligned}$$

According to the orthogonality relation of the shifted Legendre polynomials, all terms in the integral in Equation B-8 will vanish except for $r = 0$.

$$\begin{aligned}
 C_{pr}(\theta'_o, \sigma) &= \frac{4\pi^2}{2r+1} \delta_{r0} \\
 &\times \underbrace{\int_0^{\pi/2} \frac{\sin \theta'_i}{\cos \theta'_i} \exp \left[- \left(\frac{\theta'_i}{2\sigma} \right)^2 \right] \tilde{P}_p(\cos \theta'_i) d\theta'_i}_{A_p(\theta'_o, \sigma)}
 \end{aligned} \tag{B-9}$$

A closed form of $A_p(\theta'_o, \sigma)$ can be derived using the series representation of shifted Legendre polynomials [199] and the exponential function, followed by variable substitution $z = \cos \theta'_i$ and Taylor series expansion of $\cos^{-1} z$. Nonetheless the closed form of the (shifted) Legendre polynomials suffer from numerical instabilities with high orders [60]. In this case, it is more efficient to use stable recurrence relations [99].

Therefore, the specular integral can be written as,

$$B_{pr}(\theta'_o, \sigma) = \begin{cases} 2\pi A_0(\theta'_o, \sigma) & p = 0, r = 0 \\ 2\pi \sqrt{\frac{2p+1}{2}} A_p(\theta'_o, \sigma) & p > 0, r = 0 \\ 2\pi \sqrt{\frac{2r+1}{2}} A_r(\theta'_o, \sigma) & p = 0, r > 0 \\ 0 & \text{otherwise} \end{cases} \tag{B-10}$$

This implies that $a_{p0} = a_{0p}$, thus a single index for spectrum coefficients can be used, where

$$a_p = 2\rho_d \delta_{p0} + \frac{\rho_s}{\pi \sigma^2 \cos \theta'_o} \frac{F(\cos \theta'_o; \eta)}{F(1; \eta)} B_p(\theta'_o, \sigma) \tag{B-11}$$

where

$$B_p(\theta'_o, \sigma) = \begin{cases} 2\pi A_0(\theta'_o, \sigma) & p = 0 \\ 2\pi \sqrt{\frac{2p+1}{2}} A_p(\theta'_o, \sigma) & p > 0 \end{cases} \tag{B-12}$$

According to the definition of the isotropic Helmholtz HSH-based basis , it can be

shown that $\mathcal{H}_{p0}^0 = \mathcal{H}_{0p}^0$, hence the TS-BRDF can be represented as,

$$\begin{aligned}
 f_r^{TS}(\vec{\omega}'_i, \vec{\omega}'_o) &= a_0(\theta'_o) + \sum_{p=1}^{\infty} a_p(\theta'_o) \mathcal{H}_{p0}^0(\vec{\omega}'_i, \vec{\omega}'_o) \\
 &+ \sum_{p=1}^{\infty} a_p(\theta'_o) \mathcal{H}_{0p}^0(\vec{\omega}'_i, \vec{\omega}'_o) \\
 &= a_0(\theta'_o) + 2 \sum_{p=1}^{\infty} a_p(\theta'_o) \mathcal{H}_{p0}^0(\vec{\omega}'_i, \vec{\omega}'_o)
 \end{aligned}$$

□

APPENDIX C: APPROXIMATION

ACCURACY OF DIRECTIONAL

HEMISPHERICAL BASIS

In this appendix, the total energy of a surface BRDF under fixed incident direction $\vec{\omega}_i^{o'}$ will be derived, where the total squared energy of the BRDF can be obtained by integrating $[f_r(\vec{\omega}_i^{o'}, \vec{\omega}'_o)]^2$ over the cartesian product of the incoming and outgoing hemispheres¹.

$$\begin{aligned} \int_{\Omega'_o} \int_{\Omega'_i} [f_r(\vec{\omega}_i^{o'}, \vec{\omega}'_o)]^2 d\omega'_i d\omega'_o &= \int_{\Omega'_i} d\omega'_i \int_{\Omega'_o} [f_r(\vec{\omega}_i^{o'}, \vec{\omega}'_o)]^2 d\omega'_o \\ &= 2\pi \int_{\Omega'_o} [f_r(\vec{\omega}_i^{o'}, \vec{\omega}'_o)]^2 d\omega'_o \end{aligned} \quad (\text{C-1})$$

Since, the BRDF can be represented in terms of the expansion coefficients a_{0p}^{0q} in Equation 4.58 as follows,

$$f_r(\vec{\omega}_i^{o'}, \vec{\omega}'_o) = \frac{1}{2\pi} a_{00}^{00}(\vec{\omega}_i^{o'}) + \frac{1}{2\sqrt{\pi}} \sum_{r=1}^{\infty} \sum_{s=-r}^r a_{0r}^{0s}(\vec{\omega}_i^{o'}) \left[H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}_i^{o'}) \right] \quad (\text{C-2})$$

The integral in Equation C-1 can be expanded as follows,

$$\begin{aligned} &\int_{\Omega'_o} \int_{\Omega'_i} [f_r(\vec{\omega}_i^{o'}, \vec{\omega}'_o)]^2 d\omega'_i d\omega'_o \\ &= 2\pi \int_{\Omega'_o} \left\{ \frac{1}{2\pi} a_{00}^{00}(\vec{\omega}_i^{o'}) + \frac{1}{2\sqrt{\pi}} \sum_{r=1}^{\infty} \sum_{s=-r}^r a_{0r}^{0s}(\vec{\omega}_i^{o'}) \left[H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}_i^{o'}) \right] \right\} \\ &\quad \left\{ \frac{1}{2\pi} a_{00}^{00}(\vec{\omega}_i^{o'}) + \frac{1}{2\sqrt{\pi}} \sum_{r'=1}^{\infty} \sum_{s'=-r'}^{r'} a_{0r'}^{0s'}(\vec{\omega}_i^{o'}) \left[H_{r'}^{s'}(\vec{\omega}'_o) + H_{r'}^{s'}(\vec{\omega}_i^{o'}) \right] \right\} d\omega'_o \\ &= 2\pi \left\{ \left(\frac{1}{2\pi} a_{00}^{00}(\vec{\omega}_i^{o'}) \right)^2 \int_{\Omega'_o} d\omega'_o \right. \end{aligned}$$

¹Having a fixed incident direction does not eliminate the integration with respect to it.

$$\begin{aligned}
 & + 2 \frac{1}{2\pi} \frac{1}{2\sqrt{\pi}} a_{00}^{00}(\vec{\omega}'_o) \sum_{r=1}^{\infty} \sum_{s=-r}^r a_{0r}^{0s}(\vec{\omega}'_o) \int_{\Omega'_o} \left[H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i) \right] d\omega'_o \\
 & + \left(\frac{1}{2\sqrt{\pi}} \right)^2 \sum_{r=1}^{\infty} \sum_{s=-r}^r \sum_{r'=1}^{\infty} \sum_{s'=-r'}^{r'} a_{0r}^{0s}(\vec{\omega}'_o) a_{0r'}^{0s'}(\vec{\omega}'_o) \\
 & \quad \left. \int_{\Omega'_o} \left[H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i) \right] \left[H_{r'}^{s'}(\vec{\omega}'_o) + H_{r'}^{s'}(\vec{\omega}'_i) \right] d\omega'_o \right\} \quad (\text{C-3})
 \end{aligned}$$

Using the orthonormality of HSH basis Equation 4.6 and the fact that $H_0^0(\vec{\omega}'_o) = \frac{1}{\sqrt{2\pi}}$, it can be shown that,

$$\begin{aligned}
 & \int_{\Omega'_o} \left[H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i) \right] d\omega'_o \\
 & = \int_{\Omega'_o} H_r^s(\vec{\omega}'_o) d\omega'_o + \int_{\Omega'_o} H_r^s(\vec{\omega}'_i) d\omega'_o \\
 & = \sqrt{2\pi} \int_{\Omega'_o} H_r^s(\vec{\omega}'_o) H_0^0(\vec{\omega}'_o) d\omega'_o + 2\pi H_r^s(\vec{\omega}'_i) \\
 & = \sqrt{2\pi} \delta_{r0} \delta_{s0} + 2\pi H_r^s(\vec{\omega}'_i) \quad (\text{C-4})
 \end{aligned}$$

Also,

$$\begin{aligned}
 & \int_{\Omega'_o} \left[H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i) \right] \left[H_{r'}^{s'}(\vec{\omega}'_o) + H_{r'}^{s'}(\vec{\omega}'_i) \right] d\omega'_o \\
 & = \int_{\Omega'_o} H_r^s(\vec{\omega}'_o) H_{r'}^{s'}(\vec{\omega}'_o) d\omega'_o + \int_{\Omega'_o} H_r^s(\vec{\omega}'_o) H_{r'}^{s'}(\vec{\omega}'_i) d\omega'_o \\
 & \quad + \int_{\Omega'_o} H_r^s(\vec{\omega}'_i) H_{r'}^{s'}(\vec{\omega}'_o) d\omega'_o + \int_{\Omega'_o} H_r^s(\vec{\omega}'_i) H_{r'}^{s'}(\vec{\omega}'_i) d\omega'_o \\
 & = \delta_{rr'} \delta_{ss'} + \sqrt{2\pi} H_{r'}^{s'}(\vec{\omega}'_i) \delta_{r0} \delta_{s0} \\
 & \quad + \sqrt{2\pi} H_r^s(\vec{\omega}'_i) \delta_{r'0} \delta_{s'0} + 2\pi H_r^s(\vec{\omega}'_i) H_{r'}^{s'}(\vec{\omega}'_i) \quad (\text{C-5})
 \end{aligned}$$

Setting $r = r'$ and $s = s'$ yields,

$$\begin{aligned}
 & \int_{\Omega'_o} \left[H_r^s(\vec{\omega}'_o) + H_r^s(\vec{\omega}'_i) \right] \left[H_r^{s'}(\vec{\omega}'_o) + H_r^{s'}(\vec{\omega}'_i) \right] d\omega'_o \\
 & = 1 + 2\sqrt{2\pi} H_r^s(\vec{\omega}'_i) \delta_{r0} \delta_{s0} + 2\pi \left[H_r^s(\vec{\omega}'_i) \right]^2 \quad (\text{C-6})
 \end{aligned}$$

Embedding Equation C-4 and Equation C-6 into Equation C-3 leads to,

$$\begin{aligned}
 & \int_{\Omega'_o} \int_{\Omega'_i} [f_r(\vec{\omega}'_i, \vec{\omega}'_o)]^2 d\omega'_i d\omega'_o \\
 & = 2\pi \left\{ \frac{1}{2\pi} \left(a_{00}^{00}(\vec{\omega}'_i) \right)^2 \right.
 \end{aligned}$$

$$\begin{aligned}
 & + \frac{1}{\sqrt{\pi}} a_{00}^{00}(\vec{\omega}'_o) \sum_{r=1}^{\infty} \sum_{s=-r}^r a_{0r}^{0s}(\vec{\omega}'_o) H_r^s(\vec{\omega}'_o) \\
 & + \left(\frac{1}{2\sqrt{\pi}} \right)^2 \sum_{r=1}^{\infty} \sum_{s=-r}^r \left(a_{0r}^{0s}(\vec{\omega}'_o) \right)^2 \left(1 + 2\pi \left[H_r^s(\vec{\omega}'_o) \right]^2 \right) \Big\} \\
 = & \left(a_{00}^{00}(\vec{\omega}'_o) \right)^2 \\
 & + \sum_{r=1}^{\infty} \sum_{s=-r}^r \left\{ 2\sqrt{\pi} a_{00}^{00}(\vec{\omega}'_o) a_{0r}^{0s}(\vec{\omega}'_o) H_r^s(\vec{\omega}'_o) \right. \\
 & \left. + \left(a_{0r}^{0s}(\vec{\omega}'_o) \right)^2 \left(\frac{1}{2} + \pi \left[H_r^s(\vec{\omega}'_o) \right]^2 \right) \right\} \tag{C-7}
 \end{aligned}$$

Considering the P th order function approximation, such an approximation captures a certain amount of the function's total energy defined as the cumulative sum of energies maintained by individual harmonic terms, (this cumulative energy can also be referred to as *function approximation accuracy* since the more energy the harmonic terms capture from the function to be approximated, the more accurate the harmonic expansion is). Thus the approximation accuracy (or the energy captured by the P th order approximation) can be defined as,

$$\begin{aligned}
 Acc_B(R; \vec{\omega}'_o) = & \begin{cases} \frac{\left(a_{00}^{00}(\vec{\omega}'_o) \right)^2}{\int_{\Omega'_o} \int_{\Omega'_i} [f_r(\vec{\omega}'_o, \vec{\omega}'_o)]^2 d\omega'_i d\omega'_o} & R = 0 \\ \frac{\left(a_{00}^{00}(\vec{\omega}'_o) \right)^2 + \sum_{r=1}^R \sum_{s=-r}^r \left\{ 2\sqrt{\pi} a_{00}^{00}(\vec{\omega}'_o) a_{0r}^{0s}(\vec{\omega}'_o) H_r^s(\vec{\omega}'_o) + \left(a_{0r}^{0s}(\vec{\omega}'_o) \right)^2 \left(\frac{1}{2} + \pi \left[H_r^s(\vec{\omega}'_o) \right]^2 \right) \right\}}{\int_{\Omega'_o} \int_{\Omega'_i} [f_r(\vec{\omega}'_o, \vec{\omega}'_o)]^2 d\omega'_i d\omega'_o} & R > 0 \end{cases} \tag{C-8}
 \end{aligned}$$

□

APPENDIX D: ANALYTIC

DERIVATION OF $C_{npr}^q(\alpha)$

This appendix derives the analytic expression of the integral in Equation 5.9 which is used to define the irradiance harmonics of the proposed reflectance basis after moving terms which are independent of the local incident direction out of the integral, where

$$\mathcal{C}_{npr}^{m'q}(\alpha) = \int_{\Omega'_i} Y_n^{m'}(\vec{\omega}'_i) \mathcal{H}_{pr}^q(\theta'_i, \alpha, |\phi'_i - \pi|) \cos \theta'_i d\vec{\omega}'_i \quad (\text{D-1})$$

The solution of this integral encodes how much the spherical harmonics basis functions $Y_n^{m'}$ resembles the proposed isotropic Helmholtz HSH-based reflectance basis \mathcal{H}_{pr}^q . The real spherical harmonics (SH) basis is defined as,

$$Y_n^{m'}(\vec{\omega}'_i) = Y_n^{m'}(\theta'_i, \phi'_i) = \Theta_n^{m'}(\theta'_i) \Phi_{m'}(\phi'_i) \quad (\text{D-2})$$

with $n \geq 0$ denoting the harmonic order, $m' \in [-n, n]$ denoting the harmonic degree, $\theta'_i \in [0, \pi]$ and $\phi'_i \in [0, 2\pi]$. The *polar part* is given by $\Theta_n^{m'}(\theta'_i) = N_n^{|m'|} P_n^{|m'|}(\cos \theta'_i)$ where $P_n^{m'}(z)$ is the associated Legendre polynomials defined over the interval $[-1, 1]$. They are defined in terms of the m' th derivative of the Legendre polynomials as¹,

$$P_n^{m'}(z) = \frac{1}{2^n n!} (1 - z^2)^{m'/2} \frac{d^{m'}}{dz^{m'}} (z^2 - 1)^n, \quad n \geq 0, m' \in [0, n] \quad (\text{D-3})$$

¹In [56], the associated Legendre polynomials is defined with an additional $(-1)^{m'}$, which seems a complication that can be ignored at this point. It will be included in the definition of the spherical harmonics later.

The *azimuthal part* is defined as,

$$\Phi_{m'}(\phi) = \begin{cases} \sqrt{2} \cos(m' \phi'_i) & m' > 0 \\ 1 & m' = 0 \\ (-1)^{m'} \sqrt{2} \sin(-m' \phi'_i) & m' < 0 \end{cases} \quad (\text{D-4})$$

The normalization factor $N_n^{m'}$ which accounts for basis orthonormality is defined as,

$$N_n^{m'} = \sqrt{\frac{2n+1}{4\pi} \frac{(n-m')!}{(n+m')!}} \quad (\text{D-5})$$

Analogical to spherical harmonics definition, the real hemispherical harmonics can be written as [20]¹,

$$H_p^q(\vec{\omega}') = H_p^q(\theta', \phi') = \tilde{\Theta}_p^q(\theta') \Phi_q(\phi') \quad (\text{D-6})$$

The *polar part* is given by $\tilde{\Theta}_p^q(\theta') = \tilde{N}_p^{|q|} \tilde{P}_p^{|q|}(\cos \theta')$ where $\tilde{P}_p^q(z) = P_p^q(2z-1)$ is the shifted associated Legendre polynomials defined over the interval $[0, 1]$, the *azimuthal part* is defined in Equation D-4 and \tilde{N}_p^q being defined as the normalization factor,

$$\tilde{N}_p^q = \sqrt{\frac{2p+1}{2\pi} \frac{(p-q)!}{(p+q)!}} \quad (\text{D-7})$$

where the harmonic order $p \geq 0$, the harmonic degree $q \in [-n, n]$, $\theta' \in [0, \pi/2]$ and $\phi' \in [0, 2\pi]$.

Thus the reflectance basis for isotropic surfaces can be defined in terms of the polar and azimuthal parts of the hemispherical harmonics as,

$$\mathcal{H}_{pr}^q(\theta'_i, \theta'_o, \Delta\phi') = N_{pr}^q \left[\tilde{\Theta}_p^q(\theta'_i) \tilde{\Theta}_r^q(\theta'_o) + \tilde{\Theta}_r^q(\theta'_i) \tilde{\Theta}_p^q(\theta'_o) \right] \Phi_q(\Delta\phi') \quad (\text{D-8})$$

where the normalization factor is defined as,

$$N_{pr}^q = \frac{1}{\sqrt{2+2\delta_{pr}}} \begin{cases} \sqrt{2} & q > 0 \\ 1 & q = 0 \\ 0 & q < 0 \end{cases} \quad (\text{D-9})$$

¹Note the difference between us and [20] in the factor of $(-1)^q$, since it is preferred to include this factor in the harmonics definition rather than including it in the associated Legendre polynomial definition.

The solution of the integral in Equation D-1 can be derived as follows.

$$\begin{aligned}
 \mathcal{C}_{npr}^{m'q}(\alpha) &= N_{pr}^q \int_0^{2\pi} \Phi_{m'}(\phi'_i) \Phi_q(|\phi'_i - \pi|) d\phi'_i \\
 &\times \int_0^{\pi/2} \Theta_n^{m'}(\theta'_i) \left[\tilde{\Theta}_p^q(\theta'_i) \tilde{\Theta}_r^q(\theta'_o) + \tilde{\Theta}_r^q(\theta'_i) \tilde{\Theta}_p^q(\theta'_o) \right] \\
 &\quad \times \cos \theta'_i \sin \theta'_i d\theta'_i
 \end{aligned} \tag{D-10}$$

Since $q \geq 0$ (the normalization factor N_{pr}^q for negative degree basis is zero), thus the azimuthal part becomes [32],

$$\begin{aligned}
 &\int_0^{2\pi} \Phi_{m'}(\phi'_i) \Phi_q(|\phi'_i - \pi|) d\phi'_i \\
 &= \int_0^{2\pi} \Phi_{m'}(\phi'_i) \cos(q(\phi'_i - \pi)) d\phi'_i \\
 &= \delta_{m'q} \pi (-1)^q \sqrt{2(1 + \delta_{q0})}
 \end{aligned} \tag{D-11}$$

Thus all terms in Equation D-10 vanish except for $m' = q$.

$$\begin{aligned}
 \mathcal{C}_{npr}^{m'q}(\alpha) &= \mathcal{C}_{npr}^q(\alpha) = \pi N_{pr}^q (-1)^q \sqrt{2(1 + \delta_{q0})} \\
 &\quad \left\{ \tilde{\Theta}_r^q(\theta'_o) N_n^{|q|} \tilde{N}_p^{|q|} \mathcal{J}_{np}^q + \tilde{\Theta}_p^q(\theta'_o) N_n^{|q|} \tilde{N}_r^{|q|} \mathcal{J}_{nr}^q \right\}
 \end{aligned} \tag{D-12}$$

where,

$$\mathcal{J}_{np}^q = \int_0^{\pi/2} P_n^{|q|}(\cos \theta'_i) \tilde{P}_p^{|q|}(\cos \theta'_i) \cos \theta'_i \sin \theta'_i d\theta'_i \tag{D-13}$$

The associated Legendre polynomial in its explicit form [98] is written as,

$$\begin{aligned}
 P_n^{|q|}(\cos \theta'_i) &= \frac{(n + |q|)!}{2^n n!} \sin^{|q|} \theta'_i \sum_{k=0}^{\lfloor \frac{n-|q|}{2} \rfloor} (-1)^{|q|+k} \\
 &\quad \times \binom{n}{k} \binom{2n-2k}{n+|q|} \cos^{n-|q|-2k} \theta'_i
 \end{aligned} \tag{D-14}$$

The shifted associated Legendre polynomials is given by,

$$\tilde{P}_p^q(\cos \theta'_i) = P_p^q(2 \cos \theta'_i - 1) \tag{D-15}$$

Since $\theta'_i \geq 0$ and $|q| \geq 0$, $\sin^{|q|} \theta'_i$ in Equation D-14 can be expressed as $(1 - \cos \theta'_i)^{|q|/2}$,

thus the shifted associated Legendre can be expressed explicitly as,

$$\begin{aligned} \tilde{P}_p^{|q|}(\cos \theta'_i) &= \frac{(p+|q|)!}{2^p p!} [1 - (2 \cos \theta'_i - 1)^2]^{|q|/2} \\ &\times \sum_{l=0}^{\lfloor \frac{p-|q|}{2} \rfloor} (-1)^{|q|+l} \binom{p}{l} \binom{2p-2l}{p+|q|} \\ &\quad [2 \cos \theta'_i - 1]^{p-|q|-2l} \end{aligned} \quad (\text{D-16})$$

Using variable substitution, $z = \cos \theta'_i \rightarrow dz = -\sin \theta'_i d\theta'_i$, the explicit forms in Equation D-14 and Equation D-16 can be re-written as follows.

$$\begin{aligned} P_n^{|q|}(z) &= \frac{(n+|q|)!}{2^n n!} (1-z^2)^{|q|/2} \sum_{k=0}^{\lfloor \frac{n-|q|}{2} \rfloor} (-1)^{|q|+k} \\ &\times \binom{n}{k} \binom{2n-2k}{n+|q|} z^{n-|q|-2k} \end{aligned} \quad (\text{D-17})$$

and

$$\begin{aligned} \tilde{P}_p^{|q|}(z) &= \frac{(p+|q|)!}{2^p p!} [4z(1-z)]^{|q|/2} \sum_{l=0}^{\lfloor \frac{p-|q|}{2} \rfloor} (-1)^{|q|+l} \\ &\times \binom{p}{l} \binom{2p-2l}{p+|q|} [2z-1]^{p-|q|-2l} \end{aligned} \quad (\text{D-18})$$

Putting all together, the closed form of the integral \mathcal{J}_{np}^q can be derived as follows,

$$\begin{aligned} \mathcal{J}_{np}^q &= \frac{(n+|q|)!(p+|q|)!}{2^{n+p} n! p!} \sum_{k=0}^{\lfloor \frac{n-|q|}{2} \rfloor} \sum_{l=0}^{\lfloor \frac{p-|q|}{2} \rfloor} (-1)^{2|q|+k+l} \\ &\times \binom{n}{k} \binom{2n-2k}{n+|q|} \binom{p}{l} \binom{2p-2l}{p+|q|} \\ &\times \int_0^1 (1-z^2)^{\frac{|q|}{2}} z^{n-|q|-2k+1} \\ &\quad [4z(1-z)]^{\frac{|q|}{2}} (2z-1)^{p-|q|-2l} dz \end{aligned} \quad (\text{D-19})$$

The integral in Equation D-19 can be re-written as;

$$2^{|q|} \int_0^1 (1-z)^{|q|} (1+z)^{|q|/2} z^{n-\frac{|q|}{2}-2k+1} (2z-1)^{p-|q|-2l} dz \quad (\text{D-20})$$

Using the binomial theorem which restricts q to be even,

$$(1+z)^{|q|/2} = \sum_{i=0}^{|q|/2} \binom{|q|}{i} z^i \quad (\text{D-21})$$

and

$$(2z-1)^{p-|q|-2l} = \sum_{j=0}^{p-|q|-2l} \binom{p-|q|-2l}{j} (2z)^j (-1)^{p-|q|-2l-j} \quad (\text{D-22})$$

Substituting in Equation D-20,

$$2^{|q|} \sum_{i=0}^{|q|/2} \sum_{j=0}^{p-|q|-2l} 2^j \binom{|q|}{i} \binom{p-|q|-2l}{j} \int_0^1 z^{n-\frac{|q|}{2}-2k+1+i+j} (1-z)^{|q|} dz \quad (\text{D-23})$$

Using the Euler Beta integral [99] leads to,

$$2^{|q|} \sum_{i=0}^{|q|/2} \sum_{j=0}^{p-|q|-2l} 2^j \binom{|q|}{i} \binom{p-|q|-2l}{j} \frac{(n-\frac{|q|}{2}-2k+1+i+j)! (|q|)!}{(n+\frac{|q|}{2}-2k+2+i+j)!} \quad (\text{D-24})$$

Thus the closed form of \mathcal{J}_{np}^q is given by,

$$\begin{aligned} \mathcal{J}_{np}^q &= \frac{(n+|q|)! (p+|q|)!}{2^{n+p} n! p!} \sum_{k=0}^{\lfloor \frac{n-|q|}{2} \rfloor} \sum_{l=0}^{\lfloor \frac{p-|q|}{2} \rfloor} (-1)^{2|q|+k+l} \\ &\times \binom{n}{k} \binom{2n-2k}{n+|q|} \binom{p}{l} \binom{2p-2l}{p+|q|} \\ &\times 2^{|q|} \sum_{i=0}^{|q|/2} \sum_{j=0}^{p-|q|-2l} 2^j \binom{|q|}{i} \binom{p-|q|-2l}{j} \\ &\frac{(n-\frac{|q|}{2}-2k+1+i+j)! (|q|)!}{(n+\frac{|q|}{2}-2k+2+i+j)!} \end{aligned} \quad (\text{D-25})$$

□

APPENDIX E: IRRADIANCE SIGNAL

ENERGY CONTENT

The energy of the irradiance signal perceived by the camera is given by $\int_{\Omega} [E(\alpha, \beta)]^2 d\vec{\omega}$, where the integral is over the surface normals constructing the visible hemisphere Ω . Assuming directional light source (modeled as a delta function with non-decayable spectrum), the decay of the irradiance spectrum will be due to the spectrum coefficients of the surface BRDF. Eliminating the effect of the light direction, consider all possible 3D rotations applied on a directional light source, the average irradiance energy is then obtained as,

$$\varepsilon = \int_{\Omega} \int_{SO(3)} [E(\alpha, \beta)]^2 dR d\vec{\omega} \quad (\text{E-1})$$

When a light source, with spectrum coefficients \tilde{l}_n^m , undergoes a 3D rotation R , the coefficients becomes,

$$l_n^m = \sum_{o=-n}^n \tilde{l}_n^o \mathcal{D}_{mo}^n(R) \quad (\text{E-2})$$

where \mathcal{D}^n are Wigner's rotation matrices [59]. Thus the average signal energy of each illumination mode over all possible 3D rotations can be written as,

$$\begin{aligned} \int_{SO(3)} [l_n^m]^2 dR &= \int_{SO(3)} l_n^m l_{n'}^{m'} dR \\ &= \sum_{o=-n}^n \sum_{o'=-n'}^{n'} \tilde{l}_n^o \tilde{l}_{n'}^{o'} \\ &\quad \times \int_{SO(3)} \mathcal{D}_{mo}^n(R) \mathcal{D}_{m'o'}^{n'}(R) dR \end{aligned} \quad (\text{E-3})$$

From the orthogonality relation of the \mathcal{D}^n matrices [11], it can be shown that,

$$\int_{SO(3)} \mathcal{D}_{mo}^n(R) \mathcal{D}_{m'o'}^{n'}(R) dR = \frac{1}{2n+1} \delta_{nn'} \delta_{mm'} \delta_{oo'} \quad (\text{E-4})$$

where $dR = \frac{1}{8\pi} \sin \alpha d\alpha d\beta d\gamma$. Thus Equation E-3 simply becomes,

$$\int_{SO(3)} [l_n^m]^2 dR = \frac{1}{2n+1} \sum_{o=-n}^n (\tilde{l}_n^o)^2 \quad (\text{E-5})$$

The average energy content of the irradiance signal E over all possible directional light sources can be written as,

$$\begin{aligned} \varepsilon = & \sum_{\substack{nm \\ n'm'}} \int_{SO(3)} l_n^m l_{n'}^{m'} dR \\ & \times \sum_{\substack{prq \\ p'r'q'}} a_{pr}^q a_{p'r'}^{q'} \\ & \int_{\Omega} \mathcal{B}_{npr}^{mq}(\alpha, \beta) \mathcal{B}_{n'p'r'}^{m'q'}(\alpha, \beta) d\vec{\omega} \end{aligned} \quad (\text{E-6})$$

According to Equation E-4, all terms in Equation E-6 vanish except $n = n'$ and $m = m'$,

$$\varepsilon = \sum_{n=0}^{\infty} \left(\frac{\sum_{o=-n}^n \tilde{l}_n^o}{2n+1} \right) \mathcal{A}_B(n) \quad (\text{E-7})$$

where,

$$\begin{aligned} \mathcal{A}_B(n) = & \sum_{\substack{prq \\ p'r'q'}} a_{pr}^q a_{p'r'}^{q'} \\ & \times \sum_{m=-n}^n \int_{\Omega} \mathcal{B}_{npr}^{mq}(\alpha, \beta) \mathcal{B}_{np'r'}^{m'q'}(\alpha, \beta) d\vec{\omega} \end{aligned} \quad (\text{E-8})$$

with¹,

$$\mathcal{B}_{npr}^{mq}(\alpha, \beta) = \mathcal{D}_{mq}^n(\alpha, \beta) \mathcal{C}_{npr}^q(\alpha) \quad (\text{E-9})$$

Thus the average energy now depends on the illumination order n where each terms is the product of the average energy content of the illumination band n and the corresponding attenuation factor \mathcal{A}_B imposed by the surface BRDF on the illumination signal.

¹Note that $\mathcal{C}_{npr}^q(\alpha)$ vanishes when q is odd.

The \mathcal{D}^n matrices¹ can be expressed in a compact form [32],

$$\begin{aligned} \mathcal{D}_{mq}^n(\alpha, \beta, \gamma) &= \frac{\text{sgn}(q)}{2} \{ [\Phi_m(\beta)\Phi_q(\gamma) - \Phi_{-m}(\beta)\Phi_{-q}(\gamma)] \\ &\quad \times P_n^{m,q}(\cos \alpha) \\ &\quad + [\Phi_m(\beta)\Phi_q(\gamma) - \Phi_{-m}(\beta)\Phi_{-q}(\gamma)] \\ &\quad \times (-1)^q P_n^{m,-q}(\cos \alpha) \} \end{aligned} \quad (\text{E-10})$$

with

$$\text{sgn}(x) = \begin{cases} 1 & x \geq 0 \\ -1 & x < 0 \end{cases} \quad (\text{E-11})$$

where $P_n^{m,q}$ are the generalizations of the associated Legendre functions [200] which are given by the Rodriguez formula,

$$\begin{aligned} P_n^{m,q}(z) &= \frac{(-1)^{n-m}}{2^n} \sqrt{\frac{(n+m)!}{(n-q)!(n+q)!(n-m)!}} \\ &\quad \times (1+z)^{-\frac{(m+q)}{2}} (1-z)^{\frac{q-m}{2}} \\ &\quad \times \frac{d^{n-m}}{dz^{n-m}} [(1-z)^{n-q} (1+z)^{n+q}] \end{aligned} \quad (\text{E-12})$$

Since isotropic surfaces is of concern, γ can be arbitrary set to 0 and $q \geq 0$, thus the azimuthal functions that depend on γ in Equation E-10 become,

$$\Phi_q(\gamma) = \begin{cases} \sqrt{2} & q \geq 0 \\ 1 & q = 0 \end{cases} \quad (\text{E-13})$$

and

$$\Phi_{-q}(\gamma) = \begin{cases} 0 & q \geq 0 \\ 1 & q = 0 \end{cases} \quad (\text{E-14})$$

Therefore, in case of isotropic surfaces, the \mathcal{D}^n matrices can be written as,

$$\mathcal{D}_{mq}^n(\alpha, \beta) = \begin{cases} \frac{1}{\sqrt{2}} [\Phi_m(\beta)P_n^{m,q}(\cos \alpha) \\ \quad + (-1)^q \Phi_m(\beta)P_n^{m,-q}(\cos \alpha)] & q \geq 0 \\ \Phi_m(\beta)P_n^{m,0}(\cos \alpha) & q = 0 \end{cases} \quad (\text{E-15})$$

¹ γ is set to zero for isotropic surfaces.

In a more compact form,

$$\begin{aligned}\mathcal{D}_{mq}^n(\alpha, \beta) &= \frac{P_n^{m,q}(\cos \alpha) + (-1)^q P_n^{m,-q}(\cos \alpha)}{\sqrt{2(1 + \delta_{q0})}} \Phi_m(\beta) \\ &= \Delta_{m,q}^n(\alpha) \Phi_m(\beta)\end{aligned}\tag{E-16}$$

This leads to a *polar-azimuth* factorization of the basis functions in Equation E-9 such that,

$$\mathcal{B}_{npr}^{mq}(\alpha, \beta) = \Theta_{npr}^{mq}(\alpha) \Phi_m(\beta)\tag{E-17}$$

where,

$$\Theta_{npr}^{mq}(\alpha) = \Delta_{m,q}^n(\alpha) \mathcal{C}_{npr}^q(\alpha)\tag{E-18}$$

Thus the integral in $\mathcal{A}_B(n)$ Equation E-8 can be re-written as follows where the orthogonality relation of the azimuth functions eliminates the dependency on the azimuthal angle β ,

$$\begin{aligned}\int_{\Omega} \mathcal{B}_{npr}^{mq}(\alpha, \beta) \mathcal{B}_{np'r'}^{mq'}(\alpha, \beta) d\vec{\omega} &= \\ &\underbrace{\int_0^{2\pi} \Phi_m(\beta) \Phi_m(\beta) d\beta}_{2\pi} \\ &\times \int_0^{\pi/2} \Theta_{npr}^{mq}(\alpha) \Theta_{np'r'}^{mq'}(\alpha) \sin \alpha d\alpha\end{aligned}\tag{E-19}$$

Moving the summation over m inside the integral of Equation E-8,

$$\begin{aligned}\mathcal{A}_B(n) &= \sum_{\substack{prq \\ p'r'q'}} 2\pi a_{pr}^q a_{p'r'}^{q'} \int_0^{\pi/2} \mathcal{C}_{npr}^q(\alpha) \mathcal{C}_{np'r'}^{q'}(\alpha) \\ &\quad \times \left[\sum_{m=-n}^n \Delta_{m,q}^n(\alpha) \Delta_{m,q'}^n(\alpha) \right] \sin \alpha d\alpha\end{aligned}\tag{E-20}$$

where,

$$\begin{aligned}
 & \sum_{m=-n}^n \Delta_{m,q}^n(\alpha) \Delta_{m,q'}^n(\alpha) \\
 &= \frac{1}{2\sqrt{(1+\delta_{q0})(1+\delta_{q'0})}} \\
 &\times \left\{ \sum_{m=-n}^n P_n^{m,q}(\cos \alpha) P_n^{m,q'}(\cos \alpha) \right. \\
 &+ (-1)^{q'} \sum_{m=-n}^n P_n^{m,q}(\cos \alpha) P_n^{m,-q'}(\cos \alpha) \\
 &+ (-1)^q \sum_{m=-n}^n P_n^{m,-q}(\cos \alpha) P_n^{m,q'}(\cos \alpha) \\
 &\left. + (-1)^{q+q'} \sum_{m=-n}^n P_n^{m,-q}(\cos \alpha) P_n^{m,-q'}(\cos \alpha) \right\}
 \end{aligned} \tag{E-21}$$

However, the $P_n^{m,q}$ functions obey the following relation [32],

$$\sum_{m=-n}^n P_n^{m,q}(z) P_n^{m,q'}(z) = \delta_{qq'} \tag{E-22}$$

Noting that $q, q' \geq 0$, hence $\delta_{q,-q'} = \delta_{-q,q'} = \delta_{qq'}\delta_{q0}$ and $\delta_{-q,-q'} = \delta_{qq'}$, thus the summation of Equation E-23 boils to¹,

$$\begin{aligned}
 & \sum_{m=-n}^n \Delta_{m,q}^n(\alpha) \Delta_{m,q'}^n(\alpha) \\
 &= \frac{1}{2\sqrt{(1+\delta_{q0})(1+\delta_{q'0})}} \delta_{qq'} \\
 &\times \left\{ 1 + (-1)^{q'} \delta_{q0} + (-1)^q \delta_{q0} + (-1)^{q+q'} \right\} \\
 &= \frac{1}{2\sqrt{(1+\delta_{q0})(1+\delta_{q'0})}} \delta_{qq'} \{2 + 2\delta_{q0}\} \\
 &= \frac{(1+\delta_{q0})}{\sqrt{(1+\delta_{q0})(1+\delta_{q'0})}} \delta_{qq'} = \delta_{qq'}
 \end{aligned} \tag{E-23}$$

¹ $(-1)^q \delta_{q0} = \delta_{q0}$

Thus all terms in Equation E-20 vanish except for $q = q'$,

$$\begin{aligned} \mathcal{A}_B(n) &= \sum_{\substack{prq \\ p'r' \\ q \text{ is even}}} 2\pi a_{pr}^q a_{p'r'}^q \\ &\quad \times \int_0^{\pi/2} \mathcal{C}_{npr}^q(\alpha) \mathcal{C}_{np'r'}^q(\alpha) \sin \alpha d\alpha \end{aligned} \quad (\text{E-24})$$

Using Equation D-12, the polar integral in Equation E-24 can be expanded as follows,

$$\begin{aligned} &\int_0^{\pi/2} \mathcal{C}_{npr}^q(\alpha) \mathcal{C}_{np'r'}^q(\alpha) \sin \alpha d\alpha \\ &= 2\pi^2 (1 + \delta_{q0}) N_{pr}^q N_{p'r'}^q \left(N_n^{|q|} \right)^2 \\ &\quad \times \left\{ \tilde{N}_p^{|q|} \tilde{N}_{p'}^{|q|} \mathcal{J}_{np}^q \mathcal{J}_{np'}^q \int_0^{\pi/2} \tilde{\Theta}_r^q(\alpha) \tilde{\Theta}_{r'}^q(\alpha) \sin \alpha d\alpha \right. \\ &\quad + \tilde{N}_p^{|q|} \tilde{N}_{r'}^{|q|} \mathcal{J}_{np}^q \mathcal{J}_{nr'}^q \int_0^{\pi/2} \tilde{\Theta}_r^q(\alpha) \tilde{\Theta}_{p'}^q(\alpha) \sin \alpha d\alpha \\ &\quad + \tilde{N}_r^{|q|} \tilde{N}_{p'}^{|q|} \mathcal{J}_{nr}^q \mathcal{J}_{np'}^q \int_0^{\pi/2} \tilde{\Theta}_p^q(\alpha) \tilde{\Theta}_{r'}^q(\alpha) \sin \alpha d\alpha \\ &\quad \left. + \tilde{N}_r^{|q|} \tilde{N}_{r'}^{|q|} \mathcal{J}_{nr}^q \mathcal{J}_{nr'}^q \int_0^{\pi/2} \tilde{\Theta}_p^q(\alpha) \tilde{\Theta}_{p'}^q(\alpha) \sin \alpha d\alpha \right\} \end{aligned} \quad (\text{E-25})$$

where the orthogonality relation of the shifted associated Legendre polynomials causes the Θ -integral to simply become¹,

$$\begin{aligned} &\int_0^{\pi/2} \tilde{\Theta}_r^q(\alpha) \tilde{\Theta}_{r'}^q(\alpha) \sin \alpha d\alpha \\ &= \tilde{N}_r^{|q|} \tilde{N}_{r'}^{|q|} \int_0^{\pi/2} \tilde{P}_r^{|q|}(\cos \alpha) \tilde{P}_{r'}^{|q|}(\cos \alpha) \sin \alpha d\alpha \\ &= \left(\tilde{N}_r^{|q|} \right)^2 \frac{(r + |q|)!}{(2r + 1)(r - |q|)!} \delta_{rr'} = \frac{1}{2\pi} \delta_{rr'} \end{aligned} \quad (\text{E-26})$$

¹Using variable substitution $z = \cos \alpha \rightarrow dz = -\sin \alpha d\alpha$.

Therefore, Equation E-25 becomes,

$$\begin{aligned}
 & \int_0^{\pi/2} \mathcal{C}_{npr}^q(\alpha) \mathcal{C}_{np'r'}^q(\alpha) \sin \alpha d\alpha \\
 &= \pi(1 + \delta_{q0}) (N_{pr}^q)^2 (N_n^{|q|})^2 \\
 &\times \left\{ (\tilde{N}_p^{|q|})^2 (\mathcal{J}_{np}^q)^2 + (\tilde{N}_r^{|q|})^2 (\mathcal{J}_{nr}^q)^2 \right. \\
 &\left. + 2 \tilde{N}_p^{|q|} \tilde{N}_r^{|q|} \mathcal{J}_{np}^q \mathcal{J}_{nr}^q \delta_{pr} \right\} \delta_{pp'} \delta_{rr'}
 \end{aligned} \tag{E-27}$$

such that q is even. Therefore, the BRDF attenuation factors can be written as,

$$\begin{aligned}
 \mathcal{A}_B(n) &= 2\pi^2 \sum_{\substack{prq \\ q \text{ is even}}} (1 + \delta_{q0}) (a_{pr}^q)^2 (N_{pr}^q)^2 (N_n^{|q|})^2 \\
 &\left\{ (\tilde{N}_p^{|q|})^2 (\mathcal{J}_{np}^q)^2 + (\tilde{N}_r^{|q|})^2 (\mathcal{J}_{nr}^q)^2 \right. \\
 &\left. + 2\tilde{N}_p^{|q|} \tilde{N}_r^{|q|} \mathcal{J}_{np}^q \mathcal{J}_{nr}^q \delta_{pr} \right\}
 \end{aligned} \tag{E-28}$$

□

APPENDIX F: PROOF OF ANALYTIC BILINEAR PCA THEOREM

Proof. Based on Alternating Least Squares [107], given \mathbf{A}_V , the weighting matrix \mathbf{A}_U maximizes the objective function given by,

$$\Psi_{\mathbf{Y}} = E\{\|\mathbf{Y} - \bar{\mathbf{Y}}\|_F^2\} \quad (\text{F-1})$$

Using the property of matrix trace and the linearity property of the expectation operator, it can be shown that,

$$\Psi_{\mathbf{Y}} = \text{trace} \left(E \left\{ (\mathbf{Y} - \bar{\mathbf{Y}}) (\mathbf{Y} - \bar{\mathbf{Y}})^T \right\} \right) \quad (\text{F-2})$$

where,

$$\begin{aligned} \mathbf{Y} - \bar{\mathbf{Y}} &= \mathbf{A}_U \left(\sum_s c_s \mathbf{C}_s^B \right) \mathbf{A}_V^T - E \left\{ \mathbf{A}_U \left(\sum_s c_s \mathbf{C}_s^B \right) \mathbf{A}_V^T \right\} \\ &= \mathbf{A}_U \left(\sum_s (c_s - \bar{c}_s) \mathbf{C}_s^B \right) \mathbf{A}_V^T \end{aligned} \quad (\text{F-3})$$

with $\bar{c}_s = E\{c_s\}$. Therefore,

$$\begin{aligned} (\mathbf{Y} - \bar{\mathbf{Y}}) (\mathbf{Y} - \bar{\mathbf{Y}})^T &= \\ \mathbf{A}_U \left(\sum_s (c_s - \bar{c}_s) \mathbf{C}_s^B \right) \mathbf{A}_V^T \mathbf{A}_V \left(\sum_{s'} (c_{s'} - \bar{c}_{s'}) (\mathbf{C}_{s'}^B)^T \right) \mathbf{A}_U^T \end{aligned} \quad (\text{F-4})$$

Thus, for a given \mathbf{A}_V , $\Psi_{\mathbf{Y}}$ is maximized if and only if $\text{trace}(\mathbf{A}_U \Upsilon_V \mathbf{A}_U^T)$ is maximized where,

$$\Upsilon_V = \sum_s \sum_{s'} E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} \mathbf{C}_s^B \mathbf{A}_V^T \mathbf{A}_V (\mathbf{C}_{s'}^B)^T \quad (\text{F-5})$$

Since \mathbf{A}_U and Υ_V are square matrices, this is equivalent to maximizing $\text{trace}(\mathbf{A}_U^T \Upsilon_V^T \mathbf{A}_U)$ which is obtained if and only if \mathbf{A}_U consists of the H' -eigenvectors of Υ_V^T .

Similarly, given \mathbf{A}_U , the weighting matrix \mathbf{A}_V maximizes the objective function given by,

$$\begin{aligned} \Psi'_Y &= E\{\|\mathbf{Y}^T - \bar{\mathbf{Y}}^T\|_F^2\} \\ &\quad \text{trace}\left(E\left\{(\mathbf{Y}^T - \bar{\mathbf{Y}}^T)(\mathbf{Y}^T - \bar{\mathbf{Y}}^T)^T\right\}\right) \end{aligned} \quad (\text{F-6})$$

where,

$$\begin{aligned} (\mathbf{Y}^T - \bar{\mathbf{Y}}^T)(\mathbf{Y}^T - \bar{\mathbf{Y}}^T)^T &= \\ \mathbf{A}_V \left(\sum_s (c_s - \bar{c}_s) (\mathbf{C}_s^B)^T \right) \mathbf{A}_U^T \mathbf{A}_U \left(\sum_{s'} (c_{s'} - \bar{c}_{s'}) \mathbf{C}_{s'}^B \right) \mathbf{A}_V^T \end{aligned} \quad (\text{F-7})$$

Thus, for a given \mathbf{A}_U , Ψ'_Y is maximized if and only if $\text{trace}(\mathbf{A}_V \Upsilon_U \mathbf{A}_V^T)$ is maximized where,

$$\Upsilon_U = \sum_s \sum_{s'} E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} (\mathbf{C}_s^B)^T \mathbf{A}_U^T \mathbf{A}_U \mathbf{C}_{s'}^B \quad (\text{F-8})$$

which is equivalent to maximizing $\text{trace}(\mathbf{A}_V^T \Upsilon_U^T \mathbf{A}_V)$ that is attained if and only if \mathbf{A}_V consists of the W' -eigenvectors of Υ_U^T .

In case of full projection, $H' = H$ and $W' = W$. In order not to abuse notation, let $(\cdot)^*$ denote matrices with the full dimensions H and W as opposed to $(\tilde{\cdot})$. Thus \mathbf{U}^* and \mathbf{V}^* become square matrices where $\mathbf{U}^{*-1} = \mathbf{U}^{*T}$ and $\mathbf{V}^{*-1} = \mathbf{V}^{*T}$. As a result,

$$\mathbf{A}_V^{*T} \mathbf{A}_V^* = (\mathbf{V}^{*T} \mathbf{V}_B^*)^T \mathbf{V}^{*T} \mathbf{V}_B^* = \mathbf{V}_B^{*T} \mathbf{V}^* \mathbf{V}^{*T} \mathbf{V}_B^* = \mathbf{I}_W \quad (\text{F-9})$$

where \mathbf{I}_W is an identity matrix of size $W \times W$. Similar argument for $\mathbf{A}_U^{*T} \mathbf{A}_U^* = \mathbf{I}_H$. Thus Υ_V^* and Υ_U^* becomes independent of the weighting matrices such that,

$$\Upsilon_V^* = \sum_s \sum_{s'} E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} \mathbf{C}_s^{*B} (\mathbf{C}_{s'}^{*B})^T \quad (\text{F-10})$$

and,

$$\Upsilon_U^* = \sum_s \sum_{s'} E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} (\mathbf{C}_s^{*B})^T \mathbf{C}_{s'}^{*B} \quad (\text{F-11})$$

The optimal \mathbf{A}_U^* and \mathbf{A}_V^* are then obtained, directly with no iterations, as the matrices containing the eigenvectors of Υ_V^{*T} and Υ_U^{*T} , respectively.

□

APPENDIX G: PROOF OF ANALYTIC PCA THEOREM

Proof. Consider the frequency-space representation of image irradiance in the vector space,

$$\mathbf{e} = \sum_s c_s \mathbf{b}_s = \tilde{\mathbf{W}}_B \sum_s c_s \mathbf{c}_s^B \quad (\text{G-1})$$

The weighting matrix \mathbf{A}_W maximizes the objective function given by,

$$\Psi_{\mathbf{y}} = E\{\|\mathbf{y} - \bar{\mathbf{y}}\|_F^2\} = \text{trace} \left(E \left\{ (\mathbf{y} - \bar{\mathbf{y}}) (\mathbf{y} - \bar{\mathbf{y}})^T \right\} \right) \quad (\text{G-2})$$

where,

$$\mathbf{y} = \tilde{\mathbf{W}}^T \mathbf{e} = \mathbf{A}_W \underbrace{\tilde{\mathbf{W}}_B^T \tilde{\mathbf{W}}_B}_{\mathbf{Id}_{D'}} \left(\sum_s c_s \mathbf{c}_s^B \right) \quad (\text{G-3})$$

and,

$$\bar{\mathbf{y}} = E\{\mathbf{y}\} = \mathbf{A}_W \left(\sum_s E\{c_s\} \mathbf{c}_s^B \right) \quad (\text{G-4})$$

Thus the objective function in Equation G-2 can be rewritten as,

$$\Psi_{\mathbf{y}} = \text{trace} \left(E \left\{ \mathbf{A}_W \Upsilon \mathbf{A}_W^T \right\} \right) \quad (\text{G-5})$$

where,

$$\Upsilon = \sum_s \sum_{s'} E\{(c_s - \bar{c}_s)(c_{s'} - \bar{c}_{s'})\} \mathbf{c}_s^B (\mathbf{c}_{s'}^B)^T \quad (\text{G-6})$$

Therefore, the maximum of $\Psi_{\mathbf{y}}$ is attained when the matrix \mathbf{A}_W contains the D' -eigenvectors of the matrix Υ^T . □

APPENDIX H: (HEMI)SPHERICAL FUNCTIONS

Definition

A unit sphere \mathcal{S}^2 in \mathbb{R}^3 is a 2-dimensional manifold which refer to the set of points whose Euclidean distance from the origin is unity, *i.e.* ,

$$\mathcal{S}^2 = \{\mathbf{x} \in \mathbb{R}^3 | \mathbf{x}^T \mathbf{x} = 1\} \quad (\text{H-1})$$

The unit sphere can be parameterized by the spherical polar coordinates. If $\mathbf{x} = (x, y, z)$ are the Cartesian coordinates in \mathbb{R}^3 , then the spherical polar coordinates can be defined by the polar angle $\theta \in [0, \pi]$ from the z -axis and the azimuthal angle $\phi \in [0, 2\pi]$ in the xy -plane. The distance from the origin of the coordinate system is defined by the radius r , where $r = 1$ in case of normalized coordinates which lie on a unit sphere, see Figure 10.1 for illustration. Cartesian coordinates are related to their spherical ones by,

$$\begin{aligned} x &= r \sin \theta \cos \phi \\ y &= r \sin \theta \sin \phi \\ z &= r \cos \theta \end{aligned} \quad (\text{H-2})$$

In case of hemispherical functions, the south-pole of the unit sphere is mapped to the xy -plane, where the domain of the polar angle is modified such that $\theta \in [0, \pi/2]$. Thus a unit hemisphere Ω in \mathbb{R}^3 is a 2-dimensional manifold defined as,

$$\Omega = \{\mathbf{x} \in \mathbb{R}^3 | \mathbf{x}^T \mathbf{x} = 1 \text{ and } z \in [0, 1]\} \quad (\text{H-3})$$

A *spherical function* f^s provides a mapping of points on the surface of a unit sphere \mathcal{S}^2 to a scalar value, it can be defined as $f^s \in L^2(\mathcal{S}^2)$ to be the Hilbert space¹ of square integrable² functions on the two dimensional sphere \mathcal{S}^2 , such that, assuming real-valued functions, $f^s : \mathcal{S}^2 \rightarrow \mathbb{R}$. Similarly, a *hemispherical function* f^h can be defined as a mapping of points on the surface of a unit hemisphere Ω to a real number, *i.e.* $f^h : \Omega \rightarrow \mathbb{R}$, where $f^h \in L^2(\Omega)$.




Figure 10.1: The spherical polar coordinate system: it is defined by the polar angle from the z -axis θ and the azimuthal angle ϕ in the xy -plane. The distance from the origin of the coordinate system is defined by the radius r

Visualization

A spherical function $f^s(\theta, \phi)$ can be visualized in one of two ways: (1) A textured unit sphere \mathcal{S}^2 can be displayed such that a point $\mathbf{x} = (x, y, z) = \text{cart}(1, \theta, \phi)$ on its surface has the intensity of $f^s(\theta, \phi)$, see Figure 10.2(b). (2) The points on the surface of a unit sphere can be deformed such that $\mathbf{x} = (x, y, z) = \text{cart}(|f^s(\theta, \phi)|, \theta, \phi)$, *i.e.* each point is scaled radially by the absolute value of the function, see Figure 10.2(a). Note that $\text{cart}(r, \theta, \phi)$ denotes the operator which converts spherical to cartesian coordinates, it is defined by Equation H-2. Similarly, a hemispherical function with $\theta \in [0, \pi/2]$ can be visualized as in Figure 10.3.

¹Hilbert space, a generalization of Euclidean space, is a vector space having the structure of inner product and allowing the measurement of length and angles.

²An L^2 -function defined over a domain \mathcal{D} such that $f : \mathcal{D} \rightarrow \mathbb{R}$ is called a square integrable function if $\int_{\mathcal{D}} |f|^2 d\mu < \infty$ *w.r.t.* some measure μ .




Figure 10.2: An example of spherical function, $f^s(\theta, \phi) = \frac{1}{4} [\cos^3(6\theta) + \sin^4(\phi) + 1]$ visualized as (a) a unit sphere being deformed such that the distance between the sphere's center and the surface points encodes the function absolute values at different spherical coordinates, this is accomplished by the relation between the spherical and cartesian coordinates, (b) a textured unit sphere where the color of its surface points encodes the function values. Note that $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi]$.




Figure 10.3: An example of hemispherical function, $f^h(\theta, \phi) = \frac{1}{4} [\cos^3(6\theta) + \sin^4(\phi) + 1]$ visualized as (a) a unit hemisphere being deformed such that the distance between the hemisphere's center and the surface points encodes the function absolute values at different spherical coordinates, this is accomplished by the relation between the spherical and cartesian coordinates, (b) a textured unit hemisphere where the color of its surface points encodes the function values. Note that $\theta \in [0, \pi/2]$ and $\phi \in [0, 2\pi]$.

Integration

The integration over a spherical domain can be viewed as summing infinitesimal patches on the surface of a unit sphere. Using spherical polar coordinates causes these patches to be bigger as the equator is approached and vanishes as towards the poles, such an effect can be encoded by defining the element of integration as $\sin \theta d\theta d\phi$. Thus an integral over a unit sphere can be explicitly written in terms of the spherical coordinates as,

$$\int_{S^2} f^s(\vec{\omega}) d\vec{\omega} = \int_0^{2\pi} \int_0^\pi f^s(\theta, \phi) \sin \theta d\theta d\phi \tag{H-4}$$

where $\vec{\omega}$ is a unit vector of polar angle θ and azimuthal angle ϕ . In a similar manner, an integral over the surface of a unit hemisphere can be defined as,

$$\int_{\Omega} f^{\mathfrak{h}}(\vec{\omega}) d\vec{\omega} = \int_0^{2\pi} \int_0^{\pi/2} f^{\mathfrak{h}}(\theta, \phi) \sin \theta d\theta d\phi \quad (\text{H-5})$$

In spherical coordinates, the *inner product* on the surface of the unit sphere is given by,

$$\langle f^{\mathfrak{s}}, g^{\mathfrak{s}} \rangle = \int_0^{2\pi} \int_0^{\pi} f^{\mathfrak{s}}(\theta, \phi) g^{\mathfrak{s}}(\theta, \phi) \sin \theta d\theta d\phi \quad (\text{H-6})$$

While the inner product on the surface of the unit hemisphere can be defined analogically as,

$$\langle f^{\mathfrak{h}}, g^{\mathfrak{h}} \rangle = \int_0^{2\pi} \int_0^{\pi/2} f^{\mathfrak{h}}(\theta, \phi) g^{\mathfrak{h}}(\theta, \phi) \sin \theta d\theta d\phi \quad (\text{H-7})$$

Note that the complex conjugate is dropped in case of real-valued functions.

Convolution

Consider the convolution on the real line; the convolution of two signals $f_1, f_2 : \mathbb{R} \rightarrow \mathbb{R}$ can be defined as,

$$f_1 * f_2(x) = \int_{-\infty}^{\infty} f_1(\tau) f_2(x - \tau) d\tau \quad (\text{H-8})$$

In the context of group representation theory, $f_1, f_2 \in L^2(\mathbb{R})$ are square integrable functions on the group of real numbers with addition as the group composition operator. In general, the convolution of two functions $f_1(g)$ and $f_2(g)$ on a general compact group \mathfrak{G} can be defined as [201],

$$f_1 * f_2(g) = \int_{\mathfrak{G}} f_1(\varsigma) f_2(\varsigma^{-1}g) d\varsigma \quad (\text{H-9})$$

where the convolution is not, in general, commutative, *i.e.* $f_1 * f_2(g) \neq f_2 * f_1(g)$.

Let $f_1, f_2 \in L^2(\mathcal{S}^2)$ be square integrable functions defined over the spherical domain, the (left) convolution operator for the function f_1 can be defined as sweeping the entire sphere by taking all rotations of the north pole $\eta = (0, 0, 1)$,

$$f_1 * f_2(\vec{\omega}) = \int_{R_{\alpha, \beta, \gamma} \in SO(3)} f_1(R_{\alpha, \beta, \gamma} \eta) f_2(R_{\alpha, \beta, \gamma}^{-1} \vec{\omega}) dR_{\alpha, \beta, \gamma} \quad (\text{H-10})$$

where $dR_{\alpha, \beta, \gamma} = \frac{1}{8\pi^2} \sin \alpha d\alpha d\beta d\gamma$ in terms of the Euler angle coordinates [202].

It worth mentioning that a popular method to evaluate spherical convolution is to first project a discrete version of the spherical functions onto the subspace spanned by spherical harmonics basis functions (to be discussed later) and then perform the convolution in the fourier domain via simple multiplication. Theoretical guarantees for such approach has been presented in [202] where the sampling theorem ensures that there is no aliasing caused by function discretization. Yet, such guarantee only holds for the latitudinal-longitudinal sampling (*i.e.* gridding) where samples are not evenly distributed over the sphere (more samples near the poles).

APPENDIX I: PROJECTION ONTO BASIS - MONTE CARLO INTEGRATION

Orthogonal Functions

A set of functions $\{B_k(x)\}$ of degree $k \geq 0$ are said to be *orthogonal* over an interval $[a, b]$ [203] if and only if $\forall k, k' \geq 0$,

$$\int_a^b w(x)B_k(x)B_{k'}(x)dx = c_k\delta_{kk'} = c_{k'}\delta_{kk'} \quad (\text{I-1})$$

for real-valued functions, where $c_k = c_{k'}$ is the constant of the orthogonality relation and $\delta_{kk'}$ denotes the Kronecker delta, defined by,

$$\delta_{kk'} = \begin{cases} 1 & k = k' \\ 0 & k \neq k' \end{cases} \quad (\text{I-2})$$

$w(x)$ is a non-negative weighting function. If $c_k = c_{k'} = 1 \forall k, k' \geq 0$, this set obeys the *orthonormality* criterion.

If the set $\{B_k(x)\}$ satisfies the orthogonality/orthonormality criterion, *i.e.* linearly independent, it is said to form an *orthogonal/orthonormal basis* which spans the space of functions defined over the interval $[a, b]$. Thus any piecewise continuous function f over the interval $[a, b]$ can be represented as a linear combination of an infinite series of these basis functions. This means that the basis functions B_k can be considered as small pieces of information, which can be scaled and combined to produce either the original function f or a band-limited approximation \hat{f} , depending on the number of basis used for function representation (*i.e.* infinite versus finite). Examples of basis are polynomial

basis, *e.g.* Legendre polynomial, and sinusoidal basis, *e.g.* Fourier basis.

In contrast to non-orthonormal basis, using orthonormal ones for function representation has two-fold benefits [30]: projecting the functions onto the basis becomes easier and the integration of the functions product in the basis can be computed efficiently. This can be realized in the following.

In order to approximate a given function f , it is required to compute the expansion coefficients a_k which encodes how much a basis function B_k is similar to f . This can be accomplished by integrating the L^2 -inner product of the function and the basis over the full domain of f . This process is called *projection* or *expansion*.

$$a_k = \int_a^b f(x)B_k(x)dx \quad (\text{I-3})$$

The expansion coefficients can be used to scale the corresponding basis in a linear combination to *reconstruct* either the original function or a band-limited version of the function where a truncated series of basis is used.

$$f(x) = \sum_{k=0}^{\infty} a_k B_k(x) \quad \text{or} \quad \hat{f}(x) = \sum_{k=0}^K a_k B_k(x) \quad (\text{I-4})$$

If two functions $f(x)$ and $g(x)$ are represented into the basis $\{B_k\}$ by two sets of coefficients $\{a_k\}$ and $\{b_k\}$ respectively. Efficient integration can be achieved by taking advantage of the orthonormality criterion in Equation I-1 with $w(x) = 1$ and $c_k = 1$. Hence the integral of the two functions in the basis can be evaluated by simple inner product of their respective coefficients.

$$\begin{aligned} \int_a^b f(x)g(x)dx &= \int_a^b \left(\sum_k a_k B_k(x) \right) \left(\sum_{k'} b_{k'} B_{k'}(x) \right) dx \\ &= \sum_k \sum_{k'} a_k b_{k'} \int_a^b B_k(x) B_{k'}(x) dx \\ &= \sum_k a_k b_k \end{aligned} \quad (\text{I-5})$$

In case of *orthonormal*¹ basis functions, considering the K th order function approximation, the first $(K + 1)$ terms are of interest, thus such approximation captures a certain amount of the function's total energy defined as the cumulative sum of energies

¹Orthogonal basis functions can be normalized to be orthonormal by scaling each basis by a normalization factor given by $\sqrt{\frac{1}{c_k}}$

maintained by individual series terms, (this cumulative energy can also be referred to as *function approximation accuracy* since the more energy the expansion terms capture from the function to be approximated, the more accurate the expansion is). Thus the approximation accuracy can be defined as¹,

$$e^K = \frac{\sum_{k=0}^K a_k^2}{\int_a^b |f(x)|^2 dx} \quad (\text{I-6})$$

In the following subsections, the orthogonal/orthonormal polynomials are investigated which are used to form the (hemi)spherical orthonormal basis whose definition is divided into polar and azimuthal parts, where (shifted) associated Legendre polynomials are used to define the polar part while sine and cosine functions are used to define the azimuthal part.

It is important to note that the integral in Equation I-3 is defined in the continuous domain. Approaches to evaluate such integral can be categorized in two broad classes according to the form of the function at hand. In case of having an analytical form of the function $f(x)$, Monte Carlo based methods can be used which utilize the power of randomness to evaluate the expected value of the integral. While, on the other hand, if the function values are only known at specific sample points, *i.e.* scattered data, least-squares based methods are available to find the best estimate for the expansion coefficients, defining the series Equation I-4 which passes near to the sampling points in the least squares sense. This dissertation is interested in integrals which involve analytic forms, as such in this appendix, the Monte Carlo integration process is reviewed. But first an example which demonstrate projecting a 1D function onto the subspace spanned by Legendre polynomials functions is presented.

¹This can be proved by representing the function f as a linear combination of the basis functions B_k and embed this into the integral $\int_a^b |f(x)|^2 dx$ where the orthonormality property can be used to find a closed form for the integral.

Example: Legendre Polynomials

Most of the special functions of mathematical physics can be generated/defined by a generating function¹ of the form [204],

$$g(t, x) = \sum_k B_k(x)t^k \quad (\text{I-7})$$

Consider Legendre polynomials $P_n(x)$ which construct a complete, orthogonal set of functions over the interval $[-1, 1]$, whose generating function is given by [204],

$$g(t, x) = (1 - 2xt + t^2)^{-1/2} = \sum_{n=0}^{\infty} P_n(x)t^n, \quad |t| < 1 \quad (\text{I-8})$$

where n denotes the polynomial order. These polynomials can be explicitly defined in terms of a power series defined as [204],

$$P_n(x) = \sum_{k=0}^{\lfloor n/2 \rfloor} (-1)^k \frac{(2n - 2k)!}{2^n k! (n - k)! (n - 2k)!} x^{n-2k} \quad (\text{I-9})$$

Starting with $P_0(x) = 1$ and $P_1(x) = x$, a three-term recurrence relation can be used to efficiently evaluate such polynomials² instead of direct evaluation of the series in Equation I-9 (Table 10.1 gives the explicit formula of Legendre polynomials up to the eighth order)[204],

$$(n + 1)P_{n+1}(x) = (2n + 1)xP_n(x) - nP_{n-1}(x) \quad n = 1, 2, 3, \dots \quad (\text{I-10})$$

Legendre polynomials satisfy the orthogonality condition with weighting function $w(x) = 1$, over the interval $x \in [-1, 1]$, such that,

$$\int_{-1}^1 P_n(x)P_{n'}(x)dx = \frac{2}{2n + 1} \delta_{nn'} \quad (\text{I-11})$$

This property can be proved using the Rodrigues' formula of Legendre polynomials

¹A generating function is a formal power series in one indeterminate, whose coefficients construct a set of orthogonal basis functions.

²Recurrence relations are usually used to evaluate series-defined functions due to their numerical robustness and linearity in the total number of basis functions [205]. On the other hand, direct evaluation of functions' series can suffer from catastrophic floating-point cancellation.

Table 10.1: Legendre Polynomials up to 8-th order

$P_0(x) = 1$
$P_1(x) = x$
$P_2(x) = \frac{1}{2}(3x^2 - 1)$
$P_3(x) = \frac{1}{2}(5x^3 - 3x)$
$P_4(x) = \frac{1}{8}(35x^4 - 30x^2 + 3)$
$P_5(x) = \frac{1}{8}(63x^5 - 70x^3 + 15x)$
$P_6(x) = \frac{1}{16}(231x^6 - 315x^4 + 105x^2 - 5)$
$P_7(x) = \frac{1}{16}(429x^7 - 693x^5 + 315x^3 - 35x)$
$P_8(x) = \frac{1}{128}(6435x^8 - 12012x^6 + 6930x^4 - 1260x^2 + 35)$

defined as [204],

$$P_n(x) = \frac{1}{2^n n!} \frac{d^n}{dx^n} (x^2 - 1)^n \quad (\text{I-12})$$

According to Equation I-11, the Legendre polynomials can be normalized to construct orthonormal basis for function approximation as follows,

$$p_n(x) = \sqrt{\frac{2n+1}{2}} P_n(x) \quad (\text{I-13})$$

In order to comprehend the projection/reconstruction process, let's consider the function $f(x) = 1.75x^3 + 0.6x^2 + 0.21x - 0.06$ for $x \in [-1, 1]$. Since it is a third-degree polynomial, the normalized Legendre polynomials up to third-order are used (see Table 10.1). Figure 10.4 shows the projection process where the L^2 -inner product of the function f and each normalized Legendre polynomial $p_n(x)$ is integrated over the domain of f to obtain the expansion coefficients a_n , which were used to scale such polynomials to reconstruct the original function (see Figure 10.5).

Figure 10.6 shows the energy captured by individual expansion coefficient, which is the square of the corresponding coefficient divided by the total energy maintained by the function to be approximated. The function approximation accuracy e^N is also shown to reach to 100% at third-degree function approximation.




Figure 10.4: The projection of $f(x) = 1.75x^3 + 0.6x^2 + 0.21x - 0.06$ in red onto normalized Legendre polynomials up to the third order, where the expansion/projection coefficients are evaluated using Monte Carlo integration.

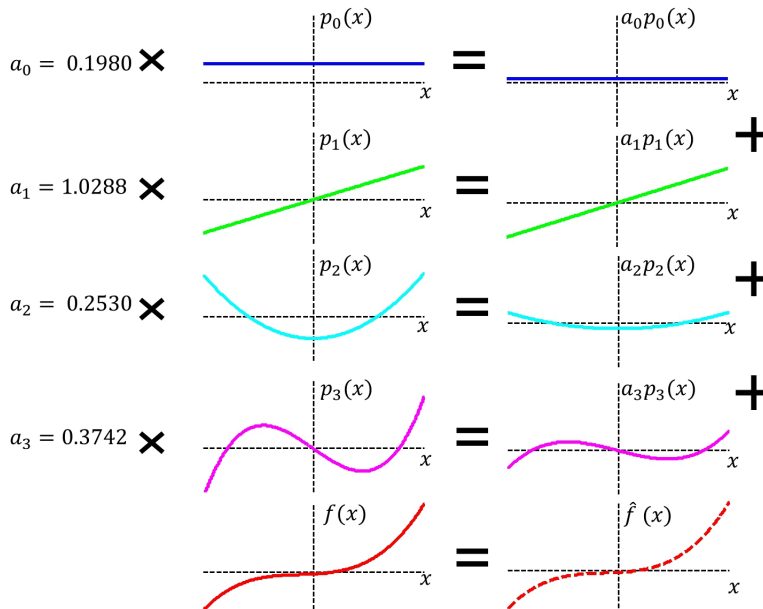


Figure 10.5: The reconstruction of $f(x) = 1.75x^3 + 0.6x^2 + 0.21x - 0.06$ (dashed red) using the normalized Legendre polynomials up to the third order.

Monte Carlo Integration

Finding an analytic form for the integral in Equation I-3 is a challenge or even worse when the function is not continuous. Consider the 1D case where the integral $\int_a^b f(x)dx$ is needed to be evaluated, one way to go around this is to equally divide the interval $[a, b]$ into J intervals with width $(b - a)/J$ and summate the areas covered by the function at each interval to approximate the integral value. This is called the *quadrature integration*.




Figure 10.6: Left: the amount of energy captured by each individual expansion coefficient, evaluated as $\frac{a_n^2}{\int_{-1}^1 |f(x)|^2 dx}$. Right: function approximation accuracy e^N of N -order function approximation using normalized Legendre polynomials, where $e^0 = 0.0301$, $e^1 = 0.8433$, $e^2 = 0.8925$ and $e^3 = 1.00$.

In general, the error decreases rapidly as J increases, however this is not the case when it comes to higher dimension.

Monte Carlo integration becomes advantageous at higher dimensions where the error bound is always function of the number of intervals/samples regardless of dimensionality, in addition to its robustness in case of discontinuities. It uses the power of randomness to sample the domain of integration according to a probability density function¹, where the expected value of the Monte Carlo estimate converges to the actual value of the integral.

The axioms of probability imply that the probability density function $p(x)$ should have non-negative values such that,

$$\int_{-\infty}^{\infty} p(x) dx = 1 \quad \text{s. t.} \quad p(x) \geq 0 \quad (\text{I-14})$$

The expected value of a function of a random variable can be defined as,

$$E\{f(x)\} = \int_{\mathfrak{D}} f(x)p(x) dx \quad (\text{I-15})$$

According to the *Law of Large Numbers*, the mean of a large number of random samples J from the function $f(x)$ converges towards the expected value of the function as $J \rightarrow \infty$, *i.e.* ,

$$E\{f(x)\} \approx \frac{1}{J} \sum_{j=0}^{J-1} f(x_j) \quad (\text{I-16})$$

¹Probability Density Function (PDF) $p(x)$ is the possibility of a specific outcome of a random variable X .

Using Equation I-15 and Equation I-16, a Monte Carlo estimator for the function integral can be defined as follows,

$$\begin{aligned} \int_{\mathfrak{D}} f(x)dx &= \int_{\mathfrak{D}} \frac{f(x)}{p(x)}p(x)dx = E \left\{ \frac{f(x)}{p(x)} \right\} \\ &\approx \frac{1}{J} \sum_{j=0}^{J-1} \frac{f(x_j)}{p(x_j)} = \frac{1}{J} \sum_{j=0}^{J-1} w(x_j)f(x_j) \end{aligned} \quad (\text{I-17})$$

In case of spherical functions, where the integration is over the surface of a unit sphere, the integration domain can be uniformly sampled such that,

$$\int_{S^2} p(\vec{\omega})d\vec{\omega} = c \int_{S^2} d\vec{\omega} = 1 \quad (\text{I-18})$$

Thus $p(\vec{\omega}) = c = 1/4\pi$, where $\int_{S^2} d\vec{\omega} = 4\pi$. Consider two independent random variables ζ_1 and ζ_2 which follow uniform distribution $U[0, 1]$, their cartesian product $U[0, 1] \times U[0, 1]$ is mapped to the spherical coordinates using the following transformation,

$$\begin{aligned} \theta &= 2 \cos^{-1} \left(\sqrt{1 - \zeta_1} \right) \\ \phi &= 2\pi\zeta_2 \end{aligned} \quad (\text{I-19})$$

In order to guarantee a uniform sampling of the unit sphere, stratified sampling [206] is used where the domain $[0, 1] \times [0, 1]$ is divided into $J \times J$ sample cells where a random sample is picked inside each cell. Thus the spherical integration can be evaluated as,

$$a_k^s \approx \frac{4\pi}{J^2} \sum_{j=0}^{J^2-1} f^s(\theta_j, \phi_j) B_k^s(\theta_j, \phi_j) \quad (\text{I-20})$$

where a_k^s is the expansion coefficient encoding how much a spherical basis function B_k^s is similar to f where both have been evaluated at the randomly picked samples $\{\theta_j, \phi_j\}$.

In case of hemispherical functions, the integration domain (unit hemisphere) can be

uniformly sampled such that,

$$\int_{\Omega} p(\vec{\omega}) d\vec{\omega} = c \int_{\Omega} d\vec{\omega} = 1 \quad (\text{I-21})$$

Thus $p(\vec{\omega}) = c = 1/2\pi$, where $\int_{\Omega} d\vec{\omega} = 2\pi$. Consider two independent random variables ζ_1 and ζ_2 which follow uniform distribution $U[0, 1]$, their cartesian product $U[0, 1] \times U[0, 1]$ is mapped to the hemispherical coordinates using the following transformation,

$$\begin{aligned} \theta &= \cos^{-1} \left(\sqrt{1 - \zeta_1} \right) \\ \phi &= 2\pi\zeta_2 \end{aligned} \quad (\text{I-22})$$

Using stratified sampling, the hemispherical integral can be evaluated as,

$$a_k^{\mathfrak{h}} \approx \frac{2\pi}{J^2} \sum_{j=0}^{J^2-1} f^{\mathfrak{h}}(\theta_j, \phi_j) f_k^{\mathfrak{h}}(\theta_j, \phi_j) \quad (\text{I-23})$$

where $a_k^{\mathfrak{h}}$ is the expansion coefficient encoding how much a hemispherical basis function $B_k^{\mathfrak{h}}$ is similar to f where both have been evaluated at the randomly picked samples $\{\theta_j, \phi_j\}$.

NOMENCLATURE

Greek Symbols

(α, β, γ)	The angular coordinates of the surface orientation. Noting that γ is only important for anisotropic surfaces, which controls the rotation of the local tangential plane about the surface normal.
(r, θ, ϕ)	Spherical polar coordinates, where r is the radial distance from the origin of the coordinate system, θ is the polar angle from the z -axis and ϕ is the azimuthal angle from the x -axis in the xy -plane.
ϵ	A sparse error vector occurred due to some contiguous irradiance corruption.
$\delta_{kk'}$	The Kronecker delta where $\delta_{kk'} = 1 \iff k = k'$.
ϵ	Radius of a small sphere.
η	Surface refractive index.
κ_1, κ_2	Scalars used in the process of linearly transforming the domain of a function/polynomial.
Φ	radiant flux/power, measured in <i>Watt</i> (W) having the dimension of energy per unit time or <i>joules/second</i> .
Ω	Surface of a unit hemisphere.
Ω'	A unit hemisphere of directions, located at surface point \mathbf{x} and oriented by the surface normal at this point, where primed coordinates are used to denote local reference frame. Subscripts i and o are used to refer to incoming and outgoing hemispheres, respectively.
Ω'_o	Outgoing upper hemisphere.

ω_s	Distant light source direction defined in the global reference frame.
ϕ'	The azimuth/tilt angle of the orthogonal projection of a light ray $\vec{\omega}'$ arriving to or leaving from the surface point \mathbf{x} on the surface tangent plane at this point measured from a fixed reference direction on that plane <i>e.g.</i> \vec{t} , where primed coordinates are used denote local reference frame. Subscripts i and o are used to refer to incoming and outgoing angles, respectively.
$\Psi_{\mathbf{Y}}$	The total matrix scatter.
$\rho(\mathbf{x})$	The surface albedo/texture at surface point \mathbf{x} .
ρ_d	Diffuse reflection or surface albedo.
σ	Scale parameter.
θ'	The inclination/slant angle of a light ray $\vec{\omega}'$ arriving to or leaving from the surface point \mathbf{x} , measured from the surface normal \vec{n} , where primed coordinates are used denote local reference frame. Subscripts i and o are used to refer to incoming and outgoing angles, respectively.
θ_j, ϕ_j	Sampling points in the (hemi)spherical domain with $0 \leq j < J$.
s_n	The integral of the HSH polar part of order n over the incoming or outgoing upper hemisphere.
$\vec{\omega}$	A unit vector of polar angle θ and azimuthal angle ϕ , <i>i.e.</i> $\vec{\omega} = \text{cart}(1, \theta, \phi)$.
$\vec{\omega}'$	A light ray (direction) arriving to or leaving from the surface point \mathbf{x} , where primed coordinates are used to denote local reference frame. Subscripts i and o are used to refer to incoming and outgoing light rays, respectively.
$\vec{\omega}'_h$	The half-way vector.
$\vec{\omega}'_i = (\theta'_i, \phi'_i)$	Incident light ray, measured with respect to the local coordinate frame of a surface point \mathbf{x} , where $\theta'_i \in [0, \pi/2]$ and $\phi'_i \in [0, 2\pi]$.
$\vec{\omega}'_o = (\theta'_o, \phi'_o)$	Outgoing/reflected light ray, measured with respect to the local coordinate frame of a surface point \mathbf{x} , where $\theta'_o \in [0, \pi/2]$ and $\phi'_o \in [0, 2\pi]$.

$\vec{\omega}'_s$	A unit vector defining the point light direction as perceived locally from the surface patch centered at \mathbf{x} .
$\vec{\omega}'_s$	The ideal specular reflection.
ζ	A parameter which influences the interaction between neighboring pixels.
ζ_1, ζ_2, \dots	Independent random variables following uniform distribution $U[0, 1]$.

Roman Symbols

(x, y, z)	Cartesian coordinates in the three-dimensional Euclidean space \mathbb{R}^3 .
$\bar{\mathbf{a}}$	The average albedo vector.
$\bar{\mathbf{h}}$	The average harmonic projection vector.
$\bar{\mathbf{s}}$	The average shape vector.
λ_a	A vector containing the albedo coefficients.
λ_h	A vector containing the harmonic projection coefficients.
λ_s	A vector containing the shape coefficients.
\hat{f}	A band-limited function approximation of a piecewise continuous function f defined over the interval $[a, b]$.
\mathbf{a}	A albedo vector obtained from the albedo model.
\mathbf{A}_k	The albedo information of data sample k being stacked into a vector.
$\mathbf{A}_U, \mathbf{A}_V$	The weighting matrices, in $\mathbb{R}^{H' \times H'}$ and $\mathbb{R}^{W' \times W'}$ respectively, encoding the relation of principal components spanning the row and column subspaces of the image space with those spanning the row and column subspaces of the irradiance harmonics.
\mathbf{A}_W	The weighting matrix which relates the principal components of the image space to that of the irradiance harmonics in the vector space.
\mathbf{B}_s	The matrix representation of the s -th irradiance harmonics of the visible surface normals.

\mathbf{b}_s	The vector representation of the s -th image irradiance harmonics \mathbf{B}_s .
\mathbf{c}	Subspace coefficients
\mathbf{C}_s^B	A lower-dimensional representation in $\mathbb{R}^{H' \times W'}$ of the s -th irradiance harmonics \mathbf{B}_s .
\mathbf{c}_s^B	Lower-dimensional representation of the image irradiance harmonics vector \mathbf{b}_s .
\mathbf{E}	A matrix representation in $\mathbb{R}^{H \times W}$ of the image irradiance of the visible surface normals to the viewer such that H denotes height and W denotes width.
\mathbf{e}	The vector in \mathbb{R}^D representation of the image irradiance \mathbf{E} where $D = H \times W$.
\mathbf{h}	A harmonic projection vector obtained from the harmonic projection model.
\mathbf{h}	Harmonic projection irradiance vector
\mathbf{H}_k	The harmonic projection irradiance of data sample k being stacked into a vector.
\mathbf{P}_a	A matrix having its columns as the albedo eigenvectors.
\mathbf{P}_h	A matrix having its columns as the harmonic projection eigenvectors.
\mathbf{P}_s	A matrix having its columns as the shape eigenvectors.
\mathbf{s}	A shape vector obtained from the shape model.
\mathbf{S}_k	The shape information of data sample k being stacked into a vector.
$\mathbf{U}, \mathbf{S}, \mathbf{V}$	The SVD decomposition of a given matrix.
\mathbf{x}	Surface point (spatial position) $\in \mathbb{R}^3$.
\mathbf{Y}	A lower-dimensional representation in $\mathbb{R}^{H' \times W'}$ of the image irradiance \mathbf{E} .
\mathbf{y}	Lower-dimensional representation of the image vector \mathbf{e} .

\mathcal{G}	A general compact group.
\mathcal{P}	The set of D -pixels representing the error vector.
\mathcal{S}'	Local spherical integral domain.
\mathcal{S}^2	The surface of a unit sphere in \mathbb{R}^3 .
\mathcal{T}	The family of the neighboring pixel pairs supporting the Gibbs potentials.
\mathbb{D}'	Local integral domain.
\mathbb{R}^M	M -dimensional Euclidean space.
\mathcal{D}	General domain of integration.
\mathfrak{N}	A geometric neighborhood system.
\mathbf{f}	A warping function for triangular meshes.
$\tilde{\mathbf{U}}, \tilde{\mathbf{V}}$	Projection matrices (linear transformations) in $\mathbb{R}^{H \times H'}$ and $\mathbb{R}^{W \times W'}$, respectively, which map the image space $\mathbb{R}^{H \times W}$ into a lower-dimensional subspace $\mathbb{R}^{H' \times W'}$ with $H' \leq H$ and $W' \leq W$.
$\tilde{\mathbf{U}}_B, \tilde{\mathbf{V}}_B$	Projection matrices (linear transformations) in $\mathbb{R}^{H \times H'}$ and $\mathbb{R}^{W \times W'}$, respectively, which map the irradiance harmonics space $\mathbb{R}^{H \times W}$ into a lower-dimensional subspace $\mathbb{R}^{H' \times W'}$ with $H' \leq H$ and $W' \leq W$.
$\tilde{\mathbf{W}}$	Orthonormal projection matrix in $\mathbb{R}^{D \times D'}$ which maps the original vector space \mathbb{R}^D into a vector subspace $\mathbb{R}^{D'}$ with $D' \leq D$.
$\tilde{\mathbf{W}}_B$	D' -eigenvectors of $\mathbf{B} = [\mathbf{b}_1 \dots \mathbf{b}_S]$ corresponding to its largest D' -eigenvalues.
$\vec{b}(\mathbf{x})$	Surface binomial vector at surface point \mathbf{x} , corresponding to x' -axis, where primed coordinates are used to denote local reference frame, $\vec{b}(\mathbf{x}) = \vec{n}(\mathbf{x}) \times \vec{t}(\mathbf{x})$.
$\vec{n}(\mathbf{x})$	Surface normal (unit) vector at surface point \mathbf{x} , corresponding to z' -axis, where primed coordinates are used to denote local reference frame, where $\vec{n}(\mathbf{x}) = (n_x, n_y, n_z)$.

\vec{s}	A vector from the surface point \mathbf{x} and the light source .
$\vec{t}(\mathbf{x})$	Surface tangent vector at surface point \mathbf{x} , corresponding to y' -axis, where primed coordinates are used denote to local reference frame.
$\hat{\mathbf{c}}$	<i>Estimated</i> subspace coefficients
$\{\mathbf{p}, \mathbf{q}\}$	Neighboring pixels in a geometric neighborhood system \mathfrak{N} .
a_k	Expansion/projection coefficients used to approximate/reconstruct a function f .
a_{np}^q	The expansion/projection coefficients of the <i>isotropic</i> reflectance kernel represented using the Helmholtz HSH-based basis. It is also called the isotropic bidirectional surface reflectance spectrum (IBSRS).
a_{pr}^L	The isotropic surface reflectance spectrum of Lambertian reflectance kernel.
a_{pr}^{qs}	The expansion/projection coefficients of the reflectance kernel represented using the Helmholtz HSH-based basis. It is also called the bidirectional surface reflectance spectrum (BSRS).
$a_{pr}^{S,qs}$	The spectral components (expansion coefficients) of the mirror-like reflectance kernel.
c_k	The constant of the orthogonality relation of orthogonal set of functions $B_k(x)$.
$c_s = l_n^m a_{pr}^q$	The coefficient of the s -th image irradiance harmonics.
c_s^t	The estimated s -th coefficient at iteration t .
e^K	Function approximation accuracy (<i>i.e.</i> amount of energy captured by the expansion/projection coefficients) using K -degree/order approximation of a function f .
e_B	The BRDF energy content.
e_L	The illumination signal energy content.

f	An arbitrary one dimensional piecewise continuous function defined over the interval $[a, b]$.
H', W'	Height and width of a reduced representation for an image, respectively.
H, W	Height and width of an image, respectively.
I_o	radiant intensity, the radiant power per unit solid angle, measured in $W.sr^{-1}$.
J	Number of random samples picked in a 1D space.
K	Maximum degree/order/number-of-terms used to approximate a function f as a linear combination of basis functions.
l_n^m	The illumination spectrum coefficients when projected onto the subspace spanned by the real spherical harmonics basis.
r	Distance from the origin of a (hemi)sphere and sometimes refer to the distance between a surface patch and a light source.
x_j	j -th sample drawn randomly from the function $f(x)$.
\mathcal{F}	A set of F -triangular faces of a 2-manifold triangular mesh
\mathcal{G}	2-manifold triangular mesh
\mathcal{V}	A set of V -vertices with $\mathbf{x}_v \in \mathbb{R}^3$ of a 2-manifold triangular mesh

Other Symbols

$E\{f(x)\}$	The expected value of a function of a random variable X .
$\langle f, g \rangle$	The inner product of two square integrable functions, <i>i.e.</i> L^2 -integrable, defined as the integral for their product over their domain.
$\{B_k(x)\}$	A set of orthogonal functions of degree $k \geq 0$.
$\mathbf{f}(\mathbf{x})$	The spatial support of the error ϵ .
$\mathcal{A}_B(n)$	The attenuation factors imposed by the surface BRDF on the illumination order n .

$\mathcal{B}_s(\alpha, \beta)$	The single-index notation of image irradiance harmonics \mathcal{B}_{npr}^{mq} where the index s is given by a sorting/ordering function.
$\mathcal{B}_{npr}^{mq}(\alpha, \beta)$	Image irradiance harmonics at surface point whose normal vector has spherical coordinates (α, β) .
$\mathcal{D}_{mm'}^n$	Wigner's rotation matrices for spherical harmonics.
$\mathcal{H}_{np}^q(\theta'_i, \theta'_o, \Delta\phi')$	Helmholtz HSH-based basis for <i>isotropic</i> reflectance.
$\mathcal{H}_{pr}^{qs}(\vec{\omega}'_i, \vec{\omega}'_o)$	Proposed Helmholtz HSH-based basis.
$\mathcal{R}(\vec{n}(\mathbf{x}))$	The surface radiance.
$\mathfrak{F}_{\text{neq}}(\mathbf{f})$	The relative frequency of the non-equal labels in pixel pair of \mathcal{T} .
cart	The operator which converts spherical coordinates to their corresponding cartesian ones, such that $(x, y, z) = \text{cart}(r, \theta, \phi)$.
sph	The operator which converts cartesian coordinates to their corresponding spherical ones, such that $(r, \theta, \phi) = \text{sph}(x, y, z)$.
$\Phi_m(\phi)$	The azimuthal part of the real (hemi)spherical harmonics.
$\rho(x)$	Robust estimate of the error norm.
$\Theta_n^m(\theta)$	The polar part of the real spherical harmonics.
$\tilde{\Theta}_n^m(\theta)$	The polar part of the real hemispherical harmonics.
\tilde{N}_n^m	The normalization factor of the hemispherical harmonics.
$\tilde{P}_n^m(x)$	Shifted associated Legendre polynomial of degree m and order/band-index n , where $n \geq 0$ and $m \in [0, n]$, defined over the interval $x \in [0, 1]$.
$\text{Acc}_B(P)$	The approximation accuracy of BRDF representation up to reflectance order P .
$\text{Acc}_E(N)$	The approximation accuracy of the N th order expansion of the image irradiance.
$\text{Acc}_L(N)$	The illumination representation approximation accuracy up to illumination order N .

$D(\theta'_h; \sigma)$	The facet slope distribution function.
$d\Phi_i(\mathbf{x}_i, \vec{\omega}'_i)$	Differential incident flux.
$d\vec{\omega}'$	The differential solid angle representing the angular size of a light ray $\vec{\omega}'$ as well as its direction. Subscripts i and o are used to refer to incoming and outgoing solid angles, respectively.
$d\vec{\omega}'_i, d\vec{\omega}'_o$	Differential solid angles of incident and outgoing/reflected light rays respectively.
$d^3E_{\vec{\omega}'}$	The energy transmitted by a patch/source into an infinitesimal region of solid angle $d\vec{\omega}'$ per unit time per unit foreshortened area.
d^n	A matrix defining the y -axis rotation.
dA	Area of an arbitrary small planar surface patch.
$dE_i(\mathbf{x}, \vec{\omega}'_i)$	Surface irradiance (incident), measured in $W.m^{-2}$, defined according to a differential solid angle $d\vec{\omega}'_i$.
$dL_r(\mathbf{x}_o, \vec{\omega}'_o)$	Differential reflected radiance, <i>i.e.</i> surface radiance.
$E(\mathbf{x}, \vec{\omega}'_o)$	Image irradiance which is proportional to the surface total radiance.
$E\{e_L\}$	The average illumination signal energy content.
$E_\chi(\mathbf{x})$	Light source exitance which measures the internally generated power radiated by a light source per unit area on the radiating surface.
$e_L^\mu(n)$	The average energy content maintained by the n -th illumination order.
$E_M(\mathbf{f})$	The maximum likelihood energy function.
$E_R(\mathbf{c})$	Robust estimation objective function.
$F(c; n)$	Fresnel formula for dielectrics.
f^h	Hemispherical function, $f^h : \Omega \rightarrow \mathbb{R}$.
f^s	Spherical function, $f^s : \mathbb{S}^2 \rightarrow \mathbb{R}$.
$f_r(\mathbf{x}, \vec{\omega}'_i, \vec{\omega}'_o)$	Bidirectional Reflectance Distribution Function.

$f_r(\vec{\omega}'_i, \vec{\omega}'_o)$	BRDF of homogeneous surface, where the surface has the same BRDF everywhere, thus spatial dependency is dropped.
$f_r^S(\vec{\omega}'_i, \vec{\omega}'_o)$	The BRDF of ideal specular surface.
$f_r^L(\vec{\omega}'_i, \vec{\omega}'_o)$	The BRDF of a Lambertian surface.
$f_r^{ON}(\vec{\omega}'_i, \vec{\omega}'_o)$	The Oren-Nayar BRDF model.
$f_r^{TS}(\vec{\omega}'_i, \vec{\omega}'_o)$	The Torrance-Sparrow (TS) BRDF.
$f_r^{WON}(\vec{\omega}'_i, \vec{\omega}'_o)$	The Wolff-Oren-Nayar BRDF model.
$f_r^W(\vec{\omega}'_i, \vec{\omega}'_o)$	The Wolff BRDF model.
$G(\vec{\omega}'_i, \vec{\omega}'_o)$	The geometric attenuation factor.
$g(t, x)$	Generating function used to define most of the special function of mathematical physics.
$H_n^m(\theta, \phi)$	The real-form of the hemispherical harmonics with n and m denote the order (band-index) and degree, respectively.
$L(\mathbf{x}, \vec{\omega}')$	Radiance which is the amount of energy traveling at some point \mathbf{x} in a specific direction $\vec{\omega}'$, per unit time, per unit foreshortened area, per unit solid angle, measured in $Wm^{-2}sr^{-1}$. Subscripts i and r are used to refer to incident and reflected radiance, respectively.
$L_i(\mathbf{x}, \vec{\omega}'_i)$	The incident radiance to the surface at point \mathbf{x} from the incident direction $\vec{\omega}'_i$.
$L_r(\mathbf{x}, \vec{\omega}'_o)$	Total surface radiance at specific local outgoing direction.
$L_r^S(\theta'_o, \phi'_o)$	Specular surface radiance.
N_n^m	The normalization factor of the spherical harmonics.
N_{np}^q	The normalization factor of the Helmholtz HSH-based basis for <i>isotropic</i> reflectance.
N_{pr}^{qs}	The normalization factor of the Helmholtz HSH-based basis.
$p(\alpha, \beta)$	The surface normal distribution function.

$p(\vec{\omega})$	Probability density function of a variable defined in the spherical domain.
$p(x)$	Probability Density Function (PDF) of a random variable X .
$P\{\mathcal{B}_{npr}^{mq}\}$	The average power content (i.e. second moment) of the irradiance harmonics.
$P_n(x)$	Legendre polynomial of order $n \geq 0$ where $x \in [-1, 1]$.
$p_n(x)$	Normalized Legendre polynomial of degree n used to construct orthonormal basis for function representation/approximation.
$P_n^m(x)$	Associated Legendre polynomial of degree m and order/band-index n , where $n \geq 0$ and $m \in [0, n]$, defined over the interval $x \in [-1, 1]$.
R_Z, R_Y	The rotation matrices about the z -axis and y -axis, respectively.
$R_{\alpha, \beta, \gamma}$	Euler angle representation of 3D rotations with ZYZ-convention, corresponding to three consecutive counterclockwise rotations about the z -, y - and z - axis respectively.
$S(\mathbf{x}_i, \vec{\omega}'_i, \mathbf{x}_o, \vec{\omega}'_o)$	Bidirectional Scattering Surface Reflectance Distribution Function.
$SO(3)$	Special orthogonal group whose elements are 3×3 real orthogonal matrices of unit determinant parameterized by the Euler angles (α, β, γ) with $\alpha \in [0, \pi]$, $\beta \in [0, 2\pi]$ and $\gamma \in [0, 2\pi]$.
$t_f(\theta'_i)$	Foreshortening factor.
$V_{\mathbf{pq}}(f_{\mathbf{p}}, f_{\mathbf{q}})$	A two-pixel potential controlling the spatial interaction between neighboring pixels.
$w(x)$	A non-negative weighting function.
$w(x_j)$	Weighting function at a random sample x_j .
$Y_n^m(\theta, \phi)$	The real-form of the spherical harmonics with n and m denote the order (band-index) and degree, respectively.

Acronyms

$BSRS$	Bidirectional Surface Reflectance Spectrum
--------	--

<i>IBSRS</i>	Isotropic Bidirectional Surface Reflectance Spectrum
<i>sr</i>	steradian - the unit of the solid angle.
AAM	Active Appearance Models
ASM	Active Shape Model
BD-PCA	Bidirectional Principal Component Analysis
BRDF	Bidirectional Reflectance Distribution Function.
BSSRDF	Bidirectional Scattering Surface Reflectance Distribution Function.
CAD	Computer-Aided Design
CAM	Computer-Aided Manufacturing
CT	Computer Tomography
CUReT	Columbia-Utrecht Reflectance and Texture Database
EM	Expectation-Maximization algorithm
EM	Expectation-Maximization
FMM	Fast Multipole Method
GPA	Generalized Procrustes Analysis
HDR	High Dynamic Range
HP	Harmonic Projection
HSH	Hemispherical Harmonics
ICP	Iterative Closest Point
MAE	Mean Absolute Error
Merl	Mitsubishi Electric Research Laboratories
MGRF	Markov Gibbs Random Field
ON	Oren-Nayar
PCA	Principal Component Analysis

PCR	Principal Component Regression
PDE	Partial Differential Equation
PDM	Point Distribution Model
PLS	Partial Least Squares
RBF	Radial Basis Function
RMS	Root Mean Square
SFS	Shape From Shading
SH	Spherical Harmonics
sIBL	Smart Image-based Lighting
SSFS	Statistical Shape From Shading
SSS	Sample Sample Size
SVD	Singular Value Decomposition
TS	Torrance-Sparrow
USF Database	University of South Florida database
WON	Wolff-Oren-Nayar

CURRICULUM VITAE

Contact Information

Department of Electrical and Computer Engineering

University of Louisville

Mobile: +1-502-418-3292

Computer Vision and Image Processing Laboratory

2211 South Brook

E-mail:

Lutz Hall, Room #6, Louisville, KY 40292, USA

shireen.elhabian@louisville.edu

WWW: www.cvip.uofl.edu/

Objective

Pursuit of advanced research and development in computer vision, image processing, medical imaging and/or machine learning as a postdoctoral fellow.

Research Interests

My research/development interests lie in the fields of computer vision, image understanding and pattern recognition. My primary focus is on illumination modeling, geometric and photometric object representation, subspace learning, 3D shape reconstruction, shape-from-X, statistical shape analysis, object detection and recognition, medical imaging, surface registration and facial analysis under varying pose, illumination, occlusion and expression.

Education

[University of Louisville](#), Louisville, Kentucky, USA

Ph.D. Candidate, [Electrical and Computer Engineering](#), Dissertation Proposal Defense: Summer 2011. Expected to Graduate by Fall 2012.

- Thesis Topic: *Phenomenological Modeling of Image Irradiance for Non-Lambertian Surfaces Under Natural Illumination*
- Candidacy Written Exam: *Investigating the Interplay between Human Face Reconstruction and Recognition*
- GPA : 4 out of 4
- Advisor: [Professor Aly A. Farag](#)
- Thesis Committee:
 - Professor Aly A. Farag, ECE Department and Director of CVIP-Lab@UofL
 - Professor Thomas L. Starr, Associate Dean for Research
 - Professor James H. Graham, Chairman of ECE Department
 - Professor Robert W. Cohn, Director of ElectroOptics Research Institute and Nanotechnology Center
- Area of Study: Computer Vision and Image Processing

[Cairo University](#), Cairo, Egypt

M.Sc., [Faculty of Computers and Information](#), April 2005

- Thesis Topic: *Object Detection and Tracking in Video Sequences*
- Advisors: Professor Khaled Mostafa, Professor Sanaa ElOla Ahmed
- Area of Study: Image Processing

Pre-Masters Postgraduate Study, [Faculty of Computers and Information](#), May 2003

- Project Topic: *Landmine Detection in Egypt*
- Advisors: Professor Khaled Mostafa, Professor Sanaa ElOla Ahmed
- Area of Study: Image Processing
- GPA: 4 out of 4 (Excellent With Honor)

B.Sc., [Faculty of Computers and Information](#), May 2002

- GPA: 3.95 out of 4 (Excellent With Honors - First of Class)
- Project: Virtual Reality Toolkit with educational demo

Academic Appointments

Graduate Research Assistant

August 2007 to present

[Department of Electrical and Computer Engineering,](#)
[University of Louisville](#)

Research and Development:

- Human jaw surface reconstruction from optical imaging for orthodontic treatments
- 3D colon segmentation for early detection of colorectal cancer
- Human identification at a distance for indoor and outdoor environments
- Detection and classification of pulmonary nodules for early lung cancer diagnosis
- Distinction between autistic corpus callosum and normal ones for autism diagnosis
- Monitoring the Earth's surface through change detection techniques

Help in Teaching and Preparing Lecture Notes for:

- ECE 523: Introduction to Biometrics
- ECE 600: Introduction to Shape Analysis
- ECE 619: Computer Vision
- ECE 643: Introduction to Biomedical Computing

Materials are available at [CVIP Lab website](#)

Course Instructor

September 2005 to July 2007

[Faculty of Computer Studies,](#)
[Arab Open University](#)

Activities:

- Lecturing class
- Grading students' assignments
- Grading Midterm and final exams
- Helping students throughout their study (office hours)

Graduate Teaching Assistant

June 2002 to July 2007

Faculty of Computers and Information,
Cairo University

Activities:

- Lecturing classes and labs.
- Helping students throughout their study (office hours)
- Helping students with their thesis projects
- Grading students' assignments
- Grading midterm exams
- Participating in the faculty's projects such as e-learning projects

Employment

E-Learning Course Developer

July 2006 to July 2007

Avicenna Knowledge Center (AKC) ,
Cairo University

Activities:

- Preparing course notes
- Designing course Web layout

Virtual Reality Application Designer

June 2002 to September 2003

Virtual Reality Department,
Suzan Mubarak Exploration Center for Science

Activities:

- Enhancing employees' programming skills
- Lecturing programming courses in C++
- Participating in research projects
- Problem analysis and solving

System Analyst and Junior Programmer

May 2003 to October 2003

Bullet Proof Soft Company

Activities:

- Problem Solving
- Desktop applications development

Apprenticeships and Training

Train The Trainer Workshop

July 2005

Faculty of Computers and Information,
Cairo University

Content:

- Training process
- Training objectives
- Trainer and trainees styles
- Trainer role
- Learning phases
- Training tools
- Presentation skills

E-learning in TVET Processes

August 2005 to September 2005

InWent - Capacity Building International, Germany, sponsored by the Federal Ministry for Economic Cooperation and Development on behalf of the Government of the Federal Republic of Germany

Content:

- Learning management systems
- Content development
- Promotion of internet competencies in general and vocational education
- Promotion and implementation of E-Learning and life long learning in an international organization
- E-Learning in the field of further education
- Support structures in the field of educational media

Introduction to C# Programming

January 2002

Microsoft in association with Faculty of Computers and Information,

Cairo University

Content:

Introducing .Net technology

C# programming language syntax

Java Programming July 2001

Faculty of Computers and Information, Cairo University

Network Setup, Simulation and Management July 2001

Telecommunication National Institute

Content:

- Basic experience in dealing with Opnet application
- Knowledge enhancement about the practical part of networks

Object Oriented Programming (C++) September 1999

Faculty of Computers and Information, Cairo University

Refereed Journal Publications

- **Shireen Elhabian** and Aly Farag. Towards Efficient Image Irradiance Modeling of Convex Lambertian Surfaces Under Single Viewpoint and Frontal Illumination. IET-Computer Vision, under review.
- **Shireen Y. Elhabian**, Khaled M. ElSayed, Sanaa ElOla H. Ahmed. Moving Object Detection in Spatial Domain using Background Removal Techniques - State-of-Art. *Recent Patents on Computer Science*, Volume 1, Issue 1, January, 2008, pp. 32-54.
- **Shireen Y. Elhabian**, Khaled M. ElSayed, Sanaa ElOla H. Ahmed. Sperm Tracking using Backward Projection. *Egyptian Informatics Journal*, June 2006.
- **Shireen Y. Elhabian**, Khaled M. ElSayed, Sanaa ElOla H. Ahmed. Innovated Approach for Automated Snakes Initialization. *Egyptian Informatics Journal*, Vol. 5, No. 1, June 2004, pp. 103-115.

Conference Publications

Image Irradiance Modeling

- **Shireen Elhabian** and Aly Farag. Modeling Image Irradiance Under Natural Illumination And Isotropic Surface Reflectance. International Conference of Image Processing (ICIP), 2012, accepted to appear.
- **Shireen Elhabian**, Ham Rara and Aly Farag. Towards Efficient and Compact Phenomenological Representation of Arbitrary Bidirectional Surface Reflectance. British Machine Vision Conference (BMVC), vol. 89, pp 1-11, Aug 2011.
- **Shireen Elhabian**, Ham Rara and Aly Farag. Towards Accurate and Efficient Representation of Image Irradiance of Convex-Lambertian Objects Under Unknown Near Lighting. International Conference of Computer Vision (ICCV), pp. 1732 - 1737, Nov 2011.
- **Shireen Elhabian**, Ham Rara, Aly Farag. Modeling Lambertian Surfaces Under Unknown Distant Illumination Using Hemispherical Harmonics. Eighth Canadian Conference on Computer and Robot Vision (CRV), pp.293-300, 2011.
- **Shireen Elhabian**, Ham Rara and Aly Farag. On The Use of Hemispherical Harmonics For Modeling Images of Object Under Known Distant Illumination. International Conference of Image Processing (ICIP), pp. 1109 - 1112, Sept 2011.

Shape Recovery

- **Shireen Elhabian**, Eslam Mostafa, Ham Rara, Aly Farag. Non-Lambertian Model-based Facial Shape Recovery from Single Image Under Unknown General Illumination. Nineth Canadian Conference on Computer and Robot Vision (CRV), pp. 252-259, May 2012.
- **Shireen Elhabian**, Ham Rara, Asem Ali, Aly Farag. Illumination-invariant Statistical Shape Recovery with Contiguous Occlusion. Eighth Canadian Conference on Computer and Robot Vision (CRV), pp. 301-308, 2011.
- Ham Rara, **Shireen Elhabian**, Thomas Starr, and Aly Farag. 3D Face Recovery from Intensities of General and Unknown Lightning Using Partial Least Squares.

Proc. of 2010 IEEE International Conference on Image Processing (ICIP), pp. 4041-4044, 2010

- Ham Rara, **Shireen Elhabian**, Thomas Starr, and Aly Farag. Model-Based Shape Recovery From Single Images Of General And Unknown Lighting. IEEE International Conference on Image Processing (ICIP), Nov. 7 - Nov. 10, 2009, Cairo, Egypt
- Ham Rara, **Shireen Elhabian**, Thomas Starr, and Aly Farag. Face reconstruction and recognition using a statistical model combining shape and spherical harmonics. IEEE SOUTHEASTCON, March 2009
- Amal Farag, **Shireen Elhabian**, Abdelrehim Ahmed, and Aly Farag. Noise Analysis of a SFS Algorithm Formulated under Various Imaging Conditions. 4th International Symposium on Visual Computing, December 1-3, 2008, Las Vegas, Nevada, USA
- Ham Rara, **Shireen Elhabian**, Thomas Starr, Aly Farag. A Statistical Model Combining Shape and Spherical Harmonics for Face Reconstruction and Recognition. Biomedical Engineering Conference, CIBEC2008. Cairo International, 1:1-4, 18-20 Dec. 2008.

Biometrics

- Mostafa Abdelrahman, Asem Ali, **Shireen Elhabian**, Ham Rara, and Aly A. Farag. A Passive Stereo System for 3D Human Face Reconstruction and Recognition at a Distance. IEEE CVPR Workshop on Biometrics, pp. 17 - 22, June 2012.
- Ham Rara, **Shireen Elhabian**, Asem Ali, Mike Miller, Thomas Starr, and Aly Farag. Face Recognition at-a-Distance using Texture and Sparse-Stereo Reconstruction. Proc. of IEEE Fourth International Conference on Biometrics: Theory, Applications and Systems (BTAS), pp. 1221-1224, 2010.
- Ham Rara, Asem Ali, **Shireen Elhabian**, Thomas Starr, and Aly A. Farag. Face Recognition at-a-Distance using Texture, Dense- and Sparse-Stereo Reconstruction. Proceedings of the International Conference on Pattern Recognition (ICPR), pp. 1221-1224, 2010.
- Ham Rara, **Shireen Elhabian**, Asem Ali, Travis Gault, Mike Miller, Thomas Starr, and Aly Farag. A Framework for Long Distance Face Recognition using Dense- and

Sparse-Stereo Reconstruction. 5th International Symposium on Visual Computing (ISVC09), Nov. 30 - Dec. 2, 2009, Las Vegas, Nevada, USA.

- Ham Rara, **Shireen Elhabian**, Asem Ali, Mike Miller, Thomas Starr, and Aly Farag. Face recognition at a distance based on sparse-stereo reconstruction. IEEE CVPR Biometrics Workshop, 2009.
- Ham Rara, **Shireen Elhabian**, Asem Ali, Mike Miller, Thomas Starr, and Aly Farag. Distant Face Recognition Based On Sparse Stereo Reconstruction. IEEE International Conference on Image Processing (ICIP), Nov. 7 - Nov. 10, 2009, Cairo, Egypt

Medical Imaging

- Marwa Ismail, **Shireen Elhabian**, Aly Farag, Gerald Dryden, Albert Seow. 3D Automated Colon Segmentation for Efficient Polyp Detection. 6th Cairo International Conference on Biomedical Engineering (CIBEC), 2012, accepted to appear.
- Aly Farag, **Shireen Elhabian**, Aly Abdelrahim, Wael Aboelmaaty, Allan Farman and David Tazman. Model-based Human Teeth Shape Recovery from a Single Optical Image With Unknown Illumination. MICCAI Medical Computer Vision Workshop (MCV) 2012, accepted to appear.
- Marwa Ismail, **Shireen Elhabian**, Aly Farag and Gerald W. Dryden. Fully Automated 3D Colon Segmentation for Early Detection of Colorectal Cancer based on Convex Formulation of the Active Contour Model. IEEE CVPR Workshop on Medical Computer Vision, pp. 58 - 63, June 2012.
- Aly Abdelrahim, Aly Farag, **Shireen Elhabian**, Eslam Mostafa and Wael Aboelmaaty. Occlusal Surface Reconstruction Of Human Teeth From A Single Image Based On Object And Sensor Physical Characteristics. International Conference of Image Processing (ICIP), 2012, accepted to appear.
- Amal Farag, Asem Ali, **Shireen Elhabian**, James Graham, Aly Farag, and Robert Falk. Feature-Based Lung Nodule Classification. Proc. of International Symposium on Visual Computing (ISVC), pp. 79-88, 2010.

- Ahmed Farag, **Shireen Elhabian**, Mostafa Abdelrahman, James Graham, Aly Farag, Dongqing Chen, and Manuel F. Casanova. Surface Modeling of the Corpus Callosum. Proc. of International Symposium on Visual Computing (ISVC), pp. 9-18, 2010.
- Amal Farag, **Shireen Elhabian**, James Graham, Aly Farag, and Robert Falk. Toward Precise Pulmonary Nodule Descriptors for Nodule Type Classification. Proc. of the 13th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI), pp. 626-633, 2010.
- Amal Farag, **Shireen Elhabian**, James Graham, Aly Farag, Salwa Elshazly, Robert Falk, Hani Mahdi, Hossam Abdelmunim, and Sahar Al-Ghaafary. Modeling of the Lung Nodules for Detection in LDCT Scans. Proc. of the 32nd IEEE Engineering in Medicine and Biology Society (EMBC), 2010.
- Ahmed Farag, **Shireen Elhabian**, Mostafa Abdelrahman, James Graham, Aly Farag, Dongqing Chen, and Manuel F. Casanova. Shape Modeling of the Corpus Callosum. Proc. of the 32nd IEEE Engineering in Medicine and Biology Society (EMBC), pp. 4288-4291, 2010.
- **Shireen Elhabian**, Hossam Abd EL Munim, Salwa Elshazly, Aly A Farag, and Mohamed Aboelghar. Experiments on sensitivity of template matching for lung nodule detection in low dose CT scans. ISSPIT 2007, 1040-1046.

Others

- Amal A. Farag, Asem M. Ali, **Shireen Elhabian** and Aly A. Farag. Probability Density Estimation By Linear Combinations Of Gaussian Kernels Generalizations And Algorithmic Evaluation. International Conference on Multimedia Technology (ICMT), pp. 6491-6494, 26-28 July 2011.
- Mostafa Abdelrahman, Asem Ali, **Shireen Elhabian** and Aly Farag. Solving Geometric Co-Registration Problem of Multi-Spectral Remote Sensing Imagery Using SIFT-based Features Toward Precise Change Detection. 7th International Symposium on Visual Computing (ISVC). Las Vegas, Nevada, USA, September 26-28, 2011, vol.2, pp 607-616.

Technical Reports

- **Shireen Elhabian** and Aly Farag. *Tutorial on Topology*. Computer Vision and Image Processing Laboratory, CVIP Lab, University of Louisville, June 2009.
- **Shireen Elhabian** and Aly Farag. *Tutorial on Computational Geometry*. Computer Vision and Image Processing Laboratory, CVIP Lab, University of Louisville, July 2009.
- **Shireen Elhabian** and Aly Farag. *Tutorial on Curves and Surfaces*. Computer Vision and Image Processing Laboratory, CVIP Lab, University of Louisville, August 2009.
- **Shireen Elhabian**, Amal Farag and Aly Farag. *Tutorial on CT Reconstructions*. Computer Vision and Image Processing Laboratory, CVIP Lab, University of Louisville, January 2009.
- **Shireen Elhabian**, Amal Farag and Aly Farag. *Tutorial on Object Registration*. Computer Vision and Image Processing Laboratory, CVIP Lab, University of Louisville, pp. 120, January 2009.
- **Shireen Elhabian** and Aly Farag. *Tutorial on MRI Reconstructions*. Computer Vision and Image Processing Laboratory, CVIP Lab, University of Louisville, pp. 107, February 2009.
- **Shireen Elhabian**, Amal Farag and Aly Farag. *Tutorial on Camera Calibration*. Computer Vision and Image Processing Laboratory, CVIP Lab, University of Louisville, February 2008.
- **Shireen Elhabian** and Aly Farag. *Tutorial on Stereo Reconstruction*. Computer Vision and Image Processing Laboratory, CVIP Lab, University of Louisville, March 2008.
- **Shireen Elhabian** and Aly Farag. *Tutorial on Shape from Shading*. Computer Vision and Image Processing Laboratory, CVIP Lab, University of Louisville, April 2008.
- **Shireen Elhabian** and Aly Farag. *2D Face Recognition Using PCA, ICA and*

LDA. Computer Vision and Image Processing Laboratory, CVIP Lab, University of Louisville, pp. 110, November 2008.

- **Shireen Elhabian** and Aly Farag. *Biometrics in a glimpse*. Computer Vision and Image Processing Laboratory, CVIP Lab, University of Louisville, September 2007.

Papers in Preparation

- **Shireen Elhabian** and Aly Farag. Phenomenological Modeling of Image Irradiance of Arbitrary Surfaces. *Journal of the Optical Society of America*.
- **Shireen Elhabian** and Aly Farag. Analytic Bilinear Subspace Construction for Modeling Image Irradiance Under Natural Illumination and Non-Lambertian Reflectance. *Image and Vision Computing*, Elsevier.
- **Shireen Elhabian**, Ham Rara and Aly Farag. Non-Lambertian Model-based Shape Recovery Under Natural Illumination. *Computer Vision and Image Understanding*.
- **Shireen Elhabian**, Aly Farag, Aly Abdelrahim, Wael Aboelmaaty, Allan Farman and David Tazman. Clinical Crowns Shape Reconstruction - An Image-based Approach. *Transaction of Medical Imaging*.

Referee Service

- Reviewer at *Journal of Pattern Recognition Letters*, Elsevier
- Reviewer at *IET Computer Vision*

Honors and Awards

- IEEE CVPR Doctoral Consortium grant, 2012
- ECE Outstanding Graduate Student Award 2009, Department of Electrical and Computer Engineering, University of Louisville, Louisville, Kentucky, USA.
- 1st place, Engineering EXPO 2009, Department of Electrical and Computer Engineering, University of Louisville, Louisville, Kentucky, USA.
- Computer Vision Summer School, Sicily, Italy, 2008
- Best Teaching Assistant Award, March 2007, Teaching Members Club, Cairo University, Egypt

- Best Teaching Assistant Award, April 2004, Faculty of Computers and Information, Cairo University, Egypt
- Evaluation Certificate, April 2003, Cairo University, Egypt
- Superiority Award, April 2003, Chamber of Information Technology and Communications (CIT) Egyptian Industries Union
- Superior Student Certification, August 2002, Defense Ministry, Egypt
- Egypt Science Festival Certification, December 2002, Cairo University, Egypt
- Top student in class throughout my undergraduate education in Egypt

Teaching Experience

University of Louisville, Louisville, Kentucky USA

- *Geometric Computer Vision*: shape from shading, stereo reconstruction, camera calibration, non-linear optimization techniques.
- *Biometrics*: face recognition (2D/3D), feature extraction, data reduction, iris recognition.

Cairo University, Cairo, Egypt

- *Virtual Reality Laboratory* (Labs: WTK, and OpenGL under VC++)
- *Digital Image Processing* (Labs: MATLAB and Image Processing Toolbox)
- *Signals and Systems* (Labs: MATLAB and Signal Processing Toolbox)
- *Digital Signal Processing* (Labs: MATLAB and Signal Processing Toolbox)
- *Computer Graphics* (Labs: Java and VC++)
- *Multimedia* (Implementation of different compression algorithms using VC++, C# and .Net)
- *Speech Processing* (Labs: Speech processing using MATLAB)
- *Advanced Digital Image Processing* (Labs: Image processing using VC++)
- *Advanced Mathematics* (Labs: MATLAB)
- *Fuzzy Logic and Neural Networks* (Labs: MATLAB Fuzzy Toolbox, and Neural Network Toolbox)
- *Introduction to C++ programming*

Application Areas

Image-based rendering, biometrics, object detection and recognition, 3D shape recovery, object classification and/or medical imaging

Skills

Programming Languages and Toolkits

C and C++	Highly Competent
MATLAB	Highly Competent
Parallel Programming with CUDA	In Progress
C#	Fair and In Progress
Java	Basic
Python	Getting Started
OpenCV	Fair and In Progress
Visualization Toolkit (VTK)	Fair and In Progress
Insight Toolkit (ITK)	Getting Started
Point Cloud Library (PCL)	Getting Started

Software

Microsoft Office	Highly Competent
Latex	Good
Microsoft Visual Studio	Highly Competent
Mathematica	Basic

Others

Presentations preparation skills
 Writing skills
 Course preparation skills
 Instructing skills

Areas of Research Focus

Mathematics:

- Real and complex analysis, differential geometry, linear algebra, multilinear algebra, spherical and hemispherical harmonics, Gabor wavelets and group theory

Signal Processing:

- Signal and image processing, probability, random variables, stochastic processes, information theory, subspace learning, pattern recognition and machine learning

Computer Vision:

- Shape-from-shading, statistical shape-from-shading, stereo reconstruction, image formation and biometrics

Computer Graphics:

- Frequency-space rendering, physically-based rendering (PBRT), illumination representation and reflectance modeling

Medical Imaging:

- Statistical shape analysis, image segmentation and image processing

Languages

Arabic	Mother Tongue
English	Fluent - TOEFL score 600
French	Basic