On The Development and Error Analysis of a High Dynamic Range Imaging System for Luminance Measurements

by

Mio Jane Stanley

B.S., University of Colorado Boulder, 2016

A thesis submitted to the Faculty of the Graduate School of the University of Colorado in partial fulfillment of the requirements for the degree of Masters of Science Department of Civil, Environmental, and Architectural Engineering

2016

This thesis entitled: On The Development and Error Analysis of a High Dynamic Range Imaging System for Luminance Measurements written by Mio Jane Stanley has been approved for the Department of Civil, Environmental, and Architectural Engineering

Professor C. Walter Beamer IV

Gregor Henze

Mark Jongewaard

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline. Stanley, Mio Jane (M.S., Architectural Engineering)

- On The Development and Error Analysis of a High Dynamic Range Imaging System for Luminance Measurements
- Thesis directed by Professor C. Walter Beamer IV

This document provides the documentation of the development of a new High Dynamic Range Imaging (HDRI) system for the University of Colorado Boulder, testing of the HDRI system with a novel calibration system, and the evaluation of the most commonly used High Dynamic Range (HDR) image creation software options. One of the primary goals of this document is to provide a comprehensive literature search of published research related to HDRI and its applications. While this particular research project will not cover all of these topics, they are included here as a resource and guide for future HDRI research projects to be conducted at the University. As a part of that goal, this project will develop a new HDRI system for the University to use in future research projects. This involves determining the correct calibration procedures necessary to obtain useful data through camera response function recovery and the creation of vignetting correction filters. A novel calibration system using small integrating spheres and LEDs is developed and tested. An introduction of this system is provided here in addition to some preliminary tests of the new HDRI system to determine the applicability of the calibration device for further HDRI research. Finally, an evaluation of various HDR image software options is performed. The results of these tests will inform which software options are best suited for the needs of future HDRI research to be conducted at the University.

Acknowledgements

I would like to thank my advisor, Dr. C. Walter Beamer, for initially piquing my interest in lighting and setting me on the path towards a career in illumination engineering and lighting design. Thank you for the encouragement and guidance you provided throughout both my undergraduate and graduate careers, and for being patient and supportive of me as I concluded this project.

I would also like to thank the Pacific Northwest National Laboratory (PNNL) and the Civil, Environmental, and Architectural Engineering Department here at the University of Colorado Boulder for providing funding. Thank you in particular to Dr. Robert Davis of PNNL for taking an interest in my project and becoming involved as a committee member.

Thank you to Mark Jongewaard for your involvement and help throughout this project in addition to being part of my committee. Thank you also to LTI Optics for allowing me to use your equipment and lab space for testing.

To Dr. Gregor Henze, thank you for agreeing to be a part of my committee with such short notice and for providing a fresh perspective outside of lighting.

To Professor Sandra Vasconez, thank you for your continued support and instilling in me the passion for lighting design.

Finally, I would like to thank Adam and my family for providing constant love and support throughout all of my endeavors, especially this one.

Contents

Chapter

| 1 | Intro | oductio | n | 1 |
|----------|-------|----------|---|----|
| | 1.1 | Motiva | ation | 1 |
| | 1.2 | Conte | nts | 3 |
| 2 | Lite | rature S | Search | 4 |
| | 2.1 | Lighti | ng Definitions | 5 |
| | 2.2 | Image | Formation | 8 |
| | | 2.2.1 | Lens Systems | 8 |
| | | 2.2.2 | CCD and CMOS Sensors | 12 |
| | | 2.2.3 | RAW and JPEG Image File Types | 16 |
| | 2.3 | Backg | round on High Dynamic Range Imaging | 18 |
| | | 2.3.1 | Motivation for HDR Photography | 18 |
| | | 2.3.2 | Overview of the HDRI Process | 21 |
| | | 2.3.3 | HDRI Experimentation and Validation | 23 |
| | 2.4 | Applic | cations of HDRI | 31 |
| | | 2.4.1 | Illuminance Analysis | 32 |
| | | 2.4.2 | Sky Dome Imaging and Image-Based Lighting | 32 |
| | | 2.4.3 | Daylighting Control | 34 |
| | | 2.4.4 | Photogrammetry | 35 |

| | | 2.4.5 | Glare Evaluation | 37 |
|---|------|---------|--|-----|
| | | 2.4.6 | Photometry | 41 |
| | 2.5 | HDR I | Image Calibrations | 44 |
| | | 2.5.1 | Camera Response Function and Creating HDR Images | 44 |
| | | 2.5.2 | Vignetting | 55 |
| | | 2.5.3 | Point Spread Function and Lens Flare | 61 |
| | 2.6 | Near I | Field Illuminance Calculations | 65 |
| | | 2.6.1 | Proposed Near-Field Photometry Methods | 67 |
| | | 2.6.2 | Near-Field Goniometer | 72 |
| | 2.7 | Currei | nt Challenges and Research Gaps | 73 |
| 3 | Deve | elopmer | nt of New HDRI System | 76 |
| | 3.1 | Equip | ment | 76 |
| | | 3.1.1 | Camera Body | 76 |
| | | 3.1.2 | Camera Lenses | 77 |
| | | 3.1.3 | Other Equipment | 77 |
| | 3.2 | Softwa | are | 80 |
| | | 3.2.1 | Tethering | 80 |
| | | 3.2.2 | HDR Image Creation | 82 |
| | | 3.2.3 | HDR Image Analysis | 85 |
| | 3.3 | Deterr | mining Camera Response Functions | 87 |
| | | 3.3.1 | Methodology | 87 |
| | | 3.3.2 | Results | 90 |
| | | 3.3.3 | Discussion | 100 |
| | 3.4 | Deterr | mining Vignetting Functions | 101 |
| | | 3.4.1 | Background | 101 |
| | | 3.4.2 | Methodology | 102 |

vi

| | | 3.4.3 | Results | ó |
|---|-------|---------|--|---|
| | 3.5 | Vignet | tting Correction Check |) |
| | | 3.5.1 | Methodology 111 | L |
| | | 3.5.2 | Results | F |
| | | 3.5.3 | Discussion |) |
| 4 | Integ | grating | Sphere Calibration Rig 123 | 3 |
| | 4.1 | Design | $1 \ldots \ldots$ | 3 |
| | 4.2 | Testin | g Lambertian Nature of Lens | 5 |
| | | 4.2.1 | Methodology 125 | 5 |
| | | 4.2.2 | Results | 7 |
| | | 4.2.3 | Discussion | 3 |
| | 4.3 | Discov | vered Issues with HDRI 130 |) |
| | | 4.3.1 | Lens Flare |) |
| | | 4.3.2 | Color LED Issues | F |
| 5 | Soft | ware Ev | valuation 140 |) |
| | 5.1 | Backg | round |) |
| | 5.2 | Metho | odology | L |
| | 5.3 | Result | s |) |
| | | 5.3.1 | Visual Aesthetic Evaluation |) |
| | | 5.3.2 | Luminance HDR | 2 |
| | | 5.3.3 | Picturenaut | Ś |
| | | 5.3.4 | Luminance-Based Performance | 7 |
| | | 5.3.5 | Comments on Usability | 7 |
| | 5.4 | Discus | ssion | 3 |

vii

| 6 | Conclusion | | | | |
|---|------------|------------------------|-----|--|--|
| | 6.1 | Final Remarks | 170 | | |
| | 6.2 | Future Research Topics | 171 | | |

viii

173

Bibliography

Appendix

| \mathbf{A} | Glos | sary | 177 |
|--------------|------|---|-----|
| | A.1 | List of Symbols | 177 |
| | A.2 | List of Acronyms and Initialisms | 178 |
| | A.3 | Lighting Definitions | 180 |
| | A.4 | Other Definitions | 182 |
| в | Cam | era Response Functions | 184 |
| | B.1 | Nikon D5200 with Sigma Fisheye Lens Daylight White Balance \ldots | 184 |
| | B.2 | Nikon D5200 with Sigma Fisheye Lens $$ Fluorescent White Balance $$ | 185 |
| | B.3 | Nikon D5200 with Sigma Prime Lens Daylight White Balance | 186 |
| | B.4 | Nikon D5200 with Sigma Prime Lens Fluorescent White Balance \ldots | 187 |
| | B.5 | Nikon D5200 with Nikkor Zoom Lens, 24mm Daylight White Balance | 188 |
| | B.6 | Nikon D5200 with Nikkor Zoom Lens, 24mm $$ Fluorescent White Balance $$. | 189 |
| | B.7 | Nikon D5200 with Nikkor Zoom Lens, 18mm Daylight White Balance | 190 |
| | B.8 | Nikon D5200 with Nikkor Zoom Lens, $18\mathrm{mm}$ Fluorescent White Balance $$. | 191 |
| \mathbf{C} | Vign | netting Functions | 192 |
| | C.1 | MATLAB Vignetting Correction Script | 193 |
| | C.2 | Vignetting Error Charts | 194 |
| | | C.2.1 Fisheye Lens | 194 |

| | | C.2.2 | Prime Lens | 196 |
|---|------|---------|---|------------|
| | | C.2.3 | Zoom 18mm Lens | 198 |
| | | C.2.4 | Zoom 24mm Lens | 200 |
| D | Soft | ware Ev | aluation Results | 202 |
| | D.1 | Lumin | ance HDR Results | 203 |
| | | D.1.1 | RAW-Based HDR Images, Classroom with Low Fluorescent Light $+$ | |
| | | | Daylight | 203 |
| | | D.1.2 | JPEG-Based HDR Images, Classroom with Low Fluorescent Light $+$ | |
| | | | Daylight | 204 |
| | D.2 | Picture | enaut Results | 207 |
| | | D.2.1 | Classroom, Fluorescent + Daylight | 207 |
| | | D.2.2 | Classroom High Fluorescent | 208 |
| | | D.2.3 | Classroom Low Fluorescent | 209 |
| | | D.2.4 | Classroom Window | 210 |
| | | D.2.5 | Spheres Far | 210 |
| | | D.2.6 | Spheres Close | 211 |
| | D.3 | hdrgen | , $raw2hdr$, and Bracket Results | 212 |
| | | D.3.1 | Classroom, Fluorescent + Daylight | 212 |
| | | D.3.2 | Classroom High Fluorescent | 213 |
| | | D 3 3 | Classroom Low Fluorescent | 214 |
| | | D 3 4 | Classroom Window | 211 |
| | | D.9.5 | Spheros Far | 210 915 |
| | | D.3.3 | | 210 |
| | | D.3.6 | Spheres Close | 216 |

Tables

Table

| 2.1 | CCD vs CMOS Characteristics | 15 |
|-----|---|-----|
| 3.1 | Minimum Measurement Area for LS-110 | 78 |
| 3.2 | HDR Image Creation Software Compatibilities | 82 |
| 3.3 | HDR Image Creation Software Accepted Input and Output Files Types $\ . \ .$ | 83 |
| 3.4 | Camera Response Function Image Sets | 89 |
| 3.5 | Prime Lens, Fluorescent White Balance Red Channel Curve | 97 |
| 3.6 | Vignetting Functions - Nikon D5200 | 107 |
| 3.7 | Nikkor Zoom 18mm Luminance Measurements | 113 |
| 3.8 | Nikkor Zoom 24mm, Sigma Prime, Sigma Fisheye Luminance Measurements | 113 |
| 4.1 | LED Sphere Properties | 125 |
| 4.2 | Lens Uniformity Test Results | 127 |
| 4.3 | Lens Lambertian Test Results | 128 |
| 5.1 | Horizontal Illuminance Levels for Test Scenes | 141 |
| 5.2 | Luminance HDR Profiles | 149 |
| 5.3 | $\label{eq:Luminance} \ {\rm HDR} \ {\rm Profile} \ {\rm Errors} \ {\rm for} \ {\rm Classroom} \ {\rm Scene} \ {\rm with} \ {\rm Fluorescent} \ + \ {\rm Daylight}$ | 153 |
| 5.4 | Matched-Pairs t-Test Results for Picturenaut HDR Images | 157 |
| 5.5 | Average Errors | 158 |

| 5.6 | Correlation Test Results between Error and Luminance for Picturenaut HDR | |
|-----|---|-----|
| | Images, 99% Confidence Interval | 165 |
| 5.7 | Matched Pair t-Test Results for Bracket HDR Images | 166 |
| 5.8 | Matched Pair t-Test Results for <i>hdrgen</i> and <i>raw2hdr</i> HDR Images | 166 |

Figures

Figure

| 2.1 | Luminous Efficacy Curves | 5 |
|------|---|----|
| 2.2 | Irradiance (a) and Radiance (b) $[1]$ | 7 |
| 2.3 | Simple Lens Model [1] | 9 |
| 2.4 | Diagram of Lens System Aperture | 10 |
| 2.5 | Thick Lens Model $[1]$ | 11 |
| 2.6 | Vignetting $[1]$ | 11 |
| 2.7 | Bayer Pattern [2] \ldots | 13 |
| 2.8 | Dynamic Range $[3]$ | 19 |
| 2.9 | High Dynamic Range Image Example | 20 |
| 2.10 | HDR False Color Luminance Example | 21 |
| 2.11 | HDR Process [4] \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots | 23 |
| 2.12 | HDR Photometry Setup $[5]$ | 44 |
| 2.13 | Image Acquisition Pipeline [4] | 45 |
| 2.14 | Optical System Diagram | 47 |
| 2.15 | Camera Response Function Example [6] | 55 |
| 2.16 | Vignetting, Full Frame and Fisheye [7] | 56 |
| 2.17 | Vignetting [8] | 56 |
| 2.18 | Examples of Starburst PSF, Ghosting, and Haze [9] | 62 |
| 2.19 | PSF as a Function of Eccentricity [10] | 63 |

| 2.20 | MTF Target [11] | 64 |
|------|--|-----|
| 2.21 | Results of Near-Field Illuminance Test [12] | 71 |
| 3.1 | Konica-Minolta LS-110 Luminance Meter | 78 |
| 3.2 | Konica-Minolta LS-110 Luminance Meter Response Curve | 79 |
| 3.3 | X-Rite Color Checker Classic and Greyscale Cards | 79 |
| 3.4 | Camera Response Function Scene | 88 |
| 3.5 | Prime f/3.5, Daylight White Balance | 91 |
| 3.6 | Fisheye f/22, Daylight White Balance \ldots | 92 |
| 3.7 | Zoom 24mm f/22, Daylight White Balance | 93 |
| 3.8 | Zoom 24mm f/9, Daylight White Balance | 94 |
| 3.9 | Zoom 24mm f/5.6, Daylight White Balance | 95 |
| 3.10 | Prime 24mm Red Curve, Daylight White Balance | 96 |
| 3.11 | Prime 24mm Red Curve, Fluorescent White Balance | 97 |
| 3.12 | Averaged Red CRF, Fluorescent White Balance | 98 |
| 3.13 | Averaged Green CRF, Fluorescent White Balance | 98 |
| 3.14 | Averaged Blue CRF, Fluorescent White Balance | 98 |
| 3.15 | Final Averaged CRF, Fluorescent White Balance | 100 |
| 3.16 | Vignetting Test Setup | 103 |
| 3.17 | Fisheye Vignetting Test Image | 103 |
| 3.18 | Fisheye Lens Vignetting Functions | 108 |
| 3.19 | Prime Lens Vignetting Functions | 108 |
| 3.20 | Zoom 18mm Lens Vignetting Functions | 109 |
| 3.21 | Zoom 24mm Lens Vignetting Functions | 109 |
| 3.22 | Vignetting Correction Example, Sigma Fisheye Lens at $f/2.8$ | 110 |
| 3.23 | Vignetting Verification Test Setup | 112 |
| 3.24 | Vignetting Target Numbering | 112 |

xiii

| 3.25 | Fisheye Lens f/2.8 Percent Error | 114 |
|---|---|--|
| 3.26 | Zoom 24mm Lens, Original Images Error | 115 |
| 3.27 | Zoom 24mm Lens, Corrected Images Error | 115 |
| 3.28 | Zoom 18mm Lens, Original Images Error | 116 |
| 3.29 | Zoom 18mm Lens, Corrected Images Error | 116 |
| 3.30 | Prime Lens, Original Images Error | 117 |
| 3.31 | Prime Lens, Corrected Images Error | 117 |
| 3.32 | Fisheye Lens, Original Images Error | 118 |
| 3.33 | Fisheye Lens, Corrected Images Error | 118 |
| 3.34 | Fisheye Lens, Reduction in Error | 120 |
| 3.35 | Prime Lens, Reduction in Error | 120 |
| 3.36 | Zoom 18mm Lens, Reduction in Error | 121 |
| 3.37 | Zoom 24mm Lens, Reduction in Error | 121 |
| | | |
| 4.1 | Back Half of Luminance Sphere | 124 |
| 4.1 4.2 | Back Half of Luminance Sphere | 124 124 |
| 4.1 4.2 4.3 | Back Half of Luminance Sphere Luminance Spheres Lambertian Test Setup | 124 124 126 |
| 4.1 4.2 4.3 4.4 | Back Half of Luminance Sphere | 124 124 126 127 |
| 4.1 4.2 4.3 4.4 4.5 | Back Half of Luminance Sphere | 124 124 126 127 129 |
| 4.1 4.2 4.3 4.4 4.5 4.6 | Back Half of Luminance Sphere | 124 124 126 127 129 131 |
| 4.1 4.2 4.3 4.4 4.5 4.6 4.7 | Back Half of Luminance Sphere | 124 124 126 127 129 131 132 |
| 4.1 4.2 4.3 4.4 4.5 4.6 4.7 4.8 | Back Half of Luminance Sphere | 124 126 127 129 131 132 133 |
| 4.1 4.2 4.3 4.4 4.5 4.6 4.7 4.8 4.9 | Back Half of Luminance Sphere | 124 126 127 129 131 132 133 134 |
| 4.1 4.2 4.3 4.4 4.5 4.6 4.7 4.8 4.9 4.10 | Back Half of Luminance Sphere | 124 126 127 129 131 132 133 134 136 |
| 4.1 4.2 4.3 4.4 4.5 4.6 4.7 4.8 4.9 4.10 4.11 | Back Half of Luminance Sphere | 124 126 127 129 131 132 133 134 136 |
| 4.1 4.2 4.3 4.4 4.5 4.6 4.7 4.8 4.9 4.10 4.11 4.12 | Back Half of Luminance Sphere | 124 126 127 129 131 132 133 134 136 137 |

| 5.1 | Classroom Scene Targets | 143 |
|--|---|--|
| 5.2 | Classroom Window Scene Targets | 143 |
| 5.3 | Calibration Spheres Targets | 144 |
| 5.4 | Classroom, Fluorescent Light Only (High and Low) | 144 |
| 5.5 | Classroom, Fluorescent plus Daylight | 145 |
| 5.6 | Classroom Window, Daylight Only | 145 |
| 5.7 | Calibration Spheres, Far Apart | 146 |
| 5.8 | Calibration Spheres, Close Together | 146 |
| 5.9 | Example HDR Images: a) hdrgen, b) raw2hdr, c) Picture naut JPEG, d) Pic- | |
| | turenaut RAW, e) Luminance HDR JPEG, f) Luminance HDR RAW, g) | |
| | Luminance HDR Robertson, h), default Bracket | 151 |
| 5.10 | eq:Percent Error for Plateau Weighting, Classroom Fluorescent + Daylight JPEG | |
| | Scene | 154 |
| 5.11 | Luminance HDR Profiles, Classroom Fluorescent + Daylight JPEG Scene | 155 |
| | | |
| 5.12 | Picturenaut Comparisons | 156 |
| 5.12 5.13 | Picturenaut Comparisons Column Chart of Average Errors | 156 159 |
| 5.125.135.14 | Picturenaut Comparisons | 156 159 160 |
| 5.12 5.13 5.14 5.15 | Picturenaut Comparisons | 156 159 160 161 |
| 5.12 5.13 5.14 5.15 5.16 | Picturenaut Comparisons | 156 159 160 161 |
| 5.12 5.13 5.14 5.15 5.16 | Picturenaut Comparisons | 156 159 160 161 |
| 5.12 5.13 5.14 5.15 5.16 5.17 | Picturenaut Comparisons | 156 159 160 161 162 162 |
| 5.12 5.13 5.14 5.15 5.16 5.17 5.18 | Picturenaut Comparisons | 156 159 160 161 162 162 163 |
| 5.12 5.13 5.14 5.15 5.16 5.17 5.18 5.19 | Picturenaut ComparisonsColumn Chart of Average ErrorsErrors in Classroom Scene with High Fluorescent LightClassroom High Fluorescent Percent Errors without Luminance HDRClassroom High Fluorescent + Daylight Percent Errors without LuminanceHDRClassroom Low Fluorescent Percent Errors without Luminance HDRClassroom Window Percent Errors without Luminance HDRSpheres Far Apart Percent Errors without Luminance HDR | 156 159 160 161 162 162 163 163 |

Chapter 1

Introduction

1.1 Motivation

High dynamic range (HDR) photography is a relatively new tool in the building research field, popular for its affordability, ease of use, and ability to acquire large amounts of data in a small amount of time. Lighting researchers primarily use it to acquire luminance data, a lighting metric that can best be related to the human perception of brightness. The luminance maps generated from HDR images can provide luminance data of the pictured scene on a pixel level, making it very simple and quick for researchers or designers to evaluate the lighting conditions across the entire scene.

There are many applications and ongoing research topics that are based on the use of high dynamic range imaging (HDRI), so the development and refinement of HDRI technologies and methods would be greatly beneficial to the building research community. It has proven itself useful in standard architectural applications, but there is still progress to be made in other applications, including daylight glare analysis and even lighting control.

One such application of HDRI is Near-field photometry. Near-field photometry is the practice of measuring and defining light sources within a distance where they cannot be treated as a point source. Conventional photometry measures light intensity distributions of sources at a sufficiently far distance such that the light source can be assumed to act as a point source. There are some oversimplifications and assumptions made during this process that can yield significant errors in certain applications, namely near-field applications. Conventional far-field photometry practices cannot account for patterns of inhomogeneity that may occur across the luminous face of a light source, like the presence of baffles or distinct lamp shapes. Although some effort has been made to establish methods of near-field photometry that can overcome these challenges, these methods often require new equipment that may be costly or difficult to use. There is a possibility that HDRI may be able to bridge the gap between near-field photometry theory and conventional practice.

While this particular research project will not begin to cover any experimentation and testing related to near-field photometry, the initial motivating factor that spurred this project was the hope of developing new near-field photometry practices using HDRI technology. Before any progress can be made in near-field applications, an HDRI system needs to be developed and calibrated. One goal of this project is to establish the HDRI system that will eventually be used to investigate this near-field photometry problem or any other HDRIrelated topics.

While the HDRI system was being developed, a novel calibration system was also designed and built to test the limits of the HDRI system. A secondary goal of this project is to document the preliminary testing done with the HDRI system using the newly designed calibration system. This calibration system consists of five small integrating spheres that are illuminated with LEDs which cover a wide range of luminance values.

A comprehensive software evaluation is also included as a part of this project to inform the selection of software used in future research projects. Prior to this project, the University had no previous experience with HDRI-related research, so it was unknown which of the available software options that deal with HDRI are best-suited for the lighting research needs of the University. This experiment will test the capabilities and performance of a handful of HDR image creation software in regards to luminance mapping.

1.2 Contents

This document is broken up into six chapters, beginning with this introduction. The next chapter is the literature search. This chapter includes a comprehensive summary of HDR-related research papers relevant to the current and future research topics to be pursued here at the University. The literature search is divided into several sections, beginning with a basic introduction of lighting-related terms. The other sections include image formation, a background on HDRI, the possible applications of HDRI, the processes involved to calibrate an HDRI system, near-field illuminance and photometry, and lastly an overview of current research gaps. Chapter 3 includes the documentation of the development of a new HDRI system for the University to use in current and future research projects. This chapter will go over the equipment and software used in the following experiments and the camera response function recovery process. Vignetting correction filters are also derived and provided for future use. Chapter 4 discusses the development and testing of a novel calibration system for HDRI systems. This calibration rig consists of five small integrating spheres illuminated with LEDs to provide points of calibration in an HDR image. Some preliminary tests and observations are included here. The evaluation of a number of software options capable of creating HDR images is provided in Chapter 5. Five individual software packages will be tested and compared to evaluate their usefulness in HDRI research. This will inform and justify the use of particular software used in current and future HDRI research. Chapter 6 will conclude this document with final remarks and ideas for future research projects.

Chapter 2

Literature Search

The Literature Search will begin with a brief section to define fundamental lighting terms, which will be used frequently throughout the paper. An understanding of these terms is crucial to understanding the content in this paper. Additionally, these lighting terms and other key terms will be defined in an attached Glossary for the reader's reference. The Glossary will also include a section to define commonly used symbols, acronyms, and initialisms. Next, the basics of image and camera technology will be introduced. This section will cover lens systems, CCD and CMOS image sensor technology, and RAW and JPEG image file types. Understanding these concepts will shed light onto the causes of numerous challenges associated with HDRI. The following section will cover background information on HDR, including motivating factors, a brief overview of the HDR image creation process, and summaries of published papers that have validated HDRI as a luminance measurement method. Additional papers that discuss the many applications of HDRI in the lighting research industry are also included to provide some context as to why HDRI is such a desirable technology to develop further. The HDRI creation and calibration processes will be explained in further detail, going over the specifics of camera response function recovery, vignetting correction, and lens flare or point spread function effects. Lastly, the concept of near-field photometry will be covered to give context for the intended purpose of the proposed HDRI system.

2.1 Lighting Definitions

There are two groups of terms defined here that are very similar. The first set of terms deal with radiant energy. However, lighting designers and researchers are interested mostly in the visible spectrum of radiant energy, so the second set of terms are introduced to deal specifically with visually evaluated radiant energy, the photometric quantities. These two sets of terms are related by a function called the action spectrum of human vision, which describes the visually evaluated portion of radiant energy as a function of wavelength. The action spectrum will actually vary between adaptation states of the eye and even among individuals, but two action spectrums have been defined as standards by the CIE called the photopic, $v(\lambda)$, and scotopic, $v'(\lambda)$, luminous efficiency curves. This document will deal only with photopic vision. The photopic, in red, and scotopic, in blue, luminous efficiency curves are shown in Figure 2.1 [13].



Figure 2.1: Luminous Efficacy Curves

Radiant Energy, Q_e , is the amount of electromagnetic energy that can be emitted, transferred, or received in the form of radiation. It is symbolized with a Q_e and expressed in units of Joules. [13]

Radiant Flux, Φ_e , is the amount of radiant power, or time rate flow of energy, from a source. Its units are Joules per second, and is symbolized with Φ_e . [13]

Irradiance, E_e , is the density of radiant energy incident on a surface, expressed in terms of Watts per area (square-meters or square-feet). In photography, this value is used to describe image brightness, or the amount of radiant energy that strikes the image plane [1]. It is typically symbolized with the letter E, but in this document E will be reserved for illuminance, so a subscript $_e$ will indicate that it is a radiometric quantity.

Radiance, L_e , is the amount of radiant energy a scene emits in the direction of the viewing point, sometimes thought of as scene brightness. It is expressed in Watts per area per steradian, where steradian is a unit of solid angle. Unlike irradiance, radiance depends on the direction of the viewing angle. Radiance is symbolized with the letter L, but again, the subscript $_e$ will be used to indicate that it is the radiometric quantity.

Pictorial representations of irradiance (a) and radiance (b) are included in Figure 2.2. Note that these images can also be used to demonstrate illuminance (a) and luminance (b).





Luminous Flux, Φ , is the flow of photopic luminous power from a source. The unit for luminous flux is the lumen [lm]. Luminous flux is represented with the Greek letter Φ . Alternatively, scotopic luminous flux is symbolized with Φ' . [13]

Illuminance, E, is the density of luminous flux incident on a surface per unit area. Illuminance is the photometric equivalent of irradiance. The units for illuminance are expressed in lumens per area, Footcandles [FC] for square-feet or Lux [lx] for square-meters. This will be symbolized with the letter E.

Luminous Intensity, I, is the light emitting power of a point source in a particular direction, or the density of luminous flux in space in that direction. Intensity does not depend on the distance from the source. The units for luminous intensity are expressed in Candelas [Cd]. Luminous intensity will be symbolized with the letter I. [13]

Luminance, L, is the photometric equivalent of radiance. It is expressed in Candelas

per unit area, where candela is a unit of luminous intensity. Luminance is also commonly referred to as the lighting metric that is best associated with the perception of brightness. One must be cautious when making this comparison as the perception of brightness is only a subjective quality, while luminance is a measurable quantity.

2.2 Image Formation

Before proceeding with the topic of the HDRI process, it is important to understand some fundamentals of how images are formed. This section will briefly cover the basics of image formation through lenses, image sensor types, and image file types.

2.2.1 Lens Systems

To begin, an image is defined as a two-dimensional pattern of brightness [1]. Images taken with a camera are created by the projection of a three-dimensional scene onto a twodimensional plane through a lens system. A digital image is simply a two-dimensional matrix of values that contains image data on a pixel level. Understanding the general premise of how these images are created via a traditional camera lens system sheds light on some of the issues that arise during the HDRI calibration process.

The purpose of a lens is to collect a finite amount of light and focus those rays of light onto an image plane. An example of a standard lens model is shown in Figure 2.3. An important property of a lens is its focal length, a fixed distance related to the distance between the image plane and the lens, z', and the distance at which an object can be clearly focused, -z. The equation for focal length, f, is defined as [1]:

$$\frac{1}{f} = \frac{1}{z'} + \frac{1}{-z}$$





The optical axis of the camera and lens system runs perpendicular to the image plane through the center of the optical aperture. In an ideal theoretical system, the aperture would be an infinitely small pinhole. However, in true applications, the aperture must have a finite diameter that is nonzero. Due to the wave nature of light, light rays will diffract at the edge of a very small aperture, causing the light to spread across the image. As aperture size decreases, the magnitude of deflection from the incoming ray increases. Figure 2.4 shows the optical axis relative to an aperture of diameter d with an incoming light ray at an angle of Φ . In photography, aperture sizes are designated with the letter f followed by a backslash and a number, where smaller numbers are associated with larger physical openings and larger numbers are associated with smaller openings in diameter. For example, f/3.5 is an example of a wide aperture and f/22 is a very narrow aperture. The reasoning for this seemingly counterintuitive numbering system is because the f-number of the lens is a ratio of the focal length to the aperture diameter [1]. Therefore, for a constant focal length, a larger aperture diameter results in a smaller ratio and a smaller f-number. Figure 2.4: Diagram of Lens System Aperture



In a composite lens system, the optical axis of each lens component should align. Creating a composite lens with multiple lenses will improve the quality of the final image by minimizing the effects of the defects and aberrations of each individual lens. A composite lens system can then be modeled with the thick lens model, Figure 2.5. This consists of two principal planes, which are perpendicular to the optical axis. The intersection points of the optical axis and the two principal planes are called the nodal points. A ray entering the first nodal point will exit through the second nodal point without changing direction. If these nodal points are coincident, then the lens is a thin lens.

Unfortunately, a perfect lens is impossible to create and there are defects and aberrations to be aware of. One such defect is vignetting, illustrated Figure 2.6. Vignetting is the smooth, gradual fall-off of image brightness toward the boundaries of the image. It is caused by the apertures blocking part of an off-axis light beam as it passes through the lens system. The farther from the optical axis a point in the image is, the less benefit it receives from the light gathering power of the lens. Vignetting correction will be discussed in greater detail in a later section.





Figure 2.6: Vignetting [1]



Aberrations of a lens will increase in magnitude as a power of the angle between the incident ray and the optical axis, so points near the optical axis will be more focused while points farther away will get "smeared" [1]. A diaphragm may improve the quality of the image periphery by blocking light that is entering from steep angles, but will also cause an increase in vignetting effects.

At the image sensor, it has been determined that the optimal size of the sampling area, or the sensor elements, should have dimensions equal to the spacing between each element. This lends itself to an optimal sensor configuration that is completely covered with sensor elements, leaving no gaps or overlaps where photons are missed or counted twice. This makes camera image sensors easy to manufacture in an optimal way.

2.2.2 CCD and CMOS Sensors

When selecting a camera for the HDRI system, one needs to consider what type of image sensor is best suited for the application. There are two types of image sensors available in consumer-grade cameras currently; a charge-coupled device (CCD) or a complementary metal oxide semiconductor (CMOS). Although the CMOS sensor was invented first in 1967, the CCD sensor has dominated in performance since its introduction in 1970 [14]. However, technological improvements beginning in the 1990s have allowed CMOS sensors to make significant improvements, and now CCD and CMOS sensors are equal competitors in the image sensor market. There are still strengths and weaknesses to each type that are inherent to their technology and structure that need to be considered when selecting the best option.

The basic architecture of CCD and CMOS image sensors will be briefly described here to understand some inherent performance differences that are attributed to their physical structure. Janesick and Putnam's 2003 paper [15] goes into further detail on the specific types of pixel and readout architectures used by CCD and CMOS image sensors, but that level of detail will be avoided here.

For both image sensor types, the entire area of the sensor is made up of individual elements, oftentimes called pixels, arranged in some pattern. For color image sensors, there are multiple element types that each have a unique filter to collect only a specific portion of light. Typically, these elements are equipped with red, green, and blue filters, although some may use cyan, magenta, and yellow, or some other fourth color [2]. For typical red, green, and blue arrays, the sensor elements are arranged in a specific pattern, called the Bayer pattern, Figure 2.7, composed of 50% green, 25% red, and 25% blue elements. During the imaging processing, the final color data for each pixel is derived by interpolating the color information from its neighboring pixels. Therefore, it is important to remember that the information provided in a single pixel in the final image does not directly correlate to information from a single pixel on the sensor array, but rather is an average or interpolation from a group of neighboring sensor pixels.





Image capture begins with photon collection at each pixel, for both CCD and CMOS type sensors. Each pixel collects these photons and converts them into a proportional amount of electrical charge via the photoelectric effect [16] [15]. In a CCD array, these electrical charges are moved through CCD shift registers to the readout node at the edge of the array, where the electrical charge is converted to a voltage signal. This voltage signal is then buffered and sent to the signal processing circuit outside the chip. The electrical charge is moved along by an applied series of pulses that are strictly timed to avoid signal loss [14].

CMOS arrays are different in that the charge-to-voltage conversion and signal buffering is performed within each pixel. This feature allows for memory-like organization that is addressable and requires only one power source, leading it to consume much less power than its CCD counterparts [14] [17]. The organization of the CMOS array does not require charge transfer over long distances as CCDs often do, which gives it more resilience to radiation damage [14]. This makes CMOS sensors more suitable for high-radiation environments, like space [15].

Historically, CCD sensors have dominated the image sensor market due to their superior sensitivity to light and low image noise levels. The relatively simple architecture of the CCD pixel allows the entire area to be sensitive to light. On the other hand, the complex architecture of CMOS sensors reduces the amount of light-sensitive area of the CMOS pixel [14]. The additional components in the CMOS pixel leads to greater absorption losses that are not present in CCD pixels [15]. Therefore, a CCD pixel of the same size will have a larger ratio of light-sensitive area to total pixel area, giving it a higher sensitivity, or in other words, an increased ability to collect photons. These architectural differences are also the reason why CCD sensors have the advantage of less noise. For the CMOS sensor, the on-chip amplifiers and transistors that allow for charge-to-voltage conversion and addressing capabilities come at the expense of additional noise [14]. Other benefits of CCD sensors include a larger dynamic range, higher resolution, better response uniformity across the array, and minimal dark current [14] [15].

Despite the aforementioned side effects, the integration of signal and image processing features onto the chip is considered a unique strength of the CMOS sensor. This on-chip integration eliminates the need for external circuits and devices, which decreases the volume and weight of the sensor [14]. This also helps the CMOS sensor achieve low power consumption in contrast to the CCD. In fact, CMOS sensors generally consume 1/8th the power that a CCD sensor would. This is because CCD sensors require several operating voltages for electron transfer and other signal processing steps. CMOS pixels operate with a single supply voltage. The other notable strength of the CMOS sensor is its high read-out speed, which can reach rates of 1000 Mpixels/s. In contrast, typical CCD speeds are less than 70 Mpixels/s [14]. The on-chip integration, low power-dissipation, and high readout speeds make CMOS sensors suitable for compact, portable, and video applications, like in cell phones and consumer-grade cameras.

Table 2.1 briefly summarizes the performance characteristics of both image sensor types[15][14].

| Performance Characteristic | ССД | CMOS |
|-------------------------------|---|--|
| Sensitivity | Entire CCD pixel area is sensitive to light | Reduced sensitivity because of additional on-chip components |
| Noise | Lower levels of image noise | More sources of image noise |
| Speed | Slow readout speed, 70 Mpixels/s | High readout speeds, up to 1000 Mpixels/s |
| Power usage | Requires multiple voltages and consumes more power | Requires only one supply voltage, 1/8th the power consumption of a CCD |
| Integration | No on-chip integration, external devices required | On-chip integration, reduces sensor size and weight |
| Blooming | Requires custom designs and fabrication to prevent blooming | Anti-blooming is intrinsic to CMOS structure |
| Dynamic Range | Twice the dynamic range of CMOS | Half the dynamic range of CCD |
| Dark Current | Less dark current sources | More dark current sources, not suitable for use in hot environments |
| Image Lag | Little to no image lag problems | Greater image lag problems |

Table 2.1: CCD vs CMOS Characteristics

CCD sensors have been considered a mature technology for several years, but there are continued efforts to enhance performance. Efforts are being made to reduce the power consumption of CCD sensors to be more competitive with CMOS sensors and restore performance losses that resulted from decreased pixel size [15]. CMOS sensors are still undergoing many developments to compete with CCD sensors for high-performance applications. Such developments include removing sources of dark current during the fabrication process, improving dynamic range with transistors and high-voltage operation, reducing pixel cross-talk (or bleeding) with shielding implants, and utilizing charge-transfer processes to eliminate image lag. There are also new developments in creating hybrid sensors, which combine the best features of CCD and CMOS technologies. This can be done by bonding a CCD pixel array to a CMOS signal-processing array, or fabricating a CMOS pixel array and CMOS readout array separately. Isolating the pixels and signal-processing circuits helps reduce image noise. Hopefully, with the continued progress and improvement of camera image sensors, consumer-grade cameras may become more reliable tools for lighting analysis.

2.2.3 RAW and JPEG Image File Types

Now that the basics of image formation and collection are known, the task of saving this data into a usable format can be addressed. There are typically two file output options on a camera, JPEG or RAW images. The most common and familiar file type is JPEG, which uses the .jpg file extension and is the default output for consumer grade digital cameras. JPEG images are compressed files to make file storage more manageable. However, those who have dealt with cameras more in depth will likely know there is another output option, called a RAW file. These files are uncompressed and take up considerable amounts of digital storage space. Unlike JPEG files, there is no single standard format for a RAW file, but rather there are a wide number of proprietary formats, including Canon's .crw and Nikon's .nef.

RAW image files contain exactly what the name implies; raw image data. The file contains unprocessed electrical charge information directly from the image sensors, either CCD or CMOS types as previously described. The only camera settings that affect the raw data are the ISO speed, the aperture size, and the shutter speed. Any additional color correction settings, like white balance or gamma, do not affect the raw pixel data. A key point of RAW images is that they contain greyscale image values, although the incident light

17

is filtered through the red, green, or blue sensor elements. What each sensor element reads out is a grayscale value of light within its respective spectrum, red, green, or blue [2].

In addition to the pixel data, RAW and JPEG files both contain metadata in the form of an EXIF (Exchangeable Image Format) header. This header typically contains information about the camera model, the shutter speed, aperture, focal length, and white balance used to create the image [2]. It is here that color correction settings can be used for later use without altering the raw data. RAW files will also include additional metadata needed to translate the raw data into an RGB image using a RAW converter.

The processing steps a RAW converter applies to a RAW image are as follows; demosaicing, white balancing, colorimetric interpretation, gamma correction, noise reduction, antialiasing, and sharpening. The metadata normally includes a decoder ring which conveys the arrangement of the color filters of the sensor to perform demosaicing, the interpolation of color information using data from each pixel and its neighbors. As mentioned earlier, the raw data is unaffected by the white balance setting in the camera, so it may be applied afterwards in a RAW converter, or the RAW converter may use its own white balance algorithm. RAW data also assumes a linear gamma response, so RAW converters will typically apply a gamma correction to match the gamma response of human vision. Edge-detecting and antialiasing processes in RAW converters will help minimize color artifacts and image noise that demosaicing cannot handle. The specific algorithms to perform each step vary between RAW converters, so the same RAW image file may take on very different appearances between them.

The difference with JPEG images is that these RAW conversion steps are all performed in the camera and then the resulting image is compressed using the JPEG algorithm before being output to the user. Some of the information contained in the raw data is thus thrown away and cannot be retrieved once the image is saved in the JPEG format. While JPEG compression is decent at preserving luminance data, it applies heavy compression to color data. RAW files commonly have 12-bits of information per pixel, equating to 4096 tonal levels of information per pixel, while JPEG pixels have 8-bits per red, green, and blue channel, equating to only 256 levels of information for each color [2]. The compression to 8-bits is what causes the loss of color information. There is clearly more flexibility in RAW files, as the data is left for the user to interpret instead of allowing the camera-integrated algorithms to throw away information that the user may want during JPEG compression. When using image data in HDRI applications like the one about to be proposed, using RAW data is preferred as there is less tampering done to the data that may be difficult or impossible to back out.

2.3 Background on High Dynamic Range Imaging

2.3.1 Motivation for HDR Photography

The human visual system is capable of perceiving a large range of luminance values, from dark, starlit nights to bright, cloudless sunny days. The reported range of exact luminance values varies slightly between sources, but it is generally agreed that the range covers 12-14 orders of magnitude, typically from 0.000,001 to 1,000,000 Cd/m^2 [13] [4] [3]. It is common to report the range of perceivable luminance values as a ratio between the maximum value and the minimum value, referred to as the dynamic range of the system. In a single scene, it is reported that the dynamic range of the human visual system is about 100,000:1, depending on the lighting conditions and adaptation state of the eye [3].

In contrast, digital cameras are limited to a dynamic range of about 100:1 to 1,000:1 on the high-end [18] [4]. This translates into images that will lack information in dark areas, which become underexposed images, or images that appear washed out in bright areas, or overexposed images. This severe limitation of camera dynamic range to capture images that match the range of human vision has led some researchers to develop techniques for creating high dynamic range (HDR) images. HDR images are created from a series of exposurebracketed standard low dynamic range (LDR) images to create a composite image that covers the total dynamic range of the LDR images in a single image. Figure 2.8 provides a graphical representation of the dynamic ranges of the human visual system, standard LDR images, and HDR images [3].



Figure 2.8: Dynamic Range [3]

Figure 2.9 shows an abbreviated example of how bracketed LDR images (top 3 images) at varying exposures can combine into a single composite HDR image, shown at the bottom. The composite HDR image was created from more images than the three provided, but only three are shown to show the extreme cases of exposures. The image on the far left shows an image that is under exposed, but captures high luminance details in the sky. The center image is a "properly" exposed image, capturing all of the details in the subject of the picture. On the far right, is the over exposed image. Although this example does not have prominent shadowed areas, over exposed images will capture details in darker, shadowed areas in the scene. Comparing the HDR image to the properly exposed image demonstrates the strength of HDR images, which is the ability to capture and display a wider range of luminance values to better represent human vision.



Figure 2.9: High Dynamic Range Image Example

While the obvious benefit of HDR imagery is to create more realistic or aesthetically appealing images like the one above, the resulting HDR images also contain pixel data that is now representative of the actual scene brightness of the pictured scene. When the HDR image is calibrated such that the pixel values represent luminance values, the HDR images may be called luminance maps. This is of particular interest to lighting designers and building researchers since luminance is a useful lighting metric for evaluating the quality and performance of lighting systems. Traditional point-by-point luminance measurements taken with a standard luminance meter can be time-consuming and are often too coarse for analyzing lighting distributions, so HDR luminance maps provide an alternate method for acquiring luminance data. A luminance map may come in a false-color form to provide a very quick evaluation of luminance across the entire scene, or it can be opened in software intended specifically for luminance mapping that will read out luminance values on a per pixel basis. An example of the false-color luminance map of the house image is provided in Figure 2.10.



Figure 2.10: HDR False Color Luminance Example

2.3.2 Overview of the HDRI Process

The reason why HDR photography is so accessible and affordable is because the only equipment truly required for high dynamic range imaging (HDRI) is a consumer-grade digital camera capable of manual control of aperture and shutter speeds. While not required, a simple tripod and a laptop equipped with camera tethering software will make the dataacquisition process much easier. One caveat is that a luminance meter with which to make
absolute reference measurements will be necessary in order to scale or calibrate the data to true luminance values. Although calibrated luminance meters are quite expensive, this device is a fundamental tool owned by many lighting designers and researchers, so it is not normally considered an extra accessory. Rather, it is assumed that those interested in HDRI already own a luminance meter, and HDRI with a consumer-grade camera will be intended to supplement this measurement device.

The HDRI process begins with taking a series of exposure-bracketed images, meaning that multiple LDR images are taken at varying exposures, usually in ascending or descending order. Using a tripod will minimize movement between the images to reduce errors due to image misalignment. Tethering the camera to a laptop with software capable of remote camera control will also help minimize movement and speed up the process.

Once enough images are taken to sufficiently cover the dynamic range of the scene, these images can be brought into a software program capable of generating HDR images. The process involves the determination of a camera response function to relate the LDR image pixel values to the estimated scene brightness values in the final HDR image. The details of this process will be discussed in a later section, Camera Response Functions. Once the HDR image is generated, there are a number of additional calibrations to perform to ensure accurate and reliable results. Absolute calibration is a simple linear scaling done in accordance to a known reference point measured with a luminance meter. More challenging effects to account for include vignetting, lens flare, and point spread function, and possibly image distortion.

Once the HDR image is calibrated, the final step in the HDRI process is tonemapping. Tonemapping is a process that compresses the full dynamic range of an HDR image to fit within the dynamic range of the display device, such as a computer monitor. Since tonemapping only affects the aesthetic value of HDR images, tonemapping will not be discussed in this document.

Figure 2.11, created by Jacobs [4], describes the HDR process pictorially.





2.3.3 HDRI Experimentation and Validation

Since the introduction of HDR photography techniques, there have been several attempts to utilize HDR images in lighting analysis applications. Since HDR photography was not created for the purpose of luminance data acquisition, it is necessary to evaluate and validate the appropriateness of this technology for luminance mapping. This section will cover the most notable papers that validate the appropriateness of HDR imaging for lighting analysis.

2.3.3.1 CapCalc

The earliest example of utilizing a camera to acquire luminance data appears to be the CapCalc system developed by Rea and Jeffrey in 1990 [19]. The CapCalc, short for capture and calculate, is a system intended for luminance measurement and image analysis. The systems consists of a video camera equipped with a photopic filter to acquire luminance information with the same response as the human visual system. The images captured by the camera are then sent to a computer for storage and analysis. While this paper does not utilize HDRI in the sense of fusing multiple LDR images into a single HDR image, the concepts behind the luminance data acquisition with a camera are the same.

The camera calibration for the CapCalc system is handled somewhat differently than by current HDR systems. The authors used a CCD camera to take advantage of the assumed linear response of the CCD image sensor. The automatic gain control of the camera was manually disabled to ensure the linearity of the camera response, keeping the input to output ratio constant. Based on the spectral response of the camera as provided by the manufacturer, the authors were able to design a custom filter to ensure that the light incident on the image sensor was filtered to match the spectral sensitivity of human vision, meaning the values in the images are equal to true luminance values.

They also established a correct zero baseline in their camera by taking a number of images in complete darkness to determine the minimum noise threshold. Ideally, these images should produce pixels with zero output, but some image noise causes some pixels to read higher than they should. From this calibration, the authors were able to determine appropriate values to subtract out of the images to eliminate clipping of data on the lowrange.

Vignetting was determined by taking images of the inside of an integrating sphere with the camera focused at infinity to minimize effects of surface imperfections. Images were obtained for every aperture-focal length combination of the camera so that the inverse of the images could be used at a later point to accommodate for vignetting. It was also pointed out that the quality of the optical components of the camera affect light diffraction and spreading effects. They were able to determine that objects must occupy at least 2% of the image frame to avoid errors.

While current HDR systems do not rely on a single image for luminance data acquisition or assume perfectly linear responses from the camera, this paper brings up many important issues that need to be addressed when trying to extract data from camera images. These issues include vignetting, hot pixels, lens flare, and spatial resolution.

2.3.3.2 Margin of Error of HDRI Luminance Measurements

Much later, in 2005, Anaokar and Moeck [11] published experimental data to verify that HDRI is an appropriate tool for measuring luminance data using a CCD camera and Photosphere to generate HDR images. Their experiments tested the accuracy of HDR luminance measurements under varying conditions for colored surfaces and lamp spectra. Munsell cards of 6 colors (gray, red, yellow, green, blue and purple) were used as the colored surfaces. These cards were then tested under three different lamp types with different spectral properties; fluorescent, mercury vapor, and metal halide. With each lamp type, high and low illuminance level tests were performed at 478 lx and 88 lx, respectively.

An N6 gray card was used in all of the tests as the reference point for absolute calibration of the final HDR images. An N9.5 white card was used to calibrate white balance under each lighting condition. Optical vignetting was determined by photographing a sheet of white paper of uniform illuminance with a small grey card in the center of the image frame. Vignetting was determined as a function of radial distance from the center of the image frame. Spatial resolution, or the point spread function, was determined by taking photographs of a printed image of white and black bands of varying widths. This allowed the authors to determine the minimum spatial resolution between two distinct objects in the image, which was found to be 4 pixels, or 0.3241 cm, for their specific target distance.

The finals results from this test showed significant errors in cool colors, like blue, green, and purple, while warmer colors, like red and yellow, showed lower magnitudes of error. The errors were calculated between the known reflectance values of the Munsell cards and the derived reflectance values from the HDR image. The magnitude of error was also shown to increase with color saturation and dark surfaces showed a tendency to be overestimated. The errors were also shown to be independent of illuminance levels and the spectrum of the light source. The overall conclusion from these tests is that HDR images of surfaces of low chroma and saturation, meaning neutrally colored surfaces, could be measured within a 20% margin of error.

Moeck [20] performed a second test of HDRI accuracy using a mirror ball and a CMOS camera a few years later. The use of the mirror ball is very unique to this study, and no other study seems to have replicated this method. The purpose of the mirror ball is to overcome the limitations of fisheye lenses, which include vignetting, lens flare, and a limited hemispherical view. The primary disadvantage to the mirror ball is that the reflected environment has small resolution. While JPEG images were used in this study due to memory constraints, the author suggests using RAW image files to avoid errors and information loss associated with quantization that occurs during image compression.

Rather than replicating the author's previous test with an indoor scene, this test involved an outdoor scene under sunny conditions, located on a roof in Arizona, to achieve uniform illuminance conditions. The targets used included 16 matte gray cards and 140 color cards. Two different gray cards were used as reference points, an N5 gray card and a lighter N7.5 gray card. The errors were calculated for each reference point, in order to compare the effects of scaling HDR images to a particular reference point. This time, the error calculations were based on illuminance values, which were derived from the luminance values measured, rather than reflectance values.

The results from these tests showed agreement with the authors previous test, in that dark surfaces with low reflectances tend to be overestimated. Conversely, it showed that light, high reflectance surfaces tend to be underestimated. Again, errors increased with color saturation and cooler colors. Surfaces with mid-range reflectances from about 17% to 60% showed errors within a 10% margin, independent of the reference point or hue. This provides confidence that HDRI shows promise as an acceptable luminance data acquisition technique for typical architectural spaces. The comparison between the calculations done for each reference point also showed that local calibration points will improve errors, meaning that for a scene with light surfaces, a light-colored reference point should be used and vice versa.

The CMOS sensor showed shortcomings in estimated luminance values for saturated

blues and greens, but comparisons to the previous HDRI tests done with a CCD camera showed a significant improvement in error from the 2005 results [11] with the CCD camera to the 2007 results [20] with the CMOS camera. The maximum error reported for the CCD camera is 32.95%, while a 16.89% maximum error was reported for the CMOS camera. Unfortunately, due to the different conditions under which each test was performed and because the tests were performed years apart, this cannot serve as a true comparison of the performance between the CCD- and CMOS-based cameras, although it may suggest that improvements in technology over time will bring forth more reliable results and improve the accuracy of luminance measurements using HDRI.

Around the same time, Inanici [10] performed another test to verify the accuracy and applicability of HDRI to lighting analysis using a Nikon camera equipped with a fisheye lens. Like Moeck and Anaokar, Inanici produced HDR images using Photosphere. Luminance data was extracted from the HDR images using proprietary MATLAB code, named HDRLab, to derive CIE XYZ values from the RGB values of the image pixels. HDRLab was also used to apply a vignetting filter to correct for vignetting effects, which was shown to have a maximum luminance loss of 23% at the periphery for the selected aperture of f/4.

Images were taken in a wide variety of test scenes, including a black room with no daylight, a typical office space, and an outdoor scene. In all test scenes, a card with 24 varying greyscale patches, a set of four contrasting grey-black-white targets, and a Macbeth ColorChecker card were used as reference points. The reference points were measured absolutely with a Minolta LS-110. These measurements were compared to the minimum, maximum, mean, and standard deviation of the corresponding pixel areas in the HDR images.

In interior tests, seven different light sources were investigated for effects of light source spectra on luminance measurements. The seven light sources tested included an incandescent lamp, a tungsten lamp, a 6500K T12 fluorescent lamp, a 3500K T8 fluorescent lamp, a 3000K T5 fluorescent lamp, a metal halide lamp, and a high pressure sodium lamp. The temperature designations indicate the Correlated Color Temperature (CCT) of the light. Even with the wide variance of spectral content of the light sources, the margin of error underneath all sources was within 10%.

It was concluded that HDR images are a reasonably accurate tool for collecting luminance data across a large field quickly and inexpensively and for generating false-color luminance maps within 10% accuracy, but is not a replacement for the traditional luminance meter. The author also acknowledges that the Minolta LS-110 luminance meter is not without its errors, but these errors are defined [10]. In the HDR images, greyscale targets showed an average error of 5.8%, while color targets showed an increased average error at 9.3%. Darker targets tended to be overestimated, likely due to light scattering effects in the lens. These increased errors for saturated and dark colors are consistent with the results obtained from Anaokar and Moeck's report [11].

2.3.3.3 Calibration Factor

Calibration factor refers to a scalar number that can be applied as a multiplier to the values in an HDR image to linearly scale all of the values so that a specified group of pixels matches the measured value of the correlating region in the physical scene as taken by a luminance meter. It is determined by taking the ratio of the measured luminance value to the luminance value produced in the HDR image. Chung [21] conducted an experiment to verify that this calibration factor is independent of ambient daylight levels in an interior daylit scene.

The tests were performed on 6 different days in a classroom illuminated only with natural daylight, under varying clear and cloudy sky conditions. The camera used for taking the HDR images was a Canon EOS 350D equipped with a Sigma zoom lens. An X-Rite Color Checker card was used as the reference target to determine the calibration factors for each HDR image. The measurement trials for each HDR image took 10 minutes each, with 3 sets of physical luminance measurements taken before the image capture process and 3 more sets of luminance measurements taken after. The 6 luminance measurements were averaged for each color to determine the reference luminance value of each color square. The illuminance was also measured at each corner of the color checker card as an indication of the illuminance level in the classroom. The HDR luminance values were calculated based on their RGB pixel values and the following luminance equation:

$$L = 179 * (0.265 * R + 0.670 * G + 0.065 * B)$$

The calculated ratio between the averaged luminance measurements and the HDR pixel values were recorded as the calibration factors for each color, in each HDR image.

A correlation test was performed to determine if there was any statistically significant relationship between the calibration factors and their associated daylight levels. Using a 99% confidence interval, none of the reference color squares showed any correlation between the ambient illuminance levels and the resulting calibration factors. The authors concluded that the calibration factors are independent of illuminance levels, so an overall calibration factor was determined by averaging all of the measurements taken across all of the different tests. The authors suggest using a specific color as the calibration point if the scene has a dominant color, but otherwise averaging the calibration factors across all of the colors works for a neutral or balanced scene. The overall averaged calibration factor showed that the variation of error was within 5.7% [21].

2.3.3.4 Camera Settings Optimization

Determining the proper camera setting and scene conditions to take the best pictures for HDRI can be difficult. While the papers mentioned here often include recommendations, they tend to vary between author and camera. In 2011, Cai and Chung [22] attempted to determine the optimal camera settings for the purpose of HDRI, specifically camera aperture, focal length, ambient light levels, and the number of LDR images used to create the HDR image. The specific camera used for their tests was a Canon 350D, fitted with a Sigma zoom lens. The authors investigated six aperture sizes, three focal lengths, two ambient light levels, and four ranges of LDR exposures.

The tests were performed in a neutrally colored interior space, lit only by fluorescent troffers. Three sets of greyscale targets were placed at varying depths in the scene (foreground, middle ground, and background), all placed perpendicular to the camera and arranged so that they were located within the center of the image to avoid vignetting errors. A color checker card was placed in the middle ground. A target of four background-foreground color combinations were also placed in the scene to measure point spread function (PSF) effects. The camera was focused to the background greyscale targets.

For all tests, the camera was set to Fluorescent White Balance, ISO 100, and saved into large JPEG files. The six aperture sizes tested were f/4, f5.6, f/8, f/11, and f/22. The three focal lengths tested were 10mm, 14mm, and 20 mm. The two ambient light levels were determined by full output of the fluorescent troffers or a low output. This resulted in 32 different tests to be carried out. The four sets of LDR images used meant using 4-5, 8-9, 12-14, or all 18 LDR images to compose the final HDR image across a shutter speed range from 1/4000 to 30 seconds. Therefore, for each of the 32 tests, four HDR images were created, giving a total of 128 HDR images to analyze.

All 128 images were calibrated based on four luminance measurements taken at the front, middle, middle-back, and back targets. 16 vignetting curves were also determined for the different lens and focal length combinations. The 5-order polynomials are all provided in the paper [22]. To calculate the magnitude of error for each HDR image, luminance measurements were compared to the 73 absolute measurements taken of the targets.

Regarding the number of LDR images used to create the HDR images, there was a strong correlation between decreased errors and the number of images used. This meant that using only 4-5 exposures consistently resulted in the greatest errors. The difference between 12-14 and 18 exposures was minimal.

For aperture size, the smallest aperture size, f/22, yielded the most significant errors,

while f/5.6 showed the smallest error for all focal lengths and ambient lighting conditions.

For apertures equal to or larger than f/5.6, focal length showed no impact on the magnitude of error. For small aperture sizes, longer focal lengths appear to reduce the magnitude of error. The authors claim that this is due to an increase in the aperture diameter to alleviate diffraction effects. However, the use of longer focal lengths reduce the accuracy for measuring luminance from light-emitting surfaces.

For targets closest to the camera, lower ambient light levels appeared to have a negative effect in accuracy. All other targets showed no significant difference between light levels. Local calibration factors proved to provide the least error, but averaged global calibration factors showed no significant difference for the mean error, except for the closest set of targets. The best results occurred with high ambient light levels and local calibration factors.

This study is yet another example that demonstrates the applicability of HDR imaging for luminance measurement of architectural spaces, showing an acceptable error range within 10%. The mean error for greyscale targets was 2.8%, 1.5% for black targets, 10.1% for colored targets, and 6.6% for light-emitting surfaces. The authors recognize that this test does not encompass all of the factors that may influence the accuracy of HDR luminance measurements, and a second round of tests dealing with light source spectral effects is to be carried out and published at a later time.

2.4 Applications of HDRI

Having been shown the validity and appropriateness of HDRI as a lighting analysis tool, many researchers have investigated into the further uses of HDRI. Luminance mapping is the most basic feature, but it can be used as a fundamental tool for more advanced applications. These applications include image-based rendering and lighting simulations for daylighting and electric lighting, daylight glare analysis, advanced photosensors for electric light control in buildings, illuminance calculations, and deriving luminance intensity distributions of luminaires.

2.4.1 Illuminance Analysis

Following their study to validate the applicability of HDRI for lighting analysis, Moeck and Anaokar [23] proposed using HDRI for illuminance analysis of architectural scenes. More specifically, the authors wished to determine the contributions of particular objects or features in the scene, like a specific window or wall feature, to the final illuminance at the camera. The claim is that such a technique may allow researchers or engineers to develop new lighting metrics in regards to lighting quality or glare.

Using the luminance maps generated from the HDRI process, the authors grouped pixels together based on the geometry of the space and isolated them using image masks. The illuminance contribution of each feature is derived from the luminance of each of the grouped pixels, L_i , the incident angle between the pixels of interest and the image center, δ , and the solid angle subtended by the pixels, $d\omega_i$. The equation for the calculation is included below:

$$E = \int L_i * \cos(\delta) * d\omega_i$$

The accuracy of calculated illuminance was verified in a controlled lab scene. The example used to test and present their results was the interior of the dome of the State Capitol building in Harrisburg, Pennslyvania. The specific features they were interested in included the surface of the dome, the gilded features that decorated the dome, and the crown molding around the dome. Their results found that the gilded features contributed the most to the illuminance at the camera, and that the pictured light sources contributed much less than expected.

2.4.2 Sky Dome Imaging and Image-Based Lighting

With the challenges associated with mathematical models in daylighting simulations, the concept of using HDR images to model sky dome luminance distributions is an appealing idea. Currently, daylighting simulations are limited to 15 sky models developed by CIE that cover a range of sky cover conditions, from perfectly overcast to completely cloudless. However, these mathematical models do not consider geographic location or other site-specific factors. These models are also quite limited in data resolution and do not have the ability to render cloud boundaries. Researchers believe that HDR sky dome images can overcome some of these challenges by capturing high-resolution, site- and time-specific sky data.

With Image Based Lighting (IBL) programs, a calibrated HDR image, or luminance map, can be used by the IBL simulation program as a light source, where each pixel provides luminance data of the source. The experiment [24] involved taking HDR images of the sky dome from the roof of the test space and HDR images of the interior of the space to create two models, one standard physically based rendering (PBR) and IBL, to simulate the space. The two models using the CIE sky models and the captured sky images were compared to evaluate the validity of HDR+IBL sky modeling.

Due to the high dynamic range of the sky dome with a visible solar corona, two apertures in combination with a neutral density filter must be used for the camera to capture the dynamic range. The absolute calibration for the sky dome images was done by measuring the horizontal illuminance at the camera during the image capture. The luminance values of the hemispherical image must be scaled so that it produces the same horizontal illuminance value at the camera position. Macbeth Color charts were used to correct for color shift caused by the neutral density filter.

The models were created in Radiance, which is able to utilize both the CIE sky models and the HDR images for IBL. In the PBR model, a mathematical sky model and a sun model were used in combination with geographical and time data. In the IBL model, the fisheye image of the sky dome was projected onto a hemispherical surface, which acts as a light-emitting surface in the model. To avoid errors in the calculations of the sun as an indirect source, the sun should be extracted from the HDR image and placed into the model as a separate, concentrated light source.

A strong advantage of the IBL method is the ability to account for surrounding struc-

tures and vegetation without explicit modeling. Explicitly modeling surrounding structures and vegetation in traditional modeling can be quite labor intensive and computationallyexpensive. For existing buildings, rooftop images may not accurately depict surrounding objects that may influence the lighting within the space, so accuracy can be improved by taking vertical fisheye images at the window surface. This then provides nearly the same view of the outdoors from the window.

2.4.3 Daylighting Control

HDR-based lighting controls is another area that some researchers have started to develop. Current daylighting control systems depend on photocells to report the lighting levels of a space back to the controller. However, there are many known complications with the execution of installing and commissioning photocells so that they function properly. Sarkar and Mistrick [25] aimed to prove the concept of using an inexpensive CMOS camera image sensor and HDR images to replace the traditional photocell in a daylight control system. Their paper describes the calibration procedure and control algorithm developed in their control system, named CamSensor.

The control system is composed of a small camera module equipped with a CMOS sensor, a computer that oversees the operation of the entire system, the CamSensor software that contains the algorithms for both the HDR capture and lighting control, and a DALI controller. The primary advantage of using the camera as a sensor is the ability to simultaneously monitor multiple target points within the camera's field of view (FOV). This system is also capable of tolerating direct views of luminous sources, accounting for surface reflectance changes, and detecting motion or occupancy. The downside of this system is the intensive and careful calibration process required for the system to function properly.

This system was tested in a large classroom with 12 target points. The camera was positioned on the ceiling such that it had unobstructed views of all 12 targets. At the time of the experiment, the system was manually operated, but the finalized product would be a fully-automated process. The camera was calibrated in a similar manner to other HDR systems, deriving the camera response function and obtaining absolute measurements at each of the 12 target points. However, the absolute calibration process is much more important and detailed for the CamSensor system. The absolute calibration was performed under multiple lighting conditions, including electric light only and typical daylight conditions. Each target was also tested with a Lambertian patch and a non-Lambertian patch. These target points were used to mark the output percentage levels of the electric lights.

Once calibrated, the control system works with a cyclical 5-step process. The first step is for the camera to acquire images of the existing conditions. The luminaire ballasts report their dimming level associated with the conditions of the image at the same time. The daylight contribution is derived by subtracting the expected electric light contribution from the reported dimming levels. New dimming levels are determined based on the target illuminance level and the projected illuminance under new dimming conditions. Once the appropriate level is found, the new dimming levels are sent to the fixtures to update. This process is repeated on specified time intervals.

Their results showed increased errors when luminance values were low, due to poor Signal-to-Noise (S/N) ratios, and greater fluctuations of dimming levels for target points farther away from the camera. This suggests higher-quality cameras should be used, or to add additional constraints to the algorithm to prevent too many fluctuations in dimming levels. The system also works by minimizing deviations of all 12 target points, so all 12 points may not reach their target illuminance levels simultaneously. The success of this system is also limited by the accuracy of the HDR imaging itself, the camera location, and the need for automation of the system.

2.4.4 Photogrammetry

Another interesting capability of photographic images that may prove useful to building and lighting researchers is photogrammetry. Photogrammetry is the use of photographic images to obtain physical 3D coordinates of objects and points within an image. Cai [26] recently published a paper to combine the techniques of HDR luminance measurements with photogrammetry measurements to obtain pixel level data that include both physical and luminance information of a scene. While there are dedicated devices for photogrammetry, they are not suitable for obtaining luminance measurements, so HDRI methodology with consumer-grade digital cameras provides an accessible and affordable method for obtaining light and geometry data simultaneously. The ability to capture luminance and geometry measurements synchronously will result in decreased measurement time and errors associated with misalignment. The author's personal belief is that such a system may help bridge the gap for obtaining more effective glare metrics.

There is some additional required equipment to capture photogrammetry data from a consumer grade camera. The author designed a special tripod capable of measuring the yaw, pitch, and roll angles of the camera, which are important for establishing a properly oriented coordinate system. A laser distance meter was also needed to obtain absolute geometry measurements to which the image must be geometrically calibrated to, much like luminance maps need a luminance measurement. Four reference points should be measured ideally at the four corners of the target plane. Extensive preliminary tests were performed in the lab to calibrate the camera system before taking measurements of a test scene.

The entire process is summarized into 11 basic steps. These included field preparation, setup and leveling, determining the reference coordinate system, taking the bracketed images necessary for the HDR imaging, absolute field measurements, generation of HDR images, luminance calibration of HDR images, luminance data retrieval, lens distortion correction, photogrammetric calibration, and finally collection of the synchronous data. The details of this process can be found in [26]. The author validated this process in a controlled lab setting. The lab results showed luminance errors between 6.2 to 1.8% and increased errors for smaller and oblique (surface that are not perpendicular to the camera) target areas. The average error for geometric measurements was 18.2 mm, but could be as large as 132.4 mm for smaller targets. The average error was deemed acceptable as it is close to expected errors of conventional photogrammetry tools. It was also found that prime and zoom lenses perform much better than fisheye lenses due to minimal image distortion.

Again, the author clarifies that this topic is still in need of further development. Six major challenges were identified, including measurements of curved surfaces, limited recording ability of camera sensors, increased errors for far-range objects, errors due to noise and defects of low-quality cameras, laborious calibration processes, and limited software capabilities for data treatment. Despite these challenges, there appears to be promise for this type of measurement method to facilitate research topics that require this type of data collection.

2.4.5 Glare Evaluation

Glare evaluation appears to be one of the most desired applications for HDR luminance mapping. Glare is a notoriously difficult phenomenon to quantify and predict, and there are currently a multitude of glare metrics that attempt to do so. However, there seems to be no consensus as to which is the preferred glare metric to adopt as a standard. It is the hope of many researchers that luminance maps derived from HDR images will provide the key to obtaining a more accurate and reliable method to quantifying and evaluating glare in a useful way.

Before any new luminance-based metrics or control systems can be created, a solid understanding of what luminance conditions human occupants prefer in their built environment is needed. Developing a better understanding for what building occupants prefer and can tolerate will help guide the lighting industry to make informed decisions to create effective lighting solutions. In 2010, Van Den Wymelenberg, et al [27] conducted a study with 18 student participants in a typical sidelit office space to begin this process. The students were asked to perform basic computer and paper tasks in the office for 20-30 minute sessions. Only natural daylight was considered, and the occupants were allowed to control the interior shading devices to create their preferred luminous environment for the tasks, but also create what they would consider just disturbing. Once the participants created the appropriate luminous environment, HDR images were taken of the room from the point of view of the participants.

The 36 images (2 conditions for each participant) were then analyzed to identify any relationships between the preference ratings of each participant and the luminance patterns in the space. Conventional glare metrics based on absolute thresholds, scene-based mean luminance thresholds, and task-based mean luminance thresholds were tested to determine if they could accurately identify the difference between comfortable and disturbing scenes. An interesting find of this study showed that while many daylighting guides recommend against direct sun patches in a space, occupants seemed to prefer some direct sunlight in their working space. Direct sunlight only caused disturbing scenes when it covered the task plane. Results also showed a weak correlation between ceiling illuminance and occupant preference, and the strongest correlation with illuminance measured at the monitor.

Regarding the performance of the glare metrics, absolute threshold-based metrics were more useful when a percentage of allowable area for above-threshold luminance was introduced. In a real scene, sources or areas that are of high-luminance may not necessarily be considered a source of glare to an occupant and may be desirable highlight or sparkle. A discriminating factor between glare and sparkle then becomes the spatial extent of the source, so introducing an allowable percentage provides some forgiveness to small, tolerable sources of high luminance and directs focus to larger, more overwhelming sources of glare. Of the three-types of luminance-based glare metrics, task-based mean luminance thresholds showed the strongest correlation to predicting occupant discomfort. This is because the human visual system can be quite insensitive to large luminance differences across the entire field of view (FOV), but is increasingly sensitive towards the foveal, or central, region of the eye.

While the study shows some interesting results that should help guide further research involving luminance-based metrics and the use of HDR to evaluate glare, the authors admit that the extent of this study is rather limited. This study only involved participants who are young adults with good vision and only tested daylit environments in a south-facing office during the winter. Further research should incorporate different age groups, other times of day and year, and other environment configurations.

A later study [28] focused on evaluating current glare metrics for assessing discomfort glare of non-uniform fluorescent sources using HDR photogrammetry. These results were then compared to subjective evaluations collected from 67 human participants. The four glare metrics used for this study were Visual Comfort Probability (VCP), the British Glare Index (BGI), the CIE Glare Index (CGI), and the Unified Glare Rating (UGR). The challenge with these glare metrics is that they were developed on the assumption of uniform light sources, yet are still used for non-uniform sources.

The experiments were conducted in a windowless room with three parabolic louvered fluorescent luminaires, which were tilted towards the subject to exaggerate glare effects. For the control test, the louvers were removed, and a translucent paper covered the opening of the fixtures in an effort to create uniform sources. Two different levels of ambient light were tested, high at 137 lux and low at 51.5 lux. The low levels were used for the control test while the higher level of light was used for the HDR photogrammetry tests. Measures were taken to calibrate the HDR images for both the luminance and photogrammetry measurements. With each lighting level, four different light source configurations were tested; one setup for each luminaire turned on individually and the fourth test with all three luminaires on simultaneously.

For the subjective evaluations, each participant sat through two 20 minute sessions, carried out on different days. The first session was performed under the control conditions. The participants were seated so that their point of view matched that of the cameras. The participants were then asked to focus on a target at the back of the room and rank the level of glare on a scale from 0 to 8, where 0 was imperceptible and 8 was intolerable.

The luminance map data in conjunction with the photogrammetry data allowed the

authors to calculate the glare ratings for each source on a per-pixel basis. The per-pixel glare ratings were summed up to find the glare ratings of each testing scenario. Each of the four glare metrics were assessed on their additivity, sub-divisibility, and predictability. If a glare metric holds under additivity, it means that if multiple discrete sources were evaluated individually, their sum should yield the same results as if the sources were evaluated as a whole. Sub-divisibility means that if a single large source is divided into subsections, it should yield the same glare ratings for any number of subdivisions. All four glare metrics proved to have valid additivity properties, but only UGR and CGI ratings held under subdivision. The subjective results varied significantly, and statistical tests showed significant differences between the calculated ratings and the perceived ratings. The results for all four metrics showed an overestimation of perceived glare in both the non-uniform test and the uniform control test.

It should be noted that one of the authors of this study also published the paper on HDR photogrammetry, and it is a belief of the authors that the use of such technology will help overcome some of the challenges of glare assessment, just as they aimed to prove with this study. These experiments identified problems with current glare metrics for overestimating glare sensations for young adults and the invalid sub-divisibility of BGI and VCP.

Another group of researchers [3] published a preliminary study to develop a new daylight glare analysis method. This new method proposes breaking up glare into two categories; absolute and relative glare factors. Absolute glare is caused by an excessively bright source that causes discomfort or even damage to the visual system and is based on absolute values instead of contrast ratios. Relative glare factor is caused by a glare source that can be adapted to with the appropriate background conditions and is based on contrast ratios. There are some standards and guidelines that attempt to define appropriate thresholds for both absolute luminance thresholds and contrast thresholds, including NUTEK and the IES.

This study used five glare metrics, Daylight Glare Probability (DGP), Daylight Glare Index (DGI), UGR, VCP, and CGI along with HDRI measurements and Radiance simulations to evaluate an existing space under daylight conditions. The Radiance simulations were used alongside the HDR images, not as a comparison between the two but rather to test consistency between glare ratings when observed using different techniques. Both the Radiance simulations and the HDR images were analyzed in a software called Evalglare to detect glare sources. The results showed inconsistent ratings between the five glare ratings and even for the same metric between the HDR image and the Radiance simulation. This study did not validate the predictability of the glare ratings to subjective evaluations. The results help identify some major problems of current glare metrics, including their inconsistencies, their unrelated thresholds and ranking systems, their complexity that limits users from understanding how to use them properly, and their inability to generate consistent results for the same space rendered in HDR images or Radiance simulations.

As an alternative to the conventional glare metrics, the authors propose dividing glare into absolute and relative components. It is suggested that a previous recommendation of $2000 \ Cd/m^2$ is too low of a threshold for absolute glare factors, but experimental validation is required before the authors can identify a more appropriate threshold value. At the time of publishing, the experiment was still in a preliminary, conceptual phase, further experimentation with human subjects and subjective responses to glare conditions will be required to validate the new proposed method.

2.4.6 Photometry

Recently, HDRI was used to evaluate the uniformity of luminance across an LED chip [29]. Unlike previous studies that have been focused on analyzing architectural spaces, this paper uses HDRI to capture a direct, up-close view of a very small and high-luminance light source in an effort to characterize its properties. The luminance values derived from the HDR image were compared to conventional measurements from a spot-luminance meter and derived calculations based on direct illuminance measurements. The problem with the spotmeter was that the LED was smaller than the acceptance area of the meter, which will cause lower readings, and the luminance of the LED exceeded the range of the meter, so a neutral density filter was required. Because of the unreliable spot measurements, the true luminance value was obtained by using the inverse square cosine law and illuminance measurements.

The authors [29] establish that RAW image files are the preferred file format for HDRI analysis to avoid errors that may result from JPEG compression. Once the HDR images were generated, the analysis was performed using pvalue, a program in Radiance that reads pixel-by-pixel values from an HDR image. The luminance pixel values were sorted into a frequency plot to demonstrate the luminance distribution across the LED. The results showed only a 2.9% difference from the mean luminance value from the HDR image to the calculated luminance value, indicating that HDR images are a viable method for measuring luminance of small, high-intensity sources. HDRI also proves its unique capability of visualizing the luminance pattern across the luminous surface that cannot be detected with any other methods.

Photometry based on imagery is also an emerging application for HDRI. Bellia, et al, in Italy [5] utilized a Canon 20D camera with a CMOS sensor and a Canon zoom lens fitted with a photopic filter to capture the images. This camera system was then used to derive luminous intensity distributions of a luminaire by taking images of the beam pattern on a perfectly diffusing, or Lambertian, surface. This method relies on the assumption that the surface obeys Lambert's law of diffuse reflection, meaning that luminance is now constant with viewing direction and depends only on the illuminance measured at the surface, E, and the reflectance of the surface, ρ .

$$L = \frac{\rho * E}{\pi}$$

The measurement process for obtaining the luminous intensity distributions (LID) requires careful setup. Stray light must be eliminated, or at the very least minimized, in the testing space, which was done by covering all other surfaces in the dark room with a fabric having 3.0% reflectance. The authors' test space was also limited to a distance of 1 meter

between the light source and the Lambertian surface, so the only light sources that can be reliably measured in their set up must be confined to a 10-cm maximum dimension. The light source and camera were mounted side-by-side on a testing bench, both aimed towards a rectangular Lambertian screen such that their optical axes strike the same point on the screen. An image of the setup is included in Figure 2.12 [5].

The camera captures the luminance map of the beam pattern produced on the screen by the light source, which can then be converted to an illuminance map using Lambert's Law. From the illuminance map, luminous intensity values can be derived by rearranging the inverse square cosine law, $E = \frac{I * \cos(\theta)}{d^2}$. The data obtained from the camera system was compared to the the luminous intensity distribution as provided by the lamp manufacturer, which was obtained with a standard goniophotometer. The light source tested in this paper was an Osram LED PAR lamp with a 25° beam, which is sufficiently small enough to fit within the size constraints of their testing apparatus. The results found a maximum deviation between the HDRI-obtained LID and the manufacturer-provided LID to be within 20%. The authors conclude that this method for photometry measurement is not intended to replace standard goniophotometry as it is not as accurate, but it does provide a quick and affordable way to generate approximate LIDs of small light sources during the design and construction phases of light source development.



Figure 2.12: HDR Photometry Setup [5]

2.5 HDR Image Calibrations

The knowledge of how images are formed and saved digitally shed light onto why the following calibrations are important to perform before using HDR image data for lighting analysis.

2.5.1 Camera Response Function and Creating HDR Images

The first step in creating HDR images from JPEG images is to determine the camera response function, which describes the mapping of scene radiance to the pixel values of the digital image. Ideally, this relationship would be linear, meaning the pixel value of the final digital image is directly proportional to the amount of light received at the corresponding sensor element, which is true only in RAW image files. At the sensor (CCD or CMOS) level, the relationship is, for the most part, linear with exceptions at the saturation and minimum thresholds. However, a number of nonlinearities, including analog-to-digital conversion (ADC) and gamma correction, are introduced before the image is stored as a useable JPEG file. It is the purpose of the camera response function to "un-do" these nonlinearities. Camera response functions are unique for each individual camera, but the response function only needs to be determined once for a given camera, as the in-camera post-processes should remain constant. Once this function is known, it can be applied to a set of LDR images, which can then be used to create a composite HDR image. Because RAW image files do not undergo these same nonlinear processes, there is no need to derive a camera response function for HDR images created from RAW images.

Figure 2.13: Image Acquisition Pipeline [4]



Figure 2.13 [30] is a visual representation of the image acquisition pipeline for film and digital cameras. It starts with the scene of interest. The pictured objects in the scene send energy towards the camera lens, which is the radiance, L_e , of the scene. After passing through the camera lens, the energy hits the camera sensor. The incident energy onto the sensor is called irradiance, symbolized as E_e . The amount of light exposure, X, on the sensor will depend on the shutter speed. Lower exposures will result from faster shutter speeds, while slower shutter speeds will allow more radiant energy into the sensor for higher exposures. The incident energy is translated into an electrical charge, which is linearly proportional to the amount of photons received, up until the saturation point above which all values are set as the same maximum value. The electrical signals are then converted to digital values, in a process referred to as analog-to-digital conversion, ADC. There is a potential for added nonlinearity during this process. Finally, the camera manufacturers are likely to impose additional nonlinear mappings in an effort to mimic a traditional camera film response and make the final image more visually appealing, such as gamma correction. These digital values, Z, are then saved and stored into the final digital image. The camera response function is responsible for translating the final digital values, Z, back into exposure values, X.

Image noise is another important component to consider when looking at camera response functions. Noise is a term for unwanted or false signals that appear in data. Image noise in particular comes in the form of random brightness or discoloration patterns that do not exist in the real scene. There are a number of sources for image noise, including dark current, quantization and dequantization, and ADC noise [31] [32] [33]. Dark current is a common phenomenon in digital cameras where electric current flows through the photodiodes of the image sensor, even in the absence of received photons. This creates a false signal of additional photons that are not present. Quantization is the process of constraining a continuous range of data values, the scene radiance values, into a much smaller set of discrete bins, such as the 0-255 integer values of 8-bit digital pixels. Dequantization is simply the reverse of this process. This binning can result in the loss of data. For example, in especially bright or dark regions of an image, scene radiance values may vary significantly but may get binned into the same integer value [31]. These differences are not retrievable in the dequantization process, thus creating image noise.

It is the entire transformation process between the values of X and Z that is described by the camera response function. Figure 2.13 reveals that there are actually a few components to the camera response function, including the linear response of the camera sensor, ADC, and whatever nonlinear mappings, like gamma correction, that are added by the manufacturer [33]. Fortunately, determining each intermediate step is not of any particular importance, so it is acceptable to consider only the overall function, f. Figure 2.14: Optical System Diagram



Once the response function is recovered, the process for deriving the desired scene radiance values, L_e , is relatively simple. The inverse of the response function is applied to the digital values, Z, to obtain the exposure values, X. The irradiance values at the sensor, E_e , are easily computed if the exposure values, X, and exposure time is known. In the final step, it is assumed that the scene radiance values, L_e , are proportionally related to the computed irradiance, E_e , values. For imaging systems, the relationship between image irradiance, E_e , and scene radiance, L_e , is given by the following equation, where d is the aperture size, h is the focal length, and ϕ is the angle subtended by the principal ray from the optical axis [32].

$$E_e = L_e * \frac{\pi}{4} * \left(\frac{d}{h}\right)^2 * \cos^4(\phi)$$

There are a large number of published papers regarding the recovery of camera response functions, but a good number are based on the assumption that the response function is linear. This assumption is severely limiting and is not applicable to conventional cameras, film or digital, as they have been shown to have nonlinear film responses. The four papers discussed here treat the camera response function as being nonlinear, and are therefore the most accepted sources in camera response function recovery to this day.

Mann appears to have been one of the first to introduce the concept of combining

digital images at a conference in 1993 [34] and published a paper with Picard a year later [31]. The goal of this paper was to create a fully-automated algorithm that would combine multiple images of varying exposures into what they called a "true image." The "true image" is a collection of "analog photometric quantities", meaning each pixel value is treated as a light measurement taken by its corresponding sensor element, the effective "light meter" [31]. This idea coincides with what many lighting designers intend to use HDR photography for today.

The Mann and Picard algorithm is comprised of four main steps. First, the nonlinear mapping is determined point-wise using the "self-calibration" method, as explained within [31]. This generates a single response curve per image. Next, the response curves are mapped into one final response curve. A certainty function can then be found by differentiating the final response curve. This certainty function describes the level of confidence that information can be accurately recovered, where the steeper portions of response curves correspond to values that can be more accurately recovered. The final pixel values are computed as weighted averages in accordance with the certainty function. The weighting will provide proper gradients across the final image, as opposed to a final composite image that was constructed with a "copy-and-paste" method of 'properly' exposed pixel values. While Mann and Picard's algorithm is generally no longer used due to its limitations and restrictions, it sparked further development into camera response function recovery algorithms that are a fundamental step of the HDRI process today.

In 1997, Paul Debevec and Jitendra Malik proposed another self-calibrating method to recover a nonlinear camera response function but challenged the parametric form assumed by Mann and Picard [30]. This algorithm is based on the assumption of reciprocity, which states that there is an inverse relationship between the sensor response and total exposure, as determined by the sensor irradiance and shutter speed. This assumption of reciprocity holds for CCD sensors as long as each sensor element measures all of the absorbed photons during the integration time, meaning none of the pixels are oversaturated or overexposed [30].

This reciprocity equation can be simply written as Z = f(X), or as $Z = f(E_e * \Delta t)$, knowing that exposure, X, is the product of irradiance, E_e , and exposure time, Δt . The nonlinear camera response function, f, is safely assumed to be monotonic, or strictly increasing, so the inverse is well-defined [30]. The pixel values, Z, are indexed spatially by i and temporally by j, so the second form of the reciprocity equation can be rewritten as:

$$Z_{ij} = f(E_{e,i} * \Delta t_j)$$

Assuming that f is invertible, the values of interest, E_e , can be determined by:

$$f^{-1}(Z_{ij}) = E_{e,i} * \Delta t_j$$

Taking the natural logarithm of the equation yields:

$$\ln(f^{-1}(Z_{ij})) = \ln(E_{e,i}) + \ln(\Delta t_j)$$

Debevec and Malik then simplify the notation by defining a new function, g, as $g = \ln(f^{-1})$. Again, g is assumed to be monotonic and smooth.

$$g(Z_{i,j}) = \ln(E_{e,i}) + \ln(\Delta t_j)$$

Debevec and Malik note here that g(z) has a finite domain, $[Z_{min} - Z_{max}]$, the minimum and maximum integer values the pixel may take. If N is equal to the number of pixels in each image and P is the number of photographs in the given image set, the problem is to minimize the following quadratic objective function:

$$O = \sum_{i=1}^{N} \sum_{j=1}^{P} \left[g(Z_{ij}) - \ln(E_{e,i}) - \ln(\Delta t_j) \right]^2 + \lambda \sum_{z=Z_{min}+1}^{Z_{max}-1} g''(z)^2$$

The first term of the equation, the double summation, is to find the solution for g and E in a least-squares sense, while the second term is a smoothness term. The scalar λ is a weighting factor that accounts for the amount of noise expected in the digital values, Z [30].

It is noted that the solutions for g(z) and the irradiance values can only be up to some scale factor, α , established by the constraint $g(Z_{mid}) = 0$. The objective function is unchanged by this. The objective function is further modified by introducing a weighting function, w(z), to emphasize smoothness in the middle of the curve. This comes from the idea that the data will fit more poorly at the extreme ends of the curve and better towards the steeper middle portion. The objective function is now written as:

$$O = \sum_{i=1}^{N} \sum_{j=1}^{P} \left[w(Z_{ij}) \left[g(Z_{ij}) - \ln(E_{e,i}) - \ln(\Delta t_j) \right] \right]^2 + \lambda \sum_{z=Z_{min}+1}^{Z_{max}-1} \left[w(z) * g''(z) \right]^2$$

The algorithm is simplified a little further by reducing the number of pixels used in determining the solution. A sufficiently overdetermined system satisfies $N(P-1) > (Z_{max} - Z_{min})$, so that in an example of an 11-photograph (P=11) set with a pixel value range of 255, using 50 pixels would be more than enough [30]. Debevec and Malik mention that the selected pixels should ideally be evenly distributed both spatially in the image and in the pixel value range between Z_{max} and Z_{min} . The pixels should also be taken from areas that have low variance in intensity-values to minimize the effect of optical blurring. However, this process is not automated in their algorithm at the time of publishing.

Similarly to Mann and Picard's algorithm, final pixel values are computed using a weighted average of all available pixels. The weighting is done so that pixel values closer to the middle portion of the response curve are given more weight. Saturated pixels are ignored to reduce the effect of blooming. This weighted average helps reduce noise and imaging artifacts in the final pixel values.

A single response curve is only appropriate for greyscale images. For color photographs, an independent response curve needs to be determined for each RGB channel. Absolute calibration then requires three individual scaling factors for each color channel in order to get both the radiance and color value correct. Debevec and Malik's default algorithm chooses a scaling factor such that a pixel with the value Z_{mid} has unit exposure. This means that any pixel with the RGB values of $(Z_{mid}, Z_{mid}, Z_{mid})$ is achromatic. To calibrate the color channels manually, the RGB scaling factors should be calibrated to a source of known color [30]. It seems that many current HDR software overlook this calibration process and assume the default algorithm is good enough for correct color calibration.

The recommended criteria for determining the minimum number of photographs needed to recover a response function is only 2. However, it is also stated that the images need to have sufficient overlap in the working range, or middle portion, of their respective response curves. Increasing the number of photographs used will also improve noise sensitivity. For creating radiance maps, the number of images required is dependent on the dynamic range of the scene of interest and the camera being used. This value can be found using R/F, where R is the dynamic range of the scene and F is the working range of the camera. Another popular iteration of the camera response algorithm came just two years later from Mitsunaga and Nayar [32]. Their algorithm is very similar to Debevec and Malik's, but eliminates the need for precise exposure values to be known. Instead, rough estimates of exposure ratios, taken from the F-number readings, are sufficient for their algorithm to recover an accurate response function. Treating the exposure ratios as rough estimates provides some forgiveness for any repeatability errors accompanied with aperture size and rounding errors for shutter speeds of the camera [4].

The model used is an N-order polynomial (notation changed to agree with Debevec and Malik's equation given previously):

$$X = g(Z) = \sum_{n=0}^{N} c_n * Z^n$$

This algorithm is a flexible parametric model with a finite number of parameters, which

is what allows it to accurately recover the response function without exact exposure values. An additional pre-processing algorithm also rejects images that have large vignetting effects or temporal changes to reduce the amount of noisy input [32].

To create HDR images, each pixel value is mapped to a scaled, relative radiance value using the computed response function. These relative radiance values are normalized by a single scaled exposure value to give all the pixels the same effective exposure. Final pixel values are computed as weighted averages, but a different weighting approach is taken in this algorithm. Instead of using the certainty functions, Mitsunaga and Nayar base their weighting function on the signal-to-noise ratio (SNR), where values are given more confidence, or weight, when the SNR is maximum, or the relative amount of noise is minimal. They handle the color correction problem by assuming the three separate response curves preserve the chromaticity of the scene [32].

Robertson, et al., [33] suggested another improvement upon Debevec and Malik's method by proposing a probabilistic approach in weighting the input values and accounting for quantization effects. This method assumes a higher confidence in data values taken from higher exposures, and weights those pixels more heavily. This method also includes a noise term that accounts for sources of noise that may include dequantization uncertainty, ADC noise, and dark current.

$$X_{Z_{ij}} = E_{e,i} * \Delta t_j + N_{ij}$$

Since exactly determining the noise terms would be difficult, the noise terms are modeled as independent Gaussian random variables, with variances σ_{ij}^2 . Again, the variances are difficult to characterize, so they are determined by certainty functions. The variances can be converted to weights, $w_{ij} = \frac{1}{\sigma_{ij}^2}$. The certainty function is found by taking the derivative of the response function with respect to a logarithmic exposure axis, and normalized so that the max value is 1.

A joint probability density function is used to determine the HDR image values by the

maximum-likelihood approach, which find solutions that maximize the probability. Alternatively, taking the natural log of the joint probability density function will yield the following objective function:

$$O(E) = \sum_{i,j} w_{ij} * \left(X_{Z_{ij}} - E_{e,i} * \Delta t_j \right)^2$$

This function is minimized and becomes:

$$\hat{E}_{e,i} = \frac{\sum_{j} w_{ij} * \Delta t_j * X_{Z_{ij}}}{\sum_{j} w_{ij} * \Delta t_j^2}$$

The Δt_j term in the numerator is what causes data from longer exposure times to be weighted more heavily. This accounts for quantization effects and reduces noise by averaging across all input data accordingly. It is important to remember that the final recovered pixel values following any of the mentioned algorithms are not true radiometric or photometric quantities, despite it being referred to as radiance or luminance in these papers. These pixel values are only relative, and require further calibration to be considered absolute measurements of radiance or luminance. The most common method for absolute calibration requires taking a true luminance measurement with a dedicated luminance meter of an area captured within the HDR image. The corresponding pixels of the HDR image can then be scaled to match the known luminance measurement, and the rest of the pixel values will also be linearly scaled based on the calibration value.

For any researcher or designer looking to determine the camera response function for his or her own HDR camera system, a step-by-step procedure has been provided in [35]. The process is as follows:

 Use a tripod to mount the camera to ensure image alignment. Tethering the camera to a laptop for remote control is preferred to minimize camera movement and improve image stability.

- (2) The image exposures should be varied by changing the shutter speed instead of the aperture. Maintaining a fixed aperture will keep vignetting effects consistent between each shot.
- (3) The camera's white balance setting needs to remain fixed during the image capture process. While the book specifically suggests "Daylight" white balance, other papers suggest that this is not necessary.
- (4) Additional image color or contrast optimization modes on the camera need to be kept off during the image acquisition process to ensure the response function does not vary between images.
- (5) The image scene to be captured needs to have large grey or white surfaces with continuous gradients for sampling. The entire scene should be neutrally colored to reduce problems with color transforms.
- (6) The scene should have areas that are very bright and very dark. The bracketed image exposures should be separated by 1 EV (a factor of two in exposure time). The darkest image should have no RGB values greater than 200 and the lightest image should have no RGB value below 20, which can be checked with histograms. Using images beyond this range to create the camera response function may be detrimental.
- (7) To perform an absolute calibration, use a luminance meter to take a measurement of a reference point in the image scene, like a grey card.

The camera response functions generated from *hdrgen* and WebHDR can be saved into a text file with the extension .rsp. Three lines of numerical values are provided, each of which describes the red, green, and blue channels, respectively. The first integer value in the line indicates the order of the polynomial, and the following values are the coefficients of the function. An example from WebHDR [6] is provided below, along with its corresponding plot in Figure 2.15. The horizontal, x, axis indicates the logarithm of the sensor signal (exposure) and the vertical, y, axis correlates to the resulting image pixel value.

Figure 2.15: Camera Response Function Example [6]

```
3
1.57501
-1.01875
0.462603
-0.0188579

3
1.54919
-1.01298
0.480414
-0.0166318

3
1.49544
-0.897716
0.414424
-0.0121498
```



2.5.2 Vignetting

Vignetting is the effect of image brightness fall-off towards the periphery of the image. An example of vignetting is provided in Figure 2.16. It is well-documented that this effect increases in severity with increasing aperture size [7] [8]. The technical definition of vignetting specifically refers to the optical phenomenon of light occlusion due to the lens components, as mentioned previously in the Image Formation section. It is caused by the inability of light rays that are off-axis to the camera to completely fill the aperture. Decreasing the aperture size then allows the light to contribute more to the final image. This is demonstrated in Figure 2.17, where the light rays coming in at angle are unable to fill the AA' aperture but completely fills the smaller BB' aperture. A test performed by a team at the University of Illinois indicated that vignetting effects become negligible for aperture sizes smaller than f/4.0 [8].





Figure 2.17: Vignetting [8]



HDR papers tend to use the term vignetting rather loosely, and what most HDR researchers correct for in their so-called vignetting filters also account for luminance loss due to two other effects called the cosine-fourth effect and pupil aberration [36]. The cosine-fourth effect is a result of the Gaussian thick-lens model, in which irradiance off-axis to the camera fall off as a cosine-fourth of the angle between the light ray and the cameras optical axis. Pupil aberration, as discovered and defined by the team at the University of Illinois [8], refers to the variation of light allowed to pass through the lens as a function of angular position, due to nonlinear refraction through the lens components. Although actual vignetting effects may be absent for apertures smaller than f/4.0, there will still be the same pattern of luminance loss due to the cosine-fourth and pupil aberration effects. Distinguishing these three effects from each other is unnecessary during vignetting correction for HDRI, so the remainder of this paper will continue to use the term vignetting as an encompassing term for all three effects.

It is undisputed that vignetting correction is important to consider and apply before utilizing HDR images for luminance measurements. However, there are a variety of techniques used to find the vignetting function and correct for it, and there are advantages and challenges to each method. The simplest method uses a single, large uniformly lit surface that covers the entire FOV of the camera. This method relies on the assumption that the luminance is the same across the entire FOV and that the center point experiences no vignetting. This allows both the camera and the target to stay stationary and only one set of images to be taken for each camera-lens-aperture combination being tested. All points in the image are compared to the center value to determine the magnitude of luminance loss due to vignetting as a function of radial distance from the image center. Some researchers have simply used matte-white surfaces, like sheets of paper, [21] [5] [11], although the preferred method is to use an integrating sphere, especially for ultra-wide angle or fisheye lenses [29] [19].

Since uniform luminance is a very difficult condition to achieve, other research teams
have proposed alternate methods to test for vignetting. These methods use small targets of known luminance, measured by a luminance meter. The most common of these methods utilizes a single stationary target and rotate the camera, typically at 5° increments, so that the target covers a discrete number of points across the center-horizontal of the image frame, to represent the entire FOV [22] [10] [26]. The downsides to this method are that a much smaller set of data points are available for the vignetting analysis, and there is a potential for added error due to camera movement.

Alternatively, one group [7] has decided to maintain a completely stationary setup by constructing an array of targets attached to a semi-circular frame and taking photos of all of the targets at once with a camera positioned in the center of the semi-circle frame. The 49 alternating white and grey target cards were placed at 5° increments, except towards the end and center where the angular resolution increased to 2.5° . The camera was positioned to sit in the center of the semi-circle, while aimed at the center target. This allows for only one set of images to be taken for each aperture size being tested, as opposed to a set for each 5° increment. Still, only a small discrete set of points are available for analysis compared to the uniform-luminance method, and this particularly setup is much more labor-intensive.

Another unique but complicated approach utilized a fixed-in-place constructed light box and a camera mounted to a motorized arm that was capable of tilting and panning [18]. The light box was constructed from a cylindrical container, the inside of which was painted a matte-white. A CFL was the light source mounted inside the cylindrical container. To create well-defined edges for the light source, a cardboard mask with a square cutout covered with translucent paper was attached to the end of the cylinder. The arm was operated by a computer program to maintain high repeatability. The camera took a complete set of bracketed images from 61 different locations such that the light source covered the entire field of view. While this overcomes some of the limitations of the single-target methods by allowing for any number of points to be collected across the entire image frame, the construction and setup are quite labor intensive, and sophisticated robotic equipment may be too costly or inaccessible to other researchers.

Regardless of which specific method is chosen to acquire the necessary data, the analysis methods are similar. Pixel data should be normalized such that the radius from the central pixel to the farthest applicable pixel (either the corner pixels for a full-frame image or to the boundary of a circular fisheye image) is equal to one. The radius can be calculated for each pixel using the x and y coordinates with respect to the center of the image. Absolute luminance calibration is also unnecessary as relative luminance values are sufficient to determine the vignetting effects [18]. The effects of vignetting loss are determined as a function of radial distance from the center of the image, and the magnitude of luminance loss is determined by comparing the luminance values in the HDR image to the true measured luminance or the assumed uniform luminance. Pixels in the center of the image should be assumed to experience no vignetting effect, and all of the pixel data should be normalized accordingly. Once the vignetting loss function is known, the HDR images are corrected by applying the inverse of this function, otherwise known as the vignetting filter [7].

In Jacobs and Wilson's [18] vignetting paper, they chose to create the HDR images in both *hdrgen* and *pfshdrcalibrate* to compare the two software's HDR image generating algorithms. The results showed no statistically significant difference between the images generated by the two different HDR-creating programs. For the vignetting curves, the authors tested 2nd, 4th, and 6th order polynomials to fit the data, but the results showed no significant improvements with the increased order. 2nd-order polynomials were deemed sufficient to model vignetting loss curves for their particular camera. Their results indicated a maximum luminance loss of 31% at the periphery, proving that the vignetting effect can be quite significant and needs to be considered when taking luminance measurements with cameras.

Thus far, the papers that have discussed vignetting treat the vignetting correction function as a unique function that each individual user of an HDRI system needs to define on their own. One paper [7] investigates the similarity of vignetting functions found using two identical Canon 40D camera bodies, and two identical Sigma fisheye lenses, resulting in four camera-lens combinations that are ideally identical. The goal of this experiment is to validate the assumption that a vignetting function determined for a specific cameralens-aperture combination can be used by anyone with the same exact camera-lens-aperture combination, which may benefit the community of HDR researchers.

Cauwerts et al., [7] designed and implemented the semi-circular array of targets described earlier to determine vignetting. For each of the four camera-lens combinations, a full set of bracketed images were taken for each possible aperture size. The HDR images were generated in *hdrgen*, using a camera response function derived for the cameras prior to the vignetting test. The authors also checked radial symmetry by comparing vignetting curves derived in the first and second quadrants of the HDR image. The vignetting curves were assessed for similarity on the basis of the difference between the root mean square errors. Differences of less than 2% were deemed sufficiently similar.

The resulting vignetting curves were created from 6th-order polynomials fitted to the data. The four camera-lens combinations proved to have a difference less than 2% for all of the apertures tested, suggesting that vignetting curves between identical camera body-lens-aperture combinations may be assumed to be the same. For radial symmetry, vignetting curves for the same image but in different quadrants were also deemed similar, with differences being less than 2% for all but apertures smaller than f/16. However, for the apertures smaller than f/16, the magnitude of the differences was still within 4%. The authors also acknowledge that a possible source of error may be imperfect alignment of the targets across the HDR image.

Their results showed a maximum luminance loss of 71% at the largest aperture, f/2.8, but aperture sizes smaller than f/7.1 showed a maximum loss of less than 5%. After applying the vignetting filters to all of the HDR images, their results showed no significant improvements in accuracy for apertures smaller than f/5.6. For larger apertures (f/2.6 f/5.6), the magnitude of error could be reduced from about 30% down to 3%, another clear indication

that applying vignetting filters is necessary for large apertures. Their results agree with Aggarwal's claim that vignetting effects are negligible for smaller aperture sizes, although the exact aperture size threshold is different between the two.

2.5.3 Point Spread Function and Lens Flare

One of most difficult challenges in HDRI currently is finding a reliable method to eliminate lens flare, light scattering, and diffraction. Light scattering effects are inherent to the optical structures of the camera, and while higher quality lenses aim to minimize some of these effects with specially designed coatings, it is impossible to eliminate these effects completely. Diffraction is a result of the finite size of the aperture, and increases in severity with smaller aperture sizes [9] [1] [22]. Ghosting is a type of lens flare that occurs when reflections appear in the rear group of lens elements, creating reflected images of the aperture on the optical image [9]. Another type of lens flare, veiling flare, is a result of reflections that appear in the front lens elements, resulting in a haze that covers the image. Starburst PSF artifacts are determined to be a result of triangular apertures.

Light scattering and diffraction is often described by a point spread function (PSF), which describes the radial effect of light information of a single pixel spilling or scattering into neighboring pixels. This results in loss of luminance information of the central pixel and an increase of false luminance data in the neighboring pixels. Many of the aforementioned papers have acknowledged the errors associated with these light scattering effects, but few have established a method to correct for them. Some have determined that lens flare and diffraction effects improve with larger aperture sizes, although increasing aperture size also comes at the cost of increasing vignetting effects. These authors suggest finding a balance between these two effects by selecting mid-range aperture sizes.



Figure 2.18: Examples of Starburst PSF, Ghosting, and Haze [9]

Inanici [10] attempted to quantify the PSF by taking an image of a point light source so that it covered only one pixel of the digital image. In image b) of Figure 2.19, the outlined box represents the size of the original point source. The amount of illuminated pixels surrounding that pixel clearly illustrates the effect of PSF, which is more pronounced in this image with a single light source than in a typical image. PSF is shown to be dependent on aperture size, exposure time, and distance from the optical center of the image (eccentricity). The eccentricity effects are shown in image d) of Figure 2.19, with the image to the far left being perfectly centered and growing with radial distance towards the right. The PSF effect is also more prominent when there is high contrast between the background and target, specifically a dark target against a bright background. This results in overestimation of dark targets, which is evident in past experimental results [10] [11]. Although Inanici was able to observe the effects of PSF, it was deemed infeasible to quantify this effect in a general way to correct for light scattering across the entire image.



Figure 2.19: PSF as a Function of Eccentricity [10]

Rea [19] and Anaokar and Moeck [11] took a different approach by attempting to define the Modulation Transfer Function (MTF), which is the Fourier transform of the PSF. The MTF is used to describe the spatial resolution of the image at which the PSF effects are not so significant to experience a complete loss in determining boundaries between distinct objects in the image. Both papers used a target of alternating black and white bars of varying widths, like the one included in Figure 2.20.



Rea [19] calculated the minimum spatial resolution from contrast ratios between the white and black areas of the target. The results determined that objects must occupy at least 2% of the image to achieve sufficient spatial resolution. Moeck and Anaokar [11] found the minimum spatial resolution by measuring the pixel width of the transition between the black and white bars. The minimum spatial resolution was found to be 4 pixels. At their specified target distance, this equated to a physical distance of 0.3241 cm. Again, this method does not provide any way to correct the final images for PSF, but attempts to define a minimum spatial resolution threshold at which objects can be reliably measured. This method appears to be limited to non-light-emitting surfaces that are uniformly lit.

2.6 Near Field Illuminance Calculations

Near-field photometry has been a highly debated topic in the lighting industry for several years. The concept of near-field photometry was introduced as early as 1987, and since then, there have been several papers that have made the argument for the need and applicability of near-field photometry. The key concept of near-field photometry is that the luminous intensity distribution of the luminaire is no longer independent of distance, and the specific viewpoint of the luminaire becomes significant in determining the contribution of light onto the point of interest. There are a number of lighting applications in which the availability of near-field photometric data would be beneficial, including wall-washing, cove lighting, task lighting, and systems that utilize wall-mounted or indirect fixtures. Despite these arguments, there has yet to be adoption of near-field photometry practices in the industry, nearly 30 years later. Reasons for resistance to adopt near-field photometry procedures likely include the need for additional, expensive testing equipment, and the lack of a standard format in which to provide such data in a concise way.

To understand the benefits of near-field photometry, the basics of far-field photometry and the conventional method of photometric calculations need to be understood as well. The term far-field is used to describe conditions in which a light source is far enough away to be considered a point-source, and so the luminous intensity is independent of distance to the source. A simplified rule of thumb to determine an appropriate distance at which this condition applies is commonly referred to as the "five-times rule," which states that the distance must be at least five-times greater than the largest dimension of the luminaire [37] [38] [39]. For example, if the luminaire is a standard 2'x4' fluorescent troffer, a distance of 20' is considered sufficient for far-field photometry to apply. However, an important caveat to this rule is that the light source must be Lambertian in distribution [37]. In fact, for non-Lambertian sources, this distance increases to a least 15 times the largest source dimension [40] [41]. In the cases in which the calculation plane is within this "five-times" distance, the luminaire is simply broken up into smaller pieces, the size determined by the conditions sufficient to meet the five-times rule. The major assumption made in conventional calculation procedures based on far-field photometry is that each small piece possess the same luminous intensity distribution, based at its new photometric center, except that the magnitude of intensities are scaled down accordingly [42]. This simplification gives rise to many errors in near-field calculations, especially when the luminaire is inhomogeneous. This has driven some researchers to propose alternate photometric procedures to account for near-field conditions.

To demonstrate the magnitude of error caused by the far-field and homogeneity assumptions, and to provide evidence that near-field photometry is necessary, Mistrick and English performed tests with an indirect/direct 4'x1' linear fluorescent fixture and an indirect metal halide fixture in typical indirect lighting applications. This was accomplished by testing the luminaires in a dark room, to minimize the contribution of interreflected light, and taking illuminance measurements at 6" increments along a designated 6' by 6' grid pattern [37]. These near-field illuminance measurements were then simulated using the conventional calculation method, based on standard far-field photometric data for each luminaire and breaking up the luminaire into smaller homogenous pieces.

For the fluorescent luminaire, the comparison between the conventionally calculated illuminance values and the true measurements showed a maximum error of 56% along the lamp axis, and 33% perpendicular to the lamp axis. The results for the metal halide fixture showed smaller discrepancies, but the maximum error still reached 32% perpendicular to the lamp axis [37]. The magnitudes of these errors help demonstrate the need for an alternate method, but the authors did not utilize their own near-field calculation procedures to demonstrate improvements. The extent of this test is also very limited, and the authors acknowledge that other tests need to be performed to demonstrate the applicability of such concepts to applications besides indirect lighting.

2.6.1 Proposed Near-Field Photometry Methods

One of the earliest introductions of the concept of near-field photometry was published in 1987 by P.Y. Ngai. This method involves the use of luminance distribution functions, which describe the luminance at any location p on the luminaire in direction r as a function L(p,r). From this function, many other photometric quantities may be derived, such as the intensity distribution or illuminances on a surface due to the luminaire. It also accounts for viewpoint dependence of luminaire luminance and can describe unique, inhomogeneous light patterns across the luminaire [42]. Since it is often more convenient to deal with intensity values rather than luminance values, an equation to derive the intensity distribution from the luminance distribution function is provided:

$$dI_{dA}(p,r) = L(p,r) * dA_p$$

where dA is the differential area of the luminaire at which the intensity is emitting from, and dA_p is the projected area of this patch in the direction of r [42]. The luminaire is thought of as a collection of pieces, dA, that each possess their own unique luminous intensity distributions. Ngai also provides equations for calculating surface illuminance and luminance using the luminance distribution function.

In order to determine the luminance distribution function of a luminaire, the intensity distribution function needs to be obtained empirically. This can be done by taking illuminance measurements on a plane near the luminaire. The intensity distribution function will depend on the physical location of the luminaire point of interest (x, y, z) and the direction from which it is being viewed (ψ, θ) . The author chooses to simplify the problem here by considering the two sets of variables separately. This means the luminaire will be considered as a collection of luminaire pieces, each described by its own photometric center (x, y, z), to have its own unique intensity distribution, $\Delta I_{dA}(x, y, z)(\psi, \theta)$. Taking advantage of luminaire symmetry can reduce the number of intensity distribution functions to be determined, and focuses on obtaining intensity distribution functions where they are most significantly different.

Solving for the intensity distribution function becomes a system of equations that can be shown in the matrix forms below:

$$[E_s] = \left[\frac{\cos\Phi}{d^2}\right][I]$$
$$[I] = \left[\frac{\cos\Phi}{d^2}\right]^{-1}[E_s]$$

Where E_s is the vector of illuminance measurements, d is the distance between the point on the luminaire and the point at which the measurement is being taken, and Φ is the angle between the direction r and the surface normal of the measurement point. These equations require choosing a discrete number of pairs for (ψ, θ) , which should be chosen carefully.

Once the intensity distribution functions are determined, the luminance distribution function is found using:

$$L(p,r) = \frac{dI_{dA}(p,r)}{dA_p}$$

In response to the challenge of near-field photometry adoption, the author believes there are three fundamental steps that need to be accomplished. First, the theory and mathematical concepts required to obtain and use near-field photometry need to be understood. This paper makes an attempt to accomplish this task. The second step is to develop a data gathering technique that can collect large amounts of necessary data points in a convenient, practical, and quick manner. Lastly, the mathematics to reconstruct usable luminance distribution functions from the data needs to be developed [42].

Another near-field photometry proposal, named "Application Distance Photometry," was proposed later in 1990. This method is specifically designed to deal with suspended indirect fixtures. The process to obtain the necessary near-field measurements begins with mounting an indirect luminaire close to a ceiling plane. Illuminance measurements are then taken at a number of specific points on the ceiling to derive luminous intensity values. Application distance photometric reports are provided for a given suspension length, or distance from the ceiling. This data can be formatted into a standard .ies file for designers to use in lighting modeling software. This process does not require discretization of the luminaire since the data already accounts for the luminaires close proximity, and so the luminaire is still treated as a point source [38].

The authors go on to describe the construction of a rig capable of taking application distance photometric data. The general concept is based on pre-existing designs of goniophotometers, with the added capability of a photocell track, to measure illuminance points at various points in the presumed ceiling plane. This method was then demonstrated with a 6" by 4' linear fluorescent fixture, suspended 12" from the ceiling plane [38]. The results show increased accuracy of near-field illuminance calculations using their proposed method in contrast to conventional methods.

While this paper is important in the history of developing the foundation for near-field photometry, this exact method appears to have a number of shortcomings. For one, this method is developed specifically for suspended indirect luminaires and is not universal for all near-field lighting applications. The authors do explain the potential to extend this method for wall-mounted fixtures, but with some added complexities. The photometric files obtained are also specific to suspension length, which may be very limiting. Application distance photometry also does not account for interreflections or shadowing form the luminaire [38].

The challenge with the previous near-field photometry methods is that they require new photometric equipment or new calculation procedures. In 1995, DiLaura and Chu proposed a simplified but robust method that could simulate near-field illuminances without the need for new equipment or calculation procedures. The primary goal of this method was to augment existing conventional far-field photometry files with as little information as possible that would allow for reasonable near-field illuminance calculations [12]. The method begins by considering the luminaire as a collection of individual elements, each with its own unique luminous intensity distribution and photometric center. This assumption of uniqueness is important as it challenges the conventional assumption of distribution homogeneity that is used in conventional illuminance calculations. However, luminaire symmetry can be used to simplify and reduce the number of unique intensity distributions needed to characterize the luminaire. The entire collection of unique luminaire pieces should yield the same results as conventional far-field photometry would, at sufficiently far distances. Within near-field distances, these unique distributions will account for inhomogenous patterns caused by lamp presence, beam patterns, and shadowing by fixture components, like baffles [12].

The individual luminous intensity distributions, otherwise known as candela distributions, are found using the simplex linear optimization method and a small number of luminance scans taken for each photometric center. This optimization is done with linear constraints that are based on realistic photometric characteristics, such as how individual distributions interact with neighboring ones and how they behave as a whole. The detailed list of constraints used is given in [12]. Only three luminance scans were taken at each photometric center, at angles of 0°, 25°, and 45°. Intensity values can be derived from these luminance measurements.

The resulting solutions come in the form of a collection of intensity values for each photometric center. The solutions are then scaled so that the total emitted lumens is equal to that as determined by conventional far-field photometric methods.

This method was validated by testing two different quadrilaterally-symmetric linear fluorescent troffers. The luminance scans were taken with a Minolta LS-100 luminance meter, attached to a special rig to maintain accurate and precise aiming. The near-field illuminance calculations derived from the proposed method and the conventional far-field method were compared to illuminance measurements taken with a Minolta T-1M illuminance meter for verification. The results showed that the calculations derived from the proposed method matched more closely to the true measurements than the conventional calculations, as can be seen in Figure 2.21 [12]. The plot labeled "lumen-micro" are the results for the conventional far-field methods and "3-scan-simplex" is the proposed method.

Figure 2.21: Results of Near-Field Illuminance Test [12]



ILLUMINANCE (h=18 inches)

A strength of this near-field calculation method is that it is not sensitive to the accuracy of the luminance scans. Luminaire luminance measurements are very difficult to obtain accurately, so the robustness of this method is beneficial. The authors were able to demonstrate this feature by adding 10% random noise to the luminance scan data to obtain satisfactory results [12]. Some suggestions for improvement that were not yet included in this paper, included the need for more luminance scans and treating the lamp ends as a separate luminaire piece unique from the other pieces that lie along the lamp. However, these improvements would come at the cost of added complexity.

While this technique does not create full near-field photometric characterizations of luminaires that can be provided as additional data in the .ies file for lighting designers to use, it does provide an intermediate step. It should also help demonstrate the need for improved methods to deal with near-field lighting situations, which are currently overlooked in many lighting design calculations and industry practice.

2.6.2 Near-Field Goniometer

A recent paper [40] on near-field photometry discusses the use of a new near-field goniometer with narrow-beam LED arrays. The near-field goniometer consists of an illuminance meter and a luminance camera that revolve around the light source. The near-field goniometer creates ray files based on the luminance maps created by the camera and the absolute luminous flux data provided by the illuminance meter. Ray files are a list of rays that are each described by a starting point on the source surface, direction of propagation, and luminous flux. This approach allows for illuminance calculations that are independent of source dimensions and beam angle. The authors go on to derive an expression starting from luminance, instead of luminous intensity.

The disadvantage of this method is that ray-tracing methods will produce noisy results at large distances as a result of fewer rays incident on the measurement area. This error is proportional to the square root of the number of incident rays, and will increase with distance. Therefore, ray-tracing methods are best reserved for near-field applications.

The proposed luminance-based near field photometry methods are then tested with an LED array in two configurations; all 5 LEDs on, and only the outer 2 on. Their results showed that far-field conditions were met at a distance far greater than the "15-times" distance as suggested. For the second configuration, the distance to achieve far-field conditions is even greater. The authors believe that these results clearly demonstrate the need for appropriate near-field photometry techniques to define narrow-beam optics, including LED arrays and car headlights. Near-field photometry can also provide better descriptions of light distribution patterns when the luminaire is inhomogeneous.

2.7 Current Challenges and Research Gaps

After study of all of the previously mentioned papers relating to HDRI and near field photometry, it is clear that there are still many research gaps to be covered. Such topics include the performance differences between CCD and CMOS cameras, corrections for saturated color errors in luminance readings, corrections for lens flare and light diffraction effects, and a lack of discussion of the software alternatives available for HDRI processing and analysis. This document will not attempt to cover all of these topics, but they will be provided here to expose the topics that are important to this field of research.

One topic that this document will attempt to cover is the lacking discussion of HDRrelated software alternatives that are currently available. From the point of view of a new researcher unfamiliar with the available tools for HDRI processing and analysis, it can be difficult to make a decision as to which software is best to invest in. A strong majority of the papers mentioned here rely solely on a program named Photosphere, or *hdrgen*, to create the HDR images. However, there are many additional software options that have HDRI capabilities. There is little to no discussion provided in any of the reviewed research papers that justify the use of Photosphere or *hdrgen* over any other available software. This document will attempt to provide a comprehensive evaluation of the capabilities and performance differences between the available software. The hope is that this document will help educate HDRI researchers choose the right software for their purposes. It will also act as a documentation and justification for the software selection made for future research projects conducted here, based on which software options best fit the needs for scientific lighting analysis.

The following topics will not be addressed in the tests and experiments described in this document, but are included here as a list of potential topics for future research projects.

Currently, there appears to be little discussion in the HDRI research community on the performance differences between CCD and CMOS sensors. The only HDRI paper that addresses the CCD and CMOS sensor dilemma is Axel Jacob's 2007 paper [4] that showed a preference for CMOS sensors. Past papers are about equally divided in the use of CCD or CMOS sensors, but no paper provides a justification for their camera or image sensor type selection. Both sensor types have their own unique benefits in regards to HDRI, so it is difficult to make a decision based on these features alone. A research study seems necessary on the performance between CCD and CMOS-based cameras to help HDRI researchers make more informed decisions for purchasing the most suitable camera for their HDRI systems.

Saturated color issues have also yet to be solved. There are already several papers that have uncovered this issue and proven that HDR luminance measurement errors are significant for saturated color targets, especially for blue and purple hues. Although most HDRI applications do not involve the study of saturated colored surfaces, it would improve the reliability of HDRI luminance maps if a solution to mitigate such errors could be found.

As previously mentioned, lens flare and light diffraction effects are also in need of further investigation. It is seemingly impossible to compose an algorithm that can reliably remove lens flare for improved HDR luminance measurements in a robust way, but perhaps it would be more meaningful to understand the magnitude of error associated with these effects.

In addition to a lack of comprehensive discussion of available software alternatives, there also appears to be gaps in the capabilities of the available software. Regarding HDR image analysis, most researchers appear to have created their own in-house MATLAB scripts instead of using publicly available programs. This appears to demonstrate the need for new or improved versions of current software that are capable of handling the necessary analysis most HDRI researchers depend on for meaningful results. The HDRI community would likely make faster progress with publicly available software that can cover all of the needs of HDRI creation and analysis, preferably in a single package.

Near-field photometry, although an already well-established research topic, is still in need of further improvements. It may be possible that HDRI is a tool that could help bridge the gap between near-field photometry and conventional practice. A method that utilizes standard DSLR cameras to acquire luminance information to create near-field photometry effects is worth investigating. HDRI methods may not be able to create true near-field photometry data like a near-field goniometer would, but it may provide an intermediate step for applications where investing in a dedicated near-field goniometer is out of scope or budget.

Chapter 3

Development of New HDRI System

3.1 Equipment

3.1.1 Camera Body

The camera used in the following experiments and tests is the Nikon D5200, which uses a CMOS image sensor. A quick summary of the settings maintained for all tests is provided below, along with explanations as to what each setting controls.

- Active D-Lighting: This camera has a feature called "Active D-Lighting" that is designed to optimize image quality. Active D-Lighting attempts to balance image contrast levels by applying additional processing to certain areas of the image, creating effects similar to HDR imaging [43]. By default this feature is on, so it was disabled for the images taken for the HDRI tests to minimize the number of variables affecting image exposure and quality.
- **ISO:** Camera ISO describes the level of sensitivity of the camera image sensor to light, where sensitivity increases with the value of the ISO number [44]. While higher light sensitivity may sound desirable, it comes at the cost of added image noise in the form of image grain. Therefore, for most photography applications, it is highly recommended to shoot using the lowest ISO setting possible. High ISO settings are reserved for low-light situations or action shots. Since the following experiments do

not deal with fast-moving or low light scenes, the ISO setting was kept at ISO 100 for all of the images taken.

Metering Mode: The camera metering mode is not important for HDRI, so the setting was left in the default Matrix metering mode. Metering mode is simply a tool to indicate proper exposures, but this is not necessary since multiple images at multiple exposures are going to be taken without any preference for which is the best for the scene [45].

3.1.2 Camera Lenses

Three different lenses were used in conjunction with the D5200 camera body in the following tests. The three lenses included a standard Nikkor 18-55mm zoom lens, a Sigma 4.5mm, F2.8 EX circular fisheye, and a Sigma 24mm, F1.8 prime lens. The fisheye was purchased for architectural applications, and the prime lens was purchased to determine if there is any difference in image quality between a basic zoom lens and a fixed focal length lens. The EX designation with the fisheye lens indicates that this particular lens is part of Sigma's professional-grade line of products, designed for superior construction and optical qualities. Additional details and technical specifications of the Sigma lenses may be found on the manufacturer's website [46].

3.1.3 Other Equipment

Additional equipment used in the HDRI testing included a Konica-Minolta LS-110 luminance meter, calibration or reference targets, and tripods. The luminance meter was calibrated on August 25, 2015 to ensure that the absolute luminance measurements taken are reliable and accurate. The LS-110 model, pictured in Figure 3.1, has an acceptance angle of $1/3^{\circ}$ and a minimum measuring distance of 40" [47]. The measurable luminance range is reported to be 0.01 Cd/m^2 to 999,900 Cd/m^2 . The measurement area of the luminance meter has a defined finite size, and the object of interest must fill the entire measurement area to provide useful readings, with the minimum measurement area defined by a diameter of 4.8 mm at a distance of 1014 mm, or 40". Minimum measurement area diameters for different measurement distances are provided in Table 3.1. The spectral response of the LS-110 is designed to match the spectral response of human vision as closely as possible, with the achieved response curve matching within 8% of the CIE standard. A plot of the two response curves is provided in Figure 3.2 [48]. It can be seen here that there are are larger discrepancies in the lower half of the spectrum, which may affect results. The accuracy of the meter is reported to be within 2% +/- 2 digits of the displayed value for measurements under 10 Cd/m^2 and within 2% +/- 1 digit for measurements of 10 Cd/m^2 or higher.

Figure 3.1: Konica-Minolta LS-110 Luminance Meter



Table 3.1: Minimum Measurement Area for LS-110

| Measuring Distance [in] | 40 | 48 | 60 | 72 | 96 | 120 |
|--------------------------------|--------|--------|--------|--------|--------|--------|
| Measurement Area Diameter [in] | 0.1890 | 0.2272 | 0.2840 | 0.3408 | 0.4544 | 0.5680 |



Figure 3.2: Konica-Minolta LS-110 Luminance Meter Response Curve

Six X-rite color checkers, three classic color charts and three greyscale charts, were used for reference targets in the tests. Examples of the color checker cards are shown in Figure 3.3. Basic tripods were also used to stabilize the camera and luminance meter during these tests.



Figure 3.3: X-Rite Color Checker Classic and Greyscale Cards

3.2 Software

The process of creating and utilizing HDR images for the purpose of lighting analysis can be thought of in three phases; image acquisition, HDR image creation, and HDR image analysis. For each phase, there are several available software options designed to make the process easier. For image acquisition, tethering software allows the user to remotely control a camera from a computer or laptop. This is typically quicker and more efficient than manual camera operation, and minimizes camera movement between images, which is important for image alignment during the HDR image creation step. For HDR image creation, the previously acquired LDR images will be used to produce a single HDR image based on one of the camera response function algorithms explained earlier. The specific algorithm used will vary between each software package. Lastly, the composed HDR images need to be viewed and analyzed. There have been a few software options designed to view HDR images specifically, allowing for luminance value readouts and other image analysis options important for lighting evaluations.

The software options that were tested and used in the following experiments are briefly explained here, but this list should not be considered an exhaustive list of all possible HDR software options. An in-depth review of the performance capabilities and applicability of HDR creation software will be included in a later section.

3.2.1 Tethering

Tethering software is not necessarily a requirement, but it greatly simplifies the data collection process by automating the image bracketing sequence and minimizes camera movement to reduce errors associated with image misalignment. The list of software studied in this report includes *SmartShooter*, *digiCamControl*, and *Sofortbild*. Again, this is not an exhaustive list of tethering software, but a small number of options that fit the needs of the following experiments and were compatible with the Nikon D5200. Smart Shooter is available for both Mac and Windows operating systems. Unfortunately, this is not available by free download. Smart Shooter is compatible with several Canon EOS and Nikon DSLR cameras. The full list of supported cameras can be found on the developer's website [49]. Smart Shooter comes with a number of pre-written scripts for applications like HDR, timelapse, and bracketing. This was initially the software of choice until there was a problem encountered with the provided HDR script which would automatically bracket exposures. It was discovered that with version 3.4, the auto-bracketing feature did not work well with repeated runs. As a result of this issue, other tethering software options were sought out. However, with update v3.17, the HDR script had been corrected, and the script worked well with no errors. Another feature of Smart Shooter that may prove useful to other HDRI researchers is the ability to load in custom scripts written in Tcl (Tool Command Language) for tethered camera control, but this feature was not utilized in these experiments.

Another Windows compatible tethering software is *digiCamControl* [50]. This software is free to download, and based on donations. This software is also compatible with both Canon and Nikon, and the full list of supported cameras is included on the website. The user interface allows the user to look through the images taken, complete with RGB histograms. *digiCamControl* also comes equipped with auto-bracketing features, with four different modes. The option used in these experiments was the manual exposure bracketing, which allows the user to specifically choose which shutter speeds to include. Again, another small problem was experienced with this auto-bracket feature with repeated and skipped exposures in version 1.2.0.0. It appears that the software changes the shutter speeds too quickly for the camera to respond, which may be easily fixed with a time delay between shutter speeds. With update 2.0.0.0 the shutter speed selection feature with the exposure bracketing has been removed.

Sofortbild [51], version 1.3, was downloaded as the Mac alternative to digiCamControl. This is a Mac exclusive app that works only with Nikon cameras. This app is also free to download. The interface is simple, with its options given in a bottom and top ribbon. It does not appear that the app allows the user to scroll through images previously taken, displaying only the most recent image. This too has an auto-bracket feature that works by asking the user for the slowest and fastest shutter speeds, the step increment, and time delay between each shutter speed. So far, this auto-bracket feature appears to work well with the Nikon D5200. No problems were encountered with the shutter speeds not incrementing correctly and it was easy to run the feature repetitively, therefore this was the most frequently used tethering software in the following experiments.

3.2.2 HDR Image Creation

There is a number of available software for the purpose of producing HDR images. WebHDR [6] provides a list of free software that includes *hdrgen*, *raw2hdr*, LuminanceHDR, Bracket, and Picturenaut. A brief overview of the capabilities and features of each software tested and used during the following experiments are summarized below. Table 3.2 provides a quick snapshot view of the compatibilities of each of the mentioned software options. Table 3.3 shows the accepted input and output file types of each software.

| | | Camera Response Algorithm | | | OS Compatibility | | | |
|--------------|---------|---------------------------|-----------|----------------------|------------------|--------|-------|--|
| | | | | | | | | |
| Software | Version | Debevec & Malik | Robertson | Mitsunaga & Nayar | Windows | MacOSX | Linux | |
| hdrgen | | Х | | X | | Х | Х | |
| raw2hdr | | | N/A | | | Х | Х | |
| Bracket | v1.0.0 | | Unknown | | Х | Х | Х | |
| LuminanceHDR | 2.4.0 | Х | Х | | Х | Х | Х | |
| Picturenaut | 3.2 | | | Х | Х | | | |

 Table 3.2: HDR Image Creation Software Compatibilities

| | Accep | ted Inp | ut File | | | | | | |
|--------------|-------|---------|---------|-----------------------------|---------------|---------------|------------------|------------------|------------|
| | Types | | | Generated Output File Types | | | | | |
| | | | | | | | Tiff- | Tiff- | |
| | | | | Radiance | JPEG HDB | Open EVD | HDR, | HDR, | PFS |
| Software | JPEG | TIFF | RAW | (.hdr) | HDR (.jpg) | LAK (.exr) | LogLuv (.tif) | 52-611 (.tif) | (.pis) |
| hdrgen | Х | Х | | Х | Х | Х | Х | | |
| raw2hdr | | | Х | | | | | | |
| Bracket | Х | Х | | Х | | Х | Х | Х | |
| LuminanceHDR | Х | Х | Х | Х | | Х | Х | Х | Х |
| Picturenaut | Х | Х | Х | Х | | Х | Х | Х | |

Table 3.3: HDR Image Creation Software Accepted Input and Output Files Types

Easily, the most popular option for HDR image creation is *hdrgen*, created by Greg Ward of Lawrence Berkley National Laboratories [52]. *hdrgen* has no user interface and must be run from the command window. Popular HDR programs like Photosphere and WebHDR both provide user interfaces with the *hdrgen* engine to create HDR images. A limiting aspect of *hdrgen* is that it is exclusive to Mac operating systems. Therefore, any researcher interested in using *hdrgen* must own a Mac computer or go through WebHDR to create HDR images. *hdrgen* may be downloaded for free from the Anyhere Software website [52]. A key feature of *hdrgen* is the ability to write and read in camera response functions in simple text files with the file extension .rsp. Other features include automatic exposure alignment, exposure adjustment, over- and under-exposure image removal, lens flare removal, and ghost removal. A full list of the possible commands and how to use them are included in a text file with the program download. The HDR image creation algorithm is based on Debevec and Malik's [30] and Mitsunaga and Nayar's [32] published algorithms.

raw2hdr is also created by Greg Ward and has the same functionality as hdrgen, except that it has the capability to accept RAW image files as input for creating HDR images. The same Mac operating system limitation holds for raw2hdr. The same commands and features available in hdrgen are also available for raw2hdr, with the exception of writing and reading in camera response function files. This is an unnecessary feature for RAW images. Bracket [53] is capable of creating HDR images from JPEG image files only. When generating HDR images, it provides options for using a recovered response curve or generating a new curve. It is unclear how the response curves are derived. Bracket can read out luminance and RGB values on a per pixel basis, also giving the X and Y coordinates within the image. There is also a feature for absolute luminance calibration based on a single pixel and histograms are provided with options for luminance or RGB values. Basic tonemapping can be achieved by altering the gamma level in the images.

LuminanceHDR [54] is another software package designed to create HDR images. When creating HDR images, the program prompts the user with six default profiles to choose from. These profiles define the weighting function and response curve shape that the HDR image creation will be based on. The three weighting function options are labeled as triangular, plateau, or guassian, and the two possible response curve options are linear and gamma. Alternatively, the user can specify customized options. An option is listed for loading in a response curve from a file, but this feature does not appear to work with version 2.4.0. There is also an option to use the Debevec and Malik or Robertson, et al. models when creating the HDR images. LuminanceHDR has many tonemapping features, but is not capable of providing luminance readouts or performing luminance-based calibration.

A program named Picturenaut [55] is yet another option for HDR image creation. Camera response curves can be calculated using the default option, or the user can load an already defined curve from a .crv file. There are some tonemapping options, but Picturenaut does not read out luminance values or allow for luminance-based calibration.

WebHDR [6] is not a software program, but rather a web-based interface that allows anyone with internet access to upload a set of images for HDR image creation. WebHDR was not used for any of the following experiments, but is included here as an alternative tool and excellent resource for beginners in HDRI. As previously mentioned, the engine behind the HDR image creation is *hdrgen*, so WebHDR may be a sufficient workaround for those who do not own Mac computers but would like to use the *hdrgen* software. However, as a web-based program, there are a few limitations. Perhaps the most obvious is the 12MB limit for the upload size. It is also restricted to JPEG images, so RAW images are not accepted at this time. The same features of *hdrgen*, including camera response function writing and reading, are available in WebHDR.

3.2.3 HDR Image Analysis

Lastly, for HDR image viewing and analysis, Photosphere and hdrscope were used. Most of the previously mentioned HDR creation software can be used for viewing HDR images as well, although some of them are not designed to produce information useful specifically for lighting analysis.

Photosphere is essentially a user interface for hdrgen and can also be downloaded from the Anyhere Software website [52]. Since hdrgen and raw2hdr do not have image viewing capabilities, Photosphere should be installed in conjunction with these programs. The user may choose to create HDR images through the Photosphere interface instead of through the command terminal. Through Photosphere, HDR images can be opened for viewing and basic luminance analysis. The user may click the image for luminance readings of the pixel, or select a rectangular region for the mean luminance value across the selected pixels. Photosphere also allows the user to calibrate the image by providing an input luminance value for a selected region of pixels. Photosphere also has false-color mapping capabilities and a histogram feature. The histogram feature can provide luminance data or RGB integer values, and will display the minimum, median, maximum, mean, and standard deviation of the selected pixels.

hdrscope is a program created by a student, Kumaragurubaran [56], from the University of Washington in conjunction with Inanici, an established author of other HDRI research papers. The intended purpose of *hdrscope* was to create an easy-to-use, user-friendly interface for HDR image analysis by combining the capabilities of Photosphere and Radiance and private user-create MATLAB scripts into one package. Since utilizing HDR analysis tools

like Radiance or personalized MATLAB scripts require knowledge of command-line programming, there is sometimes a steep learning curve that prevents entry-level lighting researchers from being able to properly analyze their results. The capabilities of *hdrscope* include perpixel and regional statistical luminance analysis, glare analysis, false-color mapping, tonemapping, automated and remote camera control, vignetting correction, and luminance- or illuminance-based calibration. Currently, *hdrscope* runs only on Windows operating systems and requires Radiance to be pre-installed on the computer. It can be downloaded for free from its website [57].

Vignetting filters can be applied through *hdrscope* for vignetting correction, but it is currently not capable of determining the vignetting function. Basic luminance- or illuminancebased absolute calibrations can be performed just like in Photosphere by selecting a region of pixels and setting them to be equal to the desired input luminance or illuminance value.

Like Photosphere, *hdrscope* provides false-color mapping and luminance statistical analysis outputs. A unique feature of *hdrscope's* luminance analysis tool is the ability to select circular and custom polygon areas in addition to basic rectangular shapes. This provides significant flexibility in the statistical analysis of luminance distributions across any shaped surface in the HDR image. In the selected region, *hdrscope* will also provide basic statistical data like the minimum, maximum, mean, median, standard deviation, and a frequency plot, or histogram, of all of the selected pixels' luminance values. In addition to the basic statistical summary, there are options to calculate percentile ratios or percentage of pixels that meet specified luminance criteria. These results can then be exported into a comma separated values, .csv, file for later use in programs like Excel or MATLAB.

hdrscope is also equipped with glare evaluation, tone-mapping, and camera tethering features, but none of these were utilized in the following tests. *hdrscope's* glare analysis tool uses Radiance's Evalglare program, which is capable of determining daylight glare probability (DGP), daylight glare index (DGI), unified glare rating (UGR), visual comfort probability (VCP), and CIE glare index (CGI) ratings. Two of the most commonly applied tonemapping operators are included in *hdrscope*, but there is also a batch processing feature that allows users to apply their own set of Radiance commands as a saved macro file. There are basic editing features like cropping, rotating, and resizing. The camera tethering feature is equivalent to HDRcapOSX and is compatible with Canon cameras produced after 2007.

3.3 Determining Camera Response Functions

3.3.1 Methodology

The first step to establishing the HDR system is to determine the camera response functions (CRFs) for the camera to be used. The primary camera in this new HDRI system will be the Nikon D5200. The Nikon camera body was tested with three different lenses, a standard zoom lens (Nikkor AF-S DX Zoom 18-55mm f/3.5), a prime lens with a fixed focal length of 24 mm (Sigma 24mm f/1.8 DG), and a circular fisheye lens (Sigma 4.5mm f/2.8 EX DC HSM Circular Fisheye). With the zoom lens, 18 mm and 24 mm focal lengths were tested.

The camera response functions describe the mapping between the image sensor response and the pixel value of the final image. This process is applicable only to compressed image formats, like JPEGs, so finding a camera response function is not necessary when using RAW image files. This is because RAW image files contain raw signal data that does not undergo the same nonlinear mapping that JPEG images do. Camera response functions are a property of the specific camera and are unique among camera manufacturers, camera models, and even with each individual camera. This response function should be independent of the camera lens, aperture size, and focal length. This initial test will verify that this assumption is indeed true so that a reliable camera response function for the camera can be used for future HDRI tests and applications.

To ensure that a reliable camera response function is generated, the photographed scene should be chosen carefully. As explained previously in the Literature Search, it is recommended [35] that the scene contains large neutrally colored surfaces that provide continuous gradients for sampling. The less colorful the scene is, the less likely problems with color transformations will occur. It is also important to include areas that are very dark and very bright to capture a wide range of luminance values. However, the darkest exposure should have no RGB values greater than 200 and the lightest exposure should have no RGB values lower than 20. Using a luminance meter to take absolute luminance measurements of points in the photographed scene will allow for absolute calibration of the HDR images later. An HDR image of the scene used for determining the camera response functions is provided in Figure 3.4.



Figure 3.4: Camera Response Function Scene

| CAMERA BODY | LENS | APERTURE | WHITE BALANCE |
|-------------|--------------|----------|--------------------|
| NIKON D5200 | Fisheye | f/3.5 | Daylight |
| | Fisheye | f/3.5 | Fluorescent - Cool |
| | Fisheye | f/5.6 | Daylight |
| | Fisheye | f/5.6 | Fluorescent - Cool |
| | Fisheye | f/9 | Daylight |
| | Fisheye | f/9 | Fluorescent - Cool |
| | Fisheye | f/14 | Daylight |
| | Fisheye | f/14 | Fluorescent - Cool |
| | Fisheye | f/22 | Daylight |
| | Fisheye | f/22 | Fluorescent - Cool |
| NIKON D5200 | Prime – 24mm | f/2.2 | Daylight |
| | Prime – 24mm | f/2.2 | Fluorescent - Cool |
| | Prime – 24mm | f/3.5 | Daylight |
| | Prime – 24mm | f/3.5 | Fluorescent - Cool |
| | Prime – 24mm | f/5.6 | Daylight |
| | Prime – 24mm | f/5.6 | Fluorescent - Cool |
| | Prime – 24mm | f/9 | Daylight |
| | Prime – 24mm | f/9 | Fluorescent - Cool |
| | Prime – 24mm | f/14 | Daylight |
| | Prime – 24mm | f/14 | Fluorescent - Cool |
| NIKON D5200 | Zoom – 18mm | f/3.5 | Daylight |
| | Zoom – 18mm | f/3.5 | Fluorescent – Cool |
| | Zoom – 18mm | f/5.6 | Daylight |
| | Zoom – 18mm | f/5.6 | Fluorescent – Cool |
| | Zoom – 18mm | f/9 | Daylight |
| | Zoom – 18mm | f/9 | Fluorescent – Cool |
| | Zoom – 18mm | f/14 | Daylight |
| | Zoom – 18mm | f/14 | Fluorescent – Cool |
| | Zoom – 18mm | f/22 | Daylight |
| | Zoom – 18mm | f/22 | Fluorescent – Cool |
| NIKON D5200 | Zoom – 24mm | f/3.8 | Daylight |
| | Zoom – 24mm | f/3.8 | Fluorescent – Cool |
| | Zoom – 24mm | f/5.6 | Daylight |
| | Zoom – 24mm | f/5.6 | Fluorescent – Cool |
| | Zoom – 24mm | t/9 | Daylight |
| | Zoom – 24mm | t/9 | Fluorescent – Cool |
| | Zoom – 24mm | t/14 | Daylight |
| | Zoom – 24mm | t/14 | Fluorescent – Cool |
| | Zoom – 24mm | t/22 | Daylight |
| | Zoom – 24mm | f/22 | Fluorescent – Cool |

Table 3.4: Camera Response Function Image Sets

To minimize movement during the tests, the camera was mounted on a tripod and tethered via a USB cable to a laptop for remote camera control. Sofortbild was used to control the bracketing of the Nikon camera. With each of the 5 camera-lens combinations, 4-5 different aperture sizes between f/2.2 and f/22 were tested. Furthermore, for each aperture, two different white balance settings were used; daylight and fluorescent. The complete list of image sets taken are provided in Table 3.4. In total, 40 image trials were taken, which will give 40 sets of camera response functions to compare. Each set of images was run through hdrgen to generate and save the camera response functions in a simple text file format.

There is conflicting advice provided in past HDRI papers regarding the optimal white balance setting, if any recommendation or justification is provided at all. A book written by Reinhard, et al, suggests using daylight white balance [35], a commonly adopted practice in several HDR research papers [10] [7] [21], while other papers use white balance settings that are appropriate for the lighting in the scene [22]. The predominant lighting source in this test scene was fluorescent, so the cool-white fluorescent white balance mode was used to compare against the commonly recommended daylight white balance setting.

3.3.2 Results

The camera response functions generated from each set of images are provided in the Appendix. All of the functions are polynomials, typically of 3rd-order, although there are a few 2nd-order and 4th-order curves. These functions were plotted and compared in Mathematica. Based on these quick visual evaluations, it was evident that the camera response functions that used the Daylight white balance mode were very inconsistent and showed significant errors. An example of such an error is shown in Figure 3.5. The corresponding camera response function shows strange behavior in the green channel curve, which is likely the cause for the green hue in the resulting HDR image. There are a few other HDR images obtained from this test with similar saturated color hues, including red (Figure 3.6), yellow-green (Figure 3.7), orange (Figure 3.8), and yellow (Figure 3.9). These too have abnormal camera response functions that appear to correspond with the color shift in the HDR image. The cause for these color abnormalities is not known.



Figure 3.5: Prime f/3.5, Daylight White Balance

Nikon-Prime f/3.5, Daylight White Balance





Figure 3.6: Fisheye f/22, Daylight White Balance

Nikon-Fisheye f/22, Daylight White Balance





Figure 3.7: Zoom 24mm f/22, Daylight White Balance

Nikon-Zoom 24mm f/22, Daylight White Balance




Figure 3.8: Zoom 24mm f/9, Daylight White Balance

Nikon-Zoom 24mm f/9, Daylight White Balance





Figure 3.9: Zoom 24mm f/5.6, Daylight White Balance

Nikon-Zoom 24mm f/5.6, Daylight White Balance







Figure 3.10: Prime 24mm Red Curve, Daylight White Balance

Even among the images that appeared to have appropriate color mapping, the camera response functions determined using the Daylight white balance sets showed significant variations. A plot of the red channel curves obtained from the prime lens image trials is included in Figure 3.10 to demonstrate the significant differences between some of the curves. The only differing parameter of these curves is the aperture size, which should not affect the camera response function.

In stark contrast to these results, the curves found using the Fluorescent white balance images appeared to be consistent across the different aperture sizes and lenses. Out of the 20 image sets for Fluorescent white balance, only 1 curve showed significant deviation. With the zoom lens at 18mm, the camera response function generated at an aperture of f/22 was a 4th-order polynomial, while all of the 19 other Fluorescent CRFs are 3rd-order polynomials. The Fluorescent White Balance equivalent of the plot in Figure 3.11 is provided in Figure 3.11, to demonstrate how closely the red channel curves match across different aperture sizes for the prime lens image sets when using Fluorescent white balance instead.



Figure 3.11: Prime 24mm Red Curve, Fluorescent White Balance

For each camera-lens combination using fluorescent white balance, the coefficients of the RGB channel curves across the 4-5 aperture sizes tested were averaged to create average red, green, and blue channel curves. An example calculation is shown in Table 3.5.

| f/2.2 | 1.911 | X^3 | -1.35305 | \mathbf{X}^2 | +0.450926 | Х | -0.00887 |
|----------|----------|----------------|----------|----------------|-----------|---|----------|
| f/3.5 | 1.94658 | X^3 | -1.41377 | X^2 | +0.477306 | Х | -0.01012 |
| f/5.6 | 1.81632 | \mathbf{X}^3 | -1.256 | \mathbf{X}^2 | +0.447009 | Х | -0.00732 |
| f/9 | 1.89372 | \mathbf{X}^3 | -1.34643 | \mathbf{X}^2 | +0.461174 | Х | -0.00847 |
| f/14 | 1.87197 | X^3 | -1.35904 | \mathbf{X}^2 | +0.496803 | Х | -0.00973 |
| Averaged | 1.887918 | X ³ | -1.34566 | X^2 | +0.466644 | Х | -0.0089 |

Table 3.5: Prime Lens, Fluorescent White Balance Red Channel Curve

The averaged RGB curves for each camera-lens combination for the Nikon camera were plotted together to determine how well the curves matched across lens types, shown in Figures 3.12-3.14. It is clear from these plots that, once again, these curves are all nearly the same. This provides confidence that the camera response functions generated from the fluorescent white balance setting will provide consistent color-space results when creating future HDR images. It also supports the idea that camera response functions should be independent of camera lens, aperture size, and focal length.





Figure 3.13: Averaged Green CRF, Fluorescent White Balance



Figure 3.14: Averaged Blue CRF, Fluorescent White Balance



Now that the Fluorescent curves have been shown to be very similar, the 19 3rd-order polynomials are used to derive an averaged camera response function for the Nikon D5200 for all future HDRI testing. By calculating the mathematical means of the coefficients of the 19 camera response functions found, the resulting final average camera response function is as follows:

$$Red: 1.964674 * x^{3} - 1.44927 * x^{2} + 0.495151 * x - 0.01055$$

$$Green: 1.492516 * x^{3} - 0.82879 * x^{2} + 0.33797 * x - 0.00169$$

$$Blue: 1.539917 * x^{3} - 0.873 * x^{2} + 0.333861 * x - 0.00078$$

To verify that calculating the mathematical means of the coefficients is a valid averaging method for these polynomials, an additional curve fitting procedure was carried out in MATLAB. This was accomplished by creating 100 points of data for each of the 19 camera response function sets found and curve fitting a 3rd-order polynomial the the data points using MATLAB's curve fitting tool. The three (Red, Green, Blue) functions derived using the curve fitting tool are as follows:

$$Red: 1.965 * x^{3} - 1.449 * x^{2} + 0.4952 * x - 0.01055$$

$$Green: 1.493 * x^{3} - 0.8288 * x^{2} + 0.338 * x - 0.00169$$

$$Blue: 1.54 * x^{3} - 0.873 * x^{2} + 0.3339 * x - 0.00078$$

These results are very nearly the same as the results obtained by calculating the mathematical means of the coefficients, with minor rounding differences. The R^2 values of the curve fits are 0.9997, 0.9998, and 0.9997 respectively, indicating very good fits for all three of the curves. This verifies that the camera response function found by calculating the means of the coefficients is a representative average of the 19 camera response functions found. The final averaged camera response function was used in creating the HDR images in the following experiments. The final averaged camera response function is plotted in Figure 3.15.



Figure 3.15: Final Averaged CRF, Fluorescent White Balance

3.3.3 Discussion

The results indicate that the Daylight white balance did not provide useful camera response functions curves for the Nikon D5200, as some HDR research suggests. Instead, it was found that the curves generated from the images that used Daylight white balance resulted in extreme abnormalities and inconsistencies. There was no difference in the way that the Daylight and Fluorescent white balance images were processed through *hdrgen*. Since the results were very consistent with the Fluorescent camera response functions, the following experiments will proceed to use the averaged camera response function derived from the fluorescent-based curves when using the Nikon D5200.

It is recommended that others working on developing their own HDRI system to determine their camera response functions perform a similar test to identify their best white balance setting. The commonly recommended Daylight white balance setting does not appear to behave best for all camera makes and models.

3.4 Determining Vignetting Functions

3.4.1 Background

Vignetting is another important correction to perform before proceeding with HDRI lighting analysis. As mentioned earlier, vignetting describes the effect of diminishing image brightness toward the image boundaries. In HDRI research papers, the term vignetting is used rather loosely to include not only vignetting effects, but also the cosine-fourth and pupil aberration effects. Please refer back to the Literature Search section on vignetting for details on these light loss effects. For the purpose of HDRI, these three effects need not be isolated and can be corrected for all at once with a single polynomial as a function of normalized image radius. This function describes the pattern of diminishing image brightness, and will be referred to here as the vignetting function to match already established HDRI conventions, but please keep in mind that this function also includes the cosine-fourth and pupil aberration effects.

The inverse of the vignetting function can be applied to the original image as a vignetting filter to correct for the luminance losses. As long as the bracketed LDR images used to create the HDR image were all taken with the same aperture size, vignetting filters do not need to be applied to each individual LDR image. The vignetting filter can be applied once to the final composite HDR image because it is assumed that the vignetting losses are consistent for a fixed lens-aperture combination. However, in instances where aperture size needed to be changed during the bracketing, the appropriate vignetting filters should be applied separately. This will not be the case in the following tests, so the vignetting filters will be applied to the final HDR images to save time.

As a quick review, there are a number of methods proposed and used in previous HDRI papers for determining the vignetting function. The premise is the same in all methods; use a target of known luminance and observe the change in portrayed luminance across different locations in the HDR image. The simplest method involves one large, uniform-luminance target that covers the entire field of view of the camera, which remains stationary during the shots. Alternatively, small targets of known or measurable luminance can be arranged to cover the FOV of the camera either in a rotating or stationary setup.

3.4.2 Methodology

The uniform luminance method was used to define the vignetting curves for this HDRI system since there was access to an integrating sphere. The surface of the integrating sphere is a more reliable surface for uniform lighting than a piece of paper or a painted wall and requires less setup than the other methods. The sphere used was about 4-feet in diameter. The sphere's surface was illuminated with a collimated fiber optic beam, which meant that there was a spot of direct light on the surface that needed to be avoided in the images. All other ambient light was kept off during the image acquisition process. LTI Optics in Westminster, CO was kind enough to lend their integrating sphere, equipment, and their testing space to carry out these vignetting tests. An image of the testing apparatus is included in Figure 3.16. For the fisheye images, it was impossible to capture an image without camera shadow or the light beam, but the setup was arranged so that the camera shadow covered as little area as possible in the image. An example of a fisheye image is provided in Figure 3.17. During the vignetting analysis, these pixels were avoided for acquiring data points to obtain the vignetting functions.



Figure 3.17: Fisheye Vignetting Test Image



The uniformity of the integrating sphere surface was checked by taking measurements with the LS-110 luminance meter. The luminance measurements ranged from 2.052 Cd/m^2 to 2.3335 Cd/m^2 , with an averaged luminance value of 2.194 Cd/m^2 . The magnitude of difference between the minimum and maximum measurements is 0.283 Cd/m^2 , meaning a difference in percentage is 12.90%. The absolute difference is quite small, which provides confidence that the sphere is sufficiently uniform, but the percent difference is higher than would be desirable. It seems that the low magnitude of the luminance readings across the entire sphere surface may result in higher SNR and might be more easily skewed by noisy data. In fact, the LS-110 manual reports that measurements taken under 10 Cd/m^2 have a larger window of error (2% +/- 2 digits) [47]. The noise floor for the luminance meter is not reported, but for future reference, it is suggested that similar tests for vignetting should utilize higher luminance levels on the integrating sphere surface to provide stronger signals that are more robust to noise and error.

The HDR images taken were created with *hdrgen* and raw2hdr. The resultant HDR images were analyzed with a custom MATLAB script. The MATLAB script is designed to obtain vectors of data points starting from the geometric center of the image to the image boundary. The script also normalizes the data based on the maximum radius of the image (either the radius of the circular fisheye image or the ray from the center of the image to the corner pixel of the image for non-fisheye images) and an averaged luminance value across a 12x12 pixel square in the center of the image. For each vignetting function, four vectors of data were collected, one in each quadrant. This allowed for an average of the data to be determined across the entire image without any dependency on location in the quadrants. Using multiple rays of data also allows for a check of the assumption of radial symmetry. For the circular fisheye images, 316 points of data were collected for each vector, and 722 points were collected for each vector for the non-circular images. These four vectors of data were plotted and averaged for each image. The vignetting functions are polynomials fitted to the average data points.

Only a relatively small number of different lens-aperture combinations were tested for vignetting. The apertures were selected to cover a representative range of apertures possible with each lens. The list of lens-aperture combinations successfully tested is provided below.

- Sigma Circular Fisheye Lens
 - * f/2.8
 - * f/3.5
 - * f/5.6
 - * f/9
 - * f/14
 - * f/22
- Sigma Prime Lens
 - * f/1.4
 - * f/2.2
 - * f/3.5
 - * f/5.6
 - * f/9
 - * f/14
- Nikkor Zoom Lens at 18mm
 - * f/3.5
 - * f/5.6

105

* f/9

- Nikkor Zoom Lens at 18mm (Cont.)
 - * f/14 * f/22
- Nikkor Zoom Lens at 24mm
 - * f/3.8
 * f/5.6
 * f/14
 * f/22

The vignetting functions obtained for each of the listed lens-aperture combinations for the Nikon D5200 are provided in Table 3.6. The R^2 value and max loss are also provided. The R^2 values are provided to show the quality of fit of the obtained functions to the averaged data points. The max loss values are shown to demonstrate the magnitude of light loss across the different aperture sizes and lenses. Plots of these vignetting curves are provided in Figures 3.18-3.21, categorized by lens model.

| CAMERA | | | | MAX |
|---------|----------|---|----------------|--------|
| LENS | APERTURE | VIGNETTING FUNCTION | \mathbb{R}^2 | LOSS |
| SIGMA | f/2.8 | $-17.133 * x^{6} + 38.535 * x^{5} - 31.188 * x^{4} + 10.819 * x^{3} - $ | 0.9985 | - |
| FISHEYE | | $1.5659^{*}x^{2} - 0.339^{*}x + 1$ | | |
| | f/3.5 | $\textbf{-16.385*x^6+33.624*x^5-22.412*x^4+4.4126*x^3+}$ | 0.9982 | - |
| | | $0.1325^{*}x^{2} - 0.2264^{*}x + 1$ | | |
| | f/5.6 | $-39.877^{*}x^{6} + 96.717^{*}x^{5} - 85.986^{*}x^{4} + 33.619^{*}x^{3} - $ | 0.9636 | - |
| | | $5.2142^*x^2 + 0.199^*x + 1$ | | |
| | f/9 | $-27.279^{*}x^{6} + 69.379^{*}x^{5} - 65.094^{*}x^{4} + 27.291^{*}x^{3} - $ | 0.8535 | 20.51% |
| | | $4.801^{*}x^{2} + 0.2989^{*}x + 1$ | | |
| | f/14 | $-4.871^{*}x^{6} + 12.194^{*}x^{5} - 10.951^{*}x^{4} + 4.0227^{*}x^{3} - $ | 0.9264 | 0.90% |
| | | $0.3939^{*}x^{2} - 0.0017^{*}x + 1$ | | |
| | f/22 | $-1.2593 * x^{6} + 1.1803 * x^{5} + 1.8107 * x^{4} - 2.97 * x^{3} +$ | 0.9034 | 0.34% |
| | | $1.4321^*x^2 - 0.1972^*x + 1$ | | |
| SIGMA | f/1.4 | $0.0638 x^4 + 0.2027 x^3 - 0.5878 x^2 - 0.1743 x + 1$ | 0.9989 | 49.56% |
| PRIME | | | | |
| | f/2.2 | $-0.3571 x^4 + 0.6725 x^3 - 0.574 x^2 + 0.1135 x + 1$ | 0.9874 | 14.51% |
| | f/3.5 | $-0.3501 x^4 + 0.7457 x^3 - 0.6655 x^2 + 0.1419 x + 1$ | 0.9897 | 12.80% |
| | f/5.6 | $-0.0994 x^4 + 0.2694 x^3 - 0.3978 x^2 + 0.0866 x + 1$ | 0.9904 | 14.12% |
| | f/9 | $-0.3133 x^4 + 0.6665 x^3 - 0.5973 x^2 + 0.1188 x + 1$ | 0.9901 | 12.53% |
| | f/14 | -0.2727 x ⁴ + 0.6156 x ³ - 0.5998 x ² + 0.1223 x + 1 | 0.9881 | 13.46% |
| NIKKOR | f/3.5 | $-14.841 x^{6} + 36.901 x^{5} - 32.848 x^{4} + 12.496 x^{3} -$ | 0.999 | 57.26% |
| ZOOM – | | $2.4891 \text{ x}^2 + 0.2085 \text{ x} + 1$ | | |
| 18MM | | | | |
| | f/5.6 | -0.3802 x ⁴ + 0.6936 x ³ - 0.6851 x ² + 0.1078 x + 1 | 0.9967 | 26.39% |
| | f/9 | $-0.1721 x^4 + 0.3334 x^3 - 0.4789 x^2 + 0.0916 x + 1$ | 0.9962 | 22.60% |
| | f/14 | $-0.1231 x^4 + 0.2561 x^3 - 0.4386 x^2 + 0.0701 x + 1$ | 0.9961 | 23.55% |
| | f/22 | $-0.2051 x^4 + 0.4206 x^3 - 0.5328 x^2 + 0.0907 x + 1$ | 0.9935 | 22.66% |
| NIKKOR | f/3.8 | $0.1951 x^4 + 0.0186 x^3 - 0.7624 x^2 + 0.1878 x + 1$ | 0.9959 | 36.09% |
| ZOOM – | | | | |
| 24MM | | | | |
| | f/5.6 | $-0.7989 x^4 + 1.4433 x^3 - 1.0855 x^2 + 0.1993 x + 1$ | 0.9945 | 24.18% |
| | f/14 | $-0.2759 x^4 + 0.5775 x^3 - 0.5905 x^2 + 0.1056 x + 1$ | 0.993 | 18.33% |
| | f/22 | $-0.3331 x^4 + 0.7421 x^3 - 0.7347 x^2 + 0.115 x + 1$ | 0.9917 | 21.07% |

Table 3.6: Vignetting Functions - Nikon D5200

Figure 3.18: Fisheye Lens Vignetting Functions



Figure 3.19: Prime Lens Vignetting Functions





Figure 3.20: Zoom 18mm Lens Vignetting Functions

Figure 3.21: Zoom 24mm Lens Vignetting Functions



3.5 Vignetting Correction Check

The example below, Figure 3.22, demonstrates the effectiveness of applying the vignetting correction function. This example is the vignetting test image for the Sigma Fisheye lens at an aperture of f/2.8. The original image is shown on the left, showcasing a significant example of luminance loss towards the image boundary due to vignetting effects. The image to the right is the same image after the vignetting filter was applied through MATLAB, which appears to be much more uniform in brightness across the entire image.



Figure 3.22: Vignetting Correction Example, Sigma Fisheye Lens at f/2.8

To further verify that these vignetting functions are effective in correcting luminance measurements, a scene with 11 known-luminance targets was captured using the listed lensaperture combinations. The HDR images were processed using the MATLAB script to correct for vignetting. The original and corrected images were compared to the known measurements to evaluate the effectiveness of the vignetting filters in reducing measurement error.

3.5.1 Methodology

The test scene was a white wall, lit from above with incandescent wall-washers. 11 square targets were created using painters tape in a grid pattern across the wall, as illustrated in Figure 3.23. The numbering of the targets is included in Figure 3.24. These targets were measured with the LS-110 luminance meter mounted on a tripod next to the camera. The first two sets of measurements, provided in Table 3.7, were taken for the Nikkor zoom lens at a focal length of 18 mm. The camera was moved closer to the wall for the Nikkor zoom lens and the Sigma prime lens, both at a focal length of 24 mm, in order to accommodate for the new field of view. Measurements 3 through 8 were taken at this distance and are provided in Table 3.8. The camera was moved even closer to the wall for the Sigma fisheye lens. Measurement 9 was taken at this distance.

Ideally, the matte wall should reflect light such that luminance is unvarying with viewing angle, but measurements were taken at each camera location to verify this assumption. It was found that the measurements 1 and 2 taken at the first location farthest from the wall varied slightly from measurements 3 through 9, which were taken closer to the wall. Such differences may be attributed to imperfections in the painted surface or slight glossiness of the paint finish. Measurement 9 was consistent with measurements 3 through 8, so it was assumed that measurements 3 through 9 were representative of the images taken with the Nikkor zoom at 24 mm, the Sigma Prime, and the Sigma fisheye lenses.

The original and corrected HDR images were all calibrated to the average measured value of Target 1 (135 Cd/m^2), which is placed in the center of the image where vignetting effects are non-existent or negligible. The HDR images were analyzed through *hdrscope* to obtain averaged luminance readings for each target area. These measurements from the HDR images are compared to the corresponding averaged luminance meter measurements to obtain error percentage values. It is expected that the corrected images will yield smaller magnitudes of error compared to the original images.



Figure 3.23: Vignetting Verification Test Setup

Figure 3.24: Vignetting Target Numbering



| | Measurement [cd/m ²] | | | | | |
|--------|----------------------------------|-------|--------|--|--|--|
| Target | 1 | 2 | AVG | | | |
| 1 | 135.3 | 134.3 | 134.8 | | | |
| 2 | 122.2 | 120.9 | 121.55 | | | |
| 3 | 104.4 | 104.9 | 104.65 | | | |
| 4 | 145.8 | 146.6 | 146.2 | | | |
| 5 | 152.5 | 153.2 | 152.85 | | | |
| 6 | 190.3 | 190.4 | 190.35 | | | |
| 7 | 102.8 | 102.9 | 102.85 | | | |
| 8 | 213.5 | 213.8 | 213.65 | | | |
| 9 | 114.3 | 113.9 | 114.1 | | | |
| 10 | 80.93 | 81 | 80.965 | | | |
| 11 | 125.4 | 126.2 | 125.8 | | | |

Table 3.7: Nikkor Zoom 18mm Luminance Measurements

Table 3.8: Nikkor Zoom 24mm, Sigma Prime, Sigma Fisheye Luminance Measurements

| | Measurement [cd/m ²] | | | | | | | | | |
|--------|----------------------------------|-------|-------|-------|-------|-------|-------|----------|--|--|
| Target | 3 | 4 | 5 | 6 | 7 | 8 | 9 | AVG | | |
| 1 | 135.7 | 135.2 | 135.5 | 135.1 | 134.8 | 134.3 | 134.3 | 134.9857 | | |
| 2 | 113.8 | 113.9 | 113.5 | 112.9 | 112 | 112.8 | 114.2 | 113.3 | | |
| 3 | 90.15 | 90.49 | 90.9 | 89.69 | 89.41 | 89.59 | 91.16 | 90.19857 | | |
| 4 | 146.9 | 147.2 | 149.9 | 148.3 | 147.3 | 149.5 | 148.7 | 148.2571 | | |
| 5 | 149.7 | 150 | 150.2 | 149.5 | 148.6 | 149.1 | 149 | 149.4429 | | |
| 6 | 210.5 | 211.4 | 211.5 | 210.5 | 210.6 | 210.1 | 211.7 | 210.9 | | |
| 7 | 89.71 | 90.52 | 91.92 | 91.19 | 91.21 | 90.64 | 85.71 | 90.12857 | | |
| 8 | 243.4 | 246.2 | 245.9 | 245.7 | 243.6 | 245.3 | 242.6 | 244.6714 | | |
| 9 | 103 | 103.5 | 103.2 | 103 | 102.5 | 102.5 | 104.6 | 103.1857 | | |
| 10 | 70.09 | 70.37 | 70.48 | 69.62 | 69.11 | 69.15 | 70.05 | 69.83857 | | |
| 11 | 110.9 | 111.3 | 110.3 | 111 | 109.4 | 109.4 | 112 | 110.6143 | | |

3.5.2 Results

The results for all lens-aperture combinations show significant improvements in luminance readings in the HDR images after the vignetting filter is applied. Targets 8-11, which are located farthest from the central target, exhibit the most significant improvement with the applied vignetting filter. An example column chart comparison of the percent errors between the original and vignetting-corrected image for the fisheye lens at an aperture of f/2.8 is shown in Figure 3.25 as one of the worst-case scenarios. The blue bars show the percent error for the original image, which show errors as high as 30% for one of the peripheral targets. The orange bars show the measurements from the corrected images, and demonstrate significant reductions in percent error. All of the corrected measurements are within 10% error.

Results for the other lens-aperture combinations are provided in the following pages. Detailed tables of the luminance readings and the percent errors found with each image are included in the Appendix for reference. The results are grouped by lens and presented in order of largest aperture to smallest. The percent errors for the images taken with the Nikkor zoom lens at 18 mm are based on the average values provided in Table 3.7, and all other images are evaluated based on the averages listed in Table 3.8.



Figure 3.25: Fisheye Lens f/2.8 Percent Error



Figure 3.26: Zoom 24mm Lens, Original Images Error

Figure 3.27: Zoom 24mm Lens, Corrected Images Error





Figure 3.28: Zoom 18mm Lens, Original Images Error

Figure 3.29: Zoom 18mm Lens, Corrected Images Error





Figure 3.30: Prime Lens, Original Images Error

Figure 3.31: Prime Lens, Corrected Images Error





Figure 3.32: Fisheye Lens, Original Images Error

Figure 3.33: Fisheye Lens, Corrected Images Error



3.5.3 Discussion

For the Sigma 4.5mm f/2.8 circular fisheve lens, previous research [7] suggests that vignetting filters diminish in effectiveness for aperture sizes smaller than f/5.6. The results found in this experiment appear to agree with the conclusion from Cauwert, et al.'s research. This is demonstrated in Figure 3.34, a column chart showing the change in percent error from the original to the correct HDR image (% Error of Original Image - % Error of Corrected Image). Positive values indicate improvements, whereas the negative values indicate that the corrected image produced a higher error than the original image. The blue and orange bars represent the largest two apertures, f/2.8 and f/3.5, respectively. For these aperture sizes, the bars indicate significant improvements up to 30% when the vignetting filter is applied. For apertures f/5.6 and smaller, the changes in error are all under 5%. This appears to agree with the statement that vignetting filters have a very small effect in improving the luminance measurements with apertures smaller than f/5.6, even for the peripheral targets 8-11. Targets 6 and 8 even show a slight increase in error when the vignetting filter is applied, although this increase is small and always within 4%. The results from this experiment show that all of the unaltered measurements for apertures f/5.6 and smaller are within 10%, even for the peripheral targets, suggesting that vignetting filters may not be necessary for acceptable errors.

However, for the non-fisheye images, the same pattern is not observed. The results for the other three lenses show that there are still significant reductions in error, about 6-15%, for peripheral targets. Figures 3.35 - 3.37 show these changes in error for the prime and zoom lenses. Despite suggestions and implications in past research that vignetting effects are negligible for smaller apertures, the results here indicate that this is not an applicable assumption for all lenses. Lenses that produce full-frame images do experience significant luminance losses even at smaller apertures, requiring the use of vignetting filters to obtain results that are within an acceptable range.



Figure 3.34: Fisheye Lens, Reduction in Error

Figure 3.35: Prime Lens, Reduction in Error





Figure 3.36: Zoom 18mm Lens, Reduction in Error

Figure 3.37: Zoom 24mm Lens, Reduction in Error



Another finding from the Aggarwal, et al., research [8] suggests that true vignetting effects (luminance loss due to occlusion of off-axis light rays by the physical size of the aperture) is negligible for apertures smaller than f/4.0. The results from this experiment appear to show increased errors for apertures of f/3.5 and larger, which is a likely indication that vignetting is affecting only these larger apertures, as suggested. Apertures f/5.6 and smaller appear to have similar magnitudes of error with no dependence on aperture size, which suggest that the luminance loss is due to the cosine-fourth effect and pupil aberration instead. The cosine-fourth effect produces luminance loss that is not dependent on aperture size, it is an effect of the Gaussian thick-lens model [8].

Overall, the results of applying the vignetting filters provide confidence in the applicability and validity of the previously determined vignetting functions. In every lens-aperture combination, applying the vignetting filters reduced the percent error of all targets to be well within 10%, as demonstrated in Figures 3.27, 3.29, 3.31, and 3.33. The results for the Sigma fisheye lens also show agreement with the results from Cauwert, et als research, which suggests that vignetting filters are not necessary for aperture sizes f/5.6 and smaller while using the Sigma 4.5mm f/2.8 circular fisheye lens. The results for the other three lenses indicate that vignetting filters may be necessary for all aperture sizes to improve luminance readings to be within the accepted 10% error threshold. These findings inform the decision to use vignetting filters for all of the following experiments.

Chapter 4

Integrating Sphere Calibration Rig

A multi-sphere rig was designed and built by Mark Jongewaard prior to these experiments to be utilized for HDRI calibration and testing. The rig's intended purpose is to act as a flexible set of five targets providing a wide range of luminance values with which to calibrate HDR images. The targets are the lenses of five small integrating spheres illuminated with LEDs. The rig is constructed with black aluminum 80/20 framing. The framing members are attached to each other such that they can be easily moved and manipulated, allowing the five targets to cover wider or smaller areas as needed.

4.1 Design

The rig consists of five small integrating spheres made of black anodized aluminum, each with a 1" diameter opening on one face and an opening for an LED on the back. Each sphere is fitted with a unique LED with varying output levels and color temperatures. Diffuse acrylic lenses were cut to fit into the 1" ports to create a uniform luminance target area. The acrylic material used for the lenses is ACRYITE's Satinice WD008 DF at a 2 mm thickness, which has a transmittance of 63.6%. The inside surface of the spheres are painted with a special white paint. The paint was made with barium sulfate and standard flat white latex paint mixed at about a 1:1 ratio, with some added water to improve consistency. The barium sulfate was added to give the white paint a diffuse finish. To block the direct view of the LED from the opening, a small painted mask is mounted inside the sphere. This ensures that the light emitting from the lenses is indirect, diffuse reflected light from the sphere and not due to the direct beam from the LED itself. An image of an open half of a sphere is provided in Figure 4.1, showing the mounted LED and the mask. The spheres are mounted to flat plates that allow them to attach to the bars of the rig. Figure 4.2 shows the fully assembled rig with all five spheres.



Figure 4.1: Back Half of Luminance Sphere

Figure 4.2: Luminance Spheres



A table describing the properties of the five LED spheres is provided below. The LEDs were selected to provide a large range of luminance values, from 10 to $120,000 + Cd/m^2$. Later, neutral density filters were taped on to cover the bottom half of each lens, giving 10 points of different luminance values instead of five. One of the LEDs is an RGBW, meaning that it can be set to emit red, green, or blue light in addition to standard white. The saturated colored LEDs can be used to study how well HDR images can estimate the luminance for sources that emit light in a very limited spectrum.

| Sphere # | LED Model | Current (A) | Voltage | Watts | Lumens | Lumen Output Ratio (Relative to lowest) | Average Luminance (Cd/m²) |
|-------------|---|----------------|---------|-------|--------|---|---------------------------------|
| 1 | Cree CXA 2590 CXA2590-0000- 000R00BB50H | 1.800 | 71.5 | 128.7 | 12350 | 537.0 | 127000 |
| 2 | Citizen CLU034- 1208B8- 503M1A2 | 0.720 | 35.7 | 25.7 | 3475 | 15.1 | 31250 |
| 3 | Cree XP-L | 3.0 | 3.4 | 10.1 | 1010 | 43.9 | 6650 |
| 4 | Cree XM-L RGBW | 0.850 | 3.3 | 2.8 | 216 | 9.4 | 1700 |
| 5 | Nichia NFSL757DT-V1 | 0.065 | 2.9 | 0.2 | 23 | 1.0 | 190 |

Table 4.1: LED Sphere Properties

4.2 Testing Lambertian Nature of Lens

4.2.1 Methodology

Before proceeding with the testing of the LED calibration spheres, one of the spheres was tested alone to determine how well the acrylic lens performs in creating a diffuse, or Lambertian, distribution of light. A Lambertian distribution defines a perfectly diffuse surface that reflects or emits light such that luminous intensity diminishes with the cosine of the viewing angle and luminance is constant over viewing angle. It was desired for the lenses to create a very diffuse port of light so that off-axis viewing angles would not alter the luminance measured at the camera.

To test this, one of the small integrating spheres was mounted to a rotating platform on an optical bench in the lab. The sphere was mounted such that the center of rotation of the sphere was aligned with the front face of the acrylic lens. A luminance meter was set up three feet away from the face of the lens. Figure 4.3 is a photograph of the test setup. The LED was allowed to warm-up for over an hour before measurements were taken. All ambient light was kept off during the measurement process, and black cloth covered the optical bench to minimize reflected light.

Two sets of measurements were taken. The first set measured the uniformity of luminance across the lens face. The measurement point locations are illustrated in Figure 4.4. The second set of measurements tested the uniformity at rotated viewing angles to check the Lambertian distribution. Measurements were taken at 5° intervals. For both tests, four rounds of measurements were taken.



Figure 4.3: Lambertian Test Setup



Figure 4.4: Measurement Points for Lens Uniformity

 Table 4.2: Lens Uniformity Test Results

| LOCATION | 1ST | 2ND | 3RD | 4TH | AVERAGE | UNITS |
|---------------------|-------|-------|-------|-------|---------|-------------------|
| CENTER | 58000 | 58050 | 58180 | 58080 | 58077.5 | Cd/m ² |
| ТОР | 57800 | 57800 | 57990 | 57920 | 57877.5 | Cd/m ² |
| BOTTOM | 57350 | 57500 | 57500 | 57510 | 57465 | Cd/m ² |
| LEFT | 57920 | 57900 | 57980 | 57970 | 57942.5 | Cd/m ² |
| RIGHT | 57900 | 58000 | 58230 | 58130 | 58065 | Cd/m ² |
| TOP-RIGHT | 57850 | 57800 | 57910 | 58000 | 57890 | Cd/m ² |
| TOP-LEFT | 57810 | 57940 | 57980 | 58100 | 57957.5 | Cd/m ² |
| BOTTOM-LEFT | 57710 | 57750 | 57820 | 57860 | 57785 | Cd/m ² |
| BOTTOM-RIGHT | 57500 | 57800 | 57850 | 57040 | 57547.5 | Cd/m ² |

| ANGLE | 1ST | 2ND | 3RD | 4TH | AVERAGE | UNITS |
|-------|------------|-------|-------|-------|---------|-------------------|
| | | | | | | |
| 0 | 58100 | 58240 | 58160 | 58090 | 58147.5 | Cd/m ² |
| 5 | 58190 | 58200 | 58260 | 58100 | 58187.5 | Cd/m ² |
| 10 | 58330 | 58200 | 58350 | 58300 | 58295 | Cd/m^2 |
| 15 | 58390 | 58240 | 58400 | 58380 | 58352.5 | Cd/m ² |
| 20 | 58200 | 57830 | 58250 | 58000 | 58070 | Cd/m ² |
| 25 | 57580 | 57270 | 57570 | 57430 | 57462.5 | Cd/m ² |
| 30 | 56510 | 56430 | 56540 | 56440 | 56480 | Cd/m^2 |
| 35 | 54940 | 55110 | 54940 | 54880 | 54967.5 | Cd/m ² |
| 40 | 53360 | 53040 | 52930 | 53160 | 53122.5 | Cd/m^2 |
| 45 | 50700 | 51090 | 50520 | 50800 | 50777.5 | Cd/m ² |
| 50 | 48430 | 48770 | 48420 | 48650 | 48567.5 | Cd/m ² |
| 55 | 46430 | 45860 | 45510 | 45800 | 45900 | Cd/m ² |
| 60 | 43510 | 43300 | 43040 | 42600 | 43112.5 | Cd/m ² |
| 65 | 40550 | 40250 | 40030 | 39650 | 40120 | Cd/m ² |
| 70 | 37770 | 37250 | 37360 | 37090 | 37367.5 | Cd/m^2 |

Table 4.3: Lens Lambertian Test Results

4.2.3 Discussion

The average luminance of the lens is calculated to be 57,845 Cd/m^2 for this particular LED. The average range of values across the face of the lens is 612.5 Cd/m^2 , which is a difference of 1.06%. With the percent difference across the lens being so small, it can be safely assumed that the acrylic lenses in combination with the integrating sphere provide sufficiently uniform luminance distributions for measurement.

Regarding the check for Lambertian distribution of the lens, the results are normalized in accordance with the 0° measurement (lens perpendicular to the meter) and plotted over viewing angle in Figure 4.5. With an ideal Lambertian distribution, it should be observed that the luminance is unvarying over angle. As demonstrated in Figure 4.5, the measured luminance value of the lens port shows to be nearly uniform for angles less than 20°. Even up to 30° , the averaged luminance reading is 97.13% of the reading at 0° . After 30° , there is a sudden decrease in luminance over angle. Measurements cease after 70° because the LS-110 luminance meter has a finite measuring area, and it is at this point that the projected area of the lens is no longer able to fully cover the measurement area. As mentioned before, the minimum possible measuring area with the LS-110 luminance meter is 4.8 mm in diameter [47].

For the intended application of the calibration spheres, $0 - 30^{\circ}$ was deemed as an acceptable range for luminance uniformity over viewing angle. The spheres are intended to be placed as perpendicular to the camera as possible, to maximize the projected viewing area of the lens. Sufficient uniformity up to 30° provides some flexibility and forgiveness in where the spheres can be located in the scene so that they do not need to be perfectly angled towards the camera.



Figure 4.5: Luminance Over Viewing Angle of Lens
4.3 Discovered Issues with HDRI

During initial testing of the newly acquired HDRI system and the calibration sphere rig, a few issues of HDRI were uncovered. The predominant two issues included dramatic lens flare artifacts and extreme errors in luminance readings of light-emitting surfaces of saturated color. It was decided that further investigation of these issues was outside of the scope of this particular project, but initial observations are included here as an introduction for possible topics to pursue in future research.

4.3.1 Lens Flare

When taking images of the calibration spheres with no ambient light, dramatic starburst lens flare artifacts appeared in the images. During bracketing, these starburst flares occurred in the higher exposures, but were absent in the low exposures. The hope was that when fusing the bracketed images into the final HDR image, the low exposures would prevent the starburst flares. Unfortunately, the lens flare effects are still present in the final HDR image and hinder the ability to obtain reliable and accurate luminance values of the targets. An example HDR image of these starburst effects is shown in Figure 4.6. This image was taken with the Nikkor zoom lens at a focal length of 18 mm and an aperture size of f/22.

Different aperture sizes were then tested to determine if aperture size affected the severity of the lens flare effects. Images of the same scene are taken with the Nikkor zoom lens at a focal length of 18mm and apertures of f/3.5, f/9, and f/22. The three images are provided in Figure 4.7. Based on these preliminary evaluations, it appears that larger apertures reduce the severity of the star-like lens flare. Further research should investigate whether or not these lens flare effects also influence luminance readings in those portions of the image and if increased aperture size mitigates the luminance measurement errors as well.



Figure 4.6: Lens Flare, No Ambient Light, f/22 $\,$



Figure 4.7: Lens Flare and Aperture Size

It was also investigated if ambient lighting levels affect the severity of lens flare. The lab space in which the spheres were set up for this scene is equipped with four rows of direct-indirect fluorescent luminaires. Two different ambient light levels were tested, one configuration with two rows of lights on, "Half On," and one with all four rows of the lights on, "Full On." The images created using the Half On and Full On lighting configurations for aperture f/22, the most severe case of lens flare, are included in Figure 4.8.

Figure 4.8: Lens Flare and Ambient Light Levels



The images clearly show that the presence of ambient light greatly reduces the visibility of the lens flare. The same observations were made in the images taken with the smaller aperture sizes. The difference between the Half On and Full On images is rather subtle, so further testing will need to be done with additional steps of ambient lighting levels to observe the relationship between the ambient light levels and severity of the lens flare. It should also be investigated if ambient lighting improves the luminance readings. It may be that the lens flare effects are less visible with ambient light, but still contribute the same magnitude of error in the luminance readings in the HDR image.

It was concluded that if these calibrating spheres are to be used as a subject of other HDR images, the images need to be taken with larger aperture sizes and ambient light to avoid the starburst lens flare issues.

4.3.2 Color LED Issues

Sphere 4, as previously mentioned, is capable of switching between white, red, green, and blue light output. In regular testing, the white mode was used. Images of the red, green, and blue modes were taken as a preliminary test to observe how well HDR images can predict the luminance of saturated color from a light emitting surface. Significant errors were found; even more so than the errors found in previous research [11] [10] involving saturated colored surfaces that are non-light emitting.

Figure 4.9: Sphere 4 RGBW Color Options



The bracketed LDR images were taken with the Sigma prime lens at an aperture of f/5.6. A larger aperture was selected to avoid lens flare, but apertures that were too large were avoided because of vignetting losses. The appropriate vignetting filter was applied to all of the final HDR images before analysis. The shutter speed range was from 1/4000 second to 1 second. The HDR images were created with *hdrgen*, *raw2hdr*, Bracket, Picturenaut, and Luminance HDR. A more thorough evaluation of these software options is provided in the next section. All of the HDR images were calibrated to the top half of Sphere 5, at a value of 189 Cd/m^2 . This was the selected calibration point since it is a white light target and it yielded the most consistent readings and has significant digits to two decimal places.

The errors of the luminance measurements for the red, green, and blue LEDs are plotted in a column chart in Figure 4.10. The top half of the spheres are designated with the letter 'a,' while the bottom half, which is covered with the neutral density filter, is designated with the letter 'b.' Figures 4.11, 4.12, and 4.13 break out the percent error for each colored LED by software type.

Some initial observations are discussed here, but none of these statements should be interpreted as conclusions. The intended purpose of the following discussion points is to prompt and guide additional research in this specific area of study.

Figure 4.13 illustrates that the blue LED clearly exhibits the largest errors in comparison to the red and green LED. This is likely due to the poor sensitivity camera image sensors have in the blue spectrum of visible light, especially with CMOS sensors [15]. However, it should also be noted that there is error from the LS-110 luminance meter. While the response curve of the LS-110 luminance meter is designed to closely match that of photopic vision, there are discrepancies in the shorter wavelengths, including blue (See Figure 3.2 in Chapter 3). It is also documented in the LS-110 manual that color-correction factors (CCF) are needed for red (CCF=0.995), green (CCF=1.018), and blue light sources (CCF=1.123) [47]. The provided CCFs still do not account for the significant differences between the luminance meter and HDRI measurements.



Figure 4.10: Percent Error for Red, Green, and Blue LEDs

Figure 4.11: Percent Error for Red LED





Figure 4.12: Percent Error for Green LED

Figure 4.13: Percent Error for Blue LED



For the blue LED, *hdrgen* shows the most extreme errors at over 400% for both readings. Picturenaut with RAW images appears to yield the most consistent results with the luminance meter measurements, with errors between 30% and 45%. For both JPEG and RAW-based HDR images created with Picturenaut, the errors between the top and bottom half also appear to be very similar in magnitude, which means that the HDR images created with Picturenaut best captures the high dynamic range between the two halves.

For the green and blue LEDs, errors were much improved with the RAW-based alternatives than the JPEG-based images. This may be a result of the loss of valuable color information that occurs during JPEG compression that is necessary for interpreting the luminance in colored areas of the image. The color information is untampered with in RAW images, which may explain its improved errors.

On average, Picturenaut RAW images produce the least amount of error for all of the colored targets at 31.995% while *raw2hdr* comes in second with an average error of 67.93%. The only errors that were within 10% came from the *hdrgen* generated image of the red LED and the Bracket generated image of the green LED, but only for the top half. The red LED also produced the smallest error on average across all software options with 41.52% for the top half and 51.24% for the bottom.

No conclusions can be drawn now, but these preliminary results have demonstrated the need for additional research involving the use of saturated colored light. Dealing with saturated colored light is a difficult challenge that conventional luminance meters and cameras are not designed to accommodate for, so these initial measurements may not be reliable for absolute luminance analysis. Tools that are designed for colored light applications should be investigated or developed. It is not the intent of the previous statements regarding the performance of the different software options to recommend one over the other. More extensive experimentation and research needs to be done to conclude which software and which tools are best suited for color light applications. These initial observations simply uncover the wide discrepancies in HDR luminance maps and luminance meter measurements when colored light is involved. This may affect lighting designers or researchers who wish to use HDR luminance maps to analyze architectural scenes with colored light. Until further research is conducted, additional images of colored light will be avoided in this project.

Chapter 5

Software Evaluation

5.1 Background

As mentioned earlier, there are several software options for HDR image generation. A few of those options include *hdrgen*, *raw2hdr*, Bracket, Luminance HDR, and Picturenaut. Previous HDRI research papers appear to deal almost exclusively with *hdrgen* and *raw2hdr*. The preference for these software options is not justified or explained in any of these publications, so it was of interest to determine if these options are truly the best performing options available or if there are other viable HDR image creation options that are suited for the needs of lighting research.

pfstools is another available tool for HDR image creation that was not tested during this project. This software is a command-based program that utilizes the Robertson, et al. [33] algorithm for determining the camera response function. Some functionality of pfstmo (tone mapping functions) and pfscalibration (camera response recovery and HDR image creation) is said to be available in Luminance HDR [58]. Unfortunately, pfstools was not successfully compiled during the timeframe of this project, so it is not included in the software evaluation. However, Jacobs and Wilson [18] tested the capabilities of pfscalibration, named pfshdrcalibrate at the time, in their 2007 paper on vignetting. It was concluded that their results between hdrgen and pfstools were equivalent.

5.2 Methodology

To test each software option, two different scenes were selected as the subject of the HDR images. These scenes included a classroom under high fluorescent light, low fluorescent light, and low fluorescent light plus daylight. Another image of this classroom was taken from a different viewpoint to include a view out the window, with only daylight as the light source. The calibration sphere rig was used for the second scene. The rig was photographed in two configurations, one in which the five spheres were spread out as far as the rig would allow and a second where the spheres were clustered in a tighter group. It was verified that the targets were within the acceptable 30° range established previously in Chapter 4, with the farthest sphere being at an angle of approximately 17.5° from the camera's optical axis.

In all scenes, horizontal illuminance measurements were taken to verify that the light levels in the space were consistent during the measurement process. These measurements were taken at desk height, approximately 2.5' off the floor, near the camera. These illuminance values were used only to validate the consistency of the ambient light levels within the space during the image acquisition process and are not used for any calculation or calibration. The averaged horizontal illuminances for each scene are listed in Table 5.1.

| | Illuminance | |
|---|-------------|----------|
| Scene | [FC] | Variance |
| Classroom, High Fluorescent Light | 59.5 | 0.0067 |
| Classroom, Lower Fluorescent Light | 52.4 | 0.0067 |
| Classroom, Low Fluorescent Light plus Daylight | 65.8 | 0.027 |
| Classroom and Window View, Daylight Only | 7.83 | 0.18 |
| Spheres with Ambient Light (All configurations) | 39.54 | 0.0023 |

Table 5.1: Horizontal Illuminance Levels for Test Scenes

Horizontal

All images were taken with the Sigma Prime lens at an aperture size of f/5.6. The prime lens was selected due to its superior optical quality and fixed focal length. The fisheye was avoided for these tests to ensure that the target areas in the image were sufficiently large for analysis. The aperture size was selected to be within the middle of the range of possible apertures for the prime lens to avoid significant vignetting effects at larger apertures and the diffraction effects that occur at smaller apertures. The vignetting filter for this lens-aperture combination was applied to every image through the custom MATLAB script before any image analysis.

$$VignettingFunctionSigmaPrimeLens, f/5.6:$$

-0.0994 * x^4 + 0.2694 * x^3 - 0.3978 * x^2 + 0.0866 * x + 1

In the classroom scenes, the X-Rite Greyscale and Color Checker and basic grey and white card targets were placed in various locations in the scene as reference points. Only the greyscale patches of the colored cards were used since it has already been established that colored surfaces result in greater errors in HDRI analysis [11] [10] and are not reliable points for comparison. The large grey card located in the center of the image frame, Target 1, was used as the calibration point for all of the classroom images. The regular classroom scenes have 29 calibration targets, labeled in Figure 5.1. The classroom scene with the window view has 10 targets, labeled in Figure 5.2.

In the images of the calibration spheres, the reference targets used are the 10 points created by the lens ports and the neutral density filters. The target labeling is provided in Figure 5.3. All of the images are calibrated to Sphere 5 without the filter. This was selected as the calibration point since this target produced consistent measurements and was measured to two decimal places. While the measurements for Spheres 1 and 2 were also very consistent, the meter is only able to provide four significant digits, meaning it is only significant to the hundredths or tenths place.

Figures 5.4 - 5.8 show examples of the tested scenes.



Figure 5.1: Classroom Scene Targets

Figure 5.2: Classroom Window Scene Targets





Figure 5.4: Classroom, Fluorescent Light Only (High and Low)





Figure 5.5: Classroom, Fluorescent plus Daylight

Figure 5.6: Classroom Window, Daylight Only





Figure 5.8: Calibration Spheres, Close Together



The same set of LDR images for each of the six scenes were fused in *hdrgen*, *raw2hdr*, Bracket, Picturenaut, and Luminance HDR for comparison. All HDR images were saved in the Radiance RGBE format with the .hdr extension so that they can be opened and analyzed through *hdrscope*.

In *hdrgen*, the HDR images were all created using the previously determined Camera Response Function, saved in an .rsp file.

$$Red: 1.964674 * x^{3} - 1.44927 * x^{2} + 0.495151 * x - 0.01055$$

Green: 1.492516 * x³ - 0.82879 * x² + 0.33797 * x - 0.00169
Blue: 1.539917 * x³ - 0.873 * x² + 0.333861 * x - 0.00078

All other settings and features were left in the default mode, meaning image alignment and exposure adjustment were left on and over-exposed/under-exposed image removal, lens flare removal, and ghost removal were kept off.

In raw2hdr, no previously determined CRF is needed due to the nature of RAW image files. Again, the same default settings used when creating the HDR image in hdrgen was used when creating images in raw2hdr.

Two JPEG-based images were created using Bracket. One used the pre-determined recovered response function option, and the other was created by recovering a new response based on the input images. It is unknown where this predetermined camera response function is obtained from. For the option to generate a new response curve, there is no option to write this newly generated curve into a file for later use, nor is there an option to upload a user-defined camera response function file. The options to reduce noise and align images were kept off.

Four HDR images were generated though Picturenaut, three using the JPEG image files and one using the RAW image files. The default settings in Picturenaut keep ghost removal, image alignment, and color balancing off. Exposure correction is on by default. The default settings were maintained for all outputs. For the response curve, the default option is to generate a new curve based on the LDR image inputs. There are also two weighting options, the standard default and the other weighting with an emphasis on middle-range values. Both weighting options were used with the JPEG images to create two HDR images. There is an option to load a user defined curve, but it must be in a .crv file, so this feature was not used. A linear curve and a "standard" gamma 2.2 curve can also be used to generate the HDR image, but the linear option was avoided since it is already known that the response curve for the Nikon D5200 does not have a linear response. The standard gamma curve was used to create the last JPEG-based HDR image to see how well it compares to the recovered response curves. The RAW-based HDR image was created with the standard weighting and a new response curve.

Luminance HDR also allows JPEG and RAW input files, so both JPEG- and RAWbased HDR images were created in Luminance HDR as well. For HDR image creation, there are six default profiles for the user to select from. To save time and data storage space, all six default profiles were evaluated with two scenes only; the classroom scene under lower fluorescent light and the classroom scene with fluorescent plus daylight. The results from this test informed the decision to use only the best performing profile for the rest of the scenes, in both JPEG- and RAW-based versions. The features of the six profiles are listed in Table 5.2. The two variables for each profile are the weighting function and the response curve. There are three options for the weighting function; triangular, plateau, and guassian. Traingular weighting emphasizes the information derived from middle exposures. Guassian weighting is similarly weighted, but with different weighting factors. Plateau weighting treats pixel information across all exposures the same [59]. The two response curve options assume linear behavior or follow a nonlinear gamma curve. Gamma curves follow a simple power-law curve, $A * x^{\gamma}$, where A is an arbitrary constant and γ is the gamma level desired [60].

| | Weighting | Response |
|-----------|------------|----------|
| | Function | Curve |
| Profile 1 | Triangular | Linear |
| Profile 2 | Triangular | Gamma |
| Profile 3 | Plateau | Linear |
| Profile 4 | Plateau | Gamma |
| Profile 5 | Gaussian | Linear |
| Profile 6 | Gaussian | Gamma |

| Table | 5.2: | Luminance | HDR. | Profiles |
|-------|------|-----------|------|----------|
| TUDIC | 0.4. | Lummanoo | | T TOHICD |

There is also an option for custom configuration. In custom configuration, the user may choose to upload a user-defined camera response curve, or generate and save a new response function based on the Robertson algorithm [33]. The features for uploading a user-defined response function and saving a new response function were not functioning in the current version of the software. A JPEG-based HDR image was created for each scene using the Robertson response function algorithm and HDR image generating model.

For the following HDR image analysis, *hdrscope* was used. Photosphere's luminance readout provides only three significant digits while *hdrscope* provides luminance values up to two decimal places with the analysis tool. Higher precision was desired for the luminancebased performance evaluation.

5.3 Results

Some observations made on the visual quality of the images is provided prior to the luminance analysis. These images provide visual clues on the differences between each software option. Before proceeding with the software evaluation across *hdrgen*, *raw2hdr*, Bracket, Picturenaut and Luminance HDR, the options created using Luminance HDR and Pictureanut need to be narrowed down first. First, the six default profiles in Luminance HDR will be compared to each other. The best performing of the six profiles will be used in the overall evaluation later. Next, the four different configurations created in Picturenaut will be narrowed down. The default configuration created with JPEG images will be compared against HDR images using RAW images, the mid-emphasis weighting function, and the standard gamma curve. Once these HDR images are narrowed down, all software options will be compared against each other to determine how well each software provides reliable luminance measurements.

5.3.1 Visual Aesthetic Evaluation

A sample batch of the HDR images generated for the Classroom Window scene is included for visual comparison, Figure 5.9. Based on aesthetics alone, some general observations can be made on the performance differences between the software options.

Luminance HDR consistently produces poor quality images that appear very flat and grey. The dynamic range and contrast is severely lacking and colors are not well-preserved. Using the tonemapping operators did not improve the quality of the images. Even changing between the different profile options and using the custom Robertson algorithm made no improvements in the image quality. The cause of the poor image quality is unknown.

Between the JPEG-based HDR images created in *hdrgen*, Bracket, and Picturenaut, the image quality appears to be similar. There are no obvious or alarming differences in contrast or color quality. Any slight differences are likely attributed to differences in tonemapping.

There are some noticeable color differences between the RAW-based and JPEG-based HDR images between *hdrgen* and *raw2hdr* and in Picturenaut. The RAW-based images have a tendency to appear warmer in tone. The cause for the color shift is unknown, although it is thought to be a difference in how the RAW algorithms handle white balance.

However, since visual aesthetic is not of significant importance to the research project at hand, none of these observations will be used to judge the usefulness of one software option over the other. A better performance measure is the ability of producing accurate luminance readings. These aesthetic differences will not be the basis for decision making for future research conducted here, but these observations have been included for those who may be interested in the visual aspect of HDRI.

Figure 5.9: Example HDR Images: a)*hdrgen*, b)*raw2hdr*, c)Picturenaut JPEG, d) Picturenaut RAW, e) Luminance HDR JPEG, f) Luminance HDR RAW, g) Luminance HDR Robertson, h), default Bracket



5.3.2 Luminance HDR

As a quick reminder, there are six default profile options available in Luminance HDR. These six profiles are defined by two variables, the weighting function and the response curve. The two scenes used to test the six profiles were the classroom scene under fluorescent and daylight, and the classroom scene under low fluorescent light. Both JPEG and RAW image files were used to create the HDR images for the scene with fluorescent and daylight, resulting in a total of 18 HDR images to analyze (6 JPEG fluorescent + daylight, 6 RAW fluorescent + daylight, and 6 JPEG low fluorescent).

To compare the linear curve to the gamma curve, the profiles were paired by weighting function type (Profiles 1 & 2, Profiles 3 & 4, and Profiles 5 & 6). Nearly all pairs had slightly smaller errors on average when using the gamma response curve. Only two of the 18 pairs tested showed slight increases in error with the gamma curve. Table 5.3 shows the percent error found for the six JPEG-based HDR images in the Classroom with Fluorescent and Daylight scene for Targets 1-23, as labeled in Figure 5.1. Targets 24-29 were not used because of noisy data. Similar tables using the RAW-based HDR images for the same scene and the JPEG-based HDR images for the Classroom scene with Low Fluorescent Lighting are included in the Appendix. Figure 5.10 shows the percent errors between the linear and gamma response curves with the Plateau weighting function only. Additional figures for the Triangular and Gaussian weighting schemes are included in the Appendix, for the Classroom with Fluorescent and Daylight scene only. The other scenes show similar patterns with no unique differences. For most targets, the difference between the linear and gamma response curve is very small. There are a few instances where the error using the gamma response curve is much lower than with the linear response curve, which is responsible for driving the average down. The targets are ordered from lowest luminance to highest, so there is no apparent correlation between the magnitude of luminance and change in percent error.

| | | TRIANGULAR | | PLAT | PLATEAU | | GAUSSIAN | |
|--------|-----------|------------|---------|---------|---------|---------|----------|--|
| | Avg. | | | | | | | |
| TARGET | Luminance | Linear | Gamma | Linear | Gamma | Linear | Gamma | |
| 23 | 3.64 | 0.05% | 0.05% | 0.22% | 0.16% | 0.16% | 0.11% | |
| 6 | 4.77 | 84.76% | 84.77% | 80.77% | 80.05% | 85.42% | 85.71% | |
| 5 | 4.89 | 84.22% | 84.12% | 80.48% | 79.80% | 85.10% | 85.36% | |
| 14 | 5.39 | 31.57% | 31.54% | 27.07% | 24.77% | 33.25% | 33.95% | |
| 20 | 6.47 | 174.23% | 169.12% | 147.85% | 91.00% | 252.97% | 233.74% | |
| 7 | 12.67 | 280.24% | 180.73% | 153.28% | 91.06% | 268.09% | 245.67% | |
| 19 | 17.18 | 48.07% | 43.25% | 35.75% | 29.44% | 48.15% | 46.41% | |
| 1 | 18.25 | 25.18% | 24.86% | 20.01% | 17.21% | 26.14% | 27.50% | |
| 22 | 21.79 | 59.24% | 58.87% | 54.29% | 53.08% | 59.44% | 60.28% | |
| 8 | 24.97 | 73.29% | 74.45% | 69.65% | 68.96% | 74.67% | 75.30% | |
| 4 | 26.89 | 82.43% | 82.04% | 78.06% | 77.32% | 82.88% | 83.30% | |
| 13 | 29.28 | 85.10% | 85.13% | 81.17% | 80.49% | 85.68% | 86.00% | |
| 18 | 34.66 | 37.50% | 37.43% | 32.31% | 30.40% | 37.84% | 39.04% | |
| 9 | 45.41 | 159.34% | 159.16% | 135.58% | 85.64% | 228.03% | 213.00% | |
| 17 | 63.56 | 86.93% | 86.93% | 83.40% | 82.75% | 87.70% | 87.92% | |
| 10 | 72.40 | 80.66% | 81.04% | 77.17% | 76.38% | 81.95% | 82.29% | |
| 21 | 95.49 | 69.68% | 71.11% | 66.08% | 65.12% | 71.22% | 71.76% | |
| 16 | 101.70 | 47.44% | 47.44% | 41.90% | 40.43% | 47.58% | 48.59% | |
| 11 | 107.63 | 8.00% | 8.05% | 8.00% | 7.18% | 7.59% | 6.66% | |
| 3 | 121.17 | 133.11% | 132.96% | 112.41% | 76.11% | 179.15% | 170.34% | |
| 2 | 125.43 | 78.97% | 78.97% | 74.98% | 74.35% | 79.74% | 80.13% | |
| 12 | 127.00 | 8.66% | 8.66% | 4.71% | 2.10% | 10.31% | 11.23% | |
| 15 | 150.47 | 225.57% | 225.30% | 199.45% | 100.73% | 385.33% | 334.19% | |
| AVE | RAGE | 85.40% | 80.70% | 72.37% | 58.02% | 100.80% | 96.02% | |

Table 5.3: Luminance HDR Profile Errors for Classroom Scene with Fluorescent + Daylight

For the weighting function type, the Plateau weighting function consistently yielded the smallest errors on average. Figure 5.11, created from the percent error values found using the JPEG-based HDR image for the Classroom scene with low fluorescent light and daylight, also shows Profile 4 as having the smallest error across all targets. Additional figures for the other two tested configurations are included in the Appendix, and both show the same behavior for Profile 4.

It was decided that Profile 4, which utilizes the Plateau weighting function and gamma response curve, was the best performing profile available and would be used for the other test scenes. It should be noted that the errors obtained with Profile 4 are still large.



Figure 5.10: Percent Error for Plateau Weighting, Classroom Fluorescent + Daylight JPEG Scene

Figure 5.11: Luminance HDR Profiles, Classroom Fluorescent + Daylight JPEG Scene



5.3.3 Picturenaut

The results for the four Picturenaut HDR images were also analyzed separately to make the overall software evaluation simpler. All six scenes were used to compare the four Picturenaut HDR configurations. Based on overall averages, the only configuration with consistent results across all six scenes is the gamma curve-based HDR image, which always had the greatest error on average. Between the JPEG-Standard Weighting, JPEG-Mid Weighting, and RAW-Standard Weighting, the results appear to be mixed and random. This can be seen in Figure 5.12. Detailed tables of the errors obtained with Picturenaut are included in the Appendix.



Figure 5.12: Picturenaut Comparisons

The relationships between these configurations are less obvious and pronounced than with the Luminance HDR profiles, so these results were analyzed statistically to verify that the gamma-based HDR images produce the largest error, and that the other options perform equally. A matched-pairs t-test was performed to test the equality of mean errors ($\mu 1 = \mu 2$) between the standard-weighted JPEG-based HDR image and the other three alternatives. The results are summarized in Table 5.4. The r_{12} value is the correlation coefficient between the two sets of data. The values under Correlation Test are the p-values of the correlation test to check for valid correlation, and the values under Paired t-Test are the p-values for the test for equal means. A confidence interval of 99% was selected, meaning p-values of less than 0.01 are required for the hypotheses of the tests to be rejected. The data pairs used were from the results of the classroom scene under high fluorescent light since it had the most pairs of data points, 29. Using a confidence interval of 99%, there were no statistically significant difference in means between the standard JPEG-based HDR image and the standard RAWbased HDR image (p=0.625) and the standard JPEG-based HDR image and mid-weighted JPEG-based HDR image (p=0.011). There was a statistically significant difference in means between the standard JPEG-based HDR image and the gamma curve JPEG-based HDR image $(p=0.000^*)$. This implies that using the gamma curve produces higher errors on average. There is no significant difference between using JPEG or RAW images, and between using the standard weighting or the mid-emphasis weighting.

Since the gamma-based HDR image is shown to be the worst performing option of the four Picturenaut configurations tested, it will be excluded from the overall software evaluation. The mid-weighted HDR image will also be excluded since it has not shown to provide any improvements from the default standard weighting function. The RAW-based standard weighted HDR images will be kept in the overall software evaluation to compare against the other RAW-based HDR images.

| | | | NONLINEAR VS |
|------------------|---------|--------------|--------------|
| | | STANDARD VS | GAMMA |
| | JPEG VS | MID-EMPHASIS | RESPONSE |
| | RAW | WEIGHTING | CURVE |
| r 12 | 0.9860 | 0.9997 | 0.9841 |
| CORRELATION TEST | 0.000* | 0.000* | 0.000* |
| PAIRED T-TEST | 0.625 | 0.011 | 0.000* |

Table 5.4: Matched-Pairs t-Test Results for Picturenaut HDR Images

5.3.4 Luminance-Based Performance

The overall software evaluation looks at HDR images created in hdrgen, raw2hdr, Bracket with previously-determined and newly recovered response functions, Picturenaut with JPEG and RAW images with standard weighting, Luminance HDR with Profile 4 with JPEG and RAW images, and lastly Luminance HDR with the Robertson algorithm. Table 5.5 lists the errors on average for each scene and each HDR output configuration as a big picture view of the results. These averages are also shown in a column chart in Figure 5.13. Detailed tables of the errors and luminance values for each target in each scene are included in the Appendix. An example of the plotted percent errors for each target in ascending order of luminance is provided in Figure 5.14 for the classroom scene under high fluorescent light.

| | HDRGEN | RAW2HDR | BRACKET | BRACKET, NEW CRF | PIC. JPEG | PIC. RAW | LUM. HDR JPEG | LUM. HDR RAW | LUM. HDR ROBERTSON |
|----------------------|--------|---------|---------|---------------------|-----------|----------|------------------|-----------------|-----------------------|
| Class Fluor.+Day | 5.48% | 3.33% | 10.14% | 5.95% | 11.37% | 11.51% | 58.02% | 59.47% | 54.45% |
| Class High Fluor. | 3.77% | 3.66% | 3.97% | 5.25% | 12.10% | 11.90% | 69.92% | 65.78% | 66.78% |
| Class Low Fluor. | 4.86% | 4.27% | 5.65% | 5.75% | 12.35% | 12.65% | 65.16% | 65.54% | 61.87% |
| Class Window | 3.49% | 4.28% | 5.41% | 6.45% | 8.00% | 5.85% | 46.03% | 64.35% | 42.01% |
| Spheres Far | 11.87% | 7.47% | 25.43% | 25.56% | 29.80% | 31.29% | 164.29% | 147.89% | 168.70% |
| Spheres Close | 12.10% | 6.50% | 24.13% | 24.15% | 31.86% | 32.13% | 175.10% | 153.51% | 181.36% |

Table 5.5: Average Errors



Figure 5.13: Column Chart of Average Errors

Figure 5.14: Errors in Classroom Scene with High Fluorescent Light



Right away, it is very obvious that Luminance HDR is consistently the worst performing option in luminance measurement in all configurations. In the classroom scenes, the errors with Luminance HDR reached as high as 200%. In the scenes with the calibrating spheres, the errors were even more extreme, reaching over 1000% in some cases. As could be expected with the visual evaluation, the Luminance HDR images appear to have a very limited dynamic range. The luminance readings obtained from the Luminance HDR images are within a very close range of the calibration point. So while luminance values that are close in magnitude to the calibration point, points that are either much lower or much higher than this luminance point experience extreme errors. This increase in error as the luminance diverges from the calibration point is illustrated in Figure 5.14. Because of these extreme errors, Luminance HDR is eliminated from further evaluation to narrow down the list of viable options for luminance analysis.

With the elimination of Luminance HDR, the results of the other HDR images are easier to differentiate between the remaining six configurations. A new plot of the errors obtained in the classroom scene under high fluorescent light with the Luminance HDR data sets removed provide a more telling picture on the performance of the remaining six options, Figure 5.15. Similar plots for the other five scenes are also included here.



Figure 5.15: Classroom High Fluorescent Percent Errors without Luminance HDR

Figure 5.16: Classroom High Fluorescent + Daylight Percent Errors without Luminance HDR





Figure 5.17: Classroom Low Fluorescent Percent Errors without Luminance HDR

Figure 5.18: Classroom Window Percent Errors without Luminance HDR





Figure 5.19: Spheres Far Apart Percent Errors without Luminance HDR

Figure 5.20: Spheres Close Together Percent Errors without Luminance HDR



The next obvious pattern is that the Picturenaut images produce higher errors on average, particularly for the classroom scenes. The Picturenaut results also show a relationship between error and an increase in luminance. The correlation between error and luminance is tested for statistical significance using a one sample t-test for correlation, $\rho = 0.0$. Since the behavior between the Picturenaut JPEG and RAW images have already been determined to be similar enough with statistical significance, only the JPEG-based HDR images will be tested for correlation. The results of the correlation tests are provided in Table 5.6. The r values are the Pearson coefficient of correlation, where a value of 1.0 indicates a perfect relationship, and the r^2 value indicates the magnitude of statistical significance, where a value of 1 is the most significant. The one-sample t-test is based on a confidence interval of 99%, meaning the p-value needs to be greater than 0.01 for the hypothesis to be rejected. From the results in Table 5.6, all six scenes show a statistically significant positive correlation between error and luminance.

Table 5.6: Correlation Test Results between Error and Luminance for Picturenaut HDR Images, 99% Confidence Interval

| | r | r ² | P-VALUE |
|-------------------|--------|----------------|----------------|
| CLASS FLUOR. + | 0.9597 | 0.921 | 0.000 |
| DAYLIGHT | | | |
| CLASS HIGH FLUOR. | 0.8694 | 0.7559 | 0.000 |
| CLASS LOW FLUOR. | 0.8958 | 0.8025 | 0.000 |
| CLASS WINDOW | 0.7757 | 0.6017 | 0.008 |
| SPHERES FAR | 0.8695 | 0.756 | 0.001 |
| SPHERES CLOSE | 0.8362 | 0.6992 | 0.003 |

For the classroom scenes, raw2hdr produces the smallest error on average for all of the scenes, except for the scene with the window view. Conversely, *hdrgen* is the second best performing for these scenes, except for in the window scene, in which it is the best. Both of the Bracket HDR images are also relatively good performing options after raw2hdr and *hdrgen*, with the average error being within 10.2%. In the scenes with the calibrating spheres,

raw2hdr has the smallest error on average with hdrgen close behind again. Both Bracket images produce average errors between 24% and 32% for the calibrating sphere scenes. Since raw2hdr and hdrgen consistently produce lower errors than Bracket, Bracket will be ruled out as the best performing software option. Before moving on to compare only hdrgen and raw2hdr, the difference between the two Bracket outputs will be checked.

Between the two Bracket HDR images, it is unclear which of the two options is the best performing. One option utilizes a predetermined camera response function based on the camera body information provided in the EXIF headers of the image, but the source from which this function is obtained is unknown. The other option is to generate a new response function based on the image inputs. Again, a matched pairs t-test with a 99% confidence interval was used to determine if the mean errors are equivalent. For most scenes, these two options perform equally well with no statistically significant difference in means. Only for the classroom scene with low fluorescent light plus daylight did the newly recovered response function file perform significantly better than the HDR image created with the predetermined response curve ($p=0.000^{*}$). From these results, no conclusion can be made over which method is better for generating more reliable luminance readings. However, because of the inability to save and reuse the newly generated response curves in Bracket, using the predetermined response curve is likely to provide more consistent results by using the same known response curve every time.

Table 5.7: Matched Pair t-Test Results for Bracket HDR Images

| | r 12 | CORRELATION | PAIRED |
|--------------------------------|-------------|-------------|---------------|
| | | TEST | T-TEST |
| CLASS FLUOR. + DAYLIGHT | 0.718 | 0.000 | 0.000* |
| CLASS HIGH FLUOR. | 0.5216 | 0.004 | 0.024 |
| CLASS LOW FLUOR. | 0.3192 | 0.138 | 0.885 |
| CLASS WINDOW | 0.8891 | 0.001 | 0.077 |
| SPHERES FAR | 0.9981 | 0.000 | 0.887 |
| SPHERES CLOSE | 0.9988 | 0.000 | 0.971 |

Lastly, it will be checked to determine if there is a statistically significant difference in the errors produced by the *hdrgen* and raw2hdr programs. Based on average errors, it is expected that raw2hdr is the best performing option with the least amount of error, as it has been shown to be the lowest average for five out of the six scenes tested.

| | ľ 12 | CORRELATION | PAIRED |
|--------------------------------|-------------|-------------|---------------|
| | | TEST | T-TEST |
| CLASS FLUOR. + DAYLIGHT | 0.0895 | 0.685 | 0.017 |
| CLASS HIGH FLUOR. | 0.7543 | 0.000* | 0.773 |
| CLASS LOW FLUOR. | 0.5985 | 0.003* | 0.221 |
| CLASS WINDOW | -0.3382 | 0.339 | 0.675 |
| SPHERES FAR | 0.5443 | 0.104 | 0.201 |
| SPHERES CLOSE | 0.1189 | 0.744 | 0.181 |

Table 5.8: Matched Pair t-Test Results for hdrgen and raw2hdr HDR Images

Contrary to what the average errors appear to suggest, the matched pair t-tests show no statistically significant difference between the errors produced with hdrgen and raw2hdr. It cannot be concluded that one produces smaller errors on average over the other. Regarding the indoor classroom scenes, this conclusion seems fair. When looking at the errors obtained in the calibrating sphere scenes, raw2hdr is the only software option that produces average errors within 10%. It may be that challenging scenes with large dynamic ranges, such as those with the calibrating spheres (approximately 12,700:1), is where uncompressed RAW image file data and raw2hdr are suited best. Unfortunately, these two scenes are not enough to draw any conclusions. Additional scenes with large dynamic range, like mid-afternoon scenes with a view outdoors and a complete outdoor scene, may start to show a pattern where raw2hdr produces less errors than hdrgen.
5.3.5 Comments on Usability

During the creation process of the many HDR images used in this experiment, a few observations were made regarding the usability of each software. Some comments of note are included here to inform other researchers what to expect when attempting to use the software options mentioned here.

Without any interface, hdrgen and raw2hdr are certainly the most difficult software options to learn and use. They cannot be download as simply as other programs and require compiling before use. As command-line programs, they are also unforgiving and extremely sensitive to typos. Such typos will prevent the program from running, so the user must be very meticulous when typing out the commands. When running large batches of images, the use of commands to generate HDR images can be an extremely tedious process. However, once the user is familiar with hdrgen and raw2hdr, these programs work very well. No computer crashing or stalling has been experienced while using these programs, and they also execute relatively quickly. raw2hdr takes considerably more time than hdrgen because of the larger input files, but run time is still within a couple of minutes.

Picturenaut was easy to use and worked well. The user interface makes it simple to select the LDR images and customize the output options. No crashing or computer stalling was experienced with Picturenaut, and run times were always within a few minutes, with RAW images again taking more time to upload.

Bracket worked well for most cases, but there appears to be an upload size limitation that causes the program to crash and quit unexpectedly. Whenever this error would occur, fewer LDR images needed to be selected to create the HDR image. This was not ideal since it was desired to minimize the amount of variables affecting the HDR image quality between software options, but it could not always be avoided. Otherwise, the user interface is easy to use and navigate.

Luminance HDR, while easy to use and navigate, ran poorly when attempting to

create many HDR images in a row. With JPEG images, the HDR images could be created in a comparable amount of time as the other programs, but severe computer stalling would occur after a handful of HDR images had already been created. There were also a few instances of unexpected crashing. When creating as many HDR images as was required in this experiment, the stalling and crashing became problematic. HDR images had to be created in small batches, ending with a complete exit of the program before starting another batch. This suggests that there may be a problem with memory management in the software code. With RAW images, the upload time and image creation time is significantly increased.

5.4 Discussion

If the only concern for HDR image creation is aesthetics, *hdrgen*, Bracket, and Picturenaut are all fair options, but the difficulty of using *hdrgen* may warrant the use of Bracket or Picturenaut for an application in which true luminance is not needed. Bracket and Picturenaut also have more tonemapping capabilities, so these programs appear to be better tailored for aesthetic and visual applications. It is also recommended to use JPEG files over RAW files because the benefit of uncompressed color information does not appear to provide any noticeable difference in image quality to counteract the large file sizes.

For future research projects to be conducted at the University of Colorado Boulder, it is strongly suggested to continue using hdrgen and raw2hdr. They have been proven to produce reasonable and acceptable errors on average within 10% for nearly all cases. The ability to write and then read in user-defined camera response functions may be the most important deciding factor to use hdrgen over Picturenaut and Bracket. The camera response function has already been determined for the Nikon D5200, and using the same response function to create all of the HDR images used in future research projects will ensure some level of consistency.

Although raw2hdr has also been shown to produce small errors and is argued to maintain color information that is lost in compressed JPEG images, the exclusive use of raw2hdr is not recommended. These experiments show no statistically significant advantage of raw2hdr over hdrgen, so it seems the drawbacks of raw2hdr outweigh the benefits in most applications. In experiments that involve running several batches of images, like this experiment, the storage of RAW image files can become problematic. The size of RAW image files can range from about 19 MB up to 25 MB, whereas a large JPEG image might reach up to 7 MB, but is typically smaller. It is also slower to create RAW-based HDR images, which can become a crucial factor in experiments that involve creating several HDR images. It is suggested that the use of raw2hdr is reserved for more important experimentation where higher accuracy and precision are required. raw2hdr is likely more suitable for color applications too, although additional testing with colored targets is needed to verify this claim.

Chapter 6

Conclusion

6.1 Final Remarks

An HDRI system for future research use is now established for the University of Colorado Boulder. The camera response function of the Nikon D5200 camera body has been determined and is included below for future use. This response function should be used to create all JPEG-based HDR images when using *hdrgen*, Photosphere, or WebHDR.

$$Red: 1.964674 * x^{3} - 1.44927 * x^{2} + 0.495151 * x - 0.01055$$

$$Green: 1.492516 * x^{3} - 0.82879 * x^{2} + 0.33797 * x - 0.00169$$

$$Blue: 1.539917 * x^{3} - 0.873 * x^{2} + 0.333861 * x - 0.00078$$

A list of vignetting filters and a MATLAB script that is capable of applying these filters have also been established for future use. The list of the vignetting functions are provided in Table 3.6 of Chapter 3. The MATLAB script is included in the Appendix. It has been determined that vignetting filters should be applied in almost all cases to ensure minimized errors in the image periphery due to vignetting losses. Images taken with the Sigma fisheye lens at apertures smaller than f/5.6 (greater in number) are exempt from this recommendation.

The development of a novel calibration system using integrating spheres and LEDs has been documented. Preliminary tests of the HDRI system with the calibration spheres as the subject of the HDR images has uncovered a number of important issues that need to be addressed in future research projects. Large errors have been discovered when trying to measure the luminance of a colored light-emitting surface, especially for blue light. Starburst lens flare issues have also been discovered, and further investigation needs to be done to quantify the effects on luminance measurements.

Finally, a handful of HDR image creation software has been tested for their application towards luminance analysis and lighting research. The results suggest that *hdrgen* and *raw2hdr* are currently the best suited for HDRI research, which provides some verification and validation as to why nearly all published HDRI paper have used these two software options exclusively. These two options are strongly suggested for future use in HDRI-related research. Bracket and Picturenaut both showed fair performance in luminance measurement, but their capabilities appear to be tailored more for visual and aesthetic purposes and less for scientific luminance analysis. Many problems were encountered with Luminance HDR in its current version, which suggest that there are bugs or other improvements that need to happen on the software development side.

6.2 Future Research Topics

As mentioned in the Introduction, the long-term goal of this HDRI research project is to pursue a method using HDRI to simulate near-field photometry. This will be modeled to build off of DiLaura and Chu's [12] ideas and use luminance data to supplement pre-existing far-field photometry to predict near-field illuminance distributions.

Based on the results of the preceding experiments, *pfstools* should undergo similar tests to determine its performance and capabilities in relation to *hdrgen*, *raw2hdr*, Bracket, Picturenaut, and Luminance HDR. Additional testing of *hdrgen* and *raw2hdr* should also be conducted to include colored targets. Previous research has already shown poor luminance measurements of saturated colored surfaces using JPEG files and *hdrgen*, but a similar testing procedure using RAW files and *raw2hdr* has yet to be published. This would help validate several claims that have been made that RAW images are the best option for colorimetric applications.

Additional testing of the calibrating spheres also needs to be done to further investigate the errors associated with colored light-emitting surfaces and lens flare artifacts. Using the saturated colored LEDs in the calibration sphere tests uncovered the issue of extreme errors in HDR luminance measurements of surfaces that emit light in a very limited, non-white spectrum. The calibration sphere tests also uncovered issues with severe lens flare or PSF effects that could possibly affect the accuracy of luminance measurements. Extensive testing needs to be done with both of these topics to be able to make any concluding statements.

Other potential topics included at the end of the Literature Search are also viable options for future research projects involving the HDRI system developed here. These topics include the performance differences between CCD- and CMOS-based camera models in measuring luminance in architectural scenes, or the development of new all-in-one HDRI software that combines the capabilities of HDR creating software and HDR analysis software.

Bibliography

- [1] Berthold Horn. Robot Vision. MIT, 1986.
- [2] Bruce Fraser. Understanding Digital Raw Capture. Adobe Systems Inc., 2004.
- [3] Jae Yong Suk, Mark Schiler, and Karen Kensek. Development of new daylight glare analysis methodology using absolute glare factor and relative glare factor. *Energy and Buildings*, 64:113–122, 2013.
- [4] Axel Jacobs. High dynamic range imaging and its application in building research. Advances in Building Energy, 1(1):177-202, 2007.
- [5] L. Bellia and G. Spada. Photometric characterisation of small sources with high dynamic range illuminance mapping. *Lighting Research Technology*, 46:329–340, 2013.
- [6] Axel Jacobs. WebHDR. http://www.jaloxa.eu/webhdr/index.shtml.
- [7] Coralie Cauwerts, Magali Bodart, and Arnaud Deneyer. Comparison of the vignetting effects of two identical fisheye lenses. *LEUKOS*, 8(3):181–203, January 2012.
- [8] M. Aggarwal, H. Hua, and N. Ahuja. On cosine-fourth and vignetting effects in real lenses. In *Eight IEEE International Conference on Computer Vision*, pages 472–479, 2001.
- [9] Feng Xiao, Jeffrey M. Dicarlo, Peter B. Catrysse, and Brian A. W. High dynamic range imaging of natural scenes. In *The Tenth Color Imaging Conference*, pages 337–342. Elsevier, Morgan Kaufmann, 2002.
- [10] MN Inanici. Evaluation of high dynamic range photography as luminance data acquisition system. *Lighting Research and Technology*, 32, 2006.
- [11] Smita Anaokar and Martin Moeck. Validation of high dynamic range imaging to luminance measurement. *LEUKOS*, 2, 2005.
- [12] Wai-Lam Chu and David L. DiLaura. Improved near-field illuminance calculations using far-field photometry and luminance scans. *Journal of the Illuminating Engineering Society*, 24(2):3–7, 1995.

- [13] David L. DiLaura, Kevin W. Houser, Richard G. Mistrick, and Gary R. Steffy. Illuminating Engineering Society, The Lighting Handbook, Tenth Edition: Reference and Application. Illuminating Engineering Society of North America, 10th edition, 2011.
- [14] Lihua Zhang, Yongjin Jin, Lin Lin, Jijun Li, and Yungang Du. The comparison of CCD and CMOS image sensors. In 2008 International Conference on Optical Instruments and Technology: Advanced Sensor Technologies and Applications. SPIE, 2009.
- [15] James Janesick and Gloria Putnam. Developments and applications of high-performance CCD and CMOS imaging arrays. The Annual Review of Nuclear and Particle Science, (53):263–300, 2003.
- [16] CCD sensors. http://www.canon.com/technology/s_labo/light/003/04.html.
- [17] CMOS sensors. http://www.canon.com/technology/s_labo/light/003/05.html.
- [18] Axel Jacobs and Mike Wilson. Determining lens vignetting with HDR techniques. In XII National Conference on Lighting, Light'2007, 2007.
- [19] Mark Rea and I.G. Jeffrey. A new luminance and image analysis system for lighting and vision i. equipment and calibration. The Journal of the Illuminating Engineering Society, 19(1):64-72, 1990.
- [20] Martin Moeck, PhD. Accuracy of luminance maps obtained from high dynamic range images. LEUKOS, 4(2):99–112, October 2007.
- [21] TM Chung and Roger T.H. NG. Variation of calibration factor over time for high dynamic range photography in a single daylit interior scene. Journal of Light & Visual Environment, 34(2):87–93, 2010.
- [22] H Cai and TM Chung. Improving the quality of high dynamic range images. Lighting Research Technology, 43:87–102, 2011.
- [23] Martin Moeck and Smita Anaokar. Illuminance analysis from high dynamic range images. LEUKOS, 2(3):211–228, January 2006.
- [24] Mehlika Inanici. Evaluation of high dynamic range image-based sky models in lighting simulation. LEUKOS, 7(2):69–84, 2010.
- [25] Abhijit Sarkar and Richard G. Mistrick. A novel lighting control system integrating high dynamic range imaging and dali. *LEUKOS*, 2(4):307–322, 2006.
- [26] H. Cai. High dynamic range photogrammetry for synchronous luminance and geometry measurement. Lighting Research Technology, 45:230-257, 2013.

- [27] Kevin Van Den Wymelenberg, Mehlika Inanici, and Peter Johnson. The effect of luminance distribution patterns on occupant preference in a daylit office environment. *LEUKOS*, 7(2):103–122, 2010.
- [28] H. Cai and T. Chung. Evaluating discomfort glare from non-uniform electric light sources. Lighting Research Technology, 45:267–294, 2013.
- [29] Yulia Tyukhova and Clarence Waters. An assessment of high dynamic range luminance measurements with led lighting. *LEUKOS*, 10(2):87–99, 2014.
- [30] Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '97, pages 369–378, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.
- [31] S. Mann and R.W. Picard. Being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures. Technical Report 323, M.I.T. Media Lab Perceptual Computing Section, Boston, Massachusetts, 1994.
- [32] Tomoo Mitsunaga and Shree K. Nayar. Radiometric self calibration. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 374–380, June 1999.
- [33] Mark A. Robertson, Sean Borman, and Robert L. Stevenson. Estimation-theoretic approach to dynamic range enhancement using multiple exposures. *Journal of Electronic Imaging*, 12(2):219–228, April 2003.
- [34] S. Mann. Compositing multiple pictures of the same scene. In Proceedings of the 46th Annual IST Conference. The Society of Imaging Science and Technology, 1993.
- [35] Eric Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec, Wolfgang Heidrich, and Karol Myxzkowski. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting.* 2nd edition, 2010.
- [36] Seon Joo Kim and Marc Pollefeys. Robust radiometric calibration and vignetting correction. Pattern Analysis and Machine Intelligence, IEEE Transactions on, 30(4):562–576, April 2008.
- [37] Richard G. Mistrick and Cheryl R. English. A study of near-field indirect lighting calculations. *Journal of the Illuminating Engineering Society*, 18(2):103-112, 1990.
- [38] Steve Stannard and John Brass. Application distance photometry. Journal of the Illuminating Engineering Society, 19(1):39–46, 1990.
- [39] Joseph B. Murdoch. Inverse square law approximation of illuminance. Journal of the Illuminating Engineering Society, 10(2):96–106, 1981.

- [40] V Jacobs, MSc, S Forment, PhD, P Rombauts, PhD, and P Hanselaer, PhD. Near-field and far-field goniophotometry of narrow-beam LED arrays. *Lighting Research Technology*, 47(4):470–482, 2014.
- [41] Commission Internationale de l'Eclairage. CIE 121 -1996 the photometry and goniophotometry of luminaires.
- [42] P.Y. Ngai. On near-field photometry. Journal of the Illuminating Engineering Society, 16(2):129-136, 1987.
- [43] Active D-lighting. http://www.nikonusa.com/en/learn-andexplore/article/fsqd6p6h/active-d-lighting.html.
- [44] Nasim Mansurov. Understanding ISO a beginner's guide. https://photographylife.com/what-is-iso-in-photography.
- [45] Nasim Mansurov. Understanding metering and metering modes. https://photographylife.com/understanding-metering-modes.
- [46] Sigma lenses. http://www.sigmaphoto.com/lenses.
- [47] Luminance Meter LS-100 LS-110. Manual.
- [48] Konica Minolta Light Meters. Manual.
- [49] Smart shooter. http://kuvacode.com/.
- [50] digiCamControl. http://digicamcontrol.com/.
- [51] Sofortbild. http://www.sofortbildapp.com/.
- [52] Greg Ward. Anyhere software. http://www.anyhere.com/.
- [53] bracket. http://www.ceng.metu.edu.tr/ akyuz/bracket/bracket.html.
- [54] Luminance HDR. http://qtpfsgui.sourceforge.net/.
- [55] Picturenaut. http://www.hdrlabs.com/picturenaut/index.html.
- [56] Viswanathan Kumaragurubaran and Mehlika Inanici. hdrscope: High dynamic range image processing toolkit for lighting simulations and analysis. In International Building Performance Simulation Association 2013 Conference, pages 3400–3407, 2013.
- [57] Viswanathan Kumaragurubaran and Mehlika Inanici. hdrscope. http://courses.washington.edu/hdrscope/index.html.
- [58] pfstools. http://pfstools.sourceforge.net/index.html.
- [59] Qtpfsgui notes: The front end HDR maps, June 2009. June 2009.
- [60] Wikipedia. "Gamma Correction".

Appendix A

Glossary

A.1 List of Symbols

- Q_e : Radiant Energy
- Φ_e : Radiant Flux
- E_e : Irradiance
- L_e : Radiance
- Φ : Luminous Flux
- E: Illuminance
- *I*: Luminous Intensity
- L: Luminance
- X: Exposure
- Z: Digital Image Pixel Value

A.2 List of Acronyms and Initialisms

- ADC: Analong-to-digital conversion
- **BGI:** British Glare Index
- **BRDF:** Bi-directional Reflectance Distribution Functions
- **CCD:** Charge-Coupled Device
- **CCT:** Correlated Color Temperature
- ${\bf CF}{\bf :}$ Calibration Factor
- ${\bf CGI:}\ {\rm CIE}\ {\rm Glare}\ {\rm Index}$
- **CMOS:** Complementary Metal Oxide Semiconductor
- ${\bf CRF}{\bf :}$ Camera Response Function
- **DOF:** Depth Of Field
- $\mathbf{DGI:}\ \mathrm{Daylight}\ \mathrm{Glare}\ \mathrm{Index}$
- **DGP:** Daylight Glare Probability
- **DSLR:** Digital Single-Lens Reflex
- ${\bf FOV:}\ {\bf Field}\ {\bf Of}\ {\bf View}$
- HDR: High Dynamic Range
- HDRI: High Dynamic Range Imaging
- **IBL:** Image Based Lighting
- LDR: Low Dynamic Range

LED: Light Emitting Diode

LID: Light Intensity Distribution

MTF: Modulation Transfer Function

PBR: Physically Based Rendering

 \mathbf{PSF} : Point Spread Function

RGB: Red Green Blue

 \mathbf{SNR} : Signal-to-noise ratio

UGR: Unified Glare Rating

VCP: Visual Comfort Probability

A.3 Lighting Definitions

- Radiant Energy: Amount of electromagnetic energy that can be emitted, transferred, or received in the form of radiation. Units: Joules Symbol: Q_e
- Radiant Flux: Amount of radiant power, or time rate flow of energy, from a source. Units: Joules/sec Symbol: Φ_e
- Irradiance: Density of radiant energy incident on a surface per unit area.

Units: Watts per unit area Symbol: E_e

- Solid Angle: Defines spatial extent, or a three-dimensional angle formed in a cone shape that originates from a point. Measured in steradians. Units: Steradians, sr Symbol: ω
- Radiance: Amount of radiant energy emitted in a unit of solid angle, dependent on viewing angle.

Units: Watts per unit area per steradian Symbol: L_e

Luminous Flux: Time rate flow of luminous power from a source. Can be expressed in photopic or scotopic luminous power. Units: Lumens [lms] Symbol: Φ (Photopic), Φ' (Scotopic)

oo por anno

Illuminance: Density of incident luminous flux on a surface per unit area.

Units: lms/m^2 (Lux) or lms/ft^2 (Footcandles) Symbol: E

Luminous Intensity: Light emitting power of a point source in a particular direction, or the density of luminous flux in space in that direction. Does not depend on distance from source. Units: Candela [Cd]

Symbol: I

Luminance: Local surface density of light emitting power in a particular direction, or the amount of luminous flux per unit solid angle emitted from a surface element in a particular direction.

Units: Cd/m^2 or Cd/ft^2

Symbol: L

- **Reflectance:** Ratio of reflected luminous flux to incident luminous flux on a surface. Φ_{off}/Φ_{on} Symbol: ρ
- **Transmittance:** Ratio of luminous flux that emits from a surface to total incident luminous flux.

Symbol: τ

- Absorptance: Ratio of luminous flux that is absorbed by a material to total incident flux. Symbol: α
- **Photopic:** Describes the adaptation state of the human visual system for luminance values higher than 10 Cd/m^2 , daylight adaptation state.
- Scotopic: Describes the adaptation state of the human visual system for luminance values lower than 0.001 Cd/m^2 , nighttime adaptation state.

A.4 Other Definitions

- **Analog-to-Digital Conversion:** The process in which analog voltage signals are converted to digital values for storage.
- **Blooming:** A photographical aberration that occurs when highly saturated sites of the imaging surface spills excess information onto neighboring sites, resulting in the neighboring sites giving higher readings than it should.
- **Certainty Function:** The derivative of the camera response function, with respect to a logarithmic exposure axis. Used to give higher confidence, or weight, to values in the steep portion of the response function and little or no weight to the values in the flat, extreme ends of the response function.
- Signal-to-Noise Ratio: A numerical value that describes image quality by comparing the magnitude of the signal (the desired information) to the magnitude of the noise (unwanted or false information). Higher SNR values correlate to better quality images.
- **Vignetting:** In photography, it describes the effect of light falloff towards the periphery of the image in a radial direction. More specifically, it is caused by the occlusion of off-axis light rays through an optical device caused by the physical size of the aperture. Vignetting effects increase with aperture size.
- **Point Spread Function:** Describes light scattering in the lens by characterizing the radial effect of light spill from a pixel to its neighbors. It is affected by aperture size, exposure time, and eccentricity.
- Eccentricity: Radial distance from the optical center in an image.
- **Camera Response Functions:** A function, typically nonlinear, that describes the mapping from scene luminance to image pixel values.

Quantum Efficiency: Ratio of electron flux to incident photon flux.

Photometric Center: The central point in the luminous opening of a luminaire.

- **Dynamic Range:** The ratio between the lightest and darkest luminance values in a scene or image.
- **ISO:** Describes the light sensitivity level of a camera, where lower ISO values correspond to lower sensitivity and vice versa.
- Noise: Unwanted or false signals of information. In an image, it takes the form of random brightness or coloration patterns , i.e. image grain, that does not actually exist in the image.
- **Dark Current:** Electric current flowing through the photo-diode without the presence of photons. A form of image noise, because it is a false electrical signal.
- Quantization: Constraining a continuous set of values into a relatively small, discrete set, i.e. binning luminance values that may range from 0 to 100,000 Cd/m^2 into 256 integer values, 0 to 255.
- **Dequantization Uncertainty:** The uncertainty, or lost information, that results from reversing the process of quantization.
- **Correlated Color Temperature:** Describes the color appearance of white light, defined by the temperature in Kelvin of an ideal blackbody radiation emitting light of equal chromaticity. 'Cool' white correlates to high CCTs, like 8000K, and 'warm' white light have lower CCTs, like 3000K.
- Lambertian: Describes light behavior that is perfectly diffuse, meaning that the surface or source reflects or emits light such that luminous intensity diminishes with the cosine of the viewing angle and luminance is constant over viewing angle

Appendix B

Camera Response Functions

B.1 Nikon D5200 with Sigma Fisheye Lens Daylight White Balance

f/3.5 *
$$R: 3.6128 * x^3 - 4.29728 * x^2 + 1.87106 * x - 0.18626$$

*
$$G: 1.77676 * x^3 - 1.26003 * x^2 + 0.502562 * x - 0.01929$$

*
$$B: 1.70503 * x^3 - 0.96641 * x^2 + 0.260452 * x - 0.000924$$

$$\begin{array}{l} {\rm f/5.6} & * \ R: 2.63797 * x^3 - 2.46681 * x^2 + 0.84649 * x - 0.01765 \\ & * \ G: 1.64864 * x^3 - 1.03363 * x^2 + 0.389322 * x - 0.00432 \\ & * \ B: 1.77244 * x^3 - 1.11108 * x^2 + 0.342171 * x - 0.00353 \end{array}$$

$$\begin{array}{ll} {\rm f}/9 & * \ R: 2.65304 * x^3 - 2.44569 * x^2 + 0.8030592 * x - 0.01094 \\ & * \ G: 1.68141 * x^3 - 1.05312 * x^2 + 0.374797 * x - 0.00309 \\ & * \ B: 1.80464 * x^3 - 1.13065 * x^2 + 0.328983 * x - 0.00297 \end{array}$$

$$\begin{array}{ll} {\rm f}/14 & * \ R: 2.46074 * x^3 - 2.22962 * x^2 + 0.78263 * x - 0.01376 \\ & * \ G:: 1.60907 * x^3 - 1.00253 * x^2 + 0.396702 * x - 0.003232 \\ & * \ B: 1.96178 * x^3 - 1.34821 * x^2 + 0.390079 * x - 0.00365 \end{array}$$

$$\begin{array}{l} {\rm f}/22 & * \ R: 121.014*x^5-354.17*x^4+399.305*x^3-215.249*x^2+55.5326*x-5.43326*x^2+36.266*x^5+85.6441*x^4-72.4707*x^3+28.6783*x^2-4.94971*x+0.363962x^2+8:7.22935*x^5-10.1434*x^4+4.83415*x^3-0.98529*x^2+0.068106*x-0.00289x^2+0.068106*x^2+0.00289x^2+0.068106*x^2+0.00289x^2+0.068106*x^2+0.00289x^2+0.068106*x^2+0.068106*x^2+0.068106*x^2+0.068106*x^2+0.068106*x^2+0.068106*x^2+0.068106*x^2+0.068106*x^2+0.068106*x^2+0.068106*x^2+0.068106*x^2+0.068106*x^2+0.068106*x^2+0.068106*x^2+0.068100*x^2+0.06810x^2+0.06810x^2+0.06810x^2+0.06810x^2+0.06810x^2+0.06810x^2+0.06810$$

$$\begin{array}{l} {\rm f}/3.5 & * \ R: 1.9489 * x^3 - 1.40336 * x^2 + 0.470692 * x - 0.00922 \\ \\ & * \ G: 1.49326 * x^3 - 0.82432 * x^2 + 0.332916 * x - 0.00186 \\ \\ & * \ B: 1.55344 * x^3 - 0.88187 * x^2 + 0.329494 * x - 0.00106 \end{array}$$

$$\begin{array}{l} {\rm f/5.6} & * \ R: 1.9003 * x^3 - 1.35832 * x^2 + 0.467647 * x - 0.00963 \\ \\ & * \ G: 1.4891 * x^3 - 0.81777 * x^2 + 0.330558 * x - 0.00189 \\ \\ & * \ B: 1.52805 * x^3 - 0.8457 * x^2 + 0.31848 * x - 0.00083 \end{array}$$

$$\begin{array}{rl} {\rm f/9} & * \ R: 1.89987 * x^3 - 1.33012 * x^2 + 0.438048 * x - 0.0078 \\ & * \ G: 1.56332 * x^3 - 0.88395 * x^2 + 0.322729 * x - 0.00209 \\ & * \ B: 1.65142 * x^3 - 0.97377 * x^2 + 0.324229 * x - 0.00188 \end{array}$$

$$\begin{array}{ll} {\rm f}/14 & * \ R: 2.02588 * x^3 - 1.56794 * x^2 + 0.55653 * x - 0.01447 \\ & * \ G: 1.48384 * x^3 - 0.83646 * x^2 + 0.35446 * x - 0.00184 \\ & * \ B: 1.47989 * x^3 - 0.81719 * x^2 + 0.337311 * x - 0.000008 \\ & {\rm f}/22 & * \ R: 1.97166 * x^3 - 1.47437 * x^2 + 0.514266 * x - 0.01156 \\ & * \ G: 1.48073 * x^3 - 0.82191 * x^2 + 0.342744 * x - 0.00157 \end{array}$$

*
$$B: 1.50183 * x^3 - 0.83217 * x^2 + 0.3305 * x - 0.00016$$

$$\begin{aligned} f/2.2 & * R : 2.62212 * x^3 - 2.45527 * x^2 + 0.865334 * x - 0.03218 \\ & * G : 1.58064 * x^3 - 0.93923 * x^2 + 0.362565 * x - 0.00398 \\ & * B : 1.62315 * x^3 - 0.89947 * x^2 + 0.275252 * x + 0.001074 \end{aligned}$$

$$\begin{array}{l} \texttt{f/3.5} & \ast \ R: 33.5253 \ast x^3 - 76.8377 \ast x^4 + 63.4258 \ast x^3 - 22.2751 \ast x^2 + 3.26953 \ast x - 0.10789 \\ & \ast \ G: -28.2225 \ast x^5 + 57.7966 \ast x^4 - 38.6213 \ast x^3 + 10.7609 \ast x^2 - 0.74575 \ast x - 0.032047 \\ & \ast \ B: 9.98505 \ast x^5 - 16.9544 \ast x^4 + 10.1847 \ast x^3 - 2.45738 \ast x^2 + 0.245407 \ast x - 0.00333 \\ \end{array}$$

f/5.6 *
$$R: 2.05082 * x^2 - 1.36998 * x + 0.319154$$

* $G: 1.4762 * x^2 - 0.58715 * x + 0.110947$
* $B: 1.09906 * x^2 - 0.12726 * x + 0.028196$

f/9 *
$$R: 2.1937 * x^2 - 1.55588 * x + 0.362186$$

* $G: 1.45499 * x^2 - 0.56298 * x + 0.107987$
* $B: 1.06391 * x^2 - 0.09269 * x + 0.028785$

$$\begin{array}{rl} {\rm f}/14 & * \ R: 4.30473 * x^3 - 5.79098 * x^2 + 2.86113 * x - 0.37489 \\ & * \ G: 1.90413 * x^3 - 1.43117 * x^2 + 0.551869 * x + 0.02483 \\ & * \ B: 1.46854 * x^3 - 0.64179 * x^2 + 0.168384 * x + 0.004866 \end{array}$$

$$\begin{aligned} \mathbf{f}/2.2 & * \ R: 1.911 * x^3 - 1.35305 * x^2 + 0.450926 * x - 0.00887 \\ & * \ G: 1.45876 * x^3 - 0.77286 * x^2 + 0.315778 * x - 0.00167 \\ & * \ B: 1.51066 * x^3 - 0.82706 * x^2 + 0.317388 * x - 0.00099 \end{aligned}$$

$$\begin{array}{l} {\rm f/3.5} & * \ R: 1.911*x^3-1.35305*x^2+0.450926*x-0.00887\\ \\ & * \ G: 1.47307*x^3-0.7964*x^2+0.324964*x-0.00164\\ \\ & * \ B: 1.53137*x^3-0.85491*x^2+0.324363*x-0.00083 \end{array}$$

$$\begin{array}{l} {\rm f}/5.6 & * \ R: 1.911*x^3-1.35305*x^2+0.450926*x-0.00887 \\ \\ & * \ G: 1.46936*x^3-0.80605*x^2+0.33809*x-0.0014 \\ \\ & * \ B: 1.57698*x^3-0.91985*x^2+0.343684*x-0.00082 \end{array}$$

$$f/9 \quad * R : 1.911 * x^{3} - 1.35305 * x^{2} + 0.450926 * x - 0.00887$$

$$* G : 1.46373 * x^{3} - 0.79854 * x^{2} + 0.336269 * x - 0.00146$$

$$* B : 1.56623 * x^{3} - 0.91478 * x^{2} + 0.349887 * x - 0.00134$$

$$f/14 \quad + R : 1.011 + x^{3} - 1.25205 + x^{2} + 0.450026 + x - 0.00887$$

$$f/14 \quad * R : 1.911 * x^{3} - 1.35305 * x^{2} + 0.450926 * x - 0.00887$$
$$* G : 1.47236 * x^{3} - 0.82681 * x^{2} + 0.355827 * x - 0.00138$$
$$* B : 1.54178 * x^{3} - 0.89209 * x^{2} + 0.350662 * x - 0.00035$$

$$f/3.8 \quad * R: 3.05227 * x^3 - 3.37454 * x^2 + 1.45373 * x - 0.13147$$

$$* G: 1.67935 * x^3 - 1.11032 * x^2 + 0.443439 * x - 0.01247$$

$$* B: 1.48427 * x^3 - 0.69522 * x^2 + 0.206646 * x - 0.004298$$

$$f/5.6 \quad * R: 61.0782 * x^5 - 160.404 * x^4 + 150.012 * x^3 - 72.2011 * x^2 + 15.7204 * x - 1.21685$$

$$\begin{array}{l} 1/5.6 & * \ R: 61.0782 * x^{5} - 160.404 * x^{4} + 159.013 * x^{3} - 73.2011 * x^{2} + 15.7304 * x - 1.21685 \\ & * \ G: 1130.18 * x^{5} - 2525.11 * x^{4} + 2029.91 * x^{3} - 743.51 * x^{2} + 116.814 * x - 7.28488 \\ & * \ B: 7.06295 * x^{5} - 9.79305 * x^{4} + 4.58342 * x^{3} - 0.9076 * x^{2} + 0.05691 * x - 0.00263 \end{array}$$

$$\begin{array}{l} {\rm f/9} & * \ R: 8.9173 * x^4 - 15.6545 * x^3 + 9.65609 * x^2 - 2.08016 * x + 0.16129 \\ \\ & * \ G: 4.24861 * x^4 - 5.88934 * x^3 + 3.09662 * x^2 - 0.4991 * x + 0.043197 \\ \\ & * \ B: 8.35229 * x^4 - 10.418 * x^3 + 3.55045 * x^2 - 0.48762 * x + 0.002922 \end{array}$$

$$\begin{array}{ll} {\rm f}/14 & * \ R: 2.18802 * x^2 - 1.52094 * x + 0.332921 \\ \\ & * \ G: 1.50003 * x^2 - 0.61245 * x + 0.112425 \\ \\ & * \ B: 1.07001 * x^2 - 0.09936 * x + 0.029352 \end{array}$$

$$\begin{array}{l} {\rm f}/22 & * \ R: 96.5261*x^5-272.122*x^4+294.119*x^3-151.159*x^2+37.0592*x-3.42325 \\ & * \ G: -242.254*x^5+575.538*x^4-500.453*x^3+202.076*x^2-36.4756*x+2.56885 \\ & * \ B: 7.71569*x^5-10.9939*x^4+5.30851*x^3-1.10261*x^2+0.075899*x-0.00356 \end{array}$$

B.6 Nikon D5200 with Nikkor Zoom Lens, 24mm Fluorescent White Balance

$$\begin{array}{rl} {\rm f}/3.8 & * \ R: 1.87726 * x^3 - 1.32763 * x^2 + 0.459332 * x - 0.00896 \\ & * \ G:: 1.47128 * x^3 - 0.80064 * x^2 + 0.331091 * x - 0.00173 \\ & * \ B: 1.60008 * x^3 - 0.9411 * x^2 + 0.342383 * x - 0.00136 \end{array}$$

$$\begin{array}{l} {\rm f/5.6} & * \ R: 1.94027 * x^3 - 1.38963 * x^2 + 0.458274 * x - 0.00892 \\ \\ & * \ G: 1.49574 * x^3 - 0.81629 * x^2 + 0.322125 * x - 0.00157 \\ \\ & * \ B: 1.55518 * x^3 - 0.87926 * x^2 + 0.325117 * x - 0.00104 \end{array}$$

$$\begin{array}{ll} {\rm f}/9 & * \ R: 1.96328 * x^3 - 1.4211 * x^2 + 0.467974 * x - 0.01015 \\ \\ & * \ G: 1.49544 * x^3 - 0.81196 * x^2 + 0.318202 * x - 0.00168 \\ \\ & * \ B: 1.54322 * x^3 - 0.8589 * x^2 + 0.316611 * x - 0.00093 \end{array}$$

$$\begin{array}{ll} {\rm f}/14 & * \ R: 2.0034 * x^3 - 1.49655 * x^2 + 0.503447 * x - 0.01029 \\ & * \ G: 1.49708 * x^3 - 0.84869 * x^2 + 0.353361 * x - 0.00175 \\ & * \ B: 1.55249 * x^3 - 0.90875 * x^2 + 0.357349 * x - 0.00109 \\ & {\rm f}/22 & * \ R: 1.91974 * x^3 - 1.40324 * x^2 + 0.493565 * x - 0.01006 \\ \end{array}$$

*
$$G: 1.46983 * x^3 - 0.81253 * x^2 + 0.344138 * x - 0.00144$$

*
$$B: 1.50961 * x^3 - 0.84842 * x^2 + 0.339276 * x - 0.00047$$

$$\begin{array}{ll} {\rm f}/3.5 & * \ R: 2.13837 * x^2 - 1.46661 * x + 0.328233 \\ & * \ G: 1.46606 * x^2 - 0.57338 * x + 0.10732 \\ & * \ B: 1.07434 * x^2 - 0.10251 * x + 0.028177 \\ & {\rm f}/5.6 & * \ R: 4.11332 * x^3 - 5.25407 * x^2 + 2.42743 * x - 0.28668 \\ & * \ G: 1.80665 * x^3 - 1.30932 * x^2 + 0.525739 * x - 0.02308 \end{array}$$

*
$$B: 1.58347 * x^3 - 0.83147 * x^2 + 0.245857 * x - 0.002148$$

$$\begin{array}{ll} {\rm f/9} & * \ R: 4.48128 * x^3 - 5.79125 * x^2 + 2.613 * x - 0.30304 \\ \\ & * \ G: 1.99685 * x^3 - 1.55203 * x^2 + 0.584971 * x - 0.02978 \\ \\ & * \ B: 1.4998 * x^3 - 0.70242 * x^2 + 0.199256 * x - 0.003364 \end{array}$$

$$f/14 \quad * R: 4.91767 * x^{3} - 6.92269 * x^{2} + 3.48253 * x - 0.47752$$

$$* G: 1.82519 * x^{3} - 1.37655 * x^{2} + 0.578975 * x - 0.02761$$

$$* B: 1.6954 * x^{3} - 0.9415 * x^{2} + 0.244581 * x - 0.00152$$

$$f/22 \quad * R: 5.61248 * x^{3} - 8.22711 * x^{2} + 4.2271 * x - 0.61246$$

$$\begin{array}{l} 1/22 & * \ R: 5.61348 * x^3 - 8.22711 * x^2 + 4.2271 * x - 0.61346 \\ & * \ G: 1.81678 * x^3 - 1.33869 * x^2 + 0.546568 * x - 0.02466 \\ & * \ B: 1.65132 * x^3 - 0.88664 * x^2 + 0.233134 * x - 0.002194 \end{array}$$

B.8 Nikon D5200 with Nikkor Zoom Lens, 18mm Fluorescent White Balance

$$\begin{array}{ll} {\rm f}/3.5 & * \ R: 2.13582 * x^3 - 1.68962 * x^2 + 0.566861 * x - 0.01306 \\ \\ & * \ G: 1.51646 * x^3 - 0.86793 * x^2 + 0.353145 * x - 0.00178 \\ \\ & * \ B: 1.51793 * x^3 - 0.86229 * x^2 + 0.345047 * x - 0.00069 \end{array}$$

$$\begin{array}{rl} {\rm f}/5.6 & * \ R: 2.1131 * x^3 - 1.66868 * x^2 + 0.569682 * x - 0.01409 \\ \\ & * \ G: 1.51094 * x^3 - 0.85981 * x^2 + 0.350626 * x - 0.00176 \\ \\ & * \ B: 1.48889 * x^3 - 0.821 * x^2 + 0.332433 * x - 0.00033 \end{array}$$

$$\begin{array}{ll} {\rm f}/9 & * \ R: 2.02268 * x^3 - 1.4928 * x^2 + 0.481025 * x - 0.01091 \\ \\ & * \ G: 1.54205 * x^3 - 0.86494 * x^2 + 0.325058 * x - 0.00217 \\ \\ & * \ B: 1.56098 * x^3 - 0.87484 * x^2 + 0.315066 * x - 0.00012 \end{array}$$

$$\begin{array}{ll} {\rm f}/14 & * \ R: 2.17406 * x^3 - 1.78453 * x^2 + 0.627313 * x - 0.01684 \\ & * \ G: 1.51145 * x^3 - 0.8793 * x^2 + 0.369352 * x - 0.00151 \\ & * \ B: 1.4884 * x^3 - 0.83313 * x^2 + 0.344077 * x - 0.000651 \end{array}$$

$$\begin{array}{l} {\rm f}/22 & * \ R: 5.03978 * x^4 - 7.15419 * x^3 + 3.44202 * x^2 - 0.35203 * x + 0.024418 \\ & * \ G: 1.80978 * x^4 - 1.43339 * x^3 + 0.473913 * x^2 - 0.146438 * x + 0.003258 \\ & * \ B: -8.52857 * x^4 + 14.6124 * x^3 - 6.24263 * x^2 + 1.16793 * x - 0.0091 \end{array}$$

Appendix C

Vignetting Functions

C.1 MATLAB Vignetting Correction Script

```
%Vignetting correction takes a raw HDR image and a vignetting correction
%function and applies the correction to output a new HDR image corrected
%for vignetting.
Filename='Example HDR Image';
type=2;
%Please specify if the image is a fisheye or not. 1= fisheye or 2=rectangular
image are the accepted options.
image=hdrread(sprintf('%s.hdr',Filename));
%Reads in the HDR image to a 3D matrix named 'image'.
[rows, columns, pages] = size (image);
%Gets dimensions of HDR image matrix.
VF='Example Vignetting Function.txt';
[coefs]=textread(VF, '%f');
%Read in vignetting function text file and assigns the values to variables
size=length(coefs);
if type==1
    cp=[3000,2006];
    maxRad=1575;
else
    cp = [round(columns/2),round(rows/2)];
    maxRad = sqrt(cp(1)^2 + cp(2)^2);
end
VCimage = zeros(rows,columns,pages);
for i=1:rows
    for j=1:columns
       x=j-cp(1);
        y=i-cp(2);
        r=sqrt(x^2+y^2)/maxRad;
        %Radius of pixel.
        V=0;
        Starts calculation for correction factor.
        for k=1:size
            V=V+coefs(k) *r^ (size-k);
            Calculates the vignetting loss for calculated radius.
        end
        CF=1/V;
        %Calculates the correction factor, which is the reciprocal of the
        %vignetting loss.
        VCimage(i,j,1)=CF*image(i,j,1);
        VCimage(i,j,2)=CF*image(i,j,2);
        VCimage(i,j,3)=CF*image(i,j,3);
        &Applies correction factor to pixel.
        if VCimage(i,j,1)<0||VCimage(i,j,2)<0||VCimage(i,j,3)<0
            VCimage(i,j,:)=0;
        end
    end
end
new=sprintf('%s%s',Filename,'-VCorrected.hdr');
hdrwrite(VCimage, sprintf('%s%s',Filename,'-VCorrected.hdr'))
```

C.2 Vignetting Error Charts



C.2.1 Fisheye Lens







































Appendix D

Software Evaluation Results
D.1 Luminance HDR Results

D.1.1 RAW-Based HDR Images, Classroom with Low Fluorescent Light + Daylight

| | | Trian | gular | Plateau | | Gaussian | |
|--------|-----------|---------|---------|---------|---------|----------|---------|
| | Avg. | | Ĩ | | | | |
| Target | Luminance | Linear | Gamma | Linear | Gamma | Linear | Gamma |
| 23 | 3.64 | 0.22% | 0.05% | 0.05% | 0.11% | 32.22% | 31.89% |
| 6 | 4.77 | 81.49% | 81.93% | 81.94% | 77.99% | 89.75% | 90.26% |
| 5 | 4.89 | 81.22% | 81.60% | 81.60% | 77.98% | 89.49% | 90.01% |
| 14 | 5.39 | 27.56% | 28.75% | 28.75% | 24.28% | 54.07% | 54.93% |
| 20 | 6.47 | 144.38% | 168.71% | 168.92% | 103.27% | 138.04% | 127.20% |
| 7 | 12.67 | 157.05% | 182.19% | 181.98% | 108.03% | 147.84% | 135.27% |
| 19 | 17.18 | 34.10% | 36.62% | 36.62% | 23.76% | 0.95% | 0.24% |
| 1 | 18.25 | 20.93% | 21.98% | 21.98% | 15.21% | 49.81% | 50.65% |
| 22 | 21.79 | 55.15% | 55.56% | 55.76% | 47.11% | 72.41% | 72.96% |
| 8 | 24.97 | 69.72% | 70.47% | 70.50% | 63.72% | 82.53% | 83.14% |
| 4 | 26.89 | 78.87% | 79.37% | 79.36% | 75.05% | 88.07% | 88.62% |
| 13 | 29.28 | 81.80% | 82.22% | 82.23% | 78.36% | 89.91% | 90.46% |
| 18 | 34.66 | 31.83% | 32.99% | 33.03% | 26.50% | 57.31% | 58.40% |
| 9 | 45.41 | 135.02% | 155.26% | 154.89% | 93.81% | 121.29% | 112.56% |
| 17 | 63.56 | 83.93% | 84.36% | 84.37% | 80.35% | 91.32% | 91.78% |
| 10 | 72.40 | 77.94% | 78.41% | 78.40% | 74.16% | 87.39% | 87.94% |
| 21 | 95.49 | 66.41% | 67.01% | 66.99% | 61.22% | 80.30% | 80.76% |
| 16 | 101.70 | 41.03% | 42.01% | 42.04% | 33.56% | 64.03% | 64.98% |
| 11 | 107.63 | 6.42% | 8.17% | 6.66% | 4.50% | 27.23% | 27.40% |
| 3 | 121.17 | 110.40% | 125.39% | 125.08% | 75.95% | 86.61% | 84.14% |
| 2 | 125.43 | 75.55% | 76.04% | 76.04% | 70.52% | 86.03% | 86.47% |
| 12 | 127.00 | 7.15% | 7.47% | 7.47% | 4.67% | 39.50% | 39.55% |
| 15 | 150.47 | 212.65% | 252.80% | 253.07% | 147.75% | 229.97% | 194.78% |
| | Average: | 73.08% | 79.10% | 79.03% | 59.47% | 82.87% | 80.19% |

D.1.2 JPEG-Based HDR Images, Classroom with Low Fluorescent Light + Daylight

| | | Trian | gular | Plateau | | Gaussian | |
|--------|-----------|---------|---------|---------|---------|----------|---------|
| | | | | | | | |
| | Avg. | | | | | | |
| Target | Luminance | Linear | Gamma | Linear | Gamma | Linear | Gamma |
| 29 | 2.10 | 0.19% | 0.11% | 0.19% | 0.11% | 0.11% | 0.03% |
| 23 | 2.61 | 87.84% | 86.92% | 83.96% | 83.58% | 87.83% | 87.78% |
| 5 | 2.67 | 82.50% | 81.84% | 78.52% | 77.94% | 82.46% | 82.44% |
| 14 | 3.44 | 19.32% | 19.64% | 17.72% | 16.50% | 19.38% | 19.38% |
| 20 | 4.50 | 328.46% | 249.06% | 197.38% | 124.72% | 338.95% | 339.70% |
| 28 | 5.35 | 85.16% | 84.31% | 81.47% | 80.94% | 85.19% | 85.18% |
| 27 | 11.97 | 33.75% | 33.80% | 32.41% | 31.64% | 34.11% | 34.01% |
| 19 | 12.37 | 256.34% | 202.61% | 152.95% | 105.61% | 259.83% | 260.41% |
| 1 | 12.41 | 88.06% | 87.15% | 84.33% | 83.70% | 88.06% | 88.05% |
| 4 | 15.62 | 82.40% | 81.75% | 78.61% | 78.02% | 82.36% | 82.36% |
| 22 | 15.76 | 72.08% | 71.11% | 68.23% | 67.87% | 72.05% | 72.01% |
| 13 | 19.37 | 48.41% | 47.54% | 45.29% | 44.38% | 48.41% | 48.41% |
| 26 | 23.37 | 0.84% | 1.89% | 3.34% | 4.96% | 1.24% | 1.24% |
| 18 | 24.36 | 174.24% | 146.71% | 112.51% | 82.98% | 174.69% | 175.35% |
| 25 | 40.69 | 80.92% | 80.20% | 76.76% | 76.08% | 80.94% | 80.91% |
| 17 | 44.51 | 14.55% | 15.18% | 12.52% | 11.06% | 14.80% | 14.68% |
| 24 | 63.36 | 359.59% | 263.68% | 206.91% | 124.81% | 371.87% | 373.40% |
| 21 | 69.36 | 78.82% | 78.41% | 74.81% | 73.99% | 78.92% | 78.92% |
| 3 | 71.08 | 66.72% | 66.60% | 63.77% | 63.28% | 66.95% | 66.95% |
| 16 | 71.47 | 41.89% | 41.55% | 39.71% | 39.11% | 42.23% | 42.15% |
| 12 | 85.22 | 13.06% | 14.48% | 12.90% | 13.98% | 12.73% | 12.81% |
| 2 | 101.00 | 154.42% | 130.14% | 100.44% | 74.84% | 153.30% | 153.49% |
| 15 | 106.37 | 428.93% | 306.36% | 251.51% | 138.47% | 457.55% | 457.55% |
| | Average: | 112.98% | 95.26% | 81.57% | 65.16% | 115.39% | 115.53% |









D.2 Picturenaut Results

| Target | Avg | Picturenaut | Picturenaut | Picturenaut | Picturenaut |
|--------|----------|-------------|-------------|-------------|-------------|
| | | JPEG | RAW | JPEG Mid- | JPEG |
| | | | | Weighting | Gamma |
| | | | | | Curve |
| 1 | 18.25 | 0.22% | 0.33% | 0.16% | 0.38% |
| 2 | 125.4333 | 22.22% | 24.32% | 21.35% | 23.83% |
| 3 | 121.1667 | 21.08% | 23.46% | 21.10% | 22.41% |
| 4 | 26.89 | 3.46% | 2.98% | 4.09% | 6.14% |
| 5 | 4.89 | 9.00% | 10.63% | 8.18% | 4.50% |
| 6 | 4.773333 | 4.96% | 7.89% | 4.75% | 1.19% |
| 7 | 12.67 | 3.63% | 1.42% | 3.39% | 1.82% |
| 8 | 24.96667 | 5.95% | 4.15% | 5.67% | 7.60% |
| 9 | 45.41333 | 9.89% | 6.17% | 9.81% | 13.35% |
| 10 | 72.4 | 16.37% | 15.36% | 15.95% | 19.30% |
| 11 | 107.6333 | 21.00% | 22.11% | 20.66% | 22.88% |
| 12 | 127 | 23.67% | 25.81% | 23.04% | 24.83% |
| 13 | 29.28 | 6.90% | 4.99% | 7.04% | 9.39% |
| 14 | 5.386667 | 5.07% | 7.49% | 5.07% | 2.29% |
| 15 | 150.4667 | 25.62% | 28.26% | 25.24% | 26.87% |
| 16 | 101.7 | 20.92% | 20.29% | 20.47% | 23.19% |
| 17 | 63.56 | 15.31% | 13.85% | 15.17% | 18.93% |
| 18 | 34.66333 | 9.15% | 6.13% | 9.18% | 11.75% |
| 19 | 17.17667 | 4.06% | 2.43% | 4.06% | 3.47% |
| 20 | 6.473333 | 2.42% | 4.43% | 2.42% | 1.80% |
| 21 | 95.49 | 19.26% | 19.43% | 19.04% | 22.60% |
| 22 | 21.78667 | 4.90% | 3.66% | 4.94% | 5.45% |
| 23 | 3.636667 | 6.42% | 9.17% | 6.14% | 4.86% |
| | Average | 11.37% | 11.51% | 11.17% | 12.12% |

D.2.1 Classroom, Fluorescent + Daylight

| Target | Avg | Picturenaut JPEG | Picturenaut RAW | Picturenaut JPEG Mid- Weighting | Picturenaut JPEG Gamma Curve |
|--------|----------|---------------------|--------------------|---------------------------------------|---------------------------------------|
| 1 | 15.12 | 0.13% | 0.07% | 0.00% | 0.13% |
| 2 | 103.77 | 24.55% | 27.28% | 23.91% | 28.00% |
| 3 | 89.39 | 23.02% | 25.59% | 22.53% | 27.18% |
| 4 | 19.57 | 4.84% | 2.74% | 4.79% | 6.12% |
| 5 | 3.41 | 13.01% | 10.67% | 13.01% | 17.42% |
| 6 | 4.43 | 6.17% | 5.50% | 6.40% | 12.05% |
| 7 | 11.59 | 1.52% | 1.01% | 1.35% | 2.99% |
| 8 | 21.81 | 2.43% | 0.78% | 2.57% | 5.46% |
| 9 | 39.37 | 9.04% | 8.41% | 8.92% | 13.79% |
| 10 | 61.04 | 14.62% | 16.45% | 14.19% | 19.19% |
| 11 | 93.36 | 22.18% | 24.28% | 21.58% | 25.80% |
| 12 | 103.43 | 25.46% | 28.04% | 25.00% | 28.80% |
| 13 | 23.71 | 6.47% | 4.10% | 6.68% | 9.84% |
| 14 | 4.27 | 8.52% | 6.87% | 8.75% | 14.38% |
| 15 | 123.70 | 27.72% | 31.31% | 27.18% | 30.36% |
| 16 | 83.42 | 22.71% | 24.89% | 22.06% | 26.62% |
| 17 | 51.90 | 16.42% | 16.05% | 15.99% | 21.06% |
| 18 | 28.34 | 8.92% | 6.23% | 8.88% | 12.30% |
| 19 | 14.21 | 0.87% | 1.43% | 1.01% | 0.30% |
| 20 | 5.25 | 7.55% | 5.46% | 7.93% | 12.69% |
| 21 | 79.30 | 21.20% | 23.98% | 20.86% | 25.45% |
| 22 | 18.23 | 3.99% | 3.33% | 3.93% | 4.43% |
| 23 | 3.04 | 14.27% | 10.65% | 14.60% | 16.58% |
| 24 | 78.25 | 21.89% | 24.24% | 21.40% | 25.88% |
| 25 | 50.36 | 15.58% | 15.76% | 15.30% | 19.79% |
| 26 | 28.95 | 8.04% | 5.52% | 7.90% | 11.70% |
| 27 | 15.02 | 2.75% | 3.88% | 2.68% | 1.95% |
| 28 | 6.78 | 3.24% | 0.44% | 3.24% | 7.96% |
| 29 | 2.71 | 13.79% | 10.10% | 13.79% | 15.64% |
| | Average: | 12.10% | 11.90% | 11.95% | 15.31% |

D.2.2 Classroom High Fluorescent

| Target | Avg | Picturenaut JPEG | Picturenaut RAW | Picturenaut JPEG Mid- Weighting | Picturenaut JPEG Gamma Curve |
|--------|----------|---------------------|--------------------|---------------------------------------|---------------------------------------|
| 1 | 12.41333 | 0.05% | 0.43% | 0.30% | 0.03% |
| 2 | 101 | 25.27% | 29.30% | 25.32% | 30.51% |
| 3 | 71.08 | 21.71% | 24.62% | 21.05% | 27.25% |
| 4 | 15.61667 | 4.40% | 4.14% | 4.59% | 5.61% |
| 5 | 2.67 | 12.36% | 10.86% | 12.73% | 13.11% |
| 12 | 85.21667 | 24.70% | 28.17% | 24.23% | 30.01% |
| 13 | 19.36667 | 6.64% | 4.68% | 6.64% | 9.33% |
| 14 | 3.443333 | 9.49% | 7.45% | 9.78% | 13.26% |
| 15 | 106.3667 | 27.83% | 31.04% | 27.07% | 32.54% |
| 16 | 71.47 | 22.54% | 25.28% | 22.01% | 28.00% |
| 17 | 44.51333 | 16.05% | 16.34% | 15.82% | 21.84% |
| 18 | 24.36333 | 9.08% | 6.38% | 8.84% | 13.31% |
| 19 | 12.36667 | 1.59% | 2.40% | 1.51% | 1.43% |
| 20 | 4.503333 | 4.81% | 4.37% | 5.26% | 10.58% |
| 21 | 69.35667 | 21.45% | 24.35% | 20.90% | 27.39% |
| 22 | 15.76333 | 4.91% | 5.16% | 4.97% | 5.92% |
| 23 | 2.606667 | 10.10% | 8.18% | 10.49% | 9.72% |
| 24 | 63.36333 | 20.60% | 21.86% | 20.14% | 26.17% |
| 25 | 40.69 | 15.07% | 14.97% | 14.72% | 20.30% |
| 26 | 23.37 | 8.39% | 5.78% | 8.09% | 11.17% |
| 27 | 11.96667 | 3.15% | 4.07% | 3.23% | 1.89% |
| 28 | 5.353333 | 1.25% | 0.50% | 1.43% | 6.85% |
| 29 | 2.096667 | 12.56% | 10.65% | 12.56% | 8.27% |
| | Average: | 12.35% | 12.65% | 12.25% | 15.41% |

D.2.3 Classroom Low Fluorescent

| Target | Avg | Picturenaut | Picturenaut | Picturenaut | Picturenaut |
|--------|----------|-------------|-------------|-------------|-------------|
| | | JPEG | RAW | JPEG Mid- | JPEG |
| | | | | Weighting | Gamma |
| | | | | | Curve |
| 1 | 3.033333 | 0.11% | 0.44% | 0.11% | 0.55% |
| 2 | 0.453333 | 12.50% | 16.91% | 12.50% | 16.18% |
| 3 | 0.976667 | 3.41% | 7.51% | 3.41% | 3.75% |
| 4 | 2.176667 | 2.14% | 0.15% | 2.60% | 3.52% |
| 5 | 4.086667 | 5.30% | 3.59% | 5.30% | 5.79% |
| 6 | 6.623333 | 8.96% | 4.88% | 8.96% | 9.86% |
| 7 | 9.89 | 13.04% | 3.54% | 13.04% | 14.76% |
| 8 | 1.503333 | 8.43% | 9.09% | 9.09% | 5.10% |
| 9 | 6.886667 | 4.16% | 0.68% | 3.58% | 4.89% |
| 10 | 29.40333 | 21.95% | 11.71% | 21.74% | 24.12% |
| | Average: | 8.00% | 5.85% | 8.03% | 8.85% |

D.2.4 Classroom Window

D.2.5 Spheres Far

| Target | Avg | Picturenaut JPEG | Picturenaut RAW | Picturenaut JPEG Mid- | Picturenaut JPEG |
|--------|----------|---------------------|--------------------|--------------------------|---------------------|
| | | | | Weighting | Gamma |
| | | | | | Curve |
| 1a | 126900 | 82.60% | 74.68% | 82.65% | 98.45% |
| 1b | 5653 | 33.30% | 41.35% | 33.43% | 69.31% |
| 2a | 31117.5 | 60.73% | 58.31% | 60.46% | 93.25% |
| 2b | 1373.75 | 23.47% | 30.19% | 23.34% | 38.66% |
| 3a | 6422.75 | 27.51% | 35.04% | 26.83% | 69.74% |
| 3b | 337.6 | 8.02% | 10.53% | 8.91% | 13.54% |
| 4a | 1677.75 | 18.20% | 25.82% | 96.94% | 38.53% |
| 4b | 112.4 | 18.05% | 16.14% | 18.59% | 17.01% |
| 5a | 187 | 0.16% | 0.71% | 0.24% | 0.71% |
| 5b | 11.34 | 26.01% | 20.11% | 23.28% | 48.06% |
| | Average: | 29.80% | 31.29% | 37.46% | 48.73% |

| Target | Avg | Picturenaut | Picturenaut | Picturenaut | Picturenaut |
|--------|----------|-------------|-------------|-------------|-------------|
| | | JPEG | RAW | JPEG Mid- | JPEG |
| | | | | Weighting | Gamma |
| | | | | | Curve |
| 1a | 127066.7 | 82.40% | 74.07% | 82.45% | 84.26% |
| 1b | 5721.667 | 35.51% | 40.45% | 34.50% | 39.04% |
| 2a | 31410 | 61.58% | 58.20% | 61.04% | 61.10% |
| 2b | 1383.333 | 23.75% | 29.13% | 23.07% | 30.09% |
| 3a | 6939.333 | 35.60% | 40.54% | 34.23% | 38.29% |
| 3b | 357.0667 | 12.75% | 12.82% | 11.46% | 14.58% |
| 4a | 1759 | 22.28% | 28.24% | 21.32% | 28.00% |
| 4b | 104.75 | 10.31% | 5.70% | 9.89% | 7.52% |
| 5a | 189.0667 | 0.04% | 0.27% | 0.23% | 0.46% |
| 5b | 10.71667 | 34.37% | 31.85% | 32.32% | 63.86% |
| | Average: | 31.86% | 32.13% | 31.05% | 36.72% |

D.3 hdrgen, raw2hdr, and Bracket Results

| Target | Avg. | hdrgen | raw2hdr | Bracket | Bracket, |
|--------|--------|--------|---------|---------|----------|
| | Lum. | | | | New |
| | | | | | CRF |
| 1 | 18.25 | 0.44% | 0.00% | 0.27% | 0.16% |
| 2 | 125.43 | 8.75% | 3.93% | 15.34% | 7.45% |
| 3 | 121.17 | 7.49% | 5.22% | 13.56% | 6.23% |
| 4 | 26.89 | 2.68% | 1.49% | 5.76% | 4.54% |
| 5 | 4.89 | 3.27% | 6.75% | 3.48% | 6.75% |
| 6 | 4.77 | 1.33% | 6.84% | 10.13% | 3.49% |
| 7 | 12.67 | 3.24% | 0.71% | 7.42% | 5.13% |
| 8 | 24.97 | 5.15% | 0.05% | 8.40% | 6.19% |
| 9 | 45.41 | 5.42% | 1.33% | 7.56% | 4.98% |
| 10 | 72.40 | 7.93% | 3.15% | 12.07% | 7.33% |
| 11 | 107.63 | 9.32% | 3.25% | 14.79% | 7.57% |
| 12 | 127.00 | 10.53% | 5.81% | 16.98% | 8.70% |
| 13 | 29.28 | 5.26% | 2.12% | 8.54% | 6.73% |
| 14 | 5.39 | 0.25% | 4.89% | 6.99% | 3.59% |
| 15 | 150.47 | 10.42% | 4.70% | 17.72% | 9.75% |
| 16 | 101.70 | 9.82% | 4.09% | 14.32% | 7.78% |
| 17 | 63.56 | 8.68% | 3.85% | 13.45% | 8.21% |
| 18 | 34.66 | 6.04% | 2.20% | 9.65% | 8.09% |
| 19 | 17.18 | 3.30% | 0.19% | 4.99% | 4.29% |
| 20 | 6.47 | 0.51% | 4.12% | 8.86% | 0.72% |
| 21 | 95.49 | 9.77% | 3.42% | 13.74% | 7.79% |
| 22 | 21.79 | 4.71% | 0.29% | 6.46% | 5.81% |
| 23 | 3.64 | 1.83% | 8.07% | 12.83% | 5.59% |

D.3.1 Classroom, Fluorescent + Daylight

| Target | Avg. Lum | hdrgen | raw2hdr | Bracket | Bracket, New CRF |
|--------|-------------|--------|---------|---------|------------------------|
| 1 | 15.12 | 0.13% | 0.13% | 0.07% | 0.20% |
| 2 | 103.77 | 1.44% | 5.18% | 4.80% | 5.04% |
| 3 | 89.39 | 2.96% | 4.17% | 6.74% | 6.54% |
| 4 | 19.57 | 3.10% | 0.65% | 2.33% | 1.65% |
| 5 | 3.41 | 9.78% | 9.20% | 4.50% | 10.37% |
| 6 | 4.43 | 4.37% | 3.92% | 0.75% | 4.14% |
| 7 | 11.59 | 0.66% | 1.44% | 1.35% | 0.81% |
| 8 | 21.81 | 0.46% | 1.93% | 0.87% | 0.55% |
| 9 | 39.37 | 0.51% | 1.35% | 0.10% | 1.63% |
| 10 | 61.04 | 1.32% | 0.86% | 2.27% | 1.94% |
| 11 | 93.36 | 0.72% | 1.95% | 3.48% | 3.10% |
| 12 | 103.43 | 3.65% | 4.44% | 7.19% | 6.82% |
| 13 | 23.71 | 4.19% | 0.98% | 4.44% | 5.33% |
| 14 | 4.27 | 5.70% | 7.58% | 1.48% | 7.11% |
| 15 | 123.70 | 3.44% | 5.31% | 8.26% | 7.53% |
| 16 | 83.42 | 3.43% | 3.83% | 6.70% | 7.02% |
| 17 | 51.90 | 4.74% | 3.83% | 5.47% | 6.69% |
| 18 | 28.34 | 4.72% | 2.32% | 5.49% | 6.31% |
| 19 | 14.21 | 1.29% | 0.82% | 1.36% | 1.92% |
| 20 | 5.25 | 6.22% | 6.41% | 3.17% | 5.27% |
| 21 | 79.30 | 3.85% | 3.09% | 8.01% | 7.87% |
| 22 | 18.23 | 3.82% | 0.53% | 3.82% | 4.15% |
| 23 | 3.04 | 9.33% | 10.98% | 0.11% | 10.98% |
| 24 | 78.25 | 5.22% | 3.56% | 9.02% | 8.94% |
| 25 | 50.36 | 5.20% | 2.91% | 7.84% | 7.72% |
| 26 | 28.95 | 5.00% | 1.37% | 5.83% | 6.41% |
| 27 | 15.02 | 3.82% | 0.42% | 5.15% | 4.48% |
| 28 | 6.78 | 1.77% | 4.72% | 2.36% | 1.18% |
| 29 | 2.71 | 8.62% | 12.32% | 2.09% | 10.47% |

D.3.2 Classroom High Fluorescent

| Target | Avg. Lum. | hdrgen | raw2hdr | Bracket | Bracket, New CRF |
|--------|--------------|--------|---------|---------|------------------------|
| 1 | 12.41333 | 0.03% | 0.05% | 0.21% | 0.35% |
| 2 | 101 | 4.46% | 5.15% | 5.92% | 4.05% |
| 3 | 71.08 | 5.01% | 3.94% | 8.38% | 6.18% |
| 4 | 15.61667 | 4.08% | 2.22% | 4.65% | 4.72% |
| 5 | 2.67 | 7.87% | 9.36% | 0.75% | 10.11% |
| 12 | 85.21667 | 5.93% | 4.81% | 8.88% | 6.37% |
| 13 | 19.36667 | 4.37% | 2.00% | 5.04% | 5.77% |
| 14 | 3.443333 | 6.58% | 7.45% | 2.23% | 7.45% |
| 15 | 106.3667 | 7.05% | 6.01% | 9.50% | 7.09% |
| 16 | 71.47 | 6.35% | 5.18% | 9.33% | 7.04% |
| 17 | 44.51333 | 5.94% | 4.61% | 7.67% | 6.14% |
| 18 | 24.36333 | 5.47% | 3.17% | 5.97% | 5.02% |
| 19 | 12.36667 | 1.83% | 0.11% | 2.16% | 2.24% |
| 20 | 4.503333 | 3.48% | 4.81% | 0.30% | 3.48% |
| 21 | 69.35667 | 6.69% | 5.32% | 10.29% | 7.42% |
| 22 | 15.76333 | 5.10% | 2.18% | 5.92% | 5.22% |
| 23 | 2.606667 | 4.35% | 7.42% | 3.71% | 7.80% |
| 24 | 63.36333 | 6.93% | 4.16% | 9.62% | 7.63% |
| 25 | 40.69 | 6.29% | 3.24% | 9.68% | 7.96% |
| 26 | 23.37 | 5.39% | 2.48% | 6.89% | 6.08% |
| 27 | 11.96667 | 3.15% | 0.72% | 4.23% | 3.23% |
| 28 | 5.353333 | 0.12% | 2.74% | 4.92% | 0.31% |
| 29 | 2.096667 | 5.41% | 11.13% | 3.66% | 10.65% |

D.3.3 Classroom Low Fluorescent

| Target | Avg. | hdrgen | raw2hdr | Bracket | Bracket, |
|--------|----------|--------|---------|---------|----------|
| | Lum. | | | | New |
| | | | | | CRF |
| 1 | 3.033333 | 0.11% | 0.44% | 0.11% | 0.11% |
| 2 | 0.453333 | 1.47% | 14.71% | 3.68% | 8.09% |
| 3 | 0.976667 | 1.71% | 8.53% | 1.71% | 3.41% |
| 4 | 2.176667 | 4.44% | 3.37% | 3.52% | 3.52% |
| 5 | 4.086667 | 5.55% | 1.55% | 5.30% | 6.28% |
| 6 | 6.623333 | 5.79% | 1.01% | 10.32% | 9.41% |
| 7 | 9.89 | 5.56% | 1.31% | 9.50% | 11.32% |
| 8 | 1.503333 | 4.43% | 7.76% | 7.10% | 7.76% |
| 9 | 6.886667 | 0.68% | 2.37% | 4.89% | 4.02% |
| 10 | 29.40333 | 5.15% | 1.75% | 7.94% | 10.62% |

D.3.4 Classroom Window

D.3.5 Spheres Far

| Target | Avg. Lum. | hdrgen | raw2hdr | Bracket | Bracket, New |
|--------|--------------|--------|---------|---------|-----------------|
| | | | / | | CKF |
| la | 126900 | 8.59% | 3.93% | 98.08% | 97.05% |
| 1b | 5653 | 13.73% | 2.96% | 5.08% | 5.29% |
| 2a | 31117.5 | 19.50% | 4.18% | 91.61% | 85.49% |
| 2b | 1373.75 | 0.36% | 5.28% | 3.92% | 5.07% |
| 3a | 6422.75 | 33.27% | 10.33% | 17.27% | 22.61% |
| 3b | 337.6 | 2.67% | 11.42% | 0.22% | 0.41% |
| 4a | 1677.75 | 11.84% | 3.03% | 13.44% | 12.96% |
| 4b | 112.4 | 26.25% | 31.71% | 22.79% | 23.86% |
| 5a | 187 | 0.06% | 0.05% | 0.02% | 0.10% |
| 5b | 11.34 | 2.38% | 1.76% | 1.85% | 2.73% |

| Target | Avg. | hdrgen | raw2hdr | Bracket | Bracket, |
|--------|----------|--------|---------|---------|----------|
| | Lum. | | | | New |
| | | | | | CRF |
| 1a | 127066.7 | 3.18% | 4.98% | 96.77% | 97.14% |
| 1b | 5721.667 | 19.70% | 3.76% | 10.56% | 6.46% |
| 2a | 31410 | 28.73% | 5.32% | 86.36% | 87.93% |
| 2b | 1383.333 | 3.92% | 2.02% | 3.52% | 3.26% |
| 3a | 6939.333 | 25.81% | 0.19% | 12.22% | 11.04% |
| 3b | 357.0667 | 2.33% | 12.42% | 1.18% | 1.28% |
| 4a | 1759 | 13.75% | 6.07% | 9.16% | 8.52% |
| 4b | 104.75 | 19.55% | 24.65% | 15.79% | 16.97% |
| 5a | 189.0667 | 0.31% | 0.17% | 0.37% | 0.84% |
| 5b | 10.71667 | 3.76% | 5.44% | 5.35% | 8.06% |

D.3.6 Spheres Close