

Spring 1-1-2017

A Comparison of Techniques Used for the Removal of Lens Flare Found in High Dynamic Range Luminance Measurements

Sarah Frances Safranek
University of Colorado at Boulder, SSafranek@me.com

Follow this and additional works at: https://scholar.colorado.edu/cven_gradetds

 Part of the [Architectural Engineering Commons](#), and the [Optics Commons](#)

Recommended Citation

Safranek, Sarah Frances, "A Comparison of Techniques Used for the Removal of Lens Flare Found in High Dynamic Range Luminance Measurements" (2017). *Civil Engineering Graduate Theses & Dissertations*. 92.
https://scholar.colorado.edu/cven_gradetds/92

This Thesis is brought to you for free and open access by Civil, Environmental, and Architectural Engineering at CU Scholar. It has been accepted for inclusion in Civil Engineering Graduate Theses & Dissertations by an authorized administrator of CU Scholar. For more information, please contact cuscholaradmin@colorado.edu.

A COMPARISON OF TECHNIQUES USED FOR THE REMOVAL OF LENS FLARE
FOUND IN HIGH DYNAMIC RANGE LUMINANCE MEASUREMENTS

by

SARAH FRANCES SAFRANEK

B.S., University of Colorado Boulder, 2015

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
Of the requirement for the degree of
Master of Science
Department of Civil, Environmental, and Architectural Engineering
2017

This thesis entitled:

A Comparison of Techniques Used for the Removal of Lens Flare Found in High Dynamic
Range Luminance Measurements

written by Sarah Frances Safranek

has been approved for the Department of Civil, Environmental, and Architectural Engineering

C. Walter Beamer IV

Sandra Vásconez

Moncef Krarti

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.

ABSTRACT

Safranek, Sarah Frances (M.S., Civil, Environmental and Architectural Engineering)

A Comparison of Techniques Used for the Removal of Lens Flare Found in High Dynamic Range Luminance Measurements

Thesis Directed by Professor C. Walter Beamer IV, Ph.D.

The presence of lens flare has been identified as a limitation associated with luminance measurements obtained using High Dynamic Range Imaging (HDRI) technology. Current documentation is lacking on the magnitude of error lens flare has on luminance values as well as the methods for lens flare removal. The purpose of this study is to provide an evaluation of the effectiveness of two existing lens flare removal techniques. The first technique involves the implementation of blind deconvolution during image postprocessing within commonly used HDRI software. The second technique aims to separate and extract lens flare from luminance values which are true to the scene using a high frequency occlusion mask. The results from this study indicated that while visually problematic, the error in luminance values resulting from lens flare effects a limited number of pixels located near the source. Of those pixels exhibiting significant error, current lens flare removal techniques are not sufficient at correcting this error nor do they adequately remove the visual appearance of lens flare found in HDR images. Future correction efforts may be best served by taking a step away from existing methods and working to create a more robust solution through the development of a Convolutional Neural Network (CNN).

Acknowledgments

Thank you to Pacific Northwest National Laboratory for their sponsorship throughout this entire experiment. I was very fortunate to glean inspiration from such a talented group of people. A special thank you is extended to Bob Davis for his guidance and encouragement during the thesis process.

I would also like to thank my advisor, Dr. Beamer, for his mentorship throughout my academic career as well as my graduate research. Many of my academic pursuits were greatly influenced by your guidance and your investment in student success.

Finally, I would like to acknowledge my parents for their unwavering encouragement throughout my academic career. Mom, thank you for challenging me to be my most ambitious self. Dad, thank you for supporting me through those challenges which proved to be most trying.

CONTENTS

CHAPTER 1: INTRODUCTION	1
1.1 Motivation.....	1
1.2 Report Overview.....	2
CHAPTER 2: LITERATURE REVIEW.....	3
2.1 Relevant Camera Definitions	3
2.2 High Dynamic Range Photography	6
2.2.1 Overview	6
2.2.2 Relevant Lighting Applications	8
2.2.3 Accuracy	10
2.2.4 Lens Flare Related Error	12
2.3 Lens Flare Luminance Error	14
2.3.1 Types of Imperfections Resulting from Lens Optics	14
2.3.2 Lens Construction and Performance	20
2.4 Attempts at Lens Flare Correction	23
2.4.1 Camera/Lens Improvements	24
2.4.2 The Point Spread Function.....	24
2.4.3 Lens Flare Removal through Blind Deconvolution.....	26
2.4.4 Lens Flare Removal by Introduction of Occlusion Mask.....	29
2.4.5 Artifact Removal Using a Neural Network	36
CHAPTER 3: EVALUATION OF CURRENT FLARE REMOVAL TECHNIQUE – BLIND DECONVOLUTION	40
3.1 Equipment.....	40
3.1.1 Camera/Lens System.....	40
3.1.2 Integrating Sphere and Light Source.....	41
3.2 Methodology.....	42
3.2.1 Set-up	42
3.2.2 Vignetting Correction.....	45
3.2.3 Removal Techniques.....	46
3.3 Technique Evaluation	48
3.3.1 Effectiveness of Removal.....	48
3.3.2 Conclusions of Evaluation	56
CHAPTER 4: ADAPTATION OF HIGH FREQUENCY OCCLUSION MASK FOR LUMINANCE MEASUREMENTS	57
4.1 Methodology.....	57
4.2 Post Processing Procedure	59
4.3 Preliminary Results and Conclusions	63
4.4 Restructured Results and Conclusion	66
CHAPTER 5: CONCLUSION	70

5.1 Study Conclusion	70
5.2 Recommendations for Future Research Using Neural Network	72
APPENDIX A: Vignetting Correction Documentation	78
APPENDIX B: Luminance Images	84
APPENDIX C: Luminance Error Data	88

LIST OF TABLES

Table 1: Descriptions of common lens aberrations [16].....	16
Table 2: List of maintained camera settings	40
Table 3: Vignetting Correction Functions by Aperture for Sigma Prime Lens	45
Table 4: Error bins of 5ft condition	50
Table 5: Error Bins of 10ft Condition.....	51
Table 6: Average and Maximum Error, 5 foot Condition	53
Table 7: Average and Maximum Error, 10 foot Condition	53
Table 8: Steps for Removing Lens Flare Using Occlusion Mask.....	61
Table 9: Pixel Measurements for Corrupted and Corrected HDRIs	67
Table 10: Comparison of Corrupted and Flare Reduced HDRIs by Aperture.....	67
Table 11: Vignetting Correction Functions for Sigma Prime Lens	83
Table 12: Vignetting Correction Functions for Zoom 18mm Lens	83
Table 13: Vignetting Correction Function for Zoom 24mm Lens.....	83

LIST OF FIGURES

Figure 1 : Simple Lens Model Describing Focal Length [4]	3
Figure 2: Depiction of change in aperture setting [1]	4
Figure 3 : Example of LDR images used to produce HDR image [5]	7
Figure 4: Visual representation of HDR process [5]	8
Figure 5: Lens flare effects varying by aperture with falsecolor renders used to determine extent of flare [7]	9
Figure 6: HDRI lens flare resulting from high luminous sources [5]	13
Figure 7: Left, graphic depicting the origin of ghost images [17]. Right, example of ghost artifacts resulting from solar source [1].	17
Figure 8: Diffraction spot surrounding point source [19]	18
Figure 9: Example of lens flare artifacts [16]	19
Figure 10: Example of Sigma Prime lens construction [2]	21
Figure 11: Comparison of sphere and aspheric lenses [1]	22
Figure 12: Example of testing resolutions for Nikon "FX" camera [1]	23
Figure 13: Example of MTF charts for Sigma 24mm, F1.8 prime lens	23
Figure 14: Example of point source used to determine PSF [14]	25
Figure 15: Left, PSF as a function of shutter speed. Right, PSF as a function of eccentricity. [14]	26
Figure 16: Annuli drawn around hot pixels in HDR image	28
Figure 17: Flare removed in an HDR image using blind deconvolution [22]	29
Figure 18: Graphic of glare removal using high frequency occlusion mask [23]	31
Figure 19: Occlusion mask removal steps: (a) Unoccluded original. (b) Single occluded capture, r_{ϕ} . (c) Glare-only image, g_{ϕ} . (d) Glare-free estimate after Gaussian interpolation, g_{ϕ} . (e) Glare-free estimate of corrupt occluded image, s_{ϕ} . (f) Final composite image after all glare-free estimates had been stitched together, s	33
Figure 20: Experimental set-up where 1 is the occlusion mask, 2 is the scene and 3 is the tethered camera.	34
Figure 21: Final results after flare removal using a blind deconvolution method (left) and the high frequency occlusion mask (right)c	35
Figure 22: Example of simplified neural network structure [25]	37
Figure 23: Image corrupted with rain droplets (left) and the same image after correction through CNN (right) [26]	38
Figure 24: Example of integrating sphere containing LED source for experimentation	42
Figure 25: View of sphere with source located 5 feet from the camera	43
Figure 26: Example of circular blocker placed 1 foot from source	44
Figure 27: Examples of source unblocked (left) and blocked (right) at 5 feet (top) and 10 feet (bottom) from the camera	44
Figure 28: Clockwise from top left: Images of the original scene luminance, blocked scene luminance, flare only, and normalized luminance values captured at f/11	48
Figure 29: Left, Normalized luminance of 10ft, f/16 condition. Right, Psuedo color render of normalized luminance of 10ft, f/16 condition	49

Figure 30: Histogram of error at 5 foot condition with f/1.4	51
Figure 31: Histogram of error at 10 foot condition with f/16	52
Figure 32: Maximum Luminance Error Comparison, 5 ft	54
Figure 33: Maximum Luminance Error Comparison, 10ft	54
Figure 34: 5ft, f/1.4 Radial Percent Error	55
Figure 35: 10ft, f/16 Radial Percent Error	55
Figure 36: Example of setup with occlusion mask	58
Figure 37: HDRI image with occlusion mask	59
Figure 38: Halos resulting from thickness of physical occlusion mask	60
Figure 39: Left: Example of single flare estimate, Center: Stitched flare estimate, Right: Final flare estimate after Gaussian interpolation to reduce period artifacts	66
Figure 40: Integrating sphere set-up with camera in place	81

CHAPTER 1: INTRODUCTION

1.1 Motivation

Lens flare is a recurrent problem throughout high dynamic range imaging (HDRI). Not only does it hinder the visual perception of an image, but it has also been identified as an issue when obtaining precise HDRI luminance measurements. The presence of flare in High Dynamic Range (HDR) captures can lead to inaccuracies in absolute and relative luminance measurements. It has been established that the conditions which foster the presence of lens flare are scenes containing high luminance sources against dark, low luminance backgrounds. These are the same settings in which HDRI is often employed as a tool to carry out glare analysis through luminance measurements. Such scenarios of high contrast often occur in scenes with high concentrations of daylight present in interior settings or stadium lighting. There have been studies related to mitigating lens flare as a visual impedance, however, it is also important to consider how effectively these solutions translate when obtaining accurate luminance measurements are being considered. The lighting industry is currently lacking a quantification of the error surrounding lens flare in luminance measurements as well as a solution for its removal visually in severe lighting settings.

This study identifies the primary causes of lens flare and assesses the most readily available lens flare correction solution as it relates to collecting accurate luminance measurements. An alternative technique for removing flare is proposed and evaluated on its effectiveness of reporting accurate luminance measurements.

1.2 Report Overview

This report is separated into five chapters. The second chapter, following this chapter's project overview, contains a detailed literature review. The literature review outlines the relevant research surrounding HDRI's use as a tool for the lighting community as well as how the optical phenomenon of lens flare affects obtaining accurate luminance measurements. The third chapter evaluates the method currently available for lens flare removal on its ability to reduce luminance measurement error. After the evaluation, a new lens flare removal technique is proposed and analyzed in chapter four. The final chapter concludes the topics discussed in this study and offers suggestions for future lens flare related studies.

CHAPTER 2: LITERATURE REVIEW

2.1 Relevant Camera Definitions

A series of relevant definitions and descriptions specific to cameras and lenses are provided in this section to allow for an improved understanding of the terminology referenced throughout this document. While this is not an extensive list of definitions, more information can be found at both the Nikon and Sigma websites [1] [2]. For definitions specific to lighting terminology and metrics, please reference the IES Lighting Handbook [3].

Focal Length: The focal length of a lens refers to the distance between the lens and the sensor when the subject is in focus. Focal length is unique to each lens and can influence the level of magnification of a subject. Wide angle lenses with a short focal length provide low magnification with a wide field of view while telephoto lenses have long focal lengths allowing the subject to appear magnified. Lenses with a fixed focal point are called prime lenses.

Mathematically, focal length (f) can be defined by Eq. 1:

$$\frac{1}{f} = \frac{1}{z'} + \frac{1}{-z} \quad \text{Eq. 1}$$




Figure 1 : Simple Lens Model Describing Focal Length [4]

Aperture: A camera’s aperture setting corresponds to the opening of the diaphragm of a lens, which is responsible for the amount of light allowed to pass through to the camera’s sensor. The iris, the mechanism used to adjust the diaphragm, is comprised of a series of blades. The variable diameter of the aperture is expressed in f-numbers (i.e., f/5.6) also referred to as f-stops. More precisely, the f-stop indicates the focal length of the lens (f) divided by the diameter of the aperture (Eq. 2.) The higher the f-stop, the smaller the opening and the less light that is passed through the lens. Each step down in aperture corresponds to halving the amount of light allowed through the diaphragm [1].

$$\frac{f}{D} = \frac{\text{focal length}}{\text{Diameter of aperture}} \quad \text{Eq. 2}$$

Additionally, varying the lens aperture influences the portion of the image that appears sharp in definition, also known as the depth of field of an image. Adjusting the aperture size manipulates the level of detail found beyond the focused subject of a scene. Again, this means, the higher the f-stop, the smaller the aperture’s opening, and resultantly a greater depth of field. Adjusting this setting may be valuable to eliminate the unwanted detail found in the background of an image’s subject.




Figure 2: Depiction of change in aperture setting [1]

Shutter Speed: While the aperture is the size of the lens' opening, the shutter speed is the amount of time that this opening is allowed to let light reach the camera's sensor. The camera's shutter, located behind the aperture diaphragm, acts as a "closed curtain" and only opens during image capture for a specified amount of time. The shutter speed can range anywhere from fractions of a second up to 30 seconds on advanced cameras. The longer the shutter is open, the more light that is allowed through to the sensor.

ISO: After light enters the lens, passes through the aperture's diaphragm and the shutter, it is then gathered and translated into a digital image via the camera sensor. The ISO setting allows control over the level of sensitivity of the sensor's response to light. A low ISO number corresponds to a lower sensitivity to light. In situations of low light levels, a higher ISO setting may be used to compensate. However, with higher sensor sensitivity comes increased granularity of the image.

Exposure Value: A combination of aperture, shutter speed, and ISO settings contribute to the overall exposure level of an image and are often referred to as the exposure triangle. For every photograph, there is a corresponding exposure value (EV) that can be calculated based on the exposure triangle's values as seen in Eq. 3. The EV reports the level of brightness displayed in the image where an EV of 0 would mean that the image is "properly exposed." An EV of -1 or 1 would indicate that the photo is underexposed with half the perceived brightness or overexposed with double the perceived brightness respectively. It is possible to achieve the same EV through multiple combinations of shutter speed, aperture and ISO. EV can be calculated as follows:

$$EV = \log_2 \frac{f/stop^2}{shutter\ speed} \quad \text{Eq. 3}$$

Camera Response Function: The camera response function is a nonlinear function which relates the intensity value to the scene radiance at any pixel located within a captured image. In RAW image formats, this function is linear. In the case of JPEG images, the camera response function can be used to counteract nonlinearities recorded in the image during JPEG compression. Each camera/lens combination has a unique camera response function.

2.2 High Dynamic Range Photography

2.2.1 Overview

High dynamic range imaging (HDRI) is a technique prevalent in the photography industry today with its capabilities found anywhere from high-quality camera systems to the everyday smartphone device. This method of photo-capture increases the dynamic range, or the difference between the lightest and darkest elements, of a scene in an effort to create an image which most closely represents the level of detail we are capable of seeing with the human eye. The HDR technique utilizes two or more low dynamic range (LDR) photographs, taken consecutively, of a single scene. During the capture process, the exposure level of the camera is adjusted such that the series of bracketed LDR images are able to capture the brightest and dimmest elements in the scene (Figure 3). The pertinent information from these LDR images are then fused together into a single HDR image that contains a greater level of detail than that of a single image. While HDRI was originally designed for artistic photography, there has been an attempt in the lighting industry to adapt the technology as a tool for luminance measurements due to the level of accessibility and efficiency of HDR systems.



Figure 3 : Example of LDR images used to produce HDR image [5]

The only equipment required during HDR capture is a consumer-grade digital camera capable of manual control and shutter speeds [5]. Use of a tripod, and tethering software is recommended to keep consistency in measurement. In the case of luminance mapping, a luminance meter is also required in order to calibrate the image for analysis. While there are numerous compilation applications available, for the purpose of luminance mapping, *hdrgen* and *raw2hdr* have been recorded to have minimal error as well as negligible difference between JPEG and RAW file formats for broadband light sources [5]. Absolute luminance value calibration is done by linearly scaling recorded pixel values according to a known reference point which would need to be recorded at the time of photo-capture with a luminance meter [5]. Finally, the HDR file undergoes a tone-mapping process which compresses the full dynamic range of the image into the dynamic range of the visual display monitor. A visual representation of the entire HDRI process can be seen in Figure 4.




Figure 4: Visual representation of HDR process [5]

2.2.2 Relevant Lighting Applications

Glare studies are one area in which HDR imaging and luminance mapping have the potential for significant impact. The level of variable complexity associated with glare makes it one of the more challenging lighting concepts to apply metrics to. Current glare metrics include Visual Comfort Probability (VCP), the CIE Glare Index (CGI) and Unified Glare Rating (UGR) for electric lighting and Daylight Glare Index (DGI) and Daylight Glare Probability (DGP) for glare caused by daylight. Each of these metrics requires the input of luminance values and viewing angle. Due to the number of luminance measurements required by traditional lighting specific meters, HDRI lends itself as a more efficient option. With a single HDR photograph, luminance data can be provided for every pixel in the scene, all while maintaining a single viewing angle. This can be particularly beneficial in quickly capturing lighting scenes with ample daylight, which may be constantly changing throughout a longer measurement process [6] [7].

In several daylight studies conducted by Van Den Wymelenberg, HDRI proved to be a useful tool to assess the level of glare experienced by occupants in various office settings [6] [8]. Based on the data gathered in these studies, common lighting design metrics were investigated leading to the justification of employing DGP over luminance based glare metrics in offices with an abundance of daylight because of their ability to better predict the subjective human responses associated with glare. Similarly, a study carried out by Jakubiec incorporated varying HDRI capture methods in the accuracy of luminance measurements in daylight interior scenes [7]. During the experimental process, high levels of daylight found in the scene resulted in “luminous overflow.” The HDRIs captured (Figure 5) of the direct sunlight indicated the severity of lens flare varied based on the aperture setting of the camera. Falsecolor luminance renderings of the scene displayed provided a greater level of detail for determining the magnitude of the lens flare artifacts than the HDRIs.




Figure 5: Lens flare effects varying by aperture with falsecolor renders used to determine extent of flare [7]

Near-field photometry

HDR photography has the potential to benefit photometry, particularly in near-field applications. Conventional measurement methods comply with the “five-times rule [9]” such that photometry must be performed at a distance at least five times greater than the maximum projected dimension of the source. The benefits of using HDRI in near-field settings was demonstrated in one study, where HDRI was utilized to assess the uniformity of a luminance across a single LED chip [10]. Due to the small dimensions of the LED in the study, typical spot measurements using current meters were unable to capture the source in its entirety, resulting in low luminance readings. A pixel-by-pixel analysis validated the use of HDR as an alternative tool for measuring small, high intensity sources. Additional research and validation of this method under a variety of experimental settings is still needed before HDRI can be implemented for regular near-field photometry applications.

2.2.3 Accuracy

While there is much excitement surrounding HDR photography as a lighting tool, users must keep in mind that this technology was not specifically designed for luminance measurement. It is therefore crucial that all of the limitations and potential inaccuracies be fully understood before interpreting HDR luminance results as absolute data.

Prior to HDR capture, there are a series of camera settings that must be maintained to increase the overall accuracy. It is important that the white balance setting matches that used to determine the camera response function. A daylight white balance setting is commonly recommended but a recent University of Colorado study found that the daylight setting may contribute to inconsistencies and abnormalities during capture, particularly in the color channels of the image [5]. While the proper white balance setting may vary from camera to camera, in the

case of the aforementioned study, a fluorescent white balance proved to be most accurate and was used for the remainder of the experimentation. The lowest ISO setting available (typically 100) is recommended for most applications as it decreases the level of noise added to the image [5]. Only in situations of low-light or action shots is a higher ISO setting recommended.

Immediate results obtained from HDRI represent the scene's relative luminance. Calibration from a known reference luminance within the scene is required at the time of capture to linearly translate these relative values to absolute within HDR software. To improve the accuracy of the calibration, it is important that the reference luminance have a value that is representative of the majority of the scene's luminance [11]. It is also recommended that when lens flare is present in instances of high contrast scenes, the calibration point should be somewhere in the scene where the flare does not visually affect the image [10].

The currently available HDRI compilation and analysis software also has proven to be problematic when interpreting measured values. In a previous University of Colorado study, a comprehensive evaluation was conducted of relevant HDRI software in which a majority of the programs being currently available yielded unacceptable levels of error in luminance values [5]. Of the software examined, *hdrgen* and *raw2hdr* proved to be the best options for HDR image creation and were recommended for HDRI analysis with a higher level of precision offered by the latter program [5].

Throughout numerous HDRI related studies, vignetting has been acknowledged as an inherent issue. Resulting from the optical structure of the lens, vignetting appears as a fall-off of brightness towards the edge of an image. The level of severity depends on the camera-lens combination as well as the aperture setting. Several studies have offered solutions to vignetting [10]. One such solution, discussed further in Chapter 3, determined a series of correction filters

that should be applied during post-processing to the HDR image depending on the lens and aperture setting used during photo capture [5].

The use of HDRI in scenes involving color, particularly colored LEDs, has also been proven to cause significant errors in luminance measurement. Previous findings have recommended using neutrally colored surfaces as significant errors occur in instances of cool colors like blue, green and purple as well as in colors with higher levels of saturation [12]. Additionally, due to the compression of information in JPEG creation and the associated color problems, RAW file format is advised [10]. An extensive study assessing the level of error found in luminance measurements throughout the visible spectrum by breaking down a broadband source into narrow bandwidth sections confirmed the previous color related issues [13]. This study also introduced further concern over the accuracy of photometric measurements for LED sources, particularly when absolute measurement is desired. Incandescent sources or those with a similar visual spectrum will yield more accurate results than most LED sources [13].

2.2.4 Lens Flare Related Error

Lens flare and diffraction effects, which will be discussed extensively in the following section, have been identified as a challenging inaccuracy associated with HDRI luminance measurements in a number of studies ([5] [7] [10] [14] [15].) Due to the level of variability surrounding lens flare, it is difficult to quantify the magnitude of error associated with these effects or to offer a simple solution to counter lens flare [5].

During the HDRI capture process, lens flare may be present in some bracketed images more so than others as the exposure time affects the amount of time the sensor is exposed to the scene's light [14]. As a result, some of the LDR bracketed images with prolonged exposures

record high levels of flare while shorter exposures may leave the flare undetectable. When the image is fused into a HDRI a balance of the flare will be displayed.

Consideration must be taken when selecting the proper aperture size for HDR capture. In the case of glare studies, the smallest aperture with any given lens is recommended if only the maximum luminance value of the scene is to be considered. This setting allows for minimal vignetting and the highest dynamic range but will also result in more dramatic lens flare [7]. Apertures as high as $f/22$ result in lens flare that potentially fill the entire picture [5] [7] as seen in Figure 6.




Figure 6: HDRI lens flare resulting from high luminous sources [5]

While there has been some work to mitigate the visual impedance caused by lens flare in HDR images, the lighting community is currently lacking an evaluation of how effective these methods are at accurately reporting final luminance values. Despite the ambiguity surrounding

lens flare error, it has been expressed by industry professionals that quantification and mitigation of lens flare effects would be valuable given HDRI's presence as a glare evaluation tool. The scenes that are the subject of glare analysis often contain levels of high luminance contrast either through small, bright sources against dark background or by ample daylight. The lighting conditions in these scenarios are also prime candidates for increased levels of lens flare in their HDR captures [14]. A consistent solution to this problem is currently lacking in the lighting industry and limits the potential HDRI has to offer in scientific applications.

2.3 Lens Flare Luminance Error

2.3.1 Types of Imperfections Resulting from Lens Optics

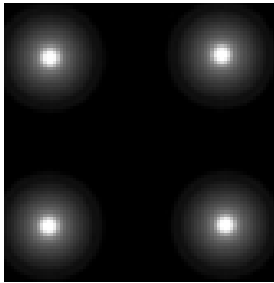
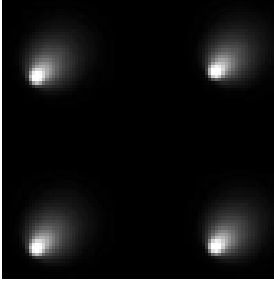
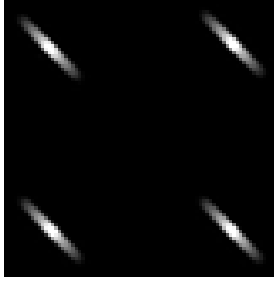

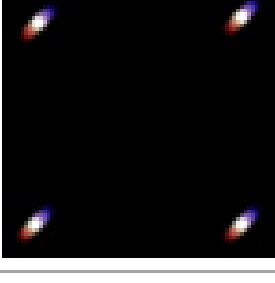
In order to offer solutions for removing lens flare, it is important to understand the origin of flare as well as how to distinguish it from other known imperfections resulting from the optical structure of the camera lens. Flare is defined as the unwanted or stray light in an optical system resulting from diffraction of light rays. An entirely transmissive lens system is impossible and as a result, there are imperfections exhibited through the camera's captured images. These imperfections include lens aberrations, ghosting and lens flare. There are a variety of solutions offered to combat each of these imperfections but due to the optical properties of the lens, they will never be eliminated entirely.

2.3.2.1 Aberration

Light rays experience refraction effects, or changes in wave direction, when they travel between transparent mediums. In a perfect optical system, all rays of light from a source in a scene would converge to the same point in the image plane. The influence which cause the rays to converge at different points are called aberrations and keep light from being brought into a sharp focus [16]. It is not possible to eliminate aberrations entirely but they are well understood

and with careful efforts they can be controlled and balanced during the lens design process to allow for the best possible results. Examples and existing correction methods for the most common aberrations found in photography are listed in Table 1.

Table 1: Descriptions of common lens aberrations [16]

Aberration	Character	Example	Correction
Spherical aberration	Monochromatic On- and off-axis Image blur		Bending, high index Aspherics Gradient doublet
Coma	Monochromatic Off-axis only, blur		Bending Spaced doublet with central stop
Oblique astigmatism	Monochromatic Off-axis blur		Spaced doublet with stop
Curvature of the field	Monochromatic Off-axis		Spaced doublet
Chromatic aberration	Heterochromatic On- and off-axis Blur		Contact doublet Spaced doublet

2.3.2.2 Ghost Artifacts

A second form of flare often identified in HDR photographs is that of ghost images resulting from reflections projected on the image plane from lens surfaces (Figure 7). The more glass surfaces within the lens increase the likelihood of ghost artifacts appearing in final images. Ghost images may occur singularly or in sets and can be acquired from inside or outside of the scene being captured. Often times, light rays entering the lens at varying and steep angles result in ghosting [16]. This particular issue can be mitigated by introducing a lens hood but cannot always be eliminated entirely. Ghosting can be differentiated from lens flare in that the artifacts will move about the image plane as their source angle changes [16].




Figure 7: Left, graphic depicting the origin of ghost images [17]. Right, example of ghost artifacts resulting from solar source [1].

2.3.2.3 Lens Flare

As a beam of light passes through an opaque object, such as a camera lens, it results in a slight lateral spreading known as diffraction [16]. This phenomenon is particularly problematic after the formation of HDR images and can be identified as star-like artifacts or extreme blurring at locations of high contrast between a bright source and a dark background. These two forms of diffraction can both be characterized based on the f-number of the aperture setting [18].

Described previously in Chapter 1, the mechanism responsible for controlling the aperture setting is a bladed iris. When light passes through the diaphragm of a lens, it is spread out perpendicularly to the diaphragm's blade. Ideally, the diaphragm would be circular as it allows all light passing through the lens to be scattered uniformly in all directions. Increased number of blades located at the diaphragm, lend to a more circular aperture, resulting in less, more symmetrical flare spread [16]. If an image were taken of a point source with a perfectly circular aperture diaphragm, the result would consist of the source itself as well as a “diffraction spot” or small circular spot surrounding the source (Figure 8).



Figure 8: Diffraction spot surrounding point source [19]

At larger apertures a more spectacular phenomenon occurs in the form of a bright “starburst” emanating from the light source. As the diaphragm's shape shrinks and deviates from a circle, light waves passing through the opening will be diffracted perpendicularly to each blade. An example of this can be seen in Figure 9 with dramatic star artifacts appearing throughout the image as a result of a film camera with a five-bladed iris. The use of a lens with a “wide open” iris or an iris with so many blades that its shape at larger f-stops is nearly a circle would resolve this issue [16].



Figure 9: Example of lens flare artifacts [16]

Similar, yet less severe, instances of lens flare are still present with more modern digital cameras. Previous research pertaining to HDRI and luminance measurements align with these notions of lens diffraction. Larger aperture sizes result in more dramatic starburst flares at locations of bright light sources [18]. During the capture process, diffraction is more prevalent in LDR images captured with a longer shutter speed as the sensor is exposed to direct light rays for longer periods of time than short shutter speeds. As a result, when the bracketed images are combined of LDR images with a range of shutter speeds, the lens flare prevents luminance data containing incorrect information about the actual light in the scene [5]. In same study, the incorporation of ambient light was tested to understand its effect on the resulting flare. Increased ambient light resulted in a decrease of visual severity in diffraction from the sources. However, it was inconclusive whether the ambient light truly minimized the error in luminance measurements are or simply minimized the visually perceived flare.

2.3.2 Lens Construction and Performance

Understanding the standard construction of lenses and their elements are the first step in countering the inaccuracies resulting from lens flare. The purpose of a lens is to collect a finite amount of light and focus those rays of light onto an image plane, turning a three-dimensional projection into a two-dimensional image plane [16]. The ideal image by lenses must fulfill three conditions, the first of which is: all light from the point object must be focused to a single point on the image plane. Second, when the object plane is perpendicular to the optical axis, the image plane must also be perpendicular. Lastly, the object and the image must closely resemble each other [1].

2.3.2.1 Grouped elements

Proper lens design requires the design and specification of a number of lens parameters involving the curvature of surfaces, the thicknesses, the air spaces, the diameters of components, and the types of glass to be used. Lens specifications provided by manufacturers often list the number of grouped elements included in their design. A higher number of lens elements can help control optical defects, however, each element introduced into the lens provides an opportunity for increased reflections and scatterings within. In composite lens systems, it is important that the optical axis of each lens component aligns perfectly [5]. Properly designed composite lens systems reduce the presence of defects and aberrations found in the final image. There is no direct correlation between the number of lens elements and the quality of the lens. An example of a typical lens specification provided by a manufacture can be seen in Figure 10.




Figure 10: Example of Sigma Prime lens construction [2]

2.3.2.2 Shape of Transparent Mediums

Traditionally, typical camera lenses were spherical; comprised of concave and complex glass surfaces with an optical axis that contains all centers of curvature for each surface. However, due to the increased level of manufacturing precision, aspherical lenses are becoming more commonplace. Aspherical lenses have non-spherical curves with counter curves along the edges of the frame and are capable of performing the function of two or more conventional spherical elements [2] (Figure 11). Each aspheric element has its own independent axis which must be made to coincide with all other centers of curvature within the lens system [20]. Such lenses require high precision manufacturing techniques but can effectively reduce aberrations and increase sharpness towards the frame edges throughout all aperture settings [1]. Additionally, the introduction of aspheric lenses results in fewer lens elements, allowing for more compact lens designs [2].




Figure 11: Comparison of sphere and aspheric lenses [1]

2.3.2.2 Lens Evaluation

The modulation transfer function (MTF) is the most comprehensive method for evaluating the optical performance of lenses [20]. MTF charts are often provided by lens manufacturers and contain information regarding the resolving power of the lens, the contrast in the image of coarse objects and effects from diffraction. These charts plot the contrast and resolution of a lens from its center to its edges against a “perfect” lens (a lens which is 100% transmissive). The data for these charts is created by opening the lens to maximum aperture and reading the contrast of sagittal and meridional line pairs at various points away from the center of the lens. The sagittal and meridional lines run parallel and perpendicular, respectively, to a central diagonal line from the center of the lens and are plotted at two resolutions: 30 lines/mm and 10 lines/mm (Figure 12). The data for 10 lines/mm specifies the capability of the lens to reproduce low spatial frequency or resolution while 30 lines/mm indicates reproduction of higher spatial frequency or resolution. The higher and straighter the plotted data, the better the performance of a lens. It should be noted that while these MTF charts are commonly provided by lens manufacturers, the testing methods used to produce the chart data vary making it difficult to compare across manufacturers. An example of diffraction and geometrical MTF charts for the Sigma 24mm, F1.8 prime lens used in Chapters 3 and 4 is depicted in Figure 13.




Figure 12: Example of testing resolutions for Nikon "FX" camera [1]




Figure 13: Example of MTF charts for Sigma 24mm, F1.8 prime lens

2.4 Attempts at Lens Flare Correction

Providing an all-encompassing solution to combat lens flare is extremely challenging due to its seemingly erratic behavior between aperture settings. There are three levels at which correction efforts can take place: at the lens, at the scene or during HDRI post-processing. Several different solutions have been offered through various studies for each of these levels however, none act as a definitive end to flare effects.

Furthermore, there is scarce documentation available when the conversation is shifted from lens flare as a visual impedance to lens flare as a cause of error in luminance measurements. The HDR capture of light that is not physically present in the scene may lead to incorrect reporting of absolute luminance values. One particular study, which evaluated HDR imaging as a luminance data acquisition system, was able to cite the error associated with flare as less than 10% on average [14]. While this value seems somewhat promising, there was little discussion surrounding where this error occurs or how it can be reduced. A better understanding of the characteristics and error associated with lens flare is necessary before HDRI technology can be implemented as a luminance measurement tool.

2.4.1 Camera/Lens Improvements

There are multiple methods for countering lens flare that attempt to minimize it directly at the optical elements of the camera/lens system. Oftentimes, this is by improving the quality of the lens elements through higher quality lens coatings that can potentially minimize the reflections occurring within. Another, more drastic, camera-oriented solution is the introduction of cameras which attempt to imitate the characteristics of the human eye [21]. This is done by employing liquids within the camera in order to reduce reflections from the lens surfaces. While these are viable attempts at mitigation, the potential for lens flare to occur will always remain.

2.4.2 The Point Spread Function

If a point source were to be placed far enough from a camera such that its light occupies a single pixel, the resulting HDR image should, ideally, display an image with that single pixel illuminated. However, this is rarely the case as the HDR image will often have multiple pixels surrounding the source illuminated as demonstrated by Figure 14. This “idealized, radially

symmetric characterization of the light fall-off surrounding a point of light in a perfectly dark surrounding” is also referred to as the Point Spread Function [22].



Figure 14: Example of point source used to determine PSF [14]

The Point Spread Function (PSF), which is commonly referenced in studies involving Fourier optics, microscopy and astronomical and medical imaging, is used to characterize the response of an imaging system to a point source or point object. Due to the physical characteristics of the optical systems, some portion of a pixel’s information is due to contributions from surrounding pixels and is ghost light not actually present in the scene [14]. The PSF varies not only depending on the lens/camera combination, location, and camera settings but also less manageable factors like oil and dust built up within the lens. As a result, it is difficult to establish properly calibrated PSFs for applications other than the most essential instances [22].

The controllable factors contributing to lens flare and point spreading are aperture size, exposure time and the distance from the optical center (also known as eccentricity). One particular study conducted at Lawrence Berkeley National Laboratory sought to characterize the

PSF in HDR luminance captures as a function of exposure time and eccentricity [14]. LDR captures with prolonged exposure times exhibited greater spillage of light around the source than shorter exposure times (Figure 15, left). Additionally, rotating the camera from normal to 90° on the point source’s periphery (using increments of 15°) exacerbated the pixel area effected by the point spreading (Figure 15, right). Conclusions from the study indicated that it is not feasible to quantify the general scattering resulting from lens flare but cited a margin of error of less than 10% on average. Those using absolute luminance measurements acquired through HDRI should factor this error into their final analysis [14].

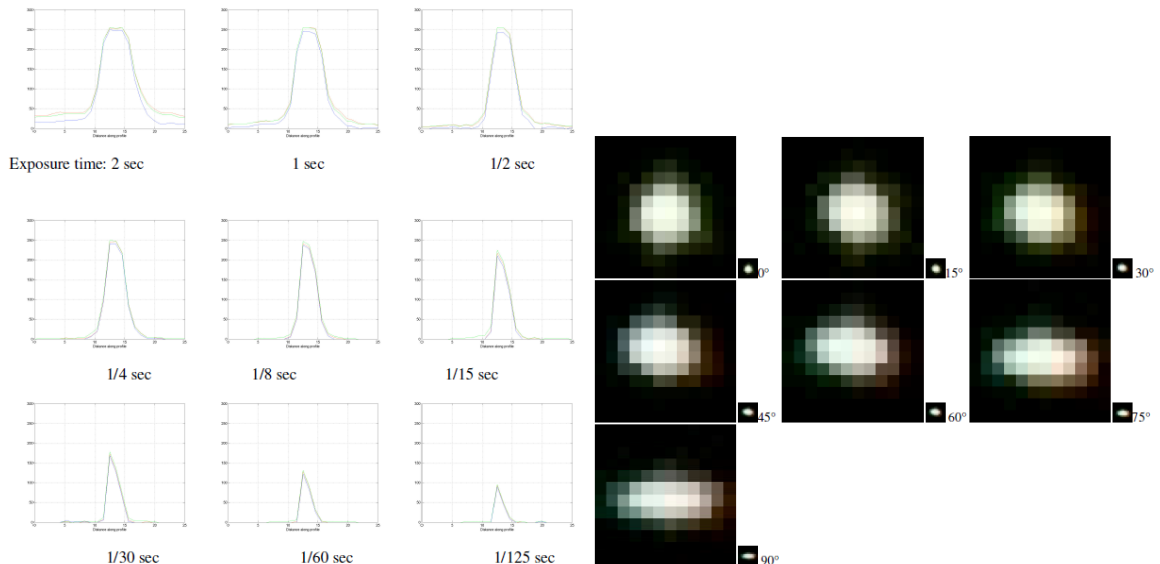


Figure 15: Left, PSF as a function of shutter speed. Right, PSF as a function of eccentricity. [14]

2.4.3 Lens Flare Removal through Blind Deconvolution

Due to the number of variables associated with lens flare and the spread of light surrounding a point source, it is most practical to determine a dynamic solution which can be employed on the final capture of an image. One such solution proposed by HDR experts involves the estimation of the PSF surrounding an image’s hot pixels [22]. This method is considered to

be a form of blind deconvolution making it similar to other methods used in medical and astronomical imaging [23].

This post-process lens flare solution works under a couple of assumptions; the first being that somewhere within the image there exists dark pixels located near bright, “hot”, pixels. If this is not the case, it is likely that the lens flare will go unnoticed and there would be no need to remove it [22]. It also is assumed that the nature of the flare is perfectly and radially symmetrical, which is not generally the case. As a result, any asymmetrical artifacts corresponding to the flare will remain present in the image. This assumption is somewhat crude as the starburst effects seen in lens flare are often dramatic and known to be asymmetrical, especially at higher apertures.

Before carrying out the blind convolution, “hot pixels”, or pixels with values greater than 1,000 times the minimum pixel value, are identified in a grey-scaled version of the HDR scene. Annuli are drawn around each of the designated hot pixels with equal radiuses to identify a minimum pixel value near the hot pixel (Figure 16). Based on the minimum ratio of dark/hot pixels for each annuli, a least squared approximation is computed for the PSF. Flare is then removed by subtracting the PSF estimate multiplied by the value of the hot pixels throughout a colored HDR image of the designated scene. An extensive description of this procedure can be found in Reinhard’s book, *High Dynamic Range Imaging* [22].

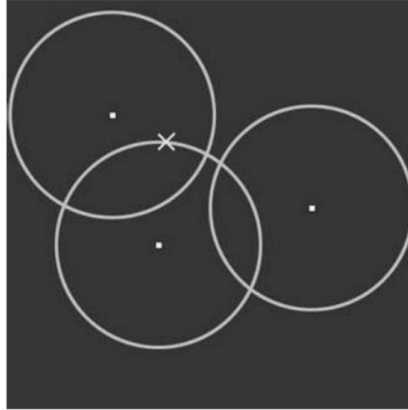


Figure 16: Annuli drawn around hot pixels in HDR image

The method detailed above can be applied during the HDR compilation processes provided in *raw2hdr* and Photosphere. By selecting the “remove lens flare” or `-f` option, the command will “attempt to estimate the camera/lens flare and remove it, but cannot be relied upon to remove this source of error completely” [10]. The image in Figure 17 demonstrates an ideal example for flare removal applied to an architectural scene using this method. An assessment of this deconvolution method as it pertains to lens flare in luminance measurements will be included in a following chapter.



Figure 17: Flare removed in an HDR image using blind deconvolution [22]

2.4.4 Lens Flare Removal by Introduction of Occlusion Mask

A unique correction method carried out at Stanford University [23] aimed to remove veiling glare from HDR images by placing a high frequency occlusion mask with perforated holes between the camera and the capture scene. In doing so, the mask facilitated in establishing the difference between direct light rays (the light physically present in the scene) and indirect light rays (the flare that is captured as a result of the optics of the camera lens). The indirect rays are then used to develop a glare estimate which can be subtracted from the original image, producing a final glare free image.

2.4.4.1 Theory

The primary goal of this method is to reconstruct a $(n \times m)$ corrupted image, $R(x,y)$ to produce a “fixed” $(n \times m)$ image, $S(u,v)$. The final, fixed image represents the scene captured using an ideal lens system, free from the effects of lens flare. The linear transform of glare from

S to R can be described by a 4D tensor, $\mathbf{L}(x,y,u,v)$. Each 2D plane of \mathbf{L} with constant u and v , or $\mathbf{L}_{u,v}(x,y)$, expresses the PSF value (referred to as the glare transfer function (GTF) in this particular report) for pixel (u,v) in the glare free image S . Similarly, each 2D plane of \mathbf{L} with constant x and y , $\mathbf{L}_{x,y}(u,v)$, describes the amount of glare pixel (x,y) in R receives from each pixel in S . This notation can be simplified by converting the corrupted and fixed images to $(nm \times 1)$ vectors $\mathbf{r}(x)$ and $\mathbf{s}(u)$ and the glare transform to a $(nm \times nm)$ matrix $L(x,u)$. The final equation for glare formation can then be carried out through simple matrix multiplication:

$$\mathbf{r} = \mathbf{L}\mathbf{s} \quad \text{Eq. 4}$$

By inverting Eq. 4, a corrupted image can be corrected when supplied with an appropriate flare estimate. Some knowledge is then required about the characteristics of the flare to generate an appropriate estimate. This is accomplished through the introduction of a high frequency occlusion mask as it provides some information about the direct light component of the scene while omitting a portion in a predictable fashion. The occlusion mask can be translated horizontally and vertically during a series of HDR captures such that every pixel in the camera's frame of view has information pertaining to the direct and indirect lighting components of the scene. By extracting the indirect luminance values for each pixel across the entire scene, a flare estimate can be stitched together through computerized post-processing.

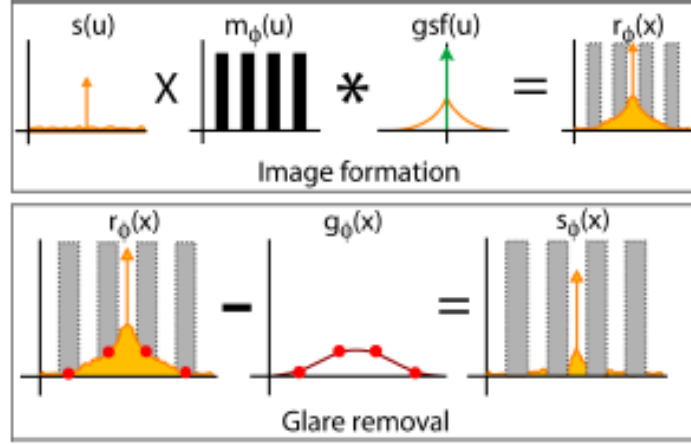


Figure 18: Graphic of glare removal using high frequency occlusion mask [23]

The frequency pattern of the occlusion mask must partially attenuate light without blocking it or introducing any additional glare, shine or scattering [23]. To achieve this, it is recommended that a large binary mask be constructed using high-quality optical material. While any high frequency pattern can be used, for simplicity, a rectilinear grid of square holes were cut into the mask. In doing so, the mask can be easily translated between captures such that every pixel receives exposure. With a mask occlusion factor of α (Eq. 5), it can be expected that the estimated glare at each pixel after reconstruction will be reduced by a factor of α [24]. The number of mask adjustments and HDR captures required to fully cover the image area is $\frac{1}{\alpha}$ although more can be taken for increased accuracy.

$$\alpha = \frac{\text{area of mask hole}}{\text{area of mask blocker}} \quad \text{Eq. 5}$$

Once the required number of photos are recorded, two post processing masks (\mathbf{m}_ϕ^+ and \mathbf{m}_ϕ^-) are constructed. The first mask, \mathbf{m}_ϕ^+ , contains all pixel information concerning the mask's holes in the HDR image while returning zeros for pixel values describing the occlusion mask.

Conversely, \mathbf{m}_ϕ^- omits pixel information belonging to the holes while keeping the values for the occlusion mask, revealing the glare that is present. For each mask position, the negative mask matrix, \mathbf{m}_ϕ^- , is multiplied elementwise with each corrupted capture, \mathbf{r}_ϕ , to form a flare estimate $\widehat{\mathbf{g}}_\phi$.

$$\widehat{\mathbf{g}}_\phi = \mathbf{m}_\phi^- \times \mathbf{r}_\phi \quad \text{Eq. 6}$$

For each glare estimate, the luminance values pertaining to the holes of the occlusion mask will be equal to zero. Interpolating $\widehat{\mathbf{g}}_\phi$ across the occluded regions using a weighted Gaussian blur will provide a complete image with a glare estimate for every pixel, including those that had been occluded. This step is represented in Eq. 7 where f is the Gaussian blur kernel. Again, the division in this equation is element wise resulting in a low frequency glare estimate, \mathbf{g}_ϕ , for the entire area of the scene.

$$\mathbf{g}_\phi = \frac{f * (\mathbf{m}_\phi^- \times \mathbf{r}_\phi)}{f * \mathbf{m}_\phi^-} \quad \text{Eq. 7}$$

Finally, the estimated glare from the previous equation is subtracted from the original corrupted image. The pixel information pertaining to the mask holes are then saved to create a composite glare-free image \mathbf{s}_ϕ , Eq.8

$$\mathbf{s}_\phi = \mathbf{m}_\phi^+ \times (\mathbf{r}_\phi - \mathbf{g}_\phi) \quad \text{Eq. 8}$$

This process is repeated independently for all captures and the result is a full set of glare-removed HDR images which can be stitched together to reveal a complete glare-free image.

After stitching, these photos are blended together using a weighted average to produce the final results. The entire process is depicted in Figure 19.




Figure 19: Occlusion mask removal steps: (a) Unoccluded original. (b) Single occluded capture, \mathbf{r}_Φ . (c) Glare-only image, $\widehat{\mathbf{g}}_\Phi$. (d) Glare-free estimate after Gaussian interpolation, \mathbf{g}_Φ . (e) Glare-free estimate of corrupt occluded image, \mathbf{s}_Φ . (f) Final composite image after all glare-free estimates had been stitched together, \mathbf{s} .

2.4.4.2 Methodology

For the formal experiment, an occlusion mask was created from a black, opaque, cardboard sheet with a 30x30 grid of 4mm square holes spaced 1mm apart for an occlusion factor α of 0.16. The mask was placed on a mechanical tripod as close to the scene as possible in an effort to keep the camera from defocusing (Figure 20). The scene depicted in this study had a uniformly lit backdrop with an object of interest placed in front of them. The contrast ratio

between the background and subject resulted in a veiling glare across the object of interest. Between HDR captures, the mask was translated horizontally and vertically in increments of 1cm into 36 different positions. The 36 captures were used to create a 6 x 6 grid of information for every pixel in the camera's view. Due to the number of HDR images required to reconstruct a glare free image, complete scene capture took between 30 and 60 minutes. After image capture, the 36 images were run through the post-processing steps described previously to stitch together a new, fixed image with the removal of glare. Finally, the results from the stitching process were compared against the same scene corrected using a blind deconvolution method to assess the change in dynamic range of the fixed images.




Figure 20: Experimental set-up where 1 is the occlusion mask, 2 is the scene and 3 is the tethered camera.

2.4.4.3 Conclusions

After stitching together the glare-free images, the uniformly veiled regions of the scene were recovered in full detail. The method was less successful at recovering scene content in

regions located near strong luminance edges. In central areas away from the strong edges, the dynamic range of the image was seen to improve from 560:1 to as much as 22400:1. An example of the removal can be seen in Figure 21 (right) with a comparison to the same scene corrected using a blind deconvolution method similar to that described by Reinhard et. al [22] in Figure 21 (left). The reconstruction using the occlusion mask improved the dynamic range of the scene while maintaining a limited amount of noise when compared to the deconvolution method.




Figure 21: Final results after flare removal using a blind deconvolution method (left) and the high frequency occlusion mask (right)c

The authors concluded that while the high frequency occlusion mask was successful in removing glare in HDR images, the time required to obtain the required HDR captures limits its application to smaller, static scenes. Additionally, the mask must be placed near the scene and remain nearly in focus for the best results, prohibiting the occlusion mask from being employed as a filter attached to the camera lens. While not ideal for every application, this experiment

offers a promising solution to reduce the presence of lens flare in HDR captures and increase the dynamic range of the image. This method has yet to be assessed for its ability to obtain accurate luminance measurements and its adaption for the use of correcting flare in such measurements may prove useful to the lighting community.

2.4.5 Artifact Removal Using a Neural Network

Machine learning through Artificial Neural Networks (ANNs) is an extensive and ever-developing topic and this section will provide an abbreviated description of ANNs as they may pertain to the development of a lens flare solution. In short, an ANN is a computational processing system modeled after the biological nervous system. The goal of these networks is to optimize a final output given a multidimensional vector of inputs through a series of layers and parameters. Hidden layers within the network, comprised of interconnected computational nodes, make decisions based upon a previous layer and determine how a stochastic change within the layer will achieve the final output (Figure 22) [25]. Neural networking offers a method of dealing with large data sets making it well-suited for pixelated images.




Figure 22: Example of simplified neural network structure [25]

Convolutional Neural Networks (CNNs), a type of ANN, are designed to reduce the number of neural connections made between layers and are employed as a tool for image and language recognition and classification. Instead of each computational node, or neuron, making connections to every other neuron in each layer, CNNs are designed such that neurons connect to sub regions in previous layers. Weight is given to pixels located within the same region as they are more likely to be related to one another [25]. Multiple layers fill the network attempting to detect different features such as color or directionality with the goal being that the network will learn image statistics for recognition. Through proper program architecture and training the networks are capable of complex processes such as face or object recognition.

Impressive image restoration has been demonstrated through the introduction of CNNs trained to perform blind inpainting on images with presence of significant noise. One such success was achieved in a NYU study which aimed to correct images that were corrupted with

dirt and water drops present between the camera lens and the capture scene (Figure 23) [26]. This form of localized corruption required the neural network to identify the presence and location of the noise at which point the network would alter those pixels to produce a clean image.



Figure 23: Image corrupted with rain droplets (left) and the same image after correction through CNN (right) [26]

In the case of the rain droplet removal, a set of images were collected of multiple scenes with and without noise present. For the clean images, a static scene was captured through a pane of anti-reflective glass, without any water droplets present. Corrupted images were taken of same scene after the effects of rain were simulated by lightly spraying the glass. The final network training set contained 6.5 million examples of 64 x 64 pixel clean/corrupted paired images. By

training the network on the difference between the pixel information for a corrupted patch and the same corresponding pixel information for a clean patch, the network begins to develop an algorithm for detecting noise and predicting the appearance of a clean output. After the network had been sufficiently trained by minimizing the mean squared error over the entire data set, the network could be applied to unique images which were corrupted by non-simulated rain droplets.

Image restoration through neural networks is a well-studied problem but because of the level of knowledge and experience required to model machine learning algorithms, it has remained somewhat of a niche subject. There is no single way of forming a network architecture and computation time and success can vary significantly depending on the level of complexity of the network. One of the largest limitations of implementing a neural network is the size of the data set required to properly train a network. While expanding the diversity of the training data allows for a more robust solution, it can be more taxing on the capture and computational processes. However, with enough motivation, the capabilities of machine learning and neural networks offer a compelling alternative for combating the variable complexity inherent with lens flare.

CHAPTER 3: EVALUATION OF CURRENT FLARE REMOVAL TECHNIQUE – BLIND DECONVOLUTION

Currently there are only a couple of lens flare removal techniques which have been offered to the HDRI community but an analysis of their effectiveness has not yet been recorded. In particular, a lens flare removal algorithm is provided through the HDR compilation programs, Photosphere and *raw2hdr*. The goal of the experiment included in this section was to evaluate the effectiveness of these removal options on a simple lighting scene. For this evaluation, measurements were taken in a darkened room at two distances (5 feet and 10 feet) from a source and compared before and after lens flare removal was carried out for both HDR programs.

3.1 Equipment

3.1.1 Camera/Lens System

The camera used throughout this study, as well as previous HDRI related studies at the University of Colorado [5] [13], is a Nikon D5200. It is equipped with complementary metal oxide semiconductor (CMOS) sensors which are recommended as the primary equipment for HDR photography [13]. The relevant settings maintained throughout the image capture process are listed below and discussion of each can be found in a previous study by Mio Stanley [5].

Table 2: List of maintained camera settings

ISO	100
Active D-Lighting	Off
Metering Mode	Matrix
White Balance	Fluorescent
File Format	JPEG Fine + RAW
Photo Size	S - (2000x3000 pixels)

The lens chosen for this experiment was the Sigma 24mm, F1.8 prime lens. This lens, which has also been used in previous University of Colorado HDRI studies [5] [13], was selected because of its fixed zoom feature. It is constructed with 15 lens elements in 11 groups. The diaphragm of the lens is made of 9 blades contributing to aperture settings ranging from $f/1.3$ - $f/16$. Aspheric lens elements in the lens design specifically eliminate the sagittal coma flare that accompanies large lens diameters and effectively suppresses lens flare [2].

3.1.2 Integrating Sphere and Light Source

The light source used for the tests in this study was a 60W CILED which offered an average luminance of $31,250 \text{ cd/m}^2$ and significant flare for analysis. Previous research into HDRI has revealed inaccuracies in absolute luminance measurements that are coupled with using sources other than those exhibiting an incandescent spectrum [13]. The values of interest during the experimentation and analysis of this particular study are strictly relative values and therefore the LED source being used is considered sufficient. For future studies evaluating lens flare for the purpose of absolute luminance measurements, it is crucial that care be taken when selecting the light source.

The source was placed in a small, 6'' diameter, integrating sphere (Figure 24) with a 1'' opening on the opposite side. The opening was covered with diffuse acrylic and tested to prove Lambertian luminance distribution across the opening [5]. A small, white mask was placed in the center of the sphere to keep the source from directly illuminating the acrylic opening. The integrating sphere was designed by Mark Jongewaard for the University of Colorado's HDRI research and has been used in previous HDRI studies ([5] [13]).




Figure 24: Example of integrating sphere containing LED source for experimentation

3.2 Methodology

3.2.1 Set-up

For experimentation, the source and sphere were placed in a neutrally colored laboratory setting with the absence of any stray light. The camera was mounted atop a tripod at 5 feet and 10 feet normal to the light source with its height adjusted such that the source was centered within the camera's view to avoid any abnormalities that may occur towards the edges of the image. To reduce movement during photo capture, the camera was tethered to a laptop and controlled remotely using the program *Sofortbild*. Six different apertures ranging from f/1.4-f/16 were captured at each distance. The shutter speed was set to cover the minimum speed (1/4000s) to the maximum speed (30s) in steps of nine resulting in seven bracketed photos for each HDR image. Prior to photo capture for each condition, several LDR photos were taken with the source turned off and ambient light present (Figure 25). These photos were used to document the exact pixel location of all significant objects in the scene and were utilized during post-processing.




Figure 25: View of sphere with source located 5 feet from the camera

The above procedure was repeated to obtain the same number of images for a control group. This time a 1 inch, opaque circle, laser cut from black poster board was placed 1 foot in front of the light source (Figure 26). The circle acted as a blocker prohibiting any light rays from directly reaching the camera lens. In doing so, the boundary of the source could be identified and accurate luminance values for the scene could be recorded without the contribution of stray light resulting from lens flare. Multiple LDR photographs were taken with the blocker in place to ensure that the entire source was directly blocked from the camera's view (Figure 27). It is worth noting that the introduction of the blocker may contribute to a small level of error captured in the scene. The specific poster board selected for this study was picked for its low reflectance properties, however, its introduction to the scene may cause the luminance values for the control images to differ slightly from what can technically be considered "true" values.




Figure 26: Example of circular blocker placed 1 foot from source




Figure 27: Examples of source unblocked (left) and blocked (right) at 5 feet (top) and 10 feet (bottom) from the camera

3.2.2 Vignetting Correction

As mentioned previously, vignetting effects from the camera/lens combination can produce errors in luminance measurement values. The source was purposely placed in the center of the captured scene to minimize the errors coupled with vignetting. To ensure additional accuracy, a vignetting correction MATLAB script, developed by the University of Colorado, was applied to the HDR captures prior to their analysis. All of the following HDR images in this study have been corrected for vignetting based on their corresponding correction filter.

This series of vignetting correction filters were created for each aperture for the Nikon D5200, Sigma Prime lens combination (Table 3). The filters were developed by placing the camera system inside of a 4-foot diameter integrating sphere maintained at uniform luminance. The resulting HDR images captured for each aperture were used to create functions which quantify the fall of light towards the edges of the image. Applying the inverse of the function across the pixels of an HDR image with the given aperture will counter the effects of vignetting. More detailed documentation surrounding the experiment and procedures used to obtain the vignetting correction functions can be found in Appendix A.

Table 3: Vignetting Correction Functions by Aperture for Sigma Prime Lens

f/number	Vignetting Correction Function
1.4	$1.59x^6 - 3.81x^5 + 2.52x^4 + 0.16x^3 - 0.78x^2 - 0.20x + 1.00$
3.5	$2.77x^6 - 6.94x^5 + 5.88x^4 - 0.85x^3 - 0.03x^2 + 0.02x + 1.00$
5.6	$2.06x^6 - 4.86x^5 + 3.58x^4 - 0.66x^3 - 0.32x^2 + 0.65x + 1.00$
9	$2.74x^6 - 6.63x^5 + 5.22x^4 - 1.23x^3 - 0.30x^2 + 0.05x + 1.00$
11	$1.97x^6 - 4.70x^5 + 1.57x^4 - 1.78x^3 - 1.02x^2 + 0.10x + 1.00$
16	$0.31x^6 + 0.59x^5 - 2.84x^4 - 2.94x^3 - 1.30x^2 + 0.16x + 1.00$

3.2.3 Removal Techniques

The most readily available lens flare removal technique, a blind deconvolution algorithm [22], can be applied through the HDR generation tools Photosphere and *raw2hdr*. As discussed previously, this method locates hot pixels and applies a counter function to the radial fall off of light surrounding these pixels. The algorithm found in both programs is fundamentally the same, however removal results may vary between the two based on the different file formats of the supplied images (JPG vs. RAW) and so both will be considered for this experiment. For each program, a lens flare removal option can be selected prior to the fusing of the LDR images. To evaluate the effectiveness of the deconvolution's ability to remove lens flare, HDR images were compared with and without the lens flare removal option selected for both programs. The final image set for this method included 48 total corrupted HDR photographs with 12 control HDR photographs.

After LDR compilation, each HDR image was converted at the pixel level from color format (RGB values) to luminance values in MATLAB using the following equation [27] :

$$L = \frac{179*0.2126R+0.7152G+0.0722B}{EV} \quad \text{Eq. 9}$$

The MATLAB luminance values were confirmed to match those calculated within the analysis software, HDRscope. The control HDR images with correct luminance values for the scene were subtracted from the corrupted HDR images of the same aperture to create a flare only image. Because each HDR image created is 3000 x 2000 pixels, the amount of data being processed can be quite extensive and difficult to handle. In an effort to reduce the volume of data, the HDR

images were cropped to include only the pixel area impacted by lens flare. The entirety of the source luminance data fit within a 600 x 600 pixel crop, centered around the source.

For the flare only images, a zeroed mask covering the source was created in MATLAB and overlaid atop the luminance pixel values. The exact size and location of the source was determined using a single LDR image of the scene with the source off and ambient light present and confirmed to cover the same area as the cardboard blocker installed in the control group images. This mask allowed for the differentiation between source and scene and aided in the calculation of any values of interest pertaining to the source or background. Finally, all luminance values in the flare image were converted to a normalized scale with the average luminance of the source equal to 1 in an effort to focus on the relative properties of lens flare. The corrupted luminance, blocked/control luminance, flare only and normalized images are displayed in Figure 28.




Figure 28: Clockwise from top left: Images of the original scene luminance, blocked scene luminance, flare only, and normalized luminance values captured at $f/11$

3.3 Technique Evaluation

3.3.1 Effectiveness of Removal

Initial observations of the luminance images corrupted by lens flare depict a scene where lens flare has a dramatic impact on a large number of pixels surrounding the source. A large circle with ill-defined edges can be seen surrounding the source. Additionally, 18 “legs” emanate from this circle brought about by the perpendicular diffraction of the direct rays from the source off of the 9 aperture blades within the lens. Introducing the cardboard blocker between the

camera and source successfully prohibits light from directly striking the lens and eliminates the lens flare. The blocked image set is used to represent the “true” luminance values of the scene without lens flare present. Both the unblocked and blocked image sets are normalized to evaluate the flare on a relative scale, at which point its presence in the unblocked images nearly disappears. This indicates that while lens flare appears to be dramatic visually, significant error only occurs at a small percentage of the pixels surrounding the source.

When viewing the normalized luminance images generated through MATLAB, there appears to be extreme lens flare present in the scene (Figure 29 left). Such an image can be somewhat misleading as demonstrated by Figure 29 right, which displays the same values on a relative scale. The luminance values represented by the first image are not visually weighted proportional to their brightness. This means that a pixel with a magnitude of 1.0 looks to be of the same brightness as a pixel with a magnitude of 100. It is recommended that psuedo color renders of the scene are referenced when considering the magnitude of luminance values.

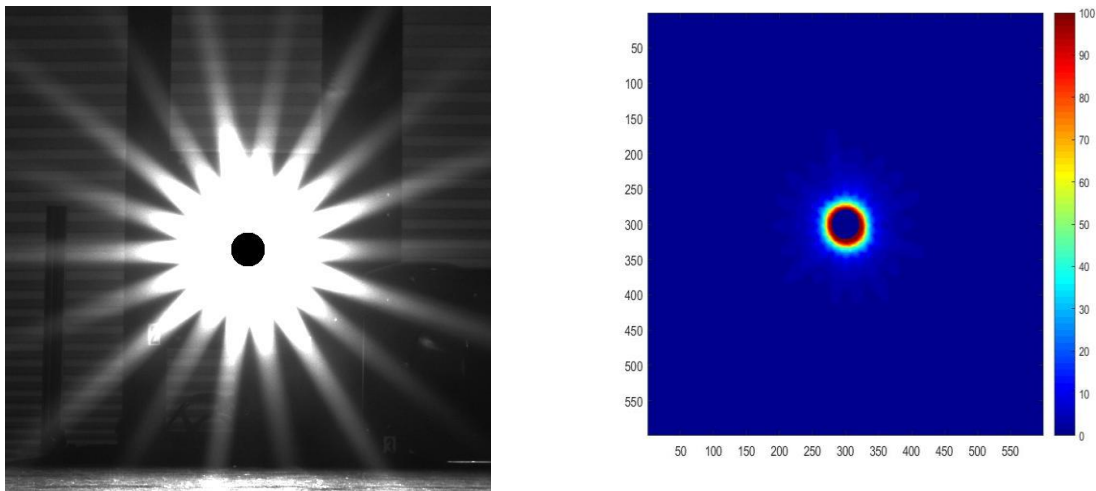


Figure 29: Left, Normalized luminance of 10ft, f/16 condition. Right, Psuedo color render of normalized luminance of 10ft, f/16 condition.

The percentage of error at each pixel was calculated for each of the 48 unblocked and normalized HDRIs using Eq. 10 to evaluate the effectiveness of the currently available lens flare removal algorithms. For the images exhibiting flare, the luminance values at each pixel were compared against the true value at each pixel, represented by the blocked image set.

$$\text{Percent Error} = \frac{|\text{Lens Flare Pixel } L - \text{True Pixel } L|}{\text{True Pixel } L} \times 100 \quad \text{Eq. 10}$$

Of the 36,000 data points, the number of pixels exceeding 10% error was found to be less than 1% of the cropped image in even the most severe apertures (Tables 4 and 5).

Histograms of the apertures with the highest level of error for each distance are included in Figure 30 and Figure 31 with the remaining histograms and corresponding false color renders located in the appendix.

Table 4: Error bins of 5ft condition

5 Foot - Error Bins											
f/#	< 1	1-10	11-20	21-30	31-40	41-50	51-60	61-70	71-80	81-90	90-100
1.4	358033	335	106	71	68	63	47	40	30	8	0
	99.454%	0.093%	0.029%	0.020%	0.019%	0.018%	0.013%	0.011%	0.008%	0.002%	0.000%
3.5	358653	127	14	7	0	0	0	0	0	0	0
	99.626%	0.035%	0.004%	0.002%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
5.6	358703	98	0	0	0	0	0	0	0	0	0
	99.640%	0.027%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
9	358413	388	0	0	0	0	0	0	0	0	0
	99.559%	0.108%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
11	358358	443	0	0	0	0	0	0	0	0	0
	99.544%	0.123%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
16	357875	926	0	0	0	0	0	0	0	0	0
	99.410%	0.257%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%

Table 5: Error Bins of 10ft Condition

10 foot – Error Bins												
f/#	< 1	1-10	11-20	21-30	31-40	41-50	51-60	61-70	71-80	81-90	90-100	> 100
1.4	340021	16483	1182	532	327	176	57	23	0	0	0	0
	94.45%	4.579%	0.328%	0.148%	0.091%	0.049%	0.016%	0.006%	0.000%	0.000%	0.000%	0.000%
3.5	316119	37485	2065	1062	639	412	321	252	203	144	87	12
	87.81%	10.413%	0.574%	0.295%	0.178%	0.114%	0.089%	0.070%	0.056%	0.040%	0.024%	0.003%
5.6	336719	19070	1543	604	410	203	117	73	46	16	0	0
	93.53%	5.297%	0.429%	0.168%	0.114%	0.056%	0.033%	0.020%	0.013%	0.004%	0.000%	0.000%
9	346838	10811	719	254	114	53	12	0	0	0	0	0
	96.344%	3.003%	0.200%	0.071%	0.032%	0.015%	0.003%	0.000%	0.000%	0.000%	0.000%	0.000%
11	345257	12263	859	291	106	25	0	0	0	0	0	0
	95.91%	3.406%	0.239%	0.081%	0.029%	0.007%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
16	321098	33153	924	1145	717	434	325	310	261	254	158	22
	89.19%	9.209%	0.257%	0.318%	0.199%	0.121%	0.090%	0.086%	0.073%	0.071%	0.044%	0.006%

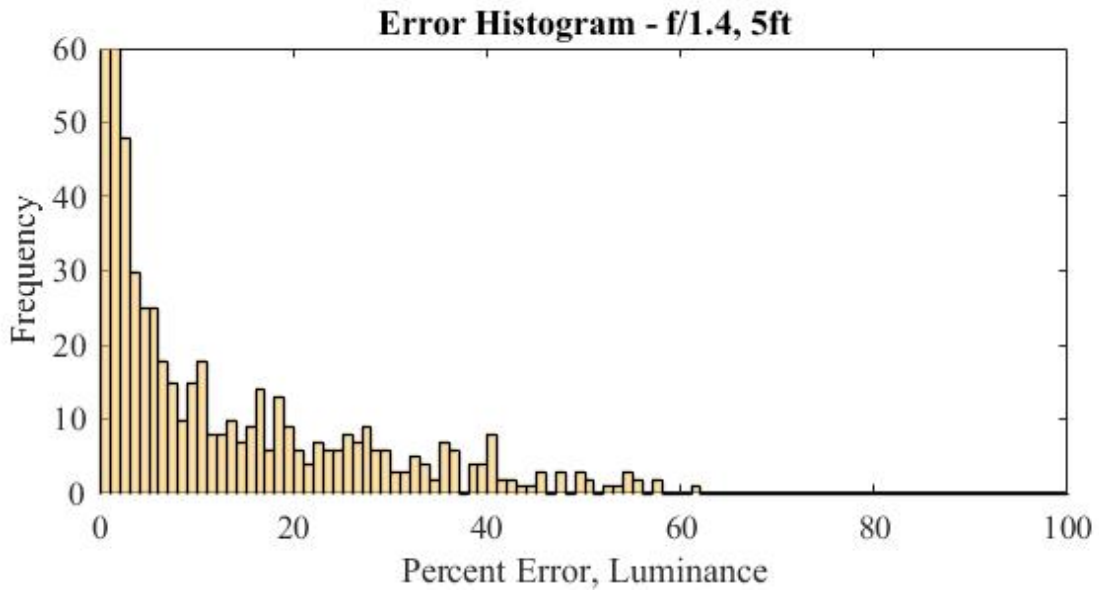


Figure 30: Histogram of error – at 5 foot condition with f/1.4

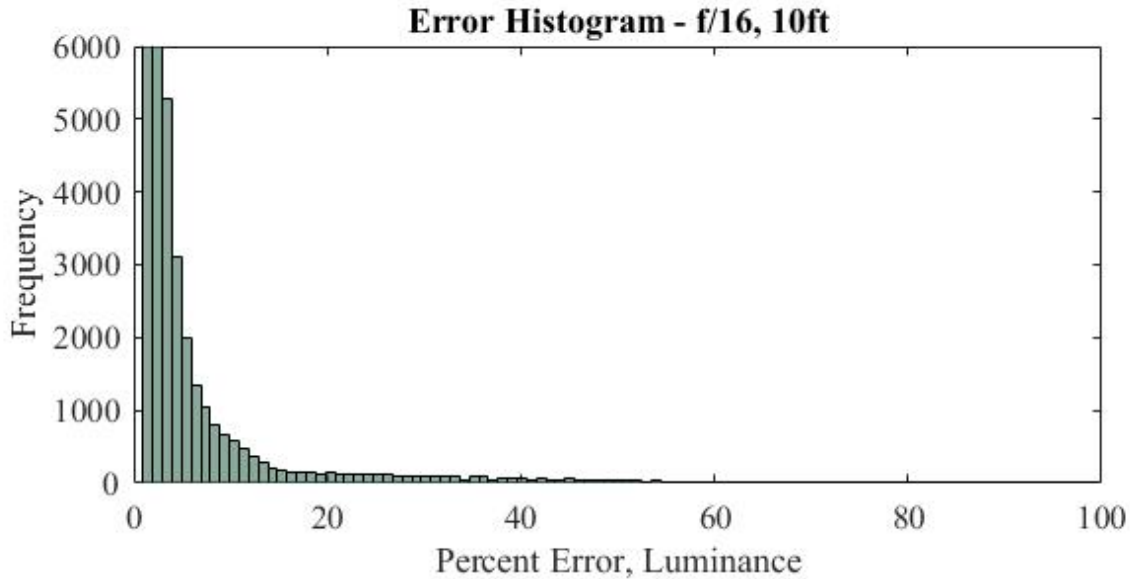


Figure 31: Histogram of error at 10 foot condition with f/16

The pixels exhibiting error were then used to determine the average and maximum luminance error for each photograph (Tables 6 and 7). While no image recorded an average error over 1.1%, the maximum error for all conditions was considered to be significant. The f/1.4 setting at the 5 foot condition and the f/16 aperture setting at the 10 foot condition reported the highest maximum error for each distance at 61.97% and 100.71% respectively. Referencing Figure 32 and Figure 33, the maximum error associated with *raw2hdr* (R2H) was almost always less severe than that of Photosphere (PS) at both of the captured distances. Interestingly, the application of the lens flare removal options (LFR) did not result in much, if any, improvement of the error conditions. In a several instances at the at both distances, the lens flare removal increased the maximum error found in the scene. This indicates that for the simple scene used in for this experiment all apertures were subject to lens flare. Finally, the average error was calculated as a function of radial distance from the source (Figure 34 and Figure 35). Significant fall off of error occurred within 5 pixels for the 5 foot condition and within 35 pixels for the 10

foot condition. The increase in distance between the source and the camera results in larger values for average and maximum error across all apertures likely due from the increased severity of the diffraction effects resulting from direct light rays hitting the diaphragm of the lens.

Table 6: Average and Maximum Error, 5 foot Condition

5ft								
f/#	Average				Maximum			
	PS	PS, LFR	R2H	R2H, LFR	PS	PS, LFR	R2H	R2H, LFR
1.4	0.0315%	0.0281%	0.0185%	0.0171%	61.97%	61.98%	55.81%	55.81%
3.5	0.0086%	0.0072%	0.0012%	0.0004%	44.44%	45.02%	14.11%	14.11%
5.6	0.0770%	0.0057%	0.0037%	0.0047%	25.82%	25.93%	14.82%	14.82%
9	0.0113%	0.0110%	0.0436%	0.0449%	22.29%	22.05%	11.25%	11.29%
11	0.0131%	0.0111%	0.0647%	0.0660%	24.60%	24.60%	11.35%	11.36%
16	0.0185%	0.0180%	0.2387%	0.2407%	38.52%	38.52%	38.52%	36.80%

Table 7: Average and Maximum Error, 10 foot Condition

10ft								
f/#	Average				Maximum			
	PS	PS, LFR	R2H	R2H, LFR	PS	PS, LFR	R2H	R2H, LFR
1.4	0.6177%	0.4361%	0.0255%	0.0214%	70.45%	59.22%	98.84%	97.87%
3.5	1.2109%	0.9988%	0.0035%	0.0033%	97.65%	100.75%	72.76%	72.76%
5.6	0.6977%	0.5502%	0.0031%	0.0034%	76.36%	85.15%	63.06%	63.06%
9	0.3681%	0.2307%	0.0301%	0.0304%	53.51%	56.87%	53.45%	53.45%
11	0.2976%	0.2714%	0.0533%	0.0535%	48.29%	48.16%	44.29%	44.23%
16	1.0938%	1.0339%	0.1724%	0.1827%	100.71%	100.65%	69.76%	69.45%




Figure 32: Maximum Luminance Error Comparison, 5 ft

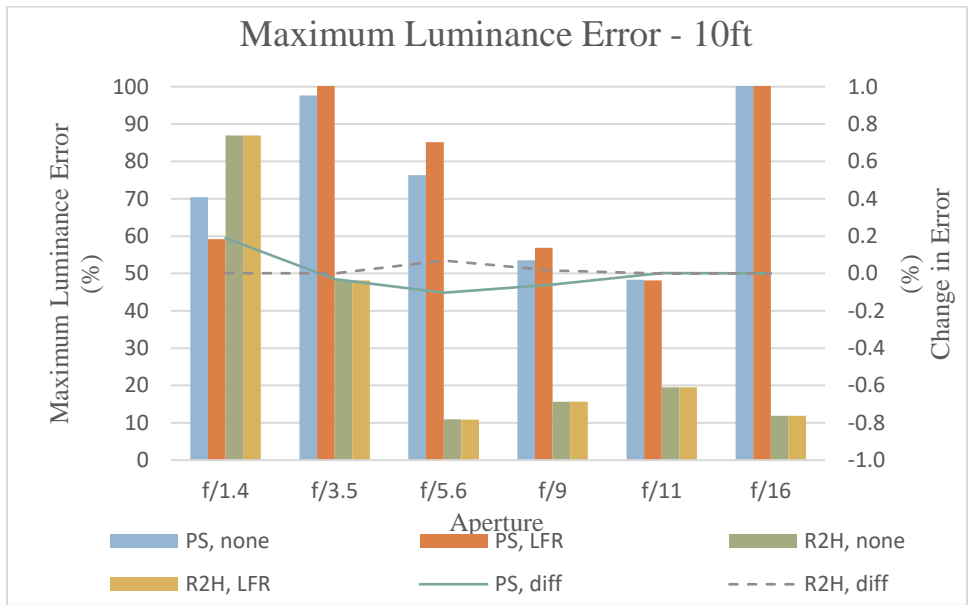


Figure 33: Maximum Luminance Error Comparison, 10ft




Figure 34: 5ft, f/1.4 Radial Percent Error




Figure 35: 10ft, f/16 Radial Percent Error

3.3.2 Conclusions of Evaluation

This experiment established that while the pixel area affected by flare may not encompass a large portion of the image, those pixels which do experience flare are subject to significant error. For the simple scene used for this analysis, the lens flare removal techniques found in the HDR programs Photosphere and *raw2hdr* offered little to no improvement in maximum error correction with minor average error reduction. Despite these minor reductions, the lens flare removal options should not be relied on entirely. In instances where the light source occupies a large percentage of the image, such as the potential application of near field photometry, the existing removal algorithm does not offer a worthy solution. There is reason to believe that these removal options may have more impact when applied to smaller flare areas and more investigation into this topic could be beneficial. *Raw2hdr* is recommended for future HDR compilation as the RAW file inputs allow for a more robust HDR output with slightly lower magnitudes of lens flare error. Error analysis across a more diverse set of scenes is recommended for determining the absolute extent of error associated with lens flare in HDR luminance measurements. It is also important to keep in mind while viewing HDR luminance values, particularly those displayed in MATLAB, that the weight assigned to the brightness of the values is not necessarily proportionate. False color renders on a relative scale provide for a more visually accurate representation of luminance values.

CHAPTER 4: ADAPTATION OF HIGH FREQUENCY OCCLUSION MASK FOR LUMINANCE MEASUREMENTS

The method of employing a high frequency occlusion mask at the scene during HDR capture proved to significantly improve the dynamic range of images when compared to the traditional blind deconvolution methods [23]. This method allows for light true to both the scene and source information to be accurately separated from light resulting from lens flare. While the experiment described in Chapter 2 indicated successful removal of veiling flare in HDR images, it did not identify whether the method was effective for producing accurate luminance measurements. Adapting this method specifically for the purpose of obtaining such measurements could prove to be a powerful solution to lens flare in the HDRI and lighting communities.

4.1 Methodology

4.1.1 HDRI Acquisition Procedure

Using a similar image acquisition method as described in Chapter 3, a series of images were captured using the Nikon D5200 camera in combination with the Sigma 24mm, F1.8 prime lens. The camera, tethered to a laptop for automated HDR capture, was placed 5 feet from an integrating sphere with the 60W CITILED source installed. The camera was adjusted such that the sphere's 1 inch opening was positioned in the center of the viewfinder, normal to the source.

A physical occlusion mask was constructed from black, opaque poster board with $\frac{1}{4}$ inch holes with a period of $\frac{1}{4}$ inch for a final occlusion factor of $\alpha = 0.25$. To achieve complete coverage of the scene, $\frac{1}{\alpha}$ or 4 HDR images were necessary. Unlike the original study, the mask was adjusted manually each time instead of with a mechanical tripod. To account for any error accompanying the manual adjustments, the image set was increased to 9 requiring a 3 x 3

translation of the mask. The occlusion mask was placed 6 inches from the source allowing the scene to remain almost entirely focused (Figure 36). Rulers and clamps were attached to the occlusion mask to aid in correctly translating it vertically and horizontally in 1/12 inch increments.



Figure 36: Example of setup with occlusion mask

The camera settings mentioned in Chapter 3, Table 2, (ISO, white balance, metering, etc.) were maintained for this series of tests. Seven bracketed LDR images were captured for each mask position with the shutter speed ranging from 30 seconds to 1/4000 seconds. These LDRs were then transferred into Photosphere where the final HDR images were fused. Figure 37 shows an example of a final HDR image at f/9 with the occlusion mask present in the scene.




Figure 37: HDRI image with occlusion mask

4.2 Post Processing Procedure

Once a full set of images were captured and compiled into HDRIs, a MATLAB script was used to translate RGB values to luminance values using Eq. 9 and then cropped for workability. The exposure value for each of the HDR images was acquired through HDRscope. Images taken using the same apertures were confirmed to have identical exposure values and, as a result, luminance values displayed during post-processing did not require calibration as they fell on the same relative scale.

Before any HDR images were fed through the reconstruction process, two continuous binary matrix masks were created to represent m_{ϕ}^{+} and m_{ϕ}^{-} . These masks, designed to cover the holed regions and blocked regions of the physical occlusion mask respectively, could be adjusted and rotated in any direction to extract relevant luminance information. During the capture process, it was observed that the thickness of the poster board used (1/16 inch) contributed to

small halos around the edges of the holes in the occlusion mask (Figure 38). The haloed regions returned skewed luminance values and as a result, m_{ϕ}^{+} and m_{ϕ}^{-} were specifically sized and placed to omit these regions.


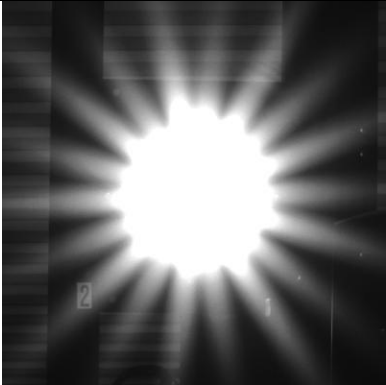
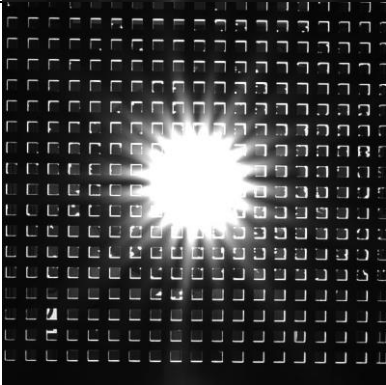
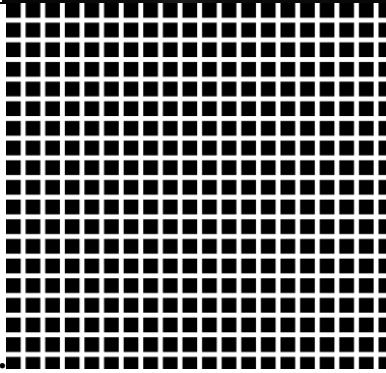
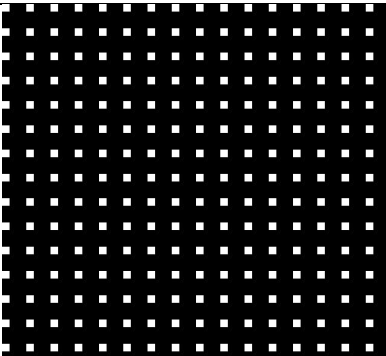


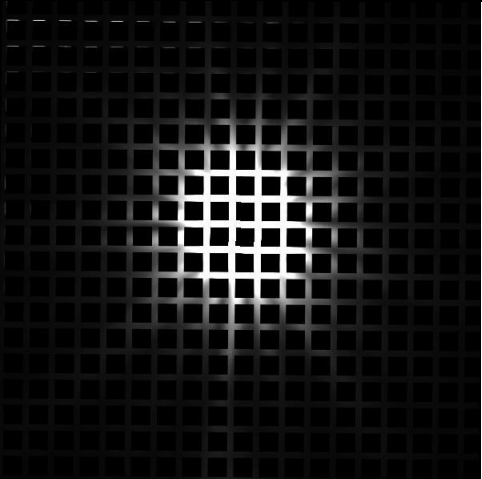

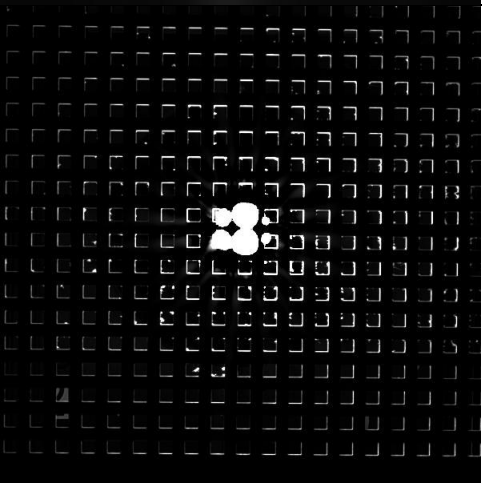
Figure 38: Halos resulting from thickness of physical occlusion mask

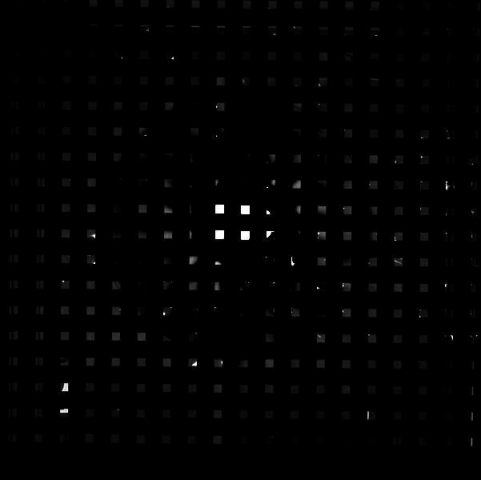

Next, a flare estimate, g_{ϕ} , was created by overlaying the binary mask m_{ϕ}^{-} atop the corrupted image, r_{ϕ} and applying a Gaussian interpolation across the zeroed pixels. The Gaussian blur kernel, f , specifies the extent of the interpolation and was set to 10 pixels (one third the mask pattern period [23]). Finally, the flare estimate was subtracted from the original corrupted image. The pixel information corresponding to the holes of the occlusion mask were then gathered to form a final glare free stitch, s_{ϕ} . It was the intent that by stitching together the relevant pixel information for all 9 HDR photographs used to capture the scene, a final image, s , would be formed. The post-processing steps are arranged sequentially in Table 8 with the luminance output visuals for each step found on the left and the mathematical process described on the right.

-

Table 8: Steps for Removing Lens Flare Using Occlusion Mask

		<p>1. CORRUPTED ORIGINAL UNBLOCKED</p> <p style="text-align: center;">o</p>
		<p>2. CORRUPTED ORIGINAL BLOCKED WITH OCCLUSION MASK</p> <p style="text-align: center;">r_ϕ</p>
		<p>3. BINARY OCCLUSION MASK WITH PIXEL INFO FOR HOLES = 0</p> <p style="text-align: center;">m_ϕ^-</p>
		<p>4. BINARY OCCLUSION MASK WITH PIXEL INFO FOR HOLES = 1</p> <p style="text-align: center;">m_ϕ^+</p>

	<p>5. FLARE ESTIMATE</p> $\widehat{g}_\phi = m_\phi^- \times r_\phi$
	<p>6. FLARE ESTIMATE FOR ENTIRE SCENE USING GAUSSIAN INTERPOLATION</p> $g_\phi = \frac{f * (m_\phi^- \times r_\phi)}{f * m_\phi^-}$ <p>f = Gaussian blur kernel (10 pixels)</p>
	<p>7. FIXED IMAGE WITH OCCLUSION MASK</p> $(r_\phi - g_\phi)$

	<p>8. FIXED IMAGE SHOWING ONLY PIXEL INFORMATION FOR MASK HOLES</p> $s_{\phi} = m_{\phi}^+ \times (r_{\phi} - g_{\phi})$ <p>Steps 2-8 are repeated 8 more times, translating the screen such that there are s_{ϕ} values for every pixel in scene</p>
	<p>9. FIXED IMAGE, STITCHED TOGETHER WITH ALL s_{ϕ}</p> <p style="text-align: center;">S</p>

4.3 Preliminary Results and Conclusions

Positioning the high frequency occlusion mask, between the source and camera, successfully identifies which pixels contain luminance values true to the scene and which have values introduced through the optical elements within the lens. Because the occlusion mask does not have any light hitting the surface seen by the camera, it should be easily distinguished against the source. However, this is not the case as false luminance values bleed across the blocked pixels. Step 5 of Table 8 identifies exactly which pixels contain luminance values that are not true to the scene and can then be used to create a glare estimate.

While this method has the potential to be highly effective in theory, there were several complications with its adaption for luminance measurements resulting from the scene being captured in this particular experiment. In previous research implementation, the occlusion mask was proven to increase the dynamic range in images suffering from veiling glare and was not as successful in regions where there was sharp contrast between luminances [23]. This limitation remained true for this experiment as the single source in the blackened room exhibited a sharp contrast. It proved to be problematic during step 6 of Table 8, during the Gaussian interpolation, as the luminance values of the pixels near the edge of the source were overestimated. Once the interpolated glare estimate was subtracted from the corrupted original, the overestimated pixel values near the edge of the source returned negative luminance values which were known to be fundamentally incorrect and were immediately discarded.

The stitching process also proved challenging and failed to produce an entirely accurate fixed image. It was somewhat difficult to shift the screen the appropriate amount ($1/12^{\text{th}}$ of an inch) between each HDRI capture. For this particular experiment, the screen was translated manually but it would have been beneficial to employ a mechanical tripod for the translations. The final stitched image bore a resemblance to the scene being captured but minute inaccuracies when adjusting the occlusion mask were visible throughout the image. Increasing the number of HDR captures used to stitch the flare-free estimate is another alternative to ensure that every pixel will have luminance information for the scene at least once. However, in increasing the number of HDRIs, the capture process becomes much lengthier and begins to accumulate a cumbersome set of data.

The final results using the methods in this experiment were poor and do not offer a sufficient solution to mitigate lens flare. It can be concluded that the introduction of a high

frequency occlusion mask fails in scenes of extreme luminance contrast to distinguish between those values that are real and those which are an artificial byproduct of the camera/lens system. This method is not recommended for applications such as near-field photometry or glare analysis as such applications are prime candidates for extreme contrast in luminances. Although it may be difficult to implement, use of the occlusion mask is best suited for scenes with veiling glare resulting from diffuse light sources such as windows.

4.3 Reconstructed Methodology

While the high frequency occlusion mask method did not yield a successful stitched image, it provided a promising step in the right direction. What may be more important to focus on moving forward is the development of an accurate flare estimate which may be subtracted from the initial corrupted image (Step 5 of Table 8). Moving from Step 5 to Step 6 in Table 8 is where the occlusion mask method begins to break down due to the large areas of zeroed pixels which the Gaussian interpolation is intended to estimate. Within these zeroed pixel areas are sharp edges between source and scene and an interpolated estimate over these pixels results in overestimations or underestimations in luminance values.

Rearranging and adjusting the post-processing steps detailed in Table 8 offers a more versatile solution for the scene being examined. Instead of stitching together nine “fixed” images to form a single “fixed” image, the stitching process was moved between Steps 5 and 6 of Table 8. By stitching together the flare estimates for all nine HDR photographs, a full flare estimate may be established. The Gaussian interpolation is still implemented afterwards but the size of the blur kernel is cut from 10 to 5 pixels as the interpolation is no longer responsible for estimating across zeroed pixels. Instead, the interpolation with decreased blur kernel (f) is used to determine intermediate values between pixels to reduce the severity of periodic artifacts (Figure 39, right).

Finally, the flare estimate is directly subtracted from the original corrupted HDR image to produce a “flare-reduced” final HDR image.




Figure 39: Left: Example of single flare estimate, Center: Stitched flare estimate, Right: Final flare estimate after Gaussian interpolation to reduce period artifacts

4.4 Restructured Results and Conclusion

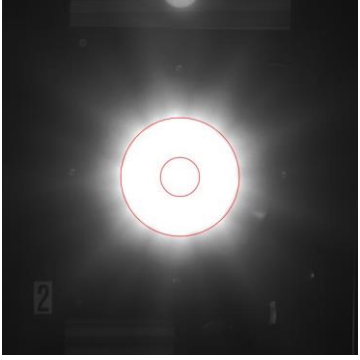
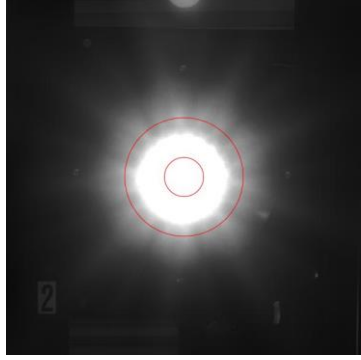
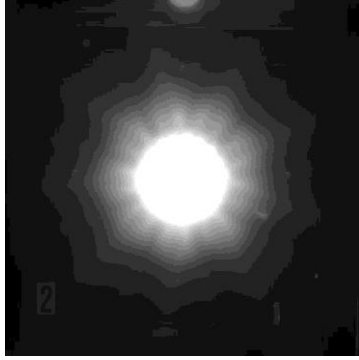
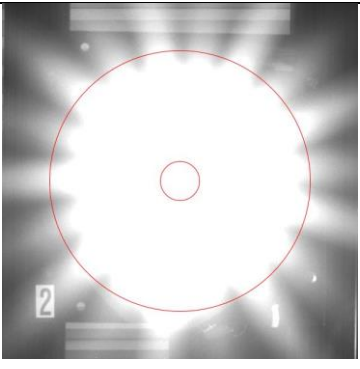
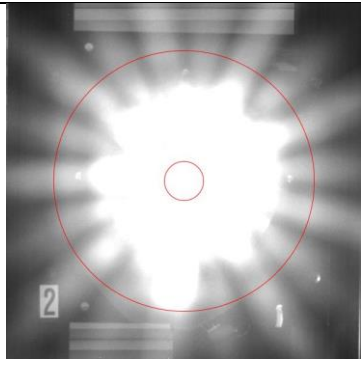

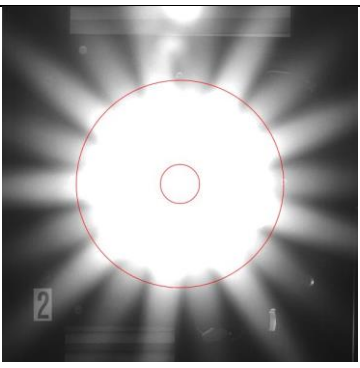
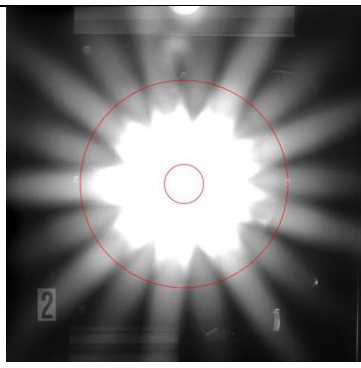
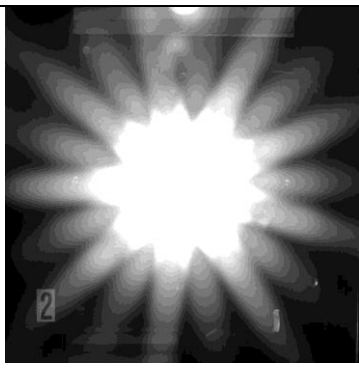
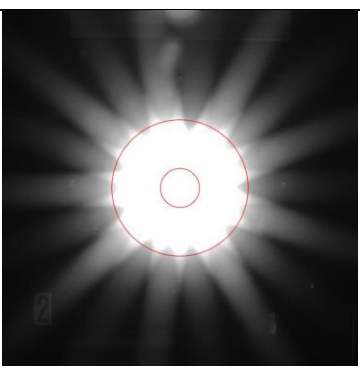
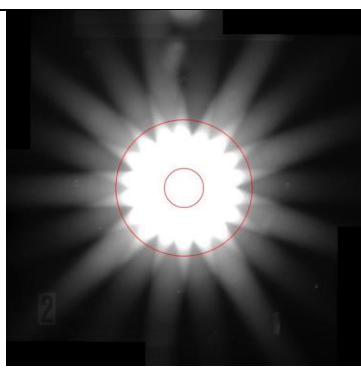
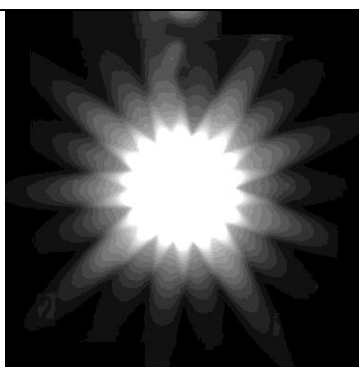
The results of this study, while not without their limitations, offer a step in the right direction for correcting for lens flare in HDR images. Once the flare estimate was directly subtracted from the corrupted original, the pixel diameter of the body of the flare was measured. All apertures experienced considerable reduction in the body of the flare ranging from 32% (f/16) to 56% (f/11) (Table 9 and Table 10). The outline of the source and the extent of the original flare were provided to act as a visual aid in the comparison of corrupted and “fixed” images.

Table 9: Pixel Measurements for Corrupted and Corrected HDRIs

f/#	Corrupted Image		Fixed Image		Pixel Area Reduction	Percent Reduction
	Pixel Diameter	Pixel Area	Pixel Diameter	Pixel Area		
1.4	226	156,612	187	106,010	50602	33%
3.5	246	186,268	182	100,214	86055	47%
5.6	200	121,815	155	71,628	50187	42%
9	421	552,971	326	330,027	222943	40%
11	350	380,997	235	169,646	211351	56%
16	218	145,453	182	100,214	45239	32%

Table 10: Comparison of Corrupted and Flare Reduced HDRIs by Aperture

f/#	Corrupted Image with Source Outlined	Fixed Image with Source and Original Corruption Outlined	Flare Reduced Image
1.4			
3.5			

5.6			
9			
11			
16			

The body of the flare experiences significant reduction but the “legs” of the flare often remain. This minor limitation is explained by the introduction of the occlusion mask to the scene. Once the mask is placed directly in front of the source, it attenuates the amount of light allowed through reducing the length of the legs of the flare as well as the area of the body. Because of the attenuation, all “fixed” HDRIs are considered to be underestimates. Moreover, the legs of the flare, while visually hindering, are recorded as having an error of less than 1% on average.

Unfortunately, there were no significant trends in flare reduction correlating to the change of camera aperture. In most instances, the reduction process left the flare symmetrical. However, in the case of the $f/9$ aperture, the reduction left the flare seemingly less manageable than the corrupted image.

CHAPTER 5: CONCLUSION

5.1 Study Conclusion

Perhaps the most significant finding of this study was the magnitude of error associated with lens flare in HDR images. Lens flare has been cited as a recurrent visual impediment found in instances of high luminance contrast between source and background. However, when examining the error histograms associated with lens flare measurements for the simple scene used in this report, less than 1% of the pixel luminance values exceed an error of 10% for the pixel area considered. Of the few pixels that experience luminance measurement error over 10%, their location is consistently within 5-35 pixels radially from the edge of the source. Of the various errors associated with HDR photography as a luminance measurement tool, lens flare may not be the most problematic.

Nonetheless, there has been expressed interest by industry professionals and those in the HDRI community to mitigate the visual imperfections resulting from lens flare while still maintaining accurate luminance values. It was therefore valuable to evaluate the effectiveness of the currently available methods for lens flare removal offered. The most easily implemented methods are the blind deconvolution algorithms found in Photosphere and *raw2hdr*. Based on the simple scene used in this experiment, the deconvolution algorithms failed to provide any increased accuracy in luminance value. There is reason to believe that both removal methods are more successful with increased distance between the camera and the source as the flare occupies a much smaller area of the sensor. The flare removal options in both Photosphere and *raw2hdr* should not be relied on as a catch-all solution for most applications especially for the use of near-field photometry.

In hopes to offer an alternative solution to lens flare removal, the proven method of introducing a high frequency occlusion mask was tested. The initial tests using the occlusion mask failed due to the characteristics of the scene. The method had been proven to work in scenarios with veiling glare, however the sharp defined edges of the circular source in this scene proved to be undetectable through the series of calculations. While this method did not work for this experiment, adapting its methodology allowed for an alternative solution to be developed. The new solution aimed to create a flare-free estimate which was then subtracted from the corrupted original HDRI. This method was able to reduce the area pertaining to the body of the flare as much as 56% depending on the aperture setting.

The high frequency occlusion mask, while exhibiting promising results, is not without its limitations. The attenuation of light caused by placing the occlusion mask in front of the source produces a conservative flare estimate. As a result, the legs of the flare remain as a visual impedance. This is not as problematic for luminance measurement purposes as the legs have been shown to have error of less than 1%. There were no significant trends in flare reduction in relation to the change of camera aperture although the smaller apertures maintained a sharper depth of field. The sharper images made it easier to manage the occlusion mask and the final results indicated increased focus than the larger apertures ($f/1.4$ and $f/3.5$).

One of the largest drawbacks of using the frequency occlusion mask is the amount of time and data required to amend a single corrupted image. Even if the results yield a perfect solution, it is not necessarily practical to carry out this process for all instances of lens flare. The duration of capturing and processing a single set of 16 HDR images took roughly an hour. Improved accuracy may be achieved through increasing the number of captures and mask translations however, in doing so, the amount of time needed also increases. Additionally, the

mask is not realistically applied to complex architectural scenes like those found in glare analysis studies. The occlusion mask must be placed near the source of the flare to keep the camera properly focused and therefore a mask applied to the camera like a filter would not be sufficient.

It is important to acknowledge that the conclusions resulting from this study are specific to the simple scene used for HDRI capture. The goal of this research was to provide a comparison of relevant removal techniques to aid in long term lens flare removal rather than an absolute quantification of the error associated with lens flare. Based on the results of this report, lens flare poses the largest problem in applications of near-field photometry and the current correction methods, particularly the blind deconvolution algorithms found in current HDRI software, can not necessarily be relied upon as a perfect solution. If there is considerable interest in eliminating lens flare entirely, a more dynamic solution than those discussed in this study is necessary. Otherwise, for more generic applications, such as glare analysis, the error associated with lens flare in a simple scene is insignificant enough to justify using luminance values as true. More work should be carried out with more complex scenes to confirm the low magnitude of error associated with lens flare.

5.2 Recommendations for Future Research Using Neural Network

Future research surrounding lens flare removal in HDR photography may be better served by taking a step away from the correction methods evaluated in this report and instead considering the potentials of implementing a CNN. While the complexity of developing a neural network exceeds the scope of this project, there has been proven success in the removal capabilities of CNNs applied to images exhibiting noise. With sufficient motivation and knowledge, a more robust solution to lens flare could be developed with a network architecture

similar to that executed in a NYU study [26] correcting for localized noise. The goal of a working CNN would be to detect the presence and magnitude of lens flare given a single, corrupted HDR image as input. The network would then offer an educated correction for the corrupted pixels and output a “clean” image. The use of a blind inpainting technique has been proven to identify and remove corruption patterns from an image without knowing the specific region of corruption [26] and a similar method is recommended for future lens flare correction.

To achieve a functional network, a large number of training images would be required. These images would depict a simple scene with the presence of lens flare and a corresponding image of the same scene with the absence of the flare. One conceivable method for obtaining a large, diverse image training set is through the use of an optical design software capable of simulating the diffraction effects associated with lens flare (Photopia is one such software for lens simulation [28]). An optical model of the specific camera lens assembly would need to be created based on the lens specifications retrieved from the manufacturer. Such specifications should detail the antireflection properties applied to each of the lens mediums. An identical model should be made depicting a theoretically perfect lens without any diffraction effects present. A series of simulations would be carried out for a given scene by first running the imperfect lens construction followed by running the perfected construction.

The diversity of the training set could be increased by altering the scene slightly across simulations. The scene should begin simple with a uniformly luminous disk to isolate the flare effects, similar to the integrating sphere used for this report. The lens flare in the images should start small, only impacting a few number of pixels (less than 100 x 100). If the simple training images prove to be successful for developing a neural network, the complexity of the data set can increase to include images across varying apertures, source sizes, source shapes, source

magnitude, and increased background information. An increased level of diversity in training images will result in a more accurate correction system. Considerable effort allocated to the development of an intelligent network could allow for a long-term solution to lens flare easily accessed and implemented by a number of HDRI users. The success of a simple network could eventually lead to the tailoring of similar networks to correct lens flare found in a variety of unique architectural scenes and lighting applications.

REFERENCES

- [1] Nikon, "Nikon," [Online]. Available: <http://www.nikonusa.com/>.
- [2] Sigma. [Online]. Available: www.sigmaphoto.com/lenses.
- [3] DiLaura and D. L. e. al., *The Lighting Handbook: Reference and Application*, 2011.
- [4] B. Horn, *Robot Vision*, Cambridge: MIT, 1986.
- [5] M. J. Stanley, "On the development and error analysis of a high dynamic range imaging system for luminance measurements," s.l. : University of Colorado Boulder, 2016.
- [6] K. Van Den Wemelenberg and M. Inanici, "Evaluating a New Suite of Luminance-Based Design Metrics for Predicting Human Visual Comfort in Offices with Daylight," *LEUKOS*, pp. 113-138, 2015.
- [7] J. A. Jakubiec and e. al., "Accurate measurement of daylit interior scenes using high dynamic range photography," in *CIE 2016 "Lighting Quality & Energy Efficiency"*, Melbourne, 2016.
- [8] K. Van Den Wymelenberg and M. Inanici, "A Critical Investigation of Common Lighting Design Metrics for Predicting Human Visual Comfort in Offices with Daylight," *LEUKOS*, pp. 145-164, 2014.
- [9] J. B. Murdoch, *Illuminating Engineering: From Edison's Lamp to the LED.*, Second ed., s.l.: Visions Communications, 2003.
- [10] Y. I. Tyukhova, "The assessment of HDR luminance measurements with LED Lighting," University of Nebraska, s.l., 2012.
- [11] S. Anaokar and M. Moeck, "Accuracy of lumiannce maps obtained from high dynamic range images," *LEUKOS*, vol. 4, no. 2, pp. 99-1122, 2007.
- [12] S. Anaokar and M. Moeck, "Validation of High Dynamic Range Imaging to Luminance Measurement," *LEUKOS*, pp. 133-144, 2005.
- [13] K. L. Teman, "Method for quantifying the spectral based error in luminance measurements," s.l. University of Colorado Boulder, 2017.
- [14] M. Inanici, "Evaluation of high dynamic range photography as a luminance mapping technique," Lawrence Berkeley National Laboratory, s.l., 2004.
- [15] M. Moeck and S. Anaokar, "Illuminance Analysis from High Dynamic Range Images," *LEUKOS*, vol. 2, no. 3, pp. 211-228, 2006.

- [16] R. Kingslake, *Optics in Photography*, Bellingham: SPIE Optical Engineering Press, 1992.
- [17] H. Keppler, "Lens Flare, Ghosts and Falloff," *Popular Photography*, 1997.
- [18] F. Xiao, J. M. Di Carlo, P. B. Catrysse and B. A. Wandell, "High Dynamic Range Imaging of Natural Scenes," *Conference Paper*, 2002.
- [19] "Hyperphysics," [Online].
- [20] R. Kingslake and B. Johnson, *Lens Design Fundamentals*.
- [21] P. Boynton and E. Kelley, "A Liquid-Filled Camera for the Measurement of High-Contrast Images".
- [22] E. Reinhard, G. Ward, S. Pattanaik, P. Debevec, W. Heidrich and K. Myszkowski, *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, Burlington: Morgan Kaufmann, 2010.
- [23] E.-V. Talvala, A. Adams, M. Horowitz and M. Levoy, "Veiling Glare in High Dynamic Range Imaging".
- [24] S. Nayar, G. Krishnan, D. Grossberg and R. Raskar, "Fast Separation of Direct and Global Components of a Scene using High Frequency Illumination," *ACM SIGGRAPH 2006 Papers*, pp. 935-944, 2006.
- [25] K. O'Shea and R. Nash, "An Introduction to Convolutional Neural Networks," 12 2 2015. [Online]. Available: https://white.stanford.edu/teach/index.php/An_Introduction_to_Convolutional_Neural_Networks. [Accessed 20 6 2017].
- [26] D. Eigen, D. Krishnan and R. Fergus, "Restoring An Image Taken Through A Window Covered with Dirt or Rain," in *ICCV*, Sydney, 2013.
- [27] V. Kumaragurubaran, "High Dynamic Range Image Processing Toolkit for Lighting Simulations and Analysis," University of Washington , 2012.
- [28] "Photopia for Solidworks," LTI Optics, [Online]. Available: <http://www.ltioptics.com/en/photopia-general-2017.html>.
- [29] A. Jacobs, "High Dynamic Range Imaging and Its Application in Building Research," *Advances in Building Research*, vol. 1, no. 1, pp. 177-202, 2007.
- [30] M. Aggarwal, H. Hua and N. Ahuja, "On Cosine-fourth and Vignetting Effects in Real Lenses," in *IEEE International Conference on Computer Vision*, 2001.

- [31] C. Cauwerts, M. Bodart and A. Denever, "Comparison of the Vignetting Effects of Two Identical Fisheye Lenses," *LEUKOS*, vol. 8, no. 3, pp. 181-203, 2012.
- [32] M. Laikin, *Lens Design*, Fourth Edition, Boca Raton: CRC Press, 2007.
- [33] M. Jongewaard, "LTI Optics," [Online]. Available: Mark.Jongewaard@colorado.edu.

APPENDIX A: Vignetting Correction Documentation

A vignetting correction experiment, prior to the work done for this report, was conducted by University of Colorado students Katie Teman, Sarah Safranek, and Shelby White. The goal of this experiment was to obtain a method for reversing the effects of vignetting through identifying the unique vignetting function for a camera/lens combination. This function can be inversely applied to the luminance value at every pixel in an HDR image to counter the radial fall off of light towards the edges of the image. The following writeup is to act as supporting documentation for the vignetting correction procedure. The correction methods detailed here are the same used to reverse vignetting effects in Chapter 3 and 4 of this study as well as other University of Colorado HDRI studies [13].

Introduction

High dynamic range images (HDRI) were created to capture a scene more closely to how a human eye would truly experience that setting. The human visual system is able to adapt to light levels ranging from 12-14 order of magnitude, and up to 5 orders of magnitude in a single scene [29]. Digital camera sensors are only able to capture a range of 3-4 orders of magnitude of light levels at a time [29]. The HDR process extends this range by taking multiple photographs from the same viewpoint with over and underexposed settings, then combining the accurately exposed portions of each image. The compiled HDRI ensures that both details in shadowed regions, as well as extremely bright subjects, are able to be seen.

The HDRI photography technique has recently been adapted for lighting applications to attain luminance values at a high resolution very quickly. Potential applications of this method for the lighting industry include luminance and glare analysis, sky model creation for daylighting, and near-field photometry [5]. A reference point of known luminance is used to

calibrate the relative luminance of each pixel to absolute measurements. While there already exist multiple software options to make this tool available to lighting designers, manufacturers, and researchers alike, there are still sources of error inherent to the HDRI technology that affect the accuracy of luminance values derived from this method. Lens flare, saturated colors, and vignetting are all known to produce errors in HDRI luminance measurements, but the severity of these errors is not yet entirely understood. This goal of this research is to better understand vignetting specifically and to correct for this effect in HDR images.

Vignetting is the fall-off of brightness towards the edge of an image and is due to the optical process of camera lenses [5]. Light rays that enter the lens off-axis are unable to fill the entire aperture which results in fewer photons reaching the edges of the camera sensor. Two similar effects are the cosine-fourth effect and pupil aberration. While the technicalities of these effects differ slightly from vignetting, the results of all three are often classified as one issue simply referred to as vignetting [5]. The severity of this effect is specific to a particular combination of camera, lens, and aperture. Past research has demonstrated that brightness fall-off towards the edge of an image increases with increased lens apertures [30] [31].

A recent study at the University of Colorado Boulder developed an HDRI system and began to calibrate and test the accuracy of the resulting HDRI luminance values [5]. An integrating sphere with even illumination was photographed to study the effects of vignetting on luminance measurements. Multiple HDR software packages for image compilation and analysis were compared to select which software would work best with this new system. After extensive evaluation, hdrngen was selected as the ideal software for Mac users to compile JPEG images and photosphere was recommended for HDR image analysis [5]. For Windows operating system

users, WebHDR and HDRscope were also considered acceptable options for image compilation and analysis respectively [5].

Methodology

Experimental Process

For HDR image capture, a Nikon D5200 with CMOS sensor was used in combination with three different lenses; the Sigma Prime lens with fixed zoom, the Nikkor Zoom 18mm, and the Nikkor Zoom 24mm. An integrating sphere, provided by Mark Jongewaard at LTI Optics, was used to ensure constant luminance levels were maintained within each of the lens' field of view. To confirm the sphere's luminance levels, a luminance meter mounted on a tripod was pointed directly at the center of the sphere. Because the sphere is meant to be a perfectly diffuse surface, the values captured using the luminance meter should be the same regardless of the angle at which the meter was directed. To ensure this, luminance measurements were recorded starting at the center of the sphere and moving upwards and then downwards along the sphere's centerline. Average luminance of the sphere was found to be 8.78 cd/m^2 within an error of 1%.

Once constant luminance levels were confirmed, the luminance meter tripod was replaced with a camera tripod, mounted at the same height and distance (8.5'' from the edge of the sphere to the flash of the camera.) An example of this setup can be seen in Figure 40. The camera system was connected via USB cable to a laptop for remote camera control utilizing the tethering software, *Sofortbild*. The tethering setup minimizes movement during image capture and allows for HDR image bracketing to be controlled directly on the laptop. For each lens/camera combination, the recommended settings were used based on a study conducted by Stanley at University of Colorado [5]. Between four to six apertures were carried out for photograph

capture for a given lens (also based on the previous work done by Stanley). After all images were captured, the luminance meter mounted tripod was put back into place and luminance measurements were recorded again for the sphere in the same manner mentioned previously.



Figure 40: Integrating sphere set-up with camera in place

Vignetting Correction Process

To counteract the vignetting effect in the images, it was first necessary to find the unique “vignetting function” for each lens/aperture combination. Each of the compiled HDR images correlated to a specific lens/aperture combination. Because the images were taken inside of an integrating sphere, they are theoretically perfectly uniform, thus, any deviance in luminance that can be seen in the image is a result of the vignetting effect caused by the camera. These images

were used to determine the vignetting function, the mathematical function that describes the rate at which the image brightness drops off radially from the center of the image to the edge of the image. When the specific lens/aperture function is inversely applied to the image, the vignetting effect can be negated. This vignetting function can be applied to any image to counteract the vignetting effect so long as the image was taken with the same lens/aperture combination.

To find the vignetting functions for each lens/aperture combination, each image was passed through a custom MATLAB script. The MATLAB script first read in the image, calculated the image bounds and located the center of the image (recall that the center of the image is where the image is brightest and no vignetting is occurring). Then, the script located specific pixels radially from the center of the image in a “pinwheel” type distribution and calculated the luminance corresponding to each point.

The pixel locations could be varied by altering the distance between each “spur” of the pinwheel and the radial distance between each point on each spur. The more points sampled, the more accurate the eventual function would be, but the longer it would take to process the image. After optimizing the number of points needed, the image was processed and the polar locations and luminance value data at each point were recorded. Then, a curve of best fit was applied to each “spur” determined through polynomial curve fitting functionality in MATLAB. In theory, since the surface should have been perfectly diffuse, the vignetting effect should have been even across the image. This mostly held true for our experiment as well (validating the credibility of the integrating sphere process) with minor error likely resulting from imperfections in the diffuse coating of the sphere and/or imperfections in the lens of the camera. The functions found for each spur were averaged together to get a single best fit function for the entire image. This became the vignetting function which, once applied to the

original image, substantially reduced the vignetting effect present throughout. Any remaining irregularities are considered to be negligible for most luminance mapping purposes.

Table 11: Vignetting Correction Functions for Sigma Prime Lens

f/number	Vignetting Correction Function – Sigma Prime Lens
1.4	$1.59x^6 - 3.81x^5 + 2.52x^4 + 0.16x^3 - 0.78x^2 - 0.20x + 1.00$
3.5	$2.77x^6 - 6.94x^5 + 5.88x^4 - 0.85x^3 - 0.03x^2 + 0.02x + 1.00$
5.6	$2.06x^6 - 4.86x^5 + 3.58x^4 - 0.66x^3 - 0.32x^2 + 0.65x + 1.00$
9	$2.74x^6 - 6.63x^5 + 5.22x^4 - 1.23x^3 - 0.30x^2 + 0.05x + 1.00$
11	$1.97x^6 - 4.70x^5 + 1.57x^4 - 1.78x^3 - 1.02x^2 + 0.10x + 1.00$
16	$0.31x^6 + 0.59x^5 - 2.84x^4 - 2.94x^3 - 1.30x^2 + 0.16x + 1.00$

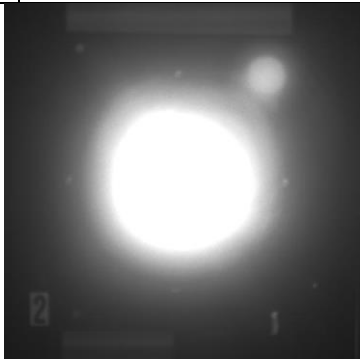

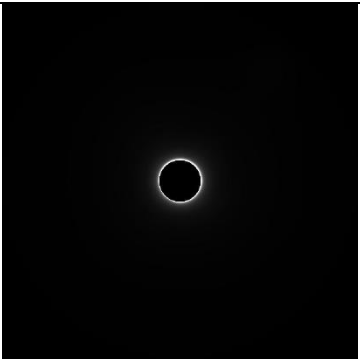
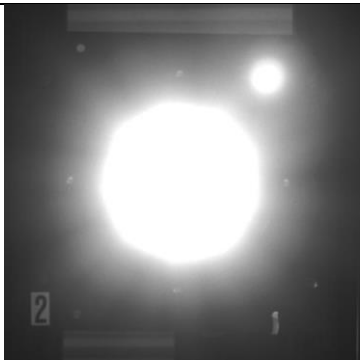
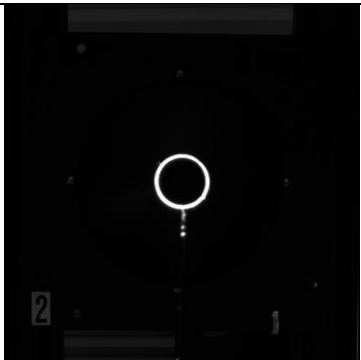
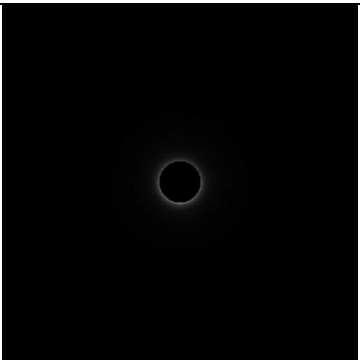


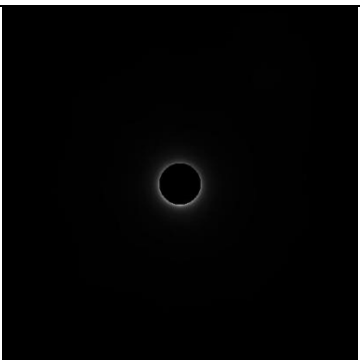
Table 12: Vignetting Correction Functions for Zoom 18mm Lens

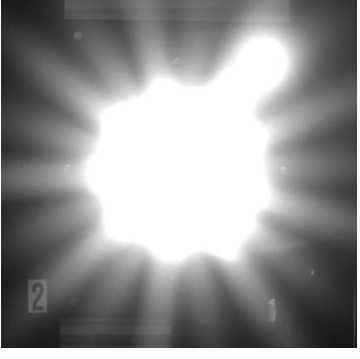


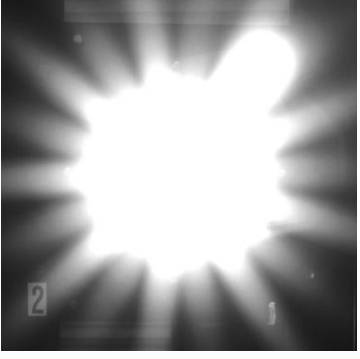


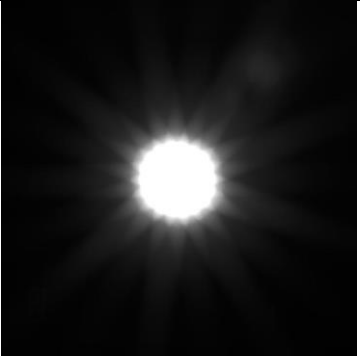
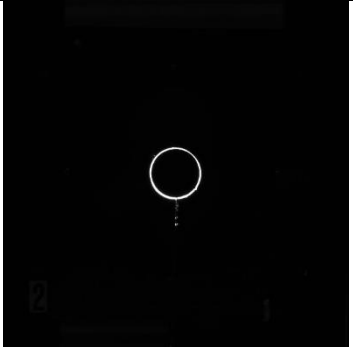
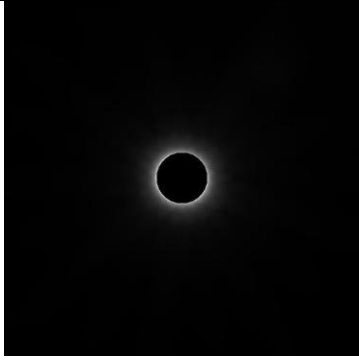
f/number	Vignetting Correction Function – Zoom 18mm Lens
3.5	$-14.5x^6 + 38.39x^5 - 37.11x^4 + 15.59x^3 - 3.122x^2 + 0.20x + 1.00$
5.6	$-7.516x^6 + 21.14x^5 - 22.27x^4 + 10.74x^3 - 2.599x^2 + 0.1995x + 1.00$
9	$-3.585x^6 + 11.74x^5 - 14.35x^4 + 8.018x^3 - 2.293x^2 + 0.265x + 1.00$
14	$-4.138x^6 + 10.22x^5 - 9.177x^4 + 3.732x^3 - 0.9205x^2 + 0.0621x + 1.00$
22	$-2.255x^6 + 7.09x^5 - 8.272x^4 + 4.376x^3 - 1.283x^2 + 0.1001x + 1.00$

Table 13: Vignetting Correction Function for Zoom 24mm Lens



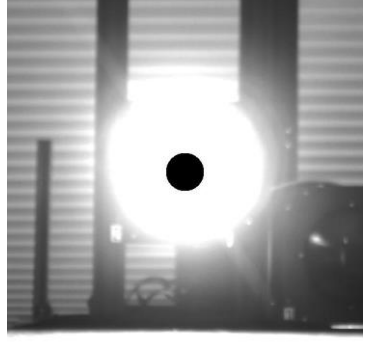


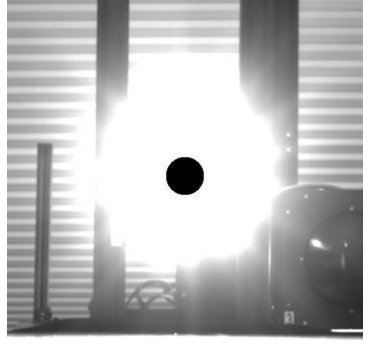


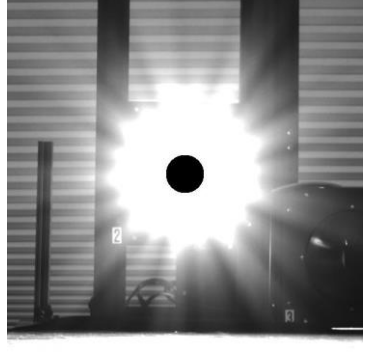


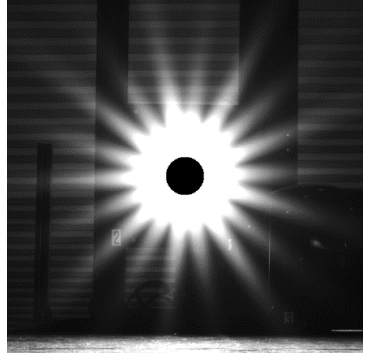
f/number	Vignetting Correction Function – Zoom 24mm Lens
3.5	$-1.72x^6 + 3.683x^5 - 1.257x^4 - 1.666x^3 + 0.7213x^2 - 0.1662x + 1.00$
5.6	$3.104x^6 - 9.896x^5 + 11.77x^4 - 6.803x^3 + 1.849x^2 - 0.2955x + 1.00$
14	$1.551x^6 - 7.570x^5 + 6.720x^4 - 2.761x^3 + 0.4122x^2 - 0.080x + 1.00$
22	$0.825x^6 - 1.362x^5 + 0.3823x^4 + 0.1987x^3 - 0.1988x^2 - 0.041x + 1.00$



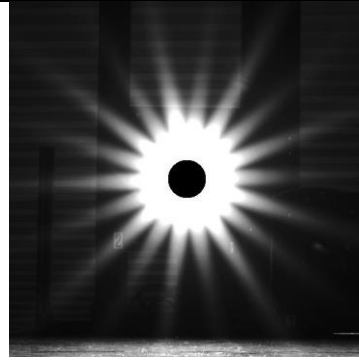
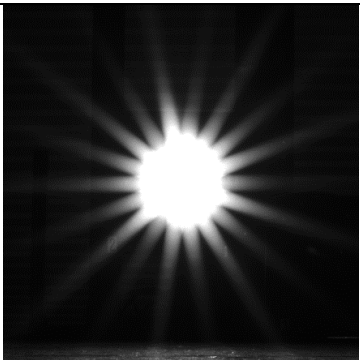
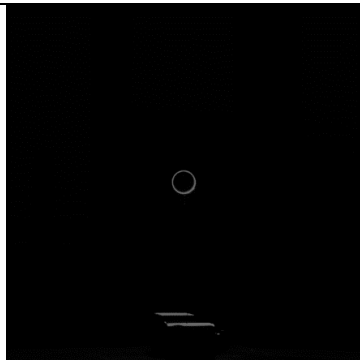
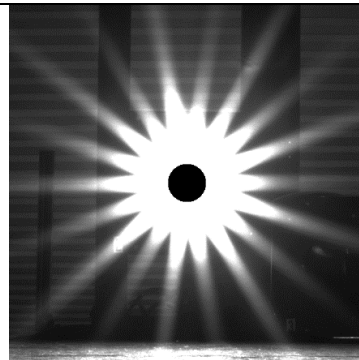
APPENDIX B: Luminance Images

LUMINANCE VISUALS, 5ft			
f/#	With Flare	Flare Blocked	Flare Normalized
1.4	 A large, bright, circular luminance spot with a soft, glowing halo and a smaller secondary spot to the upper right. A small '2' is visible in the bottom left corner.	 A dark image showing a bright white ring with a small vertical line extending downwards from its center. A small '2' is visible in the bottom left corner.	 A dark image showing a bright white ring with a small vertical line extending downwards from its center. A small '2' is visible in the bottom left corner.
3.5	 A large, bright, circular luminance spot with a soft, glowing halo and a smaller secondary spot to the upper right. A small '2' is visible in the bottom left corner.	 A dark image showing a bright white ring with a small vertical line extending downwards from its center. A small '2' is visible in the bottom left corner.	 A dark image showing a bright white ring with a small vertical line extending downwards from its center. A small '2' is visible in the bottom left corner.
5.6	 A large, bright, irregularly shaped luminance spot with a soft, glowing halo and a smaller secondary spot to the upper right. A small '2' is visible in the bottom left corner.	 A dark image showing a bright white ring with a small vertical line extending downwards from its center. A small '2' is visible in the bottom left corner.	 A dark image showing a bright white ring with a small vertical line extending downwards from its center. A small '2' is visible in the bottom left corner.

9			
11			
16			

LUMINANCE VISUALS, 10ft

f/#	With Flare	Flare Blocked	Flare Normalized
1.4	 A dark scene with a bright circular light source in the center, creating a soft, glowing flare.	 The same scene as the first column, but the light source is blocked, leaving a dark spot with a small white ring.	 The scene with the flare normalized, showing a bright white circle with a black center, surrounded by a soft glow.
3.5	 A dark scene with a bright circular light source, showing more defined rays than the f/1.4 image.	 The same scene as the first column, but the light source is blocked, leaving a dark spot with a small white ring.	 The scene with the flare normalized, showing a bright white circle with a black center, surrounded by a soft glow.
5.6	 A dark scene with a bright circular light source, showing distinct rays emanating from the center.	 The same scene as the first column, but the light source is blocked, leaving a dark spot with a small white ring.	 The scene with the flare normalized, showing a bright white circle with a black center, surrounded by a soft glow.
9	 A dark scene with a bright circular light source, showing very distinct and sharp rays.	 The same scene as the first column, but the light source is blocked, leaving a dark spot with a small white ring.	 The scene with the flare normalized, showing a bright white circle with a black center, surrounded by a soft glow.

11			
16			

APPENDIX C: Luminance Error Data

