A NEW SPATIO-ANGULAR LIGHT-FIELD DISPLAY FOR IMPROVING WELLBEING IN WINDOWLESS SPACES

by

Robin Atkins

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF

THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES

(Interdisciplinary Studies)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

March 2020

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the dissertation entitled:

A NEW SPATIO-ANGULAR LIGHT-FIELD DISPLAY FOR IMPROVING WELLBEING IN WINDOWLESS SPACES

submitted by Robin Atkins in partial fulfillment of the requirements for

the degree of Doctor of Philosophy

in Interdisciplinary Studies

Examining Committee:

Lorne Whitehead, Professor, Physics and Astronomy, UBC
Co-supervisor

Ronald Rensink, Professor, Psychology and Computer Science, UBC
Co-supervisor

Rabab Ward, Professor, Electrical and Computer Engineering, UBC
Supervisory Committee Member

Karen Bartlett, Professor, Population and Public Health, UBC
University Examiner

Steven Rogak, Professor, Mechanical Engineering, UBC
University Examiner

Additional Supervisory Committee Members:

Supervisory Committee Member

Supervisory Committee Member

Abstract

This project was motivated by the desire to improve windowless indoor environments by providing an enhanced sense of space to the occupants. Photographs of distant landscapes help surprisingly little. The approach taken for this project was to establish the requirements of a display system capable of providing a convincing sense of depth, while being sufficiently affordable for widespread use. A significant breakthrough in this project was the discovery that high resolution is neither practical nor necessary to do this. This led to the invention of a novel intermediate-resolution display capable of being manufactured using modern manufacturing technology. The research to validate the design followed two separate, synergistic paths: One was a psychophysical study to verify that human perception can accurately determine depth from somewhat blurry scenes. The study confirmed that this was indeed the case, finding that an angular diffusion, or point spread function, of 2.2 degrees can produce scenes with as many as 31 individual depth planes. The other path involved the invention of a novel display system capable of producing intermediate-resolution light fields. Optical simulations indicate that the design can produce images with the desired horizontal and vertical angular diffusion (about 3.4 degrees). In combination, these two results have established the feasibility of an affordable intermediate-resolution display system that could produce an experience indistinguishable from viewing a distant scene through a partially diffusing window.

Lay Summary

The research goal was to determine if it is possible to create, in a practical way, the visual sensation of depth provided by a window overlooking a distant space. This was achieved by developing a design for an affordable display capable of showing a scene indistinguishable from that obtained by looking through a partially diffusing window overlooking a distant scene. The results of this project can help inform the design of indoor environments such as shopping malls, working spaces, and living spaces, especially in underground and other windowless locations, in such a way as to provide the occupants with a greater sense of space than would otherwise be experienced.

Preface

This dissertation is an original intellectual product of the author, R. Atkins. The user study in Chapter 3 was designed, conducted, and analyzed by me. The novel optical design in Chapter 4 was invented by both L. Whitehead and I. The simulations and analysis of the optical design in Chapter 4 are my original work.

A description of the user study of Chapter 3 and the optical design of Chapter 4 has been published: Robin Atkins, and Lorne Whitehead (2019), Journal of the Society for Information Display. Robin Atkins conducted all the research, testing, design, and wrote the manuscript.

The user study described in Chapter 3 was conducted under approval of the UBC Ethics Review Committee, approval number H17-02613.

Table of Contents

Abstract iii

Lay Summary iv

Preface v

Table of Contents vi

List of Tables ix

List of Figures x

Glossary xvi

Chapter 1: Introduction 1

1.1 Motivation and Outline 2

1.2 Basic Concept 4

1.3 Contributions 5

Chapter 2: Background 6

2.1 Fundamentals of Depth Perception 7

2.1.1 Monocular Depth Cues 7

2.1.2 Binocular Depth Cues 10

2.1.3 Observer Motion Depth Cues 12

2.1.4 Enhancement and Interference of Depth Cues 13

2.1.5 Summary of Depth Cues 15

2.2 Mechanisms for Binocular Parallax and Motion Parallax 17

2.2.1 Binocular Parallax 17

2.2.2 Motion Parallax 18

2.2.3 Effect of Spatial Resolution on Depth Perception 21
2.3 Perception of Space 24
2.4 Fundamentals of Optical Design 27
2.4.1 Lens Aberrations 28
2.4.2 Lenses and Lens Systems 31
2.5 Light Fields 38
2.5.1 Light-Field Displays 39
2.5.2 The Resolution Challenge 42
2.6 Chapter Summary 44
Chapter 3: Perception of Depth in Blurred Displays 45
3.1 General Approach 45
3.2 Displaying the Stimuli 52
3.3 Adjusting the Position 54
3.4 Choosing the test positions 55
3.5 Pilot Studies 58
3.5.1 Pilot Study 1: Lack of Monocular Depth Cues 59
3.5.2 Pilot Study 2: Motion vs. Binocular Parallax 60
3.5.3 Pilot Study 3: Comparing Strength of Diffusion 61
3.6 Main Study 63
3.7 Chapter Summary 71
Chapter 4: Optical Design 73
4.1 Design Concept 74
4.1.1 Initial Design 76

4.1.2 Improved Lens System 77

4.1.3 Dual Diffuser System 83

4.1.4 Improved Efficiency Illumination Source 88

4.2 Simulation of the Optical Design 91

4.2.1 Imaging Stage 94

4.2.2 Illumination Stage 102

4.3 Performance analysis 107

4.4 Chapter Summary 113

Chapter 5: Conclusion and Future Directions 114

5.1 Implications for architectural and urban design 115

5.2 Future work 119

5.3 Conclusion 121

References 123

Appendices 131

Post Study Questionnaire 131

List of Tables

Table 1: Summary of resolution requirements for depth cues 15

Table 2: Age ranges of main study participants..... 65

Table 3: Number of depth planes as a function of visual acuity..... 68

Table 4: Summary of lens element performance characteristics 80

List of Figures

Figure 1: Illustration of a partially diffusing window (left) compared to a clear window (right) .. 1

Figure 2: Illustration of a partially diffusing window seen from three viewing positions 4

Figure 3: Illustration of the terms used in calculating a just noticeable difference Δz 11

Figure 4: Faubert study comparing mean visual depth acuity for both stereo and motion parallax. The error bars show the standard error. 20

Figure 5: Faubert study comparing the mean positioning error for binocular stereopsis and motion parallax. The error bars show the standard error. 21

Figure 6: Rogers study comparing average visual depth acuity as a function of spatial frequency. The error bars show standard errors..... 22

Figure 7: Spherical aberration (left), reduced with an aspherical lens (right) 28

Figure 8: Chromatic aberration of a non-achromatic biconvex lens..... 29

Figure 9: Field curvature aberration 30

Figure 10: Barrel distortion (left) and pincushion distortion (right)..... 30

Figure 11: Plano-convex (left) and biconvex lenses (right)..... 31

Figure 12: Wollaston meniscus lens 32

Figure 13: Petzval portrait lens system. From left to right: biconvex achromat, meniscus, biconvex lenses 32

Figure 14: Harrison and Schnitzer globe lens. Formed of two identical meniscus lenses placed symmetrically on both sides of a central aperture 33

Figure 15: Beck Hill Sky fisheye lens 34

Figure 16: Asphere lens. The surface of the lens deviates from a perfect sphere of radius R, thus minimizing spherical aberrations 36

Figure 17: Fresnel Lens. The surface profile of a conventional biconvex lens (right) is divided up into many smaller facets to form a Fresnel lens (left), decreasing the thickness and material required for the lens 37

Figure 18: Illustration of a light-field display viewed from two viewing positions 41

Figure 19: Apparatus for measuring diffusion 47

Figure 20: Diffusion profiles for light (1.4 degrees), medium (1.8 degrees), and strong (2.6 degrees) diffusions used in pilot study 48

Figure 21: Angular diffusion profile for "strong" diffuser used in main study. The standard deviation is 2.2 degrees. 49

Figure 22: Simulated array of lens apertures 50

Figure 23: Subjects participating in the user study, showing controller and relative position of the viewport. Lights were turned off for the study. 51

Figure 24: Illustration of the experimental apparatus 53

Figure 25: Trial responses for a single participant for a single test case. Each correct answer reduces the separation while each incorrect answer increases the separation until 30 trials have been completed. The dashed line shows the final JND threshold for 75% correct. 56

Figure 26: Timing diagram of each user study 57

Figure 27: Results from first pilot study verifying that monocular depth cues have a negligible contribution in this experimental design, and that the experimental apparatus results in similar visual acuity to prior research (29.2 ± 0.2 seconds of arc). Error bars show the standard error for three observations. 60

Figure 28: Results from second pilot study. The Mono+Motion case shows that motion parallax alone is a useful depth cue, although not as effective as binocular parallax for viewers with normal vision. The Stereo+Motion case shows that little improvement to stereo acuity is achieved when adding motion parallax to binocular parallax for viewers with normal sterescopic vision. Error bars show the standard error for three observations. 61

Figure 29: Results of the third pilot study showing how the depth JND increases with increasing strength of diffusion. Error bars show the standard errors..... 62

Figure 30: The experiment stimulus with light diffusion 63

Figure 31: The stimulus with the strong diffuser..... 64

Figure 32: The stimulus with the strong diffuser and array of apertures..... 64

Figure 33: Main study results showing the average visual acuity for light diffusion, strong diffusion, and aperture and strong diffusion. Error bars show the standard error for 22 observers. 66

Figure 34: Box plot showing the range of data between the 25th and 75th percentiles within the rectangular regions, as well as the full range of the data within the error bars. 67

Figure 35: Relationship between absolute depth and number of distinguishable depth planes, for a visual acuity of 30 seconds of arc. 68

Figure 36: Difficulty level of user study task from post-study questionnaire. Participants generally reported the “Light Diffusion” between Easy and Medium, the “Strong Diffusion” case as Challenging, and the “Aperture and Strong Diffusion” case as between Challenging and Very Challenging. 69

Figure 37: Confidence level of users' results from post-study questionnaire. Participants generally reported high confidence for "Light Diffusion", and low confidence for both "Strong Diffusion" and "Aperture and Strong Diffusion"..... 70

Figure 38: Illustration of a view of a scene through a window from two different viewing positions. In the top pair, the window is transparent; in the bottom pair the window is partially diffusing (simulated using a Gaussian low-pass filter). 73

Figure 39: Light-field display concept. Light emitting from different sets of pixels P1 and P2 are projected at different angles from the display..... 75

Figure 40: Illustration of first design iteration using an array of Fresnel lenses in front of an LCD 76

Figure 41: Optimized lens profiles. Clockwise from top-left: Fresnel, Plano-Convex, Meniscus, Globe..... 79

Figure 42: Illustration of the second design iteration using an array of globe lenses..... 81

Figure 43: Illustration of an array of projection units using globe lenses 82

Figure 44: Diffusion system with two diffusers 84

Figure 45: Optimal angular diffusion profile of first diffuser from a single light source (left) and the sum of multiple adjacent sources (right)..... 86

Figure 46: Illustration of third design iteration, using globe lenses and dual-diffusion system. The first diffuser is adjacent to the lens array, and the second diffuser is separated by a gap..... 87

Figure 47: Low optical efficiency due to globe lens array 88

Figure 48: Improved optical efficiency using an array of LEDs and Fresnel field lenses, with the facet sizes of the field lenses exaggerated for illustrative purposes 89

Figure 49: Illustration of the design concept using an array of field lenses to focus light through the apertures of the globe lens array, followed by the dual-diffusion system 90

Figure 50: Illustration of optical path, with the Imaging Stage shaded in red and the Illumination Stage shaded in yellow..... 91

Figure 51: Illustration of a 2D light-field with a uniform spatial distribution (left), and uniform angular distribution (right)..... 92

Figure 52: An illustration of the image (top) and light-field distribution (bottom) of a distant object seen through a conventional clear window (left) and a partially diffusing window (right) 93

Figure 53: Illustration of optical path for the imaging stage 94

Figure 54: 2D light field from four pixels corresponding to the optical axis of four lenses 95

Figure 55: 2D light field after propagating to the globe lens..... 96

Figure 56: 2D light field at aperture of globe lens. The spatial distribution is limited by the aperture, illustrated as a grey shaded area 97

Figure 57: 2D light field after globe lens. The angular distribution is reduced to sigma = 0.35 degrees 98

Figure 58: 2D light field after the first diffuser, showing increased angular distribution 99

Figure 59: 2D light field before the second diffuser, showing how the spatial distribution has diverged as a function of the angle. 100

Figure 60: 2D Light field after the second diffuser, showing both spatial and angular uniformity 101

Figure 61: Illustration of optical path for the illumination stage. Yellow shading illustrates the envelope of light rays emitted by the LED array. 102

Figure 62: 2D light field of the modelled LED light sources 103

Figure 63: 2D light field before the field lens. The angle is unchanged from the light source, but the light has diverged spatially due to the propagation through free space. 104

Figure 64: 2D light field after the field lens. The angular distribution has been reversed from the angle incident on the lens to focus through the aperture of the lens array..... 105

Figure 65: 2D light field at the apertures of the globe lenses, indicated by the gray shaded areas. 106

Figure 66: Angular distribution of the light field for paraxial light. The standard deviation of the paraxial light through the optimized optical system is 3.4 degrees 108

Figure 67: Angular distribution of the light field for three viewing angles. 109

Figure 68: Spatial distribution of light field across the display. The horizontal axis is the position across the display, and the vertical axis is the difference in intensity of the light, represented as Just Noticeable Differences..... 110

Figure 69: Optical components of an intermediate-resolution light-field display module..... 113

Glossary

2AFC:	Two Alternate Forced Choice: An experimental methodology whereby subjects are forced to choose between two different stimuli.
4D Light Field:	The amount of light travelling in every direction through a plane.
Arc-Seconds:	A measure of visual angle. Each second of arc corresponds to 1/3600 degrees
CSF:	Contrast Sensitivity Function. The human visual system's sensitivity to changes in spatial patterns of light
JND:	Just Noticeable Difference. The amount that some stimulus must be changed so that observers can detect the difference half the time.
LCD:	Liquid Crystal Display.
LED:	Light Emitting Diode.
Light Field:	The amount of light travelling in every direction through every point in space.
Light Field Display:	A display device that can display a 4D Light Field, by producing different images at multiple angles
OLED:	Organic Light Emitting Diode
PSF:	Point Spread Function. The shape that light produces as it passes through a diffusing medium
Psychometric Function:	A model of the relationship between a value of a stimulus and an observer response.

Chapter 1: Introduction

The goal of this thesis was to determine if it is possible to create a display that can simulate the sense of depth associated with looking through a window at a vast open space. The motivation to do so arises from studies that have shown that windows provide various psychological benefits, including one due to the "perception of spaciousness" (Markus, 1967). This project's primary goal is to show that a practical intermediate-resolution light-field display, despite being somewhat blurry, can nevertheless provide a stimulus that is indistinguishable from that of a partially diffusing window, as illustrated in Figure 1, and so can provide an observer with a substantial perception of spaciousness.

Figure 1: Illustration of a partially diffusing window (left) compared to a clear window (right)

1.1 Motivation and Outline

As human populations continue to move into more urban environments and smaller living and working spaces there arises a question of vital importance: How can human well-being be optimized in an increasingly dense environment? One estimate (Seto, Fragkias, Guneralp, & Reilly, 2011) is that by 2030 the global urban population will reach 5 billion, with three out of five people living in cities. Increasing urban population raises real estate value, which in turn applies downward pressure on the size of living space. This can have the effect of reducing access to large open spaces, fewer windows, and windows overlooking short distances rather than distant views. This lack of perceived openness can have a negative psychological impact (Markus, 1967). The research goal for the work described here is to determine if it is possible to create, in a practical way, the same visual stimulus as is found when looking through a window overlooking a distant space- i.e., the *light field* passing through the window. The conclusion reached is that this can be realized using a novel form of light field display.

Chapter 2 summarizes related work in psychology and optics, leading to the conclusion that truly displaying the light field seen through a window would require much greater display resolution than is available today. This limitation motivated the idea of generating the desired sense of spaciousness via an intermediate-resolution light-field display. A display of this type simulates the effect of looking through a partially diffusing window (a commonly used form of privacy window) at a distant scene.

Chapter 3 describes a psychophysical experiment that measures the perception of depth that can be provided by an intermediate-resolution light field. Results indicate that the human visual system can perceive a sense of depth even with significant amounts of blur.

Chapter 4 describes a practical design for displaying an intermediate-resolution light field using conventional components that are available today.

Chapter 5 discusses the architectural implications of using such an intermediate-resolution light-field display to enhance a perceived sense of openness, and describes some possible future directions based on the results of this thesis.

1.2 Basic Concept

The key concept for simulating a distant view is to present an image for which objects appear to have a fixed location in the three-dimensional world as the viewer moves around. Consider walking under a blue sky with a few clouds: the direction of the clouds and sun with respect to the three-dimensional world remains the same, and the brain interprets this as indicative of a distant view. In contrast, an object that is close by changes its relative direction as the viewer moves. This latter effect, called motion parallax, is one of the effects encountered when looking through a window. In addition, some portions of the image may be revealed as the user changes position, providing a second source of information about the location of objects, and further reinforcing the sense of depth. This effect is illustrated in Figure 2, which shows how the view of a scene through a partially diffusing window changes from different viewing positions.

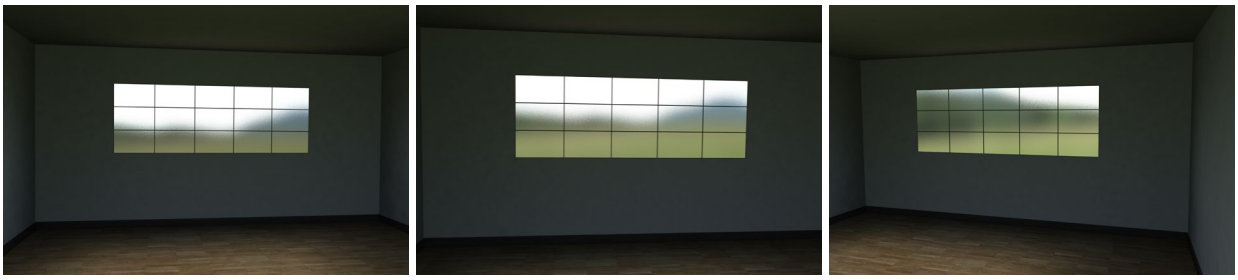

Figure 2: Illustration of a partially diffusing window seen from three viewing positions

This thesis is based upon three hypotheses:

- A view of a distant scene can provide a sense of space convincing enough to make people feel better (Chapter 2)
- Perception of depth based on the cues that occur when an observer moves around remains possible even with blurry images (Chapter 3)

- It is possible to build an economical display that shows a large number of distinct views from different positions (termed a light field-display), each of which is blurry (Chapter 4)

1.3 Contributions

The contributions of this thesis are (i) an enhanced understanding of depth perception, (ii) a novel optical design for an inexpensive intermediate-resolution light field, and (iii) real-world applications that relate to such a display. Chapter 2 includes a review of depth perception, from the perspective of how depth cues are affected by spatial and angular resolution; this then forms the basis for the psychophysical experiment described in Chapter 3 to measure the limits of human depth perception for intermediate-resolution stimuli. Chapter 4 describes the novel configuration of conventional optical elements to create an intermediate-resolution light-field display that can be built in a practical, low-cost manner.

Chapter 2: Background

This chapter begins by reviewing the basic cues of depth perception, with an emphasis on how contrast and resolution may impact the effectiveness of each. Two such cues, binocular parallax and motion parallax, are then described in greater detail, since they form the basis for the experimental procedure described later. The chapter then reviews the human perception of space and openness, and the psychological and physiological impact of being exposed to open spaces. Next is a review of several optical principles that form the basis of the display design proposed in this thesis, followed by a summary of the state of the art in light-field displays. This chapter concludes that the resolution required to achieve a high-resolution light-field display is probably decades away from being achievable, at least for practical purposes. It is this gap between available resolution and resolution requirements that this thesis project circumvents.

2.1 Fundamentals of Depth Perception

Depth perception is the ability of the human visual system to determine the three-dimensional structure of the outside world from the two-dimensional pattern of light on the eyes. It includes determining the distance of objects from the observer, the direction they are travelling, and their speed in the world. From an evolutionary perspective, the ability to do this was often critical for survival. In order to perceive depth, the visual system relies on multiple cues. Alone, any single cue may yield unreliable or ambiguous perceptions, but the combination of multiple depth cues can yield an accurate estimate of the geometry of the scene around us (Wismeijer, Erkelens, Van Ee, & Wexler, 2010). Depth cues can be placed into one of three categories: monocular, binocular, and observer-motion cues. A novel contribution of this thesis is to categorize how each type of cue is expected to be affected by spatial and angular resolution.

2.1.1 Monocular Depth Cues

Monocular depth cues are those that can be derived from a single stationary viewpoint. Such cues have been exploited by artists for hundreds of years to provide viewers with a sense of depth in a picture or painting. However, these cues work only to a limited extent; anyone who has looked at a painting, photograph, or conventional display screen is able to see that it is not actually a window. Some of these cues are:

1. Occlusion. This cue occurs when an opaque object obstructs another object in the line of sight (Gillam & Borsting, 1988). The obstructing object is perceived as closer. This cue does not require high spatial resolution, provided the occluding object is large enough to be seen in the display.

2. Accommodation. This arises from the tension of the ciliary muscle that adjusts the focal length of the lens of eye, bringing into focus objects at different distances (Baumeister & Kohnen, 2008). The reciprocal of the focal length of the eye is termed a diopter, labeled D. By definition, an object at a distance of 1 meter requires 1D of optical power (1/1m=1D); observing an object at a distance of 0.1m requires 10 diopters. Focusing on an object between 1 m and infinity requires a range of 0 to 1 diopters. A distance used commonly in ophthalmology as the "distant visual acuity" is 6m (Cline, Hofstetter, & Griffin, 1997), corresponding to 0.16 diopters – beyond this range the difference in lens convexity is considered negligible. Accommodation is most effective when viewing scenes with high spatial and angular resolutions therefore requires sharp focus. For low resolution images, accommodation has been found to be a weak cue, even for objects closer than the 6 m of the distant visual acuity (Hiura, Komine, Arai, & Mishina, 2017).
3. Aerial Perspective or Atmospheric Haze. Objects that are perceived to be obstructed by atmospheric haze (lower contrast and lower color saturation) are interpreted as located at greater distances from the observer (O'Shea, Blackburn, & Ono, 1994). A common example is how the mountains in a mountain range appear successively bluer and more hazy as their distance increases. This is a well-known and commonly exploited technique in paintings, for example in work by Leonardo da Vinci. This cue is expected to require only low to medium spatial resolution provided the objects are sufficiently large.
4. Lighting and Shadows. Lighting and shadows also contribute to the perceptions of 3D shape and depth of objects (Lipton, 1982), (Dee & Santos, 2011). This cue is independent of resolution provided the objects in the scene appear large enough to cast visible shadows.

5. Linear Perspective. This cue involves the apparent convergence of parallel lines in the image. For example, paintings and photographs often include converging lines in buildings, walls, fences, roads, or railways, to create a sense of depth (Cucker, 2013). For this cue to be most effective, the scene must contain recognizable parallel lines, which requires fairly high spatial resolution.
6. Texture Gradient. Smoother and denser textures are perceived as further away, especially if a constant surface spans the distance from near to far (Weinstein, 1957), (Degelman & Rosinski, 1976), (Maureen & Stéphane, 2002). This requires sufficient spatial resolution to display a recognizable texture – a low resolution image will provide little useful information about detailed textures.
7. Familiar Size. The brain unconsciously compares the apparent angular size of an object with knowledge of its physical size to estimate its distance. For example, hikers commonly judge distances by the apparent size of distant trees. This can fail in environments without objects of known size, or sizes that differ from what have previously been learned. To be useful as a depth cue, the objects must be recognizable, which in turn requires fairly high spatial resolution.
8. Relative velocity. Distant moving objects shift more slowly in their viewing direction than do closer ones (Swanston & Gogel, 1986). For example, if two airplanes are flying in the sky, the more distant one changes its retinal position more slowly, so it is perceived as further away (as they are assumed to be moving in approximately the same speed in the real world). This depth cue is expected to be independent of spatial resolution, provided the objects are sufficiently large. A quickly moving foreground and a slowly moving background is

expected to produce a convincing sense of depth as seen when observing a landscape through the side window of a moving vehicle, even if the view is blurry.

9. Optical expansion. As the distance between an observer and an object decreases, the retinal size of the object expands, leading to the perception of a change in its distance (Swanston & Gogel, 1986). This cue is also expected to be independent of spatial resolution; objects near the plane of the display may expand rapidly while distant objects may expand slowly or not at all. This is expected to invoke a compelling sense of depth even if the objects are large and blurry, for example by observing a landscape through a blurry front window of a moving vehicle.

2.1.2 Binocular Depth Cues

Binocular depth cues are formed by comparing the stimuli from the left and right eyes, each of which views the scene from a slightly different vantage point. Determining depth from two such points requires finding the corresponding image features between the two eyes, which requires a minimal amount of spatial resolution. There are three different kinds of binocular depth cue:

1. Binocular Parallax. This is the difference in angle of points between the two images; the distance between the point is termed binocular disparity (Howard & Rogers, 1995). By comparing two images of the same scene obtained from two slightly different positions (illustrated as points A and B separated by an inter-ocular distance I in Figure 3), the brain triangulates the object distance (z) with a higher degree of accuracy than for most other depth cues.

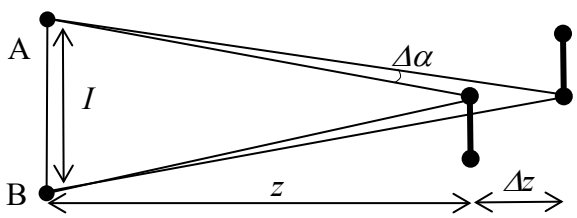

Figure 3: Illustration of the terms used in calculating a just noticeable difference Δz

The relative distance *Δz* between two points at distance *z* can be converted to an angular difference *Δα* (which Howard termed the *parallactic angle*), measured in radians, using Equation (1) (Howard H. J., 1919).

$$\Delta a = \frac{\Delta z \cdot I}{z^2 + (\Delta z \cdot z)} \quad \textbf{(1)}$$

Since *Δz* is generally much smaller than *z*, Equation (1) shows that the binocular disparity produced by a given depth separation decreases roughly as the square of the distance to the observer. The dependence of this cue on angular resolution is low; this is discussed in greater detail in section 2.2. Most conventional "3D" displays use binocular parallax to help provide a sense of depth, typically by using special eyewear to provide a different image to both left and right eyes. Parallactic angle *Δα* is an effective depth cue at distances as large as 250m (Palmisano, Gillam, Govan, Allison, & Harris, 2010).

2. Binocular Convergence. This is an indicator of depth that arises from the difference in the direction of the two eyes when looking at an object of interest (Cassin & Solomon, 1990). For close objects the angle between the eyes is quite large, while for distant objects it is small; convergence is negligible for objects farther than 10 meters (Okoshi, 2012). This cue is expected to require similar angular resolution as binocular parallax for objects that are

within 10 meters, provided the objects are sufficiently largewell as sufficiently high angular resolution.

3. Da Vinci Stereopsis. This relatively weak cue arises when an object occludes different parts of another object when viewed from the left eye as compared to the right eye (Nakayama & Shimojo, 1990). This cue requires sufficiently high spatial and angular resolution for there to be difference in the retinal image between the left and right eyes.

2.1.3 Observer Motion Depth Cues

The motion of the observer can create two kinds of depth cue, depending on whether the motion is perpendicular or parallel to the direction of view. Both correspond to cues formed by moving objects viewed by a stationary observer; the main difference is that for a moving observer motion signals exist throughout the field of view:

1. Motion Parallax. This is the difference in images caused by the movement of the observer with respect to a stationary rigid scene; it is the observer equivalent of the relative velocity cue. It operates under a similar principle as binocular parallax (Ferris, 1972). For example, when observing a stationary landscape from a moving automobile, close objects will undergo a more rapid shift of viewing direction than distant objects. This cue will create a sense of depth for low resolution objects, so long as (a) the observer is in motion and (b) their motion causes them to see a changing image. It is this second requirement that is not met by today's conventional 2D or binocular displays and is only achieved for very small ranges of movements for the latest auto stereoscopic displays. Like the relative velocity monocular cue, this is expected to invoke a convincing sense of depth even with low resolution.

2. Motion Expansion. This is the equivalent of the monocular optical expansion cue, where the relative size of all objects within a stationary scene is affected by the observer's distance from the scene. As the observer increases the distance to the scene close objects should reduce their relative size more rapidly than distant objects. Like the optical expansion cue, the motion expansion cue is expected to provide a convincing sense of depth even with low spatial and angular resolution, as an observer moves towards or away from a blurry scene.

2.1.4 Enhancement and Interference of Depth Cues

To produce a realistic sense of depth in an image, both artists and display engineers strive to use as many cues as possible while minimizing conflicts between them. An imaging system that accurately reproduces all these cues would produce a light field indistinguishable from that of an actual window. However, it is not necessary to reproduce all the cues to provide a satisfactory sense of depth - it is only necessary to ensure reasonable compatibility between the depth cues that are present, and to avoid strongly conflicting ones (Banks, Akeley, Hoffman, & Girshick, 2008).

Although monocular and binocular depth cues are effective for a single viewing position, some of these can be enhanced by changes in viewing position as the observer moves around. Of particular interest is the occlusion depth cue, being the strongest of the monocular cues (Ware, 2010). From different viewing positions a close object overlaps and obstructs different portions of a more distant object. This has the same relationship to da Vinci stereopsis as motion parallax does to binocular parallax. It is expected that for an intermediate-resolution light-field display, moving around the display to view it from different positions results in an enhanced sense of

depth. In contrast, for a painting, the strongest sense of depth may be achieved by viewing it from a single position, and possibly with only one eye; otherwise, the cues may conflict, and lead to a lessened sense of depth. It is common for a painting, drawing, or photograph to be captured from a specific viewing position. Each of the depth cues mentioned above will be consistent from that specific viewing position but may be inconsistent when viewed from another viewing position.

Cue conflict occurs when two or more cues provide different and incompatible information – for example, as an observer approaches a painting or conventional display, an object displayed in it will expand at a rate reflecting the distance of the image surface rather than the distance of the object in the scene. In a real scene, a distant object will appear to expand very little, while a close object may appear to expand significantly. By contrast, a depiction of an object in a painting or on a conventional display will appear to expand by the same factor regardless of the intended distance of the object. Conflicting cues may result in higher cognitive load as attention processes are required to comprehend incompatible depth cues, and at the most extreme have been found to cause discomfort and even nausea (Hoffman, Girshick, Akeley, & Banks, 2008)

2.1.5 Summary of Depth Cues

Table 1 summarizes the various depth cues by type and resolution requirements, and highlights those that are expected to be most effective for an intermediate-resolution display. Of these, the binocular and observer motion cues are of greatest interest to this project, as they are not supported by a conventional 2D display.

Table 1: Summary of resolution requirements for depth cues

Depth Cue	Type	Expected Spatial Resolution Required	Expected Angular Resolution Required
Occlusion	Monocular	Low	N/A
Accommodation	Monocular	High	High
Aerial Perspective	Monocular	Medium	N/A
Lighting and Shadows	Monocular	Medium	N/A
Linear Perspective	Monocular	High	N/A
Texture Gradient	Monocular	High	N/A
Familiar Size	Monocular	High	N/A
Relative velocity	Monocular	Low	N/A
Optical Expansion	Monocular	Low	N/A
Binocular Parallax	Binocular	N/A	Low
Convergence	Binocular	N/A	Medium
Da Vinci Stereopsis	Binocular	High	Medium
Motion parallax	Observer Motion	Low	Low
Motion expansion	Observer Motion	Low	Low

For an intermediate-resolution display, high resolution is not available, and so many cues are not anticipated to play a significant role in depth perception (Table 1). The notable exceptions that are anticipated to operate at low resolution and can take advantage of a light-field display are occlusion, motion parallax, and binocular parallax. These three cues are therefore of the greatest interest here. The next section explores the mechanisms of the binocular stereopsis and motion parallax depth cues in greater detail.

2.2 Mechanisms for Binocular Parallax and Motion Parallax

When viewing a blurry image of a scene, as depicted in Figure 1, both binocular and motion parallax cues can play a significant role in the perception of depth. This section begins by reviewing the early experiments that measured binocular depth perception, which inspired the design of the experiment described in Chapter 3; it then compares motion parallax against binocular parallax, referencing several recent studies that compared the effectiveness of the two cues. This section then summarizes how these cues are affected by image resolution.

2.2.1 Binocular Parallax

The first efforts to study binocular vision were conducted by Charles Wheatstone in 1838, who for this purpose invented the stereoscope – a device that can present different images to each eye to convey a sense of depth (Wheatstone, 1838). It was nearly a century later that binocular acuity (the ability of the visual system to resolve depth from binocular vision) was actually measured, for navy pilots (Howard H. J., 1919). The test determined a Just Noticeable Difference (JND) – the separation in the distance whereby an observer can barely determine if one object (in this study, a wooden peg) is closer than another; this is called the visual depth acuity. In order to isolate binocular parallax, the experimenters eliminated as many monocular depth cues as possible (although retinal image size and the visual separation of the two objects were still visible). To handle this, the test was conducted twice, once with one eye and then again with two eyes, with the binocular JND obtained by comparing the two results. Preliminary testing revealed that the monocular JND was more than 20x greater than the binocular JND, and so could be ignored. The apparatus consisted of a board into which two 1 cm diameter objects were inserted at various distances from an observer. Observers were asked to report which of the two objects

appeared nearer. The distance between the objects was varied, decreasing upon correct answers and increasing with incorrect answers. Based on the results of 106 participants, Howard proposed that a “normal” judgement of distance should be less than 8.0 seconds of arc.

Since navy pilots may not be representative of the general population, Coutant & Westheimer (1993) examined 117 college students, and found a median angular JND of 13 seconds of arc; they also found that 80% of observers could detect an angular disparity of 30 seconds of arc. In addition, the study found that 5% of the population is essentially unable to use binocular cues alone to sense depth; however, this “stereo-blind” population is still perfectly able to function by using other depth cues. Due to the relatively young age and access to eye and vision care of the populations tested in those studies, these values are not necessarily indicative of the average visual acuity of the general population. Nonetheless, a visual depth acuity angle of 30 seconds of arc can be considered a suitable reference point for verifying the depth perception experiment described in Chapter 3.

2.2.2 Motion Parallax

Motion parallax is similar to binocular stereopsis, but based on comparing two different viewpoints at slightly different times and positions instead of slightly different viewpoints. The greater an object’s apparent velocity (assuming it is stationary in the world), the closer it appears to the viewer. One difference between binocular stereopsis and motion parallax is that with binocular stereopsis, disparity is determined by the separation of the two eyes – 6.3 cm for the average adult observer (Dodgson, 2004); whereas with motion parallax, disparity is determined by the velocity of the observer relative to the scene. If the relative positions of the observer and

the scene remain static, the observer will not experience a motion parallax depth cue, while a fast-moving observer may experience a very strong and unambiguous one. It is interesting to note that both binocular stereopsis and motion parallax are used by many animals, with some animals emphasizing motion parallax to improve their depth acuity, for example the praying mantis and the cobra, and others emphasizing stereopsis, for example the hammerhead shark which has a considerably increased inter-ocular distance (Kral, 2003), (McComb, Tricas, & Kajiura, 2009).

To compare the acuity of binocular vision and motion parallax, Equation (1) can be used to convert a just noticeable difference in depth Δz to visual angle $\Delta\alpha$ using an assumed mean interocular distance I of 6.3 cm (Dodgson, 2004) or by replacing I with the displacement between the two viewpoints (in the case of motion parallax). Faubert used this analysis to compare the effectiveness of binocular stereopsis with motion parallax (Faubert, 2001), building upon previous work by Rogers and Collett (Rogers & Collett, 1989). The retinal size of the stimuli provided a monocular depth cue that was assumed to be negligible. The study was conducted under two conditions: binocular stereopsis only (by fixing the head position in a chin rest), and motion parallax only (by covering one eye and moving the chin rest by 6.7 cm and back over one second). The results, shown in Figure 4, found that stereo acuity outperformed motion parallax by roughly an order of magnitude. Faubert noted that subjects in the binocular test had healthy binocular vision so may never have needed to develop the use of motion parallax for depth judgments. To explore this idea, the test was repeated with three subjects who had recently lost the use of one eye. These subjects had motion parallax based stereo acuity that was nearly as good as normal binocular stereo acuity, indicating improved performance for motion

parallax compared to observers with normal binocular vision. This indicates that motion parallax can be used to form judgments of relative depth when binocular stereopsis is not available.

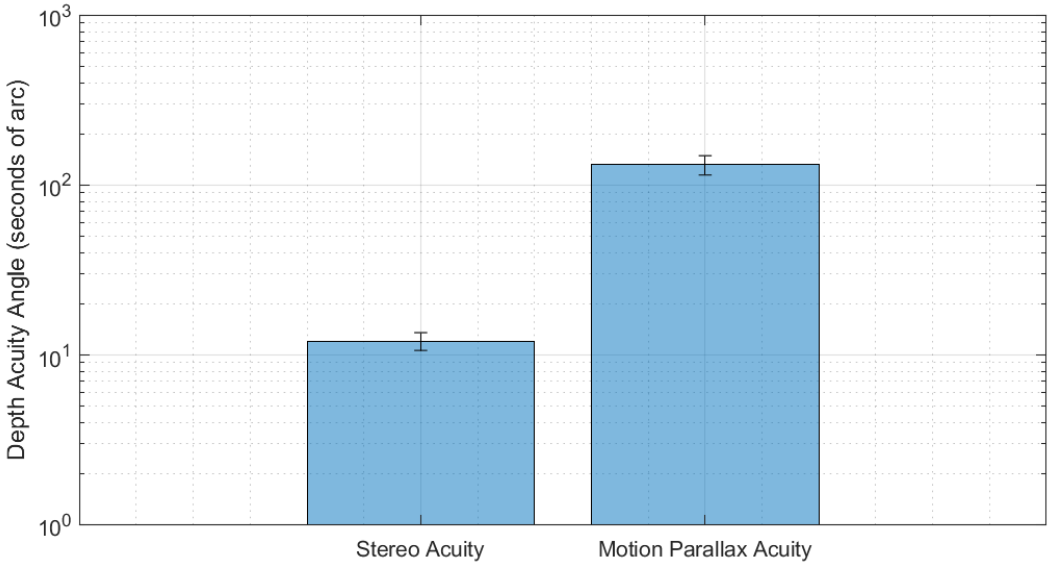

Figure 4: Faubert study comparing mean visual depth acuity for both stereo and motion parallax. The error bars show the standard error.

Faubert also experimented with combining binocular and motion parallax. As depicted in Figure 5, observers using only one eye were able to judge depth differences more accurately when allowed to adjust their viewing position to provide motion parallax, compared to a stationary viewing position. When observers used both eyes, there was still an improvement between a fixed viewing position and motion parallax, although the improvement was smaller. In all cases motion parallax alone was sufficient to judge relative distances in depth between two objects.

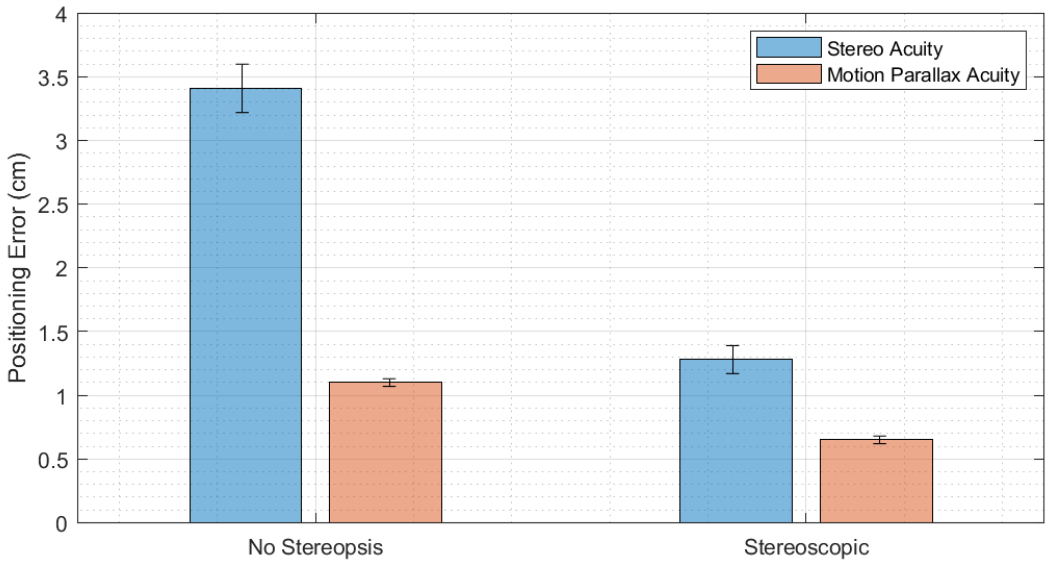

Figure 5: Faubert study comparing the mean positioning error for binocular stereopsis and motion parallax. The error bars show the standard error.

2.2.3 Effect of Spatial Resolution on Depth Perception

Binocular stereopsis and motion parallax depth cues are expected to have significant impact on the ability of an intermediate-resolution light field display to create an enhanced sense of depth. This section reviews previous studies which explored how each of these two depth cues is impacted by the spatial resolution of a display.

The human visual system is known to be more sensitive to certain spatial frequencies than others (Barten, 1999). Its sensitivity to contrast, described by the contrast sensitivity function (CSF), is greatest at frequencies near 4 cycles per degree, and less at both higher and lower frequencies. The relationship between depth acuity and spatial frequency has been studied by comparing both

binocular stereopsis and motion parallax at different spatial frequencies (Rogers & Graham, 1982). The average results for three observers are shown in Figure 6.

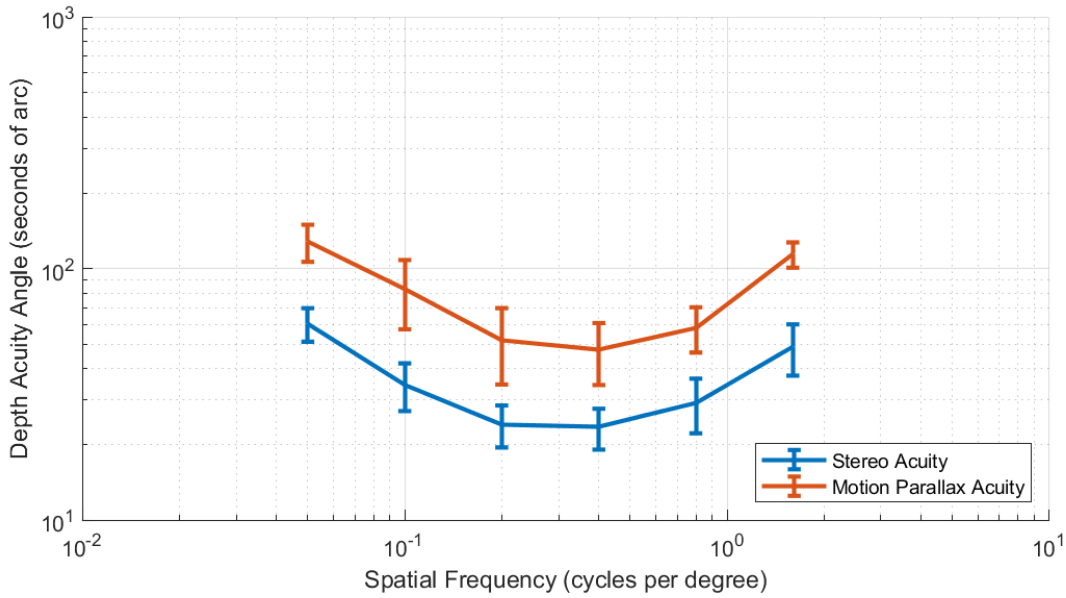

Figure 6: Rogers study comparing average visual depth acuity as a function of spatial frequency. The error bars show standard errors.

Rogers and Graham found that the peak sensitivity for depth acuity is close to 0.4 cycles per degree, roughly an order of magnitude lower than the peak contrast sensitivity. Comparing motion parallax and binocular stereopsis shows very similar shapes for the curves – likely due to similar mechanisms between the two effects. However, two observers had binocular acuity that was roughly double that of motion parallax, whereas another observer had about the same for both. This is consistent with the Faubert study, which found that binocular stereopsis works well for some observers, but no better than motion parallax for others. The sensitivity for stereo acuity was found to decrease to half of the peak sensitivity when frequencies were above 1.0 cycles per degree and below 0.1 cycles per degree. These results provide some guidance into the spatial

resolution requirements for a light-field display. Most depth acuity is provided by spatial frequencies in the range of 0.1 to 1.0 cycles per degree. Lower and higher spatial frequencies would not be expected to further contribute to the depth shown by the display. This is very promising for an intermediate-resolution (0.1-0.4 cycles per degree) light-field display because it suggests that high resolution (>0.4 cycles per degree) is not required. An intermediate-resolution display would be fully capable of supporting the perception of depth in an observer from the disparity and motion parallax depth cues alone.

2.3 Perception of Space

This section reviews the subjective experience of spaciousness. This experience involves higher-level cognitive processing beyond the purely visual – specifically how the brain interprets estimates of depth to form an impression of the scale of the local environment. Previous research studied this by showing subjects a series of photographic images and asking for the subjective impression of the scene, including terms such as "vastness" and "openness" (Klatzky, Thompson, Stefanucci, D., & McGee,, 2017). Although these studies were conducted with 2D images, several of the conclusions are relevant to a light-field display. Klatzky et al. define a "vast" space as "*extending to very far distances, seemingly without limit, making the viewer feel as a small element within the space*". The study concluded that scenes with a clear horizon line are perceived as the most vast, while forest and close views of vegetation are the least. The authors hypothesized that the sensation of a "vast" space may correspond to an extreme level of perceived depth. They also noted that subjects do not need to be within the space in order to judge that it is vast, an important consideration for a light-field display. An expanded sense of space has also been implicated as a precursor to the emotion of awe, which in turn has been linked to well-being, social engagement, and helping behaviors (Rudd, Vohs, & Aaker, 2012).

Another study correlated spatial properties of a scene with subjective ratings of "openness" (Oliva & Torralba, 2001). It found that spatial frequencies as low as 4 to 8-cycles/image provided enough information for fast recognition of various environments even when the shape and the identity of the objects in them could not be recovered, and that the perception of "openness" appears to be built on a low-resolution spatial configuration.

One area that has received considerable attention is the effect of including windows in working and living spaces. A recent study (Matusiak & Klöckner, 2016) asked participants to rate how pleasing they found a view from a window, and discovered a positive correlation between the subjective quality of a view and the view distance (the distance from the observer to the furthest point in the scene). A light-field display capable of showing great distances is expected to provide a similar positive benefit.

Collins (Collins, 1975) concluded from a review of previous studies that the role of windows includes relief from feelings of claustrophobia. Other researches have measured the physiological benefits of windows. For example, Kahn et al. (Kahn, et al., 2008) measured the heart rate of participants after inducing a mild stress in one of three environments: one equipped with a plasma television display of a scene, one with a real window, and one a plain blank wall. By measuring the time for heart rate to return to normal levels they found that the plasma display was not much different than a blank wall, while a real window was much better. Their study also tracked the eye gaze of the participants and found that participants looked just as frequently at the plasma window as a real window, but for a much smaller duration on average, indicating less interest.

The hypothesis that motion parallax plays a key role in all of this has also been investigated. For example, Radikovic et al. (Radikovic, Leggett, Keyser, & Ulrich, 2005) asked participants to report their mood after occupying one of two rooms for a period of time. Both rooms were equipped with an LCD display showing an image of an outdoor scene. In one room the image was static while in the other room a head tracking system positioned above the display changed

the view with respect to the viewer position to simulate the motion parallax of a real window. The study found that viewers ranked the motion parallax window much higher than a static display in terms of mood as well as realism. The authors concluded that a simulated window could have both psychological and physiological benefits, such as an improved sense of wellbeing. Another study (IJsselsteijn, Oosting, Vogels, Kort, & E.J. Loenen, 2006) examined the impact of monocular depth cues including motion parallax, blur, and occlusion on the realism of a virtual window. As in the Radikovic study, a head tracking system was used to display the correct perspective according to the viewer's position. Results showed that the largest improvement arose from motion parallax. The addition of occlusions at the plane of the window frame further enhanced the sensation by further reinforcing the depth cues. IJsselsteijn et al. also discussed the possible commercial opportunities for virtual windows in several industries, including health care, underground transit, and residential.

The literature reviewed in this section generally supports the hypothesis that a view of a distant scene can produce the desirable characteristic of providing a greater sense of space, and that the greater sense of space may help people to feel better.

2.4 Fundamentals of Optical Design

Optical design is a broad term used to describe the science and engineering of controlling the propagation of light. This section primarily concerns imaging – optical design with the goal of capturing or producing images. To further narrow the scope, only refractive optical systems are described, with an emphasis on the characteristics most relevant to a practical light-field display. Generally, practical optical design for imaging requires balancing resolution against cost. High-resolution imaging requires more complex lens designs, sophisticated materials, and more physical space, all of which increase cost. Conversely, less expensive optical elements tend to cause greater optical aberration, poorer resolution and/or non-uniform resolution across the field of view. Nevertheless, for a given budget, innovative designs can yield better resolutions, and at lower cost.

This section briefly reviews the history of photographic lenses for a practical reason: early lens designs were simpler than today's multi-element lenses (Kingslake, 1989). Since a light-field display may require many hundreds or thousands of lenses, it is likely more practical for it to be based on simple lens designs that can be manufactured at low cost, provided they still give acceptable performance. This section begins with the fundamental lens aberrations, and then discusses how early lens designs balanced the intrinsic tradeoffs, and how they can be beneficially adapted for a modern light-field display.

2.4.1 Lens Aberrations

Before reviewing the various types of lens systems, it is useful to lay out the criteria that makes a good lens. This section provides an overview that focuses on the requirements of an intermediate-resolution light-field display:

1. Spherical Aberration: In early lens fabrication, it was only feasible to grind a glass surface into a plane or a portion of a sphere. Unfortunately, a partial sphere is not an ideal surface shape for focusing light, since light parallel to the optical axis has a focal length that depends on its distance to the optical axis, as illustrated in Figure 7. The result is called spherical aberration. Many modern lenses are made non-spherical (aspherical) to avoid this. For a light-field display it is reasonable to assume that lenses can have an aspherical surface profile, as they are likely to be created by a numerically controlled manufacturing process. By increasing the degrees of freedom for the shape of the lens, the lens surface can be designed to be nearly free of spherical aberrations, as illustrated in Figure 7.

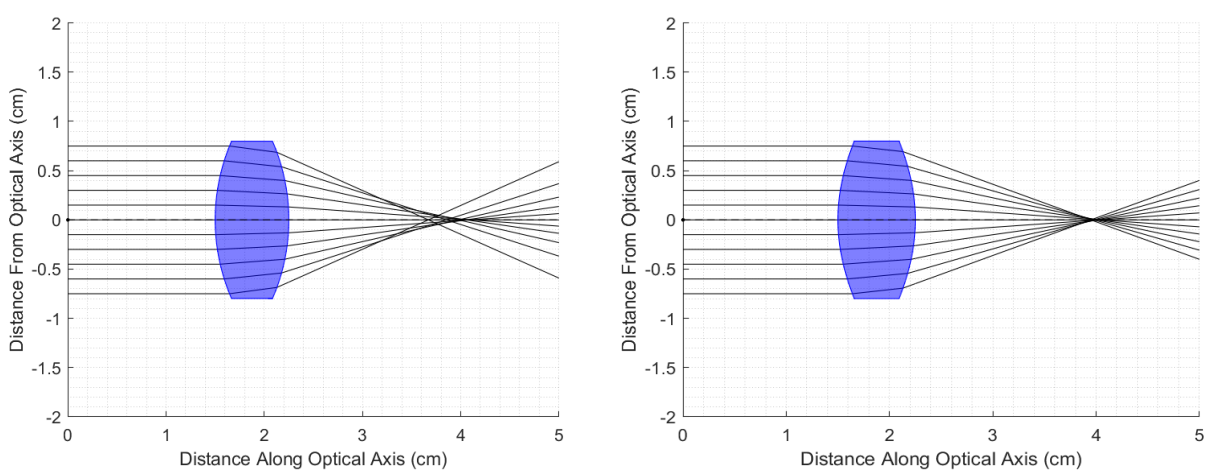

Figure 7: Spherical aberration (left), reduced with an aspherical lens (right)

2. Chromatic aberration: Different wavelengths of light are refracted differently by the lens material, and so have different focal planes, as illustrated in Figure 8. The earliest lenses to solve this problem, achromats, mitigated this form of aberration by using two different types of glass to approximately balance the focal length for all wavelengths. The use of multiple materials adds great complexity and cost to the lens manufacturing process, so would be a limiting factor for a light-field display requiring thousands of lenses.

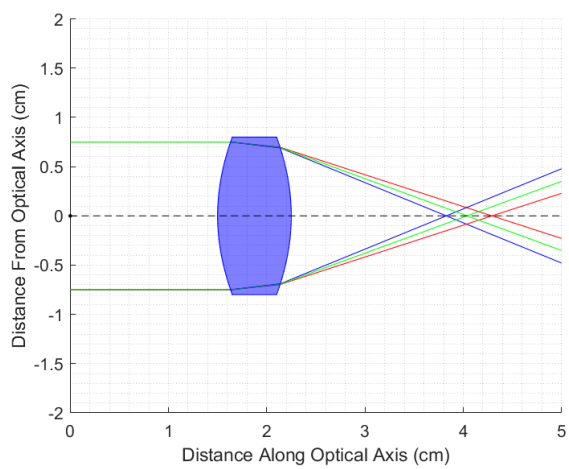

Figure 8: Chromatic aberration of a non-achromatic biconvex lens

3. Field Curvature: Field curvature refers to the phenomena whereby the focal locations of a lens form a curved surface rather than a plane, as illustrated in Figure 9. This is significant for the design of a light-field display as the image formation surface is likely to be a flat plane. The better the lens can focus light onto a plane, a property known as low field curvature, the better the performance of the display; this is especially useful for wide fields of view. For this reason, a key factor in the lens design for the light-field display is the field curvature.

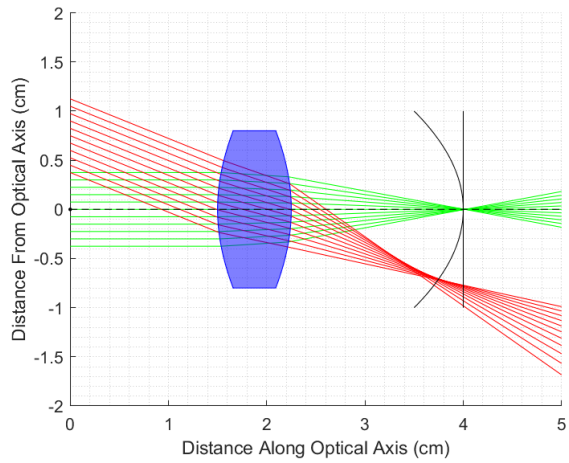

Figure 9: Field curvature aberration

4. Image distortion: Distortion is a term used to describe how parallel lines in a scene appear in the image formed by a lens. They may be either pincushion or barrel, as illustrated in Figure 10. This was a concern for early lens designs, since many early photographs were of buildings; hence, correcting for distortion was a high priority. For a light-field display it is less critical to minimize lens distortions because they can be compensated for by applying an inverse distortion to the digital image to be displayed.

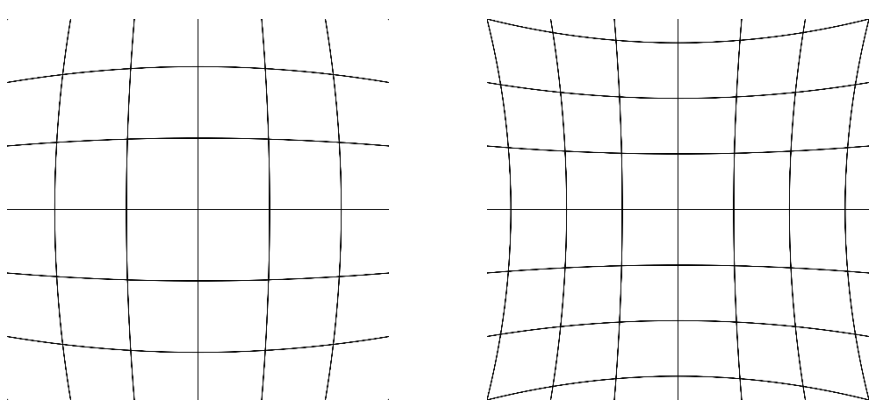

Figure 10: Barrel distortion (left) and pincushion distortion (right)

2.4.2 Lenses and Lens Systems

A lens system comprises multiple individual lens surfaces, which function to minimize the distortions described in Chapter 2.4.1. This section describes the early history of several systems that are relevant to the design of any light-field display.

1. Single-element plano-convex and biconvex lens: The earliest photographic lens was the biconvex lens, comprised of a single solid surface with both sides shaped with a convex surface (Figure 11). It was discovered that these lenses were inadequate for high quality photography, primarily due to field curvature – the inability to focus an image over a large film plane. Furthermore, they also suffered from nearly all of the other aberrations discussed in Chapter 2.4.1. Despite their shortcomings, plano-convex lenses are often integrated into autostereoscopic and other 3D displays today, primarily due to their simplicity and low cost.

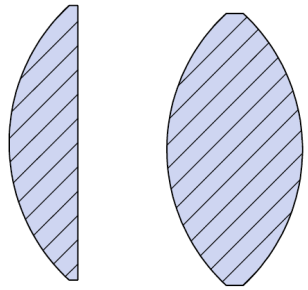

Figure 11: Plano-convex (left) and biconvex lenses (right)

2. Meniscus lens: A subsequent improvement in lens design was the meniscus lens, which significantly reduced the curvature of field. The Petzval Field Curvature (Kingslake, 1989) is defined as the sum of all the radii in a series of lenses. By having one "positive" radius

followed by a “negative” radius this curvature can be minimized. Figure 12 illustrates the Wollaston Meniscus lens, circa 1804.

Figure 12: Wollaston meniscus lens

3. Multi-element lens systems: The next stage in the evolution of lens system design was to increase the number of elements. This was driven by the need to focus more of the light onto the imaging surface, so that the shutter could be open for less time. By increasing the number of elements, the aperture of the lens system (not illustrated) can be increased, thus transmitting more light. The multi lens system enabled early photographers to significantly reduce the length of time that the subject needed to remain perfectly still, from as much as 30 minutes to 2 minutes. The Petzval portrait was the first lens of this style, illustrated in Figure 13 (United States Patent No. US2500046A, 1950).

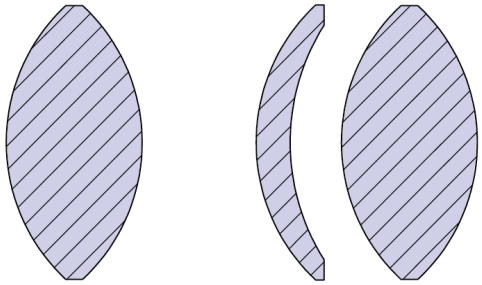

Figure 13: Petzval portrait lens system. From left to right: biconvex, meniscus, biconvex lenses

4. Globe lenses: The portrait lens system performs well over a narrow field of view (~20 degrees), but a wider angle of view is desirable for landscape photography. The Harrison & Schnitzer Globe lens of 1862 (Figure 14) was the first wide-angle lens with an 80 degree field of view (United States Patent No. 35605, 1862). One key characteristic of the design is symmetry, which reduces distortion and chromatic aberration. This lens is the basis for the optical design of the intermediate-resolution light-field display proposed in the next chapter, due to its simple design and good performance over wide angles.

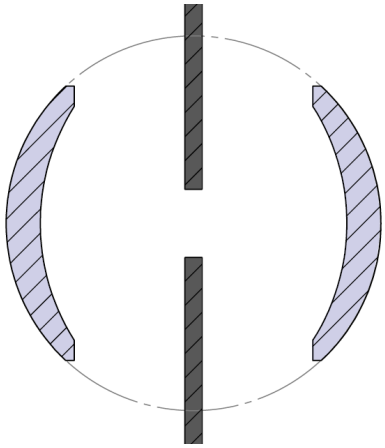

Figure 14: Harrison and Schnitzer globe lens. Formed of two identical meniscus lenses placed symmetrically on both sides of a central aperture

5. Apertures: Most lens systems employ an aperture (Figure 14) to block light rays at the periphery of a lens, where aberrations tend to be strongest. In general, the smaller the aperture, the better the image quality, but the resulting decrease in light that is transmitted through the system means lower optical efficiency. Furthermore, if the aperture becomes too small, diffraction effects around the edge of the aperture begin to dominate performance. When considering an array of lenses as in the design of a light-field display, there is an additional consideration: the gaps between the apertures, which provide no light. The

resultant non-uniform light field introduces a specific challenge: to effectively hide the array of apertures by either ensuring the gap is too small to be noticed, or to spread light laterally between the gaps. This challenge is a critical issue for an intermediate resolution light-field display, and is further addressed in detail in Chapter 4.

6. Fish-eye Lenses: The fish-eye lens was introduced in 1924 (Hill, 1924). These lenses have a very wide field of view obtained by having a first meniscus lens compress as much as a 180 degree field of view to a smaller 60 degree field of view (for example), followed by a regular lens, shown in Figure 15. It may be possible to employ an array of fish-eye lens systems to increase the useable angle of a light-field display. However, due to their complexity and cost these are not deemed practical for the initial stages of this research. These lenses suffer from extreme geometric distortions, which can be corrected by digitally compensating the image to be displayed. The Luneburg Lens and Maxwell fish-eye (Luneberg, 1944) (Niven, 1890) are theoretical designs that obtain excellent performance but require a graduated index of refraction throughout the lens material. Both the types of elements and the materials required for the fish-eye configuration would significantly increase the manufacturing cost of a lens array for a light-field display.

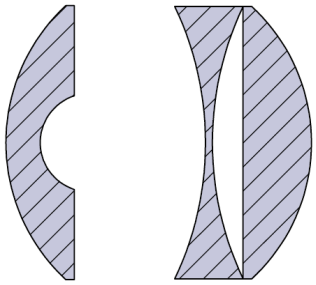

Figure 15: Beck Hill Sky fisheye lens

7. Aspherical Lenses: Up until the 1950s, it was prohibitively expensive to mass-produce lenses with anything but a spherical surface profile, which suffered from spherical aberration. However, modern lenses can be cut or molded using a surface shape termed *asphere* to minimize distortions, as illustrated in Figure 16. The surface profile is commonly represented as a conic section modified by an Nth - degree polynomial symmetric about the optical axis (Pruss, Garbusi, & Osten, 2008). However, given the high degree of correlation between a spherical term and the even power terms of the polynomial equation, it is sufficient to omit the spherical term and simply write the function as an Nth – degree polynomial with only even powers, as shown in Equation (2).

$$y = c_6x^6 + c_4x^4 + c_2x^2 + c_0 \qquad \textbf{(2)}$$

where x is the distance from the optical axis, and y is the distance along the optical axis When optimizing the aspherical lens surface in Equation (2), a high degree of correlation between the polynomial coefficients has been found (Forbes, 2007). Trying to optimize the surface often results in many local minima, which adds complexity to the optimization process. It is also difficult to measure the performance gain of adding additional polynomial terms to the surface profile, as the surface must be completely re-optimized for each degree of the polynomial. To address these limitations, the polynomial surface profile is sometimes re-factored as Q-polynomials, a set of orthogonal polynomial basis functions designed to minimize the correlation between coefficients.

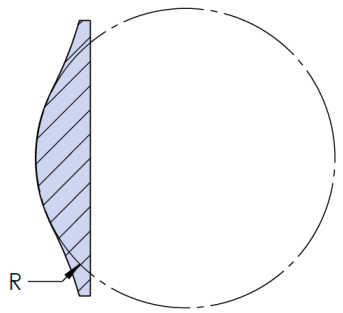

Figure 16: Asphere lens. The surface of the lens deviates from a perfect sphere of radius R, thus minimizing spherical aberrations

8. Fresnel Lenses: Fresnel lenses divide a spherical or aspherical surface profile into many smaller facets having the same curvature but smaller thickness. These lenses were first used in early lighthouses, from around 1823 (Watson B. , 1999), due to the need for a lens with large aperture and short focal length (Figure 17). They are now commonly found in many applications requiring a thin, low-profile lens, ranging from solar collectors to overhead projectors and automobile headlights. The combined features of large aperture, short focal length, and thin profile also make these very interesting for the application of light field displays. They can be formed in tiled arrangements on large flat sheets of transparent polymer at very low cost. Their major limitation is the diffraction caused by the multiple edges, reducing resolution. For an intermediate-resolution light field display, however, resolution may be of less importance than focal length, cost, wide-angle-performance, and size. The other limitations of a single Fresnel lens are the same as those encountered with the bi-convex lens, namely curvature of field, chromatic aberration, coma, etc. In Chapter 4 the use of Fresnel lenses is discussed as a viable approach for an intermediate-resolution light

field display, especially when used in combination with the Globe Lens configuration to mitigate the curvature of field and other distortions.

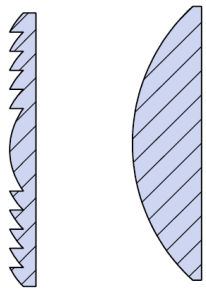

Figure 17: Fresnel Lens. The surface profile of a conventional biconvex lens (right) is divided up into many smaller facets to form a Fresnel lens (left), decreasing the thickness and material required for the lens

2.5 Light Fields

Arun Gershun coined the term "light field" to describe the amount of light travelling in every direction through every point in space (Gershun, 1939). In order to impose some practical limits on the infinite spatial and angular extent of the real-world light field, Adelson and Bergen (1991) refined this concept into a "plenoptic function" of only seven dimensions: three spatial coordinates, two dimensions of the view direction, plus wavelength and time. This was further refined by Levoy and Hanrahan (Levoy & Hanrahan, 1996) into a "4D light field" of two spatial dimensions (I_s and I_t) plus two angular dimensions (I_u, and I_v), by dropping both color and time. The 4D light field has also been called a "Spatio-Angular" representation *I(u,v,s,t):* the pixel at spatial location *(s,t)* is sometimes referred to as a "superpixel" as it can be considered to contain all the information of light passing through point *(s,t)* in the 2D image plane. A conventional 2D image is the subset of this 4D light field that is viewed from a single position. In certain cases (such as images with radial symmetry), the four independent dimensions can be reduced to a single spatial dimension and a single angular dimension, or a 2D Light Field:

$$I_{st} = \sqrt{I_s^2 + I_t^2} \qquad \textbf{(3)}$$

$$I_{uv} = \sqrt{I_u^2 + I_v^2}$$

The resulting plot of I_{st} vs I_{uv} can be considered a 2D spatio-angular light field plot, where the x-axis represents the spatial dimension, and the y-axis represents the angular dimension.

2.5.1 Light-Field Displays

A light-field display is able to simulate, to varying degrees of accuracy, the complete visual information encountered in the world. This section summarizes the background knowledge required to understand the capabilities and limitations of light-field display technologies. Specifically, limitations are discussed in terms of how they may be avoided, or at least mitigated, for an intermediate-resolution display of the kind envisioned in this thesis.

One way to understand the purpose of light-field display technology is to compare it to the more familiar stereoscopic technology commonly used to view "3D" movies. By wearing special eyewear, these displays provide a separate view to both left and right eyes, thereby providing the binocular cues needed to create a strong sensation of depth. However, several limitations hinder performance and comfort for long term-viewing, owing to conflicts between depth cues. Firstly, although these displays provide convergence and binocular parallax, the viewer's eyes are still focused on the screen rather than the 3D object. This causes a mismatch between accommodation and convergence, which is especially severe for close objects (Hoffman, Girshick, Akeley, & Banks, 2008). Secondly, stereoscopic displays must either assume or adjust to the separation of the viewer's eyes. If incorrect, the views to each eye are no longer consistent with real-world experience. An additional challenge is the assumption of *Epipolar Geometry* – that the viewer's head is horizontal. If the head tilts even slightly, an incorrect view is observed by each eye, which causes discomfort and interferes with depth estimation (Banks, Read, Allison, & Watt, 2012). Finally, as a viewer moves around a 3D display, the viewer's perspective of the display changes, but the perspective of the image shown on the display does not change in the same way, resulting in a conflict with motion parallax. An autostereoscopic display removes

the need for eyewear by adding an optical element that projects each of the two views at different angles, aligning with the viewer's eyes. In practice this requires little or no movement of the viewer from within specific viewing regions.

An enhancement to stereoscopy is multiscopy, where more than two images are presented, with not only different images to each eye but also different images from different viewing positions. In addition to convergence and binocular parallax, this technology also provides motion parallax as the viewer's head moves. Multiscopic display technologies are divided into two classes. The first uses head tracking to ensure that both eyes receive the correct image. This does not meet the needs for this project since it is limited to a single viewer, and requires high precision, high accuracy, high speed, and highly reliable viewer position tracking. The second class provides multiple views simultaneously, eliminating the need for viewer tracking. With sufficient density of viewpoints, the accommodation depth cue can be provided, thus eliminating more potential conflicts between depth cues and so increasing the likelihood of a comfortable, natural viewing experience indistinguishable from that of the real world (Balram & Tošić, 2016).

The first of this latter class of displays was invented by Gabriel Lippmann, who coined the term *Integral Imaging* (Lippmann, 1908). His idea was to use an array of tiny plano-convex lenses in front of a photographic emulsion to capture a scene from a large number of slightly different angles. When developed and illuminated through the same array of lenses, the system works as an array of projectors, each projecting light from a slightly different viewing position for each point of the image. When observed with the human eye, this should effectively reproduce the full light field as captured at the scene. While limited in viewing angle and image quality, this

technique formed the basis for what is today referred to as *light field imaging* (discussed in more detail in the next section). Figure 18 illustrates the basic principles. An array of lenses is placed in front of an LCD display, or alternatively, a printed or transparent image. Each lens projects a different image in several different directions. When viewed from different positions A and B, a different image is displayed to the viewer.

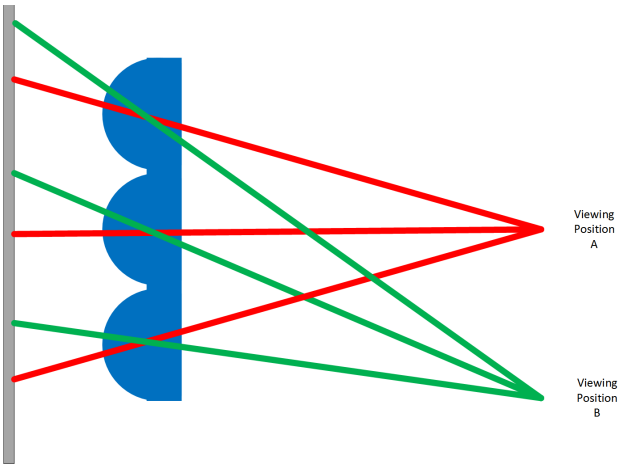

Figure 18: Illustration of a light-field display viewed from two viewing positions

Light-field displays hold great promise in minimizing conflicts between depth cues by providing the accommodation and motion parallax depth cues in addition to the binocular parallax and convergence cues, and without the need for special eyewear. However, this technology is still in its infancy and faces many technical hurdles before high quality images can be achieved. The most significant technical challenge is dealing with the massive increase in information density required – what could be considered as the *resolution challenge*.

2.5.2 The Resolution Challenge

While a conventional 2D display produces only a single image of *SxT* pixels, a light-field display must produce an image of *UxV x (SxT)* pixels, where *U* and *V* are the number of horizontal and vertical angular dimensions, between one to three orders of magnitude depending on the desired quality of the display. This is the resolution problem: the information density required for the capture, processing, and display of a light-field display is orders of magnitude greater than the equivalent 2D display, and orders of magnitude greater than can be achieved in a practical manner by today's existing technologies.

To accommodate this challenge, light-field displays today allow a slight reduction of the spatial resolution *(SxT)* in order to enable a corresponding increase in angular resolution *(UxV)*. For most display applications, the resulting compromise to the 2D image quality is one of the major hurdles to wider adoption of light field technologies today and is likely to remain until the resolution of 2D capture and display devices exceeds the capabilities of the human visual system. In order to produce both a high spatial and angular resolution image with a value of *UxV* on the order of ten views per degree (Ludé, 2018), corresponding to 1200 x 1200 angular views, the resolution of imaging systems must improve by over three orders of magnitude in each direction from where they are today. Much greater angular resolution is required to meet the capabilities of the human visual system, which could exceed 120 views per degree to achieve the binocular acuity discussed in Chapter 2.2. Furthermore, the wavelength of visible light (4.0×10^{-7} m to 7.0×10^{-7} m) places a fundamental limit on the size of optical and electro-optical elements, further inhibiting very high-resolution imaging.

In addition to image resolution, light-field displays require a matching high-resolution array of lenses. These lenses should ideally be invisible at a reasonable viewing distance, and therefore should be a fraction of a millimeter in size. However, such lenses would have their resolution substantially limited by diffraction. Small arrays of high-resolution lenses of the required size have been made, but never at the size of a typical window.

The resolution challenge further extends into other technical areas, such as capturing very small numbers of photons, and storing, distributing, and processing the same four orders of extra information. The net result is that at least for the immediate future, light-field displays may have niche applications in specific industries, but the benefit of the enhanced angular resolution does not outweigh the additional cost or the degraded image quality for general-purpose viewing and everyday use.

2.6 Chapter Summary

The key observation that drives this thesis is that although high-resolution light-field displays are not yet feasible, for at least one application - that of providing an enhanced sense of depth - a high spatial resolution field may not actually be necessary. Although an intermediate resolution light-field display may not be suitable for viewing high-resolution information, it can provide the majority of the depth cues when viewing the real world through a moderately blurry window of the type architects often select for the purpose of privacy.

A second observation is that previous research into the perception of space as summarized in Chapter 2.3 indicates that one of the key benefits of having a window may be the provision of a distant view, and that a "virtual window" that provides a sensation of a distant view via multiple depth cues may offer much the same psychological and physiological benefits as a real window.

To verify these initial observations, Chapter 3 describes an experiment to confirm whether, as suggested by perception studies, a sensation of depth can be perceived in the presence of significant angular diffusion. The key is that depth perception can operate at lower spatial frequencies than required for 2D spatial processing generally.

Chapter 3: Perception of Depth in Blurred Displays

The research reviewed in Chapter 2 indicated that peak sensitivity of human vision for binocular and motion parallax depth cues occurs for significantly lower spatial frequencies than those for 2D patterns of luminance change. To evaluate how this relates to depth perception in an intermediate-resolution display, an experiment was devised to measure depth acuity for humans under varying levels of optical diffusion. This chapter describes the experimental design and results, which are then used in Chapter 4 to guide the design of an intermediate-resolution light-field display.

The purpose of the depth acuity experiment was to evaluate how depth perception is influenced by both blur and an array of apertures placed between the observer and the viewed scene. The goal was two-fold: (i) assess possible parameters of the optical design, and (ii) verify that a practical, low-cost optical design could in theory provide a sense of depth. To characterize depth perception, the experiment measured the just noticeable difference (JND) in depth between two objects under various conditions. The JND can help determine the number of depth planes that a light-field display should ideally provide. Conventional 2D screens only provide one depth plane; a real scene may comprise an infinite number, although the depth resolution of the human visual system will set an upper bound on this. The number of perceivable depth planes of an intermediate-resolution light-field display should be somewhere between these two extremes.

3.1 General Approach

The approach used here was based on an apparatus that moved a test stimulus some distance in front or behind a reference stimulus. An advantage of this approach is that it allows free

movement of the participant (and so enabling motion parallax) without the need for special eyewear or viewer tracking systems. It also allows natural viewing, where the viewer could move as fast as they wanted, in any direction they wanted, and without any latency or lag between movement and the stimulus update. The drawbacks for this system are that the reference distance is limited to the order of 1.0 m, and the range of movement of the test stimulus was limited as well. These disadvantages, however, were not felt to be prohibitive. A pilot study confirmed that the monocular depth cues had been mitigated or nullified, and verified that the range of movement of the physical apparatus was sufficient for nearly all observers. The results of this pilot study are discussed in Section 3.4.

To simulate the optical diffusion of an intermediate-resolution display, various diffusers were installed between the viewer and the stimulus. Several commercially available diffusers were tested, producing an angular diffusion close to the amount of blur anticipated in the display system. For the pilot study, three diffusers manufactured by Rosco (Rosco) were used. The angular diffusion for each material was measured with the apparatus illustrated in Figure 19.

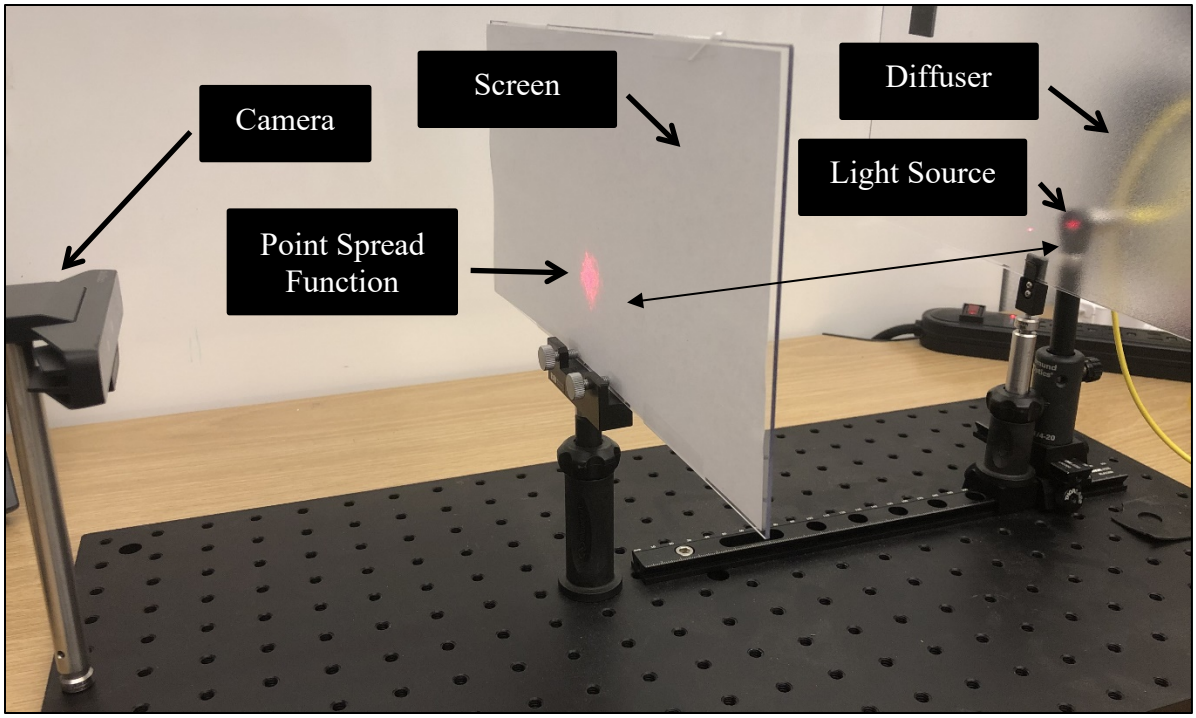

Figure 19: Apparatus for measuring diffusion

The characteristics of each diffuser was measured by directing the beam from a laser light source through it. The diffused light propagated 25 cm to an imaging screen, and a calibrated camera measured the relative intensity - and thus angular breadth - of the point spread function (PSF). A least-squares optimization then calculated the best-fit between the measured data (the PSF) and a Gaussian function with radius *r;* the radius was then converted to angular diffusion $\sigma = atan(r / 0.25m)$. The measurements for all three diffusers are illustrated in Figure 20, with the measurements in blue and the best-fit Gaussian profile in red. Although the diffusers do not exhibit an exactly Gaussian distribution, the standard deviation provides a reasonable estimate for average amount of diffusion. The best-fits for the three types of diffusers are:

- Light diffuser: *Light Opal Tough Frost (*part 3020) σ = 1.4 degrees
- Medium diffuser *Opal Tough Frost (*part 3010) σ = 1.8 degrees
- Strong diffuser *Light Tough Frost (*part 3009) σ = 2.6 degrees

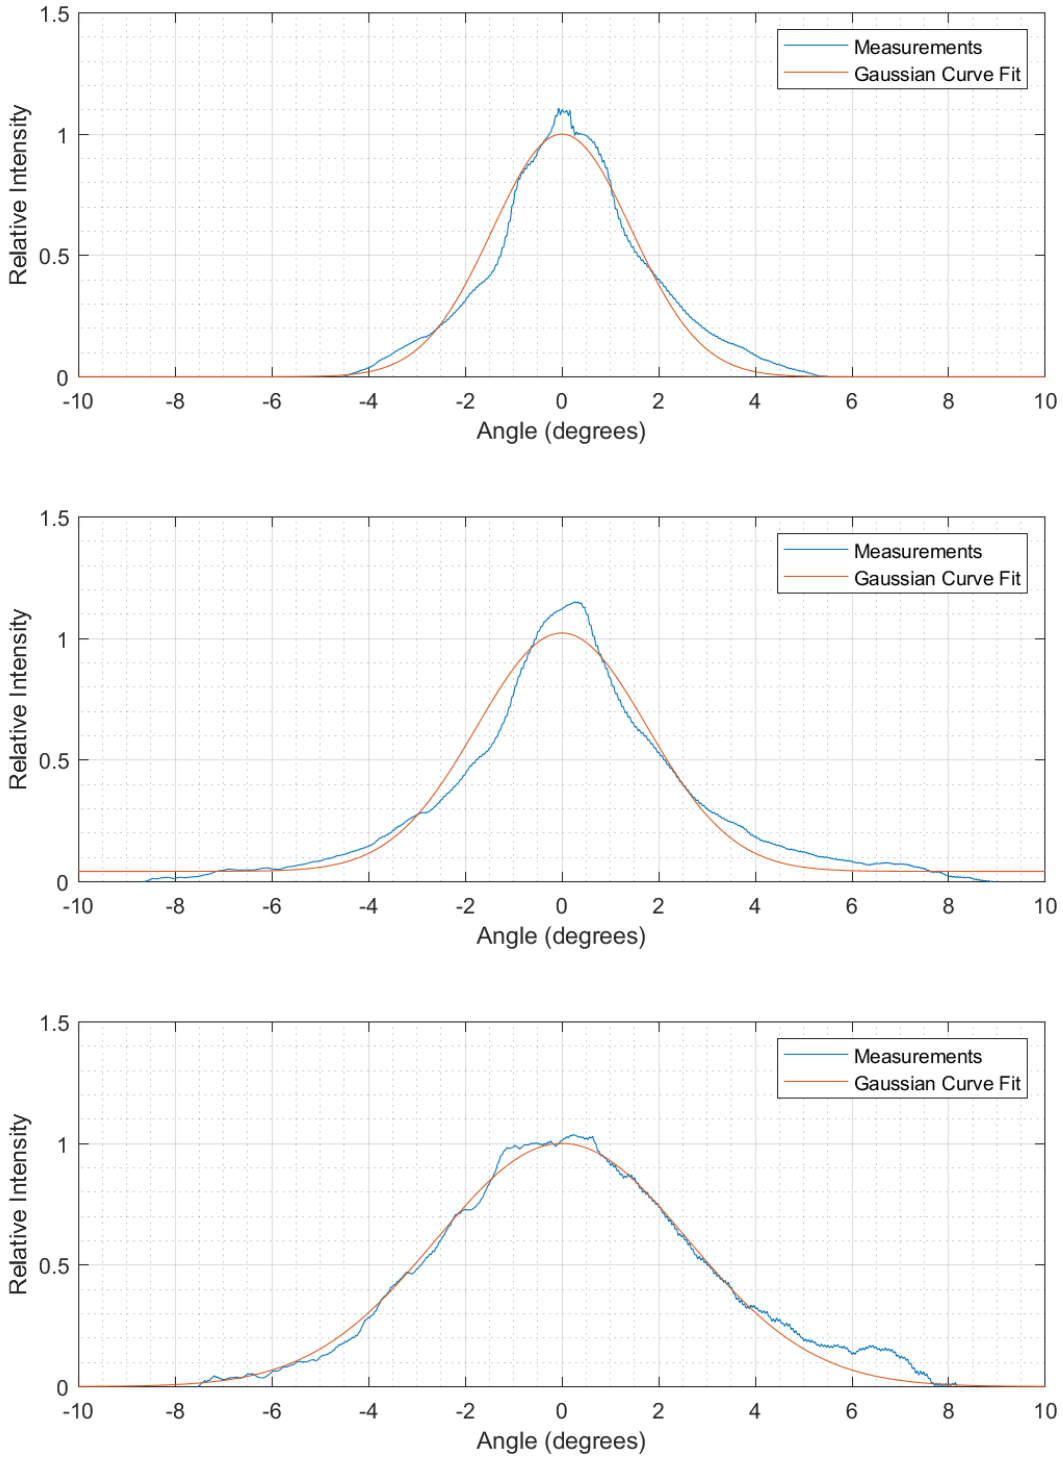

Figure 20: Diffusion profiles for light (1.4 degrees), medium (1.8 degrees), and strong (2.6 degrees) diffusions used in pilot study

For the main study, only two diffusers were selected. The “light” diffusion was the same as in the pilot study, which was found to be the minimal diffusion needed to make invisible the reflected light from the experimental apparatus. The “medium” and “strong” diffusers in the pilot study were found to yield more angular blur than was desirable. What was used instead was the SXR-9829 Glassfrost diffuser, manufactured by Decorative Film (Decorativefilms); its diffusion profile is illustrated in Figure 21. It is reasonably uniform for angles up to 1.0 degrees, and then falls rapidly until about 4.0 degrees, beyond which there is very little diffusion. The standard deviation of the diffuser was 2.2 degrees.

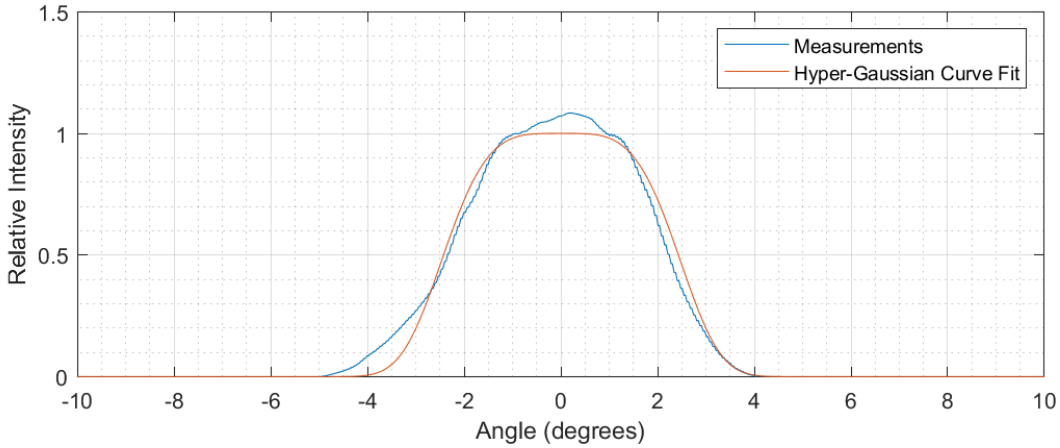

Figure 21: Angular diffusion profile for "strong" diffuser used in main study. The standard deviation is 2.2 degrees.

To simulate the lens array, an additional layer was installed between the viewer and the stimulus during the main study. This layer consisted of a printed transparency of a regular array of perforated holes on a black plastic sheet, illustrated in Figure 22. The size and spacing of the holes were determined by estimating that a practical dimension of the final design would result in roughly 6.3 mm (1/4”) apertures spaced at 2.5 cm (1”) intervals, and that a practical minimum

viewing distance would be 1.0 m. The spacing of the lenses determines the trade-off between spatial resolution and angular resolution: greater spacing results in lower spatial resolution and higher angular resolution. This trade-off is further discussed in Chapter 4. For the experiment this was scaled down, resulting in 1.6 mm (1/16”) apertures spaced at 6.3 mm (1/4”) intervals, and reduced viewing distance to 25 cm.

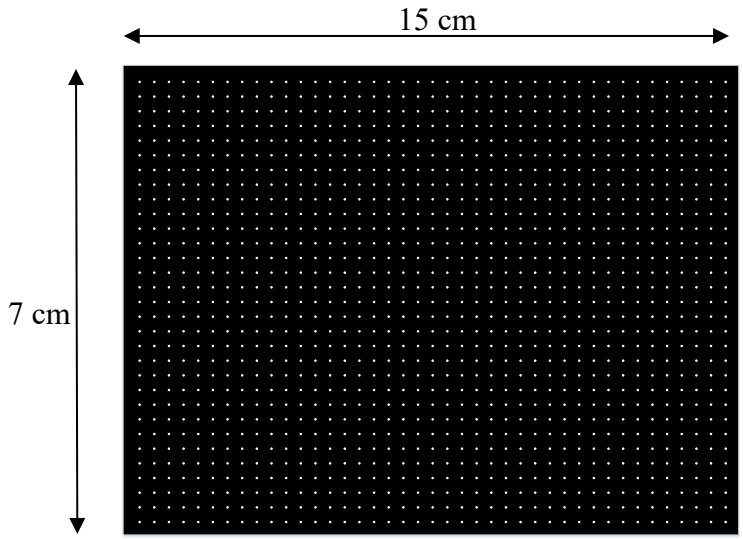

Figure 22: Simulated array of lens apertures

Observers were able to view both stimuli through a viewport measuring 15 cm x 7 cm, as illustrated in Figure 23.

Figure 23: Subjects participating in the user study, showing controller and relative position of the viewport. Lights were turned off for the study.

3.2 Displaying the Stimuli

The experiment apparatus was based on the apparatus used by Howard and Dolman (Howard H. J., 1919). The viewer was presented with two mobile phones positioned side-by-side, as illustrated in Figure 24. Each was an Apple iPhone X (iPhone X), chosen because the Organic Light Emitting Diode (OLED) can display fully black pixels, avoiding any unwanted depth cues about the surface of the screen. Both phone screens were adjusted to have the same luminance (100 cd/m^2). For the duration of the experiment, the screens were configured to remain as stable as possible, by disabling ambient light adjustment and adaptive power management.

The reference stimulus was a gray disk 2.5cm in diameter, placed on the observer's right at a distance of 1m. The test stimulus was initially at the same plane as the reference stimulus, separated laterally by 16cm. The test stimulus, controlled by the experimenter, was actuated by a stepper motor, with a range of movement of ±25cm. To minimize the depth cues associated with relative size, the size and position of the test stimulus was continually re-computed to maintain a constant visual angle from the observer. For the main study, the observer was free to move to produce motion parallax. The observer could move laterally by an unlimited amount, and longitudinally (forwards and back) between 0.9m to 1.5m from the reference stimulus. Observers could only look through the viewport.

The test administrator was present in the room for each test subject, to answer questions and ensure that each test was conducted properly. The stimuli were shown in a completely dark environment. All light sources were blacked out including power and signal LEDs. The stimuli were clearly visible as circles, and nothing else was visible through the viewport.

The experimental design was verified by placing a camera at the nominal viewing position and capturing an image of the stimuli at the zero position as well as both positive and negative maximum ranges (±25cm). In each position it was impossible to measure any difference in the size, position, or luminance of the two stimuli, confirming that the monocular depth cues had been correctly accounted for, with the exception of accommodation. Fortunately, given the diffusion, the accommodation depth cue would not have any practical impact as it requires high spatial frequencies.

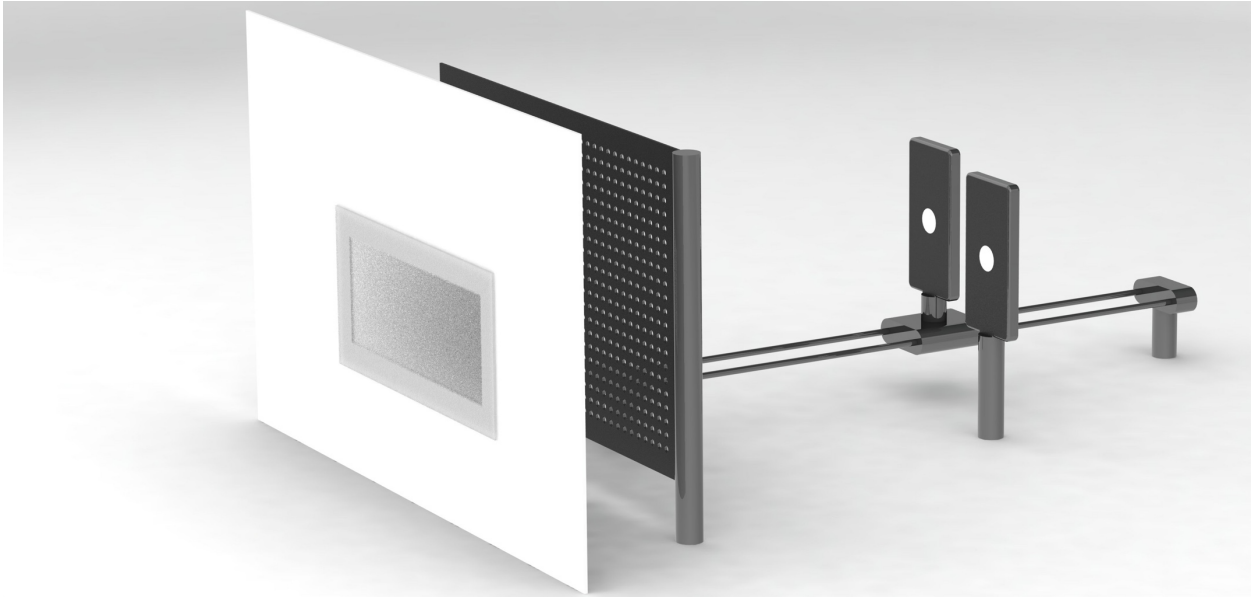

Figure 24: Illustration of the experimental apparatus

3.3 Adjusting the Position

The test stimulus was moved along two guide rails using an Axis360 linear stage manufactured by Cinetics (Cinetics Axis 360). Movement was controlled by a belt which in turn was controlled by a rotational stepper motor. The smallest step for the motor was 3.0×10^{-2} degrees, which corresponded to a distance of 6.5×10^{-6} m. At the 1m reference distance this allowed testing of visual depth acuity down to 8.6×10^{-5} degrees, which was less than 1% of the anticipated JND. The repeatability of the linear stage was verified by executing movements back and forth over the whole range (±25cm) 100 times and confirming that there was no measurable drift.

The motor emitted a mild but audible sound proportional to the speed of rotation. To prevent this from providing a clue as to the position of the stimuli, the test stimulus was moved to each new position in two steps. The first returned it to the zero position, and the second moved it either forwards or backwards to the new position. The participant was unable to determine if the second movement was moving the test stimulus forwards or backwards, and so could not use the sound to guess at the answer (the duration of the sound revealed how far the test stimulus moved, but not in which direction). Participants responded using a controller (Contour ShuttleXpress) connected to the host computer. Two buttons were clearly labeled, either left to select the left stimulus as closer, or right to select the right stimulus as closer.

3.4 Choosing the test positions

The task of the subject was to select which of the two stimuli appeared closer. Each study began with a large separation of depth, and after each response the separation was decreased or increased for a correct or incorrect response respectively. The test separations used during each trial were selected by the QUEST procedure (Watson & Pelli, 1983). This method works by modeling the psychometric function and choosing samples values to iteratively refine the function parameters to obtain an optimal fit and estimate of the JND for a threshold of 75% correct. This was a Two-Alternate- Forced-Choice (2AFC) study, so when the separation is less than a just noticeable difference the participant is expected to respond correctly less than 50% of the time. A pilot study determined initial parameters for the experiment:

1. Initial guess of threshold: 15 cm
2. Initial guess of standard deviation: 10 cm
3. Beta: 1.6
4. Delta: 1/25
5. Gamma: 0.5

Beta, delta, and gamma are dimensionless parameters of a Weibull psychometric function (Fréchet, 1927). Beta controls the steepness of the psychometric function, while delta is the fraction of trials on which the observer selects blindly or makes a mistake when well above the threshold; gamma is the fraction of trials that are correct when below the JND threshold. At the start of each test, an initial separation was chosen to be easily visible for nearly all participants (15cm), with the test stimulus randomly selected to be closer or further from the observer compared to the reference. The separation was decreased for each correct answer and increased

for each incorrect answer, up to a maximum of 30 responses. A distinct sound indicated if the response was correct or incorrect, to enable the participant to adjust their future responses if desired. Figure 25 shows the absolute separation distance obtained by one example test.

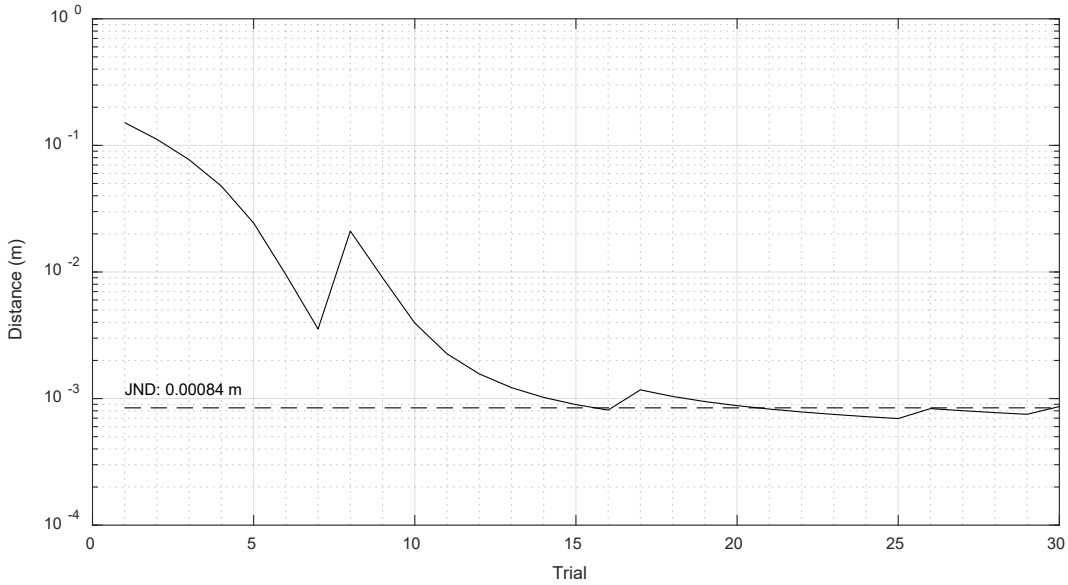

Figure 25: Trial responses for a single participant for a single test case. Each correct answer reduces the separation while each incorrect answer increases the separation until 30 trials have been completed. The dashed line shows the final JND threshold for 75% correct.

If after more than ten responses, the JND is greater than 95% of the maximum range of the slider (25cm), the study exits immediately and is recorded as “stop early”. This is to prevent frustration for the participant, since the depth JND in this case may be greater than the maximum range of movement of the experiment. When this occurs, the maximum range of the slider is reported, in this case 3300 seconds of arc. This stop early condition only occurred during the first pilot study. For all other pilot studies as well as the main study the experimental range was sufficient to measure a JND.

For each trial, the new stimulus was shown for 1s before the participant could respond. There was no time limit for a participant to respond to any one stimulus. Figure 26 shows the timing diagram of the study.

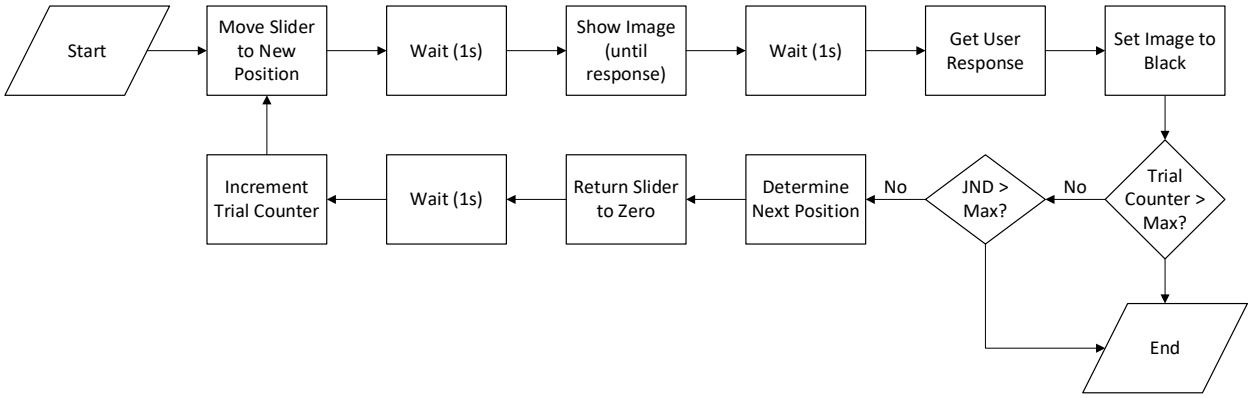

Figure 26: Timing diagram of each user study

3.5 Pilot Studies

A single subject participated in three pilot studies to validate and refine the main experiment design. The first pilot study was designed to verify that (i) there were no unintended depth cues, and (ii) that the experimental procedure could determine JNDs within the same range as the previous work discussed in Chapter 2.2.1. This study tested two conditions: no head movement with monocular vision, and no head movement with binocular vision. For both cases the head location was fixed using a chin rest, and the light diffuser was used to mask the experimental apparatus. In the monocular vision case the participant wore an eye patch over one eye. The second pilot study compared Stereopsis and Motion Parallax. This study tested three conditions: no head movement with binocular vision, head movement with monocular vision, and head movement with binocular vision. This study aimed to ascertain the relative visual acuity due to the binocular and motion parallax depth cues, in order to guide the number of test conditions required for the main study. The third pilot study compared the visual acuity for various strengths of diffusion (light, medium, and strong diffusion). For this study there were no restrictions on head movement and binocular vision.

The positions from the experiment are specified as distances. Since previous work reported visual acuity in terms of visual angle, the distance JND has been converted to visual angle in radians, using Equation (1), then reported in seconds of arc. This procedure requires knowing or assuming the inter-ocular distance *I*. 6.3 cm was selected as a representative value, to be used throughout the analysis, as the equivalent inter-ocular distance for a standard adult observer (Dodgson, 2004).

3.5.1 Pilot Study 1: Lack of Monocular Depth Cues

The first hypothesis for this experiment was that for monocular vision and a fixed head position, the subject would be unable to determine which stimulus was closer, due to the lack of any monocular, binocular, or motion depth cues. The results, illustrated in Figure 27, confirmed this hypothesis. For the first ten trials the average correct response rate was within the expected random variation of around 50%. After ten trials, the distance had reached the maximum range of the linear slider and the "stop early" condition was triggered, meaning that the actual JND was beyond the maximum range of the experiment. This result established that the experiment design had successfully controlled against the monocular depth cues.

The second hypothesis was that by adding binocular vision the experiment would produce a visual depth acuity in reasonable agreement with previous work. The results also confirmed this hypothesis, with an estimated JND for a single observer at (29.9 ± 0.2 seconds of arc), consistent with existing literature. This result established that the study design can accurately measure stereo acuity, and that the pilot subject had good stereopsis. With a larger sample size of participants, some subjects with poor stereopsis would be expected to be unable to perform this task, as has been found with previous experiments and discussed in Chapter 2.2.

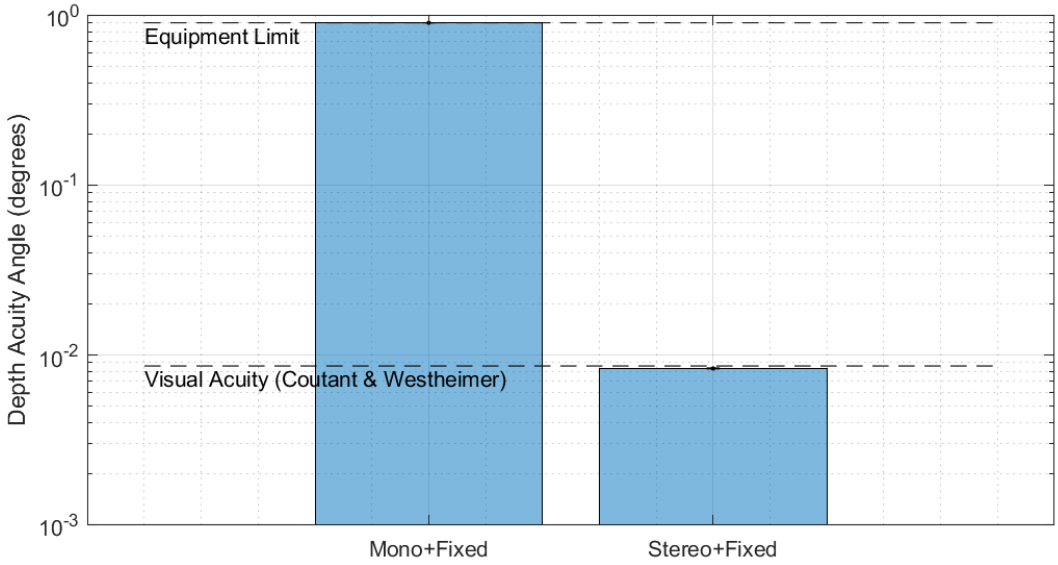

Figure 27: Results from first pilot study verifying that monocular depth cues have a negligible contribution in this experimental design, and that the experimental apparatus results in similar visual acuity to prior research (29.2 ± 0.2 seconds of arc). Error bars show the standard error for three observations.

3.5.2 Pilot Study 2: Motion vs. Binocular Parallax

The second pilot study probed the relative performance of motion and binocular parallax. A single participant repeated the experiment three times. The results (Figure 28) showed that motion parallax alone produced a larger JND (91.7 ± 22.4 seconds of arc) than binocular vision alone (29.9 ± 0.2 seconds of arc), and that the combination of both made a negligible impact on acuity compared to binocular vision alone (31.9 ± 4.0 seconds of arc). This is consistent with results found in previous work (Chapter 2.2). With multiple participants, participants with poor binocular stereopsis might be encountered, and would likely show improved depth acuity when allowed free head movements.

These results show that for participants with good stereo acuity the main study could be conducted with binocular vision alone. However, in order to test more natural viewing conditions and to be inclusive of those participants with poor stereo acuity the main study was designed to allow both binocular parallax and free head movement.

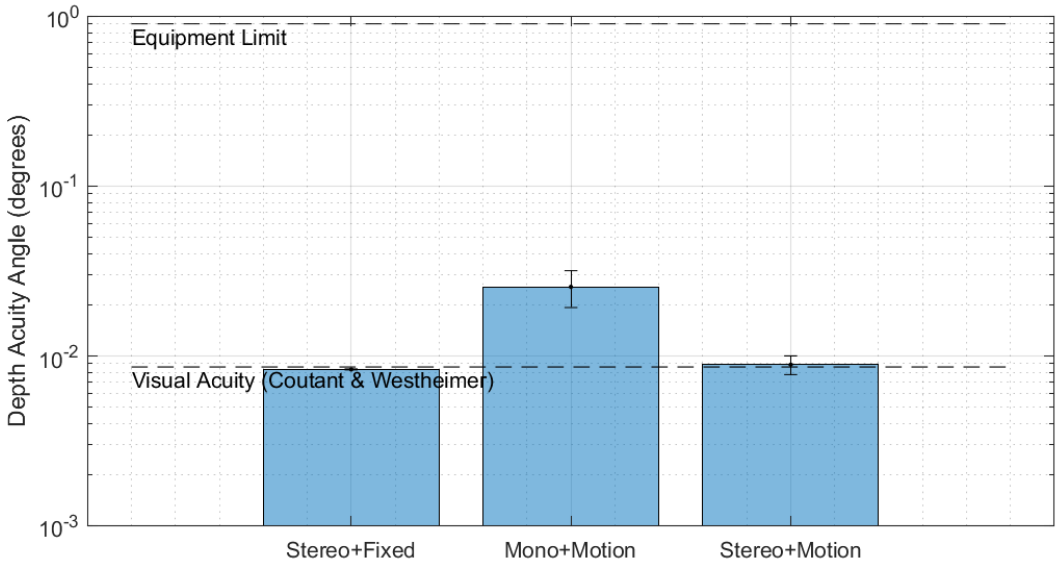

Figure 28: Results from second pilot study. The Mono+Motion case shows that motion parallax alone is a useful depth cue, although not as effective as binocular parallax for viewers with normal vision. The Stereo+Motion case shows that little improvement to stereo acuity is achieved when adding motion parallax to binocular parallax for viewers with normal sterescopic vision. Error bars show the standard error for three observations.

3.5.3 Pilot Study 3: Comparing Strength of Diffusion

The third pilot study was designed to establish useful levels of diffusion for the main study. With excessive diffusion, the range of movement of the experiment may not be sufficient to determine a JND. The hypothesis for this experiment is that increasing the diffusion reduces the amount of angular information in the signal, thereby reducing the effectiveness of binocular and motion

parallax cues. The results (Figure 29) support this hypothesis. The “light diffusion” case results in the same depth acuity as in the second pilot study, 31.9 ± 4.0 seconds of arc for a single observer, which is in line with previous experiments reported by Coutant & Westheimer (30 seconds of arc). The “medium diffusion” and “strong diffusion” in contrast, significantly reduce depth acuity, as expected. The “medium diffusion” case was found to result in a visual depth acuity of 109 ± 27 seconds of arc, and the “strong diffusion” case resulted in 217 ± 23 seconds of arc. The significant increase of the JND from “light diffusion” to “medium diffusion” is due to the reduction of spatial frequencies at and below the peak sensitivy of depth acuity as discussed in Chapter 2.2.

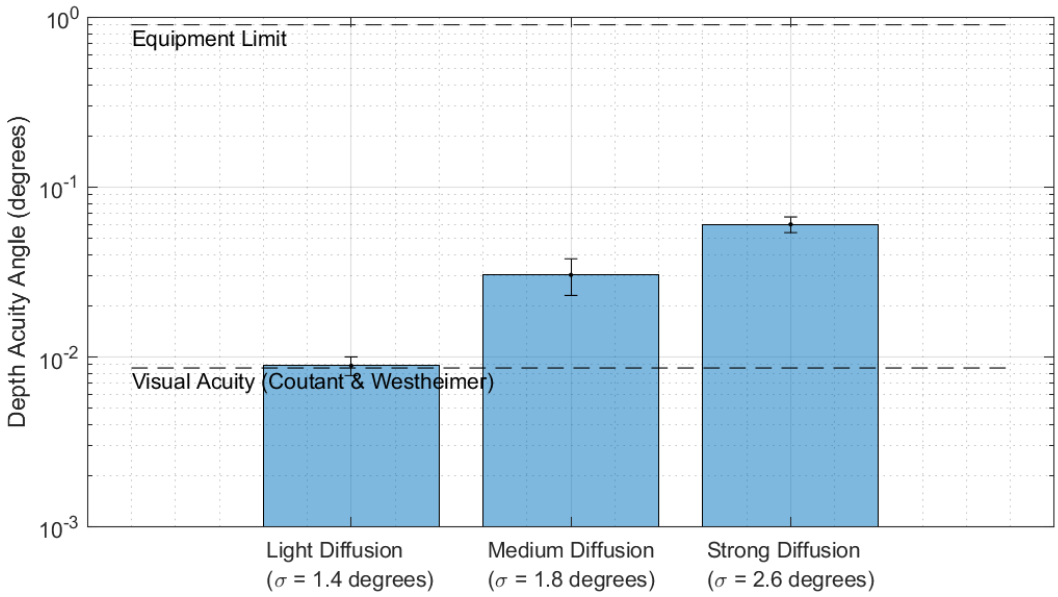

Figure 29: Results of the third pilot study showing how the depth JND increases with increasing strength of diffusion. Error bars show the standard errors.

3.6 Main Study

Based on results of the pilot studies, it was decided to allow free head movement and to test two levels of diffusion for the main study. The first diffusion level was selected to be the minimal amount of diffusion that could obscure the experimental apparatus ($\sigma = 1.4$ degrees), identical to the minimal diffusion used in the third pilot study. The resulting image, as seen through the viewing port, is shown in Figure 30. The second diffusion level ($\sigma = 2.2$ degrees) was achieved using a diffusion profile closer to that anticipated for the final display, consisting of weak wide-angle diffusion but strong low-angle diffusion, such as found on common partially diffusing glass, shown in Figure 31. A third study was conducted with the strong diffusion as well as an additional optical element to simulate the array of circular apertures according to some versions of the proposed optical design, as shown in Figure 32.

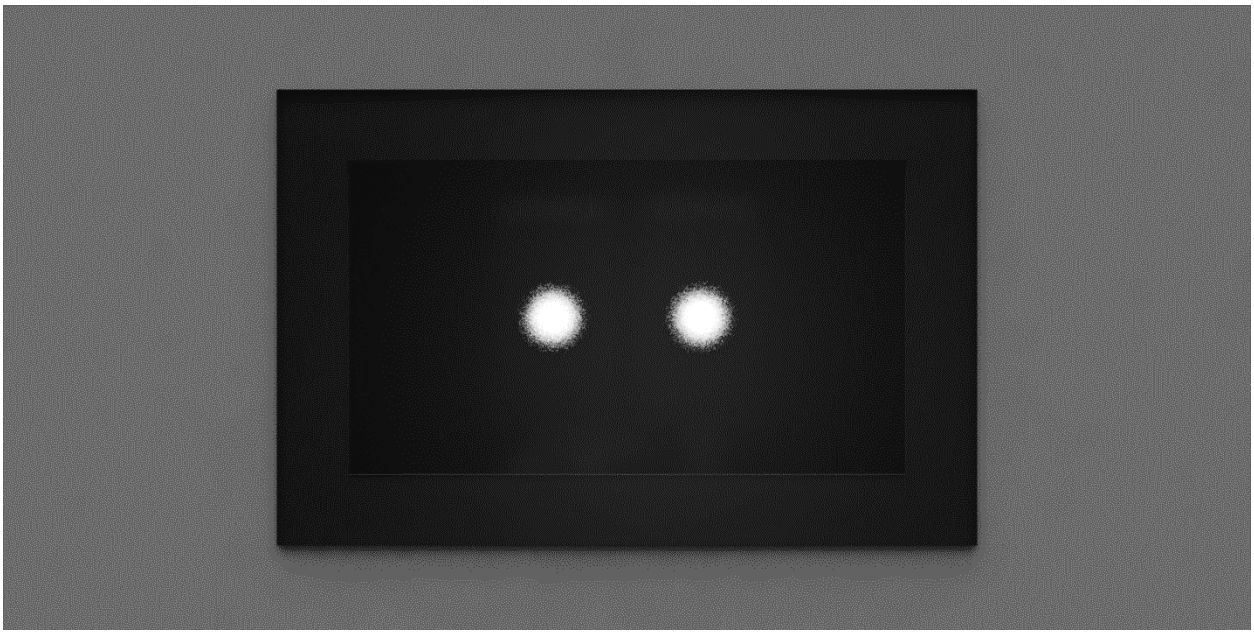

Figure 30: The experiment stimulus with light diffusion

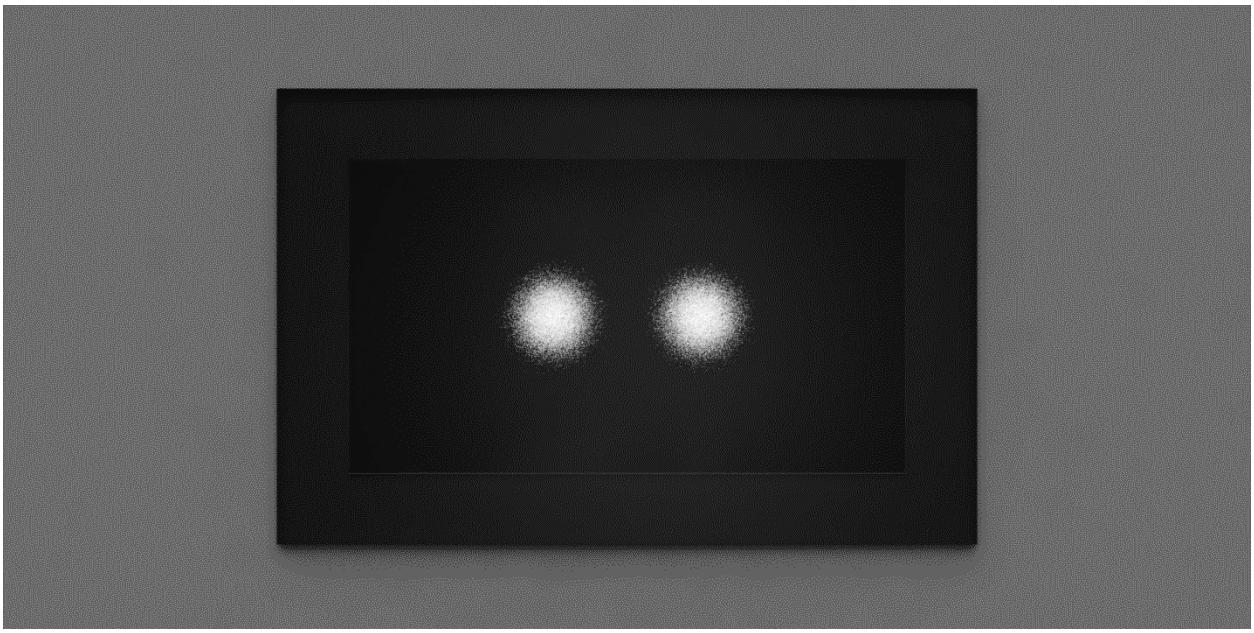

Figure 31: The stimulus with the strong diffuser

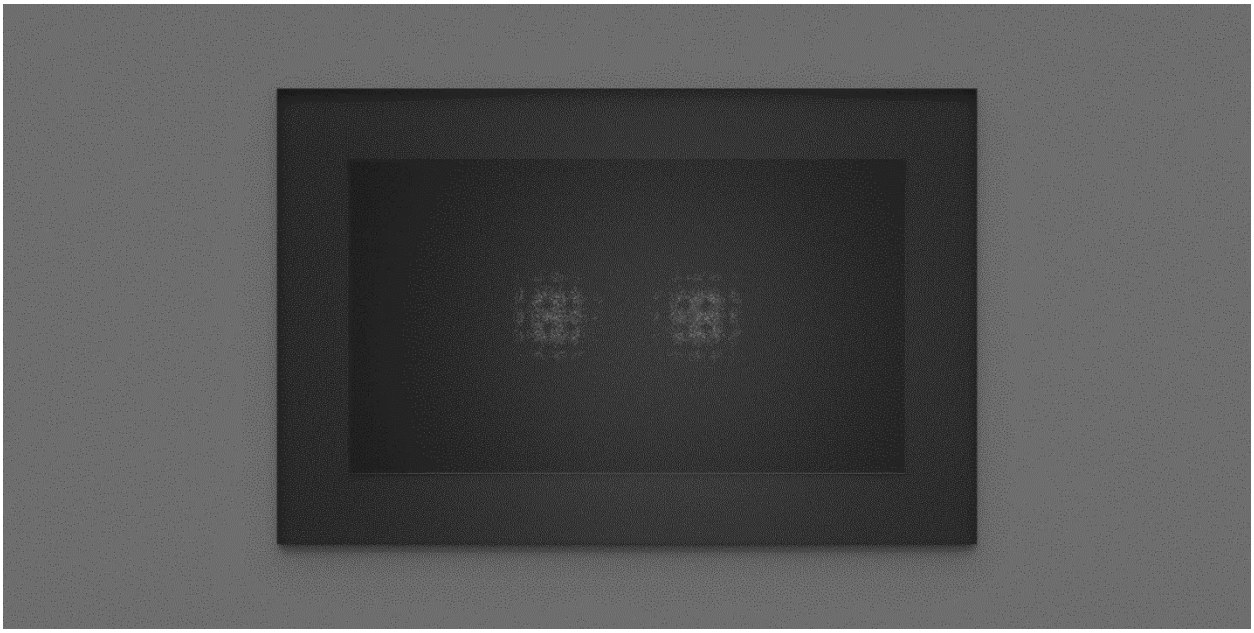

Figure 32: The stimulus with the strong diffuser and array of apertures

Before starting, instructions were shown to each participant on an LCD screen located to the left of the main experiment. The instructions explained the purpose of the experiment and the use of

the controls. Each participant then underwent a training round identical to the main experiment but with the number of responses reduced to two. Participants were free to ask questions during the training round to clarify the goals and methods.

The main study was approved through the UBC Ethics Approval process. Participants completed a consent form, then the user study training followed by the main study, followed by an optional post-study questionnaire. After the main experiment, participants were given the opportunity to complete a brief demographics survey, in which all questions were optional. No participants opted-out of the survey or any questions within it. Participant age was recorded in decades, as depicted in Table 2, with an average age of 30.

Table 2: Age ranges of main study participants

Age	Number of Participants
Under20	2
20-29	12
30-39	3
40-49	2
Over 50	3

Twenty-two participants were used. Nine of them were male and thirteen female. Ten participants reported corrected-to-normal vision while six reported normal vision without correction.

The results are shown in Figure 33. For the case of light diffusion, the average depth acuity as calculated by Equation (1) was 78 ± 12 seconds of arc. With strong diffusion, this increased to 502 ± 83 seconds of arc. Adding the aperture to strong diffusion increased this slightly to 534 ± 85 seconds of arc.

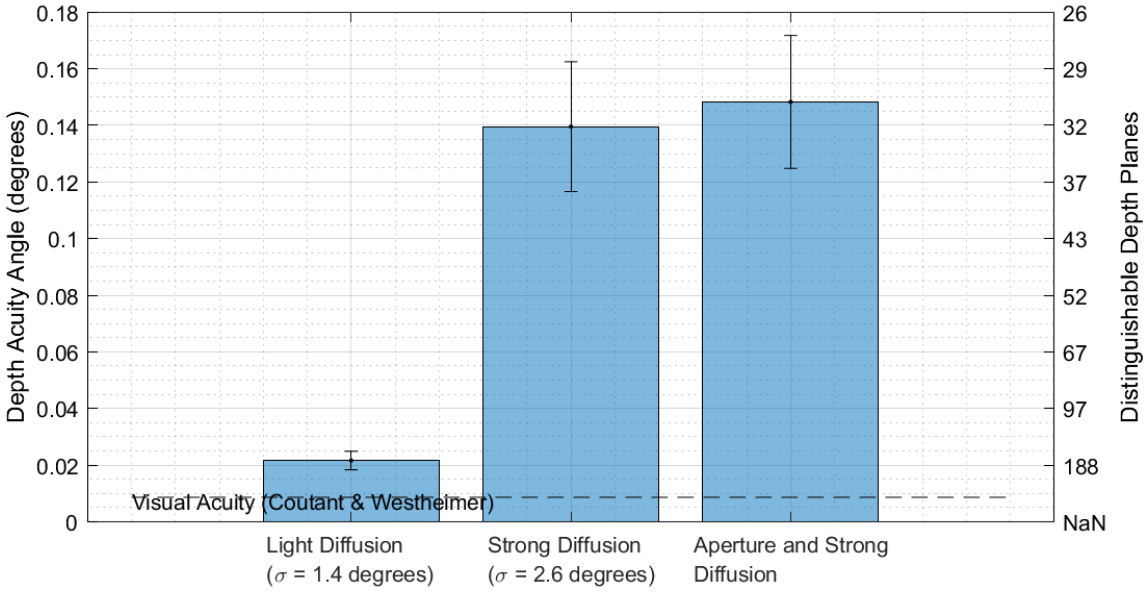

Figure 33: Main study results showing the average visual acuity for light diffusion, strong diffusion, and aperture and strong diffusion. Error bars show the standard error for 22 observers.

To portray the data in more detail, Figure 34 shows a box plot (McGill, Tukey, & Larsen, 1978), which describes the range of data that falls within the 25th and 75th percentiles: the error bars represent the full range. There is a wide range of depth acuity between subjects, even for the light diffusion case. This range is consistent with previous work described in Chapter 2.2.

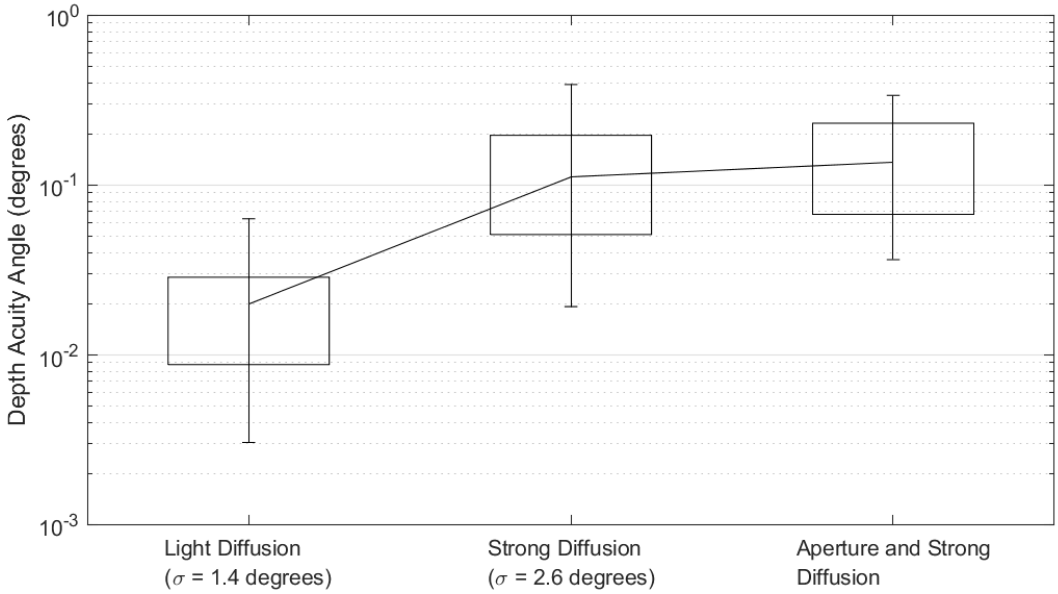

Figure 34: Box plot showing the range of data between the 25th and 75th percentiles within the rectangular regions, as well as the full range of the data within the error bars.

A key objective of the main study was to determine how many depth planes could be presented by an intermediate-resolution light-field display. This can be done by sequentially adding depth planes separated by a JND, derived from the known visual depth acuity angle. Starting at a first depth plane at 1m viewing distance, the angular threshold is converted to a distance Δz_1, using Equation (1). The distance of 1m + Δz_1 is then counted as the second unique depth plane. This procedure is then repeated, by computing a new distance Δz_2 from a new reference distance of (1+Δz_1), and counting the result as the third depth plane, and so on. This process is repeated, calculating the distance from each unique depth plane to the next, and counting the number of planes. For practical reasons, the incrementing must stop at some maximum depth; here, 1km was used, which accounts for 99.6% of the depth planes up to a point one thousand times further.

Figure 35 shows the relationship between absolute depth (measured in meters) and the number of distinguishable depth planes. For the visual acuity found by Coutant (30 seconds of arc), this corresponds to an average person being able to distinguish around 440 depth planes when looking through a window, even in the absence of any other depth cues. For the acuities found in the main study here, the number of distinguishable depth planes falls to 31-174 (Table 3).

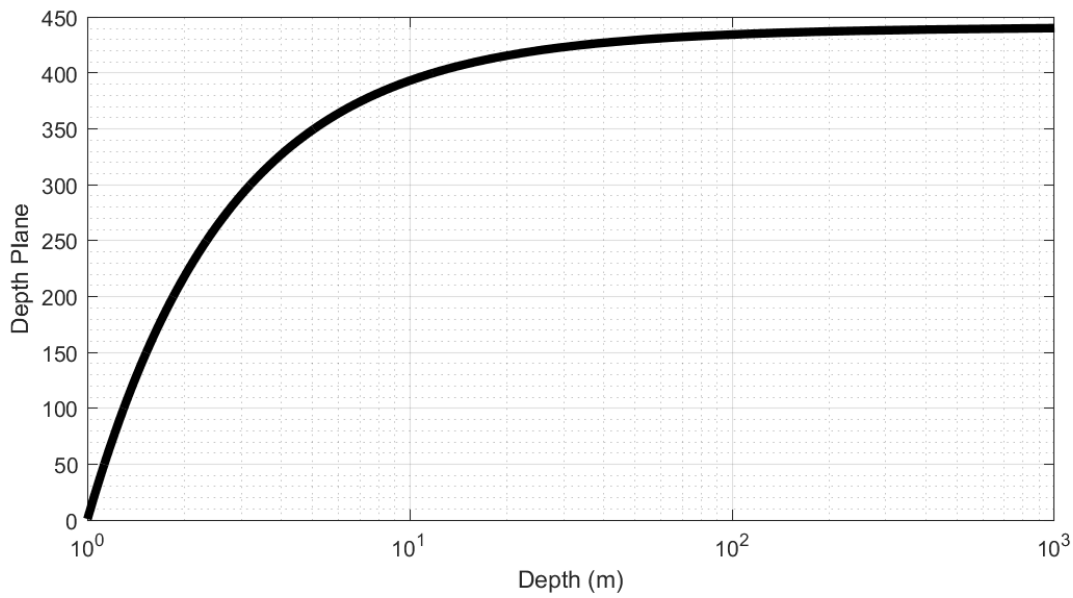

Figure 35: Relationship between absolute depth and number of distinguishable depth planes, for a visual acuity of 30 seconds of arc.

Table 3: Number of depth planes as a function of visual acuity

Viewing Conditions	Visual Acuity (seconds of arc)	Number of depth planes between 1m and 1km
Typical Visual Depth Acuity	30	440
Main Study Results with Light Diffusion	78	174
Main Study Results with Strong Diffusion and Aperture	534	31

Participants completed a post-study questionnaire (Appendix A), which asked them to report how challenging they found each experiment, and to estimate their confidence in their results, using a 5-point Likert Scale (Likert, 1932). Figure 36 reports the difficulty for each of the three test cases. Participants generally reported the "Light Diffusion" case as between Easy and Medium (average 2.4/5), the "Strong Diffusion" case as Challenging (average 4.0/5), and the "Aperture and Strong Diffusion" case as between Challenging and Very Challenging (average 4.4/5). This is not surprising, since the lack of monocular depth cues and the challenge of determining near-JND thresholds made this a challenging task.

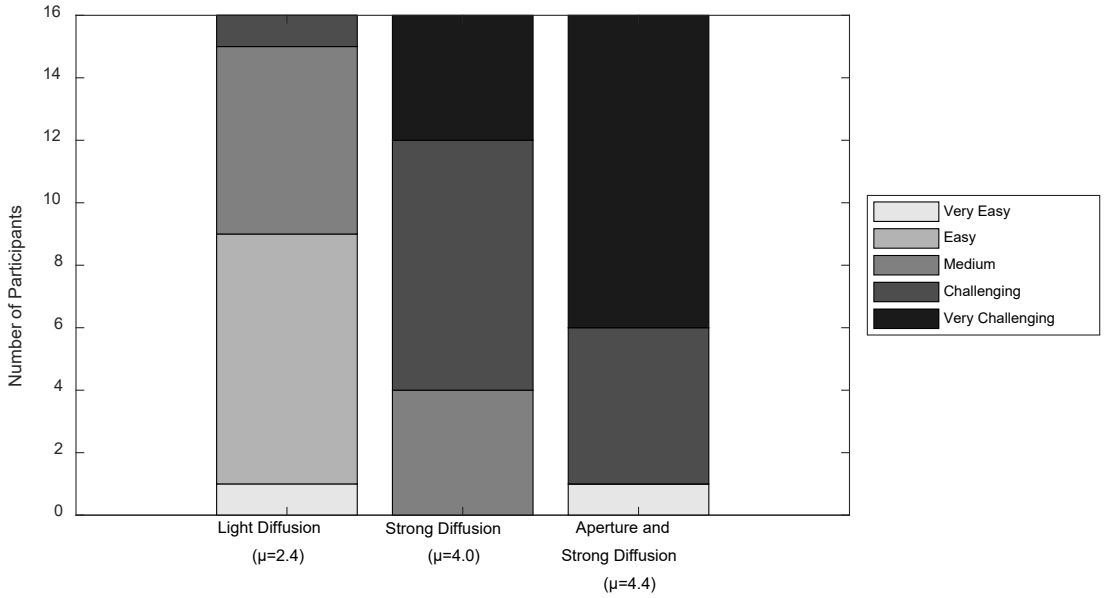

Figure 36: Difficulty level of user study task from post-study questionnaire. Participants generally reported the "Light Diffusion" between Easy and Medium, the "Strong Diffusion" case as Challenging, and the "Aperture and Strong Diffusion" case as between Challenging and Very Challenging.

Figure 37 shows how confident participants were in their responses for each of the three test cases. Participants generally reported higher confidence in their responses for the "Light

Diffusion" cases (2.3/5). In contrast, roughly half of the participants reported having low confidence (3.6/5 and 3.8/5) in their responses for the both the "Strong Diffusion" and "Aperture and Strong Diffusion" cases. This can also be explained due to the increased difficulty of the task. As might be expected, a high correlation (99.8%) was found between the difficulty and the confidence reported on the questionnaire.

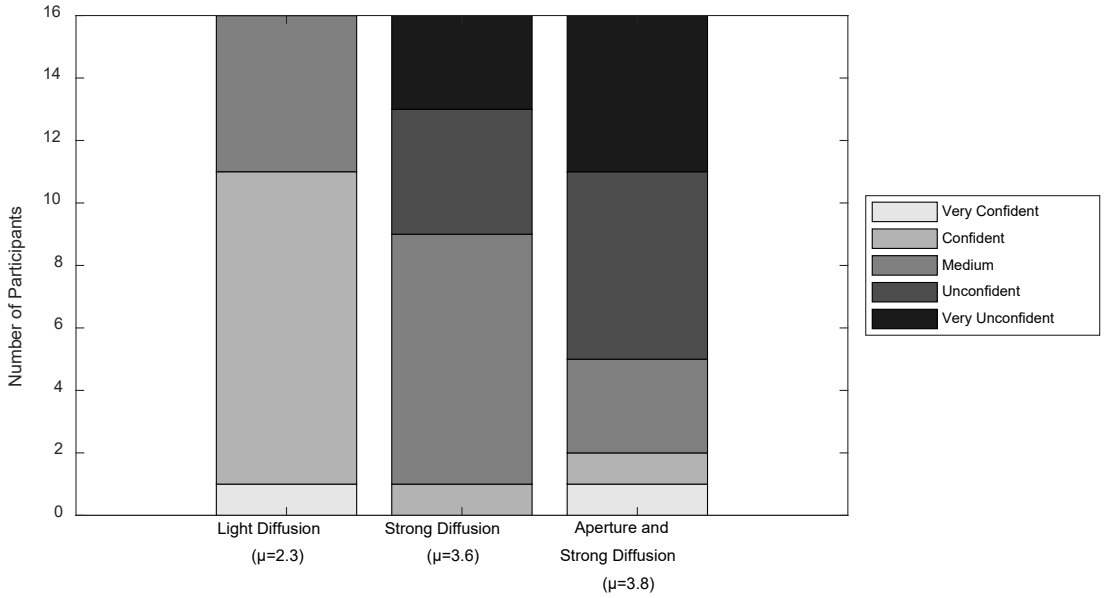

Figure 37: Confidence level of users' results from post-study questionnaire. Participants generally reported high confidence for "Light Diffusion", and low confidence for both "Strong Diffusion" and "Aperture and Strong Diffusion".

The study was designed so that the difficulty increased for each of the three experiments, thereby allowing the participant to learn as they go (each participant completed all three experiments in the same order). Some participants reported learning the study technique more easily than others, especially for the first and second experiments, where the participant was improving their ability to judge depth for this apparatus.

3.7 Chapter Summary

This study examined visual depth acuity under conditions of a simulated intermediate-resolution light-field display. A pilot experiment showed that acuity was greatly reduced by low resolution as well as by the simulated array of lenses of the display; however, viewers were still able to perceive depth. The main study showed that while the number of depth planes is greatly reduced from what a viewer would typically see through a window, dozens of depth planes can still be distinguished even with an intermediate-resolution light-field display corresponding to the "strong diffusion" cases. With a diffusion with standard deviation of 2.2 degrees, the average visual acuity was 534 seconds of arc, which corresponds to 31 depth planes.

These results reveal that not only can the visual system estimate depth even for very blurry images, but that it can do so very well. It may have been critical for humans to function well under low light conditions where only the intermediate-resolution (0.1-0.4 cycles per degree) vision is active (Nes, Koenderink, Nas, & Bouman, 1967), and would also enable navigation in some foggy situations. Furthermore, the results show that visual depth acuity is not entirely compromised even when looking at an object through an array of apertures. This could also have been important for early humans, who may have had to observe prey or predators through trees or bushes, which form natural apertures. An additional observation is that for both levels of diffusion found that the average visual acuity is less than 10% of the amount of angular diffusion, calculated for the strong diffusion case by comparing the visual acuity limit (0.15 degrees) with the level of diffusion (2.2 degrees). This may indicate that processes are being used to resolve depth differences beyond the Nyquist limit, a phenomena known, in other contexts, as Vernier, or hyper-acuity (Westheimer, 1975).

The study explored the role of stereopsis and motion parallax in depth perception. For an intermediate-resolution light-field display, these cues suffice to provide a sensation of depth. However, they can be enhanced by adding additional monocular cues to the images. For example, occlusion can be exploited to give unambiguous information about the depth ordering of objects. Portions of distant objects can also be revealed or hidden as the viewer moves around the display, providing information in addition to the motion parallax and binocular cues. Another strong cue is the difference between the image on the display and the high-resolution border/bezel that surround it, acting much like the frame of a window. This can be enhanced by adding decorative elements to the plane of the display to increase changes of occlusion, stereopsis, and motion parallax. The types of images that are most effective are those with strong low-frequency components in the background, and that take advantage of color and/or intensity gradients to reinforce the occlusion of the background by foreground objects or by the bezel. The next chapter describes how these findings can be used to inform the design of a display device capable of producing an intermediate-resolution light field.

Chapter 4: Optical Design

The experiment in Chapter 3 confirmed that viewers can perceive both binocular and motion parallax cues when viewing a scene through a partially diffusing window. This chapter proposes an optical design that can produce a light field indistinguishable from that observed when looking through such a window, as illustrated in Figure 38.

Figure 38: Illustration of a view of a scene through a window from two different viewing positions. In the top pair, the window is transparent; in the bottom pair the window is partially diffusing (simulated using a Gaussian low-pass filter).

The chapter begins by describing the design concept at a high level. This is followed by a description of the software simulation used to analyze how the light field propagates through the optical system, using the light-field described in Chapter 2.5. The chapter then analyzes the performance of the optical design, by quantifying the deviation of the light field produced by the display from that of a partially diffusing window.

4.1 Design Concept

The primary goal of a light-field display, as described in Chapter 2, is to create the desired 4D light field, defined as the intensity of light passing through a plane as a function of two spatial dimensions (I_s and I_t) plus two angular dimensions (I_u, and I_v) (Levoy & Hanrahan, 1996). Existing displays (Dimenco, 2019) attempt to accomplish this by using an array of lenses positioned in front of an LCD panel. Each lens in such an array is a single projection unit, as illustrated in Figure 39. For each lens, light from different LCD pixels (P_1 and P_2) is refracted towards different directions. Pixels P_1 near the optical axis of each lens are refracted substantially parallel to the axis, as shown in green, while a neighboring set of pixels P_2 are refracted at a different angle, as shown in orange. A light-field display can be constructed by arranging each set of neighboring pixels – corresponding to the set of views at that location – into a regular array. This technique works for narrow viewing angles. As the viewing angle increases, simple lens systems introduce increasing levels of distortion to the image. Additionally, wide angle light from pixels behind an adjacent lens can “leak through” an incorrect lens. Some displays take advantage of this to produce a repeating set of narrow views over a wider range of viewing angles. However, in practice this allows only a small range of viewer movement and is not suitable for the application of a wide-angle light-field display.

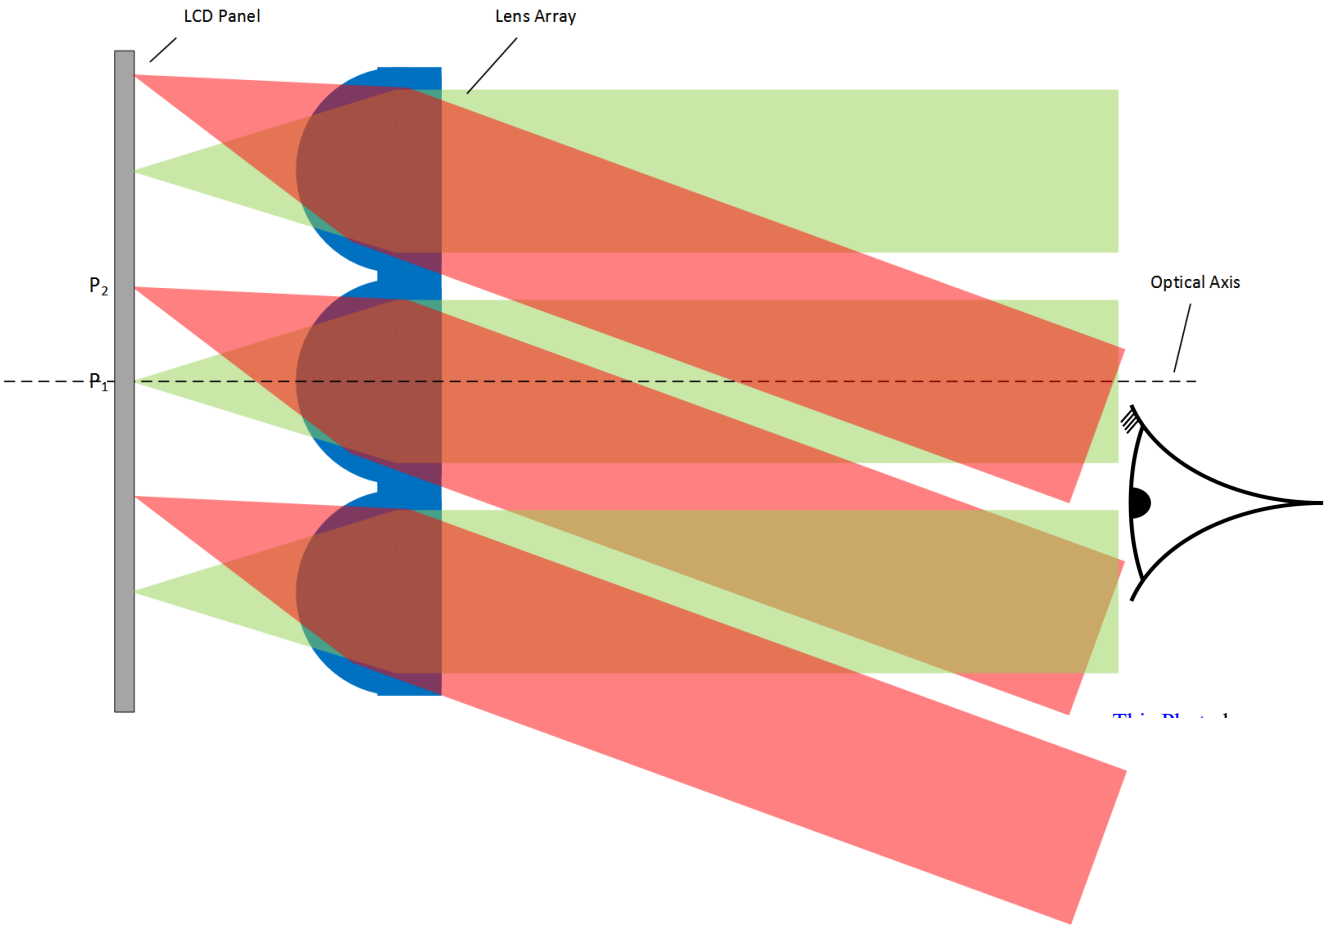

Figure 39: Light-field display concept. Light emitting from different sets of pixels P1 and P2 are projected at different angles from the display.

4.1.1 Initial Design

For the initial prototype, an array of 2.5cm aspherical Fresnel lenses were tiled as a rectangular array in front of an LCD panel to form an array of 23 x 13 projection units, as illustrated in Figure 40.

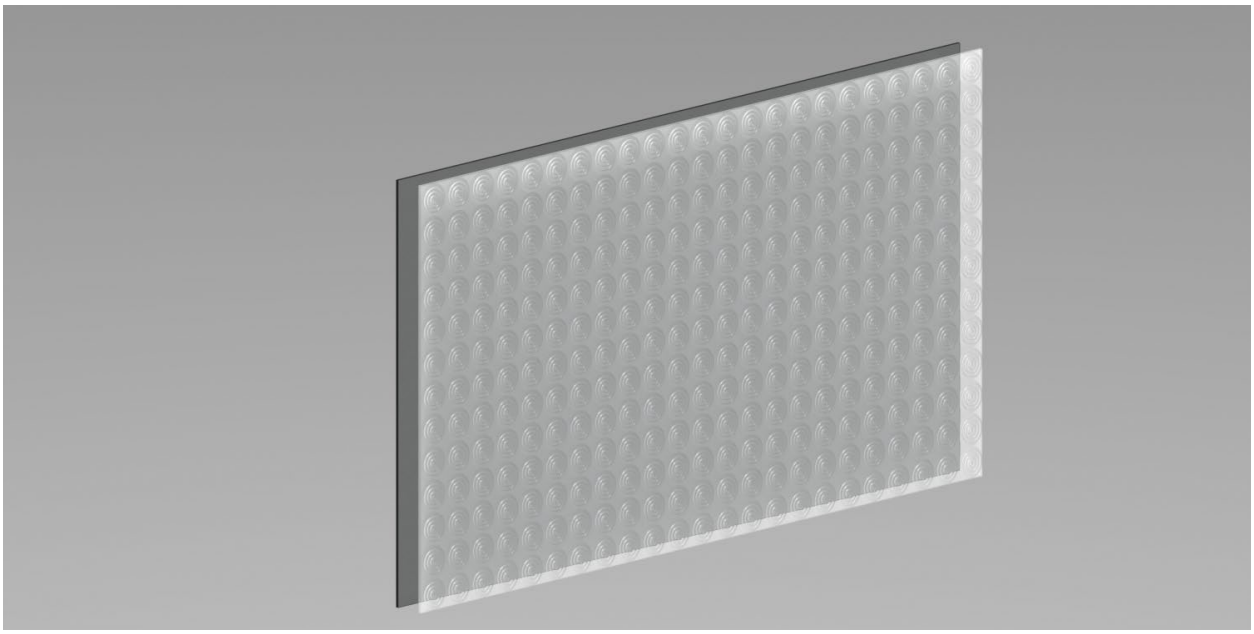

Figure 40: Illustration of first design iteration using an array of Fresnel lenses in front of an LCD

As expected, this simple design was able to produce a reasonable image within a very narrow field of view. However, the quality rapidly deteriorated at greater viewing angles. This confirmed that, as expected, a different lens is needed, one that offers improved performance over a wide field of view, while still being practical to manufacture using cost-effective molding techniques. An additional challenge was found in the discontinuities between adjacent lenses, which need to be masked in order to simulate the continuous light field produced by a diffusing window.

4.1.2 Improved Lens System

The role of the lens array in a light-field display is to direct the light from each pixel in the LCD display to a unique direction, proportional to the pixels' distance from the center of each lens (illustrated as the green and orange shaded regions in Figure 39). The lens is positioned at the focal distance from the imaging plane, so that light emitted from individual pixels is converted to a collimated beam of light travelling at an angle determined by the ratio of the distance of the pixel from the optical axis to the focal length of the lens. In practice, this works well for paraxial light (light that is near to and propagating nearly parallel to the optical axis). For pixels further from the optical axis, the lens aberrations discussed in Chapter 2.4 (dominated by curvature of field) cause progressively worse distortions.

The performance of configurations based on the early photographic lens designs described in Chapter 2.4 were analyzed and compared to a Fresnel lens array. More sophisticated designs that used increasing numbers of optical elements were not considered economical, since they would require many hundreds or thousands of lenses. The Fresnel lens was compared to a plano-convex lens, a meniscus lens, and a globe lens as described in Chapter 2.4.2.

To evaluate each lens design, a ray tracing simulation programmed in Matlab (Mathworks Matlab) calculated the standard deviation from the mean of the angles of rays projected both on the optical axis, (0 degrees from the optical axis), as well as at the widest viewing angle desired for the application (40 degrees). A performance metric was defined as the greater of these two standard deviations. Ideally, all light would be collimated, so the standard deviation would be zero for both. Real lenses, however, always have some distortion, so the standard deviation is

greater than zero. With this metric, it was found that the optimal solution is for both paraxial and non-paraxial rays to have approximately the same standard deviation, which has the added advantage of yielding a consistent degree of angular consistency over the desired range of viewing angles.

For each of the ray tracing simulations, the optical power was held constant so that a point located at half the spacing between adjacent lenses (1.27cm from the optical axis) would be collimated at a mean off-axis angle of 40 degrees. Additionally, all lenses were modelled with the same aperture, having a diameter of 4.2mm to achieve an f-number of approximately 6, typical of many photographic lenses. Only light rays that passed through the aperture were considered. The curvature of the surface, the position of the lens and aperture, and the thickness of each lens design were optimized using Matlab's implementation of the Levenberg–Marquardt non-linear solver to minimize the performance metric. The optimized lens shapes are illustrated in Figure 41, and the resulting performance metrics tabulated in Table 4.

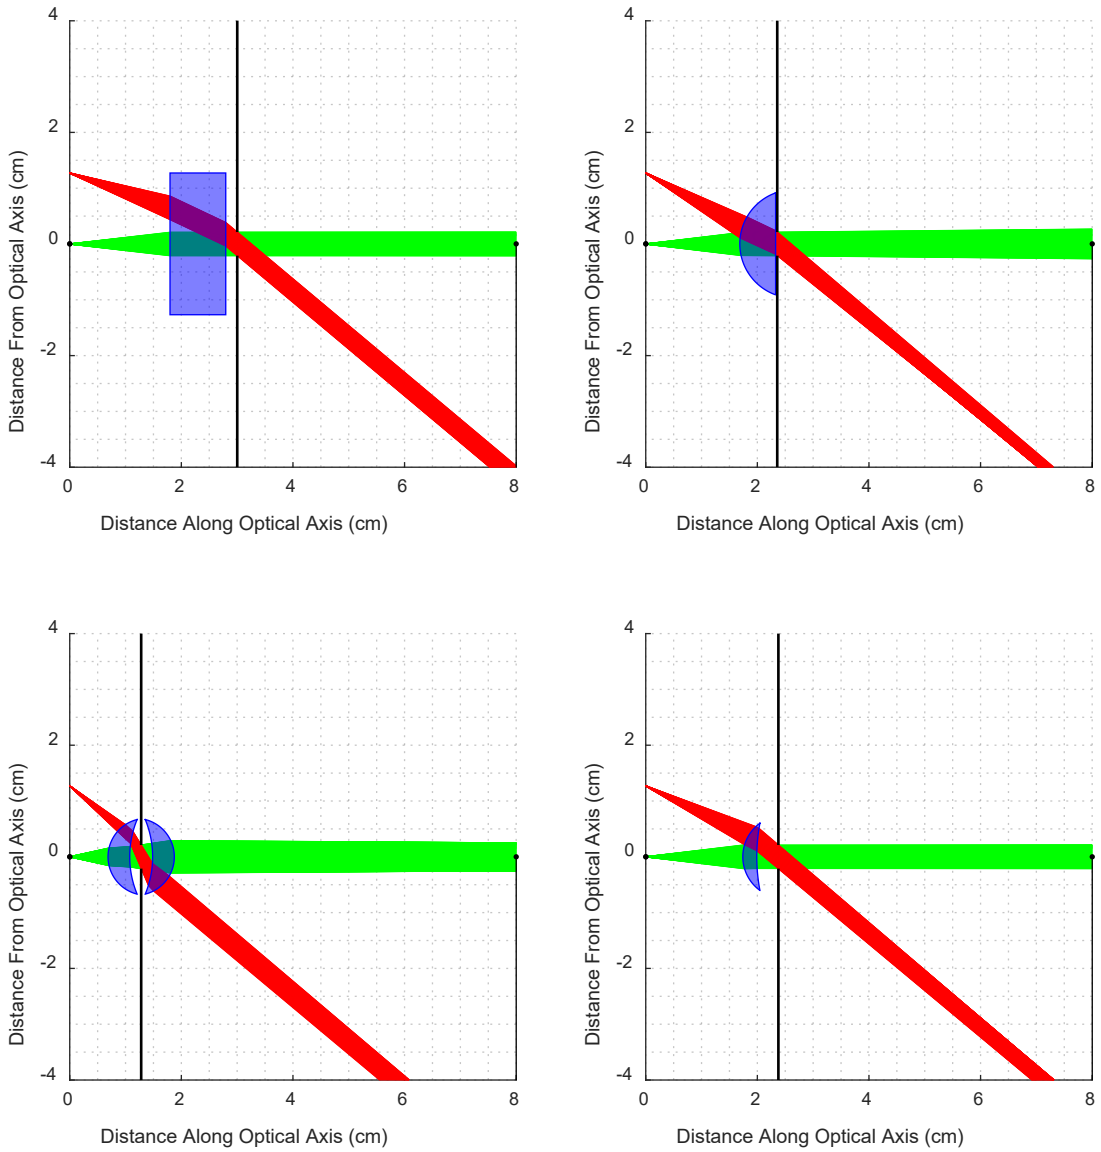

Figure 41: Optimized lens profiles. Clockwise from top-left: Fresnel, Plano-Convex, Meniscus, Globe

Table 4: Summary of lens element performance characteristics

Lens	Maximum of Angular Standard Deviation
Fresnel	0.24 degrees
Plano-Convex	0.43 degrees
Meniscus	0.19 degrees
Globe	0.15 degrees
Ideal Lens	0 degrees

As shown in Table 4, the Fresnel lens from the first design iteration has an optimized angular standard deviation of 0.24 degrees. This significantly outperforms the optimized Plano-Convex lens, due in part to the aspherical surface profile of the Fresnel Lens, which allows more degrees of freedom to be optimized. The meniscus lens is a slight improvement over the Fresnel lens, due to having two curved surfaces which allow for finer-grained optimization. As expected, the globe lens outperforms all other designs for this application, with an angular standard deviation of 0.15 degrees. For imaging applications requiring high spatial resolution, the performance of the lens is critical in producing a sharp image. For an intermediate-resolution light-field display, it is less important to produce high resolutions than to have excellent consistency across a wide range of viewing angles. Replacing the Fresnel lens array from the first prototype with an array of the optimized globe lenses, as illustrated in Figure 42, results in a more consistent light field over a wide field of view.

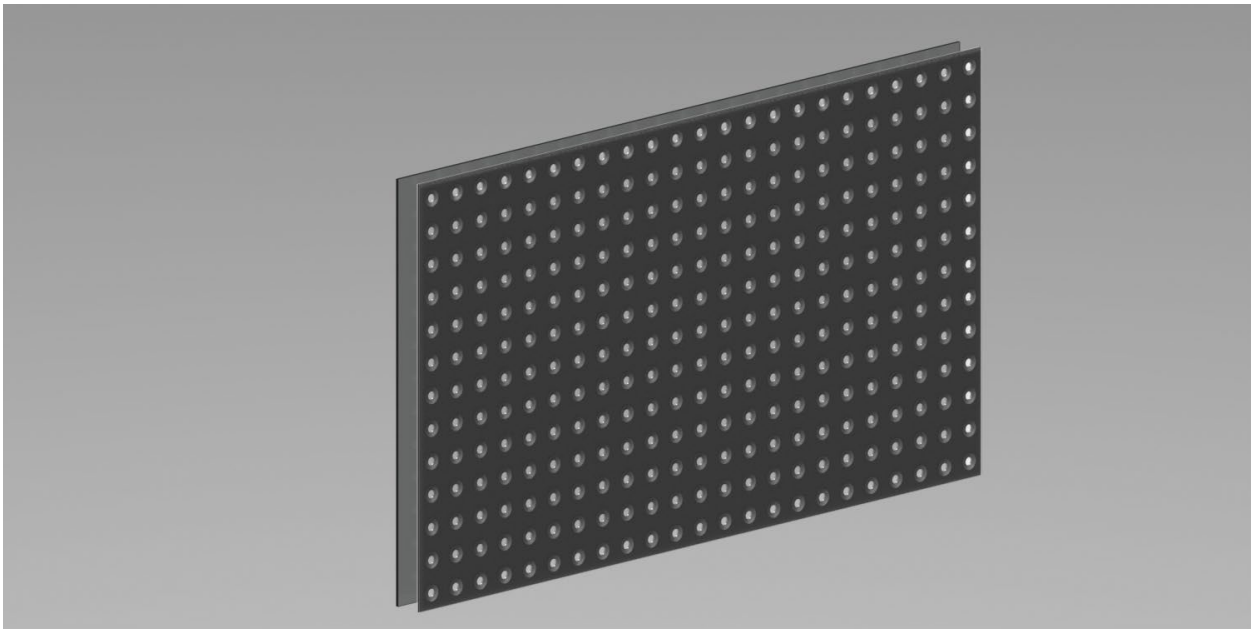

Figure 42: Illustration of the second design iteration using an array of globe lenses

While the wide-angle performance of the globe lens system is more consistent than the other lens designs that were considered, it has two drawbacks. The first, illustrated in Figure 43, is that there is an even greater need to mask the discontinuities between the lenses; this prompts the development of a diffusion system to hide these gaps. The second drawback is that the central aperture blocks a significant amount of light, greatly decreasing the optical efficiency of the system. This results in the need for an improved light source of the LCD panel to obtain the desired optical efficiency.

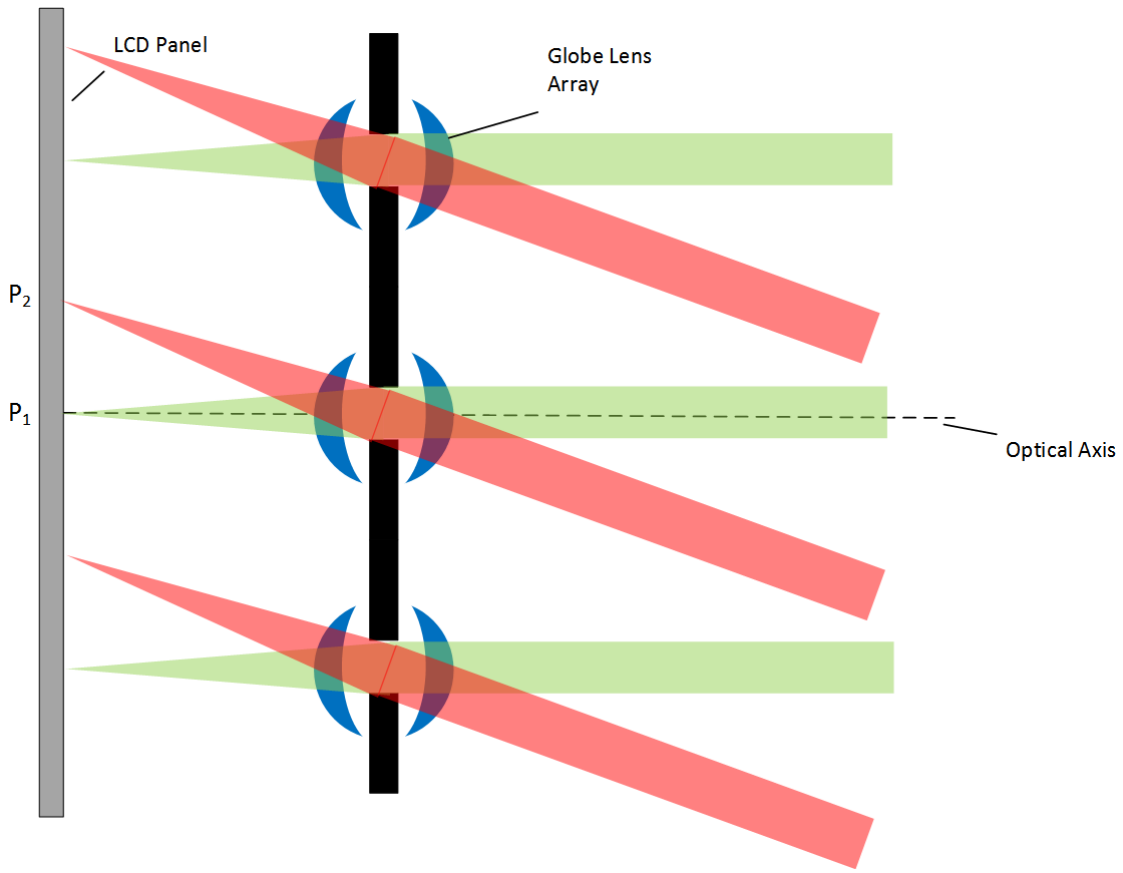

Figure 43: Illustration of an array of projection units using globe lenses

4.1.3 Dual Diffuser System

As mentioned previously, the relatively small apertures of the globe lens array results in a highly non-uniform light field. At close viewing distances, the light field appears as an array of beams of light with dark areas between them. A diffusion system is needed to fill these dark areas. Adding a highly diffusing layer could mask the gaps between the lenses, but would also increase the angular diffusion to an extent where the display could no longer provide a sense of depth. A level of diffusion is required that produces sufficient lateral spreading of the light to mask the gaps, but has a smaller amount of overall angular diffusion.

A solution was found using two diffusers, separated by a gap, as illustrated in Figure 44. The first diffuser applies a mild scattering of light emitted from the globe lens. Then, critically, the scattered light diverges laterally in the diffuser gap between the two diffusers. It is this lateral spreading of the light that masks the dark areas between the lenses without the use of heavy diffusion. The second diffuser then applies a second mild scattering of the light, to hide any residual non-uniformity. The light field produced by this combination of two diffusers separated by an air gap can therefore be made spatially uniform, while introducing only a modest amount of angular diffusion.

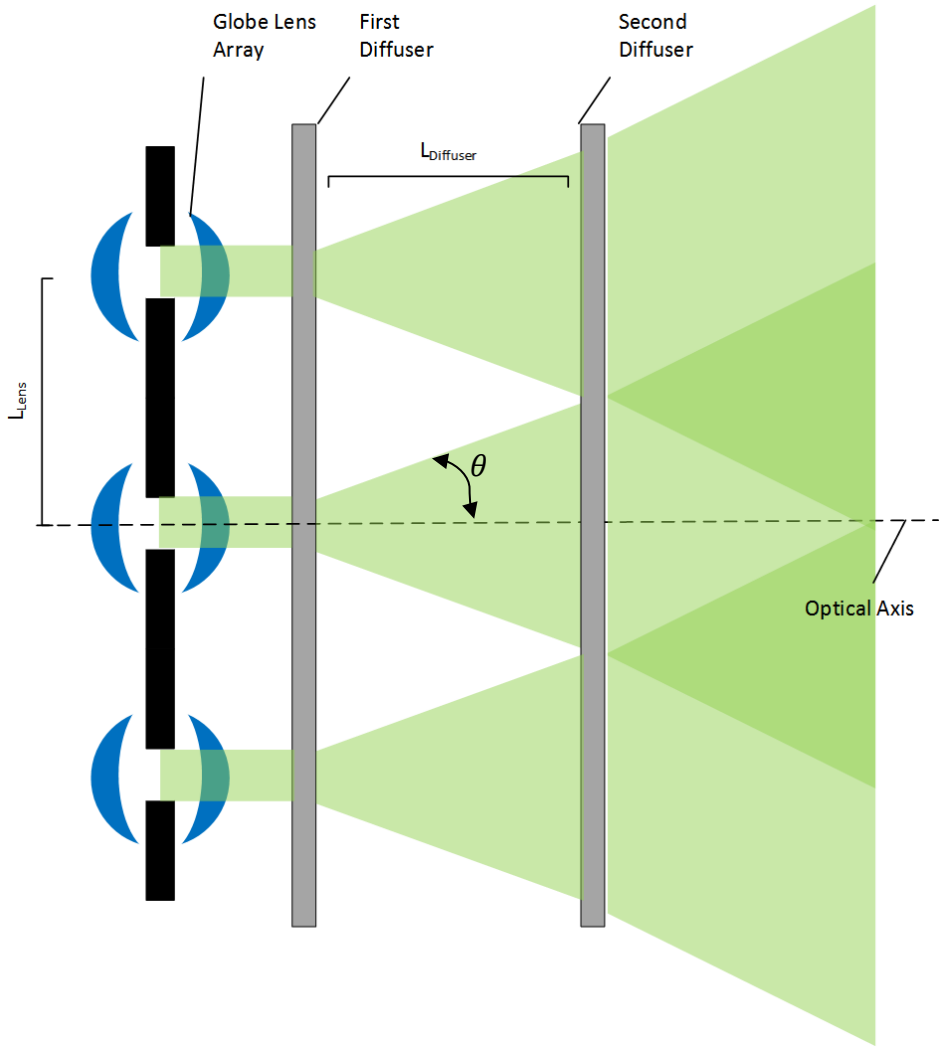

Figure 44: Diffusion system with two diffusers

The net effect of this dual-diffuser system is a combination of directional and spatial diffusion. The total amount of angular diffusion introduced by this design is essentially the sum of the scattering done by both first and second diffusers, while the amount of spatial diffusion is proportional to the diffuser gap. A large gap results in only very little directional scattering, while a narrow gap requires greater scattering to achieve the desired lateral spreading.

This design can be optimized by specifying the diffusion distribution of both first and second diffusers. The desired diffusion profile, or point spread function, of the first diffuser should be designed to ensure that the overlap from adjacent apertures of the globe lens array sums to a uniform spatial intensity, with a minimum of diffusion, and avoiding discontinuities that would require extremely precise alignment. A suitable diffusion profile with these characteristics is described in (United States Patent No. US 9534746, 2017). The equation for the angular diffusion as a function of angle θ is described by Equation (4):

$$L(\theta) = \begin{cases} \dfrac{{3s^2}/{4} - \theta^2}{s^3} & for\ |\theta| \leq {s}/{2} \\ \dfrac{\left(|\theta| - {3s}/{2}\right)^2}{2s^3} & for\ {s}/{2} < |\theta| \leq {3s}/{2} \\ 0 & for\ |\theta| > {3s}/{2} \end{cases} \qquad \textbf{(4)}$$

$$where\ s = \tan^{-1}\left(\frac{L_{Lens}}{L_{Diffuser}}\right)$$

L_{Lens} is the distance separating the centers of adjacent lenses in the lens array, and $L_{diffuser}$ is the size of the gap separating the two diffusers. The angle formed by $L_{diffuser}$ and L_{Lens} is

defined as s, the angle where the projected light intensity from two adjacent lenses is equal. The diffusion profile described in Equation (4) is illustrated in Figure 45.

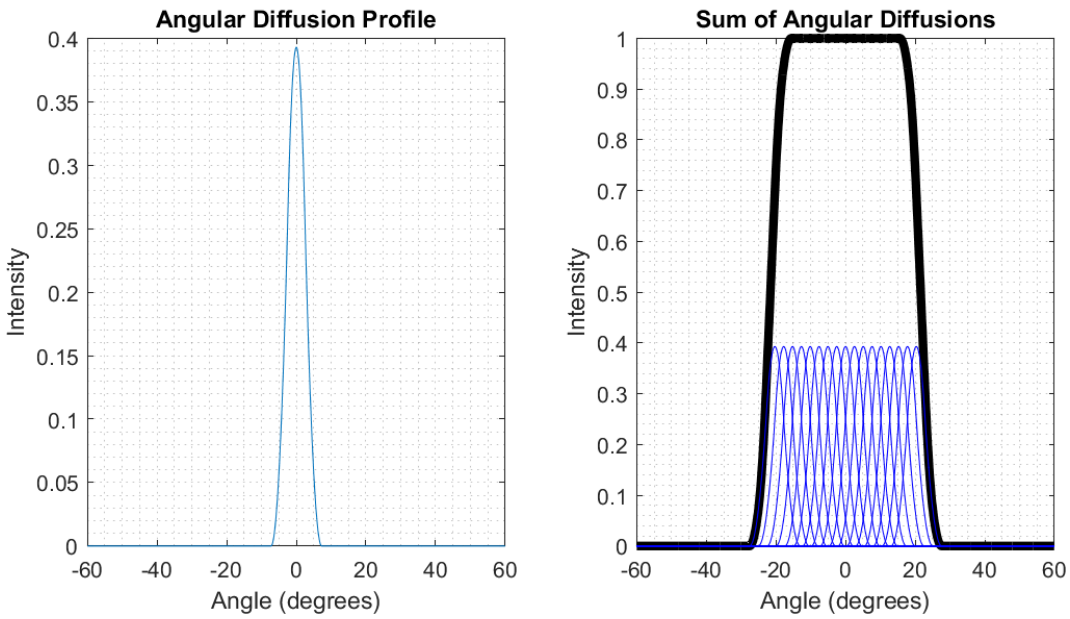

Figure 45: Optimal angular diffusion profile of first diffuser from a single light source (left) and the sum of multiple adjacent sources (right).

When the scattered light from the first diffuser has propagated through the gap separating the two diffusers, the intensity, when integrated over all angles, is approximately spatially uniform. However, the intensity distribution still has directional non-uniformity. The second diffuser must mask this by applying sufficient directional diffusion. The desired PSF can be described in a similar way as for the first diffuser, again using Equation (4). After being scattered by the second diffuser, the resulting light field is uniform spatially and smoothly varying in the directional component, achieving the objective of appearing indistinguishable from the light field perceived through a partially diffusing window.

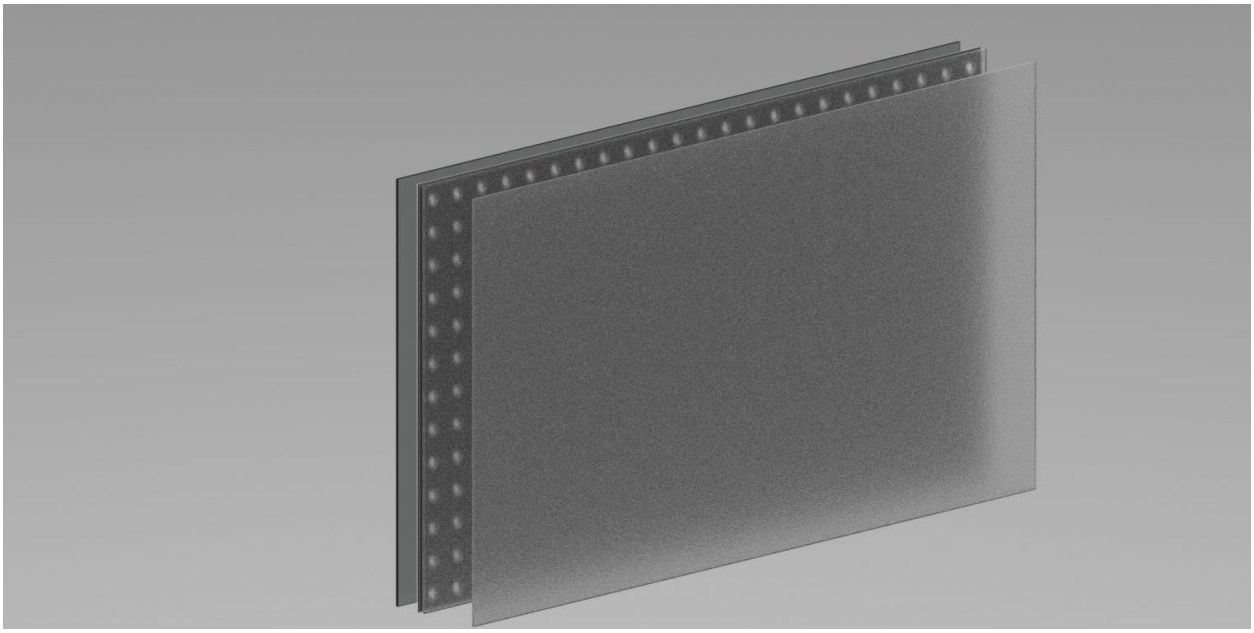

Figure 46: Illustration of third design iteration, using globe lenses and dual-diffusion system. The first diffuser is adjacent to the lens array, and the second diffuser is separated by a gap.

The resulting design was able to produce the target light field, over a wide field of view, with good angular consistency and spatial uniformity, within a practical cost. However, higher optical efficiency from the illumination source is still needed to overcome the limitation imposed by the small apertures of the globe lens.

4.1.4 Improved Efficiency Illumination Source

For the next design iterations involving a globe lens array, begin by noting that the light emitted by the LCD pixels is emitted over a fairly wide angle, most of which is blocked by the apertures in the array of globe lenses, as illustrated in Figure 47.

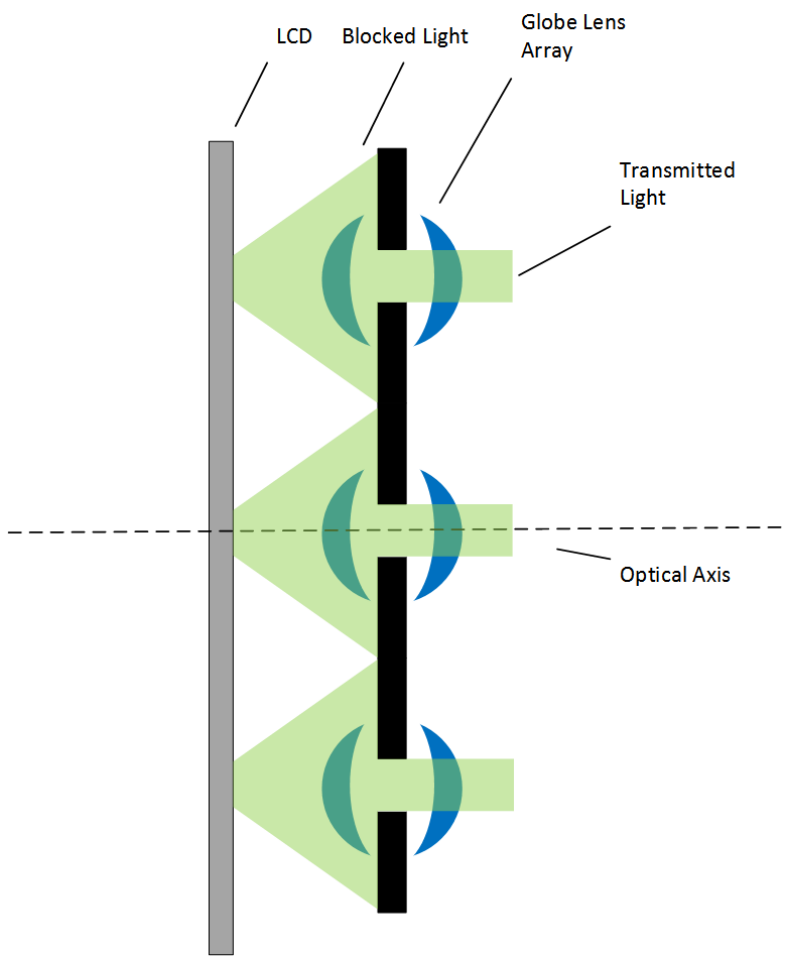

Figure 47: Low optical efficiency due to globe lens array

To overcome this limitation, an improved backlight was designed to focus light through the apertures of the globe lenses. The key concept is to employ an array of LED light sources, one

per lens, each of which is considerably smaller than the lens-to-lens separation. Light from each LED is focused by a corresponding field lens, through the LCD and then through the aperture of the corresponding globe lens, as illustrated in Figure 48. An array of double-sided Fresnel lenses is a good choice for the field lens, due to the required low thickness and required low f-number. The Fresnel lenses were designed to have a diameter equivalent to the spacing between lenses, and are positioned half-way between the LED light sources and the globe lens.

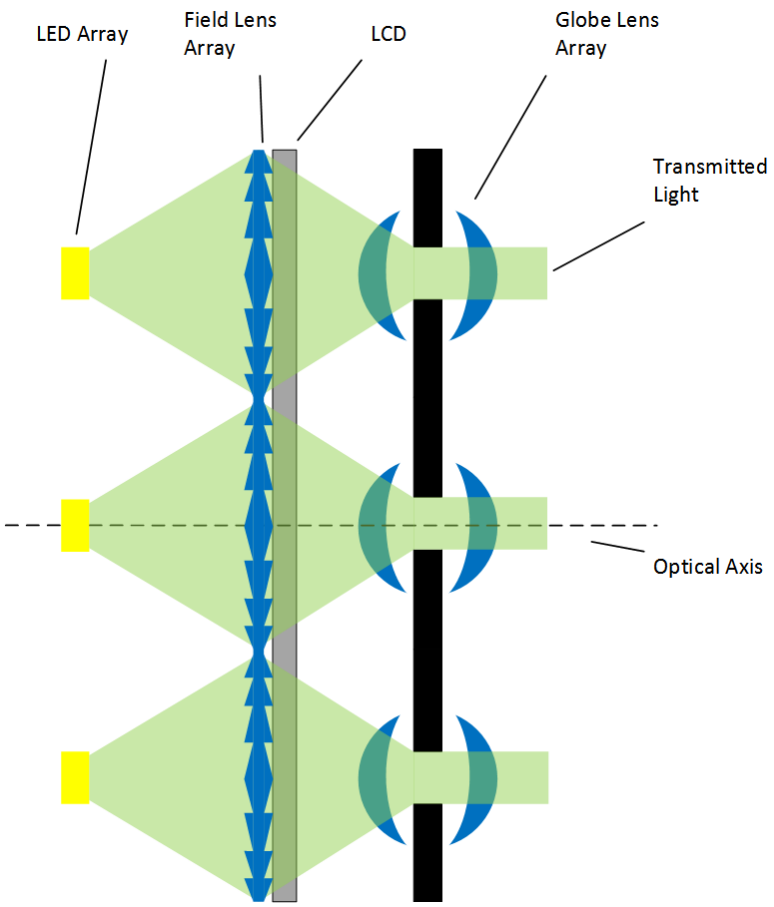

Figure 48: Improved optical efficiency using an array of LEDs and Fresnel field lenses, with the facet sizes of the field lenses exaggerated for illustrative purposes

The resulting design is able to produce the desired intermediate-resolution light field, with a wide field of view, good angular consistency, good spatial uniformity, good optical efficiency, and a practical manufacturing cost. Figure 49 illustrates the complete design, including the improved illumination source, LCD, array of globe lenses, and dual-diffuser system. This 3D model is an illustration of a single module for the proposed intermediate-resolution light-field display.

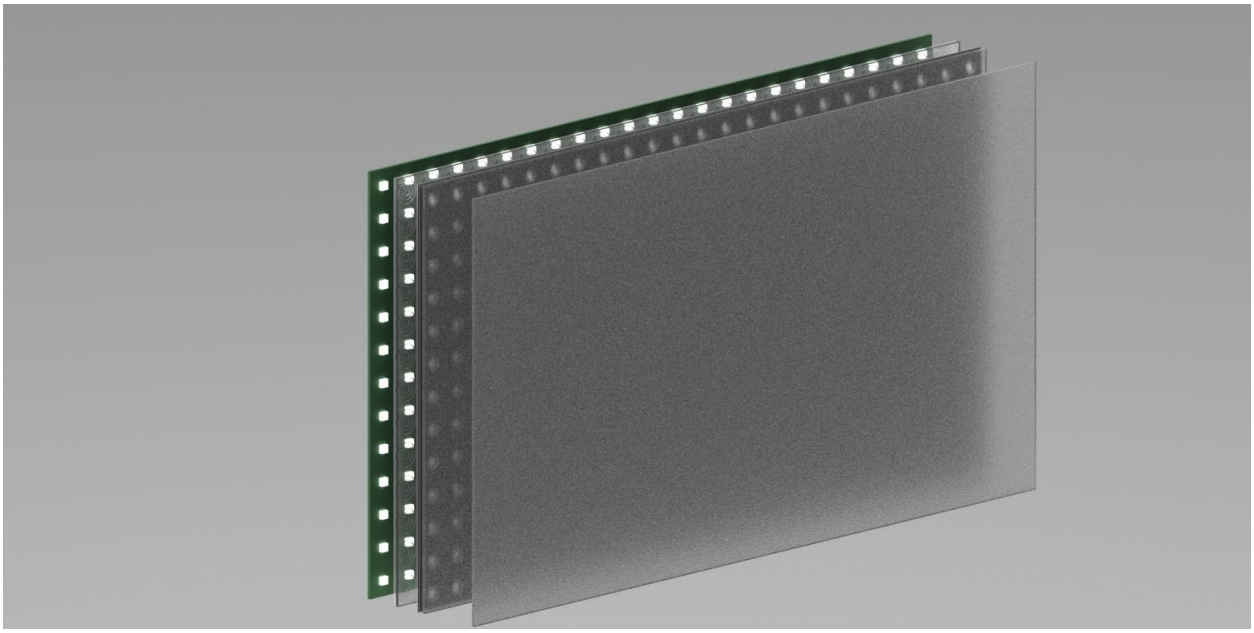

Figure 49: Illustration of the design concept using an array of field lenses to focus light through the apertures of the globe lens array, followed by the dual-diffusion system

4.2 Simulation of the Optical Design

An overview of the optical design, showing only two adjacent projection units within the array, is illustrated in Figure 50. The design is divided into two stages. An Imaging Stage, shaded in red, is responsible for producing the light field. The Illumination Stage, shaded in yellow, is responsible for providing an efficient source of illumination for the Imaging Stage.

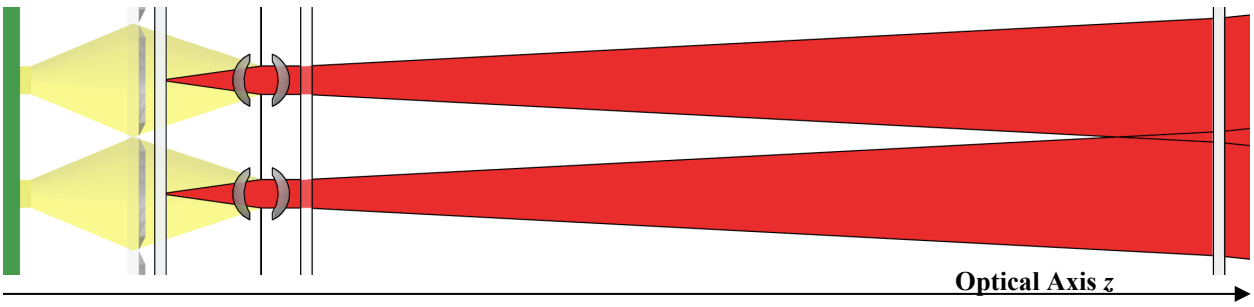

Figure 50: Illustration of optical path, with the Imaging Stage shaded in red and the Illumination Stage shaded in yellow.

The optical design was simulated by ray tracing using Matlab (Mathworks Matlab). Rays were represented as having a position along the optical axis z, one spatial dimension y representing the distance from the optical axis, and one angular dimension θ representing the angle formed with the optical axis, together resulting in a three-dimensional location for each ray. The spatial dimensions y and angular dimensions θ are representative of the 4D light field introduced in Chapter 2. The software simulation simulated 100,000 rays for each of the projection units, each with unique starting coordinates z_0, y_0, and θ_0. The analysis was also repeated with both half and twice the number of rays and found to be consistent within the precision reported in this thesis. The plane of the LED array was defined as $z=0$. As each ray intersects an optical surface along the optical axis, it is modelled as being either refracted, blocked by an aperture, or scattered by a

diffuser. Refraction and scattering only resulted in a change of the angular component θ, and propagation along the optical axis z only resulted in a change of the spatial component y.

In order to analyze the propagation of the rays through the optical system, the spatial and angular components of each ray were plotted at each optical surface along the optical axis. Each ray was represented as a semi-transparent dot at the position corresponding to its spatial and angular coordinates. Darker regions indicate a higher density of rays, which corresponds to a higher average intensity. As illustrated in Figure 51, a uniform spatial distribution would appear as a horizontal line with constant density across a wide range of spatial positions. Likewise, a uniform angular distribution would appear as a vertical line of constant density.

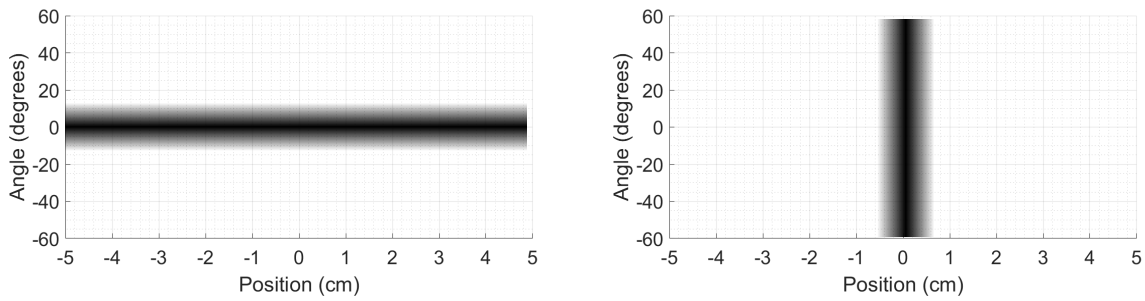

Figure 51: Illustration of a 2D light-field with a uniform spatial distribution (left), and uniform angular distribution (right)

Figure 52 compares the light-field distribution profile of a distant light source such as the sun as viewed through both a conventional clear window as well as a partially diffusing window. The light field of the conventional clear window (left side of Figure 52) has a uniform spatial distribution, indicating that the same pattern of intensities is perceived from a wide range of viewing positions, causing it to be perceived as a distant object. The standard deviation of the

angular distribution is essentially zero. By contrast, the light field of the partially diffusing window shows the same uniform spatial distribution, but with wider angular distribution due to the diffusion. The angular diffusion has a peak intensity at $\theta=0$ and a gradual reduction at greater angles. By comparison, a conventional display is limited to representing a distant image of a sun with a wide angular distribution and narrow spatial distribution – significantly different from what is encountered in a real scene.

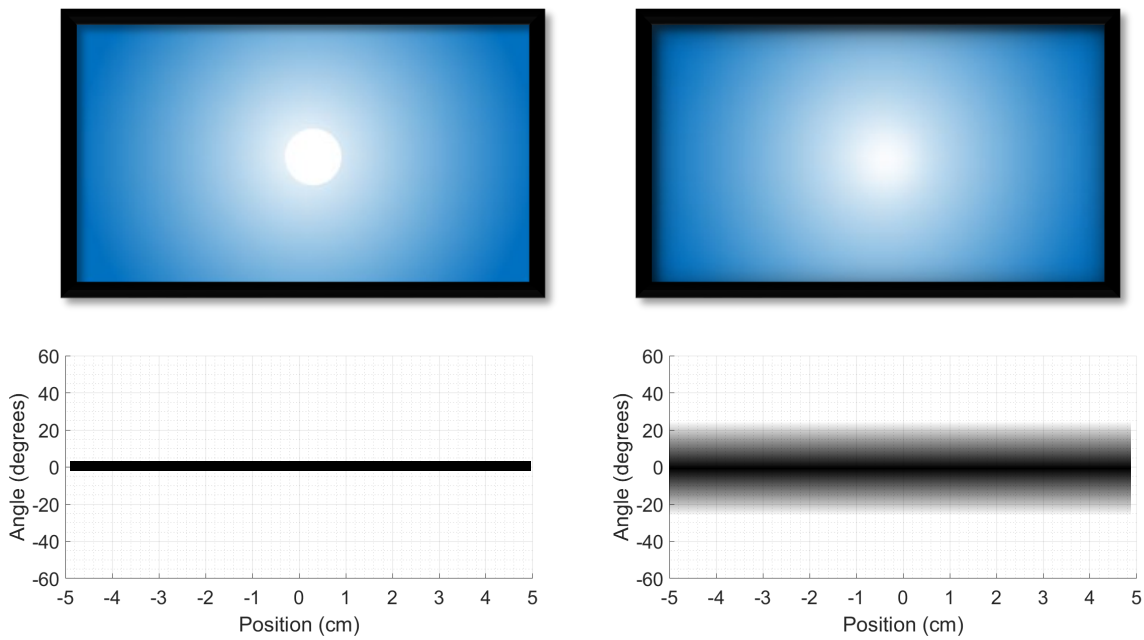

Figure 52: An illustration of the image (top) and light-field distribution (bottom) of a distant object seen through a conventional clear window (left) and a partially diffusing window (right)

4.2.1 Imaging Stage

The imaging stage of the optical design, shown in Figure 53, is primarily responsible for producing the light field. As described in Section 4.1, this stage comprises an LCD, an array of globe lenses, and a diffusion system. The light from LCD pixels close to the optical axis of each projection unit are illustrated in red, while the light from pixels far from the optical axis are illustrated in green.

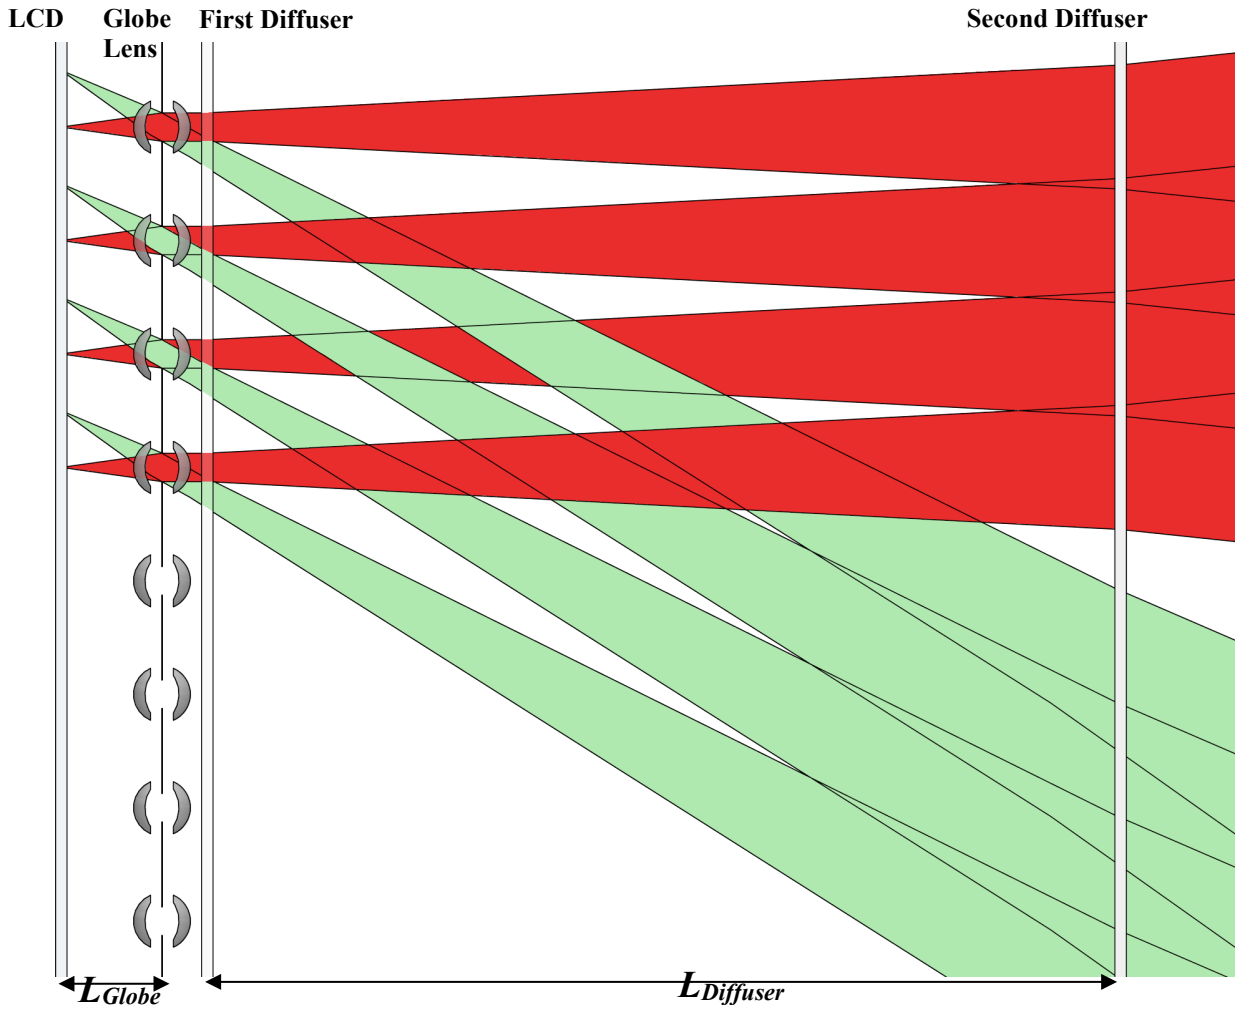

Figure 53: Illustration of optical path for the imaging stage

Starting at the light emitted from the on-axis pixels of the LCD panel, Figure 54 illustrates the 2D light field of the light transmitted through four LCD pixels, each on the optical axis of four adjacent projection units. The light distribution from each LED is represented as a Gaussian spatial distribution with sigma = 3.2 x 10^{-4} m (which is the size of a pixel in a typical 1.4 m diagonal consumer television with 3840 horizontal pixels). The angular distribution was likewise modeled as a Gaussian distribution (sigma = 11 degrees), directed towards the aperture of the globe lens, as an expected output from the field lens array. As previously mentioned, darker shading corresponds to a greater intensity of light. As can be seen, the four individual pixels sources are distinctly separated in the spatial dimension by a distance of y_0=2.5cm.

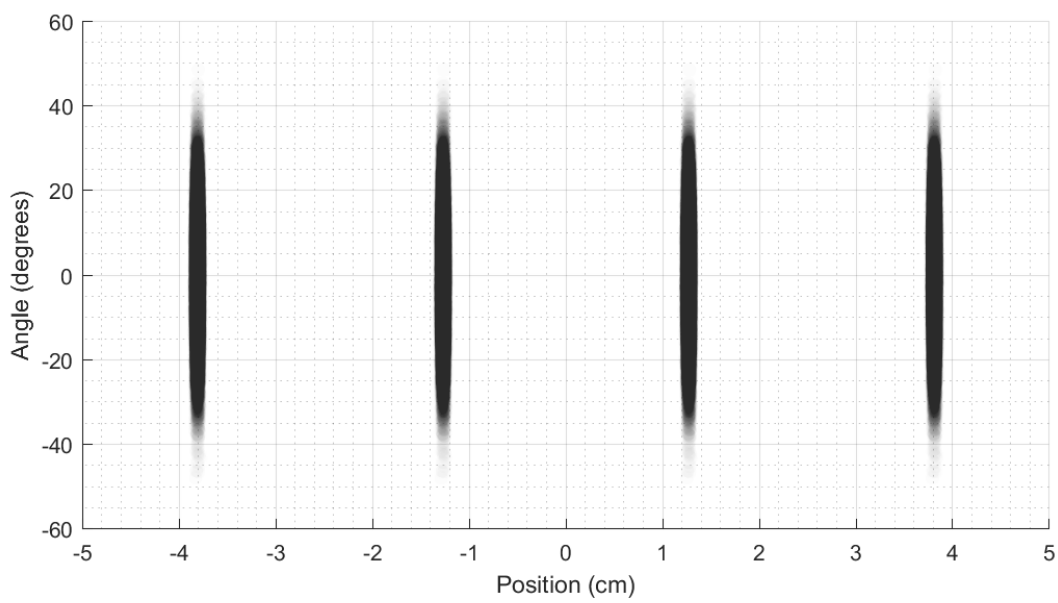

Figure 54: 2D light field from four pixels corresponding to the optical axis of four lenses

Figure 55 shows how the distribution has changed as the light travels a distance $z=L_{globe}$ along the optical axis. As these rays of light propagate away from the LCD array along the optical axis z, they shift differently along the spatial dimension y according to the arc-tangent of the initial

starting angle θ_0, as described by Equation (5). The angular distribution is unchanged, but the light has spread spatially as it propagates through air to the array of globe lenses.

$$(y - y_0) = \Delta z \operatorname{atan}(\theta_0) \quad \textbf{(5)}$$

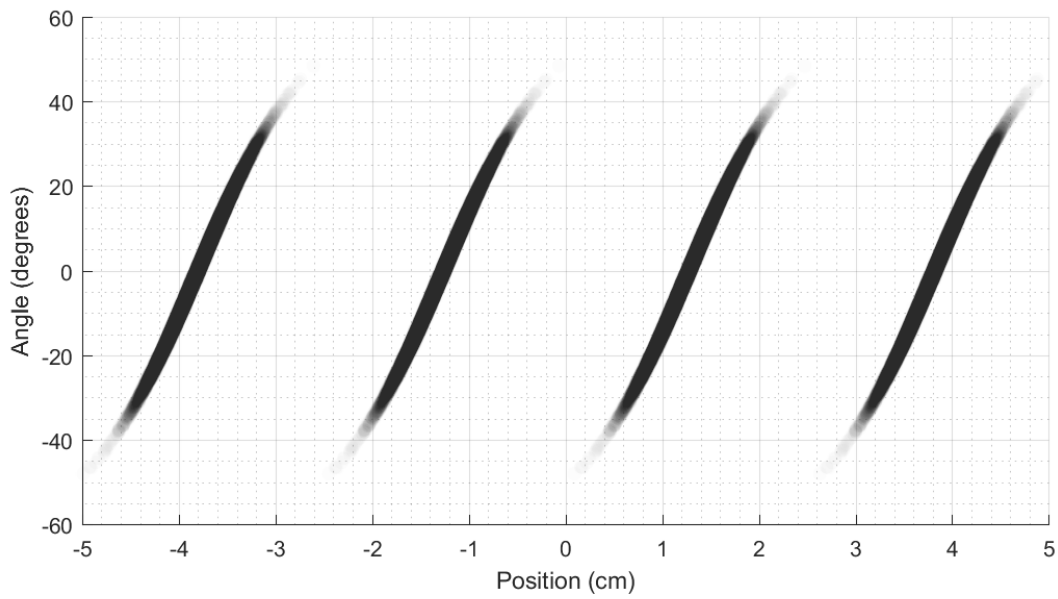

Figure 55: 2D light field after propagating to the globe lens

The role of the globe lens array is to collimate the diverging light from the LCD array into a narrow beam of projected light. This is achieved by the two meniscus lenses forming the globe lens as well as the spacing between them. Light incident on the first meniscus lens of the globe lens is partially collimated and focused slightly through the aperture at its center. Figure 56 shows the light field at the aperture, marked by a shaded grey area. The spatial distribution is limited by the aperture of 6 mm. Rays not within the spatial bounds of the aperture are absorbed.

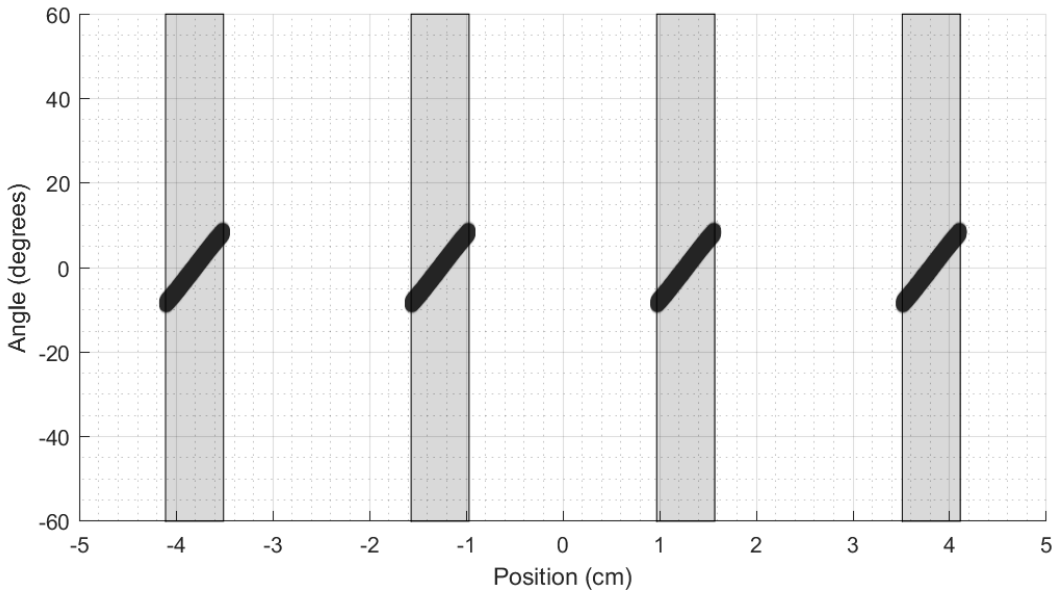

Figure 56: 2D light field at aperture of globe lens. The spatial distribution is limited by the aperture, illustrated as a grey shaded area

The second meniscus element of the globe lens completes the collimation of the light from the LCD panel. The optimized lens design produces an array of collimated beams of light with angular standard deviation of 0.35 degrees. Figure 57 shows the result.

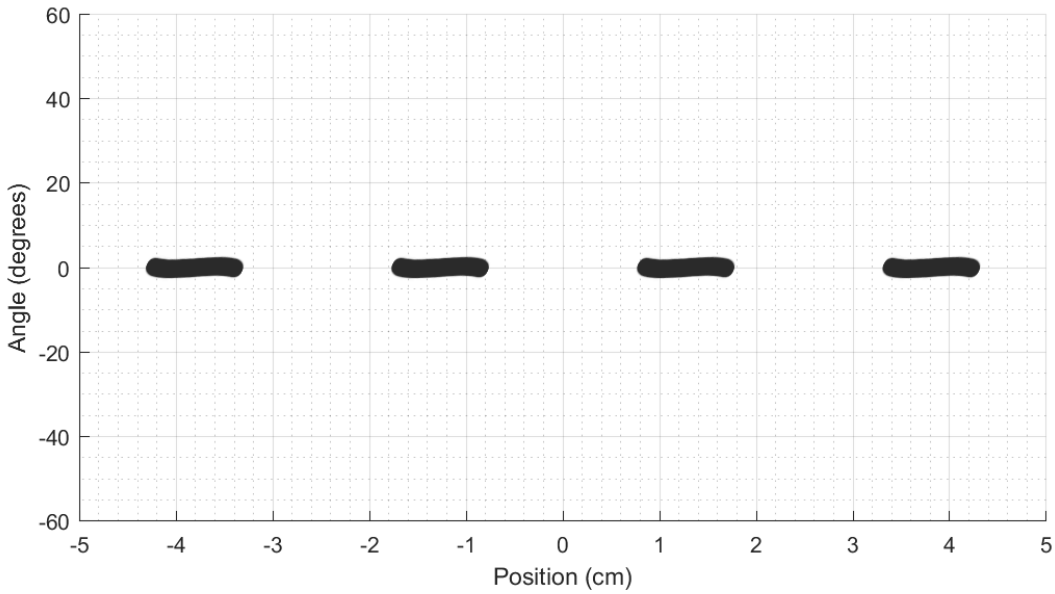

Figure 57: 2D light field after globe lens. The angular distribution is reduced to sigma = 0.35 degrees

Comparing Figure 54 and Figure 57 illustrates how the initial angular distribution of light emitted by a single pixel of the imaging surface has been collimated to a single angular direction by the globe lens. The light field projected by the display at this point is an array of discrete beams of light, shining out from the center of each of the individual lenses of the globe lens array, with one beam for each pixel projected with a different direction and amplitude, depending on the position and intensity of the pixel in relation to the optical axis of each lens. From a far enough viewing distance, the limited resolution of the human eye would cause the discrete beams of light to blend together, blurring the beams and masking the shadow between each lens, so that the resulting light field would appear uniform. However, from closer viewing distances typical of indoor environments, the shadows between each lens will become increasingly visible, resulting in a non-uniform appearance of the light field. The role of the diffusion system is to

mask these gaps between the lenses and produce a uniform light field, while introducing minimal angular diffusion in order to produce the desired number of unique depth planes.

As described in Section 4.1.1, this can be done using two diffusers separated by a gap $L_{Diffuser}$ (Figure 53). The first diffuser is immediately adjacent to the globe lens array, and slightly diffuses the collimated light emitted by the lens array. In order to model the diffusion in the ray tracing simulation, the angle of each ray of light that passes through the diffusers is randomly changed according to a probability distribution for the optimized diffuser as described in Equation (4). After being scattered by the first diffuser, the 2D light field has a greater angular spread but the same spatial extent, as shown in Figure 58. The standard deviation of the angular distribution has been increased from the highly collimated beams of light produced by the globe lens array, with standard deviation equal to 0.35 degrees, to a modestly diverging beam of light with an angular distribution of 2.5 degrees.

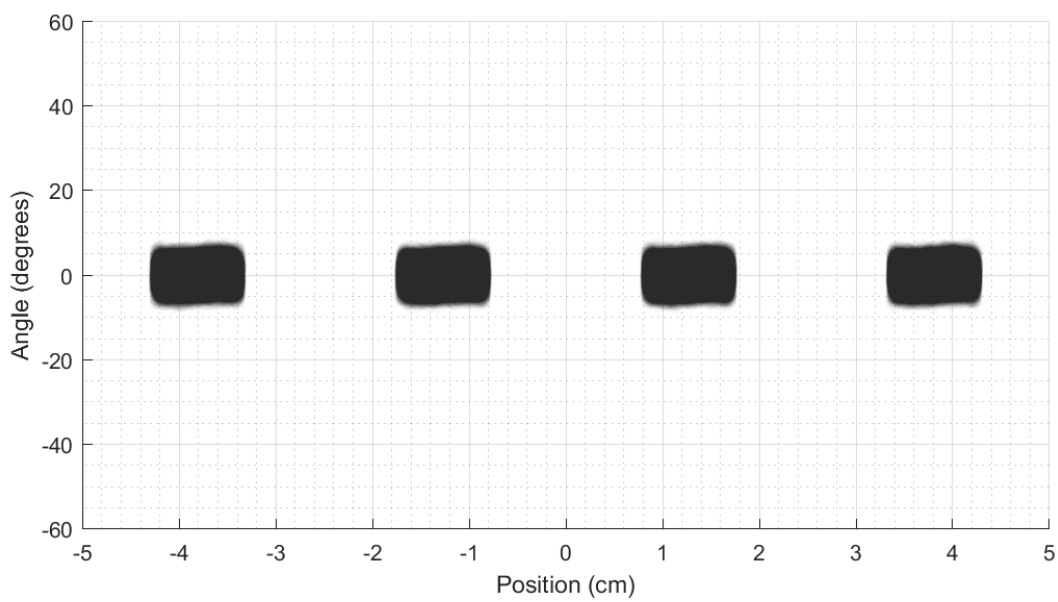

Figure 58: 2D light field after the first diffuser, showing increased angular distribution

The spatial diffusion is now achieved as the light propagates between the first and second diffusers, since the spatial distribution increases as a function of the distance along the optical axis and the tangent of the diffused angle. The net result is a diagonal pattern in the 2D light field, as illustrated in Figure 59. If the intensity were now integrated over all viewing angles, a uniform light field would be observed. However, the light field at this stage does not yet have the characteristics of a partially diffusing window, since the intensity of light is not smoothly varying across viewing angle.

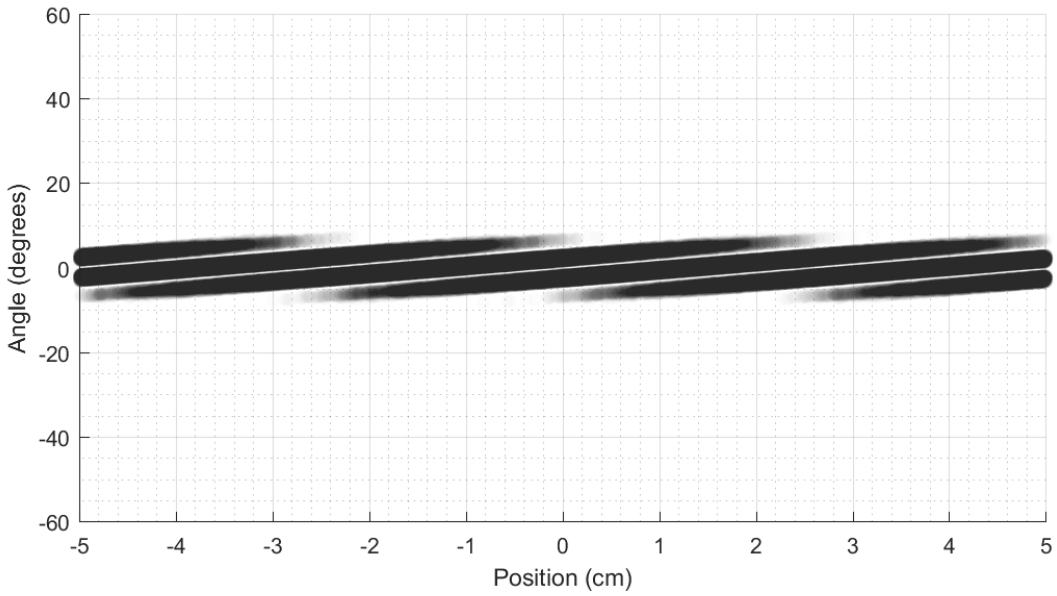

Figure 59: 2D light field before the second diffuser, showing how the spatial distribution has diverged as a function of the angle.

The role of the second diffuser is to mask the directional non-uniformity by applying sufficient angular diffusion to produce a smoothly varying light field in the angular direction. The desired

PSF can be described in a similar way as for the first diffuser, again using Equation (4). The resulting light field is illustrated in Figure 60. It exhibits a uniform distribution across spatial positions and a smoothly varying distribution across angular directions, mimicking that produced by a partially diffusing window from a distant source of light. The final angular distribution of the modelled light field has a standard deviation of 3.4 degrees.

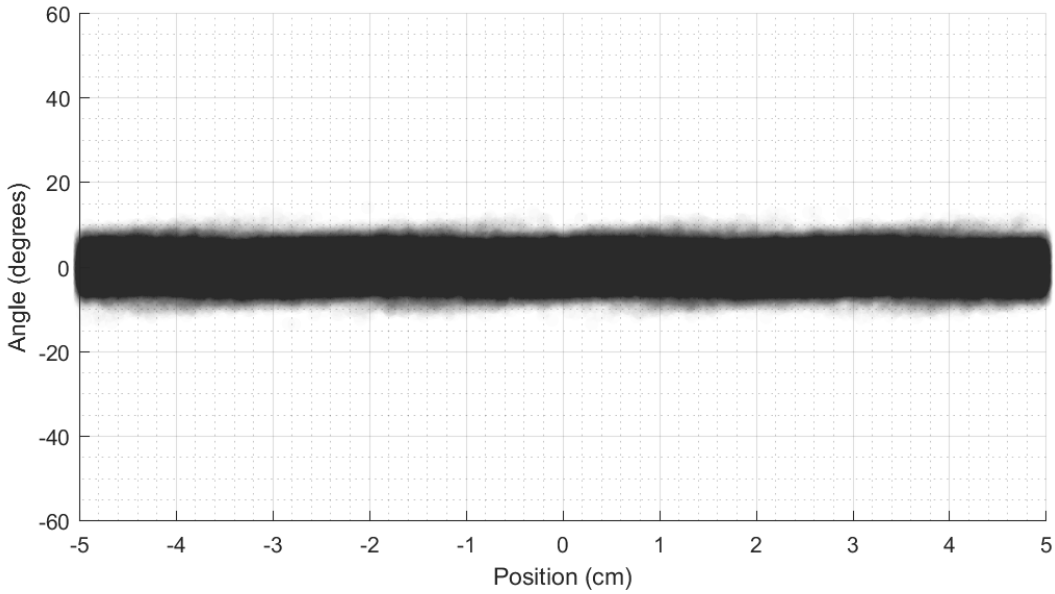

Figure 60: 2D Light field after the second diffuser, showing both spatial and angular uniformity

Comparing this to the desired light field of a partially diffusing window as shown in Figure 52 illustrates how the Imaging Stage of the optical design is able to achieve the objective of reproducing the light field of a partially diffusing window.

4.2.2 Illumination Stage

The goal of the illumination stage is to provide an efficient source of illumination for the Imaging Stage. The optical design uses an array of LEDs and an array of corresponding field lenses, configured so that light rays emitted from the LED array diverge to their maximum width $y=L_{Lens}$ at a distance along the optical axis $z=L_{Globe}/2$, half-way between the LED array and the array of globe lenses. For best efficiency, the LED would emit light with a uniform angular distribution up to $atan(L_{Lens} / L_{Globe})$, although this is not a requirement of the design. At this half-way point a field lens array focuses the diverging light into the aperture of the globe lens array. With this optical configuration, most of the light traveling through the LCD can make its way through the globe lens apertures, ensuring good overall system efficiency above 50%.

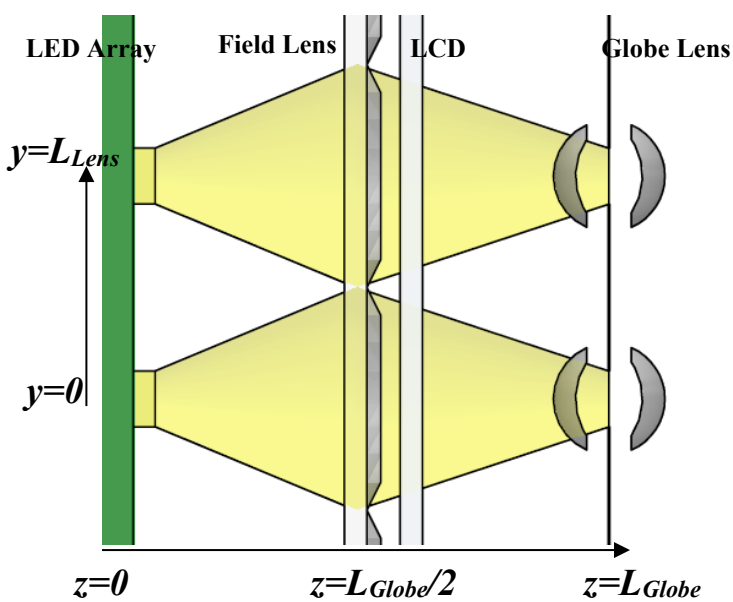

Figure 61: Illustration of optical path for the illumination stage. Yellow shading illustrates the envelope of light rays emitted by the LED array.

Figure 62 illustrates the 2D light-field distribution from four adjacent LEDs. The LED light source is specified to be at position $z=0$ on the optical axis. The distribution of light from each LED is represented as a Gaussian spatial distribution with sigma = 3.3 x 10^{-4} m, and a Gaussian angular distribution with sigma = 14 degrees. Darker shading corresponds to a higher density of rays, and thus a greater intensity of light. As can be seen, the four individual LED light sources are distinctly separated in the spatial dimension by a distance of y=2.5cm. The purpose of the illumination stage is to focus, as much as possible, all of the light emitted by the LED through the aperture of the lens array at a distance of $z=L_{Globe}$ along the optical axis.

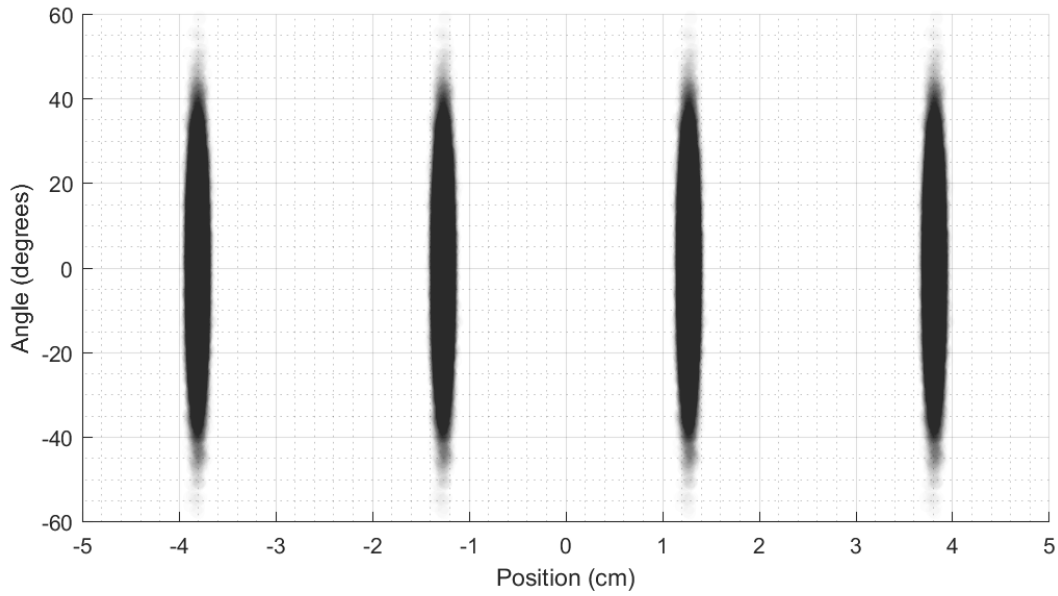

Figure 62: 2D light field of the modelled LED light sources

Figure 63 shows how the distribution changes as the light travels a distance $z=L_{Globe}/2$ along the optical axis. As these rays of light propagate away from the LED array, they diverge spatially according to the arctangent of the initial starting angle, according to Equation (5).

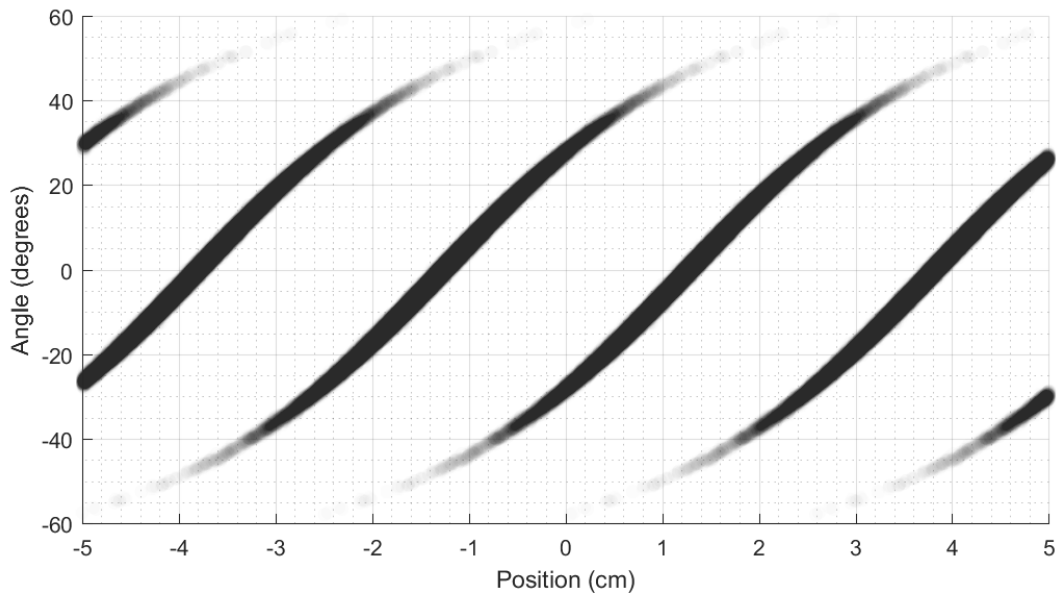

Figure 63: 2D light field before the field lens. The angle is unchanged from the light source, but the light has diverged spatially due to the propagation through free space.

At a distance half-way between the LED array and the lens array, $z=L_{Globe}/2$, the array of field lenses focuses the rays towards the corresponding apertures of the lens array. The parameters of the field lens thus have a focal length of $f=L_{Globe}/2$, diameter $d=L_{Lens}$, and $f\# = 1$. The field lens is positioned at the mid-point between the aperture of the LED array and the aperture of the lens array such that the focal points f are in the plane of the apertures. Given the requirements of large diameter and a short focal length, an array of Fresnel lenses is well suited to this application.

After refraction by the field lens array, the angular distribution is reversed to converge towards the aperture of the globe lens, as illustrated in Figure 64.

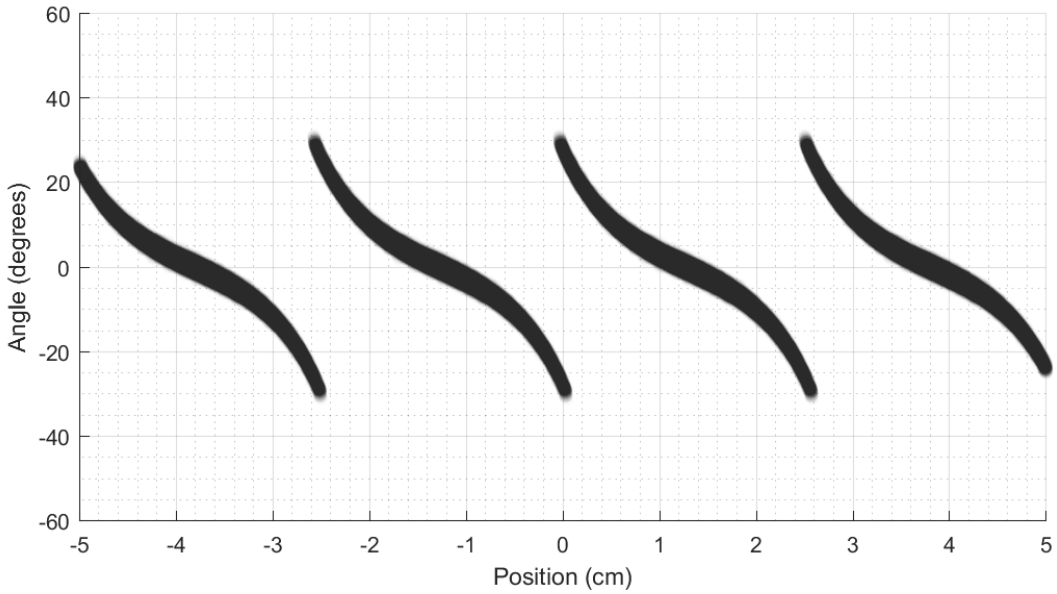

Figure 64: 2D light field after the field lens. The angular distribution has been reversed from the angle incident on the lens to focus through the aperture of the lens array

When the light from the LED reaches the aperture of the lens array, it is mostly focused through the aperture, as illustrated in Figure 65. The grey shaded area corresponds to the width of the aperture of the globe lens. The improvement to the efficiency due to the Illumination stage can be estimated by calculating the ratio of rays that propagate through the aperture of the globe lens array. Using this analysis, the ray tracing model simulates that 95% of the light transmitted by the LCD is focused through the aperture of the globe lens. Without the field lens, the simulated optical efficiency decreases to 22%.

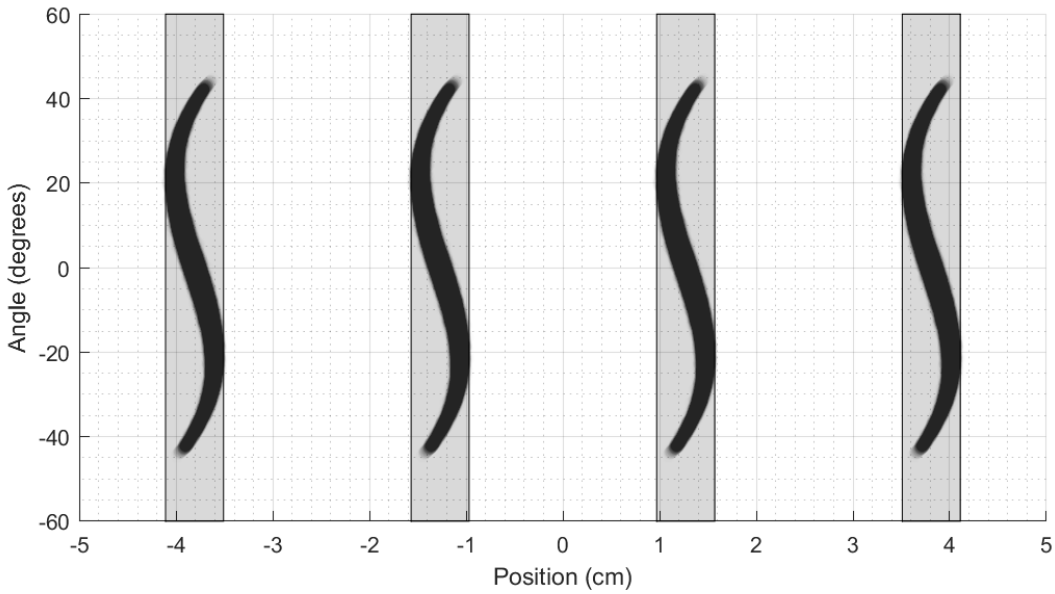

Figure 65: 2D light field at the apertures of the globe lenses, indicated by the gray shaded areas.

4.3 Performance analysis

To quantify the performance of the light-field display, the ray tracing simulation was used to determine the characteristics of the 2D light field produced by the modeled light-field display. Six key design goals were identified during the initial design iterations described in Section 4.1, required to simulate a partially diffusing window. These are summarized below:

A. Sufficiently high angular resolution

The first characteristic of a partially diffusing window is that the viewer is able to distinguish individual objects through the window as illustrated in Figure 38 and have a sense that they are separated from the window and from each other by some depth. As reported in the background literature of Chapter 2, and confirmed by the experiment of Chapter 3, greater angular diffusion leads to lower visual depth acuity and hence fewer depth planes that can be produced by the display. Although as few as two distinguishable depth planes are sufficient to convey a sense of depth, a greater number of depth planes leads to a more compelling experience. It is interesting to note that fewer than ten planes have been used to enhance the sense of realism and depth in animated cartoons (Disney, 1938).

The angular resolution can be estimated by computing the standard deviation of the direction of paraxial rays produced by the display, as illustrated in Figure 66. The standard deviation of the angular distribution is 3.4 degrees. This is on the same order of magnitude as the diffusion tested during the depth perception experiment (2.2 degrees), indicating that the optical system is able to produce on the order of 30 depth planes to provide both stereoscopic and motion parallax depth cues.

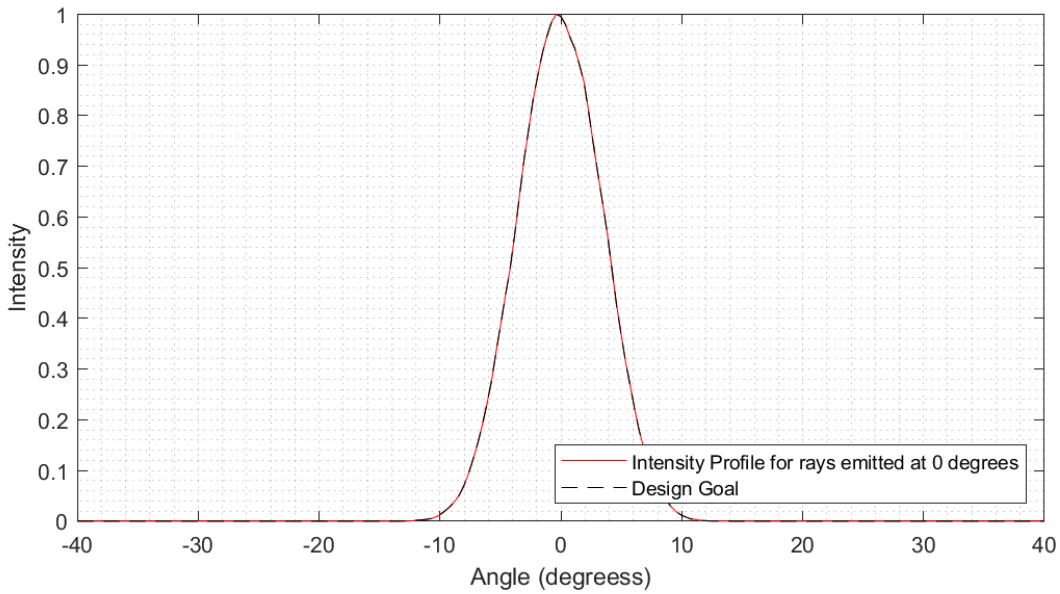

Figure 66: Angular distribution of the light field for paraxial light. The standard deviation of the paraxial light through the optimized optical system is 3.4 degrees

B. Consistent angular distribution throughout the field of view

The second characteristic of a partially diffusing window is that the image has a consistent diffusion profile across a wide range of viewing angles. To quantify this objective, the software simulation was used to calculate the intensity of one hundred thousand light rays emitted at two viewing angles (16 degrees and 27 degrees) in addition to the paraxial distribution used in evaluating the first design goal. The angular standard deviation shown in Figure 67 for the two additional angles (3.42 degrees and 3.44 degrees, respectively) was found to be within 1% of the distribution of the paraxial rays, indicating that the optical system has achieved very good consistency across a wide range of viewing angles.

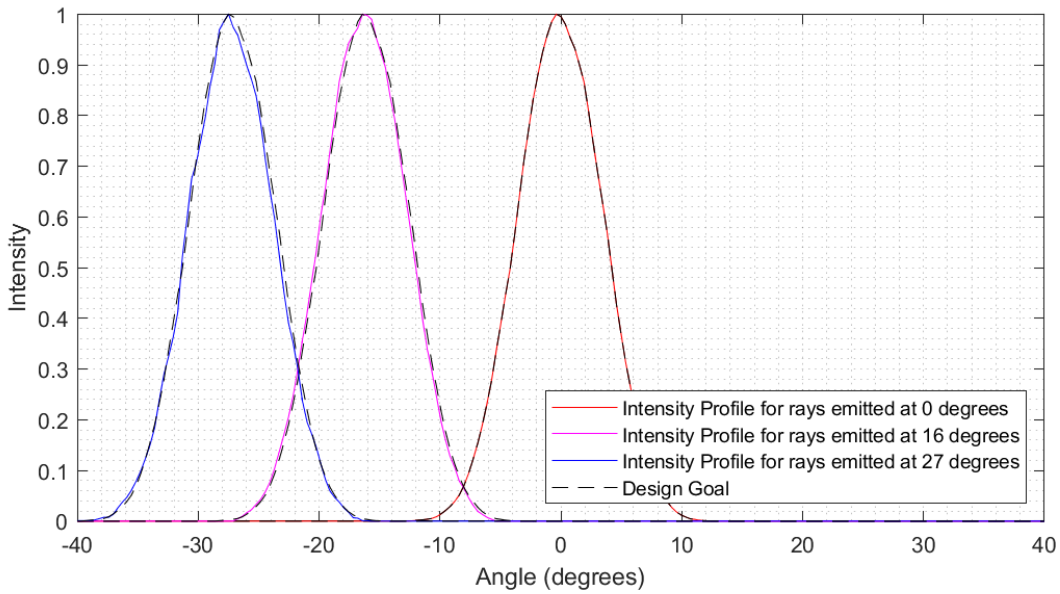

Figure 67: Angular distribution of the light field for three viewing angles.

C. Spatial Uniformity

A partially diffusing window produces a spatially uniform light field for a distant light source. Achieving spatial uniformity with a light-field display requires that there is no visible repeating pattern due to the projection of light by an array of lenses. To evaluate this the software simulation analyzed intensity as a function of spatial position with each lens of the array emitting a paraxial beam of light. The result is shown in Figure 68.

In order to accurately predict the visual quality of the light field simulations, the analysis accounts for the nonlinear sensitivity to light of the human visual system. This was modelled using a Perceptual Quantization, or *PQ* function to convert the simulated optical intensity to a non-linear representation similar to that of the human visual system (Nezamabadi, Miller, Daly, & Atkins, 2014). To convert linear intensity to perceptually-weighted intensity, an average

luminance of 50 cd/m^2 is assumed, which is typical for the brightness of indoor walls (Zumtobel, 2018). The intensity distribution is then compared to the average intensity. As can be seen, while the resulting light field is not perfectly uniform, it is within a single JND of the average, and thus is just on the threshold of appearing uniform to human vision. The spatial frequency of the pattern is determined by the spacing between the lenses and the viewing distance: for 2.5cm spacing between lenses viewed from 1m this is 1.5 cycles per degree, the upper limit of the spatial resolution of the proposed design.

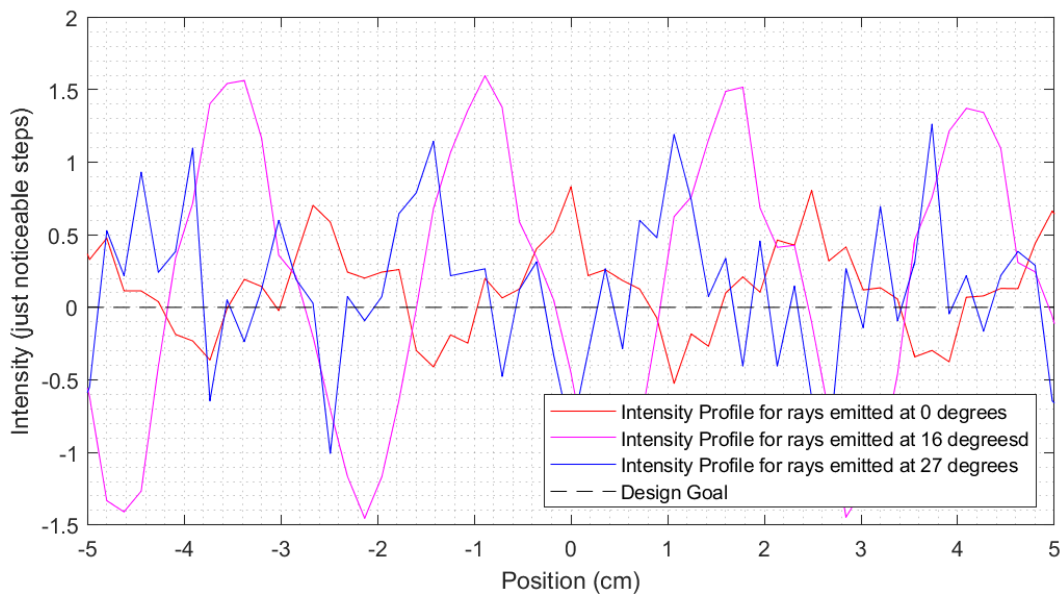

Figure 68: Spatial distribution of light field across the display. The horizontal axis is the position across the display, and the vertical axis is the difference in intensity of the light, represented as Just Noticeable Differences

D. Wide field of view

The objective for this project is a design of a light-field display with a consistent performance over a reasonable range in both horizontal and vertical directions. The range achieved by the

design described in this thesis is of ±40 degrees in both horizontal and vertical directions, which is significantly greater than existing auto-stereoscopic displays today (±6 degrees), although far less than a conventional clear window. Some applications of an intermediate-resolution light-field display may require greater viewing angle, potentially up to nearly ±90 degrees. Such extreme angles require sophisticated optical designs that are beyond the scope of this project, but are discussed in Chapter 5. Depending on the diffuser, the diffusion from the proposed optical design will cause light to be emitted at lower intensities at even wider angles, however this light will not have depth information.

E. Good optical efficiency

The optical efficiency of a partially diffusing window is determined by the proportion of incident light that is transmitted. This is less than 100% as some light is reflected or absorbed by the window; however it is typically very high. The optical efficiency of the intermediate-resolution light-field display can be similarly defined as the fraction of light emitted by the light source that is transmitted through the optical system, ignoring losses in the LCD. While not a complete metric to measure overall performance, this technique measures the percentage of light that is blocked by apertures within the optical system, which is a useful measure when designing and optimizing the optical design. By this measure, the optical efficiency was calculated to be approximately 64%. The final optical efficiency would be expected to be lower than this, scaled roughly by the optical efficiency of a conventional luminaire or UHD display, for both static and dynamic displays respectively (dynamic displays have a lower efficiency due to the physics of how the dynamic image is formed).

F. Reasonable cost

A partially diffusing window can be installed for extremely low cost – essentially that for a pane of glass. A reasonable cost for an intermediate-resolution light-field display is substantially more. There are two main categories for the cost of a light-field display, depending on whether the image it produces is static or dynamic. A static display shows a constant image at any fixed viewing position; the image only changes as the viewer moves around the display. A dynamic display can have the image change over time even at a fixed viewing position. A static display using transparent printed films in lieu of an LCD may be of similar cost to a conventional luminaire, while a dynamic display is closer to the cost of a standard high-resolution 1.4 m (55”) Ultra-High-Definition (UHD) display. The cost for the optical design proposed here will likely fall within these limits. Both the field lens array and the globe lens array could be manufactured from a molded array of plastic using conventional molding techniques. In addition to the cost of the intermediate-resolution light field module, there is the installation, power, and maintenance cost. This is expected to be on the same order as a conventional luminaire, and is further discussed in Chapter 5.

4.4 Chapter Summary

This chapter developed an optical design for an intermediate-resolution light-field display that could produce a light field indistinguishable from that of a partially diffusing window. Specifically, this design produces an angular diffusion with standard deviation of 3.4 degrees. This is roughly comparable to the diffusion used for the study in Chapter 3, where it was found that an angular diffusion with standard deviation of 2.2 degrees resulted in 31 unique depth planes.

Optical modelling confirmed consistent angular distribution and good spatial uniformity over a wide range of viewing angles (±40 degrees). The optical efficiency was calculated to be approximately 64%. By using a minimal number of readily manufacturable optical elements, the optical design can be produced at reasonable cost, similar to conventional 2D displays.

A 3D model was built with Solidworks (SolidWorks, 2016) and used to illustrate a single module of the design concept, illustrated in Figure 69. This shows the components of the optical system layered from the LED array to the diffusion system.

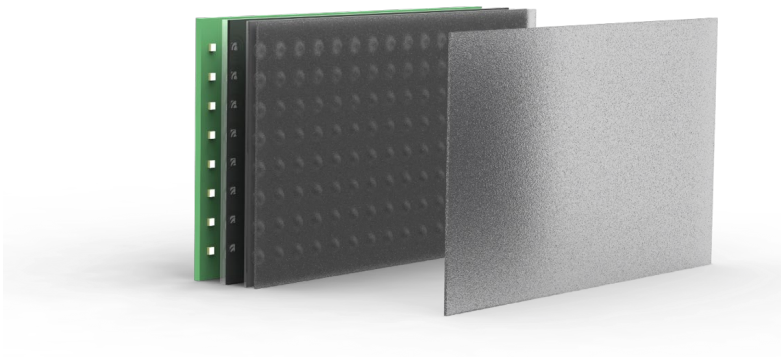

Figure 69: Optical components of an intermediate-resolution light-field display module

Chapter 5: Conclusion and Future Directions

This project set out to create a novel display system design that could create a perception of spaciousness in environments where it may be impractical to have a window. This work established that this objective could be accomplished in a practical manner in the form of a novel intermediate-resolution light-field display, and that while a physical display prototype was not actually built, it is shown that such a display could be built at reasonable cost using components based on current practical manufacturing methods.

These results were obtained in several stages. Chapter 3 described a psychophysical study to quantify the degree of depth perception that such a display could provide, and Chapter 4 developed a theoretical design that could produce the required intermediate-resolution accurate 3D light field. Several metrics were developed to compare the light field produce by the optical system with that of a partially diffusing window. Based on those metrics, and a software model of the display optics, it was shown that an intermediate-resolution display could show an image indistinguishable from what one would see through a partially diffusing window.

These results show that a light-field display system of the type developed here could be useful for improving windowless indoor environments such as shopping malls, working spaces, and living spaces, to provide the occupants with an improved sense of space. Such a display system could be fabricated as individual modules, and tiled to cover entire walls and ceilings, replacing traditional lighting systems while providing a sense of being in a much larger space.

5.1 Implications for architectural and urban design

This thesis describes a novel system which can make a windowless space more enjoyable to live in. It works by generating a display that mimics how a vast open area would appear if viewed through a partially diffusing window. As the occupant moves around they would see a changing pattern of light, just as they would if viewing a real-world scene. The mind interprets this pattern of light as implying a much greater space then there really is, which may lead to an increased perception of spaciousness, even when the occupant is not physically present in the space. Architects, researchers, and urban planners have long attempted to provide a greater sense of space and openness in windowless environments and have studied the psychological benefits of doing this. The intermediate-resolution light-field display described here is a new tool that these professionals can use to enhance the sense of space in ways that were not possible before, likely resulting in significant benefits to the occupants.

An important example is in health care. Many health care environments are windowless rooms. Artificial skylights have been installed in some hospitals and other medical facilities to calm a patient during operations and even improve post-operation recovery times (Mehrotra, Basukala, & Devarakonda, 2015). These typically display either high resolution 2D static or dynamic images on the ceiling, without stereoscopic or motion depth cues. Providing depth cues, even with a lower resolution image, may further improve their effectiveness.

Beyond health care, a virtual skylight or window could be installed in a variety of residential, commercial, or industrial environments. This could be highly beneficial within enclosed spaces such as shopping malls or offices, especially where the height of the ceiling is limited. Artificial

skylights have been installed in numerous public and private spaces to provide an aesthetically pleasing light source; lighting of this kind has been found to improve the occupant's mood or sense of calm (Canazei, et al., 2016). As in the health care example, these are typically high-resolution 2D images. Providing motion parallax to such virtual skylights might further enhance the sense of spaciousness.

In addition to an enhanced perception of spaciousness, an intermediate-resolution light-field display can also act as a source of illumination. Although light sources are most commonly placed overhead, a substantial amount of light is reflected by light colored walls, and many rooms are equipped with floor standing lamps or windows to provide a source of illumination that does not come from directly overhead. A light field display could further add to interesting alternative to ceiling-mounted illumination; by spanning an entire wall one could function both as a light source and as a window into a virtual environment. Rather than equipping a room with floor-standing lamps, a light-field wall could simulate a scene at various distances behind the wall, enabling a new tool for designers, and enabling a concept of a "virtual space".

Indoor spaces such as the Venetian Casino (Venetian, n.d.) or venues such as amusement park rides and theatre sets are often decorated with an artificial sky or background to provide a sense of "virtual reality" – producing the sensation of being in a vast outdoor space instead of an indoor room. In the Venetian Casino the sky is a painting but is slightly recessed from the foreground buildings, so that it provides a small amount of motion parallax. Providing true motion parallax as well as changing content may enhance the feeling of space and immersion in such environments.

Another interesting application is in walkways. Visitors to Chicago O'Hare airport may be familiar with the underground walkway between United Airline's B and C concourses (Underground Pedestrian Walkway at O'Hare International Airport). This walkway employs a low-tech "virtual window" to provide motion parallax. As an observer moves along the walkway, the apparent directions of structures behind the diffuse windows change according to the viewing position. This is achieved by placing objects at some physical distance behind diffusing screens, which limits how far away the objects can appear. Nonetheless it is a compelling experience. Substituting intermediate-resolution light-field displays for these panels could provide a similar experience, but with much greater "virtual depth" that also requires much less "real depth".

Transit vehicles might also be augmented by a light-field display. One interesting proposal is to reduce the incidence of motion sickness in underground subways (Chatzitsakyris, Ducla-Soares, & Zulas, 2004). Instead of looking out at a concrete subway tube, the researchers propose showing an artificial view of an above-ground virtual world that is passing by. By synchronizing the speed of the virtual world to the motion of the train, visual and vestibular motion cues can be matched, resulting in less motion sickness. An intermediate-resolution light-field display could improve this by adding stereopsis and motion parallax cues. The same technology could be potentially used in other environments known to induce motion sickness, including ships, elevators, cars, and airplanes.

Depending on the application, there may even be a cost savings associated with installing an intermediate-resolution light-field display in lieu of a window. A building envelope may be considerably less expensive without the need for a waterproof cutout as required for a window, and the thermal insulation may also be improved, thereby lowering heating and cooling expenses. Furthermore, enabling interior building spaces to have a greater degree of flexibility and functionality by providing a simulated partially diffusing window may result in more useable indoor space, thus potentially decreasing overall construction costs.

5.2 Future work

During the course of this project, several potential areas of future work were identified that could further refine the concepts developed in this thesis. Broadly speaking, these fall into two categories: perception studies and optical design.

With regards to perception, Chapter 2.2 and Chapter 2.3 described research into depth perception and the impact of providing a view of a far-away scene. However, no previous work was found that included the effect of a partially diffusing, or intermediate-resolution view. This could be a valuable direction to explore, for example by extending the study by Kahn (Kahn, et al., 2008) or Radikovic (Radikovic, Leggett, Keyser, & Ulrich, 2005), to compare the experience of a room with a real window, a partially diffusing window, and no window at all. This could be done inexpensively by simply bonding a diffusing film such as listed in Chapter 3 to an actual window. The results of this thesis suggests that the partially diffusing window could provide a significant benefit over no window at all. Another interesting possibility would be an experimental study of "open space" – in terms of the measurable characteristics of a light-field display. By correlating spatial resolution, angular resolution, and uniformity to the perceived amount of "open space", design decisions could be made more objective.

Another interesting avenue of research is to study the effect of simulating non-Lambertian reflections using a light field display. Even at very low spatial and angular resolutions, a light field display could be used as a tool for better reproducing the reflections of common building materials such as semi-gloss paints and other textures that do not contain fine detail. As spatial

and angular resolutions improve over time, this capability will increase to allow improved simulation of reflections from more complex materials.

With regards to optical design, Chapter 4.3 outlined the design for an intermediate-resolution light-field display. A potential next step would be to use this design to build a physical prototype, ideally of similar size to a conventional window. The primary challenge of building such a prototype would be the one-time costs associated with manufacturing the array of lenses. The cost of custom tooling and machining would be significant. Once the prototype display is manufactured and calibrated, it could be compared to a partially diffusing window, to verify if it is impossible for an observer to distinguish between the two.

Improvements to the proposed optical design are also possible. As discussed in Chapter 4.6, one area of potential improvement is to increase the viewing angle beyond the ±40 deg described in this thesis. Displays with viewing angles of nearly a full hemisphere would enable a greater range of potential applications. One potential technology to explore for this application is the use of an array of fish-eye lenses, as described in Chapter 2. An alternate approach could be to insert additional refractive or reflective optics to trade off vertical range for increased horizontal range, if vertical range is less important than horizontal range for a given application. It may also be similarly possible to optimize for the trade-off between reduced angular resolution and angular range.

An additional area of research is the capture or generation of images for light-field displays. Capturing a light field is significantly more difficult than capturing a single or stereoscopic

image, not just because of the amount of additional information but also due to the need to capture images from many different viewpoints simultaneously. This can be done relatively easily for computer-generated content by rendering a scene from the desired viewpoints. For real-world capture, this requires a large array of cameras partially surrounding the scene. Each camera must be carefully aligned, synchronized, exposed, and color balanced in order to produce a high-quality image. There are very few environments outside of carefully constructed studios where such a configuration is feasible, and the amount of information produced is orders of magnitude larger than a single 2D capture from a single camera. There is a significant amount of research being done to enable the capture of a scene from multiple viewing positions, and increasing the number of intermediate viewpoints via interpolation (Intel True View). Capturing an intermediate resolution has a distinct advantage that it does not require high angular sampling density, and that each viewpoint can be much lower resolution that a conventional 2D image. These characteristics of an intermediate-resolution light field are expected to greatly facilitate the capture of suitable intermediate-resolution light field images.

5.3 Conclusion

This thesis has established, both experimentally and theoretically, the feasibility of a design concept for an intermediate-resolution light-field display that could simulate the perception of depth as seen through a partially diffusing window. As technology improves over time, the spatial and angular resolutions of light-field displays can also increase, which in turn will further increase the set of practical use cases. The ultimate goal, requiring much higher resolution imaging displays and optics than are available today, is not an intermediate-resolution light-field

display indistinguishable from a partially diffusing window, but a high-resolution light-field display that is indistinguishable from a window that is completely clear.

References

Balram, N., & Tošić, I. (2016). Light Field Imaging and Display Systems. *Information Display, 32*(4), 6-13.

Banks, M. S., Akeley, K., Hoffman, D. M., & Girshick, A. R. (2008, July). Consequences of Incorrect Focus Cues in Stereo Displays. *Journal of the Society for Information Display, 24*(7), 7.

Banks, M. S., Read, J. C., Allison, R. S., & Watt, S. J. (2012). Stereoscopy and the Human Visual System. *SMPTE Motion Imaging Journal, 121*(4), 24-43.

Barten, P. G. (1999). *Contrast Sensitivity of the Human Eye and its Effects on Image Quality.* Eindhoven: Technische Universiteit Eindhoven.

Baumeister, M., & Kohnen, T. (2008, June). Accommodation and Presbyopia : Part 1: Physiology of Accommodation and Development of Presbyopia. *Ophthalmologe, 105*, 597-608.

Canazei, M., Laner, M., Staggl, S., Pohl, W., Ragazzi, P., Magatti, D., . . . Trapani, P. D. (2016). Room- and Illumination-Related Effects of an Artificial Skylight. *Lighting Research and Technology, 48*(5), 539-558.

Cassin, B., & Solomon, S. (1990). *Dictionary of Eye Terminology.* Gainesville: Triad Publishing Company.

Chatzitsakyris, P., Ducla-Soares, G., & Zulas, A. (2004). IN-Visible: Perceiving Invisible Urban Information Through Ambient Media. *Extended Abstracts on Human Factors in Computing Systems*, 1348-1350.

Cinetics Axis 360. (n.d.). Retrieved 7 19, 2019, from Cinetics Axis 360: https://cinetics.com/axis360-pro/

Cline, Hofstetter, & Griffin. (1997). *Dictionary of Visual Science.* Boston: Butterworth-Heinemann.

Collins, B. L. (1975). *Windows and People: A Literature Survey: Psychological Reaction to Environments With and Without Windows.* U.S. Dept. of Commerce, National Bureau of Standards.

Contour ShuttleXpress. (n.d.). Retrieved 7 19, 2019, from Contour ShuttleXpress: https://www.contourdesign.com/product/shuttle/

Cucker, F. (2013). *Manifold Mirrors: The Crossing Paths of the Arts and Mathematics.* Cambridge University Press.

Decorativefilms. (n.d.). Retrieved 7 19, 2019, from Decorativefilms: https://www.decorativefilm.com/

Dee, H., & Santos, P. (2011). The Perception and Content of Cast Shadows: an Interdisciplinary Review. *Spatial Cognition and Computation, 11*(3), 226-253.

Degelman, D., & Rosinski, R. (1976). Texture Gradient Registration and the Development of Slant Perception. *Journal of Experimental Child Psychology, 21*, 339-348.

Dimenco. (2019, 12 10). *Dimenco*. Retrieved 12 10, 2019, from https://www.dimenco.eu/

Disney, W. (1938). *Multiplane Educator Guide*. Retrieved 1 5, 2020, from www.waltdisney.org: https://www.waltdisney.org/sites/default/files/MultiplaneGuideCurriculumPacket_Final.pdf

Dodgson, N. (2004). Variation and extrema of human interpupillary distance. *Stereoscopic Displays and Virtual Reality Systems XI. 5291*, pp. 36-46. Proceedings of SPIE - The International Society for Optical Engineering.

Faubert, J. (2001, February). Motion Parallax, Stereoscopy, and the Perception of Depth: Practical and Theoretical Issues. *Proc. SPIE 10298, Three-Dimensional Video and Display: Devices and Systems: A Critical Review, 1029809* .

Ferris, S. H. (1972). Motion Parallax and Absolute Distance. *Journal of Experimental Psychology, 95*, 258-263.

Forbes, G. W. (2007). Shape Specification for Axially Symmetric Optical Surfaces. *Optical Society of America, 15*, 5218-5226.

Fréchet, M. (1927). Sur la loi de Probabilité de l'Ecart Maximum. *Annales de la Société Polonaise de Mathematique, 6*, 93-116.

Gershun, A. (1939). The Light Field. *Journal of Mathematics and Physics, 18*(1-4). doi:10.1002/sapm193918151

Gillam, B., & Borsting, E. (1988). The Role of Monocular Regions in Stereoscopic Displays. *Perception, 17*(5), 603-608.

Harrison, C. C., & Schnitzer, J. (1862, June 17). *United States Patent No. 35605.*

Hill, R. (1924). A Lens for Whole Sky Photographs. *Quarterly Journal of the Royal Meteorological Society, 50*, 227-235.

Hiura, H., Komine, K., Arai, J., & Mishina, T. (2017). Measurement of static convergence and accommodation responses to images of integral photography and binocular stereoscopy. *Optics Express, 25*(4).

Hoffman, D., Girshick, A. R., Akeley, K., & Banks, M. S. (2008, March). Vergence–accommodation conflicts hinder visual performance and cause visual fatigue. *Journal of Vision, 8*(3), 33.

Howard, H. J. (1919). A Test for the Judgment of Distance. *Americal Journal of Ophthalmology, 2*, 656-675.

Howard, I. P., & Rogers, B. J. (1995). *Binocular Vision and Stereopsis.* New York: Oxford University Press.

IJsselsteijn, W., Oosting, W., Vogels, I., Kort, Y., & E.J. Loenen, v. (2006). Looking At or Looking Out: Exploring Monocular Cues to Create a See-Through Experience with a Virtual Window. *PRESENCE 2006 : Proceedings of the 8th Annual International Workshop on Presence* (pp. 83-92). Cleveland: ISPR. Retrieved 3 13, 2019, from https://narcis.nl/publication/recordid/oai:pure.tue.nl:publications/8c13760b-e52e-455b-8eec-6c6d177bffc5

Intel True View. (n.d.). Retrieved 7 19, 2019, from Intel True View: https://www.intel.com/content/www/us/en/sports/technology/true-view.html

iPhone X. (n.d.). Retrieved 7 19, 2019, from Apple: https://www.apple.com/iphone-xs/

Kahn, P., Friedman, B., Gill, B., Hagman, J., Severson, R., Freier, N., . . . Stolyar, A. (2008). A Plasma Display Window? - The Shifting Baseline Problem in a Technologically Mediated Natural World. *Journal of Environmental Psychology, 28*, 192-199.

Kingslake, R. (1989). *A History of the Photographic Lens.* Academic Press.

Klatzky, R., Thompson, W., Stefanucci, J., D., G., & M. K. (2017). The Perceptual Basis of Vast Space. *Psychonomic Bulletin & Review, 24*(6), 1870-1878.

Kral, K. (2003). Behavioural-Analytical Studies of the Role of Head Movements in Depth Perception in Insects, Birds and Mammals. *Behavioural Processes, 64*(1), 1-12.

Levoy, M., & Hanrahan, P. (1996). Light Field Rendering. *SIGGRAPH*, (pp. 31-42).

Likert, R. (1932). A Technique for the Measurement of Attitudes. *Archives of Psychology*, 140: 1-55.

Lippmann, G. (1908). Épreuves Réversibles. Photographies Intégrales. *Comptes Rendus de l'Académie des Sciences, 13*, 245-254.

Lipton, L. (1982). *Foundations of the Stereoscopic Cinema - A Study in Depth.* Van Nostrand Reinhold.

Ludé, P. (2018). New Standards for Immersive Storytelling. *SMPTE 2018 Annual Technical Conference & Exhibition.*

Luneberg, R. K. (1944). *Mathematical Theory of Optics.* Rhode Island: Brown University.

Markus, T. A. (1967). The Function of Windows— A Reappraisal. *Building Science, 2*(2), 97-121. Retrieved 3 13, 2019, from https://sciencedirect.com/science/article/pii/0007362867900126

Mathworks Matlab. (n.d.). Natick, Massachusetts: The MathWorks Inc.

Matusiak, B., & Klöckner, C. A. (2016). How we Evaluate the View Out Through the Window. *Architectural Science Review, 59*(3), 203-211. Retrieved 3 13, 2019, from http://tandfonline.com/doi/full/10.1080/00038628.2015.1032879

Maureen, C., & Stéphane, M. (2002). The Texture Gradient Equation for Recovering Shape from Texture. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 24*(4), 536-549.

McComb, D. M., Tricas, T. C., & Kajiura, S. M. (2009). Enhanced visual fields in hammerhead sharks. *Journal of Experimental Biology, 212*(24), 4010-4018.

McGill, R., Tukey, J. W., & Larsen, W. A. (1978, Feb). Variations of Box Plots. *The American Statistician, 32*(1), 12-16.

Mehrotra, S., Basukala, S., & Devarakonda, S. (2015, Nov). Effective Lighting Design Standards Impacting Patient Care: A Systems Approach. *Journal of Biosciences and Medicines, 3*(11), 54-61.

Nakayama, K., & Shimojo, S. (1990). da Vinci Stereopsis: Depth and Subjective Occluding Contours from Unpaired Image Points. *Vision Research, 30*(11), 1811-1825.

Nes, F. V., Koenderink, J., Nas, H., & Bouman, M. (1967, September). Spatiotemporal Modulation Transfer in the Human Eye. *Journal of the Optical Society of America, 57*(9), 1082-1088.

Nezamabadi, M., Miller, S., Daly, S., & Atkins, R. (2014). Color Signal Encoding for High Dynamic Range and Wide Color Gamut Based on Human Perception. *Proc. SPIE 9015, Color Imaging XIX: Displaying, Processing, Hardcopy, and Applications*, *90150C.*

Niven, W. (1890). *The Scientific Papers of James Clerk Maxwell.* New York: Dover Publications.

Okoshi, T. (2012). *Three-Dimensional Imaging Techniques.* Academic Press.

Oliva, A., & Torralba, A. (2001). Modeling the Shape of the Scene: A Holistic Representation of the Spatial Envelope. *International Journal of Computer Vision, 42*(3), 145-175.

O'Shea, R. P., Blackburn, S. G., & Ono, H. (1994, Jun). Contrast as a Depth Cue. *Vision Research, 34*(12), 1595-1604.

Palmisano, S., Gillam, B., Govan, D., Allison, R., & Harris, J. (2010). Steresocopic Perception of Real Depths at Large Distances. *Journal of Vision, 10*(19), 19.

Pruss, C., Garbusi, E., & Osten, W. (2008). Testing Ashperes. *Optics & Photonics News, 19*(4), 24-29.

Radikovic, A. S., Leggett, J. J., Keyser, J., & Ulrich, R. S. (2005). Artificial Window View of Nature. *Human Factors in Computing Systems.* Portland, Oregon.

Rogers, B. J., & Collett, T. S. (1989). The appearance of surfaces specified by motion parallax and binocular disparity. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 41*, 697–717.

Rogers, B., & Graham, M. (1982). Similarities Between Motion Parallax and Stereopsis in Human Depth Perception. *Vision Research, 22*, 261-270.

Rosco. (n.d.). Retrieved 07 19, 2019, from https://us.rosco.com.

Rudd, M., Vohs, K. D., & Aaker, J. (2012). Awe Expands People's Perception of Time, Alters Decision Making, and Enhances Well-Being. *Psycological Science, 23*, 1130-1136.

Schade, W. (1950). *United States Patent No. US2500046A.*

Seto, K., Fragkias, M., Guneralp, B., & Reilly, M. (2011). A Meta-Analysis of Global Urban Land Expansion. *Plos One, 6*(8), e23777.

SolidWorks. (2016). *SolidWorks.* Tennessee: Dassault Systemes.

Swanston, M. C., & Gogel, W. C. (1986, Sep). Perceived Size and Motion in Depth From Optical Expansion. *Perception & Psychophysics, 39*, 309-326.

Underground Pedestrian Walkway at O'Hare International Airport. (n.d.). Retrieved 7 19, 2019, from Wikimedia Commons: https://commons.wikimedia.org/wiki/Category:Underground_pedestrian_walkway_at_O%27Hare_International_Airport

Venetian. (n.d.). *The Venetian Resort*. Retrieved 7 19, 2019, from https://www.venetian.com/

Ware, C. (2010). *Visual Thinking for Design.* Elsevier.

Watson, A. B., & Pelli, D. G. (1983). QUEST: A Bayesian Adaptive Psychometric Method. *Perception & Psychophysics, 33*(2), 113-120.

Watson, B. (1999). Science Makes a Better Lighthouse Lens. *Smithsonian Magazine, 30*(5), 30.

Weinstein, S. (1957, Dec). The Perception of Depth in the Absence of Texture-Gradient. *The American Journal of Psychology, 70*(4), 611-615.

Westheimer, G. (1975). Editorial: Visual Acuity and Hyperacuity. *Investigative Ophthalmology & Visual Science, 14*, 570-572.

Wheatstone, C. (1838, Dec). On Some Remarkable, and Hitherto Unobserved, Phenomena of Binocular Vision. *Contributions to the Physiology of Vision, 128*, 371-394.

Whitehead, L. (2017). *United States Patent No. US 9534746.*

Wismeijer, D. A., Erkelens, C. J., Van Ee, R. v., & Wexler, M. (2010, Jun). Depth Cue Combination in Spontaneous Eye Movements. *Journal of Vision, 10*(6), 25.

Zumtobel. (2018). *The Lighting Handbook.* Retrieved 7 19, 2019, from The Lighting Handbook: https://www.zumtobel.com/PDB/teaser/EN/lichthandbuch.pdf

Appendices

Post Study Questionnaire

How would you rate the difficulty?	Very Easy	Easy	Medium	Challenging	Very Challenging
How confident are you in the results?	Very Confident	Confident	Medium	Unconfident	Very Unconfident
Did you develop a strategy for this study?	Yes No	If Yes, can you describe it?			
Any other comments or observations?					