

Adapting Vision Correcting Displays to 3D

*Jacob Holesinger
Brian A. Barsky, Ed.
Boubacar Kanté, Ed.*



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/Eecs-2020-181

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2020/Eecs-2020-181.html>

September 21, 2020

Copyright © 2020, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

I would like to thank my mentor Professor Brian Barsky for his guidance and expertise. I would like to thank Professor Boubacar Kante for his feedback and help with revising my report. I also want to thank my fellow lab mates. Specifically I would like to thank the members of the 3D correction team. The team included master of engineering students Evelyn Yang, Sundi Xiao, Wei-Chun Hwang, Yaying Zhao, Yizhen Ding and Zhuoming Zhang as well as undergraduate researchers Jimmy Xu, Michael Qi and William He. I worked with the 3D correction team in developing the shared code used between our two projects as well as the ideas for correcting in 3D. Finally, I would like to thank my parents for their love and support.

Adapting Vision Correcting Displays to 3D

Jacob Holesinger

September 2020

Submitted to the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, in partial satisfaction of the requirements for the degree of Master of Science, Plan II. Approval for the Report and Comprehensive Examination:

Committee

Research Advisor

Brian Barsky

Second Reader

Boubacar Kante

Acknowledgements

I would like to thank my mentor Professor Brian Barsky for his guidance and expertise. I would like to thank Professor Boubacar Kante for his feedback and help with revising my report. I also want to thank my fellow lab mates. Specifically I would like to thank the members of the 3D correction team. The team included master of engineering students Evelyn Yang, Sundi Xiao, Wei-Chun Hwang, Yaying Zhao, Yizhen Ding and Zhuoming Zhang as well as undergraduate researchers Jimmy Xu, Michael Qi and William He. I worked with the 3D correction team in developing the shared code used between our two projects as well as the ideas for correcting in 3D. (Ding et al., 2020) Finally, I would like to thank my parents for their love and support.

Contents

1	Introduction	1
1.1	Vision Correcting Displays	1
1.2	Working in 3D	1
2	Background	2
2.1	Modeling the Eye	2
2.2	Point Spread Function	3
2.3	Modeling Aberrations	4
2.4	Light Fields	8
2.5	Realizations of Light Field Displays	8
2.6	Crosstalk (Violating the One to One Assumption)	11
2.7	Spatial and Angular Resolution	12
3	Past Work	13
3.1	Single Layer Displays	13
3.2	Multi Layer Displays	14
3.3	Light Field Displays	15
3.4	Error Proofing	18
4	Approach	19
4.1	Implementation	19
4.2	Code	20
4.3	Optics Calculations	20
4.4	Coordinate System	21
4.5	Retina	21
4.6	Sampling in precomputation	23
4.7	Picking Mask Hardware Parameters	24
4.8	Pinhole Diameter	27
4.9	Mask Offset	27
4.10	Simulation	29
4.11	Pupil	31
4.12	Mask Effects	31
5	Results	32
5.1	Experiments	32
5.2	The Double Image Artifact and Adding an Aperture	32
5.3	Perpendicular View	33
5.4	Angled Display	34
5.5	Varying Levels of Refractive Error	35
5.6	Error Resiliency to Angle Changes	36
6	Discussion	37

7	Future Work	39
7.1	Adding Global Optimization	39
7.2	Precise Pose Estimation	39
7.3	Accounting for Both Eyes	39
7.4	Study with Video and Continuous Changes	40
8	Conclusion	40

Abstract

Past works in vision correcting displays have demonstrated that it is possible to computationally correct errors in a viewer's eye without the need of additional lenses. This result has exciting implications both for the convenience of seeing up close without reading glasses and for the possibility of correcting aberrations that are not possible fix with corrective lenses such as glasses or contacts. While there has been a great improvement in the quality of images created by vision correcting display algorithms, the approach has yet to be extended out of a 2D setting where it is assumed that the view of the display is perpendicular. Past works have used this assumption because it is both convenient for doing optics calculations and covers a wide variety of viewing conditions. Still, in practice there are many situations where it is easier or even necessary to view the display at an angle. The purpose of this project is to allow vision correcting display algorithms to work in full 3D where correction can be accomplished under any viewing condition. This fifth year masters report was part of a collaborative effort with a Masters of engineering capstone project (Ding et al., 2020). Over the course of the year, we implemented a 3D version of past vision correcting display algorithms. We found that apart from the added difficulty of accounting for angles in the optics calculations, there were places where the techniques from past works had to be modified or augmented. We additionally conducted preliminary experiments to help gauge the added challenges of working in 3D as well as the efficacy of our modifications to the original approaches. The results demonstrate a first step in the realm of 3D correction and we hope that these changes to the existing methods will help build towards the final goal of making vision correcting displays practical for every day use.

1 Introduction

1.1 Vision Correcting Displays

Vision correcting displays allow viewers to see clear images from their device without the need of additional corrective lenses. The motivation for computational vision correction is twofold. First, while lenses are a convenient solution for most vision problems, there are specific conditions that they can not easily solve, namely higher order aberrations. Cameras and other optical systems that are able to achieve higher order correction do so using a complex chain of specially designed lenses. Computational correction would reduce the physical setup drastically. The second point of interest is that many people possess some form of farsightedness and as a result have difficulty seeing close. Vision correction in display form could add a great amount of convenience because a large amount of close tasks are now conducted on phones and tablets.

The guiding principle of computational correction is that by taking into account specific knowledge of an eye's optical aberrations, the images to be shown on a display can be altered to pre correct for errors that the eye will introduce. This process has been referred to as tailoring, precorrection or prefiltering. If successful, the viewer looking at the pre corrected image on the display with their naked eye will see an in-focus image as someone with perfect vision would see when they are looking at a normal display.

Progress has been made in computation correction through both improved algorithms to alter images as well as display hardware. Despite the difficulty inherent in the problem, this progress has been substantial. Computational correction has been demonstrated in both physical experiments using a camera to mimic a viewer's eye and in user trials. (Huang et al., 2014), (Pamplona et al., 2012). Success and discoveries from past work are discussed later in section 3.

1.2 Working in 3D

The goal of this work is to extend the approaches into a 3D setting. Past work has reduced the complexity of the problem by aligning all optical elements along the same axis. This formulation has been a natural choice to pioneer the idea of computation correction because it makes optics calculations simpler while still covering a wide range of viewing conditions. Now that a correction has been demonstrated in this setting, we hope to build on past success by relaxing assumptions on display angle.

Displays will often be at some slight angle to the eye's optical axis. The simplest reason is that the viewer will want to have the display positioned in between their eyes meaning that neither eye will view perfectly perpendicular. A second reason is that angled viewing is convenient in many situations such as hand held displays. Finally, viewers of the display will naturally want to shift as they view the display over a longer period of time.

By allowing for varied display angle and position, we hope to bring vision correcting display technology closer to being available for everyday use. The project has two main parts. First, we identify the components of existing approaches that need modification to work in full 3D and second, we test the performance of the algorithms in this new setting.

2 Background

Vision correcting display research lies at the intersection of optics and optometry. This section provides a mix of background information to motivate ideas discussed in past work and in this project. The first background sections are concerned with modeling the eye and its aberrations. Section 2.2 discusses point spread functions which are helpful in understanding optical errors within the eye and early approaches to computational correction with single layer displays discussed in section 3.1. Afterwards, sections 2.4 and 2.5 discuss light fields and light field displays which have become core concepts in the forefront of vision correcting display research.

2.1 Modeling the Eye

There are many well researched eye models in the field of optometry. Among this variety are models that include a high level of complexity and accuracy as well as more approximate models that give close results while being much easier to implement. Finite models are an example of a more complicated approach. These eye models account for an extensive collection refractive elements within the eye including details like the gradient index nature of the crystalline lens and the spacing between the components such as the pupil, cornea and back of the eye.

A class of commonly used approximate eye models are Paraxial models. Paraxial eyes use the assumption that the eye's viewing target is near the optical axis and the field of view is small. In reality, our eyes have a relatively wide field of view when compared to many optical systems. The paraxial approximation allows ray calculations to be done using the small angle assumption $\sin(\theta) \approx \theta$. Furthermore, paraxial models are comprised of a much smaller set of refractive surfaces each of a constant index of refraction.

For this project, it made sense to use a Paraxial model. While a core goal of our project was to reduce the assumptions required on viewing conditions, the Paraxial model's approximation is close enough to results that a more complicated model could produce while being easier to implement and less computationally costly. To this point, it is fair to assume that the eye will always be directed at roughly the center of the display and thus that the viewed object will be close to the optical axis.

Of the possible paraxial models, we chose the Emsley reduced eye. The model consists of a single refractive surface transitioning from air to a material with index of refraction at $\frac{4}{3}$. The reduced eye model is commonly used in optometry examples and illustrations because it is comprised of a single refractive element.

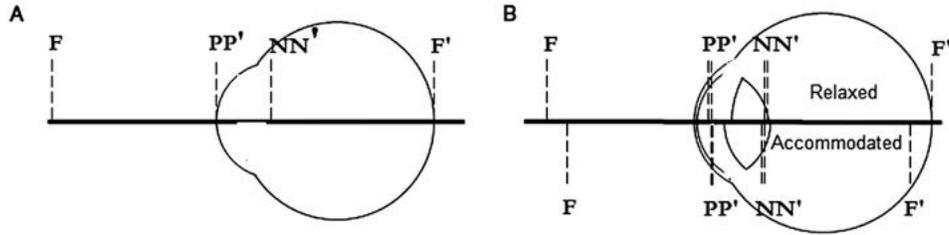


Figure 1: Two common paraxial eye models. The front and back focal points are labelled F and F' , the nodal points NN' and the principal point PP' . On the left is an example of the Emsley reduced with a single refracting surface and on the right is an example of the more complex Gullstrand–Emsley simplified eye which includes 3 refractive surfaces. Images from (Atchison and Thibos, 2016)

2.2 Point Spread Function

There are many possible sources of optical error in any imaging system. A convenient and accurate way of representing the overall combination of them all is with a point spread function (PSF). The PSF of an imaging system is the response to a single point source of light. If the imaging system is perfectly error free, the image of a point should be another point. For example, if one was to view the stars with good vision they would all appear as points of light. However, to someone with refractive error without their corrective lenses each of these points in the sky might appear as blotches of blur. To a good approximation, knowing how an imaging system acts on a point tells you how it will act on imaging a whole object. Accounting for how the PSF of the eye changes at different depths, it is possible to simulate how an eye with refractive errors would see an otherwise un-distorted image or scene. (Barsky, 2004) Due to its ability to summarize error in a simple form, PSFs are commonly used to visualize optical error. They also served an important role in early computational correction works which hoped to correct for the errors the PSF represented by applying an "inverse" of that PSF to the image in the precomputation step. This is discussed more in section 3.

An example of a point spread function is shown in figure 2. In this series of images, an increasing level of coma aberration increases the area of the point spread function which indicates a larger level of blur and error in the optical system.



Figure 2: The effect of varying levels of coma aberration shown using the point spread function. The level of coma increases from left to right.

2.3 Modeling Aberrations

Optical problems in vision occur when the eye does not properly focus light on the retina. This can be the result of both monochromatic and chromatic aberrations. Chromatic aberrations are due to the fact that our eyes bend smaller wavelengths of light such as blue light more than longer wavelengths of light such as red light. The effect on our overall vision from chromatic aberration is low compared to the effects of Monochromatic aberration. Monochromatic aberrations, which occur even in light of a single wavelength, can be classified further into the two groups of higher order aberrations and lower order aberrations.

The classification of aberrations into higher order and lower order comes from the perspective of Zernicke polynomials, a mathematical tool for representing the aberrations in the eye's wavefront. Similar to how a Taylor series could be used to approximate any function, any complex of aberration in the eye's vision can be described by a weighted sum of Zernicke polynomials. Furthermore, each polynomial represents a specific kind of optical aberration and the coefficient weighting the polynomial corresponds to how prevalent it is in the overall combined effect. The polynomials are indexed by a radial order and an angular frequency. Higher order aberrations are aberrations represented by Zernicke polynomials of radial order 3 and above while lower order aberrations are 2 and below. The Zernicke polynomials are commonly displayed in a pyramid as shown in figure 3.

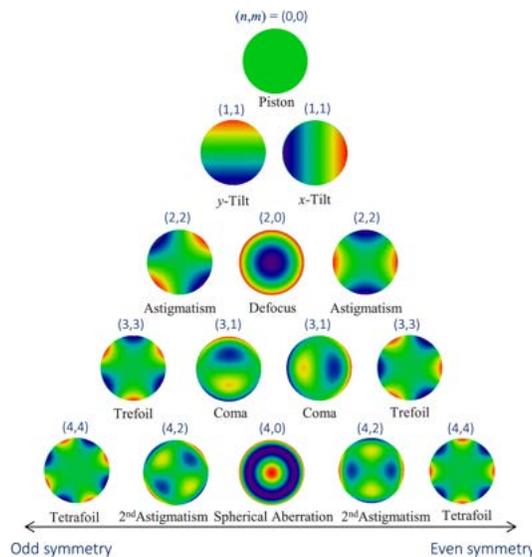


Figure 3: The 4th order and lower Zernicke polynomials shown as wavefront maps. Lower order aberrations consist of terms of 2nd order and 1st order. Higher order terms are 3rd order and higher. Diagram from (Hsieh et al., 2020)

Lower order aberrations account for on average $\geq 90\%$ of the aberration in the human eye. The most common manifestations of lower order aberrations are myopia (nearsightedness) and hyperopia (a form of farsightedness). These two effects are similar in that they both create defocus aberration represented by the Zernicke polynomial of index 2,0. These specific conditions are discussed more later on. (Iom, 2010)

An important difference between the two categories is that higher order aberrations can

not be easily corrected with traditional glasses or contacts. For an example of a higher order aberration we can look at coma aberration. The point spread function of coma aberration was shown earlier in figure 2. In single lenses, coma aberration is more present in the light entering off-axis. The only ways of correcting for single lens coma aberration are by precisely tuning the curvature of the lens to work for a specific viewing distance or by introducing a tight aperture to block off-axis light from entering. (Hecht, 1987)

Optimizing the lens for a very specific viewing condition in this way is not enough for most applications. For this reason, more inventive fixes have been studied. One example are multi-lens designs such as the Cooke triplet which have been optimized to correct for coma aberration. This has solution has the downside of a impractical form factor, certainly not something that could be conveniently worn for every day use. Another method of correcting coma specifically in eyes is designing specialty contacts with contours to fit irregular shaped corneas. This has been researched to help people with Keratoconus, a condition that is caused by the thinning of the cornea which results in an irregular refractive surface. The eyes of people with Keratoconus have a higher than average amount of higher order aberrations, particularly coma aberration. As such Keratoconus patients have blurred vision that is a challenge to correct. Past work has demonstrated how contacts specially designed for the person's prescription and cornea surface can correct for higher order aberrations in a convenient form. (Klien and Barsky, 1995).

Like all other higher order aberrations, the fact that coma aberration can not be corrected in a single lens means that people who have the aberration present in their vision are without a convenient solution.

For most people, the error due to higher order aberrations is relatively small. Still, most eyes have some level of higher order aberrations present and these can increase with age and in low light environments where the pupil dilates. (Navarro, 2019). There are also eyes such as those of people with Keratoconus where higher order aberrations have a significant effect on the person's vision even in the most optimal viewing conditions.

In contrast to higher order aberrations, lower order aberrations can fortunately be fixed with single lenses for essentially all viewing conditions. The most common manifestations of lower order aberrations are Myopia, Presbyopia, Hyperopia and Astigmatism. Myopia, illustrated in figure 4, occurs when light is focused in front of the retina.

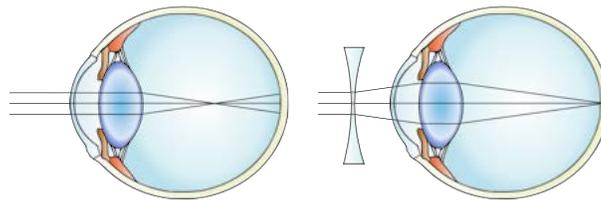


Figure 4: A myopic eye (left) corrected by a diverging lens (right).

Hyperopia and Presbyopia are the result of light being focused on a plane behind the retina. While presbyopia has the same effect on vision as Hyperopia, it is distinct in that it is a condition common with aging due to the hardening of the eyes crystalline lens. Both result in the same condition of "farsightedness" which will be discussed later along with the accommodation of the eye.

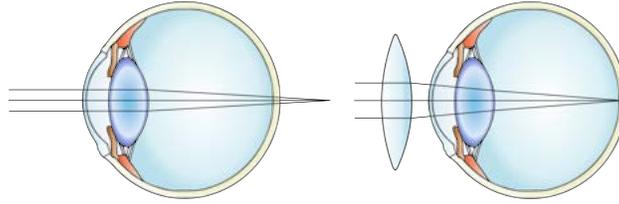


Figure 5: A hyperopic eye (left) corrected by a converging lens (right).

Another common lower order aberration that has a more complex effect is astigmatism, an aberration where light is bent more along one axis of the lens than the other (shown below in figure 6) Like the other lower order aberrations, astigmatism can be corrected by a single lens, however it can not be modeled as easily as the others, because the difference in focusing power of the two axes makes it impossible to define a single focal point.

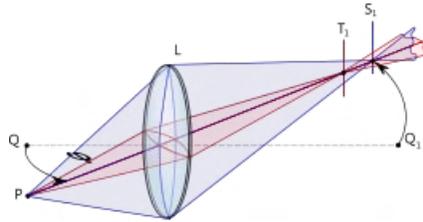


Figure 6: An illustration of astigmatism. The bend given to light along the vertical axis shown in blue is different than the bend given on the horizontal axis. As a result, the lens has an effective two mismatched focal points T and S for the two axes.

To simplify our investigation, we left higher order aberrations for later and focused on correcting specifically the aberrations of Myopia, Hyperopia and Presbyopia which all result in a defocus error. These are both easy to model and represent the most common forms of vision problems. Additionally, our hope is that implementing 3D correction for lower order aberrations as a first step, will make it easier to eventually reach the more complicated higher order aberrations.

We needed a way to create defocus error in the reduced eye model in order to simulate Myopia or Presbyopia and Hyperopia. To do this, we change the focal length of our eye model to create a mismatch between the eye's focal plane and the position of the display. In the situation where the incident angles are assumed to be small, the relationship between the distance to the object D_O , the distance to the image plane D_I and the focal length f is given by the lens maker's formula:

$$\frac{1}{f} = \frac{1}{D_O} + \frac{1}{D_I}$$

. Figure 7 shows the quantities in the lens maker's formula for an in-focus image.

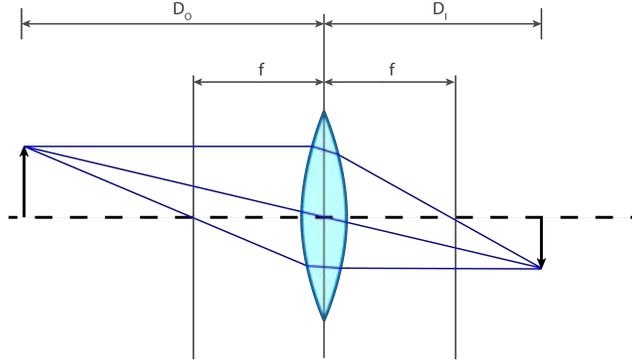


Figure 7: The relationships between the object plane and image plane for a thin lens.

The focal lengths for imaging objects at a given distance with a specific defocus error can be calculated using this formula. In the reduced eye model, relaxed eyes have a power of 60 Diopters. This corresponds to a focal length of $\approx 22.22mm$ (GuangMing, 2008)

$$f' = \frac{\eta'}{P} = \frac{\frac{4}{3}}{60} \approx 22.22mm$$

A decent approximation for eye length is $24mm$. There is a range of values amongst different people most between $22mm$ and $25mm$.(Bhardwaj et al., 2013) Combining this and the focal length for a relaxed eye we can calculate the object distance:

$$\frac{1}{f} = \frac{1}{D_i} + \frac{1}{D_o}$$

$$D_o = \frac{D_i f}{D_i - f} \approx 300mm$$

Displays at this distance would be in focus. To simulate Myopia for example, an adjustment of $-2D$ could be added to the $60D$ power that the original relaxed eye had. This would give a new focal length $f \approx 21.5mm$ and an object distance of $D_o \approx 206.4mm$. If the display was left at its original distance of $300mm$, the resulting image would have a defocus error of $300mm - 206.4mm \approx 93.6mm$

While describing this using the eye's power in Diopters matches what a real world prescription would look like, it was easier to parameterize the focus of the eye by the position of its focal plane. Putting the focal plane to far forward past the display would simulate Myopia and similarly putting it too near would simulate Presbyopia and Hyperopia. To accomplish this alternate formulation, a similar calculation can be done where the desired level of blur is found by changing the distance to the image D_I and leaving the eye's focal length fixed.

It is worth noting that this formulation does not account for the eye's ability to accommodate. Our eyes have muscles to stretch their crystalline lens which changes their focusing power to focus on different depths. In reality, there is a range of depths that we can see in focus. For Myopic people, accommodation is enough to move the focusing of light back onto the retina at viewing distances closer than their near point thus giving them "nearsightedness". For people with Hyperopia, their eye's accommodation can push the focusing forward to the retina at viewing distances farther than their far point giving them "farsightedness".

Presbyopia can be described as the loss of the ability to accommodate. As the crystalline lens hardens with age it loses its ability to stretch. While viewing distant objects is unaffected, the loss of accommodation makes it so that the eye cannot adapt to view close objects resulting in a "farsighted" effect similar to Hyperopia. Presbyopia is the reason why many older people wear reading glasses.

2.4 Light Fields

Light fields are an important concept in the research of vision correcting displays. Past research including our own 3D correction algorithm relies on creating light fields to be shown by light field displays. This section covers what light fields are and why light field displays are helpful in the context of vision correction, section 2.5 talks about various ways of constructing light field displays and finally section 3 provides an overview of past work organized by the display setups used.

In a normal display, each pixel would be described as emitting light in every direction. Similarly, in an image captured from a normal camera, each pixel in the resulting image would be the result of collecting light that has entered many different points on the camera's aperture. In light fields, each direction is treated separately. Light field displays control both the color at a point on the display as well as the direction at which it leaves the display. Light field cameras capture data for both the color at a given point on the image plane as well as the direction that the light was traveling in (equivalently where it entered the camera's aperture).

The additional information about light directions in the data captured by light field cameras allows for a slew of post processing techniques such as refocusing. Light field displays take advantage of the control over how light can escape the display in order to manipulate how the light enters the user's eye. This has been used to achieve many interesting effects including creating virtual focus planes offset from the display's actual location. (Kramida, 2015)

The light emitted from a light field display has been channeled into precise directions. As a result when color is given to a light field ray, that color will only enter the user's eye at a specific aperture point and a specific incident angle. With information about the position of the user and the aberration of their eyes, the colors along light field rays can be designed so that the light field forms the desired image on the user's retina.

2.5 Realizations of Light Field Displays

Light displays need to control both the origin and direction of their emitted light. There are two main designs commonly used in vision correcting displays which accomplish this requirement. The first are mask based displays and the second are lenslet array displays. An example of each is shown in figure 8

In our 3D approach, we used a pinhole mask based light field display. While this design does have its drawbacks which will be discussed more later, it is easier to construct and cheaper than other designs. We chose it for these reasons with the goals of eventually running physical experiments.

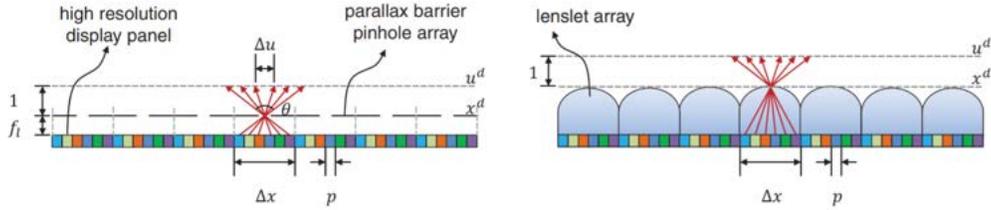


Figure 8: Light field displays created using a mask based display (left) and a lenslet array based display (right). Diagram from F. Huang (Huang, 2013) (Huang et al., 2014)

Mask based displays are comprised of two layers. The first layer is for emitting light and can be comprised of a Liquid Crystal Display (LCD) or organic light emitting diodes. A second mask layer is placed on top of the light emitting layer and is used to direct the light emitted from the bottom layer. The openings in the mask layer become the "pixels" of the light field display in that the user can only see light escaping these openings. For any given pixel in the bottom layer, emitted light can only escape through the above opening and as a result is channelled into specific directions. Figure 9 shows how the direction of light exiting a mask opening can be changed. Here the 4D nature of the emitted light field can be seen. Each light field ray is described by the combination of a 2D mask opening index and the 2D position within the grid of pixels below.

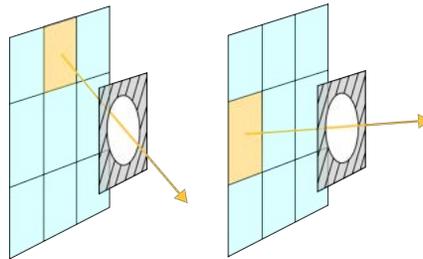


Figure 9: An illustration of a single mask opening showing how mask based light field displays can control the direction of the light they emit. Each opening in the mask represents one "pixel" of the light field display. The color as well as the direction of light coming from the opening can be manipulated. Here for example, changing the location of the orange pixel changes the direction of the outgoing orange light.

Figure 10 shows an example of the pixel values that would be set on the bottom layer of a mask based display. Each square of 5x5 pixels falls under a single mask opening allowing 25 different directions to escape that opening. Each direction of light escaping contributes to one "view" of the display. A 5x5 display such as this one could have up to 25 different views although due to the eye's size and its aperture only portions of a few views will be able to make it to the retina. To visualize the different "views" of the light field display, we can take one pixel from each 5x5 grid in the overall image. Example images corresponding to the different views are shown in the right of figure 10.

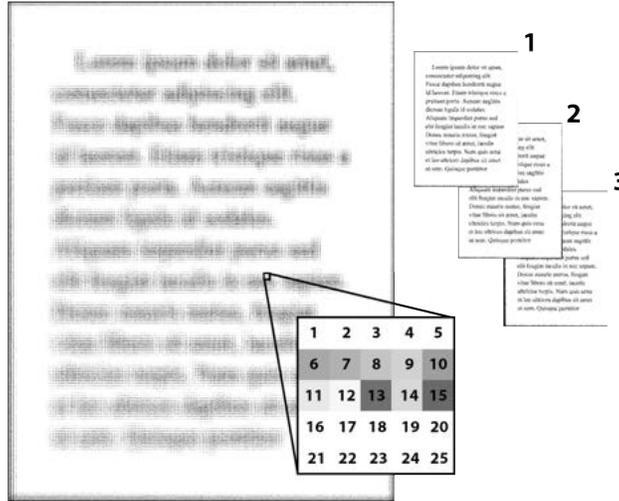


Figure 10: An example of the pixels that would be set on the bottom layer of a mask based light field display in order to create the corrected perceived image for the onlooking eye. After taking one pixel from each 5x5 grid the composite will represent one of the 25 "views" of the display.

Past works have constructed the mask layer out of an array of pinholes (Huang et al., 2014) or using a second LCD screen that blocks the light of the bottom LCD (Pamplona et al., 2012). The benefit of using an LCD for the mask is that the pitch of the mask openings can be dynamically changed to fit different viewing conditions. Pamplona et al. discuss adjusting mask parameters and show how to optimize for different situations such as viewing up close on an e-reader or far away from a television. (Pamplona et al., 2012) Pinhole based mask displays have been researched to a greater extent because they have in practice achieved similar results as LCD based masks with lower cost and complexity.

The second kind of light field display employed in the field of vision correcting displays is lenslet or lenticular based displays. Lenslet displays have the same bottom layer used to emit light. The second layer takes the form of an array of micro-lenses in place of a mask. Each micro-lens directs light from pixels in the bottom layer as a mask opening would. The benefit of using micro-lenses is that much more light is able to escape the second layer creating a brighter perceived image. While pinhole mask based light field displays are much cheaper, they have the issue of creating dim perceived images. An illustration of the different amounts of light passed by the second layer is shown in figure 11.

The ideal choice for the focal length of the micro-lenses and the offset of the second layer is so that the plane of focus for the lenses is lined up with the user's eye. Using the lens maker's formula with the offset between the layers as $depth$, the focal length as f and the distance to the plane of focus as d_F this gives:

$$\frac{1}{f} = \frac{1}{depth} + \frac{1}{d_F}$$

. The easier of the two variables to change is the offset: $depth$. Specific applications such as reading up close should be given a different offset than applications where the display is far away. In practice, the offset is very small compared to the distance to the viewer so the

approximation of setting $depth = f$ and having the focus of the micro-lenses be at infinity $d_F = \infty$ gives a good fit across most viewing distances. (Pamplona et al., 2012)

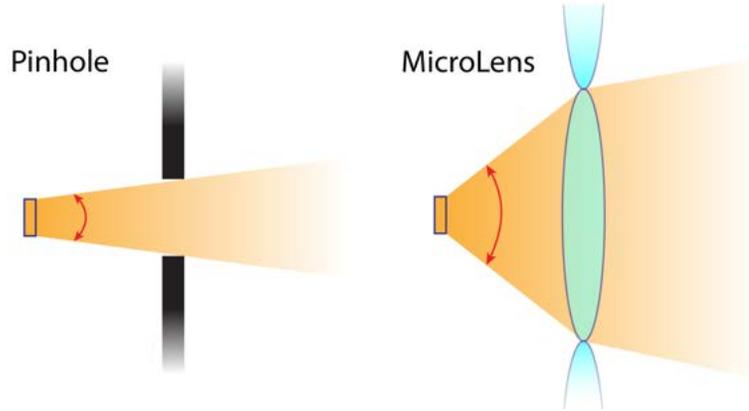


Figure 11: A comparison of the light escaping between a single pinhole and a microlens shown in 2D.

2.6 Crosstalk (Violating the One to One Assumption)

A situation that is important to avoid in light field display correction is when the light from a single pixel is able to reach the eye through more than one mask opening or micro-lens. In this situation, the light being sent in multiple directions will enter the eye at different points and angles on the aperture likely resulting in two different retina landing locations. This can be described as cross talk between mask openings because light reaches the eye through one opening originating from an incorrect pixel belonging to a neighboring mask opening.

The assumption that this situation doesn't occur has been called the one to one assumption (not to be confused with one to one sampling described later in figure 20). Pamplona et al. noted that cross talk had the effect of introducing noticeable double image errors in their results. (Pamplona et al., 2012) Their work around was tuning their light field mask parameters to minimize the number of occurrences. Huang et al.'s 2014 work implicitly made the one to one assumption in their least squares optimization. Each pixel of the underlying display was counted as a single light field ray in the optimization. The occurrences of cross talk or one to one violations are tied to the level of angular resolution employed. Thus it is possible that Huang et al didn't have as significant of an issue with cross talk as Pamplona et al. because their results were generated with lower angular resolution displays. For Huang et al's average viewing distance, the average number of views from each mask opening was reported to be around 1.66. (Huang et al., 2014)

Y. Zhen demonstrated in his masters thesis that the effect of seeing two different directions from the same pixel only occurs when the display is close enough to the viewer. (Zhen, 2019) He found that for a 3x3 display mask the minimum distance to see this effect was 460cm and 160cm for a 5x5 mask. Holding the display closer than this minimum distance made it possible for violations of the one to one assumption to occur. Since the average working distance from phones is around 36cm this effect is very likely present. (Bababekova et al., 2011)

In our results, we experienced a similar double image that can be attributed to cross talk or violation of the one to one assumption. Similar to what was found in previous work, the double image artifact was likely due to too many light field views reaching the eye’s retina. By introducing a tighter aperture, we were able to block portions of conflicting views and eliminate the artifact and the cost of darkened perceiving images. This is shown in section 5.2 of our results.

2.7 Spatial and Angular Resolution

An inherent obstacle in mask based vision correcting displays is a tradeoff between spatial and angular resolution where to increase one the other must decrease.

Angular resolution is the number of different directions that can be emitted from a single mask opening. Having a higher angular resolution means that there are more ”views” of the light field. In Huang et. al’s 2014 paper, it was shown that increasing angular resolution results in a more ”well posed” problem meaning that created light fields have a higher potential to match the ideal image when they reach the retina.

Spatial resolution on the other hand is what would be referred to as ”resolution” in the context of normal displays: how tight the pixels are packed. When viewing a light field display, viewers only see light exiting from the second layer openings. For each grid of pixels in the bottom display layer underlying an opening, the viewer only sees one source of light. Thus the spatial resolution of the light field display is the resolution of the mask openings.

Creating a higher angular resolution would mean that there are more bottom layer pixels for each top layer mask opening so that more outgoing directions can be chosen. Assuming that the size of the bottom layer display remains constant, increasing angular resolution by changing the top layer would require having fewer mask openings to cover the same bottom layer pixels and would create a lower spatial resolution. Conversely, having a lower angular resolution would mean fewer bottom layer pixels for each mask opening and a higher spatial resolution.

To give an example, a 2x2 pinhole mask, where there is a grid of 4 pixels under each pinhole, would have lower angular resolution than that of a 3x3. On the other hand, the 3x3 mask configuration would have lower spatial resolution than a 2x2 because it would drop the resolution of the underlying LCD display by 3x as opposed to 2x. This example is illustrated in figure 12.

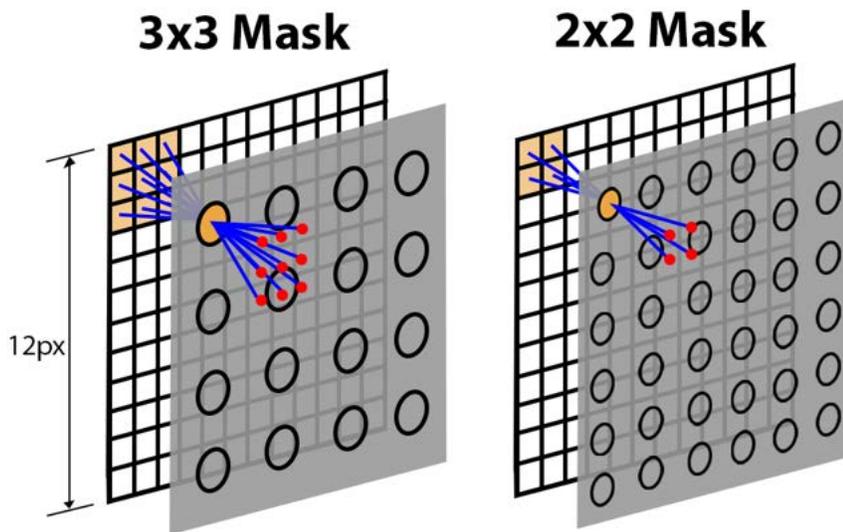


Figure 12: An illustration of the spatial vs. angular resolution tradeoff. On the left is a 3x3 mask light field display which has higher angular resolution (9 directions per opening) and lower spatial resolution (16 total openings). On the right is a 2x2 mask with lower angular resolution (4 directions per opening) and higher spatial resolution (36 openings).

3 Past Work

Works leading up to current vision correcting displays have been lead by the possibility of being able to alter images in a way that undoes the errors introduced by a viewer’s eye. Before computational correction research, optical errors in the eye had been well studied and work had even been done in simulating what an eye with a given set of optical errors would perceive in place of a correct image. (Barsky, 2004) Computational correction was the first time that researchers looked at the inverse problem of accounting for errors in images before they enter the eye. This process of altering images ahead of time is referred to as a pre-computation, prefiltering or tailoring process.

3.1 Single Layer Displays

Early designs were crafted for use with traditional displays. In this setting, researchers investigated ways to change the pixel values within a 2D image to reverse the blur that optical errors would introduce. One approach by Alonso et. al used a Wiener filter to incorporate information about the eye’s point spread function. (Alonso and Barreto, 2003) Convolution an unaltered image with the point spread function representative of the eye creates a version of the image where that eye’s optical errors have been introduced. Alonso et. al hoped to achieve the opposite effect in their pre processing with a Wiener filter. They deconvolved the image with an inverse of the point spread function ahead of time to cancel out the effects represented by that point spread function. Alonso et. al tested this same approach with point spread functions representing higher order aberrations in a later work

(Alonso and Barreto, 2004)

F. Huang and B. Barsky provide an overview of past single layer research and investigate some single layer approaches in the preliminary sections of their 2011 report (Huang and Barsky, 2011). In this report, Huang and Barsky also outline two key issues with using a single layer display. The first is that results from single layer approaches all have a dynamic range that is much larger than what can be shown in standard displays and frequently have negative pixel values. When the range is adjusted by scaling pixel values the images lose a significant amount of contrast. The second and more inescapable issue with single layer displays is that the blur from optical errors is a low pass filter which drops certain spatial frequencies. There is no way to alter the pixels of the input image so that these frequencies are not lost in the eye's optics. This makes creating an inverse of the eye's blur a fundamentally ill-posed problem.

3.2 Multi Layer Displays

To cope with the impassable obstacles of correction in the single layer setting, Huang and Barsky introduced a new setup that uses multiple layers of displays and an optimization based approach to setting the pixel values at each layer. (Huang and Barsky, 2011) The frequencies that would be dropped when emitted from a single layer display were dependent on the depth of the display from the viewer. Thus by incorporating multiple layers at slightly different depths, the dropped frequencies of one layer can be compensated for in another. As a result, the new multi layer scheme was able preserve all important frequencies and make inverse blurring into a well posed problem.

The images placed on each layer of the multi layer screen are optimized by the quality of the viewer's perceived image. In this process, the original image is decomposed by its frequency content and repartitioned across the layers. Huang et al. explore different methods of partitioning image content such as evenly across the display layers and a winner take all approach where each spatial frequency is assigned to the single layer that is best at reproducing it in the final perceived image. The final choice was a greedy partition function that approximates the winner take all partition while having a much faster run time. (Huang et al., 2012)

While Huang and Barsky's multi-layer approach represented a tremendous improvement over past single display approaches and created a new well posed setting for future research, the results still had some issues. The biggest was a lack of brightness. Although this could potentially be adjusted for by using display panels with brighter backlights, this added to the difficulty of setting up hardware. In addition, the brightness could fluctuate across different runs of the algorithm which created a flickering artifact when applied to the frames of video. Another desired improvement was even better contrast on top of the improvement shown over past single display results. A final drawback was in the added complexity of the hardware setup. Multiple layers created an undesirable form factor for the display and the depth between the layers had to be tuned to a combination of the eye's prescription and viewing distance.

The findings from Huang and Barsky's 2011 report (Huang and Barsky, 2011) were later published in the work (Huang et al., 2012). In addition to analyzing multi-layer displays and their improvements over single layer displays, this work tied its analysis to light field displays that were gaining popularity and had been investigated in Pamplona et al.'s 2012 work (Pamplona et al., 2012). Huang et al. showed that to make inverse blurring a well formed problem the display had to have the behavior of being multi layered. Light fields were

a promising direction because their control over the directions of emitted light allowed them to create multiple "virtual" display planes. The drawbacks to using a light field display are that in order to create many virtual views the display needs to have a high angular resolution which comes at the cost of reduced spatial resolution. This was discussed in background section 2.7. In addition, light field displays have a finite depth of field that is dependent on their construction. (Huang et al., 2012)

3.3 Light Field Displays

Inspired by Huang and Barsky's multi-layer display and past work with light fields, Pamplona et al. (Pamplona et al., 2012) created a correction algorithm that employed light field displays. Their tailoring process would take as input a map of the eye's aberration and output a light field that would form the correct image on the back of the viewer's eye. The authors demonstrated that the existing success of correcting for aberration in light field cameras (Ren and Pat, 2006) could also be found in the dual problem of correcting the viewer's vision with a light field display. Using the added control of direction, light could be redirected through specific parts of the viewer's eye. This allowed the display to create an in focus image despite having to pass through the possibly asymmetric and varied refraction interface of the eye. In this same vein, the authors also demonstrated the ability to map light around cataracts. (Pamplona et al., 2012).

The tailoring algorithm in this work consisted of two steps. First, a correspondence was created between light field rays and retina positions. Creating the correspondence is essential for light field based correction and has been a constant feature in following works.

Once the correspondence had been created a normalization step was run to assign colors to light field rays. This took into account the fact that more emitted angles would be able to reach the eye from the center region of the display and the fact that multiple display rays would reach the same point on the retina. The authors mention that using a one to one normalization where only one light field ray is allowed for each retina position works. However, it is advantageous to allow multiple display rays to contribute since this leads to a brighter perceived image.

Pamplona et al. tested three different light field designs: a pinhole mask based display, a two LCD layer display where the second LCD acted as a mask and finally a lenticular (lenslet) based display. The dual layer LCD allowed for fine grained control over the pitch and sizing of pinholes in the mask layer at the cost of added construction difficulty. Lenticular based displays were promising because they allowed more light from underlying pixels and created brighter perceived images. This is discussed more in background section 2.5.

Following the previous work of Huang et al. (Huang et al., 2012) and Pamplona et al. (Pamplona et al., 2012), a 2014 paper by Huang et al. (Huang et al., 2014) developed an approach that improved on Pamplona et al.'s light field tailoring. Huang et al. refers to the process of creating the desired light field as prefiltering in accordance with past multi-layer display work in contrast to the phrase "tailoring" from Pamplona et al. Huang et al's 2014 work solves for a correspondence between light field rays and retina positions in a similar manner as Pamplona et al. did. However, instead of performing a normalization step to determine the correct pixel colors, Huang et al. set up a least squares optimization.

The least squares optimization minimizes the error between the projected light field and the ideal image. The 4D light field comprised of the many light field rays is unraveled into a 1-dimensional vector, l^d . The ideal image i is also unrolled into a 1-dimensional vector i . The correspondence between light field rays and retina locations is used to build a projection

matrix P which maps the 4D light field at the display to the projected 2D image on the retina. The optimization then takes the form:

$$\begin{aligned} & \text{minimize } \|i - Pl^d\|_2 \\ & \text{subject to } 0 \leq l_i^d \leq 1, \text{ for } i = 1, \dots, N \end{aligned}$$

Huang et al. use the condition number of the matrix P to investigate the well-posedness of the problem. Lower condition numbers correspond to a more invertible matrix and better solutions. They found that when the angular resolution was increased, the condition number dropped. Exactly how adding display layers helped in the multi layer display setup, having a higher angular resolution means that more views of the display reach the eye. Here each different possible angle of light that can exit a display opening constitutes part of a different view.

The least squares optimization by Huang et al. effectively solves for all light field ray colors simultaneously. In contrast to Pamplona et al’s method, this means that the color is assigned not only based off of that specific ray’s retina position but also with respect to a global error metric of how well that color will fit into the overall ideal image.

The benefit this global color setting is that the portions of the ideal image are better partitioned across the different views of the light field so that a good quality in the perceived image can be achieved with fewer views. Huang et al. estimate that as few as 2-3 views all reaching the eye would be required for computational correction. This means that Huang et al’s algorithm can operate with displays with lower angular resolution and therefore higher spatial resolution. The improvement is visible in both simulations and photography experiments comparing the two methods (figure 13)

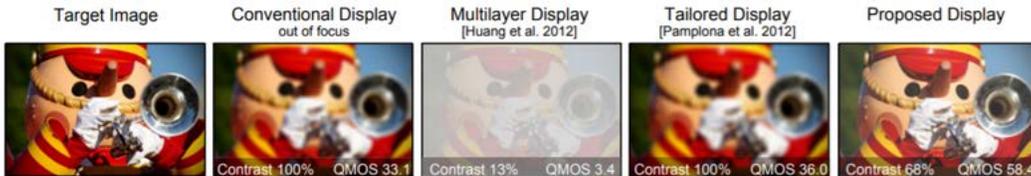


Figure 13: A comparison of simulated results of past work taken from Huang et al.’s 2014 paper (Huang et al., 2014). Left to right the images are: the target image that is the ideal to strive for, the conventional display showing the blur that would be experienced viewing a normal display without correction, the multilayer display from Huang et al. (Huang et al., 2012) which is crisp but suffers from lost contrast, the tailored display from Pamplona et al. (Pamplona et al., 2012) and finally the result from Huang et al’s 2014 paper. The new light field tailoring algorithm introduced in the 2014 paper showed noticeable improvement over Pamplona et al.s 2012 light field results.

Zehao Wu built off of Huang et al’s work in his 2016 masters thesis. One important contribution was modifications to the method for finding the light field retina position correspondence. Wu’s change to the correspondence solving step was to swap the direction that rays are traced. Previously in Huang et al’s algorithm, ray tracing started at positions on the retina. Rays were drawn from those initial positions randomly sampled as passing through the eye aperture and were then traced backwards to the display screen. The word

backwards is used because this is the reverse of physical reality where light is emitted from the display. Rays which were able to pass through the display mask had to represent light field rays that could escape the light field display. Thus, the correspondence was created by assigning the original retina position to those rays that passed through the display mask.

While accurate, the backwards method of solving the correspondence was not computationally efficient. Many of the rays traced backwards from the retina were wasted samples as they would not be able to pass the display mask and did not represent emitted light field rays. In Wu's new correspondence calculations, coined the forward method, rays were traced forwards from the display to the retina. Ray samples were generated as those passing through the display mask and as such were guaranteed to represent parts of the light field display's emitted light field. A diagram of the calculations for a single ray sample is shown below in figure 14. Equation 1 shows the math for finding the locations of various points such as the ray retina intersection p_R . (Wu, 2016)

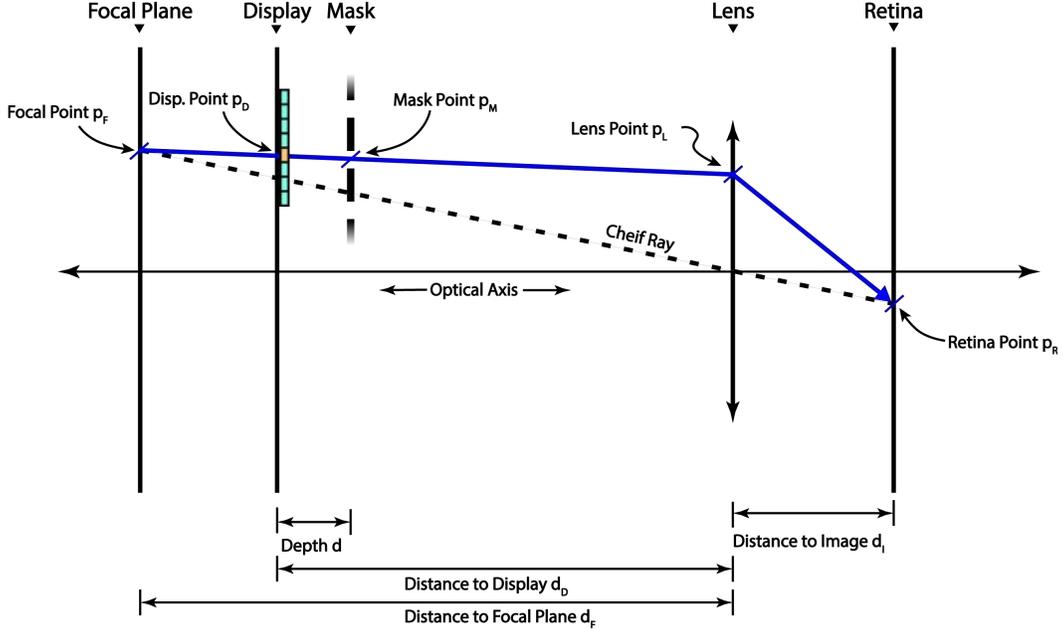


Figure 14: An illustration of a correspondence calculation from Z. Wu's forward method. Like other light field tailoring algorithms, each ray direction from each pixel needs to be tied to a location on the retina. In this case, the ray is traced back to a point on the focal plane, p_F before finally being traced to the retina landing point p_R .

$$p_F = \frac{p_M - p_D}{d} (d_D - d_F) + p_D \quad (1)$$

$$p_L = p_f + \frac{p_M - p_D}{d} d_F \quad (2)$$

$$p_R = -\frac{p_F}{d_F} d_I \quad (3)$$

Other notable research in vision correcting displays has been conducted in recent years by Charles Ding (Ding, 2017), Yirong Zhen (Zhen, 2019) and Luxin Yang (Yang, 2020) in their masters theses. Zhen analyzed artifacts in generated results through testing the one to one assumption. This is detailed in the next section 3.4. Ding investigated improving runtime performance and quality of the correspondence solver, improvements to the software simulation of the user’s perception and finally the effects of different pixel arrangements and the effects of diffraction. Studying these details are important to making sure vision correcting displays work well in practice. Luxin Yang did work optimizing the mask design for light field displays and investigating applications of vision correcting displays in near eye displays such as VR headsets. The VR setting is a very promising direction for future work because the algorithm is easier to run with the continual fixed position of the display. In addition, there is a need for correcting farsightedness in headsets due to the inconvenience of wearing glasses underneath. This research topic is touched on again in the future work section 7.3.

3.4 Error Proofing

Light field tailoring approaches to vision correction are reliant on knowing the correct position of the display in order to make accurate calculations of how light will enter the eye. If there is an error in the measurement of the display’s location or if the display changes position after the calculation has run there will be an expected drop in perceived image quality. Given a fast enough refresh rate and a very accurate estimate for eye location this error could be fixed by recomputing each time the viewer moves. However, adding the requirement of a high frame rate complicates the approach and estimates on pose will not always be perfect. Thus it is paramount that the correction algorithm tolerate some level of change in display position.

Error proofing in vision correcting displays was addressed before in F. Huang’s PhD thesis (Huang, 2013) and in a following publication (Huang et al., 2014). They tested errors in display position in two cases. First, when the display was at a closer or farther depth than the one assumed and second, when the display was at the correct depth but had been shifted laterally off of the optical axis. In both cases, there was a clear falloff in image quality as the tested display was moved.

To solve this problem, F. Huang et al. took advantage of the least squares optimization that was already included in their approach. Starting with the initial optimization problem set up for the correct position, they appended the corresponding optimization problems for neighboring shifts. Conglomerating multiple optimizations had the effect of over constraining the problem and as a result forced the solver to find solutions minimizing the error in all the included positions. While there is a slight drop in the quality of the original position, quality over the neighboring shifts is maintained. A downside is that the larger size of the optimization increases the run time of the algorithm.

In the future, error proofing will be an essential part of the 3d adaptation. Adding in changes in roll, pitch and yaw greatly increases the dimension of possible changes in display position making the optimization demanding computationally. In this work we did not include a method of error proofing hoping instead to implement the vision correcting algorithm in a fixed setting as a first step in the 3D realm.

4 Approach

4.1 Implementation

The starting place for our implementation is the process pioneered by Pamplona et al. and Huang et al. of creating a correspondence between light field rays and retina positions before setting the colors of the light field rays to optimize the eye’s perceived image. To calculate the correspondence we made a 3d version of Z. Wu’s forward direction approach. (Wu, 2016) For assigning colors we use a technique, also from Huang et al. and Pamplona et al’s works, where we pretend that the ideal image is placed on the retina. With the correspondence we can map light field rays to retina positions and then pull colors from the ideal image at that those retina positions. After assigning the pulled colors to our ray sample, we average all the rays originating from each pixel to decide its final color.

Forward method ray tracing in the correspondence step is separated into propagation between four optical planes: the display screen, the mask, the lens of the eye and the back of the eye. The 3d setup is shown in figure 15 illustrating how a retina landing point could be calculated from a ray sample.

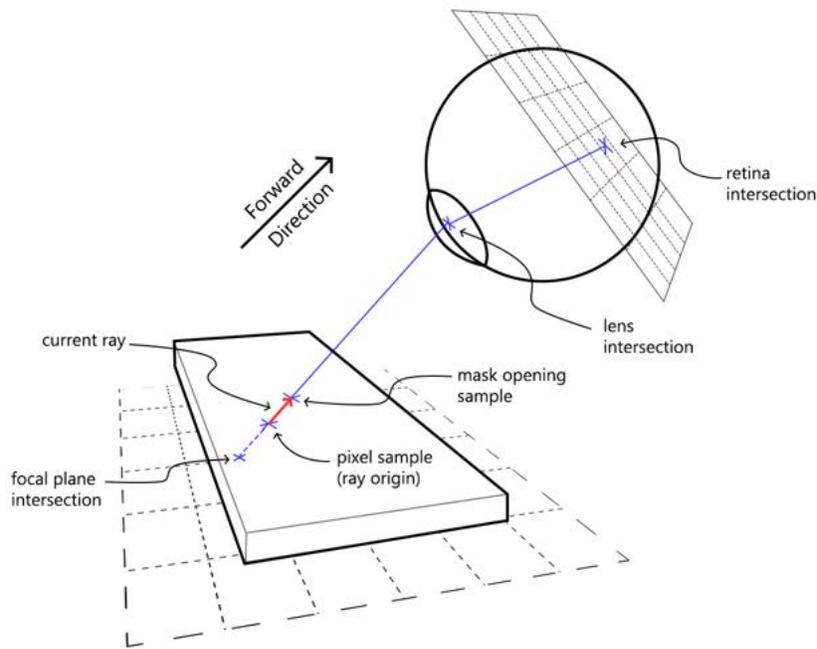


Figure 15: The forward ray trace in 3D.

The end result after assigning back colors is the precomputed image. Figure 16 shows an example of the original input image and the precomputed image generated by the correction algorithm. The disks that comprise the precomputed image each represent an above pinhole and each disk contains a variety of colors for the different directions that light can escape the pinhole.



Figure 16: Source image courtesy of Julia Sakelli (pexels). The original image is on the left and the result of running the precomputation algorithm is shown on the right.

4.2 Code

We coded our implementation for both the precomputation and our simulation in C++ for performance in the many ray calculations that have to be run. The simulation to validate our results is discussed more in section 4.10. While the run time for the precomputation step is dependent on the number of pixels in the input image and the sample rate, decent results can be calculated at lower sample rates (16-32 samples) with average run times between roughly 20 and 40 seconds. Running at a very high sample rate (100-200 samples) drops the run time to as slow as 200 to 300 seconds. This is discussed more in section 4.6. The algorithm eventually needs to run in real time in order to keep up with the frame rates of video and also update to changes in viewer position. Improvements could be made in run time by parallelizing the per pixel calculations. In addition, configurations in display mask parameters that result in fewer samples being blocked by apertures require fewer samples overall.

4.3 Optics Calculations

The first change we had to make was changing 2d optics calculations into 3d ray tracing through ray plane intersections. In 2d the planes of the display, the eye's lens and the retina are all parallel so that ray intersections only depend on the ray direction and the distance between the planes. Intersections could be calculated in a similar manner to the equations in Z Wu's forward method shown in figure 14 and equation 1. In 3D, the intersection also needs to incorporate the angle of the plane. To accomplish this we switched to ray plane intersections commonly used in computer graphics. An example is shown in figure 18

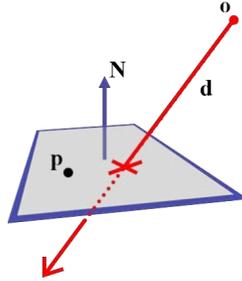


Figure 17: An example ray plane intersection. The plane’s orientation is represented by the 3d normal vector $\vec{N} = (x_N, y_N, z_N)$. The ray is given by origin $o = (x_o, y_o, z_o)$ and direction $\vec{d} = (x_d, y_d, z_d)$. The calculation also requires a point located somewhere on the plane, p . The intersection point can be calculated as $o + (\frac{(p-o) \cdot \vec{N}}{\vec{d} \cdot \vec{N}}) \vec{d}$ (Marschner and Shirley, 2018)

4.4 Coordinate System

It is easiest to have the optical axis aligned perpendicularly in many situations throughout the precomputation algorithm. For example, it is best to have the x, y axes parallel to the display with the z axis perpendicular when finding the nearest pinhole mask opening. At the same time, it is hard to calculate refraction at the eye lens if the z axis is not parallel with the lens’ optical axis.

To make calculations easier with both of these two contradictory requirements, we introduced a second coordinate system and transformed vectors and points between the two as ray tracing calculations moved between them. Figure 18 shows the two coordinate systems we employed. On the left the display coordinates represent the angle of the display and serve to make optics calculations at the display easier. On the right, the eye’s coordinate system makes calculating refraction at the lens easier. The i,j axes show how images on the display and retina can be indexed. This is flipped on the retina side for the convenience of creating results that are right side up when we simulate the image that the eye would perceive.

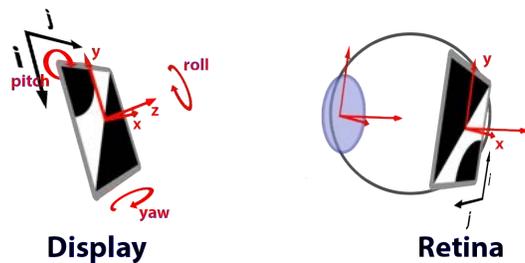


Figure 18: An illustration of the two coordinate systems introduced.

4.5 Retina

In previous approaches, the retina plane was parallel to the display so that light would always fall within a rectangle. This guarantee meant that the retina could be modelled as a

regular rectangular image and that positions could be directly converted into a pixel index before sampling a color.

In the 3D case the projected shape on the retina could be any quadrilateral. This meant that the mapping between the pixels of the ideal image and what should appear on the back of the eye was more complicated. To account for this in the 3D implementation, the tailoring step first calculates a homography, a transformation commonly applied to calculate perspective warps. (Hartley and Zisserman, 2003) This creates a map between indexes in the ideal image and positions within the quadrilateral of where the image would appear on the retina. Solving for this transformation requires a correspondence of eight 2d points. We used the four corners of the display pixels: $(0, 0)$, $(width, 0)$, $(width, height)$, $(0, height)$ and the four corners on the retina found by tracing orthogonal rays from the physical location of the display corners. We used OpenCV's solver for the computation. (Bradski, 2000). This process is illustrated for two different viewing conditions in figure 19.

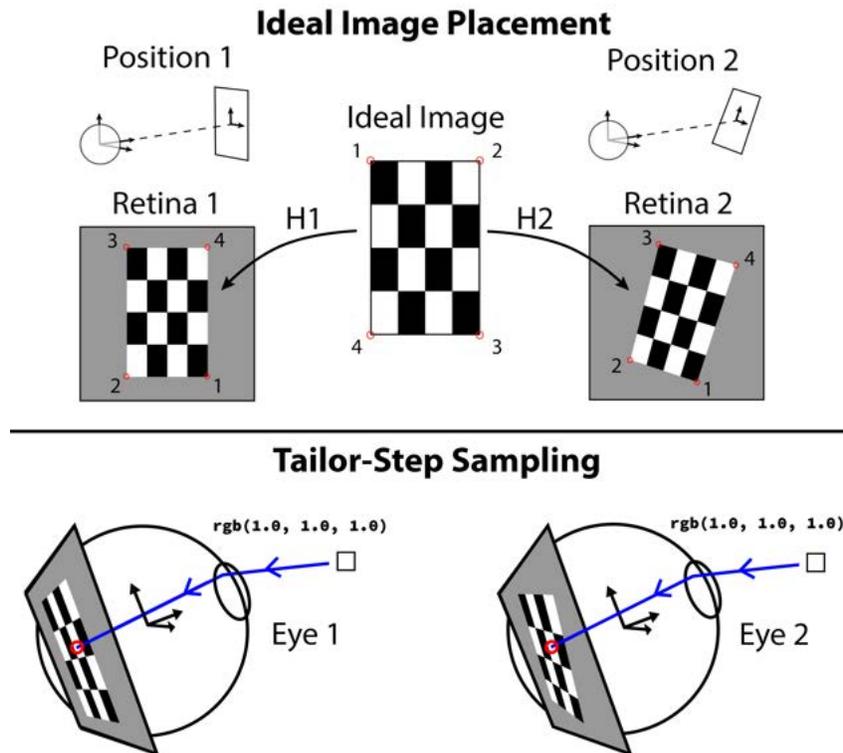


Figure 19: Two examples of placing the ideal image on the retina (shown in top) before the light field precomputation step (shown in bottom). The two perspective warp transforms (homographies) used to warp the ideal image are $H1$ and $H2$. The corners passed into the solver for these transformations numbered 1-4. In the bottom, a single ray sample is being traced from the display (not shown) to the back of the eye.

The computed homography gives a map from the ideal image pixels to physical locations on the retina and the inverse homography gives a map from retina locations to the pixels of the ideal image. Finding the correct color for a given ray sample was done by applying

the inverse homography its retina intersection then pulling a color from the ideal image at that index mapped to by the homography. An added advantage of using the transformation is that the resolution of the ideal image is separate from the size of the display’s projected area. Changes in viewing distance and angle can be accounted for by a quick updating computation of the four projected corners and the homography.

4.6 Sampling in precomputation

The ray samples used to determine the value of any given precomputed image pixel have to give an accurate estimate of how light would escape the display mask. First, the correct opening in the mask is calculated as the nearest opening in the mask to the line between the origin pixel’s location and the center of the eye. In this way, as the display changes angle, the nearest opening will correctly be reflected as the one that is most likely to direct light towards the eye. After this mask opening has been found, ray samples are drawn in between the underlying pixel and the circular opening.

An estimate could be created taking a single ray sample for each pixel and opening pair. The color of the ideal image at where the ray landed would be an indicator of the best color for that opening and ray direction since we know roughly that the light field ray generated from that pixel will land at the calculated retina location. This estimate can be improved by taking multiple samples for each pixel. We can randomly choose sample ray origins within the square of the pixel and random sample ray directions passing through the opening in the above display mask. In reality, the light from a pixel exiting an opening could land within a small area of the retina rather than a single point. Taking multiple samples inside the pixel and the opening leads to a better approximation of the colors that would result in the warped ideal image on the retina. This technique, derived from past approaches in 2D by Z. Wu (Wu, 2016), C. Ding (Ding, 2017), L. Yang (Yang, 2020) and Y. Zhen (Zhen, 2019), is illustrated in figure 20

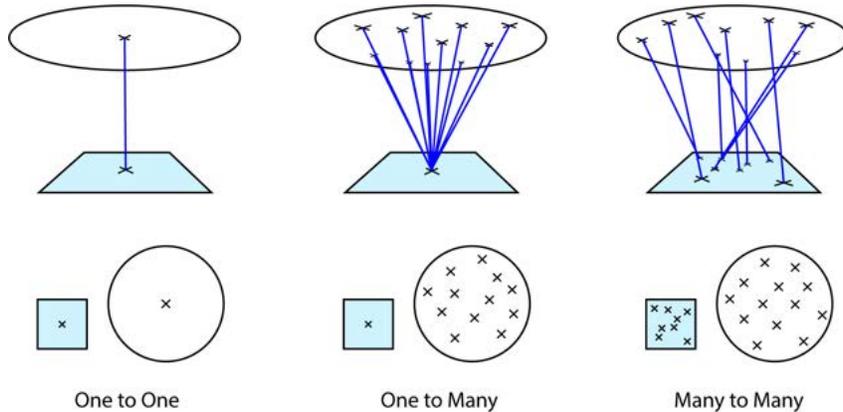


Figure 20: An illustration of sampling schemes for the precomputation algorithm. Rays are shown in blue traced from points on the pixel below through the above mask opening. The three designations described by Y.Zhen are "one to one", to "one to many" and "many to many" are optional changes to the number of ray origins and ray directions taken respectively. (Zhen, 2019)

The most accurate estimate is given by "many to many" however, the other options require fewer samples. The results in this project were all created using the "many to many" scheme.

Figure 21 shows the improvement in perceived images that comes with increased sample rates. (These images were generated using our simulation which is discussed in the next section.) There is a diminishing return after a certain number of samples which is shown in the similarity between the 256 sample rate image and the 1024 sample rate image. However, the first increases in the sample rate have a clear effect on the quality of the results shown in the improvement between the 4, 16, and 36 sample rate images. In each of these examples the number of samples per display pixel and the number of samples per pinhole were equal. The number listed for each example is the samples per display pixel times the sample per pinhole. For example 4 samples indicates 2 samples per display pixel and 2 samples per pinhole.



Figure 21: The corrected simulation at various sample rates and the uncorrected simulation on the far right.

The number of samples required to reach a good estimate in the precomputation step depends on both the configuration of the hardware mask as well as the orientation of the display. In situations where rays are more scattered or have high chances of being blocked by the eye's pupil, it takes a greater number of samples for the estimate to converge. One case is when pinhole diameters are larger so that the light exiting is spread across a greater range of angles. Another is when the pupil's diameter is smaller so that fewer rays may enter.

4.7 Picking Mask Hardware Parameters

The settings for pinhole mask angular resolution, mask offset and pinhole diameter all effect the quality of results generated by the correction algorithm. A series of experiments were conducted in order to determine the best values to use. In this section, the effect of changing these settings is discussed as well as the final motivation for choosing certain values.

The first parameter to consider is the angular resolution of the mask to address the angular vs. spatial resolution trade off discussed earlier in section 2.7.

If the angular resolution of the mask is changed without also making changes to the underlying display layer there is a drop in the spatial resolution of the mask. Figure 22 illustrates this for a 3x3, a 5x5 and a 7x7 pinhole mask. While the underlying display keeps the same 150x194pixel resolution in all three images, the number of pinholes to cover the same display screen drops from a grid of 50x64 pinholes for the 3x3 mask, to a grid of 30x38 pinholes for the 5x5 mask and finally to a grid of 21x27 for the 7x7 mask. The low pixel resolution of 150x194 was used to better illustrate the drop in quality with lower

pinhole spatial resolution. While light from neighboring pinhole apertures is blurred together somewhat hiding the effect of low resolution, the results created with 5x5 and 7x7 pinhole masks are noticeably more distorted and less clear compared to the results created with a 3x3 mask.

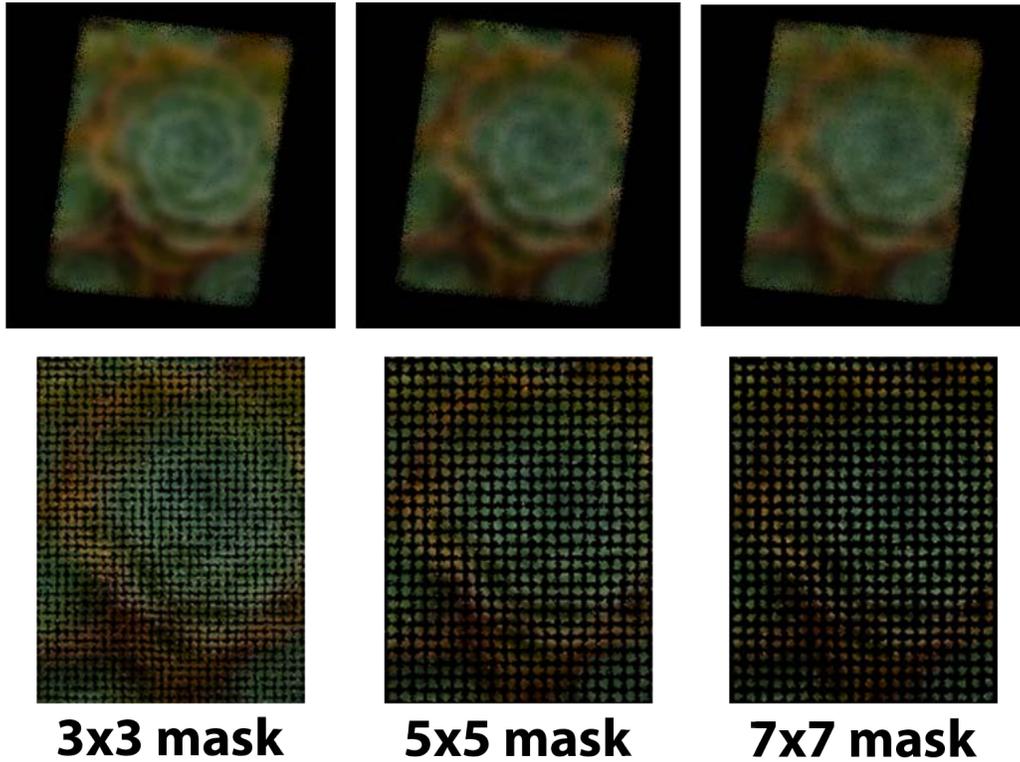


Figure 22: A comparison between using a 3x3 mask, a 5x5 mask or a 7x7 mask to cover the same underlying 150x194 pixel display layer.

The effect is similar to how lower resolution images become more pixelated except that the pinholes openings blur together instead of having blocky pixel edges. To maintain the same brightness across the varying angular resolutions the results were generated with an increasing pinhole diameter setting: 0.11mm for the 3x3 mask, 0.18mm for the 5x5 and 0.256 for the 7x7 mask. The effects to consider when choosing pinhole diameters are discussed later in on in this section.

As angular resolution increases, the spatial resolution, given by the number of pinhole openings across the mask, drops. The resolution of the underlying display can be increased to compensate for the loss in spatial resolution. In figure 23, as angular resolution is increased, the underlying display resolution is increased as well while the pixel length is decreased.

While the precomputed images shown in figure 23 each have a different size in pixels, they have been re scaled for comparison. The effect of decreasing pixel length while pixel dimensions were increased is that the precomputed images still represent the same physical area on the display. Additionally, each of the precomputed images work with the same number of pinholes, yet each pinhole has a different number of pixels beneath it.

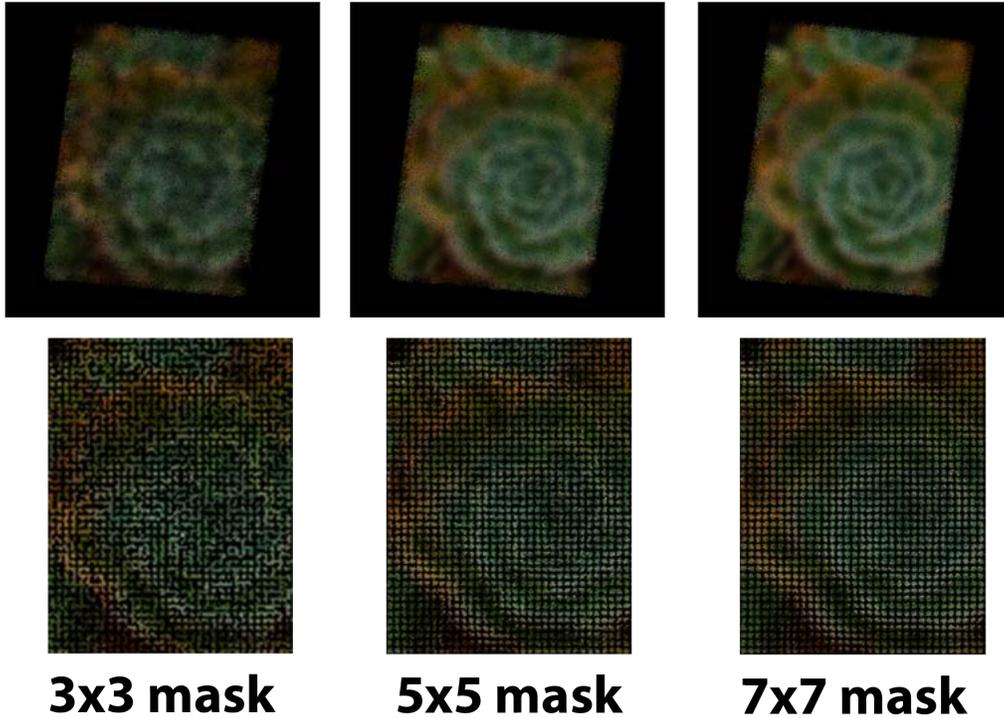


Figure 23: A comparison between a 3x3, a 5x5 and a 7x7 pinhole mask while increasing the resolution of the bottom display layer to compensate. The 3x3 mask result was run with a 105x135 pixel and 180 PPI bottom layer, the 5x5 mask result was run with a 175x225 pixel and 300 PPI bottom layer and the 7x7 mask result was run with a 245x315 pixel and 420 PPI bottom layer. All displays had the same dimensions of 15mmx19mm.

The lower angular resolutions in figure 23, have dark patches where no light directions were able to reach the eye. At the higher angular densities, there is enough angular resolution for light to reach the eye from most pinholes. These results have been exaggerated to show how having a higher resolution bottom display layer can compensate to maintain the same effective resolution while achieving higher angular resolutions.

Display manufacturers don't have much incentive to create high resolutions (> 326PPI) since the human eye viewing normal displays can't distinguish past this point. However, there is potential for high resolution displays to help create high angular resolution light field displays if the technology becomes more popular.

That being said, in our results, not having a high enough angular resolution was rarely the problem. Huang et. al. had a similar finding that a high quality correction only requires a few views of the light field (lower angular resolutions). The artifacts in the lower angular resolution displays of figure 23 are so egregious partly because it was run with a small display, placed far away and the virtual eye had a high defocus error. In our results, a the best correction was most often created using 3x3 and 5x5 masks.

4.8 Pinhole Diameter

To choose the pinhole diameter of our mask, we considered a balance between creating dimer images and the increased sharpness that comes with dropping the pinhole diameter. A large enough pinhole diameter is required to keep the perceived images from being too dark but having it too large will create a blur effect on the perceived image. An additional consideration is that a 5x5 pinhole mask can tolerate a larger diameter than a 3x3 pinhole mask (assuming the same bottom layer) because there is a larger area underneath each individual pinhole in a 5x5 mask. This tradeoff is shown in figure 24 for a 5x5 pinhole mask. Here images have been generated within the range of diameters most commonly used for this report. To help with scale, for a 326 PPI display each pixel has a width of around 0.0779mm so the below pinhole diameters are on the order of 1-2 pixels.

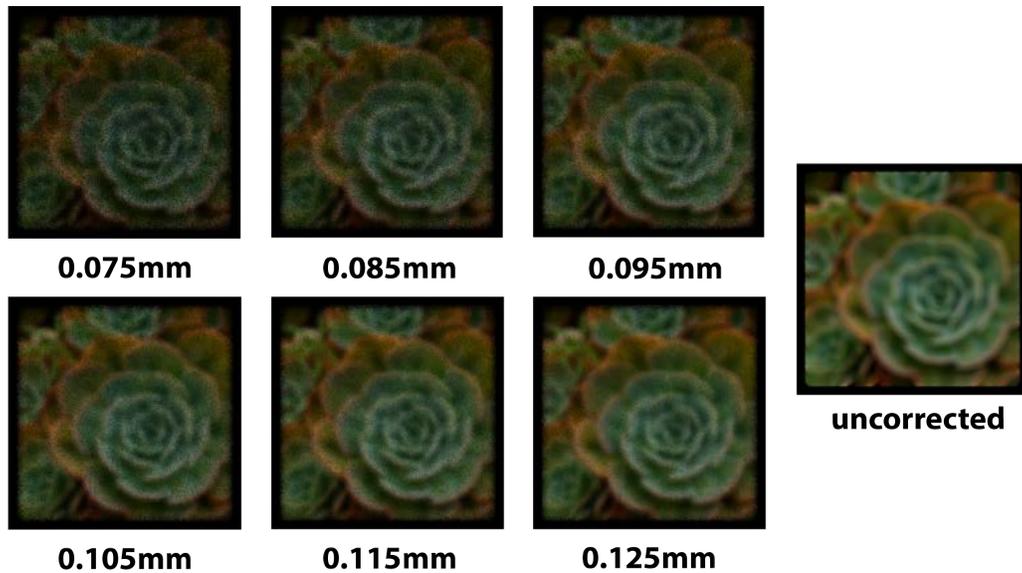


Figure 24: Tests using varying pinhole diameters. For this experiment, the display mask was a 5x5 pinhole mask with a 6mm offset. The uncorrected simulation is on the right for comparison.

4.9 Mask Offset

Another parameter that was considered to improve our corrected results was the offset of the pinhole mask from the underlying display screen. Larger mask offsets such as 10-14mm give a better correction than smaller offset such as 6-8mm. The benefit is related to the fact that sharper images come from using smaller pinholes. By increasing the mask offset, it is possible for light to reach the retina from a specific pinhole from each of the pixels in the underlying grid of 9 pixels in a 3x3 mask or 25 pixels in a 5x5 mask. Still, smaller offsets are preferable in practice because any offset creates an undesirable form factor for the display. Figure 25 shows the effect of different mask offsets on the corrected result.

Keeping both the ideal form factor of a smaller offset while also maintaining quality in corrections can be accomplished by starting with a small offset and tuning the angular

resolution of the display in conjunction with pinhole diameter. Optimal results will occur when the patch of pixels visible by the eye below each pinhole is large enough to cover the whole 3x3 or 5x5 grid. Unfortunately, as discussed in (Pamplona et al., 2012), this optimal setting is dependent on the distance from the viewers eye to the display. For faraway displays it is better to have a large offset with small pinholes because light will spread more as it travels to the eye. For near displays, it is better to use larger pinholes and a smaller offset so that the result isn't darkened and so that all emitted angles are able to reach the retina.

While the true optimal settings are dependent on the viewing distance, in practice the loss of picking a middle value for both the display offset and the pinhole diameter is acceptable. It is important to pick these two values with the angular resolution of the mask in mind since the area of bottom layer pixels that should be visible through each pinhole changes with this configuration. For 3x3s mask we most often used a 0.075mm pinhole diameter and a 6mm offset. For 5x5 masks we increased the pinhole diameter to around 0.1mm and used a larger 8-12mm offset.

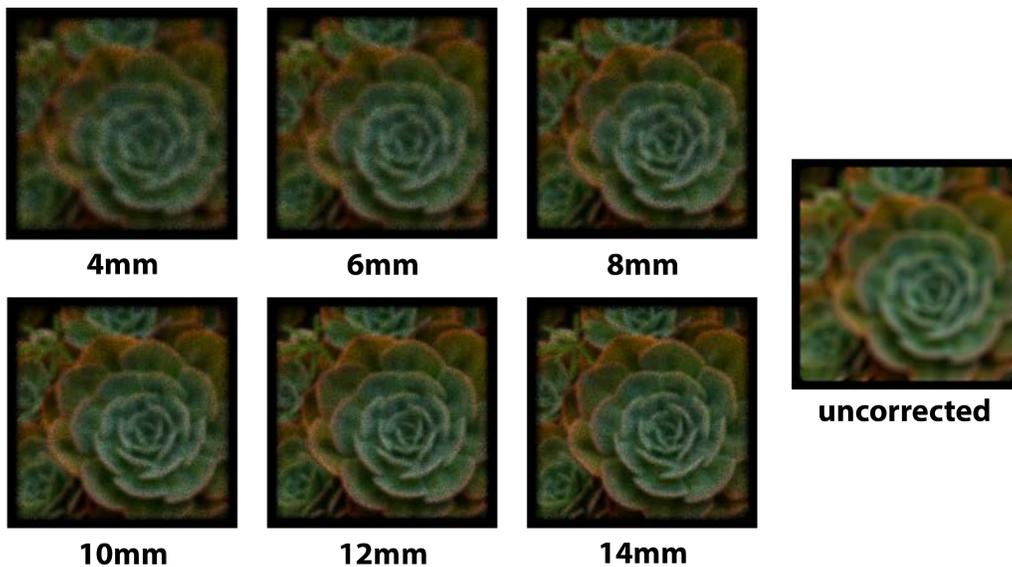


Figure 25: Tests using varying mask offsets. For this experiment the display mask was a 5x5 pinhole mask with 0.105mm diameter pinholes.

Poorly chosen mask parameters can allow the user's perceived image to remain blurred or become severely darkened. However, these parameters can be chosen to work well in a variety of different settings. Extra performance can be gained if it is known that the display is going to work at a very specific viewing distance. A final detail is that knowing how bright the viewer's environment is can give some indication of how dilated their pupils will be. For example, when outdoors viewers will have smaller pupil diameters of 2 to 4mm. Using this information, display brightness can be optimized with larger pinhole diameters in bright viewing environments or decreased diameters in low light environments.

4.10 Simulation

In place of a physical experiment, the performance of our light field tailoring algorithm was tested through simulation. Using simulation has been common throughout past works in vision correcting displays because it is much faster to perform experiments, it is much cheaper and can model optical effects that would require special lenses to test in the physical world. (Wu, 2016) Following this approach, we made our own version of a simulation that would function in the presence of display angles. In similar fashion to past works, we took a ray tracing approach, treating the eye as a single lens camera and creating an estimate of the light that would enter to create an image.

The simulation creates an estimate of the light that would land on a users retina. This estimate of the perceived image is generated pixel by pixel. For each pixel on our virtual retina, rays are cast into the scene passing through the aperture of the eye. After calculating the refraction through the lens and the intersection with the mask, the ray is checked to see if it passes through a pinhole. If so, the color on the display where the ray lands is added to an ongoing average for the color of the retina pixel being calculated. This procedure gives an estimate of the light that could reach the current retina position. Figure 26 shows a diagram of the simulation ray tracing approach.

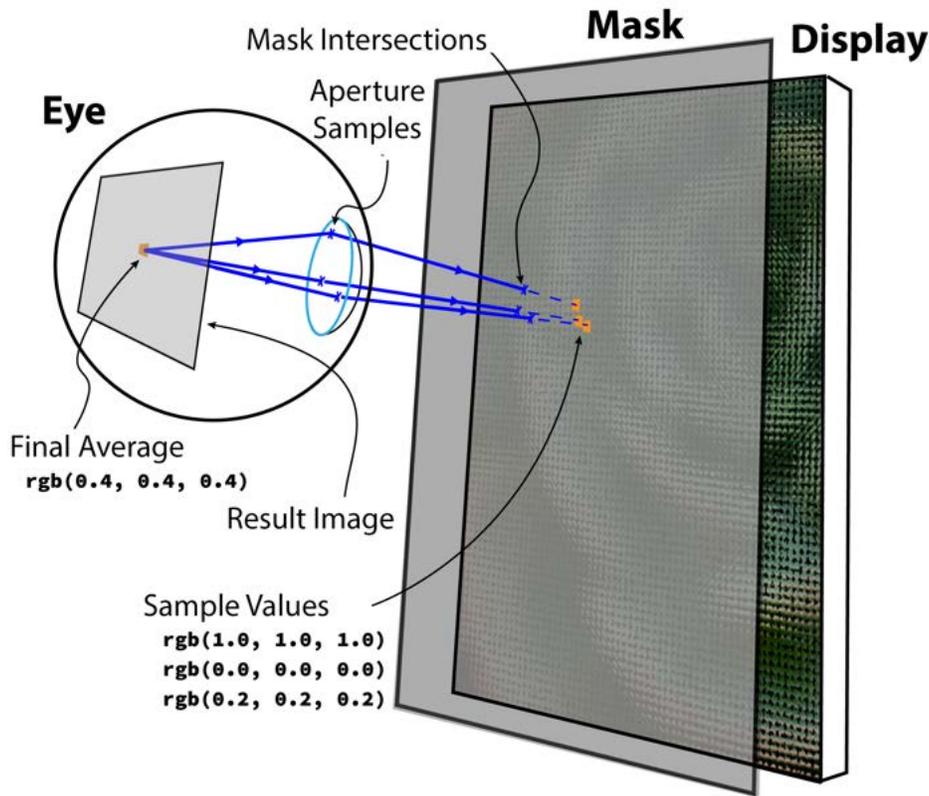


Figure 26: A diagram of the simulation algorithm.

Rays are traced from the eye to the display despite the physical reality that light would

travel in the opposite direction. The flipped tracing direction of eye to display allows easy generation of sample rays that are important to the estimate, namely light paths with the potential to reach the retina.

Many of the calculations are shared between parts of the precomputation step and the simulation. Both require tracing rays through the planes of the retina, lens, display and focal plane. The main differences are in the direction that the rays are traced and the sampling schemes used. While it is important that the run time of the tailoring step is low for real time performance requiring the number of samples per display pixel to also be kept low, a much higher sample rate is used in the simulation to more accurately approximate the image that would be perceived by the eye.

The unfortunate side effect of using a pinhole based light field display is that it can create issues with the brightness of perceived images being too low. This was discussed earlier in section 2.5. Without any added fix, the resulting simulated images are quite dark as shown in the middle left image of figure 27. One way to make the results visible is to boost the brightness in post processing. This is shown in the middle right image. Another fix is to weight rays that made it through the mask higher in the final average. This fix shown in the far right image is similar to boosting the brightness of the simulated display. The third fix we explored is to only count rays that made it through the mask in the final average. This is shown in the far left of figure 27. While this third method is giving a biased estimate of the light in the scene, it is far less noisy than the results from the other two fixes. For this reason, we chose to use this method in our results.

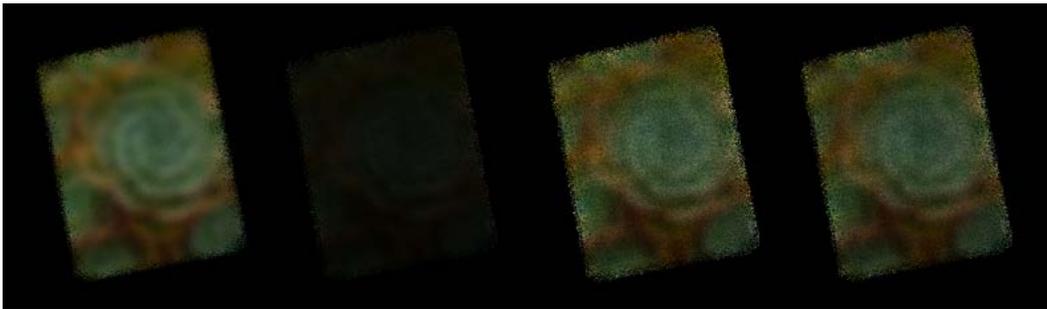


Figure 27: A comparison of different ways of account for display brightness. The far left image is generated using only rays that are able to reach the bottom layer of the display. The middle left is generated including an equal weighted black from rays that were blocked by the display mask as well as the colors from rays that passed through. This image is much darker as a result. The middle right is the darker image (middle left) after its brightness has been increased in post processing. The far right is the result of including the black from blocked rays but counteracting the darkening effect by giving a higher weight (4x) to the contribution of rays that passed the mask and found a color.

The simulation runs in a virtual scene with the display placed before a single eye. The center of the scene's coordinate system is set at the eye's lens with the eye's optical axis as the $-z$ axis. The display is parameterized by an (x,y,z) offset from the center with its angle described by a pitch, yaw and a roll. The display was modeled as two thin planes, one representing the pixels of the display and one serving as the mask. The retina was specified as a square of a given area (usually 1.5-4mm on each side) with a retina pixel resolution

usually set higher than that of the display. Reducing the retina area gave the effect of a "zoom" in the resulting simulated image and was used to maximize the display's footprint in the results.

4.11 Pupil

Increasing the eye's pupil size has the effect of exacerbating the refractive errors that are present. This is true of both higher order aberrations discussed in section 2.3 and the lower order aberrations that this work examined. Human pupils on average have a diameter of 2-4mm in bright light and can dilate to roughly 8mm in dark settings. For most of our trials, the pupil diameter was set between 4mm and 6mm. The effect of increasing pupil diameter is shown in figure 28.



Figure 28: A comparison of different pupil diameters starting at 3mm on the left and increasing to 8mm on the right. In each, the eye has the same defocus error of 130mm created by setting the focal plane at 380mm away while the display is 250mm away.

4.12 Mask Effects

When defocus error is close to zero so that the distance to the focal plane is close to the distance to the display, the pinholes of the mask can be readily seen in the prefiltered image. This artifact occurs because the display is almost in focus so that light exiting from pinholes doesn't spread as far across the retina. Figure 29 shows the artifact.



Figure 29: Various corrected simulations showcasing artifacts of the pinhole mask. The display in the far left image is angled so that the pinholes show up in the strip of the display is in the plane of eye's focal distance. The middle left is the view of an eye with no defocus error. The middle right image was created with a defocus error of 10mm and the far right was the result of a -10mm defocus error.

5 Results

5.1 Experiments

The first result section, section ?? showcases a double image artifact that was common in early results. The next section 5.3 shows the results of simulating defocus error without any display angle. Section 5.4 shows the correction running in the presence of display angles, showcasing the ability to run in full 3D. Section 5.5 shows the correction at varying levels of refractive error. Finally, section 5.6 shows the rate at which perceived image quality drops when the display angle is moved from where the precomputation step assumed.

Experiments involved simulating two perceived images. The first is an estimate of what the eye would see viewing the precomputed image behind the display mask. This is the corrected simulation. Results were generated using the mask parameter choices discussed earlier in section 4.7.

To provide a baseline to compare the corrected simulation against, the uncorrected simulation is generated as an estimate of what the eye would see viewing the original image on a normal unaltered display.

In this way, the corrected and uncorrected images have the same retina footprint and have an identical treatment from the simulation. Figure 30 shows a flowchart of creating the two simulated images.

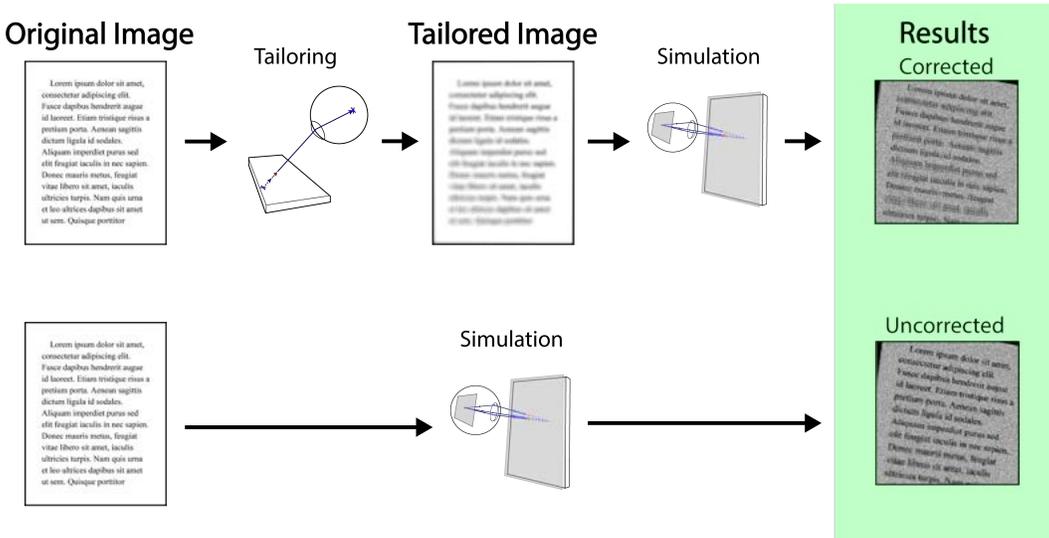


Figure 30: A flowchart showing the light field tailoring algorithm and corrected simulated result on the top and the baseline uncorrected simulation on the bottom. For the corrected result, the simulation is run with the pinhole mask above the pixel display layer and for the uncorrected only the pixel display layer is used.

5.2 The Double Image Artifact and Adding an Aperture

When either the pupil diameter or pinhole diameters are too large, a double image artifact appears in the results due to multiple conflicting views of the display reaching the retina. This effect is shown in figure 31.



Figure 31: Images showcasing the double image artifact. In the text image every other line appears twice. The checkerboard image shows how the artifact is periodic and occurs both vertically and horizontally through the image.

After experimenting with different hardware parameters and display setups, the solution was found to be using tighter pupil and pinhole apertures. A side effect of adding a pupil aperture is that the perceived image is darker and has a vignette around the borders at smaller apertures. Reducing pinhole diameters has the effect of also darkening the result and adds more noise because more light is blocked.

5.3 Perpendicular View

To compare against existing 2D implementations, we ran our experiment with the display perfectly perpendicular to the eye's optical axis. This result was generated using a pinhole mask with 0.075mm diameter pinholes, pinholes spaced every 5 pixels and a 10mm mask offset. The eye's focal point is 380mm while the display is 250mm away creating a 130mm defocus error. The simulation sample rate was 512 samples per retina pixel and the prefiltering sample rate was 18 samples per display pixel and 18 samples per pinhole.

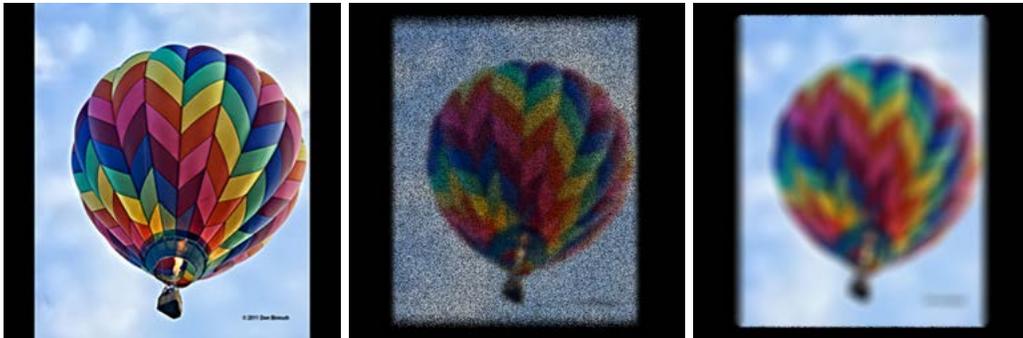


Figure 32: On the left is the original ideal image, the center is the corrected simulation and the right is the uncorrected simulation.

The above result was run with very high sample rates for the precomputation step and the simulation in order to provide a high fidelity result to compare with past works.

5.4 Angled Display

Figures 33, 34 and 35 show the algorithm running with the display at varying levels of roll, pitch and yaw. In each figure, the top row is the corrected result to be compared against the uncorrected result in the bottom row. The tests shown below were run at varying display angles with the following configuration: The display mask was 6mm off of the display screen with 3x3 pinhole density and a 0.1mm pinhole diameter. The display was 200mm away from the eye which had a defocus error of 80mm and a pupil diameter of 4mm. The simulation sample rate was 256 samples per pixel and the precomputation sample rates were 8 samples per per computed pixel and 8 samples per mask opening.

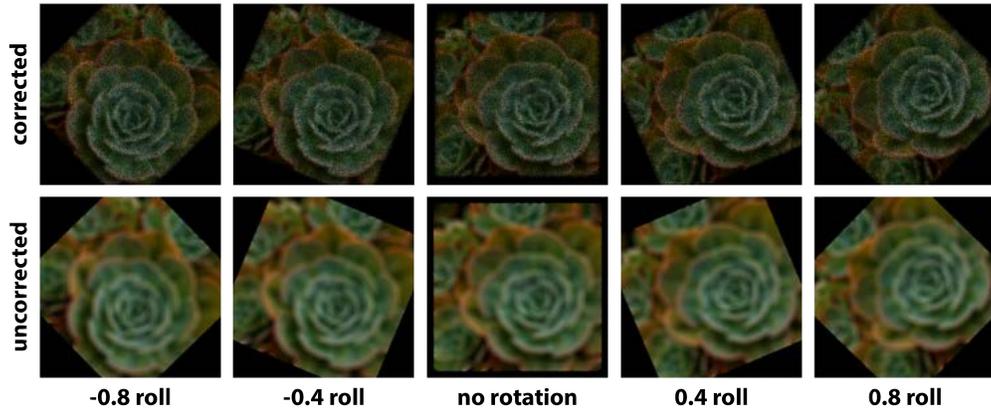


Figure 33: Tests using varying display roll

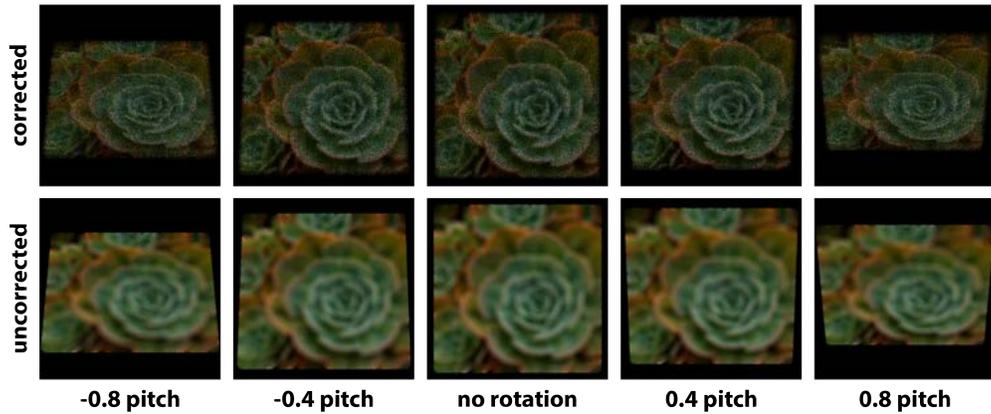


Figure 34: Tests using varying display pitch

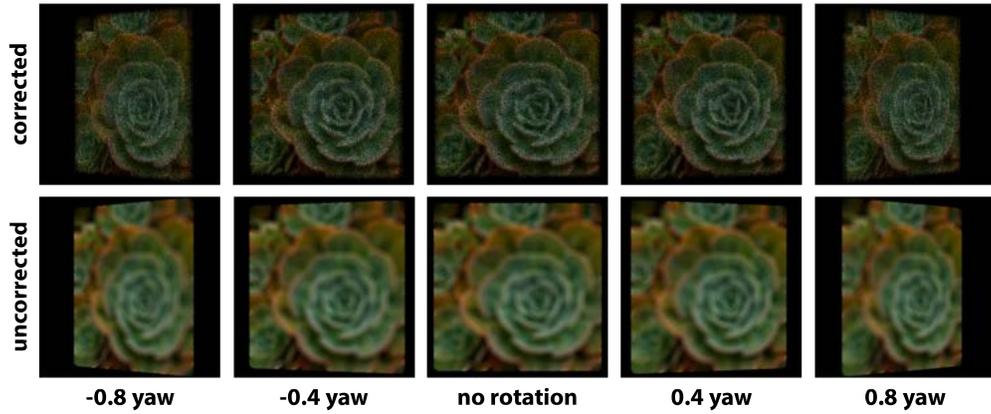


Figure 35: Tests using varying display yaw

5.5 Varying Levels of Refractive Error

The tests below in figure 36 were generated with varying levels of defocus error created in the manner described in background section 2.3. In each pair of images the top is the corrected simulation and the bottom is the simulation of a regular display. The configuration used to generate these results was a 150x194 pixel display with a density of 326ppi. For corrected results, a pinhole display mask was used with an angular resolution of 3 offset a distance of 6mm from the display screen. The eye's pupil diameter was set at 6mm. The display was set a distance of 200mm away from the eye with a roll of 0.2 radians and a yaw of -0.3 radians. The error is represented as the difference (eye focal distance - display distance). Positive levels of defocus error correspond to a Myopic eye and negative levels correspond to a Hyperopic or Presbyopic eye.

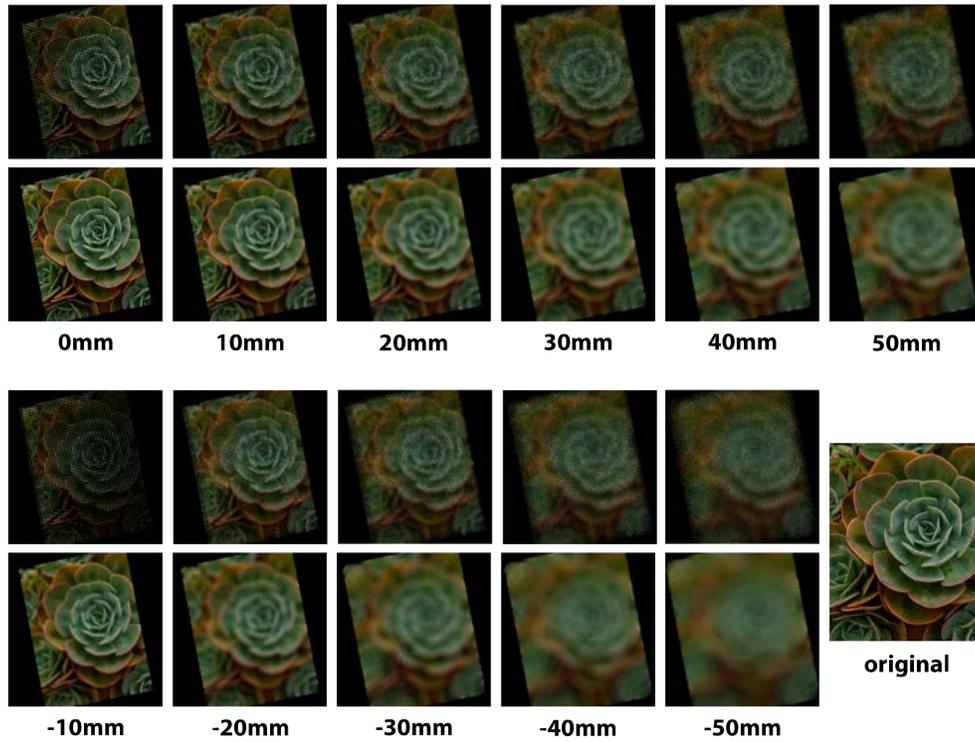


Figure 36: Tests showing varying levels of defocuss error.

5.6 Error Resiliency to Angle Changes

We ran the precomputation once with the display viewed perpendicular before then simulating the eye's view of that precomputed image in angled displays.

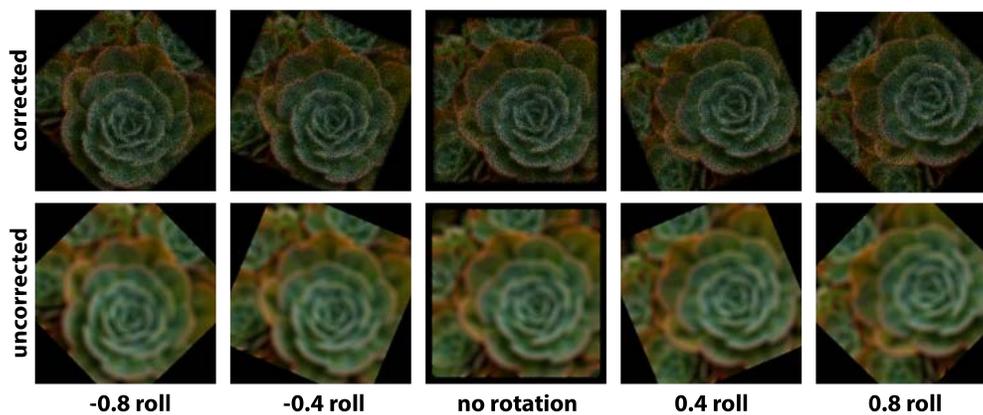


Figure 37: Different roll angles than the one used in precomputation.

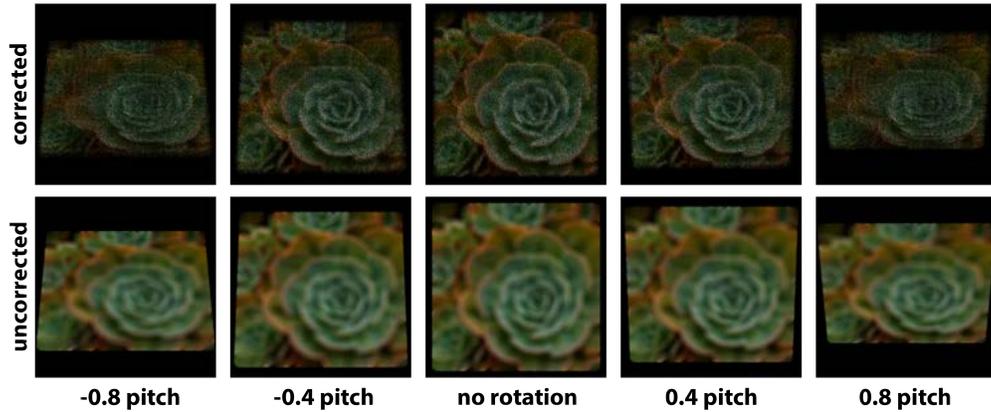


Figure 38: Different pitch angles than the one used in precomputation.

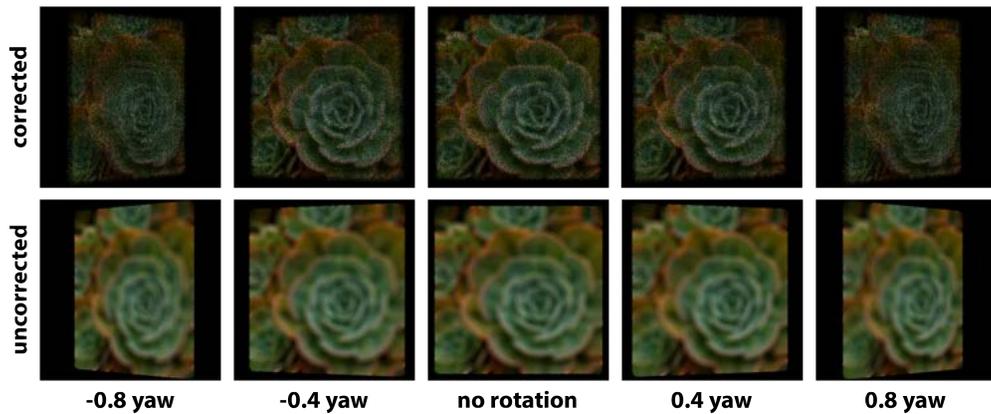


Figure 39: Different yaw angles than the one used in precomputation.

6 Discussion

Incorporating past work and new modifications the 3D implementation was able to demonstrate correction for defocus. The first step along this path was figuring out how to rewrite past calculations in 3D. After finding what ideas worked and debugging, the next step was conducting a series of experiments to determine how hardware parameters could effect the resulting perceived images and which values to take. This was discussed in the earlier approach section 4.7. Finally we conducted experiments to compare the correction across the many different viewing conditions possible with levels of defocus error and display angle.

After experimenting with a variety of different hardware settings it was found that the double image artifact shown in section 5.2 was fixed by adding tighter pupil and pinhole apertures. This gives evidence that the double image artifact experienced in this implementation is the same as the one noted in past works. (Pamplona et al., 2012) (Huang et al., 2014) The reason tighter apertures can provide a solution is that they can eliminate

portions of the conflicting light field views caused by cross talk or violation of the one-to-one assumption as discussed earlier in section 2.6.

The biggest takeaway from testing in 3d was that more care has to be taken in sampling and in choosing display parameters. To get good results in the presence of large pitch and yaw angles, tight pupil apertures and pinhole diameters had to be used since the double error artifact was more likely to occur with these settings. In addition, with the presence of a display angle, the pinholes which would direct light to the retina were no longer the pinholes directly. Without extra consideration taken to find the correct mask opening, results were often too dark or completely black from lost sample rays.

The quality of our implementation was best at medium levels of refractive error. In the grid of images shown in figure 36, artifacts due to the pinhole mask can be seen at low levels of refractive error ($|error| < 20mm$) and the corrected images are somewhat noisier and blurrier at high levels of refractive error ($50mm$). When the display is close to in focus light doesn't spread very far from the originating pinholes so that a grid pattern due to pinhole placement can be seen. This is a definitive downside of using the pinhole mask to construct the light field display. While the results with high levels of refractive error are worse overall, they still show an improvement over the uncorrected versions. It is possible that using tighter pinhole apertures in conjunction with a higher sample rate could reduce some of the blur and noise. Still, the loss in quality demonstrates that when using a ray tracing based sampling scheme such as ours, the difficulty of getting an accurate estimate for the precomputation grows harder as the level of aberration increases.

An unexpected result was that a higher angular resolution wasn't necessary if care was taken in sampling. Since angled displays would require rays emitted at more extreme angles, it was reasonable to expect that a higher angular resolution would help. However, we found that even in low angular resolution 3x3 masks, good results could be generated by casting samples through the mask openings along the path to the eye. This is shown in the results from section 5.4. It was also expected that with low angular resolutions, the quality of the perceived image would drop sharply at specific intervals as the angle varied. However, we found that when testing a continuous range, between -0.8 radians and 0.8 radians for example, there weren't significant drops in quality at specific angles.

Section 5.6 of our results shows that while changing roll has negligible effect, changing pitch and yaw, without updating the precomputed image introduces errors in the perceived image. Previous works had shown this occurrence with mismatched defocus error. (Huang et al., 2014) In order to counteract this error a combination of error proofing and frequent update calculations will be required in the future.

The specific error introduced by inaccurate pitch and yaw is due to the fact that a larger portion of the display is located at a different depth than where it was assumed to be during the precomputation. The result is that a center band of the image remains in focus while the peripheries lose quality. This can be seen in figures 38 and 39.

The fact that roll angle did not play a significant roll in the quality of the perceived image indicates that the distance from the viewer is the most important factor in correcting for defocus error. This is expected because defocus aberration is symmetrical by nature. Thus the eye entry point doesn't matter as long as the rays enter the along the same radius from the eye's center and at the same angle. In the future when higher order aberrations are included in 3D correction it is likely that an accurate measure of the display's roll will become important to image quality. While the symmetry of lower order aberrations means that only the radius from the center of the pupil needs to be right, higher order aberrations are asymmetric and changing the point of entry can change how light is bent.

7 Future Work

7.1 Adding Global Optimization

Our computation correction algorithm relied on averaging samples in place of the global optimization used by (Huang et al., 2014). Zehao Wu, showed that direct assignment from samples creates results that are very close to what can be achieved using Huang et. al.'s optimization. (Wu, 2016) Still, we believe that adding the global optimization could bring an improvement to the quality of results and adds a convenient way to conduct error proofing which is especially important in 3D.

Correction in 3D needs to operate without the assumption that the retina is rectangular and has a similar resolution to the display. While this makes it impossible to use the exact scheme that Huang et. al. employed, one workaround is to treat a warped version of the input image as the ideal image in the optimization. Finding the correct warp would be a similar procedure to how we calculated the correct homography in our precomputation step.

7.2 Precise Pose Estimation

The topic of pose estimation has been widely researched in many areas of computer science and there are readily available tools that can be used to tackle this problem in practice. A key consideration for using this technology in vision correcting displays is that it must have a high level of accuracy in the estimate. In Huang et. al.'s work in the 2D setting they tested shifting the display laterally while maintaining the same viewing position. Without including neighboring shifts in the optimization step of their tailoring algorithm, they found that shifts on the order of 5mm could result in a drop in image quality. When viewed at even a close distance of 260mm away, this corresponds to a shift of 1° in the angle offsetting the display from the viewers eye. Thankfully, with the inclusion of neighboring shifts, the drop in image quality was drastically reduced. Added accuracy to the pose estimate could improve image quality further without adding complexity to the precomputation algorithm.

7.3 Accounting for Both Eyes

This work focused on the appearance of the display from a single eye. The reality is that the user will perceive a image formed from both eyes. Since the vision correcting display algorithm is so view dependent, it is necessary to simultaneously display two separate light fields for each eye accounting for the change in position between the two. Limited to a single screen to display two distinct light fields, this presents a significant challenge that, similarly to working in 3D, needs to be solved before vision correcting displays can become practical. Some research has been done investigating parallax barriers as a way to accomplish this task. (Huang, 2013) As an alternative to simultaneously displaying two image on a single display, near eye displays used in VR already have distinct sections for each eye. L. Yang investigated using vision correcting displays in VR headsets in her masters thesis. (Yang, 2020). The issues of adapting to 3D and simultaneous display are removed in this domain. These factors along with the problem many people encounter of not being able to wear their corrective glasses inside of a headset make this a very promising direction for future research.

7.4 Study with Video and Continuous Changes

In this work we did not test our implementation on continuously changing display positions or video. In both situations, we expect that more artifacts will arise since we have only considered fixed images. While continuously recalculating the tailored light field will help remove drops in image quality from a mismatch between the position and angle of the display and the calculation assumptions, there is no guarantee that these separately correct images will look continuous when played back to back. For example, as angle changes a certain pixel may become visible from a different mask opening causing a shift in where the light escapes the display mask. We anticipate similar issues with video input. Since our tailoring algorithm relies on random sampling, frame to frame there could be flickering artifacts. An additional consideration should be given to any optimization based error proofing in the case of video as well. In this work, we did not include a global optimization. However, in F. Huang et. al.'s implementation which did include an optimization, the overall brightness of the image was found to change frame to frame. (Huang et al., 2014)

8 Conclusion

This work extends past approaches into 3D serving as a view into the changes required and a test of the quality that can be achieved. First we chose the parts of past approaches that would function best in this setting. Next, we made modifications to those past approaches where necessary. Finally we tested the performance of our implementation in a variety of different display settings.

Modifications we made to work in 3D were using standard ray tracing techniques such as ray plane intersections in place of 2D optics calculations, using homographies to warp desired images onto the retina, employing multiple coordinate systems to make calculations easier and altering sampling and display mask parameters. Employing the homography perspective warp was a necessary change since retina footprints of the display could change shape drastically. The ways we employed 3D calculations and a second coordinate system made the implementation much easier both to code and to debug.

Our results showed that with proper modifications to the vision correcting display algorithm, when correcting for defocus the main source of error introduced by allowing display angles is differences in depth across the display. We additionally found that a higher angular sampling rate was not required to correct in the presence of display angle. However, care needs to be taken in choosing display configuration parameters such as pinhole aperture diameter so that sampling done in the precomputation step can gain an accurate estimate of how light will reach the retina.

The extension of existing algorithms into 3D is needed to make vision correcting displays practical in a real world setting. Over the course of this project we were able to prototype a correction algorithm that could successfully correct for defocus in the presence of arbitrary display angles. In doing so we were able to test some of the challenges of correction with display angles and lay a groundwork for future works operating in 3D.

References

(2010). Wave aberration of human eyes and new descriptors of image optical quality and visual performance. 36:313–331.

- Alonso, J. and Barreto, A. (2003). Pre-compensation for high-order aberrations of the human eye using on-screen image deconvolution. *Engineering in Medicine and Biology Society*, 1:556 – 559.
- Alonso, M. and Barreto, A. (2004). An image processing approach to pre-compensation for higherorder aberrations in the eye. *Systemics, Cybernetics AND Informatics*.
- Atchison, D. A. and Thibos, L. N. (2016). Optical models of the human eye. *Clinical and Experimental Optometry*, 99(2):99–106.
- Bababekova, Y., Rosenfield, M., Hue, J., and Huang, R. (2011). Font size and viewing distance of handheld smart phones. *Optometry and Vision Science*, 88(7):795–797.
- Barsky, B. (2004). Vision-realistic rendering: simulation of the scanned foveal image from wavefront data of human subjects. *Applied Perception in Graphics and Visualization*, pages 73 – 81.
- Bhardwaj, V., Rajeshbhai, G., and Clin Diagn, R. J. (2013). Axial length, anterior chamber depth-a study in different age groups and refractive errors. *Journal of clinical and diagnostic research*, 7(10):2211–2212.
- Bradski, G. (2000). The OpenCV Library. *Dr. Dobb’s Journal of Software Tools*.
- Ding, C. (2017). Algorithms and applications of the vision correcting display. Master’s thesis, University of California, Berkeley.
- Ding, Y., Xiao, S., Yang, E., Zhao, Y., and Zhang, Z. (2020). Extending vision correcting display algorithm to allow for flexible viewer position. Master’s thesis, University of California, Berkeley.
- GuangMing, D. (2008). *Wavefront Optics for Vision Correction*. SPIE Press.
- Hartley, R. and Zisserman, A. (2003). *Multiple View Geometry in Computer Vision*. Cambridge University Press, USA, 2 edition.
- Hecht, E. (1987). *Optics. 2nd Ed.* Addison-Wesley.
- Hsieh, Y., Yu, Y., Lai, Y., Hsieh, M., and Chen, Y. (2020). Integral-based parallel algorithm for the fast generation of the zernike polynomials. *Opt. Express*, 28:936–947.
- Huang, F. (2013). *A Computational Light Field Display for Correcting Visual Aberrations*. PhD thesis, University of California, Berkeley.
- Huang, F., Lanman, D., Barsky, B., and Raskar, R. (2012). Correcting for optical aberrations using multilayer displays. *ACM Trans. Graph (SIGGRAPH Asia)*, 31(6):185:1–185:12.
- Huang, F., Wetzstein, G., Barsky, B., and Raskar, R. (2014). Eyeglasses-free display: Towards correcting visual aberrations with computational light field displays. *ACM Proc. of SIGGRAPH*, 33(4):1–12.
- Huang, F.-C. and Barsky, B. A. (2011). A framework for aberration compensated displays. Technical Report UCB/EECS-2011-162, EECS Department, University of California, Berkeley.

- Klien, A. and Barsky, B. A. (1995). Method for generating the anterior surface of an aberration-free contact lens for an arbitrary posterior surface. *Optometry and Vision Science*, pages 816–820.
- Kramida, G. (2015). Resolving the vergence-accommodation conflict in head-mounted displays. *IEEE transactions on visualization and computer graphics*, 22:1912 – 1931.
- Marschner, S. and Shirley, P. (2018). *Fundamentals of Computer Graphics*. A K Peters/CRC Press.
- Navarro, R. (2019). The optical design of the human eye: a critical review. *Journal of Optometry*, 2:3–18.
- Pamplona, V. F., Oliveira, M. M., Aliaga, D. G., and Raskar, R. (2012). Tailored displays to compensate for visual aberrations. *ACM Trans. Graph.*, 31(4).
- Ren, N. and Pat, H. (2006). Digital correction of lens aberrations in light field photography. In Gregory, G. G., Howard, J. M., and Koshel, R. J., editors, *International Optical Design Conference 2006*, volume 6342, pages 441 – 454. International Society for Optics and Photonics, SPIE.
- Wu, Z. (2016). Investigating computational approaches and proposing hardware improvement to the vision correcting display. Master’s thesis, EECS Department, University of California, Berkeley.
- Yang, L. (2020). An investigation into microlens array based light field display for vision aberration correction and applications to virtual reality. Master’s thesis, EECS Department, University of California, Berkeley.
- Zhen, Y. (2019). New algorithms for the vision correcting display. Master’s thesis, University of California, Berkeley.