Towards Parallax-Based Unencumbered Displays

by

Ryan A. Shanks

A thesis submitted in conformity with the requirements
for the degree of Master of Science
Faculty of Graduate Studies (Computer Science)
University of Ontario Institute of Technology

Supervisor(s): Dr. Mark Green and Dr. Faisal Z. Qureshi

Abstract

Towards Parallax-Based Unencumbered Displays

Ryan A. Shanks

Master of Science

Faculty of Graduate Studies

University of Ontario Institute of Technology

2015

We present an analysis of content adaptive parallax barriers used for multi-layer unencumbered displays. We explain the techniques involved in detail from the input light field to the output masks and how to display them. The content adaptive parallax barrier masks are produced by decomposing a matrix, which is created by applying a set of constraints to the input light field, using non-negative matrix factorization. We compare a number of matrix factorization methods, including a novel technique developed in this work. We provide a detailed exploration of design parameters for the multi-layer display to produce proper autostereoscopic results. A number of datasets are used to produce simulated results for a comparison of factorization techniques and other elements of the process in creating an unencumbered display using content adaptive parallax barriers.

Acknowledgements

I would like to extend my sincere appreciation to my supervisors, Dr. Mark Green and Dr. Faisal Qureshi, for their guidance and inspiration. They taught me a lot about research and how to properly communicate my ideas.

I would like to also thank my family and friends for all their support. My parents, Judy and Roy, and my sister, Megan, for their unwavering support.

Contents

List of Tables

List of Figures

List of Symbols

N_p	Number of pixels	59, 65
P_F	Front LCD pixel pitch, or pixel size	53, 59, 65
P_G	Back LCD pixel pitch, or pixel size	53, 59
α	Represents the line-search values used in FNMA_e and WFNMA_e, which designates how far to advance on the line in each dimension	21, 50
$\vec{\boldsymbol{a}}_L$	A ray representing the left most aligned pixels in a content adaptive parallax barrier display	62, 63, 65–67, 104
$\vec{\boldsymbol{a}}_R$	A ray representing the right most aligned pixels in a content adaptive parallax barrier display	62, 63, 65, 66, 104
$\vec{\boldsymbol{m}}$	The closest point at which a specific view can be perceived when vewing a content adaptive parallax barrier display	61–68, 105
$\vec{\boldsymbol{o}}$	A direction vector	66, 67, 104
$\vec{\boldsymbol{p}}_{LB}$	Left most pixel on the back LCD panel involved in this V. Used to determine $\vec{\boldsymbol{a}}_L$	65, 66, 104
$\vec{\boldsymbol{p}}_{LF}$	Left most pixel on the front LCD panel involved in this V. Used to determine $\vec{\boldsymbol{a}}_L$	65, 66, 104
$\vec{\boldsymbol{p}}_{RB}$	Right most pixel on the back LCD panel involved in this V. Used to determine $\vec{\boldsymbol{a}}_R$	65, 66, 104
$\vec{\boldsymbol{p}}_{RF}$	Right most pixel on the front LCD panel involved in this V. Used to determine $\vec{\boldsymbol{a}}_R$	65, 66, 104

λ	A parameter for the plenoptic function signifying the wavelength of light	12
ϕ	A parameter for the plenoptic function signifying the viewing angle	12, 13, 29
τ	A parameter for the plenoptic function signifying time	12
θ	A parameter for the plenoptic function signifying the viewing angle	12, 13
d_p	Distance between the front LCD and the back LCD	53, 59, 67, 104–106
t	Parameter which defines the intersection point between nd	66, 67, 105
D	Hessian Approximation	20, 21, 23, 50
F	Matrix containing front masks	15–22, 24–27, 34–37, 46–49
G	Matrix containing back masks	15–22, 24–27, 34–37, 46–49
L	Light field matrix	15–19, 21–27, 34–39, 42, 43, 46, 47, 49, 50, 101
W	Weighting matrix	22, 27, 49, 50
i	A specific row on the UV-plane	14, 39, 42, 49
j	A specific column on the UV-plane	14, 42, 49
k	A specific row on the ST-plane	14, 39, 42

l	A specific column on the ST-plane	14, 39, 42
m	The number of columns in the light field matrix, **L**	16, 24, 25, 27
n	The number of rows in the light field matrix, **L**	16, 24, 25, 27
r	The rank of the decomposition during matrix factorization	16, 17, 27
s	The number of rows on the ST-plane	13–15, 34, 39, 42, 47, 55
t	The number of columns on the ST-plane	13–15, 34, 42, 47, 55
u	The number of rows on the UV-plane	13–15, 34, 35, 39, 42, 55
V	A specific view shown by a content adaptive parallax barrier display	61, 62, 64, 67
v	The number of columns on the UV-plane	13–15, 34, 35, 42, 55
x	A coordinate along the x-axis	12, 13
y	A coordinate along the y-axis	12, 13

z	A coordinate along the z-axis	12

List of Abbreviations

BFGS	Broyden-Fletcher-Goldfarb-Shanno optimization update rule	20, 21
$FNMA_e$	Fast Non-negative Matrix Approximation	iv, 18, 21, 47, 49, 73–76, 82, 87, 92, 97
FOV	Field of View	106
IPD	Interpupillary Distance	vi, 56, 57, 59, 63, 65, 66, 103–108
LCD	Liquid Crystal Display	16, 28, 43, 46, 50–52, 54–56, 59, 60, 63, 65, 105, 108, 111, 112
MAX	The Maximum Possible Value	24
MSE	Mean Squared Error	24, 25
NMF	Non-negative Matrix Factorization	16, 34, 76, 113

MVD	Minimum Viewing Distance	65, 66, 105
PSNR	Peak Signal-to-Noise Ratio	24, 73, 75
WFNMA_e	Weighted Fast Non-negative Matrix Approximation	76

Chapter 1

Introduction

Leonardo DaVinci discussed the importance of producing images that cater to the binocular property of the human visual system [16]. This is a reference to the fact that we see with two eyes and our brain interprets the separate images perceived, as a single three dimensional (3D) image. Binocular vision brings some fascinating aspects such as depth cues and occluders. These aspects become particularly interesting when creating artificial images because we need to be very precise in how the images are produced or the viewer will feel removed from the presented environment or could even feel discomfort. Proper depth cues, distances, the space between the eyes and many more factors must be taken into account in 3D imaging.

The first device used for binocular, or stereoscopic, viewing of images was the stereoscope [72]. The original stereoscope, created by Charles Wheatstone, was a large device used for displaying a pair of stereographic images using mirrors to control what was visible to each eye [73]. Other researchers, such as Brewster and Holmes, contributed to the early study of stereoscopic viewing and created alternate stereoscopes [11, 12, 35]. Some of these early stereoscopes are seen in Figure 1.1. The stereoscope has been modernized over time to create devices such as the slide viewer (Figure 1.2a) or the Oculus Rift (Figure 1.2b) [31, 53, 60, 68]. There are many forms of 3D displays other than the stere-

oscope such as anaglyphic, polarized, and active shutter glasses, as well as volumetric, holographic, and parallax barrier displays [5, 6, 59].

Anaglyphic, polarization and active shutter displays all require glasses to be worn for 3D images to be seen. Anaglyphic 3D was first introduced by Louis du Hauron [21]. Anaglyphs are best known for the red and cyan glasses, where opposite colours restrict which portions, of a pair of overlayed images, are displayed to each eye [22, 70]. The colour pair, red and cyan, is one of many that can be used for producing and displaying anaglyph images. Anaglyph 3D is known to have a ghosting effect, where one eye will see some of the image intended for the other eye, which is caused by crosstalk, or the incomplete isolation of the image to its intended colour channel [74].

Polarized glasses are the grey glasses commonly seen when viewing 3D movies in the theatre. Polarization is a property of waves that can oscillate in more than one direction. Light, an electromagnetic wave, has an electric field and a magnetic field where oscillations of these fields determine the wave's polarization. There are different types of polarization; the main two are circular and linear. Polarized glasses can filter light, allowing only certain polarized light to pass through. Linearly polarized glasses have one lens which filters all light that does not have a horizontal polarization and the other lens filters all light without a vertical polarization. The pitfall is that if the viewer rotates their head the displayed images will become darker. Circular polarization allows for some head rotation, but has a lower maximum brightness. Circularly polarized glasses can be seen in Figure 1.4(b). Using these glasses, each individual image from a stereo image pair, is displayed with a different polarization, resulting in a single eye viewing a particular image [9]. The images are displayed rapidly so that the human visual system will combine the stereo pair into a single scene.

Active shutter glasses work by rapidly displaying stereo image pairs, but actively block one of the images to one of the eyes [52]. The glasses have the ability to block all light to either eye and are calibrated with the display to coordinate the blocking of

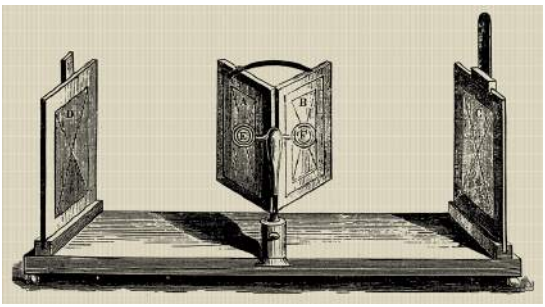

(a) A Wheatstone Bridge[1]

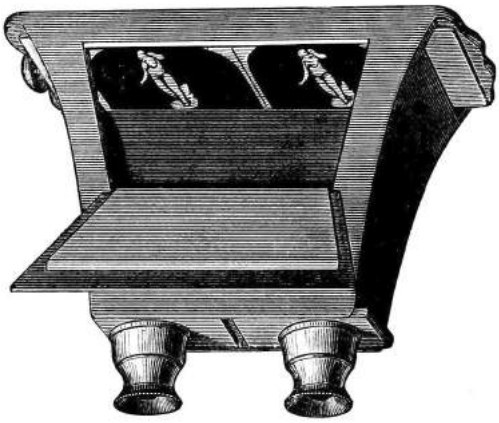

(b) The Brewster Stereoscope[2]

(c) The Holmes Stereoscope[3]

(d) A stereogram[4], which would be used in (b) or (c)

Figure 1.1: Early examples of stereoscopes (a)-(c) and a stereogram they might be used to view (d). (a) A Wheatstone Bridge. (b) The Brewster Stereoscope. (c) The Holmes Stereoscope. (d) A stereogram, which would be used in (b) or (c).

[1]Charles Wheatstone-mirror stereoscope XIXc (https://upload.wikimedia.org/wikipedia/commons/1/1f/Charles_Wheatstone-mirror_stereoscope_XIXc.jpg) Accessed on 14 July 2015. Public Domain.

[2]The Brewster stereoscope 1849 (https://upload.wikimedia.org/wikipedia/commons/b/b2/PSM_V21_D055_The_brewster_stereoscope_1849.jpg) Accessed on 14 July 2015. Public Domain.

[3]Holmes stereoscope, Dave Pape (https://upload.wikimedia.org/wikipedia/commons/f/f6/Holmes_stereoscope.jpg) Accessed on 14 July 2015. Public Domain.

[4]Stereograph as an educator, Underwood & Underwood (https://upload.wikimedia.org/wikipedia/commons/0/04/Stereograph_as_an_educator.jpg) Accessed on 14 July 2015. Public Domain.

(a) A slide viewer[1] (b) The Oculus Rift[2]

Figure 1.2: Modernized versions of the stereoscope.

[1]View-Master Model E, (c)ThePassenger (https://upload.wikimedia.org/wikipedia/commons/c/c9/View-Master_Model_E.JPG) Accessed on 14 July 2015. CC BY-SA 3.0.

[2]Oculus Rift - Developer Version - Front, (c)Sebastian Stabinger (https://upload.wikimedia.org/wikipedia/commons/a/ae/Oculus_Rift_-_Developer_Version_-_Front.jpg) Accessed on 14 July 2015. CC BY 3.0.

light [32]. Although not as widely used as polarized glasses, active shutter technology is used in higher end cinemas and in home theatres.

Volumetric displays form images in physical space, where other technologies project images onto a two-dimensional plane rather than using all three dimensions. One type of volumetric display uses a concept known as flicker fusion, where the human visual system integrates across rapidly displayed images [62]. This is done by rotating a set of Light Emitting Diodes (LEDs) quickly and displaying pieces of the intended image at the appropriate locations in space. When this is done fast enough human eyes will interpret the pieces as a whole image.

Holographic displays project an interference pattern [49], which represents a light field. The interference pattern is produced from differences between a laser and its reflection off of an object or scene [50]. When the original laser is used to project the interference pattern the recorded scene is perceived. During projection the interference pattern diffracts the laser producing the image that was stored in the pattern.

Parallax barriers are placed in front, and block a portion, of a display with respect

(a) A hologram of a mouse[1]

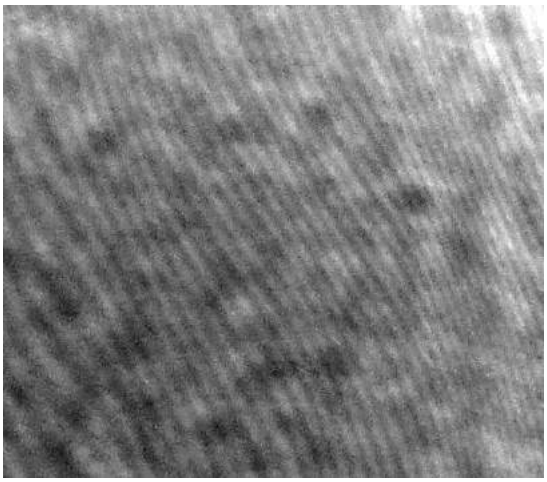

(b) A hologram diffraction pattern[2]

(c) A volumetric display[3]

Figure 1.3: Some examples of autostereoscopic displays.

[1]Holo-Mouse, (c)Georg-Johann Lay (https://upload.wikimedia.org/wikipedia/commons/5/5f/Holomouse2.jpg) Accessed on 28 August 2015. Public Domain.

[2]Holographic recording, (c)Epzcaw (https://upload.wikimedia.org/wikipedia/commons/a/a5/Holographic_recording.jpg) Accessed on 28 August 2015. Public Domain.

[3]UCSI Volex, (c)Andreba (https://upload.wikimedia.org/wikipedia/commons/7/7e/UCSI_Volex.jpeg) Accessed on 15 August 2015. CC BY 3.0.

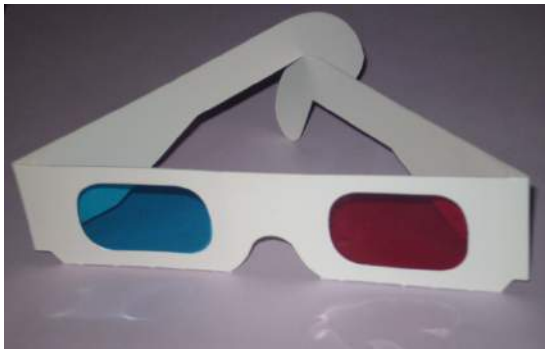

(a) Anaglyphic glasses[1]

(b) Polarized glasses[2]

(c) An anaglyphic image[3]

(d) Active shutter glasses[4]

Figure 1.4: Different forms of 3D glasses.

[1]Anaglyph glasses, (c)Snaily (https://upload.wikimedia.org/wikipedia/commons/e/e7/Anaglyph_glasses.png) Accessed on 14 July 2015. CC BY-SA 3.0.

[2]Real D glasses, Midori iro (https://upload.wikimedia.org/wikipedia/commons/4/47/REALD.JPG) Accessed on 14 July 2015. Public Domain.

[3]Stereograph as an educator - anaglyph, Dave Pape (https://upload.wikimedia.org/wikipedia/commons/0/0c/Stereograph_as_an_educator_-_anaglyph.jpg) Accessed on 14 July 2015. Public Domain.

[4]Xpand LCD shutter glasses, (c)Amidror1973 at English Wikipedia (https://upload.wikimedia.org/wikipedia/commons/9/9c/Xpand_LCD_shutter_glasses.jpg) Accessed on 14 July 2015. CC BY-SA 3.0.

to each eye. The display is divided such that each portion only shows parts of the image intended for a single eye. Refer to Section 3.1 for more information on parallax barriers.

Unencumbered, or autostereoscopic, displays, are those that display a 3D image without the use of any hardware that can burden the user. These could be glasses or some device used for tracking the user's position. Volumetric, holographic and parallax barrier displays can each be unencumbered. Parallax barriers have existed for a long time, but a recent development in the area of computational displays has given rise to content adaptive parallax barriers, which is a promising step forward towards commercial unencumbered displays [39, 45].

1.1 Contributions

The main contributions of this thesis are:

1. A complete and rigorous explanation of content adaptive parallax barriers
2. A weighted version of Fast Non-negative Matrix Approximation (FNMA_e)
3. An analysis of autostereoscopic design for Content Adaptive Parallax Barriers
4. An alternative light field sampling method named Pixel Barriers

1.2 Overview

The rest of the thesis is organized as follows. Background material is discussed in the next chapter. Content adaptive parallax barriers and the contributions of this thesis are presented Chapter 3. Chapter 4 discusses implementation details. Results of the work are presented in Chapter 5. Chapter 6 provides conclusions and discussion.

Chapter 2

Related Works

Content adaptive parallax barriers replace conventional parallax barriers through the generation of mask pairs from a light field. These masks produced are displayed on a dual-layer Liquid Crystal Display (LCD) [45]. Tensor displays build upon content adaptive parallax barriers. These produce mask sets containing more than two masks. Consequently, these need an n-LCD setup for viewing purposes. Polarization fields improve both tensor displays and content adaptive parallax barriers by treating LCD panels as polarization rotators, which allow for increased brightness by decreasing the number of intermediate polarizer layers. All of these methods process light fields, using non-negative matrix (or tensor) factorization machinery, to construct masks that are displayed using two or more LCD panels for viewing purposes.

2.1 Light Fields

A light field is a continuous four dimensional function describing a snapshot of a single wavelength of light in a scene from a subset of the possible viewing positions, $L = L(u, v, s, t)$. Light fields are often represented by a set of images, where the three colour channels, red, green and blue (RGB), each represent a single wavelength of light. As seen in Figure 2.1, each image displays the scene from a different viewing angle.

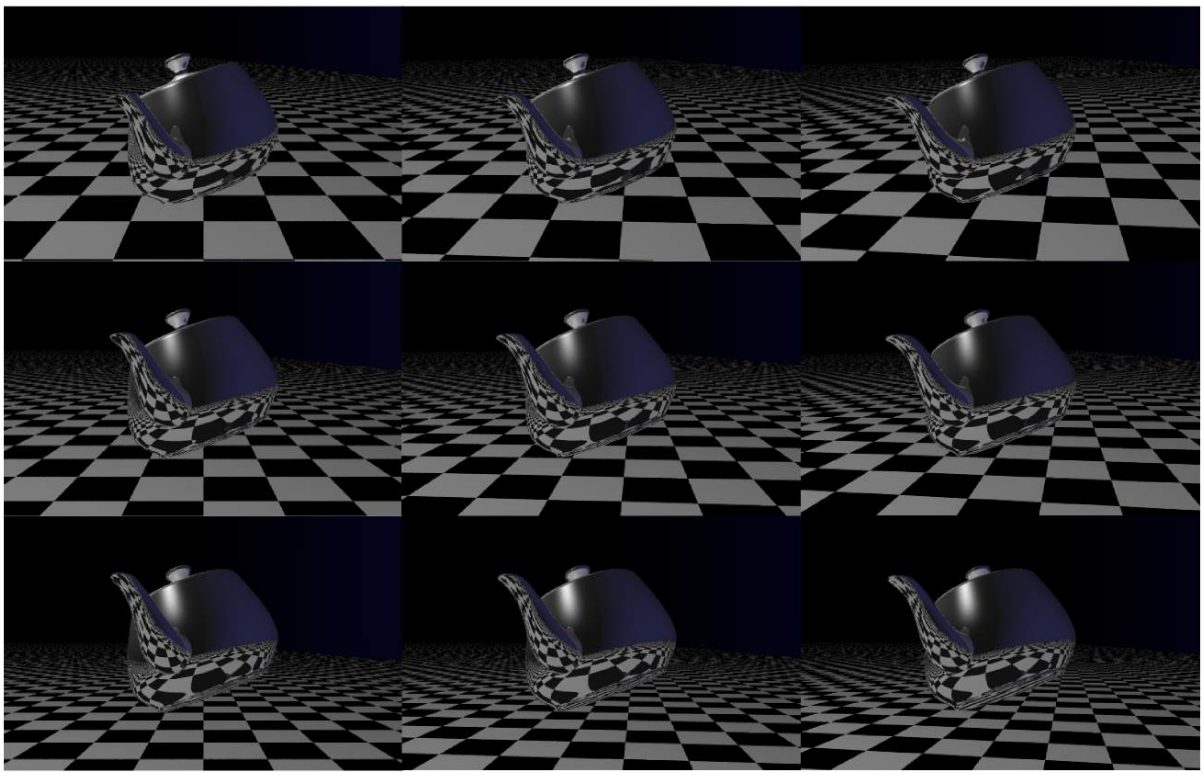

Figure 2.1: A sample light field of a teapot containing 3x3 views.

2.1.1 The Plenoptic Function

The light field is derived from the seven dimensional plenoptic function, which describes all the light in a scene from all viewing positions and angles [2,30]. A viewing position in three dimensional space ($\mathbb{R}^3$) requires a value for each dimension (x, y, z) and angles in $\mathbb{R}^3$ can be represented using spherical coordinates (θ, ϕ). The intensity of light is described by its wavelength (λ), but this may change so a specific time (τ) must be defined. All of this results in the plenoptic function, as seen in Figure 2.2(a):

$$P = P(\theta, \phi, \lambda, \tau, x, y, z).$$

2.1.2 Light Field Derivation

A light field is a four dimensional subset of the plenoptic function. More specifically, a light field describes a snapshot in time, of a single wavelength, in a scene without

occluders [30] (space without occluders is also known as free-space [51]). We can thus eliminate the wave length and time dimension present in the plenoptic function. This results in the following five dimensional function as seen in Figure 2.2(b):

$$P = P(\theta, \phi, x, y, z).$$

We eliminate the final dimension with an assumption of free space. Free space implies the convex hull of the scene is always viewed, meaning that from each viewpoint there is a single uninterrupted ray to any point in the scene [51]. One of the viewing position dimensions can be eliminated by mapping all of the rays to a single plane outside of the convex hull of the scene. As seen in Figure 2.2(c), the current description describes a position on the XY-plane and a viewing direction which gives the following function:

$$L = L(\theta, \phi, x, y).$$

The light field is a collection of rays. A ray can be represented in many ways. The current definition of a light field represents rays as having a starting position, (x, y), and a direction, (θ, ϕ). A ray could instead be represented using two points, where the first point is the starting position of the ray and the second point describes the direction to travel towards, from the start. The two-point representation of a ray provides an alternate light field representation, known as the two-plane parameterization [30, 51]. Figure 2.2(d) shows that a ray's starting position, (u, v), is located on the first plane and another point on the ray, (s, t), sits on the second plane. The light field represents the set of all rays that pass through the front plane, the UV-plane, and the back plane, the ST-plane. This yields the desired light field equation:

$$L = L(u, v, s, t).$$

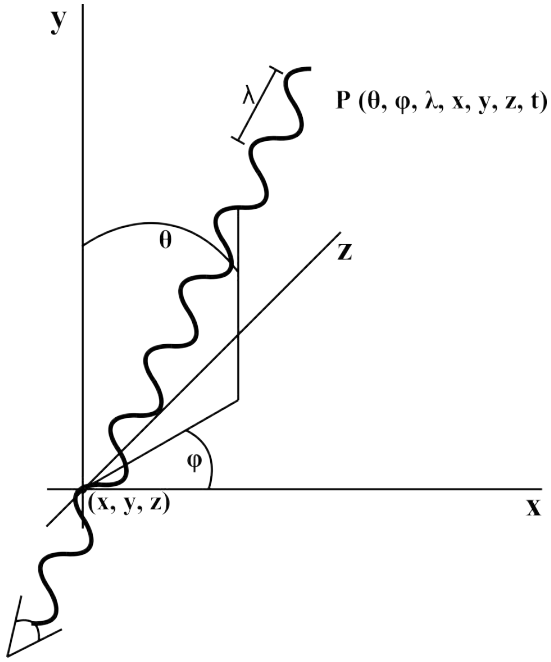

(a) Plenoptic function

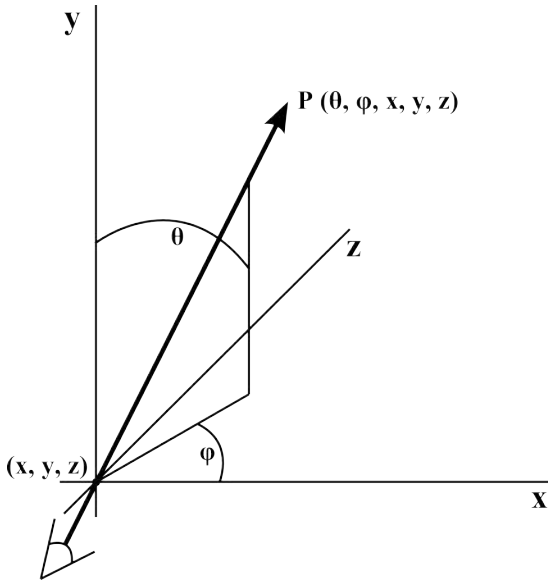

(b) Plenoptic function with set wavelength and time

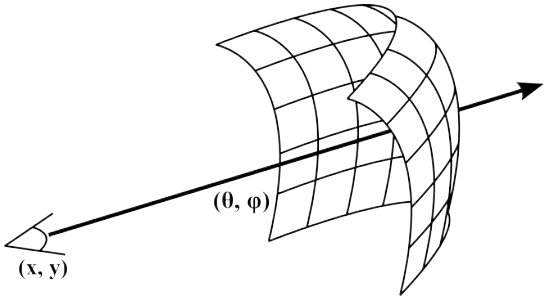

(c) Light field with polar coordinates

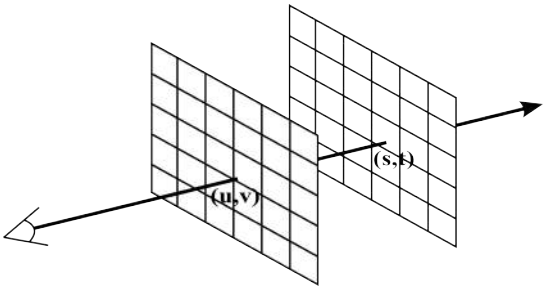

(d) Light field two-plane parameterization

Figure 2.2: Visualized subsets of the plenoptic function depicting the derivation of the two plane parameterization of the light field. [2, 51, 54]

2.1.3 Light Field Representation

A light field is a continuous function, but for practical use it needs to be represented discretely. When the two-plane parameterization is examined we see that a light field stores views of a scene from different angles. This observation allows us to move from $L(u, v, s, t)$ to $L(i, j, k, l)$, where i, j, k and l are discrete values each bounded by u, v, s and t, respectively. $L(i, j, k, l)$ gives a defined set of pixels, or resolution, to the front and back planes of the two-plane parameterization. The resolution of the front plane, represented by u and v, is the number of images that are present in the light field. The resolution of the back plane, represented by s and t, is the number of pixels in each image. A dense light field is one in which the maximum of each parameter is equal, $u = v = s = t$.

2.2 Parallax-Based Unencumbered Displays

An unencumbered display is one that works without any supporting devices that encumber the user. An example of an encumbering device is the glasses used with some three dimensional displays, which can be seen in Figure 1.4. Parallax barriers, discussed in Section 3.1, are a method to display three dimensional images without any other assisting devices.

2.2.1 Content Adaptive Parallax Barriers

Content adaptive parallax barriers use time-multiplexed masks, approximating a light field, to display three dimensional content on a dual-layer LCD setup. An input light field is approximated using a non-negative matrix factorization technique. The rank of the approximation defines the number of mask pairs produced. The mask pairs are displayed rapidly to take advantage of flicker fusion, which allows for displaying a better approximation of the light field. The objective function for content adaptive parallax

barriers is:

$$\underset{0\leq \mathbf{F}(i,j),\mathbf{G}(i,j)\leq 1}{\text{minimize}} \quad ||\mathbf{L}-\mathbf{FG}||^2.$$

Any weighted non-negative matrix factorization technique can be used, but a modified version of the Lee & Seung method was the original choice [45]. A complete explanation of content adaptive parallax barriers can be found in Chapter 3.

2.2.2 Tensor Displays

Building upon content adaptive parallax barriers, tensor displays allow for a better approximation of the input light field by accomodating an n-layer display, as opposed to only a dual-layer display [71]. This concept employs a non-negative tensor factorization scheme, which produces mask sets, rather than mask pairs. At the core, the concept is exactly the same in that a light field is decomposed into a set of masks and the masks are displayed on a multi-layered LCD setup. The main difference lies with the factorization procedure where tensors, as opposed to matrices, are input and the update rules applied are modified for tensor operations rather than just matrix operations. The objective function for tensor displays is:

$$\underset{0\leq \mathbf{F}^{(n)}\leq 1}{\text{minimize}} \quad ||\mathcal{L}-\mathcal{W}\circ\mathcal{T}||^2$$

where $\mathcal{T}=[\mathbf{F}^{(1)},\mathbf{F}^{(2)},...,\mathbf{F}^{(n)}]$ with $\mathbf{F}^{(n)}$ being the nth-layer mask, $\mathcal{L}$ is a tensor created from the input light field, and $\mathcal{W}$ is a weighting tensor [71]. The update rules applied for this display are:

$$\mathbf{F}^{(n)} \leftarrow \mathbf{F}^{(n)} \circ \left(\frac{(\mathbf{W}_{(n)} \circ \mathbf{L}_{(n)})\mathbf{F}^n_{\odot}}{(\mathbf{W}_{(n)} \circ (\mathbf{F}^{(n)}(\mathbf{F}^n_{\odot})^{\mathbf{T}}))\mathbf{F}^n_{\odot}} \right),$$

where $\mathbf{F}^n_{\odot}$ is defined by the following expression:

$$\mathbf{F}^n_{\odot} \leftarrow \mathbf{F}^n \circ \left(\mathbf{F}^{(N)} \odot \cdots \odot \mathbf{F}^{(n+1)} \odot \mathbf{F}^{(n-1)} \odot \cdots \odot \mathbf{F}^{(1)} \right).$$

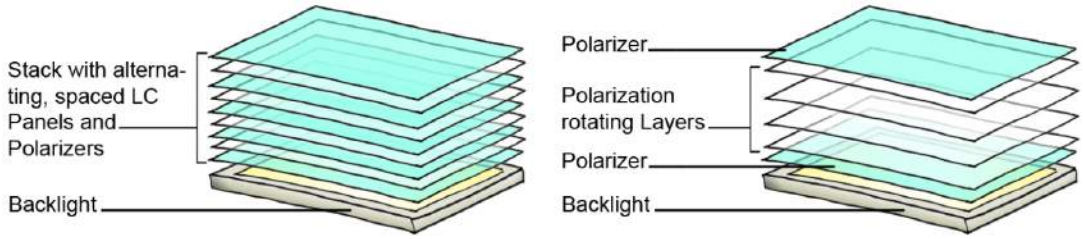

Figure 2.3: Displays the reduced number of polarizers in the polarization field design. (left) Content adaptive parallax barriers require polarizers between each LCD (right) Polarization fields only need two polarizers regardless of the number of layers. [46]

The symbol $\odot$ represents the Khatri-Rao product, defined for matrices $\mathbf{F} \in \mathbb{R}^{m \times r}$ and $\mathbf{G} \in \mathbb{R}^{n \times r}$ as:

$$\mathbf{F} \odot \mathbf{G} = [f_1 \otimes g_1 \quad f_2 \otimes g_2 \quad \cdots \quad f_k \otimes g_k],$$

where $\otimes$ is the Kronecker product and f_i and g_j represent the i^{th} and j^{th} columns of $\mathbf{F}$ and $\mathbf{G}$ respectively [71].

2.2.3 Polarization Fields

An extension of both tensor displays and content adapative parallax barriers, polarization fields can be used on a dual-layer or n-layer LCD setup. This concept improves on the others by changing the display design slightly to allow for an increased brightness. Polarization fields treat the LCD panels as polarization rotators, which means that polarizers between LCD panels can be removed so all that is needed is a single pair of vertical and horizontal polarizers, one prior to the panels and one after them as shown in Figure 2.3 [46]. Polarization fields use matrix factorization or tensor factorization, depending on the number of layers available in the display, but the input to the factorization process is different. The masks must produce optimal values for approximating the light field, but must do so while keeping the light at the proper phase as it passes through that layer. The passage of light through layered materials, such as LCD panels, is modelled using Jones calculus [40].

Polarization fields can use the same methods as either content adaptive parallax

barriers or tensor displays. The objective function for polarization fields is:

$$\underset{\phi_{min} \leq \phi_{ij} \leq \phi_{max}}{\text{minimize}} \quad ||\boldsymbol{\theta} - \mathbf{P}\boldsymbol{\phi}||^2,$$

where $\boldsymbol{\theta}$ is the polarization field, $\mathbf{P}$ is a projection matrix in which each element is given by the normalized area of overlap between pixels and rays, and $\boldsymbol{\phi}$ is the set of polarization state rotations [46].

2.2.4 Afterword

This area of research is very new, and a consequence is that there has been a limited amount of published work. Polarization fields are the most advanced approach as it reduces the necessary materials required for the display to work and increase the image brightness, which is a limiting factor as the rank and number of layers is increased. Tensor displays introduce the concept of many-layers which improves quality, but requires more hardware and the algorithms are more complex in comparison to content adaptive parallax barriers. Our plan was to investigate tensor displays by first looking at the less complex content adaptive parallax barriers as an introduction. We quickly discovered that an indepth explanation and analysis of content adaptive parallax barriers is needed to help promote research in this area.

Chapter 3

Content Adaptive Parallax Barriers

3.1 Parallax Barrier

Parallax is defined as the displacement of an observed object due to a change in the position of the observer. This perceived movement, as a consequence of a change in viewing position, is emphasized by a relative perception of the surrounding area. For example, if an object is in a completely blank environment the object may not appear to move significantly if it is at a distance, but if the object is near the observer it will appear to move rapidly. A demonstration of parallax can be viewed in Figure 3.1.

A parallax barrier is a layer which blocks some light, but allows other light to pass through [39]. Some parallax barrier patterns can be seen in Figure 3.2. The black sections will block light, while the white sections will allow light to pass through. Due to the fact that the eyes have some distance between them, known as the interpupillary distance (IPD), each eye will perceive different information when viewing a parallax barrier setup. This perception of different data is what allows a parallax barrier to provide the appearance of three dimensions. This is possible because the barrier allows different light to pass through based on the viewing position and angle. Intuitively, the parallax barrier will display one side of a stereo image to the left eye and the other side

Figure 3.1: Demonstration of parallax. The three sets of images are produced by moving the viewing position from left to right. The top and bottom images are identical other than the background[1] added.

[1]Buck Mountain Grand Teton NP1, (c)Acroterion (https://upload.wikimedia.org/wikipedia/commons/8/8b/Buck_Mountain_Grand_Teton_NP1.jpg) Accessed on 19 September 2015. CC BY-SA 3.0.

of a stereo image to the right eye as seen in Figure 3.3. For a parallax barrier to function the image being viewed must be some distance behind the barrier.

Parallax barrier setups have the potential to repeat the same 3D image to multiple viewers or show a different 3D perspective to each viewer [14,55]. Figure 3.4 demonstrates the repeated left and right eye viewing zones. Two repeated views are enough to display a 3D image, but [20] shows that a higher number of images can reduce artifacts. A higher number of views increases the quality of the perceived images, but reduces the number of potential unique viewing zones.

3.2 Content Adaptive Barriers

Content adaptive parallax barriers function in a similar manner to parallax barriers, in that they use parallax and limit the passage of light to display 3D content. However, they appear fundamentally different. There is no clear separation between the barrier and the content. Instead there are front and back masks, both of which work in conjunction to display the desired content. Different views are displayed by altering the alignment of

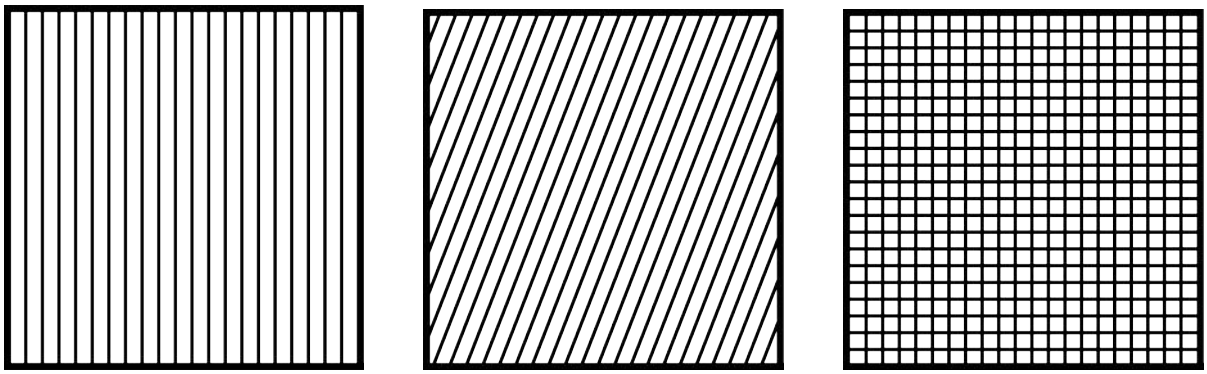

(a) Vertical pattern for horizontal parallax

(b) Diagonal pattern for reduced artifacts

(c) Cell pattern for horizontal and vertical parallax

Figure 3.2: Sample parallax barrier patterns.

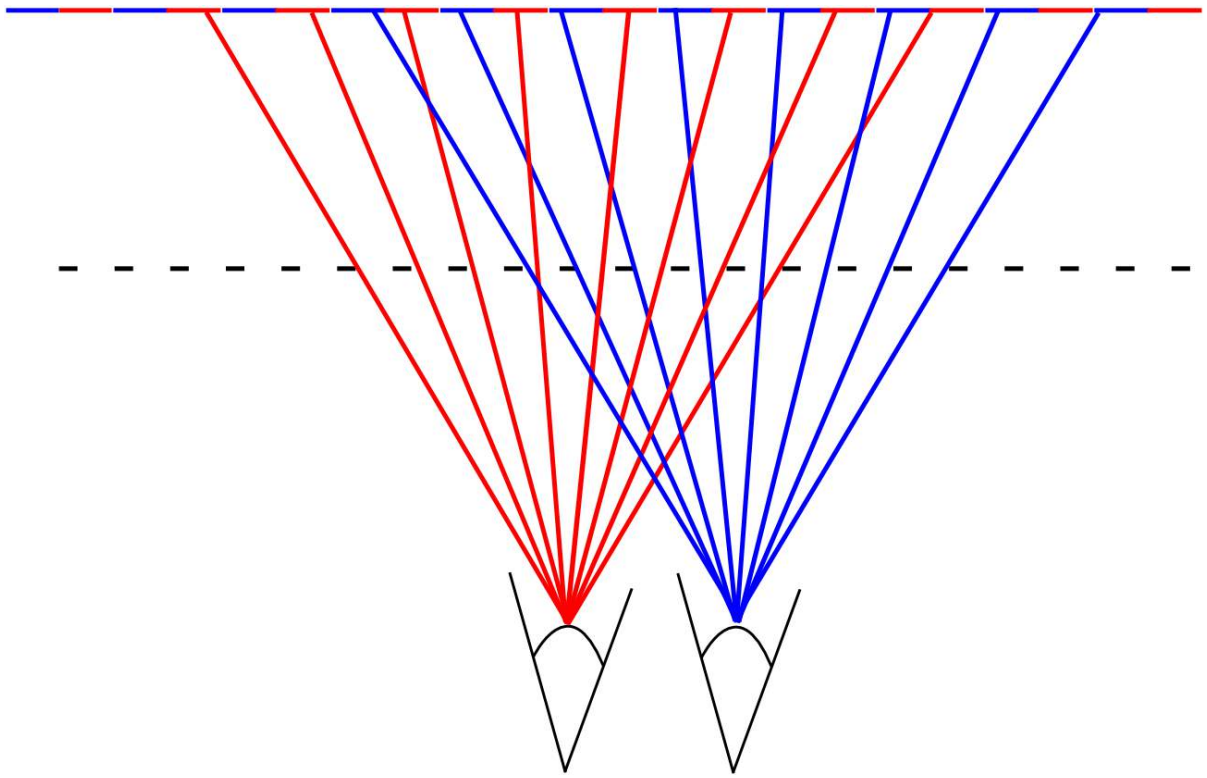

Figure 3.3: Eye separation means that each eye will view a scene from slightly different angles. The parallax barrier exploits the different viewing angles by blocking content for the left eye from the right eye and vice versa. This allows for viewing stereogram images without the encumbrance of glasses or other devices.

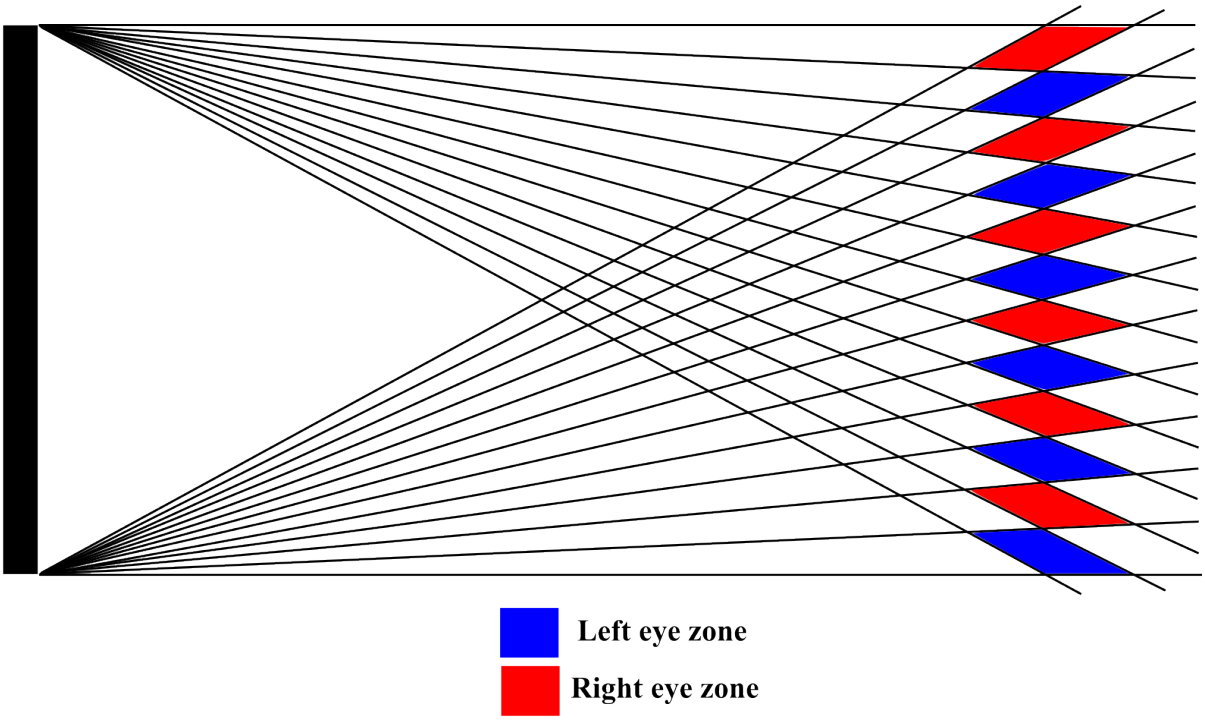

Figure 3.4: Demonstrates how a parallax barrier display can be designed to show the same 3D image to many users through repeated viewing zones.

the two masks. These masks are generated from a light field, an example of which can be seen in Figure 2.1. A light field matrix **L** is created from the input light field and the mask pairs are produced using NMF, which is described in Appendix A and later on in this chapter. The mask pairs are time-multiplexed, which means they are displayed rapidly and the human visual system will integrate them into a single image as long as the speed of the display exceeds the flicker fusion threshold.

3.3 Sampling

Content adaptive parallax barriers are a set of mask pairs produced from a collection of images. Specifically, the objective is to show each image in the collection by changing how the masks are aligned with respect to an observer. Looking at Figure 3.5a, the provided light field is seen where there are $u \times v$ images, $I_{u,v}$, and each image has $s \times t$ pixels, $p_{s,t}$. The light field matrix, **L**, is what needs to be approximated by **FG** to produce the

desired behaviour from the display. **L** and **FG** have the same dimensions, which means that a single entry of **FG** must approximate a single entry, in fact the same entry, of **L**. For a rank-1 decomposition, this means that one entry from **F** and one entry from **G** must multiply to produce an entry from **L**. This situation causes very strict constraints during NMF as it can be seen that each entry of **F** contributes to an entire row of **FG** and each entry of **G** contributes to an entire column of **FG**.

We want the masks to produce the center image, $I_{\frac{u}{2},\frac{v}{2}}$, when they are perfectly aligned, meaning each pixel from the front mask is aligned to its corresponding pixel in the back mask. The desired behaviour for the display uses the center pixel of the back mask as a pivot. We want new images to be viewed, as different pixels of the front mask become aligned with the center pixel of the back mask. All of the pixel alignments are changed, but the specific pixel from the front layer that is aligned to the back layer's center pixel determines which view is perceived. The change in alignment is driven by the viewing position of the user, relative to the display. So as they move their head around, different 3D images will be perceived.

The light field matrix, **L**, is constructed from the input light field images, but the method of construction, or how values from the light field images are picked for entries of **L**, is determined by the desired behaviour of the display. This means that the interaction of the masks, **F**' and **G**' shown in Figure 3.5c, provides a set of constraints that determine what values of **F** and **G** are needed to approximate values of the input light field. Since entries of **FG** directly correspond to entries of **L**, once we know what values from the light field correspond to entries of **FG**, that light field value can simply be entered into **L**.

3.3.1 Simple 2D Example

As a simple example, we will examine the desired behaviour in a two-dimensional system. A two-dimensional light field equates to a row of images, each with a single row of pixels,

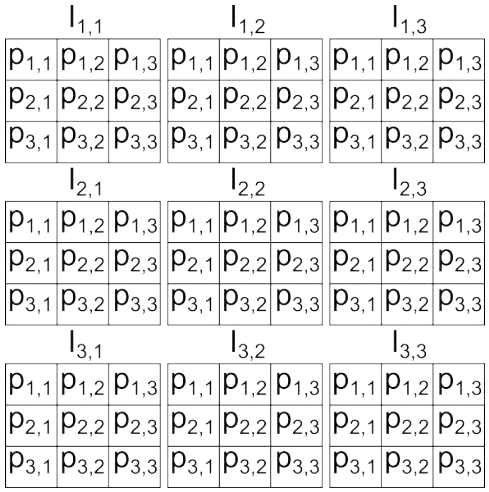

(a) The light field, represented by a two-dimensional array of images.

L

$l_{2,2}p_{1,1}$	$l_{2,1}p_{1,2}$		$l_{1,2}p_{2,1}$	$l_{1,1}p_{2,2}$				
$l_{2,3}p_{1,1}$	$l_{2,2}p_{1,2}$	$l_{2,1}p_{1,3}$	$l_{1,3}p_{2,1}$	$l_{1,2}p_{2,2}$	$l_{1,1}p_{2,3}$			
	$l_{2,3}p_{1,2}$	$l_{2,2}p_{1,3}$		$l_{1,3}p_{2,2}$	$l_{1,2}p_{2,3}$			
$l_{3,2}p_{1,1}$	$l_{3,1}p_{1,2}$		$l_{2,2}p_{2,1}$	$l_{2,1}p_{2,2}$		$l_{1,2}p_{3,1}$	$l_{1,1}p_{3,2}$	
$l_{3,3}p_{1,1}$	$l_{3,2}p_{1,2}$	$l_{3,1}p_{1,3}$	$l_{2,3}p_{2,1}$	$l_{2,2}p_{2,2}$	$l_{2,1}p_{2,3}$	$l_{1,3}p_{3,1}$	$l_{1,2}p_{3,2}$	$l_{1,1}p_{3,3}$
	$l_{3,3}p_{1,2}$	$l_{3,2}p_{1,3}$		$l_{2,3}p_{2,2}$	$l_{2,2}p_{2,3}$		$l_{1,3}p_{3,2}$	$l_{1,2}p_{3,3}$
			$l_{3,2}p_{2,1}$	$l_{3,1}p_{2,2}$		$l_{2,2}p_{3,1}$	$l_{2,1}p_{3,2}$	
			$l_{3,3}p_{2,1}$	$l_{3,2}p_{2,2}$	$l_{3,1}p_{2,3}$	$l_{2,3}p_{3,1}$	$l_{2,2}p_{3,2}$	$l_{2,1}p_{3,3}$
				$l_{3,3}p_{2,2}$	$l_{3,2}p_{2,3}$		$l_{2,3}p_{3,2}$	$l_{2,2}p_{3,3}$

(b) **L**, which is the matrix being approximated during NMF.

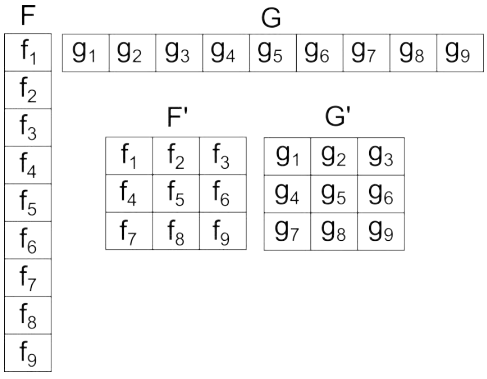

(c) **F** and **G** store a number of mask pairs equal to their rank. The rank shown here is one and the resulting mask pair is **F**' and **G**'.

FG

f_1g_1	f_1g_2	f_1g_3	f_1g_4	f_1g_5	f_1g_6	f_1g_7	f_1g_8	f_1g_9
f_2g_1	f_2g_2	f_2g_3	f_2g_4	f_2g_5	f_2g_6	f_2g_7	f_2g_8	f_2g_9
f_3g_1	f_3g_2	f_3g_3	f_3g_4	f_3g_5	f_3g_6	f_3g_7	f_3g_8	f_3g_9
f_4g_1	f_4g_2	f_4g_3	f_4g_4	f_4g_5	f_4g_6	f_4g_7	f_4g_8	f_4g_9
f_5g_1	f_5g_2	f_5g_3	f_5g_4	f_5g_5	f_5g_6	f_5g_7	f_5g_8	f_5g_9
f_6g_1	f_6g_2	f_6g_3	f_6g_4	f_6g_5	f_6g_6	f_6g_7	f_6g_8	f_6g_9
f_7g_1	f_7g_2	f_7g_3	f_7g_4	f_7g_5	f_7g_6	f_7g_7	f_7g_8	f_7g_9
f_8g_1	f_8g_2	f_8g_3	f_8g_4	f_8g_5	f_8g_6	f_8g_7	f_8g_8	f_8g_9
f_9g_1	f_9g_2	f_9g_3	f_9g_4	f_9g_5	f_4g_6	f_9g_7	f_9g_8	f_9g_9

(d) **FG**, the product of **F** and **G**, approximates **L** through NMF

Figure 3.5: Matrices representing the input and output of the content adaptive parallax barrier process.

as seen in Figure 3.6a. Figure 3.6b, demonstrates the possible alignment of pixels. First we ensure that the center view is produced when all of the pixels are aligned to their counterpart in the other mask, f_1 to g_1, f_2 to g_2 all the way to f_5 and g_5 in this case. When the user moves right the pixels between masks no longer align with their counterparts, instead the pixels from the back mask align to the counterparts of their right neighour and instead of viewing the center image we want to view the next image to the right. Similarly, when we move left the back mask pixels align with the counterpart of their left neighbour and we want to view the image to the left. By simply inspecting how the masks align we can see which entries of **FG** correspond to image pixels from the light field. Figure 3.6c demonstrates the resulting light field matrix **L**.

So for a 2D light field we can construct **L** using the following constraints, which were illustrated in 3.6c.

$$\mathbf{L}(i,k) = \begin{cases} I_{(\frac{u}{2}+i-kl)}p_{(k)}, & \frac{-u}{2} \leq i-k \leq \frac{u}{2}, \\ 0, & otherwise, \end{cases}$$

Where, $0 \leq i,k \leq s$.

3.3.2 Practical 4D Example

Before we can go deeper with a four dimensional example, there must be a clear understanding of how pixel alignment occurs, in terms of which pixels align with each other and how the alignment changes. Pixel alignment changes by physically changing the viewing position. In other words, as the Viewer's head is moved to different positions, a different pixel alignment may occur. Figure 3.7 helps to understand which pixels align to with another. One of the pixels from the front mask, **F**', will always align to the center pixel of the back mask, g_5 of **G**'. The other pixels that align do so in respect to g_5, and the pixel from the front mask which it is aligned to. In the case of Figure 3.7, the center pixel of the front mask, f_5, is aligned to the center pixel of the back mask, g_5. The consequence

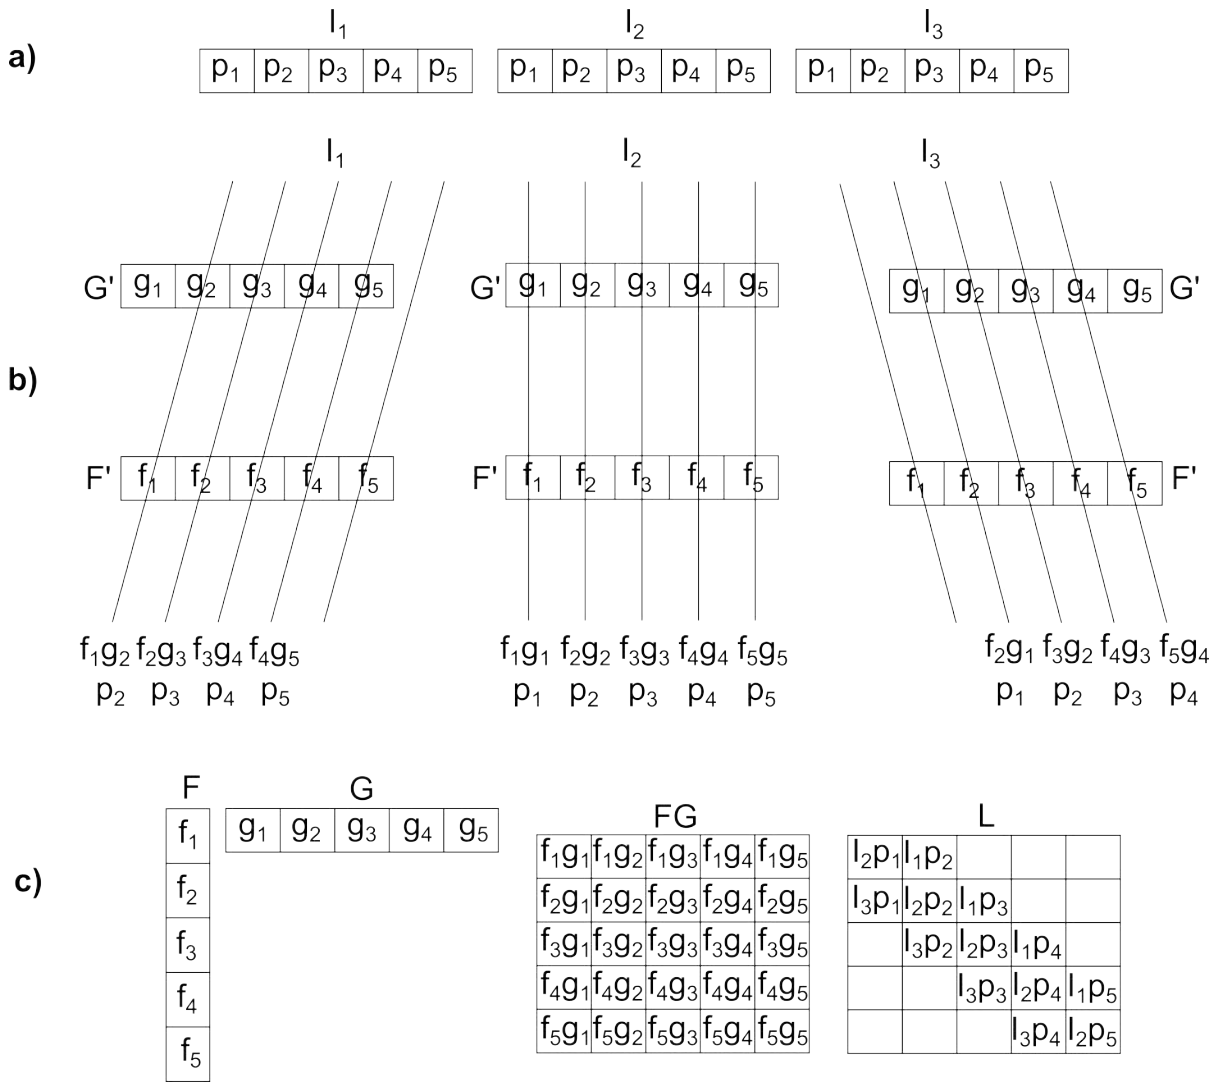

Figure 3.6: A two-dimensional illustration of display behaviour and construction of **L** (a) The input 2D light field. (b) The alignment of pixels between the front and back masks producing different views, where these alignments are produced by the viewer physically moving their head around the display. (c) The resulting matrices from the illustrated behaviour.

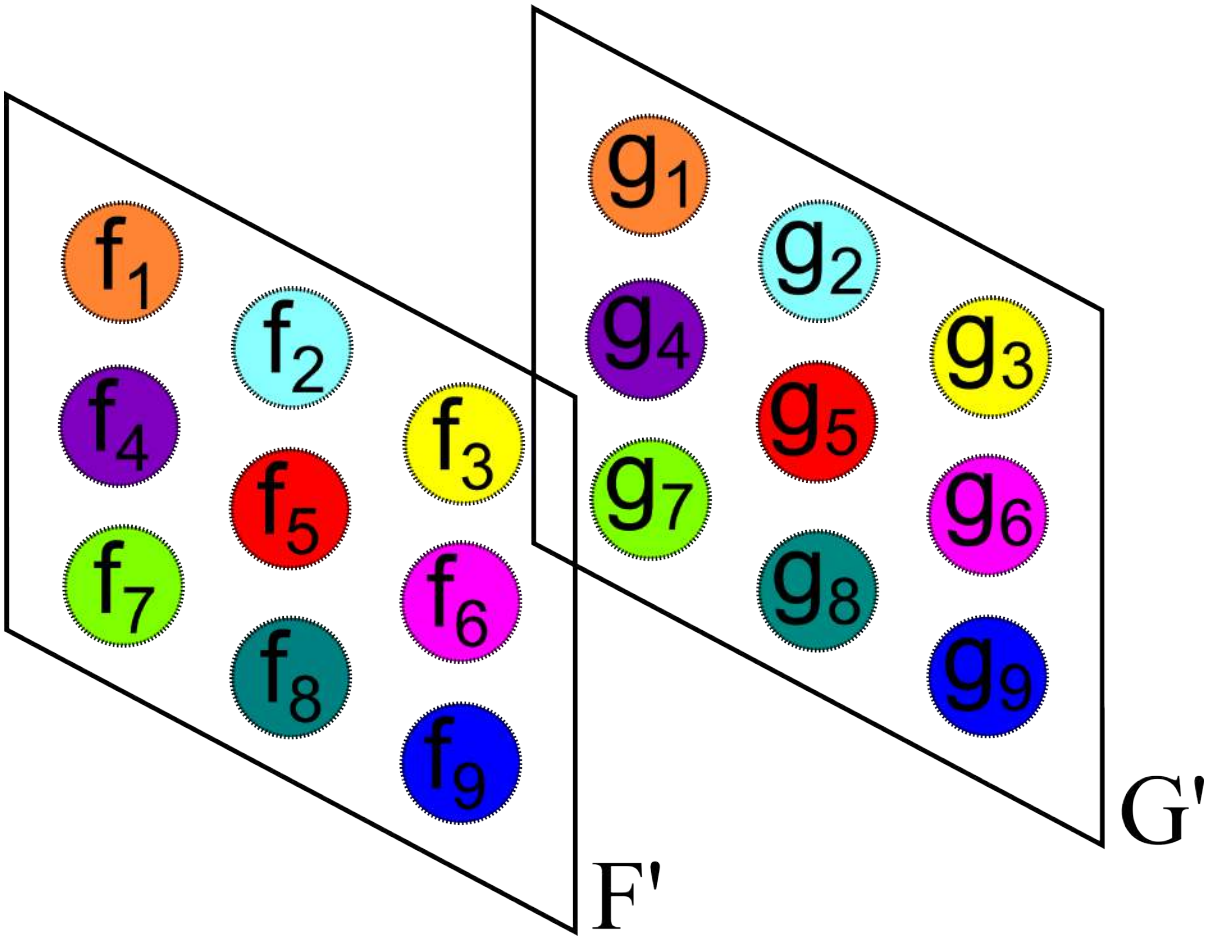

Figure 3.7: Visualization helping to describe which pixels align under different circumstances. The values displayed correspond to those present in Figure 3.5.

is that the pixel to the right of the center front mask pixel, f_6, is aligned to the pixel to the right, of the center back mask pixel, g_6. It is is important to note that the pixel aligned to g_5, will not always be f_5. For every different image the display is set to show, there will be a different pixel from the front mask aligning to g_5. No matter which pixel from the front mask aligns to the center back mask pixel, the neighboring pixels, of the front mask pixel, will align to the neighboring pixels of the centre back mask pixel. An example is that if f_4 is aligned to g_5 then the pixel immediately to the right of f_4, which is f_5, will align to the pixel immediately to the right of g_5, which is g_6.

Figure 3.8 illustrates the alignment of pixels between the front mask and back mask as well as the intended image to be displayed based on the relative viewing direction.

Relative viewing direction refers to the direction necessary to produce the desired alignment of pixels, where the center view is produced with an exact alignment of the masks. Other views are produced by shifting which pixels of the front mask align with the back mask. In reference to the centre view alignment, images adjacent to the center will have the pixel alignments shifted by one, images two away from the center will have the pixel alignments shifted by two and so on. The aligned pixels and their corresponding values from the light field are displayed in Figure 3.9. The viewing direction possibilities are determined by how the physical system is setup and is addressed in Section 3.5.

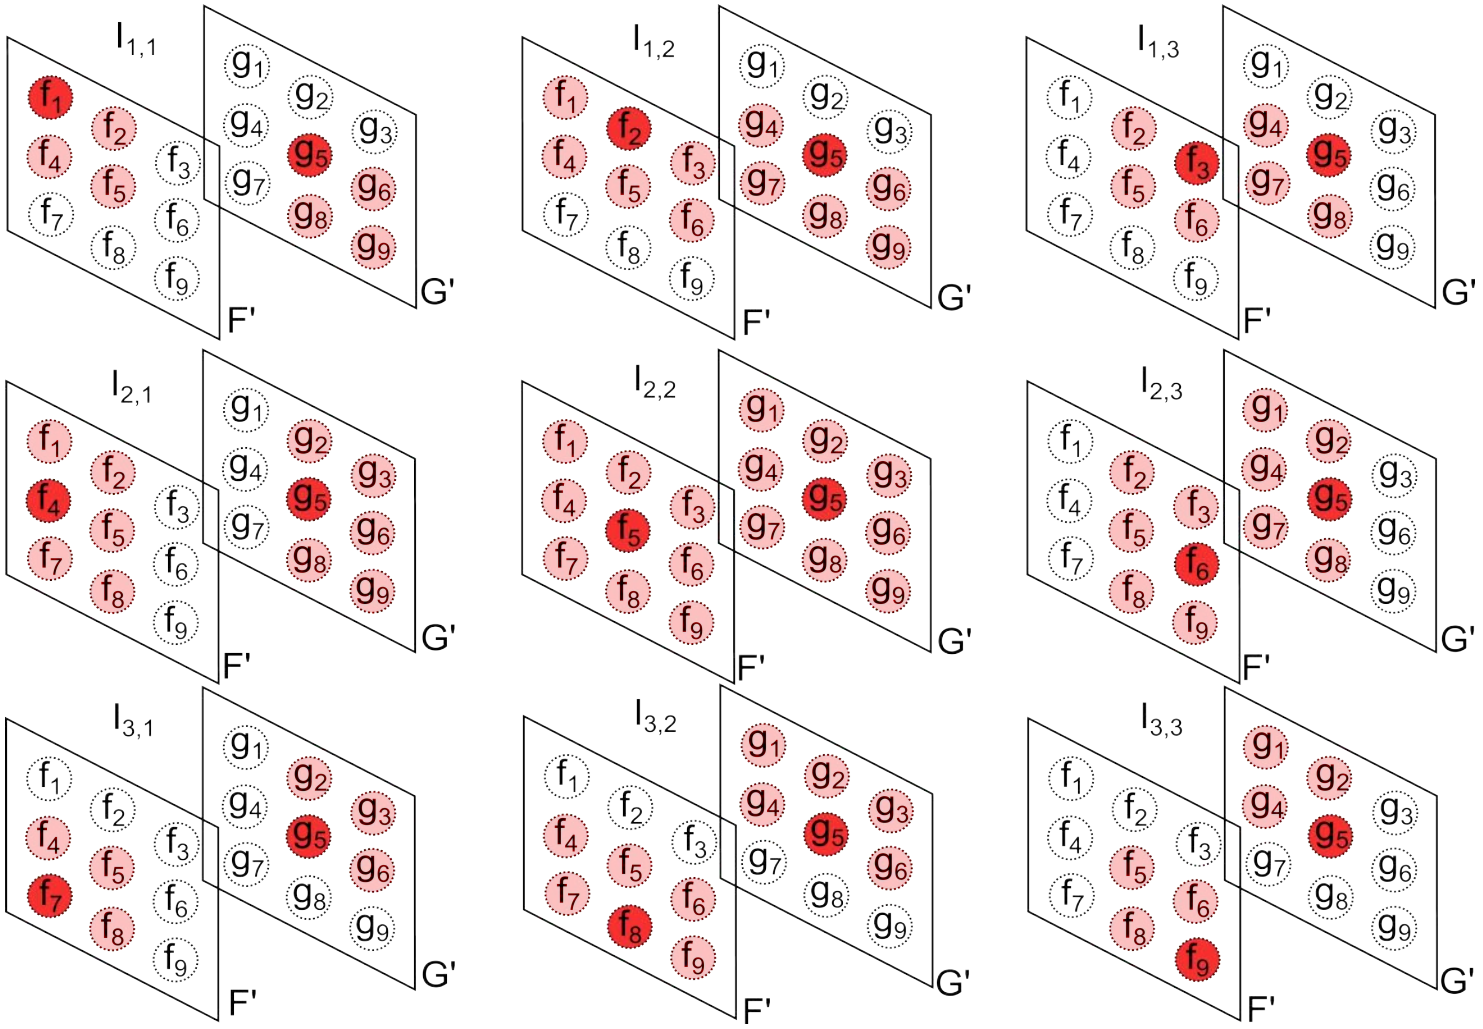

Figure 3.8: Visualization of how the display layers align to produce desired images. The circles represent pixels. The red pixels align with each other and signify the centre pixel of the image being displayed. The pink pixels align with each other, relative to the red pixel. White pixels do not align properly with the other layer and do not display any useful values for that view. The values displayed correspond to those present in Figure 3.5.

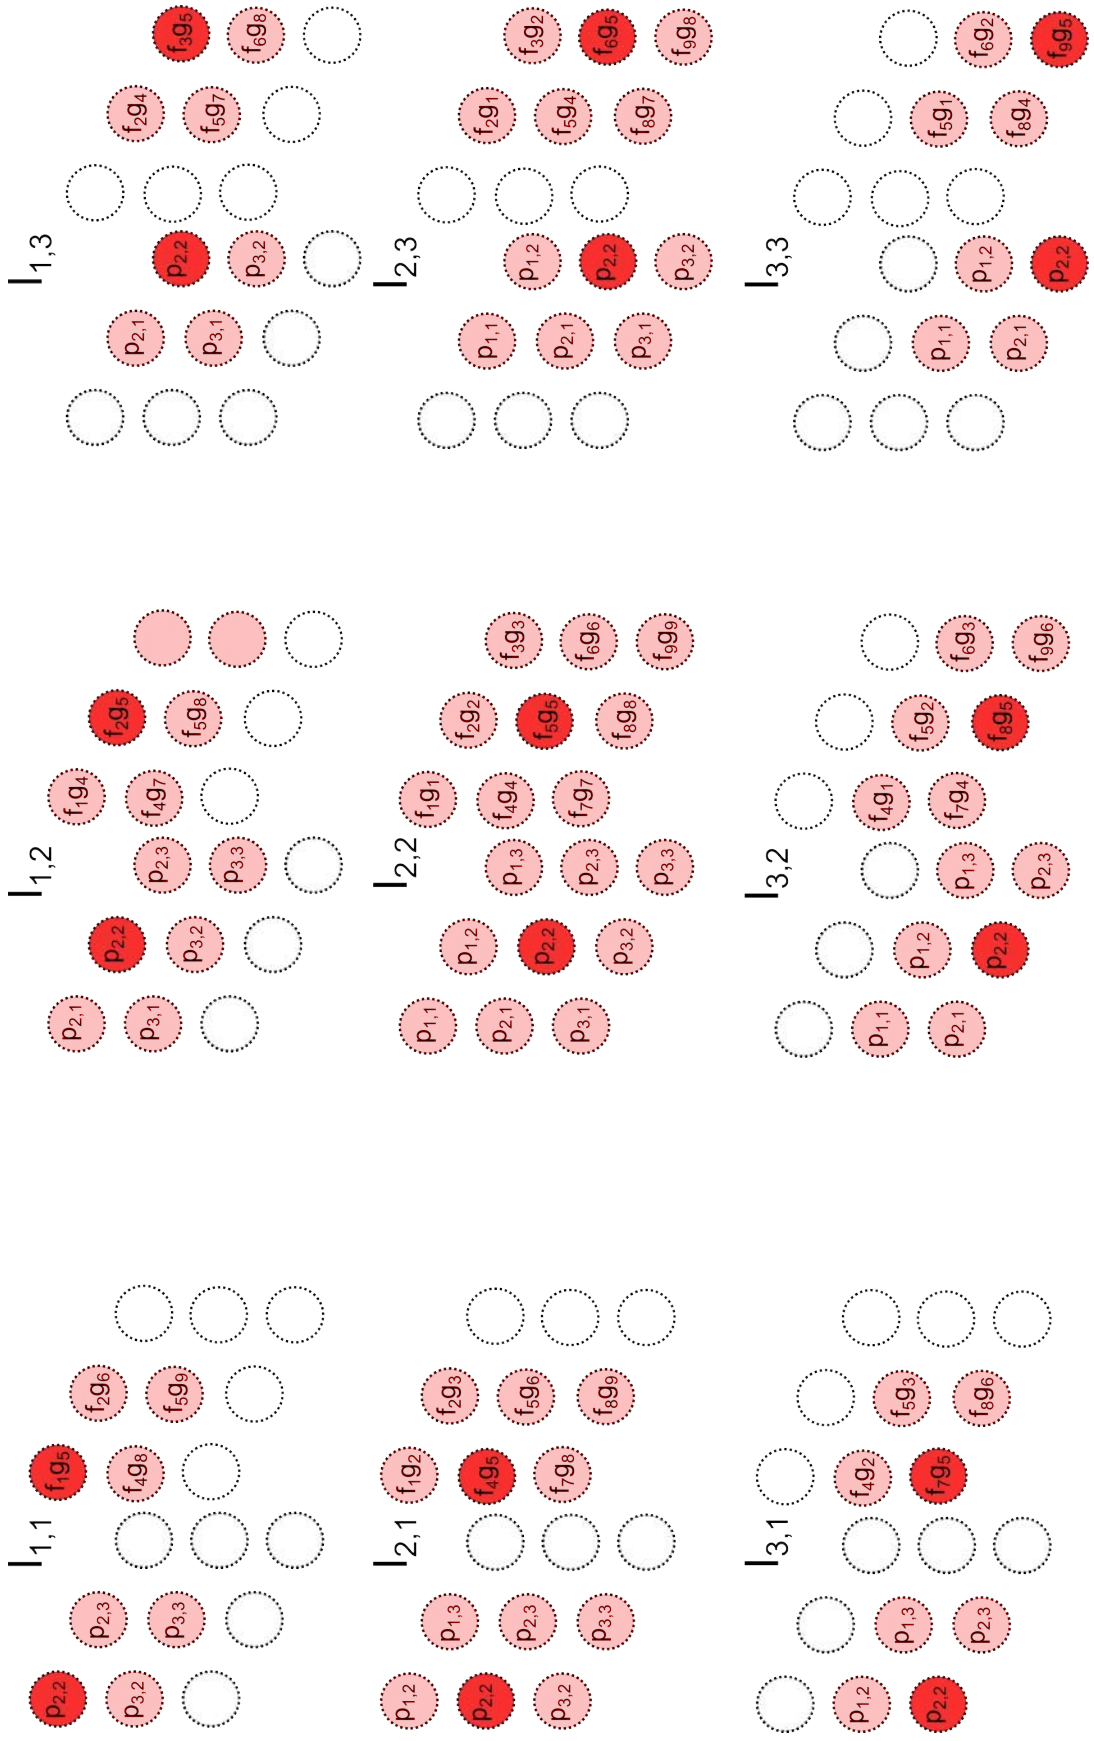

Figure 3.9: Visualization of how desired behaviour sets up a mapping between the input light field to produce the light field matrix to be approximated. The circles represent the pixels present for a particular light field image based on how the two layers are aligned. The alignments can be seen in Figure 3.8. The values displayed correspond to those present in Figure 3.5.

After observing the desired effect and inspecting the resulting light field matrix, a more formal method can be expressed for generating the light field matrix directly from the input light field. The first observation, from **L** in Figure 3.5, is that the pixel location does not change throughout columns. So even though the image being referenced changes, a column of **L** will always list the same pixel location. The second observation is that as you move left across columns or down through rows, the image being referenced inscreases by one. This means that a reduction in column value or an increase in row value each mean an increase in the image entry. The final observation is that the entire center view is listed across the diagonal of **L**. All of these observations put together produce the following equation for obtaining **L** from the input light field images.

$$\mathbf{L}(i,j,k,l) = \begin{cases} I_{(\frac{u}{2}+i-k,\frac{v}{2}+j-l)p(k,l)}, & \frac{-u}{2} \leq i-k \leq \frac{u}{2} \quad and \quad \frac{-v}{2} \leq j-l \leq \frac{v}{2}, \\ 0, & otherwise, \end{cases}$$

where $0 \leq i,k \leq S$, $0 \leq j,l \leq t$, and $\mathbf{L}(i,j,k,l)$ is equivalent to $\mathbf{L}(is+j, ks+l)$.

3.3.3 Pixel Barrier

During the implementation of content adaptive parallax barriers we discovered a novel light field sampling approach. The multi-view images are reordered as column vectors and placed in order to form a dense matrix. This creates an alternate application of the context adaptive parallax barrier which creates an opportunity to approximate the complete light field, as opposed to a down-sampled version, such as the one in the current approach. The viewer must be close enough to view the entire back plane, with each eye peering through a single pixel of the front plane as seen in Figure 3.11. This gives rise to some interesting artistic applications. One such application could be a form of interactive visual story-telling, which would work well in a museum or exhibit environment.

There are many drawbacks to this particular implementation, which include the mem-

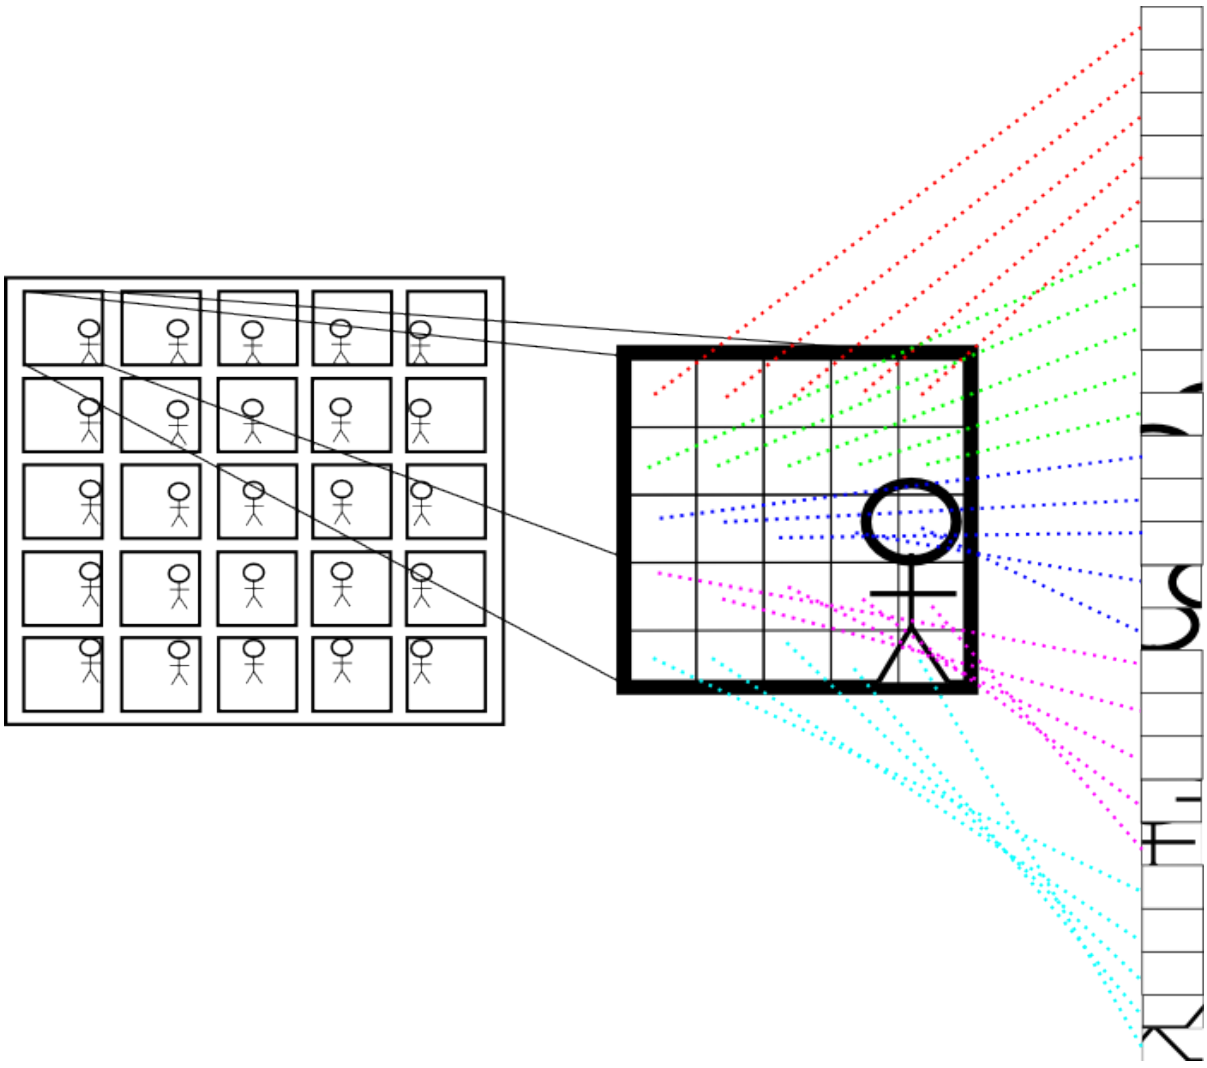

Figure 3.10: A depiction of reordering a light field into the proper matrix form to construct a pixel barrier. (left) the given light field matrix (middle) A single image from the light field (right) The image reorderedinto a column vector. This process is repeated for all images in the light field and the columns make up the input matrix for factorization.

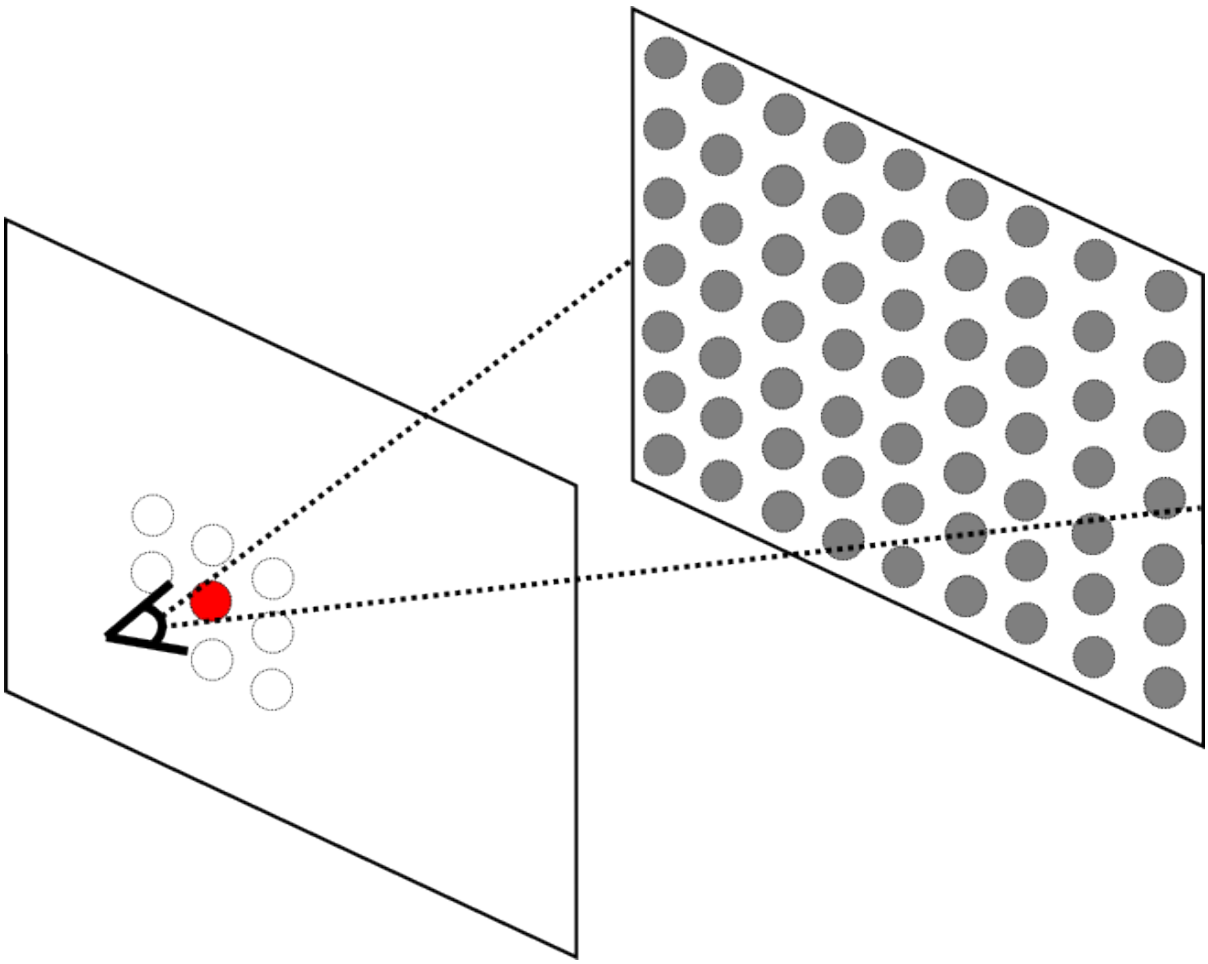

Figure 3.11: An illustration of a single eye using the pixel barrier display. The entire back mask must be viewed through a single pixel of the front mask.

ory necessity during the matrix factorization process, the poor quality of the final result due to too many highly coupled variables and the fact that this does not allow for a similar unencumbered display to the content adaptive parallax barrier setup. This approach requires more computational power and memory due to its version of **L** being dense. The sparsity of content adaptive parallax barriers significantly reduces the necessary resources.

The final result would not be displayed on a television or monitor setup, because it creates a result where each front mask pixel is approximating a view across the entire back mask, as seen in Figure 3.11. The viewer would need to look at the back mask through an individual pixel of the front mask, which is an unrealistic use of a monitor setup. An alternative worth investigating in the future would be to use a LCD panel for the back plane and have a projector display the front plane onto some form of translucent media. Another idea is to back project the mask through the LCD panel, which may allow for the desired stereoscopic viewing and would have the added benefit of the light source from the projector.

3.4 Decomposition

Originally, a weighted version of the Lee & Seung non-negative matrix factorization method [47] was used to decompose the input matrix into **F** and **G** [45] due to its ease of implementation. Any NMF method could be used, but this must be done very carefully. The input matrix cannot simply be placed into any set of update rules and produce visually acceptable results. Due to the sparsity of **L**, non-weighted matrix factorization techniques attempt to approximate all the zeros, which results in an almost black image when viewing the masks in a time-multiplexed manner. The NMF procedure is attempting to optimize across an extremely large dataset and the sparsity of the matrix must be exploited. If a procedure is used that only attempts to approximate the non-zero values

of the input matrix, there will be a much higher fitness obtained on the desired values. For a more detailed look at NMF please refer to Appendix A.

The described sampling method encodes the content adaptive mask pairs into $\mathbf{F}$ and $\mathbf{G}$ as columns and rows, respectively. The columns of $\mathbf{F}$ and the rows of $\mathbf{G}$ must be separated and reordered into matrices of size $s \times t$. The number of mask pairs produced is equal to the factorization rank. A higher rank factorization allows for a better approximation of the light field. The reason for the better approximation can be seen in Figure 3.12 where it is shown that there are more values in $\mathbf{F}$ and $\mathbf{G}$ which means there are more values contributing to each entry of $\mathbf{FG}$. More values contributing to $\mathbf{FG}$ relaxes the constraints on any one entry in $\mathbf{F}$ or $\mathbf{G}$ due to the fact that, for a rank one approximation, each entry of $\mathbf{F}$ is used in an entire row of $\mathbf{FG}$ and consequently is used to approximate an entire row of $\mathbf{L}$, similarly for $\mathbf{G}$ except its entries are used in columns of $\mathbf{FG}$. A higher rank means that instead of having one value being stretched along an entire row, or column, multiple values share the task which relaxes the constraints for any one entry of $\mathbf{F}$ or $\mathbf{G}$. Sets of mask pairs are displayed rapidly and rely on the human visual system to integrate the different pairs through flicker fusion.

3.4.1 Weighted FNMA_e

The FNMA_e technique, described in Appendix A, is not well suited for use with content adaptive parallax barriers because it is unweighted which means it is designed to minimize all variables involved equally, which is shown in Chapter 5. This means the product of $\mathbf{F}$ and $\mathbf{G}$ will attempt to approximate all values in $\mathbf{L}$ without bias, even though there are many zero values that we are not interested in viewing on the display. Every additional value from $\mathbf{L}$ that needs to be approximated by $\mathbf{F}$ and $\mathbf{G}$ adds to the complexity of the matrix factorization. Individual values from $\mathbf{F}$ and $\mathbf{G}$ need to approximate multiple values from $\mathbf{L}$ so less $\mathbf{L}$ entries simplifies the problem and allows for a better approximation. We devised a weighted version of FNMA_e to produce masks that put emphasis on the

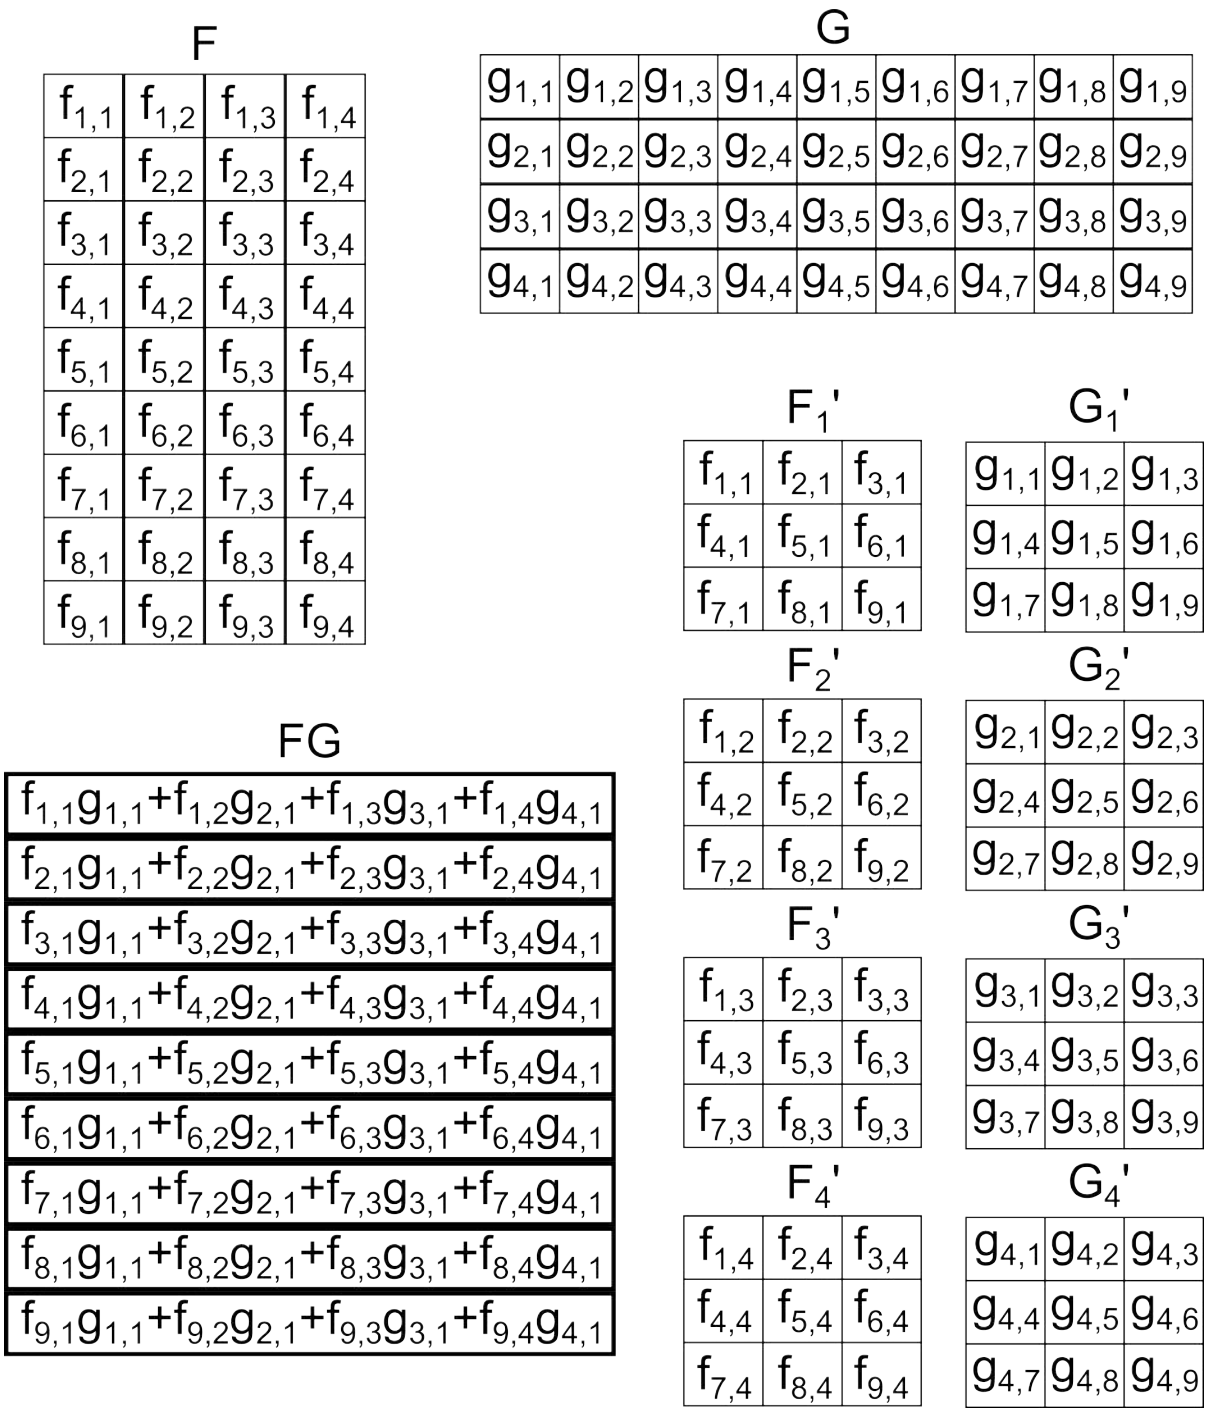

Figure 3.12: Matrices depicting the structure of **F** and **G** with a rank higher than one. (top-left) rank 4 **F** matrix. (top-right) rank 4 **G** matrix. (bottom-left) the first column of the **FG** matrix. (bottom-right) the four sets of mask pairs produced by the rank four decomposition.

desired light field entries. The objective function of weighted FNMA_e, or WFNMA_e, is:

$$\begin{aligned}\underset{p\leq\mathbf{F}(i,j),\mathbf{G}(i,j)\leq q}{\text{minimize}} \quad &= \frac{1}{2}||\mathbf{L}-\mathbf{FG}||_{\mathbf{W}}^2 \\ &= \frac{1}{2}\sum_{ij}[\mathbf{W}\circ(\mathbf{L}-\mathbf{FG})\circ(\mathbf{L}-\mathbf{FG})](i,j).\end{aligned}$$

Where, $p=0$ and $q=1$ for our application. The gradient calculation is modified to the following:

$$\nabla\mathbf{F} = ((\mathbf{FG}-\mathbf{L})\circ\mathbf{W})\mathbf{G}^T,$$
$$\nabla\mathbf{G} = \mathbf{F}^T((\mathbf{FG}-\mathbf{L})\circ\mathbf{W}).$$

The Hessian calculation and line search method needed to be modified, not only to apply the appropriate weights, but also to work with sparse matrices. While this is possible with dense matrices, in practice the matrix dimensions are so large that simply iterating through the dimensions is a very time consuming task. So a modification to leverage a sparse representation was necessary. The new Hessian update rule is

$$\begin{aligned}\mathbf{D}^{k+1} =&\mathbf{D}^k - \frac{\mathbf{D}^k\mathbf{A}^T(\mathbf{W}\circ\mathbf{B})\mathbf{X}^T+\mathbf{X}(\mathbf{W}\circ\mathbf{B}^T)\mathbf{A}\mathbf{D}^k}{\mathbf{W}\circ\mathbf{B}^T\mathbf{B}} \\ &+\left(1+\frac{(\mathbf{W}\circ\mathbf{B}^T)\mathbf{A}\mathbf{D}^k\mathbf{A}^T(\mathbf{W}\circ\mathbf{B})}{\mathbf{W}\circ\mathbf{B}^T\mathbf{B}}\right)\frac{\mathbf{X}\mathbf{X}^T}{\mathbf{W}\circ\mathbf{B}^T\mathbf{B}},\end{aligned}$$

where $\mathbf{X}=\mathbf{X}^{k+1}-\mathbf{X}^k$ and $\mathbf{B}=\mathbf{A}\mathbf{X}$ [7, 43]. The modified line search is

$$\alpha = \frac{-\mathbf{D}^T\overline{\mathbf{A}}^T(\mathbf{W}\circ(\overline{\mathbf{A}}\mathbf{X}-\mathbf{L}))}{\mathbf{D}^T\overline{\mathbf{A}}^T\overline{\mathbf{A}}\mathbf{D}}\circ\mathbf{I},$$

where α is a diagonal matrix containing search directions for each column of $\mathbf{X}$.

(a) A front view of the dual layer LCD display.

(b) A side view of the dual layer LCD display.

(c) The front LCD after removed from the original monitor.

(d) The removed backlight from the front LCD.

Figure 3.13: The dual layer LCD display used for content adaptive parallax barriers, during construction.

3.5 Display

The hardware is easily accessible as the LCD panels can be removed from an LCD monitor or ordered seperately. The main difficulty during construction is removing the polarization layers, which are manufactured onto the panels themselves. Figure 3.13 shows some stages during the construction of the display. Figure 3.13c shows one of the monitors that were broken during construction. With a finished device the mask pairs produced through matrix factorization can be displayed directly onto their corresponding LCD panel. Some considerations for the display are how far apart the panels are from each other and the capabilities of the panel itself in terms of refresh rate and resolution.

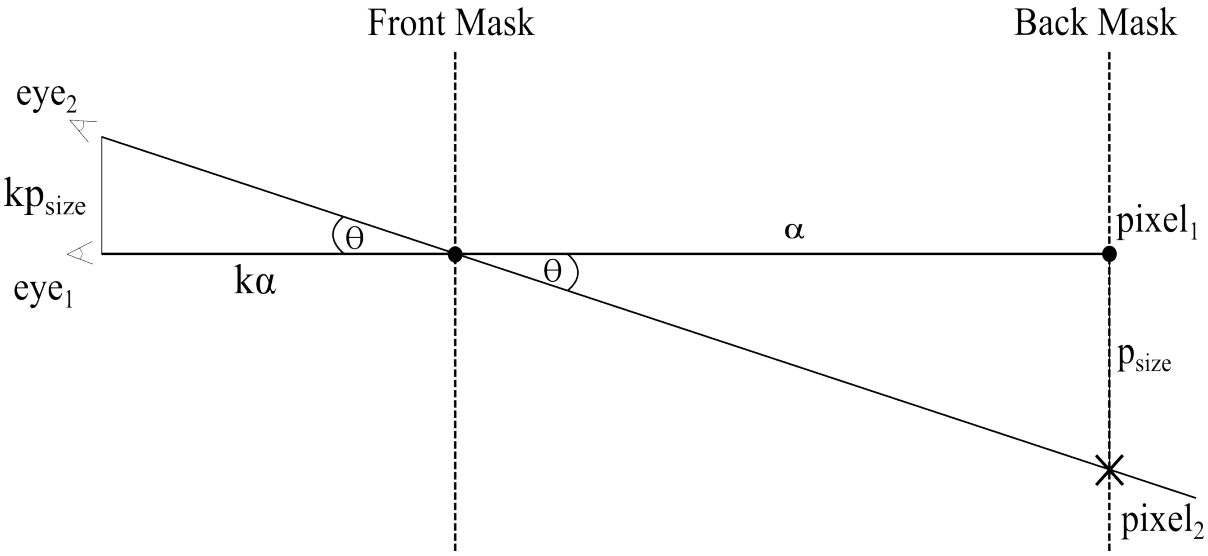

Figure 3.14: Visualization of the trigonometry involved in determining proper plane spacing.

3.5.1 Plane Spacing

Content adaptive parallax barriers work through the proper alignment of pixels on the front mask to the pixels on the back mask. To switch from one view to another in a given set of masks, the viewer must physically move so that the pixels between the front and back LCD panels align differently. The distance that must be moved is a function of the pixel size on the LCD panels, the distance between the panels and the distance the viewer is from the display. Figure 3.14 illustrates how to determine a proper spacing between the planes.

This is all that is required for content adaptive parallax barriers to function as proposed. Unfortunately, this is not enough to produce a proper autostereoscopic display. For proper stereo other factors must be accounted for such as the interpupillary distance, which is the distance between a user's eyes [19].

For conventional parallax barrier setups only two views are necessary for stereo vision. Multiple viewers are supported by the fact that each of these views repeat. So assuming each viewer is in a proper viewing zone, they will perceive the same 3D image. Referring

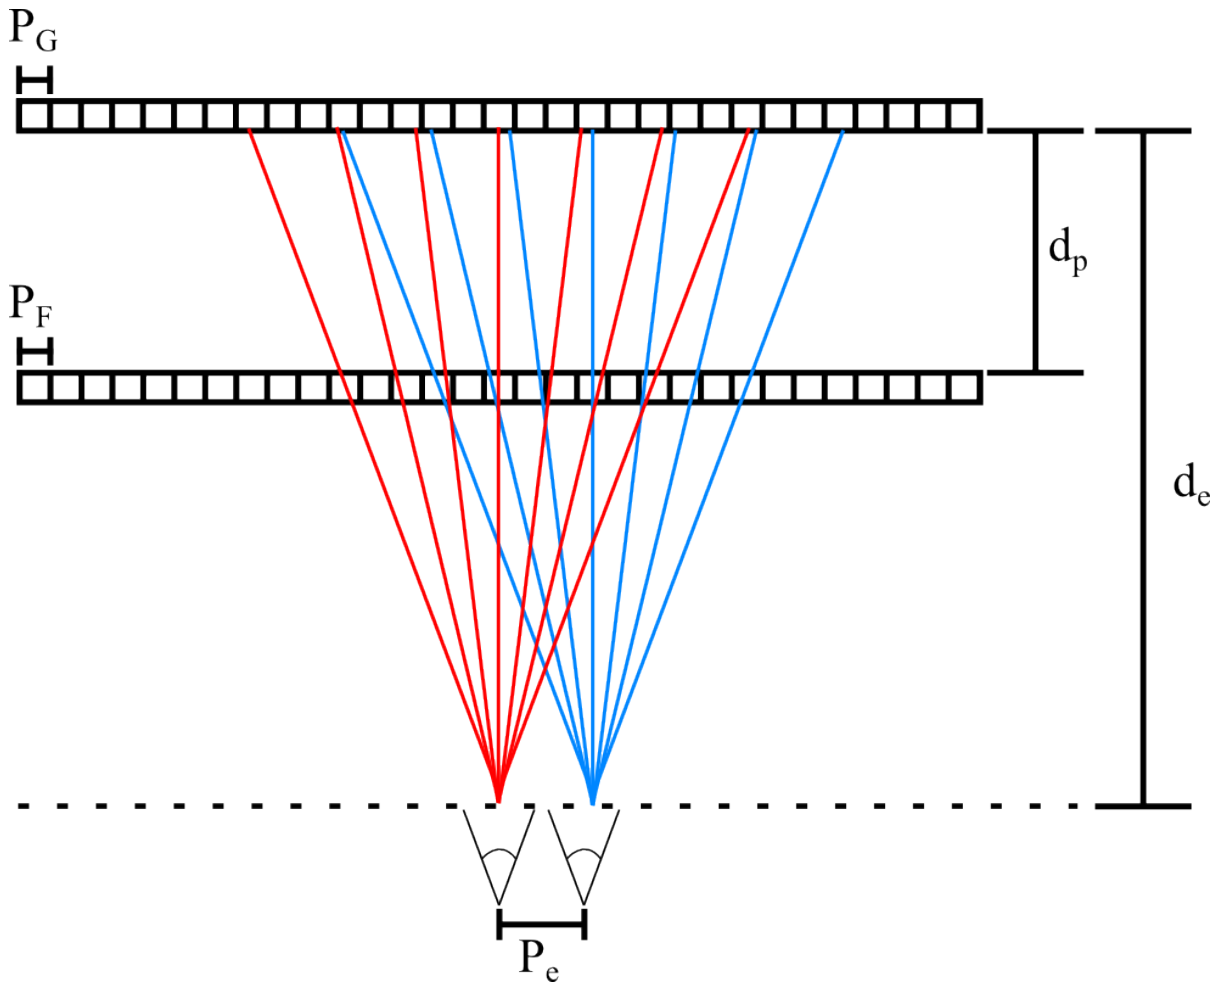

Figure 3.15: The basic parallax barrier setup.

to Figure 3.15, the proper viewing distance, d_e, is given by the formula

$$d_e = \frac{2P_G d_p}{2P_G - P_F},$$

where P_F is the pitch for the parallax barrier, P_G is the pitch for the display device and d_p is the distance between the barrier and the display [75]. Addtionally,

$$P_F = \frac{2P_e P_G}{P_e + P_G},$$

relates the interpupillary distance, P_e, to the other values, which is critical since the viewing zones must be designed for the proper eye spacing [75]. Content adaptive parallax barriers do not have the same fundamental structure as conventional parallax barriers so they cannot, without proper planning, support multiple users as easily. For more on plane spacing and other factors that effect content adaptive parallax barriers from being autostereoscopic refer to section 3.5.4.

3.5.2 Resolution

The resolution of the LCD panels is an important factor alongside the panel size since that will define the pixel pitch, or size of each pixel. Pixel pitch is the largest factor in determining optimal viewing distance due to its role in having proper pixel alignment which is discussed in Section 3.5.4. Resolution is also important because it dictates the maximum number of views that can possibly be displayed. As views move away from the center view, more pixels of the LCD panels become unaligned, where pixels from the front panel do not align with pixels from the back panel. The consequence is that eventually all of the pixels become unaligned and no more images can possibly be displayed, which is why the LCD panel resolution dictates the maximum number of possible views that can be displayed. Ideally, the number of possible views should be maximized for the best viewing experience, but this is limited by the refresh rate of the LCD panels and

potentially the human visual system which is discussed in section 6.1.2. An increased resolution of the light field across any of u, v, s or t will increase the memory requirements and the compute time for masks, but increasing resolution will lead to a better viewing experience.

3.5.3 Refresh Rate

A panel's refresh rate determines how many different images can be displayed in a certain time frame. Hertz (Hz) is the unit of measure for refresh rate, where 1Hz means a panel can display only a single image per second. The refresh rate limits how many mask pairs can be displayed without creating disomcfort for the viewer. If the refresh rate is too low there will be a visual flicker, which can be uncomfortable.

Assuming an unlimited refresh rate on the LCD panels, the maximum number of mask pairs has yet to be determined. The research concerning flicker fusion has been focused on what a person can identify or remember in some set of data. For the purposes of content adaptive parallax barriers, the set of mask pairs represents a single image, or frame in the case of video. The human visual system integrates across all the mask pairs to determine what exactly we perceive as the displayed image. So in this case it is not about what we can identify within the mask pairs, it is about what the human visual system can perceive and there has been limited research done to determine how fast our eyes can perceive data.

3.5.4 Autostereoscopic Content Adaptive Parallax Barriers

Content adaptive parallax barriers, as presented, do not support proper stereo vision for multiple viewers without assuming that the input light field images were created specifically for this purpose. This assumption would be ill-conceived as current light fields are not generated for this specific purpose in mind. Additionally, traditional parallax barriers are structurally different than content adaptive parallax barriers. While designing a con-

ventional parallax barrier display one of the choices is the pitch, or spacing between light permitting slits. Content adaptive parallax barriers do not get the choice of the pitch for the front mask as it is defined by the LCD panel which means it will also normally be equal to the pitch of the back mask, assuming the same LCD panel is used for each layer.

For a light field to be used in an autostereoscopic manner, by the content adaptive parallax barrier system, there are two requirements that must be met. The first requirement is that the cameras used to capture the light field must be spaced with interpupillary distance (IPD) in mind. IPD is the distance between an individual's two pupils. All parallax barrier systems result in set viewing zones from which the best quality image will be seen. If these zones are not spaced according to the IPD that will result in a lower maximum quality result because there will be an increase in perceived noise. The second requirement is that the difference in viewing angle, between the light field images, must be less than or equal to the angle between the viewing zones of the display. Figure 3.16 shows a comparison of the viewing angles between light field generation and viewing on the display. If the viewing angles between light field images are less than the angles between display viewing zones, the images viewed result in negative parallax. Negative parallax occurs when the perceived image is closer to the eyes than the display, which can lead to eye strain [4].

Each of the requirements, of the input light field, can be mitigated by Image Based Rendering (IBR). IBR can be used to completely meet these requirements, but only at the sacrifice of resolution. If image based rendering is used to meet the viewing angle requirement it will reduce the angular resolution of what is shown on the display. Artificially reducing the angular resolution will allow for viewing stereo images comfortably, without eye strain, but creates a situation where the objects being viewed will appear further than intended due to less disparity between the images [69].

Due to the structural differences between traditional and content adaptive parallax

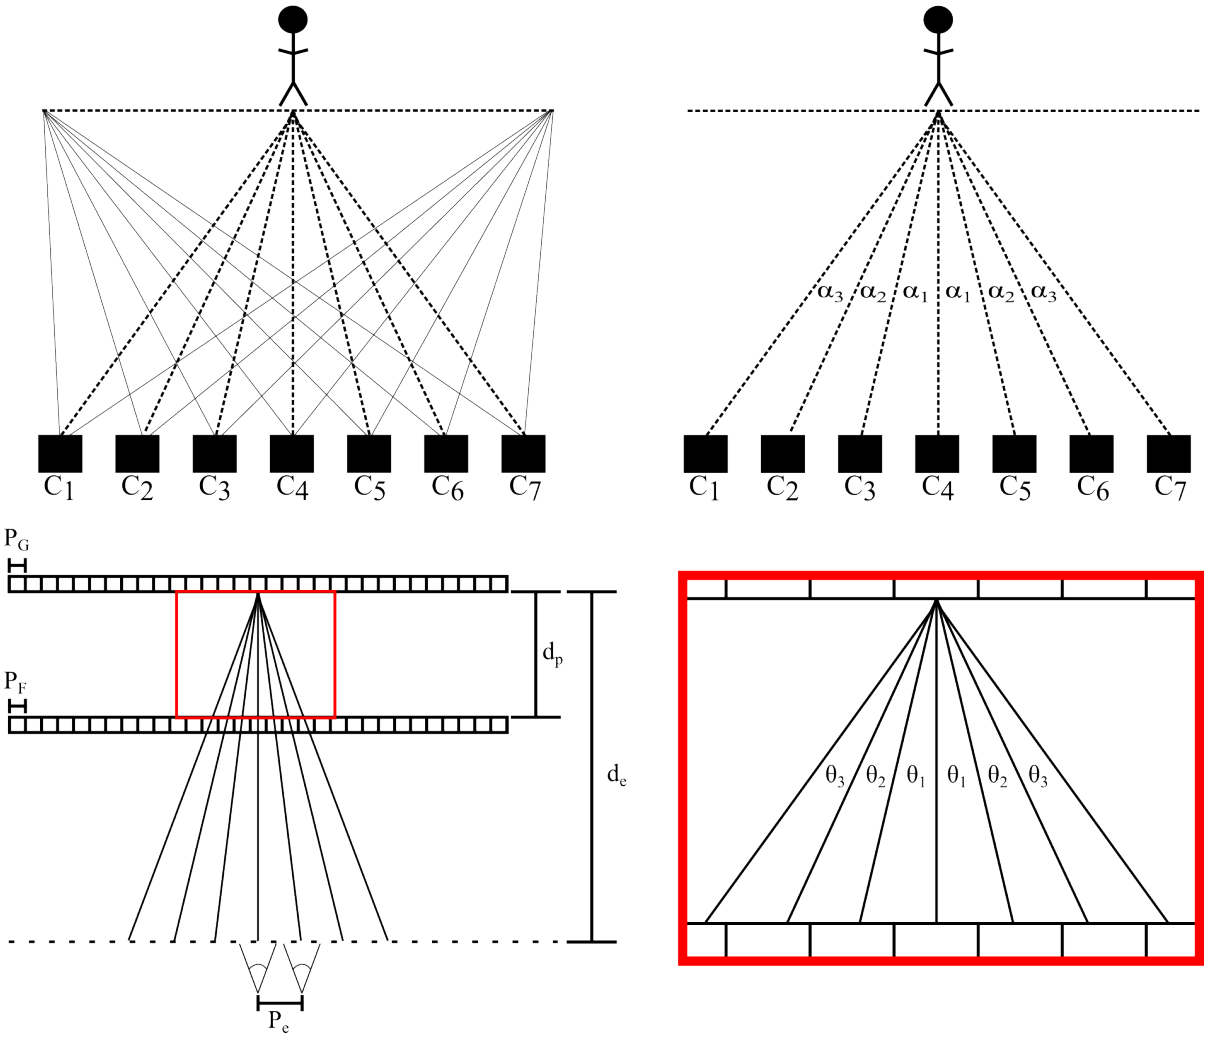

Figure 3.16: Depicts a comparison between the viewing angles used when generating the light field images, α_n, and the viewing angles present when using the multi-layer display, θ_n. It is acceptable for these angles to be different, but this would imply that either the percieved images would be disproportionate or some method to compensate for the difference must be employed.

barriers, an examination of the parameters for proper stereo viewing is necessary. The most important factor when designing an autostereoscopic display is accounting for IPD. In terms of parallax barriers, the intended viewing areas must be spaced at a distance equal to the IPD of the viewer. Generally, the design of parallax barrier systems accounts for the average or median IPD, but a device can easily be adapted to an individual. Average IPD values, for adults, are between 50mm and 75mm, where the median is 63mm [19].

To determine the actual distance between viewing zones requires knowledge of the LCD panels being used for the display, but relative distances can be examined initially. Figure 3.17 shows some of the parameters necessary for calculations regarding the viewing zones. The most significant parameters shown are the distance between planes, d_p, the number of pixels, N_p, and the pixel pitch, P_F and P_G, or the pixel size.We assume, for ease of analysis, that the front and back LCD panels are of the same make which means P_F and P_G are equal, so moving forward only P_F will be discussed. In Figure 3.17, locations on the LCD panel are addressed by their distance from the middle of the centre pixel. Movement along a panel is set as the x-axis, where distances are stated in units of P_F and movement towards or away from the panels is set as being on the z-axis, where distances are stated in units of d_p. Now a relative coordinate system is established where the origin is the middle of the centre pixel of the back LCD panel.

Content adaptive parallax barriers work under an assumption of viewing with an orthographic projection, or that any viewers are distant enough from the display that all of the pixels of the back panel are aligned to pixels of the front panel along the exact same angle. This results in the relative distance, between a viewer and the display, being much larger than the inter-display distance, or the distance between LCD panels. Figure 3.18 demonstrates how someone would need to perceive the different views being shown by the display. This assumption does not relate to reality because humans do not see in an orthographic manner, instead we see a perspective projection. The expected alignment

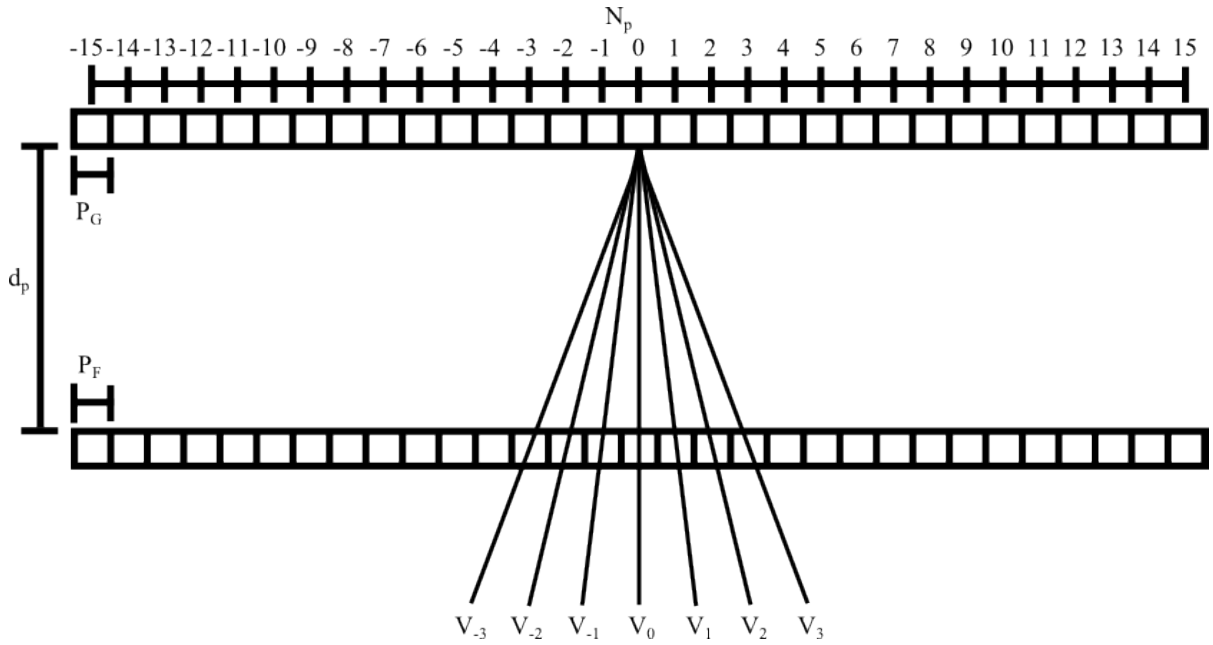

Figure 3.17: A set of one dimensional LCD panels help to describe a coordinate system that allows for a theoretical analysis of the content adaptive parallax barrier setup. There happens to be 31 pixels in this example, but there can be an arbitrary number.

angle is determined by which pixel on the front panel is aligned with the center pixel of the back mask. The pixels that will have the largest disparity from the expected alignment angle will be the outer most pixels. For this reason, only the outer most pixels need to be discussed because all of the others will be in a more favourable alignment.

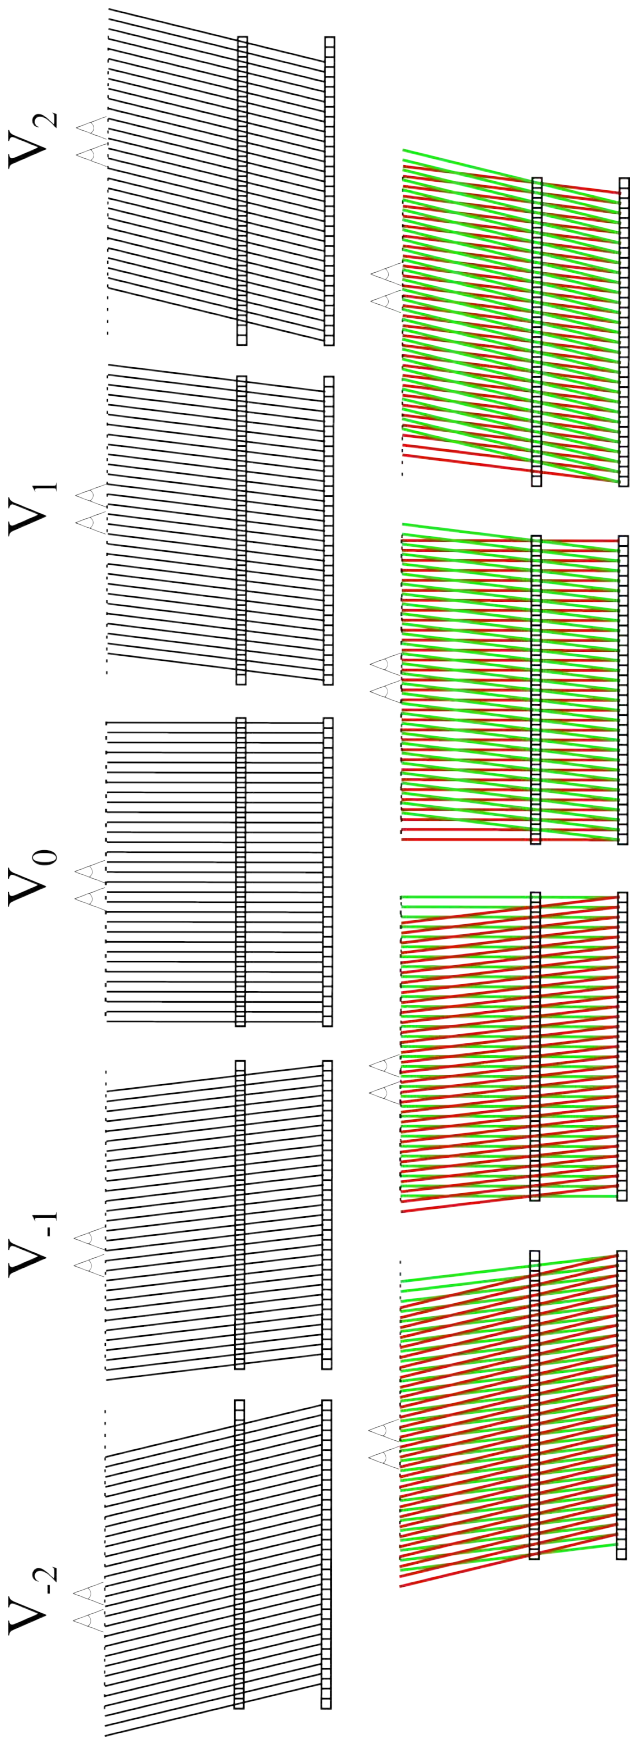

Figure 3.18: Shows the orthographic assumption of content adaptive parallax barriers. (top) Each view is seen in an orthographic manner from a specific viewing angle (bottom) Each of the views is one piece of a stereo pair where one view will be seen by the left eye and the other by the right eye.

The goal is to determine where the viewing zones occur for each image being displayed. For this reason we want to determine the start of these zones or the point of minimum viewing distance, $\vec{\boldsymbol{m}}$, for each view, $V(n)$. To determine where $\vec{\boldsymbol{m}}$ occurs we need to define an alignment behaviour for the display. We will consider a one dimensional display, where each panel can be thought of as a horizontal row of pixels. When a pixel on the back panel aligns with a pixel on the front panel, a ray is defined by the position of each pixel. For each view there is a different set of rays, as the pixels from the back panel align with different pixels from the front panel. Lets define the leftmost ray for $V(n)$, or the ray formed by the alignment of the leftmost valid pixel on the back panel and the leftmost valid pixel on the front panel, as $\vec{\boldsymbol{a}}_L$. Similarily we define $\vec{\boldsymbol{a}}_R$ is defined as the rightmost ray for $V(n)$. If these rays intersect the center of their corresponding pixels, on each panel, an orthographic system would be created, which is unrealistic. We need $\vec{\boldsymbol{a}}_L$ and $\vec{\boldsymbol{a}}_R$ to converge and the point of convergence is $\vec{\boldsymbol{m}}$. The closest possible $\vec{\boldsymbol{m}}$ occurs when $\vec{\boldsymbol{a}}_L$ passes through the leftmost point on the back panel's pixel and the rightmost point on the front panel's pixel and when $\vec{\boldsymbol{a}}_R$ passes through the rightmost point on the back panel's pixel and the leftmost point on the front panel's pixel, which is shown in Figure 3.19. This behaviour results in $\vec{\boldsymbol{a}}_L$ and $\vec{\boldsymbol{a}}_R$ being in the worst valid alignment possible, where any worse would actually be part of another view, $V(n+1)$ or $V(n-1)$. Figure 3.20 demonstrates how the described alignment behaviour allows for the calculation of $\vec{\boldsymbol{m}}$. Now we can use the established coordinate system to derive formulas for resulting IPDs and required viewing distances, where viewing distance is the distance between $\vec{\boldsymbol{m}}$ and the centre pixel of the back LCD panel.

With our coordinate system a simple ray intersection calculation is used to determine each $\vec{\boldsymbol{m}}$. With the location of each $\vec{\boldsymbol{m}}$, the distance between $\vec{\boldsymbol{m}}$ and $\vec{\boldsymbol{m}}(n+1)$ is calculated using the Euclidean distance formula $d = \sqrt{x^2 + y^2}$. If all of the distances between adjacent minimum viewing zones are within an acceptable range of the IPD then the display will produce autostereoscopic results with minimal artifacts. The minimum

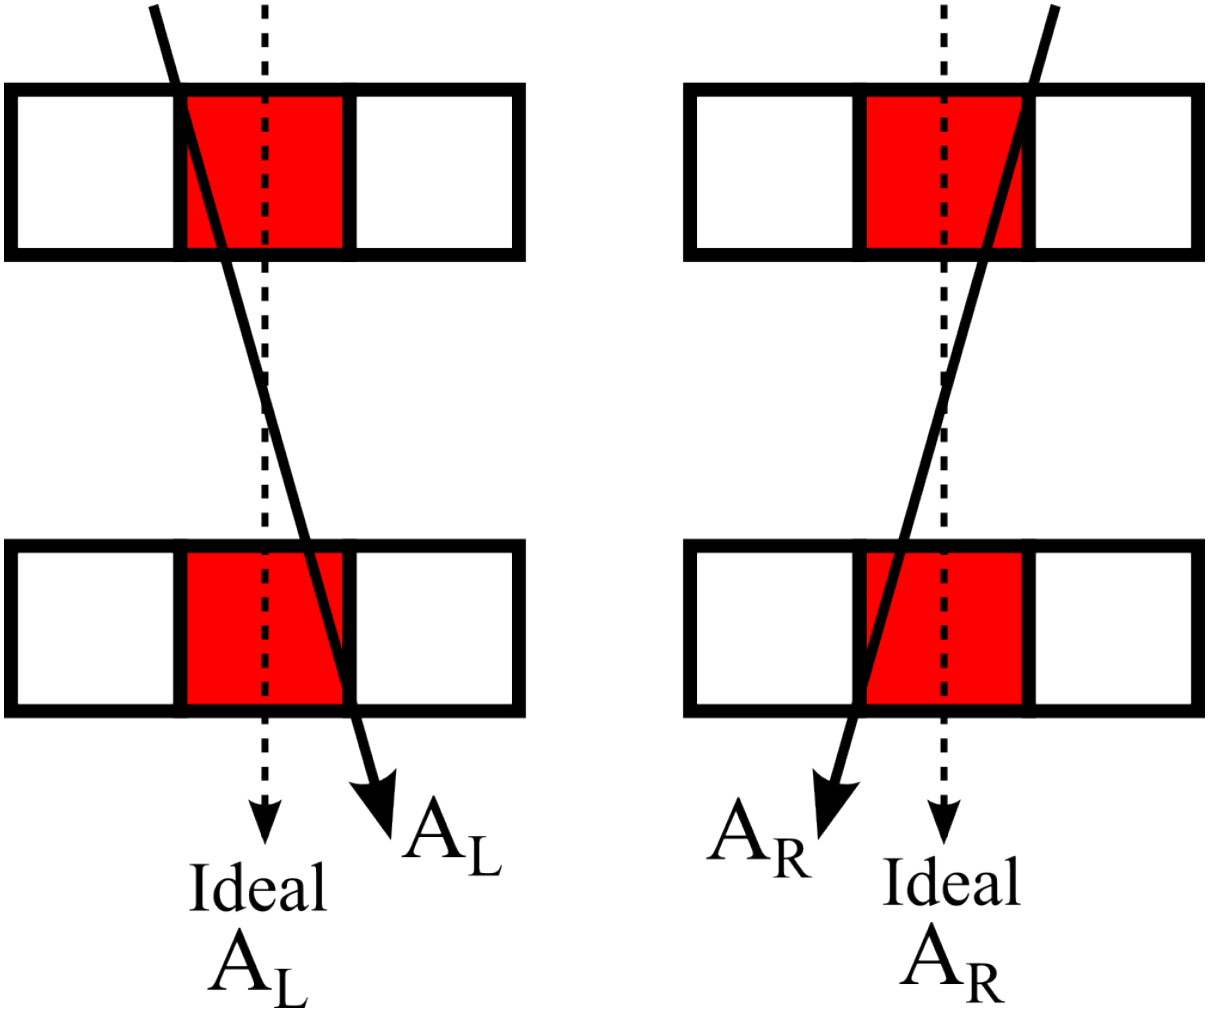

Figure 3.19: Shows how the leftmost ray, $\vec{a}_L$, and the rightmost ray, $\vec{a}_R$, intersect the panel's pixels for the analysis of the minimum distance viewing zones.

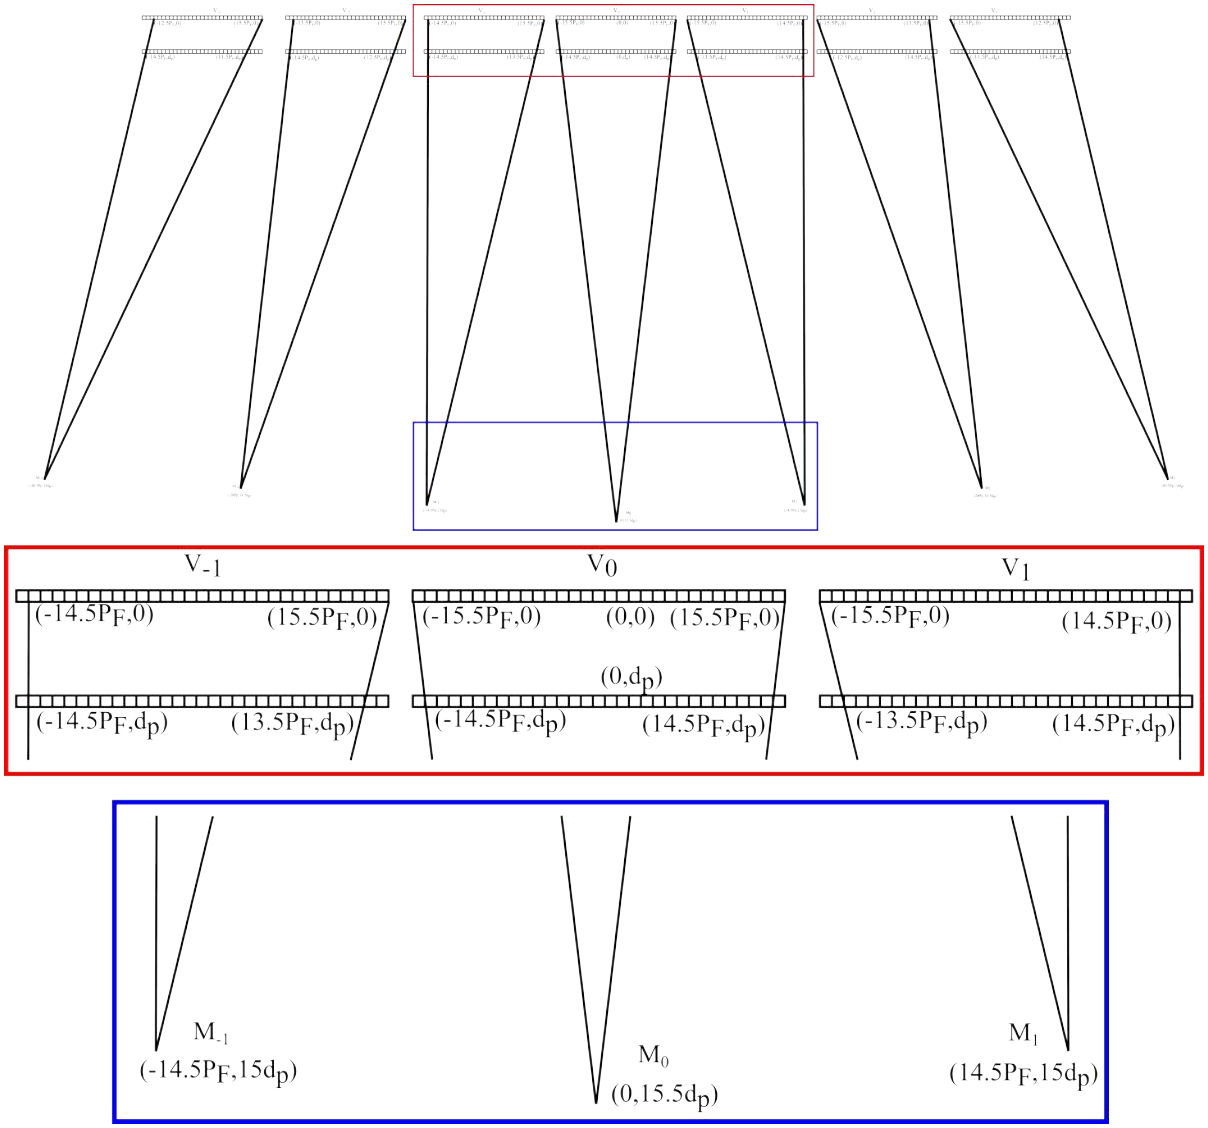

Figure 3.20: Shows the practical expected alignment of content adaptive parallax barriers, where each $V(n)$ (top) corresponds to the theoretical behaviour of those shown in 3.18. The points where the rays intersect the planes are described by the established coordinate system (middle). The minimum viewing distance points, $\vec{\boldsymbol{m}}$, are calculated by intersecting the outer most alignment rays (bottom). The set of all $\vec{\boldsymbol{m}}$ for corresponding $V(n)$ (right).

required viewing distance for z viewing zone is another Euclidean distance calculation between the middle of the center pixel on the back LCD, which is the origin of the coordinate system, to $\vec{\boldsymbol{m}}$. $\vec{\boldsymbol{m}}$ is the result of $\vec{\boldsymbol{a}}_L = \vec{\boldsymbol{a}}_R$ for a given view. $\vec{\boldsymbol{a}}_L$ and $\vec{\boldsymbol{a}}_R$ are formed from sets of intersection points on the front and back mask so if we define those points in our relative coordinate system then we can directly calculate the resulting IPDs and minimum viewing distances (MVD). Let the two intersection points contributing to $\vec{\boldsymbol{a}}_L$ be $\vec{\boldsymbol{p}}_{LB}$ and $\vec{\boldsymbol{p}}_{LF}$ and the points forming $\vec{\boldsymbol{a}}_R$ be $\vec{\boldsymbol{p}}_{RB}$ and $\vec{\boldsymbol{p}}_{RF}$. Now the ray intersections and Euclidean distance calculations give us the following formulas. Equations for determining the pixels on each plane that will contribute to $\vec{\boldsymbol{a}}_L$ and $\vec{\boldsymbol{a}}_R$.

$$\vec{\boldsymbol{p}}_{LB}(n) = \begin{cases} P_F(-\frac{N_p}{2} - n) & n < 0 \\ -\frac{N_p}{2} P_F & n \geq 0 \end{cases},$$

$$\vec{\boldsymbol{p}}_{RB}(n) = \begin{cases} \frac{N_p}{2} P_F) & n < 0 \\ P_F(\frac{N_p}{2} - n) & n \geq 0 \end{cases},$$

$$\vec{\boldsymbol{p}}_{LF}(n) = \begin{cases} (-\frac{N_p}{2} - 1)P_F & n < 0 \\ P_F(-\frac{N_p}{2} - 1 - n) & n \geq 0 \end{cases},$$

$$\vec{\boldsymbol{p}}_{RF}(n) = \begin{cases} P_F(\frac{N_p}{2} - 1 - n) & n < 0 \\ (\frac{N_p}{2} - 1)P_F & n \geq 0 \end{cases},$$

Calculate $\vec{\boldsymbol{a}}_L$ and $\vec{\boldsymbol{a}}_R$

$$\begin{aligned} \vec{\boldsymbol{a}}_L(n) &= [\vec{\boldsymbol{a}}_L(\vec{\boldsymbol{o}}, n), \vec{\boldsymbol{a}}_L(x, n)] \\ &= [\vec{\boldsymbol{p}}_{LF}(n) - \vec{\boldsymbol{p}}_{LB}(n), \vec{\boldsymbol{p}}_{LB}(n)], \end{aligned}$$

$$\begin{aligned}\vec{\boldsymbol{a}}_R(n) &= [\vec{\boldsymbol{a}}_R(\vec{\boldsymbol{o}}, n), \vec{\boldsymbol{a}}_R(x, n)] \\ &= [\vec{\boldsymbol{p}}_{RF}(n) - \vec{\boldsymbol{p}}_{RB}(n), \vec{\boldsymbol{p}}_{RB}(n)],\end{aligned}$$

Calculate parameter value, t, for $\vec{\boldsymbol{a}}_L$ and $\vec{\boldsymbol{a}}_R$ intersection used in determining $\vec{\boldsymbol{m}}$,

$$t(n) = \frac{\vec{\boldsymbol{a}}_R(x, n) - \vec{\boldsymbol{a}}_L(x, n)}{\vec{\boldsymbol{a}}_L(\vec{\boldsymbol{o}}, n) - \vec{\boldsymbol{a}}_R(\vec{\boldsymbol{o}}, n)},$$

Calculate interpupillary distance

$$IPD(\vec{\boldsymbol{m}}(n), \vec{\boldsymbol{m}}(n+1)) = \sqrt{(\vec{\boldsymbol{m}}(n_x) - \vec{\boldsymbol{m}}(n+1_x))^2 + (\vec{\boldsymbol{m}}(n_z) - \vec{\boldsymbol{m}}(n+1_z))^2},$$

Calculate minimum viewing distance

$$MVD_n = \sqrt{\vec{\boldsymbol{m}} \cdot \vec{\boldsymbol{m}}}.$$

Calculate minimum viewing distance position for $V(n)$

$$\begin{aligned}\vec{\boldsymbol{m}} &= [\vec{\boldsymbol{m}}(x), \vec{\boldsymbol{m}}(z)] \\ &= [\vec{\boldsymbol{a}}_L(\vec{\boldsymbol{o}}, n)t(n) + \vec{\boldsymbol{a}}_L(x, n), d_p t(n)].\end{aligned}$$

Each $\vec{\boldsymbol{m}}$ describes a point at the minimum distance where a particular view can be perceived, which means a greater distance in the same zone will result in a higher quality viewing experience. The consequence of analysing the display design in this biased manner is that, if the results show promise, the design criteria can be altered to produce a more favourable result. Figure 3.21 depicts the viewing zones produced from the set of $\vec{\boldsymbol{m}}$ values shown in 3.20.

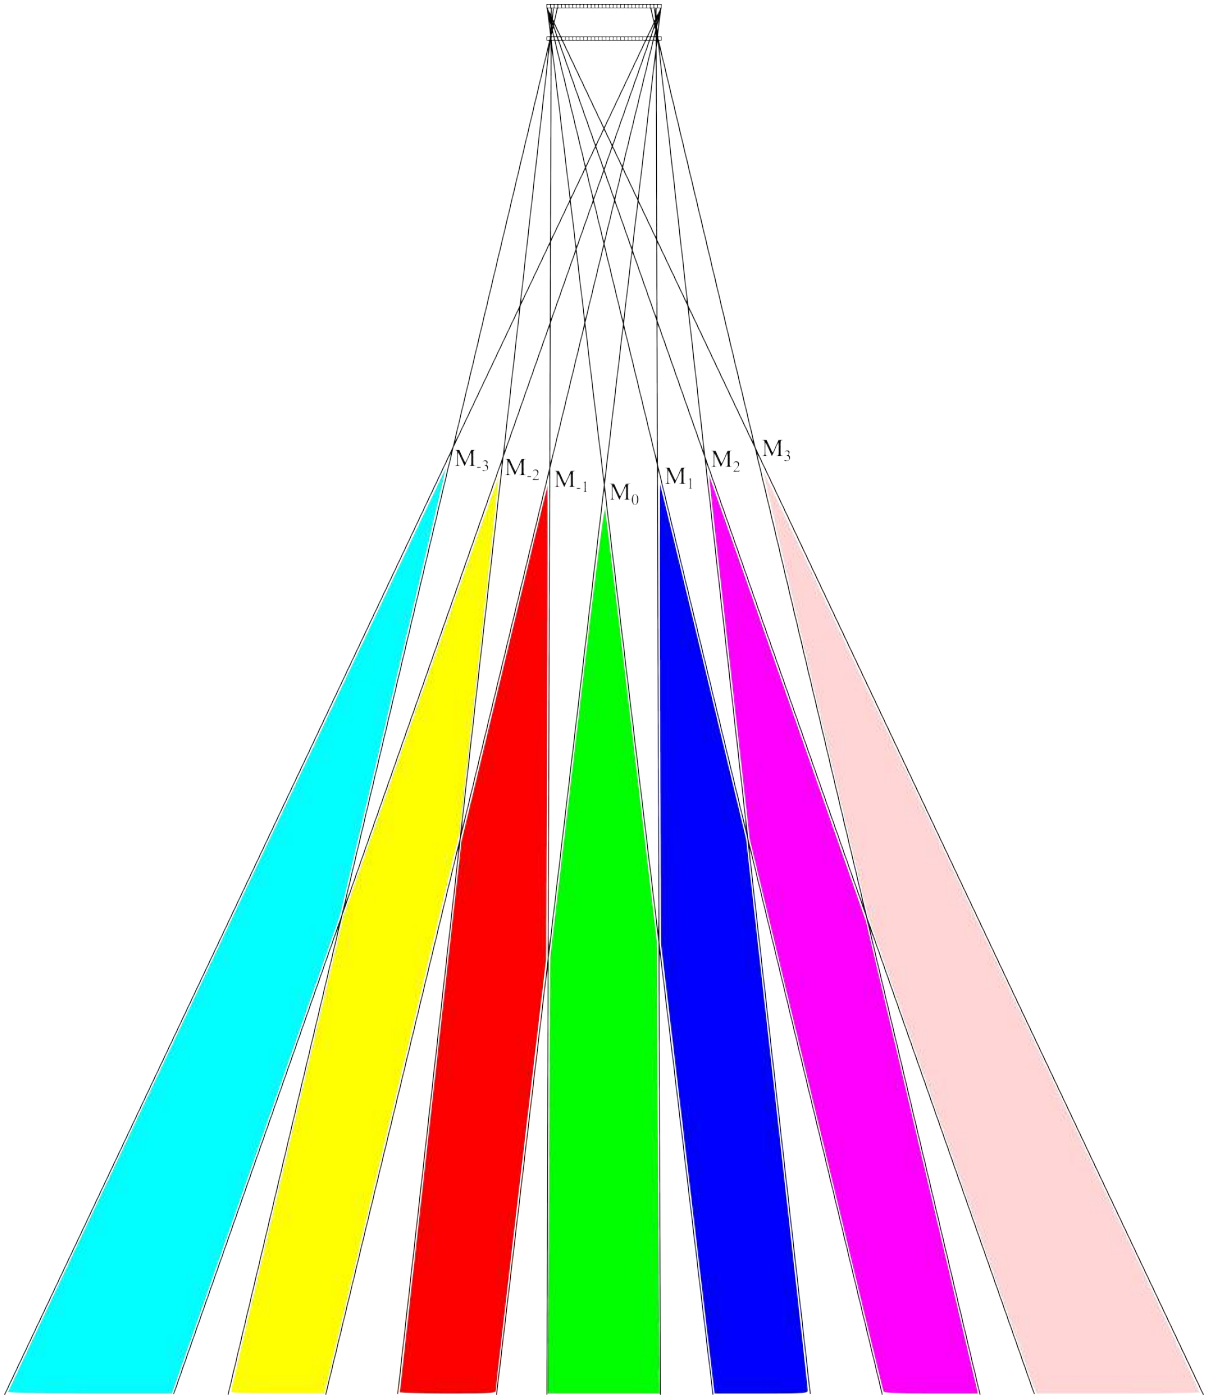

Figure 3.21: the viewing zones produced from the set of $\vec{\boldsymbol{m}}$ values shown in 3.20. Coloured zones are proper viewing areas and white zones are areas that will have crosstalk between vewing zones. Viewing from a higher distance reduces crosstalk and other artifacts further, but a significant increase in distance will result in a perceived lower resolution image simply because pixels will be too small to perceive all of them.

Chapter 4

Implementation

The code is written in C++, which uses the Boost library to manage arguments and directory operations [10]. The program takes a light field as input, forms a matrix, decomposes the matrix to produce mask pairs and simulates a multi-layer LCD setup. Any light field can be used as input, but test cases are generated using POV-Ray [56]. Input light fields are represented using an array of images. Reading and writing images is done using the FreeImage library [25]. After loading, the light field is rearranged into a matrix using the sampling method describes in Section 3.3. The light field matrix form is used as the input for one of three matrix approximation techniques which are a weighted version of Lee and Seung's approach from [47], FNMA_e shown in [42] and a weighted version of FNMA_e shown in Section 3.4.1. Matrix operations are performed using the Eigen linear algebra library [23]. The matrix factorization results are reorganized into mask pairs, which can be displayed in the simulator. The simulator uses GLFW to manage windows, OpenGL to display graphics and GLM to help with graphical math operations [26, 27, 57].

4.1 Sparse Matrix

The sampling procedure used in content adaptive parallax barriers produces a sparse matrix. A sparse matrix data structure allows for significantly reduced memory usage and number of operations performed. The sparsity alone carries benefits, but it is also the reason that weighted NMF is used. Weighted NMF combined with the sparse input matrix allows for zero-entries to be ignored entirely. Zero entries are ignored in a sparse matrix data structure, but NMF would still attempt to approximate those zeros. Weighted NMF allows for avoiding any calculation that will directly involve a zero entry in **L**. This is advantageous because it means that all calculations between **F** and **G** that produce an entry of **FG**, which relates to a zero entry of **L**, are never performed. The benefits are that higher resolution images can be used and the compute time is significantly reduced.

4.2 Simulator

A multi-layer LCD setup is simulated using OpenGL by exploiting the two-plane parameterization of a light field. Given a mask pair, each mask is rendered seperately, as a texture on a quad. The quads are rendered with one positoned in front of the other, representing the front and back LCD panels in a dual-layer display. Each quad is rendered seperately and saved as a texture. To produce the image that will actually be perceived by the display, corresponding pixel values of each of the rendered images are multiplied. The multiplication is performed in parallel using a GLSL shader and the result is rendered to screen. This entire process is repeated for each mask pair and the result is displayed as it is completed. The calculation occurs relatively fast as on a 120Hz monitor a frame rate of between 119 and 120 frames per second is achieved. Ray tracing is an alternate method for simulation.

The simulation represents the result a single eye would be viewing on the multi-layer

LCD setup. For this reason the simulator does not produce stereo results, but this type of modification is straight forward in the event that a display device can be used for stereo viewing, such as the oculus rift. On the multi-layer LCD device stereo results are unecessary as the device is autostereoscopic, if designed correctly.

The simulator allows for dynamically changing a number of parameters such as switching between an orthographic and perspective projection, change the distance between quads and altering the camera position and direction.

Chapter 5

Results

The three main sets of results are the numerical results obtained during matrix factorization, the visual results obtained while simulating the display and the values determined

5.1 Numerical Results

As seen in Figure 5.1, the unweighted form of FNMA_e produces the highest PSNR consistently. This is interesting as the method proposed for content adaptive parallax barrier is a weighted form of the Lee and Seung multiplicative update rules from [47], but neither of the weighted versions appear to perform as well. This result is misleading as discussed in Section 5.2. An interesting note is that at much higher ranks the WFNMA_e approach has a declining PSNR.

We compare the execution times of each NMF approach in Figure 5.2, where the power of sparse matrices in revealed by observing the growth rate of FNMA_e as opposed to the other two techniques. Higher ranks were timed, but the amount of time that FNMA_e took was far too long for a useful visualization.

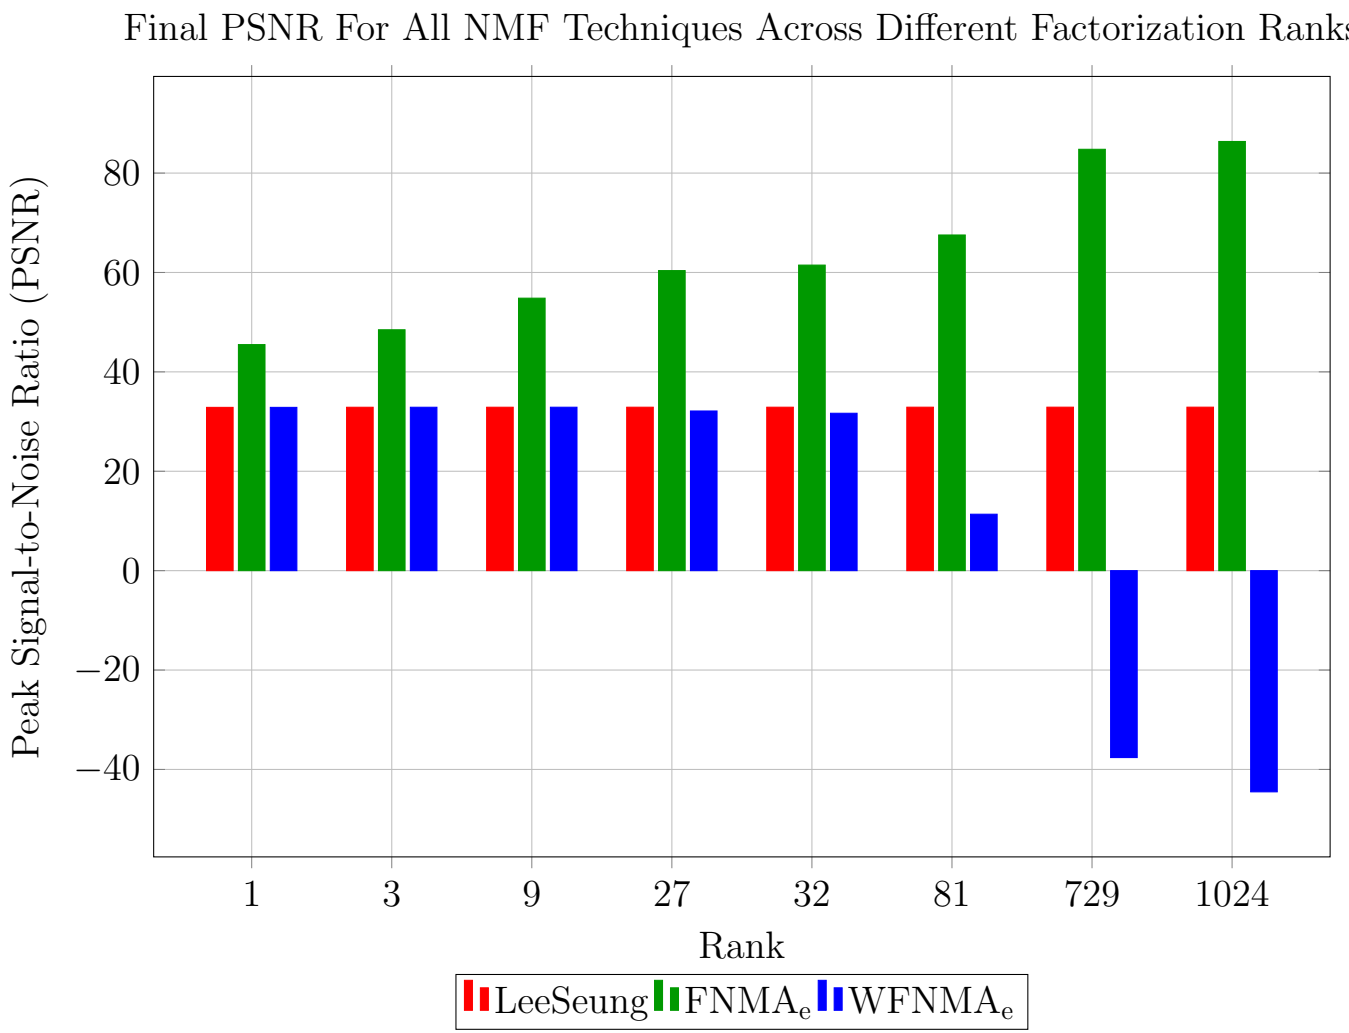

Figure 5.1: A listing of the PSNR achieved after the final iteration for each NMF technique across a number of different ranks. This is a low resolution dataset of 32x32 images each with 32x32 pixels.

Figure 5.2: The processing time for each matrix factorization technique across fifty iterations. The difference in computation time so great that a logarithmic scale must be used to properly note the relative calculation times. This is a low resolution dataset of 32x32 images each with 32x32 pixels.

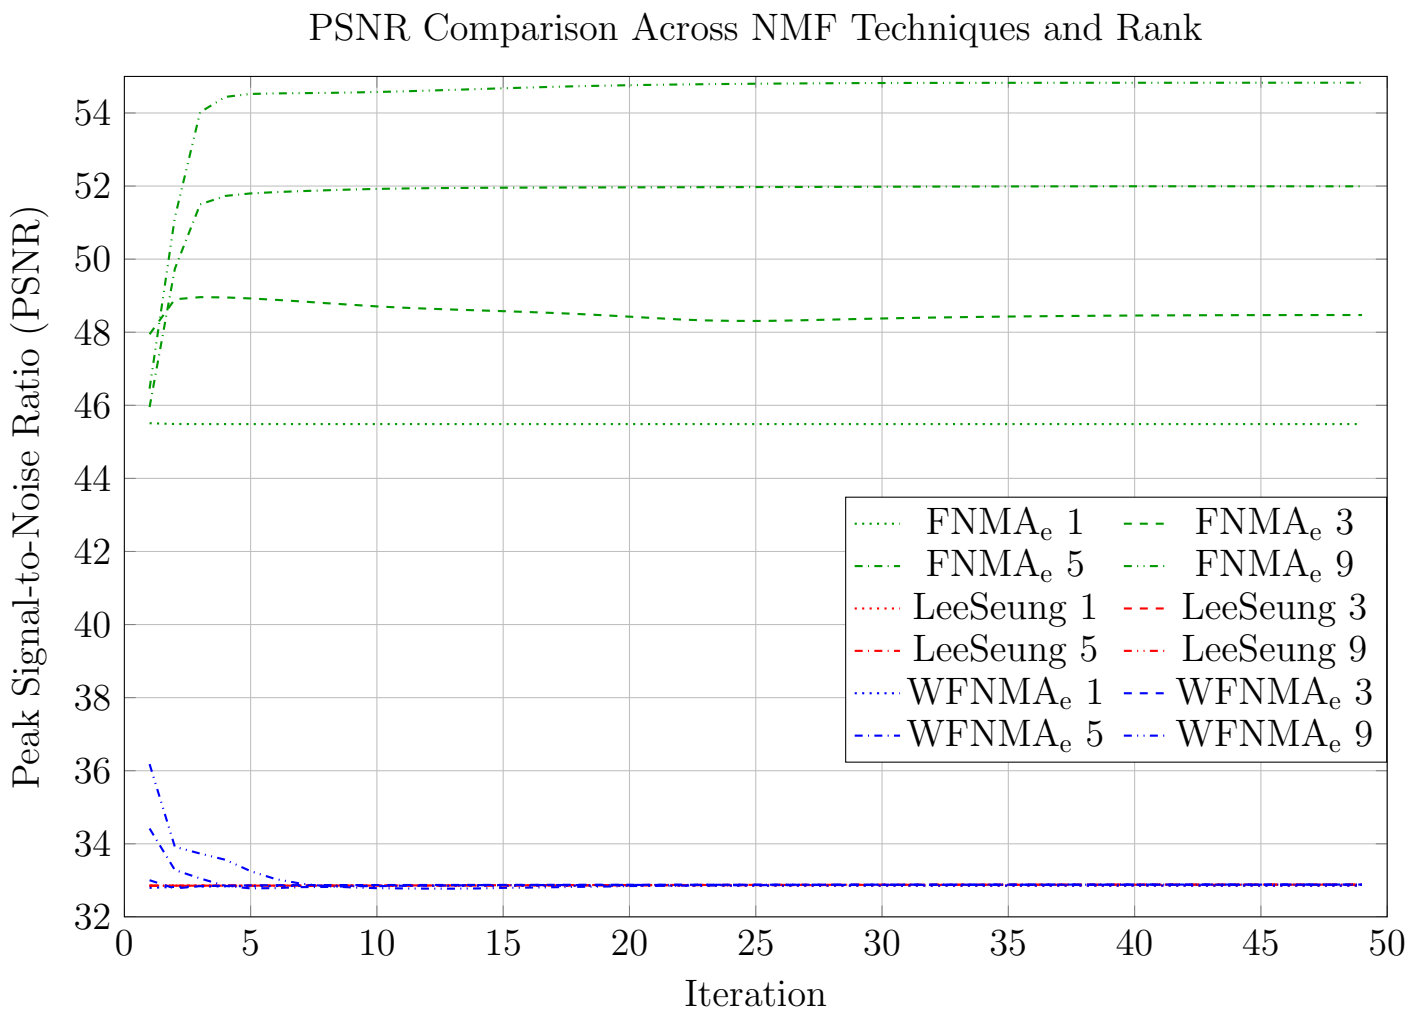

Figure 5.3: The PSNR obtained across each iteration for all factorization techniques across a number of different ranks. The legend lists the technique and the rank associated with that data series. This is a low resolution dataset of 32x32 images each with 32x32 pixels.

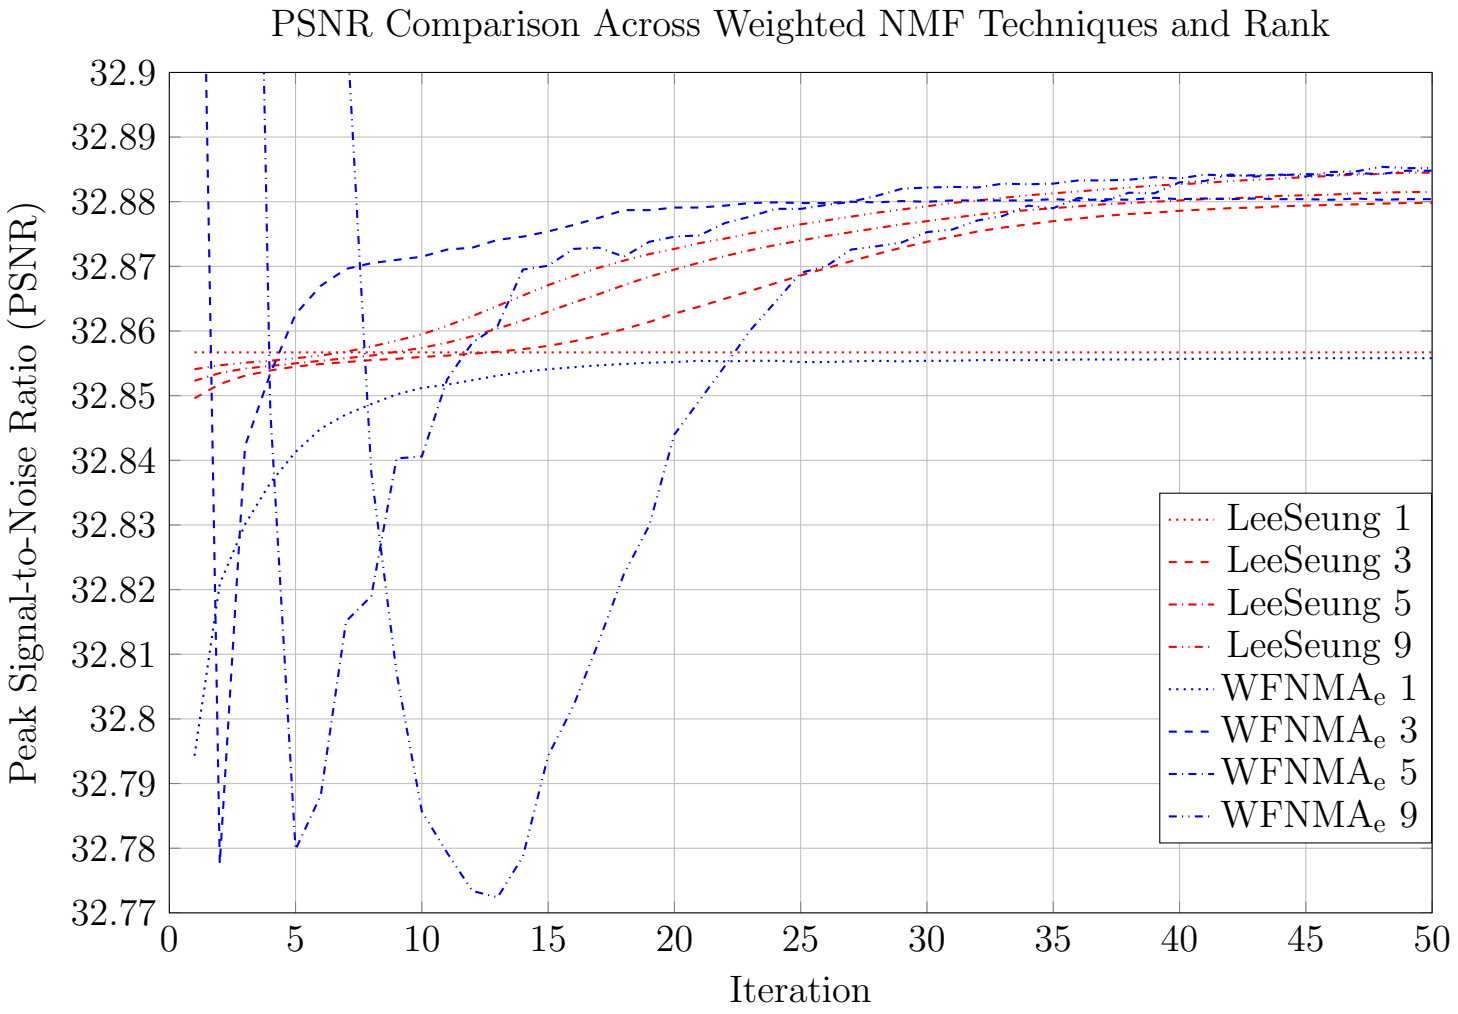

Figure 5.4: The PSNR obtained across each iteration for all weighted factorization techniques across a number of different ranks. The legend lists the technique and the rank associated with that data series. This is a close up of some of the data present in Figure 5.3. This is a low resolution dataset of 32x32 images each with 32x32 pixels.

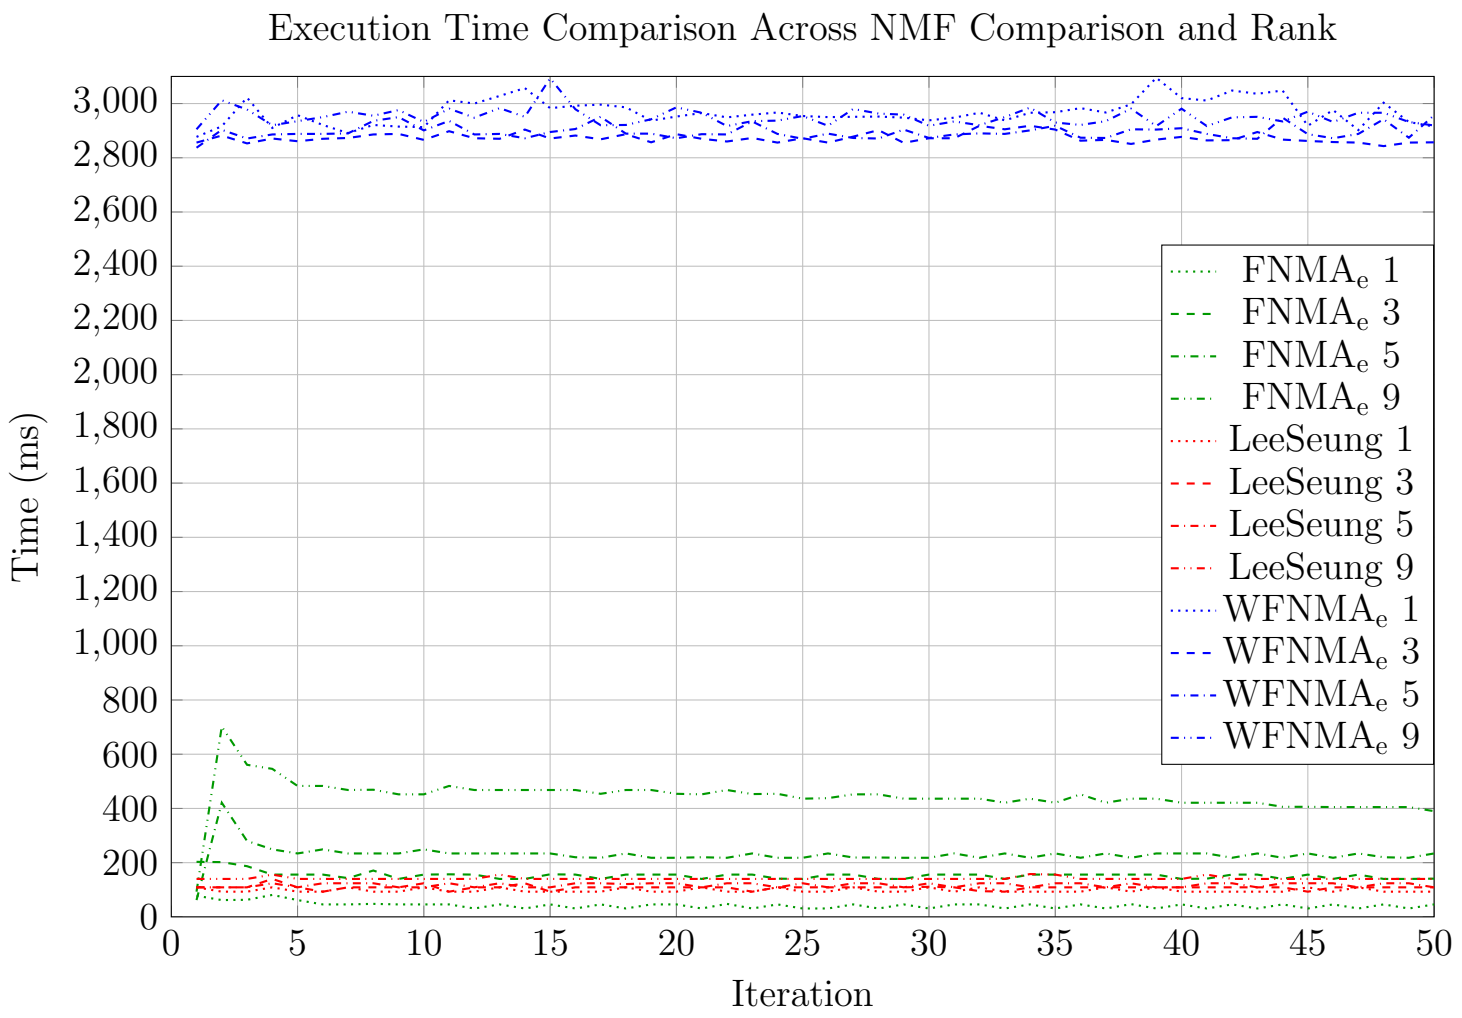

Figure 5.5: The processing time for each matrix factorization technique across fifty iterations. This is a low resolution dataset of 32x32 images each with 32x32 pixels.

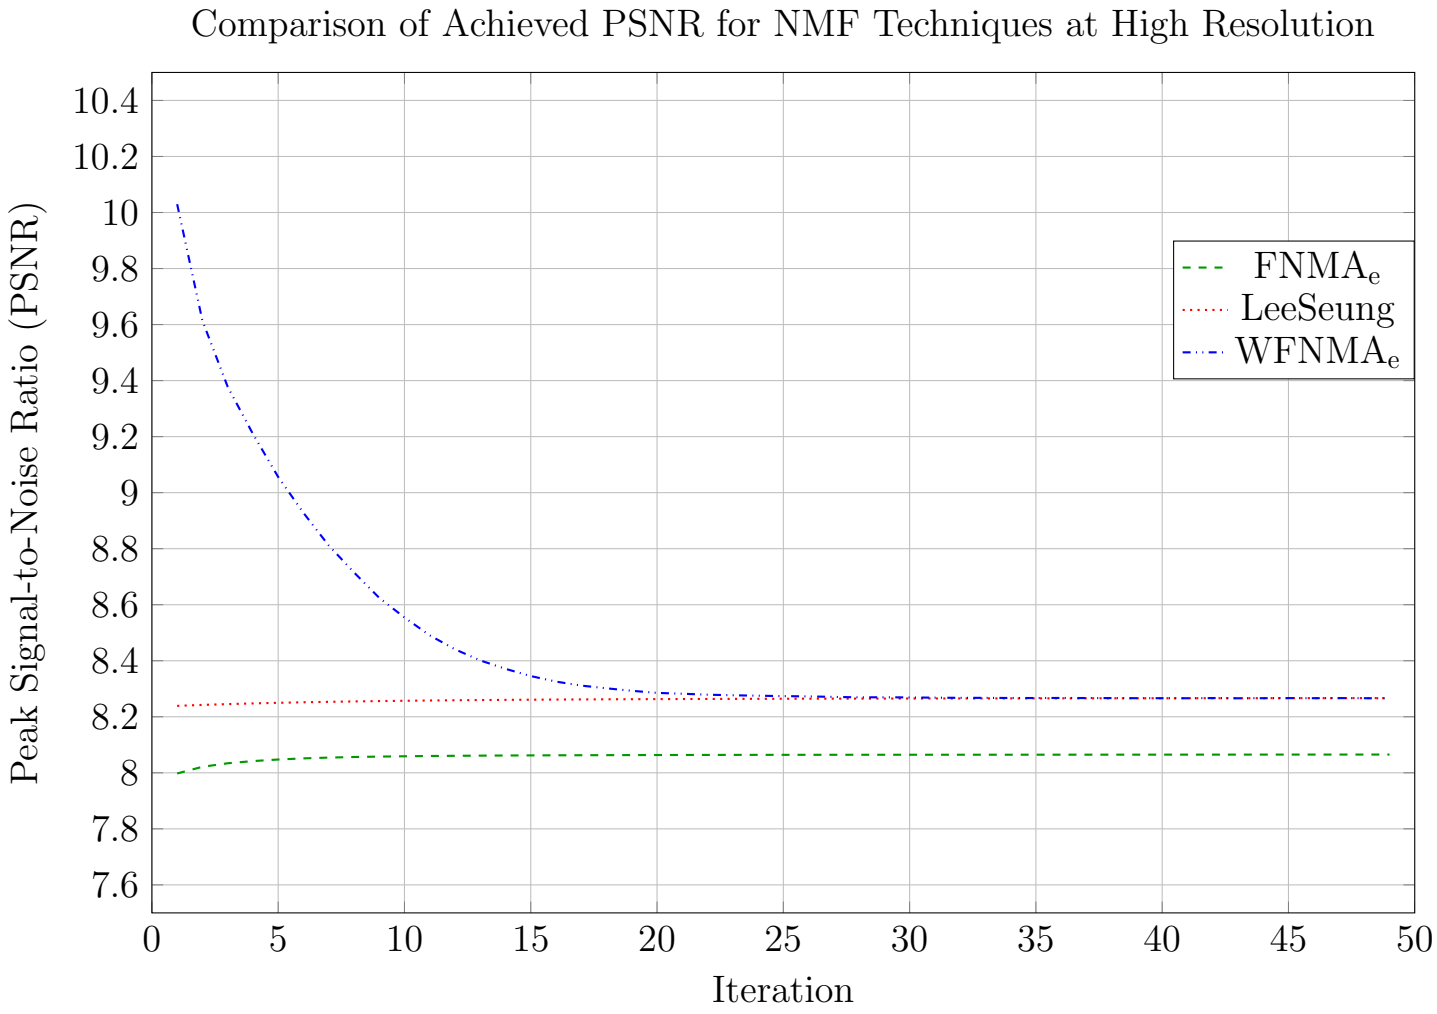

Figure 5.6: The PSNR for each matrix factorization technique across fifty iterations. This is a higher resolution dataset of 9x9 images each with 266x150 pixels.

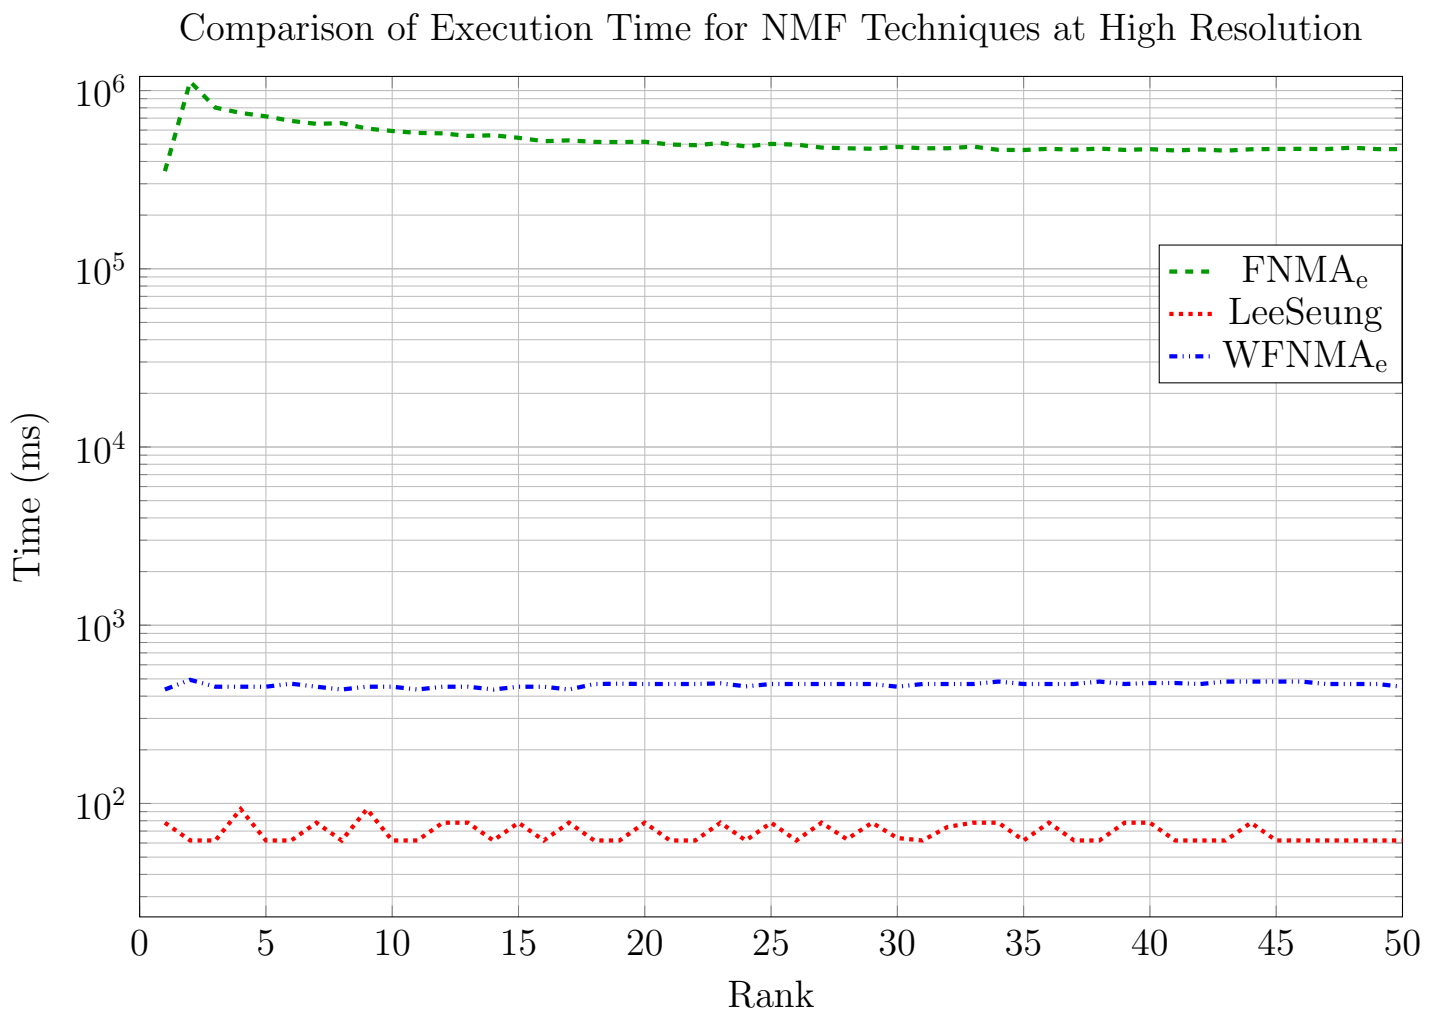

Figure 5.7: The processing time for each matrix factorization technique across fifty iterations. The difference in computation time is so great that a logarithmic scale must be used to properly note the relative calculation times. This is a higher resolution dataset of 9x9 images each with 266x150 pixels.

Figure 5.3, shows the unweighted $FNMA_e$ climb quickly and then hits a plateau where the PSNR is still rising, but much more slowly. The weighted $FNMA_e$ actually declines swiftly at first until it hits a similar PSNR of the weighted Lee and Seung method [8,17]. Figure 5.4 shows the behaviour of both weighted approaches in finer detail, where we can see, after the initial decline, the weighted $FNMA_e$ method will reliably obtain a better PSNR value than the weighted Lee and Seung approach from [8,17]. Figure 5.5 shows the execution time per iteration corresponding to the data displayed in Figures 5.3 and 5.4.

Due to the large amount of resources, compute power and memory, necessary for $FNMA_e$ at higher resolutions most of the graphs represent a low resolution dataset. Figures and display a comparison for the PSNR and execution time, respectively, across each iteration. The resolution of this daaset is 9x9 images at 266x150 pixels each. This was the maximum resolution our computer could support for $FNMA_e$ due to its dense matrix representation.

5.2 Technique Comparison

The visual results are a stark contrast to the results shown in Section 5.1. The Lee and Seung approach from [8,17] appears to produce the cleanest result. During simulation the $WFNMA_e$ is slightly noisier, but gives a relatively good visual result. The unweighted $FNMA_e$ approach simulates an almost black image for all views. In the following sections a comparison of the different NMF techniques are shown across a number of datasets. Due to the dense representation of $FNMA_e$ these images are relatively low resolution as this technique has a much higher memory requirement in comparison to the sparse representation of $WFNMA_e$ and the Lee & Seung method [47].

5.2.1 Teapot

The teapot light field was created using POVRay [56]. The input images and masks have a resolution of 266×150. This was the main dataset we used for the majority of testing. This was done because this light field is a set of multi-view images that does not have a very high amount of parallax or a high frequency. That said there are some elements with a higher frequency, and parallax is present. The variety, with a lack of extremes, allows for easily examining a broad number of factors without running many different test cases or worrying that the applied techniques will only work in this specific circumstance. In addition this set of images was not taken from another source and the teapot is a classically used model in the demonstration of computer graphics related results.

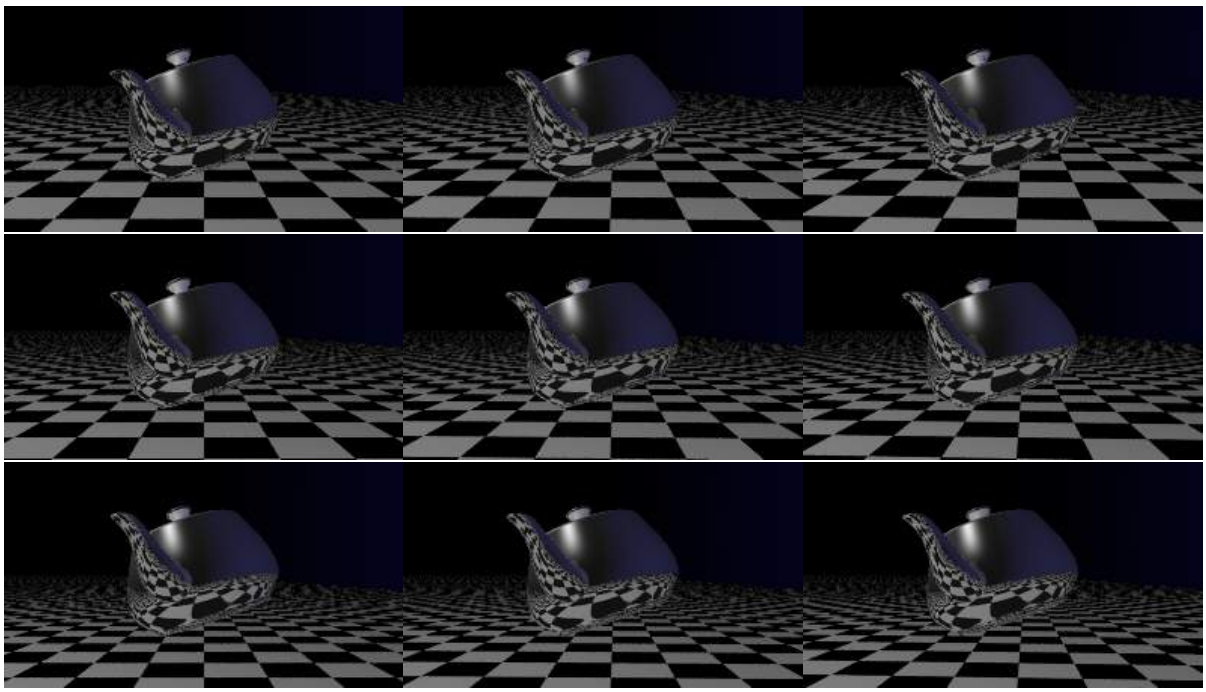

Figure 5.8: Teapot light field. All results in Section 5.2.1 are produced from this light field.

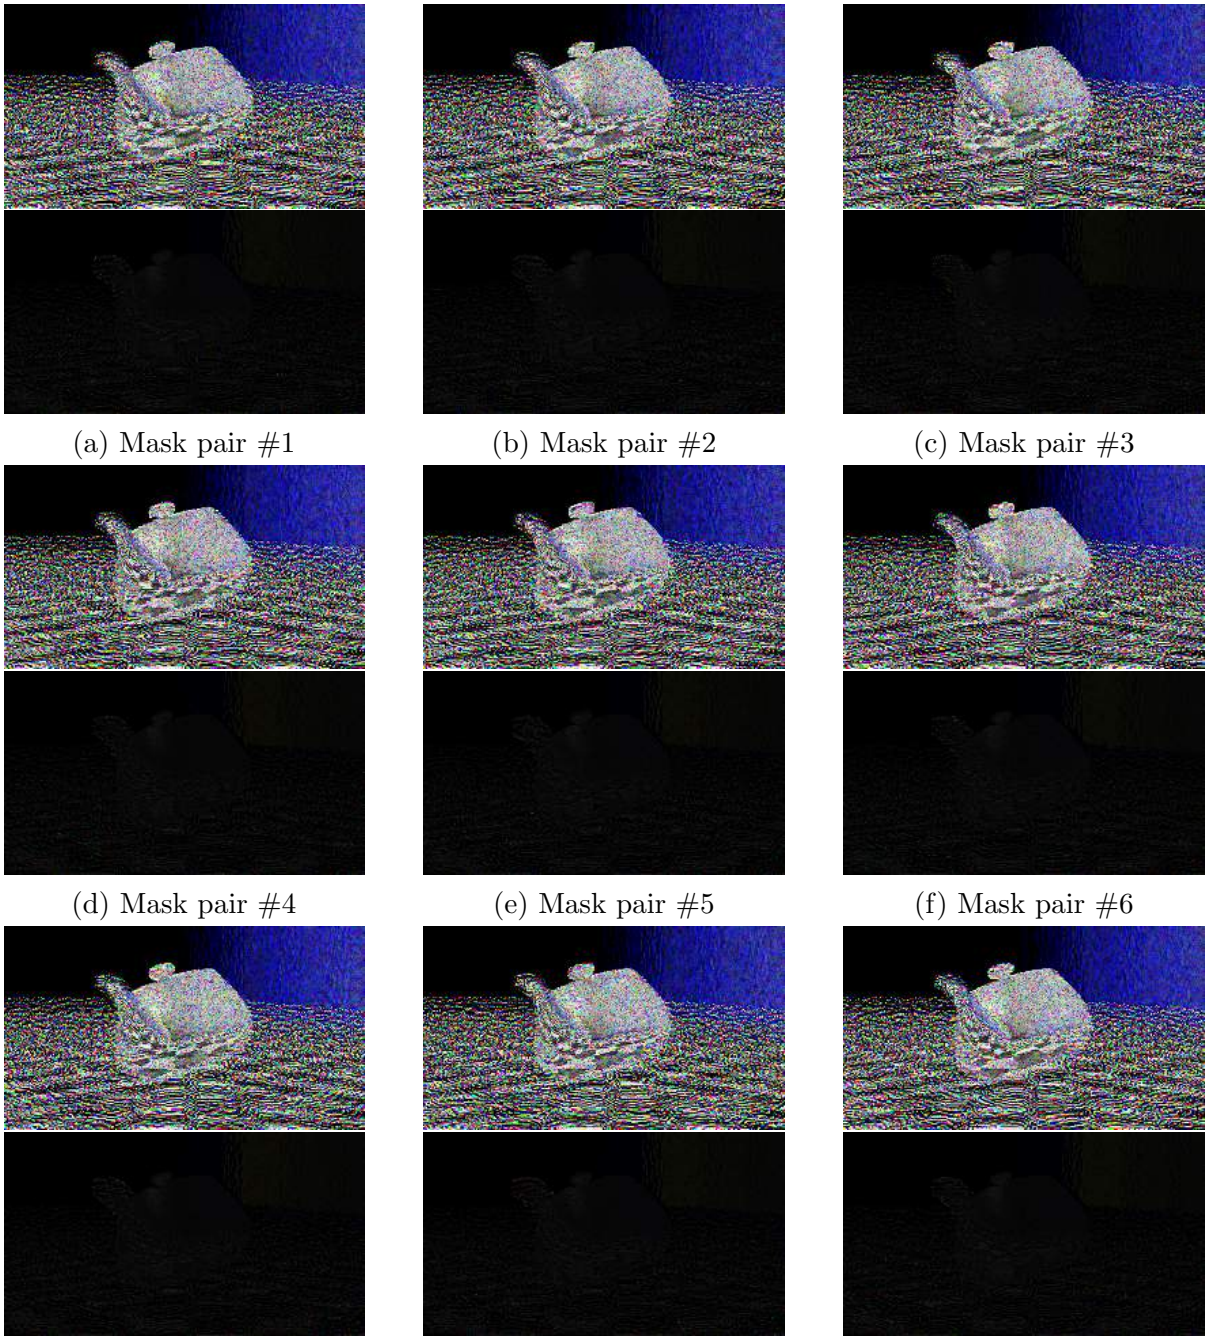

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.9: Masks produced by the Lee & Seung method, from [47], for the teapot light field.

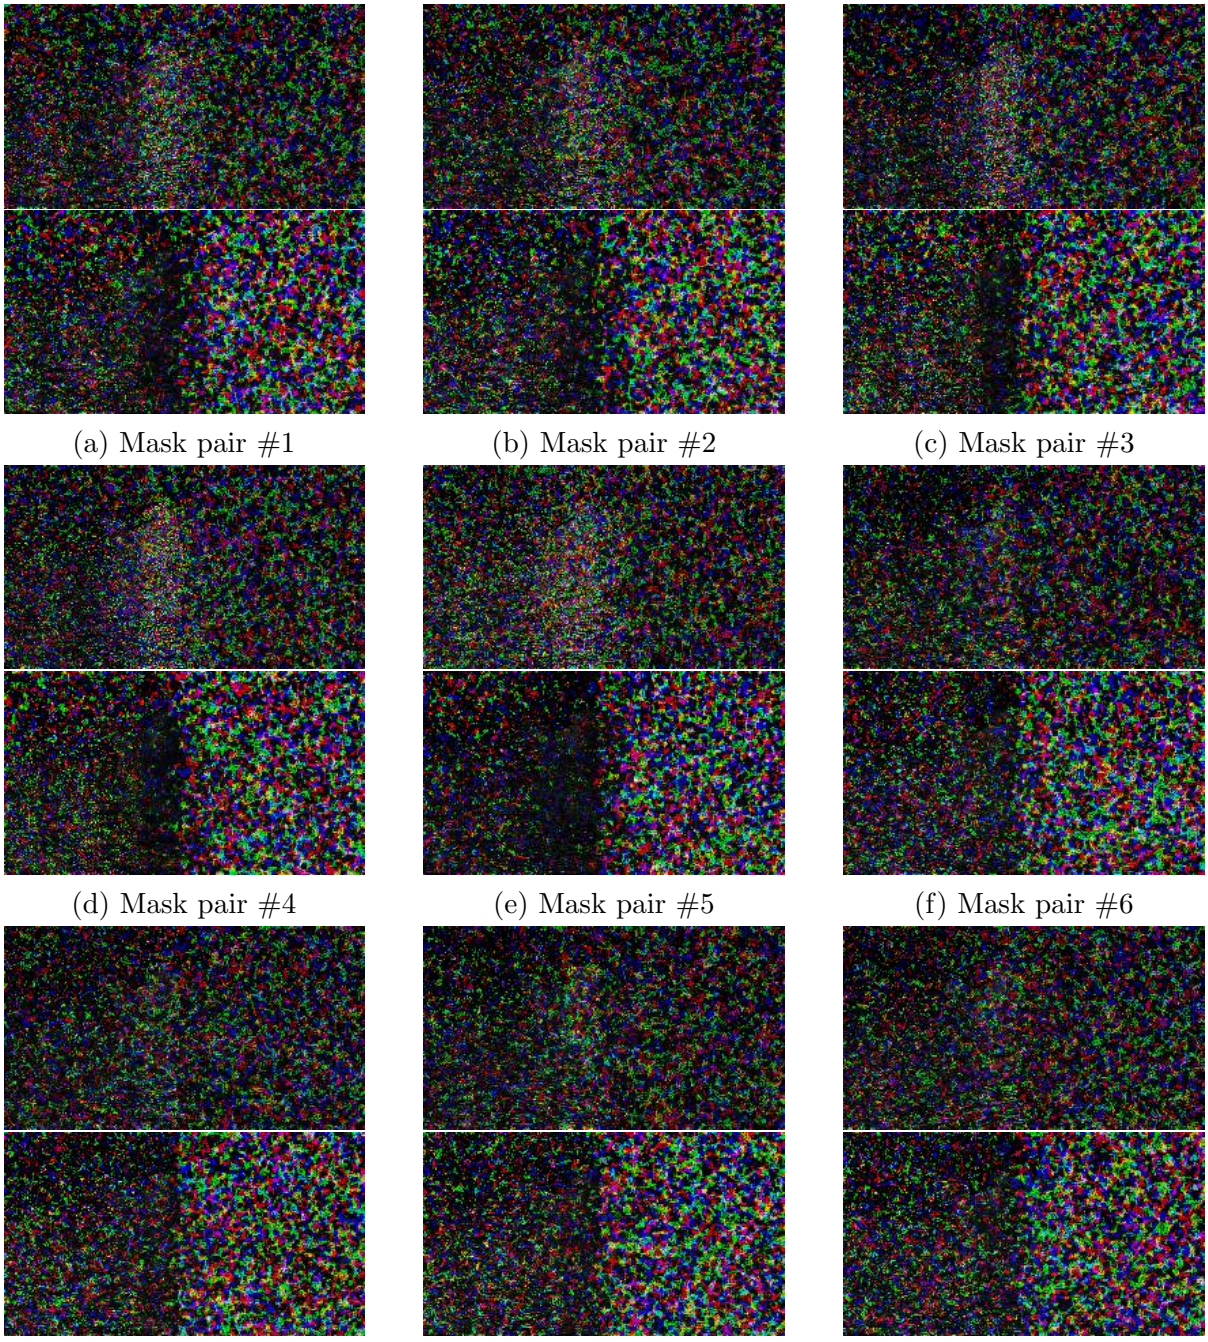

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.10: Masks produced from the WFNMAe method for the teapot light field.

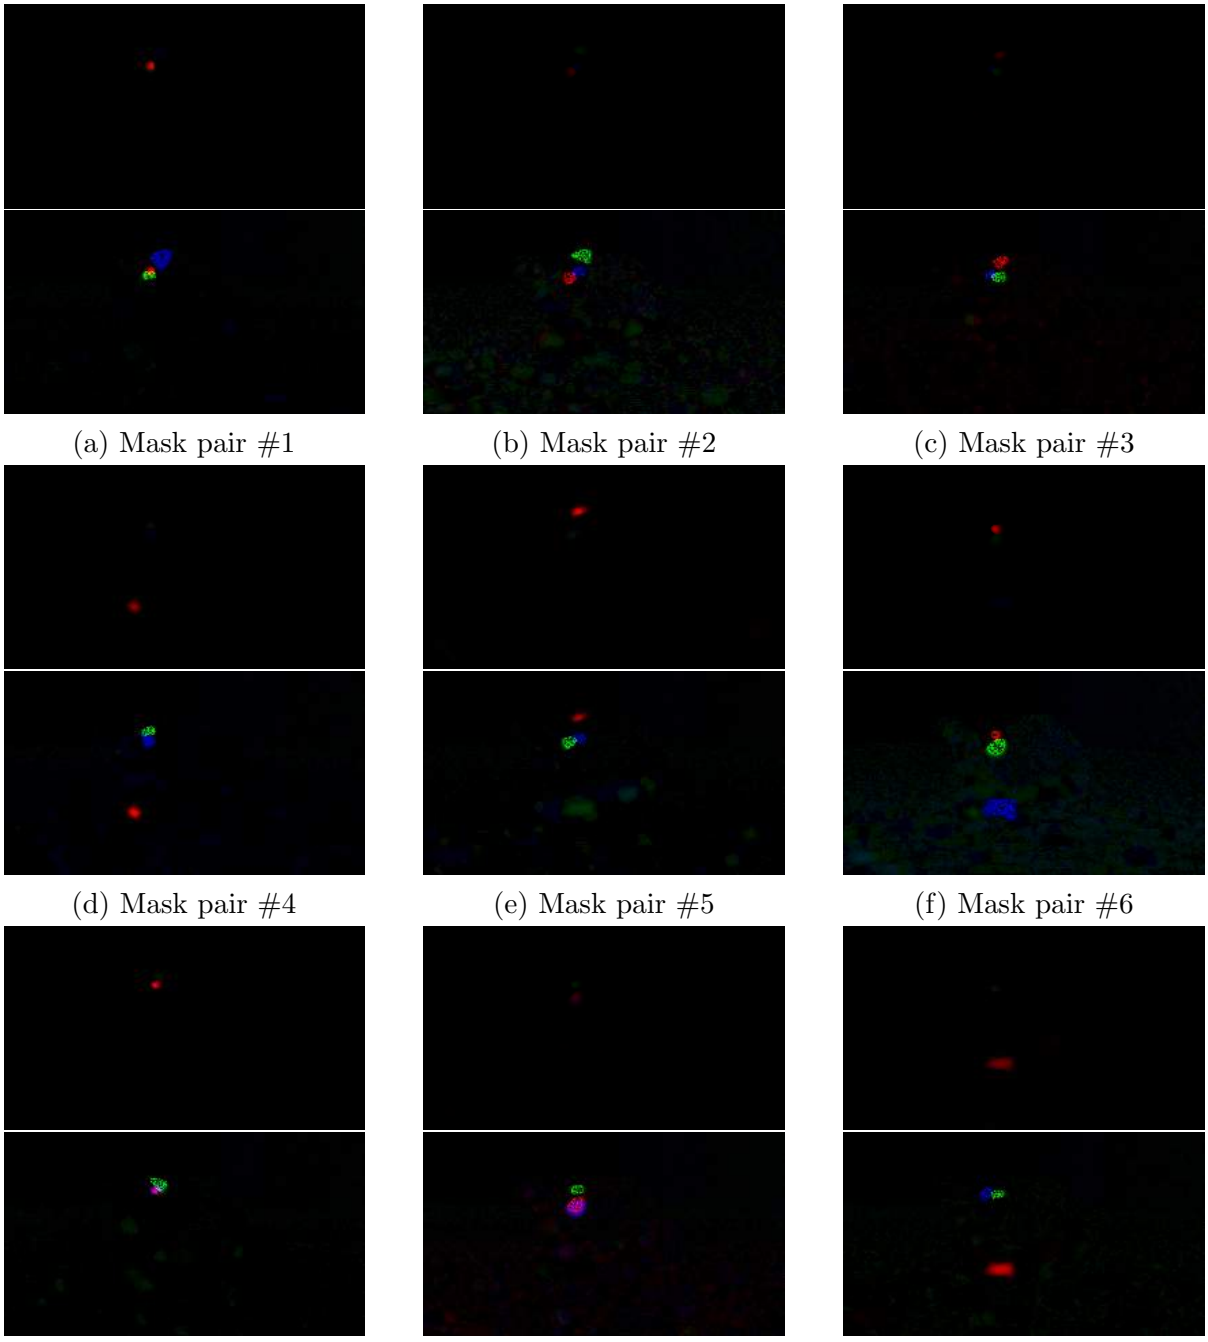

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.11: Masks produced by the FNMAe method, from [42], for the teapot light field.

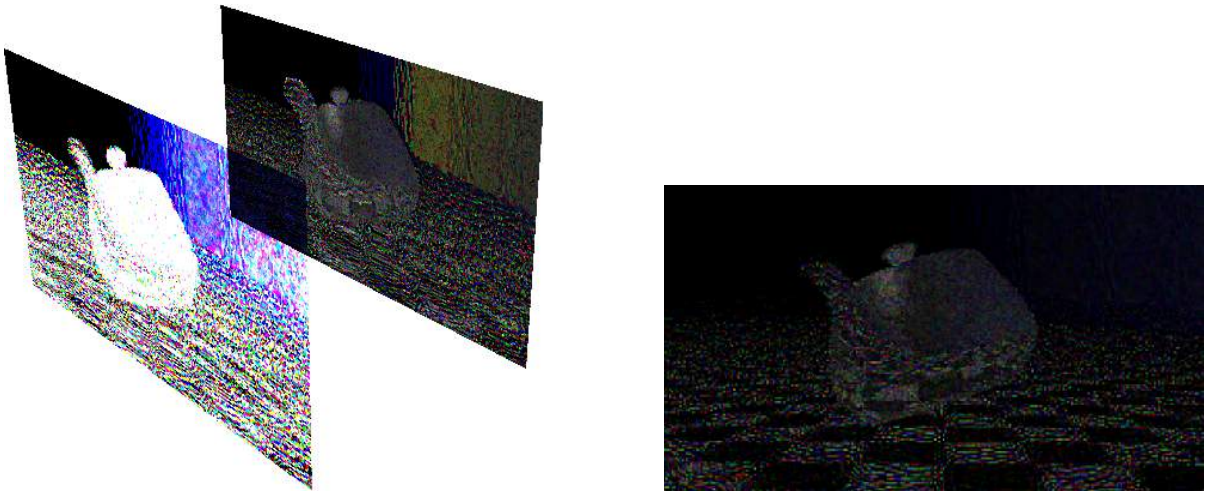

(a) Lee & Seung

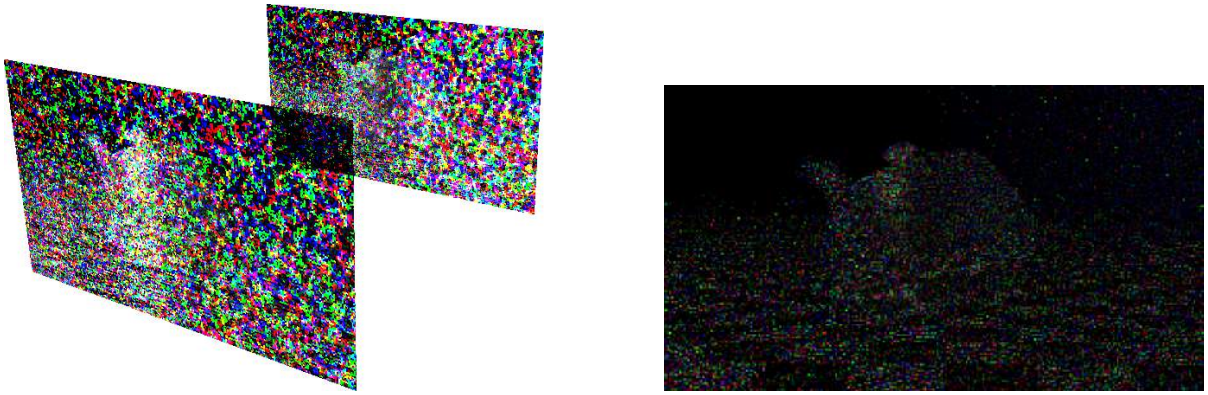

(b) WFNMAe

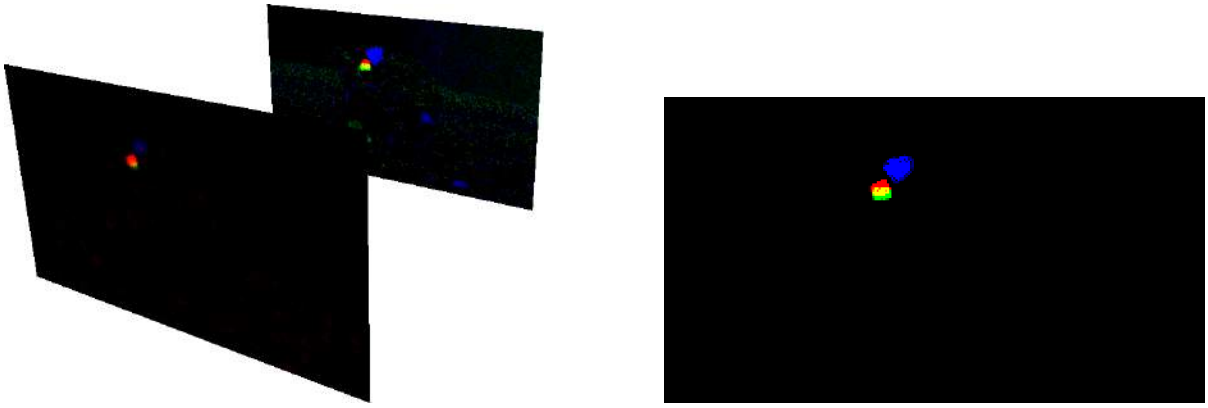

(c) FNMAe

Figure 5.12: Simulated results comparison for the teapot light field. (left) The front and back masks on their respective planes. (right) The simulated result

5.2.2 Blocks

The blocks light field was created by the content adaptive parallax barriers team at MIT [1]. The input images were resized, but the aspect ratio was preserved, to accomodate the requirements of FNMA_e. The images and masks have a resolution of 240×150. This light field was chosen due to it being used by the original authors in [45] as well as it having a particarly strong demonstration of parallax. Using the blocks light field allows for the demonstration that we have accurately reproduced the original results with the original methods as well as giving a comparison of the other NMF methods that we explored. In addition, the strong parallax, which is produced by the multiple translucent objects moving in front of each other, is very useful for testing in the context of autostereoscopic displays.

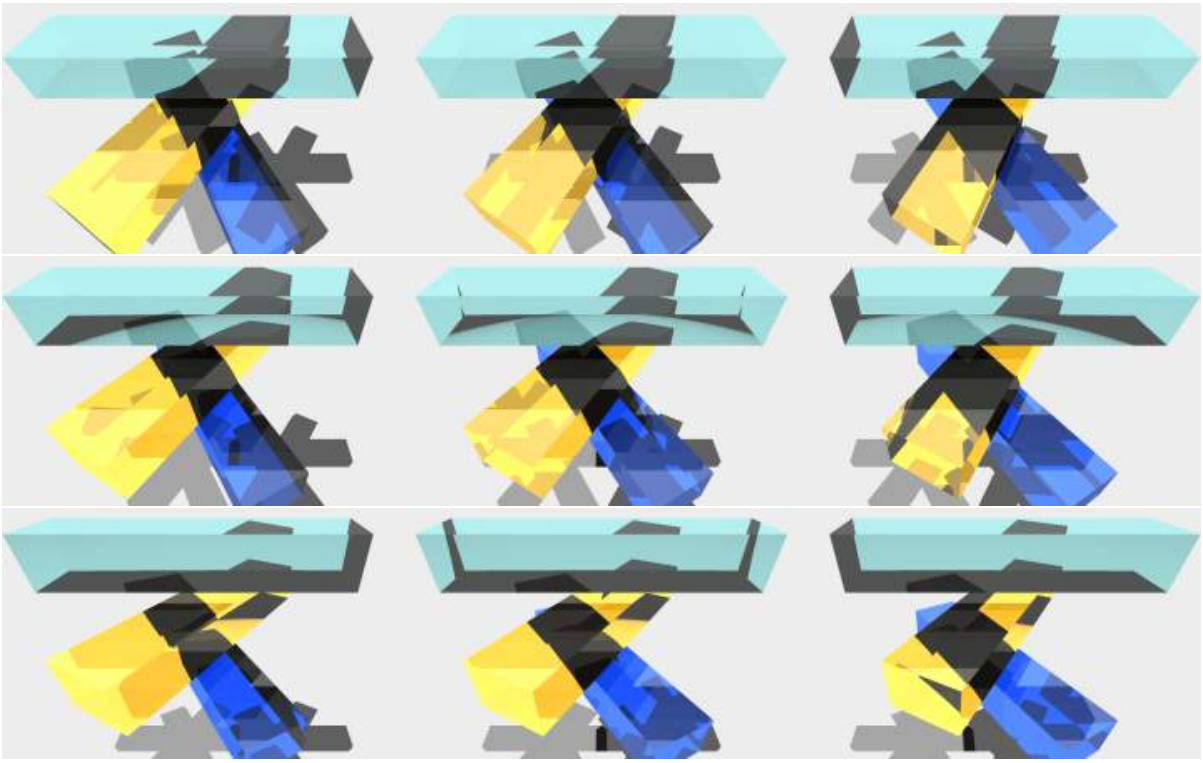

Figure 5.13: Blocks light field. All results in Section 5.2.2 are produced from this light field.

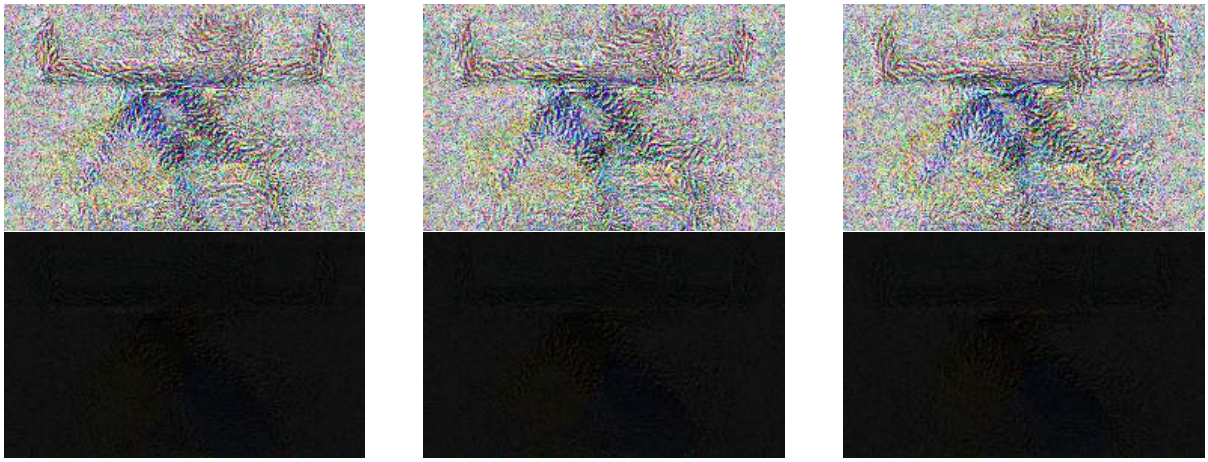

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

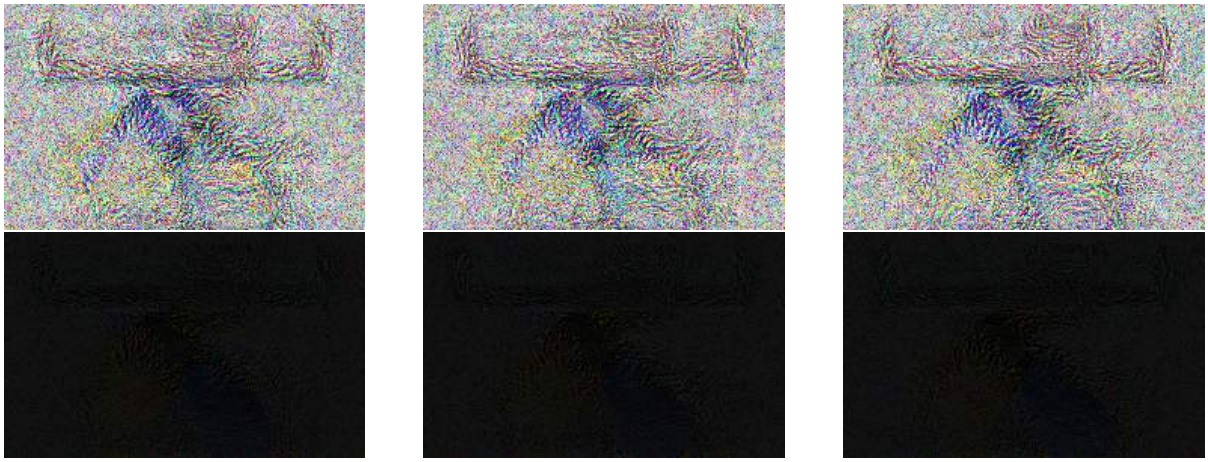

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.14: Masks produced by the Lee & Seung method, from [47], for the blocks light field.

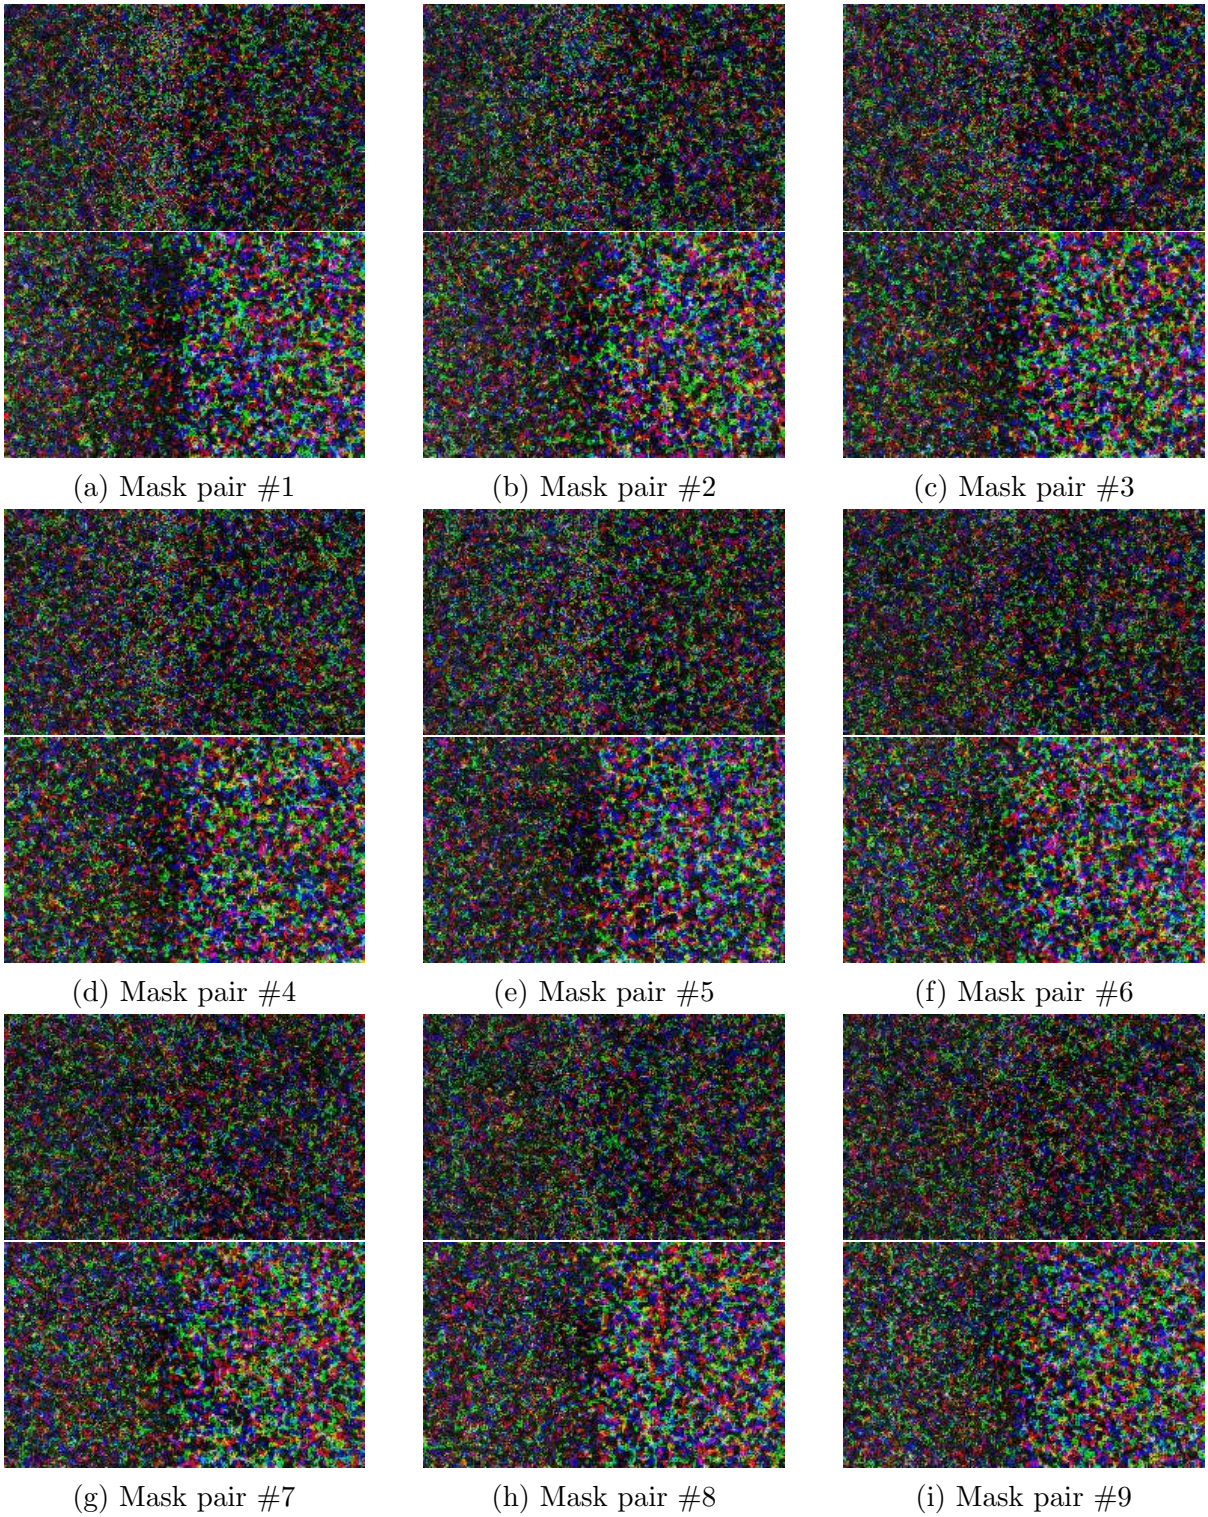

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.15: Masks produced from the WFNMAe method for the blocks light field.

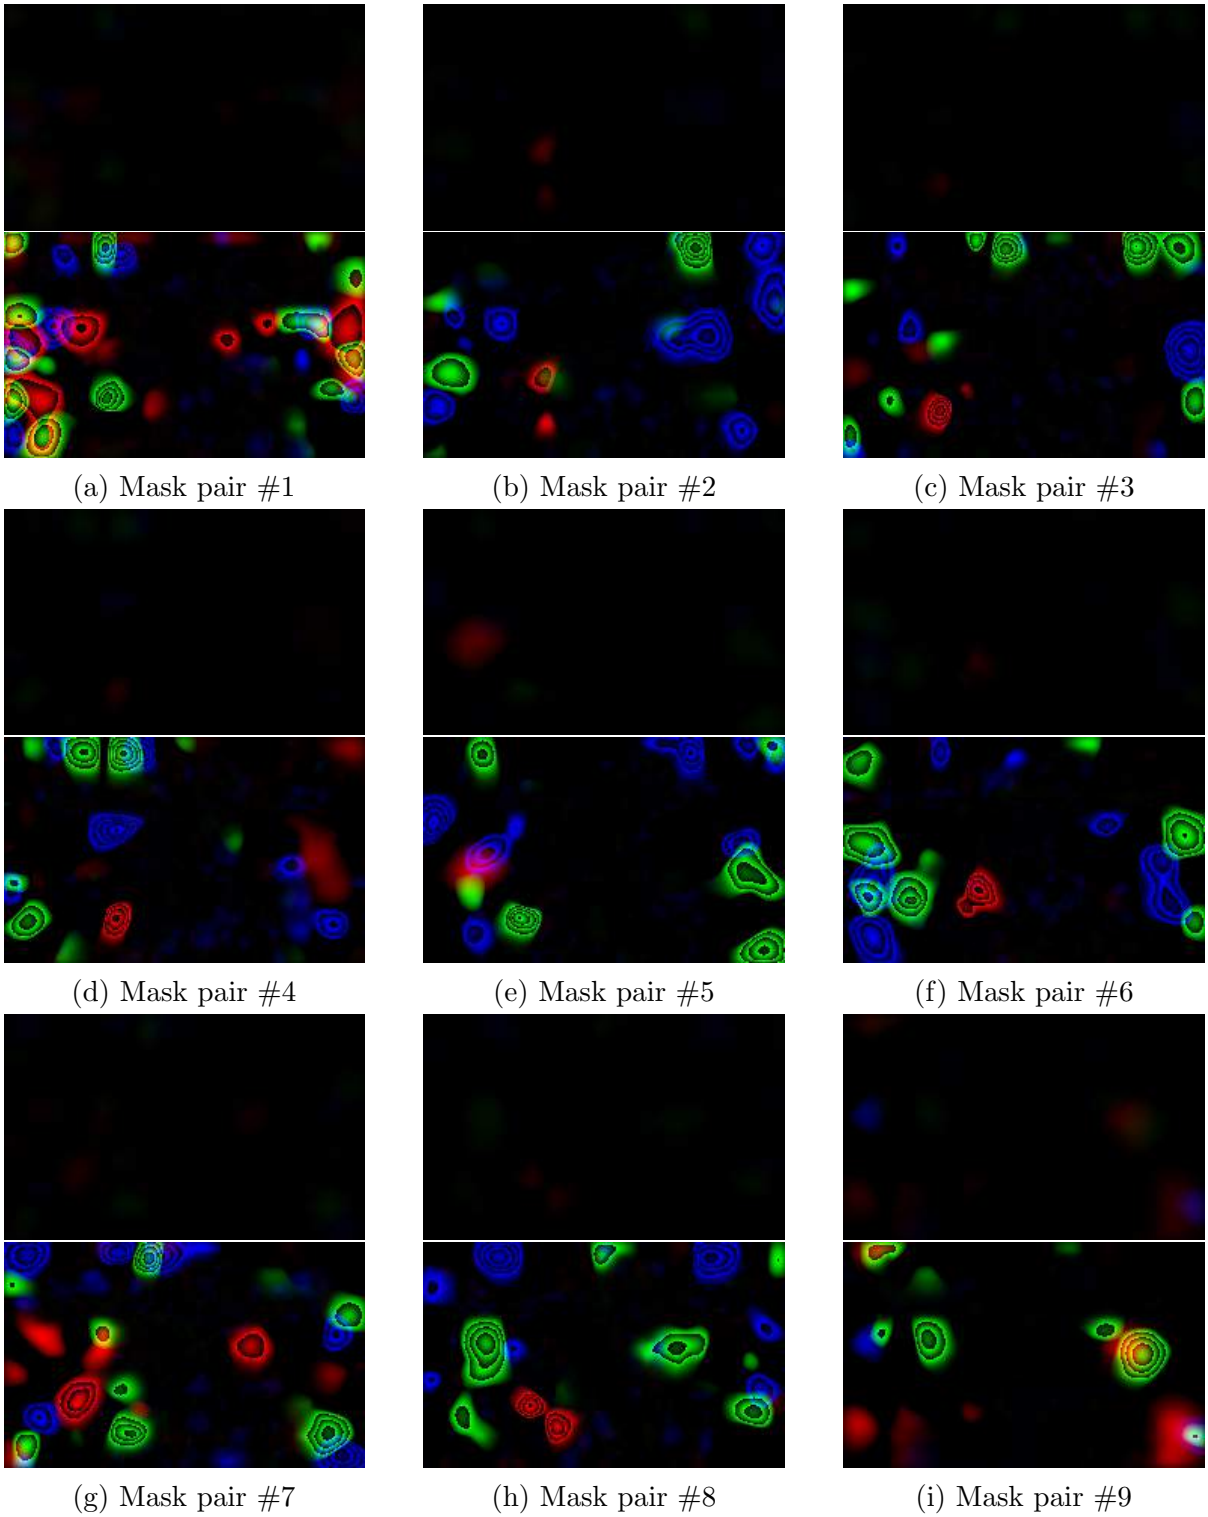

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.16: Masks produced by the FNMAe method, from [42], for the blocks light field.

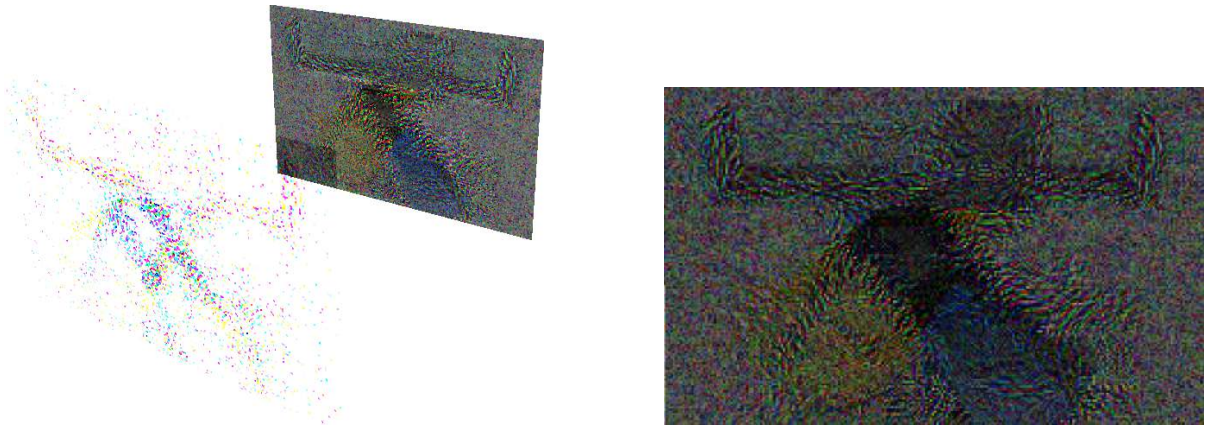

(a) Lee & Seung

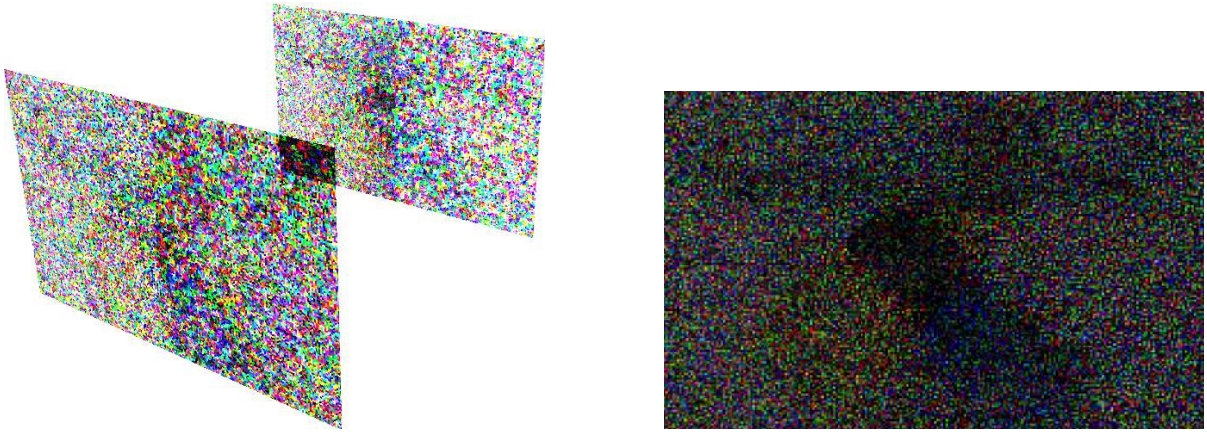

(b) WFNMAe

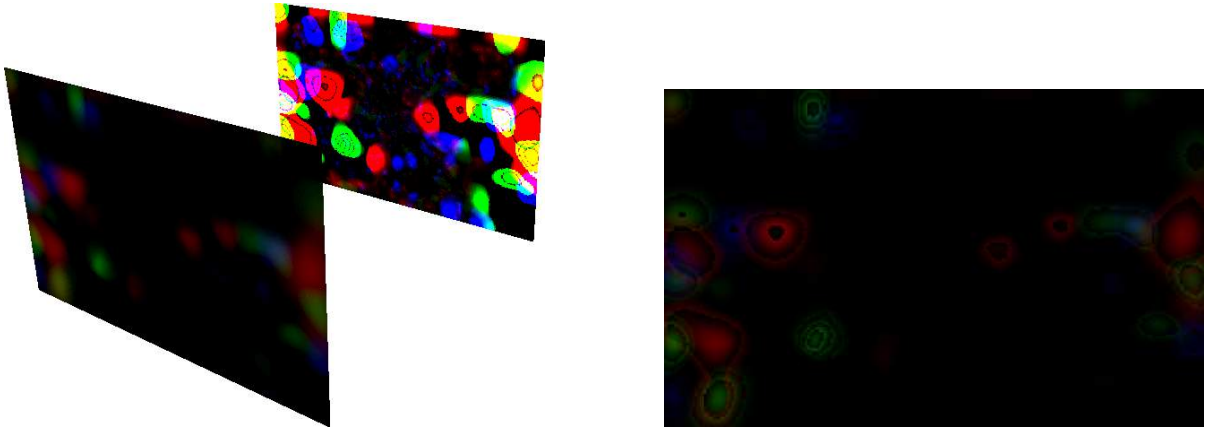

(c) FNMAe

Figure 5.17: Simulated results comparison for the blocks light field. (left) The front and back masks on their respective planes. (right) The simulated result

5.2.3 Numbers

The numbers light field was created by the content adaptive parallax barriers team at MIT [1]. The input images were resized, but the aspect ratio was preserved, to accomodate the requirements of FNMA_e. The images and masks have a resolution of 200×150. A light field where all the images are simply a different number allows for a simple demonstration of viewing different views from different viewing angles. With a proper light field, meaning one that is composed of images taken from different viewpoints of a scene, it can be difficult to percieve the difference between neighboring images as the difference between them can be relatively small.

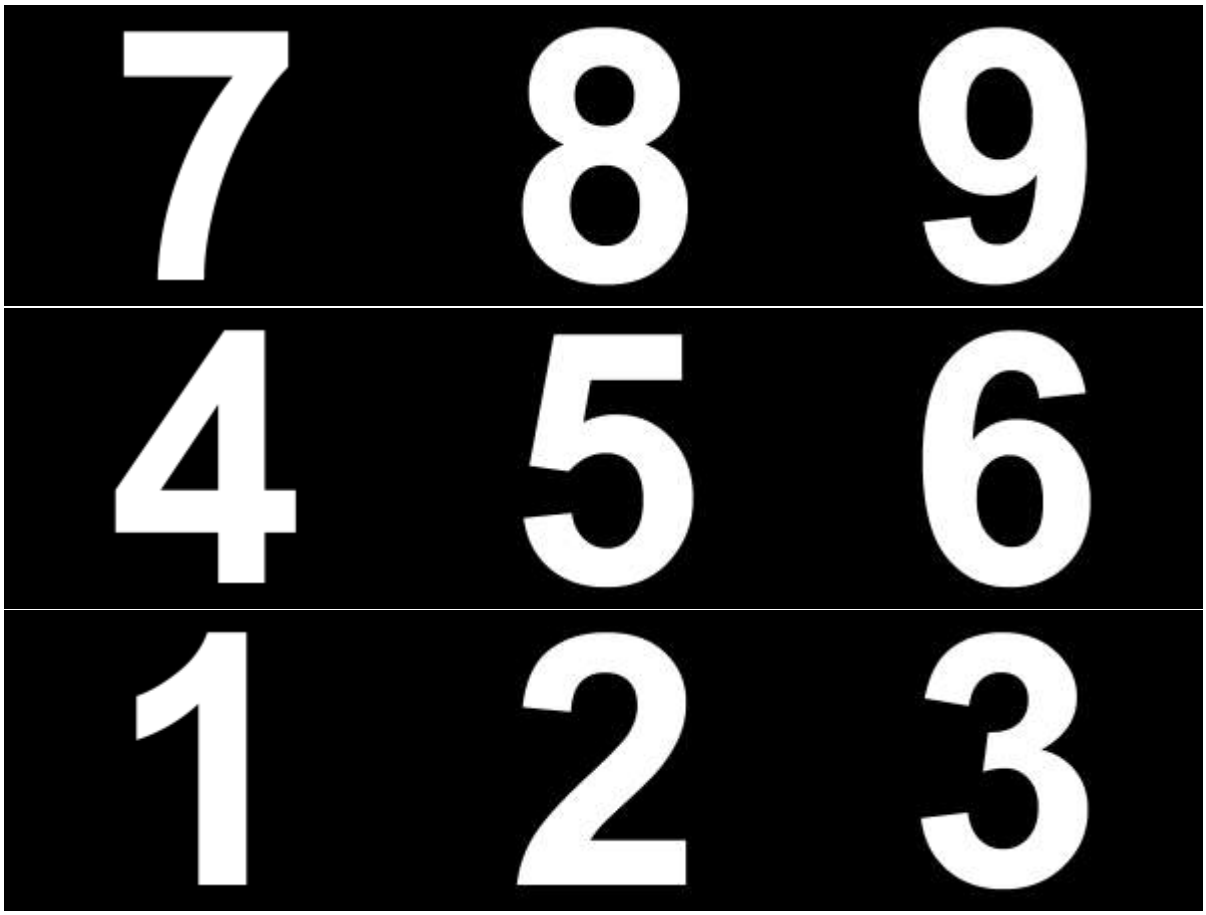

Figure 5.18: Numbers light field. All results in Section 5.2.3 are produced from this light field.

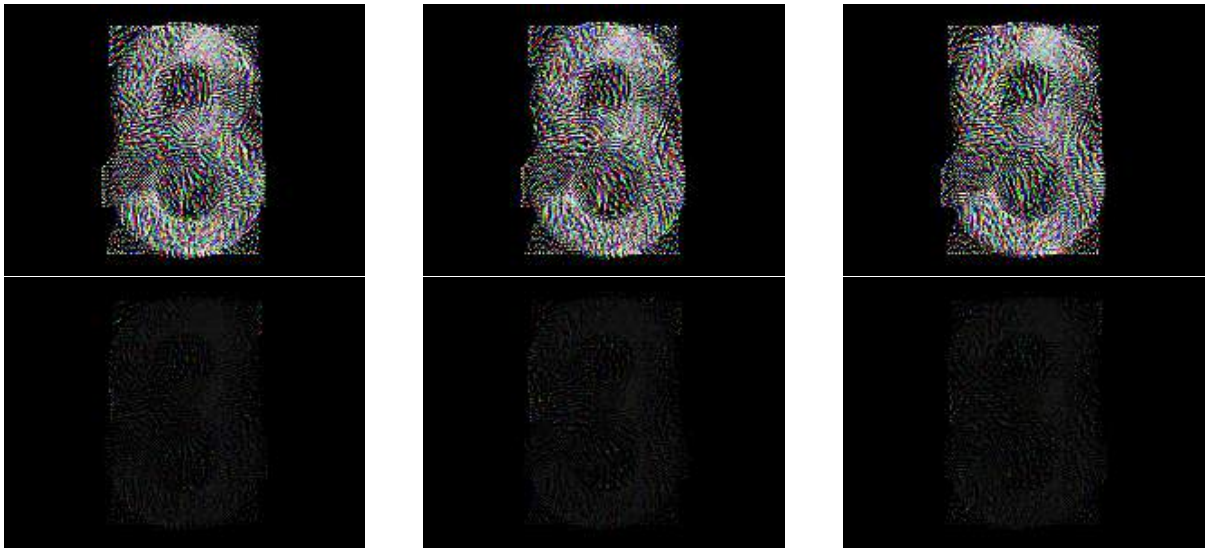

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

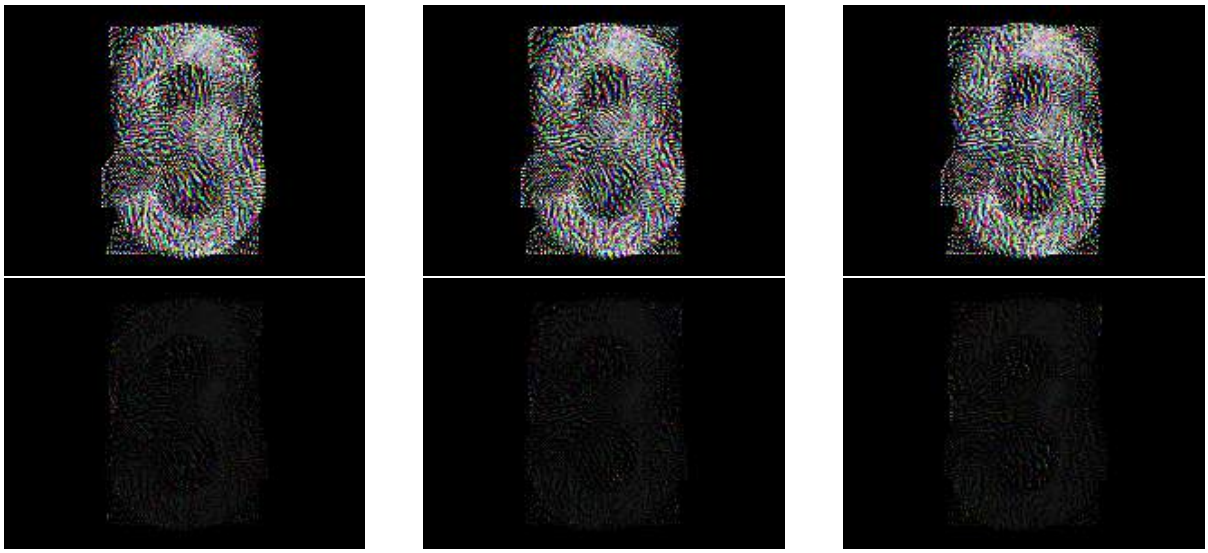

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

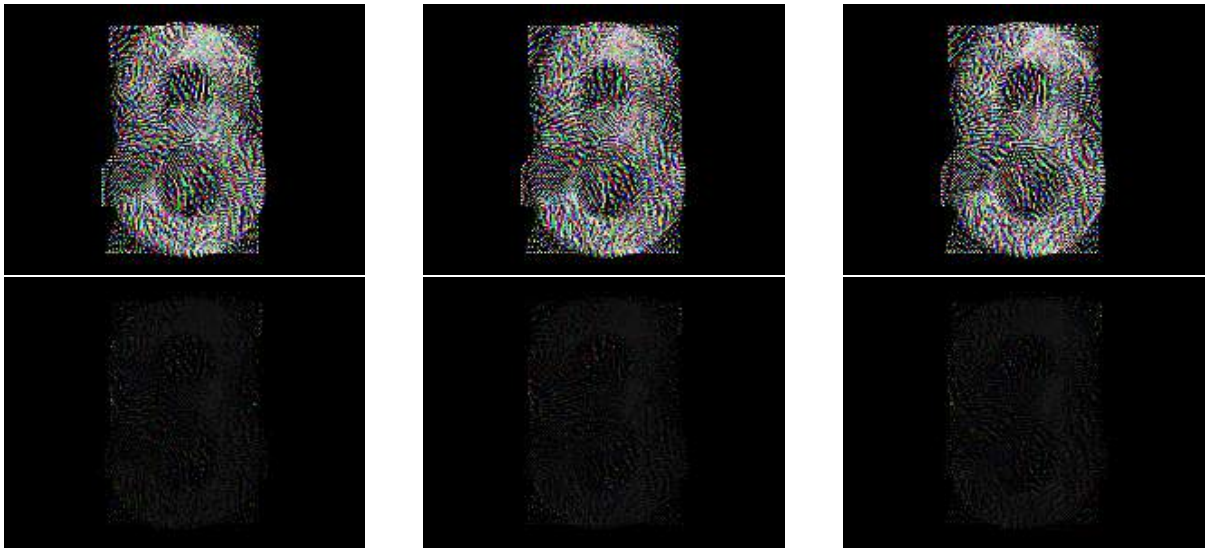

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.19: Masks produced by the Lee & Seung method, from [47], for the numbers light field.

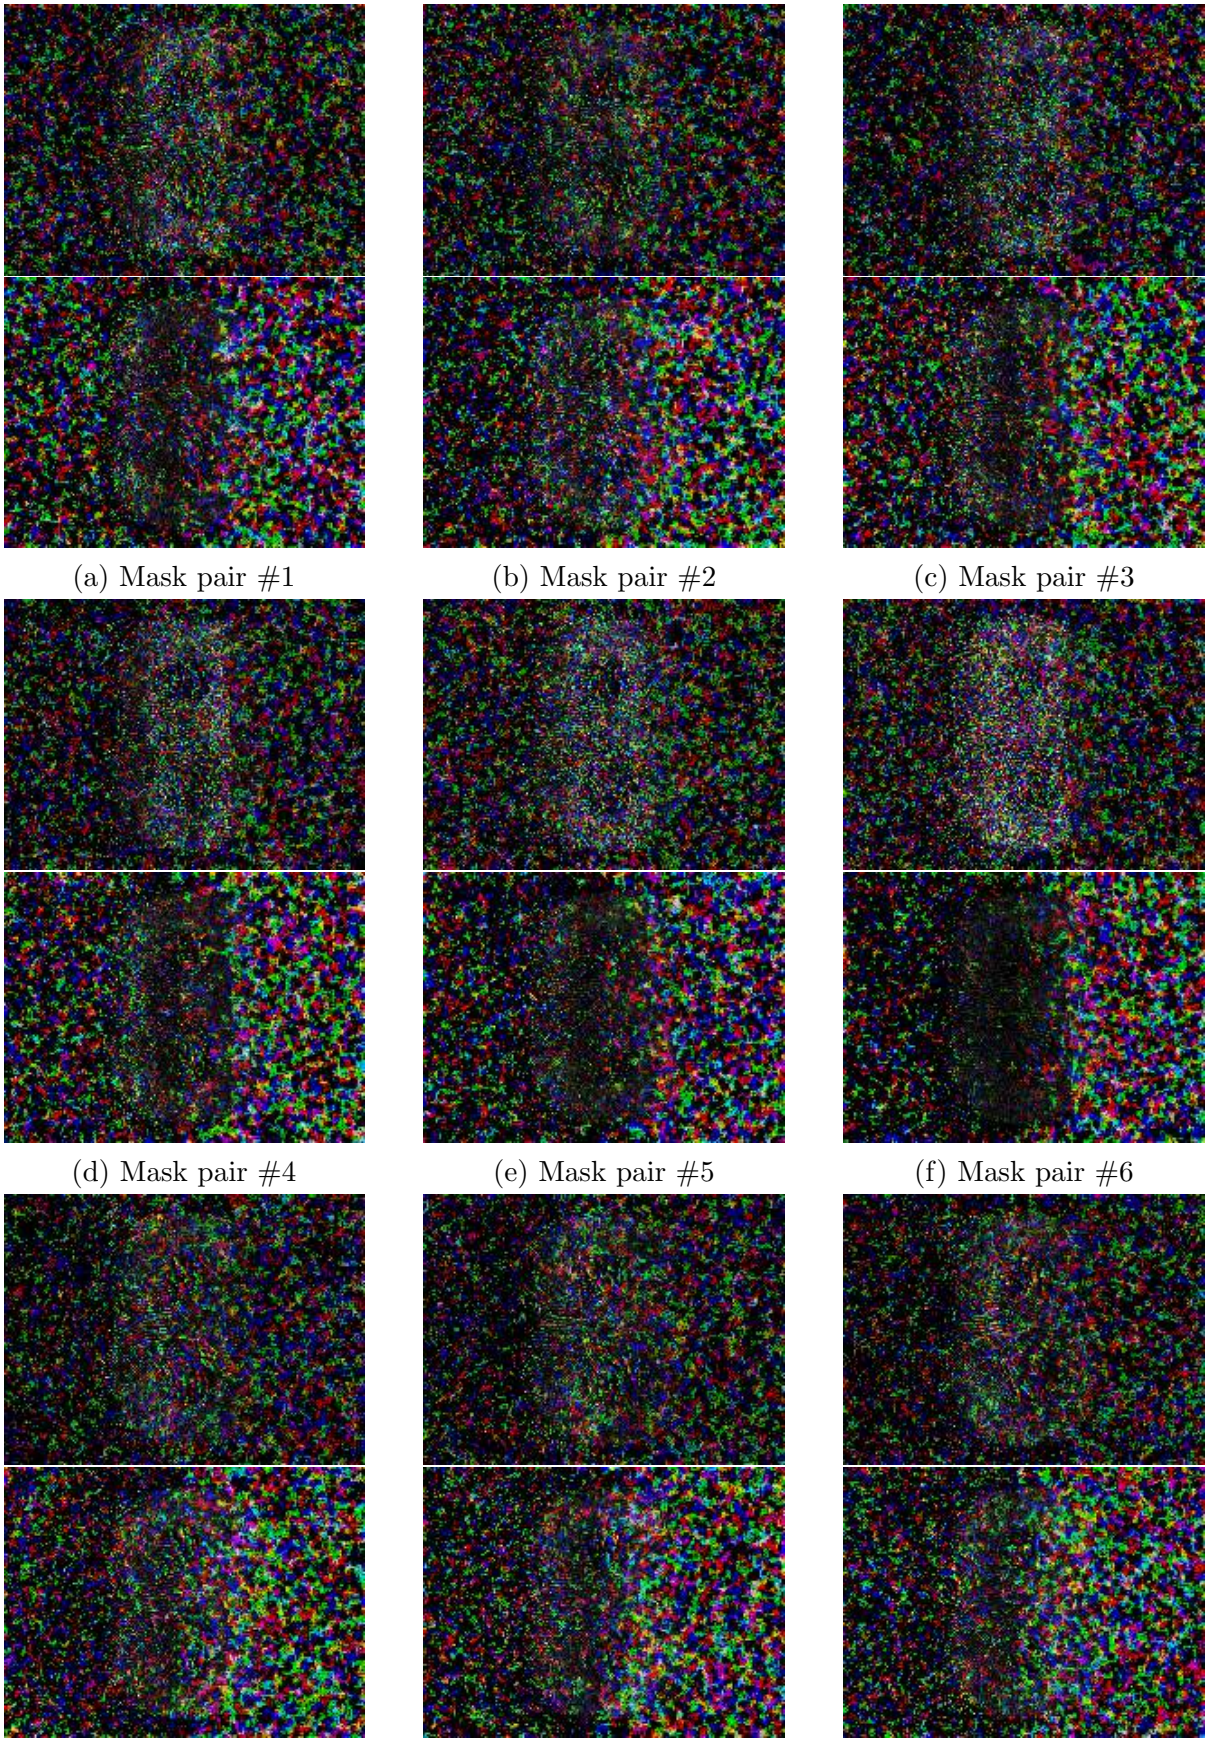

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.20: Masks produced from the WFNMAe method for the numbers light field.

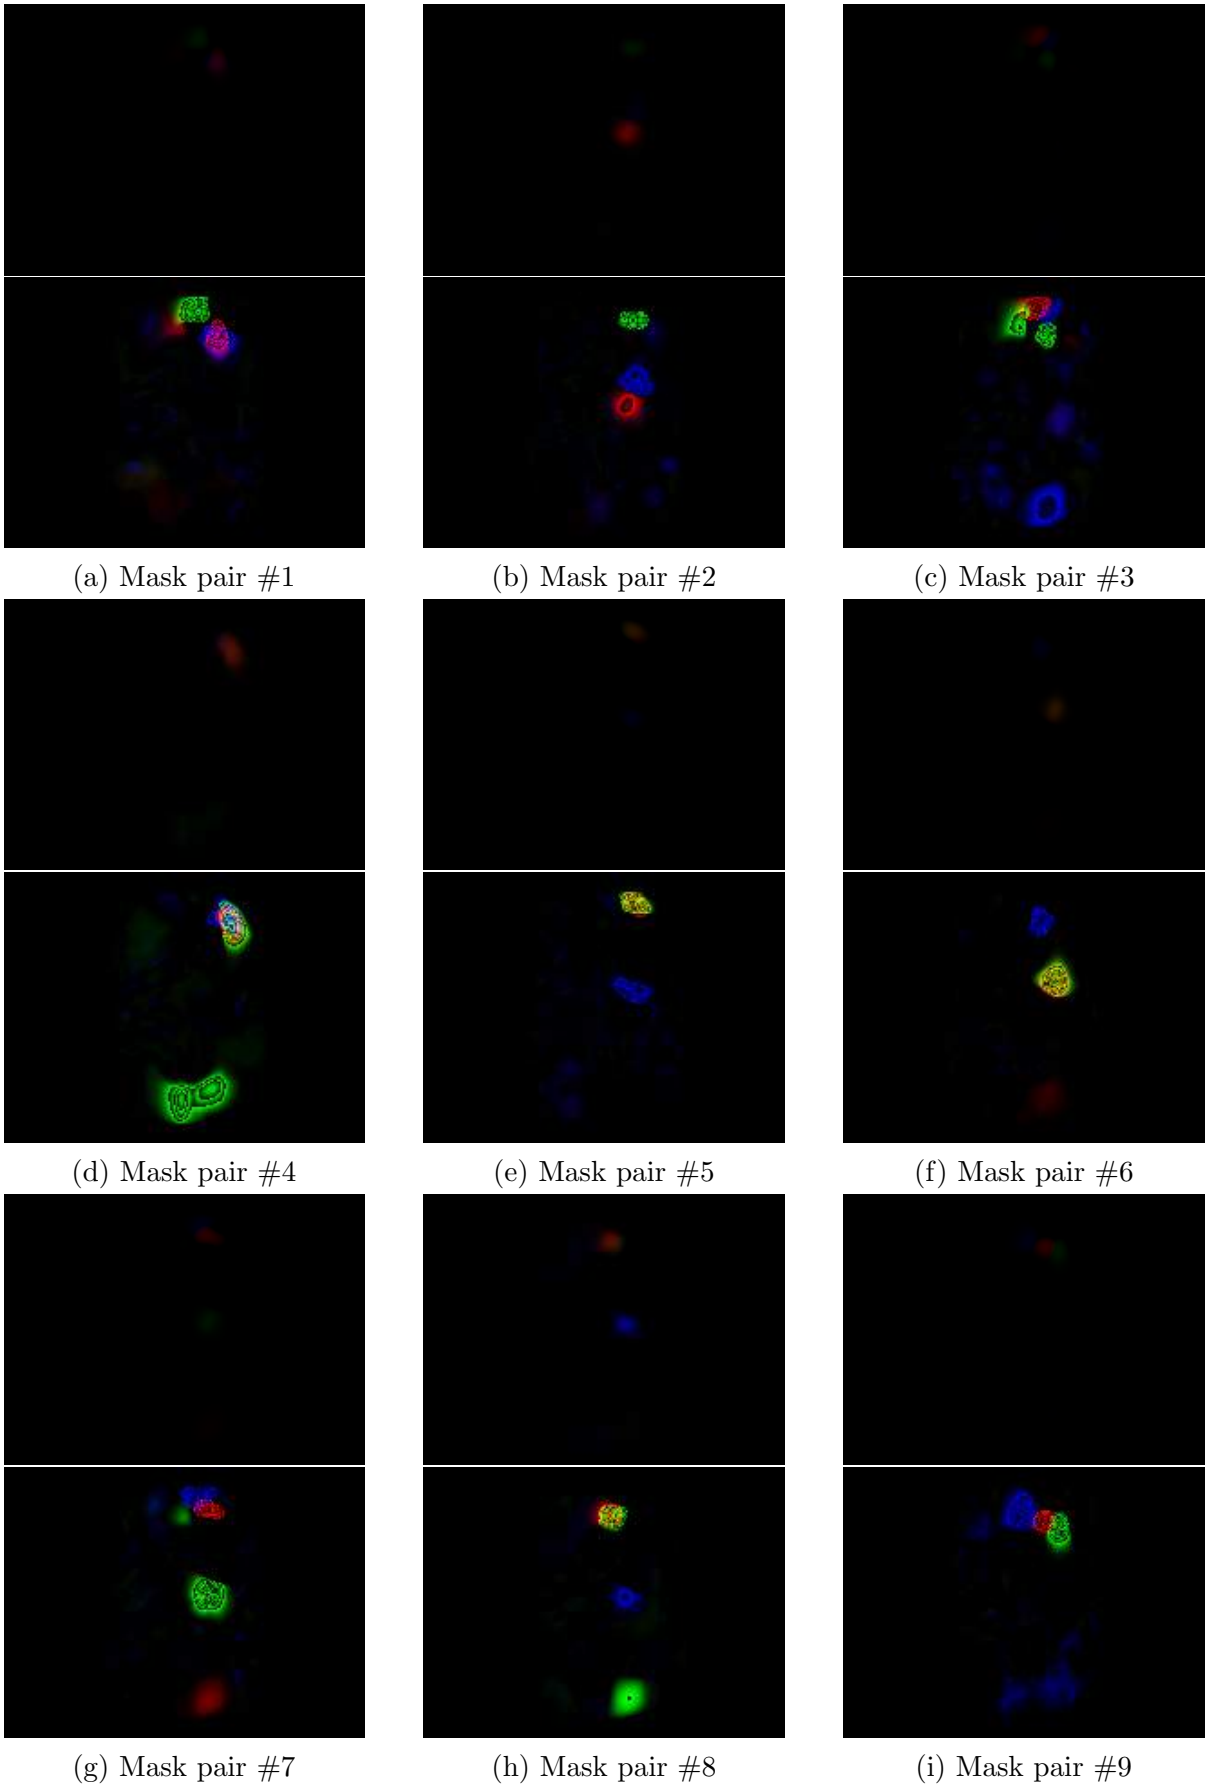

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.21: Masks produced by the FNMAe method, from [42], for the numbers light field.

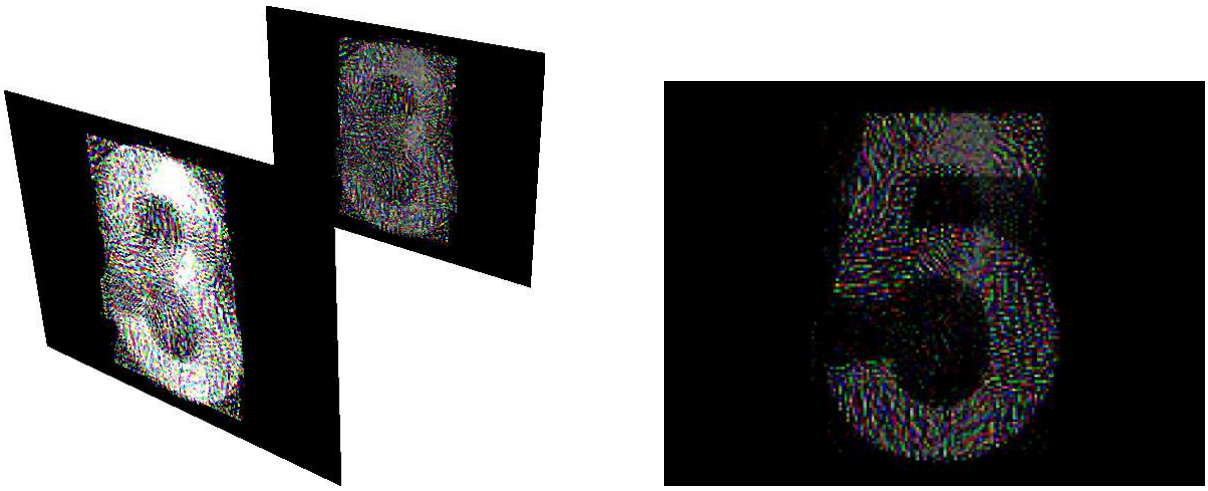

(a) Lee & Seung

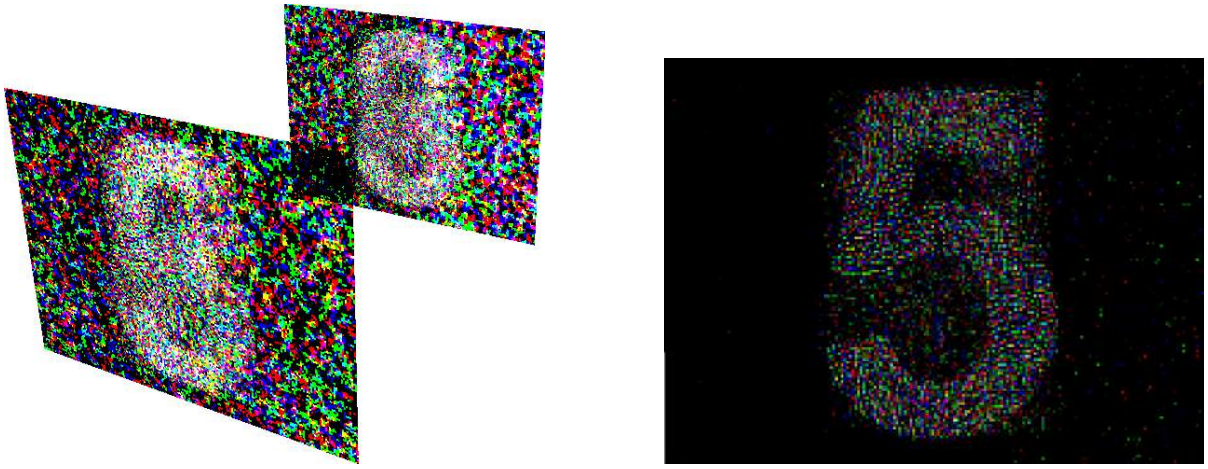

(b) WFNMAe

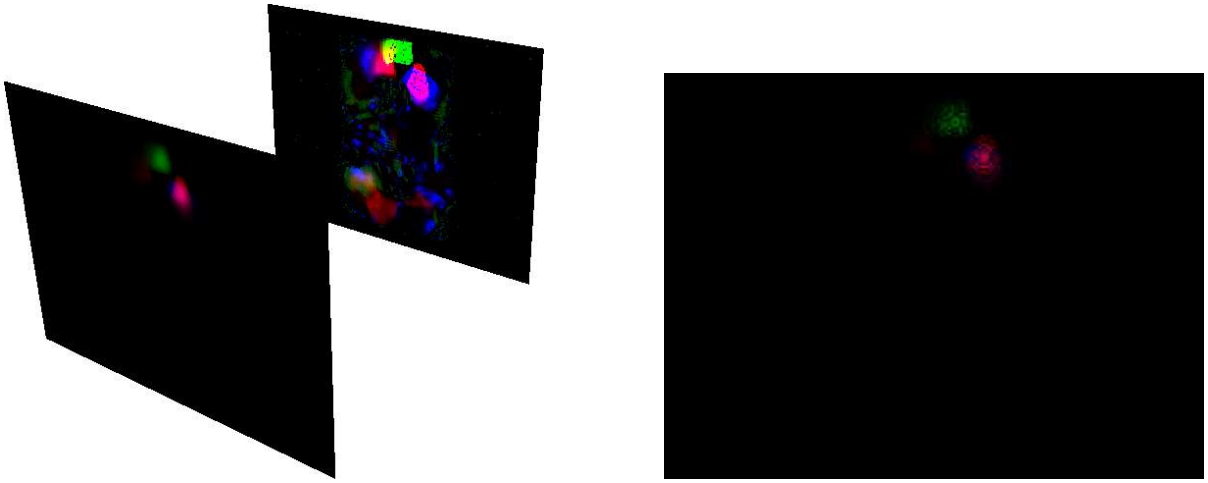

(c) FNMAe

Figure 5.22: Simulated results comparison for the numbers light field. (left) The front and back masks on their respective planes. (right) The simulated result

5.2.4 Spheres

The spheres light field was created by the content adaptive parallax barriers team at MIT [1]. The input images were resized, but the aspect ratio was preserved, to accomodate the requirements of FNMA_e. The images and masks have a resolution of 240×150. This light field was chosen due to it being used for some of the main results presented by the original authors in [45]. Using the spheres light field allows for the demonstration that we have accurately reproduced the original results with the original methods as well as giving a comparison of the other NMF methods that we explored.

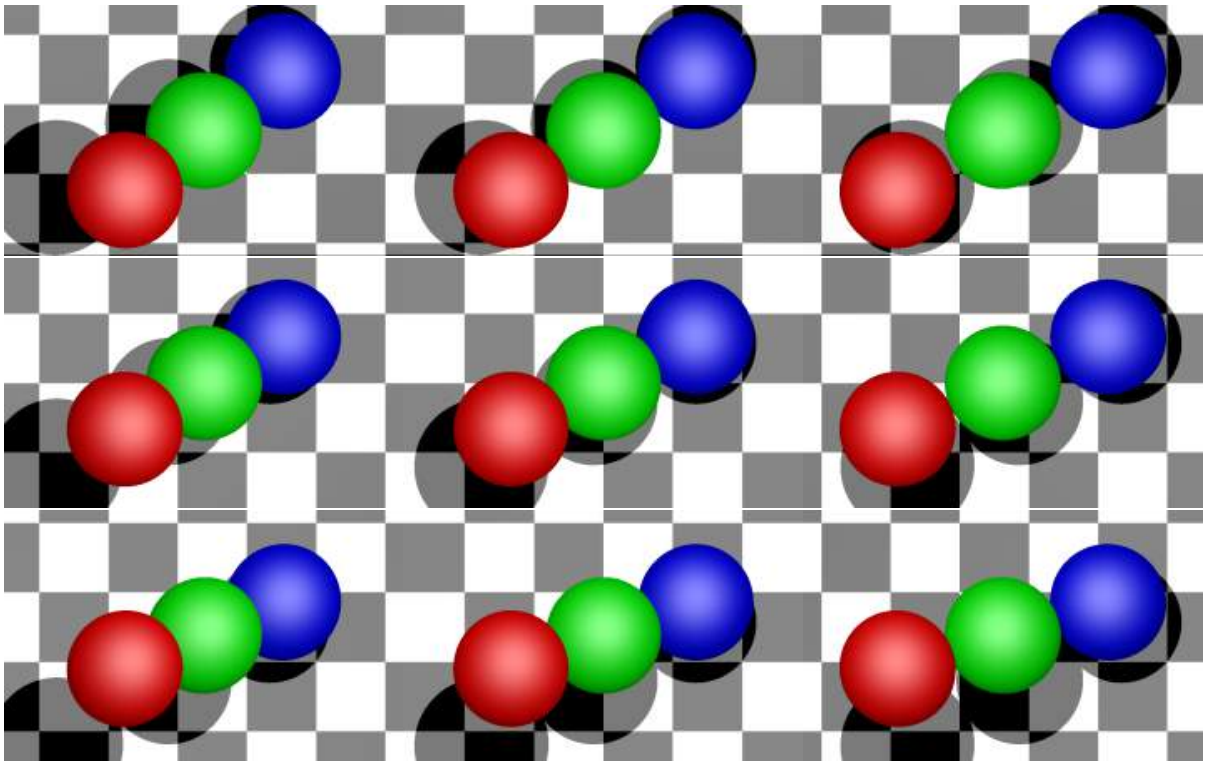

Figure 5.23: Spheres light field. All results in Section 5.2.4 are produced from this light field.

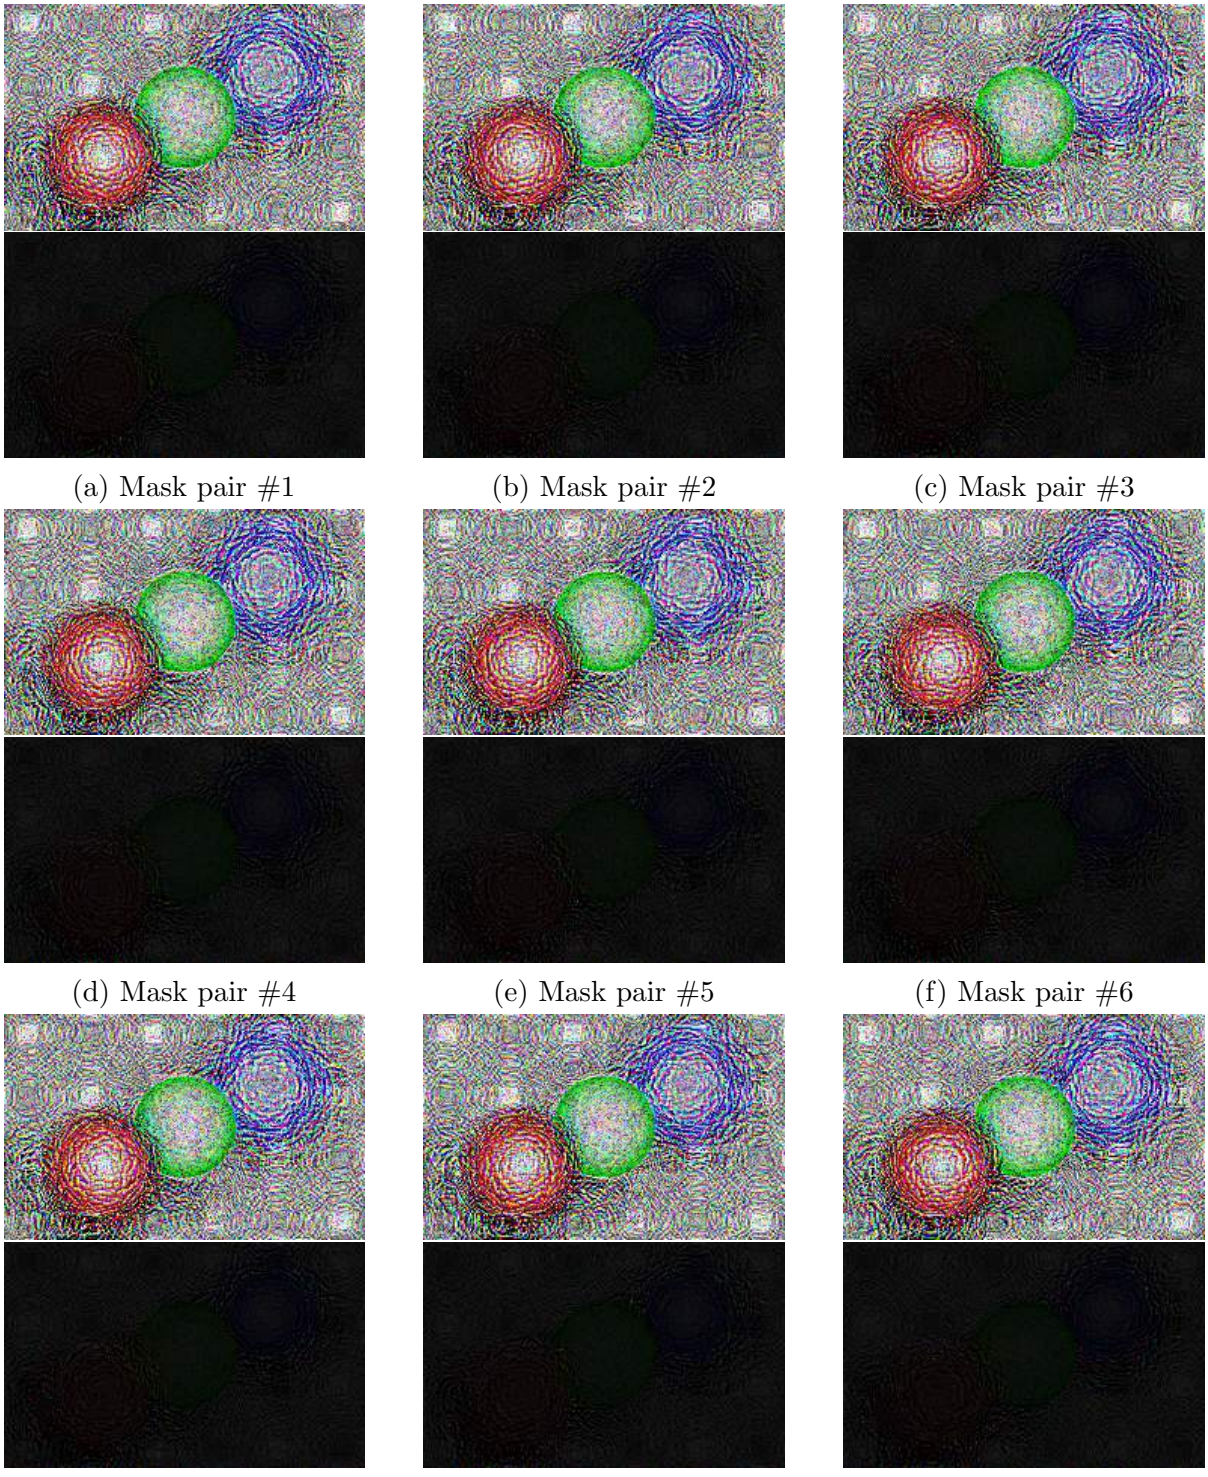

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.24: Masks produced by the Lee & Seung method, from [47], for the spheres light field.

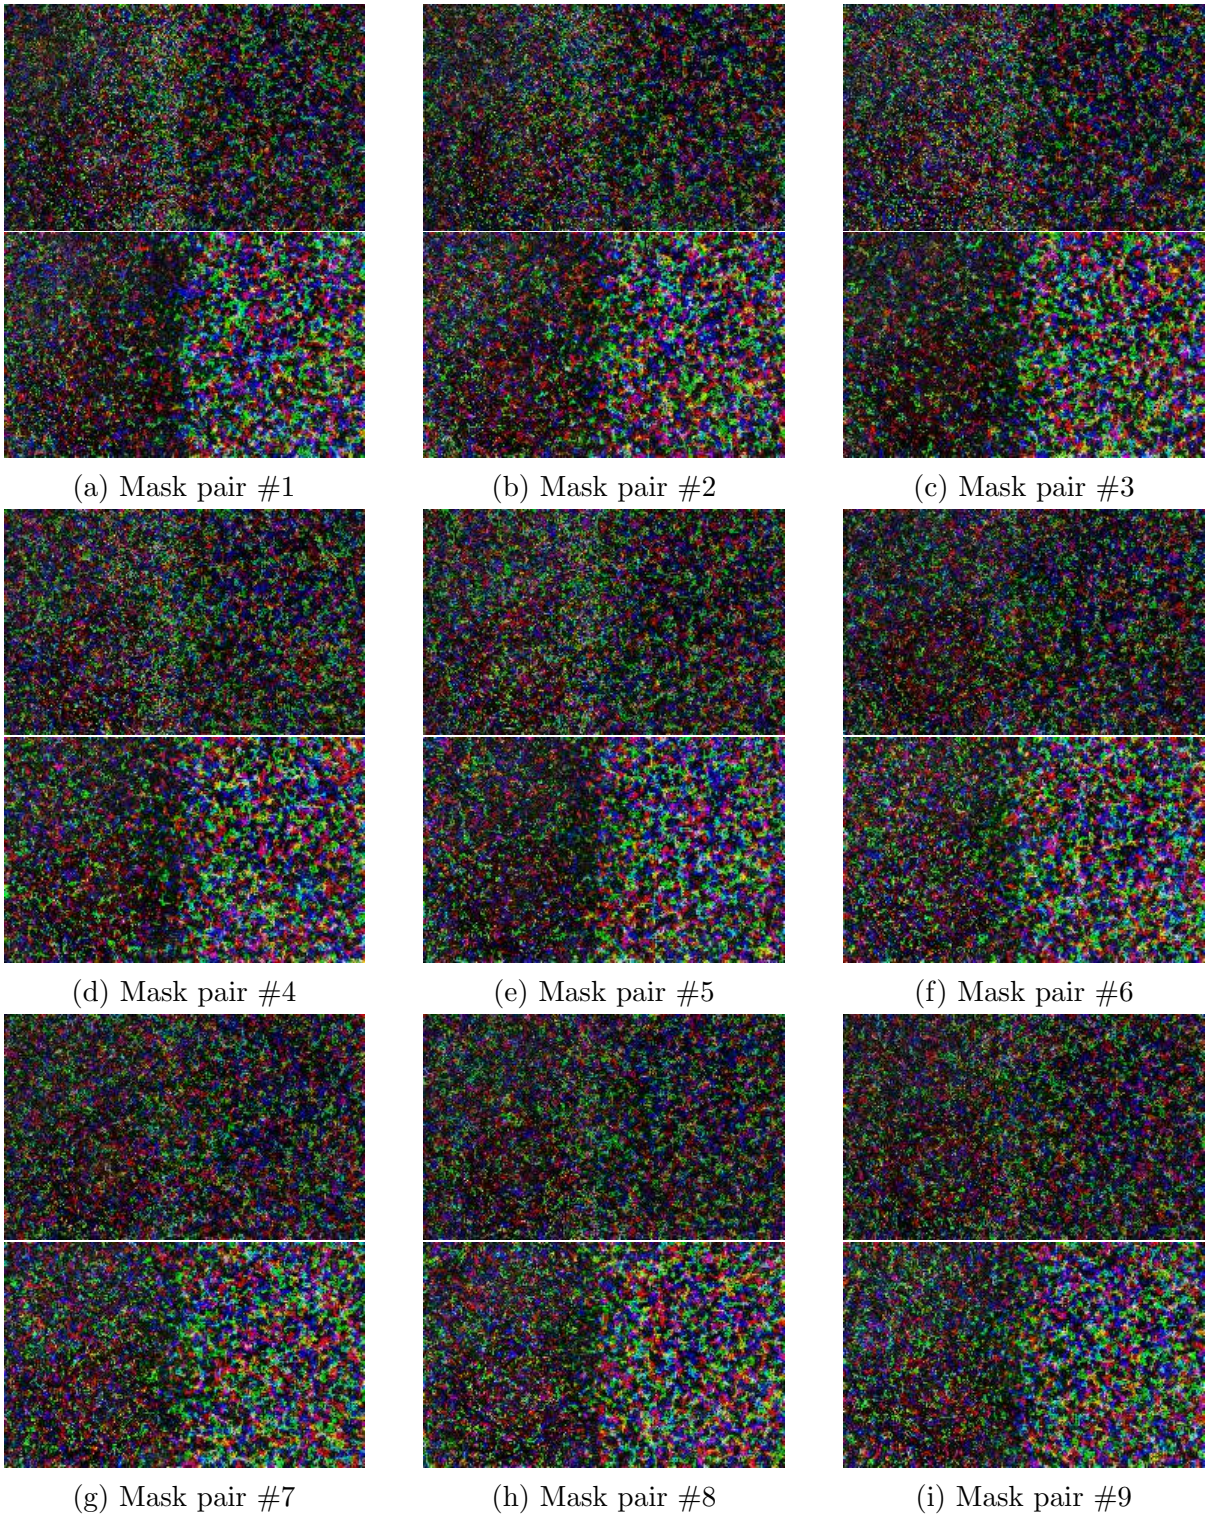

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.25: Masks produced from the WFNMAe method for the spheres light field.

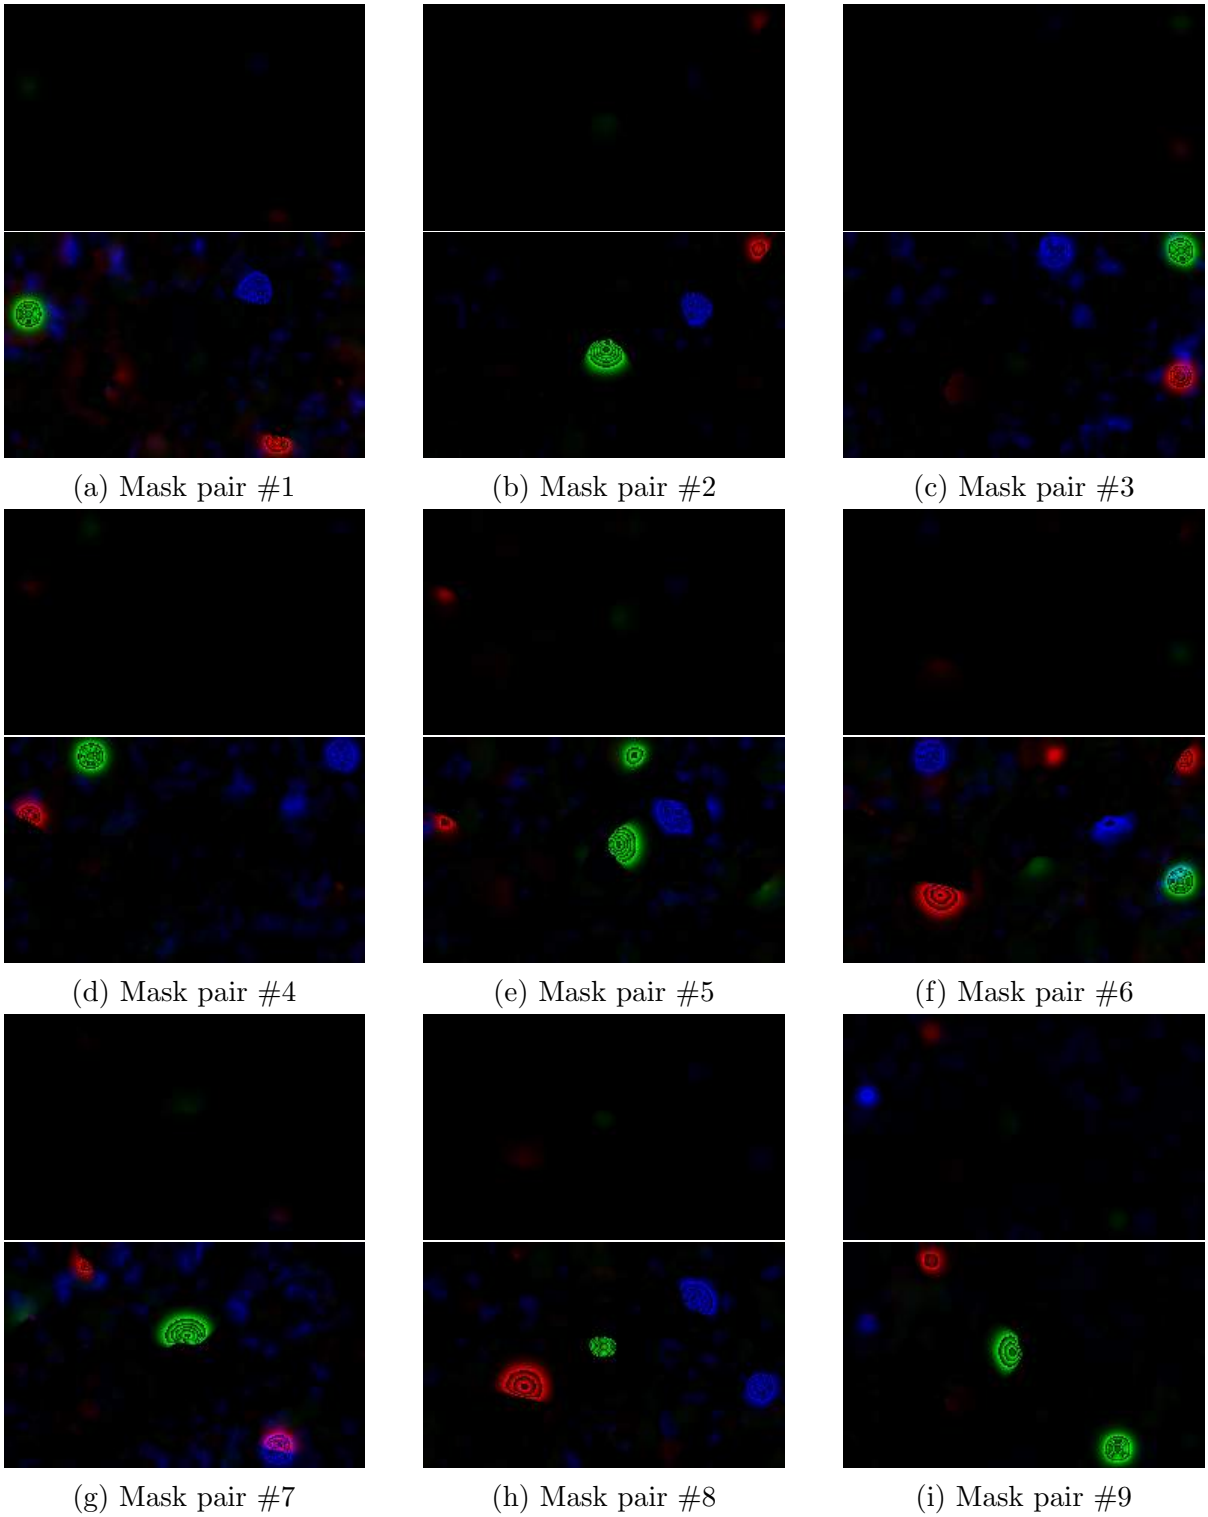

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.26: Masks produced by the FNMAe method, from [42], for the spheres light field.

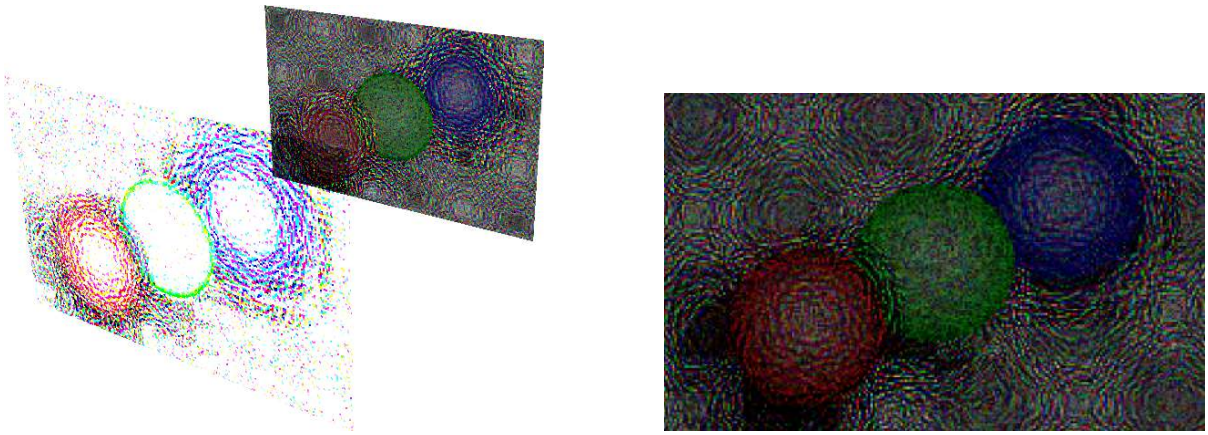

(a) Lee & Seung

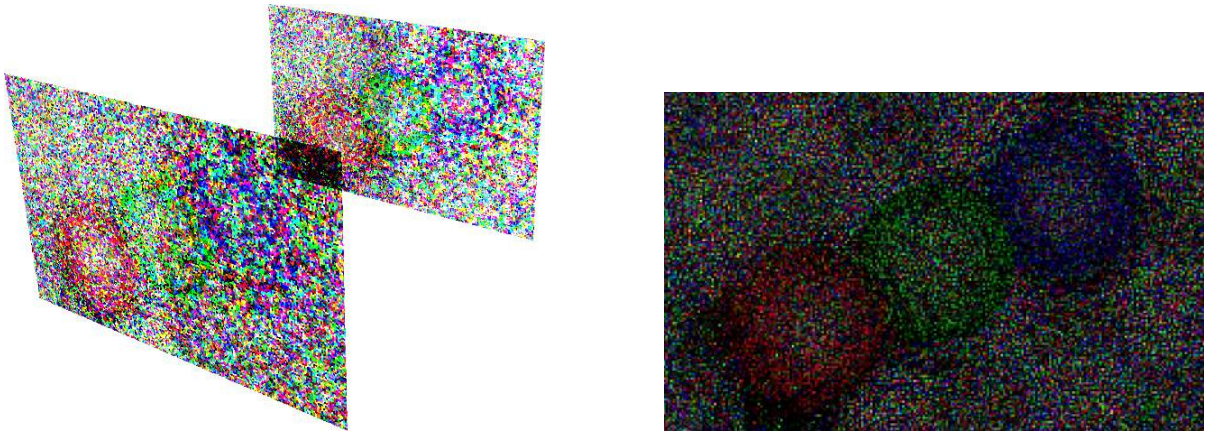

(b) WFNMAe

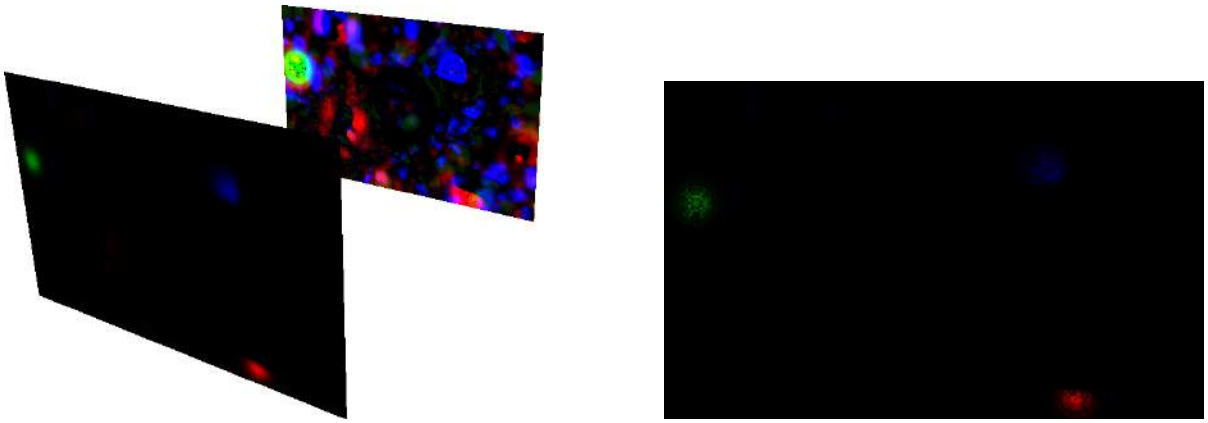

(c) FNMAe

Figure 5.27: Simulated results comparison for the spheres light field. (left) The front and back masks on their respective planes. (right) The simulated result

5.2.5 Distinct Images

The distinct images light field, seen in Figure 5.28, was created by obtaining unrelated images from a public domain website [34] to use instead of an expected set of multi-view images. The input images were resized to accomodate the memory requirements of FNMA_e. The images and resultant masks have a resolution of 266×150. We used this dataset to highlight that the current technique does not take advantage of the fact that the input is a light field, which has a large amount of redundant data due to shared pixel values between similar views. An approach that takes advantage of the redundancy in light fields would not support a light field composed of distinct images as it has a distinct lack of redundancy.

The masks produced, using the Lee and Seung method [47], for this light field are unique when comparing them to the masks, for the same technique on other datasets. The masks produced for the other datasets resemble the images present in the light field, whereas in this set they do not match any particular image from the distinct images light field. This observation further highlights the redundancy that is present in a light field that would be expected in most practical applications of the content adaptive parallax barrier approach.

The distinct images helps with identifying issues present in the current work that are not as easily identifiable with a normal light field. One example is that our current approach has issues with the frequency of the input images. During simulation we can see that the eclipse image, which is a relatively low frequency image, is very difficult to perceive without already knowing what it is supposed to be. In addition the water and bacteria images are less identifiable than a more balanced imaged. The frequency issue is less easily noticed in a proper light field because all of the images are similar which means that any small pitfalls may be masked across their similarity.

Figure 5.28: Distinct images light field. All results in Section 5.2.5 are produced from this light field. (bottom-left) Singapore Cultural Centre[1] (bottom-middle) Tank[2] (bottom-right) Flower Vase[3] (middle-left) Water[4] (middle) Golf Ball Galaxy[5] (middle-right) Orion Nebula[6] (top-left) Eclipse[7] (top-middle) Bacteria[8] (top-right) Cheetahs[9]

[1]Singapore cultural centre at night public domain image, Andrew McMillan (http://www.public-domain-image.com/free-images/architecture/city-downtown/singapore-cultural-centre-at-night) Accessed on 16 September 2015. Public Domain.

[2]M24 chaffee light tank free stock image, Greg Goebel (http://www.public-domain-image.com/free-images/transportation-vehicles/tanks/m24-chaffee-light-tank/attachment/m24-chaffee-light-tank) Accessed on 16 September 2015. Public Domain.

[3]Flowers tables vase royalty free stock photograph, Jon Sullivan (http://www.public-domain-image.com/free-images/still-life/flowers-tables-vase/attachment/flowers-tables-vase) Accessed on 16 September 2015. Public Domain.

[4]Water wallpaper public domain wallpaper, Private / anonymous (http://www.public-domain-image.com/free-images/wallpapers/water-wallpaper/attachment/water-wallpaper) Accessed on 16 September 2015. Public Domain.

[5]Golf ball galaxy an artistic image free stock photo, Luke (http://www.public-domain-image.com/free-images/computer-arts/3d-computer-graphics/golf-ball-galaxy-an-artistic-image/attachment/golf-ball-galaxy-an-artistic-image) Accessed on 16 September 2015. Public Domain.

[6]Orion nebula space galaxy royalty free stock photograph, Private / anonymous (http://www.public-domain-image.com/free-images/space/orion-nebula-space-galaxy/attachment/orion-nebula-space-galaxy) Accessed on 16 September 2015. Public Domain.

[7]Sun eclipses royalty free stock image, Jon Sullivan (http://www.public-domain-image.com/free-images/miscellaneous/sun/sun-eclipses/attachment/sun-eclipses) Accessed on 16 September 2015. Public Domain.

[8]Grouping of gram negative anaerobic borrelia burgdorferi bacteria free picture, Jamice Haney Carr, Claudia Molins, USCDCP (http://www.public-domain-image.com/free-images/science/microscopy-images/borrelia-burgdorferi/grouping-of-gram-negative-anaerobic-borrelia-burgdorferi-bacteria/attachment/grouping-of-gram-negative-anaerobic-borrelia-burgdorferi-bacteria) Accessed on 16 September 2015. Public Domain.

[9]Two cheetahs African animals acinonyx jubatus facing each other copyright friendly photo, Stolz Gary M, U.S. Fish and Wildlife Service (http://www.public-domain-image.com/free-images/fauna-animals/cheetahs-leopards-jaguars-panthers-pictures/two-cheetahs-african-animals-acinonyx-jubatus-facing-each-other/attachment/two-cheetahs-african-animals-acinonyx-jubatus-facing-each-other) Accessed on 16 September 2015. Public Domain.

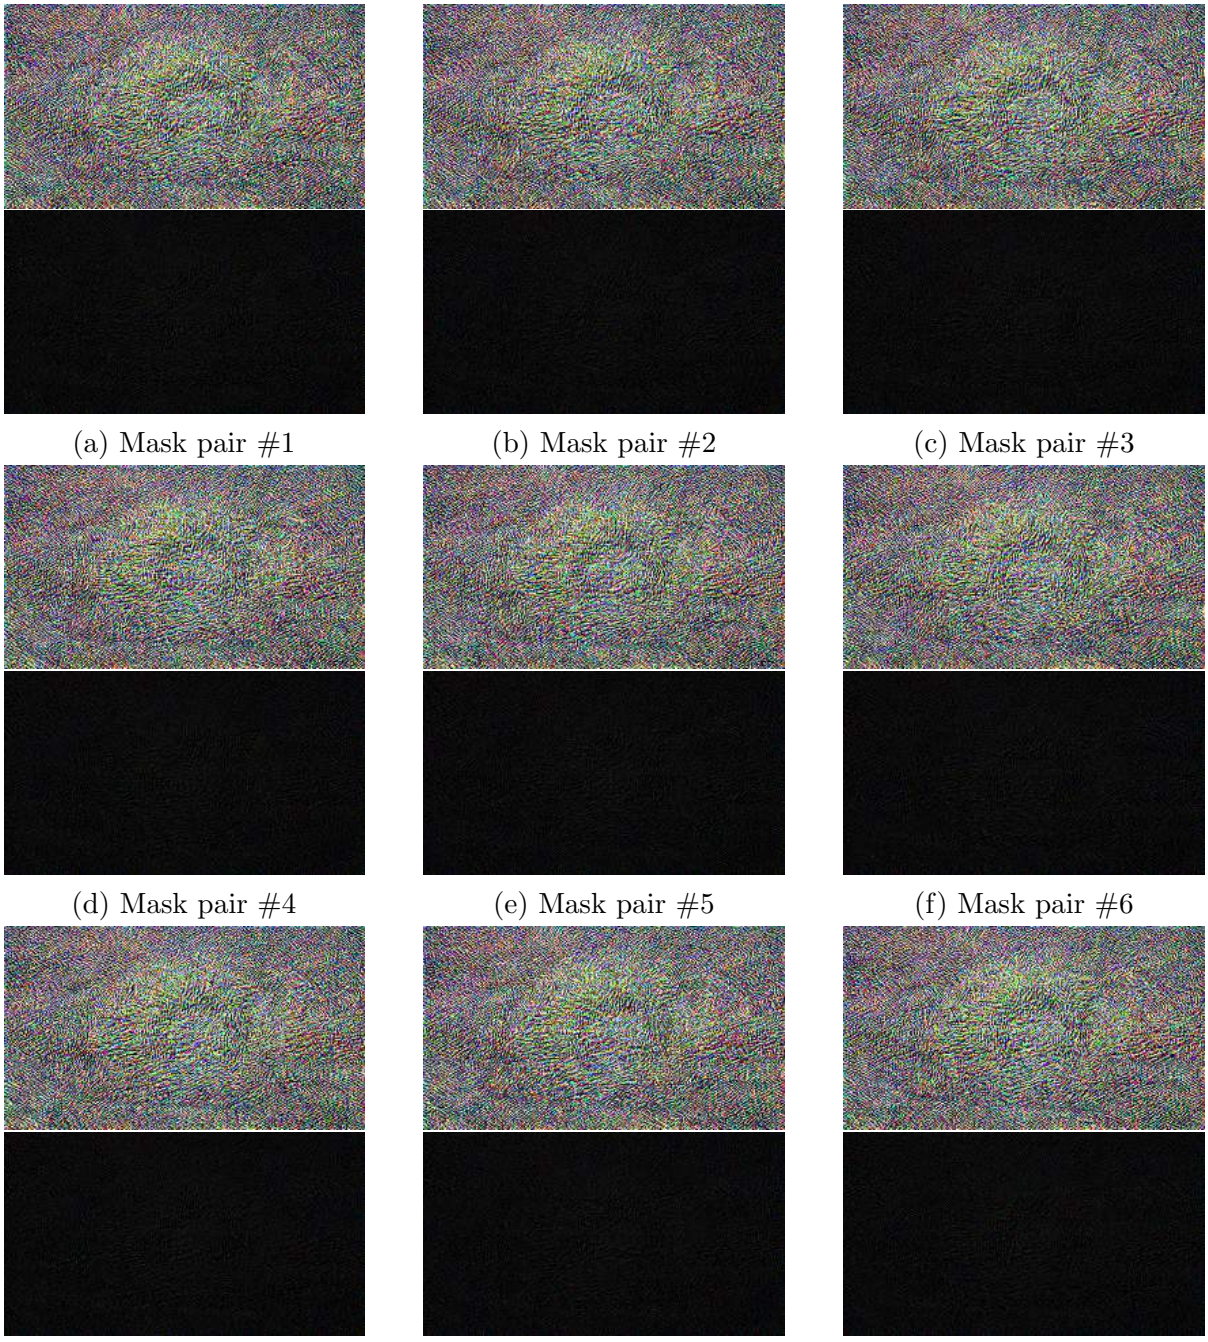

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.29: Masks produced by the Lee & Seung method, from [47], for the teapot light field.

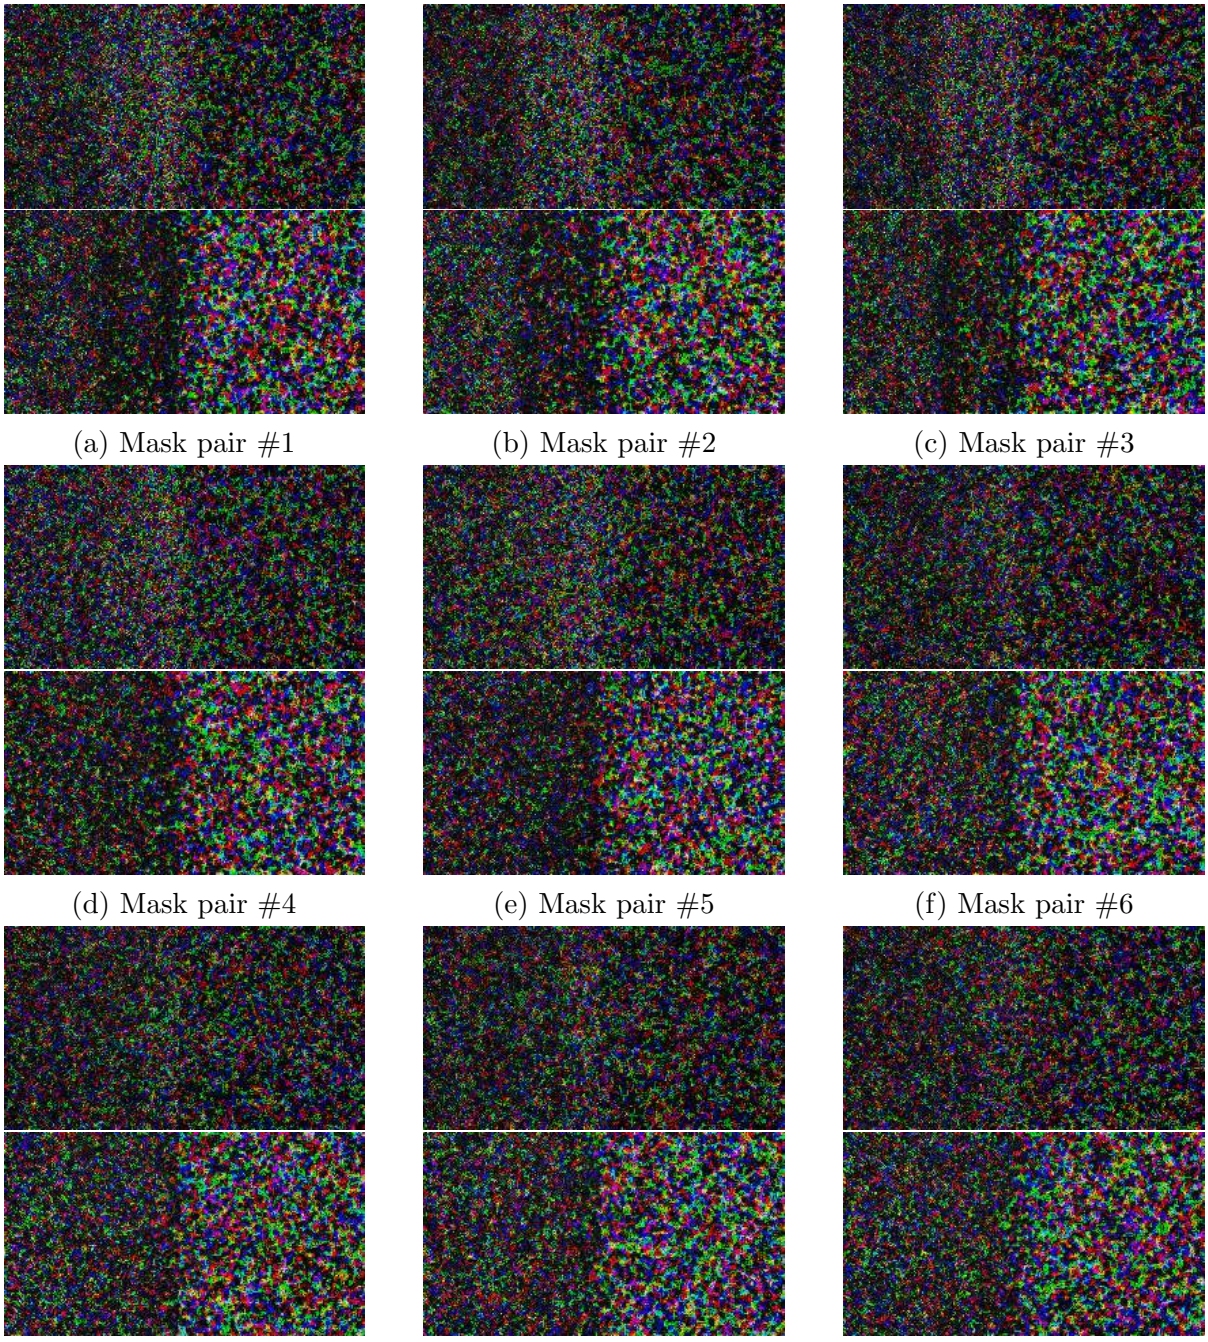

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.30: Masks produced from the WFNMAe method for the teapot light field.

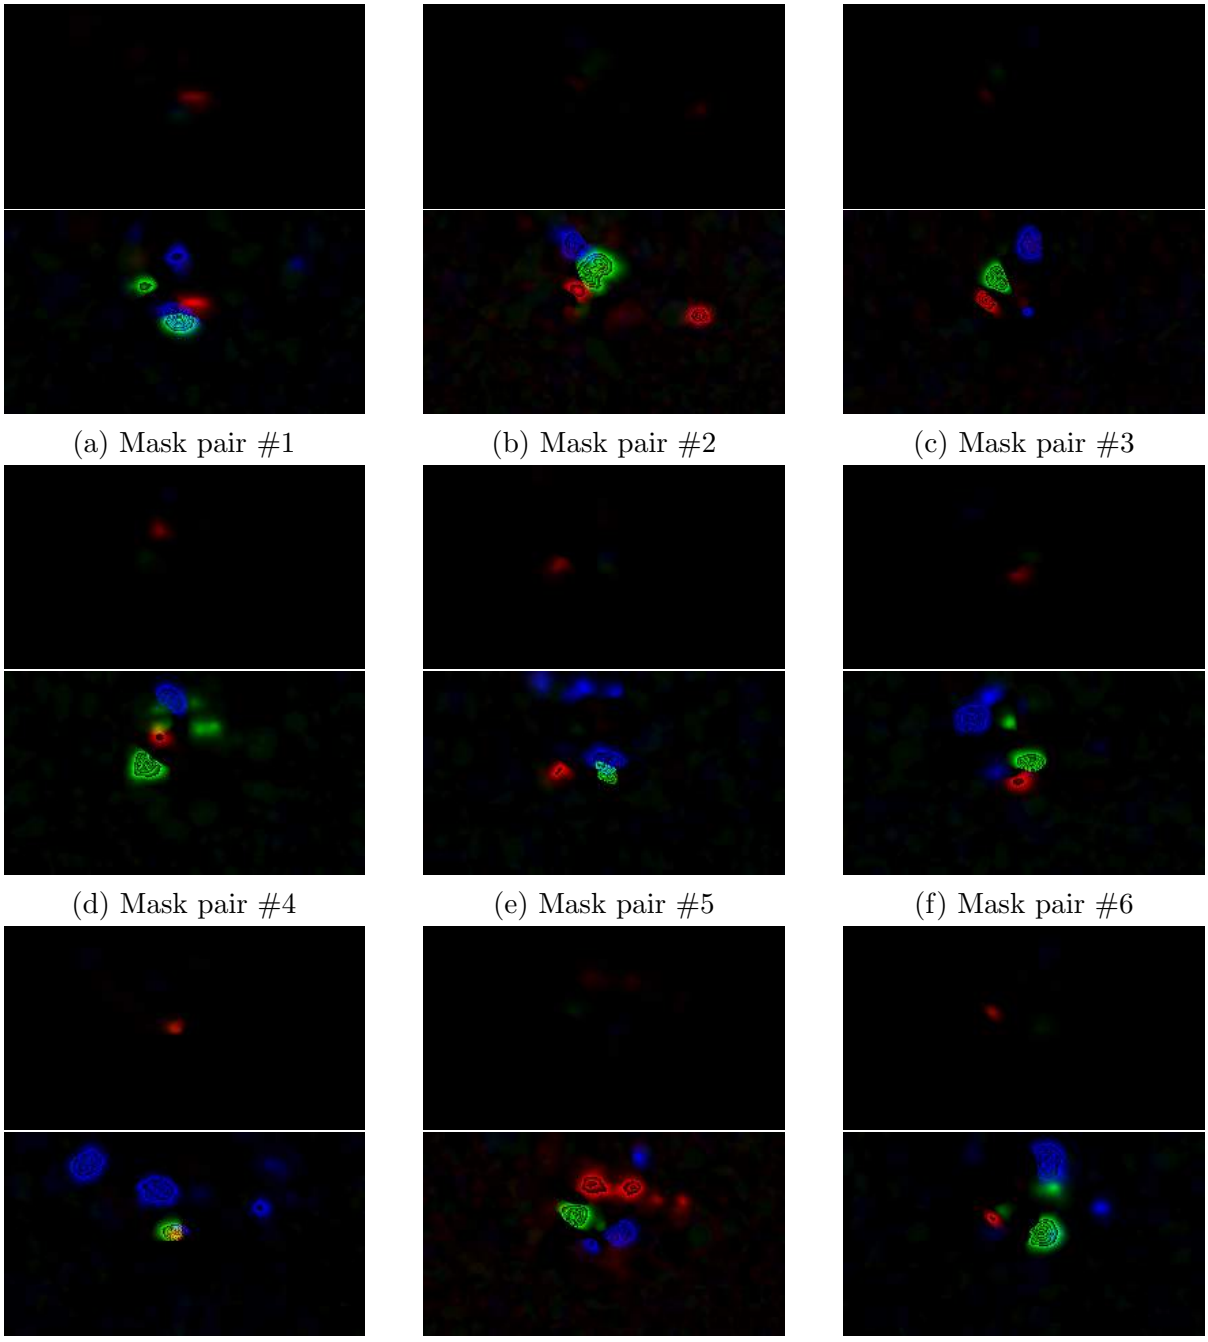

(a) Mask pair #1 (b) Mask pair #2 (c) Mask pair #3

(d) Mask pair #4 (e) Mask pair #5 (f) Mask pair #6

(g) Mask pair #7 (h) Mask pair #8 (i) Mask pair #9

Figure 5.31: Masks produced by the FNMAe method, from [42], for the teapot light field.

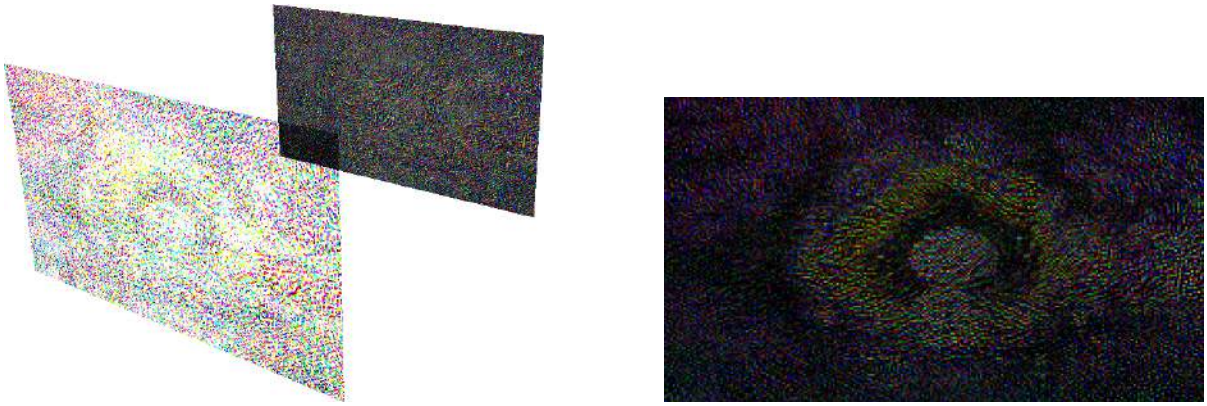

(a) Lee & Seung

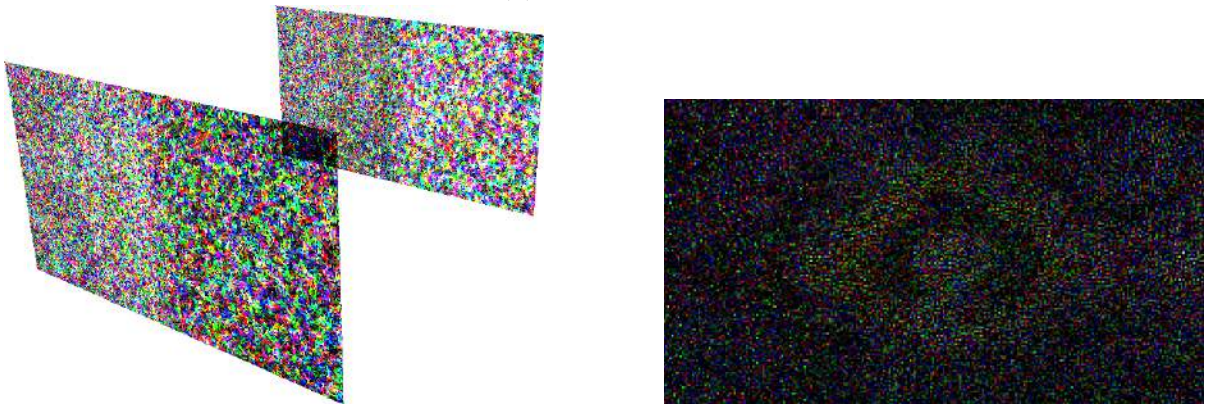

(b) WFNMAe

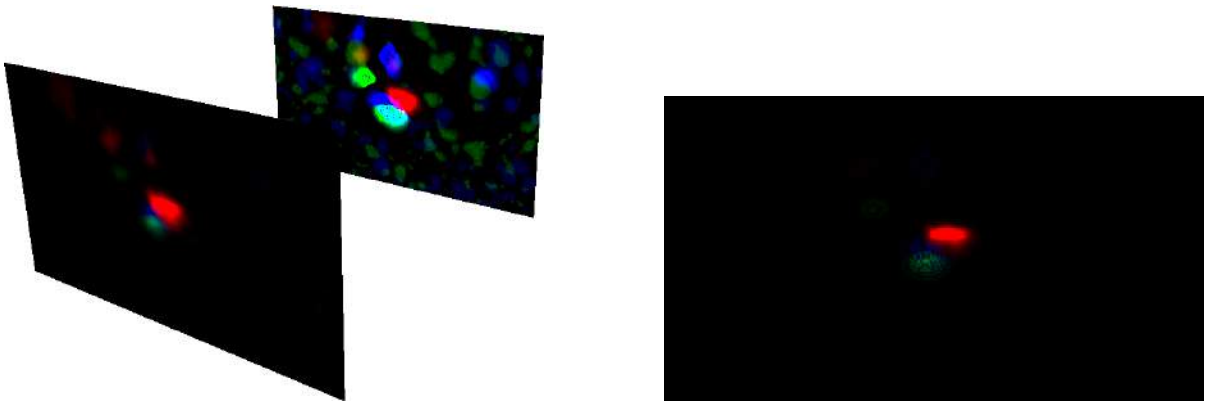

(c) FNMAe

Figure 5.32: Simulated results comparison for the teapot light field. (left) The front and back masks on their respective planes. (right) The simulated result

5.3 Orthographic vs. Perspective

We previously stated that there is an orthographic assumption in content adaptive parallax barriers when creating **L**. Our simulator has the capability to view multi-layer displays using either an orthographic or perspective projection. This allows a visual comparison between the intended and more realistic results. Figure 5.33 demonstrates the nine views produced by the distinct images light field, seen in Section 5.2.5, using an orthographic projection and Figure 5.34 shows these same views using a perspective projection.

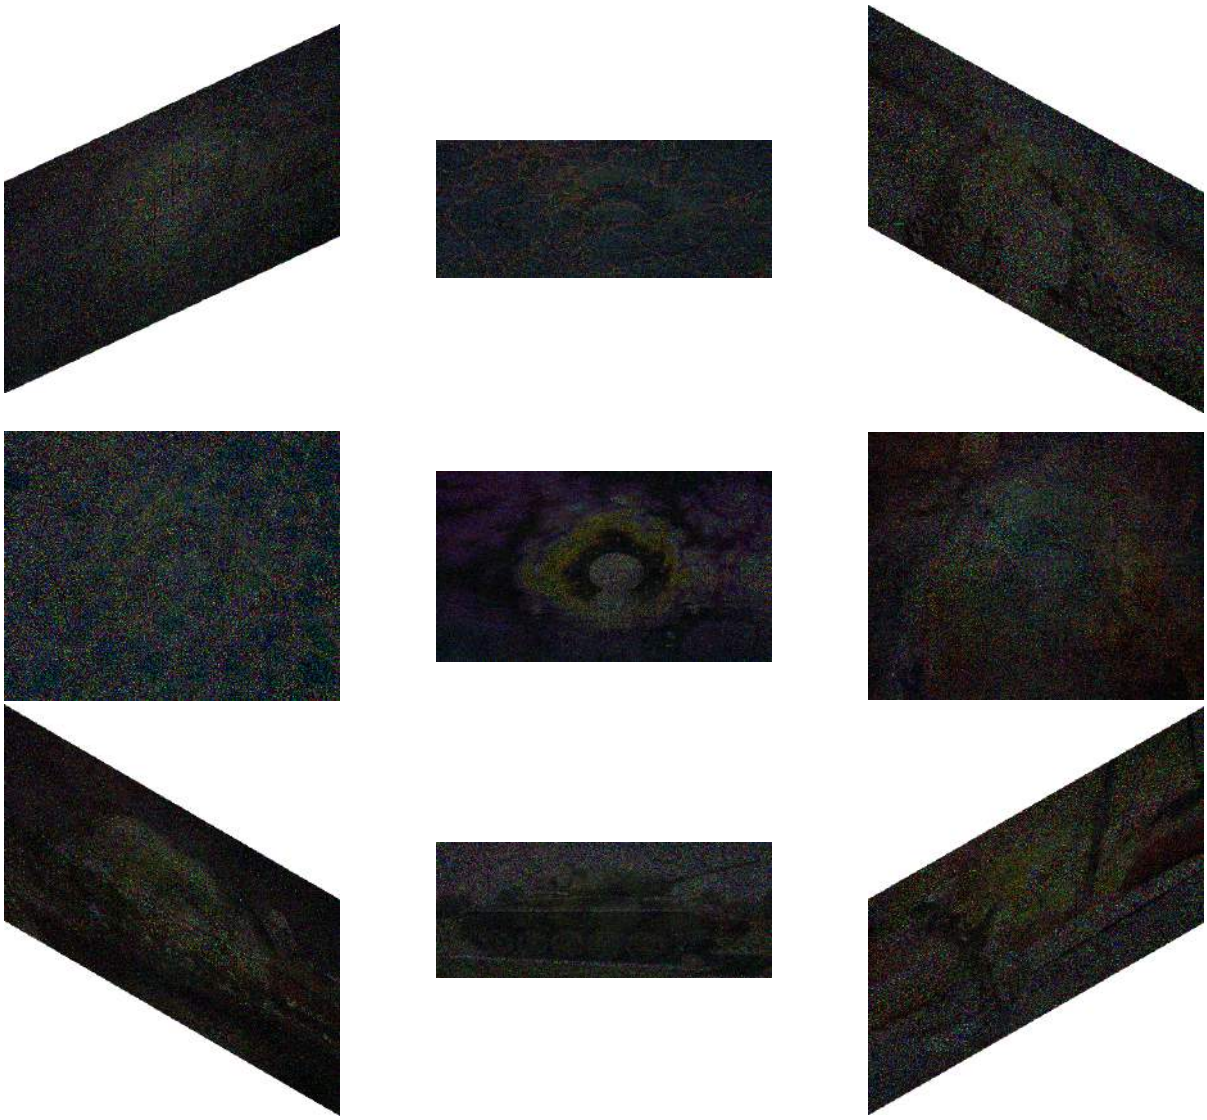

Figure 5.33: Simulated results using an orthographic projection. The input is from Figure 5.28.

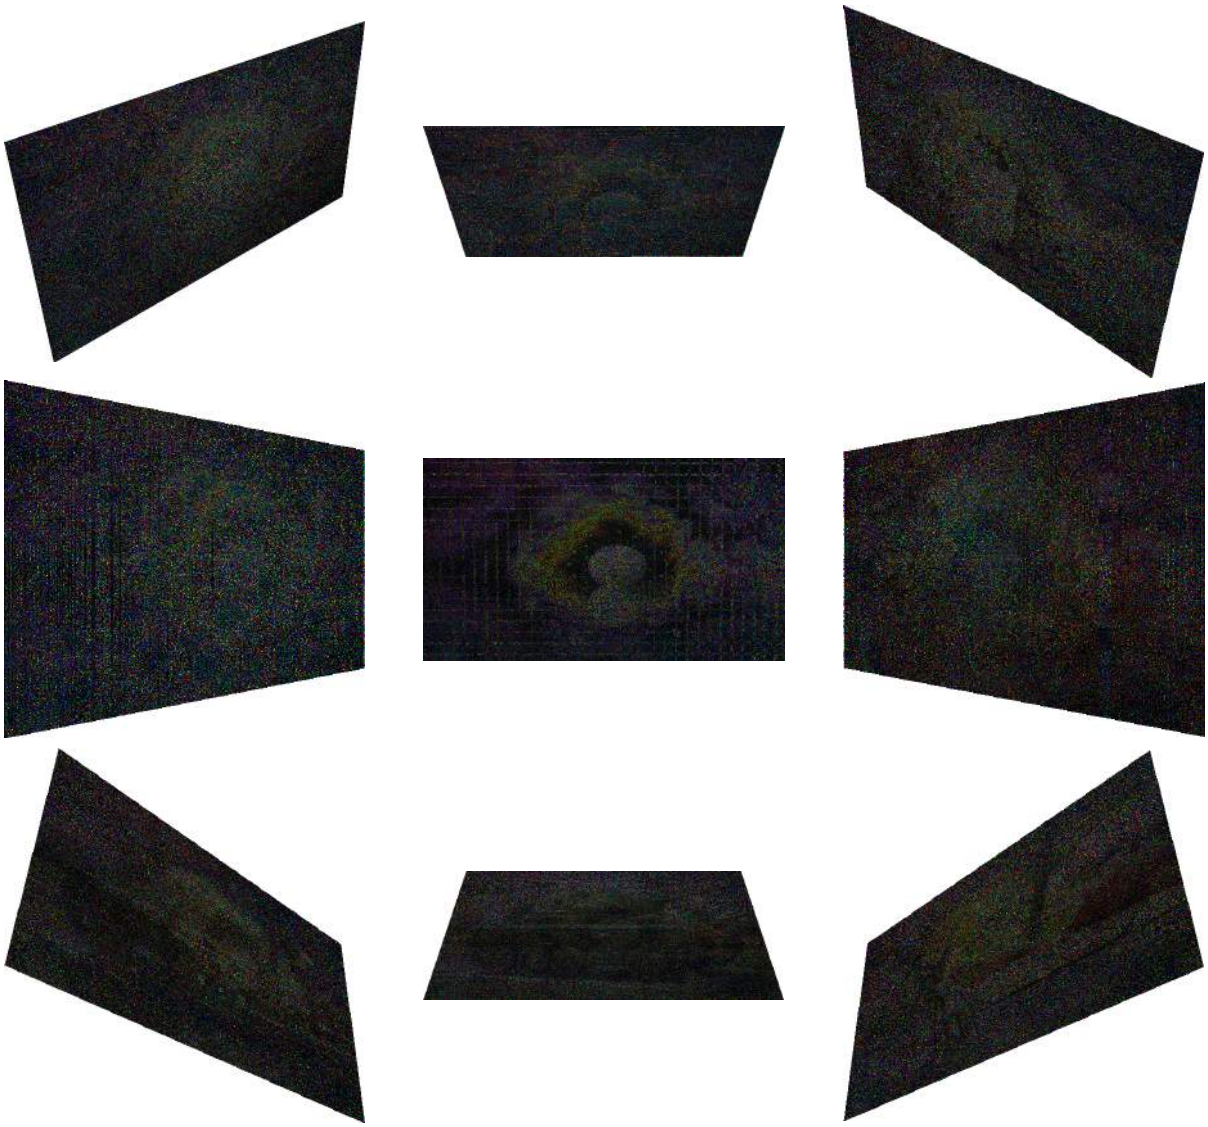

Figure 5.34: Simulated results using a perspective projection. The input is from Figure 5.28.

5.4 Autostereoscopic Design

The values shown in Table 5.1 and Table 5.2 are based on an ACER GD235HZ monitor, which can display a resolution of 1920 by 1080 pixels and has dimensions of 508mm by 287mm. This results in a horizontal pixel pitch of 0.265mm, which, with a plane spacing of 8mm, gives a minimum IPD of 253.2 mm and a minimum required viewing distance of 7.68m. Since the average adult IPD spans from 50mm to 75mm these are unacceptable

n	$\vec{\boldsymbol{p}}_{LB}(n)$	$\vec{\boldsymbol{p}}_{LF}(n)$	$\vec{\boldsymbol{p}}_{RB}(n)$	$\vec{\boldsymbol{p}}_{RF}(n)$	$\vec{\boldsymbol{a}}_{L}(\vec{\boldsymbol{o}}, n)$	$\vec{\boldsymbol{a}}_{L}(n)$	$\vec{\boldsymbol{a}}_{R}(\vec{\boldsymbol{o}}, n)$	$\vec{\boldsymbol{a}}_{R}(n)$
3.0	-254.0	-252.9	253.2	253.7	1.058	-254.0	0.5292	253.2
2.0	-254.0	-253.2	253.5	253.7	0.7937	-254.0	0.2646	253.5
1.0	-254.0	-253.5	253.7	253.7	0.5292	-254.0	0.0	253.7
0.0	-254.0	-253.7	254.0	253.7	0.2646	-254.0	-0.2646	254.0
-1.0	-253.7	-253.7	254.0	253.5	0.0	-253.7	-0.5292	254.0
-2.0	-253.5	-253.7	254.0	253.2	-0.2646	-253.5	-0.7938	254.0
-3.0	-253.2	-253.7	254.0	252.9	-0.5292	-253.2	-1.058	254.0

Table 5.1: Calculated Values (mm) Using ACER GD235HZ Parameters with d_p of 8mm

n	$t(n)$	$\vec{\boldsymbol{m}}(n, x)$	$\vec{\boldsymbol{m}}(n, z)$	MVD	$\vec{\boldsymbol{m}}$	$\vec{\boldsymbol{m}}(n+1)$	IPD
3.0	958.5	760.4	7668.0	7705.0			
2.0	959.0	507.2	7672.0	7688.0	2	3	253.2
1.0	959.5	253.7	7676.0	7680.0	1	2	253.5
0.0	960.0	0.0	7680.0	7680.0	0	1	253.8
-1.0	959.5	-253.7	7676.0	7680.0	-1	0	253.8
-2.0	959.0	-507.2	7672.0	7688.0	-2	-1	253.5
-3.0	958.5	-760.4	7668.0	7705.0	-3	-2	253.2

Table 5.2: Calculated Values (mm) Using ACER GD235HZ Parameters with d_p of 8mm (Continued)

values [19]. Furthermore, with this pixel pitch, the required distance for each viewing zone increases rapidly as we move away from the center view. So lets determine what pixel pitch would allow for acceptable IPD and viewing distance.

Table 5.3 displays some values obtained by varying the plane spacing, but maintaining the pixel pitch from the example LCD panel.

Tables 5.4, 5.5 and 5.6 list ranges of results all of which produce acceptable IPD values. Pixel pitch is the primary parameter contributing to IPD. Plane spacing is extremely important for setting up a proper and consistent viewing distance. The field of view of the display is also affected by plane spacing, but the number of input images is also an important factor. The monitor from our examples clearly does not have a small enough pixel pitch. Even the smallest pitch of the Apple retina display is 0.0877mm, which is

IPD (mm)	d_p (mm)	FOV	StdDev	dAvg
252.0	0.26458333333333334	168.6	747.0	1370.0
252.0	4.0	66.93	249.6	4117.0
252.0	8.0	36.58	124.4	7810.0
252.5	20.0	15.06	26.75	19209.0
253.5	50.0	6.05	56.48	47893.0

Table 5.3: Design Values from ACER GD235HZ monitor

Pixel Pitch (mm)	Plane Spacing (mm)	Field of View (degrees)	Standard Deviation of Minimum Viewing Distances (mm)	Average Minimum Viewing Distance (mm)
0.0525	0.0525	168.5	148.0	271.0
0.0525	4.0	14.9	5.2	3841.0
0.0525	8.0	7.5	7.3	7665.0
0.0515	20.0	2.9	29.0	19150.0
0.0455	50.0	1.04	77.2	47869.0

Table 5.4: Design Values from 50mm IPD and 1920x1080 Resolution

Pixel Pitch (mm)	Plane Spacing (mm)	Field of View (degrees)	Standard Deviation of Minimum Viewing Distances (mm)	Average Minimum Viewing Distance (mm)
0.0655	0.0655	168.6	185.0	339.0
0.0655	4.0	18.6	11.1	3848.0
0.0655	8.0	9.4	4.7	7668.0
0.06475	20.0	3.7	28.0	19151.0
0.06	50.0	1.37	76.7	47870.0

Table 5.5: Design Values from 62.5mm IPD and 1920x1080 Resolution

Pixel Pitch (mm)	Plane Spacing (mm)	Field of View (degrees)	Standard Deviation of Minimum Viewing Distances (mm)	Average Minimum Viewing Distance (mm)
0.0785	0.0785	168.6	221.7	406.0
0.0785	4.0	22.2	18.43	3856.0
0.0785	8.0	11.2	3.2	7672.0
0.078	20.0	4.5	26.5	19153.0
0.0775	50.0	1.8	76.0	47871.0

Table 5.6: Design Values from 75mm IPD and 1920x1080 Resolution

too large for the highest expected IPD [38]. LCD panels do exist with a small enough pixel pitch, but are not very affordable currently, which means some advances in the hardware are still necessary to allow for a commercial display using content adaptive parallax barriers.

Chapter 6

Discussion and Future Work

A detailed explanation and analysis of content adaptive parallax barriers has been presented. The steps necessary to go from a set of multi-view images to time-multiplexed masks are outlined as well as how each of those steps are performed. The necessary design of the physical device is also examined which assists in identifying some of the current weaknesses of the approach as well as some of the parameters that must be kept in mind when creating the optimal setup. We are hopeful that stating how content adaptive parallax barriers work, in a straight forward and rigorous manner, will help promote further related research. In addition, the analysis of the current capabilities of this technology reveal some areas where its limits can be pushed, some of which are detailed in Section 6.1.

6.1 Future Work

6.1.1 Head Tracking

The current proposed setup for content adaptive parallax barriers results in the need for a large number of images to allow for an acceptable field of view of the device. A high number of images results in the need for a higher rank. Increasing the rank is the least

recommended parameter change due to its effects on the number of computations, the necessary memory and the limits of the hardware. Doubling the rank roughly quadruples the number of computations and needed memory. Current LCD panels do not have a high enough refresh rate to adequately support the rank necessary for a large number of images. Head tracking has been used effectively in the past for autostereoscopic displays [61, 67]. There is potential to isolate where viewing zones are necessary by tracking the location of viewers and generating custom masks only in the needed areas. The result is a sever reduction in the necessary rank which means that this technology could support multiple viewers without significant increases to the required hardware.

6.1.2 Future Human Perception and Flicker Fusion

A higher rank factorization results in a better approximation of the input light field matrix. If the refresh rate of the LCD panels is not high enough there are two consequences. Firstly, what is displayed is extremely dark as the light of the scene is divided among too many masks, which do not contribute to a single view when displayed slowly. Second the human visual system cannot integrate the images if they are not viewed within a short enough time frame, which means a lower quality reasult is percieved. Theoretically with an infinite refresh rate there is the potential of having any number of masks, which would result in the highest quality approximations and produced images. That said, there have not been many tests focusing on what the human eye can possibly percieve. There has been research into what we can conciously see and remember, but that is different from what can be naturally percieved [15, 33, 36]. This opens up the possibility for research into how fast the human eye can percieve, which would allow for, among other things, the maximum necessary refresh rate under the content adaptive parallax barrier system.

6.1.3 Alternative Sampling Methods

The masks current produced, with the described procedure, do not take into account the fact that a light field is given as input. Instead it treats the input as only a set of multi-view images. This means there is room for an alternative sampling method which takes into account the redundancy between similar views and significantly reduces the number of computations and memory necessary to produce the masks.

6.1.4 Other Factorization Techniques

There are many matrix factorization techniques [65]. The different types can produce significantly different masks, but all techniques can accomplish the task. Trying newer NMF and weighted factorization methods has the potential for higher quality approximations with reduced compute time [44]. The greatest benefit that can come from new techniques or better implementations are a reduction in memory and compute time.

6.1.5 Additional Parameters

Our simulation used a fairly simple model where overlapping pixel values were multiplied to give the final result. There are many other parameters that could be used to help assess the performance of the display prior to construction. One example is the loss of light during attenuation through display planes. In a constructed display some amount of light is lost as it passes through each LCD panel and adding such behaviourwould allow for a more realistic result during simulation.

Another addition could be different layer types, such as polarizers. In simulation we decide the behaviour of light so, for our purposes, it behaves in an ideal manner, where there is no light loss and it all travel in the directon we choose. A more realistic behaviour is that everytime light interacts with a surface some of it is lost on the surface itself, some of it scatters in all directions and some of it moves in the direction we want

it to go. Polarizers are necessary in the physical device and their contribution to the loss of light could easily be added in simulation. If the simulation is upgraded to a physically based rendering example we could go even further in simulating actual behaviour prior to construction.

We could also look at changing parameters for some of the existing layers. For example, currently we assume that all the layers are the same LCD panel,meaning the same resolution, size and refresh rate. This assumption greatly affects our design analysis. If we allow for each panel to have varying parameters this would assist in determining if there is potential for better performance in a non-uniform panel display. It might be the case that a lower resolution front panel provides the desired behaviour, but gives more flexibility on viewing distance and mask computation.

6.1.6 Validation

We described some of the steps taken in validating that the NMF algorithms were working as we intended them to, but it would be very useful if that procedure was improved. If there was a standard dataset that any researcher in this area could use to validate that their algorithm is working that would make it easier to move on to later stages of research with complete confidence that the underlying algorithm does work for this application. This dataset would be made up of a number of standard inputs with expected outputs. The inputs would probably work best as a set of simple images such as fully black, fully white, various vertical and horizontal striped patterns with variable stripe sizes.

A generalized dataset might also allow for characterizing the NMF method being used. As seen in our presented results, different NMF approaches produce widely different masks. The products from a generalized set might characterize how a particular type of factorization technique will perform or it may help in determining what situations the technique is better or worse at handling.

Other than a generalized dataset, our work has revealed there is a need for a better

method of results comparison. PSNR does not adequately predict visual quality of the displayed masks. It might be that an adequate metric already exists, but very listtle research has been done to determine what metrics will relaibly predict visual quality in addition to reconstruction accuracy. It is potentially the case that a single metric might not be able to do the trick and a method must be used to predict and compare results, such as some sort of computer vision analysis. In any case a proper approach in comparing results and predicting their quality would be a valuable addition to this area.

Appendix A

Matrix Factorization

Matrix factorization decomposes one matrix into a product of matrices. Matrix factorization is widely used in problems ranging from face analysis to data compression, in signal processing to producing content adaptive parallax barriers, as seen in Chapter 3 [47]. There are many forms of matrix factorization, all of which decompose a matrix into a product of matrices [29, 47, 64]. The general form of matrix factorization is

$$\mathbf{M} = \prod_i \mathbf{M}_i,$$

where $\mathbf{M}$ is the input matrix and each of $\mathbf{M}_i$ multiply to produce an approximation to $\mathbf{M}$.

One specific form of matrix factorization is to decompose the input matrix into two approximating matrices.

$$\mathbf{L} = \mathbf{FG} + \mathbf{E}$$

where,

$\mathbf{L} \in \mathbb{R}^{m \times n}$ is the input, or observation matrix,

$\mathbf{F} \in \mathbb{R}^{m \times r}$ is the basis matrix,

$\mathbf{G} \in \mathbb{R}^{r \times n}$ is the weighting matrix,

$\mathbf{E} \in \mathbb{R}^{m \times n}$ is the error matrix,

and r is known as the rank.

This form of matrix factorization is the one used to decompose a light field into content adaptive parallax barriers which, when displayed using a multi-layered LCD device, gives the illusion of 3D. Even more specifically, content adaptive parallax barriers are produced using this form in the context of non-negative matrix factorization (NMF). We therefore restrict our discussion to NMF techniques of this form. We refer the kind reader to [65] for a more detailed review of matrix factorization techniques.

A.0.1 Non-Negative Matrix Factorization

Non-negative matrix factorization (NMF) is a subset of matrix factorization, which restricts the elements of the input matrix and approximating matrices to the set of positive real numbers.

$$\mathbf{L}(m,n), \mathbf{F}(m,r), \mathbf{G}(r,n) \geq 0,$$

where m is the number of rows of $\mathbf{F}$ and $\mathbf{L}$, n is the number of columns of $\mathbf{G}$ and $\mathbf{L}$, and r is both the number of columns of $\mathbf{F}$ and rows of $\mathbf{G}$. A strictly positive representation can be advantageous for a number of applications [65]. The principle benefit of a non-negative representation, for content adaptive parallax barriers, is that images are stored using positive RGB values.

The restrictions introduced with NMF are only one possible constraint that can be placed on the matrix factorization procedure. Another family of matrix factorization restrictions is known as box constraints. NMF is only bounded on one end, $\mathbf{F}(i,j) \geq 0$, box contraints are bounded on both ends, $\mathrm{p} \leq \mathbf{F}(i,j) \leq \mathrm{q}$. In fact, content adaptive parallax barriers require box constraints, since RGB values have a minimum and maximum value. NMF techniques are a convenient starting point because they inherently have a lower bound due to the non-negativity constraint and can be modified to include

an upper bound.

A.0.2 Lee & Seung

One of the first non-negative matrix factorization approaches was the Lee and Seung approach [47]. It uses an iterative multiplicative update rule, which minimizes the following objective function:

$$\underset{\mathbf{F}(i,j),\mathbf{G}(i,j)\geq 0}{\text{minimize}} \quad \frac{1}{2}||\mathbf{L}-\mathbf{FG}||^2.$$

The Lee & Seung update rules are:

$$\mathbf{F} \leftarrow \mathbf{F} \circ \left[\frac{\mathbf{LG}^T}{\mathbf{FGG}^T}\right] \quad \text{and} \quad \mathbf{G} \leftarrow \mathbf{G} \circ \left[\frac{\mathbf{F}^T\mathbf{L}}{\mathbf{F}^T\mathbf{FG}}\right],$$

where $\circ$ denotes the Hadamard product and the division is element-wise [48]. In the event that box constraints are necessary, an easy way to enforce the constraints is to truncate every element of $\mathbf{F}$ and $\mathbf{G}$, on each iteration, to be within the desired bounds [45].

A.0.3 FNMA$_e$

Fast Non-Negative Matrix Approximation (FNMA$_e$) is a quasi-Newton non-negative matrix approximation technique [42]. Quasi-Newton optimization is any technique that directly replaces the first derivative of Newton optimization with some other approach, quantity, or approximate derivative. Newton optimization uses the first derivative to determine the optimum value of some function, which in the context of matrices is known as the Jacobian. FNMA$_e$ replaces the Jacobian with the Hessian, which is the matrix form of the second derivative. In general, all the steps involved in FNMA$_e$ must be repeated for each of $\mathbf{F}$ and $\mathbf{G}$. So where the steps are identical, we will only describe them for $\mathbf{F}$.

The objective function for non-negativity is the same as the Lee and Seung method

from [47]:

$$\underset{\mathbf{F}(i,j),\mathbf{G}(i,j)\geq 0}{\text{minimize}} \quad \frac{1}{2}||\mathbf{L}-\mathbf{FG}||^2.$$

There are two gradient matrices, $\nabla\mathbf{F}$and $\nabla\mathbf{G}$, which are defined as:

$$\nabla\mathbf{F} = \mathbf{FGG}^T - \mathbf{LG}^T,$$

$$\nabla\mathbf{G} = \mathbf{F}^T\mathbf{FG} - \mathbf{F}^T\mathbf{L}.$$

The core concept of FNMA_e is that a set of free variables and fixed variables are identified during each iteration [42]. The fixed set, $\mathcal{I}$, is used to determine what elements meet the active constraints and have a positive derivative. For simple non-negativity the fixed set is calculated as follows:

$$\mathcal{I}_+ = \{(i,j)|\mathbf{F}(i,j) = 0, \nabla\mathbf{F}(i,j) > 0\}.$$

If box constraints are necessary the fixed set, given $p < \mathbf{F}(i,j) < q$ is:

$$\mathcal{I}_\Omega = \{(i,j)|\mathbf{F}(i,j) = p, \nabla\mathbf{F}(i,j) > 0) \quad or \quad (\mathbf{F}(i,j) = q, \nabla\mathbf{F}(i,j) < 0)\}.$$

The zero-out operator, $\mathcal{Z}[\cdot]$, on $\mathbf{F}$ is:

$$\mathcal{Z}[\nabla\mathbf{F}(i,j)] = \begin{cases} \mathbf{F}(i,j), & (i,j) \notin \mathcal{I}_+ \\ 0, & otherwise. \end{cases}$$

The projection, $\mathcal{P}[\cdot]$, is also dependant on the application. For the non-negativity constraint this is simply a projection onto the non-negative orthant, of $\mathbf{F}$:

$$\mathcal{P}_+[\mathbf{F}(i,j)] = \begin{cases} \mathbf{F}(i,j), & \mathbf{F}(i,j) \geq 0; \\ 0, & otherwise. \end{cases}$$

For box constraints the projection becomes:

$$\mathcal{P}_\Omega[\mathbf{F}(i,j)] = \begin{cases} p, & \mathbf{F}(i,j) \leq p; \\ \mathbf{F}(i,j), & p < \mathbf{F}(i,j) < q; \\ q, & q \leq \mathbf{F}(i,j). \end{cases}$$

One of the benefits of FNMA$_e$ is that the matrix $\mathbf{D}$ is used to approximate the Hessian inverse, rather than calculating the inverse itself, which provides a significant performance boost. The Hessian inverse is used to find the proper search direction when determining new values for our approximating matrices, $\mathbf{F}$ and $\mathbf{G}$. $\mathbf{D}$ is initialized as the identity matrix and is updated, if necessary, with the Broyden-Fletcher-Goldfarb-Shanno (BFGS) update rule

$$\mathbf{D}^{k+1} = \mathbf{D}^k - \frac{\mathbf{D}^k\mathbf{A}^T\mathbf{A}\mathbf{X}\mathbf{X}^T + \mathbf{X}\mathbf{X}^T\mathbf{A}^T\mathbf{A}\mathbf{D}^k}{\mathbf{X}^T\mathbf{A}^T\mathbf{A}\mathbf{X}} + \left(1 + \frac{\mathbf{X}^T\mathbf{A}^T\mathbf{A}\mathbf{D}^k\mathbf{A}^T\mathbf{A}\mathbf{X}}{\mathbf{X}^T\mathbf{A}^T\mathbf{A}\mathbf{X}}\right)\frac{\mathbf{X}\mathbf{X}^T}{\mathbf{X}^T\mathbf{A}^T\mathbf{A}\mathbf{X}},$$

where $\mathbf{X}$ is either $\mathbf{F}^T$ or $\mathbf{G}$ depending on which matrix is being updated [13, 24, 28, 63]. It is important to note that $\mathbf{F}$ and $\mathbf{G}$ are updated seperately, but we can use the same BFGS. Specifically, we can update $\mathbf{G}$ as normal, but the new $\mathbf{F}$ is obtained by updating $\mathbf{F}^T$ then and taking the transpose of the result. Referring to the previous equation, if $\mathbf{F}$ is being updated then $\mathbf{X}$ is $\mathbf{F}^T$ and $\mathbf{A}$ is $\mathbf{G}^T$. If $\mathbf{G}$ is being updated the roles switch, so $\mathbf{X}$ is $\mathbf{G}$ and $\mathbf{A}$ is $\mathbf{F}$. It is important to note that updating $\mathbf{D}$ is not always necessary depending on the current iteration and the fitness of the currently updating matrix. The update rule for $\mathbf{D}$, in the case that $\mathbf{F}$ is currently updating, is:

$$\mathbf{D}^{k+1} = \mathbf{D}^k - \frac{\mathbf{D}^k\mathbf{G}\mathbf{G}^T\mathbf{F}^T\mathbf{F} + \mathbf{F}^T\mathbf{F}\mathbf{G}\mathbf{G}^T\mathbf{D}^k}{\mathbf{F}\mathbf{G}\mathbf{G}^T\mathbf{F}^T} + \left(1 + \frac{\mathbf{F}\mathbf{G}\mathbf{G}^T\mathbf{D}^k\mathbf{G}\mathbf{G}^T\mathbf{F}^T}{\mathbf{F}\mathbf{G}\mathbf{G}^T\mathbf{F}^T}\right)\frac{\mathbf{F}^T\mathbf{F}}{\mathbf{F}\mathbf{G}\mathbf{G}^T\mathbf{F}^T}.$$

In the case that $\mathbf{G}$ is currently updating the update rule for $\mathbf{D}$ is:

$$\mathbf{D}^{k+1} = \mathbf{D}^k - \frac{\mathbf{D}^k\mathbf{F}^T\mathbf{F}\mathbf{G}\mathbf{G}^T + \mathbf{G}\mathbf{G}^T\mathbf{F}^T\mathbf{F}\mathbf{D}^k}{\mathbf{G}^T\mathbf{F}^T\mathbf{F}\mathbf{G}} + \left(1 + \frac{\mathbf{G}^T\mathbf{F}^T\mathbf{F}\mathbf{D}^k\mathbf{F}^T\mathbf{F}\mathbf{G}}{\mathbf{G}^T\mathbf{F}^T\mathbf{F}\mathbf{G}}\right)\frac{\mathbf{G}\mathbf{G}^T}{\mathbf{G}^T\mathbf{F}^T\mathbf{F}\mathbf{G}}.$$

As we mentioned earlier, the Hessian is used to determine the search direction, but a line search must be used to determine an appropriate length to advance in that direction. Any standard line search method is appropriate, but the one we use for FNMA$_e$ is a simplified version of the limited minimization rule used for constrained optimization [41, 42]. The search length, α, is found with

$$\alpha = \frac{-\mathbf{D}^T\overline{\mathbf{A}}^T(\overline{\mathbf{A}}\mathbf{X} - \mathbf{L})}{\mathbf{D}^T\overline{\mathbf{A}}^T\overline{\mathbf{A}}\mathbf{D}},$$

where $\overline{\mathbf{A}}$ is the appropriate values of $\mathbf{A}$ corresponding to the set of free variables and, as mentioned before, $\mathbf{X}$ is $\mathbf{F}^T$ or $\mathbf{G}$ and $\mathbf{A}$ is $\mathbf{G}^T$ or $\mathbf{F}$ while $\mathbf{F}$ or $\mathbf{G}$ are being updated, respectively. During an update of $\mathbf{F}$ the equation for α is

$$\alpha = \frac{-\mathbf{D}^T\overline{\mathbf{G}}(\mathbf{W}^T \circ (\overline{\mathbf{G}^T}\mathbf{F}^T - \mathbf{L}^T))}{\mathbf{D}^T\overline{\mathbf{G}}\overline{\mathbf{G}^T}\mathbf{D}},$$

where $\mathbf{L}^T$ is used as we are updating $\mathbf{F}^T$ and taking the transpose of the result for the new $\mathbf{F}$. During an update of $\mathbf{G}$ the equation for α is

$$\alpha = \frac{-\mathbf{D}^T\overline{\mathbf{F}}^T(\overline{\mathbf{F}}\mathbf{G} - \mathbf{L})}{\mathbf{D}^T\overline{\mathbf{F}}^T\overline{\mathbf{F}}\mathbf{D}}.$$

The entire FNMA$_e$ algorithm is outlined in Algorithm 1.

A.0.4 Weighted Matrix Factorization

A subset of matrix factorization uses a weighting matrix to focus on a certain set of values within the input [58]. One example of where weighted factorization could be used is in recommender systems [66], which are systems that reccomend products or services to you based on your previous behaviour. One example of a recommender system is in the streaming service Netflix, which tracks what you watch and suggests other content that might be appropriate. Content adaptive parallax barriers require weighted

Algorithm 1 FNMA_{e}

Input: $\mathbf{L} \in \mathbb{R}_+^{m \times n}, K$ s.t. $1 \leq K \leq min\{m, n\}$
Output: $\mathbf{F} \in \mathbb{R}_+^{m \times r}, \mathbf{G} \in \mathbb{R}_+^{r \times n}$
1. Initialize $\mathbf{F}_0, \mathbf{G}_0, t = 0, \mathbf{D} = \mathbf{I}$.
repeat
    2. $\mathbf{F} \leftarrow \mathbf{F}_t, \mathbf{G}_{old} \leftarrow \mathbf{G}_t$
    repeat
        3.1. Compute the gradient matrix $\nabla \mathbf{F}$.
        3.2. Compute the fixed set $\mathcal{I}_+$ for $\mathbf{G}_{old}$.
        3.3. Compute the step length vector α using line-search.
        3.4. Update $\mathbf{G}_{old}$ as
        $U \leftarrow \mathcal{Z}_+[\nabla \mathbf{F}(i, j)]$
        $U \leftarrow \mathcal{Z}_+[\mathbf{DU}]$
        $\mathbf{G}_{new} \leftarrow \mathcal{P}_+[\mathbf{G}_{old} - \mathbf{U} \cdot diag(\boldsymbol{\alpha})]$
        3.5. $\mathbf{G}_{old} \leftarrow \mathbf{G}_{new}$.
        3.6.Update $\mathbf{D}$ if necessary.
    until $\mathbf{G}_{old}$ converges
    4. $\mathbf{G}_{t+1} \leftarrow \mathbf{G}_{old}$.
    5. $\mathbf{G} \leftarrow \mathbf{G}_{t+1}$, $\mathbf{F}_{old} \leftarrow \mathbf{F}_t$.
    repeat
        6.1. Compute the gradient matrix $\nabla \mathbf{G}$.
        6.2. Compute the fixed set $\mathcal{I}_+$ for $\mathbf{F}_{old}$.
        6.3. Compute the step length vector α using line-search.
        6.4. Update $\mathbf{F}_{old}$ as
        $U \leftarrow \mathcal{Z}_+[\nabla \mathbf{G}(j, i)]$
        $U \leftarrow \mathcal{Z}_+[\mathbf{UD}]$
        $\mathbf{F}_{new} \leftarrow \mathcal{P}_+[\mathbf{F}_{old}^T - diag(\boldsymbol{\alpha}) \cdot \mathbf{U}]$
        6.5. $\mathbf{F}_{old} \leftarrow \mathbf{F}_{new}^T$.
        6.6.Update $\mathbf{D}$ if necessary.
    until $\mathbf{F}_{old}$ converges
    7. $\mathbf{G}_{t+1} \leftarrow \mathbf{G}_{old}$.
    8. $t \leftarrow t + 1$.
until Stopping criteria

matrix fatorization to function correctly. FNMA_e is adapted to be used as a weighted factorization technique in Section 3.4.1.

A weighted version of the Lee & Seung update rules, as presented in [8] and [17], are:

$$\mathbf{F} \leftarrow \mathbf{F} \circ \left[\frac{(\mathbf{W} \circ \mathbf{L})\mathbf{G}^T}{(\mathbf{W} \circ (\mathbf{F}\mathbf{G}))\mathbf{G}^T} \right] \quad \text{and} \quad \mathbf{G} \leftarrow \mathbf{G} \circ \left[\frac{\mathbf{F}^T(\mathbf{W} \circ \mathbf{L})}{\mathbf{F}^T(\mathbf{W} \circ (\mathbf{F}\mathbf{G}))} \right].$$

where $\mathbf{W}$ is a weighting matrix which enforces a set of constraints on the factorization process. $\mathbf{W}$ will differ according to the particular application, but the detailed procedure for calculating $\mathbf{W}$ for content adaptive parallax barriers is outlined in Section 3.3.

A.0.5 Technique Validation

To ensure that the matrix factorization techniques we chose were in fact behaving as intended, prior to looking at content adaptive parallax barrier masks, we used some simpler test cases. The first case was to set $\mathbf{L}$ to be equal to an image. This was used because it was very simple to modify our existing implementation to approach a single image, rather than an entire light field. In addition, once the factorization is complete we can simply multiply $\mathbf{F}$ and $\mathbf{G}$ and if the image looks like the input image then we have some verification that the technique is working.

Using an image as the input is useful we are reasonably sure that the technique is correct, but if the resultant approximation doesn't resemble the the image then we need to go further in diagnosing the problem. When this situation arose we still used an input image, but instead of a picture taken by a camera we used a very small image where we defined each pixel value manually. This allows for a direct comparison by hand to see what values are not being approximated or is there is some pattern in the results showing where the technique is not perfming as expected.

These methods of validation are rudimentary and could be greatly improved upon. They did, however, give us a reasonable confidence that our NMF implementations were

working so that we could move on to investigate other avenues of research.

A.0.6 Error Metrics

Error metrics are used to evaluate the accuracy of the approximation that **F** and **G** achieve in regards to the input **L**. The most popular metric, in this area, is the Peak Signal-to-Noise Ratio (PSNR), which uses the Mean Squared Error (MSE) to relate the best possible value to the approximation. This is easily transferrable to matrix approximation because the input is the best possible result and the approximation is a, hopefully slightly, warped version of the input. PSNR is measured in decibels (dB) so a larger value is better. PSNR can be stated as follows:

$$PSNR = 10 \log_{10} \frac{MAX^2}{MSE}$$

or

$$PSNR = 20 \log_{10} MAX - 10 \log_{10} MSE,$$

where MAX is the maximum possible value for that element, which for eight bit images is two hundred and fifty-five, and MSE is the mean squared error calculated as follows:

$$MSE = \frac{1}{mn} \sum_{mn} |\mathbf{L}_{(m,n)} - \mathbf{FG}_{(m,n)}|^2.$$

An important note is that the PSNR is not strictly a proper visual quality metric, as in it will not reliably give information as to how good the approximation appears visually to a human observer. However, PSNR has been shown to increase proportionally with visual quality metrics when examining similar datasets [37]. A typical PSNR for lossy compression is 30-50dB [3, 18]. PSNR is chosen for its ease of calculation and for comparison to related works.

Bibliography

[1] Build your own 3d display. http://alumni.media.mit.edu/~mhirsch/byo3d/index.html. Accessed: 2015-2-7.

[2] Edward H. Adelson and James R. Bergen. The plenoptic function and the elements of early vision. *Computational Models of Visual Processing*, pages 3–20, 1991.

[3] Luciano Alparone, Fabrizio Argenti, and Tiziano Bianchi. Multiresolution analysis for image compression. *Document and Image Compression*, 2006.

[4] John Baker. Generating images for a time-multiplexed stereoscopic computer graphics system. In *OE LASE'87 and EO Imaging Symp (January 1987, Los Angeles)*, pages 44–52. International Society for Optics and Photonics, 1987.

[5] Stephen A. Benton. *Selected papers on three-dimensional displays*, volume 162. SPIE-Society of Photo-optical Instrumentation Engineers, 2001.

[6] Philip Benzie, John Watson, Phil Surman, Ismo Rakkolainen, Klaus Hopf, Hakan Urey, Ventseslav Sainov, and Christoph Von Kopylow. A survey of 3dtv displays: techniques and technologies. *Circuits and Systems for Video Technology, IEEE Transactions on*, 17(11):1647–1658, 2007.

[7] Dimitri P. Bertsekas. Projected newton methods for optimization problems with simple constraints. *SIAM Journal on control and Optimization*, 20(2):221–246, 1982.

[8] Vincent D. Blondel, Ngoc-Diep Ho, and Paul Van Dooren. Weighted nonnegative matrix factorization and face feature extraction. *submitted to Image and Vision Computing*, 2007.

[9] Ulrich Bonne. Stereoscopic display system. *SPIE Milestone Series MS*, 162:177–185, 2001. US Patent 3,858,001 - 1974.

[10] Boost. Boost c++ libraries. Version 1.55, November 11 2013.

[11] David Brewster. An account of a new stereoscope. *SPIE Milestone Series MS*, 162:111–112, 2001. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science - 1844.

[12] David Brewster. On the law of visible position and binocular vision, and of the representation of solid figures by the union of dissimilar plane pictures on the retina. *SPIE Milestone Series MS*, 162:27–53, 2001. Report of the Nineteenth Meeting of the British Association for the Advancement of Science - 1850.

[13] Charles G. Broyden. The convergence of a class of double-rank minimization algorithms 2. the new algorithm. *IMA Journal of Applied Mathematics*, 6(3):222–231, 1970.

[14] Robert B. Collender. Three dimensional unaided viewing method and apparatus. *SPIE Milestone Series MS*, 162:268–283, 2001. US Patent 3,178,720 - 1965.

[15] Karen Davranche and Aurélien Pichon. Critical flicker frequency threshold increment after an exhausting exercise. *Journal of Sport and Exercise Psychology*, 27(4):515, 2005.

[16] Leonardo de Vinci. *A Treatise of Painting*. London : Printed for Senex, J. and Taylor, W., 1721.

[17] Inderjit. S. Dhillon and Suvrit Sra. Generalized nonnegative matrix approximations with bregman divergences. *Advances in neural information processing systems*, pages 283–290, 2005.

[18] Ahmet E. Dirik and Nasir Memon. Selective robust image encryption for social networks. *Multimedia Communications, Services and Security*, pages 70–81, 2013.

[19] Neil A. Dodgson. Variation and extrema of human interpupillary distance. In *Electronic Imaging 2004*, pages 36–46. International Society for Optics and Photonics, 2004.

[20] Neil A. Dodgson. On the number of viewing zones required for head-tracked autostereoscopic display. In *Electronic Imaging 2006*, pages 60550Q–60550Q. International Society for Optics and Photonics, 2006.

[21] Louis D. Duhauron. Estampes, photographies et tableaux stereoscopiques, produisant lair effet en plein jour, sans l'aide du stereoscope. *SPIE Milestone Series MS*, 162:138–145, 2001. French Patent 216,465 - 1891.

[22] Louis D. Duhauron. Stereoscopic print. *SPIE Milestone Series MS*, 162:146–148, 2001. US Patent 544,666 - 1895.

[23] Eigen. Eigen linear algebra library. Version 3.2.2, August 4 2014.

[24] Roger Fletcher. A new approach to variable metric algorithms. *The computer journal*, 13(3):317–322, 1970.

[25] FreeImage. Freeimage open source image library. Version 3.16.0, March 23 2014.

[26] GLFW. Open graphics library framework. Version 3.0.4, December 31 2013.

[27] GLM. Open graphics library mathematics. Version 0.9.7.0, February 8 2015. http://glm.g-truc.net/0.9.7/index.html.

[28] Donald Goldfarb. A family of variable-metric methods derived by variational means. *Mathematics of computation*, 24(109):23–26, 1970.

[29] Gene H. Golub. Least squares, singular values and matrix approximations. *Aplikace matematiky*, 13(1):44–51, 1968.

[30] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The lumigraph. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 43–54. ACM, 1996.

[31] Willhelm B. Gruber. Stereoscopic viewing device. *SPIE Milestone Series MS*, 162:116–118, 2001. US Patent 2,189,285 - 1940.

[32] Laurens Hammond. Stereoscopic motion-picture device. *SPIE Milestone Series MS*, 162:171–176, 2001. US Patent 1,506,524 - 1924.

[33] Selig Hecht and Emil L. Smith. Intermittent stimulation by light vi. area and the relation between critical frequency and intensity. *The Journal of general physiology*, 19(6):979–989, 1936.

[34] Matthew Hirsch and Douglas Lanman. Public domain images. http://www.public-domain-image.com. Accessed: 2015-9-16.

[35] Oliver W. Holmes. History of the "american stereoscope". *SPIE Milestone Series MS*, 162:113–115, 2001. Philedelphia Photographer - 1869.

[36] Larry E. Humes, Thomas A. Busey, James C. Craig, and Diane Kewley-Port. The effects of age on sensory thresholds and temporal gap detection in hearing, vision, and touch. *Attention, Perception, & Psychophysics*, 71(4):860–871, 2009.

[37] Q. Huynh-Thu and M. Ghanbari. Scope of validity of psnr in image/video quality assessment. *Electronics letters*, 40(13):800–801, 2008.

[38] Apple Inc. Compare ipad models. http://www.apple.com/ca/ipad/compare/. Accessed: 2015-10-2.

[39] Frederic E. Ives. Parallax stereogram and process of making same. *SPIE Milestone Series MS*, 162:146–148, 2001.

[40] R Clark Jones. A new calculus for the treatment of optical systems. *JOSA*, 31(7):488–493, 1941.

[41] Dongmin Kim, Suvrit Sra, and Inderjit S. Dhillon. *A new projected quasi-newton approach for the nonnegative least squares problem.* Citeseer, 2006.

[42] Dongmin Kim, Suvrit Sra, and Inderjit S. Dhillon. Fast newton-type methods for the least squares nonnegative matrix approximation problem. *SDM*, 7:343–354, 2007.

[43] Dongmin Kim, Suvrit Sra, and Inderjit S. Dhillon. Tackling box-constrained optimization via a new projected quasi-newton approach. *SIAM Journal on Scientific Computing*, 32(6):3548–3563, 2010.

[44] Yong-Deok Kim and Seungjin Choi. Weighted nonnegative matrix factorization. In *IEEE International Conference on Acoustics, Speech and Signal Processing, 2009. ICASSP 2009.*, pages 1541–1544. IEEE, 2009.

[45] Douglas Lanman, Matthew Hirsch, Yunhee Kim, and Ramesh Raskar. Content-adaptive parallax barriers: Optimizing dual-layer 3d displays using low-rank light field factorization. *ACM Transactions on Graphics*, 29(6):163, 2010.

[46] Douglas Lanman, Matthew Hirsch, Gordon Wetzstein, and Ramesh Raskar. Polarization fields: Dynamic light field display using multi-layer lcds. *ACM Transactions on Graphics*, 30(6):186, 2011.

[47] Daniel D. Lee and H. Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.

[48] Daniel D. Lee and H. Sebastian Seung. Algorithms for non-negative matrix factorization. *Advances in neural information processing systems*, pages 556–562, 2001.

[49] Emmett N. Leith and Juris Upatnieks. Reconstructed wavefronts and communication theory. *Journal of the Optical Society of America*, 1962.

[50] Emmett N. Leith and Juris Upatnieks. Wavefront reconstruction with diffused illumation and three-dimensional objects. *Journal of the Optical Society of America*, 1964.

[51] Marc Levoy and Pat Hanrahan. Light field rendering. *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 31–42, 1996.

[52] Lenny Lipton, Michael R. Starks, James D. Stewart, and Lawrence D. Meyer. Stereoscopic television system. *SPIE Milestone Series MS*, 162:151–170, 2001. US Patent 4,523,226 - 1985.

[53] Henry J. De N. McCollum. Stereoscopic television apparatus. *SPIE Milestone Series MS*, 162:119–121, 2001. US Patent 2,388,170 - 1945.

[54] Leonard McMillan and Gary Bishop. Plenoptic modeling: An image-based rendering system. *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, 1995.

[55] Edmond H. V. Noaillon. Art of making cinematographic projections. *SPIE Milestone Series MS*, 162:260–267, 2001. US Patent 1,772,782 - 1930.

[56] Persistence of Vision Raytracer Pty. Ltd. Persistence of vision raytracer. Version 3.7, November 19 2007. http://www.povray.org/.

[57] OpenGL. Open graphics library. Version 4.5, 2015. https://www.opengl.org/.

[58] Pentti Paatero, Unto Tapper, Pasi Aalto, and Markku Kulmala. Matrix factorization methods for analysing diffusion battery data. *Journal of Aerosol Science*, 22:S273–S276, 1991.

[59] Siegmund Pastoor and Matthias Wöpking. 3-d displays: A review of current technologies. *Displays*, 17(2):100–110, 1997.

[60] John A. Roese. Liquid crystal stereoscopic television system. *SPIE Milestone Series MS*, 162:186–194, 2001. US Patent 3,821,466 - 1974.

[61] Alfred Schwartz. Head tracking stereoscopic display. *SPIE Milestone Series MS*, 162:241–245, 2001. IEEE Tansactions on Electron Devices - 1986.

[62] J. A. Segner. De raritate luminis. pages 1–12, 1740. Vandenhoeck.

[63] David F. Shanno. Conditioning of quasi-newton methods for function minimization. *Mathematics of computation*, 24(111):647–656, 1970.

[64] Haipeng Shen and Jianhua Z. Huang. Sparse principal component analysis via regularized low rank matrix approximation. *Journal of multivariate analysis*, 99(6):1015–1034, 2008.

[65] Suvrit Sra and Inderjit S. Dhillon. *Nonnegative matrix approximation: Algorithms and applications*. Computer Science Department, University of Texas at Austin, 2006.

[66] Nathan Srebro and Tommi Jaakkola. Weighted low-rank approximations. In *ICML*, volume 3, pages 720–727, 2003.

[67] Phil Surman, Ian Sexton, Klaus Hopf, Wing K. Lee, Richard Bates, Wijnand IJsselsteijn, and Edward Buckley. Head tracked single and multi-user autostereoscopic displays. In *Visual Media Production, 2006. CVMP 2006. 3rd European Conference on*, pages 144–152. IET, 2006.

[68] Ivan E. Sutherland. A head-mounted three dimensional display. *SPIE Milestone Series MS*, 162:122–132, 2001. In Joint Computer Conference [Proceedings of the AFIPS] 1968.

[69] Wa J. Tam, Filippo Speranza, Sumio Yano, Koichi Shimono, and Hiroshi Ono. Stereoscopic 3d-tv: visual comfort. *Broadcasting, IEEE Transactions on*, 57(2):335–346, 2011.

[70] Alfred F. Watch. The anaglyph: a new method of producing the stereoscopic effect. *Journal of the Franklin Institute*, 140(6):401–419, 1895.

[71] Gordon Wetzstein, Douglas Lanman, Matthew Hirsch, and Ramesh Raskar. Tensor displays: Compressive light field synthesis using multilayer displays with directional backlighting. *ACM Transactions on Graphics*, 31(4):80, 2012.

[72] Charles Wheatstone. Contributions to the physiology of vision - part the first. on some remarkable and hitherto unobserved phenomena of binocular vision. *SPIE Milestone Series MS*, 162:3–26, 2001. Philosophical Transactions of the Royal Society of London 1838.

[73] Charles Wheatstone. Contributions to the physiology of vision - part the second. on some remarkable and hitherto unobserved phenomena of binocular vision (continued). *SPIE Milestone Series MS*, 162:54–70, 2001. Philosophical Transactions of the Royal Society of London 1852.

[74] Andrew J. Woods, Chris R. Harris, Dean B. Leggo, and Tegan M. Rourke. Characterizing and reducing crosstalk in printed anaglyph stereoscopic 3d images. *Optical Engineering*, 52(4):043203–043203, 2013.

[75] Hirotsugu Yamamoto, Yoshio Hayasaki, and Nobuo Nishida. Large stereoscopic led display by use of a parallax barrier. In *Three-dimensional Imaging, Visualization, and Display*, pages 183–205. Springer, 2009.