UNIVERSITY OF CALIFORNIA
SANTA CRUZ

ADVANCES IN DISPLAY TECHNOLOGY

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

COMPUTER SCIENCE

by

Steven Jonathan Scher

September 2012

The Dissertation of Steven Jonathan Scher
is approved:

Professor James Davis, Chair

Professor Roberto Manduchi

Professor Suresh Lodha

Tyrus Miller
Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures

Abstract

Advances in Display Technology

by

Steven Jonathan Scher

We live in a digital age. As more of our interactions are mediated by computers, their visual displays become our windows to the world. Every year, new and 'more-realistic' displays bring the world to our fingertips, and we recognize the quaint limits of last year's printers, televisions, and phones. This dissertation pushes display technology on four fronts, each immersing us more deeply in the scene we're shown, letting us forget the display itself.

First, we introduce printed photos that respond to ambient light as would the depicted 3D scene. Shadows and shading in a printed image change to match its new lighting environment. With a 3D object and a print side by side, a light to their right will casts shadows to the left on both. Shiny, dimpled 'Reflectance Paper' - with opaque ink on an overlaid transparency - selectively blocks light from each direction independently.

Second, we look to 3D displays, where stereo glasses funnel different images to each eye. Typical 3D displays show a confusing double-image to those not wearing stereo glasses. Our 3D+2DTV gives a 3D view to those with the glasses, but presents a single, 2D view to those without. For those without glasses, we

cancel one of the two stereo images with an inverted third image that only they see.

Third, we increase the fidelity of mobile displays' power-saving backlight dimming. With maximum brightness reduced, intermediate brightness levels are remapped to preferentially retain detail at strong edges in the image.

Fourth, we personalize photography. Displaying photos with content-aware image manipulations and eyetrack-based saliency subtly preserves photographers' intentions as recorded in their eye gaze history.

Display technology's relentless evolution enriches our increasingly digitized world. This dissertation's four contributions raise the immersiveness of our digital portals, heightening our view of the world and our communication with one another.

For my inspiring parents,

Barry and Linda Scher

Acknowledgments

This dissertation was made possible by the help and support of my family, friends, and colleagues. Thank you.

Thanks to my advisor, James Davis, for his encouragement and direction, and for sharing his relentless drive to continually improve the world. Thanks to my constant companions through code and experiments, Prabath Gunawardane and Oliver Wang, and my new colleagues Jing Liu and Rajan Vaish.

Special thanks also to Tom Malzbender and Ramin Samadani at HP Labs and to Sriram Swaminarayan at LANL for their friendship, expertise, and enthusiasm that kept me on-course over many years. Thanks to Roberto Manduchi and Suresh Lodha for their time and advice, as well as Yi Zhang, Manfred Warmuth, Arnav Jhala, Bruce Bridgeman, and all my other professors for always steering me in the right direction, to Tracie Tucker for righting the ship in any storm, and to Serge Belongie for showing me the sea.

Gratitude and love to my parents and family for their unwavering confidence, and the many friends who kept my spirit strong and curiosity piqued. Thank you, Rion, CJ, Corrie, Heather, for inspiring me to dream big and enjoy discovery.

Thanks to my colleagues, co-authors, and collaborators at UC Santa Cruz, HP Labs, MatterPort, LANL, Stanford, National Semiconductor, and Microsoft Research. Thanks also to the UC Regents, Intel, Los Alamos National Lab, and

Hewlett Packard for partially funding my research, and to the anonymous reviewers who have improved it.

Chapter 1

Introduction

The twentieth century has wrought enormous change on the world's inhabitants, and few technologies have wielded greater impact than the massively improved means of communication now available. The digital globe brings people together, connecting the world's citizens more closely than ever before. With computers mediating more and more of our interactions and defining our experiences of the world, the last mile of this connection - the visual display - grows ever more important.

Pens gave way to the printing press 500 years ago; paintings to photographs 200 years ago. Televisions appeared 100 years ago, and 50 years ago, digital printers. These improvements in the display of information accelerate understanding and communication, as well as the capacity for artistic conveyance of more deeply-felt

truths.

Each technological upheaval raises our expectation of the limit with which reality can be counterfeited on a canvas or a screen. We recalibrate our expectations to maintain an inner certainty that we understand the divide between the real, physical world that we perceive directly, and the imagined world so deftly painted for us. We guard this chasm, and yet, as the virtual world on screen becomes less easily discernible from the physical world, we find ourselves more deeply immersed in it. We step more easily between the two. New displays startle and astonish us at first with their lifelikeness; after novelty subsides and we accustom ourselves to the new medium, we begin to depend on its wider lanes of communication to more accurately record facts and capture dreams, and to better share them with others.

More than ever before, we live surrounded by the products of computers. Printed pages and pixelized screens fill our vision, bringing us knowledge and beauty. New tools imbue communication with new resources and new roles. This dissertation presents four new tools for the world's artists and consumers. We introduce a new printing method and an improved 3DTV that bring photographed or synthetic objects one shade closer to reality, as well as enhancing mobile displays and pushing a new direction in personalized photography.

This chapter reviews the impact of the introduction of several historical de-

velopments in display technology, notes a range of recent work in the field, and introduces our contributions.

1.1 See More of the World, See a New World

Many technological developments in display technology bring with them new modes of communication and expression. Human communication takes place in the context of expectations about the world; expansions of 'what is possible' force us to reexamine our model of the world. As one generation's display technology leapfrogs the last, the new technology comes to be regarded as a proxy for reality, while the old technology is redefined as a quaint genre.

New mechanisms of communication enabled by a new display are fertile grounds for exploration. When new toys appear, what new games can be played? Are they merely different, or truly better?

The camera obscura arrived in fifteenth-century Florence as a delightful toy, but quickly changed every aspect of the city's art: what is worth painting, how to depict it, how to paint, and even whose paintings are valued. The camera obscura projects the image of a 3D scene onto a 2D surface, as in Figure 1.1. Details of perspective, occlusion, shadows, and shading previously visible to only an expertly-trained eye were suddenly plain to all. Paintings incorporating accurate perspective effects were captivatingly realistic, in contrast to the previously popu-

(a) The Camera Obscura

(b) *Presentation of the Virgin in the Temple*, Paolo Uccello, 1434

Figure 1.1: The Camera Obscura made plain the effects of perspective. Italian Renaissance artists quickly took to perspective in their paintings.

lar style wherein subjects' relative sizes and spatial relations conferred importance and relationships, rather than describing physical scenes.

After a millennium of obscurity, works featuring perspective proliferated from countless artists. Since the camera obscura could project an image directly onto a canvas, subtle details could be directly traced by even moderately skilled artists; 21st-century artists use a projector for the same purpose. Audiences voted with their purchases, and the perspective style dominated. This was no passing fad of fickle taste. The lifelike perspective effects amazed us, yet after the initial novelty subsided, they still convey a sense of depth that draws the viewer in.

The advent of the telescope in the 17th century brought the furthest and smallest worlds within view. Tracing planets' movements allowed mathematics to weigh in on their governing physics, but meticulous orbital trajectories were

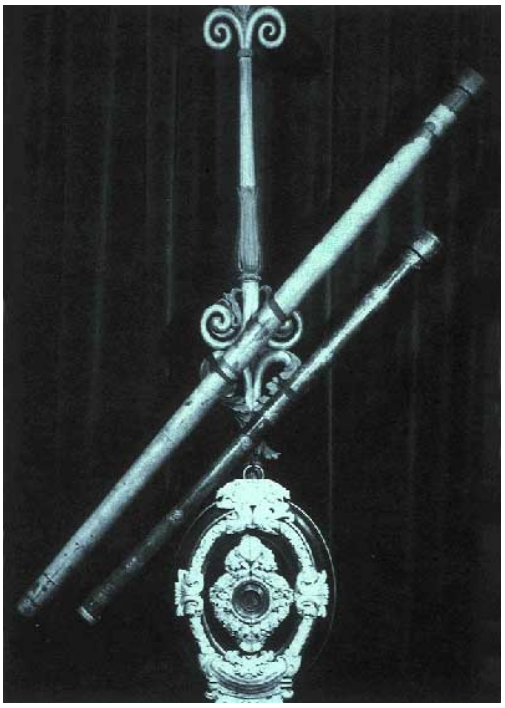

(a) Galileo's First Telescope

(b) Galileo's Sketches of the Earth's Moon

Figure 1.2: Early Telescopes allowed the first detailed views of the cosmos

not the only images made by early telescopes like that in Figure 1.2. Planets, the wandering stars, were for the first time more than dots of light. The moon's craters were noticed to be astonishingly similar to earthly rocks. Copernican heliocentrism rests not only on calculated orbits, but on the proposition that Earth is one of several specks of light that orbit the sun. Newton's gravitation posits the heavenly bodies to be made of the same stuff as the dirt under our feet. The telescopes amazing views of the solar system astonished us, assimilated into our lives, and left us richer for the sight.

The development of the photograph at the turn of the 19th century brought a

Figure 1.3: In *The Treachery of Images, 1928*, Rene Magritte paints a pipe, and below it writes (in French) "This is not a pipe." We know this is a painting; the recent growth of photography leaves viewers acutely aware that this painting is of a different class than a real pipe. And yet, the admonition *'this is not a pipe; this is a painting of a pipe'* is jarring, bringing those subconscious expectations to the forefront. Artists use many tools to let us believe in their imaginary worlds. In this dissertation, we supply a few new tools.

new sense of 'objective reality' en vogue. Moments in time were caught on paper; the past persisted. Photographs stretched out to touch lives across time and across continents. We didn't know how much we loved them until we had them. As cameras proliferate into every phone and social networks simplify photo distribution, photos displace stories as the shibboleth of truth. We now demand "photos, or it didn't happen" where someone's word would have been good enough. This is not limited to our friends' embarrassing exploits: juries have accustomed themselves to explicit evidence and have accordingly raised the expected evidentiary bar.

Photography also had a profound effect on painters. Relieved of their former place as authoritative documenters of the world, some embraced a wider view of their medium, precipitating Modernism and Surrealism [14]. Magritte's *Treachery*

of Images, in Figure 1.3 confronts the limits of the artifice of illustration. It entrances viewers by forcing us to reconcile our recognition of the depicted pipe with its obvious distinction: it is not a pipe. Magritte is reputed to have reproached doubters to stuff the painted pipe with tobacco.

With each step forward, new technologies let artists hypnotize their viewers, captivating them with unexpectedly 'real' images. Viewers find their incredulity forestalled a moment longer, and sink more deeply into the depicted world. Though time dulls our wonder at the marvels of illustration available to us, they do not lose their power to enthrall us with imagined lands.

1.2 Four New Displays

Display technology must be constantly reinvented and expanded. Existing technologies lose their miraculous wonder, and become recognized as proper objects (ink on paper), rather than fooling the viewer into believing the image printed is more real. I want to help suspend viewers' disbelief, if only for a moment more.

This dissertation expands the forefront of modern displays in four areas. We describe a 3DTV that remains readable, in 2D, for those unable to wear 3D glasses; a new genre of printers that increments the printable page from two dimensions to four; a mobile LCD display with reduced power consumption; and a camera whose photographs capture not only an image of the subject, but the intent of

the photographer.

1.2.1 3D+2DTV

(a) Standard 3DTV (b) Our 3D+2D TV (c) A 3rd Frame Cancels R

Figure 1.4: (a) A typical 3DTV shows a different image to each eye through the stereo glasses at the bottom of the image; without glasses both images are superimposed. (b) Our 3D+2DTV shows only one of the two images to those without glasses. (c) A negative 3rd image is visible only without glasses.

3D displays are increasingly popular in consumer and commercial applications. Many such displays show 3D images to viewers wearing special glasses, while showing an incomprehensible double-image to viewers without glasses, as in Figure 1.4a. We demonstrate a simple method that provides those with glasses a 3D experience, while viewers without glasses see a 2D image without artifacts, as in Figure 1.4b.

In addition to separate Left and Right images in each frame, we add a third image, invisible to those with glasses, as in Figure 1.4c. In the combined view seen by those without glasses, this cancels the Right image, leaving only the Left.

If the Left and Right images are of equal brightness, this approach results in

low contrast for viewers without glasses. Allowing differential brightness between the Left and Right images improves 2D contrast. We determine experimentally that viewers with glasses maintain a strong 3D experience, even when one eye is significantly darker than the other.

Since viewers with glasses see a darker image in one eye, they experience a small distortion of perceived depths due to the Pulfrich Effect. This produces illusions similar to those caused by a time delay in one eye. We experimentally find that a 40% brightness difference cancels an opposing distortion caused by the typical 8 millisecond delay between the Left and Right images of sequential active-shutter stereoscopic displays.

Our technique is applicable to displays using either active-shutter glasses or passive glasses. The key feature needed is a third channel of images seen only by those without glasses. We demonstrate a prototype system using a combination of active shutter projectors and a polarizing filter.

1.2.2 Printing Reflectance Functions

The reflectance function of a scene point captures the appearance of that point as a function of lighting direction. We present an approach to printing the reflectance functions of an object or scene so that its appearance is modified correctly as a function of the lighting conditions when viewing the print. For example, such

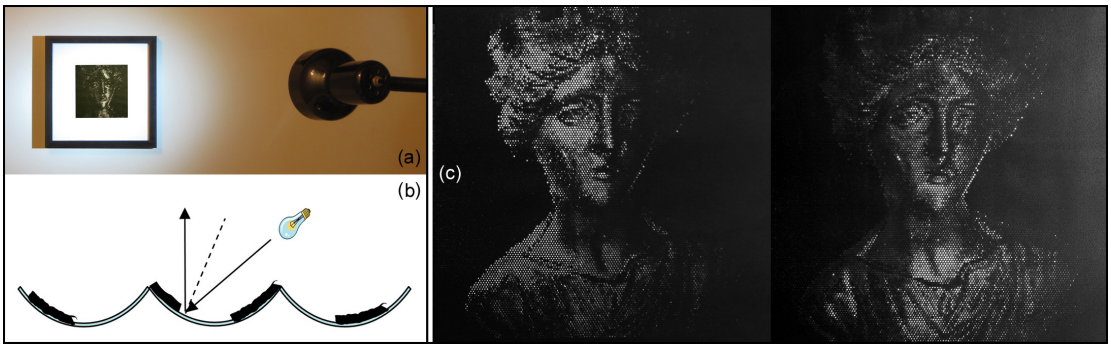

Figure 1.5: (a) We introduce a method for printing reflectance functions, images which correctly respond to the direction from which they are illuminated. (b) A microgeometry reflective substrate returns a specular highlight unless opaque ink is deposited on the point corresponding to a particular incident illumination direction. (c) Our prototype is shown displaying different images in response to two different illumination conditions. Image used with the permission of Cultural Heritage Imaging.

a photograph of a statue printed with our approach appears to cast shadows to the right when the photograph is illuminated from the left. Viewing the same print with lighting from the right will cause the statue's shadows to be cast to the left. Beyond shadows, all effects due to the lighting variation, such as Lambertian shading, specularity, and inter-reflection can be reproduced. We achieve this ability by geometrically and photometrically controlling specular highlights on the surface of the print. For a particular viewpoint, arbitrary reflectance functions can be built up at each pixel by controlling only the specular highlights and avoiding significant diffuse reflections. Our initial binary prototype uses halftoning to approximate continuous grayscale reflectance functions.

1.2.3 Higher-Contrast Low-Power Displays

Original Image | 50% Brightness | 50% Brightness with Contrast Preserved

Figure 1.6: We preserve contrast at salient edges when an image's brightness is reduced.

Mobile devices with LCD screens may save power by reducing the backlight brightness, however this reduces the available dynamic range. We choose an image-specific gamma-curve that preserves contrast at edges under reduced dynamic range.

We form a 2D histogram of the gray levels that form edges. This histogram is used to produce a gamma curve that preserves contrast at edges while allowing contrast to diminish between gray level pairs that are not spatially adjacent.

The algorithm is simple, elegant, fast, and requires little memory, making it suitable for implementation in hardware (e.g. FPGA) aboard mobile devices.

1.2.4 Personalized Photographs

Photography provides tangible and visceral mementos of important experiences. Recent research in content-aware image processing to automatically improve photos relies heavily on automatically identifying salient areas in images. While automatic saliency estimation has achieved estimable success, it will always face inherent challenges where saliency involves semantic judgments involving relationships between people or objects in the scene and the unseen photographer.

(a) Original image (b) Saliency from Eye Tracks (c) Content-Aware Resizing

Figure 1.7: Saliency maps derived from eyetracks allow automatic content-aware resizing to personalize photographs to reflect the photographer's view.

Tracking the photographer's eyes allows a direct, passive means to estimate scene saliency. We instrument several content-aware image processing algorithms with eye track based saliency estimation, producing personalized photos that accentuate the parts of the image important to one particular person.

1.3 Discussion

The evolution of visual media has allowed communication to accelerate dramatically, with new means of understanding accompanying each new technological press forward. We present here several new nudges to push the boundaries of display technology.

Our reflectance-function printer and 3D+2DTV help shepherd in a new genre of more flexible displays, whose additional artifices create more-vivid illusions. Our higher-contrast low-power displays and personalized photographs similarly transmit a closer approximation to an intended, rich image within a traditional display paradigm.

Chapter 2

3D+2DTV

2.1 Introduction

Stereoscopic displays provide different images to the viewer's right and left eyes to produce a three-dimensional (3D) percept. These displays' falling prices have caused them to grow from a niche product to mass market acceptance with applications in entertainment, medical imaging, and engineering visualization.

The most popular 3D display paradigm shows a pair of images on the same screen, intended for the viewers' left and right eyes. The lenses of special "stereo glasses" pass images to the correct eye. A viewer not wearing these glasses sees both images superimposed, creating a "ghosted" double-image where two copies of objects appear overlaid [Figure 2.1a].

(a) Standard 3DTV (b) Our 3D+2D TV (c) A 3rd Frame Cancels R

Figure 2.1: (a) A typical glasses-based 3DTV shows a different image to each eye of viewers wearing stereo glasses, visible through the glasses at the bottom of the figure, while those without glasses see both images superimposed, visible directly on the screen at the top of the figure. (b) Our 3D+2DTV likewise shows a different image to each eye of viewers wearing stereo glasses, but shows only one of these images to those without glasses, removing the "ghosted" double-image. (c) We accomplish this by displaying a 3rd image to those not wearing glasses that is not visible to those wearing glasses, canceling out one image of the stereo pair.

It is not always desirable to require that all viewers wear stereo glasses. They can be prohibitively expensive, or may interfere with other activities. It would be preferable to allow those not wearing stereo glasses to see a single, unghosted image [Figure 2.1b].

We accomplish simultaneous viewing of 3D and 2D images by replacing the pair of images [Left, Right] with a triplet [Left, Right, Neither], where those wearing stereo glasses see the Neither image with neither eye; only those *without* stereo glasses can see it. The Neither image is the negative of the Right image [Figure 2.1c] so that they cancel when superimposed, leaving only the Left.

Unfortunately, this raises the minimum black level for viewers without stereo glasses, drastically decreasing the contrast ratio. This can be mitigated by reduc-

ing the brightness of the Right image, α_R, to $\alpha_R < 100\%$, but maintaining the Left's full brightness.

If this adjustment is small, the effect on the 3D experience of viewers with stereo glasses is negligible, but the increase in contrast ratio for viewers without glasses is also modest. If this reduction is larger, the improved contrast ratio for viewers without glasses will be significant, but if too large, the 3D experience of viewers with glasses will deteriorate. We conduct experiments identifying the acceptable range of α_R for both viewers with and without glasses, and find in Sections 2.4.1 and 2.4.2 that both are satisfied when $20\% \leq \alpha_R \leq 60\%$.

When viewers wearing stereo glasses see a brighter image with one eye than the other, they soon become accustomed to this and report an acceptable 3D experience. However, they also report that horizontally-moving objects appear at different depths than stationary or vertically-moving objects with the same disparity. This small, but measurable, phenomenon is known as the "Pulfrich Effect" and is similar to a time-delay of several milliseconds in their perception of the darker image.

We conduct experiments to quantify this effect. We also measure a depth-distortion of similar magnitude caused by the 8 millisecond delay between the Left and Right images in a 120Hz display. The distortion is small enough that it is typically ignored by 3D content creators. These two effects cancel each other

when one eye's brightness is 40% that of the other eye.

The primary contribution of this work is a simple method to allow simultaneous viewing of 3D content by viewers with glasses, and 2D content by viewers without. We support this contribution with experiments measuring: viewer preferences among 2D degradation options, viewer ability to perceive 3D when one eye is dimmed, and the magnitude of the Pulfrich Effect in this system. Lastly, we demonstrate a prototype built using two commercial 3D projectors.

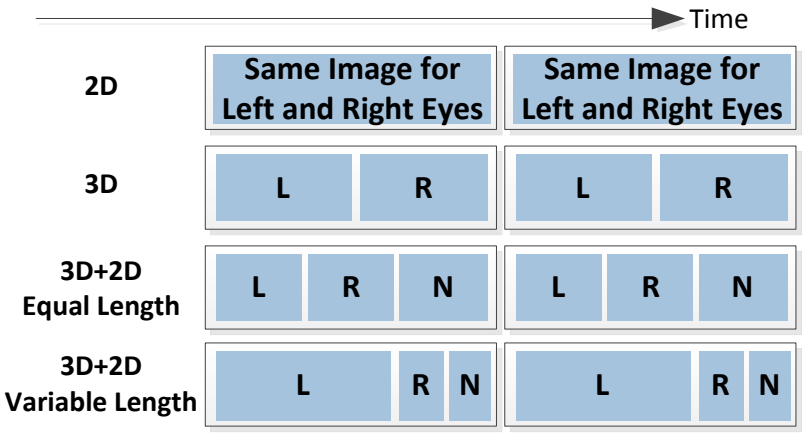

Figure 2.2: We propose a new sequence of frames. (1st row) A traditional 2D display shows a single image to both eyes. (2nd row) Each frame in a traditional sequential stereo display shows a distinct image to the Left (L) and Right (R) eyes of a viewer with glasses, while a viewer without glasses sees both images overlaid, with both eyes. (3rd row) Our 3D+2D display adds a third image (N) to each frame, shown to neither eye of the viewer with glasses, but seen by both eyes of a viewer without glasses. This third image is used to display the negative of the Right image, leaving them a low-contrast version of the Left image. (4th row) A 3D+2D display may allot more time to the left image to improve contrast, shortening the R and N images accordingly.

2.2 Related Work

Didyk et al. have also considered the problem of displaying a 3D image to a viewer wearing glasses while creating an acceptable 2D image for those without glasses, which they refer to as "Backward Compatible Stereo" [27, 28]. They reduce the disparity between objects in the left and right images to a minimal threshold, preferentially retaining high-frequency components. Smaller disparities make the 2D composite image more acceptable to viewers without glasses, but a ghost image remains. Reducing disparity also improves tolerance to cross-talk, where images intended for 3D viewers' left and right eyes are not entirely hidden from the other eye [79].

Anaglyph stereo uses two color channels with passive glasses to provide different views to each eye, while sacrificing color fidelity and showing a double-image to viewers not wearing stereo glasses. The most common example uses red and cyan filters, but amber and blue filters have been used to reduce ghosting seen by viewers not wearing glasses [80, 67].

Projection on an arbitrary textured object such as a brick wall is possible by adding a color cancellation term to the projected image [36, 37, 4, 12]. We use the same principle, treating one of the stereo channels as a texture to be canceled. Projecting one image, and neutralizing it with a compensation image, may also be used to project coded patterns visible to high speed cameras, but invisible to

human observers whose eyes integrate the images at high projection frame rates, while simultaneously projecting a desired image [68, 38].

The undesirable ghosting seen by viewers not wearing stereo glasses can also be avoided by using an autostereoscopic 3D display that does not require special glasses. Several techniques have been used to create such displays [29]. For example, a parallax barrier blocks light from reaching proscribed directions [63], and a lenticular array bends light toward the desired direction [55]. Autostereoscopic displays are generally more complex than glasses-based 3D displays and therefore more expensive.

2.3 Methods

A 3D+2D display is not restricted to a single stereo display technology. The key feature required is a third channel of information visible only to those not wearing glasses. In this section, we first discuss implementation options for an additional channel. We then discuss several options for the content of the third channel, which impacts the quality of the composite 2D image.

2.3.1 Implementing a Third Channel

Active-shutter displays show each image of the two-image frame packet sequentially, while the lenses of special stereo glasses become transparent or opaque

in synchrony to block each eye from seeing images not intended for it. The temporal pattern can easily include more channels, to support our method, or uses such as additional stereo viewpoints [1] [56]. We expect our method to be most popular with active shutter displays, whose stereo glasses are typically much more expensive than passive stereo glasses, costing $100 or more. Figure 2.2 illustrates possible temporal patterns supporting our method for both equal and variable length frames.

Several types of passive glasses may be used to build 3D displays. The most common glasses contain polarizing filters of orthogonal polarizations, while the display produces matching polarized images for each eye [50]. An alternate option uses lenses with orthogonal spectral filters, each allowing different narrow bands of red, green, and blue wavelengths to pass [49]. Passive stereo systems may produce the two images simultaneously with a pair of projectors, on interleaved rows of a flat-screen display, or sequentially with a projector and filter wheel.

The third channel we require could be provided by a single method, such as a third mutually orthogonal spectral filter, or by combining methods, such as using polarization and spectral filters together to produce four orthogonal channels.

Our prototype implementation combines polarization with active shutter projectors to achieve the necessary three channels.

2.3.2 Brightness of the Composite 2D Image

Our 3D+2D display shows three images each frame: L, R, and N. Viewers not wearing stereo glasses see only L, because N is chosen as the inverse of R such that N+R yields a uniform gray. The gray field raises the black level of the display: the brightness of the darkest pixel of the screen. If the three images are of equal brightness, the brightest pixel will be only twice as bright as the darkest pixel: a terrible contrast ratio. Allowing the L image to be brighter than the R and N images increases the contrast ratio. Several options are available to produce the N image, with different effects on contrast. We now analyze three possible options, depicted in Figure 2.3, and their impact, quantified in Figure 2.4.

Throughout, let L, R, and N be vectors of image pixels, containing all possible brightnesses. Let the functions $MAX(\cdot)$ and $MIN(\cdot)$ find the maximum or minimum element in the vector. Let $max_L = MAX(L)$ be the maximum possible brightness for any pixel in L, and similarly define max_R. Let $\alpha_R = max_R/max_L \leq 100\%$ refer to the brightness of the darker image R relative to L. Let $max_{2D} = MAX(L + R + N)$ be the brightness of the composite 2D image seen by viewers without glasses, and let its darkest possible pixel be $min_{2D} = MIN(L + R + N)$. Since N will be chosen to cancel out R, that is, $R + N = max_R$, we find that

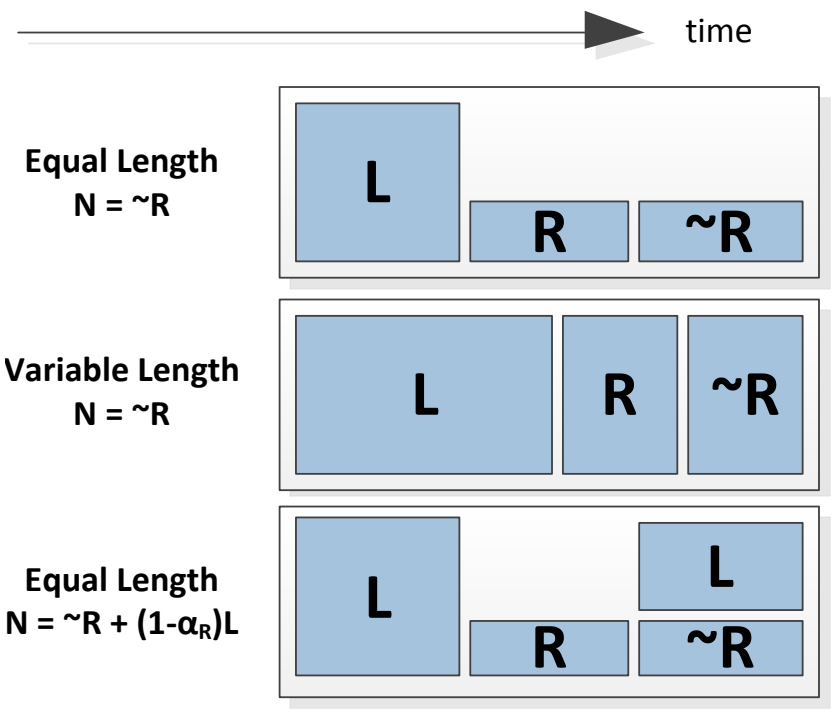

Figure 2.3: (Top Row) When all three frames have Equal Length and N = ˜R, some available light is wasted. (Middle Row) Variable-Length Frames waste no light, improving contrast for 2D viewers. (Bottom Row) Equal-Length frames may be improved by setting N=˜$R + (1 - \alpha_R) \cdot L$, wasting less light. For brevity we refer to the inverse of R as: ˜R= $(alpha_R \cdot max_L - R)$.

$$min_{2D} = MIN(R + N) = max_R = \alpha_R \cdot max_L \tag{2.1}$$

We now analyze how max_{2D} varies with α_R. First, in the simplest Equal-Length implementation of our technique, the three frames [L,R,N] are accorded

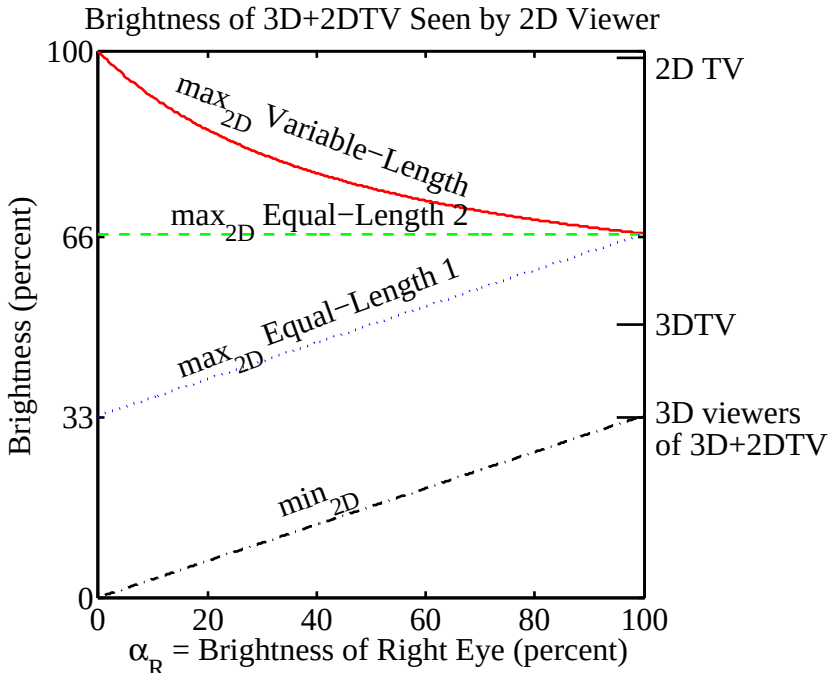

Figure 2.4: The contrast between the brightest pixel (max_{2D}) and the darkest pixel (min_{2D}) of the composite 2D image seen by viewers not wearing stereo glasses improves when α_R is decreased. Variable-Length frames produce a brighter 2D image than Equal-Length frames. For comparison, we have noted the brightness of a standard 2D TV, a standard 3D TV, and the 3D view of a 3D+2DTV.

equal time by the display. In this case, to darken the R and N images, α_R is reduced but the brightness of the L image max_L remains unchanged. To cancel R with N, we constrain N to: $R + N = max_R$, trivially achieved by setting $N = max_R - R = \alpha_R \cdot max_L - R$. Since the left frame is allotted one third of the display's photons, $max_L = 1/3$, so the total brightness of the composite image is then:

$$
\begin{aligned}
max_{2DE1} &= MAX(L + R + (\alpha_R \cdot max_L - R)) \\
&= max_L \cdot (1 + \alpha_R) \\
&= 1/3 + \alpha_R/3
\end{aligned}
\tag{2.2}
$$

Second, a Variable-Length display may instead dim the R and N images by affording them a smaller fraction of the total time in comparison to the L image. For example, plasma and DLP displays typically form each frame from shorter microframes, which could be reapportioned unequally among the L, R, and N images. Alternatively, more whole frames may be devoted to L than to R or N. For example, the variable-length frame case of $\alpha_R = 50\%$ may be achieved with an equal-length frame sequence (L,L,R,N). In this case, darkening the R and N images allows a corresponding increase in the brightness of L. Thus, while $N = max_R - R$, as before, max_L is now constrained as:

$$
\begin{aligned}
max_L &= 1 - 2 \cdot max_R \\
&= 1 - 2 \cdot \alpha_R \cdot max_L \\
&= 1/(1 + 2 \cdot \alpha_R)
\end{aligned}
$$

We thus find max_{2D} in this case as:

$$\begin{aligned} max_{2DV} &= MAX((1 - 2 \cdot max_R) + R + (max_R - R)) \\ &= 1 - maxR \\ &= 1 - \alpha_R \cdot max_L \\ &= 1 - \alpha_R \cdot 1/(1 + 2 \cdot \alpha_R) \\ &= (1 + \alpha_R)/(1 + 2 \cdot \alpha_R) \end{aligned} \tag{2.3}$$

Third, the Equal-Length implementation may be improved. Observe that, as initially described, with $\alpha_R < 100\%$ the N frame never shines with full brightness. Its unused brightness can be repurposed to duplicate L. Before, we had set $N = (max_R - R) = (\alpha_R \cdot max_L - R)$. Now, we add L in the unused portion of N:

$$N = (\alpha_R \cdot max_L - R) + (1 - \alpha_R) \cdot L$$

With $max_L = 1/3$ as before, the brightness of the composite image in this case is:

$$
\begin{aligned}
max_{2DE2} &= MAX(L + R + (\alpha_R \cdot max_L - R) \\
&\quad + (1 - \alpha_R) \cdot L) \\
&= max_L + (\alpha_R \cdot max_L) \\
&\quad + (1 - \alpha_R) \cdot max_L \\
&= 2 \cdot max_L \\
&= 2/3
\end{aligned}
\tag{2.4}
$$

This constant brightness falls roughly halfway between the two simpler implementations, as seen in Figure 2.4.

Viewers wearing stereo glasses also experience lower brightness when this system is employed, because precious display time is devoted to emitting photons for viewers without glasses that never reach either eye of those wearing glasses. When using equal length frames, 3D viewers experience a brightness reduction of 33%. With variable length frames the maximum 3D brightness might increase or decrease depending on the choice of α_R.

2.4 Experiments

We conducted several experiments to assess our design's viability. Subjects viewed a 42" plasma 3DTV from a distance of approximately 1.3m; at this distance, 1 pixel subtends 0.0215 degrees.

In all 3D experiments, subjects wore 120Hz shutter glasses, and the brightness of the right eye, α_R, was progressively lowered. Subjects were instructed to respond only when they felt their eyes had adjusted, typically answering immediately for α_R near 100% and waiting about a minute for α_R near 0%. Partial adaptation to new light levels occurs immediately, with additional adaptation over approximately 20 minutes [81]. We measured only the immediate effects; significant effects have previously been found when allowing subjects to adapt to differential light levels for as little as 1 minute [2] or 3 minutes [82, 30]. Ambient lighting was moderate; the Pulfrich effect has been observed under these conditions [30]. We measured subjects' depth perception and did not survey visual comfort. Kooi et al tested viewer comfort with $\alpha_R = 75\%$ and found it acceptable, while Beldie et al tested several values of α_R with two photos and found $\alpha_R =$ 40% and 60% acceptable [51, 11].

All confidence intervals $c.i._{95\%} = \pm 2 \cdot \sigma/\sqrt{n}$ were calculated for each α_R by finding the mean m_s for each of the n subjects s and computing the standard deviation σ over the means $[m_s]$.

2.4.1 2D Viewer Preferences

Standard 3DTVs show a double-image to 2D viewers. A 3D+2DTV removes this ghosting, but also reduces the contrast. This experiment investigates at what level of contrast viewers prefer the original, ghosted image to a lower-contrast image without ghosting.

We presented subjects with two images on a standard 2D TV, and asked them to choose which they prefer, as in Figure 2.5(a). On the left, we simulated the double-image L+R seen on a traditional active-shutter stereoscopic display; on the right, we simulated L+R+N in accordance with the analysis of the previous section. We conducted experiments simulating simple Equal-Length frames [Figure 2.3(top)] and Variable-Length frames [Figure 2.3(middle)] with different subjects for each.

Each subject viewed ten test images at eleven values of α_R, with n=10 subjects participating in each experiment.

At high contrast levels viewers nearly uniformly prefer our method. Only at very low 2:1 contrast do viewers find contrast reduction equally objectionable as ghosting [Figure 2.5(b)]. For a display with $\alpha_R = 20\%$ and equal-length frames, 80% of our subjects prefer low-contrast images without ghosting. A display capable of producing variable-length frames is able to provide a higher contrast for the same value of α_R. With this design, the preference for our system rises to 95% at

$\alpha_R = 20\%$.

2.4.2 3D Viewer Depth Perception

We display a brighter image to 3D viewers' left eyes than to their right eyes. However, stereoscopic vision is degraded when the images seen by each eye become disimilar [23]. Small brightness differences may be imperceptible, but an all-black right-eye image obviously precludes stereoscopic vision. This experiment quantifies depth perception between these extremes.

We presented subjects wearing shutter glasses, with a stereoscopic display of a 7x3 array of wooden boxes, as in Figure 2.6(a). The top and bottom rows were identical and unchanged throughout the experiment, with each box in the row at a different depth, in the range of [-0.13 to +0.13] degrees disparity. The middle row of boxes were all at the same depth; this depth was varied in each trial. Subjects were asked to identify which column of the top-and-bottom-row boxes was at the same depth as the boxes in the middle row. They answered e.g. "The boxes in the middle row are at the same depth as the top and bottom boxes in column three."

$n = 6$ subjects each undertook 130 trials, randomly varied across 5 different brightness levels and 13 possible depths.

We find that depth perception is surprisingly robust against differences in

image brightness between the two eyes, and is not significantly affected until α_R falls below 25% [Figure 2.6(b)].

In a second experiment, we showed subjects a set of five vertical sticks, as seen in Figure 2.7(a), while again the brightness of the image seen by their left and right eyes differed. One of the three central sticks was displayed with a different disparity than the other 4 sticks, so that it was perceived as lying at a different depth. The subject was asked to identify which stick was at a different depth than the other sticks. Each subject made judgements with α_R varied to 27 levels, binned to 14 levels in Figure 2.7(b). n=15 subjects participated in the experiment. Individual trial depths were chosen randomly in the range of [0.02 to 0.15] degrees of disparity.

Viewers' ability to perceive depth differences was not impaired until the brightness of the darker eye became very dark, similarly to the previous experiment. Accuracy fell slowly from the equal-brightness case until α_R fell below 10%. When $\alpha_R < 10\%$, subjects answered as if guessing randomly.

2.4.3 Moving 3D Objects and the Pulfrich Effect

When subjects' left and right eyes view images of unequal brightness, they perceive accurate depths for stationary and vertically-moving objects, but the depths they report for objects moving horizontally show a predictable distortion.

This illusion, known as the Pulfrich effect [65, 57], has been used to produce 3D effects on television by distributing tens of millions of paper glasses with one dark lens [83]. Similar glasses are also used clinically to diagnose and treat depth perception difficulties [26, 41].

The Pulfrich effect is approximately linearly dependent on horizontal speed, and therefore consistent with a time delay in the darker image reaching the brain. While the undimmed eye sees the present world, the darkened eye sees the world as it was a few milliseconds ago. The undimmed eye sees a horizontally-moving object's current position, but the darker eye its previous position, changing its apparent disparity, and thus its apparent depth [32].

The human eye contains two types of light-sensing cells: rods and cones. Rods are far more sensitive to low light levels, and also have a longer latency than cones in transmitting images to the brain. The Duplicity Theory of Vision holds that light-adapted vision involves chiefly the cones, while dark-adapted vision relies more on rods than cones, precipitating the delay in perception [64].

In addition to the virtual time delay caused by a dimmed eye, all sequential 3D displays also experience a real time delay: on a 120Hz display that shows (left,right) image pairs at 60Hz, the image shown to the right eye will always lag behind the image shown to the left eye (or vice-versa) by 1/120th of a second (8 milliseconds). This time delay is accompanied by an attendant distortion of depth.

3D Content creators often ignore this distortion: the Blu-ray 3D specification, and Nvidia developer advice treats the left and right frames as simultaneous [86, 35].

If the virtual time delay caused by the Pulfrich effect is of a similar magnitude as the actual time delay of sequential-frame stereo displays, we expect that it can likewise be safely ignored. We conducted an experiment to measure the impact of the Pulfrich effect on our system. We find the two effects to be of similar magnitude, and that one effect may be used to cancel out the other.

We showed subjects a 3D scene containing two rows of seven stationary boxes, as in the first experiment of section 2.4.2. The boxes lie at different depths, with the left-most pair of boxes furthest away, and the right-most pair closest. A moving box passed horizontally between the two rows of stationary boxes, as in Figure [2.8(a)].

The subjects were asked to identify the stationary box whose depth most closely matched the depth of the moving box. The stationary boxes were unaltered throughout the experiment, but the moving box's speed, direction, and disparity (true depth) were randomly varied. Due to the depth distortions of the Pulfrich Effect, the subjects estimated a consistently different depth for the moving object, depending on its speed and the brightness of each eye.

This experiment was conducted at each of nine levels of α, with the center box moving both left and right at each of five speeds. n=3 subjects undertook six trials

at each condition, with the center box's disparity randomly permuted between (0, 0.13, or 0.26) degrees; each data point is thus averaged over 18 responses.

The illusory depths seen by the subjects are shown in Figure 2.8(b). Higher speeds create larger distortions, and movements in opposite directions (right/left) produce opposite depth illusions (closer/further), consistent with the time-delay theory.

Errors are smallest when the left eye is dimmed to approximately 40% the brightness of the right eye, rather than when both eyes have equal brightness. At this value, the virtual time delay caused by the Pulfrich Effect largely cancel out the 8ms delay in the sequential display of left and right stereo images. Greater brightness differentials produce larger illusions, consistent with longer latencies.

If the virtual delay were large, it would need to be offset by delaying the dimmed video feed. This appears unnecessary; the Pulfrich Effect poses no obstacle to dimming one eye of a 3D+2DTV.

2.5 Prototype

Our protoype uses two projectors and a polarization-preserving screen, seen in Figure 2.9(a). A standard, unpolarized, 3D (120Hz) projector synced to LCD active-shutter glasses displays the images L and R seen by the left and right eyes of the viewer wearing glasses. The second projector displays the 3rd image, N,

and is linearly polarized. The LCD glasses contain an orthogonal linear polarizer that blocks the N image from the second projector.

Note that the first projector spends half its light on the L frame and half on the R frame. The second projector spends all its light on the N frame, but half of this light is lost to the linear polarizer. This leaves all three frames with similar brightness. Geometric and photometric calibration are required to align the images and correct non-linearities in projected brightness [15].

We evaluated our system by displaying images in standard 3D, as well as using our 3D+2D method. Figure 2.11 shows a number of examples at $\alpha_R = 30\%$, together with the third channel that we introduced. The example images were captured by photographing the projection screen through a pair of shutter glasses that reveals the images delivered to the left and right eyes of 3D viewers, while the region outside the glasses shows the experience of viewers without glasses. In our implementation, only very minor ghosting is visible in the 2D region, and the third channel is blocked by shutter glasses.

We repeated the 2D viewer preference test described in Section 2.4.1 with our prototype, testing only $\alpha_R < 50\%$, since this is the range in which our prototype and method are most useful. Viewer preferences on the prototype closely matched preferences for simulated Equal-Length Frames on a 2D display, as seen in Figure 2.9(b).

We compared our prototype to Didyk et al.'s "Backward Compatible Stereo," which improves the 2D viewing experience by reducing stereo disparity to the minimum that still maintains a perception of depth for 3D viewers [27, 28]. The original stereo image is shown in Figure 2.10a. The results of Didyk et al. in Figure 2.10b and our method in Figure 2.10c demonstrate that both methods greatly reduce ghosting. This is apparent in the reduced blurriness of the overall image, at left, and explicitly in the closeup, at right. While our method completely removes ghosting, it also reduces contrast. Using the results of Didyk et al. as the input to our method, as in Figure 2.10d, provides similar results to using the original inputs.

2.6 Limitations

Our prototype's screen exhibits significant specular reflection, limiting radiometric calibration accuracy, so that some residual ghosting remains. Its preservation of polarization is also imperfect, leaving the N image slightly visible through stereo glasses. Higher quality components and calibration would rectify these issues.

This work has focused on completely eliminating ghosting, and we have compared 2D viewer preferences regarding ghosting and contrast under this assumption. However, when the ghost is relatively dim, it is not as objectionable [79].

Further study might reveal an optimum tradeoff between ghosting and contrast reduction by only partially canceling the ghost image.

Several factors known to affect perceived image quality and stereo fusion have been left unexplored as beyond the scope of this work. Studies investigating the effects of image contrast, image content, eye dominance, and dark adaptation over time may prove insightful.

2.7 Discussion

Many current 3D displays require that viewers wishing to see the 3D scene wear special stereo glasses; viewers without glasses not only do not see a 3D scene, but see an unappealing and confusing double-image. 3D display technology has proven popular and useful in entertainment, engineering, and medical applications where it may be impractical or undesirable to require all viewers to wear stereo glasses.

We have demonstrated a method to produce 3D displays where viewers wearing glasses see a 3D scene, while those without glasses see a single 2D scene. We have shown that reducing the brightness of one of the images shown to the 3D viewer does not interfere with depth perception, while allowing acceptable contrast for the 2D viewer. We have also demonstrated that depth-distortions caused by the Pulfrich effect are only of similar magnitude to other distortions present in all sequential-stereo displays, and can offset them.

(a)

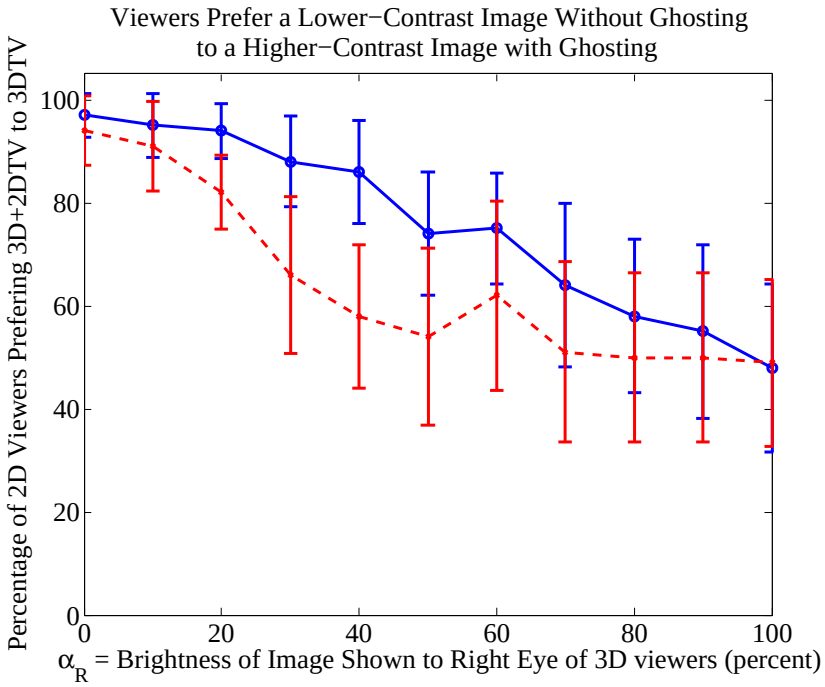

(b)

Figure 2.5: (a) We showed viewers two versions of an image and asked which they prefer: (left) the ghosted double-image they would see on a typical 3D display if they did not wear stereo glasses, or (right) the lower-contrast image without ghosting that they would see on our display. In this example, $\alpha_R = 30\%$. (b) We asked subjects to choose between the image they would see without glasses on a traditional 3D TV and on our display. Note that our display is preferred by a majority of users. As the brightness of one eye decreases, the contrast ratio increases, and a greater percentage of viewers prefer our display. 95% confidence intervals shown.

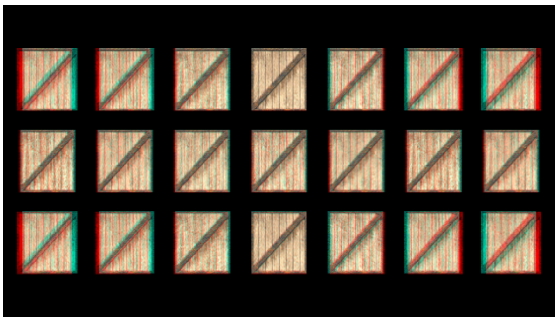

(a)

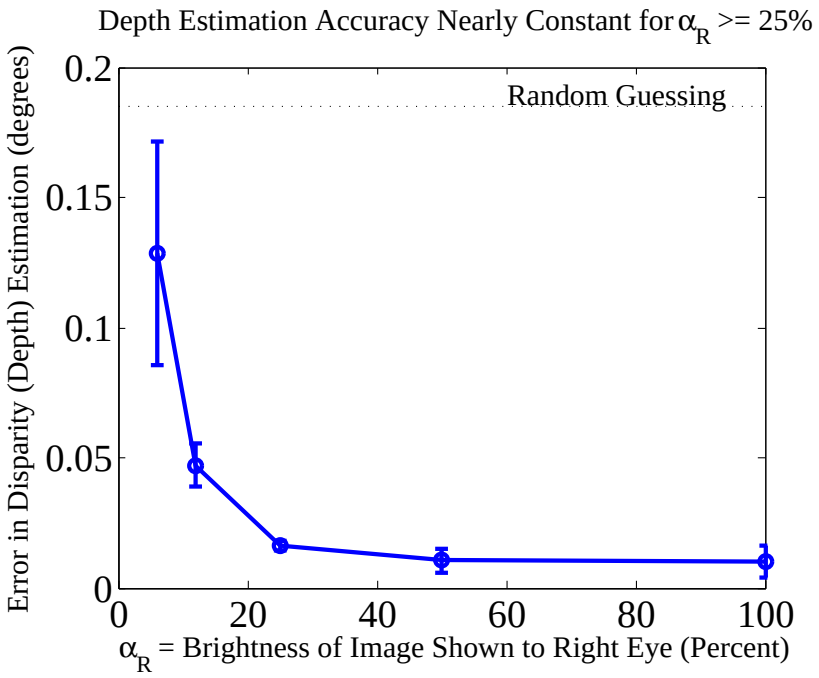

(b)

Figure 2.6: (a) This experiment quantified viewers' ability to perceive depth in static images on a stereoscopic display when one eye is presented with a darker image than the other eye. The subject was shown 3 rows of boxes, reproduced here in anaglyph format for illustrative purposes. The top row and bottom row are identical, featuring 7 boxes of different disparities, with the left-most box appearing furthest away and the right-most box closest. The middle row contains 7 boxes, all with the same disparity. The subject was asked which column in the top and bottom rows is at the same depth as the boxes in the middle row. (b) As one eye's brightness decreases, viewers' ability to perceive depth was excellent for $\alpha_R \geq 25\%$. 95% confidence intervals shown.

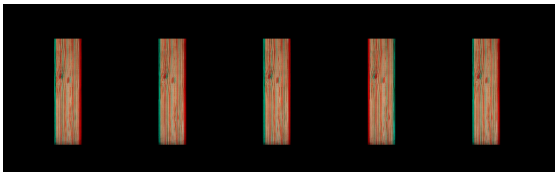

(a)

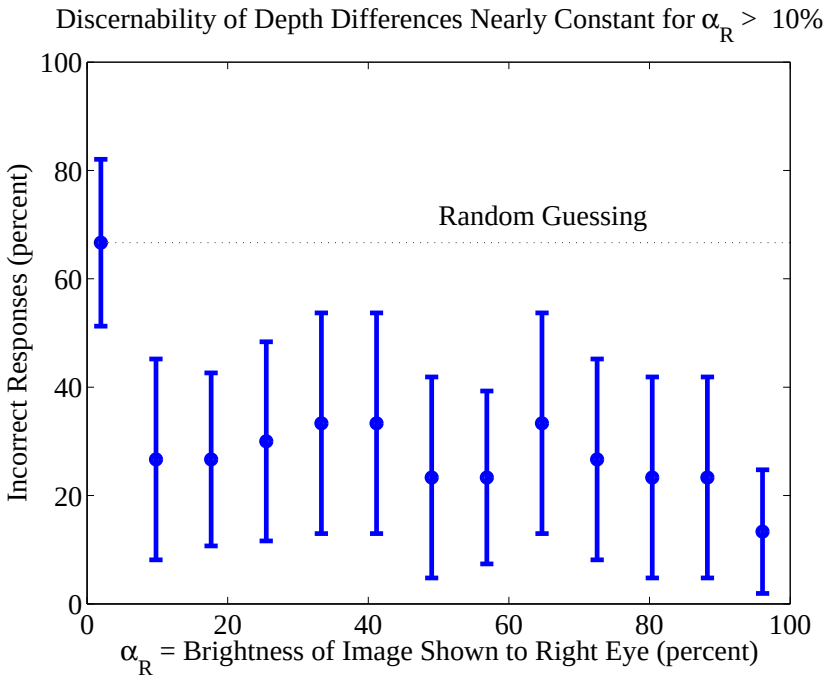

(b)

Figure 2.7: (a) In this experiment, subjects viewed a stereo display where one eye viewed a brighter image than the other. Subjects viewed five sticks, with one displayed at a different disparity than the other four. This screen shot has been converted to anaglyph form. (b) Viewers' ability to perceive depth differences was quite good for $\alpha_R \geq 10\%$. 95% confidence intervals shown.

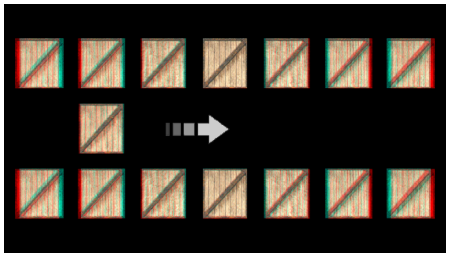

(a)

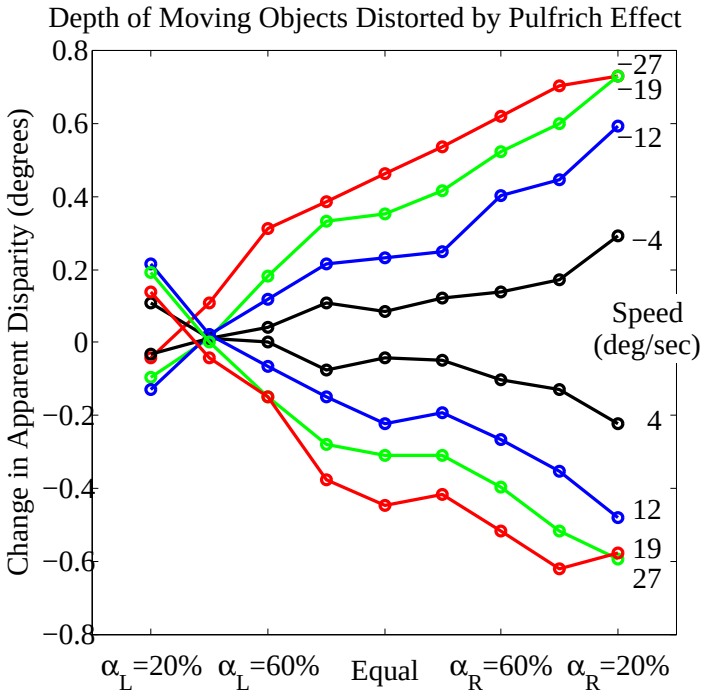

(b)

Figure 2.8: (a) This experiment quantifies the impact of the Pulfrich Effect on depth perception. Subjects viewed two identical rows of seven boxes, with the left-most boxes furthest and the right-most boxes closest. A moving box passed between the two rows, and the subject chose which stationary box was at the same depth as the moving box. Negative speeds denote movement in the opposite direction. (b) When one eye is brighter than the other, the depth of moving objects is misperceived. Faster objects have a greater distortion in their apparent depth. A larger difference between the brightness of the two eyes also causes a greater distortion in the perceived depth. Minimal distortion is found at $\alpha_L = 40\%$, not at equal brightness, because of the frame-sequential stereo display. 95% confidence interval shown only for one speed, for clarity.

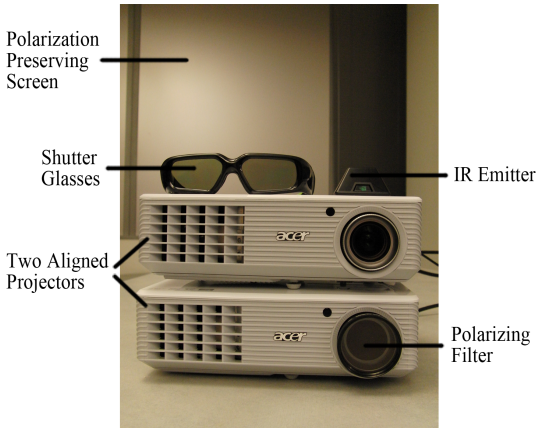

(a)

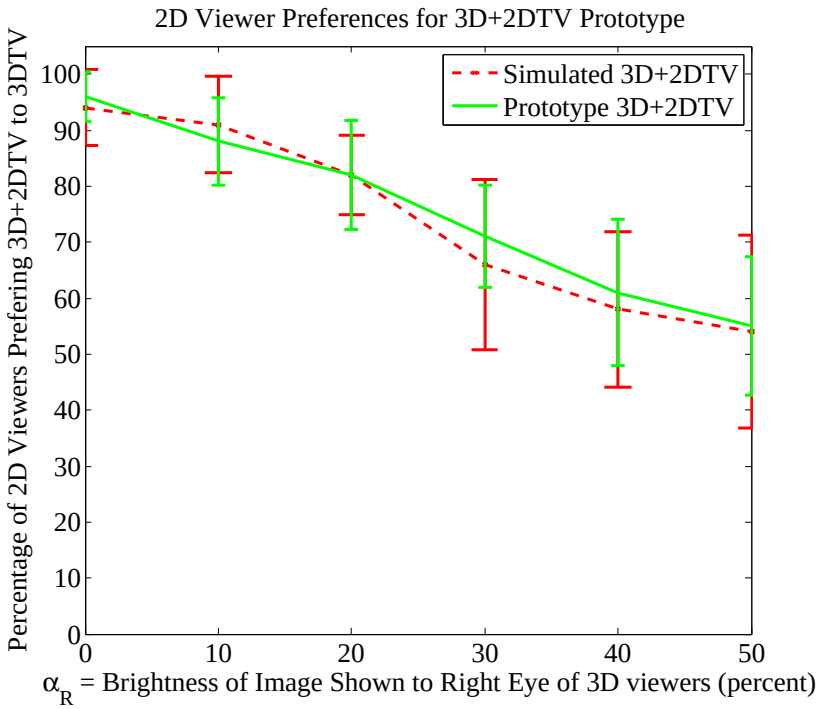

(b)

Figure 2.9: (a) Our prototype uses two projectors and a polarization-preserving screen. An unpolarized 3D projector synchronized with active-shutter LCD 3D glasses shows the L and R images. A second, linearly polarized projector shows the N image. LCD shutter glasses contain a linear polarizing filter that blocks the light from the polarized projector. (b) We repeated the 2D Viewer preference test with our prototype and obtained very similar results to the initial experiment. 95% confidence intervals shown.

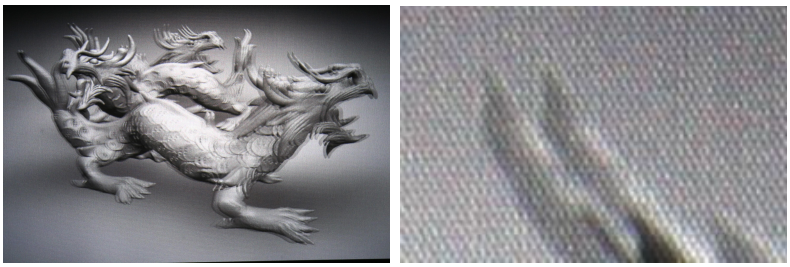

(a) Original Images on a Standard 3DTV

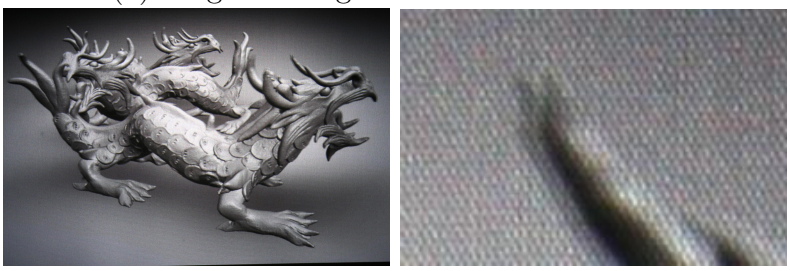

(b) Didyk et al's Result on a Standard 3DTV

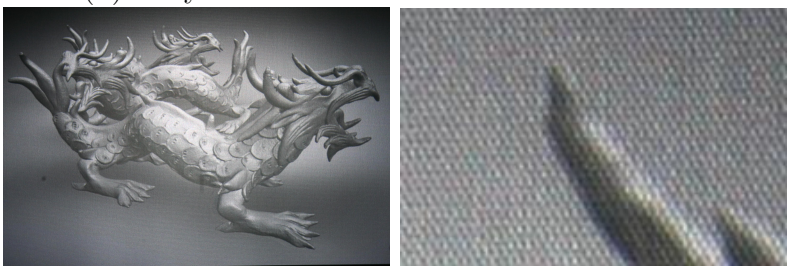

(c) Original Images on our 3D+2DTV

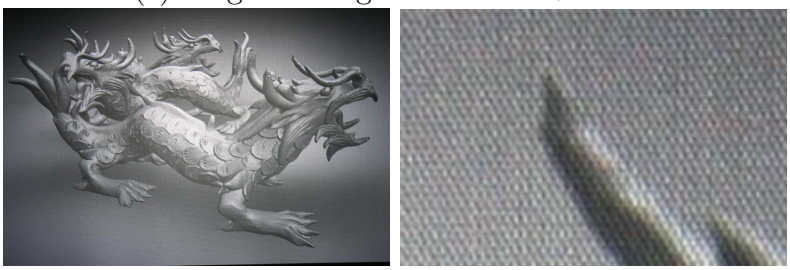

(d) Didyk et al.'s Result on our 3D+2DTV

Figure 2.10: Here we compare our method, using a 3rd frame to cancel out one of the two stereo images, to the technique of Didyk et al, wherein the disparity between the two images is reduced. The left column shows the overall image, while the right column shows a closeup of the dragon's horns, one of the largest areas of ghosting. (a) 2D view of unmodified stereo image. (b) Didyk et al. (c) our method with $\alpha_R = 30\%$. (d) applying Didyk et al, then our method.

3DTV 3D+2DTV 3rd Channel

Figure 2.11: Here we show several examples of our prototype in use. In each image, the projector screen is visible directly at the top of the image, and through each lens of the stereo glasses at the bottom of the image. (Left Column) A typical 3D display (Middle Column) our 3D+2D prototype, with $\alpha_R = 30\%$ (Right Column) the 3rd channel we display to cancel out the right-eye image. Images used with permission from Flickr users GammaMan, Trondheim_Byarkiv, and Isaiah-v.

Chapter 3

Printing Reflectance Functions

3.1 Introduction

An object's appearance changes when lit from different directions. Shadows are cast and specular highlights shine. Even diffuse surfaces change appearance, providing a valuable perceptual cue to object shape and material. Unfortunately, when printed on paper, this variability is lost. A photographic print represents just one appearance, regardless of the ambient lighting when the print is viewed.

A real scene's appearance at each point depends upon the incident illumination, the surface Bi-directional Reflectance Distribution Function (BRDF), [60] and the local surface orientation. Typical printers can not print images which react to incident illumination because they use inks with a limited range of BRDFs

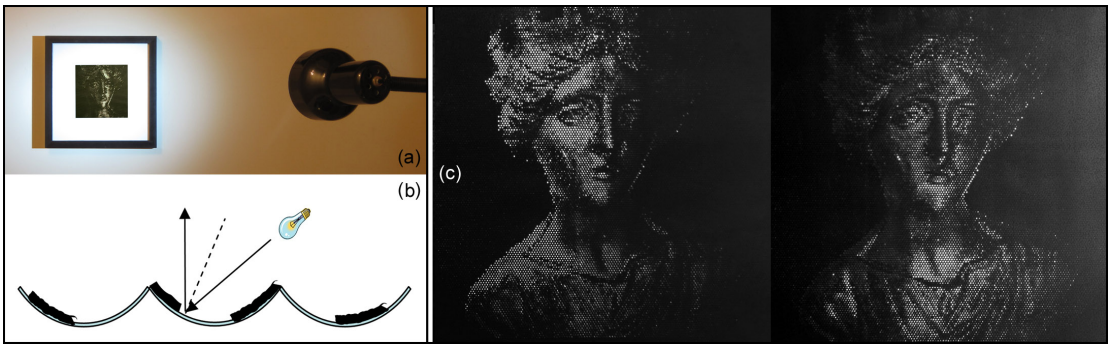

Figure 3.1: (a) We introduce a method for printing reflectance functions, images which correctly respond to the direction from which they are illuminated. (b) A microgeometry reflective substrate returns a specular highlight unless opaque ink is deposited on the point corresponding to a particular incident illumination direction. (c) Our prototype is shown displaying different images in response to two different illumination conditions. Image used with the permission of Cultural Heritage Imaging.

and print on paper which is flat, having a single orientation everywhere.

Simply expanding the range of specularity of the printer's inks is insufficient. Surface orientation plays an important role in controlling the observed intensity of reflected light. To correctly represent this interaction with light, the local surface orientation of the paper needs to match that of the original object.

A conceptually simple solution would be to orient the local surface normal of each pixel appropriately, and to print the object's surface BRDF onto this pixel. However this requires changing the physical shape of the underlying paper at the time of printing. This is an expensive process that requires equipment far more complex than just depositing ink onto a surface.

We introduce a method of printing reflectance functions [25] which makes use

of paper with a static microgeometry structure, shown in Figure 3.1b. The paper consists of a hexagonal array of spherical depressions that have a specularly reflective surface. By selectively printing opaque or partially opaque ink on portions of that surface, we can control whether a specific incident lighting direction returns a specular highlight, and to what degree. Because both the paper and ink are designed to minimize diffuse reflections, the specular return fully controls the appearance of this reflectance paper and is used to control its appearance as a function of lighting direction. Note that this scheme gives sufficient expressive power to specify two dimensions of the 4D BRDF. We choose to specify, for a single fixed viewing direction, how much light will be reflected from each incident lighting direction, thereby building up an arbitrary reflectance function.

The primary contribution of this work is a method for printing reflectance functions using existing printers and special paper. We support this contribution with a ray traced simulation and analysis of errors, as well as with an initial prototype implementation.

3.2 Related Work

Several research groups have demonstrated passive physical surfaces and devices to generate controlled reflectance functions. Fuchs et al. share our goal of generating the reflectance field of a scene, and built an optical assembly that

produces lighting-dependent imagery using a lenticular array [34]. Light entering a lenslet from behind the print will pass through a different location depending on its entering angle. A more complex optical assembly is presented that has both view- and lighting-dependence, albeit at low resolution. These methods generate a back-lit image as opposed to our printing process, which generates controlled reflectance functions designed to be viewed in a reflective manner, like a conventional print.

Matusik et al. use a printing process, as we do, to generate images with a varying BRDF by linear blends of a set of inks and foils [54]. In their work, BRDFs are associated with normals perpendicular to the surface of the paper, and thus do not afford enough control to mimic the reflectance functions of an arbitrary 3D scene.

Alexa et al. share one of our primary goals, namely to develop a mechanism to generate arbitrary images as a function of varying lighting direction by controlling the reflectance function of a surface [3]. In their work, diffuse reflection is used as a basis for generating images, which results in the ability to specify two samples of the reflectance function in practice, and thus two approximate output images. By using only specular reflections to build up an arbitrary reflectance functions and a 32x32 printed matte overlaying each dimple, our design can achieve greater independence between many more images.

Regg et al. also build up an image using only specular highlights, as we do [69]. Their goal is not to mimic the reflection functions of a real scene, but to create a set of points with a controlled 3d percept. A curved scratch on acrylic generates highlights at two separate locations along the scratch, for the viewpoints associated with each eye. A set of such curves are etched under computer control, each yielding a perceived 3D point location to build up percept of the desired 3D shape.

Several researchers have also explored the possibility of milling a homogeneous material's surface to achieve a desired effect. The milling induces a variation of the material's appearance by setting the orientation of each surface patch to reflect more or less light in desired directions. Such methods can be used to automate the design of Bas Relief in an approach that preserves important 3D depth cues [88]. Weyrich et al. mill a custom array of microfacets to control the overall BRDF of a patch of aluminum, effectively treating it as a single pixel [89]. In addition, the authors validate their ability to generate a complex BRDF by viewing reflected light off that patch on a secondary surface. Our method extends their ability to control reflectance from a single large pixel to a 2D image and differs in that our pixels are three orders of magnitude smaller in area, and are produced with a printer rather than a milling machine.

Lastly, Nayar et al. present an active approach to varying the appearance

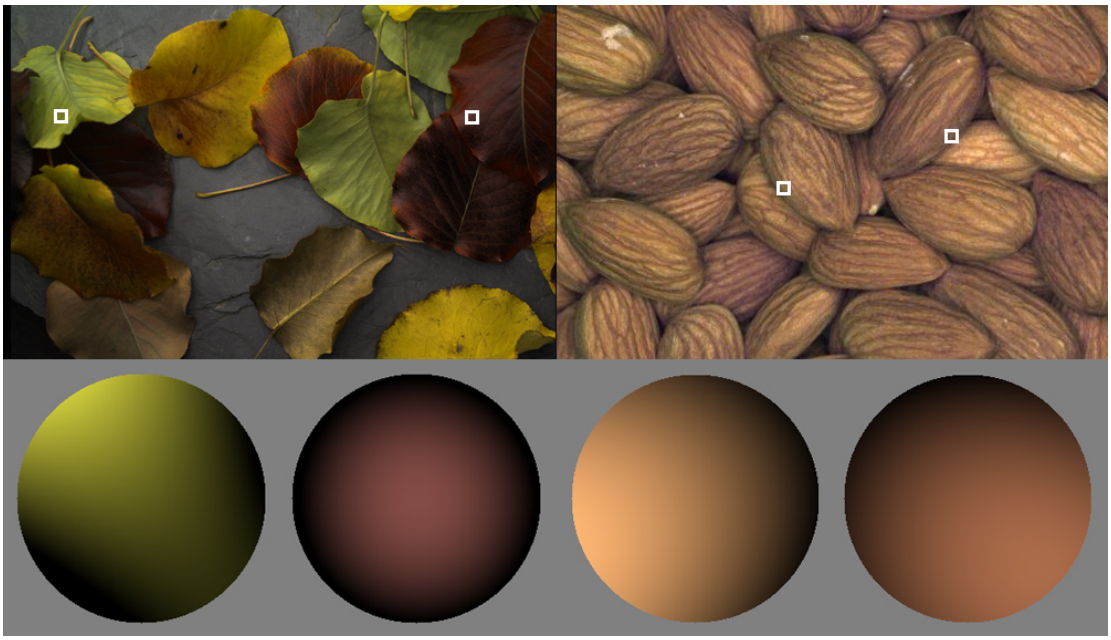

Figure 3.2: The shape of a reflectance function (bottom) is dependent on the orientation and shadowing of the underlying surface. For example, the indicated pixels of the almonds on the right are chosen on opposite edges, so respond most brightly to light from opposing angles.

of a photograph under lighting changes by measuring the 2D illumination field, and then generating correspondingly lit images [59]. We strive to achieve similar functionality with a passive surface.

3.3 Methods

3.3.1 Reflectance Functions

2D reflectance function can be measured easily for real scenes by taking photographs lit from many different lighting directions, and a wide variety of devices

have been constructed to do so [25, 52, 61, 87]. In principle this function can contain high frequency components from specular highlights and hard shadows. However, in practice even low order approximations are sufficient for a variety of presentation and evaluation purposes [40, 66]. Manipulations of reflectance functions have proved useful especially in both cinema applications [25, 62] and archeological applications [33, 58].

Figure 3.2 shows several examples of reflectance functions for individual pixels, parameterized by the incoming illumination angle. On the left is a set of colored leaves, with the reflectance of specific marked pixels shown below. On the right are pixels chosen on opposite edges of almonds. Note that shadows are present in both cases at extreme lighting angles, and this shows up as dark regions in the reflectance function. Although the reflectance functions are very smooth, they are oriented differently at different pixels, and this is sufficient to provide a convincing perception of object shape when the incident light angle is changed. Flat paper printed with standard inks can only represent centered isotropic functions similar to the red pixel in this example. Our design can represent arbitrary functions. Note that for the photographically collected reflectance functions we employ, global illumination effects such as shadows and interreflections are present at the time of image capture, and thus are represented in the reflectance function. Whether such secondary effects are modeled is dependent on the generation of the

reflectance data that is input to our system, whether captured or simulated.

3.3.2 Reflectance Paper

We designed our reflectance paper substrate with spherical dimples, which we metalized to produce a mirrored finish. Since the surface is entirely specular, light incident on the paper from some particular direction will strike the entire surface, but only light that strikes points with a corresponding particular orientation will be reflected toward the viewer. Opaque ink covering this point on the dimple prevents the dimple from being illuminated from this direction, as seen in Figure 3.1b.

Note that achieving sensitivity to the full hemisphere of possible lighting directions requires a reflecting surface consisting of only 1/4th spheres, not hemispheres. To simplify the process of printing on a geometrically complex surface even further, we constructed media consisting of 1/5th spheres, subtending approximately 70 degrees, which will respond to 140 degrees of the possible 180 degrees of incoming lighting directions. Lighting outside of this range will not change the appearance of the print measurably.

Our initial prototype considers only binary, black and white printing using fully opaque ink. With this constraint, levels of gray may be approximated through dithering, as discussed in Section 3.3.7. To achieve true continuous tone one could

vary the thickness or transparency of the ink to modulate the reflectivity instead of providing a binary opaque mask as we did.

We have described this design using an orthographic directional light source for simplicity of description. The design itself is theoretically correct and functions correctly under arbitrary illumination environments due to the linearity of light transport. This includes area lights and near field point illuminants, as well as more general lighting environments.

Our reflectance paper builds arbitrary reflectance functions from specular highlights off a dimpled substrate. This is a dramatically different mechanism from conventional paper, which uses primarily diffuse reflections off paper and ink. The question arises how brightness compares between the two cases. For conventional white paper, assuming ideal diffuse reflection, all incoming light energy from any direction is distributed uniformly across possible viewing directions. This is also the case in our dimpled reflective substrate, assuming ideal specular reflections. Across a single dimple, our facet microgeometry distributes all incoming light uniformly across viewing direction. In both cases ink serves to attenuate the brightness spatially. We therefore expect reflectance paper prints to have comparable brightness to prints on conventional white paper.

Contrast ratio in our design is limited by the performance of the inks. In particular, any diffuse reflections introduced by either the dimpled metallic substrate

or the inks is distributed across all viewing directions, and thus serves to decrease the contrast ratio of our passive 'display'.

3.3.3 Avoiding Ink Specularity

One design deposits ink directly on the dimples, blocking the specular highlights of the reflective substrate, without introducing diffuse contributions that would lower the contrast ratio. Dry electrophotographic toner, used in laserjet printers, or liquid electrophotographic (LEP) inks, used in commercial printing, are unfortunately quite specular themselves. We were able to reliably deposit LEP inks successfully in a controlled manner on our dimpled substrate, even on large 2.3 mm. diameter dimples, but the specularity of the ink itself causes specular returns from lighting directions the ink is intended to block.

To achieve a more matte finish, matte topcoats could be applied, but as these operate by introducing microfacets that spread the surface reflections to a wider range of reflection angles; they create new problems. Although effective at reducing specular highlights off the ink, this also distributes the energy from the uncovered specular region in a diffuse manner, which unacceptably lowers the contrast ratio of our reflectance paper.

We have developed another approach to avoid the specular return from the ink. Instead of printing directly on the metalized dimples themselves, we use a

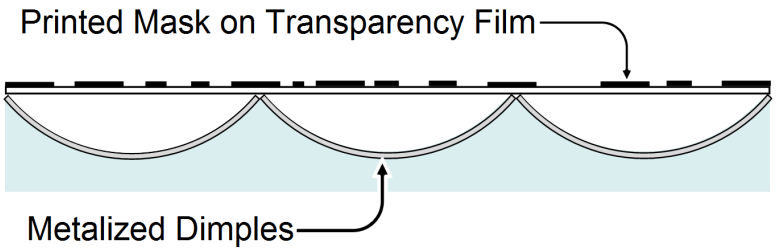

Figure 3.3: To avoid artifacts due to ink specularity we print reflectance functions on an overlaid transparency instead of directly on the metalized dimples themselves.

standard laser printer to print a mask on a transparent film that is placed on top of the dimpled surface, as shown in Figure 3.3. In this approach, specular highlights off the inks on the transparency film are in the reflection direction given by the normal of the film, not the normals associated with the shape of the metalized dimple. These unwanted specular highlights are generally never directed towards the viewer, and therefore avoided. This allows us to avoid the difficulty of inventing inks that do not have significant specular return.

3.3.4 Geometry of the Specular Dimples and Transparency Film

The geometry for a single dimple is shown in Figure 3.4. A right-handed coordinate system with origin o is placed at the center of the sphere. The dimple shape is a portion of the lower hemisphere, satisfying:

$$x^2 + y^2 + z^2 = r^2 \tag{3.1}$$

$$z \leq z_0 \tag{3.2}$$

with the top of the grid of dimples, where the transparency film may be placed, at $z = z_0$. The value $z_0 = -cos(\eta/2)$ is determined by the angle η defining the proportion of hemisphere formed by o and two opposite points on the dimple at the grid surface ($z = z_0$). Using the implicit representation of the sphere, $F(x, y, z) \equiv x^2 + y^2 + z^2 - r^2 = 0$, the (upward pointing) normal n at a point $s = (x, y, z)^T$ on the dimple is given by:

$$n = \frac{-\Delta F(s)}{\|\Delta F(s)\|} = \frac{-s}{r} = (\frac{-x}{r}, \frac{-y}{r}, \frac{-z}{r})^T \tag{3.3}$$

To simplify notation for the subsequent derivations we set $r = 1$ without loss of generality. For a given view direction v, any point s on the specular dimple surface responds to a single lighting direction, $l(s)$. The direction l corresponds to a specular reflection of v through the normal vector n at s. In other words,

$$l(s) = (2nn^T - I)v = (2ss^T - I)v \tag{3.4}$$

where, with $r = 1$, we obtain the right-hand term in Equation 3.4 by substituting $n = -s$ from Equation 3.3. In Equation 3.4, the matrix I is the 3x3 identity

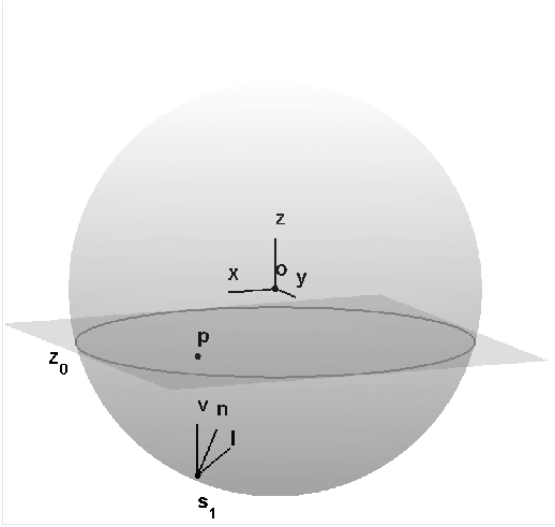

Figure 3.4: The geometry of a single metalized dimple with an overlaid transparency.

matrix, and the vectors l, v, and s are unit vectors, with the convention that v and l point towards the viewer and the light, respectively. For use in the subsequent discussion, we define the matrix $M \equiv 2ss^T - I$.The components of $l(s) = (l_x, l_y, l_z)^T$, for a general view direction $v = (v_x, v_y, v_z)^T$, are given by

$$\begin{pmatrix} l_x \\ l_y \\ l_z \end{pmatrix} = \begin{pmatrix} 2x^2 - 1 & 2xy & 2xz \\ 2xy & xy^2 - 1 & 2yz \\ 2xz & 2yz & 2z^2 - 1 \end{pmatrix} \begin{pmatrix} v_x \\ v_y \\ v_z \end{pmatrix} \tag{3.5}$$

Since each point s on the dimple responds to a single lighting direction, each individual dimple samples the *lighting space* or reflectance function at one spatial location, and the set of dimples on the hexagonal grid provide a *spatial sampling* of the reflectance function.

Our prints are designed to be viewed from the direction normal to the top surface of the grid of dimples. In the coordinate system shown in Figure 3.4, the normal viewing direction is $v = (0, 0, 1)^T$. When this view direction is substituted into Equation 3.5, the resulting lighting directions, as a function of the dimple point coordinates, are given by:

$$l(s) = \begin{pmatrix} l_x \\ l_y \\ l_z \end{pmatrix} = \begin{pmatrix} 2xz \\ 2yz \\ 2z^2 - 1 \end{pmatrix} \tag{3.6}$$

This equation implies rotational symmetry for the normal viewing direction. Rotating any point on the dimple around the z axis rotates the corresponding lighting direction by the same amount around the z axis. This rotational symmetry means any vertical slice through the origin, for example $x = 0$, has all the information needed for analyzing the system.

3.3.5 Viewpoint Transformation

Our system and prototypes are designed for a fixed view direction normal to the top surface of the grid of dimples. Real world, near field viewing implies that the viewing direction of some dimples will be slightly away from normal. Empirically we find that changing the view direction degrades the experience very gracefully.

This section confirms this empirical finding by analytically showing that a rotation of view direction has the effect of smoothly warping the lighting space function. Using Equation 3.4 we see that a rotation of viewpoint $w = R(\theta)v$, where $R(\theta)$ is the rotation matrix (for example around the x axis), results in a modification of lighting space given by $l_R = MR(\theta)v$. For the normal view direction, the components of lighting direction l are given by Equation 3.6. Applying a rotation around the x axis to the normal view direction results in transformed lighting directions given by:

$$l_R(s) = \begin{pmatrix} 2xz \cdot cos(\theta) + 2xy \cdot sin(\theta) \\ 2yz \cdot cos(\theta) + (2y^2 - 1) \cdot sin(\theta) \\ (2z^2 - 1) \cdot cos(\theta) + 2yz \cdot sin(\theta) \end{pmatrix} \tag{3.7}$$

One can show that for points where $x = 0$, $l_R = R(\theta)l = MR(\theta)v$ so that rotation by θ of the viewpoint results in a rotation in the opposite direction for the lighting space function (for those points where $x = 0$). The off-axis points where $x \neq 0$ result in the smooth transformations of x, y, z shown in Equation 3.7.

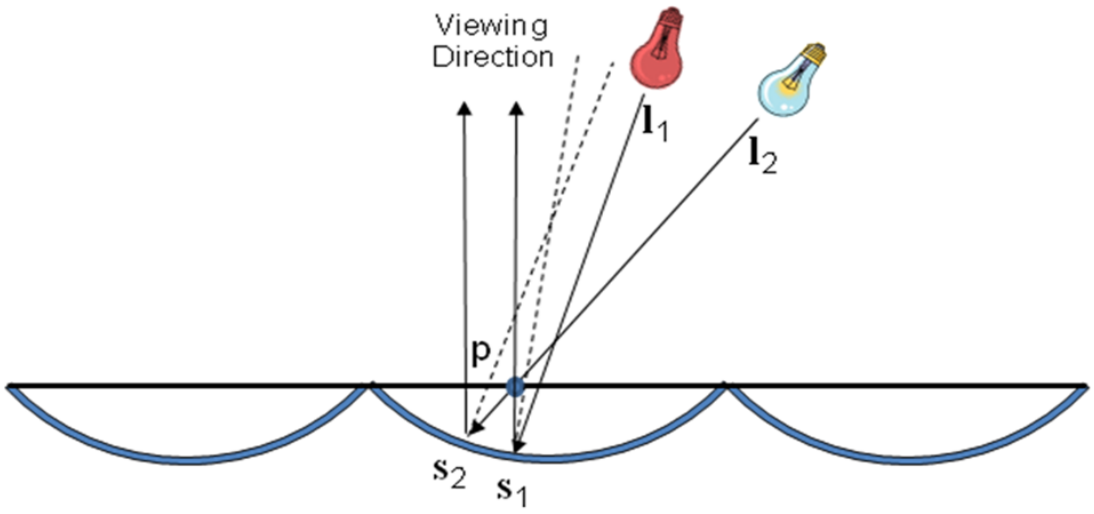

Figure 3.5: A single ink patch on a transparent overlay is effective at attenuating incoming light from direction l_1, but inadvertently also attenuates a single other light direction, that from l_2.

3.3.6 Echo in Lighting Space for Printed Films

When printing on films placed on the top surface of the grid of dimples, we project the reflectance function from the dimple surface onto the transparency. For view direction normal to the transparency film surface, the remapping of reflectance values simply applies the reflectance value of the point on the dimple with the same x and y coordinates as the point on the transparency. Modifying the transparency of a point $p = (x_0, y_0, z_0)$ on the film affects the transmission of light from exactly two points on the dimple, s_1 and s_2, and therefore two lighting directions. By convention, we set the transparency of point p to control the lighting direction labeled l_1 in Figure 3.5, but because of the geometry, this constrains our control over lighting direction labeled l_2, which we call the echo

direction. In the dithering section we will show an adaptive, data dependent algorithm to reduce the effects of this echo constraint.

The two points on the dimple satisfy independent conditions,

$$s_1 = p - k_1 v \tag{3.8}$$

$$s_2 = p - k_2 l_2 \tag{3.9}$$

for constants k_1 and k_2 that determine the intersection points on the dimple. Intuitively, s_1 is found by looking from viewpoint v through point p to the intersecting point s_1 on the dimple. The second point, s_2, is determined indirectly by the condition that the blocked lighting direction l_2 goes from the intersecting point s_2 on the dimple to the point p. Thus a single point on the transparency designed to block lighting direction l_1 also blocks lighting direction l_2, as shown in Figure 3.5. Note that this echo occurs in lighting space, within a dimple, not spatially across dimples. Figure 3.6 plots the angular deviation between the primary and echo directions. For the dimples used in our prototype, which subtend 70 degrees, the maximum angle between the primary and echo lighting angles is 12.9 degrees. We control the echo by designing the dither pattern to take them into account.

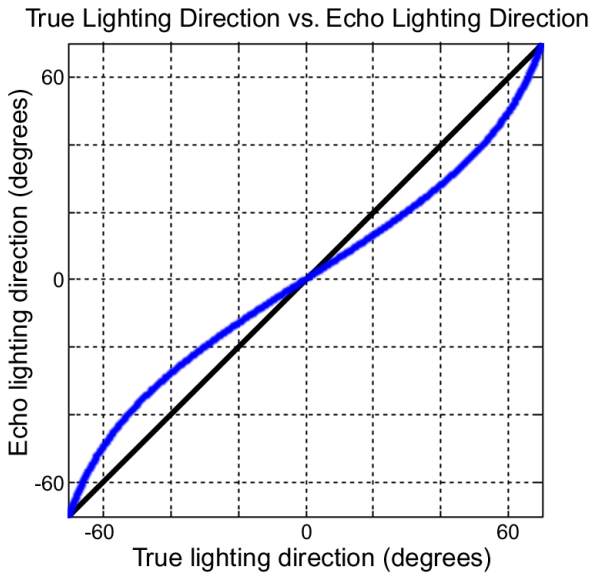

Figure 3.6: The maximum offset between a desired, blocked light ray, l_1 and its echo', l_2 is at worst 12.9 degrees in our roughly 1/5th sphere prototype.

3.3.7 Dithering

Our current method of printing is limited to black ink, with no access to true grayscale inks. Simply thresholding the reflectance function to black and white, as was done in our initial prototype, produces images which relight correctly, but appear metallic, as shown in Figure 3.7. This section describes halftoning techniques for better rendering on the reflectance paper system of Figure 3.3. Halftoning techniques are well developed for traditional printing [7, 85] and have been employed in reflectance functions [54]. The reflectance paper halftoning problem is interesting because the reflectance function has statistics different from those of traditional images and the halftoning may be used to reduce the effects

Figure 3.7: Simply thresholding the grayscale reflectance function (top) to black and white (bottom) , produces images which relight correctly, but with the expected quantized binary response.

of the echo shown in Figure 3.5.

3.3.7.1 Spatial and Lighting Space Dither

Within each dimple in the printed film of Figure 3.3 are samples from different lighting directions, and the set of dimples arranged spatially in a hexagonal grid sample the reflectance field spatially. The film encoding reflectance functions can be halftoned represented by binary reflectance values that preserve local grayscale average values. These are 4D printed functions, however, and care

must be taken in choosing the halftone patterns. Halftoning can occur either spatially across dimples, or in the lighting domain within a single dimple, or jointly. Halftoning spatially uses a single threshold within each dimple, but the thresholds vary between the different dimples (according to a predetermined dither matrix). This results in grayscale reflectance functions when averaged spatially across dimples. When an area light source is present, halftoning in the lighting domain by varying thresholds within a dimple may improve the spatial resolution of the rendered grayscale values by averaging within a single dimple. This only occurs for area light sources. For a point light source, halftoning in the lighting domain within a single dimple has no effect since only a single ray is reflected from each dimple.

We implement spatial halftoning by using dispersed dot dither matrices designed for hexagonal grids. The dither matrices specify spatially varying thresholds that we apply across the dimples. Within a single dimple, a single threshold is used for all lighting angles, but across dimples, the threshold is spatially varying and periodic, as with traditional image halftones. We implemented a small hex mask with 3 threshold values, and a larger one with 27 threshold values [85].

Figure 3.8 compares grayscale, thresholding, and dithering simulated for a single lighting direction. The images are shown at the resolution of our physical prototype, illuminated from a single lighting angle.

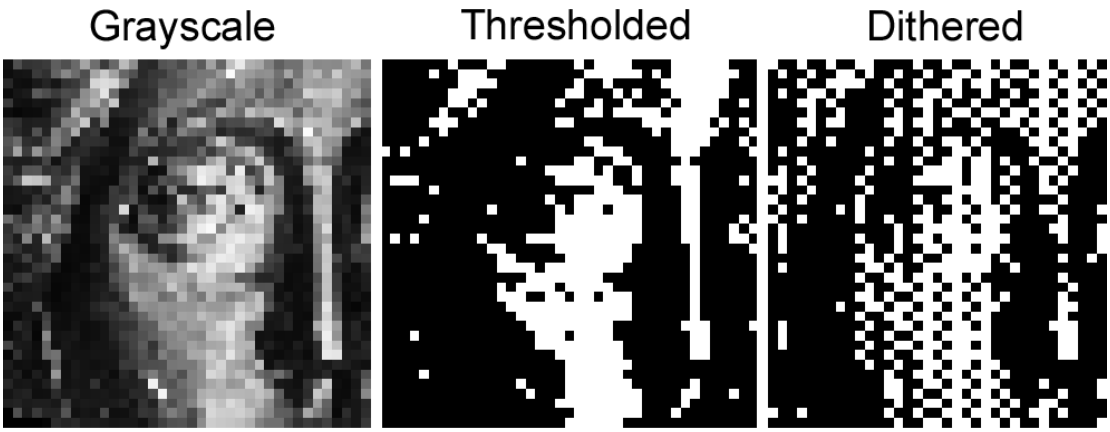

Figure 3.8: Comparison of grayscale, black and white thresholding, and dithering, calculated at the dot pitch of our manufacturing process.

Figure 3.9 illustrates the difference between masks produced through black and white thresholding and spatial dithering. The left image of the figure shows a portion of the original reflectance of a synthetically generated sphere. The center image in the figure shows the results of thresholding and the right image the results of halftoning (yellow borders outline the individual dimples). The hexagonal grid structure is aligned with our underlying dimples, at a 1mm scale. The thresholding shown in the center of the figure does not preserve grayscale because of the fixed threshold values but the halftone version on the right results in grayscale preservation by dithering across space with varying threshold values, producing the observed dot size differences.

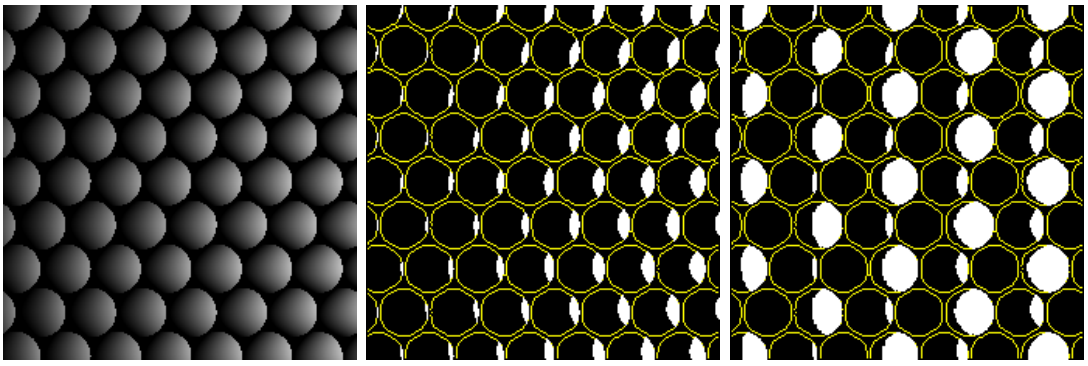

Figure 3.9: Comparison of the attenuation masks produced by b) binary thresholding and c) an ordered dither pattern for a) the original reflectance function.

3.3.7.2 Echo

For the ideal dither function printed directly on the dimples, there is no echo and each lighting direction could independently take values of 0 (opaque) or 1 (transparent). When printing on the transparency film, however, the geometry of Figure 3.5 shows that exactly two lighting directions are controlled by each single point on the printed film. The plot of Figure 3.6 shows the mapping relationship in angle between the two lighting directions controlled by point p, and this mapping determines which lighting directions interact. The plotted function is circularly symmetric so the single vertical slice shown represents the entire function. Whether an error actually occurs due to the echo interaction depends in a simple non-linear manner on whether the lighting function dither matrix values are transparent or opaque.

To reduce the echo' in lighting space, we modify the results of the standard

halftone, which does not take echo into account, in a lighting function adaptive iterative second stage. We approximate the geometry of Figure 3.5 by using point samples, spatially quantized pixel locations on the printing film, and nearest pixel intersection rules for determining which pixels positions interact together for the echo' and proceed to reduce the echo. Section 3.8 provides pseudo-code for the two stage halftoning algorithm that first generates a standard halftone and then reduces the echo in the second iterative stage.

A transparent primary halftone value does not block the echo halftone value, so that no change in the primary halftone value is needed. Only for those interacting pixels where the ideal dithering solution would result in primary lighting direction opaque and echo lighting direction transparent, we examine errors against the ideal first stage halftone values and flip the dither matrix element to transparent when this lowers the error. This process is iterated (very few iterations are needed in practice), and the comparison is to the ground truth dither that would be possible without echo when printing directly on the dimples.

In the binary search (DBS) halftoning of [5], a halftone post-processing step of halftone value flips and swaps are used in conjunction with a printer and vision model to iteratively improve halftone quality. Our approach differs in motivation, we aim to reduce echo, and in details, but at a conceptual level there are some similarities between our approach and DBS.

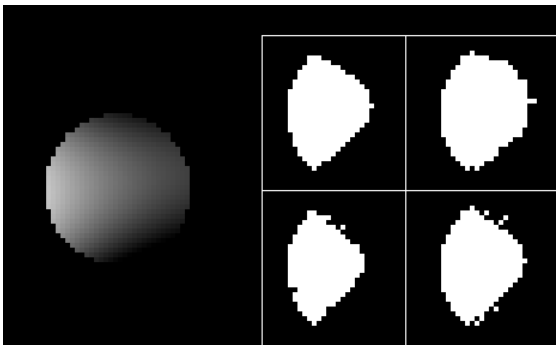

Figure 3.10: Printed masks and simulated output for all lighting directions for a single dimple. a) Standard halftone reflectance function for this dimple. b) Echo compensated halftone, blocking fewer lighting directions; c) Simulation output of halftone pattern in (a). d) Simulation output for halftone pattern in (b). In both (c) and (d) the simulation shows that echoes erode the desired reflectance function. The echo compensated halftone (b) is designed so that (d) approximates (a) more closely than (c) does.

Figure 3.10 and 3.11 illustrate the effect of accounting for the echo interactions in the halftone design. Figure 3.10 maps a single dimple's appearance when lit from each lighting direction, while Figure 3.11 shows an entire image lit from a single direction. Figure 3.10 shows the original grayscale reflectance function, along with printed masks and the simulated outputs. The binary image in 3.10a shows the standard halftone for this dimple, determined by thresholding the grayscale with a fixed threshold (for this dimple). Figure 3.10b shows the echo compensated halftone pattern that results by applying the second stage of the algorithm shown in Section 3.8. In the bottom row of Fig. 3.10 the outputs of simulations that model the echo interactions are shown. Figure 3.10c shows the simulation results when we apply the standard halftone, and Fig 10d shows the simulation results

when echo compensation is applied. Figure 3.10d approximates Figure 3.10a more closely than does Figure 3.10c, with 9 pixels different as opposed to 36.

Figure 3.11 illustrates the effect of accounting for the echo for a fixed single lighting direction for the full image. In the top row, Figure 3.11a shows the standard halftone for this viewpoint, and Figure 3.11b shows the echo compensated halftone for this viewpoint. In the bottom row of Figure 3.11 the outputs of simulations that model the echo interactions are shown. Figure 3.11c shows the simulation results for the standard halftone and Figure 3.11d shows the simulation results for the halftone designed with echo compensation. Figure 3.11d approximates Fig 3.11a much more closely than does Figure 3.11c.

3.4 Simulation

We have built both an interactive viewer which shows the function we intend to print on the transparency, and a full raytraced simulation of the microgeometry. The interactive viewer calculates the amount of light which we expect the viewer to observe as a function of incoming light angle. This is the value obtained by evaluating the reflectance function, and represents ground truth. The viewer allows real time interaction with the incident illumination, as well as switching between color, grayscale, black and white, and several versions of dithering.

The raytraced simulator allows us to verify that the microgeometry behaves

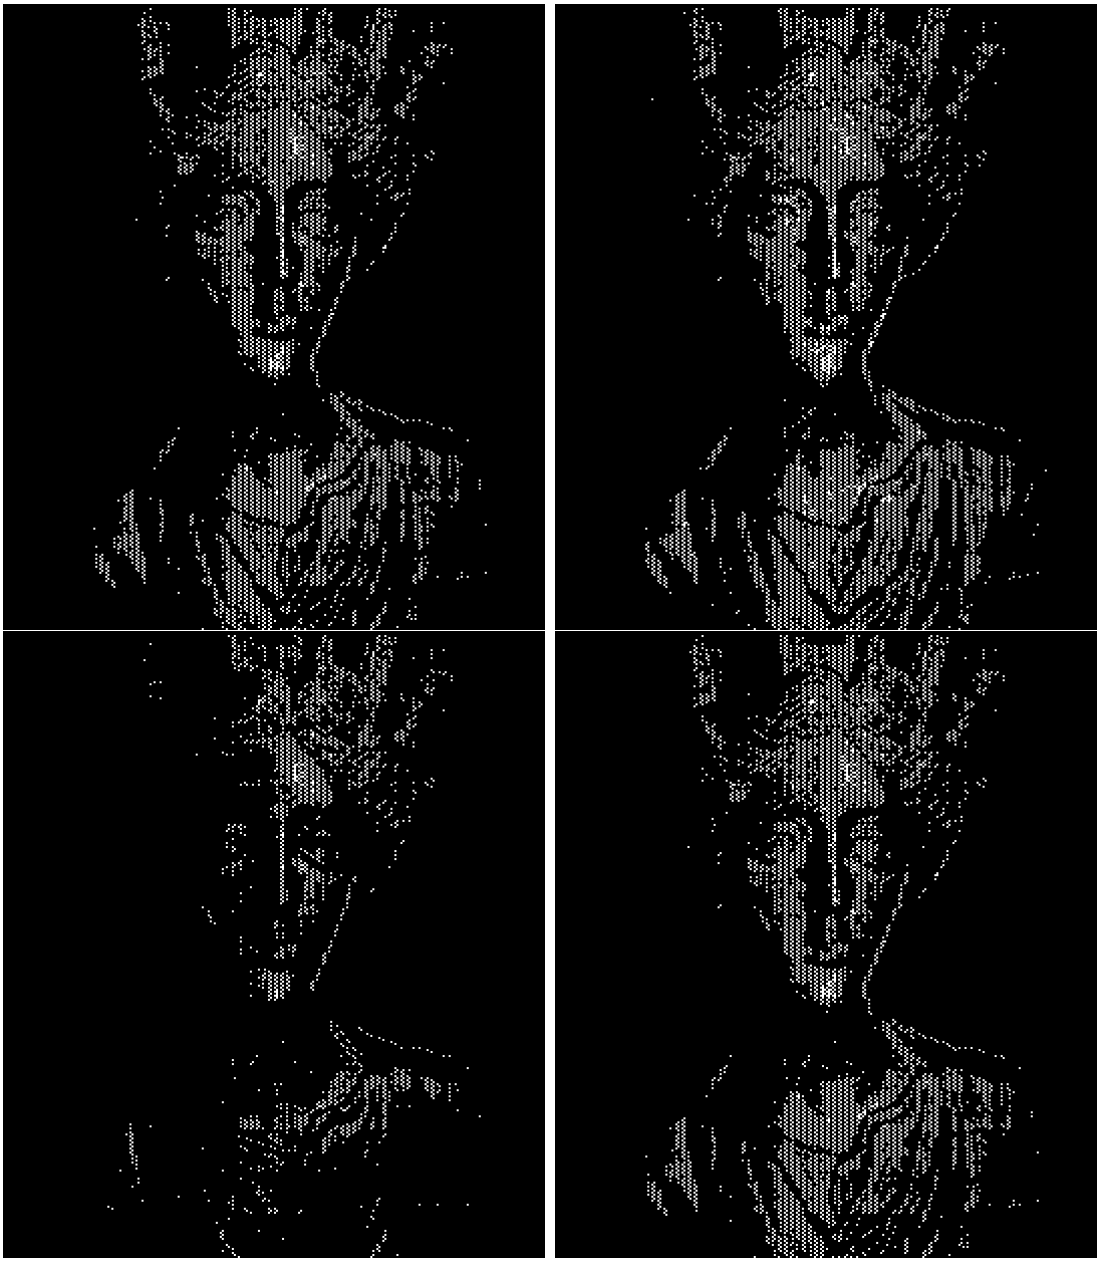

Figure 3.11: Printed masks and simulated outputs for a fixed viewpoint and fixed lighting direction. a) Standard halftone. b) Echo compensated halftone; c) Simulation output of halftone pattern in (a). d) Simulation output for halftone pattern in (b). As in Figure 10, the simulation in (d) approximates (a) more closely than the simulation in (c) does.

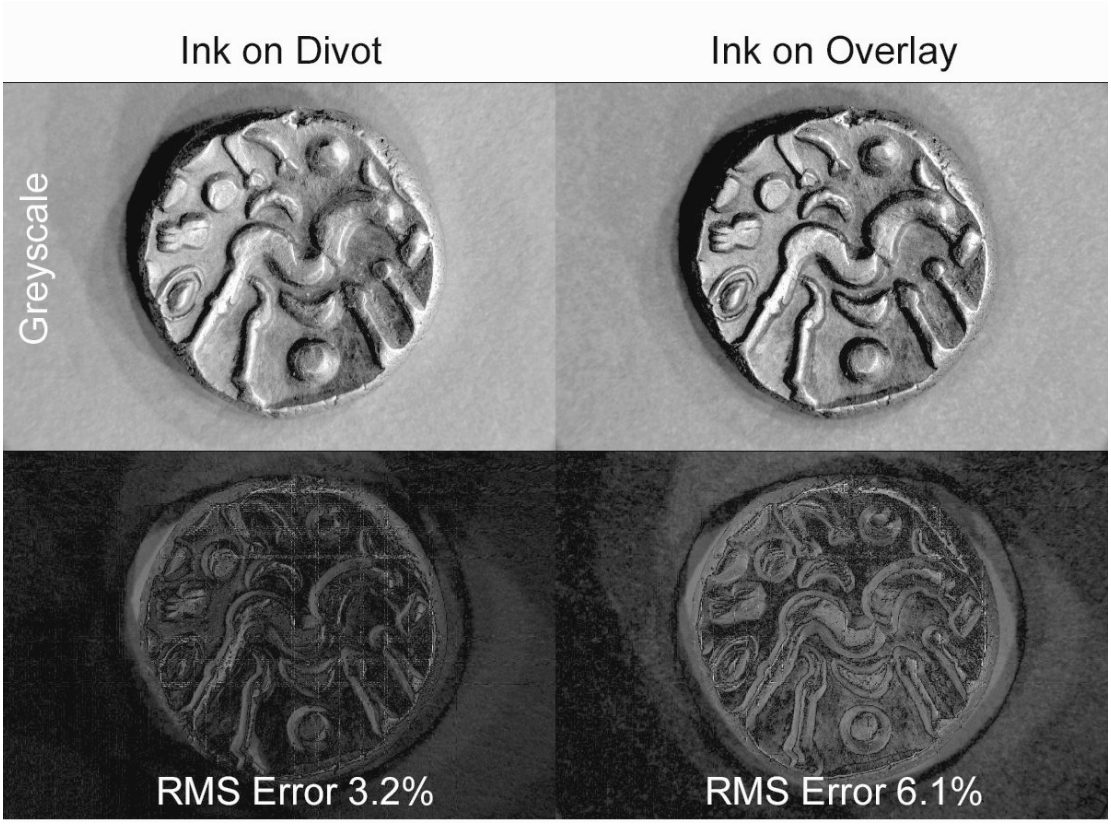

Figure 3.12: Using an overlay mask is expected to introduce some errors. Using a raytraced simulation of our microgeometry we analyze this error, finding that the overlay has only a modest effect, when compared against ground truth. The residual RMS Error of 3.2% is due to using a finite specular lobe width in the simulation. The error images in the bottom row have been contrast enhanced to make them visible when printed. Image used with the permission of Cultural Heritage Imaging.

as intended, and to investigate what affect printing on a transparency mask may have. The simulator uses a 10x10 grid of ink applied to either the dimple or transparency, approximately one dimple per rendered pixel, and 100 supersamples per pixel. We verified that sampling was sufficient by rendering a few comparisons with 2500 supersamples. Black/white dithering is not used, thus simulated ink is grayscale, in which case it ideally attenuates light without introducing unwanted scattering or specular contribution. The color example attenuates light separately in each of the RGB channels. The underlying surface is specular, but not a perfect mirror. The simulation accounts for multiple reflections between the transparency and surface, as well as shadows, although we do not expect these to be relevant, due to the specific design of our dimples.

In order to evaluate the effect of printing on a transparency, as opposed to directly on the dimples, we compare raytraced simulations. Figure 3.12 shows raytraced images. Absolute error images are also shown. Error is evaluated as the RMS difference from the reflectance image we intended to produce. As expected, printing on a transparency has a greater error than printing directly on dimples, however the error is within the bounds predicted by our analysis, and the images themselves appear similar. In order to verify that our design correctly reproduces variable illumination, we simulate a light moving from left to right over the surface. Figure 3.13 shows the intended ground truth effect compared against the simulated

Figure 3.13: To validate our design using an overlay across multiple lighting directions we simulate the effect of moving a light across a 90 degree arc of illumination angles. Five frames from the resulting sequence are shown here. The average RMS error across all frames is 5.0%, and nowhere does the error exceed 6.3%, remaining relatively consistent across the full range of lighting angles.

results from our design, including printing on a transparency layer. The average RMS error over all frames is 5.0%, small enough that we consider this a validation that our design is functional.

3.5 Prototype

We have created several prototype implementation of reflectance paper. We produced two sets of dimpled substrates, one with a hexagonal array of 2.3 mm diameter dimples and another higher resolution set with 1.0 mm hexagonally arranged dimples. In both cases the substrates were produced by using a commercial array of microlenses as a mold master. These substrates were then metalized using silver or aluminum vacuum sputtered on the substrates to produce a mirror-like reflective surface. We validated that the shape of the dimples were spherical cross

sections subtending 70 degrees by using a profilometer, as shown in Figure 3.14. We were successful at registering the reflectance data with the substrate geometry and printing directly on the dimples with an Indigo commercial printer, but found the specularity of the inks far too high to effectively block specular highlights from the dimples themselves, as described in Section 3.3.3.

To avoid artifacts due to ink specularity, we printed our binary reflectance pattern on a transparency using a standard laser printer, which was manually registered with the underlying reflective dimpled substrate. The reflectance function that defines the transparency mask was printed at a resolution of 32x32 samples per dimple, approximately the resolution limit of the printer. The lighting dome that captured the original real world photographs samples far fewer lighting directions, so we used a polynomial approximation to generate a continuous reflectance function prior to re-sampling for printing. Figure 3.15 shows the physical prototype as well as macro views of the dimples and printed transparency.

In order to evaluate whether our prototype faithfully represents the printed

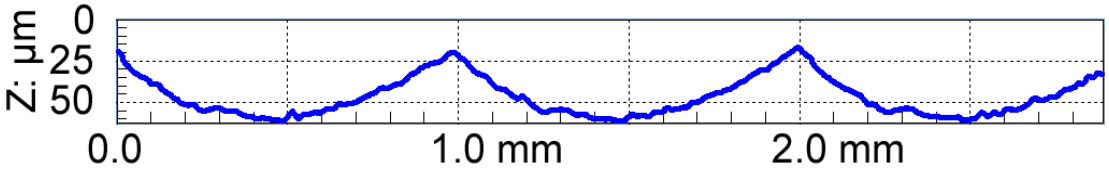

Figure 3.14: Our manufacturing process is validated by using a profilometer to measure a cross section of our mirrored substrate. Dimples have a depth on the order of 50 microns, and deviations are below the resolution of our profilometer.

reflectance function, we illuminated it from several different angles, using a light bulb to approximate a point source. Figure 3.16 shows a thresholded black/white print illuminated from three different directions. The figure also shows the ground truth binary reflectance function from approximately the same lighting configuration. Note that the actual printed reflectance paper has appearance very close to the intended print. Note also that the illumination appears to come from different angles, as expected.

We have explored several methods for binarization of reflectance functions, spanning black and white thresholding, spatial dithering, and dithering in the lighting domain. We printed examples of each, and compared the appearance of the physical prototypes, as shown in Figure 3.17. As expected, spatial dithering obtains graylevels at the cost of spatial resolution and works best with low frequency spatial content. Lighting domain dithering obtains graylevels at the cost of lighting domain resolution, and works best with low frequency lighting. We thus use area light sources to illuminate this prototype.

Figure 3.18 shows a print with spatial dithering illuminated from a point source at many different lighting angles. Note that there is full 2D control over the hemisphere of incident lighting angles. When viewing the physical prototype, the appearance varies smoothly as the light source moves. Representing different objects only requires printing a new transparency and aligning it to the dimpled

substrate. Figure 3.19 shows a shell illuminated from several lighting conditions.

3.6 Limitations

The early physical prototypes shown in this work have many limitations: they lack color, the ink and transparency both have a strong specular component, and the mask and substrate are aligned by hand. We believe that design iterations and better engineering will remedy many of these deficiencies.

In our current implementation, incident illumination from the viewpoint will theoretically generate a specular return from the ink, negating the desired masking effect. Since the light source is rarely extremely close to the viewer, this has not been observed in practice.

Our design is optimized for a single viewing direction. This limitation prevents extreme viewpoint changes. However, in our design, viewpoint changes are approximately equivalent to rotation of the lighting environment, and we have noticed that the appearance is well preserved from nearby positions in practice.

3.7 Discussion

We have developed a method for printing reflectance functions using a standard printer and specially prepared paper. We have produced several physical

prototypes, and validated that they properly display reflectance images. We have also developed a method for producing custom halftones which takes into account our design using a transparent overlay.

Future work on developing less reflective black inks and reducing the size of the microgeometry would be useful. A full-color extension is also possible by having triads of reflective red, green and blue specular divots on the media. Much like colored Christmas ornaments, these would still be mirrored, producing negligible diffuse return, but reflective in only a predetermined color band. The relative proportion of each primary would then be selectively controlled by spatially varying the transmissivity of the ink placed over this medium. Using conventional colored inks overlying a metalized mirrored substrate would be ineffective, as this would introduce unwanted diffuse contributions having the effect of lowering the contrast ratio of our prints.

A more speculative future direction is reversing the geometry so that the lighting angle is fixed and the images change in response to changing viewpoint. In this case we might encode animations or other effects, rather than illumination variation, and we envision such effects being potentially useful on billboards.

```
Algorithm 1
Halftone with Echo Reduction:
1:  for each dimple do
2:      get the constant Threshold value for this dimple
3:      for each Lighting Direction do {generate Standard halftone}
4:          if Reflection > Threshold then
5:              Halftone = Transparent
6:          else
7:              Halftone = Opaque
8:          end if
9:      end for{each Lighting Direction}
10:     for i = 1 to K do {Echo reduction iterations}
11:         for each Lighting Direction do
12:             if Halftone = Opaque then
13:                 Compute quantized Echo Direction
14:                 Retrieve HalftoneEcho value from the Echo Direction
15:                 if HalftoneEcho = Transparent then
16:                     if ErrorTransparent < ErrorOpaque then
17:                         Halftone = Transparent
18:                     end if
19:                 end if
20:             end if
21:         end for{each Lighting Direction}
22:     end for{Echo reduction iterations}
23: end for{each dimple}
```

3.8 Algorithm Details

For completeness we provide pseudo-code for the spatial halftoning algorithm described in Section 3.3.7. This pseudo-code includes the first stage labeled Standard halftone (described in Section 3.3.7.1), and a second stage called Echo Reduction Iteration (described in Section 3.3.7.2). The thresholds used for the dimples are periodic on the hexagonal grid of the reflection paper dimples, but for each

individual dimple there is a constant threshold whose value is determined by the index of the dimple.

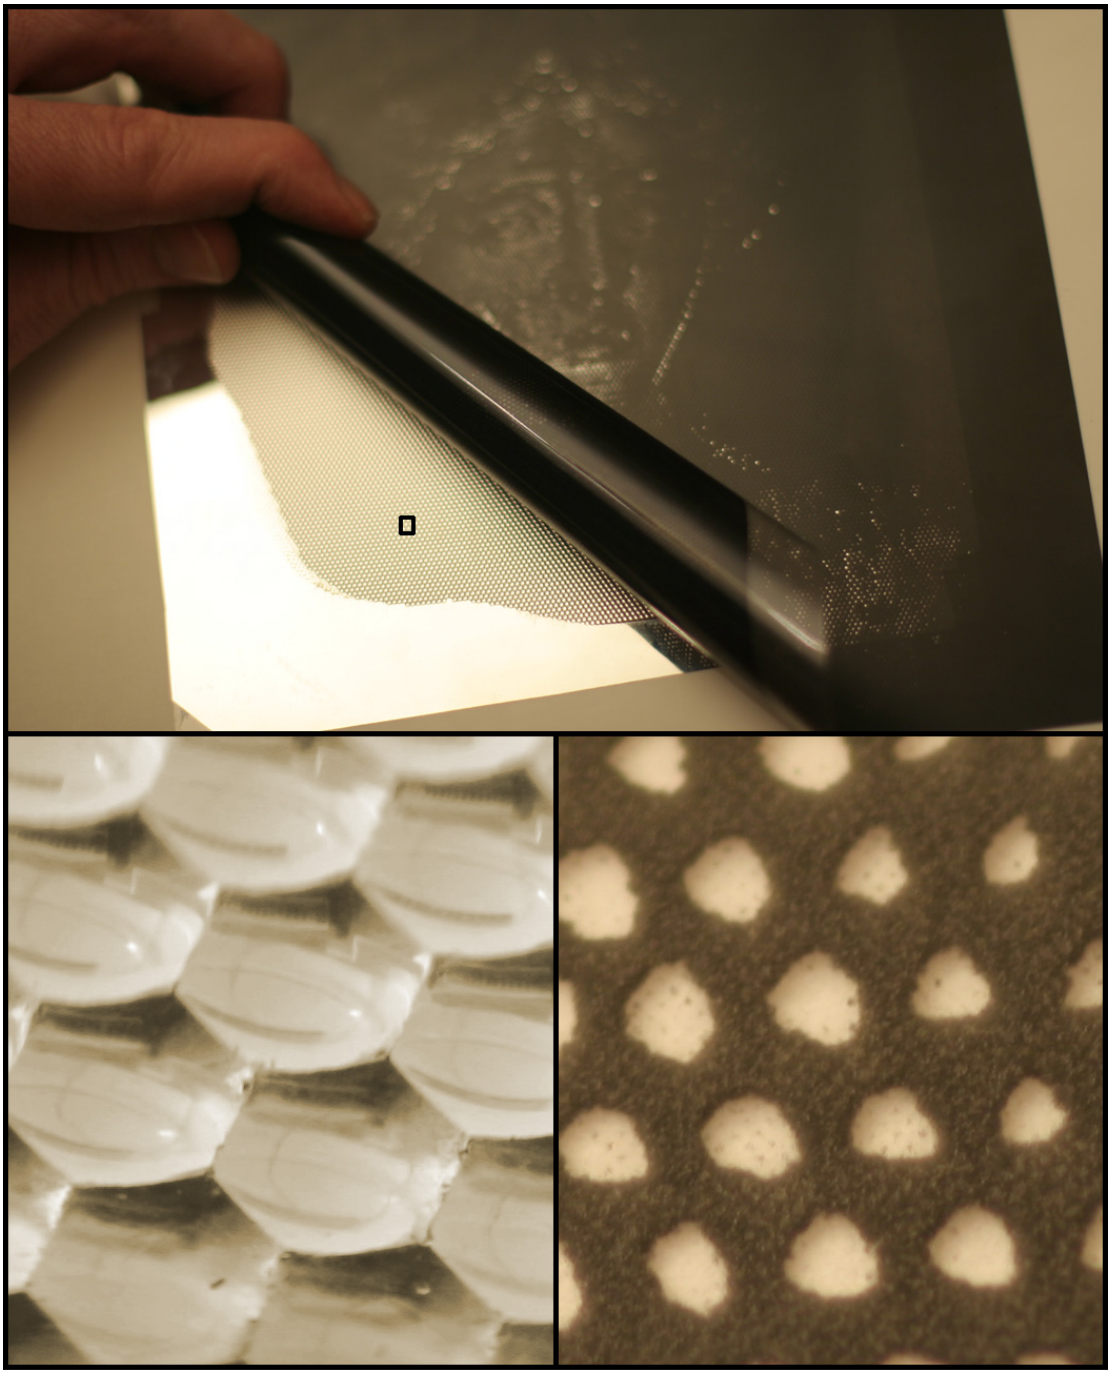

Figure 3.15: (top) Our physical prototype consists of an array of mirrored spherical depressions covered with a mask printed on a transparency. (bottom left) Close up view of the mirrored geometry inside the indicated box. (bottom right) Close up view of the printed mask.

Figure 3.16: Photographs of our physical prototype under three different illumination conditions are compared against the intended black and white print from the same lighting angle.

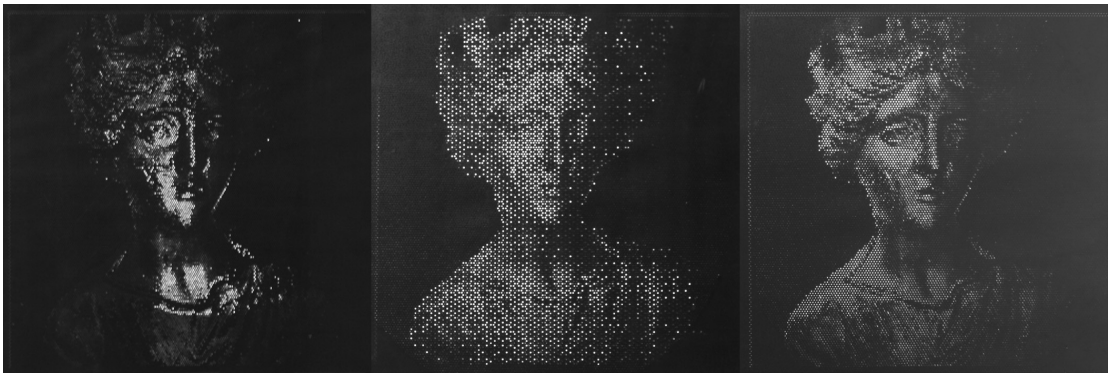

Figure 3.17: We have produced several prototypes. (left) Black and white masks simply threshold the measured reflectance function. (middle) Introducing spatial dithering produces much better grayscale reproduction, at the cost of resolution. (right) Dithering in the lighting domain is capable of producing grayscale at the full spatial resolution of our mask, but requires area lights for best effect.

Figure 3.18: A single print with spatial dithering illuminated from a point source at many different lighting angles. Note that there is full 2D control over the hemisphere of incident lighting angles.

Figure 3.19: A shell printed with spatial dithering, illuminated with several different lighting conditions.

Chapter 4

Higher-Contrast Low-Power Displays

4.1 Introduction

Mobile devices go to great lengths to maximize useful life from limited battery capacity. Since much of this energy is expended by their LCD screens' backlights, saving power by dimming the backlight to the minimum level necessary is common. Unfortunately, a dimmer backlight degrades perceived image quality, through the loss of both brightness and contrast.

It would be desirable to adjust the image displayed to preserve contrast despite the loss of dynamic range, in order to minimize the reduction in image quality.

This is, essentially, a tone mapping problem. However, mobile devices typically have limited computational resources available for image processing, placing severe constraints on viable algorithms. Ideally, the algorithm should run on an FPGA.

We find that significant contrast may be preserved by performing a global change to the device's gamma curve, taking advantage of the software-definable gamma curve present in some mobile display hardware. Our task is then to choose the gamma curve that maximally preserves contrast under backlight dimming, without exceeding available computational resources.

We present a new algorithm for this purpose, 2D Gradient Histogram Equalization. Standard histogram equalization identifies the gamma curve that devotes equal area on-screen to each brightness level. Its core data structure is a histogram of image brightness values. We replace this data structure with a 2D Gradient Histogram, producing a different gamma curve that better preserves contrast at edges in the image.

We are motivated by the observation that removing strong edges from the image is a very undesirable distortion, while creating ambiguity between the relative brightness of nonadjacent regions is a more acceptable distortion. We limit the introduction of ambiguity between pairs of gray levels that form edges, forcing the added ambiguity to occur between pairs of gray levels present in few adjacent pairs of pixels. In essence, since some pairs of gray levels must become more

similar, we choose those pairs that minimally distort the image.

Consider a grayscale image with discrete gray levels $G_1, G_2, ..., G_N$. An image brightness histogram tallies the number of pixels with each gray level G_i. Similarly, a 2D Gradient Histogram counts the number of pixels with gray level G_i that are adjacent to pixels with gray level G_j, filling a 2D NxN table. We project the information in the table onto its off-diagonal in order to generate a gamma curve with we can perform histogram equalization.

The primary contribution of this work is the use of the weighted 2D histogram data structure to identify a gamma curve that preserves contrast under LCD backlight dimming.

The remainder of the chapter is organized as follows: section 4.2 reviews related prior work. Section 4.3 describes Histogram Equalization with the 2D histogram data-structure and section 4.4 presents results.

4.2 Related Work

Backlight dimming for mobile power savings has received extensive attention over the past eight years. Most work, as we do, focuses on choosing the optimal gamma curve that translates desired pixel brightness to LCD transparency. No prior work, however, has optimized the gamma curve to preserve image edges using a gradient histogram.

The simplest option applies a linear transformation to the gamma curve, chosen to maximize some image quality metric [18], [47], [46], [17]. Similarly, the γ parameter of an x^{γ} transformation may be optimized [6].

To avoid color changes, rather than a single gamma curve, some groups apply separate gamma curves to each color channel [10], separate curves to the luminance and chromaticity [20] or transform to an alternate color space [72].

This gamma curve adjustment may be made by changing each pixel's assigned brightness value in software [24], with a dedicated hardware Image-Processing Unit [72], or, for linear adjustments, by setting new LCD reference voltages [22], [19].

In the work most similar to ours, Irlani et. al's "Histogram Equalization for Backlight Scaling" (HEBS) modifies the Programmable LCD Reference Driver to apply more general piecewise-linear gamma curve adjustments [44]. They compute a brightness histogram for the image and choose a desired gamma curve that would best perform Histogram Equalization, then approximate this with a piecewise-linear curve. We approach the problem very similarly, but our work focuses exclusively on choosing the desired gamma curve to preferentially preserve image edges.

Reggiero et al compute a "Multi-Histogram" by forming histograms for each block of a 6x6 block subdivision of the image, in order to better capture the

spatial distribution of brightness values, and use this information to choose a linear transformation of the gamma curve. [72]

While most methods consider only the distribution of image brightness, Tsai et al. use a non-global, content-dependent tone-mapping operation to improve contrast under backlight dimming [84]. Mantiuk et al. perform a piecewise-linear gamma curve adjustment that optimizes a metric of image contrast over a gaussian pyramid [53].

While some papers, like us, address only the choice of gamma curve for a single image, others deal with the changing gamma curve and backlight dimming over time needed to display video. Bartolini et al. find optimal linear transformations for an image sequence [9]. Several time-varying extensions of the HEBS algorithm have also been demonstrated [16], [45]. Hsiu et al. find the optimal time-series of acceptable backlight-dimming-levels with a dynamic programming algorithm, building onto any single-frame backlight-dimming algorithm [42].

Power savings have also been achieved through alternative manipulations of the backlight. Cheng et al. replace the typical backlight's steady glow with a pulse-width modulated signal [21].

Transforming an image's gamma curve in tandem with a changing backlight level has proven useful for applications outside power savings. Wu et al. apply a linear gamma curve adjustment, not to save power, but to compensate for reduced

image brightness when viewing an LCD at a non-orthographic viewing angle, using a camera to determine the viewer's eye position [91]. Mantiuk et al. adjust the gamma curve to obtain a similar image on different displays, or in different lighting environments [53].

4.3 Methods

An LCD displays an image by shining a backlight through an array of semi-transparent cells. In mobile devices, the backlight level is uniform across the display. The backlight brightness B ranges in $0 \leq B \leq 1$. The transparency T of each cell in the LCD is individually controlled, and also lies in the range $0 \leq T \leq 1$.

The luminace $L : 0 \leq L \leq$ emanating from a particular pixel depends upon both the global backlight setting B (which is the same for all pixels) and that pixels transparency T: $L = B * T$

A given luminance L may be achieved at different levels of backlight brightness B by appropriately setting the transparency T. For example, the same luminance will be achieved if the backlight is reduced to half its previous luminance if the transparency of the LCD is doubled. This is limited, however, by the finite range of backlight and transparency settings. Note also that if the backlight is kept constant, transparency directly controls luminance.

The luminance of each pixel in an image is typically not specified directly. Rather, each pixel is assigned one of several (e.g. 64 or 256) discrete gray levels G, and these gray levels are mapped to the available range of transparencies, $T = f(G)$. The mapping is typically nonlinear and is referred to as a "gamma curve" because a common choice is $f(G) = G^{\gamma}$ with e.g. $\gamma = 2.5$.

4.3.1 Histogram Equalization

Histogram equalization is a commonly used image processing technique. It is a global adjustment to the gamma curve that improves local contrast.

A normalized histogram $H_{gray-levels}$ of an image's pixel brightness levels may be formed by simply counting the fraction of pixels set to each gray level $G_1, G_2, ..., G_N$. Some gray levels usually occupy a higher fraction of the image than others; the image may be mostly dark, or mostly light. The histogram is trivially computed in time o(# pixels) and produces an array of length N, with $0 \leq H(x) \leq 1 \; \forall \; x$.

Consider any range of gray levels G_i to G_j which together make up some fraction $F = \sum_{x=i}^{j} H(x)$ of the pixels in the image. One view of histogram equalization's goal is to use a fraction F of the range of available luminances to represent the gray levels between G_i and G_j, for all ranges of gray levels. From another perspective, the difference in luminance between pixels with gray levels G_i and G_j should be approximately F.

This may be accomplished simply by setting the gamma curve $f(G)$, relating gray level to transparency, to be the integral of the histogram $H(x)$:

$$T(G_i) = f(G_i) = \sum_{j=1}^{i} H(G_j) \tag{4.1}$$

Histogram equalization makes the luminances of some pairs of gray levels more different from each other, and the luminances of other pairs of gray levels more similar. Histogram equalization with $H_{gray-levels}$ asserts that gray levels G_i and G_{i+1} should be made more distinguishable if many pixels have brightness G_i, and more generally, that pixels with gray level G_i and G_j should be made distinguisable if many pixels' gray levels fall between G_i and G_j.

In the next two sections, we explore an alternative heuristic: gray levels G_i and G_{i+1} should be made more distinguishable if many pixels have gray level G_i *and are spatially adjacent to pixels with gray level* G_{i+1}, and similarly for the more general case of the distinguishability of G_i and G_j.

4.3.2 The 1D and 2D Gradient Histogram

We form the 1D Gradient Histogram $H_{Gradient-1D}$ as before, but this time counting only pixels with gray level G_i if they have a spatially-adjacent neighbor with gray level G_{i+1}. The histogram is normalized to sum to 1.

We can perform Histogram Equalization on the image using Equation 4.1

with $H_{Gradient-1D}$ rather than $H_{gray-levels}$. Doing so devotes a large fraction of the available luminance range to the sets of gray levels that are spatially adjacent. This has the effect of emphasizing edges in the image.

Calculation of $H_{Gradient-1D}$ requires o(# pixels) computation time. It is roughly as taxing as $H_{gray-levels}$, requiring fewer arithmetic operations but more memory accesses. The entire image need not be accessible in memory; only a buffer three times the length of one row or column of the image is required.

We now form a 2D Gradient Histogram of all gray level pairs, $H_{Gradient-2D}$. Every pair of adjacent pixels are examined, counting how many times a pixel of gray level G_i is spatially adjacent to a pixel of gray level G_j. The histogram forms a 2D table of size NxN and is normalized to sum to 1.

This data structure $H_{Gradient-2D}$ has several notable attributes. Its main diagonal is substantially similar to the brightness histogram $H_{gray-levels}$ for most images, since many adjacent pixels will have the same gray level. We remove the diagonal to remove this influence. Furthermore, its first off-diagonal is exactly the 1D Gradient Histogram $H_{Gradient-1D}$. An example 2D Gradient Histogram is seen in Figure 1.

Also, note that it is symmetric, and thus only half of it need be considered. Thus we henceforth keep only the upper triangular portion, that is, elements $H_{Gradient-2D}(G_i, G_j)$ with $i < j$.

Like the 1D Gradient Histogram, the 2D Gradient Histogram may be found in time linear in the number of pixels, and requires a 3-row or 3-column memory buffer.

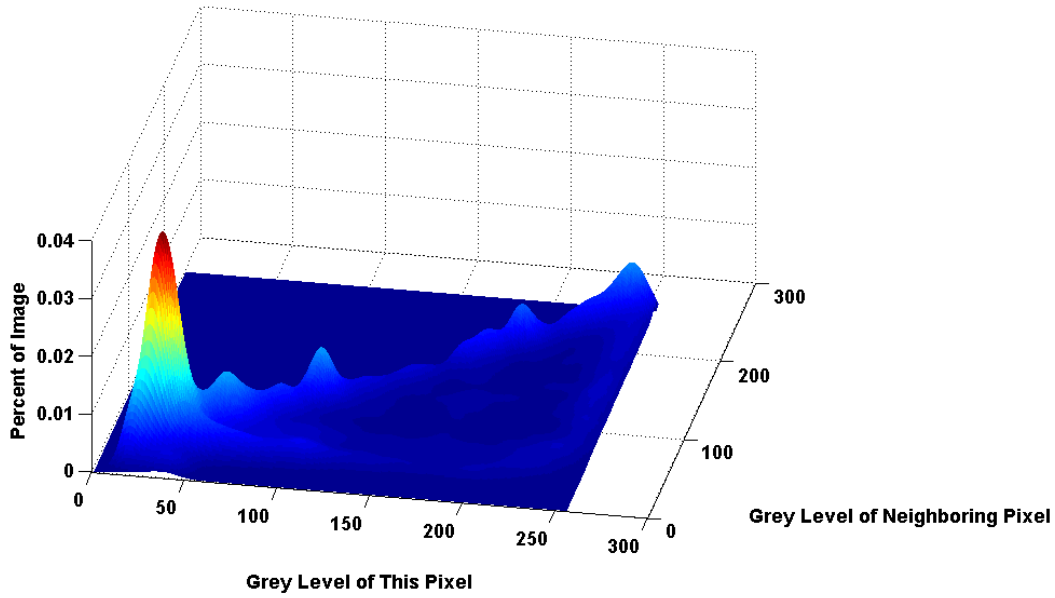

Figure 4.1: An example of a 2D Gradient Histogram.

We wish to extract from $H_{Gradient-2D}$ a 1D histogram for use in histogram equalization with Equation 4.1. The simplest method is to take the off-diagonal, corresponding to the 1D Gradient Histogram, and to ignore the rest of the 2D histogram. This method gains nothing from the additional information available.

An improved means of collapsing the histogram to 1D is to project each element of the 2D histogram onto the off-diagonal. This projection of the element at (G_i, G_j) might, for example, be mapped to the lowest element of the pair, G_i,

or the highest element of the pair, G_j. Alternatively, it could be mapped to the mean of the two, $G_{(i+j)/2}$.

More generally, the contribution could be spread over the off-diagonal between G_i and G_j. These projection methods all preserve edges between spatially adjacent, nonconsecutive gray levels better than merely using the off-diagonal and ignoring the rest. The contribution of a single histogram element at (G_i, G_j) toward enhancing the distinguishability of G_i and G_j is precisely the same for all of them. They differ in their effects on intermediate gray levels between G_i and G_j. Mapping to G_i increases the distinguishability of G_i from G_{i+1}.

Taking inspiration from the bilateral filter, we distribute the value of the element at (G_i, G_j) over the off-diagonal between G_i and G_j unequally, weighted by the previous values of that region of the off-diagonal. In this rich-get-richer approach, when choosing between equally-effective means of enhancing the contrast between G_i and G_j, we choose the means that best preserves other edges in the image.

Various extensions to this algorithm are possible. When calculating the 1D or 2D Gradient Histogram, any desired measure of image saliency may be used to weight the influence of each pixel, in order to cause one region of the image to affect the choice of gamma curve more than another region.

In order to ovecome noise and aliasing, the restriction of *neighbor* to directly-

adjacent pixels may be relaxed, consider neighbors within some finite distance, . A similar effect can be achieved by summing together the 2D Gradient Histograms from multiple scales of a Gaussian Pyramid representation of the image. Noise can also be addressed by applying a gaussian filter to the 2D Gradient Histogram itself after it has been calculated.

4.4 Results

We apply gamma curves derived from Histogram Equalization on several images in Figure 4.2. The first two columns of each row shows the original image at full brightness, and reduced to 50% brightness, with no change to the gamma curve. The third and fourth columns show the image reduced to 50% brightness after histogram equalization with either the Gray-Level Histogram, or the 2D Gradient Histogram. The fifth column shows the gamma curves applied. The red dotted curve is derived from the Gray-Level Histogram, while the blue solid curve is derived from the 2D Gradient Histogram.

Note that, in the first two rows, the 2D Gradient Histogram's result preserves much more detail near the central subjects's faces than the Gray-Level histogram, and more effectively sets them off from the background. This is because the red curve reaches higher than the blue curve at the left edge, raising the brightness of the darkest pixels because there are so many dark pixels. In contrast, the blue

curve is much lower at the left edge, because there are few strong edges between two dark pixels.

In the third and fourth rows, note the difference between the red and blue curves at the high (right) edge. Since there are many very bright pixels in this image, gray-level histogram equalization attempts to reveal the differences between quite-bright and very-bright pixels, whereas the 2D Gradient Histogram Equalization focuses nearly exclusively on pixels of medium brightness, where strong edges occur in the image.

4.5 Limitations

We have focused in this work on the basic problem of choosing a gamma curve, which we apply to all three color channels simultaneously. We formed the 2D Gradient Histogram 4 times - for the red, green, and blue channels separately - and for the luminance channel, and averaged them. This generally performed well, but lead to occasional color shifts in dark regions of the image. More careful treatment of color spaces may be required.

This method is highly susceptible to image compression block-encoding artifacts, e.g. JPEG. We detected and ignored these, lest they be accentuated by our desire to increase the prominence of edges.

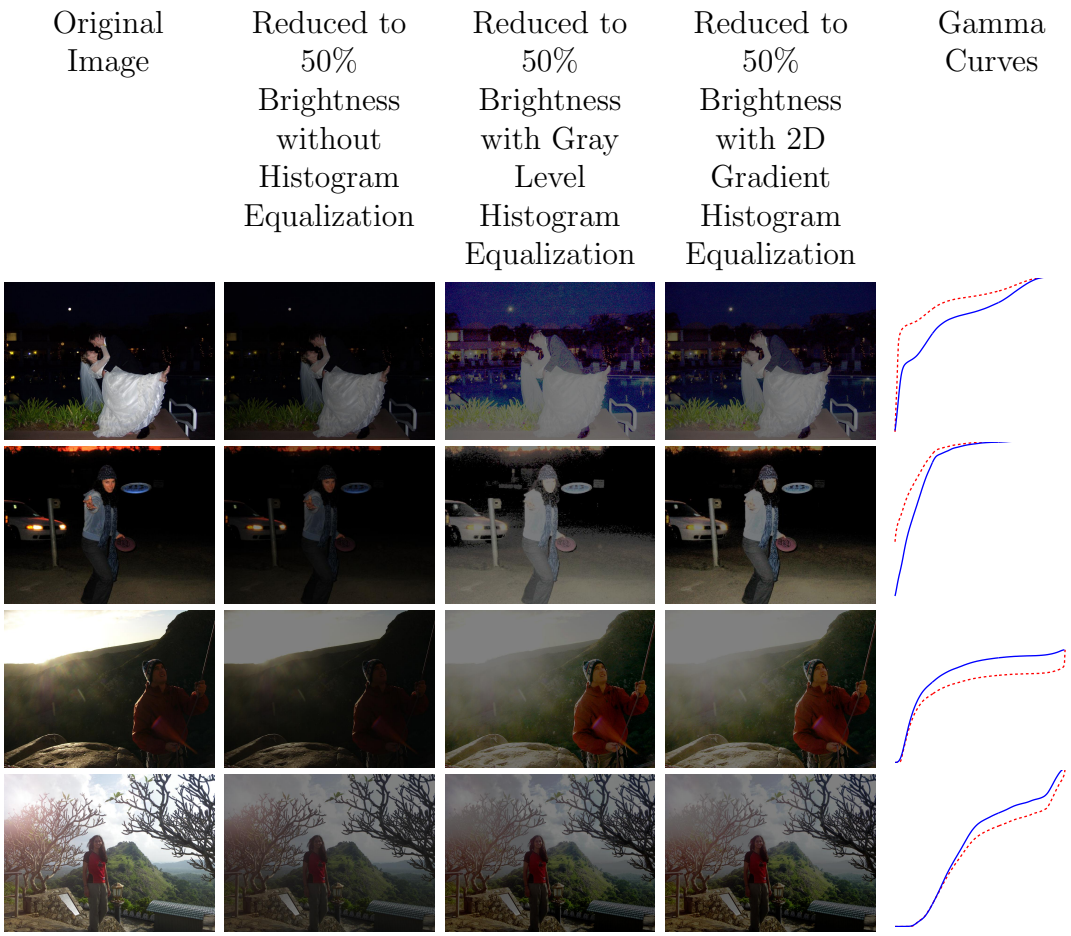

Figure 4.2: Each row shows the original image and the result of applying Histogram Equalization with various histograms. Column 2 shows the image reduced to 50% brightness with no adjustment to the gamma curve, leaving a very dark image. Column 3 uses a standard gray-level histogram, and column 4 uses the 2D Gradient Histogram. Column 5 shows the gamma curves used to perform Histogram Equalization. The dashed red curve results from gray-level Histogram Equalization. The solid blue curve results from 2D Gradient Histogram Equalization.

4.6 Discussion

Mobile devices may save power by dimming their backlights, but must offset the lost dynamic range by remapping the gamma curve to restore lost contrast.

This is typically done through linear scaling of the gamma curve, or histogram equalization based on a Gray Level Histogram. We have presented an alternative gamma curve remapping that performs histogram equalization with a 2D Gradient Histogram, better preserving image edges.

By changing only the global curve, and choosing that curve through a simple 2D Gradient Histogram, this method is easily implementable in mobile devices. This allows substantial power savings with minimal loss to image quality.

Chapter 5

Personalized Photographs

5.1 Introduction

Photos and videos are a powerful medium for capturing a moment's fleeting experience and later sharing it with others. The best photography does not merely faithfully document the scene in front of the camera. Rather, the photographer uses various artifices to influence the viewer's perception of the scene, directing the viewer to notice certain aspects of the image. Choices at the time of image capture set up the photo's framing, exposure, and focus, while further adjustments are made afterward with image editing software. Increasing automation has broadened the base of photographers able to avail themselves of these means of expression.

A casual photographer, while wishing to preserve 'what they noticed,' has historically settled for simply recording an accurate portrait of what is in front of them. Recent research in content-aware image processing has dramatically improved the ability of the amateur photographer to apply software that automatically or semi-automatically modifies their photo to accentuate some region of the photo.

Many such algorithms rely crucially on an estimated saliency map of the image: which regions are important, and which are not? Automatic saliency estimation faces two important challenges. First, determining important and unimportant regions of some photos requires high-level scene analysis beyond current capabilities. Second, objective saliency may be elusive when two photographers disagree as to the salient parts of the same scene. The two photographers may have different motives in taking their pictures, differing knowledge of the semantic content scene, or different relationships to the subjects of the photo.

While objective saliency struggles amidst ambiguity, personal saliency is more tractable. The viewers eyes could be subtly directed to the same parts of the image that the photographer most noticed. This is made possible by recording photographers' eye movements to identify the parts of the scene to which they attend. Images are then manipulated to draw viewers' eyes to those same regions. Photographs of the same object, taken from the same place, with the same camera,

should differ depending on the photographer, and what caught *their* eye. Lacking such a camera, we conduct experiments in a laboratory setting to demonstrate its feasibility and explore various image processing algorithms.

Where automatic saliency algorithms can fail to account for semantic scene content, eye tracking may supply useful, personalized saliency maps. Content-aware image processing algorithms using these saliency maps provide a new means of communicating one's experience.

The primary contribution of this work is our experimental demonstration of eyetrack data's applicability to estimating personalized image saliency for content-aware image processing algorithms that emphasize those parts of the scene that most struck the viewer's eye.

The remainder of this chapter is organized as follows. Section 5.2 reviews prior work in deriving saliency from eye tracks, and in performing content-aware image processing based on saliency. Section 5.3.1 demonstrates the semantic ambiguity that frustrates automatic saliency estimation and motivates personalized, eye tracking based saliency. Section 5.3.2 presents the results of experiments integrating eye tracking based saliency with content aware image processing algorithms, and Section 5.4 discusses future research directions.

5.2 Related Work

Eye tracking has to date recieved little attention with regard to personalized saliency estimation in content aware image processing.

Santella et al [73] created a user interface allowing a computer user to semi-automatically crop an image by recording their eye tracks while using image editing software. Our intended application targets photographers at image capture time, and considers several content aware image procesing techniques rather than cropping.

In another line of research, Santella et. al. [74], [75],[76] strive toward an artistic goal, seeking to automate the creation of stylized cartoons. Conversely, we seek to preserve the appearance of an authentic image while redirecting a new viewer's eye to match the experience of the photographer. Our efforts are similar in that both use eyetracks of individuals to identify regions of interest in images, and use this information to modify the image.

In order to biometrically mark a photograph's author, Blythe et al. [13] embed a small camera within an SLR viewfinder, in order to document the photographer's iris, embedding their identity in a digital watermark. Hua et al. [43] designed a head-mounted augmented-reality display that includes eye tracking. A good survey of additional eye tracking applications is available [31].

We also note that while tomorrow's augmented-reality glasses may feature eye

tracking, embedding eye tracking in cameras is not strictly a technology of the future. Canon included an "Eye Controlled Focus" option in several film-based SLR cameras from 1992 to 2004: an eye-tracker built into the viewfinder directed the camera's autofocus.

Several content aware image processing techniques may be used to direct a viewer's attention in an image. For example, the brightness, contrast, and color saturation may be selectively diminished or enhanced, or the image may be cropped the image to limit the viewer's attention to the areas desired. More flexible tools of recent interest are content-aware resizing algorithms, such as Seam Carving [8] or related methods [70] [78] [71] [90] [77] that selectively enlarge or shrink different regions of the image.

Content-aware resizing has received extensive attention, particularly over the past 5 years. Most work focuses on one of two distinct challenges. First, a saliency map must be constructed to determine which parts of the image should be emphasized, and which de-emphasized or removed. Second, and separately, the image is nonuniformly resampled to remove those image regions deemed least important, leaving the important regions behind. We respond to the first challenge. While typical automatic methods find strong edges or high-frequency content [39],[48], [8], passively-collected eye tracks allow a new answer. What does the photographer want the viewer to see? What the photographer saw.

5.3 Methods

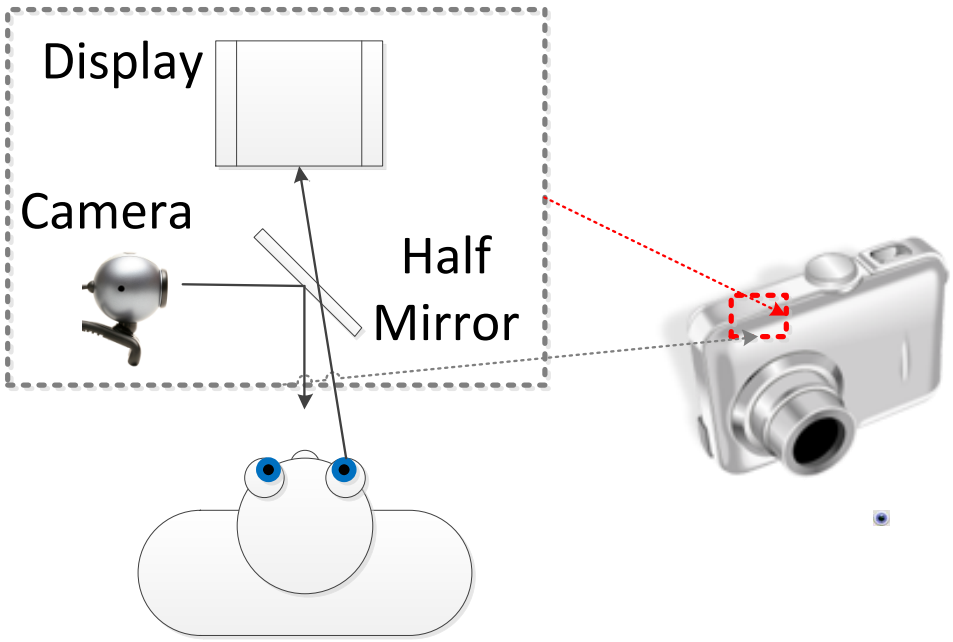

Figure 5.1: Subjects viewed a screen through a beam splitter, so that an eye-tracking camera may monitor their eye movements. This experiment simulates an eye tracker deployed within a camera's viewfinder.

5.3.1 Saliency

Tracking a photographer's eye movements allows the construction of a saliency map indicating the parts of the scene most noticed. Looking ahead, we expect to find future cameras equipped with eye trackers built directly into their viewfinders. In order to conduct experiments investigating the utility of this configuration, we simulate this scenario with off-the-shelf equipment in a laboratory setting.

Rather than a camera's viewfinder, subjects peered through a half-mirror to see a computer monitor while a Bouis infrared eye tracker recorded their eye movements, as in Figure 5.1. The eye tracker contains an infrared light source and a small array of infrared photosensors. Before viewing a photo, the subject viewed a sequence of 25 calibration images consisting of points on a 5x5 grid. This calibration typically provided an accuracy of 20-50 pixels on an 800x600 screen.

We sample eye gaze directions at 1kHz and estimate the average time spent looking at each pixel by convolving with a Gaussian filter. The filter width was chosen to spread the contribution of each measurement over an area matched to the measured accuracy of the gaze-direction estimation while viewing this photo. Santella et al [73] used a more sophisticated methodology to better segment complex objects from their backgrounds, but we have found our simple technique sufficient for the tasks at hand.

We compare the observed saliency maps to two automatic methods. The 'Itti' algorithm [48] begins by applying a filter bank to the image. These filter responses are then normalized and averaged.

The Graph Based Visual Saliency (GBVS) algorithm [39] constructs a fully-connected graph with a node for each pixel, with directed edges weighted according to the dissimilarity between the pixels' responses to filters and their distance. The stationary distribution is obtained through the power method to find 'interesting'

pixels. A new graph is then constructed, also with a node for each pixel, with connections only between neighboring nodes, and weighted by the similarity of their interestingness (as found by the first graph). The power method is again used to find the stationary distribution, concentrating the mass into localized regions. The authors [39] have kindly made available implementations of the GBVS and Itti algorithms.

Figure 5.2 compares the GBVS and Itti algorithms to saliency maps derived from recorded eye tracks. Note that in this case all three subjects' recorded eye tracks focus on the person riding the camel, while both saliency algorithms distributed their attention over a large region of the photo. The visual cues that make the camel's rider so interesting to human viewers are high level semantic cues difficult for any automatic saliency algorithm to identify.

5.3.2 Content Aware Image Processing

Content aware image resizing distorts the sizes of different parts of an image, enlarging or shrinking some more than others in order to emphasize salient regions. Differing saliency maps will emphasize different areas in the resulting image. In the popular *seam carving* algorithm, a subset of pixels in the original image is chosen to appear in the resulting image. To achieve this, the original image is iteratively shrunk by one row or one column. Rather than an intact column, a

seam is removed - a set of pixels that are all diagonally or vertically adjacent, with one pixel from each row. The seam is chosen to preserve the parts of the image weighted highly by the saliency map and remove the parts given low weight.

Attention can also be drawn to one part of an image by selectively defocusing other parts. This effect is commonly used by photographers when capturing photos, by using a shallow depth of field to keep their subject in focus while other objects are out of focus. A similar effect can be approximated after image capture by blurring some parts of the image with a Gaussian filter. We applied a different level of Gaussian blur at each pixel, with the kernel's width smaller for more salient pixels. We find that this subtly de-emphasizes overlooked regions of the image.

We now compare saliency maps from viewers with distinct ideas of what in a scene is salient. In the previous section's Figure 5.2, the three human subjects showed remarkable agreement that the camel rider was the most interesting part of the photo. In contrast, the subject in Figure 5.3(c) attended to each of the fish and a rock, while the subject in Figure 5.3(e) concentrated only on the large blue fish. What is "interesting" varies from person to person. This difference in judged saliency leads to two very different seam carved results. Figure 5.3(d) includes all four fish and regions from the top of the photo, while Figure 5.3(f) centers tightly around the blue fish. The GBVS algorithm's saliency in Figure 5.3(b), meanwhile,

encompasses a large part of the image.

Consider the scene of four ultimate frisbee players in Figure 5.4. While many viewers will find the players more salient than the background, viewers will disagree as to whether some players are more important to the photo than others. To demonstrate the ability of selective defocus to capture the photographer's experience, a subject was asked to look at each of four players in the photo, in turn. Their eye tracks were recorded, giving four separate saliency masks, and four selectively defocused images. Each leaves a different player in focus while the rest of the image is slightly blurred, drawing the viewer's eye toward the player to whom the original viewer directed their attention.

5.4 Discussion

Content-aware image processing provides exciting and useful tools to photographers, and depends crucially on estimating image saliency. We have demonstrated that passively tracking the eyes of photographers would provide personalized saliency maps for use in such algorithms.

(a) Original image

(b) Saliency from eye track of subject 1

(c) Saliency from eye track of subject 2

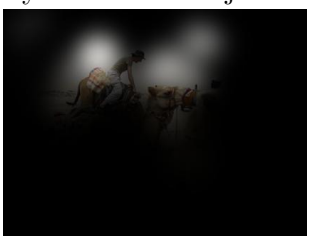

(d) Saliency from eye track of subject 3

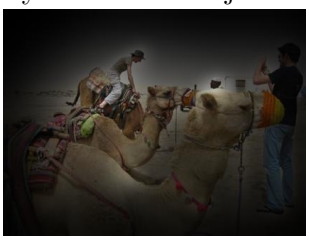

(e) Saliency from Itti automatic algorithm

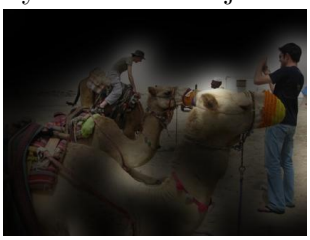

(f) Saliency from GBVS automatic algorithm

Figure 5.2: Saliency of an image is estimated from recorded eye tracks, and from two automatic saliency estimation algorithms. Note that the automatic algorithms find most of the image salient, while all three subjects' eyes concentrate on the camel's rider.

(a) Original image

(b) Saliency from automatic GBVS algorithm

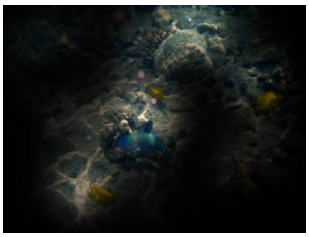

(c) Saliency from eyetracks of subject 3

(d) Content aware resizing based on subject 3

(e) Saliency from eyetracks of subject 4

(f) Content aware resizing based on subject 4

Figure 5.3: Saliency maps derived from eyetracks of two subjects distinctly differ, and the result of content aware resizing thus differs as well. In this case, the automatic saliency algorithm finds most of the image to be salient.

(a) Original image

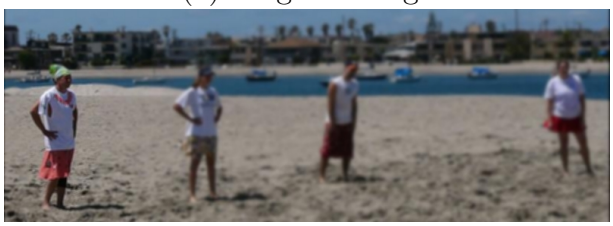

(b) Accentuating player 1 (far leftt)

(c) Accentuating player 2

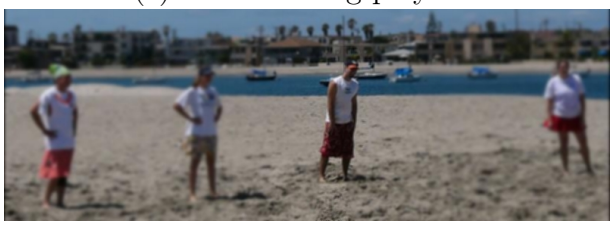

(d) Accentuating player 3

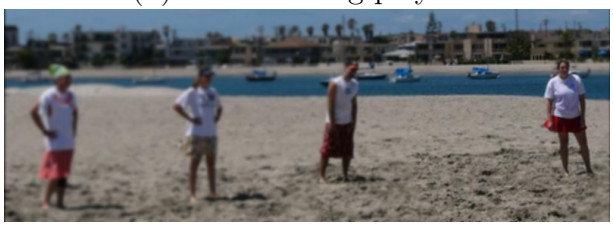

(e) Accentuating player 4 (far right)

Figure 5.4: Viewers may disagree with regard to the salient parts of an image. This image contains four players, any or all of whom may be salient, depending on the viewer. To simulate this, a subject was asked to look at each of the four people in the photo, in turn. Eye movements during each of those glances were recorded separately, and were used to render four different images, each drawing attention to one person by selectively defocusing the non-salient regions.

Chapter 6

Conclusions

Every year, our eyes linger ever-longer over a digitized world. Cellphones and laptops travel with us everywhere; monitors and televisions fill our lives at work and at home; even the unelectrified world is wrapped in digitally printed paper. The state of display technology defines a large part of our experience. The coming decade will see this trend advance, and scientists developing its underlying technologies must play the vanguard for new displays with quality commensurate to their expanding role.

One promising avenue to improve these disparate but pervasive displays is an ever-increasing fidelity of displays to the scenes photographed and displayed on them. Rising resolutions and pixel density are part of this goal, but advances in many other aspects are also needed. In this dissertation, we address this question

in four areas.

First, we have developed a 3D+2DTV, that eases the barrier to usefulness of glasses-based 3D displays. While a traditional 3DTV disturbs those not wearing appropriate stereo glasses, our 3D+2DTV coexists easily in a hybrid scenario where only some viewers wear stereo glasses. Only those with glasses obtain a 3D experience, but those without find the display palatable, unlike a traditional 3DTV.

Second, we introduce a printing process able to produce arbitrary reflectance reflectance functions on a passive print. With our mirrored, dimpled Reflectance Paper with a transparency overlay, a standard printer may print photos that respond to ambient light when viewed to give the appearance of a 3D scene. This vastly larger gamut of reflectance functions allows prints to more faithfully recreate the perception of an photographed or imagined object lying behind the ink.

Third, we have improved the quality of mobile displays. As cellphones become our inseparable companions, we stare deeply into their screens to find our own lives. Largely limited in their ever-shrinking size by the batteries needed to power large screens, cell-phones save power by dimming their screens, with an attendant decrease in image quality. 2D Gradient Histogram Equalization preserves contrast where it's needed most by globally adjusting the display's gamma curve, in a way

that preserves local contrast at salient edges.

Fourth, we explore personalizing photographs to give viewers an experience more closely tied to the photographer than to the camera. Content-aware image manipulation based on a saliency metric derived from the photographer's eyetracks embeds a personalized experience into photos that reflects the person behind the camera, rather than only the scene in front of it.

These efforts constitute our contribution toward a future of improved visual displays that transcend their present shackles. We hope that further work will follow on all four fronts to implement these technologies in commercial applications, as well as theoretical development of their undergirding bases. We look forward to a thriving world of advanced displays.

Bibliography

[1] Maneesh Agrawala, Andrew C. Beers, Ian McDowall, Bernd Fröhlich, Mark Bolas, and Pat Hanrahan. The two-user responsive workbench: support for collaboration through individual views of a shared space. In *Proceedings of SIGGRAPH 1997*, Annual Conference Series, pages 327–332. ACM Press/Addison-Wesley Publishing Co., 1997.

[2] T.S. Aiba and S.S. Stevens. Relation of brightness to duration and luminance under light-and dark-adaptation. *Vision Research*, 4(78):391 – 401, 1964.

[3] M. Alexa and W. Matusik. Reliefs as images. *ACM Trans. Graphics (Proc. ACM Siggraph)*, 2010.

[4] Daniel G. Aliaga, Yu Hong Yeung, Alvin Law, Behzad Sajadi, and Aditi Majumder. Fast high-resolution appearance editing using superimposed projections. *ACM Trans. Graph.*, 31(2):13:1–13:13, April 2012.

[5] J. Allebach. Dbs: retrospective and future directions. *Proc. Color Imag-*

ing: Device Independent Color, Color Hardcopy, and Graphic Arts VI, SPIE, 2000.

[6] Bhojan Anand, Karthik Thirugnanam, Jeena Sebastian, Pravein G. Kannan, Akhihebbal L. Ananda, Mun Choon Chan, and Rajesh Krishna Balan. Adaptive display power management for mobile games. In *Proceedings of the 9th international conference on Mobile systems, applications, and services*, MobiSys '11, pages 57–70, New York, NY, USA, 2011. ACM.

[7] G Arce. *Modern Digital Halftoning*. 2008.

[8] Shai Avidan and Ariel Shamir. Seam carving for content-aware image resizing. *ACM Transactions on Graphics, (Proceedings SIGGRAPH 2007)*, 26(3), 2007.

[9] Andrea Bartolini, Martino Ruggiero, and Luca Benini. Hvs-dbs: human visual system-aware dynamic luminance backlight scaling for video streaming applications. In *Proceedings of the seventh ACM international conference on Embedded software*, EMSOFT '09, pages 21–28, New York, NY, USA, 2009. ACM.

[10] Andrea Bartolini, Martino Ruggiero, and Luca Benini. Visual quality analysis for dynamic backlight scaling in lcd systems. In *Proceedings of the Conference on Design, Automation and Test in Europe*, DATE '09, pages 1428–1433,

3001 Leuven, Belgium, Belgium, 2009. European Design and Automation Association.

[11] Ion Paul Beldie and Bernd Kost. Luminance asymmetry in stereo tv images. *Stereoscopic Displays and Applications II*, (1457):242–247, August 1991.

[12] Oliver Bimber, Daisuke Iwai, Gordon Wetzstein, and Anselm Grundhöfer. The visual computing of projector-camera systems. In *ACM SIGGRAPH 2008 classes*, SIGGRAPH '08, pages 84:1–84:25, New York, NY, USA, 2008. ACM.

[13] Paul Blythe and Jessica Fridrich. Secure digital camera. In *Proceedings of Digital Forensic Research Workshop (DFRWS*, pages 17–19, 2004.

[14] D. Britt. *Modern Art: Impressionism to Post-Modernism*. Thames & Hudson, 2008.

[15] M. Brown, A. Majumder, and R. Yang. Camera-based calibration techniques for seamless multiprojector displays. *IEEE Transactions on Visualization and Computer Graphics*, 11(2):193–206, 2005.

[16] Naehyuck Chang, Inseok Choi, and Hojun Shim. Dls: dynamic backlight luminance scaling of liquid crystal display. *IEEE Trans. Very Large Scale Integr. Syst.*, 12(8):837–846, August 2004.

[17] Liang Cheng, Shivajit Mohapatra, Magda El Zarki, Nikil Dutt, and Nalini Venkatasubramanian. Quality-based backlight optimization for video playback on handheld devices. *Adv. MultiMedia*, 2007(1):4–4, January 2007.

[18] W-C. Cheng and M. Pedram. Transmittance scaling for reducing power dissipation of a backlit tft-lcd.

[19] Wei-Chung Cheng and Chain-Fu Chao. Minimization for led-backlit tft-lcds. In *Proceedings of the 43rd annual Design Automation Conference*, DAC '06, pages 608–611, New York, NY, USA, 2006. ACM.

[20] Wei-Chung Cheng and Chain-Fu Chao. Perception-guided power minimization for color sequential displays. In *Proceedings of the 16th ACM Great Lakes symposium on VLSI*, GLSVLSI '06, pages 290–295, New York, NY, USA, 2006. ACM.

[21] Wei-Chung Cheng, Chih-Fu Hsu, and Chain-Fu Chao. Temporal vision-guided energy minimization for portable displays. In *Proceedings of the 2006 international symposium on Low power electronics and design*, ISLPED '06, pages 89–94, New York, NY, USA, 2006. ACM.

[22] Wei-Chung Cheng and M. Pedram. Power minimization in a backlit tft-lcd display by concurrent brightness and contrast scaling. *Consumer Electronics, IEEE Transactions on*, 50(1):25 – 32, feb 2004.

[23] Stevenson S. B. Cormack, L. K. and C. M. Schor. Interocular correlation, luminance contrast and cyclopean processing. *Vision Research*, 31:2195–2207, 1991.

[24] Radu Cornea, Alex Nicolau, and Nikil Dutt. Software annotations for power optimization on mobile devices. In *Proceedings of the conference on Design, automation and test in Europe: Proceedings*, DATE '06, pages 684–689, 3001 Leuven, Belgium, Belgium, 2006. European Design and Automation Association.

[25] P. Debevec, T. Hawkins, C. Tchou, H. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. *Proceedings of SIGGRAPH 2000, Computer Graphics Proceedings, Annual Conference Series*, 2000.

[26] Charles J.M. Diaper. Pulfrich revisited. *Survey of Ophthalmology*, 41(6):493 – 499, 1997.

[27] Piotr Didyk, Tobias Ritschel, Elmar Eisemann, Karol Myszkowski, and Hans-Peter Seidel. A perceptual model for disparity. *ACM Transactions on Graphics*, 30(4):96:1–96:10, August 2011.

[28] Piotr Didyk, Tobias Ritschel, Elmar Eisemann, Karol Myszkowski, and Hans-Peter Seidel. Apparent stereo: The cornsweet illusion can enhance perceived

depth. In *Human Vision and Electronic Imaging XVII, IS&T/SPIE's Symposium on Electronic Imaging*, pages 1–12, Burlingame, CA, 2012.

[29] Neil Dodgson. Autostereoscopic 3d displays. *IEEE Computer*, 38:31–36, 2005.

[30] P.C. Dodwell, G.S. Harker, and I. Behar. Pulfrich effect with minimal differential adaptation of the eyes. *Vision Research*, 8(11):1431 – 1443, 1968.

[31] Andrew Duchowski. A breadth-first survey of eye-tracking applications. *Behavior Research Methods*, 34:455–470, 2002. 10.3758/BF03195475.

[32] V. Dvorak. ber analoga der personlichen differenz zwischen beiden augen und den netzhautstellen desselben auges. *Prag: Sitzber. d. k. bhm. Gesellsch. d. Wiss*, page 6574, 1872.

[33] T. Freeth, Y. Bitsakis, X. Moussas, J. Seiradakis, Tselikas A., H. Mangou, Zafeiropoulou. M., R. Hadland, D. Bate, A. Ramsey, M. Allen, A. Crawley, P. Hockley, T. Malzbender, D. Gelb, W. Ambrisco, and M. Edmunds. Decoding the ancient greek astronomical calculator known as the antikythera mechanism. *Nature*, 444:587–591, 2006.

[34] M. Fuchs, R. Raskar, H.P. Seidel, and H. Lensch. Towards passive 6d reflectance function displays. *ACM Trans. Graphics (Proc. ACM Siggraph)*, 27, 2008.

[35] Samuel Gateau and Robert Neuman. Stereoscopy, from xy to z. *Short Course Siggraph Asia 2010*, 2010.

[36] M.D. Grossberg, H. Peri, S.K. Nayar, and P.N. Belhumeur. Making one object look like another: Controlling appearance using a projector-camera system. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004.

[37] A. Grundhofer and O. Bimber. Real-time adaptive radiometric compensation. *IEEE Transactions on Visualization and Computer Graphics*, 14(1):97–108, 2008.

[38] Anselm Grundhöfer, Manja Seeger, Ferry Hantsch, and Oliver Bimber. Dynamic adaptation of projected imperceptible codes. In *Proceedings of the 2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality*, ISMAR '07, pages 1–10, Washington, DC, USA, 2007. IEEE Computer Society.

[39] J. Harel, C. Koch, and P. Perona. Graph based visual saliency. *Proceedings of Neural Information Processing Systems (NIPS)*, 2006.

[40] T. Hawkins, J. Cohen, and P. Debevec. A photometric approach to digitizing cultural artifacts. *Proceedings of the 2001 conference on Virtual reality, Archeology, and Cultural Hertiage (VAST 01)*, 2001.

[41] G Heron and G.N. Dutton. The pulfrich phenomenon and its alleviation with a neutral density filter. *Br J Ophthalmol.*, Dec 1989.

[42] Pi-Cheng Hsiu, Chun-Han Lin, and Cheng-Kang Hsieh. Dynamic backlight scaling optimization for mobile streaming applications. In *Proceedings of the 17th IEEE/ACM international symposium on Low-power electronics and design*, ISLPED '11, pages 309–314, Piscataway, NJ, USA, 2011. IEEE Press.

[43] Hong Hua. Integration of eye tracking capability into optical see-through head-mounted displays. volume 4297, pages 496–503. SPIE, 2001.

[44] Ali Iranli, Hanif Fatemi, and Massoud Pedram. Hebs: Histogram equalization for backlight scaling. In *Proceedings of the conference on Design, Automation and Test in Europe - Volume 1*, DATE '05, pages 346–351, Washington, DC, USA, 2005. IEEE Computer Society.

[45] Ali Iranli, Wonbok Lee, and Massoud Pedram. Backlight dimming in power-aware mobile displays. In *Proceedings of the 43rd annual Design Automation Conference*, DAC '06, pages 604–607, New York, NY, USA, 2006. ACM.

[46] Ali Iranli, Wonbok Lee, and Massoud Pedram. Hvs-aware dynamic backlight scaling in tft-lcds. *IEEE Trans. Very Large Scale Integr. Syst.*, 14(10):1103–1116, October 2006.

[47] Ali Iranli and Massoud Pedram. Dtm: dynamic tone mapping for backlight scaling. In *Proceedings of the 42nd annual Design Automation Conference*, DAC '05, pages 612–617, New York, NY, USA, 2005. ACM.

[48] L. Itti, C. Koch, and E. Niebur. A model of saliency based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine*, 1998.

[49] H. Jorke and M. Fritz. Stereo projection using interference filters. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 2006.

[50] Seung-Cheol Kim and Eun-Soo Kim. A new liquid crystal display-based polarized stereoscopic projection method with improved light efficiency. *Optics Communications*, 249(1-3):51–63, 2005.

[51] F Kooi and A Toet. Visual comfort of binocular and 3D displays. *Displays*, 25(2-3):99–108, 2004.

[52] T. Malzbender, D. Gelb, and H. Wolters. Polynomial texture maps. *Proceedings of Siggraph 2001, Computer Graphics Proceedings, Annual Conference Series*, pages 519–528, 2001.

[53] RafałMantiuk, Scott Daly, and Louis Kerofsky. Display adaptive tone mapping. *ACM Trans. Graph.*, 27(3):68:1–68:10, August 2008.

[54] W. Matusik, B. Ajdin, J. Gu, J. Lawrence, H. Lensch, F. Pellacini, and S. Rusinkiewicz. Printing spatially-varying reflectance. *ACM Trans. Graphics (Proc. ACM Siggraph)*, 28(5), 2009.

[55] Wojciech Matusik and Hanspeter Pfister. 3d tv: a scalable system for real-time acquisition, transmission, and autostereoscopic display of dynamic scenes. *ACM Transactions on Graphics*, 23(3):814–824, August 2004.

[56] Ian E. McDowall, Mark T. Bolas, Dan Corr, and Terry C. Schmidt. Single and multiple viewer stereo with DLP projectors. *Proc. SPIE 4297*, 2001.

[57] M.J. Morgan and P. Thompson. Apparent motion and the pulfrich effect. *Perception*, 4(1):3–18, 1975.

[58] M. Mudge, T. Malzbender, C. Schroer, and M. Lum. New reflection transformation imaging methods for rock art and multiple-viewpoint display. *7th International Symposium on Virtual Reality, Archeology and Cultural Heritage (VAST 2006)*, 2006.

[59] S. Nayar, P. Belhumeur, and T. Boult. Lighting sensistive display. *ACM Trans. Graphics (Proc. ACM Siggraph)*, 23(4), 2004.

[60] F. Nicodemus, J. Richmond, and J. Hsai. Geometrical considerations and nomenclature for reflectance. *U.S. Dept. of Commerce, National Bureau of Standards*, 1977.

[61] P. Peers, T. Hawkins, and P. Debevec. A reflective light stage. *ICT (USC Institute for Creative Technologies) Technical Report, ICT-TR-04.2006*, 2006.

[62] P. Peers, N. Tamura, W. Matusik, and P. Debevec. Post-production facial performance relighting using reflectance transfer. *ACM Trans. Graphics (Proc. ACM Siggraph)*, 26(3), 2007.

[63] Ken Perlin, Salvatore Paxia, and Joel S. Kollin. An autostereoscopic display. In *Proceedings of SIGGRAPH 2000*, Annual Conference Series, pages 319–326. ACM Press/Addison-Wesley Publishing Co., 2000.

[64] Joan Pollack. Reaction time to different wavelengths at various luminances. *Attention, Perception, & Psychophysics*, 3:17–24, 1968. 10.3758/BF03212706.

[65] Carl Pulfrich. Die stereoskopie im dienste der isochromen und heterochromen photometrie. *Die Naturwissenschaften*, 1922.

[66] R. Ramamoorthi and P. Hanrahan. An efficient representation for irradiance environment maps. *In Proceedings of SIGGRAPH 2001, Computer Graphics Proceedings, Annual Conference Series.*, 2001.

[67] Jerome Ramstad, Monte. Patent WO 2011031326. Interference filters for viewing anaglyphs, 2011.

[68] Ramesh Raskar, Greg Welch, Matt Cutts, Adam Lake, Lev Stesin, and Henry Fuchs. The office of the future: a unified approach to image-based modeling and spatially immersive displays. In *Proceedings of Siggraph 98*, Annual Conference Series, pages 179–188, 1998.

[69] C. Regg, S. Rusinkiewicz, W. Matusik, and M. Gross. Computational high-light holography. *ACM Transactions on Graphics (Proc. Siggraph Asia)*, 29(6), 2010.

[70] Michael Rubinstein, Ariel Shamir, and Shai Avidan. Improved seam carving for video retargeting. *ACM Transactions on Graphics, (Proceedings SIGGRAPH 2008)*, 27(3), 2008.

[71] Michael Rubinstein, Ariel Shamir, and Shai Avidan. Multi-operator media retargeting. *ACM Transactions on Graphics, (Proceedings SIGGRAPH 2009)*, 28(3), 2009.

[72] Martino Ruggiero, Andrea Bartolini, and Luca Benini. Dbs4video: dynamic luminance backlight scaling based on multi-histogram frame characterization for video streaming application. In *Proceedings of the 8th ACM international conference on Embedded software*, EMSOFT '08, pages 109–118, New York, NY, USA, 2008. ACM.

[73] Anthony Santella, Maneesh Agrawala, Doug Decarlo, David Salesin, and

Michael Cohen. Gaze-based interaction for semi-automatic photo cropping. In *CHI '06: Proceedings of the SIGCHI confernce on Human Factors in computing systems*, pages 771–780, 2006.

[74] Anthony Santella and Doug DeCarlo. Abstracted painterly renderings using eye-tracking data. *Non-Photorealistic Animation and Rendering 2002)*, pages 75–82, 2002.

[75] Anthony Santella and Doug DeCarlo. Stylization and abstraction of photographs. *ACM Transactions on Graphics, (Proceedings SIGGRAPH 2002)*, pages 769–776, 2002.

[76] Anthony Santella and Doug DeCarlo. Visual interest and npr an evaluation and manifesto. *Non-Photorealistic Animation and Rendering 2004*, pages 71–78, 2004.

[77] Vidya Setlur, Ramesh Raskar, Saeko Takagi, Michael Gleicher, and Bruce Gooch. Automatic image retargeting. In *Mobile and Ubiquitous Multimedia (MUM), ACM.* Press, 2005.

[78] Ariel Shamir and Shai Avidan. Seam carving for media retargeting. *Commun. ACM*, 52(1):77–85, 2009.

[79] M Siegel and S Nagata. Just enough reality: comfortable 3-D viewing via

microstereopsis. *IEEE Transactions on Circuits and Systems for Video Technology*, 10(3):387–396, 2000.

[80] Hansen Per Skafte Sorensen Nils Lykke Sorensen, Svend Erik Borre. Patent US 6687003. Method for recording and viewing stereoscopic images in color using multichrome filters, 2004.

[81] L. Standing, P. Dodwell, and D. Lang. Dark adaptation and the pulfrich effect. *Attention, Perception, & Psychophysics*, 4:118–120, 1968. 10.3758/BF03209521.

[82] J. C. Stevens and S. S. Stevens. Brightness function : Effects of adaptation. *J. Opt. Soc. Am.*, 53(3):375–385, Mar 1963.

[83] Eric A. Taub. Still thinking outside the box. In *The New York Times*. July 18 2002, New York, NY, USA, 2002.

[84] Pei-Shan Tsai, Chia-Kai Liang, Tai-Hsiang Huang, and Homer H. Chen. Image enhancement for backlight-scaled tft-lcd displays. *IEEE Trans. Cir. and Sys. for Video Technol.*, 19(4):574–583, April 2009.

[85] R. Ulichney. *Digital Halftoning*. MIT Press, 1987.

[86] By Anthony Vetro, Thomas Wiegand, and Gary J Sullivan. Overview of the

stereo and multiview video coding extensions of the h . 264 / mpeg-4 avc standard. *Proceedings of the IEEE*, 99(4):626–642, 2011.

[87] A. Wenger, Gardner A., Tchou C., Unger J., Hawkins T., and Debevec P. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Trans. Graphics*, 24(3), 2005.

[88] T. Weyrich, J. Deng, C. Barnes, S. Rusinkiewicz, and A. Finkelstein. Digital bas-relief from 3d scenes. *ACM Trans. Graphics (Proc. ACM Siggraph)*, 26(5), 2007.

[89] T. Weyrich, P. Peers, W. Matusik, and S. Rusinkiewicz. Fabricating microgeometry for custom surface reflectance. *ACM Trans. Graphics (Proc. ACM Siggraph)*, 28(5), 2009.

[90] Lior Wolf, Moshe Guttmann, and Daniel Cohen-Or. Non-homogeneous content-driven video-retargeting. In *Proceedings of the Eleventh IEEE International Conference on Computer Vision (ICCV-07)*, 2007.

[91] Chih-Nan Wu and Wei-Chung Cheng. Viewing direction-aware backlight scaling. In *Proceedings of the 17th ACM Great Lakes symposium on VLSI*, GLSVLSI '07, pages 281–286, New York, NY, USA, 2007. ACM.