Perceptual Considerations for Displays under Varying Ambient Illumination

by

Allan G. Rempel

M.Sc., The University of British Columbia, 1997
B.Sc. (Great Distinction), University of Saskatchewan, 1993
B.Comm. (Great Distinction), University of Saskatchewan, 1993

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

The Faculty of Graduate Studies

(Computer Science)

The University Of British Columbia
(Vancouver)

January, 2012

Abstract

The capabilities of visual display devices such as television sets and computer monitors have advanced dramatically over the past decade. In particular, prototype high dynamic range (HDR) displays have emerged during that time, which provide significantly greater output brightness and dynamic range than conventional currently available displays. These increased capabilities open up two new avenues of research: investigating how to take advantage of the characteristics of these displays to improve the appearance of images and video, and observing and characterizing the responses of the human visual system to the new kinds of imagery these displays are capable of producing.

The contributions of this dissertation advance our understanding along both of these lines. We present a set of experiments in which subjects viewed video content on HDR displays, and observe that subjects preferred higher display brightness settings in higher ambient illumination environments, and that subjects generally preferred the highest contrast available, irrespective of brightness, with effectively no visual fatigue (compared to conventional displays) during extended viewing sessions. We then present a novel technique for enhancing the contrast of low dynamic range (LDR) legacy content to take advantage of the capabilities of HDR displays, which is an underconstrained problem with several solutions of varying effectiveness. Our technique boosts the brightness of saturated image regions while keeping darker image regions dark. Next, we present a set of experiments in which we use a prototype display with precise colour control to determine the effect of different colours of glare on vision in low-light environments. The data from these experiments provide insight into human vision in low-light environments which can improve the design of displays for such environments. Finally, we present a set of experiments showing the effect of contrast on depth perception in monocular viewing environments, and how

this effect can be enhanced by adjusting the contrast between highlights and mid-tones in HDR imagery.

Collectively, these contributions advance our ability to display imagery on new generations of high-contrast and high-brightness displays for more satisfying viewing experiences, and can aid in the design of future displays.

Preface

This dissertation includes work that I have previously published with various co-authors.

The work on video viewing preferences in Chapter 3 is a minor revision of the work published as an article[1] entitled "Video Viewing Preferences for HDR Displays Under Varying Ambient Illumination" in the Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization (APGV) published by ACM. I was the primary author of that work and designed and conducted all the experiments. I also presented the work at APGV 2009 [Remp 09].

The work on dynamic range expansion in Chapter 4 is a minor revision of the work published as an article[2] entitled "Ldr2Hdr: On-the-fly Reverse Tone Mapping of Legacy Video and Photographs" in the journal *Transactions on Graphics* published by ACM. Several people including myself contributed to the ideas behind this work, as indicated by the seven authors on the publication, but I developed all the code that was eventually used for the final published algorithm and produced all the images. I also presented this work at the ACM SIGGRAPH 2007 conference [Remp 07].

The description of the first two experiments on vision in low-light environments in Chapter 5 is a minor revision of the work published as an article[3] entitled "Display Considerations for Night and Low-Illumination Viewing" in the Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization (APGV) published by ACM. The idea behind this work originated with Rafał Mantiuk; however I contributed to the design of the work from its inception, conducted the majority of the experiments, and presented the work at APGV 2009 [Mant 09]. The description

[1] http://doi.acm.org/10.1145/1620993.1621004

[2] http://doi.acm.org/10.1145/1275808.1276426, http://doi.acm.org/10.1145/1276377.1276426

[3] http://doi.acm.org/10.1145/1620993.1621005

of the third experiment in Chapter 5 is a minor revision of the work published as an article[4] entitled "Display Considerations for Improved Night Vision Performance" in the Proceedings of the 19th Color Imaging Conference [Remp 11c].

The work on depth perception in Chapter 6 was previously published by UBC as a technical report entitled "The Role of Contrast in the Perceived Depth of Monocular Imagery." This work was my idea, and I designed and conducted all the experiments and wrote the paper [Remp 11a] and a poster which I presented at APGV 2011 [Remp 11b].

All images in Chapters 3, 4, and 5 are reprinted with the permission of ACM, except Figure 5.10 which is in the public domain.

Ethical approval for the experiments described in this dissertation was obtained from the UBC Office of Research Services under certificate numbers H07-02516 and H09-02278.

[4] http://www.imaging.org/IST/store/epub.cfm?abstrid=45157

Table of Contents

List of Tables

List of Figures

Acknowledgements

First, I must of course thank my supervisor, Wolfgang Heidrich, without whom none of this work would have been possible. Thanks also to my supervisory committee members, Kelly Booth and David Lowe, for their valuable insight and feedback on my work in progress, and to my examination committee members Michiel van de Panne and Boris Stoeber and external examiner James Ferwerda, all of whose comments helped me to improve the final dissertation.

Academic colleagues are an enormously valuable resource, and so for countless fruitful discussions, support, and friendship, I would like to thank my colleagues in the lab: Abhijeet Ghosh, Derek Bradley, Matt Trentacoste, Brad Atcheson, Cheryl Lau, Gordon Wetzstein, Nasa Rouf, Anika Mahmud, Ivo Ihrke, and Lukas Ahrenberg. I would especially like to thank Rafał Mantiuk for his close collaboration, and for sharing his knowledge and experience which dramatically helped to improve the quality of my work. Thanks also to Leah Findlater and Karyn Moffatt for their advice and experience in statistical processing and ethical review procedures. Special thanks to Deb Wilson for helping me learn to navigate the vast network of resources for which the library is the gateway; libraries continue to suffer funding pressure, and while the internet may be valuable, professional academics still sometimes need the professional-grade tools in libraries rather than the merely consumer-grade tools on the internet. And thanks to countless other fellow grads for walking with me on this journey.

I would like to thank my colleagues at BrightSide Technologies and Dolby Canada for their tremendous help in working with the prototype HDR displays that they developed, for teaching me to use their light measurement devices, and for providing me with a practical context in which to apply the theory behind this research. Special thanks to Helge Seetzen, Jason Harrison, and Peter Longhurst, without whose participation in the NSERC IPS program that collaboration may not

have happened. Thanks also to Thomas Wan, Henry Ip, Damir Wallener, Dave Tan, Neil McPhail, Gerwin Damberg, Vincent Kwong, Hiroe Li, Anders Ballestad, Anne Webster, and the rest of the folks there.

Thanks to all my friends for their support and encouragement throughout this process. Special thanks to Dawna Tong, Cheryl Enns, Heidi Epp, Annie Ding, Mirela Andronescu, Jennifer Stewart, and my fellow radical reformers at Peace Mennonite Church.

In any list of acknowledgements where individuals are named, especially for a process that lasts as long as a PhD program, and despite best efforts to remember all those who helped along the way, there is always the possibility of having forgotten someone important, and if I have done so, I sincerely apologize.

Thanks to my family, Stephen and Heidi Nighswander-Rempel, Bret and Angela RempelEwert, Otto and Doris Rempel-Dirks, Reg and Donna Sandberg, Mike and Ang Sandberg, Pete and Tanya Sandberg, and the many cousins, aunts, uncles, nieces, and nephews who have provided so much support and encouragement throughout this process.

Finally I would like to dedicate this work to my new family: Stephanie, Grayson, and any other future additions that may occur. Thank you Stephanie for all your love and support and encouragement, and I look forward to the adventure that awaits us.

Chapter 1

Introduction

The capabilities of visual display devices such as television sets and computer monitors have advanced dramatically over the past decade. These advances include increases in brightness and contrast, resolution, and the availability of 3D displays that take advantage of binocular disparity, sending different images to the two eyes through the use of glasses which filter the images being sent to each eye. During this time, displays have made increasing use of light-emitting diodes (LEDs), which have become more practical and less expensive, and which are capable of producing higher luminance more efficiently than the cold-cathode fluorescent lighting (CCFL) that has traditionally been used in liquid-crystal display (LCD) television sets and monitors. Most commercially available displays that use LEDs as a backlight for the LCD panel use the LEDs as a simple replacement for the CCFL layer and provide illumination that is as uniform as possible to the LCD layer. However, a new generation of high dynamic range (HDR) displays is emerging, which uses more complex methods to individually control the LEDs to provide significantly greater output luminance and dynamic range than has previously been available on conventional displays, which brings us closer to being able to meet the capabilities of the human visual system.

Figure 1.1 shows a photo of the CCFL backlight in a conventional LCD display, while Figure 1.2 shows a diagram of an array of LEDs which replaces the uniform fluorescent backlight in the current generation of HDR displays.

HDR Displays. The dynamic range of a signal is the ratio between the largest and smallest values of that signal. In the realm of displays, the term is often used synonymously with luminance contrast or contrast ratio, which is the ratio between the highest and lowest luminance levels. Luminance is a measure of the amount of visible light emitted or reflected from a surface, which

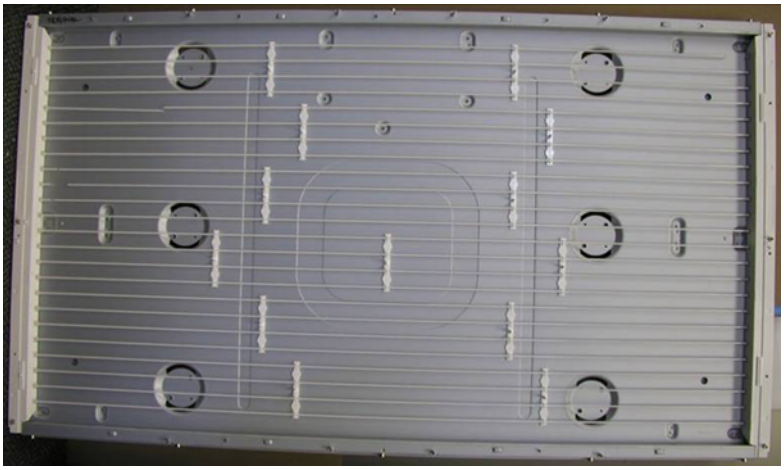

Figure 1.1: Photo of CCFL backlight in conventional LCD display. The light from the fluorescent tubes shines through a diffuser panel which evens it out across the display area. The uniform light then shines through the LCD panel which colours the light and produces the image. [Image courtesy H. Seetzen]

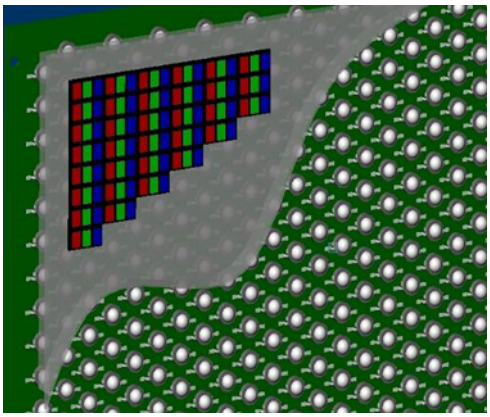

Figure 1.2: Diagram of low resolution LED array in current HDR displays. The LEDs can be individually controlled to produce a varying backlight, which when coupled with an LCD panel can produce HDR imagery. [Image courtesy H. Seetzen]

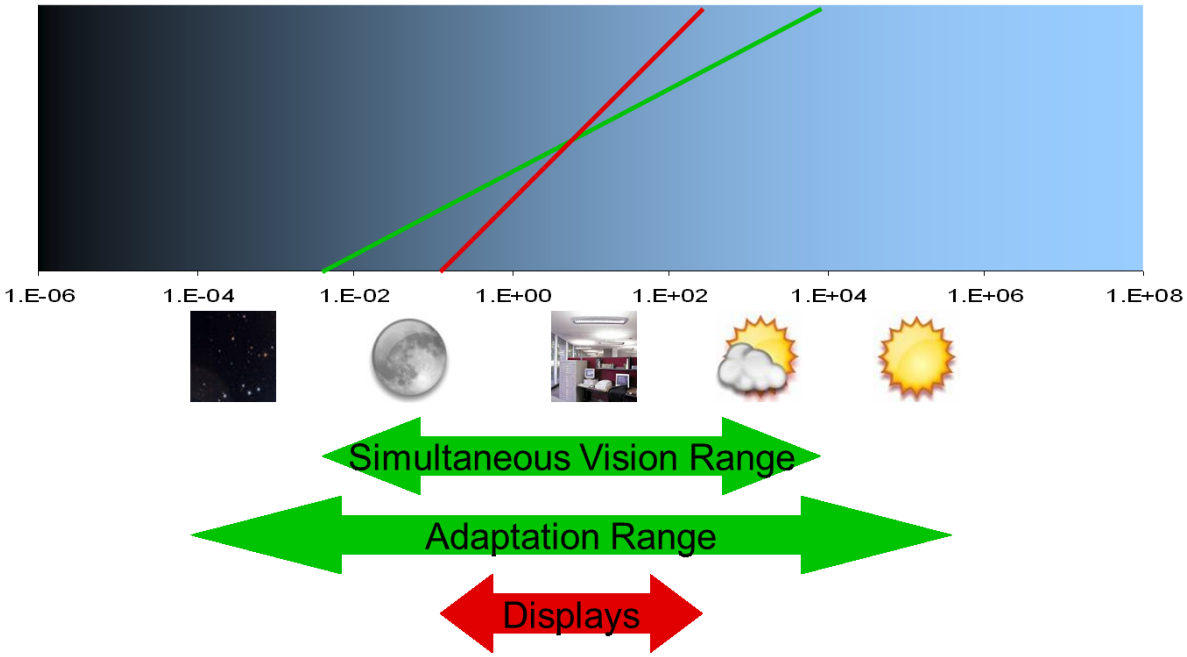

Figure 1.3: Dynamic ranges of displays, the human visual system, and the universe as visible from the earth. [Image courtesy H. Seetzen]

for our purposes is always greater than zero because even very dark rooms in which a display is in use always have at least some light reflecting off surfaces and scattering through the air. The dynamic range perceivable by the human visual system (HVS) is vastly greater than that which can be emitted by conventional imaging pipelines. Where the observable universe has a dynamic range of approximately $10^{14} : 1$, the HVS has a simultaneous vision range of approximately $10^{6} : 1$ which can be increased to $10^{10} : 1$ through the adaptation mechanism. Conventional low dynamic range (LDR) displays have a dynamic range on the order of only $10^{3} : 1$. However, the technique of dual modulation (modulating both LED and LCD layers) has enabled the development of prototype HDR displays whose dynamic range is on the order of $10^{5} : 1$, much closer to meeting the capabilities of the HVS than previous displays have been. Figure 1.3 shows these different ranges. A good overview of HDR technology and the capabilities of the HVS, including these and other details, can be found in [Seet 09].

All displays are limited in their ability to achieve the dynamic range of which they're capable, by the discretization or quantization inherent in the design of the displays themselves or the signals

which they are fed. Most modern conventional (LDR) images are discretized into 256 steps for each of the red, green, and blue colour channels, and modern displays (including the LCD panel that forms part of an HDR display) assume input images of that form. This *bit depth* of 8 bits may not always be sufficient to show a smoothly varying gradient between lighter and darker regions, and *banding* artifacts may result when there is a visible line between two adjacent regions whose transition is supposed to be smooth. This problem may be more pronounced when the dynamic range is higher, so modern HDR pixel values are often represented as floating-point numbers with a bit depth of 32 bits per colour channel[1], providing a vastly greater range of quantization levels [Debe 97].

A good overview of dynamic range and bit depth can be found in Seetzen et al. [Seet 04]. They discuss *just noticeable differences* (JNDs) which in the luminance domain refer to the smallest luminance difference that is perceived as being different, at a given baseline luminance level. They note that above approximately 10 cd/m^2, a linear increase in JND steps requires an exponential increase in luminance levels to be visually detected. This is consistent with Stevens' Law which describes an exponential relationship between the physical magnitude of a stimulus and the perceived intensity of that stimulus across many modes of human perception, including brightness [Stev 63]. In the context of HDR displays, as the luminance range increases so significantly, so too does the need to have a greater range of quantization levels to represent all the JNDs.

The increased capabilities of HDR displays coupled with the increased controllability due to the second modulator (the LED panel) create new challenges for the display of content in general, and existing low dynamic range (LDR) content in particular. The adaptation of LDR content for display on HDR devices is an underconstrained problem, as it involves the mapping of a small space into a much larger space, and there are many algorithms that have been proposed for that task (e.g. [Bant 06, Meyl 06, Meyl 07, Akyu 07, Didy 08]), each of which has its own features and limitations. In addition, the display of HDR content, or any other type of content, on HDR displays is quite recent and gives rise to questions about how that content is being perceived.

[1] Kunkel et al. have investigated the influence of bit depth on image quality and developed a high bit depth pipeline for future vision research [Kunk 11].

HDR display devices also provide a powerful new platform for research into the human visual system (HVS) and perceptually motivated design considerations for future generations of displays. The fine controllability of the intensity and colour of the backlight allows for the construction of experimental environments over a vastly broader range of viewing conditions (brightness and contrast conditions of viewing stimuli) than is possible with conventional LDR displays, which can be used to reveal characteristics of the HVS and suggest how those characteristics can be used to design better displays. These can be single-display experiments in which the HDR display occupies most or all of the subjects' attention, or multiple-display experiments in which subjects must divide their attention between multiple environments. The former can be used to simulate entertainment or office environments, while the latter can be used to simulate the cockpit environment of a car or an airplane, where a driver or pilot must maintain attention outside the vehicle while affecting the necessary controls inside the vehicle.

As these new generations of displays are being developed, they introduce two new avenues of research: investigating how to take advantage of the characteristics of these displays to improve the appearance of images and video, and observing and characterizing the responses of the human visual system to the new kinds of imagery these displays are capable of producing. The contributions of this dissertation advance our understanding along both of these research avenues, and inform the future development of both HDR and other types of displays to provide improved viewing experiences. In particular, this dissertation presents four projects, which are introduced below.

Video Viewing Preferences. HDR image capture technologies (e.g. [Debe 97]) have enabled the development of HDR content for HDR devices. However in the near term, legacy LDR content will continue to represent the vast majority of available content for HDR displays, and will in any event continue to be present just as old black-and-white film and television content continues to be present. Moreover, even as HDR pipelines become standard for new content creation, it may still be the case that some new content will be created with LDR technologies either because of the simplicity of the content or because of artistic choice. Therefore, the appropriate display of LDR

content on HDR devices will continue to be an important and persistent issue.

In Chapter 3, we describe a set of experiments in which subjects watched video content on an HDR display in several different ambient illumination environments and were asked to adjust the brightness and black level of the display to their preference. Subjects were also given questionnaires to document their observations and subjective preferences as well as any visual fatigue they may have experienced. A brief summary of our findings is that:

- users tend to minimize the black level settings within the physical limits of the HDR display we used in our experiments. This behaviour is independent of the ambient illumination.
- there were no signs of visual fatigue in any of the subjects, even with high contrast settings and in dark environments.
- the majority of subjects prefer lower display brightness for darker environments. A minority of subjects preferred close-to-maximum brightness independent of the ambient illumination levels.

Ldr2Hdr. In Chapter 4, we present *Ldr2Hdr*, a new reverse tone mapping algorithm that is specifically tailored toward on-the-fly processing of video streams without precomputation. Tone mapping is the process of mapping one set of colours or intensities to another set, and is often used to describe the mapping of HDR data to fit within the limitations of LDR devices; hence *reverse tone mapping* expands LDR data to fit HDR devices. We assume that the input LDR content has been optimized for viewing on regular TV screens, which is the case both for footage shot on video, and for film content that has been re-mastered for DVD distribution. While offline reverse tone mapping is an interesting research topic, our work is focused on methods that can be integrated directly into the display hardware to deal with video streams from conventional video cameras or DVD players. To support this goal, our method was designed to have the following properties:

- The algorithms are well suited for implementation on GPUs, as well as signal processors or field programmable gate arrays (FPGAs), which can be located in the display itself. The

algorithms are efficient enough to work in real-time on dynamic HDTV resolution video streams.

- No user input is required. All parameters can be chosen up front based on the hardware characteristics of the display.
- The method is robust in the sense that it does not produce disturbing artifacts. This is confirmed by the Visible Differences Predictor which showed a very low probability of visibly detectable artifacts being introduced.
- In particular, the output video stream is temporally coherent. Colours and intensities do not change abruptly unless they do so in the input image.

Night Displays. Chapter 5 discusses the next contribution, in which we conducted a set of psychophysical studies to determine how to improve display usability in dark environments, which involves using a spectral light distribution that is perceived well by human daylight vision photoreceptors (cones), but has very little effect on night vision photoreceptors (rods). By employing long-wavelength light (amber to red), the display can be seen by the daylight vision mechanism, which offers sufficient resolution to read information, while the night vision mechanism is mostly unaffected by the light emitted from the display. As a result, the display does not interfere with other tasks, such as driving. We then show that displays that employ separate red, green and blue LEDs in the backlight are capable of producing desirable emission spectra for night viewing so that no other hardware modifications in the display design are necessary to implement our approach. We:

- conducted a threshold experiment to demonstrate that achromatic contrast detection is mostly separable between rods and cones and therefore, long wavelength light almost invisible to rods has little effect on the night vision mechanism;
- show how this property of visual system can be used in the design of displays intended for use in the dark; and

- measured the desirable brightness of displays used in dark environments.

Contrast as a Depth Cue. The final contribution, discussed in Chapter 6, describes a set of experiments we conducted to observe the relationship between display contrast and depth perception. Since high dynamic range (HDR) displays have been shown at conferences (e.g. SIGGRAPH 2004), they have been confused with 3D displays by some observers. In this work, we explore this perceptual connection by conducting a series of experiments to examine the effect that contrast has on depth perception. In particular, we consider the contrast both of large-scale features and of small-scale features, both independently and in concert. We found that in each of three experiments, subjects perceived increases in contrast to correspond with increases in perceived depth. Our findings indicate that we can simulate sensations of depth by manipulating contrast, particularly that of highlights within images, and that modern high-contrast displays can simulate greater sensations of depth.

Chapter 2

Related Work

In this chapter we describe previous work related to the various components of this dissertation.

2.1 High Dynamic Range Imaging and Displays

It is widely acknowledged that the dynamic range of the traditional imaging pipeline (from the photographic capture of real world images to the digitization, storage, and display of those images) is severely limited compared to the abilities of the human visual system. This observation has led to active research on high dynamic range (HDR) image sensors (e.g. [Acos 00]), image capture (e.g. [Debe 97]), tone mapping (e.g. [Rein 02]), video standards such as the 10-bit log H.264 (AVC), and HDR file formats (e.g. [Ward 05, Mant 04a, Mant 06]). HDR displays were first demonstrated by Seetzen and co-workers [Seet 04, Seet 03]. Over the past few years, major display manufacturers such as LG Philips, Samsung, and AUO have started to work on this technology, and have presented prototype displays (e.g. [Phil 06]), indicating that HDR display technology is now getting ready for the consumer market.

The design of the current generation of HDR displays is based on the technique of *dual modulation* [Seet 09, ch. 6][Tren 06, ch. 3][Tren 07]. The use of two modulators in series results in output images whose contrast is the product of the contrasts of the two modulators. There are several types of displays that can be constructed by employing this technique. Two notable designs are described in [Seet 09]. One is a projection system in which a light source is shone through two modulators in series. The other is the use of a panel of individually controllable LEDs that provides the backlight to an LCD panel, which is the design used for the Brightside/Dolby DR-37P prototype display as well as other displays that have been developed for the consumer market.

Displays such as the DR-37P require the use of specialized image processing algorithms that set the LEDs and LCD to appropriate values for the input imagery. When the LEDs are set to varying levels in the brighter and darker regions of an image, the resultant nonuniform light field requires adjusting the image that would otherwise have been sent to the LCD panel, as that image had assumed a uniform light field. Figure 2.1 shows the flowchart for this process. The original input image is downsampled to the resolution of the LED array, to determine the desired intensity levels at each LED. The resultant LED levels are then multiplied by the *point spread functions* of the LEDs to determine how the light from each LED would illuminate its neighbourhood, and how at each pixel the backlight intensity is potentially composed of light from multiple neighbouring LEDs. The original input image is then divided by this resultant simulated backlight intensity map to arrive at the final image that is sent to the LCD panel.

2.2 Viewer Preferences Under Varying Ambient Illumination

LDR to HDR Expansion. Following the introduction of HDR displays [Seet 04], researchers started developing methods to improve the contrast of legacy low dynamic range images and video to make use of the improved dynamic range of HDR displays [Bant 06, Meyl 06, Meyl 07, Remp 07, Didy 08]. This process can be seen as an inverse to conventional (forward) tone mapping operators (TMOs, e.g. [Tumb 93, Rein 02, Patt 98]). While the technical details of the reverse or inverse TMOs differ, and only some of them aim to provide a true mathematical inverse to a (forward) TMO, they share the common goal of improving image contrast without introducing disturbing artifacts. Akyüz et al. [Akyu 07] demonstrated in a user study that even a simple linear scaling of the content can be successful in achieving this goal.

Perceptual Studies. Other researchers have also turned their attention to user studies to evaluate some of the forward operators and to generally learn about viewers' responses to HDR imagery. Ledda et al. [Ledd 05] used an HDR display in combination with two LDR displays to determine which TMOs produced images that better represented a reference HDR image. They found that of the TMOs available at the time, Reinhard's photographic operator [Rein 02] and Johnson and

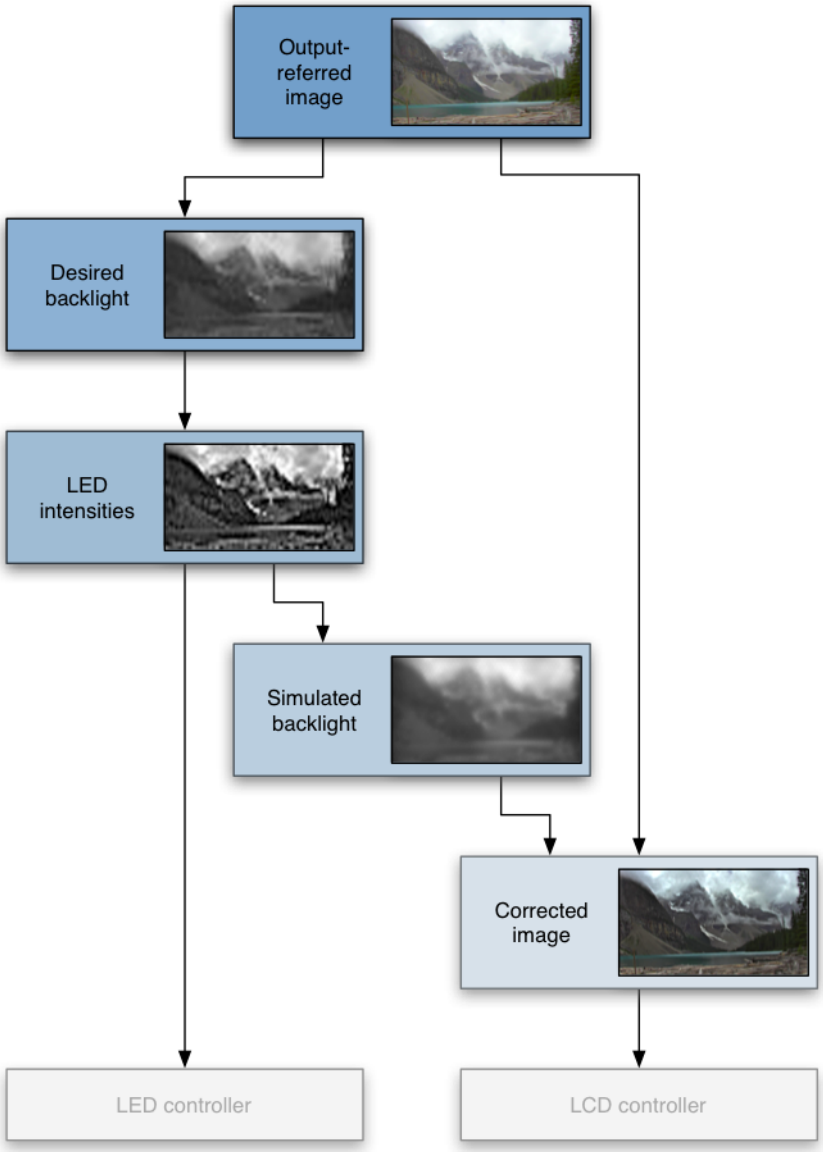

Figure 2.1: Flowchart of LED/LCD imaging algorithm. The input image is used to determine the desired backlight levels, which determine LED driving values. The LEDs shine through a diffuser panel, and the point spread functions of the LEDs shining through that panel are composed to simulate a map of the overall backlight intensity. The "Corrected image" which is shown on the LCD panel is then the input image adjusted to compensate for this varying intensity. Specifically, because the LED backlight intensity variations are blurred (in order to avoid possible mis-alignment problems with the LCD image), the LCD image must be sharpened to preserve the contrast of edges in the overall image. [Image courtesy M. Trentacoste [Tren 06, p. 51]]

Fairchild's iCAM [Fair 02] colour appearance model produced the best images. Ashikhmin and Goyal [Ashi 06] evaluated five TMOs using real physical scenes to determine which operators produced images that were considered real, pleasing, or representative of the original scenes, and rather than finding a clear winner, concluded that subjective image quality determination benefits substantially from comparison with the real physical scene rather than simply a sense of abstract realism. More recently, Yoshida et al. [Yosh 07] and Čadík et al. [vCad 08] each evaluated TMOs using perceptual attributes such as brightness, contrast, and detail reproduction, and while the former found that local operators tended to outperform global operators, the latter found the opposite to be true, leading them to conclude that image quality is more complex and not so easily reduced to a single metric. Akyüz et al. [Akyu 07] evaluated six reverse operators on a number of images to determine which produced the most pleasing results, and found that a simple linear scaling operator performed remarkably well. In their study, Seetzen et al. [Seet 06a] obtained preferred parameters for luminance, contrast, and amplitude ratio for HDR displays, while Yoshida et al. [Yosh 06] studied the preferred brightness, contrast, and colour saturation characteristics of a TMO. These studies show an increasing value placed on psychophysical or perceptual validation of HDR imaging techniques than has formerly been the case, and demonstrate a variety of experimental design scenarios. However, many issues have remained unexplored, such as viewer preferences while watching video rather than still images, adaptation to the ambient illumination, visual fatigue from extended HDR viewing, and the relationship between ambient illumination and display brightness and contrast.

Ambient Illumination. Ambient illumination plays an important role in the perception of visual imagery. This relationship has been analyzed in research dating back decades, notably as one of Stevens' famous psychophysical observations [Stev 61, Stev 63] in which he determined a power law relationship between the physical strength of a stimulus and the perception of that stimulus, across a broad range of modes of human sensory perception. Novick [Novi 69] determined observers' preferred tone-reproduction curve gamma values, which govern the amount of curvature in the function relating input value to output luminance, on cathode-ray tube television monitors

under several levels of surround luminance, using both tungsten and artificial daylight ambient illumination. He observed that as ambient illumination increases, viewers' preferred gamma value markedly decreases, but that gamma is largely independent of either the colour of the ambient illumination or whether the image is displayed in full colour or green monochrome. This agreed with earlier work by Bartleson and Breneman [Bart 67] that used slide projection. De Marsh [De M 72] built upon these results and determined approximate gamma values of 1.0 for a bright surround (20 fL (foot-lambert) or 69 cd/m^2), 1.2 for a dim surround (4 fL or 14 cd/m^2), and 1.5 for a dark surround. More recently, Devlin et al. [Devl 06] considered the effect of ambient illumination on contrast perception for still images being shown on a cathode-ray tube monitor, and developed functions to correct for that, which they validated perceptually through an extensive series of experiments.

With the popularization of mobile devices, renewed attention has been given to the automatic adjustment of display brightness under varying ambient illumination. Merrifield and Silverstein [Merr 88] described a general model for the relationship between display brightness and ambient illumination, consisting of a log-log linear relationship in the photopic luminance range and a constant relationship in the scotopic luminance range. More recently, Swinkels et al. [Swin 08] described a technique for adjusting the display brightness of a mobile device in a rapidly fluctuating ambient light environment without introducing flicker, which is the rapid vascillation of colour or brightness in all or parts of an image.

Adaptive Display. Several patents describe devices that modulate the output of a display (usually LCD) in response to variation in ambient illumination. Some (e.g. [Pitt 76, Kalm 93]) simply modulate the display brightness in response to ambient illumination for day or night viewing, while others have a more complex mechanism to adjust the output to compensate for reflected ambient light (e.g. [Capp 97]) or to correct for spatial non-uniformities that can be caused by a number of factors including ambient light (e.g. [Katy 01]). However, these all deal with low rather than high dynamic range displays. Further, our results suggest a much more smoothly and finely defined variance in display brightness than the coarse two-level brightness control of the Pittman

and Kalmanash patents, and suggest a more global tone mapping operator than the Cappels or Katyl patents.

Mantiuk et al. [Mant 08a] have developed a tone mapping operator that modulates its output based on the capabilities of the display and the ambient lighting in which the output is to be displayed, using image statistics and a model of the human visual system. Ghosh et al. [Ghos 05] took the opposite approach by developing a system in which the ambient lighting of the room was computer-controlled and varied with the scene being displayed. An informal survey showed that participants who played a racing video game strongly preferred the dynamic lighting over static lighting. Another similar technology is the line of "Ambilight" flat-panel television displays which create varying ambient lighting by illuminating the wall behind the display based on the intensity of the displayed imagery [Phil 08].

2.3 Dynamic Range Expansion and Ldr2Hdr

Advances in imaging technology result in continual improvements of image quality for new content. However, once data has been captured and encoded, its quality is frozen and determined by the state-of-the-art of imaging technology at the time of capture. Often it is undesirable, or even impossible, to later recapture scenes with new hardware or techniques, so the existing data must be enhanced in order to take advantage of improved output devices. Prominent examples are the colourization of black-and-white imagery [Wels 02, Levi 04], removal of scratches and repair of damaged image regions, or the removal of unwanted objects [Hira 96, Born 02, Bert 00, Sun 05].

The problem of estimating HDR images from LDR photographs has also recently received some attention. Li et al. [Li 05] discuss the problem of HDR *companding*, i.e. the process of first compressing the contrast through tone-mapping, and later expanding the contrast back to the original range, which is similar to the companding process outlined by Clark et al. [Clar 28]. The expansion step assumes that the intermediate LDR image has been created with the specific contrast reduction algorithm they propose. Meylan et al. [Meyl 06, Meyl 07] detect specular highlights in LDR images and boost the intensity in these regions. The stability of this segmentation across

frames in video sequences is unclear, and presents a potential source of artifacts. It is also worth noting that current HDR displays built on the principles described by Seetzen et al. [Seet 04] are not very good at reproducing very small bright areas such as isolated highlights.

In work parallel to ours, Akyüz et al. [Akyu 07] determined through psychophysical studies that in many circumstances a linear contrast scaling works surprisingly well for mapping LDR content onto HDR screens. Unlike Meylan et al., this approach does not attempt to specifically boost the contrast in saturated regions such as highlights.

Banterle et al. [Bant 06] found a way to approximately invert Reinhard et al.'s photographic tone mapping operator [Rein 02]. While their approach produces beautiful HDR images, it is not suitable for our purposes for a variety of reasons. First, it is too slow for real-time purposes. Second, it requires the image-specific adjustment of parameters, which cannot be tolerated in our setting. Perhaps most importantly, Banterle et al.'s algorithm relies on segmentation of the images into regions, which can result in temporal artifacts due to image noise, compression artifacts, and other minute differences between frames.

2.3.1 More Recent Work on Dynamic Range Expansion

Since the publication of our Ldr2Hdr paper [Remp 07], there has been some additional work on dynamic range expansion and related technologies.

The creation of new techniques for the production of high dynamic range images opens up a need for metrics to evaluate the quality or fidelity of those images. The visible difference predictor [Daly 93], which had already seen extensions to HDR [Mant 05] saw further extensions [Aydi 08, Mant 11] based on newer more accurate models of human vision and data from perceptual experiments.

Another approach to the evaluation of reverse tone mapping operators, which was employed by Masia et al. [Masi 09] is to observe how the performance of the various techniques degrades as they are fed input images with varying levels of under- or over-exposure. They conducted perceptual experiments with photographs taken at each of five different exposure levels, where the highest and lowest exposures always had brighter and darker details (respectively) washed out

and indistinguishable. They observed that spatial artifacts were more disturbing than inaccuracies in the expanded intensities, and that less aggressive techniques (ie. conservative techniques with smoother operators and smaller contrast boosting factors) might perform better when the quality of the input imagery is unknown.

In addition to evaluating operators, work has also continued on the development of new operators for enhancing the dynamic range of images. One approach is to attempt to mathematically invert the tone mapping operator that is estimated to have created the original low dynamic range images of the high dynamic range scene [Bant 08]. Another approach is to use models of the human visual system, such as those used for the visible difference predictors, to develop a technique that adapts to the capabilities and environment of the display on which the imagery is being viewed [Mant 08a].

The expansion of dynamic range brings with it also the expansion of colour gamuts. Heckaman and Fairchild [Heck 06] investigated this effect by "pushing down" the white point of HDR imagery and expanding the perceived lightness, chroma, brightness, and colourfulness of the imagery. They also noted the increases in bit depth necessary to represent these expanded gamuts.

2.4 Display Considerations for Night and Low-Illumination Viewing

The development of displays that are suitable for viewing in low-illumination environments follows on previous work in several areas, discussed here. Disability and discomfort glare can be responsible for diminishing the legibility of displays as well as the ability to see beyond the display in dark environments. One common approach to improving the ability to see beyond the display is by reducing the intensity of the backlight. Overall, the luminance range applicable to displays viewed in the dark falls in the realm of mesopic vision, in which the predictions of previous visual models often do not hold, and so the results of this section can aid in the development of new more comprehensive visual models.

Disability and Discomfort Glare. A display seen in the dark will always cause a certain amount of disability and discomfort glare. Disability glare is due to the light that is scattered in the eye optics and on the retina, which elevates adaptation luminance and reduces contrast by increasing the retinal luminance of the stimuli[1]. Discomfort glare is observed when a bright source of light either is distracting or evokes a dazzling effect that causes the eyes to squint or avert (sometimes referred as dazzling glare) [Vos 03]. Vos and van den Berg [Vos 99] proposed a comprehensive model of disability glare including a set of equations collectively covering a range of glare angles from 0° to 100°.

Backlight Dimming. The majority of the work on dimming the display backlight is motivated by energy saving rather than improved display usability at low light. The basic idea involves compensating for the dimmer backlight with increased transparency of the LCD layer, so that the difference between the undimmed image and the image with dimmed backlight is minimal [Chan 04]. The compensation methods may account for both temporal aspects to reduce flicker visibility [Iran 06], and spatial aspects to reduce contours due to hard clipping [Kero 07].

Mesopic Vision. Mesopic vision, in which both cones and rods are active, is an important factor in display design because displays intended to be used at low ambient light levels need to operate in the mesopic luminance range. The mesopic range starts at about $10^{-3}\,\mathrm{cd/m^2}$(photopic) and ends at about $5\,\mathrm{cd/m^2}$, although reported ranges vary [Wysz 82, p. 406]. Interactions between cones and rods are complex and not fully understood. A good review of the work on mesopic vision can be found in Hess et al. ([Hess 90, p. 34–37]) and Wyszecki and Stiles ([Wysz 82, p. 549–552]). To quantify visual performance at scotopic light levels, a practical model for mesopic photometry has been proposed to CIE for standardization by Elohalma and Halonen [Eloh 05]. The model assumes that the mesopic luminance is a linear combination of scotopic and photopic luminance, which has been shown to well approximate the measurement data for relatively broad-band light sources (about 100 nm). Elohalma and Halonen also note that mesopic luminance for narrow-

[1]Contrast reduction is caused by the scattered light L_s, which elevates background luminance L: $\Delta L/(L+L_s) < \Delta L/L$, where ΔL is a luminance difference and $\Delta L/L$ is the contrast without scattering.

band light sources is too complex to be modeled as a linear combination of scotopic and photopic luminance, but more complex models do not provide a better match to the data for broad-band light sources[2].

Visual Models. Several visual models that account for both cone and rod photoreceptor performance have been proposed in computer graphics [Ferw 96, Patt 98, Patt 00]. These models can potentially predict the phenomena relevant for night displays. However, the models proposed so far have not been well calibrated and validated on the psychophysical data. While these models give good qualitative predictions, their quantitative predictions are not reliable in the scotopic luminance range which is of interest to us. In our work, we obtain new data which can be used to develop a visual model that includes an emphasis on the prediction of intra-ocular light scatter, such as the very recently developed HDR-VDP-2 [Mant 11].

2.5 The Role of Contrast in the Perceived Depth of Monocular Imagery

One of the first descriptions of contrast as a depth cue dates back to Leonardo da Vinci, who observed the phenomenon of aerial perspective, in which atmospheric haze reduces the contrast of distant objects and alters the colour to be more blue [Rich 39, p. 210-212, 234-241].

More recently, many experiments have been conducted to determine the effects of brightness and aerial perspective on depth perception.

Brightness. Miles [Mile 53] observed the interaction between brightness and depth perception by noting that binocular viewing of two images with unequal brightnesses produced anomalous depth perception. Gilchrist [Gilc 80] conducted an experiment in which patches of paper at varying depths and with varying luminance were seen monocularly through a pinhole, and found that

[2]We attempted to use their model to analyze our experimental results, but found that the range of photopic to scotopic luminance ratios described by the model is too small to apply the model on our data.

the perceived lightness (from a Munsell chart) of a constant-luminance patch differed dramatically depending on how far away the patch was perceived to be. Gilchrist also suggested that the causality between perceived lightness and perceived depth may be bidirectional. Schirillo et al. [Schi 90] built on Gilchrist's experiments, separating the concepts of brightness and lightness, and using binocular stereoscopic viewing of a computer monitor in place of the physical setup Gilchrist had used.

Contrast. O'Shea et al. [OShe 94] investigated aerial perspective by conducting an experiment which showed using a computer monitor that objects with higher-contrast edges are deemed to be nearer to the viewer than objects with lower-contrast edges. Ichihara et al. [Ichi 07] extended this work by differentiating area contrast (that of large-scale features) from texture contrast (that of small-scale features) and observing that both have an effect on perceived depth.

Colour. Triesman [Trie 62] observed that both contrast and colour could affect depth cues in the viewing of stereoscopic scenes, but that the strongest cues were intensity differences at the edges of objects in the scenes. More recently, Troscianko et al. [Tros 91] observed that while a colour gradient between red and green did not significantly affect depth, a colour saturation gradient between red and grey did.

Integration of Depth Cues. There are a variety of cues that provide indications of depth within visual imagery. The binocular cues include stereopsis (the fusion of disparate retinal images into a three-dimensional mental model) and oculomotor effects, the cues from the strain of ocular muscles. The host of monocular cues includes parallax, perspective, the relative sizes of objects, the sizes of familiar objects, occlusion, texture gradients, blur, haze, and others. These cues may agree or disagree with each other and may vary in the strength of their effects, and that variance is often based on image content. Wijntjes and Pont [Wijn 10] observed that binocular stereo could improve the depth perception in images, but it is not a foregone conclusion that stereo is always better as not all images in their study benefited from a stereo representation. Held et al. [Held 10] manipulated focus and blur to create an effect similar to tilt-shift photography, in which the lens is tilted

to create a shallow depth of field effect, for simulating a miniature scene. Cipiloglu et al. [Cipi 10] developed a framework based on fuzzy logic for enhancing depth perception of imagery using many different depth cues.

Chapter 3

Viewer Preferences Under Varying Ambient Illumination

Conventionally available "low dynamic range" (LDR) display technologies have a very limited dynamic range compared to the abilities of the human visual system (HVS). This has led to the development of "high dynamic range" (HDR) technologies, from HDR image capture (e.g. [Debe 97]), through image transmission and HDR file formats (e.g. [Mant 04a, Ward 05, Mant 06]), to the HDR displays[1] that ultimately convey the imagery back to the HVS [Seet 03, Seet 04]. Over time, native HDR content will likely continue to be developed, but in the near term, the vast majority of available content will continue to be conventional LDR content.

Adapting existing content for display on HDR devices is an underconstrained problem, and one that has received increased attention in recent years. Much of this work has centered around the transformation of imagery between HDR and LDR formats through tone mapping (e.g. [Patt 98, Ward 05]) and reverse tone mapping (e.g. [Bant 06, Remp 07]). However, user studies are often helpful to provide validation of techniques as well as other information about how displays should be set up for optimal viewing. Recently, HDR displays have been the focus of a number of user studies (e.g. [Ledd 05, Seet 06a, Yosh 06, Akyu 07], but these have only taken into consideration still images, not video, and have not considered the possibility of visual fatigue under long term use, the adaptation of the eyes to ambient light or the relationship between ambient light and user preferences for the presentation of HDR imagery.

[1] Some technologies such as 35 mm film have higher dynamic range than conventional LDR television sets and computer monitors, but our focus here is on HDR displays which could be used to replace those LDR display technologies, which begin with Seetzen et al.

The research presented in this chapter analyzes these factors. In order for HDR displays to emerge as viable alternatives to standard televisions, visual fatigue, such as double vision or headaches, needs to be explicitly ruled out even for extended viewing periods. Similarly, it is necessary to understand if higher brightness and dynamic range are preferred by users in standard television environments, and how preferences for these parameters might depend on the surroundings. To that end, we constructed the viewing environment pictured in Figure 3.1 to determine how users' preferences are affected by dark and bright ambient surroundings.

Many consumer-grade devices such as displays and camcorders have special "demo modes", in which colour saturation, sound volume, and similar parameters are artificially enhanced. Such modes give the impression of better image or sound quality in side-by-side comparisons with other devices on a showroom floor, but they can be ill-suited for extended use if the viewing environment is not well matched to the display settings, set for the showroom floor. A key goal of our work is to analyze whether enhanced brightness and contrast produce a similar "demo effect", or whether they yield sustained improvements in perceived image quality.

3.1 Experimental Design

We conducted two experiments to explore visual fatigue, viewer preferences, and the relationship between ambient lighting and preferred brightness and contrast. The first study, Experiment 3.1, was primarily designed to obtain information about visual fatigue in standard low-light settings, similar to the environments commonly used for watching television. This experiment used longer (1.5 hour) movies.

The second study, Experiment 3.2, tested a wider range of lighting configurations, including one comparable to bright office lighting. The goal of this experiment was to analyze the dependency of viewer preferences on ambient light levels. This experiment used shorter (22 minute) TV shows.

3.1.1 Setup

Our goal was to construct a viewing environment as realistic and free from distractions as possible. We wanted participants to come as close as possible to the experience of having a high dynamic range home theater environment in which to make their adjustments and comments.

Display. We used a Dolby DR-37P 37" (94 cm diagonal, 46 cm high, 82 cm wide) prototype high dynamic range display with an LCD panel and LED backlight. We measured the maximum and minimum intensities and contrast of the display with a Minolta LS-100 luminance meter. Maximum luminance was 4000 cd/m^2, while the minimum was below the detection threshold of the luminance meter (0.001 cd/m^2). The maximum simultaneous contrast (ratio between the lightest and darkest regions displayable at the same time) was approximately 90,000:1. The response curves of the individual colour channels were measured and accounted for in all experiments.

Viewers sat in an armchair at a viewing distance of about 1.5 m, which is approximately the recommended HDTV viewing distance of three times the height of the display (3H) [ITU 90], with the viewer's eyes level with the center of the display.

Acoustics. Since the prototype HDR display produces a significant amount of fan noise, acoustic damping was required for our experiments with video viewing. In the first experiment, we enclosed the display in a sound dampening box and conducted the experiment in a room designed for a Noise Criterion level of NC 25 to NC 30 [Bera 57]. (For reference, theatres without amplification systems have a level of NC 20-25, while movie theatres have a level of approximately NC 30.) In the (shorter) second experiment, sealed headphones were used.

Illumination. In the first experiment, ambient illumination in the room was provided by six black incandescent torchiere floor lamps standing 184 cm high. Each was fitted with a Philips "Natural Light" 150 W incandescent bulb. The lamps were positioned as shown on the left of Figure 3.1. The light from each of the torchiere cones shone upward and outward, providing a diffuse ambient light that did not directly illuminate either the subject or the display.

Figure 3.1: Layouts of Experiments 3.1 and 3.2 (left); viewer in experiment 3.2 dark (center) and bright (right) ambient light.

Our first experiment focused on the issue of potential fatigue, and therefore we designed that experiment to use low ambient illumination levels (see Table 3.1). As we shall see in Section 3.2, these relatively small variations in ambient illumination did not result in statistically significant differences in user preferences for brightness or contrast.

In the second experiment, we increased the ambient illumination range by a factor of 10. To this end, we redesigned the lighting environment as shown on the right side of Figure 3.1. In addition to the torchieres from the first experiment, we used two 40 W fluorescent bulbs positioned on the floor on either side of the armchair, and two 500 W directional halogen lamps.

Amb. light, exp. 1	<0.01(*)	0.75	8.5	28	74
Amb. light, exp. 2	<0.01(*)	70	700		

Table 3.1: Ambient light levels in lux for Experiments 3.1 and 3.2, measured at the location of the viewer's head. (*) Ambient lighting was off entirely for the lowest setting.

We measured the five different ambient light settings of the first experiment and the three settings of the second experiment, with an illuminance meter positioned approximately where the viewer's head would be positioned, oriented toward the display. The measured levels (in lux) are shown in Table 3.1. For comparison, typical office lighting is approximately 500 lux [Karw 01]. All measurements were taken with the HDR display switched off. The display itself contributes an average of 4-7 lux to the overall ambient illumination, depending on the specific content being shown, which is why, later in this chapter, we show the lowest ambient illumination level in Ex-

periment 3.2 as 7 lux. Figure 3.1 shows a viewer in the setup of Experiment 3.2 in both the lowest and highest ambient lighting conditions.

Video Content. All video content originated from commercial DVDs, and was adjusted for viewing on the HDR displays. For the first experiment, we used 5 feature-length movies ("Blades of Glory", "Cats & Dogs", "Charlotte's Web", "The Princess Bride", and "Zoolander"), while the second experiment used 3 shorter TV episodes ("Friends" Season 5, episodes 11, 16, and 19).

Currently, there are two methods for automatically and efficiently applying reverse tone mapping to video sequences: the simple linear-luminance scaling proposed by Akyüz et al. [Akyu 07], and the more sophisticated method by Rempel et al. [Remp 07]. Of these two methods, we chose the simple linear scaling for our experiments, for two reasons. First, Rempel et al.'s method is designed to be very conservative, and thus it rarely uses the full display intensity. This makes it difficult to use that method for analyzing brightness preferences. Second, linear scaling provides a worst case scenario for the amplification in contrast. The gradients in natural images follow a heavy tail distribution [Rude 94, Dror 04]. Linear scaling of contrast shifts the whole curve, while more sophisticated methods [Remp 07, Meyl 07, Bant 06] have a more localized effect. For our experiment, Equation 3.1 shows the output intensity I, for a linearized pixel value $x \in [0 \dots 1]$ in terms of a user selected peak brightness and black level:

$$I(x) = (\text{peak} - \text{black}) \cdot x + \text{black}. \tag{3.1}$$

User Controls. Instead of directly exposing the peak brightness and black level as user controlled parameters, we opted to remap these two parameters to controls that the subjects would be more familiar with as they are similar to the controls on normal televisions. The two controls we used were a "brightness" parameter b and a "contrast" parameter c, from which the peak luminance and black level were derived as follows:

$$\text{peak} = I_{max} \cdot b \qquad \text{black} = \text{peak} \cdot (1 - c).$$

The test subjects used a wireless USB input device to adjust brightness and contrast settings.

After each adjustment, the system briefly displayed a line of text indicating one of "BRIGHTNESS" or "CONTRAST" and one of "++", "−−", "MIN", or "MAX", denoting an increment, a decrement, or that the user had hit the bottom or top stops, respectively. To avoid subjects remembering positions, subjects did not receive any other absolute indication of their locations on the brightness or contrast continua.

A possible concern with a study like ours was that subjects might be distracted by the content itself, and might neglect to focus on the given task of optimizing the parameter settings. In order to counter this effect, and help subjects to find optimal settings quickly, we implemented a *staircase* procedure, as follows. At any point in time we kept track of a *minimum* and a *maximum* value for each of the two parameters (brightness and contrast), which represent the range of acceptable values for those parameters. The range would be stretched if the viewer selected values outside the range for an extended period. After 3 minutes of inactivity, the ranges were reduced in size by one half and the parameter settings were automatically changed to the opposite end of the new range, prompting the subject to make further adjustments while at the same time converging to a desired "optimal" parameter setting. An example run for a typical subject is shown in Figure 3.2. The squares show the automatic settings made by the staircase procedure after 3 minutes of inactivity, while the diamonds show the last value the viewer selected (i.e. the viewer's preferred value) prior to the next 3-minute period of inactivity. The sequence shows the viewer consistently returning to the same levels after the automatic settings changes.

3.1.2 User Studies

Ten subjects (4 female, 6 male, aged 19–71) participated in Experiment 3.1, and seventeen subjects (10 female, 7 male, aged 19–79) in Experiment 3.2. Each of them had normal (20/20) or corrected-to-normal vision and normal colour vision, which was confirmed through the administration of a Snellen visual acuity test and the Ishihara [Ishi 07] colour deficiency tests. All subjects were naive to the purpose of the experiment. Table 3.2 lists the age/gender breakdown of the subjects.

Prior to the viewings, subjects were briefed on the task to be performed and were instructed to adjust the brightness and contrast of the HDR display to a level they found pleasing. Subjects were

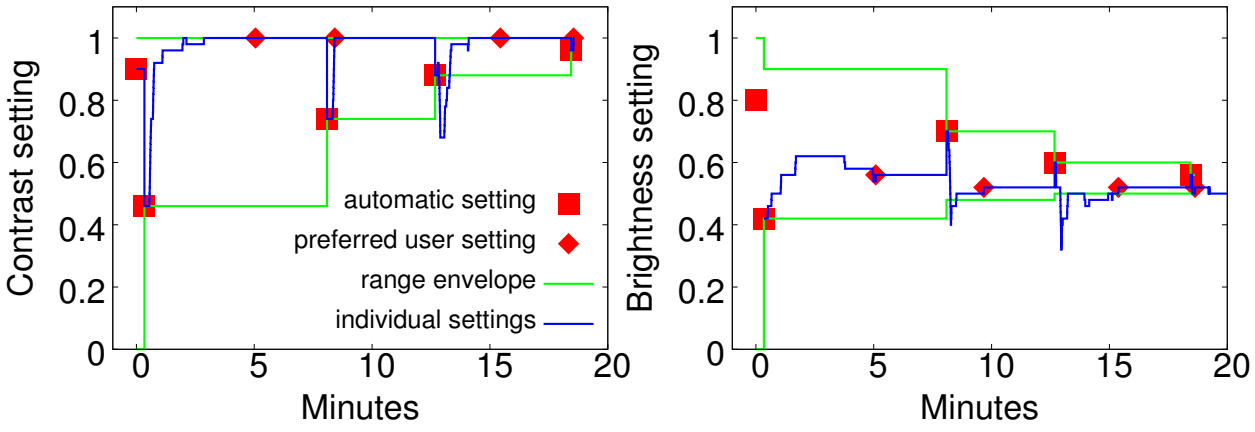

Figure 3.2: Automatic and user-selected settings for a typical subject. Each square is an automatic setting change; each diamond is the subject's last selected point before quiescence. The blue line shows each change made by the subject.

encouraged to continue to adjust settings throughout the experiment as desired. Subjects were also told that they might occasionally find brightness and contrast changing spontaneously, in which case they were to continue to adjust brightness and contrast to their desired levels.

In the first experiment, each subject watched the same five movies of approximately 1.5 hours duration each. Each movie was seen on a different day, rather than all on the same day, to avoid introducing fatigue due to extremely long sessions. In the second experiment, each subject watched the same three 22-minute television episodes. The shows were arranged in a random order for each subject. For each subject, each show was seen at a different ambient light level chosen from the set described above. The ordering of the light levels was also random, and independent of the order of the shows. In the second experiment, where each of the 6 orderings of shows and lighting levels occurred multiple times, we adjusted the order slightly to obtain a more uniform distribution of the conditions and ensure that no two subjects had both the same ambient light ordering and the same show ordering.[2]

The data collected consisted of a time-stamped log of all brightness and contrast changes made

[2]On hindsight, we could have constructed a *Latin square* design with $6 \times 6 = 36$ different conditions even though we did not have 36 subjects, and simply used as much of the Latin square as we had subjects for. However, our approach has a similar stratification effect.

Subject	Age	Gender
1	25	M
2	30	M
3	25	M
4	45	M
5	19	M
6	55	F
7	65	F
8	47	F
9	23	M
10	71	F

Subject	Age	Gender
1	(*)	F
2	(*)	F
3	(*)	F
4	23	M
5	37	F
6	55	F
7	19	M
8	28	F
9	42	F
10	36	M

Subject	Age	Gender
11	32	M
12	79	M
13	77	F
14	31	F
15	26	M
16	(*)	F
17	(*)	M

Table 3.2: Breakdown of subjects in Experiments 3.1 (left) and 3.2 (center, right). (*) indicates subjects who did not give us their ages.

during a session as well as a visual fatigue questionnaire and a more general questionnaire. The visual fatigue questionnaire asked users to identify on an 11-point Likert [Like 34] scale whether they experienced any of 10 symptoms of visual fatigue. These symptoms, listed in Table 3.3, are widely used in the human factors literature to self-report visual fatigue arising from the use of video display units [Dill 95]. Likert scales are commonly used in user studies to determine the strength of a subject's perception or opinion of some factor, and they may measure intensity from "none" to "extreme" or from "strongly disagree" to "neutral" to "strongly agree." The general questionnaire consisted of both subjective questions about their experiences during the experiment and questions which asked the viewer to compare the HDR display to other types of displays on a 7-point Likert scale.

3.2 Analysis and Results

Visual Fatigue. Our Likert questionnaires showed that viewers experienced remarkably little visual fatigue even after the longer (1.5 hour) movie sessions. Over half the subjects selected 0 ("none") for all symptoms, and over 95% of individual responses, including those of the older participants, indicated 0. In the remaining responses where subjects indicated more than 0 visual

Possible symptoms of visual fatigue	Worst case score (out of 10)
Double vision	0
Problems in focusing	1
Burning/pricking sensation in the eyes	2
Blurred vision	1
Tearing/Watery eyes	2
Pain around the eyes	0
Headache	0
Image break-up	0
Image floating	0
Colour change	0

Table 3.3: Possible visual fatigue symptoms with highest scores reported on questionnaire for subjects after each viewing, in Experiment 3.1.

fatigue, the average level was 1.18 out of 10.0. The overall average reported level of visual fatigue across all questions was 0.0325 out of 10.0. Five one-sample t-tests were conducted across the five different ambient illumination settings. The mean fatigue score was not statistically different from 0 in each case (t(9) = 1.5, t(9) = 1.809, t(9) = 1.5, t(9) = 1.809, t(9) = 1.406) and the low scores showed no correlation with the ambient condition. Only four of the symptoms received any scores greater than 0. Table 3.3 shows the highest score received for each symptom. For each of the four symptoms with scores above 0, the worst-case visual fatigue score (1 or 2 out of 10) was only reported in 1 out of 40 responses. Another set of visual fatigue scores obtained after the shorter sessions with higher ambient lighting in Experiment 3.2 was also not statistically different from 0 (t(16) = 1.3237).

These visual fatigue results were obtained at the viewers' preferred brightness settings. However, we do not expect significant visual fatigue at other settings or in other viewing environments either. The factors that have been found to contribute to aesthenopia (eye strain) in the human factors literature include viewing distance, duration of use, age, gender, and lighting and glare [Dill 96]. Our visual fatigue results are uniformly low in both high and low ambient illumination environments, among a demographically diverse group of subjects, for both shorter and

longer session durations, and at a relatively short viewing distance that would tend to exhibit high rather than low levels of aesthenopic complaints.

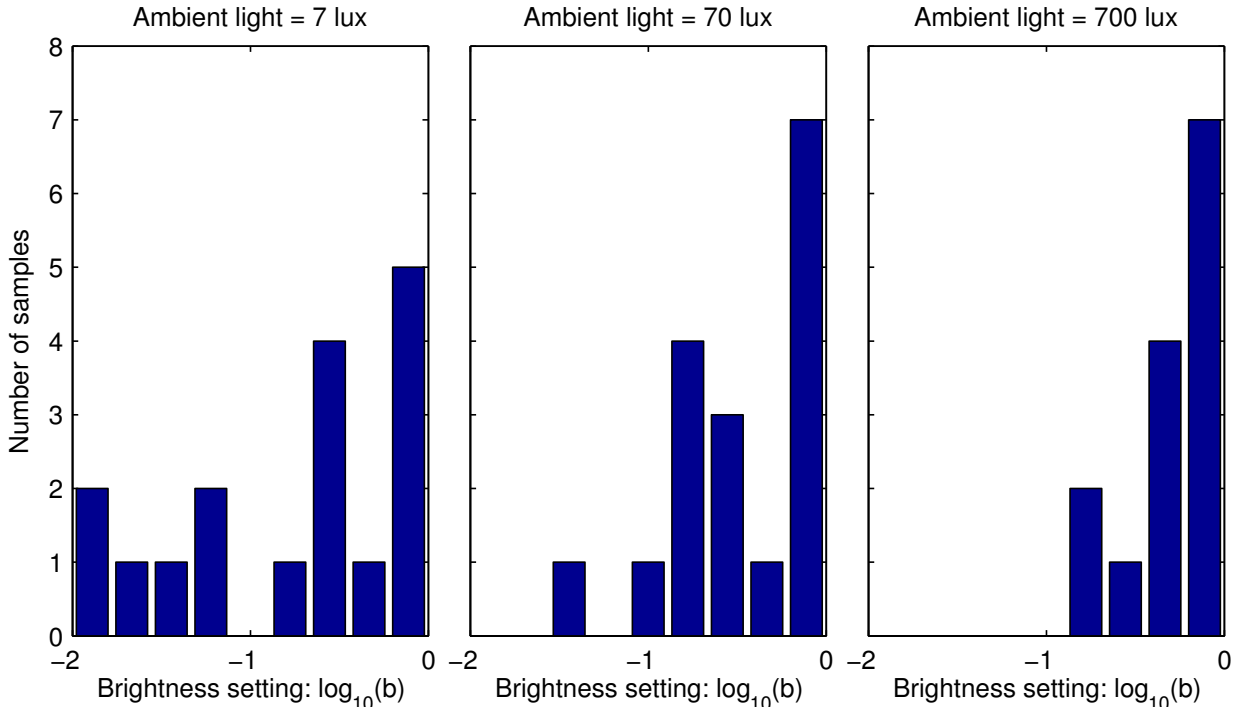

Figure 3.3: Histograms of the preferred brightness settings for the three ambient light levels.

Brightness and Black Level Preference. The adjustment patterns depicted in Figure 3.2, as well as similar graphs for other subjects shown at the end of this chapter (Figures 3.9, 3.10, and 3.11)[3] demonstrate that the staircasing procedure is effective in helping the subjects zero in on a preferred brightness setting over the course of a session. Further evidence for such a convergence is provided by the decreasing frequency of adjustments over time within each session (see below, and Figure 3.7). In our analysis, we therefore use the last setting for each user in each session as the preferred setting for that user in the respective ambient conditions.

The differences in the preferred brightness settings for the three ambient light levels are shown

[3]The contrast settings graphs for the full set of subjects showed uniformly high settings with insignificant differentiation and are omitted.

as histograms in Figure 3.3. As the ambient light increased, subjects tended to avoid the lower brightness settings. The histograms show significant individual differences in user preference for low ambient illumination, but a strong preference for bright screens in bright, office-like environments. The contrast preference is unaffected by the ambient illumination, with all subjects adjusting contrast to near maximal levels in all settings, effectively lowering the display black level to its minimum physically achievable value.

The rightmost (700 lux) histogram in Figure 3.3 shows fewer data points than the other two histograms. This is because in those missing cases, subjects happened to make frequent enough brightness adjustments that the automatic readjustment after 3 minutes of inactivity never occurred. In those cases, we calculated the time-weighted average of the subjects' settings during those sessions and determined that they fall within the same four histogram bins as the other data points, and thus do not alter our results.

We also analyzed aggregate preference over all subjects. Because our data fails the assumptions of normal distribution and equal variance, we use the non-parametric Kruskal-Wallis test rather than ANOVA to test for factor significance. The test indicated a significant difference in median brightness settings between the three ambient light levels ($\chi^2 = 6.25,\ p < 0.05$). The episodes were selected to vary little in their overall brightness, but some differences could not be avoided (compare cumulative histograms in Figure 3.5). We analyzed whether adjustments may be affected by the video content. This effect was found to not quite reach statistical significance at $p = 0.05$ ($\chi^2 = 5.69,\ p = 0.058$). This suggests that the subjects could compensate with the brightness setting for the overall video brightness. We did not find any effect of gender ($\chi^2 = 0.4$, $p < 0.05$), but we found that, of those subjects who told us their ages, the older group of subjects (36–79 years) chose the preferred brightness to be on average 0.3 log-10 units higher than for the younger group of subjects (19–35 years, $\chi^2 = 5.76$, $p < 0.05$).

The difference in the median brightness settings for the three ambient light levels, shown in Figure 3.6, was surprisingly low. If the subjects tried to fully compensate for the difference in the mean luminance of the surround, the slope of the curves in Figure 3.6 would be 1. This is because increasing the illuminance level by 1 log unit, results in roughly 1-log unit higher luminance of

all diffuse surfaces. Instead, the slope of the curve between the 70 and 700 lux setups was ≈0.26. This suggests that the large contrast between the luminance of the surround and the display content was not distracting and had moderate effect on the brightness settings.

The difference in the brightness settings can be partly explained by the increased reflectance of light from the screen and therefore loss of contrast. The light setup was designed in such a way that direct reflections of the light sources on the screen were avoided. For estimating the indirect illumination bouncing off the display, the LCD panel can be modeled as a diffuse reflector with a reflectance of about 1%. In a 700 lux environment, the luminance reflected off the screen is therefore approximately

$$0.01 \cdot \frac{700}{\pi} \frac{\text{lux}}{\text{sr}} \approx 2.2 \,\text{cd/m}^2. \tag{3.2}$$

If the original contrast shown on the screen is $\Delta L/L$, the contrast reduction caused by the reflected light, $L_{reflected}$, is

$$\frac{\Delta L}{L + L_{reflected}} = \alpha \cdot \frac{\Delta L}{L}. \tag{3.3}$$

In Figure 3.4 we plot how much contrast is lost $(1-\alpha)$ due to ambient light reflection for a range of luminance values produced on the screen. The vertical lines indicate 10th and 90th percentile of the video content at the preferred brightness levels for a particular ambient light setup. In low ambient light environments shown by the red and green dashed curves, the loss of contrast is low where the curves meet the lines denoting preferred display luminance. But as the ambient light increases to the blue curve, the contrast loss increases and the preferred display luminance levels have to increase (vertical lines shift to the right) in order to mitigate the contrast loss in darker image regions. The histograms in Figure 3.3 also show that subjects entirely avoided lower brightness settings for the high ambient light setup. Therefore, the loss of contrast in the dark regions could prompt subjects to elevate brightness settings for higher ambient light levels.

Since our prototype display could show luminance levels almost an order of magnitude higher than a standard TV display, we could check whether current consumer-grade TV displays are bright enough to meet viewers' preferences. The peak luminance of a typical LCD-TV display

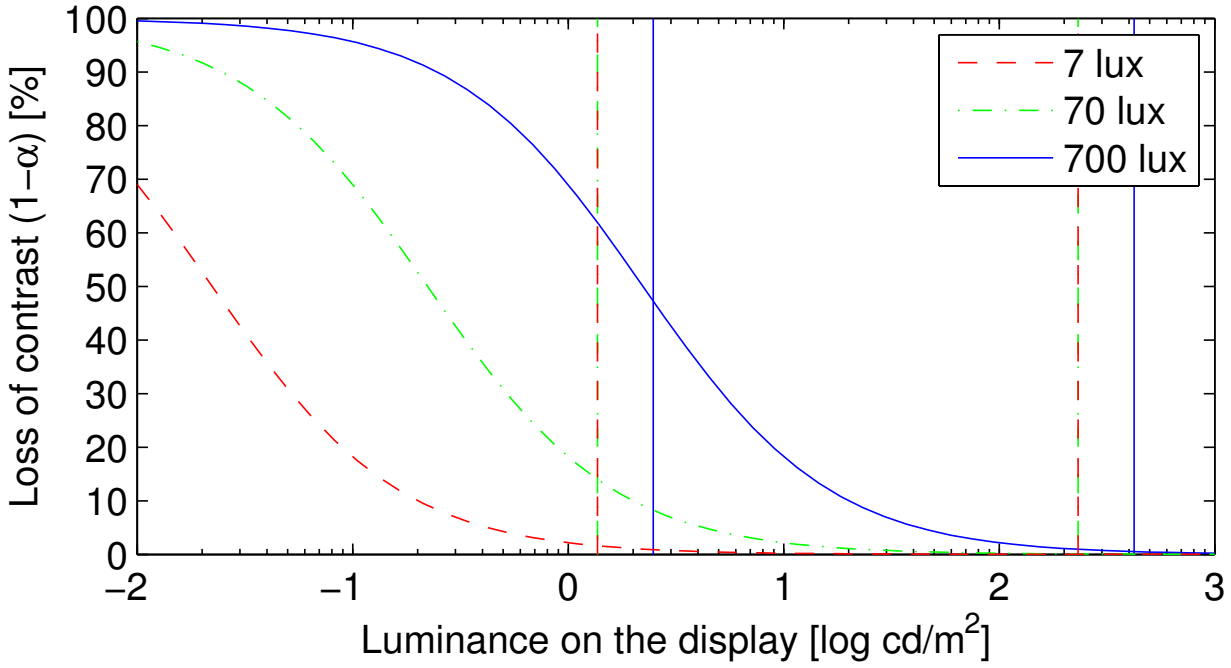

Figure 3.4: Contrast loss on the display due to reflections from the panel at the three ambient illumination levels used in Experiment 3.2. The vertical lines represent 10th and 90th percentiles of the video content for the preferred display brightness levels at a particular ambient light level (refer to Figure 3.6).

usually does not exceed 500 cd/m^2. Since the median pixel value for all video frames was 47 (see Figure 3.5), the corresponding luminance values on a display with the 500 cd/m^{2}peak luminance is ≈12 cd/m^2(L_{peak}·(47/255)$^{2.2}$). This value is lower than the median luminance of the video content shown on our prototype display (peak 4000 cd/m^2), which was 37.3 cd/m^{2}for the darkest and 67.7 cd/m^{2}for the brightest ambient light level, as shown in Figure 3.6. It must be noticed, however, that there were a number of subjects selecting the highest possible brightness setting, especially for the two brighter ambient light setups (see histograms in Figure 3.3). We did not allow the brightness to be further increased to avoid clipping of the video content. Nevertheless, we can expect that a brighter display could result in preferences for even higher luminance levels. This suggests that viewers may prefer displays offering higher peak luminance levels than those offered by the majority of displays, even for relatively dim environments.

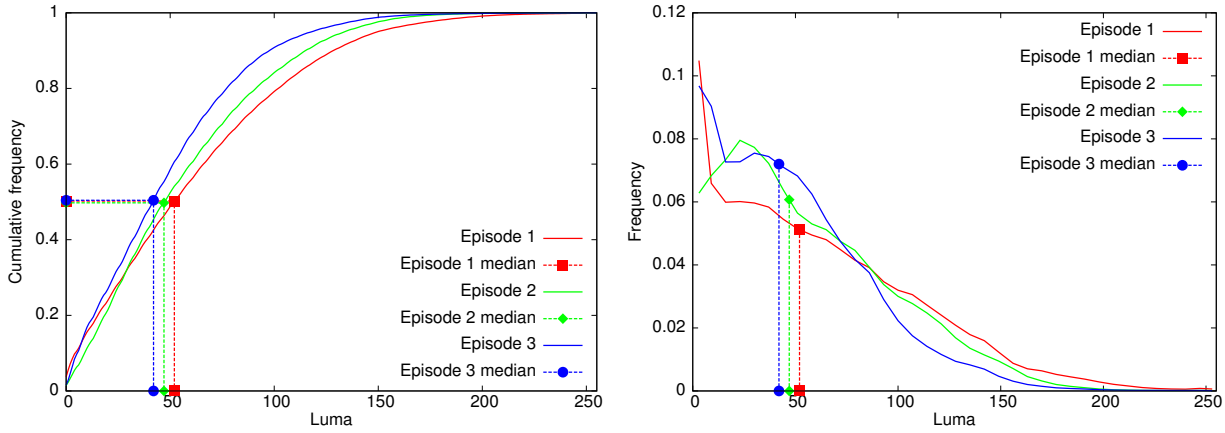

Figure 3.5: Cumulative (left) and non-cumulative(right) histograms of the luma values for the three episodes. The dashed lines indicate the median luma value for each episode.

This matches the observation by Seetzen et al. [Seet 06a] that perceived image quality increases with higher display peak luminance, as long as the contrast ratio is sufficiently high. In our experiment, all subjects tended to select the highest available contrast levels (i.e. the lowest available black levels) and most selected a brightness level significantly higher than what is possible on conventional displays. Seetzen et al. had high dynamic range still-image content. We can postulate that subjects in our study might have selected even higher luminance levels if the study had used true HDR content rather than adjusted LDR footage.

Both the slope and the shape of the brightness adjustment curve are very similar to that found in a different application area: cockpit displays. Merrifield and Silverstein [Merr 88] measured preferred display brightness adjustment for an aircraft cockpit display viewed under a range of ambient illumination levels. They reported the slope of the manual brightness adjustment to be 0.276 over the photopic range of ambient illumination, but found no changes in the brightness settings over the scotopic range of ambient illumination. This remarkably resembles our results shown in Figure 3.6, even though our task did not involve directing viewers' attention outside the display, as was the case for the cockpit display measurements.

Adjustment Frequencies. To further analyze the reliability of our results, we also considered the frequencies of adjustments both across sessions and within a single session. Figure 3.7 shows

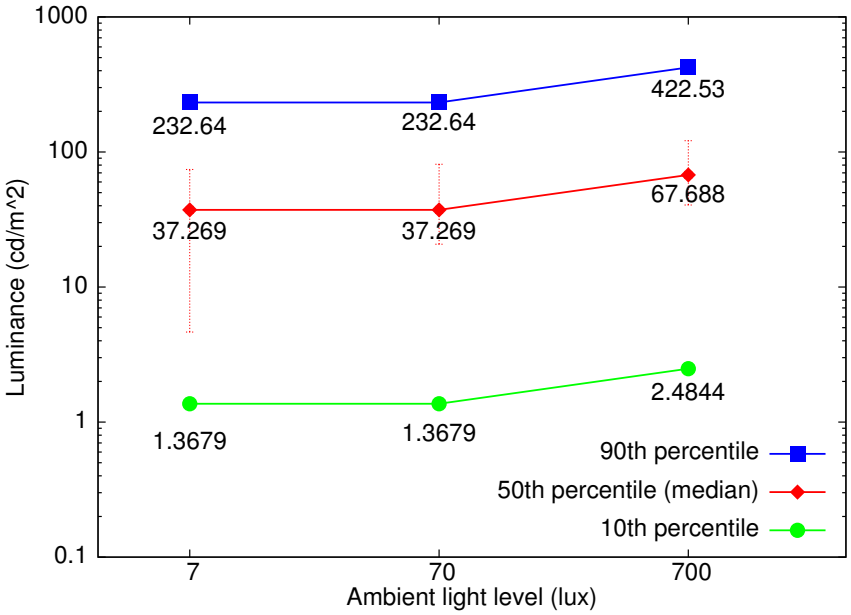

Figure 3.6: Preferred display brightness levels by ambient lighting level. The lines represent luminance of the 10th, 50th and 90th percentiles of the video content (computed for all the frames). The error bars show the 25th and 75th percentiles of the brightness settings. The preferred brightness setting was found by computing the median of all subject settings.

the resulting statistics for our first experiment. The second experiment produced similar results. As one might expect, subjects tended to experiment more with the parameter settings in the first session than in the remaining ones. However, our experimental design forced subjects to make a significant number of adjustments (200 on average) even for subsequent sessions. Within each session, subjects required time to find suitable settings from the initial random starting point, but then quickly settled on parameters they found most suitable. This convergence shows that final settings are reliable and not scene dependent.

Comparison to Other Types of Displays. In addition to the formal experiments described above, we also asked subjects to rate the experience of HDR displays in comparison to other display technologies they had previously experienced. In rare cases where subjects did not know the name of the type of display, we tried to help them figure it out by asking them about characteristics such as the size and age of their tv. Subjects showed a strong preference for watching movies on

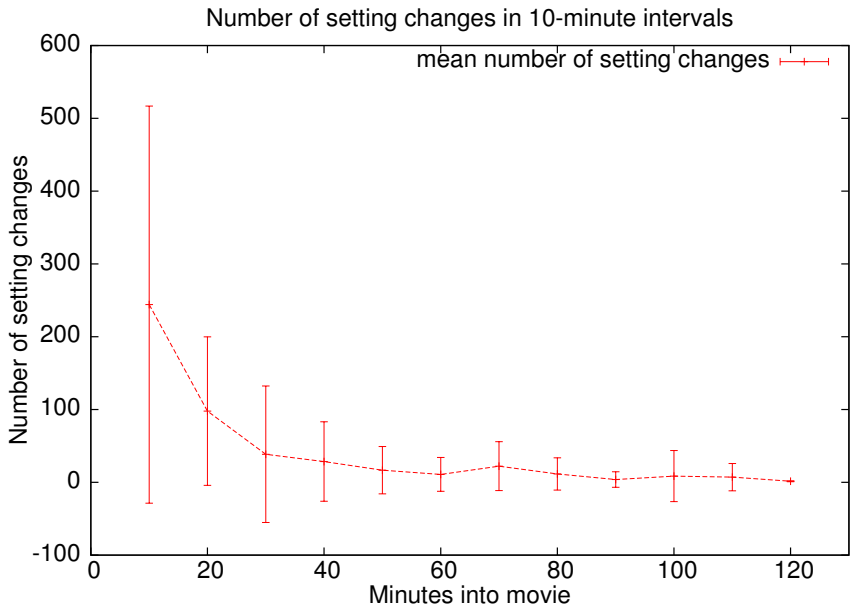

Figure 3.7: Average setting changes made during a first session. Subsequent sessions showed similar patterns with approximately half the overall number of changes per session.

the HDR display over CRT displays, and slightly weaker preferences for the HDR display over rear projection televisions and LCD displays. Six subjects were indifferent between the HDR display and a movie theater and only three preferred the movie theater. Figure 3.8 shows all the types we identified and subjects' preferences among them.

3.3 Discussion

We conducted two user studies in which we analyzed viewing preferences and the potential for visual fatigue on next generation HDR displays. Our experiments show that visual fatigue is not a serious concern even in dark environments. While subjects tended to always maximize the available display contrast, we found a sub-linear relationship between the preferred display brightness and the level of ambient illumination. These results are consistent across a wide demographic spectrum.

The results of this study could be the first step in designing HDR television sets with automatic

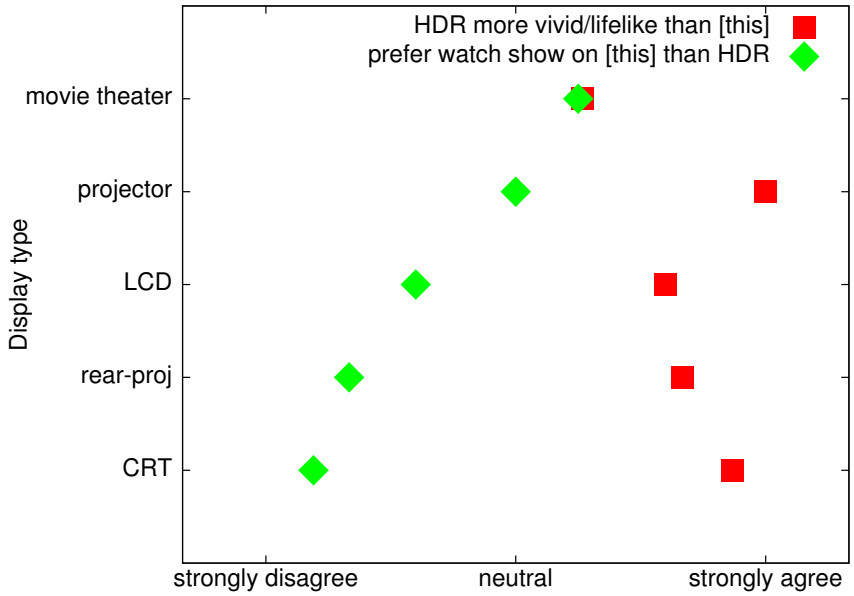

Figure 3.8: Subjects preferred HDR display and found it more vivid.

brightness controls to provide a more pleasurable viewing environment under a variety of ambient viewing conditions.

In the current study we have focused on live action content without drastic illumination effects. This choice was based on the need to limit the number of parameters in the study. Now that visual fatigue has been ruled out and we have gained a better understanding of brightness and contrast preferences under different ambient light levels, we would like to conduct further studies that analyze the effect of different types of content, such as flat-shaded animation and very dark or highly stylized live action footage.

The preference that viewers express for maximizing the available contrast calls for techniques that enhance the contrast of existing LDR content in sensible ways for display on HDR devices. The next chapter describes a technique that accomplishes this.

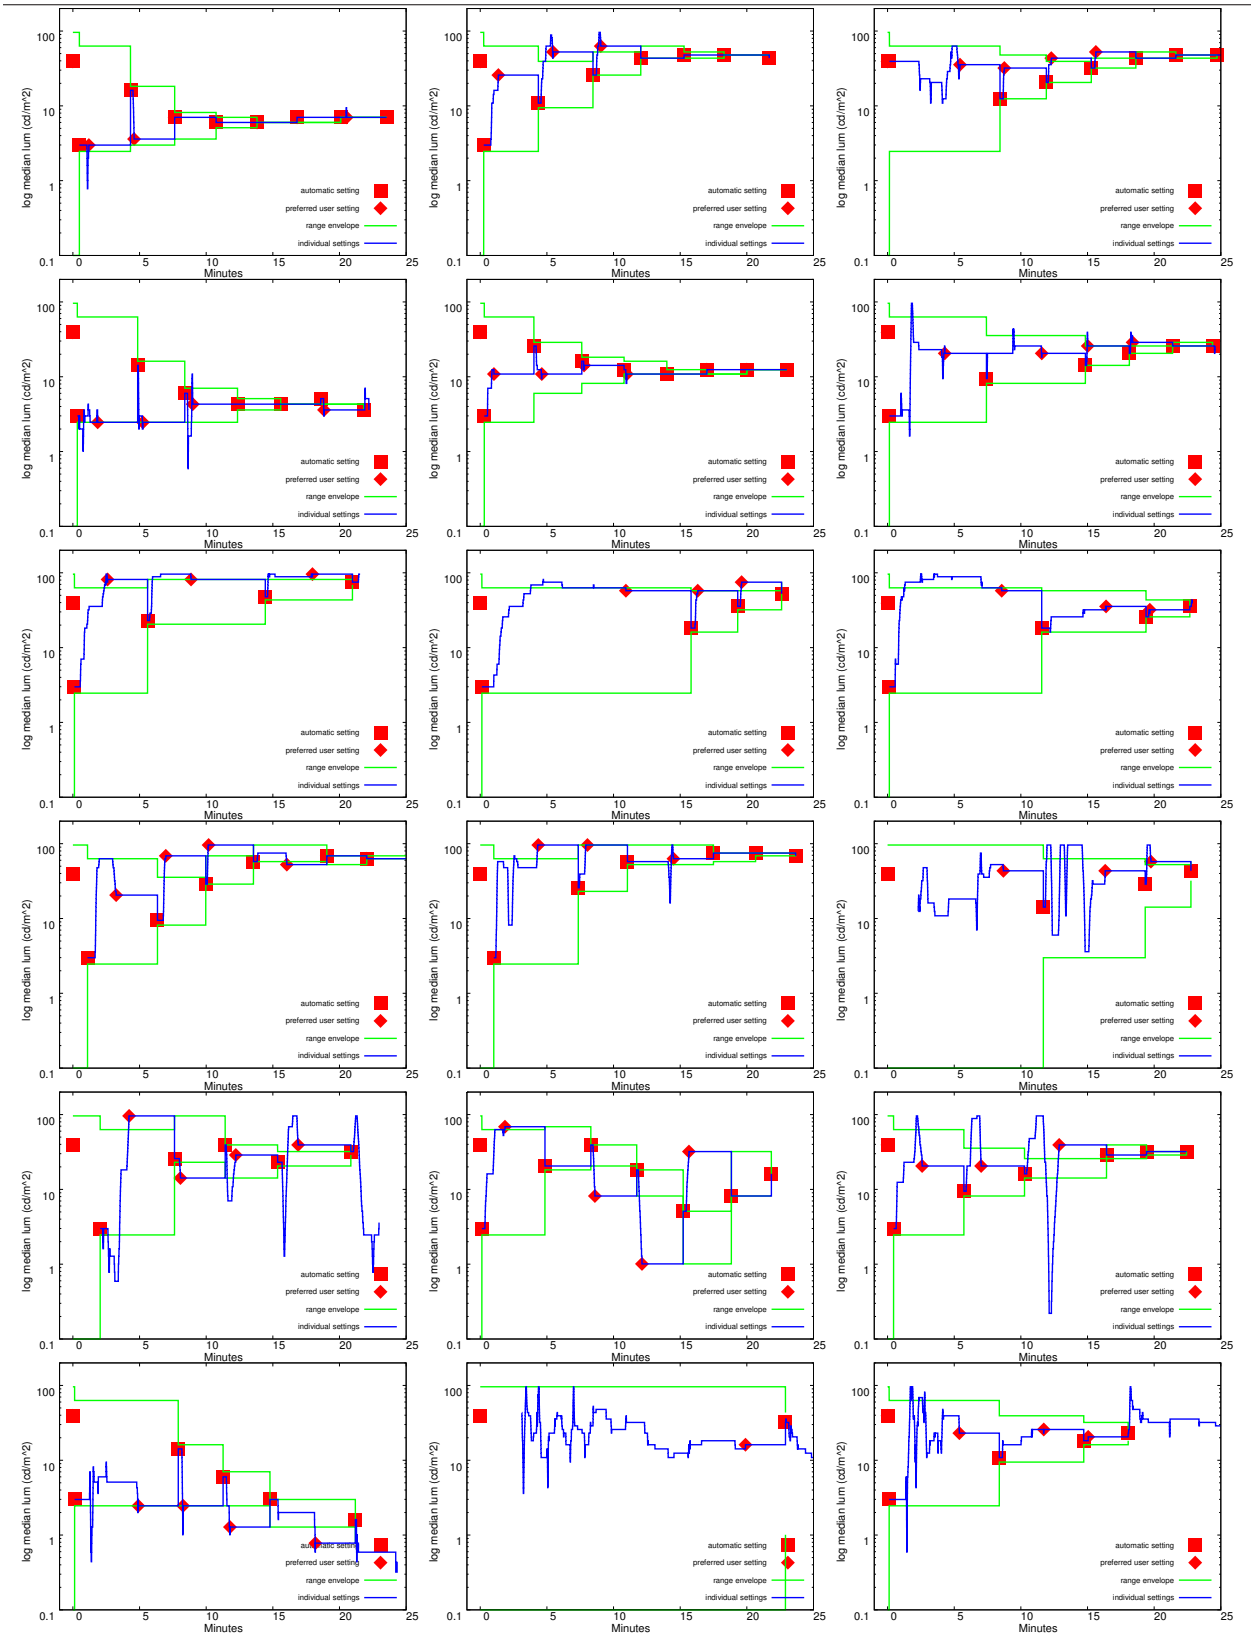

Figure 3.9: Automatic and user-selected brightness settings for 7 (left), 70 (center), and 700 lux (right), for the first 6 subjects in Experiment 3.2.

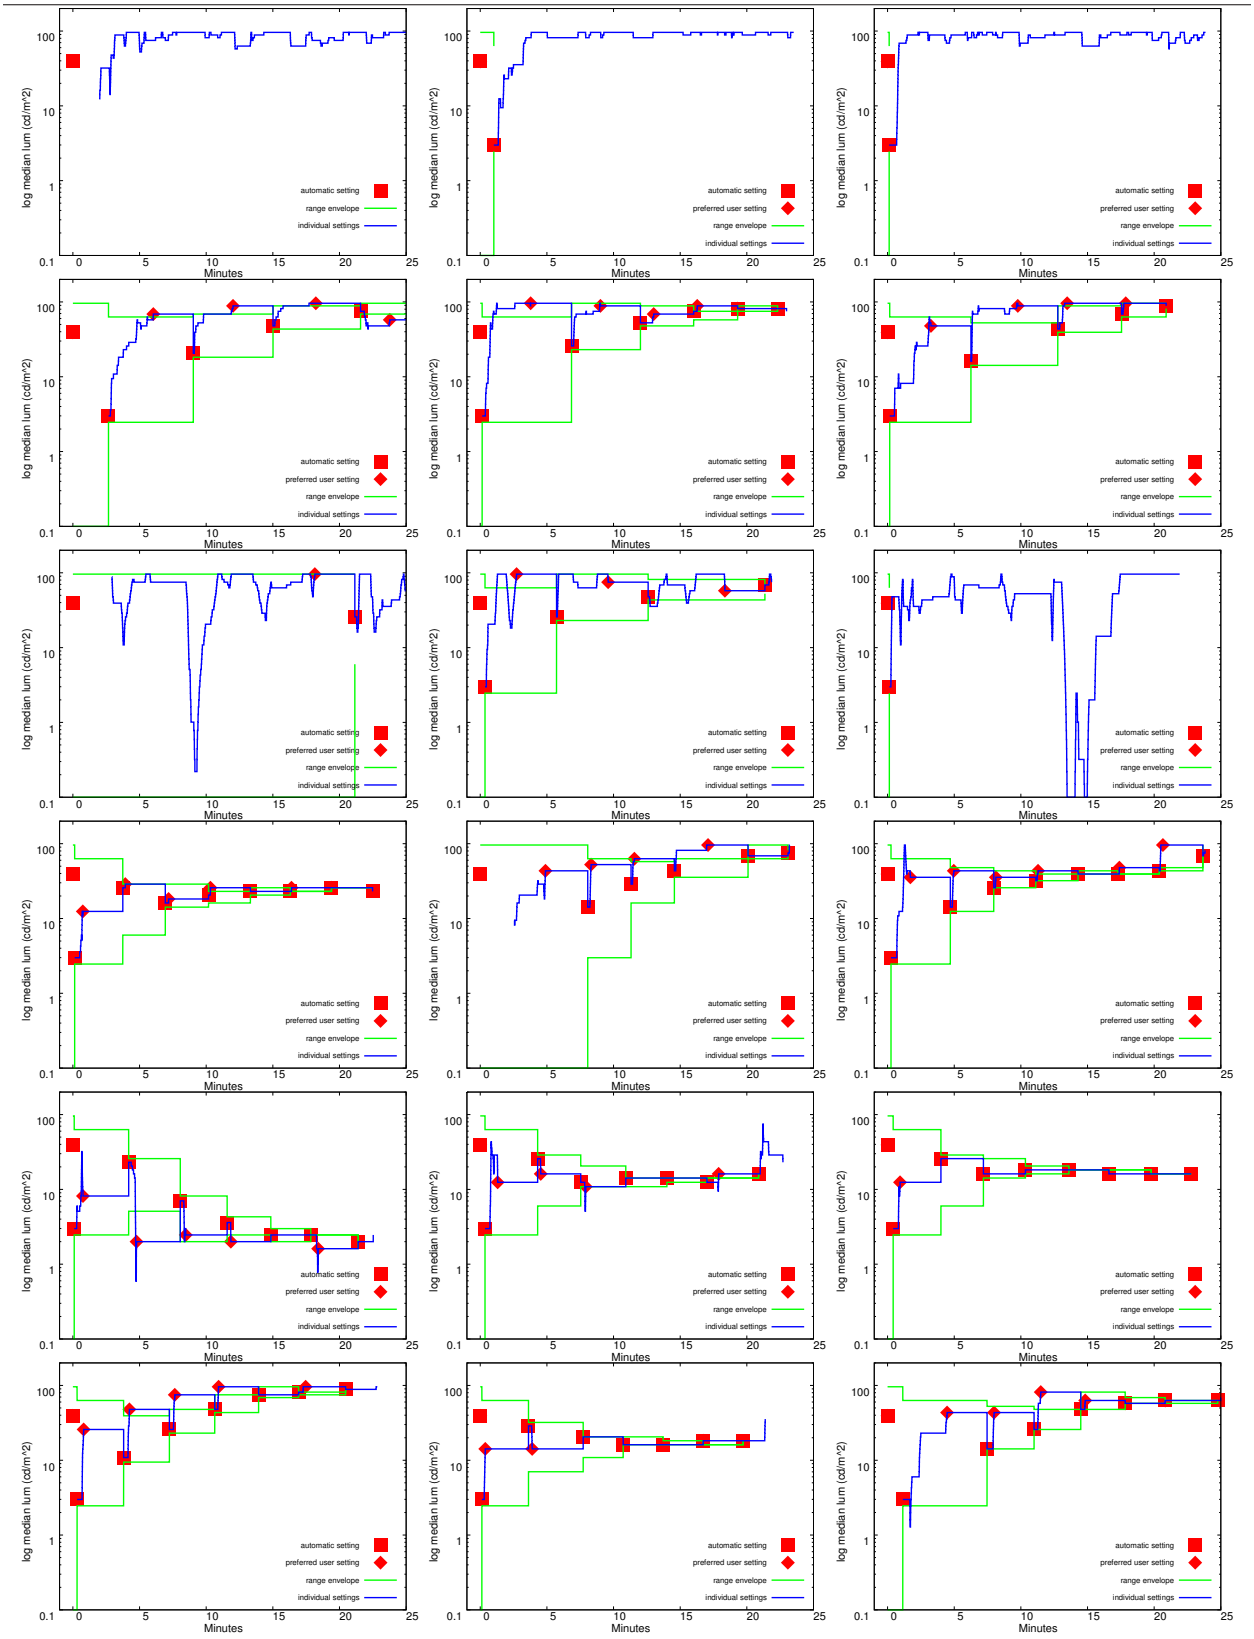

Figure 3.10: Automatic and user-selected brightness settings for 7 (left), 70 (center), and 700 lux (right), for the second 6 subjects in Experiment 3.2.

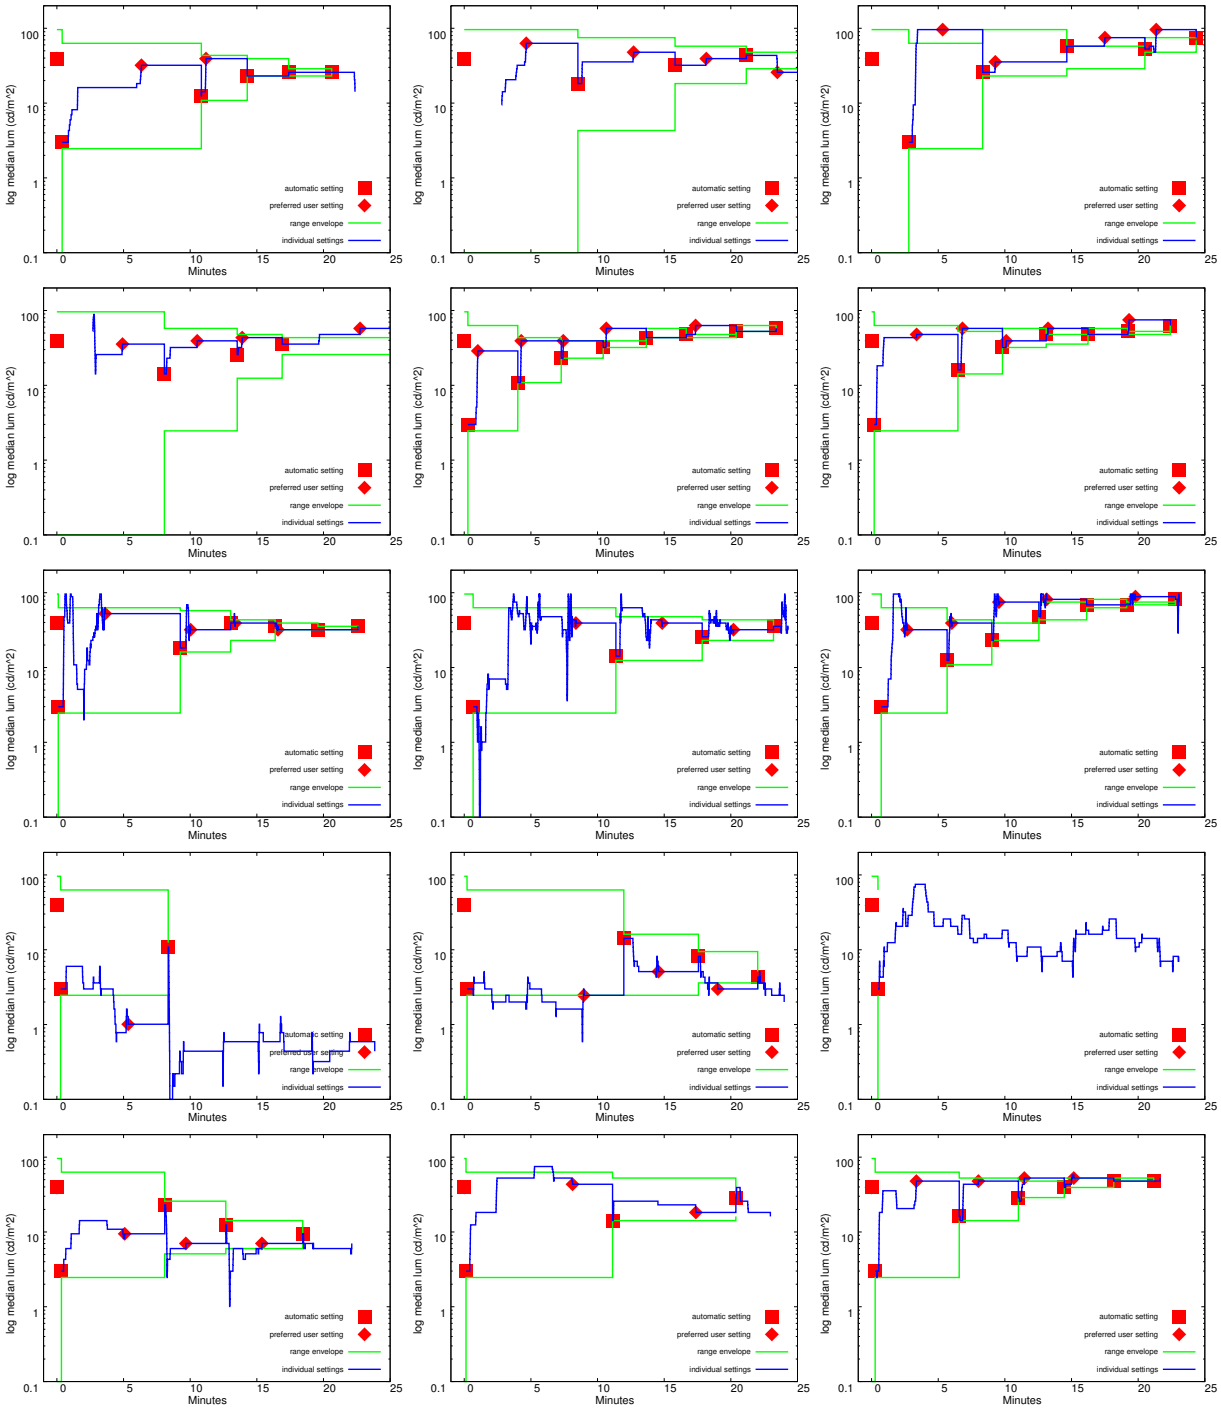

Figure 3.11: Automatic and user-selected brightness settings for 7 (left), 70 (center), and 700 lux (right), for the last 5 subjects in Experiment 3.2.

Chapter 4

Dynamic Range Expansion and Ldr2Hdr

In the previous chapter, we saw that viewers preferred the maximum available contrast when watching video content for entertainment on high dynamic range displays. In this chapter, we describe an algorithm that enhances the contrast of LDR imagery without introducing disturbing artifacts, and which is suitable for real-time implementation on GPUs or other display hardware where it can process video input streams.

Despite the increasing availability of HDR content, legacy low dynamic range (LDR) images and video will represent the majority of content in the near term. Therefore, methods must be found for displaying such content on HDR screens. Recent user studies show that perceived image quality increases with *simultaneous* increases of *both* brightness and contrast [Seet 06b, Yosh 06]. This suggests that there should be better solutions for presenting LDR content on HDR screens than by emulating LDR display hardware. Instead, we desire a method that expands the dynamic range of legacy images and video to more closely match the capabilities of the new output devices. As observed by Banterle et al. [Bant 06], this task can be thought of as a 'reverse tone mapping' operator, which takes an LDR image and estimates the HDR image it represents.

Figure 4.1: Output of our on-the-fly reverse tone mapping algorithm. Left: low dynamic range input image. Center left: a visualization of the brightness enhancement function computed by our method. Center right: two virtual exposures of the resulting HDR image with a contrast of $9300:1$. Right: the same image shown on an HDR display, using a 10% semi-transparent filter to show details in bright regions.

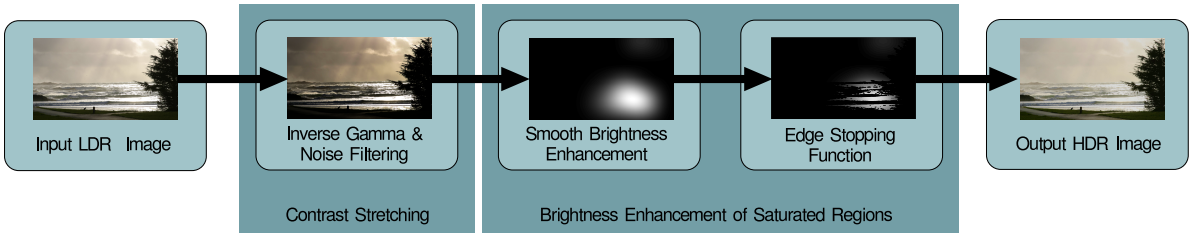

Figure 4.2: Overview of our on-the-fly reverse tone mapping method.

In most cases, it will be impossible for the reverse tone mapping to be an exact mathematical inverse: the HDR information that may have been present in the original scene was lost during capture or encoding, and thus cannot be recovered in full detail. However, it is possible to use heuristics to generate plausible images with better visual quality than the LDR input image. In this sense, reverse tone mapping is similar to other restorative techniques, including colourization of greyscale movies and images and scratch removal.

4.1 Reverse Tone Mapping

The design of our reverse tone mapping algorithm is inspired by both the image encoding process in LDR cameras, as well as the darkroom techniques of dodging and burning [Adam 83], that have also inspired Reinhard et al.'s tone mapping algorithm [Rein 02] and its inverse [Bant 06]. Dodging refers to the process of blocking light from certain image areas during the exposure of a print, while illuminating other regions. On photographic paper, which has a negative response, the dodged image regions will be reproduced in brighter tones compared to the unprocessed image regions. The sharpness of the dodging process can be controlled by adjusting the distance between the mask and the print material: smaller distances result in sharper transitions [Adam 83]. The process of burning, on the other hand, allows more light to expose regions of the negative, making those regions darker in the final print.

Our method consists of two components, as depicted in Figure 4.2. First, an inverse gamma stage maps pixel values into linear luminance in order to compensate for the non-linear encoding

of intensities in LDR images. We also apply an image filter to suppress noise, as well as quantization and image compression artifacts. The second component is the computation of a *brightness enhancement* function, which is used to increase brightness in image regions containing saturated pixels, i.e. pixels with at least one of the R, G, or B channels near the maximum value. The brightness enhancement function is mostly smooth, but may be constrained by an *edge-stopping* function that ensures that the brightness enhancement ends at strong edges. In the terminology of darkroom techniques, one can think of the smooth brightness enhancement as a dodging operation with a smooth boundary, while the edge-stopping function introduces a sharp boundary at strong edges between dark and bright regions.

Both components are designed to be temporally coherent and robust under noise, as detailed in the following.

4.1.1 Inverse Gamma and Contrast Scaling

The first step of our method is to map the non-linearly encoded pixel values from the LDR image into a linear luminance range with absolute pixel luminance values. We then map the pixel values to absolute output intensities on the HDR display. An optional image filter reduces noise and quantization artifacts.

Linearization of Pixel Values. The linearization of pixel values involves compensating for non-linearities in the LDR representation by applying the inverse of the encoding function. We are primarily interested in applying our method to video footage that has been processed for viewing on normal television sets. Standard video and television formats use a gamma curve of 2.2 [ITU 90, Stok 96, Poyn 03], which itself is meant to compensate for non-linearities in conventional display technologies such as CRTs[1]. By inverting this gamma curve we obtain pixel values that are approximately proportional to the luminance in the original scene.

Although we focus on the processing of dynamic content like video and television material,

[1]Newer display technologies such as LCD panels usually simulate a gamma curve comparable to that of CRTs for backwards compatibility.

we also experimented with digital photographs. Unlike most digital video sources, digital photo cameras often apply additional non-linear transformations on the typically linear raw sensor measurements. The goal of these transformations is to simulate response curves similar to those of analog film [Adam 83], as well as to compress the contrast of digital cameras (often 1 000 : 1 or more) to that of conventional screens ($300:1-400:1$). To account for the additional non-linearities in digital photographs, one could analyze the image header, which typically encodes camera type and settings. Using this information, one could choose between a set of precomputed response functions. In our experiments, we found that a simple gamma curve produces credible images without visible artifacts even for photographic material.

Contrast Scaling. Once we have linearized the pixel values, we map them to absolute luminance values to be shown on the HDR display. On conventional LDR screens, the contrast of the imagery is usually stretched to the full capabilities of the display device, subject to user preference. As a consequence, the same content can be presented with a contrast of 150 : 1 on one display, and 500 : 1 on another display; the contrast is simply uniformly stretched over the feasible range of the output device.

This observation suggests that we can increase the dynamic range for presentation on HDR screens simply by further linear stretching of the image contrast. A recent perceptual study independently performed by Akyüz et al. [Akyu 07] confirms this concept. However, we find that this stretch must not be too large, otherwise images will start looking unnatural. In particular, bright objects may appear to 'glow' in certain situations for large increases in contrast.

We tested a large number of images with different scaling factors, and determined among ourselves that stretching the contrast up to about 5 000 : 1 yields an improved image quality without introducing artifacts. This threshold is conservative: for many images much larger scale factors produce outstanding results. However, above this threshold, a small percentage of images will start degrading in visual quality. Since we desire a robust and temporally coherent algorithm without the need for user controlled parameters, we select a conservative mapping. On our reference HDR display, we fix the black level of the image at about 0.3 cd/m^2, which provides a deep black under

normal viewing conditions. We then scale the white point to 1 200 cd/m^2, significantly brighter than regular displays. These numbers may have to be adjusted for other HDR displays, but they are independent of the images, thereby ensuring temporal coherence without user input.

Noise Filtering and Quantization Reduction. The contrast stretching and non-linear mapping of pixel values amplify quantization artifacts and noise. LDR input images are usually quantized to 256 pixel values, while over 1 000 different values are required to cover the dynamic range of HDR displays at the precision of Just Noticeable Difference (JND) steps [Seet 04]. Lossy video compression can further reduce the number of available intensity levels in a local image region.

In the past, sophisticated algorithms for addressing such artifacts have been developed, for example the method by Daly and Feng [Daly 04]. Due to our real-time constraints, we employ a simple modified spatial Bilateral filter [Toma 98] to alleviate this situation without blurring sharp features. We use Gaussian functions for both the spatial and the photometric term in the Bilateral filter. However, we stretch the variance of the photometric term linearly with the stretch introduced by the non-linear intensity mapping for the local pixel value, such that the photometric variance is always 2 quantization levels. The effect of this approach is similar to performing the Bilateral filter with uniform variance *before* the intensity mapping. However, our ordering makes better use of the available bit depth (bits per pixel) so that the operations can be performed in fixed point arithmetic. Despite this optimization, the Bilateral filter remains an expensive operation, so that its spatial extent needs to be kept to around 4 pixels, depending on the computing platform.

4.1.2 Brightness Enhancement

The second component of our method is a brightness enhancement function, which we introduce with the goal of increasing the luminance of the output image in regions where at least one colour channel is saturated. In our work, pixels are said to be *saturated* if their values are at or above the level considered to represent the highest meaningful level for that image format. In these regions, information has probably been lost because the scene intensity was outside the capabilities of the camera or recording medium at the time of capture. While this information cannot be restored

exactly, we can attempt to approximate the visceral response associated with the higher contrast and overall brightness in the original scene.

Our approach for achieving this goal is to compute a function that produces a monochromatic (white) image that can be multiplied with the linearized intensity image to form a high dynamic range image. This has to be done carefully in order to avoid introducing spatial or temporal artifacts, as described below. Our brightness enhancement function is primarily smooth, but may contain sharp edges in areas of strong image gradients in the original image. The result is an increase in brightness not just for pixels with saturated colour channels, but for a whole neighbourhood surrounding such pixels, much like the traditional darkroom technique of dodging with a mask positioned far from the photographic paper.

Smooth Brightness Enhancement. The smooth portion of the brightness enhancement function is computed by first determining a binary mask of pixels where at least one colour channel exceeds a certain threshold intensity. Video formats typically use a white level of 235, meaning that full white of reflective objects corresponds to that pixel value [ITU 90]. However, video streams also contain larger, 'super-saturated' pixel values corresponding to specular highlights or light sources. We found that using a threshold value of 230 works well for separating the saturated from the super-saturated regions in the presence of lossy video compression. For photographs, we found a threshold of 254 to be adequate in the presence of artifacts introduced by lossy compression.

From this mask, we derive a smooth brightness enhancement function by blurring the mask with a large kernel of approximately Gaussian shape. The exact size of the blur kernel depends on the display dimensions and anticipated range of viewing distance. For example, on a $37''$ (94 cm diagonal) HDR display with a resolution of 1920×1080 pixels, we use a blur with standard deviation of 150 pixels, which corresponds to 1.2° at a viewing distance of 3 m. As a result, the spectrum of the blur filter primarily contains low angular frequencies of 0.5 cycles per degree or less, to which the human visual system is not very sensitive [van 67, Patt 98]. By designing the brightness enhancement function in this way, we make sure that it can be used to increase the intensity in the neighbourhood of saturated image regions without introducing visible artifacts.

These parameters apply specifically to the Dolby DR-37P which has been the dominant display in HDR viewing environments. However, we can generalize this technique to displays of arbitrary size, resolution, and viewing distance keeping the standard deviation of the blur at or above 1.2° of visual angle to keep the angular frequencies of the blur filter at low levels to which the human visual system is not very sensitive.

We apply the smooth brightness enhancement function by linearly mapping its values to a range of $[1 \dots \alpha]$ and then multiplying it onto the result from the inverse gamma stage. The value of the brightness amplification factor α is chosen based on the capabilities of the target HDR display. In our experiments with the Brightside DR37-P, we used a value of 4, corresponding to a peak intensity of $4 \cdot 1\ 200 = 4\ 800\ \mathrm{cd/m^2}$. Due to the large blur radius, this peak intensity is only reached for large saturated regions.

Edge-Stopping Function. The smooth brightness enhancement function by itself stretches the global contrast, and yields images that appear more crisp than the stretched contrast images when viewed on an HDR display. However, this function cannot enhance local contrast around sharp edges. To further improve appearance under such conditions, we introduce an edge-stopping function that limits the influence of the brightness enhancement to image regions that are not separated by strong edges from the saturated pixels.

We compute this binary edge stopping function using a flood fill algorithm that uses the initial binary saturation mask as a seed, as follows. The flood fill proceeds outwards from these saturated pixels until it reaches pixels with a large gradient magnitude, or the boundary of the area of influence for the smooth brightness enhancement function. We estimate image gradients using divided differences, but for robustness we use a wide baseline of 5 pixels, such that we obtain thick edges that reliably prevent the flood fill algorithm from leaking across the edges.

The final edge stopping function is cleaned up with a morphological *open* operator and blurred slightly to suppress aliasing, before it is multiplied with the smooth brightness enhancement function. The *open* operator consists of a combination of erosion and dilation operators that erodes the sizes of holes in the saturation extension area and then dilates those holes that survive, in order

to preserve the strong edges that should persist while reducing the likelihood of weak or intermittent edges creating possible temporal artifacts. As before, the result is linearly mapped to a range of $[1 \ldots \alpha]$, and then multiplied onto the colour image. Figure 4.3 shows an LDR input image, the smooth brightness enhancement function, and that same function including the edge stopping term.

Implementation Using Image Pyramids. Although the brightness enhancement algorithm calls for an image blur with a large radius, it can be implemented in a highly efficient manner using image pyramids [Burt 83]. Figure 4.4 illustrates this hierarchical version of the algorithm. The large Gaussian blur is implemented by successively downsampling the mask representing the saturated pixel regions (1), and then upsampling it again with nearest-neighbour interpolation, while applying a small (3×3 pixel) Gaussian blur at each level (2). To compute the edge stopping function, we first generate the edge image at the highest resolution as described above, and then downsample into an image pyramid (3). The actual edge stopping function is then created through a sequence of upsampling and dilation operations from the lowest resolution binary saturation mask (4). As with the smooth brightness enhancement function, the upsampling for the edge stopping function uses nearest neighbour interpolation. The dilation operators use a 3×3 binary mask (4), but stop at pixels that are marked as an edge in the edge image of the corresponding resolution (3).

Note that the radius of the dilation at each level (4) is the same as that of the blur on the corresponding level of the upsampling pyramid (2), so that the blur and the edge stopping function propagate outwards at the same speed. Note also that the edge stopping function can have hard edges in smooth image regions. However, these edges are outside the area of influence of the blur function, and thus do not create discontinuities in the final image.

4.2 Evaluation and Results

The algorithms presented in this chapter are based on experiments with a large set of sample images, verified by viewing the resulting images on a commercial HDR display, the Brightside DR37-P.

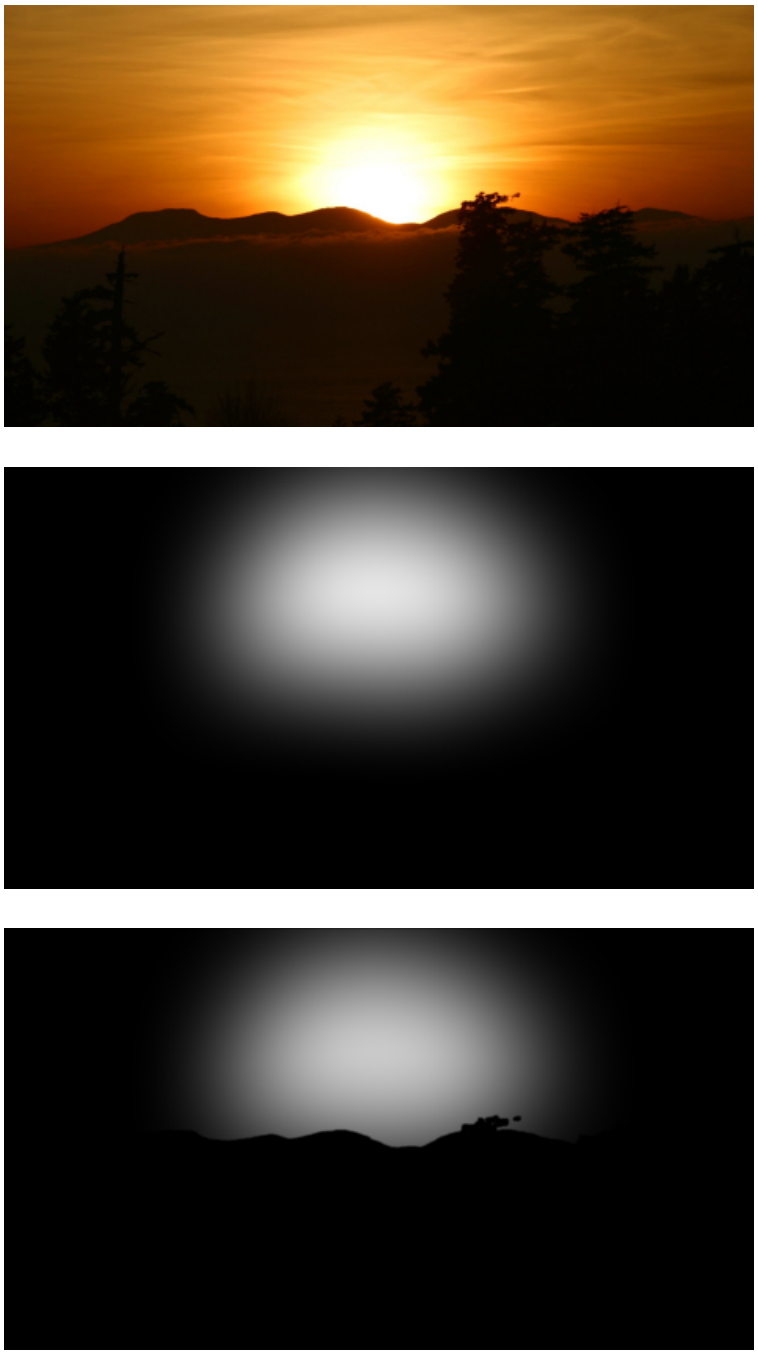

Figure 4.3: Smooth brightness enhancement function (center) and brightness enhancement with edge stopping function (bottom) for an input image (top).

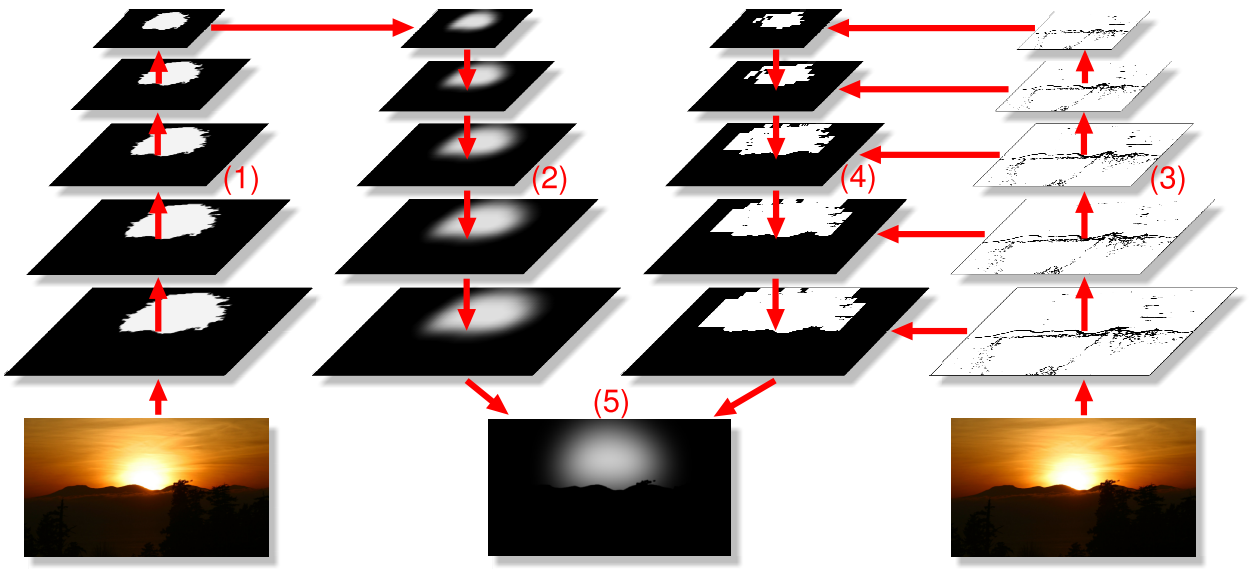

Figure 4.4: Diagram of the brightness enhancement algorithm using image pyramids (see text for details).

For a quantitative analysis we applied the visible difference predicator (VDP) [Mant 04b], an HDR variant of the original VDP work by Daly et al. [Daly 93]. The VDP by Mantiuk et al. computes the likelihood of a human observer detecting the difference between two images, based on a side-by-side comparison of individual image regions. In contrast to other image metrics, the HDR VDP accounts for HDR-specific effects such as veiling glare which is the glare of bright regions that obscures nearby dark regions. The performance of the VDP has been demonstrated with a perceptual study [Mant 04b].

Since the VDP assumes adaptation to the individual image regions, it is possible to directly compare the linearized input LDR image to the HDR image created by reverse tone mapping. In this fashion, we can validate that the saturation extension does not introduce artifacts. Figure 4.5 shows some examples of the comparison, where image regions have been coloured according to their detection probability.

Note that the VDP is not suitable for identifying the exact shape of the differences; it simply indicates regions in which problems may occur. As expected, the probability of detection for the smooth saturation extension is low, even in direct side-by-side comparison, as measured by the

Figure 4.5: Coloured results of the VDP comparison between LDR and HDR image.The red areas show regions that contain pixels that have a 25% probability of being detected as different.

VDP. On the other hand, the edge stopping function does enhance the local contrast, which can be detectable in side-by-side comparisons, with detection probabilities of 25% or more for 0.27% of the pixels. Visual inspection shows that these local contrast enhancements do not negatively affect image quality, but in many cases actually improve it.

The individual stages of the algorithm are designed for robustness and temporal coherence. For example, while we use thresholding of the image to determine a saturation mask, this mask is aggressively low-pass filtered before use as a brightness enhancement function. This approach not only ensures that the spatial frequencies are not disturbing to humans, but it also makes the method extremely stable under temporal fluctuations caused by noise or compression artifacts.

For the edge stopping function, the wide baseline gradient estimation proves robust under noise and thus ensures that the edge information is effective at stopping the flood fill algorithm. The specific edge shape varies with the chosen edge threshold, but both our visual evaluation and the experiments with the VDP indicate that the precise shape is perceptually masked by significant veiling glare. The key characteristic of the edges is temporal coherence, which we ensure with the robust edge thresholding.

Since the brightness enhancement function is monochromatic, we always reproduce the colours from the LDR image. In the future, it may be interesting to develop methods that adapt the colour in saturated regions when different colour channels saturate in different regions. However, this is a potentially unstable operation, which is why we opted to maintain the colours from the LDR image instead.

In Figure 4.6 we show a subset of the test images we used in our comparison. Due to limitations of the print medium, we cannot present the generated HDR imagery here directly. Instead, we show split-image representations with two virtual exposures, as well as images tone-mapped with Reinhard et al.'s photographic tone mapping operator [Rein 02]. The split-image representation shows the contrast in the output images, while the tone-mapped representation is another way of verifying that our method does not introduce artifacts, since high spatial frequencies introduced by our algorithm would be preserved by this operator.

4.3 Discussion

In this chapter we have presented a method for on-the-fly expansion of the dynamic range of legacy, low dynamic range, video content for viewing on HDR displays. The method is robust and temporally coherent, and does not require image-specific parameter adjustment. As such the method is well-suited for integration directly in HDR display hardware, where it can be used to process video streams from legacy sources such as television or DVDs.

We have focused on the space of real-time solutions to the problem of expanding the dynamic range of LDR imagery. More sophisticated, albeit slower methods would be interesting to explore in the future.

The Ldr2Hdr algorithm has also had significant industrial impact. As part of my graduate work under a NSERC-MITACS Industrial Postgraduate Scholarship (IPS), I wrote implementations of the algorithm for our industrial partner, Dolby Canada. We also collaborated with Dolby on the filing of patents for the technique [Remp 08a, Remp 08b, Remp 08c].

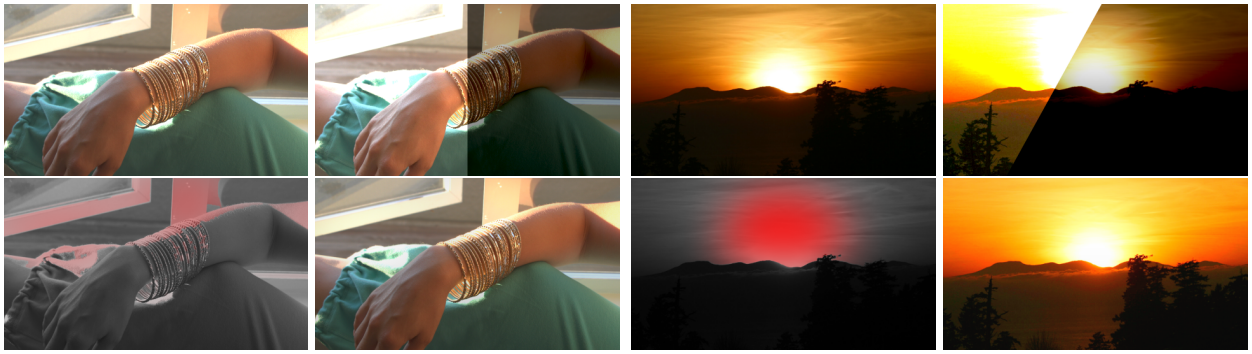

(a) (b)

(c) (d)

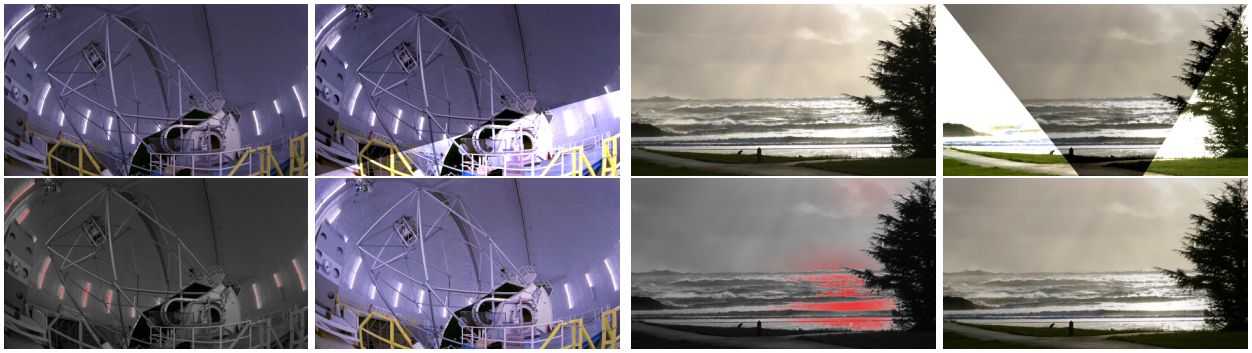

(e) (f)

Figure 4.6: Examples of our algorithm. In each group, we have in clockwise order, from the top left: the original LDR image; a split-screen representation of the HDR output image; a tone-mapped version of the HDR output (using Reinhard et al.'s operator [2002]); and a visualization of the brightness enhancement function overlaid with a grayscale version of the image.

Chapter 5

Display Considerations for Night and Low-Illumination Viewing

In the previous two chapters, we looked at how to adapt existing LDR content for display on HDR devices in order to provide more satisfying viewing experiences. Now we turn our attention to how HDR displays have given us new ways to examine perceptual characteristics of the human visual system. The results of this work can be used both for the development of new models of human visual perception and also to guide the development of future generations of displays.

Most digital displays are primarily designed to be used under office lighting. This is one reason why the prevalent colour scheme uses black letters on a white background, which reduces the impact of ambient light reflections on the screen. In order to minimize the strain caused by frequent adaptation changes between the display and environment, the peak luminance of displays is designed to closely match the luminance of a diffuse reflective white colour under normal office lighting. These steps make the displays easier to use in bright environments, but at the same time make them less usable in dark environments, in which we need to rely on our night vision. Unfortunately displays are rarely optimized for viewing under such conditions. As a result, most mobile displays can be quite unpleasant to use, if not dazzling, at night.

When a display is intended to be at the center of a viewer's attention, for example when watching a movie or playing a video game, the display brightness is usually set high, in order to extend the perceived colour gamut and therefore improve image quality. For such applications, retaining good vision of the environment where the display is used is not necessary. In this chapter we focus on another group of applications, in which retaining good vision outside the display is essential. Such applications may include navigation system displays used while driving, mobile phones used

at night, cockpit displays, monitoring instruments, electronic book readers, and augmented telescopes. We consider the performance of a display that should be the least obtrusive when used in the dark. Such a display may need to be dim to reduce adaptation strain, disability and discomfort glare, but at the same time it should be bright enough to be legible.

The two common approaches for improving display usability in the dark are backlight dimming and the use of colour schemes that reduce emitted light. Some devices that employ the first approach are equipped with a light sensor that can detect dark environments and dim the display backlight accordingly. The second approach involves changing the colour scheme to negative text polarity (bright letters on dark background) so that the least amount of light is emitted from the screen and the glare or fatigue caused by the display is reduced.

In this chapter we study another method of improving display usability in dark environments, which involves using a spectral light distribution that is perceived well by human daylight vision photoreceptors (cones), but has very little effect on night vision photoreceptors (rods). By employing long-wavelength light (amber to red), the display can be seen by the daylight vision mechanism, which offers sufficient resolution to read information, while the night vision mechanism is mostly unaffected by the light emitted from the display. As a result, the display has a smaller effect on the user's dark adaptation, causes less fatigue, and does not interfere with other tasks, such navigating around in dark environments. This property of the visual system has been well known since World War II [Hech 45, Hulb 51], and utilized in vehicles such as submarines, whose control rooms were illuminated in red to preserve the night vision necessary for lookout or use of a periscope at night [Schr 00, p.44]. Red illumination is also common in civilian aviation, to illuminate instruments, charts, and the rest of the cabin area, without diminishing night vision. This was particularly necessary for celestial navigation, which was commonly used by some pilots before the widespread use of radio navigation aids. Interestingly, cockpit displays for military aircrafts must not emit long-wavelengths because such wavelengths may interfere with night vision goggles [Abil 93].

Although the advantage of long-wavelength light at night seems to be well known, the effect has not been well quantified in the literature. This is especially true for the case in which the observed scene contains large variations in luminance levels, including a mixture of photopic (seen

by cones) and scotopic (seen by rods) stimuli. This is the case when a much brighter display is used in dark environments, and so we design our experiments to mimic this situation. The purpose of the work in this chapter is to confirm and quantify the advantages of the long-wavelength colour scheme at night and give guidelines for the design of a display that adapts to an ambient light level.

In Section 5.1 we show that displays that employ separate red, green and blue LEDs in the backlight are capable of producing desirable emission spectra for night viewing so that no other hardware modifications in the display design are necessary to implement our approach. In Experiment 5.1 (Section 5.2) we test how glare due to a chromatic light source affects night vision. The goal is to show that different colours cause different amount of veiling glare, and that in particular long-wavelengths light is the least prone to cause disability glare. In Section 5.3 we demonstrate that the measurements of the photophobic reactions to coloured light indicate that long-wavelength light also causes the least amount of the discomfort glare. The study discussed in Section 5.4 shows that red coloured letters are also more legible at low luminance levels. We further investigate the design of a display for dark environments by conducting Experiment 5.2 on the range of preferable display brightness settings (Section 5.5). Our results can serve as a guideline for selecting display brightness and colour scheme.

5.1 Display Considerations

The requirements for displays viewed in bright and dark lighting are very different. Displays intended for bright lighting must reduce the effect of ambient light that is reflected from the screen. This is usually achieved by boosting display brightness or using positive polarity of the text (dark letters on bright background) [Spen 94]. Contrary to that, displays intended for viewing under low ambient lighting must be dark to reduce disability glare and discomfort glare, which is associated with dazzling or distraction caused by a bright source of light (a display in our case) [Vos 03]. At the same time, such displays must be bright enough to activate cones, since the low acuity of rod vision would render the display useless at scotopic light levels. Reading text or the fine details of a map with low-resolution rod vision is very difficult and only possible with large magnifications.

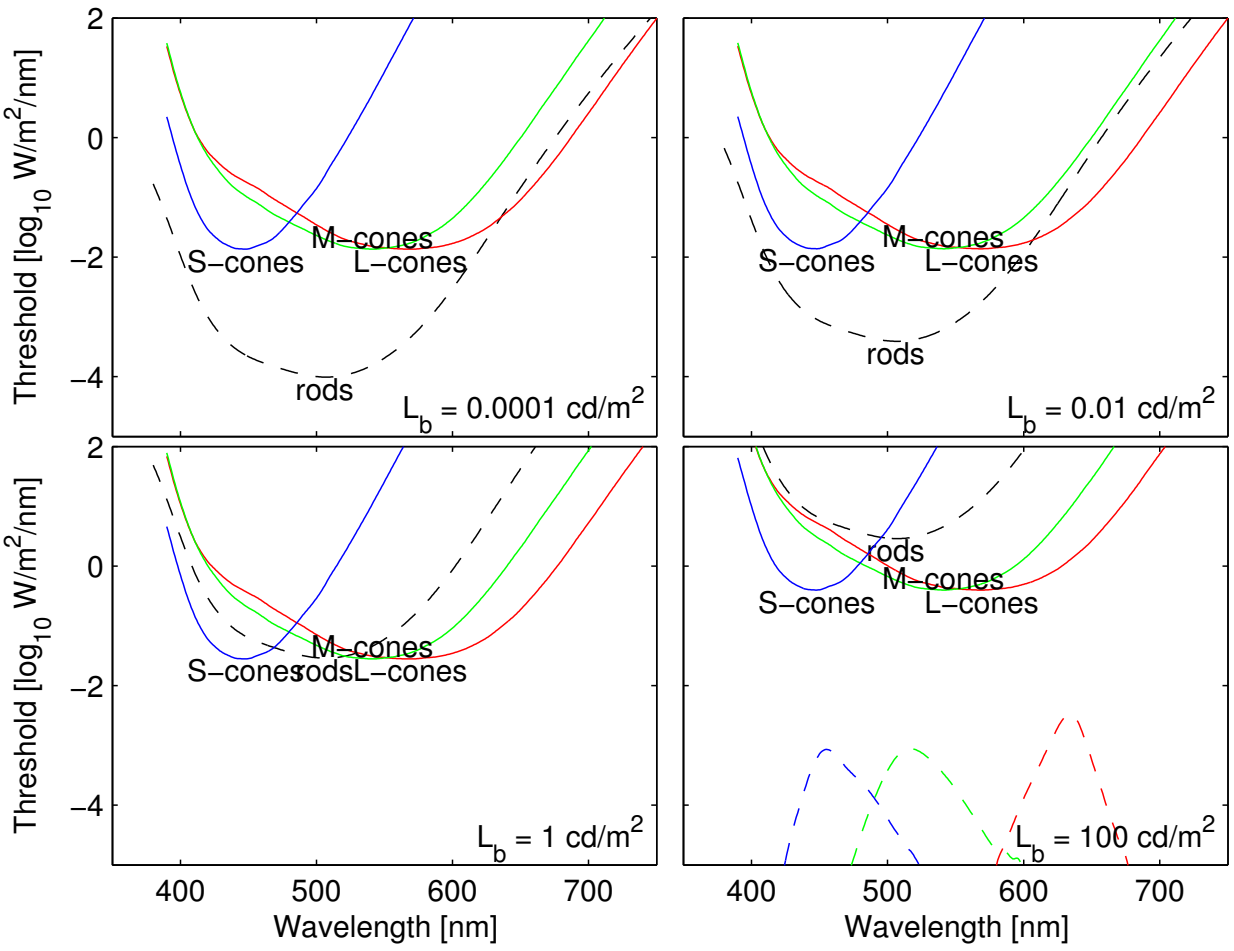

Figure 5.1: Absolute thresholds for four types of photoreceptors in the human retina. Each plot corresponds to a different background luminance L_b. Cone sensitivities are from the Stockman and Sharpe [Stoc 00] measurement and rod sensitivities are based on CIE 1951 scotopic luminosity function. Absolute thresholds are adjusted for the scotopic and photopic t.v.i. functions [Hess 90]. The lower right plot includes also the spectral emission curves for the red, green and blue LED backlight display used in the experiments.

Overall, displays viewed in the dark must minimize the effect of disability and discomfort glare while maximizing photopic luminance for legibility.

An increasing number of LCD displays use coloured (red, green and blue) light emitting diodes (LEDs) for the backlight illumination. LEDs offer several advantages over conventional cold cathode fluorescent lamp (CCFL) backlights: they result in more saturated colour primaries due to their

narrower spectral emission bands, can consume less power, can be built into thinner displays and can be dimmed to very low light levels, which is especially important for our application. This is a significant advantage over CCFL backlights, which can be dimmed to only 20 % of their peak brightness. To match the colour primaries of the ITU-R BT.709-5 (sRGB) standard, the spectral peak of the red LEDs is selected to be within the 620–640 nm range. This property, which allows the display to emit long-wavelength narrow band light, happens to make displays with LED backlights ideal for our night mode scheme. This observation suggests that a nighttime display can utilize existing LED display hardware, but with only the red primary active. In fact, we employ a prototype of such a display for our following experiments.

5.2 Experiment 5.1: Disability Glare

Since the display must emit photopic luminance levels to be legible, it is going to cause disability glare in dark environments. The disability glare is caused by intra-ocular light scatter, which is considered to be mostly wavelength-independent [Whit 93][1]. Therefore, we can expect that the wavelength effects are explained by the rod-cone sensitivity curves shown in Figure 5.1. For a dark adapted eye, rods can be over 100 times more sensitive than cones (as seen in the upper left pane in Figure 5.1). However, the cone thresholds above 610 nm are either below or very close to the rod thresholds, suggesting that rods in this wavelength range are less sensitive than cones. Therefore, a display emitting only long-wavelength light should remain legible (i.e. be seen by cones) and should reduce the undesired glare in scotopic vision (i.e. be less visible to rods).

Some studies conducted with red and blue filtered light arrive at the opposite conclusion, indicating that the spectral light composition of the source of glare has no effect on disability glare [Stee 93]. The hypothesis for our experiments, which we show in the following, is that these results are valid only for the photopic luminance range, since the target used in these experiments had a mean luminance level of $23\,\mathrm{cd/m^2}$, which was above the level of rod saturation. These ex-

[1] van den Berg et al. [Berg 91] report a small wavelength dependency of the intra-ocular light scatter caused by the selective transmission through the iris and the sclera. We do not consider chromatic aberration in the eye, which is very localized and does not cause disability glare.

periments indicated a lack of wavelength dependency on intra-ocular light scatter, but not a lack of wavelength dependency on disability glare.

Still, the rod-cone sensitivity curves (Figure 5.1) alone do not prove convincingly that long wavelength light reduces disability glare at scotopic light levels. Since rods and cones share the same visual pathways, scattered red light that is sensed by L-cones could potentially inhibit rod signals. To examine if spectral light composition has an effect on disability glare, we conducted a detection experiment in the presence of a coloured source of glare.

Stimuli. Figure 5.2 shows the setup used for this experiment, which was conducted in a dark room (< 0.1 lux) with no illumination other than that described here.

The device used to create the glare was a prototype display consisting of a backlight of coloured LEDs and a 23-inch 1920×1200 LCD panel. The display was suitable for the experiments because of its narrow-band emission spectra, shown as the dashed curves in the lower right pane of Figure 5.1, and its red LEDs emitting long-wavelength light with the peak close to 640 nm. In this experiment, the image displayed on the LCD display was a uniform field of white, while the LEDs were modulated to produce the necessary coloured stimuli.

The detection target was a 7×7 cm Gabor patch with a spatial frequency of 1.0 cycles per degree. The spatial frequency was selected to maximize sensitivity at scotopic light levels. The patch was displayed on a conventional 18-inch 1280×1024 LCD monitor. The brightness setting on the monitor was set to the minimum, and the Gabor patch was shown through a neutral density (ND) filter with an optical density of 3.0 D. The part of the monitor not showing the patch was covered with black matte paper. The baseline of the patch was set at the medium gray value and the amplitude of the patch was adjustable by keyboard control.

All the experiments in this chapter were written in Matlab using the Psychophysics Toolbox extensions [Brai 97].

Subjects. The subjects in this experiment were myself and Rafał Mantiuk, the first author of the paper [Mant 09] in which this work was first published. The use of authors as subjects is typical in many psychophysical experiments such as this which can be especially tedious. A pilot

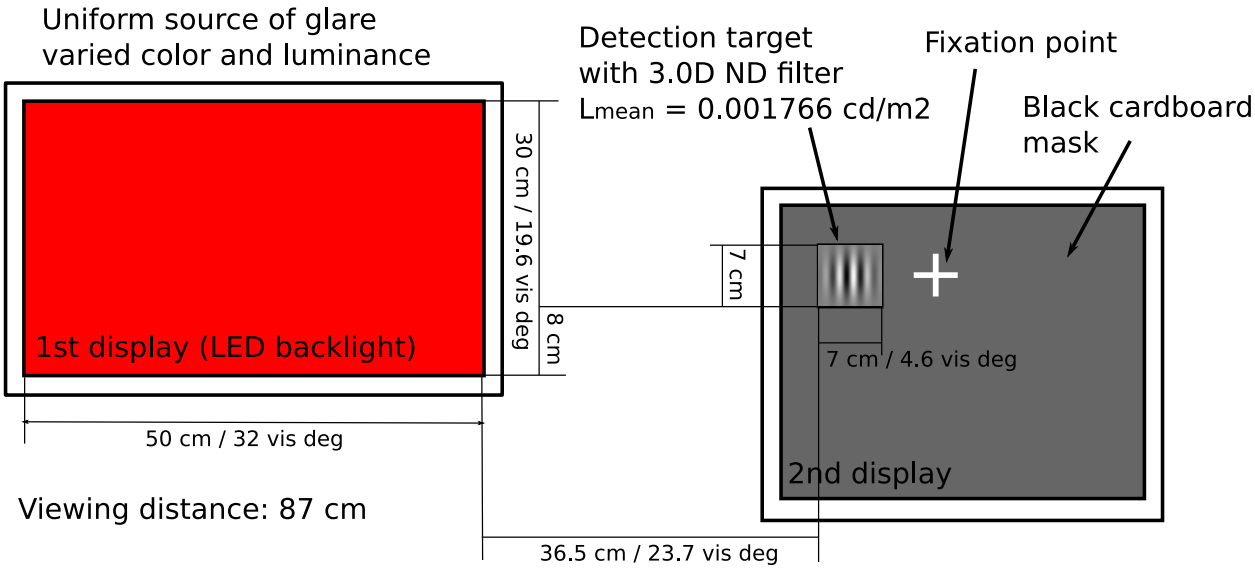

Figure 5.2: Experimental setup for Experiment 5.1.

study confirmed that the results for two subjects are consistent and repeatable. We did not find it necessary to include more subjects in this study since our goal was to measure the relative effect, rather than to estimate the mean detection thresholds for a larger population. Both subjects have normal or corrected-to-normal vision and normal colour vision.

Experimental Procedure. The task was to adjust the amplitude of the Gabor patch to a point where it was just barely visible (i.e. distinguishable from a uniform field)[2]. For each glare setting, the subjects performed this task twice, once from a point of maximum amplitude and once from a point of minimum amplitude. Prior to the experiment, the subjects spent approximately 5 minutes in the dark room in order to adapt to the scotopic environment. The subjects then performed the experiment with the source of glare turned off to measure the absolute detection threshold, and then again with the source of glare in each of 4 colours (red, green, blue, white) and set to each

[2]This technique is known as *method of adjustment*, and produced adequate results for our purposes. However, subsequent discussions about this experiment suggested that we might have been able to obtain more accurate results with a different technique, such as indicating which of a set of images was seen to contain the most visible Gabor patch. This could also have made the experiment less tedious which might have enabled us to use a larger subject pool, further improving the accuracy and strength of the results.

of 4 intensity levels, given in Table 5.1, presented in order of increasing luminance to minimize adaptation effects, for a total of 17 trials.

Levels	1	2	3	4
Red	2.9	6.47	13.5	29.4
Green	3.2	7.08	14.8	32.2
Blue	0.35	0.77	1.60	3.48
White	4.4	9.59	20.1	43.6

Table 5.1: Luminances of 4 colours of backlight at 4 intensity levels in Experiment 5.1. The levels are approximately evenly spaced on a log scale to provide perceptually uniform differences in brightness.

Results. Figure 5.3 shows the detection thresholds as a function of the scotopic and photopic luminance of the glare source. The left plot shows that the disability glare is wavelength dependent, as different colours of glare result in different thresholds at the same photopic luminance level. The long-wavelength light (red) results in the lowest detection threshold for a given photopic luminance. Therefore, a display that emits long-wavelength red will cause the least scotopic glare compared to other colours at the same photopic luminance level. Comparison of both plots shows that the detection thresholds across all colours are more strongly correlated with the scotopic rather than photopic luminance of the glare. Therefore, neural interactions between rod and cone signals have only minimal effect on the disability glare in the scotopic luminance range. In the context of display design it means that the displays intended for viewing in the dark should minimize scotopic luminance to reduce glare while maximizing photopic luminance to improve legibility.

5.3 Discomfort Glare

Discomfort glare is observed when a bright source of light evokes a dazzling effect that causes the eyes to squint or avert (sometimes referred as a dazzling glare), or distraction by a bright source of light [Vos 03]. Experiment 5.1 confirmed that long-wavelength red induces the least amount of scotopic disability glare. But does it also reduce discomfort glare? The subjective measures of

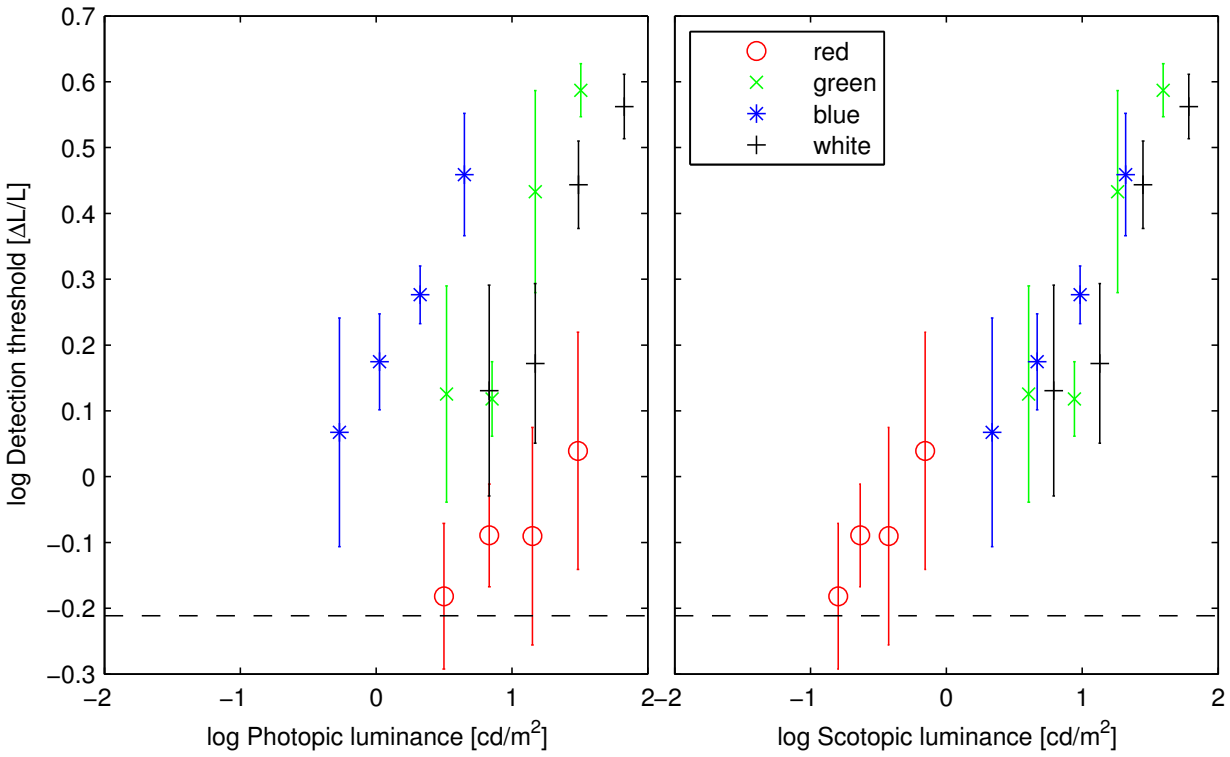

Figure 5.3: Detection thresholds in the presence of glare. Colours indicate different colours (spectral composition, shown in Figure 5.1 lower right) of the glare sources and the x-axis represents its photopic (left) and scotopic (right) luminance, computed using the respective luminosity efficiency functions. The dashed line represents the absolute detection threshold without a glare source.

discomfort glare, based on the 9-point Boer glare rating scale (from unnoticeable to unbearable), are usually very noisy and are not good at describing the wavelength dependency of the discomfort glare [Feke 06]. Therefore, we refer to the data from more accurate objective measurements based on electromyography, in which muscular activity associated with squinting is tested. Stringham et al. [Stri 03] measured the action spectrum of photophobia, which is a similar condition to dazzling glare. In Figure 5.4 we plot their results represented as the photophobia sensitivity P together with the photopic luminous efficiency function V and the ratio V/P. The ratio peaks at about 620 nm, suggesting that the long wavelength red produces the least amount of discomfort glare at the same luminance level. This result must, however, be interpreted with care, since the data for narrow

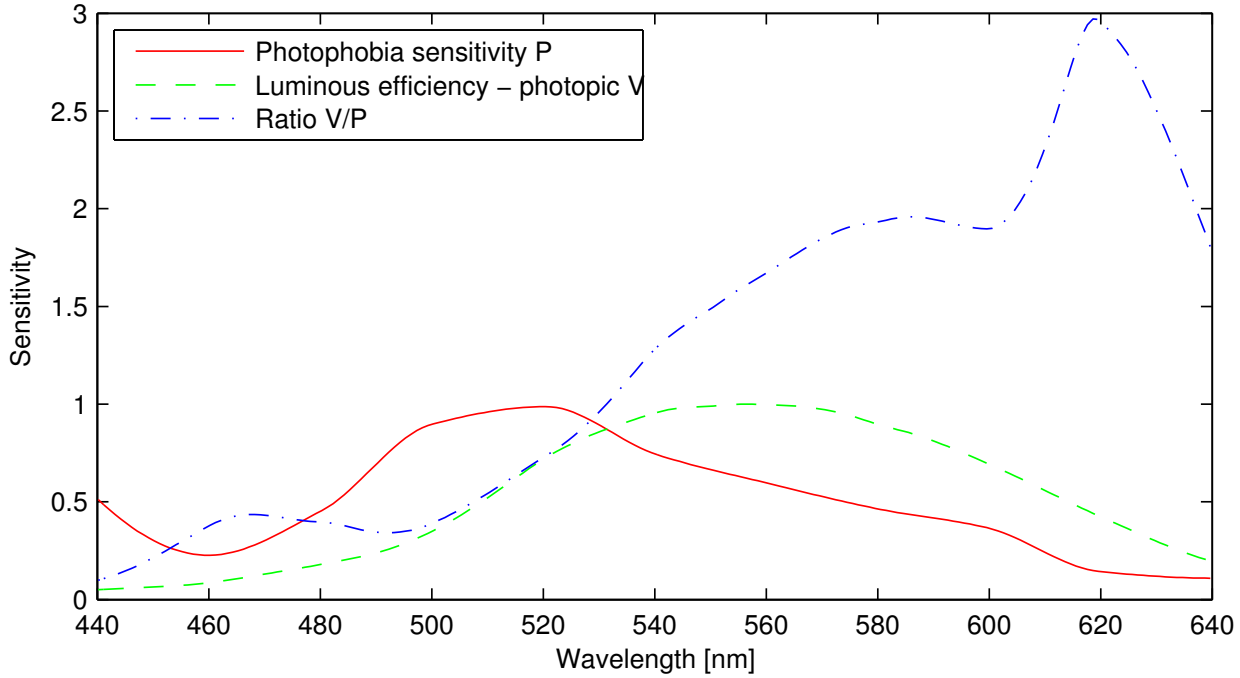

Figure 5.4: Photophobia sensitivity compared with the photopic luminous efficiency function. Wavelengths close to 620 nm evoke the least discomfort glare when adjusted for their photopic luminance level. Photophobia sensitivity data from [Stri 03].

band stimuli, used to measure the action spectrum of photophobia, is not necessarily valid for broad-band glare sources.

5.4 Recognition Rate

Okabayashi et al. [Okab 06] measured the legibility of Snellen figures shown on a CRT display at low luminance levels (0.3–10 $\mathrm{cd/m^2}$) and under low ambient light (2 lux). In Figure 5.5 we plot their data showing character recognition rate for the three primary colours of the display. In addition to photopic luminance, we found the corresponding scotopic luminance values assuming typical CRT phosphor emission spectra. Their results indicate that red figures had a higher rate of correct responses than green and blue figures, even when shown at the same photopic luminance level. Their results bring forward another argument showing that long-wavelength light is more suitable for displaying information at low luminance levels.

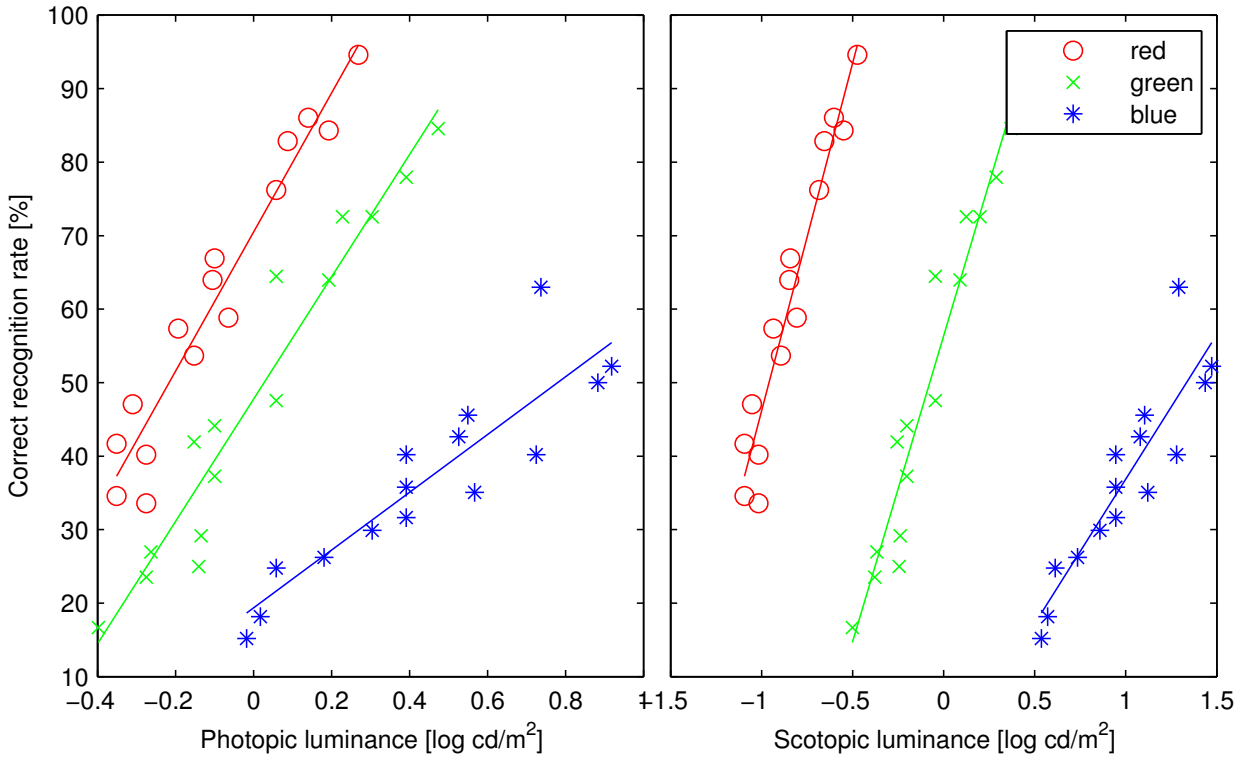

Figure 5.5: Character recognition rate for colour Snellen figures shown on a CRT monitor. Red colour gives the highest recognition rate at the same luminance level as green and blue colours. Data from [Okab 06].

5.5 Experiment 5.2: Preferred Brightness

After demonstrating that long-wavelength light is the least obtrusive for night vision, the next question to answer is that of how bright the display intended for viewing in the dark should be. We measured the preferred brightness (backlight intensity) in the following experiment.

Stimuli. This experiment used one display, which is the same prototype display used for the previous experiment. A monochrome full-screen image was displayed on the LCD panel, while subjects used keyboard controls to adjust the intensity of the LED backlight to a level they found preferable. The backlight could brighten or dim the display from $\approx$2000 cd/m^{2}to $\approx$0.5 cd/m^2. Two images were used for this experiment, both in their native (positive) black-on-white polarity (Figure 5.6) and also in negative white-on-black polarity (Figure 5.7).

ACT III

PROLOGUE

Enter Chorus

Chorus

Thus with imagined wing our swift scene flies
In motion of no less celerity
Than that of thought. Suppose that you have seen
The well-appointed king at Hampton pier
Embark his royalty, and his brave fleet
With silken streamers the young Phoebus fanning:
Play with your fancies, and in them behold
Upon the hempen tackle ship-boys climbing;
Hear the shrill whistle which doth order give
To sounds confused; behold the threaden sails,
Borne with the invisible and creeping wind,
Draw the huge bottoms through the furrow'd sea,
Breasting the lofty surge. O, do but think
You stand upon the ravage and behold
A city on the inconstant billows dancing;
For so appears this fleet majestical,
Holding due course to Harfleur. Follow, follow:
Grapple your minds to sternage of this navy,
And leave your England, as dead midnight still,
Guarded with grandsires, babies and old women,
Either past or not arrived to pith and puissance;
For who is he, whose chin is but enrich'd
With one appearing hair, that will not follow
These cull'd and choice-drawn cavaliers to France?
Work, work your thoughts, and therein see a siege;
Behold the ordnance on their carriages,
With fatal mouths gaping on girded Harfleur.
Suppose the ambassador from the French comes back;
Tells Harry that the king doth offer him
Katharine his daughter, and with her, to dowry,
Some petty and unprofitable dukedoms.
The offer likes not: and the nimble gunner
With linstock now the devilish cannon touches,

Alarum, and chambers go off

And down goes all before them. Still be kind,
And eke out our performance with your mind.

Exit

SCENE I. France. Before Harfleur.

Alarum. Enter KING HENRY, EXETER, BEDFORD, GLOUCESTER, and Soldiers, with scaling-ladders

KING HENRY V

Once more unto the breach, dear friends, once more;
Or close the wall up with our English dead.
In peace there's nothing so becomes a man
As modest stillness and humility:
But when the blast of war blows in our ears,
Then imitate the action of the tiger;
Stiffen the sinews, summon up the blood,
Disguise fair nature with hard-favour'd rage;
Then lend the eye a terrible aspect;
Let pry through the portage of the head
Like the brass cannon; let the brow o'erwhelm it
As fearfully as doth a galled rock
O'erhang and jutty his confounded base,
Swill'd with the wild and wasteful ocean.
Now set the teeth and stretch the nostril wide,
Hold hard the breath and bend up every spirit
To his full height. On, on, you noblest English.
Whose blood is fet from fathers of war-proof!
Fathers that, like so many Alexanders,
Have in these parts from morn till even fought
And sheathed their swords for lack of argument.
Dishonour not your mothers; now attest
That those whom you call'd fathers did beget you.
Be copy now to men of grosser blood,
And teach them how to war. And you, good yeoman,
Whose limbs were made in England, show us here
The mettle of your pasture; let us swear
That you are worth your breeding; which I doubt not;
For there is none of you so mean and base,
That hath not noble lustre in your eyes.
I see you stand like greyhounds in the slips,
Straining upon the start. The game's afoot:
Follow your spirit, and upon this charge
Cry 'God for Harry, England, and Saint George!'

Exeunt. Alarum, and chambers go off

SCENE II. The same.

Enter NYM, BARDOLPH, PISTOL, and Boy

BARDOLPH

On, on, on, on, on! to the breach, to the breach!

NYM

Pray thee, corporal, stay: the knocks are too hot;
and, for mine own part, I have not a case of lives:
the humour of it is too hot, that is the very
plain-song of it.

PISTOL

The plain-song is most just: for humours do abound:
Knocks go and come; God's vassals drop and die;
And sword and shield,
In bloody field,
Doth win immortal fame.

Boy

Would I were in an alehouse in London! I would give
all my fame for a pot of ale and safety.

PISTOL

And I:
If wishes would prevail with me,
My purpose should not fail with me,
But thither would I hie.

Boy

As duly, but not as truly,
As bird doth sing on bough.

Enter FLUELLEN

FLUELLEN

Up to the breach, you dogs! avaunt, you cullions!

Driving them forward

PISTOL

Be merciful, great duke, to men of mould.
Abate thy rage, abate thy manly rage,
Abate thy rage, great duke!
Good bawcock, bate thy rage, use lenity, sweet chuck!

NYM

These be good humours! your honour wins bad humours.

Exeunt all but Boy

Boy

As young as I am, I have observed these three
swashers. I am boy to them all three: but all they
three, though they would serve me, could not be man
to me; for indeed three such antics do not amount to
a man. For Bardolph, he is white-livered and
red-faced; by the means whereof a' faces it out, but
fights not. For Pistol, he hath a killing tongue
and a quiet sword; by the means whereof a' breaks
words, and keeps whole weapons. For Nym, he hath
heard that men of few words are the best men; and
therefore he scorns to say his prayers, lest a'
should be thought a coward: but his few bad words
are matched with as few good deeds; for a' never
broke any man's head but his own, and that was
against a post when he was drunk. They will steal
any thing, and call it purchase. Bardolph stole a
lute-case, bore it twelve leagues, and sold it for
three half pence. Nym and Bardolph are sworn
brothers in filching, and in Calais they stole a
fire-shovel. I knew by that piece of service the
men would carry coals. They would have me as
familiar with men's pockets as their gloves or their
handkerchers: which makes much against my manhood,
if I should take from another's pocket to put into
mine, for it si plan pocketing up of wrongs. I
must leave them, and seek some better service:
their villany goes against my weak stomach, and
therefore I must cast it up.

Exit

Re-enter FLUELLEN, GOWER following

GOWER

Captain Fluellen, you must come presently to the
mines; the Duke of Gloucester would speak with you.

FLUELLEN

To the mines! tell you the duke, it is not so good
to come to the mines; for, look you, the mines is
not according to the disciplines of the war: the
concavities of it is not sufficient; for, look you,
the adversary, you may discuss unto the duke, look
you, is digt himself four yard under the
countermines: by Cheshu, I think a' will plough up
all, if there is not better directions.

GOWER

The Duke of Gloucester, to whom the order of the
siege is given, is altogether directed by an
Irishman, a very valiant gentleman, i' faith.

FLUELLEN

It is Captain Macmorris, is it not?

GOWER

I think it be.

FLUELLEN

By Cheshu, he is an ass, as in the world. I will
verify as much in his beard: he has no more
directions in the true disciplines of the wars, look
you, of the Roman disciplines, than is a puppy-dog.

Enter MACMORRIS and Captain JAMY

GOWER

Here a' comes, and the Scots captain, Captain Jamy, with him.

FLUELLEN

Captain Jamy is a marvellous falourous gentleman,
that is certain; and of great expedition and
knowledge in th' aunchient wars, upon my particular
knowledge of his directions: by Cheshu, he will
maintain his argument as well as any military man in
the world, in the disciplines of the pristine wars
of the Romans.

JAMY

I say gud-day, Captain Fluellen.

FLUELLEN

God-den to your worship, good Captain James.

GOWER

Figure 5.6: The map and text images used in Experiment 5.2.

ACT III

PROLOGUE

Enter Chorus

Chorus

This with imagined wing our swift scene flies
In motion of no less celerity
Than that of thought. Suppose that you have seen
The well-appointed king at Hampton pier
Embark his royalty, and his brave fleet
With silken streamers the young Phoebus fanning:
Play with your fancies, and in them behold
Upon the hempen tackle ship-boys climbing;
Hear the shrill whistle which doth order give
To sounds confused, behold the threaden sails,
Borne with the invisible and creeping wind,
Draw the huge bottoms through the furrow'd sea,
Breasting the lofty surge. O, do but think
You stand upon the ravage and behold
A city on the inconstant billows dancing;
For so appears this fleet majestical,
Holding due course to Harfleur. Follow, follow:
Grapple your minds to sternage of this navy,
And leave your England, as dead midnight still,
Guarded with grandsires, babies and old women,
Either past or not arrived to pith and puissance;
For who is he, whose chin is but enrich'd
With one appearing hair, that will not follow
These cull'd and choice-drawn cavaliers to France?
Work, work your thoughts, and therein see a siege;
Behold the ordnance on their carriages,
With fatal mouths gaping on girded Harfleur.
Suppose the ambassador from the French comes back;
Tells Harry that the king doth offer him
Katharine his daughter, and with her, to dowry,
Some petty and unprofitable dukedoms.
The offer likes not: and the nimble gunner
With linstock now the devilish cannon touches,

Alarum, and chambers go off

And down goes all before them. Still be kind,
And eke out our performance with your mind.

Exit

SCENE I. France. Before Harfleur.

Alarum. Enter KING HENRY, EXETER, BEDFORD, GLOUCESTER, and Soldiers, with scaling-ladders

KING HENRY V

Once more unto the breach, dear friends, once more;
Or close the wall up with our English dead.
In peace there's nothing so becomes a man
As modest stillness and humility:
But when the blast of war blows in our ears,
Then imitate the action of the tiger;
Stiffen the sinews, summon up the blood,
Disguise fair nature with hard-favour'd rage;
Then lend the eye a terrible aspect;
Let pry through the portage of the head
Like the brass cannon, let the brow o'erwhelm it
As fearfully as doth a galled rock
O'erhang and jutty his confounded base,
Swill'd with the wild and wasteful ocean.
Now set the teeth and stretch the nostril wide,
Hold hard the breath and bend up every spirit
To his full height. On, on, you noblest English.
Whose blood is fet from fathers of war-proof!
Fathers that, like so many Alexanders,
Have in these parts from morn till even fought
And sheathed their swords for lack of argument:
Dishonour not your mothers; now attest
That those whom you call'd fathers did beget you.
Be copy now to men of grosser blood,
And teach them how to war. And you, good yeoman,
Whose limbs were made in England, show us here
The mettle of your pasture; let us swear
That you are worth your breeding; which I doubt not;
For there is none of you so mean and base,
That hath not noble lustre in your eyes.
I see you stand like greyhounds in the slips,
Straining upon the start. The game's afoot:
Follow your spirit, and upon this charge
Cry 'God for Harry, England, and Saint George!'

Exeunt. Alarum, and chambers go off

SCENE II. The same.

Enter NYM, BARDOLPH, PISTOL, and Boy

BARDOLPH

On, on, on, on, on! to the breach, to the breach!

NYM

Pray thee, corporal, stay: the knocks are too hot,
and, for mine own part, I have not a case of lives:
the humour of it is too hot, that is the very
plain-song of it.

PISTOL

The plain-song is most just: for humours do abound:
Knocks go and come; God's vassals drop and die;
And sword and shield,
In bloody field,
Doth win immortal fame.

Boy

Would I were in an alehouse in London! I would give
all my fame for a pot of ale and safety.

PISTOL

And I.
If wishes would prevail with me,
My purpose should not fail with me,
But thither would I hie.

Boy

As duly, but not as truly,
As bird doth sing on bough.

Enter FLUELLEN

FLUELLEN

Up to the breach, you dogs! avaunt, you cullions!

Driving them forward

PISTOL

Be merciful, great duke, to men of mould.
Abate thy rage, abate thy manly rage,
Abate thy rage, great duke!
Good bawcock, bate thy rage; use lenity, sweet chuck!

NYM

These be good humours! your honour wins bad humours.

Exeunt all but Boy

Boy

As young as I am, I have observed these three
swashers. I am boy to them all three: but all they
three, though they would serve me, could not be man
to me; for indeed three such antics do not amount to
a man. For Bardolph, he is white-livered and
red-faced; by the means whereof a' faces it out, but
fights not. For Pistol, he hath a killing tongue
and a quiet sword; by the means whereof a' breaks
words, and keeps whole weapons. For Nym, he hath
heard that men of few words are the best men, and
therefore he scorns to say his prayers, lest a'
should be thought a coward: but his few bad words
are matched with as few good deeds; for a' never
broke any man's head but his own, and that was
against a post when he was drunk. They will steal
any thing, and call it purchase. Bardolph stole a
lute-case, bore it twelve leagues, and sold it for
three half-pence. Nym and Bardolph are sworn
brothers in filching, and in Calais they stole a
fire-shovel. I knew by that piece of service the
men would carry coals. They would have me as
familiar with men's pockets as their gloves or their
handkerchers: which makes much against my manhood,
if I should take from another's pocket to put into
mine, for it is plain pocketing up of wrongs. I
must leave them, and seek some better service:
their villany goes against my weak stomach, and
therefore I must cast it up.

Exit

Re-enter FLUELLEN, GOWER following

GOWER

Captain Fluellen, you must come presently to the
mines; the Duke of Gloucester would speak with you.

FLUELLEN

To the mines! tell you the duke, it is not so good
to come to the mines; for, look you, the mines is
not according to the disciplines of the war: the
concavities of it is not sufficient; for, look you,
the adversary, you may discuss unto the duke, look
you, is digt himself four yard under the
countermines: by Cheshu, I think a' will plough up
all, if there is not better directions.

GOWER

The Duke of Gloucester, to whom the order of the
siege is given, is altogether directed by an
Irishman, a very valiant gentleman, i' faith.

FLUELLEN

It is Captain Macmorris, is it not?

GOWER

I think it be.

FLUELLEN

By Cheshu, he is an ass, as in the world. I will
verify as much in his beard: he has no more
directions in the true disciplines of the wars, look
you, of the Roman disciplines, than is a puppy-dog.

Enter MACMORRIS and Captain JAMY

GOWER

Here a' comes, and the Scots captain, Captain Jamy, with him.

FLUELLEN

Captain Jamy is a marvellous falourous gentleman,
that is certain, and of great expedition and
knowledge in th' ancient wars, upon my particular
knowledge of his directions: by Cheshu, he will
maintain his argument as well as any military man in
the world, in the disciplines of the pristine wars
of the Romans.

JAMY

I say gud-day, Captain Fluellen.

FLUELLEN

God-den to your worship, good Captain James.

GOWER

Figure 5.7: The map and text images (reverse polarities) used in Experiment 5.2.

Subjects. Nine subjects (5 male, 4 female, ranging in age from 19 to approximately 45) were used for this experiment. Each had normal or corrected-to-normal vision and normal colour vision.

Experimental Procedure. Each subject was shown each of the 4 images in Figures 5.6 and 5.7 in random order, with the backlight set to each of 4 colours (red, green, blue, and white) in random order. For each of those 16 conditions, the subject was asked to set the backlight intensity 3 times: first, to their most preferred luminance; second, to a luminance just too dark; third, to a luminance just too bright and thus uncomfortable to read.

Results. The first of each set of three preferred brightness settings was omitted from the results as the range provided by the light and dark settings provided a more reliable and less noisy measurement. The range generally contained the first setting, though the location of that setting within the range varied considerably. Figure 5.8 shows the luminance ranges from "just too dark" to "just too bright" with respect to colour, content (map or text) and polarity (positive: black on white or negative: white on black). The peak luminance levels are all in the photopic luminance range ($> 5\,\mathrm{cd/m^2}$) and relatively high for the dark environment. This suggests that the lower limit for dimming the display can be the transition point between photopic and mesopic vision. Despite large variations, typical for preference experiments, all factors showed statistical significance in the ANOVA test ($\log(L_{photopic})$ for colour $\times$ map/text $\times$ polarity; colour: $F(3,159) = 7.58$, $p < 0.01$; map/text: $F(1,159) = 4.4$; $p < 0.05$; polarity: $F(1,159) = 4.81$; $p < 0.05$). Moderate photopic luminance variations across colours may suggest that the brightness adjustments are mostly determined by photopic luminance. The luminance of the red map or red text was adjusted so that the emitted photopic luminance was comparable with other colours, while the scotopic luminance was much lower. Maps were adjusted to be brighter than text. The probable reason for this is that reading a map involves distinguishing between several luminance levels, and this task can be better performed by the more sensitive photopic vision. Positive polarity images were set to lower luminance than the negative polarity images (except white text) but the difference between both polarities was surprisingly small.

The range of preferred display luminance found in this experiment can be used as a guideline

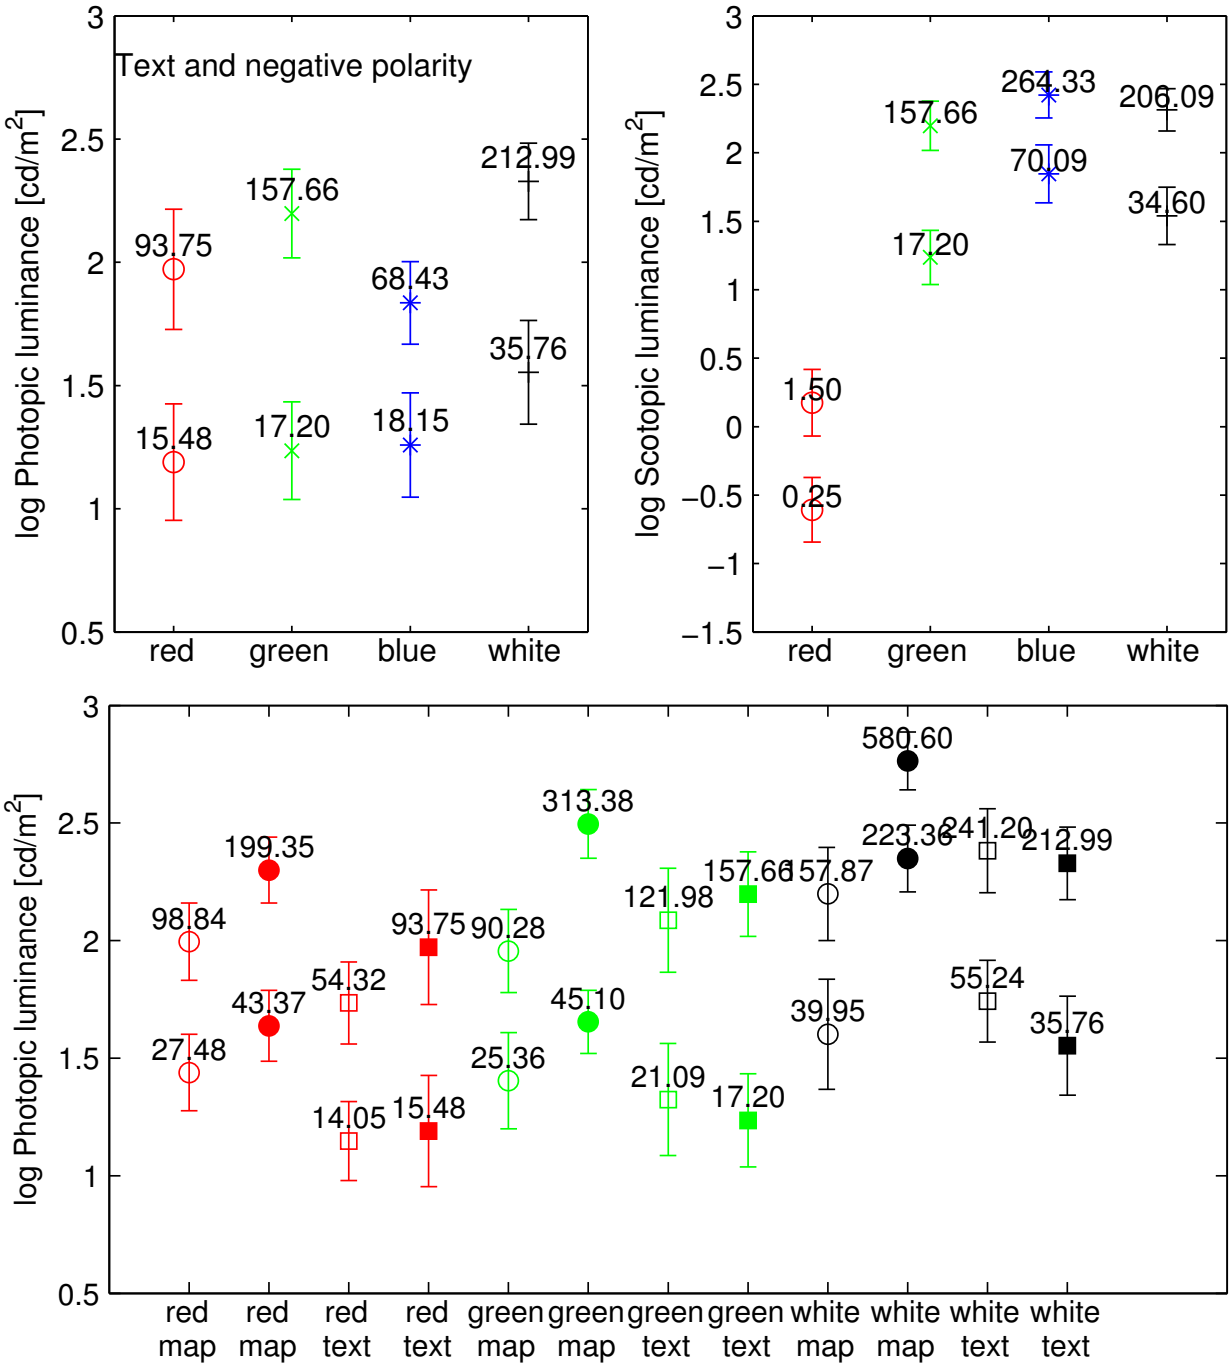

Figure 5.8: Results of Experiment 5.2 – preferred brightness range. For easier comparison the top 2 graphs show the photopic and scotopic luminance settings only for the negative polarity text conditions. The bottom graph shows the complete results in photopic luminance units. The empty markers represent positive and the filled markers negative polarity. The upper and lower markers denote the just too bright and just too dark settings respectively. The numbers represent either photopic or scotopic peak display luminance in cd/m^2. Error bars denote the standard error of the means.

for selecting the brightness of a display used in dark environments. The lower end of that range should be used for displays that should interfere the least with night vision, while the upper end should be used for entertainment displays.

5.6 Experiment 5.3: Reading Performance

Having determined visibility thresholds and viewer preferences in our previous two experiments, the next step is to measure performance differences due to colour and polarity through a more objective experimental procedure.

Stimuli. This experiment used one display, which is the same prototype display used for the previous experiment. A series of images was presented to the subjects, where each image was composed of four smaller images, shown in Figure 5.9, arranged in a 2×2 pattern. Three of the four images were identical, but the fourth contained an additional word "not" in the left column, which is highlighted in the figure. The fourth would appear in any of the four quadrants at random. Both positive polarity and negative polarity images were used.

For each image displayed, the intensity of the LED backlight was set to the mean preferred value for that image, which was determined in the previous experiment. In that way, comparisons between colours were made between the optimal values for each of those colours.

Subjects. Five subjects (4 male, 1 female, ranging in age from approximately 25 to approximately 45) participated in this experiment. Each had normal or corrected-to-normal vision and normal colour vision.

Experimental Procedure. Each subject was shown 30 images, half of which were positive polarity and half were negative. For each polarity each subject was shown images with red, green, and white backlight. For each condition, each subject was shown 5 images. The orderings of polarity and colour were randomized.

At the beginning of the task, subjects were familiarized with the images and the placement of

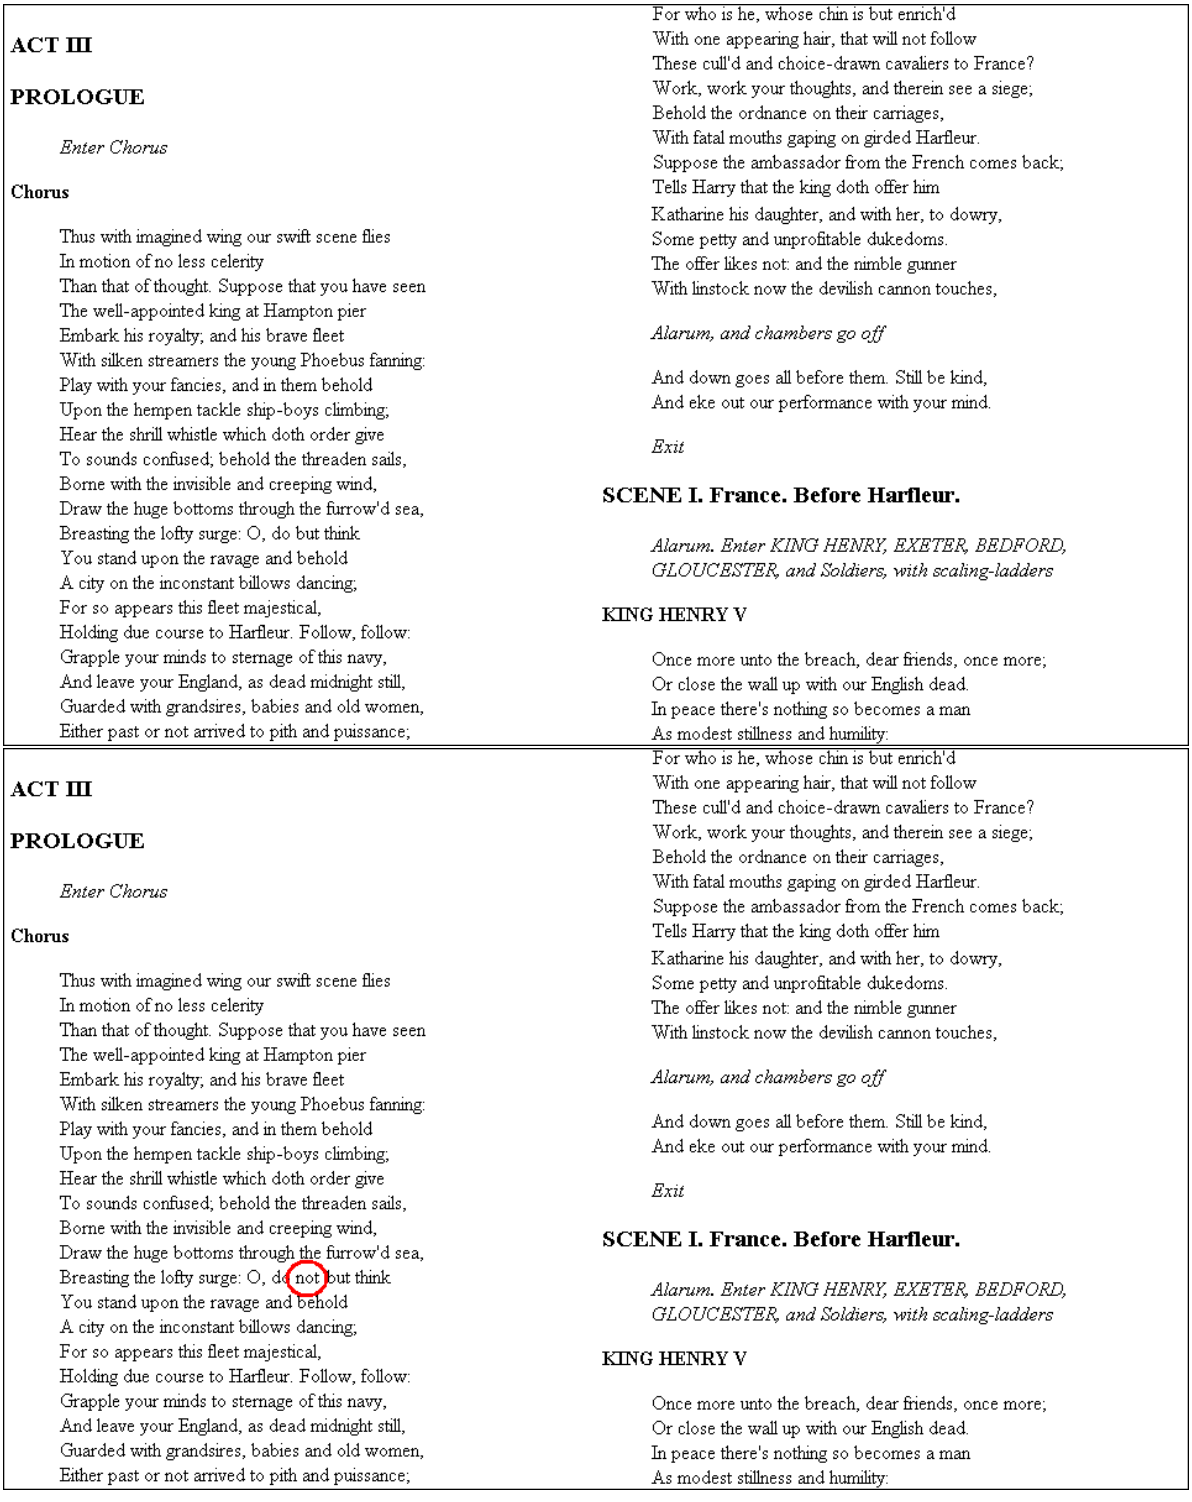

ACT III

PROLOGUE

Enter Chorus

Chorus

Thus with imagined wing our swift scene flies
In motion of no less celerity
Than that of thought. Suppose that you have seen
The well-appointed king at Hampton pier
Embark his royalty; and his brave fleet
With silken streamers the young Phoebus fanning:
Play with your fancies, and in them behold
Upon the hempen tackle ship-boys climbing;
Hear the shrill whistle which doth order give
To sounds confused; behold the threaden sails,
Borne with the invisible and creeping wind,
Draw the huge bottoms through the furrow'd sea,
Breasting the lofty surge: O, do but think
You stand upon the ravage and behold
A city on the inconstant billows dancing;
For so appears this fleet majestical,
Holding due course to Harfleur. Follow, follow:
Grapple your minds to sternage of this navy,
And leave your England, as dead midnight still,
Guarded with grandsires, babies and old women,
Either past or not arrived to pith and puissance;
For who is he, whose chin is but enrich'd
With one appearing hair, that will not follow
These cull'd and choice-drawn cavaliers to France?
Work, work your thoughts, and therein see a siege;
Behold the ordnance on their carriages,
With fatal mouths gaping on girded Harfleur.
Suppose the ambassador from the French comes back;
Tells Harry that the king doth offer him
Katharine his daughter, and with her, to dowry,
Some petty and unprofitable dukedoms.
The offer likes not: and the nimble gunner
With linstock now the devilish cannon touches,

Alarum, and chambers go off

And down goes all before them. Still be kind,
And eke out our performance with your mind.

Exit

SCENE I. France. Before Harfleur.

Alarum. Enter KING HENRY, EXETER, BEDFORD, GLOUCESTER, and Soldiers, with scaling-ladders

KING HENRY V

Once more unto the breach, dear friends, once more;
Or close the wall up with our English dead.
In peace there's nothing so becomes a man
As modest stillness and humility:

ACT III

PROLOGUE

Enter Chorus

Chorus

Thus with imagined wing our swift scene flies
In motion of no less celerity
Than that of thought. Suppose that you have seen
The well-appointed king at Hampton pier
Embark his royalty; and his brave fleet
With silken streamers the young Phoebus fanning:
Play with your fancies, and in them behold
Upon the hempen tackle ship-boys climbing;
Hear the shrill whistle which doth order give
To sounds confused; behold the threaden sails,
Borne with the invisible and creeping wind,
Draw the huge bottoms through the furrow'd sea,
Breasting the lofty surge: O, do not but think
You stand upon the ravage and behold
A city on the inconstant billows dancing;
For so appears this fleet majestical,
Holding due course to Harfleur. Follow, follow:
Grapple your minds to sternage of this navy,
And leave your England, as dead midnight still,
Guarded with grandsires, babies and old women,
Either past or not arrived to pith and puissance;
For who is he, whose chin is but enrich'd
With one appearing hair, that will not follow
These cull'd and choice-drawn cavaliers to France?
Work, work your thoughts, and therein see a siege;
Behold the ordnance on their carriages,
With fatal mouths gaping on girded Harfleur.
Suppose the ambassador from the French comes back;
Tells Harry that the king doth offer him
Katharine his daughter, and with her, to dowry,
Some petty and unprofitable dukedoms.
The offer likes not: and the nimble gunner
With linstock now the devilish cannon touches,

Alarum, and chambers go off

And down goes all before them. Still be kind,
And eke out our performance with your mind.

Exit

SCENE I. France. Before Harfleur.

Alarum. Enter KING HENRY, EXETER, BEDFORD, GLOUCESTER, and Soldiers, with scaling-ladders

KING HENRY V

Once more unto the breach, dear friends, once more;
Or close the wall up with our English dead.
In peace there's nothing so becomes a man
As modest stillness and humility:

Figure 5.9: Quadrants of the text image used in the reading performance experiment. The lower image with the word "not" was in one of the four quadrants (at random) while the upper image without the "not" was in the other three quadrants.

the word "not". Subjects were told that this was a timed task, and that they were to locate the quadrant and press the appropriate button as quickly as possible for each image. For ease of use, the buttons used were F1, F12, left-Ctrl, and right-Ctrl.

Results. The ANOVA test on the response time for colour $\times$ polarity indicated no statistically significant difference for displays of different colour ($F(2,329) = 0.9$, $p = 0.4$), but negative polarity was found to be statistically significantly better than positive polarity ($F(1,329) = 8.77$, $p < 0.01$). This result did not indicate any disadvantage of using coloured display as compared to a white display in a text reading task. The response time for negative polarity (white on black, 2.71 s average response) was on average 0.37 second shorter than for positive polarity (3.08 s). This suggests that using negative polarity colour scheme for dark environments may improve text legibility.

5.7 Discussion

In our experiments, we could test only the three primary colours of our prototype display. However, the rod-cone sensitivity curves suggest that a hypothetical display with red primaries as in our prototype display as well as amber primaries with the peak close to 610 nm can also display colours that do not disrupt scotopic vision. When paired with the red primary, the amber primary forms the colour gamut that maximizes photopic and minimizes scotopic luminance thus reducing disability glare. Figure 5.10 shows the CIE chromaticity diagram with the D65 white point and the three points that denote the primary colours of the sRGB colour space, which is the standard most commonly used in modern digital cameras and displays. The *gamut* of displayable colours is the polygonal region bounded by those three points. The fourth amber primary is often considered in the display design because it can extend the gamut of distinguishable colours and is more power-efficient than red in terms of luminance output. Our results indicate that such a 4-primary display would be beneficial under low ambient light. Such a display can adapt to ambient light by switching between 2-primary dark (red, amber) and 4-primary bright (red, green, blue, amber) modes, where the dark mode minimizes glare in the scotopic vision range.

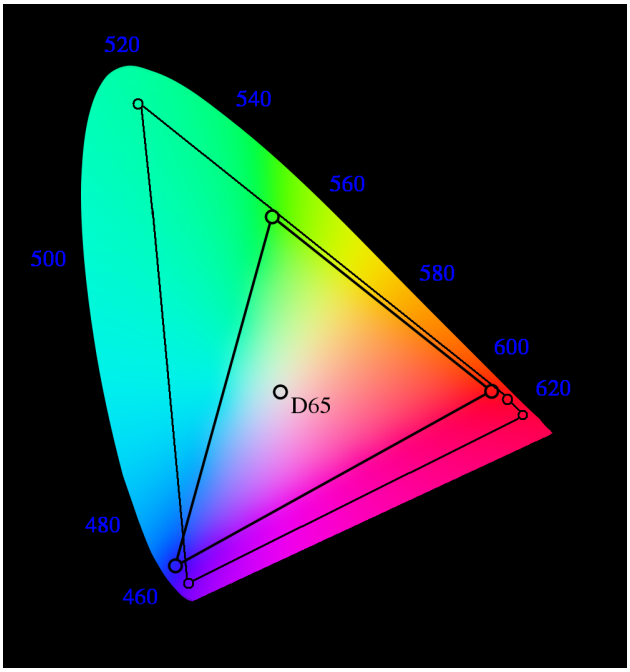

Figure 5.10: CIE chromaticity diagram and sRGB primaries (inner polygon) and RGB LED primaries with amber (outer polygon). The gamut of displayable colours is the polygonal area bounded by the points.

In this study we did not investigate how viewing a coloured display affects adaptation state, that is, whether a quick glance at the navigation display can temporarily impair scotopic vision. However, assuming that the adaptation in dark environments is determined by scotopic luminance when detecting scotopic targets, we can speculate also that in this regard a long-wavelength display should have the least impact on visual performance.

Although we see many benefits of a display design that preserves rod vision, it may not be an essential requirement for some applications. For example, the study of Flannagan[Flan 07] shows that the rods play little role in detecting pedestrians in a night driving scenario. However, that study does not preclude that the rod vision may influence other tasks, such as enhancing the saliency of emergency signals in the periphery.

Since the publication of the work described in this chapter, the data from these experiments has already found use in the development of new version of the HDR Visible Difference Predictor [Mant 11]. It is similar to the VDP that was used in Chapter 4; however it is built upon a

much more comprehensive visual model, and its predictive ability is much higher than that of other metrics.

Chapter 6

The Role of Contrast in the Perceived Depth of Monocular Imagery

Since ancient times, scientists have considered the cues that provide us with sensations of depth as the light from three-dimensional scenes is projected onto our curved two-dimensional retinae. These phenomena have been well known already for many decades [Bori 42, ch. 8].

Optical illusions provided some early opportunities to study the different types of cues and observe when they produce accurate 3D percepts and when they do not. For instance, the *Moon Illusion*, in which the size of the moon at the horizon appears larger than when higher up in the sky, was known to the ancient Greeks. von Helmholtz [Helm 25, v. III, p. 360-362] discusses the history of explanations of this phenomenon from Ptolemy (150 A.D.) through to his time, with the conclusion being that the intermediate objects along the horizon between the viewer and the moon lead the viewer to perceive the moon to be farther away at the horizon than it is higher up, and since it occupies the same angular size, it is perceived to be larger.

New technologies can also expose characteristics of depth perception. For instance, one can

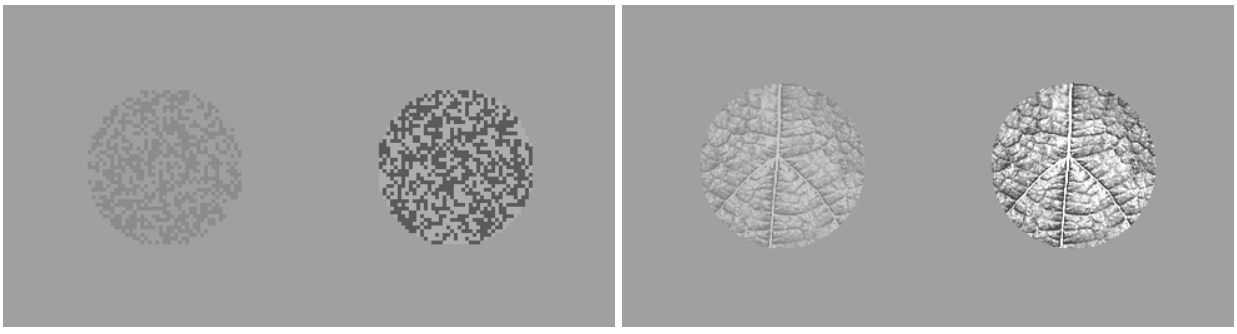

Figure 6.1: The perceived depth of an object is related to the contrast both between the object and the background, and within the object itself.

observe a difference in the strength of 3D percepts between television programs viewed on smaller and larger displays, or between standard-definition and high-definition displays. In fact, my personal observations of this difference in the context of sports telecasts as well as the animated 3D graphic title sequence of the program "Numb3rs" provided some of the motivation for this work. These suggest that depth perception may be related to display characteristics such as resolution, brightness, contrast, or size. These characteristics can reduce the gap between the retinal image produced by the display and that which would have been produced by viewing the actual scene.

Some depth cues such as *parallax* and *oculomotor* effects can normally be perceived only by viewing the actual 3D scene, or with a sufficiently accurate simulation using 3D stereo viewing equipment. Parallax refers to having two different views of a scene, and can come from motion (of the object or the observer) or from the *stereopsis* feature of human binocular vision, which is the ability to fuse two slightly different retinal images into a single 3D image. Oculomotor cues refer to the sensation of movement of the ocular muscles, and are much more subtle.

In addition, there are many other, generally monocular, depth cues that may be captured in a painting or photograph or other 2D image of a 3D scene, such as perspective, relative sizes of objects, familiarity with sizes of objects, occlusion, and *aerial perspective* which includes contrast, colour saturation, and haze. Contrast and brightness are of particular interest to us since they can be manipulated through a much greater range on high dynamic range displays than is possible on conventional displays.

HDR displays such as those demonstrated by Seetzen and colleagues [Seet 04] have a dramatically increased dynamic range over conventionally available *low dynamic range* (LDR) display technologies, and come much closer to meeting the abilities of the human visual system (HVS). When viewing demonstrations of these displays, many observers have remarked that they appear to be able to represent 3D to a much greater degree than conventional displays [Seet 07]. In this work, we examine some facets of 3D (depth) perception and conduct psychophysical experiments which examine the nature of this depth percept and whether the effect is more prominent on HDR displays.

6.1 Experiments

We conducted three experiments to analyze the relationship between contrast and perceived depth of objects in an image. Our first experiment was designed to reproduce the results obtained by Ichihara [Ichi 07] for random dot patterns (Section 6.1.1), and to extend them to natural textures (Section 6.1.2). Our second experiment (Section 6.1.3) was designed to analyze the impact of a higher dynamic range on depth perception. Finally, our third experiment (Section 6.1.5) analyzes the impact of tone curves on depth perception for natural scenes.

6.1.1 Experiment 6.1a: LDR Texture Contrast and Area Contrast

The experiments of Ichihara [Ichi 07] are well designed to show the effect of brightness and contrast on depth perception in the context of conventional LDR displays. These are the depth cues upon which HDR technology can have the most significant effect, so they are the focus of this study. Our Experiment 6.1a mimics closely their Experiment 2, with the exception that ours was conducted on an HDR display simulating LDR by using a uniform backlight, while theirs was conducted on a conventional LDR display.

Subjects. Ten subjects (8 male, 2 female, aged 22–42) participated in Experiments 6.1 (both parts a and b) and 6.2. Each of them had normal (20/20) or corrected-to-normal vision, which was confirmed through the administration of a Snellen visual acuity test. Each session lasted approximately one hour, after which subjects filled out a questionnaire about their experience and perceptions.

Stimuli. The stimuli consisted of random-dot disks, shown horizontally side-by-side, on a uniform background, and were presented on a Dolby DR-37P HDR display at a viewing distance of 100 cm, at eye level in the center of the subject's field of view. Some examples of random-dot disks are shown on the left side of Figure 6.1. The distance between the centers of the disks was 5.7° of visual angle, and each disk had a diameter of 3.2° (132 pixels). Each disk was composed of an equal number of light and dark dots, arranged randomly throughout the disk. Each dot was a

3×3 pixel square.

Each pair of disks included a reference disk and a test disk. The reference disk had a constant texture contrast of 0.5 and a constant average luminance of 30 cd/m^2. For the test disk, texture contrast was varied from 0.1 to 0.9 in steps of 0.1, and the average luminance was varied from 10 to 50 cd/m^{2}in steps of 10 cd/m^2. For this experiment, the Michelson contrast metric was used. Texture contrast was defined as the contrast between light and dark dots within a disk ($(L_{light} - L_{dark})/(L_{light} + L_{dark})$), while area contrast was defined as the contrast between the average luminance of a disk and the uniform background luminance ($(L_{disk} - L_{background})/(L_{disk} + L_{background})$). Background luminance levels of 20 and 40 cd/m^{2}were used.

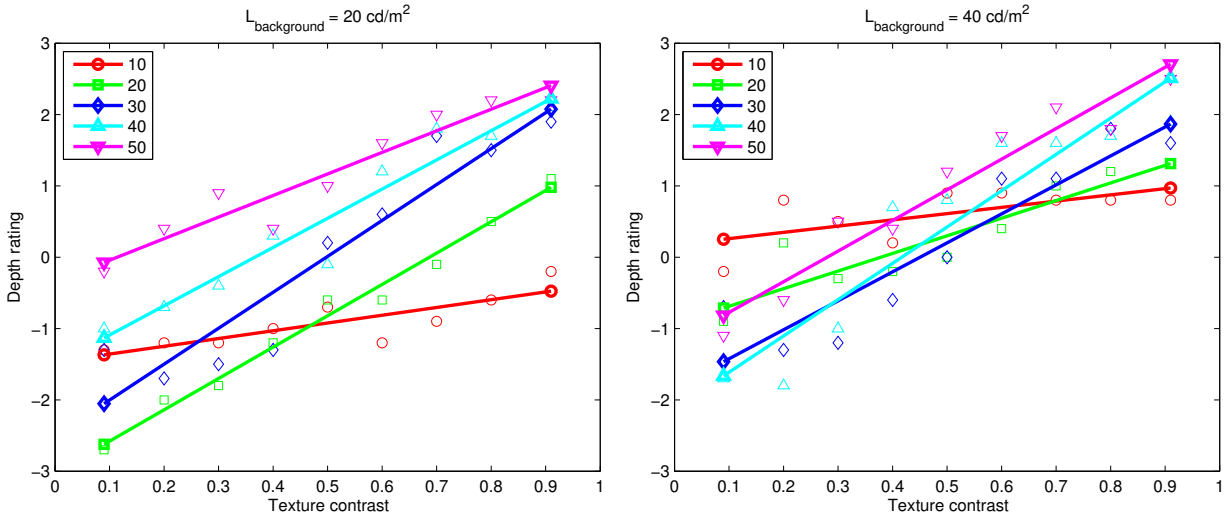

Figure 6.2: Perceived closeness increases with texture contrast in Experiment 6.1a.

Experimental Procedure. Subjects were shown a series of pairs of random-dot disks, and asked to indicate on the keyboard for each pair which disk appeared closer (or larger) and by how much. The F, D, and S keys were used to indicate that the left disk was closer by a low, medium, or high amount respectively, while the J, K, and L keys were used to indicate similarly that the right disk was closer. The space bar was used to indicate that the subject perceived the disks to be at about the same depth.

A practice run of 10 pairs of disks covering a representative range of contrasts was first con-

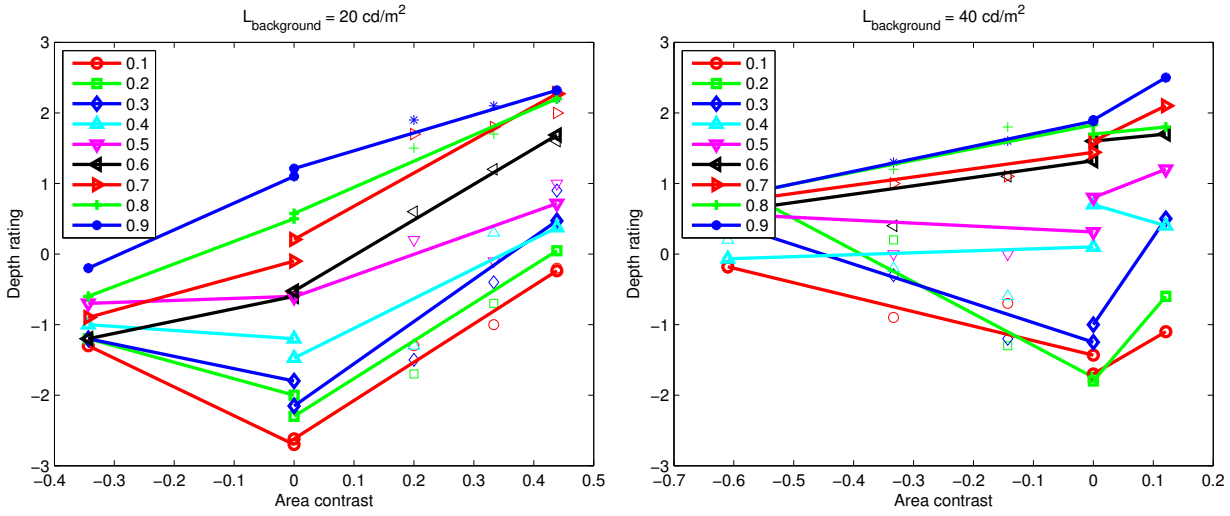

Figure 6.3: Perceived closeness increases with area contrast in Experiment 6.1a.

ducted to familiarize subjects with the mechanics of the experiment. This was then followed by 90 trials: 2 background luminance levels $\times$ 5 average disk luminance levels $\times$ 9 texture contrast levels. The full set of 90 trials was randomized, and the positions of the reference and test disks were exchanged at random. Subjects went through the experiment at their own pace.

Results. The results are substantially similar to those of Ichihara et al. Increases in texture contrast corresponded with increases in the perceived closeness of a disk, as shown in the left plot of Figure 6.4 which combines the results for background levels of 20 and 40 cd/m^{2}for clarity and brevity, as well as Figure 6.2 which corresponds with Figure 4 in [Ichi 07, p. 692]. Further, the effect of texture contrast was significantly more pronounced at low levels of area contrast than at high levels, as shown in the left plot of Figure 6.5. This corresponds with Figure 5 in [Ichi 07, p. 692].

We also used the experimental data to analyze two questions not addressed by Ichihara et al.: how does depth perception vary with area contrast at constant texture contrast levels, and how does varying the texture contrast affect those curves? Figure 6.3 as well as the right plot of Figure 6.4 show that perceived closeness generally increases as the area contrast of a disk increases, with texture contrast being held constant, when the area contrast is positive (bright area on a dark

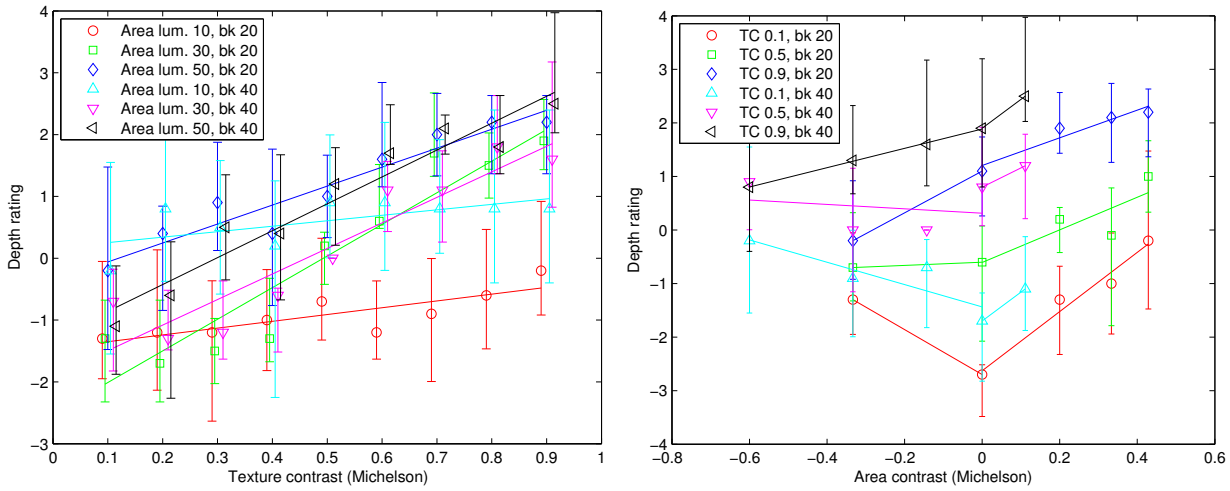

Figure 6.4: Data from random-dot disk images: Mean depth judgements as a function of texture contrast (left) at different levels of average disk luminance (10-50 cd/m^2), and area contrast (right) at different levels of texture contrast (0.1-0.9), at two levels of background luminance (20,40 cd/m^2). The vertical axis shows the depth rating score, where 3 indicates that the disk seemed closer by a large amount and -3 indicates that the disk seemed farther by a large amount. The solid lines show best fit using least-squares approximation, while the error bars show the standard deviations at each point. The points are offset slightly in x to improve clarity.

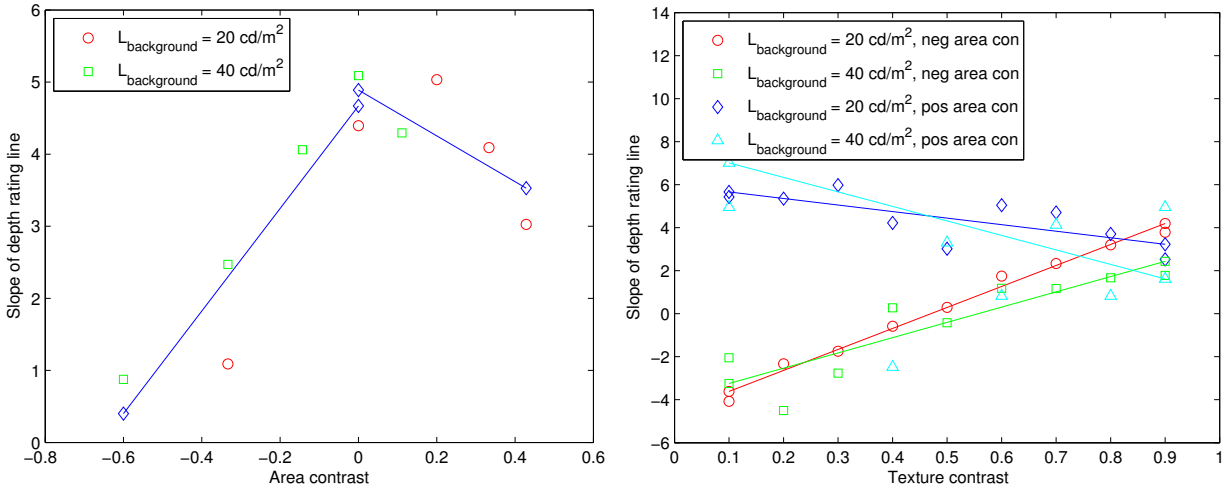

Figure 6.5: Data from random-dot disk images: The relationship between area contrast and the degree to which texture contrast affects depth perception (left), and the relationship between texture contrast and the degree to which area contrast affects depth perception (right). The blue lines show best fit using least-squares approximation.

background). When the area contrast is negative (dark area on a brighter background), perceived closeness increases as the area contrast increases at low levels of texture contrast, but decreases as area contrast increases at high levels of texture contrast. The effects of increases in texture contrast, mitigating the effect of positive area contrast and reversing the effect of negative area contrast, are shown in the right plot of Figure 6.5.

An ANOVA test on depth rating $\times$ texture contrast showed statistical significance in all cases except those with the lowest area luminance (10 cd/m^2) as indicated by the $F_{8,81}$ values in the upper ("dots") row of Table 6.1. The same is true for the "leaf" row which shows the results for Experiment 6.1b (Section 6.1.2). The ANOVA test on depth rating $\times$ area contrast was also statistically significant; the $F_{3,36}$ values are given in the bottom row of Table 6.1. Ichihara et al. found statistical significance in the difference between the slopes of the lines relating depth rating and texture contrast on their random-dot images, (the left plot of Figure 6.4) and so did we on both our random-dot and leaf images ($F_{4,45} = 21.10$ (dots) and $F_{4,45} = 10.48$ (leaf) at 20 cd/m^2, $F_{4,45} = 14.20$ (dots) and $F_{4,45} = 9.44$ (leaf) at 40 cd/m^2, $p < 0.01$).

Back lum.	20			40		
Area lum.	10	30	50	10	30	50
dots	n.s.	55.39	9.05	n.s.	31.55	20.25
leaf	n.s.	62.90	26.70	3.27	70.31	38.43
Tex con.	0.1	0.5	0.9	0.1	0.5	0.9
dots	10.93	5.23	5.50	$3.28^{\dagger}$	$3.06^{\dagger}$	n.s.
leaf	$3.70^{\dagger}$	28.95	n.s.	7.45	7.80	n.s.

Table 6.1: F values from ANOVA performed on depth rating vs. texture contrast at varying area luminance (above) and depth rating vs. area contrast at varying texture contrast (below), with $p < 0.01$, or $p < 0.05$ where denoted by (†). Some results, denoted by "n.s.", did not reach statistical significance.

6.1.2 Experiment 6.1b: Natural Textures

For the purposes of seeing how contrast affects natural scenes, it is useful to know whether these relationships extend to images of naturally occurring texture in addition to random-dot images.

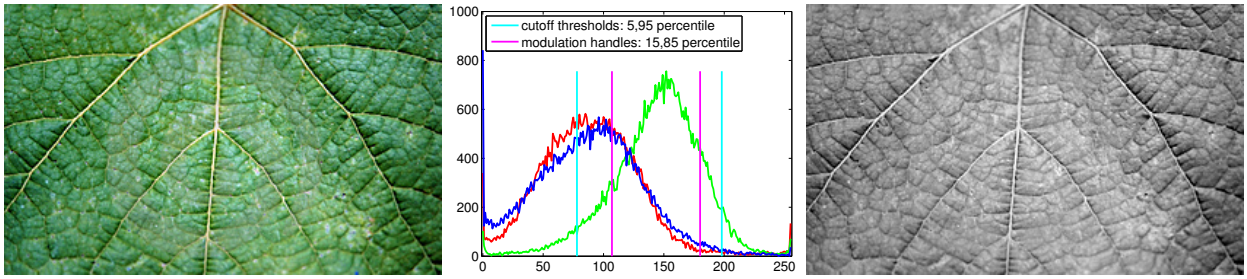

Figure 6.6: Left: the original leaf colour photograph image; Middle: the histograms for R, G, and B, with thresholds based on G; Right: the desaturated image where $R = G$ and $B = G$.

We tested this by running the same experiment as above, but using a leaf texture adapted from a close-up photograph instead of the random-dot images. For convenience, we integrated the two textures (dot and leaf) into the same experimental run, doubling the total number of trials from 90 to 180, where each trial could (at random) be either a leaf or dot texture. Therefore, the same subjects participated in both parts of the experiment.

Our experiments in this work relate to intensity and contrast, but not colour, so it was necessary to desaturate the colour from the leaf image. There are a variety of standard techniques for this, such as the luma computations of ITU-R Recommendations BT.601 ($Y = 0.299R + 0.587G + 0.114B$) or BT.709 ($Y = 0.2126R + 0.7152G + 0.0722B$), but since the leaf image was predominantly green, we used the simpler technique of setting the red and blue channels to be equal to the green channel (i.e. $Y = 0R + 1G + 0B$). The differences between these techniques would normally manifest themselves in brightness and/or contrast differences, but since our experiment controls both of those characteristics, there should be no difference in our experiment between those techniques. To confirm that, we tested all three techniques under a broad range of experimental conditions and observed no visible differences between the three techniques.

The histogram of the image shows a roughly Gaussian shape as one would expect, which is significantly different from the binary distribution of the random-dot pattern. We adjusted the distribution of the leaf image to make the texture contrast and overall area luminance consistent with that of the random-dot distribution through a two-step process. First, we clamped pixel values that exceeded low and high thresholds which we defined by the 5 and 95 percentile levels of the

image histogram (i.e. of the green channel). This prevented outlier pixels (such as the unnaturally large number of clamped pixels at 0 and 255) from unduly increasing the contrast of small high-frequency details to a level out of proportion with the random-dot patterns at the corresponding texture and area contrast levels. Second, we established handles at the 15 and 85 percentile levels of the image histogram, which we moved (along with the image pixel values) to match the low and high levels of the binary distribution in the dot images in the various texture and area contrast configurations.

Multiple visual comparisons (over a broad range of contrast settings) between the random-dot and leaf images at the same contrast levels showed them to have similar levels of overall brightness and contrast. In addition, the measured luminance levels of the leaf images at different settings were commensurate with the measured luminance levels of the random-dot images. We also made ramp images consisting of all the pixels in a leaf image, sorted by value. That enabled us to measure (with a Minolta LS-100 luminance meter) the luminance of the dark and light areas of that leaf image, which were also commensurate with the measured luminance of dark and light areas of the corresponding random-dot images.

Figure 6.6 shows the leaf texture photograph [Andr 09], the histograms of the red, green, and blue channels, and the monochromatic version we created to use in our tests.

Results. The relationships between area contrast, texture contrast, and depth were substantially the same as those observed in the random-dot experiment, and are shown in Figure 6.7. The results were statistically significant, as discussed in Section 6.1.1.

6.1.3 Experiment 6.2: HDR Texture Contrast and Area Contrast

Our next experiment was designed to analyze how the results of the previous experiment change when the range of contrasts is stretched out of the LDR domain and into the HDR domain.

The design of currently available HDR displays is based on the principle of dual modulation, which uses both high-resolution and low-resolution modulators [Seet 03]. A consequence of this design is that high dynamic range is only available with images that contain blocks of thousands

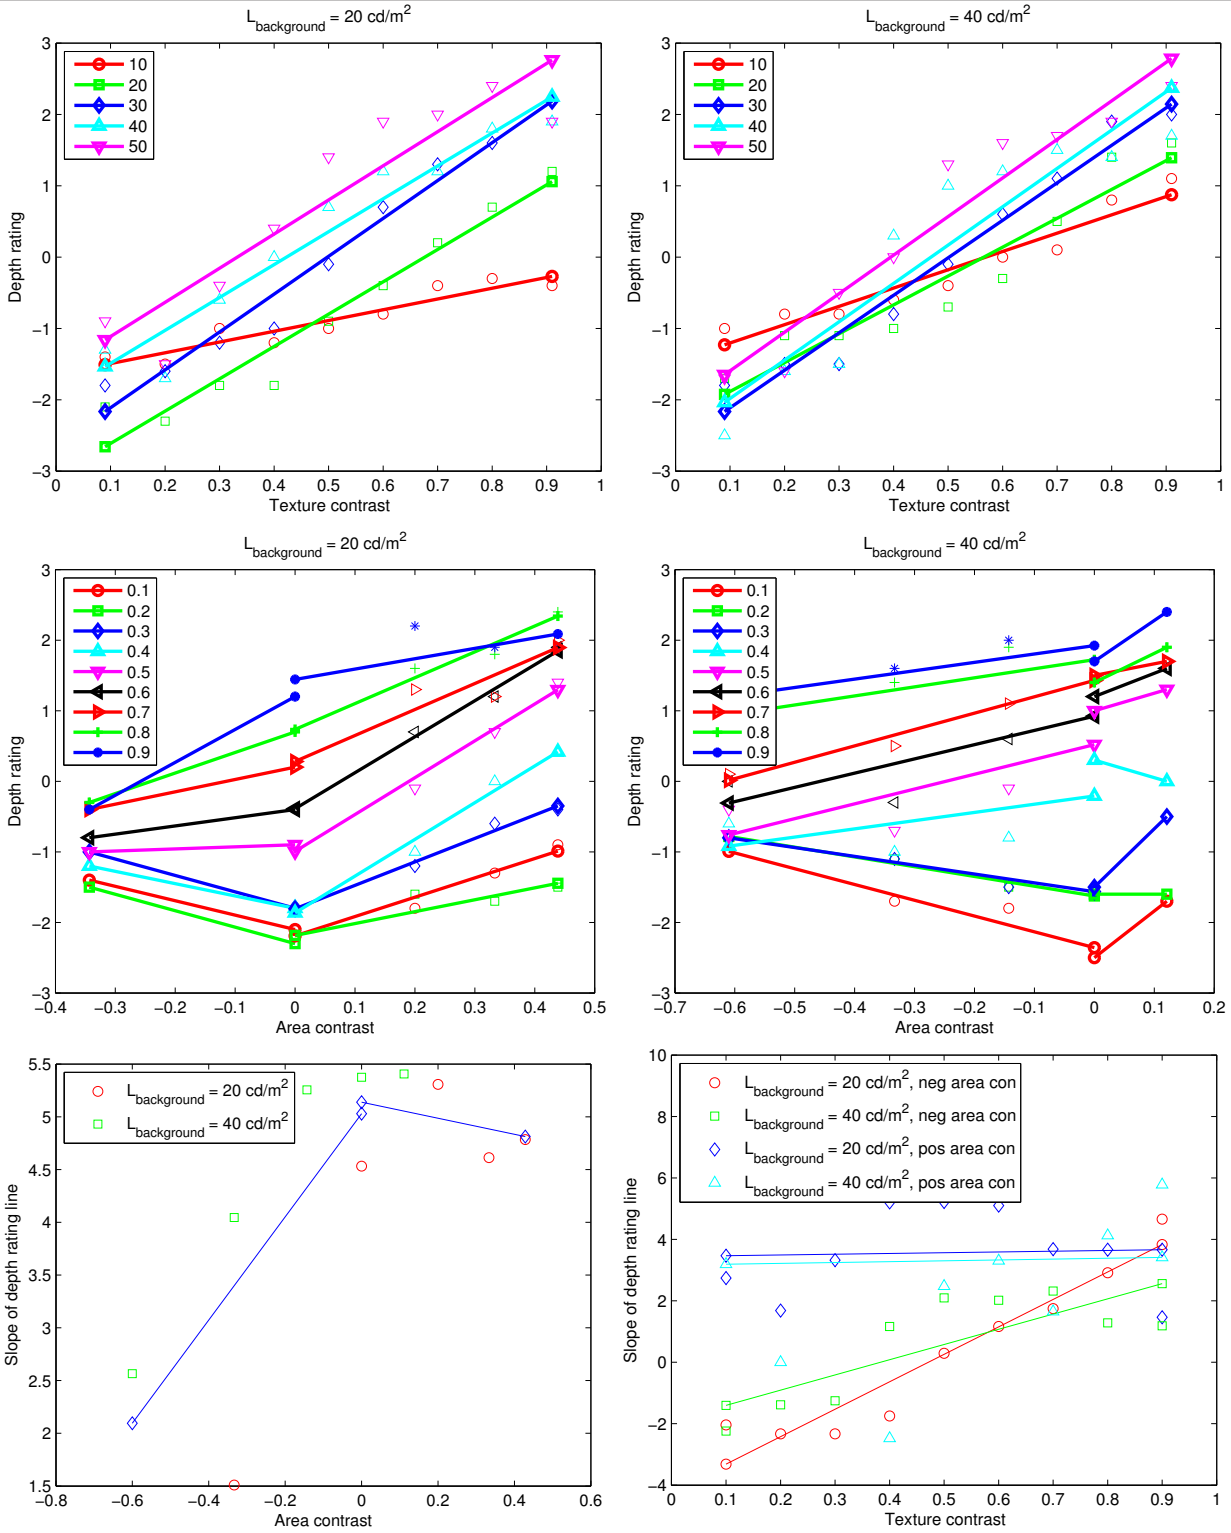

Figure 6.7: Experiment 6.1b: LDR leaf texture images (detailed results). The upper and middle plots show how perceived closeness increases with texture contrast and area contrast respectively. The lower plots show how the texture contrast effect varies with area contrast, and vice versa.

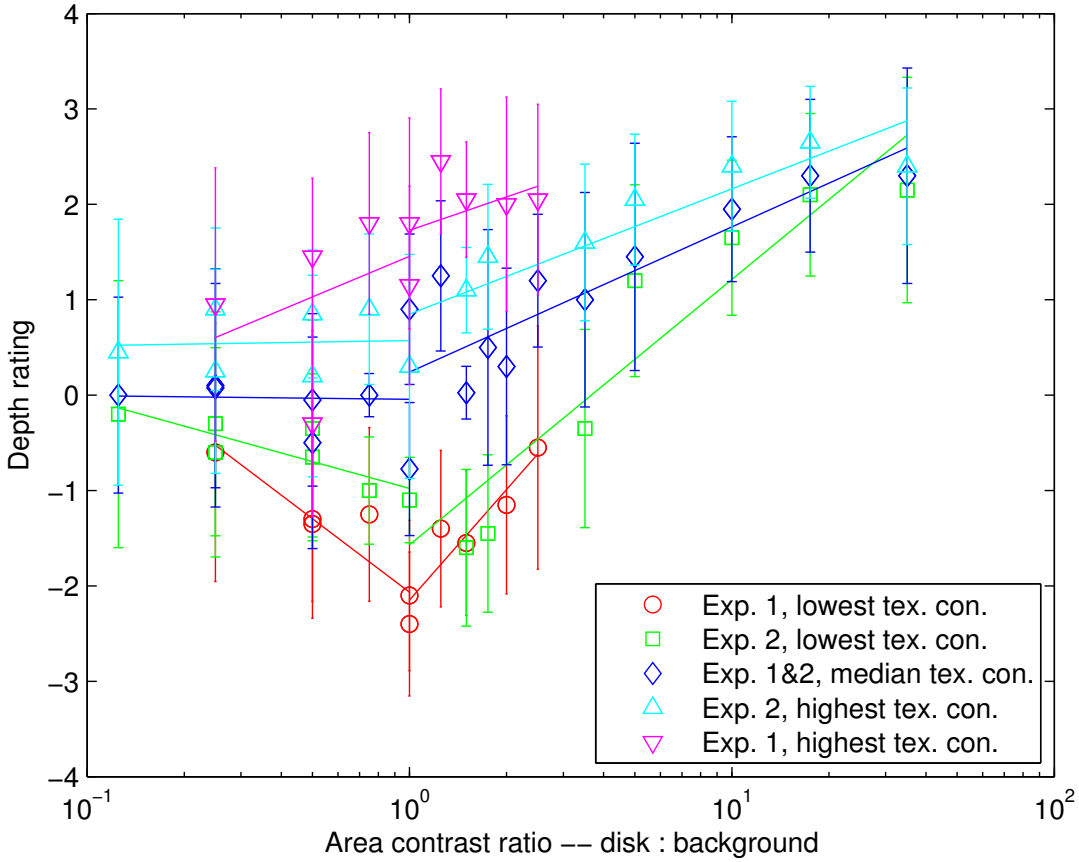

Figure 6.8: The effect of varying area contrast ratios on the perceived depth of disks, at low, medium, and high levels of texture contrast. The error bars show the standard deviations around the means, while the solid lines are best-fit lines for their corresponding data points.

or more contiguous pixels of similar brightness. In the context of the Ichihara experiments, it is therefore impossible to achieve texture contrasts beyond those available on LDR displays. However, high area contrasts can easily be achieved, and this is a significant component of the increased depth perception reported on HDR displays. However, as we will see in Section 6.1.5, it only accounts for part of the overall effect.

Subjects. The same subjects participating in Experiment 6.1 also participated in Experiment 6.2, which followed Experiment 6.1 in the same session.

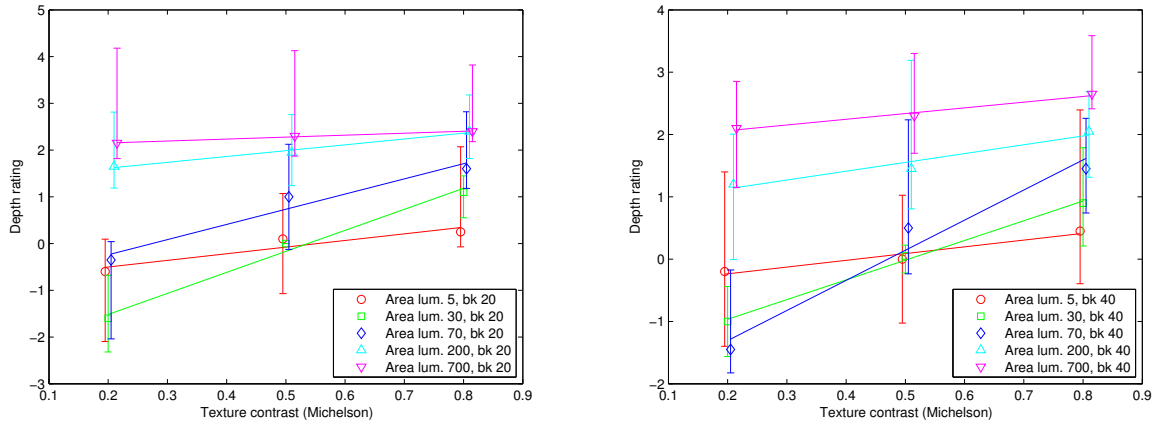

Figure 6.9: Data from HDR random-dot disk and leaf (combined) images: Median depth judgements as a function of texture contrast at varying levels of average disk luminance and background luminance. The solid lines show best fit using least-squares approximation, while the error bars show the standard deviations at each point. The points are offset slightly to improve clarity.

Stimuli. The stimuli used for this experiment were similar to those of the previous experiment. Again, random-dot disks and leaf texture images of the same size were used, but this time the distance between the disks was increased to 8.5° to allow the disks to be maximally coincident with an approximately circular block of LEDs. The full range of area contrasts would not have been achievable with the disk spacing used in Experiment 6.1. There is also precedent for the use of varying spacing between disks as Ichihara et al. used both 10.1° and 5.7° and achieved similar results in both conditions. The disk spacing was necessitated by the dual modulation scheme of the Dolby DR-37P HDR display that was used for this experiment, which employs an LCD panel and a backlight comprised of a hexagonal grid of LEDs. We placed the disks in order to fit as closely as possible to a block of 3 rows of LEDs with 2 LEDs in the top and bottom rows and 3 LEDs in the middle row. The grouping together of LEDs allowed for a greater range of area contrasts.

The range of luminances and contrasts was significantly greater than that used in Experiment 6.1. Background luminances remained at 20 and 40 cd/m^2, but disk area luminances were extended to a range from 5 to 700 cd/m^{2}using texture contrast levels of 0.2, 0.5, and 0.8. The reference disk remained at a constant texture contrast of 0.5 and a constant average luminance of

30 cd/m^2.

Experimental Procedure. The procedure was very similar to that of the previous experiment. There were a total of 54 trials: 3 background luminance levels $\times$ 3 average disk luminance levels $\times$ 3 texture contrast levels $\times$ 2 types of stimuli (random-dot disk and leaf image). Again, the full set of conditions was randomized, the positions of the reference and test disks were exchanged at random, and subjects completed the experiment at their own pace. The experiment was preceded by a practice run of 8 pairs of disks.

Results. The results of this experiment generally showed the same relationships between area contrast, texture contrast, and perceived depth as were observed in Experiment 6.1. Increases in texture contrast generally corresponded with increases in the perceived closeness of a disk, more so at low area contrast levels than at high (positive or negative) area contrast levels. Additionally, the increase in positive area contrast yielded an increase in the perceived closeness of a disk, particularly at lower texture contrast levels. This can be seen in Figure 6.9 by the consistently high depth ratings for the 200 cd/m^{2}and 700 cd/m^{2}lines, which are noticeably higher than those for lower area luminances. These strong depth ratings persist even when the background luminance is increased from 20 cd/m^{2}to 40 cd/m^2.

Figure 6.8 also shows the effect of area contrast on perceived depth, and the increased effect of the high levels of area contrast achievable on HDR displays over the lower levels of area contrast achievable on LDR displays. In this plot, we use the contrast ratio metric (L_{light}/L_{dark}) which is better able to represent higher dynamic ranges than is the case with Michelson contrast. The plot includes the results using both the leaf textures and the random-dot textures, since the results from the two types of textures were very similar to each other. The full results are shown in Figures 6.10 and 6.11. The ANOVA test showed statistical significance for most of these relationships, as shown in Table 6.2.

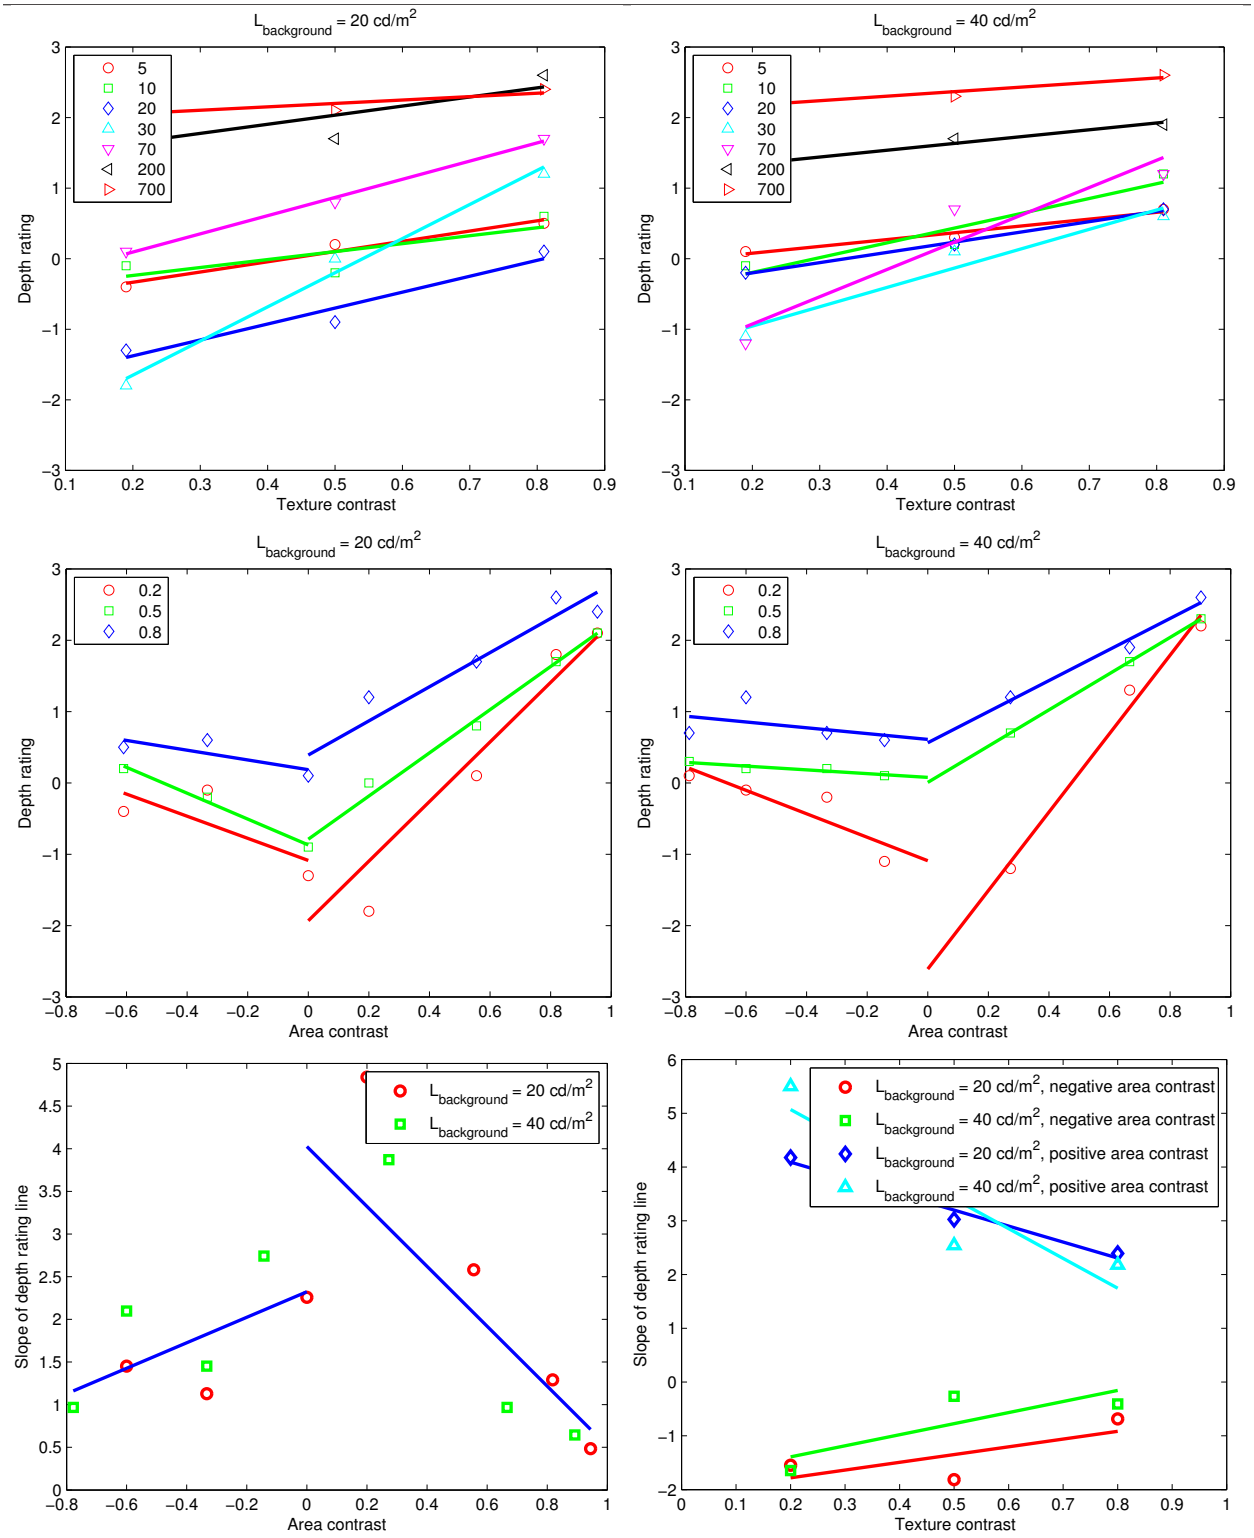

Figure 6.10: Experiment 6.2a: HDR random-dot disk images (detailed results). The upper and middle plots show how perceived closeness increases with texture contrast and area contrast (separated into positive and negative) respectively. The lower plots show how the texture contrast effect varies with area contrast, and vice versa. In all plots, the lines show best fit for those points.

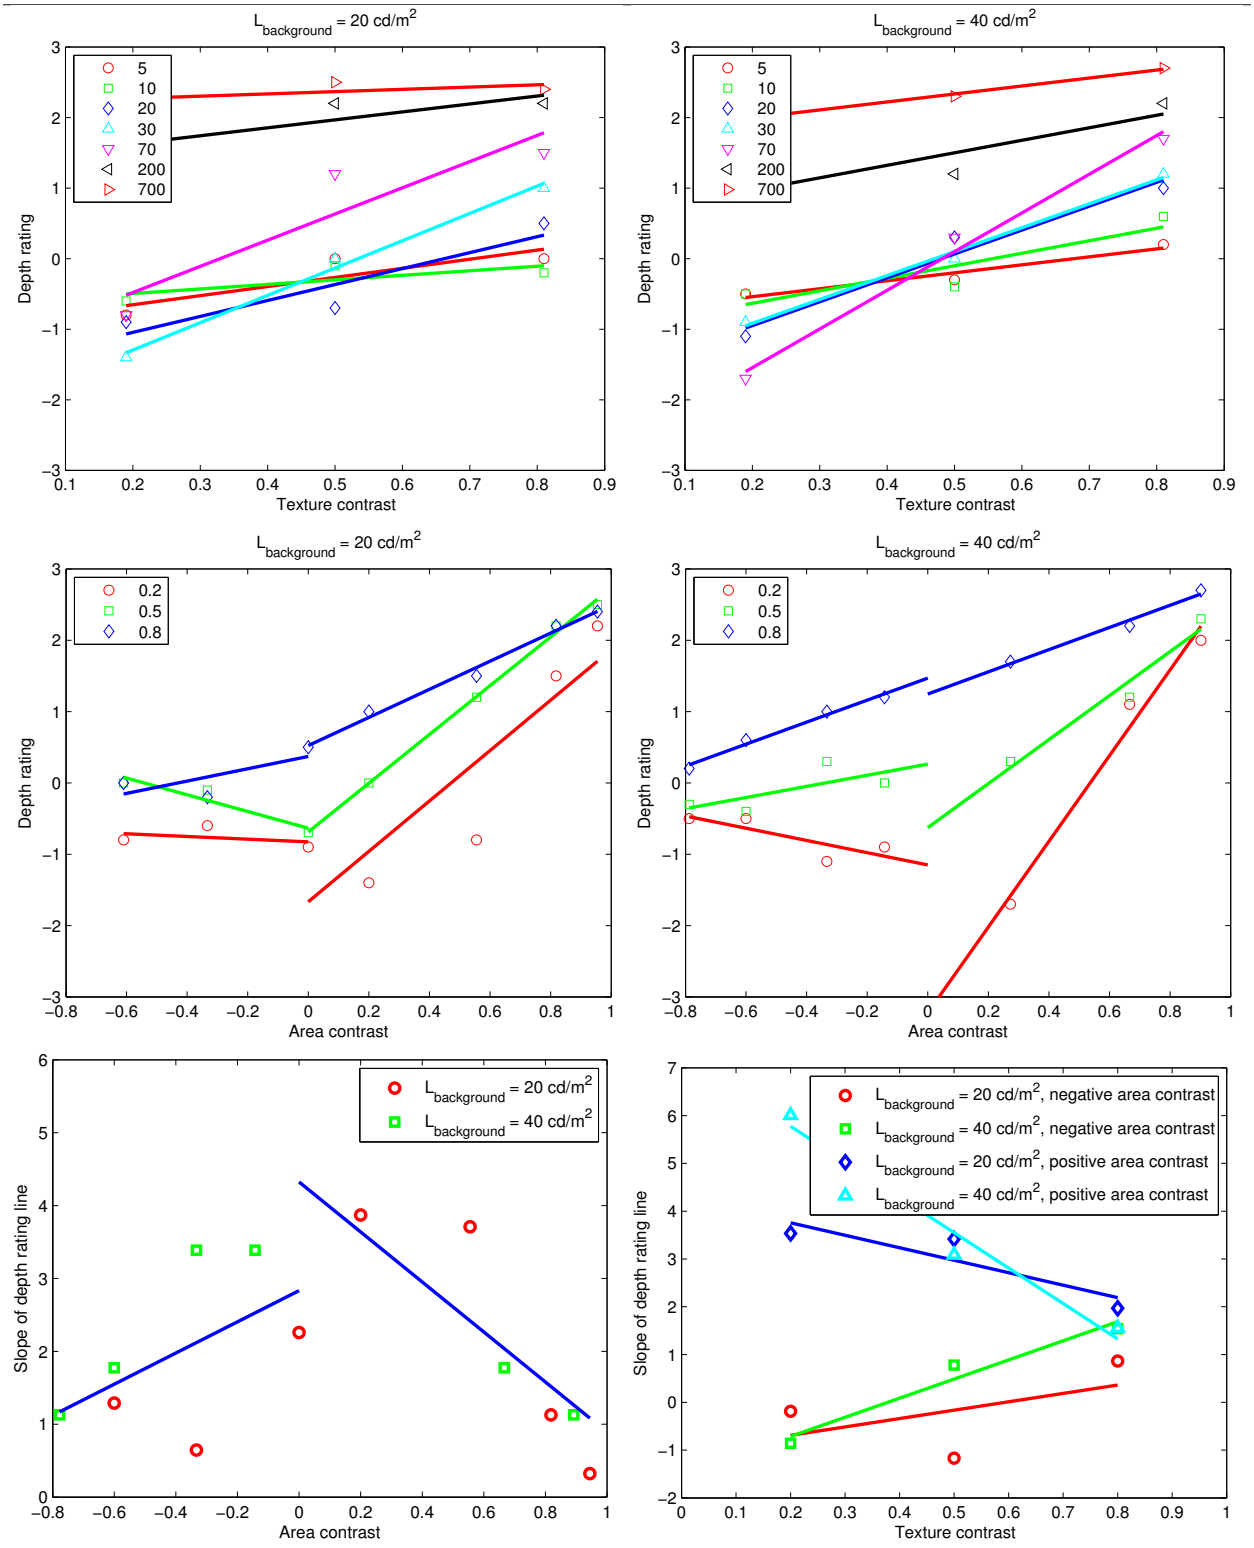

Figure 6.11: Experiment 6.2b: HDR leaf texture images (detailed results). The upper and middle plots show how perceived closeness increases with texture contrast and area contrast (separated into positive and negative) respectively. The lower plots show how the texture contrast effect varies with area contrast, and vice versa. In all plots, the lines show best fit for those points.

	$CR \leq 1$	$CR \geq 1$
Exp. 1, lowest tex. con.	$F_{3,116} = 14.90$	$F_{4,115} = 13.10$
Exp. 2, lowest tex. con.	$F_{4,135} = 3.13^{\dagger}$	$F_{7,152} = 65.20$
Exp. 1&2, median tex. con.	n.s.	$F_{10,269} = 25.32$
Exp. 2, highest tex. con.	n.s.	$F_{7,152} = 20.95$
Exp. 1, highest tex. con.	$F_{3,116} = 6.22$	$F_{4,115} = 3.79$

Table 6.2: F values from ANOVA performed on depth rating vs. contrast ratio between the disk area luminance and the background luminance, with $p < 0.01$, or $p < 0.05$ where denoted by (†). These correspond to the plots in Figure 6.8. Some results, denoted by "n.s.", did not reach statistical significance.

6.1.4 Questionnaire Results

In the questionnaire, subjects were asked whether they observed a relationship between contrast and the perceived depth of objects in the images, by indicating on a 7-point Likert scale the degree to which they thought disks of different characteristics were closer or farther. Figure 6.12 shows the characteristics and the subjects' responses. Note that subjects indicated strongly that they perceived brighter disks and higher-contrast disks to be closer than others. They also indicated very weakly that disks that were darker than the background seemed farther. This qualitative observation corresponds well to the quantitative observation in the right plot of Figure 6.5. We surmise that the weakness of this indication is likely due to a possible ambiguity about the meaning of "darker"; it could refer to dark disks with a high contrast against the background which would seem closer, or disks that are darker than the (closer-seeming) bright disks and hence seem farther.

Subjects were also asked (on a 7-point Likert scale) whether they found it easier to make depth judgements with the random-dot disks or with the natural-image leaf-patterned disks, and indicated that it was much easier with the leaf-patterned disks.

6.1.5 Experiment 6.3: Contrast Modulation in Natural Scenes

Experiment 6.2 demonstrated the importance of area contrast for depth perception. In Experiment 6.3, we analyzed the role of texture contrast for scenes with an already high area contrast. In

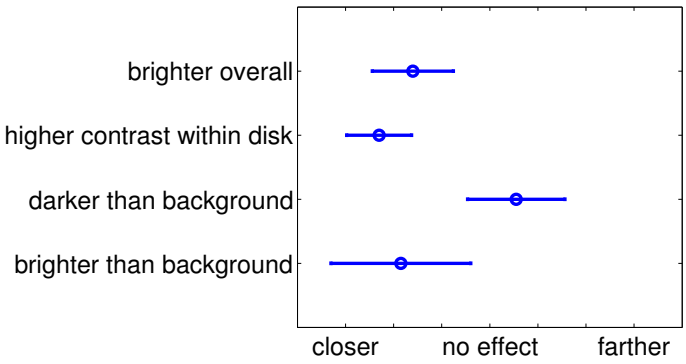

Figure 6.12: The degree to which subjects perceived disks of different characteristics to be closer or farther.

particular, we observe how modulating the contrast of a texture patch through the application of a tone curve can affect the perception of depth within that scene.

Subjects. Eleven subjects (all male, aged 22–34) participated in the experiment. Two of them had been subjects in the previous two experiments, while the remaining nine had no prior experience with this research. Each of them had normal (20/20) or corrected-to-normal vision, which was confirmed through the administration of a Snellen visual acuity test.

Stimuli. The stimuli in this experiment were textured regions that had been cropped from two HDR images. The purpose of working with cropped regions rather than full images was to eliminate depth cues other than contrast, including large-scale shadows, occlusions, scale recognition from familiar large-scale features, and so forth. One consequence of this is that the horizon line (and consequently the slant of surfaces in the image) would be different between the original image and the cropped region. But since all depth comparisons were made between images whose corresponding surfaces have the same slant, we can still isolate the relationship between perceived depth and contrast. Figure 6.13 shows the larger images (above) and the cropped regions (below). The images were displayed on the same HDR display that was used in the previous experiments.

Each image was rendered at multiple different contrast levels in which the intensity of both the brighter and darker parts of the image were independently modulated. Within the texture

Figure 6.13: Top: The full-sized HDR images. Bottom: The regions of the above images that were used in Experiment 6.3.

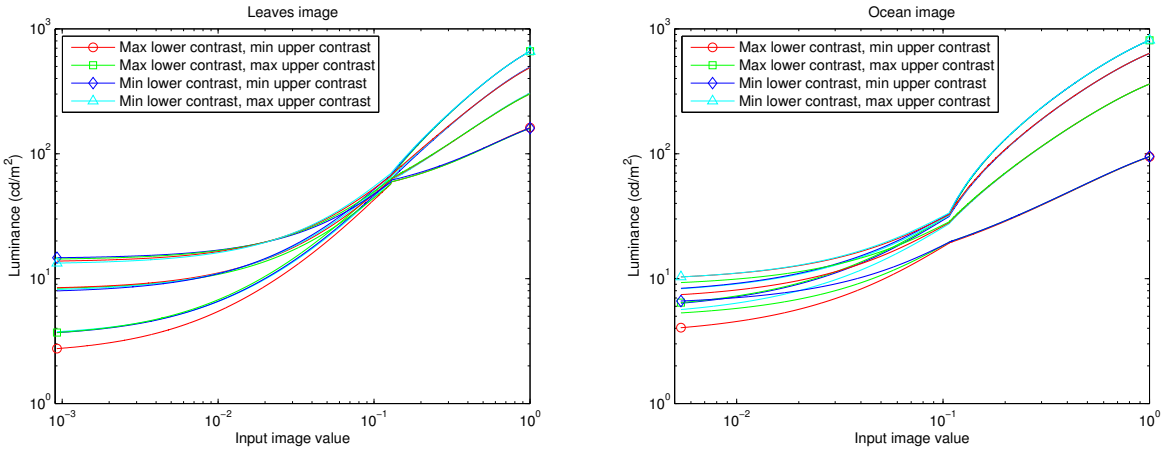

Figure 6.14: Logarithmic plots of the tone curves for the two images used in Experiment 6.3. Each plot shows 12 curves, corresponding to the 12 combinations of 4 upper contrast levels and 3 lower contrast levels. At high image input values the 4 upper contrast levels are consistent and easily distinguished. At low image input values, the 3 lower contrast levels are mostly consistent for the Leaves image, but are less consistent for the Ocean image as the upper contrast levels affected them somewhat. The colours of the curves are arbitrarily selected and serve only to differentiate them. To avoid clutter, the legend shows only the 4 extreme combinations of upper and lower contrast, omitting the intermediate levels. The markers show the differences between the minimum and maximum upper contrast levels and between the minimum and maximum lower contrast levels.

patches, the mid-tones, defined as the intensities in the range of the median input level in the image, were kept relatively constant throughout these contrast manipulations. The patches where shown against a dark background, with the goal of keeping area contrast uniformly high across the experiment. Sigmoidal functions similar to the ones used by Mantiuk and Seidel [Mant 08b] were used to reduce the contrast in the dark and bright regions without causing discontinuities in the tone curve. Sigmoid or S-shaped curves have been commonly used in traditional photography to adjust the intensity distributions at both the high and low ends while avoiding intensity discontinuities [Rein 02]. The upper intensities were scaled to four different levels while the lower intensities were scaled to three different levels. Figure 6.14 shows logarithmic plots of the tone curves that were used for the two images (Leaves and Ocean) in this experiment.

Experimental Procedure. Subjects were shown pairs of images in which the two differed only in the scaling levels of the upper and lower intensities. For each pair, the subject was asked to indicate on the keyboard which of the two showed a greater sense of depth within the scene. The screen was then cleared to the background level for 0.3 s before the next pair was displayed.

A practice run of 8 pairs of images selected at random from the full set was first conducted to familiarize subjects with the mechanics of the experiment. This was then followed by 156 trials: all combinations of 12 contrast variations (4 upper contrast levels $\times$ 3 lower contrast levels) of an image taken 2 at a time ($C_2^{12} = 66$), plus 12 pairs where the same image was present on both sides, times 2 images. Subjects went through the experiment at their own pace. There was a two-minute break at the half-way point of the experiment. The whole experiment took less than 25 minutes.

Results. Figure 6.15 shows the subjects' perception of the comparative depth of scenes with different levels of contrast between the highlights and the mid-tones. As in Experiments 6.1 and 6.2, we use the contrast ratio (L_{light}/L_{dark}) for our contrast calculations. Each point in this plot represents 1 of the 12 contrast variations for an image, where the horizontal axis represents the ratio between the bright point and the mid-tone for that variation. The vertical axis represents the mean depth rating score, relative to the other versions of the image; the error bars indicate the standard deviations. The vertical range is normalized to (0,1), where 0 indicates that all other variations

across all subjects were seen to have more depth than the current variation, while 1 indicates that all other variations were seen to have less depth, and 0.5 indicates that an equal split between the number of images that were perceived as containing more vs. less depth. As indicated in the plot, subjects indicated a greater sense of depth as the contrast of the highlights was increased, for both images. The ANOVA was statistically significant for both the Leaves ($F_{11,714} = 14.57$) and Ocean ($F_{11,714} = 24.76$) images ($p < 0.01$). However, we did not find a similar correlation between the sense of depth and the contrast of the darker parts of the images.

The smaller range of highlight contrast in the Leaves image is due to the higher mid-tone level of that image, which limits the available high-end contrast. The mid-tone levels of the Leaves image ranged between 56-70 $\mathrm{cd/m^2}$while the mid-tone levels of the Ocean image ranged between 19-34 $\mathrm{cd/m^2}$. In all cases, the background level was set to 0, resulting in luminance levels between 0.1-0.2 $\mathrm{cd/m^2}$.

The significance of the highlights in conveying depth is consistent with the observations of [Berb 83], as well as [Meyl 06] who found that bright specular highlights lead to a "more natural impression." Meylan et al. also pointed out that the strength of this effect can vary between images, which is also consistent with our results. These observations relate back to Gabriel Lippmann's challenge that a photographic print might someday appear as a window into the world [Lipp 08]. If the evolution of photography and the rendering of natural images is toward making it appear as if one were looking out a window, the sensations of both depth and a natural impression should be heightened as we make forward progress. Since the contrast of highlights conveys both, the use of those along with higher-contrast displays brings us closer to our goal.

6.2 Discussion

The results of our experiments show that the effect of depth within a scene can be heightened by increasing the contrast, both of small-scale features as well as large-scale features, within the scene. Experiments 6.1 and 6.2 demonstrated the effect of area contrast on depth and showed how HDR displays can enhance the depth effect even when only boosting area contrast. Experiment

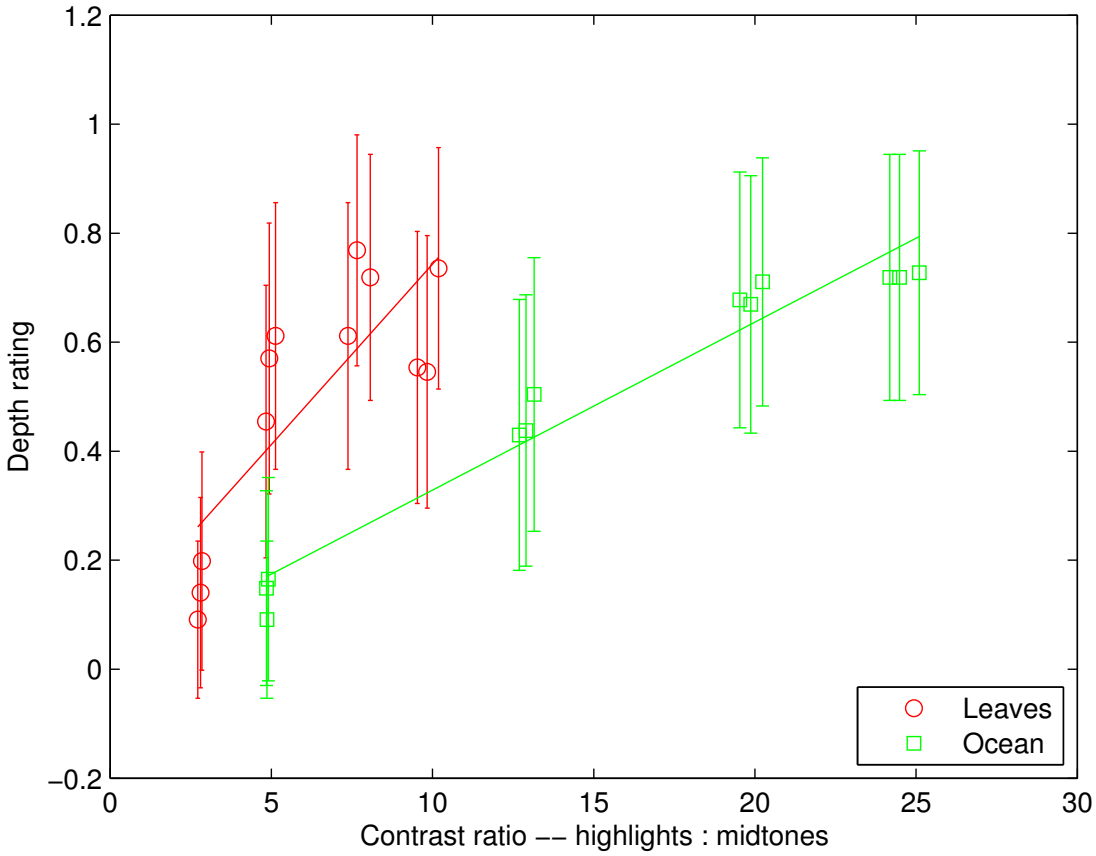

Figure 6.15: Subjects' perception of the comparative depth of scenes with different levels of highlight contrast. The depth rating indicates the probability (based on our experimental results) that a viewer would identify that image as having greater depth than other images within the range of our study.

6.3 then showed that even when area contrast is set to already-high levels such as 200:1 to 500:1 against a black background, it is still possible to obtain an even stronger sense of depth by boosting the contrast of the highlights.

Multiplying the highlight contrast with the mid-tone area contrast, we obtain total contrast ratios of up to 5000:1 in our Experiment 6.3. Even at those contrast levels, we were able to obtain noticeable differences in subjects' perception of depth. By taking advantage of the full range of available contrast in current and future generations of high-contrast display devices, we should be able to improve viewers' sensations of depth still further, particularly when the depth information from contrast is in line with other depth cues that are present in images. On the other hand, if the depth from contrast conflicts with the other depth cues, high-contrast display devices may diminish rather than enhance viewing experiences, just as inconsistent brightness was observed by Miles to diminish binocular viewing experiences [Mile 53].

Chapter 7

Conclusion

As new generations of displays are being developed which have dramatically greater capabilities than existing displays, we need to know how best to take advantage of the characteristics of those displays to create satisfying viewing experiences, and to observe how the human visual system responds to imagery displayed on such devices. The work presented in this dissertation contributes significantly to our understanding of both of these areas, and will be valuable in guiding the design of future generations of displays.

7.1 Summary of Contributions

In Chapter 3 we described a study in which we analyzed viewing preferences and the potential for visual fatigue on HDR displays. Our experiments show that visual fatigue is not a serious concern even in dark environments. We also found that subjects preferred higher levels of display brightness in higher ambient illumination environments, and that this relationship is strongly similar to that observed in previous work with cockpit display illumination. The results of this study could be the first step in designing HDR television sets with automatic brightness controls to provide a more pleasurable viewing environment under a variety of ambient viewing conditions.

In Chapter 4 we presented a method for on-the-fly expansion of the dynamic range of legacy, low dynamic range, video content for viewing on HDR displays. The method is robust and temporally coherent, and does not require image-specific parameter adjustment. As such the method is well-suited for integration directly in HDR display hardware, where it can be used to process video streams from legacy sources such as television or DVDs. This contrast enhancement method is also motivated by the observation in Chapter 3 that viewers prefer the highest contrast available

when viewing entertainment content.

In Chapter 5 we showed that a display that emits long-wavelength light (amber to red) offers several advantages over green and blue displays that are used under low ambient light. A red-coloured display least affects visual performance due to disability glare as it is the least likely to cause dazzling or eye aversion, even if emitting light at high photopic luminance levels. Other studies indicate that red coloured letters are the most legible at low luminance levels. We also showed that a display in dark environments must emit at least low photopic luminance levels ($\approx$20 cd/m^2) to be comfortable to read, and higher luminance ($\approx$40 cd/m^2) if displayed content, such as maps, requires distinguishing between several brightness levels.

In Chapter 6 we conducted three experiments in which we observed that increases in contrast resulted in increases in perceived depth in a displayed scene. We confirmed the findings of previous work which explored this relationship with conventional (LDR) displays, and expanded upon their analysis in three ways. We extended the luminance and contrast range to that available on HDR displays, we used image textures of natural scenes in addition to abstract textures, and we showed that perceived depth increases with area contrast just as it does with texture contrast. We then showed in a third experiment that the contrast between the highlights and the mid-tones in natural scenes is a strong determinant of the perceived depth within those scenes. Together, our experiments show the depth effect that can come from high-contrast imagery, and that the contrast of both large-scale features and small-scale features contribute significantly to that overall effect.

Collectively, the work presented in this dissertation significantly advances our ability to display imagery on advanced high-contrast and high-luminance display devices to create satisfying viewing experiences. Our results can be used in the design of future generations of television sets, as well as other types of displays such as cockpit displays, to enhance both enjoyment and functionality.

7.2 Limitations and Future Work

The work presented in this dissertation opens up several avenues for possible future research and development. In this section, we discuss some of those.

Video Viewing Preferences. In our work on video viewing preferences we focused on live action content without drastic illumination effects. This choice was based on the need to limit the number of parameters in the study. Now that visual fatigue has been ruled out and we have gained a better understanding of brightness and contrast preferences under different ambient light levels, future work could explore the effect of different types of content, such as flat-shaded animation and very dark or highly stylized live action footage.

Ldr2Hdr. Our work on dynamic range expansion of LDR imagery for HDR devices was focused on solutions that could have real-time implementations. More sophisticated, albeit slower methods would be interesting to explore in the future. Our work also focused only on the luminance domain; it would be interesting to explore more complex algorithms that consider the colour channels as well, either independently or in concert. As new displays are being developed that use coloured LEDs, the consideration of colour in dynamic range expansion will become of even greater value.

In collaboration with a colleague, Cheryl Lau, who has expertise in colour science, I have begun some exploratory work on adapting the Ldr2Hdr algorithm to make better use of the colour of the saturated input pixels. Some input images such as sunsets have image regions that saturate progressively in one colour channel after another. When the colour channels for an individual pixel are unevenly clipped, the colour of that pixel may be different than the real-world colour that pixel was supposed to represent, and extending the dynamic range through an achromatic luminance-based technique would retain that colour difference. To address this, we are exploring techniques that selectively extend the dynamic range depending on pixel colour, such as the following.

Consider the RGB cube that ranges in value from (0,0,0) to (1,1,1) and encodes the R, G, and B values of an input LDR image on its 3 axes. When one channel is saturated, the pixel will lie on a face of the cube incident to the (1,1,1) point; when two channels are saturated the pixel will lie

on an edge incident to (1,1,1); when all three channels are saturated, the pixel will lie at the point (1,1,1). In all three cases, the actual pixel values will be clipped versions of what their true values would have been, which we can attempt to recapture. To unclip those values, we can extend them along a direction perpendicular to the part of the cube on which those values lay, using the values in the saturation extension image to determine the degree to which to extend them. Further work is required to determine the value of techniques such as this one, but it has been an intriguing avenue of research so far.

Depth Perception. The current surge of interest in depth perception in visual imagery, at the same time as 3D representations of theatrical feature films are also finding renewed popularity, indicates motivation for further research in this area. In the future, we plan to conduct further studies to examine the relative strengths of different cues, and better quantify the degree to which HDR imagery can have an impact on depth perception. Also, since our work thus far has focused only on the luminance domain, we plan to extend the analysis to colour and examine the effect of colour contrast on depth perception.

The depth perception experiments conducted for this dissertation were all subjective, in that individuals consciously indicated their sensations of depth in the imagery. In future depth experiments, it would be valuable to use more objective performance-based methodologies that are more robust, less noisy, and less susceptible to individuals' whims. This will become more feasible with the use of specialized viewing equipment as discussed in the next paragraph.

It would also be interesting to compare depth from contrast with depth from binocular disparity. One approach to this would be to use a specialized 3D projection system to compare the viewing of 2D vs. 3D stimuli with different levels of contrast. Another approach is to explore how the ability to perceive the hidden shapes in autostereograms (Magic Eye pictures) is affected by display factors including contrast.

Bibliography

[Abil 93] A. Abileah. "Night vision goggle compatible liquid crystal display device". U.S Patent 5,262,880, Nov. 1993.

[Acos 00] P. Acosta-Serafini, I. Masaki, and C. Sodini. "Single-chip imager system with programmable dynamic range". U.S Patent 6,977,685, Feb. 2000.

[Adam 83] A. Adams. *The Print*. Little, Brown, and Co., 1983.

[Akyu 07] A. Akyüz, E. Reinhard, R. Fleming, B. Riecke, and H. Buelthoff. "Do HDR displays support LDR content? A psychophysical evaluation". *ACM Trans. Graph. (Proc. ACM Siggraph)*, Vol. 26, No. 3, pp. Article 38, 7 pages, 2007.

[Andr 09] Andreas G. "Suede leaf". http://www.dreamstime.com/suede-leaf-imagefree191102, accessed Dec. 17, 2009, 2009.

[Ashi 06] M. Ashikhmin and J. Goyal. "A Reality Check for Tone-Mapping Operators". *ACM Trans. Applied Perception*, Vol. 3, No. 4, pp. 399–411, 2006.

[Aydi 08] T. O. Aydin, R. Mantiuk, K. Myszkowski, and H.-P. Seidel. "Dynamic range independent image quality assessment". *ACM Trans. Graph.*, Vol. 27, No. 3, p. Article 69, 2008.

[Bant 06] F. Banterle, P. Ledda, K. Debattista, and A. Chalmers. "Inverse Tone Mapping". In: *Proc. of GRAPHITE '06*, pp. 349–356, 2006.

[Bant 08] F. Banterle, P. Ledda, K. Debattista, and A. Chalmers. "Expanding low dynamic range videos for high dynamic range applications". In: *Proc. of SCCG '08*, pp. 33–41, 2008.

[Bart 67] C. J. Bartleson and E. J. Breneman. "Brightness Reproduction in the Photographic Process". *Photographic Science and Engineering*, Vol. 11, No. 4, pp. 254–262, 1967.

[Bera 57] L. L. Beranek. "Criteria for Noise in Buildings". *NOISE Control*, Vol. 3, No. 1, pp. 19–27, 1957.

[Berb 83] K. Berbaum, D. Tharp, and K. Mroczek. "Depth perception of surfaces in pictures: looking for conventions of depiction in Pandora's box". *Perception*, Vol. 12, pp. 5–20, 1983.

[Berg 91] T. van den Berg, J. IJspeert, and P. de Waard. "Dependence of intraocular straylight on pigmentation and light transmission through the ocular wall.". *Vision Res*, Vol. 31, No. 7-8, pp. 1361–7, 1991.

[Bert 00] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. "Image Inpainting". In: *Proc. of ACM Siggraph 2000*, pp. 417–424, 2000.

[Bori 42] E. G. Boring. *Sensation and perception in the history of experimental psychology*. D. Appleton-Century company, 1942.

[Born 02] R. Bornard, E. Lecan, L. Laborelli, and J.-H. Chenot. "Missing Data Correction in Still Images and Image Sequences". In: *Proc of MULTIMEDIA '02*, pp. 355–361, 2002.

[Brai 97] D. H. Brainard. "The Psychophysics Toolbox". *Spatial Vision*, Vol. 10, pp. 433–436, 1997.

[Burt 83] P. Burt and E. Adelson. "The Laplacian Pyramid as a Compact Image Code". *IEEE Transactions on Communication*, Vol. 31, No. 4, pp. 532–540, 1983.

[Capp 97] R. D. Cappels and M. Hernandez. "System and method for adjusting the output of an output device to compensate for ambient illumination". U.S Patent 5,670,985, Sep. 1997.

[Chan 04] N. Chang, I. Choi, and H. Shim. "DLS: dynamic backlight luminance scaling of liquid crystal display". *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, Vol. 12, No. 8, pp. 837–846, 2004.

[Cipi 10] Z. Cipiloglu, A. Bulbul, and T. Capin. "A Framework for Enhancing Depth Perception in Computer Graphics". *Proc. 6th ACM Symp. Applied Perception in Graphics and Visualization (APGV)*, pp. 141–148, 2010.

[Clar 28] A. B. Clark, H. Nyquist, and D. K. Gannett. "Electrical Picture-Transmitting System". U.S Patent 1,691,147, Nov. 1928.

[Daly 04] S. Daly and X. Feng. "Decontouring: Prevention and removal of false contour artifacts". In: *Proc. of Human Vision and Electronic Imaging IX*, pp. 130–149, 2004.

[Daly 93] S. Daly. "The Visible Differences Predictor: An Algorithm for the Assessment of Image Fidelity". *Digital Images and Human Vision*, pp. 179–206, 1993.

[De M 72] L. E. De Marsh. "Optimum Telecine Transfer Characteristics". *Journal of the Society of Motion Picture and Television Engineers*, Vol. 81, pp. 784–787, 1972.

[Debe 97] P. Debevec and J. Malik. "Recovering High Dynamic Range Radiance Maps from Photographs". In: *Proc. of ACM SIGGRAPH*, pp. 369–378, 1997.

[Devl 06] K. Devlin, A. Chalmers, and E. Reinhard. "Visual Calibration and Correction for Ambient Illumination". *ACM Trans. Applied Perception*, Vol. 3, No. 4, pp. 429–452, 2006.

[Didy 08] P. Didyk, R. Mantiuk, M. Hein, and H. Seidel. "Enhancement of Bright Video Features for HDR Displays". In: *Eurographics Symposium on Rendering*, p. to appear, 2008.

[Dill 95] T. W. Dillon and H. H. Emurian. "Reports of Visual Fatigue Resulting From Use of a Video Display Unit". *Computers in Human Behavior*, Vol. 11, No. 1, pp. 77–84, 1995.

[Dill 96] T. W. Dillon and H. H. Emurian. "Some Factors Affecting Reports of Visual Fatigue Resulting From Use of a VDU". *Computers in Human Behavior*, Vol. 12, No. 1, pp. 49–59, 1996.

[Dror 04] R. Dror, A. Willsky, and E. Adelson. "Statistical characterization of real-world illumination". *Journal of Vision*, Vol. 4, No. 9, pp. 821–837, 2004.

[Eloh 05] M. Eloholma and L. Halonen. "Performance based model for mesopic photometry". Tech. Rep., Helsinki University of Technology, Lighting Laboratory, 2005.

[Fair 02] M. Fairchild and G. M. Johnson. "Meet iCAM: A Next-Generation Color Appearance Model". In: *Proc. IS&T/SID 10th Color Imaging Conference*, pp. 33–38, 2002.

[Feke 06] J. Fekete, C. Sik-Lanyi, and J. Schanda. "Spectral discomfort glare sensitivity under low photopic conditions". *Ophthalmic and Physiological Optics*, Vol. 26, No. 3, pp. 313–317, 2006.

[Ferw 96] J. Ferwerda, S. Pattanaik, P. Shirley, and D. Greenberg. "A Model of Visual Adaptation for Realistic Image Synthesis". In: *Proc. of SIGGRAPH 96*, pp. 249–258, 1996.

[Flan 07] M. Flannagan. "Vision in Night Driving: The Roles of Rod and Cone Photoreceptors". In: *Proc. of 4th International Driving Symposium on Human Factors in Driver Assessment, Training, and Vehicle Design*, 2007.

[Ghos 05] A. Ghosh, M. Trentacoste, H. Seetzen, and W. Heidrich. "Real Illumination from Virtual Environments". In: *Eurographics Symposium on Rendering*, pp. 243–252, 2005.

[Gilc 80] A. L. Gilchrist. "When does perceived lightness depend on perceived spatial arrangement?". *Perception & Psychophysics*, Vol. 28, No. 6, pp. 527–538, 1980.

[Hech 45] S. Hecht and Y. Hsia. "Dark Adaptation Following Light Adaptation to Red and White Lights". *Journal of the Optical Society of America*, Vol. 35, No. 4, pp. 261–267, 1945.

[Heck 06] R. L. Heckaman and M. D. Fairchild. "Expanding Display Color Gamut Beyond the Spectrum Locus". *Color Research and Application*, Vol. 31, No. 6, pp. 475–482, 2006.

[Held 10] R. Held, E. Cooper, J. O'Brien, and M. Banks. "Using Blur to Affect Perceived Distance and Size". *ACM Trans. Graph.*, Vol. 29, No. 2, p. Article 19 (16 pages), 2010.

[Helm 25] H. von Helmholtz. *Physiological Optics / translated from 3d German ed. originally published in 1866*. The Optical Society of America, 1924-25.

[Hess 90] R. Hess, L. Sharpe, and K. Nordby. *Night Vision: Basic, Clinical and Applied Aspects*. Cambridge University Press, 1990.

[Hira 96] A. Hirani and T. Totsuka. "Combining Frequency and Spatial Domain Information for Fast Interactive Image Noise Removal". In: *Proc. of ACM Siggraph '96*, pp. 269–276, 1996.

[Hulb 51] E. O. Hulburt. "Time of Dark Adaptation after Stimulation by Various Brightnesses and Colors". *Journal of the Optical Society of America*, Vol. 41, No. 6, pp. 402–404, 1951.

[Ichi 07] S. Ichihara, N. Kitagawa, and H. Akutsu. "Contrast and depth perception: Effects of texture contrast and area contrast". *Perception*, Vol. 36, pp. 686–695, 2007.

[Iran 06] A. Iranli, W. Lee, and M. Pedram. "Backlight dimming in power-aware mobile displays". In: *Proceedings of the 43rd annual conference on Design automation*, pp. 604–607, ACM New York, NY, USA, 2006.

[Ishi 07] S. Ishihara. *Ishihara's Tests for Colour Deficiency*. Kanehara Trading Inc., 2007.

[ITU 90] ITU. "ITU-R BT.709, Basic Parameter Values for the HDTV Standard for the Studio and for International Programme Exchange". Standard Recommendation 709, International Telecommunication Union, 1990.

[Kalm 93] M. H. Kalmanash. "Backlight for liquid crystal devices". U.S Patent 5,211,463, May 1993.

[Karw 01] W. Karwowski. *International encyclopedia of ergonomics and human factors*. Taylor & Francis, 2001.

[Katy 01] R. H. Katyl, R. G. Greene, B. Yost, and J. P. Krusius. "Correction methods for brightness in electronic display". U.S Patent 6,275,825, Aug. 2001.

[Kero 07] L. Kerofsky and S. Daly. "Brightness Preservation for LCD Backlight Dimming". *Sharp Technical Journal*, No. 95, pp. 50–57, 2007.

[Kunk 11] T. Kunkel, G. Damberg, and L. Johnson. "A High Bit Depth Digital Imaging Pipeline for Vision Research". *Poster in Proc. 8th Symp. Applied Perception in Graphics and Visualization (ACM APGV)*, 2011.

[Ledd 05] P. Ledda, A. Chalmers, T. Troscianko, and H. Seetzen. "Evaluation of Tone Mapping Operators Using A High Dynamic Range Display". *ACM Trans. Graph. (Proc. ACM Siggraph)*, Vol. 24, No. 3, pp. 640–648, 2005.

[Levi 04] A. Levin, D. Lischinski, and Y. Weiss. "Colorization Using Optimization". *ACM Trans. Graph. (Proc. ACM Siggraph)*, Vol. 23, No. 3, pp. 689–694, 2004.

[Li 05] Y. Li, L. Sharan, and E. Adelson. "Compressing and Companding High Dynamic Range Images with Subband Architectures". *ACM Trans. Graph. (Proc. ACM Siggraph)*, Vol. 24, No. 3, pp. 836–844, 2005.

[Like 34] R. Likert, S. Roslow, and G. Murphy. "A Simple and Reliable Method of Scoring the Thurstone Attitude Scales". *Journal of Social Psychology*, Vol. 5, pp. 228–238, 1934.

[Lipp 08] M. G. Lippmann. "Reversible Prints". *Communication at the French Society of Physics: Session of March 20, 1908, J. Phys.*, Vol. 4, No. 7, pp. 821–825, 1908.

[Mant 04a] R. Mantiuk, G. Krawczyk, K. Myszkowski, and H.-P. Seidel. "Perception-Motivated High Dynamic Range Video Encoding". *ACM Trans. Graph. (Proc. of ACM Siggraph)*, Vol. 23, No. 3, pp. 733–741, 2004.

[Mant 04b] R. Mantiuk, K. Myszkowski, and H. Seidel. "Visible Difference Predicator for High Dynamic Range Images". In: *Proc. IEEE International Conference on Systems, Man and Cybernetics*, pp. 2763–2769, 2004.

[Mant 05] R. Mantiuk, S. Daly, K. Myszkowski, and H. Seidel. "Predicting visible differences in high dynamic range images - model and its calibration". In: *Human Vision and Electronic Imaging X, vol. 5666, SPIE Proceedings Series*, pp. 204–214, 2005.

[Mant 06] R. Mantiuk, A. Efremov, G. Krawczyk, K. Myszkowski, and H.-P. Seidel. "Backward Compatible High Dynamic Range MPEG Video Compression". *ACM Trans. Graph. (Proc. of ACM Siggraph)*, Vol. 25, No. 3, pp. 713–723, 2006.

[Mant 08a] R. Mantiuk, S. Daly, and L. Kerofsky. "Display Adaptive Tone Mapping". *ACM Trans. Graph. (Proc. of ACM Siggraph)*, Vol. 27, No. 3, pp. Article 68, 10 pages, 2008.

[Mant 08b] R. Mantiuk and H.-P. Seidel. "Modeling a Generic Tone-mapping Operator". In: *Computer Graphics Forum (Proc. Eurographics)*, pp. 699–708, 2008.

[Mant 09] R. Mantiuk, A. G. Rempel, and W. Heidrich. "Display Considerations for Night and Low-Illumination Viewing". *Proc. 6th Symp. Applied Perception in Graphics and Visualization (ACM APGV)*, pp. 53–58, 2009.

[Mant 11] R. Mantiuk, K. J. Kim, A. G. Rempel, and W. Heidrich. "Calibrated visual metrics for visibility and quality predictions in all luminance conditions". *ACM Trans. Graph.*, Vol. 30, No. 3, p. Article 40 (12 pages), 2011.

[Masi 09] B. Masia, S. Agustin, R. W. Fleming, O. Sorkine, and D. Gutierrez. "Evaluation of reverse tone mapping through varying exposure conditions". *ACM Trans. Graph.*, Vol. 28, No. 5, p. Article 160, 2009.

[Merr 88] R. Merrifield and L. D. Silverstein. "The ABC's of automatic brightness control". In: *Society for Information Display Internatiational Symposium Digest of Technical Papers*, pp. 178–181, 1988.

[Meyl 06] L. Meylan, S. Daly, and S. Süsstrunk. "The reproduction of specular highlights on high dynamic range displays". In: *Proc. of the 14th Color Imaging Conference*, 2006.

[Meyl 07] L. Meylan, S. Daly, and S. Süsstrunk. "Tone mapping for high dynamic range displays". In: *Proc. of Human Vision and Electronic Imaging XII*, 2007.

[Mile 53] P. W. Miles. "Anomalous binocular depth perception due to unequal image brightness". *AMA Archives of Ophthalmology*, Vol. 50, No. 4, pp. 475–478, 1953.

[Novi 69] S. B. Novick. "Tone Reproduction from Colour Telecine Systems". *British Kinematography Sound And Television*, Vol. 51, No. 10, pp. 342–347, 1969.

[Okab 06] S. Okabayashi, H. Miura, N. Sugie, and T. Hatada. "Driver's Perception of Images in Automotive Multicolor Display System". In: *Proc. of IEEE Industry Applications Conf., 41st IAS Annual Meeting*, 2006.

[OShe 94] R. O'Shea, S. Blackburn, and H. Ono. "Contrast as a depth cue". *Vision Research*, Vol. 34, No. 12, pp. 1595–1604, 1994.

[Patt 00] S. Pattanaik, J. Tumblin, H. Yee, and D. Greenberg. "Time-Dependent Visual Adaptation for Realistic Image Display". In: *Proc. of SIGGRAPH'00*, pp. 47–54, 2000.

[Patt 98] S. N. Pattanaik, J. A. Ferwerda, M. D. Fairchild, and D. Greenberg. "A multiscale model of adaptation and spatial vision for realistic image display". In: *Proc. of ACM SIGGRAPH '98*, pp. 287–298, 1998.

[Phil 06] Philips Inc. "Philips showcased latest LCD backlighting improvements at FPD 2006". http://preview.tinyurl.com/yhc4nb, 2006.

[Phil 08] Philips Inc. "Philips Ambilight". http://www.misc.philips.com/ambilight/uk/index2.html, accessed Nov. 19, 2008, 2008.

[Pitt 76] C. R. Pittman. "Ambient light responsive illumination brightness control circuit". U.S Patent 3,962,600, June 1976.

[Poyn 03] C. Poynton. *Digital Video and HDTV: Algorithms and Interfaces*. Morgan Kaufmann, 2003.

[Rein 02] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. "Photographic Tone Reproduction for Digital Images". *ACM Trans. Graph. (Proc. of ACM Siggraph)*, Vol. 21, No. 3, pp. 267–276, 2002.

[Remp 07] A. G. Rempel, M. Trentacoste, H. Seetzen, H. D. Young, W. Heidrich, L. Whitehead, and G. Ward. "Ldr2Hdr: On-the-fly Reverse Tone Mapping of Legacy Video and Photographs". *ACM Trans. Graph. (Proc. of ACM Siggraph)*, Vol. 26, No. 3, pp. Article 39, 6 pages, 2007.

[Remp 08a] A. G. Rempel, W. Heidrich, H. Seetzen, G. J. Ward, and L. A. Whitehead. "Enhancing Dynamic Range of Images". Patent PCT/CA2008/001413, filed July 30, 2008.

[Remp 08b] A. G. Rempel, W. Heidrich, H. Seetzen, G. J. Ward, and L. A. Whitehead. "Enhancing Dynamic Range of Images". Patent US12/182,121, filed July 29, 2008.

[Remp 08c] A. G. Rempel, W. Heidrich, H. Seetzen, G. J. Ward, and L. A. Whitehead. "Enhancing Dynamic Range of Images". Patent US12/183,033, filed July 30, 2008.

[Remp 09] A. G. Rempel, W. Heidrich, and R. Mantiuk. "Video Viewing Preferences for HDR Displays Under Varying Ambient Illumination". *Proc. 6th Symp. Applied Perception in Graphics and Visualization (ACM APGV)*, pp. 45–52, 2009.

[Remp 11a] A. G. Rempel, W. Heidrich, and R. Mantiuk. "The Role of Contrast in the Perceived Depth of Monocular Imagery". Tech. Rep., The University of British Columbia, 2011.

[Remp 11b] A. G. Rempel, W. Heidrich, and R. Mantiuk. "The Role of Contrast in the Perceived Depth of Monocular Imagery". *Poster in Proc. 8th Symp. Applied Perception in Graphics and Visualization (ACM APGV)*, 2011.

[Remp 11c] A. G. Rempel, R. Mantiuk, and W. Heidrich. "Display Considerations for Improved Night Vision Performance". *Proc. 19th Color Imaging Conference*, pp. 191–194, 2011.

[Rich 39] J. P. Richter, I. A. Richter, and R. C. Bell. *The literary works of Leonardo da Vinci*. Oxford university press, 1939.

[Rude 94] D. L. Ruderman. "The statistics of natural images". *Network: Computation in Neural Systems*, Vol. 5, pp. 598–605, 1994.

[Schi 90] J. Schirillo, A. Reeves, and L. Arend. "Perceived lightness, but not brightness, of achromatic surfaces depends on perceived depth information". *Perception & Psychophysics*, Vol. 48, No. 1, pp. 82–90, 1990.

[Schr 00] P. R. Schratz. *Submarine Commander: A Story of World War II and Korea*. University Press of Kentucky, 2000.

[Seet 03] H. Seetzen, L. Whitehead, and G. Ward. "A High Dynamic Range Display Using Low and High Resolution Modulators". In: *Society for Information Display International Symposium Digest of Technical Papers*, pp. 1450–1453, 2003.

[Seet 04] H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, and A. Vorozcovs. "High Dynamic Range Display Systems". *ACM Trans. Graph. (Proc. ACM Siggraph)*, Vol. 23, No. 3, pp. 760–768, 2004.

[Seet 06a] H. Seetzen, H. Li, L. Ye, W. Heidrich, L. Whitehead, and G. Ward. "Observations of Luminance, Contrast, and Amplitude Resolution of Displays". In: *Society for Information Display (SID) Digest*, pp. 1229–1233, 2006.

[Seet 06b] H. Seetzen, H. Li, L. Ye, G. Ward, L. Whitehead, and W. Heidrich. "Guidelines for Contrast, Brightness, and Amplitude Resolution of Displays". In: *Society for Information Display (SID) Digest*, pp. 1229–1233, 2006.

[Seet 07] H. Seetzen. "Personal communication". At trade shows, HDR images were said by observers to look more "plastic" or "three-dimensional". This was true both for experts as well as a naive audience. However, the effect was found to be image-dependent. In particular, images adapted from LDR content often did not exhibit this effect as strongly., 2007.

[Seet 09] H. Seetzen. "High Dynamic Range Display and Projection Systems". Ph.D. Thesis, The University of British Columbia, 2009.

[Spen 94] G. Spenkelink and K. Besuijen. "Brightness: Highest luminance or background luminance". *WWDU94. Book of short papers. Institute of Occupational Health, University of Milan, Milano*, Vol. 2, 1994.

[Stee 93] R. Steen, D. Whitaker, D. Elliott, and J. Wild. "Effect of filters on disability glare". *Ophthalmic and Physiological Optics*, Vol. 13, No. 4, pp. 371–376, 1993.

[Stev 61] S. S. Stevens. "To Honor Fechner and Repeal His Law". *Science*, Vol. 133, No. 3446, pp. 80–86, 1961.

[Stev 63] J. C. Stevens and S. S. Stevens. "Brightness Function: Effects of Adaptation". *Journal of the Optical Society of America*, Vol. 53, No. 3, pp. 375–385, 1963.

[Stoc 00] A. Stockman and L. Sharpe. "The spectral sensitivities of the middle-and long-wavelength-sensitive cones derived from measurements in observers of known genotype". *Vision Research*, Vol. 40, No. 13, pp. 1711–1737, 2000.

[Stok 96] M. Stokes, M. Anderson, and R. Motta. "Standard default color space for the Internet". Tech. Rep., World Wide Web Consortium, 1996.

[Stri 03] J. Stringham, K. Fuld, and A. Wenzel. "Action spectrum for photophobia". *Journal of the Optical Society of America A*, Vol. 20, No. 10, pp. 1852–1858, 2003.

[Sun 05] J. Sun, L. Yuan, J. Jia, and H.-Y. Shum. "Video Enhancement using Per-Pixel Virtual Exposures". *ACM Trans. Graph. (Proc. ACM Siggraph)*, Vol. 24, No. 3, pp. 861–868, 2005.

[Swin 08] S. Swinkels, I. Heynderickx, D. Yeates, and M. Essers. "Ambient Light Control for Mobile Displays". In: *Society for Information Display Internatiational Symposium Digest of Technical Papers*, pp. 1006–1009, 2008.

[Toma 98] C. Tomasi and R. Manduchi. "Bilateral Filtering for Gray and Color Images". In: *Proc. of ICCV '98*, p. 839, 1998.

[Tren 06] M. Trentacoste. "Photometric Image Processing for High Dynamic Range Displays". M.Sc. Thesis, The University of British Columbia, 2006.

[Tren 07] M. Trentacoste, W. Heidrich, L. Whitehead, H. Seetzen, and G. Ward. "Photometric Image Processing for High Dynamic Range Displays". *J. Visual Communication and Image Representation*, Vol. 18, No. 5, pp. 439–451, 2007.

[Trie 62] A. Triesman. "Binocular rivalry and stereoscopic depth perception". *The Quarterly Journal of Experimental Psychology*, Vol. 14, No. 1, pp. 23–37, 1962.

[Tros 91] T. Troscianko, R. Montagnon, J. L. Clerc, E. Malbert, and P.-L. Chanteau. "The role of colour as a monocular depth cue". *Vision Research*, Vol. 31, No. 11, pp. 1923–1930, 1991.

[Tumb 93] J. Tumblin and H. Rushmeier. "Tone Reproduction for Realistic Images". *IEEE Computer Graphics and Applications*, Vol. 13, No. 6, pp. 42–48, 1993.

[van 67] F. van Nes and M. Bouman. "Spatial Modulation Transfer in the Human Eye". *Journal of the Optical Society of America*, Vol. 57, pp. 401–406, 1967.

[vCad 08] M. Čadík, M. Wimmer, L. Neumann, and A. Artusi. "Evaluation of HDR Tone Mapping Methods Using Essential Perceptual Attributes". *Computers & Graphics*, Vol. 32, pp. 330–349, 2008.

[Vos 03] J. Vos. "Reflections on glare". *Lighting Research & Technology*, Vol. 35, No. 2, p. 163, 2003.

[Vos 99] J. Vos and T. van den Berg. "Report on disability glare". *CIE Research Note*, Vol. 135, No. 1, 1999.

[Ward 05] G. Ward and M. Simmons. "JPEG-HDR: A Backwards-Compatible, High Dynamic Range Extension to JPEG". In: *Proc. of Color Imaging Conference (CIC) '05*, 2005.

[Wels 02] T. Welsh, M. Ashikhmin, and K. Mueller. "Transferring Color to Greyscale Images". In: *Proc. of ACM Siggraph '02*, pp. 277–280, 2002.

[Whit 93] D. Whitaker, R. Steen, and D. Elliott. "Light scatter in the normal young, elderly, and cataractous eye demonstrates little wavelength dependency". *Optometry and Vision Science*, Vol. 70, No. 11, pp. 963–968, 1993.

[Wijn 10] M. W. A. Wijntjes and S. C. Pont. "Pointing in Pictorial Space: Quantifying the Perceived Relative Depth Structure in Mono and Stereo Images of Natural Scenes". *ACM Trans. Appl. Percept.*, Vol. 7, No. 4, p. Article 24 (8 pages), 2010.

[Wysz 82] G. Wyszecki and W. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulae*. 1982.

[Yosh 06] A. Yoshida, R. Mantiuk, K. Myszkowski, and H.-P. Seidel. "Analysis of Reproducing Real-World Appearance on Displays of Varying Dynamic Range". In: *Proc. of Eurographics 2006*, pp. 415–426, 2006.

[Yosh 07] A. Yoshida, V. Blanz, K. Myszkowski, and H.-P. Seidel. "Testing tone mapping operators with human-perceived reality". *Journal of Electronic Imaging*, Vol. 16, No. 1, pp. 013004–1 – 013004–14, 2007.