

Evaluating Object Visibility with Simulated Visual Impairment
Using Real and Rendered Scenes

A DISSERTATION
SUBMITTED TO THE FACULTY OF
UNIVERSITY OF MINNESOTA
BY

Christopher Scott Kallie

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Advisers:
Gordon E. Legge Paul R. Schrater

December, 2012

© Christopher Scott Kallie 2012

Acknowledgements

NIH Grant EY017835 supported this research.

Thanks to my advisors, dissertation reviewers, and defense committee members for leadership, encouragement, and feedback during the writing and presentation of this research, including: Gordon E. Legge (primary advisor, committee member, & reviewer), Cheryl A. Olman (committee chair & reviewer), Paul R. Schrater (advisor, committee member, & reviewer), and Albert Yonas (committee member & reviewer).

Thanks to Gordon E. Legge, Cheryl A. Olman, Jennifer M. Scholz, Paul R. Schrater, and Deyue Yu for editing advice.

Thanks to Kathleen H. Briggs, Muzi Chen, and Heejung Park for assistance with research participant recruitment and data collection.

Thanks to Allen M.Y. Cheong, Amy A. Kalia, and Deyue Yu for assistance and advice with designing the physical objects experiment.

Thanks to Amy A. Kalia, MiYoung Kwon, Cheryl A. Olman, and Michael-Paul Schallmo for experimental software and hardware setup advice on the desktop computer experiments.

Thanks to Paul J. Beckmann, Sandeep Dhull, Shane R. Hoversten, Amy A. Kalia, MiYoung Kwon, Cynthia Murphy, Paul R. Schrater, Robert A. Shakespeare, William B. Thompson, and Gregory J. Ward for computational advice and computer code.

Thanks to Daniel J. Kersten, Gordon E. Legge, and Paul R. Schrater for advice on experimental design, data analysis, and image processing.

Thanks to Lawrence G. Kallie and Holly J. Kallie for supplying computer hardware, software, and photographic equipment.

Thanks to Brian J. Allen, Jason G. Kallie, and Kari A. Sundem-Allen for providing shelter and hospitality during the writing of this dissertation.

Thanks to many supporting University of Minnesota entities who made this research possible, including: The Minnesota Laboratory for Low-Vision Research; The Computational Perception and Action Lab; The College of Liberal Arts—Office of Information Technology (especially, Kemal Badur, Carla A. Bates, Joshua Buysse, Chris Neville, Anthony Odufuye, and Eric Peterson); The Minnesota Supercomputing Institute (especially Jim Johnson and Benjamin J. Lynch); The LCSE-MSI Visualization Laboratory (especially Michael R. Knox); The Department of Psychology Administration (especially Guillermo A. De Paz), Accounting (especially Kamran Motevaze), and Facilities Management.

There are many more students, staff, faculty, and friends who deserve credit for supporting me during my academic work and education at the University of Minnesota. Thanks to numerous individuals from an array of departments, labs, courses, offices, centers, and facilities who have contributed their time and talent for my enrichment.

Thanks to my friends, family, and countless individuals who have offered kindness and support during this difficult and enlightening journey. There are many who deserve an abundance of my personal gratitude for keeping me well—I feel truly blessed by your kindness and support. Each of you has been a guiding luminary and spiritual inspiration for me, and I am grateful for receiving such an abundance of unconditional love. The

earning of a Ph.D. was a direct consequence of many, many encouraging and caring people. Without every bit of your effort, this journey would not have been successful.

Dedication

This dissertation is dedicated to those who respect humanity, life, and all of Nature by promoting the return and preservation of our water to its pure and natural existence.

Abstract

Visual accessibility is "the effectiveness with which vision can be used to travel safely through the space and to pursue the intended activities in the space (Legge, Yu, Kallie, Bochslar, & Gage, 2010)." A major issue for visual accessibility involves the visibility of objects under visually challenging conditions, especially those faced by people with impaired vision. This research investigates visibility of objects under conditions of artificially induced blur, and concludes with simple design principles for improving object visibility when vision is compromised by reductions in acuity and contrast sensitivity. Additionally, this research bridges a gap between object visibility studies in physical environments, and studies of visibility using image processing and screen displays of the same objects and scenes. Psychophysical measurements of visibility of objects shown on the desktop display accounted for 58% of the variability in visibility measurements of the same scenes and objects in the physical environment. The technical aspects of virtual modeling—from stimulus conception to photometrically and geometrically accurate displaying of stimulus images—are described.

Chapter 1 of this dissertation provides an overview of this research.

In Chapter 2, experiments are described in which object visibility was tested in a physical environment. The objects were boxes and cylinders. The subjects were normally sighted young adults who wore blur goggles that simulate the effects of visually impaired acuity and contrast sensitivity loss. Interactions between illumination source (i.e., luminaire location), surface reflectance, and object shape led to the conclusion that

curved objects were more visible than flat objects with blurry vision. However we wished to confirm this conclusion for a wider range of viewing conditions.

Chapter 3 extended the work of the physically based study to a computer desktop displayed environment. Several models for digitally simulating physical blur goggles were examined. A high dynamic range (HDR) image of the point spread function (PSF) of the physical blur goggles, convolved with HDR images of the stimuli, plus a slight contrast attenuation, provided the best agreement between visibility performance of subjects viewing the objects on a computer display and the performance of the subjects who viewed the corresponding real objects (Chapter 2). A comparison between photographs and renderings showed no difference between the two image types. The main effects and interaction findings of the physical experiment (from Chapter 2) were largely reproduced with our desktop model. The desktop model was used to study novel conditions (i.e., new object locations and orientations), where it was found that under diffuse overhead illumination near the center of the room, the curved objects were actually less visible than flat objects. In all other cases that were tested, curved objects still provided enhanced visibility over flat objects.

Chapter 4 contains a technical report on the methods used to create photometrically accurate HDR images based on physically based 3D objects, materials, and scenes. It illustrates the 3D modeling, rendering, photography, and photometry methods for reproducing photometrically and geometrically accurate scenes on a calibrated computer display. In addition to modeling and calibration techniques, Chapter 4 illustrates methods for spatial filtering (i.e. visual blur), and non-spatial image processing, including color

balancing and contrast reduction, required for accurate reproduction of the physically based experiment. Example code for all major steps, including spatial and non-spatial filtering, is provided.

Collectively, this series of experiments demonstrated a simple heuristic for designing visually accessible spaces: namely that adding curvature to objects can enhance object visibility under reduced acuity and reduced contrast sensitivity situations. This enhancement, however, is lost under diffuse illumination, where other methods for increasing contrast within or around the object would be more reliable.

Moreover, this research demonstrates a complete workflow from conception of a physical environment to complete virtual modeling of the physical space, as well as digitization of models for visual impairment. This research provides a foundation for further study of visibility in more complex and realistic environments (including motion, as demonstrated in this dissertation's supplemental video archive), linking virtual desktop models to physical environments.

Table of Contents

| | |
|---|-------------|
| Acknowledgements | i |
| Dedication | iv |
| Abstract..... | v |
| Table of Contents | viii |
| List of Tables | xii |
| List of Figures..... | xiv |
| List of Equations | xvi |
| Chapter 1: Introduction—Dissertation Overview | 1 |
| Chapter 2: Identification and Detection of Simple 3D Objects with Severely Blurred Vision..... | 7 |
| Introduction | 8 |
| Illumination | 9 |
| Distance | 11 |
| Color | 12 |
| Height | 13 |
| Shape | 13 |
| Interactions | 13 |
| Contrast | 13 |
| Methods | 14 |
| Participants | 14 |

| | |
|---|-----------|
| Materials..... | 15 |
| Environment | 15 |
| Blur | 15 |
| Stimuli | 19 |
| Illumination | 19 |
| Procedure..... | 20 |
| Photometry | 22 |
| Analysis..... | 22 |
| Results..... | 23 |
| ANOVA Main Effects on Identification | 23 |
| ANOVA Interaction Effects on Identification | 25 |
| Summary of ANOVA Findings..... | 27 |
| Confusion Matrix Analysis | 27 |
| Correlation Analysis..... | 28 |
| Discussion | 30 |
| | |
| Chapter 3: Comparing Physical, Photographed, and Rendered Image Presentations for Object Recognition with Optically and Digitally Blurred Stimuli..... | 37 |
| Introduction | 38 |
| General Method | 43 |
| Scene and Object Modeling and Photorealistic Image Formation | 43 |
| Apparatus and Experimental Protocol..... | 45 |
| Participants | 47 |
| Analysis..... | 48 |
| Experiment 1—Effects of Gaussian Blur | 52 |

| | |
|--|-----------|
| Participants | 52 |
| Method | 53 |
| Results | 54 |
| Discussion | 59 |
| Experiment 2—Effects of Digital Bangerter Blur | 60 |
| Participants | 61 |
| Method | 61 |
| Results | 64 |
| Discussion | 67 |
| Experiment 3—Effects of Digital Bangerter Blur with Contrast Reduction | 69 |
| Participants | 69 |
| Method | 69 |
| Results | 71 |
| Discussion | 79 |
| Experiment 4—Effects of Orientation and Position..... | 80 |
| Participants | 81 |
| Method | 81 |
| Results | 83 |
| Discussion | 86 |
| Summary and Concluding Discussion | 87 |
| Chapter 4: Generating and Displaying Photometrically Accurate Stimulus Images | |
| | 92 |
| Introduction | 93 |
| 4.1. HDR Image Formation | 98 |

| | |
|---|------------|
| 4.1.1. Generating the Radiance Model and Rendering Images | 99 |
| 4.1.2. Obtaining Photometrically Accurate Photographs | 104 |
| 4.2. Calibration Issues and Techniques | 107 |
| 4.2.1. Accurate 8-bit Monitor Calibration | 107 |
| 4.2.1.1. Note on Displaying 32-bit Images | 107 |
| 4.2.1.2. Note on Camera Versus Display Calibration | 107 |
| 4.2.1.3. Note on Recalibration | 108 |
| 4.2.2. Recalibrating and Displaying 8-bit Converted Images | 109 |
| 4.2.2.1. Recalibration Step 1 | 113 |
| 4.2.2.2. Recalibration Step 2 | 113 |
| 4.2.2.3. Recalibration Step 3 | 113 |
| 4.2.2.4. Recalibration Step 4 | 114 |
| 4.2.2.4.1. Note on Coefficient Analysis | 115 |
| 4.2.2.4.2. Note on Tone Mapping | 115 |
| 4.2.2.4.3. Note on Matlab Scripting with Photoshop Javascript | 116 |
| 4.2.2.5. Recalibration Step 5 | 117 |
| 4.3. Spatial Image Filtering Techniques | 118 |
| 4.3.1. Gaussian Blur Filter | 118 |
| 4.3.2. Digital Bangerter Blur Filter | 120 |
| 4.4. Non-Spatial Image Filtering Techniques..... | 125 |
| 4.4.1. Color Balance (White Balance)..... | 125 |
| 4.4.2. Contrast Adjustment..... | 129 |
| Bibliography | 132 |

List of Tables

| | |
|---|-----|
| Table 2.1. Main Effects (Marginal Means) and Selected Interaction Means | 26 |
| Table 2.2. Identification ANOVA | 27 |
| Table 3.1. Coefficients of Determination, Regression, and Eigenvector For Identification and Detection (Experiments 1 - 3) | 55 |
| Table 3.2. Gaussian Blur Means (Experiment 1)..... | 57 |
| Table 3.3. Gaussian Blur Identification ANOVA (Experiment 1) | 57 |
| Table 3.4. Digital Bangerter Blur Means (Experiment 2) | 65 |
| Table 3.5. Digital Bangerter Blur Identification ANOVA (Experiment 2) | 66 |
| Table 3.6. Digital Bangerter Blur Contrast Reduction Means (Experiment 3) | 72 |
| Table 3.7. Digital Bangerter Blur Contrast Reduction Identification ANOVA (Experiment 3) | 72 |
| Table 3.8. Coefficients of Determination, Regression, and Eigenvector for Object Identification with and without Window-Illuminated Boxes | 78 |
| Table 3.9. Digital Bangerter Blur Orientation and Position Means (Experiment 4) | 84 |
| Table 3.10. Digital Bangerter Blur Orientation and Position Identification ANOVA (Experiment 4) | 85 |
| Table 4.1. Sample <i>rif</i> File for Rendering Radiance Images..... | 102 |
| Table 4.2. HDR Image to Recalibrated JPEG Image Script..... | 111 |

| | |
|--|-----|
| Table 4.3. Gaussian Blur Filter Function..... | 120 |
| Table 4.4. Digital Bangerter Blur Function | 123 |
| Table 4.5. Color Balance (White Balance) Function..... | 127 |
| Table 4.6. Contrast Adjustment Function..... | 130 |

List of Figures

| | |
|--|----|
| Figure 2.1. Experimental Setup | 16 |
| Figure 2.2. Stimuli | 17 |
| Figure 2.3. Bangerter Filtered, and False Colored Stimuli | 18 |
| Figure 2.4. Interaction Plots..... | 24 |
| Figure 2.5. Confusion Matrix False Color Maps (by Shape, Height, and Color)..... | 29 |
| Figure 2.6. Lobby..... | 35 |
| Figure 3.1. Layout of Environment and Stimuli | 40 |
| Figure 3.2. Response GUI (Experiments 1-3) | 48 |
| Figure 3.3. Gaussian Blur Stimulus Samples (Experiment 1) | 54 |
| Figure 3.4. Gaussian Blur Identification Interaction Plots (Experiment 1) | 58 |
| Figure 3.5. Digital Bangerter Blur Stimulus Samples (Experiment 2) | 63 |
| Figure 3.6. Digital Bangerter Blur Identification Interaction Plots (Experiment 2) | 67 |
| Figure 3.7. Digital Bangerter Blur with Contrast Reduction Stimulus Samples (Experiment 3) | 71 |
| Figure 3.8. Interaction Plots for Digital Bangerter Blur with Contrast Reduction (Experiment 3) | 73 |
| Figure 3.9. Michelson Contrast for Boxes under Window Illumination | 76 |

| | |
|--|-----|
| Figure 3.10. Identification performance for Physical versus the Bangerter Contrast = 40% Model | 79 |
| Figure 3.11. Rendered Image Samples for Digital Bangerter Blur Orientation and Position (Experiment 4) | 82 |
| Figure 3.12. Interaction Between Object Type and Illumination (Experiment 4) | 86 |
| Figure 3.13. Lobby with Various Digital Filters..... | 91 |
| Figure 4.1. Desktop Luminance Recalibration and Verification | 95 |
| Figure 4.2. HDR Image to 8-bit Display Image Conversion Steps | 97 |
| Figure 4.3. Camera Calibration Apparatus | 105 |
| Figure 4.4. Digital Bangerter PSF..... | 121 |

List of Equations

| | |
|--------------------|-----|
| Equation 3.1 | 70 |
| Equation 4.1 | 110 |
| Equation 4.2 | 113 |
| Equation 4.3 | 129 |
| Equation 4.4 | 129 |
| Equation 4.5 | 129 |
| Equation 4.6 | 130 |
| Equation 4.7 | 130 |
| Equation 4.8 | 130 |
| Equation 4.9 | 130 |

Chapter 1: Introduction—Dissertation Overview

"The effectiveness with which vision can be used to travel safely through the space and to pursue the intended activities in the space" is how visual accessibility is defined (Legge, Yu, Kallie, Bochler, & Gage, 2010). Estimates of the prevalence of visual impairment in the US range between 3.5% (The Eye Diseases Prevalence Research Group, 2004) and 6.4% (Vitale, Cotch, & Sperduto, 2006). As of 2000, over 3.3 million Americans had low vision (The Eye Diseases Prevalence Research Group, 2004). Visual impairments are known to impede mobility, and a large body of research has focused on this issue. As noted by Arditi and Brabyn (2000), making objects more visually salient for pedestrians with visual impairment improves visual accessibility.

The research in this dissertation tested the visibility (detection and identification) of simple convex objects (boxes and cylinders) in a physical classroom environment with participants wearing blurring goggles to simulate visual impairment. Boxes and cylinders were chosen because they are primitive shapes that possess two basic surface types (flat and curved). Photometrically accurate 3D virtual models of the same stimuli, along with digital models of the blurring goggles, were created to test the visibility of objects in the virtual model for comparison with visibility of the objects in the physical environment. Desktop computer models are desirable because they provide an alternative approach to building physical models, which can be cost prohibitive. Furthermore, it is faster to perform psychophysical research on a computer as compared to a physical environment. The desktop model was then used to answer further questions about object visibility as a function of shape, position, and orientation, along with interactions with position and type of illumination within the virtual testing environment.

Chapter 2 describes experiments in which subjects detected and identified objects in a physical environment. A classroom was converted into a testing environment, with major features including the existing overhead fluorescent lighting, and a pair of moveable artificial windows. A 4-foot wide by 40-foot long by 16-inch high sidewalk, painted gray, along one side of the classroom was constructed. Targets were placed near the viewing end of the sidewalk (6 feet from the edge) and subjects sat at various viewing distances along the sidewalk (See Figure 2.1). Target objects for visibility testing consisted of boxes and cylinders constructed out of expanded polystyrene and painted either gray or white (See Figure 2.2).

Subjects wore Bangerter occlusion foils as a model for visual impairment. Viewing was monocular with an effective acuity of 20/900. The Bangerter blur model for low vision was chosen because it reduces visual function by blurring the vision of normally sighted individuals in a predictable and consistent manner. Furthermore, the blur foil is commercially available and has been used in other studies of visual impairment.

The visibility of simple convex objects (boxes and cylinders) was studied with the blur goggles, under two lighting conditions (overhead and artificial window), three distances (10, 17, & 24 feet), two object colors (gray and white), and three object heights (2, 4, & 6 feet). Environmental conditions were chosen to be representative of illuminations, distances, colors, and heights of objects that may be encountered in an indoor institutional setting. The most notable main finding was that cylinders were more visible than boxes. There were also informative interactions between object shape, reflectance, and illumination—suggesting that these three variables interact predictively to produce

contrast variations within and around the object, which, in turn, enhance or mask its visibility under blurred vision.

Chapter 3 describes a study of visibility of objects displayed on a desktop computer screen. It extends the findings of Chapter 2 in two ways: 1) an entirely digital model of the environment (including objects, materials, and lighting) was built and psychophysically tested to determine whether findings in physical reality could be replicated on a computer desktop environment; 2) a digital model of the blur goggles was created and tested. An empirically validated desktop model has advantages over a physical environment for visibility testing, including safety and convenience, rapid novel scene and object prototyping, and faster data collection.

The physical classroom environment, stimulus objects, and blur goggles were entirely simulated in the computer modeling effort. Renderings of the original scene and stimulus objects were produced using Radiance software. Photographs of the same stimulus objects were collected in the original physical environment. Three digital models of the physical blur goggles were created and tested, with each successive model producing visibility (i.e. identification and detection) results closer to those found in the physical experiment (Chapter 2). The digital blur models included: 1) Gaussian blur, 2) digital Bangerter point spread function (PSF), and 3) digital Bangerter PSF with contrast attenuation. It was found that photographs and renderings produced statistically similar results, indicating that rendered images are as good as photographs for studies on visibility.

New room configurations and object positions were rendered, and the digital Bangerter blur model was used to test the generalizability of the earlier finding that cylinders were more visible than boxes. It was found that under suboptimal configurations between observer, background, and indirect overhead illumination, cylinders were in fact less visible than boxes under blurred vision. This finding expands earlier assumptions about object shape, and further implicates the 3-way interaction between shape, reflectance, and illumination for object visibility under blurred vision.

The research presented in Chapter 3 lays groundwork for some of the necessary features required for accurate and reliable desktop reproductions of physical reality for visibility studies. The chapter ends with an example demonstrating its modular approach to visibility studies. It is hoped that more sophisticated models of visual impairment can be modeled digitally with the system described here as a foundation for further research on visual accessibility.

This dissertation ends with Chapter 4, titled "Generating and Displaying Photometrically Accurate Stimulus Images." The chapter describes the technical methods used to generate, filter, and display stimulus images with photometric accuracy compared to the physical environment. Included in the chapter are: 1) methods for rendering and photographing existing and conceivable scenes using a personal computer; 2) calibrating camera and display systems to make images on a desktop screen luminance accurate, as compared to physical scenes; 3) spatial filtering techniques, including a device and method for obtaining the optical transfer function (OTF) of a physical filter medium; and 4) non-spatial filtering techniques—including adjusting the color balance of one image

type to match the color balance of another, and a simple contrast reduction model—both based on the HSV color space.

The broader purpose of Chapter 4 is to enable future investigators to extend studies of visibility in the physical world on a desktop computer environment, with photometrically accurate image presentations that match those found in the real (i.e. physical) world.

Chapter 2:

Identification and Detection of Simple 3D Objects with Severely Blurred Vision

This chapter has been published as: Kallie, C. S., Legge, G. E., & Yu, D. (In Press).

Identification and Detection of Simple 3D Objects with Severely Blurred Vision.

Investigative Ophthalmology & Visual Science.

Introduction

The visual accessibility of a space has been defined as "the effectiveness with which vision can be used to travel safely through the space and to pursue the intended activities in the space (Legge et al., 2010)." A major issue for visual accessibility is the degree to which the objects within spaces are easily and safely identifiable and detectable by people with visual impairment. Visual accessibility includes two major issues: safety and navigability. Much of today's landscape architecture and building architecture has been built with insufficient consideration for the increasing visually impaired population. Simple means for enhancing visibility of objects and improving the accessibility of architectural spaces have been described by Arditi and Brabyn (2000). For example, these authors noted that enhancing edge contrast of steps and other obstacles would improve visibility for those with impaired vision. An ultimate goal of our research on visual accessibility is to create design tools and theoretical support for the creation and optimization of visually accessible spaces. This would include architectural design principles for making key features of spaces more visible for a majority of visually impaired users, and would provide suggestions to designers for how to create and evaluate existing, retrofitted, and new construction projects in residential and public spaces.

The research presented here examines the effects of illumination, viewing distance, color, height, and shape on the identification and detectability of simple convex objects—boxes and cylinders—with severely reduced visual acuity. We tested normally sighted participants with artificially reduced acuity (wearing goggles fitted with Bangerter

occlusion foils) rather than low-vision subjects in order to minimize individual variability. Related research in our lab is examining the generalization of results obtained with artificial acuity reduction to real low vision (Bochsler, Legge, Gage, & Kallie, In Press). Our purpose was to address the impact of important visual variables on object visibility under conditions of low spatial resolution and image contrast, which are generic factors that are often a consequence of low vision. We chose simple convex objects, a simple uniform gray background and gray floor layout, and artificial acuity reduction in this experiment to avoid some of the complexity associated with natural environments and real visual deficits. We reasoned that once we understood the results presented here, we would be better prepared to perform future experiments with visually impaired participants, and understand the effects of more complex real or realistically rendered environments.

We tested the effects of two environmental variables—illumination and viewing distance—and three object variables—color, height, and shape. We briefly review the rationale for examining these factors in the following subsections.

Illumination. Overall changes of illumination within a scene can have major effects on the visibility of objects. Kuyk et al. (1996) studied the performance of 88 low-vision subjects who walked through an obstacle course. When overall illumination was reduced from photopic to mesopic levels, the mean time to complete the obstacle course and the mean number of contacts with obstacles doubled. Another study found that under low illumination, negotiating a step was more difficult for elders than for younger individuals,

even though the elders had good clinical visual acuities (Alexander, Ashton-Miller, Giordani, Guire, & Schultz, 2005).

Because we are concerned with architectural design principles, we focused on lighting arrangement rather than overall illumination level. Architects make decisions on the location and type of lighting in indoor spaces while keeping overall illumination at required photopic levels. The Lighting Handbook (DiLaura, Houser, Mistrick, & Steffy, 2011) recommends a wide range of illuminance values for different spaces and circumstances, including: safety and environmental requirements, adaptation, usage and performance requirements, activity levels, and the most common age of users. As an example, recommended illuminance values for educational building lobbies range between 25 and 200 lux (DiLaura et al., 2011, p. 22.6).

We reasoned that the differences in illumination (diffuse illumination versus localized directional illumination) would produce an effect on object identification and detectability. Specifically, we predicted that the even illumination provided by overhead fluorescent luminaires would produce less object visibility than the direct illumination from our artificial windows. Many people have the intuition that flat ambient illumination is better than directional lighting, and in some cases (e.g. to minimize glare) this intuition is likely to be correct. However, diffuse illumination might lower contrasts for objects. Directional lighting is more likely to produce visible luminance gradients on curved surfaces from self-occlusions between the surface and light source. In our previous study on the visibility of ramps and steps, we found relatively weak effects of lighting arrangement, but one directional lighting condition (far window condition) produced

slightly better performance than overhead lighting (Legge et al., 2010, fig. 5). In the present study, we predicted that there would be overall increases of contrasts in and around the target objects with window illumination, due to the directional nature of the light source, and consequently, we predicted that window illumination would produce better human performance.

Distance. Viewing distance effects are important for walking. As illustrated by Sedgwick, it is ecologically important for people to estimate the distances of objects that are resting on the ground (Sedgwick, A, 1986). Pedestrians must be able to see an object (or obstacle) early enough to avoid a collision and to plan a route around it. A number of studies showed that people slow down (Alexander et al., 2005; Fuhr, Liu, & Kuyk, 2007; Klein, Klein, Lee, & Cruickshanks, 1998), or display more careful stepping behaviors (Vale, Scally, Buckley, & Elliott, 2008), when experiencing reductions in vision. Those findings suggest that people slow down in order to maintain necessary reaction times to avoid collisions. The detection distance problem has also been studied in the context of orientation and mobility (O&M) training efficacy (Goodrich & Ludt, 2003; Ludt & Goodrich, 2002). Goodrich and Ludt noted the importance of detecting obstacles early enough to avoid incidents. In one training study (Ludt & Goodrich, 2002), the authors were able to train the majority of visually impaired pedestrians to detect drop-offs, obstacles placed on the ground, and overhangs at sufficient distances to avoid collisions.

Given low-resolution viewing, visibility at distances that give reasonable response times during walking are important for visual accessibility. We predicted that closer

objects would be more detectable and recognizable than distant objects, due to their increased angular size on the retina.

Color. Color can also play an important role in low-vision object recognition (Wurm, Legge, Isenberg, & Luebker, 1993), and possibly object detection. In a study comparing normally sighted individuals and age-related macular degeneration (AMD) participants, those with AMD showed an increased ability to identify objects in scenes with color (versus gray-scale equivalents) (Boucart, Desprez, Hladiuk, & Desmettre, 2008). Similarly, a meta-analysis (examining a variety of recognition tasks including both normal and low-vision experiments) of color (versus gray-scale) and object recognition showed that color enhances object recognition (Bramão, Reis, Petersson, & Faisca, 2011). In the two aforementioned studies, color referred to hue and saturation. In the present study, "color" refers to gray and white objects.

In our study, we tested the effects of gray and white objects against a gray background, predicting that white objects would be more detectable, on average, than gray objects. We made this prediction because white objects create higher edge contrast with our gray background (under the majority of our lighting conditions). Internal contrast may also be more visible on white objects than the darker (gray) objects. All else being equal, luminance ratios, and hence internal contrasts, should be the same across gray and white objects. However, it is possible that the same contrast on a brighter surface may be slightly more visible than on a darker surface. For example, using the Atick and Redlich contrast sensitivity model (Atick & Redlich, 1992), for a sinusoidal

contrast grating of 10 c/deg, a 10% contrast gradient is easily visible at 10 cd/m², but at 1 cd/m², the gradient is close to threshold visibility, and becomes invisible at 0.1 cd/m².

Height. In the present study, the objects were placed on the ground plane. They varied in height from 2 to 6 feet (0.61 to 1.83 meters). Similar to our prediction for distance, we predicted that taller objects would be more identifiable and detectable than shorter objects, due to their larger visual angles.

Shape. We restricted attention to convex objects (i.e. objects without any concavities) to avoid complications of self-occlusion within concavities. We chose boxes and cylinders because they possess two different basic surface types: flat and curved. When flat objects are illuminated from afar (e.g. a box illuminated by a luminaire located a few meters away), each of the reflecting flat surfaces of the box will have relatively uniform surface luminance and hence low internal contrast. However, a cylinder of the same size and position, seen from the same viewpoint will produce a higher luminance gradient across its curved surface, potentially producing an internal contrast cue. For this reason, if internal contrast were relevant to object detection, we would expect cylinders to be more identifiable and detectable than boxes.

Interactions. We have made predictions about the main effects of our independent variables. In addition, from informal observations of our stimuli, we expected to find noteworthy interactions, which will be discussed in the Results section.

Contrast. We based some of our predictions listed above on contrast because of its presumptive importance in identification and detection. Even mild contrast reductions in everyday scenes—especially those with veiling light producing glare—can have

detrimental effects on object detection (Brabyn, Schneck, Haegerstrom-Portnoy, & Lott, 2004). Arditi and Brabyn (2000) noted that universal design principles for visual accessibility should address (and increase) contrast for obstacles and landmark features. They specifically noted that using lighting and color to increase contrast for steps and curbs is necessary to make those objects more visible for both visually impaired and normally sighted people.

Contrast between object and background has been studied in AMD patients, showing that enhanced contrast between object and background had a positive effect on detection of objects in 2D images (Bordier, Petra, Dauxerre, Vital-Durand, & Knoblauch, 2011; Tran, Guyader, Guerin, Despretz, & Boucart, 2011). Kuyk et al. (1996) noted that their low-vision subjects made more contacts with low-contrast objects than high-contrast objects while walking through an obstacle course.

To summarize, we were interested in the impact of low resolution (low acuity) on real-world object visibility. We investigated the impact on object recognition and detection of important environmental variables (lighting arrangement and viewing distance) and object variables (shape, color, and size). In making predictions, we were guided by the expectation that stimulus contrast and angular size would determine performance.

Methods

Participants

Eleven normally sighted young adults (college undergraduates) were recruited on campus, and were paid for their participation. All participants had normal or corrected to

normal visual acuity and no known eye problems. Participants gave informed consent, the experimental protocol was approved by the University of Minnesota's IRB, and it followed the tenets of the Declaration of Helsinki.

Materials

Environment. A 4-foot wide by 32-foot long by 16-inch high (1.22-meter wide by 9.75-meter long by 0.41-meter high) sidewalk was constructed out of gray painted 4 by 8 feet (1.22 by 2.44 meter) wooden staging platforms. Gray paint was Valspar satin light gray porch and floor enamel. The platforms were placed lengthwise along the long side of a 10.13 by 5.66 meter classroom in a layout shown in Figure 2.1 (schematic of room, objects and viewing positions), and Figure 2.2 (photographs from the 24-foot, or 7.32-meter viewing location). Polyester "silver gray" colored felt fabric was hung at the end of the sidewalk forming the visual background for the target objects. The felt was nearly the same color as the gray paint. A 1.37 by 2.29 meter piece of felt was hung at the far end, and a 1.27 by 2.29 meter piece was hung on the adjacent side wall near the end (seen in Figures 2.1, 2.2, & 2.3). The purpose of the felt was to create a uniform background, reducing contrast cues brought on by irregular background features.

Blur. Using 2 overlapping Bangerter occlusion foils (The Fresnel Prism and Lens Company, model number: <0.1), blur goggles were constructed from a pair of welding goggles. Individual foils were found to decrease acuity and contrast sensitivity (Odell, Leske, Hatt, Adams, & Holmes, 2008), and our double layering of the foils produced an approximate Snellen acuity of 20/900, determined from measurements with the Lighthouse Distance Visual Acuity Chart, and Pelli-Robson contrast sensitivity of 0.6

previously measured in our lab (Legge et al., 2010). This level of severe acuity reduction would be categorized as "profound low vision" by ICD-9-CM (*The International classification of diseases, 9th revision, clinical modification, 1986*), and by the ICO (International Council of Ophthalmology, Colenbrander, 2010). We chose to study severe blur for two reasons. First, a population-based study of 2520 older adults indicated that people were not disabled in mobility tasks unless acuity was poorer than 20/200 (West et al., 2002).




Figure 2.1. Experimental Setup. This diagram shows the dimensions and essential objects in our testing space.

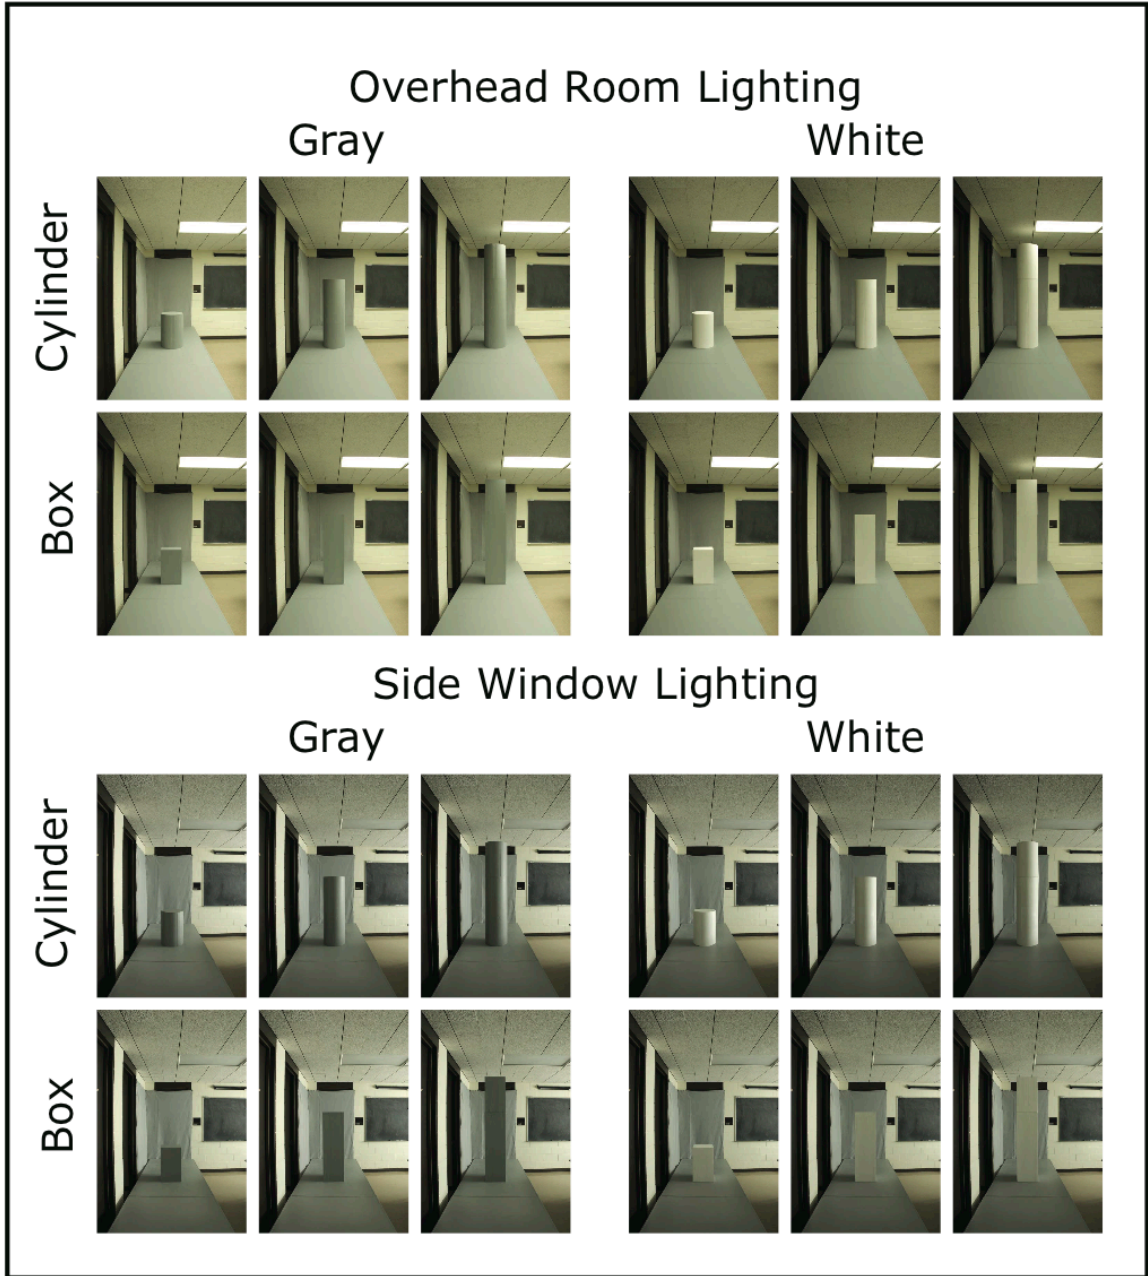


Figure 2.2. Stimuli. This figure shows the experimental stimuli from one viewing distance, including the 2 lighting conditions, 2 colors, 2 shapes, and 3 heights.

Second, pilot testing with milder blur (one Bangerter occlusion foil) with an equivalent acuity of 20/140 yielded performance levels too close to ceiling in our task to be useful. We note that although people with profound low vision often use nonvisual mobility aids such as a white cane or dog guide, they also use their remaining vision. Orientation and mobility specialists have often noted the preference of people to rely on residual vision, sometimes to their detriment (Ludt & Goodrich, 2002).

Digital approximations of two objects and scenes viewed through the goggles are shown in Figure 2.3.

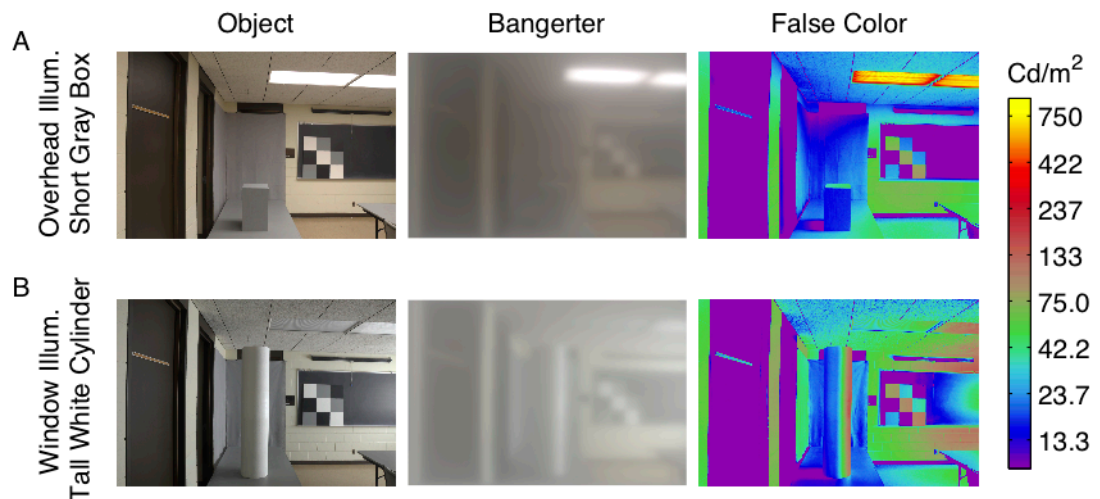


Figure 2.3. Bangerter Filtered, and False Colored Stimuli. Panel A shows one of the least detectible conditions (Overhead Illumination, Short Gray Box), and panel B shows one of the most detectible conditions (Window Illumination, Tall White Cylinder). The same scenes are shown through a digital Bangerter Filter approximation (for demonstration purposes in the center column), illustrating what our research subjects saw through the goggles. Finally, luminance mapped false color images are shown in the right column. All images were created using Anyhere Software (Ward Larson, n.d.-a); blur filtered images were created using a Nikon SLR, Anyhere Software, and Matlab (*MATLAB R2011a*, 2011).

Stimuli. Solid boxes and cylinders were constructed out of expanded polystyrene using a nichrome hot wire foam cutter. We reasoned that for a fair comparison of the effects of shape, boxes and cylinders should be matched for volumes and heights. In the real world, for example, a 120-liter trash container has the same volume and height regardless of whether it has a rectangular or circular footprint. Boxes had square bases with 0.365-meter sides, and cylinders were 0.406 meters in diameter. Three object heights were 2, 4, and 6 feet (0.61, 1.22, and 1.83 meters). Our short objects were similar in size to a typical bench, while medium height objects were similar in size to trash receptacles, and the tall objects were representative of larger objects such as people and structural building components. Photographs of the twelve objects are shown under the two lighting conditions (described in the next subsection) in Figure 2.2.

Box orientation was frontal planar, providing minimal internal object contrasts when illuminated from 2 meters or more. Cylinders, by nature, reveal almost 180 degrees of their profile from any viewing angle at moderate distances, providing an opportunity for self-occlusions with direct lighting, thereby maximizing internal contrast values. We chose only one orientation for boxes (frontal planar) because it produced a worse case scenario for internal contrast. Different viewing angles of boxes would improve internal contrast measurements.

Illumination. Many residential, commercial and institutional spaces make use of both artificial and natural lighting. Objects and spaces under diverse lighting conditions vary dramatically in their visual appearance over different lighting situations. We reasoned that two common lighting conditions would include typical institutional overhead

fluorescent illumination, and side illumination from windows (configurations shown in Figure 2.1). Overhead lighting was preexisting, fitted with recessed acrylic prismatic 4 lamp SP41 fluorescent luminaires.

We simulated window lighting with two movable box-shaped luminaires built out of sheet metal, containing 12 25-watt fluorescent tubes in each, painted flat white on the inside. Calendared acrylic sheet was fixed to the 0.91 by 0.91-meter apertures, which were 0.25 meters in front of the fluorescent tubes. Tubes were SP65 fluorescents. The windows produced room illumination similar to north-facing windows overlooking a sunny, snow-covered landscape. Mean luminance of each window was 785 candelas per square meter.

Luminance levels for two of our viewing conditions, including targets and background are shown in Figure 2.3, where false-color images show pixel-based luminance levels of the target objects and backgrounds. The checkerboard pattern to the side of the target objects was used for photometric calibration and analysis. The average luminance for our experimental stimuli ranged between 34 and 77 cd/m^2 .

Procedure

Before the experiment, participants were shown all 12 objects without goggles. There were 2 lighting conditions and 3 distances, making 6 blocks. Viewing distances were 10, 17, and 24 feet (3.05, 5.18, and 7.32 meters). Prior to the first trial of each block, at least one object of each height, size, and color were presented to the participant without blur goggles at the seated testing position. This preview was intended to ensure that the participant was familiar with the appearance of objects for the condition in question.

Following the preview, the participant put on earmuffs with amplified white noise and the goggles. The earmuffs attenuated sound by 29db, and the white noise further masked any sounds made by the experimenter while placing objects on the stage. Viewing was monocular, with the dominant eye chosen by a pointing task. An opaque lens on the blur goggles occluded the other eye. We chose monocular vision because many visually impaired people rely on their better eye for visual tasks.

Before each trial, the participant was asked to turn his/her head toward the wall while the next test object was placed in position. Once the object was positioned, the participant was asked to turn his/her head and view the object for up to 4 seconds. After approximately 4 seconds, the participant was asked to turn away and respond, which rarely happened because participants almost always responded before the end of the 4-second viewing interval.

There were 15 trials in each block, including 12 objects (3 heights, by 2 colors, by 2 shapes) and 3 catch trials with no objects present. Trials in each block were randomized for each subject, and blocks were randomized between subjects to reduce order effects. There were four judgments on each trial—first a present/absent judgment for detection, second a confidence rating on detection (1 not confident to 5 very confident), third a shape choice (box or cylinder) for identification, and fourth a height choice (short, medium, or tall) for identification. Although confidence ratings were recorded and d' was measured for the detection task, they were not analyzed in the present paper. We found that d' measurements correlated highly with hit rate, and in order to examine the pattern of results more closely, we opted for a confusion matrix analysis

instead. If the subject gave an "absent" response, they were not required to report on shape or height. Feedback was not given.

Photometry

A Minolta CS-100 Colorimeter was used to sample luminance values across the viewable surfaces of the objects. A luminance profile was measured and Michelson contrasts were estimated for four boundary locations around the object (top, bottom, left, and right) and also internal to the facing surface of the object.

Analysis

Proportions correct for identification were calculated for each subject and factor combination (illumination, distance, color, height, & shape). Identification proportions correct were arcsine transformed for a 5-factor, within subjects interaction model ANOVA. We restricted our model to only the 5 main factors and selected interactions that we found particularly interesting. Interaction plots and means tables were constructed for graphical and tabular viewing and post-hoc analysis. Confusion matrix false-color maps were created to examine the distribution of confusable object types, and to show errors of identification that still had correct object detections. Identification proportion correct was correlated with 6 Michelson contrast measurements and angular size for objects. Internal contrast was measured from brightest and dimmest luminance samples within each object. Boundary contrasts were measured between the object and background from adjacent luminance samples at several locations on the perimeter of the object at 4 locations (top, bottom, left, & right).

Results

We focused primarily on object-identification performance, represented by the proportion of correctly identified objects. An identification response was scored correct only if the subject identified both the shape (box or cylinder) and the height (short, medium, or tall) correctly. We will first discuss the identification results in detail, examining marginal means (Table 2.1), and an analysis of variance (ANOVA, Table 2.2) for the 5 main factors and selected interactions. Later, we will examine the data with a confusion matrix analysis, and finally show significant correlations between identification performance and the important variables of contrast and visual angle.

ANOVA Main effects on Identification. There were significant main effects for viewing distance, $F(2, 766) = 11.19, p < .0001$, and height, $F(2, 766) = 5.80, p = .0031$ (Table 2.2 & Figure 2.4). Performance decreased with increasing distance. Mean proportions correct for 10, 17, and 24 feet (3.05, 5.18, and 7.32 meters) were .42, .31, and .26 respectively. Mean proportions correct for short, medium, and tall objects were .26, .36, and .36 respectively. See Figure 2.4A, Table 2.1A, and Table 2.2 for graphical and numerical results. We predicted increasing detectability for increasing heights and decreasing distances. The main effects support our predictions. No significant interaction effects for distance and height were found, $F(4, 766) = 0.21, p = .9303$. We will return to distance and height in the correlation analysis below, where we will combine distance and height in an analysis of the effect of visual angle.




Figure 2.4. Interaction Plots. Panel A shows identification proportion correct versus distance and height. Panels B & C show proportion correct versus illumination, shape, and color. Error bars represent ± 2 SE.

Surprisingly, our prediction of a main effect of lighting arrangement (i.e., illumination) was not confirmed, $F(1, 766) = 0.55, p = .4587$; identification performance

did not differ significantly for overhead and "window" lighting. But there were strong interaction effects for illumination by color, and illumination by shape. These interactions can be seen in Table 2.1, as well as Figure 2.4. We will return to these informative interaction effects, and report their statistics, after addressing the remaining main effects.

The main effect for object color, $F(1, 766) = 132.04, p < .0001$, gray or white, was evident, as shown in Figure 2.4. For example, in 3 out of the 4 illumination by shape pairings, white was easier to detect than gray (Figures 2.4B & 2.4C). Mean proportions correct for white and gray were .49 and .17, respectively, as shown in Table 2.1A.

Similarly, a main effect for shape, $F(1, 766) = 119.57, p < .0001$, was also evident. Side-by-side comparison of two graphs (Figures 2.4B & 2.4C) shows that cylinders (Figure 2.4C) were easier to detect than boxes (Figure 2.4B). Mean proportions correct for boxes and cylinders were .18 and .48, respectively, as shown in Table 2.1A.

A main effect for subject, $F(10, 766) = 3.70, p = .0001$, was also observed.

ANOVA Interaction Effects on Identification. Three 2-way interactions: illumination by color, $F(1, 766) = 33.01, p < .0001$; illumination by shape, $F(1, 766) = 39.71, p < .0001$; and color by shape, $F(1, 766) = 19.79, p < .0001$; along with the 3-way interaction, illumination by color by shape, $F(1, 766) = 9.93, p = .0017$; were all significant (Figure 2.4B-C). Most notable are the 2-way interaction between shape and illumination (2 red dashed lines in Figure 2.4B-C), and the 3-way interaction between color, shape, and illumination (4 solid blue lines in Figure 2.4B-C). All 8 values for the 3-way interaction are shown in Table 2.1C and Figure 2.4B-C. The three-way interaction is noteworthy because it represents a simple demonstration of the effects of three important

variables (illumination, color, and geometry), which we argue must be considered together to predict visibility. We will return to this important issue in the Discussion.

| Main Effects Means (A) | Category | Effect | Mean | N-trials | | | | |
|--|----------------------|--|------------------------------|--------------|----------|----------|------|----------|
| | Illumination | overhead | | 0.32 | 396 | | | |
| | | window | | 0.34 | 396 | | | |
| | Distance | 10 | | 0.42 | 264 | | | |
| | | 17 | | 0.31 | 264 | | | |
| | | 24 | | 0.26 | 264 | | | |
| | Color | white | | 0.49 | 396 | | | |
| | | gray | | 0.17 | 396 | | | |
| | Height | medium | | 0.36 | 264 | | | |
| tall | | | 0.36 | 264 | | | | |
| short | | | 0.27 | 264 | | | | |
| Shape | box | | 0.18 | 396 | | | | |
| | cylinder | | 0.48 | 396 | | | | |
| Selected 2-way Interaction Effects Means (B) | | | | | | | | |
| Selected 2-way Interaction Effects Means (B) | Illumination x Shape | Illumination | Shape | Mean | N-trials | | | |
| | | overhead | box | 0.26 | 198 | | | |
| | | overhead | cylinder | 0.38 | 198 | | | |
| | | window | box | 0.11 | 198 | | | |
| | | window | cylinder | 0.58 | 198 | | | |
| | Distance x Height | Distance | Height | Mean | N-trials | | | |
| | | 10 | medium | 0.44 | 88 | | | |
| | | 10 | tall | 0.45 | 88 | | | |
| | | 10 | short | 0.35 | 88 | | | |
| | | 17 | medium | 0.33 | 88 | | | |
| | | 17 | tall | 0.35 | 88 | | | |
| | | 17 | short | 0.26 | 88 | | | |
| | | 24 | medium | 0.32 | 88 | | | |
| | | 24 | tall | 0.28 | 88 | | | |
| | | 24 | short | 0.18 | 88 | | | |
| | | Selected 3-way Interaction Effects Means (C) | | | | | | |
| | | Selected 3-way Interaction Effects Means (C) | Illumination x Color x Shape | Illumination | Color | Shape | Mean | N-trials |
| | | | | overhead | white | box | 0.47 | 99 |
| | | | | overhead | white | cylinder | 0.64 | 99 |
| overhead | gray | | | box | 0.04 | 99 | | |
| overhead | gray | | | cylinder | 0.13 | 99 | | |
| window | white | | | box | 0.08 | 99 | | |
| window | white | | | cylinder | 0.76 | 99 | | |
| window | gray | | | box | 0.13 | 99 | | |
| window | gray | | | cylinder | 0.39 | 99 | | |

Table 2.1. Main Effects (Marginal Means) and Selected Interaction Means. N refers to the number of trials used to calculate the mean.

| Source | Sum Sq. | d.f. | Mean Sq. | F | Prob>F |
|--------------------|---------|------|----------|--------|--------|
| subject | 0.13059 | 10 | 0.01306 | 3.7 | 0.0001 |
| illumination | 0.00194 | 1 | 0.00194 | 0.55 | 0.4587 |
| distance | 0.079 | 2 | 0.0395 | 11.19 | 0 |
| color | 0.46609 | 1 | 0.46609 | 132.04 | 0 |
| height | 0.04098 | 2 | 0.02049 | 5.8 | 0.0031 |
| shape | 0.42208 | 1 | 0.42208 | 119.57 | 0 |
| illumination*color | 0.11652 | 1 | 0.11652 | 33.01 | 0 |
| illumination*shape | 0.14017 | 1 | 0.14017 | 39.71 | 0 |
| distance*height | 0.00303 | 4 | 0.00076 | 0.21 | 0.9303 |
| color*shape | 0.06984 | 1 | 0.06984 | 19.79 | 0 |
| illum*color*shape | 0.03504 | 1 | 0.03504 | 9.93 | 0.0017 |
| Error | 2.70398 | 766 | 0.00353 | | |
| Total | 4.20926 | 791 | | | |

Table 2.2 *Identification ANOVA*. Data were arcsine transformed. This ANOVA table is also reported in paragraph form in the Results section.

Summary of ANOVA Findings. The most important findings from the ANOVA were: 1) cylinders were always easier to identify than boxes. 2) Although white was generally easier to see than gray, under window lighting, gray boxes were easier to identify than white boxes—even though the background was always gray. 3) As expected, performance was better for larger objects and nearer viewing distances. 4) Performance did not differ between overhead and “window” lighting.

Confusion Matrix Analysis. The preceding analysis dealt with correct identification of the stimuli. In this section, we report all possible stimulus/response combinations with confusion matrix false-color maps (Figure 2.5). The confusions include correct identifications (which were analyzed above, and labeled Hits here), Correct Rejections, False Alarms, Misses, and Identification Errors. The remaining cells include shape confusions, height confusions and joint shape and height confusions, which were left as unlabeled cells in Figure 2.5.

There are three main findings from the confusion matrix analysis. First, detection errors were mostly misses and not false alarms, that is, subjects sometimes failed to see objects that were present, but rarely reported seeing objects on the catch trials. Second, most of the confusions were between shapes, and few confusions were between heights, or between shapes and heights. Third, it is interesting to see how the miss rate changed with lighting condition, object shape, and object color. In the overhead matrices (Figures 2.5A, 2.5B, & 2.5C), there were many misses for gray objects, and few misses for white objects. However, in the window matrices (Figures 2.5D, 2.5E, & 2.5F), the pattern of misses is different, where boxes have many misses and cylinders have far fewer misses. This finding further indicates illumination, color, and geometry—collectively—as important variables to predict visibility.

Correlation Analysis. We computed the correlations between identification proportion correct and several measures of stimulus contrast. We found that 58.3% of the variability in proportion correct performance was accounted for by Michelson contrast. Not surprisingly, the left edge Michelson contrast accounted for very little (2.46%), due to the fact that illumination came primarily from the right side of the objects, which can be seen in both overhead and window illumination (see Figure 2.2). Michelson contrasts at the top edge and bottom edge of the objects accounted for 43.1% and 42.1% of the variance, respectively. Michelson contrast on the Right edge of the objects and Internal to the front surface of the object accounted for 57.7% and 57.5% of the variance respectively. The variance accounted for by the maximum of the measured Michelson

contrast values was 58.3%. Maximum Michelson contrast refers to the greatest of the five measured contrast values for a given target and viewing condition.

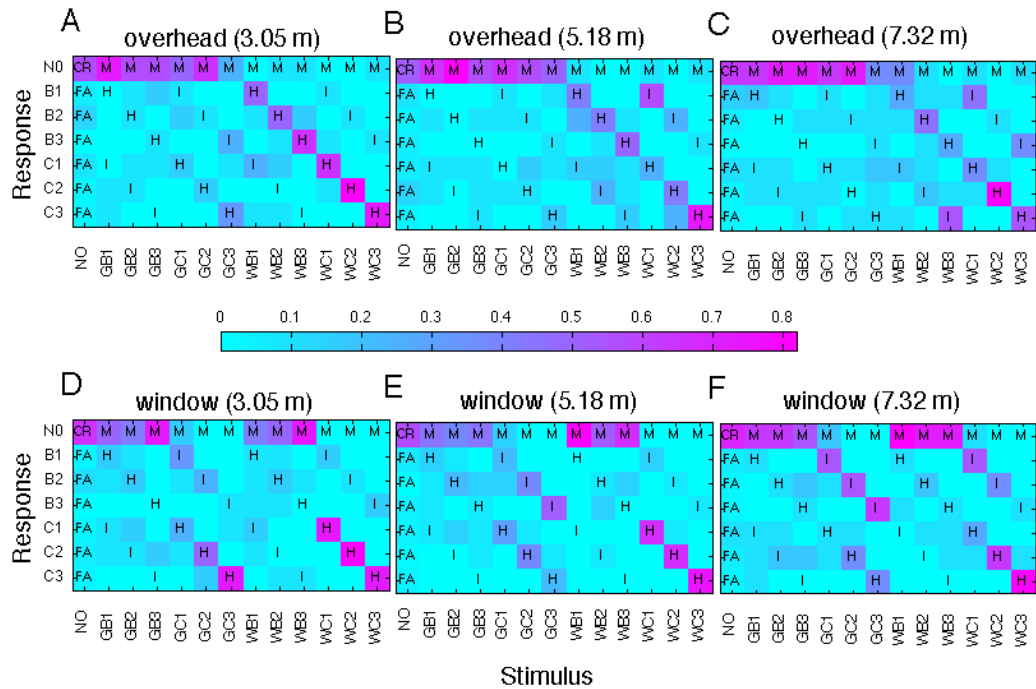


Figure 2.5. Confusion Matrix False Color Maps (by Shape, Height, and Color). This figure shows confusion matrix color maps for proportion correct for the two object types, three heights, and two colors. Key: H = Hits, CR = Correct Rejections, FA = False Alarms, M = Miss, & I = Identification errors; NO = No Object, G = Gray and W = White. B = Box, C = Cylinder, 1 = Short, 2 = Medium, & 3 = Tall.

Visual angle, computed from the height and distance of each object, accounted for 5.8% of the variability in identification. A stepwise regression analysis of all seven predictors (i.e., six contrasts plus visual angle) showed that two predictors: internal contrast plus visual angle, accounted for 63.8% of the variability in identification

performance. The remaining contrast values were highly correlated with one another, and therefore did not explain significantly more of the variability.

Discussion

We are interested in object recognition in low-resolution (low-acuity) viewing because of its potential relevance to low vision and visual accessibility.

We identified several variables which we expect to be important in low-resolution object recognition, including illumination, distance, color, height, and shape. In addition to main effects, we have identified important interactions that occur in real spaces. Our findings suggest three relatively simple heuristics for designing visually accessible spaces, as noted in the following three paragraphs.

First, and of most importance, contrast is a dominating factor for detection and identification of convex objects with blurred vision. The higher the contrast, the better the visibility.

However, the emphasis on contrast raises a second more subtle point. We would typically think of producing good visibility of an object by selecting its surface reflectance (color) to be different from the background e.g., the white objects seen against the gray background in our experiments. However, contrast between object and background is also strongly dependent on the lighting arrangement. Consider, for example, the case in our experiment in which performance is actually better for a gray object on gray background, than for a white object on gray background (results shown in Figure 2.4B-window). The interesting point here, from a design perspective, is that the typical notion that target/background contrast is always improved by increasing the

reflectance difference between object and background can sometimes be wrong. This is why the interaction effects are important to consider for visual saliency. Examples shown in Figure 2.2 illustrate this interesting and instructive reversal. This type of situation can occur in a variety of overhead and window illumination settings, such as the common occurrence when a low-reflectance object has more direct illumination than a high-reflectance object.

Third, for detection and identification of objects with blurred vision, the angular size of objects plays an ecologically important role. In our experiment, we found that larger and closer objects are easier to see. In some cases, it may be possible to predict the viewing distance at which a person with a given level of acuity would be able to recognize an object or object feature of a given angular size. Given information on typical pedestrian travel paths and walking speeds, such predictions could be useful in assessing the likelihood that hazards would go undetected visually. We included a concrete example in Legge et al. (2010, fig. 11), where we illustrated the relationship between viewing distance and acuity for detecting a step down.

Kuyk et al. (1996) and Hassan et al. (2002) both examined the effects of number of collisions as a function of object height in studies in which low-vision subjects walked through an obstacle course. This paradigm differs from ours in which subjects were asked to recognize objects from a distance. Kuyk et al. (1996) tested a heterogeneous group of veterans. They found between-subject main effects and interactions for number of collisions as a function of object type (head level objects, floor level walk around objects, and step-over objects). Hassan et al. (2002) tested AMD patients and found no significant

differences from a group of age-matched normally sighted controls in walking speed and number of obstacle contacts. The lack of difference between the groups and lack of an effect of obstacle height may indicate that the acuity of the AMD subjects was adequate to detect the obstacles when closely approached. The stronger effect of object height in our study is likely due to the greater importance of object angular size in measures of object visibility at a distance.

In our study, the data indicate that cylinders are much more visible than boxes (i.e., curved surfaces are more visible than frontal-planar flat surfaces). Our finding that cylinders were more visible than boxes with overhead and window illumination suggests that curves may be more generally visible than planes, as long as the surface normal of the object curves away from the source of illumination and that the illumination of the space is not completely diffuse. Further studies on orientations of flat objects and placements of curved objects in a variety of illuminations should be performed in order to determine the generalizability of our findings.

Do we expect our findings to generalize to low vision? Findings from our previous studies of the visibility of ramps and steps suggest the affirmative. In our initial study with the steps and ramps (Legge et al., 2010), 48 normally sighted subjects were tested on the recognition of five targets—step up, step down, ramp up, ramp down, and a flat surface. They wore blurring goggles, constructed from Bangerter occlusion foils that reduced acuity to about 20/140 (mild blur) or 20/900 (severe blur). The effects of variation in lighting were milder than expected. Performance declined for the largest viewing distance, but exhibited a surprising reversal for nearer viewing. Of relevance to

pedestrian safety, the step up was more visible than the step down. In a second study (Bochsler, Legge, Kallie, & Gage, 2012), we investigated the impact of two additional factors expected to facilitate the recognition of steps and ramps during low-acuity viewing—surface texture and locomotion. We found that coarse texture on the ground plane reduced the visibility of ramps and steps, but walking enhanced visibility. In a third study, we conducted similar measurements on a group of 16 low-vision subjects with heterogeneous eye conditions and a wide range of acuities (Bochsler et al., In Press). As expected, overall performance decreased with acuity. More importantly, the qualitative pattern of results was very similar for the low-vision subjects and the normals with artificial acuity reduction. In particular, the low-vision subjects showed the same performance differences across target type (e.g., Step Up was much more visible than Step Down), benefited from locomotion, and exhibited poorer performance with texture. These parallel findings between normal and low vision encourage us to expect that similar parallels would exist for object recognition.

While we expect the qualitative features of our results to generalize to people with reduced acuity associated with low vision, we mention three caveats, as noted previously in Bochsler, et al. (2012). First, the Bangerter blur foils reduce acuity and contrast sensitivity for normally sighted subjects, but are not necessarily representative of any particular form of low vision. Second, we studied monocular viewing to simplify the optical arrangements for our subjects, and to simplify potential extension of the findings to low vision. Many people with mild or severe low vision have unequal vision status (acuities and other visual characteristics) of the two eyes, with performance determined

primarily by the better eye (see, e.g., Kabanarou & Rubin, 2006). Third, our subjects knew that one of the five targets (step up, step down, ramp up, ramp down, or flat) was present in each trial, and where to look for it, but low-vision pedestrians navigating unfamiliar locations in the real world do not always know when and where obstacles will appear in their path. Such uncertainties pose challenges for mobility not present in our study. We also note that our study focused on severely reduced acuity, but it is known that severe field restriction has a major impact on low-vision mobility (cf., Kuyk et al., 1996).

A surprising result in our study was the lack of an effect of the lighting arrangement. We think that this null result was due to the limited range of lighting conditions tested. As a counterpoint to our null finding, we conclude with an example illustrating how variations in natural lighting (sunlight) in a public space can dramatically affect the visibility of key features (Figure 2.6). Using high dynamic range (photometric) photography (Ward Larson, n.d.-a), Figure 2.6 shows a lobby at 4 different times during a single July morning. The Michelson contrast between the floor mat and the floor is 0.67, while the Michelson contrast along the edge of the shaft of sunlight is 0.88. The pattern of sun and shade produces high contrasts, which may interfere with visual detection of object features such as the floor mat. Of more importance to safety, the complex pattern of light and dark may make it difficult to identify the open stairwell. This example illustrates another limitation of our study: the potential interfering effects on object detection and identification from high-contrast cast shadows.

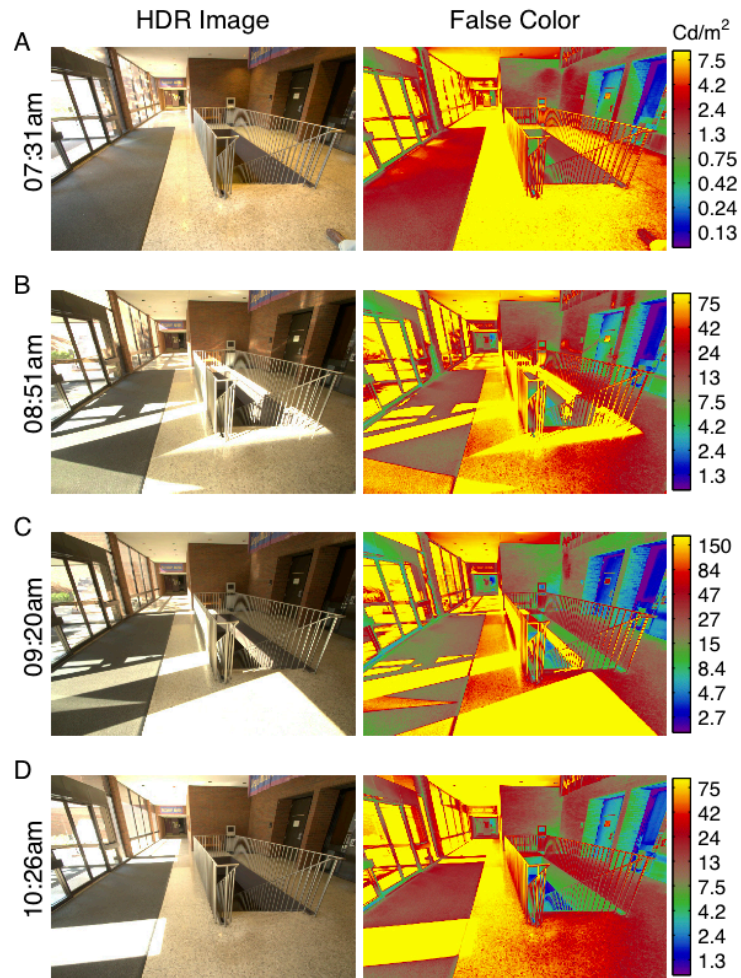


Figure 2.6. Lobby. This figure shows a lobby outside the experimental testing room. Images were recorded on a sunny morning in late July. Before direct sunlight reaches the lobby (Panel A), the dominating contrast features on the floor are along the carpet and around the stairwell, thereby indicating the presence of each. However, as the rising Sun passes through the lobby (Panels B-D), the dominating contrasts are along the borders of the shafts of sunlight, which can mask informative features of the space, such as the carpet on the left and the hazardous stairwell on the right. High dynamic range and false color images were created using Anywhere Software (Ward Larson, n.d.-a).

Do we expect our results to generalize to more complex and realistic environments and obstacles? In this study, we have articulated how to manipulate simple design

features that can greatly enhance object visibility. We expect that the design features that accounted for much of the variability in identification performance in the present study—internal and edge contrast, and object angular size—would carry over to more complex and realistic environments. Objects with simple contours are found throughout the architectural landscape. Placing them in an environment such as the one showcased in Figure 2.6 would be a good step toward examining object visibility in naturally complex environments.

In summary, the reflectance of objects and their immediate backgrounds, object shape, illumination and lighting arrangement work interactively to affect the visibility of objects, and should be considered collectively in evaluating visual accessibility.

Chapter 3:
Comparing Physical, Photographed, and Rendered Image Presentations
for Object Recognition with Optically and Digitally Blurred Stimuli

Introduction

We define visual accessibility as "the effectiveness with which vision can be used to travel safely through the space and to pursue the intended activities in the space (Legge et al., 2010)." A major requirement for visual accessibility is the ability to detect and identify objects that may be in a person's intended walking path. Reductions in acuity and contrast sensitivity are two major factors limiting object visibility, and are responsible for reductions in mobility performance (Goodrich & Ludt, 2003). Many authors, including Kuyk et al. (1996), also include field loss as a major factor, although the present psychophysical experiments do not address visual field loss.

It is difficult to predict when obstacles or other targets will be visible for people with different levels of acuity and contrast sensitivity. In a recent series of experiments, we constructed physical models of common environmental obstacles, including ramps and steps (Bochsler et al., 2012; Legge et al., 2010) and simple 3D objects (Chapter 2), using artificially induced blur to simulate visual impairment. We recently confirmed substantial similarity in the performance of sighted subjects with artificial acuity reduction and the performance of subjects with low vision (Bochsler et al., In Press).

Previous research with visually impaired pedestrians has linked illumination, contrast, and object type to mobility performance—and more fundamentally—to the visibility of physical objects (Kuyk et al., 1996). In our recent study of the visibility of physical 3D objects (Chapter 2), we examined the effects of lighting arrangement (overhead and artificial window), viewing distance, object shape (boxes and cylinders), height, and color (gray and white) with simulated visual impairment using blur goggles and normally

sighted subjects. However, construction of the physical environment was labor intensive. Human testing was difficult for the experimenter because of the need to move large physical test objects, and time consuming for subjects who must wait for placement of physical targets between trials. Because it was difficult to construct and manipulate physical environments, we aimed to extend the study of physical objects and environments to computer based testing. In the study presented in this chapter, our central research question asks whether the visual accessibility of physical spaces can be evaluated on a desktop computer. Specifically, can image based recreations of objects be displayed on a calibrated computer system to yield visibility performance (i.e., identification and detection) of objects equal to visibility performance measured in physical environments?

In this series of experiments, we recreated a physical testing environment using image capturing and image rendering technologies to evaluate the efficacy of the desktop display for testing object visibility. The layout was in a classroom environment with major scene props including: controlled illumination, a sidewalk elevated 16 inches, classroom furniture, and surrounding classroom features. Displays were created representing different subject viewpoints and different locations for the target objects. The general layout of the environment is shown in Figure 3.1. The Figure shows all the viewing locations and the two object locations that were studied, as well as major elements of the scene. The actual scene description is realistically created, as shown in rendered sample images in Figure 2.2.




Figure 3.1. Layout of Environment and Stimuli. This figure shows objects, illuminations, target locations, sample target shapes and orientations, and viewing positions for all experiments. Experiments 1-3 are shown along the top of the figure (i.e., along the length of the sidewalk). In Experiment 4 there were two object locations and two viewpoints: one along the sidewalk (shown in white and magenta), and the second along the center of the room (shown in solid magenta).

In the original physical objects experiment (Chapter 2), three notable findings associated with blurred viewing were: 1) objects with curvature (i.e. cylinders) were more visible than objects with flat surfaces facing the observer (i.e. boxes); 2) interactions between illumination, color, and shape played a vital role in object visibility; and 3) Michelson contrast measurements strongly predicted visibility performance.

In the study presented in this chapter, we asked two types of questions following the results of the original (Chapter 2) study. The first question was about whether we can take advantage of the computer desktop format of presentation—i.e., can we extend our

research to the desktop environment? Secondly, if we were successful, we wanted to ask more questions about the nature of our objects and lighting. Specifically, we wanted to know whether cylinders would be more visible than boxes under a variety of viewing positions and orientations. The first question is addressed here in Experiments 1 through 3. The second question is addressed in Experiment 4.

Producing computer images that are photometrically similar to physical environments is technically challenging. Gregory J. Ward has developed a collection of specialized imaging, modeling, and rendering software packages and procedures for the purpose of recreating realistic scenes with photometric and visual accuracy. In the study presented here, a computer model and rendered images using the Radiance software suite (Ward Larson, n.d.-b), with the methods described in the Radiance textbook (Ward Larson & Shakespeare, 2011), were created. Through accurate numerical modeling of geometric and reflective material surface properties, and accurate reproduction of photon flux of light sources, the Radiance software suite renders the physical lighting of scenes with photometric accuracy. The Radiance rendering engine implements a number of deterministic and probabilistic light modeling algorithms, while optimizing computational speed with physically accurate results, to create images that look realistic and are numerically equivalent to real-world photometry measurements. Most other 3D rendering software packages do not preserve veridical representations of light energy across a scene.

In addition to renderings, we took photographic images of the original stimuli that were photometrically accurate, using software and methods developed by the author of

the Radiance software suite (Ward Larson, n.d.-a, n.d.-c). Photographs are different in that they are images of a real environment, collected by an imaging apparatus.

Renderings come from computer models of environments, and are not collected by cameras, but rather are simulated by computer models and algorithms.

A desktop environment avoids many of the problems and pitfalls of human subject testing in a physical environment. Collecting photometrically accurate photographs of existing scenes requires modest computer and camera abilities, and several minutes of interactive desktop computer use per image. Creating a Radiance model requires programming skills and a capable personal computer. Otherwise, with Radiance it is less expensive, safer, and less time consuming to produce and test images that are photometrically similar to real environments. Rendered images for this experiment took on average 1/2 hour per image on a modern laptop computer. Further advantages include the ability to model and render existing environments, add novel objects, positions, orientations, and viewpoints to existing environments, or model and render completely novel environments that do not presently exist in physical reality. These advantages are essential features of the Radiance software suite—Radiance provides the ability to prototype and examine the photometric physical qualities of environments, rapidly, inexpensively, and independent of physical construction. Finally, computer modeling allows precise control over the 3-dimensional states of objects and viewpoints. In the physical environment, the positions of objects and observers are variable. However, in 3D modeling, such variability can be precisely controlled.

General Method

In the present study, we created a virtual model of the physical setup described in Chapter 2, and tested whether desktop computer display experiments could replicate the performance found in the original physical environment study. In an effort to improve the performance of the virtual model, we developed and tested three digital models of the original Bangerter occlusion foil blur goggles used in the physical experiment.

In Experiment 1, we tested whether a Gaussian low-pass filter would provide a satisfactory model for the blur goggles in order to replicate the pattern of results obtained in the physical experiment. In Experiment 2, we photographically obtained the point spread function (PSF) of our Bangerter blur goggles, and tested whether it could replicate our physical objects results. In Experiment 3, we tested whether addition of a stage of contrast reduction to the Experiment 2 model would improve the match between the desktop results and the original physical experiment. Finally, in Experiment 4, we extended our findings of the digital Bangerter blur model to novel environmental conditions. We wished to expand the findings from our original physical objects study (Chapter 2) and to demonstrate the utility of the digital imaging and rendering approach.

Scene and Object Modeling and Photorealistic Image Formation

In the original physical experiment (Chapter 2), there were 78 unique stimulus conditions, including: 2 illumination conditions (Overhead & Window), 3 viewing distances (10, 17, & 24ft), 2 object shapes (Boxes & Cylinders), 3 object heights (2, 4, & 6ft), and 2 object colors (Gray & White)—yielding 72 objects; plus 6 'no object' stimuli (one for each of the 2 illumination by 3 distance conditions). The layout of the scene,

objects, and viewing positions are shown in Figure 3.1, and the objects from one viewing distance are shown in Figure 2.2.

Using Radiance software, a model of the physical space and objects in the original physical testing environment was created. Radiance is a 3D spatial modeling, image rendering, and image analysis toolbox capable of accurate simulation of environmental geometry and lighting (Ward Larson & Shakespeare, 2011, p. xviii). The Radiance software suite is unique among modeling and rendering software in that images can look realistic and can accurately recreate intended luminance values in the scene. Our Radiance model of the original physical environment was created by Rob Shakespeare, and includes details about lighting, geometry, and reflectance properties of all major components of the original physical testing environment. Physical properties of the original physical scene were collected on sight, and used for the accurate creation of the 3D Radiance model. Physical properties include: 3D scene geometry, object shapes and positions, reflective properties of surfaces, and physical lighting characteristics.

Using Shakespeare's Radiance model, the 78 unique stimulus conditions were modeled and rendered according to the methods described in Chapter 4. The photometric accuracy of the Radiance model was confirmed by the calibration and validation methods from renderings, as described in Chapter 4.

In addition to the HDR rendered images, HDR photographs of each of the 78 stimulus conditions were created according to the methods described in Chapter 4. Like the renderings, the photographs were also made to be photometrically accurate and were photometrically validated according to the methods described in Chapter 4.

Rendered and photographed HDR images were filtered and converted to JPEG images and displayed according to the methods described in Chapter 4. A flow diagram showing the process of going from HDR image to displayed image is shown in Chapter 4, Figure 4.2. HDR image formation, calibration issues and techniques, spatial image filtering techniques, and non-spatial image filtering techniques, along with sample code developed for this series of experiments, are documented and explained in Chapter 4.

Apparatus and Experimental Protocol

Experiments were written in Matlab (*MATLAB R2010a*, 2010), using Psychophysics Toolbox extensions (Brainard, D. H., 1997; Kleiner, M., Brainard, D. H., & Pelli, D. G., 2007; Pelli, D. G., 1997). Images were shown on a 27-inch computer display that was luminance calibrated according to the method described in Chapter 4. The room was darkened to reduce ambient light reflections from interfering with the display.

Angular dimensions, aspect ratios, and pixel counts for rendered and photographed images were controlled in order to keep the visual angles on the desktop the same as the visual angles seen in the original physical objects experiment. The 27-inch (2560 by 1440 pixel) display was viewed from a distance of 16 inches, creating a viewing angle of 73 by 45 degrees. A chin rest was used to keep the subjects at the desired viewing distance. Both rendered and photographed images were centered on the display area with gray surrounding pixels. Rendered images were 2112 by 1314 pixels, and the viewing angle of the rendered images was 66 by 44 degrees. The camera images were 2144 by 1424 pixels, and the viewing angle of the camera with 18-mm lens was approximately 64 by 42

degrees. The angular field of view of the original physical blur goggles was approximately 48 degrees, and was circular.

Similar to each illumination and viewing distance block in the physical objects experiment, for Experiments 1-3, six blocks of trials were formed according to 2 illumination types (overhead and window) by 3 viewing distances (10, 17, & 24 ft.). Experiment 4 had a slightly different block design, where 2 illumination types and 2 viewing positions created 4 blocks. The order of blocks and individual trials within blocks were randomized for each subject. Blocks included trials of a complete set of object stimuli (including boxes, cylinders, and catch trials without any object present). Details of the number of trials are covered in individual experiment sections below.

At the beginning of each block, subjects were first presented with a series of 1-second previews of the unfiltered, display calibrated stimulus images (image preparation is described in Chapter 4). The preview image (i.e., the non spatially filtered image) sequence (for Experiments 1-3) was presented at the beginning of each block as follows: no object, gray boxes from short to tall, no object, gray cylinders, no object, white boxes, no object, and finally white cylinders. In other words, they saw all 12 objects, once each, in increasing height, plus the no object stimulus four times interleaved. The preview images were similar to the sequences of images shown in Figure 2.2. Experiment 4 was similar to the first 3 experiments, except subjects saw boxes of different orientations, along with the cylinders and the no-object conditions. After the preview cycle ended, trials began as noted in the next paragraph. Breaks were allowed and encouraged between blocks.

A single trial consisted of a 4-second viewing of a single image on the 27-inch display, followed by a response graphical user interface (GUI) displayed on the 13-inch laptop display (the GUI for Experiments 1-3 is shown in Figure 3.2). After each stimulus exposure, the 27-inch screen returned to gray, and the GUI response box was presented for the subjects to respond by mouse clicks. After the subject responded, the next 4-second trial was displayed on the 27-inch screen until all trials in a block were finished.

For Experiments 1-3, GUI responses for each trial were for Object Present ('Yes' or 'No'), Confidence (1-5, with '5' being certain). If the subject responded 'No' to the first question, then the next trial began after a confidence rating response. Otherwise, the next answer required was for Shape ('Box' or 'Cylinder'), and finally Height ('Tall', 'Medium', or 'Short'), in which case the next trial began after the height response. Figure 3.2 shows the response GUI with radio buttons for mouse-clickable answers. Experiment 4 was simplified, where the subject had only a single response to choose from, including: 'Box', 'Cylinder', or 'None'.

Participants

Participants included 35 college undergraduates who were recruited on campus. Ages were from 20 to 26, and 12 were male. Participants had normal or corrected to normal visual acuity and had no known eye health or cognitive problems. Eight participants earned research experience points (REP) for class credit, and the remainder received \$20-\$25 payment by university issued check. Experiments lasted 2-2.5 hours. Participants gave informed consent, the experimental protocol was approved by the University of Minnesota's IRB, and it followed the tenets of the Declaration of Helsinki.

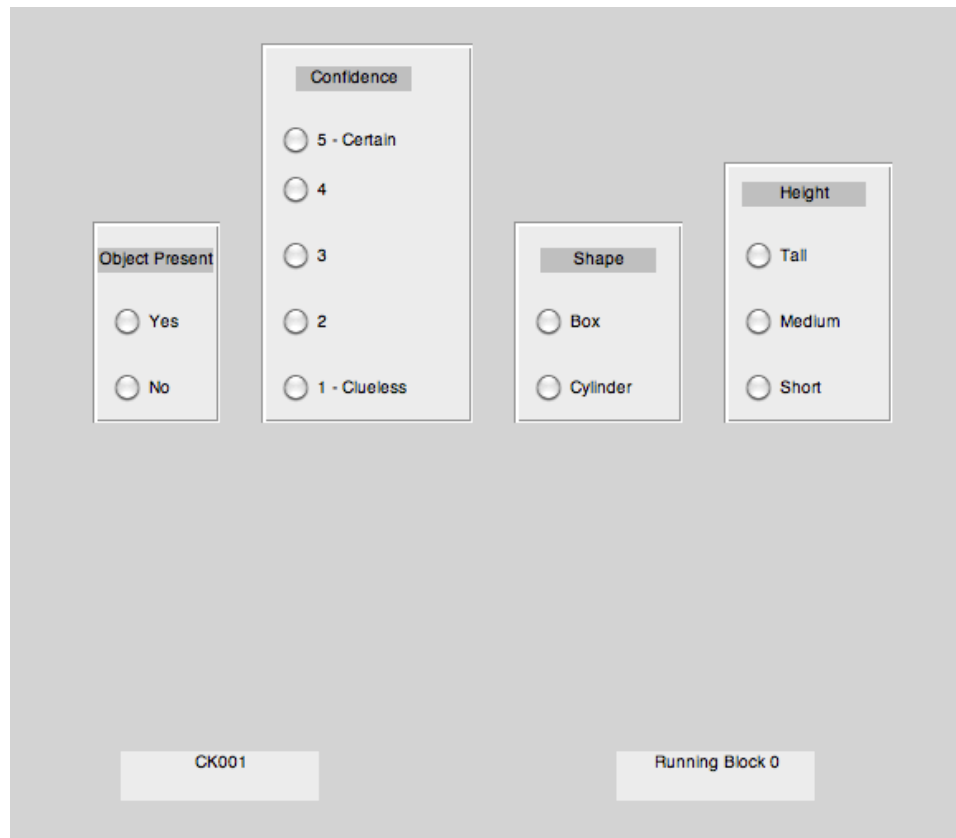


Figure 3.2. Response GUI (Experiments 1-3). This figure shows the mouse-clickable interface and response choices subjects made following each stimulus presentation.

Analysis

Identification is defined as reporting the correct shape and height for an object present. Detection is defined as reporting 'Yes' to 'Object Present' when an object was in fact present in the stimulus. Since results for both visibility measurements were very similar, much of the analysis here (e.g., ANOVAs and interaction plots) is with identification; some tables show both identification and detection statistics.

Main effects, and selected interaction effects, for Experiments 1-3, were compared to the effects found in the original physical objects experiment (Chapter 2). Original

physical objects data are shown numerically in Table 2.1 and selected interactions are shown graphically alongside the interaction plots for Experiments 1-3. ANOVAs for each of the desktop models (differing in the formulation of the digital blur filter) are examined and compared to the ANOVA of the original physical objects experiment (Chapter 2, Table 2.2).

Finally, coefficients of determination and regression coefficients for correspondence and predictive power of each desktop model were produced and compared. Data points for the coefficient analysis included object identifications and correct rejections. Table 3.1 (in the Results section) shows the coefficients, and gives quantitative results of how each model predicts the physical objects experiment. Correlations were computed between performance levels (proportion correct) for each of the 78 viewing conditions in the original physical experiment and the corresponding performance levels in desktop Experiments 1-3. Coefficients of determination (squared correlation coefficients) and regression coefficients were computed to evaluate the predictive power the models tested in Experiments 1-3. For a perfect model, the coefficient of determination, would be $R^2 = 1$, and the regression coefficients would be, intercept $\beta_0 = 0$, and slope $\beta_1 = 1$. This would indicate three fundamental properties of a perfect model, as follows: 1) $R^2 = 1$ would indicate that the model perfectly predicts real world performance; 2) $\beta_0 = 0$ would indicate that when performance is at 0% in the model, it is also at 0% in the real world; and 3) $\beta_1 = 1$ would indicate that the gains in percentage performance in the model are identical to the gains in percentage performance in the real world. (Note that a similar R^2 and β analysis was used to characterize the relationship between pixel luminance and

physical luminance in Chapter 4. The reader should be sure not to confuse the modeling coefficients for psychophysics here to the modeling coefficients in the photometric calibration and proof described in Chapter 4.)

In addition to the correlation and regression analysis as described above, an analysis of eigenvectors is reported (Results section, Table 3.1). The eigenvector analysis is similar to the regression analysis, however the assumptions for the relationship between variables are slightly different. In the case of the regression analysis, it is assumed that the dependent variable is performance in the physical environment, while the predictor variable is performance on the desktop. However, the eigenvector analysis is based upon the assumption that both physical performance and desktop performance are dependent upon a third, independent variable—*visibility*. Since both performance measurements are dependent upon *visibility*, it is appropriate to compare the two dependent variables with eigenvectors, which produce unbiased estimates of the relationship, or the covariance between the desktop model and the physical experiment. For the sake of easy comparison, a function was written that transforms the eigenvectors to coefficients of a regression line that pass through the eigenvectors on the Cartesian plane. The eigenvector coefficients are labeled with the Greek letter ψ in Table 3.1.

The following paragraphs describe the method and rationale for including the eigenvector analysis in more detail.

The angles and positions of the eigenvectors were translated into slope-intercept form—the values of which are represented by coefficients, ψ_1 and ψ_0 , respectively. The ψ coefficients are unbiased estimators of the correspondence between performance for

physical objects versus performance for desktop objects. The regression analysis, on average, produces slopes that are biased downward and intercepts that are biased upward, which is due to the fact that performance on physical objects were regressed upon performances on desktop objects. The underlying assumption of regression is that one variable is dependent upon the other—which is not truly accurate in our case. These biases are small, but can be seen when viewed alongside the data and compared to their corresponding eigenvector lines, as shown in Figure 3.10.

Alternatively, the eigenvector analysis assumes that both measures are dependent, weighting both performance measures equally, thereby finding the vector that best represents the relationship between those the two dependent measures (i.e., visibility in the physical environment and visibility on the desktop environment). If only one performance measure (i.e., physical) was dependent on the other (i.e., desktop), then the regression analysis would have been the most appropriate model for comparison. However, since both measures are dependent upon a third (unknown or unmeasured) variable—namely visibility—the regression analysis is biased as noted in the preceding paragraphs. Both methods are presented in this dissertation in order to satisfy the greatest number of readers.

It is notable that although the regression is slightly biased, the coefficients between the regression and the eigenvector models are very highly correlated.

In summary, the most notable difference between the eigenvector and regression analyses is that the coefficients derived from the eigenvector analyses are unbiased between two dependent measures of visibility. If one wishes to compare gains or

thresholds in visibility performances between multiple experiments—each having a related dependent measure of visibility—then the eigenvector ψ coefficients are unbiased, as compared to the regression β coefficients.

Finally, ANOVAs of arcsine transformed identification proportions correct were produced to examine main and interaction effects relevant to the study (ANOVAs were also examined for detection performance, but they were similar to identification performance, so they were left out of this paper). Two-way (distance by height) and 3-way (color by shape by illumination) interaction plots were produced (Experiments 1, 2, & 3) for comparison to the physical experiment results. Selected means tables, hypothesis tests, and figures were produced for examining the important differences or similarities of each of the 4 desktop experiments, relative to the original physical objects experiment.

Experiment 1—Effects of Gaussian Blur

In Experiment 1, we tested whether visibility (i.e. detection and identification) of Gaussian blurred images of objects resemble the performance of subjects viewing corresponding physical objects. We wanted to test whether a simple Gaussian blur model could replicate the effects of visibility found with the physical Bangerter blur goggles. If this were the case, then the modeling effort for simulating Bangerter blur foils would be simple to describe and implement.

Participants

Eighteen subjects were recruited according to the description in General Methods above.

Method

The 78 rendered images were used for this experiment, and were prepared according to the method described in Chapter 4. No photographs were used in this experiment. The function used for Gaussian filtering is shown in Chapter 4, Table 4.3. Blocks and trials were presented according the description in the General Methods section. Each image was filtered to 5 levels of Gaussian blur (samples are shown in Figure 3.3); all levels of blur were shown in each block. The blur levels in SD pixel units were $\sigma = 30, 45, 60, 75,$ and 90 pixels. Sigma pixel values converted to degrees were: $\sigma_{\text{deg}} = 0.971, 1.46, 1.94, 2.43,$ and 2.91 degrees, and the corresponding cut-off frequencies were: $f_c = 1.03, 0.69, 0.51, 0.41, 0.34$ cyc/deg, respectively, corresponding to Snellen acuities ranging from 20/600 to 20/1800, bracketing the nominal acuity of the Bangerter blur goggles which we measured to be around 20/900 (Legge et al., 2010). This range of blur levels was examined in order to find a level of blur whose pattern of performance replicated the pattern of performance found in the physical objects experiment.

In Experiment 1, there were 72 object stimuli (2 illumination x 3 distance x 2 color x 3 height x 2 shape) and 6 non-object stimuli (2 illumination x 3 distance). Subjects were shown each object once, and each non-object 3 times, making 90 unfiltered image presentations (72 objects + 6 non-objects x 3 = 90). With 90 unfiltered images and 5 levels of blur, there were 450 trials per subject. There were 18 subjects: 450 trials per subject x 18 Subjects = 8100 total trials. Trials were blocked and randomized as noted in the General Methods section, with the illumination by distance combination as a block, yielding 6 total blocks.



Figure 3.3. Gaussian Blur Stimulus Samples (Experiment 1). This figure shows four sample rendered images and Gaussian blurred images. The first row shows 4 sample rendered images without blur. Rows 2-6 show 5 levels of Gaussian blur. Objects shown are the gray box with overhead lighting (columns 1 and 3), and the white cylinder with window lighting (columns 2 and 4) at 10 feet (columns 1-2) and 24 feet (columns 3-4).

Results

Correlations were computed to compare performance in this experiment with corresponding conditions in the original physical objects experiment. For all levels of blur and for both detection and identification, the coefficients of determination (R^2)

ranged from .11 to .26 (Table 3.1 A & D). The Gaussian blur that was most predictive of physical world performance was $\sigma = 60$ pixels ($\sigma_{deg} = 1.94$, or $f_c = 0.51$). This means that 26% of the variance in the physical world experiment can be accounted for by the Gaussian $\sigma_{deg} = 1.94$ model. The regression coefficients for that model included intercept $\beta_0 = .04$, and slope $\beta_1 = .54$. Regression coefficients are also listed in Table 3.1A & D.

| Identification (A, B, & C) | | Model | | | Coef of Det | Regression Coefs | | | Eigenvector Coefs | | |
|----------------------------|------------------|-------------------|---------------|----------------------|--------------|------------------|-------------|---------------|-------------------|-------------|--------------|
| | | σ -pix | σ -deg | f_c | R^2 | β_0 | β_1 | $\Sigma\beta$ | ψ_0 | ψ_1 | $\Sigma\psi$ |
| Gaussian Blur (A) | 30 | 0.971 | 1.03 | 0.13 | -0.14 | 0.6 | 0.46 | -2.36 | 3.3 | 0.94 | |
| | 45 | 1.46 | 0.69 | 0.25 | -0.1 | 0.62 | 0.52 | -0.75 | 1.5 | 0.75 | |
| | 60 | 1.94 | 0.51 | (Fig3.4) 0.26 | 0.04 | 0.54 | 0.57 | -0.29 | 1.09 | 0.8 | |
| | 75 | 2.43 | 0.41 | 0.17 | 0.15 | 0.44 | 0.58 | -0.19 | 1.13 | 0.95 | |
| | 90 | 2.91 | 0.34 | 0.22 | 0.18 | 0.51 | 0.69 | -0.05 | 1.18 | 1.14 | |
| Banger-ter (B) | Contrast Percent | | | | | | | | | | |
| | 100 | | | (Fig3.6) 0.51 | -0.02 | 0.62 | 0.59 | -0.15 | 0.82 | 0.67 | |
| Contrast Reduction (C) | Contrast Percent | | | | | | | | | | |
| | 5 | | | 0.08 | 0.3 | 0.33 | 0.62 | 0.06 | 1.61 | 1.66 | |
| | 10 | | | 0.17 | 0.25 | 0.44 | 0.69 | 0.07 | 1.19 | 1.27 | |
| | 20 | | | 0.28 | 0.21 | 0.52 | 0.73 | 0.08 | 0.96 | 1.05 | |
| | 40 | | | 0.52 | 0.11 | 0.62 | 0.73 | 0.03 | 0.81 | 0.84 | |
| | 60 | | | 0.52 | 0.05 | 0.62 | 0.67 | -0.05 | 0.81 | 0.76 | |
| | 80 | | | (Fig3.8) 0.58 | 0 | 0.68 | 0.68 | -0.09 | 0.85 | 0.76 | |
| | 100 | | | 0.52 | -0.04 | 0.69 | 0.66 | -0.18 | 0.95 | 0.77 | |
| Detection (D, E, & F) | | Gaussian Blur (D) | | | R^2 | β_0 | β_1 | $\Sigma\beta$ | ψ_0 | ψ_1 | $\Sigma\psi$ |
| | | σ -pix | σ -deg | f_c | | | | | | | |
| Gaussian Blur (D) | | 30 | 0.971 | 1.03 | 0.11 | -0.29 | 0.97 | 0.68 | -7.1 | 7.98 | 0.88 |
| | | 45 | 1.46 | 0.69 | 0.17 | -0.01 | 0.71 | 0.7 | -2.23 | 3.1 | 0.86 |
| | | 60 | 1.94 | 0.51 | 0.24 | 0.09 | 0.65 | 0.74 | -0.86 | 1.76 | 0.9 |
| | | 75 | 2.43 | 0.41 | 0.24 | 0.2 | 0.59 | 0.79 | -0.47 | 1.46 | 0.99 |
| | | 90 | 2.91 | 0.34 | 0.19 | 0.32 | 0.52 | 0.84 | -0.28 | 1.47 | 1.19 |
| Banger-ter (E) | Contrast Percent | | | | | | | | | | |
| | 100 | | | 0.45 | 0.19 | 0.61 | 0.8 | 0 | 0.86 | 0.86 | |
| Contrast Reduction (F) | Contrast Percent | | | | | | | | | | |
| | 5 | | | 0.05 | 0.55 | 0.24 | 0.79 | 0.1 | 1.36 | 1.46 | |
| | 10 | | | 0.07 | 0.53 | 0.27 | 0.8 | 0.21 | 0.99 | 1.21 | |
| | 20 | | | 0.1 | 0.49 | 0.31 | 0.81 | 0.17 | 0.95 | 1.12 | |
| | 40 | | | 0.39 | 0.28 | 0.59 | 0.87 | 0.07 | 0.91 | 0.98 | |
| | 60 | | | 0.45 | 0.16 | 0.68 | 0.84 | -0.07 | 1 | 0.93 | |
| | 80 | | | 0.44 | 0.11 | 0.7 | 0.82 | -0.18 | 1.08 | 0.91 | |
| | 100 | | | 0.41 | -0.03 | 0.83 | 0.8 | -0.58 | 1.5 | 0.92 | |

Table 3.1 *Coefficients of Determination, Regression, and Eigenvector For Identification and Detection*

(Experiments 1 - 3). This table shows how well each desktop model predicted performance in the physical

experiment. All 78 stimuli were used for computing each statistic. Statistics in bold have corresponding interaction plots with corresponding figure numbers in parentheses. Coefficients of determination estimate the precision between the desktop models and physical performance. Regression and eigenvector coefficients determine accuracies for models predicting physical performance, where eigenvector coefficients are unbiased estimators of the accuracy. Regression terms estimate how desktop performance predicts physical performance, without worrying about whether the predictors are themselves related or redundant. Eigenvectors look for correlations between dimensions, and find independent axes that align with the principal directions of variation. The eigenvector and correlation analyses together provide an unbiased view of the relationship between desktop and physical visibility. Regression coefficients (β_0 & β_1) and eigenvector coefficients (ψ_0 & ψ_1) describe a linear relationship: y (physical) versus of x (model) in terms of slope and intercept. Summation values ($\Sigma\beta$ & $\Sigma\psi$) give the predicted physical performances at perfect performance in the physical world for each model. Note that eigenvector coefficients have been transformed from (unreported) eigenvector values, into (reported) slope-intercept form as ψ_1 (slope) and ψ_0 (intercept).

The marginal means for detection and identification for varying levels of Gaussian blur are shown in Table 3.2. These data demonstrate the trend in decreasing performance as a function of increasing blur.

Table 3.3 shows the arcsine transformed identification ANOVA for the $\sigma_{deg} = 1.94$ model. Comparing this ANOVA to the ANOVA from the physical objects experiment (shown Table 2.2), there are similarities across all main effects and interaction effects, with one important exception: an effect for illumination (overhead vs. window) was found in this experiment, but it was not found in the physical objects experiment.

| σ_{pix} | σ_{deg} | f_c | Identification | $\text{SEM}_{\text{ident}}$ | Detection | $\text{SEM}_{\text{detect}}$ | N-trials |
|-----------------------|-----------------------|-------|----------------|-----------------------------|-----------|------------------------------|----------|
| 30 | 0.971 | 1.03 | 0.81 | 0.0109 | 0.97 | 0.0047 | 1296 |
| 45 | 1.46 | 0.69 | 0.72 | 0.0125 | 0.93 | 0.0072 | 1296 |
| 60 | 1.94 | 0.51 | 0.57 | 0.0138 | 0.85 | 0.0098 | 1296 |
| 75 | 2.43 | 0.41 | 0.45 | 0.0138 | 0.76 | 0.0118 | 1296 |
| 90 | 2.91 | 0.34 | 0.30 | 0.0128 | 0.62 | 0.0135 | 1296 |

Table 3.2 *Gaussian Blur Means (Experiment 1)*. Selected means show proportions correct for 5 blur levels for detection and identification performance.

| Source | Sum Sq. | d.f. | Mean Sq. | F | Prob>F |
|--------------------|----------|------|----------|--------|--------|
| Blur | 3.0796 | 4 | 0.7699 | 318.32 | 0 |
| Subject | 0.6319 | 17 | 0.03717 | 15.37 | 0 |
| Illumination | 0.39266 | 1 | 0.39266 | 162.35 | 0 |
| Distance | 2.1783 | 2 | 1.08915 | 450.31 | 0 |
| Color | 0.37788 | 1 | 0.37788 | 156.24 | 0 |
| Height | 0.09391 | 2 | 0.04696 | 19.41 | 0 |
| Shape | 0.0293 | 1 | 0.0293 | 12.11 | 0.0005 |
| Illumination*Color | 0.09863 | 1 | 0.09863 | 40.78 | 0 |
| Illumination*Shape | 0.08598 | 1 | 0.08598 | 35.55 | 0 |
| Distance*Height | 0.00547 | 4 | 0.00137 | 0.57 | 0.6876 |
| Color*Shape | 0.19279 | 1 | 0.19279 | 79.71 | 0 |
| Illum*Color*Shape | 0.03573 | 1 | 0.03573 | 14.77 | 0.0001 |
| Error | 15.58333 | 6443 | 0.00242 | | |
| Total | 22.78549 | 6479 | | | |

Table 3.3 *Gaussian Blur Identification ANOVA (Experiment 1)*. Data were arcsine transformed.

Figure 3.4 shows the identification interaction plots for the $\sigma_{\text{deg}} = 1.94$ model, and compares the same interactions from the physical objects experiment. Three major differences can be noted by comparing these two sets of data. First, the mean identification performance for the Gaussian $\sigma_{\text{deg}} = 1.94$ model is higher than the physical objects mean performance (Gaussian Identification = .57, Physical Identification = .33). There is a Gaussian model that has mean performance closer to the mean performance of the physical objects experiment ($\sigma_{\text{deg}} = 2.91$ Identification = .30), but the pattern of results is not as similar, as is evident by the R^2 analysis (for $\sigma_{\text{deg}} = 1.94$, $R^2 = .24$, and for

$\sigma_{\text{deg}} = 2.91$, $R^2 = .19$), indicating that the Gaussian blur is not a good model for the Bangerter blur goggles. Also, boxes were less visible with windows than overhead illumination in the physical objects experiment, but the interaction was opposite in the $\sigma_{\text{deg}} = 1.94$ model, as shown in Figure 3.4.

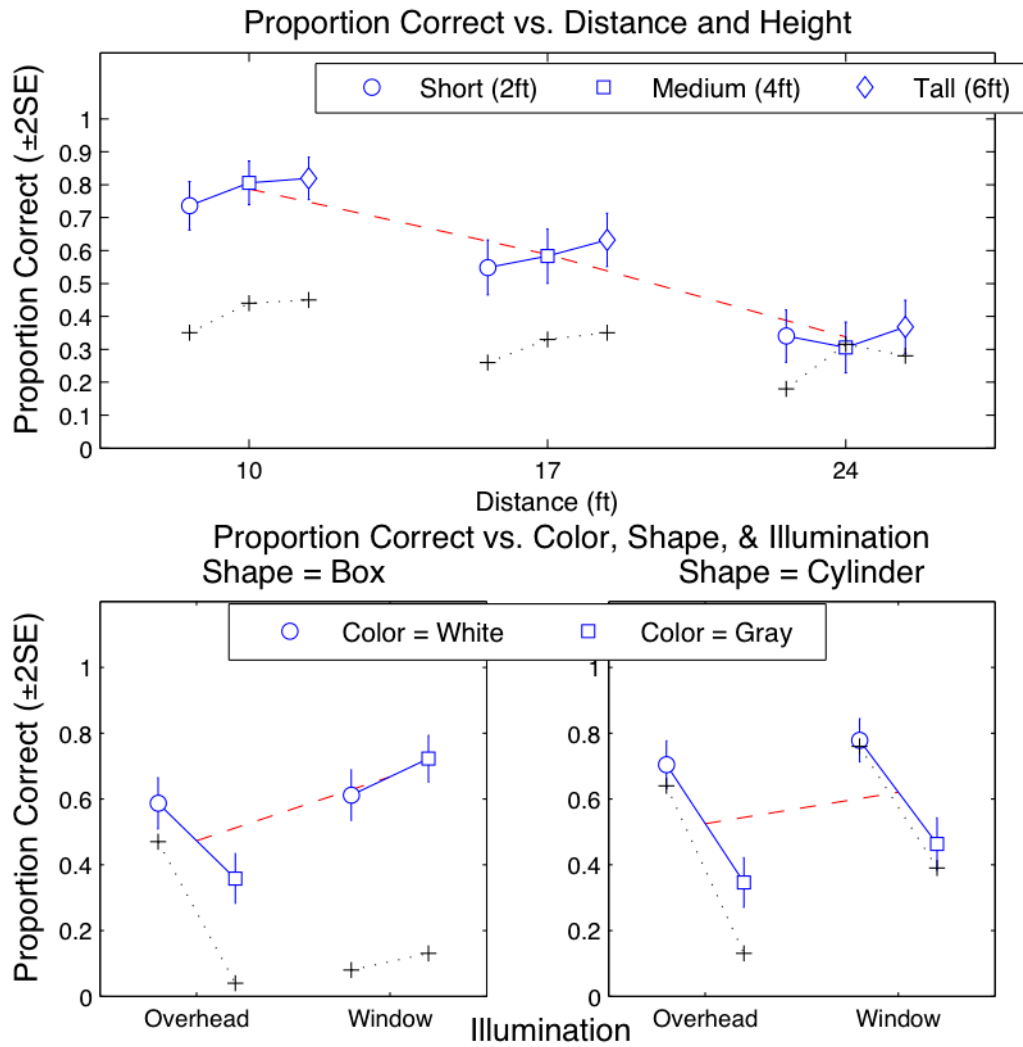


Figure 3.4. Gaussian Blur Identification Interaction Plots (Experiment 1). This figure shows the identification performance for the highest correlating Gaussian blur ($\sigma\text{-deg} = 1.94$), i.e., the best Gaussian

match for identification. The coefficient of determination between this model and the physical experiment is: $R^2 = .26$. Black crosses show original physical data.

In summary, the pattern of results in visibility between Experiment 1 and physical objects (Chapter 2) is different—as noted by the coefficient and graphical analyses. The ANOVAs and interaction plots show the details of where major differences in performance occur—namely with boxes and window illumination.

Discussion

From the coefficient of determination analysis alone, it is evident that the Gaussian blur model is insufficient to simulate a physical Bangerter occlusion foil goggles. The similarities in subject performance between the Gaussian blur model and the physical experiment are however promising. The coefficients of determination were well above chance, and the regression coefficients, although far from ideal, show a desirable pattern of results, which appear to be in the direction of the physical objects results. These similarities are further confirmed by the pattern of results shown in the interaction plots. The data in those plots show a pattern that has identifiable similarities to the physical objects data plots.

The differences, however, are too great to ignore, and compel us to consider improved modeling for the images, and for the Bangerter occlusion foil goggles. We considered two key improvements over the Gaussian blur model. First, the Gaussian blur was not a close enough approximation to the Bangerter occlusion foils' filtering characteristics. Second, the Gaussian blur filter was applied to the JPEG images and not

the HDR images, which could have caused luminance and contrast-related artifacts within the displayed blurred images.

Additionally, although there were photometrically accurate renderings of the physical space, we had no way of knowing with certainty whether artifacts of the renderings were causing some unknown difference between the HDR images and the physical environment. We sought to handle these issues with a new filter model, improved image processing methods, and a direct comparison of performance with rendered and photographed images, in Experiment 2.

Experiment 2—Effects of Digital Bangerter Blur

It is unclear a priori how refined a model of blur would be necessary to accurately simulate the visibility performance associated with the Bangerter occlusion foils. As a first approximation, we tried Gaussian blur as a generic low-pass filter. We varied the size of the convolution kernel with the hope of finding a level of blur that would yield performance roughly equivalent to the performance in the physical experiment with the Bangerter goggles. Experiment 1 showed that the Gaussian filter does not provide an adequate model for Bangerter blur filters.

The next step was to derive a more accurate filtering function for the Bangerter blur goggles. At least one study has characterized the modulation transfer characteristics of several commercially available Bangerter occlusion foils (Pérez, Archer, & Artal, 2010). The models that were presented were graphical representations of the modulation transfer function (MTF). Based on our research, graphical MTF data alone was insufficient for accurate characterization of the Bangerter occlusion foils. To solve this, we characterized

the Bangerter blur goggles by measuring the point spread function (PSF) of the blur goggles using the methods described in Chapter 4. The Fourier transform of the PSF is the optical transfer function (OTF), which includes the MTF, plus the phase transfer function (PTF). In our experiment, it became evident that the more complete camera based OTF was capable of sufficiently accurate characterization of the physical blur goggles for visibility testing.

In Experiment 2, we tested whether a digitally imaged Bangerter filter can produce visibility effects that match performance in the physical world. Additionally, in this experiment, two different types of images were examined, including renderings and photographs, in order to test whether there are differences between rendered and photographed images, in terms of visibility performance. HDR photographs were collected according to the methods described in Chapter 4.

Participants

Seven participants were recruited according to the description in General Methods.

Method

We photometrically calibrated a Nikon D90 DSLR camera, as noted by the method described in Chapter 4. This effectively created a megapixel photometer that could be used for imaging scenes and for creating a digital model of the Bangerter blur goggles. We used the camera to obtain HDR photographs of the original 78 stimulus conditions, according to the method described in Chapter 4. We then built an apparatus to form a digital image of the point-spread function (PSF) of the Bangerter blur goggles, according to the method described in Chapter 4. The PSF was used to derive the optical transfer

function (OTF) for the Bangerter blur goggles and derive a digital filter. A tone-mapped image of the PSF is shown in Chapter 4, Figure 4.4. The OTF is the Fourier transform of the PSF, which occurs in the filtering process shown in Chapter 4, Table 4.4. The spatial imaging filtering process is described in Chapter 4, Section 4.3.2.

We created unfiltered, color balanced (i.e. white balanced), and luminance calibrated images of both Radiance renderings and HDR photographs, and tested and verified their photometric accuracies according to the methods described in Chapter 4. Filtered stimulus images were then formed using the color and luminance accurate settings made in the recalibration phase, and then filtered—all according to the methods described in Chapter 4.

In summary, we made both Photographed and Rendered images and filtered them with a digitally obtained Bangerter filter (see middle panel in Chapter 4, Figure 4.2 for a flow diagram of the process). We tested subjects in the same way as described in the general methods and Experiment 1 methods. The two goals were to test the merits of a measured transfer function for the Bangerter occlusion foil goggles, and to test the empirical equivalence of the rendered and photographed images. Samples of the unfiltered rendered and photographed images, as well as digitally filtered Bangerter images, are shown in Figure 3.5. There are very subtle differences between the rendered and photographed images, but the differences are nearly perceptually indistinguishable. Moreover, the two image types are likely to be indistinguishable for visibility testing of the filtered images—a hypothesis that is being tested here.




Figure 3.5. Digital Bangerter Blur Stimulus Samples (Experiment 2). This figure shows rendered (rows 1 & 3) and photographed (rows 2 & 4) image samples for unfiltered and Bangerter filtered objects.

In Experiment 2, there were the same 78 unfiltered rendered images (including 72 objects plus 6 no-object stimuli) as described in Experiment 1, plus 78 corresponding (nearly identical) photographed images (see Figure 3.5 for comparison between rendered and photographed images). There were 2 repeated measures for each subject (plus the extra repeats for the no-object stimuli, as noted in General Methods), totaling 360 trials per subject. There were 7 subjects, summing to 2520 trials. There were 12 blocks (6 blocks x 2 repeated measures) for this experiment. Blocks and trials were randomized as noted in general methods. Rendered and photographed images were randomly sequenced within each block.

Results

The coefficients of determination comparing performance with the digital Bangerter model to physical world performance for object detection and identification were $R^2 = .45$ and $R^2 = .51$, respectively, which is an improvement over the Gaussian blur model.

The marginal means for the main effects are listed in Table 3.4 (which can be compared to physical objects means in Table 2.1). The table shows that windows provided about 10% more visibility than overhead. Detection did not change for heights, but identification did improve for taller objects. Cylinders were more visible than boxes for both detection and identification, as expected.

In addition to the marginal means table, an ANOVA (Table 3.5) was conducted, which can be compared to the ANOVA of the physical objects experiment (Table 2.2). The ANOVA revealed the significance of the main effects listed in Table 3.4—including an effect of illumination, which was not found with the physical objects experiment.

The ANOVA also revealed that there was no sign of a main effect for photographed versus rendered image type, and no evidence for interaction effects for image type (i.e. rendered vs. photographed) versus illumination, color, or shape. The lack of these main and interaction effects regarding image type (rendered vs. photographed) indicate that both rendering and photographic methods for simulating the physical environment are equivalent for purposes of object visibility measurements under the conditions of this study.

| Main Effects Means (A) | Category | Effect | Mean | SEM | N-trials |
|------------------------|--------------|----------|--------|--------|----------|
| | Illumination | overhead | 0.5 | 0.0158 | 1008 |
| | | window | 0.68 | 0.0147 | 1008 |
| | Distance | 10 | 0.74 | 0.017 | 672 |
| | | 17 | 0.6 | 0.0189 | 672 |
| | | 24 | 0.43 | 0.0191 | 672 |
| | Color | gray | 0.41 | 0.0155 | 1008 |
| | | white | 0.77 | 0.0132 | 1008 |
| | Height | 2 | 0.56 | 0.0192 | 672 |
| 4 | | 0.6 | 0.0189 | 672 | |
| 6 | | 0.62 | 0.0188 | 672 | |
| Shape | box | 0.49 | 0.0158 | 1008 | |
| | cylinder | 0.69 | 0.0146 | 1008 | |

| Selected 2-way Interaction Effects Means (B) | Illumination x Shape | Illumination | Shape | Mean | SEM | N-trials |
|--|----------------------|--------------|----------|--------|--------|----------|
| | | overhead | box | 0.44 | 0.0221 | 504 |
| | | | cylinder | 0.55 | 0.0222 | 504 |
| | | window | box | 0.54 | 0.0222 | 504 |
| | cylinder | | 0.83 | 0.0168 | 504 | |
| | Distance x Height | 10 | 2 | 0.75 | 0.0292 | 224 |
| | | | 4 | 0.72 | 0.0301 | 224 |
| | | | 6 | 0.75 | 0.029 | 224 |
| | | 17 | 2 | 0.53 | 0.0334 | 224 |
| | | | 4 | 0.64 | 0.0322 | 224 |
| | | | 6 | 0.63 | 0.0324 | 224 |
| | | 24 | 2 | 0.39 | 0.0326 | 224 |
| 4 | | | 0.43 | 0.0332 | 224 | |
| 6 | | | 0.47 | 0.0334 | 224 | |

| Selected 3-way Interaction Effects Means (C) | Illumination x Color x Shape | Illumination | Color | Shape | Mean | SEM | N-trials | |
|--|------------------------------|--------------|----------|----------|----------|--------|----------|-----|
| | | overhead | gray | box | 0.1 | 0.0185 | 252 | |
| | | | | cylinder | 0.25 | 0.0273 | 252 | |
| | | | white | box | 0.78 | 0.0262 | 252 | |
| | | | | cylinder | 0.86 | 0.0221 | 252 | |
| | | | window | gray | box | 0.54 | 0.0315 | 252 |
| | | | | | cylinder | 0.74 | 0.0278 | 252 |
| | | white | | box | 0.54 | 0.0315 | 252 | |
| | | | cylinder | 0.92 | 0.0171 | 252 | | |

Table 3.4 *Digital Bangerter Blur Means (Experiment 2)*. Identification performance.

Notable similarities between Experiment 2 and the original physical objects experiment (Chapter 2) include the following: Closer and taller objects were more visible than further and shorter objects. White objects were more visible than gray, and cylinders were more visible than boxes. All of these findings are consistent with the physical objects experiment.

| Source | Sum Sq. | d.f. | Mean Sq. | F | Prob>F |
|-------------------|---------|------|----------|--------|--------|
| Subject | 0.13848 | 6 | 0.02308 | 17.64 | 0 |
| ImgType | 0.00089 | 1 | 0.00089 | 0.68 | 0.4105 |
| Illum | 0.16329 | 1 | 0.16329 | 124.82 | 0 |
| Distance | 0.28866 | 2 | 0.14433 | 110.33 | 0 |
| Color | 0.61924 | 1 | 0.61924 | 473.36 | 0 |
| Height | 0.01191 | 2 | 0.00596 | 4.55 | 0.0106 |
| Shape | 0.19195 | 1 | 0.19195 | 146.73 | 0 |
| ImgType*Illum | 0.00116 | 1 | 0.00116 | 0.89 | 0.3469 |
| ImgType*Color | 0.00355 | 1 | 0.00355 | 2.71 | 0.0998 |
| ImgType*Shape | 0.00355 | 1 | 0.00355 | 2.71 | 0.0998 |
| Illum*Color | 0.35463 | 1 | 0.35463 | 271.08 | 0 |
| Illum*Shape | 0.03503 | 1 | 0.03503 | 26.78 | 0 |
| Distance*Height | 0.01053 | 4 | 0.00263 | 2.01 | 0.0902 |
| Color*Shape | 0.00355 | 1 | 0.00355 | 2.71 | 0.0998 |
| Illum*Color*Shape | 0.0197 | 1 | 0.0197 | 15.06 | 0.0001 |
| Error | 2.6033 | 1990 | 0.00131 | | |
| Total | 4.44941 | 2015 | | | |

Table 3.5 *Digital Bangerter Blur Identification ANOVA (Experiment 2)*. Data were arcsine transformed.

Figure 3.6 shows the interaction plots, which can be compared to the interactions for the experiment with the physical stimuli. The pattern of results in Figure 3.6 is closer to the physical data than the results of Gaussian blur (Figure 3.4), although there is still an unexpected discrepancy for boxes under window illumination. Furthermore, in the physical experiment, no main effect for illumination was found. However, the effect for illumination was strong here, and the differences were over 10% according to Table 3.4 (well outside the SEM, which was on the order of 1% for this experiment). The most notable difference is the overall mean proportion correct for identification here was 59%, and for the physical objects it was 33%. This difference in performance is also reflected in the slope coefficient in the regression analysis (detection $\beta_1 = 0.61$, and identification $\beta_1 = 0.62$), indicating that the gain on increasing object visibility is greater on the desktop than in the physical world.




Figure 3.6. Digital Bangerter Blur Identification Interaction Plots (Experiment 2). This figure shows identification performance for digital Bangerter blur. The coefficient of determination between this and original physical experiment is $R^2 = .51$. Black crosses show original physical data.

Discussion

The results of the experiment with the digital Bangerter filter are encouraging. The pattern of results is closer to the results in the physical world than a simple Gaussian blur model, indicating that the fundamental way the Bangerter blur foils filter light is

important to model with the precision that is provided by the optical transfer function (OTF) of the physical blur foils.

The remaining differences in performance between this desktop model and performance in the physical experiment need to be addressed. For example, the difference in illumination effect is substantial, along with the subtler differences found in the 3-way interactions between illumination, shape, and color. These effects could be partially due to the greater statistical power in this experiment, but that is unlikely the only cause. We suspect that the lack of effect of illumination in the physical experiment could have been due to a relatively weak effect being lost in the data noise. It is also possible that the illumination effect could be somewhere in the filtering process, but we have not been able to identify a likely source within the algorithm or code implementation.

Although 51% of the variability in identification of objects in the physical experiment was explained by this model, that only gets us half way to our goal of accounting entirely for the performance in the physical experiment. What are the sources of the unexplained variability? We speculated that another possibility might lie in the optical properties of the Bangerter occlusion foils.

Our measured PSF of the Bangerter foils did not account for glare. When viewing the world through Bangerter occlusion foils, glare enters into the filter from the surrounding light sources, and effectively lowers contrasts seen through the filter. In our earlier study of the visibility of ramps and steps (Legge et al., 2010), we found that adding an occlusion tube to the end of the blur goggles did in fact improve visibility performance. Since we recorded a PSF of the foils using a point of light surrounded by darkness (see

Chapter 4 for methods), we speculated that our digital Bangerter model did not take into account contrast reduction associated with scattered light from glare. We speculated that the nature of glare in terms of photon flux from surrounding environment would manifest itself, or could be modeled as, a simple reduction of contrast—especially for the purposes of object visibility. We designed a new experiment to test this hypothesis, Experiment 3.

Experiment 3—Effects of Digital Bangerter Blur with Contrast Reduction

Given the results of Experiment 2, we hypothesized that the digital Bangerter model did not account for glare, and that glare could be modeled, for visibility purposes, as contrast reduction within the stimulus images. In this experiment, we tested whether a simple contrast reduction model (methods are shown in Chapter 4, Section 4.2) applied to the digitally filtered Bangerter images could account for some of the differences observed between Experiment 2 and the physical experiment (Chapter 2) for detection and identification performance.

Participants

Ten subjects were recruited according to the description in General Methods.

Method

Rendered images were prepared in the same fashion as Experiment 2. A contrast reduction filter was made to vary contrast between 5% and 100% of the original images. The following equation shows how the luminance values were adjusted between 100% contrast (i.e., the original luminance/contrast values of the input image), and lesser contrast (i.e., down to 5% contrast of the original input images):

$$PixelLum_{out} = ContrastProp * ...$$

$$(PixelLum_{in} - MeanPixLum) + MeanPixLum \quad \text{Equation 3.1}$$

Through this function, an image of 100% contrast is an image that has no contrast attenuated pixel values. In other words, it is one whose contrast proportion is equal to 1 in the above equation. Conversely, a 0% contrast image would be one whose pixel values are all equal to the mean luminance of the input image pixels. In other words, its contrast proportion is equal to 0 in the above equation. Images between 0% and 100% contrast, therefore, lie between a uniform image and the original full contrast (i.e., non-contrast-manipulated, or non-attenuated) image. The mathematical derivation of the above equation, and further discussion of the upper and lower limits are shown in Chapter 4, Section 4.2. The contrast reduction filter (Chapter 4, Table 4.6) was applied in HDR, as a last step before the 8-bit JPEG transform (as shown in Chapter 4, Figure 4.2). Figure 3.7 shows samples of images filtered with the contrast reduction function.

In Experiment 3, there were 90 rendered images, 7 levels of contrast, and 10 subjects, yielding 6300 trials. Blocks and trial randomizations were as noted in general methods. The different contrast levels were included in randomized trials within blocks.




Figure 3.7. Digital Bangerter Blur with Contrast Reduction Stimulus Samples (Experiment 3). This figure shows two objects (overhead illuminated short grey box—left column, and window illuminated tall white cylinder—right column) at 10 feet. The 4 rows represent 4 (100, 60, 20, & 5%) of the 7 contrast levels tested.

Results

Correlations were computed to compare performance in this experiment with corresponding conditions in the original physical objects experiment. For very low levels of contrast (as defined by the contrast function, Equation 3.1), the coefficients of determination were low (see Table 3.1C). The best coefficient of determination comparing the digital Bangerter contrast reduction model to performance in the physical

experiment for identification was $R^2 = .58$, which occurred at 80% input into the contrast function (Equation 3.1—Matlab image processing function is shown in Table 4.6). In other words, the best match occurred when the contrast attenuation function was set to 80%. Table 3.6 shows mean performances for varying contrast level. Table 3.7 shows the ANOVA. Figure 3.8 shows the interaction plot for identification at 80% function contrast, which shows a direct comparison to the physical interactions.

| Contrast % | Identification | SEM _{ident} | Detection | SEM _{detect} | N-trials |
|------------|----------------|----------------------|-----------|-----------------------|----------|
| 5 | 0.13 | 0.0126 | 0.37 | 0.0180 | 720 |
| 10 | 0.18 | 0.0144 | 0.40 | 0.0183 | 720 |
| 20 | 0.23 | 0.0158 | 0.47 | 0.0186 | 720 |
| 40 | 0.36 | 0.0179 | 0.62 | 0.0181 | 720 |
| 60 | 0.47 | 0.0186 | 0.72 | 0.0168 | 720 |
| 80 | 0.50 | 0.0186 | 0.77 | 0.0158 | 720 |
| 100 | 0.55 | 0.0186 | 0.83 | 0.0140 | 720 |

Table 3.6 *Digital Bangerter Blur Contrast Reduction Means (Experiment 3)*. Shows mean performance of identification and detection for each contrast level.

| Source | Sum Sq. | d.f. | Mean Sq. | F | Prob>F |
|--------------------|----------|------|----------|--------|--------|
| Contrast | 3.21904 | 6 | 0.53651 | 137.73 | 0 |
| Subject | 1.17306 | 9 | 0.13034 | 33.46 | 0 |
| Illumination | 0.64923 | 1 | 0.64923 | 166.67 | 0 |
| Distance | 1.48318 | 2 | 0.74159 | 190.38 | 0 |
| Color | 3.53472 | 1 | 3.53472 | 907.44 | 0 |
| Height | 0.13403 | 2 | 0.06701 | 17.2 | 0 |
| Shape | 0.44013 | 1 | 0.44013 | 112.99 | 0 |
| Illumination*Color | 0.08769 | 1 | 0.08769 | 22.51 | 0 |
| Illumination*Shape | 0.01265 | 1 | 0.01265 | 3.25 | 0.0716 |
| Distance*Height | 0.02824 | 4 | 0.00706 | 1.81 | 0.1235 |
| Color*Shape | 0.08234 | 1 | 0.08234 | 21.14 | 0 |
| Illum*Color*Shape | 0.02509 | 1 | 0.02509 | 6.44 | 0.0112 |
| Error | 19.51139 | 5009 | 0.0039 | | |
| Total | 30.38079 | 5039 | | | |

Table 3.7 *Digital Bangerter Blur Contrast Reduction Identification ANOVA (Experiment 3)*. Data were arcsine transformed.




Figure 3.8. Interaction Plots for Digital Bangert Blur with Contrast Reduction (Experiment 3). This figure shows the contrast reduction model that produced the highest correlation with the physical objects experiment. The coefficient of determination between this and original experiment was $R^2 = .58$. Contrast level was 80% of original contrast. Black crosses show original physical data.

A reduction of contrast (i.e., scaling the contrast values to 80% of their original values through the contrast function) improved the coefficient of determination from $R^2 = .51$ (Experiment 2, Table 3.1B) to $R^2 = .58$ (Experiment 3, Table 3.1C). There was a

slight improvement in regression parameter estimates for identification at 80% function contrast as well ($\beta_0 = 0.00$, $\beta_1 = 0.68$). The second term indicates that the gain on increasing visibility is still different on the desktop than with the physical experiment, but it is slightly closer to unity. Mean performance for identification was also reduced to 50% (Table 3.6), closer to the performance level of 33% found in the original physical experiment.

The interaction plot in Figure 3.8 (as well as those in the previous two experiments—Figures 3.6 and 3.4) revealed that boxes under window illumination were systematically more visible than window-illuminated boxes in the original physical experiment. This combination of object type and illumination seems to be driving the majority of the differences in performance between desktop and physical experiments, as desktop performance for boxes under overhead illumination and cylinders under either illumination were much more closely matched to performance in the physical experiment. We hypothesized that the visibility of the window-illuminated boxes were more highly sensitive to slight perturbations in orientation and position than the other three stimulus types (i.e., window-illuminated boxes were more sensitive to perturbations than overhead-illuminated boxes or cylinders under either illumination). In the physical experiment (Chapter 2), all objects were physically placed into position, by the experimenter, for each trial. Since the front face of the boxes were almost in perfect alignment with the window luminaire (as shown in Figure 2.1), any small perturbations in box orientation or position in the physical environment could have had drastic effects on

trial by trial visibility. This effect would not have been as drastic for boxes under overhead illumination or cylinders under either illumination.

To test whether slight perturbations in orientation and position could have affected visibility of boxes illuminated by artificial windows, a series of new window-illuminated boxes with slight orientation and position perturbations were rendered. Michelson contrasts were then computed between the front face of each box and the background immediately behind and to the right side of each box.

Figure 3.9 shows Michelson contrast values for a subset of the newly rendered window-illuminated boxes. The data in Figure 3.9 reveal that slight perturbations in orientation (0 to 3 degrees toward the artificial window) and position (0 to 6 inches away from the viewing location) do, in fact, drastically change Michelson contrast. In other words, these particular stimuli are highly sensitive to very slight object perturbations—due to the fact that the frontal planar surface of the flat objects were nearly in perfect alignment with the main source of illumination.

If placement of the target location and orientation in the physical experiment (Chapter 2) were slightly more variable, from trial to trial, or slightly but systematically biased compared to the digital experiments—on the order of a few degrees or inches—then the most affected images, in terms of psychophysical responses for visibility, would be those very same boxes under window illumination. If this speculation were to be correct, then we would expect very slight random or systematic errors to affect visibility performances of the window-illuminated boxes, and the comparisons of visibility between desktop and physical experiments for window-illuminated boxes would be less reliable. The results

shown in Figure 3.9 suggest that, indeed, very small perturbations in object position and orientation can greatly change objective visibility measures (in this case, Michelson contrast), and it is highly likely that the Michelson contrast visibility effect would correspond to psychophysically measured visibility performance differences as well.




Figure 3.9. Michelson Contrast for Boxes under Window Illumination. Under artificial window illumination, the Michelson contrast between boxes and the immediate background changed as a function of small changes in orientation and location—on the order of a few degrees and inches, respectively.

In Experiments 1-3, we systematically found that a majority of the discrepancy between desktop and physical visibility performance was with the window-illuminated boxes. This systematic difference, along with the analysis described immediately above and shown in Figure 3.9, suggest that there may have been slight but significant random location and orientation errors with stimulus presentations within the physical experiment, and/or slight but systematic errors between physical and desktop stimuli. Because window-illuminated boxes were the most susceptible to visibility changes, we reanalyzed the coefficients of determination between physical versus the contrast reduction models examined in this section, this time excluding the visibly unstable window-illuminated boxes.

The reanalysis (i.e. with and without window-illuminated boxes) revealed that the highest coefficient of determination for identification was $R^2 = .77$ when contrast (as defined by the contrast function, Equation 3.1) was 40% and window-illuminated boxes were removed. Table 3.8 shows coefficients of determination, regression, and eigenvectors for identification for all levels of contrast, with all objects (shown in Table 3.8A, which is a reprinted portion of Table 3.1), as well as coefficients ignoring window-illuminated boxes (Table 3.8B).

Based on the contrast = 40% model (i.e., the highest R^2 criterion excluding window-illuminated boxes), Figure 3.10 shows identification performance of physical versus the Contrast = 40% model for all image types. As noted in Table 3.8B and in Figure 3.10, lines for regression and eigenvectors for the contrast = 40% model are close to the equality line, indicating a very good fit of this model to the physical data.

| | Contrast Level | R ² | β ₀ | β ₁ | Σβ | ψ ₀ | ψ ₁ | Σψ |
|---|----------------------|----------------|----------------|----------------|-------------|----------------|----------------|-------------|
| Identification with All Objects (A) | 5 | 0.08 | 0.3 | 0.33 | 0.62 | 0.06 | 1.61 | 1.66 |
| | 10 | 0.17 | 0.25 | 0.44 | 0.69 | 0.07 | 1.19 | 1.27 |
| | 20 | 0.28 | 0.21 | 0.52 | 0.73 | 0.08 | 0.96 | 1.05 |
| | (Fig 3.10) 40 | 0.52 | 0.11 | 0.62 | 0.73 | 0.03 | 0.81 | 0.84 |
| | 60 | 0.52 | 0.05 | 0.62 | 0.67 | -0.05 | 0.81 | 0.76 |
| | 80 | 0.58 | 0 | 0.68 | 0.68 | -0.09 | 0.85 | 0.76 |
| | 100 | 0.52 | -0.04 | 0.69 | 0.66 | -0.18 | 0.95 | 0.77 |
| | | | | | | | | |
| Identification without Window-Illuminated Boxes (B) | 5 | 0.22 | 0.34 | 0.52 | 0.87 | 0.22 | 1.27 | 1.49 |
| | 10 | 0.34 | 0.3 | 0.6 | 0.9 | 0.2 | 1.06 | 1.26 |
| | 20 | 0.48 | 0.25 | 0.65 | 0.9 | 0.18 | 0.91 | 1.09 |
| | (Fig 3.10) 40 | 0.77 | 0.12 | 0.74 | 0.86 | 0.08 | 0.83 | 0.91 |
| | 60 | 0.72 | 0.07 | 0.7 | 0.77 | 0.02 | 0.8 | 0.82 |
| | 80 | 0.68 | 0.04 | 0.7 | 0.74 | -0.03 | 0.82 | 0.79 |
| | 100 | 0.61 | 0.01 | 0.71 | 0.72 | -0.09 | 0.88 | 0.78 |

Table 3.8. *Coefficients of Determination, Regression, and Eigenvector for Object Identification with and without Window-Illuminated Boxes.* Table A is a reprinted portion of Table 3.1. Table B excludes window-illuminated boxes in the analyses. Coefficients in bold are reflected in Figure 3.10.

In summary, the pattern of results with the digital Bangerter with contrast reduction model is closer to the pattern of results of the physical objects experiment, namely: 58% of the variability in the physical objects experiment can be accounted for by the digital Bangerter contrast = 80% model. The most variance explained was 77%, when discounting the window-illuminated boxes and choosing the contrast = 40% model.




Figure 3.10. Identification performance for Physical versus the Bangerter Contrast = 40% Model.

Coefficients for these data are shown in bold in Table 3.8.

Discussion

By adding a stage of contrast attenuation, we improved the correlation between performance of subjects in the physical experiment and subjects viewing the desktop images with the digital Bangerter blur. The contrast reduction model improved the Bangerter blur model, especially for identification at an 80% function contrast level. The

improvement in the R^2 value over digital Bangerter without contrast reduction was significant. We tested the hypothesis that for identification, the 80%-contrast function model outperformed the 100%-contrast (i.e., non-attenuated) model at accurately predicting performance in the physical world. Subtracting the physical results from each model's results, we computed the deviations of the models from real for each stimulus. We then performed a 1-tailed t-test on those differences, because we knew that the 80% model was better than the 100% model. By comparing the differences between the two models, we found that in identification performance for Experiment 3, the 80% contrast model was better at predicting the physical data than the 100% model. A 1-tailed t-test found this difference to be significant ($t(77) = 2.481, p = .0076$). Furthermore, the improvement of the Bangerter blur contrast reduction model over the original simple Gaussian blur model was drastic: 58% of the variance accounted for compared with 26%. The Bangerter blur contrast reduction model provides our best approximation to the performance values obtained in the physical experiment. We will discuss potential improvements to this model later in the General Discussion.

Experiment 4—Effects of Orientation and Position

The general finding, thus far in the physical and desktop experiments, was that cylinders are more visible than boxes. However, we did not test enough viewing conditions to ensure the generalizability of this finding. To ensure that superior cylinder visibility was not a result of the specific details of the physical layout of our earlier experiments, we rendered and tested more box orientations and new object locations for boxes and cylinders. We hypothesized that cylinders would always be at least as visible

as boxes of different orientations, because cylinders have internal contrast characteristics that boxes showing only one or two flat vertical surfaces would not have. In other words, we expected that cylinders would have contrast values equal to or greater than boxes, and would therefore be as visible or more visible than boxes, under blurry vision—in any orientation or position.

We used the advantages of rendering over physical testing in order to expand the conditions examined in the original physical objects experiment. Ultimately, we wanted to determine whether different object orientations and positions in the room would interact with object visibility for different shapes. The additional viewing and object locations are shown in magenta in Figure 3.1. We extended our Radiance model of the environment with the Bangerter blur model (without contrast attenuation) to test the effects of these novel stimulus conditions.

Participants

Seven subjects in Experiment 4 were the same as those used in Experiment 2.

Method

Newly rendered HDR images were created in the same fashion as described in previous experiments. Radiance code was generated to create new orientations and positions for the objects (4-foot high gray and white boxes and cylinders which were used before). A secondary object location and viewpoint were added (shown in Figure 3.1). Viewing distance from viewer to object was 17 feet for both viewing and target locations. Figure 3.11 shows a selection of the new unfiltered and blurred images. The images in this experiment included boxes of 5 new box orientations, along with the original frontal

planar facing boxes, as well as cylinders at the 17-foot viewing location. Furthermore, images from a novel viewpoint and object location were created (shown in Figure 3.1). Rendering of fifty new images was performed; the total rendering time was approximately 24 hours on a laptop computer, or approximately 1/2 hour per image.



Figure 3.11. Rendered image samples for Digital Bangerter Blur Orientation and Position (Experiment 4). Four sample images show two new box orientations with the original viewpoint and original object location (left 2 images), and two boxes at new object and viewing location (right 2 images). One is a rotated medium height white box rotated 15 degrees with overhead illumination (column 1), and a gray box rotated 75 degrees under artificial window illumination (column 2). Columns 3 and 4 are the same object colors (white and gray), and the same orientations, but shifted to a new part of the room, with a new viewing location. Row 1 shows the original rendered and 8-bit converted images; row 2 shows the digital Bangerter filtered images.

There were 4 blocks in this experiment, including two different illumination conditions (overhead and window lighting), and two different geometrical object/viewing locations. The two locations involved testing visibility for objects along the wall (same as Experiments 1-3), and in the center of the room (viewing and stimulus locations are colored in magenta in Figure 3.1). Subjects viewed the 6 different box orientations, plus

cylinders and catch trials (i.e., images with no object) in each block. The 6 box orientations were 0, 15, 30, 45, 60, and 75 degrees, with 0 degrees being frontal planar and rotation occurring around the vertical axis (counterclockwise from the top). Figure 3.1 shows the arrangement of lighting, object location, orientation, and viewing position; for clarity, only the 0-degree and 45-degree boxes, along with cylinders are shown orthographically in Figure 3.1. Total trials for all subjects sum to 1680. Blocking and randomization is described earlier in the General Methods section.

Results

Main effects for Illumination and color were found (See Table 3.9A for means and Table 3.10 for the ANOVA). There was no main effect for position in the room, as indicated by the ANOVA and the means table. The main effect for object type (including the cylinders and the different orientations of boxes) was evident. Additionally, Table 3.9B shows selected 2-way interaction means. White objects were much more visible than gray in overhead illumination, but the effect was smaller for window illumination (which was also observed in Experiments 1-3). Additional interactions between illumination, position, color, and object type were observed. Significant findings included: illumination by color, illumination by object type, and color by object type. Object type by position was significant at the $\alpha = .05$ level (see Table 3.10 for ANOVA).

An illumination by color post hoc analysis indicated that white objects were equally visible under overhead illumination. Gray objects were slightly less visible than white under window illumination and seriously less visible under overhead illumination (see Table 3.9 for selected means).

| Main Effects | Category | Effect | Mean | SEM | N-trials |
|----------------------------|----------|----------|------|--------|----------|
| | Illum | overhead | gray | 0.38 | 0.0173 |
| window | | | 0.58 | 0.0176 | 784 |
| Color | gray | white | 0.31 | 0.0165 | 784 |
| | | white | 0.65 | 0.017 | 784 |
| Position | center | left | 0.48 | 0.0179 | 784 |
| | | left | 0.47 | 0.0178 | 784 |
| Object Type | Box = 0 | Box = 0 | 0.49 | 0.0335 | 224 |
| | | Box = 15 | 0.47 | 0.0334 | 224 |
| | | Box = 30 | 0.49 | 0.0335 | 224 |
| | | Box = 45 | 0.56 | 0.0333 | 224 |
| | | Box = 60 | 0.37 | 0.0323 | 224 |
| | | Box = 75 | 0.45 | 0.0333 | 224 |
| | | Cylinder | 0.52 | 0.0335 | 224 |
| Illum x Color | overhead | gray | 0.07 | 0.0128 | 392 |
| | | white | 0.68 | 0.0236 | 392 |
| | window | gray | 0.54 | 0.0252 | 392 |
| | | white | 0.62 | 0.0245 | 392 |
| Illum x Position | overhead | center | 0.36 | 0.0243 | 392 |
| | | left | 0.39 | 0.0246 | 392 |
| | window | center | 0.6 | 0.0247 | 392 |
| | | left | 0.56 | 0.0251 | 392 |
| Color x Position | gray | center | 0.32 | 0.0235 | 392 |
| | | left | 0.3 | 0.0231 | 392 |
| | white | center | 0.65 | 0.0241 | 392 |
| | | left | 0.65 | 0.0242 | 392 |
| Illumination x Object Type | overhead | Box = 0 | 0.42 | 0.0468 | 112 |
| | | Box = 15 | 0.38 | 0.046 | 112 |
| | | Box = 30 | 0.42 | 0.0468 | 112 |
| | | Box = 45 | 0.44 | 0.0471 | 112 |
| | | Box = 60 | 0.38 | 0.046 | 112 |
| | | Box = 75 | 0.42 | 0.0468 | 112 |
| | | Cylinder | 0.18 | 0.0364 | 112 |
| | window | Box = 0 | 0.55 | 0.0472 | 112 |
| | | Box = 15 | 0.57 | 0.047 | 112 |
| | | Box = 30 | 0.56 | 0.0471 | 112 |
| | | Box = 45 | 0.68 | 0.0443 | 112 |
| | | Box = 60 | 0.37 | 0.0457 | 112 |
| | | Box = 75 | 0.48 | 0.0474 | 112 |
| | | Cylinder | 0.86 | 0.0332 | 112 |
| Position x Object Type | center | Box = 0 | 0.47 | 0.0474 | 112 |
| | | Box = 15 | 0.52 | 0.0474 | 112 |
| | | Box = 30 | 0.55 | 0.0472 | 112 |
| | | Box = 45 | 0.6 | 0.0465 | 112 |
| | | Box = 60 | 0.36 | 0.0455 | 112 |
| | | Box = 75 | 0.43 | 0.047 | 112 |
| | | Cylinder | 0.46 | 0.0473 | 112 |
| | left | Box = 0 | 0.5 | 0.0475 | 112 |
| | | Box = 15 | 0.43 | 0.047 | 112 |
| | | Box = 30 | 0.43 | 0.047 | 112 |
| | | Box = 45 | 0.52 | 0.0474 | 112 |
| | | Box = 60 | 0.38 | 0.0462 | 112 |
| | | Box = 75 | 0.47 | 0.0474 | 112 |
| | | Cylinder | 0.57 | 0.047 | 112 |

Table 3.9 *Digital Bangert Orientation and Position Means (Experiment 4)*. This table shows identification performance for selected main effects and 2-way interactions of interest.

| Source | Sum Sq. | d.f. | Mean Sq. | F | Prob>F |
|-----------------------|---------|------|----------|--------|--------|
| Subject | 0.19047 | 6 | 0.03175 | 21.19 | 0 |
| Illumination | 0.15263 | 1 | 0.15263 | 101.87 | 0 |
| Color | 0.42396 | 1 | 0.42396 | 282.98 | 0 |
| ObjectType | 0.04197 | 6 | 0.00699 | 4.67 | 0.0001 |
| Position | 0.00058 | 1 | 0.00058 | 0.39 | 0.5334 |
| Illumination*Color | 0.25647 | 1 | 0.25647 | 171.19 | 0 |
| Illum*ObjectType | 0.15351 | 6 | 0.02559 | 17.08 | 0 |
| Illumination*Position | 0.00393 | 1 | 0.00393 | 2.62 | 0.1055 |
| Color*ObjectType | 0.02808 | 6 | 0.00468 | 3.12 | 0.0048 |
| Color*Position | 0.00021 | 1 | 0.00021 | 0.14 | 0.7086 |
| ObjectType*Position | 0.02238 | 6 | 0.00373 | 2.49 | 0.0212 |
| Error | 2.29372 | 1531 | 0.0015 | | |
| Total | 3.56791 | 1567 | | | |

Table 3.10 *Digital Bangarter Orientation and Position Identification ANOVA (Experiment 4)*. Data were arcsine transformed.

In the most surprising finding, revealed by a post hoc analysis of illumination by object type, it was shown that cylinders were the most visible under window illumination, yet they were the least visible under overhead illumination. An interaction plot (Figure 3.12) and the 2-way interaction means (shown in Table 3.9B), demonstrate this surprising result and the large effect size. In fact, as revealed by the post hoc analysis and Figure 3.12, cylinders were statistically worse than all boxes under overhead illumination, but statistically better than all but one of the box orientations (45-degrees) in window illumination (standard error bars in Figure 3.12 demonstrate the significance and large effect size).

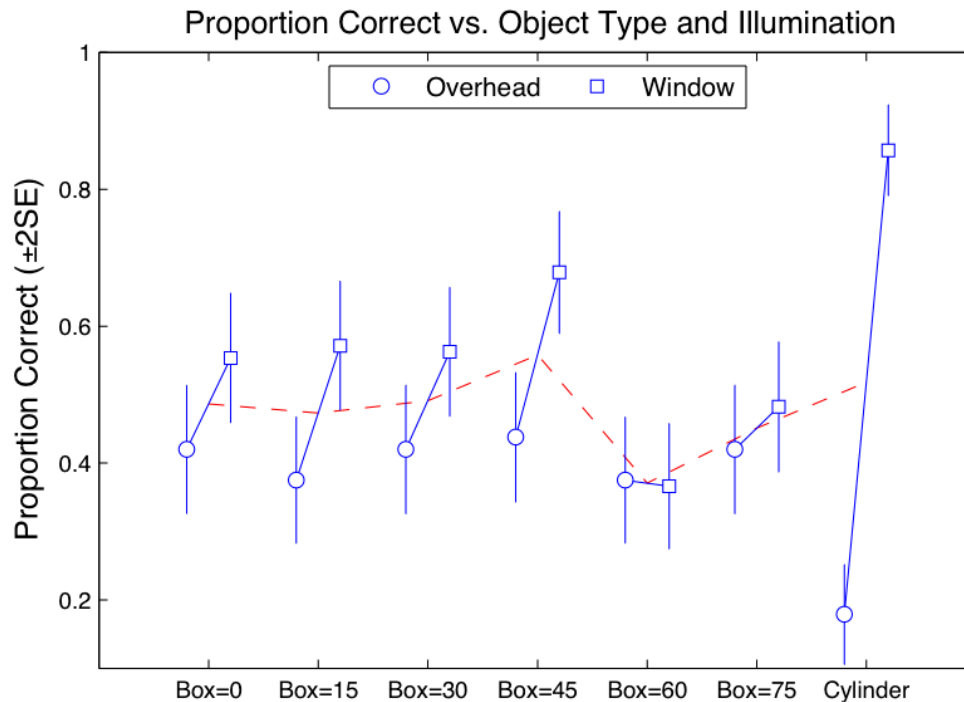


Figure 3.12. Interaction Between Object Type and Illumination (Experiment 4). This figure shows the mean identification performance across 7 object types for both types of illumination.

Discussion

In Experiment 4, we hypothesized that cylinders would be at least as visible as boxes. It was found that there is at least one case where this generalization is untrue. It was believed a priori, that since cylinders produced a curved surface that always included the 2 surface normals of a box in any orientation, that a cylinder of the same size (i.e., volume and height) and color would always be at least as visible, and often more visible, than any box. We found that this was not the case. In fact, controlling for color and size, all of our boxes were more visible in overhead illumination than the volume-matched cylinder. This is a case where cylinders are less visible than boxes, which is contrary to our general interpretive findings in previous experiments. One likely explanation is that

the contrast gradient on the surface of a cylinder is low spatial frequency, while the gradient of a box corner is a high frequency step, and some frequencies may survive the effects of certain blur filters better than others.

The most important take-home message is that we were able to test a new set of questions of physical space visibility under a virtual desktop simulation of the environment and simulated visual impairment, without the use of a physical setup and time consuming physically based data collection.

Summary and Concluding Discussion

There are a number of noteworthy advantages to visibility testing on a desktop computer display rather than using physical models of the environment. The most notable is that the desktop is much more economical in terms of capital investment, time required to build prototypes, and efficiency of psychophysical data collection. Going from scene concept to model building to scene/object rendering and testing is much quicker and cheaper than building and manipulating physical scenes and objects. It is also much safer to construct virtual objects as opposed to physical objects. As we have demonstrated, this approach allows visibility of objects in possible physical configurations to be compared with rendered computer models before actually building the environment.

In Experiment 2, we found that photographed and rendered images were no different in terms of detection and identification (i.e., visibility) performance. This leaves us to consider the bridge to performance from desktop to physical spaces. We saw improvement in each step of the filtering methods developed in this research (i.e., from the models in Experiment 1 through Experiment 3). These results apply to conditions of

severe blur, and may not extrapolate to visibility studies involving the higher acuities of normal vision.

The digital Bangerter model with contrast reduction explained 58% of the variability of visibility of the physical experiment. This is a modest but promising level of agreement between desktop and physical visibility performance. What are the remaining sources of variability not accounted for and how can the desktop model be improved?

It is possible that the range of original stimulus conditions were limited, resulting in a somewhat compressed range of performance levels. Correlations are weaker when the range of dependent variables is small. Furthermore, there may be inherent variability between individual subjects and experiments.

Some possibilities for the overall higher performance of subjects in the desktop experiments include: 1) Subjects may have had greater knowledge of where to look in the desktop environment versus the physical environment. Increasing the uncertainty of target locations in the desktop environment might reduce overall desktop performance and improve the match to the physical experiment. 2) There may have been glare effects that are nonlinear in the Bangerter goggles, not captured by our PSF and contrast reduction modeling, which could be tested with a series of digital Bangerter PSF images that include a variety of ambient illuminations during filter digitization. 3) Testing was limited to dynamic range capabilities of display device.

In summary, three categories of possible factors limiting the empirical equivalence between the desktop experiments and the original physical experiments include: 1) what subjects knew (e.g., priors, where to look, and other group differences); 2) technology

limitations (e.g., display, PSF); and 3) inherent noise in this type of testing. It is likely that a combination of factors limited our results.

The greatest limitation of our desktop approach is that the scope of visual impairment modeling was severely limited by a simple pair of physical blurring goggles. In other words, the Bangerter blur foils simply cannot simulate a wide variety of visual impairments. The Bangerter blur model used here only tested a single prototype for visual deficit. The Bangerter blur model has been shown to simulate some forms of visual impairment, as discussed by these authors (Odell et al., 2008; Pérez et al., 2010), and as demonstrated in our own lab (Bochsler et al., In Press). However, many more models of visual impairment should be explored—especially models that are more representative of real visual deficits. Since our filtration method is modular, it is intended that different blur, contrast, and field restriction models can be exchanged with our blur model prototype. In other words, if one can conceive of an optical transfer function which models visual impairment, then the model can be tested using the modular methods described here.

Another limitation of the study was a lack of testing realistically complex scenes. This problem can be overcome by developing more complicated scenes with the same photographic or rendering approaches used here (or a combination of real, photographed, and rendered as demonstrated by this research). In addition to static images, modeling motion is necessary for realistic models of visually impaired navigation. This can be accomplished with video cameras, or with Radiance's rendering functions, as demonstrated in this dissertation's accompanying video archive. In the video, a simulated

10-foot walk down a 4-foot-wide walkway with railings approaches an obstacle. The room is the same that was used in the experiments described in this dissertation, however the lighting was changed to recessed spot lights (instead of fluorescents), and the target object was a step (instead of boxes or cylinders). This demonstration was performed with the same Radiance model used in this research, with the addition of using the animation features provided by the Radiance software suite.

In Figure 3.13, we show a complex scene that was photographed using the methods described in Chapter 4, Section 4.1.2. A scene of this complexity could also have been obtained by Radiance modeling and rendering. It shows a sunlit lobby with a variety of distracting features and obstacles (Figure 3.13, upper left panel). Our digital Bangerter blur model was applied to the top right image, and the contrast reduction model was applied to the bottom left image. This is what a person would see if wearing our physical Bangerter blur goggles in the actual lobby. To further extend our digital model, and to demonstrate the modularity of the approach used in this research, we have added a simple visual field loss (i.e., a digitally created scotoma) to the bottom right image. It is hoped that future research adapts the inherent modularity in our simulation of visual impairment. Instead of using a Bangerter PSF and a simple contrast reduction model, we hope that future research can replace the blur models described here with models of visual impairment that are more inline with physically measured visual impairments. The methods presented here, it is hoped, can facilitate that goal.




Figure 3.13. Lobby with Various Digital Filters. The top-left image is of a sunlit lobby with realistically complex illumination, objects, and hazards. The top-right image shows the same scene with a 100% contrast digital Bangerter filter applied. The bottom-left shows the contrast reduced digital Bangerter filter applied. The bottom-right shows a contrast reduced digital Bangerter filter with a simple scotoma model applied to the filter, which simulates the effects of visual field loss.

Chapter 4:

Generating and Displaying Photometrically Accurate Stimulus Images

Introduction

In this research, the word *image* is used in two ways: Sometimes, *image* refers to the pattern of luminance presented to a human observer. And sometimes, *image* refers to an array of numbers stored in memory on a camera or computer. Stimulus images were generated and displayed to be geometrically and photometrically accurate 2D replications of the physical world. While geometric accuracy was handled by optical and trigonometric calculations (as noted in Chapter 3, General Methods), photometrically accurate reproductions of the scenes and stimuli require a number of steps, which are explained here in Chapter 4.

We have often commented on the photometric accuracy of the digital images as representations of scene luminances. We ensured photometric correspondence between digitally and physically projected scenes and objects by making direct comparisons of image luminance values on the calibrated computer monitor with luminance values of corresponding areas in the physical space. The test pattern used for ensuring accurate luminance and contrast reproductions is, in fact, found within each of the stimulus images used in this research. The test pattern was made from a 3-ft by 3-ft black, gray, and white colored checkerboard, which is visible in Figures 3.3, 3.5, and 3.11.

Using the checkerboard test pattern, we sampled luminance values from the physical viewing positions for both the computer-displayed images and the real-world (i.e., physical) images. The luminance values for physical and desktop images are shown as blue crosses in Figure 4.1. They are the main source of the photometric calibration technique described in this chapter—used to match the desktop stimuli to the physical

stimuli. Concurrently, the statistical properties of the resulting analysis (shown in Figure 4.1) represent the physically based source of confirmation of luminance and contrast reproduction fidelity for the desktop stimuli. The coefficients of determination all exceed .996, which tells us that greater than 99.6% of the variation in sampled luminance values in the physical images are accounted for by the computer images. The intercepts and slopes of the regression analysis are equally impressive (i.e., close to 0 and 1, respectively), which tells us the parameters of the recalibrated image patches are not only highly correlated with the physical image patches, but they are also one-to-one with the measured luminance values. We realize these numerical comparisons are only directly measurable and comparable using a photometer for the non-spatially filtered images, but this is an important validation step for subsequent filtering steps in the process. Indirect validation methods for filtered images are discussed in the next paragraph.

Moreover, and of equal importance, for filtered images where measurements were not obtainable by direct physical observation (i.e., by a photometer), mathematical induction was used as a guide for filter design and validation—ensuring luminance and contrast fidelity and consistency with an ideal physical filtering medium. In other words, where direct measurements were unobtainable, indirect methods were used to ensure the filtered luminance values were attenuated in a predictable, well controlled, and repeatable manner. In summary, the spatial filters were carefully designed to modulate simulated photon flux in an intentional, predictable, controlled, and repeatable fashion. These filtering methods are noted in detail in this chapter.

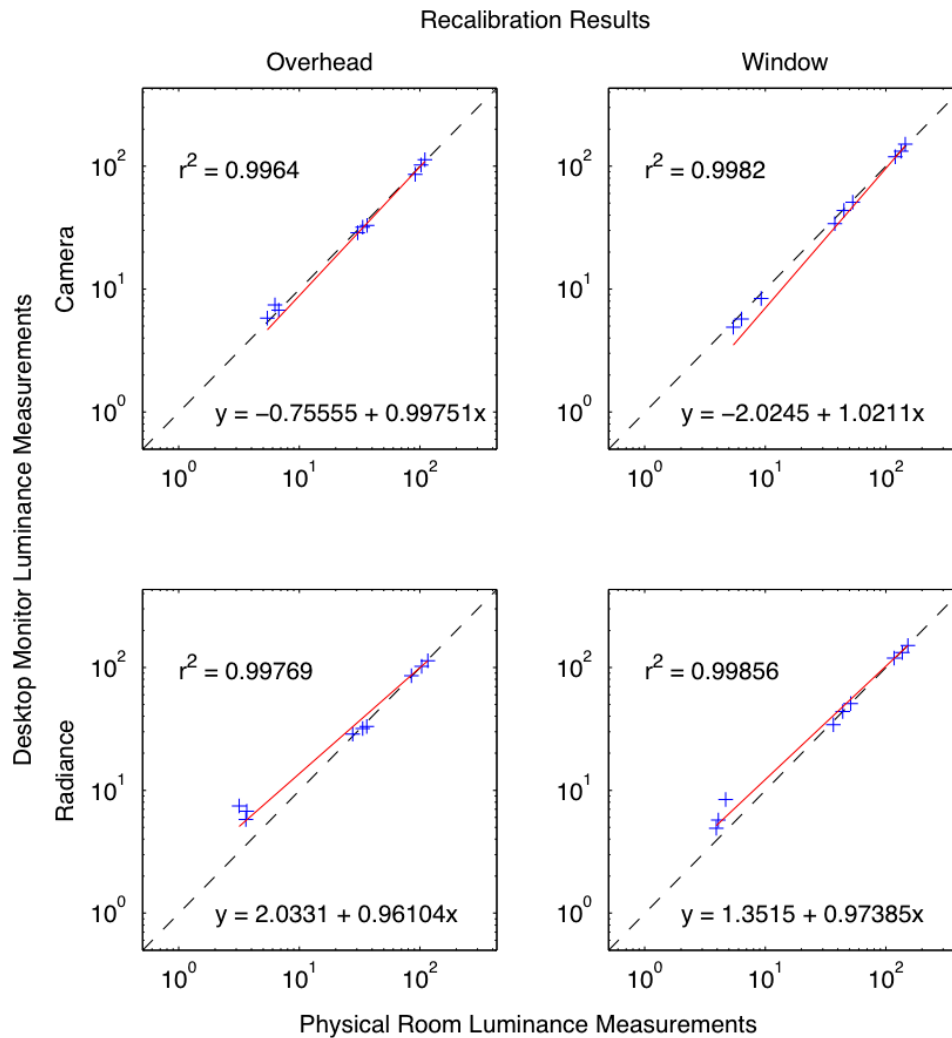


Figure 4.1. Desktop Luminance Recalibration and Verification. Each plot represents the correspondence between desktop versus physical output luminance values for the checkerboard test pattern measurements for each of the two image types (Photographed and Rendered), and for each of the two illumination types (Overhead and Window). Every physical, photographed, and rendered stimulus image had this checkerboard test pattern in the same physical location (as can be seen in Chapter 3, Figures 3.3, 3.5, and 3.11). The checkerboard test pattern had three each of black, gray, and white flat (i.e., closely Lambertian) painted squares. The domain and range of these plots are the same as the capable luminance range of the 27-inch Cinema Display used in the desktop experiments. Coefficients of determination and linear regression equations are shown for each image type. The plots were scaled in \log_{10} (cd/m^2) for easy data

point viewing and for examining the ratios between x and y. This allowed us to verify luminance and contrast equivalence between physical and desktop images.

In general, in order to display images with the same luminance properties as the physical world, the first step was to create photometrically accurate HDR digital images. The HDR images were then transformed and filtered in various steps to display JPEG output images on a screen with the same luminance and contrast values as the original physical scenes. Figure 4.2 shows a flowchart from HDR image to display JPEG image. Note that an HDR image is a representation of dynamic ranges beyond the capabilities of current display systems, and cannot be displayed or viewed properly without some form of numerical transform, such as from HDR to an 8-bit JPEG (as is standard for current display systems, and as was used in these experiments). The major limitation of the image reproduction method described here is the dynamic range and quantized character of the display system. Since HDR display systems are more expensive and less common than off-the-shelf 8-bit display systems, we chose to experiment only with a widely available and commonly used 8-bit LDR display system.

To ensure luminance accuracy, a well-calibrated display system was made, including hardware, software, and a collection of scripts, as noted in this chapter. Each step described in this chapter (including spatial filtering) affects the displayed luminance and contrast values for each pixel of each stimulus image. Because of its importance on *visibility*, and because controlling luminance and contrast requires a number of carefully planned and measured steps, a separate chapter was dedicated to the issue. This chapter is a record of the essential process developed for this experiment. The process described

here can be reused and refined for calibrating and displaying luminance and contrast accurate images in future experiments.

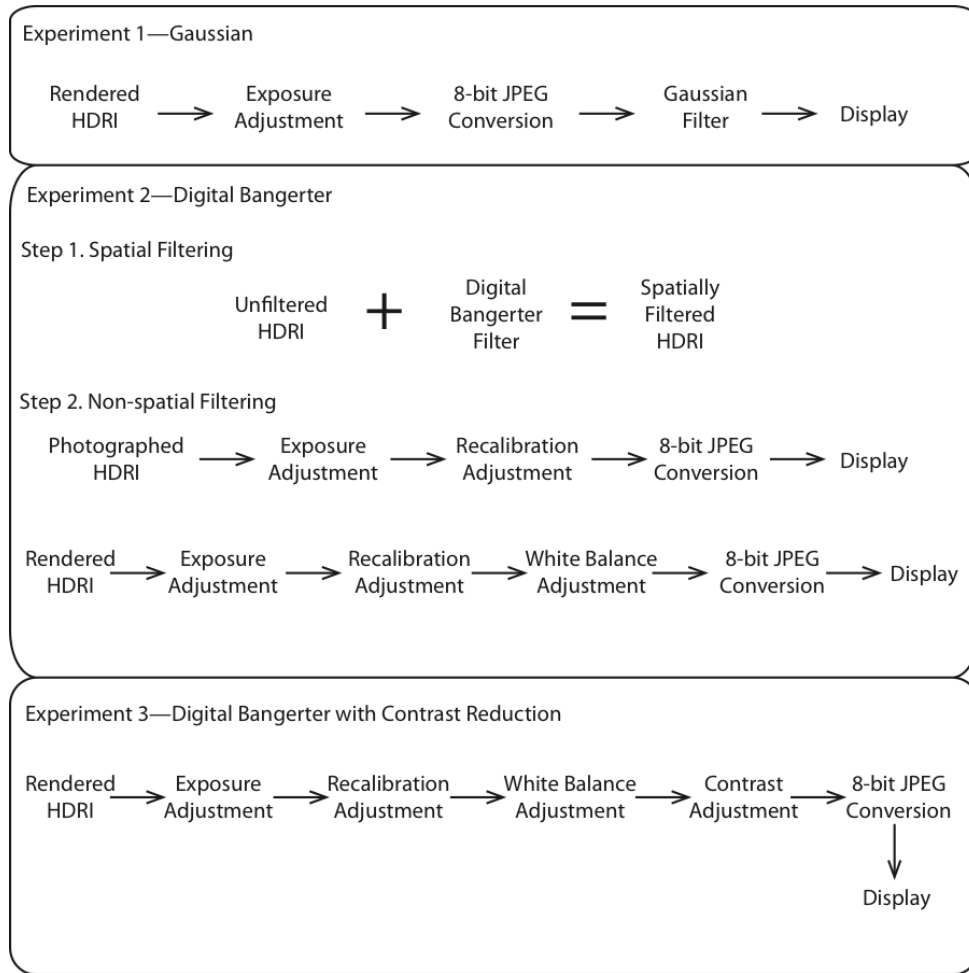


Figure 4.2. HDR Image to 8-bit Display Image Conversion Steps. Image preparation flow diagram for display image preparation for all experiments (Experiment 4 used the same method as Experiment 2 in this diagram).

The process from HDR conception to photometrically accurate displayed output image is described in four sections, including: 1) high dynamic range (HDR) image

formation—Section 4.1; 2) calibration issues and techniques—Section 4.2; 3) spatial image filtering techniques—Section 4.3; and finally 4) non-spatial image filtering techniques—Section 4.4. Together, these steps take the reader from scene and image conception to photometrically accurate, viewable stimulus images on a calibrated computer display. As noted earlier, computer-displayed to physical-world photometric validation (i.e., numerical proof) of this process is shown in Figure 4.1.

4.1. HDR Image Formation

In the original scene, as well as in renderings and photographs, a white, gray, and black, 3-foot by 3-foot checkerboard pattern was included in every stimulus image. This pattern was used for calibration purposes throughout the calibration process (i.e., at every step of image development), in order to ensure photometric precision and accuracy.

Two types of digital HDR images were created—including rendered and photographed images. Renderings were made from models using the Radiance software suite, while photographs were taken from the physical testing space. Rendered and photographed image generation is described in the following two subsections (4.1.1. & 4.1.2), respectively. All images were stored in Radiance 32-bit HDR image format, which is capable of storing the wide ranges of light levels found in natural scenes, including both color and luminance information. Throughout the process, 32-bit HDR and 8-bit JPEG image pixels were inspected, and luminance sampling and verification was performed using Radiance (Ward Larson, n.d.-b), Photosphere (Ward Larson, n.d.-c), Matlab (*MATLAB R2011a*, 2011), and/or a Minolta CS100 Colorimeter, where necessary.

4.1.1. Generating the Radiance Model and Rendering Images. Rendered images were created from a virtual model of our original physical testing space using the Radiance software suite (Ward Larson, n.d.-b). Rob Shakespeare created the model of our physical testing space—the methods for which are described in the currently available *Rendering With Radiance* textbook (Ward Larson & Shakespeare, 2011), which is the textbook that is available from the book's current publisher as of the writing of this dissertation. Shakespeare's model included accurate geometrical modeling of the physical objects, accurate material reflectance properties of the objects, and accurate modeling of the light sources. These values were obtained from direct photometric measurements of reflective materials, and from lighting specifications provided by lighting manufacturers.

We supplemented Shakespeare's model of the classroom space with key features of our particular setup, including: lighting arrangements, additional physical objects (including boxes, cylinders, and the test checkerboard pattern), viewpoints that were used in the original experiment (Chapter 2), and novel experimental conditions (Chapter 3, Experiment 4). As noted in the General Methods section of Chapter 3, geometry was computed for the viewpoints such that the output images subtended nearly the same visual angles as in the physical experiment when viewed by subjects from a fixed viewing point.

Once the Radiance model was prepared, rendering was performed using the Radiance executive program, called *rad*. The rendering *quality* was set to *medium*, which *rad* uses to set a number of rendering variables. The details of the *quality* settings are technical and beyond the scope of this report, however a brief and informative summary is described in

the following pages, along with a brief description of how to execute the Radiance rendering engine.

The *medium quality* setting allows a balance between rendering *quality* and computational time required to perform the operation on a computer. *Low quality* renderings take about a minute on a modern personal computer, but lack the realism required to get photometrically accurate images. For example, an image created from the *low quality* setting made areas such as the ceiling appear too dark. Images created from the *high quality* setting are photometrically accurate and realistic, but they take 1-2 days to render on a modern personal computer. The *medium quality* was the best setting that could render the hundreds of images required for this dissertation—in a timely fashion (about 30 minutes each), with realistic appearance, and high photometric fidelity across the entire image plane. There are many more settings that *rad* controls, given a specific *quality* setting. These settings are described in detail in the Rendering with Radiance book (Ward Larson & Shakespeare, 2011). The chapters on direct calculation (Chapter 11) and indirect calculation (Chapter 12) give detailed technical descriptions for the *quality* settings used in the renderings created here.

Forming a Radiance model description is well beyond the scope of this dissertation. However, the rendering process is briefly covered here to describe both how and why we chose the rendering settings used in this research. Here is a brief top-level programmatic description of how HDR images were rendered, using the prepared Radiance model description: The *rad* executive program creates a cascade of automated commands and files. The basic sequence of commands and files are as follows. First, an *ocnv* command

forms an *octree* file, followed by commands *rpict*, *pfilt*, and *rm*, which render and filter an image, and finally remove the temporary image file, respectively. The example commands shown here are what a typical *rad* executive program compiles and executes to create a single image:

```
>> oconv material.rad room.rad lights.rad object.rad > octree/example.oct
>> rpict -t 60 -vu 0 0 1 -vf views/example.vf -x 4224 -y 2816 -dp 512 -ar 54 -ms
6.1 -ds .3 -dt .1 -dc .5 -dr 1 -sj .7 -st .1 -ab 2 -aa .15 -ad 800 -as 128 -av 0.01 0.01
0.01 -lr 8 -lw .002 -av .01 .01 .01 -ds .1 -ps 6 -pt .08 octree/example.oct >
pics/example.unf
>> pfilt -r .6 -r .6 -x /2 -y /2 pics/example.unf > pics/example.hdr
>> rm -f pics/example.unf
```

These commands appear complex, but they were actually generated from the output of the *rad* command using a prototypical *rif* file as input. Using a verbose setting in *rad*, the above commands were extracted from the Bash shell's output, in order to produce multiple images within a single Bash script, as described in the following paragraph (note that a number of different shells could be used with Radiance, depending on operating system and computer settings). However, for one-at-a-time image prototyping, it is more convenient to use *rad* on a *rif* file, snippets of which are shown in Table 4.1. The *rif* file is the text file that possesses textual links to all the materials and objects for a particular scene, and is the input for the *rad* executive program. Note that the many individual materials and object descriptions (in the form of many additional individual text files) are not shown here—most of those files were provided by Shakespeare, and the balance were

written for this particular set of experiments. Furthermore, note that some of the key *rpict* and *pfilt* command settings displayed above are a result of the *medium quality* setting within the *rif* file (shown in Table 4.1).

```
###rif file
DETAIL= med
INDIRECT= 2
PENUMBRAS= false
PICTURE= pics/example
QUALITY= med
REPORT= 1
RESOLUTION= 2112 1408
UP= Z
VARIABILITY= med
ZONE= Interior -1 400 -5 226 0 101
materials= material.rad
render= -av .01 .01 .01 -ds .1
scene= room.rad
.
.
scene= object.rad
objects= room.rad lights.rad ...
.
.
objects= object.rad ...
.
.
pfilt= -r .6
view= example -vf views/example.vf
```

Table 4.1. *Sample rif File for Rendering Radiance Images*. Snippets of a sample *rif* file show essential settings and code used to render images. Vertical and horizontal dots indicate snipped code.

Since the output of the *rad* command with a given *rif* file actually produces the shell commands (as shown above the previous paragraph), the output can be copied and pasted into an executable Bash (or other shell script) text file. There, changes can be made to each argument, and ultimately a complete list of Radiance images can be produced from a single executable script containing all the slightly varied sets of the example commands

required for a complete image dataset. For details see the Radiance manual pages (“Radiance Manual Pages,” 1997) or the Rendering With Radiance Text (Ward Larson & Shakespeare, 2011).

The reason for paying special attention to the *medium quality* setting is not to relay the details of the process of creating an image. Rather, the critical idea for the research presented here is that any modern inexpensive personal computer is capable of rendering many photometrically accurate images without requiring many days or weeks to perform the necessary computations. In other words, one does not need access to a supercomputer to make images of high enough quality for low-vision visibility studies. A supercomputer (or lots of time) may be required to obtain a large collection of pixel-perfect images required in higher resolution domains of research (e.g. studying texture, specular reflections, or other fine visual details), however, with the images used here, there were no *visibility* differences between renderings of *medium quality* and corresponding high-resolution photographs. This idea was confirmed by experimental hypothesis testing described in Chapter 3. In summary, balancing fidelity with obtaining data within reasonable time is an important take-home message demonstrated by this research.

For Experiments 1-3 (Chapter 3), 78 rendered images were created, and they were examined for photometric accuracy, meaning: 1) the matrix of screen output luminance values of the checkerboard patterns were compared to the photometric values produced by the checkerboard in the physical space; and 2) photometric measurements of the screen output were compared to photometric measurements of the physical space. In regards to this section (i.e., Section 4.1), photometric accuracies were inspected in the

HDR image matrix. In the final analysis (the results of which are shown in Figure 4.1, and are described in the recalibration Section, 4.2), display pixel luminance measurements were compared directly to physical luminance measurements in the real space, and were exceptionally accurate, as noted in the statistical decomposition shown in Figure 4.1.

4.1.2. Obtaining Photometrically Accurate Photographs. Creating photometrically accurate images from physical environments requires three steps, including: camera calibration, collecting scene images, and finally forming HDR images. The necessary steps are described in the following paragraphs.

In order to ensure that the HDR images were linear with respect to photometry in the physical world, a Nikon D90 digital SLR camera was calibrated according to the calibration method described by Axel Jacobs (Jacobs, 2011a). Note that much of the workflow calibration methods used in this research (i.e., for calibrating HDR images to the display) were inspired by the calibration and imaging methods described in Jacobs' WebHDR software instructions his accompanying theory (Jacobs, 2011b). A camera calibration apparatus (shown in Figure 4.3) was built to conveniently and reliably calibrate the digital SLR camera. The calibration apparatus was built out of black foam core, with gray and white patches of construction paper. The box was placed beneath a sunlit southward-facing opened window in a darkened room where a shaft of direct, unfiltered sunlight could reflect off the top of the box. Inter-reflections and the dark ambient lighting of the room, along with the shape and reflectance properties of the box, allowed calibration of the camera across an enormous range of luminance values—

because the apparatus provided large patches of variable reflectance at substantial intervals across the luminance range. In the calibration image shown in Figure 4.3, luminance values were between 5 and 20,000 cd/m^2 , which were again due to the design of the box and its placement within the environment. Figure 4.3 shows a tone-mapped JPEG and false color version of the HDR image that was used to calibrate the camera.




Figure 4.3. Camera Calibration Apparatus. The left image shows the calibration apparatus, and the right image maps luminance information in false color in nits (i.e., cd/m^2), which are the reflectance values estimated by the camera calibration algorithm. Images are tone mapped here for *viewability*.

The HDR calibration image shown in Figure 4.3 was formed as follows: Exposure bracketed images of the calibration device were collected using Sofortbild Mac Tethered Camera Shooting software (Hafeneger, 2009). The image sequence was then fed into Ward's *hdrgen* software (Ward Larson, n.d.-a). The command used to create the camera calibration file from the bracketed images was (shown using the Bash shell):

```
>> hdrgen -r camera.rsp -o camera.hdr camera/*.JPG
```

This procedure allowed the camera to produce a linear response curve across the entire range of luminances recorded by the camera. In addition to the camera response

function (i.e., the 'camera.rsp' output of the above command), a calibration factor (CF) will be used to ensure that the luminance values in the image are one-to-one with physical luminance values. The calibration factor method and equation is discussed by Jacobs here (Jacobs, 2011a), however we performed the calibration process (or more specifically, the *image recalibration process*) at the last stage of output image formation. In this regard, a variation of Jacobs' calibration factor computation step was used at a latter stage of image formation (i.e. at the desktop computer monitor output stage), in order to accurately recalibrate images (described in Section 4.2.2). In summary, the camera was calibrated to correlate with luminance, but the range was not yet adjusted (again, ranging is performed later, and is described in Section 4.2.2).

After the camera was calibrated, the second step was to produce HDR images of the stimuli. Exposure bracketed images were recorded using the same tethered camera shooting software as above (Hafeneger, 2009), and bracketed images were collected according to the bracketed shooting instructions described by Jacobs (Jacobs, 2011a).

In the third and final step, the exposure bracketed stimulus images were then fed into the *hdrgen* software, along with the camera response function obtained earlier. The *hdrgen* command for creating images is similar to the *hdrgen* command used to calibrate the camera, as noted above, except the camera response function (called *camera.rsp* in the above example) is taken as an input rather than produced as output. The image output of the *hdrgen* command is a Radiance format, photometrically linear HDR image. As with the 78 original stimulus renderings, there were 78 photographed stimulus images that were recorded—i.e., one for each unique stimulus. The camera's visual angle,

viewing positions (locations and orientations), and scene/object positions were the same as those used for rendered images. Details for the geometrical settings of the HDR photographed images created here (and for renderings from Section 4.1.1) are noted in the General Methods section of Chapter 3.

4.2. Calibration Issues and Techniques

In order to ensure correct on-screen pixel luminance values were exact with respect to physical luminance, the computer system was calibrated (Section 4.2.1), and the output JPEG images were adjusted to incorporate the luminance response of the output display (Section 4.2.2).

4.2.1. Accurate 8-bit Monitor Calibration. Computer systems typically include displays with 8 bits of intensity resolution for each of the RGB (red, green, and blue) color channels. Modern off the shelf color profiling and calibration systems are widely available, and are well suited for calibrating the RGB profile of computer display systems. A 27-inch Apple Cinema Display was color calibrated with a Spyder 3 Pro and the accompanying DataColor monitor color profile software (*DataColor Spyder 3 Utility*, 2007). This procedure ensured that the luminance and color values of the 8-bit display system were linear with respect to 8-bit images being fed into the display.

4.2.1.1. Note on Displaying 32-bit Images. The 32-bit Radiance images cannot be displayed on 8-bit display systems. They must first be transformed and converted to an 8-bit image format, however this procedure will be discussed later in section 4.2.2.

4.2.1.2. Note on Camera Versus Display Calibration. To understand the theme or metaphor of calibration, the Spyder 3 Pro calibration procedure is to the

display, what the camera calibration steps in Section 4.1.2 are to the camera. Both calibration techniques ensure a linear response between HDR image and physical luminance in a numerical sense. In other words, the calibration procedures in both cases ensure a nearly perfect correlation between image types.

4.2.1.3. Note on Recalibration. In order to ensure the gains are correctly adjusted for the specific display system being used (in this case, the desktop monitor), an additional step, termed *recalibration*, is required. In other words, the numerically accurate images must be readjusted so that they display numerically accurate luminance values on a screen of known dynamic range. I have introduced the term *recalibration* to distinguish that procedure from the presently discussed calibration steps. Recalibration will be described in Section 4.2.2. For now we are only concerned with linear calibration of physical devices to ensure high correlations. Ranging is the issue that recalibration deals with in Section 4.2.2.

Using a Minolta CS-100 Colorimeter and a 3-parameter robust curve fitting procedure, the monitor's dynamic range and gamma response curve were measured. The minimum monitor luminance was estimated to be 0.5 cd/m², and the maximum luminance was 432 cd/m². This procedure allowed us to know the darkest and brightest capabilities of the monitor. After obtaining the dynamic range of the calibrated monitor, we then looked back to the original stimulus images, and found that the only scene locations that exceeded our display's dynamic range were at or very near the luminaires found within the image. Those areas were not physically close to the stimuli (i.e. boxes and cylinders) in the 2D images, so we reasoned that the resulting clipping of bright areas

(during the 32-bit to 8-bit conversion, described in Section 4.2.2) would probably have little impact on the visibility of the objects being tested. Note that we will use the photometer again later for *recalibrating* images to set the gains, but the current step was simply required to ensure that the monitor was capable of showing our stimuli without clipping the bright areas on or around them. If our boxes and cylinders had specularities, this clipping issue would have required addressing. In other words, if specularities were present within the objects, then their dynamic ranges may have exceeded the brightness of the monitor, and a different experimental approach would have been required to properly display the objects.

In summary, the monitor was calibrated quickly and easily with an inexpensive and easy to use color calibrator. The rest of Section 4.2.1 (i.e. subsections 4.2.1.1, 4.2.1.2, & 4.2.1.3) was used to clarify terminology and methods. By this step in the calibration, we have now ensured highly correlated luminance correspondences between the physical world, rendered images, photographed images, and display output. Additionally, however, absolute correspondences between physical and displayed luminances are required. In the next section, we will describe how to convert HDR images into displayable 8-bit images, and furthermore how to readjust (i.e., *recalibrate*) the pixel values during conversion—for a specific calibrated monitor in such a way as to ensure that the luminance values shown on the monitor are identical to the luminance values of the stimuli in the physical world.

4.2.2. Recalibrating and Displaying 8-bit Converted Images. The digital HDR images used in this research contain 32-bit RGB (i.e., red, green, & blue) pixel values.

Physical luminance values can be derived from Radiance HDR images using the following formula copied directly from Ward's Radiance File Formats Report (Ward Larson, n.d.-d, p. 33):

$$\text{luminance} = 179 * (0.265*R + 0.670*G + 0.065*B) \quad \text{Equation 4.1}$$

Radiance HDR images can be inspected or analyzed numerically and graphically using Photoshop, Matlab, Radiance commands, or Radiance based graphical utilities such as Photosphere (Ward Larson, n.d.-c). The 32-bit HDR images do not simply have scaled values from 0 to 1 (or, more technically, 0-255 as with 8-bit image formats). For this reason, the 32-bit HDR images must be transformed into 8-bit images that are viewable on 8-bit display systems. This procedure is essentially how graphical programs like Photosphere and Photoshop transform HDR images to make them viewable on a typical display system. Often, the process is referred to as *tone mapping*, but that term is loaded because it is commonly associated with nonlinear transformations from 32-bit to 8-bit images. Moreover, oftentimes when images are transformed using common *tone mapping* algorithms, the output image is optimized for *viewability*, but less often for *visibility* (i.e., photometric accuracy). For the purposes of our study, photometric accuracy (i.e., *visibility*) was a necessary quantitative goal. As long as the display device (i.e. our calibrated monitor) can display the range of pixel intensities we are interested in, then photometric accuracy and *visibility* will co-emerge during the formation and display of calibrated (and, moreover, *recalibrated*) images.

To make images according to the guidelines described in the preceding paragraphs, a Matlab/Photoshop/Javascript function was developed for transforming our 32-bit HDR images to photometrically accurate display images (shown in Table 4.2).

```

fid = fopen('Image.hdr');           % get image header information
for i = 1:10
    d{i} = fgetl(fid);
end
Exposure = eval(d{j}(10:end)); % get exposure value from header
I = hdrread('Image.hdr'); % open image pixel matrix in Matlab
I = I*(1/2^Exposure); % adjust pixel levels to exposure value
I = I*RecalFactor; % adjust pixel levels to recalibration factor
I = ColorAdjust(I,ImageType); % color balance (rendered only, Exp 2-4)
I = ContrastAdjust(I,ContrastLevel); % contrast adj (Experiment 3)
psopendoc([pwd, '/Image.hdr']); % open image in Photoshop
pssetpixels(I); % set Photoshop pixels to Matlab values
pstxt = 'activeDocument.bitsPerChannel = BitsPerChannelType.EIGHT;';
psjavascriptu(pstxt); % change to 8 bit in Photoshop
pstxt = ['var jpegQuality = 10;'. ... % jpeg quality
'var saveFile = "',pwd, '/Image.jpg";'. ... % save location
'SaveJPEG( new File(saveFile), jpegQuality );'. ...
'function SaveJPEG(saveFile, jpegQuality){'. ...
'jpegSaveOptions = new JPEGSaveOptions();'. ...
'jpegSaveOptions.embedColorProfile = true;'. ... % color profile
'jpegSaveOptions.formatOptions = FormatOptions.STANDARDBASELINE;'. ...
'jpegSaveOptions.matte = MatteType.NONE;'. ...
'jpegSaveOptions.quality = jpegQuality;'. ...
'activeDocument.saveAs(saveFile, jpegSaveOptions);'. ...
'app.activeDocument.close(SaveOptions.DONOTSAVECHANGES);'. ...
'}'];
psjavascriptu(pstxt); % send command to PS

```

Table 4.2. *HDR Image to Recalibrated JPEG Image Script*. This example script reads an HDR image in Matlab and Photoshop, adjusts the image to the exposure, readjusts the image to a recalibration factor set by the user, and calls Photoshop to convert from 32-Bit RGB floating-point HDR image to 8-bit RGB JPEG. If the recalibration factor is set correctly (according to a measured ratio of luminance values measured on the desktop display over the luminance values measured in the physical scene), then the output image will be photometrically accurate on the display (i.e., in relation to the actual photometric values in the physical scene image). Note that a single recursion of this script is required to get the *recalibration* set correctly, with *RecalFactor* = 1 on the first pass. On the second pass, *RecalFactor* must be set as noted in the text in Section 4.2.2.3. Clipping of bright values will occur at maximum luminance of

the display. (Limitations of the display would not be addressed with this script. For example, luminance linearity at or near the blacks is not usually accurate with an LCD display system. Also, color calibration devices and/or functions built into the display system or written by the researcher must maintain correct gamma correction—but this is generally automatic with modern computer systems and color calibration hardware and software, as noted in Section 4.2.1.) To use this script, the Photoshop Matlab Toolbox setup notes are here (*Adobe Photoshop MATLAB Toolbox*, 2011), and the Photoshop Javascript code was derived from here (*Adobe Photoshop CS5 Scripting Guide*, 2010). Portions of the code in this example script were taken from these cited Adobe documentation sources. Dependencies are described later in this chapter.

The process noted above and shown in Table 4.2 is called *recalibration* for two reasons. First, the recalibration process must not be confused with the many calibration steps required in the formation of HDR images (Section 4.1), or with calibrating the display system (Section 4.2.1). Second, recalibration is an iterative process, which requires (in our case) displaying the images once for estimating a calibration factor (CF, a term borrowed and slightly modified from Jacobs, 2011a); and then for displaying them again for quality assurance, and ultimately using them in the psychophysical experiments. This entire process must be performed, and experiments must be deployed, codependently, on the calibrated rendered and photographed images, as well as on the fully calibrated display system as described in earlier sections of this chapter.

In summary, a recalibration procedure was used to convert the 32-bit HDR images to 8-bit JPEG images in a way that guaranteed the luminance values of the displayed images were the same as the luminance values in the physical world. The steps used to perform this procedure are as follows:

4.2.2.1. Recalibration Step 1. 8-bit JPEG images for all four HDR image types (described next) were formed (using the specialized Matlab/Photoshop/Javascript function shown in Table 4.2) with a default calibration factor (CF) equal to one. The four image types were: rendered overhead, rendered window, photographed overhead, and photographed window. (The correct CF will then be computed as described in Section 4.2.2.2, and then the correct CF is used in Section 4.2.2.3, to make a final adjustment on the images.)

4.2.2.2. Recalibration Step 2. The four output images were displayed on the calibrated screen, and a Minolta CS-100 Colorimeter was used to measure the nine checkerboard spots on each displayed output image. The luminance values of the checkerboard patterns were divided into the luminance values of the checkerboard in the physical environment (for each of the two illumination types: overhead and window). The ratios were computed by the formula shown next—to minimize luminance deviations and ultimately ensure the luminance values on screen were exactly the same as the physically measured luminance values. The CF technique was derived from Jacobs HDR calibration technique described here (Jacobs, 2011a), and the CF formula is:

$$CF = \text{Luminance}_{\text{Real}} / \text{Luminance}_{\text{HDR}} \quad \text{Equation 4.2}$$

4.2.2.3. Recalibration Step 3. A second set of 8-bit images was formed using the same script as before (Table 4.2), but this time with calibration factors obtained from 4.2.2.2 for each image type. Once the CF was applied to the (now recalibrated) images, they were again displayed on the calibrated display.

4.2.2.4. Recalibration Step 4. The Minolta CS-100 Colorimeter was used one more time to ensure that the adjusted luminance values on the calibrated screen were the same as the luminance values in the physical environment. The desktop outputs were, in fact, very close to the physical environment as shown by Figure 4.1, showing the graphical (and numerical) luminance values for each of the four image types. In each of the corresponding subplots of Figure 4.1, the nine values obtained from the checkerboard stimulus were plotted, a line of equality was drawn, and the domain and range of the plot was set to the dynamic range of the calibrated display. Furthermore, the coefficients of determination were computed (as shown in the subplots), regression lines were estimated, and intercept/slope for each regression analysis is shown in equation form (see subplots of Figure 4.1). If everything lies on the line of equality, then we would have perfect ratio reproduction, and perfect absolute reproduction—of the tested luminance values—thereby indicating perfect correspondences for luminances and contrast values. The statistics tell us how far from perfect the reproductions actually were. As can be seen from Figure 4.1, there is very little deviation from physical measurements—indicating very good fidelity for luminance and contrast reproduction. In summary, the estimated coefficients of determination were all above .996, and the estimated coefficients of regression were close to zero for intercepts and close to one for slopes (See Figure 4.1 for exact values). These values are the mark of high fidelity in recreating photometric accuracy of the physical stimuli on our computer-displayed images.

Besides the statistical analysis described in the last paragraph (and shown in Figure 4.1), the *recalibration* step and photometric validation is essentially finished. The rest of

this subsection (i.e., Sections 4.2.2.4.1—4.2.2.4.3, next) includes notes for additional clarification, and additional notes on the scripting techniques that were used.

4.2.2.4.1. Note on Coefficient Analysis. Whereas all previous image processing was performed to get the coefficients of determination between image matrices and physical reality close to unity (i.e., $R^2 = 1$), the recalibration process ensured that the actual ratio of luminance values were close to $y = x$ (i.e. the regression coefficients between display and physical objects close to $\beta_0 = 0$ and $\beta_1 = 1$). In other words, the image formation steps discussed in Section 4.1 and Section 4.2.1 were made to ensure high correlations, but the *recalibration* step in this section (i.e., Section 4.2.2) ensured a one-to-one luminance and contrast correspondence with the physical-objects stimuli.

Recall that coefficients here should not be confused with the analysis used in human psychophysics experiments. It is also worth noting that instead of applying several calibration steps along the way from HDR to display output, we were able to perform a single recalibration (using a single Calibration Factor, or CF) at the very last conversion step from HDR to JPEG. Performing a single CF reduced the number of recalibration steps to the minimum—thereby removing any and all unnecessary calibration steps along the procedure.

4.2.2.4.2. Note on Tone Mapping. The details of the *recalibration* script (Table 4.2) are important and worthy of further discussion. The algorithm used to ensure proper 32-bit to 8-bit conversion is available in Photoshop. It is under the realm of *tone mapping* algorithms, but as noted earlier, *tone mapping* is a loaded term. *Tone mapping* is often associated with nonlinear transformations of HDR images to compress the dynamic range of the image so that it can be viewed on a

display possessing lower dynamic range capabilities (e.g., the image in Figure 4.3 is tone mapped to be viewed in this dissertation). In other words, *tone mapping* is for *viewability*. However, we are more interested in accurate *visibility* as opposed to *viewability*. For the purposes of *visibility*, any nonlinear *tone mapping* approach would be fatally wrong. Rather than perform nonlinear *tone mapping*, a linear transform was applied using the Photoshop commands described above. Likewise, a point and click approach to the script shown in Table 4.2 can be performed with the following mouse clicks in Photoshop CS5.1 Extended (Mac version):

Image->Mode->8 Bits/Chanel -> Method -> Exposure and Gamma ->
Exposure 0.00 -> Gamma 1.00.

Additionally, the Exposure value in this point-and-click command sequence can be changed from 0.00 to another number, effectively adjusting the Calibration Factor (CF) as described earlier. Gamma must remain untouched, as it would nonlinearize the outputs with respect to the inputs (In fact, the monitor's gamma response function is handled by the operating system and by the monitor calibration device, and need not be manually adjusted).

4.2.2.4.3. Note on Matlab Scripting with Photoshop Javascript. The process used in this research (i.e., the script shown in Table 4.2) was written and performed using the Photoshop Javascript feature (*Adobe Photoshop CS5 Scripting Guide*, 2010) available in Photoshop Extended. Moreover, Photoshop Extended uses a Photoshop Matlab toolbox (*Adobe Photoshop MATLAB Toolbox*, 2011), which interfaces Photoshop with Matlab, and gives Matlab the ability to manipulate images in Photoshop (and vice versa). For simplicity, only

the essential commands are shown in the example code shown in Table 4.2 (the caption also explains some of the essential details of the script).

4.2.2.5. Recalibration Step 5. Finally, the four estimated *recalibration* factor values were applied to the remainder of the four HDR image types (rendered overhead, rendered window, photographed overhead, and photographed window), and all corresponding stimulus images were filtered using the appropriately obtained CF. Note that during HDR image formation, the exposure settings for each of the four HDR image types was constant, resulting in each image type possessing the same luminance values across the pixel plane (except, of course, for the differences accounted for by camera location, stimulus objects, and interreflections within the scene). Prepared images of each image type were randomly chosen and photometrically checked to ensure the reliability of exposure constancy within each image type. This process resulted in a complete set of 8-bit recalibrated images for all preview (i.e. non-spatially filtered) images. The same calibration factors were also used for adjusting spatially blurred images, as described later in Section 4.3.

In summary, we first created JPEG images with $CF = 1$. Then we displayed the images on the calibrated screen and took photometric samples of areas with known luminance using the Minolta CS-100 Colorimeter. We then divided the screen-measured luminance values into the actual physical luminance values that were measured previously using the same Minolta CS-100 Colorimeter. The resulting quotient is the correct CF (calibration factor) for *recalibrating* the image for the display. The *recalibrated* images were then verified to be linear and nearly exact within the dynamic

range limits of the display (as demonstrated by the graphical and statistical results shown in Figure 4.1).

This method of 8-bit conversion uses the full capacity of the display system. Because all filtering occurs in HDR, any quantization artifacts would be equivalent to the artifacts created by the range and quantization characteristics of the 8-bit display system. In other words, artifacts are limited to the statistical variance shown in the analysis in Figure 4.1, and the capabilities of the display system—particularly in how the system handles and reproduces 8-bit images.

4.3. Spatial Image Filtering Techniques

After the calibration techniques described above, various filters were applied according to the requirements of each experiment, as discussed in the Methods section of Chapter 3. The details of the two spatial filters reported here include the Gaussian blur and the digital Bangerter blur filters. The flow diagrams for each experiment are shown in Figure 4.2, and the details for each filter, respectively, are described here.

4.3.1. Gaussian Blur Filter. The Gaussian blur function (Chapter 3, Experiment 1) was applied to rendered images only. No photographs were used in Experiment 1. The filter was an attempt at simplicity, and was therefore only applied to recalibrated JPEG images (in other words, the filter was not applied in high dynamic range, or HDR). The reason for this was to try the simplest possible filtering design that could be performed with minimal computing requirements (including software, hardware, memory, and system complexity). We essentially tested whether the application of a simple filtering

algorithm to common JPEG images would result in visibility measures similar to those found in the physical experiment.

The Gaussian blur filter was written within the data collecting Matlab script, and filtering was performed during Experiment 1. In other words, spatial filtering occurred in real-time during the presentation of stimuli to subjects. The original 78 (rendered only) recalibrated JPEG images were read in by the Matlab experimental script (not shown), and filtered using the function shown in Table 4.3. Filtering was performed in the Fourier domain using next power of two square padded gray images, in order to reduce edge-fading and bleed-over between the edges of the images (details are within the code). The script then cropped the image back to its original size and the output was displayed to the subject. Note that the mean luminance of each output images was always equal to the mean luminance of each input image. We reasoned that mean luminances which may be absorbed by a real filter could be ignored, although in previous research we estimated that the physical Bangerter occlusion foils transmit about half of the available light (Legge et al., 2010, p. 5). Note that an attenuation step can be added here (by multiplying the output image by a constant between 0 and 1, in order to reduce the luminance values of each pixel by a linear and predetermined attenuation factor. For example, multiplying the output image at this step by 0.5 would result in each pixel having exactly half the luminance of the non-attenuated pixel.

In summary, for Experiment 1 (Chapter 3), Gaussian blur filtering was performed on JPEG images.

```

function L = filt_img(I,sig)
    sf = 4;
    I = imresize(I, 1/sf);
    J = ones(2^nextpow2(max(size(I))),...
            2^nextpow2(max(size(I))),3).*mean(I(:));
    edgeLoc = fix([size(J,1)/2-size(I,1)/2+1 ...
                  size(J,1)/2+size(I,1)/2 ...
                  size(J,2)/2-size(I,2)/2+1 ...
                  size(J,2)/2+size(I,2)/2]);
    J(edgeLoc(1):edgeLoc(2),edgeLoc(3):edgeLoc(4),:) = I;
    fs = fspecial('gaussian',[size(J,1) size(J,2)],sig/sf);
    K = zeros(size(J));
    for i2 = 1:3
        K(:,:,i2) = real( fftshift( ifft2( fft2(J(:,:,i2)).*...
            fft2(fs,size(J,1),size(J,2)) ) ) );
    end
    L = K(edgeLoc(1):edgeLoc(2),edgeLoc(3):edgeLoc(4),:);
    L = imresize(L,sf);
end

```

Table 4.3. *Gaussian Blur Filter Function.* This function was used on JPEG images during the Gaussian blur experiment.

4.3.2. Digital Bangerter Blur Filter. The digital Bangerter filter (Experiments 2, 3, & 4, in Chapter 3) was created using HDR photography of a Bangerter occlusion foil lens. The blur lens used for the physical objects experiment had two layers of Bangerter occlusion foils (Fresnel Lens Company part number <0.1), with one foil mounted to each side of a circular cut piece of 1/8-inch clear acrylic. The lens was placed 20cm in front of the calibrated camera, through the use of threaded extension rings, in order to match the distance between the foil and eye of a subject wearing the physical goggles. In the center of a 3-foot by 5-foot foam-core board, a pinhole was made with a sewing needle. Behind the pinhole, a bright white LED (RadioShack part number 276-0017) was installed and illuminated. The foam-core apparatus with pinhole and LED is essentially a physical manifestation of a 2D spatial delta signal. The camera was positioned in front of the

foam-core board so the pinhole was perpendicularly centered on the camera's imaging plane, and the majority of the black foam-core surface covered the camera's view. The camera was centered and focused on the pinhole, and the double Bangerter blur foil was fixed to the front of the camera. Given the camera settings, the projected image of the pinhole was smaller than the point spread effects of the anti-aliasing filter built into the camera, thereby ensuring the purest possible PSF image (with the camera's optics being the limiting quality factor). The optical system was of high enough quality that lens flare and other optical aberrations were minimal and negligible, as determined by a separate HDR image analysis of the 2D spatial delta signal without Bangerter filtering (not shown). The Bangerter blur filter image was then formed using the HDR camera imaging method described in Section 4.1.2, in total darkness, with the exception of the illuminated LED. Figure 4.4 shows a viewable (i.e., *tone mapped*) JPEG image of the resulting HDR.




Figure 4.4. Digital Bangerter PSF. Bangerter PSF Images (Tone Mapped and False Color, Experiments 2, 3, & 4). The *tone mapped* gray scale image (left) is a Radiance approximation for how the blur foil appears when viewing a pinpoint of white light surrounded by darkness. The false color image (right) maps color to luminance in cd/m^2 (or nits); the false color image is in log 2 scale with a cutoff of 100 cd/m^2 . This was

adjusted to be viewable on print or a computer monitor. In other words, *tone mapping* and log scaling is for *viewability* purposes on this document.

The HDR Bangerter blur foil image shown in Figure 4.4 is the point spread function (PSF) of the 2D spatial delta signal imaged through the Bangerter blur foils. Using camera images obtained and the HDR image formation method described in Section 4.1.2, the Bangerter HDR PSF was created with exposure-bracketed images as follows:

```
>> hdrgen -o bangerter.hdr -r camera.rsp bangerter/*.JPG
>> pvalue -o -b bangerter.hdr > luminance.dat
>> pvalue -o -b -r luminance.dat > BangerterLuminance.hdr
```

The first command, *hdrgen*, created the HDR image in 32-bit RGB color (recall this technique was used for generating images in Section 4.1.2). The camera response function obtained earlier was also used as before to ensure photometric linearity with the digital filter. The remaining *pvalue* commands removed the RGB components and saved an accurate luminance map of the Bangerter blur PSF. In other words, the chroma (i.e., hue and saturation) values were removed, leaving behind a luminance-accurate gray-scale HDR image.

The output luminance-only image was then fed into the HDR filtering script shown in Table 4.4. In summary, a Radiance format 32-bit HDR image of the Bangerter blur filter was created by physical convolution with a physical 2D spatial delta signal—i.e., the PSF of the blur apparatus. The Fourier transform of that image is the OTF (optical transfer function) of the blur apparatus. A *tone mapped* PSF of the filter is shown in Figure 4.4

```

%% Main function
F = hdrread('BangerterLuminance.hdr'); % Digital Bangerter filter
F(:, :, 2:3) = []; % remove redundant gray scale data
FPad = fftprepLumFiltZeros(F); % pad filter with zeros
I = hdrread('Image.hdr'); % Unfiltered HDR Input image
[IPad edgeLoc] = fftprep(I); % pad image with mean gray
JPad = single(ones(size(IPad))); % allocate space for speed
for i = 1:3 % filter each rgb channel
    JPad(:, :, i) = real( fftshift( ifft2( fft2(IPad(:, :, i)).*...
        fft2(FPad) ) ) );
end
J = fftcrop(JPad, edgeLoc); % crop padded image
K = J ./ (sum(J(:)) / sum(I(:)) ); % rescale to original luminance
hdrwrite(K, 'Filtered.hdr'); % write image

%% called from Main function
% pad filter image for fft filtering
function [J, edgeLoc] = fftprepLumFiltZeros(I)
    J = zeros(2^nextpow2(max(size(I))), ...
        2^nextpow2(max(size(I))));
    edgeLoc = fix([size(J,1)/2-size(I,1)/2+1 ...
        size(J,1)/2+size(I,1)/2 ...
        size(J,2)/2-size(I,2)/2+1 ...
        size(J,2)/2+size(I,2)/2]);
    J(edgeLoc(1):edgeLoc(2), edgeLoc(3):edgeLoc(4), :) = I;
end
% pad input image for fft filtering
function [J, edgeLoc] = fftprep(I)
    J = ones(2^nextpow2(max(size(I))), ...
        2^nextpow2(max(size(I))), 3).*mean(I(:));
    edgeLoc = fix([size(J,1)/2-size(I,1)/2+1 ...
        size(J,1)/2+size(I,1)/2 ...
        size(J,2)/2-size(I,2)/2+1 ...
        size(J,2)/2+size(I,2)/2]);
    J(edgeLoc(1):edgeLoc(2), edgeLoc(3):edgeLoc(4), :) = I;
end
% post fft filtering image crop
function J = fftcrop(I, edgeLoc)
    J = I(edgeLoc(1):edgeLoc(2), edgeLoc(3):edgeLoc(4), :);
end

```

Table 4.4. *Digital Bangerter Blur Function*. The code snippets here include the filters that were used to create HDR digital Bangerter blurred images from HDR renderings and photographs. The main script reads in the HDR digital Bangerter filter, and a rendered or photographed HDR scene image. Reading, padding, Fourier filtering, cropping (removing the padding), and image exporting functions are shown.

Note that a 2D OTF (optical transfer function) could have been reported instead of the PSF, but since it includes the MTF (modulation transfer function) and the PTF (phase

transfer function), there is no benefit to presenting the OTF numerically or graphically. MTF functions have been reported in previous literature (Pérez et al., 2010), however we found MTF functions to be insufficient for robust optical characterization of our Bangerter blur goggles—in other words, we found high-frequency phase information to be important for digitally characterizing this filter. Since the OTF is the Fourier transform of the PSF, either is sufficient to completely characterize the filter, however the PSF is easier to report and the OTF is faster to perform mathematical operations with.

The first part of this section (Section 4.3.2) explained how we created the digital version of the Bangerter occlusion foil goggles, followed by how we filtered scene images using the newly formed digital Bangerter blur filter. The rest of this section includes a brief discussion on the methods chosen, and human performance, as it relates to the digital Bangerter blur filter.

Similar to the Gaussian Blur function, the mean luminance values of the output (i.e., digital Bangerter filtered) images were made to be equal to the mean luminance values of the input images (Table 4.4, Line 13). In other words, the input and output stimulus images had the same mean luminance. Furthermore, the filtering here was performed in HDR, instead of with 8-bit JPEG images (as was the case with the Gaussian blur model described in Section 4.3.1). We switched from LDR to HDR filtering in order to ensure accuracy across the luminance ranges encountered in physical scenes (i.e., to avoid aliasing in the luminance domain). The process is described in Figure 4.2 (lower panel), which includes the *recalibration* step as noted in Section 4.2.2.

We expected that the factor of two difference between mean luminance transmitted in the physical Bangerter goggles and the digital Bangerter model would have relatively little effect on performance. Ultimately, the computer desktop subjects performed better than the subjects who viewed the physical stimuli. The lower luminances seen by the subjects with in the physical experiment may have had some impact in depressing their performance compared to the digital experiments, and a luminance attenuation factor of 0.5 as described earlier in this chapter would be worthy of further exploration, although it was not pursued in the present study.

In summary, a Digital Bangerter blur filter was developed and applied to HDR images, and these images were converted to JPEGs using the method described by Section 4.2.2.

4.4. Non-Spatial Image Filtering Techniques

Finally, two non-spatial filtering techniques were used for adjusting the color balance (i.e. white balance) of the rendered images (Experiments 2, 3, & 4), and for the contrast reduction model used in Experiment 3. Both of these non-spatial filters are described in the following two subsections.

4.4.1. Color Balance (White Balance). The original Radiance renderings were prepared to be accurate according to the color profiles of the luminaires used in the physical experiment. In other words, the color temperatures of the overhead and artificial window luminaires were programmed into the Radiance model. The camera-based photographs were obtained by setting the appropriate white balance manually on the camera—i.e., by matching the camera's white balance to the color temperatures of the

overhead and artificial window fluorescent tubes (set in degrees Kelvin). Due to these Radiance and camera settings, the photographed images appeared color neutral on the calibrated display; however the rendered overhead images appeared too yellow (Kelvin warm), and the rendered artificial window images appeared too blue (Kelvin cool). In order to make all images appear properly color balanced during the psychophysical experiments, a color balancing function was developed (Table 4.5). The function used the checkerboard samples in photographed and rendered images to estimate a color correction factor for the overhead and window lighting conditions. It made small color adjustments to the rendered images in the HSV color space, so that the white balance of the rendered images matched the white balance of the photographed images. This process was incorporated into the HDR to JPEG recalibration algorithm described in Section 4.2.2. The color balancing function (Table 4.5) is called from the executive HDR-to-JPEG function, Table 4.2 (Line 9).

This approach for white balancing ensured that the luminance values of the pixels remained unchanged, because the adjustment was performed in HDR immediately before *recalibration* (i.e., before Step 4.2.2), and because the adjustment factors were computed in the HSV color space. This assumption was confirmed by hypothesis testing as described in Chapter 3, where the ANOVA found no difference between photographed and rendered images in terms of psychophysical *visibility* performance (see Chapter 3, Table 3.5—image type main effect and image type interactions). In summary, it was confirmed that the small correctional rotations of color in the HSV space did not affect luminance or contrast in any psychophysically measureable way. A description of the

HSV color space is noted in Palmer's introduction on color spaces (Palmer, 1999, sec.

3.1.2).

```
function J = ColorAdjust(I,n)
% color samples
D90Over = [246 226 209; 251 231 213; 253 234 215];
D90Wind = [224 225 224; 234 237 235; 241 245 244];
RadOver = [244 216 159; 254 233 172; 254 236 174];
RadWind = [197 217 209; 209 231 223; 219 241 232];
% reshape
D90Over = reshape(D90Over,1,3,3)./2^8;
D90Wind = reshape(D90Wind,1,3,3)./2^8;
RadOver = reshape(RadOver,1,3,3)./2^8;
RadWind = reshape(RadWind,1,3,3)./2^8;

% transform with hsv, v = 1
D90OverH = rgb2hsv(D90Over); D90OverH(:,:,3) = ones;
D90WindH = rgb2hsv(D90Wind); D90WindH(:,:,3) = ones;
RadOverH = rgb2hsv(RadOver); RadOverH(:,:,3) = ones;
RadWindH = rgb2hsv(RadWind); RadWindH(:,:,3) = ones;

TO2 = hsv2rgb(D90OverH)./hsv2rgb(RadOverH);
TW2 = hsv2rgb(D90WindH)./hsv2rgb(RadWindH);

TO3 = mean(TO2,2);
TW3 = mean(TW2,2);

if n == 'o'
    T3 = TO3;
elseif n == 'w'
    T3 = TW3;
end
for i = 1:3
    J(:,:,i) = I(:,:,i).*T3(:,:,i);
end
end
```

Table 4.5. *Color Balance (White Balance) Function*. The Color Adjustment function shown here was created to match the Radiance rendered color balance to those taken by the camera. Color balance matching allowed the rendered stimuli to appear color accurate on the image display. Without this color balance adjustment on the Radiance rendered images, they would have appeared too blue or too yellow to seem realistic. To avoid any subjective interpretations of the stimulus images, they were color corrected to match one another with this script. The RGB values (lines 3-6) were estimated by inspecting the checkerboard

RGB values using Photosphere (Ward Larson, n.d.-c). The RGB values were converted to HSV (lines 14-17), and the chroma (i.e., hue and saturation) values were ratio adjusted according to the details shown in the script (lines 19-23). See Palmer for theory on why using an HSV color balancing works without significantly affecting luminance intensities (Palmer, 1999, sec. 3.1.2). Furthermore, note that hypothesis testing in Chapter 3, Table 3.5 ensured that there were no *visibility* performance differences due to these small HSV transformations (i.e., rotations). See main text in Section 4.4.1 for possible improvements that could be made to this script.

Note, however, that an improvement over this method (i.e., instead of using Matlab's *rgb2hsv* function in Table 4.5) may be to use a different color space for temperature adjustments, such as Matlab's *rgb2ycbcr*, or another physiologically valid color space altogether. The author did not discover Matlab's *rgb2ycbcr* function until after data collection was completed, but it is likely that the YCbCr color space handles luminance more robustly than the HSV color space. Furthermore, it is likely that a shift along a more appropriate color space (e.g., a chromaticity or Kelvin-defined color space) could be used, which would not only allow necessary color shifts to maintain accurate luminance values, but would further improve the color fidelity of the digital images.

Nevertheless, recall that as indicated earlier, the coefficients of determination, and the regression coefficients, along with the graphical analysis shown and described in Figure 4.1, give compelling evidence for luminance and contrast equality between real and displayed images. Since this filter (i.e., the color balancing filter) and all other non-spatial filters were incorporated into the HDR-to-JPEG image filtering algorithm, the entire system was *recalibrated* in accordance with the accuracy and precision demonstrated in Figure 4.1. Ultimately, the psychophysical results (ANOVA, Chapter 3, Table 3.5) found

no differences in *visibility* performance between photographed and rendered images, which gives compelling evidence that the small white balance corrections used here did not affect the *visibility* of objects in rendered (as compared to photographed) images.

In summary, for Experiments 2-4, color balancing was applied to rendered HDR images, before JPEG conversion.

4.4.2. Contrast Adjustment. For the contrast reduction experiment (Experiment 3), an algorithm was developed to make contrast reduced HDR images. The Matlab function is shown in Table 4.6. The function ensured that the mean luminance of each output image was equal to the mean luminance of each input image. For a 100% contrast setting, the pixels of the output image are identical to the pixels of the input image. For a 0% contrast setting, all pixels of the output image are equal to the mean luminance (and mean chroma) value of the input image. The mathematical formula for contrast attenuation is shown here:

Given Input luminance L_1 , mean luminance μ , input image contrast C_1 , contrast reduction proportion ρ , output contrast C_1 , and output luminance L_2 , and given the following equalities:

$$\text{Input image pixel luminance values: } L_1 = L_{x,y} \quad \text{Equation 4.3}$$

$$\text{Mean luminance: } \mu = \frac{1}{XY} \sum_{x=1}^X \sum_{y=1}^Y L_{x,y} \quad \text{Equation 4.4}$$

$$\text{And input contrast function: } C_1 = C_{x,y} = \frac{L_{x,y} - \mu}{\mu} = \frac{L_1 - \mu}{\mu} \quad \text{Equation 4.5}$$

To find the output luminance L_2 :

Reduce contrast C_1 by proportion: $0 \leq \rho \leq 1$

$$C_2 = C_1 \cdot \rho = \frac{L_2 - \mu}{\mu} = \frac{L_1 - \mu}{\mu} \cdot \rho \quad \text{Equation 4.6}$$

The output luminance, by substitution, and simplified, is:

$$L_2 = \mu(C_1 \cdot \rho) + \mu = \rho(L_1 - \mu) + \mu \quad \text{Equation 4.7}$$

And finally, the limits are:

$$\lim_{\rho \rightarrow 0} L_2 = \mu, \text{ and } \lim_{\rho \rightarrow 1} L_2 = L_1 \quad \text{Equations 4.8 \& 4.9}$$

The formula in Equation 4.7 is written in the 6th line of the Matlab function shown in Table 4.6. As with the white balance adjustment, a possible improvement of this function would be to perform the contrast attenuation in the YCbCr color space.

```
function outrgb = ContrastAdjust(inrgb,rho)
    inhsv = rgb2hsv(inrgb);           % input hsv values
    muhsv = repmat(mean(mean(inhsv,1),2),...
        size(inhsv,1),size(inhsv,2)); % mean hsv vals
    pctmat = ones(size(inhsv)).*rho;  % percent matrix
    outhsv = pctmat.*(inhsv-muhsv)+muhsv; % output hsv values
    outrgb = hsv2rgb(outhsv);        % output rgb values
end
```

Table 4.6. *Contrast Adjustment Function*. The contrast adjustment function shown here allows contrast adjustment from 0 to 100% contrast, with the mean output image luminance remaining equal to the mean input image luminance.

The formula described in the Equation 4.7 and the Matlab function (Table 4.6, Line 6) compresses the values of the contrast function by factor ρ (rho). As noted in the results of Experiment 3, our best fit (i.e., the model that best explained performance in the physical experiment) occurred when $\rho = 0.8$. In other words, the best fit occurred when ρ was set to 0.8. Finally, the filtering function (shown in Table 4.6) was performed after the

digital filtering operation, as shown in the executive script in Table 4.2, Line 10. The conversion steps are shown in the bottom panel of Figure 4.2.

In summary, for Experiment 3, contrast reduction was applied to HDR Bangerter filtered images, before JPEG conversion.

Bibliography

- Adobe Photoshop CS5 Scripting Guide*. (2010). Adobe Systems Incorporated.
Retrieved from www.adobe.com/devnet/photoshop/scripting.html
- Adobe Photoshop MATLAB Toolbox*. (2011). Adobe Systems Incorporated. Retrieved from http://help.adobe.com/en_US/photoshop/cs/using/WS298C49A0-AC53-43c5-B19D-554CE1D87C3Ea.html
- Alexander, N. B., Ashton-Miller, J. A., Giordani, B., Guire, K., & Schultz, A. B. (2005). Age Differences in Timed Accurate Stepping With Increasing Cognitive and Visual Demand: A Walking Trail Making Test. *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences*, *60*(12), 1558 -1562.
doi:10.1093/gerona/60.12.1558
- Arditi, A., & Brabyn, J. (2000). Signage and Wayfinding. In B. Silverstone, M. A. Lang, B. Rosenthal, & E. E. Faye (Eds.), *Lighthouse Handbook on Vision Impairment and Vision Rehabilitation*. New York: Oxford University Press.
- Atick, J. J., & Redlich, A. N. (1992). What Does the Retina Know about Natural Scenes? *Neural Computation*, *4*(2), 196-210. doi:10.1162/neco.1992.4.2.196
- Bochsler, T. M., Legge, G. E., Gage, R., & Kallie, C. S. (In Press). Recognition of Ramps and Steps by People with Low Vision. *Investigative Ophthalmology & Visual Science*.
- Bochsler, T. M., Legge, G. E., Kallie, C. S., & Gage, R. (2012). Seeing Steps and Ramps with Simulated Low Acuity: Impact of Texture and Locomotion. *Optometry and Vision Science*, *89*(9).

- Bordier, C., Petra, J., Dauxerre, C., Vital-Durand, F., & Knoblauch, K. (2011). Influence of Background on Image Recognition in Normal Vision and Age-Related Macular Degeneration. *Ophthalmic and Physiological Optics*, 31(3), 203-215. doi:10.1111/j.1475-1313.2011.00820.x
- Boucart, M., Desprez, P., Hladiuk, K., & Desmetre, T. (2008). Does Context or Color Improve Object Recognition in Patients with Low Vision? *Visual Neuroscience*, 25(5-6), 685-691. doi:10.1017/S0952523808080826
- Brabyn, J., Schneck, M., Haegerstrom-Portnoy, G., & Lott, L. (2004). Functional Vision: "Real World" Impairment Examples from the SKI Study. *Visual Impairment Research*, 6, 35-44. doi:10.1080/13882350490517644
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 433-436.
- Bramão, I., Reis, A., Petersson, K. M., & Faisca, L. (2011). The Role of Color Information on Object Recognition: A Review and Meta-Analysis. *Acta Psychologica, In Press, Corrected Proof*. doi:16/j.actpsy.2011.06.010
- Colenbrander, A. (2010). Assessment of functional vision and its rehabilitation. *Acta Ophthalmologica*, 88(2), 163-173. doi:10.1111/j.1755-3768.2009.01670.x
- DataColor Spyder 3 Utility. (2007). Integrated Color Solutions, Inc. Retrieved from spyder.datacolor.com
- DiLaura, D., Houser, K., Mistrick, R., & Steffy, G. (Eds.). (2011). *The Lighting Handbook: Reference & Application* (10th ed.). Illuminating Engineering Society of North America.

- Fuhr, P. S., Liu, L., & Kuyk, T. K. (2007). Relationships Between Feature Search and Mobility Performance in Persons With Severe Visual Impairment. *Optometry and Vision Science, 84*(5), 393-400. doi:10.1097/OPX.0b013e31804f5afb
- Goodrich, G. L., & Ludt, R. (2003). Assessing Visual Detection Ability for Mobility in Individuals with Low Vision. *Visual Impairment Research, 5*(2), 57-71. doi:10.1076/vimr.5.2.57.26265
- Hafeneger, S. (2009). *Sofortbild*. Retrieved from www.sofortbildap.com
- Hassan, S. E., Lovie-Kitchin, J. E., & Woods, R. L. (2002). Vision and mobility performance of subjects with age-related macular degeneration. *Optometry and Vision Science: Official Publication of the American Academy of Optometry, 79*(11), 697-707.
- Jacobs, A. (2011a). Camera Calibration. Retrieved from <http://www.jaloxa.eu/webhdr/calibrate.shtml>
- Jacobs, A. (2011b). WebHDR. Retrieved from <http://www.jaloxa.eu/webhdr/index.shtml>
- Kabanarou, S. A., & Rubin, G. S. (2006). Reading with central scotomas: is there a binocular gain? *Optometry and Vision Science: Official Publication of the American Academy of Optometry, 83*(11), 789-796. doi:10.1097/01.opx.0000238642.65218.64
- Klein, B. E. K., Klein, R., Lee, K. E., & Cruickshanks, K. J. (1998). Performance-Based and Self-Assessed Measures of Visual Function as Related to History of Falls,

- Hip Fractures, and Measured Gait Time: The Beaver Dam Eye Study. *Ophthalmology*, 105(1), 160-164. doi:16/S0161-6420(98)91911-X
- Kleiner, M., Brainard, D. H., & Pelli, D. G. (2007). What's new in Psychtoolbox-3? *Perception 36 ECVF Abstract Supplement*.
- Kuyk, T., Elliott, J. L., Biehl, J., & Fuhr, P. S. (1996). Environmental Variables and Mobility Performance in Adults with Low Vision. *Journal of the American Optometric Association*, 67(7), 403-409.
- Legge, G. E., Yu, D., Kallie, C. S., Bochsler, T. M., & Gage, R. (2010). Visual Accessibility of Ramps and Steps. *Journal of Vision*, 10(11), 8. doi:10.1167/10.11.8
- Ludt, R., & Goodrich, G. L. (2002). Change in Visual Perceptual Detection Distances for Low Vision Travelers as a Result of Dynamic Visual Assessment and Training. *Journal of Visual Impairment & Blindness*, 96(1), 7-21.
- MATLAB R2010a*. (2010). MathWorks. Retrieved from <http://www.mathworks.com/products/matlab/>
- MATLAB R2011a*. (2011). MathWorks. Retrieved from <http://www.mathworks.com/products/matlab/>
- Odell, N. V., Leske, D. A., Hatt, S. R., Adams, W. E., & Holmes, J. M. (2008). The Effect of Bangerter Filters on Optotype Acuity, Vernier Acuity, and Contrast Sensitivity. *Journal of American Association for Pediatric Ophthalmology and Strabismus*, 12(6), 555-559. doi:16/j.jaapos.2008.04.012
- Palmer, S. E. (1999). *Vision Science: Photons to Phenomenology*. MIT Press.

- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437-442.
- Pérez, G. M., Archer, S. M., & Artal, P. (2010). Optical Characterization of Bangerter Foils. *Investigative Ophthalmology & Visual Science*, 51(1), 609 -613.
doi:10.1167/iovs.09-3726
- Radiance Manual Pages. (1997). *RADIANCE comprehensive index*. Retrieved from http://radsite.lbl.gov/radiance/whatis_comp.html
- Sedgwick, A. (1986). Space Perception. In Boff, K R, Kaufman L, & Thomas, J P (Eds.), *Handbook of Perception and Human Performance* (pp. 21-1—21-57). John Wiley & Sons.
- The Eye Diseases Prevalence Research Group. (2004). Causes and Prevalence of Visual Impairment Among Adults in the United States. *Archives of Ophthalmology*, 122(4), 477-485. doi:10.1001/archopht.122.4.477
- The International classification of diseases, 9th revision, clinical modification : ICD-9-CM*. (1986). (Annotated.). Ann Arbor, Mich: Commission on Professional and Hospital Activities.
- Tran, T. H. C., Guyader, N., Guerin, A., Despretz, P., & Boucart, M. (2011). Figure Ground Discrimination in Age-Related Macular Degeneration. *Investigative Ophthalmology & Visual Science*, 52(3), 1655 -1660. doi:10.1167/iovs.10-6003
- Vale, A., Scally, A., Buckley, J. G., & Elliott, D. B. (2008). The Effects of Monocular Refractive Blur on Gait Parameters when Negotiating a Raised Surface.

Ophthalmic and Physiological Optics, 28(2), 135-142. doi:10.1111/j.1475-1313.2008.00543.x

Vitale, S., Cotch, M. F., & Sperduto, R. D. (2006). Prevalence of Visual Impairment in the United States. *JAMA: The Journal of the American Medical Association*, 295(18), 2158 -2163. doi:10.1001/jama.295.18.2158

Ward Larson, G. (n.d.-a). *hdrgen*. Retrieved from www.anywhere.com

Ward Larson, G. (n.d.-b). *Radiance*. Retrieved from <http://radsite.lbl.gov/radiance/>

Ward Larson, G. (n.d.-c). *Photosphere*. Retrieved from www.anywhere.com

Ward Larson, G. (n.d.-d). *Radiance File Formats*. Retrieved from radsite.lbl.gov/radiance/refer/filefmts.pdf

Ward Larson, G., & Shakespeare, R. (2011). *Rendering With Radiance: the art and science of lighting visualization*. Randolph M Fritz. Retrieved from <http://www.radiance-online.org/pipermail/radiance-general/2011-August/008080.html>

West, S. K., Rubin, G. S., Broman, A. T., Muñoz, B., Bandeen-Roche, K., & Turano, K. (2002). How does visual impairment affect performance on tasks of everyday life? The SEE Project. Salisbury Eye Evaluation. *Archives of Ophthalmology*, 120(6), 774-780.

Wurm, L. H., Legge, G. E., Isenberg, L. M., & Luebker, A. (1993). Color Improves Object Recognition in Normal and Low Vision. *Journal of Experimental Psychology. Human Perception and Performance*, 19(4), 899-911.

