

Sensitivity to luminance and chromaticity gradients in a complex scene

Alexa I. Ruppertsberg

Bradford Optometry Colour and Lighting Lab,
School of Life Sciences, University of Bradford,
Bradford, UK



Marina Bloj

Bradford Optometry Colour and Lighting Lab,
School of Life Sciences, University of Bradford,
Bradford, UK



Anya Hurlbert

Institute of Neuroscience, Newcastle University,
Newcastle Upon Tyne, UK



Image gradients—smooth changes in color and luminance—may be caused by intrinsic surface reflectance properties or extrinsic illumination phenomena, including shading, shadowing, and inter-reflections. In turn, image gradients may provide the visual system with information concerning the origin of these factors, such as the orientation of surfaces with respect to the light source. Color gradients induced by mutual illumination (MI) may play a similar role to that of luminance gradients in shape-from-shading algorithms; it has been shown that 3D shape perception modulates the influence of MI on surface color perception (M. G. Bloj, D. Kersten, & A. C. Hurlbert, 1999). In this study, we assess human sensitivity to changes in color and luminance gradients that arise from changes in the light source position, within a complex scene. In [Experiment 1](#), we tested whether observers were able to discriminate between gradients due to different light source positions. We found that observers reliably detected a change in the gradient information when the light source position differed by only 4 deg from the reference scene. This sensitivity was mainly based on the luminance information in the gradient ([Experiments 2 and 3](#)). Some observers make use of the spatial distribution of chromaticity and luminance values within gradients when discriminating between them ([Experiment 4](#)). The high sensitivity to gradient differences supports the notion that gradients contain information that may assist in the recovery of 3D shape and scene configuration properties.

Keywords: chromaticity, luminance, gradients, mutual illumination, inter-reflections, shading

Citation: Ruppertsberg, A. I., Bloj, M., & Hurlbert, A. (2008). Sensitivity to luminance and chromaticity gradients in a complex scene. *Journal of Vision*, 8(9):3, 1–16, <http://journalofvision.org/8/9/3/>, doi:10.1167/8.9.3.

Introduction

An important task of human vision is to extract depth information from the two-dimensional (2D) image of the world projected onto the retina. This 2D pattern of light is the result of scene illumination, surface reflectance properties of objects and their spatial configuration. Variations in illumination arise from three main phenomena: shading, shadowing, and inter-reflections, also called “mutual illuminations” (MI). Shading depends on the surface orientation with respect to a light source and on the relative distance between the surface and the source. Shadowing depends on the visibility of the light source from different surface points. Inter-reflections arise from light reflected between surfaces (see [Figure 1](#) for examples of these phenomena).

These illumination phenomena cause spatial variations in luminance and chromaticity within the scene, even where the surface reflectance is perfectly uniform. Therefore, we

label these sources of spatial variation “extrinsic” to the object and distinguish them from “intrinsic” factors such as inhomogeneities in surface reflectance. We use the term “gradient” for any smooth, continuous change in image chromaticity and/or luminance due to these extrinsic or intrinsic factors and the terms “color gradient” and “luminance gradient” for changes in either alone. Image gradients in turn provide information about the extrinsic and intrinsic factors from which they arise.

Luminance gradients have been intensely analyzed in computer vision, while color gradients have been relatively neglected as sources of quantitative information about the scene. Luminance gradients which arise from shading provide clues to the local and global orientation of surfaces with respect to the light source and other objects and are therefore crucial in enabling the visual system to recover 3D information. This notion is well known and forms the basis of shape-from-shading algorithms. These algorithms typically discard chromatic information and compute 3D surface curvature from

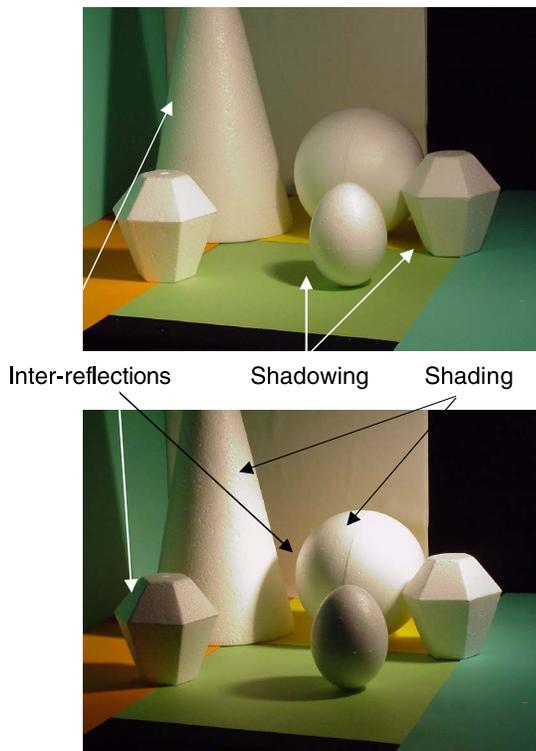


Figure 1. Examples of the illumination phenomena “shading,” “shading,” and “inter-reflections” illustrated by two photographs. All objects are white, positioned on a patterned floor and in front of a patterned background, and are illuminated by one light source. The egg and polyhedral object cause shadowing on the table surface because they are between the table surface and the light source. Shading is particularly visible on the cone and sphere, whose surface orientations change smoothly with respect to the light source. Examples of inter-reflections are visible on the left side of the cone, the nearby polyhedral object, and the sphere. The cone and polyhedral object receive light from the green wall and the sphere receives light from the illuminated side of the cone, which causes a patch of increased brightness on the far side of the sphere. The reader may find more examples. The only difference between the two scenes is the illumination direction, which is strongly signaled by cast shadows and smooth shading changes.

image intensity values under strong constraints, assuming, for example, direct illumination only, no shadows or inter-reflections, and an achromatic world of known albedo (Horn & Brooks, 1986; Ikeuchi & Horn, 1981; Pentland, 1984).

Extrinsic color gradients arise largely from shadows or inter-reflections, not shading. Most shape reconstruction algorithms explicitly consider either shadows or inter-reflections, not both, an artificial dissociation as shadows tend to occur where inter-reflections are maximal (Langer, 1999; Ruppertsberg & Bloj, 2007). Those few that consider both effects typically operate in a monochromatic world (Forsyth & Zisserman, 1990; Haddon & Forsyth,

1998). Other algorithms which consider color images first remove the effects due to inter-reflections (Funt & Drew, 1993; Funt, Drew, & Brockington, 1992), even though it has been demonstrated that both the luminance and chromatic components of inter-reflection-induced gradients can contribute substantially to image radiance (Forsyth & Zisserman, 1991; Langer, 2001) and may be used to recover estimates of surface reflectance (e.g., Funt, Drew, & Ho, 1991). Intrinsic color gradients, on the other hand, have received more attention in computer vision, particularly in color image segmentation (e.g., Hoang, Geusebroek, & Smeulders, 2005). Thus, while it is acknowledged that color gradients are often significant in natural images, their relationship to luminance gradients has been little explored (see Ben-Shahar and Zucker, 2004) and their potential use for information recovery has not been exploited.

In the study of human visual perception, color gradients have likewise been little studied. Although more is known about the perception of luminance gradients, the extent to which luminance gradients alone support shape discrimination is debated. Shading information on its own in a 2D image is ambiguous (Curran & Johnston, 1996) and additional sources of information, such as outline (Knill & Kersten, 1991), occluding boundaries (Ramachandran, 1988), surface texture (Curran & Johnston, 1994), cast shadows (Erens, Kappers, & Koenderink, 1993b), or prior assumptions about the probability of scene properties such as light source position (Adams, 2007; Sun & Perona, 1998) must be supplied to specify 3D shape. On the other hand, the information contained in “intrinsic” shadows (shadows formed by an object on itself) is used by the human visual system to recover the direction of illumination as well as object shape, while “extrinsic” shadows provide information about relative positions and orientations of objects (Knill, Mamassian, & Kersten, 1997).

While the ability to discriminate image gradients does not necessarily predict the ability to use the information they contain for other tasks such as shape discrimination, we argue that it is unlikely that the visual system relies on information in gradients which it cannot discriminate. Therefore, the sensitivity of the human visual system to gradients is relevant in considering their potential contribution to other perceptual tasks. For example, brightness discrimination ability does not predict but is superior to depth discrimination (Langer & Bülthoff, 2000). For luminance gradients, it has been found that both detection (Bijl, Koenderink, & Toet, 1989; McCann, Savoy, Hall, & Scarpetti, 1974) and the ability to determine gradient direction (Erens, Kappers, & Koenderink, 1993a) depend on within-image contrast and not on field size. Other studies (Curran & Johnston, 1996; Todd & Mingolla, 1983) suggest that small changes in gradients influence the human visual system’s interpretation of shape even when other sources provide parallel or conflicting information. In particular, perceived surface curvature depends on light source position and orientation with respect to the

surface, and these in turn determine the magnitude as well as the position of gradients on a fixed 3D surface.

Other perceptual tasks may rely on the use of image gradients. Both inter-reflections and extrinsic shadows have been shown to enable discrimination of contact between an object and the ground plane (Madison, Thompson, Kersten, Shirley, & Smits, 2001). In the chromatic Mach card (Bloj et al., 1999), the perceived color of one side is influenced by the perceived 3D shape of the card, which governs the interpretation of the image gradients induced by inter-reflections. It is not known whether the converse effect occurs, i.e., whether the perceived chromatic gradient influences the perceived shape. But because the magnitude and extent of inter-reflection-induced gradients vary with the 3D spatial configuration of the scene (Funt & Drew, 1993; Ruppertsberg, Hurlbert, & Bloj, 2007), it is possible that the information contained in these gradients enables recovery of 3D shape and scene configuration properties, as well as detection of the mutual illumination itself. To use this information, the human visual system must be able to detect and discriminate these gradients in their naturally occurring strength.

A preliminary study has looked at discrimination and detection thresholds for isolated color gradients (Bloj, Wolf, & Hurlbert, 2002). Much is also known about the discrimination of discrete luminance and color differences (see Pokorny & Smith, 1986 for a review). These studies, however, have used very simple and artificial stimuli. In this study, we assess human sensitivity to changes in color and luminance gradients that arise from changes in the light source position, within a single complex scene. The color gradient is due solely to inter-reflections, whereas the luminance gradient is due to combined sources. The changes in gradients correspond to a concrete physical manipulation of the environment, and therefore this is a first step to determining the role played by gradients in natural scene perception. To address the possibilities properly, we use radiosity techniques to generate images of physically accurate gradients in complex scenes (Ruppertsberg & Bloj, 2006). Figure 2a illustrates the scene configuration and the nature of the light source position manipulation: we rotate the light source in the vertical plane that includes the viewing axis (the “slant” of the light source direction).

In the first part of this paper, we present an analysis of the luminance and chromaticity information in gradients as a function of a wide range of light source positions and card opening angles. In the second part, we then focus on the sensitivity of human observers to the luminance and chromaticity information in gradients as a function of light source position at a fixed card opening angle. We use light source position changes as the best method of altering gradients without changing other key parameters in the scene. The specific questions we address are

1. Are observers able to discriminate between gradients caused by small changes in light source position?

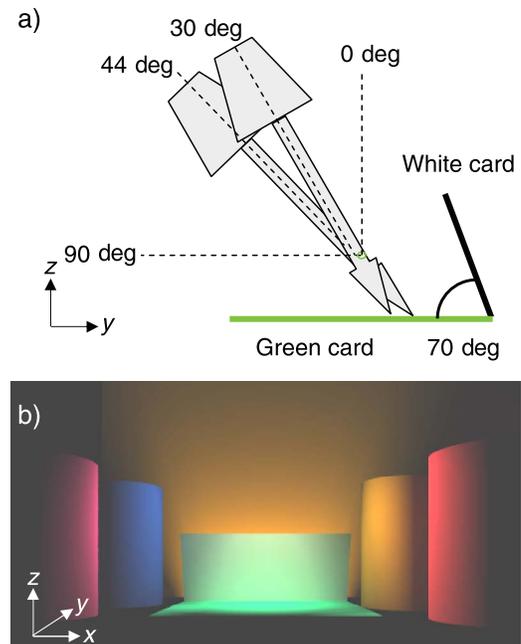


Figure 2. Setup of gradient scene. (a) Schematic side view. The card opening angle represented in the figure is the one used in the psychophysical experiments (70 deg). In the analyzing gradients section, this angle can take values between 50 and 90 deg, in 10 deg steps. The figure also indicates light source positions of 44 and 30 deg as well as the directions corresponding to light source positions of 0 and 90 deg. For the analyzing gradients section, the light source position varies between 20 and 90 deg, in 10 deg steps while for the psychophysical experiments it varies between 30 and 44 deg in 1 deg steps. (b) Scene rendered in RADIANCE looking at the white card; light source position 37 deg, card opening angle 70 deg. This is the reference scene corresponding to Experiment 1.

2. Are differences in chromatic and luminance gradients independently discriminable?
3. Do observers make use of the spatial distribution information when discriminating between gradients?

Methods

Gradients

The central object in our stimulus (described in detail below) is a “corner” formed from two cards, one white and one colored, joined in an acute to right-angled opening angle. The resultant gradient on the white card is a mixture of illumination effects. The color gradient is solely due to inter-reflections between the two hinged cards, whereas the luminance gradient is a combination of direct and indirect illumination caused by the orientation of the card with respect to the light source and by inter-reflections. Figure 2a shows a side view of a 70-deg-corner made

from a green and a white card illuminated by a tungsten light source. A light source position of 0 deg corresponds to direct illumination of the green card, whereas 90 deg corresponds to direct illumination of the white card. We chose a green card because hue discrimination studies showed good performance in this part of the visible spectrum (Wright & Pitt, 1934). Figure 2b shows an example of our experimental stimuli; a computer-rendered image of this corner centrally placed in a scene with additional objects on either side. The light source is at 37 deg. The gradient on the white card partly due to inter-reflected light from the green card is clearly visible. It has two components, a vertical and a horizontal one, the latter being a direct consequence of the beam size of the illuminating light source.

For reference, we use a Cartesian coordinate system centered on the observer with xz as the image plane and positive y pointing away from the observer (see Figure 2). All corners that we rendered consisted of a green card measuring 70 by 60 cm ($W \times H$) and a white card measuring 70 by 30 cm ($W \times H$). The green card lay 77 cm above the ground parallel to the xy -plane. To its far end, the white card was attached at a given card opening angle, making it frontoparallel to the observer, when the card opening angle was 90 deg. The card opening angle was varied between 50 and 90 deg. The front end of the green card was 1.34 m away from the observer. On either side of the corner additional objects enriched the scene. The tilt of the light source was kept constant at 0 deg (tilt is the angle between the z -axis and the projection of the light source vector L onto the xz -plane), and the slant was varied between 20 and 90 deg (slant is the angle between the projection of the light source vector L onto the yz -plane and the z -axis). We will refer to the slant as “light source position.” It is determined by rotating a 1-m-long vector around a point in space, which was 1.64 cm away from the observer (at the center of the green card) and at 1.07 m above the ground (at the top of the white card; see Figure 2a). The end of this vector corresponds to the light source position.

All scenes in this study were created using RADIANCE (Ward, 1994), a physical rendering software package that simulates the interaction of light and objects. We have previously shown that RADIANCE can accurately model the light-object interaction when a spectral rendering method is used (Ruppertsberg & Bloj, 2006). All surfaces were modeled to have Lambertian surface properties; all material colors were described by surface reflectance functions and the colors of the illuminant by spectral power distributions. The surface reflectance functions of colored materials and the modeled light source (a low voltage spotlight: Altman MR 16 Micro Ellipse, with a 75-W, 36 deg reflector) were based on measurements taken in our lab. We used a spectral rendering method with 81 wavebands from 380 to 780 nm, whose output is a hyperspectral image that can be converted to a three-layered XYZ image by applying color-matching functions

(Ruppertsberg & Bloj, 2008). We used five light bounces for all renderings as this is a sensible compromise between accuracy and speed. For simple geometrical scenes similar to the one used in this study, we have not found a significant difference between renderings with two and 50 bounces.

Stimuli for analysis of gradients

For the first part of the study, where we analyzed the luminance and chromaticity information in gradients, we varied the card angle between 50 and 90 deg and the light source position between 20 and 90 deg, both in 10 deg steps, yielding 40 different scenes. All other parameters and objects of the scenes were identical to the ones used in the scenes for the psychophysical experiment.

Stimuli for psychophysical experiments

For the psychophysical experiments, we fixed the card opening angle at 70 deg. The tilt of the light source was kept constant at 0 deg and its slant varied between 30 and 44 deg in 1 deg steps resulting in 15 different scenes. The scene with a light source position (slant) of 37 deg (“the 37 deg scene”) was defined as the reference scene and is shown in Figure 2b.

All calculated gradient values were within the gamut of the monitor and gamma-corrected. Because a pixel-by-pixel comparison of the computed values of the gradients and the actual displayed gradient on the monitor is technically impossible, we measured 1 deg samples with a spectroradiometer (PR650) at several points along the gradient and obtained a one-to-one correspondence with averages of computed values for the same area. Measured luminance values ranged from 19 to 57 cd/m^2 ; in the figures, we report calculated values.

Each image was 308 by 420 pixels ($H \times W$), corresponding to 3.2×6.3 deg of visual angle, and was presented on a black surround. The gradient extended horizontally over 2.1 deg at the top and 2 deg at the bottom and vertically over 0.9 deg of visual angle (see Figure 2b).

To prevent participants from using cast shadow information in the images, we pasted each gradient into the reference scene. Thus, each scene from 30 to 44 deg differed only in the gradient on the white card.

Adaptation stimulus, fixation cross and mask

The adaptation stimulus, the fixation cross and the mask had the same overall size as the scenes and each consisted of two regions: the central region, which corresponded to the location and spatial extent of the gradient in the scene, and the surround, which covered the

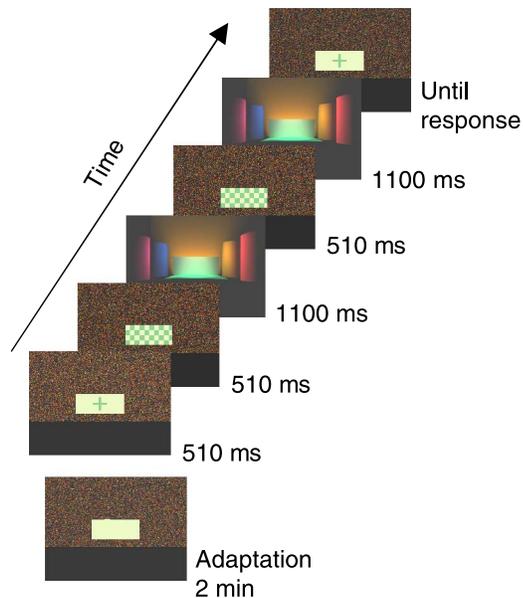


Figure 3. Presentation sequence.

rest of the scene (see Figure 3). The surround was the 1 by 1 pixel scramble of the reference scene's surround (37 deg scene). This was an attempt to keep the luminance and chromaticity as similar as possible to the stimulus scenes. The central region for the adaptation stimulus was a homogenous patch of green with the same mean luminance as the gradient of the reference scene. The central region of the fixation cross had the same homogenous patch as the adaptation stimulus with a darker fixation cross, and the central patch of the mask was a checkerboard made of the same colors used in the fixation cross stimulus.

Apparatus

Psychophysical experiments were run on a standard PC with a 24-bit graphics card. The stimulus presentation was controlled with Matlab (Mathworks®) and stimuli were presented on a 17-in CRT-monitor (NEC FE700+), which was calibrated with a spectroradiometer (Photo Research PR650, Glen Spectra Ltd., Middlesex, UK). The monitor was switched on for at least an hour before an experiment. Participants responded via the keyboard and were seated 114 cm from the screen with a head- and chin-rest in a dark experimental room.

Procedure

We used a roving 2-interval same-different design with the differencing rule, which does not require the observer to have knowledge of the stimulus range and decision is based purely on the difference between two observations

(Macmillan & Creelman, 2005). In each trial, two gradients were shown (one in each interval, see Figure 3 for timings and example images). In one of the intervals, the gradient corresponded to the reference scene, while the other gradient alternated (or roved) between the gradients corresponding to a light source position between 30 and 44 deg (in 1 deg steps). The order in which the reference and other gradients were displayed was randomized. We measured d' between the reference scene and all other scenes. Figure 3 shows the presentation sequence of a trial. A trial contained two intervals and each interval consisted of a fixation period (510 ms), a mask period (510 ms), and the stimulus with a gradient (1100 ms). Participants were instructed to compare the two gradients and decide whether they were the same or different. They responded by pressing a “same” or “different” labeled button on the keyboard and were instructed to give a correct answer, not a fast answer. One block of 120 trials assessed observers' sensitivity between the reference gradient and gradients with a light source orientation change of $+x$ and $-x$ deg ($1 \leq x \leq 7$), yielding two pairs of hit and false alarm rates from which two respective d' values were computed (Macmillan & Creelman, 2005, Table A5.4). Each block started with a 2-minute adaptation phase and an acoustic signal informed the participant of the immediate start of the experiment.

Participants observed for 30 minutes a day and data were collected over a six-week period. Although we present the results in different experiments, the order of sessions was randomized across experiments to minimize possible learning effects.

Participants

Four observers, three females (two naïves and one author) and one male (naïve) between 23 and 35 years old, with normal or corrected-to-normal acuity and functional color vision, assessed by the Farnsworth 100 Hue test participated in this study. They gave their signed consent prior to the experiments and naïve participants were paid for their time.

Results

Analyzing gradients: How does the luminance and chromaticity information in a gradient change as a function of light source position and card opening angle?

In this section, we describe the characteristics of our gradients in more detail and conclude with an outline of our psychophysical experiments.

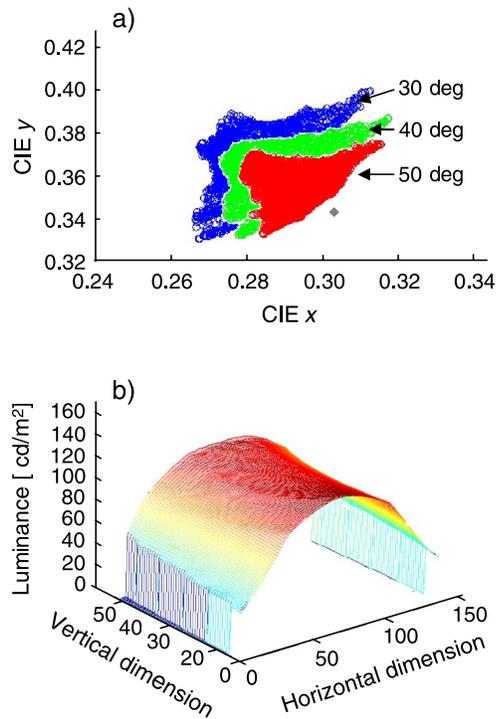


Figure 4. (a) Chromaticity values of all pixels on the white card (card opening angle 70 deg) for three different light source positions (30, 40, and 50 deg). The diamond indicates the color of the illuminant. (b) Spatial luminance distribution of all pixels on the white card (opening angle 70 deg) for a light source position of 50 deg. The horizontal dimension of the gradient is plotted along the x-axis and the vertical dimension of the gradient along the y-axis.

In Figure 4a, we show the distribution of CIE x and y chromaticity values of all pixels on the white card (opening angle 70 deg) for three different light source positions (30, 40, and 50 deg). Each distribution contains a broad spectrum of chromaticity values, which shifts systematically along a straight trajectory in the chromaticity plane as the light source position changes. For a light source position of 30 deg, a greater number of pixels in the gradient have greenish chromaticity values than for 40 and 50 deg, which contain more pixels with whitish chromaticity values. This difference is due to an increase in the strength of inter-reflections between the two cards as the light source rotates away from direct illumination of the white card. The diamond indicates the color of the illuminant (CIE $(x, y) = (0.3031, 0.3430)$). For the luminance values, it is possible to display information about the spatial distribution; Figure 2b shows the luminance distribution for a light source position of 50 deg.

The luminance values of all pixels in the white card are plotted so that those near the junction between the green and white card are at the back of the graph (near 50 on the vertical dimension axis) and those from the top edge of the white card are at the front of the graph (near 0 on the vertical dimension axis). We plotted the graph in this way

to avoid obscuring the lower luminance values. The plot for other light source angles looks similar in shape, but because the proportion of direct illumination on the white card increases with greater light source angles, all luminance values increase.

To assess the effects of light source position change in a more systematic manner, we approximated each gradient by its vertical profile. The vertical profile of the gradient was calculated by averaging horizontally over five central columns, running up from the bottom to the top of the white card. A profile therefore contains 55 values, with each value representing the mean of 5 pixels in a single row at the corresponding vertical position. Figures 5 and 6 show the chromaticity and luminance profiles for different card opening angles (50 to 90 deg) and light source positions (20 to 90 deg).

Note that light source position angles are with respect to the surface normal of the green card, whereas the card opening angles are with respect to the plane of the green card. From Figure 5, we can see that pixels on the white card are greener the more acute the card opening angle due to an increased capture of the inter-reflected light. Also the spread of chromaticity values is largest for acute card opening angles. As the card opening angle increases, the chromaticity spread for each light source position is not only shifted away from green toward white (color of the illuminant), but it is also compressed. Hence, scenes with right-angled objects will contain less color spread. Yet, different light source positions and card opening angles can lead to a very similar spread of chromaticity values. The luminance profiles (Figure 6) of more acute card opening angles tend to have lower luminance values. As the light source position increases, the luminance values increase because the amount of direct illumination on the white card increases. However, the slope of the profile changes too. For a light source position of 90 deg, i.e., when the white card is directly illuminated (Figure 2a), the bottom of the card receives less light than the top part. This is due to the radial fall-off of the light beam. For a card opening angle of 50 deg, the lower half of the white card is darker for light source positions of 20 and 30 deg; this is due to shadowing of the white card.

Outline of experiments

The variation in chromaticity and luminance values in these gradients caused by changes in the environment (light source position and card opening angle) is notable. If observers are able to use information in gradients to recover 3D shape and scene configuration properties, then discrimination between gradients induced by different environmental factors must be possible. We carried out a psychophysical study to determine whether human observers are actually able to distinguish between gradients as a function of light source position for a fixed card opening angle. For all experiments, we focused on

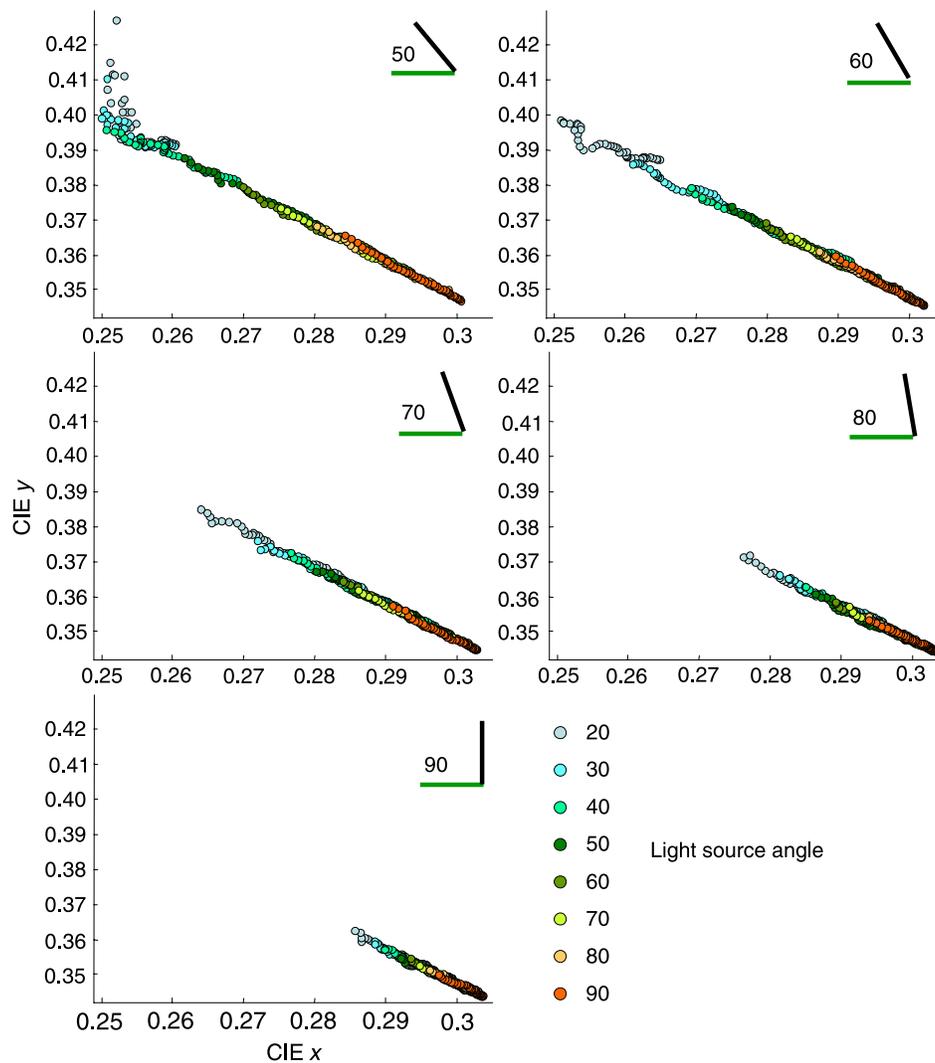


Figure 5. Chromaticity profiles of the white card for different card opening angles, 50 to 90 deg, indicated in each plot by the icons on the top right representing a schematic side view of the card. For a given card opening angle, the light source could take positions between 20 and 90 deg in 10 deg steps (see also Figure 2a).

one card opening angle (70 deg) and studied light source positions ranging between 30 and 44 deg.

Experiment 1 establishes the baseline sensitivity of human observers to changes in gradients. In **Experiments 2** and **3**, we explore the specific contributions of chromaticity and luminance information by presenting gradients differing only in their chromaticity or luminance distribution. In **Experiment 4**, we assess whether observers make use of the spatial distribution of the gradient.

Experiment 1: Sensitivity to luminance and chromaticity gradients in a simulated complex scene

Experiment 1 assessed how sensitive human observers are to changes in gradients, which were created by altering

the position of the light source. Observers were presented with two gradients that differed in their chromaticity and luminance distributions. Light source positions varied by $\pm 1, 2, 3, 4, 5, 6,$ and 7 deg from the reference scene.

Results

Figure 7 shows the d' values over different light source positions for all four observers from **Experiment 1**. As one would expect, sensitivity increases with larger differences between test and reference scenes. To determine the threshold for reliable discrimination in our task, we set the criterion to a percent correct rate of 80%, corresponding to $d' = 3$ in this experimental design (Macmillan & Creelman, 2005, Table 9.1). The double-arrows in the figure extend to the light source position for which this criterion was reached first and the numbers indicate the position difference in degrees to the reference scene.

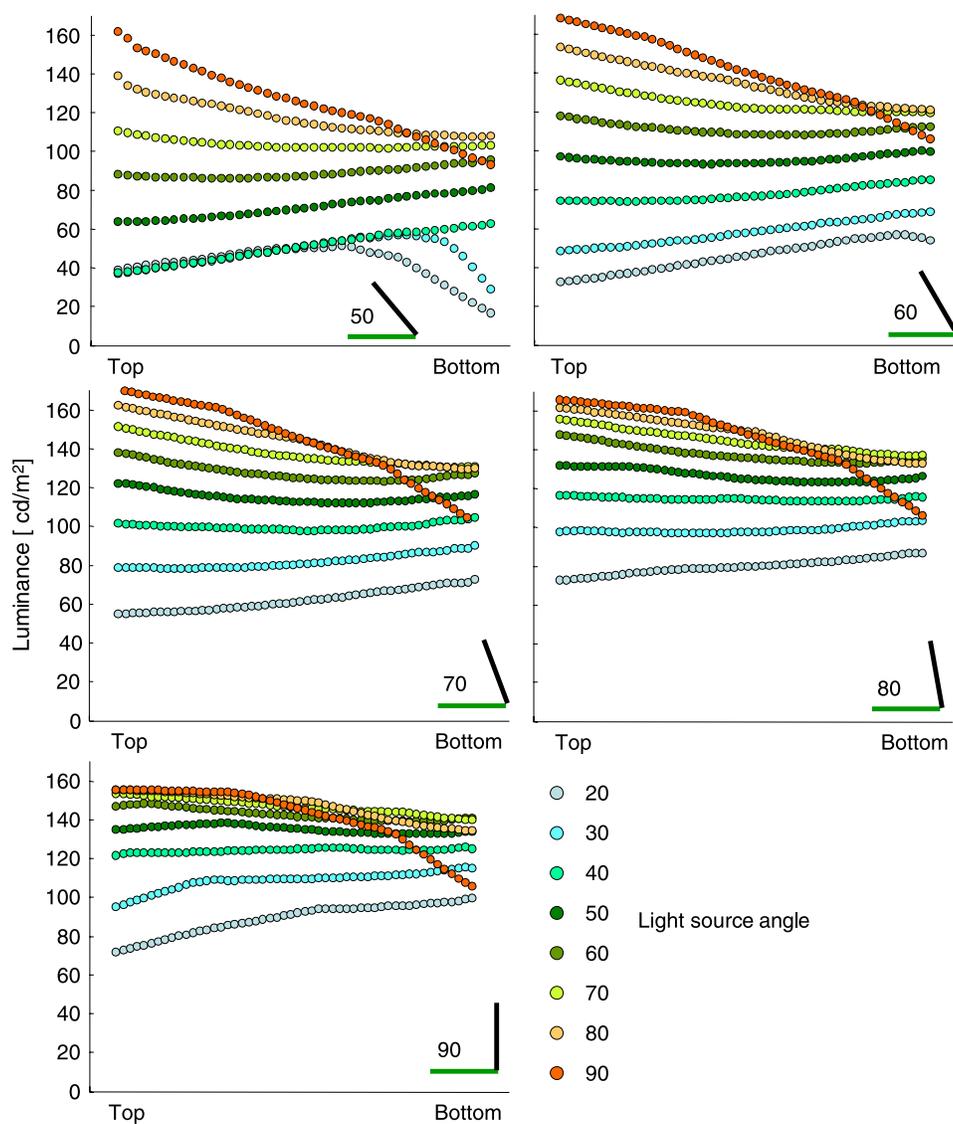


Figure 6. Luminance profiles of the white card for different card opening angles, 50 to 90 deg, indicated in each plot by the icons on the top right representing a schematic side view of the card. For a given card opening angle the light source could take positions between 20 and 90 deg in 10 deg steps (see also Figure 2a).

Across observers, the mean light source position difference that was reliably discriminated was +4 deg and -4.25 deg.

Experiment 2: Sensitivity to gradients differing only in their chromaticity

In Experiment 2, observers were presented with two gradients that differed only in their chromaticity distribution and contained the same luminance distribution as the reference scene. Because of the spectral rendering method used (Ruppertsberg & Bloj, 2008), we stored our stimuli as XYZ images. To produce a test stimulus that contained its own chromaticity distribution but the luminance distribution of the reference scene, we converted the

XYZ values of the test scene to chromaticity values xY_{test} and used the Y values of the reference scene: $xY_{\text{test}} Y_{\text{ref}}$. These values were then converted back to XYZ (for conversions, see Wyszecki & Stiles, 2000, p. 139) and further to RGB to be displayed on the monitor. We studied light source position changes of ± 5 , 6, and 7 deg.

Results

Figure 8 shows the d' values over different light source positions for all four observers from Experiment 2. The maximum d' values ranged from 1.3 to 3.5. The dotted line marks the criterion of performance ($d' = 3$), which observers consistently failed to reach. A d' value of less than 2 cannot be regarded as reliable performance as this corresponds to less than 67% correct (Macmillan &

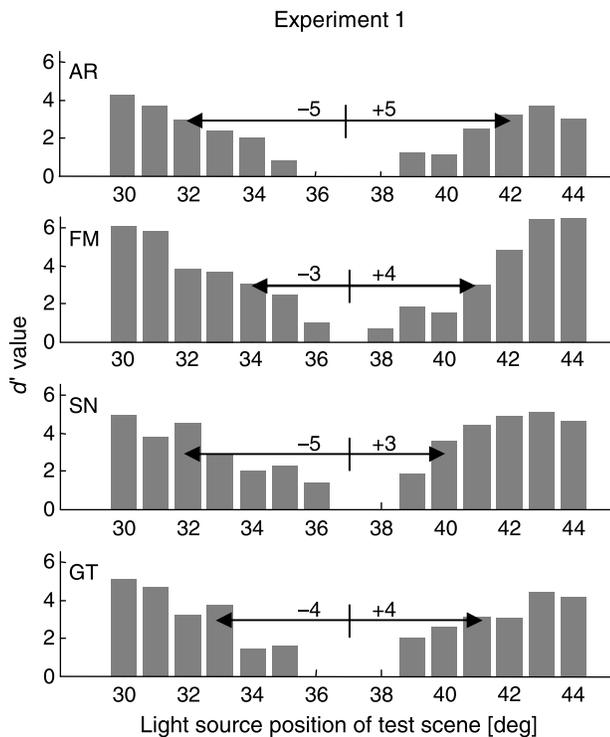


Figure 7. Results from Experiment 1. The test stimuli varied in their chromaticity and luminance distribution from the reference scene. The arrows indicate a performance level of $d' = 3$. Where x-axis ticks are visible, the d' value is zero.

Creelman, 2005, Table 9.1). Observers FM and GT achieved d' values close to 3 only for the test scene that had the largest difference in light source position from the reference scene (-7 deg).

Experiment 3: Sensitivity to gradients differing only in their luminance

In Experiment 3, observers were presented with two gradients that differed only in their luminance distributions. Both gradients contained the chromaticity distribution of the reference scene. To produce a test stimulus that contained its own luminance distribution but the chromaticity distribution of the reference scene, we converted the XYZ values of the reference scene to chromaticity values $x_{y_{ref}}$ and used the Y values of the test scene: $x_{y_{ref}} Y_{test}$. These values were then converted back to XYZ and further to RGB to be displayed on the monitor. We studied light source position changes of $\pm 1, 2, 3, 4, 5, 6,$ and 7 deg.

Results

Figure 9 shows the d' values over different light source positions for all four observers from Experiment 3. Across observers, the mean light source position difference at

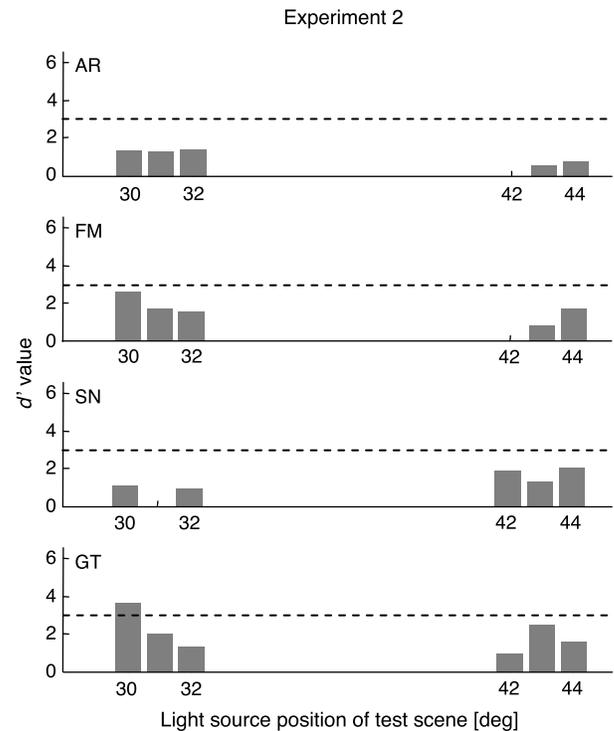


Figure 8. Results from Experiment 2. The test stimuli varied only in their chromaticity distribution from the reference scene. The dashed line indicates the performance criterion $d' = 3$. Where x-axis ticks are visible, the d' value is zero.

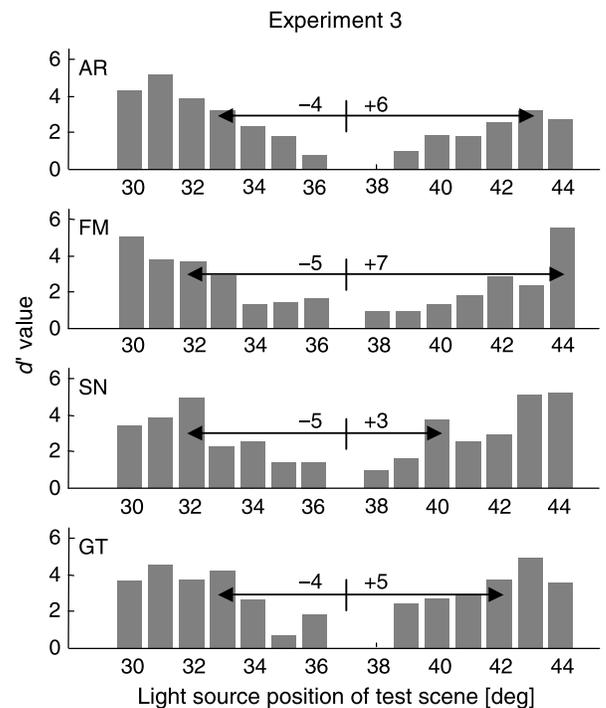


Figure 9. Results from Experiment 3. The test stimuli varied only in their luminance distribution from the reference scene. The arrows indicate a performance level of $d' = 3$. Where x-axis ticks are visible, the d' value is zero.

Observer	AR	FM	GT	SN
$t(13)$	-1.254	3.545	-0.579	1.209
p	0.232	0.004	0.573	0.248

Table 1. Results of the paired t test comparing performance across Experiment 1 and 3 for each observer (two-tailed).

Observer	AR	FM	GT	SN
$t(13)$	-1.758	3.344	1.174	3.97
p	0.102	0.005	0.262	0.002

Table 2. Results of the paired t test comparing performance across Experiment 1 and 4a for each observer (two-tailed).

which observers achieved a d' value of 3 was +5.25 deg and -4.5 deg.

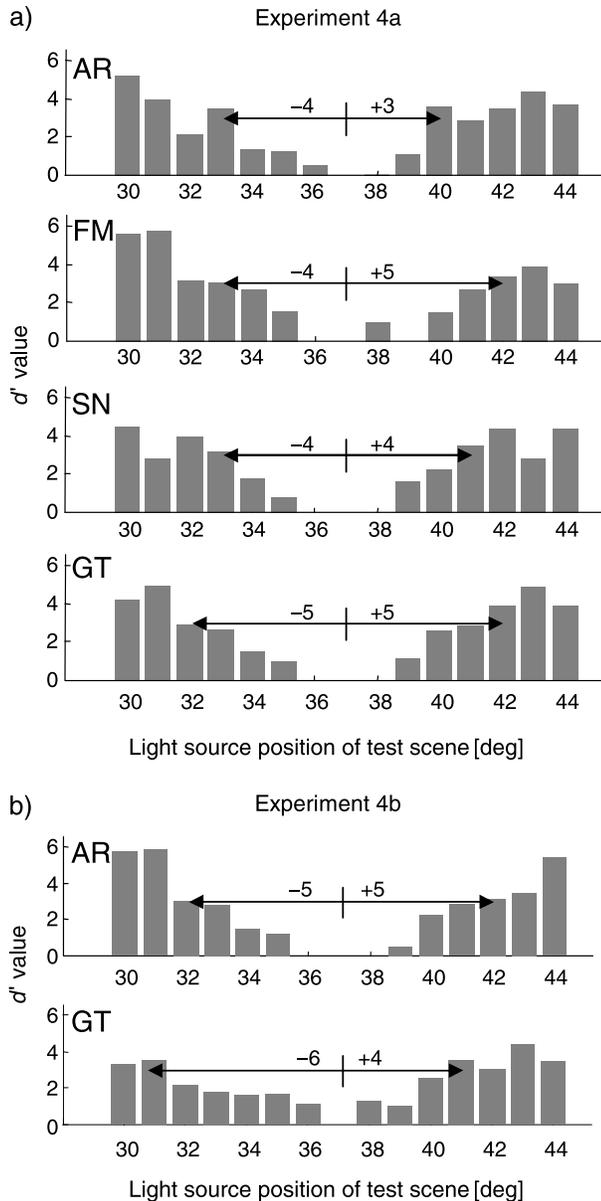


Figure 10. Results from Experiment 4a and Experiment 4b. (a) Test and reference gradients were identical to those in Experiment 1, but pixel scrambled. (b) This was a control experiment for two observers. Test and reference scenes contained a homogenous patch with the average chromaticity and luminance of the original gradients from Experiment 1.

Comparisons across Experiments 1, 2, and 3

The poor performance of observers in Experiment 2 indicates that behavior in Experiment 1 was not based on chromaticity but rather on luminance information. This was also suggested by the similarities in performance for Experiments 1 and 3. A paired t test (two-tailed) for each observer showed that there was no significant difference between the results of Experiment 1 and 3 for three of the four observers (Table 1). For observer FM, the difference reached significance with a lower performance in Experiment 3 than in Experiment 1. This significance might be due to FM's extremely high mean performance in Experiment 1 (Table 4). Because FM showed some indication of using chromatic information in Experiment 2, it is possible that the gradients in Experiment 3 were impoverished for FM and hence performance dropped. GT's performance in Experiment 2 was similar to FM, but GT had a lower mean performance in Experiment 1 than FM; hence, no significant difference was found between Experiment 1 and 3 for GT.

Gradients do not consist merely of numerous different luminance and chromaticity values but these values also have a particular spatial distribution. In the next two experiments, we address the question whether observers make use of the spatial distribution when distinguishing between gradients.

Experiment 4a: Sensitivity to scrambled luminance and chromaticity gradients

To eliminate the spatial pattern but to keep the same range of luminance and chromaticity values as in Experiment 1, gradients were scrambled in 1-by-1 pixels (for each presentation of a gradient, the scrambling was randomized). To produce the scrambled versions of the

Observer	AR	GT
$t(13)$	-1.966	1.218
p	0.071	0.245

Table 3. Results of the paired t test comparing performance across Experiment 1 and Experiment 4b (two-tailed).

gradients, our starting point is the images used in [Experiment 1](#): the central section corresponding to the white card with the green gradient (see [Figure 2b](#)) was altered by taking each pixel in that area from its original position and placing it in a random location on the white card disrupting the spatial organization of the gradient. Similarly to [Experiment 1](#), the test and reference gradients differed in their chromaticity and luminance values, but in this experiment were devoid of their original spatial distribution. We studied light source position changes of $\pm 1, 2, 3, 4, 5, 6,$ and 7 deg.

Results

[Figure 10a](#) shows the d' values over different light source positions for all four observers from [Experiment 4a](#). Across observers, the mean light source position difference at which observers achieved a d' value of 3 was ± 4.25 deg. A paired t test between the results of [Experiment 1](#) and [4a](#) (two-tailed; [Table 2](#)) showed that the spatial distribution in gradients was used by some observers (FM and SN) but not by others (AR and GT).

A possible strategy used by AR and GT could be to average across the gradient and compare the mean luminance estimate in order to solve the task. This strategy would lead to similar performance levels for the original gradients and the scrambled gradients. To further test this idea, spatial averages of the original gradients were presented to observers AR and GT in [Experiment 4b](#).

Experiment 4b: Sensitivity to homogenous luminance and chromaticity gradients

In [Experiment 4b](#), observers AR and GT were presented with scenes showing spatial averages of the original gradients (i.e., homogenous/uniform patches) used in [Experiment 1](#).

Results

[Figure 10b](#) shows the d' values over different light source positions for AR and GT from [Experiment 4b](#).

	Experiment 1	Experiment 3	Experiment 4a	Experiment 4b
AR	2.2	2.4	2.6	2.7
FM	3.6	2.5*	2.6*	–
GT	2.8	2.9	2.6	2.5
SN	3.3	3.0	2.6*	–

Table 4. Mean d' values for each observer and experiment. The asterisk indicates a significant difference to the performance in [Experiment 1](#) and reflects the results of the paired t tests (significance level $p \leq 0.05$).

Across observers, the mean light source position difference at which observers achieved a d' value of 3 was $+4.5$ and -5.5 deg. A paired t test (two-tailed) showed that there was no significant difference between the results of [Experiment 1](#) and [4b](#) ([Table 3](#)).

These results are consistent with observers AR and GT employing an averaging strategy that does not make use of the spatial distribution of luminance values in a gradient.

[Table 4](#) lists mean performance for all observers and experiments.

Discussion

It is rare in the natural world to find a uniform surface perfectly homogeneously illuminated. Illumination not only changes with position relative to the light source but also with inter-reflections from other surfaces. The resulting color signal reaching our eyes is actually that of a color gradient of which we are rarely aware. These color gradients may play a similar role to that of luminance gradients in shape-from-shading algorithms because they vary with the 3D spatial configuration of surfaces and with the position of the direct light source. Inter-reflections may change the color signal by as much as 20 CIELAB units (Langer, 2001) and are a cue for shape perception (Bloj et al., 1999) and spatial layout (Madison et al., 2001).

As a first step to determining the role that gradients play in natural scene perception, we analyzed the luminance and chromaticity information in gradients as a function of a particular physical manipulation of the environment, namely the change of light source position. In the psychophysical experiments, we then investigated how sensitive human observers are to these changes. There are two options to study this question: one may either create real gradients in the real world or simulate the setup with a rendering program. Enormous progress has been made in computer graphics over the years and radiosity rendering techniques that allow spectral rendering will come close to an accurate simulation of physical reality.¹ We opted to use a physical rendering software package and a spectral rendering technique (Ruppertsberg & Bloj, 2008; Yang & Maloney, 2001) that lead to highly accurate simulation results (Ruppertsberg & Bloj, 2006).

In the first part of the paper, we analyzed the chromaticity and luminance properties of gradients occurring from two uniform surfaces as a function of light source position and card opening angle. We found that changes in light and setup geometry had clear and direct effects on both the luminance and chromaticity part of the gradient as shown in [Figures 5](#) and [6](#). Hence, color gradients contain information that can support recovery of 3D shape and scene configuration properties (Forsyth & Zisserman, 1991).

The results of our psychophysical experiments then demonstrate that observers are able to detect these changes in gradients. [Experiment 1](#) shows that observers are able to discriminate between two gradients that differ due to different light source positions, in the absence of all other directional information, such as cast shadows or shading changes on other surfaces. For our scene, we found that observers reliably detected a change in the gradient information when the light source position differed by only 4 deg from the reference scene. This corresponds to a difference in luminance contrast (computed as the contrast between the top and bottom of the gradient profile; see [Appendix A](#)) of 4%² compared with reported threshold luminance contrasts of approximately 6% on the same measure, for simple pedestal changes (Chaparro, Stromeyer, Huang, Kronauer, & Eskew, 1993). In our stimuli, the luminance component of total cone contrast within the gradient ranged from 2% to 13% for light source positions from 30 to 44 deg and for each 1-deg-step the luminance contrast component linearly increased by 0.93%. The total cone contrast ranged from 4.6 to 13.6% for light source positions from 30 to 44 deg and for each 1-deg-step it linearly increased by 0.71%. Hence, increasing the light source position corresponds to a linear increase in luminance and total cone contrast ($R^2 = 0.95$ for both). While the chromatic cone contrasts ranged from 4.1% to 4.5% for light source positions from 30 to 44 deg, there was no linear relationship between chromatic contrast and light source position ($R^2 = 0.08$). Chaparro and colleagues (1993) reported threshold chromatic contrasts of approximately 0.7% on the same measure, for simple pedestal changes.

In terms of angular difference of the illuminant direction, our thresholds are smaller than those obtained from an assessment of observers' ability to judge the illuminant direction directly, where tilt and slant were estimated accurately with an error of ± 6 deg (Pentland, 1982). However, we acknowledge that these are different tasks. A similar error of ± 6 deg was found when observers were asked to judge the gradient direction, but only when the luminance contrast (measured as the standard Michelson contrast) was 64% (Erens et al., 1993a). By this measure of luminance contrast, our reference scene had a luminance contrast of 2.2% and a 1.5% contrast difference yielded reliable discrimination. McCann and colleagues (1974) measured the threshold for visibility of a gradient as a function of a different contrast measure (dividing the luminance difference by the maximum luminance) and found it to be 23% (for a performance of 80% correct responses). For comparison, using the contrast measure of McCann and colleagues, our reference gradient had 4.3% within-image luminance contrast. The scenes that were reliably distinguished from the reference scene (33 and 41 deg scenes) had 1.1% and 7% contrast, respectively.

Using a more pragmatic color difference measure such as delta E (CIELUV), we find that the difference between the test and reference scene at the top (or bottom) of the

vertical profile predicts well the observed behavior in our psychophysical experiments. The scenes that were reliably distinguished from the reference scene (33 and 41 deg scenes) in [Experiments 1](#) and [3](#) differed by 3.7 and 3.4 delta E units at the top and 2.6 and 2.3 delta E units at the bottom end of the vertical profile. For the gradients in [Experiment 2](#), the maximum difference did not exceed 2.5 for the top end and 3.4 delta E units for the bottom end of the vertical profile (the white point was set to $XYZ = [100, 100, 100]$). Applying a color difference formula like delta E to images might stretch its intended use as it was developed for single color patches under direct illumination. But this approach has been applied to photographs and was found to correlate well with visual perceptibility and tolerances (Stokes, Fairchild, & Berns, 1992) and Zhang and Wandell's (1996) development of S-CIELAB addressed the need to compare entire images. We have also found that pixelwise delta E value computation for entire images can provide useful insights for predicting where scenes will be distinguishable (Ruppertsberg & Bloj, 2007).

In [Experiments 2](#) and [3](#), we disentangled the different roles that the chromaticity and luminance information played in discriminating between gradients. We constructed stimuli in which the gradients differed in their chromaticity distribution but shared the luminance distribution of the reference scene ([Experiment 2](#)), and another set of stimuli in which gradients differed in their luminance distribution but shared the chromaticity distribution of the reference scene ([Experiment 3](#)). Observers were unable to distinguish between gradients that differed solely in chromaticity. Only when the light source position difference was extreme (-7 deg) were some observers able to distinguish between the test and the reference gradient. Potentially, chromaticity information plays a larger role with larger illuminant direction differences. For luminance-different gradients ([Experiment 3](#)), observers showed similar discrimination thresholds as for the original gradients ($+5.25$ and -4.5 deg). These results confirm that observers' behavior in [Experiment 1](#) was mainly governed by the luminance differences.

Analyzing the stimuli in terms of total cone contrast supports these results. In [Experiment 2](#), where gradients differed only in their chromaticity component, total cone contrast varied between 7.5% and 7.9%, but no linear relationship existed between light source position and total cone contrast ($R^2 = 0.38$). In [Experiment 3](#), where gradients differed only in their luminance component, total cone contrast was very similar to values in [Experiment 1](#). Total cone contrast varied between 4.2% and 13.7% and for each 1-deg-step total cone contrast linearly increased by 0.73% ($R^2 = 0.95$).

We can also estimate performance in [Experiment 1](#) using information summation (Madison et al., 2001):

$$d'_{\text{sum}} = \sqrt{(d'_{\text{lum}})^2 + (d'_{\text{chrom}})^2}, \quad (1)$$

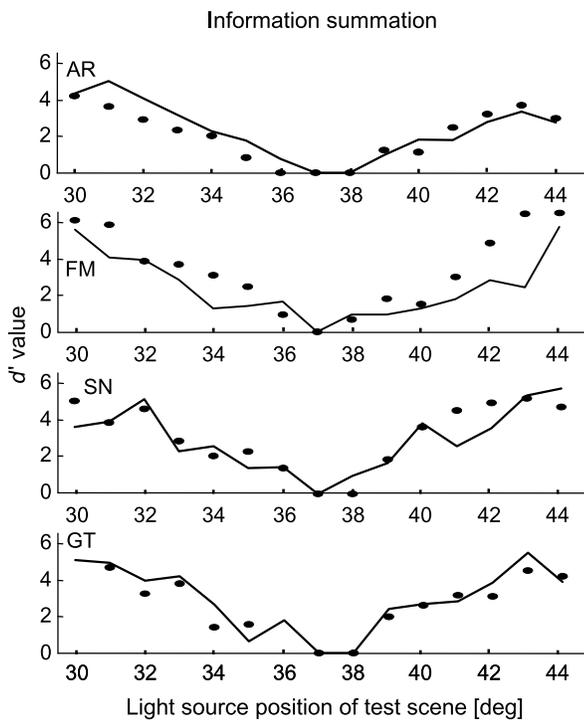


Figure 11. Predictions from information summation compared to results from Experiment 1. The points re-plot the data from Figure 7 and the solid lines indicate the predicted d' based on information summation and d'_{chrom} measured in Experiment 2 and d'_{lum} measured in Experiment 3.

where d'_{lum} is the sensitivity measure from Experiment 3 (luminance only condition) and d'_{chrom} is the sensitivity measure from the chromaticity only condition (Experiment 2). As can be seen from Figure 11 performance in the case where both luminance and chromaticity are present is well predicted by information summation.

To exclude the possibility that the constant luminance gradient in Experiment 2 overshadowed the chromaticity gradient, we ran a control experiment, in which we replaced the luminance distribution with an average luminance value and combined it with the chromaticity distribution; observers still failed to discriminate between gradients.

Our stimuli engaged both the luminance and chromatic detection system. Grating detection studies (Mullen, 1985) have shown that the contrast sensitivity function (CSF) for luminance has a band-pass and for chromatic gratings a low-pass characteristic. For spatial frequencies lower than 0.5 cycles/deg, the CSF for chromatic gratings stays constantly high. Our gradient extended vertically over 0.9 deg of visual angle, which corresponds to 0.55 cycles/deg if we assume that a gradient is half a cycle. This coincides with Mullen's (1985) threshold value for which the CSF for luminance gratings starts to be greater than for chromatic gratings. Hence, our stimulus was well chosen to engage both detection systems.

In Experiment 4a, we tested whether observers used the two-dimensional pattern of the gradient when

discriminating between them. We scrambled the gradient to disrupt the pattern while keeping the same pixel values as in the original image. If observers used the spatial layout, this manipulation should severely disrupt their performance. If, however, they used an averaging strategy when assessing gradients, scrambling would have no consequences. All four observers could perform the task but performance of two observers significantly dropped with this manipulation, showing that they had used the spatial layout of the gradient. When the other two observers were presented with averages of the gradients (Experiment 4b; homogenous patch), they showed no performance difference, indicating that they did not use the pattern of the gradient and that an averaging strategy was sufficient to describe their behavior. Interestingly, these two observers (AR and GT) scored lower d' values in general as can be seen in Table 4, which lists all observers' mean d' values from all experiments. Observers FM and SN were highly sensitive in Experiment 1 and dropped in their performance when the gradient was manipulated.

Since the gradients were always presented within the same reference scene, one could imagine that the discrimination performance was governed by the contrast of the edge of the white card to the local background and not by the within-gradient contrast. For luminance gradients, Erens and colleagues (1993a) ruled out this possibility by surrounding the gradient with an annulus that had a sinusoidal luminance pattern. Results from our Experiment 4a—in which we pixel scrambled the gradient—indicate that the local edge contrast was not crucial for observers' performance. Similarly, when we scrambled (1-by-1 pixel) the surround of the scene in a preliminary experiment, we obtained the same 4-deg difference for reliable discrimination performance (Ruppertsberg, Mayat, Hurlbert, & Bloj, 2004). We cannot completely rule out the influence of the local edge contrast, but it is certainly not the main factor governing gradient discrimination performance.

Due to our experimental design, the reference scene ("37 deg scene") was presented in every trial and it is possible that over-exposure to the same scene might have been responsible for observers' high sensitivity to small differences between gradients. To prove the generality of our results we ran a control experiment (Appendix A), in which the reference scene changed. Observers still exhibited a high sensitivity to gradient differences, demonstrating that the human visual system is highly sensitive to gradient differences.

Conclusions

Observers are able to discriminate between gradients in a complex scene caused by as little as a 4-deg change in light source position. This performance is mainly due to

the discriminability of the luminance component in the gradient and some observers also use the spatial pattern of the gradient to aid discrimination.

Appendix A

Cone contrast

Total cone contrast is calculated as:

$$TCC = \sqrt{\left(\frac{\Delta L}{L_{ave}}\right)^2 + \left(\frac{\Delta M}{M_{ave}}\right)^2 + \left(\frac{\Delta S}{S_{ave}}\right)^2}, \quad (A1)$$

where $\Delta L = L_{top} - L_{bottom}$, L_{top} is the average L cone value of the top of the vertical profile of a 5-by-5 pixel area and L_{bottom} is the average L cone value of the bottom of the vertical profile of a 5-by-5 pixel area. L_{ave} is the average across all L cone values of the vertical profile and similarly for M and S cone values. Thus, the total cone contrast represents the differential stimulus to the three cones caused by the total difference in cone excitation across the gradient, against the baseline cone excitation. To calculate the effective contrast of the gradient seen by the luminance and chromatic mechanisms, we use the method of Eskew, McLellan, and Guilianini (1999), and project the total cone contrast vector onto the luminance and chromatic mechanisms separately, and calculate the magnitude of each of these components independently.

Control experiment

In a control experiment, we tested whether a potential overexposure to a reference scene might account for the high sensitivity observed in our experiments. We used the same temporal 2-interval same-different design as before but tested four image pairs that differed by 5 deg in light

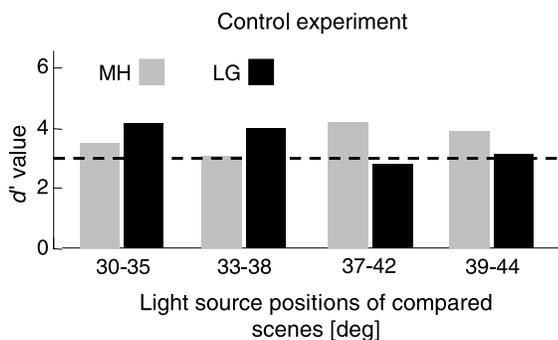


Figure A1. Results of control experiment for two new naïve observers. The dashed line indicates a performance level of $d' = 3$; the numbers identify the comparison pair (30 deg–35 deg, 33 deg–38 deg, 37 deg–42 deg and 39 deg–44 deg).

source position: 30 vs. 35 deg, 33 vs. 38 deg, 37 vs. 42 deg, and 39 vs. 44 deg. For each image pair, we presented observers with the possible combinations [A B], [B B], [B A], or [A A]. Hence, each image appeared equally often. All conditions and combinations were randomly interleaved and presented to two new naïve observers. From our previous results, we expected that observers should have a high sensitivity to a 5 deg difference in light source position. As can be seen from Figure A1 observers were able to reliably tell the images apart ($d' > 3$) in all cases.

Acknowledgments

We wish to thank our summer student Fazila Mayat (supported by the Nuffield Foundation) for her contributions and help with the experiments.

This work was supported by the EPSRC, Grant No. GR/S 13231 to MB, and was presented in part at the Fourth Symposium on Applied Perception in Graphics and Visualization in Tübingen, Germany, July 2007.

Commercial relationships: none.

Corresponding author: Marina Bloj.

Email: m.bloj@bradford.ac.uk.

Address: Division of Optometry, School of Life Sciences, University of Bradford, Bradford, BD7 1DP, UK.

Footnotes

¹RGB coding will give correct results if all participating surfaces and lights have flat spectra, i.e., are achromatic, or if only one of them is not flat.

²We computed the luminance component of total cone contrast by projecting the total cone contrast vector onto the putative chromatic and luminance mechanisms (Eskew et al., 1999); see Appendix A for details.

References

- Adams, W. J. (2007). A common light-prior for visual search, shape, and reflectance judgments. *Journal of Vision*, 7(11):11, 1–7. <http://journalofvision.org/7/11/11/>, doi:10.1167/7.11.11. [PubMed] [Article]
- Ben-Shahar, O., & Zucker, S. W. (2004). Hue geometry and horizontal connections. *Neural Networks*, 17, 753–771. [PubMed] [Article]
- Bijl, P., Koenderink, J. J., & Toet, A. (1989). Visibility of blobs with a Gaussian luminance profile. *Vision Research*, 29, 447–456. [PubMed]

- Bloj, M. G., Kersten, D., & Hurlbert, A. C. (1999). Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature*, *402*, 877–879. [PubMed]
- Bloj, M., Wolf, K., & Hurlbert, A. (2002). The perception of colour gradients [Abstract]. *Journal of Vision*, *2*(7):154, 154a, <http://journalofvision.org/2/7/154/>, doi:10.1167/2.7.154.
- Chaparro, A., Stromeyer, C. F., 3rd., Huang, E. P., Kronauer, R. E., & Eskew R. T., Jr. (1993). Colour is what the eye sees best. *Nature*, *361*, 348–350. [PubMed]
- Curran, W., & Johnston, A. (1994). Integration of shading and texture cues: Testing the linear model. *Vision Research*, *34*, 1863–1874. [PubMed]
- Curran, W., & Johnston, A. (1996). The effect of illuminant position on perceived curvature. *Vision Research*, *36*, 1399–1410. [PubMed]
- Erens, R. G., Kappers, A. M., & Koenderink, J. J. (1993a). Estimating the gradient direction of a luminance ramp. *Vision Research*, *33*, 1639–1643. [PubMed]
- Erens, R. G., Kappers, A. M., & Koenderink, J. J. (1993b). Perception of local shape from shading. *Perception & Psychophysics*, *54*, 145–156. [PubMed]
- Eskew, R. T., Jr., McLellan, J. S., & Guilianini, F. (1999). Chromatic detection and discrimination. In K. R. Gegenfurtner & L. T. Sharpe (Eds.), *From genes to perception* (pp. 345–368). Cambridge: Cambridge University Press.
- Forsyth, D., & Zisserman, A. (1990). Shape from shading in the light of mutual illumination. *Image and Vision Computing*, *8*, 42–49.
- Forsyth, D., & Zisserman, A. (1991). Reflections on shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *13*, 671–679.
- Funt, B. V., & Drew, M. S. (1993). Color space analysis of mutual illumination. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *15*, 1319–1326.
- Funt, B. V., Drew, M. S., & Brockington, M. (1992). *Recovering shading from color images*. Paper presented at the Second European Conference on Computer Vision. [Article]
- Funt, B. V., Drew, M. S., & Ho, J. (1991). Color constancy from mutual reflection. *International Journal of Computer Vision*, *6*, 5–24. [Article]
- Haddon, J., & Forsyth, D. (1998). *Shading primitives: Finding folds and shallow grooves*. Paper presented at the Proceedings of the Sixth International Conference on Computer Vision, Bombay, India.
- Hoang, A. A., Geusebroek, J.-M., & Smeulders, A. W. M., (2005). Color texture measurement and segmentation. *Signal processing*, *85*, 265–275.
- Horn, B. K., & Brooks, M. J. (1986). The variational approach to shape from shading. *Computer Vision, Graphics and Image Processing*, *33*, 174–208. [Article]
- Ikeuchi, K., & Horn, B. K. (1981). Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, *17*, 141–184. [Article]
- Knill, D. C., & Kersten, D. (1991). Apparent surface curvature affects lightness perception. *Nature*, *351*, 228–230. [PubMed]
- Knill, D. C., Mamassian, P., & Kersten, D. (1997). Geometry of shadows. *Journal of the Optical Society of America A, Optics, Image Science, and Vision*, *14*, 3216–3232. [PubMed]
- Langer, M. S. (1999). When shadows become interreflections. *International Journal of Computer Vision*, *34*, 193–204.
- Langer, M. S. (2001). A model of how interreflections can affect color appearance. *Color Research and Application*, *26*, 218–221.
- Langer, M. S., & Bülthoff, H. H. (2000). Depth discrimination from shading under diffuse lighting. *Perception*, *29*, 649–660. [PubMed]
- Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide* (2nd ed.). Mahwah, NJ: Lawrence Erlbaum Associates.
- Madison, C., Thompson, W., Kersten, D., Shirley, P., & Smits, B. (2001). Use of interreflection and shadow for surface contact. *Perception & Psychophysics*, *63*, 187–194. [PubMed]
- McCann, J. J., Savoy, R. L., Hall, J. A., Jr., & Scarpetti, J. J. (1974). Visibility of continuous luminance gradients. *Vision Research*, *14*, 917–927. [PubMed]
- Mullen, K. T. (1985). The contrast sensitivity of human colour vision to red-green and blue-yellow chromatic gratings. *The Journal of Physiology*, *359*, 381–400. [PubMed] [Article]
- Pentland, A. P. (1982). Finding the illuminant direction. *Journal of the Optical Society of America*, *72*, 448–455.
- Pentland, A. P. (1984). Local shading analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *6*, 170–187.
- Pokorny, J., & Smith, V. C. (1986). Colorimetry and color discrimination. In K. R. Boff, L. Kaufmann, & J. P. Thomas (Eds.), *Handbook of perception* (pp. 8-1–8-51). New York: John Wiley & Sons.
- Ramachandran, V. S. (1988). Perception of shape from shading. *Nature*, *331*, 163–166. [PubMed]
- Ruppertsberg, A. I., & Bloj, M. (2006). Rendering complex scenes for psychophysics using RADIANCE: How accurate can you get? *Journal of the Optical Society of America A, Optics, Image Science, and Vision*, *23*, 759–768. [PubMed]

- Ruppertsberg, A. I., & Bloj, M. (2007). Reflecting on a room of one reflectance. *Journal of Vision*, 7(13):12, 1–13, <http://journalofvision.org/7/13/12/>, doi:10.1167/7.13.12. [PubMed] [Article]
- Ruppertsberg, A. I., & Bloj, M. (2008). Creating physically accurate visual stimuli for free: Spectral rendering with RADIANCE. *Behavior Research Methods*, 40, 304–308. [PubMed]
- Ruppertsberg, A. I., Hurlbert, A., & Bloj, M. (2007). On seeing and rendering colour gradients. *Proceedings of the 4th Symposium on Applied Perception in Graphics and Visualization, APGV* (pp. 75–82). New York: ACM.
- Ruppertsberg, A. I., Mayat, F., Hurlbert, A., & Bloj, M. (2004). Sensitivity to colour gradients and its dependence on complexity of surround. *Perception*, 34, 248–249.
- Stokes, M., Fairchild, M. D., & Berns, R. S. (1992). Precision requirements for digital color reproduction. *ACM Transactions on Graphics*, 11, 406–422.
- Sun, J., & Perona, P. (1998). Where is the sun? *Nature Neuroscience*, 1, 183–184. [PubMed] [Article]
- Todd, J. T., & Mingolla, E. (1983). Perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 583–595. [PubMed]
- Ward, G. J. (1994). The RADIANCE lighting simulation and rendering system. Paper presented at the SIGGRAPH 94, Orlando, Florida.
- Wright, W. D., & Pitt, F. H. G. (1934). Hue-discrimination in normal colour-vision. *Proceedings of the Physical Society*, 46, 459–473.
- Wyszecki, G., & Stiles, W. S. (2000). *Color science* (2nd ed.). New York: John Wiley & Sons.
- Yang, J. N., & Maloney, L. T. (2001). Illuminant cues in surface color perception: Tests of three candidate cues. *Vision Research*, 41, 2581–2600. [PubMed]
- Zhang, X., & Wandell, B. A. (1996). A spatial extension of CIELAB for digital color image reproduction. *SPIE Proceedings of the Society for Information Display*, 27, 731–734.