Published in final edited form as:

J Opt Soc Am A Opt Image Sci Vis. 2009 April ; 26(4): 949–961.

Contrast, constancy, and measurements of perceived lightness under parametric manipulation of surface slant and surface reflectance

Sarah R. Allred* and David H. Brainard

Department of Psychology, University of Pennsylvania, 3401 Walnut Street, 302C, Philadelphia, Pennsylvania 19104, USA

Abstract

Across many scenes, local contrast provides a valid cue to surface reflectance, but it is not the only such cue. To generalize beyond theories of lightness that rely exclusively on local contrast, we need to know which other potential cues matter. We had observers make lightness matches between two scene locations, and varied the surface slant and local surround reflectance of one of the locations. When local contrast was a valid cue to reflectance, all observers were approximately lightness constant. When it was not, observers' lightness matches were intermediate between contrast matching and lightness constancy. For most observers, surface slant exerted an effect on perceived lightness beyond that explainable by local contrast.

1. INTRODUCTION

For perceived surface color to be a useful guide to object identity, it should correlate with surface reflectance. This is difficult to achieve because the sensory signal that reaches the eye confounds surface reflectance with the illuminant. For example, the light reflected from a ripe banana under bright midday sunlight is very different than it is under cloudy, late afternoon sunlight. The ability of the visual system to maintain a stable perception of surface color across changes in viewing conditions is called color constancy. When attention is restricted to grayscale stimuli, the term lightness constancy is often used instead.

A large body of literature confirms that human vision exhibits approximate color constancy across changes in illumination, e.g., [1–6]. A feature of most experiments is that a test object is viewed in the context of a broader scene, and the illuminant is manipulated while the objects surrounding the test are held fixed. Under these conditions, approximate color constancy may be achieved if the brain codes color through some sort of ratio between the light reflected from the test and that reflected from nearby objects [7–9]. For example, the ratio of cone responses to the light reflected from neighboring surfaces is approximately invariant with respect to changes in illumination [10–13].

The notion that perceived color and lightness at an image location depend on a ratiolike comparison between the stimulus at that location and at neighboring locations is at the core of many theories of color and lightness perception [7,14–18]. We will refer to this general idea as *local contrast coding*. Local contrast coding provides an intuitive explanation for some

illusions (e.g., simultaneous contrast; see [19,20] for discussion). It is also invoked to explain data from single-unit electrophysiology in the retina, the LGN, and the primary visual cortex [21].

If all that ever changed in a scene were the illuminant, then local contrast would always provide a valid cue to object surface reflectance. Indeed, if the surfaces in the viewed scene never vary, achieving constancy would not be a challenging computational problem. What makes constancy difficult is that both the illuminant and the contextual surfaces in the scene can change. For example, bananas fall from the plant to the ground. When the objects surrounding an object of interest change, local contrast is not necessarily a valid cue to surface reflectance. And the fact that local contrast does not always predict appearance is evident in various visual illusions (e.g., White's illusion; again see [19,20]). Experimental studies that explicitly separate effects of changing the illuminant from those of changing the local surround also show that knowing contrast alone is not in general sufficient to predict perceived color and lightness [22–24].

Although it is clear that we need to generalize beyond theories that rely solely on local contrast as the explanatory construct, the empirical foundations for such generalization remain to be established. An important agenda is to understand what stimulus factors produce effects beyond those explainable in terms of local contrast. To this end, a fruitful experimental approach is to co-vary local contrast and other aspects of the stimulus [23–25]. Here we adopt this approach to study effects of object pose in three-dimensional scenes.

Hochberg and Beck ([26]; see also [27–29]) showed that manipulations that change the perceived pose of a surface relative to a directional light source while holding the stimulus constant change its perceived lightness. This allowed them to demonstrate that the lightness effect was driven by the perceived scene layout, with local contrast held constant. More recent work has studied this type of effect parametrically [30–32], providing data that allow development and evaluation of quantitative models [30,32,33]. This parametric work does not, however, separate effects of local contrast from those of geometry. To clarify the interaction of these two factors, we report experiments that combine manipulation of surface pose and of local contrast.

2. METHODS

A. Observers

Observers were six adults between 20 and 35 years of age. Observers FP, HB, and IY were paid volunteers who were naive to the purposes of the experiment and had little experience in psychophysical observations. The other observers (SRA, DBH, RTO) were lab members with varying degrees of familiarity with the experimental design and aims. Note that observer DBH is not the second author (DHB).

B. Apparatus

Observers looked through an aperture into an experimental chamber to view two stages, as shown in Fig. 1. Observers were seated 1.3 m from the stages. Ambient illumination was provided via a single incandescent theater bulb mounted above and to the left of the observer. Light from the bulb passed through a blue filter, and the voltage to the bulb was computer-controlled. In addition, illumination to part of the booth was manipulated via a hidden projector (EPSON PowerLite 8200i) that was also computer-controlled. At RGB settings of [0, 0, 0], the projector cast some light, and this was included in our calculations of the ambient illuminant. With the projector at [0, 0, 0] and the incandescent light at its normal experimental settings, a white card on the left stage reflected light of CIE xyY coordinates of (0.41,0.41,402

cd/m²); an identical card on the right stage reflected light with CIE xyY coordinates of (0.41,0.41,261 cd/m²). A box covered with black felt was placed in the booth and mounted on four adjustable feet, two of which can be seen in Fig. 1. A small rectangular slit was cut in the box, and the box was carefully adjusted until the edge of the light generated by the hidden projector vanished through the thin slit. This light trap served to minimize observers' awareness of the hidden projector.

C. Stimuli

Stimuli were constructed by printing a standard gray surface (nominal reflectance=0.12) of size 6 cm by 6 cm centered on a grayscale Mondrian pattern (18 nominal reflectance values, range: 0.02 to 1). We defined the white surface produced by the printer as having a nominal reflectance of 1, and took the nominal reflectance of the other surfaces to be the ratio of the luminance of light reflected from them to the luminance of light reflected from the white surface under the same illumination. In the rest of this paper, we will refer to nominal reflectance simply as reflectance. Identical Mondrians (see Fig. 1) were mounted on two rotatable stages; a reference stage on the left and a match stage on the right. At standard viewing distance (1.3 m) the central surfaces subtended 2.6° of visual angle.

Simulating surfaces—Background surfaces of different reflectance were simulated by using the hidden projector in the following fashion. First, we created six achromatic surfaces (reflectances 0.12, 0.20, 0.31, 0.57, 0.72, 1). With the experimental lights on and the projector at its minimum settings [0, 0, 0], we took radiometer readings (PR650) of each surface at the reference location at 0° slant. We then took radiometer readings of each surface at the match location at each of five different surface slants (0°, 5°, 10°, 15°, 20°). The measured chromaticity changed little with location or surface slant. We fit the measured CIE xyY values as a function of the reflectance of the surface under question. In each case the data were well fit by a second-order polynomial. From this function, we could calculate the predicted CIE xyY values for a surface of any reflectance. We repeated this procedure for each of the match slants. From these fits, we could predict the CIE xyY values at each surface slant for a surface of any reflectance.

We simulated different surfaces by combining the standard surface (reflectance=0.12) with projector RGB settings calibrated to produce the CIE xyY values measured for the real surfaces. To do this, we took radiometer measurements of the standard surface at each slant under different projector RGB settings. We used the measurements together with software provided by the Psychophysics Toolbox [34] to determine the projector RGB settings required to produce the desired CIE xyY values at each slant.

The range of surface reflectances that could be simulated was limited by the gamut of the projector. Because the standard surface had reflectance (0.12) we could not in principle simulate surfaces less reflective than that. In practice, we were able to simulate adequately surfaces of reflectance 0.13 to 0.94. Because the intensity of the projector changes in discrete steps, rather than continuously, the precision with which surfaces could be simulated was also limited. We report measured CIE xyY values and simulated reflectance values actually obtained. The obtained reflectance was generally within 0.5% of the requested value.

To ensure good alignment between the projected light at the match location and the underlying standard surface, the location of the standard surface in the pixel coordinates of the projector was measured daily for each slant. The small black felt strip around the outer edge of the standard surface (Fig. 1) served to hide any residual misalignment.

In the experiments, spots of different reflectance were simulated by projecting onto the reference and match squares. At 0° slant, spots were circles with a radius of 1.25 cm and at

standard viewing distance they subtended 1.1° of visual angle. As with the geometry of the projected squares, the geometry of the projected circles was manipulated with changing slant to simulate a physically rotating circle.

Hereafter, we refer to manipulations of match and reference simply as reflectance changes rather than as simulated reflectance changes. All reported luminance values were measured in situ.

D. Psychophysical Task

Observers initiated a block of trials by pressing a key, after which a shutter opened to reveal the experimental booth (Fig. 1). On each trial, observers performed the following 2AFC psychophysical task. Two spots were presented, one at the center of the reference surround (gray square on the left stage, Fig. 1) and one at the center of the match surround (gray square on the right stage). Observers were instructed to move a joystick to indicate which spot appeared lighter. Reference and match spots were presented for 1500 ms accompanied by a 250 Hz tone. Across all blocks of trials, the reference surround reflectance was fixed at 0.16 and the reference stage was fixed at 0° slant.

Match surround reflectance and match slant were varied between blocks of trials, but remained fixed within a block of trials. Match surround reflectance took values of 0.16, 0.25, 0.34, 0.44, 0.56, and match slant took values of 0° , 10° , and 20° . Match surround reflectance and slant were parametrically varied, yielding 15 possible match conditions. Because the projector could not simulate the correct chromaticity for one match surround reflectance/slant condition (viz., reflectance=0.56, slant= 0°) it was not tested. This left 14 match conditions.

Within each block of trials (one match surround reflectance/slant condition), we calculated a point of subjective equality (PSE) for five different reference spots (reflectance values=0.18, 0.20, 0.22, 0.26, 0.32). One match surround condition (reflectance=0.34) was added midway through the experiment, and two subjects (IY, HB) were not tested in this condition. All reference spots were increments.

The procedure for each reference spot was as follows: On each trial, the reflectance of the match spot was determined by implementing an adaptive staircase calculated by the QUEST algorithm [35] as implemented in the Psychophysics Toolbox [34]. For each of the five reference spots, we ran three interleaved staircases (10 trials each) with different target response probabilities (25%, 50%, 75%). In each experimental session, observers ran between four and seven blocks of 150 trials each. A block lasted approximately six minutes, and included trials for one match slant and one match surround reflectance.

Between blocks, the shutter closed while the experimenter initiated a new block of trials with a different match slant and match surround reflectance. Observers were offered the opportunity to take a break between each block of trials, and each experimental session lasted between 35 min and 1 h.

Observers were instructed to judge the lightness of the reference and match spots. The instructions were intended to cause observers to match apparent reflectance, as opposed to their apparent luminance (often referred to as brightness) or the apparent contrast between the spots and their immediate surround. To this end, each observer underwent an induction procedure at the start of the experiment [32]. In a separate experimental room, observers were seated in front of a box that had been divided in two, with each side illuminated by a single directional light source. The right side of the box contained a paint palette, and the left side of the box contained three cubes, each of which was painted a different shade of gray. While viewing the cubes, observers were told that in the experiment, they would be matching painted surfaces or

simulations of such surfaces. Observers were instructed to hold the painted cubes and view them in different orientations and locations within the box. Subsequently, observers were shown fixed cubes with only one painted surface and asked to pick the same paint from the palette. A more detailed description of the induction procedure and instructions is provided as part of the supplemental material available at http://color.psych.upenn.edu/supplements/slant_contrast.

E. Data Analysis and Predictions

Within a block of trials, we fit the probability of reporting that the match spot was lighter as a function of match spot reflectance with a 4-parameter cumulative Gaussian. Fits were obtained using a maximum-likelihood method [36] implemented by the *psignifit* toolbox in Matlab (see <http://bootstrap-software.org/psignifit/>). Two parameters α , β determine the shape of the cumulative Gaussian, and there is a floor parameter γ and a ceiling parameter λ . The point of subjective equality (PSE) was defined as the reflectance at which observers reported the match spot as lighter on 50% of trials. Each PSE was thus based on 30 forced-choice trials.

Figure 2 shows the data and fit for one reference spot in one match condition. Pilot data indicated that within a subject, such PSEs collected on different days were highly consistent, and in fact were often identical within the reflectance resolution of our hidden projector. Because of this consistency, we generally collected two PSEs per subject per condition. For a few observers and conditions, only one PSE was collected. The full set of observer by condition by session data is tabulated in the supplement (http://color.psych.upenn.edu/supplements/slant_contrast).

3. RESULTS

We measured the perceived lightness of small match spots across parametric changes of slant and local surround reflectance. First, in Subsection 3.A we document that lightness constancy was relatively high when match surround reflectance and slant were identical to reference surround reflectance and slant. In Subsections 3.B and 3.C, we examine lightness constancy under manipulations of match slant (Subsection 3.B), where local contrast is a valid cue to reflectance, and under manipulations of match surround reflectance (Subsection 3.C), where local contrast is not a valid cue. Finally, in Subsection 3.E, we examine interactions between surround reflectance and slant.

A. Equal Slant and Surround

Figure 3 shows average PSEs for all six observers as a function of reference spot reflectance when the match surround reflectance was equal to the reference surround reflectance. The data shown, as well as all other data reported in this paper, are tabulated as part of the supplemental material (http://color.psych.upenn.edu/supplements/slant_contrast.) In this condition, the background squares on the left and the right have the same reflectance (see Fig. 1). Because the light source is to the left of the observer and angled across the booth, the incident illumination at the left location (reference) is about twice the incident illumination at the right location (match). If observers were lightness constant, then by definition the reflectance of each PSE would be identical to the reflectance of the reference spot. In other words

$$R_{\text{match_spot}}(\text{PSE}) = R_{\text{reference_spot}}, \quad (1)$$

where R indicates reflectance values. The predictions of lightness constancy are shown as solid curves in Fig. 3. In general, observers exhibited good lightness constancy for this condition, with some individual variation.

To understand the deviations from constancy, it is helpful to consider the pattern that would be shown by an observer who matched the luminance of the spots rather than their reflectance. This prediction is obtained by

$$L_{\text{match_spot}}(\text{PSE}) = L_{\text{reference_spot}}, \quad (2)$$

where L indicates reflected luminance. Because the illuminant intensity is less at the match location than at the reference location, a much higher reflectance (PSE) is needed to equate luminance of the match spot and the reference spot. The predictions of luminance matching are shown as the dotted curves in Fig. 3.

When match surround reflectance and reference surround reflectance are the same, local contrast is a valid cue to the reflectance of the match spot; in other words, the predictions of local contrast matching are the same as the predictions of lightness constancy (solid curve in Fig. 3).

To quantify the degree of constancy, we calculated an error-based constancy index (after [31]) for each subject by comparing the difference between the measured data point and the predictions made from both lightness constancy and luminance matching, as follows:

$$CI_{\text{error}} = \frac{\sqrt{\epsilon_{\text{luminance}}^2}}{\sqrt{\epsilon_{\text{luminance}}^2} + \sqrt{\epsilon_{\text{constancy}}^2}}. \quad (3)$$

Here each ϵ^2 is calculated as the sum of the squared error between each observed PSE and the relevant prediction. The index can range from 0 to 1, where 1 represents perfect lightness constancy (data along solid curve) and 0 represents luminance matching (data along dotted curve, a failure of lightness constancy). Intuitively, the index characterizes where the data fall with respect to the two different predictions. In this condition, the CI values of observers ranged from 0.58 to 0.97, as indicated in the individual plots.

B. Slant Manipulation

Observers were approximately lightness constant with respect to the illumination gradient present across the experimental chamber. To determine whether observers retained lightness constancy across illumination changes mediated by other scene variables, we manipulated match slant by rotating the stage on which the match card was mounted (see Fig. 1). Under this manipulation, we kept the match surround reflectance the same as the reference surround reflectance. Rotating the stage changed the effective illumination incident at the match location. However, since the reflectance of the surfaces did not change, local contrast remained a valid cue to surface reflectance.

Figure 4 and Figure 5 plot PSE as a function of reference spot reflectance when the match slant was 10° (Fig. 4) and 20° (Fig. 5). If subjects were lightness constant, then the reflectance of their PSEs should be unaffected by changing the match slant. This prediction is shown by the solid curves, which are unchanged across Fig. 3–Fig. 5. Since rotating the stage changes the illumination incident on the match, however, the luminance-matching predictions do change. If perceived lightness followed luminance rather than surface reflectance, PSEs would increase with slant (dotted curves in Fig. 3–Fig. 5).

Observers' PSEs were relatively constant across changes in slant; the data remain close to the solid curves rather than deviating further toward luminance matching. The CI values also reveal this constancy.

Together, Fig. 3–Fig. 5 show two effects for each observer. The first (Fig. 3) is how much constancy the observer shows across a spatial illumination gradient. The second (compare Fig. 3 with Fig. 4 and Fig. 5) is how much constancy the observer shows with respect to a change in slant, in addition to the illumination gradient. We can separate these two effects by normalizing the PSE at each match slant by the PSE at 0 match slant. We then plotted PSE as a function of match slant for each reference spot (Fig. 6). This normalization preserves information about relative constancy across changes in slant but discards information about absolute constancy.

Observers exhibited very high degrees of lightness constancy across changes in slant: for each reference spot (each different color), the normalized PSE stayed essentially the same as slant changed (slopes of colored curves are near 0). If subjects were perfectly lightness constant, PSE should not change with slant; that is, the slope of a line through the points should be 0. However, if perceived lightness followed luminance rather than surface reflectance, PSEs would increase with match slant (angled dashed black curve). We quantified the degree of relative constancy using the same constancy index as in Fig. 3–Fig. 5 applied to the normalized PSE values. Constancy index values were close to one for all six observers.

C. Reflectance Manipulation

When match slant was manipulated and local surround reflectance held fixed, perceived lightness followed surface reflectance rather than luminance. However, local contrast under this slant manipulation remained a valid cue to surface reflectance. Next, we examined whether observers would show similar degrees of lightness constancy across a manipulation where local contrast did not predict the reflectance of the match spot.

To do so, we held match slant fixed at 0° and varied match surround reflectance. The reference surround reflectance was always 0.16. We used match surround reflectances of 0.25, 0.34, 0.44 and 0.56.

To examine how perceived lightness of the match spot changed with match surround, we again normalized PSEs by the PSEs in the condition where match surround reflectance and match slant were the same as reference surround reflectance and reference slant (data in Fig. 3). The normalized PSEs as a function of match surround reflection are shown in Fig. 7. These data are also shown in unnormalized form in a supplemental figure (http://color.psych.upenn.edu/supplements/slant_contrast).

If observers were lightness constant across changes in match surround reflectance, then their PSEs would not change and the data would fall along the dashed horizontal lines in Fig. 7. Changing the match surround reflectance affects the local contrast of the match spot. If perceived lightness followed local contrast, then PSE would increase with increasing match surround reflectance (angled dashed curves in Fig. 7).

The data show that local contrast affected perceived lightness. As match surround reflectance increased, PSEs also increased. However, although PSEs were affected by local contrast, they were not completely determined by it. Normalized PSEs were significantly lower than contrast matching predictions for all observers at each match surround reflectance ($p < 0.05$, paired t -test) except one (observer FP, match surround 0.25). As in Fig. 6, we calculated an error-based constancy index, where data were compared to lightness constancy predictions and contrast matching predictions. As with the previous index, one represents complete lightness constancy.

When match surround reflectance was varied, the mean constancy index was 0.50, with individual observer values ranging from 0.29 (observer FP) to 0.61 (observer RTO).

For some of the match surround manipulations, full lightness constancy would have entailed matching a decrement to a reference spot that was an increment. Comparing increments and decrements is a perceptually difficult task [37], so we wondered whether this might have intruded upon the data. We repeated the analysis after excluding data points where the reflectance of the PSE was a decrement or a very small increment (within 0.03 of the reflectance of the match surround.) The results were similar, with a mean constancy index of 0.47 and an individual observer range of 0.16 to 0.81.

D. Intermediate Discussion

Figure 8 summarizes the data reported above. When reference slant and surround reflectance were the same as match slant and surround reflectance, observers were fairly lightness constant across the illumination gradient, as seen by the high constancy index values in the top left panel of Fig. 8. Observers were also constant across changes in match slant (top right panel, Fig. 8) when local contrast was a valid cue to the surface reflectance of the match spot. They were less constant across changes in match surround reflectance (bottom left panel, Fig. 8) when local contrast was not a valid cue.

An additional effect may be seen by closer examination of Fig. 7: the effect of local contrast depends on the reflectance of the reference spot. This is seen in Fig. 7 by noting the spread in the data for the separate reference reflectances. For each observer, the magenta points (highest reflectance reference spots) are closer to lightness constancy (dashed horizontal line) than are the red points (lowest reflectance reference spot).

To document this effect, we compared two different model fits to the data. We assumed that the normalized PSEs in Fig. 7 could be modeled as a linear function of match surround reflectance constrained to go through the normalization point of (0.16, 1). We then compared the model predictions made when PSEs for each reference spot were fit separately (colored curves in Fig. 7) to predictions made when all PSEs were modeled by one line (not shown). We used the AIC (An Information Criterion, sometimes called Akaike's Information Criterion) [38] to compare the two models. This criterion assigns scores to different models, with a lower score meaning that a model is preferred. Briefly, the likelihood of the data L given a maximum-likelihood fit to the model is calculated, and the model score decreases with increasing likelihood. The model is then penalized by the number of its free parameters K :

$$\text{AIC} = -2 \ln [L(\theta|y)] + 2K. \quad (4)$$

Two models can then be pitted against each other with the difference between AIC scores (ΔAIC) determining the extent to which one model is preferred over the other. ΔAIC values of greater than 10 are taken to indicate that one model is significantly preferred [39].

All subjects showed high ΔAICs in the direction of preferring the model that fit each reference reflectance separately. This means that the effect of manipulating the match surround was dependent on the reflectance of the spot to which PSEs were being made (Fig. 9, right panel).

For comparison, a similar analysis revealed that the effect of match slant was not dependent on the reflectance of the reference spot for all but one observer (see low ΔAICs in the left panel of Fig. 9). These ΔAICs are also consistent with the observation that, except for observer IY, the best-fit curves for each reference spot shown in Fig. 6 tend to lie on top of each other.

E. Interactions between Contrast, Slant, and Reflectance

Manipulations of slant and surround reflectance each change the luminance surrounding the match spot. Made independently, these luminance changes had different effects on perceived lightness. That is, subjects were more lightness constant when luminance changes were induced by manipulating slant than when luminance changes were induced by changing match surround reflectance. To explore more completely the relationship between local contrast and perceived lightness, we varied match slant and match surround reflectance parametrically and investigated whether the effects of slant and surround reflectance could be modeled with one function.

Figure 10 presents PSEs for three observers as a function of match surround luminance for all 14 match conditions and two different reference spots. Match slants are distinguished on the plot by color, and match surround reflectances are distinguished by symbol shape. All data points on a single panel were matches made to a single reference spot. Since both manipulations (slant and reflectance) change the luminance of the local surround, we characterized the PSE as a function of a single variable, the luminance of the match surround. Because the analysis presented in Fig. 9 suggested that the effects of match surround reflectance were dependent on the reference spot reflectance, we considered separately PSEs made to each reference spot.

Figure 10 confirms the conclusion we drew from the data in Fig. 7: the full range of data is not explainable as contrast matches (predictions shown as black solid curve). It could be, however, that the deviations from contrast matching are completely accounted for by the photometric properties of the surround with no additional effect of slant. In this case, when data are plotted as a function of surround luminance, they should fall on a common curve, independent of slant. To test whether this was the case, we compared fits made to all data simultaneously (black dashed curves, Fig. 10) with fits made separately at each slant (colored curves, Fig. 10). This is the same type of model comparison used in Subsection 3.D to test whether or not the effect of match surround reflectance on perceived lightness was dependent on reference spot reflectance.

To fit the data in Fig. 10, we used a three-parameter Naka–Rushton function:

$$L_{\text{match_spot}}(\text{PSE}) = M \frac{(gL_{\text{match_surround}})^n}{(gL_{\text{match_surround}})^n + 1}. \quad (5)$$

Parameters that best fit the data were determined using parameter search in Matlab. Of primary interest in modeling the data is whether a single function of luminance accurately described PSEs in all 14 match conditions simultaneously, or whether the PSEs must be separated into groups by slant in order to be well fit. Our choice of parametric function was somewhat arbitrary. The Naka–Rushton function has often been used because its parameters are intuitively related to relevant psychophysical, physiological, or physical variables; here its use is dictated by its utility in describing the data.

To determine whether all PSEs for a particular subject and reference spot reflectance could be well fit by a single function, we compared a model where PSEs for all 14 slant/surround conditions were fit simultaneously (1 function/3 parameters) to a model where PSEs were separated into three groups by slant (3 functions/9 parameters).

For some observers and some reference spots, it was apparent that slant mattered. For example, in the lower right panel of Fig. 10, the three match surrounds outlined by the black box had roughly equal luminance, though they differed in both slant and reflectance. For this reference spot, the PSEs were clearly distinct, indicating that perceived lightness was dependent on which

combination of reflectance and slant determined a particular match surround luminance. However, for the same observer in the same match condition, slant played a less important role in perceived lightness when matches were made to a lower reflectance reference spot (top right panel, Fig. 10). Though contrast alone could explain PSEs made to this reference spot, the function of contrast that fit the data well (dashed black curve) was not a simple ratio of local luminance values (solid black curve). These are two clear examples of when slant either mattered (bottom right panel) or did not (top right panel). However, the degree to which slant mattered was not as obvious for other observers and reference spot reflectances.

We compared AIC scores of model fits made to all data simultaneously with AIC scores of model fits made to data at each slant separately. Figure 11 shows the Δ AIC between the model in which slant mattered and the model in which it did not. The height of the bar represents the degree to which the model that takes into account slant was preferred. Δ AICs greater than 10 are thought to indicate statistical preference. For most observers, slant tended to play a more important role for higher reflectance reference spots. Although there was variability between observers, we found statistical support for the full model at high contrast reference spots for all observers except one (FP).

To verify that the Naka–Rushton function characterized the data well, we also fit each PSE with its own mean, so that PSEs were described by 14 parameters. We then calculated the relevant AIC score. Δ AICs between the 14-parameter model (each point fit by its own mean) and the 9-parameter model (PSEs separated by slant and fit with Naka–Rushton function) were less than 10 for all but one reference spot for one subject (FP, data not shown). Small Δ AICs mean that moving to 14 parameters from 9 parameters does not capture any more variability in the data, indicating that the 9-parameter Naka–Rushton model provided a good description of the data.

4. DISCUSSION

A. Central Findings

We measured perceived lightness across parametric changes in match slant and match surround reflectance. These two manipulations both changed the luminance of the immediate surround of the match spot in our experiments, but a perfectly lightness constant system should treat them differently. We found that when local contrast was a valid cue to reference spot reflectance, all observers were approximately lightness constant. When local contrast was not a valid cue to reference spot reflectance, observers' lightness matches were intermediate between predictions of contrast matching and those of constancy. Of central interest, however, is that quantitative model comparison showed that surface slant exerts an effect on perceived lightness beyond that explainable by the photometric properties of the local surround.

Our results generalize the classic report of Hochberg and Beck ([26]; see also [28]), who showed that perceived scene layout affects lightness when the retinal image is held fixed. Our stimuli differ from those used in the early demonstrations; in our experiments there were well-defined local surrounds coplanar with the reference and match spots being judged. Moreover, the local surrounds themselves were embedded in larger coplanar Mondrians. Under our conditions, we show that the local surround exerts a strong effect on surface lightness. Even under these conditions, the data show an independent effect of surface slant.

Note that when match surround reflectance was fixed, observers showed good constancy across changes in surface slant (Fig. 6; Fig. 8, top right panel). Indeed, the degree of constancy with respect to slant is higher than that found in previous quantitative studies [30,31]. In the earlier studies there was not a valid local contrast cue to constancy; presumably this is the reason for the lower constancy found there. The observation that we find deviations from constancy when

we manipulate local contrast while holding slant constant (Fig. 8, bottom left panel) is consistent with this view.

B. Effect of Instructions

It is clear that the instructions provided to observers can have profound effects on color and lightness matches [4,23,40–44]. Indeed, for some stimulus configurations, distinct matches are found depending on whether observers are instructed to match (i) apparent luminance (brightness), (ii) apparent surface reflectance (lightness), or (iii) apparent contrast [42,43]. Although it is clear that instructions can affect matches, they do not always do so [23,25,31, 43]. Why instructions matter sometimes and not others is an important issue, as is the question of what instructional effects tell us about the nature of the underlying perceptual representation. On the former point, one hypothesis is that an explicit perception of illuminant change is a critical factor in whether or not instructions affect matches [43]. On the latter point, it remains unclear whether instructional effects indicate something about the nature of the perceptual representation itself or instead indicate the action of cognitive processing applied to a perceptual representation that is itself cognitively impenetrable [45]. We look forward to the elaboration of experimental methods that can distinguish between possible causes of instructional effects [46,47].

In the meantime, we agree with the recent admonition of Blakeslee *et al.* [43] that it is important for investigators to be explicit about what instructions were provided to their observers. We gave our observers lightness instructions and accompanied these instructions with an induction procedure [32] designed to clarify what we meant by this.

There were individual differences in performance. Despite our efforts to be clear in our instructions, it remains possible that the individual differences reflect different interpretations of the instructions by different observers. In previous work [31] we found that instructional effects were small relative to individual observer variation within different instructional groups. Our current lightness instructions were more explicit than those used by Ripamonti *et al.* [31], but we have no compelling way to verify that observers understood the instructions as we intended. Indeed, even in conversations with experts it is difficult to be confident that the same words mean the same thing to all involved.

C. Local Contrast and Other Scene Variables

Previous authors have studied how local contrast interacts with global scene variables in the perception of surface lightness [19,23,42,43,47]; also see [20,24]. That work studied the effect of remote image regions, but did not explicitly manipulate 3D scene geometry. Consistent with our current findings, however, the general conclusion to be drawn from that work is that factors other than local contrast can affect surface lightness. Less understood is how local contrast interacts with other image factors. In one study [23], observers showed nearly perfect lightness constancy with respect to variation in surround luminance for incremental (but not for decremental) test stimuli. Our data deviate more from constancy, although our spots are also increments. There are enough differences in stimuli to make identifying the reason for this difference difficult. One possibility is that the thin black border between the local surrounds and the coplanar Mondrians, which was present in our stimuli for technical reasons, “insulated” the reference and match spots from the stabilizing context provided by the Mondrian and thus led to a larger effect of local contrast in our experiments. On the other hand, Gilchrist *et al.* [48] indicate that insulation effects of this sort are provided by light but not dark borders, and in a color constancy experiment Brainard [6] did not find a substantial effect of a similar thin dark border.

Both our current experiments and the work reviewed in the previous paragraph underline the importance of determining how best to frame and model the interaction between local contrast and other scene variables in the perception of lightness. A number of alternative approaches are available in the literature. One is Gilchrist's anchoring theory [19,48], which would account for effects of scene geometry through the process by which the image is parsed into local and global frameworks.

Another approach, which we have advocated, draws on an analysis of the inverse-optics computation required to achieve lightness constancy. In this tradition, Boyaci *et al.* [30] and Bloj *et al.* [33] proposed equivalent illuminant models to account for the effect of surface slant on lightness. The key explanatory variable in these models is the observer's implicit estimate of the geometry of the illumination. This estimate is referred to as the equivalent illuminant, and for each observer, parameters describing the equivalent illuminant location and directionality provide a quantitative account of the variation in matched lightness as a function of slant. A structural feature of the equivalent illuminant models, as formulated, is that the observer's estimate of illumination geometry is taken to be constant across experimental changes in the slant of the surface being judged. Because we studied only three slants, our current data do not have sufficient degrees of freedom to test the models of Boyaci *et al.* and Bloj *et al.* The fact that local surround can affect perceived lightness when slant is held constant, however, does imply that to be successful these models will have to be generalized to specify how the equivalent illuminant is affected by changes in the reflectance of objects within scenes of fixed geometry.

ACKNOWLEDGMENTS

David Harwood helped with data collection. This work was supported by the National Institutes of Health (NIH) RO1 EY10016 and P30 EY001583.

REFERENCES

1. Helson H, Jeffers VB. Fundamental problems in color vision, ii. hue, lightness, and saturation of selective samples in chromatic illumination. *J. Exp. Psychol* 1940;26:1–27.
2. Breneman EJ. Corresponding chromaticities for different states of adaptation to complex visual fields. *J. Opt. Soc. Am. A* 1987;4:1115–1129. [PubMed: 3598755]
3. McCann JJ, McKee SP, Taylor TH. Quantitative studies in retinex theory. A comparison between theoretical predictions and observer responses to the “color Mondrian” experiments. *Vision Res* 1976;16:445–458. [PubMed: 941426]
4. Arend L, Reeves A. Simultaneous color constancy. *J. Opt. Soc. Am. A* 1986;3:1743–1751. [PubMed: 3772637]
5. Brainard DH, Brunt WA, Speigle JM. Color constancy in the nearly natural image. I Asymmetric matches. *J. Opt. Soc. Am. A* 1997;14:2091–2110.
6. Brainard DH. Color constancy in the nearly natural image. II. Achromatic loci. *J. Opt. Soc. Am. A* 1998;15:307–325.
7. Land EH, McCann JJ. Lightness and retinex theory. *J. Opt. Soc. Am* 1971;61:1–11. [PubMed: 5541571]
8. Brill MH. A device performing illuminant-invariant assessment of chromatic relations. *J. Theor. Biol* 1978;71:473–478. [PubMed: 642543]
9. Brill MH. Further features of the illuminant-invariant trichromatic photosensor. *J. Theor. Biol* 1979;78:305–308. [PubMed: 491718]
10. West G, Brill MH. Necessary and sufficient conditions for von Kries chromatic adaptation to give color constancy. *J. Math. Biol* 1982;15:249–258. [PubMed: 7153671]
11. Worthey JA, Brill MH. Heuristic analysis of von Kries color constancy. *J. Opt. Soc. Am. A* 1986;3:1708–1712. [PubMed: 3772633]

12. Brainard DH, Wandell BA. Analysis of the retinex theory of color vision. *J. Opt. Soc. Am. A* 1986;3:1651–1661. [PubMed: 3772627]
13. Foster DH, Nascimento SM. Relational colour constancy from invariant cone-excitation ratios. *Proc. R. Soc. London* 1994;257:115–121.
14. Wallach H. Brightness constancy and the nature of achromatic colors. *J. Exp. Psychol* 1948;38:310–324. [PubMed: 18865234]
15. Cornsweet, TN. *Visual Perception*. Academic; 1970.
16. Land EH. Recent advances in retinex theory. *Vision Res* 1986;26:7–21. [PubMed: 3716215]
17. Walraven J. Discounting the background-the missing link in the explanation of chromatic induction. *Vision Res* 1976;16:289–295. [PubMed: 1266073]
18. Shevell SK. The dual role of chromatic backgrounds in color perception. *Vision Res* 1978;18:1649–1661. [PubMed: 726318]
19. Gilchrist, A. *Seeing Black and White*. Oxford U. Press; 2006.
20. Adelson, EH. Lightness perception and lightness illusions. In: Gazzaniga, M., editor. *The New Cognitive Neurosciences*. Vol. 2nd ed. MIT Press; 1999.
21. Shapley R. The importance of contrast for the activity of single neurons, the vep and perception. *Vision Res* 1986;26:45–61. [PubMed: 3716213]
22. Katz, D. *The World of Colour*. Paul, Kegan; Trubner, Trench, editors. 1935.
23. Arend LE, Spehar B. Lightness, brightness, and brightness contrast: 2. reflectance variation. *Percept. Psychophys* 1993;54:457–468. [PubMed: 8255708]
24. Kraft JM, Brainard DH. Mechanisms of color constancy under nearly natural viewing. *Proc. Natl. Acad. Sci. U.S.A* 1999;96:307–312. [PubMed: 9874814]
25. Delahunt PB, Brainard DH. Does human color constancy incorporate the statistical regularity of natural daylight? *J. Vision* 2004;4:57–81.
26. Hochberg JE, Beck J. Apparent spatial arrangement and perceived brightness. *J. Exp. Psychol* 1954;47:263–266. [PubMed: 13152306]
27. Gilchrist AL. Perceived lightness depends on perceived spatial arrangement. *Science* 1977;195:185–187. [PubMed: 831266]
28. Gilchrist AL. When does perceived lightness depend on perceived spatial arrangement? *Percept. Psychophys* 1980;28:527–538. [PubMed: 7208266]
29. Bloj MG, Kersten D, Hurlbert AC. Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature (London)* 1999;402:877–879. [PubMed: 10622251]
30. Boyaci H, Doerschner K, Snyder JL, Maloney LT. Surface color perception in three-dimensional scenes. *Visual Neurosci* 2006;23:311–321.
31. Ripamonti C, Bloj M, Hauck R, Kiran M, Greenwald S, Maloney SI, Brainard DH. Measurements of the effect of surface slant on perceived lightness. *J. Vision* 2004;4:747–763.
32. Doerschner K, Boyaci H, Maloney LT. Human observers compensate for secondary illumination originating in nearby chromatic surfaces. *J. Vision* 2004;4:92–105.
33. Bloj M, Ripamonti C, Mitha K, Hauck R, Greenwald S, Brainard DH. An equivalent illuminant model for the effect of surface slant on perceived lightness. *J. Vision* 2004;4:735–746.
34. Brainard DH. The psychophysics toolbox. *Spatial Vis* 1997;10:433–436.
35. Watson AB, Pelli DG. Quest: a Bayesian adaptive psychometric method. *Percept. Psychophys* 1983;33:113–120. [PubMed: 6844102]
36. Wichmann FA, Hill NJ. The psychometric function: I. Fitting, sampling, and goodness-of-fit. *Percept. Psychophys* 2001;63:1293–1313. [PubMed: 11800458]
37. Whittle P, Challands PD. The effect of background luminance on the brightness of flashes. *Vision Res* 1969;9:1095–1110. [PubMed: 5350377]
38. Akaike H. A new look at the statistical model identification. *IEEE Trans. Autom. Control* 1974;19:716–723.
39. Burnham, KP.; Anderson, DR. *Model Selection and Multi-Model Inference*. Springer; 2002.
40. Arend LE, Goldstein R. Simultaneous constancy, lightness, and brightness. *J. Opt. Soc. Am. A* 1987;4:2281–2285. [PubMed: 3430215]

41. Bauml KH. Simultaneous color constancy: how surface color perception varies with the illuminant. *Vision Res* 1999;39:1531–1550. [PubMed: 10343819]
42. Arend LE, Spehar B. Lightness, brightness, and brightness contrast: 1. Illuminance variation. *Percept. Psychophys* 1993;54:446–456. [PubMed: 8255707]
43. Blakeslee B, Reetz D, McCourt ME. Coming to terms with lightness and brightness: effects of stimulus configuration and instructions on brightness and lightness judgments. *J. Vision* 2008;8:1–14.
44. Reeves AJ, Amano K, Foster DH. Color constancy: phenomenal or projective? *Percept. Psychophys* 2008;70:219–228. [PubMed: 18372745]
45. Pylyshyn Z. Is vision continuous with cognition? the case for cognitive impenetrability of visual perception. *Behav. Brain Sci* 1999;22:341–365. [PubMed: 11301517]discussion, pp. 366–423
46. Logvinenko AD, Petrini K, Maloney LT. A scaling analysis of the snake lightness illusion. *Percept. Psychophys* 2008;70:828–840. [PubMed: 18613631]
47. Hillis JM, Brainard DH. Distinct mechanisms mediate visual detection and identification. *Curr. Biol* 2007;17:1714–1719. [PubMed: 17900902]
48. Gilchrist A, Kossyfidis C, Bonato F, Agostini T, Cataliotti J, Li X, Spehar B, Annan V, Economou E. An anchoring theory of lightness perception. *Psychol. Rev* 1999;106:795–834. [PubMed: 10560329]

Fig. 1.
Observer's view of experimental setup. The circular stages on which the cards rested could be rotated.

Fig. 2.

Example psychometric function for observer DBH. On each trial, the reference spot reflectance was 0.32, the reference surround reflectance was 0.16, the reference slant was 0° , the match surround reflectance was 0.16, and the match slant was 0° . The match spot reflectance was selected on each of 30 trials by the staircase procedure. For visualization purposes, the 30 trials have been divided into six bins of five trials each. In each bin, the reflectance values and responses were averaged to get the x and y values for each plotted data point. The horizontal bars show the standard error of the mean match spot reflectance for that bin. The black curve represents the best-fit cumulative Gaussian to the data. The black vertical line represents the extracted point of subjective equality (PSE), or where the best fit curve reaches 50%.

Fig. 3.

PSE as a function of reference spot reflectance for all six observers in one condition (reflectance of match surround=0.16, match slant=0°). Error bars represent standard error of the mean across sessions. Solid curves represent PSE predictions for both contrast matching and lightness constancy, which coincide for this condition. Dotted curves represent luminance matching predictions. The error-based constancy index is reported in the lower right corner of each panel.

Fig. 4.

PSE as a function of reference spot reflectance for all six observers in one condition (reflectance of match surround=0.16, match slant=10°). Error bars represent SEM across sessions. Solid curves represent PSE predictions for contrast matching and lightness constancy, which coincide for this condition. Dotted curves represent luminance matching predictions. The error-based constancy index is reported in the lower right corner of each panel.

Fig. 5.

PSE as a function of reference spot reflectance for all six observers in one condition (reflectance of match surround=0.16, match slant=20°). Error bars represent SEM across sessions. Solid curves represent PSE predictions for contrast matching and lightness constancy, which coincide for this condition. Dotted curves represent luminance matching predictions. The error-based constancy index is reported in the lower right corner of each panel.

Fig. 6.

Normalized PSE as a function of match slant for all six observers. Each color represents a different reference spot reflectance. Reference surround was equal to match surround (0.16). At each slant, the PSE was normalized by the PSE at slant 0. Colored curves represent the best-fit line through the data points when the curve was constrained to go through the point (0, 1). Horizontal dashed black curves represent predictions of lightness constancy and contrast matching, which are the same for this condition. Angled dashed black curves show predictions of luminance matching. Constancy index values are calculated using normalized PSEs.

Fig. 7.

Normalized PSE as a function of match surround reflectance for all six observers when match slant was equal to reference slant (0°). Each color represents a different reference spot reflectance. At each match surround, the PSE was normalized by the PSE obtained in the condition where match surround and reference surround were of equal reflectance (0.16). Colored curves represent the best-fit line through the data points, when the curve was constrained to go through the point (0.16, 1). Horizontal dashed black curves represent predictions of lightness constancy and luminance matching, which are the same for this condition. Angled dashed black curves show predictions from contrast. Constancy index values are calculated using normalized PSEs.

Fig. 8.

Constancy index (CI) values for each subject. Each panel represents CI values calculated in across a different condition. Top left: match slant and surround reflectance are the same as reference slant and surround reflectance. Values reported in Fig. 3. Top right: match surround reflectance and reference surround reflectance are the same, match slant is varied. Values reported in Fig. 6. Bottom left: match slant and reference slant are the same, match surround reflectance is varied. Values reported in Fig. 7.

Fig. 9.

Difference in the AIC score between a model in which normalized PSEs for all reference spots are predicted by one line and a model in which normalized PSEs are predicted by five lines, one for each reference spot. Each bar is a different subject. Δ AICs are shown for slant manipulation (left panel) and match surround reflectance manipulation (right panel). Horizontal black line indicates a Δ AIC of 10.

Fig. 10.

PSE as a function of match surround luminance for three observers for the lowest reflectance reference spot (top panels, reference spot reflectance=0.16) and the highest reflectance reference spot (bottom panels, reference spot reflectance=0.32). Each color represents a different slant (red=0°; blue=10°; green=20°) and each symbol represents a different match surround reflectance (cross=0.16; circle=0.25; star=0.34; diamond=0.44; square=0.56). Black solid curves are predictions from contrast matching, and colored solid curves are predictions for full lightness constancy at each slant. Dashed curves represent maximum-likelihood fits of the data to the modified Naka–Rushton function. Colored dashed curves fit data separately by slant, black dashed curves fit data from all slants simultaneously.

Fig. 11.

Difference in AIC score between the two models for each reference spot reflectance (dark blue=0.18, light blue =0.20, green=0.22, orange=0.26, red=0.32) and each observer. Higher bars indicate that nine parameters were required to fit the data. Horizontal black bar represents ΔAIC of 10.