

massachusetts institute of technology — artificial intelligence laboratory

Role of Low-level Mechanisms in Brightness Perception

Pawan Sinha and Antonio Torralba

AI Memo 2001-017
CBCL Memo 203

August 2001

ABSTRACT

Brightness judgments are a key part of the primate brain's visual analysis of the environment. There is general consensus that the perceived brightness of an image region is based not only on its actual luminance, but also on the photometric structure of its neighborhood. However, it is unclear precisely how a region's context influences its perceived brightness. Recent research has suggested that brightness estimation may be based on a sophisticated analysis of scene layout in terms of transparency, illumination and shadows. This work has called into question the role of low-level mechanisms, such as lateral inhibition, as explanations for brightness phenomena. Here we describe experiments with displays for which low-level and high-level analyses make qualitatively different predictions, and with which we can quantitatively assess the trade-offs between low-level and high-level factors. We find that brightness percepts in these displays are governed by low-level stimulus properties, even when these percepts are inconsistent with higher-level interpretations of scene layout. These results point to the important role of low-level mechanisms in determining brightness percepts.

INTRODUCTION

Two kinds of accounts have been proposed for explaining brightness perception phenomena. For want of better terminology, we shall refer to these as the 'low-level' and 'high-level' accounts. The teleological motivations for both accounts are similar – they are meant to explain how the visual system can cope with varying illumination and transparency in the natural world. However, they differ in the level of sophistication of their underlying neural mechanisms (and, hence, in their likely loci along the visual pathway). While the low-level account posits simple filter-like mechanisms operating perhaps as early as the retina¹⁻⁵, the high-level account invokes comparatively more sophisticated processing strategies such as junction analysis^{6, 7}, memory based reasoning^{8, 9}, three-dimensional shape recovery¹⁰⁻¹⁴ and gestalt-like principles of perceptual organization^{13, 15, 16, 17}. The historical antecedents of the two accounts can be traced back to the late 19th century. Hering¹⁸ championed low-level mechanisms, such as lateral inhibition, acting upon the stimulus array very early along the visual pathway. Helmholtz¹⁹, on the other hand, proposed that brightness percepts are based on a sophisticated analysis of scene layout in terms of transparency, illumination and shadows.

Differences between these two approaches are illustrated in the explanations they offer for a classical brightness illusion – simultaneous contrast induction²⁰ - shown in figure 1. The small gray squares inside the large black and white squares have identical image luminance. Perceptually, however, the small

square on the black background appears lighter. According to the low-level account, the differences in perceived brightness of the two inner squares are due to the different amounts of inhibition induced by the black versus the white surround. According to the high-level account, the dark side of the simultaneous contrast display is interpreted by the visual system as a region in shadow or as a region overlaid with dark transparent film. For the inner square in the dark region to project the same image luminance as that of the one on the light side despite the shadow or the overlaid dark film, its ‘real luminance’ must be higher (to compensate for the shadow or transparency induced attenuation). This inference induces a change in the perceived brightness – the square on the dark side appears lighter than the one on the white side.

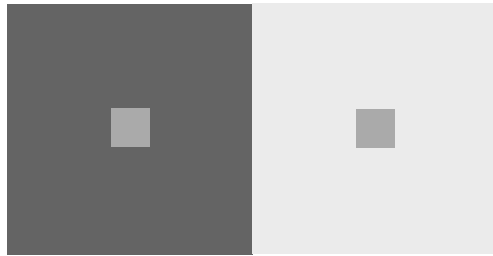


Figure 1. A simultaneous-contrast display. The two small squares appear to have different brightness even though they have identical luminance.

High-level explanations have proved remarkably versatile in accounting for several brightness phenomena^{7-11, 13, 14, 16, 21-24}. In fact, even for brightness phenomena that have traditionally been explained via low-level accounts, such as Mach bands, high-level accounts provide viable explanations. The extensive catalogue of brightness phenomena that has accumulated over the past century can be broadly divided into two categories – phenomena such as simultaneous contrast induction that can be explained by both low-level and high-level accounts and phenomena such as Benary’s cross (and several recent elegant demonstrations^{7-11, 13, 14, 16, 21-24}) for which only high-level accounts seem to provide plausible explanations. It appears, therefore, that the domain of applicability and explanatory power of high-level accounts subsumes the domain of low-level accounts. This observation leads naturally to the question that if high-level processing can adequately explain brightness phenomena, then why should low-level mechanisms be invoked at all in accounts of brightness perception? One plausible answer is that low-level mechanisms need to be invoked insofar as high-level accounts may be implemented, at least in part, via low-level mechanisms. However, at present this hypothesis lacks direct experimental support. Consequently, it provides inadequate grounds to attribute a necessary role to low-level mechanisms in brightness perception.

A less speculative and more direct way of addressing the question is to determine whether there exist any brightness phenomena for which low-level accounts predict the observed percepts while high-level accounts do not. Finding such phenomena would provide unequivocal evidence for a role of low-level mechanisms in brightness perception and would establish that such mechanisms are a necessary, though not sufficient, component of a comprehensive account of brightness perception. With this motivation, we have

devised three displays for which low-level and high-level accounts produce qualitatively and quantitatively different predictions about observers' brightness percepts. The displays also allow us to carefully study potential trade-offs between low-level and high-level factors. In what follows, we describe results of psychophysical experiments conducted with each of these three displays. We find that in these displays, it is the predictions of the low-level account that are consistent with observers' brightness percepts.

RESULTS

We report results from three experiments. Experiment 1 explores whether the perception of a region as a shadow patch or as a painted patch influences the perceived brightness of a region within it. Experiments 2 and 3 compare the roles of perceived versus physical photometric contexts in determining region brightness.

Experiment 1:

Figure 2(a) shows the setup we used for the first set of experiments. Two opaque patches stuck to a clear plexi-glass sheet cast shadows on an opaque white surface. A thick black outline was drawn over the penumbral boundary of one of these shadow patches. This, as Hering discovered more than a century ago¹⁸, dramatically changed the interpretation of the dark region. Instead of being perceived as a shadow patch, the region appeared to be painted with a uniform shade of gray. The display thus had two patches of exactly the same luminance, but one of them was perceived as being dark due to shadow while the darkness of the other was attributed to reduced surface reflectance. To insure that the observers did indeed perceive the two regions as different (one in shadow, the other painted dark), an opaque screen was positioned so as to hide part of the plexi-glass sheet with the patch casting a shadow in the 'paint region'. This screen did not obscure the observers' view of the shadows. When presented with this display, all subjects described perceiving the two regions differently – one as a shadow and the other as gray pigment.

These different interpretations provided an opportunity to test for scene-level effects on the perceived brightness of probe patches. Our experiment was designed to test whether two identical gray patches, when placed within the two regions would appear to have different brightness. For each pair of probes, subjects were asked to respond 'same' or 'different' on the basis of perceived brightness.

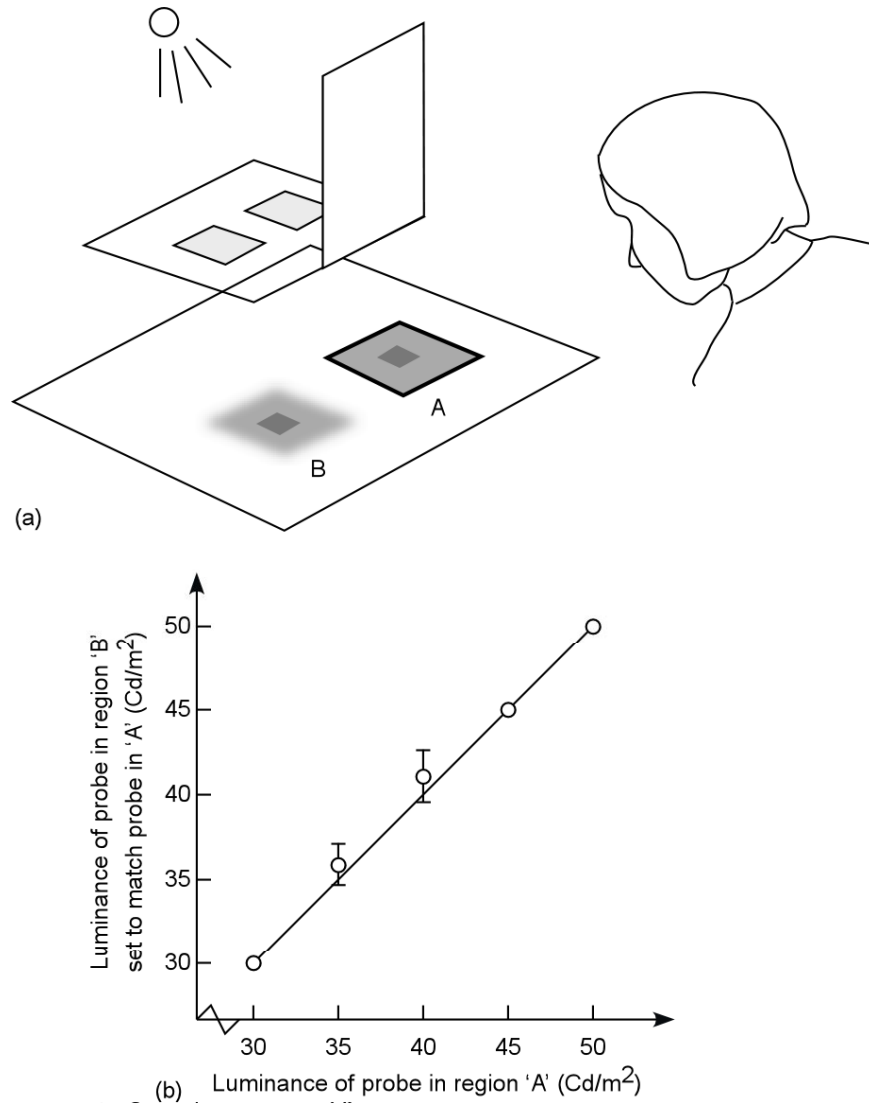


Figure 2. (a) The setup used in our first set of experiments. A black border changes the interpretation of one of the shadows to be a lower reflectance region. Subjects were asked to match the brightness of the two gray patches placed within the two regions. (b) Results from brightness matching experiments. Subjects' matches are nearly veridical for several patch reflectances indicating the lack of a brightness induction phenomenon.

A high-level account would suggest that the probe placed in the shadow region would appear to be brighter than the one placed in the 'paint' region, since the visual system would be expected to compensate for the attenuation due to shadowing. The low-level account, on the other hand, would predict no difference in the appearance of the two probe patches (or a small one in the opposite direction due to a small decrease in the average luminance caused by the black border). In order to quantitatively assess the influence of high-level factors on brightness percepts in this display, we compared brightness-matching results obtained under three conditions which differed in the appearance of the two regions (we will refer to the regions as 'A' and 'B'). The three conditions were: 1. both 'A' and 'B' seen as shadows (neither of the two had a

black outline), 2. both regions seen as pigment (black outlines drawn around both), and 3. region 'B' seen as shadow and region 'A' as pigment (black outline around 'A'). The results were plotted in a graph (figure 2(b)) where the abscissa denoted the luminance of the probe in region 'A' and the ordinate denoted the luminance of the probe in region 'B' chosen by subjects as having the same perceived brightness as the probe in 'A'. If the visual system resorts to high-level analysis and compensates for illumination attenuation due to shadowing, the data-points for condition 3 would be expected to lie below the line of unit slope passing through the origin while data from conditions 1 and 2 would lie along the line. The sum of the distances of the points from the unit-slope line and the slope of the regression line would provide a quantitative measure of the influence of high-level factors on brightness percepts. We find that data-points from conditions 1 and 2 lie along the line and those from condition 3 are either on the line or slightly above it (regression analysis revealed the slope of the best-fitting line for data from condition 3 to be 0.98), indicating nearly veridical brightness matching uninfluenced by the very different perceptual interpretations assigned to their backgrounds (figure 2(b) shows data from condition 3). The slight deviations away from the line that do exist are in a direction contrary to the predictions of the high-level account.

Experiment 2:

Figure 3(a) shows the general structure of the display used in the second experiment. It comprised a thin rectangle embedded in a larger one. Both rectangles could be assigned precisely controlled luminance gradients along their lengths. In this display, the perceived brightness profile of the inner rectangle is governed by two factors: its actual luminance gradient and the gradient induced by the spatially varying luminance profile of the enclosing rectangle. To create our experimental display, we assigned the inner rectangle a luminance gradient having a magnitude and direction such as to precisely null the induced gradient from the surround. As a consequence, the inner rectangle perceptually appeared to have uniform brightness throughout its extent, even though it actually possessed a non-zero luminance gradient (figure 3(b)). Our experiments involved placing two small horizontally separated probes with identical luminance within the inner rectangle. Subjects were asked to adjust the brightness of one of the probes to have it match the brightness of the other one.

For this display, a high-level account would be expected to predict one of two outcomes. If the visual system infers illumination distributions based on the appearance of the outer rectangle, then the two probes would appear to have different brightness, with the one on the right looking lighter (since it is in the low illumination zone). On the other hand, if the illumination distribution is inferred based on the appearance of the inner strip, the two probes would be expected to look similar since they are embedded in a perceptually uniform field and there is no manifest cause for them to appear different. A low-level mechanism, which relies on comparisons of the actual image luminances, would predict that the probes

would be affected by the physical (but perceptually non-apparent) gradient of the inner strip, causing the probe on the right to appear lighter.

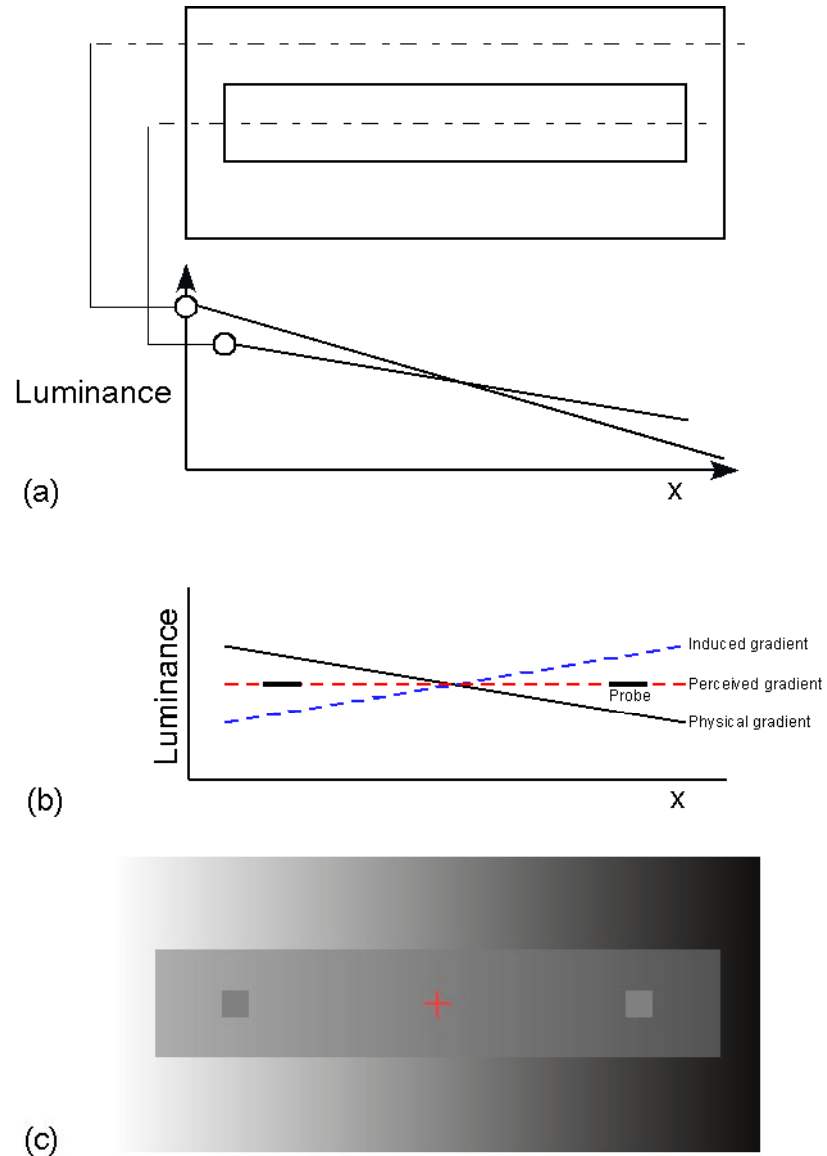


Figure 3. (a) The basic display design used in our second experiment. Both rectangles have linear luminance gradients. The magnitude of the outer gradient is fixed but that of the inner gradient is under experimental control. (b) By setting the inner gradient to precisely null the induced gradient from the surrounding rectangle, the inner strip can be made to appear as having homogenous brightness throughout its length. (c) Two identical probe squares placed within the perceptually uniform strip appear to have different brightness.

As the reader can verify from figure 3(c), the two physically identical probe squares embedded in a perceptually homogenous field are perceived as having different brightness. Figure 4(a) shows the quantitative differences in the perceived brightness of the probes averaged across five observers. Thus, the

two probe squares look different even though there is no apparent scene-level cause within the inner rectangle to motivate such a difference. These results are consistent with the operation of low-level mechanisms on raw image luminances.

However, as we indicated above, these results are also consistent with a high-level analysis of the outer rectangle, rather than the inner one. In other words, the visual system may be able to infer the prevailing illumination conditions on the basis of the large enclosing rectangle. This possibility is made especially plausible given that the outer rectangle (a) is the largest surface in the display, (b) has the highest luminance, and (c) encloses the inner rectangle. All of these are principles that have been used to define surface whites and provide information about the prevailing illumination conditions¹⁶. Thus, it is not clear whether the observed brightness percept is due to low-level mechanisms operating on raw image luminances or high-level inferences about illumination conditions based on the outer rectangle. We tested this issue in two ways – first by examining the effects of removing the enclosing rectangle on the perceived brightness of the probe squares and second by removing the inner rectangle (replacing it with a uniform black area) and thus exploring whether induction from the outer rectangle on its own could account for the observed brightness difference between the probe squares. Results from these two manipulations are shown in figures 4(b) and (c). It is evident that removal of the outer rectangle does not significantly alter the perceived brightness difference between the two probe squares. Furthermore, the outer rectangle on its own is inadequate to induce a substantial brightness difference between the two probes. On the basis of these results, we conclude that the observed brightness difference is due to the actual, but perceptually non-apparent, gradient in the inner strip. In other words, the brightness percepts here are engendered by mechanisms operating before the final perceptual output and remain, to a large extent, impervious to higher-level percepts.

The resistance of these brightness percepts to higher-level influences is also indicated by the results of an additional experiment. The experiment makes use of the fact that small changes in the magnitude of the outer gradient around the ‘equilibrium’ state (when the inner strip appears perceptually uniform) can be used to induce marked changes in the appearance of the inner strip. Depending on whether the outer gradient is made slightly steeper or slightly shallower relative to the equilibrium state, the inner strip is imparted a perceptual gradient in one or the other direction. This provides a convenient way for exploring the influence of high-level factors on brightness percepts, while keeping the low-level factors constant. We find that these reversals of perceived gradient direction in the inner strip (and the corresponding changes in the high-level inferences regarding illumination or transparency gradients) do not alter the perceived brightness of the probe squares. Figure 4(c) shows the quantitative experimental data.

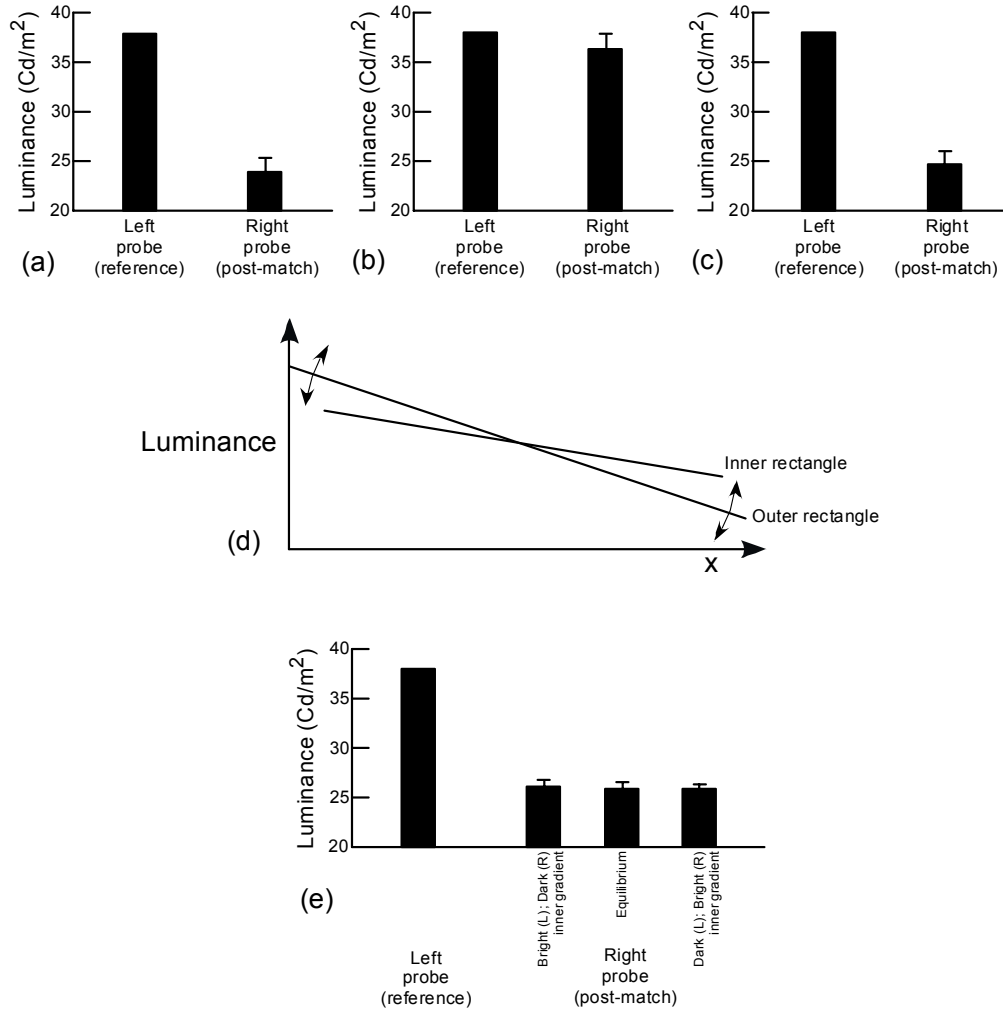


Figure 4. (a) Results from brightness matching experiments performed with the display shown in figure 3(c). (b) Brightness matching results with inner strip removed. (c) Brightness matching results with outer enclosing rectangle removed. (d) The experimental strategy for assessing the trade-off between high-level and low-level factors. By changing the gradient magnitude in the enclosing rectangle about the equilibrium state, the inner rectangle can be imparted perceptual brightness gradients in opposite directions while leaving its physical luminance values unchanged. (e) Results from brightness matching experiments under three conditions: 1. The inner rectangle appears to have a gradient that is brighter on the left, 2. The inner rectangle looks perceptually uniform, and 3. The inner rectangle appears to have a gradient that is brighter on the right. Brightness matching results stay unchanged across the three conditions pointing towards their being governed by the unchanged physical luminance structure of the inner strip.

Experiment 3:

For our third experiment, we devised a variant of the Craik-O'Brien-Cornsweet (COBC) effect^{2, 25, 26}. Unlike a conventional COBC display where the physical luminances of the regions a little distance from the central wedge are identical (figure 5(a)), in our version we set them to be slightly different in order to

produce a display where the side that was actually of higher luminance was perceived as being darker and vice-versa (figure 5(b)). The display had four equal-width regions – the two outer flanks of uniform luminance (set equidistant above and below middle gray) and the two inner flanks with shallow linear luminance ramps. We investigated how this display would affect the perceived brightness of two identical probe squares, placed one on either side of the wedge. The probes were placed in the center of the outer flanks and initially had luminance corresponding to middle gray. Subjects were asked to change the brightness of one of the probes to match the other one.

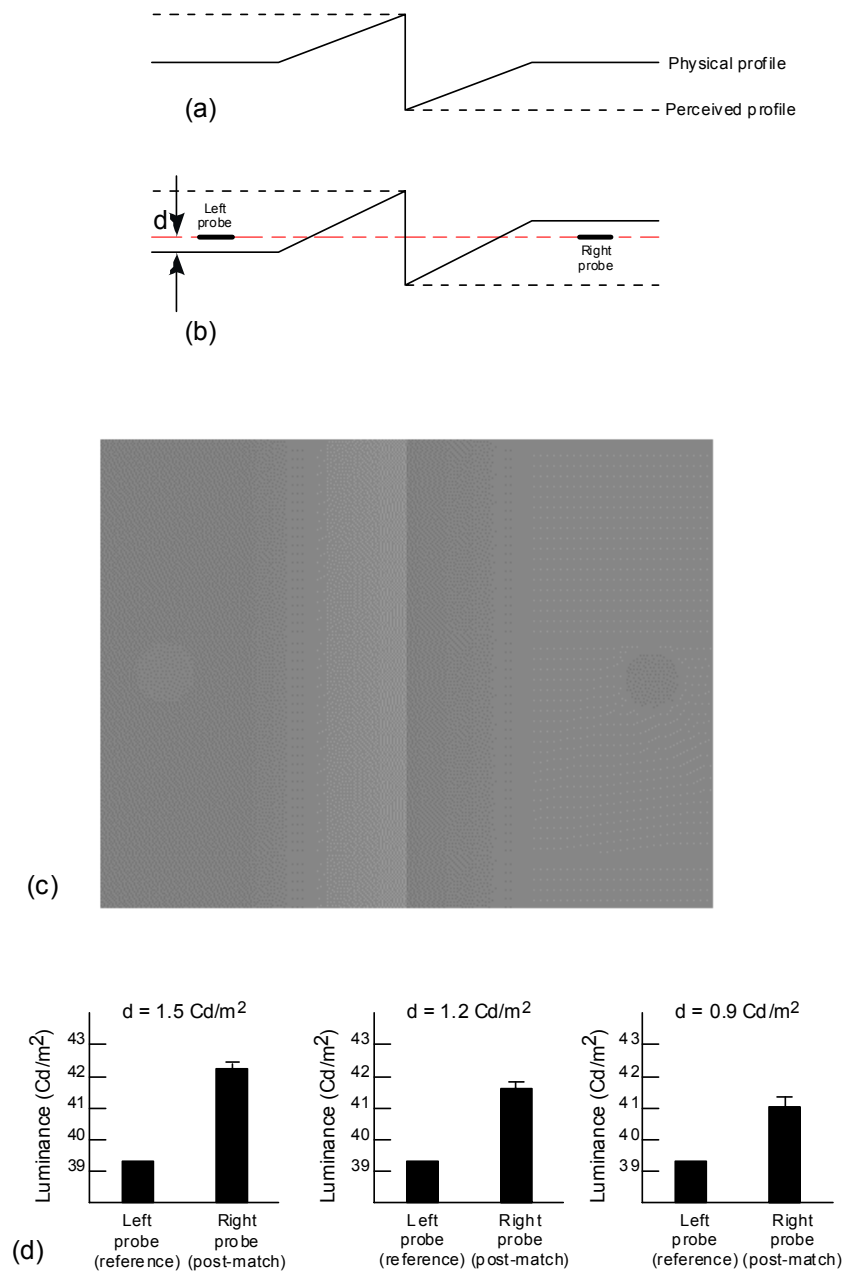


Figure 5. (a) Intensity profile across a conventional Craik-O' Brien-Cornsweet (COBC) display. Though the two flanks away from the center have identical image luminance, the gradients in the middle cause one side to appear lighter than the other. (b) Unlike a conventional COBC display, the flanks in our display differed

from each other so that the side perceived as dark actually had higher luminance than the other side. Two identical probes with their luminance set at the mean value of the flanks' luminances were placed one on each flank. (c) The appearance of the display used in our experiments. The probe on the seemingly lighter flank is perceived as being lighter than the one on the darker flank. (d) Brightness matching results from our experiments. The three sets of results correspond to a quantitative assessment of potential trade-offs between low-level and high-level factors in governing the brightness percepts in this display. We systematically reduced the luminance difference between the outer flanks to determine whether at very small differences, high-level factors would be able to overwhelm the low-level ones. We found that for all values of luminance difference tested, the brightness percepts were consistent with the low-level factors.

A high-level account (which would attribute the perceived darkness to shadow or attenuating transparency) would suggest that the probe square on the perceptually darker side would appear lighter than the other one (as in a conventional simultaneous contrast display shown in figure 1). A low-level account would predict an effect in the other direction. Figure 5(c) shows the appearance of the display. The experimental results, shown in figure 5(d), are consistent with the prediction of the low-level account. All subjects perceived the probe on the perceptually darker (but physically brighter) side as being darker than the other one. In order to quantitatively assess potential trade-offs between low and high-level factors, we investigated whether scene-level analysis could overwhelm the low-level influences if the latter were weakened by reducing the luminance difference between the flanks. As figure 5(d) shows, we found that for all values of luminance difference tested, the perceived brightness of the probe squares remained consistent with low-level factors.

DISCUSSION

Taken together, data from the three sets of experiments reported above provide compelling evidence for a role of low-level mechanisms in brightness perception. It may seem unusual for us to present this 'back to the basics' result as our main conclusion. After all, researchers have argued for a role of low-level mechanisms in estimating brightness values in various displays for more than a century. However, what is notable, and a motivating factor for the current work, is that for the set of brightness phenomena studied so far, it is possible to propose high-level explanations that can supplant traditional low-level accounts. For instance, even the illusion of Mach bands, which is generally thought of as arising out of low-level processes, admits a high-level account²⁹. Thus, the experiments to date do not place sufficient constraints to preclude purely high-level accounts of brightness perception. Not surprisingly, this ambiguity has led researchers to reconsider conventional ideas regarding the role of low-level mechanisms in brightness perception. Indeed, some recent papers have argued for a purely high-level theory^{9, 21}. The reason such extreme positions are tenable is that so far it has not been conclusively shown that low-level factors are necessary in brightness perception. This ambiguity represents a fundamental gap in the field. The contribution of our experiments lies in resolving the ambiguity by using displays for which low-level

and scene-level accounts yield different predictions. Furthermore, by allowing independent manipulation of the image cues relevant for low-level and high-level mechanisms, our displays provide a convenient tool for studying the trade-offs between low-level and high-level mechanisms.

It can be argued that the reason for the observed lack of high-level influences in experiments two and three is that the displays may not be readily interpretable in terms of scene-level factors such as illumination gradients or three-dimensional structure (in experiment 1, subjects did not report any difficulties in interpreting the gray regions as shadow or reflectance changes – a testimony to the compelling quality of Hering’s illusion). To address this issue, we have created additional variants of the displays used in our experiments. These new versions are designed to permit an easy interpretation of the displays in terms of the scene’s three-dimensional characteristics and illumination distributions. At issue is whether by facilitating such analysis regarding a scene, the influence of high-level factors in governing brightness percepts may become evident. Figures 6 (a) and (c) show two of the variants we created corresponding to the displays in figures 3(c) and 5(c) respectively. The gradients in figure 6(a) are easily interpretable as illumination variations as is the difference in brightness of the two lower faces of the cube in figure 6(c). However, notwithstanding the inclusion of cues for aiding high-level scene analysis, we find that the brightness percepts obtained with these displays (results shown in figures 6(b) and 6(d)) are unchanged relative to those in the original ones. This finding is consistent with the hypothesis that the brightness percepts are governed by low-level aspects of these displays.

While our results demonstrate a role of low-level mechanisms in brightness perception, they do not imply that such mechanisms constitute a comprehensive account of all brightness phenomena. Complementary to our demonstrations reported here are several ingenious displays which show that non-local scene configuration plays an important role in determining brightness or color percepts^{7-11, 13, 14, 16, 21-24, 27, 28}. Our current efforts focus on understanding how the visual system arbitrates between the low and high-level factors in order to arrive at a unified brightness percept.

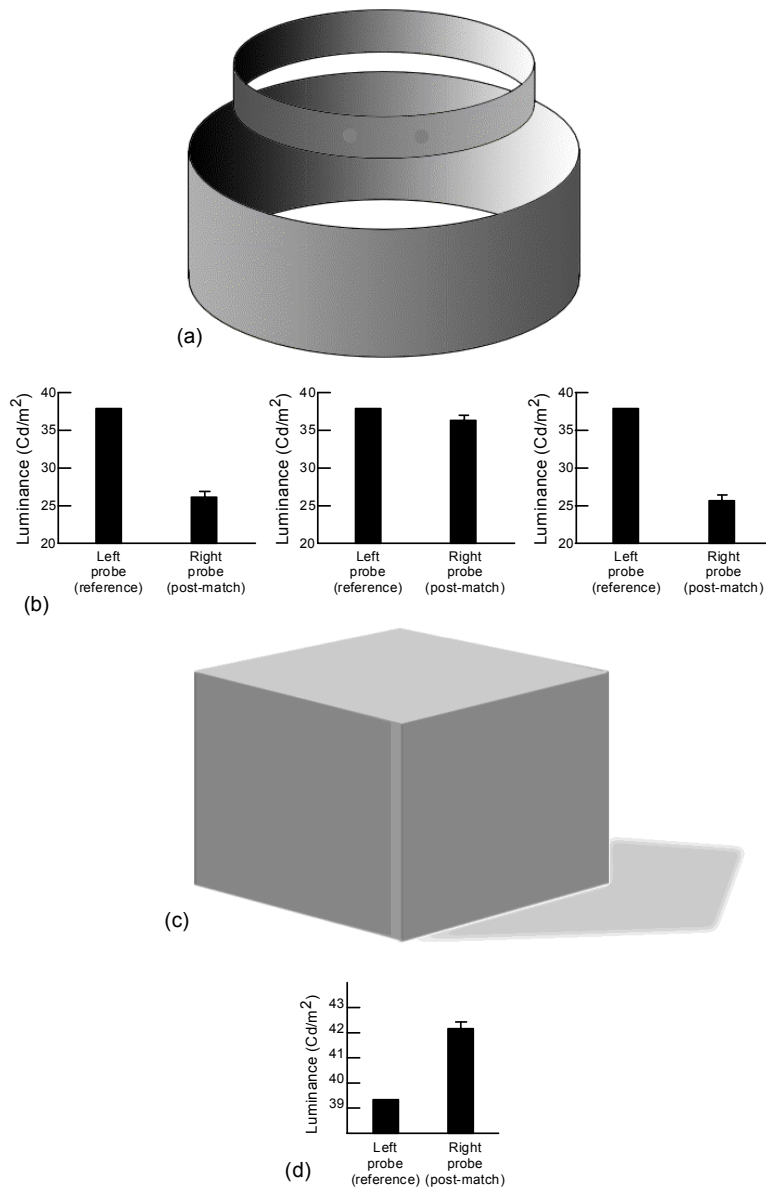


Figure 6. Variations of the displays used in our experiments, intended to facilitate their interpretation in terms of scene-level factors such as illumination gradients and three-dimensionality. (a) A variant of the display shown in figure 3(c). The gradients are now readily interpretable as gradual illumination changes with the light source positioned on the left. The brightness percepts of the two probe dots are unchanged relative to figure 3(c) (please note that the gradient directions here are left-right flipped relative to those in figure 3). (b) Brightness matching results with the display shown in (a) and its two variants. The middle panel shows matching results after making the outer surface of the top ring black and the right panel shows results after removal of lower ring. (c) A variant of the modified COBC display shown in figure 5(c). The two fields of figure 5(c) are now interpreted as two differently illuminated faces of a three-dimensional cube. The probe on the right face, which is perceived to be in shadow, is seen as being darker than the other – a result that runs counter to predictions from high-level accounts. (d) Brightness matching results with the COBC display shown in (c).

METHODS

Subjects: 15 individuals, naïve as to the purpose of our experiments, participated in this study (5 subjects per experiment). All had normal or corrected to normal vision. The experiments were conducted individually with each subject.

Apparatus: Illumination in experiment 1 was provided by an incandescent bulb and a diffuse ambient, yielding a total illuminance of 620 lux. Experiments 2 and 3 were conducted on a 19" Trinitron color monitor with antireflective coating connected to a 500 MHz Pentium III computer. The monitor had 24-bit color depth and a resolution of 1024x768 pixels. Black and white shades on the monitor corresponded to luminance values of 0 and 76 Candela/m² respectively. The gamma setting was adjusted to provide a linear mapping between gray-level values and luminance. It was our experience, however, that the perceptual effects reported above were robust against significant changes in the Gamma settings.

Stimuli:

Experiment 1: The luminance of the non-shadowed and shadowed parts of the display were 130 and 85 Candela/m² respectively. The opaque patches were placed on a clear plexi-glass sheet and cast shadows on a white paper screen. Each shadow patch subtended 10 degrees of visual angle at a viewing distance of 60 cm. The probe patches subtended 1.5 degree at the same distance. The probes were cut out from neutral density sheets manufactured by ColorAid Corporation (New York). For each pair of probes, subjects were asked to respond 'same' or 'different' on the basis of perceived brightness.

Experiment 2: The display subtended 10 degrees of visual angle at a viewing distance of 60 cm. The enclosing rectangle had a linear luminance gradient going from black to white. The inner rectangle was initially set to the mean luminance of the outer rectangle. Subjects were allowed to change the gradient of the inner rectangle to render it perceptually homogeneous. The initial luminance of the probe squares was set to middle-gray. In the experiment designed to assess the effect of changes in the outer rectangle's gradient on the perceived brightness of the probes, the end points of the gradient were shifted through +/- 7.5 Cd/m².

Experiment 3: The display had four equal-width regions – the two outer flanks of uniform luminance (set equidistant above and below middle gray) and the two inner flanks with small linear luminance ramps. The entire display subtended 15 degrees at 60 cm. The probes (1 deg. each) were placed in the center of the outer flanks and initially had luminance corresponding to middle gray. Subjects were asked to change the brightness of one of the probes to match the other one.

All experiments were performed in compliance with the guidelines set down by MIT's Committee on the Use of Humans as Experimental Subjects.

ACKNOWLEDGEMENTS

The authors would like to thank Professors Alan Gilchrist, Mriganka Sur, Ted Adelson, and Tomaso Poggio for useful discussions regarding this work. This work was supported in part by the Alfred P. Sloan Foundation Fellowship in neuroscience to PS.

REFERENCES

1. Wallach, H. (1963) The perception of neutral colors, *Scientific American*, **208**: 107-116.
2. Cornsweet, T. N. (1970) *Visual Perception*. New York: Academic Press.
3. Frisby, J. P. (1979) *Seeing: Illusion, brain and mind*. Oxford University Press.
4. Land, E. H., McCann, J. J. (1971) Lightness and retinex theory. *J Opt Soc Am*, **61**:1-11.
5. Land, E. H. (1986) Recent advances in Retinex theory. *Vision Research*, **26**: 7-21.
6. Todorovic, D. (1997) Lightness and junctions. *Perception*, **26**:379-394.
7. Adelson, E. H. (1999) Lightness perception and lightness illusions. In *The Cognitive Neurosciences*. Edited by Gazzaniga M. Cambridge, MA: MIT Press.
8. Purves, D., Shimpf, A., Lotto, R. B. (1999) An empirical explanation of the Cornsweet effect, *J. Neurosci*,**19**:8542-8551.
9. Lotto, R. B., Purves, D. (1999) The effects of color on brightness. *Nat Neurosci*, **2**:1010-1014.
10. Gilchrist, A. L. (1981) When does perceived lightness depend on perceived spatial arrangement? *Percept Psych*, **28**:527-538.
11. Knill, D. C., Kersten, D. (1991) Apparent surface curvature affects lightness perception. *Nature*, **351**:228-230.
12. Sinha, P. & Adelson, E. H. (1993) Recovering reflectance and illumination in a world of painted polyhedra. *Proceedings of the Fifth International Conference on Computer Vision*, Berlin, Germany, IEEE Computer Society Press, Los Alamitos, CA, pp. 156-163.
13. Adelson, E. H. (1993) Perceptual organization and the judgment of brightness. *Science*, **262**:2042-2044.
14. Bloj, M. G., Kersten, D., Hurlbert, A. C. (1999) Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature*, **402**, 877-879.

15. Agostini, T., Proffitt, D. R. (1993) Perceptual organization evokes simultaneous lightness contrast. *Perception*, **22**:263-272.
16. Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li X., Spehar, B. Annan, V., Economou, E. (1999) An anchoring theory of lightness perception. *Psychological Review*, **106**:795-834.
17. Koffka, K. (1935) *Principles of Gestalt psychology*. New York: Harcourt, Brace & World.
18. Hering, E. (1874) *Outline of a Theory of the Light Sense*, translated from the German by L. Hurvich and D. Jameson. Harvard University. Cambridge (1964)
19. Helmholtz H. von (1867) *Handbuch der Physiologischen Optik*. Leipzig: Voss.
20. Katz, D. (1935) *The world of colour*. London: Kegan Paul, Trench, Trubner & Co.
21. Williams, S. M., McCoy, A. N., Purves, D. (1998) The influence of depicted illumination on brightness. *Proc Natl Acad Sci USA*, **95**:13296-13300.
22. Logvinenko, A. D. (1999) Lightness induction revisited. *Perception*, **28**:803-816.
23. Kingdom, F. (1999) Old wine in new battles? Some thoughts on Logvinenko's Lightness induction revisited'. *Perception*, **28**:929-934.
24. Paradiso, M. (2000) Visual neuroscience: Illuminating the dark corners. *Current Biology*, **10**:15-18.
25. Craik, K. (1966) *The Nature of Psychology*. Cambridge, UK: Cambridge Univ. Press.
26. O'Brien, V. (1959) Contrast by contour-enhancement. *Am J Psychol*, **72**:299-300.
27. Rossi, A. F., Rittenhouse, C. D., Paradiso, M. A. (1996) The representation of brightness in primary visual cortex. *Science*, **273**:1391-1398.
28. Rossi, A. F., Paradiso, M. A. (1999) Neural correlates of perceived brightness in the retina, lateral geniculate nucleus, and striate cortex. *J Neurosci*, **19**:6145-6156.
29. Lotto, R. B., Williams, S. M. and Purves, D. (1999) Mach bands as empirically derived associations. *Proc. Natl. Acad. Sci.*, **96**: 5245-5250.