

An empirical explanation of color contrast

R. Beau Lotto* and Dale Purves

Duke University Medical Center, Box 3209, Durham, NC 27710

Contributed by Dale Purves, August 7, 2000

For reasons not well understood, the color of a surface can appear quite different when placed in different chromatic surrounds. Here we explore the possibility that these color contrast effects are generated according to what the same or similar stimuli have turned out to signify in the past about the physical relationships between reflectance, illumination, and the spectral returns they produce. This hypothesis was evaluated by (i) comparing the physical relationships of reflectances, illuminants, and spectral returns with the perceptual phenomenology of color contrast and (ii) testing whether perceptions of color contrast are predictably changed by altering the probabilities of the possible sources of the stimulus. The results we describe are consistent with a wholly empirical explanation of color contrast effects.

vision | visual perception | visual association | visual psychophysics | color perception

A central problem throughout the history of attempts to understand color vision has been the relationship between the spectral content of light returned to the eye from a visual target and the color sensation elicited. A fact that has been particularly difficult to rationalize is that the same spectral stimulus can appear to be differently colored when viewed against different chromatic backgrounds (1–8). Thus, an object's color is somehow determined by the spectral characteristics of the scene in which it appears (Fig. 1).

These contextual effects, generally referred to as simultaneous color contrast or chromatic induction, present a challenge to theories of color vision. Most textbook accounts present color contrast as an incidental consequence of mechanisms that allow the spectral return from a surface to generate similar color sensations under different illuminants (i.e., color constancy; see, for example, ref. 9, pp. 119–120). The neural circuitry thought to underlie these anomalous effects is either lateral interactions among chromatically sensitive neurons at the input stages of the visual system (6, 10–16) or neuronal interactions promoting a more general adaptation of the color-processing system to the predominant spectral return (and/or spectral contrasts) in the scene (12, 14, 16–23). Despite the ability of these interpretations to predict some aspects of color contrast, other investigators have found such explanations to be unsatisfactory (see refs. 24–27).

Given the ongoing uncertainty about the explanation of color contrast effects—and indeed about color sensation generally—we here examine an alternative theory, namely that perceptions of color represent the accumulated experience of the species and the observer with the reflectances and illuminants that generated the same or similar spectral profiles in the past.

Experimental Methods

Construction and Presentation of Computer Graphics. All test stimuli were created with a Power Macintosh G3 computer, using Adobe PHOTOSHOP 5.0 software and the standard Macintosh 32-bit true color palette. The stimuli were displayed on a calibrated 48-cm (diagonal) color monitor (Sony Multiscan 300sf; monitor resolution 1024 × 768; scan rate 75 Hz; noninterlaced). The computer interface for each test routine was created with DIRECTOR 6.0 (Macromedia, San Francisco, CA).

The luminances of computer-generated colors were measured

photometrically with an optical power meter (model 371R; Graseby Optronics, Orlando, FL) under the relevant test conditions (values are given below); the average spectral content of the stimuli was determined by calculating the average activation of the red (R), green (G), and blue (B) guns. The colors of the surrounds used in the experiments reported in Fig. 3B correspond to Munsell hue red (5R), orange (5YR), yellow (5Y), yellow/green (5GY), green (5G), cyan (5BG), blue (5B), purple/blue (5PB), purple (5P), or red/purple (5RP), each at three levels of chromaticity (chroma 3, 6, and 9) and at 41 cd/m². The colors of the targets correspond to Munsell hue red (5R), yellow (5Y), green (5G), or blue (5B), presented with a chroma of 6 and a luminance of 28 cd/m².

Subjects and Testing Procedures. Subjects with normal or corrected acuity of 20/20 or better and normal trichromatic vision observed the test stimuli on the monitor screen from a distance of 60 cm in an otherwise darkened room after adaptation to the ambient light (approximately 0.5 cd/m²). The dimensions of the targets were 1° × 1° for the stimuli used to determine color contrast effects (see Fig. 3B) and 1.75° × 1.75° in the more complex scenes (see Fig. 5). Before testing, each subject was adapted to the testing conditions for 5 min. The task for the experiments described in Fig. 3B was to adjust the perceived hue, saturation, and brightness of the right target on the achromatic surround until it matched the color of the left target on the chromatic surround, using a series of “buttons” at the bottom of the scene. Similarly, for the experiments described in Fig. 5, subjects adjusted the hue, saturation, and brightness of the “adjusted targets” until they matched the color of their corresponding “test target” in the scene. Adjustments in perceived hue, saturation, and brightness were generated by changing the relative output of the red, green, and blue guns of the monitor. When the colors of the two targets appeared identical (or as similar as it was possible to make them), the subjects clicked the “match” button, which recorded the hue, saturation, and brightness values of the adjusted target for subsequent analysis, and called up the next scene. Subjects repeated each test twice for the experiments in Figs. 3B and three times for the experiments reported in Fig. 5. The values in hue–saturation–brightness computer color space for each condition were averaged and compared (see *Results*). Statistically significant differences in the experiments described in Fig. 5 were determined by Student's *t* test.

Results

Rules Describing the Interaction of Reflectances, Illuminants, and Spectral Returns. If color contrast effects such as those in Fig. 1 do indeed arise from human experience with the typical sources of spectral stimuli, then the phenomenology of these perceptions should always accord with the rules that describe the physical interaction between reflectances, illuminants, and spectral re-

*To whom reprint requests should be addressed. E-mail: lotto@neuro.duke.edu.

The publication costs of this article were defrayed in part by page charge payment. This article must therefore be hereby marked “advertisement” in accordance with 18 U.S.C. §1734 solely to indicate this fact.

Article published online before print: *Proc. Natl. Acad. Sci. USA*, 10.1073/pnas.210369597. Article and publication date are at www.pnas.org/cgi/doi/10.1073/pnas.210369597




Fig. 1. Examples of simultaneous color contrast. The spectral returns from the square targets in the centers of the red and yellow circular surrounds are identical. The color sensations elicited by the same targets on differently chromatic backgrounds are obviously different. The purpose of the present work was to examine how and why color perception is influenced by spectral context.

turns. A first step, therefore, was to examine the various ways in which reflectances and illuminants can be combined to generate spectral returns.

Any spectral stimulus is the product of the reflectance efficiency function of a surface multiplied by the spectrum of the illuminant. Although spectral stimuli in natural scenes vary widely (28), the relationship between the underlying reflectance efficiency functions of surfaces and their illuminants follows straightforward rules. These relationships are illustrated in Fig. 2. We have used a sine wave for the sake of simplicity, but the systematic changes in spectral returns that occur as a function of different illuminating spectra are equally apparent for the reflectance efficiency function of any surface and the spectrum of any illuminant. These relationships are indicated in Fig. 3*A* (1–3) and can be summarized as follows.

(i) When a surface is illuminated by a spectrum similar to the surface’s reflectance efficiency function (i.e., by wavelengths that are well reflected by the surface, as in Fig. 2*A*), the variance in the *y* axis of the distribution returned from the surface is increased (i.e., the bandwidth of the distribution between 400 and 700 nm narrows), and its overall power is also increased. The shape of the distribution, however, remains highly correlated with the return from the same surface under an equally intense illuminant whose power is uniformly distributed.

(ii) When a surface is illuminated by a spectrum that is opposite in its distribution of power to the surface’s reflectance efficiency function (i.e., by wavelengths poorly reflected by the surface, as in Fig. 2*B*), the variance in the *y* axis of the spectral distribution returned from the surface is decreased (i.e., the bandwidth of the distribution between 400 and 700 nm broadens) and its overall power is also decreased. The shape of the distribution again remains highly correlated with the return from the same surface under an equally intense illuminant whose power is uniformly distributed.

(iii) When a surface is illuminated by a spectrum that is neither the same nor the opposite of the surface’s reflectance efficiency function (as in Fig. 2*C* and *D*), the variance in *y* (i.e., bandwidth) and overall power of the distribution are affected to an intermediate degree. In distinction to the minimal effects in *i* and *ii*, however, the distribution of power in the spectral return becomes less correlated with the return from the same surface under neutral illumination, the return shifting along the *x* axis in the direction of the spectral profile of the illuminant.

(iv) Finally, as indicated by the three different curves in each of 1–3 in Fig. 3*A*, all of these effects are influenced by the variance (“purity”) of the spectral distribution of the illuminant: the narrower the distribution of power in the illuminant relative

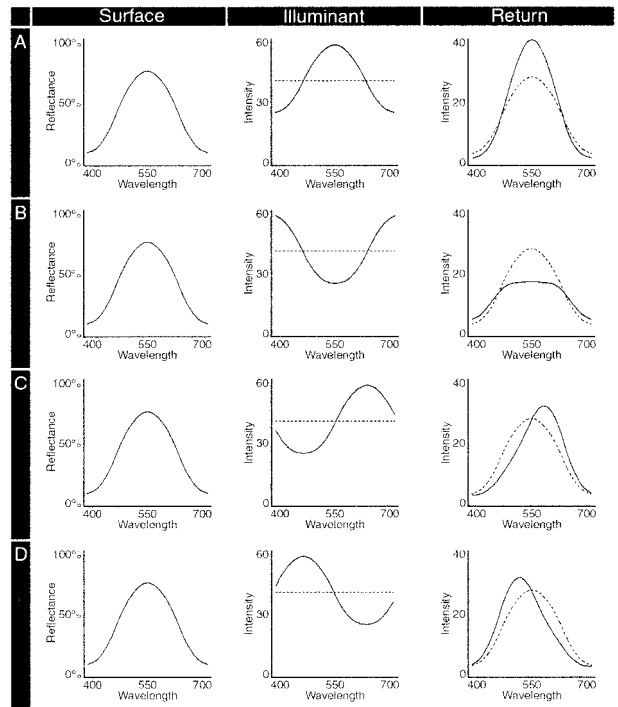


Fig. 2. The interaction of reflectance and illumination in generating spectral returns. A “surface” with a particular reflectance efficiency function, illustrated here as a simple sine wave, is “illuminated” by spectra that are the same (A) or the opposite (B) or are shifted toward long (C) or short (D) wavelengths. The spectral returns given by the product of the reflectance efficiency function and the illuminant vary systematically relative to the return generated by the same surface under an equally intense illuminant whose power is uniformly distributed (dotted line) (see text for further explanation).

to the reflectance efficiency function of the surface, the greater are all of the effects described above.

These several rules thus describe the way in which spectral returns from any surfaces are affected by different illuminants. If, as we hypothesize, spectral stimuli elicit percepts as a result of visual experience with the typical provenance of spectral returns, then when someone is presented with any chromatic combination in a color contrast stimulus, the colors perceived should accord with the rules of these physical relationships.

Rules Describing Perceptions of Color Contrast. To explore whether the colors perceived in response to stimuli such as those shown in Fig. 1 are predicted by the physical relationships illustrated in Fig. 3*A*, a series of chromatic test targets were presented to subjects on differently chromatic surrounds (see diagram in Fig. 3*B* Top) (as described in *Experimental Methods*, the spectral returns of both the target and background were selected to elicit a color sensations that ranged broadly over perceptual color space). By having subjects adjust the perceived hue, saturation, and brightness of a target on an achromatic background until its apparent color matched that of a physically identical target on an isoluminant, chromatic background, we could measure the change induced in the color of the target by the chromatic surround.

The results of these tests are shown in Fig. 3*B* (1–3) and can be summarized as follows.

(i) When a target is presented on a chromatic background whose spectral return is similar to that of the target, the apparent color of the target relative to its appearance on a neutral surround decreases in saturation and brightness, with little change in hue.

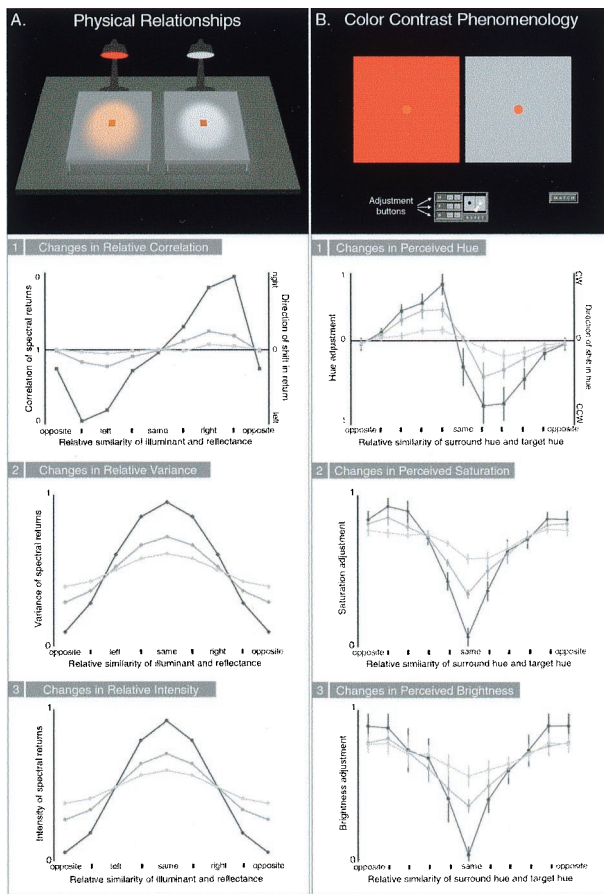


Fig. 3. Comparison of the physical relationships between surface reflectance and spectral illumination (A) and the phenomenology of color contrast (B). (A) The cartoon at the top indicates the comparison being made between the return from surfaces under illumination with light whose power is uniformly distributed (right) and illumination of the same surface with chromatic light (left). Changes are shown in the correlation (1), variance in y (or bandwidth in x) (2), and overall intensity (3) of the spectral return from the target surface relative to the return generated by the same surface under an illuminant with uniformly distributed power at three different levels of variance [indicated by the light (broad bandwidth), medium, and dark (narrow bandwidth) gray lines] (see Fig. 2). The distance along the abscissa indicates the relative similarity of the power distribution of the illuminant and the reflectance efficiency function of the surface. The left ordinate shows the normalized change in the correlation, variance, and intensity of the returns (note that increasing variance in y represents a narrowing of the bandwidth of the spectral profile); the right ordinate in 1 also shows the direction in which the spectral return is shifted relative to the return from the same surface under neutral illumination (rightward indicates a shift toward longer wavelengths, and leftward, a shift toward shorter wavelengths). (B) The cartoon at the top shows an example of the test paradigm (see *Experimental Methods*). Graphs show average adjustments made by subjects in the hue (1), saturation (2), and brightness (3) of the target on the neutral background to match the color sensation elicited by the same target on a chromatic (but equiluminant) background. Each stimulus was presented at three levels of saturation [indicated, as in A, by light (low saturation), medium, and dark (high saturation) gray lines]. The distance along the abscissa indicates the relative similarity of the hue of the surround and the hue of the target. The ordinate shows the normalized change in the hue, saturation, and brightness of the target on the chromatic surround relative to the target on the achromatic surround. CW on the right ordinate in 1 represents shifts in hue that are clockwise in Munsell color space, whereas CCW represents hue shifts that are counterclockwise. The data are the average of the responses of the two authors and one naïve subject to all contrast stimuli; bars are standard errors.

(ii) When the test target is presented on a chromatic background whose spectral return is complementary to that of the target, the apparent color of the target relative to its appearance




Fig. 4. Empirical explanation of the colors perceived in response to the color contrast stimuli in Fig. 1. When placed under an “orange” illuminant, a surface that appears yellowish in white light (Left) generates the same spectral return as a reddish surface under yellowish light (Right; cf. Fig. 1). To the extent that the two scenes in Fig. 1 are likely to have been generated by the different conditions illustrated, the left and right targets in Fig. 1 should, according to the empirical theory under consideration here, appear relatively yellowish and reddish, respectively.

on a neutral surround increases in saturation and brightness, with little change in hue.

(iii) When the test target is presented on a chromatic background whose spectral return is neither the same as nor complementary to that of the target, the hue of the target relative to its appearance on a neutral surround changes maximally (i.e., away from the hue of the surround), with little change in saturation and brightness.

(iv) Finally, as indicated by the three curves in each of 1–3 in Fig. 3B, these contextual effects on hue, saturation, and brightness are all enhanced by decreasing the variance of the spectrum of the chromatic surround.

Notice that these observed changes in the perception of hue, saturation and brightness of a target in different chromatic surrounds are in each instance complementary to the way spectral returns from surfaces change under different chromatic illuminants (compare Fig. 3A and B).

An Empirical Explanation of Color Contrast Effects. Although the results shown in Fig. 3 do not rule out explanations of color contrast based on lateral interactions between chromatically sensitive neurons, the complementary relationship between perceptions of color contrast and the physical behavior of reflectance, illumination, and spectral returns summarized in Fig. 3 suggests an alternative explanation. Rather than being an incidental consequence of receptive field properties, this relationship is what would be expected if color contrast phenomena were generated empirically by the past significance of spectral stimuli. Consider, for example, the color contrast effects elicited by the stimuli in Fig. 1. In empirical terms, because the average spectral profile of a scene is always shifted toward the spectrum of the illuminant (see Fig. 3A), the return from the surrounds in Fig. 1 increases the probability that the left and right panels are under different chromatic illuminants, as illustrated in Fig. 4. Given this probable difference in illumination, the underlying sources of the identical returns from the two targets would have been differently reflective surfaces. For instance, under this implied illuminant, the return from the left target in Fig. 1 could only have been generated by a surface that would appear yellow in white light, and the return from the right target in Fig. 1 could only have been generated by a surface that would appear orange in white light (Fig. 4). As a result of this past experience with the underlying sources of stimuli such as those in Fig. 1, the percepts generated by the target returns are different from those that would be generated by the same targets on a neutral background. In particular, the appearance of the targets would be expected to differ in a manner complementary to the physical interactions summarized in Fig. 3A.

As shown in Fig. 3B, this is, in fact, what occurs when targets are placed on different chromatic backgrounds.

The Effects of Altering the Empirical Significance of Color Contrast Stimuli. If this explanation of color contrast is correct, then the perceived color of the target in a stimulus should change in a predictable manner as the spectral characteristics of the returns from the scene are manipulated so as to make the stimulus more or less consistent with different combinations of reflectances and illuminants. Moreover, this effect should be apparent, even if the spatial chromatic average of the stimulus is unchanged.

Increasing the probability that identical target returns arise from different conditions of chromatic illumination increases perceived color differences. We therefore examined the responses of subjects to additional stimuli in which the empirical significance of the targets was changed by manipulating the characteristics of their surroundings while keeping the average spectral returns from the scenes the same. In the stimuli shown in Fig. 5A and B, for instance, subjects were asked to report the apparent color difference between identical chromatic targets embedded in uniformly colored (Fig. 5A) or multicolored (Fig. 5B) surrounds. The perceived differences in the qualities of the target color in these two circumstances were measured by having the observers adjust the apparent hue, saturation, and brightness of the same targets set against a uniform gray surround in the two test boxes beneath the scene until the differences between them matched the perceived differences between the two targets in the color contrast stimuli. The rationale for this comparison was that increasing the number of different spectral returns consistent with the two identical targets being differently illuminated should increase the probability of the target returns being generated by differently reflective surfaces. As a result, the perceived color difference of the targets should increase, in keeping with this change in the probable significance of the stimulus.

When subjects were presented with the scene in Fig. 5A, a relatively modest adjustment was required to match the two targets, the direction of the adjustment being that expected from the empirical information in the surround about the possible underlying sources of the stimulus (see Fig. 5D). When, however, the subjects were presented with the multicolored backgrounds in Fig. 5B (which have the same average spectral content as the corresponding surrounds in Fig. 5A), the color difference between the two targets (mostly hue) was perceived to be more than twice as great (see Fig. 5D). Relative to the stimulus in Fig. 5A, the additional spectral information in Fig. 5B makes the left and right panels more consistent with the past experience of red and blue illumination, respectively.

Decreasing the probability that identical returns arise from different conditions of chromatic illumination decreases perceived color differences. If the perceived color difference of two identical targets is increased by a constellation of spectral returns consistent with the targets being under different chromatic illuminants, then spectral information surrounding the targets that is inconsistent with this experience should decrease the color differences of the targets.

To test this prediction, subjects were presented with the stimuli in Fig. 5C. In this case, the constellation of spectral returns was chosen so that the left and right panels were less likely to be under “red” and “blue” illumination than those in Fig. 5B. Thus, the spectra of the “red” squares in the left array were shifted to the left (toward the central part of the spectrum), and their variance in y was decreased (i.e., the bandwidth of the distribution was broadened), whereas the variance of the “yellow” and “green” squares was increased (i.e., the bandwidths of the distributions were narrowed) and their spectral power was increased. These changes are inconsistent with the past experience of the visual system with “red” illumination, which would

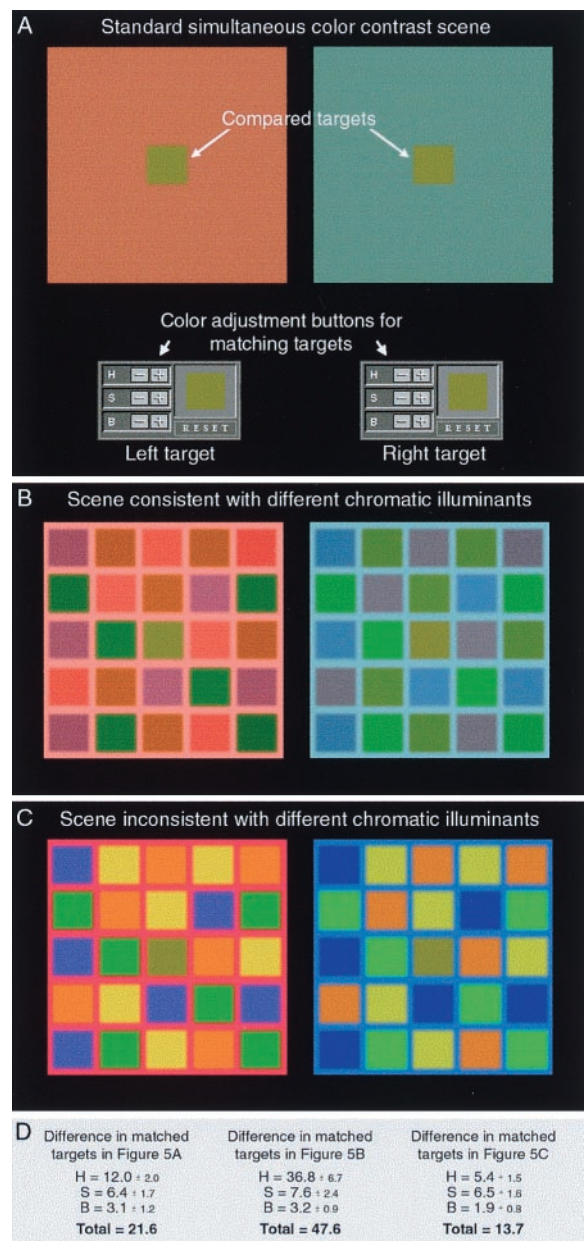


Fig. 5. Comparison of standard color contrast (A) and more complex scenes which provide spectral cues that either increase (B) or decrease (C) the probability that the two arrays are differently illuminated. Using the “buttons” provided, the subjects’ task was to adjust the apparent hue (H), saturation (S), and brightness (B) of the targets in comparison boxes below the stimulus until the colors in the comparison boxes matched the apparent colors of the corresponding targets in the test scene. The adjusted targets were presented below the scenes to preserve the empirical significance of the arrays. The differences in these adjustments are shown in D. The values given represent the average responses of the two authors and eight naive subjects to the stimuli indicated; all differences are statistically significant ($P < 0.01$).

have decreased the luminance and broadened the distributions of the returns from “green” and “yellow” squares, at the same time narrowing the distribution of the returns from “red” squares (see Fig. 3A). Similarly, the right array in Fig. 5C was made less consistent with a “blue” illuminant by both increasing the luminance and narrowing the distribution from the “yellow” and “orange” squares and decreasing the luminance of the “blue” squares. Again, both changes are opposite the spectral

differences that would be generated if the array were under a “blue” illuminant.

As a result of these manipulations, perceptions of target hue, saturation, and brightness were more similar in response to the stimuli in Fig. 5C than to those in Fig. 5B (see Fig. 5D). Indeed, the targets appeared more similar in Fig. 5C than when presented on the uniform surrounds in Fig. 5A, showing that the perceptual differences in response to the stimuli in Fig. 5A and B do not arise from structural differences in the complexity of the surrounds. Nor can these results be explained by differences in the spectral profile of the stimulus, because the spatial chromatic average is the same in Fig. 5A–C. Finally, the reduction in color contrast in Fig. 5C cannot be attributed to the increased saturation of the red and blue grid immediately surrounding the targets because, according to conventional theories of color contrast, such a change should have increased, not decreased, the effects of color contrast.

Thus in each of these challenges, the perceived colors of the targets changed according to whether the stimulus was more consistent with the targets being differently reflective objects under different illuminants or similarly reflective objects under similar illuminants.

Discussion

The spectral composition of visual stimuli is determined by both the reflectance of objects and their illumination (as well as a host of other more subtle factors not considered in this study). Consequently, the underlying sources of the light that interacts with retinal receptors are inevitably ambiguous. Successful behavioral responses to spectral stimuli nonetheless depend on a proper evaluation of the relative contributions to the stimulus of reflectance and illumination: the response will be inappropriate if the contribution from illumination is mistaken for the contribution of surface reflectances, or *vice versa*. Because an observer cannot determine these relative values from the stimulus as such, we have proposed that the visual system resolves this dilemma by seeing colors wholly according to the past significance of spectral returns (8, 29). The present study has sought to extend this general hypothesis of visual perception to the perception of color contrast.

To this end, we first determined the physical relationships between reflectances, illuminants, and spectral returns and compared these relationships with the effects of placing targets in a range of different chromatic contexts (Fig. 3). The complementary relationship between the physical rules governing reflectance, illumination, and the resulting spectral returns and the perceptual phenomenology of color contrast supports the conclusion that color contrast effects are an empirical consequence of the way the spectral returns from surfaces change with illumination. Consistent with this conclusion, changing the empirical significance of identical targets alters the perception of color contrast in the expected directions. To the extent that a color contrast stimulus was consistent with the returns from the targets arising from differently reflective objects under different illuminants, the perceived difference in color between two identical targets increased (compare Fig. 5A and B). Conversely, to the extent that a color contrast stimulus was inconsistent with this possibility (and therefore more consistent with the source being similarly reflective surfaces under similar illuminants), the difference in target color was decreased (cf. Fig. 5B and C). These different perceptions of identical targets in stimuli with the same average chromatic values suggest that “illusions” of color contrast are based on past experience. If color contrast were a consequence of (i) lateral interactions between chromatically sensitive neurons, (ii) the adaptation of retinal receptors to the predominant wavelengths surrounded the targets, or (iii) chromatic ratios across contrast boundaries in the scene, the

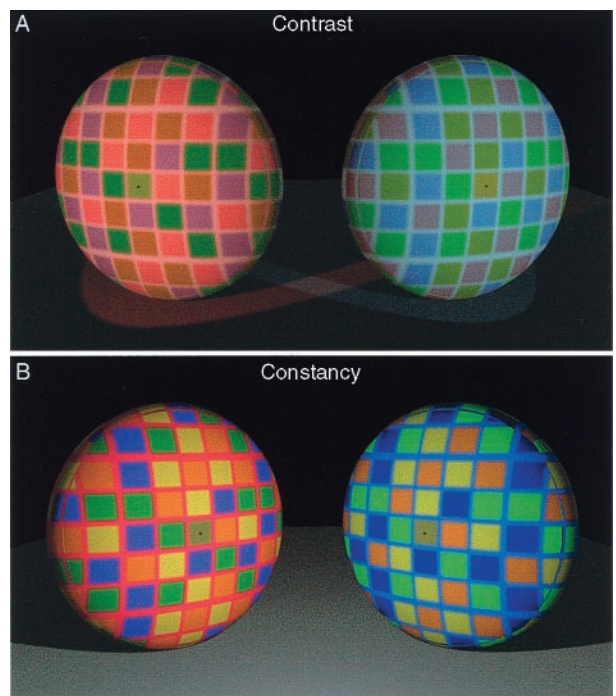


Fig. 6. The similar empirical basis of color contrast (A) and constancy (B). (A) Although the spectral returns from the central squares (indicated by the black dots) are identical, they elicit different sensations of color because the stimuli increase the probability that the two returns originate from differently reflective surfaces under different illuminants. (B) When, however, the spectral information in the scene increases the probability that the returns from the targets originate from similarly reflective surfaces under similar illumination, the central squares (indicated by black dots) elicit relatively similar sensations of color. Thus, despite the fact that the spatial average of the return from the discs surrounding the targets is the same in A and B, the color of two targets appears either different (color contrast) or similar (color constancy), according to the empirical significance of the stimuli.

appearance of the two identical targets in these stimuli should always have been the same.

These results speak not only to the basis of color contrast, but to the seemingly opposite phenomenology of color constancy. Color contrast effects have long been considered illusions arising from mechanisms aimed at allowing the observer to see more or less the same object colors in different illuminants [or, for some investigators, as a means of maximizing visual differences (15, 30)]. The evidence we present here, however, suggests that these contextual effects on color perception are neither illusory nor remedial, but the signature of a basic strategy of color vision, namely to generate chromatic percepts empirically according to how often in past experience the stimulus has signified a particular combination of reflectance and illumination. As a result, a stimulus that has more often signified different surfaces under different illuminants enhances color contrast (Fig. 6A), whereas a stimulus that has more often signified similar surfaces under similar illumination diminishes color contrast (Fig. 6B).

More specifically, color contrast effects are generated if the stimulus (as in Fig. 6A) increases the probability that the underlying sources of the two identical target returns (central squares of the discs in Fig. 6A, indicated with a black dot) are differently reflective surfaces in different illuminants. In this case, the returns from the target patches appear patently different in color because the significance of the targets provided by the constellation of returns from their surrounds reflexively generates a percept that accords with differently reflective surfaces under different lights. Conversely, if the stimulus is

more consistent with the underlying sources of the returns from two targets being similarly reflective surfaces under similar illuminants (as is the case for the central squares of the discs in Fig. 6B, indicated by black dots), then the perceptual response elicited will be the same or a more similar color (i.e., color constancy). Thus, the distinction between color constancy and color contrast in terms of two targets returning different spectra to the eye (but appearing the same) or similar spectra to the eye (but appearing different) is meaningless. What distinguishes constancy from contrast is simply the similarity or difference of the empirical significance of the targets in the stimulus, not their physical attributes. When considered in these empirical terms, color constancy and contrast are simply manifestations of the same probabilistic strategy for generating perceptions of color.

Conclusion

This same general argument (and much evidence) has recently been invoked to explain other aspects of visual perception,

including simultaneous brightness contrast (31, 32), the effects of color on brightness (8), Mach bands (33, 34), the filling in associated with Cornsweet edges (35), the Chubb illusion (36), and the perception of oriented lines (37). In each case, the perceived visual qualities appear to be manifestations of what the stimuli have typically turned out to be. With respect to color, this strategy implies that perception will always accord with the frequency of the underlying stimulus sources, and that manipulating chromatic stimuli will change the colors perceived in a manner predicted by the empirical significance of the altered stimuli.

We are grateful to Mark Changizi for his participation in the color matching experiments; to Sid Simon for technical advice; and to Mark Changizi, David Fitzpatrick, Surajit Nundy, Tom Polger, and Zhiyong Yang for helpful criticisms. This work was supported by the National Institutes of Health (NS29187).

1. da Vinci, L. (1651) *The Notebooks of Leonardo da Vinci*, ed. Richter, I. A. (1952) (Oxford Univ. Press, Oxford, U.K.), p. 136.
2. Chevreul, M.-E. (1838) *The Principles of Harmony and Contrast of Colours, and Their Applications to the Arts*, trans. Martel, C. (1955) (Longmans, London), 2nd Ed.
3. Hering, E. (1864) *Outlines of a Theory of the Light Sense*, trans. Hurvich, L. M. & Jameson, D. (1964) (Harvard Univ. Press, Cambridge, MA).
4. von Helmholtz, H. L. F. (1924) *Helmholtz's Treatise on Physiological Optics*, trans. Southall, J. P. C. (1924–1925) (The Optical Society of America, Rochester, NY), Vols. I–III.
5. Evans, R. M. (1948) *An Introduction to Color* (Wiley, New York).
6. Land, E. H. (1986) *Vision Res.* **26**, 7–21.
7. Brainard, D. H. (1999) in *The Science of Color*, ed. Shevell, S. (The Optical Society of America, Washington, DC), 2nd Ed., in press.
8. Lotto, R. B. & Purves, D. (1999) *Nat. Neurosci.* **2**, 1010–1014.
9. Palmer, S. E. (1999) *Vision Science: Photons to Phenomenology* (MIT Press, Cambridge, MA).
10. Daw, N. W. (1967) *Science* **158**, 942–944.
11. Land, E. H. & McCann, J. J. (1971) *J. Opt. Soc. Am.* **61**, 1–11.
12. De Valois, K. K. (1977) *Vision Res.* **17**, 1057–1065.
13. Jameson, D. & Hurvich, L. (1989) *Annu. Rev. Psychol.* **40**, 1–22.
14. Creutzfeldt, B., Lange-Malecki, B. & Dreyer, E. (1990) *J. Opt. Soc. Am.* **7**, 1644–1653.
15. Zaidi, Q., Yoshimi, B., Flanigan, N. & Canova, A. (1992) *Vision Res.* **32**, 1695–1707.
16. Hurlbert, A. (1996) *Curr. Biol.* **6**, 1381–1284.
17. von Kries, J. (1905) in *Handbuch der Physiologie des Menschen*, ed. Nagel, W. (Vieweg, Braunschweig, Germany), pp. 109–282.
18. D'Zmura, M. & Lennie, P. (1986) *J. Opt. Soc. Am.* **3**, 1662–1672.
19. Creutzfeldt, B., Crook, J., Kastner, S., Li, C.-Y. & Pei, X. (1991) *Exp. Brain Res.* **87**, 3–21.
20. Creutzfeldt, B., Kastner, S., Pei, X. & Valberg, A. (1991) *Exp. Brain Res.* **87**, 22–45.
21. Chichilnisky, E.-J. & Wandell, B. A. (1995) *Vision Res.* **35**, 239–254.
22. Walsh, V. (1995) *Curr. Biol.* **5**, 703–705.
23. Webster, M. A. & Mollon, J. D. (1995) *Nature (London)* **373**, 694–698.
24. Ware, C. & Cowan, W. B. (1982) *Vision Res.* **22**, 1353–1362.
25. Jenness, J. W. & Shevell, S. K. K. (1994) *Vision Res.* **35**, 797–805.
26. Brown, R. O. & MacLeod, D. I. A. (1997) *Curr. Biol.* **7**, 844–849.
27. Kraft, J. M. & Brainard, D. H. (1999) *Proc. Natl. Acad. Sci. USA* **96**, 307–312.
28. Burton, G. J. & Moorhead, I. R. (1987) *Appl. Opt.* **26**, 157–170.
29. Purves, D., Lotto, R. B., Williams, S. M., Nundy, S. & Yang, Z. (2000) *Philos. Trans. R. Soc. London B*, in press.
30. Hardin, C. L. (1988) *Color for Philosophers* (Hackett, Indianapolis).
31. Williams, S. M., McCoy, A. N. & Purves, D. (1998) *Proc. Natl. Acad. Sci. USA* **95**, 13301–13306.
32. Williams, S. M., McCoy, A. N. & Purves, D. (1998) *Proc. Natl. Acad. Sci. USA* **95**, 13296–13300.
33. Lotto, R. B., Williams, S. M. & Purves, D. (1999) *Proc. Natl. Acad. Sci. USA* **96**, 5239–5244.
34. Lotto, R. B., Williams, S. M. & Purves, D. (1999) *Proc. Natl. Acad. Sci. USA* **96**, 5245–5250.
35. Purves, D., Shimpi, A. & Lotto, R. B. (1999) *J. Neurosci.* **19**, 8542–8551.
36. Lotto, R. B. & Purves, D. (2000) *J. Cog. Neurosci.*, in press.
37. Nundy, S., Lotto, R. B., Coppola, D., Shimpi, A. & Purves, D. (2000) *Proc. Natl. Acad. Sci. USA* **97**, 5592–5597.