

Mechanisms of color constancy under nearly natural viewing

(color appearance/adaptation/natural scenes/retinex)

J. M. KRAFT* AND D. H. BRAINARD

Department of Psychology, University of California, Santa Barbara, CA 93106

Communicated by Robert Merrill Boynton, University of California, San Diego, CA, November 6, 1998 (received for review August 13, 1998)

ABSTRACT Color constancy is our ability to perceive constant surface colors despite changes in illumination. Although color constancy has been studied extensively, its mechanisms are still largely unknown. Three classic hypotheses are that constancy is mediated by local adaptation, by adaptation to the spatial mean of the image, or by adaptation to the most intense image region. We measure color constancy under nearly natural viewing conditions, by using a design that allows us to test these three hypotheses directly. By suitable stimulus manipulation, we are able to titrate the degree of constancy between 11% and 83%, indicating that we have achieved good laboratory control. Our results rule out all three classic hypotheses and thus suggest that there is more to constancy than can be easily explained by the action of simple visual mechanisms.

As a computational task, seeing is difficult because there is no simple relation between object properties and the retinal image (1, 2). In the case of color, the difficulty arises because the light reflected from a surface depends both on its reflectance and on the illumination (3). The term color constancy refers to our ability to perceive stable surface colors despite changes in illumination. Although color constancy has been studied extensively (4–12), the mechanisms that mediate it are not yet well-understood.

Color constancy can be linked to chromatic adaptation (5, 13–15), wherein the visual system adjusts its sensitivity to a light according to the context in which the light appears. To explain how chromatic adaptation could mediate constancy in complex scenes, various theorists have suggested that the state of adaptation at an image location is set by its local surround (16–19), by the spatial average of the image (20–23), or by the most intense image region[†] (6, 24–26). It remains uncertain, however, whether any of these ideas can actually account for constancy under natural viewing conditions or whether more complex explanations are required. In our view, the fundamental open question of color constancy is what aspects of complex images govern the visual system's sensitivity. Recent experiments with simple stimuli indicate that the classic ideas described above may not be sufficient (9, 14, 27–30).

Herein we present experiments that measure color constancy for moderately complex images. Our experiments are designed to test explicitly how well adaptation to the local surround, spatial mean, or most intense image region can account for constancy. Two features of our design are worth noting. (i) Constancy experiments usually use simple abstract stimuli or simulations of uniformly illuminated surfaces. Although these stimuli can be controlled easily, they do not embody the complexity of the visual environment. We used stimuli consisting of actual illuminated surfaces, con-

figured in three dimensions. By using rich stimuli, we hoped to study constancy as it operates under natural viewing. (ii) In a typical constancy experiment, the (usually simulated) surfaces are held constant across illumination changes. Such a design eliminates the illuminant-surface ambiguity that makes achieving constancy a difficult computational task. Indeed, when the surfaces are held fixed, most models of constancy make similar predictions. In our experiments, we independently manipulate both surfaces and illuminants. This design allows us to change the illuminant while silencing the action of specific mechanisms and thus permits strong tests.

To understand our logic, consider the local surround. In general, surrounding a test stimulus with a background reduces the amount of the background's hue perceived in the test (31, 32). If we hold a collection of surfaces fixed and add long-wavelength light to the illuminant, the light reflected to an observer from all the surfaces will contain more power at long wavelengths. In the absence of any adjustment by the visual system, all the surfaces would appear redder. Adaptation to the local surround, however, tends to counteract the effect of the physical change and help maintain stable object color appearance. To test whether adaptation to the local surround can fully explain constancy, we measured color appearance under conditions where we varied both the illumination and the surface collection to equate the local surround of a test patch. If the local surround is the key aspect of the image, then constancy would be obliterated for these conditions because the light from the surround does not change with the illuminant.

METHODS

Overview. In our experiments, observers looked into a chamber in which the spectrum of the illuminant and the spectral reflectance of all visible surfaces could be independently controlled (Fig. 1). The appearance of a test patch in the chamber could be adjusted through the use of a projection colorimeter. The observers' task was to adjust the chromaticity of this test patch until it appeared achromatic (somewhere on the continuum from black to gray to white) (9, 12, 33–36). The difference between the achromatic settings across an illuminant change indicates how well the observer adjusted to the difference between the illuminants. Our methods are similar to those described in detail elsewhere (12).

Abbreviation: CIE, Commission Internationale de l'Eclairage.

*To whom reprint requests should be addressed.

[†]In this discussion and throughout, we do not explicitly distinguish ideas, and experiments related to achromatic color (i.e., studies of lightness constancy) from those related to full color. The full color version of adaptation to the most intense image region is typically cast in terms of three image regions, each the most intense as seen by one class of cone photoreceptors. In this paper we will use the phrase "most intense image region" as a shorthand for the more general concept.

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.

© 1999 by The National Academy of Sciences 0027-8424/99/96307-6\$2.00/0
PNAS is available online at www.pnas.org.




FIG. 1. Experimental stimuli. Each picture is a view of the experimental chamber as configured for one experimental condition. (*Left*) Neutral-illuminant conditions. (*Right*) Orange-red, pale-red, and yellow-illuminant conditions for the three experiments. The location of the test panel is indicated in the upper left picture. In the pictures for the local surround and spatial mean experiments, the projection colorimeter was turned off. Thus in these pictures, the light reflected from the test panel indicates the change in illuminant. In the pictures for the maximum flux experiment, the projector was turned on and set so as to equate the tristimulus coordinates of the light reaching the observer from the test patch. The pictures shown were rendered from carefully calibrated hyperspectral images of the experimental stimuli (see *Appendix*).

We present three experiments. In the local surround experiment, we equated the light reflected from the background surrounding the test patch across two conditions

having different illuminants. In the first condition, the background of the chamber was gray cardboard and it was illuminated by a neutral light. In the second condition, the

background was blue cardboard and it was illuminated by an orange-red light.[‡] Other objects were also visible in the apparatus (see Fig. 1).

In the spatial mean experiment, we equated the mean light from the entire apparatus across two conditions having different illuminants. The surfaces in the chamber were the same in this experiment as for the local surround experiment, but the orange-red illuminant was replaced by a pale-red illuminant.

In the maximum flux experiment, we equated the light reflected from a frame surrounding the test patch across two conditions having different illuminants. In each condition, the light reflected from the frame caused the greatest number of quantal absorptions in all three cone classes. In the first condition, a yellow frame was illuminated by a neutral light, and in the second condition, a magenta frame was illuminated by a yellow light. The rest of the apparatus was lined with dark gray cardboard in both conditions.

Observers. Four observers participated in the experiments. Two were naive (EAH, a female, age 19 years; MTR, a male, age 20 years), and two were informed (DHB, a male, age 37 years; PBE, a male, age 33 years). All observers were color normal as defined by results from the American Optical Company H-R-R (Hardy, Rand, and Rittler) Pseudoisochromatic plates and the Ishihara color plates.

Stimuli. The experimental chamber (102 cm high \times 70 cm wide \times 73 cm deep) was illuminated by red, green, and blue theater stage lamps. The intensity of the three lamps determined the chromaticity and luminance of the illumination. Details of the illuminant hardware and control software are given elsewhere (11). The chamber was lined with matte cardboard that could be changed between sessions. The test patch (10 cm wide \times 15 cm high; $6.3^\circ \times 10.1^\circ$) was made of dark gray Munsell paper [Local Surround and Spatial Mean experiments, N 3/; Maximum Flux experiment, N 2.5/; the letter "N" indicating neutral, and the number indicating lightness on a scale from 1 (black) to 10 (white)].

The test patch was spot illuminated by a projection colorimeter (11) that controlled the chromaticity and luminance of the light reaching the observer from the patch. To minimize the visibility of the projected light, the test patch was surrounded by a thin black border (5 mm; 0.3°). In the maximum flux experiment, the border was reduced in size (1.6 mm; 0.1°) and was barely visible. Control experiments conducted in our laboratory indicate that the presence of such a border does not materially affect the appearance of the test patch for conditions similar to ours (12).

A variety of objects were placed in the apparatus in the local surround and spatial mean experiments to provide cues to the illuminant: a Macbeth Color Checker, a cylinder covered in wrinkled aluminum foil, three objects made from gray cardboard, and one wall lined with gray cardboard.

Procedure. Observers viewed the test patch through a rectangular aperture in the front of the chamber. In each session, observers made four achromatic settings at each of two luminances by adjusting knobs that controlled the CIE (Commission Internationale de l'Éclairage) Lab a^* and b^* chromaticity coordinates of the stimulus. A short break separated two blocks within a session. Each observer made settings in two sessions per condition. Observers were encouraged to look around the apparatus before finalizing each setting.

[‡]For descriptive clarity, we use color names to describe physical stimuli. For illuminants, these names correspond to the appearance of the illuminant reflected from a nonselective surface and seen in isolation. For surfaces, these names describe the appearance when the surface was purchased in an art supply store. Table 1 gives the chromaticities and luminances of our illuminants, as well as the descriptive color names.

To ensure that the test patch itself did not provide a cue to the illuminant, the starting chromaticity of the test for the first adjustment in each block was chosen at random from a region of color space centered on the chromaticity of the local surround (local surround and maximum flux experiments) or on the chromaticity of the spatial mean of the image (spatial mean experiment). On subsequent adjustments in a block, the starting chromaticity was chosen as a function of the preceding settings (12). In each experiment, this procedure equates the average initial chromaticity of the test patch for the two experimental illuminants.

RESULTS

Human color vision is trichromatic, so a set of three-dimensional tristimulus coordinates can specify a light completely. If the absolute intensity of a light can be neglected, then only two-dimensional chromaticity coordinates are needed. We plot our results by using the standard (two-dimensional) 1931 CIE chromaticity diagram (37). Points on this diagram represent the relative spectral power distribution of the light reflected to the observer, taking human trichromacy into account.

Local Surround. The Fig. 2 *Top* shows data for one observer in the local surround experiment. The triangle to the left indicates the mean (across session) achromatic setting under the neutral illuminant and the triangle to the right indicates the mean achromatic setting under the orange-red illuminant. The two corresponding illuminant chromaticities are shown by circles.

To interpret our data in terms of constancy, it is helpful to consider how the achromatic settings would behave for a color constant visual system. (i) Note that for such a system, the same surface should appear achromatic under both illuminants. (ii) Note that the chromaticity of the light reflected to the observer from a fixed surface varies with the illuminant. Thus, for a color constant system, we expect the achromatic settings to shift with the illuminant. Indeed, good constancy is indicated when the shift in achromatic chromaticity has approximately the same direction and magnitude as the shift in illuminant chromaticity (12). In the data for the local surround experiment, the achromatic chromaticity shifts in the same direction as the illuminant chromaticity, but the magnitude of the shift is considerably less. This result indicates partial but not complete color constancy.

It is useful to summarize our data with a single number that indexes the degree to which the achromatic settings indicate constancy. To do so, we use a constancy index for which zero indicates no constancy and one indicates perfect constancy. In essence, our index compares the shift in achromatic chromaticity with the shift in illuminant chromaticity. Its actual computation is described in detail elsewhere (12).[§] The constancy indices for the three data sets shown in Fig. 2 are provided in the legend.

Fig. 3 provides the constancy indices for all four observers in the local surround experiment. The mean index is 0.53, significantly greater than zero ($t = 7.85$; one-tailed $P < 0.01$). Because the design of the experiment silences mechanisms of local color contrast, the data indicate that other mechanisms must play a role in human color constancy. This conclusion is in accord with work in the achromatic domain (27, 38) and recent full-color experiments (12, 39).

Spatial Mean. The Fig. 2 *Middle* shows data for one observer in the spatial mean experiment. Again the results indicate partial color constancy, although in this experiment the effect

[§]The computation of the index requires choosing one of the experimental illuminants as a standard. The indices we report are the average of those obtained by choosing each of the two experimental illuminants as the standard.

is smaller than for the local surround experiment. The mean constancy index for the spatial mean experiment is 0.39. The data indicate constancy significantly greater than zero ($t =$




FIG. 2. Experiment results. Illuminant (circles) and achromatic (triangles) chromaticities are shown for one observer for each of the three experiments. For each experiment, the figure shows data for the observer whose constancy index was closest to the mean index. The SEM (computed across sessions) for each datum are smaller than the plotted point. Local surround: Observer DHB, constancy index = 0.45; spatial mean, DHB, 0.29; maximum flux, MTR, 0.33.

3.93; one-tailed $P < 0.05$) and thus reject adaptation to the spatial mean as the sole mechanism for constancy. This conclusion is in agreement with previous qualitative observations (25).

Note that in this experiment, adaptation to the local surround opposed color constancy. This is because the chromaticity shift of the surround was opposite to the shift of the illuminant. Thus the experiment rules out any reasonable possibility that the joint action of local contrast and adaptation to the spatial mean mediate human color constancy (40). In addition, our results cannot be easily explained by assuming that the visual system uses a small pooling area for computation of the spatial mean, because the spatial mean of image subregions centered on the test were generally biased against constancy (see Appendix).

Maximum Flux. The Fig. 2 Bottom shows data for one observer in the maximum flux experiment. The data do indicate that the intense yellow surround has a considerable effect. The achromatic chromaticity under the neutral illuminant is shifted relative to its position in the other experiments. At the same time, the data reveal significant constancy, as the achromatic chromaticities under the two illuminants do not superimpose. The mean constancy index for this experiment is 0.33, and the data reject adaptation to the most intense image region as the sole mechanism for constancy ($t = 14.14$; one-tailed $P < 0.001$).

Also note that mechanisms of local contrast were equated in this experiment, as we constructed the surface with maximum flux in the area immediately surrounding the test patch. Thus the data also rule out the possibility that the local surround and most intense image region jointly mediate human color constancy.

Controls. Constancy measured in our main experiments was substantial but by no means complete. In these experiments, we intentionally manipulated the stimuli to silence several potential constancy mechanisms. In a control experiment, we established that when these mechanisms are not silenced, constancy is considerably more complete. In this first control, we used the traditional design: we lined the chamber with the same gray cardboard in both conditions and compared achromatic settings made under the neutral and orange-red illuminants. For this experiment, mechanisms sensitive to the local surround, spatial mean, and most intense image region all support constancy. The mean constancy index for this control experiment was 0.83, significantly more than that observed in our main experiments (matched-pairs, one-tailed t-tests: local surround, $t = 5.96$ and $P < 0.01$; spatial mean, $t = 5.07$ and $P < 0.01$; maximum flux, $t = 23.71$ and $P < 0.001$). This lends




FIG. 3. Constancy indices. The figure shows the constancy indices for individual observers in each experiment.

support to the idea that the classic mechanisms do play some role in supporting human color constancy.

To show that the constancy observed in our experiments was not artifactual, we ran a second control experiment with a simplified stimulus. In each condition of this experiment, we lined the chamber with a single background surface. In one condition, we used the gray cardboard and the neutral illuminant. In the other condition, we used the blue cardboard and the orange-red illuminant. As in the local surround experiment, the light reflected to the observer from the background near the test was identical in the two conditions. All other objects were removed from the chamber, so that the two conditions presented similar stimuli to the observer. The mean constancy index for this control experiment was 0.11, significantly less than those for our three main experiments (matched-pairs, one-tailed *t* tests: local surround, *t* = 5.49 and *P* < 0.01; spatial mean, *t* = 2.59 and *P* < 0.05; maximum flux, *t* = 7.61 and *P* < 0.01).

Although the stimuli in the two conditions of our second control experiment were matched in the neighborhood of the test patch, they differed elsewhere. This is because the spatial distribution of illumination within the chamber is affected by light reflected from the background back into the chamber, an effect referred to as mutual illumination. It has been suggested that mutual illumination can provide a cue for color constancy (41). Consistent with this idea, the small residual constancy measured in our control experiment was significantly greater than zero (*t* = 12.77; one-tailed *P* < 0.001), a result in accord with observations by Gilchrist for similarly reduced conditions (42).

DISCUSSION

Our experiments show that in their standard form, classic theories cannot account for human color constancy under natural viewing conditions. In three experiments, we silenced the contribution of adaptation to the local surround, to the spatial mean, and to the most intense image region. Each of these experiments showed substantial residual constancy.

The power of our experiments arises because we manipulate both the surfaces and illuminants. By covarying these two factors, we can separate how the visual system adjusts to changes in illumination from the action of various mechanisms of adaptation. Indeed, across our entire set of experiments, we are able to titrate the degree of constancy from 11% to 83% by manipulating the information content of the stimuli. Thus our preparation allows for the systematic experimental study of constancy in a way that has not previously been possible.

The stimuli in our experiments are actual illuminated surfaces seen in three-dimensional configurations. Thus our results show that the stimulus arrangements necessary to reveal effects beyond those mediated by classic mechanisms can be readily achieved under fairly natural viewing conditions.

If the classic theories cannot account for constancy, how can it be explained? Others have recently observed that under simple viewing conditions, color appearance is affected by the range of luminances and chromaticities present in the image (9, 14, 28–30). One reasonable interpretation of our results is that, as suggested by Webster and Mollon (14), the mechanisms underlying these phenomena contribute to constancy, because in our experiments the stimulus range covaried with the illuminant. It might also be possible to construct a theory in which the classic mechanisms acting in concert mediated constancy (40, 43). As noted above, our current results do rule out the possibility that local contrast with adaptation to the spatial mean or local contrast with adaptation to the surface of maximum flux can account for constancy.

An alternative approach is to take seriously the idea that the visual system attempts to use all available information in the scene to form an estimate of the illuminant. As we noted in the

introduction, this is a difficult computational problem, and it is not surprising that the visual system accomplishes it only approximately. At the same time, computational analyses of color constancy have shown that the optimal illuminant estimation requires using more than just the local surround, spatial mean, or surface of maximum cone flux (44). Although we used rich stimuli, our hyperspectral images provide a complete specification. Thus we are positioned to compare directly human performance and that of computational algorithms.

APPENDIX: STIMULUS DETAILS

This appendix provides additional information about the stimuli for each experiment. This information supplements that provided in *Methods* and Table 1. In the *Appendix* and elsewhere in this paper chromaticity (*x y*) is specified with respect to the CIE 1931 colorimetric system and luminance (*Y*) is specified as cd/m².

Local Surround. Achromatic settings were made at 4.8 and 10.8 cd/m². The neutral-illuminant conditions of the local surround and spatial mean experiments are identical except for a slight difference in test luminances. None-the-less, observers DHB, EAH, and MTR made settings for this condition in both experiments. Because we observed no effect of the change in test luminance for these observers, PBE made settings only for one of these two conditions. Thus the test luminances for observer PBE in the neutral-illuminant condition for the local surround experiment were 3.3 and 7.4 cd/m².

Spatial Mean. Colorimetric specifications for the spatial mean of the entire image were (*x y Y*) = (0.33 0.33 3.2) for the neutral-illuminant condition and (0.34 0.33 3.6) for the orange-red-illuminant condition. For smaller rectangular regions centered on the test and subtending 13° by 17°, 19° by 23°, 25° by 29°, and 31° by 35° the local spatial means for the neutral-illuminant condition were (*x y Y*) = (0.33 0.33 3.4), (0.33 0.33 3.3), (0.34 0.33 3.3), and (0.34 0.33 3.2), and for the orange-red-illuminant condition they were (0.32 0.32 3.5), (0.33 0.32 3.5), (0.34 0.32 3.5), and (0.35 0.33 3.7). Colorimetric specifications for the immediate surround were (*x y Y*) = (0.34 0.32 4.2) and (0.29 0.30 3.0), respectively, for the neutral-illuminant and orange-red-illuminant conditions. Achromatic settings were made at 3.3 and 7.4 cd/m².

Maximum Flux. Colorimetric specifications for the most intense image region were (*x y Y*) = (0.42 0.42 14.6) for the neutral-illuminant condition and (0.42 0.42 15.8) for the yellow-illuminant condition. Achromatic settings were made at 2.6 and 3.5 cd/m².

Control 1 and 2. Achromatic settings were made at 4.6 and 10.7 cd/m².

Hyperspectral Camera. To measure the spatial means and maximum flux values of our stimuli, we used a hyperspectral camera (see <http://color.psych.ucsb.edu/hyperspectral/>). The

Table 1. Chromaticities and luminances of our illuminants and local surround areas, as well as descriptive color names for illuminants and corresponding cardboard surfaces

| Experiment | Illuminant name | Card-board name | Illuminant | | | Surround | | |
|----------------|-----------------|-----------------|------------|----------|----------|----------|----------|----------|
| | | | <i>x</i> | <i>y</i> | <i>Y</i> | <i>x</i> | <i>y</i> | <i>Y</i> |
| Local surround | Neutral | Gray | 0.32 | 0.32 | 10.1 | 0.33 | 0.33 | 4.1 |
| Local surround | Orange red | Blue | 0.45 | 0.37 | 34.9 | 0.34 | 0.33 | 4.6 |
| Spatial mean | Neutral | Gray | 0.33 | 0.32 | 11.1 | 0.34 | 0.32 | 4.2 |
| Spatial mean | Pale red | Blue | 0.39 | 0.35 | 20.8 | 0.29 | 0.30 | 3.0 |
| Maximum flux | Neutral | Yellow | 0.33 | 0.33 | 17.4 | 0.42 | 0.42 | 12.8 |
| Maximum flux | Yellow | Magenta | 0.39 | 0.44 | 56.0 | 0.41 | 0.42 | 13.8 |
| Control 1 | Neutral | Gray | 0.31 | 0.33 | 11.2 | 0.33 | 0.34 | 3.8 |
| Control 1 | Orange red | Gray | 0.45 | 0.37 | 40.0 | 0.46 | 0.38 | 12.9 |
| Control 2 | Neutral | Gray | 0.31 | 0.33 | 11.2 | 0.33 | 0.34 | 3.8 |
| Control 2 | Orange red | Blue | 0.44 | 0.37 | 34.8 | 0.33 | 0.33 | 3.9 |

camera provided the full spectrum (400 nm to 700 nm in 10-nm steps) at each image pixel. The camera was based on a scientific grade monochrome charge-coupled device camera (Photometrics PXL, 2000 by 2000 spatial resolution at 12 bits/pixel, direct digital interface, electronic control of temporal integration, liquid cooling) interfaced to a Macintosh host. Each hyperspectral image consists of 31 monochromatic image planes, with each plane acquired through a narrowband (roughly 10-nm bandwidth full-width at half-height) interference filter. The image planes were calibrated by comparing image data with direct spectral measurements (Photo Research PR-650) of a reference surface. Each hyperspectral image was used to compute an XYZ image providing the 1931 CIE XYZ tristimulus values at each pixel. Spatial means were calculated from the XYZ images and then converted to chromaticity and luminance. The images shown in Fig. 1 are based on these rendered versions of hyperspectral images. The hyperspectral images were also used to compute long-, middle-, and short-wave sensitive cone images from which the maximum flux values were calculated (45).

This research was supported by National Institutes of Health Grants EY10016 and EY06745.

- Landy, M. S. & Movshon, J. A., eds. (1991) *Computational Models of Visual Processing* (MIT Press, Cambridge, MA).
- Knill, D. & Richards, W., eds. (1996) *Perception as Bayesian Inference* (Cambridge Univ. Press, Cambridge, U.K.).
- Brainard, D. H., Wandell, B. A. & Chichilnisky, E.-J. (1993) *Curr. Directions Psychol. Sci.* **2**, 165–170.
- Helson, H. (1938) *J. Exp. Psychol.* **23**, 439–476.
- Burnham, R. W., Evans, R. M. & Newhall, S. M. (1957) *J. Opt. Soc. Am.* **47**, 35–42.
- McCann, J. J., McKee, S. P. & Taylor, T. H. (1976) *Vision Res.* **16**, 445–458.
- Arend, L. E. & Reeves, A. (1986) *J. Opt. Soc. Am. A* **3**, 1743–1751.
- Brainard, D. H. & Wandell, B. A. (1992) *J. Opt. Soc. Am. A* **9**, 1433–1448.
- Bauml, K. H. (1994) *J. Opt. Soc. Am. A* **11**, 531–542.
- Lucassen, M. P. & Walraven, J. (1996) *Vision Res.* **36**, 2699–2711.
- Brainard, D. H., Brunt, W. A. & Speigle, J. M. (1997) *J. Opt. Soc. Am. A* **14**, 2091–2110.
- Brainard, D. H. (1998) *J. Opt. Soc. Am. A* **15**, 307–325.
- Jameson, D. & Hurvich, L. (1989) *Annu. Rev. Psychol.* **40**, 1–22.
- Webster, M. A. & Mollon, J. D. (1995) *Nature (London)* **373**, 694–698.
- Kaiser, P. K. & Boynton, R. M. (1996) *Human Color Vision* (Opt. Soc. Am., Washington, DC), pp. 507–522.
- Wallach, H. (1948) *J. Exp. Psychol.* **38**, 310–324.
- Walraven, J., Benzschawel, T. L., Rogowitz, B. E. & Lucassen, M. P. (1991) in *Testing the Contrast Explanation of Color Constancy*, eds. Valberg, A. & Lee, B. B. (Plenum, New York), pp. 369–377.
- Shapley, R. M. (1986) *Vision Res.* **26**, 45–62.
- Valberg, A. & Lange-Malecki, B. (1990) *Vision Res.* **30**, 371–380.
- Buchsbaum, G. (1980) *J. Franklin Inst.* **310**, 1–26.
- Land, E. H. (1983) *Proc. Natl. Acad. Sci. USA* **80**, 5163–5169.
- Land, E. H. (1986) *Vision Res.* **26**, 7–21.
- Brainard, D. H. & Wandell, B. A. (1986) *J. Opt. Soc. Am. A* **3**, 1651–1661.
- Land, E. H. & McCann, J. J. (1971) *J. Opt. Soc. Am.* **61**, 1–11.
- McCann, J. J. (1994) *Proceedings of the 47th IS&T Annual Meeting*, pp. 397–401.
- Cataliotti, J. & Gilchrist, A. (1995) *Percept. Psychophys.* **57**, 125–135.
- Adelson, E. H. (1993) *Science* **262**, 2042–2044.
- Singer, B. & D'Zmura, M. (1994) *Vision Res.* **34**, 3111–3126.
- Jenness, J. W. & Shevell, S. K. (1995) *Vision Res.* **35**, 797–805.
- Brown, R. O. & MacLeod, D. I. A. (1997) *Curr. Biol.* **7**, 844–849.
- Evans, R. M. (1974) *The Perception of Color* (Wiley, New York).
- Albers, J. (1975) *Interaction of color* (Yale Univ. Press, New Haven, CT).
- Helson, H. & Michels, W. C. (1948) *J. Opt. Soc. Am.* **38**, 1025–1032.
- Werner, J. S. & Walraven, J. (1982) *Vision Res.* **22**, 929–944.
- Fairchild, M. D. & Lennie, P. (1992) *Vision Res.* **32**, 2077–2085.
- Chichilnisky, E. J. & Wandell, B. A. (1996) *Visual Neurosci.* **13**, 591–596.
- CIE (1986) in *Colorimetry*, (Bureau Central de la CIE), 2nd Ed.
- Gilchrist, A. L. (1988) *Percept. Psychophys.* **43**, 415–424.
- Kuriki, I. & Uchikawa, K. (1998) *J. Opt. Soc. Am. A* **15**, 2263–2274.
- Helson, H. (1964) *Adaptation-Level Theory: An Experimental and Systematic Approach to Behavior* (Harper & Row, New York).
- Funt, B. V. & Drew, M. S. (1993) *IEEE Transactions On Pattern Analysis And Machine Intelligence* **15**, 1319–1326.
- Gilchrist, A. & Jacobsen, A. (1984) *Perception* **13**, 5–19.
- Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., Spehar, B., Annan, V. & Economou, E. *Psychol. Rev.*, in press.
- Brainard, D. H. & Freeman, W. T. (1997) *J. Opt. Soc. Am. A* **14**, 1393–1411.
- Smith, V. & Pokorny, J. (1975) *Vision Res.* **15**, 161–171.