17

Color and brightness induction: from Mach bands to three-dimensional configurations

Qasim Zaidi

The interplay of lights and objects in natural or man-made settings creates fascinating visual effects for the acute observer (Minnaert, 1993). Many of these effects can be explained satisfactorily by physical processes outside the observer or by the physics of light absorption by photopigments, but some are due to peculiarities of the visual system. To function effectively in a large variety of settings, the human visual system contains neural mechanisms that perform edge enhancement, spectral differencing, contrast induction, adaptation to steady and temporally varying lights, inference of motion and three-dimensional shape, and perceptual constancy. All of these mechanisms can be conceived of as performing particular transformations on the retinal image. These transformations generally fulfill functionally important roles, but they can also lead to blatantly nonveridical percepts, as shown by the literature on visual illusions (e.g., Luckiesh, 1965; Frisby, 1980; Griffiths & Zaidi, 1999).

A particularly important task for an observer is to segregate different objects from each other and from the background. A number of visual cues, including brightness, color, texture, motion, shape, stereo-disparity, and occlusion, facilitate this task. A number of studies have shown that the perceived magnitude of almost every visual modality is influenced by the magnitude of that modality in the surround. In fact, Chevreul's (1839) Law of Simultaneous Contrast of Colors can be safely generalized: "In the case where the observer sees at the same time two contiguous fields, they will appear as dissimilar as possible in almost every modality." Obviously, a visual process that enhances perceived differences between contiguous fields will facilitate object segregation. In a three-dimensional world, some segregations, like figure from ground, confer greater functional advantages than others.

A gray region appears brighter when viewed on a darker enclosing surround and darker when viewed on a brighter surround (Chevreul, 1839). When the surround is variegated in brightness, the perceived brightness of the test region can depend on a multitude of factors. For example, in Fig. 17.1 (Zaidi, Spehar, & Shy, 1997) the gray regions in the centers of the two configurations were made of identical materials and are of identical luminance (a purely physical measure), yet the gray region in the picture on the right appears considerably lighter than the gray region in the picture on the left. Figure 17.1 consists of a photograph of two figurally identical three-dimensional configurations, each consisting of an H shape in the foreground, with the gray square being the horizontal bar between the two vertical bars, and a larger square flap at an angle behind the H. If the perceived difference in the brightness of the two gray patches is due to induced contrast, then on the right the brightness induced by the dark background is of greater magnitude than the darkness induced by the light vertical bars. Similarly, the inducing effect of the light background prevails in the configuration on the left. In both pictures, each gray region has an equal perimeter

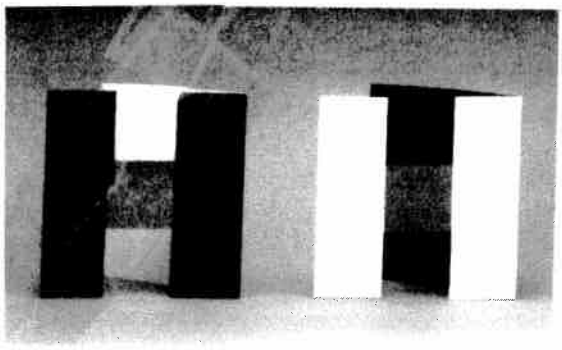

Figure 17.1: Two-dimensional pictures of a three-dimensional configuration. The two gray regions are of equal luminance, so any perceived difference in brightness is due to a combination of induced lightness from dark surrounds and induced darkness from light surrounds. The relative brightness of the two gray regions can be used to infer the strength of the induced effect from the background versus the effect from the flanking surrounds on the sides (Zaidi, Spehar, & Shy, 1997).

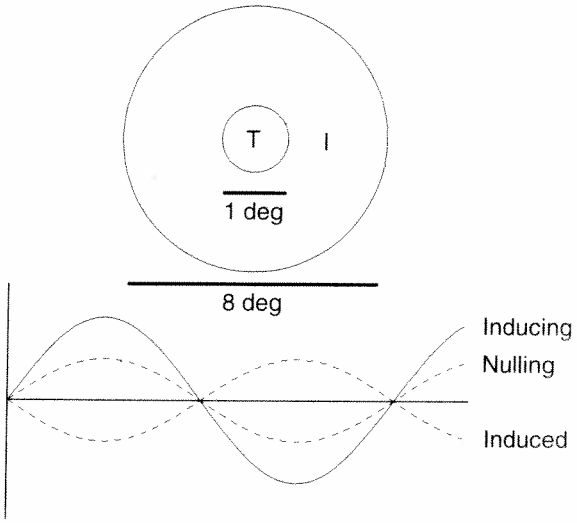

Figure 17.2: Nulling technique for measuring the magnitude of color induction. On the top, I indicates the inducing stimulus and T the test stimulus. On the bottom the solid line marked "Inducing" depicts the variation in time of the color of the inducing annulus. The dashed line marked "Induced" depicts the modulated appearance of the test disk when it is, in fact, not modulated. The dotted line marked "Nulling" depicts the real modulation of the disk required to make it appear steady (Krauskopf, Zaidi, & Mandler, 1986).

flanked by dark and light, and the visible areas of the background flaps are smaller than the visible areas of the vertical flanks. In spite of this, induced contrast from the flaps dominates that from the vertical flanks.

In this paper we will concern ourselves with the neural mechanisms that are responsible for color induction and the manner in which they affect color appearance. The main substrate for color induction is lateral connections between neural elements, most probably inhibitory. These connections range from mechanisms early in the visual stream that are responsible for edge enhancements, to higher level mechanisms that infer three-dimensional shape from two-dimensional cues in retinal images. Color appearance is also influenced by mechanisms of visual adaptation, and we will discuss models and experiments that isolate induction from adaptation effects, and explicitly examine their interaction.

Methods

Most studies on color induction have used measurements based on asymmetric matching or memory, despite Helmholtz's (1924, Vol. 2, pp. 264–271) detailed exposition of the drawbacks of these methods. To overcome these limitations, Krauskopf, Zaidi, and Mandler (1986) introduced a nulling method of measuring the magnitude of induction (Fig. 17.2). When an observer fixates on a white disk, the disk appears greenish when surrounded by a red annulus and reddish when surrounded by a green annulus. The effects of the annulus can be neutralized by adding some red light to the disk in the first case and some green light to the disk in the second. The amount of light needed to make the disk appear white may be taken as a measure of induction. Performing these operations in successive experiments requires the observer to keep a standard white in mind. On the other hand, if the color of the annulus is modulated from red to green, the disk appears to be modulated in counterphase to the annulus. If a modulated component is added to the disk, the disk can be made to appear steady. The nulling modulation can be used as a measure of induction. If induc-

tion is strictly complementary, the modulation added to the disk must be along the same line in color space as that applied to the annulus. Nulling modulations have been measured using methods of adjustment (Krauskopf, Zaidi, & Mandler, 1986; Chubb, Sperling, & Solomon, 1989), 2AFC staircases (Spehar, DeBonet, & Zaidi, 1996; DeBonet & Zaidi, 1997), and 2AFC methods of constant stimuli (Singer & D'Zmura, 1994, 1995). A spatial analogue of the method involving the nulling of induced sinusoidal gratings by superimposed physical gratings was devised by McCourt (1982).

Temporal and spatial nulling methods have the following advantages: (i) During each trial, the observer fixates on the central test. Successive contrast effects are kept to a minimum, because, except for small eye movements, the portion of the retina that receives the induced modulation is never exposed to the inducing field. (ii) The observer is required only to signal the absence or presence of modulation without extracting any subjective qualities of the test. At the null point, the induced effect of all phases of the surround modulation are equally canceled. No memory demands are made. The nulling method thus meets the requirements for Brindley's (1960) class A experiments. (iii) In every trial, the chromaticity and luminance of every point of the display, averaged over a cycle of modulation, can be kept the same for all conditions. The steady-state adaptation is therefore constant throughout the experiments. This enables induction effects on appearance to be separated from adaptation effects.

Neural locus of color induction

For a long time there was a debate concerning the neural locus of color induction. Induction implies that the effects of a stimulus falling on one part of the retina are modified by stimuli falling on another part of the retina at some level in the visual system. For example, it had been supposed that the stimulation of cones of one class in one part of the retina results in the desensitization of cones of the same class in other parts of the retina (e.g., Evans, 1948; Alpern, 1964). Alternatively, a similar form of lateral interaction had been proposed between second-stage opponent mechanisms (e.g., Jameson & Hurvich, 1964; Shevell, 1978; Guth, Massof, & Benzschawel, 1980; Ware & Cowan, 1982). The debate was resolved by Krauskopf, Zaidi, and Mandler (1986), who showed a cortical locus for color induction involving color mechanisms beyond the linear-opponent stage.

They first measured the amplitude of the induced modulation as a function of the amplitude of the inducing modulation for each of the three cardinal axes of color space (Krauskopf, Williams, & Heeley, 1982; Zaidi, 1992).[1] The amplitude of the nulling modulation versus the amplitude of the surround modulation is shown in Figs. 17.3A and B for two observers for each of the three cardinal directions. The lines are the best-fitting cubic polynomials of the form:

$$(1) \quad N(A) = aA - bA^3,$$

where N is the required nulling modulation and A is the surround modulation amplitude. The curve for the luminance direction is the steepest, followed by that for the YV-direction, and that for the RG-direction is the shallowest. This experiment revealed that the nulling amplitudes are sufficiently different functions of inducing amplitudes in the three cardinal directions to allow a critical test of the locus of induction effects. This was carried out in a second experiment, in which the amplitude and direction of the modulation needed to null the effect of an inducing stimulus in various noncardinal directions was measured.

The second experiment was concerned with the question of whether induction takes place solely within independent opponent mechanisms. If this were true, induction would not, in general, result in changes of appearance in a strictly complementary direction. In the stimulus domain, modulation in a direction inter-

[1] A number of equivalent designations have been used for the three cardinal axes, including (Luminance, Constant B or S, Constant R&G or L&M), (L+M+S, L–M, S), amd (Light-Dark, Red-Green, Yellow-Violet). In the text we will use (LD, RG, YV), although some graphs taken from previous publications may use one of the others.

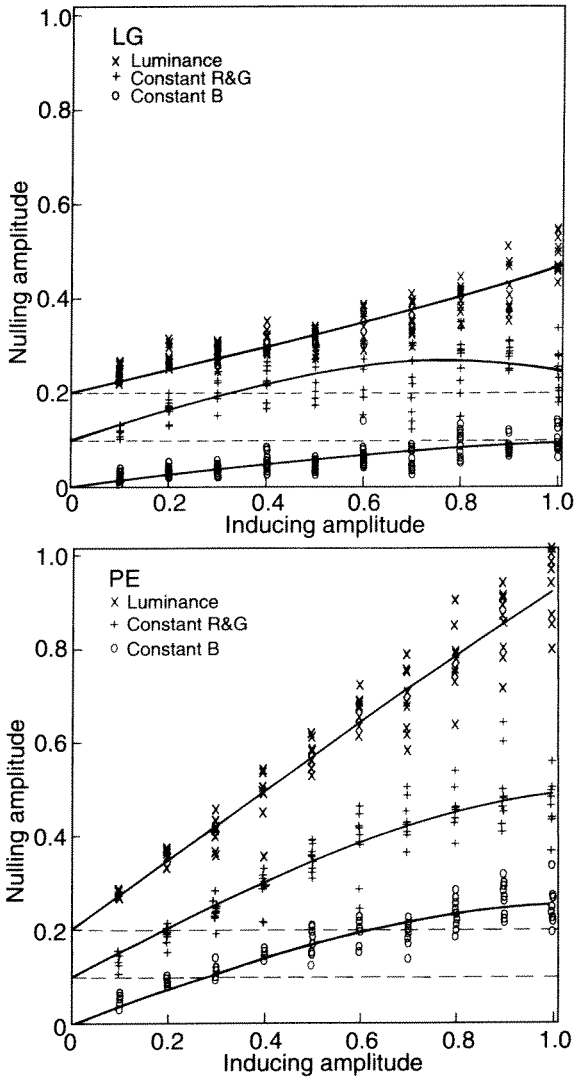

Figure 17.3: Nulling amplitude as a function of inducing amplitude for two observers. In each panel the results from top to bottom are for modulation along the LD-axis, along theYV-axis, and along the RG-axis. Note the displaced ordinates. Lines fitted to the data are best fits of cubic polynomials (Krauskopf, Zaidi, & Mandler, 1986).

mediate between two cardinal directions is equivalent to the vector sum of modulations of the appropriate amplitudes in the two cardinal directions. If induction took place solely within cardinal mechanisms, then the induced effect of this modulation could be nulled by a modulation that simultaneously nulled the induced effect in each of the cardinal directions. However, because the nulling functions are not identical for the cardinal directions, the induced effects in the two directions will be different multiples of the projected amplitudes of the inducing modulation. The predicted nulling modulation will therefore not be in the same color direction as the intermediate inducing direction. If induction took place in higher order color mechanisms that were tuned maximally to the inducing direction, then the best nulling modulation would be in the same direction as the inducing modulation. In forced-choice judgments for six intermediate directions in the isoluminant plane, the null in the inducing direction was consistently preferred to the null predicted from the cardinal directions. For isoluminant stimuli, the results were consistent with the complementary nature of chromatic induction.

The results ruled out interactions at both the cone and opponent mechanism levels as an explanation of induction. The observers preferred the settings in the direction of the inducing stimuli to those predicted from the cardinal-direction settings. Because modulation along the RG-axis varied the input only to the L- and M-cones, whereas modulation along the constant YV-axis only varied the input to the S-cones, explanations in terms of independent cone mechanisms also make the wrong prediction that the vector sum of cardinal-direction settings will be preferred. In fact, since the cardinal directions are linear combinations of cone absorptions, rejection of the vector-summation rule for chromatic induction also ruled out explanations in terms of any other linear combinations of cone absorptions. One possibility is that chromatic induction reflects lateral interactions beyond the stage of independent linear cardinal mechanisms within mechanisms preferentially tuned to many different directions of color space (Krauskopf et al., 1986). Recordings in the parvocellular layers of the macaque LGN have revealed two major classes of cells that are preferentially tuned to the chromatic cardinal axes (Derrington, Krauskopf, & Lennie, 1984), whereas similar experiments on the visual cortex reveal a diversity of cells. The cells that respond to pure chromatic modulation

have maximal responses in many directions in color space, rather than only in the cardinal directions (Lennie, Krauskopf, & Sclar, 1990). It seems likely that lateral interactions between cortical cells cause chromatic induction.

Edge enhancement effects

In a stunning intellectual feat, from phenomenological observations of bright and dark bands at luminance edges where – based on physical considerations alone – none were expected, Mach (1865, 1866a, 1866b, 1868) inferred the retinal mechanisms of lateral inhibition and an excellent approximation of the mathematical form of center-surround receptive fields. These bands enhance the local lightness of the lighter side of the edge and the local darkness of the darker side. Mach bands have also been claimed to exist for chromatic stimuli at saturation edges (Pease, 1978). Consequently, it is worth examining the extent to which they contribute to brightness and chromatic induction (Bekesy, 1968).

The effect of Mach bands on appearance is illustrated nicely by the phenomenon called "grating induction." McCourt (1982) and Foley and McCourt (1985) showed that when a narrow uniform strip is inserted within a sinusoidal grating, an induced grating is perceived in the strip (Fig. 17.4A). Zaidi (1989) created a series of grating configurations to examine the role that proximal and distal elements of the inducing stimulus play in grating induction. Figures 17.4A–C show that gratings of different orientations can be induced by combining Mach bands created at the edges of inducing gratings of the same spatial frequencies and orientations, offset by different fractions of a cycle across the test region. In Fig. 17.4D there is no induced grating, and the Mach bands can be seen clearly. In Figs. 17.5A–D induced gratings of the same orientation and frequency have been created by offsetting inducing gratings of different orientations and spatial frequencies so that the row of pixels directly above the test field is identical in all four figures, as is the row directly below the test field. These local identities lead to similarity in the induced percept, demonstrating the primacy of edge effects in visual grating induction. However, the perceived contrast of the induced grating decreases progressively from Fig. 17.5A–D. This impression was confirmed by nulling measurements that showed that even though Mach bands are confined to local regions of the test, they are affected by more extended regions of the surround. For equiluminant chromatic stimuli, the amplitude of induced modulation was less than for luminance stimuli; otherwise, the effects were qualitatively similar.

In a grating induction display such as Fig. 17.4A, Zaidi and Sachtler (1991) noticed that if the real grating is drifted to the left, the induced grating also appears to drift to the left. If, after prolonged viewing of the drifting stimulus, the motion is abruptly stopped, the real grating appears to move to the right, as would be expected from the classical motion aftereffect (Wohlgemuth, 1911). Simultaneously, in the uniform central gap there appears a grating in counterphase to the surround that also appears to move to the right. It seemed as though motion adaptation could take place in areas of the retina that were not directly exposed to the physical adapting stimulus. To test whether the adaptation was caused by the induced grating per se, Zaidi and Sachtler measured the effect of prolonged viewing of aligned (Fig. 17.4A) and offset (Fig. 17.4D) inducing gratings on contrast thresholds for moving gratings confined to the physically uniform test region. The results showed that the desensitizing effect of offset surrounds was roughly equal to that of aligned surrounds. Therefore, the induced percept of cohesive gratings did not have a greater desensitization effect on test thresholds than did induced Mach bands. For both types of surrounds, the desensitizing effect was highest close to the edges of the adapting stimulus, and gradually decreased with distance from the edge.

Even though Mach bands can be combined to form induced gratings, whether edge enhancing mechanisms are important influences on the appearance of spatially extended tests is a different question. Zaidi, Yoshimi, and Flannigan (1991) presented data that bear on this question by separating the effects of area,

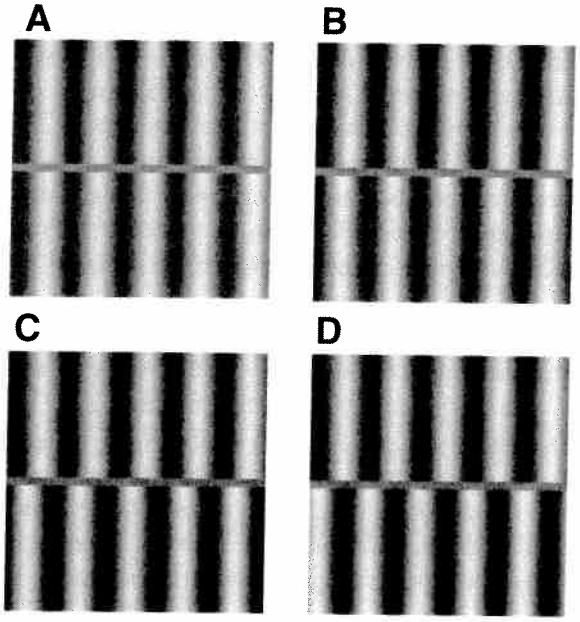

Figure 17.4: Displays showing sinusoidal grating induction in the central homogeneous test field. The orientation of the induced grating depends on the phase of the inducing grating above the test field relative to the phase of the grating below the test field. The relative shifts are (A) 0.0 cycles (in-phase), (B) 0.125 cycles, (C) 0.25 cycles, and (D) 0.5 cycles (counterphase). In (D) the induced percept consists of light and dark patches. Note: Mask inducing gratings to see actual homogeneity of test patch (Zaidi, 1989).

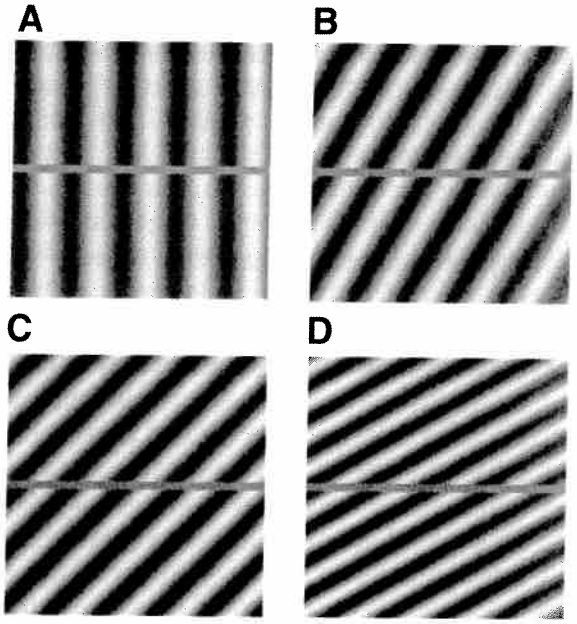

Figure 17.5: Displays showing induced gratings with the same orientation and spatial frequency. The inducing gratings have orientations of (A) 90 deg, (B) 60 deg, (C) 45 deg, and (D) 30 deg. The spatial frequency of the inducing gratings increases from (A) to (D) (Zaidi, 1989).

shape, and perimeter on the change in perceived color of a test region. The test regions were multiple-lobed shapes like those in Fig. 17.6A. The purposes of this experiment required tests whose shapes, areas, and perimeters could be varied independently. In addition, it was considered desirable to use shapes with smooth contours and without sharp corners. The analytic expression for Fourier descriptors provided by Zahn and Roskies (1972) proved to be suitable for generating the test shapes. The number of lobes was 3, 4, 5, 6, 7, 8, 9, 10, 20, or 40. The area of each test was equal to the area of a circle of radius 1 deg, that is, equal to π deg^2. The length of the perimeter of a test was 2.5π, 3.5π, 4.5π, or 5.5π deg. The test field was surrounded by a disk with an outer diameter of 5 deg. The nulling method described above was used to measure induction in the three cardinal directions. Since edge enhancements occur on the border between the test and the surround, for tests subtending equal areas and having comparable shapes, if the magnitude of induced contrast was proportional to the perimeter length, it would indicate that appearance was affected by the summation of Mach bands along the interior perimeter of the test. The results (Figs. 17.6B–D) showed that the magnitude of induced contrast does not depend on the length of the perimeter of the test. In all three panels the slope of the best-fitting regression line was not significantly different from zero. Hence, local edge enhancements do not have a significant effect on the appearance of spatially extended regions.

Spatially extended effects

The spatial extent of color induction has generally been studied with disk–annulus configurations of dif-

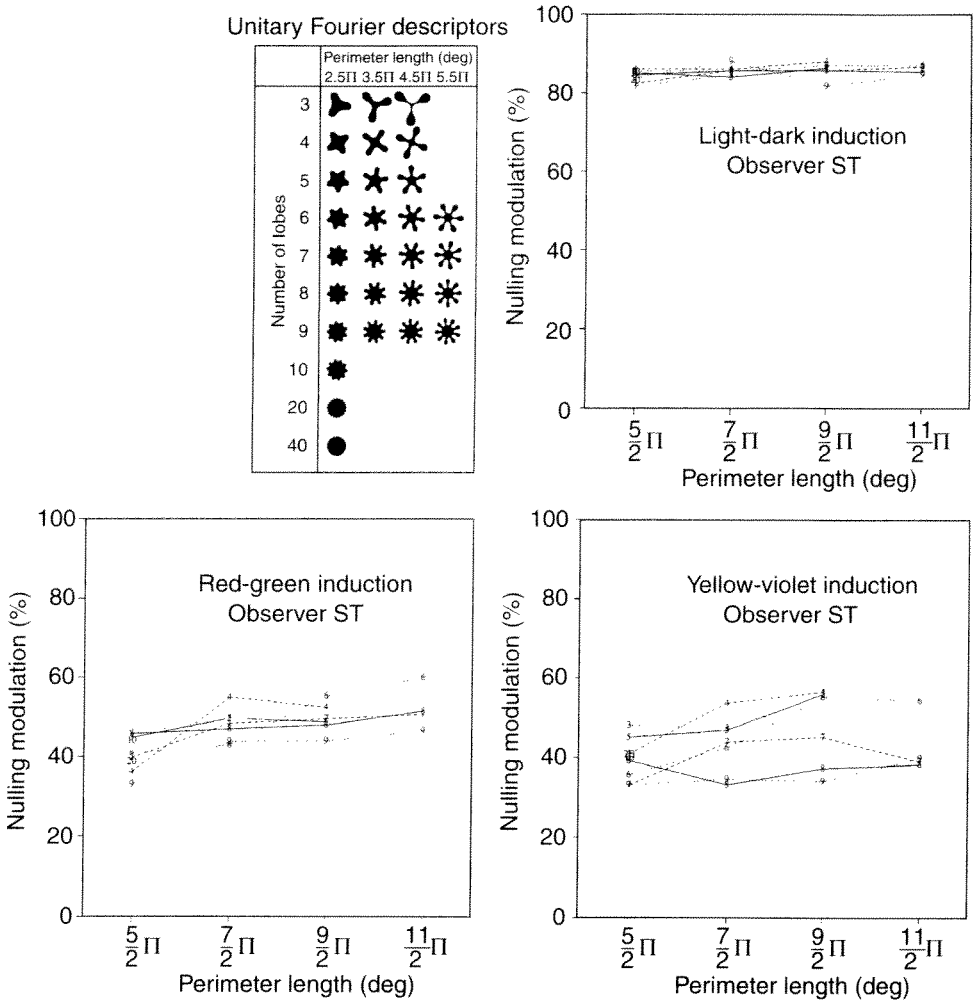

Figure 17.6: (A) Unitary Fourier descriptors of equal area, used as test stimuli. (B–D) Amplitude of nulling modulation in each cardinal direction expressed as a percentage of inducing modulation versus the length of the perimeter of the test. Each test stimulus is represented by the number of lobes that range from 3 to 40. Curves connect nulling amplitudes for tests of the same number of lobes (Zaidi, Yoshimi, & Flannigan, 1991).

ferent sizes (e.g., Yund & Armington, 1975; Reid & Shapley, 1988). These methods generally yield data where the amount of induction is a nonlinear function of parameters such as surround size. In these configurations, it is not possible to test whether the nonlinearity is due to nonlinear spatial interactions or just to a nonlinear decrease in effectiveness with distance from the test. To overcome these limitations, Zaidi et al. (1992) studied the differential effects of surrounding regions at different distances from the test by using stimuli similar to those shown in Fig. 17.7. The central disks are the test regions. In the surround, along every radial line through the center, the color of the surround varies as a spatial sine wave along a color line through mid-gray, resulting in a blurred bulls-eye target whose concentric rings vary sinusoidally in color. Each of the rows shows the color-varying surrounds of a single spatial frequency that were used in a single condition.

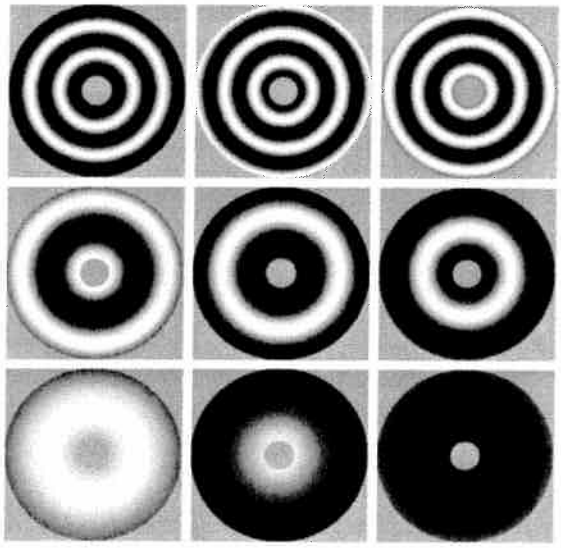

Figure 17.7: Single-frequency induction stimuli. Surrounds are annuli whose luminance varies sinusoidally along each radius. Each row depicts one spatial frequency condition. Across each row, the surround is shown at three different phases. The central disks are the tests, and in all nine pictures they are at the same mid-gray luminance; the differences in apparent brightnesses are due to different amounts of induced brightness from the surround (Zaidi et al., 1992).

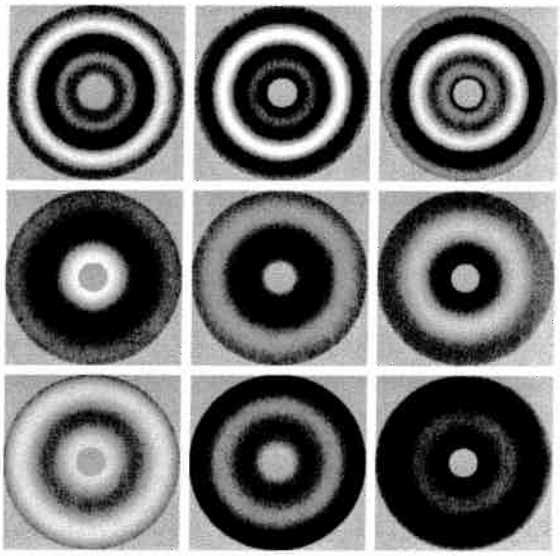

Figure 17.8: Compound-frequency induction stimuli. Surrounds are annuli whose luminance varies as the sum of two sine waves of different frequencies along each radius. The top row shows a high-frequency sine wave added to a medium-frequency sine wave. In the middle row, two different intermediate frequency sine waves are added. In the bottom row, a low-frequency sine wave is added to a medium-frequency sine wave. Across each row, the surround is shown at three different phases. The central disks are the tests, and in all nine pictures they are at the same mid-gray luminance; different apparent brightnesses are due to different amounts of induced brightness from the surround (Zaidi et al., 1992).

Each row shows three different phases of the sine wave with respect to the inner edge of the surround. The central disks in all nine pictures are physically the same mid-gray but appear to have different brightnesses. The surround consisted of a single sine wave of one of seven different spatial frequencies, and its color varied along one of the three cardinal directions of color space. Two aspects of the phenomenal appearance of the central test are directly relevant. First, within each row, as the phase of the surrounding sine-wave changes, the appearance of the test changes. The appearance of the test is roughly complementary to the appearance of the inner edge of the surround. In the experiments, as the phase of the surround (with respect to the inner edge) was changed uniformly in time, so that the sine wave appeared to drift toward the center at a constant velocity, the appearance of the center changed cyclically in time. The induced modulation was nulled by the addition of a real modulation to the test field, and the nulling modulation was used as the measure of the induced effect. Second, the magnitude of the change in the appearance is least in the top row, which has the surround with the highest spatial frequency, and largest in the bottom row, which has the surround with the lowest spatial frequency.

The measured amplitudes of the nulling modulation, when plotted against the spatial frequency of the surrounding sine waves, formed low-pass functions for all three color directions, LD, RG, and YV. For a variety of spatial models of induction, these low-pass functions indicate that the effects of elements of the surround decrease monotonically with distance from the test. Zaidi et al. (1992) and DeBonet and Zaidi (1997) showed that this spatial weighting function

could be estimated from the Fourier transform of the nulling amplitude versus spatial frequency function for brightness induction but not for chromatic induction. The reason is that, for induced achromatic brightness, the total effect of the surround could be described as the sum of the induced effects of individual elements of the surround. The spatial summation inference was based on the results of superposition tests, that is, the brightness induced by every pair of surrounds presented simultaneously was equal to the sum of the brightnesses induced by each component presented singly. Chromatic induction failed the spatial superposition tests.

The spatial superposition assumption was tested by comparing the induced effect of surrounds composed of pairs of circularly symmetric sine waves to the sum of the induced effects of the constituent sine waves. The paired sine waves were set to be in identical phases at the inner edge of the surround. These compound stimuli are shown varying in luminance in Fig. 17.8. The central disks are the test regions set at the mean luminance level. Each row shows surrounds consisting of the sum of two spatial frequencies windowed by the edges of the surround. The three rows consist of the same medium frequency paired with a high (top row), medium (middle row), and low frequency (bottom row). Across each row, three phases (with respect to the inner edge) of the paired sine waves are shown. Although the central disks are all at the same luminance level, they appear to be different, depending on the frequencies and phases of the surrounding sine waves. Sine waves of each of the spatial frequencies used in the previous experiment were paired with each other. The amplitude of each constituent sine wave was 0.5 to give a maximum amplitude modulation of 1.0.

Figures 17.9A–C show the data for the paired sine waves. The ordinate of each point is the amplitude of the required nulling modulation. Each curve in the figures connects the data for a particular spatial frequency when paired with the spatial frequencies corresponding to the abscissa. Different line types have been used to distinguish the curves. The key to identifying the curves is to begin at the leftmost point where the curves are ordered from top to bottom in the same

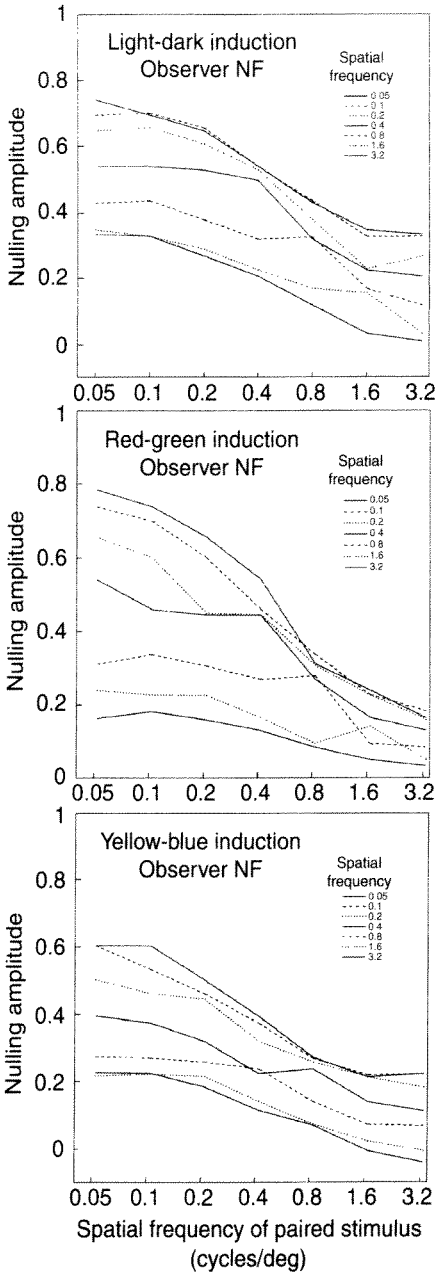

Figure 17.9: Results of superposition for color induction. The modulation required to null color induction is plotted against the spatial frequency of one sinusoidal component of the surrounding compound wave, with the other spatial frequency as the curve parameters indicated by the line types (Zaidi et al., 1992).

order as the spatial frequencies shown in the caption.

All of the curves for brightness induction are roughly parallel and have similar shapes, indicating that the amount of modulation required to null the induced effect decreases as each frequency is paired with progressively higher frequencies. Additionally, the incremental induced effect of adding sine waves of different spatial frequencies (represented by the curves) is fairly independent of the paired spatial frequencies (indicated on the abscissa). The main departure from parallelism is the noticeable upward protuberance in each curve where the two paired frequencies were equal, that is, the surround consisted of a single sine wave. At other pairings, the amplitude required to null the effect of a compound surround is a little less than the sum of the amplitudes required to null the effects of the two components. The slight subadditivity in nulling amplitude of the compound surround could be due to a phase mismatch in the modulation induced by the components. Because superposition holds, a punctate summation model can be used to fit the brightness induction data.

On the other hand, the curves for chromatic induction are not parallel. The range covered by the left-hand side of the curves (pairs that include the lowest spatial frequency) is considerably greater than the range covered by the right-hand side (pairs that include the highest spatial frequency). The effect of adding a high spatial frequency chromatic component severely attenuates the magnitude of induced contrast, especially for the RG condition.

Spatial variations in the region surrounding a central test disk, can be decomposed into variations along radial lines through the center of the disk (Figs. 17.7 and 17.8), and variations tangential to the radial lines. Zaidi and Zipser (1993) used radial patterns to examine the effect of surrounds whose color varied sinusoidally along concentric rings. The surrounds appeared to be composed of spokes, each spoke being of uniform color and the color of adjacent spokes varying continuously in a sinusoidal manner. The results showed that for both lightness and chromatic induction, on tests at the mean level, the induced effect of radially varying surrounds summed to zero, that is, to the sum of the induced effects of individual elements of the surround.

The results for brightness induction were interpreted in terms of a model that postulates a weighted spatial integration of induced effects. The perceived brightness at a point in visual space has two components, one due to the luminance of the light emanating from that point and the second due to the total induced effect of surrounding points. The model makes three assumptions about the induced effect. First, the induced effect of any surrounding point is in the complementary direction from the surround luminance relative to the test, with a magnitude proportional to the difference between the surround point and the mean level of the whole surround. Second, the induced effect of each surrounding point is weighted by a decreasing function of spatial distance from the test point. Third, the total induced effect is simply the sum of the induced effects of individual surrounding points. Algebraically, this model is defined by

$$(2) \qquad y = -\int_0^{2\pi} \frac{\int_0^{\infty} g(s)A(\Omega, s)s\,ds}{2\pi} d\Omega ,$$

where y is the total induced effect at the test point, Ω is the angular orientation, s is the spatial distance between the test and induction point, $g(s)$ is the monotonically decreasing spatial weighting function of s, and $A(\Omega,s)$ is the signed magnitude of the luminance difference between the inducing point at (Ω,s) and the surround mean level.

The stimuli used in these experiments were circularly symmetric and varied only along radial lines; therefore, if the weighting function is assumed to be isotropic, Eqn. (2) can be reduced to a function of just the radial distance. For a surround consisting of a drifted single sinusoid of spatial frequency equal to ϕ cycles/deg, the induced effect at time t for the center point of the circular test can be expressed as:

$$(3) \quad y(t, \phi_i) = -A\int_L^X g(s)\cos[2\pi(\rho_0 t - \phi_i s + \phi_i L)]s\,ds.$$

where A is the amplitude of the surround sine wave, L

is the inner edge of the surround (i.e., the radius of the test disk), X is the outer edge of the surround, and ρ_0 is the temporal frequency of the drift (in cycles/s). For the present model, given that the test is uniform in contrast and luminance, the induced effect of all points inside the test on the test center is zero. Therefore, Eqn. (3) is expressed solely in terms of the effect of the surround. In the case of compound sine wave stimuli, with a second sinusoid of spatial frequency ϕ_j cycles/deg, the total induced effect is equal to:

$$y(t, \phi_i, \phi_j) = 0.5\, y(t, \phi_i) + 0.5\, y(t, \phi_j).$$

Since the induced contrast and brightness modulations could be suitably nulled with the addition of real sinusoidal modulation of the same temporal frequency as the inducing modulation, it is sufficient to describe the inducing, induced, and nulling modulations in terms of their amplitude and phase. For each condition, the amplitude and phase of the induced modulation were obtained from the Fourier transform of Eqn. (3) in the temporal frequency domain. By exploiting the fact that the drift was at a constant velocity given by ρ_0 divided by ϕ_i, the Fourier transform was simplified to

$$(4) \quad Y(\rho_0, \phi_i) = -\frac{A}{2}[\delta(\rho - \rho_0) + \delta(\rho + \rho_0)] \cdot e^{i2\pi\phi_i L} \cdot \int_L^X g(s) e^{i2\pi\phi_i s} s\, ds,$$

where $Y(\rho_0, \phi_i)$ is the Fourier transform of induced modulation for a surround of spatial frequency ϕ_i drifted toward the test point at a temporal frequency of ρ_0, and δ is the Dirac delta function.

If the three assumptions of the model are satisfied, then, given the proper choice of $g(s)$, Eqn. (4) should fit the data. Since many smooth monotonic functions can be approximated by exponential functions, it was assumed that the spatial weighting function could be approximated by a negative exponential function of the form:

$$(5) \quad g(s) = \kappa e^{-\alpha s}.$$

Equation (5) was substituted into Eqn. (4), and solved to obtain the following expression:

$$(6) \quad Y(\rho_0, \phi_i) = \frac{-A\kappa}{(\alpha + i2\pi\phi_i)^2} \cdot [(1 + \alpha X + i2\pi\phi_i X) \cdot \exp(i2\pi\phi_i L - \alpha X - i2\pi\phi_i X) - (1 + \alpha L + i2\pi\phi_i L) e^{-aL}].$$

Expressions for the amplitude and the temporal phase of the induced modulation were then derived by transforming the right-hand side of Eqn. (6) into the polar form:

$$(7) \quad Y(\rho_0, \phi_i) = Amplitude \cdot e^{iPhase}.$$

This model provided a good fit to the contours of the brightness induction curves. The spatial weighting functions, $e^{-\alpha s}$, can be expressed in terms of the space constant $1/\alpha$ in degrees of visual angle, which gives the distance from the test at which the effectiveness of a surround point has fallen to $1/e$ of the maximum. The best estimates of the brightness induction space constants for two observers were 0.74 and 0.29 deg (DeBonet & Zaidi, 1997).

The failure of the superposition test for chromatic induction may not indicate a failure of weighted spatial summation if the nonlinearity in the data is due to punctate amplitude nonlinearities. Amplitude nonlinearities would also explain the nonlinear nulling functions in Fig. 17.3. To test this possibility, it was assumed that local chromatic signals from each point in the image pass through an amplitude compression in the visual system prior to the stage of lateral interactions responsible for chromatic induction. Such a nonlinearity will have two effects on the spatial summation model. Inside the surround, the nonlinearity will reduce the effective contrast of the surrounding wave. Inside the test, the nonlinearity will reduce the effectiveness of the nulling modulation. Mathematically, this is represented by

$$(8)\quad y = -\zeta[N] = -\int_0^{2\pi} \frac{\int_0^{\infty} g(\Omega, s)\zeta[A(\Omega, s)]s\,ds}{2\pi} d\Omega,$$

where ζ is an odd-symmetric nonlinear compressive function, y is the actual induced modulation, and N is the measured nulling amplitude. The equation

$$(9)\quad \zeta[A] = A - cA^3$$

was used as the odd-symmetric compressive function of amplitude A.When applied to sinusoidal stimuli, ζ generates higher order harmonics. However, the optimal choice of the compressive nonlinearity for the fits to the present data generated higher harmonic energy that was less than 10% of the energy in the fundamental. Substituting $\zeta[A]$ and $g(s)$ into the one-dimensional form of Eqn. (8) yields an instantaneous induction level given by:

$$(10)\quad y(t) = -N(t) + cN^3$$
$$= -\int_X^L \kappa e^{-\alpha s}[A(s, t) - cA^3(s, t)]s\,ds,$$

where $A(s, t)$ is the amplitude at radius s and time t of the surround. By using Fourier transforms similar to those used for the analysis of the linear model and removing higher order harmonics, an expression was derived for the amplitude of the inducing stimulus. However, numerical simulations showed that whereas adding a pointwise compressive nonlinearity to the model helped predict the compression in Fig. 17.3, it actually changed the predicted pattern of the results slightly in the direction opposite to the nonadditivity in the superposition measurements in Figs. 17.9B and C. The compressive nonlinearity had the effect of slightly attenuating predicted induction from low-frequency surrounds rather than from high-frequency surrounds. Consequently, the nonadditivity of chromatic induction is not due to a pointwise contrast compression. The causes of the subadditivity are discussed further in the section on habituation.

Adaptation mechanisms and brightness induction

The studies in the previous section isolated lateral effects from contamination by the effects of steady-state adaptation by keeping the time- and space-averaged mean luminance of all points in the stimulus equal. The results showed that brightness induction can be characterized as a summation process with a negative exponential spatial weighting function. Figure 17.10 from Zaidi, DeBonet, and Spehar (1995), however, demonstrates a failure of spatial summation. This figure consists of three vertical surround segments filled with a random binary texture. The space-averaged luminance of the three segments is equal, and the spatial contrast progressively decreases from left to right with values of 1.0, 0.3, and 0.0 Michelson contrast. Centered in each of the surround segments are five spatially uniform diamonds decreasing in luminance from top to bottom. The diamonds across each row have identical luminances, yet most observers see them as increasing in lightness from left to right in the top rows and from right to left in the bottom rows. This display shows that even from a nonfigural variegated surround, brightness induction depends in a complex manner on the relative luminance of the test and individual regions of the surround. A surround made of equal numbers of light and dark squares makes a test at higher than its average luminance appear darker than does a spatially uniform surround at the average luminance, and it makes a test at lower than its average luminance appear lighter than does the spatially uniform average luminance surround. The empirical ranked lightnesses of diamonds within each row are presented in Table 17.1, along with predicted rankings from the model described later in this section.

Other reports of additivity failure (Brown & MacLeod 1991; Schirillo & Shevell, 1996) also consisted of stimuli where the test was at a different level than the mean luminance and color of the surround, thus making it imperative to explicitly consider spatially local and extended adaptation mechanisms. In a steady display it is impossible to separate the effects of adaptation from induction. Spehar, DeBonet, and Zaidi

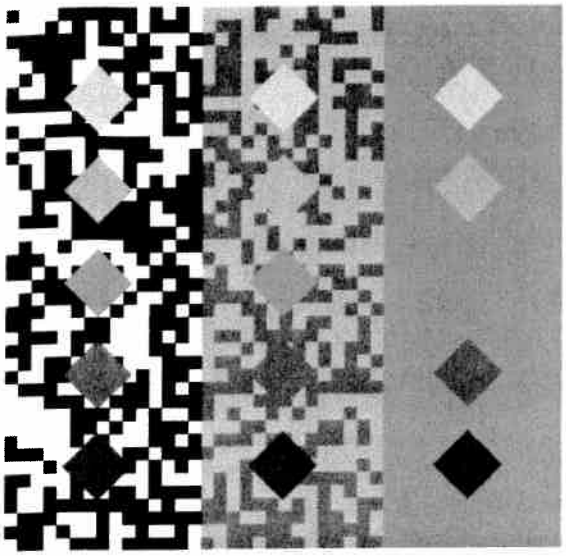

Figure 17.10: Brightness induction from random, binary textured surrounds. The three vertical surround segments have equal spatially averaged luminance, while the spatial contrast progressively decreases from left to right (1.0, 0.33, and 0.0). Centered in each of the surround columns are five spatially uniform diamonds with luminance decreasing from top to bottom. The luminance of the diamonds in the middle row is equal to the mean luminance of the surround segments. Diamonds across each row are of identical luminance, but their perceived lightnesses differ.

Observer JS			Observer BS		
3 (3)	2 (2)	1 (1)	3 (3)	2 (2)	1 (1)
1 (1)	2 (3)	3 (2)	3 (2)	2 (3)	1 (1)
1 (1)	2 (2)	3 (3)	1 (1)	2 (2)	3 (3)
1 (1)	2 (2)	3 (3)	1 (1)	2 (2)	3 (3)
1 (1)	2 (2)	3 (3)	1 (1)	2 (2)	3 (3)

Table 17.1: For two observers the empirical (and predicted) ranked lightness of diamonds within each row are presented in a similar configuration as the display shown in Fig. 17.10 (Zaidi, DeBonet, & Spehar, 1995).

(1996) used the time-varying nulling method to generate a general model for brightness induction from such variegated nonfigural surrounds to identify the conditions under which the induced effect can be described as spatially additive, and to delineate the processes that lead to failures of additivity.

The stimuli used were similar to those shown in Fig. 17.11A, in which a foveally fixated, spatially uniform disk was surrounded by an annulus filled with a binary random texture that was composed of equal numbers of two sets of randomly intermixed, equal-sized square elements. The luminance of each set of elements was modulated sinusoidally in time. The mean level, amplitude, and phase of temporal modulation were independently controlled for each set. Temporal modulation of the luminance of the surround elements resulted in an induced modulation of the brightness of the test. The amplitude of the nulling modulation was used as the measure of the overall induced effect.

In Experiments 1 and 2, the total induction on the test was measured when the luminance of both sets of texture elements was modulated sinusoidally at 0.5 Hz. The luminance modulation amplitudes of one set were 0.0, 0.2, 0.4, 0.6, 0.8, or 1.0, paired with modulation amplitudes of the other set of 1.0, 0.5, 0.0, –0.5, or –1.0, where a positive or negative amplitude denotes modulation in the same or opposite phase as the paired modulation. One cycle of each of these combinations of luminance modulations is shown schematically in Fig. 17.11B. In Experiment 1, to test for spatial additivity, the test and all elements of the surround had the same space- and time-averaged mean luminance. In Experiment 2, to test for local gain controls, the time-averaged luminance of the test disk was set at either 0.5 or 1.5 times this value. Experiment 3 tested for spatially extended gain controls by using the same spatial configuration to examine whether modulation in spatial contrast of the surround could produce brightness induction in the test. This was achieved by modulating the luminances of the two surround components sinusoidally at 0.5 Hz, with equal amplitudes, in opposite phase, around different mean luminance levels, as depicted in Fig. 17.11C. Experiment 4 measured induction from spatially uniform surrounds at three different mean levels on tests at the same three levels. This experiment tested for the effect of local gain controls in the surround, and whether spatially uniform

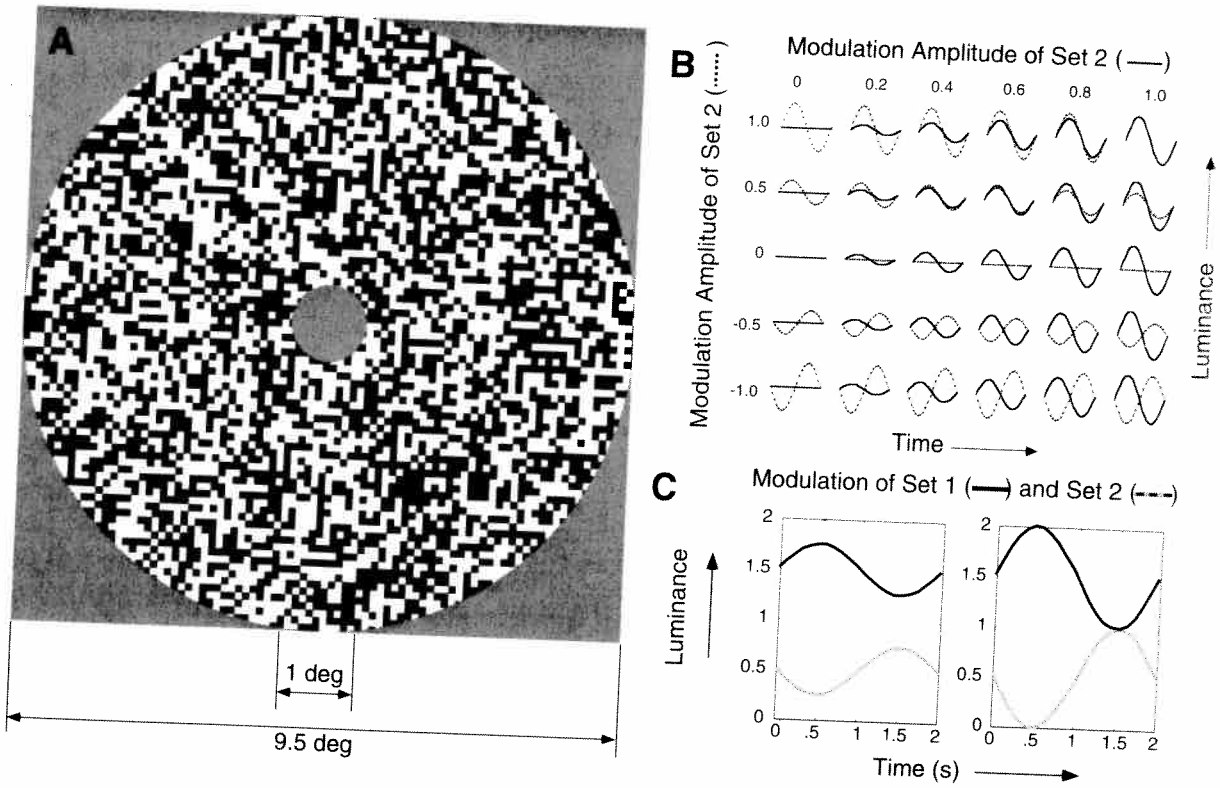

Figure 17.11: (A) Stimulus configuration: spatially uniform, 1-deg disk surrounded by 9.5-deg annulus composed of a binary random texture. (B) The luminance modulations of each set of the surround elements in Experiments 1 and 2. The luminances of both sets of texture elements were modulated sinusoidally at 0.5 Hz. The luminance modulation of one set was set at an amplitude of 0, 0.2, 0.4, 0.6, 0.8, or 1.0, paired with a modulation amplitude of the other set at 1.0, 0.5, 0, –0.5, or –1.0, where a negative sign denotes modulation in the opposite phase. (C) The luminance modulation of each set of the surround elements in time (sinusoidal modulation with frequency of 0.5 Hz) in Experiment 3. The two surround components were modulated with equal amplitudes in counterphase around different mean luminance levels. As a result, the space-averaged luminance of the surround was constant while the spatial contrast was modulated with amplitudes of 0.5 (left panel) or 1.0 (right panel) (Spehar, DeBonet, & Zaidi, 1996).

surrounds had similar brightness induction effects as textured surrounds.

Results from all four experiments, for one observer, are presented in Figs. 17.12A–E. The curves are the best fits from the model described below. For Experiments 1 and 2, the amplitude of nulling modulation is plotted as a function of the amplitude of the modulation of one set with the amplitude of modulation of the other set as a curve parameter. For each modulation level of one set, the magnitude of nulling modulation is a linear function of the amplitude of the paired set. In addition, the five curves for each observer are parallel and equally spaced, indicating that the amplitude of nulling modulation is a linear function of the amplitude of each of the surround sets. These results unambiguously support a model that postulates the spatial summation of lateral effects.

A comparison of Figs. 17.12A–C shows that the amount of required nulling modulation increases as the mean luminance of the test increases. It is important to

note that the empirically measured nulling modulation amplitude is plotted, and not the amplitude of the induced modulation. It is well established (e.g., Watson, 1986) that the threshold for the detection of temporal modulation at 0.5 Hz is an increasing function of the mean luminance of the field. Craik (1938) conceptualized this fact in terms of a gain factor for the test modulation, set by the test mean luminance. In Experiments 1 and 2, even if the brightness induced from the surround were independent of the luminance level of the test, because of the gain set by the mean luminance of the test, the amount of real modulation needed to null the induced modulation should increase as a function of the test mean. The results of these experiments imply that local adaptation mechanisms in the test field should be incorporated into a general model of brightness induction.

For Experiment 3, the magnitude of nulling modulation is plotted as a function of the amplitude of contrast modulation. The three sets of symbols in each graph represent the data for the tests at the three different mean luminance levels. Contrast modulation of the surround does not produce any significant brightness induction for tests at the same mean luminance level as the surround (middle points). However, the results for test luminance levels at 0.5 and 1.5 show a nulling modulation of approximately equal amplitude but opposite sign. Phenomenally, this can be described in the following way: The test at a mean luminance of 0.5 appears lighter on the higher contrast surround and darker on the lower contrast surround; the opposite happens for the test at a mean luminance of 1.5. The change in sign of the required nulling modulation as a function of test mean level indicates that brightness induction is not a function of contrast modulation per se, and that the results may be better understood if the inducing effects of the two sets of surround elements are considered separately. A positive sign indicates that the nulling modulation was in the same phase as the modulation of the surround set with the higher mean luminance level in Fig. 17.11C, and a negative sign indicates that the nulling modulation was in phase with surround set of the lower mean luminance. The surround set whose mean luminance is closer to the mean level of the test seems to have a greater inducing effect. In the case where the test level is equidistant from the mean levels of the two surround sets, the induced effects cancel out and there is roughly zero induced modulation. These results suggest that the magnitude of the difference between the luminance level of the test and the mean luminance level of each surround element should be considered in modeling the total induced effect from complex surrounds. The model below postulates that there are pairwise lateral connections between points in the test and the surround, and that the magnitude of the induction signal between them is a decreasing function of the mean luminance difference between them.

For Experiment 4, the magnitude of nulling modulation is plotted as a function of the surround mean luminance level with the test luminance level as a curve parameter. For all three test levels the magnitude of the nulling modulation was the highest when the surround modulation was at the same mean luminance as the test. The magnitude of the nulling modulation decreased monotonically as a function of the difference between the test and the surround mean luminance levels. When the mean luminance levels of the test and the surround were equal (i.e., the highest points for each surround luminance) the magnitude of nulling modulation was approximately constant, thus indicating that the effect of the local gain control set by the surround mean luminance on the inducing signal roughly balances the effect on the nulling modulation of the gain set by the test mean.

To account for their results, Spehar, DeBonet, and Zaidi (1996) generalized the model in the previous section so that the induced effect from each point in the surround is proportional to its luminance attenuated by two gain controls and a spatial weighting function:

$$(11)\quad I(t) = -\int_0^{2\pi} \frac{\int_0^{\infty} W(s)\Gamma_D(\Omega, s)\Gamma_s(\Omega, s)L(\Omega, s, t)s\,ds}{2\pi} d\Omega.$$

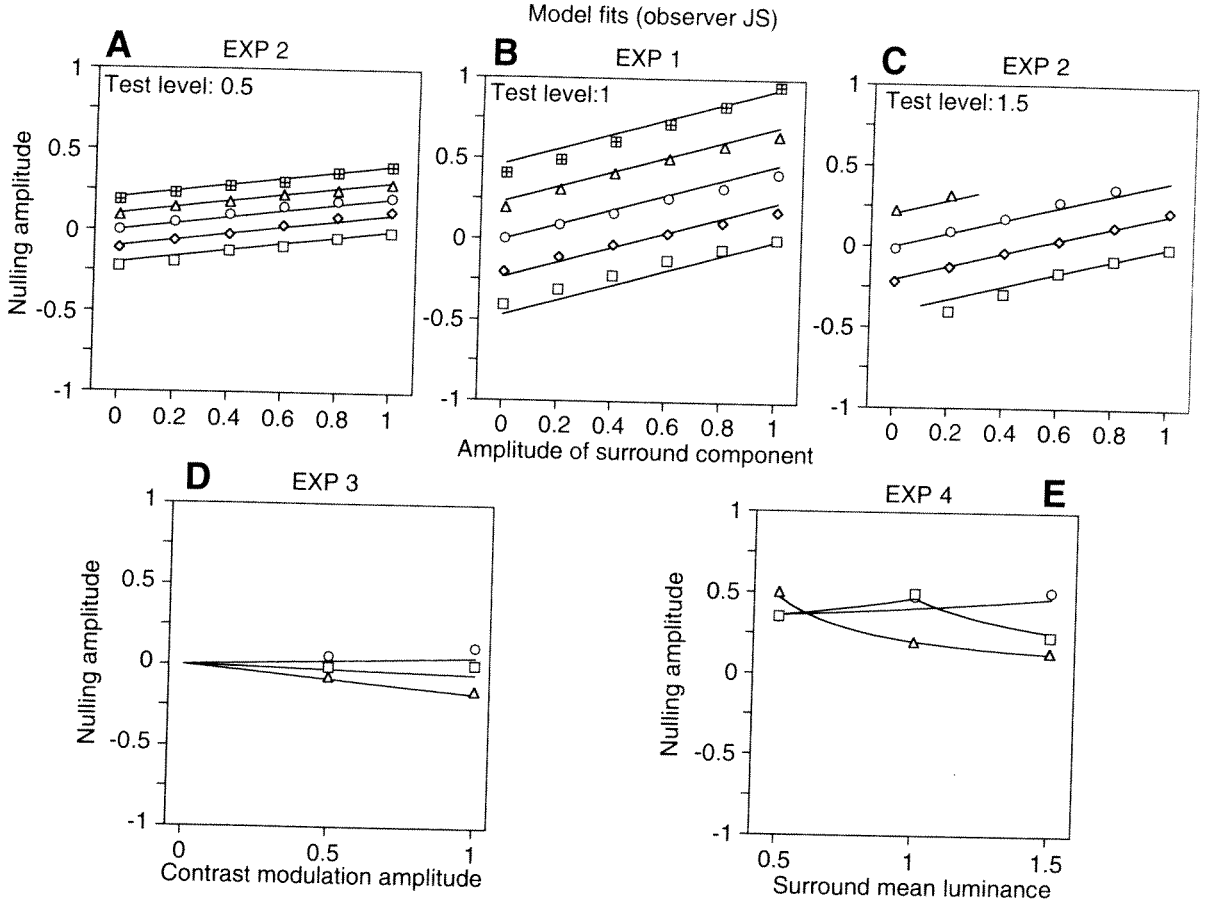

Figure 17.12: (A–C) Results of Experiments 1 and 2. The amplitude of nulling modulation (ordinate) is plotted as a function of the magnitude of the modulation of one set of the surround elements (abscissa) with the amplitude of modulation of the paired set as a curve parameter: (□: −1.0, ◊: -0.5, O: 0, Δ: 0.5, ⊞: 1.0). Each data point is the average of two 2AFC staircases (10 turns each). Test mean luminance equaled 0.5, 1.0, and 1.5 times 25 cd/m^2 for (A)–(C). (D) Results of Experiment 3. The magnitude of nulling modulation (ordinate) is plotted as a function of the amplitude of contrast modulation (abscissa). The curves represent data for tests at three different mean luminance levels. Triangles represent data for the test with a mean luminance of 0.5. Squares represent data for the test with a mean luminance of 1.0. Circles represent data for the test with a luminance level of 1.5. (E) Results of Experiment 4. The magnitude of nulling modulation is plotted as a function of the surround mean luminance level with the test mean luminance level as a curve parameter. Triangles represent data points for the test with a mean luminance equal to 0.5, squares test mean luminance 1.0, and circles test luminance level 1.5. Curves represent the best-fitting predictions of the model (Spehar, DeBonet, & Zaidi, 1996).

$I(t)$ is the total induced effect on the test patch at time t. (Ω, s) are the polar coordinates of a surround point, where Ω is the angular direction in radians and s the spatial distance from the test in degrees of visual angle. $L(\Omega, s, t)$ is the luminance at that point at time t. $\Gamma_S(\Omega, s)$ and $\Gamma_D(\Omega, s)$ are two gain control factors that affect the signals from (Ω, s). $W(s)$ is a negative exponential spatial weighting function of s (Zaidi et al., 1992).

In the model the response of the visual system to a luminance signal at every point is gain controlled by a factor that depends solely on the mean luminance level

at that point. Following tradition (e.g., von Wiegand, Hood, & Graham, 1995), hyperbolic gain control functions were used to calculate the gain factor for each point in the surround:

$$(12)\quad \Gamma_s(\Omega, s) = \frac{\gamma_S}{\gamma_S + \int L(\Omega, s, t)dt},$$

where the parameter γ_S is a constant for each observer. By incorporating this gain control, the induction model is able to predict that surround modulation around a low mean luminance level generates more induction than the modulation of the same amplitude around a higher level.

To allow for the possibility that the local gain factor for the centered test field could be different from the surround gain, the factor Γ_C for the test was calculated by:

$$(13)\quad \Gamma_C = \frac{\gamma_C}{\gamma_C + \int L(0, 0, t)dt},$$

where $\int L(0, 0, t)dt$ is the time-averaged luminance of the center of the test and γ_C is a constant parameter for each observer.

The experimental results show that induced modulation depends on the pairwise differences between the mean levels of the test and individual surround points. This was modeled by attenuating the induction from each point by a gain factor set by the absolute difference between the time-averaged luminance at that point and the time-averaged luminance of the test:

$$(14)\quad \Gamma_D(\Omega, s) = \frac{\gamma_D}{\gamma_D + \left|\Gamma_s(\Omega, s)\int L(\Omega, s, t)dt - \Gamma_C\int L(0, 0, t)dt\right|},$$

where the parameter γ_D is constant for each observer. For conditions that include different mean levels, the local adaptation mechanism operating on the test field will influence the effectiveness of the added nulling modulation. Therefore, the true null will be achieved when the real modulation, after being gain-controlled by the test mean level, is equal and opposite to the induced modulation:

$$(15)\quad \Gamma_C N(t) = -I(t),$$

where $N(t)$ is the luminance modulation required to counteract the induction at time t. Thus the complete expression for the null is:

$$(16)\quad N(t) = \frac{1}{\Gamma_C} \cdot \int_0^{2\pi} \frac{\int_0^{\infty} W(s)\, \Gamma_D(\Omega, s)\, \Gamma_s(\Omega, s)\, L(\Omega, s, t)s\,ds}{2\pi} d\Omega.$$

For the spatial and temporal configurations used by Spehar, DeBonet, and Zaidi (1996), the general model can be simplified. Because the spatial composition of the binary texture in the surround was a uniform random distribution, it is sufficient to consider the effects of identical numbers and distribution of the two sets of surround elements, instead of considering each surround point individually. Therefore, instead of determining an observer's spatial weighting function, it is sufficient to estimate its aggregate effect on each of the two types of surround elements. In addition, the integrals in Eqn. (16) can be replaced by the sum of the independent effects of the two surround sets. Further simplification can be achieved because of the nature of the temporal modulation. The luminances of the two surround components were always modulated sinusoidally with the same frequency, either in phase or in the opposite phase. The model predicts that the induced modulation should also be sinusoidal with the same frequency and in the opposite phase with either one or both of the surround components. Therefore, for these conditions, it is sufficient to describe the inducing and nulling stimuli by just their signed amplitudes of modulation instead of considering instants of the modulating waveform.

As a result, for all of the conditions in the present study, Eqn. (16) can be simplified to predict the ampli-

tude of the required nulling modulation N by the equation

$$(17) \quad N = \frac{w}{\Gamma_C} \sum_{i=1}^{2} \frac{\Gamma_{D_i} \cdot \Gamma_{s_i} \cdot A_i}{2},$$

where w incorporates the effect of the integrated spatial weighting function over the surround, A_i is the signed amplitude of luminance modulation of the ith component, and Γ_{D_i} and Γ_{S_i} are the gain controls that apply to the set i. This simplified model has only four free parameters: the three gain control constants γ_S, γ_D, and γ_C, the spatial weighting parameter w that scales the amplitude of induction for each observer.

The entire set of each observer's data was fit simultaneously with Eqn. (17). Figure 17.12 shows that the model's predictions fit the data extremely well. At all fixed mean levels of test and surround, the model predicts a linear relationship between the amplitude of the modulation of the surround components and the nulling amplitude. It also accounts for the changes in the nulling modulation amplitude due to variations in the mean luminance level of the test. The model correctly predicts the relative amplitude and phase of brightness induction from the contrast-modulated textured surround, and that the magnitude of brightness induction for different luminance levels of uniform surrounds and tests is a monotonically decreasing function of the difference between their luminance levels.

There is a large amount of psychophysical and physiological evidence for the spatially local gain controls used (Shapley & Enroth-Cugell, 1984; Chen, MacLeod, & Stockman, 1987). These adaptation mechanisms are known to occur relatively early in the visual system. The novel suggestion in this model is the pairwise spatially extended gain control on lateral interactions. Since the spatial weighting function for brightness induction falls off steeply as a function of distance from the test, these pairwise connections can be restricted to fairly short distances in retinal or cortical coordinates. A static compressive nonlinearity on these pairwise connections is not a viable alternative to this spatially extended gain control, because the predictions from a static nonlinearity depart significantly from the straight lines required to fit the data from Experiments 1 and 2. Another model for color induction that explicitly considers the role of adaptation mechanisms is by Courtney, Finkel, and Buchsbaum (1995).

This model decouples the inducing signal from a region and the signal induced into that region. The decoupling removes the need to make a recursive model, like Grossberg and Todorovic (1988), and results in computational simplicity. The assumption is that the incoming induced signals affect appearance but would not affect modulation thresholds or the nulling modulation. There is no psychophysical or physiological evidence that the outgoing inducing signal from a point is affected by in-coming induced signals. On the other hand, Spehar and Zaidi (1997a) have shown that the steady luminance level of the surround influences temporal contrast sensitivity only by presenting a contrast pedestal at the edge of the test and does not affect luminance modulation thresholds inside the test. Similar inferences were drawn on the basis of increment threshold measurements by Cornsweet and Teller (1965).

To judge how well the model performs for static displays, the effect of local adaptation on the appearance of the test was incorporated into the model by the equation:

$$(18) \quad P_C = \Gamma_C \cdot C + I_C,$$

where P_C, the predicted perceived gray level, is equal to C, the luminance of the test, multiplied by the gain factor for that luminance level, Γ_C [Eqn. (13)], plus I_C the total induced brightness on the test [Eqn. (11)]. Using the parameters estimated for the time-varying measurements, predictions were made for perceived brightness in Fig. 17.10. The predicted rankings, shown in Table 17.1, differed somewhat between observers, yet they agreed almost perfectly with the actual rankings made by each observer.

The success of the present model shows that perceived gray levels can be predicted in complex achromatic configurations by incorporating the effects of local and spatially extended adaptation mechanisms

and linear summation of the induced effects of individual elements of the surround. The model consists of a simple nonrecursive integral equation with the only independent variables being the physical luminances of individual pixels, making it easy to implement for arbitrary, nonfigural, achromatic displays.

Habituation and chromatic induction

The failures of superposition tests for chromatic induction (Fig. 17.9) showed that high spatial frequency chromatic variations inhibit the inducing power of low-frequency surrounds. Since any nonlinear weighting function just multiplies the effectiveness of individual pixels, it is distributive over the addition of sine waves and is not the cause of the failure. Static response nonlinearities were ruled out earlier, and since the superposition tests were equated for the mean levels of all pixels in the display, gain control mechanisms cannot be the cause of the failure either. Even though it is difficult to conceive of hard-wired, nonlinear lateral connections that could lead to this particular failure of superposition, it is always a possibility. An alternative possibility is that high spatial frequency chromatic variations mask low-frequency variations. Some support for this alternative is provided by the study of Zaidi, Spehar, and DeBonet (1997), who showed that adapting to high-frequency chromatic texture will elevate thresholds for detecting low-frequency chromatic changes, but that this does not happen for luminance changes.

They used steady adapting fields subtending 14.14 deg horizontally and 10.63 deg vertically. The fields were covered with random binary and quaternary distributions of uniform-sized squares, 8.52 squares per deg^2 of visual angle. Three types of binary texture were used, which will be termed LD, RG, and YV for mnemonic purposes. Each type of texture consisted of equal numbers of randomly intermixed squares of two different colors, whose chromaticities and luminances were equal to points halfway between *W* and the extreme points on the corresponding cardinal axis. The three types of quaternary textures, LDRG, RGYV, and YVLD, were formed by adding the corresponding pairs of binary textures. For example, the LDRG texture consisted of light-red, dark-red, light-green, and dark green squares. In all six types of textures, the space-averaged chromaticity and luminance were equal to *W*. Thresholds for detecting full-field changes on the textured fields were compared to thresholds for detecting changes parallel to and toward one or the other end of the three cardinal axes, that is, the R, G, Y, V, L, and D color directions, from a uniform achromatic field whose chromaticity and luminance were equal to the space average of the texture. The chromaticity or luminance of all of the pixels of the screen was changed over 3 s as a half-cycle of a sinusoid. To control for criterion effects, each trial also included another interval in which the illuminant was not altered. The observers indicated the interval in which they perceived a color change. The observer adapted to the background for 2 min at the initiation of each session and readapted for 2 s after each trial.

The results for two observers are shown in Fig. 17.13. The chromatic content of the background texture is indicated on the abscissa. The log threshold elevation for detecting a change in each color direction as compared to the baseline threshold is plotted on the ordinate. The results are systematic and similar for the two observers. The presence of chromatic spatial variations makes it less likely that full-field chromaticity changes will be perceived, but thresholds for detection of full-field luminance changes are not affected by the presence of spatial variations. With one exception, changes toward a chromatic direction are affected only when there is spatial contrast along the same axis. There was no systematic effect of superimposing spatial contrast along a color axis orthogonal to the color direction of the simulated illumination change. The results indicate that the masking effect of spatial contrast is relatively independent within each of the two chromatic mechanisms.

The observers in these experiments were instructed to fixate on the center of the screen, but small eye movements are unavoidable when trying to maintain fixation (Carpenter, 1988). If the main effect of eye-movements were integration over space within recep-

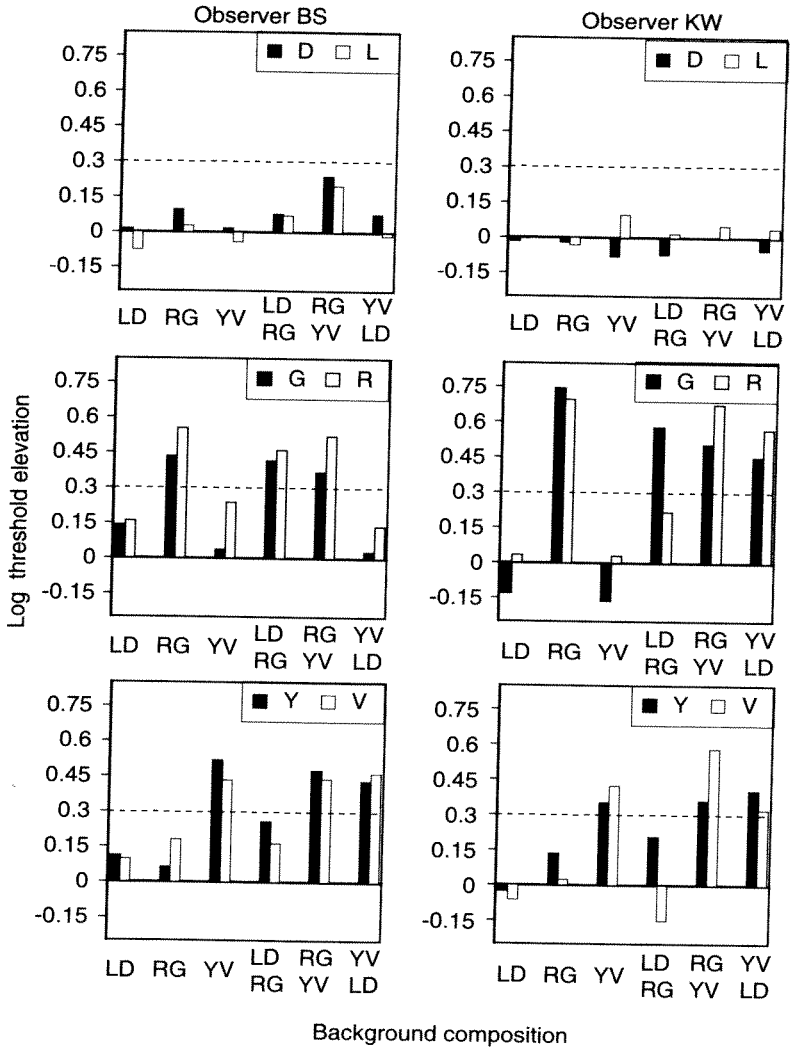

Figure 17.13: Results of adapting to high-frequency texture on full-field luminance and chromatic changes. The log of the threshold for detecting a change in each color direction minus the log of the baseline threshold for that color direction is plotted against the chromatic content of the background texture (see text). Symbols representing the color direction of the test are shown in the insets. A dashed horizontal line is drawn at 0.3 to indicate a doubling of threshold magnitude (Zaidi, Spehar, & DeBonet, 1997).

tive fields, as has been proposed by, for example, D'Zmura and Lennie (1986) and Fairchild and Lennie (1992), the adaptation level would be set by the mean chromaticity and luminance. Since the space-averaged colors of all of the backgrounds were identical to *W*, the presence of threshold elevations on the textured backgrounds rules out spatial integration as a major factor. In our view, however, eye movements lead to transient stimulation of receptive fields at the borders of the squares, thus creating a temporal modulation of stimulation to individual neurons, and prolonged temporal modulation has been shown to cause chromatically selective elevations of thresholds (Krauskopf, Williams, & Heeley, 1982; Zaidi & Shapiro, 1993).

The best clue for explaining why habituating to textures raises thresholds for large-field chromatic changes but not for luminance changes is provided by the study of Krauskopf and Zaidi (1985), which showed that habituating to modulation of a large disk raised thresholds for a concentric smaller disk in the chromatic case but not in the luminance case. Habituation to luminance modulation occurred only when the habituating stimulus shared the edge of the test stimulus. In the visual system, beginning from ganglion

cells in the retina, neurons are spatially bandpass for luminance variations and hence are insensitive to variations that are uniform over their receptive fields (Shapley & Lennie, 1985). If detection of the large-field luminance changes on all of the backgrounds occurs at the edges of the display, then habituation to luminance modulations will not alter luminance thresholds. Neurons with receptive fields wholly within the boundary do not participate in the detection of luminance changes at the boundary, and habituation of neurons that are stimulated by eye movements across the boundary will be common to all conditions. On the other hand, since chromatically sensitive cortical neurons are responsive to chromatic variations that are uniform over their receptive field (Lennie, Krauskopf, & Sclar, 1990), large-field chromatic changes are detected inside the boundary and most probably near the fixation point. Habituation of neurons in the central field by eye movements across the internal edges in the texture will therefore raise chromatic thresholds. If the squares are too large, there will be few neurons whose receptive fields oscillate across boundaries; and if the squares are too small, there may be too much integration within receptive fields for there to be a substantial modulation of responses. Therefore, receptive field sizes and amplitudes of eye movements will jointly determine the sizes of the squares that elevate thresholds the most.

Using probe-flash measurements, Zaidi, Spehar, and DeBonet (1998) showed that the effects of adapting to textured fields were qualitatively similar to habituating to prolonged temporal modulations (Shapiro & Zaidi, 1992; Zaidi & Shapiro, 1993). Adapting to high spatial frequency textures, however, elevated thresholds considerably more than habituating to temporal modulations of spatially uniform fields.

There are a few differences between the texture adaptation conditions and the drifting compound sine-wave conditions, and a test of habituation by specific spatial frequencies should be done under drifting conditions. However, the most probable explanation for the spatial subadditivity in chromatic induction is that habituation to high-frequency chromatic variations reduces the total inducing signal from the surround to the test. A number of studies using other methods have shown that adding chromatically different dots or rings to spatially uniform surrounds reduces induction more than by adding stimuli differing in luminance (Shevell & Wesner, 1990; Jenness & Shevell, 1995). For example, Shevell and Wesner (1990) showed that the induced effect of a spatially uniform, red annulus on a central test is more severely attenuated by a thin equiluminant white ring embedded in the surround than by a dark ring. If their results are analyzed in terms of spatial frequency, a thin white ring equiluminant with the red surround adds high-frequency chromatic components, whereas a thin dark ring adds high-frequency luminance components. The results of Shevell and Wesner's study show that adding high spatial frequency chromatic components to a low-frequency chromatic surround attenuates the induced effect considerably more than adding high spatial frequency luminance components, which is consistent with the habituation results above.

Perceptual organization and object segmentation

Figure 17.1 is just one of many pictures in the literature showing a relationship between inferred three-dimensional organization and perceived brightness. Others can be found in Gilchrist (1977, 1980), Knill and Kersten (1991), Adelson (1993), and Sinha and Adelson (1993). In addition, Benary (1924), White (1979), and Zaidi (1990) presented brightness illusions that are incompatible with all extant low-level models of brightness induction, including the ones in this chapter. Inferred perceptual organizations have been suggested as explanations for these illusions (e.g., Benary, 1924; Spehar, Gilchrist, & Arend, 1992; Spehar, Gilchrist, & Arend, 1995; Taya, Ehrenstein, & Cavonius, 1995).

One important point of contention is whether the relevant perceptual organization is in terms of surface properties like appurtenance (Benary, 1924), belonging (Kanisza, 1979), transparency (Adelson, 1993; Taya, Ehrenstein, & Cavonius, 1995; Spehar, Gil-

christ, & Arend, 1995), and so on, or whether the important factors are midlevel detections of junctions formed where edges meet in the retinal image. Junctions have been employed in recovering three-dimensional configurations from line drawings by Huffman (1971), Clowes (1971), Sugihara (1984), and Kanade (1980), and they have also been found to be useful in recovering shape from images, especially from raw range data. For example, in Sugihara's (1987) knowledge-guided system for range data analysis, a junction dictionary is used for the extraction and organization of edges and vertices, by consulting it to predict positions, orientations, and physical types of missing edges. These predictions guide the system as to where to search and what kinds of edges to search for, as well as how to label the extracted edges into an interpretation. An analysis of brightness illusions in terms of junctions affecting transparency and depth interpretations has previously been found useful by Adelson (1993) and Pessoa and Ross (1996).

Zaidi, Spehar, and Shy (1997) used three-dimensional configurations and two-dimensional pictures of the configurations to test whether there is a difference in induced contrast from surrounds that differ in perceived depth relationships, for example, background, coplanar, or occluding, but that are equal in retinal adjacency to the test. Each retinal image of a three-dimensional scene can be approximated by a two-dimensional projection, and retinal images on the two eyes are similar for objects at reasonably large distances from the observer. The use of pictures enables concentration on visual cues that are important in inferring three-dimensional organization from retinal images when stereoscopic disparity is not available. By using a variety of orthographic projections they were able to identify visual cues that enhance or inhibit induced contrast independent of the viewing angle and the presence or absence of other depth or perspective cues.

In the pictures in Fig. 17.1, to every observer, the gray test regions seemed to form part of the H, that is, to belong with the vertical bars. Domination of the induced effect by the background thus refutes Benary's notion that induced contrast is enhanced by

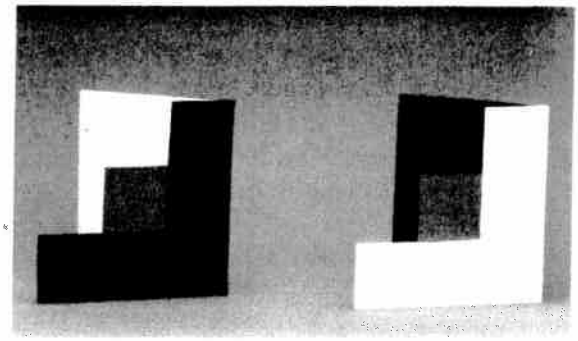

Figure 17.14: This depiction of the Benary cross illusion refutes explanations that are based on "belonging" or coplanarity (e.g., Benary, 1924), because the L-shaped region has less induced effect despite appearing coplanar and "belonging" with the test region (Zaidi, Spehar, & Shy, 1997).

appurtenance. Since the test regions were seen as parts of coplanar H's, these pictures also provide a counterexample to Gilchrist's (1977, 1980) coplanar ratio hypothesis. Another example where a larger inducing region has less effect despite appearing coplanar and belonging with the test region is shown in Fig. 17.14. The figure consists of a photograph of two figurally identical three-dimensional configurations, each consisting of an L shape in the foreground, a gray square embedded in the crook of the L, and a larger square flap at an angle behind the L. Induced contrast from the background flaps dominates that from the flanking L. In Figs. 17.1 and 17.14, the phenomenal effect was the same whether the three-dimensional configurations were viewed monocularly or binocularly. In terms of a junction-based analysis, both of these figures demonstrate that the largest effect on induced contrast is an attenuation across the edge of the test that is adjacent to the top of each T junction. Zaidi, Spehar, and Shy (1997) found that this rule held for all of the pictures that they tested, irrespective of any inferred surface properties or depth planes. They also showed that this rule could explain White's illusion and all of Benary's illusions.

In a world composed of three-dimensional objects, the top edges of the T junctions are likely to signify

occluding or more distant surrounds. In both of these cases, it would be advantageous to increase contrast from proximate backgrounds to facilitate figure-ground segregation. In fact, given the high utility of this segregation, especially for objects at a distance at which there are no other cues to the three-dimensional configuration, it may be useful for the visual system to always function as if T junctions separate the foreground from the background. The cost of illusory brightness differences in perceptually coplanar displays is probably much smaller than the benefits of a simple strategy that enhances figure-ground segregation.

Two pieces of evidence that were already discussed further support Chevreul's notion that the main functional role of color induction is to perceptually enhance marginal figure-ground differences. When the test and surround are at the same mean level, the amplitude functions in Fig. 17.3 show that fractional induced chromatic contrast is greatest for the smallest excursions in the color of the surround. In addition, the results in Fig. 17.12 show that induction accentuates differences most from those elements of the surround that are at luminance levels closest to the mean level of the test.

Figures 17.1 and 17.14 show that figural cues have to be factored into any adequate explanation of color induction. It is possible that incorporating the influence of midlevel perceptual cues like T junctions on spatial summation and lateral gain controls will eventually lead to mechanistic models that can explain brightness induction in natural settings.

Induction and color constancy

In most settings, even though changes in the spectral reflectances are possible (Nassau, 1983), particularly for living organisms (Zaidi & Halevy, 1993), in general a change in perceived colors in objects is going to be due to a change in illumination. A visual system can be said to possess the property of color constancy if the color percepts assigned to individual objects are invariant across illumination conditions. In terms of the responses of neurons or signal processing units, color constancy results if, at some stage of the visual system, the neurally transformed signals from objects in a scene vary by less than a discriminable difference across varying illuminations (Ives, 1912b). The logical first step in an analysis of color constancy is to specify the problem in terms of changes in neural signals caused by illumination changes.

In Fig. 17.15 are shown the excitations of the three cone classes and the cardinal mechanisms from each of 170 natural and man-made objects (Vrhel, Gershon, & Iwan, 1994) under zenith skylight (Z) and direct sunlight (T) (Taylor & Kerr, 1941). The spectrum labeled Z was measured by pointing the measuring instrument at the sky; of all of the phases of sunlight, it has the highest relative energy in the short wavelengths due to Rayleigh scattering. Direct sunlight at ground level (T) has the least relative energy in the short wavelengths and most in the long wavelengths. The plots in Fig. 17.15 represent the extreme case of comparing signals from an object under direct sunlight to that same object shadowed from the sun and reflecting pure skylight. In each panel, the solid line along the diagonal is the locus of equal signals under the two illuminants. Each point represents an individual object. The open circles at the top right corners in each of the L, M, and S panels represent direct cone absorptions of the two daylights. The most noticeable aspect of all three of the top panels is that the points lie close to straight lines, that is, there is a strong correlation ($r^2 > 0.998$) between the quanta absorbed by each cone type from different objects across daylight illumination changes. Comparable correlations between cone absorptions from other samples of objects and illuminants were found by Dannemiller (1993) and Foster and Nascimento (1994).

In the bottom left panel, the L/(L+M)-axes represent hues going from greenish to reddish (left to right and bottom to top). The points representing the objects all lie below the diagonal, indicating that the chromaticities of all of the objects have shifted toward green under illuminant Z as compared with under illuminant T. The S/(L+M)-axes represent hues going from yellowish to violet (left to right and bottom to top). The

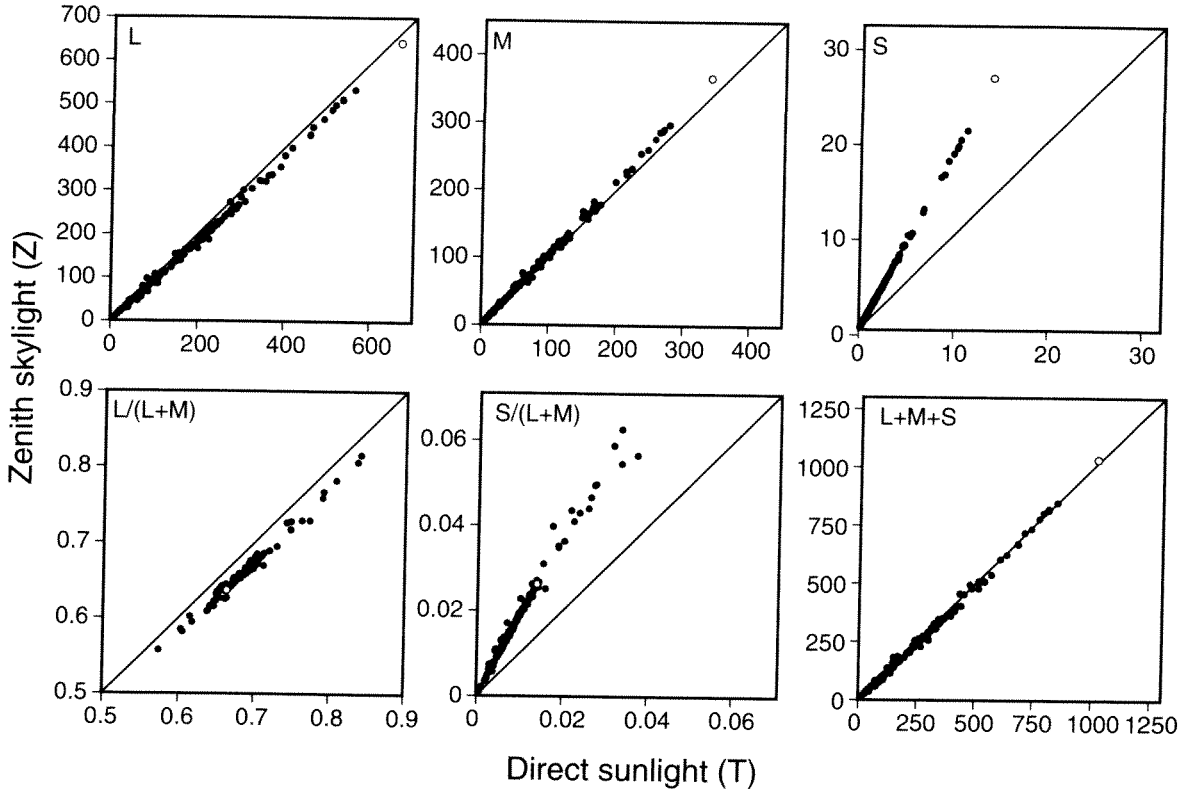

Figure 17.15: The excitation of the L-, M-, and S-cones (top three panels) and the exclusive excitation of the chromatic and luminance mechanisms along the RG, YV, and LD cardinal axes (bottom three panels) from each of the 170 objects from Vrhel, Gershon, and Iwan (1994) under illuminants T (abscissa) and Z (ordinate). The open circles represent the illuminants (Zaidi, Spehar, & DeBonet, 1997).

effect of an illuminant change is a shift in all of the chromaticities by approximately the same multiplicative factor, indicating that the chromaticities at the violet end are shifted the most. The L+M+S panel represents radiance changes at constant hue and saturation, and is dominated by L- and M-cone absorptions. Because the daylight spectra were equated for illuminance, a shift from T to Z causes almost no change in the total cone absorptions.

The illuminant-caused shifts in chromatic signals correlate well with everyday observations. Grass and bricks look appreciably bluer in shadows and yellower in sunshine. More systematic documentation is available in paintings made in the open air from the second half of the nineteenth century. A good example is Corot, who first used to sketch the scene and then painted it patch by patch, reproducing colors in each patch. Corot's painting, "Island of Saint Bartollemeo," shows walls of buildings in sunlight and shade. In cases where a wall is both lighted and shadowed, only on the shadowed bricks can be found traces of blue pigment. Da Vinci had remarked that a white dress appeared blindingly bright on the side turned toward the sun, but bluish on the other side. Similarly, Delacroix (1937) wrote: "The true colour of the flesh can be seen only in the sun and in the open air. If a man puts his head out of a window, its colouring is quite different from what it is indoors. Which shows the absurdity of studies done in a studio, where each one does his best to reproduce the wrong colour."

Given the above observations, the preoccupation with color constancy is perplexing. However, since the

colors of objects depend not only on physical absorptions, but also on neural transformations (Monge, 1789; Chevreul, 1839; Ives, 1912b), it is worth asking whether any neural transformations help in discounting the effect of illumination changes. Because the shifts in Fig. 17.15 are systematic, there is a chance that the human visual system can attenuate their perceptual effects through the use of simple adaptation strategies without having to estimate reflectance or illumination spectra, as in the schemes proposed by Maloney and Wandell (1986) and D'Zmura and Iverson (1993a; b). An adaptation process that modifies cone signals so that they fall along the diagonals in the top panels of Fig. 17.15, or modifies second-stage signals so that they fall along the diagonals in the bottom panels, will lead to color constancy. As a corollary, any empirically measured color constancy will be consistent with models involving either type of these mechanisms, and to identify whether cone or second-stage adaptation mechanisms are responsible for constancy will require additional experiments like those done by Pugh and Mollon (1979) or Zaidi, Shapiro, and Hood (1992).

Historically, computational schemes for human color constancy have involved early adaptation mechanisms. In terms of our linking hypothesis, color constancy would be achieved at an early stage if neural processes equated first- or second-stage signals from each object in a scene across illumination conditions. As a result of this processing, the outputs of at least the second-stage mechanisms (and possibly even the first stage) should be transformed under each illuminant in a manner that, when plotted similar to Fig. 17.15, all of the points should fall on the diagonal of unit positive slope. The simplest mechanism that has been proposed for accomplishing this purpose is Von Kries's adaptation (Ives, 1912b; Brill, 1995), where each photoreceptor signal is gain-controlled by its own time-integrated signal, that is. for each object the signals (L, M, S) are transformed to:

$$\left(\frac{L}{\int L dt/L_E}, \frac{M}{\int M dt/M_E}, \frac{S}{\int S dt/S_E}\right). \quad (19)$$

For a steady uniform field, the value of each integral is equal to the cone absorption from that field, and the ratio of cone absorptions for the integrated value is transformed to be equal to the ratios for an equal-energy light (L_E:M_E:S_E). This transformation could thus provide a simple explanation for the progressive desaturation of the perceived color of a continuously viewed, uniformly colored field (Vimal, Pokorny, & Smith, 1987). In his numerical simulations of color constancy, Ives (1912b) assumed that the integral for each photoreceptor was equal to the quantal absorptions by that class of receptors from the steady illuminant. Thus the result of the transformation was to make the illuminant appear achromatic. When Ives's assumption was applied separately for the two illuminants to signals from each object in Fig. 17.15 by Zaidi, Spehar, and DeBonet (1997), the result was rigid rotations of the lines between (0,0) and the open circles representing the illuminants in the S, M, and L panels, in a manner that the open circles were shifted to the unit diagonals. Since all of the points representing individual objects lie on or close to these lines, the transformed chromatic signals from individual objects were also fairly well equated across the illumination conditions, thus predicting color constancy. However, there are a number of conceptual problems in accepting this transformation as an explanation of human color constancy. First, the values of the integrals for free viewing of a variegated scene are difficult to predict a priori. In viewing a variegated scene, the ratios of time-integrated cone absorptions will only be equal to the absorption ratios of the illuminant spectrum if the integrated object reflectance spectrum for each photoreceptor is uniform, a condition that is unlikely for most natural scenes, even with the spatial averaging of reflectances due to active scanning (Brown, 1994). In reality, the spatially local values of the integrals will vary across the visual field; and to the extent that the gain for each photoreceptor is set by the spatially local signal that it receives from the particular region imaged on it, the transform in Eqn. (19) will shift the chromaticity of that object toward the achromatic point. A realistic version of this transform will thus not lead to color constancy. The second problem

has to do with the stage in the visual system that is important for color constancy transformations. It is difficult to imagine why an equal-energy light should have a privileged status for an individual photoreceptor, that is. there is no theoretical justification for the L_E, M_E, and S_E terms in the denominators of Eqn. (19). On the other hand, in color-opponent cells, the achromatic signal can have a privileged position as the zero point toward which the response of the system is shifted by a high-pass temporal filter after opponent combination. However, in a variegated scene, discounting the integrated values of opponent signals creates problems similar to those discussed in the context of the integrals in Eqn. (19). Since it is unlikely that the integrals of the opponent signals will be proportional to the values from the illuminant, spatially local adaptation will shift all colors toward the achromatic point. This would be consistent with and an alternative explanation for the progressive desaturation of a colored scene that is stabilized on the retina, but similar to local photoreceptor adaptation it could only equate chromatic signals across illuminants at the cost of losing all perceived color differences in the scene. Third, and most importantly, the empirical results reviewed in this chapter show that when viewing a variegated field neither the appearance of colors nor the limen of discrimination are determined by the space-averaged level of stimulation. Consequently, models of color adaptation or constancy that rely on spatial- and/or temporal-integrated levels as the controlling parameters (Judd, 1940; West & Brill, 1982; Land, 1983; Worthey, 1985; Dannemiller, 1989; Brill, 1990; Valberg & Lange-Malecki, 1990; Brainard & Wandell, 1992; Finlayson, Drew, & Funt, 1993, 1994; Finlayson & Funt, 1996) may be consistent with some sets of empirical data, but are not sufficient to explain the results of experiments that isolate individual color mechanisms and adaptation processes (Zaidi, Spehar, & DeBonet, 1997).

It has often been proposed that color induction can lead to color constancy (e.g., Valberg & Lange-Malecki, 1990; Walraven et al., 1991). This assertion has usually been based on the results of studies that measure perceived shifts in colors of just one test patch rather than over the whole scene, and it seems irreconciliable with the finding that juxtaposing two patches shifts their appearances in complementary color directions (Chevreul, 1839; Krauskopf, Zaidi, & Mandler, 1986). Simultaneous color induction will shift the signals from juxtaposed objects in opposite directions, and therefore we cannot discount the effect of an illumination change by shifting signals from all objects in a scene in a correlated fashion, as, for example, toward the diagonals in Fig. 17.15. In some cases it is possible that induced contrast will counter a shift in the spectrum of the illuminant. This discounting is most likely to occur for unsaturated hues that are surrounded by more saturated hues. For saturated hues, the induced shift is more likely to be in a direction that exacerbates the effect of the illuminant change. Constancy of the appearance of an individual test patch could therefore be due to color induction but is unlikely to be a good measure of color constancy over the extent of a variegated scene. This objection applies particularly to methods that measure the achromatic loci of a test patch under different illuminants (e.g., Brainard & Speigle, 1994).

Zaidi (1998) argued that, since objects do appear to be of different colors under different illuminants, the nervous system could potentially derive information that aids in the recognition of both objects and illuminants by comparing the altered appearances of objects under different illuminants. Figure 17.15 shows that changes in the spectral composition of the illumination on a set of objects lead to affine transformations of the set of object chromaticities. Affine transformations have well-defined invariants, and these invariants can be used to derive the transformation parameters. This was accomplished by two shape-alignment types of algorithms that succeeded in identifying objects with identical reflectance functions, and also derived the relative chromaticities of the two illuminants. Because information about objects and illuminants is useful in many different tasks, it would be more advantageous for the visual system to use such algorithms to extract both sorts of information from retinal signals than to automatically discount either at an early neural stage.

Neural mechanisms of color induction

To paraphrase Koffka (1963), the central questions of visual perception are why things look the way they do, and why they look the way they are. Induced contrast and its contribution to appearance are broad enough topics, that this chapter presents only one of a number of essentially nonoverlapping treatments (compare, e.g., Gilchrist, 1994). The construction of appearance by the visual system involves memory, hard-wired perceptual priors, and cognitive inferences (Griffiths & Zaidi, 1999) as well as successive stages of neural transformations of photoreceptoral signals. We have concentrated on developments from the last decade that have led to progressively more general mechanistic models that delineate processes occurring at various stages of the visual system.

The result that color induction is due to lateral connections neither at the cone nor at the opponent mechanism level, but rather at the level of mechanisms tuned preferentially to many different color directions, implicates the visual cortex as the site of these interactions. It seems not to be generally recognized that human observers have exquisite memory for colors (Sachtler & Zaidi, 1992) and are intolerant of spurious hue changes (Li & Zaidi, 1997), and that such changes would occur frequently if appearance was determined by only two independently adapting color-opponent mechanisms.

The fact that Mach bands do not contribute substantially to the induced effect on spatially extended tests also points to a postretinal site, as do the estimated brightness induction space constants of 0.74 and 0.29 deg, centered on the fovea, for two observers. These estimates, however, are smaller than the space constants of 1.23 and 0.74 deg estimated for contrast–contrast induction for the same observers by DeBonet and Zaidi (1997).

The spatial summation of brightness induction is interesting because it indicates suprathreshold spatial additivity across a fairly large region of visual space. Early stages of the visual system are composed of neurons that have small receptive fields and fairly narrow bandpass sensitivity in the spatial frequency domain (Shapley & Lennie, 1985). In addition, psychophysical results are consistent with a small number of classes of spatial frequency selective mechanisms functioning independently at threshold (Graham, 1989). These mechanisms are generally thought to have spatial frequency bandwidths that are less than one octave wide. In brightness induction, the spatial additivity covered spatial scales over a range of six octaves, from 0.05 to 3.2 cycles/deg. Thus, at some stage of the visual system that is important for the computation of brightness contrast, the outputs of these mechanisms must be summed in a fairly simple manner to lead to the point-by-point additivity of lateral effects.

The spatial summation results may lead one to think that brightness induction could simply be the outcome of processing by neurons with large center-surround receptive fields that add or average the brightness of the surround. However, the results of induction from textured surrounds rule out this possibility and instead require gain controls on the connections between neurons covering spatially circumscribed areas. The effects of figural inferences on induced contrast also argue against simple lateral inhibition over extended cortical areas. In the present state of knowledge about visual neuro-physiology, it is not possible to even speculate about possible neural mechanisms for extracting T junctions or other figural features. It is clear, though, that perceived brightness and color are computed from the retinal image in conjunction with other perceptual attributes.

Acknowledgments

This work was supported by NEI grant EY07556. Michael Shy helped with the production of this manuscript.