PERGAMON

Vision Research 40 (2000) 2885–2903

Vision Research

www.elsevier.com/locate/visres

Control of chromatic adaptation: signals from separate cone classes interact

Peter B. Delahunt *, David H. Brainard

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

Received 7 December 1999; received in revised form 11 April 2000

Abstract

Match stimuli presented on one side of a contextual image were adjusted to have the same appearance as test stimuli presented on the other side. Both full color and isochromatic contextual images were used. Contextual image pairs were constructed that had identical S-cone image planes, while their L- and M-cone image planes differed. The data show that the S-cone component of the matches depends on the L- and M-cone planes of the contextual image. This dependence means that matches obtained using isochromatic stimuli (lightness matches) may not be used directly to predict full color matches. © 2000 Elsevier Science Ltd. All rights reserved.

Keywords: Colour; Constancy; Adaptation; Appearance; Lightness

1. Introduction

Color appearance depends on context. A classic example is simultaneous color contrast, where the immediate surround of an image region affects its color appearance (see for example Evans, 1948). The immediate surround, however, is not the only relevant contextual factor (e.g. Gilchrist, 1980, 1988; Wesner & Shevell, 1992; Adelson, 1993; Kraft & Brainard, 1999). Understanding color context is daunting because there are an astronomical number of possible contexts, even when the regularities of natural images are taken into account. To make progress, it is necessary to develop principles that allow us to predict the effect of many contexts from measurements of a few (Krantz, 1968; Brainard & Wandell, 1992; Baüml, 1995; Chichilnisky & Wandell, 1995, 1999).

Color context effects have been studied in two somewhat distinct traditions. In one tradition, the stimuli are restricted to be isochromatic, so that they vary only along the physical dimension of intensity. In these studies, judgments are made of a single attribute of color appearance, lightness (or sometimes brightness). Restricting the stimuli in this way simplifies stimulus specification and control and therefore allows easier exploration of spatially rich contexts. We refer to studies where the stimuli are restricted to be isochromatic as lightness experiments (e.g. Land & McCann, 1971).

It is not clear how results from lightness experiments generalize when both intensity and chromaticity vary. Investigators in a second tradition have worked directly with full color stimuli. The generality of this approach is purchased at the cost of increased stimulus complexity and in general full color studies have explored simpler spatial configurations than lightness studies.

Can results from lightness experiments be used to predict what will happen in full color experiments? If there is a principle that lets us generalize lightness results the number of contexts that must be studied is greatly reduced. At the first stage of color processing by the human visual system, the image is registered by the L-, M-, and S-cones. At this stage, therefore, the neural representation of the image consists of three distinct image planes. Von Kries (1902) suggested that perhaps context effects occur independently within each of these

* Corresponding author. Tel.: +1-805-8935108; fax: +1-805-8934303.

E-mail address: delahunt@psych.ucsb.edu (P.B. Delahunt).

PII: S0042-6989(00)00125-5

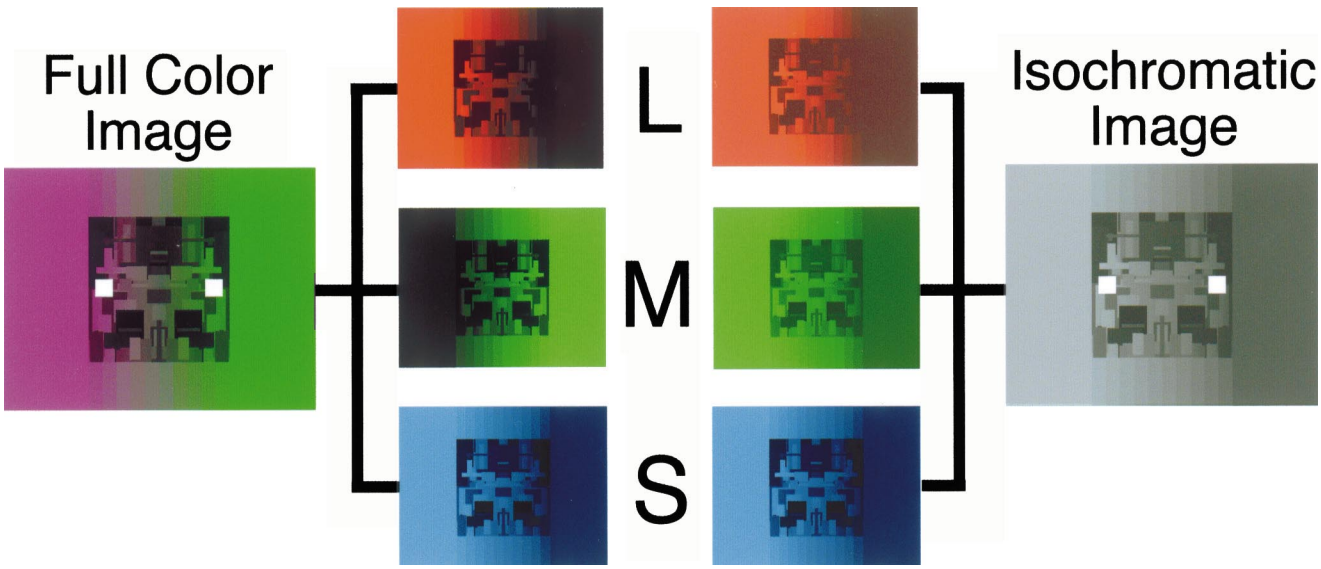

Fig. 1. This figure illustrates the contextual images used and the idea that any image may be represented in terms of three image planes. A full color contextual image is shown on the left along with its L-, M-, and S-cone planes. An isochromatic image is shown on the right, along with its three cone image planes. Note that the S-cone planes of both contextual images are identical. The L- and M-cone image planes differ. For the isochromatic image, they were constrained so that the image was isochromatic. The test and match stimulus locations are shown by the white squares in the contextual images.

image planes[1]. If so, then it is sufficient to characterize the effect of context for each image plane separately. We refer to this general idea as the *independence hypothesis*. The independence hypothesis is fundamental to Land's retinex theory (Land & McCann, 1971; McCann, McKee & Taylor, 1976; Land, 1986; see Brainard & Wandell, 1986). The notion of independence was also a key element in Stiles (1949, 1959) analysis of increment threshold data, although in his hands independence applied to the π-mechanisms rather than to cone signals.

Failures of the independence hypothesis have been demonstrated using threshold data. For example, Pugh and Mollon showed that there are two sites that regulate signals from S-cones. At the second site, S-cone thresholds are affected by signals that originate in the L- and M-cones (Pugh & Mollon, 1979). Failures of independence are also indicated by experiments that employ appearance tasks (Cicerone, Krantz & Larimer, 1975; Werner & Walraven, 1982; Singer & D'Zmura, 1995; Baüml & Wandell, 1996; Chichilnisky & Wandell, 1996; Poirson & Wandell, 1996; see Section 8). There may, however, be conditions of interest where independence holds to good approximation. Creutzfeldt, Lange-Malecki and Dreyer (1990) reported that lightness matches obtained separately with red, green, and blue isochromatic stimuli could be used to predict full color matches in the context of full color images. Chichilnisky and Wandell (1995) showed that independence holds for a large set of asymmetric (haploscopic) color matches. Since much recent work on color context effects has been motivated by an interest in color constancy, we wanted to examine the independence hypothesis carefully for stimuli typical of those used in color constancy research.

Here we report experiments where subjects judged color appearance in the context of both full color and isochromatic images. The images were designed to allow us to test the independence hypothesis directly, with minimal assumptions about how context affects appearance. More specifically, we studied context effects for Mondrian-style images (McCann et al., 1976; Arend & Reeves, 1986; see Fig. 1) presented on a computer controlled color monitor and investigated whether signals from L- and M-cones influence signals from S-cones. In addition, we used our data to test parametric models of how context regulates signals originating in the cone photoreceptors.

2. General methods

2.1. Overview

Asymmetric matching was used to measure how context affects appearance. Test and match stimuli were placed in Mondrian-style contextual images (see Fig. 1). The test stimulus was placed towards one side of the

[1] There are two variants of the von Kries hypothesis. Both variants assert that any effect of context acts by changing the gains on signals from the three classes of cones. The strong form of the hypothesis asserts that the gains for each cone class are determined only by signals from the same cone class, while the weak form is less restrictive and allows for the possibility that signals in one cone class can affect the gains of another. In this paper, we are primarily concerned with a generalization of the strong form.

image while the match stimulus was placed towards the other. The images produced a context effect, so that test and match stimuli that were physically identical appeared different. The observer's task was to adjust the match stimulus so that it appeared identical to the test stimulus, and this task was repeated for multiple test stimuli. The light reaching the observer from the test and match stimuli was characterized by *LMS* cone coordinates, specified here by column-vectors $\boldsymbol{R}_t = [R_{tL}\ R_{tM}\ R_{tS}]^T$ and $\mathbf{R}_m = [R_{mL}\ R_{mM}\ R_{mS}]^T$, respectively. Here the superscript T denotes vector transpose. Thus for each contextual image the observer established pairs of stimuli $\{\mathbf{R}_t, \mathbf{R}_m\}$ that appeared identical when viewed in the context of that image.

As illustrated by Fig. 1, any image can be decomposed into three image planes, one for each of the three cone types. Given this, it is possible to construct a pair of contextual images, one full color and one isochromatic, that have identical S-cone planes but differ in their L- and M-cone planes. For such a *contextual image pair*, the independence hypothesis makes a clear prediction about the relation between asymmetric matches collected in the context of the full color image and those collected in the context of the isochromatic image.

Suppose that test stimuli with the same S-cone components are used with both the full color and isochromatic contextual images. Let $\{\mathbf{R}_t^a, \mathbf{R}_m^a\}$ represent a match set in the full color image and $\{\mathbf{R}_t^b, \mathbf{R}_m^b\}$ represent a match set in the isochromatic image. The independence hypothesis implies that when the S-cone planes of the contextual images are identical, then the S-cone components of $\mathbf{R}_m^a$ and $\mathbf{R}_m^b$ must be equal. This implication is derived formally in the Appendix and is the empirical property tested in this paper.

2.2. Stimuli

Fig. 1 illustrates the type of contextual image used in the experiments. The contextual image was presented either on a computer controlled monitor (Experiments 1–3) or using a back-projection system (Experiment 4). Contextual image pairs were constructed as follows. First, a full color image was created as described below. Then a second, isochromatic, contextual image was derived from the first. The second contextual image was constructed to have the properties required for our tests of the independence hypothesis. For example, given a full color contextual image we could extract the S-cone image plane $\mathbf{I}_S$ and create L- and M-cone image planes whose spatial structure was identical to the S-cone image plane: $\mathbf{I}_L = w_L\ \mathbf{I}_S$ and $I_M = w_M\ \mathbf{I}_S$. The image resulting from combining these $\mathbf{I}_L$, $\mathbf{I}_M$, and $\mathbf{I}_S$ image planes is isochromatic, can be made to have a desired chromaticity through appropriate choice of weights w_L and w_M, and has an S-cone image plane identical to the full color contextual image from which it was derived.

For our purposes, the first image in each contextual image pair need not satisfy any particular constraints. We did think it desirable, however, that these images be similar to images used in many studies of color constancy (see Section 1) and that they produce measurable context effects. The central portion of the contextual images consisted of simulations of illuminated surfaces, with the spectral properties of the illuminant varying from the left to the right side of the image. This central portion was embedded in a surround intended to enhance the context effect.

For the central portion, the simulated surfaces were flat, matte, of rectangular shape, and spectrally non-selective. The width of each rectangle was chosen randomly from the range 1–129 pixels; the height was chosen randomly from the range 1–257 pixels (100 pixels = 36 mm). The surface reflectance of each rectangle was chosen randomly from the range 0–1. The surfaces on the right side of the central portion were then obtained by mirror reversal of the left-side, so that the simulated surfaces were left–right symmetric in the image.

The simulated illuminant had a spatial gradient from the left to the right side of the image. Because we simulated only non-selective surfaces, we use cone coordinates rather than spectral functions to specify the simulated illuminants. The illumination for each image was determined by the *LMS* coordinates of the illuminant at the left and right edges of the central portion. The *LMS* coordinates of the illuminant at any horizontal location within the central portion were then obtained by linear interpolation. The *LMS* coordinates of the illuminant were constant from top-to-bottom within the central portion.

The central portion of the image was embedded in a surround. The surround may also be thought of as a simulation of illuminated surfaces. In the case of the surround, however, all simulated surfaces had a reflectance of 1. The area to the left of the central portion was illuminated by the same simulated illuminant used at the left most edge of the central portion. The area to the right was illuminated by the same simulated illuminant used at the right most edge of the central portion. Above and below the central portion, the illumination was varied in ten equal width steps from left to right. This gave the surround a simulated illuminant gradient from left to right similar to that of the embedded central portion.

It is convenient to specify the contextual image in terms of the average *LMS* coordinates at the locations where the test and match stimuli were placed. These coordinates were obtained by averaging the values contained within rectangles of the same width and horizontal positions as the test and match areas, and with the same height as the full contextual image. These values are provided in Table 1 for each contextual image used and also indicated on the data figures.

Table 1
The *LMS* and CIE *xyY* coordinates for each contextual image are provided[a]

		L	*M*	*S*	*X*	*y*	*Y*
Full color							
Experiments 1 & 3(a)	Test	0.0153	0.0118	0.0217	0.2804	0.1961	9.67
	Match	0.0257	0.0254	0.0065	0.2947	0.4944	17.97
Experiments 2 & 3(b)	Test	0.0363	0.0328	0.0090	0.3385	0.4615	24.56
	Match	0.0369	0.0333	0.0420	0.2613	0.2455	24.74
Experiment 4	Test	0.9533	0.6138	1.1410	0.3147	0.1846	563.64
	Match	0.9356	0.9029	0.3606	0.2901	0.4463	701.74
Supplemental	Test	0.0308	0.0280	0.0062	0.3441	0.4875	20.88
	Match	0.0090	0.0082	0.0219	0.2230	0.1468	5.83
Isochromatic							
Experiments 1 & 3(a)	Test	0.0313	0.0284	0.0217	0.2864	0.3214	21.12
	Match	0.0094	0.0085	0.0065	0.2874	0.3222	6.34
Experiments 2 & 3(b)	Test	0.0128	0.0115	0.0086	0.2904	0.3242	8.60
	Match	0.0602	0.0543	0.0415	0.2884	0.3216	40.54
Experiment 4	Test	1.4026	1.1692	1.1380	0.2936	0.2562	863.03
	Match	0.3358	0.2727	0.3575	0.2906	0.2579	228.29
Supplemental	Test	0.0313	0.0284	0.0217	0.2864	0.3214	21.12
	Match	0.0094	0.0085	0.0065	0.2874	0.3222	6.34

[a] The values were obtained from the image regions near the test and match locations, as described in the text. Luminance (Y) is in units of cd/m^2.

All experiments were run in darkened rooms. For experiments 1–3 the stimuli were rendered on a computer controlled monitor (Power Macintosh 7200/120; ViewSonic PT770 17 in. color monitor; refresh rate 75 Hz). Head position was stablized using a chin rest placed 40 cm from the monitor. The square central portion of the image was of side 132 mm, which subtended a visual angle of 18.7°. The match and test stimuli were square and of side 18 mm (visual angle of 2.5°). They were placed 110 mm apart, corresponding to a center-to-center separation of 15.7°. The outer horizontal and vertical dimensions of the surrounding part of the contextual image were 300 and 235 mm, respectively. The monitor was viewed monocularly through a square aperture in a black reduction screen placed 20 cm from the monitor. The screen measured 72 cm (w) × 91 cm (h) and the square aperture was of side 12 cm. The aperture was positioned so that the observer was only able to view the internal area of the monitor.

In Experiment 4, the images were displayed using a back projection system (based on a General Electric Model LCD 36 projector). This apparatus was developed by Rasengane, Palmer and Teller at the University of Washington and is described in detail elsewhere (Rasengane, Palmer & Teller, 2000). They kindly allowed us to use it for Experiment 4. Although the spatial and chromatic control provided by this system is less precise than that obtained with a typical CRT, the projection system provided high luminance. Observer head position was stabilized with a chin rest placed 77 cm from the projector screen. The square central portion of the image was of side 200 mm, which subtended a visual angle of 14.75°. The match and test stimuli were of side 26.5 mm and subtended a visual angle of 1.93°. They were placed 130 mm apart, corresponding to a center-to-center separation of 9.78°. The outer horizontal dimension of the surrounding part of the contextual image was 590 mm. Because of the limited size of the projection screen, the vertical portion (that part above and below the central portion) of the surrounding part of the contextual image was not visible. The back projection screen was viewed monocularly.

2.3. Observer's task

The observer's task was to adjust the match stimulus so that it appeared the same as the test stimulus. The test and match were shown in alternation for one second each, and the observer was instructed to fixate back and forth between the test and match locations. The perceptual criterion used by the observer to evaluate appearance can influence the results in an asymmetric matching experiment. For example, Arend and Reeves (1986) and Baüml (1999) distinguished appearance matches, where the observer was instructed to judge the appearance of the light reaching the eye, from paper matches, where the observer was instructed to judge the identity of the simulated surface. Our current understanding of what stimulus conditions support an effect of instructions on asymmetric matches is not very advanced. Our instructions were in terms of appearance and did not ask observers to consider surface identity.

2.4. Procedure

At the start of the experiment, the observer adapted to the contextual image for 1 min. The observer used a game controller to adjust the match stimulus. Adjustments varied the CIELAB $L^*a^*b^*$ coordinates of the match. Approximately red/green (a^*) and yellow/blue (b^*) adjustments were made by pushing the appropriately colored buttons on the game controller, while luminance adjustments were made using the joy stick control. The XYZ coordinates of the white point of the monitor (RGB values set to maximum) were used as the white point for the conversion to $L^*a^*b^*$ coordinates. The observer could at any time choose one of three step sizes for the adjustments (see Table 2). Note that the observer could adjust all three coordinates for both full color and isochromatic contextual images.

For each test stimulus and contextual image, matches were obtained in two sessions. Between five and eight test stimuli were used in each session. Each test was matched once in a session in random order. We refer to these as the normal matches. After making matches to each test in random order, a second set of matches was made to test for symmetry. The test and match locations were swapped in the contextual image, and the initial matches were used as the tests. The order of the second set of matches was also randomized within each session. We refer to the second set of matches made in each session as symmetry-check matches. These symmetry-check matches are not well-suited for testing the independence hypothesis, since the test stimuli are not constrained to be the same for the two contextual images. The symmetry-check matches are not considered further in this paper. They are, however, available along with the normal matches on the world wide web at URL http://color.psych.ucsb.edu/independence/.

In the second session for each test stimulus and contextual image, the contextual image was reversed from left to right on the monitor, with a corresponding reversal of test and match locations. This was done to counterbalance the data against any left–right asymmetries in the apparatus or observers' visual systems. Note that this manipulation is independent of the manipulation of matching direction (normal versus symmetry-check) within a contextual image. Data from the two sessions were averaged.

Table 2
The table provides the adjustment step sizes used for each of the $L^*a^*b^*$ coordinates

	L^*	a^*	B^*
Coarse	4.50	6.00	6.00
Medium	0.75	1.00	1.00
Fine	0.15	0.20	0.20

2.5. Calibration

Conversion between cone coordinates and monitor settings was achieved using the general model of monitor performance and calibration procedures described by Brainard (1989). Calibrations were performed every 2 weeks using a Photo Research PR-650 spectraradiometer. Spectral measurements were made at 4 nm increments between 380 and 780 nm but interpolated with a cubic spline to the CIE recommended wavelength sampling of 5 nm increments between 380 and 780 nm. To correct the data for any small violations of the calibration assumptions, test and match stimuli were replayed after each session and measured directly using the radiometer. For some experiments direct measurements were made of the contextual image as well. The Smith–Pokorny cone fundamentals were used to calculate cone coordinates from spectral light measurements. For calculations, each cone spectral sensitivity was normalized to a peak of one and spectra were expressed in units of Watts/(sr-m^2-nm). CIE XYZ and xyY coordinates were computed with respect to the CIE 1931 color matching functions.

The same procedures were used to calibrate the projection system.

3. Experiment 1

3.1. Specific methods

The contextual images are specified in Table 1 (see Section 2). For the full color image the L- and M-cone coordinates increased from test-to-match while the S-cone coordinates decreased. For the isochromatic image, the coordinates of all three classes of cones decreased from test-to-match. The S-cone plane for both images was identical up to small deviations due to quantization and limits on calibration precision.

Eight test stimuli were used for each contextual image. The eight tests used for each contextual image had a constant chromaticity and varied in luminance (see Table 3). The S-cone components of the tests were the same for the full color and isochromatic images. The tests used for the two contextual images had different chromaticities from each other, and the tests used for the isochromatic image did not have the same chromaticity as the image itself.

Three males served as observers for this experiment. One observer (PBD, age 34) was the first author, a second was experienced but naive as to the specific questions under consideration in this study (JMK, age 32), and a third was naive (MTR, age 21). All were color normal as defined by results from the American Optical Company H–R–R (Hardy, Rand and Rittler) pseudoisochromatic plates and the Ishihara color plates.

Table 3
The LMS and CIE xyY values are provided for the test stimuli used in each experiment[a]

	Test color	No. of tests	Full color			Isochromatic		
			x	y	Y	x	y	Y
Experiment 1	Gray	8	0.285	0.320	4–32	0.333	0.333	5–40
Experiment 2	Gray	8	0.285	0.320	4–32	0.285	0.320	4–32
Experiment 3(a)	Gray	5	0.285	0.320	8–26	0.285	0.320	8–26
	Red	5	0.360	0.290	8–26	0.360	0.290	8–26
	Green	5	0.290	0.430	8–26	0.290	0.430	8–26
	Blue	5	0.220	0.220	8–26	0.220	0.220	8–26
	Yellow	5	0.400	0.430	8–26	0.400	0.430	8–26
Experiment 3(b)	Gray	5	0.285	0.320	15–40	0.285	0.320	15–40
	Red	5	0.360	0.290	10–30	0.360	0.290	10–30
	Green	5	0.290	0.430	15–40	0.290	0.430	15–40
	Blue	5	0.220	0.220	4–20	0.220	0.220	4–20
	Yellow	5	0.400	0.430	15–40	0.400	0.430	15–40
Experiment 4	Gray	8	0.285	0.320	400–2000	0.285	0.320	400–2000

[a] Luminance (Y) is in units of cd/m^2.

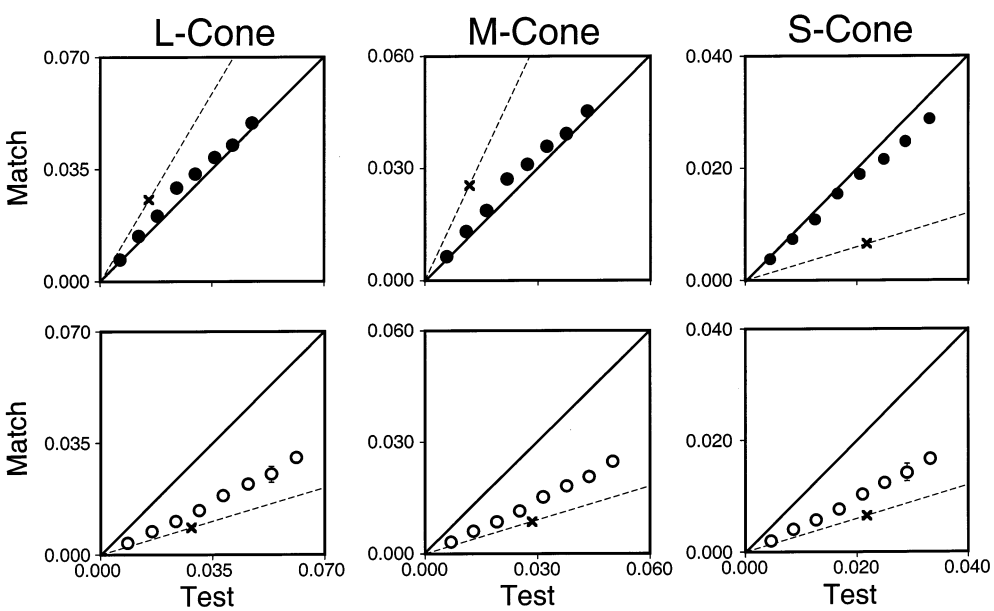

Fig. 2. This figure shows the mean L-, M- and S-cone components of the matches made in Experiment 1 for Observer JMK. The matches set by the other two observers were similar. The top panels shows the settings for the full color contextual image, while the bottom panels are for the isochromatic contextual image. The coordinates of the match stimulus are plotted against those of the corresponding test stimulus. The solid line along the positive diagonal indicate where the data would plot if there were no context effect. The thin dashed lines provide a description of the contextual image. These lines pass through the origin and a point (indicated by the x) which specifies the cone coordinates of the contextual image averaged over a region near test and match locations (see description in text and numerical values in Table 1). Where visible, error bars indicate ± 1 SEM.

3.2. Results

The results for Observer JMK are shown in Fig. 2. The top panels plot results for the full color contextual image, while the bottom panels plot results for the isochromatic image. Each panel shows test and match data for a single cone type. In these plots, a context effect is revealed if the match and test coordinates differ for at least one cone type; that is if the plotted points in at least one panel fall off the positive diagonal. A context effect is readily seen for the isochromatic contextual image. There was also an effect for the full color image, although it appears small when plotted in these coordinates. The effect of context for the full color image is more apparent when we plot the data in chromaticity coordinates. Fig. 3 shows the full color

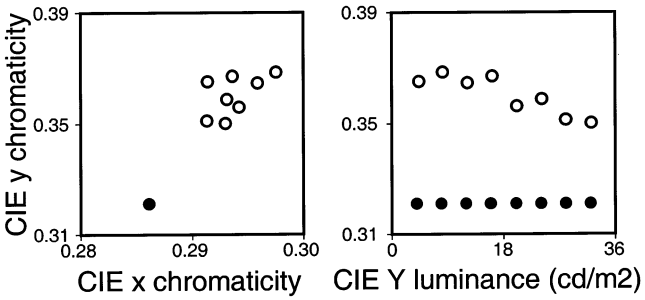

Fig. 3. The left panel shows xy chromaticity coordinates of the test and match stimuli for Observer JMK in Experiment 1. The right panel shows the same data, but y chromaticity is plotted against luminance. ● show test stimuli and ○ show matches. No error bars are shown.

data for this observer. One panel plots CIE 1931 y chromaticity against x chromaticity, while the other plots y chromaticity against luminance. The solid circles indicate the test stimuli while the open circles show the mean matches. In this coordinate system it is clear that context produces an effect — the test and match stimuli have different chromaticities. There is also a systematic effect of test luminance on match chromaticity. The matching data for all the experiments are available in numerical form on the world wide web at URL http://color.psych.ucsb.edu/independence/.

To use the data from Experiment 1 to test the independence hypothesis, we compare the S-cone component of the settings made in the full color and isochromatic contextual images. This comparison is made for all three observers in Fig. 4. Each panel plots the S-cone matches against the S-cone tests for the full color (closed circles) and isochromatic (open circles) images. If the independence hypothesis held, the closed and open circles would overlay. This is clearly not the case. The data falsify the independence hypothesis in its most general form. Signals from S-cones are affected by the L- or M-cone image planes of the contextual image. (See Section 6 below for a control that rules out the possibility that rod signals are mediating the effect.)

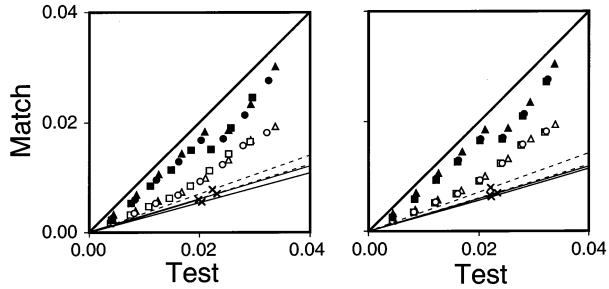

Fig. 5. The left panel of this figure illustrates the effect of shifting the S-cone sensitivity functions for one observer (MTR) using the results obtained in Experiment 1. The general format is the same as for Fig. 4. Matches set for the full color contextual image are shown by filled symbols, while those for the isochromatic contextual image are shown by open symbols. The circles replot the data from Fig. 4. The squares show coordinates computed with respect to an S-cone spectral sensitivity shifted by -10 nm while the triangles show the coordinates when the spectral sensitivity is shifted by $+10$ nm. The thin solid lines and associated x's describe the two contextual images for the -10 nm shift, while the thin dashed lines and associated x's are for the $+10$ nm shift. The figure shows that shifting the S-cone sensitivity has little effect on the basic features of the data or on the agreement between full color and isochromatic contextual images. In the latter case, note that the relevant comparison is of the two thin solid lines with each other or two thin dashed solid lines with each other. The right panel shows the same analysis for variants of S-cone sensitivity obtained by assuming different macular pigment densities (see text for details).

One concern in interpreting the results of experiments that rely on precise matching of the images seen by each class of cones is whether the conclusions are robust with respect to uncertainty in the estimates of cone spectral sensitivity. Here we are concerned with equations of images seen by S-cones. To check the dependence of the conclusion on the particular choice of S-cone sensitivities, we re-analyzed the data in two ways. First, we used shifted versions of the estimate of S-cone spectral sensitivity. The peak sensitivity of the S-cones was shifted by -10 and $+10$ nm while leaving the shape of the curve constant when plotted in energy units on a wavelength scale. This single parameter perturbation was chosen for convenience and not to model the diverse set of factors that might lead to

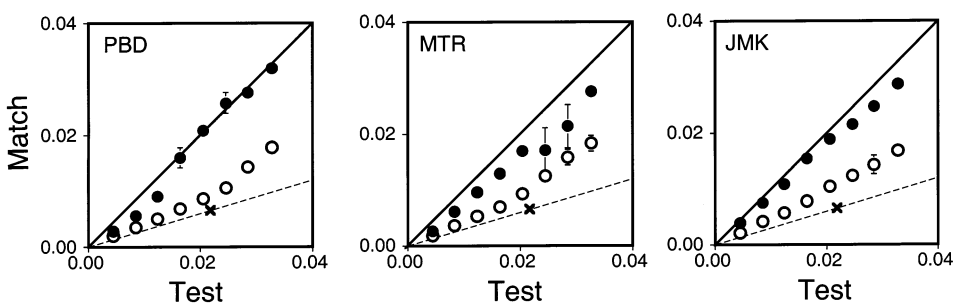

Fig. 4. The S-cone component of the matches is shown for each observer for Experiment 1. Matches set for the full color contextual image are shown by ●, while those for the isochromatic contextual image are shown by ○. Other aspects of the figure are in the same format as Fig. 2. Where visible, error bars indicate ± 1 SEM.

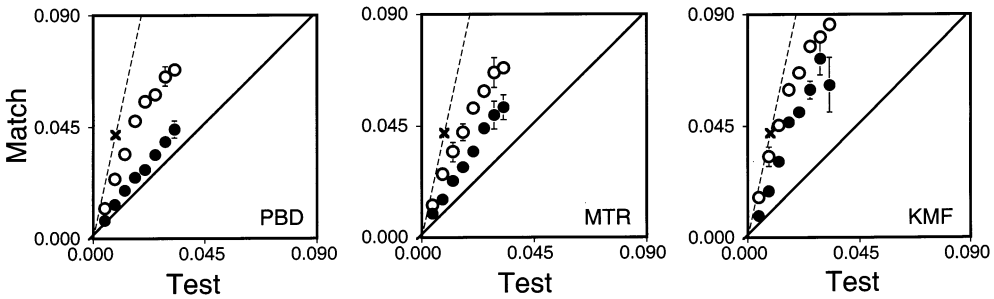

Fig. 6. The S-cone component of the settings is illustrated for each observer for Experiment 2. Same format as Fig. 4.

individual variation in S-cone spectral sensitivity (see Webster & MacLeod, 1988). The left panel of Fig. 5 shows, for one observer, the S-cone settings of the matches for full-color and isochromatic contextual images for three different versions of S-cone spectral sensitivity: the standard Smith–Pokorny estimate and the two versions obtained by shifting. Note that the separation between the settings for the two contextual images is not eliminated by perturbation of S-cone spectral sensitivities. The figure also indicates that perturbing the S-cone spectral sensitivities has a very small effect on the equality of the S-cone contextual image planes in the full color and isochromatic images. Second, we created two new versions of S-cone spectral sensitivity by assuming different densities of macular pigment. Here we adjusted the Smith–Pokorny estimate of S-cone spectral sensitivity by assuming that it incorporated the Vos estimates of macular pigment optical density (Vos, 1972) and then adjusting the density up or down by a factor of 3.16 at every wavelength. The right panel of Fig. 5 shows the same data as in the left panel but using these latter estimates of S-cone sensitivity rather than the ones obtained by shifting. Again the effect of perturbing S-cone spectral sensitivity is small relative to the experimental effect of interest.

4. Experiment 2

4.1. Specific methods

To verify the generality of the results obtained in Experiment 1, we repeated the experiment using a different contextual image pair. In this case, the full color image had a test-to-match gradient only for the S-cones (see Table 1). The same test stimuli were used for both the full color and isochromatic images, and the chromaticity of these stimuli were closely matched to the chromaticity of the isochromatic image itself.

Two of the male observers (PBD, MTR) from Experiment 1 and a naive female observer (KMF, age 21) participated in this experiment. All were color normal as defined by results from the American Optical Company H–R–R (Hardy, Rand and Rittler) pseudoisochromatic plates and the Ishihara color plates.

4.2. Results

Fig. 6 shows the data from Experiment 2 in the same form as Fig. 4. Again the independence hypothesis is clearly falsified.

5. Experiment 3

5.1. Specific methods

In Experiments 1 and 2 all test the stimuli had the same chromaticity. In Experiment 3 we used test stimuli with five different chromaticities. When presented in a neutral surround, tests of these chromaticities appeared gray, red, green, blue and yellow. The test stimuli are specified in Table 3.

One of the male observers from Experiments 1 and 2 (PBD) and two naive female observers (SIM, age 19 and LMD, age 29) participated in this experiment. All were color normal as defined by results from the American Optical Company H–R–R (Hardy, Rand and Rittler) pseudoisochromatic plates and the Ishihara color plates.

In pilot observations with the test stimuli used in this experiment, we noticed that gamut limitations sometimes prevented observers from being able to set satisfactory matches. Thus in this experiment, observers rated the perceptual quality of each match as good, medium or bad. In analyzing the results, the quality rating (good, medium or bad) of each match was checked. No matches were judged bad by any of the observers. Matches that were near the lower or upper end of the monitor gamut (one or more 8-bit settings less than 5 or greater than 250) were discarded unless they were judged good. The number of discarded matches was 0, 1 and 3% for observers PBD, SIM, and LMD respectively. In the same order, these observers rated 99, 97 and 67% of the matches as good.

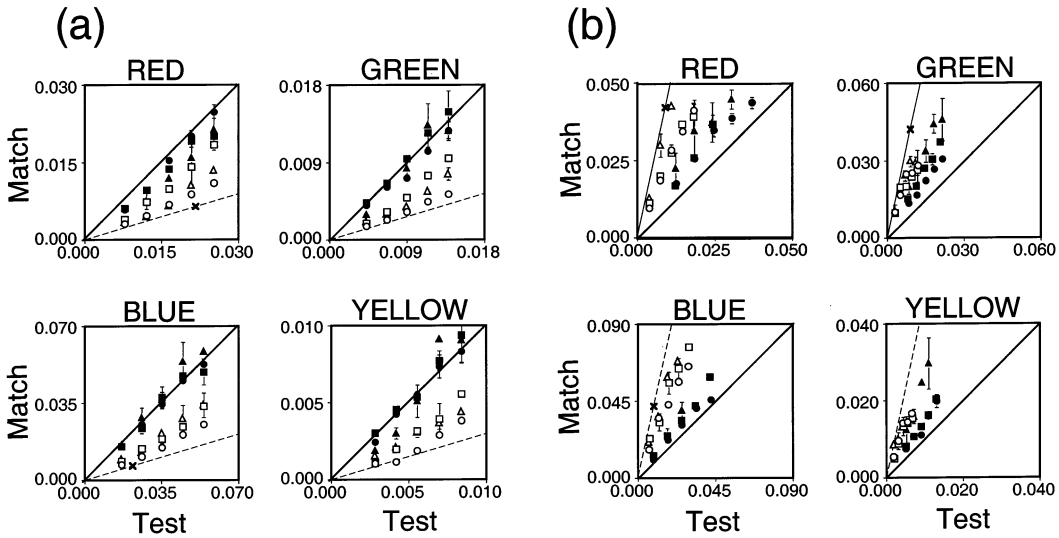

Fig. 7. The S-cone component of the matches for Experiment 3 is shown. Part (a) shows the matches made to the four sets colored test stimuli using the same contextual images as in Experiment 1. Part (b) shows the matches made to the same test stimuli using the same contextual images as Experiment 2. Each panel shows data for tests of a different chromaticity. The full color matches are illustrated by filled symbols and the isochromatic matches are shown by open symbols. The different symbols represent the three observers (PBD circles, SIM squares, LMD triangles). Other aspects of the figure are in the same format as Fig. 2. Where visible, error bars indicate ± 1 SEM.

5.2. Results

Fig. 7a shows scatter plots of the S-cone component of the matches made by all three observers for the contextual image pair used in Experiment 1. Each panel shows results for a single test chromaticity. Results are shown for the red, green, blue and yellow tests. Results for the gray tests were similar to those already shown for Experiments 1 and 2. The full color and isochromatic images produce different S-cone match components for all observers and test chromaticities, further confirming the generality of the results from Experiment 1. Similar results hold for the contextual image pair used in Experiment 2, as shown in Fig. 7b.

6. Experiment 4

6.1. Specific methods

To conclude definitively that Experiments 1–3 falsify independence, we must rule out the possibility that the observed effects are due to a rod-cone interaction. The light levels used in Experiments 1–3 are low enough (see Table 1) that rod activity could in principle affect color appearance (see for example Shapiro, Pokorny & Smith, 1996). To check for rod influence, we repeated Experiment 1 using a back-projection system that presented stimuli at considerably higher light levels. The contextual images were designed to have the same relative cone excitation coordinates as those used in Experiment 1. The test stimuli had the same chromaticity as those used for the full color image in Experiment 1. The overall luminances of both images and test stimuli, however, were much higher (see Tables 1 and 3).

One male observer (PBD) and one female observer (HLC, age 28) participated in this experiment. Observer HLC was color normal as confirmed by the Ishihara color plates.

As with Experiment 3, matches that were near the edge of the monitor gamut were discarded unless they were judged good. The number of discarded matches was 0 and 2% for observers PBD and HLC, respectively. In the same order, these observers rated 100 and 77% of the matches as good. Neither observer rated any matches as bad.

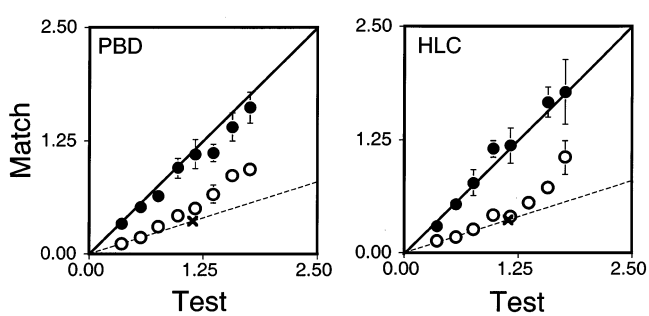

Fig. 8. The S cone component of the settings is illustrated for each observer for the rod control experiment (Experiment 4). Same format as Fig. 4. One data point for observer HLC was lost because of computer failure.

6.2. Results

Fig. 8 shows the results for both observers. They are very similar to those obtained in Experiment 1. We estimated the scotopic retinal illuminance of the contextual images using measured spectral power distributions and the formula for conversion between photopic luminance and pupil diameter provided by Pokorny and Smith (1997). Across the test and match locations in the full color and isochromatic contextual images, the minimum retinal illuminance was 3.5 log scotopic trolands. At this level, increment threshold data (Aguilar & Stiles, 1954; see Shapiro et al., 1996) indicate that the rods are well saturated: Experiment 4 shows that the results observed in Experiments 1–3 are not a consequence of rod signals.

7. Modeling the results

7.1. Preamble

The results of Experiments 1–4 show that the independence hypothesis fails. In this section, we consider the data more quantitatively. There are two distinct ways independence could fail. First, it might be the case that there is no site where signals from each cone class are regulated separately. This would occur for example, if substantial signal regulation occurs after cone signals are recombined at opponent sites. Alternatively, signals from each class of cones may be regulated separately, but the regulation may depend on signals from the other two cone classes. To separate these possibilities, we investigated whether a model consistent with separate regulation can account for the large asymmetric matching data sets collected in Experiment 3.

7.2. Models

A simple model consistent with separate regulation is the diagonal model. In this model, the match stimuli are predicted from the test stimuli through the equation

$$\mathbf{R}_m = \mathbf{D}\ \mathbf{R}_t \tag{1}$$

where **D** is a diagonal matrix. The entries of **D** capture the effect of cone specific gain changes caused by the change of context. The diagonal model expresses the weak form of von Kries' hypothesis.

Previous authors have noted that the diagonal model is inadequate to explain data collected on uniform fields. The diagonal model predicts that a plot of match versus test coordinates should always fall along a line through the origin. This prediction fails both because the best fitting line may be offset from the origin (Jameson & Hurvich, 1964; Walraven, 1976; Shevell, 1978) and because a line does not adequately describe the data (Richards & Parks, 1971; Mausfeld & Niederee, 1993; Chichilnisky & Wandell, 1996, 1999; Mausfeld, 1998; Schirillo, 1999a,b). These failures are seen in our data (e.g. Figs. 2 and 6; see also Fig. 3).

In their octant model, Mausfeld and Niederee provide a generalization of the diagonal model that can account for both types of failure (Mausfeld & Niederee, 1993; Mausfeld, 1998). We adopted the form of this model used by Chichilnisky and Wandell (1999) and formulated it to predict asymmetric matches. We refer to it here as the increment-decrement diagonal (IDD) model. In this model, each cone component of the match is predicted from the corresponding component of the test, so the model remains consistent with separate regulation. For the L-cone component, the match coordinate R_{mL} is predicted from the test coordinate R_{tL} as follows:

$$\begin{aligned} R_{mL} &= D_{L1}\ (R_{tL} - R_{t0L}) + R_{m0L}, \quad R_{tL} < R_{t0L} \\ R_{mL} &= D_{L2}\ (R_{tL} - R_{t0L}) + R_{m0L}, \quad R_{tL} \geq R_{t0L}. \end{aligned} \tag{2}$$

Similar expressions apply for the M- and S-cones. In this formulation, there are two gains (D_1 and D_2) associated with each cone class. There is also a parameter R_{t0} that defines the breakpoint between effective increments and decrements and a parameter R_{m0} that describes an additive effect of context. In Chichilnisky and Wandell's (1999) terms, parameters R_{t0} and R_{m0} describe the *neutral point* at the test and match locations, respectively. In our modeling, there is no necessary relation between the neutral point parameters and any physical properties of the stimulus.

The IDD model can be made consistent with the independence hypothesis by constraining its parameters. When image planes for one cone class are matched across a contextual image pair, we require that the parameters for that cone class match. When constrained in this way, we call the model the *independent increment–decrement diagonal model* (IDD-Ind model).

We fitted the models by using numerical search procedures (MATLAB Optimization Toolbox) to find the model parameters that minimize the difference between model predictions and individual matches. What metric to use for evaluating differences depends on the goal of the model fit. To evaluate overall quality of model fit, it is useful to use a perceptual error metric. To relate prediction error to conclusions about processing in separate cone pathways, it is more appropriate to minimize prediction error in *LMS* cone coordinates. We fitted the models in both ways.

Fig. 9 summarizes model performance for each subject and each contextual image pair used in Experiment 3. The fits used to produce this figure were obtained by minimizing mean CIELAB ΔE^* between model predictions and individual matches. The bars show the mean CIELAB ΔE^* prediction error obtained with the IDD

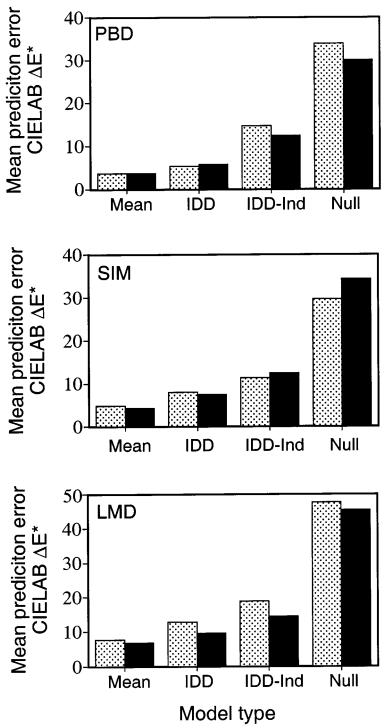

Fig. 9. The model fit errors are shown for each observer. See description in text. A table providing the numerical values for these fit errors, along with those obtained using the diagonal model of Eq. (1) and a version of the diagonal model consistent with the independence hypothesis, is available on the world wide web at URL http://color.psych.ucsb.edu/independence/.

model and the IDD-Ind model. Along with the model fits, the bars show comparison measures that indicate (a) the precision of the observers' matches (mean model) and (b) the prediction error obtained if one predicts the matches with the cone coordinates of the test stimuli (null model). To evaluate the precision, we took the mean match to each test stimulus (mean taken in CIELAB $L^*a^*b^*$ coordinates) and used this mean to predict the individual matches. The mean model provides a lower bound on the quality of fit that could be obtained with any model, while the null model indicates the size of the context effect to be explained. For each observer, fit quality is shown separately for the two contextual image pairs used in Experiment 3. The error plotted is averaged across the full color and isochromatic members of each contextual image pair. The parameters for the IDD model are available on the world wide web at URL http://color.psych.ucsb.edu/independence/.

For each observer, the IDD model captures most but not all of the effect of context. This is seen by comparing the prediction error for this model to the precision of the data (mean) and the size of the effect (null). For each observer, the IDD prediction errors exceed the precision of the data. To provide some intuition about the size of this model failure, we can proceed as follows. We take the size of the context effect to be the difference between the null and mean model error, and the amount of the effect explained to be the difference between the null and IDD model error. Given these measures, the IDD model captures, on average, 90% of the context effect.

We did not see any systematic pattern to the prediction residuals for the IDD model. Fig. 10 shows these residuals in *LMS* coordinates for one observer and contextual image pair, along with the residuals for the mean model. To make this figure, IDD model parameters were obtained by minimizing error in the *LMS* coordinate system, and the means for the Mean model were computed in the *LMS* coordinate system. Although the IDD residuals are larger than those of the mean model, there is no obvious pattern that is consistent across observers. Additional residuals for the S-cone component may be seen in Fig. 11.

We can compare the performance of the IDD model to that of the IDD-Ind model. The overall CIELAB ΔE^* prediction errors for the IDD-Ind model are greater than those for the IDD model for all observers and contextual image pairs. On average, the IDD-Ind model explains only 72% of the context effect. More importantly, the S-cone prediction residuals for the IDD-Ind model are highly systematic. Fig. 11 shows S-cone prediction residuals for the IDD and IDD-Ind models when the model parameters were determined by minimizing error in the *LMS* coordinate system. In each case, it is clear that the IDD-Ind model cannot account for the difference in the S-cone component of the matches obtained with the full color and isochromatic contextual images. The IDD model, on the other hand, handles both data sets fairly well.

The IDD model as fit above assumes that signal regulation occurs before signals from the different classes of cones are combined. Following Chichilnisky (1994) we fit variants of the IDD model to the data. In each variant, a randomly chosen linear transformation was applied to the *XYZ* coordinates of the test and match settings and the IDD model was fit in the transformed color space. The predictions of the model were then transformed back to the *XYZ* space. For each variant, the entries of the 3×3 transformation matrix were drawn from a uniform distribution over the range -1 to 1. A total of 1000 variants were fit, each with its own linear transformation. Fig. 12 shows a histogram of the resulting prediction errors, averaged over all three observers. The vertical dashed line in the figure shows the fit error obtained when the linear transformation is selected to transform the *XYZ* coordinates to *LMS* coordinates. None of the randomly chosen transformations provided a better fit than the cone space transformation, suggesting that to the extent that the data can be modeled by regulation at a single stage in visual processing, that stage is the initial *LMS* encoding.

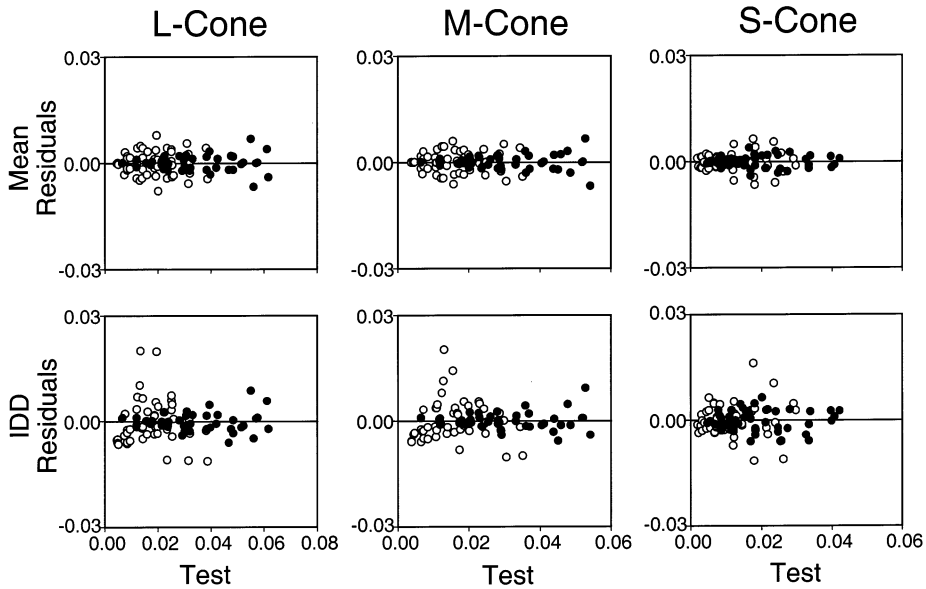

Fig. 10. Model fit residuals in *LMS* coordinates are shown for one observer (SIM) for the results obtained in Experiment 3(b). The filled circles represent the full color condition and the open circles represent the isochromatic condition. Residuals are shown for the mean model (top panels) and IDD model (bottom panels).

8. Discussion

8.1. Summary

The results from Experiments 1, 2 and 3 show that S-cone context effects are influenced by the L- and M-cone contextual image planes. The effect is large and consistent across observers and contextual image pairs. The results from Experiment 4 rule out the possibility that this effect is due to rod signals. Our modeling suggests that the failure of independence is largely due to an influence of L- and M-cone contextual image planes on signal processing in S-cone pathways.

8.2. Other failures of independence

We have shown that context effects are not computed independently for the S-cone plane. It is also of interest to know if S-cone signals have an effect on L- and M-cone processing. We created a new full color image from the isochromatic contextual image used in Experiments 1 and 3. This full color image had the same L- and M- cone stimulation as the isochromatic image, but paired with a different S-cone image plane (see Table 1 for the coordinates)[2]. For this full color image, two observers (PBD and SIM) set matches for tests with the same chromaticity as the gray tests used in Experiment 3. Fig. 13 compares these matches with those obtained previously in Experiment 3 for the same tests and the isochromatic contextual image. The results suggest that there is an influence of the S-cone contextual image on

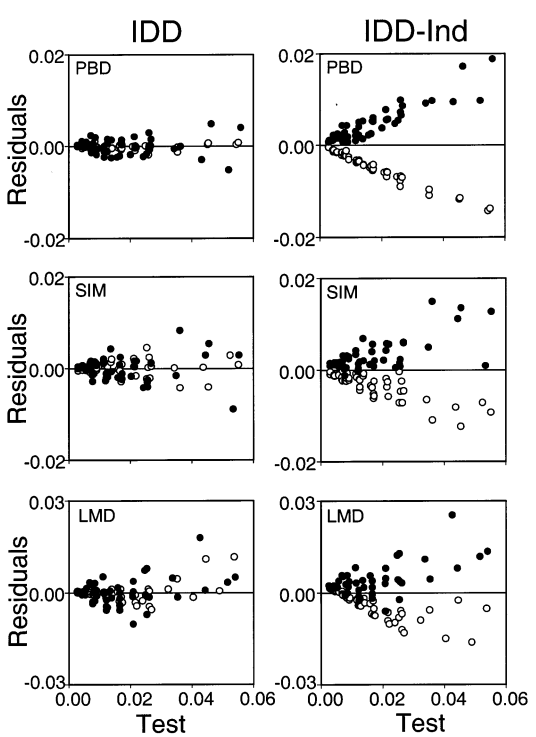

Fig. 11. S-cone model fit residuals are shown for all three observers in Experiment 3(a) for both the IDD and IDD-Ind models. The ● represent the full color condition and the ○ represent the isochromatic condition. Left panels: IDD model. Right panels: IDD-Ind model.

[2] For the comparison presented, the tests used for the full color image did not have the same luminances as the tests used for the isochromatic image. This violates a formal requirement of our test of independence (see Appendix). In practice this is not problematic as the data are regular as a function of luminance.

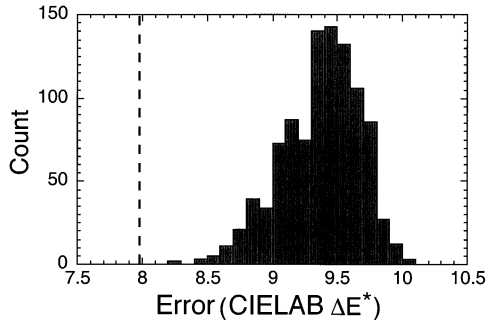

Fig. 12. This histogram shows the errors for variants of the IDD model fit in randomly chosen color spaces (see text for details). One thousand color spaces were chosen. The dashed line on the left of the histogram shows the fit error obtained when the IDD model is fit in cone space.

the processing of L- and M-cone signals, although this effect is small. We conclude provisionally that S-cones do influence the processing of L- and M-cone signals, although we not yet studied this effect in any detail. Similar measurements for a second contextual image pair, where again the L- and M-cone image planes were held constant while the S-cone image plane was varied, did not show an effect of the S-cone manipulation on the L- and M-cone components of the matches. Thus the influence of S-cones on L- and M-cone processing may be smaller or less robust than the influence of L- and M-cones on S-cone processing.

Are there interactions between L- and M-cones? Because of the large overlap between L- and M-cone spectral sensitivities, the range over which the L- and M-cone image planes can be independently manipulated is small, particularly when the S-cone image plane must also be held constant. In pilot experiments, we were not able to show any influence of the L-cone contextual image on the processing of M-cone signals or vice-versa. This failure may be due to the small size of the contextual manipulations we were able to create.

8.3. Relation to other work

8.3.1. Appearance measurements

Other authors have used appearance tasks to study the independence hypothesis. Cicerone et al. (1975) rejected independence for L- and M-cones based on studies of how spectral unique yellow varied as a function of an adapting field. Their conclusion rested on the assumption that unique yellow occurs when the output of a linear red/green mechanism is zero. More recent data falsify this linearity assumption (Werner & Wooten, 1979; Ikeda & Ayama, 1980; Burns, Elsner, Pokorny & Smith, 1984; Ayama, Kaiser & Nakatsue, 1985; Ejima & Takahashi, 1985; Chichilnisky & Wandell, 1999) and the status of their conclusion is thus uncertain.

Werner and Walraven (1982) and Chichilnisky and Wandell (1996) examined the independence hypothesis using achromatic settings made on different backgrounds. One conclusion drawn by Werner and Walraven was that signals from L- and M-cone pathways influence S-cone gains and that this effect is mediated at a non-opponent site. Our conclusions drawn from asymmetric matching experiments are in agreement with theirs (see below). Chichilnisky and Wandell's data indicate that the ratio of L- to M-cone contrast on the incremental achromatic locus depends on the S-cone

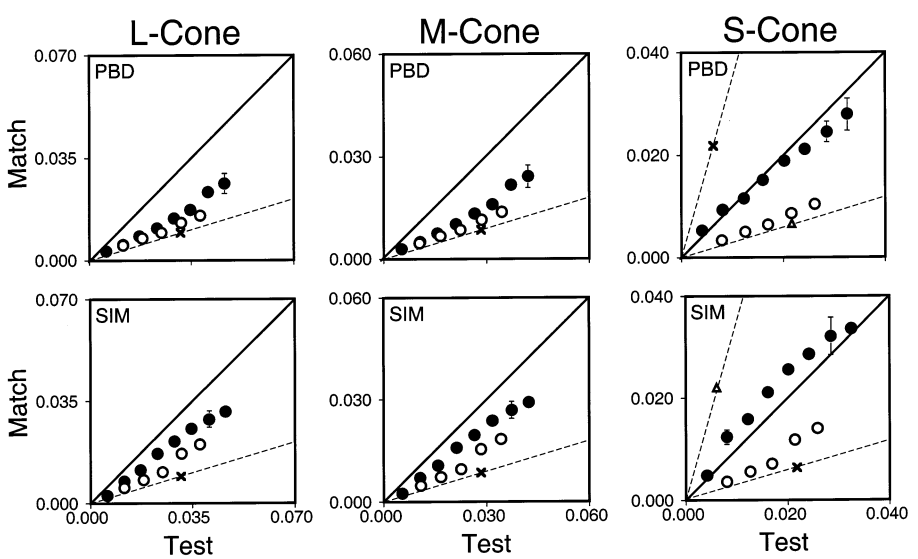

Fig. 13. Data from the supplementary experiment (see Section 8). Matches for the full color contextual image are shown by ● while isochromatic condition settings are shown by ○. The L- and M-cone planes were the same for both contextual images. In the panels showing the data for the S-cones, the two dashed lines indicate the properties of the S-cone image planes. The dashed lines containing the triangles are for the full color contextual image while the dashed lines containing the x's are for the isochromatic contextual image. Where visible, error bars indicate ± 1 SEM.

component of the background against which the achromatic settings were made. This result can be explained as a violation of independence where signals from S-cones affect processing of signals in L- and M-cone pathways. It is also possible to reconcile their data with independence if S-cone gains are regulated separately at different opponent sites, as allowed for in our general formulation (see Appendix A).

Creutzfeldt et al. (1990) on the other hand, were able to predict full color results from three (red, green and blue) isochromatic experiments. In this sense, their data support the independence hypothesis. Our results indicate that their prediction success does not generalize. One difference between studies is that our stimuli and procedures were designed specifically to reveal failures of independence. Chichilnisky and Wandell (1995) did not observe failures of independence in asymmetric matching experiments. Differences between their experiments and ours are that they used haploscopic matching and spatially uniform backgrounds whereas we used simultaneous matching within contextual images that had considerable spatial structure.

Poirson and Wandell (1996) (see also Baüml & Wandell, 1996) showed that varying the spatial structure of a test stimulus affects its color appearance. The most parsimonious account of the their data involves gain changes at opponent sites and is not consistent with independence. We do not think it likely that our experiments tap the same mechanisms studied by Poirson and Wandell. First, although our stimuli are spatially complex, the test and match squares were the same size as each other and did not vary in size across our experiments. Second, most of the variance in our data is captured by models (e.g. the IDD model) where signal regulation occurs at cone-specific sites. A more quantitative investigation of whether Poirson and Wandell's model can account for our data is beyond the scope of this paper.

Singer and D'Zmura (1995) studied chromatic contrast induction using a nulling task and showed that signals from one cone class could regulate signals from another. Modulating the L- and M-cone contrast in an annulus can produce an apparent modulation of contrast in a central disk. Under certain conditions, this apparent modulation can be nulled by modulating the S-cone contrast in the central disk. Their results are inconsistent with the independence hypothesis and their modeling incorporates gain control at opponent sites. It therefore seems likely that their experiments tap mechanisms different from the ones we studied here.

8.3.2. Threshold measurements

Threshold experiments have been used to argue that signal regulation occurs at opponent sites (Pugh & Mollon, 1979; Wandell & Pugh, 1980a,b; Stromeyer & Sternheim, 1981; Stromeyer, Cole & Kronauer, 1985; Krauskopf, Williams & Heeley, 1982; Mollon, 1982; Poirson & Wandell, 1996; Chen, Foley & Brainard, 2000a,b; for a review see Eskew, McLellan & Giulianini, 1999; see also Ahn & MacLeod, 1993). Note that adaptation at opponent sites per se does not necessarily falsify the independence hypothesis in its general form, as this form allows signals from each cone class to be regulated separately as they flow to multiple opponent sites (see the Appendix.) Thus measurements that show only that adaptation occurs at opponent sites are silent with respect to the independence hypothesis.

Influence of signals from one cone class on sensitivity to signals from other cone classes has also been examined. These results show a clear influence of signals from L- and M-cones on detection thresholds mediated by S-cones (Pugh & Mollon, 1979; Mollon, 1982; see Eskew et al., 1999), as well as an interaction between L- and M-cones (Wandell & Pugh, 1980a,b; Stromeyer & Sternheim, 1981; Stromeyer et al., 1985; see Eskew et al., 1999). These results are consistent with our general conclusions, although it is not clear whether the threshold effects are mediated at the same sites tapped by our measurements.

The fact that signals from L- and M-cones affect thresholds mediated by S-cone signals does not necessarily imply that violations of independence will be revealed by appearance experiments. First, it is generally accepted that thresholds are affected both by signal gains and by noise levels (Green & Swets, 1966). Threshold effects mediated by changes in noise levels would not necessarily be revealed in appearance experiments.

Second, even for fixed noise, thresholds can be affected by response compression. Suppose that the threshold effects reviewed above occur because L- and M-cone signals push an S-cone opponent site towards saturation. Although this can raise threshold for S-cone signals, the effect need not be revealed in an asymmetric matching experiment: for a match to occur the opponent site must be in the same (saturated) state for both test and match stimuli. A MATLAB program that demonstrates this latter point for a simplified visual system is available at URL http://color.psych.ucsb.edu/independence/.

Adaptation at opponent sites is inconsistent with the IDD model. It is possible that the variance in our data that is not accounted for by the IDD model is due to adaptation at opponent sites. If so, the fact that the IDD model failures are small suggests that adaptation at opponent sites plays at most a small role for our conditions. If this is the case, then our effects may have a different origin than the effects that have been studied in the detection literature.

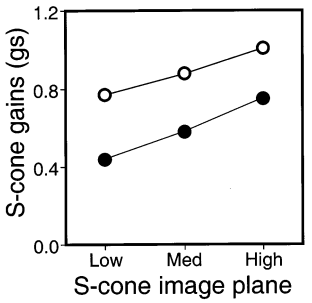

Fig. 14. The gains g_S computed from the data for each of the six contextual images are shown. The closed circles show gains for the low L- and M-cone level, while the open symbols show gains for the high L- and M-cone level.

8.3.3. Physiology

Dacey, Lee, Stafford, Pokorny and Smith (1996) have shown that the HII horizontal retinal cells sum signals from L-, M- and S-cones. If these cells mediate gain control of signals passed by S-cones, then they provide a candidate site for the interactions we observe.

8.4. Observer's task

Under some stimulus conditions, instructions play an important role in how observers match color stimuli (e.g. Arend & Reeves, 1986; Baüml, 1999). We believe, however, that there also exist stimulus conditions (e.g. flat matte surfaces seen under spatially diffuse illumination) for which observers cannot dissociate the color appearance of a surface from a judgment of its reflectance properties. For our stimuli, both authors found it difficult to make two separate matching judgments. A precise specification of when instructions can affect asymmetric matches and consensus as to how such effects should be modeled awaits further research. It remains an open question as to how well our models and conclusions generalize to stimulus conditions where instructions play a more salient role.

8.5. Implications for understanding color context effects

Our results show that asymmetric matches made for isochromatic stimuli cannot be used in a simple way to predict full color matches. In particular, the results undermine the central idea of the retinex theory (Land & McCann, 1971; McCann et al., 1976; Land, 1986) which is based on the assumption that contextual processing occurs independently for the three classes of cones.

If asymmetric matches made separately in isochromatic images are insufficient to predict full color matches, is there some other way to simplify the empirical problem of characterizing general color context effects? Suppose one has a parametric model that predicts matching stimuli from test stimuli for any particular context, so that the same model works for any contextual image pair as long as the model parameters are properly chosen for that pair. (For our conditions, the IDD model can serve in this role.) Then the problem of understanding full color context effects is solved if we determine the relation between context and the model parameters (Krantz, 1968; Brainard & Wandell, 1992; Baüml, 1995; Chichilnisky & Wandell, 1999). In a study of color constancy that employed asymmetric matching, for example, Brainard and Wandell (1992) used a diagonal model and parameterized contexts in terms of the illuminant spectral power distribution. They showed that across a set of contexts the diagonal model parameters could be determined as a linear function of the illuminant. Thus effects for all illuminants within a linear model could be predicted from measurements for a small set of basis functions.

A similar approach might be taken for our paradigm. Suppose we make measurements of how S-cone matching depends on the S-cone image plane for fixed L- and M-cone contextual image planes, and also on how S-cone matching depends on the L- and M-cone image planes for a fixed S-cone image plane. Then is it possible to predict what happens when we measure for various combinations? We have not yet studied this issue in detail, and the following development is intended to be illustrative rather than definitive.

We created six contextual image pairs. The six contextual image pairs were obtained by crossing three S-cone image planes with two sets of L- and M-cone image planes. We refer to the three levels of S-cone image plane as high, medium and low in reference to the slope of the line determined by the S-cone coordinate of the contextual image at the test and match locations (e.g. the slope of the thin dashed lines shown in the data figures). Similarly, we refer to the two levels of L- and M-cone image plane as high and low. For each image pair, observer PBD made asymmetric matches using the test stimuli from Experiment 1. For each contextual image pair, we predicted the S-cone matches using a simple one parameter model $R_{mS} = g_S R_{tS}$. Although this model is not correct in general (see above), it is adequate for illustration. We estimated g_S from the slopes of the test versus match plots for the S-cone component of the matches. We then examined how g_S depended on the contextual images. Fig. 14 plots g_S as a function of the three levels of S cone context and two levels of L- and M-cone context. The plot makes it clear that although both S-cone and L- and M-cone image planes affect g_S, they combine their influence in a simple additive manner.

If the absence of an interaction shown by Fig. 14 generalizes, then it is possible to build a complete theory of how the S-cone components of full color asymmetric matches depend on context from independent measurements of how the L-, M- and S-cone contextual image planes affect the S-cone matches.

Also of note in Fig. 14 is that when we express the context in terms of the average image gradient for each contextual image plane, the influence of L- and M-cones on the S-cone gains has the same sign as the influence of S-cones. This provides further support for the notion that our effects are not mediated at opponent sites.

Acknowledgements

We thank E. Adelson, E. Chichilnisky, D. Dacey, J. Kraft, J. Loomis, D. MacLeod, J. Palmer, D. Teller, B. Wandell, Q. Zaidi for useful discussions. J. Tietz provided technical support. S. Maloney helped conduct experimental sessions. D. Teller and J. Palmer at the University of Washington allowed us to use their apparatus to conduct Experiment 4, and H. L. Chien volunteered to observe in this experiment. Supported by NEI EY10016.

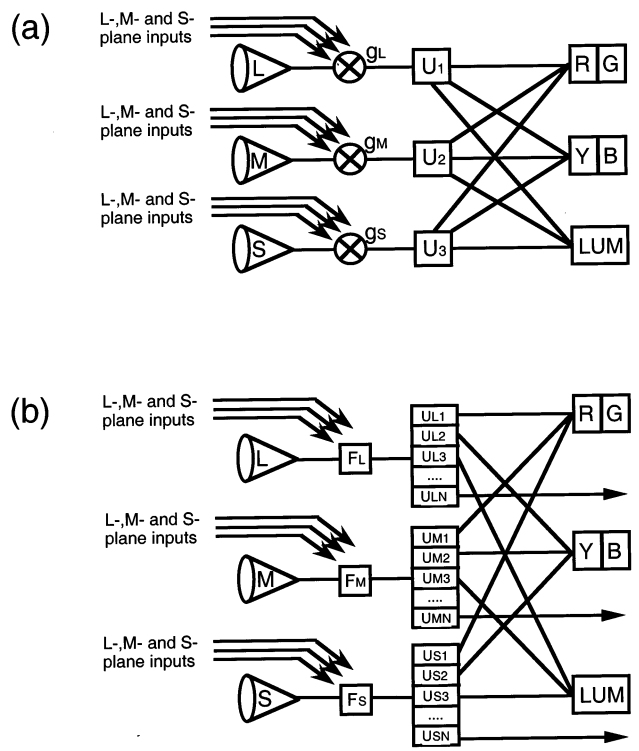

Fig. 15. In the top panel, the weak form of the von Kries Hypothesis is illustrated. The cone signals are subjected to a gain which is controlled by inputs that may arise in all three cone classes. The gain-adjusted signals (U_1, U_2, U_3) are then sent on for further processing. In the figure, this processing consists of an opponent stage as incorporated in most modern models of color vision. Gain control at the opponent sites is not allowed in this model. The bottom panel illustrates a generalized version of the von Kries hypothesis. The cone signals are subjected to an unspecified transform (F). This function can be different for each cone class and can have multiple independent outputs (e.g. (U_{L1} … U_{LN}), so that the signals from a single cone class may be regulated differently as they are passed to separate post-receptoral sites.

Appendix A. Formal development

In this appendix we formalize our analysis of the relation between asymmetric matches and the independence hypothesis. We consider contextual image pairs that are identical for one cone type and differ for the other two. Denote the two contextual images as a and b. Following the notation developed in the text, let the cone coordinates of matches established in contextual image a be denoted by $\{\mathbf{R}_t^a\ \mathbf{R}_m^a\}$ and those of matches established in contextual image b be denoted by $\{\mathbf{R}_t^b\ \mathbf{R}_m^b\}$. Let the vectors $\mathbf{I}_L^a$, $\mathbf{I}_M^a$, $\mathbf{I}_S^a$ parameterize the L-, M- and S-cone image planes of image a and similarly for $\mathbf{I}_L^b$, $\mathbf{I}_M^b$, $\mathbf{I}_S^b$. We will consider the case where $\mathbf{I}_S^b = \mathbf{I}_S^a$ and show that the independence hypothesis requires that if the S-cone component of $\mathbf{R}_t^a$ equals the S-cone component of $\mathbf{R}_t^b$ then the S-cone components of the corresponding $\mathbf{R}_m^a$ and $\mathbf{R}_m^b$ must be equal. The same argument can be applied to derive analogous predictions for the L- and M-cones when $\mathbf{I}_L^b = \mathbf{I}_L^a$ and $\mathbf{I}_M^b = \mathbf{I}_M^a$.

We follow Stiles (Stiles, 1967; see also Krantz, 1968; Brainard & Wandell, 1992) and assume that an asymmetric match occurs when signals corresponding to the test and match region are identical at some critical site along the visual pathways. Consider the relation between the cone coordinates $\mathbf{R}$ of a stimulus and the corresponding responses $\mathbf{U}$ at the critical site. This relation will depend both on the location of the stimulus and on the contextual image, and we can write it in the rather general form

$$\mathbf{U} = \mathbf{F}(\mathbf{R};\ \mathbf{x},\ \mathbf{I}_L,\ \mathbf{I}_M,\ \mathbf{I}_S) \tag{A1}$$

where $\mathbf{x}$ denotes the image location, and $\mathbf{I}_L$, $\mathbf{I}_M$, and $\mathbf{I}_S$ are parametric descriptions of the L-, M-, and S-cone planes of the contextual image. We do not know the specific form of $\mathbf{F}(\,)$, but we will assume that when location ($\mathbf{x}$) and contextual image ($\mathbf{I}_L$, $\mathbf{I}_M$, $\mathbf{I}_S$) are held fixed, each $\mathbf{R}$ corresponds to a unique $\mathbf{U}$. We do not, however, constrain the dimension of $\mathbf{U}$ to be the same as that of $\mathbf{R}$, so that some states of the critical site may not be attainable for a particular location and contextual image.

The idea of the independence hypothesis is that all effects of context occur independently in the three cone classes. To explain this restriction, first consider the specific hypotheses articulated by von Kries. The weak form of the von Kries hypothesis specifies that the only effect of context is gain changes in cone-class specific pathways (see Fig. 15). This idea can be expressed by placing restrictions on the form of Eq. (A1). Write $\mathbf{U} = [U_1\ U_2\ U_3]^T$ and require that for any context each entry of $\mathbf{U}$ is obtained from the quantal absorptions in a single cone class through a simple gain control operation:

$$U_1 = g_L(\mathbf{x},\mathbf{I}_L,\ \mathbf{I}_M,\ \mathbf{I}_S)R_L, \quad U_2 = g_M(\mathbf{x},\ \mathbf{I}_L,\ \mathbf{I}_M,\ \mathbf{I}_S)R_M,$$
$$U_3 = g_S(\mathbf{x},\ \mathbf{I}_L,\ \mathbf{I}_M,\ \mathbf{I}_S)R_S. \tag{A2}$$

Here R_L, R_M and R_S are the L-, M-, and S-cone entries of R and g_L, g_M, and g_S are scalar gains that depend both on location and the contextual image. Although the weak form of the von Kries hypothesis specifies that context acts through changes in cone-class specific pathways, this form does not necessarily satisfy the independence hypothesis: across contexts the gains in a single cone-class may be affected by signals from all three cone classes. For consistency with the independence hypothesis, we must require further that the gain for each cone class depends only on location and the contextual image plane for that cone class:

$$U_1 = g_L(\mathbf{x},\mathbf{I}_L)R_L, \quad U_2 = g_M(\mathbf{x},\mathbf{I}_M)R_M,$$
$$U_3 = g_S(\mathbf{x},\mathbf{I}_S)R_S. \tag{A3}$$

This is the strong form of the von Kries hypothesis.

Note that neither form of the von Kries hypothesis contradicts color opponency. As illustrated in the figure, these hypotheses only require that opponency occur after the critical site (**U**) before which all context effects occur.

It is possible to develop tests of the independence hypothesis by assuming that the weak form of the von Kries hypothesis holds and then evaluating whether the dependence of the gains on context satisfy Eq. (A3) above. We wish, however, to develop tests that require as few assumptions as possible. To do so, we use the independence hypothesis to place restrictions on Eq. (A1) without assuming the specific form of the von Kries hypothesis in Eq. (A2).

Let $U_1 \ldots U_N$ be the entries of **U**. Rather than assuming that Eq. (A2) holds, we assume that we can define a fixed (across all locations and contexts) partition of **U** into three sets of entries, $\mathbf{U}_L$, $\mathbf{U}_M$, and $\mathbf{U}_S$. This partition should have the property that when **x**, $\mathbf{I}_L$, $\mathbf{I}_M$, and $\mathbf{I}_S$ are held fixed the entries of $\mathbf{U}_L$ depend only on R_L and not on R_M and R_S. Similarly, the partition should have the property that $\mathbf{U}_M$ depends only on R_M and $\mathbf{U}_S$ depends only R_S when **x**, $\mathbf{I}_L$, $\mathbf{I}_M$, and $\mathbf{I}_S$ are held fixed. This requirement is illustrated in the bottom panel of Fig. 15. It produces an analog of the weak form of the von Kries hypothesis. When the context is held fixed, each entry of **U** depends on signals originating in only one cone class, but we no longer require that the entry be obtained by a simple gain control operation, nor do we require that signals from each cone class propagate along a single pathway. Thus this form is consistent both with subtractive adaptation (Jameson & Hurvich, 1964; Walraven, 1976; Shevell, 1978), with differential gains for incremental and decremental signals (Mausfeld & Niederee, 1993; Chichilnisky & Wandell, 1996; Mausfeld, 1998), and with the possibility that context affects signals originating in a single cone class differently at different post-receptoral sites (Ahn & MacLeod, 1993).

Like the weak form of the von Kries hypothesis, the hypothesis illustrated in the bottom panel of Fig. 15 does not necessarily satisfy the independence hypothesis. We can further require, however, that the relation between each entry of **U** and the cone signal that drives it is independent of the contextual image planes for the other two cone classes:

$$\mathbf{U}_L = \mathbf{F}_L(R_L;\ \mathbf{x},\ \mathbf{I}_L) \tag{A4}$$

with similar restrictions holding for $\mathbf{U}_M$ and $\mathbf{U}_S$. In this form of the hypothesis, the functions $\mathbf{F}_L(\,)$, $\mathbf{F}_M(\,)$, and $\mathbf{F}_S(\,)$ characterize the effect of context on color appearance. We assume that $\mathbf{F}_L(\,)$ has the property that each R_L corresponds to a unique $\mathbf{U}_L$ when location and context are held fixed, and similarly for $\mathbf{F}_M(\,)$ and $\mathbf{F}_S(\,)$.

We must now show that the restrictions of Eq. (A4) above make a testable empirical prediction. Suppose that Eq. (A4) holds and consider asymmetric matches in context images a and b that satisfy the property $\mathbf{I}_S^a = \mathbf{I}_S^b = \mathbf{I}_S$. This property ensures that the two context images have the same spatial extent and we assume that the test stimulus is placed at location **t** in both images and that the match stimulus is placed at location **m** in both images. Suppose that the test stimuli used in the two images have the same S-cone component, so that $R_{tS}^a = R_{tS}^b = R_{tS}$. From Eq. (A4) we have for an asymmetric match in context image a

$$\mathbf{U}_S^a = \mathbf{F}_S(R_{tS};\ \mathbf{t},\ \mathbf{I}_S) = \mathbf{F}_S(R_{mS}^a;\ \mathbf{m},\ \mathbf{I}_S) \tag{A5}$$

and for context image b that

$$\mathbf{U}_S^b = \mathbf{F}_S(R_{tS};\ \mathbf{t},\ \mathbf{I}_S) = \mathbf{F}_S(R_{mS}^b;\ \mathbf{m},\ \mathbf{I}_S). \tag{A6}$$

By inspection of Eqs. (A5) and (A6) for the test location we obtain that $\mathbf{U}_S^a = \mathbf{U}_S^b$ which gives us

$$\mathbf{F}_S(R_{mS}^a;\ \mathbf{m},\ \mathbf{I}_S) = \mathbf{F}_S(R_{mS}^b;\ \mathbf{m},\ \mathbf{I}_S) \tag{A7}$$

Since each R_S corresponds to a unique $\mathbf{U}_S$, Eq. (A7) allows us to conclude that $R_{mS}^a = R_{mS}^b$. This is the implication of the independence hypothesis that we test empirically.

Finally, we wish to show that when the restrictions of Eq. (A4) hold, it is possible to simplify the problem of predicting full color asymmetric matches. What we have shown above is that in this case, S-cone components of asymmetric matches made in any contextual image are independent of R_{tL}, R_{tM}, $\mathbf{I}_L$, and $\mathbf{I}_M$. Thus we can use asymmetric matching with these variables held fixed to study the relation between $\mathbf{I}_S$ and the S-cone component of the asymmetric matches. Once we understand this relation, the restrictions of Eq. (A4) ensure that the same relation will hold for any choice of R_{tL}, R_{tM}, $\mathbf{I}_L$, and $\mathbf{I}_M$. A similar argument applies to the relation between $\mathbf{I}_L$ and the L-cone component of asymmetric matches and to the relation between $\mathbf{I}_M$ and the

M-cone component of asymmetric matches. Thus the restrictions allow us to characterize full color matches through independent measurements of the effect of each contextual image plane on the corresponding match component. If we use an isochromatic contextual image and test stimuli with the same chromaticity as the contextual image, we can study the effect of each contextual image plane on the corresponding image component in a single experiment.

References

Adelson, E. H. (1993 (December 24)). Perceptual organization and the judgment of brightness. *Science*, *262*, 2042–2044.

Aguilar, M., & Stiles, W. S. (1954). Saturation of the rod mechanism of the retina at high levels of illumination. *Optica Acta*, *1*, 59–65.

Ahn, S. J., & MacLeod, D. I. A. (1993). Link-specific adaptation in the luminance and chromatic channels. *Vision Research*, *33*, 2271–2286.

Arend, L. E., & Reeves, A. (1986). Simultaneous color constancy. *Journal of the Optical Society of America*, *A*, *3*, 1743–1751.

Ayama, M., Kaiser, P. K., & Nakatsue, T. (1985). Additivity of red chromatic valence. *Vision Research*, *25*, 1885–1891.

Baüml, K. H. (1995). Illuminant changes under different surface collections: examining some principles of color appearance. *Journal of the Optical Society of America*, *A*, *12*(2), 261–271.

Baüml, K. H. (1999). Simultaneous color constancy: how surface color perception varies with the illuminant. *Vision Research*, *39*, 1531–1550.

Baüml, K. H., & Wandell, B. A. (1996). Color appearance of mixture gratings. *Vision Research*, *36*(18), 2849–2864.

Brainard, D. H. (1989). Calibration of a computer controlled color monitor. *Color Research and Application*, *14*(1), 23–34.

Brainard, D. H., & Wandell, B. A. (1986). Analysis of the retinex theory of color vision. *Journal of the Optical Society of America*, *A*, *3*, 1651–1661.

Brainard, D. H., & Wandell, B. A. (1992). Asymmetric color-matching: how color appearance depends on the illuminant. *Journal of the Optical Society of America*, *A*, *9*(9), 1433–1448.

Burns, S., Elsner, A. E., Pokorny, J., & Smith, V. C. (1984). The Abney effect: chromaticity coordinates of unique and other constant hues. *Vision Research*, *24*, 479–489.

Chen, C. C., Foley, J. M., & Brainard, D. H. (2000a). Detection of chromoluminance patterns on chromoluminance pedestals I: threshold measurements. *Vision Research 40*, 773–788.

Chen, C. C., Foley, J. M., & Brainard, D. H. (2000b). Detection of chromoluminance patterns on chromoluminance pedestals II: model. *Vision Research 40*, 789–803.

Chichilnisky, E. J. (1994). *Perceptual measurements of neural computations in color appearance*. Unpublished Ph.D., Stanford University, Palo Alto, CA.

Chichilnisky, E. J., & Wandell, B. A. (1995). Photoreceptor sensitivity changes explain color appearance shifts induced by large uniform backgrounds in dichoptic matching. *Vision Research*, *35*(2), 239–254.

Chichilnisky, E. J., & Wandell, B. A. (1996). Seeing gray through the on and off pathways. *Visual Neuroscience*, *13*(3), 591–596.

Chichilnisky, E. J., & Wandell, B. A. (1999). Trichromatic opponent color classification. *Vision Research*, *39*, 3444–3458.

Cicerone, C. M., Krantz, D. H., & Larimer, J. (1975). Opponent-process additivity – III. Effect of moderate chromatic adaptation. *Vision Research*, *15*, 1125–1135.

Creutzfeldt, O., Lange-Malecki, B., & Dreyer, E. (1990). Chromatic induction and brightness contrast: a relatistic color model. *Journal of the Optical Society of America*, *A*, *7*, 1644–1653.

Dacey, D. M., Lee, B. B., Stafford, D. K., Pokorny, J., & Smith, V. (1996). Horizontal cells of the primate retina: cone specificity without spectral opponency. *Science*, *271*, 656–659.

Ejima, Y., & Takahashi, S. (1985). Interaction between short- and longer-wavelength cones in hue cancellation codes: nonlinearities of hue and cancellation as a function of stimulus intensity. *Vision Research*, *25*, 1911–1922.

Eskew, R. T., McLellan, J. S., & Giulianini, F. (1999). Chromatic detection and discrimination. In K. Gegenfurtner, & L. T. Sharpe, *Color vision: from molecular genetics to perception* (pp. 345–368). Cambridge: Cambridge University Press.

Evans, R. M. (1948). *An introduction to color*. New York: Wiley.

Gilchrist, A. L. (1980). When does perceived lightness depend on perceived spatial arrangement? *Perception and Psychophysics*, *28*, 527–538.

Gilchrist, A. L. (1988). Lightness contrast and failures of constancy: a common explanation. *Perception and Psychophysics*, *43*, 415–424.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.

Ikeda, M., & Ayama, M. (1980). Additivity of opponent chromatic valence. *Vision Research*, *20*, 995–999.

Jameson, D., & Hurvich, L. M. (1964). Theory of brightness and color contrast in human vision. *Vision Research*, *4*, 135–154.

Kraft, J. M., & Brainard, D. H. (1999). Mechanisms of color constancy under nearly natural viewing. *Proceedings of the National Academy of Sciences USA*, *96*(1), 307–312.

Krantz, D. (1968). A theory of context effects based on cross-context matching. *Journal of Mathematical Psychology*, *5*, 1–48.

Krauskopf, J., Williams, D. R., & Heeley, D. W. (1982). Cardinal directions of color space. *Vision Research*, *22*, 1123–1131.

Land, E. H. (1986). Recent advances in retinex theory. *Vision Research*, *26*, 7–21.

Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, *61*, 1–11.

Mausfeld, R. (1998). Color perception: from Grassman codes to a dual code for object and illumination colors. In W. G. K. Backhaus, R. Kliegl, & J. S. Werner, *Color vision — perspectives from different disciplines* (pp. 219–250). Berlin: de Gruyter.

Mausfeld, R., & Niederee, R. (1993). An inquiry into relational concepts of colour, based on incremental principles of colour coding for minimal relational stimuli. *Perception*, *22*(4), 427–462.

McCann, J. J., McKee, S. P., & Taylor, T. H. (1976). Quantitative studies in retinex theory: a comparison between theoretical predictions and observer responses to the 'Color Mondrian' experiments. *Vision Research*, *16*, 445–458.

Mollon, J. D. (1982). A taxonomy of tritanopias. *Documenta Ophthalmologica Proceedings Series*, *33*, 87–101.

Poirson, A. B., & Wandell, B. A. (1996). Pattern-color separable pathways predict sensitivity to simple colored patterns. *Vision Research*, *36*(4), 515–526.

Pokorny, J., & Smith, V. C. (1997). How much light enters the retina? In C. R. Cavonius, *Colour vision deficiencies XIII* (pp. 491–511). Dordrecht: Kluwer.

Pugh, E. N. J., & Mollon, J. D. (1979). A theory of the πl and π3 color mechanisms of Stiles. *Vision Research*, *19*, 293–312.

Rasengane, T. A., Palmer, J., & Teller, D. Y. (2000). Infants show Weber's law at high retinal illuminances. *Vision Research* (under review).

Richards, W., & Parks, E. A. (1971). Model for color conversion. *Journal of the Optical Society of America*, *61*(7), 971–976.

Schirillo, J. A. (1999). Surround articulation. I. Brightness judgments. *Journal of the Optical Society of America*, *16*(4), 793–803.

Schirillo, J. A. (1999). Surround articulation. II. Lightness judgments. *Journal of the Optical Society of America*, *16*(4), 804–811.

Shapiro, A. G., Pokorny, J., & Smith, V. C. (1996). Cone-rod receptor spaces with illustrations that use CRT phosphor and light-emitting-diode spectra. *Journal of the Optical Society of America*, *A*, *13*(12), 2319–2328.

Shevell, S. K. (1978). The dual role of chromatic backgrounds in color perception. *Vision Research*, *18*, 1649–1661.

Singer, B., & D'Zmura, M. (1995). Contrast gain control — a bilinear model for chromatic selectivity. *Journal of the Optical Society of America*, *12*(4), 667–685.

Stiles, W. S. (1949). Increment thresholds and the mechanisms of colour vision. *Documenta Ophthalmologica*, *3*, 138–163.

Stiles, W. S. (1959). Color vision: the approach through increment threshold sensitivity. *Proceedings National Academy of Sciences* (*USA*), *45*, 100–114.

Stiles, W. S. (1967). Mechanism concepts in colour theory. *Journal of the Colour Group*, *11*, 106–123.

Stromeyer, C. F. III, Cole, G. R., & Kronauer, R. E. (1985). Second-site adaptation in the red–green chromatic pathways. *Vision Research*, *25*, 219–237.

Stromeyer, C. F. III, & Sternheim, C. E. (1981). Visibility of red and green spatial patterns upon spectrally mixed adapting fields. *Vision Research*, *21*(3), 397–408.

Von Kries, J. (1902). Chromatic adaptation,: Originally published in Festschrift der Albrecht-Ludwigs-Universitat. In D. L. MacAdam, *Sources of color vision* (pp. 145–148) Cambridge: MIT Press.

Vos, J. J. (1972). Literature review of human macular absorption in the visible and its consequences for the cone receptor primaries TNO Report: The Netherlands: Institute for Perception, Soesterberg.

Walraven, J. (1976). Discounting the background: the missing link in the explanation of chromatic induction. *Vision Research*, *16*(3), 289–295.

Wandell, B. A., & Pugh, E. N. (1980). Detection of long-duration, long-wavelength incremental flashes by a chromatically coded pathway. *Vision Research*, *20*, 625–636.

Wandell, B. A., & Pugh, E. N. (1980). A field-additive pathway detects brief-duration, long-wavelength incremental flashes. *Vision Research*, *20*, 613–624.

Webster, M.A., MacLeod, D.I.A. (1988). Factors underlying individual differences in the color matches of normal observers. Journal of the Optical Society of America, A, 5, 1722–1735.

Werner, J. S., & Walraven, J. (1982). Effect of chromatic adaptation on the achromatic locus: the role of contrast, luminance and background color. *Vision Research*, *22*(8), 929–944.

Werner, J. S., & Wooten, B. R. (1979). Opponent chromatic mechanisms: relation to photopigments and hue naming. *Journal of the Optical Society of America*, *69*, 422–434.

Wesner, M. F., & Shevell, S. K. (1992). Color perception within a chromatic context-changes in red green equilbria caused by non-contiguous light. *Vision Research*, *32*, 1623–1634.