BACKGROUNDS AND ILLUMINANTS:
THE YIN AND YANG OF COLOR CONSTANCY

Richard O. Brown
Center for Human Information Processing
University of California at San Diego
La Jolla, CA 92093-0109

Current address:
The Exploratorium
3601 Lyon Street
San Francisco, CA 94123-1099

e-mail robrown@exploratorium.edu
tel. (415) 353-0483
FAX (415) 561-0370

0. ABSTRACT

Colored objects normally maintain stable color appearances across a wide range of viewing conditions, even though these changing conditions introduce large variations into the corresponding visual signals. This is the phenomenon known as “color constancy”. Two major types of variation arise from changes in the light illuminating the objects, and from changes in the backgrounds against which the objects are seen. Most studies of color constancy have focused on the challenge posed by variations in illumination, and many of the best-known models of color constancy were developed primarily to solve this problem. Color constancy with changing backgrounds has received relatively little attention, but is increasingly being recognized as an equally important problem. There is an interesting complementarity between these two aspects of color constancy, in that simple mechanisms that would tend to maintain excellent color constancy for one of these types of variation, tend to fail quite badly for the other. In particular, many color constancy models rely on the space-averaged light from scenes to estimate illumination, but such models generally misinterpret colored backgrounds as colored illuminants.

There may not be a general solution that achieves color constancy with both changing illuminants and changing backgrounds. It is argued that instead of seeking a general, computational solution, color constancy should be studied in terms of the actual properties of the visual system and of the ecological color signals it evolved to see. An analysis of measured natural reflectances and illuminants has led to several hypotheses about the mechanisms involved in biological color constancy: (1) The popular “Grey World” models of color constancy, which interpret changes in the space-averaged light as

changes in illumination, will generally fail, as variations in the chromaticity of space-averaged light are at least as likely to arise from changing backgrounds as changing illuminants; (2) Linear models based on 3-dimensional representations of illuminants and surfaces are inadequate to capture important variations in ecological color signals; (3) The relative variations due to changing backgrounds and illuminants are highly asymmetric in the luminance and color-opponent channels, with variations in the luminance channel corresponding largely to changes of illumination intensity, while variations in the chromatic channels primarily represent varying reflectances. (4) Asymmetries in the known physiological and psychophysical properties of the luminance and color-opponent channels may represent important adaptive tuning to these measured asymmetries in the corresponding ecological signals, suggesting that color constancy is not simply “lightness constancy x 3”.

1. COLOR VISION

When we open our eyes, our world of stably colored objects seems to appear instantly and automatically. The reliability and seeming effortlessness of color perception belies the difficulties of this achievement. It remains a major goal for perceptual science to understand how biological systems process visual signals from the external world to generate perceived colors. Our naïve, direct experience of color is that color is simply a property of colored objects, just as size and shape are, and reflecting this view one approach to color science has sought to identify color with the physical properties of objects. But since Newton's famous spectral analysis of light, color research has predominantly focused on the relationship between colors and the spectral characteristics of light, with objects viewed more as the modifiers of colored light than as themselves the source of color (see Finger 1994, Mausfeld 1998.) A competing line has focused on the important role of contrast in color vision, and the identification of colors with ratios of light signals from different parts of a visual scene (Zeki 1993, Whittle, this book?). A more biologically oriented approach, and the one I favor, is to identify color with the color sensations of seeing organisms, and to try to understand color in terms of the relationship between the physical signals captured by eyes and the neural processes which generate colors.

Studies and discussions of color often suffer from the many ambiguities and confusion in our language for color. . J. J. Gibson (cited in Eco 1985) observed that "The meaning of the term colour is one of the worst muddles in the history of science!" The term color, and other color words, have been variously used to describe light, objects, contrasts, and perceptions. Under ordinary circumstances, these are all so tightly linked that the confusion is minimal. But the scientific study of color typically involves the deliberate dissection of these variables into competing influences, and so requires more precise terminology to distinguish all these aspects of color. Here the term color will be used to refer to the sensation of color in an organism with color vision, while objects and lights will be described in terms of their spectral reflectances and spectral power distributions, respectively. Color, so defined as a sensation, need not correspond to any physically measurable property of objects or lights, and in this sense there is no defined measure of the "veridicality" of color perception. Another language ambiguity arises from the frequent exclusion of blacks, whites and greys from the domain of color; here color will refer to the full range of colors including these achromatic colors.

OBJECT COLORS: It is generally assumed that the primary purpose of color vision is to support detection and identification of objects in the environment. Helmholtz (1866) noted that "Colors have their greatest significance for us in so far as they are properties of bodies and can be used as marks of identification of bodies," and on this point Hering (1920) agreed, "In vision, we are not concerned with perceiving light rays as such, but with perceiving the external objects mediated by these radiations; the eye must inform us, not about the momentary intensity or quality of the light reflected from external objects, but about these objects themselves." This emphasis on colored objects as the business of color vision makes color constancy paramount, so the perceived colors can be reliably associated with appropriate colored objects. But it should be kept in mind that seeing stably colored objects need not be the only purpose of color vision. For instance, perceiving the changing colors of the sky may be important for keeping time or predicting weather changes, and some organisms appear to use color sensitivity primarily to detect such environmental changes.

The perception of object colors is mediated by the light reflected from objects to the eye. Figure 1 is a very simplified schematic of the basic problem of color vision. The banana represents a typival colored object of interest to the visual system. For the purposes of color vision, its surface may be approximated by a spectral reflectance function showing the fraction of incident light at each wavelength reflected from the surface. In this case, the banana is ripe, and absorbs most of the short wavelengths of light while reflecting much of the middle and long wavelength light. The banana is illuminated by daylight, which may be represented by a spectral power distribution, showing its relative power at each wavelength. The light reflected from the banana, which provides the proximal signal available for color vision, is determined at each wavelength by the banana's spectral reflectance multiplied by the illuminant's soectral power distribution. Note that the proximal light signals for color vision depend equally on the reflectances of objects, and the incident illumination on those objects. In human eyes, these proximal light signals would normally be sensed by three types of cones, reducing a high-dimensional proximal light signal (that is, one which may vary independently at each wavelength) to a 3-dimensional neural signal. The disarmingly simple yet elusive goal of color science is to understand how this set of 3 cone responses is processed to reliably generate the perceived color of the banana.

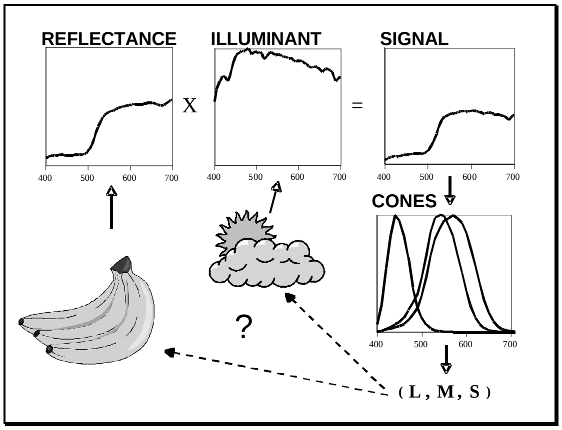

*********************** Figure 1 about here ************************

<u>COLOR CONTRAST:</u> When isolated uniform spots of light are viewed against completely dark backgrounds, their perceived colors (sometimes called “aperture colors”) can be reliably predicted from the physical composition of light in the patch. But the same uniform patch of light may generate a very different color appearance as soon as it is juxtaposed with, surrounded by, or preceded by, other spots of light. Hering observed, “In general, one and the same ray can be seen, according to the circumstances, in all possible color hues,” and Delacroix (see Evans 1964) boasted that he could paint the skin of Venus from the dirtiest mud, provided he could surround it with appropriate contrast colors. Such color contrast phenomena provide compelling evidence that perceived colors are not determined locally, by just the light signals from each point in a scene, but are relativistic, and involve comparisons of light signals across space and time. Color contrast effects are often treated as illusions or defects of visual processing. Kaiser & Boynton (1996), for example, refer to the aperture color of light as its “objective color”, while the colors perceived in an identical spot of light in another context are called “subjective colors”. On the other hand, many color researchers consider contrast the essential mechanism for achieving color constancy. Whether color contrast works to

favor or hinder color constancy depends on how it relates to changes in illumination and changes in backgrounds.

2. COLOR CONSTANCY

Color constancy is the tendency for objects to maintain stable color appearances, despite considerable variations in the physical and neural signals mediating color vision. Color constancy is just one of many perceptual constancies that allow us to recognize and maintain stable perceptual representations of the external world, despite changes in the received signals informing us about the world. Because the fixed, intrinsic properties of objects are likely to have greater behavioral significance than the fluctuating signals available for their perception at any moment, the constancies are considered essential perceptual achievements. A general definition of perceptual constancy, after Hochberg 1988, is “the constancy of perception of the fixed properties of distal objects, despite variations in the proximal stimuli from the objects”. For example, a friend’s face may be uniquely recognized across a wide range of distances, positions, lighting, motions, and facial expressions, which all generate very different retinal images. Note that just as this “face constancy” need not preclude us from simultaneously recognizing faces and perceiving all these variations, color constancy need not imply that we are blind to the variations, such as changing backgrounds and illumination.

Applying Hochberg’s above formulation to color constancy generates the following general definition of object color constancy: “the constancy of the perceived colors of objects, despite variations in the proximal light signals.” This differs importantly from the most common formulation, which restricts color constancy to ‘the constancy of the perceived colors of objects despite variations in the illumination'. The general definition includes variations in the illumination, as well as many other types of variations that may pose challenges for color constancy. These include changes of the backgrounds against which objects are seen, changes in the atmospheric conditions through which light signals travel, and other changes such as changes in the viewing geometry, and changes in the neural sensitivity to the signals. Each of these challenges will be considered in turn below, with examples of the failures of color constancy that arise from each.

Variations of Illumination. The best known challenge to object color constancy comes from changes in the light illuminating objects. The difficulties caused by changing illumination can be appreciated from the schematic of color signals shown in Figure 1. Because the light reflected from the banana depends on the reflectance multiplied by the illuminant, any change in the illuminant will cause a proportionate change in the reflected light. For example, the light reflected from a banana may become greenish when the banana is an unripe green under a white illuminant, or when the banana is ripe yellow but its illuminant is greenish, perhaps from being filtered through green leaves. The reflected light reaching the eye from the banana may be identical in these two cases, but determining the true state of the banana may be important to a hungry animal. The changes of illumination to be considered should include not only temporal changes in the average illumination across an entire scene, which is what most color constancy references to changes of "the" illuminant imply, but also to the spatial variations of illumination across different parts of the scene or even across individual objects

In everyday experience, most colored objects do seem to maintain approximate color constancy across diverse illuminants, suggesting that our visual systems are largely successful at unscrambling objects' reflectances from illuminants. This is generally considered the great achievement of color constancy; Hering (1920) wrote,"The approximate constancy of the colors of seen objects, in spite of large quantitative or qualitative changes of the general illumination of the visual field, is one of the most noteworthy and most important facts in the field of physiological optics." Examples of failure of color constancy with changing illumination can be enjoyed in the clouds at sunset, as they run through a dramatic range of colors due entirely to changing illumination, or on a uniform white movie screen, which fortunately does not maintain the constant appearance of a large white rectangle as the projector changes its illumination.

Variations of Background. Besides changing illumination, objects may also be seen against a variety of different backgrounds. In ordinary viewing, this rarely seems to affect their perceived colors. And because the local light signals from ordinary, opaque objects depend on the object's surface and its illuminant as shown in Fig 1, but not on its background, this phenomenon of color constancy with variations of background may at first seem a trivial "achievement". But the well-known effects of color contrast certainly do affect color appearances. Helson (1940) pointed out that "Hue, lightness, and saturation depend not only upon composition and intensity of light from an object but fully as much upon the reflectance of background and other objects" Whittle and Challands (1969) and Gilchrist (1979) noted that the reliance of the visual system on edges and contrasts makes constancy with changing backgrounds an important challenge for lightness constancy. Hamilton (1979) made the provocative suggestion that the evolution of trichromacy in our primate line may have been driven not by the reflectances of objects of interest, but by the need to identify them against a variety of different backgrounds. Failures of color constancy with changing background include the well-known simultaneous contrast effects. One striking example is the moon's appearance as bright white against the deep black background of the heavens, even though its surface is dark grey rock.

Variations of Atmospheric conditions. In the usual idealizations of color signals, such as the cartoon in Figure 1, the proximal signals are equated with the light reflected from objects. In fact, the atmosphere through which light is transmitted can greatly attenuate and distort this light. These filtering effects can arise from vision through fog, haze, precipitation, smog, water or other transparent media. In general, these atmospheric effects increase with viewing distance; the "aerial perspective" of Leonardo, in which distant objects were painted with more blue and less contrast than near objects, takes advantage of this common effect of atmospheric haze. Color constancy with variations in atmospheric conditions has not been extensively studied, but there is evidence that visual processes provide partial compensation for the reduced contrasts in scenes viewed through "veiling illuminants" (Gilchrist & Jacobsen 1983) or underwater (Emmerson and Ross 1987). Brown & MacLeod (1997) studied the dependence of color appearance on the variance [roughly, the contrast] of surrounding colors, and suggested that the observed perceptual compensation for reduced contrast may be involved in color constancy with variations in atmospheric conditions. A typical failure of color constancy with atmospheric conditions occurs when objects become greyer and eventually disappear as the fog rolls in.

Other Variations. In addition to the above, there are a variety of other sources of variation in the color signals from objects that pose difficulties for color constancy. Changes in viewing geometry, such as may arise from relative motion of the object, illuminant or viewer, or from gradients of surface orientation within objects, may have important effects. Many of these effects are obscured by the canonical one-dimensional approximations of reflectances and illuminants (such as in Fig 1) widely used in color constancy models; this oversimplification of complex physical properties led Wandell (1996) to call this type of canonical model “a ruse”. For example, changes in the surface orientation with respect to both the direction from the illuminant, and the direction to the observer, will generally affect both the intensity and the spectral distribution of the received light signals. The changing colors of iridescent butterfly wings as they move provides one example of such a failure of color constancy with changes in viewing geometry. And variations in the spatial distribution of illumination, such as from a spotlight to diffuse lighting, may also have dramatic effects. Specular highlights, and variations of shading with shape, are 2 common manifestation of these effects of viewing geometry on color signals; note that these both also provide valuable information about the objects and their illuminants.

Another effect of viewing geometry lies not in the external signals, but in anisotropies in the eye itself, such as spatial variations in its optics, resolution, and light and spectral sensitivities. These may produce large changes in object colors as they are seen in different parts of the visual field. For example, colored objects appear decreasingly saturated and eventually turn grey as their retinal images move toward the periphery. Colored patterns such as stripes may also appear to change color with viewing distance as the color contrast effects seen at close range become assimilation and eventually a homogeneous mixing at further distances. Even brightly colored objects will look grey under dim illumination, as the cone signals become too weak to support color vision, even though the physical chromatic contrasts are unaffected. Finally, dynamic variations in the sensitivity of the visual system, including changing pupil size and neural adaptations to previous stimuli, also affect perceived colors, with colored afterimages providing a dramatic example. Although such neural sources of variation in early color signals are not typically considered part of the domain of color constancy, these are the proximal signals for subsequent brain processes, and so they should also be kept in mind.

3. ILLUMINANTS VERSUS BACKGROUNDS

While all of the above types of variations may pose difficulties for color constancy, and must be included in any eventual comprehensive model, this chapter focuses just on the problem of achieving color constancy when both backgrounds and illuminants may be varying. These two sources of variation have distinct but largely complementary effects on color signals, as illustrated in Figure 2. Consider a reference spot seen on a particular background, under a particular illuminant, as in the top stimulus. If only the background reflectance changes, as in the bottom left stimulus, the light received from the spot itself would be unaffected, but the contrast of the spot relative to its background may change dramatically - even reversing sign, as in the example shown. On the other hand, if only the illuminant changes, as in the bottom right, this directly affects the light reflected from the test spot itself, but usually has only a minor effect on the contrast of the test spot relative to its background. Thus, if the color of the test spot was entirely determined by just the light from the spot itself, as in a purely local model of color vision, there would be excellent color constancy when the background changed, but none when the illumination changed (as the spot color would change in proportion to the illuminant change). On the other hand, if the color of the test spot depended entirely on the contrast of the spot relative to its background, as in a number of color constancy models, there would be good color constancy when the illumination changed, but none when the background reflectance changed (as now the spot would change color in proportion to the background change). Thus while may be easy if only one of these can change, the challenge for color constancy to succeed in the real world is to simultaneously handle both types of change. Brainard & Wandell (1986) remind us that "Human vision maintains approximate color constancy despite variation in the spectral reflectance functions of nearby objects and despite variation in the spectral power distribution of the ambient light." Whittle and Challands (1969) suggested that this requires two types of constancy: "first with respect to changes of background alone such as occur with relative movement of the object and its background, and second with respect to changes of illumination of the object and its background together."

In general, when surround signals change as in Figure 2, the proximal signal is ambiguous whether this change arose from a change in illumination, a change in background reflectance, or some combination of the two. In a typical color constancy experiment, a subject is shown a fixed reference spot and surround as in the top of Figure 2, and asked to adjust the spot in a test surround, such as the those in the bottom of Figure 2, until the test spot matches the reference spot. Note that although the two surrounds shown in the bottom of Figure 2 provide identical proximal signals, constancy would require the subject to somehow discriminate them and make very different spot settings in each. When such experiments are simulated with emissive displays controlled by computers, not just the proximal signals but in fact the entire stimulus is physically identical for the two conditions; whether the change in the surroud is considered a change of background reflectance or a change of illumination, and in consequence the degree of color constancy assigned to the subject’s response, depends only on which occult software variable was changed. In practice, owing to the commonly restricted formulation of color constancy in terms of changing illumination, such experiments usually equate color constancy with constant contrast responses (as if compensating for changing illumination). But perhaps such experiments should best be interpreted more neutrally, in terms of spatial interactions affecting color, rather than as measures of color constancy.

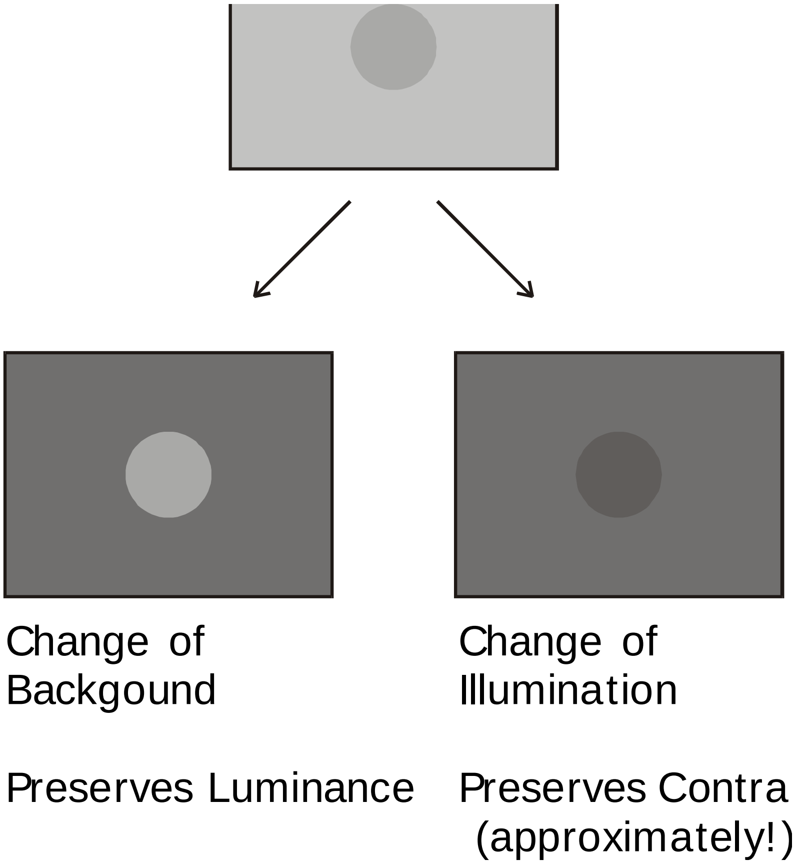

********************** Figure 2 about here ***********************

4. MODELS OF COLOR CONSTANCY

With the paradigm shift from color as a property of objects to color as a response to light, the mystery of color constancy arose. Because the intensity and spectral composition of light from objects changes with their illumination, shouldn't their perceived colors also change? So the challenge for color constancy was early formulated in terms of invariant object colors despite varying illumination. The local model of color vision, in which the colors seen at each point correspond just to the composition of light at the same point, has been the default model of color vision, and probably still corresponds to the notions of most people with a modicum of science education. It became the reference point for color constancy models to improve upon, as it offers zero color constancy with changing illumination. But this same local model can provide perfect color constancy with changing backgrounds, and so it is useful to include it in the mix of color constancy models. Indeed, based on evidence discussed in section 5, under natural conditions it may actually be a more successful model of chromatic color constancy than some of the more celebrated color constancy models, which provide excellent color constancy with changing illuminants, but at the expense of zero color constancy with changing backgrounds.

Most models are designed to first estimate the illuminant, in order to then compensate for it. Note that estimating the illuminant may be neither necessary nor sufficient for achieving color constancy. For example the models of Wallach (1940) and Cornsweet (1970) can achieve lightness constancy with changing illuminants by using only signals based on the ratios of light across edges, and need not ever estimate or represent the absolute intensities of illuminants in order to compensate for them. But if estimating the illuminant is the goal, it would seem at first that the most sensible thing to do would be to simply look up at the sky, and measure the illuminant directly. But for all the varying and often complex indirect schemes that have been suggested for estimating the illuminant, this direct approach does not seem to be part of any serious color constancy models. But two approaches are closely related: Land and McCann’s original Retinex model (see McCann 1989) used the light reflected from white surfaces to directly estimate the illuminant, and others have used specular reflectances within the scene the same way.

Most recent models of color constancy can be divided into two broad (but partially overlapping) classes. The first group of models takes advantage of known visual mechanisms such as contrast and adaptation to estimate or normalize to changing illuminants. The second group is based on linear models of the physical properties of illuminants and reflectances, and constructs algorithms to reconstruct the most likely surface reflectances from the proximal signals.

The visual models have generally been directed toward achieving color constancy with respect to variations in illumination, but not changes in background reflectance. Given one proximal light signal and two unknowns (background and illuminant), the easy solution is to try to eliminate one of the unknowns. So generally a strong assumption is made which severely constrains or eliminates changes of background reflectance, leaving only illumination changes to account for. Lightness constancy. The most common of these assumptions is some form of the “Grey World” hypothesis: that the space-averaged reflectance of visual scenes is a neutral middle grey, and thus that the space-averaged light from each scene is proportional to its illuminant. Grey World models include those based on von Kries adaptation to the entire visual scene (e.g. Ives, Helson 1943, Lennie & Fairchild), and the more recent Retinexes (Land 1983). A modification of this model assumes only that the space-average reflectance from scenes is known, but is not necessarily neutral grey (Buchsbaum 1980). When the Grey World assumption holds true, adaptation to the space-average light would tend to achieve color constancy despite changes in illumination. The critical assumption is that all deviations from the standard in the space-averaged light from scenes represent changes of illumination. If this Grey World assumption is violated, such as might happen in a forest of predominantly green surfaces, these models will sutomaticall normalize as if the illuminant were green and the leaves grey, and consequently generate strong failures of color constancy. Brainard & Wandell (1986) studied this problem in Land's Retinex, and concluded that "the algorithm is too sensitive to changes in the color of nearby objects to serve as an adequate model of human color constancy."

The second general class of color constancy models involve constructing linear models of the physical properties of reflectances and illuminants, and devising computational solutions to estimate the most likely reflectances from these models. These models, unlike the visual models, may bear little resemblance to the known properties of biological visual systems, and their design may be oriented more toward machine vision than human vision. (Many are even designed to reconstruct the full spectral reflectance functions of surfaces; this seems akin to suggesting that the goal of olfaction is to reconstruct 3-dimensional molecular models of odorants.)

The linear models are based on the hypothesis that low-dimensional linear models are sufficient to represent both the spectral reflectances of surfaces and the spectral power distributions of illuminants for color vision (Cohen 1964, *****). Such models are usually limited to 2 or 3 dimensions, corresponding to the trichromacy of human color vision. Then under specified conditions, it may be possible to precisely reconstruct the reflectances and illuminants of scenes. These models would not be expected to succeed if their assumptions about the low dimensionality of natural color signals are violated.

Both types of color constancy models can perform extremely well in artificial model worlds which incorporate their assumptions. (For the visual models, these are most commonly that changes in space-averaged background reflectances are small compared to change sin space-averaged illumination; for linear models, that natural reflectances and illuminants have low dimensionality.) In fact, these model are commonly “tested” only in conditions that incorporate their own assumptions, and seldom under natural viewing conditions in the “real world”. So how well do their assumptions represent the natural color environments for which color vision evolved? This becomes an empirical question, which may be addressed by analyzing the actual properties of natural surfaces and illuminants.

5. ECOLOGY OF COLOR SIGNALS

An evolutionary view of perception holds that our sensory and cognitive processes are tuned to the relevant ecological signals and challenges that affect survival and reproduction (von Uexkull 1909, Vollmer 1984, Delbruck 1986). From this perspective, the purpose of color vision would be not to solve some general computational problem, such as reconstructing all the physical reflectance functions, but to reliably inform the organism about behaviorally relevant color signals in the environment. To understand color vision, then, it may be more valuable to study the properties of natural color signals than the physics of light. Measuring and analyzing ecological color signals serves two purposes here: One is to test the assumptions and constraints built into current color constancy models against the real world; and the other is to seek insights from these color signals into the challenges and possible solutions of color constancy.

Individual Natural Reflectances

The interactions of a surface with light are surprisingly complex. To properly characterize surfaces requires high-dimensional functions that include not just the dependence on the wavelengths of incident light, but the incident angle, the reflected angle, and in the case of fluorescent surfaces the various wavelengths of emitted light as well. In most models of color vision, surface reflectances are approximated by one-dimensional functions of reflectance versus wavelength. (Indeed, in lightness models, they are often just a scalar!) The measured surface reflectances to follow are such one-dimensional approximations, and so they presumably underestimate the diversity and dimensionality of reflectances.

Cohen (1964) analyzed the spectral reflectances of 433 Munsell color chips and found them to be well characterized by a 3-dimensional basis set. But the Munsell chips are artificial surfaces, explicitly designed to be well-behaved for human color vision, and so may not represent very well the true diversity of natural surfaces. Another commonly studied set of natural reflectances is Krinov's (1947) measurements of terrains, but as discussed below this should not be considered a set of individual object reflectances. A third approach to collecting natural light signals makes simultaneous spectral analyses of the light from many points in a scene (Webster and Mollon 1997, Ruderman et al. 1998), but these light signals are not true reflectances, but the products of reflectance and illumination, and there is no practical way to analyze just the reflectance components from them. Lacking an appropriate data set for natural spectral reflectances, I undertook to compile one using a portable Photo Research PR-650 spectrophotometer. By measuring the spectral power distributions of light reflected from various natural surfaces, and dividing by the light reflected from an artificial white standard held in the same location and orientation, one-dimensional approximations of surface reflectances were obtained.

Ideally one hopes to study a set of surface reflectances representative of the natural surfaces for which our color vision evolved to see. For practical reasons, I settled for measuring the surface reflectances of natural objects found within a few miles of the University of California, San Diego. Also for practical reasons, I did not randomly sample points from across natural scenes, but tried to capture the gamut of natural colors by selecting a wide variety of colored objects. So brightly colored flowers and fruits are vastly over-represented relative to their actual frequency in nature, while large expanses of sand and leaves were very undersampled relative to their spatial extent. But note that relatively uncommon but salient color signals, particularly of colored fruits, may have been a dominant driving force in color vision (Polyak 1957), while the variations in background colors of leaves may have been spectral noise that our color vision actually evolved NOT to see (Nagle and Osorio 1993).

A total of 563 such natural objects were measured relative to the white standard to obtain a new set of natural object reflectance data (Brown 1994). These reflectances were highly diverse and contained far more spectral variation than was found in the Krinov set. To represent the chromatic range in this set, the chromaticities of these objects, under simulated illumination by CIE Source C, are plotted in Fig3. The same data are shown on both the standard CIE chromaticity plot (left), and on the more physiologically relevant MacLeod-Boynton chromaticity plot (right), in which the axes correspond to the axes of the color-opponent channels. Interestingly, the chromaticities almost all lie to one side of a line through the white point, and are biased toward the red, yellow and green regions, with little representation of blues or turquoises. This chromatic bias in individual reflectances suggests that space-averaged backgrounds are also unlikely to be neutral grey, as the Grey World hypothesis requires.

Averaged Natural Reflectances: It is also useful to measure the space-averaged reflectances of various scenes. The best-known set of such measurements are Krinov's (1947) spectral data from Russian terrains. These terrains have sometimes been described in color constancy studies as natural objects, which in one sense they are, but they are certainly not small objects such as leaves and fruits, but rather very large expanses of terrain, measured while rotating the spectrophotometer or even from an airplane. As Krinov noted, "Thus the data obtained refer basically to average natural backgrounds." These spectral measurements correspond to virtual reflectances of the whole terrain. (They are not true averages of individual reflectances, because the actual illuminant may vary from point-to-point across the terrain, and particularly in the shadows, while this approximation assumes a uniform illuminant.) The gamut spanned by the chromaticities of these average natural backgrounds, again assuming Source C illumination, is shown in Fig 3. Krinov's set clearly spans a much smaller gamut of chromaticities than the individual object reflectances, suggesting that treating Krinov's set as represenative of natural object reflectances leads to a great underestimate of the actual diversity of object reflectances.. On the other hand, treating Krinov's set as representative of space-averaged background reflectances, it is hardly the tight cluster around the white point that the Grey World hypothesis assumes. On average, these

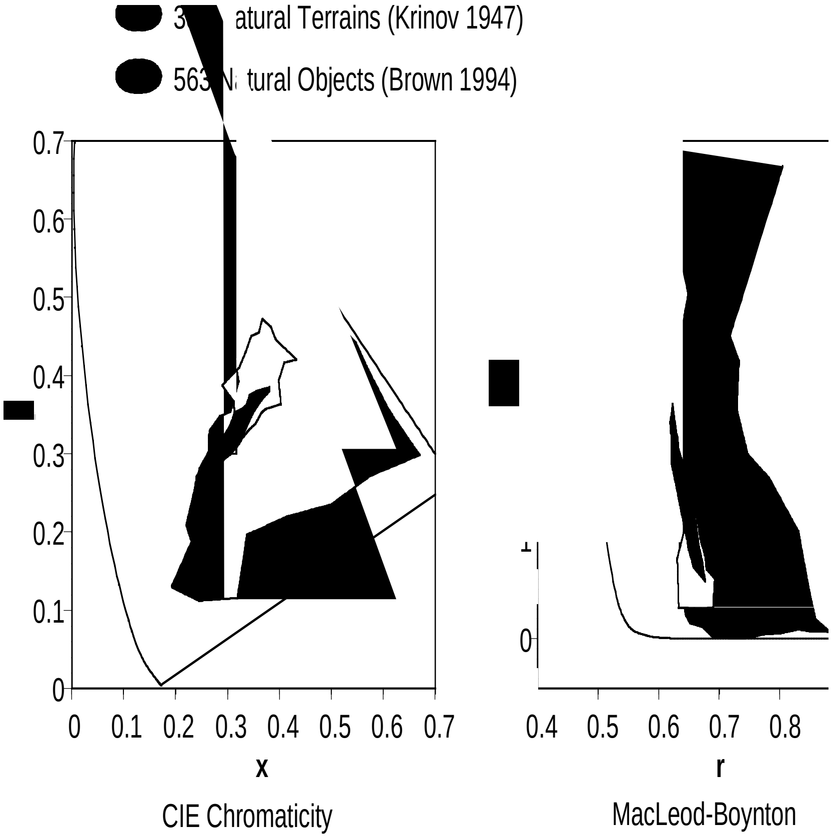

terrains were a dark desaturated yellow-green, demonstrating that the world is not grey. (I made similar space-averaged measurements in San Diego, using a calibrated diffuser over the spectrophotometer, and obtained a gamut of terrains similar to Krinov's but relatively lacking in the green regions.)

********************** Figure 3 about here ***********************

Individual and Averaged Natural Illuminants:

Illuminants, like reflectances, are more complex than the one-dimensional spectral power distribution functions commonly used to represent them. The spatial distribution of all the sources of illumination, and the polarization, are two important factors omitted from such approximations. Moreover, the natural illuminants studied are usually just the daylight from the sky, while the actual light reflected from objects may come from a variety of other sources including previously reflected or filtered light with significantly different spectral characteristics.

There have been many studies of the spectral power distributions of natural illuminants (Henderson 1970). One of the most influential was that of Judd, MacAdam, & Wyszecki (1964), who analyzed the data from 3 other sets of measurements of daylight. They plotted the chromaticities of each sample, and the gamut spanned by these chromaticities is replotted in Fig 3. These daylights form a rather compact range of chromaticities, somewhat smaller than the range of average natural backgrounds, and oriented along a blue-yellow axis. According to the Grey World Hypothesis, the range of daylight variance should be larger than that of space-averaged backgrounds. If in fact the opposite holds, as appears to be the case here, normalizing to the average of the background will likely backfire, by misinterpreting the chromaticity of the background as that of the illuminant.

The daylights studied by Judd, MacAdam, & Wyszecki and plotted in Fig 3 included both measurements of the integrated daylight from the entire sky, and measurements taken in the shade and representing essentially the chromaticity of the sky minus the sun. Such restricted daylights do often occur in the shadows of scenes, but they do not represent space-averaged illuminants a scene would be likely to have. Two other sets of correlated color temperature measurements of daylight which compared these two types of daylight measurements (full sky, and sky-minus-sun) were also obtained from Henderson (1970) and replotted in Fig 4. This clearly shows that almost all of the chromatic variance in mixed daylight sets (such as that in Fig 3) was restricted to the sky-minus-sun measurements, reflecting the well-known fact that the sky is blue. But when sunlight is included in measurements of daylight, there is remarkably little chromatic variance in these daylights.

This implies that most of the chromatic variance in daylights is found in the spatial variations between sunlit and shady regions of a scene, particularly when the sky (and thus the shadows) is blue, and not in the temporal variations of space-averaged illumination across entire scenes. Since it is the space-averaged illuminant that most color constancy models seek to estimate or adapt to, it would seem that the problem for color constancy of varying chromaticities in natural illumination has been greatly overestimated. On the other hand, illumination does vary from point-to-point across scenes, and this will pose a serious challenge for color constancy. Endler has analyzed the illumination in natural forests on a fine scale, and characterized many additional contributions, such as reflections and filtering through leaves, in addition to light from the sky. So while the Judd, MacAdam & Wyszecki set overestimates the chromatic variance in space-averaged illuminants, it underestimates the chromatic variance in spatially varying illumination.

*********************** Figure 4 about here ************************

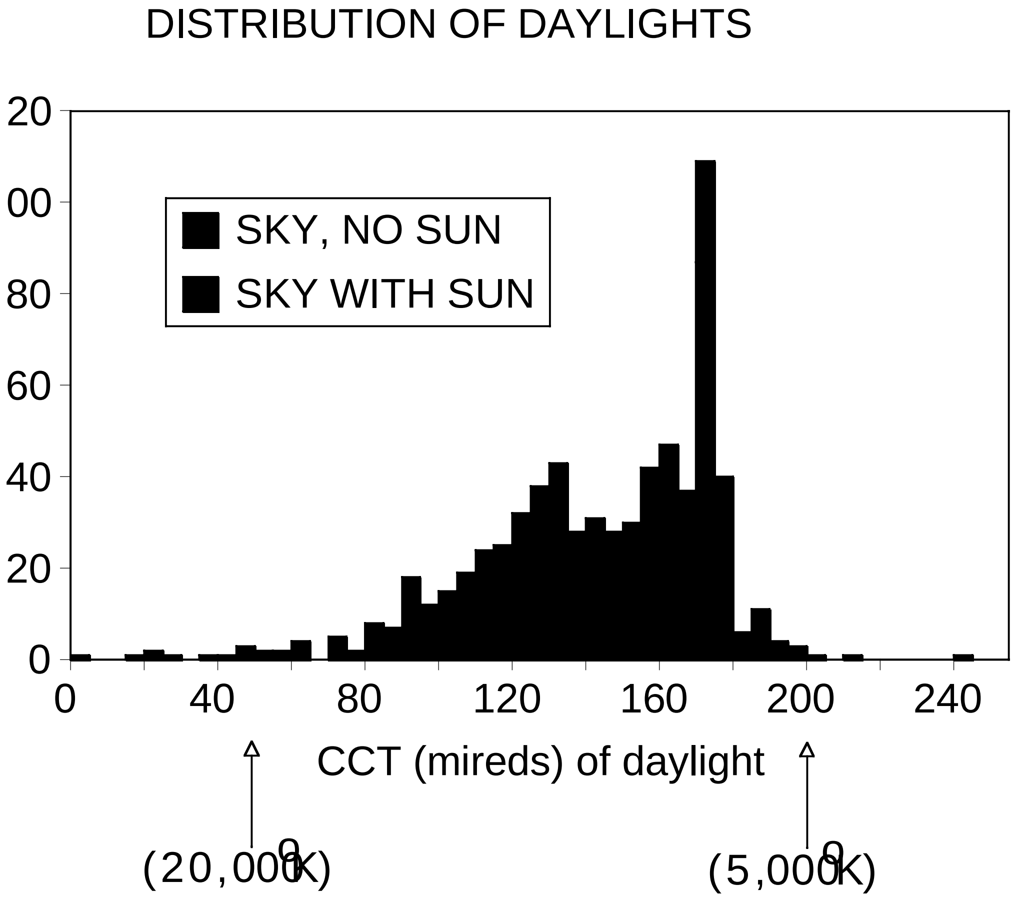

Analysis of Spectral Variance:

The linear models approach to color constancy relies on the assumption that natural reflectances and illuminants can be represented by low-dimensional basis sets. Notably, Maloney reported that Cohen’s 3 principal components derived from the reflectances of the Munsell set could account for >99% of the variance in the reflectances of Krinov’s natural terrains set. But Maloney used an unusual metric, of variance not from the mean of the set but from absolute zero reflectance (black), and consequently the bulk of the variance accounted for represented just the DC signal in these all-positive reflectance functions. It would be interesting to know just how much of the chromatic variance (roughly, the variance in the shapes of the reflectance functions) could be accounted for by such a low-dimensional set. Applying another metric, of the variance from the mean accounted for in normalized reflectance functions, yielded starkly different results: now Cohen's Munsell set could account for just 43% of the chromatic variance in the Krinov set, and 65% of the variance in the individual object reflectances. A new 3-dimensional basis set derived from natural reflectances was also developed, and this improved these values of variance accounted for to 76% for the Krinov set, and 87% for the individual reflectances. This suggests that at least for chromatic constancy, the low-dimensional linear models will generate much larger errors than was previously appreciated. Endler also analyzed the variances in hi measurements of microilluminants, and found that while 3 basis functions could capture most of the variance in illumination within each scene, a different set of 3 basis functions was generally required for each different scene. It has been suggested that the trichromacy of human color vision might represent an optimal adaptation to an underlying 3-dimensionality in natural color signals, of surfaces and/or illumiannts (Cohen 1964, Shepard 1992). The present analysis suggests that natural signals contain considerably more than 3 dimensions of variance. Of course, the fact that many or most nonmammalian vertebrate animals have evolved 4 or more dimensions of spectral sensitivity also provides strong evidence that there remain natural color signals for which we humans are colorblind.

In addition to the chromatic variances shown in Figs 3 and 4, reflectances and illuminants have important achromatic variances in amplitude. Fig 5 illustrates the range of achromatic variance in for illuminants (the Judd, MacAdam & Wyszecki set) and reflectances (the Brown set), in comparison with the chromatic variances along the r-g and y-b color opponent axes (corresponding to the axes of the MacLeod Boynton chromaticity diagram in Fig 3). Natural illuminants span an enormous range of intensities, roughly 10 orders of magnitude from starlight to bright sunlight, while the range of total reflectances of natural objects spans less than 2 orders of magnitude from deep black to white. This disparity in the range of intensities is indicated on the logarithmic scale of luminance at left. But with this intensity variance removed, the natural reflectances span a much larger range of the chromatic axes than do the natural illuminants. Apparently, the chromatic and achromatic dimensions of color vision face very different challenges from changing illuminants versus changing backgrounds. Might the visual system have evolved distinct strategies for achieving achromatic and chromatic color constancy, in tune with these differences in the ecological color signals?

*********************** Figure 5 about here ************************

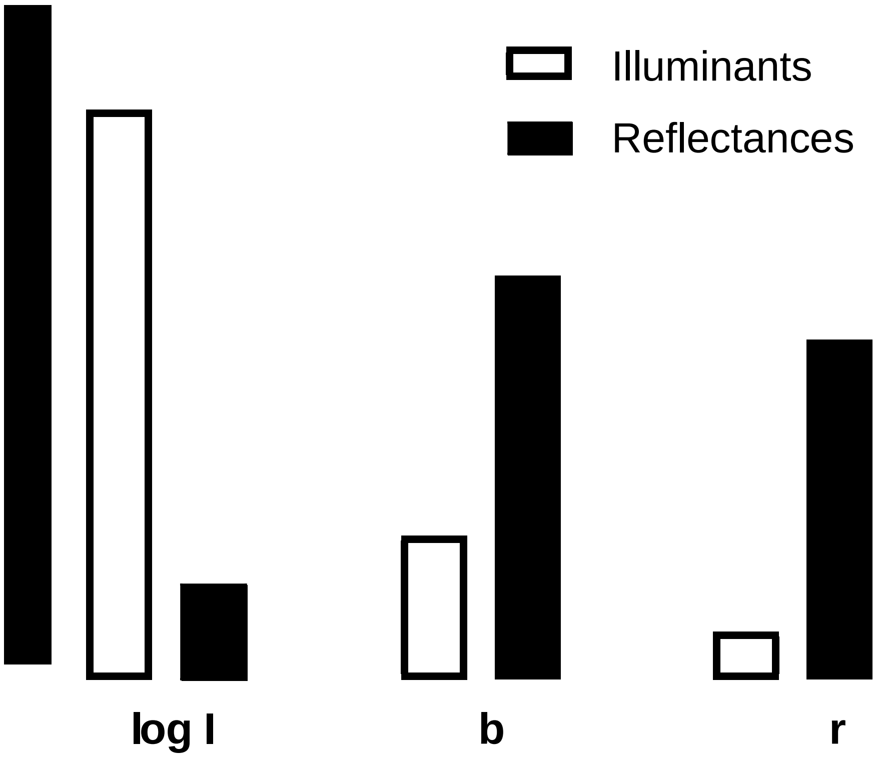

6. ASYMMETRIES IN THE OPPONENT CHANNELS

The initial human visual response to color signals occurs in the 3 types of cones, which each respond to the intensity of illumination weighted by different spectral sensitivities. Each type of cone is subject to almost the same 10 billion-fold range of illuminant intensities, and to the 100-fold range of reflectances, as the luminance signal of Fig 5, and so the 3 cone channels are practically symmetrical in their relative variance due to illuminants and backgrounds. Thus the cones do not seem well suited to take advantage of the achromatic/chromatic asymmetry in natural color signals.

The retina transforms the initial cone signals into the luminance and color-opponent signals of retinal ganglion cells which the brain uses for subsequent visual processing, including color vision. This transformation places almost all of the intensity variation into the luminance channel, while the remaining chromatic signals are in the 2 color-opponent channels. This suggests that the opponent transformation places almost all of the variance in illumination into the luminance channel, and leaves the 2 color-opponent channels to deal with variations due primarily to reflectances. While there may be many other reasons for this retinal transformation into opponent channels (see MacLeod and von der Twer, this book?), the observation that it largely separates the illuminant variation from the reflectance variation suggests that one advantage it may offer is to facilitate color constancy with changing backgrounds and illuminants.

The transformation from cone responses into opponent channels has not generally been regarded as particularly relevant for color constancy. According to David Hubel (1988), "the two ways of handling color – r, g, and b on the one hand and b-w. r-g, and y-b on the other- are really equivalent." But there are a number of intriguing asymmetries in the properties of the luminance and color-opponent channels which might represent adaptations to their different distributions of natural color signals. Some of these asymmetries are discussed below, with speculations on their possible relevance for color constancy.

(1) Saturation of chromatic but not achromatic induction at low contrast

One of the striking but often overlooked differences between achromatic chromatic color induction is that the strength of achromatic induction increases monotonically with increasing contrast, while the strength of chromatic induction saturates at surprisingly low chromatic contrasts, and is flat or even deceasing for higher contrasts. A nice demonstration of this (from Meyer, described in Helmholtz 1866) is the tissue paper effect: a grey piece of paper on a saturated color background generally appears lightly tinged with an induce complementary color, but desaturating the background by overlaying a piece of white tissue paper may dramatically increase the induced color. Kirschmann (1892) reported diminishing returns with increasing chromatic contrast in his classic studies of color induction. De Valois et al. (1986) also reported an asymmetry between achromatic and chromatic induction with achromatic induction continuing to increase over a much larger range of contrasts than chromatic induction. Might this asymmetry in induction relate to a corresponding difference between luminance and achromatic variance? Helmholtz suggested that Meyer's Effect could be interpreted in terms of the likelihood that the light from the inducing surround corresponds to the color of the illuminant. Since the chromatic range of natural illuminants is quite small, it might make sense to "bet" that only desaturated surrounds are likely to represent the chromaticity of the illuminant. On the other hand, for achromatic induction there would be practically no limit to the range of surround intensities corresponding to illuminant intensities. It must be noted however that humans also achieve color constancy over a large range of artificial illuminants, spanning a much larger gamut of chromaticities than natural illuminants, and the simplistic assumption that illuminants rarely vary in chromaticity cannot account for this.

(2) **Chromatic but not achromatic sensitivity to diffuse stimulation**.

It is commonly noted that the early visual system responds primarily to local contrasts, and thus loses sensitivity to uniform changes. Whittle and Challands wrote that "loosely speaking, the visual system differentiates the input, and to achieve certain perceptual goals we have to integrate it. Because of this contrast dependence, there is remarkably little effect of changing the overall intensity of ilumination (Walraven et al. 1990), and has been taken as the key to lightness oconstancy. But there is considerabl;e evidence that the color-opponent channels are not so dependent on spatial contrast. In De Valois et al (1958) study of primate LGN cells, strong responses were found to diffuse monochromatic light. The center-surround structure of retinal ganglion cells and LGN cells also supports this distinction. Typical Type II cells, with a cone-selective center and a complementay or non-selective antagonistic surround, provide the same cells sensitivity to spatial variations in luminance contrast while maintaining sensitivity to spectral variations in uniform illumination (Hubel 1993). A 2-DG study by Tootell et al. 1988 also found much stronger cortical responses to spatially diffuse chromatic variation than to spatially diffuse luminance variations. And a number of psychophysical studies have found a corresponding effect in which contrast sensitivity requires much greater spatial contrast for luminance than for chromaticity (Mullen 1985). So perhaps the visual system does not "differentiate" the chromatic signals as much as it does the luminance signals. Since the space-averaged chromaticity of light from a scene is more likely to represent surface reflectances than illuminants, allowing the DC chromatic signals to pass might be an important contributor to chromatic color constancy, possibly accounting for why the green forest continues to look green and not a normalized grey. At the same time, the empirical observations that chromatic ganzfelds lose their perceived color (Cohen 1964), that retinally stabilized stimuli fill in with the surrounding color (Iarbus 1967), and that dichoptic contrast matches are determined almost entirely by contrast (Whittle, this volume) indicate that under these circumstances local chromatic signals may not reach the brain.

(3) Dependence of chromatic induction on luminance contrast.

Another interesting asymmetry among the opponent channels lies in the interaction between chromatic induction and achromatic contrast. Kirschmann (1892) reported maximal color induction occurred when the inducing surround and the induced test spot had equal brightness. Others have found that induction is strongest for surrounds equal or greater in luminance, and falls off when the surround is dimmer than the test spot (Hurvich and Jameson 1959, Kinney 1962). What implications might this have for color constancy? Recall that most of the chromatic variance in daylights was associated with the spatial changes from sun to shadow under blue skies (Fig 5). If the strongest chromatic deviations in natural illuminants are associated with deep shadows, it may be valuable to link chromatic induction to luminance relations. The experiment shown in Fig 6 provides a hint for how this interaction may be used to promote color constancy with spatially varying illumination across a scene. The surround consisted of sectors varying in luminance and chromaticity. Induction was measured into test spots of variable luminance. The result was that induction into the test spot was strongest from sectors having the same luminance as the test spot, consistent with Kirschman's Law. In other words, the dark test spots were most affected by the dark sectors of the surround, and the light test spots by the light sectors of the surround. This mechanism, if it operates similarly when viewing natural scenes, would tend to segregate chromatic interactions within shadows from those within sunlit areas. These results also provide another strong challenge to models based on the average light from surrounds: The surround with large dark green sectors and small bright purple sectors had the identical space-averaged light as the complementary surround with large dark purple sectors and small bright green

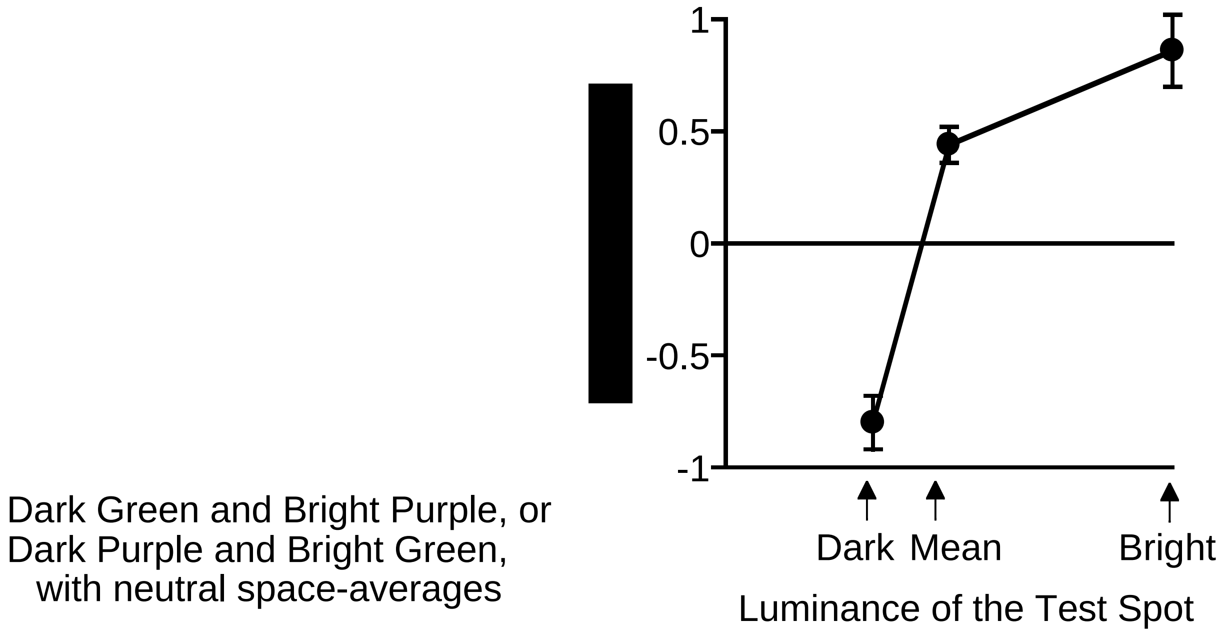

Color Induction with Unequal Luminance Surround

sectors, yet they had opposite color induction effects on the test spots. Both surrounds also had essentially the same luminance variances and chromatic variances, but differed in the covariances of luminance and chromaticity.

*********************** Figure 6 about here ***********************

In considering the possible contribution of the opponent channel transformation to color constancy, it is necessary to assess the likelihood that the neural processing underlying color constancy actually involves opponent representations. Many models of color constancy, especially those involving von Kries adaptation, are based on cone-based representations. The cone signals are transformed by retinal processing to generate the luminance and color-opponent retinal ganglion cell signals, which are sent to the

LGN and thence to the primary visual cortex for further processing. Both the LGN and the primary cortex maintain opponent representations, and so it seems a reasonable assumption that postretinal color processing occurs largely with opponent representations. Therefore, evidence that the adaptations and spatial interactions involved in color constancy are postretinal would support that the likelihood that they occur on an opponent representation, and not in a cone space.

A number of lines of evidence does suggest a cortical site for much of the processing involved in color constancy. There is the anatomical evidence that receptive fields in precortical visual areas are generally quite small, and not likely to support the long-range interactions involved in color constancy. Zeki's (1983) study of color responses to Mondrian displays under changing illumination found that cells even in V1 were responding based primarily on the wavelengths and not the color appearances, while cells inV4, with much larger receptive fields and long-range connections, had responses resembling color constancy. Land et al (1983) studied a patient with a severed corpus callosum, and found that color induction effects were restricted to half visual fields, suggesting that the spatial interactions involved must be cortical. De Valois et al. (1986) suggested that the temporal properties of color induction also indicate a cortical locus. Alan Gilchrist (1977) demonstrated the effects of depth, which presumably involve cortical processing, on color induction. Olson and Boynton (1984) found that color induction in stimuli presented to one eye can be masked by stimuli presented to the other eye, and wrote, "It is concluded that the basis of the chromatic induction is largely or entirely nonretinal." And two studies (Thompson & Latchford 1986, and Webster Day, & Willenberg 1988) used the McCollough Effect of orientation-contingent color adaptation to demonstrate that the McCollough Effect adapts to the local physical composition of light, not its perceived color; this implied that the perceived color was generated after the presumably cortical locus of the orientation-dependent effect. The above experiments all point to the likelihood that much of the processing involved in color constancy has a cortical locus, and so are likely to involve opponent representations.

This hypothesis suggests that the visual system evolved to take advantage of the asymmetries in the achromatic and chromatic dimensions of color signals for changing backgrounds and illuminants by confining almost all of the achromatic variance into one channel, leaving primarily reflectance variance in the other two. The asymmetries and other idiosyncracies of color induction in these channels may be tuned to take advantage of this difference, and make the problem of color constancy easier. But by no stretch does this solve color constancy, as each dimension still has to deal with varying illuminants and backgrounds, just with different distributions of these. The luminance channel must deal with essentially the problem of "lightness constancy", over a range of 10 orders of magnitude of illumination and 2 orders of magnitude of reflectance. And while the color-opponent channels have relatively small chromatic variance in natural illumination to contend with, this can still be significant. For example, Maksimov (check ref) pointed out that the leaf of a dandelion in direct sunlight has approximately the same physical chromaticity as a dandelion flower in deep shade, yet the flowers look yellow and the l;eaves green in both sun and shade.

7. SUMMARY

The seeming immediacy and reliability of color vision belies the inherent ambiguity of color signals, and the complexity of the neural processing involved in color perception. The effort to understand human color vision has been one of the major enterprises of perceptual study, engaging the efforts of generations of top researchers including many of history's most celebrated scientists, such as Leonardo, Newton, Young, Helmholtz, Maxwell, Mach, Schrodinger, and Crick. Despite this enormous effort, we still lack a model of color vision which can successfully predict the perceived colors even in simple 2-dimensional "Mondrians".

Much of the difficulty arises from two essential but complementary aspects of color vision: its dependence on the spectral power distribution of light from each point in a scene, reflected in the cone quantum catches, and its dependence on comparisons between the light signals from across the scene, reflected in postreceptoral contrast signals. The local signals involve equally the surface reflectance properties of objects and the spectral power distribution of illuminants, and so make it difficult to maintain object color constancy across varying illuminants. Contrast representations are largely invariant with changing illuminants, but introduce problematic sensitivity to changing background reflectances.

Two classes of color constancy models make strong assumptions about the properties of natural color signals in order, which were tested in an analysis of natural spectral reflectances and illuminants. The frequent assumption made by visual models of color constancy that the space-averaged light from scenes may be used to estimate the illuminant was shown to be inconsistent with the natural data for chromatic variance, but a more plausible hypothesis for intensity variations. The assumption made by linear models of color constancy that natural color signals are well represented by 2- or 3-dimensional basis sets was challenged by the inability of such models to represent the chromatic variance in natural reflectances.

A further analysis of the natural reflectances and illuminants found an interesting asymmetry between the luminance and chromatic dimensions of color vision. The luminance channel must handle an enormous range of illuminant intensities, and a much smaller (though still large) range of reflectances, corresponding to the problem of lightness constancy. But the chromatic variance in the color-opponent channels is considerably larger for natural reflectances than natural illuminants. Evidence from neuroscience and psychophysics suggests that much of the processing for color constancy occurs at or beyond this opponent processing stage. This raises the possibility that the visual system has adopted distinct strategies for achieving color constancy in the luminance and color-opponent channels, rather than approaching color constancy as simply “lightness constancy X 3”. Asymmetries in the properties of luminance and color-opponent channels, including the saturation of chromatic but not achromatic induction at low contrast, the sensitivity of chromatic but not achromatic mechanisms to diffuse stimulation, and the dependence of chromatic induction on achromatic contrast, are all possible manifestations of such adaptive responses to natural color signals.

FIGURE LEGENDS

Figure 1. A simplified view of color vision.

A colored object, such as the banana shown in this cartoon, may be roughly characterized by its surface reflectance function, which plots the proportion of incident light reflected at each wavelength. The light illuminating an object is represented by its spectral power distribution, which shows its intensity at each wavelength. The light reflected from the object at each wavelength is given by the product of these two measures. In the human visual system, the proximal light signal is sampled by 3 types of cones, each of which may be characterized by its spectral sensitivity to light of different wavelengths. The resulting triplet of cone responses, L, M and S, provides the initial neural signal available for color vision. Although these cone responses depend as much on the illuminants as on the reflecting surfaces, to achieve color constancy the visual system must generate perceived colors which depend only on the object.

Figure 2. Change of Background versus Change of Illumination.

This illustrates the complementary effects on visual signals caused by two types of variation which pose challenges for color constancy. The reference stimulus, shown above, is a middle grey spot (reflecting, say, 50% of the incident light) on a lighter grey background. If the same spot is seen under the same illuminant but against a darker background, it will still reflect 50% of the incident light, and thus have the same luminance as before, but its contrast relative to the background changes. On the other hand, if the spot is seen under a dimmer illuminant but against the original background, it will now have a proportionally lower luminance (reflecting 50% of a dimmer light), but its contrast relative to the background will not change. To achieve color constancy for both types of change, the spot must produce the same perceived color in all three conditions shown.

Figure 3. Chromatic distributions of natural color signals.

The plots show the range of chromaticities for 3 sets of natural color signals, plotted in both the standard CIE chromaticity diagram (left), and the physiologically-based MacLeod-Boynton chromaticity diagram (right). To plot the chromaticities corresponding to reflectance data, diffuse illumination by CIE Source C (represented by the white cross) was assumed.

Natural Objects (individual black circles): These represent the chromaticities of 563 individual natural reflectances. Data is from Brown's (1994) measurements of the spectral reflectances of 563 natural objects in San Diego.

Natural Terrains (lined region): This outlines the gamut of chromaticities spanned by 337 large, space-averaged natural backgrounds. Data is from the Russian data reported by Krinov, 1947, based on wide-field measurements of spectral reflectances of natural terrains. The data in electronic form were kindly provided by Larry Maloney, and corrected against Krinov's original published data.

Daylights (black region): This outlines the gamut of chromaticities spanned by daylight illumination, from part or all of the sky. These chromaticities were taken from the data published in Judd, MacAdam & Wysezcki 1964. (Some of the original published points could not be resolved, but these were generally near the center of the distribution, so their omission does not affect the gamut shown here.)

Figure 4. Variances of Natural Color Signals in the Opponent Channels.

This illustrates the relative ranges of variation in illuminants and reflectances, for each dimension of MacLeod-Boynton color space (corresponding to the physiological opponent-color channels). For luminance, the variation of intensity in illuminants is many orders of magnitude greater than for reflectances, showing a dramatic difference even on the logarithmic scale of luminance. But for purely chromatic variations, after luminance variation is removed, this relation is reversed, with natural reflectances showing greater variation than natural illuminants in both the blue-yellow (b) and red-green (r) dimensions. Illuminants data from Judd et al. 1964, and reflectances data from Brown, as in Fig. 3

Figure 5. Distribution of Daylights.

This plot shows that almost all of the chromatic variation in measured daylight illuminants was obtained from measurements taken in shadow (grey bar), and corresponds to the varying colors of patches of sky. When the sun is included in integrated measurements of total daylight illumination (black bars), there is very little chromatic variation among measurements. Data from the daylight measurements of Henderson & Hodgkiss (1964) and Winch et al. (1966) were pooled and plotted on a mired scale. (Mireds are an inverse measure of correlated color temperature, which provides a more perceptually uniform scale of colors than color temperature. Data were obtained from Henderson 1975.)

Figure 6. Dependence of Color Induction on Luminance of the Induced Spot.

This summarizes data from an experiment by Brown (1995). When a background varied in both luminance and chromaticity, the color induced into test spots depended strongly on the luminance of the test spots, and was primarily determined by regions of the surround having luminances equal to or greater than the test spots.

In this study, a small test spot of variable luminance was embedded in surrounds varying in both luminance and chromaticity. The 2 complementary surrounds used had identical space-averaged luminances and neutral chromaticity, but one consisted of 3 large, dark green sectors and 3 small, bright purple sectors, while the other consisted of 3 large, dark purple sectors and 3 small, bright green sectors. (The sizes of the sectors were inversely related to their luminances to maintain the neutral space-average. Purple and green were chosen to maximize the saturation available on monitor.) Subjects cancelled the induced color by adjusting the chromaticity of the test spot along a purple-green axis until it appeared neutral grey. Results from 18 subjects (9 for each surround) were pooled, and data . Net color induction is indicated , combined from the two complementary disk constructions, show that when the test spot had the same luminance as the dark sectors, its color appeared complementary to the dark sectors. Test spots with luminances equal to either the bright test spot or the mean luminance of the disk appeared complementary to the bright sectors.

ACKNOWLEDGEMENTS

I am very grateful to Donald MacLeod for his kind support and encouragement of this research. I thank Rainer Mausfeld as well as all the participants in the ZiF workshop for their companionship and their valuable comments and criticisms. This work was funded by NIH EY01711 to DIAM.

REFERENCES

Olson, Conrad X., and Boynton, Robert M. (1984) Dichoptic metacontrast masking reveals a central basis for monoptic chromatic induction. Perception & Psychophysics, 35 (4):295-300

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

Brown, R.O. (1993). Integration of nonlinear contrasts: implications for color constancy. Perception 22 (Supplement):14 (ECVP).

Brown, R.O. (1993). A cone-based linear model of spectral reflectances. Optical Society of America Annual Meeting Technical Digest:252.

Brown, R.O. (1994). The world is not grey. Investigative Ophthalmology and Visual Science 35:2165 (ARVO).

Buchsbaum, G. (1980). A spatial processor model for object-color perception. J. Franklin Institute, 310, 1-26.

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

Cohen, Jozef (1964). Dependency of the Spectral Reflectance Curves of the Munsell Color Chips. Psychonomic Science 1 (12):369-370.

The Committee on Colorimetry of the Optical Society of America (1963). The Science of Color, Optical Society of America, Washington DC

Cornsweet, T. N. (1970). Visual Perception, Academic Press, New York.

Delbruck, M. (1986). Mind from matter- an essay on evolutionary epistemology, Blackwell Scientific, Palo Alto.

De Valois, R.L., Smith, C.J., Kitao, A.J., and Kita, S. (1958). Responses of single cells in different layers of the primate lateral geniculate nucleus to monochromatic light. Science, 127:238-239.

De Valois RL; Webster MA; De Valois KK; Lingelbach B. (1986). Temporal properties of brightness and color induction. Vision Research, 1986, 26(6):887-97

D'Zmura, M.; Lennie, Peter (1986). Mechanisms of color constancy. Journal of the Optical Society of America. A, 3 (10):1662-1672

D'Zmura, Michael; Iverson, Geoffrey (1993). Color constancy: I. Basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces. Journal of the Optical Society of America. A 10: 2148-2165.

Eco, Umberto (1985): 'How Culture Conditions the Colours We See'. pp. 158-75. In Blonsky, Marshall (Ed.) (1985): On Signs: A Semiotics Reader. Oxford: Blackwell

Emmerson, P. G.; Ross, H. E. (1987). Variation in colour constancy with visual information in the underwater environment. Acta Psychologica, 65 (2):101-113.

Endler, J.A. (1993). The color of light in forests and its implications. Ecological Monographs, 63, 1-27.

Evans (1948). An introduction to color, Wiley, New York.

Finger, S. (1994). Origins of Neuroscience, Ch. 7 Color Vision. Oxford University Press, Oxford

Gilchrist, A. L.; Jacobsen, A. (1983). Lightness constancy through a veiling luminance. Journal of Experimental Psychology: Human Perception & Performance, 9, 936-944.

Gilchrist, Alan L. (1988). Lightness contrast and failures of constancy: A common explanation. Perception & Psychophysics, 43, 415-424

Gilchrist, A. L. (1977). Perceived lightness depends on perceived spatial arrangement. Science, 195, 185-187.

Gilchrist, A. L. (1979). Perception of Surface Blacks and Whites. Scientific American 240(3): 112-124.

Hamilton, W.J. (1979) Are Selection Pressures Different? Discussion of Snoddrely (1979) In The Behavioral Significance of Color (E.H. Burtt, Jr., ed.) Galrland STPM Press, New York 282-283

Helmholtz Treatise on Physiological Optics, 1866

Helson, H.; Jeffers, V. B. Fundamental problems in color vision. II. Hue, lightness, and saturation of selective samples in chromatic illumination. Journal of Experimental Psychology, 26,1-27.

Helson, H. (1943). Some factors and implications of color constancy. Journal of the Optical Society of America, 33, 555-567.

Henderson, S.T.(1970). Daylight and its spectrum. American Elsevier Publishing Company, New York.

Hering, E. Outlines of a theory of the light sense ,trans. by Leo Hurvich and Dorothea Jameson, Harvard Univ Press, Cambridge, 1964

Hochberg, Julian (1988). Visual Perception. In Stevens' handbook of Experimental Psychology, Vol. 1, Perception and Motivation. (ed. R. C. Atkinson, Herrnstein, R. J., Lindzey, G. and Luce, R. D.). John Wiley & Sons, New York.

Hubel, D.H. (1988). Eye, Brain and Vision. Scientific American, New York.

Hurlbert, Anya.

Colour vision: Is colour constancy real?

Current Biology, 1999 Jul-Aug, v9 (n15):R558-R561

Hurlbert, Anya (1986). Formal connections between lightness algorithms. Journal of the Optical Society of America. A, 3 (10):1684-1693

Hurvich, L.M. (1981). Color Vision, Sinauer Associates Inc., Sunderland, MA

IArbus, A. L. (1967). Eye Movements and Vision, by Alfred L. Yarbus. Translated from Russian by Basil Haigh. Translation editor: Lorrin A. Riggs. Plenum Press, New York.

Jameson, D. and Hurvich, L. M. (1959). Perceived color and its dependence on focal, surrounding, and preceding stimulus variables. J. Opt. Soc. Am., 51, 46-53.

Judd, D. B., MacAdam, D. L. and Wysezcki, G. (1964). Spectral distribution of typical daylight as a function of correlated color temperature. J. Opt. Soc Am. 54, 1031-1040.

Kaiser, P.K. & Boynton, R.M. (1996). Human Color Vision, 2nd edition, 1996, OSA: Washington

Kinney, J. A. S. (1962). Factors affecting induced color. Vision Res., 2, 503-525.

Kirschmann, A (1892). Some effects of contrast. American Journal of Psychology, 4, 542-557.

Krinov, E.L. (1947). Spectral Reflectance Properties of Natural Formations. Technical Translation, National Research Council of Canada TT:439, 1947

Land, Edwin H.; et al. (1983). Colour-generating interactions across the corpus callosum. Nature, 303, 616-618.

.

Maloney, Laurence T. (1986). Evaluation of linear models of surface spectral reflectance with small numbers of parameters. Journal of the Optical Society of America A, 1986 Oct, v3 (n10):1673-1683.

Maloney, L. T.; Wandell, Brian A. Color constancy: A method for recovering surface spectral reflectance. Journal of the Optical Society of America A, 3, 29-33.

Marr, D. (1982) Vision : a computational investigation into the human representation and processing of visual information, W. H. Freeman, San Francisco.

Mausfeld, R. (1998) Color Perception: From Grassman Codes to Dual Code for Object and Illuminayion Colors. In Color Vision – Perspectives from Different Disciplines, Walter de Gruyter & Co., Berlin

McCann. J.J. (1989). The role of simple nonlinear operations in modeling human lightness and color sensations. SPIE v. 10777: 355-363.

Mullen, K.T. (1985). The contrast sensitivity of human colour vision to red-green and blue-yellow chromatic gratings. Journal of Physiology, 359: 381-400.

Nagle, M.G. and Osorio, D. (1993). The tuning of human photopigments may minimize red-green chromatic signals in natural conditions.. Proc. Roy. Soc. Lond. B, 252, 209-213.

Olson, Conrad X.; Boynton, Robert M. (1984). Dichoptic metacontrast masking reveals a central basis for monoptic chromatic induction. Perception & Psychophysics, v35 (n4):295-300.

Polyak, S. L. ((1957). The vertebrate visual system, edited by H. Kluver, University of Chicago Press, Chicago.

Ruderman, D.L.; Cronin, T.W.; Chuan-Chin Chiao (1998). Statistics of cone responses to natural images: implications for visual coding. Journal of the Optical Society of America A 15 (8):2036-45

Shepard, Roger N. (1992). The perceptual organization of colors: An adaptation to regularities of the terrestrial world? IN: The adapted mind: Evolutionary psychology and the generation of culture. Jerome H. Barkow, Ed; Leda Cosmides, Ed; et al. Oxford University Press, New York, NY, USA. pp. 495-532

Snodderly, D. M. (1979). Visual Discriminations Encountered in Food Foraging By a Neotropical Primate: Implications for the Evolution of Color Vision. In The Behavioral Significance of Color (E.H. Burtt, Jr., ed.) Garland STPM Press, New York. Ch. 6, 237-279

Thompson, P.; Latchford, G. (1986). Colour-contingent after-effects are really wavelength-contingent. Nature, v320 (n6062):525-526.

Tootell RB; Silverman MS; Hamilton SL; De Valois RL; Switkes E. (1988). Functional anatomy of macaque striate cortex. III. Color. Journal of Neuroscience, 8(5):1569-93.

von Uexküll, J. 1909. Umwelt und Innenwelt der Tiere. Springer-Verlag. Excerpted in Foundations of comparative ethology . Series title: Benchmark papers in behavior ; v. 19. ed. G. Burghardt (Van Nostrand Reinhold, 1985)

Vollmer, G. (1094). Mesocosm and objective knowledge. IN Concepts and Approaches in Evolutionary Epistemology (F.M. Wuketits ed.), 69-121.

Wallach, H. (1948). Brightness constancy and the nature of achromatic colors. J. exp. Psychology,, 38: 310-324.

Walreven, J., Enroth-Cugell, C., Hood, D.C., MacLeod, D.I.A., and Schnapf, J.L. (1990). The Control of Visual Sensitivity. In Visual Perception: The Neurophysiological Foundations, edited by Lothar Spillmann and John S. Werner. San Diego : Academic Press.

Wandell, Brian A. (1995). Foundations of vision . Sinauer, Sunderland, Mass.

Webster, Michael A.; Mollon, J. D. (1997). Adaptation and the color statistics of natural images. Vision Research, 37 (23):3283-3298.

Webster, W. R.; Day, R. H.; Willenberg, K. (1988). Orientation-contingent color aftereffects are determined by real color, not induced color. Perception & Psychophysics, 44, 43-49.

Whittle, P. and Challands, P.D.C. (1969). The effect of background luminance on the brightness of flashes. Vision Research, 9, 1095-1110.

Wysezcki, G, and Stiles, W.S. (1982). Color Science: Concepts and Methods, Quantitative Data and Formulas, 2nd edition. Wiley, New York.

Zeki S. (1983). Colour coding in the cerebral cortex: the responses of wavelength-selective and colour-coded cells in monkey visual cortex to changes in wavelength composition. Neuroscience, 9(4):767-81.

Zeki, S. (1993). A Vision of the Brain. Blackwell Scientific Publications, Oxford.