C. L. HARDIN

THE VIRTUES OF ILLUSION*

1. THE PROBLEM

If we ask why color vision evolved, we are really asking two questions. First, what advantages did animals gain by being able to detect, extract and exploit wavelength information in their environment? Second, what advantages did animals, particularly human beings, gain by representing wavelength information as *color*, that is, in the form of hue qualities? These are two distinct questions, although they are often confounded. Colors are not wavelengths, spectral reflectances, or any of the physical characteristics of the world that cause human beings to have color perceptions. The central reason is this: For human beings and a variety of other ·animals, there are exactly four perceptually basic unitary hues, of which all other hues are composed in pairwise fashion, but there is no such thing as a privileged set of exactly four particular basic wavelengths or spectral reflectances of which the other wavelengths or spectral reflectances are composed. The perceptual structure of.colors thus has no counterpart in the domain of wavelengths of light, even though we normally see those colors because we are stimulated by light that has an appropriate wavelength configuration [Hardin, 1988].

But if this is so, one who attempts to answer the second question is confronted with a puzzle: if colors have causes but no counterparts in the physical world, how could it be advantageous for animals like us to represent the world to ourselves as colored? If red, yellow, green, and blue are creatures of our brains, how do they give us knowledge of the world? Or, at the very least, how do they help us to get along in the world? Colors are surely not just ornamental baubles, installed in us by Providence to soften the harsh realities of life and give employment to painters, interior decorators and crayon manufacturers. We must explain how it is that, paradoxically, configuring an animal so that it represents the world as being composed of colored objects — even

though it isn't — gives the deluded creature an advantage in coping with its environment.

The explanation that I would like to put forward for your consideration is this: By representing the reflectances in the environment by means of basic visual qualities analogous to what we humans see as red, yellow, green, or blue, or pairwise combinations of these, animals like ourselves that have restricted information-processing capabilities are better able to go about the business of recognizing objects and receiving signals from other organisms. To justify this claim, it will be necessary to tell a tale of evolution and ecology. The tale is, at the present state of knowledge, decidedly speculative, but I would urge that it is at least a plausible, and perhaps even a likely story.

Let us begin the story by considering an animal that perceives the world as without hue, that is, *achromatically*. We shall first see how achromatic vision is physically realized, and then ask what an animal that already has achromatic vision gains by adding color — *chromatic* — vision[1] to its perceptual repertoire. Following the discussion in Jacobs [1981], we may consider the potential advantages of chromatic vision under three headings: object detection, object recognition, and the reception of signals.

2. ENGINEERING A WAVELENGTH DETECTOR

To detect an object visually is to distinguish it from its background. To register the physical differences of lightness and darkness between object and background, an animal requires at least two visual receptor cells, the one to register the energy coming from one side of the object-background boundary, the other to register the energy coming from the other side, along with some means of comparing the output from the two cells. Visual receptors such as rods or cones always respond to a whole range of wavelengths, but, they do so *univariantly*, so that a few photons at a wavelength to which the cell is maximally sensitive will elicit the same response as a large number of photons at a wavelength for which it is less sensitive. This means that a population of receptor cells with the same spectral sensitivity will not generally be able to distinguish the intensity of a particular light signal (as measured by the number of photons) from its spectral composition. If these two types of

information are to be disentangled, our two-cell boundary detector will have to be replaced by a system consisting of two pairs of cells. To distinguish wavelength differences from intensity differences, the members of each pair must differ from each other in spectral sensitivity, and there must be a mechanism for comparing the outputs of the members of each pair with each other. To detect the differences in intensity across the visual border, there must also be a different mechanism for comparing the output of the first pair with the output of the second pair.

It is inevitable that an increase in the dimensions of an animal's response involves an increase in the size and complexity of its underlying neural machinery. An animal can gain a capability of finer spectral discrimination by increasing the variety of receptoral spectral sensitivities available in a given retinal area as well as decreasing the bandwidth of wavelengths to which each is tuned. But each receptor takes up space in the retina, and retinal space is always at a premium: space that is devoted to accommodating the variety of paired receptoral types required for improved wavelength analysis of a small patch of light is space that is not available for finer resolution of the spatial image, for this requires subdivision of that patch by increasing the density of receptors of the same type. In turn, image resolution and spectral analysis both require *comparisons* of the outputs of different receptors, so their requirements are antagonistic to the conditions for high sensitivity to light of low intensity, which is obviously best served by *pooling* receptoral outputs. Every animal eye therefore represents a compromise between the desiderata of sensitivity, acuity, and wavelength discrimination. The nature of the compromise depends upon the nature of the animal and the character of its ecology.

A widely favored compromise between the demands of acuity on the one hand and wavelength resolution on the other, the one adopted by our own species and many other vertebrates, is to use three photopigments with overlapping spectral sensitivities. Why three rather than two? Following a suggestion by Lythgoe [1979], let us consider a fish, swimming in water which has a peak transmittance at a particular wavelength. The transmission efficiency declines at both longer and shorter wavelengths. To take maximum advantage of the available light, it would behoove the fish to have a photopigment whose spectral

efficiency matches those transmission characteristics. This would enable the fish to detect a target that is darker than the background. And, indeed, those species of teleost fish whose vision has been most carefully studied do have cone photopigments that effect such an environmental match.

It would be a further advantage to our fish to have two other photopigments offset with respect to the first, one responding best to longer wavelengths, one responding best to shorter ones. This would facilitate detection of predators and prey that reflect more light than their backgrounds at wavelengths either higher or lower than the wavelength at which light is most easily transmitted through the medium. In the absence of such detectors, the reduced sensitivity of the single photopigment at those offset wavelengths would counterbalance the luminosity of the targets, rendering them difficult to distinguish from the background.[2] Whether a particular species sports all three classes of photopigments or only two, one tuned to and one offset from the prevailing spectrum, depends on the prevailing target reflectances. In either case, the development of multiple photopigments and multiple receptor types supplies the *first* of the necessary conditions for color vision.

The comparison of information between the responses of receptor types that differ in spectral sensitivity is the *second* necessary condition for color vision. A frequently encountered configuration of visual systems is to have the outputs of their photoreceptor classes summed and differenced to yield two chief types of postreceptoral channels. Roughly speaking, the *summed* outputs yield an *achromatic* channel that reads variations in lightness level. Signals that bear information about the *difference* in response of different photoreceptor types subserving the same retinal region are conveyed on opponent *chromatic* channels. On a chromatic channel, the differences are coded as deviations — either excitations or inhibitions — from a base rate which tracks the average light level. Just as the individual receptors lose wavelength information because of their univariant response, higher-order visual cells lose information about absolute signalling rates of receptors. Instead, they preserve differences in the signalling rates of the receptors of different types that subserve the same visual area (the information required for wavelength discrimination), as well as differ-

ences between the signalling rates of cells subserving adjoining areas
(the information required for spatial discrimination).

We have already seen that at the receptoral level, many visual
systems, including ours, trade off a measure of spectral discrimination
in order to gain increased spatial resolution. What is gained by this
second, post-receptoral loss of information about the absolute signalling
rates of receptors? The answer becomes plain if we regard this arrange-
ment as a solution to the problem of the *optic bottleneck*. In human
beings, for example, information from 120 million receptors must be
conveyed to the brain by an optic nerve of 1.2 million fibers. At first
blush, this looks like a tenfold information loss. But in fact, much of
that information is redundant, and it is possible to recode it so as to
minimize the loss of the portion of the information that is biologically
significant. Nature has solved this problem of limited channel capacity
by coding chromatic information only as differences, and piggybacking
it on an existing achromatic channel which was already designed to
carry out the function of conveying high-resolution spatial and temporal
information and wide-ranging lightness contrasts. It seems plausible,
then, that opponency came about in the first instance as an efficient
coding expedient to add wavelength information to an effective achro-
matic system that was already in place without requiring an unwieldy
increase in channel capacity [Lennie, 1984].

It is amusing to remind ourselves of what happened when, in the
1940s, engineers tackled the problem of designing a color television
system that would not render obsolete a large installed base of black-
and-white sets. The problem they faced was how to add color informa-
tion to an existing black-and-white signal without degrading that signal
or enlarging the transmission bandwidth to an unacceptable size. Their
solution was, of course, to let the black-and-white signal continue to
handle resolution and lightness contrast, and to "piggyback" pure color-
difference information onto the achromatic signal. Long ago, nature had
solved the analogous problem in an analogous way.

3. OBJECT RECOGNITION

We have so far been concerned to describe the development of color
vision as it contributes to the *detection* of targets, and this is surely the

initial reason why color vision gained a foothold in the animal world. But, as Jacobs [1981] points out, the advantage of chromatic over achromatic vision in detecting targets is hardly overwhelming: an object that can be picked out from its background in a color photograph is almost always distinguishable, although sometimes not as readily so, when the photograph is in black-and-white. The benefits of wavelength sensitivity for visibility are doubtless real, but they are also comparatively marginal.[3] Color comes into its own in the *recognition* of objects and in conveying specific biological information.

In order to appreciate properly the benefits that color brings to object recognition tasks, it is important to bear in mind the distinction between a target's *visibility* and its *accessibility by visual search*. In one sense of the term, a person's face may be fully visible in a crowd, but only if it is pointed out to you. The presence of so many other faces distracts your gaze, and makes it hard to locate the face of interest. Now let the person in question paint her face green, and your task will be made simpler by far, provided, of course, that the others have not followed suit. Contrast is, of course, one of the operative principles here, but it is by no means the only one. Human factors research has thrown light on what some of these other factors are: Just rendering a visual display in color rather than black and white shortens the time required to find a particular object in the display by as much as one-third. Giving the object a distinctive shape or other achromatic attribute shortens search time less than giving it a distinctive color. The same holds for accuracy: "The gain in accuracy with the use of colors can be at least 176% better than size, 32% better than brightness, and 202% better than shape" [Christ, 1975]. If the task is to find a particular colored object among a number of colored distractors, and if the target object is *categorically* different in color from the distractors, the time required to find the target is independent of the number of distractors [Nagy and Sanchez, 1988]. So color is a powerful aid to object recognition, and categorical color difference seems particularly helpful. Let us now take a closer look at what color categories are, what is known of their role in perception, and what their biological basis might be.

4. CATEGORIES AND FOCAL COLORS

As they are found in human beings, color categories are clusters of resembling colors in a standard color space that center around salient prototypical or "focal" colors. Color categories have two notable features. First, they involve qualitative differences that give rise to boundaries between categories. Color sample A might be k just-notice-able differences from B, and B be k jnd's from C, and yet A and B be seen to be in one color category and C in another. For instance, high-Value (lightness) and high Chroma (saturation) Munsell color samples of Hue 5Y, 5GY, and 5G are approximately equispaced perceptually, but normal observers will see the first as predominantly yellow, and the second and third as predominantly green. Second, within each color category some samples will be seen as "best" (or "focal" or "prototypical") examples of the category, and others will be seen as "poorer" examples. The "best" examples will be found in regions of the color solid in which quality changes from one sample to the next are perceived to be least, whereas the "poorer" examples are in those regions in which the quality changes are perceived to be greatest.

In what follows, it is important to bear in mind the distinction between *color* categories and *hue* categories. According to generally accepted estimates, there are about ten million distinguishable colors, but there are only four psychologically basic hues and about 160 distinguishable hues. The focal *colors* red, yellow, green, and blue are of unitary hue, or very nearly so, and of high saturation. They exhibit their hue constituents of the same name most pronouncedly. Human chromatic color categories are thus, in the first instance, founded on human hue categories, although lightness and saturation have a very important role to play. So, for instance, a crimson and a pink will both be of red hue, although they are quite different colors and, indeed, fall into different color categories.

There are eleven prototype-centered color categories [Rosch, 1975].[4] Three of them — black, white, and (middle) gray — are of course achromatic. The eight chromatic clusters cover limited regions of color space, and, with the exception of brown and pink, are centered in regions of high saturation. Four of the chromatic focal colors — red, yellow, green, and blue — are particularly salient, because their hues

are perceptually unmixed and their chromatic component strong. The remaining high-saturation focal colors are orange, and purple. Orange and purple are mixtures in which the two hue components are strong and approximately equal.[5] With pink and brown, there is an interaction between the chromatic and achromatic systems: low-lightness focal red when whitened and lightened appears as pink,[6] and high-lightness focal orange when blackened and darkened appears as brown [Boynton, 1987].

The salience of these eight chromatic and three achromatic focal colors when compared with all of the other available colors is shown in a number of ways: they are most readily picked out of piles of color chips by young children, and they are more easily remembered and more quickly named by both adults and children. In languages with well-developed color vocabularies their names are shorter, more frequently used, and more consistently and reliably applied by native informants. These eleven prototypical colors are the focal referents of the "basic color terms" that Berlin and Kay investigated in their pathbreaking study of semantic universals [Berlin and Kay, 1967]. They are also the earliest and some of the best examples of natural prototypes studied by Eleanor Rosch [Harnad, 1987]. The properties of color categories are the subjects of continuing research, particularly by Robert Boynton and his associates.[7]

So in visual search and recognition tasks — at least for human beings — colors add speed, accuracy, and memorability, and focal colors seem to serve better for each of these tasks than non-focal colors. What gives colors these desirable properties is a nice balance of qualitative diversity and simplicity. Red, yellow, green, and blue hues are qualitatively distinct from each other; there is no quantitative variation that carries one into the other. But the diversity is sharply limited; there are only four such hue qualities of which one needs to keep track. The set of evolutionary engineering compromises that gave us three photopigments and a way around the "optic bottleneck" also provided us with the elegant four-part chromatic system. The fourfold hue classification is founded upon the functional structure of the color vision system, and in turn provides the foundation for the prototype-centered color categories.

5. BIOLOGICAL SIGNALLING

The workings of the visual system provide us with an account of the categorical scheme of color perception, and that scheme proves to be particularly useful in visual search and identification. What now of the third potential advantage of color vision, the ability to receive specific bits of biological information? There is a very large body of evidence that among particular species of animals in particular situations, specific colors induce specific behavioral patterns [Burtt, 1979; Humphrey, 1972]. One could cite many examples, but there is little need to do so. The pervasiveness and biological importance of color signals is scarcely to be doubted. To understand the role that a simple system of color categories can play in the successful reception of signals, let us contrast, in the human case, color discrimination with color identification and signalling.

It has been estimated that under optimal conditions human beings can *distinguish* about ten million colors. But if they are asked to *identify* colors in the absence of a reference standard, even the experts cannot reliably do so for more than 30 or 40 colors. In a recent experiment in Robert Boynton's laboratory, subjects were shown, a few at a time, a total of 424 distinct colors. When asked, midway in the experiment, to estimate the number of distinct colors they had seen, some subjects guessed that they had seen as few as 20 different colors. In the reliable identification of signal lights, even smaller numbers of colors are involved. Distinguishing lights rapidly and without error is obviously essential for train crews and users of control panels in aircraft and nuclear power plants. The maximum number of distinct lights that may safely be used in such crucial situations is five [Hardin, 1988].; [Boynton, 1987]; [Boynton and Olson, 1987]. For surface colors, six may be employed.

In cases like these in which the number of acknowledged colors ranges from ten million in some circumstances to, say, 30 in others, large numbers of colors that could be distinguished when set side-by-side are, on the occasions in which recognition tasks must be performed, lumped together into a relatively small number of equivalence classes. Such tolerance of variations in coloration is a desirable attri-

bute, but if the signal is not to be misidentified, there must be a gap between one equivalence class and the next. For traffic signs, one must not, for instance, use those oranges that might, at a quick glance, be confused with reds on the one hand or yellows on the other. In natural settings, the displaying animal or plant must employ a range of colors that will be easily recognizable by the intended receiver and not readily confused with other colors attaching to similar shapes in the environment.

There is a further advantage to an animal's having categorical color perception. Even though its visual system does not adapt perfectly to the prevailing illumination, such an animal will tend to respond to recognition tasks as if it perceived surface colors with an extremely high degree of constancy. The relative coarseness of the categorical scheme involved in the recognition of surface colors renders the system tolerant of minor variations in illumination, just as it discounts normal variations in object spectral reflectances and in the color perceptions among individuals of the same species.

6. CONCLUSION

The moral is plain enough. For purposes of distinguishing an object from surroundings of similar lightness, one needs to have a receptive system that is sensitive to small wavelength differences. For signal-reception and identification purposes, one needs to have a receptive system engendering a few basic categories that can ignore minor variations and lump stimuli together into a small number of salient, memorable equivalence classes. With elegant economy, our visual systems does both by drawing on the same neural resources. The system of qualitative classification that this involves need not match any analogous set of structures outside the organism in order to provide real advantages to the animal that uses it. The hues that we human beings see express our system of coding wavelength information rather than some set of properties of reflecting surfaces. But the form of the coding is not just a bit of non-functional adornment freeloading on the serious business of visual information processing. We must see it, rather, as supplying the means by which a rich amount of sensory information can be rapidly and efficiently represented by cognitive

machinery of limited capacity [Miller, 1957]. With color, the medium is the message.

NOTES

* This is a revised version of Hardin [1990], with permission of SPIE.
[1] To call this "color" vision is not to presuppose that the animal experiences qualities that we would, were we to experience them, call "colors". It is to suggest that what they experience is functionally analogous to what we experience. For an intriguing discussion of chromatic ecology and alien color qualities, see Thompson et al. [1992].
[2] It is interesting that the maximal sensitivity of human vision is to wavelengths that correspond pretty nearly to the average spectrum of the ambient light near a forest floor.
[3] Cf. Jacobs [1981]. However, the differences are quite important if the task is to pick out a target of relatively uniform spectral reflectance against dappled or variegated surroundings, as in a forest with splashes of sunlight coming through the vegetation. See Mollon [1989].
[4] Whether there may be more than eleven is at present controversial, but this seems to be the maximum number represented by basic color terms in any current natural language.
[5] Cf. Hardin [1988]. I now regard my discussion of color categories and basic color terms in this book to be flawed, although not in respects relevant to the present paper.
[6] However, the best pinks are somewhat blueish.
[7] See, for instance, Boynton and Olson [1987] and its accompanying bibliography.

REFERENCES

Berlin, B. and Kay, P.: 1967, Basic Color Terms (University of California Press, Berkeley and Los Angeles).

Boynton, R. M.: 1987, 'Evaluating Color Appearance Using the OSA Uniform Color Scales Samples and a Color-Naming Method', Color Appearance, Optical Society of America Technical Digest Series, 15.

Boynton, R. M. and Olson, C. X.: 1987, 'Locating Basic Colors in the OSA Space', Color Research and Application, v. 12, no. 2, 94—105.

Burtt, E. H. (ed.) : 1979, The Behavioral Significance of Color (Garland STPM Press, New York and London).

Christ, R. E.: 1975, 'Review and Analysis of Color Coding Research for Visual Displays', Human Factors, 17, 542—570.

Hardin, C. L.: 1988, Color for Philosophers: Unweaving the Rainbow (Hackett Publishing Company, Indianapolis).

Hardin, C. L.: 1990, 'Why Color?' in Brill, M. (ed.) : 1990, Perceiving, Measuring and Using Color, The Society of Photo-optical Instrumentation Engineers, 1250, 293—300.

Harnad, S. (ed.) : 1987, Categorical Perception (Cambridge University Press, Cambridge and New York).

Humphrey, N. K.: 1972, '"Interest" and "Pleasure": Two Determinants of a Monkey's Visual Preferences', Perception 1, 395—416.

Jacobs, G.: 1981, Comparative Color Vision (Academic Press, New York), ch. 6.

Lennie, P.: 1984, 'Recent Developments in the Physiology of Color Vision', Trends in Neurosciences, v. 7, n. 7, 243—249.

Lythgoe, J. N.: 1979, The Ecology of Vision (Oxford University Press, New York and London).

Miller, G. A.: 1957, 'The Magic Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information', Psychological Review, 63, 81—97.

Mollon, J. D.: 1989, '"Tho' She Kneel'd In That Place Where They Grew . . .": The Uses and Origins of Primate Colour Vision', Journal of Experimental Biology, 146, 21—38.

Nagy, A. and Sanchez, R.: 1988, 'Color Difference Required for Parallel Visual Search', Optical Society of America Technical Digest Series 11, 42.

Rosch, E.: 1975, 'The Nature of Mental Codes for Color Categories', Journal of Experimental Psychology: Human Perception and performance 1, 303—322.

Thompson, E., Palacios, A. and Varela, F. J.: 1992, 'Ways of Coloring: Comparative Vision as a Case Study for Cognitive Science', Behavioral and Brain Sciences, v. 15, n. 1, 1—26.

Department of Philosophy
Syracuse University
Syracuse, NY 13244
USA