

Ways of coloring: Comparative color vision as a case study for cognitive science

Evan Thompson

Department of Philosophy, University of Toronto, Toronto, Ontario, Canada M5S1A1

Adrian Palacios

Department of Biology, Yale University, New Haven, CT 06511; Institut des Neurosciences (CNRS-Paris VI), 75005 Paris, France

Electronic mail: apalac@yalevm.bitnet

Francisco J. Varela

Institut des Neurosciences (CNRS-Paris VI) and CREA, Ecole Polytechnique, 75005 Paris, France

Electronic mail: fv@frunip62.bitnet

Abstract: Different explanations of color vision favor different philosophical positions: Computational vision is more compatible with objectivism (the color is in the object), psychophysics and neurophysiology with subjectivism (the color is in the head). Comparative research suggests that an explanation of color must be both experientialist (unlike objectivism) and ecological (unlike subjectivism). Computational vision's emphasis on optimally "recovering" prespecified features of the environment (i.e., distal properties, independent of the sensory-motor capacities of the animal) is unsatisfactory. Conceiving of visual perception instead as the visual guidance of activity in an environment that is determined largely by that very activity suggests new directions for research.

Keywords: adaptation; color vision; comparative vision; computation; ecological optics; evolution; objectivism; ontology; qualia; sensory physiology; subjectivism

The study of color vision provides a microcosm of research in cognitive science: Each of the disciplines that compose cognitive science has made significant contributions to our understanding of color. Neuroscientists have had some success in uncovering the anatomical and physiological correlates of color vision in the visual system, primarily in primates (DeValois & DeValois 1975; Livingstone & Hubel 1984; Zeki 1983); cellular biologists have characterized the retinal basis of sensitivity (Brown & Wald 1964; Dartnall et al. 1983; Svaetichin & MacNichol 1958); molecular biologists have isolated and sequenced the genes for the three different types of color-sensitive photopigments in the human eye (Nathans et al. 1986); psychophysicists have contributed quantitative models for human visual performance (Hurvich 1985; Hurvich & Jameson 1957; Jameson 1985); cognitive psychologists have provided models of the structure of human color categories (Heider 1972; Rosch 1973); linguists have shown that human languages contain a limited number of "basic" color terms (Berlin & Kay 1969) and have provided models to derive these semantic universals from properties of the visual system (Kay & McDaniel 1978); researchers in computational vision and artificial intelligence have devised computational models and algorithms for color constancy (Gershon 1987; Hurlbert 1986; Land 1983; Maloney 1985; Maloney & Wandell

1986); and finally, philosophers have discussed the ontological status of color and its implications for theories of mind and perception (Hardin 1988; Hilbert 1987; Thompson 1989).

This target article is intended as a contribution to this ongoing interdisciplinary effort. We propose to offer here a new empirical and philosophical perspective on color vision, one based on recent experimental research in *comparative* color vision – studies of color vision in various animal species. We do not intend to provide a detailed scientific review of current research on this topic (see Goldsmith, 1990; Jacobs 1981; and Nuboer 1986). Rather, we wish to draw on this material, especially recent research on fishes, birds, and insects, to cast new light on some fundamental questions in visual perception, cognitive science, and the philosophy of mind.

Our presentation has three stages. In the first, we provide an overview of various types of explanation for color vision in contemporary visual science, showing how particular types of explanation have been used to motivate various views about what color is, that is, about the ontology of color. As we shall see, those who favor objectivism about color, the view that colors are perceiver-independent physical properties (Hilbert 1987; Matthen 1988), rely on computational vision, whereas those who favor subjectivism, the view that colors are internal sen-

sory qualities (Hardin 1988), rely on psychophysics and neurophysiology. In the second stage, we propose a broader comparative and evolutionary perspective on color vision. We present what we call "the comparative argument," which purports to show that an adequate account of color must be *experientialist* (unlike objectivism) and *ecological* (unlike subjectivism). In the third stage, we explore the implications of the comparative argument for vision research. We argue that the typical emphasis in computational vision on optimally "recovering" prespecified features of the environment (i.e., distal properties whose specification is thought to be independent of the sensory-motor capacities of the animal) is unsatisfactory. Instead, visual perception is better conceived as the visual guidance of activity in an environment that is constituted largely by that very activity. Here we present what we call an "enactive" approach to perception (proposed originally by Varela 1979; 1984; 1989; 1991a; and developed subsequently by Varela et al. 1991b). We then suggest some directions for further research that follow from our discussion.

1. Explanation in visual science and the ontology of color

1.1. Levels of explanation: A brief overview

A central concern in contemporary visual science (indeed throughout all cognitive science) is the relation among various levels of generalization and explanation. Following Churchland and Sejnowski (1988), we can distinguish several notions of "level" at work in cognitive science: levels of analysis, of organization, and of operation ("processing"). Because these notions will prove to be of use in our discussion of color vision, we review them briefly here.

In vision research, the notion of levels of analysis is most familiar from the work of Marr and Poggio (1977). In their framework, vision requires analysis and explanation at three different levels: (i) the level of computational theory; (ii) the level of algorithm; and (iii) the level of physical implementation. The computational level is an abstract analysis of the problem or task, which for early vision, according to Marr and Poggio, is the recovery of three-dimensional scenes from ambiguous two-dimensional projections, otherwise known as "inverse optics" (Marr 1982; Poggio et al. 1985). For color vision, the inverse optics problem is to recover the invariant surface spectral reflectances of objects in a scene. The algorithmic level is concerned with the specific formal procedures required to perform a given computational task. Finally, the level of physical implementation is concerned with how the algorithms are physically realized in biological or artificial systems.

It is well known that Marr (1982) claimed that these three levels of analysis were largely independent. In the study of biological vision, Marr also supposed that the algorithmic level corresponds to psychophysics and to parts of neurophysiology, whereas the implementational level corresponds to most of neurophysiology and neuroanatomy (1982, p. 26). This conception of explanation in visual science, especially as applied to the study of natural vision, has generated considerable discussion and debate. Among other things, many dispute Marr's (1982)

claim that the three levels of analysis are largely independent. Some favor a more "bottom up" approach to the explanation of visual processes, and some criticize Marr's assumption of optimality at the computational level, that is, that "what is being computed is optimal in some sense or is guaranteed to function correctly" (1982, p. 19) [see also Schoemaker, "The Quest for Optimality: A Positive Heuristic of Science?" *BBS* 14(2) 1991; and Anderson, "Is Human Cognition Adaptive?" *BBS* 14(3) 1991.] We do not intend to review all of these controversies here.¹ We mention them, rather, as pointers toward some of the issues that will arise shortly when we discuss models of color vision, and when we present our alternative "enactive" approach to visual perception in section 3.

In contrast to the notion of levels of analysis, the notion of levels of organization is relatively straightforward. In the nervous system, we find highly organized structures at many different scales from molecules to synapses, neurons, neuronal ensembles, neural networks, maps, systems, and so on. Each level has properties specific to it, which in turn require different techniques for their investigation. Such organizational complexity is certainly evident in color vision, ranging from the chemical properties of receptor photopigments to the network properties of retinal and cortical cells.

Finally, in addition to these levels of organization, we find many levels of operation in the nervous system. How these levels are to be assigned, however, is considerably less clear than it is for levels of organization. The typical procedure is to order the levels hierarchically from peripheral (lower) to central (higher) areas (measured in terms of synaptic distance from sensory stimulation), thereby suggesting that "processing" in the nervous system proceeds sequentially. We wish, however, to dissociate the notion of levels of operation from the idea that processing among the levels is sequential. If (as we and many others believe) "higher" levels can significantly affect the processing in "lower" levels, then the notion of sequential processing will be of limited application, or at least will have to be modified considerably. To cite just one example that is relevant for our discussion here: Although the visual system is typically described as carrying out sequential processing from retina to lateral geniculate nucleus (LGN) to visual cortex, it is also well known that there are massive back-projections from all areas of the cortex to the thalamic nuclei (Steriade & Deschenes 1985). In the case of the visual system, there are actually more fibers going down from the visual cortex to the LGN than go in the reverse direction (Robson 1983). This organization suggests that neuronal activity in central levels may considerably modulate the activity at peripheral levels, an idea that is also supported by some recent experiments (e.g., Varela & Singer 1987). We set this issue aside here. However the relations among levels of operation must ultimately be conceptualized, it is obvious that there are various levels to be distinguished. For example, in primate color vision, we need to understand at the very least the two-way interactions between operations in the retina, thalamus, striate (V1) and peristriate (V4) visual cortex.

With these three notions of "level" in hand we can now turn specifically to color vision. In the remainder of section 1 we give a brief overview of the types of explanation offered for color vision, showing how they have been

used to motivate contrasting philosophical positions on the ontology of color.

1.2. Color space: Psychophysics and neurophysiology

In general, psychophysics and neurophysiology have taken as their point of departure what is known as "color space." This is the closed space formed by the three semi-independent dimensions of color known as hue, chroma or saturation, and value or brightness (Figure 1).² Hue obviously refers to the redness, greenness, yellowness, or blueness of a given color. Saturation refers to the proportion of hue in a given color relative to the achromatic (white-black) dimension: Saturated colors have a comparatively greater degree of hue, whereas desaturated colors are comparatively closer to gray. Brightness refers

to the achromatic or white-black dimension. In this space, colors can be seen to exhibit certain relations among themselves, such as the hue-opponency of red and green, and yellow and blue. These kind of relations compose what we call the *phenomenal structure of color space*, or simply the *phenomenal structure of color*.

How are we to explain the generation of this phenomenal structure? Why does (our) color space have this phenomenal structure and not some other? It is primarily this question that has motivated the psychophysical and neurophysiological study of color vision. Rather than review this enormous field, we present merely a few points that are relevant for our purposes in this paper.³ The basic idea is to provide a mapping from the phenomenal color space of Figure 1 into a new color space whose coordinates correspond to psychophysical and/or neurophysiological processes relevant for color vision. We call the axes of these new color spaces "color channels." Strictly speaking, channels are specified psychophysically and so are not isomorphic with unique neuronal pathways (Boynton 1988, p. 77), but we intend to use the term "channel" both in this psychophysical sense and somewhat more loosely to refer to underlying neurophysiological processes (such as color opponent receptive field properties) that can be studied at various levels of analysis, organization, and operation.

We should note that the following maps of color space are idealized. We do not intend to suggest that they provide full-fledged "linking propositions" needed to identify chromatic perceptual states and states of the visual substrate (see Teller 1984; 1990; Teller & Pugh 1983). Visual science is still far from being able to provide the full story of how the activity in multiple neuronal areas becomes integrated to form our experience of color. Our intention, however, is simply to provide some illustrations of the kinds of covariance that have been established between aspects of the phenomenal and the biological.

We begin with the three kinds of retinal cones, short-wave ($S^1 = S$), medium-wave ($S^2 = M$), and long-wave ($S^3 = L$), which respond with a differential sensitivity to wavelength according to the photopigment they carry in their outer segments (Figure 2a, bottom). At this level, we can construct a rudimentary map of color space whose coordinates correspond to the relative activity of the cones, which are present (in various proportions) at each point of the visual field. This map corresponds to a vector $\vec{S}_r = \langle S^1, S^2, S^3 \rangle / \langle \vec{r} \rangle$, where \vec{r} is the surface coordinate. A convenient representation is a (Maxwell) triangle (Figure 2b) instead of a 3-D graph to depict the spectral loci of monochromatic lights: The three kinds of cone receptor appropriately adjusted in activation are required to match a test-light of any spectral composition and intensity. This property corresponds to the *trichromacy* of normal human color vision.

We refer to this mapping of color space as "rudimentary" because it takes into account only the relative absorptions of the cone photopigments. Postreceptor cells, however, both combine and compare (subtract) cone signals, thereby giving rise to three new types of color channels: two opponent chromatic channels (R-G, Y-B) and one nonopponent achromatic channel (Wh-Bk), which can be found in primates at the retinal and thalamic levels (Figure 2c). These new channels result, then, from




Figure 1. The phenomenal structure of human color space. (a) The three-dimensional space of hue, saturation or chroma, and brightness. We use here the standard Munsell color space. (b) A slice in color space for the purple sector (Munsell values 3P-9P), using the ISCC-NBS color names (or equivalent centroid numbers).




a**b****c**

Figure 2. Biological mappings of human color space at various levels. (a) The spectral absorption of human cone pigments and the overall photopic sensitivity curve for humans (in discontinuous lines). In the upper part a plot of the wavelength discrimination curve (reciprocal of the just noticeable difference in wavelength necessary for wavelength discrimination). Modified from Nubser (1986). (b) For any light, the relative Long/Medium/Short wave length absorptions (L, M, S) can be plotted as relative activity on a Maxwell triangle with orthogonal axes of cone length. The loci of pure spectral colors are shown calculated by normalizing the pigment spectra for equal

areas and computing the quantum catch. Equal absorption for all three areas is labelled as w . From Goldsmith (1990). (c) On the left, opponent neural system obtained by the weighted excitatory and inhibitory responses between retinal elements. On the right, a diagram showing how each of these channels can be used to provide a set of axes for color space. The Wh-Bk (White-Black) channel receives excitatory input mostly from L and M cones, whereas the two antagonistic channels receive both excitatory and inhibitory inputs: $+L - M$ for the R-G (Red-Green) channel, and $(M + L) - S$ for the Y-B (Yellow-Blue) channel. From Hurvich (1981), and Ingling and Tsou (1977).

d




Figure 2(d) Chromaticity diagram computed from the three CIE (Commission Internationale de l'Eclairage) 1931 color matching functions. As in (b) the relative activities have been normalized so that two values suffice to locate all points in this modified triangle.

linear combinations of the receptor activations, which can be written vectorially as $\vec{c}_r = M \cdot \vec{s}_r$, $\vec{c}_r = \langle C^1, C^2, C^3 \rangle(\vec{r})$, with $C^1 = Wh-Bk$, $C^2 = R-G$, $C^3 = Y-B$. The matrix M for constructing the channels is at the core of various color vision theories (Wyszecki & Stiles 1982); we return to provide an algorithm for its determination below in section 2. These three color channels proper can be used to provide a set of axes for color space (Figure 2c). This diagram thus displays the trivariance of human color vision mapped onto the three dimensions of color space at the physiological level.

This mapping, too, has limitations, for it relies on an analogy between the existence of opponent colors and the existence of chromatically opponent cells, which though obviously promising nonetheless neglects many details of the fit between properties of color perception and the properties of these peripheral cells (Hood & Finklestein 1983; Teller 1990). It also does not take into account the multilevel neuronal interactions in the visual pathway that somehow constitute our entire experience of color (Livingstone & Hubel 1984; Zeki 1983).

To obtain axes more appropriate for this level we need to use psychophysical global response functions. The functions chosen will depend on which aspect of color experience we are interested in quantifying more precisely, for example, chromatic sensitivity, discrimination, or color mixing (see Wyszecki & Stiles 1982). For our purposes here, one useful standard set comprises the empirically determined CIE (Commission Internationale de l'Eclairage) color-matching functions $\langle x, y, z \rangle(\lambda)$, which specify the well-known human chromaticity diagram (Figure 2d). To project our three-dimensional color space into this two-dimensional plane we equalize for brightness. At the center is the achromatic neutral (white/gray) point; movements away from the point indicate an increase in saturation, with the maximally saturated hues along the periphery. In this chromaticity diagram, we find complementary colors at the opposite ends of the

space, the pure spectral locus at the outermost boundary, and the purple range as "nonspectral" loci. Thus *any* color we perceive can be matched in this space by an appropriate value of the underlying channels (global response functions). It should be made clear that although strictly speaking the CIE diagram is a *stimulus* space, the stimuli are useful in characterizing color experience by mapping its extent, and so the CIE diagram can be read as a mapping of one aspect of color experience. (This point will also apply to the relation between the stimulus spaces and color perception in other animals as discussed below).

These color-matching functions were derived mostly for purely psychophysical and industrial purposes and so are not particularly useful as a guide to underlying neuronal processes. Nevertheless, they do provide a more precise and quantitative way of mapping some aspects of the phenomenal structure of color space. It should be noted that alternative color channels motivated by computational algorithms (Land 1983) and perceptual and neurobiological data (Zeki 1980; 1983; 1985) have been proposed by taking into account the global integration properties of visual mechanisms, because a local description (i.e., independent of the rest of the visual field) violates perceptual evidence and neurophysiological data. For example, Land (1983) proposes three lightness "indicators," which result from discarding the illuminant from the receptor activity after long-range integration of local values.

We have now provided enough illustrations of the various neurophysiological and psychophysical mappings of color space for this target article. Henceforth, we will use the term "color space" to refer to this kind of composite representation in which the phenomenal structure of color and the structure of the visual system covary. For our purposes here, then, color space refers to the following multilevel description:

- (i) receptor space: the raw array of local activity under given illumination conditions in a scene;
- (ii) "lightness" indicator space: the globally integrated activity at various levels after discarding the illuminant from the receptor space via lateral interactions;
- (iii) physiological channels space: the local activity of subtraction and addition of integrated values to conform to antagonistic mechanisms obtained from indicator values;
- (iv) psychophysical channel space: the perceptual, high-level integration into separable mechanisms, obtained on the basis of the underlying physiological activity relevant to color channels;
- (v) phenomenal space: the color appearance space of hue, saturation, and brightness.

The main task for the psychophysical and neurophysiological study of color vision is to uncover the appropriate biological processes underlying all these levels and to formulate, test, and establish the "linking propositions" needed to relate the various levels (Teller 1984; 1990; Teller & Pugh 1983). As a matter of general principle, however, it is clear that the phenomenal structure of color covaries with the structure (and, as we shall see later, the ecological interactions) of the perceiver. If we wish to explore this fact, we should determine whether and how changes in the structure (and ecological interactions) of the perceiver can be correlated with changes in the phenomenal structure that color exemplifies. This is the main task of this target article.

1.3. Computational color vision

A full explanation of color vision requires that we also understand how color appearances remain relatively stable or constant in natural light environments – a phenomenon known as *color constancy*. Computational color vision is particularly concerned with this phenomenon. Because the retinal activity from a given point hopelessly confounds the illumination with the reflectance properties of surfaces, the core problem is to disentangle these variables and assign colors that correlate with surface properties.

In what follows, the ideas we present naturally take a mathematical form, which cannot be simplified without losing some important features. The nonmathematical reader can skip the formulae, however, without losing the basic points we need for our argument here. (Our presentation throughout this section is indebted to the more detailed treatment in Hurlbert 1989, ch. 3).

Most computational models simplify the overall situation by considering only the surface reflectance (or albedo), which depends on object properties, not on viewing geometry, as in [1]:

$$I(\lambda, \tilde{r}) = \rho(\lambda, \tilde{r})E(\lambda, \tilde{r}) \quad [1]$$

where I is the irradiance, λ is wavelength, E is the surface illumination, and ρ is the reflectance, and \tilde{r} , as before, is the surface coordinate. The irradiance affects an array of sensors which have a specific nonlinear response function $R^i(\lambda)$, $i = 1, \dots, n_{recep}$, comprising a number of different cone classes (e.g., for primates $n_{recep} = 3$). Under these conditions the raw receptor response corresponding to a point \tilde{r} from a surface is the integral:

$$S^i(\tilde{r}) = \int_{\text{visible window}} d\lambda R^i(\lambda) \rho(\lambda, \tilde{r}) E(\lambda, \tilde{r}) \quad [2]$$

We have written “visible window” simply to indicate the extent of the wavelength sensitivity depending on the species being considered. For primates, this “visible” range is approximately 400–700 nm; for insects it shifts down to approximately 310–590 nm; in birds it broadens to approximately 350–720 nm.

As we mentioned above, the computational approach to color constancy is a prime example of inverse optics – the recovery of what are taken to be objective attributes of three-dimensional scenes from ambiguous two-dimensional projections. In the case of color vision, the problem is to discard the source illuminant E and retain the invariant spectral reflectances ρ of object surfaces given only the retinal activity S . This problem – like inverse optics problems generally – is underscored or ill-posed. (Poggio et al. 1985). To solve an ill-posed problem one must restrict the class of admissible solutions by introducing constraints; these constraints are said to “regularize” the problem. In the case of color vision, a combination of empirical evidence and task-level analysis has shown that these constraints are basically of three kinds: (i) low-dimensional models of lights and reflectance; (ii) global computations; and (iii) spatial segmentation.

1.3.1. Low dimensionality. Naturally occurring illuminants and object reflectances can be described as lying within a low-dimensional space: A few basis functions, when added together in the correct proportions, suffice to

span the entire diversity of actual lights and reflectances. Formally expressed this becomes:

$$E(\lambda, \tilde{r}) = \sum_{j=1}^{n_{illum}} \varepsilon_j(\tilde{r}) E_j(\lambda)$$

$$\rho(\lambda, \tilde{r}) = \sum_{k=1}^{n_{reflec}} \zeta_k(\tilde{r}) \rho_k(\lambda) \quad [3]$$

where E_j and ρ_k are basis functions, and ε_j and ζ_k are spatially varying coefficients. On the basis of empirical evidence from measurements of typical ambient conditions and object reflectances, n_{illum} and n_{reflec} are usually taken to be 3 or at most 6 (Maloney 1985). It is typically further assumed that illumination is quite uniform over space,

$$E(\lambda, \tilde{r}) = E(\lambda), \quad \varepsilon_j(\tilde{r}) = \varepsilon_j \quad [4]$$

and that reflectance is invariant under changes in viewing geometry (i.e., a Lambertian reflection model). Thus, computational color vision is fundamentally constrained by the low dimensionality of *both* the stimuli and the receptor types, because these are known to come in small numbers.

Under these conditions the irradiance equation [2] takes the general form:

$$S^i(\tilde{r}) = T^{ik}(\varepsilon) \zeta_k(\tilde{r}) \quad [5]$$

where

$$T^{ik}(\varepsilon) = \sum_j \varepsilon_j \tau_{jik}, \quad \tau_{jik} = \int d\lambda R^i(\lambda) \rho_k(\lambda) E_j(\lambda) \quad [6]$$

The matrix, $T^{ik}(\varepsilon)$ (which Maloney calls the “light transformation matrix”), depends on the illuminance and reflectance basis functions, the sensory sensitivities, which are fixed, and the illuminant, which is variable. Clearly, since this matrix and the $\zeta_k(\tilde{r})$ are, in general, not entirely known, the equations are underdetermined, and to find solutions further constraints need to be introduced. These take various forms. For example, Buchsbaum (1980) requires that a weighted average of all reflectances in a given scene be known. In contrast, Maloney and Wandell (1986), and Yuille (1984), assume that there is at least one more sensor type than there are reflectance components. This assumption obviates the need for the previous ones, and exploits instead the various sensors at each location. For instance, using only photoreceptors, these algorithms would recover constant colors for materials that can be described using no more than two basis functions. Finally, D’Zmura and Lennie (1986) introduce eye movements and light adaptations to recover illuminants.

1.3.2. Global computations. The foregoing discussion focused on the quality of the light signals and the number of receptors, but it did not take into account the way the local activity of a photoreceptor, $S^i(\tilde{r})$, is not the most significant variable. More relevant to account for color constancy and chromatic induction is the interaction of receptor activity over distant places in the visual scene, which transforms luminance (a quantity which expresses a local level of activity) into “lightness,” $L^i(\tilde{r})$, a level of activity closer to reflectance (and one that is relative to

other levels of activity in the scene). These global interactions can all be understood as a manifestation of the lateral interactions and reentrant circuits typical of both the retina and parts of the visual system, which lead to internally specified values rather than raw sensory values.

There are a number of equivalent "lightness" algorithms (Hurlbert 1986). In general, though, lightness algorithms proceed by (i) taking a differential of values of the intensity over different locations of space; (ii) applying a threshold operation that eliminates small values because of smooth changes in the illumination and retains large values resulting from abrupt changes in reflectance at the borders between patches; and (iii) integrating the result of this operation back into reflectance values for each position in space. For example, one of the first algorithms, proposed by Horn (1974), obtains lightness by simulating a diffusion of the activity of one receptor over the entire layer. Mathematically, this is expressed as a solution to the Poisson equation on the receptor activity:

$$\nabla^2 L(\vec{r}) = \theta[\nabla^2 S(\vec{r})] \quad [7]$$

where $\theta[\cdot]$ is a thresholding operation performed on the Laplacian operator ∇^2 , which embodies the neural lateral interactions. When the sensor array is finite, and surrounded by a constant boundary condition, [7] can be solved explicitly. More recently, global computations have been approached by noticing that each sensory receptive field has an excitatory center and an inhibitory periphery that can be seen as a filter for the light signal. It is also known that these receptive fields have various sizes and degrees of steepness. Thus, a family of recent algorithms assumes that the sensor array is convolved through a center-periphery profile at each point of the visual scene (Land 1986) and at various scales (Hurlbert 1986). In this case:

$$L(\vec{r}) = \int_{\text{Scale range}} d\mu \int_{\text{Region}} d\vec{q} \theta[\nabla^2 G(\vec{r} - \vec{q}; \mu) S(\vec{q})] \quad [8]$$

where G is a Gaussian function, with a continuum of parameters μ . The Laplacian of G is roughly equivalent to a difference of Gaussians (DOG function), similar to receptive fields. Expression [8] degrades when approximated in the discrete case, but sums over ten scales of μ yield reasonable lightness values for so-called "Mondrian" scenes (displays consisting of about 100 different colored papers arranged arbitrarily that resemble the paintings of Piet Mondrian).

1.3.3. Spatial segmentation. Even with low-dimensional constraints and network global computations, reflectances are still underdetermined. One missing key element is the way a scene is segmented into the relevant patches on which the calculation of reflectance will be performed. Some extra assumptions about surfaces (abruptness of change, distributed averages, etc.) must therefore be brought to bear. These assumptions in part serve the purpose of color vision, which is presumably important in object discrimination and identification (D'Zmura & Lennie 1986), a point to which we return in later sections. The overall effect of segmentation is to make reflectance values correspond not to local scene coordinates \vec{r} , but to regions σ , yielding lightness values

over regions, $L_\sigma^i, \sigma \in \Sigma$. One of the better known segmentation algorithms is the one from Rubin and Richards (1982; 1988), which seeks to determine where material changes occur in a scene using only spectral intensity responses in separate points on the retina. For example, we can state this idea by considering two different receptors, which will have a spectral crosspoint on opposite sides of an edge when:

$$[S^i(\vec{r}_1) - S^i(\vec{r}_2)][S^j(\vec{r}_1) - S^j(\vec{r}_2)] < 0 \quad [9]$$

If the product is negative, one channel increases while the other decreases. Such a crosspoint will be produced only by material changes, under such simplifying assumptions as uniform illumination. In addition, one can consider the signs of the slopes of each response function (opposite slope sign condition). There are conditions in which these algorithms will not segment a scene into material discontinuities but will give false positives because of shadows, occlusions, or illuminant variations.

Another model for segmentation is provided by D'Zmura and Lennie (1986). In this model, mechanisms of light adaptation ("a multiplicative change in sensitivity in the independent cone mechanisms followed by an adaptive linear transformation of scaled cone signals at color-opponent sites" [p. 1670]), combined with eye movements that expose the eyes to the average light reflected from the field of view, are used to evaluate and discount the illuminant, thereby recovering reflectance designators. This scheme does not rely on a prior segmentation of the scene. Instead, the designators are transformed to yield estimates of hue, which is, compared to saturation and lightness, relatively independent of object shape and viewing geometry. These hue estimates can then be used in the task of segmentation.

These three elements – low-dimensional constraints, global integration, and scene segmentation – must come together for artificial systems to regularize the ill-posed problem of recovering reflectance. Since the assumptions introduced are about the natural world, they can be expected to fail when they are not satisfied in the world. For example, with a few exceptions (e.g., D'Zmura & Lennie 1986) most current computational algorithms do not perform well in the presence of significant specular components; the algorithms require a virtually uniform illumination, and the collection of surface reflectances must average to the same "gray" in every scene.

Our purpose is not to provide a comprehensive discussion of computational color vision. Enough has been said to indicate that there are at present different approaches to color vision, which focus on different respective kinds of color phenomena: On the one hand, we have computational theories of color constancy, on the other hand, psychophysical and neurophysiological investigations of a range of such chromatic phenomena as constancy, contrast, color matching, color blindness, and so on. At the present stage of research, the question of how these approaches to color vision might be related does not admit of a clear and nonpartisan answer. Only a handful of studies explore possible links among the various levels of analysis and kinds of phenomena – for example, Zeki's (1980; 1983; 1985) neurophysiological studies of the cortical mechanisms underlying color constancy, or Buchsbaum & Gottschalk's (1983) formal analysis of opponent color mechanisms. This question does raise a number of

conceptual and empirical issues, however, which in turn have considerable implications for the ontology of color.

1.4. Current ontologies

To discern these issues consider that human color vision exhibits only *approximate* color constancy. Many factors can affect constancy; among the best documented are the effects of sensitivity to the spectral quality of the illumination (Helson 1938; Helson & Jeffers 1940; Judd 1940). As Jameson and Hurvich (1989, p. 7) note in a recent review: "Departures from perfect color constancy with changes in the spectral quality of illumination . . . imply that perceived contrast between objects of different surface reflectance varies with the level and kind of illumination in which they are seen and to which the visual system is adapted."

From the standpoint of the computational level of analysis, the approximate constancy of human color vision is not surprising. For example, in Maloney & Wandell's (1986) model, surface reflectance can be completely recovered only if there are more sensor types than degrees of freedom in reflectance. Since naturally occurring reflectances require 3 to 6 degrees of freedom for their full specification (Maloney 1985) and human color vision is trichromatic (3 receptor types), Maloney & Wandell's model predicts that there are chromatic differences among naturally occurring surface reflectances that cannot be detected by a trichromatic system (assuming, of course, that no other kind of disambiguation is available).

So far, then, we have an overall agreement between psychophysics and computational vision. The problems arise when we ask how the approximate constancy of natural color vision is to be *explained*. It is in the kind of answer given to this question that we find the motivation for current views on the ontology of color.

Starting from the computational level of analysis as outlined above, we assume that the function of color vision is the achievement of color constancy, defined as the recovery of the invariant surface spectral reflectances in a scene. We then are led to explain approximate color constancy as a departure from ideal or perfect color constancy, the implication being that such a departure constitutes a visual shortcoming or error (cf. Maloney & Wandell 1986, p. 32). Obviously, once such a conceptual framework is in place, it is natural to suppose that color is simply the property of surface spectral reflectance. Thus, consider the following passage from Maloney (1985):

The analyses of Chapter 2 [those presenting finite-dimensional linear models of lights and reflectances] used data appropriate to human environments and suggested that what we call color corresponds to an objective property of physical surfaces. Depending on the lights and surfaces present in a scene, we succeed or fail in estimating these properties. Failures of color constancy, from this viewpoint, can be considered as visual illusions. We misestimate true color as we might misestimate true height in an Ames room (p. 119).⁴

We might wonder, however, whether this "top-down" computational approach, although consistent with the approximate constancy of natural color vision, should be accorded the status of an explanation. If we wish to design a visual system that exhibits complete constancy, and the system exhibits only approximate constancy, then we are

justified in saying that the system does not perform optimally, that it fails to achieve the task *for which it was designed*. But because natural color vision presumably resulted from evolutionary tinkering involving "trade-offs" rather than optimal design, why should the approximate constancy of natural color vision be explained by appealing to such a strong, engineering notion of optimality? Furthermore, even if natural color vision is in some sense optimal (relative to a given species and its niche), it might exhibit approximate constancy for biological and ecological reasons that preclude designating this kind of constancy as involving visual error. For example, most computational approaches seem to assume that color vision is concerned primarily with the reflecting properties of surfaces. As a result, illumination conditions are treated merely as something to be "discounted" in the task of recovering reflectance. Natural color vision appears to be concerned with illumination conditions in their own right, however, for these provide indications about weather conditions, time of day, and so forth (Jameson & Hurvich 1989). To emphasize color constancy at the expense of sensitivity to the illumination in its own right would therefore seriously prejudice the behaviors that natural color vision serves.

Consider, then, what happens if we proceed in a more "bottom-up" direction by taking the performance of natural color vision and its biological embodiment as our reference point. Here our point of departure is color space and its dimensions. That color constancy is only approximate provides an example of how these dimensions (hue-saturation-brightness) can shift depending on the state of the perceiver and the conditions of viewing. We therefore give more attention to the local, context-dependent features of perception than to the high-level, physically invariant properties of the environment. Furthermore, because our point of departure is color understood phenomenally, we are less likely to play favorites among the different ways colors can be encountered. For example, afterimage colors as well as surface colors require explanation. These both count as genuine color phenomena because they exhibit the three dimensions of hue, saturation, and brightness. It therefore becomes natural to identify color with this phenomenal structure. And because this structure does not reduce to properties of either light waves or surface reflectance (more on this later), we will probably be led to embrace subjectivism. Thus, consider the following passage from the conclusion of Zeki's (1983) pioneering study of cortical cell responses to both surface colors and after-image colors:

The results described here . . . suggest that *the nervous system, rather than analyze colors, takes what information there is in the external environment, namely, the reflectance of different surfaces for different wavelengths of light, and transforms that information to construct colors*, using its own algorithms to do so. In other words, it constructs something which is a property of the brain, not the world outside. (Emphasis in original, p. 764)

If we compare the above passage from Maloney (1985) with this remark of Zeki's, we can see that despite the considerable advances made in the study of color vision in recent years, disagreement remains among vision researchers on the ontology of color. Ontology is more the specialty of philosophers, but they do not agree either. In

tact, the two positions just outlined, with their respective links to computational vision and to neurophysiology, correspond precisely to the most recent discussions by philosophers: Hilbert (1987) and Matthen (1988) defend objectivism largely on the basis of computational color vision (Maloney & Wandell's model and Land's retinex theory); Hardin (1988) defends subjectivism largely on the basis of neurophysiology and psychophysics (opponent-process theories).

Two basic claims constitute Hilbert's version of objectivism. First, the centerpiece of his position is the typical objectivist claim that we must distinguish between color as an objective property of the world and color as we perceive it. For Hilbert, each objective color is identical with a distinct spectral reflectance. Objects that have identical surface spectral reflectances have the same color; objects that have different surface spectral reflectances have different colors. Second, Hilbert claims that since our color perception and color terms are indeterminate with respect to surface reflectance, they give us only "anthropocentrically defined colors and not colors themselves" (p. 27). For Hilbert, "red," "green," "yellow," and "blue" do not name determinate spectral reflectances; rather, they name indeterminate kinds of spectral reflectance whose specifications are arbitrary from a purely physical standpoint, but nonetheless of interest in relation to the structure of the human visual system. Hilbert accordingly calls his position "anthropocentric realism."

Matthen (1988) defends a similar view by first developing a theory of perceptual content. In his view, perceptual states have content because they have the function to detect things of a certain type. Matthen then argues on the basis of Land's retinex theory (Land 1977; 1983) that the function of color vision is to detect surface reflectance. Because Matthen identifies the contents of types of chromatic perceptual states with the distal property they supposedly have the function to detect, he is naturally led to claim that color simply is that distal property, namely, surface reflectance.

Hardin (1988), on the other hand, develops an extensive argument against objectivism, which consists of two basic points: First, surface spectral reflectance is only one of the many kinds of stimuli that can give rise to color experience; second, the properties of color—for example, the uniqueness and binariness of hue and hue-opponency—cannot be found in properties of the (distal or proximal) physical stimuli for color vision. The second is the more important point, for it consists in the claim that there is no mapping from physical stimuli to phenomenal color space that is sufficient to ground objectivism. As we saw above, however, there are mappings from color space to the visual system at various levels of organization and operation. Hardin relies precisely on these kinds of mappings, especially opponent-process theories, to support his subjectivist view that there are no "extradermal" colored objects; there are only chromatic perceptual states. In his words: "Colored objects are illusions, but not unfounded illusions. We are normally in chromatic perceptual states, and these are neural states . . . We are to be eliminativists with respect to color as a property of objects, but reductionists with respect to color experiences" (pp. 111–112).

We refer to these two positions as "computational

objectivism" and "neurophysiological subjectivism," respectively, thus highlighting the link between current color ontologies and explanation in visual science. The debate between these two positions has so far proceeded with computational objectivists downplaying the phenomenal structure of color and neurophysiological subjectivists responding by emphasizing the context-dependent, approximate constancy of human surface color perception.

Our intention in the remainder of this target article is to move beyond this debate by offering a broader empirical and philosophical perspective grounded in comparative color vision. Before we proceed, let us lay our cards on the table. With respect to the debate as outlined so far, we are fundamentally in agreement with Hardin's claim that "every attempt . . . to type-identify chromatic sensory states in terms of their stimuli is fundamentally misguided" (1989, p. 3). Nonetheless, we believe that Hardin's neurophysiological subjectivism is far too restrictive, for there are dimensions of color vision that do not yield to analysis purely in terms of the neurophysiological structure of the perceiver. These dimensions are, we argue, ecological. Hardin (1990) has recently begun to emphasize some of these dimensions, but we believe he has not gone far enough. On the other hand, although computational objectivism does emphasize the environmental context of color vision, it usually does so in a profoundly unbiological and unecological way by making animal-independent, distal properties the ultimate point of reference.

2. The comparative argument

2.1. Overview

Two pervasive phenomena of natural color vision form the basis for the comparative argument:

1. Animals whose neural apparatuses have little in common beyond the peripheral photoreceptor level (e.g., insects, fishes, birds, and primates), and that inhabit considerably diverse environmental contexts, nonetheless possess color vision.

2. Despite this commonality, color vision varies across species and animal groups. Among the most salient variations are the *type* (dimensionality) and *amount* (sensitivity) of color vision and its neural substrates. These variations imply different phenomenal color spaces, some of which are incommensurable.

These two phenomena constrain any attempt to explain color vision and the ontology of color. Our claim is that they constrain such a theory to be *experientialist* and *ecological*. Color can be understood only in relation to the visual perception of a given individual or species (contrary to objectivism); but such visual experience can be understood only in the context of its ecological embodiment (contrary to subjectivism).

Consider (1) first. Why do so many species of invertebrates, nonmammalian vertebrates, and mammals possess color vision? To answer we must appeal not only to comparative physiology, but also to the evolutionary histories of seeing animals (probably at several levels of selection), to common features among the diverse environmental contexts of color vision, and to changes in the

environment that are a function of animal-environment coevolution (we mention examples later).

These ecological dimensions of color vision have generally been taken to support computational objectivism. Among computational visual scientists, the argument (which usually goes unstated) is that because color vision is biologically pervasive, the evolution of color vision must consist in various species devising their own unique "solutions" to the information-processing problem of recovering surface reflectance in their respective environments. Among philosophers, this argument takes the form we reviewed above: The contents of perceptual states are to be type-identified by the (distal) properties they have the function to detect; the function of color vision is to detect surface reflectance; therefore color can be identified with surface reflectance (Matthen 1988).

In this section, we argue at some length against this view on the basis of (2) above, which develops into three related points:

(i) The properties of color, especially of different and sometimes incommensurable color spaces, cannot be modelled on the basis of properties of physical stimuli like surface reflectance. Hence color cannot be identified with surface spectral reflectance.

(ii) Given considerable variation in the dimensionality and sensitivity of color vision, and given a role for color vision in determining the boundaries of surfaces, the segmentation of the visual scene and therefore of what counts as a distinct surface to be perceived may in fact be relative to the structure of the perceiving animal. Thus surfaces may themselves be relational like color, providing no animal-independent anchor for objective color as surface spectral reflectance.

(iii) Natural color vision is concerned not just with detecting surfaces but also with a variety of other tasks in various terrestrial, aquatic, and aerial contexts. Among these are the discrimination of illumination conditions and the generation of a set of perceptual categories that have "cognitive significance" for animals in a variety of interactions. For these reasons, it is a mistake to suppose that the one and only (or even primary) function of color vision is the recovery of surface spectral reflectance.

Because each of these points rests on the idea of differences in color space, we begin by discussing the evidence for the existence of different kinds of color space among perceiving animals.

2.2. The color space of other animals

It is tempting to assume that our visual abilities provide the norm for understanding color vision. This assumption might be justifiable if humans – or our primate relatives – were unique in possessing color vision. In reality, however, color vision is widespread throughout the animal world. Indeed, it seems that virtually every animal class has some species with trichromatic vision (Jacobs 1981, p. 153). But it would also be a mistake to take trichromacy as the norm. Many animals are *dichromats* (e.g., squirrels, rabbits, tree shrews, some fishes, possibly cats and dogs, some New World monkeys); others appear to be *tetrachromats* (e.g., goldfish, the Japanese dace, turtles), perhaps even *pentachromats* (pigeons, ducks).

Before discussing the evidence for higher dimensional color spaces, it is important to consider how color vision

also varies considerably in its amount or sensitivity as determined by the spectral sensitivity, wavelength discrimination, and colorimetric purity functions. By measuring these functions for various animals, one can compare their overall sensitivities to spectral stimuli, their abilities to discriminate on the basis of wavelength, and whether spectral stimuli appear more or less saturated.

Each of these functions will differ for color vision of different dimensionality. The wavelength discrimination curve is of particular interest here, for it can also be taken as an indication of the type of color vision system: A maximum or minimum is expected where there is a crossover between two primary responses. For example, our three primaries are revealed in our wavelength discrimination curve, which has two maxima (Figure 2a, top). These maxima correspond to the two regions in the spectrum where our hue discrimination is finest (580nm and 470nm). The curve for the goldfish, however, shows three regions of best hue discrimination at 610nm, 500nm, and 400nm (Neumeyer 1985; 1986). This finding suggests that the goldfish has four active primaries and so is potentially a tetrachromat (Crawford et al. 1990; Neumeyer 1988). In contrast, the wavelength discrimination curve for the pigeon shows four regions of best hue discrimination at 390nm, 450nm, 540nm, and 600nm (Emmerton & Delius 1980), suggesting that the pigeon has five active primary mechanisms, and so is potentially a pentachromat. Three of the minima not including the one at 390nm in the UV region are shown in Figure 4c (Palacios et al. 1990a).

The three functions also differ among animals that have color vision of the same dimensionality, among "normal" and "anomalous" individuals, and even among "normal" individuals. To cite examples of each kind of variation: (i) Humans and forager honey bees are both trichromats, but bee color vision is shifted toward the ultraviolet, with the points of best hue discrimination at about 400nm and 490nm (Menzel 1979; 1989), as can be seen in Figure 3, which also shows the receptor-level and the opponent channel color space for these insects. (ii) For normal human trichromats, spectral sensitivity peaks at about 555nm; the spectral sensitivity of deuteranomalous trichromats, however, is shifted toward longer wavelengths, whereas that of protanomalous trichromats is shifted toward shorter wavelengths. (iii) Finally, each of the three functions can differ slightly among "normal" individuals: For example, men and women appear to differ in their color mixtures (Neitz & Jacobs 1986).

Now that we have introduced the idea of variations in the dimensionality and sensitivity of color vision, we can turn to what these variations tell us about color space. We focus first on color vision in birds, for as J. K. Bowmaker remarked some years ago: "The true culmination of the evolution of color vision in vertebrates is probably to be found in the highly evolved diurnal animals, perhaps best represented by diurnal birds, and it is within these species that we should look for color vision significantly more complex than our own and utilizing more of the available spectrum" (1980b, p. 196).

As we mentioned above, evidence that is now being accumulated indicates that such diurnal birds as the pigeon and the duck are at least tetrachromats, perhaps even pentachromats (Jane & Bowmaker 1988; Burkhardt 1989; Chen et al. 1984; Goldsmith 1990; Palacios 1991;

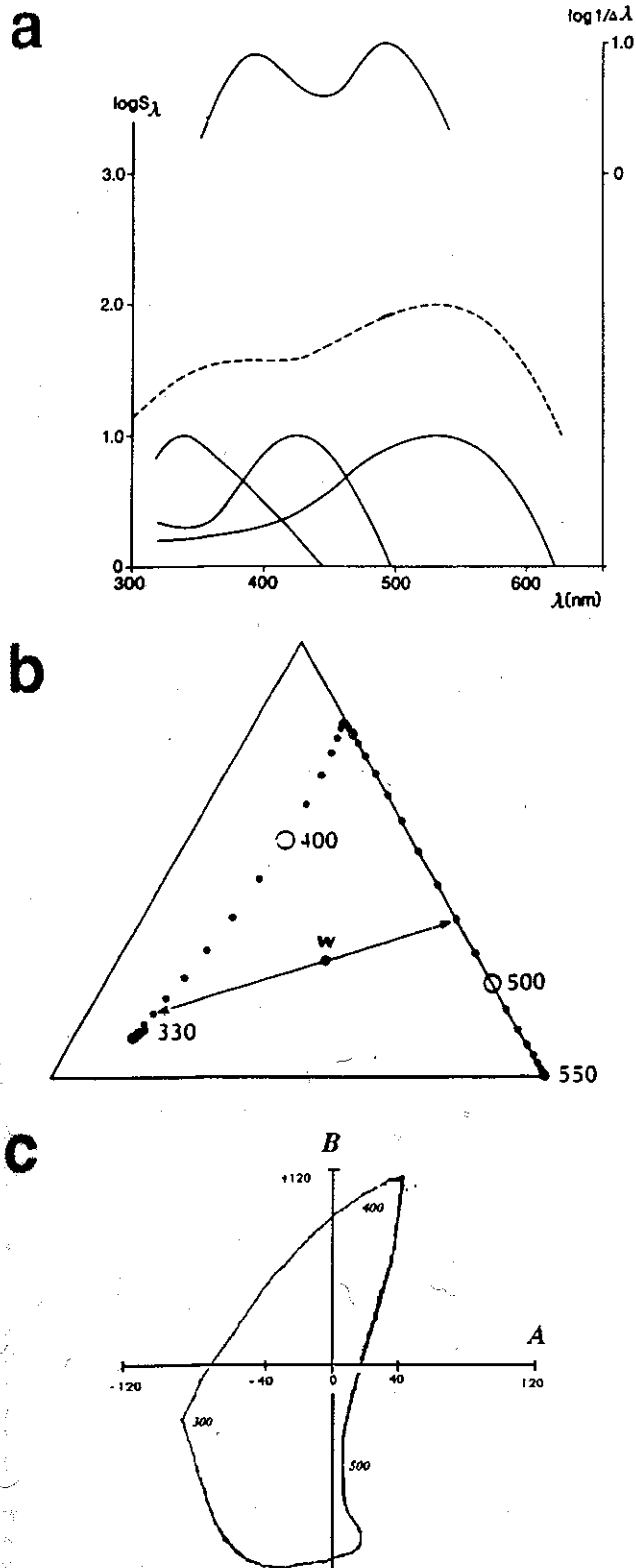


Figure 3. The honey bee: an alternative trichromatic color space. (a) Spectral sensitivities of the three pigments present in the forager honey bee's retina (*Apis mellifera*), with the overall sensitivity curve (discontinuous line) and the two-peaked wavelength discrimination function. Modified from Nuboer (1986). (b) Maxwell triangle for photoreceptor activities, indicating the spectral loci. Compare with Figure 2b. From Goldsmith (1990). (c) Chromaticity diagram for constant brightness, calculated on the basis of two antagonistic channels postulated for the bee from physiological findings: $A = +UV - B - G$; $B = UV - B + G$. From Menzel (1989).

Palacios & Varela, in press; Palacios et al. 1990b; Varela et al. 1991a). This evidence is derived from a variety of experiments with species ranging over various families within each order. The evidence also pertains to several levels, from the photoreceptor and retinal constitution, to the neurophysiological, and psychophysical or behavioral levels.

To begin at the retinal level, *five* different types of cone-oil droplet combinations have been described in the retinas of various birds such as pigeons, ducks, and penguins; passerines have at least four such combinations (Bowmaker 1977; Chen & Goldsmith 1986; Chen et al. 1984; Jane & Bowmaker 1988). As can be seen from Figure 4a, the "visible" spectral range available to diurnal birds includes that available to humans, but it also extends considerably further into the short-wave region. Indeed, it is now generally agreed that many birds have color vision in the near-ultraviolet region. For example, Wright (1972) found that the removal of an ultraviolet component changes the color of certain stimuli for the pigeon; and Goldsmith (1980) found that hummingbirds can distinguish near-ultraviolet light (370nm) from darkness, and from white light lacking wavelengths below 400nm. Humans cannot perform either of these tasks.⁵

The cones in the avian retina, unlike those in mammals and insects, also possess oil droplet inclusions, which appear to act as cut-off filters, thereby increasing in number the combination of receptor sensitivities (Bowmaker 1980b). Oil droplets are also found in the retinas of some fishes, amphibians, and reptiles. In the pigeon retina, for example, there are up to four types of colored oil droplets in combination with three types of cone photopigment for the long-wave region alone (Figure 4b) (Bowmaker 1977). Furthermore, this information about retinal organization is regional, because in birds like the pigeon there are two foveal regions that mediate different behavioral roles (Bloch & Martinoya 1983; Maldonado et al. 1988); evidence indicates that sensitivity and discrimination are different in these two visual regions (Nuboer & Wortel 1987; Remy & Emmerton 1989). These regional differences increase even more the complexity of pigeon color vision, for the color perceived depends on the visual field being attended.

Turning now to psychophysics, wavelength discrimination curves, as we have already mentioned, show four distinct minima (Emmerton & Delius 1980; Palacios et al. 1990a). Color-mixture experiments for the pigeon provide direct evidence for tetrachromacy (Palacios & Varela 1991; Palacios et al. 1990b). A definitive proof of pentachromacy would require five-way color-mixture experiments, which have yet to be performed.

There is unfortunately little evidence at present about the neural basis for avian chromatic channels in general (see Maturana & Varela 1982; Varela et al. 1983). It is nevertheless possible to form an educated guess about the possible shape of the pigeon's color channels, comparable to those shown in Figures 2c and 3c for humans and bees, respectively. The basic idea, introduced by Buchsbaum and Gottschalk (1983), is to obtain the weighted combination of mutual excitation and inhibition that maximally *decorrelates* the primary photoreceptor responses (see Appendix A). In their original calculations, Buchsbaum and Gottschalk (1983) used the Vos-Walraven primary responses for humans, which are psy-

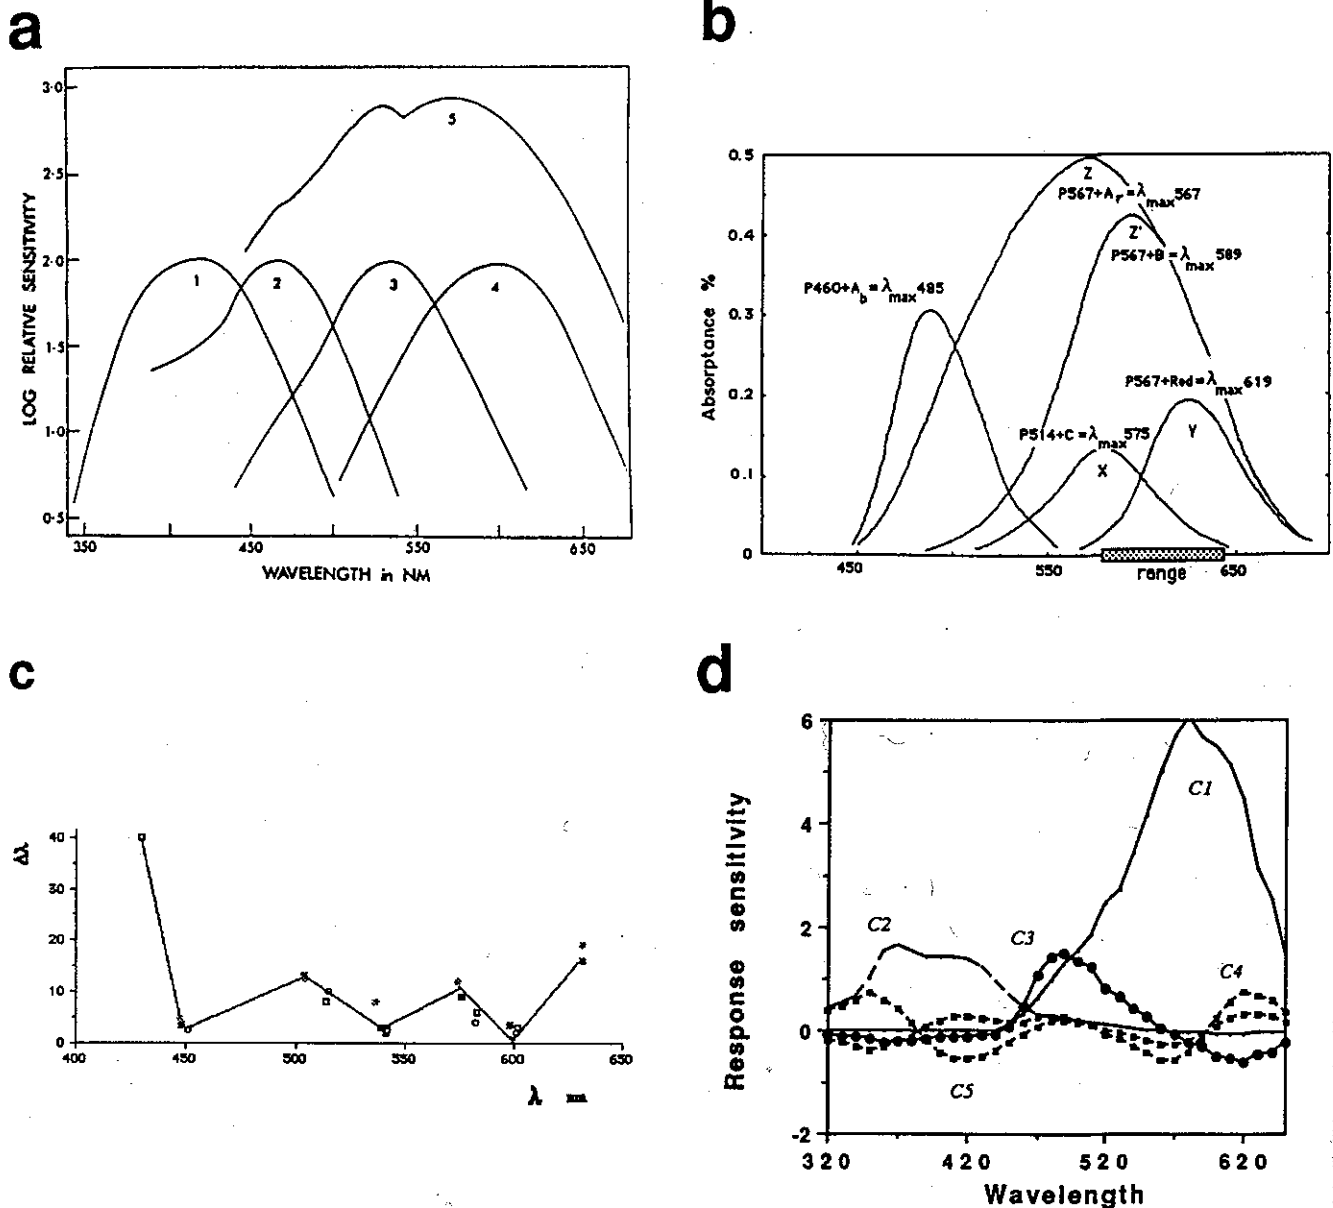


Figure 4. The color hyperspace of diurnal birds. (a) Relative spectral sensitivities of the five major cone classes of the mallard duck (*Anas platyrhynchos* L.), calculated by taking into account both oil droplet and ocular media absorption. The fifth curve has been shifted up one log unit arbitrarily for clarity. From Jane and Bowmaker (1988). (b) The cone photopigments (maxima at 460nm, 514nm, and 567nm) and oil droplet (50% cutoff transmission at 476nm, 476nm, 554nm, 610nm, and 570nm, respectively) combinations actually present in the pigeon's retinal "red" field. The sector designated "range," the long-wave window 580–640nm, is compatible with trichromacy according to color-mixture data. From Palacios, Martinoya, Bloch and

Varela (1990), based on data from Bowmaker (1977). (c) A behavioral determination of the wavelength discrimination function for the pigeon (*Columbia livia*) not including the UV region, showing three conspicuous minima. Data from Palacios, Bonnardel & Varela (1990a). (d) Proposed chromatic opponent channels for the pigeon, based on weighted subtractions and additions that maximally decorrelate primary responses (see text). The primary responses considered were maxima at 360nm, 415nm, 520nm, 580nm, and 620nm. These channels can adequately predict the known photopic sensitivity, wavelength discrimination curve, and color mixture in pigeons. From Palacios (1991).

chophysiologically derived. The resulting channels correspond remarkably well with the Wh-Bk (White-Black), R-G (Red-Green), and Y-B (Yellow-Blue) channels known to the psychophysicist, and to color-opponent profiles at the retinal or geniculate level known to the neurophysiologist.

We have applied this same procedure to other species, as explained in Appendix A. Unlike the human data, the animal data are incomplete; at present, the best one has to

work with are raw microspectrophotometric data. The proposed channels can be validated by their capacity to predict known behavioral evidence, such as sensitivity, discrimination, and color mixture. In Figure 4d we show the result of the decorrelation of the primary responses of the pigeon, thereby giving a set of putative channels. These channels adequately predict the known data on sensitivity, wavelength discrimination, and color mixture (Palacios 1991). We typically find that we need five

channels to account for the available data: one achromatic luminance channel (C1) and four color-opponent channels (C2, C3, C4, C5) with different zero crossings (Figure 4d).⁶

It is important to realize that such an increase in chromatic dimensionality does not mean that pigeons exhibit greater sensitivity to the monochromatic hues that we see. For example, we should not suppose that since the hue discrimination of the pigeon is best around 600nm, and since we see a 600nm stimulus as orange, pigeons are better at discriminating spectral hues of orange than we are. Indeed, we have reason to believe that such a mapping of our hue terms onto the pigeon would be an error: In an experiment designed to determine whether and how pigeons group spectral stimuli into hue categories, Wright and Cummings (1971) found that pigeons treat wavelengths to either side of 540nm as falling into different hue categories, whereas humans do not. As Jacobs (1981, p. 118) notes in his discussion of this experiment: "Among other things, this result strongly emphasizes how misleading it may be to use human hue designations to describe color vision in non-human species."

This point can be made even more forcefully, however, when it is a difference in the *dimensionality* of color vision that we are considering. An increase in the dimensionality of color vision indicates a fundamentally different kind of color space. We are familiar with trichromatic color spaces such as our own, which require three independent axes for their specification, given either as receptor activation or as color channels (Figure 2). A tetrachromatic color space obviously requires four dimensions for its specification. It is thus an example of what can be called a *color hyperspace*.

The difference between a tetrachromatic and a trichromatic color space is therefore not like the difference between two trichromatic color spaces: The former two color spaces are *incommensurable* in a precise mathematical sense, for there is no way to map the kinds of distinctions available in four dimensions into the kinds of distinctions available in three dimensions without remainder. One might object that such incommensurability does not prevent one from "projecting" the higher-dimensional space onto the lower; hence the difference in dimensionality simply means that the higher space contains more perceptual content than the lower. Such an interpretation, however, begs the fundamental question of how one is to choose to "project" the higher space onto the lower. Because the spaces are not isomorphic, there is no unique projection relation. Furthermore, to pass from one space to another, one needs to specify the appropriate axes (color channels), which differ according to the animal (even for animals that have color vision of the same dimensionality).

To mark this kind of difference in color space, consider the color space of the forager honey bee, which we presented in Figure 3 above. If bees are able to enjoy the experience of hue, the hues they perceive are likely to be different from ours, because they match wavelengths to which we are also sensitive with lights drawn from the near-ultraviolet region of the spectrum, which we cannot see. In spite of this difference, there is a sense in which bee color space and human color space can be said to be commensurable, for the dimensionalities of the spaces

are the same, and so a precise correspondence can be provided between these two perceptual spaces. In the case of tetrachromats or pentachromats such a correspondence is not possible.

This incommensurability can be more easily envisioned with the help of the evidence for tetrachromacy in teleost fishes, especially the goldfish (Neumeyer 1988). In Figure 5a, we present the pigment triangle for the goldfish, which should be compared with the previously presented triangles for humans (Figure 2b) and the bee (Figure 3b). Here the familiar Maxwell triangle has been doubled to accommodate the additional coordinate needed to map spectral loci. This representation was first proposed by Neumeyer (1988) for goldfish and independently by Burkhardt (1989) for the color vision of birds. We also reproduce here Goldsmith's (1990) suggestive rendering of the same idea as a pigment *tetrahedron* for an "imaginary" turtle with a retina whose photoreceptors have no oil droplets (Figure 5b).


To generate these kinds of color hyperspaces at the physiological and psychophysical levels, we need at least four channels. For teleost fish, these channels can be obtained in a manner similar to those for the pigeon – by maximally decorrelating the primary responses as known from microspectrophotometric and physiological data (Harosi & Hashimoto 1983; Neumeyer 1988). In this case, the transformation matrix from primary responses to channels is explicitly given in Appendix A. Here again the putative channels can correctly predict the known behavioral evidence for sensitivity, wavelength discrimination, and color mixture for the goldfish (Palacios 1991). It is interesting to note that Neumeyer and Arnold (1989) have recently shown that the goldfish switches from trichromatic to tetrachromatic modes depending on light conditions – an indication that the ecological embeddedness of the animal is quite pertinent even at this level of description.

This complex of issues can also be approached from a frequency analysis of color signals and responses that could provide a way to ascertain the dimensionality of color space directly. The basic idea was first proposed by Barlow (1982; see also Bowmaker 1983), but since it is rather novel, we have relegated our treatment to an appendix (see Appendix B).

In this section, we have presented an array of evidence for the existence of different kinds of color space among perceiving animals. The evidence includes the diversity of kinds of photopigments and sensitivity functions, wavelength discrimination and categorical perception, color mixture, physiological processes, and ethology. Although each form of evidence taken in isolation might be unconvincing, taken as a whole it makes a strong case for the existence of significant variations in the dimensionalities of color space among perceiving animals. We must now wonder what these differences might mean in experiential or phenomenal terms. What do these comparative variations in color space imply for our understanding of color experience?

2.3. Novel hues and diversity in color experience

Since some readers may be sceptical about attributing color experience to creatures other than ourselves (or our primate relatives), let us first attempt to distinguish more



77. (c) A discrimination of the UV spectrum by the pigeon. The response of the pigeon is high in the UV region (below 400 nm) and decreases towards longer wavelengths. The response of the human is low in the UV region and peaks around 550 nm.

data. The capacity to sensitivity, we show responses of channels. data on mixture need five




a**b****c**

Figure 5. Tetrachromatic color hyperspaces. (a) Pigment color space for the goldfish, requiring four relative values to span the entire space. In contrast to Figures 2 and 3, the representation demands an extra dimension, which is obtained by doubling the color triangle. The spectral loci thus obtained from pigment absorption data no longer fall on the boundaries of the space, but inside them. From Neumeyer (1988). (b) Pigment tetrahedron for an "imaginary" turtle with no oil droplets in its cones, with visual pigment maxima at 370nm, 450nm, 520nm, and 620nm. In this format it is easy to see that the animal is likely to have three nonspectral stimulus regions (see Figure 2) which would fall along the dashed lines. From Goldsmith (1990). (c) Proposed

precisely between color vision and color experience. Although it is difficult to draw a principled distinction between mere wavelength-specific behavior and color vision (Menzel 1979), color vision is sometimes defined as the ability to discriminate wavelengths independent of their relative intensities. This ability would not seem to entail the enjoyment of color experience, however, for it seems possible to imagine the former without the latter. Unfortunately, it is also not clear how to draw a principled distinction between color vision as wavelength discrimi-

chromatic opponent channels for the goldfish, by maximal decorrelation of the pigment data from Harosi and Hashimoto (1983). An achromatic channel and three opponent channels are found. These four channels adequately predict sensitivity, wavelength discrimination, and color-mixture data, as well as physiological data from the fish retina. From Palacios (1991). To the right, the putative channels are used as axes of a color hyperspace of four dimensions, with opponencies between three sets of dimensions (labelled here simply as +C1 - C1, +C2 - C2, etc.), plus the achromatic or brightness axis (C1 = Wh). Compare with Figure 2c and 3c for the trichromatic spaces of humans and the bee respectively.

nation and the full-fledged perceptual experience of color.

A more satisfactory approach to this problem is to hold that color perception involves at least three important phenomena: additive color mixture (hue, saturation, and brightness matches for spectral stimuli), color contrast (simultaneous and successive) and color constancy. In particular, it seems reasonable to suppose that color contrast and color constancy are necessary for color experience.⁷ It is therefore interesting to note that these

chromatic phenomena have now been demonstrated for a variety of species. Color constancy and color induction have been found in bees (Neumeyer 1980; 1981), goldfish (Ingle 1985), and pigeons (Varela et al. 1991a). In the case of pigeons, we have also seen that these animals group adjacent wavelengths into categories, which though different from the groupings humans perform, nonetheless seem to be categories of hue (Wright & Cummings 1971). Given these psychophysical results, then, it does not seem unreasonable to suppose that these animals, especially birds and fishes, experience color.

Let us now consider color vision of higher dimensionality than our own. When they hear of the evidence for tetrachromacy or pentachromacy many people respond by asking: "Well, what are the extra colors that these perceivers see?" This question is understandable, but somewhat naive, for, as pointed out above, we should not suppose that tetrachromats or pentachromats are simply better at seeing the colors that we see. On the contrary, to see in four or five dimensions, as it were, is not to discriminate more finely in three dimensions. In other words, tetrachromats and pentachromats should not be conceived as perceivers who simply make finer hue-saturation-brightness discriminations among, say, blue and green or red and yellow (like perceivers who simply see finer shades of our colors). Such an ability would not amount to an increase in the *dimensionality* of color space; it would consist only in a relative increase in hue-saturation-brightness sensitivity within the dimensions of our trichromatic color space. If we wish to understand what tetrachromacy and pentachromacy imply for color experience, we must instead ask what the possession of additional dimensions to make chromatic distinctions could mean in experiential or phenomenal terms.

At this point, we can offer only imaginative speculation, for we still lack knowledge of the post-retinal neuronal processes involved in tetrachromatic and pentachromatic perception and we obviously do not know what such perception is *like* from the point of view of the goldfish or the pigeon.⁸ By returning to consider our color space, however, and by asking how this space would be transformed by the addition of a new dimension, we can perhaps achieve an indirect appreciation of what a tetrachromatic color hyperspace might be like. Recall, then, that because our visual system has two chromatically opponent channels, we are able to experience four unique hues (red, green, yellow, and blue) and their binary combinations (orange, purple, etc.). A tetravariant visual system, however, like that suggested for the goldfish (Figure 5), would contain three chromatically opponent channels. We are therefore entitled to speculate that these three channels (call them r-g, y-b, and p-q) would enable a tetrachromat to experience *six* basic hue components (r, g, y, b, p, q), binary combinations of these hues (e.g., r + y, y + p, etc.), and *ternary* combinations as well (e.g., r + y + p, g + y + p). Thus the color hyperspace of a tetrachromat might reflect a phenomenal structure composed not only of two new basic hue components, which would combine to form novel binaries, but also an entirely new kind of hue not found in the phenomenal structure of our color space, namely, ternary hues.⁹ These ternary hues would correspond to the additional kind of chromatic distinction available to a tetrachromat, but not to a trichromat.

This point about novel colors can be made more accessible with the help of the diagrams presented in Figures 2, 4, and 5. To represent tetrachromatic stimulus mixtures we found that we had to add an additional axis to the plane so that it became a volume (Figure 5). Therefore, as Burkhardt (1989) notes: "While in man's chromaticity diagram there is only one intermediate color which does not occur in the daylight spectrum, namely, purple, in tetrachromatic vision there would be three intermediate colors which are not present in the daylight spectrum, namely, mixtures of red and blue (purple), of green and UV, and of red and UV . . ." (pp. 794-95). Similar kinds of novel, nonspectral stimulus mixtures are indicated by the dashed lines in Goldsmith's (1990) pigment tetrahedron for an imaginary turtle (Figure 5b). If such novel, nonspectral stimulus mixtures can be shown to be treated as colors by the animal through its behavior, then the existence of color hyperspace as a domain of behaviorally significant distinctions would be strongly reinforced.

2.4. Computational objectivism revisited

We now pursue the implications of our comparative discussion, beginning with the view that we call "computational objectivism." Our first task is to determine whether some sufficient subset of the properties of color can be identified with such physical properties as surface spectral reflectance. If these properties of color cannot be so identified then we have reason to reject the objectivist's claim that color is simply surface spectral reflectance.

The obvious place to begin is color space. For something to be a color it must have a location in some color space; that is, it must be specifiable in terms of hue, saturation, and brightness. By taking these three properties as our reference point, we can construct an argument against the identification of color with surface spectral reflectance, which we will call the "argument from external irreducibility." (The main features of this argument were originally proposed by Hardin [1984; 1988, pp. 66-67].)

The argument from external irreducibility

1. For something to be a (chromatic) color it must be a hue.
2. For something to be a hue it must be either unique or binary (or ternary).
3. Therefore, if hues are to be reductively identified with physical properties, these physical properties must admit of corresponding unique, binary (or ternary) divisions.
4. Organism-independent, external properties such as light-waves and spectral reflectances do not admit of such divisions.
5. Therefore, color cannot be reductively identified with such organism-independent, external properties.

Although this argument has conceptual components, we do not intend it to be primarily conceptual. We are interested not in conceptual analysis, (i.e., in making claims about the essential features of the concept of color), but in determining what color is, given the concept of color as it figures in visual science, especially in psychophysical explanation. Thus (1) and (2) should be read as consequences that follow from how color is conceptually and empirically specified in visual science.

The main empirical claim in the argument is obviously (4). To put the point another way: Given only light wavelengths or the spectral reflectance profiles for surfaces, we cannot model or state generalizations about hue. Light waves or surface spectral reflectances do not stand in relations to each other that can be described as unique or binary, or for that matter opponent or nonopponent, balanced or unbalanced, saturated or desaturated, and so forth. There is simply no mapping from such physical properties to the properties of color that is sufficient to establish the objectivist identification.¹⁰

This argument obviously depends on considerations about what properties a mapping must have to be sufficient to establish objectivism. We are supposing that such a mapping must enable us to state generalizations about features of color such as the unique/binary structure of hue and the opponent relations. The objectivist might deny this point. Such a denial would be tantamount to claiming, contrary to (1)–(3) above, that we should replace our current understanding of color in visual science with a new concept of physical color as surface spectral reflectance. It is tempting to dismiss this conceptual replacement idea out of hand: It is one thing to argue for a distinction between physical color and perceived color, but it is quite another to uphold the distinction by divesting hue, and thereby color, of those properties used in its conceptual and empirical specification in visual science. If the properties of hue, such as being unique or binary, and the opponent relations, could be successfully identified with some set of physical properties such as surface spectral reflectances, then statements about these physical properties would provide us with an alternative theoretical access to the properties of color that figure in psychophysical explanation. We would then have reason for accepting a new notion of physical color. But because we have no physical model for these properties of color, what is to motivate such a notion?

The argument from external irreducibility could be and indeed has been advanced without taking into consideration comparative color spaces (Hardin 1984; 1988). We present the argument here because it becomes even stronger when we place it in the context of comparative color vision. Consider hue as it varies across dichromatic, trichromatic, and tetrachromatic (to say nothing of pentachromatic) color spaces. As we know from studies of human color blindness (see Hurvich, 1981, for an overview), a dichromat has only one opponent-hue pair (yellow-blue or more rarely red-green). Therefore, unlike a trichromatic color space, a purely dichromatic color space contains no binary hues. Similarly, a tetrachromatic color hyperspace would contain ternary hues not found in a trichromatic color space. There are thus different *kinds* of hue to be found in each of these color spaces. The unique, binary, and ternary structures that compose these different kinds of hue do not map onto properties of surface spectral reflectance. Neither the unity among the phenomena (color *qua* hue-saturation-brightness relations) nor the relevant diversity (different dimensionalities and hence kinds of hue) is to be found at the purely physical level of spectral reflectance.

At this point, the objectivist will no doubt appeal to the idea that the contents of (types) of perceptual states should be identified according to the distal properties they have the function to detect, that the function of color

vision is to detect surface spectral reflectance, and that this functionalist type-identification is sufficient to establish the claim that color is surface spectral reflectance (Matthen 1988).

Even if the function of color vision is to detect surface reflectance, it does not follow that color is surface reflectance. In fact, computational objectivists often simply beg the question about the status of color by building objectivism into their representationist theories of perception. But there is an even more fundamental problem we wish to stress: The claim that the function of color vision is to detect surface reflectance is at best considerably misleading and at worst seriously flawed. It is misleading because a comparative ecological examination of color vision reveals that color vision has many other biological functions besides those involved in the detection of surfaces. Most notably, color vision is concerned with illumination conditions in their own right and with the perceptual significance of color in guiding behavioral interactions (we provide examples shortly). It is flawed because it is not at all clear that surfaces are themselves perceiver-independent in the way the objectivist supposes. Let us begin with this point, because it is the more controversial.¹¹

In the top-down functional decomposition characteristic of the computational level of analysis (such as inverse optics), vision is decomposed into various more-or-less modular tasks. In the case of color vision, the task is to recover information about surface spectral reflectance given a collection of objects. This statement of the "problem" of color vision assumes that the visual scene has already been segmented into areas that correspond to distinct objects and their surfaces. But this assumption begs the question of the purposes that color vision may serve. As D'Zmura and Lennie (1986, p. 1666) note: "To find the loci of responses that correspond to different objects, one must already have segmented the scene to establish which lights come from which objects. This begs the question of the purpose of color vision, which we believe plays an important role in the discrimination among objects and in their identification."

Consider the regularization constraints that we discussed above. Among these is the fact that naturally occurring illuminants and reflectances can be adequately modelled in a low-dimensional space. What, we might ask, constitutes a "naturally occurring reflectance"? If we examine these models, we see that so-called natural reflectances correspond to the surface reflectances of typical objects from our human environment (e.g., bricks, grass, buildings, etc.). Given a class of such objects, one measures their surface spectral reflectances and then determines which finite set of basis reflectance functions best models the variance in the class. The visual system, however, is never simply presented with such prespecified objects. On the contrary, the determination of what and where an object is, as well as its surface texture and orientation (hence the overall context in which color is perceived), is a complex process that the visual system must achieve.

In fact, we have already seen that color vision contributes to this process of spatial segmentation. For example, we reviewed Rubin & Richards's (1982; 1988) idea that an early goal of biological color vision is to determine where material changes occur in a scene using spectral cross-points and opposite slope signs. Another example comes

from D'Zmura and Lennie (1986), who propose an algorithm for color constancy in which the geometric stability of hue contributes to segmenting a scene whereas variations in lightness and saturation contribute to establishing an object's shape and position. Finally, at the neurophysiological level, it is generally held that distinct retinal-geniculate-cortical pathways are involved in color, form, and movement perception (De Yoe & Van Essen 1988; Lennie 1984; Livingstone & Hubel 1988), but there is disagreement over the properties of these pathways and their relations (Lennie et al. 1990; Logothetis et al. 1990; Mollon 1990; Schiller et al. 1990). Nonetheless, it seems safe to say that not only the achromatic process involving the lightness dimension of surface color, but also the chromatically opponent processes play a role in spatial segmentation (e.g., color contrasts can be used to determine borders), to say nothing of how the specification of color and shape may be combined in, say, V4 (Heywood & Cowey 1987).

This interdependency between color vision and spatial segmentation is downplayed by computational objectivism. This view simply assumes that the surfaces of objects provide a perceiver-independent peg on which to hang objective color as spectral reflectance. The objectivist presumably thinks that the assumption needs no defence, since surface spectral reflectances can be specified in physical terms. But although the *reflectance* at any point in the scene can be specified in physical terms, what counts as a *surface* may in fact involve tacit reference to a type of perceiver. This point has not been evident for several reasons which we need to explore.

First, it is usually simply assumed that surfaces are properties of the world found at a purely physical level of description. At the level of description found in perceptual theory, however, surfaces also figure as properties of the *perceptual object*, that is, the object as construed in relation to the sensory-motor capacities of perceiving animals. Here surfaces are treated in relation to the body-scaling of the animal, and, in the case of vision, to properties like visual shape, texture, orientation, lightness, and color. How surfaces at this perceptual level are to be linked to the purely physical level is precisely the issue. Simply to assume that this link can be made without implicating the sensory-motor capacities of the animal is to beg the question. For example, the visual system responds to singularities at many scales that characterize apparent contours in edges, yet these scales are integrated into a unified behavioral designation. This integration, however, is not implicit in the singularities themselves: it depends on how the neuronal processes treat them (DeValois & DeValois 1988).¹²

Second, such issues often remain hidden because many theories of perception focus on tasks (e.g., recovering reflectance) in an already well-specified or easily segmented context (e.g., Mondrian displays). Uncritical attention to visual tasks in such artificially simplified contexts makes one forget the complexities involved in spatial segmentation itself.

Third, virtually all theories of perception focus on our familiar human environment rather than the considerably different environments of, say, birds, fishes, or insects. For example, the prespecified objects in low-dimensional models of reflectance are typically middle-sized, frontally viewed, "human" objects, such as bricks, grass, build-

ings, Munsell color chips, Mondrian displays, and so forth. They are not, for example, silhouettes against the background sky, as seen frontally and laterally by birds, ultraviolet reflectance patterns of flowers, as seen by birds and bees, aquatic objects that contrast with the volume colors of the downwelling or background space light as viewed by fishes, and so on. Because of this attention to prespecified human objects, the issue of how the world comes to be segmented into a given collection of surfaces by different perceiving animals is hardly ever empirically raised, or explored philosophically (but see Stroll 1986).

Finally, if, as Gouras and Zrenner (1981, p. 172) claim, "it is impossible to separate the object sensed from its color because it is the color contrast itself that forms the object," then what counts as the perceptual object may vary considerably depending on the type of color vision system involved.¹³ Gouras and Zrenner are referring here to the perceptual object. Nonetheless, the interdependency between color and surface perception is enough to show that at the level of description relevant to perception, it is not at all evident, as computational objectivism assumes, that the specification of surfaces is not relative to the perceiver. In other words, the kinds of surfaces that populate the world as visually perceived by a given animal may depend for their specification on the processes by which that animal segments its visual scene.

Let us summarize this line of argument, which we call the "argument from perceiver-relativity":

The argument from perceiver-relativity

1. Color vision contributes to the task of segmenting the visual scene into regions of distinct surfaces and/or objects.
2. Color vision varies considerably throughout the animal world.
3. We may accordingly expect spatial segmentation to vary as well.
4. What counts as the surface of an object (for perception theory) therefore has to be specified in relation to the perceiving animal.

We take this argument to be conditional, based on a reasoned hypothesis or conjecture that requires further empirical investigation in a comparative neurophysiological and ecological context. Nonetheless, we believe its plausibility undermines much of the computational objectivist's assumption that surfaces provide a perceiver-independent anchor for color.

The comparative ecology of color vision reveals several other points where computational objectivism is inadequate. The computational objectivist typically assumes that the sole or genuine function of color vision is object detection. Matthen (1988), for example, argues that the "biological function" of color vision is object detection via the recovery of surface reflectance. He claims that non-surface modes of color appearance, such as the blue of the sky, should be explained as the result of "normal misperception," that is, as cases of visual representation that are nonveridical, but do not result from the malfunction or maladaptation of the visual system. Thus Matthen tries to ground the philosophical claim that only surfaces are genuinely colored by relying on a notion of "biological function." The irony of this proposal is that Matthen's claim about the *biological* function of color vision is advanced on the basis of a rather controversial *computa-*

tional model (Land's retinex theory) and is undermined by the actual biological and ecological operation of color vision. Color vision is not limited to the perception of surfaces; it includes the perception of the ambient lighting conditions in their own right (not merely as "information to be discounted"), for these are relevant to a variety of environmental conditions, such as weather and time of day (Jameson & Hurvich 1989). Nonsurface modes of color vision also serve to heighten contrast between aerial or aquatic backgrounds (volume colors) and foreground objects (surface colors) (Levine & MacNichol 1982; Lythgoe 1979).

Although we still lack extensive knowledge of the ecological function of color vision in various animal species, the evidence we do have is sufficient to demonstrate that speculations about color vision should not be dictated by top-down computational models that rely on a considerably simplified human perceptual context. Instead, as the following examples will illustrate, color vision must be understood within the context of the quite different behavioral repertoires available to perceiving animals.

Consider first the link in chromatic ecology between visual pigments and animal niches. Studies of aquatic visual ecologies have shown that the retinas of deep sea fishes have been reduced to one rhodopsin pigment, with a sensitivity maxima around 470–490nm (Levine & MacNichol 1979; Loew & Lythgoe 1978; Partridge et al. 1989). According to one interpretation, this range would permit the maximum of contrast sensitivity for movement (Crescitelli et al. 1985; Muntz 1975). In contrast, for species that have a bioluminescent organ, a different pigment of the porphyropsin family, whose sensitivity maximum (513–539nm) coincides with the main luminous emission of the organ, has been described (Bowmaker et al. 1988; Partridge et al. 1989). As one moves upward toward more illuminated depths, all species have a larger diversity of photopigments and retinal arrangements, from which one may conclude that different visual objects are pertinent for each species (Lythgoe 1979). This ecological link between photopigments and environments becomes even more striking in migratory fishes, which go from river to sea: Here the relative amounts of rhodopsin and porphyropsin change according to the time of day and season (Beatty 1969; 1984; Bridges 1972; Muntz & McFarland 1977; Muntz & Mouat 1984; Whitmore & Bowmaker 1989). Such polymorphism is also present among other nonmigratory species (Archer & Lythgoe 1990; Archer et al. 1987; Whitmore & Bowmaker 1989). Neumeyer and Arnold (1989) have also recently shown, as we mentioned above, that the goldfish is tetrachromatic for an ambient illumination of 25lux, but trichromatic for a lower illumination around 1.5lux. They suggest that this capacity for a dimensional shift is likely to have an ecological interpretation. This evidence for aquatic ecologies is admittedly fragmentary, yet it serves to indicate the need to link chromatic performance to the ecological setting of the animal (Muntz 1975; Wheeler 1982).

Among birds, the retinal oil droplets vary considerably even for species with similar global living conditions (Budnik et al. 1984; Jane & Bowmaker 1988; Martin 1977; Martin & Lett 1985). For example, the common tern, a predator bird, has a significant amount of red and yellow

droplets in the dorsal retina, while the barn swallow, which catches insects, has a large quantity of translucent droplets (Goldsmith et al. 1984). In fact, Partridge (1989) has shown by means of cluster analysis that the ecological niche (herbivore, fishing, etc.) is more important in predicting the kinds and distribution of oil droplets than strict phylogenetic kinship. The presence of ultraviolet pigments in birds also provides an example. These pigments can be linked to bird-fruit coevolution, including the dissemination of kernels (Snow 1971; Burkhardt 1982), and to ethological factors involving animal recognition, for bird plumages have been shown to have high frequency content, and so might require higher-dimensional color spaces for their recognition (Brush 1990; Burkhardt 1989; Durrer 1986; Hudon & Brush 1989; Weedon 1963).

Ultraviolet sensitivity in birds may also be used in aerial navigation. As we have seen, pigeons have excellent short-wave and near-ultraviolet discrimination. It is possible that, in Nuboer's (1986, pp. 370–71) words, "the excellent spectral discrimination within this range . . . represents an adaptation to the coloration of an unclouded sky. This property enables the pigeon to evaluate short-wave gradients in the sky, ranging from white at the sun's locus to highly saturated (ultra) violet at angles of 90° to the axis between observer and sun." Furthermore, since pigeon navigation is based on orientation with respect to the sun's azimuth, "the perception of colour gradients in the sky may control navigation indirectly when the sun is hidden by clouds."

A different, but perhaps even more important feature of the ecological function of color vision is to yield a set of perceptual categories that have "cognitive significance" for perceiving animals in a variety of behavioral interactions (Jacobs 1981, pp. 170–71). A color category can guide behavior in various ways depending on the things which exemplify it: In the case of fruits, it guides feeding; in the case of animal coloration, it may guide various social interactions, such as mating. Pigeons have been shown to group spectral stimuli into hue categories, and the brightly colored feathers of birds must have cognitive significance for behavior, especially behavior involving sexual recognition. Finally, although object discrimination is obviously important for these kinds of behavior, the cognitive significance of color may have an affective dimension (perhaps related to the overall hormonal/motivational level of the animal) that cannot be explained simply as a function of object discrimination (Varela et al. 1983).

Much research remains to be done on the relations among color vision, perceptual color categories, and animal behavior (Burt 1979; Hailman 1977). Although color as a perceptual category with cognitive significance obviously plays a great role in human life, there is still little evidence about this dimension of color perception in nonhuman animals, especially nonprimates. In the case of birds, however, it seems safe to conclude that this kind of color experience does exist, as we have been arguing here. In any case, the evidence that we have presented in the previous paragraphs serves to demonstrate our point that the functions of color vision should be understood in the context of the actual behavioral repertoires and visual ecologies of perceiving animals.

2.5. Neurophysiological subjectivism revisited

To emphasize the active role that color vision plays in tasks such as spatial segmentation and in guiding the interactions of perceiving animals implies an approach to color perception that is also different from neurophysiological subjectivism. To demonstrate this point, we need to consider Hardin's (1988; 1990) defense of neurophysiological subjectivism.

Hardin's strategy is to offer what we can call an "argument from internal reducibility" whose main claim is that the properties of hue (e.g., uniqueness, brightness) can be reductively identified with psychophysical and eventually neural properties of the visual system. This argument, coupled with the "argument from external irreducibility," leads Hardin to the position that there are no "extradermal" colored objects; there are only chromatic neural states. Hardin's defense of this idea is worth quoting:

We have no good reasons for thinking that such a replacement of the one [phenomenal] description by the other [neural] description would leave anything out, with a consequent loss of information. On the contrary, we have reason to expect that a proper neural description would be richer, more complete, and, in principle, more penetrable by the intellect. Problems that are intractable at the extradermal physical level or at the phenomenal level promise to yield analysis in neurological terms (1988, p. 111).

Two points appear to be contained in this remark, one ontological, the other methodological. The ontological claim is that color, or rather chromatic experience, is a type of neural state or process. The methodological claim is that color phenomena can be analyzed in neurological terms. These two claims obviously support each other: If colors are really neural states, then we have reason to pursue a neurological analysis of color phenomena; on the other hand, if we can give a neurological analysis of color phenomena (and we cannot give a comparable physical analysis), then we have reason to believe colors are neural states. We make this distinction not to be pedantic, but because it is primarily the methodological issue that we wish to address here, not the ontological one. In other words, we do not intend to evaluate Hardin's position by embarking upon a discussion of the mind-body problem for visual experience. It is, rather, the scope and limits of a purely neurological approach to color phenomena that interests us.

Our aim in this final section of the comparative argument will be to show that there are indeed phenomena that, intractable as they are at the extradermal and organism-independent physical level as well as the phenomenal level, nonetheless fail to yield to analysis in purely neurological terms. These phenomena are ecological in the broadest sense; that is, they encompass not only the extradermal world as an animal environment, but also perceiving animals as both assemblies of sensory-motor networks and as organismic unities that *shape the extradermal world into an environment in their interactions*.

Consider first the polymorphism in the color vision of the squirrel monkey and the spider monkey (Jacobs 1986). In these species, all males are dichromats, whereas

three-quarters of the females are trichromats. Several explanations have been proposed for this polymorphism (Mollon et al. 1984; Nuboer 1986). According to one, it has resulted from adaptation to the spatial heterogeneity of the environment: It is possible that different phenotypes inhabit regions of the jungle that differ in the spectral composition of their ambient light. A second proposal appeals to the hypothesis of group selection: It might be advantageous for the animal community to have members with several forms of color vision. A third proposal appeals to frequency dependent selection: There may be an ecological balance between the availability of certain fruits and the number of phenotypes that can detect them. Finally, another hypothesis holds that the colors of local fruits coevolved with the differences in color vision (Snodderly 1979).

Our second example comes from the color vision of bees. We have seen that bees have trichromatic vision that is shifted towards the ultraviolet. It has been argued that this distinctive form of trichromacy coevolved with the colors of flowers, which often have contrasting patterns in ultraviolet light (Barth 1985; Lythgoe 1979; Menzel 1989; Nuboer 1986). On the one hand, flowers attract pollinators by their food content, and so must be conspicuous and yet different from flowers of other species. On the other hand, bees gather food from flowers, and so need to recognize flowers from a distance. This mutual advantage seems to have determined a coevolution of plant features and sensory-neural capacities in the bee.

Finally, consider that the colored "objects" that animals discriminate are often (perhaps typically) other animals. Therefore, within an ecological framework our inquiry should be concerned just as much with animal coloration – indeed with the coloration of living things in general – as with animal color vision (see Burkhardt 1989; Burt 1979; Hailman 1977; Lythgoe 1979). Coloration obviously affects an animal's visibility, both to conspecifics and to members of other species in its environment. It is therefore not surprising to find coloration involved in camouflage and in many kinds of visual recognition (e.g., species recognition, sexual recognition, individual recognition, recognition of motivational state, etc.; Baylis 1979; Rowland 1979). Indeed, the ecological entanglement of color vision and animal coloration is truly astounding. Consider, for example, the variations in color vision and coloration among fishes in a tropical coral reef, perhaps one of the richest of color environments.

These kinds of phenomena indicate that a purely neural explanation for color vision is incomplete. To explain the polymorphism in spider and squirrel monkey color vision, and hence the differences in the perceptual experiences of these animals, we must appeal not simply to the neurophysiological constitution of these animals, but also to the evolutionary histories of their environmental interactions, perhaps at several levels of selection.¹⁴ Similarly, to understand why bee color vision is shifted toward the ultraviolet, and hence why the color space of the bee might comprise novel hues, we must appeal to animal-environment coevolution. Finally, to understand the relations among color vision, animal coloration, visual recognition, and animal communication, we must appeal to a broad range of physiological, ecological, and evolu-

tionary considerations, ranging from the physiological functions of pigmentation, to coordinated inter- and intraspecific animal interactions, to the coevolution of the various behavioral partners (Burt 1979).

We expect that Hardin would not deny any of these points. Indeed, Hardin has recently drawn on evolutionary – or more broadly, ecological – considerations to defend his view that although chromatic categories (red, green, yellow, and blue) have no counterparts in the extradermal world, such categories confer evolutionary advantages on perceiving organisms (Hardin 1990). Hardin argues that color vision does not represent the world as it really is, but rather “encodes information” about light, reflectance, and so forth, in a subjectively generated form that is salient, vivid, and of great practical value for the perceiver. The salience and vividness are to be explained neurophysiologically, whereas the practical value is to be explained ecologically.

We believe, however, that Hardin has not yet appreciated the moral that evolutionary and ecological considerations have for his neurophysiological subjectivism. Color vision does not merely provide practical knowledge of the environment; it also participates in the *codetermination* of perceiving animals and their environments. By *codetermination* we mean both (1) that animals select properties in the physical world relevant to their structure (body-scaling, sensory-motor capacities, etc.), shaping these properties into environments that have behavioral significance; and (2) that environments select sensory-motor capacities in the animal and thereby constrain animal activity (Levins & Lewontin 1983; 1985). Consider once again the coevolution of plant features and sensory-neural capacities in the bee (and other invertebrates). This coevolution implies not only that bee color vision is sensitive to ultraviolet because it is advantageous for bees to detect flowers that have ultraviolet reflectances, but also that flowers have ultraviolet reflectances because it is advantageous for them to be seen by bees. Thus, the evolution of bee color vision did not simply provide the bee with a practical knowledge of its environment; it also contributed to the very determination of that environment. As Barth (1985) says in his wonderful study of insects and flowers: “The colorful field of flowers is an insect environment that reflects the insects themselves (p. vii) . . . the plants and their pollinators are environment and reflection of one another” (p. 266). Such sensory-neural and environment coevolution provides, then, a particularly dramatic example of how the visual environment is not only relative to the animal, but also partly determined by the visually guided activity of the animal itself.

Such animal-environment codetermination is not limited to invertebrates. As Humphrey (1984) has observed, most of the world’s colors are organic colors carried by the pigments of plants and animals – for example, the colors of flowers and fruits, of plumages, of tropical fishes, and so on. Such organic colors have been selected because of their biological significance to those who can see them. It is interesting to note that some pigments, for example, carotenoids, play a key role both on the side of the discriminated object (plants, fruits, feathers), and on the side of the primary processes in the retina (visual pigments, oil droplets). Thus the presence of carotenoids is emblematic of the evolutionary codetermination of per-

ceiving animals and their environments (Rothschild 1975).

Hardin’s subjectivism neglects this role that visual perception plays in animal-environment codetermination. The neglect derives, we believe, from Hardin’s implicit acceptance of the subjectivist-objectivist framework for evaluating perception derived from Galileo, Newton, and Locke. Thus although Hardin has emphasized the role that color vision plays in generating chromatic categories that have intersubjective, cognitive significance for perceiving animals in their interactions, he nonetheless wishes to drive a principled wedge between, on the one hand, color construed as a subjective encoding of information about the world, and on the other hand, surface reflectances construed as objective properties of the world. He claims, for example, that colors are subjective because he supposes that if there were no perceiving animals in the world, there would be no colors; since objects and their surfaces would remain, however, these are objective (Hardin 1990). (This same argument was in fact given by Galileo in 1623: “. . . Colors and so on are no more than mere names so far as the object in which we place them is concerned, and . . . they reside only in the consciousness. Hence if the living creatures were removed, all these qualities would be wiped away and annihilated” [Drake 1957, p. 274].)

This line of argument not only overlooks but actually does violence to virtually every aspect of the ecologically entangled relations of perceiving animals and their environments. First, it overlooks the fact just mentioned that most of the world’s colors are organic colors. The evolution of color vision is inextricably linked to the evolution of organic coloration – so much so that “in a world without animals that possessed colour vision there would be very little colour” (Humphrey 1984, p. 146). It is therefore irrelevant – perhaps even somewhat perverse – to appeal to *metaphysical* intuitions about what the world would be like “if the living creatures were removed” when one’s concern is to provide a *naturalistic* explanation of perceiving animals and their environments.

Second, Hardin’s argument overlooks the role that color vision plays in spatial segmentation and hence the relational nature of the surfaces of perceptual objects, which we reviewed above. Elsewhere Hardin (1988, pp. 111–12) has himself drawn attention to a similar point: “Because perceptions of color differences and perceptions of boundaries are closely intertwined neural processes, we see colors and shapes together. Roughly speaking, as color goes, so goes visual shape.” For Hardin, however, there is an important difference between color and shape; thus he continues: “Consequently, there are no visual shapes in the ultimate sense, just as there are no colors. But visual shapes have their structural analogues in the physical world, namely, shapes *simpliciter*, and colors do not.”¹⁵ We find this point unclear, for Hardin does not tell us exactly what he means by “structurally analogous” and “shapes *simpliciter*.” We obviously agree with Hardin that colors do not have structural analogues in the physical world *in the way that objectivists have supposed* – that is, analogues that do not depend in any way upon the existence of perceivers. This point, however, does not prevent our specifying *context-dependent and interest-relative* structural analogues of

color, as the science of colorimetry and its associated color technologies clearly indicate (see Hurvich 1981, Chapters 20–21). This point might strike some as unfair, since Hardin's claim might be that there are no context-independent and non-interest-relative structural analogues for color, whereas there are for visual shape. But if this is Hardin's point, then we are not at all convinced it is true. Unlike Newton and Locke, we no longer take shape to be among the fundamental, microscopic properties of matter (cf. Priest 1989). And, as a macroscopic property, what gets picked out as a given shape may depend on the interests and capacities of those performing the specification. In this sense, surfaces as specified in terms of shapes and boundaries might be more properly thought to belong not to the physical world *per se* (the world at a purely physical level of description), but rather to what Gibson (1979) calls the "ecological environment," that is, the world as construed in relation to certain animal capacities (cf. Stroll 1986).

The moral of these considerations, we believe, is that the empirical study of color vision – indeed, of perception in general – should not be saddled with some *a priori* subjective/objective distinction. There is nothing wrong with drawing a distinction between subjective and objective, or internal and external, relative to the framework of a given neurophysiological, psychophysical, or behavioral experiment. The problems arise, rather, when we attempt to force perception theory as a whole into some absolute, subjective/objective straitjacket derived from the empiricist tradition. Hardin (1988) has already impressively demolished many of the dogmas about color in this tradition. He has rightly built his case from biological evidence, but this evidence demands a more sophisticated interactionist approach to color vision than neurophysiological subjectivism delivers. We now turn to the more constructive task of outlining such an approach.

3. Toward an enactive view of color vision

Although the shortcomings of computational objectivism and neurophysiological subjectivism are different, they are related. Computational objectivism conceives of color vision as the "recovery" of animal-independent, distal properties; neurophysiological subjectivism conceives of color vision as the "projection" of subjectively generated qualities onto a distal world of objects and their surfaces. In either case, the role that vision plays in the codetermination of animal and environment is neglected.

Consider the question: "Which came first, the world or the image?" The answer of inverse optics is given ambiguously by the names of the tasks investigated – to recover shape from shading, surface reflectance from varying illuminants, and so on. We call this stance the "chicken position":

Chicken position: The distal world can be specified independently of the animal; it casts images on the perceptual system whose task is to recover the world appropriately from them.

This position is so ingrained that we tend to think the only alternative is the "egg position":

Egg position: The perceptual system projects its own world and the apparent reality of this world is merely a reflection of internal laws of the system.

Our discussion of color vision, however, indicates that neither position is satisfactory. We have seen that colors are not already labelled properties in the world which the perceiving animal must simply recover (objectivism). On the other hand, we have seen that they are not internally generated qualities that the animal simply projects onto the world (subjectivism). Rather, colors are properties of the world that result from animal-environment codetermination. Our case study of color vision suggests that the world and the perceiving animal determine each other, like chicken and egg.

To situate our discussion of vision within the context of animal-environment codetermination, it is worth repeating the summary provided by Levins and Lewontin (1983; 1985) of how organisms "construct" their environments: (1) Organisms determine in and through their interactions what in the physical environment constitutes their relative environments; (2) organisms alter the world external to them as they interact with it; (3) organisms transduce the physical signals that reach them, and so the significance of these signals depends on the structure of the organism; (4) organisms transform the statistical pattern of environmental variation in the world external to them; and (5) the organism-environment relationship defines the "traits" selected for in evolution (cf. Oyama 1985). These five kinds of phenomena involve circular and reciprocal (though not symmetrical) processes of interaction in which the structure of the environment constrains the activity of the organism, but the activity of the organism shapes the environment, and so contributes to the constitution of the environmental constraints (cf. Odling-Smee 1988). It is on the basis of these interactive processes that Levins and Lewontin claim that "the environment and the organism actively co-determine each other" (1985, p. 89).

The implications of this codetermination of animals and their environments have been mostly neglected in perceptual theory, not only by the computational research program of inverse optics, but even by proponents of the so-called "ecological" approach to visual perception (Gibson 1979; Turvey et al. 1981). We will comment on the ecological approach presently; at the moment, we wish to delve further into the reasons for the neglect of animal-environment codetermination in the research program of inverse optics (Marr 1982; Poggio et al. 1985).

Simplifying for the purposes of brevity, inverse optics claims that the animal visually perceives by instantiating various functions that map from two-dimensional images on the receptor array (input) to perceptions of the three-dimensional world (output) via intermediate representations (and given various independent physical constraints). So stated, this account of perception has at least three important consequences that run counter to the idea that visual perception participates in animal-environment codetermination.

First, animal and environment are treated as fundamentally separate systems: The distal environment (objects, surfaces, etc.) is specified in advance; it provides a source of input that is independent of the animal. The perceiving animal, on the other hand, is treated as an input-output system whose function is to solve the ill-posed problem of recovering this prespecified environment. Second, perceptual and motor mechanisms are treated as fundamentally distinct subsystems of the ani-

mal. Since the "outputs" of perceptual systems are considered to be perceptual beliefs about the distal scene, perceptual systems form a mechanism for the fixation of belief. On the basis of its perceptual beliefs, the animal may adjust its activity, but the adjustment of activity *per se* is not treated as part of the perceptual process. Third, perception does not in any way shape the environment; it merely recovers the environment. It might be admitted that animal activity can perturb, select, or construct the environment, but since perception is considered to be fundamentally distinct from action, perception *per se* does not participate in animal-environment codetermination.

This account of perception is based in a well-established empirical research program and so should not be dismissed either on conceptual grounds or simply by adducing counterexamples. It can be challenged, however, by offering an alternative theoretical and empirical framework as a rival research program. At this point in our target article we obviously do not intend to embark on a detailed defense of such an alternative research program.¹⁶ Our intention here is simply to outline briefly a framework for understanding visual perception in which we take seriously the role of vision in the codetermination of animal and environment.

The first step for perceptual theory is to refuse to separate perception from action, or, more generally, from perceptually guided activity. This refusal is in fact common to a number of different research programs, such as the "ecological approach" of Gibson (1979) and his followers (Turvey et al. 1981), the biological approach to cognition of Maturana and Varela (1980; 1987), Freeman's view of brain processes (Freeman 1975; Freeman & Skarda 1985; Skarda & Freeman 1987), and the recent work in AI and robotics of Brooks (1986; 1987; 1989). All of these research programs take as central the fact that perception and action have evolved together – that perception is always *perceptually guided activity*. But whereas the first research program (Gibson's) chooses to focus on properties of the animal environment and optical properties of the ambient light, the others focus on the sensory-motor structure of the animal, either as neuronal networks that link sensory and motor surfaces or as "layers" of "activity producing systems" in artificial robotic "creatures."

We must encompass both the extradermal world conceived as the animal's environment and the sensory-motor structure of the animal in any adequate theory of perception. We believe that the original Gibsonian program exaggerated the role of invariances in the receptor array activity and their hypothesized specification of the environment. That program neglected not only the complex neural processes that are required to guide activity, but also how those processes contribute to shaping different environments depending on the animal. The original Gibsonian program remains unsatisfactory precisely because it does not take this further step, namely that of shifting the reference point for understanding perception from the environment to the structure of the perceiving animal, understood as the kinds of self-organizing neuronal networks that couple sensory and motor surfaces, which determine both how the animal can be modulated by environmental events and how sensory-motor activity participates in animal-environment codetermination.

[See also Ullman: "Against Direct Perception" *BBS* 3(3) 1980.] Elsewhere one of us has argued that a consistent application of this shift in perspective is tantamount to treating the animal as an autonomous self-organizing system rather than as a heteronomous input-output system (Varela 1979; 1984; 1989; see also Freeman & Skarda 1985; Skarda & Freeman 1987). We do not intend to repeat these arguments here; we mention the point because it is primarily this second step – emphasizing the autonomous organization of the animal – that marks the difference between our emphasis on perceptually guided activity and Gibson's. In contrast, many Gibsonians continue to treat perception in largely optical terms, and so attempt to build up the theory of perception almost entirely from the side of the environment. We believe this tendency is largely the result of Gibson's belief that the only alternative to the mistaken sense-data view of perception is direct realism (see Gibson 1967; Turvey 1977). Our approach, however, like that of some more recent Gibsonians (e.g., Kelso & Kay 1987), takes from Gibson the deep insight that perception must be understood within the ecological context of guided activity, but we develop this insight in two important ways: (1) by focusing on the self-organizing properties of neural networks as the proper substrate of animal activity; and (2) by treating the environment not simply as the ecological setting for animal activity, but also as something determined by that very activity. To label this concern with perceptually guided activity thus understood, we will use the term *enactive* as proposed by Varela (1989; 1991a), and as subsequently developed by Varela et al. (1991b).

The point of departure for an enactive approach to vision, then, is not the problem of recovering a pre-specified distal world. Rather, it is to specify the sensory-motor patterns that underlie the visual guidance of animal activity in its local situation. Our examination of differences in color vision led us to hypothesize that animals with different sensory-motor capacities would segment the world in different ways. As a corollary, we claim that the prespecified world we find in, say, low-dimensional models of surface reflectance is actually the world as described in relation to the sensory-motor capacities of the higher primates. It is perhaps a legitimate simplification to specify or label the world in advance when studying our own visual capacities (or those of animals very much like us). It is not legitimate, however, when studying perception in animals that differ considerably from us.

To make this point clearer, consider again the visual system of birds, which provides such a stark contrast to the visual systems of the more familiar mammals. As we mentioned above, the avian retina has two regions of high neuronal density (foveas), which give rise to distinct frontal and lateral visual fields that in turn correspond roughly to further anatomical projections in the brain – the parallel thalamo-fugal and tecto-fugal pathways. Experiments reveal interesting differences between these two visual fields: Frontal fixation is used for static and slow stimuli, and lateral fixation for fast-moving stimuli (Maldonado et al. 1988). There are also differences in accommodation, depth of focus (Bloch & Martinoya 1983), spectral sensitivity (Nuboer & Wortel 1987; Remy & Emmerton 1989), and probably chromatic vision (Varela et al. 1983). Thus, visual discrimination for birds is not a cyclopean image reconstruction but a con-

Actualized specification according to avian sensory-motor activity – a visual world-to-the-front and a visual world-to-the-side are enacted by the animal. It is the enactive behavior that actually reveals what constitutes a relevant world for the animal, not a reconstruction of the world as it appears visually to us.

This emphasis on sensory-motor patterns of activity is, of course, incompatible with abstract task-analyses for vision per se. Our objection, rather, is to the biologically implausible idea of a prespecified or already labelled world that the perceiving animal must recover appropriately. Although this assumption is built into Marr's conception of the computational level of analysis and of vision as inverse optics, it need not be accepted by those who wish to provide abstract task-analyses for vision and to build artificial visual systems. Indeed, there are models that considerably relax this assumption, such as Grossberg's (1984; Carpenter & Grossberg 1987) adaptive resonant neuronal networks and Edelman's selective recognition automata (Reeke & Edelman 1988). Similarly, in Brooks's (1986; 1987; 1989) recent works in robotics, the ongoing updating of sensory-motor activity is the key for successful design, rather than the representation of prespecified features of the world. By construing visual perception not as recovery or re-presentation, but as guided activity, these models implicitly embody the shift in perspective that we are calling "enaction."

This enactive orientation also implies an understanding of the relationship between the physical and the perceptual different from the usual one in the computational level of analysis. Inverse optics typically assumes that the task of perception is simply to recover properties of the physical world. The enactive approach suggests that perception is not about the physical world in this way. The world that a given animal perceives cannot be given a purely physical-level specification, for what an animal perceives depends on three kinds of factors: (1) physical-level constraints; (2) sensory-motor activity as constituted by neuronal processes and developmental constraints; and (3) evolutionary history. For example, such physical-level constraints as spectral reflectances and light signals are certainly ingredients of what the animal sees. They are not sufficient to determine the *perceptual object*, however, for, as we have seen, color spaces of different dimensionalities can be constructed on the basis of the same physical signals. To account for these differences and hence for the differences in color among the relevant perceptual objects, we must in addition appeal to sensory-motor activity and evolutionary history. Each of these three factors is necessary to determine the perceptual object; in the absence of any one of them, therefore, the perceptual object cannot be properly explained.

This claim about the status of the perceptual object also serves to mark the difference between enaction and subjectivism. Hardin's subjectivism implies that the perceptual object is simply "in the head," and so can be reconstructed in entirely neural terms. As he says: "The idea that suggests itself is to show how phenomena of the external field are represented in the visual cortex and then to show how descriptions of the visual field may be replaced by descriptions of neural processes" (Hardin 1988, p. 111). Our critique of neurophysiological subjectivism in the previous section implies that the perceptual object, though experiential, is also ineliminably ecologi-

cal, and so, contrary to Hardin, is not simply "in the head."

The enactive view of perceptual content is also different from both the "externalist" view that perceptual content is provided by distal physical properties and the "internalist" view that perceptual content is provided by subjective qualities (qualia). According to the enactive view, the contents of perceptual states are to be type-identified by way of the ecological properties perceived, and these ecological properties are to be type-identified by way of the states that perceive them. One should not be put off by this circularity, for it is informative. To specify perceptual content for a given animal we must investigate the relevant environmental properties, and to determine the relevant environmental properties we must investigate the sensory-motor patterns of activity that constitute the animal's perceptual states. This circularity is also empirically well-founded: Recall the discussion of how color vision and the ecological properties detected by color vision (e.g., plant and animal coloration) have in the course of evolution been selected for each other. The enactive view of perceptual content thus follows from animal-environment codetermination.

Now that we have provided an idea of the kind of conceptual space in which an enactive approach to vision could grow, let us return specifically to color vision. According to enactivism, color is neither a perceiver-independent property, as in objectivism nor is it merely a projection or property of the brain, as in subjectivism. Rather, it is a property of the enacted perceptual environments experienced by animals in their visually guided interactions. Unlike computational objectivism and neurophysiological subjectivism, this does not lead to an eliminativist position regarding color: color is not divested of its phenomenal or experiential structure in favor of spectral reflectance; nor is it divested of its extradermal locus in favor of neural states. Instead, color is a property of the extradermal world understood as an animal's environment, a world that is enacted by animal-environment codetermination. Thus we arrive at the view announced at the beginning of this paper, according to which color is both ecological and experiential.

Our view might in some respects recall Locke's (1690/1975) concept of color as a relational property, but there are significant differences. Locke held that color is relational because it is a "secondary quality," a disposition of objects to cause color sensations in a perceiver. According to the Lockean view, then, color is not merely *relational*, but also *dispositional* and *subjective* (see Bennett 1971). Ecological experientialism, however, does not imply that color is dispositional and subjective. We have not tried to explicate the relational nature of color by attempting to link dispositional properties of an organism-independent physical world, and private sensations, qualia, or sense-data. This is not feasible, we feel, despite repeated empiricist attempts (Westphal 1987; Thompson 1989). Nor does ecological experientialism rest on the distinction between primary and secondary qualities. On the contrary, our argument that not only color but also other high-level, spatial properties of the scene (object surfaces as determined by shapes and boundaries) are relational runs directly counter to the Lockean and Newtonian attempt to draw a principled distinction between color as a secondary quality, and size,

shape, and so forth, as primary qualities. Rather, we have emphasized the relational nature of the perceptual environment as a whole resulting from the enactive dimensions of visually guided activity.

Our intention in this target article has been to offer a broad, comparative framework for the ongoing, interdisciplinary effort to understand color vision and visual perception in cognitive science. This framework suggests specific directions for further research.

(i) The first concern of our comparative approach is to determine more precisely the kinds of color space there are in the animal world. For tetrachromacy, we need further evidence of four-way color mixture; to establish pentachromacy, we need evidence of five-way color mixture. Frequency modulation in the study of color vision (as described in Appendix B) might be useful in this area.

(ii) A related question concerns how the relevant color vision mechanisms and the dimensionalities of color space are related to perceptual phenomena such as constancy and segmentation of the visual scene. Of particular interest here is how color and other visual phenomena such as visual shape, texture, and space, interact to constitute different perceptual objects for various perceiving animals.

(iii) A third research objective is to determine (at least to a degree comparable to what is known of primates) the neuronal mechanisms underlying the variety of color spaces of different animal groups, especially fishes and birds. This is the key to understanding how color vision figures in the larger context of animal life and behavior.

(iv) Finally, the ecological aspects of the perceptual environment need to be investigated, for example, local illuminance and reflectance conditions, animal coloration, and animal communication. The task here is to develop further the means to describe the perceptual environment from a given animal's point of view, rather than imposing anthropocentric assumptions about such environments.

These questions have hardly begun to be addressed in detail in visual science, but we can expect their investigation to reveal even further the splendor of color as a naturalized aesthetic, or, in the words of Cézanne, "the place where our brain and the universe meet" (Merleau-Ponty 1964, p. 67).

APPENDIX A

Decorrelation procedure for calculating chromatic channels (Buchsbaum & Gottschalk 1983; Palacios, in press)

Consider responses r_i to an arbitrary illuminant belonging to a set $\{I(\lambda)\}$,

$$r_i = \int d\lambda w_i R^i(\lambda) I(\lambda)$$

where the w_i are weighting factors for each primary response (as explained below). Next construct a covariance matrix Γ between the receptor response as follows:

$$\Gamma = \begin{pmatrix} \gamma_{11} & \dots & \gamma_{1n_{recep}} \\ \vdots & \ddots & \vdots \\ \gamma_{n_{recep}1} & \dots & \gamma_{n_{recep}n_{recep}} \end{pmatrix}$$

with

$$\gamma_{ij} = \text{Ex}\{r_i r_j\} - \text{Ex}\{r_i\} \cdot \text{Ex}\{r_j\} \quad [11]$$

where Ex is the expectation operator. To achieve optimal decorrelation the obvious step is to obtain the eigenvalues φ_i and eigenvectors for Γ , and the new matrix A constituted of the eigenvectors and their transpose A^T so that:

$$A^T \circ \Gamma \circ A = \begin{bmatrix} \varphi_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \varphi_{n_{recep}} \end{bmatrix}$$

The eigenvector transformation is now well defined by:

$$\tilde{c}_r = A^T \cdot \tilde{s}_r \quad [12]$$

and the postulated chromatic channels ($C^1, \dots, C^{n_{recep}}$) thus calculated can be compared with the available experimental evidence. As Buchsbaum and Gottschalk (1983) emphasize, the covariance matrix depends not only on the shape of the primary responses, but also on the ensemble properties of the illuminants $\{I(\lambda)\}$. In fact, to arrive at an explicit expression for the correlations, we need to make some assumptions about the expectations $\text{Ex}\{I(\lambda)\}$ of the ensemble. A correlation $R(\lambda, \mu)$ and a covariance $K(\lambda, \mu)$ function can be defined as follows:

$$\begin{aligned} R(\lambda, \mu) &= \text{Ex}\{I(\lambda)I(\mu)\} \\ K(\lambda, \mu) &= R(\lambda, \mu) - \text{Ex}\{I(\lambda)\} \text{Ex}\{I(\mu)\}. \end{aligned} \quad [13]$$

If the choice is $K(\lambda, \mu) = \delta(\lambda - \mu)$ where δ is the Dirac delta function, this amounts to using monochromatic illuminants. Inserting [13] into [11] finally yields an explicit form for the entries in the covariance matrix [11]:

$$\gamma_{ij} = \int d\lambda R^i(\lambda) R^j(\lambda) \quad [14]$$

Thus the relative contributions from each class of retinal receptors need to be filled in by weighting factors previously mentioned. These are the only unknowns in our calculations; we have adjusted them so that the resulting channels have a good fit with experiments. These values should not be seen as ad hoc, however, but as proportions that should covary with neural characteristics.

For example, in the case of the goldfish this procedure yields:

C^1	0.11	0.37	0.72	0.58	R^1
C^2	0.83	-0.5	0.23	-0.12	R^2
C^3	0.26	0.68	0.15	-0.67	R^3
C^4	-0.48	-0.4	0.64	-0.44	R^4

For the putative color channels of the pigeon, see Figure 4d.

APPENDIX B

Frequency analysis of color vision (Barlow 1982; Bonnardel & Varela 1989).

The basic strategy is to consider an illuminant (or a response capacity of the visual system) in the frequency domain, that is, to examine the spectral power distribution of the signal (or the response mechanism) in terms of cycles over "visible" window. For example, we can consider the frequency response of the three chromatic channels required for human vision. The channels proposed by Hurvich and Jameson, when studied under Fourier analysis, predict that beyond 2-3 cycles/300nm there should be little response, with a peak of sensitivity for signals around 1.7 cycles/300nm. These predictions correspond well with the first measurements of such a modulation sensitivity function (MSF) obtained with a specially built apparatus that can produce sinusoidally modulated illuminants with controlled contrast, frequency, and phase (Bonnardel & Varela 1989). The conclusion is that the signals for color vision are *band-limited*, that is, bounded in both the variable ("visible" window) and the frequency (cycles per "visible" window) domain. Now, a nontrivial relationship exists between a collection of such band-limited signals and the number of significant independent samples required to reconstruct with sufficient accuracy any func-

in the collection (Buchsbau & Gottschalk 1983; Dym & Kean 1975). Specifically, it can be shown that:

$$n_{\text{sample}} = \Phi[BT] + 1 \quad [15]$$

where $\Phi[x]$ stands for the highest integer smaller than x ; and n_{sample} is the number of independent channels required to sample the space of signals limited by B (in the frequency domain) and by T (in the wavelength domain).¹⁷ For example, if we take $n_{\text{sample}} = 3$, and $T = 300\text{nm}$, as in the human trichromatic system, a band-limitation $B = 1.5$ cycle/300nm is predicted. This limitation is within the range of the measured band-limitation of human natural scenes, which contain about 1% of all reflectances within the 1.5 cyc/300nm limit (Maloney 1985) but falls a little short of the observed MTF in humans which peaks at this value. In contrast, if $n_{\text{sample}} = 4$ and $T = 300\text{nm}$, as is the case in birds, one would predict a band-limit of $B = 0.001$ cyc/nm. This result is due in part to the avian sensitivity window being large into 370nm, and on the other hand, to the pigment sensitivities being narrower: the combination of these two allows for less demodulation of the MSF at higher frequencies and corresponds to the higher frequency content directly visible in the putative channels discussed before.

Briefly stated, then, different chromatic dimensions will satisfy the sampling theorem with different combinations of the three quantities involved (n_{sample} , B and T), thus permitting quantitative comparisons of diverse color vision mechanisms. In particular, the measurements of MSF might represent a way to ascertain directly the dimensionality of a color space. On the basis of such a determination of color space, one could then undertake a comparison of the frequency-limitations of color vision mechanisms with the frequency-limitations of the reflectances of the relevant objects in the animal's environment. As we mentioned above, data have been collected concerning the frequency-limitations of human natural scenes (Maloney 1985); in the case of birds, some data have been collected about the reflectance properties of objects such as feathers, which, not surprisingly, have higher frequency contents than those of human natural objects (Burkhardt 1989). Obviously, work in this area is just beginning.

ACKNOWLEDGMENTS

We have discussed the material presented in this paper with a number of people. In particular we wish to thank: Mark Anzures, Ronald deSousa, Daniel Dennett, Lynd Ferguson, Timothy Goldsmith, C. L. Hardin, David R. Hilbert, André Kukla, Jim McGilvray, Christa Neumeyer, William Seager, and Mark Thornton.

During the preparation of this paper ET was supported by a Social Sciences and Humanities Research Council of Canada Postdoctoral Fellowship (#456-89-0236). AP was supported by the Simone et Cino del Duca, and the Philippe Foundations. The financial support to FV from Fondation de France (Chaire Scientifique), CNRS, Ministère de la Recherche et la Technologie, and the Prince Trust Fund is gratefully acknowledged.

NOTES

1. For further discussion, see Churchland and Sejnowski (1988), Sejnowski et al. (1988), and Boden (1988). Marr's optimality assumption has been criticized by Kitcher (1988); for some related points see Ramachandran (1985).

2. In visual science there is confusion and some disagreement over the use of the terms "brightness" and "lightness." According to Wyszecki and Stiles (1982, pp. 493-500), *brightness* is the "attribute of a visual sensation according to which a given visual stimulus appears to be more or less intense" (p. 493), whereas *lightness* is the "attribute of a visual sensation according to which the area in which the visual stimulus is

presented appears to emit more or less light in proportion to that emitted by a similarly illuminated area perceived as a 'white' stimulus" (p. 494). Strictly speaking, then, "brightness" refers to a dim-to-dazzling scale, whereas "lightness" refers to the gray scale of black and white. Nevertheless, many visual scientists, a large number of whom we cite in this paper, use "brightness" to refer to the white-black dimension. Furthermore, "lightness" has also been used in a related, but somewhat different sense by Land (1977; 1983), and by other researchers in computational color vision (Hurlbert 1986), to mean the psychophysical correlate of average relative reflectance (Land 1983) or scaled integrated reflectance (McCann et al. 1976). To avoid confusion, then, we use "brightness" for the achromatic dimension.

3. The material on which we are drawing here can be found in Boynton (1979), Hurvich (1981), Wyszecki and Stiles (1982), Mollon and Sharpe (1983), and Ottoson and Zeki (1985).

4. Our point in citing this passage is not that computational color vision commits one to objectivism about color - Land (1978; 1983), for example, holds distinctly nonobjectivist views. It is, rather, that the computational conception of color vision as concerned almost exclusively with the task of recovering surface reflectance suggests a form of objectivism.

5. For additional studies of near-ultraviolet sensitivity in bird vision, see Wright (1979), Delius and Emmerton (1979), Emmerton and Delius (1980), Emmerton (1983), Burkhardt (1982; 1989), and Burkhardt and Maier (1989). Cones with peak sensitivity in the near-ultraviolet have also been found in fishes: Harosi and Hashimoto (1983); Neumeyer (1985); Bowmaker and Kunz (1987).

6. Notice that channel C2 in Figure 4d has the appearance of an achromatic channel because of very small negative contributions from the long-wave primaries. This issue, as well as the general problem of the relative merits of a tetrachromatic or pentachromatic model for explaining the empirical data on pigeon color vision, require more elaborate discussion than we can provide here (see Palacios 1991).

7. This point is made clearly by Gouras (1985, p. 386), though it is a familiar theme in the history of color science. See Wasserman (1979) for a historical survey.

8. Some philosophers would no doubt go even further and argue that we cannot know what tetrachromatic or pentachromatic perception is like because the relevant facts (tetrachromatic or pentachromatic experience) are accessible only from the point of view of tetrachromatic and pentachromatic perceivers (cf. Nagel 1974). Although this claim is certainly relevant to our discussion, space constraints do not allow us to consider the various arguments here. We will therefore reserve discussion of this matter for another occasion (see Akins 1990; P. M. Churchland 1985; and Jackson 1982).

9. Hardin (1988, p. 146) notes this possibility of ternary hues by imagining a hypothetical tetrachromatic "visual superwoman," but does not extend his discussion to actual tetrachromacy among vertebrates such as birds and fishes.

10. Several objectivists (P. M. Churchland 1985; 1986; Hilbert 1987, pp. 111-18; Matthen 1988) have argued that such a mapping can be found in Edwin Land's (1977; 1983) retinex color space in which colors correspond to points in a three-dimensional space whose axes correspond to values of lightness calculated independently in each of three long-, middle-, and short-wave bands based on the sensitivities of the human (cone) photoreceptors. These arguments overlook two features of Land's model: (i) The axes of Land's color space are usually given as axes of lightness, not (average relative) reflectance. This is important because lightness is a sensation that can be measured only by a visual system, and problems arise for the straightforward identification of lightness with reflectance just as they do for color. (ii) Since the retinex color space attempts to specify colors purely in terms of lightness values it does not model the opponent relations and unique/binary structure of hue. In fact,

we cannot at present be said to understand how (chromatic) color could be generated purely from (achromatic) lightness comparisons.

11. It should also be noted that computational objectivist arguments such as Matthen's (1988) rely on a very strong notion of adaptive biological functions. We believe there are serious problems with this notion, but we will not pursue this point here. See Varela (1984); Maturana and Varela (1980).

12. We intend to investigate this issue in greater detail in another essay.

13. This claim is similar to one made by Berkeley (1710, Part I, para. 10). See also Wilson (1987).

14. A similar claim could be made for polymorphism in the evolution of human color vision, for example, red-green color blindness in human males.

15. This line of argument corresponds closely to one of Locke's (1690/1975) formulations of the primary/secondary quality distinction. Locke held that in the case of shape, our ideas (visual shapes) resemble (are structurally analogous to) their physical causes (shapes simpliciter), whereas in the case of color, they do not. Our criticism of this view as espoused by Hardin is similar to Berkeley's rejection of the view as espoused by Locke (see note 13).

16. For extensive elaboration of a research program for neuroscience in which the perceiving animal is treated not as an input-output system specified in terms of external mechanisms of control, but rather as an autonomous self-organizing system, see Maturana and Varela (1980; 1987), and Varela (1979; 1984; 1989; 1991a; 1991b), and Varela et al. (1991b).

17. This so-called sampling theorem requires, however, that the ensemble of band-limited functions have stringent averages (Brill & Benzschawel 1985).

Open Peer Commentary

Commentary submitted by the qualified professional readership of this journal will be considered for publication in a later issue as Continuing Commentary on this article. Integrative overviews and syntheses are especially encouraged.

More than mere coloring: The art of spectral vision

Kathleen A. Akins and John Lamping

Xerox Palo Alto Research Center, Palo Alto, CA 94304

Electronic mail: akins@parc.xerox.com

Suppose you were God, with a month of free Sundays, and you decided to design a visual system for a small creature. What would you do with the abundance of spectral information available from the world? How would you make use of the fact that sunlight ranges in wavelength from about 320 nm. in the ultraviolet on out past 800 nm. in the infrared (NASA 1977), or of the fact that the refraction, absorption, and reflection of light by various media alters the composition of those spectral frequencies?

The right answer, we think (as Thompson et al. would agree), is just this – it all depends. First, you would (as a less than omnipotent god) be subject to the standard "materials" constraints. These are the various facts about the way the physical world is – the amount of energy it takes to break chemical bonds, what wavelengths of light would be absorbed by potential materials for the eye, the abundance or paucity of particular

wavelengths in the atmosphere, and so on. Then there are questions about the organism itself: Where does it live and what does it do? Exactly how you would draw on spectral information would depend on what visual tasks are appropriate to the organisms tempered by the spectral information locally available.

Consider an actual example – the visual/spectral problems confronting the fish (Levine & MacNichol 1979; Munz & MacFarland 1977). First, water absorbs light – the deeper the water the less light – but exactly which wavelengths are absorbed depends on the types of suspended organic particles. So, for example, clear water with little organic matter (such as in tropical oceans) will most easily absorb red and violet light. At depths below 25 meters, essentially all red and violet light has been filtered out, leaving only the intermediate "blue" wavelengths. Hence the "Mediterranean blue" appearance of such water to us. Marshes and swamps, but contrast, contain decomposing plants, tannins, lignins, and yellow-green plankton that effectively filter out almost all light at the depth of only three meters. The light that does survive, however, is in the red-orange region of the spectrum and gives the water its dark red-brown appearance. In addition, water molecules (and the suspended particles in water) scatter light – the medium of transport is itself "colored." So, assuming that one visual task of the fish is to spot objects underwater, this detection must be done against a background of colored light.

One current hypothesis is that the rods and cones of some fish have evolved to function as *filters* that serve to highlight the contours of objects against the background space light. If the photopigment is "matched" (maximally sensitive) to the spectral range of the background light, then a dark object will be highlighted against the brighter background; if the photopigment is "offset" from the dominant wavelengths of the background light, then a bright object will be outlined against a poorly illuminated background. In fact, this scheme seems to be used in many species of fish. For example, the skipjack tuna, which spots its prey from below (dark object against a bright surface), has only one photopigment that is matched to the background light. The guppy's retina is divided into two regions. In the lower region, which looks upward and is used to spot prey, there is one photopigment that is matched to the background space light. The upper region, which looks down on the colorful mating display of the male guppy, has three spectrally diverse cone types. Finally, the walleye, bluegill, and piranha, which inhabit dark, particle-laden (hence, red-shifted) waters, are each dichromats, with one cone type matched to the near-infrared – a wavelength common to their "black" water habitat during the dusk and dawn hours of feeding.

What the above examples illustrate, we think, is that the use of spectral information could take many forms. It could be used to recover the invariant surface reflectances of opaque objects; then again, it might be used for the delineation of object from ground, to discern shape, monitor motion (Dobkins & Albright 1990), as a cue for eye movement (Ballard, forthcoming), and so on. There is, in other words, no *a priori* answer to the question, "What role does spectral information play in a visual system?"

What are the implications of this fact for color vision? Questions of consciousness aside, color vision, as defined behaviorally, is a certain *kind* of chromatic visual function – a selective response to light stimuli of different wavelengths regardless of intensity. (There seems no need to second guess the kinds of visual function that are necessary for our color experience insofar as questions of ontology hinge on the facts about function. If the ontological conclusions can be reached for any kind of color vision, then they will hold for conscious color experience as well.) Here, the form of such discriminations can vary immensely between species – the type of color vision a creature has will depend on the number of its color receptors, the response profiles of those receptors, and the additional "wiring" farther up the line. One can think of the initial "choice" of